

On the Theory and Practice of Modern Secure Messaging

Présentée le 25 avril 2024

Faculté informatique et communications
Laboratoire de sécurité et de cryptographie
Programme doctoral en informatique et communications

pour l'obtention du grade de Docteur ès Sciences

par

Daniel Patrick COLLINS

Acceptée sur proposition du jury

Prof. R. Guerraoui, président du jury
Prof. S. Vaudenay, directeur de thèse
Prof. P. Rösler, rapporteur
Dr B. Poettering, rapporteur
Prof. M. Kapralov, rapporteur

To mum and dad.

Abstract

Billions of people now have conversations daily over the Internet. A large portion of this communication takes place via secure messaging protocols that offer “end-to-end encryption” guarantees and resilience to compromise like the widely-used Double Ratchet protocol of Perrin and Marlinspike. This thesis explores secure messaging from a cryptographic perspective in both the analysis and improvement of existing messaging solutions as well as the design of protocols with new security and efficiency characteristics.

The first half of the thesis considers communication between *two* parties. We first draw our attention to the impending threat of quantum computers on Diffie-Hellman-based key exchange protocols, and in particular on the widely used Extended Triple Diffie-Hellman (X3DH) key exchange protocol of Perrin and Marlinspike. We propose a new deniable authenticated key exchange protocol, K-Waay, that is based on the relatively conservative plain learning-with-errors (LWE) assumption and is faster than previous proposals. We then consider *active attack detection* to ensure parties can detect if and when they have been compromised and impersonated by an adversary, even if messages can be delivered out-of-order like in the Double Ratchet protocol. We consider both in-band and out-of-band detection (the latter offering better security but being less convenient for users) and prove formally that immediate active attack detection is sometimes inherently expensive but, despite this, a relaxed yet meaningful notion of active attack detection can be achieved practically.

The second half of this thesis then deals with communication between a dynamic *group* of parties. Firstly, we formalise the *group administration* problem where a (dynamic) portion of a given group is entrusted with additional privileges: we identify and formalise their core role of enforcing *access control*. We propose two protocols extending the *continuous group key agreement* methodology underpinning the recent IETF Messaging Layer Security (MLS) standard, and demonstrate experimentally that administration can be achieved with very little overhead for MLS. Finally, we formalise the practical Sender Keys group messaging protocol used by WhatsApp and Signal (which in fact relies on two-party communication at its core) and prove in a new security model that the core protocol structure is sound. Through our formalisation, we report some drawbacks of Sender Keys, especially in terms of its resilience to state compromise, and propose some tweaks to overcome them using standard cryptographic primitives, each of which either incurs little overhead or in fact improves practical efficiency.

Résumé

De nos jours, des milliards de personnes utilisent régulièrement internet pour communiquer. Une grande partie de ces communications utilisent des protocoles de messagerie sécurisée qui offrent des garanties de “chiffrement de bout en bout” comme le très répandu protocole Double Ratchet de Perrin et Marlinspike. Dans cette thèse, nous explorons ainsi la notion de messagerie sécurisée. Nous étudions à la fois des solutions existantes mais nous proposons aussi de nouveaux protocoles supportant de nouvelles contraintes de sécurité et d’efficacité.

La première moitié de cette thèse porte sur la communication entre deux parties. Nous portons d’abord notre attention sur la menace imminente que les ordinateurs quantiques représentent pour les protocoles d’échange de clés basés sur le protocole Diffie-Hellman. Le protocole d’échange de clés Extended Triple Diffie-Hellman (X3DH) de Perrin et Marlinspike, qui est largement utilisé, y est particulièrement vulnérable. Nous proposons ainsi un nouveau protocole, K-Waay, qui permet l’échange de clés authentifiées tout en offrant la possibilité de déni plausible. Notre protocole est plus rapide que les propositions précédentes basées, et repose sur l’hypothèse que le problème d’apprentissage avec erreurs (LWE) est difficile. Nous examinons ensuite *la détection d’attaques* afin de garantir que les deux parties puissent détecter si un adversaire a usurpé une identité ou s’il a réussi à les compromettre, et cela même si les messages sont reçus dans le désordre, comme dans le protocole Double Ratchet. Nous considérons à la fois la détection en bande et hors bande (cette dernière offrant une meilleure sécurité mais étant moins pratique). Nous concluons que la détection immédiate d’une attaque active est parfois intrinsèquement coûteuse mais que, malgré cela, il est possible de la rendre pratique et efficace en relaxant légèrement le modèle de sécurité.

La seconde partie de cette thèse traite des communications entre un groupe dynamique de parties. Nous commençons par formaliser le problème *d’administration de groupe* où un sous-ensemble (dynamique) d’un groupe donné se voit confier des privilèges supplémentaires : nous identifions et formalisons leur rôle principal qui est de gérer le contrôle d’accès. Nous proposons deux protocoles étendant la méthodologie *d’accord de clé de groupe continu* qui étaye la nouvelle norme Messaging Layer Security (MLS) de l’IETF, et nous démontrons expérimentalement que l’administration peut être réalisée avec un surcoût faible pour MLS. Enfin, nous formalisons le protocole de messagerie de groupe Sender Keys utilisé par WhatsApp et Signal (qui repose sur une communication bipartite). Nous prouvons dans un nouveau modèle de sécurité que le protocole est robuste. Grâce à notre formalisation, nous signalons

Résumé

certaines inconvénients de Sender Keys, en particulier en matière de résilience face à un adversaire étatique, et nous proposons des ajustements basés sur des primitives cryptographiques standardisées afin de les surmonter.

Acknowledgements

I would first like to thank my advisor, Serge Vaudenay, for his continued guidance over the last four years, for taking a chance on me, and for allowing me to flourish in his lab LASEC. Thank you also to my private defence jury – Serge, Rachid Guerraoui, Michael Kapralov, Bertram Poettering and Paul Rösler – for their valuable time and feedback.

Thank you to the current members of LASEC. Loïs, my officemate of four years, thank you for our endless conversations, hilarious in-jokes and fun collaborations. Thank you Boris for your advice and company, Bénédict for our conversations and your insights, Abdullah for your kindness and for bonding with me over our shared interests, and Laurane for our various activities and many enjoyable chats. Thank you Sylvie for your support of both me and the lab.

Thank you to the former and spiritual members of LASEC. Andi, for bonding over pop culture, projects and everything in between; Khashayar, for our variety of conversations and the memorable lifts home; Ritam, for always being up for cake and to chat; Simone, for all the memories – the travel, collaborations and friendship. David and Phillip, I'm glad our collaborations and relationship did not end when you left LASEC. Thank you Subhadeep for your guidance and companionship when it was just us two in the evenings. Thank you Khanh for your friendship. Thank you Aymeric, Novak, Helen, Fatih and Gwangbae – your time at LASEC was valuable to me. Martine, thank you for your help and service over the years. Thank you to my project students for the opportunity to advise you and to the lab interns and visitors over the years.

Thank you to Rachid and Thanasis for your semester project supervision – distributed computing is still in my heart – and Jovan for your friendship.

A special thank you to Dario, Giulio, Julian and Paul for each hosting me in their labs for research visits – all of these trips were special and productive, and it was great to spend time with you and everyone else who I met during these visits. Thank you to Julian for taking me under your wing and Paul for mentoring me.

This thesis, and my research, would have been impossible without the effort and input of my collaborators, both past and present, and all those who supported me in the cryptographic community – a particular shoutout to those from Bochum, CISPA and IMDEA.

Acknowledgements

Thank you to all of the friends I met in Lausanne that helped me through my PhD since I arrived in September 2019. Guohao, for the endless conversations (we have had several since I started writing these acknowledgements) and irreplaceable friendship; Plouton, for always outdoing yourself in cooking and your warm companionship; Rodrigo, for your integrity, honesty and strength. Bakul, for teaching me how to be a roommate. ML4ED, for your companionship on the second floor of INF – thank you Jade for visiting Lois and I! Christian and Sylvain, thanks for the memories. Nicolas, coucou! Greg, thanks for your friendship. Aditya, Dina, Lars, Michal, Valentin and all those I've forgotten, merci!

Thank you to my Honours thesis advisor Vincent Gramoli for introducing me to and encouraging me in the world of research. Thank you to Tyler Crain for your formative guidance and for encouraging me to travel for my PhD, without whom this would not have happened. Thank you to all my other colleagues, peers and friends from the University of Sydney. Thank you to my teachers in school and university.

I would be remiss not to mention my friends, family and mentors who I met and grew up with in Sydney and beyond. Thank you to my mum, dad and brother for your endless love and support, to Tash and Joey, to Ash, Chris, Eric, Frank, Ian, and everyone else – I have missed you all dearly.

Finally, I would like to thank all those who I did not thank above, as there were many more who impacted my life and helped me.

Lausanne, April 9, 2024

D. C.

Contents

Abstract (English/Français)	i
Acknowledgements	v
1 Introduction	1
1.1 Context and Motivation	1
1.2 Contribution	4
1.2.1 Part I: Two-Party Communication	4
1.2.2 Part II: Group Messaging	7
1.3 List of Works	11
2 Preliminaries	13
2.1 Notation and Conventions	13
2.2 Cryptographic Primitives	15
2.2.1 Key-Encapsulation Mechanism (KEM)	15
2.2.2 Signatures	16
2.2.3 Pseudorandom Function (PRF)	17
2.2.4 Hash Function	18
2.2.5 Incremental Set Hash Function	18
2.2.6 Symmetric Encryption	19
I Two-Party Communication	23
3 K-Waay: Fast and Deniable Post-Quantum X3DH Without Ring Signatures	25
3.1 Contribution	25
3.1.1 Summary	27
3.1.2 Technical Overview	29
3.1.3 Additional Related Work	33
3.2 Split-KEM	33
3.2.1 Security	34
3.2.2 Deniability	37
3.3 Deniable Authenticated Key Exchange	38
3.3.1 Syntax	38
3.3.2 Security Model	39

Contents

3.3.3	Deniability	44
3.4	K-Waay: Post-Quantum X3DH from Split-KEM	45
3.4.1	Construction	45
3.4.2	Security	47
3.5	Deniable Split-KEM from Lattices	56
3.5.1	Lattice Toolbox	56
3.5.2	Extended-LWE	57
3.5.3	Construction	61
3.5.4	Security Analysis	61
3.5.5	Building a UNF-1KCA and IND-1BatchCCA Split-KEM	69
3.5.6	Concrete Instantiation	70
3.6	Evaluation and Discussion	72
3.6.1	Benchmarks	73
3.6.2	Advantages, Limitations and Discussion	76
4	On Active Attack Detection in Messaging with Immediate Decryption	79
4.1	Contribution	79
4.1.1	Summary	81
4.1.2	Technical Overview	82
4.1.3	Additional Related Work	85
4.2	(Authenticated) Ratcheted Communication	85
4.3	In-Band Active Attack Detection: RID	91
4.3.1	RID-Secure RC	94
4.4	Out-Of-Band Active Attack Detection: UNF	99
4.4.1	UNF-Secure ARC from a RID-Secure RC	100
4.4.2	UNF-Secure ARC from Any RC	102
4.5	Lower Bounds for Active Attack Detection	105
4.5.1	Communication Cost for r-RID Security	105
4.5.2	Communication Cost for r-UNF Security	111
4.6	Optimisations and Performance/Security Trade-Offs	111
4.6.1	On the Practicality of s-RID and s-UNF Security	112
4.6.2	Epoch-Based Optimisation for s-RID Security	113
4.6.3	Pruning for UNF Security	115
4.6.4	Lightweight Bidirectional Authentication	119
II	Group Messaging	125
5	Cryptographic Administration for Secure Group Messaging	127
5.1	Contribution	127
5.1.1	Summary	129
5.1.2	Technical Overview	130
5.1.3	Additional Related Work	132

5.2	(Administrated) Continuous Group Key Agreement	132
5.2.1	Continuous Group Key Agreement	132
5.2.2	Administrated CGKA	136
5.2.3	Correctness	137
5.2.4	Security	140
5.3	A-CGKA Constructions	145
5.3.1	Individual Admin Signatures	146
5.3.2	Dynamic Group Signature	159
5.3.3	Description	160
5.3.4	Integrating A-CGKA into MLS	169
5.4	Evaluation and Discussion	170
5.4.1	Benchmarks and Performance	171
5.4.2	Modelling in Related Work	175
6	WhatsApp with Sender Keys? Analysis, Improvements and Security Proofs	177
6.1	Contribution	177
6.1.1	Summary	179
6.1.2	Technical Overview	180
6.1.3	Additional Related Work	183
6.2	Two-Party Channels	184
6.2.1	Primitive Definition and Correctness	184
6.3	Group Messenger	191
6.3.1	Security Model	194
6.3.2	Modelling Two-Party Channel Ciphertexts	197
6.4	Sender Keys	198
6.4.1	Protocol	198
6.4.2	Security	201
6.5	Analysis and Improvements	206
6.5.1	Security Analysis and Limitations	206
6.5.2	Proposed Improvements: Sender Keys+	209
6.5.3	Sender Keys/Sender Keys+ vs CGKA	212
6.5.4	Sender Keys in Practice	213
6.6	Sender Keys and Sender Keys+: Full Protocols and Security	214
6.6.1	Protocol Descriptions	214
6.6.2	Sender Keys Security	219
6.6.3	Sender Keys+ Security	226
7	Conclusion	231
7.1	Primitive Summary	231
7.1.1	Two-Party Communication	231
7.1.2	Group Communication	232
7.2	Discussion and Future Work	233

Contents

7.2.1	Composability	233
7.2.2	Model Limitations	233
7.2.3	Unification and Verification	234
7.2.4	Deniable Post-Quantum X3DH (Chapter 3)	234
7.2.5	Active Attack Detection in Messaging (Chapter 4)	235
7.2.6	Group Administration (Chapter 5)	236
7.2.7	Sender Keys (Chapter 6)	238
A	Appendices	241
A.1	QROM Preliminaries	241
A.2	Proof of Theorem 5 (Chapter 3)	243
A.2.1	Proof in the QROM	243
A.2.2	Proof in the ROM	246
A.3	Proof of Theorem 6 (Chapter 3)	247
A.3.1	Proof in the ROM	247
A.3.2	Proof in the QROM	248
A.4	Tables for Sender Keys (Chapter 6)	252
	Bibliography	277
	Curriculum Vitae	279

1 Introduction

1.1 Context and Motivation

For thousands of years, the primary aim of cryptography has been to enable secret communication between physically separated parties. In 1976, Diffie and Hellman proposed an elegant and now ubiquitous method for enabling two parties to share a secret over an untrusted network [DH76]. As a result of the ensuing cryptographic revolution, cryptography is used today by billions of people and countless computers each day to protect communication via protocols like Transport Layer Security (TLS) [Res18] and, more recently, secure messaging applications.

Secure messaging, the focus of this thesis, has exploded in popularity and importance in the last decade. This has been enabled by two main factors. Firstly, billions of people now own and use smartphones, which enable highly convenient, fast and user-friendly digital conversations. Secondly, and in connection with this, there has been an increased awareness of mass surveillance, both in and especially outside of the security community, spearheaded most notably by Edward Snowden's whistleblowing in 2013 [Gua13]. The documents that Snowden leaked to the public revealed that the National Security Agency (NSA), among other state actors, ran (and undoubtedly continue to run) extensive national and international spying and surveillance programs. Looking ahead, this motivates our consideration of advanced attack vectors like secret state compromise in this thesis [BSJ⁺15].

Secure messaging shares many of the same security requirements as typical communication protocols like TLS. Firstly, communication must be *confidential*, preventing an adversary from deducing anything more from a communication transcript than possibly some metadata. *Authentication* ensures that two (or more) parties can be sure that they are communicating with each other and not some unintended party. Finally, *integrity* ensures that communication cannot be feasibly tampered with by a network adversary without detection.

Messaging nevertheless imposes more requirements than say TLS, motivating specialised protocol design. For one, messaging sessions between two parties are expected to last for

Chapter 1. Introduction

months or years, whereas TLS sessions are typically on the order of minutes. Due to this and the fact that smartphones, the primary medium for instant messaging, have a particularly large attack surface, state *exposure* (or *compromise* [CCG16]) is inevitable for many users. Thus, resilience to state exposure has become a central requirement both in the literature and in practice alike. Moreover, users are not expected to be online at all times, so users should be able to communicate asynchronously as smoothly as possible (thus precluding multi-round protocols for sending and receiving). In popular messaging systems like WhatsApp, an intermediate service (or *central server* for simplicity) manages communication between all parties that in particular forwards messages to users whenever they are online. Some level of plausible deniability, or simply *deniability*, is also sometimes desirable, allowing users to deny having either used a message system or simply having had a particular conversation.¹

For dealing with state exposure in secure messaging, two complementary security notions have become standard in the aftermath of the Snowden revelations. Firstly, *forward security* (FS) [Gün90, BG21] ensures that, upon state exposure, communication *prior* to the exposure is still secure. To achieve FS, secret keying material must somehow be updated over time and new keys replaced with old ones. We will assume in this thesis that secure memory/state erasures are possible since otherwise old keys are always accessible to the adversary upon compromise, preventing FS. Secondly, *post-compromise security* (PCS) [CCG16] ensures that upon state exposure, while security in the present is lost, parties can eventually and automatically *restore* security such that the effect of state compromise is reduced or kept to a minimum. Indeed, some form of PCS has been achieved in practice many years before it was studied in the context of messaging through the common “good practice” of key rotation or re-establishing secure channels after some time or on-demand.² As computing power has advanced significantly over the last few decades, however, it is now possible to cheaply and *continuously* update keying material, enabling strong FS and PCS guarantees. An important aspect of research is thus to identify inherent limitations in achieving FS and PCS and to identify salient performance/security trade-offs for practice.

In the simpler case of *two-party* messaging, we assume both parties, which we will sometimes refer to as Alice and Bob, send and receive messages between each other. Modern messaging protocols are descendants of the seminal Off-The-Record (OTR) protocol [BGB04], which provides FS and PCS through the use of continuous Diffie-Hellman key exchanges. OTR was a significant step forward at the time since popular online messaging applications like MSN Messenger and AOL Instant Messenger provided no end-to-end security guarantees for users. Inspired by OTR, the Double Ratchet [PM16, ACD19] component of the Signal protocol (previously TextSecure [FMB⁺16]) was proposed by Perrin and Marlinspike. It has since become the de-facto standard for messaging [EM19], used by applications including

¹In this thesis, we mainly consider deniability on the *protocol* level, e.g., constructing a deniable key exchange protocol; achieving deniability on the *system* level or in a fully *practical* sense is more complex and sensitive, and to some extent is more of a social or legal matter than a strictly technical one [CCHD23].

²Terms to describe post-compromise security in the literature include future secrecy, backward secrecy and self-healing [DV19].

WhatsApp, Signal, Facebook Secret Conversations and Wire. In a nutshell, the Double Ratchet extends OTR by employing a second ratchet, a forward-secure symmetric ratchet, in addition to the continuous Diffie-Hellman exchange (the *asymmetric* ratchet) present in both protocols.

In group messaging, the group is now dynamic, meaning parties can be added and removed from a given group and users must agree on the group state and messages over time, complicating protocol design and implementation. A naive but nonetheless functional approach here is for all parties to build group messaging on top of two-party channels, which necessarily incurs $O(n)$ sender communication overhead for every message sent in a group of n users. This was, and in some edge cases still is, the approach taken by the Signal application in practice. For increased efficiency, the Sender Keys protocol [Mar14] used by WhatsApp [Wha20] and Signal [M⁺16] enables senders to produce constant-sized ciphertexts that are then forwarded to all group members (via so-called *server fan-out*). In these protocols, refreshing all secrets for post-compromise security incurs $O(n^2)$ communication overhead for a group with n users. While Sender Keys was already deployed (as it still is), MLS (Message Layer Security) was being developed by the IETF in a tight collaboration between industry and academia (and was standardised in 2023 [BBR⁺23]). The MLS protocol, by contrast, achieves logarithmic-sized per-party network overhead in “good” or “fair-weather” cases [ACDT20] for group key updates.³

In either setting, parties must somehow authenticate each other before executing the messaging protocol. This is most commonly achieved by assuming that some public key infrastructure (PKI) is made available to all parties with which they can upload long and short-term keying material. In two-party communication, and sometimes group messaging, authentication is realised via authenticated key exchange (AKE) that leverages the PKI to securely establish initial keying material. Due to the combination of long-term and short-term keys used, key exchange protocols can achieve *forward secrecy*, including many based on the Diffie-Hellman key exchange mentioned above, which ensures that exposing long-term keys does not compromise previously established session keys. Nowadays, many messaging systems use the Extended Triple Diffie-Hellman (X3DH) protocol [MP16b], a descendent of the 3DH protocol [KP05], that provides features including asynchronicity, deniability and forward secrecy, but notably not post-quantum security, i.e., security in the presence of an adversary equipped with a quantum computer.⁴ In group messaging protocols like MLS, rather than performing pairwise key exchange, parties instead upload keying material to a PKI that is directly used for authentication, e.g., when new members are added to a group. In addition, given the threat of state exposure and impersonation attacks that may go undetected, we will argue later on that it is vital to continuously authenticate the conversation transcript itself.

In this thesis, we view secure messaging in the lens of *provable security*. In provable security, a security definition for some primitive or protocol is proposed and its security is reduced to

³This nonetheless degrades to *linear* overhead in the worst case given a dynamic group in which members adaptively update their keying material and the common group key over time.

⁴In 2023, Signal developed and deployed the PQXDH protocol [KS23] that provides security assuming the adversary can perform active attacks using classical computation and passive attacks using quantum computation.

Chapter 1. Introduction

the (conjectured) hardness of mathematical problems like the hardness of factoring or, more generally, to that of abstract cryptographic building blocks. From initial works on provable security in the 1980s onwards, starting from the definition of semantic security [GM82], it has become standard, at least when public-key cryptography is involved, that new constructions are analysed in this framework. Moreover, as the techniques have been more widely understood and developed, increasingly complex protocols and systems are now analysed in terms of provable security. In the context of secure messaging, security definitions here are experiments played between a challenger and an adversary driving protocol execution that can, among other attacks, expose secret keying material during execution. We therefore adopt this framework in order to precisely state and prove security properties about the protocols that we consider in this thesis.

Finally, we observe that works that study secure messaging after authentication often consider either ratcheted (or continuous) key exchange, or secure messaging proper. Throughout this thesis, we focus on either secure messaging or continuous key exchange, depending on context. To build ratcheted key exchange from messaging is straightforward in general as parties can simply sample fresh keys and send them as (encrypted) application messages. One has more freedom designing a protocol in the converse direction; for example, MLS employs a complex key schedule on top of TreeKEM to build group messaging [ACDT21a]. Thus, the two settings are closely related.

1.2 Contribution

In this section, we outline the main contributions of this thesis contained in Chapters 3 to 6. Part I of this thesis, entitled Two-Party Communication, comprises of Chapters 3 and 4, and Part II, entitled Group Messaging, comprises of Chapters 5 and 6. Before these two parts, we introduce in Chapter 2 some notation and notions we use throughout the thesis. After these two parts, we summarise and discuss the findings of this thesis in Chapter 7 before concluding.

1.2.1 Part I: Two-Party Communication

The first part of this thesis consists of Chapter 3, which is concerned with two-party authenticated key exchange suitable for messaging, and Chapter 4, which considers the detection of active attacks in two-party messaging.

K-Waay: Fast and Deniable Post-Quantum X3DH Without Ring Signatures

In Chapter 3, we first consider the problem of authenticated key exchange for bootstrapping two-party channels which is useful for both two-party and group messaging (e.g., for protocols that use two-party channels like Sender Keys). Indeed, Signal's X3DH protocol is widely used in practice and satisfies several unique requirements that make it suitable for messaging.

Apart from *asynchronicity*, it provides a form of *deniability* [VGIK20], which in the case of X3DH allows a party, who is being framed, to deny having participated in a conversation with their counterpart who is framing them. However, since X3DH essentially consists of several parallel Diffie-Hellman key exchanges, it does not provide protection against a passive polynomial-time quantum adversary, and Signal’s recent PQXDH protocol does not protect against an *active* quantum adversary.

To rectify this, Brendel et al. [BFG⁺22b] and Hashimoto et al. [HKKP21, HKKP22] concurrently proposed alternate X3DH protocols that protect against even active quantum attackers. These protocols rely on two-user ring signatures [RST01], or equivalently [HKKP22], a form of designated-verifier signatures, to achieve authentication while ensuring deniability. At present, post-quantum instantiations of these primitives are more complex and their security is less well-understood than primitives like the CRYSTALS-Kyber key encapsulation mechanism (KEM) [BDK⁺18] recently standardised by the United States National Institute of Standards and Technology (NIST) government agency. For example, existing works that rely on random oracles [BR93] tend not to contain proofs in the *quantum* random oracle model [BDF⁺11], i.e., a random oracle that the adversary can query in superposition. In addition, proposed parameters for concrete instantiations tend not to be large enough since they tend not to take the security loss of the security proofs into account as we observe in the chapter.

We propose instead to use a slightly modified variant of the Frodo key exchange protocol [BCD⁺16] to achieve deniable, post-quantum authentication. As such, its security is based on the vanilla learning-with-errors (LWE) assumption, which is more conservative than using structured lattices [ANS23]; for this same reason, the German agency BSI chose to recommend FrodoKEM, a KEM derived from Frodo key exchange, over CRYSTALS-Kyber [BSI23]. In particular, structured lattices may very well have exploitable weaknesses due to their additional mathematical structure [Tea21]. To capture this primitive generically, we significantly extend the security requirements of the split-KEM primitive introduced by Brendel et al. [BFG⁺20] as their security notions, as we show, are insufficient for building (deniable) authenticated key exchange. In more detail, a split-KEM extends KEM in that encapsulation now requires both a secret key and a public key, and decapsulation uses the corresponding public and secret keys respectively. We introduce an unforgeability notion for split-KEM, which captures the difficulty of crafting a ciphertext without any secrets, as well as a notion for deniability, capturing that split-KEM encapsulation computationally looks the same using either the intended secret and public keys *or* the decapsulation keying material as input.

We formalise an X3DH-like AKE protocol that we call K-Waay that uses split-KEM, ephemeral and long-term KEMs, and signatures which parties use to attest to ephemeral KEM and split-KEM public keys uploaded to a server to be used asynchronously by a given party’s counterpart. Our AKE security model, like that of Hashimoto et al. [HKKP21, HKKP22], supports ephemeral state exposure, although it is weaker than theirs since only the ephemeral-ephemeral portion of the key exchange (that overall uses both long-term and ephemeral keys) relies on split-KEM. As we do not model reusable semi-static keys directly, we propose a novel countermeasure in

Chapter 1. Introduction

the (quantum) random oracle model to allow for ephemeral key reuse that prevents ephemeral key exhaustion under adversarial or otherwise extreme conditions. Our deniability notion for AKE strengthens that of Brendel et al. [BFG⁺22b] by additionally giving the judge or adversary the ephemeral state of the receiver. We also benchmark our key exchange protocol and demonstrate even with our conservative choice of split-KEM instantiation and parameters which capture the security loss from the proofs, that K-Waay is between 4x and 6x faster than existing ring signature-based approaches.

On Active Attack Detection in Messaging with Immediate Decryption

In Chapter 4, we then consider two-party messaging proper where parties are assumed to have successfully completed key exchange (e.g., by using K-Waay). A notable property of Signal’s Double Ratchet core is that it supports *immediate decryption* as formalised by Alwen et al. [ACD19] and achieved by a number of recent academic protocols [ACD19, PP22, BRT23, CZ24]. That is, protocol execution should natively support receiving messages in different orders to what they were sent in, as well as fully dropped ciphertexts. This is especially important in challenging network settings like mesh networking [BRT23] and can also be useful for efficiency, e.g., the messaging service provider can choose to buffer sent videos for users.

Given the likelihood of state exposure, active attacks, where the adversary uses the state of a compromised party and impersonates them, pose a significant threat. Messaging systems, including Signal via its safety numbers protocol [Mar17], allow parties to compare their *long-term* keys out-of-band via QR code or manual comparison. Hence, given the out-of-band channel provides authenticity, this allows parties to audit the PKI and detect active attacks affecting these keys. In the event of state compromise, however, the adversary can mount an active attack and inject application messages, in which case there is no cryptographic mechanism in place to detect it. Moreover, without an additional assumption or channel, it is not possible to detect active attacks, since the adversary can block all honest messages sent by a compromised party and replace them with its own. To alleviate this, some solutions based on out-of-band communication have been previously proposed, but either do not fully authenticate the conversation transcript [DH21] or have other drawbacks like requiring multiple rounds of communication and not formally capturing immediate decryption [DGP22].

In contrast, to detect attacks *in-band*, that is, over the same channel as is used for communication, *RECOVER security* was defined by Durak and Vaudenay [DV19] and later extended by Caforio et al. [CDV21]. RECOVER security allows parties to detect that, if a single honest message is able to get through after compromise, that active attacks *can* be detected, even if the adversary has full knowledge of secret states. More precisely, RECOVER security encompasses two complementary notions. First, r-RECOVER guarantees that if Bob receives a forgery, Bob will reject all subsequent honest messages. Second, s-RECOVER guarantees that if Bob receives a forgery, his communication partner Alice will reject all messages sent *by* Bob after he received the forgery. However, their notions and constructions are not compatible

with immediate decryption, and so cannot be readily adopted by systems like Signal.

To remedy these issues, we systematically explore active attack detection for messaging with immediate decryption in both the in-band and out-of-band settings. For RECOVER security, we generalise the two notions above to support immediate decryption, and propose a baseline construction that compiles a messaging scheme into one that is secure under both of our notions. We also propose two analogous security notions assuming parties have access to an authentic out-of-band channel, and construct secure schemes in two ways: first, from a RECOVER-secure messaging scheme, incurring no additional overhead, and secondly, from any messaging scheme. We then show that in order to satisfy our r-RECOVER notion or the analogous out-of-band notion, the ciphertext size must grow linearly in the number of messages sent (and security parameter) assuming unidirectional communication. Our information-theoretic argument implies that our baseline construction, which attaches all sent ciphertexts to each ciphertext to achieve r-RECOVER security, is asymptotically optimal. To bypass this negative result, we argue that s-RECOVER security is comparatively cheaper to achieve and can be sufficient for practice as it still provides r-RECOVER-like guarantees (albeit after an honest round-trip). We propose a series of optimisations, where ultimately the communication overhead of s-RECOVER security is a single hash digest and a few indices (in proportion to how synchronised parties are). Since s-RECOVER security effectively provides ‘delayed’ r-RECOVER guarantees, we believe that this solution is affordable enough to be adopted in practice by systems like Signal.

1.2.2 Part II: Group Messaging

The second part of this thesis moves to group messaging: Chapter 5 considers the effects of group administration on security, and Chapter 6 explores the Sender Keys protocol used by WhatsApp and Signal.

Cryptographic Administration for Secure Group Messaging

In Chapter 5, we consider a problem that only arises in group messaging, namely *group administration*. In practice in a group conversation, one or more group members are generally afforded more privileges than others. Systems like Telegram offer admins diverse, fine-grained control over the conversation flow, including disallowing certain types of messages from being sent. Fundamentally, though, the core of administration, which is what we target here, is to enforce *access control*. More precisely, we consider a privileged set of users, the group *administrators* (or admins), who solely have the ability to add and remove users from the conversation.

We observe that many practical protocols do not provide adequate protection for access control, and in particular the ability to enforce group administration. The *burgle into the group* attack [RMS18] allows an adversary to trivially forge messages that add and remove parties

Chapter 1. Introduction

from a group, thereby allowing them to completely bypass any cryptography used. This issue notably affected WhatsApp in particular without the need for any state compromise.⁵ Similar attacks have been reported for the federated messaging system Matrix [ACDJ23, ADJ24].

In this chapter, we aim to formalise group administration, which provides a natural mechanism to prevent these types of attacks and enable cryptographic access control within a group. To this end, we start from the recent continuous group key agreement (CGKA) primitive [ACDT20]. In CGKA, a sequence of secrets that can be used for messaging are established over time by a dynamic group of users. CGKA captures the core of group messaging, and TreeKEM, which can be cast as a CGKA [ACDT20, ACDT21a], forms the basis of the recently standardised MLS protocol and largely determines its performance. In addition, MLS is expected to be increasingly adopted in the future [Hog23, Gie23], and MLS and CGKA have received much academic attention in recent years [ACDT20, ACJM20, KPPW⁺21, ACDT21a, HKP⁺21, AAN⁺22b, AJM22, BDG⁺22, AHKM22, AMT23]. We therefore consider group administration through the lens of CGKA. We observe that neither previous CGKA protocols, nor existing work on group messaging in general, captured group administration at the cryptographic level. The closest we are aware of here is the work of Rösler et al. [RMS18] which considers group administration as a security requirement but does not adopt the provable security methodology. Consequently, in CGKA protocols in the literature, this renders all group members formally as group administrators with the same rights.

In this chapter, we introduce *administrated CGKA*, or A-CGKA, which natively captures administration and supports the adding, removing and updating of administrators (the latter being for forward security and post-compromise security). Namely, we extend the (unauthenticated) CGKA security notion of Alwen et al. [ACDT20] to additionally ensure that only admins can change the group membership and update admin secrets, even if several non-admin users are compromised, with the exception that we allow non-admins to update their own keys and remove themselves. We provide two constructions with different efficiency and security characteristics that build on top of a CGKA. Our first construction, Individual Admin Signatures (IAS), associates each admin with an evolving signature key pair and requires messages for admin-exclusive operations to be accompanied by admin signatures. Our second construction, Dynamic Group Signature (DGS), associates the set of admins with a secondary CGKA, or *admin CGKA*, following the observation that administration need not be exercised by individually authenticated administrators. In particular, the admin CGKA is used to manage the set of admins, where admins use the dynamic CGKA secret to derive common signature key-pairs that attest to admin operations. The composition of CGKAs employed in DGS may be of independent interest more generally for managing hierarchies in large-scale messaging. We then formally prove our two protocols secure. Moving from CGKA to group messaging, we detail how to integrate IAS-like administration into MLS by leveraging the existing credential infrastructure in MLS. By extending an existing implementation of MLS, we benchmark the

⁵We are not aware of any information indicating that the attack has been fixed in WhatsApp [RMS18]. Signal, for which state exposure was required to mount the attack, has since migrated to a newer private group management system [CPZ20] that is immune to it.

resulting protocol and demonstrate that it (as expected) incurs minimal overhead over MLS without administration. We finally discuss several directions for future work in both theory and practice for group administration.

WhatsUpp with Sender Keys? Analysis, Improvements and Security Proofs

In Chapter 6, the final main chapter of this manuscript, we focus on the group messaging protocol *Sender Keys* [Mar14] used by applications like WhatsApp and Signal. The core of Sender Keys is conceptually simple and works as follows. Each group member is associated with a *sender key*, which consists of a symmetric chain key ck and a signature public/secret key pair (spk, ssk) . All group members learn each ck and spk corresponding to each group member via two-party channels (maintained by each pair of group members). To send an application message, a user hashes ck to derive a pair (ck', mk) , encrypts using mk , signs the ciphertext and sends it alongside the signature. For forward security, they also delete mk and replace ck with ck' . Fresh sender keys are shared over two-party channels, and in particular all sender keys are refreshed whenever a group member is removed from a group.

Despite its practical importance, the protocol had not previously been formally modelled or described in depth. Rösler et al. [RMS18] previously analysed WhatsApp's implementation of Sender Keys which revealed flaws like the aforementioned burgle into the group attack, but did not provide a full protocol description or general characterisation of its security. While being more complex, MLS and CGKA protocols have seen more academic attention despite not yet being widely deployed. Sender Keys nonetheless remains as a practical alternative to MLS: WhatsApp is able to support billions of daily users and supports up to 1024 users per group [Wha23]. In addition, while offering logarithmic-sized updates in “good” cases, sufficiently “bad” and concurrent executions can lead to $O(n)$ -sized update messages, which is inherent to CGKA at least when using off-the-shelf primitives [BDG⁺22].

Towards formalising Sender Keys, we first propose a primitive syntax and security model suitable for the protocol. Our security model parametrises the security of the underlying two-party channels, which allows us to observe that Sender Keys offers weaker than expected post-compromise security guarantees. In particular, key refreshes after compromise may not be secure due to the channels not having fully recovered, which can limit security in practice if the channels remain stale. We provide a full protocol description based on our understanding of Sender Keys mainly from WhatsApp's white paper [Wha20] and Signal's source code [M⁺16]. Our protocol also captures key updates without group changes that are implemented by Signal. We then prove security in our model, in which we have to notably restrict the adversary due to inherent limitations in Sender Keys. In this, we identify its deficiencies, including insecure group membership (affecting WhatsApp and to a lesser extent Signal in practice) and a lack of forward-secure authentication.

To remedy these security issues, we then propose and prove secure a new protocol that we call *Sender Keys+* that improves security without significantly affecting performance, and in

Chapter 1. Introduction

particular improves both the efficiency and security of key updates. By comparison to our baseline Sender Keys protocol, Sender Keys+ now uses signatures to attest to group changes and adds a message authentication code (or replaces encryption with an AEAD [Rog02]) in sending and receiving. For updates, first observe that fully recovering from compromise requires all sender keys to be refreshed, which incurs $O(n^2)$ communication for a group with n members, since each user needs to send a (constant-sized) sender key to all $O(n)$ users over two-party channels. We observe that, given sufficient synchronisation of parties, that to update, a user can sample a secret that is then hashed into each chain key, i.e., can update *all* sender keys at once, which restores confidentiality (assuming restored two-party channels) in just $O(n)$ communication (which is optimal). We also provide a theoretical comparison between Sender Keys(+) and CGKA-based protocols, from which we can conclude that Sender Keys, and especially Sender Keys+, may be sufficiently secure and performant for many practical deployments.

1.3 List of Works

The following works were produced while the author of this manuscript was working at EPFL. Works in **bold** appear in the body of this thesis. Full versions or preprints are also cited where relevant.

1. **Daniel Collins, Loïs Huguenin-Dumittan, Ngoc Khanh Nguyen, Nicolas Rolin, and Serge Vaudenay. K-Waay: Fast and deniable post-quantum X3DH without ring signatures. *USENIX Security 2024*. [CHDN⁺24a, CHDN⁺24b]**
2. **David Balbás, Daniel Collins, and Phillip Gajland. WhatsApp with Sender Keys? Analysis, improvements and security proofs. *ASIACRYPT 2023*. [BCG23b, BCG23a]**
3. **Khashayar Barooti, Daniel Collins, Simone Colombo, Loïs Huguenin-Dumittan, and Serge Vaudenay. On active attack detection in messaging with immediate decryption. *CRYPTO 2023*. [BCC⁺23a, BCC⁺23b]**
4. Renas Bacho, Daniel Collins, Chen-Da Liu-Zhang, and Julian Loss. Network-agnostic security comes (almost) for free in DKG and MPC. *CRYPTO 2023*. [BCLZL23, BCLZL22]
5. **David Balbás, Daniel Collins, and Serge Vaudenay. Cryptographic administration for secure group messaging. *USENIX Security 2023*. [BCV23, BCV22]**
6. Subhadeep Banik, Daniel Collins, and Willi Meier. Near collision attack against Grain V1. *ACNS 2023*. [BCM23a, BCM23b]
7. Daniel Collins, Simone Colombo, and Loïs Huguenin-Dumittan. Real world deniability in messaging. *Cryptology ePrint Archive, Paper 2023/403*. [CCHD23]⁶
8. Andrea Caforio, Daniel Collins, Subhadeep Banik, and Francesco Regazzoni. A small GIFT-COFB: Lightweight bit-serial architectures. *AFRICACRYPT 2022*. [CCBR22a, CCB22b]
9. Hailun Yan, Serge Vaudenay, Daniel Collins, and Andrea Caforio. Optimal symmetric ratcheting for secure communication. *The Computer Journal*, 2021. [YVCC23]
10. Andrea Caforio, Daniel Collins, Ognjen Glamočanin, and Subhadeep Banik. Improving first-order threshold implementations of SKINNY. *INDOCRYPT 2021*. [CCGB21a, CCGB21b]
11. Daniel Collins, Rachid Guerraoui, Jovan Komatovic, Petr Kuznetsov, Matteo Monti, Matej Pavlovic, Yvonne-Anne Pignolet, Dragos-Adrian Seredinschi, Andrei Tonkikh, and Athanasios Xytkis. Online payments by merely broadcasting messages. *DSN 2020* [CGK⁺20a, CGK⁺20b]

⁶This extended abstract was accompanied by a presentation at the IACR conference RWC 2023.

2 Preliminaries

In this chapter, we present some preliminary material that is used in subsequent chapters in this thesis. We first introduce some choices we make for notation and conventions, particularly describing how we define security. We then introduce some cryptographic primitives that either appear in several chapters or are otherwise suited to be defined here.

2.1 Notation and Conventions

We denote the empty string by ε and a blank value by \perp , the latter which we return in case of algorithm failure. Given a set S , S^* (respectively S^n) is the set of all strings of arbitrary length (resp. of length n) whose elements are in S . For two sets S and T , let $S \overset{\cup}{\leftarrow} T$ denote the reassignment of S to the set $S \cup T$, and let $S \overset{-}{\leftarrow} T$ denote the reassignment of S to the set $S \setminus T$. Let $[n] = \{1, \dots, n\}$, i.e., the set of integers between 1 and n . Let $[a, b] = \{a, \dots, b\}$, i.e., the set of integers between a and b . We consider randomised algorithms throughout the thesis: we let \mathcal{R} denote a randomness space. We write $s \overset{\$}{\leftarrow} S$ to denote uniformly sampling s from S , and $y \overset{\$}{\leftarrow} f(x)$ to denote running randomised algorithm f with input x and uniform randomness. Alternatively, we write $y \leftarrow f(x; r)$ to denote running f with fixed randomness r . Let 1^λ be the security parameter (here stated in unary), and $\text{poly}(\lambda)$ denotes a polynomial in λ . We denote the indicator function with respect to boolean P by $\mathbb{1}[P]$, which is equal 1 if P is true and 0 otherwise.

Alice and Bob. A *user*, *participant*, or *party* is an entity that takes part in a protocol. In Part I of this thesis (Chapters 3 and 4), we consider primitives that are executed between two parties. As is common, we refer these two parties as Alice and Bob, or A and B . Let P be one party (A resp. B) and \bar{P} be their partner (B resp. A).

In Part II of this thesis (Chapters 5 and 6), we will consider primitives that deal with more than two parties. Here, users are identified by a unique identity string ID , which is a public value. Groups are also uniquely identified by a public group identifier gid . We assume these are known by and agreed upon by all parties for simplicity. The core primitives of study in

Chapter 2. Preliminaries

these chapters will be stateful, thus users keep an internal *state* γ with all information used for protocol execution. This may include keys, message records, dictionaries and parameters, among other values.

Security. To capture the security of cryptographic primitives and protocols, we work within the framework of *provable security*. In provable security, a security notion is defined as an experiment played between an adversary, often \mathcal{A} , and a challenger, say C , in which \mathcal{A} is tasked to perform some task, such as outputting a forged signature or more generally breaking some cryptographic primitive in some experiment-dependent sense. The *advantage* of \mathcal{A} in the experiment will be stated as a probability that is taken over the randomness of the challenger and adversary. In general, we say a scheme is *secure* if the advantage of any *efficient* adversary is *small* (we define efficient and small below). Such a claim is proven by reduction to hardness assumptions (e.g., the hardness of factoring integers) and/or the security of some cryptographic building blocks (e.g., SUF-CMA security of a signature scheme).

We are interested in both classical and quantum security in this thesis, i.e., security against a classical or quantum adversary. We denote by an *efficient* adversary a probabilistic polynomial-time (PPT) algorithm or quantum polynomial-time (QPT) algorithm, unless explicitly stated otherwise, where we mean polynomial in the security parameter λ . Our theorem statements will therefore hold insofar as the underlying assumptions hold and/or building blocks are secure with respect to a classical or quantum adversary. We consider an advantage to be *small* if it is *negligible* in λ . A function $f : \mathbb{N} \rightarrow \mathbb{R}$ is *negligible* if for every positive c , there exists integer N_c such that for all $x > N_c$, $|f(x)| < 1/x^c$ (for our purposes f outputs a probability). We let negl be a negligible function (in λ , which we omit writing and is implicit). We nonetheless state our security bounds exactly when proving security, i.e., the advantage of the scheme being proven secure will be upper bounded by some explicitly stated function of advantages of other primitives/assumptions and other terms.

Maps. We use maps, or associative arrays, which associate keys with values: $m[\cdot] \leftarrow x$ defines a new map with values initially set to x and $m[k]$ returns the element indexed by key k . Keys are tuples of any length $n \geq 1$. We index maps with integers starting from 1; in this case, $m[a : b]$ returns the list of elements whose keys are between a and b . We access the element of a tuple using a dot notation, e.g., $\gamma.k$ denotes the value k in the tuple γ . The function $\text{length}(m)$ returns the number of keys in map m . All dictionaries can optionally be indexed by an oracle query q to represent the state of a dictionary at the time q is made, e.g., $E[ID; q]$ denotes the value of $E[ID]$ at the beginning of query q .

Keywords. In oracles and algorithms, some special predicates are used. The predicate ‘**require** P ’ enforces that a logical condition P is satisfied; otherwise the oracle/algorithm finishes immediately and returns \perp . The predicate ‘**reward** P ’ is executed in games and is such that if P holds, the adversary satisfies a winning condition in a game. Namely, in games where the advantage is defined by the probability that the adversary outputs 1, a game variable win is set to 1, and in indistinguishability games the game reveals the bit $b \in \{0, 1\}$ to the adversary.

The keyword ‘**public var**’ indicates that the adversary has read access to the variable `var`. The prefix operator $\text{Alg}(++x)$, is equivalent to writing first $x \leftarrow x + 1$ and then $\text{Alg}(x)$. Algorithms, oracle names and cryptographic parameters are denoted in sans-serif font.

(Quantum) Random Oracle Model. Some of our results will hold in the *random oracle model* (ROM) [BR93] or the *quantum random oracle model* (QROM) [BDF⁺11]. A random oracle is a uniformly random function, say H , associated with some input and output space. In the ROM or QROM, the adversary’s queries are visible to whichever entity (the challenger or some other adversary, for example) is controlling the random oracle. In the QROM, the adversary can query the random oracle in superposition. That is, the adversary provides a quantum state as input, say $|\phi\rangle = \sum \alpha_x |x\rangle$, and the quantum random oracle returns the evaluated state, in this case $\sum \alpha_x |H(x)\rangle$.

2.2 Cryptographic Primitives

We introduce some core cryptographic primitives, including their syntax as well as correctness and security notions where appropriate, which we invoke throughout this thesis as needed.

2.2.1 Key-Encapsulation Mechanism (KEM)

Definition 1 (KEM). A KEM KEM is a tuple of three efficient algorithms (KeyGen , Encaps , Decaps) defined as follows:

- $(\text{pk}, \text{sk}) \xleftarrow{\$} \text{KeyGen}(1^\lambda)$: The key generation function takes the security parameter λ as input, and outputs a pair of public/secret keys (pk, sk) .
- $K, \text{ct} \xleftarrow{\$} \text{Encaps}(\text{pk})$: The encapsulation function takes a public key pk as input, and outputs a ciphertext ct and a key K .
- $K/\perp \leftarrow \text{Decaps}(\text{sk}, \text{ct})$: The decapsulation function takes a secret key sk and a ciphertext ct as inputs, and outputs a key K or the error symbol \perp .

Finally, we say a KEM is $(1 - \delta)$ -correct if

$$\Pr \left[\begin{array}{l} (\text{pk}, \text{sk}) \xleftarrow{\$} \text{KeyGen}(1^\lambda); \\ K \neq K' : \quad K, \text{ct} \xleftarrow{\$} \text{Encaps}(\text{pk}); \\ \quad \quad \quad K' \leftarrow \text{Decaps}(\text{sk}, \text{ct}) \end{array} \right] \leq \delta .$$

Definition 2 (KEM Indistinguishability). We consider the games defined in Figure 2.1. Let \mathcal{K} be a finite key space. A KEM scheme over \mathcal{K} $\text{KEM} = (\text{KeyGen}, \text{Encaps}, \text{Decaps})$ is *IND-CPA* (resp. *IND-CCA*) if for any efficient adversary \mathcal{A} with no access to the decapsulation oracle

Chapter 2. Preliminaries

(respectively any efficient adversary \mathcal{A}) we have

$$\text{Adv}_{\text{KEM}}^{\text{ind-cpa/cca}}(\mathcal{A}) := \left| \Pr[\text{IND-CPA/CCA}_{\text{KEM}}(\mathcal{A}) \Rightarrow 1] - \frac{1}{2} \right| = \text{negl}.$$

where $\Pr[\text{IND-CPA/CCA}_{\text{KEM}}(\mathcal{A}) \Rightarrow 1]$ is the probability that \mathcal{A} wins the IND-CPA/CCA_{KEM}(\mathcal{A}) game defined in Figure 2.1.

Game IND-CPA/CCA _{KEM} (\mathcal{A})	Oracle DEC(ct)
1: $b \xleftarrow{\$} \{0, 1\}$	1: if ct = ct* : return \perp
2: $(pk, sk) \xleftarrow{\$} \text{KeyGen}(1^\lambda)$	2: $K' \leftarrow \text{Decaps}(sk, ct)$
3: $(ct^*, K_0) \xleftarrow{\$} \text{Encaps}(pk)$	3: return K'
4: $K_1 \xleftarrow{\$} \mathcal{K}$	
5: $b' \xleftarrow{\$} \mathcal{A}^{\text{DEC}}(pk, ct^*, K_b)$	
6: return $\mathbb{1}[b' = b]$	

Figure 2.1: Indistinguishability games for KEM. In the IND-CPA game, the adversary cannot query DEC.

2.2.2 Signatures

Definition 3. A signature scheme is a tuple of three efficient algorithms (KeyGen, Sign, Vrfy):

- $(pk, sk) \xleftarrow{\$} \text{KeyGen}(1^\lambda)$: The key generation function outputs a pair of keys.
- $\sigma \xleftarrow{\$} \text{Sign}(sk, m)$: The signing function takes as inputs a secret key sk and the message to sign m , and it outputs a signature σ .
- $0/1 \leftarrow \text{Vrfy}(pk, m, \sigma)$: The verification function takes as inputs a public key pk , the signed message m , and the signature σ , and it outputs either 0 or 1 (for failure and success, respectively).

Finally, we say a signature scheme is $(1 - \delta)$ -correct if for all messages m :

$$\Pr \left[\text{Vrfy}(pk, m, \sigma) = 0 : \begin{array}{l} (pk, sk) \xleftarrow{\$} \text{KeyGen}(1^\lambda); \\ \sigma \leftarrow \text{Sign}(sk, m) \end{array} \right] \leq \delta$$

Definition 4 (SUF-CMA Security). We consider the game shown in Figure 2.2. We say a signature scheme Sig is *SUF-CMA* if for all efficient adversaries \mathcal{A} , we have

$$\text{Adv}_{\text{Sig}}^{\text{suf-cma}}(\mathcal{A}) := \Pr[\text{SUF-CMA}_{\text{Sig}}(\mathcal{A}) \Rightarrow 1] = \text{negl}.$$

Game $\text{SUF-CMA}_{\text{Sig}}(\mathcal{A})$	Oracle $\text{SIGN}(m)$
1: $L \leftarrow \emptyset$	1: $\sigma \xleftarrow{\$} \text{Sign}(\text{sk}, m)$
2: $(\text{pk}, \text{sk}) \xleftarrow{\$} \text{KeyGen}(1^\lambda)$	2: $L \leftarrow L \cup \{(m, \sigma)\}$
3: $m^*, \sigma^* \xleftarrow{\$} \mathcal{A}^{\text{SIGN}}(\text{pk})$	3: return σ
4: return $\mathbb{1}[\text{Vrfy}(\text{pk}, m^*, \sigma^*) \text{ and } (m^*, \sigma^*) \notin L]$	

Figure 2.2: SUF-CMA game.

2.2.3 Pseudorandom Function (PRF)

We introduce security notions for a pseudorandom function (PRF) and two generalisations, namely dual PRF and triple PRF. For these, we consider functions of the form $F: \mathcal{K}^\ell \times D \rightarrow R$, i.e., that map $\ell \geq 1$ keys in \mathcal{K} and a value in a domain D to a value in a range R . For such a function, we are interested in the security of F as a PRF when the i -th argument acts as the PRF key for all $i \in [1, \ell]$. Thus, we will consider the ℓ inputs to the function from the perspective of the caller as the $\ell - 1$ other values in \mathcal{K} and the D inputs, and use the notation F_i as shorthand to indicate this. In fact, our notions match Giacon et al.'s formulation of a *split-key PRF* [GHP18].

Definition 5 (PRF). Let $F: \mathcal{K} \times D \rightarrow R$ be a function. We consider the game shown in Figure 2.3. We say that F is a secure *PRF* if, for all efficient adversaries \mathcal{A} , we have

$$\text{Adv}_F^{\text{prf}}(\mathcal{A}) := \left| \Pr[\text{PRF}_F(\mathcal{A}) \Rightarrow 1] - \frac{1}{2} \right| = \text{negl}.$$

Definition 6 (Dual PRF). Let $F: \mathcal{K} \times \mathcal{K} \times D \rightarrow R$ be a function. We consider the game shown in Figure 2.3. We say that F is a secure *dual PRF* or *2PRF* if, for all efficient adversaries \mathcal{A} , we have

$$\text{Adv}_F^{2\text{prf}}(\mathcal{A}) := \max_{i \in \{1, 2\}} \left| \Pr[\text{PRF}_{F_i}(\mathcal{A}) \Rightarrow 1] - \frac{1}{2} \right| = \text{negl}.$$

where F_i denotes F keyed in its i -th argument for $i \in \{1, 2\}$.

Definition 7 (Triple PRF). Let $F: \mathcal{K} \times \mathcal{K} \times \mathcal{K} \times D \rightarrow R$ be a function. We consider the game shown in Figure 2.3. We say that F is a secure *triple PRF* or *3PRF* if, for all efficient adversaries \mathcal{A} , we have

$$\text{Adv}_F^{3\text{prf}}(\mathcal{A}) := \max_{i \in \{1, 2, 3\}} \left| \Pr[\text{PRF}_{F_i}(\mathcal{A}) \Rightarrow 1] - \frac{1}{2} \right| = \text{negl}.$$

where F_i denotes F keyed in its i -th argument for $i \in \{1, 2, 3\}$.

Observe the difference between our notion of dual PRF and the dual PRF introduced by Bellare [Bel06], which is a function $F: K_1 \times K_2 \rightarrow R$ and PRF when keyed from either K_1 or K_2 (the latter referred to as the *swap PRF* case).

Game $\text{PRF}_F(\mathcal{A})$	Oracle $\text{EVAL}(a_1, \dots, a_\ell)$
1: Sample random function $G: \mathcal{K}^{\ell-1} \times D \rightarrow R$	1: require $a_\ell \notin X$
2: $k \xleftarrow{\$} \mathcal{K}$	2: $X \leftarrow X \cup \{a_\ell\}$
3: $b \xleftarrow{\$} \{0, 1\}$	3: if $b = 0$
4: $b' \xleftarrow{\$} \mathcal{A}^{\text{EVAL}}$	4: return $F_k(a_1, \dots, a_\ell)$
5: $X \leftarrow \emptyset$	5: else
6: return $\mathbb{1}[b' = b]$	6: return $G(a_1, \dots, a_\ell)$

Figure 2.3: PRF game for function F_k taking ℓ arguments a_1, \dots, a_ℓ as input where $a_1 \in \mathcal{K}, \dots, a_{\ell-1} \in \mathcal{K}$ and $a_\ell \in D$.

2.2.4 Hash Function

Definition 8. A hash function H consists of efficient algorithms KGen and Eval such that:

- $\text{hk} \xleftarrow{\$} \text{KGen}(1^\lambda)$: The key generation function outputs a hash key hk .
- $h \leftarrow \text{Eval}(\text{hk}, \text{pt})$: The evaluation function takes as inputs a hash key hk and a message $\text{pt} \in \{0, 1\}^*$ and outputs digest h .

Definition 9 (Collision Resistance). We consider the game in Figure 2.4. We say that a hash function H is collision resistant if, for all efficient adversaries, we have

$$\text{Adv}_H^{\text{cr}}(\mathcal{A}) := \Pr[\text{CR}_H(\mathcal{A}) \Rightarrow 1] = \text{negl}.$$

Game $\text{CR}_H(\mathcal{A})$
1: $\text{hk} \leftarrow \text{KGen}(1^\lambda)$
2: $(m_1, m_2) \leftarrow \mathcal{A}(\text{hk})$
3: return $\mathbb{1}[H.\text{Eval}(\text{hk}, m_1) = H.\text{Eval}(\text{hk}, m_2) \wedge m_1 \neq m_2]$

Figure 2.4: Collision resistance of a hash function H .

2.2.5 Incremental Set Hash Function

Clarke et al. [CDv⁺03] define incremental *multiset* hash functions and multiset collision resistance. That is, a hash function which takes a set of elements as input, where the digest can be updated with an operation with complexity proportional to the number of elements added/removed. For our purposes (i.e., in Chapter 4), it suffices to consider an *incremental set hash function* and *set collision resistance*.

Definition 10. An incremental set hash function H consists of the following efficient algorithms:

- $hk \xleftarrow{\$} \text{IncGen}(1^\lambda)$: The key generation function outputs a hash key hk .
- $h_S \leftarrow \text{IncEval}(hk, S = \{m_1, \dots, m_k\})$: The evaluation function for a set of messages takes a hash key hk and set S and outputs digest h_S .
- $h_{S \cup S_h} \leftarrow \text{IncEval}(hk, h, S_h = \{m'_1, \dots, m'_k\}, S = \{m_1, \dots, m_k\})$: The evaluation function for a set of messages and digest takes as input a hash key hk , digest h , set S_h (associated with h) and set S and outputs a digest $h_{S \cup S_h}$.

An incremental hash function is *correct* if for all $hk \xleftarrow{\$} \text{IncGen}(1^\lambda)$ and non-empty $S \leftarrow \{m_1, \dots, m_k\}$ it holds that

$$\text{IncEval}(hk, S) = h',$$

where $h' \leftarrow \text{IncEval}(hk, h, S_h, S')$, $S_h \cup S' = S$, $S_h \cap S' = \emptyset$ and h is the result of either 1) a call $\text{IncEval}(hk, S_h)$ or 2) a call $\text{IncEval}(hk, S_1)$ and one or more calls to $\text{IncEval}(hk, \cdot, S_2, S_3)$ for disjoint sets $S_1, S_2, S_3 \subset S_h$ s.t. each pair of sets S_2 and S_3 is strictly increasing in cardinality and the last call is s.t. $S_2 \cup S_3 = S_h$.

Definition 11 (Set Collision Resistance). Consider the game defined in Figure 2.5. A family of incremental set hash function H is set collision resistant, if for any efficient adversary, we have

$$\text{Adv}_H^{\text{SCR}}(\mathcal{A}) := \Pr[\text{SCR}_H(\mathcal{A}) \Rightarrow 1] = \text{negl}.$$

Game $\text{SCR}_H(\mathcal{A})$
1: $hk \xleftarrow{\$} H.\text{IncGen}(1^\lambda)$
2: $(S_1 = \{m_i\}_i, S_2 = \{m'_j\}_j) \xleftarrow{\$} \mathcal{A}(hk)$
3: return $\mathbb{1}[\![H.\text{IncEval}(hk, S_1) = H.\text{IncEval}(hk, S_2) \wedge (S_1 \neq S_2)\!]\!]$

Figure 2.5: Set collision resistance of an incremental set hash function H .

We refer the reader to Clarke et al. [CDv⁺03] for more details about incremental set hash functions.

2.2.6 Symmetric Encryption

Definition 12 (Symmetric Encryption). A symmetric encryption scheme $\text{SymEnc} := (\text{Gen}, \text{Enc}, \text{Dec})$ is a tuple of efficient algorithms such that:

- $k \xleftarrow{\$} \text{Gen}(1^\lambda)$: Given the security parameter 1^λ (encoded in unary) the generation algorithm returns a key $k \in \mathcal{K}$.
- $c \xleftarrow{\$} \text{Enc}(k, m)$: Given a key k and a message m , the encryption algorithm returns a ciphertext c .

Chapter 2. Preliminaries

- $m \leftarrow \text{Dec}(k, c)$: Given a key k and a ciphertext c , the decryption algorithm returns a message m .

For simplicity of exposition, we model Enc as probabilistic rather than explicitly modelling input initialisation vectors (we therefore implicitly assume they are sampled randomly and appended to the ciphertext where relevant). We say that SymEnc is correct if for any message $m \in \mathcal{M}$ and any key $k \in \mathcal{K}$ it holds that,

$$\Pr \left[\text{Dec}(k, c) = m \mid \begin{array}{l} k \stackrel{\$}{\leftarrow} \text{Gen}(1^\lambda) \\ c \stackrel{\$}{\leftarrow} \text{Enc}(k, m) \end{array} \right] = 1,$$

where the probability is taken over the random coins of Enc , and \mathcal{M} and \mathcal{K} denote the message space and key space respectively.

Definition 13 (SymEnc IND-CPA Security). We consider the game in Figure 2.6. We say that a symmetric encryption scheme SymEnc is secure under chosen plaintext attacks, or is IND-CPA, if, for all efficient adversaries \mathcal{A} , we have

$$\text{Adv}_{\text{SymEnc}}^{\text{ind-cpa}}(\mathcal{A}) := \left| \Pr[\text{IND-CPA}_{\text{SymEnc},1}(\mathcal{A}) \Rightarrow 1] - \Pr[\text{IND-CPA}_{\text{SymEnc},0}(\mathcal{A}) \Rightarrow 1] \right| = \text{negl}.$$

Game $\text{IND-CPA}_{\text{SymEnc},b}(\mathcal{A})$	Oracle $\text{ENC}(m)$
1: $k \stackrel{\$}{\leftarrow} \text{SymEnc.Gen}(1^\lambda)$	1: $c \stackrel{\$}{\leftarrow} \text{Enc}(k, m)$
2: $m_0, m_1 \leftarrow \perp$	2: return c
3: $(m_0, m_1, \text{st}) \leftarrow \mathcal{A}^{\text{ENC}}$	
4: require $ m_0 = m_1 $	
5: $c^* \stackrel{\$}{\leftarrow} \text{Enc}(k, m_b)$	
6: $b' \leftarrow \mathcal{A}(c^*, \text{st})$	
7: return b'	

Figure 2.6: IND-CPA security for symmetric encryption scheme SymEnc .

Definition 14 (PRG Security of H). We consider the game in Figure 2.7. Let H be a function $H: \mathcal{S} \rightarrow \mathcal{W} \times \mathcal{K}$. We say that H is a secure *PRG* if, for all efficient adversaries \mathcal{A} , we have

$$\text{Adv}_H^{\text{prg}}(\mathcal{A}) := \left| \Pr[\text{PRG}_H(\mathcal{A}) \Rightarrow 1] - \frac{1}{2} \right| = \text{negl}.$$

Definition 15 (Message Authentication Code). A message authentication code (MAC.Gen, MAC.Tag, MAC.Vrfy) is a tuple of efficient algorithms as follows:

- $k \stackrel{\$}{\leftarrow} \text{Gen}(1^\lambda)$: The generation algorithm takes as input the security parameter 1^λ (encoded in unary) and outputs a key $k \in \mathcal{K}$.

Game $\text{PRG}_H(\mathcal{A})$	Oracle ROR
1: $b \xleftarrow{\$} \{0, 1\}$	1: if $b = 0$:
2: $b' \leftarrow \mathcal{A}^{\text{ROR}}$	2: $(w, k) \xleftarrow{\$} \{0, 1\}^\lambda \times \{0, 1\}^\lambda$
3: return $\mathbb{1}[b = b']$	3: if $b = 1$:
	4: $s \xleftarrow{\$} \{0, 1\}^\lambda$
	5: $(w, k) \leftarrow H(s) \in \{0, 1\}^\lambda \times \{0, 1\}^\lambda$
	6: return (w, k)

Figure 2.7: Pseudorandom generator (PRG) security for H.

- $\tau \xleftarrow{\$} \text{Tag}(k, m)$: The tag generation takes as inputs a key k and a message m , and outputs a tag τ .
- $b \leftarrow \text{Vrfy}(k, m, \tau)$: The verification algorithm takes as inputs a key k , a message m and a tag τ and outputs a bit $b \in \{0, 1\}$.

We say that MAC is correct if for any message $m \in \mathcal{M}$ and any key $k \in \mathcal{K}$ it holds that

$$\Pr \left[\text{Vrfy}(k, m, \tau) = 1 \mid \begin{array}{l} k \xleftarrow{\$} \text{Gen}(1^\lambda) \\ \tau \xleftarrow{\$} \text{Tag}(k, m) \end{array} \right] = 1,$$

where the probability is taken over the random coins of Tag , and \mathcal{M} and \mathcal{K} denote the message space and key space respectively.

Definition 16 (MAC SUF-CMA Security). We consider the game in Figure 2.8. Let MAC be a message authentication scheme. We say that MAC satisfies strong existential unforgeability if, for all efficient adversaries \mathcal{A} , we have

$$\text{Adv}_{\text{MAC}}^{\text{suf-cma}}(\mathcal{A}) := \Pr[\text{SUF-CMA}_{\text{MAC}}(\mathcal{A}) \Rightarrow 1] = \text{negl}.$$

Game $\text{SUF-CMA}_{\text{MAC}}(\mathcal{A})$	Oracle $\text{TAG}(m)$
1: $k \xleftarrow{\$} \text{MAC.Gen}(1^\lambda)$	1: $\tau \xleftarrow{\$} \text{MAC.Tag}(k, m)$
2: $Q \leftarrow \emptyset$	2: $Q \leftarrow Q \cup \{m, \tau\}$
3: $(m, t) \leftarrow \mathcal{A}^{\text{TAG}}$	3: return τ
4: return $\mathbb{1}[(m, t) \notin Q \wedge \text{MAC.Vrfy}(k, m, t) = 1]$	

Figure 2.8: SUF-CMA security for message authentication code MAC.

Two-Party Communication Part I

3 K-Waay: Fast and Deniable Post-Quantum X3DH Without Ring Signatures

In this chapter, we propose a deniable, post-quantum authenticated key exchange protocol that we call K-Waay that can be used in place of Signal’s X3DH protocol [MP16b]. An extended abstract corresponding to this work appeared at USENIX Security 2024, and was joint work with Loïs Huguenin-Dumittan, Ngoc Khanh Nguyen, Nicolas Rolin and Serge Vaudenay [CHDN⁺24a].

The author of this thesis was the primary contributor to the following parts of this work: the deniable authenticated key exchange primitive, key indistinguishability security model and definition of deniability, protocol description and protocol security proofs. The author further contributed to several other parts of this work in various capacities. Moreover, this work was the result of many hours of discussion and collaboration, and so it is difficult to attribute precisely here. The contents of the full version of this paper is nonetheless included in this thesis for comprehensibility and completeness.

3.1 Contribution

Researchers for several years now have sought to build cryptographic primitives and protocols that are resistant to efficient quantum attacks [Sho94]. This is highly evidenced with the NIST Post-Quantum Cryptography competition for standardising quantum-safe key encapsulation mechanisms (KEM) and signatures, organised by the United States National Institute of Standards and Technology (NIST). Recently, four schemes were selected by NIST for standardisation, out of which three rely on algebraic lattices. Indeed, with the US National Security Agency releasing their new CNSA 2.0 Suite [US], which says that CRYSTALS-Kyber [BDK⁺18] and CRYSTALS-Dilithium [DKL⁺18] should be the main cryptographic force for communication security beginning from 2030, lattices are a natural candidate for building more advanced cryptographic primitives, such as secure messaging.

The widely used Signal protocol for secure messaging as currently deployed is not quantum-safe since it is based on Diffie-Hellman key exchange [DH76]. The protocol, used in ap-

Chapter 3. K-Waay: Fast and Deniable Post-Quantum X3DH Without Ring Signatures

plications like Signal and WhatsApp, comprises two components, namely (1) the X3DH key exchange [MP16b] which is used to bootstrap sessions of (2) the Double Ratchet messaging protocol [PM16]. The Double Ratchet has been investigated in a line of recent works [ACD19, BFG⁺22a, CJSV22] that each neatly abstract the protocol into primitives like so-called *continuous key agreement*. Fortunately, these primitives have post-quantum (PQ) instantiations that leaves the core structure and resulting security guarantees of the Double Ratchet in place.

In standard X3DH, parties use a mixture of ephemeral (one-time), semi-static (many-time but temporary) and long-term keys. First, parties upload their keying material to a central server or public key infrastructure in a so-called prekey bundle. A party can then derive a session key by downloading their partner's bundle and performing three (or four) Diffie-Hellman key exchanges with a mixture of ephemeral and long-term (resp. plus semi-static) keys, ensuring at least confidentiality if the ephemeral *or* long-term key of each party is corrupted.

Observe that X3DH does not use signatures after signed prekeys are uploaded: at that point, the DH exchanges provide (implicit) authentication guarantees. Consequently, the protocol provides a level of *deniability* [DNS04, DGK06, UG15] as was formalised by Vatandas et al. [VGIK20]: informally, a participant can deny having performed key exchange with its counterpart. This is an important privacy guarantee that prevents (at least on a cryptographic level) a conversation transcript from incriminating an unsuspecting party, particularly in situations like whistleblowing and protesting.

In 2023, Signal announced and rolled out their initial *hybrid* post-quantum key exchange solution called PQXDH [KS23]. Like in X3DH, several Diffie-Hellman key exchanges are performed at once, but in PQXDH, parties upload prekey bundles that also contain a Kyber-1024 public key that the initiator additionally encapsulates to the responder with. Moreover, prekey bundles are still signed with the same signature scheme as regular X3DH based on Curve25519 [Ber06]. Although PQXDH provides post-quantum confidentiality [BJKS23a], which is an important first step towards post-quantum security as it prevents “store-now-decrypt-later” attacks, it does not provide post-quantum *authentication* as an active quantum attacker can trivially forge pre-key bundles. It is thus prudent to design a suitable X3DH alternative that is *fully* post-quantum secure.

A natural direction for building such a protocol is to emulate X3DH's structure by replacing Diffie-Hellman key exchange with a cryptographic group action, such as CSIDH [CLM⁺18]. In order to broadly capture this protocol structure, Brendel et al. [BFG⁺20] introduce a primitive called *split-KEM* that captures the symmetry of e.g. Diffie-Hellman. In a split-KEM, a party A encapsulates to their partner B by using their own secret sk_A and their partner's public key pk_B to produce a ciphertext; B then decapsulates it using sk_B and pk_A . The authors define indistinguishability-based security notions and notice that Frodo [BCD⁺16] lattice-based key-exchange fulfills the split-KEM syntax and the weakest notion of indistinguishability they

define.¹ Although they present an X3DH-like protocol, they do not define a security model, and, looking ahead, their split-KEM security notions do not suffice to construct X3DH-like key exchange with authenticity and deniability.

In two recent works, Hashimoto et al. [HKKP21, HKKP22] and Brendel et al. [BFG⁺22b] concurrently proposed instead to construct X3DH-like key exchange using KEMs directly. Since a core feature of X3DH is its *asynchronicity*, a challenge-response protocol cannot be employed using KEMs alone to provide authentication [SSW20]. Thus to ensure deniability, two seemingly different approaches were proposed: Hashimoto et al. [HKKP21] apply ring signatures while Brendel et al. [BFG⁺22b] use a flavour of designated verifier signatures; these primitives were later shown to be equivalent [HKKP22].

As described in the aforementioned works, the currently most efficient post-quantum ring signatures [Beu20, ESZ22, LAZ19, LN22, YEL⁺21] are proven to be secure in the random oracle model (ROM) [BR93] and can enjoy signatures that are a handful of kilobytes large. Often, however, the constructions do not come with a security proof in the quantum random oracle model (QROM) [BDF⁺11]. In this vein, parameters are generally optimistically chosen as the security loss incurred by proofs in the ROM is not taken into account when setting them, without even mentioning QROM loss, which is usually much larger. Further, security notions can differ between works, making it less clear exactly when they are appropriate for use.

More generally, it is of interest to determine the cost (or overhead) that deniability incurs in (X3DH-like) key exchange. Towards this goal, Hashimoto et al. [HKKP22] provide benchmarks for their baseline, non-deniable X3DH-like protocol based on signatures and KEMs, and Brendel et al. [BFG⁺22b] consider parameter sizes for (but do not benchmark) existing ring and designated verifier signatures. As such, a more fine-grained and detailed evaluation will help inform practitioners on the overhead incurred by deniability in the post-quantum setting.

While the use of ring signatures to build PQ and deniable X3DH is at least theoretically understood, this far from exhausts the protocol design space. Motivated by this and the above discussion, we therefore ask the following research question:

- Can we design a provably-secure, efficient and deniable post-quantum X3DH alternative that does *not* require ring signatures?

3.1.1 Summary

In this chapter, we propose an efficient, deniable and post-quantum X3DH-like protocol without ring signatures that we call K-Waay. To summarise our contributions:

- Towards building our protocol, we revisit the split-KEM formalism proposed by Brendel

¹The construction can conceptually be seen as instantiating the lattice-based cryptographic group action of Beullens [Beu20].

Chapter 3. K-Waay: Fast and Deniable Post-Quantum X3DH Without Ring Signatures

et al. [BFG⁺20] and deduce that several additional properties, namely notions of authenticity and deniability, are needed to construct a secure X3DH-like deniable authenticated key exchange protocol (DAKE) (Section 3.2).

- We propose K-Waay, an X3DH-like DAKE that uses deniable and unforgeable split-KEM at its core (Section 3.4). Our protocol uses signatures to sign prekeys, and then uses ephemeral KEM, long-term KEM and split-KEM for the final key exchange step. We compare the security of our protocol with the state-of-the-art in Table 3.1.
- The main drawback of a naive version of our protocol is that parties can run out of ephemeral keys, thus making the protocol synchronous if this happens (e.g. Bob needs to wait for Alice’s fresh ephemeral key before sending a message). While such a problem would rarely occur in practice, given enough keys are uploaded on the server, we propose a simple trick that makes the reuse of ephemeral keys possible on the receiver’s side for messages they received while offline. We think this trick could be of independent interest as it – perhaps surprisingly – allows for a specific kind of key reuse for a split-KEM that is *not* IND-CCA secure.
- We prove key indistinguishability in our model that captures ephemeral key reuse and session state exposure, and prove a variant of deniability that strengthens the notion of Brendel et al. [BFG⁺22b] by additionally leaking the victim’s session state to the adversary in the security game.
- We instantiate a post-quantum split-KEM secure under our new security notions derived from the Frodo key exchange protocol (FrodoKEX) [BCD⁺16] based on the plain LWE assumption. The parameters we chose provide strong security guarantees, providing more than 192 bits of classical and quantum security for our core split-KEM security notions OW-CPA, decaps-CPA and deniability. We then use a transform in the (Q)ROM to prove it UNF-1KCA and IND-1BatchCCA (i.e., our new unforgeability and indistinguishability definitions for split-KEM). This construction incurs a security loss as usual in the (Q)ROM, but our final split-KEM still provides around 128 (resp. 64) bits of security in the ROM (resp. QROM) assuming the adversary is limited to 2^{64} (resp. quantum) random oracle queries.
- We benchmark our protocol K-Waay using our modified version of FrodoKEX (which we call FrodoKEX+) as the split-KEM, along with standard X3DH and the two previous proposals for PQ X3DH-like AKE [HKKP22, BFG⁺22b] (Section 3.6). We find that while K-Waay has larger prekeys, it is $6\times$ faster compared to these. In addition, the only non-standard primitive we use in K-Waay (FrodoKEX+) is based on both an assumption (LWE) and a scheme (FrodoKEM) that have been thoroughly scrutinised by the cryptographic community. Overall, we believe our protocol is more mature and therefore suitable for short to medium-term integration compared to previous work based on ring signatures.

3.1 Contribution

Protocol	PQ Conf	PQ Auth	KCI	FS	SSR	RR	Deniability
X3DH [MP16b, CCD ⁺ 20]	✗	✗	✓	PFS	✓	✓	Malicious
PQXDH [KS23, BJKS23a]	✓	✗	✓	PFS	✗	✗	Semi-honest+
KEM+Sigs [HKKP22]	✓	✓	✓	PFS	✓	✗	✗
HKKP [HKKP22]	✓	✓	✓	WFS	✓	✗	Semi-honest
SPQR [BFG ⁺ 22b]	✓	✓	✓	WFS	✗	✓	Semi-honest
K-Waay (Section 3.4)	✓	✓	✓	WFS	✓	✗	Semi-honest

Table 3.1: Comparison between different security properties proven for existing X3DH-like key exchange protocols, namely post-quantum confidentiality (PQ Conf), authentication (PQ Auth), resistance to key-compromise impersonation attacks when long-term keys are exposed (KCI), perfect forward secrecy (PFS) or weak forward secrecy (WFS) [Kra05], session state reveal (SSR), randomness reveal (RR) and deniability (where the judge/adversary is either honest-but-curious or is malicious and can inject messages). Protocols can be generically strengthened to handle randomness reveal by standard application of the so-called NAXOS trick [LLM07]. “KEM+Sigs” refers to the non-deniable baseline X3DH-like protocol proposed by Hashimoto et al. [HKKP21, HKKP22], and “HKKP” refers to their deniable X3DH protocol *without* NIZKs (their protocol with NIZKs implies maliciously-secure deniability w.r.t. a *classical* adversary). The security of PQXDH is based on the recent analysis of Bhargavan et al. [BJKS23a] except that Kret and Schmidt argue it also provides at least semi-honest deniability [KS23].

3.1.2 Technical Overview

X3DH-like Key Exchange. A quantum-secure X3DH-like protocol should satisfy certain properties. Apart from satisfying standard authenticated key exchange (AKE) properties like secrecy and authentication, it should also be *asynchronous*. That is, parties should be able to upload keying material to a central server, after which an initiating party can derive a session key immediately with their counterpart who may be offline. This also entails *receiver-obliviousness*, using the language of Hashimoto et al. [HKKP22], as the initial key upload should not depend on the keys of any other party. Another is *deniability*, allowing parties to claim that they plausibly did not participate in the key exchange. Note that we cannot possibly ensure that parties can claim that they never uploaded prekeys as they are signed (and using primitives like ring signatures would violate receiver-obliviousness). Finally, a deniable authenticated key exchange protocol, or *DAKE*, should, like X3DH, provide security guarantees even if the session state of a party is leaked.

Revisiting Split-KEM. In an attempt to model the primitive central to X3DH-like AKE, Brendel et al. [BFG⁺20] introduced *split-KEM*, which is similar to a standard KEM except the encapsulator can contribute to the derived key. However, we discovered that the accompanying security definitions were not sufficient to use such a primitive as the main component of a key exchange protocol. The reason is that their notions ensure that an encapsulated ciphertext will not leak information on its encapsulated key, but not that only the sender can send a “legitimate” ciphertext to the receiver (or that only the sender and receiver can derive a common key). In other words, there is no guarantee of implicit authentication. Therefore, we introduce

Chapter 3. K-Waay: Fast and Deniable Post-Quantum X3DH Without Ring Signatures

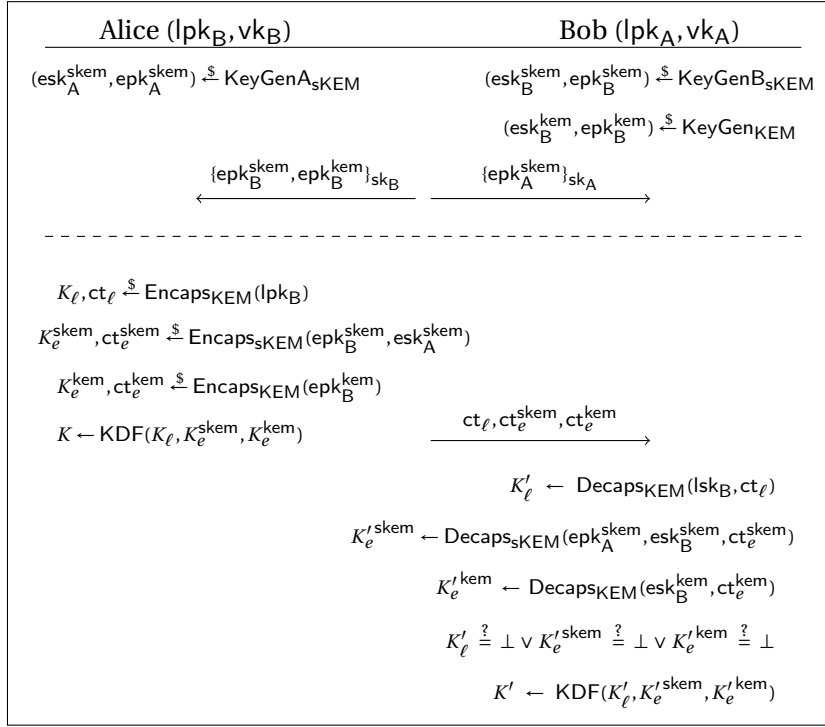


Figure 3.1: High-level overview of the K-Waay protocol. Values in brackets $\{\cdot\}_{sk}$ are signed with sk and the signature is verified upon reception. For clarity, we omit the calculation and addition of session identifier sid to KDF .

the notion of unforgeability against one known-ciphertext attacks for split-KEM (UNF-1KCA), which ensures that if Alice receives a message allegedly sent by Bob, either Bob really sent it or the decapsulation will fail. Looking ahead, this will be used in the security proof of the protocol to argue that either the adversary relayed a legitimate split-KEM ciphertext to the receiver that the adversary cannot learn the decapsulation of, or the receiver aborts as the ciphertext is forged.

We also introduce an intermediary notion that we call decaps-CPA, which enforces that an adversary should not be able to recover a key *decapsulated* by some party without knowing the sender's or receiver's secret key. We will prove that our lattice-based split-KEM satisfies this notion, then we will apply some transform in the (Q)ROM to obtain a UNF-1KCA split-KEM.

Finally, we also define a notion capturing deniability for split-KEM, which states that no judge J can be convinced that a party A sent a given ciphertext to B , even knowing B 's secret key but assuming both parties did not deviate from the protocol. This models a setting where B communicates with A and later tries to frame the latter by giving the transcript and their own secret key to J .

Construction. As any X3DH-like protocol, our construction works in 4 phases: long-term key generation, prekey generation, sending and receiving. The first observation we make is

that in X3DH, prekey bundles are signed with a long-term signing key before being uploaded to the server. This fact is often abstracted away in formal analyses as it hurts the claims one can make about the deniability of X3DH: as a signature is undeniable by definition, users cannot deny they *participated* in the protocol. Based on this, our goal was to achieve some level of peer-deniability [CF11], where parties can deny they communicated with someone in particular, and to leverage the fact that we use signatures to authenticate the prekeys. Our protocol works then as follows (see Figure 3.1 for a high-level overview). The long-term key pair consists of a KEM and signature key pair, the latter being used to sign the prekey, which consists of an ephemeral KEM key pair and ephemeral split-KEM key pair. The former is used for forward secrecy while the second is used for implicit authentication of the sender. Although usually ephemeral keys cannot be used for authentication as they are dynamic, in our case we can since they are authenticated (i.e., signed) by their owner. Then, the sender encapsulates against both KEM public keys of the receiver, and uses their own split-KEM secret key and the receiver’s public key to derive a split-KEM ciphertext. Upon decapsulation, the receiver recovers the three encapsulated keys and combines them using a PRF to derive the shared key.

Ephemeral Split-KEM Key Reuse. The way our protocol is described above works perfectly well if the split-KEM satisfies the UNF-1KCA unforgeability notion mentioned above. However, in practice, it could happen that some party, say Bob, is offline for too long and all their ephemeral split-KEM keys have been used. If that occurs, another sender would have to wait for Bob to come online and upload new keys before they can send him a message.

We fix this issue by modifying the protocol as follows: when Bob’s ephemeral public keys have run out on the server, a sender can simply reuse one of them. Then, when Bob is back online, he groups the ciphertexts corresponding to the same public key and decrypts all ciphertexts in a group *at once*. If one or more of the split-KEM decapsulations in a group fails, Bob outputs \perp for *all* ciphertexts and, e.g., restarts the protocol. Otherwise, Bob proceeds as before (and *never* decapsulates again using the same split-KEM key).

DAKE Modelling. To formally capture ephemeral key reuse, our DAKE syntax includes an algorithm `BatchReceive` that takes as input a session state and one or more messages to be received. Our key indistinguishability notion therefore has to account for `BatchReceive`, and so we extended the typical Bellare-Rogaway-style modelling [BR94] to this end (ignoring this, our model combines aspects of the models of Hashimoto et al. [HKKP21, HKKP22] and Brendel et al. [BFG⁺22b]). In particular, our notion of *partnering* between sending and receiving sessions is such that a given sending session can be partnered with several receiving sessions. Partnered sessions, in turn, are used when defining trivial attacks that our protocol (and often, but not always *any* protocol) cannot prevent (e.g., the adversary exposing both the session and long-term state of the receiver) as well as correctness checks. The adversary’s challenge query, i.e., the query that returns to them either a real or random key, returns a single key, which in the case of `BatchReceive` corresponds to only part of its output.

Chapter 3. K-Waay: Fast and Deniable Post-Quantum X3DH Without Ring Signatures

Security. We show this version of the protocol is secure assuming the split-KEM satisfies a stronger notion than IND-CPA that we call IND-1BatchCCA (in addition to UNF-1KCA security). This definition is the same as traditional IND-CPA (adapted to the split-KEM syntax), except the adversary can query a decapsulation oracle *once* with multiple public keys and ciphertexts, and the oracle returns \perp if *one or more* of the decapsulations failed, and the resulting keys otherwise. We show that one can easily build an IND-1BatchCCA split-KEM out of a CPA-secure one in the (Q)ROM, conveniently using the same transform mentioned above that builds a UNF-1KCA scheme out of a decaps-CPA one.

As in previous protocols [BFG⁺22b, HKKP22], the long-term KEM provides implicit authentication of the receiver as only they can decrypt. As mentioned above, the ephemeral KEM provides forward secrecy, and the UNF-1KCA/IND-1BatchCCA split-KEM provides implicit authentication of the sender, as it guarantees that only the sender could have sent a ciphertext that correctly decapsulates (unforgeability), and no adversary knows what is inside that ciphertext (indistinguishability), even after seeing the decapsulation of one batch of ciphertexts encapsulated against the same public key (given no decapsulation failed). We note that the sender-to-receiver authentication depends both on a long-term key (i.e., the signing key) and an ephemeral one (the split-KEM key). Consequently, our model (that allows session state exposure) is more restrictive than that of Hashimoto et al. [HKKP22], since in particular it suffices for the adversary to learn a receiver’s ephemeral state during key exchange to forge a message that the receiver accepts. Intuitively, this is because split-KEM is effectively a symmetric primitive. Nevertheless, the security that we achieve is stronger than weak forward security without session state exposure.

Deniable Split-KEM from Lattices. We provide a lattice-based split-KEM which satisfies both deniability and UNF-1KCA security. Our starting point is the Frodo key-exchange (FrodoKEX) [BCD⁺16], which was identified (among other schemes) as a split-KEM by Brendel et al. [BFG⁺20], the security of which relies on the well-known learning-with-errors (LWE) problem [Reg05]. We highlight that the vanilla construction of FrodoKEX does not enjoy the aforementioned properties.² Indeed, when looking closely at the security games of deniability and UNF-1KCA, partial information about the secret keys are revealed - thus making a reduction to LWE completely non-trivial. We circumvent this problem in two ways.

First, we reduce deniability of our scheme to a so-called Extended-LWE problem [AP12], where in addition to a standard LWE instance, the adversary is given a short random combination of the secret coefficients. We show that deniability of our scheme reduces straightforwardly to Extended-LWE, and then follow the methodology of Alperin-Sheriff and Peikert [AP12] to reduce it further to plain LWE.

Towards UNF-1KCA security, we slightly modify the Frodo split-KEM by introducing masking terms. As the name suggests, they are used to hide the partial information about secret keys. In Section 3.6.2 we discuss the necessity of this (perhaps seemingly artificial) change.

²Nevertheless, we found no practical attack on deniability/UNF-1KCA for FrodoKEX.

3.1.3 Additional Related Work

The security of X3DH has been modelled in detail by Cohn-Gordon et al. [CCD⁺20]. Vatanadas et al. [VGIK20] investigate the deniability of X3DH and similar key exchange protocols under the deniability notion of Di Raimondo et al. [DGK06], requiring strong knowledge-of-exponent-type assumptions to prove X3DH secure. Dobson and Galbraith [DG22] propose a SIDH-based X3DH-like protocol which is unfortunately now broken [CD23]. Very recently, Kiltz et al. [KPRR23] prove a simplified version of X3DH tightly-secure in the generic group model under a new multi-user assumption supporting corruptions although do not allow the adversary to expose parties' session states.

Unger and Goldberg build a number of different DAKEs [UG15, UG18]. However, the protocols do not provide post-quantum guarantees: only in their later paper [UG18] is it suggested to add a PQ KEM for post-quantum *confidentiality* and the authors do not propose a more comprehensive hybrid protocol. Nevertheless, the protocols provide relatively strong online deniability (where a judge and a party can communicate while trying to frame another party) at the expense of stronger primitives like dual-receiver encryption and non-committing encryption.

Alwen et al. [ABH⁺21] introduce the notion of authenticated key encapsulation mechanism (AKEM) and some security definitions. AKEM captures the same primitive as a split-KEM, but we opted for the syntax and language of the latter as it was meant to be used in an X3DH-like protocol.

Cremers and Feltz [CF11] introduce peer-deniability, which captures the kind of participation deniability property we are after, namely that a party cannot deny using a system but can deny communicating with a particular party. However, their security notion does not require the simulator to output the session key nor the judge/adversary to distinguish between the real and simulated key, and so composability issues may arise from using it.

Related to split-KEM is *signcryption* [Zhe97], which in syntax is roughly its encryption analogue. We are not aware of any works that consider deniable signcryption in the post-quantum setting.

3.2 Split-KEM

The primitive at the core of our protocol is split-KEM, which we present in this section. It was first defined by Brendel et al. [BFG⁺20].

Definition 17 (Split-KEM). A split-KEM sKEM is a tuple of four efficient algorithms (KeyGenA, KeyGenB, Encaps, Decaps) defined as follows:

- $(pk_A, sk_A) \xleftarrow{\$} \text{KeyGenA}(1^\lambda)$ (resp. $(pk_B, sk_B) \xleftarrow{\$} \text{KeyGenB}(1^\lambda)$): The key generation function of the first/second party takes the security parameter λ as input, and outputs a pair

Chapter 3. K-Waay: Fast and Deniable Post-Quantum X3DH Without Ring Signatures

of public/secret keys (pk_A, sk_A) (resp. (pk_B, sk_B)).

- $(K, ct) \stackrel{\$}{\leftarrow} \text{Encaps}(pk_A, sk_B)$: The encapsulation function takes the public key pk_A of party A and party B's secret key sk_B as inputs, and outputs a ciphertext ct and a key K .
- $K/\perp \leftarrow \text{Decaps}(pk_B, sk_A, ct)$: The decapsulation function takes the secret key sk_A of party A, their counterpart's public key pk_B and a ciphertext ct as inputs, and outputs a key K or the error symbol \perp .

We say a split-KEM is $(1 - \delta)$ -correct if

$$\Pr \left[\begin{array}{l} (pk_A, sk_A) \stackrel{\$}{\leftarrow} \text{KeyGenA}(1^\lambda); \\ (pk_B, sk_B) \stackrel{\$}{\leftarrow} \text{KeyGenB}(1^\lambda); \\ (K, ct) \stackrel{\$}{\leftarrow} \text{Encaps}(pk_A, sk_B); \\ K' \leftarrow \text{Decaps}(pk_B, sk_A, ct) \end{array} \right] \leq \delta .$$

Intuitively, a split-KEM is similar to a normal KEM except material from both participants is used for encapsulation (i.e., the final key will depend on both parties' secret/public keys). In an X3DH-like protocol, it can be used to implicitly authenticate the party encapsulating. In the language of Brendel et al. [BFG⁺20], our notion of split-KEM is "asymmetric", as it is assumed that B always encapsulates and A always decapsulates. This is sufficient for our purpose, but we note that all the results presented in this chapter can be adapted to a symmetric split-KEM where $\text{KeyGenA} = \text{KeyGenB}$.

3.2.1 Security

We will need several security properties from the split-KEM to prove our whole protocol secure. We first define one-wayness (OW-CPA) for sKEM, which is very similar to the usual one for KEM and another new notion called IND-1BatchCCA. Looking ahead, we will show that any OW-CPA split-KEM can easily be transformed into a IND-1BatchCCA one in the (quantum) random oracle model or (Q)ROM (c.f. Chapter 2).

Definition 18 (split-KEM OW-CPA). We consider the OW-CPA game defined in Figure 3.2. A split-KEM scheme $s\text{KEM} = (\text{KeyGenA}, \text{KeyGenB}, \text{Encaps}, \text{Decaps})$ is OW-CPA if for any efficient adversary \mathcal{A} we have

$$\text{Adv}_{s\text{KEM}}^{\text{ow-cpa}}(\mathcal{A}) = \Pr[\text{OW-CPA}_{s\text{KEM}}(\mathcal{A}) \Rightarrow 1] = \text{negl} .$$

Definition 19 (split-KEM IND-1BatchCCA). We consider the IND-1BatchCCA game defined in Figure 3.2. Let \mathcal{K} be a finite key space. A split-KEM scheme over \mathcal{K} $s\text{KEM} = (\text{KeyGenA}, \text{KeyGenB}, \text{Encaps}, \text{Decaps})$ is IND-1BatchCCA if for any efficient adversary \mathcal{A} we have

$$\text{Adv}_{s\text{KEM}}^{\text{ind-1batchcca}}(\mathcal{A}) := \left| \Pr[\text{IND-1BatchCCA}_{s\text{KEM}}(\mathcal{A}) \Rightarrow 1] - \frac{1}{2} \right| = \text{negl} .$$

<p>Game IND-1BatchCCA_{sKEM}(\mathcal{A})</p> <ol style="list-style-type: none"> 1: $b \xleftarrow{\\$} \{0, 1\}; q \leftarrow 0$ 2: $pk_A, sk_A \xleftarrow{\\$} \text{KeyGenA}(1^\lambda)$ 3: $pk_B, sk_B \xleftarrow{\\$} \text{KeyGenB}(1^\lambda)$ 4: $K_0, ct^* \xleftarrow{\\$} \text{Encaps}(pk_A, sk_B)$ 5: $K_1 \xleftarrow{\\$} \mathcal{K}$ 6: $b' \xleftarrow{\\$} \mathcal{A}^{\text{BatchDEC}}(pk_A, pk_B, ct^*, K_b)$ 7: return $\mathbb{1}[b' = b]$ 	<p>Oracle BatchDEC($\{(pk_i, ct_i)\}_{i=1}^d$)</p> <ol style="list-style-type: none"> 1: if $q = 1$: return \perp 2: else : $q \leftarrow q + 1$ 3: for $i \in \{1, \dots, d\}$ 4: if $(pk_i, ct_i) = (pk_B, ct^*)$: return \perp 5: $K'_i \leftarrow \text{Decaps}(pk_i, sk_A, ct_i)$ 6: if $K'_1 = \perp \vee \dots \vee K'_d = \perp$: return \perp 7: return (K'_1, \dots, K'_d)
<p>Game OW-CPA_{sKEM}(\mathcal{A})</p> <ol style="list-style-type: none"> 1: $pk_A, sk_A \xleftarrow{\\$} \text{KeyGenA}(1^\lambda); pk_B, sk_B \xleftarrow{\\$} \text{KeyGenB}(1^\lambda)$ 2: $K^*, ct^* \xleftarrow{\\$} \text{Encaps}(pk_A, sk_B)$ 3: $K' \xleftarrow{\\$} \mathcal{A}(pk_A, pk_B, ct^*)$ 4: return $\mathbb{1}[K' = K^*]$ 	

Figure 3.2: IND-1BatchCCA and OW-CPA games.

We also recall the different notions of indistinguishability for (asymmetric) split-KEM defined by Brendel et al. [BFG⁺20] that we do not use except to argue their insufficiency for AKE below:

<p>Game xy-IND-CCA_{sKEM}(\mathcal{A})</p> <ol style="list-style-type: none"> 1: $b \xleftarrow{\\$} \{0, 1\}$ 2: $n_x \leftarrow 0; n_y \leftarrow 0;$ 3: $pk_A, sk_A \xleftarrow{\\$} \text{KeyGenA}(1^\lambda)$ 4: $pk_B, sk_B \xleftarrow{\\$} \text{KeyGenB}(1^\lambda)$ 5: $K_0, ct^* \xleftarrow{\\$} \text{Encaps}(pk_A, sk_B)$ 6: $K_1 \xleftarrow{\\$} \mathcal{K}$ 7: $b' \xleftarrow{\\$} \mathcal{A}^{\text{ENC,DEC}}(pk_A, pk_B, ct^*, K_b)$ 8: return $\mathbb{1}[b' = b]$ 	<p>Oracle ENC(pk)</p> <ol style="list-style-type: none"> 1: if $n_y \geq y$: return \perp 2: $n_y \leftarrow n_y + 1$ 3: $K, ct \xleftarrow{\\$} \text{Encaps}(pk, sk_B)$ 4: if $(pk, ct) = (pk_A, ct^*)$: return \perp 5: return (K, ct) <p>Oracle DEC(pk, ct)</p> <ol style="list-style-type: none"> 1: if $n_x \geq x$: return \perp 2: $n_x \leftarrow n_x + 1$ 3: if $(pk, ct) = (pk_B, ct^*)$: return \perp 4: return $\text{Decaps}(pk_{\overline{p}}, sk_{\overline{p}}, ct)$
----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

Figure 3.3: xy -IND-CCA games for an “asymmetric” split-KEM from Brendel et al. [BFG⁺20], where $x, y \in \{n, s, m\}$. When doing the comparison on the first line of both oracles, we assume $n = 0, s = 1$ and $m = \infty$.

Definition 20 (split-KEM xy -IND-CCA). We consider the xy -IND-CCA game defined in Figure 3.3. A split-KEM scheme $sKEM = (\text{KeyGenA}, \text{KeyGenB}, \text{Encaps}, \text{Decaps})$ is xy -IND-CCA,

Chapter 3. K-Waay: Fast and Deniable Post-Quantum X3DH Without Ring Signatures

with $x, y \in \{n, s, m\}$ if for any efficient adversary \mathcal{A} we have

$$\text{Adv}_{\text{sKEM}}^{\text{xy-ind-cca}}(\mathcal{A}) := \left| \Pr [\text{xy-IND-CCA}_{\text{sKEM}}(\mathcal{A}) \Rightarrow 1] - \frac{1}{2} \right| = \text{negl}.$$

These indistinguishability notions range from nn -IND-CCA, which is similar to a kind of IND-CPA security as the adversary has no access to encapsulation or decapsulation oracles, to mm -IND-CCA, which captures strong IND-CCA security for split-KEMs. More generally, all notions are of the form xy -IND-CCA, $x, y \in \{n, s, m\}$, where x (resp. y) specifies the number of queries an adversary can make to the decapsulation (resp. encapsulation) oracle (none, single, or many).

On the Original Split-KEM Security Notions. We recall that the advantage of split-KEMs over normal KEMs is that they capture the fact that the party encapsulating can contribute (static) keying material towards the shared key, whereas it is not the case with KEMs, as the encapsulation function only takes the receiving party's public key as input. In particular, this means that KEMs cannot be used for implicit authentication of the encapsulator, unlike split-KEMs. However, we argue that the original xy -IND-CCA definitions for split-KEMs [BFG⁺20] do not capture implicit authentication either and thus are not suited for their purpose (building an asynchronous DAKE). In fact, any IND-CPA (resp. IND-CCA) KEM can easily be converted to an (asymmetric) split-KEM satisfying nn -IND-CCA (resp. mm -IND-CCA).

More formally, imagine a setting where Alice and Bob know each other's public key, and Bob wants to implicitly authenticate to Alice using a split-KEM. In addition, we assume a mm -IND-CCA split-KEM sKEM_0 exists (note mm -IND-CCA security is the strongest so this holds for all weaker notions). We first modify sKEM_0 such that on a special ciphertext ct^* not in the original ciphertext space, Decaps returns a constant key K^* . Let's call this modified scheme sKEM . We observe that sKEM is still mm -IND-CCA secure as no adversary can break an honestly-generated challenge ciphertext. Now, implicit authentication means that if Alice decapsulates a ciphertext and obtains a key K , then only Alice knows K . However, in our case, any adversary can send ct^* to Alice and set their own key to K^* . Both the adversary and Alice will share the same key and implicit authentication does not hold. In a way, xy -IND-CCA security does not prevent forgeries.

UNF-1KCA. This leads us to define our notion of UNF-1KCA security for split-KEMs below which, along with OW-CPA (which can be turned into IND-1BatchCCA), guarantees that only Bob (and obviously Alice) can know the result of Alice's decapsulation on some ciphertext. More precisely, UNF-1KCA ensures that no adversary can forge a valid split-KEM ciphertext for A even knowing a ciphertext that was computed with respect to a public key chosen by the adversary³, under the condition that the public key used for encapsulation and the known ciphertext are different from the pair made of A 's public key and the ciphertext output

³Looking ahead, the fact that the public key is adversarially-chosen will be useful for proving security under key-compromise impersonation attacks for our full protocol.

by the adversary. We also define a security notion called decaps-CPA that will serve as a building block to build UNF-1KCA. The decaps-CPA notion ensures that it is hard for an adversary knowing a ciphertext ct (under an adversarially-chosen public key) to come up with a ciphertext ct' (possibly equal to ct) and a key K' such that the decapsulation of ct' returns K' .

Definition 21 (split-KEM UNF-1KCA). We consider the UNF-1KCA game defined in Figure 3.4. A split-KEM scheme $sKEM = (\text{KeyGenA}, \text{KeyGenB}, \text{Encaps}, \text{Decaps})$ is UNF-1KCA if for any efficient adversary \mathcal{A} we have

$$\text{Adv}_{sKEM}^{\text{unf-1kca}}(\mathcal{A}) := \Pr[\text{UNF-1KCA}_{sKEM}(\mathcal{A}) \Rightarrow 1] = \text{negl}.$$

Definition 22 (split-KEM decaps-CPA). We consider the decaps-CPA game defined in Figure 3.4. A split-KEM scheme $sKEM = (\text{KeyGenA}, \text{KeyGenB}, \text{Encaps}, \text{Decaps})$ is decaps-CPA if for any efficient adversary \mathcal{A} we have

$$\text{Adv}_{sKEM}^{\text{decaps-cpa}}(\mathcal{A}) := \left| \Pr[\text{decaps-CPA}_{sKEM}(\mathcal{A}) \Rightarrow 1] - \frac{1}{2} \right| = \text{negl}.$$

Game UNF-1KCA _{sKEM} (\mathcal{A})	Game decaps-CPA _{sKEM} (\mathcal{A})
1: $pk_A, sk_A \xleftarrow{\$} \text{KeyGenA}(1^\lambda)$	1: $pk_A, sk_A \xleftarrow{\$} \text{KeyGenA}(1^\lambda)$
2: $pk_B, sk_B \xleftarrow{\$} \text{KeyGenB}(1^\lambda)$	2: $pk_B, sk_B \xleftarrow{\$} \text{KeyGenB}(1^\lambda)$
3: $pk, st \xleftarrow{\$} \mathcal{A}(pk_A, pk_B)$	3: $pk, st \xleftarrow{\$} \mathcal{A}(pk_A, pk_B)$
4: $K_B, ct \xleftarrow{\$} \text{Encaps}(pk, sk_B)$	4: $K_B, ct \xleftarrow{\$} \text{Encaps}(pk, sk_B)$
5: $ct' \xleftarrow{\$} \mathcal{A}(pk_A, pk_B, ct, K_B, st)$	5: $K'_A, ct' \xleftarrow{\$} \mathcal{A}(pk_A, pk_B, ct, st)$
6: if $(ct, pk) = (ct', pk_A)$: return 0	6: $K_A \leftarrow \text{Decaps}(pk_B, sk_A, ct')$
7: $K_A \leftarrow \text{Decaps}(pk_B, sk_A, ct')$	7: if $K_A = \perp$: abort
8: if $K_A = \perp$: return 0	8: return $\mathbb{1}[K_A = K'_A]$
9: return 1	

Figure 3.4: Games UNF-1KCA and decaps-CPA.

3.2.2 Deniability

Finally, we state the notion of split-KEM deniability we would like to achieve.

Definition 23 (Deniability). We consider the game shown in Figure 3.5. We say a split-KEM $sKEM$ is *deniable* if there exists a simulator Sim s.t. for all efficient adversaries \mathcal{A} , we have

$$\text{Adv}_{sKEM, \text{Sim}}^{\text{deny}}(\mathcal{A}) := \left| \Pr[\text{DENV}_{sKEM, \text{Sim}}^{\text{REAL}}(\mathcal{A}) \Rightarrow 1] - \Pr[\text{DENV}_{sKEM, \text{Sim}}^{\text{SIM}}(\mathcal{A}) \Rightarrow 1] \right| = \text{negl}.$$

Informally, the setting considered is the following. Alice and Bob use the split-KEM to establish a shared key (we assume the public keys are only used for this one exchange), and Alice (while

Game $\text{DENY}_{\text{sKEM,Sim}}^{\text{REAL}}(\mathcal{A})$	Game $\text{DENY}_{\text{sKEM,Sim}}^{\text{SIM}}(\mathcal{A})$
1: $(pk_A, sk_A) \xleftarrow{\$} \text{KeyGenA}(1^\lambda)$	1: $(pk_A, sk_A) \xleftarrow{\$} \text{KeyGenA}(1^\lambda)$
2: $(pk_B, sk_B) \xleftarrow{\$} \text{KeyGenB}(1^\lambda)$	2: $(pk_B, sk_B) \xleftarrow{\$} \text{KeyGenB}(1^\lambda)$
3: $K, ct \xleftarrow{\$} \text{Encaps}(pk_A, sk_B)$	3: $K, ct \xleftarrow{\$} \text{Sim}(pk_B, sk_A)$
4: $b \xleftarrow{\$} \mathcal{A}(pk_A, pk_B, sk_A, K, ct)$	4: $b \xleftarrow{\$} \mathcal{A}(pk_A, pk_B, sk_A, K, ct)$
5: return b	5: return b

Figure 3.5: Deniability game (we assume w.l.o.g. that B encapsulates and A simulates).

following the protocol) wants to frame Bob and prove that he did communicate with her. Therefore, after receiving Bob’s ciphertext and deriving the key, Alice gives both public keys, the derived key, the ciphertext and her own secret key to a judge (i.e., the adversary) that must decide whether Bob actually sent the ciphertext that was used to derive the key or not. The scheme is deniable if there is a simulator that, given Alice’s view, outputs a ciphertext and a key indistinguishable from the ones output by Bob. We discuss deniability further after introducing our key exchange deniability notion (Section 3.3.3) and in Section 3.6.

3.3 Deniable Authenticated Key Exchange

In this section, we describe our model for deniable authenticated key exchange (DAKE) that we tailor to the semantics and flow of X3DH.

3.3.1 Syntax

Definition 24. A DAKE DAKE is a tuple of four efficient algorithms (KeyGen, Init, Send, BatchReceive) defined as follows:

- $(pk, sk) \xleftarrow{\$} \text{KeyGen}(1^\lambda)$. This function takes as input the security parameter λ and outputs the long-term public/secret key pair of the caller.
- $(st_i, \text{prek}_i) \xleftarrow{\$} \text{Init}(sk_i, \text{role})$. This function takes as inputs a long-term secret key sk_i and a role $\text{role} \in \{\text{sender}, \text{receiver}\}$ and outputs a session state st_i and a prekey bundle prek_i . Init models the creation of key material that will be uploaded to the public key infrastructure by both parties (e.g., a prekey bundle in X3DH). The output values depend only on the public key of party i executing the function.
- $(k, m) \xleftarrow{\$} \text{Send}(sk_i, pk_j, st_i, \text{prek}_j)$. This function takes as inputs the secret key of the executing party i , the public key of the intended recipient pk_j , party i ’s session state st_i and the (claimed) prekey bundle of the intended recipient prek_j , and outputs a key k and a message m .
- $\{k_s\}_s \leftarrow \text{BatchReceive}(sk_i, st_i, \{pk_j, \text{prek}_j, m_j\}_j)$. This function takes as inputs the secret

3.3 Deniable Authenticated Key Exchange

key of the executing party i , an ephemeral state of party i st_i and a vector of size $d \geq 1$ of the form $(pk_j, prek_j, m_j)$ for party i 's session with the public key of the (claimed) sender pk_j , the (claimed) prekey bundle of party j $prek_j$ and a message m_j , and outputs a vector of d keys (k_1, \dots, k_d) , some or all of which may be \perp .

`Init` explicitly captures parties uploading ephemeral keys to a central server in the first protocol step. This contrasts with the formal modelling in some previous works on X3DH-like key exchange [BFG⁺22b, HKKP22] that model a three-move key exchange with a single initiator. As `Init` is independent of keying material from the caller's counterpart, our definition captures so-called *receiver obliviousness* [HKKP22] (sometimes *post-specified peers* [CK02]), corresponding to some, but not all, key exchange protocols in the literature.

The most novel part of our primitive is `BatchReceive` which in particular captures ephemeral key reuse when uploaded ephemeral keys are exhausted. In the case of key exhaustion, when a party comes back online, they execute `BatchReceive` several times (once per ephemeral state st_i), where the number of inputs of the form $(pk_j, prek_j, m_j)$ in a given `BatchReceive` call corresponds to how many times st_i is re-used. Otherwise, `BatchReceive` can be used as in standard AKE with a single value $(pk_j, prek_j, m_j)$ as input. Note that `BatchReceive` takes sender prekey bundles $prek_j$ as input to emphasise that they are generated via `Init` and not `Send`, although each bundle could in principle be contained in each m_j .

3.3.2 Security Model

We now describe the security model we consider for our DAKE, which extends existing models in some ways to support `BatchReceive`.

Parties and Sessions. We assume that there are n parties P_1, \dots, P_n (or $1, \dots, n$) where party P_i (resp. or i) is associated with long-term key pair (pk_i, sk_i) output by `KeyGen`. Each party runs one or more *sessions* (sometimes called *oracles* [BFG⁺20]), where the s -th session of P_i is denoted by π_i^s . Each session π_i^s is associated with the following local fields:

- `sid`, the session identifier or session id.
- `pid`, the partner identifier.
- `role` $\in \{\perp, \text{sender}, \text{receiver}\}$, the role of P_i .
- `status` $\in \{\perp, \text{accept}, \text{reject}\}$, the status of π_i^s .
- `k`, the session key.
- `st`, the session state.
- `r`, the session randomness.

Chapter 3. K-Waay: Fast and Deniable Post-Quantum X3DH Without Ring Signatures

All fields are initialised to \perp except r which is initialised to uniform randomness. A session either has role sender or receiver, and its counterpart, its *partner* pid , has the other role; note a receiver may have several counterparts (capturing ephemeral key reuse).

Fields pid , role , status and r in session π_i^s are set directly by the challenger, and the rest are (sometimes implicitly) set by the underlying DAKE algorithms called by the challenger. Moreover, in light of the definition of BatchReceive, sid , pid and k are *vectors* for a receiver ($\text{role} = \text{receiver}$); we sometimes write $\vec{\text{sid}}$, $\vec{\text{pid}}$ and \vec{k} for clarity to indicate this.

Suppose P_i is acting as a receiver. Initially, P_i calls `Init`, and then eventually calls `BatchReceive`. Before this point, one or more senders P_j (i.e., parties with $\text{role} = \text{sender}$) may call `Init` and then `Send` with respect to the output prek from P_i 's `Init` call (assuming honest message delivery), which outputs messages of the form m_j . Finally, P_i invokes `BatchReceive` with one or more m_j values as input. A party has $\text{status} = \text{accept}$ if and only if $k \neq \perp^4$, and stores any session state after calling `Init` and before setting $\text{status} \neq \perp$ due to a `Send` or `BatchReceive` call in st .

Partnering. We define partnering between two sessions to capture security using session identifiers:

Definition 25 (Partnering). For any (i, P_j, s, t) , we say that sessions π_i^s and π_j^t are *partners* if

1. $\pi_i^s.\text{role} \neq \pi_j^t.\text{role}$.
2. If $\pi_i^s.\text{role} = \text{sender}$, then $\pi_i^s.\text{pid} = j$ and $i \in \pi_j^t.\vec{\text{pid}}$. If $\pi_i^s.\text{role} = \text{receiver}$, then $j \in \pi_i^s.\vec{\text{pid}}$ and $i = \pi_j^t.\text{pid}$.
3. If $\pi_i^s.\text{role} = \text{sender}$, then $\pi_i^s.\text{sid} \in \pi_j^t.\vec{\text{sid}}$ and $\pi_i^s.\text{sid} \neq \perp$. If $\pi_i^s.\text{role} = \text{receiver}$, then $\pi_j^t.\text{sid} \in \pi_i^s.\vec{\text{sid}}$ and $\pi_j^t.\text{sid} \neq \perp$.

Looking ahead, this definition ensures that two sessions can only be partners if they both have $\text{status} = \text{accept}$. Our definition mainly differs from previous work in that there can be many senders (and thus partnered sessions) for a given receiver. Ignoring this aspect, our definition is only slightly different from that of Hashimoto et al. [HKKP22] in that we restrict sid to be not equal to \perp ; this is an artifact of the fact we model ‘four-move’ key exchange (including prekey uploading).

KIND Security Game. We first define key indistinguishability (KIND) and then define deniability separately. Following previous work, we define a KIND experiment played between a challenger C and adversary \mathcal{A} in text below. The experiment $\text{KIND}_{\text{DAKE}}^n(\mathcal{A})$ is parameterised by the DAKE DAKE and integer n , the number of parties (honest or otherwise) in the lifetime of the game’s execution. The game is divided into distinct phases defined as follows.

⁴In particular, `BatchReceive` may output several keys; as long as at least one of them is not \perp , the calling party accepts.

3.3 Deniable Authenticated Key Exchange

Setup. C first uniformly samples challenge bit $b \in \{0, 1\}$. Then, for each party P_i , C calls $(pk_i, sk_i) \xleftarrow{\$} \text{KeyGen}(1^\lambda)$ and provides $\{pk_1, \dots, pk_n\}$ and 1^λ as input to \mathcal{A} .

Phase 1. \mathcal{A} adaptively makes any number of the following queries in any order:

- $\text{EXEC}(i, s, \text{prek}, m)$: \mathcal{A} starts or runs the next step of execution in session π_i^s . In each call, C uses randomness tape $\pi_i^s.r$ as needed.
 - To start the execution in session π_i^s not previously started, \mathcal{A} calls $\text{EXEC}(i, s, \text{prek}, m)$ with special input $m = (\text{start}, \text{sender}, j)$ (resp. $(\text{start}, \text{receiver}, \vec{j})$) (where start is defined only in the context of this game) that, if not previously called, sets $\pi_i^s.\text{pid} = j$ (resp. $\pi_i^s.\text{pid} = \vec{j}$) and $\pi_i^s.\text{role} = \text{sender}$ (resp. $\pi_i^s.\text{role} = \text{receiver}$); observe input prek is ignored by C . Then, C invokes $(st_i, \text{prek}_i) \xleftarrow{\$} \text{Init}(sk_i, \text{role})$ and outputs prek_i to \mathcal{A} .
 - Given that P_i has started in π_i^s , $\pi_i^s.\text{status} = \perp$ and $\pi_i^s.\text{role} = \text{sender}$, when \mathcal{A} calls $\text{EXEC}(i, s, \text{prek}, \perp)$, C invokes $(k, m) \xleftarrow{\$} \text{Send}(sk_i, pk_j, st_i, \text{prek})$ (where $j = \pi_i^s.\text{pid}$), returns output m to \mathcal{A} and sets $\pi_i^s.\text{status}$ to reject (resp. accept) if $k = \perp$ (resp. $k \neq \perp$).
 - If $\pi_i^s.\text{role} = \text{receiver}$ and $\pi_i^s.\text{status} = \perp$, when \mathcal{A} calls $\text{EXEC}(i, s, \{s_j, \text{prek}_j, m_j\}_{j \in \vec{j}'})$, C aborts if $\vec{j}' \neq \pi_i^s.\text{pid}$ and otherwise invokes $k \leftarrow \text{BatchReceive}(sk_i, st_i, \{pk_j, \text{prek}_j, m_j\}_j)$ and outputs to \mathcal{A} \perp if BatchReceive fails (resp. nothing otherwise) and sets $\pi_i^s.\text{status}$ to reject (resp. accept).
- $\text{LTK}(i)$ outputs sk_i . P_i is hereafter *corrupted*.
- $\text{REGISTER}(pk_i, i)$ registers a new party P_i for $i > n$ not previously registered, sets their long-term public key to pk_i and distributes pk_i to all other oracles; P_i is immediately marked as *corrupted*.
- $\text{STATE}(i, s)$ outputs $\pi_i^s.\text{st}$, which is hereafter *revealed*.
- $\text{KEY}(i, s, j)$ outputs $\pi_i^s.k_j$ if $\pi_i^s.\text{role} = \text{receiver}$ and $\pi_i^s.\text{status} \neq \perp$ and otherwise outputs $\pi_i^s.k$.

Test. When \mathcal{A} decides to move to the next phase, it issues the following query TEST which (if successful) returns either a real or random key:

- $\text{TEST}(i, s, j)$: If $\pi_i^s.\text{status} \neq \text{accept}$, C returns \perp . Otherwise:
 - If $\pi_i^s.\text{role} = \text{sender}$, C aborts if $j \neq \pi_i^s.\text{pid}$, and otherwise returns either $\pi_i^s.k$ if $b = 0$ or a uniformly sampled key k if $b = 1$;
 - If $\pi_i^s.\text{role} = \text{receiver}$, C aborts if $j \notin \pi_i^s.\vec{\text{pid}}$ or $\pi_i^s.k_j = \perp$, and otherwise returns either $\pi_i^s.k_j$ if $b = 0$ or a uniformly sampled key k if $b = 1$.

At this point, π_i^s (which we say is with respect to key j if $\pi_i^s.\text{role} = \text{receiver}$) is said to be the *test session*.

Chapter 3. K-Waay: Fast and Deniable Post-Quantum X3DH Without Ring Signatures

Phase 2. \mathcal{A} adaptively issues queries as in Phase 1.

Guess, Freshness and Correctness. After Phase 2, \mathcal{A} outputs bit b' . Suppose that \mathcal{A} made query $\text{TEST}(i, s, j)$, i.e., π_i^s is the test session with respect to key j and $j \in \pi_i^s.\text{pid}$ (with equality at least when $\pi_i^s.\text{role} = \text{sender}$). The following *freshness* conditions are checked by C ; if any condition *is not satisfied*, C sets b' to a uniform bit (i.e., \mathcal{A} gains no advantage):

1. $\text{KEY}(i, s, j')$ has not been queried, where j' is arbitrary if $\pi_i^s.\text{role} = \text{sender}$ and $j' = j$ if $\pi_i^s.\text{role} = \text{receiver}$.
2. If π_i^s and π_j^t are partners, then $\text{KEY}(j, t, i')$ has not been queried, where i' is arbitrary if $\pi_i^s.\text{role} = \text{sender}$ and $i' = i$ if $\pi_i^s.\text{role} = \text{receiver}$.
3. P_i is not corrupted or $\pi_i^s.\text{st}$ has not been revealed.
4. If π_i^s and π_j^t are partners, then P_j is not corrupted or $\pi_j^t.\text{st}$ has not been revealed.
5. If π_i^s has no partner session, then P_j is not corrupted when $\pi_i^s.\text{status} = \perp$.
6. If π_i^s has no partner session, then if $\pi_i^s.\text{role} = \text{sender}$, for any session π_j^t such that prek_j was both output by $\text{Init}(\text{sk}_j, \text{receiver})$ and input to Send in π_i^s by C , P_j is not corrupted or $\pi_j^t.\text{st}$ is not revealed.
7. If π_i^s has no partner session, then if $\pi_i^s.\text{role} = \text{receiver}$, for any session π_j^t such that prek_j was both output by $\text{Init}(\text{sk}_j, \text{sender})$ and input to BatchReceive in π_i^s by C , $\pi_j^t.\text{st}$ is not revealed and $\pi_i^s.\text{st}$ is not revealed.

Then, the following *correctness* conditions are checked by C which, iterating over all relevant parties i, j, k , only consider the subset of sessions corresponding to *honest* protocol runs where \mathcal{A} faithfully follows the protocol specification. If any condition *is satisfied*, C sets $b = b'$ (i.e., \mathcal{A} wins):

1. There exist distinct sessions π_i^s and π_j^t such that $\pi_i^s.\text{role} = \pi_j^t.\text{role}$ and either (1) $\pi_i^s = \text{receiver}$ and $\pi_i^s.\text{sid}_j = \pi_j^t.\text{sid}_i$ or (2) $\pi_i^s.\text{sid} = \pi_j^t.\text{sid}$.
2. Assuming $\pi_i^s.\text{role} = \text{receiver}$, there exist sessions π_i^s with respect to key j and π_j^t that are partners such that $\pi_i^s.k_j \neq \pi_j^t.k$ (analogously when $\pi_i^s.\text{role} = \text{sender}$).
3. There exist distinct sessions π_i^s , π_j^t and π_k^u such that $\pi_i^s.\text{status} = \pi_j^t.\text{status} = \pi_k^u.\text{status} = \text{accept}$ and $\pi_i^s.\text{sid}_k = \pi_j^t.\text{sid}_k = \pi_k^u.\text{sid}$ (assuming i, j are receivers here but analogously in other cases).

Finally, the game outputs 1 if and only if $b = b'$.

Security is formally defined in Definition 26.

Definition 26 (DAKE Key Indistinguishability). We consider the KIND game described above. We say a DAKE is key indistinguishable if for all efficient adversaries \mathcal{A} and polynomially-bounded n (the total number of parties), we have

$$\text{Adv}_{\text{DAKE},n}^{\text{kind}}(\mathcal{A}) := \left| \Pr[\text{KIND}_{\text{DAKE}}^n(\mathcal{A}) \Rightarrow 1] - \frac{1}{2} \right| = \text{negl}.$$

Discussion. Following previous work, we define freshness conditions to prevent the adversary from mounting trivial attacks. Conditions 1 to 5 correspond exactly to the forward-secure variant of security in [HKKP22]. Due to the design of our DAKE K-Waay, we additionally restrict the adversary via conditions 6 and 7. The clauses in these conditions are essentially due to the fact that in K-Waay the only secret keying material required to call Send is an ephemeral split-KEM secret. For example, suppose that the tester π_i^s is the receiver. Due to the ‘symmetric’ nature of split-KEM, without these restrictions, revealing π_i^s .st allows the adversary to inject to P_i by simulating Send (akin to a key-compromise impersonation (KCI) attack using P_i ’s ephemeral state) and trivially distinguish. Consequently, we restrict session state exposure in this case.

Our model does not support randomness exposure or manipulation. As is standard, however, one can employ the NAXOS trick [LLM07] to obtain security given, e.g., randomness but not long-term keying material is exposed. Note also that we do not force `Init` to be called, e.g., by the sender or senders first or `Init` to be called by both the sender or senders and receiver before a party calls `Send` or `BatchReceive`.

Apart from the fact we make several extensions to typical AKE modelling to capture `BatchReceive`, the game is closest to that of Hashimoto et al. [HKKP22] except that we additionally enforce correctness checks as Brendel et al. [BFG⁺22b] do. To capture partnering, we consider partner and key identifiers that may be *vectors* for a receiver, such that several sender sessions may be partnered with a receiver session if, for a given sender session, it partners with a part/component of the receiver session. We do not capture semi-static keys explicitly as in Brendel et al.’s work [BFG⁺22b], although in principle they could be captured in `Init`. Like Hashimoto et al.’s [HKKP22], our game supports message injection, session state exposure or revealing (unlike Brendel et al.), session key exposure, long-term key exposure (corruption) and adversarial long-term key registration (also considered corruption). Following previous work, we allow the adversary to expose session states individually, rather than expose all states of a given party at once; security w.r.t. this weaker class of attacks is thus tightly implied by our notion. During execution, a single challenge test query is made by the adversary that reveals a real or random key output in some session. For `BatchReceive` which can output several keys, just one of the output keys are tested.

Trivial Attacks. We restrict the adversary’s behaviour to prevent ‘trivial’ attacks (e.g. directly revealing the challenge key) by defining *freshness* predicates. Due to our protocol’s design, our notion restricts more than the full forward security notion under session state exposure

Chapter 3. K-Waay: Fast and Deniable Post-Quantum X3DH Without Ring Signatures

defined by Hashimoto et al. [HKKP22]. Our freshness predicates imply weak forward secrecy and implicit authentication given session state exposure is not allowed (enforced in some recent works like [BHJ⁺15, CCG⁺19]). Brendel et al.'s model provide these guarantees but additionally protect against randomness exposure [BFG⁺22b], whereas we allow exposures on session states under some conditions unlike them.

3.3.3 Deniability

We next introduce our security notion for a deniable DAKE. To this end, we introduce security game $\text{DENY}_{\text{DAKE,Sim}}^{\text{exp}}$ in Figure 3.6.

Game $\text{DENY}_{\text{DAKE},n,\text{Sim}}^{\text{exp}}(\mathcal{A})$	Oracle $\text{CHAL}(i, j)$
1: $b \xleftarrow{\$} \{0, 1\}$	1: require $i \in [n] \wedge j \in [n]$
2: $L \leftarrow \emptyset$	2: $(k, m) \leftarrow (\perp, \perp)$
3: for $i \in [n]$:	3: $(\text{st}_i, \text{prek}_i) \xleftarrow{\$} \text{Init}(\text{sk}_i, \text{sender})$
4: $(\text{pk}_i, \text{sk}_i) \xleftarrow{\$} \text{KeyGen}(1^\lambda)$	4: $(\text{st}_j, \text{prek}_j) \xleftarrow{\$} \text{Init}(\text{sk}_j, \text{receiver})$
5: $L \leftarrow L \cup \{(\text{pk}_i, \text{sk}_i)\}$	5: if $b = 0$: $(k, m) \xleftarrow{\$} \text{Send}(\text{sk}_i, \text{pk}_j, \text{st}_i, \text{prek}_j)$
6: $b' \xleftarrow{\$} \mathcal{A}^{\text{CHAL}}(L)$	6: else : $(k, m) \xleftarrow{\$} \text{Sim}(\text{sk}_j, \text{pk}_i, \text{st}_j, \text{prek}_i, \text{prek}_j)$
7: return $1[b' = b]$	7: $T \leftarrow (\text{prek}_i, \text{prek}_j, m)$
	8: if $\text{exp} = \text{true}$: return (k, T, st_j)
	9: else : return (k, T)

Figure 3.6: Deniability game.

Definition 27 (DAKE deniability). We consider the game shown in Figure 3.6. We say a DAKE is DENY^{exp} for $\text{exp} \in \{\text{true}, \text{false}\}$ if there exists an efficient simulator Sim s.t. for all efficient adversaries \mathcal{A} and polynomially-bounded n , we have

$$\text{Adv}_{\text{DAKE,Sim,exp}}^{\text{deny}}(\mathcal{A}) := \left| \Pr[\text{DENY}_{\text{DAKE},n,\text{Sim}}^{\text{exp}}(\mathcal{A}) \Rightarrow 1] - \frac{1}{2} \right| = \text{negl}.$$

Our definition captures the following deniability property. Initially, the judge \mathcal{A} is given the long-term keys of all parties. \mathcal{A} then observes honest protocol runs between pairs of parties (via CHAL). Depending on the challenge bit b , either Send or a simulator Sim that takes as input the secret keying material of the receiver trying to frame the sender is executed in each run. Moreover, \mathcal{A} is given the prekey messages independent of b and, if the parameter exp is set to true, also the session state of the receiver in each protocol run. The goal of the adversary is to distinguish whether Send or Sim is being called.

Our notion $\text{DENY}^{\text{false}}$ corresponds most closely with that of Brendel et al. [BFG⁺22b] which was also adopted by Cremers et al. [CZ24]. Due to how Brendel et al.'s AKE primitive is defined, they also consider semi-static key pairs which are also given to the adversary. $\text{DENY}^{\text{true}}$

provides stronger deniability, corresponding in practice to a receiver who co-operates with a judge by handing over the entire contents of their device. Although incomparable formally, our DAKE would not be considered deniable under a notion like that of Brendel et al. [BFG⁺22b] since their protocol does not formally model long-term signatures. Note that our definition, like Brendel et al.'s [BFG⁺22b], can be straightforwardly converted to a “simulation-based” notion like Definition 23.

Finally, observe that our definition, like that of Brendel et al. [BFG⁺22b] does not consider deniability for the receiver but only for the sender. One could define such a notion, in which the goal is for the judge (adversary) to distinguish between the output of BatchReceive and a simulator Sim that has access to the long-term and ephemeral states of all corresponding senders and is given (honest) ciphertexts output by Send as input. Here, one could argue deniability for K-Waay using the security of the ephemeral KEM and then the KDF. A weaker definition would require the judge to distinguish between the output of Send *and* BatchReceive calls, and the output of a simulator Sim given the senders' states, which is trivial to satisfy (and thus we did not capture it) but is closer to offline semi-honest deniability (whereas the first notion sketched above is more ‘online’).

3.4 K-Waay: Post-Quantum X3DH from Split-KEM

In this section, we present our split-KEM-based protocol K-Waay before proving it satisfies our key indistinguishability and deniability security notions formalised in Section 3.3.

3.4.1 Construction

We present our DAKE K-Waay (Key-exchange With asynchrony, authentication and peer-deniability) in Figure 3.7.

Each party is associated with a long-term public/secret key pair which in K-Waay comprises of a signature and KEM key pair generated in KeyGen. In Init, ephemeral KEM and split-KEM keys for both parties are generated and the public keys are signed with the long-term signature key.

After initialisation, the sender P_i (sometimes called the initiator) invokes Send that takes the prekey prek_j output by the receiver P_j 's Init call as input. After verifying the signature in prek_j , P_i encapsulates to (1) the long-term KEM key of P_j ; (2) the ephemeral KEM key contained in prek_j ; and (3) the ephemeral split-KEM key contained in prek_j . Note that the split-KEM provides implicit authentication (without it, Send could be simulated without secrets). Moreover, the sender's prekey bundle and in particular the split-KEM public key is generated *before* Send is called for deniability. P_i then combines the encapsulated keys using a KDF and outputs the key and its message for P_j consisting of the three encapsulation ciphertexts.

Chapter 3. K-Waay: Fast and Deniable Post-Quantum X3DH Without Ring Signatures

KeyGen(1^λ)	Init(sk_i, role)
<pre> 1: // Long-term key generation 2: (kpk, ksk) $\xleftarrow{\\$}$ KeyGen_{LKEM}(1^λ) 3: (spk, ssk) $\xleftarrow{\\$}$ KeyGen_{Sig}(1^λ) 4: pk \leftarrow (spk, kpk) 5: sk \leftarrow (ssk, ksk) 6: return (pk, sk) </pre>	<pre> 1: // Prekey generation/upload 2: if role = sender : 3: (espki, esski) $\xleftarrow{\\$}$ KeyGen_{sKEM}(1^λ) 4: ekpki \leftarrow \perp 5: else : 6: (espki, esski) $\xleftarrow{\\$}$ KeyGen_{sKEM}(1^λ) 7: (ekpki, ekski) $\xleftarrow{\\$}$ KeyGen_{EKEM}(1^λ) 8: σ_i $\xleftarrow{\\$}$ Sign_{Sig}(sk_i.ssk, (espki, ekpki)) 9: preki \leftarrow (espki, ekpki, σ_i) 10: return (st_i = (esski, ekski, preki), preki) </pre>
Send($sk_i, pk_j, st_i, prek_j$)	BatchReceive($sk_i, st_i, S = \{pk_j, prek_j, m_j\}_j$)
<pre> 1: (esski, ekski, preki) \leftarrow st_i 2: (espkj, ekpkj, σ_j) \leftarrow prek_j 3: msg \leftarrow (espkj, ekpkj) 4: require Vrfy_{Sig}(pk_j.spk, msg, σ_j) 5: (K_ℓ, ct_ℓ) $\xleftarrow{\\$}$ Encaps_{LKEM}(pk_j.kpk) 6: (K_k, ct_k) $\xleftarrow{\\$}$ Encaps_{EKEM}(ekpk_j) 7: (K_s, ct_s) $\xleftarrow{\\$}$ Encaps_{sKEM}(espkj, esski) 8: m \leftarrow (ct_ℓ, ct_k, ct_s) 9: sid \leftarrow P_i P_j pk_i pk_j prek_i prek_j m 10: k \leftarrow KDF(K_ℓ, K_k, K_s, sid) 11: return (k, m) </pre>	<pre> 1: (esski, ekski, preki) \leftarrow st_i 2: fail \leftarrow false; k_j \leftarrow \perp 3: for j : (pk_j, prek_j, m_j) \in S : 4: (ct_ℓ, ct_k, ct_s) \leftarrow m_j 5: (espkj, ekpkj, σ_j) \leftarrow prek_j 6: if \negVrfy_{Sig}(pk_j.spk, (espkj, ekpkj), σ_j) : 7: k_j \leftarrow \perp 8: continue 9: K_ℓ \leftarrow Decaps_{LKEM}(sk_i.ksk, ct_ℓ) 10: K_k \leftarrow Decaps_{EKEM}(eksk_i, ct_k) 11: K_s \leftarrow Decaps_{sKEM}(espkj, esski, ct_s) 12: sid \leftarrow P_j P_i pk_j pk_i prek_j prek_i m_j 13: if K_s = \perp : fail \leftarrow true 14: if (K_ℓ = \perp) \vee (K_k = \perp) \vee (K_s = \perp) : k_j \leftarrow \perp 15: else : k_j \leftarrow KDF(K_ℓ, K_k, K_s, sid) 16: if fail : return $\perp^{ S }$ 17: else : return {k_j}_j </pre>

Figure 3.7: K-Waay: An X3DH-like DAKE from IND-CCA KEMs EKEM and LKEM, SUF-CMA signature scheme Sig and IND-1BatchCCA and UNF-1KCA split-KEM sKEM.

Receiving is analogous: receiver P_i verifies P_j 's prekey, decapsulates using its three respective secret keys and derives the session key (recall though that sender prekey bundles were not required in previous X3DH protocols). If P_i 's prekeys have run out, it is possible that multiple P_j 's have sent using the same prekey $prek_i$. In that case, P_i decapsulates for all sessions using the same secret keys but aborts if *any* split-KEM decapsulations failed in *any* of the sessions (a signature check failing does not however lead to the receiver aborting). We assume that for a

given $\text{BatchReceive}(\text{sk}_i, \text{st}_i, S)$ call, each element of S corresponds to a different party.

3.4.2 Security

Theorem 1. Consider $(1 - \delta_{\text{EKEM}})$ -correct IND-CCA KEM EKEM, $(1 - \delta_{\text{LKEM}})$ -correct IND-CCA KEM LKEM, $(1 - \delta_{\text{Sig}})$ -correct SUF-CMA signature scheme Sig and $(1 - \delta_{\text{sKEM}})$ -correct IND-1BatchCCA, UNF-1KCA split-KEM sKEM and triple PRF KDF used to build K-Waay (Figure 3.7). Then, we have that for polynomially-bounded n and every efficient adversary \mathcal{A} that makes at most q oracle queries, one can build an adversary \mathcal{B} such that

$$\begin{aligned} \text{Adv}_{\text{K-Waay}, n}^{\text{kind}}(\mathcal{A}) &\leq \frac{q}{3} \cdot (\delta_{\text{Sig}} + \delta_{\text{LKEM}} + \delta_{\text{EKEM}} + \delta_{\text{sKEM}}) + \\ &2q^2 \cdot (\epsilon_{\text{EKEM}} + \epsilon_{\text{LKEM}} + 2\epsilon_{\text{KDF}} + 2\epsilon_{\text{Sig}}) + q^3 \cdot (\epsilon_{\text{EKEM}} + \epsilon_{\text{LKEM}} + \epsilon_{\text{sKEM}} + 3\epsilon_{\text{KDF}}), \end{aligned}$$

where $\epsilon_{\text{EKEM}} = \text{Adv}_{\text{EKEM}}^{\text{ind-cca}}(\mathcal{B})$, $\epsilon_{\text{LKEM}} = \text{Adv}_{\text{LKEM}}^{\text{ind-cca}}(\mathcal{B})$, $\epsilon_{\text{Sig}} = \text{Adv}_{\text{Sig}}^{\text{suf-cma}}(\mathcal{B})$, $\epsilon_{\text{sKEM}} = \text{Adv}_{\text{sKEM}}^{\text{ind-1batchcca}}(\mathcal{B}) + \text{Adv}_{\text{sKEM}}^{\text{unf-1kca}}(\mathcal{B})$ and $\epsilon_{\text{KDF}} = \text{Adv}_{\text{KDF}}^{\text{3prf}}(\mathcal{B})$.

Proof. Our proof proceeds by constructing sequences of hybrids, which we first summarise. Let Game Γ_1 be exactly the KIND game played with respect to DAKE K-Waay (Figure 3.7). We first transition to Game Γ_2 , which differs from Game Γ_1 in that honest protocol runs, all Vrfy_{Sig} checks in BatchReceive calls are removed and Decaps calls are replaced by the output of the Encaps calls in the corresponding Send calls whenever they are consistent. To this end, we invoke the correctness of K-Waay's building blocks. Then, we transition to Game Γ_3 in which the challenger immediately outputs the session π_i^s that the adversary makes real-or-random challenge query $\text{TEST}(i, s, j^*)$ with respect to. We then partition \mathcal{A} 's possible executions of Game Γ_3 into several events.

Suppose π_i^s has a partner session (with respect to key j^* if $\pi_i^s.\text{role} = \text{receiver}$) (event E_p), say π_j^t . Observe that by definition of partnering and construction of the protocol (in particular by definition of sid), it follows that partnered sessions correspond to *honest* protocol runs. Then, considering π_i^s and π_j^t , if the receiver's session state, say $\pi_j^t.\text{st}$, is revealed (event $E_p \wedge E_{c1}$), we reduce to the IND-CCA security of the long-term KEM LKEM, since the freshness conditions imply P_j must not have been corrupted. Otherwise (event $E_p \wedge \neg E_{c1}$), we reduce to the IND-CCA security of the ephemeral KEM EKEM. After both cases, we transition to an unwinnable game by keying KDF with the now uniformly random key output by the respective KEM call, a transition we perform repeatedly and omit from this description hereafter. Otherwise (event $\neg E_p$), we consider whether party P_i in test session π_i^s has the role sender or receiver:

- $\pi_i^s.\text{role} = \text{sender}$ (event $\neg E_p \wedge E_s$): As P_j can only be corrupted after P_i accepts, we first use the SUF-CMA security of Sig to argue that P_i 's Send call in the test session must be with honestly-generated input (prek). Then, let E_{c2} be the event that P_j is corrupted. Given $\neg E_p \wedge E_s \wedge E_{c2}$, we reduce to the security of EKEM, since by freshness the state

Chapter 3. K-Waay: Fast and Deniable Post-Quantum X3DH Without Ring Signatures

π_j^t .st associated with prek_j must not have been exposed. Otherwise ($\neg E_p \wedge E_s \wedge \neg E_{c2}$) we reduce to the security of LKEM.

- π_i^s .role = receiver (event $\neg E_p \wedge \neg E_s$): As above, we first argue using SUF-CMA security that input prek_j used in the test session's BatchReceive call must have been honestly generated. Then by freshness, we know that neither π_i^s .st nor π_j^t .st associated with prek_j are revealed, in which case we first reduce to the UNF-1KCA security of sKEM to prevent injections on the split-KEM ciphertext, after which we reduce to the IND-1BatchCCA security of sKEM.

Let $\text{Adv}_{\text{DAKE},n}^{\text{gi}}(\mathcal{A})$ be the advantage of adversary \mathcal{A} in winning game Game Γ_i for relevant i which we introduce below. Furthermore, let $\text{Adv}_{\text{DAKE},n}^{\text{gi}}(\mathcal{A}, E)$ be the same advantage except restricted to event E , so in particular if $\text{Adv}_{\text{DAKE},n}^{\text{gi}}(\mathcal{A})$ is of the form $|\Pr[X] - \frac{1}{2}|$, $\text{Adv}_{\text{DAKE},n}^{\text{gi}}(\mathcal{A}, E)$ is of the form $|\Pr[X \wedge E] - \frac{1}{2}|$.

Game Γ_1 : This is the original key indistinguishability game.

Game Γ_2 : This differs from Game Γ_1 in that, in honest protocol runs, all signature verification calls in BatchReceive calls are removed and the output of Decaps calls are replaced with the output of the corresponding Encaps call in Send. It follows at this point that the three correctness checks in the KIND game evaluate to true. Since for a given BatchReceive(\cdot, \cdot, S) call there must be $|S|$ corresponding Send and Init calls, there are at most $q/3$ iterations of the **for** loop in BatchReceive (counting over all such calls in a given execution of Game Γ_1). It then follows from a standard hybrid argument and the correctness of Sig, LKEM, EKEM and sKEM that:

$$\text{Adv}_{\text{DAKE},n}^{\text{g1}}(\mathcal{A}) \leq \text{Adv}_{\text{DAKE},n}^{\text{g2}}(\mathcal{A}) + \frac{q}{3} \cdot (\delta_{\text{Sig}} + \delta_{\text{LKEM}} + \delta_{\text{EKEM}} + \delta_{\text{sKEM}}).$$

Game Γ_3 : This differs from Game Γ_3 in that the challenger immediately outputs the session π_i^s that the adversary \mathcal{A} calls TEST(i, s, j^*) with respect to. Noting that there are at most q such possible sessions and applying a standard argument, it follows that:

$$\text{Adv}_{\text{DAKE},n}^{\text{g2}}(\mathcal{A}) \leq q \cdot \text{Adv}_{\text{DAKE},n}^{\text{g3}}(\mathcal{A}).$$

Case 1: Test session π_i^s is partnered (Game Γ_{3a} and Game Γ_{3b}):

Game $\Gamma_{3a.1}$: Let E_p be the event that test session π_i^s has a partner, say π_j^t . Let E_{c1} be the event that the ephemeral state st of the receiver (in π_i^s and π_j^t) is revealed. Games Game $\Gamma_{3a.i}$ are defined given $E_p \wedge E_{c1}$. Game $\Gamma_{3a.1}$ differs from Game Γ_3 in that the game initially outputs π_j^t , the partner of π_i^s (observe that $j = j^*$ where j^* is defined in the previous hop), as well as a bit indicting whether π_i^s is the sender or receiver. By the same reasoning as

above, we have

$$\text{Adv}_{\text{DAKE},n}^{\text{g3}}(\mathcal{A}, E_p \wedge E_{c1}) \leq 2q \cdot \text{Adv}_{\text{DAKE},n}^{\text{g3a.1}}(\mathcal{A}).$$

Game $\Gamma_{3a.2}$: Game $\Gamma_{3a.2}$ differs from Game $\Gamma_{3a.1}$ in that the output key K in the call to LKEM.Encaps and the corresponding LKEM.Decaps call or calls (which are guaranteed to exist given E_p , and K is identical by definition of Game Γ_2) made in the test and partner sessions with respect to the receiver's public key and secret key, respectively, are replaced with a key k uniformly sampled by the challenger. Observe that since E_{c1} holds, by freshness, P_j cannot be corrupted, and thus we reduce to the security of LKEM.

Let \mathcal{A}' be a IND-CCA adversary who simulates for Game $\Gamma_{3a.1}$ /Game $\Gamma_{3a.2}$ adversary \mathcal{A} as follows. Let pk be the IND-CCA challenge public key, (ct^*, K^*) be the challenge ciphertext and key respectively.

In the Setup phase, \mathcal{A}' uniformly samples bit b_{sim} , calls $(\text{pk}_\ell, \text{sk}_\ell) \stackrel{s}{\leftarrow} \text{KeyGen}(1^\lambda)$ locally for $\ell \neq k$ where k is the sender, sets $\text{pk}_k \leftarrow \text{pk}$, and returns $\{\text{pk}_1, \dots, \text{pk}_n\}$ and 1^λ to \mathcal{A} . Observe here (and later for Game $\Gamma_{3b.2}$) that, since E_p holds, we have matching sid values for test session π_i^s and partner π_j^t . Note by construction of sid, the presence of substring $\text{prek}_i || \text{prek}_j$ and m in the common value sid implies that Send must have been called honestly in π_i^s and also in BatchReceive for tuple $(\text{pk}_j, \text{prek}_j, m_j)$ in π_j^t for E_p to hold. Thus, we do not need to consider injections in the test session itself (although we have to in general in the BatchReceive call).

Before proceeding, we argue that \mathcal{A}' can simulate on behalf of parties with a maliciously-registered long-term key locally, which applies here and in the rest of the proof. Since π_i^s is partnered, as argued above, π_i^s and π_j^t correspond to honest (completed) executions, and so neither P_i and P_j can be malicious. For unpartnered sessions, since Send and BatchReceive cannot be called by the game, the test session π_i^s cannot be corrupted itself (since testing requires $\pi_i^s.\text{status} \neq \perp$), and otherwise condition 5 restricts the non-tested party P_j from being corrupted, thus precluding its key from being registered maliciously. Finally, computation involving messages or prekey bundles from maliciously-registered parties does not require any secret material not already known to \mathcal{A}' .

In Phase 1, when \mathcal{A} calls $\text{EXEC}(k', \cdot, \cdot, \cdot)$ where k' corresponds to the sender in π_i^s and π_j^t and the challenger is supposed to invoke Send, \mathcal{A}' replaces the call to $\text{Encaps}_{\text{LKEM}}$ with the output (ct^*, K^*) , and otherwise simulates locally. When \mathcal{A} calls $\text{EXEC}(k, \cdot, \cdot, \cdot)$ corresponding to the receiver in π_i^s and π_j^t and the challenger is supposed to invoke BatchReceive, \mathcal{A}' replaces the output of the relevant $\text{Decaps}_{\text{LKEM}}$ calls corresponding to either i or j , depending on who is the receiver, with K^* and the output of other calls $\text{Decaps}_{\text{LKEM}}$ with the output obtained from $\text{DEC}(\cdot)$; \mathcal{A}' otherwise simulates locally.

In the Test phase, i.e., when \mathcal{A} calls $\text{TEST}(i, s, j)$, \mathcal{A}' simulates with respect to bit b_{sim} . \mathcal{A}' then simulates Phase 2 as above and the rest of the game locally, ultimately outputting the same bit as \mathcal{A} ; observe that \mathcal{A}' can efficiently evaluate the freshness conditions. Since \mathcal{A}' perfectly simulates Game $\Gamma_{3a.1}$ when playing with respect to challenge bit 0 and Game $\Gamma_{3a.2}$

Chapter 3. K-Waay: Fast and Deniable Post-Quantum X3DH Without Ring Signatures

when it is 1, it follows that:

$$\text{Adv}_{\text{DAKE},n}^{\text{g3a.1}}(\mathcal{A}) \leq \text{Adv}_{\text{DAKE},n}^{\text{g3a.2}}(\mathcal{A}) + \text{Adv}_{\text{LKEM}}^{\text{ind-cca}}(\mathcal{A}').$$

Game $\Gamma_{3a.3}$: This differs from Game $\Gamma_{3a.2}$ in that, for the test and partner sessions, the call to KDF made in Send and the corresponding calls made in BatchReceive with respect to ciphertext ct_ℓ output by Send are replaced with uniformly sampled keys. Let \mathcal{A}' be a PRF adversary playing with respect to KDF keyed in its first argument simulating for Game $\Gamma_{3a.2}$ /Game $\Gamma_{3a.3}$ adversary \mathcal{A} as follows. \mathcal{A}' simulates locally all calls except the Send and BatchReceive calls made in π_i^s and π_j^t , where it replaces the relevant calls $\text{KDF}(K_\ell, K_k, K_s, \text{sid})$ with the call $\text{EVAL}(K_k, K_s, \text{sid})$. Since K_ℓ is uniform (by definition of Game $\Gamma_{3a.2}$) and, by definition of freshness, K_ℓ is not revealed to \mathcal{A} , the simulation is perfect and we have:

$$\text{Adv}_{\text{DAKE},n}^{\text{g3a.2}}(\mathcal{A}) \leq \text{Adv}_{\text{DAKE},n}^{\text{g3a.3}}(\mathcal{A}) + \text{Adv}_{\text{KDF}}^{\text{3prf}}(\mathcal{A}').$$

Finally, we have $\text{Adv}_{\text{DAKE},n}^{\text{g3a.3}}(\mathcal{A}) = 0$ since the output of TEST is identical regardless of the challenge bit and it is not otherwise used by the challenger or leaked to the adversary.

Game $\Gamma_{3b.1}$: We now consider the case when $E_p \wedge \neg E_{c1}$, i.e., the case where the receiver's session state st in π_i^s and π_j^t is not revealed. Game $\Gamma_{3b.1}$ differs from Game Γ_3 in that the game initially outputs π_j^t , the partner of π_i^s , as well as a bit indicating whether π_i^s is the sender or receiver. Since Game $\Gamma_{3b.1}$ is exactly Game $\Gamma_{3a.1}$, we have

$$\text{Adv}_{\text{DAKE},n}^{\text{g3}}(\mathcal{A}, E_p \wedge \neg E_{c1}) \leq q \cdot \text{Adv}_{\text{DAKE},n}^{\text{g3b.1}}(\mathcal{A}).$$

Game $\Gamma_{3b.2}$: In Game $\Gamma_{3b.2}$, the output of $\text{Encaps}_{\text{EKEM}}$ and the corresponding $\text{Decaps}_{\text{EKEM}}$ call or calls in the test session are replaced with a uniformly random key k . IND-CCA adversary \mathcal{A}' simulates for Game $\Gamma_{3b.1}$ /Game $\Gamma_{3b.2}$ adversary \mathcal{A} as follows. \mathcal{A}' follows the same broad approach as the adversary defined in the hop between Game $\Gamma_{3a.1}$ and Game $\Gamma_{3a.2}$. In particular, \mathcal{A}' simulates the receiver in their session's call to Init except it uses the IND-CCA challenge public key pk , replaces the output of EKEM in the test session Encaps and the corresponding Decaps calls with the challenge ciphertext and key and replaces other Decaps calls with calls to oracle DEC . By the same reasoning as before, it follows that:

$$\text{Adv}_{\text{DAKE},n}^{\text{g3b.1}}(\mathcal{A}) \leq \text{Adv}_{\text{DAKE},n}^{\text{g3b.2}}(\mathcal{A}) + \text{Adv}_{\text{EKEM}}^{\text{ind-cca}}(\mathcal{A}').$$

Game $\Gamma_{3b.3}$: This replaces the relevant outputs of KDF in the test session with a uniformly random key. As in Game $\Gamma_{3a.3}$, this game is now unwinnable, i.e., $\text{Adv}_{\text{DAKE},n}^{\text{g3b.3}}(\mathcal{A}) = 0$. As before, we reduce to the security of KDF, except now we key KDF in the PRF game with the second

argument K_k . We then arrive at:

$$\text{Adv}_{\text{DAKE},n}^{\text{g3b.2}}(\mathcal{A}) \leq \text{Adv}_{\text{DAKE},n}^{\text{g3b.3}}(\mathcal{A}) + \text{Adv}_{\text{KDF},n}^{\text{3prf}}(\mathcal{A}').$$

Case 2: Test session π_i^s is unpartnered and $\pi_i^s.\text{role} = \text{sender}$ (Game Γ_{3c}):

Game $\Gamma_{3c.1}$: Let $\pi_i^s.\text{pid} = j$ and E_s be the event that $\pi_i^s.\text{role} = \text{sender}$. Game $\Gamma_{3c.1}$ differs from Game Γ_3 in that the challenger immediately outputs j . By a standard argument, we have:

$$\text{Adv}_{\text{DAKE},n}^{\text{g3}}(\mathcal{A}, \neg E_p \wedge E_s) \leq \min\{q, n\} \cdot \text{Adv}_{\text{DAKE},n}^{\text{g3c.1}}(\mathcal{A}).$$

Game $\Gamma_{3c.2}$: This differs from Game $\Gamma_{3c.1}$ in that the challenger aborts if the call $\text{Send}(\text{sk}_j, \text{pk}_j, \text{st}_i, \text{prek}_j)$ in the test session is such that prek_j was not previously output by a call to $\text{Init}(\text{sk}_j, \text{receiver})$. Note that by freshness condition 5 that P_j must not be corrupted until after $\pi_i^s.\text{status}$ is changed from \perp , which, by definition of E_s , means until after it is set to accept. In order for Send to accept on input $\text{prek}_j = (\text{espk}_j, \text{ekpk}_j, \sigma_j)$ not previously output by $\text{Init}(\text{sk}_j, \text{receiver})$ (and thus for the game to abort), \mathcal{A} needs to find a different prek_j such that $\text{Vrfy}(\text{pk}_j.\text{spk}, (\text{espk}_j, \text{ekpk}_j), \sigma_j)$ (by construction of Send). Using this observation, we reduce to the SUF-CMA security of Sig .

Let \mathcal{A}' be a SUF-CMA adversary simulating for Game $\Gamma_{3c.1}$ adversary \mathcal{A} . Let pk be the SUF-CMA challenge public key. In the Setup phase, \mathcal{A}' sets $\text{pk}_j = \text{pk}$ and otherwise simulates locally. In particular, unlike in previous hops, \mathcal{A}' also samples the random Game $\Gamma_{3c.1}$ bit. In each subsequent phase, for each call $\text{EXEC}(j, u, \cdot, m)$ such that $m = (\text{start}, \text{role}, \cdot)$, \mathcal{A}' replaces the $\text{Sign}_{\text{Sig}}(\text{sk}_j.\text{ssk}, (\text{espk}_j, \text{ekpk}_j))$ call in $\text{Init}(\text{sk}_j, \text{receiver})$ by a call to $\text{SIGN}((\text{espk}_j, \text{ekpk}_j))$, and otherwise simulates the call locally. When the challenger calls $\text{Send}(\cdot, \text{pk}_j, \cdot, \text{prek}_j)$ where $\text{prek}_j = (\text{espk}_j, \text{ekpk}_j, \sigma_j)$, \mathcal{A}' checks whether (1) $(\text{espk}_j, \text{ekpk}_j)$ was previously queried to SIGN which output σ_j and (2) $\text{Vrfy}_{\text{Sig}}(\text{pk}, (\text{espk}_j, \text{ekpk}_j), \sigma_j) = 1$. Given (1) and (2) both hold, \mathcal{A}' returns $(m, \sigma) = ((\text{espk}_j, \text{ekpk}_j), \sigma_j)$ to its challenger. \mathcal{A}' otherwise simulates locally, aborting if \mathcal{A} outputs a bit. The simulation is perfect and it follows that:

$$\text{Adv}_{\text{DAKE},n}^{\text{g3c.1}}(\mathcal{A}) \leq \text{Adv}_{\text{DAKE},n}^{\text{g3c.2}}(\mathcal{A}) + \text{Adv}_{\text{Sig}}^{\text{suf-cma}}(\mathcal{A}').$$

Game $\Gamma_{3c.3}$: In Game $\Gamma_{3c.3}$, the challenger initially outputs π_j^t , where π_j^t is the session that prek is output by $\text{Init}(\text{sk}_j, \text{receiver})$ and input to the Send call in test session π_i^s . By a standard failure event argument, we have:

$$\text{Adv}_{\text{DAKE},n}^{\text{g3c.2}}(\mathcal{A}) \leq q \cdot \text{Adv}_{\text{DAKE},n}^{\text{g3c.3}}(\mathcal{A}).$$

Game $\Gamma_{3c.4a.1}$: Let E_{c2} be the event that P_j is corrupted. We construct hybrid sequence Game $\Gamma_{3c.4a}$ (resp. Game $\Gamma_{3c.4b}$) to deal with the case that E_{c2} holds (resp. does not hold).

Chapter 3. K-Waay: Fast and Deniable Post-Quantum X3DH Without Ring Signatures

Game $\Gamma_{3c.4a.1}$ differs from Game $\Gamma_{3c.3}$ in that the output of $\text{Encaps}_{\text{EKEM}}$ in the Send call in test session π_i^s and of the (possible) corresponding $\text{Decaps}_{\text{EKEM}}$ calls in π_j^t are replaced with uniformly random output. By freshness, P_j 's session state $\pi_j^t.\text{st}$ associated with prek input to the test Send call is not revealed.

IND-CCA adversary \mathcal{A}' simulates for Game $\Gamma_{3c.4a.1}$ adversary \mathcal{A} as follows. Let (pk, k, ct) the challenge public key, key and corresponding ciphertext (respectively) of \mathcal{A}' . \mathcal{A}' embeds pk in session π_j^t by replacing the public key output by $\text{KeyGen}_{\text{EKEM}}$ in $\text{Init}(\text{sk}_j, \text{receiver})$ with pk , which outputs prek_j . Upon prek_j being input to Send in the test session, \mathcal{A}' replaces the output of $\text{Encaps}_{\text{EKEM}}$ with (k, ct) . When the challenger calls $\text{BatchReceive}(\text{sk}_j, \cdot, \{\cdot, \cdot, m_{j'} = (\cdot, ct', \cdot)\}_{j'})$ in session π_j^t , if $ct' = ct$, \mathcal{A}' replaces the output of $\text{Decaps}_{\text{EKEM}}$ with k ; else, \mathcal{A}' replaces the call $\text{Decaps}_{\text{EKEM}}(\cdot, ct')$ with the call $\text{DEC}(ct')$. \mathcal{A}' otherwise simulates locally and outputs the same bit as \mathcal{A} . By similar reasons to before, we have:

$$\text{Adv}_{\text{DAKE},n}^{\text{g}3c.3}(\mathcal{A}, E_{c2}) \leq \text{Adv}_{\text{DAKE},n}^{\text{g}3c.4a.1}(\mathcal{A}) + \text{Adv}_{\text{EKEM}}^{\text{ind-cca}}(\mathcal{A}').$$

Game $\Gamma_{3c.4a.2}$: This replaces the output of KDF in the Send and BatchReceive calls as before in the test session and π_j^t with uniformly random keys. Otherwise, by the exact same argument as for Game $\Gamma_{3b.3}$ (where since the two sessions are unpartnered, they must diverge in sid, and so both EVAL queries are allowed in the simulation), we have $\text{Adv}_{\text{DAKE},n}^{\text{g}3c.4a.2}(\mathcal{A}) = 0$ and

$$\text{Adv}_{\text{DAKE},n}^{\text{g}3c.4a.1}(\mathcal{A}) \leq \text{Adv}_{\text{DAKE},n}^{\text{g}3c.4a.2}(\mathcal{A}) + \text{Adv}_{\text{KDF}}^{\text{3prf}}(\mathcal{A}').$$

Game $\Gamma_{3c.4b.1}$: We assume $\neg E_{c2}$, i.e., that P_j is not corrupted. We reduce to the IND-CCA security of LKEM. The reduction follows the same high-level strategy as previous hops (embedding the challenge pk in pk_j and the challenge in the test Send call and possibly the corresponding BatchReceive call), noting that non-challenge $\text{Decaps}_{\text{LKEM}}(\text{sk}_j, \cdot)$ queries are replaced with calls to DEC. We then have:

$$\text{Adv}_{\text{DAKE},n}^{\text{g}3c.3}(\mathcal{A}, \neg E_{c2}) \leq \text{Adv}_{\text{DAKE},n}^{\text{g}3c.4b.1}(\mathcal{A}) + \text{Adv}_{\text{LKEM}}^{\text{ind-cca}}(\mathcal{A}').$$

Game $\Gamma_{3c.4b.2}$: As in Game $\Gamma_{3c.4a.2}$, this replaces the output of KDF in the Send and BatchReceive calls in π_i^s and π_j^t with a uniformly random key. As argued several times above, it follows that $\text{Adv}_{\text{DAKE},n}^{\text{g}3c.4b.2}(\mathcal{A}) = 0$ and

$$\text{Adv}_{\text{DAKE},n}^{\text{g}3c.4b.1}(\mathcal{A}) \leq \text{Adv}_{\text{DAKE},n}^{\text{g}3c.4b.2}(\mathcal{A}) + \text{Adv}_{\text{KDF}}^{\text{3prf}}(\mathcal{A}').$$

Case 3: Test session π_i^s is unpartnered and $\pi_i^s.\text{role} = \text{receiver}$ (Game Γ_{3d}):

Game $\Gamma_{3d.1}$: Game $\Gamma_{3d.1}$ differs from Game Γ_3 in that the challenger immediately out-

3.4 K-Waay: Post-Quantum X3DH from Split-KEM

puts j , the third argument in \mathcal{A} 's $\text{TEST}(i, s, j)$ call. As for Game $\Gamma_{3c.1}$, we have

$$\text{Adv}_{\text{DAKE},n}^{\text{g}3}(\mathcal{A}, \neg E_s) \leq \min\{q, n\} \cdot \text{Adv}_{\text{DAKE},n}^{\text{g}3\text{d}.1}(\mathcal{A}).$$

Game $\Gamma_{3d.2}$: This differs from Game $\Gamma_{3d.1}$ in that the challenger aborts if the call $\text{BatchReceive}(\text{sk}_i, \text{st}_i, \{\text{pk}_{j'}, \text{prek}_{j'}, m\}_{j'})$ in the test session is such that prek_j was not previously output by a call to $\text{Init}(\text{sk}_j, \text{sender})$. As in Game $\Gamma_{3c.2}$, P_j must not be corrupted until after π_i^s .status is set to accept. By reducing to SUF-CMA security essentially as in Game $\Gamma_{3c.2}$, it follows that:

$$\text{Adv}_{\text{DAKE},n}^{\text{g}3\text{d}.1}(\mathcal{A}) \leq \text{Adv}_{\text{DAKE},n}^{\text{g}3\text{d}.2}(\mathcal{A}) + \text{Adv}_{\text{Sig}}^{\text{suf-cma}}(\mathcal{A}).$$

Game $\Gamma_{3d.3}$: This differs from Game $\Gamma_{3d.2}$ in that the challenger initially outputs π_j^t , the session that generated prek_j which formed part of the input to BatchReceive in the test session π_i^s . By a standard argument we have:

$$\text{Adv}_{\text{DAKE},n}^{\text{g}3\text{d}.2}(\mathcal{A}) \leq q \cdot \text{Adv}_{\text{DAKE},n}^{\text{g}3\text{d}.3}(\mathcal{A}).$$

Game $\Gamma_{3d.4}$: This differs from Game $\Gamma_{3d.3}$ in that the challenger aborts if the $\text{Send}(\text{sk}_j, \text{pk}_i, \text{st}_j, \text{prek}_i)$ call in session π_j^t (if it exists) and the relevant component in the BatchReceive call in the test session π_i^s were not both with respect to honestly generated split-KEM keying material (namely, an honestly generated split-KEM public key from prek_i and prek_j from the previous hop) and the same split-KEM ciphertext. By freshness, neither of the two ephemeral states π_i^s .st and π_j^t .st are revealed. Consequently, we reduce to the UNF-1KCA security of split-KEM sKEM.

UNF-1KCA adversary \mathcal{A}' simulates for Game $\Gamma_{3d.3}$ /Game $\Gamma_{3d.4}$ adversary \mathcal{A} as follows. Let $(\text{pk}_A, \text{pk}_B)$ be the two challenge public keys given to \mathcal{A}' . In the $\text{Init}(\text{sk}_i, \text{receiver})$ call in session π_i^s , \mathcal{A}' simulates except replaces the call to $\text{KeyGen}_{\text{sKEM}}$ by pk_A . Similarly, in the $\text{Init}(\text{sk}_j, \text{sender})$ call in session π_j^t , \mathcal{A}' replaces $\text{KeyGen}_{\text{sKEM}}$ by pk_B . In the $\text{Send}(\dots, \text{prek})$ call in session π_j^t where $\text{prek} = (\text{espk}, \dots)$, \mathcal{A}' outputs espk to its UNF-1KCA challenger, receives $(\text{pk}_A, \text{pk}_B, \text{ct}, K_B)$ from its challenger, and replaces the call to $\text{Encaps}_{\text{sKEM}}$ with tuple (ct, K_B) . Finally, when the $\text{BatchReceive}(\text{sk}_i, \cdot, \{\cdot, \text{prek}_{j'}, m = (\cdot, \cdot, \text{ct}_s)\}_{j'})$ call in test session π_i^s is made, \mathcal{A}' outputs ct_s corresponding to $j' = j$ to its challenger. As the simulation is perfect and the probability that \mathcal{A}' wins is exactly the probability that 1) $(\text{ct}, \text{pk}_A) \neq (\text{ct}_s, \text{pk})$ and 2) relevant $\text{Decaps}_{\text{sKEM}}$ call in BatchReceive outputs $k \neq \perp$, it follows by a standard failure event argument that:

$$\text{Adv}_{\text{DAKE},n}^{\text{g}3\text{d}.3}(\mathcal{A}) \leq \text{Adv}_{\text{DAKE},n}^{\text{g}3\text{d}.4}(\mathcal{A}) + \text{Adv}_{\text{sKEM}}^{\text{unf-1kca}}(\mathcal{A}').$$

Game $\Gamma_{3d.5}$: This differs from Game $\Gamma_{3d.4}$ in that the output k of the relevant test session split-KEM decapsulation and the corresponding encapsulation (if it exists) are both replaced by a uniformly random key. Note that by definition of Game $\Gamma_{3d.4}$, \mathcal{A} can only input an honestly generated split-KEM ciphertext to the BatchReceive call in the test session from P_j

Chapter 3. K-Waay: Fast and Deniable Post-Quantum X3DH Without Ring Signatures

and that the split-KEM public key in P_j 's corresponding Send call (if it exists) must be honestly generated. We therefore reduce to the IND-1BatchCCA security of sKEM. We embed the IND-1BatchCCA keys pk_A, pk_B in the simulation as in the previous hop. When \mathcal{A} queries $\text{EXEC}(i, s, S = \{s_{j'}, \text{prek}_{j'}, m_{j'}\}_{j'})$, \mathcal{A}' replaces all $\text{Decaps}_{\text{sKEM}}$ calls involving sk_A *except* the call corresponding to the test session by the output of its query to oracle BatchDEC, replaces this final $\text{Decaps}_{\text{sKEM}}$ call with the IND-1BatchCCA challenge key and otherwise simulates locally. It follows that:

$$\text{Adv}_{\text{DAKE},n}^{\text{g3d.4}}(\mathcal{A}) \leq \text{Adv}_{\text{DAKE},n}^{\text{g3d.5}}(\mathcal{A}) + \text{Adv}_{\text{sKEM}}^{\text{ind-1batchcca}}(\mathcal{A}').$$

Game $\Gamma_{3d.6}$: This game replaces the relevant invocation of KDF in the test session's BatchReceive call by a uniformly random value. Note as usual that $\text{Adv}_{\text{DAKE},n}^{\text{g3d.6}}(\mathcal{A}) = 0$. By keying KDF in its third argument as a PRF and a standard argument it follows that:

$$\text{Adv}_{\text{DAKE},n}^{\text{g3d.5}}(\mathcal{A}) \leq \text{Adv}_{\text{DAKE},n}^{\text{g3d.6}}(\mathcal{A}) + \text{Adv}_{\text{KDF}}^{\text{3prf}}(\mathcal{A}').$$

Finally note that by the triangle inequality, we have, among other inequalities:

$$\text{Adv}_{\text{DAKE},n}^{\text{g3}}(\mathcal{A}) \leq \text{Adv}_{\text{DAKE},n}^{\text{g3}}(\mathcal{A}, E_p) + \text{Adv}_{\text{DAKE},n}^{\text{g3}}(\mathcal{A}, \neg E_p).$$

The result follows using this observation and by combining the sequences of hybrids together in a standard way. \square

We now prove the deniability of K-Waay.

Theorem 2. Consider deniable split-KEM sKEM with simulator Sim_{sKEM} used to build K-Waay (Figure 3.7). Then, we have that for every efficient adversary \mathcal{A} that makes at most q oracle queries, there exists an efficient Sim s.t. one can build an adversary \mathcal{B} such that for $\text{exp} \in \{\text{true}, \text{false}\}$ we have:

$$\text{Adv}_{\text{K-Waay}, \text{Sim}, \text{exp}}^{\text{deny}}(\mathcal{A}) \leq q \cdot \text{Adv}_{\text{sKEM}, \text{Sim}_{\text{sKEM}}}^{\text{deny}}(\mathcal{B}).$$

Proof. We construct a sequence of hybrids and reduce to the deniability of sKEM (i.e., $\text{DENY}_{\text{sKEM}, \text{Sim}_{\text{sKEM}}}$ security) in each step. Before this, we define the simulator Sim that we use in the proof, which uses the simulator Sim_{sKEM} (which must exist if sKEM is deniable) as a subroutine.

Observe in K-Waay that, given an honestly generated prek_j , any party with knowledge only of public keying material can simulate all steps in Send except for the $\text{Encaps}_{\text{sKEM}}$ call which requires sender P_i 's secret key. Thus, our simulator Sim (Figure 3.8) simulates these steps and since it takes the receiver's key sk_j as input it can also invoke the deniability simulator Sim_{sKEM} to complete the call.

Sim($sk_j, pk_i, st_j, prek_i, prek_j$)	
1:	$(essk_j, eksk_j, prek_j) \leftarrow st_j$
2:	$(espk_i, ekpk_i, \sigma_i) \leftarrow prek_i$
3:	$(espk_j, ekpk_j, \sigma_j) \leftarrow prek_j$
4:	$spk \leftarrow \text{GetPK}(sk_j.ssk)$
5:	require $\text{Vrfy}_{\text{Sig}}(spk, (espk_j, ekpk_j), \sigma_j) = 1$
6:	$(K_\ell, ct_\ell) \xleftarrow{\$} \text{Encaps}_{\text{LKEM}}(pk_j.kpk)$
7:	$(K_k, ct_k) \xleftarrow{\$} \text{Encaps}_{\text{EKEM}}(ekpk_j)$
8:	$(K_s, ct_s) \xleftarrow{\$} \text{Sim}_{\text{sKEM}}(espk_i, essk_j)$
9:	$m \leftarrow (ct_\ell, ct_k, ct_s)$
10:	$sid \leftarrow P_i P_j pk_i pk_j prek_i prek_j m$
11:	$k \leftarrow \text{KDF}(K_\ell, K_k, K_s, sid)$
12:	return (k, m)

Figure 3.8: Simulator Sim for the deniability game where we assume function GetPK(sk) that takes a signature secret key as input and outputs the corresponding public key.

Let Γ_0 be the DAKE DENY game instantiated with K-Waay. For $i \in [q]$, let Γ_i be the same as Γ_{i-1} except that in the i -th CHAL call, the call to Send is replaced with a call to Sim. Note that the steps executed in Send only differ in that it calls $\text{Encaps}_{\text{sKEM}}$ rather than Sim_{sKEM} .

For $i \in [q]$, let \mathcal{B} be a split-KEM DENY adversary with input $(pk_A, pk_B, sk_B, K, ct)$ from its challenger playing $\text{DENY}^{\text{REAL}}$ given \mathcal{A} is playing Γ_{i-1} and DENY^{SIM} if it is playing Γ_i . \mathcal{B}' locally simulates long-term public key generation and the first $i-1$ calls to CHAL. When \mathcal{A} makes their i -th call to CHAL, \mathcal{B} simulates CHAL until it reaches the **if** statement except that it replaces the output of calls KeyGenA/KeyGenB calls in Init calls with pk_A/pk_B . Then, instead of executing the **if/else** block in CHAL, \mathcal{B} simulates Sim except that it replaces the output of the call to Sim_{sKEM} with (K, ct) . \mathcal{B} then simulates locally, and returns (k, T, st_r) (where st_r contains sk_B) if $exp = \text{true}$ and returns (k, T) otherwise. \mathcal{B} continues simulating locally and finally outputs the same bit as \mathcal{A} . Noting that DAKE deniability game $\text{DENY}_{\text{K-Waay, Sim}}$ considers only honest executions of K-Waay, it follows that the simulation is perfect, and so by $\text{DENY}_{\text{sKEM}}$ security we have

$$\left| \text{Adv}_{\text{DAKE}}^{\Gamma_{i-1}}(\mathcal{A}) - \text{Adv}_{\text{DAKE}}^{\Gamma_i}(\mathcal{A}) \right| \leq \text{Adv}_{\text{sKEM, Sim}_{\text{sKEM}}}^{\text{deny}}(\mathcal{B}).$$

By application of the triangle inequality and telescoping sums:

$$\left| \text{Adv}_{\text{DAKE}}^{\Gamma_0}(\mathcal{A}) - \text{Adv}_{\text{DAKE}}^{\Gamma_q}(\mathcal{A}) \right| \leq q \cdot \text{Adv}_{\text{sKEM, Sim}_{\text{sKEM}}}^{\text{deny}}(\mathcal{B}).$$

To complete the proof, observe that $\text{Adv}_{\text{DAKE}, n}^{\Gamma_q}(\mathcal{A}) = 0$ since CHAL behaves identically independent of challenge bit b . □

3.5 Deniable Split-KEM from Lattices

In this section we build an efficient deniable split-KEM under the hardness of LWE. We start by introducing briefly several concepts of lattice-based cryptography that we use to design the scheme. We then present our split-KEM and prove that it is correct, OW-CPA, deniable and decaps-CPA. After this, we provide a generic transformation in the (Q)ROM to build UNF-1KCA and IND-1BatchCCA split-KEM assuming decaps-CPA and OW-CPA security respectively. Finally, we provide concrete parameters for our split-KEM that provide at least 128 bits of classical and quantum security.

3.5.1 Lattice Toolbox

L_∞ and L_α norms. We start by recalling what the L_∞ and L_α norms over \mathbb{Z}_q are. For an element w in \mathbb{Z}_q , we write $\|w\|_\infty$ to mean $|\langle w \rangle_q|$ (i.e., $|w \pmod{q}|$). Then, we define the L_∞ and L_α norms for $\mathbf{w} = (w_1, w_2, \dots, w_n)$ over \mathbb{Z}_q as follows:

$$\|\mathbf{w}\|_\infty = \max_{j \in [n]} \|w_j\|_\infty, \quad \|\mathbf{w}\|_\alpha = \sqrt[\alpha]{\|w_1\|_\infty^\alpha + \dots + \|w_n\|_\infty^\alpha}.$$

By default, $\|\mathbf{w}\| := \|\mathbf{w}\|_2$.

Probability Distributions. For a finite set S , we define $\mathcal{U}(S)$ to be the uniform distribution on S . We will also use the binomial distribution Bin_1 which is defined as: $\text{Bin}_1(-1) = \text{Bin}_1(1) = 1/4$ and $\text{Bin}_1(0) = 1/2$.

Rounding Functions. Given two parameters q and $B < \log q - 1$, we define the rounding function $\lfloor \cdot \rfloor_{q,2}$ and the cross-rounding function $\langle \cdot \rangle_{q,B}$ as follows:

$$\lfloor v \rfloor_{q,2} := \left\lfloor \frac{2^B}{q} \cdot v \right\rfloor \bmod 2^B, \quad \langle v \rangle_{q,B} := \left\lfloor \frac{2^{B+1}}{q} \cdot v \right\rfloor \bmod 2,$$

for any $v \in \mathbb{Z}_q$.

Reconciliation Function. We recall the (generalised) reconciliation mechanism from Bos et al. and Peikert [BCD⁺16, Pei14], which for every approximate agreement in \mathbb{Z}_q allows extracting shared bits. We refer the reader to the aforementioned works for more details. Let q be a positive integer. Let B be the number of bits we want to extract from one coefficient in \mathbb{Z}_q so that $B < \log q - 1$. Now, for any $v \in \mathbb{Z}_q$, which is represented as an integer in $[0, q)$, we define the following functions.

Definition 28 (Randomised doubling function (dbl)). For any $v \in \mathbb{Z}_q$, we define $\text{dbl}(\cdot)$ as follows:

$$\text{dbl}(v) : v \mapsto 2v - e, \quad e \stackrel{\$}{\leftarrow} \text{Bin}_1$$

Then, we have the following property which comes from [BCD⁺16, Claim 3.1].

Lemma 1. Let q be odd. If $v \in \mathbb{Z}_q$ is uniformly random and $\bar{v} \stackrel{\$}{\leftarrow} \text{dbl}(v) \in \mathbb{Z}_{2q}$, then $\lfloor \bar{v} \rfloor_{2q,B}$ is uniformly random given $\langle \bar{v} \rangle_{2q,B}$.

Now, we are ready to define the reconciliation function $\text{rec} : \mathbb{Z}_{2q} \times \mathbb{Z}_2 \rightarrow \mathbb{Z}_{2^B}$.

Definition 29 (Reconciliation function (rec)). For any $w \in \mathbb{Z}_{2q}$ and bit $b \in \{0, 1\}$, let v be the closest element to $w \in \mathbb{Z}_{2q}$ s.t. $\langle v \rangle_{2q,B} = b$. Then, we define rec as

$$\text{rec}(w, b) := \lfloor v \rfloor_{2q,B}.$$

The next result gives an important property of the reconciliation function rec , as described by Peikert [Pei14, Section 3.2].

Lemma 2. Let q be odd and $\bar{v} \stackrel{\$}{\leftarrow} \text{dbl}(v)$. If $|v - w| \leq \lfloor \frac{q}{2^{B+2}} \rfloor$ then

$$\text{rec}(2w, \langle \bar{v} \rangle_{2q,B}) = \lfloor \bar{v} \rfloor_{2q,B}.$$

Finally, we define the $\text{HelpRec} : \mathbb{Z}_q \mapsto \{0, 1\}$ function as follows:

Definition 30 (HelpRec function). On any input $v \in \mathbb{Z}_q$,

$$\text{HelpRec}(v) := \langle \bar{v} \rangle_{2q,B}, \quad \text{where } \bar{v} \leftarrow \text{dbl}(v).$$

All the functions above can be naturally generalised to take as input vectors and matrices over \mathbb{Z}_q by applying the function to each of the coefficients.

Learning-with-Errors. The security of our lattice constructions relies on the learning-with-errors (LWE) problem introduced by Regev [Reg05]. In this chapter we will consider the case where both the secret and error coefficients come from a probability distribution over \mathbb{Z} .

Definition 31 ($\text{LWE}_{n,m,\chi,q}$). Let $n, m \in \mathbb{N}$ and χ be a probability distribution over \mathbb{Z} . The LWE problem asks the adversary \mathcal{A} to distinguish between the following two cases:

1. $(\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e} \bmod q)$ for $\mathbf{A} \stackrel{\$}{\leftarrow} \mathcal{U}(\mathbb{Z}_q^{n \times m})$, a secret $\mathbf{s} \stackrel{\$}{\leftarrow} \chi^m$ and error $\mathbf{e} \stackrel{\$}{\leftarrow} \chi^n$,
2. $(\mathbf{A}, \mathbf{t}) \stackrel{\$}{\leftarrow} \mathcal{U}(\mathbb{Z}_q^{n \times m}) \times \mathcal{U}(\mathbb{Z}_q^n)$.

We say that the $\text{LWE}_{n,m,\chi,q}$ assumption holds if any efficient adversary cannot distinguish between the two distributions except with negligible probability.

3.5.2 Extended-LWE

Our proof of deniability for the split-KEM will involve a new security assumption, which we call the Extended-LWE problem (ELWE). Intuitively, it is similar to the plain LWE problem, but the adversary is now also given random linear combinations of the secrets and errors.

Chapter 3. K-Waay: Fast and Deniable Post-Quantum X3DH Without Ring Signatures

Definition 32 (ELWE $_{n,m,\bar{n},\chi,q}$). Let $n, m \in \mathbb{N}$ and χ be a probability distribution over \mathbb{Z} . The ELWE problem asks the adversary \mathcal{A} to distinguish between the following two cases:

1. $(\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{e} \bmod q, \mathbf{Z}, \mathbf{W}, \mathbf{Z}\mathbf{s} + \mathbf{W}\mathbf{e} \bmod q)$ for $\mathbf{A} \xleftarrow{\$} \mathcal{U}(\mathbb{Z}_q^{n \times m})$, secret $\mathbf{s} \xleftarrow{\$} \chi^m$, error $\mathbf{e} \xleftarrow{\$} \chi^n$, and $(\mathbf{Z}, \mathbf{W}) \xleftarrow{\$} \chi^{\bar{n} \times m} \times \chi^{\bar{n} \times n}$,
2. $(\mathbf{A}, \mathbf{t}, \mathbf{Z}, \mathbf{W}, \mathbf{Z}\mathbf{s} + \mathbf{W}\mathbf{e} \bmod q)$ for $\mathbf{A} \xleftarrow{\$} \mathcal{U}(\mathbb{Z}_q^{n \times m})$, $\mathbf{t} \xleftarrow{\$} \mathcal{U}(\mathbb{Z}_q^n)$, secret $\mathbf{s} \xleftarrow{\$} \chi^m$, error $\mathbf{e} \xleftarrow{\$} \chi^n$, and $(\mathbf{Z}, \mathbf{W}) \xleftarrow{\$} \chi^{\bar{n} \times m} \times \chi^{\bar{n} \times n}$.

We say that the ELWE $_{n,m,\bar{n},\chi,q}$ assumption holds if any efficient adversary cannot distinguish between the two distributions except with negligible probability.

This problem is a natural generalisation of the Extended-LWE problem by Alperin-Sheriff and Peikert [AP12], where now (\mathbf{Z}, \mathbf{W}) are matrices and not just vectors. Here, we also simplify the definition and assume that the coefficients of \mathbf{Z} and \mathbf{W} come from the same distribution χ as the secrets and errors.

We show in the following theorem that the hardness of this newly introduced ELWE problem reduces to the hardness of LWE.

Theorem 3. Let q be an odd prime and χ be symmetric around 0. If there is an efficient adversary \mathcal{A} which wins ELWE $_{n,m,\bar{n},\chi,q}$ with probability ε , then there also exists an efficient adversary \mathcal{B} which wins LWE $_{n+m,m,\chi,q}$ with probability at least $\delta_{\text{elwe}} \cdot \varepsilon - \text{negl}(n)$ where

$$\delta_{\text{elwe}} := \Pr[\mathbf{Z}(\mathbf{e} - \mathbf{d}) = \mathbf{0} \pmod{q} : \mathbf{Z} \xleftarrow{\$} \chi^{\bar{n} \times (n+m)}, \mathbf{e}, \mathbf{d} \xleftarrow{\$} \chi^{n+m}]. \quad (3.1)$$

Proof. We prove the statement by introducing a sequence of LWE-type games Γ_i . We start with $\Gamma^1 := \text{ELWE}_{n,m,\bar{n},\chi,q}$ and give an efficient reduction from Γ_i to Γ_{i+1} . In the end, we finish with LWE $_{n+m,m,\chi,q}$. We denote $\text{Adv}_i(\mathcal{A})$ to be the probability that \mathcal{A} wins Γ_i . Then, the proof follows by the composition of the constant number of efficient reductions.

Game Γ_1 : This is the standard ELWE $_{n,m,\bar{n},\chi,q}$ game. The adversary \mathcal{A} wins this game with probability ε .

Game Γ_2 : Here, we consider the ELWE-type game where the secret vector is uniformly random. Namely, the challenger samples the public $\mathbf{A} \xleftarrow{\$} \mathcal{U}(\mathbb{Z}_q^{(n+m) \times m})$, secret $\mathbf{s} \xleftarrow{\$} \mathcal{U}(\mathbb{Z}_q^m)$, error $\mathbf{e} \xleftarrow{\$} \mathcal{U}(\mathbb{Z}_q^{n+m})$ as well as the hint matrix $\mathbf{Z} \xleftarrow{\$} \chi^{\bar{n} \times (n+m)}$. Then it flips a bit $b \xleftarrow{\$} \mathcal{U}(\{0, 1\})$. If $b = 0$ then the challenger computes

$$\mathbf{t} := \mathbf{A}\mathbf{s} + \mathbf{e}$$

and otherwise it samples $\mathbf{t} \xleftarrow{\$} \mathcal{U}(\mathbb{Z}_q^{n+m})$. The challenger outputs $(\mathbf{A}, \mathbf{t}, \mathbf{Z}, \mathbf{Z}\mathbf{e})$.

Lemma 3. For every efficient adversary \mathcal{A} , there is an efficient adversary \mathcal{B} such that $\text{Adv}_2(\mathcal{B}) \geq \text{Adv}_1(\mathcal{A}) - \text{negl}(n)$.

Proof. The reduction follows similarly as in the one by Applebaum et al. [ACPS09]. Suppose the algorithm \mathcal{B} is given a tuple $(\mathbf{A}, \mathbf{t}, \mathbf{Z}, \mathbf{h})$ from Γ_2 . With probability at most $1/q^{(n+m)-m-1} \leq 1/q^{n-1}$, matrix \mathbf{A} is not full-rank. Let us exclude that case and assume without loss of generality that we can write

$$\mathbf{A} := \begin{bmatrix} \mathbf{A}_0 \\ \mathbf{A}_1 \end{bmatrix}, \quad \mathbf{Z} := \begin{bmatrix} \mathbf{Z}_0 & \mathbf{Z}_1 \end{bmatrix}, \quad \text{and} \quad \mathbf{t} := \begin{bmatrix} \mathbf{t}_0 \\ \mathbf{t}_1 \end{bmatrix}$$

where $\mathbf{A}_1 \in \mathbb{Z}_q^{n \times m}$ and the matrix $\mathbf{A}_0 \in \mathbb{Z}_q^{m \times m}$, which contains the first m rows of \mathbf{A} , is invertible. Thus, define $\mathbf{A}' := \mathbf{A}_1 \mathbf{A}_0^{-1} \in \mathbb{Z}_q^{n \times m}$, and $\mathbf{t}' := \mathbf{A}' \mathbf{t}_0 - \mathbf{t}_1 \in \mathbb{Z}_q^n$. Then, it runs \mathcal{A} on input

$$(\mathbf{A}', \mathbf{t}', \mathbf{Z}_0, -\mathbf{Z}_1, \mathbf{h})$$

and returns what \mathcal{A} outputs.

Suppose that $\mathbf{t} = \mathbf{A}\mathbf{s} + \mathbf{e}$ where $\mathbf{s} \in \mathbb{Z}_q^m$ and $\mathbf{e} := (\mathbf{e}_0, \mathbf{e}_1) \in \mathbb{Z}_q^m \times \mathbb{Z}_q^n$. Then

$$\mathbf{t}' = \mathbf{A}' \mathbf{t}_0 - \mathbf{t}_1 = \mathbf{A}_1 \mathbf{A}_0^{-1} (\mathbf{A}_0 \mathbf{s} + \mathbf{e}_0) - (\mathbf{A}_1 \mathbf{s} + \mathbf{e}_1) = \mathbf{A}' \mathbf{e}_0 - \mathbf{e}_1$$

which is a valid LWE instance since χ is symmetric around 0. Also, if \mathbf{A} is uniformly random among all nonsingular matrices, then \mathbf{A}' and \mathbf{B}' are statistically close to uniformly random matrices over \mathbb{Z}_q . As for the hints, note that

$$\mathbf{h} = \mathbf{Z}_0 \mathbf{e}_0 + \mathbf{Z}_1 \mathbf{e}_1 = \mathbf{Z}_0 \mathbf{e}_0 + (-\mathbf{Z}_1)(-\mathbf{e}_1),$$

so \mathbf{h} is a well-formed hint for Γ_1 .

On the other hand, if \mathbf{t} is uniformly random, then so is \mathbf{t}' . It can be argued similarly as before that all the other components follow the distribution for $b = 1$. \square

Game Γ_3 : We consider the knapsack version of ELWE. Here, the challenger samples the public $\mathbf{G} := \overset{\$}{\mathcal{U}}(\mathbb{Z}_q^{n \times (n+m)})$, secret $\mathbf{e} \overset{\$}{\mathcal{U}}(\mathbb{Z}_q^{n+m})$ and the hint matrix $\mathbf{Z} \overset{\$}{\mathcal{U}}(\chi^{\tilde{n} \times (n+m)})$. Then it flips a bit $b \overset{\$}{\mathcal{U}}(\{0, 1\})$. If $b = 0$ then the challenger computes $\mathbf{t} := \mathbf{G}\mathbf{e}$, and otherwise it samples $\mathbf{t} \overset{\$}{\mathcal{U}}(\mathbb{Z}_q^n)$. Finally, the challenger outputs $(\mathbf{G}, \mathbf{t}, \mathbf{Z}, \mathbf{e})$.

Lemma 4. For every efficient adversary \mathcal{A} , there is an efficient adversary \mathcal{B} such that $\text{Adv}_3(\mathcal{B}) \geq \text{Adv}_2(\mathcal{A}) - \text{negl}(n)$.

Proof. The reduction is similar to the proof of Micciancio and Mol [MM11, Lemma 4.9]. Suppose the algorithm \mathcal{B} is given a tuple $(\mathbf{G}, \mathbf{t}, \mathbf{Z}, \mathbf{h})$ from Γ_3 . Then, \mathcal{B} can construct a randomised matrix $\mathbf{A} \in \mathbb{Z}_q^{(n+m) \times m}$ whose columns generate the kernel of \mathbf{G} . In particular, if \mathbf{G} is uniformly random, then so are (\mathbf{A}, \mathbf{B}) , up to the constraint that they are nonsingular. Then, \mathcal{B} computes any solution \mathbf{r} such that $\mathbf{G}\mathbf{r} = \mathbf{t}$. Finally, it samples a uniformly random $\mathbf{s} \overset{\$}{\mathcal{U}}(\mathbb{Z}_q^m)$ and runs \mathcal{A} on input

$$(\mathbf{A}, \mathbf{A}\mathbf{s} + \mathbf{r}, \mathbf{Z}, \mathbf{h})$$

Chapter 3. K-Waay: Fast and Deniable Post-Quantum X3DH Without Ring Signatures

and returns what \mathcal{A} outputs.

Suppose that $\mathbf{G}\mathbf{e} = \mathbf{t} = \mathbf{G}\mathbf{r}$. By definition of the matrix \mathbf{A} , $\mathbf{G}(\mathbf{r} - \mathbf{e}) = \mathbf{0}$ implies that there exists some vector $\mathbf{x} \in \mathbb{Z}_q^m$ such that $\mathbf{r} - \mathbf{e} = \mathbf{A}\mathbf{x}$. Thus,

$$\mathbf{A}\mathbf{s} + \mathbf{r} = \mathbf{A}(\mathbf{s} + \mathbf{x}) + \mathbf{e}$$

which is a valid LWE instance since $\mathbf{s} + \mathbf{x}$ is still uniformly random over \mathbb{Z}_q^m . As for the hints, we still have $\mathbf{h} = \mathbf{Z}\mathbf{e}$ and thus \mathcal{B} correctly simulates Γ_2 for $b = 0$. The case $b = 1$ follows by arguing that \mathbf{t} is uniformly random and if \mathbf{G} is nonsingular then \mathbf{r} must be uniformly random. \square

Game Γ_4 : This game is a plain knapsack LWE problem. The challenger samples the public $\mathbf{G} \xleftarrow{\$} \mathcal{U}(\mathbb{Z}_q^{n \times (n+m)})$ and a secret $\mathbf{e} \xleftarrow{\$} \mathcal{U}(\mathbb{Z}_q^{n+m})$. Then it flips a bit $b \xleftarrow{\$} \mathcal{U}(\{0, 1\})$. If $b = 0$ then the challenger computes $\mathbf{t} := \mathbf{G}\mathbf{e}$, and otherwise it samples $\mathbf{t} \xleftarrow{\$} \mathcal{U}(\mathbb{Z}_q^n)$. Finally, the challenger outputs (\mathbf{G}, \mathbf{t}) .

Lemma 5. For every efficient adversary \mathcal{A} , there is an efficient adversary \mathcal{B} such that $\text{Adv}_4(\mathcal{B}) \geq \delta_{\text{elwe}} \cdot \text{Adv}_3(\mathcal{A})$.

Proof. We follow the proof strategy from [AP12, Theorem 1]. Suppose the algorithm \mathcal{B} is given a tuple $(\mathbf{G}_0, \mathbf{G}_1, \mathbf{t})$ from Γ_4 . Then, it samples $\mathbf{Z} \xleftarrow{\$} \chi^{\tilde{n} \times (n+m)}$, $\mathbf{d} \xleftarrow{\$} \chi^{n+m}$ and a matrix $\mathbf{V} \xleftarrow{\$} \mathcal{U}(\mathbb{Z}_q^{n \times \tilde{n}})$. Further, it sets

$$\mathbf{G}' := \mathbf{G} - \mathbf{V}\mathbf{Z} \quad \text{and} \quad \mathbf{t}' := \mathbf{t} - \mathbf{V}\mathbf{Z}\mathbf{d}.$$

Finally, it runs \mathcal{A} on input

$$(\mathbf{G}', \mathbf{t}', \mathbf{Z}, \mathbf{Z}\mathbf{d})$$

and returns what \mathcal{A} outputs.

Clearly, if \mathbf{G} (resp. \mathbf{t}) is uniformly random then so is \mathbf{G}' (resp. \mathbf{t}'). Hence, the case $b = 1$ follows directly. Suppose $b = 0$ and thus $\mathbf{t} = \mathbf{G}\mathbf{e}$. Then, we have

$$\mathbf{t}' = \mathbf{t} - \mathbf{V}\mathbf{Z}\mathbf{d} = \mathbf{G}\mathbf{e} - \mathbf{V}\mathbf{Z}\mathbf{d} = \mathbf{G}'\mathbf{e} + \mathbf{V}(\mathbf{Z}\mathbf{e} - \mathbf{Z}\mathbf{d}).$$

Hence, if $\mathbf{Z}\mathbf{e} = \mathbf{Z}\mathbf{d}$ then $([\mathbf{G}'_0 \ \mathbf{G}'_1], \mathbf{t}', \mathbf{Z}, \mathbf{Z}\mathbf{d})$ is indeed a valid knapsack ELWE tuple. This happens exactly with probability at most δ_{elwe} by definition. Otherwise, $\mathbf{V}(\mathbf{Z}\mathbf{e} - \mathbf{Z}\mathbf{d})$ is a uniformly random vector over \mathbb{Z}_q , and so is \mathbf{t}' . Thus, the tuple output by \mathcal{B} follows the case $b = 1$ for Γ_3 . The statement now follows by simple calculation. \square

Game Γ_5 : Here, we consider the plain LWE game. Recall that the challenger samples the public $\mathbf{A} \xleftarrow{\$} \mathcal{U}(\mathbb{Z}_q^{(n+m) \times m})$, secret $\mathbf{s} \xleftarrow{\$} \mathcal{U}(\mathbb{Z}_q^m)$, error $\mathbf{e} \xleftarrow{\$} \mathcal{U}(\mathbb{Z}_q^{n+m})$. Then it flips a bit $b \xleftarrow{\$} \mathcal{U}(\{0, 1\})$. If $b = 0$ then the challenger computes $\mathbf{t} := \mathbf{A}\mathbf{s} + \mathbf{e}$, and otherwise it samples $\mathbf{t} \xleftarrow{\$} \mathcal{U}(\mathbb{Z}_q^{n+m})$. At the end, the challenger outputs (\mathbf{A}, \mathbf{t}) .

Lemma 6. For every efficient adversary \mathcal{A} , there is an efficient adversary \mathcal{B} such that $\text{Adv}_5(\mathcal{B}) \geq \text{Adv}_4(\mathcal{A}) - \text{negl}(n)$.

Proof. The reduction is identical to the one of Micciancio and Mol [MM11, Lemma 4.8] which we recall for completeness. Suppose the algorithm \mathcal{B} is given a tuple (\mathbf{A}, \mathbf{t}) from Γ_5 . If \mathbf{A} is full-rank, then \mathcal{B} can construct a (randomised) matrix $\mathbf{G} \in \mathbb{Z}_q^{n \times (n+m)}$ whose rows generate all the vectors \mathbf{x} such that $\mathbf{x}^T \mathbf{A} = \mathbf{0}$. Also, if \mathbf{A} is chosen at random among all full-rank matrices, then \mathbf{G} is also distributed statistically close to a uniformly random. Then, \mathcal{B} outputs $(\mathbf{G}, \mathbf{Gt})$ to \mathcal{A} and returns what \mathcal{A} outputs.

Suppose $b = 0$ and $\mathbf{t} = \mathbf{As} + \mathbf{e}$. Then $\mathbf{Gt} = \mathbf{GAs} + \mathbf{e} = \mathbf{Ge}$, which is the correct instance of Γ_4 for $b = 0$. On the other hand, if \mathbf{t} is uniformly random, then so is \mathbf{Gt} . \square

The statement of the theorem now follows by combining all the previous lemmas using reduction composition. \square

3.5.3 Construction

We can now present our Frodo-inspired [BCD⁺16] split-KEM, which we call FrodoKEX+. The scheme is given in Figure 3.9. The key generation works as follows. The public key pk_A for party A is a pair $(\mathbf{A}, \mathbf{B}_A)$, where \mathbf{A} is a uniformly random matrix over \mathbb{Z}_q given as a public parameter, and $\mathbf{B}_A := \mathbf{AS}_A + \mathbf{D}_A$ where $\mathbf{S}_A, \mathbf{D}_A \stackrel{\$}{\leftarrow} \chi^{n \times \tilde{n}}$. The secret key becomes a pair $\text{sk}_A = (\mathbf{S}_A, \mathbf{D}_A)$. Similarly, the public key pk_B for party B is a pair $(\mathbf{A}, \mathbf{B}_B)$, where $\mathbf{B}_B := \mathbf{S}_B \mathbf{A} + \mathbf{D}_B$, while the secret key is $\text{sk}_B = (\mathbf{S}_B, \mathbf{D}_B)$, where $\mathbf{S}_B, \mathbf{D}_B \stackrel{\$}{\leftarrow} \chi^{\tilde{n} \times n}$.

Then, B samples a matrix $\mathbf{E}_B \stackrel{\$}{\leftarrow} \chi^{\tilde{n} \times \tilde{n}}$ and computes the matrix $\mathbf{V} := \mathbf{S}_B \mathbf{B}_A + \mathbf{E}_B$. Next, it computes $\text{ct} \leftarrow \text{HelpRec}(\mathbf{V})$ and $\mathbf{K} \leftarrow \text{rec}(\mathbf{V}, \text{ct})$. Then, B outputs ct . Then, party A decapsulates as follows: given $(\text{pk}_B, \text{sk}_A, \text{ct})$, it computes $\mathbf{V}' = \mathbf{B}_B \mathbf{S}_A + \mathbf{F}_A$ and $\mathbf{K}' = \text{rec}(\mathbf{V}', \text{ct})$. Finally, A returns the key \mathbf{K}' .

We note that the construction can easily be made symmetric, in the sense that A could encapsulate using B's public key by changing the order of matrices when multiplying in Encaps such that the dimensions match. Then, Decaps can be modified similarly such that B can decapsulate the resulting ciphertext.

3.5.4 Security Analysis

Lemma 7 (Correctness). Let χ be a symmetric distribution around 0 and δ_{corr} be the following probability:

$$\Pr \left[|\langle \mathbf{s}, \mathbf{d} \rangle + e + f| > \frac{q}{2^{B+2}} : \mathbf{s}, \mathbf{d} \stackrel{\$}{\leftarrow} \chi^{2n}, e, f \stackrel{\$}{\leftarrow} \chi \right]. \quad (3.2)$$

Then, sKEM defined in Figure 3.9 is $(\tilde{n}^2 \delta_{\text{corr}})$ -correct.

Chapter 3. K-Waay: Fast and Deniable Post-Quantum X3DH Without Ring Signatures

KeyGenA(1^λ)	Encaps($\text{pk}_A = (\mathbf{A}, \mathbf{B}_A), \text{sk}_B = (\mathbf{S}_B, \mathbf{D}_B, \mathbf{F}_B)$)
1: $\mathbf{S}_A, \mathbf{D}_A \xleftarrow{\$} \chi(\mathbb{Z}_q^{n \times \bar{n}})$	1: // We assume B encapsulates
2: $\mathbf{F}_A \xleftarrow{\$} \chi^{\bar{n} \times \bar{n}}$	2: $\mathbf{E}_B \xleftarrow{\$} \chi^{\bar{n} \times \bar{n}}$
3: $\mathbf{B}_A \leftarrow \mathbf{A}\mathbf{S}_A + \mathbf{D}_A$	3: $\mathbf{V} \leftarrow \mathbf{S}_B\mathbf{B}_A + \mathbf{E}_B$
4: $\text{pk}_A \leftarrow (\mathbf{A}, \mathbf{B}_A)$	4: $\text{ct} \leftarrow \text{HelpRec}(\mathbf{V})$
5: $\text{sk}_A \leftarrow (\mathbf{S}_A, \mathbf{D}_A, \mathbf{F}_A)$	5: $\mathbf{K} \leftarrow \text{rec}(2\mathbf{V}, \text{ct})$
6: return (pk_A, sk_A)	6: return (\mathbf{K}, ct)
KeyGenB(1^λ)	Decaps($\text{pk}_B = (\mathbf{A}, \mathbf{B}_B), \text{sk}_A = (\mathbf{S}_A, \mathbf{D}_A, \mathbf{F}_A), \text{ct}$)
1: $\mathbf{S}_B, \mathbf{D}_B \xleftarrow{\$} \chi^{\bar{n} \times n}$	1: $\mathbf{V}' \leftarrow \mathbf{B}_B\mathbf{S}_A + \mathbf{F}_A$
2: $\mathbf{F}_B \xleftarrow{\$} \chi^{\bar{n} \times \bar{n}}$	2: $\mathbf{K}' \leftarrow \text{rec}(2\mathbf{V}', \text{ct})$
3: $\mathbf{B}_B \leftarrow \mathbf{S}_B\mathbf{A} + \mathbf{D}_B$	3: return \mathbf{K}'
4: $\text{pk}_B \leftarrow (\mathbf{A}, \mathbf{B}_B)$	
5: $\text{sk}_B \leftarrow (\mathbf{S}_B, \mathbf{D}_B, \mathbf{F}_B)$	
6: return (pk_B, sk_B)	

Figure 3.9: Our variant of FrodoKEX [BCD⁺16] expressed as a split-KEM. The matrix $\mathbf{A} \in \mathbb{Z}_q^{n \times n}$ is assumed to be a public parameter and sampled uniformly at random.

Proof. Suppose $(\text{pk}_A, \text{sk}_A) \xleftarrow{\$} \text{KeyGenA}(1^\lambda)$ and $(\text{pk}_B, \text{sk}_B) \xleftarrow{\$} \text{KeyGenB}(1^\lambda)$, and let

$$(\mathbf{K}, \text{ct}) \xleftarrow{\$} \text{Encaps}(\text{pk}_A, \text{sk}_B) \quad \text{and} \quad \mathbf{K}' \xleftarrow{\$} \text{Decaps}(\text{pk}_B, \text{sk}_A, \text{ct}).$$

We want to prove that $\mathbf{K} = \mathbf{K}'$. By definition of encapsulation, we know that $\mathbf{K} = \text{rec}(2\mathbf{V}, \text{ct})$ where $\text{ct} = \text{HelpRec}(\mathbf{V})$ and

$$\mathbf{V} = \mathbf{S}_B\mathbf{B}_A + \mathbf{E}_B = \mathbf{S}_B\mathbf{A}\mathbf{S}_A + \mathbf{S}_B\mathbf{D}_A + \mathbf{E}_B.$$

Thus, by Lemma 2, $\mathbf{K} = \lfloor \mathbf{V} \rfloor_{2q, 2^B}$. On the other hand,

$$\mathbf{V}' = \mathbf{B}_B\mathbf{S}_A + \mathbf{F}_A = \mathbf{S}_B\mathbf{A}\mathbf{S}_A + \mathbf{D}_B\mathbf{S}_A + \mathbf{F}_A$$

which implies that $\mathbf{V} - \mathbf{V}' = \mathbf{S}_B\mathbf{D}_A + \mathbf{E}_B - \mathbf{D}_B\mathbf{S}_A - \mathbf{F}_A$. If $\|\mathbf{V} - \mathbf{V}'\|_\infty < \frac{q}{2^{B+2}}$ then by Lemma 2 we must have

$$\mathbf{K}' = \text{rec}(2\mathbf{V}', \text{HelpRec}(\mathbf{V})) = \lfloor \mathbf{V} \rfloor_{2q, 2^B} = \mathbf{K}$$

so correctness holds. Now, using the fact that χ is symmetric around 0, the probability $\|\mathbf{V} - \mathbf{V}'\|_\infty > \frac{q}{2^{B+2}}$ can be upper-bounded using the union bound as follows:

$$\Pr \left[\|\mathbf{S}_B\mathbf{D}_A + \mathbf{E}_B - \mathbf{D}_B\mathbf{S}_A - \mathbf{F}_A\|_\infty > \frac{q}{2^{B+2}} \right] \leq \bar{n}^2 \cdot \Pr \left[|\mathbf{s}_0^T \mathbf{d}_0 + \mathbf{s}_1^T \mathbf{d}_1 + e + f| > \frac{q}{2^{B+2}} \right]$$

where $\mathbf{s}_0, \mathbf{s}_1, \mathbf{d}_0, \mathbf{d}_1 \xleftarrow{\$} \chi^n$ and $e, f \xleftarrow{\$} \chi$. This concludes the proof. \square

OW-CPA Security. Next, we focus on proving OW-CPA security.

Lemma 8 (OW-CPA Security). Let $\bar{n} = O(\lambda)$ and χ be a symmetric distribution over $[-\gamma, \gamma]$ for any $\gamma > 0$. Then, under the $\text{LWE}_{n,n,\chi,q}$ and $\text{LWE}_{n+\bar{n},n,\chi,q}$ assumptions, for every efficient adversary \mathcal{A} , the probability of \mathcal{A} winning the OW-CPA game is at most $2^{-B\bar{n}^2} + \text{negl}(\lambda)$.

Proof. Let \mathcal{A} be an efficient adversary against the OW-CPA game. We prove the statement using the hybrid games described explicitly in Figure 3.10. In each game Γ_i , we define ε_i to be the probability that the efficient adversary \mathcal{A} wins the security game.

Game Γ_1 : This is the standard OW-CPA game.

Game Γ_2 : Instead of computing $\mathbf{B}_A \leftarrow \mathbf{A}\mathbf{S}_A + \mathbf{D}_A$, the experiment samples $\mathbf{B}_A \leftarrow \mathcal{U}(\mathbb{Z}_q^{n \times \bar{n}})$. One can naturally build an efficient adversary, which can solve the $\text{LWE}_{n,n,\chi,q}$ problem with probability at least $\frac{1}{\bar{n}}|\varepsilon_2 - \varepsilon_1|$. Hence, we deduce that this probability is negligible.

Game Γ_3 : Here, the experiment computes the values \mathbf{B}_B and \mathbf{V} differently. Namely, instead of computing:

$$\begin{bmatrix} \mathbf{B}_B & \mathbf{V} \end{bmatrix} := \mathbf{S}_B \begin{bmatrix} \mathbf{A} & \mathbf{B}_A \end{bmatrix} + \begin{bmatrix} \mathbf{D}_B & \mathbf{E}_B \end{bmatrix},$$

it samples

$$\begin{bmatrix} \mathbf{B}_B & \mathbf{V} \end{bmatrix} \stackrel{\$}{\leftarrow} \mathcal{U}(\mathbb{Z}_q^{\bar{n} \times (n+\bar{n})}).$$

Thus, one can naturally construct an efficient reduction which solves $\text{LWE}_{n+\bar{n},n,\chi,q}$ with probability at least $\frac{1}{\bar{n}}|\varepsilon_6 - \varepsilon_5|$.

Finally, it is easy to see that in Γ_3 the matrix \mathbf{V} is actually uniformly random over \mathbb{Z}_q . Hence by Lemma 1, for the adversary \mathcal{A} , which is given ct, the key \mathbf{K} looks uniformly random. Therefore, the probability of guessing the key is bounded by $2^{-\bar{n}^2 B}$. \square

Deniability. We will use the (transposed) matrix version of ELWE where the secrets and errors are now matrices. In particular, we will be interested in the problem of distinguishing between

$$(\mathbf{A}, \mathbf{S}\mathbf{A} + \mathbf{E} \bmod q, \mathbf{Z}, \mathbf{W}, \mathbf{S}\mathbf{Z} + \mathbf{E}\mathbf{W} \bmod q)$$

and

$$(\mathbf{A}, \mathbf{T}, \mathbf{Z}, \mathbf{W}, \mathbf{S}\mathbf{Z} + \mathbf{E}\mathbf{W} \bmod q)$$

where $\mathbf{S} \stackrel{\$}{\leftarrow} \chi^{\bar{n} \times m}$, $\mathbf{E} \stackrel{\$}{\leftarrow} \chi^{\bar{n} \times n}$ and $\mathbf{T} \stackrel{\$}{\leftarrow} \mathcal{U}(\mathbb{Z}_q^{\bar{n} \times n})$. This problem can be reduced to ELWE with reduction loss \bar{n} via a standard hybrid argument.

We are ready to prove deniability of the split-KEM based on Extended-LWE. Intuitively, matrices $(\mathbf{S}, \mathbf{E}) := (\mathbf{S}_B, \mathbf{D}_B)$ will be the secret and error constructed by party B, which are hidden from the adversary, while $(\mathbf{Z}, \mathbf{W}) := (\mathbf{D}_A, \mathbf{S}_A)$ will be the error and the secret generated by A which are

Chapter 3. K-Waay: Fast and Deniable Post-Quantum X3DH Without Ring Signatures

$\Gamma_1(\mathcal{A})$	$\Gamma_2(\mathcal{A})$	$\Gamma_3(\mathcal{A})$
1: $\mathbf{S}_A, \mathbf{D}_A \xleftarrow{\$} \chi(\mathbb{Z}_q^{n \times \bar{n}})$	1: $\mathbf{B}_A \leftarrow \mathcal{U}(\mathbb{Z}_q^{n \times \bar{n}})$	1: $\mathbf{B}_A \leftarrow \mathcal{U}(\mathbb{Z}_q^{n \times \bar{n}})$
2: $\mathbf{F}_A \xleftarrow{\$} \chi^{\bar{n} \times \bar{n}}$	2: $\mathbf{S}_B, \mathbf{D}_B \xleftarrow{\$} \chi^{\bar{n} \times n}$	2: $\mathbf{B}_B \xleftarrow{\$} \mathcal{U}(\mathbb{Z}_q^{n \times \bar{n}})$
3: $\mathbf{B}_A \leftarrow \mathbf{A}\mathbf{S}_A + \mathbf{D}_A$	3: $\mathbf{F}_B \xleftarrow{\$} \chi^{\bar{n} \times \bar{n}}$	3: $\mathbf{V} \xleftarrow{\$} \mathcal{U}(\mathbb{Z}_q^{\bar{n} \times \bar{n}})$
4: $\mathbf{S}_B, \mathbf{D}_B \xleftarrow{\$} \chi^{\bar{n} \times n}$	4: $\mathbf{B}_B \leftarrow \mathbf{S}_B \mathbf{A} + \mathbf{D}_B$	4: $\text{ct} \leftarrow \text{HelpRec}(\mathbf{V})$
5: $\mathbf{F}_B \xleftarrow{\$} \chi^{\bar{n} \times \bar{n}}$	5: $\mathbf{E}_B \xleftarrow{\$} \chi^{\bar{n} \times \bar{n}}$	5: $\mathbf{K} \leftarrow \text{rec}(2\mathbf{V}, \text{ct})$
6: $\mathbf{B}_B \leftarrow \mathbf{S}_B \mathbf{A} + \mathbf{D}_B$	6: $\mathbf{V} \leftarrow \mathbf{S}_B \mathbf{B}_A + \mathbf{E}_B$	6: $\mathbf{K}' \xleftarrow{\$} \mathcal{A}(\mathbf{A}, \mathbf{B}_A, \mathbf{B}_B, \text{ct})$
7: $\mathbf{E}_B \xleftarrow{\$} \chi^{\bar{n} \times \bar{n}}$	7: $\text{ct} \leftarrow \text{HelpRec}(\mathbf{V})$	7: return $\mathbb{1}[\mathbf{K} = \mathbf{K}']$
8: $\mathbf{V} \leftarrow \mathbf{S}_B \mathbf{B}_A + \mathbf{E}_B$	8: $\mathbf{K} \leftarrow \text{rec}(2\mathbf{V}, \text{ct})$	
9: $\text{ct} \leftarrow \text{HelpRec}(\mathbf{V})$	9: $\mathbf{K}' \xleftarrow{\$} \mathcal{A}(\mathbf{A}, \mathbf{B}_A, \mathbf{B}_B, \text{ct})$	
10: $\mathbf{K} \leftarrow \text{rec}(2\mathbf{V}, \text{ct})$	10: return $\mathbb{1}[\mathbf{K} = \mathbf{K}']$	
11: $\mathbf{K}' \xleftarrow{\$} \mathcal{A}(\mathbf{A}, \mathbf{B}_A, \mathbf{B}_B, \text{ct})$		
12: return $\mathbb{1}[\mathbf{K} = \mathbf{K}']$		

Figure 3.10: Security games for the proof of Lemma 8. The lines in blue highlight the main differences from the previous game.

given as input to the simulator. The key observation is that the additional hint provided as $\mathbf{S}\mathbf{Z} + \mathbf{E}\mathbf{W} \bmod q$ will be used to simulate the “shared key” \mathbf{V} (before applying the reconciliation function).

Theorem 4 (Deniability). Let $\bar{n} = \text{poly}(\lambda)$. Then, the sKEM defined in Figure 3.9 is deniable under the $\text{ELWE}_{n,n,\bar{n},\chi,q}$ and $\text{LWE}_{n,n,\chi,q}$ assumptions.

Proof. Let \mathcal{A} be an efficient adversary against the deniability game. We prove the statement using the hybrid games defined in Figure 3.12. In each game Γ_i , we define ε_i to be the probability that the efficient adversary \mathcal{A} outputs $b = 1$.

Game Γ_1 : This is the standard (real) deniability experiment, which we recall here. First, both $\mathbf{S}_A, \mathbf{D}_A \xleftarrow{\$} \chi^{n \times \bar{n}}$ and $\mathbf{S}_B, \mathbf{D}_B \xleftarrow{\$} \chi^{\bar{n} \times n}$ and $\mathbf{F}_A, \mathbf{F}_B \xleftarrow{\$} \chi^{\bar{n} \times \bar{n}}$ are sampled. Then, the public keys $\mathbf{B}_A = \mathbf{A}\mathbf{S}_A + \mathbf{D}_A$ and $\mathbf{B}_B = \mathbf{S}_B \mathbf{A} + \mathbf{D}_B$ are computed. The encapsulation algorithm samples $\mathbf{E}_B \xleftarrow{\$} \chi^{\bar{n} \times \bar{n}}$ and sets $\mathbf{V} \leftarrow \mathbf{S}_B \mathbf{B}_A + \mathbf{E}_B$. Finally, the experiment runs $\text{ct} \leftarrow \text{HelpRec}(\mathbf{V})$ and $\mathbf{K} \leftarrow \text{rec}(2\mathbf{V}, \text{ct})$ and eventually outputs

$$(\mathbf{A}, \mathbf{B}_A, \mathbf{B}_B, \mathbf{S}_A, \mathbf{D}_A, \mathbf{F}_A, \mathbf{K}, \text{ct})$$

to the adversary \mathcal{A} .

Game Γ_2 : The experiment is identical to the previous one, apart from the fact that

now \mathbf{V} is explicitly computed as $\mathbf{V} = \mathbf{S}_B \mathbf{D}_A - \mathbf{D}_B \mathbf{S}_A + \mathbf{B}_B \mathbf{S}_A + \mathbf{E}_B$. Clearly, $\varepsilon_1 = \varepsilon_2$ since

$$\begin{aligned} \mathbf{V} &= \mathbf{S}_B \mathbf{D}_A - \mathbf{D}_B \mathbf{S}_A + \mathbf{B}_B \mathbf{S}_A + \mathbf{E}_B \\ &= \mathbf{S}_B \mathbf{D}_A - \mathbf{D}_B \mathbf{S}_A + (\mathbf{S}_B \mathbf{A} + \mathbf{D}_B) \mathbf{S}_A + \mathbf{E}_B \\ &= \mathbf{S}_B \mathbf{B}_A + \mathbf{E}_B. \end{aligned}$$

Game Γ_3 : Here, the experiment follows Γ_2 with the only difference being that the experiment samples \mathbf{B}_B uniformly at random from $\mathbb{Z}_q^{\bar{n} \times n}$ instead of computing $\mathbf{B}_B = \mathbf{S}_B \mathbf{A} + \mathbf{D}_B$.

Lemma 9. There exists an efficient algorithm \mathcal{B} that solves the $\text{ELWE}_{n,n,\bar{n},\chi,q}$ problem with probability at least $\frac{1}{\bar{n}} |\varepsilon_3 - \varepsilon_2|$.

Proof. We provide a reduction \mathcal{B} to the (transposed) matrix version of the Extended-LWE problem as described above. Namely, the reduction is given a tuple of matrices $(\mathbf{A}, \mathbf{B}, \mathbf{Z}, \mathbf{W}, \mathbf{H})$. Then, it sets $\mathbf{S}_A := -\mathbf{W}$, $\mathbf{D}_A := \mathbf{Z}$ and $\mathbf{B}_B := \mathbf{B}$. Further, the reduction samples $\mathbf{F}_A \stackrel{\$}{\leftarrow} \chi^{\bar{n} \times \bar{n}}$ and computes

$$\mathbf{B}_A := \mathbf{A} \mathbf{S}_A + \mathbf{D}_A \quad \text{and} \quad \mathbf{V} := \mathbf{H} + \mathbf{B}_B \mathbf{S}_A + \mathbf{E}_B$$

where $\mathbf{E}_B \stackrel{\$}{\leftarrow} \chi^{\bar{n} \times \bar{n}}$. Finally, the reduction runs $\text{ct} \leftarrow \text{HelpRec}(\mathbf{V})$ and $\mathbf{K} \leftarrow \text{rec}(2\mathbf{V}, \text{ct})$ and outputs $(\mathbf{A}, \mathbf{B}_A, \mathbf{B}_B, \mathbf{S}_A, \mathbf{D}_A, \mathbf{F}_A, \mathbf{K}, \text{ct})$ to the adversary.

Suppose the input tuple received by \mathcal{B} is a true Extended-LWE instance, i.e. $\mathbf{B}_B = \mathbf{B} = \mathbf{S}_B \mathbf{A} + \mathbf{D}_B$ for $\mathbf{S}_B, \mathbf{D}_B \stackrel{\$}{\leftarrow} \chi^{\bar{n} \times n}$. This implies that $\mathbf{H} = \mathbf{S}_B \mathbf{Z} + \mathbf{D}_B \mathbf{W} = \mathbf{S}_B \mathbf{D}_A - \mathbf{D}_B \mathbf{S}_A$ and hence

$$\mathbf{V} = \mathbf{H} + \mathbf{B}_B \mathbf{S}_A + \mathbf{E}_B = \mathbf{S}_B \mathbf{D}_A - \mathbf{D}_B \mathbf{S}_A + \mathbf{B}_B \mathbf{S}_A + \mathbf{E}_B.$$

This implies that when the input tuple is the Extended-LWE instance then \mathcal{B} perfectly simulates the output of Γ_2 .⁵ On the other hand, if \mathbf{B}_B is uniformly random then \mathcal{B} perfectly simulates the output of Γ_3 . Finally the statement follows by further reducing the matrix version of ELWE to the standard one. \square

Game Γ_4 : First, we rename the variables $(\mathbf{S}_B, \mathbf{D}_B, \mathbf{E}_B) := (\mathbf{S}_{\text{sim}}, \mathbf{D}_{\text{sim}}, \mathbf{E}_{\text{sim}})$. Further, instead of picking \mathbf{B}_B uniformly at random, the experiment now samples alternative secrets/errors $\mathbf{S}_B, \mathbf{D}_B \stackrel{\$}{\leftarrow} \chi^{\bar{n} \times n}$ for \mathbf{B} and sets $\mathbf{B}_B := \mathbf{S}_B \mathbf{A} + \mathbf{D}_B$. The rest is identical as in Γ_3 .

Lemma 10. There exists an efficient algorithm \mathcal{B}' solves the $\text{LWE}_{n,n,\chi,q}$ problem with probability at least $\frac{1}{\bar{n}} |\varepsilon_3 - \varepsilon_2|$.

Proof. We describe a reduction \mathcal{B} which solves the matrix version of LWE. Then, the reduction to plain LWE follows by a hybrid argument. First, \mathcal{B} is given a tuple (\mathbf{A}, \mathbf{B}) where either $\mathbf{B} = \mathbf{S}_B \mathbf{A} + \mathbf{D}_B$ for short $\mathbf{S}_B, \mathbf{D}_B$ or \mathbf{B} is uniformly random. In either case, only given \mathbf{A} and \mathbf{B} , the reduction \mathcal{B} can simulate the rest of Γ_3 (and Γ_4). If $\mathbf{B} = \mathbf{S}_B \mathbf{A} + \mathbf{D}_B$ then this becomes Γ_4 , and when \mathbf{B}_B is uniformly random then \mathcal{B} simulates Γ_3 . \square

⁵We used the fact that χ is symmetric around 0 to argue that $\mathbf{S}_A := -\mathbf{W}$ is correctly distributed.

Sim($\mathbf{A}, \mathbf{B}_B, \mathbf{S}_A, \mathbf{D}_A, \mathbf{F}_A$)	
1:	$\mathbf{S}_{\text{sim}}, \mathbf{D}_{\text{sim}} \xleftarrow{\$} \chi^{\bar{n} \times n}$
2:	$\mathbf{E}_{\text{sim}} \xleftarrow{\$} \chi^{\bar{n} \times \bar{n}}$
3:	$\mathbf{V}_{\text{sim}} \leftarrow \mathbf{S}_{\text{sim}} \mathbf{D}_A - \mathbf{D}_{\text{sim}} \mathbf{S}_A + \mathbf{B}_B \mathbf{S}_A + \mathbf{E}_{\text{sim}}$
4:	$\text{ct} \leftarrow \text{HelpRec}(\mathbf{V}_{\text{sim}})$
5:	$\mathbf{K} \leftarrow \text{rec}(2\mathbf{V}_{\text{sim}}, \text{ct})$
6:	return (\mathbf{K}, ct)

Figure 3.11: Simulator for the deniability game.

Finally, we present the simulator in Figure 3.11. Γ_4 can now be alternatively described in the following way. The experiment first samples $\mathbf{S}_A, \mathbf{D}_A \xleftarrow{\$} \chi^{n \times \bar{n}}$ and $\mathbf{S}_B, \mathbf{D}_B \xleftarrow{\$} \mathbb{Z}_q^{\bar{n} \times n}$ and $\mathbf{F}_A \xleftarrow{\$} \chi^{\bar{n} \times \bar{n}}$. Further, the public keys are defined as $\mathbf{B}_A = \mathbf{A} \mathbf{S}_A + \mathbf{D}_A$ and $\mathbf{B}_B = \mathbf{S}_B \mathbf{A} + \mathbf{D}_B$. Finally, it runs $(\mathbf{K}, \text{ct}) \xleftarrow{\$} \text{Sim}(\mathbf{A}, \mathbf{B}_B, \mathbf{S}_A, \mathbf{D}_A, \mathbf{F}_A)$ and outputs $(\mathbf{A}, \mathbf{B}_A, \mathbf{B}_B, \mathbf{S}_A, \mathbf{D}_A, \mathbf{F}_A, \mathbf{K}, \text{ct})$. Thus, deniability follows by simply combining the previous lemmas. \square

decaps-CPA Security. Finally, we show that our split-KEM satisfies the decaps-CPA security notion (see Definition 22).

Lemma 11 (decaps-CPA Security). Let $\bar{n} = O(\lambda)$, m be such that the ciphertext space of sKEM is $\{0, 1\}^m$, and χ be a probability distribution over $[-\gamma, \gamma]$ symmetric around 0 for any $\gamma > 0$. Suppose $\text{LWE}_{n+\bar{n}, n, \chi, q}$ is hard. Then, for every efficient algorithm \mathcal{A} , the probability of winning the decaps-CPA game is at most $2^m \cdot (\delta_{\text{cpa}}^{\bar{n}^2} + \text{negl}(\lambda))$ where

$$\delta_{\text{cpa}} := \max_{\substack{\text{ct} \in \{0, 1\}^m \\ u \in \mathbb{Z}_2^B}} \Pr [\text{rec}(2w, \text{ct}) = u] . \quad (3.3)$$

Proof. Let \mathcal{A} be an efficient adversary against the decaps-CPA game. We prove the statement using the hybrid games described explicitly in Figure 3.13. In each game Γ_i , we define ϵ_i to be the probability that the efficient adversary \mathcal{A} wins the security game.

Game Γ_1 : This is the standard decaps-CPA game corresponding to the sKEM in Figure 3.9.

Game Γ_2 : In this game, the ciphertext ct is not given to the adversary anymore. Note that the first phase adversary outputting \mathbf{B} is now useless and it can be removed⁶, along with the operations needed to compute ct . Given the ciphertext space is $\{0, 1\}^m$ for some $m \in \mathbb{Z}$, we have $\epsilon_2 \geq \frac{1}{2^m} \epsilon_1$ as any adversary in Γ_2 can simulate the view of an adversary in Γ_1 by guessing ct .

Game Γ_3 : In this game, the only change is that instead of computing $\mathbf{B}_B = \mathbf{S}_B \mathbf{A} + \mathbf{D}_B$, it is picked uniformly at random from $\mathbb{Z}_q^{\bar{n} \times n}$. The indistinguishability between Γ_3 and Γ_2

⁶Formally, we would run it anyway to propagate the adversary's state st , but we omit this for simplicity.

$\Gamma_1(\mathcal{A})$	$\Gamma_2(\mathcal{A})$
1: $\mathbf{S}_A, \mathbf{D}_A \xleftarrow{\$} \chi(\mathbb{Z}_q^{n \times \bar{n}})$	1: $\mathbf{S}_A, \mathbf{D}_A \xleftarrow{\$} \chi(\mathbb{Z}_q^{n \times \bar{n}})$
2: $\mathbf{F}_A \xleftarrow{\$} \chi^{\bar{n} \times \bar{n}}$	2: $\mathbf{F}_A \xleftarrow{\$} \chi^{\bar{n} \times \bar{n}}$
3: $\mathbf{B}_A \leftarrow \mathbf{A}\mathbf{S}_A + \mathbf{D}_A$	3: $\mathbf{B}_A \leftarrow \mathbf{A}\mathbf{S}_A + \mathbf{D}_A$
4: $\text{pk}_A = (\mathbf{A}, \mathbf{B}_A)$	4: $\text{pk}_A = (\mathbf{A}, \mathbf{B}_A)$
5: $\text{sk}_A = (\mathbf{S}_A, \mathbf{D}_A, \mathbf{F}_A)$	5: $\text{sk}_A = (\mathbf{S}_A, \mathbf{D}_A, \mathbf{F}_A)$
6: $\mathbf{S}_B, \mathbf{D}_B \xleftarrow{\$} \chi^{\bar{n} \times n}$	6: $\mathbf{S}_B, \mathbf{D}_B \xleftarrow{\$} \chi^{\bar{n} \times n}$
7: $\mathbf{F}_B \xleftarrow{\$} \chi^{\bar{n} \times \bar{n}}$	7: $\mathbf{F}_B \xleftarrow{\$} \chi^{\bar{n} \times \bar{n}}$
8: $\mathbf{B}_B \leftarrow \mathbf{S}_B\mathbf{A} + \mathbf{D}_B$	8: $\mathbf{B}_B \leftarrow \mathbf{S}_B\mathbf{A} + \mathbf{D}_B$
9: $\text{pk}_B = (\mathbf{A}, \mathbf{B}_B)$	9: $\text{pk}_B = (\mathbf{A}, \mathbf{B}_B)$
10: $\mathbf{E}_B \xleftarrow{\$} \chi^{\bar{n} \times \bar{n}}$	10: $\mathbf{E}_B \xleftarrow{\$} \chi^{\bar{n} \times \bar{n}}$
11: $\mathbf{V} \leftarrow \mathbf{S}_B\mathbf{B}_A + \mathbf{E}_B$	11: $\mathbf{V} \leftarrow \mathbf{S}_B\mathbf{D}_A - \mathbf{D}_B\mathbf{S}_A + \mathbf{B}_B\mathbf{S}_A + \mathbf{E}_B$
12: $\text{ct} \leftarrow \text{HelpRec}(\mathbf{V})$	12: $\text{ct} \leftarrow \text{HelpRec}(\mathbf{V})$
13: $\mathbf{K} \leftarrow \text{rec}(2\mathbf{V}, \text{ct})$	13: $\mathbf{K} \leftarrow \text{rec}(2\mathbf{V}, \text{ct})$
14: $b \xleftarrow{\$} \mathcal{A}(\text{pk}_A, \text{pk}_B, \text{sk}_A, \mathbf{K}, \text{ct})$	14: $b \xleftarrow{\$} \mathcal{A}(\text{pk}_A, \text{pk}_B, \text{sk}_A, \mathbf{K}, \text{ct})$
15: return b	15: return b
$\Gamma_3(\mathcal{A})$	$\Gamma_4(\mathcal{A})$
1: $\mathbf{S}_A, \mathbf{D}_A \xleftarrow{\$} \chi(\mathbb{Z}_q^{n \times \bar{n}})$	1: $\mathbf{S}_A, \mathbf{D}_A \xleftarrow{\$} \chi(\mathbb{Z}_q^{n \times \bar{n}})$
2: $\mathbf{F}_A \xleftarrow{\$} \chi^{\bar{n} \times \bar{n}}$	2: $\mathbf{F}_A \xleftarrow{\$} \chi^{\bar{n} \times \bar{n}}$
3: $\mathbf{B}_A \leftarrow \mathbf{A}\mathbf{S}_A + \mathbf{D}_A$	3: $\mathbf{B}_A \leftarrow \mathbf{A}\mathbf{S}_A + \mathbf{D}_A$
4: $\text{pk}_A = (\mathbf{A}, \mathbf{B}_A)$	4: $\text{pk}_A = (\mathbf{A}, \mathbf{B}_A)$
5: $\text{sk}_A = (\mathbf{S}_A, \mathbf{D}_A, \mathbf{F}_A)$	5: $\text{sk}_A = (\mathbf{S}_A, \mathbf{D}_A, \mathbf{F}_A)$
6: $\mathbf{S}_B, \mathbf{D}_B \xleftarrow{\$} \chi^{\bar{n} \times n}$	6: $\mathbf{S}_B, \mathbf{D}_B \xleftarrow{\$} \chi^{\bar{n} \times n}$
7: $\mathbf{F}_B \xleftarrow{\$} \chi^{\bar{n} \times \bar{n}}$	7: $\mathbf{F}_B \xleftarrow{\$} \chi^{\bar{n} \times \bar{n}}$
8: $\mathbf{B}_B \leftarrow \mathcal{U}(\mathbb{Z}_q^{\bar{n} \times n})$	8: $\mathbf{B}_B \leftarrow \mathbf{S}_B\mathbf{A} + \mathbf{D}_B$
9: $\text{pk}_B = (\mathbf{A}, \mathbf{B}_B)$	9: $\text{pk}_B = (\mathbf{A}, \mathbf{B}_B)$
10: $\mathbf{E}_B \xleftarrow{\$} \chi^{\bar{n} \times \bar{n}}$	10: $\mathbf{E}_{\text{sim}} \xleftarrow{\$} \chi^{\bar{n} \times \bar{n}}$
11: $\mathbf{V} \leftarrow \mathbf{S}_B\mathbf{D}_A - \mathbf{D}_B\mathbf{S}_A + \mathbf{B}_B\mathbf{S}_A + \mathbf{E}_B$	11: $\mathbf{S}_{\text{sim}}, \mathbf{D}_{\text{sim}} \xleftarrow{\$} \chi^{\bar{n} \times n}$
12: $\text{ct} \leftarrow \text{HelpRec}(\mathbf{V})$	12: $\mathbf{V} \leftarrow \mathbf{S}_{\text{sim}}\mathbf{D}_A - \mathbf{D}_{\text{sim}}\mathbf{S}_A + \mathbf{B}_B\mathbf{S}_A + \mathbf{E}_{\text{sim}}$
13: $\mathbf{K} \leftarrow \text{rec}(2\mathbf{V}, \text{ct})$	13: $\text{ct} \leftarrow \text{HelpRec}(\mathbf{V})$
14: $b \xleftarrow{\$} \mathcal{A}(\text{pk}_A, \text{pk}_B, \text{sk}_A, \mathbf{K}, \text{ct})$	14: $\mathbf{K} \leftarrow \text{rec}(2\mathbf{V}, \text{ct})$
15: return b	15: $b \xleftarrow{\$} \mathcal{A}(\text{pk}_A, \text{pk}_B, \text{sk}_A, \mathbf{K}, \text{ct})$
	16: return b

Figure 3.12: Security games for the proof of Theorem 4. The lines in blue highlight the main differences from the previous game. The lines in gray correspond to the simulator defined in Figure 3.11.

Chapter 3. K-Waay: Fast and Deniable Post-Quantum X3DH Without Ring Signatures

follows directly from $\text{LWE}_{n,n,\chi,q}$.

Game Γ_4 : Now, instead of computing \mathbf{B}_A and \mathbf{V}' as:

$$\begin{bmatrix} \mathbf{B}_A \\ \mathbf{V}' \end{bmatrix} = \begin{bmatrix} \mathbf{A} \\ \mathbf{B}_B \end{bmatrix} \mathbf{S}_A + \begin{bmatrix} \mathbf{D}_A \\ \mathbf{F}_A \end{bmatrix},$$

the experiment samples $\mathbf{B}_A \xleftarrow{\$} \mathcal{U}(\mathbb{Z}_q^{n \times \bar{n}})$ and $\mathbf{V}' \xleftarrow{\$} \mathcal{U}(\mathbb{Z}_q^{\bar{n} \times \bar{n}})$ uniformly at random. Then, the reduction executes Lines 4 to 6 of Γ_4 . Clearly there is an efficient adversary which solves $\text{LWE}_{n+\bar{n},n,\chi,q}$ with probability at least $\frac{1}{\bar{n}}|\varepsilon_4 - \varepsilon_3|$.

Finally, since \mathbf{V}' is uniformly random, the probability that any adversary wins Γ_4 , i.e. $\mathbf{K}_A = \mathbf{K}'_A$, can be upper-bounded by $\delta_{\text{cpa}}^{\bar{n}^2}$ by definition of δ_{cpa} . The statement now follows by combining the previous hybrid games. \square

$\Gamma_1(\mathcal{A})$	$\Gamma_2(\mathcal{A})$	$\Gamma_3(\mathcal{A})$
1: $\mathbf{S}_A, \mathbf{D}_A \xleftarrow{\$} \chi(\mathbb{Z}_q^{n \times \bar{n}})$	1: $\mathbf{S}_A, \mathbf{D}_A \xleftarrow{\$} \chi(\mathbb{Z}_q^{n \times \bar{n}})$	1: $\mathbf{S}_A, \mathbf{D}_A \xleftarrow{\$} \chi(\mathbb{Z}_q^{n \times \bar{n}})$
2: $\mathbf{F}_A \xleftarrow{\$} \chi^{\bar{n} \times \bar{n}}$	2: $\mathbf{F}_A \xleftarrow{\$} \chi^{\bar{n} \times \bar{n}}$	2: $\mathbf{F}_A \xleftarrow{\$} \chi^{\bar{n} \times \bar{n}}$
3: $\mathbf{B}_A \leftarrow \mathbf{A}\mathbf{S}_A + \mathbf{D}_A$	3: $\mathbf{B}_A \leftarrow \mathbf{A}\mathbf{S}_A + \mathbf{D}_A$	3: $\mathbf{B}_A \leftarrow \mathbf{A}\mathbf{S}_A + \mathbf{D}_A$
4: $\mathbf{S}_B, \mathbf{D}_B \xleftarrow{\$} \chi^{\bar{n} \times n}$	4: $\mathbf{S}_B, \mathbf{D}_B \xleftarrow{\$} \chi^{\bar{n} \times n}$	4: $\mathbf{B}_B \xleftarrow{\$} \mathcal{U}(\mathbb{Z}_q^{\bar{n} \times n})$
5: $\mathbf{F}_B \xleftarrow{\$} \chi^{\bar{n} \times \bar{n}}$	5: $\mathbf{B}_B \leftarrow \mathbf{S}_B \mathbf{A} + \mathbf{D}_B$	5: $\mathbf{K}'_A, \text{ct}' \xleftarrow{\$} \mathcal{A}(\mathbf{A}, \mathbf{B}_A, \mathbf{B}_B)$
6: $\mathbf{B}_B \leftarrow \mathbf{S}_B \mathbf{A} + \mathbf{D}_B$	6: $\mathbf{K}'_A, \text{ct}' \xleftarrow{\$} \mathcal{A}(\mathbf{A}, \mathbf{B}_A, \mathbf{B}_B)$	6: $\mathbf{V}' \leftarrow \mathbf{B}_B \mathbf{S}_A + \mathbf{F}_A$
7: $\mathbf{B}, \text{st} \xleftarrow{\$} \mathcal{A}(\mathbf{B}_A, \mathbf{B}_B)$	7: $\mathbf{V}' \leftarrow \mathbf{B}_B \mathbf{S}_A + \mathbf{F}_A$	7: $\mathbf{K}_A \leftarrow \text{rec}(2\mathbf{V}', \text{ct}')$
8: $\mathbf{E}_B \xleftarrow{\$} \chi^{\bar{n} \times \bar{n}}$	8: $\mathbf{K}_A \leftarrow \text{rec}(2\mathbf{V}', \text{ct}')$	8: return $\mathbb{1}[\mathbf{K}_A = \mathbf{K}'_A]$
9: $\mathbf{V} \leftarrow \mathbf{S}_B \mathbf{B} + \mathbf{E}_B$	9: return $\mathbb{1}[\mathbf{K}_A = \mathbf{K}'_A]$	
10: $\text{ct} \leftarrow \text{HelpRec}(\mathbf{V})$		$\Gamma_4(\mathcal{A})$
11: $\mathbf{K} \leftarrow \text{rec}(2\mathbf{V}, \text{ct})$		1: $\mathbf{B}_A \xleftarrow{\$} \mathcal{U}(\mathbb{Z}_q^{n \times \bar{n}})$
12: $\mathbf{K}'_A, \text{ct}' \xleftarrow{\$} \mathcal{A}(\mathbf{A}, \mathbf{B}_A, \mathbf{B}_B, \text{ct}, \text{st})$		2: $\mathbf{B}_B \xleftarrow{\$} \mathcal{U}(\mathbb{Z}_q^{\bar{n} \times n})$
13: $\mathbf{V}' \leftarrow \mathbf{B}_B \mathbf{S}_A + \mathbf{F}_A$		3: $\mathbf{V}' \xleftarrow{\$} \mathcal{U}(\mathbb{Z}_q^{\bar{n} \times \bar{n}})$
14: $\mathbf{K}_A \leftarrow \text{rec}(2\mathbf{V}', \text{ct}')$		4: $\mathbf{K}'_A, \text{ct}' \xleftarrow{\$} \mathcal{A}(\mathbf{A}, \mathbf{B}_A, \mathbf{B}_B)$
15: return $\mathbb{1}[\mathbf{K}_A = \mathbf{K}'_A]$		5: $\mathbf{K}_A \leftarrow \text{rec}(2\mathbf{V}', \text{ct}')$
		6: return $\mathbb{1}[\mathbf{K}_A = \mathbf{K}'_A]$

Figure 3.13: Security games for the proof of Lemma 11. The lines in blue highlight the main differences from the previous game.

Remark. We highlight that the construction does not require a super-polynomial modulus. To see this, let us first focus on δ_{cpa} in (3.3). In the instantiation we set $B = O(1)$, $\bar{n} = O(\sqrt{\lambda})$ and $\gamma = O(1)$. Further, δ_{cpa} can be bounded by $2^{-B} + 1/q$, and if $q = \text{poly}(\lambda)$ then $\delta_{\text{cpa}} \leq (2^{-B} + 1/q)^{\bar{n}^2} = \text{negl}(\lambda)$. Hence, the winning probability in Lemma 11 is negligible for $q = \text{poly}(\lambda)$.

Now, let us move on to Equation 3.2. Recall that n is a dimension responsible for hardness of LWE, say $n = \text{poly}(\lambda)$. Suppose as in Lemma 11 that the distribution χ outputs integers between $-\gamma$ and γ . Then, to obtain $\delta_{\text{corr}} = 0$, we need to pick $q > 2^{B+2}(2n\gamma^2 + 2\gamma) = \text{poly}(\lambda)$, hence asymptotically the modulus can still be polynomial. In practice, we would allow negligible δ_{corr} (as we do for K-Waay), and thus we further decrease the modulus.

3.5.5 Building a UNF-1KCA and IND-1BatchCCA Split-KEM

KeyGen _{sKEM} (1^λ)	Encaps _{sKEM} (pk _A , sk _B)
1: (pk, sk) $\xleftarrow{\$}$ KeyGen _{sKEM₀} (1^λ)	1: $K_0, \text{ct} \xleftarrow{\$}$ Encaps _{sKEM₀} (pk _A , sk _B)
2: return (pk, sk)	2: $t \leftarrow H'(\text{pk}_A, \text{pk}_B, \text{ct}, K_0)$
	3: $K \leftarrow H(\text{pk}_A, \text{pk}_B, \text{ct}, K_0)$
	4: return $K, (\text{ct}, t)$
Decaps _{sKEM} (pk _B , sk _A , (ct, t))	
1: $K'_0 \leftarrow \text{Decaps}_{\text{sKEM}_0}(\text{pk}_B, \text{sk}_A, (\text{ct}, t))$	
2: if $H'(\text{pk}_A, \text{pk}_B, \text{ct}, K'_0) \neq t$:	
3: return \perp	
4: return $H(\text{pk}_A, \text{pk}_B, \text{ct}, K'_0)$	

Figure 3.14: T_{CH} transform for split-KEMs. We assume that pk_B can be derived from sk_B or is contained in it.

We have proven so far that the modified version of FrodoKEX given above is decaps-CPA and OW-CPA. We show now that any scheme satisfying both these properties can easily be transformed into a UNF-1KCA and IND-1BatchCCA split-KEM in the ROM and QROM. The construction is similar to the T_{CH} transform introduced by Huguenin-Dumittan and Vaudenay [HV22] translated to the split-KEM setting. We present it in Figure 3.14. Then, the following theorem states the security guarantees of the resulting split-KEM.

Theorem 5. Let sKEM₀ be any split-KEM and sKEM := $T_{\text{CH}}(\text{sKEM}_0)$ be the split-KEM obtained from applying the T_{CH} transform (Figure 3.14) to sKEM₀. Then, in the ROM, we have that for any efficient UNF-1KCA adversary \mathcal{A} , one can build efficient \mathcal{B} and \mathcal{C} adversaries s.t.

$$\text{Adv}_{\text{sKEM}}^{\text{unf-1kca}}(\mathcal{A}) \leq \frac{q_{H'}^2 + 1}{2^s} + (q_H + q_{H'} + 1) \cdot \text{Adv}_{\text{sKEM}_0}^{\text{decaps-cpa}}(\mathcal{C}),$$

where q_H and $q_{H'}$ are the number of queries made by \mathcal{A} to the random oracles H and H' , respectively, and s is the output size of both random oracles. In the QROM, the bound becomes

$$\text{Adv}_{\text{sKEM}}^{\text{unf-1kca}}(\mathcal{A}) \leq \frac{8(q_H + q_{H'})^2}{2^{2s}} + \epsilon + 2(2(q_H + q_{H'}) + 1)^2 \cdot \text{Adv}_{\text{sKEM}_0}^{\text{decaps-cpa}}(\mathcal{B}),$$

where $\epsilon := \frac{2}{2^s} + 8\sqrt{2/2^s} + \frac{40e^2(q_{H'}+2)^3+2}{2^s}$.

Chapter 3. K-Waay: Fast and Deniable Post-Quantum X3DH Without Ring Signatures

n	\tilde{n}	q	B	χ	$ t $	$ \text{pk} $	$ \text{ct} $
1452	8	31751	4	$\mathcal{U}(\{-1, 1\})$	64B	21.3KB	72B

Table 3.2: Concrete parameters for our lattice-based split-KEM. We note that in practice, we do not need to include the whole matrix \mathbf{A} in the public key pk , but rather the seed for the pseudorandom function to generate it (as is the case in this table). The ciphertext ct comprises the original split-KEM ciphertext (8B) and the tag t (64B).

$\tilde{n}^2 \delta_{\text{corr}}$ (3.2)	δ_{elwe} (3.1)	δ_{cpa} (3.3)
2^{-48}	2^{-46}	$2^{-3.9996}$

Table 3.3: Correctness and security terms.

Proof. We defer the proof to Appendix A.2. □

Similarly, we have that the T_{CH} transform makes an IND-1BatchCCA scheme out of an OW-CPA one, which is stated in the following theorem.

Theorem 6. Let sKEM_0 be any split-KEM and $\text{sKEM} := T_{\text{CH}}(\text{sKEM}_0)$ be the split-KEM obtained from applying the T_{CH} transform (Figure 3.14) to sKEM_0 . Then, in the ROM, we have that for any efficient IND-1BatchCCA adversary \mathcal{A} , one can build efficient \mathcal{B} s.t.

$$\text{Adv}_{\text{sKEM}}^{\text{ind-1batchcca}}(\mathcal{A}) \leq \frac{q_{H'}^2 + d}{2^s} + 2(q_H + q_{H'} + d) \cdot \text{Adv}_{\text{sKEM}_0}^{\text{ow-cpa}}(\mathcal{B})$$

where q_H and $q_{H'}$ are the number of queries made by \mathcal{A} to the random oracles H and H' , respectively, s is the output size of both random oracles, and d is the number of tuples submitted to the IND-1BatchCCA oracle BatchDEC. In the QROM, the previous bound becomes

$$\text{Adv}_{\text{sKEM}}^{\text{ind-1batchcca}}(\mathcal{A}) \leq \delta + \epsilon_1 + \epsilon_2 + \epsilon_3 + 2(q_H + d + q_{H'}) \sqrt{2 \text{Adv}_{\text{sKEM}_0}^{\text{ow-cpa}}(\mathcal{B})},$$

where δ is the correctness error, $\epsilon_1 = \frac{40e^2(q_{H'}+d+1)^3+2}{2^s}$, $\epsilon_2 = 8d(d+2q_{H'}+1)\sqrt{2/2^s}$ and $\epsilon_3 = \frac{4d}{2^s}$.

Proof. As the proof is nearly identical to the proof of IND-qCCA security of the T_{CH} transform for PKE/KEM [HV22], we defer it to the Appendix A.3. □

3.5.6 Concrete Instantiation

In Table 3.2 we propose a parameter set for FrodoKEX+ where we aim for 256-bit security before applying the transform and 128-bit (resp. 64-bit) security after the transform assuming 2^{64} random oracle (resp. quantum random oracle) queries. In addition, we give the security terms in Table 3.3. We show in the following how these parameters were computed, where we set $(B, \tilde{n}) = (4, 8)$.

Correctness Error and Security Loss. One of the main challenges in instantiating our FrodoKEX variant is computing δ_{corr} and δ_{elwe} from Equations 3.2 and 3.1. They are related to the correctness error and the security of loss of ELWE. To this end, we introduce a simple distribution χ which will allow us to efficiently compute these values. Namely, we set χ to be a uniform distribution over the set $\{-1, 1\}$. Clearly, it is symmetric around 0 and has standard deviation equal to 1.

Another useful property of this distribution is that a product XY , where $X, Y \stackrel{\$}{\leftarrow} \chi$, still follows the distribution of χ . Based on this observation, we have

$$\delta_{\text{corr}} = \Pr_{\substack{X_1, \dots, X_{2n+1} \stackrel{\$}{\leftarrow} \chi \\ E \stackrel{\$}{\leftarrow} \chi}} \left[\left| \sum_{i=1}^{2n+1} X_i + E \right| > \frac{q}{2^{B+2}} \right].$$

We can directly compute this term using Laurent polynomials. Namely, define

$$P(X) := \Pr_{X \stackrel{\$}{\leftarrow} \chi} [X = 1] \cdot X + \Pr_{X \stackrel{\$}{\leftarrow} \chi} [X = -1] \cdot X^{-1} = \frac{1}{2} \cdot (X + X^{-1}).$$

Then, using the convolution properties, we observe that the probability of $X_1 + \dots + X_{2n+2} = k$, for some $-2n - 2 \leq k \leq 2n + 2$, is equal to the k -th coefficient of the polynomial $P(X)^{2n+2}$. Hence, we calculate δ_{corr} by computing $P(X)^{2n+2}$ and summing all the k -th coefficients, such that $2n + 2 \geq |k| > \frac{q}{2^{B+2}}$.

We now turn into computing δ_{elwe} . The first step is the analysis of the following random variable $\mathbf{v} = \frac{1}{2} \cdot (e - d) \cdot \mathbf{z}$, where $e, d \stackrel{\$}{\leftarrow} \chi$ and $\mathbf{z} \stackrel{\$}{\leftarrow} \chi^{\bar{n}}$. We denote this distribution as \mathcal{V} . By simple calculation we get:

$$\Pr[\mathbf{v} = \mathbf{a}] = \Pr \left[\frac{1}{2} \cdot (e - d) \cdot \mathbf{z} = \mathbf{a} \right] = \begin{cases} \frac{1}{2} & \text{if } \mathbf{a} = \mathbf{0} \\ \frac{1}{2^{\bar{n}+1}} & \text{if } \mathbf{a} \in \{-1, 1\}^{\bar{n}} \\ 0 & \text{otherwise} \end{cases}$$

Then, the multivariate Laurent polynomial corresponding to \mathbf{v} has an elegant form:

$$P(X_1, \dots, X_{\bar{n}}) = \frac{1}{2} + \frac{1}{2^{\bar{n}+1}} \prod_{i=1}^{\bar{n}} (X_i + X_i^{-1}).$$

As before, we observe that δ_{elwe} is the probability that for $\mathbf{v}_1, \dots, \mathbf{v}_{n+m} \stackrel{\$}{\leftarrow} \mathcal{V}$,

$$2 \cdot (\mathbf{v}_1 + \dots + \mathbf{v}_{n+m}) = \mathbf{0} \pmod{q} \iff \mathbf{v}_1 + \dots + \mathbf{v}_{n+m} = \mathbf{0}^7.$$

In terms of the newly defined Laurent polynomials, δ_{elwe} is the constant coefficient of:

$$P(X_1, \dots, X_{\bar{n}})^{n+m} = \sum_{j=0}^{n+m} \binom{n+m}{j} \frac{1}{2^{n+m-j}} \cdot \frac{1}{2^{(\bar{n}+1)j}} \cdot \prod_{i=1}^{\bar{n}} (X_i + X_i^{-1})^j.$$

⁷This holds as long as $n + m < q/2$ since then no modulo overflow occurs.

Chapter 3. K-Waay: Fast and Deniable Post-Quantum X3DH Without Ring Signatures

We now look at the constant coefficient of each of the $n + m + 1$ terms of the sum. The first observation is that

$$(X_i + X_i^{-1})^j = \sum_{k=0}^j \binom{j}{k} X_i^{(j-k)} X_i^{-k} = \sum_{k=0}^j \binom{j}{k} X_i^{(j-2k)}.$$

Hence, the constant coefficient of the expression above is 0 if j is odd, and $\binom{j}{j/2}$ when j is even. Consequently, the constant coefficient of $\prod_{i=1}^{\bar{n}} (X_i + X_i^{-1})^j$ is either 0, for odd j , or $\binom{j}{j/2}^{\bar{n}}$ for even j . Hence, we conclude that

$$\delta_{\text{elwe}} = \sum_{j \text{ even}} \binom{n+m}{j} \frac{1}{2^{n+m-j}} \cdot \frac{1}{2^{(\bar{n}+1)j}} \cdot \left(\frac{j}{2}\right)^{\bar{n}}$$

which can then be computed efficiently for our parameters. Finally, δ_{cpa} can be straightforwardly computed for small primes⁸, such as $\approx 2^{15}$.

Hardness of Extended-LWE. We measure the hardness following the methodology used for the original FrodoKEX [BCD⁺16] for fair comparison, and refer to it for more details on the attacks. Here, the main bottleneck of setting the parameters is the reduction loss between ELWE and plain LWE. Taking this into account for the parameters proposed above, we aim for 307-bit classical LWE security.

We consider the primal and dual BKZ attacks [SE94, CN11]. As a subroutine, the BKZ algorithm with block-size b uses an algorithm for the shortest vector problem (SVP) in lattices of dimension b . As in Frodo [BCD⁺16], for precautionary purposes we only count the cost of one such call (even though in practice it will run the SVP sub-algorithm polynomially many times). The lower-bound on the time complexity of one call is given by about $b2^{cb}$ CPU cycles, where $c \approx 0.292$ for classical attacks, and $c \approx 0.265$ for quantum attacks (see Laarhoven [Laa16, Section 14.2.10]). For 307-bit classical security, this corresponds to the block size being 1018, and the root Hermite factor being ≈ 1.0020 (in the quantum setting these parameters correspond to 279 bits of security). Further, we estimate the hardness of LWE against known attacks using the LWE estimator by Albrecht et al. [APS15]. Namely, we run the estimator under both “sieving” and “enumeration”, and set the final root Hermite factor δ as the largest root Hermite factor returned by the program. Finally, we make sure that $\delta_{\text{cpa}}^{\bar{n}^2} \approx 2^{-256}$ for the decaps-CPA proof.

3.6 Evaluation and Discussion

In this section, we evaluate empirically our protocol both in terms of time and space usage, and then discuss some aspects of our modelling and protocol. Hereafter, we refer to the X3DH-like protocol of Brendel et al. [BFG⁺22b] as SPQR, and the baseline deniable protocol

⁸One can formally prove that δ_{cpa} can be bounded by $\frac{1}{2^B} + \frac{1}{q}$, but we compute the value directly instead.

Scheme	Cl. (C)	Cl. (Q)	ROM bnd	QROM bnd	Assumption
FrodoKEX+	128	64	$(q_H + d)/2^{192}$	$(q_H + d)/2^{128}$	LWE
Raptor [LAZ19]	114	103	?	✗	NTRU
DualRing-LB [YEL ⁺ 21]	(128)	(64)	?	✗	MSIS, MLWE
Falafel [Beu20]	128	64	?	✗	MSIS, MLWE

Table 3.4: Security comparison between FrodoKEX+ and several post-quantum RS. ‘Cl.’ stands for claimed number of security bits. DualRing-LB’s authors do not seem to make a clear security claim, we thus assume NIST level I. ‘?’ indicates that no bound is explicitly given for the security, ‘✗’ indicates that no proof is provided in the QROM.

(i.e., with ring signatures and without NIZKs) by Hashimoto et al. [HKKP22] as HKKP.

3.6.1 Benchmarks

Security of the Relevant Non-Standard Primitive. Like K-Waay, SPQR and HKKP can each be implemented using only (soon to be) standardised primitives, except for a single primitive in each case, we consider here the security of the relevant non-standard primitives. In the case of K-Waay it is a split-KEM, here implemented using a variant of FrodoKEX passed through the T_{CH} transform (that we call FrodoKEX+), and in the case of both HKKP and SPQR it is a ring signature scheme, or RS (or a designated-verifier signature scheme (DVS) derived from RS). The authors of both SPQR and HKKP proposed possible implementations for the RS without picking one in particular. The most efficient one for a ring of size 2 we are aware of that has an existing C implementation is Raptor [LAZ19] which we use for the benchmarks below. Other candidates would be Falafel [Beu20] or DualRing-LB [YEL⁺21].

We present in Table 3.4 a summary of the security claims, approximate leading factor in the bounds in the (Q)ROM, and assumptions for these non-standard primitives. We note that none of these primitives are proven secure in the standard model and all are based on lattices.

First, we note that parameters for these RS schemes were chosen *before* the reduction in the ROM was performed. That is, a primitive P based on lattices is built, parameters are chosen such that P satisfies the security claim, then P is used to build a RS in the ROM, which incurs a loss factor that usually depends on the number of queries to the random oracle q_H . In particular, it is common to have at least a q_H factor in the security bound (e.g. if the adversary can make 2^{32} queries to the RO, the security level is reduced by 32 bits). Therefore, the claimed security level does not match the *provable* security level. In the QROM, the security loss is usually greater: square root and q_H^2 or q_H^3 losses are quite common, however these schemes have not been proven secure in this model.

We took a different approach in designing a split-KEM with a conservative assumption (i.e., plain LWE) and choice of parameters. Therefore, FrodoKEX+ with our proposed parameters achieve 128 (resp. 64) bits of classical (resp. quantum) security *after* the (Q)ROM proof. We provide the (approximate) highest terms of both the ROM and QROM security bounds in

Chapter 3. K-Waay: Fast and Deniable Post-Quantum X3DH Without Ring Signatures

Table 3.4. These satisfy our security claims as long as $q_H \leq 2^{64}$ in the ROM and $q_H \cdot d \leq 2^{64}$ in the QROM, where d is the number of public key/ciphertext tuples allowed in the IND-1BatchCCA game. In K-Waay, d corresponds to the number of distinct users trying to communicate with an offline receiver after all prekeys have run out, and thus should typically be small. We note that the leading coefficients in all our bounds are always related to the (Q)ROM security (e.g. finding a collision) and not to the underlying hardness assumptions. Therefore, one can easily make these coefficients smaller by picking a larger digest size n , e.g. 512 bits.

The reason behind the approximations and lack of QROM proofs for PQ ring signatures is likely the youth of the field and the speed at which it is evolving. Still, we believe it is worth noting as it makes any comparison between our protocol and previous ones quite difficult.

Benchmarking. The protocols we benchmarked are: our own implementation of the X3DH protocol; Brendel et al.’s [BFG⁺22b] protocol SPQR based on PQ KEMs, a signature scheme and DVS; Hashimoto et al.’s [HKKP22] protocol HKKP based on PQ KEMs, a signature scheme and RS; a baseline protocol made only with PQ KEMs and a signature scheme similar to the *non-deniable* variant of HKKP; and our protocol K-Waay based on FrodoKEX+ as a PQ split-KEM, PQ KEMs and a signature scheme.⁹

We chose Kyber-512 as the KEM, both Falcon-512 and Dilithium2 for signatures, and Raptor for ring signatures. We implemented both HKKP and SPQR with signed prekeys as is the case in Signal’s implementation of X3DH. That is, a PQ signature key pair is part of the long-term key, and ephemeral keys uploaded to server are signed with it. Note that this is made explicit in K-Waay as the ephemeral keys are signed with the long-term one. The authors of HKKP show that this is not necessary in their protocol, however not doing so weakens perfect forward secrecy.

We built the different protocols in C using the `liboqs` library¹⁰ for Kyber, Falcon, and Dilithium, the Raptor implementation provided by the authors¹¹, and a modified version of the `lwe-frodo` library¹² with scaled parameters to properly simulate FrodoKEX+. More precisely, the modulus was set to the first power of 2 larger than the modulus in FrodoKEX+, the addition of the noise during decapsulation was also added, and the noise distributions were modified to match the ones of FrodoKEX+. We did not optimise the scheme in any way (e.g. by using AVX instructions or parallelisation) and we leave this as future work. For the sake of completeness, we also provide a reference implementation of FrodoKEX+ in Rust¹³ for the interested reader. All benchmarks were run on a virtual machine running Ubuntu 22.04 with 2 cores of an Intel i7-9750H running at 2.60GHz and 4GB of RAM allocated.

⁹SPQR and HKKP do not formally treat (regular) signatures, but we include them for fair comparison and because a practical system would use signatures for prekey bundles.

¹⁰<https://github.com/open-quantum-safe/liboqs>

¹¹<https://github.com/zhenfeizhang/raptor>

¹²<https://github.com/lwe-frodo/lwe-frodo>

¹³<https://github.com/lehugueni/frodokexp-rust>

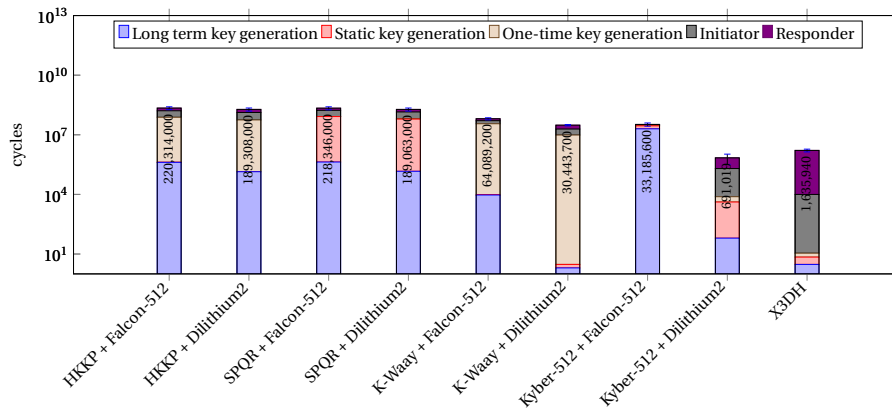


Figure 3.15: Speed benchmark for X3DH protocols.

Speed. For the speed benchmark, we measured how many cycles each protocol takes in one execution. We summarise our results in a logarithmic graph on Figure 3.15 (note that the internal division of the bars is linear).

Fixing the choice of KEM and signature scheme, our protocol K-Waay, depending on the choice of KEM and signature scheme, is between 3 and 6 times faster than the previous proposals even with our relatively conservative parameter choice. In our protocol K-Waay using Dilithium2, most cycles are spent in the ephemeral key generation, while using Falcon makes the static key generation as expensive as the ephemeral key one. Overall, one can see that Falcon, while more compact than Dilithium2, has a great impact on efficiency. For instance, K-Waay with Dilithium2 is faster than the non-deniable scheme using Kyber and Falcon.

Apart from Falcon, we see that the most time-consuming primitives are the non-standard ones, i.e., ring signatures and split-KEM. Hence, we see that the KEM+SIG protocol (HKKP’s baseline proposal) performs even better than X3DH, which shows once again that lattice-based schemes can be faster than their classical counterparts. However, we recall that it does not provide deniability. At last, we note that X3DH is the only construction that spends more time in sending and receiving than generating keys. Our protocol’s Send and Receive (i.e., BatchReceive with a single input message) procedures are very fast.

Data Size. In Table 3.5, we provide for each scheme the size of the long-term keys, the prekeys (output by Init in our DAKE syntax), and the ciphertext output by the sender. We computed for both HKKP and SPQR the size with and without long-term signatures. We see that K-Waay compares well in term of long-term public key and ciphertext size as both are smaller than signed HKKP and SPQR. However, the prekeys are much larger as one could expect from a LWE-based scheme and due to our conservative setting of parameters.

Chapter 3. K-Waay: Fast and Deniable Post-Quantum X3DH Without Ring Signatures

Scheme	$ pk $	$ prek $	$ ct $
K-Waay + Dilithium	2112	24520	1632
K-Waay + Falcon	1697	22790	1632
HKKP [HKKP22]	1700	1700	4056
HKKP [HKKP22] + Dilithium2	3012	4120	4056
HKKP [HKKP22] + Falcon	2597	2390	4056
SPQR [BFG ⁺ 22b]	3400	1632	4824
SPQR [BFG ⁺ 22b] + Dilithium2	4712	4052	4824
SPQR [BFG ⁺ 22b] + Falcon	4297	2322	4824

Table 3.5: Size comparison in bytes between K-Waay instantiated with FrodoKEX+, HKKP [HKKP22] and SPQR [BFG⁺22b]. We also computed the sizes for both HKKP and SPQR implemented with signed prekey bundles.

3.6.2 Advantages, Limitations and Discussion

Running Out of Ephemeral Keys. The main disadvantage of our protocol is that running out of ephemeral keys requires the receiver to abort if *any* of the sessions that used the same prekey is bogus. If this happens, then a malicious party could mount some kind of denial of service (DoS) attack against the user that was offline for too long by sending a bogus split-KEM ciphertext. There is an obvious trade-off between the risk of such an attack happening and the number of ephemeral keys uploaded on the server (and thus also storage). We leave the analysis and the mitigation of such a threat as future work, but we believe that if a reasonable amount of prekeys are uploaded, creating fake accounts is difficult (e.g., by requiring a phone number as in Signal), and/or users are online often enough, such an attack would be difficult to mount. Furthermore, several practical mitigations are possible. For instance, if the receiver (i.e., the victim) received a bogus ciphertext among the n ciphertexts sent for the same prekey while offline, they can restart K-Waay with each the n parties but as the initiator, which will probably succeed. The victim could also send n new prekeys to the n initiators directly, making sure the protocol will succeed at the next iteration. This would make the attack less useful as it could only *delay* communication and not prevent it.

We also think it is worth mentioning that the trick we propose might be easy to misimplement. In particular, it is crucial that no information about which split-KEM ciphertext failed leaks if such a situation occurs. That is, precautions should be taken such that leakage via side-channels in the scope of the system designer’s threat model are prevented.

Split-KEM Instead of Ring Signatures. The fact that we use a primitive similar to a post-quantum KEM allows us to leverage the extensive literature on the topic and existing safe/optimised implementations. This also gives good security guarantees as post-quantum KEMs have been heavily scrutinised as part of the NIST standardisation process. For example, as mentioned above, our proposed lattice-based implementation is based on a key-exchange variant of FrodoKEM, which is itself the PQ KEM recommended by the German Federal Office for Information Security (BSI) [fIS23]. Overall we think that a split-KEM such as FrodoKEX+

is more mature and closer to being usable in practice than ring signatures.

On the Necessity of Modifying FrodoKEX. Currently, our split-KEM significantly differs from the original FrodoKEX in two aspects: (i) the modulus for our construction has to be prime in order for our reduction from Extended-LWE to LWE to hold¹⁴, and (ii) we have to introduce additional masking terms to prove UNF-1KCA security. However, we believe that both changes are artefacts of the security proofs, and the original FrodoKEX split-KEM should be (up to a reasonable security loss) deniable.

There are alternative reduction techniques from Extended-LWE to LWE in the literature [BJRW21, BLP⁺13], which do not rely on having an odd modulus at the cost of using discrete Gaussian error distributions with large parameters. It is thus an interesting research problem to efficiently reduce Extended-LWE to LWE for even modulus with small reduction loss. In practice, the most efficient LWE attacks do not consider the structure of the modulus, so intuitively this should translate to the Extended-LWE setting¹⁵.

As for our second main modification, it is not immediately clear how to argue UNF-1KCA security without the additional masking terms. Hence, we leave deniability and UNF-1KCA security of the original FrodoKEX construction as exciting future work.

Deniability. While the signature on the ephemeral public keys might give the impression that our protocol is less deniable than X3DH or previous PQ alternatives, this is actually not the case. The reason is that prekey bundles in these protocols are signed as well, but this detail is abstracted away in the analysis (i.e., it is assumed that all parties have received and authenticated all public keys before the protocol actually starts). While this kind of analysis allows for strong deniability claims, in practice these protocols do not satisfy something stronger than some kind of peer-deniability. The exception is the ring signature based variant by Hashimoto et al. [HKKP22], where the prekey bundle is not necessarily signed. However, in this variant, the authors prove the security of their protocol in a weaker model than their non-deniable, signature-based protocol (i.e., it satisfies a weaker notion of forward secrecy). Overall, if deniability should not come at the price of security, peer-deniability seems like the best notion one can achieve in these DAKEs.

We wished to provide a transparent model for peer-deniability, where the upload of signed ephemeral keys is made explicit. We also strengthen the deniability definition of Brendel et al. [BFG⁺22b] by allowing the exposure of one of the parties (i.e., the receiving one, which would be the malicious party trying to frame the sender). While our protocol satisfies our stronger (in term of key exposure) notion of deniability, we believe both previous PQ X3DH alternatives satisfy it as well. Indeed, in these schemes, the ephemeral keys are KEM and RS keys only, which are deniable. Hence, exposing these should not harm deniability.

¹⁴Recall FrodoKEX [BCD⁺16] uses a power-of-two modulus for efficiency.

¹⁵Recently, various frameworks have been developed [DDGR20, DSGHK23], which measure concrete hardness of LWE given hints of specific form, such as linear combination of secrets with short random coefficients.

Chapter 3. K-Waay: Fast and Deniable Post-Quantum X3DH Without Ring Signatures

Hashimoto et al. [HKKP22] consider a strong notion of deniability where the adversary is malicious (i.e., can arbitrarily deviate from the protocol) and show how to modify HKKP such that it is secure against such a threat. However, such deniability comes at the expense of NIZKs, which are complex, expensive and are not always proven secure in the QROM when random oracles are used. Moreover, as in other deniable systems against malicious adversaries, non-falsifiable assumptions (i.e., knowledge-type assumptions) are required to prove the security. In addition, it seems difficult to defend against adversaries actively trying to frame a given user in messaging in practice [GPA19, CCHD23]; for example, an adversary could also simply ask questions that would identify the victim with good probability. Because of these reasons, we do not consider such a notion of deniability here.

To contextualise our results, we remark here that cryptographic deniability, which is targeted in this chapter and all previous work on deniable X3DH key exchange, translates to deniability on a *system* level only if the application preserves deniability. For example, Collins et al. [CCHD23] observe that Signal as currently deployed does not provide this kind of ‘practical’ deniability for ordinary users. Suppose Bob is trying to frame Alice and hands over their phone that contains a transcript of communication between Alice and Bob to a judge. Because Signal authenticates users (either directly or indirectly through Signal sealed sender [Lun18]), unless Bob was able to modify their phone (which depends on the technical expertise of Bob), the judge can deduce that the conversation plausibly took place as in the transcript, regardless of the cryptographic protocols employed underneath. It is interesting future work to further explore deniability on the broader system level and practical deniability [RMA⁺23, YGS23].

An Optimisation. As presented in Section 3.4, the K-Waay protocol generates a signature for each ephemeral public key uploaded. This can easily be optimised by signing the whole prekey bundle containing several ephemeral keys. This way, the server needs to store only one PQ signature for each user. The downside is that now each user needs to download the whole bundle to verify the signature. This offers a trade-off between data stored at the server and sent to clients.

4 On Active Attack Detection in Messaging with Immediate Decryption

In this chapter, we explore different trade-offs between security and performance for active attack detection in two-party messaging when messages can be delivered in a different order to what they were sent in by parties (i.e., messaging with immediate decryption). An extended abstract corresponding to this work appeared at CRYPTO 2023. The work presented in this chapter is the result of a close collaboration with Simone Colombo, noting that Loïs Huguenin-Dumittan was the primary contributor for the lower bound result which we include for completeness, with further contributions from Loïs, Khashayar Barooti and Serge Vaudenay. A full version of this work can be found on the Cryptology ePrint Archive [BCC⁺23b].

4.1 Contribution

The asynchronicity of messaging and the unreliability of networks has driven the design of ratcheting-based protocols with *immediate decryption* [ACD19, PP22, BRT23, CZ24], i.e., the support of out-of-order delivery and message loss on the *protocol* level. This property ensures that legitimate messages can be immediately decrypted by the receiver upon arrival and placed correctly among other received messages. Furthermore, communication can continue even if some messages are permanently lost. As highlighted in Chapter 1, the Signal protocol, the current de-facto messaging standard, supports immediate decryption. By contrast, many schemes in the literature fail if even a single message is lost (see [BS]⁺17, PR18, DV19, CDV21] for a non-exhaustive list).

The aforementioned security notions do not guarantee message authentication when the adversary impersonates parties, e.g., through state compromise. The lack of authentication implies that parties cannot *detect* active attacks. A recent phishing attack against Signal's phone number verification service enabled attackers to re-register accounts to another device, demonstrating the practicality of impersonation attacks via secret state compromise [Sup22]. Similar attacks that steal verification codes to hijack accounts affect a plethora of messaging applications. The proliferation of spyware such as Pegasus represents an additional—and worrying—threat for secret exfiltration [SRCM⁺22].

Chapter 4. On Active Attack Detection in Messaging with Immediate Decryption

The most widely used mechanisms for detecting active attacks use an *out-of-band* authenticated channel. All deployed mechanisms we are aware of in practice [Mar17] and many proposed in the literature [DH21, DGP22] assume such a channel. Solutions like Signal’s safety numbers [Mar17] enable parties to authenticate their long-term keys by comparing QR codes in person. However, as observed by Dowling and Hale [DH20, DH21], Signal’s approach fails to provide guarantees after a user’s state is exposed, since the safety numbers protocol only authenticates the session key output from X3DH.

To remedy this situation, Dowling and Hale [DH21] proposed to add an additional authentication key to each iteration of Signal’s asymmetric ratchet for on-demand use in out-of-band authentication. Their construction allows parties to immediately—that is, without additional communication rounds—authenticate their entire *asymmetric ratchet* out of band. However, messages forged under symmetric keys will never be detected. The only other construction in the literature to our knowledge, proposed by Dowling, Günter and Poirrier [DGP22], requires *three rounds* of in-band communication before an out-of-band hash comparison can take place. Contrary to Dowling and Hale’s solution, this approach authenticates all messages (albeit does not formally treat out-of-order messages), but imposes additional rounds, which is especially problematic in the presence of an active adversary. This raises our first research question:

1. Can we authenticate *all* messages in a *single round* of out-of-band communication to detect active attacks in the immediate decryption setting?

Out-of-band authentication is not always practical or even possible. A convenient alternative is to detect active attacks *in-band*, i.e., using the same channel as the messaging protocol. The adversary can, in the worst case, block all messages sent by honest parties, thereby forcing users to resort to out-of-band communication, but mounting such a persistent attack requires considerable resources. Durak and Vaudenay [DV19] thus introduced RECOVER security: if a party receives a forgery, then this party does not accept subsequent messages sent honestly by his counterpart. Caforio et al. [CDV21] then extended RECOVER security to enable a party to detect whether their *partner* was compromised, i.e., whether they received a forgery. By contrast to out-of-band authentication, no additional messages are required to support attack detection: in-band ciphertexts contain the authentication information. However, these notions and the corresponding constructions assume in-order message delivery and fail on message dropping. This raises a second question, first suggested by Alwen et al. [ACD19]:

2. Can we achieve extended RECOVER security—immediate in-band active attack detection—while supporting immediate decryption?

To detect active attacks, parties need to authenticate their *entire* message history: each message may be a forgery, i.e., the result of an active attack. With immediate decryption, parties cannot be sure which messages their partner has received until they receive an honest reply

from them. Intuitively, each message needs to “contain” the message history up until when it was sent. Looking ahead, we formalise and confirm this intuition, which motivates the exploration of performance/security trade-offs and optimisations. In this regard, existing protocols for both in-band and out-of-band active attack detection represent only a subset of the potential design space. Consequently we also ask:

3. What are the *communication costs* of in- and out-of-band active attack detection for messaging with immediate decryption, and what useful performance/security *trade-offs* can be made?

4.1.1 Summary

In this chapter, we explore the aforementioned questions. In more detail:

- We introduce (Section 4.2) a primitive that we call *authenticated ratcheted communication*, which captures immediate decryption and models communication through both insecure in-band and authentic out-of-band channels.
- In Section 4.3, we formalise in-band active attack detection for immediate decryption, by extending RECOVER security [DV19, CDV21], with two notions, namely r-RID and s-RID security, for detecting active attacks towards the receiver and on reception of messages from the sender after they were attacked; combined, these two notions comprise RID (recover with immediate decryption) security. We propose a scheme secure under these notions.
- We consider *out-of-band* active attack detection for immediate decryption messaging in Section 4.4. We introduce notions r-UNF and s-UNF (which combine to UNF for unforgeable), which are analogous to the notions for in-band detection. Demonstrating their similarity, we construct a UNF-secure scheme from a RID-secure scheme. We also construct a UNF-secure ARC scheme given a RC scheme.
- In Section 4.5, we prove with an information-theoretic argument that ciphertexts in a scheme satisfying either r-RID or r-UNF security must grow *linearly* in the number of messages sent, implying that attaching the set of sent ciphertexts to each message as we do is essentially optimal.
- In Section 4.6, we consider different ways to bypass the aforementioned lower bounds. First, we discuss ways to optimise the s-RID-secure scheme. We show how one can drastically reduce the overhead as long as the communication between the two parties is relatively synchronised. We believe the corresponding scheme is the most practical that we propose. We also explore optimising the *authenticated* out-of-band channel construction, by considering pruning-based optimisations, where parties securely prune messages as soon as they are authenticated. We finally propose a new three-move authentication protocol and a corresponding security notion.

4.1.2 Technical Overview

We assume a network where parties communicate over two types of channels: insecure channels and out-of-band authenticated channels. The adversary has full control over insecure channels. In particular, it can read, deliver, modify and delay messages. Over the authenticated channels, the integrity and authenticity of the messages are protected, that is, the adversary can read, deliver, duplicate and delay messages but not modify them. In the Signal application, the insecure channel is the usual network, whereas the out-of-band channel is that which is used for safety number verification [Mar17], typically in-person.

(Authenticated) Ratcheted Communication. We introduce a syntax for ratcheted communication (RC) in which sent and received messages are associated with totally ordered *ordinals* (epoch/index pairs in formalisations of the Double Ratchet [ACD19]). Ordinals enable our protocol to support *immediate decryption* [ACD19], i.e., message loss and re-ordering on the network. We build on this syntax to define *authenticated ratcheted communication*, or ARC, which comprises two additional functions AuthSend and AuthReceive. A party can use AuthSend to send an authentication tag through the out-of-band channel that the counterpart processes with AuthReceive. AuthSend outputs an ordinal that is at least as large as the last *sent* ordinal for that party. AuthReceive, if successful, should authenticate all messages up to that ordinal; this is captured in UNF security. Our secondary correctness notion ORDINALS enforces these semantics even in presence of forgeries.

RID Security. We revisit the definitions of RECOVER security [DV19, CDV21] in the immediate decryption setting. We define two complementary security notions for RID security:

- r-RID ensures that the receiver of a forgery does not accept honest messages with ordinals *larger* than that of the forgery.
- s-RID security enables a party to detect if their counterpart has ever received a forgery (i.e., a forgery from the sender).

If a scheme is both r-RID- and s-RID-secure, then it is RID-secure. These notions are orthogonal to forward security and post-compromise security and we allow the adversary to have full access to each party's secret states and control their randomness.

We propose a construction that transforms any ratcheted communication scheme into a provably RID-secure one. In the construction, both parties locally keep track of messages they have sent and received. Every time they send a message, they attach *all* messages (i.e., the ciphertexts from the underlying RC) they have sent and received so far to their ciphertext. When a party receives a message that “contradicts” what it has sent or received, it can deduce that an active attack took place.

To reduce communication costs, parties send ordinals and hashes of messages, instead of

complete ciphertexts. For r-RID security, a receiver \bar{P} compares the input message and the supposed set of sent messages contained inside it with what \bar{P} has received previously. For s-RID, it suffices for a receiver \bar{P} (who knows exactly what it sent) to check whether the sender claims to have received anything that \bar{P} did not send. Here, P only needs to send a *single* hash alongside the set of received ordinals (which are generally smaller than hashes), since \bar{P} can recompute the hash locally. Since the channel is insecure, parties need to perform a series of checks on the ciphertexts to prevent the adversary from tampering with the sets of sent and received messages sent. Both r-RID and s-RID security rely on the collision resistance of the hash function.

UNF Security. We define notions analogous to r-RID and s-RID for *authenticated* ratcheted communication schemes. The r-UNF (receiver unforgeable) notion ensures that a party does not accept authentication tags after receiving a forgery, whereas s-UNF (sender unforgeable) ensures that a party does not accept authentication tags coming from a counterpart that received a forgery. We show that a RID-secure scheme can be turned into a UNF-secure scheme. The transformation highlights the similarity between RID and UNF security. In the former, parties authenticate all messages they have sent and received in band, whereas in the latter the messages are authenticated out of band. Concretely, the transformation uses the ciphertext of a RID-secure RC scheme as the authentication tag for an ARC scheme, which, intuitively, moves authentication material to the out-of-band channel. We also provide a direct construction that is analogous to our construction for RID security, except the performance hit here is only taken when sending and receiving authentication tags, rather than for every (in-band) message.

Lower Bounds. We prove a linear lower bound on the ciphertext size of any r-RID-secure RC: assuming unidirectional communication, each ciphertext must capture all information contained in previously sent ones. In fact, the security notion requires that the receiving party is able to immediately detect if *any* subset of previous ciphertexts contains a forgery, since (1) the sender does not know what has been received and (2) ciphertexts can be arbitrarily re-ordered or dropped.

For the proof, we construct an (inefficient) encoder/decoder pair for a list of input messages (m_1, \dots, m_n) and randomness R , that uses the r-RID RC to compress the input. More precisely, the encoder uses the RC to send messages (m_1, \dots, m_n) with randomness R to produce a list of ciphertexts (ct_1, \dots, ct_n) , and outputs only (ct_n, R) . The decoder uses the RC to receive ciphertext ct_n , generates every possible ct_1 for all possible messages m_1 with randomness R and attempts to successfully receive one of them. If this succeeds, it means m_1 was the same message as the one input to the encoder (i.e., the “honest” one). Then, the decoder continues with m_2 and so on, eventually outputting (m_1, \dots, m_n, R) . Finally, by setting the distributions of the messages and randomness as uniform¹, one can argue by Shannon’s theorem that the ciphertext space must be exponentially large in $n \cdot |m_i|$. The formal proof is actually more complicated as the r-RID security does not need to be perfect and many ciphertexts might be

¹A comparable argument can be made instead in terms of the *entropy* of the messages.

Chapter 4. On Active Attack Detection in Messaging with Immediate Decryption

successfully received when decoding.

Then, following a nearly identical proof, we prove a linear lower bound on the authentication tag size of any r -UNF-secure ARC. These proofs might be of independent interest.

Practical Active Attack Detection. We explore how to overcome the linear communication complexity that r -RID and r -UNF impose. We first observe that ciphertexts can be much smaller to achieve s -RID or s -UNF security. As noted above, it suffices for P to send a single hash of all received messages with the corresponding ordinals, since partner \bar{P} can recompute the hash if it stores all messages it sends. Assuming each ordinal uses c space, ciphertexts reduce in size from $O(n(\lambda + c))$ to $O(\lambda + nc)$ given P has received n messages.

We propose an optimised protocol that achieves s -RID security. In the protocol, parties keep track of epochs [ACD19]. Party A starts with $ep = 0$ and B with $ep = 1$. While sending each message parties attach the ep alongside. If $ep_A = t$, A does not accept any messages with $ep > t + 1$ and if A receives a message with $ep = t + 1$ they update their ep value $t + 2$. The observation is that epoch values only increase when both parties have received a message. Using this fact, it can be shown that it suffices to convey information about the messages received in the last *two epochs* to provide s -RID security. Thus, the number of ordinals required in each message depends on the communication pattern: in practice, it should usually suffice to only send a few ordinals and a hash digest achieve s -RID security. To see why the optimisation works, if an honest message was sent from A to B after A received a forged message, either this forgery was received in the last 2 epochs, or there was another forgery and honest message pair after it, as otherwise the ep values would be out of sync.

We observe that parties achieve r -RID/ r -UNF-like guarantees after one round of honest communication from s -RID/ s -UNF security. If P detects that their partner \bar{P} has received a forgery, P can let \bar{P} know, and thus \bar{P} can learn that they have received a forgery (which is what r -RID/ r -UNF guarantee). We also formalise this by proposing a lightweight three-move protocol over the out-of-band channel and a corresponding security model which captures bidirectional message authentication. Participant P (resp. \bar{P}) sends their set of received messages to their partner in the first and second moves. In the second and third moves, \bar{P} (resp. P) sends a bit that indicates whether the set of received messages was consistent with what they actually sent.

Furthermore, for UNF security, we note that the authentication tags can be compressed over time by including acknowledgements in tags. Since the out-of-band channel is authentic, parties are sure that the authentication information—that is, the sets of sent and received messages—coming from the counterpart is correct. This enables parties to prune already authenticated messages.

4.1.3 Additional Related Work

A growing line of work considers the performance and security of messaging in both the two-party [BS]⁺17, PR18, JS18, DV19, CDV21, BRV20 and more general group settings [ACDT20, ACDT21a, AJM22] settings. Some of these works provide similar [JS18] and sometimes weaker [JMM19a] guarantees for in-band active attack detection assuming in-order communication. To our knowledge, in-band active attack detection has not yet been explicitly explored in group messaging, but schemes like MLS ensure that if the state of two parties is forked then their states become incompatible, in some protocol-specific sense.

Naor et al. [NRS20] introduced the concept of immediate key delivery for round-based group key exchange: if some parties go offline, the remaining ones should be able to complete it and successfully output a shared secret. This property is related but distinct from immediate decryption in this work as it focuses on keys instead of messages.

Apart from Durak and Vaudenay and Caforio et al. who introduced the RECOVER notions, Dowling et al. [DHRR22] ensure *r*-RECOVER, but not *s*-RECOVER security via signatures, while providing anonymity guarantees even upon state exposure. Dowling et al. [DGP22] frame their authentication guarantees as follows: if no long-term keys are compromised, then all messages exchanged are authentic. Otherwise, active attacks can be detected out-of-band. They achieve this by signing all messages with long-term keys. Our protocols and security notions can be adapted to achieve these guarantees. In distributed computing, the problem is formalised in terms of accountability, which enables parties to detect faulty (Byzantine) nodes [HKD07, CGG⁺22]. In multi-party computation, a line of work has explored security with *identifiable abort* [IOZ14] which ensures that if parties fail to compute a given function, they can identify the party that caused the failure.

The encoder/decoder technique that we use to prove the lower bounds in Section 4.5 have been used before in cryptography [LN18, JLN19]. While the basic idea is the same, the technical details of the proofs are not comparable as the primitives are different. Related work in group messaging achieves communication lower bounds in symbolic models of execution [BDR20, ANPPP23] and in a black-box impossibility setting [BDG⁺22].

4.2 (Authenticated) Ratcheted Communication

In this section we introduce the *ratcheted communication* (RC) cryptographic primitive and an extension *authenticated ratcheted communication* (ARC) supporting out-of-band authentication. These primitives augment classic ratcheting-based secure messaging schemes [JMM19a, ACD19, CDV21] in two ways: (1) sent and received messages are associated with *ordinals*, and, for ARC, (2) the syntax encompasses two additional stateful algorithms AuthSend and AuthReceive.

Ordinals associated with messages enable a party to (1) order incoming messages, which is

Chapter 4. On Active Attack Detection in Messaging with Immediate Decryption

essential in the immediate decryption setting; (2) keep track of how many messages have been communicated; and (3) infer which messages have been authenticated using the out-of-band channel. Ordinals of the form num can be elements of any set on which a total order is defined. In Alwen et al.'s [ACD19] and Bienstock et al.'s [BFG⁺22a] modelling of Signal's Double Ratchet protocol, an ordinal num is defined as a pair of integers (e, c) such that $(e, c) < (e', c')$ if $e < e'$ or both $e = e'$ and $c < c'$. We formally define an RC scheme below.

Definition 33 (Ratcheted communication (RC)). A *ratcheted communication* (RC) scheme comprises the following efficient algorithms:

- $\text{pp} \xleftarrow{\$} \text{Setup}(1^\lambda)$: The setup algorithm takes as input the security parameter $\lambda \in \mathbb{N}$, expressed in unary, and outputs public parameters pp .
- $(\text{st}_A, \text{st}_B, z) \xleftarrow{\$} \text{Init}(\text{pp})$: The initialisation algorithm takes as input public parameters pp and outputs a state st_P for $P \in \{A, B\}$, and public information z .
- $(\text{st}'_P, \text{num}, \text{ct}) \xleftarrow{\$} \text{Send}(\text{st}_P, \text{ad}, \text{pt})$: The send algorithm takes as inputs state st_P , associated data ad and a plaintext pt and outputs a new state st'_P , an ordinal num and ciphertext ct .
- $(\text{acc}, \text{st}'_P, \text{num}, \text{pt}) \leftarrow \text{Receive}(\text{st}_P, \text{ad}, \text{ct})$: The receive algorithm takes as inputs a state st_P , associated data ad and ciphertext ct and outputs an acceptance bit $\text{acc} \in \{\text{true}, \text{false}\}$, state st'_P , ordinal num and plaintext pt .

The Receive algorithm returns dummy $\text{st}'_P, \text{num}, \text{pt}$ which are ignored when $\text{acc} = \text{false}$.

The Double Ratchet. Signal's Double Ratchet protocol can be viewed as an RC. In the work of Alwen et al. [ACD19], a secure messaging scheme consists of an initialisation algorithm and party-specific Send and Receive algorithms with no associated data. The Receive algorithms, but not the Send algorithms, output an epoch/index pair $(e, i) \in \mathbb{N}^2$ which plays the role of an ordinal. The Double Ratchet as modelled by Alwen et al. [ACD19] can thus be considered an RC by modifying its Send algorithm to output each (e, i) pair as an ordinal and enforcing that $\text{ad} = \perp$ is always input to Send and Receive.

Authenticated Ratcheted Communication (ARC). In an ARC, parties rely on stateful AuthSend and AuthReceive algorithms to authenticate the communication using a (possibly narrowband) out-of-band authenticated channel. AuthSend outputs an authentication tag and an ordinal, whereas AuthReceive takes an authentication tag as input and outputs an authentication bit and an ordinal. Intuitively, the authentication tag is sent via the out-of-band authenticated channel and it enables the receiver to detect active attacks using the AuthReceive algorithm. Participants can decide when to invoke the algorithms and thus use the authentication tag on-demand, e.g., when an out-of-band channel is available.

AuthSend and AuthReceive outputs ordinals with the same semantics as Send and Receive. Namely, the num that AuthSend outputs is greater or equal to the last num that Send outputs;

4.2 (Authenticated) Ratcheted Communication

besides ordering authentication tags with respect to messages the party has sent or received, the ordinal indicates which messages (all up until num) the authentication tag authenticates. Similarly, for AuthReceive, the ordinal num indicates that all messages with $\text{num}' \leq \text{num}$ have been authenticated with the received tag.

Definition 34 (Authenticated ratcheted communication (ARC)). An *authenticated ratcheted communication* (ARC) scheme comprises the following efficient algorithms:

- Setup, Init, Send, Receive are defined as in RC (Definition 33).
- $(\text{st}'_P, \text{num}, \text{at}) \stackrel{\$}{\leftarrow} \text{AuthSend}(\text{st}_P)$: The authenticated send function takes as input a state st_P and outputs a new state st'_P , an ordinal num and an authentication tag at.
- $(\text{auth}, \text{st}'_P, \text{num}) \leftarrow \text{AuthReceive}(\text{st}_P, \text{at})$: The authenticated receiving function takes as inputs state st_P and authentication tag at and outputs an authentication bit $\text{auth} \in \{\text{true}, \text{false}\}$, an updated state st'_P and an ordinal num.

The AuthReceive algorithm returns dummy st'_P, num which the scheme ignores when $\text{auth} = \text{false}$.

One could alternatively define AuthSend/AuthReceive to output sets of ordinals corresponding to which messages have been authenticated, rather than single ordinals. Our security notions ensure that this information can be efficiently computed by parties using the ordinals that the algorithms output.

Correctness. We define *correctness* for an RC and ARC scheme with the CORRECT game presented in Figure 4.1. The game takes a security parameter and a schedule sched as inputs. We use a schedule to model the message flow between the participants, which can (1) send a message, (2) receive a message, and for ARC only, (3) send an authentication tag, or (4) receive a sent authentication tag. More precisely, sched is an ordered list of instructions of the form $(P, \text{"send"}, \text{ad}, \text{pt})$, $(P, \text{"rec"}, j)$, and for ARC only, $(P, \text{"authsend"})$, or $(P, \text{"authrec"}, j)$, where $P \in \{A, B\}$, ad denotes associated data, pt denotes a plaintext, and $j \in \mathbb{N}$ indicates either the (ad, ct) pair or the at to be received—that is, to be processed by Receive or AuthReceive respectively.

A correct (A)RC scheme must recover the correct plaintext from the corresponding associated data/ciphertext pair. Moreover, the scheme must satisfy the following properties.

- Subsequent calls to the Send algorithm outputs strictly increasing ordinals (line 6 in Figure 4.1).²
- Ordinals are equal for corresponding calls to Send (resp. AuthSend for ARC) and Receive (resp. AuthReceive for ARC) (lines 11 and 21).

²This could instead require Send to output strictly increasing ordinals w.r.t. Send *and* Receive calls made by P, which is satisfied in Alwen et al.'s work [ACD19], but we opted against this for generality's sake.

Chapter 4. On Active Attack Detection in Messaging with Immediate Decryption

```

Game CORRECTRC(sched)
1: pp  $\xleftarrow{\$}$  Setup( $1^\lambda$ ); (stA, stB, z)  $\xleftarrow{\$}$  Init(pp)
2: ad*[·], pt*[·], ct*[·], at*[·]  $\leftarrow \perp$ ; received[·], sent[·]  $\leftarrow$  false; sent-num*  $\leftarrow \perp$ 
3: for  $i = 1$  to length(sched) :
4:   if sched[ $i$ ] parses as (P, "send", ad, pt) for some ad, pt, P  $\in \{A, B\}$  :
5:     (stP, num, ct)  $\xleftarrow{\$}$  Send(stP, ad, pt)
6:     if  $i > 1 \wedge \text{num} \leq \text{sent-num}_P$  : return 1
7:     sent[ $i$ ]  $\leftarrow$  true; adP[ $i$ ]  $\leftarrow$  ad; ptP[ $i$ ]  $\leftarrow$  (num, pt); ctP[ $i$ ]  $\leftarrow$  ct; sent-numP  $\leftarrow$  num
8:   elseif sched[ $i$ ] parses as (P, "rec",  $j$ ) for some  $j \in \mathbb{N}$ , P  $\in \{A, B\}$  :
9:     if  $\neg \text{sent}[j] \vee \text{received}[j] \vee \text{at}_{\bar{P}}[j] \neq \perp$  : continue
10:    (acc, st'P, num, pt)  $\leftarrow$  Receive(stP, ad $\bar{P}$ [ $j$ ], ct $\bar{P}$ [ $j$ ])
11:    if  $\neg \text{acc} \vee ((\text{num}, \text{pt}) \neq \text{pt}_{\bar{P}}[j])$  : return 1
12:    received[ $j$ ]  $\leftarrow$  acc; stP  $\leftarrow$  st'P
13:   elseif sched[ $i$ ] parses as (P, "authsend") for some P  $\in \{A, B\}$  :
14:     (stP, num, at)  $\xleftarrow{\$}$  AuthSend(stP)
15:     if num < sent-numP : return 1
16:     sent[ $i$ ]  $\leftarrow$  true; atP[ $i$ ]  $\leftarrow$  (num, at)
17:   elseif sched[ $i$ ] parses as (P, "authrec",  $j$ ) for some  $j \in \mathbb{N}$ , P  $\in \{A, B\}$  :
18:     if  $\neg \text{sent}[j] \vee \text{received}[j] \vee \text{at}_{\bar{P}}[j] = \perp$  : continue
19:     (num $\bar{P}$ , at $\bar{P}$ )  $\leftarrow$  at $\bar{P}$ [ $j$ ]
20:     (auth, st'P, num)  $\leftarrow$  AuthReceive(stP, at $\bar{P}$ )
21:     if  $\neg \text{auth} \vee \text{num} \neq \text{num}_{\bar{P}}$  : return 1
22:     received[ $j$ ]  $\leftarrow$  true; stP  $\leftarrow$  st'P
23: return 0

```

Figure 4.1: Correctness game for an RC/ARC scheme RC. Highlighted statements are only executed for an ARC scheme.

- For ARC, AuthSend outputs an ordinal greater or equal to the ordinal returned by the last call to Send (line 15).

We require that these properties hold even when forgeries are received by one or both parties, and enforce them in the ORDINALS game presented in Figure 4.3. We also capture these properties for clarity in our correctness game.

We formally define correctness for an (A)RC scheme in Definition 35 below.

Definition 35 (CORRECT). Consider the correctness game CORRECT presented in Figure 4.1. An RC (resp. ARC) scheme RC is *correct* if, for all $\lambda \in \mathbb{N}$, and all sequences of the form sched with elements of the form (P, "send", ad, pt), (P, "rec", j), (resp. also of the form (P, "authsend"),

4.2 (Authenticated) Ratcheted Communication

$(P, \text{"authrec"}, j)$), for $P \in \{A, B\}$, we have

$$\Pr[\text{CORRECT}_{\text{RC}}(\text{sched}) \Rightarrow 1] = 0 .$$

Correctness states that AuthSend must output an ordinal greater or equal to the ordinal that the last call to Send returned. If AuthSend does not increase the ordinal, then it is clear which messages are authenticated; if the ordinal increases in AuthSend, the application designer must keep track of the last num that Send returned to infer what the tag authenticates. Nonetheless, the latter case may be desirable to ensure that all ordinals output by Send and AuthSend are distinct.

Oracles. Our security notions for RC and ARC build on a common set of oracles, introduced in Figure 4.2. The SEND (resp. RECEIVE) oracle enables the adversary to send (resp. receive) a message on behalf of a party P . In SEND, the caller can specify the randomness used by Send or let the challenger sample randomness uniformly. For ARC, AUTHSEND enables the adversary to send an authentication tag on behalf of a party P , whereas AUTHRECEIVE handles AuthReceive. The oracles $\text{EXP}_{\text{pt}}(j)$ and $\text{EXP}_{\text{st}}(j)$ expose plaintexts and states to the adversary, respectively.

The oracles of Figure 4.2 model a communication network composed of insecure in-band and authentic out-of-band channels. The SEND and RECEIVE oracles enable the adversary to read, deliver, modify and delay messages, but AUTHSEND and AUTHRECEIVE do not allow the modification of authentication tags.

We assume an *always-authentic* out-of-band channel. To our knowledge, all deployed solutions for out-of-band authentication and relevant literature [DH20, DGP22] assume this. One could conceivably define a stronger model where the out-of-band channel is authentic only in some cases, e.g., the tampering rate is bounded, or multiple out-of-band channels exist but the adversary can compromise only a subset of them.

Ordinals. For RC and ARC schemes, we require, even in the presence of an adversary that injects forgeries, that Send and Receive (as well as AuthSend and AuthReceive for ARC schemes) output correct ordinals. We consider these properties in CORRECT (Figure 4.1), but they must hold also in presence of forgeries. We formalise this notion with the ORDINALS game in Figure 4.3.

In this game the challenger verifies three predicates, which correspond to the conditions for correct ordinals presented above. In Definition 36 we formalise ORDINALS security for (A)RC schemes.

Definition 36 (ORDINALS). Consider the ORDINALS game in Figure 4.3. We say that an (authenticated) ratcheted communication scheme RC is ORDINALS secure if, for all possibly

Chapter 4. On Active Attack Detection in Messaging with Immediate Decryption

<p>Oracle SEND(P, ad, pt, r)</p> <hr/> <pre> 1: $i \leftarrow i + 1$ 2: if $r = \varepsilon$: $r \xleftarrow{\\$} \mathcal{R}$ 3: $(st_P, num, ct) \leftarrow \text{Send}(st_P, ad, pt; r)$ 4: $state[i] \leftarrow st_P$ 5: $plaintext[i] \leftarrow pt$ 6: $log[i] \leftarrow ("send", P, num, ad, ct)$ 7: return (num, ct) </pre>	<p>Oracle RECEIVE(P, ad, ct)</p> <hr/> <pre> 1: $(acc, st, num, pt) \leftarrow \text{Receive}(st_P, ad, ct)$ 2: if $\neg acc$: return \perp 3: $i \leftarrow i + 1$ 4: $st_P \leftarrow st$; $state[i] \leftarrow st_P$ 5: $plaintext[i] \leftarrow pt$ 6: $log[i] \leftarrow ("rec", P, num, ad, ct)$ 7: return num </pre>
<p>Oracle AUTHSEND(P)</p> <hr/> <pre> 1: $i \leftarrow i + 1$ 2: $(st_P, num, at) \xleftarrow{\\$} \text{AuthSend}(st_P)$ 3: $auth[(P, i)] \leftarrow at$ 4: $state[i] \leftarrow st_P$ 5: $log[i] \leftarrow ("authsend", P, num, at)$ 6: return (num, at) </pre>	<p>Oracle AUTHRECEIVE(P, j)</p> <hr/> <pre> 1: $at \leftarrow auth[(\bar{P}, j)]$ 2: if $at = \perp$: return \perp 3: $(auth, st, num) \leftarrow \text{AuthReceive}(st_P, at)$ 4: if $\neg auth$: return \perp 5: $i \leftarrow i + 1$ 6: $st_P \leftarrow st$; $state[i] \leftarrow st_P$ 7: $log[i] \leftarrow ("authrec", P, num, at)$ 8: return num </pre>
<p>Oracle EXP_{pt}(j)</p> <hr/> <pre> 1: $i \leftarrow i + 1$ 2: $log[i] \leftarrow ("ptexp", j)$ 3: return $plaintext[j]$ </pre>	<p>Oracle EXP_{st}(j)</p> <hr/> <pre> 1: $i \leftarrow i + 1$ 2: $log[i] \leftarrow ("stexp", j)$ 3: return $state[j]$ </pre>

Figure 4.2: Oracles which use variables $state$, $plaintext$, log , $auth$, st_* and i , all initialised in games where the oracles are used. AUTHSEND and AUTHRECEIVE are only used when considering ARC.

unbounded adversaries \mathcal{A} we have

$$\Pr[\text{ORDINALS}_{\text{RC}}(\mathcal{A}) \Rightarrow 1] = 0.$$

The ORDINALS game in Figure 4.3 is not suited to the case where ordinals can be arbitrary and in particular collide between parties. Thus, a given party and their partner must be associated with disjoint ordinals: models of practical protocols like the Double Ratchet achieve this by associating one party with even epochs and the counterpart with odd epochs.

Game ORDINALS _{RC} (\mathcal{A})
1: $pp \leftarrow \text{Setup}(1^\lambda); (st_A, st_B, z) \leftarrow \text{Init}(pp)$ 2: $\text{state}[\cdot], \text{plaintext}[\cdot], \text{log}[\cdot], \text{auth}[\cdot], st_* \leftarrow \perp$ 3: $i \leftarrow 0$ 4: $\mathcal{A}^{\text{SEND, RECEIVE, EXP}_{pt}, \text{EXP}_{st}, \text{AUTHSEND, AUTHRECEIVE}}(pp, z)$ 5: if $\exists P, \text{num}, \text{num}', \text{ad}, \text{ct}, x, y$ s.t. 6: $\text{not-increasing}(\text{log}, P, \text{num}, \text{num}', x, y) \vee \text{different}(\text{log}, P, \text{num}, \text{num}', \text{ad}, \text{ct}, \text{at}) \vee$ 7: $\text{auth-monotonic}(\text{log}, P, \text{num}', y)$ 8: return 1 9: return 0
<hr/> different($\text{log}, P, \text{num}, \text{num}', \text{ad}, \text{ct}, \text{at}$) <hr/> 1: return $((\text{"send"}, P, \text{num}, \text{ad}, \text{ct}) \in \text{log} \wedge (\text{"rec"}, \bar{P}, \text{num}', \text{ad}, \text{ct}) \in \text{log}) \vee$ 2: $((\text{"authsend"}, P, \text{num}, \text{at}) \in \text{log} \wedge (\text{"authrec"}, \bar{P}, \text{num}', \text{at}) \in \text{log})) \wedge (\text{num} \neq \text{num}')$
<hr/> not-increasing($\text{log}, P, \text{num}, \text{num}', x, y$) <hr/> 1: return $((\text{"send"}, P, \text{num}, _, _) = \text{log}[x] \vee (\text{"rec"}, P, \text{num}, _, _) = \text{log}[x]) \wedge$ 2: $(\text{"send"}, P, \text{num}', _, _) = \text{log}[y] \wedge (0 < x < y) \wedge (\text{num} \geq \text{num}')$
<hr/> auth-monotonic($\text{log}, P, \text{num}', y$) <hr/> 1: $\text{num} \leftarrow \max\{\perp, \text{num}'' : (\text{"send"}, P, \text{num}'', _, _) = \text{log}[x] \wedge 0 < x < y\}$ 2: return $(\text{"authsend"}, P, \text{num}', _, _) = \text{log}[y] \wedge (\text{num} > \text{num}')$

Figure 4.3: ORDINALS game. Highlighted statements are only considered for an ARC.

4.3 In-Band Active Attack Detection: RID

In this section we consider *in-band* active attack detection in the immediate decryption setting.

Caforio et al. [CDV21] define RECOVER security, which encompasses both r-RECOVER security and s-RECOVER security, but their notions and constructions only support in-order message delivery. Intuitively, r-RECOVER security prevents a party from being able to deliver an honest message *after* delivering a forgery, and s-RECOVER security allows a party to detect and stop communication when their partner has delivered a forgery. We extend these notions to handle out-of-order message delivery by introducing r-RID and s-RID, which we present in Figure 4.4 and illustrate in Figure 4.5. Combined, these two properties ensure RID security. Note that these definitions are orthogonal to the usual forward and post-compromise security notions that the ratcheting literature considers [BSJ⁺17, ACD19].

The winning condition in RID consists of three predicates:

Chapter 4. On Active Attack Detection in Messaging with Immediate Decryption

- forgery verifies whether a forgery was accepted by one of the participants. In the predicate, we denote the impersonated party as P and the recipient of the forgery as \bar{P} .
- $\text{bad-}\bar{P}$ checks whether the recipient of the forgery manages to detect the attack. This predicate corresponds to r-RID security.
- bad-P establishes whether P , i.e., the participant that the adversary impersonates to send the forgery, fails to detect the attack. Since \bar{P} is the recipient of the forgery, the detection of the attack by P relies on a ciphertext sent by \bar{P} and honestly delivered. This predicate corresponds to s-RID security.

The game imposes that if forgery returns true, then at least one between bad-P and $\text{bad-}\bar{P}$ must return true for the adversary to win the game.

Definition 37 (RID). An RC \mathcal{RC} is r-RID (resp. s-RID/RID) if, for all efficient adversaries \mathcal{A} , we have:

$$\text{Adv}_{\mathcal{RC}}^{\text{r-rid}}(\mathcal{A}) = \Pr[\text{r-RID}_{\mathcal{RC}}(\mathcal{A}) \Rightarrow 1] = \text{negl}$$

$$\text{(resp. } \text{Adv}_{\mathcal{RC}}^{\text{s-rid}}(\mathcal{A}) = \Pr[\text{s-RID}_{\mathcal{RC}}(\mathcal{A}) \Rightarrow 1] / \text{Adv}_{\mathcal{RC}}^{\text{rid}}(\mathcal{A}) = \Pr[\text{RID}_{\mathcal{RC}}(\mathcal{A}) \Rightarrow 1] \text{)}$$

where game r-RID (resp. s-RID/RID) is defined in Figure 4.4.

Game $\text{r-RID}_{\mathcal{RC}}(\mathcal{A})$	$\text{s-RID}_{\mathcal{RC}}(\mathcal{A})$	Game $\text{RID}_{\mathcal{RC}}(\mathcal{A})$
1: $\text{pp} \xleftarrow{\$} \text{Setup}(1^\lambda); (\text{st}_A, \text{st}_B, z) \xleftarrow{\$} \text{Init}(\text{pp})$		1: return $\text{r-RID}_{\mathcal{RC}}(\mathcal{A}) \vee \text{s-RID}_{\mathcal{RC}}(\mathcal{A})$
2: $\text{state}[\cdot], \text{plaintext}[\cdot], \text{log}[\cdot] \leftarrow \perp$		$\text{forgery}(\text{log}, P, \text{num}, \text{ad}, \text{ct}, x)$
3: $\text{auth}[\cdot], \text{st}_* \leftarrow \perp$		
4: $i \leftarrow 0$		1: return ("send", P , num, ad, ct) $\notin \text{log} \wedge$
5: $\mathcal{A}^{\mathcal{O}}(\text{pp}, z)$		2: ("rec", \bar{P} , num, ad, ct) = log[x]
6: if $\exists P, \text{num}, \text{num}', \text{ad}, \text{ct}, \text{ad}', \text{ct}', x, y$ s.t.		$\text{bad-}\bar{P}(\text{log}, P, \text{num}, \text{num}', \text{ad}', \text{ct}')$
7: $\text{forgery}(\text{log}, P, \text{num}, \text{ad}, \text{ct}, x) \wedge$		
8: $\text{bad-}\bar{P}(\text{log}, P, \text{num}, \text{num}', \text{ad}', \text{ct}') :$		1: return ("rec", \bar{P} , num', ad', ct') $\in \text{log} \wedge$
9: $\text{bad-P}(\text{log}, P, \text{num}', \text{ad}', \text{ct}', x, y) :$		2: ("send", P , num', ad', ct') $\in \text{log} \wedge$
10: return 1		3: (num < num')
11: return 0		$\text{bad-P}(\text{log}, P, \text{num}', \text{ad}', \text{ct}', x, y)$
		1: return (y > x) \wedge
		2: ("send", \bar{P} , num', ad', ct') = log[y] \wedge
		3: ("rec", P , num', ad', ct') $\in \text{log}$

Figure 4.4: r-RID, s-RID and RID games for $\mathcal{O} = \{\text{SEND}, \text{RECEIVE}, \text{EXP}_{\text{pt}}, \text{EXP}_{\text{st}}\}$.

Although r-RID may seem stronger than s-RID at first glance, the two notions are not comparable. There exist schemes which provide r-RID and not s-RID security and vice versa, e.g.,

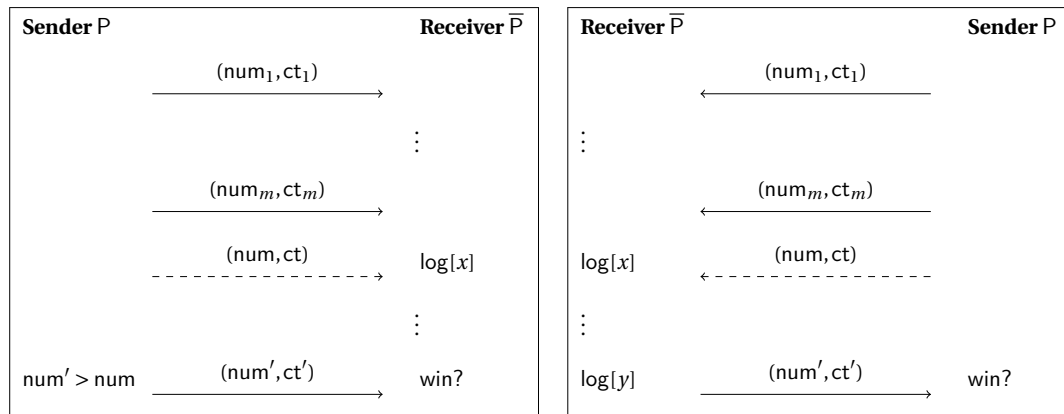


Figure 4.5: Visualising r-RID (left) and s-RID (right). Each figure showcases an adversary’s winning condition in the respective game. The dashed arrows are forged messages. If \bar{P} accepts the message at time “win?” then the adversary wins.

the scheme proposed in Figure 4.6 if the checks for either r-RID or s-RID are removed from the checks subroutine given the underlying RC is not r-RID or s-RID secure, respectively (the Double Ratchet is neither, for example).

However, we note the following link between the two notions. Suppose we use an s-RID scheme. This means that \bar{P} is able to detect that \bar{P} received a forged message. Then, if P sends an “abort” message to \bar{P} , \bar{P} would be able to detect the forgery after *one* honest round-trip of messages. In other words, s-RID RC schemes can be transformed (by adding an “abort” message) into RC schemes with a weak variant of r-RID security: r-RID after a honest round-trip.

On Fine-Grained Security. Suppose A sends 5 messages with $\text{num} \in \{1, \dots, 5\}$, B receives a forgery with $\text{num} = 1000$, and then A sends 5 messages with $\text{num} \in \{6, \dots, 10\}$. If B never sends, i.e., A is the sender and B the receiver, RID-security only guarantees that the forgery might be detected when A sends the honest message with $\text{num}' = 1001$ (observe the condition “ $\text{num} < \text{num}'$ ” in predicate $\text{bad-}\bar{P}$ in Figure 4.4). Intuitively, B should be able to detect the forgery on receipt of the honest message with $\text{num} = 6$ since this message is “independent” of the forgery with $\text{num} = 1000$. By the not-increasing predicate of the ORDINALS security, all messages that A sends after one round-trip will have $\text{num} > 1000$, so such an attack will nevertheless be eventually detected. Fine-grained security capturing these scenarios can be formalised by tracking state exposures and message delivery timing at the cost of greater definitional complexity; we leave it open to do so. Some forgeries will be defeated by our construction below but it is likely required to leverage the security of the underlying RC to build a scheme providing this kind of security. Looking ahead, this remark also applies to UNF ARC schemes defined in Section 4.4.

4.3.1 RID-Secure RC

In this subsection we build a RID-secure RC scheme that relies on an ORDINALS-secure RC scheme and a collision-resistant hash function H (Definition 9). We present our transformation in Figure 4.6.

Scheme Description. Each party P keeps track of every message it has sent and received (in S and R , respectively). This information is communicated to \bar{P} every time P calls `Send` (via variables S and R').

The `Send` procedure prepares the set R' , which contains the ordinals and a hash of all received messages (line 3). This step can be optimised by using an incremental hash function as we discuss in Section 4.6.1. Next, it calls `RC.Send` with input (ad', pt) where $ad' = (ad, S, R')$ is the associated data. The ciphertext ct contains both ct' and sets S and R' . Finally, it adds the pair (num, h) to S (line 8), where the hash h is computed as $H.Eval(st_P.hk, (ad, ct))$, where $ct = (ct', S, R')$. Intuitively, (num, h) acts as a summary of P 's state after calling `RC.Send` which can be checked by \bar{P} for inconsistency.

When \bar{P} invokes `Receive`, the procedure calls `RC.Receive`, which outputs $num \neq \perp$ if the call is successful. Since ct contains R' , \bar{P} checks that what P received so far was correct (line 3 in checks). In addition, using the S set contained in the ciphertext ct , \bar{P} can further check whether the ciphertexts it received so far have indeed been sent by P . This is verified from lines 5 to 18 of checks. Some checks detect tampering of ct by the adversary (e.g. $ct.S$ should not contain ordinals larger than the one of the current ciphertext, or if ct was sent earlier than another ciphertext already received, $ct.S$ should be consistent with messages already acknowledged, etc.). If everything verifies, `Receive` stores (num, h) in R and adds $ct.S$ to the set of acknowledged messages (lines 9 and 10).

Associated Data. The sets S and R' included in the ciphertext are also included in the authenticated data passed to the underlying RC. This is actually not needed for RID security, but for authentication and confidentiality. Although we do not define these notions here, this should be done in practice and remains as future work to formalise.

On Errors. Note that our scheme outputs a generic error symbol \perp in all cases. In particular, our construction outputs the same symbol regardless of whether the error was due to detecting an active attack, or from the situation where the adversary did not expose any states and simply sent a malformed ciphertext. The latter situation entails a denial of service attack vector if errors are treated the same way in both cases, so in practice (and in future work) they should be differentiated between.

Security Analysis. Correctness of RC_{RID} follows from the correctness of the underlying RC scheme RC and the fact that the checks always outputs false when only honest messages are received. Similarly, ORDINALS-security follows from the ORDINALS security of RC , as RC_{RID} outputs the num that RC outputs. As the next theorems state, the construction of Figure 4.6

$RC_{RID}.Setup(1^\lambda)$ <ol style="list-style-type: none"> 1: $pp_0 \xleftarrow{\\$} RC.Setup(1^\lambda)$ 2: $hk \xleftarrow{\\$} H.KGen(1^\lambda)$ 3: $hk' \xleftarrow{\\$} H.KGen(1^\lambda)$ 4: $pp \leftarrow (pp_0, hk, hk')$ 5: return pp <hr/> $RC_{RID}.Send(st_p, ad, pt)$ <ol style="list-style-type: none"> 1: $(st'_p, hk, hk', S, R, _, _) \leftarrow st_p$ 2: $nums' \leftarrow \{num' : (num', _) \in R\}$ 3: $R' \leftarrow (nums', H.Eval(hk', R))$ 4: $ad' \leftarrow (ad, S, R')$ 5: $(st_p.st'_p, num, ct') \xleftarrow{\\$} RC.Send(st'_p, ad', pt)$ 6: $ct \leftarrow (ct', S, R')$ 7: $h \leftarrow H.Eval(hk, (num, ad, ct))$ 8: $st_p.S \leftarrow S \cup \{(num, h)\}$ 9: return (st_p, num, ct) <hr/> $RC_{RID}.Receive(st_p, ad, ct)$ <ol style="list-style-type: none"> 1: $(ct', S^{\bar{P}}, R^{\bar{P}}) \leftarrow ct$ 2: $(st'_p, hk, _, _, R, S_{ack}, _) \leftarrow st_p$ 3: $ad' \leftarrow (ad, S^{\bar{P}}, R^{\bar{P}})$ 4: $(acc, st'_p, num, pt) \leftarrow RC.Receive(st'_p, ad', ct')$ 5: if $\neg acc$: return $(false, st_p, \perp, \perp)$ 6: $h \leftarrow H.Eval(hk, (num, ad, ct))$ 7: if $checks(st_p, ct, h, num)$: 8: return $(false, st_p, \perp, \perp)$ 9: $st_p.R \leftarrow R \cup \{(num, h)\}$ 10: $st_p.S_{ack} \leftarrow S_{ack} \cup S^{\bar{P}}$ 11: $st_p.st'_p \leftarrow st'_p$ 12: return (acc, st_p, num, pt) 	$RC_{RID}.Init(pp)$ <ol style="list-style-type: none"> 1: $(pp_0, hk, hk') \leftarrow pp$ 2: $(st'_A, st'_B, z') \leftarrow RC.Init(pp_0)$ 3: $max\text{-}num \leftarrow 0$ 4: $S, R, S_{ack} \leftarrow \emptyset$ 5: $st_A \leftarrow (st'_A, hk, hk', S, R, S_{ack}, max\text{-}num)$ 6: $st_B \leftarrow (st'_B, hk, hk', S, R, S_{ack}, max\text{-}num)$ 7: $z \leftarrow (z', pp)$ 8: return (st_A, st_B, z) <hr/> $checks(st_p, ct, h, num)$ <ol style="list-style-type: none"> 1: $(nums', h') \leftarrow ct.R$ 2: $R^* \leftarrow \{(num', _) \in st_p.S : num' \in nums'\}$ 3: $s\text{-}bool \leftarrow (H.Eval(st_p.hk', R^*) \neq h')$ 4: $R' \leftarrow \{(num', _) \in st_p.R : num' \leq num\}$ 5: $r\text{-}bool \leftarrow (R' \not\subseteq ct.S)$ 6: $r\text{-}bool \leftarrow r\text{-}bool \vee$ 7: $\quad (\exists (num^*, _) \in ct.S : num^* \geq num)$ 8: if $num < st_p.max\text{-}num$: 9: $r\text{-}bool \leftarrow r\text{-}bool \vee ((num, h) \notin st.S_{ack})$ 10: $r\text{-}bool \leftarrow r\text{-}bool \vee (ct.S \not\subseteq st.S_{ack})$ 11: $S_{ack}' \leftarrow \{(num', _) \in st_p.S_{ack} :$ 12: $\quad num' < num\}$ 13: $r\text{-}bool \leftarrow r\text{-}bool \vee (S_{ack}' \not\subseteq ct.S)$ 14: else : 15: $st_p.max\text{-}num \leftarrow num$ 16: $r\text{-}bool \leftarrow r\text{-}bool \vee$ 17: $\quad (\exists (num', _) \in st.S_{ack} \setminus ct.S :$ 18: $\quad num' < st_p.max\text{-}num)$ 19: return $r\text{-}bool \vee s\text{-}bool$
----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

Figure 4.6: RID-secure RC scheme RC_{RID} based on a RC scheme RC (Definition 33) and a hash function H (Definition 8). RC_{RID} requires the following variables: $max\text{-}num$ represents the largest received num ; S is the set of (num, h) pairs; R is the set of received (num, h) pairs; S_{ack} is the set of (num, h) which are expected to be received (according to the received ciphertext ct). All sets are append-only.

is r -RID-secure (Theorem 7) and s -RID-secure (Theorem 8). The construction is therefore RID-secure.

Chapter 4. On Active Attack Detection in Messaging with Immediate Decryption

Theorem 7. Consider collision resistant hash function H used to build RC_{RID} (defined in Figure 4.6). Then, we have that for every efficient adversary \mathcal{A} , one can build an adversary \mathcal{B} such that

$$\text{Adv}_{\text{RC}_{\text{RID}}}^{\text{r-rid}}(\mathcal{A}) \leq \text{Adv}_H^{\text{cr}}(\mathcal{B}).$$

Proof. Let us assume there exists an adversary $\tilde{\mathcal{A}}$ playing the r-RID game, running in time \tilde{t} and making at most \tilde{q} queries. Let us call the advantage of this adversary $\tilde{\epsilon}$, hence we have

$$\Pr[\text{r-RID}_{\text{RC}_{\text{RID}}}(\tilde{\mathcal{A}}) \Rightarrow 1] = \tilde{\epsilon}.$$

Let E be an event that occurs when $\text{r-RID}_{\text{RC}_{\text{RID}}}(\tilde{\mathcal{A}})$ outputs 1. The proof strategy is to construct an adversary \mathcal{A}^* , running in time $\approx \tilde{t}$ such that

$$\Pr[\text{CR}_H(\mathcal{A}^*) \Rightarrow 1 \mid E] = 1.$$

By definition of r-RID, E occurring means that there exist $P, (\text{num}, \text{ad}, \text{ct}), (\text{num}', \text{ad}', \text{ct}'), x$ such that $\text{bad-}\bar{P}(\log, P, \text{num}, \text{num}', \text{ad}', \text{ct}')$ and $\text{forgery}(\log, P, \text{num}, \text{ad}, \text{ct}, x)$ are both true. This means that the message with ordinal num was not sent by P but received at some point by \bar{P} ($\log[x] = (\text{"rec"}, \bar{P}, \text{num}, \text{ad}, \text{ct})$), and the message with ordinal num' was also received and was actually sent by P . Moreover, $\text{num} < \text{num}'$.

We separate the two cases where 1) the message with ordinal num (the forged message) is received before the message with ordinal num' (the honest message), and 2) the message with ordinal num' is received first. We first analyse the former case.

We argue that unless the adversary found a collision, the message with ordinal num' (the honest message) would not have been delivered. Suppose that the honest message *was* delivered. Let st_P be the state of the receiver P while receiving message num' , and $\text{st}_{\bar{P}}$ be the state of the sender \bar{P} while sending the message $(\text{num}', \text{ad}', \text{ct}')$. As $(\text{"rec"}, P, \text{num}', \text{ad}', \text{ct}') \in \log$, it means $\text{RC}_{\text{RID}}.\text{Receive}(\text{st}_P, \text{ad}', \text{ct}') \rightarrow (\text{true}, \text{st}'_P, \text{num}', \text{pt}')$, which implies $\text{checks}(\text{st}_P, \text{ct}', \text{num}', H.\text{Eval}(\text{hk}, (\text{num}', \text{ad}', \text{ct}')))$ returned false.

Note that as $\text{num} \leq \text{num}'$, we have

$$(\text{num}, H.\text{Eval}(\text{hk}, (\text{num}, \text{ad}, \text{ct}))) \in R' \text{ and } (\text{num}, H.\text{Eval}(\text{hk}, (\text{num}, \text{ad}, \text{ct}))) \in \text{ct}' . S,$$

as otherwise r-bool would have been set to true in line 5. Let $h_f := H.\text{Eval}(\text{hk}, (\text{num}, \text{ad}, \text{ct}))$. As $(\text{num}, h_f) \in \text{ct}' . S$, we have

$$(\text{num}, h_f) \in \text{st}_{\bar{P}} . S \tag{4.1}$$

as ct' was sent by \bar{P} . This would mean that \bar{P} did send a message with ordinal num , let us call

it $(\text{num}, \text{ad}_h, \text{ct}_h)$. Hence, we have that,

$$(\text{H.Eval}(\text{hk}, (\text{ad}_h, \text{ct}_h, \text{num})), \text{num}) \in \text{st}_{\bar{P}}.S \quad (4.2)$$

By combining (4.1) and (4.2) and the fact that num can appear only once in $\text{st}_{\bar{P}}$ (due to the ORDINALS security of RC_{RID}), we get that $\text{H.Eval}(\text{hk}, (\text{num}, \text{ad}_h, \text{ct}_h)) = \text{H.Eval}(\text{hk}, (\text{num}, \text{ad}, \text{ct}))$ which gives a collision. This is because $(\text{ad}_h, \text{ct}_h) \neq (\text{ad}, \text{ct})$ as $(\text{"send"}, \bar{P}, \text{num}, \text{ad}_h, \text{ct}_h) \in \text{log}$ and $(\text{"send"}, \bar{P}, \text{num}, \text{ad}, \text{ct}) \notin \text{log}$.

Now we discuss the case where $(\text{num}', \text{ad}', \text{ct}')$ is received before $(\text{num}, \text{ad}, \text{ct})$. As $\text{num} \leq \text{num}'$, this would mean that while receiving $(\text{num}, \text{ad}, \text{ct})$, $\text{max-num} \geq \text{num}' \geq \text{num}$. This would mean $(\text{num}, h) \in \text{st}_P.S_{\text{ack}}$, otherwise the condition on line 9 would have not been satisfied. As S_{ack} is only updated by adding the elements in $S^{\bar{P}}$ when a message is received, and as (num, h_f) is not in $\text{ct}^*.S$, for any honest ct^* , there should exist a forged message $(\text{num}'', \text{ad}'', \text{ct}'')$ received before $(\text{ad}, \text{ct}, \text{num})$ such that $(\text{num}, h_f) \in \text{ct}'' .S$. As we considered $(\text{num}, \text{num}')$ to be the first pair of messages violating the r-RID property, we know that $\text{num}'' > \text{num}' > \text{num}$.

We split the two cases where $(\text{num}'', \text{ad}'', \text{ct}'')$ is received before $(\text{num}', \text{ad}', \text{ct}')$ and the case where it is received after $(\text{num}', \text{ad}', \text{ct}')$. Let us consider the first case. We argue in this case $(\text{num}', \text{ad}', \text{ct}')$ (the honest message) would not be accepted. As num'' is received before num' , $\text{num}' < \text{num}'' \leq \text{max-num}$. And as $\text{r-bool} = \text{false}$, $S'_{\text{ack}} \subseteq \text{ct}' .S$ (line 12). However $(\text{num}, h_f) \in S'_{\text{ack}}$, as it was in $\text{ct}'' .S$, hence it should also be in $\text{ct}' .S$, which would mean $h_f = \text{H.Eval}(\text{hk}, \text{ad}_h, \text{ct}_h, \text{num})$ which is again a collision.

Now let us consider the case where $(\text{num}'', \text{ad}'', \text{ct}'')$ is received after the message $(\text{num}', \text{ad}', \text{ct}')$. We argue that $(\text{num}'', \text{ad}'', \text{ct}'')$ should not have been accepted. We split the cases where $\text{num}'' \geq \text{max-num}$ and $\text{num}'' < \text{max-num}$. Let us consider the later first. As $(\text{ad}'', \text{ct}'', \text{num}'')$ was accepted, and hence $\text{r-bool} = \text{false}$, $\text{ct}'' .S \subset S_{\text{ack}}$ (line 10). Now $(\text{num}, h_f) \in \text{ct}'' .S$, so $(\text{num}, h_f) \in S_{\text{ack}}$. As without loss of generality we can imagine $(\text{num}'', \text{ad}'', \text{ct}'')$ being the first message vouching for (h_f, num) , this would mean (num, h_f) was added to S_{ack} by an honest message, i.e. $h_f = h_h$ which leads to a collision again.

Finally for the case in which $\text{num}'' \geq \text{max-num}$, again, considering that $(\text{num}'', \text{ad}'', \text{ct}'')$ is the first message vouching for (h_f, num) , we have $(\text{num}, h_f) \in S_{\text{ack}} \setminus \text{ct}'' .S$ (and so r-bool would be set to true) unless $h_f = h_h$. Moreover, at this point $(\text{num}', \text{ad}', \text{ct}')$ has already been received so, $\text{max-num} \geq \text{num}' > \text{num}$. Hence, unless $h_f = h_h$, r-bool would be set to true in line 7. This concludes the proof that $(\text{num}', \text{ad}', \text{ct}')$, $(\text{num}, \text{ad}, \text{ct})$ are accepted if and only if $\text{H.Eval}(\text{hk}, (\text{num}, \text{ad}_h, \text{ct}_h)) = \text{H.Eval}(\text{hk}, (\text{num}, \text{ad}, \text{ct}))$.

Now we describe the CR adversary \mathcal{A}^* . \mathcal{A}^* runs the initialisation of RC_{RID} by replacing the sampling step of hk with the hk given by the CR_H game, then runs $\tilde{\mathcal{A}}$ as a subroutine, and computes $(\text{"rec"}, P, \text{num}, \text{ad}_f, \text{ct}_f) \in \text{log}$, and $(\text{"send"}, \bar{P}, \text{num}, \text{ad}_h, \text{ct}_h) \in \text{log}$ such that $(\text{ad}_f, \text{ct}_f) \neq (\text{ad}_h, \text{ct}_h)$ and $h_f = h_h$ if possible. Given E , this pair always exists as we have $\Pr[\text{CR}_H(\mathcal{A}^*) \Rightarrow 1 \mid E] = 1$. Moreover, as \mathcal{A}^* is just running $\tilde{\mathcal{A}}$ as a subroutine and not doing

Chapter 4. On Active Attack Detection in Messaging with Immediate Decryption

anything extra, the time it runs is also $\approx \tilde{t}$. Finally, we have

$$\begin{aligned} \Pr[\text{CR}_H(\mathcal{A}^*) \Rightarrow 1] &\geq \Pr[\text{CR}_H(\mathcal{A}^*) \Rightarrow 1 \mid E] \cdot \Pr[E] \\ &= \Pr[\text{s-RID}_{\text{RC}_{\text{RID}}}(\tilde{\mathcal{A}}) \Rightarrow 1] = \tilde{\epsilon}. \end{aligned}$$

Hence, $\tilde{\epsilon} \leq \Pr[\text{CR}_H^{\mathcal{A}^*}(1^\lambda) \Rightarrow 1]$. □

Theorem 8. Consider collision resistant hash function H used to build RC_{RID} (defined in Figure 4.6). Then, we have that for every efficient adversary \mathcal{A} , one can build an adversary \mathcal{B} such that

$$\text{Adv}_{\text{RC}_{\text{RID}}}^{\text{s-rid}}(\mathcal{A}) \leq \text{Adv}_H^{\text{cr}}(\mathcal{B}).$$

Proof. The proof strategy is essentially the same as the one taken for the proof of Theorem 7. For any adversary $\tilde{\mathcal{A}}$ playing the s-RID game, we construct an adversary \mathcal{A}^* playing the CR game with comparable complexity. We first describe the adversary \mathcal{A}^* in terms of $\tilde{\mathcal{A}}$ and proceed by proving that \mathcal{A}^* wins at least as often as $\tilde{\mathcal{A}}$.

As with the previous proof we define an event E that occurs only when $\text{s-RID}_{\text{RC}_{\text{RID}}}(\tilde{\mathcal{A}}) \Rightarrow 1$, and we prove that $\Pr[\text{CR}_H(\mathcal{A}^*) \Rightarrow 1 \mid E] = 1$.

The event $\text{s-RID}_{\text{RC}_{\text{RID}}}(\tilde{\mathcal{A}}) \Rightarrow 1$, means there exists $P, \text{num}, \text{ad}, \text{ct}, \text{num}', \text{ad}', \text{ct}', x, y$ such that $x < y$, $(\text{num}, \text{pt}, \text{ct})$ is a forged message received by \bar{P} and logged in $\log[x]$, and $(\text{num}', \text{ad}', \text{ct}')$ is an honest message sent by \bar{P} (logged in $\log[y]$) and received by P . As $(\text{num}', \text{ad}', \text{ct}')$ was received, P 's checks call returned false.

Let us define $h_f = H.\text{Eval}(\text{hk}, (\text{num}, \text{ad}, \text{ct}))$. When receiving the forged message, \bar{P} adds (num, h_f) to $\text{st}_{\bar{P}}.R$. As $y > x$, (num, h_f) is in $\text{st}_{\bar{P}}.R$ when \bar{P} sends. Hence $\text{num} \in \text{num}'$ for the honest message $(\text{num}', \text{ad}', \text{ct}')$ sent by \bar{P} . Now as $(\text{num}', \text{ad}', \text{ct}')$ was accepted, we have, due to line 3 of checks, that $H.\text{Eval}(\text{hk}', R^*) = H.\text{Eval}(\text{hk}', \text{st}_{\bar{P}}.R)$.

If $R^* \neq \text{st}_{\bar{P}}.R$ we have already found a collision. So let us assume that $R^* = \text{st}_{\bar{P}}.R$. Now as $(\text{num}, h_f) \in \text{st}_{\bar{P}}.R$, we also have that $(\text{num}, h_f) \in R^* \subset \text{st}_P.S$.

This would mean that there exists an honest message $(\text{num}, \text{ad}_h, \text{ct}_h)$ such that $H.\text{Eval}(\text{hk}, (\text{num}, \text{ad}_h, \text{ct}_h)) = h_f$. But note that as $(\text{num}, \text{ad}_h, \text{ct}_h)$ is an honest message, $(\text{"send"}, P, \text{num}, \text{ad}_h, \text{ct}_h) \in \log$ but $(\text{"send"}, P, \text{num}, \text{ad}, \text{ct}) \notin \log$ as the message was forged, hence $(\text{num}, \text{ad}_h, \text{ct}_h) \neq (\text{num}, \text{ad}, \text{ct})$, which again yields a collision pair.

Now the CR adversary \mathcal{A}^* does the following: they run the s-RID adversary $\tilde{\mathcal{A}}$ as a subroutine with the hk given by the CR challenger. They later find $P, \text{num}, \text{ad}, \text{ct}, \text{num}', \text{ad}', \text{ct}', x, y$ satisfying the condition, in case they exist. Now (for example) by exposing the states of the parties once $(\text{num}', \text{ad}', \text{ct}')$ was sent by \bar{P} they can find the collision pair described above. Hence we

have,

$$\Pr[\text{CR}_H(\mathcal{A}^*) \Rightarrow 1] \geq \Pr[\text{CR}_H(\mathcal{A}^*) \Rightarrow 1 \mid E] \cdot \Pr[E] = \Pr[\text{s-RID}_{\text{RC}_{\text{RID}}}(\tilde{\mathcal{A}}) \Rightarrow 1] \quad (4.3)$$

Observe that the run-time of \mathcal{A}^* is roughly the run-time of $\tilde{\mathcal{A}}$. This concludes the proof. \square

Optimisation. The s-RID notion imposes less overhead than r-RID: the construction can be further optimised and still provide s-RID security. We describe an optimisation in Section 4.6.2 that allows for the set of ordinals in \mathcal{R} of each party to be pruned each round trip of communication, which particularly improves communication complexity when parties are well-synchronised.

4.4 Out-Of-Band Active Attack Detection: UNF

In-band active attack detection is not always possible, as an adversary may block all honest messages sent by one or more parties. For example, modern messaging solutions in practice use a (possibly malicious) third party server to relay messages between participants, thereby introducing a single point of failure for in-band communication. This motivates us to consider out-of-band active attack detection, whereby parties can exchange authentication tags out-of-band, and to define *unforgeable* security (UNF).

An ARC scheme is *unforgeable* if, as soon as one of the two parties accepts a forgery, both parties can detect this out-of-band. We formalise this security notion through the UNF game (Figure 4.7), which, similarly to RID, encompasses r-UNF and s-UNF. The winning condition in UNF consists of three predicates: *forgery*, *bad- \bar{P}* (corresponding to r-UNF) and *bad-P* (corresponding to s-UNF) that are essentially the same as the predicates that we use to define RID security (Definition 37), except they rely on authentication tags instead of ciphertexts for forgery detection (and thus active attack detection is performed on-demand).

Definition 38 (UNF). An ARC ARC is r-UNF (resp. s-UNF/UNF) if, for all efficient adversaries \mathcal{A} , we have:

$$\begin{aligned} \text{Adv}_{\text{ARC}}^{\text{r-unf}}(\mathcal{A}) &= \Pr[\text{r-UNF}_{\text{ARC}}(\mathcal{A}) \Rightarrow 1] = \text{negl} \\ (\text{resp. } \text{Adv}_{\text{ARC}}^{\text{s-unf}}(\mathcal{A}) &= \Pr[\text{s-UNF}_{\text{ARC}}(\mathcal{A}) \Rightarrow 1] / \text{Adv}_{\text{ARC}}^{\text{unf}}(\mathcal{A}) = \Pr[\text{UNF}_{\text{ARC}}(\mathcal{A}) \Rightarrow 1]) \end{aligned}$$

where game r-UNF (resp. s-UNF/UNF) is defined in Figure 4.7.

As for RC schemes, we do not define message indistinguishability [DV19, ACD19] for ARC schemes. In our constructions, the schemes include in the authentication tag only *public* material, i.e., messages that have already been sent over the insecure channel. Since the adversary already has access to the entire transcript of the insecure channel, the authentication material should not give any additional advantage in a message indistinguishability game.

Chapter 4. On Active Attack Detection in Messaging with Immediate Decryption

Game $r\text{-UNF}_{\text{ARC}}(\mathcal{A})$	$s\text{-UNF}_{\text{ARC}}(\mathcal{A})$	Game $\text{UNF}_{\text{ARC}}(\mathcal{A})$
1: $\text{pp} \xleftarrow{\$} \text{Setup}(1^\lambda)$; $(\text{st}_A, \text{st}_B, z) \xleftarrow{\$} \text{Init}(\text{pp})$		1: return $r\text{-UNF}_{\text{ARC}}(\mathcal{A}) \vee s\text{-UNF}_{\text{ARC}}(\mathcal{A})$
2: $\text{state}[\cdot], \text{plaintext}[\cdot], \text{log}[\cdot] \leftarrow \perp$		$\text{forgery}(\text{log}, P, \text{num}, \text{ad}, \text{ct}, x)$
3: $\text{auth}[\cdot], \text{st}_* \leftarrow \perp$		1: return ("send", \bar{P} , num, ad, ct) $\notin \text{log} \wedge$
4: $i \leftarrow 0$		2: ("rec", \bar{P} , num, ad, ct) = log[x]
5: $\mathcal{A}^{\mathcal{O}}(\text{pp}, z)$		$\text{bad-}\bar{P}(\text{log}, P, \text{num}, \text{num}', \text{at})$
6: if $\exists P, \text{num}, \text{num}', \text{ad}, \text{ct}, \text{at}, x, y$ s.t.		1: return ("authrec", \bar{P} , num', at) $\in \text{log} \wedge$
7: $\text{forgery}(\text{log}, P, \text{num}, \text{ad}, \text{ct}, x) \wedge$		2: (num \leq num')
8: $\text{bad-}\bar{P}(\text{log}, P, \text{num}, \text{num}', \text{at})$:		$\text{bad-P}(\text{log}, P, \text{num}', \text{at}, x, y)$
9: $\text{bad-P}(\text{log}, P, \text{num}', \text{at}, x, y)$:		1: return (y > x) \wedge
10: return 1		2: ("authsend", \bar{P} , num', at) = log[y] \wedge
11: return 0		3: ("authrec", P, num', at) $\in \text{log}$

Figure 4.7: $r\text{-UNF}$, $s\text{-UNF}$ and UNF games for $\mathcal{O} = \{\text{SEND}, \text{RECEIVE}, \text{EXP}_{\text{pt}}, \text{EXP}_{\text{st}}, \text{AUTHSEND}, \text{AUTHRECEIVE}\}$.

4.4.1 UNF-Secure ARC from a RID-Secure RC

Highlighting the similarity between RID security and UNF security, we show in this subsection that one can directly use a RID-secure RC to build a UNF-secure ARC. The ARC scheme uses the Setup, Gen, Init, Send, Receive function of the RC. To send an authentication tag with AuthSend, the ARC scheme calls the Send function on a dummy message to generate a ciphertext ct that acts as the authentication tag. The function AuthReceive is then implemented as a Receive call on the authentication tag/ciphertext. The construction is detailed in Figure 4.8.

Then, we can show the following theorem, which also implies that the scheme of Figure 4.8 is $r\text{-UNF}$ - and $s\text{-UNF}$ -secure.

Theorem 9. Let RC_{RID} be a RC scheme and ARC_{UNF} be the ARC scheme built out of RC as shown in Figure 4.8. If RC_{RID} is RID, ORDINALS-secure and correct, then ARC_{UNF} is UNF-, ORDINALS-secure and correct.

Proof. Correctness follows from the correctness of the underlying RC_{RID} and the use of domain separation for tags and ciphertexts.

We sketch a proof showing RID security of RC_{RID} implies UNF security of ARC_{UNF} . For any adversary \mathcal{A} playing the UNF game with ARC_{UNF} , we build a RID adversary \mathcal{B} for RC_{RID} . Each query made by \mathcal{A} to the oracles SEND, RECEIVE, EXP_{pt} , EXP_{st} are forwarded by \mathcal{B} to its own corresponding oracles (and domain separation is correctly implemented where needed).

$\text{ARC}_{\text{UNF}}.\text{Setup}(1^\lambda)$ 1: return $\text{RC}_{\text{RID}}.\text{Setup}(1^\lambda)$	$\text{ARC}_{\text{UNF}}.\text{Receive}(\text{st}_P, \text{ad}, \text{ct})$ 1: $(b, \text{ct}') \leftarrow \text{ct}$ 2: if $b \neq 0$: return $(\text{false}, \perp, \perp, \perp)$ 3: return $\text{RC}_{\text{RID}}.\text{Receive}(\text{st}_P, \text{ad}, \text{ct}')$
$\text{ARC}_{\text{UNF}}.\text{Init}(\text{pp})$ 1: return $\text{RC}_{\text{RID}}.\text{Init}(\text{pp})$	$\text{ARC}_{\text{UNF}}.\text{AuthSend}(\text{st}_P)$ 1: $(\text{st}'_P, \text{num}, \text{ct}) \leftarrow \text{RC}_{\text{RID}}.\text{Send}(\text{st}_P, 0, 0)$ 2: return $(\text{st}'_P, \text{num}, (1, \text{ct}))$
$\text{ARC}_{\text{UNF}}.\text{Send}(\text{st}_P, \text{ad}, \text{pt})$ 1: $\text{ct}' \leftarrow \text{RC}_{\text{RID}}.\text{Send}(\text{st}_P, \text{ad}, \text{pt})$ 2: $\text{ct} \leftarrow (0, \text{ct}')$ 3: return ct	$\text{ARC}_{\text{UNF}}.\text{AuthReceive}(\text{st}_P, \text{at})$ 1: $(b, \text{at}') \leftarrow \text{at}$ 2: if $b \neq 1$: return $(\text{false}, \perp, \perp)$ 3: $(\text{acc}, \text{st}'_P, \text{num}, \text{pt}) \leftarrow \text{RC}_{\text{RID}}.\text{Receive}(\text{st}_P, 0, \text{at}')$ 4: return $(\text{acc}, \text{st}'_P, \text{num})$

Figure 4.8: UNF-secure ARC scheme ARC_{UNF} based on a RID-secure RC scheme RC_{RID} .

Queries of the form $\text{AUTHSEND}(P)$ are simulated by \mathcal{B} querying “ $\text{at}' \leftarrow \text{SEND}(P, 0, 0)$ ” and setting “ $\text{at} \leftarrow (1, \text{at}')$ ”, which perfectly simulates the generation of a tag in ARC_{UNF} . Finally, AUTHRECEIVE queries are simulated using the RECEIVE oracle on the tag/ciphertext. \mathcal{B} can perfectly simulate the UNF game for \mathcal{A} .

Now, let us assume that the UNF adversary \mathcal{A} wins with the forgery and bad- P predicates both evaluating to true. It means a forgery was received by a party P , then, later, that party sent a tag (i.e. a ciphertext in the RID game played by \mathcal{B}) that is honestly and successfully delivered to a party \bar{P} . That implies that in the RID game played by \mathcal{B} , a party received a forgery, then sent a message that was successfully delivered, which is a winning condition for \mathcal{B} .

The second case is when the UNF adversary \mathcal{A} wins with the forgery and bad- \bar{P} predicates both evaluating to true. This means that a forgery was received by a party P with ordinal num, then a tag with ordinal $\text{num}' \geq \text{num}$ was successfully received by \bar{P} . Note that in our ARC_{UNF} construction the tags are ciphertexts, thus the ordinals are strictly increasing, i.e., $\text{num}' > \text{num}$. Therefore, in the RID game played by \mathcal{B} , a forgery with ordinal num was received by P , then later a honest ciphertext with ordinal $\text{num}' > \text{num}$ was successfully delivered to P , making the bad- \bar{P} predicate in the RID game true.

Hence, for any adversary \mathcal{A} winning the UNF game, there exists a RID adversary \mathcal{B} that wins with at least the same probability.

Finally, ORDINALS -security follows from the ORDINALS security of RC_{RID} and the fact that $\text{ARC}_{\text{UNF}}.\text{Send}$ and $\text{ARC}_{\text{UNF}}.\text{AuthSend}$ calls directly output num from $\text{RC}_{\text{RID}}.\text{Send}$. \square

Chapter 4. On Active Attack Detection in Messaging with Immediate Decryption

4.4.2 UNF-Secure ARC from Any RC

We present an unoptimised UNF-secure ARC scheme ARC_{base} given a RC scheme (Definition 33). We present our scheme in Figure 4.9. We later describe a communication-optimised scheme in Section 4.6.3.

<p>$\text{ARC}_{\text{base}}.\text{Setup}(1^\lambda)$</p> <hr/> <p>1: $\text{pp}_0 \xleftarrow{\\$} \text{RC.Setup}(1^\lambda)$; $\text{hk} \xleftarrow{\\$} \text{H.KGen}(1^\lambda)$ 2: return (pp_0, hk)</p> <p>$\text{ARC}_{\text{base}}.\text{Init}(\text{pp})$</p> <hr/> <p>1: $(\text{pp}_0, \text{hk}) \leftarrow \text{pp}$ 2: $(\text{st}'_A, \text{st}'_B, z') \xleftarrow{\\$} \text{RC.Init}(\text{pp}_0)$ 3: $\text{num}, \text{max-num} \leftarrow \perp$; $S, R, S_{\text{ack}} \leftarrow \emptyset$ 4: $\text{st}_A \leftarrow (\text{st}'_A, \text{hk}, S, R, S_{\text{ack}}, \text{num}, \text{max-num})$ 5: $\text{st}_B \leftarrow (\text{st}'_B, \text{hk}, S, R, S_{\text{ack}}, \text{num}, \text{max-num})$ 6: $z \leftarrow (z', \text{pp})$ 7: return $(\text{st}_A, \text{st}_B, z)$</p> <p>$\text{ARC}_{\text{base}}.\text{Send}(\text{st}_P, \text{ad}, \text{pt})$</p> <hr/> <p>1: $(\text{st}'_P, \text{hk}, S, _, _, _) \leftarrow \text{st}_P$ 2: $(\text{st}_P.\text{st}'_P, \text{num}, \text{ct}) \xleftarrow{\\$} \text{RC.Send}(\text{st}'_P, \text{ad}, \text{pt})$ 3: $h \leftarrow \text{H.Eval}(\text{hk}, (\text{ad}, \text{ct}))$ 4: $\text{st}_P.S \leftarrow S \cup \{(\text{num}, h)\}$ 5: $\text{st}_P.\text{num} \leftarrow \text{num}$ 6: return $(\text{st}_P, \text{num}, \text{ct})$</p> <p>$\text{ARC}_{\text{base}}.\text{AuthSend}(\text{st}_P)$</p> <hr/> <p>1: $(_, _, S, R, _, \text{num}, _) \leftarrow \text{st}_P$ 2: $\text{at} \leftarrow (S, R, \text{num})$ 3: return $(\text{st}_P, \text{num}, \text{at})$</p>	<p>$\text{ARC}_{\text{base}}.\text{Receive}(\text{st}_P, \text{ad}, \text{ct})$</p> <hr/> <p>1: $(\text{st}'_P, \text{hk}, _, R, S_{\text{ack}}, _, \text{max-num}) \leftarrow \text{st}_P$ 2: $(\text{acc}, \text{st}'_P, \text{num}, \text{pt}) \leftarrow \text{RC.Receive}(\text{st}'_P, \text{ad}, \text{ct})$ 3: if $\neg \text{acc}$: return $(\text{false}, \text{st}_P, \perp, \perp)$ 4: $h \leftarrow \text{H.Eval}(\text{hk}, (\text{ad}, \text{ct}))$ 5: if $\text{num} \leq \text{max-num} \wedge (\text{num}, h) \notin S_{\text{ack}}$: 6: return $(\text{false}, \text{st}_P, \perp, \perp)$ 7: $\text{st}_P.R \leftarrow R \cup \{(\text{num}, h)\}$ 8: $\text{st}_P.\text{st}'_P \leftarrow \text{st}'_P$ 9: return $(\text{acc}, \text{st}_P, \text{num}, \text{pt})$</p> <p>$\text{ARC}_{\text{base}}.\text{AuthReceive}(\text{st}_P, \text{at})$</p> <hr/> <p>1: $(_, _, S, R, S_{\text{ack}}, \text{num}, \text{max-num}) \leftarrow \text{st}_P$ 2: $(S^{\bar{P}}, R^{\bar{P}}, \text{num}^{\bar{P}}) \leftarrow \text{at}$ 3: // \bar{P} received a forgery 4: if $R^{\bar{P}} \not\subseteq S$: return $(\text{false}, \text{st}_P, \text{num})$ 5: $R_{\subseteq}^P \leftarrow \{(\text{num}, _) \in R : \text{num} \leq \text{num}^{\bar{P}}\}$ 6: // P received a forgery 7: if $R_{\subseteq}^P \not\subseteq S^{\bar{P}}$: return $(\text{false}, \text{st}_P, \text{num})$ 8: $\text{st}_P.S_{\text{ack}} \leftarrow S_{\text{ack}} \cup S^{\bar{P}}$ 9: $\text{st}_P.\text{max-num} \leftarrow \max\{\text{max-num}, \text{num}^{\bar{P}}\}$ 10: return $(\text{true}, \text{st}_P, \text{num}^{\bar{P}})$</p>
---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

Figure 4.9: UNF-secure ARC scheme ARC_{base} based on a RC scheme RC (Definition 33). The scheme uses the following additional variables compared to RC: S is the set of sent values stored as (num, h) ; R is the set of received values stored as (num, h) ; S_{ack} is the set of (num, h) expected to be received (according to the received authentication tag at); num is the ordinal corresponding to the party's last *sent* message; max-num represents the largest num received in an at . For simplicity of exposition, we omit the optimisation where R is sent as a single hash and n ordinals as done in Figure 4.6 for RID security.

Scheme Description. Procedures Send and Receive make use of the respective RC procedures. Send stores the hash of (ad, ct) for the message being sent, together with the num that the

underlying RC.Send algorithm outputs. A tuple composed of num and this hash is stored in a set S , which is in turn stored by the calling party. Send also updates ordinal num in the caller's state. The Receive procedure verifies if the RC.Receive algorithm accepts the inputs and that the received message is not a forgery on a previously authenticated message, which, as we will describe, is contained in S_{ack} . If both checks pass, Receive stores the hash of (ad, ct) together with the ordinal num returned by RC.Receive in a set R .

AuthSend includes in the authentication tag (at) the hashes of the caller's sent and received messages together with the last ordinal num returned by RC.Send. Since the adversary can reorder messages both in the normal channel and in the out-of-band channel, num indicates to the recipient of the authentication tag which messages they should compare against at. AuthReceive parses the authentication tag and checks whether the messages received by the counterpart are in the caller's local set S . Then, it verifies whether its local set of received messages, excluding the messages not encompassed by at, is a subset of the messages sent by the counterpart. If one of these conditions is not satisfied, then a forgery is detected. The sent messages authenticated by the counterpart are stored in a set S_{ack} . ARC_{base} .Receive uses this set to avoid forgeries on already authenticated ordinals.

The size of the authentication tags and the state of each party in the scheme of Figure 4.9 is linear in the number of sent and received messages. We show in Section 4.6.3 that this can be reduced by pruning. Messages can nevertheless be efficiently exchanged out-of-band in practice, e.g., using Bluetooth. Otherwise, parties can send authentication information over the insecure channel and authenticate it using the out-of-band channel by hashing and comparing digests [PV06].

Security Analysis. We now analyze the security properties of the scheme in Figure 4.9. Correctness of the scheme follows from the correctness of the underlying RC scheme. Similarly, ORDINALS security follows from the ORDINALS security of RC, as the scheme of Figure 4.9 outputs the same num that RC outputs.

The UNF-security of ARC_{base} (Figure 4.9) is derived, as before, from the collision resistance of the hash function that the scheme uses, which we prove in what follows.

Theorem 10 (UNF security of ARC_{base}). Consider collision resistant hash function H used to build ARC_{base} (defined in Figure 4.9). Then, we have that for every efficient adversary \mathcal{A} , one can build an adversary \mathcal{B} such that

$$\text{Adv}_{\text{ARC}_{\text{base}}}^{\text{unf}}(\mathcal{A}) \leq \text{Adv}_H^{\text{cr}}(\mathcal{B}).$$

Proof. Assume an adversary \mathcal{A} playing the UNF game (Figure 4.7), which makes at most q oracle queries and runs in time at most t . We assume the advantage of \mathcal{A} is ϵ , hence by Definition 38 we have $\Pr[\text{UNF}_{\text{ARC}_{\text{base}}}(\mathcal{A}) \Rightarrow 1] = \epsilon$. We construct an adversary \mathcal{B} , running in time approximately equal to t , which, running \mathcal{A} as a subroutine, wins the collision resistance game for H (Definition 9), that is $\Pr[\text{CR}_H(\mathcal{B}^{\mathcal{A}}) \Rightarrow 1] = 1$.

Chapter 4. On Active Attack Detection in Messaging with Immediate Decryption

The UNF adversary \mathcal{A} wins when one party accepts a forgery (predicate forgery) and at least one of the two parties fails to detect the forgery (predicates bad-P and bad- \bar{P}). Suppose there exist $P, \text{num}, \text{num}', \text{ad}, \text{ct}, \text{at}, x, y$ such that $\text{forgery}(\log, P, \text{num}, \text{ad}, \text{ct}, x) = \text{true}$. We analyze the predicates bad-P and bad- \bar{P} separately, starting with the latter.

The forgery predicate states that $(\text{"send"}, P, \text{num}, \text{ad}', \text{ct}') \in \log$ and $(\text{"rec"}, \bar{P}, \text{num}, \text{ad}, \text{ct}) = \log[x]$ for some $x \in \mathbb{N}$, where $(\text{ad}', \text{ct}') \neq (\text{ad}, \text{ct})$. This means that $(\text{num}, \text{H.Eval}(\text{hk}, (\text{ad}', \text{ct}')) \in S^P$ and $(\text{num}, \text{H.Eval}(\text{hk}, (\text{ad}, \text{ct}))) \in R^{\bar{P}}$, otherwise the forgery is trivially detected because $(\text{num}, _) \notin R^{\bar{P}}$. Moreover, by the bad- \bar{P} predicate we know that $R_{\leq}^{\bar{P}} \subseteq S^P$ for any $\text{num} \leq \text{num}'$, which implies that $(\text{num}, \text{H.Eval}(\text{hk}, (\text{ad}', \text{ct}')) = (\text{num}, \text{H.Eval}(\text{hk}, (\text{ad}, \text{ct})))$. By ORDINALS security, (num, \cdot) can appear only once in S^P , respectively in $R^{\bar{P}}$, and by assumption $(\text{ad}', \text{ct}') \neq (\text{ad}, \text{ct})$, therefore we have a collision for $\text{H.Eval}(\text{hk}, \cdot)$.

We now analyze the bad-P predicate. The forgery predicate states that $(\text{"send"}, P, \text{num}, \text{ad}', \text{ct}') \in \log$ and $(\text{"rec"}, \bar{P}, \text{num}, \text{ad}, \text{ct}) = \log[x]$ for some $x \in \mathbb{N}$, where $(\text{ad}', \text{ct}') \neq (\text{ad}, \text{ct})$, otherwise the forgery is trivially detected. This implies that $(\text{num}, \text{H.Eval}(\text{hk}, (\text{ad}, \text{ct}))) \in R^{\bar{P}}$ when $\text{ARC}_{\text{base}}.\text{Receive}(_, \text{ad}, \text{ct}) \rightarrow (\text{true}, _, \text{num}, _)$. By the bad-P predicate we know that \bar{P} sends an authentication tag at after accepting (ad, ct) , since $(\text{"authsend"}, \bar{P}, \text{num}', \text{at}) = \log[y]$ and $y > x$, which means that $(\text{num}, \text{H.Eval}(\text{hk}, (\text{ad}, \text{ct})))$ is in the $R^{\bar{P}}$ that at contains. The rest of the argument follows the same approach as the previous paragraph.

$\mathcal{B}^{\mathcal{A}}(\text{hk})$
1: $(1^\lambda, \text{hk}_0) \leftarrow \text{hk}; \text{pp} \xleftarrow{\$} \text{ARC}_{\text{base}}.\text{Setup}(1^\lambda); (\text{pp}_0, \text{hk}') \leftarrow \text{pp}; \text{pp}' \leftarrow (\text{pp}_0, \text{hk})$
2: $(\text{st}_A, \text{st}_B, z) \xleftarrow{\$} \text{ARC}_{\text{base}}.\text{Init}(\text{pp}')$
3: $\text{state}[\cdot], \text{plaintext}[\cdot], \log[\cdot], \text{auth}[\cdot], \text{st}_* \leftarrow \perp; i \leftarrow 0$
4: $\mathcal{A}^{\mathcal{O}}(\text{pp}, z)$
5: if $\exists \text{num}, P, \text{ad}, \text{ct}, \text{ad}', \text{ct}' : (\text{"send"}, P, \text{num}, \text{ad}, \text{ct}) \in \log \wedge$
6: $(\text{"rec"}, \bar{P}, \text{num}, \text{ad}', \text{ct}') \wedge (\text{ad}, \text{ct}) \neq (\text{ad}', \text{ct}') :$
7: return $(\text{ad}, \text{ct}), (\text{ad}', \text{ct}')$
8: else abort

Figure 4.10: CR adversary \mathcal{B} where $\mathcal{O} = \{\text{SEND}, \text{RECEIVE}, \text{AUTHSEND}, \text{AUTHRECEIVE}, \text{EXP}_{\text{pt}}, \text{EXP}_{\text{st}}\}$ for the proof of Theorem 10.

We describe in Figure 4.10 the adversary \mathcal{B} which plays the collision resistance game. \mathcal{B} runs the $\text{ARC}_{\text{base}}.\text{Setup}$ procedure and replaces the hash key hk' that the procedure returns with the hk that the adversary receives from the CR challenger. After running \mathcal{A} as a subroutine, \mathcal{B} analyzes the log array to find the $(\text{ad}, \text{ct}), (\text{ad}', \text{ct}')$ pairs that represents a forgery and returns those. If \mathcal{A} wins the UNF game, then \mathcal{B} wins the CR game, that is

$$\Pr[\text{CR}_H(\mathcal{B}) \Rightarrow 1] \geq \Pr[\text{UNF}_{\text{ARC}_{\text{base}}}(\mathcal{A}) \Rightarrow 1] = \epsilon.$$

Moreover, \mathcal{B} runs \mathcal{A} as a subroutine and executes a negligible additional amount of work. \square

4.5 Lower Bounds for Active Attack Detection

We study in this section the size of (1) ciphertexts of any r -RID-secure RC and (2) authentication tags of any r -UNF-secure ARC. In particular, all of our constructions so far (and hereafter) achieving one of these properties incurs (at least in the worst case) linear growth in the ciphertext or tag length in terms of the number of messages sent. We show here that one cannot hope for better by proving two lower bounds. More precisely, we show that the ciphertext space (resp. tag space) of a r -RID RC (resp. r -UNF ARC) grows exponentially in the number of messages sent. Note that we cannot prove a lower bound on the ciphertext size directly as it is always possible that *some* ciphertext is small. However, our bounds imply that at least n bits are required to represent any ciphertext or tag in their respective domain after the n -th message.

4.5.1 Communication Cost for r -RID Security

In what follows, we consider a RC that is perfectly correct:³ for all randomness r , valid states st_P and associative data ad , the function $\text{Send}(st_P, ad, \cdot; r)$ mapping a plaintext to a ciphertext is injective.

The next theorem proves that the ciphertext size in a r -RID RC grows linearly in the number of messages sent (multiplied by either the security parameter or message size).

Theorem 11. Let Π be a perfectly correct RC, n_s and λ be fixed, and T_{λ, n_s} be the time complexity of the (efficient) adversary given on the left of Figure 4.12. In addition, let $\gamma \in \mathbb{Z}$ be such that, for all adversaries \mathcal{A} running in at most time T_{λ, n_s} which query oracle SEND at most n_s times, we have

$$\Pr[r\text{-RID}_{\Pi}(\mathcal{A}) \Rightarrow 1] \leq \frac{1}{2^\gamma}.$$

Let $\mathcal{M} = \{0, 1\}^n$ and $\mathcal{C} = \{0, 1\}^k$ (without loss of generality) be the plaintext and ciphertext space associated to Π , respectively. Then,

$$\begin{aligned} k &\geq n + (n_s - 1)(\gamma - 2), \text{ if } \gamma \leq n \\ k &\geq 2 + n_s(n - 2), \text{ if } \gamma > n. \end{aligned}$$

A third lower bound gives

$$k \geq nn_s - \frac{1}{1 - \frac{2^n n_s}{2^\gamma}},$$

which is tighter for low values of n (e.g. $n = 1, 2$) and when $\gamma > n + \log(n_s)$.

³We discuss how to relax this requirement directly after our proof.

Chapter 4. On Active Attack Detection in Messaging with Immediate Decryption

Proof. We show that if k is smaller than the given bounds, one can build an encoder/decoder for a uniform source such that the expected bit-length of a codeword is strictly lower than the entropy (i.e., the log of the size of the sampling set), contradicting Shannon's source coding theorem [Sha48].

More formally, we consider a source that samples uniformly at random from the set $\{0,1\}^{n \times n_s} \times \{0,1\}^r$, where r is the maximal number of bits (i.e., random coins) needed by the two procedures Setup and Init of Π and n_s invocations of Send. We present an encoder and decoder for such a source in Figure 4.11 (the non-boxed instructions in the encoder, and the decoder shown on the left). For the sake of explanation, assume that the RC used in the encoder/decoder has perfect r -RID security. Then, the sender sends n_s honestly generated ciphertexts ct_1, \dots, ct_{n_s} , and the receiver receives the last ciphertext ct_{n_s} . By perfect r -RID security, for any $i < n_s$, any ciphertext $ct'_i \neq ct_i$ should be rejected by the receiver. Thus, one can build an (inefficient) decoder that tests all ct'_i to find the correct one and recovers the corresponding message. In a sense, all ct_i must be encoded in the last ciphertext ct_{n_s} . The actual encoding is more complicated as if the r -RID security is not perfect, there will be a number of false positives (i.e., $ct'_i \neq ct_i$ but ct'_i is accepted by the receiver). Note that w.l.o.g., we omit the associated data throughout the proof (or assume $ad = \perp$) as it plays no role.

Lemma 12. Our encoder (Figure 4.11) is perfectly correct, i.e.,

$$\Pr[\text{Decode}(\text{Encode}(m_1, \dots, m_{n_s}, R)) = (m_1, \dots, m_{n_s}, R)] = 1.$$

Proof. The value R output by Decode is the same as the one input in Encode. Since the initial states depend only on R and Π is correct, ct_{n_s} will decrypt to m_{n_s} . The states st_A^i will be identical in both the encoding and decoding procedures as they are generated from st_A^{i-1} , the previously recovered message m_{i-1} and randomness R_i . This implies that the sets of accepting messages S_i will be the same as they depend only on st_B^1 and st_A^{i-1} . In addition, by the perfect correctness of Π , each message m_i will be in the corresponding set S_i . Hence, the decoder can recover each message m_i by reading S_i at the index given in the input. \square

Lemma 13. Let C be the random variable corresponding to the codeword length output by Encode. In addition, let $F_i := S_i \setminus \{m_i\}$ be the set of false positives, where S_i and m_i are as in Encode. Then, $\mathbb{E}[C] \leq k + r + \sum_{i=1}^{n_s-1} 1 + \log(\mathbb{E}[|F_i|] + 1)$.

Proof. By construction, the encoder outputs a codeword of $k + r + \sum_{i=1}^{n_s-1} \lceil \log(|S_i|) \rceil$ bits. Therefore, we have

$$\mathbb{E}[C] = k + r + \sum_{i=1}^{n_s-1} \mathbb{E}[\lceil \log(1 + |F_i|) \rceil] \leq k + r + \sum_{i=1}^{n_s-1} 1 + \mathbb{E}[\log(1 + |F_i|)]$$

which is upper bounded by $k + r + \sum_{i=1}^{n_s-1} (1 + \log(\mathbb{E}[|F_i|] + 1))$, by the linearity of expectation and the definition of F_i , the fact that $\lceil x \rceil \leq 1 + x$, and Jensen's inequality applied consecutively. \square

4.5 Lower Bounds for Active Attack Detection

Encode(m_1, \dots, m_{n_s}, R)	Decode(b, data, R)
<pre> 1: parse ($R_{-1}, R_0, \dots, R_{n_s}$) $\leftarrow R$; pp \leftarrow Setup($1^\lambda; R_{-1}$); ($\text{st}_A^0, \text{st}_B^0, z$) \leftarrow Init(pp; R_0) 2: for $i \in \{1, \dots, n_s\}$ do // send the n_s messages 3: ($\text{st}_A^i, \text{num}, \text{ct}_i$) \leftarrow Send($\text{st}_A^{i-1}, m_i; R_i$) 4: ($\text{acc}, \text{st}_B^1, \text{num}, m'_{n_s}$) \leftarrow Receive($\text{st}_B^0, \text{ct}_{n_s}$) // Receive ct_{n_s}: $m'_{n_s} = m_{n_s}$ by perfect correctness 5: // Collecting false positives and correct messages: 6: for $i \in \{1, \dots, n_s - 1\}$ do 7: $S_i \leftarrow \emptyset$ 8: for $m \in \{0, 1\}^n$ do 9: ($_, _, \text{ct}'$) \leftarrow Send($\text{st}_A^{i-1}, m; R_i$) 10: ($\text{acc}, _, _, m'$) \leftarrow Receive($\text{st}_B^1, \text{ct}'$) 11: if acc: 12: if $m \neq m_i$: return ($0, m_1, \dots, m_{n_s}, R$) 13: $S_i \leftarrow S_i \cup \{m\}$ 14: $L_i \leftarrow$ sort(S_i) 15: $e_i \leftarrow$ index of m_i in L_i (in binary with $\lceil \log(L_i) \rceil$ bits) 16: encode ct_{n_s} with k bits 17: return ($1, \text{ct}_{n_s}, R$) 18: return ($\text{ct}_{n_s}, R, e_0 \parallel \dots \parallel e_{n_s-1}$) </pre>	<pre> 1: if $b = 0$: 2: (m_1, \dots, m_{n_s}) \leftarrow data 3: return (m_1, \dots, m_{n_s}, R) 4: else $\text{ct}_{n_s} \leftarrow$ data 5: parse ($R_{-1}, R_0, \dots, R_{n_s}$) $\leftarrow R$ 6: pp \leftarrow Setup($1^\lambda; R_{-1}$); ($\text{st}_A^0, \text{st}_B^0, z$) \leftarrow Init(pp; R_0) 7: ($\text{acc}, \text{st}_B^1, \text{num}, m_{n_s}$) \leftarrow Receive($\text{st}_B^0, \text{ct}_{n_s}$) 8: // Collecting false positives: 9: for $i \in \{1, \dots, n_s - 1\}$ do 10: $S_i \leftarrow \emptyset$ 11: for $m \in \{0, 1\}^n$ do 12: ($_, _, \text{ct}'$) \leftarrow Send($\text{st}_A^{i-1}, m; R_i$) 13: ($\text{acc}, _, _, m'$) \leftarrow Receive($\text{st}_B^1, \text{ct}'$) 14: if acc: $m_i \leftarrow m$ 15: ($\text{st}_A^i, _, _$) \leftarrow Send($\text{st}_A^{i-1}, m_i; R_i$) 16: return ($m_1, \dots, m_{n_s}, R$) </pre>

Figure 4.11: Encoder without (resp. with) boxed instructions and decoder on the left (resp. right) for proving the first 2 (resp. third) lower bound(s) in Theorem 11.

Chapter 4. On Active Attack Detection in Messaging with Immediate Decryption

Finally, we show the following key lemma.

Lemma 14. Let F_i , $i \in [n_s - 1]$ be defined as above and n, γ as in the statement of the Theorem (Theorem 11). Then, $\mathbb{E}[|F_i|] \leq 2^{n-\gamma}$.

Proof. We proceed by contradiction. That is, we show that if $\mathbb{E}[|F_i|] > 2^{n-\gamma}$, then there exists an adversary \mathcal{A}_i s.t. $\Pr[r\text{-RID}_\Pi(\mathcal{A}_i) \Rightarrow 1] > \frac{1}{2^\gamma}$.

\mathcal{A}_i	\mathcal{A}_i
1: $m_1, \dots, m_{n_s} \xleftarrow{\$} \{0, 1\}^{n \times n_s}$	1: $m_1, \dots, m_{n_s} \xleftarrow{\$} \{0, 1\}^{n \times n_s}$
2: $R_{-1}, R_0, \dots, R_{n_s} \xleftarrow{\$} \{0, 1\}^r$	2: $R_{-1}, R_0, \dots, R_{n_s} \xleftarrow{\$} \{0, 1\}^r$
3: for $j \in \{1, \dots, n_s\}$ do	3: for $j \in \{1, \dots, n_s\}$ do
4: if $j = i$:	4: if $j = i$:
5: $\text{st}_A^{i-1} \leftarrow \text{EXP}_{\text{st}}(A)$	5: $\text{st}_A^{i-1} \leftarrow \text{EXP}_{\text{st}}(A)$
6: $(\text{ct}_j, \text{num}_j) \leftarrow \text{SEND}(A, \emptyset, m_j, R_j)$	6: $(\text{ct}_j, \text{num}_j) \leftarrow \text{SEND}(A, \emptyset, m_j, R_j)$
7: $\text{RECEIVE}(B, \emptyset, \text{ct}_{n_s})$	7: $(\text{num}_{\text{at}}, \text{at}) \leftarrow \text{AUTHSEND}(A)$
8: $m \xleftarrow{\$} \{0, 1\}^n$	8: $i_{\text{at}} \leftarrow \text{index of at}$
9: $_, _, \text{ct} \leftarrow \text{Send}(\text{st}_A^{i-1}, m; R_i)$	9: $\text{AUTHRECEIVE}(B, i_{\text{at}})$
10: $\text{RECEIVE}(B, \emptyset, \text{ct})$	10: $m \xleftarrow{\$} \{0, 1\}^n$
11: return B	11: $_, _, \text{ct} \leftarrow \text{Send}(\text{st}_A^{i-1}, m; R_i)$
	12: $\text{RECEIVE}(B, \emptyset, \text{ct})$
	13: return B

Figure 4.12: r-RID adversary for the proof of Theorem 11 (resp. Theorem 12) on the left (resp. on the right).

We present such an adversary \mathcal{A}_i on the left of Figure 4.12. The adversary samples n_s messages m_1, \dots, m_{n_s} at random and, letting A and B be the two parties, makes A send these with randomness R_1, \dots, R_{n_s} , respectively. Then, \mathcal{A}_i makes B receive the last ciphertext ct_{n_s} . Next, \mathcal{A}_i samples a random message m , sends it using state st_A^{i-1} and randomness R_i to get a ciphertext ct and makes B receive it. Now, as ct_{n_s} is sent after ct (ct_{n_s} is sent with $\text{st}_A^{n_s-1}$ and ct with st_A^{i-1}), ct_{n_s} and ct will decrypt respectively to num_{n_s} and num_i s.t. $\text{num}_{n_s} > \text{num}_i$ by correctness. Then, if $m \neq m_i$, then ct is different from the i -th ciphertext ct_i (as we assume $\text{Send}(\text{st}_A^{i-1}, \cdot; R_i)$ is injective). Therefore, if ct is accepted and $m \neq m_i$, then ct and num_i satisfy the forgery predicate of the r-RID game in Figure 4.4. In addition, as ct_{n_s} is sent and delivered honestly, forgery and $\text{bad-}\bar{P}$ in the r-RID game hold with respect to B, $\text{num}_i, \text{num}_{n_s}, \perp, \text{ct}, \perp, \text{ct}_{n_s}$, relevant x and any y (as in line 6 of Figure 4.4), so the adversary wins. We call this event win.

We now compute the probability that win happens, which is the probability that $\text{ct}_i \neq \text{ct}$ and B accepts ct . Let m_1, \dots, m_{n_s} and the whole randomness R ($R = R_{-1}, R_0, \dots, R_{n_s}$) be fixed. As before, let S_i be the set of messages m s.t. $\text{Receive}(\text{st}_B^1, \text{ct})$ accepts, for $\text{ct} = \text{Send}(\text{st}_A^{i-1}, m; R_i)$. Note that since S_i depends only (m_1, \dots, m_{n_s}, R) (which are now fixed), it is deterministic.

Therefore, conditioned on m_1, \dots, m_{n_s}, R , we have

$$\Pr_m[\text{win}] = \Pr_m[m \in S_i \wedge m \neq m_i] = \Pr_m[m \in F_i] = \frac{|F_i|}{2^n}$$

as m is sampled uniformly at random. Hence, overall

$$\Pr_{m, m_1, \dots, m_{n_s}, R}[\text{win}] = \mathbb{E}_{m_1, \dots, m_{n_s}, R}[\Pr_m[m \in F_i]] = \frac{\mathbb{E}[|F_i|]}{2^n}.$$

Note that both the source and the adversary sample m_1, \dots, m_{n_s}, R uniformly at random. Finally, if $\mathbb{E}[|F_i|] > 2^{n-\gamma}$, then $\Pr[\text{win}] > \frac{2^{n-\gamma}}{2^n} = 2^{-\gamma}$, which leads to the contradiction. \square

By the previous lemma, we have $\log(\mathbb{E}[|F_i|] + 1) \leq \log(2^{n-\gamma} + 1) \leq \max(0, n - \gamma) + 1$. Plugging this result into Lemma 13, we get $\mathbb{E}[C] \leq k + r + (n_s - 1)(\max(0, n - \gamma) + 2)$. In addition, as our encoder outputs a uniquely decodable code, we know that $n_s n + r \leq \mathbb{E}[C]$ by Shannon's source coding theorem. Hence, we get

$$k + r + (n_s - 1)(n - \gamma + 2) \geq n_s n + r \iff k \geq n + (n_s - 1)(\gamma - 2)$$

if $\gamma \leq n$ and otherwise

$$k + r + (n_s - 1)2 \geq n_s n + r \iff k \geq 2 + n_s(n - 2).$$

Now that the first two lower bounds have been shown, we prove the final bound in the following lemma.

Lemma 15. Let k, n, n_s, γ as in the statement of the theorem. Then,

$$k \geq n n_s - \frac{1}{1 - \frac{2^n n_s}{2^\gamma}}.$$

Proof. In order to prove this lemma, we build another encoder/decoder pair very similar to the previous one. They are shown in Figure 4.11 (*with* the boxed instructions for the encoder and the boxed decoder on the right). The only difference in the encoder is that if *one* false positive is found, the encoder outputs a bit set to zero and the trivial encoding of the input. Let's call this event fail. If fail does not occur, a bit set to 1, the last ciphertext, and the randomness are output.

In the decoder, either the first bit of the input is set to 0 and the input is returned straightaway, or $b = 1$ and the decoder proceeds as before. However, as there are no false positives, the m_i can be recovered without using the indices e_i (i.e., the correct message would be the only element of S_i). Overall, the expected codeword length is $\mathbb{E}[C] = 1 + \alpha n n_s + (1 - \alpha)k + r$, where $\alpha := \Pr[\text{fail}]$, as if fail occurs a trivial encoding (on $1 + n n_s + r$ bits) is used, and otherwise the

Chapter 4. On Active Attack Detection in Messaging with Immediate Decryption

encoder outputs $1 + k + r$ bits. By Shannon source coding theorem, we obtain

$$1 + \alpha n n_s + (1 - \alpha)k + r \geq n n_s + r \iff k \geq n n_s - \frac{1}{1 - \alpha}$$

In addition, we have $\alpha := \Pr[\text{fail}] = \Pr[\cup_{i=1}^{n_s-1} \{|F_i| \geq 1\}] \leq \sum_{i=1}^{n_s-1} \Pr[|F_i| \geq 1]$ as fail occurs if at least one of the sets of false positives F_i contains an element. Then, we have $\Pr[|F_i| \geq 1] \leq \mathbb{E}[|F_i|] \leq 2^{n-\gamma}$, where the first inequality follows from Markov's inequality and the second from Lemma 14. Overall, we get $\alpha \leq \frac{n_s 2^n}{2^\gamma}$. Hence,

$$k \geq n n_s - \frac{1}{1 - \alpha} \geq n n_s - \frac{1}{1 - \frac{2^n n_s}{2^\gamma}}.$$

Finally, some algebra shows that this bound is tighter than the second one when

$$2^\gamma \geq 2^n n_s \frac{2 - 2n_s}{3 - 2n_s},$$

that is, when γ is larger than $\approx n + \log(n_s)$. □

This completes the proof. □

On Imperfect Correctness. For simplicity in the proof, we only considered RC schemes which are perfectly correct. Note, however, that it should be possible to obtain a slightly worse bound for RC schemes that are *not* perfectly (computationally or statistically) correct. In more detail, perfect correctness is used twice in the proof of Theorem 11: (1) in the encoder to argue that the encoded messages will decrypt properly and (2) in the reduction (Figure 4.12) to argue that $m \neq m_i \Rightarrow \text{ct} \neq \text{ct}_i$. Then, if the probability that a correctness error arises is at most δ , we can argue as follows. In case (1), we can simply use the trick used to prove the 3rd bound (i.e., output the trivial encoding if the encoded message does not decrypt properly) to get the same bounds $-1/(1 - \delta \cdot n_s)$. Then, in case (2), we will have at the end $\Pr[\text{win}] > 2^{-\gamma} - \delta$ (the bound calculated at the end of our proof of Lemma 14) as we need to take into account the probability that m triggers a correctness error. This should incur an additional $\approx -\log(\delta)$ loss in the bound. Overall, the proof still holds with $\delta > 0$, and if it is small then the bounds remain nearly identical.

Beyond Uniformly Sampled Messages. As stated, our lower bound assumes that the n_s messages encoded are sampled uniformly at random from the message space. We note that a comparable bound holds if the distribution is different from uniform given the bound is stated in terms of the entropy of the messages.

4.5.2 Communication Cost for r -UNF Security

We consider a perfectly correct ARC (i.e., the function $\text{Send}(\text{st}_P, \text{ad}, \cdot; r)$ is injective for all randomness r , valid states st_P and associative data ad). The following theorem states that the tag size of a secure ARC grows linearly in the number of messages (times either the security parameter or message size).

Theorem 12. Let Π be a perfectly correct ARC, n_s and λ be fixed, and T_{λ, n_s} be the time complexity of the (efficient) adversary given on the right of Figure 4.12. In addition, let $\gamma \in \mathbb{Z}$ be such that for all adversaries \mathcal{A} running in at most time T_{λ, n_s} which query oracle SEND at most n_s times, we have: $\Pr[r\text{-UNF}_{\Pi}(\mathcal{A}) \Rightarrow 1] \leq \frac{1}{2^\gamma}$. Let $\mathcal{M} = \{0, 1\}^n$ and $\mathcal{T} = \{0, 1\}^k$ be the plaintext and tag space associated to Π , respectively. Then, $k \geq n_s(\gamma - 2)$, if $\gamma \leq n$, and $k \geq n_s(n - 2)$, if $\gamma > n$. A third lower bounds gives

$$k \geq nn_s - \frac{1}{1 - \frac{2^n(n_s+1)}{2^\gamma}},$$

which is tighter for low values of n (e.g. $n = 1, 2$) and when $\gamma > n + \log(n_s)$.

Proof. We only provide the idea of the proof, as it is nearly identical to the one of Theorem 11. The intuition is that if a party \bar{P} sends n_s messages and then an authentication tag at to P , then at must contain information about all the ciphertexts previously sent. This is because \bar{P} cannot know in the worst case which messages were ever received by P .

More precisely, the only difference with the proof of Theorem 11 is that we use at instead of ct_{n_s} ; the rest follows similarly. In particular, the encoder outputs the randomness, the tag and the indices of the encoded messages in the sets of correctly received messages. Then, the decoder receives the authentication tag and tries to receive all possible ciphertexts ct_1 . Among the ones that are successfully received, it extracts the correct message using the index provided by the encoder. Next it moves to receiving all possible ciphertexts ct_2 and so on, until the n_s messages are received. In addition, we give on the right of Figure 4.12 the adversary that can be used to prove an upper bound on the number of messages that are correctly received (i.e., the number of messages in the sets S_i used in the proof of Theorem 11). \square

4.6 Optimisations and Performance/Security Trade-Offs

In Section 4.5, we showed that r -RID and r -UNF impose a linear communication complexity on RC and ARC schemes respectively. In this section we explore ways to bypass these lower bounds and propose practical approaches for active attack detection. We first argue that s -RID/ s -UNF security can be achieved at a much lower cost than r -RID/ r -UNF security. Noting that ciphertexts grow without bound in the schemes presented up to this point, we propose two methods of pruning, or garbage collection, to reduce communication overhead. The first involves pruning the set R that is sent for s -RID security every full round trip in

Chapter 4. On Active Attack Detection in Messaging with Immediate Decryption

Figure 4.6. The second involves pruning unnecessary values included in authentication tags in our construction in Figure 4.9 for UNF security. We then conclude the chapter by proposing a lightweight three-move protocol that authenticates communication in both directions over the out-of-band channel.

4.6.1 On the Practicality of s-RID and s-UNF Security

Security Notion:	r-RID	s-RID
Overhead	$\mathcal{O}(n\lambda + c_n)$	$\mathcal{O}(\lambda + c_n)$
Optimised overhead (Section 4.6.2)	N/A	$\mathcal{O}(\lambda + c_{n_{\text{fresh}}})$

Table 4.1: Overhead induced by the two RID security notions. We assume that n messages are received and c_i is the space needed to encode i ordinals. The variable n_{fresh} refers to the number of messages received in the last two epochs.

We focus here on s-RID security, but similar arguments hold for s-UNF security. The scheme RC_{RID} (Figure 4.6) achieves s-RID security by sending to the counterpart the list of received ordinals and an *hash* of the set R . Recall informally that this suffices for security because a party can immediately detect when their counterpart has received a forgery (by keeping in state their sent messages and recomputing the hash). Table 4.1 summarises the overhead incurred by the two notions together with the optimisation presented in the next subsection (Section 4.6.2), where we show that it is enough to only send information about messages received during the two last epochs. This significantly reduces the overhead for the scenarios in which the communication is “balanced”.

Ordinals. The RID-secure RC of Figure 4.6 sends the set of *received* ordinals for authentication (line 2 of the RC_{RID} .Send algorithm). Since ordinals are elements of a set on which a total order is defined, a simple optimisation—that could reduce the overhead by up to 50%—consists in sending the smallest set among the set of received ordinals *or* the set of not received ordinals alongside a bit indicating which type of set has been sent. This optimisation applies to all schemes that send sets of ordinals.

One can reduce ciphertext size further by optimising for the “good case” scenario where messages are delivered in-order; in this case, ordinals can be encoded in ranges. For epochs with no lost messages, it suffices to encode only the last index. In any case, as the size of a single message in today’s secure messaging applications can be several kilobytes or more, especially when audio and video is used, the overhead that s-RID imposes seems reasonable. We leave a deeper and more concrete analysis to determine the impact of RID/UNF security in practice to future work.

Incremental Hashing. In Figure 4.6, the entire set of received messages is hashed (using a regular hash function) every time a message is sent by P. Consequently, when P receives a new message, the entire hash must be re-computed when P sends their next message. To

avoid this, the scheme can instead use an *incremental hash function* (Definition 10) such that, when a message is received, an efficient operation only depending on the new message and the previous digest can be executed to derive the new digest. Hash digests can be as small as a group element [CDv⁺03]. This enables parties to prune their set of sent/received messages in state. For example, if \bar{P} receives a message m claiming that \bar{P} has received the first k messages from P , and P has received messages for all possible ordinals that precede the ordinal of m , then P can safely store just the incrementally-hashed value corresponding to the first k messages, since \bar{P} can no longer claim to have only received a strict subset of the k messages.

The security proof then follows almost exactly as in the proof that use the collision resistance of the hash function by instead assuming *set collision resistance* (Definition 11) and arguing that a set collision can be constructed in the exact same situations as a collision when using a (regular) hash function.

4.6.2 Epoch-Based Optimisation for s-RID Security

In this subsection we show how to design an optimised s-RID-secure RC scheme given a correct and ORDINALS-secure RC scheme. A formal description is given in Figure 4.13.

At initialisation, one party is associated with $ep = 0$, say P , and the other, say \bar{P} , with $ep = 1$. Each time a party sends a message they also attach their current ep to the message. Upon receiving a message, a party P with $ep = t$ checks whether the ep attached to the message received from \bar{P} is at most $t + 1$. If the \bar{P} 's ep is exactly $t + 1$ the party P updates their ep to $t + 2$. The ep value of the parties are always one apart at each point in time. Note this is exactly how epochs were defined by Alwen et al. [ACD19] except each epoch was associated with an iteration of the Diffie-Hellman-based asymmetric ratchet.

To achieve s-RID security the sender does the following. Whenever sending a message with $ep = t$, the sender attaches the num and the accumulated hash of all messages they received during the time their ep was t and $t - 2$. That is, the parties do the same as the original s-RID construction, but only for the messages they have received in the last 2 epochs.

Although this optimisation does not change the worst-case complexity of Figure 4.6, if the direction of the conversation changes frequently enough, the overhead significantly decreases.

We continue by stating the main theorem of this subsection and providing a proof sketch.

Theorem 13. Consider collision resistant hash function H used to build RC_{s-RID} (defined in Figure 4.13). Then, we have that for every efficient adversary \mathcal{A} , one can build an adversary \mathcal{B} such that

$$\text{Adv}_{RC_{s-RID}}^{s\text{-rid}}(\mathcal{A}) \leq \text{Adv}_H^{\text{cr}}(\mathcal{B}).$$

Proofsketch. Let $[(ct_f, \cdot, t_f), (ct_h, \cdot, t_h)]$, be the closest pair of sent-received messages contra-

Chapter 4. On Active Attack Detection in Messaging with Immediate Decryption

$RC_{s-RID}.Setup(1^\lambda)$ 1: $pp_0 \xleftarrow{\$} RC.Setup(1^\lambda)$ 2: $hk \xleftarrow{\$} H.KGen(1^\lambda)$ 3: $hk' \xleftarrow{\$} H.KGen(1^\lambda)$ 4: $pp \leftarrow (pp_0, hk, hk')$ 5: return pp	$RC_{s-RID}.Receive(st_P, ad, ct)$ 1: $(ct', ep^{\bar{P}}, R^{\bar{P}}) \leftarrow ct$ 2: $(st'_P, hk, hk', S, R_{curr}, R_{prev}, ep) \leftarrow st_P$ 3: $ad' \leftarrow (ad, ep^{\bar{P}}, R^{\bar{P}})$ 4: $(acc, st'_P, num, pt) \leftarrow RC.Receive(st'_P, ad', ct')$ 5: if $\neg acc$: return $(false, st_P, \perp, \perp)$ 6: $h \leftarrow H.Eval(hk, (num, ad, ct))$ 7: if $checks(st_P, ct, h, num)$: 8: return $(false, st_P, \perp, \perp)$ 9: $st_P.R_{curr} \leftarrow R_{curr} \cup \{(num, h)\}$ 10: $st_P.st'_P \leftarrow st'_P$ 11: // Advance epochs accordingly 12: if $ep^{\bar{P}} = st_P.ep + 1$: 13: $st_P.ep \leftarrow st_P.ep + 2$ 14: $st_P.R_{prev} \leftarrow R_{curr}$ 15: $st_P.R_{curr} \leftarrow \emptyset$ 16: return (acc, st_P, num, pt)
$RC_{s-RID}.Init(pp)$ 1: $(pp_0, hk, hk') \leftarrow pp$ 2: $(st'_A, st'_B, z') \xleftarrow{\$} RC.Init(pp_0)$ 3: $ep_{ack} \leftarrow 0$ 4: $ep \leftarrow 0$ 5: $S, R_{curr}, R_{prev} \leftarrow \emptyset$ 6: $st_A \leftarrow (st'_A, hk, hk', S, R_{curr}, R_{prev}, ep)$ 7: $st_B \leftarrow (st'_B, hk, hk', S, R_{curr}, R_{prev}, ep)$ 8: $z \leftarrow (z', pp)$ 9: return (st_A, st_B, z)	$checks(st_P, ct, h, num)$ 1: $(nums', h') \leftarrow ct.R$ 2: $ep' \leftarrow ct.ep$ 3: if $ep' > st_P.ep + 1$: 4: $s\text{-bool} \leftarrow 1$ 5: $R^* \leftarrow \{(num', _) \in st_P.S : num' \in nums'\}$ 6: $s\text{-bool} \leftarrow s\text{-bool} \vee (H.Eval(st_P.hk', R^*) \neq h')$ 7: return $s\text{-bool}$
$RC_{s-RID}.Send(st_P, ad, pt)$ 1: $(st'_P, hk, hk', S, R_{curr}, R_{prev}, ep) \leftarrow st_P$ 2: $nums' \leftarrow \{num' : (num', _) \in R_{curr} \cup R_{prev}\}$ 3: $R' \leftarrow (nums', H.Eval(hk', R_{curr} \cup R_{prev}))$ 4: $ad' \leftarrow (ad, ep, R')$ 5: $(st_P.st'_P, num, ct') \xleftarrow{\$} RC.Send(st'_P, ad', pt)$ 6: $ct \leftarrow (ct', ep, R')$ 7: $h \leftarrow H.Eval(hk, (num, ad, ct))$ 8: $st_P.S \leftarrow S \cup \{(num, h)\}$ 9: return (st_P, num, ct)	

Figure 4.13: Optimised s-RID-secure RC scheme RC_{s-RID} given a correct and ORDINALS-secure RC scheme RC.

dicting the s-RID condition. That is, (ct_f, \cdot, t_f) is a forgery received by \bar{P} before they sent the honest message (ct_h, \cdot, t_h) which was received by P. We consider the time when (ct_h, \cdot, t_h) was sent by \bar{P} . As mandated by the construction (ct_h, \cdot, t_h) contained num values and accumulated hash of all messages \bar{P} has received at epochs t_h and $t_h - 2$. Following the same argument as the proof of Theorem 8 one can show that no forgeries, including (ct_f, \cdot, t_f) , were received by \bar{P} while $ep_{\bar{P}} \in \{t_h, t_h - 2\}$.

The two messages that changes $ep_{\bar{P}}$ from $t_h - 4$ to $t_h - 2$ and from $t_h - 2$ to t_h , were honest messages by P. Let us call them $(ct_{t_h-3}, \cdot, t_h - 3)$ and $(ct_{t_h-1}, \cdot, t_h - 1)$ respectively. Note that,

4.6 Optimisations and Performance/Security Trade-Offs

both these messages were received after (ct_f, \cdot, t_f) was received and before (ct_f, \cdot, t_f) was sent, as otherwise (ct_h, \cdot, t_h) would contradict (ct_f, \cdot, t_f) . Note that between sending the messages $(ct_{t_h-3}, \cdot, t_h-3)$ and $(ct_{t_h-1}, \cdot, t_h-1)$, ep_P was changed meaning P received a message with $ep = t_h-2$ that caused the change of ep_P . Let us call this message $(ct_{t_h-2}, \cdot, t_h-2)$. We argue this message should have been forged.

Let us assume by contradiction that this message was honest. Note that $(ct_{t_h-2}, \cdot, t_h-2)$ was sent after $(ct_{t_h-3}, \cdot, t_h-3)$ was received, hence after (ct_f, \cdot, t_f) was received. Now if $(ct_{t_h-2}, \cdot, t_h-2)$ is honest it forms a pair with (ct_f, \cdot, t_f) which violates the s-RID condition and has less distance from the original pair which is a contradiction. So $(ct_{t_h-2}, \cdot, t_h-2)$ must have been forged. A visualisation of the scenario can be found in Figure 4.14.

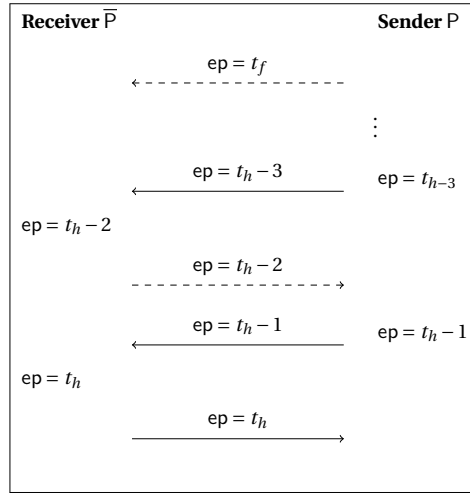


Figure 4.14: Visualising the proof sketch of Theorem 13. The dotted messages are forged and the others are honest messages. Intuitively we argue if the last message did not contradict the first message, the fourth message would have contradicted the third, therefore violating s-RID in the other direction.

One other observation is that, $(ct_{t_h-2}, \cdot, t_h-2)$ was received before $(ct_{t_h-1}, \cdot, t_h-1)$ was sent, hence, before (ct_h, \cdot, t_h) was received. This shows that the pair $[(ct_{t_h-2}, \cdot, t_h-2), (ct_{t_h-1}, \cdot, t_h-1)]$ also violates the s-RID (for \bar{P} and not P) and has less distance than the original pair. \square

4.6.3 Pruning for UNF Security

In this subsection, we present a scheme that optimises bandwidth consumption for UNF security (Figure 4.15). A complete description and security proof are given below. Our scheme takes advantage of the fact that messages sent out-of-band cannot be forged. Suppose that P sends an authentication tag to \bar{P} , then \bar{P} acknowledge the reception of the tag to P. At this point, P no longer needs to send the information that \bar{P} has already obtained. Our scheme supports out-of-order communication even on the authenticated channels. Our approach is complicated by this and the fact that parties have to keep track of, e.g., which tags their

Chapter 4. On Active Attack Detection in Messaging with Immediate Decryption

partner has received to determine what is safe to prune from state (in S_{atSeen}), which incurs relatively small overhead in typical executions.

Scheme Description. Send stores the hash of (ad, ct) for the message being sent, together with the corresponding num that the underlying RC.Send call returns. It stores (num, h) in a set S , which is in turn stored in the party's internal state. The Send algorithm also updates the ordinal num in the state.

The Receive procedure verifies whether RC.Receive accepts the inputs and verifies that the received message is not a forgery based on any information stored from previous AuthReceive calls. Given this passes, Receive stores the hash of (ad, ct) together with the ordinal num returned by RC.Receive in a set R .

The AuthSend procedure is similar to the unoptimised one, except that (1) it stores the set of sent messages S authenticated *within the current* at into an array S_{at} , indexed by counter cnt_{at} , and (2) it empties the set S_{atSeen} locally; the goal of this set is to communicate which tags that the caller has received from their counterpart.

Before processing the message, AuthReceive first verifies whether $\text{cnt}_{\text{at}}^{\bar{P}} \leq \text{max-cnt}_{\text{at}}$. The goal of this check is to avoid processing old authentication tags, since if the check is true, AuthReceive must have previously authenticated the tag's content in a newer (in terms of $\text{cnt}_{\text{at}}^{\bar{P}}$) tag. However, pruning still needs to be performed before returning in this case since the contents of set S_{atSeen} is only sent once before being flushed in a given AuthSend call.

Otherwise, given the sent and received messages of the two parties are consistent, AuthReceive first stores the counter of the input tag in S_{atSeen} , which will be sent to the counterpart in the next call to AuthSend. Then it removes the already authenticated messages from memory. The caller already authenticated the subset R_{\subseteq}^P and it can thus remove the corresponding messages from its set R , which represents now the set of currently unauthenticated received messages. Similarly, AuthReceive uses the set of authentication tags that the counterpart already processed to prune the set of sent messages.

In detail, the pruning of sent messages works as follows. When a party P receives a set of messages $S_{\text{at}}^{\bar{P}}[\text{cnt}_{\text{at}}] \leftarrow S$ sent by the counterpart with the authentication tag number cnt_{at} , it stores them in a set S_{ack} . Then, when P sends a subsequent authentication tag back to \bar{P} , it informs \bar{P} that the authentication tag cnt_{at} was received using the S_{atSeen}^P set. When this tag is delivered, \bar{P} can remove the acknowledged messages $S_{\text{at}}^{\bar{P}}[\text{cnt}_{\text{at}}]$ for all counters in S_{atSeen}^P from its set $S^{\bar{P}}$. This reduces the size of the authentication tag as, on every round-trip on the out-of-band channel, all authenticated messages can be removed from the sets S and R . To reduce the size even further, we can use the same hashing optimisation for the received set that we used in Figure 4.6.

Security Analysis. We informally argue that the scheme prunes only messages that have already been authenticated. The procedures use sets $\text{stp}.S$ and $\text{stp}.R$ to detect active attacks.

4.6 Optimisations and Performance/Security Trade-Offs

<p>ARC_{OP}.Setup(1^λ)</p> <hr/> <pre> 1: $pp_0 \leftarrow \text{RC.Setup}(1^\lambda); hk \leftarrow \text{H.KGen}(1^\lambda)$ 2: return (pp_0, hk) </pre> <p>ARC_{OP}.Init(pp)</p> <hr/> <pre> 1: (pp_0, hk) $\leftarrow pp$ 2: (st'_A, st'_B, z') $\leftarrow \text{RC.Init}(pp)$ 3: $num, \text{max-num}, \text{cnt}_{at}, \text{max-cnt}_{at} \leftarrow 0$ 4: $S, R, S_{ack}, S_{at}, S_{atSeen} \leftarrow \emptyset$ 5: $st_A \leftarrow (st'_A, hk, S, R, S_{ack}, num, \text{max-num},$ 6: $\text{cnt}_{at}, \text{max-cnt}_{at}, S_{at}, S_{atSeen})$ 7: $st_B \leftarrow (st'_B, hk, S, R, S_{ack}, num, \text{max-num},$ 8: $\text{cnt}_{at}, \text{max-cnt}_{at}, S_{at}, S_{atSeen})$ 9: $z \leftarrow (z', pp)$ 10: return (st_A, st_B, z) </pre> <p>ARC_{OP}.Send(stp, ad, pt)</p> <hr/> <pre> 1: (st'_p, hk, S, \dots) $\leftarrow stp$ 2: ($stp.st'_p, num, ct$) $\leftarrow \text{RC.Send}(st'_p, ad, pt)$ 3: $h \leftarrow \text{H.Eval}(hk, (ad, ct))$ 4: $stp.S \leftarrow S \cup \{(num, h)\}$ 5: $stp.num \leftarrow num$ 6: return (stp, num, ct) </pre> <p>ARC_{OP}.Receive(stp, ad, ct)</p> <hr/> <pre> 1: ($st'_p, hk, _, R, S_{ack}, _, \text{max-num}, \dots$) $\leftarrow stp$ 2: (acc, st'_p, num, pt) $\leftarrow \text{RC.Receive}(st'_p, ad, ct)$ 3: if $\neg acc$: return ($false, stp, \perp, \perp$) 4: $h \leftarrow \text{H.Eval}(hk, (ad, ct))$ 5: if $num \leq \text{max-num} \wedge (num, h) \notin S_{ack}$: 6: return ($false, stp, \perp, \perp$) 7: $stp.R \leftarrow R \cup \{(num, h)\}$ 8: $stp.st'_p \leftarrow st'_p$ 9: return (acc, stp, num, pt) </pre>	<p>ARC_{OP}.AuthSend(stp)</p> <hr/> <pre> 1: ($_, _, S, R, _, num, _, \text{cnt}_{at}, _, S_{atSeen}$) $\leftarrow stp$ 2: $at \leftarrow (S, R, num, \text{cnt}_{at}, S_{atSeen})$ 3: $stp.\text{cnt}_{at} \leftarrow stp.\text{cnt}_{at} + 1$ 4: $stp.S_{at}[stp.\text{cnt}_{at}] \leftarrow S$ 5: $stp.S_{atSeen} \leftarrow \emptyset$ 6: return (stp, num, at) </pre> <p>ARC_{OP}.AuthReceive(stp, at)</p> <hr/> <pre> 1: ($_, _, S, R, S_{ack}, num, \text{max-num},$ 2: $\text{cnt}_{at}, \text{max-cnt}_{at}, S_{at}, _$) $\leftarrow stp$ 3: ($S^{\bar{P}}, R^{\bar{P}}, num^{\bar{P}}, \text{cnt}_{at}^{\bar{P}}, S_{atSeen}^{\bar{P}}$) $\leftarrow at$ 4: $R_{\underline{C}}^P \leftarrow \{(num, _) \in R : num \leq num^{\bar{P}}\}$ 5: if $\text{cnt}_{at}^{\bar{P}} \leq \text{max-cnt}_{at}$: 6: $\text{prune}(stp, R^{\bar{P}}, \text{cnt}_{at}^{\bar{P}}, S_{atSeen}^{\bar{P}}, R_{\underline{C}}^P)$ 7: return ($true, stp, num^{\bar{P}}$) 8: // \bar{P} received a forgery 9: if $R^{\bar{P}} \not\subseteq S$: return ($false, stp, num$) 10: // P received a forgery 11: if $R_{\underline{C}}^P \not\subseteq S^{\bar{P}}$: return ($false, stp, num$) 12: $stp.S_{ack} \leftarrow stp.S_{ack} \cup S^{\bar{P}}$ 13: $stp.\text{max-num} \leftarrow \max\{\text{max-num}, num^{\bar{P}}\}$ 14: $stp.\text{max-cnt}_{at} \leftarrow \max\{\text{cnt}_{at}^{\bar{P}}, stp.\text{max-cnt}_{at}\}$ 15: $\text{prune}(stp, R^{\bar{P}}, \text{cnt}_{at}^{\bar{P}}, S_{atSeen}^{\bar{P}}, R_{\underline{C}}^P)$ 16: return ($true, stp, num^{\bar{P}}$) </pre> <p>prune($stp, R^{\bar{P}}, \text{cnt}_{at}^{\bar{P}}, S_{atSeen}^{\bar{P}}, R_{\underline{C}}^P$)</p> <hr/> <pre> 1: $stp.S_{atSeen} \leftarrow stp.S_{atSeen} \cup \{\text{cnt}_{at}^{\bar{P}}\}$ 2: $stp.R \leftarrow stp.R \setminus R_{\underline{C}}^P$ 3: for $i \in S_{atSeen}^{\bar{P}}$ do 4: $stp.S \leftarrow stp.S \setminus stp.S_{at}[i]; stp.S_{at}[i] \leftarrow \emptyset$ </pre>
------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

Figure 4.15: Optimised UNF-secure ARC scheme ARC_{OP} based on a RC scheme RC (Definition 33). The sets S , R and S_{ack} are as in Figure 4.6. The variable max-num represents the largest num received in an at. The counters cnt_{at} and max-cnt_{at} keep track of how many tags have been sent and largest cnt_{at} received inside of a tag, respectively. S_{atSeen} is the list of cnt_{at} of received at since the last sent one; $S_{at}[i]$ contains the content of S sent in the i th at.

Chapter 4. On Active Attack Detection in Messaging with Immediate Decryption

AuthReceive prunes $st_P.R$ by removing elements in $R_{\underline{c}}^P$; since the procedure authenticates the elements in $R_{\underline{c}}^P$ at line 11, it is safe to prune $st_P.R$. AuthReceive prunes $st_P.S$ by removing elements in $st_P[i]$ for $i \in S_{at}Seen^{\bar{P}}$. The set $S_{at}Seen^{\bar{P}}$ contains counters of the authentication tags that P sent to \bar{P} and \bar{P} correctly received. Moreover, the AuthReceive procedure updates $st.S_{at}Seen^{\bar{P}}$ at line 1, i.e., after the integrity checks. Since the AuthSend stores the set of sent messages S authenticated within the current at into the array S_{at} , pruning $st_P.S$ only removes messages that have already been received and authenticated by \bar{P} .

Recall that the adversary can only delete and replay authentication tags in the out-of-band channel. We informally discuss how the scheme handles these cases. Assume P and \bar{P} exchange some messages, P receives an authentication tag $at^{\bar{P}}$ from \bar{P} and then sends the authentication tag at^P ; the adversary removes at^P from the channel. Since the adversary removes at^P , P does not acknowledge the reception of $at^{\bar{P}}$ to \bar{P} (because AuthSend empties $st_P.S_{at}Seen$ at every invocation). Consequently, \bar{P} does not prune $st_{\bar{P}}.S$: these messages will be authenticated with the next authentication tag and security is preserved. The AuthReceive procedure handles adversarial reordering of authentication tags with counters cnt_{at} at line 5.

We formally state our security claim for ARC_{OP} below.

Theorem 14. Consider collision resistant hash function H used to build ARC_{OP} (defined in Figure 4.15). Then, we have that for every efficient adversary \mathcal{A} , one can build an adversary \mathcal{B} such that

$$\text{Adv}_{ARC_{OP}}^{\text{unf}}(\mathcal{A}) \leq \text{Adv}_H^{\text{cr}}(\mathcal{B}).$$

Moreover, ARC_{OP} is correct and ORDINALS secure.

Proof. Correctness and ORDINALS-security for the transformation of Figure 4.15 follow from the scheme in Figure 4.9.

The scheme is the same as Figure 4.9 modulo the optimisations we introduced. The proof of the theorem thus reduces to showing that the optimisations preserve the security properties of the unoptimised ARC scheme ARC_{base} (Figure 4.9). Observe that $st_P.R$ and $st_P.S$ are used to detect active attacks. We start by showing that pruning these sets does not undermine UNF security.

- The set $st_P.R$ is pruned by removing elements from $R_{\underline{c}}^P$, which was authenticated on line 11 of AuthReceive. We therefore know that messages in $R_{\underline{c}}^P$ are honest. Thus, we can stop sending them to \bar{P} hereafter.
- $st_P.S$ is pruned by all sets $st_P.S_{at}[i]$ for $i \in S_{at}Seen^{\bar{P}}$. By construction we know that $S_{at}Seen^{\bar{P}}$ contains counters for which \bar{P} *accepted* the authentication tags, since those are included in line 1 of prune. Therefore, $\{st_P.S_{at}[i]\}_{i \in S_{at}Seen^{\bar{P}}}$ contains all messages in $st_P.S$ that \bar{P} correctly received and authenticated, which the procedure stores in $st_{\bar{P}}.S_{ack}$ for future checks. Hence, P can safely stop sending those and prune S correspondingly.

We proceed by showing that UNF security still holds. By the arguments above, an authentication tag at that the AuthSend procedure generates *after* another authentication tag at', will contain only messages that have *not* been authenticated in at'. Therefore the check on line 5 preserves security.

The check on line 9 verifies whether $R^{\bar{P}} \subseteq S$. Without pruning, this property is met in the absence of forgeries as shown for ARC_{base} . Assume for contradiction that $R^{\bar{P}}$ contains a message not authenticated yet, but S does not contain this message due to pruning. This means that the message was removed from S by removing one of the values in $st_P.S_{at}$ whose counter cnt_{at} was present in $S_{at}Seen^{\bar{P}}$. Since the counter is present in $S_{at}Seen^{\bar{P}}$, we know that \bar{P} accepted the authentication tag containing cnt_{at} , i.e., \bar{P} correctly received and authenticated the message. But this means by construction that \bar{P} pruned the message from R on line 2 of prune, which leads to a contradiction. Therefore the check preserves UNF security.

The check on line 11 verifies whether $R_{\bar{C}}^P \subseteq S^{\bar{P}}$. Without pruning, this property is met in absence of forgeries as shown for ARC . Note that P removes from R only messages that have been authenticated (on line 2 of prune), therefore $R_{\bar{C}}^P$ only contains unauthenticated messages. Similarly, by the argument presented in the paragraph above, $S^{\bar{P}}$ contains messages included in at's whose counter was not included in $S_{at}Seen^P$, and therefore unauthenticated messages. We conclude that this check also preserves UNF security. \square

Note that ARC_{OP} (Figure 4.15) sends all the authentication material through the out-of-band channel. This might be impractical when the authenticated out-of-band channel is narrowband, e.g., if parties use QR-codes to authenticate the communication. We can improve the scheme by using both channels: use the insecure channel to send the authentication data and the possibly narrowband authenticated channel to verify the integrity of the data, e.g., using the protocol of Pasini and Vaudenay [PV06]. While the idea of using both channels for authentication is natural [BSSW02], some security risks might arise when the scheme does not correctly match the two messages. Since UNF security depends on both the messages, and therefore on the messages being correctly matched, it might be safer to enforce this property at the scheme level which we leave as future work to formalise.

4.6.4 Lightweight Bidirectional Authentication

In this subsection, we propose a three-move bidirectional authentication protocol. Figure 4.16 describes the protocol at a high level. Parties include in the authentication tag only the set of *received* messages. The receiver of the tag compares then the set of received messages from the counterpart with the set of sent messages. We envision this approach could be used when participants meet in person or online and can both authenticate the respective views of the conversation at the same time. Signal, among other messaging solutions, already requires such an exchange for parties executing its safety numbers protocol that are aiming for mutual verification [Mar17].

Chapter 4. On Active Attack Detection in Messaging with Immediate Decryption




Figure 4.16: Overview of the three-move authentication procedure. The boolean `at-succ` indicates whether the counterpart's set of received messages is a subset of the caller's local set of sent messages.

In what follows, we propose an appropriate security model, describe the protocol in detail and provide a security proof in our model.

Security Model. We modify the UNF game (Section 4.4) by requiring the adversary to run authentication sessions in sequence. We present the corresponding game in Figure 4.17.

When the adversary, through the `AUTHSEND'` oracle, starts the authentication protocol with initial sender P , the adversary's access to `AUTH*`, `SEND'` and `RECEIVE'` oracles is restricted until the three-move authentication session is completed between P and \bar{P} . This is encoded in the next-oob-op variable. This simplifies the exposition, but it is not necessary in particular to restrict `SEND'` and `RECEIVE'` calls in practice, even though parties who authenticate in person would generally not send messages during this time. To handle this, parties can simply buffer messages during authentication that are authenticated in the next authentication session. However, buffering messages implies that an attack carried out during the out-of-band authentication will not be detected until the next authentication protocol. For this reason we block in-band communication during the three-move protocol and we encourage this restriction to be maintained also in practice. We note also that parties are guaranteed slightly weaker security than in the UNF game. Namely, after receiving the first authentication message, the receiver can deduce that their counterpart has not received a forgery but not that they themselves have until they receive the third message in the protocol.

In Definition 39 we define 3M-UNF-security for ARC schemes.

Definition 39 (3M-UNF). An ARC ARC is 3M-UNF if, for all efficient adversaries \mathcal{A} , we have:

$$\text{Adv}_{\text{ARC}}^{3m\text{-unf}}(\mathcal{A}) = \Pr[3\text{M-UNF}_{\text{ARC}}(\mathcal{A}) \Rightarrow 1] = \text{negl},$$

where game 3M-UNF is defined in Figure 4.17.

4.6 Optimisations and Performance/Security Trade-Offs

Game 3M-UNF _{ARC} (\mathcal{A})	Oracle AUTHRECEIVE'(P, j)
1: auth-state[.] \leftarrow 0; next-oob-op \leftarrow \perp 2: play UNF _{ARC} with $\mathcal{A}^{\mathcal{O}}$ (pp, z) and predicates forgery, bad-P, bad- \bar{P}	1: if next-oob-op \neq (P, "authrec", j) : 2: return \perp 3: at \leftarrow auth[(\bar{P} , j)] 4: if at = \perp : return \perp 5: (auth, st, num) \leftarrow AuthReceive(st _P , at) 6: if \neg auth : return \perp 7: i \leftarrow i + 1 8: auth-state[P] \leftarrow auth-state[P] + 1 mod 3 9: if auth-state[P] = auth-state[\bar{P}] = 0 : 10: // Last step of the protocol completed 11: next-oob-op \leftarrow \perp 12: else 13: next-oob-op \leftarrow (P, "authsend") 14: st _P \leftarrow st; state[i] \leftarrow st _P 15: log[i] \leftarrow ("authrec", P, num, at) 16: return num
<hr/> Oracle SEND'(P, ad, pt, r)	
1: if auth-state[P] \neq 0 : return \perp 2: return SEND(P, ad, pt, r)	
<hr/> Oracle RECEIVE'(P, ad, ct)	
1: if auth-state[P] \neq 0 : return \perp 2: return RECEIVE(P, ad, ct)	
<hr/> Oracle AUTHSEND'(P)	
1: if next-oob-op \notin {(P, "authsend"), \perp } : 2: return \perp 3: i \leftarrow i + 1 4: (st _P , num, at) $\xleftarrow{\$}$ AuthSend(st _P) 5: auth[(P, i)] \leftarrow at; state[i] \leftarrow st _P 6: Init \leftarrow (auth-state[P] $\stackrel{?}{=} 0$) // Boolean 7: auth-state[P] \leftarrow auth-state[P] + 1 mod 3 8: log[i] \leftarrow ("authsend", P, num, at, Init) 9: next-oob-op \leftarrow (\bar{P} , "authrec", i) 10: return (num, at)	<hr/> bad-P(log, P, num', at, x, y)
<hr/> forgery(log, P, num, ad, ct, x)	1: return (y > x) \wedge 2: ("authsend", \bar{P} , num', at, \perp) = log[y] \wedge 3: ("authrec", P, num', at) \in log
1: As in Figure 4.7	<hr/> bad- \bar{P} (log, P, num, num', at)
	1: return num \leq num' \wedge 2: ("authrec", \bar{P} , num', at) \in log \wedge 3: ("authsend", P, num', at, false) \in log

Figure 4.17: 3M-UNF game for $\mathcal{O} = \{\text{SEND}', \text{RECEIVE}', \text{AUTHSEND}', \text{AUTHRECEIVE}', \text{EXP}_{\text{pt}}, \text{EXP}_{\text{st}}\}$. Highlighted statements correspond to differences relative to the UNF game (Figure 4.7).

The oracles in Figure 4.17 mandate that the participants send all messages in authentication via the out-of-band channel. By providing no security guarantees on the first message (i.e., delaying guarantees for the receiver until receiving the third message), it is possible to send the first message in the protocol over the in-band channel and then authenticate it in the second message with an additional hash [PV06]. Consequently, the protocol can be made essentially non-interactive out-of-band: the counterpart to the initiator can simply send the message out-of-band, and the bit can be determined easily via determining e.g. QR code scanning success/failure. By contrast, solutions like safety numbers require *both* parties to

Chapter 4. On Active Attack Detection in Messaging with Immediate Decryption

scan QR-codes out-of-band.

Scheme Description. We present a 3M-UNF-secure scheme in Figure 4.18. The AuthSend and AuthReceive procedures encode the three-move authentication protocol of Figure 4.16. To identify the different states of the bidirectional authentication, we borrow the terminology from TCP and refer to SYN, SYN-ACK, and ACK messages and roles. When a party P first calls AuthSend, it takes the SYN role and sends to \bar{P} the set of *received* messages and the current num, i.e., $at \leftarrow (R, num)$; this set is stored in a separate set R_{at} . As described below, we use R_{at} in AuthReceive to optimise the scheme. The counterpart \bar{P} replies with a SYN-ACK message, containing its set of received messages, the current ordinal num and the bit at-succ. The bit at-succ indicates whether P 's set of received messages is included in \bar{P} 's set of sent messages (line 10), i.e., at-succ indicates whether the authentication of P 's set of received messages was successful. As the counterpart, \bar{P} stores the current set of received messages in R_{at} . Upon receiving the SYN-ACK message, the initiator P checks whether $at\text{-succ}^{\bar{P}} = \text{true}$ and rejects the authentication tag otherwise. P then sends the ACK message $at \leftarrow (num, at\text{-succ})$. Finally, \bar{P} calls AuthReceive to process the ACK message. The party checks the at-succ variable to verify that the set $R^{\bar{P}}$ is a subset of S^P . If the check passes, the authentication protocol ends.

The optimisation of the scheme consists in pruning the set of received messages as soon as the counterpart authenticates them. This reduces the size of the authentication tags, since parties include in at only the received messages that have not been authenticated yet. The AuthReceive algorithm on line 11 checks whether the counterpart authenticated set of received messages R . If this is the case, all the authenticated messages are stored in R_{ack} —this set is used in the Receive algorithm to avoid replay attacks—and at the same time those messages are removed from R thanks to set R_{at} , thereby reducing the size of the next authentication tag and memory consumption. After the pruning, the R set contains only received messages that the counterpart *still* needs to authenticate.

Related Work. Dowling et al. [DGP22] propose a scheme that is broadly similar to ours. In particular, their protocol uses three moves in-band to allow parties to agree on a common set of respectively received messages R and R' . Then, to authenticate messages and detect active attacks, parties compare a hash $H(R, R')$ (with hash function H) out-of-band. Note however that they do not consider RID security and that they do not formally treat out-of-order message delivery.

The correctness of ARC_{3M} can be shown using the underlying correctness of RC.

Security Analysis. ORDINALS security is inherited from the underlying RC scheme. As usual, we argue that 3M-UNF security follows from the collision resistance of the underlying hash function.

Theorem 15 (Unforgeability of ARC_{3M}). Consider collision resistant hash function H used to build ARC_{3M} (defined in Figure 4.18). Then, we have that for every efficient adversary \mathcal{A} , one

4.6 Optimisations and Performance/Security Trade-Offs

<pre> ARC_{3M}.Setup(1^λ) ----- 1: // As in Figure 4.9 2: return ARC_{base}.Setup(1^λ) ARC_{3M}.Init(pp) ----- 1: (pp₀, hk) ← pp 2: (st'_A, st'_B, z) $\stackrel{\\$}{\leftarrow}$ RC.Init(pp₀) 3: num ← \perp 4: S, R, R_{ack}, R_{at} ← \emptyset 5: role-at, at-succ ← \perp 6: st_A ← (st'_A, hk, S, R, num, role-at, 7: at-succ, R_{ack}, R_{at}) 8: st_B ← (st'_B, hk, S, R, num, role-at, 9: at-succ, R_{ack}, R_{at}) 10: z ← (z', pp) 11: return (st_A, st_B, z) ARC_{3M}.Send(st_P, ad, pt) ----- 1: // As in Figure 4.9 2: return ARC_{base}.Send(st_P, ad, pt) ARC_{3M}.AuthSend(st_P) ----- 1: (\perp, \perp, R, num, role-at, at-succ, \perp, R_{at}) ← st_P 2: if role-at = \perp: 3: st_P.role-at ← SYN 4: at ← (R, num); st_P.R_{at} ← R 5: elseif role-at = SYN-ACK: 6: at ← (R, num, at-succ) 7: st_P.R_{at} ← R; st_P.at-succ ← \perp 8: else // role-at = ACK 9: at ← (num, at-succ) 10: st_P.role-at, st_P.at-succ ← \perp 11: return (st_P, num, at) </pre>	<pre> ARC_{3M}.Receive(st_P, ad, ct) ----- 1: (st_P, hk, \perp, R, \perp, \perp, R_{ack}, \perp) ← st_P 2: (acc, st'_P, num, pt) ← RC.Receive(st'_P, ad, ct) 3: if \negacc: return (false, st_P, \perp, \perp) 4: h ← H.Eval(hk, (ad, ct)) 5: if \exists h': (num, h') ∈ R_{ack} ∧ h ≠ h': 6: return (false, st_P, \perp, \perp) 7: R ← R ∪ {(num, h)} 8: st_P ← (st_P, hk, \perp, R, \perp, \perp, R_{ack}, \perp) 9: return (acc, st_P, num, pt) ARC_{3M}.AuthReceive(st_P, at) ----- 1: (\perp, \perp, S, R, \perp, role-at, 2: at-succ, R_{ack}, R_{at}) ← st_P 3: R^{\bar{P}} ← \emptyset; at-succ^{\bar{P}} ← true 4: if role-at = \perp: 5: role-at ← SYN-ACK; (R^{\bar{P}}, num^{\bar{P}}) ← at 6: elseif role-at = SYN: 7: (R^{\bar{P}}, num^{\bar{P}}, at-succ^{\bar{P}}) ← at 8: else // receive ACK case 9: (num^{\bar{P}}, at-succ^{\bar{P}}) ← at 10: at-succ ← (R^{\bar{P}} $\stackrel{?}{\subseteq}$ S) // Boolean 11: if at-succ^{\bar{P}}: 12: R_{ack} ← R_{ack} ∪ R_{at}; R ← R \ R_{at} 13: R_{at} ← \emptyset 14: else // failure 15: return (false, st_P, num) 16: st_P ← (\perp, \perp, S, R, num, role-at, 17: at-succ, R_{ack}, R_{at}) 18: return (at-succ, st_P, num^{\bar{P}}) </pre>
--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

Figure 4.18: Optimised 3M-UNF-secure ARC scheme ARC_{3M} based on a RC scheme RC (Definition 33). ARC_{base} refers to the unoptimised ARC defined in Figure 4.9. We assume ARC_{base}.Send updates local variable num. As before, the representation of R communicated can be optimised to contain only a single hash.

can build an adversary \mathcal{B} such that

$$\text{Adv}_{\text{ARC}_{3M}}^{3m\text{-unf}}(\mathcal{A}) \leq \text{Adv}_{\text{H}}^{\text{cr}}(\mathcal{B}).$$

Chapter 4. On Active Attack Detection in Messaging with Immediate Decryption

Proof. We proceed similarly to the proof of Theorem 10. Without loss of generality, we analyze the authentication of P , who we assume to be the initiator, towards \bar{P} . The adversary cannot call the SEND' and $\text{RECEIVE}'$ oracles once the authentication process is started, therefore the sets S and R of both parties are fixed until the completion of the protocol.

To authenticate the set of received messages R^P , P first sends $\text{at} \leftarrow (R^P, \text{num})$ to \bar{P} . To verify the authenticity of R^P , the party \bar{P} verifies whether $R^P \subseteq S^{\bar{P}}$. By the arguments of the proof for Theorem 10, this reduces to the collision-resistance of the hash function H . After receiving the first tag, \bar{P} is able to detect forgeries received by P but not by itself. This is taken into account in the 3M-UNF game (line 3 in $\text{bad-}\bar{P}$), which states that a forgery received by \bar{P} is valid only if it is not detected after receiving the second or third tag in the authentication process. Then, when receiving the second tag $(R^{\bar{P}}, \text{num}^{\bar{P}}, \text{at-succ}^{\bar{P}})$ from \bar{P} , P is able to tell if itself received a forgery if $\text{at-succ}^{\bar{P}} = \text{false}$.⁴ By the same arguments as before, P can tell whether \bar{P} received a forgery by checking $R^{\bar{P}} \subseteq S^P$. Finally, upon receiving the third tag, \bar{P} can detect a forgery using at-succ^P .

The optimisation maintains 3M-UNF-security. Recall that the goal of the optimisation (lines 11-13 in Figure 4.18) is to reduce the size of the R set by storing authenticated messages in R_{ack} . To achieve this reduction, the party executing AuthReceive removes from R the set R_{at} , which is the set of received messages authenticated by the counterpart through the at-succ variable. Since by construction all the messages in R_{at} have been already authenticated by the counterpart, removing them from R does not remove unauthenticated messages from R . \square

⁴In fact, in our formalism, unsuccessful AuthReceive calls are not modelled, so $\text{at-succ}^{\bar{P}}$ will always be true, but we argue this way to be more general.

Group Messaging Part II

5 Cryptographic Administration for Secure Group Messaging

In this chapter, we formalise administration in group messaging by extending *continuous group key agreement* to capture group administrators. An extended abstract corresponding to this work appeared at USENIX Security 2023, and was joint work with David Balbás and Serge Vaudenay [BCV23]. A full version of this work can be found on the Cryptology ePrint Archive [BCV22]. The author of this thesis was an advisor of David Balbás for his MSc thesis which contains some preliminary results [Bal21] on group administration that eventually led to this publication after significant revision.

5.1 Contribution

One of the major challenges in the design of group messaging protocols is the need to account for group *evolution* or *dynamics*, particularly to prevent membership changes from diminishing the confidentiality of sent messages. Since the list of group members may change at any point in time, complex key agreement protocols are required. Overall, securing group membership involves three main aspects: (1) *key updates*, ensuring that new members cannot read past messages and removed members cannot read future messages; (2) *membership consistency*, ensuring that all members faithfully know the list of members at any time; and (3) *securing control messages* (i.e., notifications for member addition and removal operations) from active adversaries and from the delivery service itself (i.e., ensuring authentication).

Many state-of-the-art protocols, including passively-secure continuous group key agreement (CGKA) protocols [ACDT20, KPPW⁺21], Sender Keys (WhatsApp, Signal Messenger) [Wha20, BCG23b] and Matrix [ACD⁺23], include cryptographic mechanisms for securing key updates, but provide weaker and sometimes even no guarantees for securing membership consistency and control messages. We identify membership consistency as both a correctness and a security property that is critical for confidentiality (otherwise, the sender of a message may not know the receivers) but is often ignored in the literature. Failing to secure control messages can also result in catastrophic attacks. Practical examples include the *burgle into a group attack* [RMS18], which exploits the lack of authentication of control messages to allow an

Chapter 5. Cryptographic Administration for Secure Group Messaging

adversary with partial control over the central server to enter arbitrary group chats in Signal and WhatsApp. Recent attacks on the Matrix protocol [ACD⁺23, ADJ⁺24] make use of similar vulnerabilities, enabling the server to take over the control of a group.

In order to secure group membership, we observe that there is a strong trend in practice to distinguish between at least two types of users in a group: group administrators (admins) and standard users. In groups with administrators, all group changes are either performed or approved by the admins. Therefore, we address the problem of secure group management by developing a cryptographic framework for *group administration*.

Administration in Messaging Apps. Generally, an admin has all the capabilities of a standard user plus a set of administrative rights. In practice, admins are implemented at the application level via policies enforced either by the central server or users. Examples are the popular messaging apps Signal, Telegram and WhatsApp (as of 2024).

- In WhatsApp, only admins can add and remove users, create a group invite link, and govern the admin subgroup. All groups must have at least one admin, and in particular when the last admin leaves, a user is selected randomly as the new admin.
- In Telegram, the group creator can designate other admins with diverse sets of capabilities. Besides adding and removing users, admins can impose partial bans on any user's capabilities, such as sending or receiving messages, and can even restrict the content that users can send [Tel].
- In Signal Messenger, admins can specify whether all members or only admins can add and remove users from a group (in the latter case non-admins can request to add users) and create a group invite link.

Despite administration mechanisms being widely deployed, there is little mention of admins in the literature. Existing CGKA and group messaging approaches make no formal distinction between admin and non-admin users, which results in giving admin capabilities to all users.

Security Goals. There are four main security goals that our cryptographic administrators aim to achieve. In groups where no distinction is made between admins and standard members, our solutions can be extended to the whole group by treating all members as admins; these goals nonetheless still apply:

- Reduce trust on the delivery service, such that it has no control over group administration and membership.
- Mitigate the impact of insider attacks [KS05, AJM22] on protocol execution. Insider adversaries, or compromised group members, will not be able to gain control of a group unless they are administrators¹.

¹Note that denial-of-service attacks from malicious non-admin insiders as in [ACJM20] are not necessarily

- Increase the resilience of implementations of messaging protocols, preventing pitfalls such as the *burgle into a group attack* [RMS18] or the recent attacks on Matrix [ACDJ23].
- Reduce concurrency issues, especially when the delivery service is not a central server [WKHB21], since only a reduced set of members are able to commit group changes.

Admin Capabilities. Let $G = \{ID_1, \dots, ID_n\}$ be a group of users participating in messaging or continuous key exchange and $G^* \subseteq G$ be a non-empty subset of group administrators. Unlike regular group members, the administrators $ID \in G^*$ that we consider can: (1) add and remove members from the group, (2) approve/reject join and removal requests, (3) designate other administrators, (4) give up their admin status, (5) remove the admin status of other users. For performance and security, regular group members should be able to remove themselves and make key updates without admin approval.

These correspond to the common administration features among the solutions presented above. In the case of Telegram, their “fine-grained administration” is practical due to the lack of end-to-end encryption. By default, Telegram relies on a central server that decrypts all messages, which is incompatible with our schemes.

5.1.1 Summary

In this chapter, we cast group administration as a formal *cryptographic* problem. The complexity of secure messaging requires modular constructions and proofs of security; our main goal in this chapter is to provide these. Our core contributions are as follows.

1. We introduce the *administrated CGKA* (A-CGKA) primitive in Section 5.2 by extending the continuous group key agreement (CGKA) primitive.
2. We introduce a novel game-based correctness notion for both CGKA and A-CGKA in Section 5.2.3 that emphasises the role of *group dynamics* which we argue is centrally linked to group administration.
3. Extending existing CGKA key indistinguishability security notions, we introduce a game-based security notion (Section 5.2.4) which further aims to prevent even fully compromised non-admin users from modifying group membership.
4. We present two A-CGKA constructions, *individual admin signatures* (IAS) and *dynamic group signature* (DGS), each built on top of a CGKA protocol. Each approach provides different security and efficiency properties. In Section 5.3, we formalise both protocols in detail and prove their correctness and security. We analyse their performance in Section 5.4.1.

prevented; we also remark that this family of attacks does not affect confidentiality. This issue is discussed in later sections.

Chapter 5. Cryptographic Administration for Secure Group Messaging

5. We propose an extension to MLS in Section 5.3.4 that provides efficient secure administration that we also implement and benchmark locally (Section 5.4.1).

5.1.2 Technical Overview

From CGKA to A-CGKA. Inspired by the MLS standard (and later draft versions in particular), CGKA has been increasingly formalised in the so-called *propose and commit* paradigm [ACJM20, AJM22, ACDT21a]. In CGKA, each user maintains a state which is input to and updated by local CGKA algorithms. Users in a given group can create proposal messages to propose to add or remove users, or to update their keying material for PCS reasons. Proposals are then combined by a member to form a *commit* message, which is ordered and distributed to all group members via a centralised *delivery service*. This is then *processed* by users which make the committed changes effective.

We extend CGKA to A-CGKA to support administration on the primitive level. We support additional proposal types, namely for adding and removing admins, as well as for admin key updates. In A-CGKA, only admins can make admin proposals, and moreover only admins can authorise all types of commit messages. That is, users only process group changes that have been attested by an admin (except for users leaving the group by themselves, which can always be processed).

Correctness. Our notion of A-CGKA correctness enforces that users that process the same sequence of commit messages for a given group derive consistent views of both the evolution of group members and admins, and also of the shared key. Our game explicitly checks that group membership can change only through processing *commit* messages. We also enforce that honest proposals have their intended effect when embedded in commit messages upon processing.

Early CGKA game-based correctness notions were limited to key consistency [ACDT20] or were not formally specified [KPPW⁺21]. It is only recently that group evolution guarantees have been considered, notably in the latest version of Alwen et al. [AJM20, AJM22] which works in the UC and the monolithic security definition of Alwen et al. [ACDT21b]. We emphasise this because, given the length of our IAS construction and corresponding correctness proof, it is possible that subtle bugs concerning group evolution are hidden in existing group messaging constructions and implementations.² In particular, we spotted (resolved) inconsistencies while trying to prove our own protocols correct.

Security. Our security notion captures two core guarantees. Firstly, like previous work [ACDT20, KPPW⁺21], we consider a key indistinguishability game where the adversary drives CGKA execution via oracles and may compromise parties. We prevent the adversary from winning the game trivially by restricting its behaviour using protocol-dependent clean-

²For example, in Figure 8 (SGM) and Figure 17 (RTreeKEM) in Alwen et al. [ACDT21b], commit processing is not well-defined for a user who is processing their own removal from the group.

ness predicates, similar to previous work on messaging and key exchange.

Secondly, differing from ‘standard’ CGKA, we require that the adversary is unable to forge an (admin) commit message that results in a change in group structure for the processing party, even if the adversary knows the group key. We model this by allowing the (semi-active) adversary to inject commit messages to particular parties which process the messages (albeit without updating their state). The adversary can adaptively expose the states of participants and make challenge queries. Security is ensured insofar as the adversary does not trivially compromise an administrator, so in particular they can compromise many non-admins. As in the case of key indistinguishability, we also specify a separate admin cleanness predicate to capture trivial attacks for this attack vector. Our security notion allows for FS and PCS guarantees with respect to the admin keying material.

Constructions. In this chapter, we provide two separate, modular constructions of A-CGKA from a CGKA that we describe in Section 5.3. We also introduce an extension of MLS that supports administration. The security of the authentication mechanisms in all our protocols matches the FS and PCS demands of modern group messaging.

In our first construction, individual admin signatures (IAS) (Section 5.3.1), admins keep track of their own signature key pair. Admin proposals and commits which change the group structure or admin structure or keys are signed using the committing admin’s signature key. Admins update their signature keys via admin update proposals or by crafting commit messages.

Our second construction, dynamic group signature (DGS) (Section 5.3.2), relies on a secondary CGKA and authenticates a group of admins as a whole (possibly all group members). This highlights the fact that administration can be done without authenticating single users. Instead of maintaining individual signatures, admins instead execute within this CGKA and use the common secret to derive a signature key pair for each epoch. Non-admins keep track of the signature public key over time and verify that commits are signed using it.

Proofs. We formally prove that our protocols IAS and DGS are correct and secure with respect to our A-CGKA definitions. We prove IAS secure with respect to a sub-optimal admin cleanness predicate (somewhat weak forward security). We argue that the protocol and proof can be very easily modified to satisfy optimal security using forward-secure signatures [BM99] with no asymptotic overhead; this is discussed in Section 5.3.1. By contrast, in DGS, forward security is inherited from the admin CGKA.

MLS Extension. In Section 5.3.4, we embed the MLS protocol with A-CGKA functionality more organically by making use of MLS’s credential infrastructure. We describe the main modifications needed and propose an extension of MLS that admits secure administration. Moreover, we implement and benchmark the efficiency of our MLS extension (and include a reference to the source code); we present our results in Section 5.4.1.

5.1.3 Additional Related Work

The TreeKEM protocol in MLS, initially proposed by Bhargavan et al. [BBR18], was inspired by Asynchronous Ratchet Trees [CCG⁺18]. Later, variants of TreeKEM arose like Tainted TreeKEM [KPPW⁺21], Insider-Secure TreeKEM [AJM22], Re-randomised TreeKEM [ACDT20], and Causal TreeKEM [Wei19]. MLS as a whole is studied by Brzuska et al. [BCK22] and Alwen et al. [ACDT21a].

CGKAs have been recently used to formally build full group messaging protocols [ACDT21a]. Besides TreeKEM, several CGKA variants have been proposed [ACJM20, AHKM22, WKHB21, AAN⁺22b]. Side works deal with multi-group security [CHK21], efficient key schedules for multiple groups [AAB⁺21], and concurrency [BDR20]. Separately, Rösler et al. [PRSS21] surveys group key exchange protocols. Group admins were considered by Rösler et al. [RMS18], although without a formal cryptographic approach, instead opting for property-based security notions described more informally akin to a symbolic model of security.

An alternative approach towards securing group membership was taken in the Signal Private Group System [CPZ20] which we discuss in Section 5.4.2.

5.2 (Administrated) Continuous Group Key Agreement

In this section, we introduce *continuous group key agreement* (CGKA) and then extend it to formalise our *administrated CGKA* (A-CGKA) primitive. We also introduce our correctness and security definitions for both CGKA and A-CGKA.

5.2.1 Continuous Group Key Agreement

The aim of the continuous group key agreement (CGKA) primitive [ACDT20] is to provide shared secrets (denoted by k) to dynamic groups of users over time. In CGKA, each group, labelled with a group identifier gid , is subject to additions (*add*), removals (*rem*), and user state refreshes/key updates (*upd*).

The definition of CGKA is introduced below, in the so-called *propose and commit* paradigm [AJM22, ACJM20], in which different operation proposals in a given group (for standard CGKA, including *add*, *remove* and *key update* operations) are collated into a commit message by a group member which is processed by users. The evolution of a CGKA in time is captured by *epochs*; a group member advances to a new epoch every time they successfully process a commit message, at which point there may be a change in the shared secret, in which case there may also be a change in the group structure from their perspective.

Note that the primitive is *stateful*: each user keeps their own state γ and calls each of the following algorithms locally which may update the state.

Definition 40. A continuous group key agreement (CGKA) scheme is a tuple of algorithms

5.2 (Administrated) Continuous Group Key Agreement

CGKA = (Init, Create, Prop, Commit, Proc, Prop-Info, Props) such that:

- $\gamma \stackrel{s}{\leftarrow} \text{Init}(1^\lambda, ID)$: The initialisation function takes as inputs a security parameter 1^λ and an identity ID and outputs an initial state γ .
- $(\gamma', T) \stackrel{s}{\leftarrow} \text{Create}(\gamma, \text{gid}, G)$: The group creation algorithm takes as inputs a state γ , a group identifier gid , and a list of group members $G = \{ID_1, \dots, ID_n\}$ and outputs a new state γ' and a control (welcome) message T , where $T = \perp$ indicates failure.
- $(\gamma', P) \stackrel{s}{\leftarrow} \text{Prop}(\gamma, \text{gid}, ID, \text{type})$: The proposal algorithm takes as input a state, a group identifier, an ID , and a proposal type $\text{type} \in \text{types} = \{\text{add}, \text{rem}, \text{upd}\}$, and outputs a new state γ' and a proposal message P , where $P = \perp$ indicates failure.
- $(\gamma', T, k) \stackrel{s}{\leftarrow} \text{Commit}(\gamma, \text{gid}, \vec{P})$: The commit algorithm takes as inputs a state, a group identifier, and a vector of proposals \vec{P} , and outputs a new state γ' , a control message T where $T = \perp$ indicates failure, and the (possibly new) group secret k .
- $(\gamma', \text{acc}) \leftarrow \text{Proc}(\gamma, T)$: The processing algorithm takes as inputs state and a control message T , and outputs a new state γ' and an acceptance bit $\text{acc} \in \{\text{true}, \text{false}\}$, where $\text{acc} = \text{false}$ indicates failure.
- $(\text{gid}, \text{type}, ID, ID') \leftarrow \text{Prop-Info}(\gamma, P)$: The proposal information algorithm takes as inputs a state and a proposal P , and outputs the group identifier of the proposal gid , its type type , the ID of the user *affected* by the proposal and the proposal creator ID' .
- $\vec{P} \leftarrow \text{Props}(\gamma, T)$: The proposal extractor algorithm takes as inputs a state and control message T and outputs the vector of proposals \vec{P} associated with T , where $\vec{P} = \perp$ indicates failure.

Finally, given ID 's state γ and gid , the (possibly empty) set of group members in gid from ID 's perspective is stored as $\gamma[\text{gid}].G$, and the group secret k for gid is $\gamma[\text{gid}].k$.

Protocol Execution. For simplicity, we assume that all users and groups are associated with a unique identifier ID and gid , respectively, which is, e.g., enforced by the PKI. Once every user has initialised their state using `Init`, a group is created when some party calls `Create` with some gid and a list of ID s as inputs. The `Init` algorithm can also serve to authenticate and register keys on a PKI when appropriate [ACDT20, KPPW⁺21, ACJM20]. In Section 5.3, we expand on authentication: for each protocol, we describe our assumptions on the PKI. The `Create` algorithm outputs a control message T that must be processed by prospective group members, including the group creator, to join the group gid .

In our formalism, any user can propose a member addition (`add`), member removal (`rem`) or key update (`upd`, only available for the caller) at any time. This is done via the `Prop` method, which outputs a proposal message P . Proposals encode the information needed to make a

Chapter 5. Cryptographic Administration for Secure Group Messaging

change in the group structure or keying material, but the encoded changes are not immediately applied to the group. We emphasise that only the caller of Prop can use argument type = upd to propose to update (i.e., refresh) their keying material, in which case the input ID is ignored. Following Alwen et al. [ACDT21a], we define Prop-Info which outputs proposal attributes, rather than allowing for their direct access, to support possibly encrypted proposals (e.g., as in a PrivateMessage in the MLS standard [BBR⁺23]).

Proposed changes become effective once a user commits a (possibly empty) vector of proposals $\vec{P} = (P_1, \dots, P_m)$ using Commit; we assume that proposals are suitably propagated in the group. The Commit algorithm outputs the new group key k and a control message T that contains the information needed by all current and incoming group members to process the changes. Typically, the Commit algorithm also updates the keying material of the caller. Control messages are processed via Proc, which updates the caller's state and outputs a bit acc indicating success or failure. We note that Proc and Props do not require a group identifier as input; this models the standard behaviour of a messaging protocol where, upon reception of a message, the user needs to determine which group the message corresponds to. In the event that a group member needs to parse the proposals in a commit message T without processing it, it can do so via the Props algorithm.

Example. Consider 5 parties $\{ID_1, \dots, ID_5\}$ executing a CGKA protocol. After they each initialise their states as $\gamma_i \stackrel{\$}{\leftarrow} \text{Init}(1^\lambda, ID_i)$, the following actions take place:

1. ID_1 calls $\text{Create}(\gamma_1, \text{gid}, \{ID_1, ID_2, ID_3, ID_4\})$, which updates γ_1 and outputs a control message T_0 . At this stage, the group is still empty, $G = \emptyset$.
2. Each ID_i (including ID_1) processes the group creation as $\text{Proc}(\gamma_i, T_0)$, which updates each state γ_i and outputs $\text{acc} = \text{true}$ to each user. At this stage, $G = \{ID_1, ID_2, ID_3, ID_4\}$, and the group members share a common secret k_1 .
3. Several users propose changes in the group:
 - ID_2 proposes to add ID_5 to the group by calling $(\gamma_2, P_1) \stackrel{\$}{\leftarrow} \text{Prop}(\gamma_2, \text{gid}, ID_5, \text{add})$.
 - ID_3 wants to update its keying material and thus calls $(\gamma_3, P_2) \stackrel{\$}{\leftarrow} \text{Prop}(\gamma_3, \text{gid}, ID_3, \text{upd})$.
 - ID_1 proposes to remove ID_4 from the group and calls $(\gamma_1, P_3) \stackrel{\$}{\leftarrow} \text{Prop}(\gamma_1, \text{gid}, ID_4, \text{rem})$.

This is shown in Figure 5.1. The group remains the same.

4. ID_2 collates all proposals in a commit message by calling $(\gamma_2, T_1, k_2) \stackrel{\$}{\leftarrow} \text{Commit}(\gamma_2, \text{gid}, (P_1, P_2, P_3))$. The group remains the same, since parties have not yet processed T_1 .
5. All parties process T_1 by calling $\text{Proc}(\gamma_i, T_1)$ and updating their states. Now, $G = \{ID_1, ID_2, ID_3, ID_5\}$ and these members share a new common secret k_2 , which is not

5.2 (Administrated) Continuous Group Key Agreement

known to ID_4 . In addition, ID_3 (due to the update) and ID_2 (due to the commit) have refreshed their CGKA keying material for FS and PCS.




Figure 5.1: Diagram of a sample CGKA execution with 5 parties. Parties in green (ID_2, ID_3, ID_5) will update their keying material after ID_2 commits, while ID_4 will leave the group.

Commit Semantics. We assume that proposals input to Commit are processed in some deterministic, publicly-known, a priori determined order, that we call the *policy*. It is possible to extend the syntax of A-CGKA with a dedicated policy algorithm that defines this order [Bal21].

Alternative Definitions. In some CGKA definitions (especially older ones [ACDT20, KPPW⁺21]) and older MLS drafts, group changes are made effective immediately by processing proposals. To this end, the Commit and Prop algorithms are replaced by specific ‘action’ algorithms such as Add, Remove, Update. The *propose and commit paradigm* used in this chapter was introduced in draft 8 of MLS to allow for multiple operations to be applied at the same time to improve latency and concurrency support [BBR⁺23]. As mentioned in [KPPW⁺21], the older protocols can be written in the propose and commit paradigm, which is more flexible (and also better suited for group administrators as we will see).

A relevant difference between the definition in Alwen et al. [ACDT20] (and other game-based formulations such as in [KPPW⁺21]) and ours is that we work in the multi-group setting, and so we consider group identifiers of the form gid . Multi-group CGKAs have not been fully formalised in the literature, although multi-group security has been studied [CHK21], as well as efficient key schedules for multiple groups [AAB⁺21]. According to our definition, a user can be in many groups, identified by different values gid and interleave operations from each group arbitrarily. There are formulations of CGKA in the universal composability (UC) model [ACJM20, AJM22] which use group identifiers. In fact, composability guarantees in the UC model rely on the existence of unique, a priori established ‘session identifiers’ [Can01, KT11]; these can be established in practice via a central server or a distributed protocol [BLR04]. Note Cremers et al. [CHK21] considers cross-group security for messaging but does not treat CGKA as a primitive formally.

Finally, we note that there are other small differences in the literature. The semantics of Tainted TreeKEM [KPPW⁺21] enable users to speculatively execute operations; our syntax could be

Chapter 5. Cryptographic Administration for Secure Group Messaging

modified to support this. In works on CGKA such as Alwen et al. [ACDT20], it is typical to make a distinction between standard commit and welcome messages. We implicitly incorporate this distinction in our constructions, but avoid it in the primitive syntax for simplicity. In some works [AJM22, ACJM20], additional algorithms such as GetKey are provided; we treat k as a state variable instead, which is output by Commit. As mentioned, it is also possible to conceive a policy algorithm as in [Bal21] that determines what the effect of a sequence of proposals should be (e.g., which proposals are prioritised, removing duplicates, etc.). Groups in some works [ACDT21a, AJM22], as well as the MLS standard itself, are initially of size 1: since initialising groups with several participants at once is common in practice, we believe it is of interest to model it. Different works formalise the role of the PKI to different degrees: Alwen et al. and Alwen et al. [ACDT21a, AJM22] consider an explicitly modelled PKI where users can choose their own possibly malicious keying material. Nevertheless, all works on CGKA to date formally assume that the PKI acts consistently for all users.

5.2.2 Administrated CGKA

An administrated continuous group key agreement (A-CGKA) is a CGKA where only a group G^* of ID 's, the so-called *group administrators*, can commit (and therefore make effective) changes to the group structure, such as adding and removing users. As with the group of users G in both CGKA and A-CGKA, the group of administrators G^* is dynamic.

Definition 41. An administrated continuous group key agreement (A-CGKA) scheme is a tuple of algorithms $A\text{-CGKA} = (\text{Init}, \text{Create}, \text{Prop}, \text{Commit}, \text{Proc}, \text{Prop-Info}, \text{Props})$ such that:

- Algorithms $\text{Init}, \text{Proc}, \text{Prop-Info}, \text{Props}$ are defined as for a CGKA (Definition 40).
- In Prop and Prop-Info , types is redefined as $\text{types} = \{\text{add}, \text{rem}, \text{upd}, \text{add-adm}, \text{rem-adm}, \text{upd-adm}\}$.
- $(\gamma', T) \stackrel{\$}{\leftarrow} \text{Create}(\gamma, \text{gid}, G, G^*)$ additionally takes as input a group of admins G^* .
- $(\gamma', T, k) \stackrel{\$}{\leftarrow} \text{Commit}(\gamma, \text{gid}, \vec{P}, \text{com-type})$ additionally takes as input a commit type $\text{com-type} \in \text{com-types} = \{\text{std}, \text{adm}, \text{both}\}$.

Given ID 's state γ and gid , $\gamma[\text{gid}].G$ and $\gamma[\text{gid}].k$ are defined as in Definition 40, and $\gamma[\text{gid}].G^*$ stores the set of admins in gid from ID 's perspective.

The execution of an A-CGKA is analogous to CGKA. Besides the introduction of the group of admins G^* , we introduce three additional proposal types add-adm , rem-adm , and upd-adm which concern administrative changes. Namely, an admin can propose to add another admin to the group of administrators, revoke the admin capabilities from a party, or update their administrative keying material, respectively. The commit type com-type specifies the scope of a commit operation, that is, whether it affects the general group (std), the administration of

5.2 (Administrated) Continuous Group Key Agreement

the group (adm), or both at the same time (both). For the latter, a simple example is when an admin is both adding a member (group modification) and refreshing its admin keys (admin modification).

We note that the Create algorithm, enforced by our correctness and security notions to come, will require the condition $\emptyset \subset G^* \subseteq G$. Thus, the group administrators are always a subset of the group members. We take this approach following previous CGKAs [ACDT20, KPPW⁺21, ACJM20] and group messaging protocols [BBR⁺23, ACDT21a] where only group members can perform commits or make changes in the group. In these works and ours, it is (computationally) impossible for an external user to administrate a group, since external commits are not permitted.

Real-World Administrators. A-CGKA captures the core admin features of popular applications such as WhatsApp and Signal as mentioned in the introduction. We remark that the fact that non-admins are not allowed to make changes (except for leaving a group) is a desired consequence of our formulation of A-CGKA. A more fine-grained conceivable solution could be to allow admins to send a policy change proposal, to, e.g., modify the ability of all members to call Commit to add new users.

5.2.3 Correctness

Due to their similarity, we define the correctness of CGKA and A-CGKA together. Correctness of an (A)-CGKA scheme (A)-CGKA under the notion CORR is defined by game $\text{CORR}_{(A)\text{-CGKA}, \text{C}_{\text{corr}}}$ played by adversary \mathcal{A} in Figure 5.2. The main properties captured by the game are the following:

- *View consistency:* All users who transition to the same epoch (i.e., those which process the same sequence of commit messages) have the same group view (i.e., G , G^* and key k).
- *Message processing:* The group structure (G , G^*) and k can only be modified due to calls to Proc.
- *Forking states:* If the group is partitioned into subgroups that process different sequences of commit messages (thus leading to different group views), the game ensures that members in each partition have consistent views.
- *Multiple groups:* The adversary may create groups via CREATE on behalf of different users, and interact with different IDs in multiple groups.

Separately, we ensure that a user's state is not modified whenever a particular algorithm call fails. As observed for two-party messaging [BSJ⁺17], we require incorrect inputs to not affect the functionality of the protocol, and in particular to not cause a denial of service.

Chapter 5. Cryptographic Administration for Secure Group Messaging

<p>Game $\text{CORR}_{(A)\text{-CGKA}, C_{\text{corr}}}(\mathcal{A})$</p> <pre> 1: public ep-view[·], ST[·], T[·] ← ⊥ 2: public first-crt[·] ← ⊥ 3: public prop-ctr, com-ctr ← 0 // msg counters 4: public ep[·] ← (-1, -1) // user epoch tracker 5: win ← 0 6: ST[ID] $\stackrel{\\$}{\leftarrow}$ Init($1^\lambda, ID$) $\forall ID$ 7: $\mathcal{A}^\mathcal{O}$ 8: require C_{corr} // optional predicate 9: return win // 1 if \mathcal{A} is rewarded </pre>	<p>Oracle $\text{PROP}(ID', \text{gid}, ID, \text{type})$</p> <pre> 1: require type \in types 2: $(\gamma, P) \stackrel{\\$}{\leftarrow}$ Prop(ST[ID'], gid, ID, type) 3: if $P = \perp$: return // failure 4: $(\text{gid}^*, \text{type}^*, ID^*, ID'^*) \leftarrow$ Prop-Info(γ, P) 5: reward $(\text{gid}^*, \text{type}^*, ID^*, ID'^*) \neq (\text{gid}, \text{type}, ID, ID')$ 6: CheckSameGroupState(ST[ID'], γ, gid) 7: $T[\text{gid}, \text{ep}[\text{gid}, ID'], \text{prop}, ++\text{prop-ctr}] \leftarrow P$ 8: $\text{ST}[ID'] \leftarrow \gamma$ // upd. ST of proposer ID' </pre>
<p>Oracle $\text{CREATE}(ID, \text{gid}, G, G^*)$</p> <pre> 1: $(\gamma, T) \stackrel{\\$}{\leftarrow}$ Create(ST[ID], gid, G, G^*) 2: if $T = \perp$: return 3: reward $\neg(\emptyset \neq G^* \subseteq G)$ 4: CheckSameGroupState(ST[ID], γ, gid) 5: $T[\text{gid}, (-1, -1), \text{com}, ++\text{com-ctr}] \leftarrow T$ 6: $\text{ST}[ID] \leftarrow \gamma$ </pre>	<p>Oracle $\text{COMMIT}(ID, \text{gid}, I, \text{com-type})$</p> <pre> // $I = (i_1, \dots, i_k)$ for some k 1: require com-type \in com-types 2: require $\text{ep}[\text{gid}, ID] \notin \{(-1, -1), \perp\}$ 3: $\vec{P} \leftarrow (T[\text{gid}, \text{ep}[\text{gid}, ID], \text{prop}, i])_{i \in I}$ 4: $(\gamma, T, k) \stackrel{\\$}{\leftarrow}$ Commit(ST[ID], gid, \vec{P}, com-type) 5: if $T = \perp$: return // failure 6: reward $ID \notin \text{ST}[ID][\text{gid}].G$ // no external comm. 7: CheckSameGroupState(ST[ID], γ, gid) 8: $T[\text{gid}, \text{ep}[\text{gid}, ID], \text{com}, ++\text{com-ctr}] \leftarrow T$ 9: $T[\text{gid}, \text{ep}[\text{gid}, ID], \text{vec}, \text{com-ctr}] \leftarrow \text{Props}(\gamma, T)$ 10: $T[\text{gid}, \text{ep}[\text{gid}, ID], \text{key}, \text{com-ctr}] \leftarrow k$ 11: $\text{ST}[ID] \leftarrow \gamma$ </pre>
<p>Oracle $\text{DELIVER}(ID, \text{gid}, (t, c), c')$</p> <pre> 1: require $\text{ep}[\text{gid}, ID] \in \{(t, c), (-1, -1), \perp\}$ 2: $T \leftarrow T[\text{gid}, (t, c), \text{com}, c']$ // honest delivery 3: $(\gamma, \text{acc}) \leftarrow$ Proc(ST[ID], T) 4: if $\neg \text{acc}$: return // failure 5: reward $\text{Props}(\text{ST}[ID], T) \neq T[\text{gid}, (t, c), \text{vec}, c']$ 6: reward $\neg(\emptyset \neq \gamma[\text{gid}].G^* \subseteq \gamma[\text{gid}].G)$ 7: if $(t, c) = (-1, -1)$: // create msg 8: UniqueCreatePerGID(γ, gid, c') 9: if $ID \notin \gamma[\text{gid}].G$: // ID removed 10: $\text{ep}[\text{gid}, ID] \leftarrow \perp$ 11: reward $\gamma[\text{gid}].k \neq \perp$: // key deleted 12: else : // ID in group 13: UpdateView($\gamma, \text{gid}, t, c'$) 14: reward $\gamma[\text{gid}].k \neq T[\text{gid}, (t, c), \text{key}, c']$ 15: $\text{ep}[\text{gid}, ID] \leftarrow (t+1, c')$ 16: $\text{ST}[ID] \leftarrow \gamma$ </pre>	<p>UpdateView($\gamma, \text{gid}, t, c'$)</p> <pre> 1: $v \leftarrow \text{ep-view}[\text{gid}, t+1, c']$ 2: if $v = \perp$: $\text{ep-view}[\text{gid}, t+1, c'] \leftarrow \gamma$ 3: else : CheckSameGroupState(v, γ, gid) </pre>
	<p>CheckSameGroupState($\gamma_1, \gamma_2, \text{gid}$)</p> <pre> 1: reward $\gamma_1[\text{gid}].k \neq \gamma_2[\text{gid}].k$ 2: reward $\gamma_1[\text{gid}].G \neq \gamma_2[\text{gid}].G$ 3: reward $\gamma_1[\text{gid}].G^* \neq \gamma_2[\text{gid}].G^*$ </pre>
	<p>UniqueCreatePerGID(γ, gid, c')</p> <pre> 1: if first-crt[gid] $\neq \perp$: 2: require $c' = \text{first-crt}[\text{gid}]$ 3: else : first-crt[gid] $\leftarrow c'$ </pre>

Figure 5.2: Correctness game for (A)-CGKA with respect to predicate C_{corr} for $\mathcal{O} = \{\text{CREATE}, \text{DELIVER}, \text{PROP}, \text{COMMIT}\}$. Highlighted code is executed only when considering an A-CGKA. Note that when **reward** P is true for predicate P , variable win is set to 1.

5.2 (Administrated) Continuous Group Key Agreement

Overview. The game starts by setting up several public dictionaries. The main two are $ST[\cdot]$, which is a dictionary indexed by ID which keeps the states of each of the users ID throughout the game; and $T[\cdot]$, which keeps all control messages and proposals generated by the (A)-CGKA algorithms. A message T stored in T is indexed by the corresponding gid, epoch, message type (prop, com or vec, standing for proposal, commit or proposal vector, respectively), and a message counter prop-ctr or com-ctr; keys output by Commit for some gid and epoch are also stored in T and are marked with key. After initialisation, we let the adversary \mathcal{A} interact with the oracles with respect to multiple groups. The variable win is set to 1 if one of the **reward** clauses is true, which leads to \mathcal{A} winning given the (optional) predicate C_{corr} is also true when \mathcal{A} finishes executing.

Epochs. Control messages output by successful Create and Commit calls are labelled uniquely by the challenger. For correctness, an *epoch* is a pair (t, c) , where t is an integer relative to a particular group and party which increments upon each successful Proc call while in the group, and c is the value of the global variable com-ctr when the corresponding control message was output. For a given group, each party's epoch value is stored in $ep[\cdot]$ and initialised to $(-1, -1)$, and is set to \perp when they leave a group. To this end, we model correctness in the presence of an adversary who maintains arbitrary network partitioning, so long as they provide a consistent view of messages to parties in each 'partition'.

Group Consistency. We enforce that, for each group member, that each group is only (possibly) updated upon a successful call to Proc (via CheckSameGroupState). For A-CGKA, we ensure that, for a group gid, $\emptyset \neq G^* \subseteq G$ must hold at all times. In Proc, we enforce that all users who transition to the same epoch have the same view of the group and set of admins when relevant (via UpdateView). To this end, the dictionary ep-view stores the state of the first party who transitions to a given epoch (t, c) for a gid, and the state of subsequent parties who transition to the same epoch are checked against the value stored in ep-view.

We also require that, even if there are multiple calls to Create, only one of them is processed by group members (i.e., states do not fork from the initial epoch). To achieve this, the variable first-ctr tracks the commit number of the first successfully processed create message for a given gid. This check is made in UniqueCreatePerGID.

Key Partnering. Note that new epoch keys are derived upon successful Proc calls, and that a new key k (for group members) is always derived in this case. Correctness ensures that all users who transition to the same epoch derive the same key k . The consistency between the key output by Commit and the actual epoch keys is also verified. Moreover, whenever a user derives a key $k \neq \perp$, they must be a group member (line 11 of DELIVER).

Liveness. We enforce that some algorithms, such as Prop, always succeed on 'valid' input, since without such a check, an (A)-CGKA with algorithms that always fail is considered correct. One example of a liveness check is an equality check between proposals output by Props and the input proposals in COMMIT. Since the precise semantics of (A)-CGKA vary between applications, additional checks are delegated to a correctness predicate C_{corr} which, in part,

Chapter 5. Cryptographic Administration for Secure Group Messaging

parameterises the correctness game. In the previous example, C_{corr} may depend on the protocol policy for the application of proposals. Without extra checks, the predicate should be set to $C_{\text{corr}} = \text{true}$.

Definition 42 ($\text{CORR}_{(A)\text{-CGKA}}$). Consider the correctness game CORR presented in Figure 5.2. A CGKA (resp. A-CGKA) is *correct* if, for all $\lambda \in \mathbb{N}$ and all computationally unbounded adversaries \mathcal{A} , we have $\Pr[\text{CORR}_{(A)\text{-CGKA}, C_{\text{corr}}}(\mathcal{A}) \Rightarrow 1] = 0$.

Naturally, this notion can be relaxed to consider correctness that holds with high probability w.r.t. a computationally-bounded adversary (e.g., to capture decryption failures when using lattice-based cryptography as done in Chapter 3).

5.2.4 Security

A-CGKA is a primitive that extends the functionality of a CGKA, which allows a group to establish a sequence of group keys, to provide secure administration. Therefore, any A-CGKA construction must satisfy at least ‘standard’ CGKA security (key indistinguishability).

The main additional goal of A-CGKA over standard CGKA is to prevent unauthorised (standard) users from deciding on changes to a group, capturing the security of the group structure. Note that an A-CGKA in which the adversary fully controls a standard group member is not secure with respect to key indistinguishability, but should still prevent unauthorised group changes. We define (A)-CGKA security in Definition 43 after describing our notion below.

Overview. At its core, the security game in Figure 5.3 is a key indistinguishability game that captures the security of the common group secret for a single group (we implicitly assume a fixed gid), extending the game in [ACDT20]. The game considers a partially active adversary who can make forgery attempts and schedule messages but cannot totally control message delivery. Namely, the adversary can inject a control message to a specific party ID , but this message is not stored in the array T that keeps track of all honestly generated messages after Proc is called. The main consequence of this is that injected proposals cannot be included in commits generated by the challenger. Nevertheless, the adversary can make the challenger commit on vectors of proposals created via PROP .

Informally, the adversary can win the game if it plays a clean game where it either 1) correctly guesses the challenge bit by distinguishing between correct and uniformly sampled keys or (for A-CGKA) 2) manages to forge a message which, after being processed by a user, changes its view of (G, G^*) . We assume that the dependency on the PKI is implicit in the game; we describe the PKI functionality we assume for each protocol as they are introduced. A detailed description follows.

Epochs. Messages output by successful Create , Commit , and Prop calls are uniquely labelled by the challenger via counters prop-ctr , com-ctr . Whenever such a call is made, the corresponding messages are stored in variable T with (incremented) last argument $++\text{com-ctr}$.

5.2 (Administrated) Continuous Group Key Agreement

<p>Game KIND_{(A)-CGKA, C_{cgka}, C_{adm}, C_{forgery}}(\mathcal{A})</p> <ol style="list-style-type: none"> 1: $b \xleftarrow{\\$} \{0, 1\}$ 2: $K[\cdot], ST[\cdot] \leftarrow \perp$ 3: public $T[\cdot], G[\cdot], ADM[\cdot] \leftarrow \perp$ 4: $\text{prop-ctr}, \text{com-ctr}, \text{exp-ctr} \leftarrow 0$ 5: $\text{ep}[\cdot], \text{exp}[\cdot] \leftarrow (-1, -1); C[\cdot] \leftarrow -1$ 6: $\text{chall}[\cdot], \text{forged} \leftarrow \text{false}$ 7: $ST[ID] \xleftarrow{\\$} \text{Init}(1^\lambda, ID) \forall ID$ 8: $b' \xleftarrow{\\$} \mathcal{A}^\emptyset$ 9: if $\neg C_{\text{cgka}} \wedge \neg \text{forged}$ 10: $b'' \xleftarrow{\\$} \{0, 1\}; \text{return } b''$ 11: else return $\mathbb{1}[b = b' \vee \text{forged}]$ <p>Oracle CREATE(ID, G, G^*)</p> <ol style="list-style-type: none"> 1: $(\gamma, T) \xleftarrow{\\$} \text{Create}(ST[ID], G, G^*)$ 2: if $T = \perp$ return // failure 3: $T[(-1, -1), \text{com}, ++\text{com-ctr}] \leftarrow (T, \text{both})$ 4: $ST[ID] \leftarrow \gamma$ <p>Oracle PROP(ID, ID', type)</p> <ol style="list-style-type: none"> 1: $(\gamma, P) \xleftarrow{\\$} \text{Prop}(ST[ID], ID', \text{type})$ 2: $T[\text{ep}[ID], \text{prop}, ++\text{prop-ctr}] \leftarrow P$ 3: $ST[ID] \leftarrow \gamma$ <p>Oracle COMMIT($ID, (i_1, \dots, i_k), \text{com-type}$)</p> <ol style="list-style-type: none"> 1: $\vec{P} \leftarrow (T[\text{ep}[ID], \text{prop}, i])_{i=(i_1, \dots, i_k)}$ 2: $(\gamma, T, k) \xleftarrow{\\$} \text{Commit}(ST[ID], \vec{P}, \text{com-type})$ 3: if $T = \perp$ return // failure 4: $T[\text{ep}[ID], \text{com}, ++\text{com-ctr}] \leftarrow (T, \text{com-type})$ 5: $T[\text{ep}[ID], \text{vec}, \text{com-ctr}] \leftarrow \text{Props}(ST[ID], T)$ 6: $(t_s, t_a) \leftarrow \text{ep}[ID]$ 7: if $\text{com-type} = \text{adm}$ $t_s \leftarrow t_s - 1$ 8: $\mathcal{K}[t_s + 1] \leftarrow k; ST[ID] \leftarrow \gamma$ <p>Oracle EXPOSE(ID)</p> <ol style="list-style-type: none"> 1: $\text{exp}[ID, ++\text{exp-ctr}] \leftarrow \text{ep}[ID]$ 2: return $ST[ID]$ 	<p>Oracle DELIVER($ID, (t_s, t_a), c$)</p> <ol style="list-style-type: none"> 1: require $\text{ep}[ID] \in \{(t_s, t_a), (-1, -1)\}$ 2: $(T, \text{com-type}) \leftarrow T[(t_s, t_a), \text{com}, c]$ // honest deliv. 3: if $C[(t_s, t_a)] \in \{c, -1\}, C[(t_s, t_a)] \leftarrow c$ 4: else return // bad commit for epoch 5: $(\gamma, \text{acc}) \leftarrow \text{Proc}(ST[ID], T)$ 6: if $\neg \text{acc}$ return // failure 7: if $ID \notin \gamma.G$ // ID removed 8: $\text{ep}[ID] \leftarrow (-1, -1)$ 9: else // ID in group, update dictionaries 10: $\text{ep}[ID] \leftarrow (t_s, t_a)$ 11: if $\text{com-type} \in \{\text{std}, \text{both}\}$ 12: $K[t_s + 1] \leftarrow \gamma.k$ 13: $G[t_s + 1] \leftarrow \gamma.G$ 14: $\text{ep}[ID] \leftarrow \text{ep}[ID] + (1, 0)$ 15: if $\text{com-type} \in \{\text{adm}, \text{both}\}$ 16: $ADM[t_s + 1] \leftarrow \gamma.G^*$ 17: $\text{ep}[ID] \leftarrow \text{ep}[ID] + (0, 1)$ 18: $ST[ID] \leftarrow \gamma$ <p>Oracle CHAL(t_s)</p> <ol style="list-style-type: none"> 1: require $(K[t_s] \neq \perp) \wedge \neg \text{chall}[t_s]$ 2: $\text{chall}[t_s] \leftarrow \text{true}$ 3: if $b = 0$ return $K[t_s]$ 4: if $b = 1$ return $r \xleftarrow{\\$} \{0, 1\}^\lambda$ <p>Oracle REVEAL(t_s)</p> <ol style="list-style-type: none"> 1: require $(K[t_s] \neq \perp) \wedge \neg \text{chall}[t_s]$ 2: $\text{chall}[t_s] \leftarrow \text{true}$ 3: return $K[t_s]$ <p>Oracle INJECT(ID, m, t_a)</p> <ol style="list-style-type: none"> 1: require $C_{\text{adm}} \wedge (\text{ep}[ID] = (\cdot, t_a)) \wedge (t_a \neq -1)$ 2: require $(m, \cdot) \notin T$ // external forgery 3: $(\gamma, \perp) \leftarrow \text{Proc}(ST[ID], m)$ 4: if C_{forgery} 5: $\text{forged} \leftarrow \text{true}$ // successful forgery 6: return b // adversary wins 7: else return \perp
-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

Figure 5.3: Single-group key indistinguishability (KIND) security game for (A)-CGKA, parametrised by predicate C_{cgka} (resp. and predicates C_{adm} and C_{forgery} for A-CGKA) for $\mathcal{O} = \{\text{CREATE}, \text{PROP}, \text{COMMIT}, \text{CHAL}, \text{DELIVER}, \text{REVEAL}, \text{EXPOSE}\}$ (resp. and INJECT). Highlighted code is executed only when considering an A-CGKA.

Chapter 5. Cryptographic Administration for Secure Group Messaging

The evolution of the group after parties process such control messages is modelled using epochs (differing from the epochs considered for correctness), which are each represented as an integer t_s (for CGKA) or a pair of integers (t_s, t_a) (for A-CGKA). The *standard epoch* t_s represents the time between two successive key evolutions, where a different key should be derived in each t_s . The *administrative epoch* t_a represents the time between two changes in the group administration (i.e., between two sequences of simultaneous admin updates, adds and/or removals).

For CGKA, epochs (t_s) advance every time a commit is processed. For A-CGKA, t_s advances if the commit type $\text{com-type} \in \{\text{std}, \text{both}\}$ and t_a advances if $\text{com-type} \in \{\text{adm}, \text{both}\}$. Group members can be in different epochs, captured by the variable $\text{ep}[ID]$ which stores the current standard/admin epoch pair for a given group member ID . If a participant ID is not in the group, then $\text{ep}[ID] = -1$ (CGKA) or $\text{ep}[ID] = (-1, -1)$ (A-CGKA) holds.

Challenges. At any point in the game, the adversary can challenge with respect to a standard epoch t_s by calling CHAL. In a challenge, the adversary is given the group key $\mathcal{K}[t_s]$ if the challenger's bit is $b = 0$, and a random string $r \xleftarrow{\$} \{0, 1\}^\lambda$ if $b = 1$. The adversary must try to determine the value of b by outputting a guess b' of b . A given execution is considered valid when either the *standard cleanness predicate* C_{cgka} is true or, for A-CGKA, the adversary makes a forgery and the *admin cleanness predicate* is true. Cleanness ensures that no trivial attacks on the game are possible; we elaborate on this below.

Exposure Mechanisms. In order to capture group key ratcheting (for FS and PCS), the adversary has two mechanisms to obtain secret group material: it can *expose* a user ID using EXPOSE and *reveal* the group secret k using REVEAL. An exposure leaks the entire current state of ID stored in $\text{ST}[ID]$. We keep track of the specific epochs in which each ID was exposed using the $\text{exp}[\cdot]$ variable. On the other hand, a reveal leaks the group key to the adversary on a specified epoch t_s , in which case $\text{chall}[t_s]$ is set to true to prevent the adversary from challenging on t_s (conversely, the reveal fails when $\text{chall}[t_s] = \text{true}$).

Injections. For A-CGKA, the adversary can also win the game by successfully *injecting* a forged commit. An injection can be attempted by calling $\text{INJECT}(ID, m, t_a)$, given that ID is in admin epoch t_a , where ID is the target group member and m is the forged message. Note we require $t_a \neq -1$ since the adversary could otherwise trivially invite a new user into a new group that it controls. Forgeries can only be attempted if the *administrative predicate* C_{adm} is not violated. As discussed below, C_{adm} captures administration security by excluding trivial attacks. A trivial scenario excluded by any predicate is when the adversary has exposed an administrator immediately before a forgery attempt. The adversary wins the game if the forgery is accepted by any group member $ID \in G$ and if the *forgery predicate* C_{forgery} that we define below is true.

CGKA Cleanness Predicate. The security game in Figure 5.3 is parametrised by three cleanness predicates, C_{cgka} , C_{adm} and C_{forgery} . The first predicate C_{cgka} follows approaches like [BRV20, ACDT20] to parametrise the security of the common (A)-CGKA key. Namely, this pred-

5.2 (Administrated) Continuous Group Key Agreement

icate excludes trivial attacks on indistinguishability, i.e., those that break security unavoidably such as exposing a user and issuing a challenge before its keying material has been updated. Further, it captures the exact security of the protocol (with respect to key indistinguishability), which in our case comprises forward security and post-compromise security after updates. If an independent CGKA is used to construct an A-CGKA, the predicate C_{cgka} may mostly depend on the security of the CGKA. An example, as we show in later sections, is our second construction DGS.

A more fine-grained characterisation of this predicate is to write it as $C_{\text{cgka}} = C_{\text{cgka-opt}} \wedge C_{\text{cgka-add}}$, where $C_{\text{cgka-opt}}$ is an *optimal*, generic cleanness predicate that excludes only unavoidable trivial attacks, and $C_{\text{cgka-add}}$ is an additional cleanness predicate that depends on the scheme and excludes other attacks. We define $C_{\text{cgka-opt}}$ in a similar way to the safe predicate in [ACDT20]. Namely, we exclude the following cases: (1) the group secret in challenge epoch t_s^* was already challenged or revealed, and (2) a group member ID whose state was exposed in standard epoch $t_{\text{exp}} \leq t_s^*$ did not update their keys (i.e., processed their own commit, or processed a commit in which they were the subject of an add, remove, or update proposal) or was not removed before the challenge epoch t_s^* . The optimal cleanness predicate is given in Figure 5.4 for an adversary that makes oracle queries q_1, \dots, q_n in the game.

$$\begin{array}{l}
 C_{\text{cgka-opt}} : \forall (i, ID, \text{ctr} \in (0, \text{exp-ctr}) : q_i = \text{CHAL}(t_i^*), \\
 (ID \notin G[t_i^*]) \vee \\
 (\exists (t_i, c) : (\text{tExp}(ID, \text{ctr}).t_s < t_i \leq t_i^*) \wedge \text{hasUpd}_{\text{std}}(ID, T[(t_i, \cdot), \text{com}, c], T[(t_i, \cdot), \text{vec}, c]) \wedge (C[(t_i, \cdot)] = c)) \\
 \vee (t_i^* < \text{tExp}(ID, \text{ctr}).t_s).
 \end{array}$$

Figure 5.4: Optimal CGKA predicate where the adversary makes oracle queries q_1, \dots, q_n .

The predicate is the logical disjunction of three clauses: for every exposure, adversarial challenge, and party ID , we require that either (1) ID was not a group member at the challenge time; (2) the challenge epoch precedes the exposure (forward security); or (3) ID updated between the exposure and the (subsequent) challenge (post-compromise security). To avoid cluttering the predicate, our game already enforces that only one challenge or reveal can be performed per epoch (which is optimal for our game).

We have used the following auxiliary functions. The function tExp is such that $\text{tExp}(ID, \text{ctr}) = \text{exp}[ID, \text{ctr}]$ if $\exists k : q_k = \text{EXPOSE}(ID)$, and -1 (for CGKA) or $(-1, -1)$ (for A-CGKA) otherwise. Given \vec{P} , the function $\text{hasUpd}_{\text{std}}(ID, (T, \text{com-type}), \vec{P})$ (sans com-type for CGKA) outputs true if either: (i) ID has processed a commit of his own, where $\text{com-type} \in \{\text{std}, \text{both}\}$, or (ii) ID is a user affected by an add, update, or removal proposal in \vec{P} that had an effect.

Admin Cleanness Predicate. The second predicate C_{adm} models administration security. Intuitively, this predicate should be more permissive in some aspects than C_{cgka} , since a forgery attempt should be permitted even if the adversary knows the state of a (standard) group member. Following the approach above, we can decompose C_{adm} as $C_{\text{adm}} = C_{\text{adm-opt}} \wedge C_{\text{adm-add}}$.

Chapter 5. Cryptographic Administration for Secure Group Messaging

$C_{\text{adm-opt}}$ is symmetric to $C_{\text{cgka-opt}}$ and excludes the following family of attacks: the adversary attempts a forgery on a member ID' at an administrative epoch t_a^* while having exposed the state of an administrator $ID \in G^*$ at an administrative epoch $t_{\text{exp}} \leq t_a^*$, such that ID has not updated at some point between them. The predicate is optimal, as any attack that it excludes must occur while an administrator is directly under state exposure. In the game itself, we also require that ID' is in the challenge epoch specified by the adversary, i.e., $\text{ep}[ID'] = (\cdot, t_a^*)$. Notice that this predicate is unrelated to the common group secret and standard epochs t_s , and only relates to administration dynamics.

$$\boxed{\begin{array}{l} C_{\text{adm-opt}} : \forall (i, ID, ID', \text{ctr} \in (0, \text{exp-ctr}) : q_i = \text{INJECT}(ID', \cdot, t_i^*), \\ (ID \notin \text{ADM}[t_i^*]) \vee \\ (\exists (t_i, c) : (\text{tExp}(ID, \text{ctr}) \cdot t_a < t_i \leq t_i^*) \wedge \text{hasUpd}_{\text{adm}}(ID, \text{T}[(\cdot, t_i), \text{com}, c], \text{T}[(\cdot, t_i), \text{vec}, c]) \wedge C[(\cdot, t_i)] = c) \\ \vee (t_i^* < \text{tExp}(ID, \text{ctr}) \cdot t_a). \end{array}}$$

Figure 5.5: Optimal administrative predicate where the adversary makes oracle queries q_1, \dots, q_n .

The optimal administrative predicate $C_{\text{adm-opt}}$ is captured in Figure 5.5. In the expression, the function $\text{hasUpd}_{\text{adm}}$ is defined as in $\text{hasUpd}_{\text{std}}$, except it is defined with respect to $\text{com-type} \in \{\text{adm}, \text{both}\}$ (rather than $\text{com-type} \in \{\text{std}, \text{both}\}$).

A-CGKA Forgery Predicate. For A-CGKA, we define security under active attacks performed using the INJECT oracle with respect to a predicate C_{forgery} . The predicate we describe captures the fact that if admins have not been compromised, then non-admins can only make group changes for self-removes, i.e., when non-admins want to remove themselves. Moreover, we require that self-removes cannot be forged themselves by parties that are not corrupted. If there are no self-remove operations, then the predicate reduces to the fact that non-admins cannot cause changes in the group.

The predicate C_{forgery} is defined as follows with respect to variables in INJECT and the KIND game in general. Suppose m is input to INJECT. Let $\vec{P} = \text{Props}(\text{ST}[ID], m)$. Consider $\vec{P}' = \{P \in \vec{P} : P' \in \text{T}[\text{ep}[ID], \text{prop}, \cdot] \wedge \text{Prop-Info}(\text{ST}[ID], P) = \text{Prop-Info}(\text{ST}[ID], P')\}$.³ Let $H = \{ID : (\text{gid}, \text{rem}, ID, ID) = \text{Prop-Info}(\text{ST}[ID], P) \wedge (P \in \vec{P}')\}$ and $H^* = \{ID : (\text{gid}, \text{rem-adm}, ID, ID) = \text{Prop-Info}(\text{ST}[ID], P) \wedge (P \in \vec{P}')\}$. Then C_{forgery} is true if and only if $(\text{ST}[ID].G \setminus H, \text{ST}[ID].G^* \setminus H^*) \neq (\gamma.G, \gamma.G^*)$. If there are no self-removes, i.e., $H = H^* = \emptyset$, this simplifies to the predicate $(\text{ST}[ID].G, \text{ST}[ID].G^*) \neq (\gamma.G, \gamma.G^*)$; let C_{forgery}^* be this simplified predicate.

Definition 43 (KIND). A CGKA CGKA (resp. A-CGKA A-CGKA) is KIND w.r.t. predicate C_{cgka} (resp. and $C_{\text{adm}}, C_{\text{forgery}}$) if, for all efficient adversaries \mathcal{A} , we have

$$\text{Adv}_{(\text{A})\text{-CGKA}}^{\text{kind}}(\mathcal{A}) := \left| \Pr[\text{KIND}_{(\text{A})\text{-CGKA}, C_{\text{cgka}}, (C_{\text{adm}}, C_{\text{forgery}})}(\mathcal{A}) \Rightarrow 1] - \frac{1}{2} \right| = \text{negl},$$

³The effect of the equality check with respect to Prop-Info is that a dishonest proposal P' that has the same semantics as an honest proposal P will not be considered a ‘forgery’ by C_{forgery} .

where game KIND is defined in Figure 5.3.

Limitations. Our security definition, which is based on Alwen et al.’s unauthenticated CGKA notion [ACDT20], does not allow arbitrary message injections to participants. Thus, attacks on robustness are not captured by our security model. In particular, so long as non-admins are allowed to make commits, our A-CGKA schemes will only provide as much security as the underlying core CGKA: using MLS’s TreeKEM, for example, a malicious non-admin can deny service by sending a malformed commit message that can be successfully processed only by some of the users. This can be fixed at the expense of using NIZKs within TreeKEM [ACJM20, DDF21]. In any case, we note that confidentiality is not compromised under this family of attacks, as their main consequence is to “disconnect” users from the protocol (in particular, new users cannot be added).

If only admins are allowed to commit, then our schemes (to be introduced) are safe against this attack vector for some non-strongly robust variants of TreeKEM, such as the one used in MLS [BBR⁺23]. Standard users can still attain FS and PCS guarantees, and in particular PCS when their update proposals are committed.

Among the broader family of group key agreement protocols, where long-lived sessions and PCS are not always considered, modelling fully active adversaries is common [PRSS21]. We also do not explicitly model authentication (we implicitly assume an incorruptible PKI) and randomness manipulation, and we do not explicitly model parties who do not delete their state as instructed by the algorithm and are then exposed (via a *no-deletion* oracle or similar [ACDT21a]). We leave these for future work.

Multi-group security can be captured rather easily. The main difference (besides increased notation complexity introduced by the gids as in Figure 5.2) appears in the state exposure oracle: exposing the state of a party implies a security loss in all groups that the party is a member of simultaneously. The feature is not included as our security proofs are in the single-group setting.

Our INJECT oracle does not allow the adversary to inject group creation messages. Of course, the adversary can always make a new group with whatever users it chooses. It is nonetheless possible to extend our security notion to allow for injections, such that the adversary can only create a group for ‘valid’ users, i.e., those who have registered their keys with the PKI (which would be checked when a user processes a welcome message).

5.3 A-CGKA Constructions

In this section, we first present two different constructions for A-CGKA: *individual admin signatures* (IAS) and *dynamic group signature* (DGS). In the first construction, we construct A-CGKA on top of a CGKA by adding additional administration mechanisms based on authentication via signatures. In the second construction, we use two independent but synchronised

Chapter 5. Cryptographic Administration for Secure Group Messaging

CGKAs. We selected these two approaches due to their simplicity, efficiency and flexibility to adapt to different underlying CGKAs.

A first attempt of A-CGKA is to simply require group members to keep a list of administrators over time. Whenever an admin wants to make a commit, it can simply check whether the admin-changing proposals have been made by administrators, then commit them, and the other users will verify the admin condition upon processing. This approach is functional, but not secure in our model due to a lack of admin authentication. An adversary can easily forge a commit message and impersonate an admin unless this message is authenticated (for example signed). Many notions of CGKA security [ACDT20, KPPW⁺21] do not necessarily imply such a level of authentication.

One partial fix is to require admins to sign using a key derived from a long-term identity key. Then, security cannot be recovered if the admin is compromised once, resulting in the adversary winning the A-CGKA game too. Our constructions provide FS and PCS for the relevant admin authentication mechanisms in order to circumvent this problem.

In the following two subsections, all implemented A-CGKA algorithms, including `lnit`, are stateful as if executed by the same party and, as written, *do not explicitly return the updated local state*. Instead, they modify the state during runtime. In the event of algorithm failure, the state is not modified and appropriate failure values are output.

5.3.1 Individual Admin Signatures

In our first construction, *individual admin signatures* (IAS), we build a generic and modular administration mechanism on top of an arbitrary CGKA protocol (denoted by CGKA). Each group administrator $ID \in G^*$ maintains their own signature key pair (ssk, spk) that is updated over time. Each key pair is independent from the keys used in CGKA, which is mostly used as a black-box. Group members keep track of the list of admins G^* which is (possibly) updated upon processing each control message. Proposed changes to the group and to the administration are signed using an admin's keying material.

The IAS construction is presented in Figures 5.6, 5.7 and 5.8. The first figure describes the A-CGKA algorithms, and the second and third describe helper functions and auxiliary methods. We note that the algorithms defined in Figure 5.6 are incomplete without the helper functions; therefore, the construction spans all three figures.

States. We represent the state of a participant by the symbol γ , which is in part a dictionary of states indexed by group identifiers, i.e., $\gamma[gid]$. Users further maintain a common state via $\gamma.s0$ encoding the underlying CGKA state, security parameter 1^λ in $\gamma.1^\lambda$ and the user's ID in $\gamma.ME$. For each group gid , users keep a separate variable $\gamma[gid].adminList$ that encodes the group administrators and two administration-related signature key pairs. The state also tracks the group members as $\gamma[gid].G = \gamma[gid].s0.G$ and the CGKA key as $\gamma[gid].k = \gamma[gid].s0.k$ from the underlying CGKA, as well as the admins as $\gamma[gid].G^* = \{ID : \gamma[gid].adminList[ID] \neq \perp\}$.

Description

In our functions in Figures 5.6, 5.7 and 5.8, we often omit the group identifier of the state to simplify presentation. We assume that γ refers to $\gamma[\text{gid}]$ whenever gid is a subject of the algorithm, such as when it is a parameter of the function, and sometimes omit gid when it is clear from context. We note that our scheme nevertheless supports multiple groups.

Randomness. In our construction, we make randomness used by protocol algorithms explicit, including sampled randomness $r_0 \in \{0, 1\}^\lambda$ as input. Namely, for the input randomness r_0 used in any randomised method, we apply a PRF $(r_1, \dots, r_k) \leftarrow H_k(r_0, \gamma)$ that combines the entropy of r_0 and the state γ . We do this to reduce the impact of randomness leakage and manipulation attacks [BRV20]: without prior knowledge of γ (and assuming it has a sufficient entropy), an adversary that reads or manipulates r_0 will not be able to derive a corresponding r_i value. This is an additional feature that aims to maintain certain security properties in stronger adversarial models than considered in this chapter, and does not interfere with the rest of the protocol.

PKI. IAS assumes a basic, incorruptible PKI functionality where all parties are authenticated with the PKI. The PKI provides a fresh signature public key spk for which only the party ID can retrieve the corresponding secret key ssk . This functionality is used in for two main reasons:

1. When the group of administrators expands; namely, when a party ID' crafts a group gid or makes an admin add proposal⁴; and
2. When a non-admin user wishes to remove themselves from gid (a ‘self-remove’).

For these purposes, we define a getSpk algorithm, which on input (ID, ID', gid) for subject ID and caller ID' outputs spk relevant to the context the call is made in. We also assume a method of the form $\text{getSsk}(\text{spk}, ID, \text{gid})$ that returns the ssk associated to spk when called by ID given they uploaded it. During protocol execution, parties upload signature key pairs (ssk, spk) to the PKI via an abstract $\text{registerKeys}(ID)$ method in initialisation and via $\text{registerKey}(ID)$ during the two aforementioned scenarios.⁵ Formally, the adversary is only exposed to getSpk and registerKey ; we assume the other functions are called as needed in the security game though.

Initialisation. Before the creation of a group, a participant starts by calling the init method, which initialises the state γ . In turn, init calls CGKA.init from the underlying CGKA to initialise its state $\gamma.s0$. (ssk, spk) and $(\text{ssk}', \text{spk}')$ are two signature key pairs for group administration. The first pair is the valid admin signature key pair using during protocol execution, while the second pair stores updated keys after a commit or a key update operation is performed by the participant but before it is processed (i.e., acts as a temporary variable). After successfully processing a commit message, the second key pair replaces the first.

⁴Since admin proposals are all signed, and at proposal time it is not clear if a commit will contain an admin add proposal, admins always upload their updated signature key to the PKI at proposal time. For simplicity, signature keys are also always uploaded to the PKI when sampled at commit time also.

⁵This abstraction is made to reduce notational complexity.

Chapter 5. Cryptographic Administration for Secure Group Messaging

<pre> IAS.Init($1^\lambda, ID; r_0$) 1: $(r_1, r_2) \leftarrow H_2(r_0, \gamma); \gamma.s0 \leftarrow \text{CGKA.Init}(1^\lambda, ID; r_1)$ 2: $\gamma.ME \leftarrow ID; \gamma.1^\lambda \leftarrow 1^\lambda$ 3: $\gamma[\cdot].\text{adminList}[\cdot] \leftarrow \perp$ // maps ID to spk 4: $\gamma[\cdot].\text{ssk}, \gamma[\cdot].\text{spk} \leftarrow \perp$ // active admin key pair 5: $\gamma[\cdot].\text{ssk}', \gamma[\cdot].\text{spk}' \leftarrow \perp$ // temporary key pair 6: $\text{registerKeys}(ID)$ // gen/upload keys to PKI </pre>	<pre> IAS.Create($\text{gid}, G, G^*; r_0$) 1: require $(\gamma.ME \in G^*) \wedge (G^* \subseteq G)$ 2: $(r_1, r_2) \leftarrow H_2(r_0, \gamma)$ 3: $(\gamma.s0, W_0) \leftarrow \text{CGKA.Create}(\gamma.s0, \text{gid}, G; r_1)$ 4: if $W_0 = \perp$ return \perp 5: $\text{adminList}[\cdot] \leftarrow \perp$ // this is not $\gamma.\text{adminList}$ 6: for $ID \in G^*$: 7: $\text{adminList}[ID] \leftarrow \text{getSpk}(ID, \gamma.ME)$ 8: $\gamma.\text{spk}' \leftarrow \text{adminList}[ME]$ 9: $\gamma.\text{ssk}' \leftarrow \text{getSsk}(\gamma.\text{spk}', ME)$ 10: $T_W \leftarrow (\text{'wel'}, \gamma.ME, W_0, \text{adminList})$ 11: return $(\text{gid}, \perp, T_W, \text{Sign}(\gamma.\text{ssk}', (\text{gid}, \perp, T_W); r_2))$ </pre>
<pre> IAS.Prop($\text{gid}, ID, \text{type}; r_0$) 1: $P \leftarrow \perp; (r_1, r_2, r_3, r_4) \leftarrow H_4(r_0, \gamma)$ 2: if $\text{type} = \text{*adm}$ // Note if $\text{type} = \text{upd-adm}$, keys are updated 3: require $\gamma.ME \in \gamma.G^*$ 4: $P \leftarrow \text{makeAdminProp}(\text{gid}, \text{type}, ID; r_1, r_2)$ 5: else $(\gamma.s0, P) \leftarrow \text{CGKA.Prop}(\gamma.s0, \text{gid}, ID, \text{type}; r_1)$ 6: if $(\text{type} = \text{rem}) \wedge (ID = ME) \wedge (ID \notin \gamma.G^*)$ 7: $(\text{ssk}, \text{spk}) \leftarrow \text{KeyGen}(\gamma.1^\lambda; r_3); \text{registerKey}(ID)$ 8: $P \leftarrow (P, \text{Sign}(\text{ssk}', P; r_4))$ 9: return P </pre>	<pre> IAS.Proc(T) 1: $(\text{gid}, T_C, T_W, \sigma_T) \leftarrow T; \text{acc} \leftarrow \text{false}$ 2: if $(\gamma.ME \notin \gamma[\text{gid}].s0.G) \wedge (T_W \neq \perp)$ 3: require $T_W.\text{msg-type} = \text{'wel'}$ 4: $\text{acc} \leftarrow \text{p-Wel}(\text{gid}, T_W, \sigma_T)$ // welcome helper 5: else if $(\gamma.ME \in \gamma[\text{gid}].s0.G) \wedge (T_C \neq \perp)$ 6: $(\text{msg-type}, ID, C_0, \cdot, \Sigma, \cdot) \leftarrow T_C$ 7: require $\text{msg-type} = \text{'comm'}$ 8: for $\sigma : (P, ID', \sigma) \in \Sigma$: 9: if $\neg \text{Vrfy}(\text{getSpk}(ID', ME), \sigma, P)$ $\vee (ID' \in \gamma[\text{gid}].G^*)$ return false 10: if $\sigma_C = \perp$ // no sign - check only self-removes 11: $(\gamma', \text{acc}) \leftarrow \text{CGKA.Proc}(\gamma[\text{gid}].s0, C_0)$ 12: $SR \leftarrow \{ID' : (\cdot, ID') \in \Sigma\}$ 13: if $\neg \text{acc} \vee \gamma'[\text{gid}].s0.G \cup SR \neq \gamma[\text{gid}].s0.G$ 14: return false 15: $\gamma[\text{gid}].s0 \leftarrow \gamma';$ return true 16: if $\neg [(ID \in \gamma[\text{gid}].G^*) \wedge$ $(\text{Vrfy}(\gamma[\text{gid}].\text{adminList}[ID], (\text{gid}, T_C, T_W), \sigma_T))]$ 17: return false // verification failed 18: $\text{acc} \leftarrow \text{p-Comm}(\text{gid}, T_C)$ // admin commit helper 19: return acc </pre>
<pre> IAS.Commit($\text{gid}, \vec{P}, \text{com-type}; r_0$) 1: require $\gamma.ME \in \gamma.s0.G \wedge \text{com-type} \in \{\text{adm}, \text{std}, \text{both}\}$ 2: $(r_1, \dots, r_4) \leftarrow H_4(r_0, \gamma)$ 3: $(\vec{P}_0, \vec{P}_A, \Sigma, \text{admReq}) \leftarrow \text{propCleaner}(\text{gid}, \vec{P})$ 4: require $\text{verifyPropSigs}(\vec{P}_0, \Sigma, \vec{P}_A)$ 5: if $\text{admReq} \vee (\text{com-type} \in \{\text{adm}, \text{both}\})$ 6: require $\gamma.ME \in \gamma.G^*$ 7: if $\text{com-type} \in \{\text{adm}, \text{both}\}$: $C_A \leftarrow \vec{P}_A$ 8: if $\text{com-type} \in \{\text{std}, \text{both}\}$ 9: $(C_0, W_0, \text{adminList}, k) \leftarrow \text{c-Std}(\text{gid}, \vec{P}_0, \vec{P}_A; r_1)$ 10: require $C_0 \neq \perp$ // Generate new key pair and sign new spk 11: $(\gamma.\text{ssk}', \gamma.\text{spk}') \leftarrow \text{KeyGen}(\gamma.1^\lambda; r_2); \text{registerKey}(ID)$ 12: $T_C \leftarrow (\text{'comm'}, \gamma.ME, C_0, C_A, \perp, \gamma.\text{spk}')$ 13: if $W_0 \neq \perp$ // share updated admin list 14: $T_W \leftarrow (\text{'wel'}, \gamma.ME, W_0, \text{adminList})$ 15: else $T_W \leftarrow \perp$ 16: $\sigma_T \leftarrow \text{Sign}(\gamma.\text{ssk}, (\text{gid}, T_C, T_W); r_4)$ 17: else // only self-removes - no admin sig 18: $(C_0, \perp, \perp, k) \leftarrow \text{c-Std}(\text{gid}, \vec{P}_0, \perp; r_3)$ 19: $T_C \leftarrow (\text{'comm'}, \gamma.ME, C_0, \perp, \Sigma, \perp)$ 20: $T_W \leftarrow \perp; \sigma_T \leftarrow \perp$ 21: return $((\text{gid}, T_C, T_W, \sigma_T), k)$ </pre>	<pre> IAS.Prop-Info(P) 1: if P is of the form $(P, \sigma) : (P, \sigma) \leftarrow P$ // self-removes 2: if P is a CGKA proposal 3: $(P_{\text{gid}}, P_{\text{type}}, P_{\text{ID}}, P_{\text{ID}'}) \leftarrow \text{CGKA.Prop-Info}(\gamma.s0, P)$ 4: else if P is an admin proposal 5: $(P_{\text{gid}}, P_{\text{type}}, P_{\text{ID}}, P_{\text{ID}'}, \perp, \perp) \leftarrow P$ 6: return $(P_{\text{gid}}, P_{\text{type}}, P_{\text{ID}}, P_{\text{ID}'})$ </pre>

Figure 5.6: Individual admin signatures (IAS) A-CGKA built from a CGKA CGKA, a signature scheme Sig, and PRFs $H_n : \mathcal{R} \times \text{ST} \rightarrow \mathcal{R}^n$ for $n \leq 4$, randomness space \mathcal{R} and state space ST. We let $\gamma[\text{gid}].G = \gamma[\text{gid}].s0.G$, $\gamma[\text{gid}].G^* = \{ID : \gamma[\text{gid}].\text{adminList}[ID] \neq \perp\}$, and $\gamma[\text{gid}].k = \gamma[\text{gid}].s0.k$.

<pre> valid(P) // Predicate checks validity of admin proposal 1: ($P.gid, P.type, P.ID, P.ID'$) \leftarrow Prop-Info(P) 2: $S_1 := (P.gid = gid)$ // correct group 3: $S_2 := (P.ID \in \gamma[gid].G)$ // ID member 4: $S_3 := (P.ID' \in \gamma[gid].G^*)$ // ID' admin 5: $C_1 := (P.type = rem-admin)$ 6: $S_4 := (P.ID \in \gamma[gid].G^*)$ // ID admin 7: $C_2 := (P.type = add-admin)$ 8: return $S_1 \wedge S_2 \wedge S_3 \wedge (\neg C_1 \vee S_4) \wedge \neg(C_2 \wedge S_4)$ </pre> <hr/> <pre> makeAdminProp($gid, type, ID; r_1, r_2$) 1: $P_0 \leftarrow \perp$ 2: if $type = add-admin$ 3: $spk_{pki} \leftarrow getSpk(ID, \gamma.ME)$ 4: $P_0 \leftarrow (gid, type, ID, \gamma.ME, spk_{pki})$ 5: else if $type = rem-admin$ 6: $P_0 \leftarrow (gid, type, ID, \gamma.ME, \perp)$ 7: else if $type = upd-admin$ 8: if $(\gamma.ssk', \gamma.spk') \neq (\perp, \perp)$ 9: return \perp // only one update per epoch 10: $(\gamma.ssk', \gamma.spk') \leftarrow KeyGen(\gamma.1^A; r_1)$; registerKey($ID$) 11: $P_0 \leftarrow (gid, type, \gamma.ME, \gamma.ME, \gamma.spk')$ 12: else return \perp 13: return $(P_0, Sign(\gamma.ssk, P_0; r_2))$ </pre> <hr/> <pre> c-Std($gid, \vec{P}_0, \vec{P}_A; r_1$) 1: $(\gamma.s0, (C_0, W_0), k) \leftarrow CGKA.Commit(\gamma.s0, gid, \vec{P}_0; r_1)$ 2: if $W_0 \neq \perp$ // list for new users only 3: $adminList' \leftarrow updAL(\gamma.adminList, \vec{P}_A)$ 4: return $(C_0, W_0, adminList', k)$ 5: else return (C_0, \perp, \perp, k) </pre> <hr/> <pre> verifyPropSigs($\vec{P}_0, \Sigma, \vec{P}_A$) 1: for $(P, ID, \sigma) \in \Sigma$ // non-admin self-removes 2: $spk \leftarrow getSpk(ID, \gamma.ME)$ 3: if $\neg Vrfy(sp, P, \sigma) \vee P \notin \vec{P}_0 \vee ID \in \gamma.G^*$ 4: return false 5: for $(P, \sigma_P) \in \vec{P}_A$: // admin props 6: $(\perp, \perp, \perp, ID') \leftarrow Prop-Info(P)$ 7: $spk_P \leftarrow \gamma.adminList[ID']$ 8: if $\neg(Vrfy(sp, P, \sigma_P) \wedge valid(P))$ 9: return false 10: return true </pre>	<pre> propCleaner(gid, \vec{P}) 1: $admReq \leftarrow false$; $\vec{P}_0, \vec{P}_A, \Sigma \leftarrow []$ 2: for $P \in \vec{P}$: 3: $(gid, P.type, P.ID, P.ID') \leftarrow Prop-Info(P)$ 4: if $(P.type = *-adm) \wedge valid(P)$ 5: $\vec{P}_A \leftarrow [\vec{P}_A, P]$ 6: $admReq \leftarrow true$ 7: else // \vec{P}_0 is handled by CGKA 8: if $P.type = rem$ \wedge $(P.ID = P.ID') \wedge (P.ID \notin \gamma[gid].G^*)$ 9: $(P', \sigma) \leftarrow P$ 10: $\vec{P}_0 \leftarrow [\vec{P}_0, P']; \Sigma \leftarrow [\Sigma, (P, P.ID, \sigma)]$ 11: else if $P.type \in \{add, rem\}$ 12: $admReq \leftarrow true$ 13: else $\vec{P}_0 \leftarrow [\vec{P}_0, P]$ // Admin rem from $G \implies rem$ from G^* 14: if $(P.type = rem) \wedge (P.ID \in \gamma[gid].G^*)$ 15: $P' \leftarrow makeAdminProp(gid, rem, ID; \perp)$ 16: $\vec{P}_A \leftarrow [\vec{P}_A, P']$ 17: $(\vec{P}_0, \vec{P}_A) \leftarrow enforcePolicy(\vec{P}_0, \vec{P}_A)$ 18: return $(\vec{P}_0, \vec{P}_A, \Sigma, admReq)$ </pre> <hr/> <pre> p-Wel(gid, T_W, σ) 1: $(\perp, ID, W_0, adminList) \leftarrow T_W$ 2: $(\gamma[gid].s0, acc) \leftarrow CGKA.Proc(\gamma.s0, W_0)$ 3: $acc \leftarrow acc \wedge Vrfy(getSpk(\gamma.ME, ID), (gid, \perp, T_W), \sigma)$ 4: if acc 5: $\gamma[gid].adminList \leftarrow adminList$ 6: if $adminList[ME] \neq \perp$ 7: $\gamma.spk \leftarrow adminList[ME]$ 8: $\gamma.ssk \leftarrow getSsk(sp, ME)$ 9: return acc </pre> <hr/> <pre> updAL($adminList, \vec{P}_A$) 1: for $P \in \vec{P}_A$ 2: $(gid, type, ID, \perp, spk, \perp) \leftarrow P$ 3: if $type \in \{add-admin, upd-admin\}$ 4: if $(type = add-admin) \wedge (ID = \gamma.ME)$ 5: $\gamma.spk, \gamma.ssk \leftarrow (spk, getSsk(sp, ID))$ 6: $adminList[ID] \leftarrow spk$ 7: if $type = rem-admin$ 8: $adminList[ID] \leftarrow \perp$ 9: if $(ID = \gamma.ME)$ 10: $(\gamma.ssk, \gamma.spk) \leftarrow (\perp, \perp)$ 11: return $adminList$ </pre>
----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

Figure 5.7: Helper functions for IAS in Figure 5.6 (part I).

Chapter 5. Cryptographic Administration for Secure Group Messaging

IAS.Props(T)	p-Comm(gid, T_C)
<pre> // Supports non-welcome control messages 1: (T_C, T_W, σ_T) $\leftarrow T$ 2: $\vec{P}_0 \leftarrow \text{CGKA.Props}(T_C.C_0)$; $\vec{P}_A \leftarrow T_C.C_A$ 3: return $\vec{P}_0 \parallel \vec{P}_A$ </pre>	<pre> 1: ($\perp, ID, C_0, C_A, \Sigma, \text{spk}$) $\leftarrow T_C$ // Check signatures in proposals 2: if $C_A \neq \perp$ // $C_A = \vec{P}_A$ 3: if $\neg \text{verifyPropSigs}(\perp, \perp, C_A)$ return false // Apply commit 4: if $C_0 \neq \perp$ 5: (γ', acc) $\leftarrow \text{CGKA.Proc}(\gamma.s_0, C_0)$ 6: if $\text{acc} = \text{false}$ return false 7: if $\gamma.ME \notin \gamma'.G$ // user removed 8: $\gamma[\text{gid}] \leftarrow \perp$ // reinitialise state (only for gid) 9: else $\gamma[\text{gid}].s_0 \leftarrow \gamma'$ // Set temporary updated keys 10: if ($ID = \gamma.ME$) \vee ($\exists P \in C_A : P.ID = \gamma.ME$) 11: ($\gamma.\text{ssk}, \gamma.\text{spk}$) \leftarrow ($\gamma.\text{ssk}', \gamma.\text{spk}'$) 12: $\gamma.\text{ssk}', \gamma.\text{spk}' \leftarrow \perp$ 13: $\gamma.\text{adminList} \leftarrow \text{updAL}(\gamma.\text{adminList}, C_A)$ 14: $\gamma.\text{adminList}[ID] \leftarrow \text{spk}$ // committer's key 15: return true </pre>
<pre> enforcePolicy(\vec{P}_0, \vec{P}_A) // This method can be extended to other policies 1: $\text{numAdmins} \leftarrow G^*$ 2: for $P \in [\vec{P}_0, \vec{P}_A]$ 3: ($\text{gid}, P.\text{type}, P.ID, P.ID'$) $\leftarrow \text{Prop-Info}(P)$ 4: if $\text{type} = \text{rem}$ // if duplicates, removal prevails 5: Delete any other P' s.t. $P.ID = P'.ID$ // except for rem-adm proposals 6: else if $\text{type} = \text{rem-adm}$ 7: Delete any other admin P' s.t. $P.ID = P'.ID$ 8: $\text{numAdmins} -$ 9: else if $\text{type} = \text{add-adm}$, $\text{numAdmins}++$ 10: require $\text{numAdmins} \geq 1$ // ensures $\emptyset \neq G^* \subseteq G$ 11: return (\vec{P}_0, \vec{P}_A) </pre>	

Figure 5.8: Helper functions for IAS in Figure 5.6 (part II).

Group Creation. The Create algorithm creates the group gid from the list of members G , the admin list from G^* , and outputs a (signed) control message T for the new members in G . The adminList variable maps identifiers $ID \in G^*$ to signature public keys spk_{ID} . The public signature keys are obtained via getSpk and each admin's private key can be retrieved from the PKI via getSsk while they are processing T , the control message that adds them to the group. The group creator directly stores such key pair as $(\gamma.\text{ssk}', \gamma.\text{spk}')$.

Proposals. Any group member can use Prop to create a proposal of a non-admin type; the algorithm calls CGKA.Prop in this case. Administrative proposals are restricted to admins and crafted by makeAdminProp, which includes an administrative signature in the proposal. The signature is included to prevent an (insider) adversary from forging the sender of the proposal in an attempt to impersonate an admin.⁶ Proposal creation does not have any effect on the state other than the storage of temporary keys for proposals with type $\text{type} = \text{upd-adm}$. In the case of an add-adm proposal to promote ID to admin status, the proposer $\gamma.ME$ retrieves a public signature key spk of ID from the PKI using getSpk. In the case of a rem proposal where $ID = ME$ and $ID \notin \gamma[\text{gid}].G^*$ (i.e., a non-admin self-remove), the caller samples a new signature key pair, registers the public key spk with the PKI and signs their proposal.

The Prop-Info method simply retrieves the main information of a proposal. As mentioned in

⁶Note in MLS that all proposals are in any case signed.

the previous section, it could be adapted to support CGKAs and A-CGKAs where proposals are encrypted⁷ (under some key derived from the group key, for instance).

Commits. The Commit algorithm can only be called by group administrators (except for the special case in which only key updates and self-removes are proposed, in which cases standard users can commit), and performs the following actions:

1. Clean the input vector of proposals \vec{P} , ensuring that they are well-formed. This is done via the `propCleaner` algorithm, which in turn calls the `enforcePolicy` method. For security reasons, we adopt the main features of the MLS policy (removing duplicates and prioritising removals) in our construction [BBR⁺23], but extensions to this policy can be implemented. In addition, we verify the legitimacy of the admin proposals and the fact that self-remove proposals are correctly signed via `verifyPropSigs`. Then, the predicate `valid(P)` verifies that the gid matches, that added users (respectively admins) do not belong to G (resp. G^*), that removed users do belong to G (resp. G^*), and that the proposer is an admin. Finally, we ensure that all users removed from G are also removed from the `adminList`.
2. Carry out the administrative and the standard commits and produce an administrative commit message C_A (which is the clean admin proposal vector), a standard CGKA commit C_0 , and an updated `adminList`. We split the CGKA commit in two components C_0 and W_0 as is usual in the literature [ACDT20, AJM22, KPPW⁺21, ACJM20]. If the CGKA does not allow for this, it is easy to modify the protocol without compromising security⁸.
3. Generate a new administrative signature key pair $(\gamma.ssk', \gamma.spk')$ stored as the temporary key pair.
4. Produce the final control message T which includes the new `spk'`. The message T is again split into two components: A first component T_W (for welcome) includes all the required information for incoming A-CGKA members, including the new list of admins. A second component T_C (for commit) contains the updating information for group members. Both components are signed together using the committer's current $\gamma.ssk$.

The `Props` method, given a commit, retrieves the list of proposals that it implements; this simply calls the underlying `CGKA.Props` algorithm and combines the output with the list of admin updates directly contained in a well-formed commit message.

Processing Control Messages. The `Proc` method takes a control message T as input and updates the state accordingly. The algorithm returns an acceptance bit `acc` which is true if

⁷Some precaution must be taken with respect to security when proposals are encrypted under the CGKA key, as the adversary gains access to multiple additional ciphertexts which can result in a security loss.

⁸This division is made for clarity, but it may be used to improve efficiency too. Namely, the welcome part W_0 of a commit message does not need to be processed by existing group members, so in principle C_0 can be sent only to these. In A-CGKAs such as IAS, this can also be applied if signatures are handled carefully.

Chapter 5. Cryptographic Administration for Secure Group Messaging

the processing succeeds, in which case the state is updated. Otherwise, the state remains the same. During an execution of Proc, some checks must pass before the state is updated. For newly added users, p-Wel verifies the message signature on the adminList, attempts to process the message via the underlying CGKA, and updates the state given this succeeds. For group members, p-Com verifies the administrator signature and the signatures in the admin proposals. The state is updated if all verification succeeds; a removed user deletes their state, and temporary keys are updated if necessary. The case in which the control message T is not signed is handled by Proc directly by verifying that no changes to the group structure are made except possibly for valid signed self-removals.

Features

We first note that the IAS protocol can be built over any (unauthenticated or authenticated) CGKA. Since signatures are often already present in CGKAs such as [AJM22], the extension from CGKA to A-CGKA can be more direct (and thus incur less overhead) than presented here. This also holds true for any group messaging scheme, such as the administrated MLS extension we describe in Section 5.3.4.

Commit and Propose Policies. Our construction allows standard users to perform a commit if there are no changes in the group structure or administration. This is an optional design choice that does not affect security in our model (and could be reflected in a correctness predicate), although, as previously discussed, adversarial group members may deny service if the underlying CGKA is not robust. We also enforce that standard users cannot propose administrative changes (even if these could be later ignored by admins), and similarly can be allowed as required by an application.

Security Mechanisms. The security of the group administration is provided by the admin signatures; an adversary should not be able to commit changes to the group unless it compromises the state of one of the group administrators. The update mechanism provides optimal post-compromise security in our model.

On the other hand, administrative actions are undeniable and traceable both by group members and by the message delivery service. Separately, additional protections (i.e. checking members are registered on the PKI) are needed to ensure that parties are not invited to fake groups where the list of group administrators is forged.

On Optimal Forward Security. Note that, as defined, our construction does not satisfy forward security with respect to injection queries even if the underlying CGKA provides optimal forward security. Concretely, suppose that ID makes their last update in epoch 3, and then their state is exposed in epoch 5. Then ID can trivially forge commit messages for parties that are in epochs 3 and 4 since their keying material has not been updated. A similar forward security issue is present in the MLS standard affecting confidentiality [ACDT20]. Optimal security can be straightforwardly achieved by replacing regular signatures with *forward-secure*

signatures [BM99]. Forward-secure signatures allow signers to non-interactively update their secret keys and provide forward security given state exposure. In IAS, it suffices to use forward-secure signatures such that whenever an epoch passes and an admin has not sampled a new signature key, they invoke the signature scheme's secret key update function, where new signature keys are otherwise derived as in the construction. We note that forward-secure signatures involve an overhead that may be undesirable in some cases, and also they are not used in current protocols (signatures are already used in MLS's CGKA, for instance). In Theorem 17, we characterise the exact security of IAS using standard primitives via our sub-optimal predicate. In this way, the security of both alternatives is fully characterised.

Correctness and Security

We proceed to proving correctness and security.

Theorem 16. Let CGKA be a correct CGKA (Definition 42) and Sig be a 1-correct signature scheme. Then, the IAS protocol (Figures 5.6, 5.7 and 5.8) is correct with respect to Definition 42.

Proof. We want to prove that no adversary \mathcal{A} can win $\text{CORR}_{\text{A-CGKA}, \text{C}_{\text{corr}}}$ (where we set $\text{C}_{\text{corr}} = \text{true}$) played with respect to IAS (Figures 5.6, 5.7 and 5.8) given that CGKA is correct. We analyze the different game oracles separately and sketch parts derived by direct inspection or based on CGKA correctness.

For PROP, \mathcal{A} can win the game in PROP if either Prop-Info incorrectly interprets the proposal or if the Prop call changes the view of the group. In the first case, correctness follows from the correctness of CGKA.Prop-Info if the proposal is standard, and by inspection of IAS otherwise. In the second case, the group view is never changed by Prop (as makeAdminProp only updates $\gamma.\text{ssk}'$, $\gamma.\text{spk}'$ in case of an admin proposal, and CGKA.Prop is correct in the case of a standard proposal).

CREATE and COMMIT can be proven correct by inspection in a similar way.

For DELIVER, we examine each **reward** clause. Note that in line 2 of DELIVER, T is either a commit message created by Create or by Commit.

We start with the clause $(\emptyset \neq \gamma[\text{gid}].G^* \subseteq \gamma[\text{gid}].G)$. If T is a create message, then line 1 of Create and the fact that variable adminList is populated ensures that this condition is fulfilled for T . Upon processing, and after a correct PKI retrieval, p-Wel overwrites $\gamma.\text{adminList}$ and $\gamma.s0$ (via CGKA.Proc), so the condition holds. If T is a (standard) commit made by a non-admin user – that is, one without a signature – then there are no changes to the group as checked explicitly by the Proc algorithm from line 10. Otherwise, if T an admin commit, then it is processed by p-Comm. We can distinguish some further cases depending on the proposals contained in T :

Chapter 5. Cryptographic Administration for Secure Group Messaging

- If T contains proposals of types `upd`, `upd-adm`, `add` only, then the condition is trivially met.
- If T additionally contains `rem` proposals, for every removed ID a corresponding `rem-adm` proposal is generated by an honest admin in `propCleaner`, hence $G^* \subseteq G$.
- If T additionally contains `add-adm` proposals, the `valid(P)` predicate checks that the added admins are already group members (via S_2), hence $G^* \subseteq G$.
- If T additionally contains `rem-adm` proposals, the final check in `enforcePolicy` ensures that $G \neq \emptyset$.
- Any other combination of several contradicting proposals affecting the same ID is handled by `enforcePolicy`, which prioritises removals (while preserving admin removals for the same user) which performs a final check on the size of G^* .

We conclude that, for any possible combination of proposals, the condition is always met provided that `acc = true` with respect to T .

The next cases are the **reward** `Props(ST[ID], T) \neq T[gid, (t, c), vec, c']` and the **reward** `γ [gid].k $\neq \perp$` condition given `Proc` outputs γ such that $ID \notin \gamma[\text{gid}].G$. It is straightforward to see that both conditions are met by correctness of `CGKA.Props`.

The check by `UpdateView` rewards the adversary if two users processing the same commit message (on epoch (t, c)) differ in their group view. We show correctness by induction. Suppose ID_1 and ID_2 process the same commit message T . If they are in epoch $(-1, -1)$ and process a create message, correctness is easily seen. For the inductive step, we assume that their group views were equal in (t, c) , and we want to show that they remain equal after moving to epoch $(t+1, c')$. The commit is handled by `p-Comm`, and the only parts that can change for different users are the **if** condition in line 10 and `updAL`. The behaviour of both sections of code varies only on the modification of γ 's signature keys, but not on the group structure. Hence, by the correctness of `CGKA`, we conclude that ID_1 and ID_2 end up having consistent views.

The edge case in which a user is just added to the group is handled by `p-Wel`, and follows from `CGKA` correctness and the fact that `γ .adminList \leftarrow adminList` is executed where `adminList` is directly provided in T .

Finally, the check **reward** `γ [gid].k \neq T[gid, (t, c), key, c']` follows from the correctness of the `CGKA.Commit` algorithm which outputs the new group key k . \square

Recall that IAS, which uses a (regular) digital signature scheme, does not provide optimal forward security. Therefore, we prove security with respect to a sub-optimal admin cleanness predicate C_{adm} where $C_{\text{adm}} = C_{\text{adm-opt}} \wedge C_{\text{adm-add}}$ and $C_{\text{adm-opt}}$ is defined in Section 5.2.4.

Theorem 17. Consider KIND `CGKA` `CGKA` with respect to cleanness predicate C_{cgka} , `SUF-CMA` signature scheme `Sig` and `PRF` H_4 . Then, for A-`CGKA` IAS, we have, with respect

to predicates $C_{\text{cgka}}, C_{\text{adm}}, C_{\text{forgery}}$, that for every efficient adversary \mathcal{A} that makes at most q oracle queries, one can build an adversary \mathcal{B} such that

$$\text{Adv}_{\text{IAS}}^{\text{kind}}(\mathcal{A}) \leq q \cdot \text{Adv}_{H_4}^{\text{prf}}(\mathcal{B}) + \text{Adv}_{\text{CGKA}}^{\text{kind}}(\mathcal{B}) + q^2 \cdot \text{Adv}_{\text{Sig}}^{\text{suf-cma}}(\mathcal{B}).$$

Proof idea. We first bound an adversary \mathcal{A} 's advantage in distinguishing between the $\text{KIND}_{\text{A-CGKA}}$ game and a game G_1 which replaces calls to hash functions H_i by uniformly sampling the output (modelling each H_i in IAS as a PRF). Then, we divide \mathcal{A} 's behaviour in G_1 into two events based on whether they successfully query the INJECT oracle (event E_1) or not (event E_2). Given E_1 , we reduce security via a number of SUF-CMA adversaries. Otherwise, we reduce directly with $\text{KIND}_{\text{CGKA}}$ adversary, at which point the claim follows.

We start with the definition of the predicate C_{adm} predicate (Figure 5.9) that we prove IAS secure with respect to. After proving security, we discuss how one can (easily) extend the proofs and ensure optimal security using forward-secure signatures.

$ \begin{aligned} C_{\text{adm}} : & \forall (i, ID, ID', \text{ctr} \in (0, \text{exp-ctr}) : q_i = \text{INJECT}(ID', \cdot, t_i^*), \\ & (ID \notin \text{ADM}[t_i^*]) \vee \\ & (\exists (t_i, c) : (\text{tExp}(ID, \text{ctr}).t_a < t_i \leq t_i^*) \wedge \\ & \text{hasUpd}_{\text{adm}}(ID, T[(\cdot, t_i), \text{com}, c], T[(\cdot, t_i), \text{vec}, c]) \wedge \\ & (C[(\cdot, t_i)] = c)). \end{aligned} $

Figure 5.9: Sub-optimal administrative predicate where the adversary makes oracle queries q_1, \dots, q_n .

Note that it differs from the optimal predicate (Figure 5.5) only by the lack of $(t_i^* < \text{tExp}(ID, \text{ctr}).t_a)$ condition, so it holds that $C_{\text{adm}} \wedge C_{\text{adm-opt}} = C_{\text{adm}}$. In particular, the forward security guarantees are weaker since, e.g., if a party updates their admin key in admin epoch 3 then if they are exposed in epoch 5 then the adversary can make a trivial forgery in the construction (and thus it is considered a trivial attack by C_{adm}).

Towards proving security, we prove the following lemma.

Lemma 16. Let \mathcal{A} be a $\text{KIND}_{\text{A-CGKA}, C_{\text{cgka}}, C_{\text{adm}}, C_{\text{forgery}}}^{\mathcal{A}}$ adversary playing with respect to IAS. Consider any query \mathcal{A} makes of the form $\text{INJECT}(ID, m, t_a)$ which outputs $v \neq \perp$. Then there is an efficient algorithm that, given the inputs/outputs of oracle queries \mathcal{A} makes, parses $m = (\text{gid}, T_C, T_W, \sigma_T)$ and derives pk such that $\text{Vrfy}(\text{pk}, \sigma_T, (\text{gid}, T_C, T_W)) = 1$.

Proof. Consider a given query $\text{INJECT}(ID, m, t_a)$ made by \mathcal{A} that outputs $v \neq \perp$. Given $v \neq \perp$ and by definition of INJECT, a call $(\gamma, \text{true}) \leftarrow \text{Proc}(\text{ST}[ID], m)$ was previously made by the challenger such that $C_{\text{forgery}} = \text{true}$. Note \mathcal{A} cannot register (malicious) signature keys with the PKI, and keys are assumed to be bound to the context that they are used in, e.g., for self-removals. In addition, any non-admin commit comprising of group changes that only consists of self-removes will not result in the adversary winning by construction of IAS,

Chapter 5. Cryptographic Administration for Secure Group Messaging

and if a self-remove is created then even if another ‘dishonest’ self-remove can be created for that party, C_{forgery} will consider it equivalent to the ‘original’ self-remove. Thus, self-removes cannot cause C_{forgery} to be true. Now, INJECT disallows $t_a \neq -1$, which is the case if and only if $ID \notin G$, and that unsigned control messages cannot change the group structure. Thus, we only need to consider control messages that are input to p-Comm in Proc and result in output $\text{acc} = \text{true}$ when $\sigma_T \neq \perp$. To reach p-Comm, the Proc call must be such that $\text{Vrfy}(\gamma[\text{gid}].\text{adminList}[ID], (\text{gid}, T_C, T_W), \sigma_T) = 1$. Since $\gamma[\text{gid}].\text{adminList}[ID]$ must have been previously sent in some commit or welcome message by construction of IAS, it follows that $\text{pk} = \gamma[\text{gid}].\text{adminList}[ID]$ is efficiently computable. \square

We prove IAS secure below.

Proof. Let G_0 be the $\text{KIND}_{\text{IAS}, C_{\text{cgka}}, C_{\text{adm}}, C_{\text{forgery}}}^{\mathcal{A}}$ game. Let G_1 be as in G_0 , except that all calls of the form $(r_1, \dots, r_i) \leftarrow H_i(r_0, \gamma)$ are replaced with calls of the form $(r_1, \dots, r_i) \stackrel{\$}{\leftarrow} R^i$ where R is the space of random coins used by each (A)-CGKA algorithm. Note that in IAS we always have $i \leq 4$; we assume for simplicity that each H_i for $i \leq 4$ is implemented by calling H_4 and then truncating the output as needed.

Let $G_{0,0} = G_0$. Let $G_{0,j}$ be G_0 except that the first j calls of the form $H_i(r_0, \gamma)$ that adversary \mathcal{A} makes are replaced as above, and the rest remain unchanged. Note that since every oracle query that \mathcal{A} makes results in at most one call to a function of the form $H_i(\cdot, \cdot)$, $G_{0,k} = G_1$ for some $k \leq q$.

Consider $G_{0,j-1}$ and $G_{0,j}$ where $j \geq 1$; suppose these games are played by adversary \mathcal{A} . Let \mathcal{A}' be a PRF adversary (keyed with the first argument). \mathcal{A}' simulates directly except when \mathcal{A} makes their oracle query which leads to the j -th call to a function of the form H_i . Upon this call, \mathcal{A} simulates this call by calling $\text{EVAL}(\gamma)$ for the input γ , and truncates the response (r_1, \dots, r_4) to (r_1, \dots, r_i) when necessary before continuing execution (i.e., its simulation). Clearly \mathcal{A}' perfectly simulates $G_{0,j-1}$ when \mathcal{A}' 's challenger's bit is 1. When \mathcal{A}' 's challenger's bit is 0, note that the output of \mathcal{A}' 's EVAL call, namely (r_1, \dots, r_4) , is of the form $F(\gamma)$ for a uniformly random function $F: \text{ST} \rightarrow R^4$, where ST is the A-CGKA state space. Since F is randomly chosen, this output is distributed identically to (r_1, \dots, r_4) where each r_i is uniformly sampled from R . It follows that \mathcal{A}' perfectly simulates $G_{0,j}$ when its challenge bit is 0. By combining the sequence of game hops, it follows that

$$\left| \Pr[G_0(\mathcal{A}) \Rightarrow 1] - \frac{1}{2} \right| \leq \left| \Pr[G_1(\mathcal{A}) \Rightarrow 1] - \frac{1}{2} \right| + q \cdot \text{Adv}_{H_4}^{\text{prf}}(\mathcal{A}').$$

In the following, we will simulate for a G_1 adversary via either a SUF-CMA or CGKA adversary. Without hopping from G_0 to G_1 , the simulation would not have been identical when e.g. using the SUF-CMA SIGN oracle. Since we have transitioned to G_1 , which uses randomness normally, there are no issues regarding simulation and randomness.

Let \mathcal{A} be a G_1 adversary. Let E_1 be the event that \mathcal{A} makes a query to INJECT such that INJECT outputs value $v \neq \perp$ (i.e., the challenge bit). Let E_2 be the event that this does not occur; clearly $\Pr[E_1] + \Pr[E_2] = 1$. We consider each event separately. Without loss of generality, we restrict our simulations given E_1 to the case where C_{adm} is true as otherwise the adversary cannot win, and similarly given E_2 to the case where C_{cgka} is true. It suffices to observe that these two predicates are efficiently computable such that the adversaries we construct below can abort during unclean executions.

E_1 : By construction of IAS and Lemma 16, note, given that INJECT outputs $v \neq \perp$, that a signature forgery has occurred where the signature keying material is sampled due to an oracle call. Given E_1 , we need to determine which input values ID and t_a are used by \mathcal{A} on the first INJECT call which outputs $v \neq \perp$. By construction of IAS, this can happen as a result of a query to CREATE, PROP or COMMIT. However, \mathcal{A} may make at most q injection attempts with respect to this key pair, each of which may be a winning one. Thus, the reduction has to guess both 1) the INJECT query which first outputs $v \neq \perp$ and 2) the query q_i which generates the signature key pair corresponding to this injection.

Let $E_{1,i,j}$ be the event that \mathcal{A} makes query q_i which leads to the generation of signature key pair (ssk, spk) such that query q_j is the first query to INJECT resulting in output $v \neq \perp$ ⁹. Note that IAS is such that each oracle query q_i results in at most one new signature key pair being sampled by the challenger. By the union bound, we have:

$$\Pr[G_1(\mathcal{A}) \Rightarrow 1 \wedge E_1] \leq \sum_{i,j \in \{1, \dots, q\}} \Pr[G_1(\mathcal{A}) \Rightarrow 1 \wedge E_{1,i,j}].$$

Suppose $E_{1,i,j}$ holds. Let \mathcal{A}' be an SUF-CMA adversary that simulates for G_1 adversary \mathcal{A} . We aim to show that $\Pr[G_1(\mathcal{A}) \Rightarrow 1 \wedge E_{1,i,j}] \leq \text{Adv}_{\text{Sig}}^{\text{suf-cma}}(\mathcal{A}')$.

\mathcal{A}' simulates as follows. \mathcal{A}' simulates variable initialisation as in G_1 except that it lazily simulates Init calls (ensuring it remains polynomially-bounded). \mathcal{A}' simulates the first $i - 1$ of \mathcal{A} 's oracle queries locally, i.e., simulates all relevant behaviour resulting in corresponding state and game variables being set and updated. This includes queries of the form getSsk and getSpk which \mathcal{A}' simulates via signature scheme Sig .

Consider \mathcal{A}' 's i -th oracle query q_i . Let $(\text{ssk}^*, \text{spk}^*)$ denote the signature key pair sampled by the SUF-CMA challenger; recall that SUF-CMA adversary \mathcal{A}' has access to oracle SIGN . \mathcal{A}' simulates as follows:

- If q_i is to CREATE, then $(\text{ssk}^*, \text{spk}^*)$ plays the role of $(\gamma.\text{ssk}', \gamma.\text{spk}')$. Namely, \mathcal{A}' sets the output of $\text{getSpk}(ID, \gamma.ME)$ to spk^* and that of $\text{getSsk}(\gamma.\text{spk}', ME)$ to ssk^* after embedding spk^* in $\text{adminList}[ME]$ at line 7 of Create. \mathcal{A}' also calls $\text{SIGN}((\text{gid}, \perp, T_W))$ which

⁹Note that we only consider the first such query since the simulation will end after this point.

Chapter 5. Cryptographic Administration for Secure Group Messaging

outputs σ . \mathcal{A}' otherwise simulates and returns the result to \mathcal{A} (which includes σ).

- If q_i is to PROP with input type = upd-adm, then $(\text{ssk}^*, \text{spk}^*)$ plays the role of $(\gamma.\text{ssk}', \gamma.\text{spk}')$ (line 10 of makeAdminProp). Similarly to above, \mathcal{A}' embeds spk^* in P_0 (line 11 of makeAdminProp), and then simulates the rest of the oracle call.
- Otherwise, q_i is to COMMIT. Note we assume $E_{1,i,j}$ holds. Thus, the branch at line 5 in Commit must be executed. As before, $(\text{ssk}^*, \text{spk}^*)$ plays the role of $(\gamma.\text{ssk}', \gamma.\text{spk}')$; \mathcal{A}' simulates the remainder of the call.

All other oracle queries, except to INJECT, are simulated locally by \mathcal{A}' except when signatures with respect to spk^* are required, in which case SUF-CMA oracle SIGN is used, or when spk^* is to be embedded in a message. Note that at most q SIGN queries are made by \mathcal{A}' since each oracle query \mathcal{A} makes produces at most one signature (which may or may not require SIGN to simulate).

Note by construction of IAS that each signature key pair is sampled by a single party and is never revealed/embedded in another message, and that each key pair is uniformly and independently sampled. Thus, the simulation is valid, ignoring EXPOSE queries, since the challenge key pair is independent of all other keying material and challenge secret key ssk^* is never revealed. Regarding EXPOSE queries, since C_{adm} is true (since we assume E_1), \mathcal{A} will never query EXPOSE with respect to the challenge key pair, a query which would otherwise not be able to be simulated.

For \mathcal{A} 's query $q_{j'} = \text{INJECT}(ID)$ such that $j' \neq j$, \mathcal{A}' simulates by returning \perp to \mathcal{A} . When \mathcal{A} makes query $q_j = \text{INJECT}(ID, m, t_a)$, \mathcal{A}' inspects $m = (\text{gid}, T_C, T_W, \sigma_T)$ and returns the message/signature pair $((\text{gid}, T_C, T_W), \sigma_T)$ to \mathcal{A} which, by Lemma 16, exists. These two steps are valid by definition of $E_{1,i,j}$ and that $j' < j$ in the first step since we only simulate up to query j .

It thus follows that the simulation is perfect. Noting that \mathcal{A} wins at most as often as \mathcal{A}' (since e.g., \mathcal{A} may come up with a forgery (m, σ) nevertheless rejected by Proc), we have:

$$\Pr[G_1(\mathcal{A}) \Rightarrow 1 \wedge E_{1,i,j}] \leq \text{Adv}_{\text{Sig}}^{\text{suf-cma}}(\mathcal{A}').$$

E_2 : Note first that given E_2 , we can deduce that every query to INJECT will have output \perp . Let \mathcal{A}' be a $\text{KIND}_{\text{CGKA}, \text{C}_{\text{cgka}}}^{\mathcal{A}'}$ adversary. \mathcal{A}' simulates for $\text{KIND}_{\text{A-CGKA}, \text{C}_{\text{cgka}}, \text{C}_{\text{adm}}, \text{C}_{\text{forgery}}}^{\mathcal{A}}$ adversary \mathcal{A} as follows. When \mathcal{A} makes an oracle query, \mathcal{A}' executes all code in IAS except for that which makes use of CGKA algorithms, which are processed via \mathcal{A}' 's oracles; \mathcal{A}' then returns each response to \mathcal{A} . One case we must deal with is when Proc is called at line 11; the call may succeed but the caller may ignore the state update, undoing the changes made; this happens given the group state changes as a result of the call. Note that \mathcal{A}' can determine whether this case occurs or not using the fact that commit messages are honestly delivered by construction of the KIND game and by correctness which ensures that honestly-generated and delivered

commit messages are accepted. In particular, \mathcal{A}' can use the policy to determine whether or not the call would update the group state, and only call its DELIVER when the group state is not changed. Finally, \mathcal{A}' outputs the same bit as \mathcal{A} .

To see that the simulation is perfect, first note that CGKA algorithms are used as black boxes in IAS. Moreover, except in the case dealt with above, CGKA state variables s_0 for each ID are not used except as input to CGKA algorithms, after which they are immediately overwritten, exactly as done by the CGKA KIND challenger given correctness and in particular the fact that failing algorithm calls do not update the state. Thus, it suffices to simulate CGKA code using \mathcal{A}' 's oracles. Thus:

$$\Pr[G_1(\mathcal{A}) \Rightarrow 1 \wedge E_2] = \Pr[\text{KIND}_{\text{CGKA}, C_{\text{cgka}}}(\mathcal{A}') \Rightarrow 1]$$

We then have:

$$\begin{aligned} \left| \Pr[G_1(\mathcal{A}) \Rightarrow 1] - \frac{1}{2} \right| &= \left| \Pr[G_1(\mathcal{A}) \Rightarrow 1 \wedge (\vee_{i,j} E_{1,i,j} \vee E_2)] - \frac{1}{2} \right| \\ &\leq \left| \sum_{i,j} \Pr[G_1(\mathcal{A}) \Rightarrow 1 \wedge E_{1,i,j}] + \Pr[G_1(\mathcal{A}) \Rightarrow 1 \wedge E_2] - \frac{1}{2} \right| \\ &\leq \left| \sum_{i,j} \Pr[G_1(\mathcal{A}) \Rightarrow 1 \wedge E_{1,i,j}] \right| + \left| \Pr[G_1(\mathcal{A}) \Rightarrow 1 \wedge E_2] - \frac{1}{2} \right| \\ &\leq q^2 \cdot \text{Adv}_{\text{Sig}}^{\text{suf-cma}}(\mathcal{A}') + \text{Adv}_{\text{CGKA}}^{\text{kind}}(\mathcal{A}'), \end{aligned}$$

where the second and third lines follow from the union bound and triangle inequality, respectively. The result follows by combining this with the game hop earlier. \square

Optimal Forward Security. As described above, we can achieve optimal forward security, and thus optimal admin security (i.e., security with respect to $C_{\text{adm-opt}}$) by replacing signatures with forward-secure signatures [BM99, MMM02]. The logic of the security reduction is very similar to that presented with regular signatures above. The main difference is that forward-secure signature calls are replaced by oracle calls, including possibly key exposure calls.

5.3.2 Dynamic Group Signature

In our second construction, *dynamic group signature* (DGS), the group administrators agree on a *common* signature key pair that they use for signing administrative messages on an underlying CGKA. To agree on a secret and generate a common key pair, they run a separate CGKA. As opposed to IAS, group administrators may now be opaque to group members if the concrete CGKA which is used allows it. The reason is that they authenticate admin messages using an admin signature key that is shared among all admins. Notably, group members do not need to keep track of an administrator list; admins implicitly track this via their CGKA.

5.3.3 Description

Protocol. DGS is introduced in Figures 5.10 and 5.11. In the algorithm, we refer to the primary (or standard) CGKA as CGKA, and to the administrative CGKA as CGKA*. The first CGKA allows group members to agree on a common secret and group over time as in IAS, whereas the second exists only for administrative purposes (i.e., admins deriving a common signature key). Note that CGKA* need not be implemented in the same way as the primary CGKA. This can be exploited by a protocol designer either for performance reasons or if, for instance, stronger FS and PCS guarantees are desired for the administrative CGKA. For simplicity of exposition, DGS as written does not support self-signed removal operations that non-admins can commit directly, but we note that the technique to implement them is identical to IAS.

States. Each party stores $\gamma.s0$, corresponding to the primary CGKA, as well as $\gamma.sA$, corresponding to CGKA*, which are used for each group they consider. For a given gid , let gid^* be another group identifier such that for all $gid_1 \neq gid_2$, $gid_1 \neq gid_2 \neq gid_1^* \neq gid_2^*$. We assume for DGS that gid is used by the main CGKA and gid^* by the admin CGKA. Besides these fields, the state includes the administrative public key $\gamma[gid].spk$ known by all group members (and can be a public group parameter, known for instance by a central server) to enable verification. The state variables are now $\gamma[gid].G = \gamma.s0[gid].G$, $\gamma[gid].G^* = \gamma.sA[gid^*].G$, and $\gamma[gid].k = \gamma.s0[gid].k$.

PKI. As in IAS, we assume an incorruptible PKI functionality. Similar to IAS, DGS relies on abstract helpers `registerKey(gid)` and `getSpk(gid)`, both of which now only take an identifier gid as input. We assume that admins register their admin signature public keys whenever they are updated or created, which either happens at group creation time or during a Commit call; nonetheless we only require that each user calls `getSpk(gid)` when they join group gid . Authentication could be implemented while ensuring k -anonymity such that a member authenticates his group membership but not his identity; such a feature cannot be provided by IAS without modification.

Initialisation. The `Init` procedure calls the `CGKA.Init` and `CGKA*.Init` algorithms to initialise $\gamma.s0$ and $\gamma.sA$, respectively, and sets $\gamma[\cdot].spk, \gamma[\cdot].ssk \leftarrow \perp$.

Group Creation. The `Create` algorithm creates a group for the two separate CGKAs by calling the two corresponding `Create` methods. These calls output new states $s0$ and sA , which overwrite the stored states, as well as control messages W_0 and W_A , which are collated into a *create* control message $T = T_{CR}$. We assume that an initial group signature public key is sampled and uploaded to the PKI by the caller of `Create`.

Proposals. The `Prop` algorithm generates a proposal message P by using `CGKA.Prop` the input type is standard and `CGKA*.Prop` when it is administrative (i.e., of the form $*-adm$). As in IAS, a validity check on the caller ID' and the target ID of the proposal is made using the `valid(P)` predicate (from IAS). Administrative proposals are signed with $\gamma.spk$; this is done to protect against insider adversaries that may re-send previously crafted administrative proposals (i.e.,

<p>DGS.Init($1^\lambda, ID$)</p> <hr/> <pre> 1: $\gamma.s0 \xleftarrow{\\$} \text{CGKA.Init}(1^\lambda, ID)$ 2: $\gamma.sA \xleftarrow{\\$} \text{CGKA}^*.\text{Init}(1^\lambda, ID)$ 3: $\gamma.ME \leftarrow ID; \gamma.1^\lambda \leftarrow 1^\lambda$ 4: $\gamma[\cdot].\text{spk}, \gamma[\cdot].\text{ssk} \leftarrow \perp$ </pre> <p>DGS.Create($\text{gid}, G, G^*; r_0$)</p> <hr/> <pre> 1: require $(\gamma.ME \in G^*) \wedge (G^* \subseteq G)$ 2: $(r_1, r_2, r_3, r_4) \leftarrow H_4(r_0, \gamma)$ 3: $(W_0, \gamma.s0) \leftarrow \text{CGKA.Create}(\gamma.s0, \text{gid}, G; r_1)$ 4: $(W_A, \gamma.sA) \leftarrow \text{CGKA}^*.\text{Create}(\gamma.sA, \text{gid}^*, G^*; r_2)$ 5: $T_{CR} \leftarrow \text{'create'}, W_0, W_A$ 6: $(\gamma[\text{gid}].\text{spk}, \gamma[\text{gid}].\text{ssk}) \leftarrow \text{Gen}(\gamma.1^\lambda; r_3); \text{registerKey}(\text{gid})$ 7: $\sigma_T \leftarrow \text{Sign}(\gamma.\text{ssk}, (\text{gid}, T_{CR}); r_4)$ 8: return $(\text{gid}, T_{CR}, \perp, \perp, \sigma_T)$ </pre> <p>DGS.Prop($\text{gid}, ID, \text{type}; r_0$)</p> <hr/> <pre> 1: $(r_1, r_2) \leftarrow H_2(r_0, \gamma)$ 2: if $\text{type} = *-\text{adm}$ 3: require $\gamma.ME \in \gamma.sA.G$ 4: $(\gamma.sA, P_0) \leftarrow \text{CGKA}^*.\text{Prop}(\gamma.sA, \text{gid}^*, ID, \text{type}; r_1)$ 5: $P \leftarrow (P_0, \text{Sign}(\gamma.\text{ssk}, P_0; r_2))$ 6: else 7: $(\gamma.s0, P) \leftarrow \text{CGKA.Prop}(\gamma.s0, \text{gid}, ID, \text{type}; r_1)$ 8: return P </pre> <p>DGS.Prop-Info(P)</p> <hr/> <pre> 1: $(\text{gid}, \text{type}, ID, ID') \leftarrow \text{CGKA.Prop-Info}(\gamma.s0, P)$ 2: if $\gamma.sA[\text{gid}] \neq \perp$ // Must be an admin proposal 3: $\text{type} \leftarrow \text{type} \parallel \text{-adm}$ 4: return $(\text{gid}, \text{type}, ID, ID')$ </pre> <p>DGS.Props(T)</p> <hr/> <pre> // Supports non-welcome control messages 1: $(\text{gid}, T_{CR}, T_W, T_C, \sigma_T) \leftarrow T$ 2: $\vec{P}_0 \leftarrow \text{CGKA.Props}(T_C.C_0); \vec{P}_A \leftarrow \text{CGKA}^*.\text{Props}(T_C.C_A)$ 3: return $\vec{P}_0 \parallel \vec{P}_A$ </pre>	<p>DGS.Commit($\text{gid}, \vec{P}, \text{com-type}; r_0$)</p> <hr/> <pre> 1: require $\gamma.ME \in \gamma.G$ 2: require $\text{com-type} \in \{\text{adm}, \text{std}, \text{both}\}$ 3: $(r_1, r_2, r_3) \leftarrow H_3(r_0, \gamma)$ 4: $C_0, C_A, W_0, W_A, k \leftarrow \perp$ 5: $(\vec{P}_0, \vec{P}_A, \text{admReq}) \leftarrow \text{propCleaner}(\text{gid}, \vec{P})$ 6: if $\text{admReq} \vee (\text{com-type} \in \{\text{adm}, \text{both}\})$ 7: require $\gamma.ME \in \gamma.G^*$ 8: $(\text{ssk}, \text{spk}) \leftarrow \text{getSigKey}(\gamma.sA.k)$ // old keys 9: if $\text{com-type} \in \{\text{adm}, \text{both}\}$ // update spk 10: $(\text{spk}, C_A, W_A) \leftarrow \text{c-Adm}(\text{gid}, \vec{P}_A; r_1)$ 11: if $\text{com-type} \in \{\text{std}, \text{both}\}$ 12: $(C_0, W_0) \leftarrow \text{c-Std}(\text{gid}, \vec{P}_0, k; r_2)$ 13: $T_C \leftarrow \text{'comm'}, C_0, C_A, W_A, \text{spk}$ 14: $T_W \leftarrow \text{'wel'}, W_0, \text{spk}$ 15: $\sigma_T \leftarrow \text{Sign}(\text{ssk}, (\text{gid}, T_C, T_W); r_3)$ 16: else // can be done by non-admins 17: $(C_0, \perp, k) \leftarrow \text{c-Std}(\vec{P}_0; r_1)$ 18: $T_C \leftarrow (\text{gid}, \text{'comm'}, C_0, \perp, \perp)$ 19: $T_W, \sigma_T \leftarrow \perp$ 20: if $k = \perp$ $k \leftarrow \gamma.s0.k$ 21: return $(\perp, T_C, T_W, \sigma_T, k)$ </pre> <p>DGS.Proc(T)</p> <hr/> <pre> 1: $(\text{gid}, T_{CR}, T_W, T_C, \sigma_T) \leftarrow T; \text{acc} \leftarrow \text{false}$ 2: if $T_{CR} \neq \perp \wedge T_W = T_C = \perp$ 3: if $\gamma.ME \in \gamma[\text{gid}].s0.G$ return false 4: $\text{acc} \leftarrow \text{p-Create}(\text{gid}, T_{CR}, \sigma_T)$ 5: else if $(\gamma.ME \notin \gamma[\text{gid}].G) \wedge (T_W \neq \perp)$ 6: $\text{acc} \leftarrow \text{p-Wel}(\text{gid}, T_C, T_W, \sigma_T)$ 7: else if $(\gamma.ME \in \gamma[\text{gid}].G) \wedge (T_C \neq \perp)$ 8: $\text{acc} \leftarrow \text{p-Comm}(\text{gid}, T_C, T_W, \sigma_T)$ 9: return acc </pre>
------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	--------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

Figure 5.10: Dynamic group signature (DGS) construction of an A-CGKA, built from two (possibly different) CGKAs CGKA and CGKA^* , a signature scheme Sig , and PRFs $H_n : \mathcal{R} \times \text{ST} \rightarrow \mathcal{R}^n$ for $n \leq 4$, randomness space \mathcal{R} and state space ST . We let $\gamma[\text{gid}].G = \gamma[\text{gid}].s0.G$, $\gamma[\text{gid}].G^* = \gamma[\text{gid}].sA.G$, and $\gamma[\text{gid}].k = \gamma[\text{gid}].s0.k$.

Chapter 5. Cryptographic Administration for Secure Group Messaging

<p>c-Adm(gid, $\vec{P}_A; r_1$)</p> <hr/> <pre> 1: for $P \in \vec{P}_A$ 2: if $P.type = \text{add-adm}$ 3: require $P.ID \in \gamma.G$ 4: $(\gamma.sA, (C_A, W_A), k) \leftarrow$ CGKA*.Commit($\gamma.sA, \text{gid}^*, \vec{P}_A; r_1$) 5: if $C_A = \perp$ return \perp 6: $(\text{ssk}, \text{spk}) \leftarrow \text{getSigKey}(k)$ 7: return (spk, C_A, W_A) </pre>	<p>p-Wel(gid, T_C, T_W, σ_T)</p> <hr/> <pre> 1: $(\text{msg-type}, W_0, \text{spk}) \leftarrow T_W$ 2: require $\text{msg-type} = \text{'wel'} \wedge \text{getSpk}(\text{gid}) = \text{spk}$ 3: if $\neg \text{Vrfy}(\text{spk}, (\text{gid}, T_C, T_W), \sigma_T)$ return false 4: $(\gamma', \text{acc}) \leftarrow \text{CGKA.Proc}(\gamma.s0, W_0)$ 5: if $\neg \text{acc}$ return false 6: $\gamma[\text{gid}].s0 \leftarrow \gamma'$ 7: $\gamma[\text{gid}].\text{spk} \leftarrow \text{spk}$ 8: return true </pre>
<p>c-Std(gid, $\vec{P}_0; r_2$)</p> <hr/> <pre> 1: $(\gamma.s0, C_0, W_0, k) \leftarrow \text{CGKA.Commit}(\gamma.s0, \text{gid}, \vec{P}_0; r_2)$ 2: return (C_0, W_0, k) </pre>	<p>p-Comm(gid, T_C, T_W, σ_T)</p> <hr/> <pre> 1: $(\text{msg-type}, C_0, C_A, W_A, \text{spk}) \leftarrow T_C$ 2: require $\text{msg-type} = \text{'comm'}$ 3: $\gamma' \leftarrow \gamma.sA$ 4: if $\sigma_T = \perp$ // no sig \Rightarrow check no changes to G 5: $(\gamma', \text{acc}) \leftarrow \text{CGKA.Proc}(\gamma.s0, C_0)$ 6: if $\neg \text{acc} \vee (\gamma'[\text{gid}].G \neq \gamma[\text{gid}].G)$ 7: return false 8: $\gamma[\text{gid}].s0 \leftarrow \gamma'$ 9: return true 10: else if $\neg \text{Vrfy}(\gamma[\text{gid}].\text{spk}, (\text{gid}, T_C, T_W), \sigma_T)$ 11: return false 12: if $\gamma.ME \in \gamma.G^*$ 13: $(\gamma', \text{acc}) \leftarrow \text{CGKA}^*.\text{Proc}(\gamma.sA, C_A)$ 14: if $\neg \text{acc}$ return false 15: else if $W_A \neq \perp$ 16: $(\gamma', \perp) \leftarrow \text{CGKA}^*.\text{Proc}(\gamma.sA, W_A)$ 17: if $C_0 \neq \perp$ 18: $(\gamma^\dagger, \text{acc}^\dagger) \leftarrow \text{CGKA.Proc}(\gamma.s0, C_0)$ 19: if $\neg \text{acc}^\dagger$ return false 20: $\gamma[\text{gid}].s0 \leftarrow \gamma^\dagger$ 21: if $\gamma.ME \notin \gamma^\dagger.G$ // removed user 22: $\gamma[\text{gid}] \leftarrow \perp$ 23: return true 24: $\gamma[\text{gid}].sA \leftarrow \gamma'; \gamma[\text{gid}].\text{spk} \leftarrow \text{spk}$ 25: return true </pre>
<p>propCleaner(gid, \vec{P})</p> <hr/> <pre> 1: $\text{admReq} \leftarrow \text{false}; \vec{P}_0, \vec{P}_A \leftarrow []$ 2: for $P \in \vec{P}$ 3: $(\text{gid}', \text{type}, \text{ID}, \text{ID}') \leftarrow \text{prop-info}(P)$ 4: if $\text{gid}' = \text{gid}^* \wedge P.type = \text{'*adm'} \wedge \text{IAS.valid}(P)$ 5: $\vec{P}_A \leftarrow [\vec{P}_A, P]$ 6: $\text{admReq} \leftarrow \text{true}$ 7: else if $\text{gid}' = \text{gid}$ 8: $\vec{P}_0 \leftarrow [\vec{P}_0, P]$ 9: if $\text{type} \in \{\text{add}, \text{rem}\}$ 10: $\text{admReq} \leftarrow \text{true}$ 11: // admin rem from G \Rightarrow rem also from G^* 12: if $(\text{type} = \text{rem}) \wedge (\text{ID} \in \gamma.G^*)$ 13: $P' \leftarrow \text{CGKA}^*.\text{Prop}(\gamma.sA, \text{gid}^*, \text{ID}, \text{rem})$ 14: $\vec{P}_A \leftarrow [\vec{P}_A, P']$ 15: $(\vec{P}_0, \vec{P}_A) \leftarrow \text{enforcePolicy}(\vec{P}_0, \vec{P}_A)$ 16: return $(\vec{P}_0, \vec{P}_A, \text{admReq})$ </pre>	<p>p-Create(gid, T_{CR}, σ_T)</p> <hr/> <pre> 1: $(\text{msg-type}, W_0, W_A) \leftarrow T_{CR}$ 2: require $\text{msg-type} = \text{'create'}$ 3: $(\gamma_0, \text{acc}) \leftarrow \text{CGKA.Proc}(\gamma.s0, W_0)$ 4: $(\gamma_A, \perp) \leftarrow \text{CGKA}^*.\text{Proc}(\gamma.sA, W_A)$ 5: if $\gamma_A[\text{gid}].G = \perp \vee (\emptyset \neq \gamma_A[\text{gid}].G \subseteq \gamma_0[\text{gid}].G)$ 6: $(\gamma.s0, \gamma.sA) \leftarrow (\gamma_0, \gamma_A)$ 7: else return false 8: return $\text{acc} \wedge \text{Vrfy}(\text{getSpk}(\text{gid}), (\text{gid}, T_{CR}), \sigma_T)$ </pre>
	<p>getSigKey(r)</p> <hr/> <pre> 1: $(\text{ssk}, \text{spk}) \leftarrow \text{KeyGen}(1^\lambda; H_{ro}(r))$ 2: return (ssk, spk) </pre>
	<p>enforcePolicy(\vec{P}_0, \vec{P}_A)</p> <hr/> <pre> 1: return $(\vec{P}_0, \vec{P}_A) \leftarrow \text{IAS.enforcePolicy}(\vec{P}_0, \vec{P}_A)$ </pre>

Figure 5.11: Helper functions for DGS in Figure 5.10 w.r.t. random oracle $H_{ro} : \{0, 1\}^\lambda \rightarrow \{0, 1\}^\lambda$.

those that are legitimate but correspond to a previous epoch), or even create new ones if these are sent in plaintext. The Prop-Info algorithm is fully based on the respective Prop-Info algorithms (which in principle may not necessarily output fourth value ID' if anonymous proposals are allowed); we assume Props is likewise inherited from the underlying CGKAs.

Commits. For a given gid, administrative changes are committed via $CGKA^*.Commit$ (which outputs (C_A, W_A)) and standard group changes via $CGKA.Commit$ (outputting (C_0, W_0) as usual). Note that the $CGKA^*$ secret k_{adm} associated with the $CGKA^*.Commit$ call is used to generate each admin signature key pair (ssk', spk') .

The new admin key spk' is included in the final A-CGKA commit message so that (non-admin) group members can process it. In order to prove the authenticity of the commit (and of spk'), the committer signs the whole commit message including C_A, W_A, C_0, W_0 and spk' with the old admin signing key $\gamma.ssk$. In addition, the committer must verify all proposal signatures in advance.

As before, a commit can be split into a welcome message T_W for newly added users, and a commit message T_C for group members. These are signed jointly in our construction to simplify the security proof, but may also be signed separately. In T_C , we also include the welcome messages to $CGKA^*$, since they must always be addressed to current group members (i.e., of G). In a given Commit call, one or both of $CGKA$ and $CGKA^*$ may be updated, and in particular $CGKA$ but not $CGKA^*$ (for a non-admin commit) and vice-versa (for an admin-only commit).

Processing Control Messages. The Proc method takes a control message T , determines the type of message (create, welcome, or commit) and the gid, and updates the state only if processing succeeds ($acc = true$). Newly added users verify the admin signature, process the welcome message using $CGKA.Proc$ and store the new public admin key (spk' , provided in T) in $\gamma.spk$.

Group members verify the administrator signature (if the commit requires administrative rights) using $\gamma.spk$. Then, depending on the commit type at least one of $CGKA$ and $CGKA^*$ are updated via the corresponding $CGKA.Proc$ algorithm. Given $CGKA^*$ or both $CGKAs$ are updated, the updated admin key is set as $\gamma.spk \leftarrow spk'$. In case T contains a create message T_{CR} , both $CGKAs$ process the respective welcome messages contained in T_{CR} separately.

Features

DGS allows the use of two distinct and independent $CGKA$ protocols that authenticate admins as a group, providing some notable features that differ from IAS. One can also imagine a 'hybrid' approach where users run IAS except that some IAS keys are maintained and updated via their own $CGKA$ and/or a hierarchy of $CGKAs$ is employed.

Chapter 5. Cryptographic Administration for Secure Group Messaging

Minimal Information Reveal. As opposed to IAS, the set of group administrators can be opaque to the central server and to the rest of the group (whenever the underlying CGKAs preserve the anonymity of group members with respect to external parties); this is reflected in the way DGS uses a PKI.

Limitations. A drawback of DGS is that enforcing different “levels of administration”, for which IAS can be easily extended, is not straightforward. Nevertheless, one can still implement minor policies such as muting users at an application level (as done in practice). We also note that admins may not have a reliable view of the set of admins if $CGKA^*$ is susceptible to insider attacks that violate robustness¹⁰. If these attacks are relevant, one can deploy heavier protocols such as the P-Act-Rob protocol from Alwen et al. [ACJM20]. A third limitation is that admins cannot give up their admin status immediately; they must send a self rem-adm proposal, erase their admin state, and wait for another admin to commit. This occurs generally in CGKA when a member leaves a group; to minimise this delay, the MLS standard requires users to commit immediately upon receipt of a valid proposal (in this case a removal). This problem could nevertheless be solved using the same approach as for self-removes in IAS, i.e., allowing admin self-removes without an admin signature.

Security Mechanisms. We note the conceptual simplicity of achieving PCS and FS for the group administration keys (in the adversarial model for $CGKA^*$) given the existence of secure CGKA schemes in the literature, since both properties are ensured by $CGKA^*$ itself. Update mechanisms are largely simplified due to a single admin key being used. Delegation and revocation of admin keys are also straightforward.

Correctness and Security

Theorem 18. Let $CGKA$ and $CGKA^*$ be correct CGKAs, and Sig be a 1-correct signature scheme. Then, the DGS protocol (Figures 5.10 and 5.11) is correct with respect to Definition 42.

Proof. We prove that no adversary \mathcal{A} can win $CORR_{A-CGKA, C_{corr}}$ (where we set $C_{corr} = \text{true}$) played with respect to DGS (Figures 5.10 and 5.11) given that $CGKA, CGKA^*$ are correct. We omit some details which are analogous to IAS’ correctness proof (note that IAS and DGS are designed such that they share sections of their code).

For PROP, the correctness of Prop-Info follows from the correctness of $CGKA.Prop-Info, CGKA^*.Prop-Info$ by assumption. Also, the adversary cannot win after the group membership check as Prop only modifies the state by calling $CGKA.Prop$ and $CGKA^*.Prop$; which are correct by assumption.

The group membership check must always pass in CREATE and COMMIT for identical reasons.

¹⁰This scenario is out of the scope of our security model where admins are fully trusted.

For DELIVER, we examine each **reward** clause as done previously with IAS. As before, note that in line 2 of DELIVER, T is a commit message created either by Create or by Commit. Also, it is easy to see that the CGKA.Props check in PROP holds.

The clause $(\emptyset \neq \gamma[\text{gid}].G^* \subseteq \gamma[\text{gid}].G)$ is met upon generation of any create message T (produced by CREATE) by construction (line 1 of Create). When any message is processed by p-Create, the condition is enforced again.

If T is a (standard) commit made by a non-admin user – that is, one without a signature – then there are no changes to the group in the commit as checked explicitly by the auxiliary c-Std upon commit. The message must only contain a T_C which is processed by p-Comm, which again enforces this condition.

If T an admin commit, then it is processed by p-Comm or by p-Wel. In both cases, commit messages are processed by the underlying CGKA, CGKA* methods. Therefore, correctness depends on the Commit algorithm. The case distinction follows the exact same logic as in IAS, since the algorithms propCleaner and enforcePolicy, and the predicate valid(P) enforce the same conditions.

The next case is the **reward** $\gamma[\text{gid}].k \neq \perp$ for a removed (or non-member) ID , which is enforced in DGS by p-Comm and by the correctness of the CGKAs.

For the last check by UpdateView, one can proceed by induction as in IAS. The main difference is that the admin update is not done manually as in IAS (i.e., modifying adminList), but rather by the underlying Proc algorithms of CGKA*, which yields the result easily. The last check for the consistency of the derived group key also follows as in IAS. We omit the details. \square

Theorem 19. Consider KIND CGKA CGKA with respect to cleanness predicate C_{cgka} , KIND CGKA CGKA* with respect to C_{cgka^*} , SUF-CMA signature scheme Sig and PRF H_4 . Then, for A-CGKA DGS, we have, with respect to cleanness predicates $C_{\text{cgka}}, C_{\text{adm}}, C_{\text{forgery}} = C_{\text{forgery}}^*$, that for every efficient adversary \mathcal{A} that makes at most q oracle queries, one can build an adversary \mathcal{B} such that

$$\text{Adv}_{\text{DGS}}^{\text{kind}}(\mathcal{A}) \leq \text{Adv}_{\text{CGKA}}^{\text{kind}}(\mathcal{B}) + q \cdot \left(\text{Adv}_{H_4}^{\text{prf}}(\mathcal{B}) + \text{Adv}_{\text{Sig}}^{\text{suf-cma}}(\mathcal{B}) + q_{\text{ro}} \cdot \text{Adv}_{\text{CGKA}^*}^{\text{kind}}(\mathcal{B}) + 2^{-\lambda} \right),$$

where C_{adm} is a function of C_{cgka^*} defined below and C_{forgery}^* is defined in Section 5.2.4.

Proof idea. We first describe C_{adm} . Intuitively, C_{adm} ensures that the set of safe oracle queries for DGS adversary \mathcal{A} given inject queries of the form $q_i = \text{INJECT}(ID, m, t_a)$ are those that are safe for CGKA* adversary \mathcal{A}' under essentially the same queries, replacing $\text{INJECT}(\cdot, \cdot, t_a)$ queries with queries of the form $\text{CHAL}(t_a)$. To prove security, we use a similar game-hopping argument as in IAS. We first replace H_i calls using the PRF assumption. We then consider E_1 (a successful injection is made) and E_2 (otherwise) as in IAS. Given E_2 , we can simulate directly via the CGKA adversary. Given E_1 , we simulate differently depending on whether \mathcal{A} makes a query to random oracle H_{ro} with a correct CGKA* key before the successful injection

Chapter 5. Cryptographic Administration for Secure Group Messaging

or not. Here, if \mathcal{A} is successful, we reduce security to the CGKA* adversary by intercepting the relevant random oracle query and guessing the correct bit using information from CHAL. Otherwise, we can simulate via an EUF-CMA adversary as in IAS since the signature key is now uniform from the adversary's perspective.

Predicate C_{adm} . C_{adm} is tailored to DGS and is a function of the underlying CGKA* predicate C_{cgka^*} . Intuitively, C_{adm} ensures that the set of safe oracle queries for DGS adversary \mathcal{A} given inject queries of the form $q_i = \text{INJECT}(ID, m, t_a)$ are those that are safe for CGKA* adversary \mathcal{A}' (i.e., the predicate C_{cgka^*}) under roughly the same queries, replacing at most one $\text{INJECT}(\cdot, \cdot, t_a)$ queries with a $\text{CHAL}(t_a)$ query. For example, noting the symmetry between predicates $C_{\text{adm-opt}}$ and $C_{\text{cgka-opt}}$, if CGKA* is secure with respect to CGKA predicate $C_{\text{cgka-opt}}$, then DGS is secure with respect to admin predicate $C_{\text{adm-opt}}$.

We define C_{adm} more formally. Let $Q = (q_i)_I$ be the ordered sequence of oracle queries made by the DGS adversary \mathcal{A} . To define C_{adm} , we construct an ordered sequence of queries Q^* that are made by the CGKA* adversary \mathcal{A}' in the security proof below by replacing, inserting and/or deleting queries in-order. Let $\ell \in [1, q_{inj}] \cup \{\perp\}$ where q_{inj} is the number of INJECT queries made by \mathcal{A} . To this end, consider each $q_i \in Q$ and, for each ℓ , define q_i^* to be either a single query or a sequence of queries in Q^* as follows:

- $q_i = \text{CREATE}(ID, G, G^*)$: Set $q_i^* = \text{CREATE}(ID, G^*)$.
- $q_i = \text{PROP}(ID, ID', \text{type})$: Set $q_i^* = \perp$ if $\text{type} \neq \text{* - adm}$ and $q_i^* = \text{PROP}(ID, ID', \text{type}^*)$ otherwise where $\text{type} = \text{type}^* \text{- adm}$.
- $q_i = \text{COMMIT}(ID, (i_1, \dots, i_k), \text{com-type})$: If the condition $(\text{admReq} \vee \dots)$ at line 6 of Commit is false, $\text{com-type} = \text{std}$ or ID is not currently an admin, set $q_i^* = \perp$. Otherwise, let $\{ID_1, \dots, ID_j\}$ be the (possibly empty) set of parties for which CGKA* rem proposals are introduced by propCleaner (line 12). Let \vec{P}_A be the value input to CGKA*.Commit at line 4 of c-Adm (or $\vec{P}_A = \perp$ if the line is not reached), and (i_1, \dots, i_k) the corresponding proposal indices in the CGKA* KIND game. Let q' be the REVEAL query that reveals the key k output by Commit in the COMMIT call if the corresponding signature key pair is not used for the first successful INJECT query, and \perp otherwise. Then, set q_i^* to the sequence $(\text{PROP}(ID, ID_1, \text{rem}), \dots, \text{PROP}(ID, ID_j, \text{rem}), \text{COMMIT}(ID, (i_1, \dots, i_k), q'))$.
- $q_i = \text{CHAL}(t_s)$: Set $q_i^* = \perp$.
- $q_i = \text{DELIVER}(ID, (t_s, t_a), c)$: Let $(T, \text{com-type}) = T[(t_s, t_a), \text{com}, c]$ for DGS KIND game variable T . If Proc is called by the game, $ID \in G$ holds, $T_c \neq \perp$ holds and either $ID \in G^*$ holds or W_A (contained in T) is $\neq \perp$, set $q_i^* = \text{DELIVER}(ID, t_a, c^*)$, where $c^* \leq c$ is the number of times CGKA*.Commit was called after c queries to COMMIT. Otherwise, set $q_i^* = \perp$.
- $q_i = \text{REVEAL}(t_s)$: Set $q_i^* = \perp$.
- $q_i = \text{EXPOSE}(ID)$: Set $q_i^* = q_i$.

- $q_i = \text{INJECT}(ID, m, t_a)$: Set $q_i^* = \perp$ if q_i^* is the j -th query to INJECT where $j \neq \ell$ (possibly \perp) and $\text{CHAL}(t_a)$ otherwise ($\ell \neq \perp$).

Then, C_{adm} is defined to be true if C_{cgka^*} is true for runs where the first successful INJECT query is the ℓ -th query to INJECT (where the \perp -th query denotes no successful injection).

We prove DGS secure below.

Proof. Following the proof of Theorem 17 we let G_0 be the $\text{KIND}_{\text{A-CGKA}, C_{\text{cgka}}, C_{\text{adm}}, C_{\text{forgery}}}$ game, and G_1 be as in G_0 except that all calls of the form $(r_1, \dots, r_i) \leftarrow H_i(r_0, \gamma)$ are replaced with calls of the form $(r_1, \dots, r_i) \stackrel{\$}{\leftarrow} R^i$. We transition between G_0 and G_1 exactly as in Theorem 17. That is, we define hybrids $G_{0,j}$ where $G_{0,0} = G_0$, $G_{0,j} = G_1$ when $j \geq q$ and $G_{0,j}$ differs from G_0 for appropriate $0 < j < q$ by replacing the first j calls to functions of the form H_i with uniformly sampled values by the challenger. As before, we have $|\Pr[G_{0,j-1}(\mathcal{A}) \Rightarrow 1] - \Pr[G_{0,j}(\mathcal{A}) \Rightarrow 1]| \leq \epsilon_F$ for each $j \geq 1$, where ϵ_F is the advantage of PRF adversary \mathcal{A}' , which implies also that

$$\left| \Pr[G_0(\mathcal{A}) \Rightarrow 1] - \frac{1}{2} \right| \leq \left| \Pr[G_1(\mathcal{A}) \Rightarrow 1] - \frac{1}{2} \right| + q \cdot \text{Adv}_{H_4}^{\text{prf}}(\mathcal{A}').$$

Let E_1 be the event that \mathcal{A} queries INJECT such that the oracle does *not* output \perp (i.e., it outputs the challenge bit). Let E_2 be the event that this does not occur; clearly $\Pr[E_1] + \Pr[E_2] = 1$. To prove security, we will reduce to the security of the primary CGKA, i.e., CGKA, given E_2 , and to CGKA^* and signature security given E_1 .

We first consider the simpler E_2 case where no successful injection is made (and thus INJECT calls can be easily simulated); let \mathcal{A}' be a KIND adversary w.r.t. CGKA simulating for G_1 adversary \mathcal{A} given E_2 . \mathcal{A}' simulates as follows. For each oracle query, \mathcal{A}' simulates relevant CGKA^* calls locally unless stated otherwise. In particular, \mathcal{A}' simulates $\text{Init}(1^\lambda, ID)$ for ID only as needed (i.e., lazily). Then:

- $\text{CREATE}(ID, G, G^*)$: \mathcal{A}' calls $\text{CREATE}(ID, G)$ if needed and otherwise locally simulates.
- $\text{PROP}(ID, ID', \text{type})$: \mathcal{A}' simulates locally if type is of the form $*$ -adm and simulates via $\text{PROP}(ID, ID', \text{type})$ otherwise.
- $\text{COMMIT}(ID, (i_1, \dots, i_k), \text{com-type})$: Since CGKA.Commit is called after $\text{CGKA}^*.Commit$, \mathcal{A}' can simulate CGKA^* calls locally and call $\text{COMMIT}(ID, J = (j_1, \dots, j_k))$, where J corresponds to the set of relevant CGKA proposal indices derived from (i_1, \dots, i_k) , to simulate CGKA.Commit calls.
- $\text{DELIVER}(ID, (t_s, t_a), c)$: \mathcal{A}' simulates the relevant CGKA.Proc call (there is at most one such call made by construction of DGS Proc) via $\text{DELIVER}(ID, t_s, c')$ where c' is the index of the relevant CGKA control message. \mathcal{A}' simulates locally otherwise.
- $\text{REVEAL}(t_s)$ and $\text{CHAL}(t_s)$: \mathcal{A}' simulates directly via their respective oracles.

Chapter 5. Cryptographic Administration for Secure Group Messaging

- EXPOSE(ID): \mathcal{A}' calls EXPOSE(ID) and simulates the rest of the call locally.
- INJECT: By definition of E_2 , INJECT always returns \perp , and since INJECT does not modify the state, \mathcal{A}' simply outputs \perp upon each INJECT call.

By construction, DGS inherits the (normal) CGKA cleanness predicate C_{cgka} from CGKA, and so the simulation is perfect and it follows thus that:

$$\Pr[G_1(\mathcal{A}) \Rightarrow 1 \wedge E_2] \leq \text{Adv}_{\text{CGKA}}^{\text{kind}}(\mathcal{A}').$$

Now, consider adversary \mathcal{A} playing G_1 given E_1 occurs, i.e., \mathcal{A} makes a successful INJECT call. Note that there are at most q different CGKA* epochs during a given execution and consequently at most q different CGKA* signature key pairs computed by correct parties (since each signature key pair is derived from a given CGKA* epoch secret).

Given E_1 , let F_1 be the event that G_1 adversary \mathcal{A} calls random oracle H_{ro} with input r that corresponds to the signature key pair used in the successful INJECT query (guaranteed to exist by Lemma 16) *before* the injection is made. That is, r is such that $(\text{ssk}, \text{spk}) \leftarrow \text{Gen}(1^\lambda; H_{\text{ro}}(r))$ is called by the challenger at some point. Let F_2 be the complementary event (i.e., either such a H_{ro} query is made after the successful injection or not at all). We consider the case with F_1 (by simulating via the CGKA* KIND game) and F_2 (via the SUF-CMA game) separately.

Consider F_1 . Let $F_{i,j}$ be the event that the aforementioned $H_{\text{ro}}(r)$ query is the i -th query to H_{ro} and is with respect to the j -th CGKA* key pair sampled by oracle queries during the game's execution; clearly at most $q \cdot q_{\text{ro}}$ such events occur. Let \mathcal{A}' be a CGKA* adversary who simulates for G_1 adversary \mathcal{A} as follows given $F_{i,j} \wedge E_1$. We assume KIND is such that \mathcal{A}' can reveal secrets k output by Commit from oracle COMMIT. Note that our predicate C_{adm} is designed for this simulation.

\mathcal{A}' simulates by calling relevant CGKA* oracles and simulating locally otherwise. When \mathcal{A}_i calls getSigKey while simulating Commit calls, H_{ro} is queried with input r ; \mathcal{A}_i lazily samples in this case. When \mathcal{A} queries INJECT, \mathcal{A}' simply returns \perp . Note that COMMIT invokes Commit but does not output the key k that Commit outputs. When \mathcal{A}' first derives key k from COMMIT for the x -th CGKA* key pair, for $x \neq j$, \mathcal{A} reveals secret k output by Commit; \mathcal{A} can thus simulate DGS algorithm Commit perfectly. When \mathcal{A}' makes their i -th query to H_{ro} with input r , \mathcal{A}' makes a CHAL query for the corresponding epoch, which outputs k' ; \mathcal{A}_i finishes simulating and returns bit 0 if and only if $k' = r$. The simulation is perfect, and when $b = 0$ \mathcal{A}_i wins iff \mathcal{A} wins, since \mathcal{A}_i outputs 0 only if they derive the correct key or signature public key in the simulation, and when $b = 0$ the challenge oracle outputs the correct key. The $b = 1$ case is similar except in the case that $b = 1$ and the r sampled by the game is the same as the real key (which happens with probability $\frac{1}{2^\lambda}$); it follows that

$$\Pr[G_1(\mathcal{A}) \Rightarrow 1 \wedge E_1 \wedge F_1] \leq q \cdot q_{\text{ro}} \cdot \Pr[\text{KIND}_{\text{CGKA}^*, C_{\text{cgka}^*}}(\mathcal{A}') \Rightarrow 1] + \frac{q}{2^\lambda}.$$

We consider F_2 . Let F'_i be the event, for $1 \leq i \leq q$, that a successful injection is made which uses the i -th CGKA* key pair sampled by oracle queries during the game's execution. Note that such a key pair must exist and that a signature forgery, by algorithm construction, is a necessary but not sufficient condition to make a INJECT query with a non-bottom response. Consider \mathcal{A}_i who simulates as follows. \mathcal{A}_i simulates all queries locally except that \mathcal{A}_i embeds his challenge key in the i -th such CGKA* key pair in relevant control messages and uses the SIGN oracle to produce signatures as necessary. Note that the safety predicate is such that, conditioned on F'_i , that an EXPOSE query that leaks the corresponding signature key is disallowed. In addition, \mathcal{A} does not make any H_{ro} query that would lead to a trivial exposure by definition of F_2 , and thus the distribution of the challenge key pair is uniform (i.e., correct). It follows that the simulation is perfect and that:

$$\Pr[G_1(\mathcal{A}) \Rightarrow 1 \wedge E_1 \wedge F_2] \leq q \cdot \text{Adv}_{\text{Sig}}^{\text{uf-cma}}(\mathcal{A}')$$

The proof is completed by combining the sequence of game hops considered hitherto. \square

5.3.4 Integrating A-CGKA into MLS

Some group messaging protocols already authenticate group members via signatures and public-key infrastructure. The MLS specification [BBR⁺23] relies on credentials, which are essentially public signature keys for each protocol user that are certified by a PKI; these keys authenticate messages originating from that user¹¹. Therefore, it is possible to extend the CGKA used in MLS to an A-CGKA in a more efficient way than using a compiled A-CGKA construction resembling IAS. We note that, in practice, it is feasible to support secure administration in MLS via an *MLS extension*, a feature that enables additional proposal types and actions in the protocol [BBR⁺23]. Constructing such an extension is almost straightforward and we identify three main necessary changes:

- Credentials are not necessarily refreshed in MLS, meaning that admins (and users in general) whose state is compromised at some point lack forward security and post-compromise security on their authentication keys (unless they proactively update them). Our solution is to enforce an IAS-like credential update mechanism for admin signature keys (providing post-compromise security) which may be invoked without updating the core CGKA secret.
- Group members need to keep track of the administrators (for an IAS-like extension). To this end, we introduce new admin proposal types and enforce that admin proposals are signed, alongside corresponding update policies and modifications to Commit and Proc.
- As in IAS, admins register their keys with the PKI as they are updated over time.

¹¹Signatures play an important role in MLS: “...group members can verify a message originated from a particular member of the group. This is guaranteed by a digital signature on each message from the sender's signature key. The signature keys held by group members are critical to the security of MLS against active attacks...” [BBR⁺23].

Modifications. We propose an extension of the main algorithms of the MLS protocol (in particular, of the CGKA-related Prop, Commit and Proc) in Figure 5.12, that we also benchmark in Section 5.4.1. Our goal is to show how an IAS-like protocol can be easily achieved with relatively low overhead. We follow Alwen et al. [ACDT21a] (in particular, Figure 8 in the full version [ACDT21b]), as this is the most comprehensive formalisation of MLS in the literature at the time of writing; therefore, we also work in the single-group setting and omit gids. We omit the Send and Receive algorithms as these are used to send application messages only. We note also that their Create method supports only one initial participant.

Protocol Details. The main modifications are to (1) the admins' credentials, which are regularly updated via upd-adm proposals and admin commits; and (2) in the introduction of the three additional proposal types from A-CGKA. For brevity, we omit several parts of the protocol, such as some sanity checks (e.g., some **require** predicates), functionality that we do not need to modify and details on a higher level than CGKA (like the use of a MAC by MLS). Also for simplicity, we extend the CGKA state γ to include the state variables used in IAS. Following Alwen et al. [ACDT21a], we split the processing algorithms in two – one for commit messages, and one for welcome messages – that a committer produces for each incoming user separately.

Overall, the overhead with respect to (bare-bones) MLS is minimal; we essentially only need to support the new types of proposals and to refresh admin credentials for admin updates. Most of the protocol logic relates to updating signatures and the adminList. Note that proposals are always signed in MLS so signing within makeAdminProp can be foregone. We also support self-removal proposals that can be committed by standard users.

Correctness and Security. We leave it open to formally propose and prove correctness and security for an appropriate MLS extension; we sketch here how it could be done. The modelling of messaging and MLS in particular of Alwen et al. [ACDT21a] is more complex than ours. In particular, they consider CGKA as a sub-primitive that is used to build secure group messaging (SGM) alongside several other primitives. Thus, one could re-define SGM to account for new proposal types and administration as we have done for A-CGKA, including admin correctness guarantees and security upon injections from non-admins. Since admin proposals are tightly-coupled with protocol flow, proposing a model 'on top' of theirs to provide admin security seems difficult, although modular guarantees would be ideal. Proving correctness and security then boils down to similar case analysis and reductions to theirs and ours.

5.4 Evaluation and Discussion

In this section, we describe our implementation and corresponding experimental results for our MLS extension proposed in Section 5.3.4. We then discuss how our modelling in this chapter compares to the related work.

<pre> Prop($ID, type; r_0$) 1: if $type = *-adm$ 2: require $\gamma.ME \in \gamma.G^*$ 3: $(P, \perp) \leftarrow IAS.makeAdminProp(type, ID, r_0)$ // getSpk is replaced in makeAdminProp 4: if $type \in \{add, rem, upd\}$ 5: $(\gamma, P) \leftarrow CGKA.Prop(\gamma, ID, type; r_0)$ // Added users' keys retrieved from contact list/PKI 6: $\sigma \leftarrow Sign(\gamma.ssk, P)$ // all MLS proposals are signed 7: return (P, σ) Commit($(\vec{P}_0, \vec{P}_A), com\text{-}type; r_0$) 1: $(r_1, r_2, r_3) \leftarrow H_3(r_0, \gamma)$ 2: if $com\text{-}type \in \{adm, both\}$ 3: require $\gamma.ME \in \gamma.G^*$ 4: require $IAS.verifyPropSigs(\perp, \perp, \vec{P}_A)$ 5: $C_A \leftarrow \vec{P}_A$ 6: $adminList' \leftarrow IAS.updAL(adminList, \vec{P}_A)$ 7: $(\gamma.ssk', \gamma.spk') \leftarrow KeyGen(\gamma.I^A; r_1)$ 8: $\gamma \leftarrow updSpk(\gamma, ID, spk')$ 9: if $com\text{-}type \in \{std, both\}$ 10: $(\gamma, C_0, W_0, \perp) \leftarrow CGKA.Commit(\vec{P}_0; r_2)$ 11: if $W_0 \neq \perp$ // share updated adminList 12: Prepare welcome msgs as in [ACDT21a] for \vec{W} 13: for $W \in \vec{W}$: 14: $W \leftarrow W adminList'$ 15: $\sigma \stackrel{\\$}{\leftarrow} Sign(\gamma.ssk, W)$ // rand. 16: $T \leftarrow ('com', \gamma.ME, C_0, C_A, \gamma.spk')$ 17: $\sigma \leftarrow Sign(\gamma.ssk, T; r_3)$ 18: return $((T, \sigma), \vec{W})$ </pre>	<pre> Proc-WM(W) // ID is the committer of W 1: require $W.adminList[ID] \neq \perp$ 2: Run Proc-WM(W) in [ACDT21a] 3: $\gamma.adminList \leftarrow W.adminList$ 4: Check $adminList[ID]$ with PKI Proc-CM(T, σ) 1: $(com', ID, C_0, C_A, spk') \leftarrow T$ 2: if $adminList[ID] = \perp$ 3: require $Vrfy(getSpk(ID), T, \sigma)$ 4: Run Proc-CM(T) in [ACDT21a] 5: require no membership changes to $\gamma.G$ except self-removals 6: if $adminList[ID] \neq \perp$ 7: require $Vrfy(adminList[ID], T, \sigma)$ 8: if $spk' \neq \perp$ // spk was registered 9: require $spk' = getSpk(ID)$ 10: Update keys and adminList as in IAS 11: $IAS.p\text{-}Comm(T)$ 12: Check new adminList keys with PKI getSpk(ID) // Get spk from ID's credential 1: return $Cred[ID].spk$ updSpk(γ, ID, spk') // Register spk' with the PKI 1: $\gamma \leftarrow registerPKI(\gamma, ID, spk')$ // Update ID's credential 2: $\gamma.Cred[ID].spk \leftarrow spk'$ </pre>
-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

Figure 5.12: Construction of an MLS extension that supports group administrators, effectively turning the CGKA in MLS into an A-CGKA. Highlighted lines correspond to our main modifications in the original secure group messaging (SGM) construction of Alwen et al. [ACDT21a]. $Cred[\cdot]$ denotes a dictionary that stores the credentials of all ID 's. We also use the abstract function $registerPKI$ for standard PKI functionality of registering signature keys and $getSpk$ for fetching keys as in IAS. WM (resp. CM) denotes ‘welcome message’ (resp. ‘commit message’). Some technical details are omitted.

5.4.1 Benchmarks and Performance

We implemented the protocol in Section 5.3.4 to obtain a realistic estimate of the overhead of securely administrating a real-world messaging protocol. We modified an open-source

Chapter 5. Cryptographic Administration for Secure Group Messaging

implementation of MLS in Go¹² and compare the running times of MLS (which also performs e.g. parent hashing and non-admin proposal signing), with the running times of administrated MLS in different scenarios. In particular, we analyze the Commit and Proc algorithms in Figure 5.12, where the latter includes Proc-CM and also processing proposals when relevant (done separately in the implementation). We ran our benchmarks on a laptop with a 4-core 11th Gen Intel i5-1135G7 processor and 16 GB of RAM using Go's `testing` package¹³. Core cryptographic operations were implemented as HPKE [BBLW22] with ciphersuite DHKEM(P-256, HKDF-SHA256), HKDF-SHA256, AES-128-GCM (in particular using ECDSA for signatures) from Go standard libraries. We measured the time taken for a *single group member* to perform the relevant operation. For each data point, we took the average over 100 iterations that randomised the group members and admins performing group operations, as performance can be affected by their position in MLS's TreeKEM.

Our results are displayed in Figures 5.13 and 5.14, where we show the running time of the Commit (Figure 5.13) and Proc (Figure 5.14) algorithms in different realistic scenarios. We run experiments where relevant, i.e., when there are no admin operations, using the original implementation as a baseline, as well as using our modified implementation, to demonstrate that the additional admin logic we introduce does not noticeably affect performance.

On the one hand, we present the running time of both algorithms for varying group size $|G|$ with a fixed member/admin ratio $|G|/|G^*| = 4$. In the event of updating users, there are $t = |G^*|$ users updating and/or $t/2$ admins doing an admin-update. On the other hand, we benchmark our algorithms for fixed group size $|G| = 64$, $|G^*| = 16$, while varying the number of updates t (and/or $t/2$ updating admins) in a commit.

For both cases, we compare the committing and processing times of (1) standard commits (com-type = std, omitted in fixed group size benchmarks), (2) standard commits with t update proposals, (3) standard and admin commits (com-type = both) with $t/2$ admin-update proposals but no standard-update proposals, and (4) standard and admin commits with t update and $t/2$ admin-update proposals.

Communication Overhead. In both the baseline and our implementations, proposals used 364 to 366 bytes, and admin proposals used 364 to 368 bytes (all proposals being signed). Commit message sizes in both implementations vary proportionally with the size of the group and the number of proposals. In the *baseline MLS implementation*, a typical Commit for $|G| = 8$ and $|G| = 128$ with $t = 2$ and $t = 32$ update proposals uses 1.49 KB and 17.11 KB respectively. In *our implementation*, corresponding commits use 1.56 KB and 17.17 KB respectively. If $t/2$ admin updates are added (1 and 16, respectively), commits require 1.60 KB and 17.65 KB. In general, commits in our implementation, even with admin proposals, incur only a small amount of overhead (tens of bytes) over the baseline implementation when fixing the number of proposals.

¹²The original source code is available at <https://github.com/cisco/go-mls>.

¹³<https://pkg.go.dev/testing>




Figure 5.13: Benchmark of the Commit algorithm in the following scenarios: (1) standard commits (com-type = std, omitted right), (2) standard commits with t update proposals, (3) standard and admin commits (com-type = both) with $t/2$ admin-update proposals but no standard-update proposals, and (4) standard and admin commits with t update and $t/2$ admin-update proposals. Original MLS is displayed as baseline. *Left*: running time with respect to group size $|G|$ on constant member/admin $|G|/|G^*| = 4$ ratio and constant number of updates $t = |G^*|$ ($t/2$ admin updates). *Right*: running time with respect to the number of updating users t (and $t/2$ admin updates), for fixed $|G| = 64$.

Protocols. The results above show that the additional cost (for users) of running a securely-administrated MLS is minimal. Figure 5.13 shows that the Commit algorithm involves less than a 20% overhead when up to $|G|/8$ members carry out admin updates simultaneously (note that admin updates also involve standard updates). Figure 5.14 shows that the processing time of admin and standard updates is very similar, and increases linearly in the number of updates.

Separately, we analyze the overhead of IAS and DGS for group members, both for number of operations and for message size. We note that this assumes that IAS and DGS are implemented modularly and not integrated with an existing CGKA as before. In IAS, admins generate a signature key pair and sign every time they carry out a commit or a proposal, and verify a small amount t of signatures (typically $t \leq |G^*|$) in admin proposals before a commit. If we denote the cost (time/length) of a message signature or verification by s , and the cost (time/length) of a signature key pair generation by k , we obtain the values¹⁴ in Table 5.1. Note that $s = \mathcal{O}(\lambda)$ and $k = \mathcal{O}(\lambda)$ (i.e., are constant) for security parameter λ .

The overhead of DGS depends heavily on the cost of the admin CGKA CGKA^* (an optimistic estimation can be $\mathcal{O}(\log m)$ [ACDT20, KPPW⁺21, ACJM20] but can be $\mathcal{O}(m)$ in the worst case). CGKA operations only affect administrators. Note that a DGS admin-only commit is not sent to standard members (only the signed new admin key has to). Hence, DGS is very efficient for standard users.

¹⁴We ignore the size of user identifiers IDs as we assume they are small and are each a constant or even independent function of the security parameter. In any case, the overhead from identifiers should be similar in existing application-level protocols.

Chapter 5. Cryptographic Administration for Secure Group Messaging

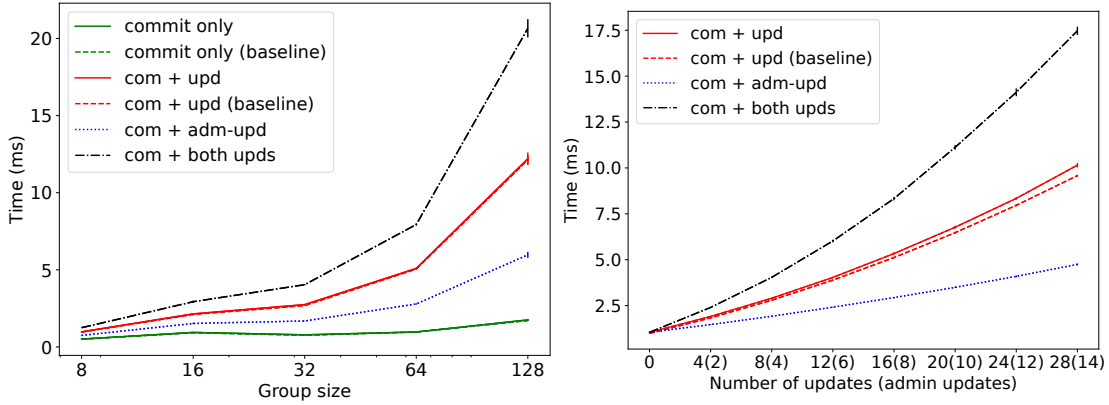


Figure 5.14: Benchmark of the Proc algorithm when processing a commit message. The different scenarios are those of Figure 5.13.

	Length (Adm)	Length (All)	Time (Adm)	Time (All)
IAS	$ts + tk$	$ts + tk$	$\mathcal{O}(ts + k)$	$\mathcal{O}(ts)$
DGS	$C + s + k$	$s + k$	$\mathcal{O}(C + s + k)$	$\mathcal{O}(s + k)$

Table 5.1: Additional cost of IAS and DGS with respect to a plain CGKA (per group) where t is the number of admin proposals and C (for DGS only) refers to the cost of running CGKA*.

The ratio of additional messages sent, which is application-specific, is hard to estimate. Admin-only commits and admin modifications are expected to be less frequent than standard operations. The number of update proposals (although individually very cheap) is expected to be at most linear in $|G|$.

Forward-secure signatures (for optimally-secure IAS) can be instantiated with essentially constant amortised overhead in space and time relative to a regular signature scheme, while supporting unbounded secret key updates [MMM02].

We conclude that IAS presents a generally affordable overhead for all users, while DGS introduces basically no cost for standard users and is more costly for administrators if $|G^*|$ is relatively large.

Admins in TreeKEM Variants. The Tainted TreeKEM protocol provides efficiency advantages if only a subset of users carry out tree-changing operations (adds and removals) [KPPW⁺21]. Tainted TreeKEM, however, is not formalised in the propose-and-commit-paradigm, which complicates the efficiency comparison; such an analysis thus remains open. When standard users are allowed to commit updates, the tree blanking issue with MLS TreeKEM is not worsened by administrators, hence efficiency should not decrease either.

5.4.2 Modelling in Related Work

CGKAs and MLS. The CGKA abstraction has deviated from MLS and has become an object of study of its own [AHKM22, BDG⁺22, AAN⁺22b], but in general still inherits important limitations from MLS. Among them, CGKAs rely on the availability of a well-behaved PKI, generally require total ordering on control messages (recent work relaxes this to causal ordering [AMT23]), and fail to capture messaging solutions that deviate from group key agreement such as Signal and WhatsApp.

In MLS, there exists a strong architectural separation between the Delivery Service (DS, usually a central server) and a so-called Authentication Service (AS) whose design is left to the infrastructure designers [BBR⁺23]. Following this separation, the Delivery Service is often modelled adversarially in CGKAs whereas the AS is abstracted as a PKI [ACDT20, KPPW⁺21, ACJM20] as in this chapter. A compromised AS allows for the corruption of user credentials, resulting in trivial user impersonation.

PKI. Both IAS and DGS rely on a PKI that we assume is incorruptible. In IAS, parties use the PKI to verify the identity of administrators and self-removing users. In DGS, incoming users use it to retrieve the current group-wide admin signature public key. As the PKI is only used to establish an initial root of trust among parties, i.e., forward and post-compromise security are ensured for existing group members without additional PKI calls, our modelling is consistent with the separation between delivery and authentication discussed above. Note that group administration aims to remove the trust in the DS (the server) but is still vulnerable to a corrupt AS. Previous CGKA work follows similar PKI abstractions [ACDT20, KPPW⁺21], or ignores the AS [CHK21]. That is, in all group messaging works we are aware of, the PKI *always* behaves consistently and correctly for all users. Two partial exceptions are Alwen et al. [AJM22] and Alwen et al. [ACDT21a], where malicious keys can be registered, although security (inevitably) degrades strongly for such users. By abstracting away the AS, our schemes are compatible with diverse authentication solutions such as out-of-band verification.

Signal Private Groups. In Chase et al. [CPZ20] and as deployed in Signal, a central server manages the membership of a group whilst hiding the set of group members from non-members (modulo metadata leaked to a network adversary). The main goals of this solution are to achieve user privacy and act as a single source of truth for the membership of a group. We believe that this approach could be extended to support secure administration; an advantage is that users no longer have to track group membership individually as in (A)-CGKA, which prevents consistency issues when users do not apply the same sequence of group updates locally. We note that Signal Private Groups however does not fully protect from server and network attacks as our A-CGKA constructions do: for example, it is possible for the server to re-add removed users. In addition, the system has not been analysed in composition with an underlying group messaging protocol (pairwise Signal) where concurrency issues can arise.

6 WhatsApp with Sender Keys? Analysis, Improvements and Security Proofs

In this chapter, we formally model and analyse the group messaging protocol Sender Keys used by WhatsApp and Signal, and propose some extensions and improvements. An extended abstract corresponding to this work appeared at ASIACRYPT 2023, and was joint work with David Balbás and Phillip Gajland [BCG23b]. A preliminary version of this work appeared at RECSI 2022 [BCG22]. A full version of this work can be found on the Cryptology ePrint Archive [BCG23a].

6.1 Contribution

In the context of group messaging, Signal [M⁺16]¹ and later WhatsApp [Wha20] have adopted the so-called *Sender Keys* protocol [Mar14], which has enjoyed widespread adoption for numerous years. Besides these, other popular solutions such as Matrix [ADJ24] and Session [Jef20] implement variants of this protocol. In Sender Keys, messages are encrypted using a user-specific symmetric key (which is then hashed forward) and then authenticated with a signature. Additionally, parties rely on secure two-party channels (instantiated in practice with the Double Ratchet) to share key material between them. Looking ahead, two-party channels will be central to determine the security attained by any instantiation of Sender Keys.

As mentioned in Chapter 5, a baseline for secure group messaging has been recently established by the IETF Messaging Layer Security (MLS) [BBR⁺23] standard, a joint effort between academia and industry². The protocol provides sub-linear complexity for group operations (adding/removing members and updating key material). Academic works have also explored so-called continuous group key agreement (CGKA) [BBR18, ACDT20, KPPW⁺21, ACJM20, ACDT21a], although these are only a component of a fully-fledged group messaging protocol. So far, in terms of complete messaging protocols, only the modular construction

¹Contrary to the folklore understanding that the Signal Messenger uses the pairwise channels approach for group messaging in small groups, Signal currently uses Sender Keys whenever possible. We refer to Section 6.5.4 for details.

²Recent academic works and ongoing discussions in mailing lists have identified and addressed several security issues that emerged during the standardisation of MLS [ACDT20, AJM22, IETF23].

Chapter 6. WhatsUpp with Sender Keys? Analysis, Improvements and Security Proofs

from [ACDT21a] building on CGKA (which includes MLS), DCGKA [WKHB21] in the decentralised setting and very recently Matrix [ADJ24] have been formalised to date.

Despite being the most complete and well-studied protocol to-date in the literature, MLS and CGKAs in general still have some drawbacks. While some exhibit sub-linear performance in specific executions (and this class of executions is not well-characterised in the literature), their performance can degrade to linear in general, which is unavoidable at least when using off-the-shelf cryptographic primitives [BDG⁺22]. Moreover, they tend to be complex, increasing their attack surface and making them more susceptible to design and implementation bugs. Finally, given the standardisation of MLS only occurred recently, MLS is yet to be widely deployed.

Hence, Sender Keys and similar approaches to group messaging remain an essential and practical alternative with different security / performance trade-offs. Firstly, Sender Keys stands out for its relative simplicity, which reduces its potential attack surface, making the protocol less susceptible to vulnerabilities in both its design and implementation. Secondly, Sender Keys offers good performance in small to moderate-sized groups, as demonstrated by its successful adoption for groups of up to 1024 parties in WhatsApp and Signal [Wha20, M⁺16]. While the main group operations (adding and removing users) respectively have $\mathcal{O}(n)$ and $\mathcal{O}(n^2)$ total communication complexity for groups of size n , concrete efficiency suffices in practice. Thirdly, Sender Keys offers forward-secure confidentiality and robust support for concurrent and out-of-order application message exchange.

Surprisingly, despite having the widest adoption and an open source implementation of its core cryptographic operations [M⁺16], Sender Keys has not been formally studied in the literature, prompting the following natural question:

Can we formalise the Sender Keys protocol in a meaningful security model?

To answer this question we start by introducing a new cryptographic primitive, along with a security model, to capture a broad class of group messaging protocols that do not necessarily employ CGKA [ACDT20] at their core. Our framework provides native support for group messaging protocols that utilise secure two-party communication channels under the hood, for which we introduce a clean level of abstraction. This novel framework proves instrumental in our analysis, as existing literature predominately focuses on CGKA-oriented models that do not suit Sender Keys and similar protocols.

Subsequently, we present a detailed description of the core Sender Keys protocol within our framework and provide a security proof validating the soundness of the protocol. In our analysis, we observe that Sender Keys presents several deficiencies that, despite not being easily exploitable flaws, prevent several desirable and fundamental security notions from being met. These include forward security under message injections, resilience against injections impacting group membership changes³, and fast recovery from state compromise. These

³Note that Signal uses a dedicated private group management solution in practice [CPZ20] that we do not

findings call into question the widespread use of the term “secure messaging” by commercial messaging solutions, motivating the need for more detailed discussion about the nuances around these protocols.

In this regard we propose an improved version of Sender Keys, that we call Sender Keys+, where we only employ readily available cryptographic primitives that have minimal impact on efficiency⁴. This addresses the following pertinent question:

How can we improve the security of Sender Keys whilst preserving its practical efficiency?

Overall, we believe that the formalisation and establishment of a provably secure variant of Sender Keys, such as the Sender Keys+ protocol proposed in this chapter, can serve as a valuable foundation for future implementations of the protocol.

6.1.1 Summary

In summary, the main scientific contributions of this chapter are the following:

- We introduce a new cryptographic primitive that we call *Group Messenger* (GM). We establish a modular security model for GM designed to capture messaging protocols like Sender Keys that are not necessarily based on group key agreement. It accounts for an active adversary capable of controlling the network and adaptively learning the secret states of different parties.
- We develop a general framework for composing two-party channels with group messaging protocols that use them. Our approach parameterises the security of the Group Messenger primitive based on the underlying two-party channels, presenting a novel perspective that, to the best of our knowledge, has not been explored previously.
- We formally describe Sender Keys, based on an analysis of Signal’s source code [M⁺16], WhatsApp’s security white paper [Wha20], and the yowsup library [Gal21].
- We prove the security of Sender Keys in our model and describe several shortcomings. These force us to restrict the capabilities of the adversary substantially for the proof to be carried out.
- We propose security fixes and improvements, several of which result in the improved protocol Sender Keys+. In particular, we secure group membership changes, improve the forward security of the protocol, and introduce an efficient key update mechanism. We also formalise the additional security guarantees in our model.

capture and is less affected by this attack vector than WhatsApp [RMS18]; we refer to Section 6.5.4 for further details.

⁴Our approach veers away from a theoretically systematic exploration to determine the “optimal” security for a Sender Keys-like protocol, as this would require non-standard primitives that considerably degrade performance [BRV20, ACJM20].

6.1.2 Technical Overview

Security in Group Messaging. Besides standard notions such as confidentiality, authenticity, and integrity of sent messages, we again consider forward security (FS) and post-compromise security in this chapter. Additionally, protocols must secure group membership updates, namely removed members must not be able to read messages sent after their removal, and newly added members must not (by default) be able to read past messages.

Most of the different formalisations of security in the literature model an adversarial Delivery Service (DS), the entity responsible for delivering messages between parties over the network. The adversary (modelling the DS) can act as an eavesdropper with extended capabilities, e.g., that can schedule messages to be consistently delivered by users, as in Alwen et al. [ACDT20], as a semi-active adversary that can schedule messages arbitrarily [KPPW⁺21], or as an active adversary that can inject messages to different degrees [ACJM20, BCV23]. In many protocols, including Sender Keys and MLS, the DS relies mainly on some centralised infrastructure (the *central server* hereafter).

Sender Keys. In a Sender Keys group G , every user $ID \in G$ owns a so-called *sender key* which is shared with all group members. A sender key is a tuple $SK = (\text{spk}, \text{ck})$, where spk is a public signature key (with a private counterpart ssk), and ck is a symmetric *chain key*. Every time ID sends a message m to the group, ID encrypts m using a *message key* mk that is deterministically derived (via a key derivation function H_1) from its chain key ck and erased immediately after being used. Upon message reception, group members derive mk to decrypt the corresponding ciphertext, which can also be delivered out-of-order as we discuss in later sections. Messages are authenticated by appending the sender's signature to them. In Figure 6.1, we provide a high-level depiction of what takes place in a three-member group $G = \{A, B, C\}$ when A sends a message that parties B and C receive.

Informally, forward security is provided by using a fresh message key for every message: every time a message is sent, the chain key is symmetrically ratcheted, i.e., hashed forward using a key derivation function H_2 . The protocol, that we describe further in Section 6.4, also requires that there exist confidential and authenticated communication channels between each pair of group members. These are used for sharing sender keys when parties are added or removed from the group, or when some party updates their key material. For example, in the event that some ID leaves the group, members erase their own sender key and start over. This mechanism provides a form of PCS when a user is removed as the key material is refreshed.

Modelling Two-Party Channels. Formally capturing the security of two-party channels is central to our analysis of Sender Keys since fresh sender keys are sent over these channels. Two-party channels that are not regularly used can undermine security. For example, if a group member ID 's state is compromised, there is no guarantee that fresh keys sent by other members (via two-party channels) are not leaked, since ID 's two-party channels may not yet have healed yet. Moreover, two-party channels can take more than one round trip to heal when using the Double Ratchet, as is the case for WhatsApp and the Signal Messenger [ACD19].

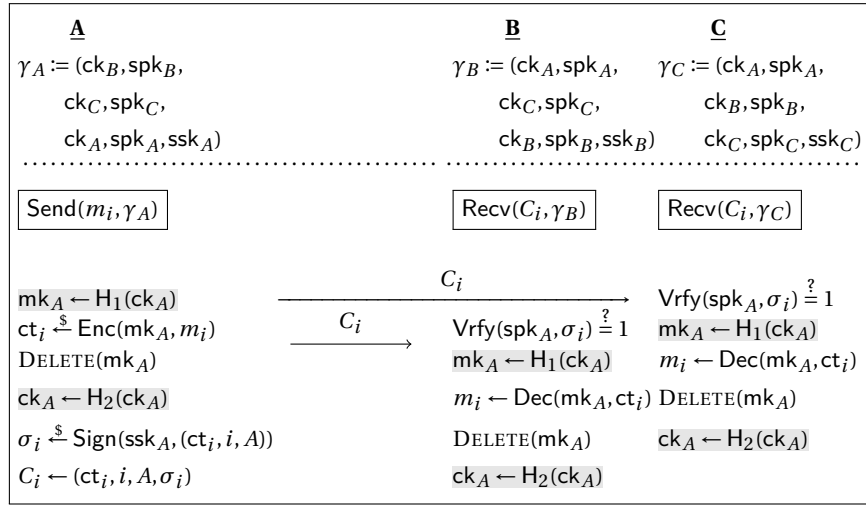


Figure 6.1: Simplified diagram for sending/receiving messages between three group members. For $ID \in \{A, B, C\}$, ID 's initial sender key is $(\text{ck}_{ID}, \text{spk}_{ID})$. The state γ_{ID} of ID contains the sender keys of all group members.

Our modelling starts in Section 6.2 with the introduction of a primitive 2PC for two-party channels. We define a two-party channel with initialisation (Init), channel initialisation (InitCh), send (Send) and receive (Recv) algorithms. Notably, InitCh allows parties to adaptively bootstrap channels, and deviates from works on ratcheting-based two-party messaging that abstract authentication away [BSJ⁺17]. Our security model captures both forward security and post-compromise security. To model PCS, we introduce a crucial parameter, denoted as Δ and referred to as the *PCS bound*. This parameter, inspired by Alwen et al. [ACD19] and Blazy et al. [BBL⁺22], serves as an upper bound on the number of synchronous communication steps or *channel epochs* required to restore security following a compromise.

Our Primitive: Group Messenger. In Section 6.3, we define a new cryptographic primitive, Group Messenger (GM), which includes five stateful algorithms that: initialise a party's state (Init), send an application message (Send), receive an application message (Recv), execute a change proposal in the group (Exec), and process a change in the group (Proc). Supported group changes are: group creation, member addition, member removal, and sender key updates. Note this contrasts with the three-phase propose/commit/process flow for updates (the so-called propose-commit paradigm [ACJM20]) used by MLS and newer CGKA protocols.

We define a game-based security notion for GM that captures a partially active adversary with control over the Delivery Service, taking inspiration from previous CGKA modelling [ACJM20, BCV23]. In our model in Section 6.3.1, we capture the security of each protocol by parameterising the game with a *cleanness predicate* (sometimes safety predicate in other work), which excludes trivial attacks, including both those which are unavoidable and those indicative of security weaknesses.

Chapter 6. WhatsApp with Sender Keys? Analysis, Improvements and Security Proofs

Security Analysis of Sender Keys. With this formalism established, in Section 6.4.2 we define cleanness predicates for Sender Keys that precisely capture its security. We define three sub-predicates that restrict the capabilities of the adversary for message *challenges*, capturing confidentiality; for message *injections*, capturing integrity and authenticity; and for re-orderings and forgeries of control messages (*concurrency*), capturing the message ordering provided by the central server.

Notably, the restrictions that we impose via our cleanness predicates are necessary for the security proof to go through and reveal several shortcomings in the protocol. Examples include:

- Sender Keys achieves only a weak form of PCS through key updates. Healing from a compromise requires multiple messages (at least $\Delta + 1$), even if a user is removed during healing.
- Control messages lack proper authentication and are malleable. An adversary with partial control over the network, such as the server, can make arbitrary changes in the group membership (such as adding new users without any member's authorisation), which is a significant practical concern.
- Forward security is sub-optimal, as messages are malleable after they are sent if a state exposure occurs.

Towards proving Sender Keys secure, we provide a somewhat detailed description of Sender Keys for pedagogical reasons in Section 6.4.1 and a complete formalisation in Section 6.6.1; for reference we also provide two tables in Appendix A.4 explaining variables used in our security games and formal specification. These may be of independent interest, as in particular we provide the first complete formal specification of Sender Keys. In Section 6.6.2, we formally prove that Sender Keys is secure under standard assumptions in our (necessarily restricted) model. Since we treat the primitives (symmetric encryption, signatures, PRG, two party channels) as black boxes, one can build, e.g., post-quantum Sender Keys following our specification.

In Section 6.5.4, we compare our description of Sender Keys with the implementations in WhatsApp and Signal, clarifying the extent to which our findings are applicable to these popular apps. We remark that our core analysis is nevertheless implementation-agnostic, and the fact that we model the underlying two-party channels in a fine-grained fashion allows us to capture their impact on security of Sender Keys in the face of state exposure (i.e., FS and PCS).

Shortcomings and Proposed Improvements. Leaving aside the security limitations that are intrinsic to the design of the protocol, we find that one can improve both its security and efficiency in several aspects. Hence, in Section 6.5 we propose modifications to the protocol with the aim of securing group membership, strengthening the (weaker than expected) forward security for authentication, and integrating efficient post-compromise security updates. Notably,

our novel PCS update mechanism improves the key update mechanism implemented by our core protocol and performed in Signal, bringing down the total communication complexity from quadratic to linear in the group size. Moreover, as a result of our modular approach with respect to modelling two-party channels, our modelling can capture the security improvement (or weakening) that results from replacing the Double Ratchet by an alternative two-party messaging protocol. Note that our formal of Sender Keys (Section 6.6.1) also contains our modifications for Sender Keys+ (they are nonetheless visually separated from the core protocol).

We extend our modelling to establish the security of our modified protocol, called Sender Keys+. The main technical step involves redefining the cleanness predicates (Section 6.5), which are strictly less restrictive compared to those used for the original protocol. Notably, the adversary is now allowed to inject control messages (given the group has recovered from any state exposures). We also allow the adversary to mount more fine-grained attacks for application message forgeries, and allow *arbitrary* challenges after some party has updated over a refreshed channel (before, we could only re-allow challenges on the *updater*). In Section 6.6.3 we formally prove security additionally assuming a secure MAC and dual PRF.

6.1.3 Additional Related Work

The formal extension of CGKA to group messaging was explored by Alwen et al. [ACDT21a], while the key schedule of MLS was proven secure by Brzuska et al. [BCK22]. We provide more thorough comparison between CGKA-based protocols and Sender Keys/Sender Keys+ in Section 6.5.3.

The work of Cong et al. [CEST22] shares some similarities with ours as it also constructs group messaging from two-party channels and achieves $\mathcal{O}(n)$ key update complexity. However, they do not model two-party channels as a standalone primitive nor dynamic groups formally, and their protocols require more interaction than ours (e.g., the initial group key agreement protocol can take several rounds in general).

Concurrency, a crucial aspect in CGKA-based protocols, has been a central topic in works such as [AAN⁺22b, AAN⁺22a, BDR20]. Secure administration in CGKAs was explored in the previous chapter (Chapter 5) of this thesis. Weidner et al. [WKHB21] adopt a Sender Keys-like approach to construct a decentralised CGKA protocol but they do not capture group messaging, and their security model does not support message injections (hence considering a passive adversary). Moreover, the theorems in their work assume a non-adaptive adversary where the adversary must announce all queries at the game's outset. Their work nonetheless extends the scope of modern messaging to decentralised networks without a central authority, diverging from existing approaches that target centralised networks. A simplified (notably lacking forward security) decentralised variant of Sender Keys is implemented by the Session app [Jef20].

Chapter 6. WhatsUpp with Sender Keys? Analysis, Improvements and Security Proofs

Also relevant to our work are secure two-party messaging protocols that propose alternatives to the Double Ratchet [MP16a], such as [JS18, PR18, DV19, ACD19, BRV20, PP22]. Inspired by more practical-oriented endeavors, we acknowledge recent cryptographic audits conducted on Telegram [AMPS22], Matrix [ACDJ23], Threema [KGP23], and WhatsApp’s backup service [DFG⁺23].

Sender Keys. While some works on Sender Keys lack formalism and security proofs, they offer valuable insights. Rösler et al. [RMS18] evaluate Sender Keys, provide a high-level description of the protocol, and examine practical vulnerabilities in WhatsApp group chats. Multi-group security and key update mechanisms for Sender Keys are informally discussed by Cremers et al [CHK21]. In [BCG22], a preliminary analysis of the security of Sender Keys is carried out. While the paper only includes informal discussions and no proofs, it serves as an initial exploration for the ideas in the present work. We remark that the scope of [BCG22] is limited, as it does not formally develop a security model, and assumes that all two-party channels used by Sender Keys are *perfectly secure*, which is unrealistic and impossible to develop in practice.

Concurrent work by Albrecht et al. [ADJ24] develops a device-oriented security model and a proof for a recent specification of Matrix (i.e., for the updated protocol that mitigates the issues described in [ACDJ23]). For group messaging, Matrix implements the Megolm protocol, which is Sender Keys-inspired but still deviates significantly from our description in this chapter, particularly regarding server interaction. Remarkably, [ADJ24] and our work arrive to similar conclusions in our analysis, such as the insecurity of group management and the challenges imposed by message ordering. Our works are complimentary and open new research directions. Examples include exploring whether the improvements behind Sender Keys+ can also be applied to Megolm, as well as extending our modelling to consider the (challenging) multi-device setting as they do.

6.2 Two-Party Channels

In this section, we present our approach towards capturing two-party channels as a standalone primitive for modelling Sender Keys.

6.2.1 Primitive Definition and Correctness

We start by defining two-party channels.

Definition 44 (Two-Party Channel). A two-party channel scheme is a tuple of efficient algorithms $2PC := (\text{Init}, \text{InitCh}, \text{Send}, \text{Recv})$ such that:

- $\gamma \xleftarrow{\$} \text{Init}(1^\lambda, ID)$: The probabilistic initialisation algorithm takes a user identity ID as input and outputs an initial state γ .
- $\text{acc} \xleftarrow{\$} \text{InitCh}(\gamma, ID^*)$: The probabilistic channel initialisation algorithm takes a state γ

and a user identity ID^* as inputs and outputs an acceptance bit $\text{acc} \in \{\text{true}, \text{false}\}$ and updates the caller's state.

- $(C, e_{2\text{pc}}, i_{2\text{pc}}) \stackrel{\$}{\leftarrow} \text{Send}(\gamma, m, ID^*)$: The probabilistic sending algorithm takes as inputs a message m , the intended message recipient ID^* and a state γ , and outputs a ciphertext C , a channel epoch-index pair $(e_{2\text{pc}}, i_{2\text{pc}})$ corresponding to m (or \perp upon failure), and updates the state.
- $(m, ID^*, e_{2\text{pc}}, i_{2\text{pc}}) \leftarrow \text{Recv}(\gamma, C)$: The deterministic receiving algorithm takes as inputs a ciphertext C and a state γ , and outputs a message m , a user identity ID^* corresponding to the sender of m and a channel epoch-index pair $(e_{2\text{pc}}, i_{2\text{pc}})$ corresponding to m (or \perp upon failure), and updates the state.

Our $2\text{PC} := (\text{Init}, \text{InitCh}, \text{Send}, \text{Recv})$ primitive captures two initialisation functions. The first function initialises the state of a party by taking its ID as input, while the second function is used to initialise a communication channel with a counterpart ID^* . Consider two parties, ID and ID^* who intend to communicate over a two-party channel. Both parties initialise their states, γ_{ID} and γ_{ID^*} using the Init function. Subsequently, ID and ID^* initiates the communication channel by invoking $\text{InitCh}(\gamma_{ID}, ID^*)$ and $\text{InitCh}(\gamma_{ID^*}, ID)$ (note we assume both parties must call InitCh).⁵ It is worth noting that, similar to DCGKA [WKHB21], our two-party channel primitive implicitly depends on some universally available and consistent public-key infrastructure.

Channel Epochs and Indices. Our notion of channel epochs is exactly the notion of epochs as defined in [ACD19] used to model the Double Ratchet protocol [MP16a] and is an example of an *ordinal* as in Section 4.2. For a given two-party channel, ID and ID' are each associated with a channel epoch $e_{2\text{pc}}$, which corresponds to how many times the direction of communication has changed, alongside an index $i_{2\text{pc}}$, indicating how many Send calls have been made by a sender or the latest message received by a receiver. Initially, the sender (say ID) sets $e_{2\text{pc}} = 0$ and the receiver ID' sets $e_{2\text{pc}} = -1$. Thereafter, party ID (resp. ID') is the sender in even (resp. odd) channel epochs, and the receiver in odd (resp. even) channel epochs. When a party is a sender in channel epoch $e_{2\text{pc}}$ and receives a message from $e_{2\text{pc}} + 1$, they advance to $e_{2\text{pc}} + 1$; likewise they advance channel epochs when they are a receiver and then send a message. When a party sends as a sender, they increment their index $i_{2\text{pc}}$, and as a receiver they set $i_{2\text{pc}}$ to the message received with the highest index. Note that $i_{2\text{pc}}$ represents the number of messages sent or received for a given channel epoch $e_{2\text{pc}}$; $i_{2\text{pc}}$ is set to 0 whenever a party has just become a sender again. Hence, we define a total ordering on channel epochs and indices $(e_{2\text{pc}}, i_{2\text{pc}})$ such that $(e_{2\text{pc}}, i_{2\text{pc}}) \leq (e'_{2\text{pc}}, i'_{2\text{pc}})$ when $e_{2\text{pc}} < e'_{2\text{pc}}$, or $e_{2\text{pc}} = e'_{2\text{pc}}$ and $i_{2\text{pc}} \leq i'_{2\text{pc}}$.

Oracles for Correctness. To capture correctness we employ several oracles that the adversary can query. We describe them briefly before describing the correctness notion.

⁵In practice, a party given pair of parties may initialise and tear down many channels between each other over time; by associating a given 'party' with several identifiers, our modelling here is without loss of generality.

Chapter 6. WhatsUpp with Sender Keys? Analysis, Improvements and Security Proofs

INIT-CH(ID, ID'): This oracle initialises the two-party channel between parties ID and ID' .

SEND(ID, ID', m): This oracle allows the sending of a message m (i.e., generating a ciphertext) from party ID to party ID' .

RECEIVE($ID, ID', e_{2pc}, i_{2pc}$): This oracle delivers the ciphertext output by ID' for channel epoch (e_{2pc}, i_{2pc}) in a previous SEND call to ID .

EXPOSE(ID): This oracle leaks the state of party ID to the adversary.

Game 2PC-CORR $_{2PC}^{\mathcal{O}}$	Oracle SEND(ID, ID', m)
1: for all ID :	1: require $(ID, ID') \in CH$
2: $\gamma[ID] \stackrel{\$}{\leftarrow} \text{Init}(1^\lambda, ID)$	2: $(C, e_{2pc}, i_{2pc}) \stackrel{\$}{\leftarrow} \text{Send}(m, ID', \gamma[ID])$
3: $M[\cdot], CT[\cdot] \leftarrow \perp$	3: if $C = \perp \vee e_{2pc} = \perp \vee i_{2pc} = \perp$
4: $CH \leftarrow \emptyset$	4: $\text{win} \leftarrow 1$
5: $\text{win} \leftarrow 0$	5: $M[ID, ID', e_{2pc}, i_{2pc}] \leftarrow m$
6: $\mathcal{A}^{\mathcal{O}}$	6: $CT[ID, ID', e_{2pc}, i_{2pc}] \leftarrow C$
7: return win	7: return $(\text{win}, C, e_{2pc}, i_{2pc})$
Oracle INIT-CH(ID, ID')	Oracle RECEIVE($ID, ID', e_{2pc}, i_{2pc}$)
1: require $(ID, ID') \notin CH$	1: require $(ID', ID) \in CH$
2: $\text{acc} \stackrel{\$}{\leftarrow} \text{InitCh}(ID', \gamma[ID])$	2: require $CT[ID, ID', e_{2pc}, i_{2pc}] \neq \perp$
3: if $\neg \text{acc}$	3: $C \leftarrow CT[ID, ID', e_{2pc}, i_{2pc}]$
4: $\text{win} \leftarrow 1$	4: $(m, ID', e'_{2pc}, i'_{2pc}) \leftarrow \text{Recv}(C, \gamma[ID])$
5: $CH \stackrel{\cup}{\leftarrow} (ID, ID')$	5: $m' \leftarrow M[ID, ID', e_{2pc}, i_{2pc}]$
6: return acc	6: if $(m', e'_{2pc}, i'_{2pc}) \neq (m, e_{2pc}, i_{2pc})$
Oracle EXPOSE(ID)	7: $\text{win} \leftarrow 1$
1: return $\gamma[ID]$	8: return win

Figure 6.2: 2PC-CORR correctness for 2PC where $\mathcal{O} = \{\text{INIT-CH}, \text{EXPOSE}, \text{SEND}, \text{RECEIVE}\}$. Lines in teal correspond only to bookkeeping and state update operations.

Correctness. We provide a correctness game played between a challenger and a computationally unbounded adversary (Figure 6.2). here, The adversary can invoke the INIT-CH oracle to establish a secure channel between ID and ID' complete; we require that $\text{InitCh}(ID, ID')$ is always successful when first called with input (ID, ID') . Subsequently, the adversary can dynamically query the SEND and RECEIVE oracles to execute the protocol. Note that the EXPOSE oracle is unrestricted, providing the adversary with leaked state information. It is worth noting that RECEIVE only allows the adversary to make the challenger call Receive on inputs that were output by the corresponding Send call in SEND.

We formally define correctness below.

Definition 45 (Correctness of 2PC). We consider the $2\text{PC-CORR}_{2\text{PC}}$ game defined in Figure 6.2. A two-party channels scheme 2PC is correct or 2PC-CORR if for any (possibly unbounded) adversary \mathcal{A} we have

$$\text{Adv}_{2\text{PC}}^{2\text{pc-corr}}(\mathcal{A}) := \Pr[2\text{PC-CORR}_{2\text{PC}}^{\mathcal{A}} \Rightarrow 1] = 0.$$

Security. The Double Ratchet protocol has been the subject of several academic works [ACD19, CCD⁺20, CJSV22, BFG⁺22a] that analyse its security on a fine-grained level for two-party communication. When used by multiple parties in a group during the execution of Sender Keys, analysing the Double Ratchet protocol becomes complex, making it difficult to replace it with other protocols. To tame this complexity, we adopt a comparatively simpler notion of two-party communication in similar complexity to the formalism of Weidner et al. for their DCGKA protocol [WKHB21]. Our modelling captures forward security, is parameterised by the post-compromise security of the underlying channels and supports out-of-order message delivery. We define security after describing the security game and introducing some predicates below.

Oracles for Security. Compared to correctness, the adversary \mathcal{A} will be given access to the following oracles in the security game:

$\text{INIT-CH}(ID, ID')$, $\text{SEND}(ID, ID', m)$, $\text{EXPOSE}(ID)$: As in correctness.

$\text{CHAL}(ID, ID', m_0, m_1)$: This oracle generates a message challenge where the adversary \mathcal{A} provides two messages m_0 and m_1 of the same length, and party ID sends m_b to party ID' .

$\text{RECEIVE}(ID, C)$: This oracle delivers the ciphertext C to party ID .

Game Description. The notion of security defined in the $2\text{PC-IND}_{2\text{PC}, b, C_{2\text{pc}}, \Delta}$ game in Figure 6.3 is parameterised by a cleanness predicate $C_{2\text{pc}}$ and a PCS bound $\Delta > 0$. Broadly, the cleanness predicate prevents the adversary from winning the game by making a trivial attack, i.e., via a bit guess (resp. forgery) based on a challenge (resp. delivery) using exposed key material. PCS after an exposure is parameterised by Δ , which is the number of message steps (i.e. number of times that the sender-receiver roles alternate) that the channel requires for healing.

The game starts by initialising the states of all parties and the dictionaries CH and M, which store challenge ciphertexts and all sent (including challenge) ciphertexts, respectively. It also initialises El as a variable that tracks the channel epoch-index pair of a given channel (ID, ID') . Then, the adversary \mathcal{A} can adaptively query all the oracles listed above. Finally, \mathcal{A} outputs a guess b' of b that the challenger also outputs given that the execution has been *clean*, i.e., that the cleanness predicate $C_{2\text{pc}}$ holds.

\mathcal{A} can win the game in two different ways. Firstly, it can make a correct guess of the bit b ; note

Chapter 6. WhatsUpp with Sender Keys? Analysis, Improvements and Security Proofs

<p>Game 2PC-IND_{2PC,b,C_{2pc},Δ}(\mathcal{A})</p> <ol style="list-style-type: none"> 1: for all ID: 2: $\gamma[ID] \stackrel{\\$}{\leftarrow} \text{Init}(1^\lambda, ID)$ 3: $\text{CH}[\cdot], \text{M}[\cdot], \text{EI}[\cdot] \leftarrow \perp$ 4: $b' \leftarrow \mathcal{A}^{\mathcal{O}}$ 5: require C_{2pc} 6: return b' <p>Oracle INIT-CH(ID, ID')</p> <ol style="list-style-type: none"> 1: $\text{acc} \stackrel{\\$}{\leftarrow} \text{InitCh}(\gamma[ID], ID')$ 2: return acc <p>Oracle SEND(ID, ID', m)</p> <ol style="list-style-type: none"> 1: $(C, e_{2pc}, i_{2pc}) \stackrel{\\$}{\leftarrow} \text{Send}(\gamma[ID], m, ID')$ 2: require $C \neq \perp$ 3: $\text{M}[ID, ID'] \stackrel{\cup}{\leftarrow} \{(C, e_{2pc}, i_{2pc})\}$ 4: $\text{EI}[ID, ID'] \leftarrow (e_{2pc}, i_{2pc})$ 5: return (C, e_{2pc}, i_{2pc}) 	<p>Oracle CHAL(ID, ID', m_0, m_1)</p> <ol style="list-style-type: none"> 1: require $m_0 = m_1$ 2: $(C^*, e_{2pc}, i_{2pc}) \stackrel{\\$}{\leftarrow} \text{Send}(\gamma[ID], m_b, ID')$ 3: require $C^* \neq \perp$ 4: $\text{CH}[ID, ID'] \stackrel{\cup}{\leftarrow} \{(C^*, e_{2pc}, i_{2pc})\}$ 5: $\text{M}[ID, ID'] \stackrel{\cup}{\leftarrow} \{(C^*, e_{2pc}, i_{2pc})\}$ 6: $\text{EI}[ID, ID'] \leftarrow (e_{2pc}, i_{2pc})$ 7: return (C^*, e, i) <p>Oracle RECEIVE(ID, C)</p> <ol style="list-style-type: none"> 1: $(m, ID', e_{2pc}, i_{2pc}) \leftarrow \text{Recv}(\gamma[ID], C)$ 2: require $m \neq \perp$ 3: if $(C, e_{2pc}, i_{2pc}) \notin \text{M}[ID', ID]$: 4: return $(b, ID', e_{2pc}, i_{2pc})$ // forgery 5: $\text{CH}[ID', ID] \stackrel{-}{\leftarrow} \{(C, e_{2pc}, i_{2pc})\}$ 6: if $(e_{2pc}, i_{2pc}) > \text{EI}[ID', ID]$: 7: $\text{EI}[ID', ID] \leftarrow (e_{2pc}, i_{2pc})$ 8: return $(\perp, ID', e_{2pc}, i_{2pc})$ <p>Oracle EXPOSE(ID)</p> <ol style="list-style-type: none"> 1: require $\text{CH}[ID', ID] = \emptyset \forall ID'$ 2: return $\gamma[ID]$
---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	-----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

Figure 6.3: Indistinguishability (2PC-IND) security for 2PC where $\mathcal{O} = \{\text{INIT-CH}, \text{SEND}, \text{CHAL}, \text{RECEIVE}, \text{EXPOSE}\}$. Lines in teal correspond only to bookkeeping and state update operations. Dictionaries CH, M and EI store challenged messages, sent messages, and channel epoch-index respectively.

(ignoring RECEIVE) the only operation that depends on b is the CHAL oracle. Ciphertexts C^* generated by $\text{CHAL}(ID, ID', m_0, m_1)$ are stored in CH, and removed from CH once they are delivered. We use CH to prevent trivial forgeries, as if there is any challenge C^* that has not yet been delivered to ID' (and such that by correctness ID' must be able to decrypt it), then \mathcal{A} cannot leak the state of ID' via $\text{EXPOSE}(ID')$. Secondly, \mathcal{A} can directly obtain b by making a successful forgery that is accepted via $\text{RECEIVE}(ID, C)$. To leak b to \mathcal{A} , RECEIVE checks that C does not correspond to a message sent by the sender ID' in epoch (e_{2pc}, i_{2pc}) , where (ID', e_{2pc}, i_{2pc}) are the outputs of the challenger's Recv call within the RECEIVE query.

We note that the challenge-and-send style of our game is analogous to the security game for GM (to be introduced). This game is also multi-user as it captures all channels at once, extending other single-user security models such as those in [ACD19, WKHB21]. Our modelling presents some similarities with the modelling of two-party channels in [WKHB21], but differs

from it in several aspects. It is also multi-challenge as CHAL can be called several times. Our `Init` algorithm does not require the public key of the counterpart as opposed to theirs (we capture this via `InitCh` for each channel), and correctness is captured as part of their security model. More importantly, their adversary is not allowed to attempt forgeries (the model only captures confidentiality of sent messages) or out of order delivery.

Predicates. The game is parametrised by the two-party channels cleanness predicate C_{2pc} , which we divide into two sub-predicates as $C_{2pc} := C_{2pc\text{-chal}} \wedge C_{2pc\text{-inj}}$. Both sub-predicates capture PCS and are additionally parametrised by the PCS bound Δ . We follow the blueprint of [ACD19] for the predicate definition.

$$\begin{array}{l}
 \boxed{C_{2pc\text{-chal}} :} \quad \forall(i, ID, ID', e_{2pc}, i_{2pc}) : q_i = \text{EXPOSE}(ID) \wedge \\
 (e_{2pc}, i_{2pc}) = \max\{\text{El}[ID, ID'; q_i], \text{El}[ID', ID; q_i]\}, \exists(e'_{2pc}, i'_{2pc}, j) : (i < j) \wedge \\
 [(q_j = \text{CHAL}(ID', ID, \cdot, \cdot) \wedge (e'_{2pc}, i'_{2pc}) = \text{El}[ID', ID; q_j] \wedge e'_{2pc} < e_{2pc} + \Delta) \vee \\
 (q_j = \text{CHAL}(ID, ID', \cdot, \cdot) \wedge (e'_{2pc}, i'_{2pc}) = \text{El}[ID, ID'; q_j] \wedge e'_{2pc} < e_{2pc} + \Delta)]
 \end{array}$$

Figure 6.4: Predicate $C_{2pc\text{-chal}}$ where \mathcal{A} makes oracle queries q_1, \dots, q_q .

Challenge (Figure 6.4). Suppose that \mathcal{A} exposes a party ID and later makes a CHAL query involving ID (either as a sender or as a receiver) and some other party ID' . Essentially, the predicate requires that the challenge message then must belong to a channel epoch e'_{2pc} that is Δ or more epochs past the channel epoch e_{2pc} , where e_{2pc} corresponds to the largest epoch value between ID and ID' at exposure time.

Injection (Figure 6.5). For an intuitive description, let C be a forged ciphertext that, when processed by `Recv` by some ID' , claims to be from sender ID and from epoch-index (e_{2pc}^*, i_{2pc}^*) . Then, the predicate requires that if \mathcal{A} exposes ID and later attempts to inject C to some ID' , both the epoch e'_{2pc} of the (ID, ID') channel at injection time and e_{2pc}^* are Δ or more epochs further from the channel epoch e_{2pc} at exposure time. Considering both epochs e_{2pc}, e_{2pc}^* prevents injections on out-of-order messages. Self-injections (i.e., where $ID = ID'$) are also restricted since the adversary can trivially mount such an attack on the Double Ratchet [ACD19].

$$\begin{array}{l}
 \boxed{C_{2pc\text{-inj}} :} \quad \forall(i, ID, ID', e_{2pc}, i_{2pc}) : q_i = \text{EXPOSE}(ID) \wedge \\
 (e_{2pc}, i_{2pc}) = \max\{\text{El}[ID, ID'; q_i], \text{El}[ID', ID; q_i]\}, \\
 \exists(e'_{2pc}, i'_{2pc}, e_{2pc}^*, i_{2pc}^*, j, C) : (i < j) \wedge (\min\{e_{2pc}^*, e'_{2pc}\} < e_{2pc} + \Delta) \wedge \\
 (e'_{2pc}, i'_{2pc}) = \min\{\text{El}[ID, ID'; q_j], \text{El}[ID', ID; q_j]\} \wedge \\
 [(q_j = \text{RECEIVE}(ID', C) = (\cdot, ID, e_{2pc}^*, i_{2pc}^*) \wedge (C, e_{2pc}^*, i_{2pc}^*) \notin M[ID, ID'; q_j]) \vee \\
 (q_j = \text{RECEIVE}(ID, C) = (\cdot, ID', e_{2pc}^*, i_{2pc}^*) \wedge (C, e_{2pc}^*, i_{2pc}^*) \notin M[ID', ID; q_j])]
 \end{array}$$

Figure 6.5: Predicate $C_{2pc\text{-inj}}$ where \mathcal{A} makes oracle queries q_1, \dots, q_q .

In order to model relevant injection predicates for our GM primitive such as $C_{sk\text{-inj}}^\Delta$ (Fig-

Chapter 6. WhatsUpp with Sender Keys? Analysis, Improvements and Security Proofs

ure 6.11), we abuse notation and refer directly to $C_{2\text{pc-inj}}$. In the context of GM, we parametrise the predicate by a ciphertext $C_{2\text{pc}}$, and simply replace the belonging to M by the equivalent condition in the GM security game $(\cdot, C_{2\text{pc}}) \notin M[ID, ID'; q_j]$.

We formally define security below.

Definition 46 (Security of 2PC). We consider the $2\text{PC-IND}_{2\text{PC}, C_{2\text{pc}}, \Delta}$ game defined in Figure 6.3. A two-party channels scheme 2PC satisfies indistinguishability or 2PC-IND security with respect to cleanness predicate $C_{2\text{pc}}$ and PCS bound $\Delta > 0$ if for any efficient adversary \mathcal{A} we have

$$\begin{aligned} \text{Adv}_{2\text{PC}, C_{2\text{pc}}, \Delta}^{2\text{pc-ind}}(\mathcal{A}) &:= \left| \Pr[2\text{PC-IND}_{2\text{PC}, 1, C_{2\text{pc}}, \Delta}(\mathcal{A}) \Rightarrow 1] - \Pr[2\text{PC-IND}_{2\text{PC}, 0, C_{2\text{pc}}, \Delta}(\mathcal{A}) \Rightarrow 1] \right| \\ &= \text{negl} \end{aligned}$$

Instantiations. By previous work [ACD19], the Double Ratchet can be seen to achieve a PCS bound of $\Delta = 3$. However, by replacing the Diffie Hellman key exchange component in the Double Ratchet (referred to as continuous key agreement [ACD19]) with a KEM, the PCS bound can be improved to $\Delta = 2$. This is optimal since if a user is exposed in channel epoch $e_{2\text{pc}}$ and acts as the sender, then they can decrypt a message from channel epoch $e_{2\text{pc}} + 1$ based on correctness requirements. While we do consider protocols with weak PCS and hence larger values of Δ , including $\Delta = \infty$ if new randomness is never injected in key derivation, protocols lacking forward security like TLS are considered insecure within our model.

For channel initialisation InitCh , an initial key exchange between the parties needs to be carried out. Typically this is done via the X3DH protocol [MP16b] and by relying on a PKI which we abstract away in this chapter but previously explored in Chapter 3.

Extensions. Our security model adopts the core requirements of the modelling in [ACD19]. Our cleanness predicates are designed to suit a large class of two-party messaging protocols parametrised by the PCS bound Δ . As a consequence, our analysis is not as fine-grained as possible, but we gain in readability and modularity.

The recent work of Blazy et al. [BBL⁺22] devises metrics (generalising our PCS bound in some sense) to classify different two-party messaging protocols based on their resilience to different adversarial behaviours and the resulting PCS guarantees. Future work could incorporate these factors into our modelling of two-party channels. Another direction would be to parametrise the PCS bound based on whether a party is exposed while acting as a sender or receiver on the channel (as opposed to considering a worst-case Δ). One could take an even more fine-grained approach in the style of [CCD⁺20] also. A bound of $\Delta = 1$ would be possible when only considering receiver exposure.

We also note that if one replaces the Double Ratchet with another protocol, either keeping [ACD19, CZ22, PP22] or dropping [JS18, PR18, DV19] support for out-of-order message delivery along the way, it is possible to consider less restrictive injection predicates. For in-

stance, by using signatures, one no longer needs to restrict self-injections. For simplicity, we also restrict injections corresponding to out-of-order delivery (all epochs $< e_{2pc} + \Delta$ in $C_{2pc-inj}$); notice that restricting winning injections only on message epochs corresponding to old ciphertexts in-transit at corruption time is sufficient as done in [ACD19].

6.3 Group Messenger

We introduce our main cryptographic primitive called *Group Messenger* that captures sending and receiving application messages between members of a dynamic group. We note that our primitive captures a single group; extending it to multiple groups is straightforward by using group identifiers (see e.g. [BCV23]).

Definition 47 (Group Messenger). A Group Messenger GM is a tuple of efficient algorithms $GM := (\text{Init}, \text{Send}, \text{Recv}, \text{Exec}, \text{Proc})$ such that:

- $\gamma \xleftarrow{\$} \text{Init}(ID)$: The probabilistic initialisation algorithm takes as input a user identity ID and outputs an initial state γ .
- $C \xleftarrow{\$} \text{Send}(\gamma, m)$: The probabilistic sending algorithm takes as inputs a message m and a state γ , and outputs a ciphertext C (or \perp upon failure) and updates the state.
- $(m, ID^*, e, i) \leftarrow \text{Recv}(\gamma, C)$: The deterministic receiving algorithm takes as inputs a ciphertext C and a state γ , and outputs a message m , an identity ID^* corresponding to the sender, a group epoch e and index i both corresponding to m (or \perp upon failure), and updates the state.
- $T \xleftarrow{\$} \text{Exec}(\gamma, \text{cmd}, IDs)$: The probabilistic execution algorithm takes as inputs a command $\text{cmd} \in \{\text{crt}, \text{add}, \text{rem}, \text{upd}\}$, a list of identities IDs and a state γ , and outputs a control message T (or \perp upon failure) and updates the state.
- $b \leftarrow \text{Proc}(\gamma, T)$: The deterministic processing algorithm takes as inputs a control message T and a state γ , and outputs an acceptance bit $b \in \{0, 1\}$ and updates the state.

Finally, the (possibly empty) set of group members is stored as $\gamma.G$

In our syntax, a distinction is made between *application* messages and *control* messages. Specifically, distinct algorithms are employed for the transmission and reception of application messages, as well as for the execution and processing of group modifications. These modifications, executed via *Exec*, are parameterised by a command *cmd* that encompasses operations including in this chapter user addition *add*, removal *rem*, group creation *crt*, or user key material update *upd*. Moreover, in scenarios where two-party messaging protocols are required for the group primitive (although not applicable to CGKAs such as TreeKEM [ACDT20]), two-party messages are formally assumed to be sent alongside or within ciphertexts or control

Chapter 6. WhatsUpp with Sender Keys? Analysis, Improvements and Security Proofs

messages. Consequently, they are abstracted away from our syntax. Looking ahead, we will model the fact that ciphertexts and control messages are sent alongside two-party channel ciphertexts (and can be received with a different ciphertext or control message) when we define security in Section 6.3.1.

Message Epochs. We define a message epoch as a pair of integers (e, i) , internal to the state γ of a party ID , that captures time and synchronisation between parties. Message epochs are central to our description of Sender Keys and security model. Each application message sent by ID corresponds to a single message epoch (e, i) , which is output by the Recv algorithm at the receiver's end. The *epoch* e advances whenever ID processes a new group change (i.e., a control message). The *index* i advances when ID sends a new message. If control messages are delivered to group members in lockstep (i.e. sequentially in the same order), parties who have the same epoch e will have the same view of the group membership. We define a total ordering $(e, i) \leq (e', i')$ when $e < e'$, or $e = e'$ and $i \leq i'$. Nevertheless, we remark that 2PC channel epochs are independent from GM message epochs.

Oracles for Correctness. We introduce a game-based notion for GM correctness (as well as for security later), where the adversary \mathcal{A} will have access to various oracles that we outline below.

CREATE(ID, IDs): creates a group by executing $\text{Exec}(\gamma, \text{crt}, IDs)$ with ID as the initiator, generating a control message T . Namely, \mathcal{A} obtains $C_b \leftarrow \text{Send}(\gamma[ID], m_b)$.

SEND(ID, m): ID sends an application-level message m using the Send algorithm, producing a ciphertext C .

RECEIVE($ID, ID' e, i$): ID receives the ciphertext sent by ID' corresponding to message epoch (e, i) by calling $\text{Recv}(\gamma[ID], C)$, where C is stored in $\text{CT}[ID', e, i]$ by the challenger (note we only consider honest delivery for correctness).

ADD(ID, ID') / REMOVE(ID, ID') / UPDATE(ID): ID adds ID' / removes ID' / refreshes ID 's secrets by calling $\text{Exec}(\gamma[ID], \text{add}, ID')$ / $\text{Exec}(\gamma[ID], \text{rem}, ID')$ / $\text{Exec}(\gamma[ID], \text{upd}, ID)$, generating control message T . The e -th control message (i.e., the epoch e control message) is stored as $T[i]$.

DELIVER(ID): ID is delivered the next control message in T , if it exists, via $\text{Proc}(\gamma[ID], T[E[ID]])$ (where ID is in epoch $E[ID]$).

EXPOSE(ID): Leaks the state γ of ID to \mathcal{A} .

Correctness. We capture correctness in a game played by an unbounded adversary in Figure 6.6. Our game captures several properties, given that all messages are generated honestly:

- *Message delivery:* Application messages (generated by SEND) must be received correctly by all group members (checked at lines 9-10 of RECEIVE).

Game M-CORR _{GM} (\mathcal{A})	Oracle ADD(ID, ID')	Oracle RECEIVE(ID, ID', e, i)
1: for all ID : 2: $\gamma[ID] \stackrel{\$}{\leftarrow} \text{Init}(1^\lambda, ID)$ 3: $T[\cdot], M[\cdot], CT[\cdot] \leftarrow \perp$ 4: $ep, E[\cdot], l[\cdot] \leftarrow 0$ 5: $G[\cdot] \leftarrow \emptyset$ 6: $win \leftarrow 0$ 7: $\mathcal{A}^\mathcal{O}$ 8: return win	1: require $E[ID] = ep$ 2: $T \stackrel{\$}{\leftarrow} \text{Exec}(\gamma[ID], \text{add}, \{ID'\})$ 3: require $T \neq \perp$ 4: $T[ep] \leftarrow T; ep \leftarrow ep + 1$ 5: $G[ep] \leftarrow G[ep - 1] \cup \{ID'\}$ 6: return T	1: require $ID \in G[e]$ 2: require $CT[ID', e, i] \neq \perp$ 3: $C \leftarrow CT[ID', e, i]$ 4: $(m', ID', e', i') \leftarrow \text{Recv}(\gamma[ID], C)$ 5: $m \leftarrow M[ID', e, i]$ 6: if $E[ID] < e$: 7: if $m' \neq \perp$ 8: $win \leftarrow 1$ 9: else if $(m', e', i') \neq (m, e, i)$ 10: $win \leftarrow 1$ 11: return win
Oracle CREATE(ID, IDs) 1: require $ID \in IDs$ 2: require $ep = 0$ 3: $T \stackrel{\$}{\leftarrow} \text{Exec}(\gamma[ID], \text{crt}, IDs)$ 4: $T[ep] \leftarrow T; ep \leftarrow ep + 1$ 5: $G[0], G[1] \leftarrow IDs$ 6: return T	Oracle REMOVE(ID, ID') 1: require $E[ID] = ep$ 2: $T \stackrel{\$}{\leftarrow} \text{Exec}(\gamma[ID], \text{rem}, \{ID'\})$ 3: require $T \neq \perp$ 4: $T[ep] \leftarrow T; ep \leftarrow ep + 1$ 5: $G[ep] \leftarrow G[ep - 1] \setminus \{ID'\}$ 6: return T	Oracle DELIVER(ID) 1: require $E[ID] < ep$ 2: require $ID \in G[E[ID]]$ 3: $acc \leftarrow \text{Proc}(\gamma[ID], T[E[ID]])$ 4: if $\neg acc$ 5: $win \leftarrow 1$ 6: $E[ID] \leftarrow E[ID] + 1$ 7: if $ID \in G[E[ID]]$ 8: if $G[E[ID]] \neq \gamma[ID].G$ 9: $win \leftarrow 1$ 10: $l[ID] \leftarrow 0$ 11: return win
Oracle SEND(ID, m) 1: require $ID \in G[E[ID]]$ 2: $C \stackrel{\$}{\leftarrow} \text{Send}(\gamma[ID], m)$ 3: if $C = \perp$ 4: $win \leftarrow 1$ 5: $l[ID] \leftarrow l[ID] + 1$ 6: $M[ID, E[ID], l[ID]] \leftarrow m$ 7: $CT[ID, E[ID], l[ID]] \leftarrow C$ 8: return C	Oracle UPDATE(ID) 1: require $E[ID] = ep$ 2: $T \stackrel{\$}{\leftarrow} \text{Exec}(\gamma[ID], \text{upd}, \{ID\})$ 3: require $T \neq \perp$ 4: $T[ep] \leftarrow T; ep \leftarrow ep + 1$ 5: $G[ep] \leftarrow G[ep - 1]$ 6: return T	
	Oracle EXPOSE(ID) 1: return $\gamma[ID]$	

Figure 6.6: Game defining M-CORR_{GM} with adversary \mathcal{A} for $\mathcal{O} = \{\text{CREATE}, \text{SEND}, \text{ADD}, \text{REMOVE}, \text{UPDATE}, \text{EXPOSE}, \text{RECEIVE}, \text{DELIVER}\}$. Lines in teal correspond only to book-keeping and state update operations.

- *Group evolution*: Group operations, namely crt (group creation), add (user addition), rem (user removal), and upd (key update), must have their intended effects on the group when received and processed (checked at lines 4-5 and 7-9 of DELIVER).
- *Group membership consistency*: The list of group members must be consistent among all group members, assuming they process the same sequence of control messages (checked at lines 7-9 of DELIVER)
- *Out-of-order delivery*: Messages corresponding to past epochs must be decryptable if

Chapter 6. WhatsUpp with Sender Keys? Analysis, Improvements and Security Proofs

delivered out-of-order (checked at lines 9-10 of RECEIVE). Messages corresponding to future epochs must be rejected upon reception (checked at lines 6-8 of RECEIVE)

We define correctness below.

Definition 48 (Correctness of GM). A Group Messenger $GM := (\text{Init}, \text{Send}, \text{Recv}, \text{Exec}, \text{Proc})$ is M-CORR if, for all (possibly unbounded) adversaries \mathcal{A} , we have

$$\text{Adv}_{GM}^{\text{m-corr}}(\mathcal{A}) := \Pr[\text{M-CORR}_{GM}(\mathcal{A}) \Rightarrow 1] = 0,$$

where game M-CORR_{GM} is defined in Figure 6.6.

6.3.1 Security Model

We introduce a game-based model of security for our Group Messenger primitive that captures the main desirable security properties of a group messaging scheme. In brief, our game $\text{M-IND}_{GM, C_{gm}}$ captures a partially active adversary who can, in particular, expose the state of users and inject (possibly malformed) messages at any time. We consider the confidentiality of application messages with FS and PCS, and we also model the out-of-order delivery of application and control messages.

Oracles for Security. Compared to correctness, the adversary \mathcal{A} will be given access to the following oracles in the security game:

$\text{CREATE}(ID, ID_s), \text{SEND}(ID, m), \text{ADD}(ID, ID'), \text{REMOVE}(ID, ID'), \text{UPDATE}(ID), \text{EXPOSE}(ID)$:
As in correctness.

$\text{CHAL}(ID, m_0, m_1)$: outputs a ciphertext C_b corresponding to the message m_b sent by ID , where b is the bit that parametrises the game. Namely, \mathcal{A} obtains $C_b \leftarrow \text{Send}(\gamma[ID], m_b)$.

$\text{RECEIVE}(ID, C)$: ID receives a ciphertext C by calling $\text{Recv}(\gamma[ID], C)$. The sender ID' is inferred from output of Recv . In the event of a successful forgery, \mathcal{A} obtains the value of b (and can thus return b to immediately win the game).

$\text{DELIVER}(ID, T)$: ID is delivered a control message T via $\text{Proc}(\gamma[ID], T)$.

Game Description. We introduce the $\text{M-IND}_{GM, C_{gm}}$ game in Figure 6.7, parameterised by bit b that has to be guessed by \mathcal{A} as in a message indistinguishability game. The adversary wins the game if it directly guesses b correctly, which it can always do if it carries out a successful forgery (since it is given b directly in this case). The game is further parameterised by a protocol-specific *cleanness* predicate C_{gm} (sometimes safety predicate [ACDT20]) that rules out trivial attacks and captures the exact security of the protocol.

Game M-IND_{GM,b,C_{gm}}(\mathcal{A}) 1: for all ID : 2: $\gamma[ID] \stackrel{\$}{\leftarrow} \text{Init}(1^\lambda, ID)$ 3: $T[\cdot], M[\cdot], CH[\cdot], SM[\cdot] \leftarrow \perp$ 4: $ep \leftarrow 0$ 5: $E[\cdot], l[\cdot] \leftarrow 0$ 6: $b' \leftarrow \mathcal{A}^{\mathcal{O}}$ 7: require C_{gm} 8: return b'	Oracle ADD(ID, ID') 1: $T \stackrel{\$}{\leftarrow} \text{Exec}(\gamma[ID], \text{add}, \{ID'\})$ 2: require $T \neq \perp$ 3: $T[ep] \leftarrow T; ep \leftarrow ep + 1$ 4: return T	Oracle EXPOSE(ID) 1: return $\gamma[ID]$
Oracle CREATE(ID, IDs) 1: require $ID \in IDs$ 2: require $ep = 0$ 3: $T \stackrel{\$}{\leftarrow} \text{Exec}(\gamma[ID], \text{crt}, IDs)$ 4: $T[ep] \leftarrow T; ep \leftarrow ep + 1$ 5: return T	Oracle REMOVE(ID, ID') 1: $T \stackrel{\$}{\leftarrow} \text{Exec}(\gamma[ID], \text{rem}, \{ID'\})$ 2: require $T \neq \perp$ 3: $T[ep] \leftarrow T; ep \leftarrow ep + 1$ 4: return T	Oracle RECEIVE(ID, C) 1: $(m, ID', e, i) \leftarrow \text{Recv}(\gamma[ID], C)$ 2: if $m \neq \perp$: 3: Update $SM[ID, ID']$ 4: if $\text{recv-forgery}(C)$: 5: return b 6: return
Oracle CHAL(ID, m_0, m_1) 1: require $ m_0 = m_1 $ // Send m_b based on b 2: $C^* \stackrel{\$}{\leftarrow} \text{Send}(\gamma[ID], m_b)$ 3: require $C^* \neq \perp$ 4: $l[ID] \leftarrow l[ID] + 1$ 5: $CH[ID, E[ID], l[ID]] \leftarrow C^*$ 6: return C^*	Oracle UPDATE(ID) 1: $T \stackrel{\$}{\leftarrow} \text{Exec}(\gamma[ID], \text{upd}, \{ID\})$ 2: require $T \neq \perp$ 3: $T[ep] \leftarrow T; ep \leftarrow ep + 1$ 4: return T	Oracle DELIVER(ID, T) 1: $\text{acc} \leftarrow \text{Proc}(\gamma[ID], T)$ 2: require acc 3: $E[ID] \leftarrow E[ID] + 1$ 4: $l[ID] \leftarrow 0$ 5: if $\text{proc-forgery}(T)$: 6: return b 7: return
Oracle SEND(ID, m) 1: $C \stackrel{\$}{\leftarrow} \text{Send}(\gamma[ID], m)$ 2: require $C \neq \perp$ 3: $l[ID] \leftarrow l[ID] + 1$ 4: $M[ID, E[ID], l[ID]] \leftarrow C$ 5: return C		

Figure 6.7: Game defining M-IND_{GM,C_{gm}} with adversary \mathcal{A} and cleanness predicate C_{gm} for $\mathcal{O} = \{\text{CREATE}, \text{CHAL}, \text{ADD}, \text{REMOVE}, \text{UPDATE}, \text{SEND}, \text{EXPOSE}, \text{RECEIVE}, \text{DELIVER}\}$. Lines in teal correspond only to bookkeeping and state update operations.

Message Epochs. We define a function $\text{m-ep}(ID, ID', q)$ that indicates the highest message epoch (e, i) , as output by the Recv algorithm, for which a user ID has received a message from ID' at the time of query q for $ID \neq ID'$. For $ID = ID'$, this indicates the local state value for $(E[ID], l[ID])$. The m-ep function reflects the view of user ID' by user ID .

Dictionaries. The challenger keeps a record of messages and game variables in several dictionaries. The state of each party is stored in $\gamma[\cdot]$ and updated when an algorithm is called on a given $\gamma[ID]$. Ciphertexts and challenged ciphertexts are stored in M and CH , respectively, each of them indexed by an ID and a message epoch (e, i) . The unique honest control message that starts a given epoch e is stored in $T[e]$, and the most recent epoch of the group is stored in variable ep (note that we implicitly assume a total ordering of control messages). The current

Chapter 6. WhatsUpp with Sender Keys? Analysis, Improvements and Security Proofs

message epoch of ID is stored in $E[ID], I[ID]$. Even if each control message in T corresponds to a single epoch, different parties can be in different epochs. We say ID is in epoch e before query q if the last control message processed by ID before query q is $T[e]$.

The message epochs corresponding to skipped messages from sender ID' stored by ID are kept in $SM[ID, ID']$. We keep SM updated in the RECEIVE oracle as follows: given a message epoch (e, i) and an ID' output by Recv, if $(e, i) \in SM[ID, ID']$ then (e, i) is erased from $SM[ID, ID']$. Otherwise, we add all pairs (e', i') such that $(e', i') < (e, i)$ and (e', i') corresponds to all messages sent by ID' not delivered to ID .

Outcome. After q oracle queries, \mathcal{A} outputs a guess b' of b if the cleanness predicate C_{gm} is satisfied (otherwise the game aborts). \mathcal{A} can win the game in three different ways: by directly guessing the challenge bit correctly, by injecting a forged application message via RECEIVE successfully, or by injecting a forged control message via DELIVER. The cleanness predicate C_{gm} parameterises the security of a given protocol by restricting the capabilities of the adversary to exclude a class of attacks. Additionally, we explicitly state predicates *recv-forgery* and *proc-forgery* in our game, which model the conditions under which a RECEIVE or DELIVER call result in a successful forgery (leaking b to \mathcal{A}). We expand on these predicates in Section 6.3.2.

We define security below.

Definition 49 (Message indistinguishability of GM). A Group Messenger $GM := (\text{Init}, \text{Send}, \text{Recv}, \text{Exec}, \text{Proc})$ is M-IND with respect to cleanness predicate C_{gm} if, for all efficient adversaries \mathcal{A} , we have

$$\text{Adv}_{GM, C_{gm}}^{\text{m-ind}}(\mathcal{A}) := \left| \Pr[\text{M-IND}_{GM, 1, C_{gm}}^{\mathcal{A}} \Rightarrow 1] - \Pr[\text{M-IND}_{GM, 0, C_{gm}}^{\mathcal{A}} \Rightarrow 1] \right| = \text{negl}.$$

where game $\text{M-IND}_{GM, C_{gm}}$ is defined in Figure 6.7.

Related Security Notions. Our security model takes inspiration from the game-based modelling developed for MLS and for CGKA (which we introduced formally in Section 5.2). Nevertheless, it is not ideal to adopt these models as they consider a single group key, which is not directly compatible with a Sender Keys (or similar) approach to group messaging. The closest security model to ours in the literature comes from that of Weidner et al. for DCGKA [WKHB21], which however does not consider message injections nor adaptive security.

Limitations. Our security game allows a single successful injection to occur, since after this point the adversary is given the secret bit for free. That is, we do not allow ‘trivial’ message forgeries that do not result in the adversary winning the game. Hence, full active security cannot be captured by our modelling. Like several other models in the literature (e.g., [ACDT20, KPPW⁺21, WKHB21]), our security model considers a single group (see [CHK21] for an analysis of cross-group security) and ignores randomness exposure or manipulation [BRV20].

6.3.2 Modelling Two-Party Channel Ciphertexts

Given that the GM protocol uses two-party channels (as Sender Keys does), these need to be modelled accurately within the GM security game, particularly to describe forgeries via the *recv-forgery* and *proc-forgery* predicates. We introduce additional notation to define how two-party ciphertexts can be sent alongside GM messages; we opt for such modelling for convenience, as in this way the adversary gets access to all two-party channels explicitly. We remark that this subsection can be skipped for GM protocols that do not employ two-party channels, since in this case, simply any ciphertext input to Proc (resp. Recv) that was not previously output by Exec (resp. Proc) would be considered a forgery.

Essentially, we want to capture the fact that an Exec or Send call can output *several* two-party channel ciphertexts, whereas Proc and Recv should only take as input a *single* two-party channel ciphertext (i.e., the one intended for the caller) for efficiency. We thus assume input/output ciphertexts and control messages for group messenger algorithms take the following form. Let C_{2pc} be a 2PC ciphertext and let T_{core} (resp. C_{core}) be the remaining part of a control (resp. application) message in the GM primitive. For output, we assume control messages output by Exec are of the form $(T_{core}, C_{2pc}^1, \dots, C_{2pc}^k)$ for some k , and ciphertexts output by Send are of the form $(C_{core}, C_{2pc}^1, \dots, C_{2pc}^k)$. For input, we assume control messages input to Exec (resp. to Recv) are of the form (T_{core}, C_{2pc}) (resp. (C_{core}, C_{2pc})).

Forgery Predicates. We define the predicates *proc-forgery* and *recv-forgery* in Figure 6.8 using the input/output semantics introduced above. Used in Figure 6.7, the purpose of these predicates is to handle ciphertext ‘splitting’ resulting from the use of two-party channels. Without accounting for this splitting, forgeries could be defined as usual. Essentially, we consider that a control message $T^* = (T_{core}^*, C_{2pc}^*)$ is a forgery whenever either T_{core}^* or C_{2pc}^* are not part of an honestly generated message (i.e. in $T[\cdot]$). Forgeries for Recv are defined analogously.

$\text{proc-forgery}(T^* = (T_{core}^*, C_{2pc}^*)) :$	$\exists \{ (T, \vec{C}), (T', \vec{C}') \} \subseteq T[\cdot] :$
$(T_{core}^*, C_{2pc}^*) \in \{ (T, C_i), (T', C_i), (T, C'_j) \} \wedge (C_i \in \vec{C}) \wedge (C'_j \in \vec{C}')$	
$\text{recv-forgery}(C = (C_{core}^*, C_{2pc}^*)) :$	$\exists \{ (C_0, \vec{C}), (C'_0, \vec{C}') \} \subseteq \text{CH}[\cdot] \cup \text{M}[\cdot] :$
$(C_{core}^*, C_{2pc}^*) \in \{ (C_0, C_i), (C'_0, C_i), (C_0, C'_j) \} \wedge (C_i \in \vec{C}) \wedge (C'_j \in \vec{C}')$	

Figure 6.8: Predicates that determine what is considered a forgery in Figure 6.7 for algorithms Proc and Recv.

The predicates imply that it is *not* considered a forgery if a two-party ciphertext is received with a different control message/ciphertext than it was sent with. That is, the adversary is allowed to mix-and-match ciphertexts, i.e., by replacing the C_i corresponding to some T (resp. C_0) by C'_j corresponding to some other T' (resp. C'_0).

6.4 Sender Keys

In this section, we explain and provide a simplified protocol description for Sender Keys. We state our security theorem for Sender Keys, describe how we model the cleanness predicates that characterise its security as a Group Messenger and then sketch its security. In Section 6.6 we provide a complete protocol description and proof.

6.4.1 Protocol

We describe the Sender Keys protocol in our GM syntax according to the details inferred from [Wha20] and [M⁺16], although we acknowledge that our interpretation may not precisely match the closed-source implementation of WhatsApp. In this subsection we present a detailed overview of the main algorithms depicted in Figure 6.9. For Exec and Proc, we only present the remove operation as it involves key refreshing and is considered the most complex, while the create, add, and update operations follow a similar approach. For the sake of clarity, we make some simplifications in this section, but the complete protocol logic can be found later in Figures 6.19 to 6.22 (in Section 6.6, where we also provide supplementary descriptions of the protocol logic).

Two-Party Channels and the Server. The Sender Keys protocol assumes the existence of authenticated and secure two-party communication channels between each pair of users, which can be achieved through the use of Signal’s Double Ratchet protocol [MP16a] also used by WhatsApp [Wha20]. Additionally, the protocol relies on a central server to distribute both control messages and application messages. We assume that the server provides a *total ordering for control messages*, ensuring that all parties process control messages in the correct order.⁶ Total ordering is not required for application messages. User authentication is initially performed via the central server (modelled here with 2PC.InitCh), after which users authenticate other group members through the underlying two-party communication channel. We note that this deviates from other work in the literature such as [AJM22] where the *authentication service* is different to the *delivery service*.

Primitives. The protocol relies on standard primitives including a symmetric encryption scheme $\text{SymEnc} = (\text{Gen}, \text{Enc}, \text{Dec})$, a signature scheme $\text{Sig} = (\text{Gen}, \text{Sign}, \text{Vrfy})$, and two different key derivation functions H_1, H_2 (our improved protocol also uses message authentication codes). We include formal definitions in Chapter 2.

State Initialisation. Each user is assumed to maintain a state γ containing: a secret key used for signing ssk , a list of current group members \mathbf{G} , the current epoch ep , the current index of their chain key i_{ck} (indicating the number of times the user’s sender key has been ratcheted forward), a list of key counters kc (indicating the number of times that a sender key has been re-sampled since ID initialised their state), a dictionary of sender keys $\text{SK}[\cdot] := (\text{spk}_{ID}, \text{ck}_{ID}, i_{\text{ck}})$

⁶We remark that total ordering is a standard assumption in the CGKA line of work [ACDT20, KPPW⁺21, ACJM20, AJM22, ACDT21a] and is assumed by MLS.

<pre> Init($ID, 1^\lambda$) 1: $\gamma.ME \leftarrow ID$ 2: $\gamma.(ssk, \mathbf{G}, ep, i_{ME}) \leftarrow \perp$ 3: $\gamma.(SK[\cdot], MK[\cdot], kc[\cdot]) \leftarrow \perp$ 4: return γ Send(m, γ) 1: require $ME \in \mathbf{G}$ 2: if $SK[ME, kc[ME]] = \perp$: // Sample fresh sender key 3: $\vec{C} \leftarrow \text{PreSendFirst}()$ 4: if $i_{ME} = 0$: 5: $\vec{C} \leftarrow (\vec{C}, \text{SendToMissing}())$ 6: $mk \leftarrow H_1(SK[ME, kc[ME]].ck)$ 7: $ct \xleftarrow{\\$} \text{SymEnc.Enc}(mk, m)$ // UpdateCK also updates i_{ME} 8: $\text{UpdateCK}(ME, kc[ME])$ 9: $M \leftarrow (ct, (ep, i_{ME}), kc[ME], i_{ck}, ME)$ 10: $\sigma \xleftarrow{\\$} \text{Sig.Sign}(ssk, M)$ 11: return $C := ((M, \sigma), \vec{C})$ Exec($\gamma, \text{cmd} = \text{rem}, ID$) 1: require $ID \in \mathbf{G}$ 2: $\vec{C}[\cdot] \leftarrow \perp$ 3: $T \leftarrow (\text{rem}, ME, ID, ep + 1)$ 4: return (T, \vec{C}) </pre>	<pre> Recv($\gamma, C = ((M, \sigma), C_{2pc})$) 1: parse M as $(ct, (e, i), kc', i_{ck}', ID)$ 2: require $ID \in \mathbf{G}$ 3: if $SK[ID, kc'] = \perp$: 4: $(SK[ID, kc'], kc^*, ep', aux, ID^*) \leftarrow \text{2PC.Recv}(\gamma, C_{2pc})$ 5: require $(ID, e, kc') = (ID^*, ep', kc^*)$ 6: $\text{DeleteOldCK}(ID, aux)$ 7: else require $C_{2pc} = \perp$ 8: require $e \leq ep$ 9: require $\text{Sig.Vrfy}(SK[ID, kc'].spk, \sigma, M)$ // Derive or fetch mk from state 10: $mk \leftarrow \text{UpdateKeysRecv}()$ 11: $m \leftarrow \text{SymEnc.Dec}(mk, ct)$ 12: return (m, ID, e, i) Proc($\gamma, (T = (\text{rem}, ID, ID', ep'), C_{2pc})$) 1: require $ID \in \mathbf{G} \wedge C_{2pc} = \perp$ 2: require $ep' = ep + 1$ 3: $\mathbf{G} \leftarrow \{ID'\}$ 4: $ep \leftarrow ep + 1; i_{ME} \leftarrow 0$ 5: for all $ID^* \in \mathbf{G}$: 6: $kc[ID^*] \leftarrow kc[ID^*] + 1$ 7: $SK[ID', \cdot] \leftarrow \perp$ 8: if $ID = ME$: 9: $\gamma \leftarrow \perp$ // delete γ 10: return true </pre>
---------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

Figure 6.9: Simplified Sender Keys protocol description (Init, Send, Recv, and example Exec and Proc operations for $\text{cmd} = \text{rem}$) built from symmetric encryption scheme SymEnc, signature scheme Sig, PRG $H = (H_1, H_2)$ and two-party channels scheme 2PC. For readability we use and update state variables without explicitly referring to γ , and assume stateful algorithms implicitly return (possibly updated) state. For conditions of the form “**require** T ” when T is false, the function outputs \perp except for the state which is not updated. The full protocol description is provided in Section 6.6 where the helper functions PreSendFirst, SendToMissing, UpdateCK, DeleteOldCK and UpdateKeysRecv that appear in this Figure are defined.

indexed by a user ID and a key counter, and a list of message keys MK. The Init algorithm initialises the state variable of users; in practice this is done by a user when they install the messaging application.

Group Creation. This occurs via $\text{Exec}(\gamma, \text{crt}, IDs)$, which takes a list of users $\mathbf{G} := \{ID_1, \dots, ID_{|\mathbf{G}|}\}$ as input; two-party channels are initialised by users upon processing the control

Chapter 6. WhatsUpp with Sender Keys? Analysis, Improvements and Security Proofs

message via `2PC.InitCh` (c.f. Section 6.6).

Message Sending. To send an application message m to the group, every $ID \in \mathbf{G}$ must have the caller's (ME) sender key. The process is as follows:

- If ME does not have a sender key, ME calls helper `PreSendFirst`, which first generates a fresh sender key $((\gamma.ssk, spk) \xleftarrow{\$} \text{Sig.Gen}(1^\lambda)$ and $ck \xleftarrow{\$} \{0, 1\}^\lambda$). The sender key is then set as $SK[ME, kc[ME]] \leftarrow (spk, ck, i_{ck})$ where $i_{ck} = 0$. ME shares this key with each $ID \in \mathbf{G}$ using `2PC.Send`, resulting in a vector of ciphertexts \vec{C} .
- If ME has a non-empty sender key but not all parties have it ME , calls helper `SendToMissing` to share the key with them via `2PC.Send`, therein updating \vec{C} .

Then, ME generates a new message key mk from their chain key $SK[ME, kc[ME]].ck$, encrypts m using mk , and ratchets its chain key forward by setting $ck \leftarrow H_2(SK[ME, kc[ME]].ck)$. Finally, ME signs the ciphertext and sends it together with \vec{C} .

Message Receiving. To receive a message from ID , ME follows these steps:

- ME checks if they have ID 's sender key $SK[ID, kc']$ corresponding to the key counter kc' indicated in the received message. If ME does not have it, they retrieve it from the two-party ciphertext C_{2pc} using `2PC.Recv`, aborting the `Recv` call if the sender key cannot be found.
- ME performs epoch consistency checks and verifies the signature on the ciphertext using the signature public key $SK[ID, kc'].spk$.
- The message key mk required to decrypt the message is computed from the chain keys as $mk \leftarrow H_1(SK[ID, kc'].ck)$, and is deleted after use.

Out-of-Order Messages. In the scenario of out-of-order message delivery (handled by helper `UpdateKeysRecv`), the following cases arise (we let $i_{ck} := SK[ID, kc'].i_{ck}$):

- If the received message comes from a past epoch $(e, i) < (ep, i_{ck})$, ME searches for the relevant skipped message key in MK .
- If $e = ep$ and $i > i_{ck}$, ME ratchets ID 's chain key $i - i_{ck}$ times, and stores the skipped message keys in MK .
- If $e > ep$, the message reception fails since ME is not synchronised with the latest group epoch and cannot (even) determine whether the sender is still a member of the group.

Handling out-of-order message delivery constitutes a significant portion of the protocol's logic. For instance, parties must keep track of (and announce) the highest i_{ck} associated with a given kc . Failing to do so can result in correctness and security issues, as parties may overlook the need to store and delete keys in MK .

Key Updates. In at least Signal’s implementation of Sender Keys (it is not mentioned in WhatsApp’s white paper [Wha20]) a simple (but somewhat weak) on-demand key update mechanism is supported. A party ME can update its key material via $\text{Exec}(\gamma, \text{crt}, ME)$. This operation lazily samples a fresh sender key $(\text{spk}, \text{ck}, 0)$ and distributes it over the two-party channels. All users sample a fresh key after processing a removal.

Membership Changes. The protocol allows individual group members to be added or removed from the group via $\text{Exec}(\gamma, \text{add}, ID)$ and $\text{Exec}(\gamma, \text{rem}, ID)$. These operations result in the distribution of a control message T to the group sent in clear. Newly added members are also sent a welcome 2PC ciphertext containing group information. Note that we model single adds/removes for simplicity but this can be extended in a straightforward manner to handle batched group changes.

Upon processing a control message via $\text{Proc}(\gamma, T)$, ME proceeds as follows:

- If some ID^* is being removed at epoch e , ME erases all sender keys corresponding to ID^* (except for skipped message keys).⁷ The sender keys corresponding to other users are replaced with new ones when receiving messages from epoch $e' \geq e$, ensuring messages sent concurrently with the removal can still be received by ME .
- If some ID^* is being added, ME initialises its 2PC with ID^* via 2PC.InitCh .
- If ME is itself removed, it erases its state. If it is added to some group (or processes a create message), it initialises two-party channels with every $ID \in \mathbf{G}$.

Note that after either updating or adding or removing a user, for performance and security reasons, new sender keys are only distributed once a party sends his first message.

6.4.2 Security

In this subsection, we argue that Sender Keys as described in Section 6.4 is secure with respect to our security model in Section 6.3.1. However, the security captured by our cleanness predicate is far from theoretically optimal since Sender Keys is relatively weak in security, and so in Section 6.5 we strengthen it by modifying the protocol in different ways. Our predicates are parameterised by the security of the underlying two-party channels. We first state our main theorem below.

Theorem 20. Consider IND-CPA symmetric encryption scheme SymEnc , SUF-CMA signature scheme Sig , PRG function $H : \{0, 1\}^\lambda \rightarrow \{0, 1\}^\lambda \times \{0, 1\}^\lambda$ and 2PC-IND two-party channels scheme 2PC with respect to PCS bound δ and cleanness predicate $C_{2\text{pc}}$ (Figures 6.4 and 6.5). Then, for Group Messenger Sender Keys (Figures 6.19 to 6.22) we have, with respect to cleanness

⁷A different deletion schedule may be applied as long as these keys are clearly marked as being no longer valid, e.g., if ID^* announces its maximum i_{ck} value over two-party channels when it processes its own removal.

Chapter 6. WhatsUpp with Sender Keys? Analysis, Improvements and Security Proofs

predicate $C_{\text{gm}} = C_{\text{sk}}^\Delta$ (Figure 6.13), that, for every efficient adversary \mathcal{A} that makes at most q oracle queries, one can build an adversary \mathcal{B} such that

$$\text{Adv}_{\text{GM}, C_{\text{gm}}}^{\text{m-ind}}(\mathcal{A}) \leq 2 \cdot \epsilon_{2\text{pc}} + q^2 \cdot (\epsilon_{2\text{pc}} + \epsilon_{\text{sym}} + q \cdot \epsilon_{\text{prg}}) + q \cdot \epsilon_{\text{sig}},$$

where $\epsilon_{\text{sym}} = \text{Adv}_{\text{SymEnc}}^{\text{ind-cpa}}(\mathcal{B})$, $\epsilon_{\text{sig}} = \text{Adv}_{\text{Sig}}^{\text{suf-cma}}(\mathcal{B})$, $\epsilon_{\text{prg}} = \text{Adv}_{\text{H}}^{\text{prg}}(\mathcal{B})$ and $\epsilon_{2\text{pc}} = \text{Adv}_{2\text{PC}, C_{2\text{pc}}, \Delta}^{2\text{pc-ind}}(\mathcal{B})$.

We define C_{sk} below and prove the theorem in Section 6.6.2. A proof sketch is provided at the end of this section. Our security notion is adaptive as users can adaptively call oracles and in particular compromise users. Security is tighter when we restrict the game to consider non-adaptive adversaries as described below.

Corollary 1. Under the same conditions of Theorem 20, and considering a non-adaptive security game, for every efficient adversary \mathcal{A} that makes at most q oracle queries, one can build an adversary \mathcal{B} such that

$$\text{Adv}_{\text{GM}, C_{\text{gm}}}^{\text{m-ind}}(\mathcal{A}) \leq 2 \cdot \epsilon_{2\text{pc}} + q \cdot (\epsilon_{\text{sym}} + q \cdot \epsilon_{\text{prg}}) + q \cdot \epsilon_{\text{sig}},$$

where ϵ_{sym} , ϵ_{sig} , ϵ_{prg} and $\epsilon_{2\text{pc}}$ are defined above.

Sender Keys and Two-Party Channels. To illustrate how the cleanness predicates for Sender Keys must depend on the underlying two-party channels, consider a strongly secure two-party channel 2PC that provides optimal FS and PCS. Now, consider an execution of Sender Keys where all parties share the same view of the group $G = \{ID_1, ID_2, ID_3\}$, in which

1. ID_1 generates a control message $(T_{\text{core}}, \vec{C})$ to remove party ID_3 ($q_1 = \text{REMOVE}(ID_1, ID_3)$),
2. ID_1 and ID_2 process T ($q_{2,1} = \text{DELIVER}(ID_1, (T_{\text{core}}, \vec{C}[ID_1]))$, $q_{2,2} = \text{DELIVER}(ID_2, (T_{\text{core}}, \vec{C}[ID_2]))$),
3. \mathcal{A} exposes ID_2 ($q_3 = \text{EXPOSE}(ID_2)$);
4. ID_1 sends an application message ($q_4 = \text{SEND}(ID_1, m)$).

Recall that in step 4, ID_1 samples a new sender key that it sends to ID_2 over 2PC, since processing remove messages results in the sender keys of all parties being refreshed. Even with optimally-secure 2PC, the adversary will be able to decrypt the key sent over 2PC (by the correctness of the channel) and thus decrypt the ciphertext output in query q_4 .

Cleanness. Our goal is to describe a suitable cleanness predicate C_{sk} for Sender Keys. The intuition behind this cleanness predicate is based on the following observations about the protocol:

- The exposure of a group member compromises the security of subsequent chain and message keys⁸ until a secure key refresh takes place. This enables the adversary to forge messages since they also gain access to the exposed signature keys.
- Control messages can be trivially forged and injected by a network adversary as they are not authenticated.
- Forward-secure confidentiality holds except for messages delivered out-of-order since parties only delete message keys after using them, so a message that is delayed forever results in the corresponding message key never being deleted.⁹
- All parties recover from state exposure (via $\text{EXPOSE}(ID)$) after security on the two-party channels is restored (considering the PCS bound Δ) and then either a) a removal is made effective, or b) all parties update their keys successfully.

To formalise the security predicate we introduce conventions for tracking the channel epochs of each user's two-party channels. We assume the game M-IND maintains the largest channel epoch-index for each user's two-party channels over time. The game obtains this information by observing the channel epoch-index pairs generated by the 2PC.Send and 2PC.Rec operations within the group messenger. Specifically, we use a variable of the form $\text{El}[ID, ID']$, where $\text{El}[ID, ID']$ represents the largest channel epoch-index pair from ID 's perspective for the channel between them and user ID' , as for two-party channels. More generally, two-party state variables that we use below can be tracked easily by an M-IND adversary such that our predicates are well-defined.

The refresh $_{\Delta}$ Predicate. We define the predicate $\text{refresh}_{\Delta}(ID, ID', q_i, e)$, parameterised by the PCS bound $\Delta > 0$ of the underlying two-party channels. Informally, given that ID' is exposed in query q_i ($q_i = \text{EXPOSE}(ID')$), $\text{refresh}_{\Delta}(ID, ID', q_i, e)$ is true if the (ID, ID') channel has healed and *then* ID has sampled a fresh sender key in or by epoch e (or will do so upon their next Send call). If the predicate is true, ID has recovered from the exposure in q_i .

More formally, let $(e_{2\text{pc}}, i_{2\text{pc}}) = \max\{\text{El}[ID', ID; q_i], \text{El}[ID, ID'; q_i]\}$. Then $\text{refresh}_{\Delta}(ID, ID', q_i, e)$ is true if a) for $(e'_{2\text{pc}}, i'_{2\text{pc}}) = \text{El}[ID', ID; q_j]$ for some $j > i$, $e'_{2\text{pc}} \geq e_{2\text{pc}} + \Delta$ holds; and b) during query q_k with $k \geq j$, member ID processes one of the following control messages corresponding to epoch e :

1. a removal of some member ID^* ,
2. an addition of ID itself,
3. a group creation message, or

⁸Although it is not captured in our model, note that the exposure of a message key alone only compromises the message it refers to and does not (computationally) leak information about the chain key or other message keys.

⁹In practice, applications like WhatsApp and Signal bound the amount of (logical) time that keys are active for and the total number of keys that can be stored at once.

Chapter 6. WhatsUpp with Sender Keys? Analysis, Improvements and Security Proofs

4. an update from ID itself.

As such, if ID executes (and processes) an update that involves sending new key material over a refreshed two-party channel, this key material should be safe. We also define a simpler predicate $\text{refresh-s}(ID, e)$ which is true if member ID processes one of the aforementioned control messages corresponding to epoch e . Observe that both refresh_Δ and refresh-s events may only happen when ID moves to a new group epoch e .

Cleanness for Sender Keys. We divide our cleanness predicate into three components (challenge, injection, concurrency) that we specify below. The final predicate (the conjunction of these three components) is defined in Figure 6.13.

Challenge (Figure 6.10). The effect of this predicate is to prevent challenges on exposed users (i.e., due to EXPOSE calls). After exposing (with query q_i) *any* user ID' , adversarial queries to CHAL are disallowed for every ID in the group until $\text{refresh}_\Delta(ID, ID', q_i, e)$ occurs for some later epoch $e > E[ID'; q_i]$. Note that this only restricts challenge queries q_j where $i < j$. To capture forward security precisely, some challenges made before an exposure ($i > j$) are also forbidden. These affect messages sent by some ID in message epochs $(e, i) \geq \text{m-ep}(ID', ID, q_i)$, which correspond to keys that ID' still stores (including skipped message keys stored at exposure time) or can derive due to being in a previous message epoch (for example if the user is offline).

$$\begin{array}{|l}
 \hline
 C_{\text{sk-chall}}^\Delta \\
 \hline
 \forall (i, j, ID, ID') : q_i = \text{EXPOSE}(ID') \wedge q_j = \text{CHAL}(ID, \cdot, \cdot), \\
 (i > j \wedge \text{m-ep}(ID', ID, q_i) > (E[ID; q_j], l[ID; q_j]) \wedge \\
 (E[ID; q_j], l[ID; q_j]) \notin \text{SM}[ID', ID, q_i]) \\
 \vee (i < j \wedge \exists e : E[ID'; q_i] < e \leq E[ID; q_j] \wedge \text{refresh}_\Delta(ID, ID', q_i, e)) \\
 \hline
 \end{array}$$

Figure 6.10: Challenge cleanness predicate for Sender Keys where the adversary makes oracle queries q_1, \dots, q_q .

Injection (Figure 6.11). Firstly, let us recall the two-party ciphertext splitting semantics defined in Section 6.3.2. Namely, a GM ciphertext C naturally splits into $C = (C_{\text{core}}, C_{2\text{pc}})$ where $C_{2\text{pc}}$ is processed by the two-party channels. An injection is said to have occurred when a message with a forged C_{core} and/or $C_{2\text{pc}}$ was successfully processed.

We define the injection predicate to prevent injections of *application messages* coming from a user that has been exposed and has not refreshed its keys (we consider control messages in the last predicate). We start with the definition for C_{core} . After exposing *a specific* user ID' with query q_i , \mathcal{A} cannot make a query $q_j = \text{RECEIVE}(ID, C)$ to impersonate ID' with a forgery ciphertext C corresponding to some epoch e^* (i.e., such that tuple (ID', e^*) is output by $\text{Recv}(\gamma[ID], C)$ in the game) in the following situations:

1. $e^* \geq E[ID'; q_i]$ and there hasn't been a $\text{refresh}_\Delta(ID', ID, q_i, e')$ event for the sender ID' at some epoch e' such that $E[ID'; q_i] < e' \leq e^*$, where the receiver ID has also processed the key update from ID' 's message at injection time, i.e., $E[ID; q_j] \geq e'$.

2. $e^* < E[ID'; q_i]$ but the signature key of ID' at epoch e^* was the same key as in the exposure epoch $E[ID'; q_i]$. Formally, this is expressed by the condition that there has not been any event $\text{refresh-s}(ID', e')$ for an epoch $e^* < e' \leq E[ID'; q_i]$.

For C_{2pc} , we directly adopt the injection cleanness predicate $C_{2pc\text{-inj}}$ used to define two-party channel security (Figure 6.5). For additional clarity, we parametrise the predicates by the ciphertexts C_{core}, C_{2pc} . We also define the auxiliary predicate $C_{sk\text{-inj-core}}^\Delta(C_{core})$ in Figure 6.11.

$C_{sk\text{-inj-core}}^\Delta(C_{core}) :$	$\forall (i, j, ID, ID') :$
$(C_{core}, \cdot) \notin M[ID', \cdot; q_j] \wedge (i < j) \wedge q_i = \text{EXPOSE}(ID') \wedge$	
$q_j = \text{RECEIVE}(ID, (C_{core}, \cdot)) \wedge (\cdot, e^*, \cdot, ID') \leftarrow \text{Recv}(\gamma[ID; q_j], (C_{core}, \cdot)) \text{ in } q_j,$	
$\exists e' : [(E[ID'; q_i] < e' \leq e^* \wedge \text{refresh}_\Delta(ID', ID, q_i, e'))$	
$\vee (e^* < e' \leq E[ID'; q_i] \wedge \text{refresh-s}(ID', e'))]$	
$C_{sk\text{-inj}}^\Delta :$	$\forall C_{core}, C_{2pc}, C_{sk\text{-inj-core}}^\Delta(C_{core}) \wedge C_{2pc\text{-inj}}(C_{2pc})$

Figure 6.11: Auxiliary core injection cleanness predicate (top) and injection cleanness predicate (bottom) for Sender Keys, where the adversary makes oracle queries q_1, \dots, q_q . The injection cleanness predicate additionally uses $C_{2pc\text{-inj}}$ (Figure 6.5).

Concurrency (Figure 6.12). This predicate ensures several properties in the protocol. Firstly, it enforces that users process control message in the same order (albeit they need not be synchronised beyond this restriction). Additionally, it prevents the injection of *all* control messages. It is important to note that control messages are not signed in the core protocol, making injections trivial. Furthermore, the predicate guarantees that every user proposing a group change (via the Exec, ADD, REMOVE or UPDATE oracles) is in the most recent epoch. In practice, this predicate ensures that there is a unique honest control message in each epoch of the game. To complete the predicate, we further enforce that the two-party channel ciphertext C_{2pc} input to each DELIVER call is either honest or is allowed by the two-party channel injection predicate $C_{2pc\text{-inj}}$ (Figure 6.5).

The concurrency predicate ensures both security and correctness by addressing scenarios where members propose concurrent group changes or process group changes in different orders. Without enforcing this predicate, the behaviour of the protocol, as specified in this work, becomes ill-defined.

$C_{sk\text{-con}} :$	$\forall (i, ID) : q_i = \text{DELIVER}(ID, (T_{core}, C_{2pc})), \exists j < i :$
$q_j = (\text{ADD or REMOVE or UPDATE or CREATE})(ID, \cdot) \wedge \exists e' : (T, \vec{C}) = T[e'] \wedge$	
$(C_{2pc} \in \vec{C} \vee \vec{C} = \perp \vee C_{2pc\text{-inj}}(C_{2pc})) \wedge (E[ID; q_i] = e' - 1 = E[ID; q_j])$	

Figure 6.12: Concurrency cleanness predicate in the ideal case where the adversary makes oracle queries q_1, \dots, q_q . This predicate additionally uses $C_{2pc\text{-inj}}$ (Figure 6.5).

Limitations and Extensions. Our cleanness predicate enforces a total ordering on control messages, in contrast to considering causal ordering such as in Weidner et al.'s mod-

$$C_{sk}^\Delta : C_{sk\text{-chall}}^\Delta \wedge C_{sk\text{-inj}}^\Delta \wedge C_{sk\text{-con}}$$

Figure 6.13: Sender Keys cleanness predicate which makes use of sub-predicates defined in Figures 6.10 to 6.12.

elling [WKHB21], or no ordering at all. This assumption is consistent with real-world protocols (as in WhatsApp) where a central server is trusted to provide such an ordering, but makes our model unsuitable for decentralised protocols. If our security model allowed for it, one could modify our cleanness predicates to allow for ‘trivial’ injections that are non-winning, by not giving the adversary the challenger’s bit b given that the forgery is trivial (i.e., it violates the injection predicate). Our concurrency predicate and security model could be strengthened to allow several Exec calls in an epoch, from which the network chooses one that is processed to all parties, which has been modelled for TreeKEM in the past [ACDT20].

Proof Sketch for Theorem 20. Towards proving the theorem (see Section 6.6.2 for the full proof), we construct a series of hybrids. We first transition to a game where injections on the two-party channels are disallowed, following from their underlying security. After that, we transition to a game where oracle RECEIVE never outputs challenge bit b , reducing the transition to SUF-CMA signature security, while still excluding trivial injections due to cleanness. Then, we move to a game where the adversary is limited to a single CHAL query, losing a factor of q in the resulting reduction. Subsequently, we transition to a game where the message key used in the CHAL query (if it exists) is replaced by a uniformly random key that remains unknown to the adversary due to cleanness, and the two-party ciphertexts that send the key’s ancestor chain key are replaced by dummy ciphertexts, which follows from the 2PC security and the PRG security of (H_1, H_2) . Finally, we directly reduce to the IND-CPA security of the symmetric encryption scheme.

6.5 Analysis and Improvements

For the proof of security of Sender Keys (Theorem 20) to go through, we need to impose severe restrictions on the adversarial behaviour through the cleanness predicate C_{sk} . Hence, even though we managed to prove Sender Keys secure, we did so in a weak model that reveals important security shortcomings of the protocol. In this section, we elaborate on these limitations and propose changes to enhance security while maintaining efficiency. Some of these findings were presented in a preliminary analysis in [BCG22].

6.5.1 Security Analysis and Limitations

Injection of Control Messages. Our first observation is that control messages lack user authentication, necessitating a high level of trust in the server to prevent the crafting of its own messages. To address this, in predicate $C_{sk\text{-con}}$ we need to enforce that every delivered control

message has been honestly generated. A server deviating from standard behavior could mount a host of attacks. We give three examples below.

Censorship attack: The server can remove any member(s) ID from \mathbf{G} such that all remaining members assume ID left the group by himself, whilst ID believes a different user ID' removed him.

- The server delivers a control message $T := (\text{rem}, ID, ID, \cdot)$ where $(T, \dots) \leftarrow \text{Exec}(\cdot, \text{rem}, ID)$ to every $ID' \in \mathbf{G} \setminus \{ID\}$.
- The server delivers a control message $T' := (\text{rem}, ID', ID, \cdot)$ where $(T', \dots) \leftarrow \text{Exec}(\cdot, \text{rem}, ID)$ to $ID \in \mathbf{G}$.

Burgle into the group attack: This attack, observed by Rösler et al. [RMS18], allows the server to add any member(s) ID to \mathbf{G} . For this, the server just delivers a control message $T := (\text{add}, \cdot, ID, \cdot)$ where $(T, \dots) \leftarrow \text{Exec}(\cdot, \text{add}, ID)$ to every $ID' \in \mathbf{G}$.

Unsafe group administration: In general, administration cannot be enforced or trusted due to the lack of authentication of control messages, as discussed in Section 5.1.

Weak Post-Compromise Security. Sender Keys offers only a limited form of PCS. Essentially, a refresh_Δ event is the only possibility for ID to recover from a state compromise. This event only occurs whenever another user is removed or whenever ID triggers an on-demand update (or trivially when ID is new to the group). On-demand updates are supported by our primitive syntax and protocol description, but it is not clear whether they are implemented in practice beyond Signal (there is no mention to them in WhatsApp's white paper [Wha20], for instance).

Moreover, the update mechanism is not satisfactory. Since only the updater ID refreshes its sender key, this allows a passive adversary to eavesdrop on messages sent by any other group member due to the adversary's knowledge of the chain keys corresponding to those members. Extending the update mechanism to the entire group in a naive manner would result in a total communication complexity of $\mathcal{O}(n^2)$.

PCS and Two-Party Channels. PCS guarantees are even weaker due to the reliance of Sender Keys on two-party channels. As parametrised by refresh_Δ , if ID sends new key material over a two-party channel with ID' that has not been healed (after Δ back-and-forth messages) since the last exposure of either ID or ID' , then such key material is still compromised. In practice, if the state of ID is compromised, both the group and the two-party sessions will be exposed. Therefore, unless parties refresh their individual two-party channels consciously (by sending each other messages), executing updates or removals in the group session will not have the desired healing effect.

In the real world, usually not all pairs of members of a group exchange private messages regularly, hence not refreshing their two-party channels. The fact that even manually triggering a

Chapter 6. WhatsUpp with Sender Keys? Analysis, Improvements and Security Proofs

key update does not necessarily heal the group from a state compromise conveys an important security limitation.

Lack of Forward Security on Authentication. Beyond PCS limitations, we observe that the forward security guarantees for authentication provided by Sender Keys are sub-optimal. Consider a simple group $G = \{ID_1, ID_2\}$ and the attack described in Figure 6.14.

- 1: $q_1 = \text{SEND}(ID_1, m)$ generates ciphertext C encrypted under mk and signed under ssk_1 .
- 2: $q_2 = \text{EXPOSE}(ID_1)$, where \mathcal{A} obtains ssk_1 , but not mk .
- 3: \mathcal{A} modifies C and signs it again under ssk_1 to create a forgery C' corresponding to the same message epoch as C .
- 4: $q_3 = \text{RECEIVE}(ID_2, ID_1, C')$, which is a successful injection.

Figure 6.14: Attack on authentication forward security in Sender Keys.

Note that q_3 is a forbidden query by C_{sk-inj} . q_3 attempts to inject a message that corresponds to key material used *before* the state exposure, hence one can envision stronger FS where queries like q_3 are allowed. This attack can occur naturally if ID_2 is offline when m is first sent by ID_1 .

An attack of a similar nature can also occur in a messaging scheme where the same signature keys are re-used across groups, and are refreshed at different times, as pointed out by Cremers et al. [CHK21].

Additional Remarks. In a Sender Keys group, each user is associated with a distinct symmetric key, resulting in a state that contains $\mathcal{O}(n)$ secret material at all times. However, the exposure of a single member compromises the keys of all group members, rendering this use of multiple keys ineffective for enhancing security. The primary advantage of employing multiple keys is for concurrency reasons. Nevertheless, in large groups, this approach can pose scalability challenges. As a result, it is possible to explore trade-offs between the level of concurrency supported and the amount of secret material that needs to be stored at any given time. At the opposite end of the spectrum, all users could maintain a single symmetric chain that is the *same* for everyone. This approach, suitable for scenarios where concurrent message transmission is unlikely, reduces the state size to $\mathcal{O}(1)$ and requires $\mathcal{O}(n)$ PCS updates. Moreover, this would improve security by reducing PCS updates to the sending and processing of a single constant-sized message, which represents an optimal solution.

Sender Keys, as described in Section 6.2.1, is susceptible to randomness exposure and randomness manipulation attacks. Namely, the adversary does not need to leak a member's state, but simply control the randomness used by the device, inhibiting any form of PCS. Protection against this family of attacks can be attained at small cost if freshly generated keys are hashed with the state [JS18] and the classic NAXOS trick used in authenticated key exchange [LLM07].

Other attacks are possible but unavoidable unless symmetric encryption for application

messages is replaced by some form of public-key encryption. Suppose that the adversary \mathcal{A} exposes user ID ($\gamma \leftarrow \text{EXPOSE}(ID)$) and then calls $C \leftarrow \text{CHAL}(ID, m_0, m_1)$. \mathcal{A} can trivially win since it can derive mk used in the CHAL query trivially from γ . Some ratcheting protocols provide strong security in that, if a user is impersonated towards, their state should ‘diverge’ and no longer be useful for decrypting messages from honest parties [PR18, JS18, BRV20]. This is not possible to achieve in a Sender Keys-like protocol (nor is achieved by the Double Ratchet or MLS for messages in the same symmetric ratchet) since the key schedule for message encryption is deterministic and independent of previously received messages.

Although our formalism does not capture multi-group security, it can be adapted to capture it, for example by using group identifiers to label different groups. In principle, reusing the same two-party channels between parties that are in several groups together could lead to an *increase* in security by faster healing (since more messages would be sent on the same channel).¹⁰ As argued by Cremers et al. [CHK21], if a party reuses the same signature key across groups (ignoring any privacy concern from doing so), PCS authenticity guarantees would improve, since replacing the key after channels are healed would immediately heal authenticity across all groups (rather than just one). We leave it open to explore cross-group security further and in particular capture multi-group security formally.

6.5.2 Proposed Improvements: Sender Keys+

We propose several improvements to Sender Keys below. Our improvements are constrained by the desire to retain the performance characteristics and structure of Sender Keys. In particular, we retain $\mathcal{O}(1)$ -sized ciphertexts, do not increase key sizes, and require only standard cryptographic primitives. Our improved version of Sender Keys, which we call Sender Keys+, is presented fully in Figures 6.19 to 6.22. We formalise security by introducing several modifications to our cleanness predicate that we describe at the end of this section.

Secure Control Messages. A simple way of resolving the attacks in Section 6.5.1 would be for users to sign their own control messages and verify signatures before processing control messages. Additional protocol logic for correctness is required, namely that users who craft a control message but have not shared their sender key yet (because they have not spoken in the group) generate a signature key pair and share their public key over the two-party channels.

By introducing this tweak, we can weaken the cleanness predicate such that it no longer enforces honest control message delivery ($C_{\text{sk}^+ \text{-con}}$). On the other hand, we need to introduce the restriction that no secret signature key ssk can be known to the adversary at delivery time, similar as in the injection predicate. We do so by introducing a new control predicate $C_{\text{sk}^+ \text{-ctr}}^\Delta$ that follows the blueprints of the injection predicate (Figure 6.11), and that we show in Figure 6.15.

¹⁰Note this would not reduce security upon state exposure since all of a given party’s state, and therefore channels, are assumed (in this thesis) to be compromised at the same time.

Chapter 6. WhatsUpp with Sender Keys? Analysis, Improvements and Security Proofs

Improved Forward-Secure Authentication. We propose two possible improvements that address the attack in Section 6.5.1 (Figure 6.14) to varying extents.

MACing from the Chain Key. The first improvement, which has minimal overhead, is to MAC the application messages with a MAC key τk that we derive via an additional $H_3(\text{ck})$. The modification is done in the Send algorithm as follows: given an unsigned ciphertext $\tilde{C} = (\text{ct}, (e, i), \text{ME})$, we obtain the MAC tag $\tau \leftarrow \text{MAC.Tag}(\tau k, \tilde{C})$. Then, we sign the ciphertext with the appended tag $\sigma \leftarrow \text{Sig.Sign}(\text{ssk}, (\tilde{C}, \tau))$. The verification of the MAC tag is easily carried out at the receiver's end. We include this simple tweak in the protocol in Figures 6.19 to 6.22. Naturally, symmetric encryption could instead be replaced with an AEAD to achieve the same effect.

The main security improvement that results from this upgrade is that, in the attack in Section 6.5.1, the adversary additionally needs knowledge of τk to forge the MAC tag. Hence, one of the following situations must occur before delivery:

- The sender ID is exposed before the message is sent. Then, both τk and ssk are compromised.
- The sender ID is exposed after the message is sent (leaking ssk), and another group member ID' is exposed before the message is delivered (leaking τk).

In particular, the attack of Figure 6.14 no longer results in a successful message delivery. The MAC key can be stored together with the message key for out-of-order messages, such that the MAC can always be verified in a correct execution of the protocol. We note that insider attacks (forgeries from other group members) cannot be prevented by MACing, but we do not model these.

The modified injection predicate that results from this improvement is shown in Figure 6.17. Essentially, we define an auxiliary predicate $C_{\text{sk}^+ \text{-inj-extra}}^\Delta$ that considers the security given by the message/MAC keys (similarly as in Figure 6.10). Then, the modified $C_{\text{sk}^+ \text{-inj}}$ is the logical disjunction of the former injection predicate with $C_{\text{sk}^+ \text{-inj-extra}}^\Delta$, and is hence strictly weaker.

Ratcheting Signature Keys. An alternative mitigation strategy for the attack of Figure 6.14 is to ratchet signature keys. Let (ssk, spk) be ID 's signature key pair, where spk is part of its sender key. Before sending a new message m to the group, ID can generate a new key pair $(\text{ssk}', \text{spk}') \stackrel{s}{\leftarrow} \text{Gen}(1^\lambda)$. Then, ID can attach the new spk' to the ciphertext corresponding to encrypting m , and sign the package using ssk . This (by now standard) countermeasure not only provides strong forward security but also post-compromise security for the authentication of messages. Nevertheless, it involves larger overhead, so it may not be desirable in all scenarios and we refrain from including it in Sender Keys+.

Efficient PCS Updates. We propose an asynchronous update mechanism to refresh all chain keys at once, recovering PCS on-demand for the whole group with a single update (and $\mathcal{O}(n)$

complexity for a group of n users). Recall that our Group Messenger primitive supports updates via $\text{Exec}(\gamma, \text{upd}, \{ID\})$.

A Naive Solution. Let ID be the updating party. ID generates a new sender key for himself as in the case of a remove operation; namely samples a fresh ck and a fresh $(\text{ssk}, \text{spk}) \xleftarrow{\$} \text{Sig.Gen}(1^\lambda)$. Additionally, ID samples randomness $r \xleftarrow{\$} \{0, 1\}^\lambda$. Then, it distributes $(\text{ck}, \text{spk}, r)$ over the two-party channels. Upon reception, every group member (including ID itself) sets $\text{SK}[ID] \leftarrow (\text{ck}, \text{spk})$; and then for every $ID' \in \mathbf{G}$, set $\text{SK}[ID'].ck \leftarrow H'(\text{SK}[ID'].ck, r)$, where $H' : \{0, 1\}^\lambda \times \{0, 1\}^\lambda \rightarrow \{0, 1\}^\lambda$ is a secure key derivation function. Since r is freshly sampled and distributed securely, all chain keys recover from exposure. Note that r must be used and erased immediately, as all updated chain keys are exposed if r is leaked at any future time.

Our Solution. The previous solution fails in out-of-sync scenarios such as the following. Suppose that ID' is in message epoch $(1, 1)$ when ID sends an update message T . Then, ID' speaks in the group before receiving T (for example, while being offline), ratcheting its key to $(1, 2)$. All group members will update the chain key $\text{ck}_{ID'}^{1,1}$ (i.e. corresponding to the message epoch $(e, i) = (1, 1)$) in $\text{SK}[ID']$, but ID' will be in message epoch $(1, 2)$ (and therefore will have erased $\text{ck}_{ID'}^{1,1}$). In general, if there are application messages in transit concurrently with the update, users will be out-of-sync.

To support asynchronicity, we propose that all parties ratchet their chain key N times forward, where N is a fixed constant that we call the *concurrency bound* (for example $N = 100$; in practice the cost of executing 100 hash function calls sequentially is negligible). In the event that ℓ messages have been sent out-of-sync, then the chain key is ratcheted $N - \ell$ times instead. Then, parties update the ratcheted chain keys with the sent randomness r . To synchronise between them and the update initiator ID , the latter sends a list with his view of the key indices of each group member (in the control message). We describe the update protocol as part of Figures 6.19 to 6.22. Note that this mechanism requires the assumption of total ordering of control messages to avoid overlapping updates.

The security improvement is reflected in the challenge cleanness predicate in Figure 6.16. The predicate follows the blueprint of the challenge predicate for Sender Keys (Figure 6.10), except that now it also suffices that *some* arbitrary member ID^* that has a healed channel with ID' updates after the exposure, and that ID processes such update before the challenge.

Efficient Remove Operations. The previous update mechanism can be extended to improve the efficiency of group removals from $\mathcal{O}(n^2)$ (everyone needs to generate and distribute a new key) to $\mathcal{O}(n)$ in terms of communication complexity. Note a removal can be made effective if the party that sends the remove message T distributes update material among all group members except for the removed party ID' . If ID' leaves, the next member that speaks in the group must also trigger an update. This tweak, like our solution above, has the drawback that the signature keys are not refreshed. To tame the written complexity of the protocol, we do not include this tweak in Sender Keys+. Furthermore, considering the minimal overhead of updates, they could potentially become the preferred method for sharing sender keys in the

Chapter 6. WhatsUpp with Sender Keys? Analysis, Improvements and Security Proofs

group under all circumstances (we also do not use this tweak when a user samples their first sender key at group creation time or upon being added to a group). This approach allows the group to achieve PCS almost for free.

Cleanness Predicates for Sender Keys+. The cleanness predicates corresponding to our improvements and described informally above are detailed in Figures 6.15 to 6.17, and the joint predicate $C_{sk^+}^\Delta$ in Figure 6.18.

$$\begin{array}{|l}
 \hline
 C_{sk^+-con} : \forall (i, ID) : q_i = \text{DELIVER}(ID, T), \exists j < i : \\
 q_j = (\text{ADD or REMOVE or UPDATE or CREATE})(ID, \cdot) \wedge \\
 \exists e' : E[ID; q_i] = e' = E[ID; q_j] \\
 \hline
 C_{sk^+-ctr-core}^\Delta(T_{core}) : \forall (i, j, ID, ID') : \\
 (T_{core}, \cdot) \notin T[\cdot] \wedge (i < j) \wedge q_i = \text{EXPOSE}(ID') \wedge q_j = \text{DELIVER}(ID, (T_{core}, \cdot)), \\
 \exists e' : [(E[ID'; q_i] < e' \leq e^* \wedge \text{refresh}_\Delta(ID', ID, q_i, e')) \vee (e^* < e' \leq E[ID'; q_i] \wedge \text{refresh-s}(ID', e'))] \\
 \hline
 C_{sk^+-ctr}^\Delta : \forall T_{core}, C_{2pc} : C_{sk^+-ctr-core}^\Delta(T_{core}) \wedge C_{2pc-inj}(C_{2pc}) \\
 \hline
 \end{array}$$

Figure 6.15: Modified concurrency predicate, additional auxiliary core control predicate, and additional control predicate for Sender Keys+.

$$\begin{array}{|l}
 \hline
 C_{sk^+-chall}^\Delta : C_{sk^+-chall}^\Delta \vee [i < j \wedge \exists e, e', k, ID^* : E[ID'; q_i] < e' < e \leq E[ID; q_j] \wedge \\
 q_k = \text{UPDATE}(ID^*) \wedge \text{ep}[q_k] = e - 1 \wedge \text{refresh}_\Delta(ID^*, ID', e', q_k)] \\
 \hline
 \end{array}$$

Figure 6.16: Modified challenge cleanness predicate for Sender Keys+.

$$\begin{array}{|l}
 \hline
 C_{sk^+-inj-extra}^\Delta(C_{core}) : C_{sk^+-inj-core}^\Delta(C_{core}) \\
 \vee [i > j \wedge m\text{-ep}(ID', ID, q_i) > (E[ID; q_j], \lfloor ID; q_j \rfloor) \wedge (E[ID; q_j], \lfloor ID; q_j \rfloor) \notin \text{SM}[ID', ID; q_i]] \\
 \vee [i < j \wedge \exists e : E[ID'; q_i] < e \leq E[ID; q_j] \wedge \text{refresh}_\Delta(ID, ID', q_i, e)] \\
 \vee [i < j \wedge \exists e, e', k, ID^* : E[ID'; q_i] < e' < e \leq E[ID; q_j] \wedge \\
 q_k = \text{UPDATE}(ID^*) \wedge \text{ep}[q_k] = e - 1 \wedge \text{refresh}_\Delta(ID^*, ID', e', q_k)] \\
 \hline
 C_{sk^+-inj}^\Delta : \forall C_{core}, C_{2pc} : C_{sk^+-inj-extra}^\Delta(C_{core}) \wedge C_{2pc-inj}(C_{2pc}) \\
 \hline
 \end{array}$$

Figure 6.17: Modified auxiliary injection predicate and injection predicate for Sender Keys+. We remark that the additional logic simply mimics the structure of $C_{sk^+-chall}^\Delta$.

$$C_{sk^+}^\Delta : C_{sk^+-chall}^\Delta \wedge C_{sk^+-inj}^\Delta \wedge C_{sk^+-con}^\Delta \wedge C_{sk^+-ctr}^\Delta$$

Figure 6.18: Modified cleanness predicate for Sender Keys+.

6.5.3 Sender Keys/Sender Keys+ vs CGKA

As remarked in the introduction, Sender Keys (and especially Sender Keys+) offers different efficiency and security trade-offs over CGKA-based protocols. We provide a detailed comparison below.

PCS. When a user ID is exposed, the confidentiality of all subsequent messages is lost in both CGKA and Sender Keys(+). For an update to take effect in Sender Keys(+), all two-party channels must have healed. In this case, an update by ID' only heals the confidentiality of messages sent by ID' in Sender Keys, as opposed to the confidentiality of messages sent by *all* members in Sender Keys+.

It is worth noting that both Sender Keys and Sender Keys+ require up to PCS bound Δ messages (or rounds) to heal after a compromise (due to the two-party channels) in addition to the update message. In contrast, a single message suffices for some CGKA protocols [ACDT20, KPPW⁺21, ACJM20, AHKM22].

Update Efficiency. In Sender Keys+, an update message requires $O(n)$ communication by the updating user, where each member is sent a constant-size message. In TreeKEM variants, or in general binary-tree-based CGKAs, updates involve best-case $O(\log n)$ size for the updating user and have to be entirely downloaded by each member, involving a total $O(n \log n)$ download overhead. Nevertheless, this can degrade to $O(n)$ per member (i.e., $O(n^2)$ total). The multi-recipient PKE approach taken by Hashimoto et al. [HKP⁺21] achieves the same asymptotic complexity as Sender Keys+, although with larger concrete costs.

Insider Security. The attack of Alwen et al. [AJM22] that highlights the need for IND-CCA (and not only IND-CPA) public-key encryption in TreeKEM also applies to symmetric encryption in Sender Keys, but can be mitigated by using a MAC (appropriately) or an AEAD. Considering the analysis of Alwen et al. [AJM22], it is not clear how to mount fake group attacks as they do, although if different users process different control messages, they may end up with different views of the group. This attack however also applies to CGKAs in general.

Separately, we note that Sender Keys(+) does not suffer from the forward security issues from MLS's CGKA [ACDT20].

6.5.4 Sender Keys in Practice

We discuss different aspects of how our formalism compares to the implementation of Sender Keys by WhatsApp and Signal in practice.

AEAD for encryption. Both WhatsApp [Wha20] and Signal¹¹ use AES-256 in CBC mode for Sender Keys, which does not provide authentication guarantees.

Control messages in WhatsApp. Add and remove operations are processed following [RMS18, Sec. 5.2.2, (p. 10)] via 'modification messages' which contain similar information to our control messages. To our knowledge, dedicated sender key updates are not supported by WhatsApp, as the feature is not mentioned in their white paper [Wha20].

¹¹https://github.com/signalapp/libsignal/blob/3b7f3173cc4431bc4c6e55f6182a37229b2db6fd/rust/protocol/src/group_cipher.rs#L43C29-L43C29

Chapter 6. WhatsUpp with Sender Keys? Analysis, Improvements and Security Proofs

Signal group management. At least on Android, Signal has since 2021 required the use of so-called v2 groups, which encompasses Signal’s private group management system [CPZ20].¹² Here, Sender Keys is used whenever possible, although occasionally pairwise channels are used as a fallback.¹³ The state of the group membership is dictated by the central server; effectively, adds and removes are totally ordered by it. Our control message abstraction captures this latter fact, but our formalism does not capture the private group system. However, our signing of control messages in Sender Keys+ provides more guarantees on group membership given total ordering on control messages: a malicious Signal server can, e.g., re-add removed users without authorisation from group members [CPZ20, p. 45/46].

Updates in Signal. Every sender key must be updated by default every two weeks and a global maximum of 90 days¹⁴; these are sent via two-party channels, and there is no central control message sent. In this chapter, we opted to model update control messages following the CGKA line of work, for clarity, and because they are particularly well-suited to our improved update mechanism from Section 6.5.2.

6.6 Sender Keys and Sender Keys+: Full Protocols and Security

6.6.1 Protocol Descriptions

We introduce the full Sender Keys and Sender Keys+ protocols in Figures 6.19 to 6.22, extending the descriptions in Section 6.4 and Section 6.5. In Tables A.1 and A.2 (Appendix A.4) we describe the variables and dictionaries used here and in this chapter in general which may be especially useful in this subsection.

Below, we make some additional remarks intended to help the reader to parse the pseudocode which is inherently complex, not least due to the additional variables and logic that is required by our modifications.

Sent and Unsent Sender Keys. When a new user ID joins and a member ME processes the message via $\text{Proc}(\cdot, T = (\text{add}, \cdot))$, ID does not receive the sender key of ME until ID speaks again. Hence, ME needs to keep track of this newly added user; it does so via the $\text{no-SK}[\cdot]$ dictionary. Namely, $\text{no-SK}[ID] = \text{true}$ in the view of ME if ME has not sent his sender key to ID yet. This functionality is captured in the SendToMissing algorithm.

Sending Control Messages without a Sender Key. A different scenario is that ME calls Exec and generates a control message T which, in Sender Keys+, needs to be signed. In the event that ME does not have a working sender key yet (e.g., due to a recent removal), ME gener-

¹²<https://github.com/signalapp/Signal-Android/blob/0775fc7ead818a8380e7e374d15898cbabccaa9c/app/src/main/java/org/thoughtcrime/securesms/jobs/PushGroupSendJob.java#L325C76-L325C76>

¹³See the ‘for’ loop at <https://github.com/signalapp/Signal-Android/blob/0775fc7ead818a8380e7e374d15898cbabccaa9c/app/src/main/java/org/thoughtcrime/securesms/messages/GroupSendUtil.java#L249C1-L249C52>

¹⁴<https://github.com/signalapp/Signal-Android/commit/35393fc33165e5b1417e7b1a7d6f85d0d7919c6f>

6.6 Sender Keys and Sender Keys+: Full Protocols and Security

<p>Init($ID, 1^\lambda$)</p> <pre> 1: $\gamma.ME \leftarrow ID$ 2: $\gamma.(ssk, G, ep, i_{ME}, last-kc) \leftarrow \perp$ 3: $\gamma.SK[\cdot, \cdot] \leftarrow \perp$ // $(ID, ck) \rightarrow (spk_{ID}, ck_{ID}, i_{ck})$ 4: $\gamma.MK[\cdot, \cdot] \leftarrow \perp$ // $(ID, (kc, i_{ck})) \rightarrow (mk, \tau k)$ 5: $\gamma.no-SK[\cdot] \leftarrow false$ // $ID \rightarrow boolean$ 6: $\gamma.max-i_{ck}[\cdot, \cdot] \leftarrow \perp$ // $kc \rightarrow i_{ck}$ 7: $\gamma.kc[\cdot] \leftarrow \perp$ // $ID \rightarrow kc$ 8: $\gamma.rs[\cdot] \leftarrow \perp$ // $kc \rightarrow random\ coin$ 9: $\gamma.\gamma_{2pc} \leftarrow 2PC.Init(ID)$ 10: return γ </pre> <p>Send(γ, m)</p> <pre> 1: require $ME \in G$ 2: $\tilde{C}[\cdot] \leftarrow \perp$ 3: if $SK[ME, kc[ME]].i_{ck} = \perp$: // Sample sender key if needed 4: $\tilde{C} \leftarrow PreSendFirst()$ 5: if $i_{ME} = 0$: 6: $\tilde{C} \leftarrow (\tilde{C}, SendToMissing())$ 7: $(mk, \tau k) \leftarrow (H_1, H_3)(SK[ME, kc[ME]].ck)$ 8: $ct \xrightarrow{\\$} SymEnc.Enc(mk, m)$ 9: $UpdateCK(ME, kc[ME])$ 10: $M \leftarrow (ct, (ep, i_{ME}), kc[ME], i_{ck}, ME)$ 11: $\tau \leftarrow MAC.Tag(\tau k, M)$ 12: $\sigma \xrightarrow{\\$} Sig.Sign(ssk, M, \tau)$ 13: return $C := (M, \tau, \sigma, \tilde{C})$ </pre> <p>UpdateCK(ID, kc')</p> <pre> 1: $SK[ID, kc'].ck \leftarrow H_2(SK[ID, kc'].ck)$ 2: $SK[ID, kc'].i_{ck} \leftarrow SK[ID, kc'].i_{ck} + 1$ 3: if $ID = ME$: 4: $i_{ME} \leftarrow i_{ME} + 1$ </pre>	<p>Recv($\gamma, C = ((M, \tau, \sigma), C_{2pc})$)</p> <pre> 1: parse M as $(ct, (e, i), kc', i_{ck}', ID)$ 2: require $ID \in G$ 3: if $SK[ID, kc'] = \perp$: // Receive ID's sender key via 2PC if needed 4: $((SK[ID, kc'], kc^*, ep', max-i_{ck}', last-kc'), ID^*, \cdot, \cdot) \leftarrow 2PC.Recv(C_{2pc})$ 5: require $ID^* = ID \wedge ep' = e \wedge kc^* = kc$ 6: $DeleteOldCK(ID, max-i_{ck}', last-kc')$ 7: else require $C_{2pc} = \perp$ 8: require $e \leq ep$ 9: require $Sig.Vrfy(SK[ID, kc'].spk, \sigma, M, \tau)$ 10: $(mk, \tau k) \leftarrow UpdateKeysRecv()$ 11: require $MAC.Vrfy(\tau k, M, \tau)$ 12: $m \leftarrow SymEnc.Dec(mk, ct)$ 13: return (m, ID, e, i) </pre> <p>UpdateKeysRecv()</p> <pre> 1: $i_{ck} \leftarrow SK[ID, kc']$ // Store skipped keys in MK given out-of-order delivery 2: while $i_{ck} < i_{ck}'$: 3: $(mk, \tau k) \leftarrow (H_1, H_3)(SK[ID, kc'].ck)$ 4: $MK[ID, (kc', i_{ck}')] \leftarrow (mk, \tau k)$ 5: $UpdateCK(ID, kc')$ 6: $i_{ck}'++$ 7: if $i_{ck} > i_{ck}'$: 8: require $MK[ID, (kc', i_{ck}')] \neq \perp$ 9: $(mk, \tau k) \leftarrow MK[ID, (kc', i_{ck}')]$ // Delete stored key for forward security 10: $MK[ID, (kc', i_{ck}')] \leftarrow \perp$ 11: else : 12: $(mk, \tau k) \leftarrow (H_1, H_3)(SK[ID', kc'].ck)$ 13: $UpdateCK(ID, kc')$ 14: return $(mk, \tau k)$ </pre>
-------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------	----------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------------

Figure 6.19: Sender Keys and Sender Keys+ description (part 1 of 4: Init, Send/Recv, two helpers). State variables are implicitly stateful where relevant for readability. Text in black colour corresponds to standard Sender Keys. Coloured text corresponds to the modifications in Sender Keys+ from Section 6.5.2: blue text corresponds to securing control messages via signatures, teal text corresponds to MACing for forward security and violet text corresponds to PCS updates.

ates an ephemeral sender key containing only a signature key spk . The key is immediately distributed over the two-party channels via `OneTimeSpk`. We remark that whenever both `SendToMissing` and `OneTimeSpk` are executed in the same algorithm (such as in `Exec`), only

Chapter 6. WhatsUpp with Sender Keys? Analysis, Improvements and Security Proofs

<p>Exec(γ, cmd = rem, ID)</p> <hr/> <pre> 1: require $ID \in G$ 2: $\tilde{C}[\cdot] \leftarrow \perp$ 3: $T \leftarrow (\text{rem}, ME, ID, \text{ep} + 1)$ 4: $\tilde{C} \leftarrow (\text{SendToMissing}(), \text{OneTimeSpk}())$ 5: return ($T, \sigma := \text{Sig.Sign}(\text{ssk}, T), \tilde{C}$) </pre> <p>Proc($\gamma$, ($T = (\text{rem}, ID, ID', \text{ep}'), \sigma, C_{2pc}$))</p> <hr/> <pre> 1: require $ID \in G$ 2: if $\text{SK}[ID, \text{kc}[ID]].\text{spk} = \perp$: 3: $((\text{SK}[ID, \text{kc}[ID]], \text{kc}^*, \text{ep}^*, \text{max-}i_{\text{ck}}', \text{last-}kc')$ $ID^*, \cdot, \cdot) \leftarrow 2\text{PC.Recv}(C_{2pc})$ 4: require $ID^* = ID$ 5: require $\text{ep}' = \text{ep}^* + 1$ 6: require $\text{kc}^* = \text{kc}[ID]$ 7: DeleteOldCK($ID, \text{max-}i_{\text{ck}}', \text{last-}kc'$) 8: else require $C_{2pc} = \perp$ 9: require $\text{Sig.Vrfy}(\text{SK}[ID, \text{kc}[ID]].\text{spk}, \sigma, T)$ 10: require $\text{ep}' = \text{ep} + 1$ 11: $G \leftarrow G \setminus \{ID'\}$ 12: $\text{ep} \leftarrow \text{ep} + 1$; $i_{ME} \leftarrow 0$ 13: for all $ID^+ \in G$: 14: $\text{kc}[ID^+] \leftarrow \text{kc}[ID^+] + 1$ 15: $\text{SK}[ID', \cdot] \leftarrow \perp$ 16: if $ID = ME$: 17: $\gamma \leftarrow \perp$ 18: return true </pre>	<p>Proc(γ, ($T = (\text{add}, ID, ID' \neq ME, \text{ep}'), \sigma, C_{2pc}$))</p> <hr/> <pre> 1: require $ID \in G$ 2: if $\text{SK}[ID, \text{kc}[ID]].\text{spk} = \perp$: 3: $((\text{SK}[ID, \text{kc}[ID]], \text{kc}^*, \text{ep}^*, \text{max-}i_{\text{ck}}', \text{last-}kc')$ $ID^*, \cdot, \cdot) \leftarrow 2\text{PC.Recv}(C_{2pc})$ 4: require $ID^* = ID$ 5: require $\text{ep}' = \text{ep}^* + 1$ 6: require $\text{kc}^* = \text{kc}[ID]$ 7: DeleteOldCK($ID, \text{max-}i_{\text{ck}}', \text{last-}kc'$) 8: else require $C_{2pc} = \perp$ 9: require $\text{Sig.Vrfy}(\text{SK}[ID].\text{spk}, \sigma, T)$ 10: require $\text{ep}' = \text{ep} + 1$ 11: $G \leftarrow G \cup \{ID'\}$ 12: $\text{ep} \leftarrow \text{ep} + 1$; $i_{ME} \leftarrow 0$ 13: $\text{acc} \leftarrow 2\text{PC.InitCh}(ID')$ 14: require acc 15: $\text{no-SK}[ID'] \leftarrow \text{true}$ 16: $\text{kc}[ID'] \leftarrow 0$ 17: return true </pre> <p>Proc(γ, ($T = (\text{add}, ID, ME, \text{ep}'), \sigma, C_{2pc}$))</p> <hr/> <pre> 1: $((G', \text{kc}', \text{ep}^*, \text{spk}), ID^*, \cdot, \cdot) \leftarrow 2\text{PC.Recv}(C_{2pc})$ 2: require $ID^* = ID$ 3: $\text{SK}[ID, \text{kc}'[ID]].\text{spk} \leftarrow \text{spk}$ 4: require $\text{Sig.Vrfy}(\text{SK}[ID, \text{kc}'[ID]].\text{spk}, \sigma, T)$ 5: require $\text{ep}^* + 1 = \text{ep}'$ 6: $(G, \text{kc}, \text{ep}) \leftarrow (G', \text{kc}', \text{ep}')$ 7: for all $ID' \in G \setminus \{ME\}$: 8: $\text{acc} \leftarrow 2\text{PC.InitCh}(ID')$ 9: require acc 10: $\text{kc}[ME], i_{ME} \leftarrow 0$ 11: return true </pre>
<p>Exec(γ, cmd = add, ID)</p> <hr/> <pre> 1: require $ID \notin G$ 2: $\tilde{C}[\cdot] \leftarrow \perp$ 3: $T \leftarrow (\text{add}, ME, ID, \text{ep} + 1)$ 4: $\tilde{C} \leftarrow (\text{SendToMissing}(), \text{OneTimeSpk}())$ 5: $\text{welcome} \leftarrow (G, \text{kc}, \text{ep}, \text{spk})$ 6: $(\tilde{C}[ID], \cdot, \cdot) \stackrel{\\$}{\leftarrow} 2\text{PC.Send}(ID, \text{welcome})$ 7: return ($T, \sigma := \text{Sig.Sign}(\text{ssk}, T), \tilde{C}$) </pre>	<p>Exec(γ, cmd = crt, IDs)</p> <hr/> <pre> 1: $G \leftarrow IDs$; $\tilde{C}[\cdot] \leftarrow \perp$ 2: $T \leftarrow (\text{crt}, ME, IDs)$ 3: $\tilde{C} \leftarrow \text{OneTimeSpk}()$ 4: return ($T, \sigma := \text{Sig.Sign}(\text{ssk}, T), \tilde{C}$) </pre>

Figure 6.20: Sender Keys and Sender Keys+ description (part 2 of 4: Proc/Exec part 1 of 2).

6.6 Sender Keys and Sender Keys+: Full Protocols and Security

Proc($\gamma, (T = (\text{crt}, ID, IDs), \sigma, C_{2pc})$)	Proc($\gamma, (T = (\text{upd}, ID, kc', ep', \text{Upd-Ind}), \sigma, C_{2pc})$)
<pre> 1: if SK[ID, kc[ID]].spk = \perp : 2: ((SK[ID, kc[ID]], kc*, ep*, ·, ·), ID*, ·, ·) ← 2PC.Recv(C_{2pc}) 3: require ID* = ID 4: require ep* = kc* = \perp 5: else require $C_{2pc} = \perp$ 6: require Sig.Vrfy(SK[ID, kc[ID]].spk, σ, T) 7: require ep = \perp 8: $G \leftarrow IDs$ 9: for all ID' $\in G$: 10: acc ← 2PC.InitCh(ID') 11: require acc 12: kc[ID'] ← 0 13: ep, i_{ME} ← 0 14: return true </pre>	<pre> 1: require Sig.Vrfy(SK[ME].spk, σ, T) 2: if ID = ME : 3: require $C_{2pc} = \perp$ 4: $r \leftarrow rs[kc']$; $rs[kc'] \leftarrow \perp$ 5: else : 6: ((SK[ID, kc[ID] + 1], kc*, ep*, max-i_{ck}', last-kc', r), ID*, ·, ·) ← 2PC.Recv(C_{2pc}) 7: require ID* = ID 8: require ep' = ep* + 1 9: require kc* = kc[ID] + 1 10: require ep' = ep + 1 11: ep ← ep'; i_{ME} ← 0 12: kc[ID] ← kc[ID] + 1 13: DeleteOldCK(ID, max-i_{ck}', last-kc') // Hash forward N times before hashing with r 14: for all ID' $\in G \setminus \{ID\}$: 15: if SK[ID', kc[ID']].ck = \perp : 16: continue (line 14) 17: while SK[ID', kc[ID']].i_{ck} < Upd-Ind[ID'] : 18: (mk, τk) ← (H₁, H₃)(SK[ID', kc[ID']].ck) 19: MK[ID', (kc[ID'], SK[ID', kc[ID']].i_{ck})] ← (mk, τk) 20: UpdateCK(ID', kc[ID']) 21: $\ell \leftarrow SK[ID', kc[ID']].i_{ck} - \text{Upd-Ind}[ID']$ 22: require $\ell < N$ 23: ck' ← SK[ID', kc[ID']].ck 24: do N - ℓ times : ck' ← H₂(ck') 25: SK[ID', kc[ID']].ck ← F(ck', r) 26: return true </pre>
<pre> Exec($\gamma, \text{cmd} = \text{upd}, ME$) 1: require ME $\in G$ 2: Upd-Ind[·] ← \perp 3: (spk, ssk) $\stackrel{\\$}{\leftarrow}$ Sig.Gen 4: $r \stackrel{\\$}{\leftarrow} \{0, 1\}^\lambda$ 5: Upd-Ind[ME] ← SK[ME, kc[ME]].i_{ck} 6: ck $\stackrel{\\$}{\leftarrow} \{0, 1\}^\lambda$ 7: SK[ME, kc[ME] + 1] ← (spk, ck, i_{ck}) 8: $m \leftarrow (\text{SK}[ME, kc[ME] + 1], kc[ME] + 1, ep,$ $\text{max-i}_{ck}[\text{last-kc}], \text{last-kc}, r)$ 9: $\vec{C}[\cdot] \leftarrow \perp$ 10: for all ID $\in G \setminus \{ME\}$: 11: ($\vec{C}[ID], \cdot, \cdot$) $\stackrel{\\$}{\leftarrow}$ 2PC.Send(ID, m) 12: Upd-Ind[ID] ← SK[ID', kc[ID]].i_{ck} 13: $rs[kc[ME]] \leftarrow r$ 14: $T \leftarrow (\text{upd}, ME, kc[ME] + 1, ep + 1, \text{Upd-Ind})$ 15: return (T, $\sigma := \text{Sig.Sign}(ssk, T), \vec{C}$) </pre>	

Figure 6.21: Sender Keys and Sender Keys+ description (part 3 of 4: Proc/Exec part 2 of 2).

one of them will output a non-blank ciphertext, depending on whether the caller's signature key $\text{SK}[ME, kc[ME]].\text{spk}$ exists or not.

Index Updates. Most of the protocol logic behind our new update mechanism is explained in Section 6.5. Due to the synchronisation issues mentioned there, the update initiator sends his view (message epoch) of everyone else's sender key. This information is stored in the $\text{Upd-Ind}[\cdot]$ dictionary, which is sent as part of the control message.

<pre> SendToMissing() // Send my sender key to new parties via 2PC 1: if SK[ME, kc[ME]].spk = ⊥: return 2: m ← (SK[ME, kc[ME]], kc[ME] ep, max-ick[last-kc], last-kc) 3: C̃[·] ← ⊥ 4: for all ID ∈ G \ {ME}: 5: if no-SK[ID]: 6: (C̃[ID], ·, ·) $\stackrel{\\$}{\leftarrow}$ 2PC.Send(ID, m) 7: no-SK[ID] ← false 8: return C̃ </pre>	<pre> PreSendFirst() // Sample and send new sender key to all via 2PC 1: (spk, ssk) $\stackrel{\\$}{\leftarrow}$ Sig.Gen 2: ck $\stackrel{\\$}{\leftarrow}$ {0, 1}^λ 3: max-ick[last-kc] ← SK[ME, last-kc].ick 4: SK[ME, kc[ME]] ← (ck, spk, 0) 5: m ← (SK[ME, kc[ME]], kc[ME], ep, max-ick[last-kc], last-kc) 6: C̃[·] ← ⊥ 7: for all ID' ∈ G \ {ME}: 8: (C̃[ID'], ·, ·) $\stackrel{\\$}{\leftarrow}$ 2PC.Send(ID', m) 9: no-SK[ID'] ← false 10: last-kc ← kc[ME] 11: return C̃ </pre>
<pre> OneTimeSpk() // Sample and send spk to all via 2PC 1: C̃[·] ← ⊥ 2: if SK[ME, kc[ME]].spk = ⊥: 3: (spk, ssk) $\stackrel{\\$}{\leftarrow}$ Sig.Gen 4: SK[ME, kc[ME]] ← (spk, ⊥, ⊥) 5: m ← (SK[ME, kc[ME]], kc[ME], ep, max-ick[last-kc], last-kc) 6: for all ID ∈ G \ {ME}: 7: (C̃[ID], ·, ·) $\stackrel{\\$}{\leftarrow}$ 2PC.Send(ID, m) 8: last-kc ← kc[ME] 9: return C̃ </pre>	<pre> DeleteOldCK(ID, max-ick', last-kc') // Store mk's until max-ick' then delete ck 1: if SK[ID, last-kc'].ck = ⊥: 2: return 3: while SK[ID, last-kc'].ick < max-ick': 4: (mk, τk) ← (H₁, H₃)(SK[ID, last-kc'].ck) 5: MK[ID, (last-kc', SK[ID, last-kc'].ick)] ← (mk, τk) 6: UpdateCK(ID, last-kc') 7: SK[ID, last-kc'].ck ← ⊥ </pre>

Figure 6.22: Sender Keys and Sender Keys+ description (part 4 of 4: remaining helpers).

Additional State Variables. The state variables $\text{max-ick}[\cdot]$ and last-kc were omitted in Figure 6.9. Essentially, these variables keep track of the maximum index $\text{max-ick}[ID, kc]$, corresponding to the sender key $\text{SK}[ID, kc]$, for which a message was sent. This is critical to determine what skipped message keys (if any) should be stored in MK so that chain keys can eventually be deleted for forward security. This synchronisation mechanism occurs in the Recv algorithm via the two-party channels, where last-kc specifies the last key counter the bound max-ick refers to.

Correctness. For completeness, we observe here that Sender Keys and Sender Keys+ are correct as in Definition 48. The argument, which we sketch below, proceeds by a similar case analysis to what we did for IAS and DGS in Section 5.3, except here we argue with respect to the relevant winning conditions of Figure 6.6:

- *Message delivery:* By construction of Sender Keys/+ and the correctness of 2PC, for a given message epoch, all active group members will derive the same message key, from which this follows from the correctness of SymEnc (note we also must invoke the

6.6 Sender Keys and Sender Keys+: Full Protocols and Security

correctness of other primitives here and below, e.g., the fact that all Sig.Vrfy checks pass).

- *Group evolution*: This follows from the consistency between how state variable \mathbf{G} , by construction of Exec and Proc , and dictionary \mathbf{G} in Figure 6.6 are updated.
- *Group membership consistency*: This follows similarly to the previous point.
- *Out-of-order delivery*: Messages from past epochs or future indices in the current epoch are handled correctly by construction of UpdateKeysRecv which stores skipped message keys (as well as the correctness of other protocol components like 2PC as above). Ciphertexts from future epochs will always be rejected since the caller of Recv compares their local variable ep with epoch value e attached to each ciphertext output by Send .

6.6.2 Sender Keys Security

We here prove Theorem 20, i.e., the security of our core Sender Keys protocol. Let \mathcal{A} be an adversary against the Sender Keys protocol that plays the GM message indistinguishability game $\text{M-IND}_{\text{GM}, \text{C}_{\text{gm}}}$ (Figure 6.7) with respect to cleanness predicate $\text{C}_{\text{gm}} = \text{C}_{\text{sk}}^\Delta$ (Figure 6.13), and let q_1, \dots, q_q be the oracle queries of \mathcal{A} in a given execution. The proof follows a series of hybrid games, where Game Γ_0 is the original game in Figure 6.7.

Exposed Keys and Key Sequences. Before diving into the details, we characterise the set of exposed chain keys $\text{ExpKeys}_{\text{ck}}$ as those keys that can be (trivially) derived by the adversary following its state exposure queries. To this end, we observe that all chain keys (and also message keys) generated during protocol execution are uniquely identified by three parameters: the epoch number e , the key index i , and the owner ID ; we label them as $\text{ck}_{ID}^{e,i}$ and $\text{mk}_{ID}^{e,i}$. Given a user ID , its chain keys form *sequences* such that the key in message epoch $(e, i+1)$ is deterministically derived from the key in (e, i) as $\text{ck}_{ID}^{e,i+1} = \text{H}_1(\text{ck}_{ID}^{e,i})$ in the chain. Formally, let $q_i = \text{DELIVER}(ID, (T_{\text{core}}, C_{2\text{pc}}))$ where T_{core} is either: a remove message for $ID' \neq ID$, a create message, a message that adds ID to the group, or an update message for ID , and $q_j = \text{DELIVER}(ID, (T'_{\text{core}}, C'_{2\text{pc}}))$ where T'_{core} is either the next remove message for any $ID' \in \mathbf{G}$ for \mathbf{G} from the perspective of ID before q_i where $j > i$, or an update message for ID ; otherwise, $q_j = q_{q+1}$ (where $\text{E}[ID; q_{q+1}]$ denotes the state of $\text{E}[ID]$ after query q_q) if no such query was made. Let also $e = \text{E}[ID; q_i]$ and $\tilde{e} = \text{E}[ID; q_j]$. Then,

$$\text{ck}_{ID}^{e,0}, \dots, \text{ck}_{ID}^{e,i_e}, \text{ck}_{ID}^{e+1,0}, \dots, \text{ck}_{ID}^{\tilde{e}-1,i_{\tilde{e}}}$$

is the chain key sequence for ID in epochs e to \tilde{e} . Note since $\text{ck}_{ID}^{e,0}$ is the first key generated after a removal or update that $\text{ck}_{ID}^{e,0}$ is generated using fresh randomness and distributed to all parties via two-party channels.

For every $q_j = \text{EXPOSE}(ID_j)$ query such that $(e_j, i_j) \leftarrow \text{m-ep}(ID_j, ID', q_j)$, all the chain keys of ID' that are exposed are exactly those in the chain key sequence of ID' starting from epoch

Chapter 6. WhatsUpp with Sender Keys? Analysis, Improvements and Security Proofs

(e_j, i_j) . If we denote this set by $\text{EpCK}_{ID'}^{(j)}$, then we have that

$$\text{ExpKeys}_{\text{ck}} = \bigcup_{j=1}^k \left\{ \text{ck}_{ID'}^{e,i} : ID' \in \mathbf{G}[ID_j; q_{i_j}] \wedge (e, i) \in \text{EpCK}_{ID'}^{(j)} \right\}.$$

where $\mathbf{G}[ID_j; q_{i_j}] = ID_j \cdot \gamma \cdot \mathbf{G}$ represents the view of the group of ID_j at the time of query q_{i_j} . Note that, for any group member, if one of its keys in a chain key sequence is in $\text{ExpKeys}_{\text{ck}}$, all the subsequent keys until either a removal or update for ID_j is processed are also in $\text{ExpKeys}_{\text{ck}}$.

Hybrid Games. We define the main sequence of games below. We then bound the corresponding advantages by a series of lemmas.

Game Γ_0 : This is the original M-IND game, parameterised by cleanness predicate $C_{\text{gm}} = C_{\text{sk}}^\Delta$.

Game Γ_1 : In this game, we remove the **return b** conditions whenever $\text{DELIVER}(ID, (T_{\text{core}}, C_{2\text{pc}}))$ or $\text{RECEIVE}(ID, (C_{\text{core}}, C_{2\text{pc}}))$ are called such that $C_{2\text{pc}}$ was not previously output in some previous oracle query that outputs a ciphertext or control message.

Game Γ_2 : In this game, we remove the **return b** condition in the RECEIVE and DELIVER oracles, so that they always returns nothing. Hence, the adversary can no longer win the game by injecting.

Game Γ_3 : In this game, we allow a single call to the CHAL oracle, as opposed to arbitrarily many calls.

Game Γ_4 : In this game, all chain keys and corresponding message keys in ID 's key sequence that includes ck such that $H_2(\text{ck}) = \text{mk}$, where mk is the message key used in the underlying Send call in the $\text{CHAL}(ID, \cdot, \cdot)$ query (if it exists), are replaced by uniformly random values. In addition, all 2PC ciphertexts that transmit the chain key ck or keys earlier in the key sequence leading to ck are replaced with encryptions of 0^ℓ where ℓ is the length of the message encrypted.

To complete the proof, Game Γ_4 is simulated by an IND-CPA SymEnc adversary. In the lemmas hereafter, we assume that an adversary \mathcal{B} simulating a hybrid for \mathcal{A} can efficiently determine whether \mathcal{A} has violated the cleanness predicate, and aborts execution since \mathcal{B} can no longer win.

Lemma 17. There exists an adversary \mathcal{B}_1 with similar running time to \mathcal{A} such that

$$\text{Adv}_{\text{GM}, C_{\text{gm}}}^{\text{g0}}(\mathcal{A}) \leq \text{Adv}_{\text{GM}, C_{\text{gm}}}^{\text{g1}}(\mathcal{A}) + 2 \cdot \text{Adv}_{2\text{PC}, C_{2\text{pc}}, \Delta}^{2\text{pc-ind}}(\mathcal{B}_1).$$

Proof. We proceed by constructing $\mathcal{B}_{1,b'}$, a $2\text{PC-IND}_{2\text{PC}, b, C_{2\text{pc}}, \Delta}$ adversary that simulates Game Γ_0 /Game Γ_1 for adversary \mathcal{A} depending on $2\text{PC-IND}_{2\text{PC}, b, C_{2\text{pc}}, \Delta}$ bit b and $\text{M-IND}_{\text{GM}, b', C_{\text{gm}}}$ bit b' . At a high level, $\mathcal{B}_{1,b'}$ simulates all $\text{M-IND}_{\text{GM}, C_{\text{gm}}}$ oracle queries using

6.6 Sender Keys and Sender Keys+: Full Protocols and Security

its own oracles and otherwise simulating locally, except every $\text{CHAL}(ID, m_0, m_1)$ is simulated as if the challenger's bit is b' . In more detail, $\mathcal{B}_{1,b'}$ replaces all:

- $2\text{PC.InitCh}(ID')$ calls from ID by the output of query $\text{INIT-CH}(ID, ID')$.
- $2\text{PC.Send}(m, ID')$ calls from ID by the output of $\text{SEND}(ID, ID', m)$.
- $2\text{PC.Recv}(C)$ calls from ID as follows: $\mathcal{B}_{1,b'}$ first calls $\text{RECEIVE}(ID, C)$, which outputs $(b, ID', e_{2\text{pc}}, i_{2\text{pc}})$. If $b \neq \perp$, $\mathcal{B}_{1,b'}$ returns b to its challenger and stops simulating. Otherwise, by cleanness and construction of Sender Keys (argued below), C must have been previously output by a $\text{SEND}(ID, ID', m)$ call. In this case, $\mathcal{B}_{1,b'}$ thus replaces the Recv call with $(m, ID', e_{2\text{pc}}, i_{2\text{pc}})$.

In addition:

- If \mathcal{A} calls $\text{EXPOSE}(ID)$, $\mathcal{B}_{1,b'}$ uses the output of its own $\text{EXPOSE}(ID)$ call and its state from locally simulating to respond to \mathcal{A} 's query.
- Finally, $\mathcal{B}_{1,b'}$ outputs the same bit as \mathcal{A} .

Note that by construction of Sender Keys, after 2PC.Recv is called in Recv and Proc calls, the output is checked so it is “appropriate” for the context it is called in (i.e., it is consistent with the received C_{core} or T_{core} (e.g., by checking ID and epoch matches with the input C)). In addition, in a given Sender Keys Recv/Proc call that does *not* invoke 2PC.Recv , $C_{2\text{pc}} = \perp$ is enforced, preventing GM forgeries that include an arbitrary $C_{2\text{pc}}$ value that is simply ignored. Thus, if $\mathcal{B}_{1,b'}$ outputs (b, \dots) from RECEIVE or DELIVER such that $b \neq \perp$, it must be that a valid 2PC forgery was made.

Using the triangle inequality, we have

$$\begin{aligned} \text{Adv}_{\text{GM}, C_{\text{gm}}}^{\text{g}0}(\mathcal{A}) &= |\Pr[G_0^1 \Rightarrow 1] - \Pr[G_0^0 \Rightarrow 1]| \\ &\leq |\Pr[G_0^1 \Rightarrow 1] - \Pr[G_1^1 \Rightarrow 1]| + |\Pr[G_0^0 \Rightarrow 1] - \Pr[G_1^0 \Rightarrow 1]| \\ &\quad + |\Pr[G_1^1 \Rightarrow 1] - \Pr[G_1^0 \Rightarrow 1]| \\ &\leq \text{Adv}_{2\text{PC}, C_{2\text{pc}}, \Delta}^{2\text{pc-ind}}(\mathcal{B}_{1,1}) + \text{Adv}_{2\text{PC}, C_{2\text{pc}}, \Delta}^{2\text{pc-ind}}(\mathcal{B}_{1,0}) + \text{Adv}_{\text{GM}, C_{\text{gm}}}^{\text{g}1}(\mathcal{A}) \end{aligned}$$

where the inequality $|\Pr[G_0^{b'} \Rightarrow 1] - \Pr[G_1^{b'} \Rightarrow 1]| \leq \text{Adv}_{2\text{PC}, C_{2\text{pc}}, \Delta}^{2\text{pc-ind}}(\mathcal{B}_{1,b'})$ holds because the simulation is perfect given RECEIVE never outputs (b, \dots) with $b \neq \perp$, and when RECEIVE does output such a (b, \dots) , $\mathcal{B}_{1,b'}$'s advantage is at least as large as \mathcal{A} 's since in this case $\mathcal{B}_{1,b'}$ always outputs the correct bit. The result follows. \square

Lemma 18. There exists an adversary \mathcal{B}_2 with similar running time to \mathcal{A} such that

$$\text{Adv}_{\text{GM}, C_{\text{gm}}}^{\text{g}1}(\mathcal{A}) \leq \text{Adv}_{\text{GM}, C_{\text{gm}}}^{\text{g}2}(\mathcal{A}) + q \cdot \text{Adv}_{\text{Sig}}^{\text{suf-cma}}(\mathcal{B}_2).$$

Chapter 6. WhatsUpp with Sender Keys? Analysis, Improvements and Security Proofs

Proof. Let E be the event that \mathcal{A} in a clean execution of Game Γ_1 (i.e., when cleanness predicate C evaluates to true) calls $\text{RECEIVE}(ID, (C_{\text{core}}, C_{2\text{pc}}))$ that outputs $m \neq \perp$ for some ID such that $(C_{\text{core}}, \vec{C})$ was not previously output by SEND , where $C_{2\text{pc}}$ is some (correct) 2PC ciphertext by definition of Game Γ_1 . Observe first that Game Γ_1 and Game Γ_2 are identical given $\neg E$. Note that at most q signature keys are sampled and sent over the two-party channels during the game's execution: the first Send call in an epoch and the (single) $\text{Exec}(\cdot, \text{crt}, IDs)$ call result in one signature key being sampled in each call.

Then, let E_i be the event that \mathcal{A} 's first call to RECEIVE with argument $(ID, (C_{\text{core}}, C_{2\text{pc}}))$ satisfying the conditions of E is such that the internal Recv call outputs ID' corresponding to the i -th signature key sampled by the challenger (i.e., the i -th sender key).

We define SUF-CMA adversary $\mathcal{B}_{ID, e}$ who simulates for Game Γ_1 adversary \mathcal{A} given E_i holds. \mathcal{B}_i simulates as follows. \mathcal{B}_i locally simulates for \mathcal{A} and responds to all of \mathcal{A} 's queries except queries involving the i -th sender key sampled: let ID be the key holder. For these queries, ID sets sk in variable SK to the SUF-CMA public key pk . \mathcal{B}_i generates all signatures associated with sk via SUF-CMA oracle SIGN .

Finally, consider when \mathcal{A} makes their first query of the form $\text{RECEIVE}(ID, C)$ that returns $m \neq \perp$ that was not previously output by SEND . Observe that $C = (C_{\text{core}} = (M, \sigma), C_{2\text{pc}})$ and $M = (c, (e, i), \text{kc}', i_{\text{ck}'}, ID)$ where σ is a signature on $(c, (e, i), i_{\text{ck}'}, ID)$. As forgeries on $C_{2\text{pc}}$ are disallowed, only C_{core} can possibly be the source of the forgery. By construction of Recv , $(c, (e, i), \text{kc}', i_{\text{ck}'}, ID)$ must be different from values previously input to Recv for a non-bottom value to be output, and, by definition of event E_i , C_{core} must differ from values previously output by SEND . Moreover, by cleanness, \mathcal{A} must not have been able to make a state exposure that enables it to access signature secret key ssk . Thus, the simulation is well-defined and signature σ is a valid forgery, and so \mathcal{B}_i extracts σ from C and returns (M, σ) to its SUF-CMA challenger.

Finally, we have

$$\begin{aligned}
 \text{Adv}_{\text{GM}, C_{\text{gm}}}^{\text{g}^1}(\mathcal{A}) &= |\Pr[G_1^1 \Rightarrow 1 \wedge \neg E] - \Pr[G_1^0 \Rightarrow 1 \wedge \neg E] + \Pr[G_1^1 \Rightarrow 1 \wedge E] - \Pr[G_1^0 \Rightarrow 1 \wedge E]| \\
 &\leq \text{Adv}_{\text{GM}, C_{\text{gm}}}^{\text{g}^2}(\mathcal{A}) + |\Pr[G_1^1 \Rightarrow 1 \wedge E] - \Pr[G_1^0 \Rightarrow 1 \wedge E]| \\
 &\leq \text{Adv}_{\text{GM}, C_{\text{gm}}}^{\text{g}^2}(\mathcal{A}) + \Pr[E] \\
 &\leq \text{Adv}_{\text{GM}, C_{\text{gm}}}^{\text{g}^2}(\mathcal{A}) + \sum_i \Pr[E_i] \\
 &\leq \text{Adv}_{\text{GM}, C_{\text{gm}}}^{\text{g}^2}(\mathcal{A}) + \sum_i \text{Adv}_{\text{Sig}}^{\text{suf-cma}}(\mathcal{B}_i) \\
 &\leq \text{Adv}_{\text{GM}, C_{\text{gm}}}^{\text{g}^2}(\mathcal{A}) + q \cdot \text{Adv}_{\text{Sig}}^{\text{suf-cma}}(\mathcal{B}_2),
 \end{aligned}$$

where the last step holds by combining each \mathcal{B}_i into \mathcal{B}_2 . □

6.6 Sender Keys and Sender Keys+: Full Protocols and Security

Lemma 19. There exists an adversary \mathcal{B}_3 with similar running time to \mathcal{A} such that

$$\text{Adv}_{\text{GM}, \text{C}_{\text{gm}}}^{\text{g}^2}(\mathcal{A}) \leq q_{\text{chal}} \cdot \text{Adv}_{\text{GM}, \text{C}_{\text{gm}}}^{\text{g}^3}(\mathcal{B}_3)$$

where $q_{\text{chal}} \leq q$ denotes the number of CHAL oracle queries made by \mathcal{A} .

Proof. We adopt the same high-level strategy as the proof of Lemma 6 in [ACDT19] (conference version [ACDT20]). Let H_0 be exactly Game Γ_2 with $b = 0$. For $i \in [1, q_{\text{chal}}]$, let H_i be exactly H_{i-1} except the i -th query $\text{CHAL}(ID, m_0, m_1)$ uses m_1 (i.e., acts as if the challenge bit is $b = 1$). Observe first that we have $\text{Adv}_{\text{GM}, \text{C}_{\text{gm}}}^{\text{g}^2}(\mathcal{A}) = |\Pr[H_0 \Rightarrow 1] - \Pr[H_{q_{\text{chal}}} \Rightarrow 1]|$. We will show, for $i \in [1, q_{\text{chal}}]$, that there exists adversary $\mathcal{B}_{3,i}$ playing Game Γ_2 such that $|\Pr[H_i \Rightarrow 1] - \Pr[H_{i-1} \Rightarrow 1]| = \text{Adv}_{\text{GM}, \text{C}_{\text{gm}}}^{\text{g}^3}(\mathcal{B}_{3,i})$. The claimed result then follows by applying the sequence of hybrids and the triangle inequality.

$\mathcal{B}_{3,i}$ simulates as follows. For \mathcal{A} 's first $i - 1$ $\text{CHAL}(ID, m_0, m_1)$ calls, $\mathcal{B}_{3,i}$ calls $\text{SEND}(ID, m_1)$ and returns the result. For \mathcal{A} 's i th $\text{CHAL}(ID, m_0, m_1)$ call, $\mathcal{B}_{3,i}$ calls $\text{CHAL}(ID, m_0, m_1)$ and returns the result. For \mathcal{A} 's subsequent $\text{CHAL}(ID, m_0, m_1)$ calls, $\mathcal{B}_{3,i}$ calls $\text{SEND}(ID, m_0)$ and returns the result. If \mathcal{A} ever makes an EXPOSE query that would trivially allow for them to decrypt any challenge ciphertext, or has previously called EXPOSE such that the resulting CHAL query would be trivially decryptable, $\mathcal{B}_{3,i}$ aborts. Note that this condition can be efficiently determined based on \mathcal{A} 's oracle queries. $\mathcal{B}_{3,i}$ processes all other queries using its own oracles.

Note that if \mathcal{A} 's (multi-challenge) execution satisfies the cleanness predicate, then so too does $\mathcal{B}_{3,i}$'s. To see this, note that $\mathcal{B}_{3,i}$ and \mathcal{A} make the same queries to all oracles except for CHAL and SEND. In particular, since $\mathcal{B}_{3,i}$ makes the same EXPOSE queries as \mathcal{A} (given $\mathcal{B}_{3,i}$ does not abort) and less CHAL queries, there are the same or possibly less opportunities for the challenge predicate to fail in $\mathcal{B}_{3,i}$'s execution as compared to \mathcal{A} 's. Moreover, the additional SEND queries that $\mathcal{B}_{3,i}$ makes do not affect any predicates. Thus, if $\mathcal{B}_{3,i}$'s challenge bit b is 0, then $\mathcal{B}_{3,i}$ perfectly simulates H_{i-1} , and similarly $\mathcal{B}_{3,i}$ perfectly simulates H_i given $b = 1$. The result follows. \square

Lemma 20. There exists an adversary \mathcal{B}_4 with similar running time to \mathcal{A} such that

$$\text{Adv}_{\text{GM}, \text{C}_{\text{gm}}}^{\text{g}^3}(\mathcal{A}) \leq q \cdot (\text{Adv}_{\text{GM}, \text{C}_{\text{gm}}}^{\text{g}^4}(\mathcal{A}) + \text{Adv}_{2\text{PC}, \text{C}_{2\text{pc}}, \Delta}^{2\text{pc-ind}}(\mathcal{B}_4) + q \cdot \text{Adv}_{\text{H}}^{\text{prg}}(\mathcal{B}_4))$$

Proof. The proof of this lemma proceeds via hybrid sub-games. Consider the (restricted) chain key sequence in an execution of Game Γ_3 starting from epoch e until key epoch (ID, e', i') corresponding to the output of Send in the $\text{CHAL}(ID, \cdot, \cdot)$ call, if it exists. If it does not exist, then Game Γ_3 adversary \mathcal{A} has no advantage as their execution is independent of the challenge bit. Otherwise, we replace all chain keys and their corresponding message keys in this sequence by uniformly random values. Note that, in the protocol in Figures 6.19 to 6.22, we model H_1 and H_2 as a PRG $\text{H} : \{0, 1\}^\lambda \rightarrow \{0, 1\}^{2\lambda}$. Namely, $\text{H}(ck_i) = (ck_{i+1}, mk_i)$ outputs an updated chain

Chapter 6. WhatsUpp with Sender Keys? Analysis, Improvements and Security Proofs

key and a new message key.

Let E_{ck} be the event where call $\text{CHAL}(ID, \cdot, \cdot)$ is made such that the underlying Send call uses message key mk iteratively derived from some ck , where $\text{ck} = \text{ck}_{ID}^{e,0}$, for some e , is the start of ID 's corresponding chain key sequence. Observe that at most q such chain key sequences are possible at the start of an execution of Game Γ_3 where \mathcal{A} makes q oracle queries, and each such ck is associated with a corresponding ID (the user who sampled ck). Denote for simplicity this chain key sequence by $\text{ck}_1, \dots, \text{ck}_m$ (where $\text{ck}_1 := \text{ck}_{ID}^{e,0}$, $\text{ck}_m := \text{ck}_{ID}^{e',i'}$ and $m \leq q$). For each possible E_{ck} , we can construct a sequence of hybrids H_i for $i = -1, 0, \dots, m$ as follows. Game H_{-1} is as in the original Game Γ_3 . Game H_0 differs from Game H_{-1} in that all two-party ciphertexts encrypting ck or ancestors in its key sequence are replaced with encryptions of dummy message 0^ℓ for messages of length ℓ . For $i \geq 1$, Game H_{i-1} differs from Game H_i in that, in the latter, we replace both mk_{i-1} and ck_i by uniformly random $r_{i-1} \xleftarrow{\$} \mathcal{K}$ and $s_i \xleftarrow{\$} \mathcal{W}$, respectively. Finally, Game H_m is Game Γ_4 ; all non-exposed keys are independent.

We first construct $2\text{PC-IND}_{2\text{PC},b,C_{2\text{pc}},\Delta}$ adversary \mathcal{B} that simulates for adversary \mathcal{A} playing (as we will argue) Game H_{-1} or H_0 depending on its challenger's bit. \mathcal{B} simulates similarly to \mathcal{B}_1 in the proof of Lemma 17 except when simulating $2\text{PC.Send}(m, ID')$ calls. Here, instead of replacing all such calls with the output of $\text{SEND}(ID, ID', m)$, \mathcal{B} replaces calls that encrypt ck or its key sequence ancestors in m with the output of $\text{CHAL}(ID, ID', m, 0^{|m|})$. \mathcal{B} otherwise simulates identically. Observe that in a clean execution of H_{-1} , chain key ck_1 must not be exposed, and that \mathcal{B} , who is parametrised by ck , can deduce exactly which 2PC.Send calls to replace with a CHAL call. It follows that \mathcal{B} simulates H_{-1} given the challenge bit is 0 and H_0 given it is 1.

Note in a clean execution of H_0 that the starting key in the chain $\text{ck}_{ID}^{e,0} \notin \text{ExpKeys}_{\text{ck}}$ is generated by ID using fresh randomness. Besides, $\text{ck}_{ID}^{e,0}$ is only sent over two-party channels that contains no information about the key due to the previous hop (noting in a clean execution that the channels must have healed if previously compromised), so it is hidden from \mathcal{A} . Now, let \mathcal{A} be an adversary that interpolates between any two Games H_{i-1} and H_i . Then, we can create an adversary \mathcal{B} against PRG indistinguishability from \mathcal{A} as follows. Since by induction in H_{i-1} the seed ck_{i-1} of the PRG is a uniformly random value, \mathcal{B} simply embeds a PRG challenge in mk_{i-1} and ck_i (recall that we consider a PRG with an expansion factor of 2 such that $\text{H}(\text{ck}_{i-1}) = (\text{ck}_i, \text{mk}_{i-1})$). Then, \mathcal{B} simulates the rest of the game locally and returns the guess of \mathcal{A} .

It follows that the simulation is perfect. By considering all events E_{ck} and the union bound, it follows that there exists \mathcal{B}_4 so that:

$$\text{Adv}_{\text{GM}, C_{\text{gm}}}^{\text{g}^3}(\mathcal{A}) \leq q \cdot (\text{Adv}_{\text{GM}, C_{\text{gm}}}^{\text{g}^4}(\mathcal{A}) + \text{Adv}_{2\text{PC}, C_{2\text{pc}}, \Delta}^{2\text{pc-ind}}(\mathcal{B}_4) + q \cdot \text{Adv}_{\text{H}}^{\text{prg}}(\mathcal{B}_4))$$

□

6.6 Sender Keys and Sender Keys+: Full Protocols and Security

Lemma 21. There exists an adversary \mathcal{B} with similar running time to \mathcal{A} such that

$$\text{Adv}_{\text{GM}, \text{C}_{\text{gm}}}^{\text{g}^4}(\mathcal{A}) \leq \text{Adv}_{\text{SymEnc}}^{\text{ind-cpa}}(\mathcal{B})$$

Proof. We reduce to the security of the encryption scheme. Let \mathcal{A} be an adversary against Game Γ_4 . Then, we can build an adversary \mathcal{B} against the IND-CPA security of the encryption scheme. Let b^* be the (hidden) bit that parameterises the IND-CPA game of \mathcal{B} . Then, \mathcal{B} simulates Game Γ_4 for \mathcal{A} except for the challenge query $q^* = \text{CHAL}(ID^*, m_0, m_1)$, where it proceeds as follows. \mathcal{B} receives m_0, m_1 from \mathcal{A} and forwards them to the IND-CPA challenger, who outputs a ciphertext c^* . Then, \mathcal{B} crafts a ciphertext C^* as if it originated from ID^* and sends it to \mathcal{A} . The simulation continues until the game finishes, and \mathcal{B} returns the same guess b' as \mathcal{A} .

As the message key mk used to encrypt the challenge message in the original Game Γ_4 is a uniformly random key, as we argued above, the simulation is perfect. Hence, the lemma follows. \square

Finally, Theorem 20 follows by combining the sequence of hybrids above.

Proof Strategy for Corollary 1. The hybrids are defined similarly except that they differ in the definition of Game Γ_1 and Game Γ_4 . Let the resulting sequence of games be denoted Game Γ_1' , Game Γ_2' , Game Γ_3' , Game Γ_4' . In Game Γ_1' , all two-party channel ciphertexts that cannot be trivially decrypted by the adversary are replaced with encryptions of dummy strings of the form 0^ℓ . Game Γ_4' differs as in Game Γ_4 except that 2PC ciphertexts are not changed. We can then essentially directly use the above lemmas except for Lemmas 17 and 20:

- For Game Γ_1' , note that since $2\text{PC-IND}_{2\text{PC}, b, \text{C}_{2\text{pc}}, \Delta}$ adversary $\mathcal{B}_{1, b'}$ is given all of \mathcal{A} 's queries in advance, it can efficiently deduce which 2PC.Send queries to replace with CHAL and SEND depending on which ciphertexts can be trivially exposed by \mathcal{A} or not. It then follows that

$$\text{Adv}_{\text{GM}, \text{C}_{\text{gm}}}^{\text{g}^0}(\mathcal{A}) \leq \text{Adv}_{\text{GM}, \text{C}_{\text{gm}}}^{\text{g}^1}(\mathcal{A}) + 2 \cdot \text{Adv}_{2\text{PC}, \text{C}_{2\text{pc}}, \Delta}^{2\text{pc-ind}}(\mathcal{B}_1)$$

- For Game Γ_4' , the reduction no longer needs to guess ck , since this information can be efficiently derived from the sequence of queries q_1, \dots, q_q initially given to \mathcal{A} . A sequence of hybrids H_i for $i \geq 0$ can then be directly constructed; note we can ignore the hop between H_{-1} and H_0 since Game Γ_1 already handles this. It then follows that

$$\text{Adv}_{\text{GM}, \text{C}_{\text{gm}}}^{\text{g}^3}(\mathcal{A}) \leq \text{Adv}_{\text{GM}, \text{C}_{\text{gm}}}^{\text{g}^4}(\mathcal{A}) + q \cdot \text{Adv}_{\text{H}}^{\text{prg}}(\mathcal{B}_4)$$

The result follows.

6.6.3 Sender Keys+ Security

In this subsection, we prove the security of our Sender Keys+ protocol. We do so with respect to the modified cleanness predicate in Figure 6.18. The proof follows similar steps as the proof for Theorem 20.

From our Sender Keys+ protocol, we model H_1, H_2, H_3 as a PRG $H : \{0, 1\}^\lambda \rightarrow \{0, 1\}^{3\lambda}$, and the KDF F used for the update operation as a dual PRF. Note the proof would still work with respect to the Sender Keys predicate C_{sk}^Δ (Figure 6.13) as the new predicate C_{sk+}^Δ (Figure 6.18) is strictly less restrictive. Towards the proof, we re-define the notion of key sequences introduced in Section 6.6.2 to capture our new update mechanism. We now consider chain key sequences for ID starting with $ck_{ID}^{e,0}$ where zero or more update operations from $ID' \neq ID$ are applied to a chain key of the form $ck_{ID}^{e',i'}$. In addition, updates from parties $ID \neq ID'$ now result in a new key sequence. We state the theorem below.

Theorem 21. Let $\text{SymEnc} := (\text{Enc}, \text{Dec})$ be a $(q, \epsilon_{\text{sym}})$ -IND-CPA $_{\text{SymEnc}}$ symmetric encryption scheme, $\text{Sig} := (\text{Gen}, \text{Sign}, \text{Vrfy})$ a $(q, \epsilon_{\text{sig}})$ -SUF-CMA $_{\text{Sig}}$ signature scheme, $H : \{0, 1\}^\lambda \rightarrow \{0, 1\}^\lambda \times \{0, 1\}^\lambda$ a $(q, \epsilon_{\text{prg}})$ -PRG $_H$ function, F a $(q, \epsilon_{\text{dprf}})$ -dPRF function, MAC a $(q, \epsilon_{\text{mac}})$ -SUF-CMA $_{\text{MAC}}$ message authentication code and 2PC a $(q, \epsilon_{2\text{pc}})$ -2PC-IND $_{2\text{PC}, C_{2\text{pc}}, \Delta}$ two-party channels scheme for PCS bound $\Delta > 0$. Then Sender Keys+ (Figures 6.19 to 6.22) is

$$(q, 2 \cdot \epsilon_{2\text{pc}} + q^3 \cdot (\epsilon_{2\text{pc}} + \epsilon_{\text{sym}} + N \cdot q \cdot \epsilon_{\text{prg}} + q \cdot \epsilon_{\text{dprf}} + q \cdot \epsilon_{\text{mac}}) + q \cdot \epsilon_{\text{sig}})\text{-M-IND}_{\text{GM}, C_{\text{gm}}}$$

with respect to predicate $C_{\text{gm}} = C_{sk+}^\Delta$ (Figure 6.18) and concurrency bound N , where the two-party channels predicate $C_{2\text{pc}}$ is defined in Figures 6.4 and 6.5.

Hybrid Games. We define the main sequence of games below.

Game Γ_0 : This is the original M-IND game, parameterised by cleanness predicate $C_{\text{gm}} = C_{sk+}^\Delta$.

Game Γ_1 : In this game, we remove the **return** b conditions whenever $\text{DELIVER}(ID, (T_{\text{core}}, C_{2\text{pc}}))$ or $\text{RECEIVE}(ID, (C_{\text{core}}, C_{2\text{pc}}))$ are called such that $C_{2\text{pc}}$ was not previously output in some previous oracle query that outputs a ciphertext or control message.

Game Γ_2 : In this game, we completely remove the **return** b condition in the DELIVER oracle and in the RECEIVE oracle for forgeries that occur when predicate $C_{sk\text{-inj-core}}^\Delta(C_{\text{core}})$ is true for ciphertexts of the form C_{core} .

Game Γ_3 : In this game, we allow a single call to the CHAL oracle, as opposed to arbitrarily many calls.

Game Γ_4 : In this game, all chain keys and corresponding message and MAC keys in ID 's key sequence that includes ck such that $H_2(ck) = mk$ and $H_3(ck) = \tau k$, where mk (resp. τk) is the message key (resp. MAC key) used in the underlying Send call in the $\text{CHAL}(ID, \cdot, \cdot)$ query (if it exists), are all replaced by uniformly random values. In addition, all 2PC ciphertexts that

6.6 Sender Keys and Sender Keys+: Full Protocols and Security

transmit the chain key ck or keys earlier in the key sequence leading to ck are replaced with encryptions of 0^ℓ where ℓ is the length of the message encrypted.

Game Γ_5 : In this game, we completely remove the **return** b condition in the RECEIVE oracles. Hence, the adversary cannot win the game by injecting. To complete the proof, Game 5 is simulated by an IND-CPA SymEnc adversary.

Lemma 22. There exists an adversary \mathcal{B}_1 with similar running time to \mathcal{A} such that

$$\text{Adv}_{\text{GM}, \text{C}_{\text{gm}}}^{\text{g}0}(\mathcal{A}) \leq \text{Adv}_{\text{GM}, \text{C}_{\text{gm}}}^{\text{g}1}(\mathcal{A}) + 2 \cdot \text{Adv}_{2\text{PC}, \text{C}_{2\text{pc}, \Delta}}^{2\text{pc-ind}}(\mathcal{B}_1).$$

Proof. The proof is essentially identical to that of Lemma 17 so we omit it. □

Lemma 23. There exists an adversary \mathcal{B}_2 with similar running time to \mathcal{A} such that

$$\text{Adv}_{\text{GM}, \text{C}_{\text{gm}}}^{\text{g}1}(\mathcal{A}) \leq \text{Adv}_{\text{GM}, \text{C}_{\text{gm}}}^{\text{g}2}(\mathcal{A}) + q \cdot \text{Adv}_{\text{Sig}}^{\text{suf-cma}}(\mathcal{B}_2).$$

Proof. The proof follows the same high-level idea as for Lemma 18. That is, we consider events E_i for $i \in [1, q']$ for some $q \leq q'$ such that the first successful forgery is made using the i -th signature key pair sampled. For this proof, we consider forgeries now over both RECEIVE and DELIVER rather than just RECEIVE. By construction of Sender Keys+, DELIVER forgeries given $C_{\text{sk-inj-core}}^\Delta(C_{\text{core}})$ is true only occur as a result of a signature forgery. Thus, by a very similar reduction to Lemma 18, the result follows. Note that unlike in the proof of Theorem 20 at this point, we have not yet completely disallowed injections on RECEIVE: we still allow forgeries that are permitted by $C_{\text{sk+}}^\Delta$ but disallowed by C_{sk}^Δ . □

Lemma 24. There exists an adversary \mathcal{B}_3 with similar running time to \mathcal{A} such that

$$\text{Adv}_{\text{GM}, \text{C}_{\text{gm}}}^{\text{g}2}(\mathcal{A}) \leq q_{\text{chal}} \cdot \text{Adv}_{\text{GM}, \text{C}_{\text{gm}}}^{\text{g}3}(\mathcal{B}_3)$$

where $q_{\text{chal}} \leq q$ denotes the number of CHAL oracle queries made by \mathcal{A} .

Proof. The proof is identical to that of Lemma 19 so we omit it. □

Lemma 25. There exists an adversary \mathcal{B}_4 with similar running time to \mathcal{A} such that

$$\text{Adv}_{\text{GM}, \text{C}_{\text{gm}}}^{\text{g}3}(\mathcal{A}) \leq q^2 \cdot (\text{Adv}_{\text{GM}, \text{C}_{\text{gm}}}^{\text{g}4}(\mathcal{A}) + \text{Adv}_{2\text{PC}, \text{C}_{2\text{pc}, \Delta}}^{2\text{pc-ind}}(\mathcal{B}_4) + N \cdot q \cdot \text{Adv}_{\text{H}}^{\text{prg}}(\mathcal{B}_4) + q \cdot \text{Adv}_{\text{F}}^{2\text{prf}}(\mathcal{B}_4))$$

where N is the concurrency bound (c.f. Section 6.5.2).

Proof. The proof diverges from the proof of Lemma 20 in order to handle the new update mechanism. As in Lemma 20, we consider the event E_{ck} where the $\text{CHAL}(ID, \cdot, \cdot)$ call invokes Send with key mk in the key sequence starting from ck . Lemma 20 then constructs hybrids H_{-1}, H_0, \dots, H_m given E_{ck} .

Chapter 6. WhatsUpp with Sender Keys? Analysis, Improvements and Security Proofs

Observe that, fixing ck , mk could have been derived from zero or more update operations initiated by $ID' \neq ID$. Now, by cleanness, after invoking the security of the two-party channels, mk is hidden from the adversary. If no update operations were made, then it must be that ck is not exposed. If there was one update operation, then by cleanness, either the update secret r or ck are hidden (or possibly both).

Let i be the i -th last update operation applied to form mk , where $i \in [0, q']$ for some $q' < q$. Let $E_{ck,i}$ be the event that the i -th last update operation is hidden from the adversary for $i \geq 1$, and $i = 0$ where ck itself is secure. Observe that $E_{ck} = \cup_{i \in [0, q']} E_{i, q}$ for some $q' < q$.

Our proof strategy then is as follows. For each $E_{ck,i}$, we first hop by invoking the security of the two-party channels to replace all two-party ciphertexts that communicate 'safe' value r or ck (or a descendent in the key sequence) with encryptions of dummy messages (hopping between H_{-1} and H_0). The simulation proceeds analogously to that in Lemma 20. Then (hopping to H_m):

- For $i = 0$, we hop by iteratively replacing all relevant H_j calls using the PRG assumption; there are at most $N \cdot q$ such queries. We replace each F call (of which there are at most q) along the way with a uniform value by the dual PRF assumption, keying F with ck' .
- For $i \geq 1$, we first hop by replacing the first call to F made by the challenger with r with a uniform value using the dual PRF assumption and keying F with r . We then hop using the PRG assumption on H_j and keying PRF F thereafter with ck' ; note there are at most $N \cdot i \leq N \cdot q$ such calls.

By a similar argument to Lemma 20, the hop with the two-party channels and the hops with the PRG and dual PRF are sound. The result follows by combining the sequence of hybrids. \square

Lemma 26. There exists an adversary \mathcal{B}_5 with similar running time to \mathcal{A} such that

$$\text{Adv}_{GM, C_{gm}}^{g^4}(\mathcal{A}) \leq \text{Adv}_{GM, C_{gm}}^{g^5}(\mathcal{A}) + q \cdot \text{Adv}_{MAC}^{\text{suf-cma}}(\mathcal{B}_5).$$

Proof. Let E be the event that a successful forgery to RECEIVE is made in a clean execution of Game Γ_4 . Note Game Γ_4 and Game Γ_5 are identical given $\neg E$. Let E_i be the event that \mathcal{A} 's first forgery with RECEIVE is with respect to the i -th key sequence starting from ck ; note $E = \cup_{i \in [1, q']} E_i$ for some $q' \leq q$. Observe by the previous game hop that the corresponding MAC key τk is uniform. We construct a SUF-CMA_{MAC} adversary \mathcal{B} that simulates for Game Γ_4 adversary \mathcal{A} given E_i holds. \mathcal{B} locally simulates all calls except MAC query using τk on message m during the simulation is replaced with the output of query TAG. Finally, when \mathcal{A} makes a successful forgery using τk , \mathcal{B} extracts the tag t from the message and returns it to its challenger. The simulation is perfect and so the result follows by a similar derivation as in Lemma 18. \square

6.6 Sender Keys and Sender Keys+: Full Protocols and Security

Lemma 27. There exists an adversary \mathcal{B} with similar running time to \mathcal{A} such that

$$\text{Adv}_{\text{GM}, \text{C}_{\text{gm}}}^{\text{g5}}(\mathcal{A}) \leq \text{Adv}_{\text{SymEnc}}^{\text{ind-cpa}}(\mathcal{B})$$

Proof. The proof is essentially identical to that of Lemma 21 so we omit it. □

7 Conclusion

In this thesis, we have considered four different aspects of messaging, the first two concerning two-party communication (Part I), namely post-quantum X3DH key exchange (Chapter 3) and active attack detection (Chapter 4), and the latter two group messaging (Part II), namely group administration (Chapter 5) and the Sender Keys protocol (Chapter 6), each of which we have reasoned about in the framework of provable security.

7.1 Primitive Summary

We introduced several primitives and abstractions over the course of this thesis to capture different aspects of messaging. Here, we summarise and compare them, and consider how they may be instantiated. As done for this thesis as a whole, we present our comparison in two parts, namely in terms of 1) two-party communication and 2) group communication.

7.1.1 Two-Party Communication

Authentication. The goal of Chapter 3 is to construct a post-quantum X3DH-like authenticated key exchange protocol, namely K-Waay. We capture this in our deniable authenticated key exchange (DAKE) primitive that we introduce in Section 3.3. The syntax is somewhat non-standard as it generalises previous work in the BatchReceive algorithm that allows a *receiver* to simultaneously finish several key exchange sessions with different *senders*. Apart from this, our security notion can be seen as a Bellare-Rogaway-style definition [BR94] that combines aspects of the models of Hashimoto et al. [HKKP22] and Brendel et al. [BFG⁺22b]. DAKE assumes a public key infrastructure (PKI) is made available to all parties for registering long-term keys (in particular, our model captures for malicious key registration).

(Authenticated) Ratcheted Communication. In Chapter 4, we define the *ratcheted communication* (RC) primitive to capture stateful message exchange between two parties where ciphertexts can be re-ordered or dropped (i.e., supports *immediate decryption* [ACD19]). To capture an additional out-of-band channel that parties use for detecting active attacks, we

Chapter 7. Conclusion

then define the more general *authenticated ratcheted communication* (ARC) primitive. These primitives abstract away authentication in the `Init` algorithm that outputs the initial states of both parties; we assume this is implemented in practice through something like K-Waay.

Syntactically, RC and ARC both have stateful `Send` and `Receive` algorithms which take associated data as input. Both algorithms output an *ordinal*, such as an index, which captures information about the order of ciphertexts. ARC further provides two algorithms, `AuthSend` and `AuthReceive`, that model sending and receiving authentication tags over the out-of-band channel. We propose several instantiations of both RC and ARC throughout the chapter with different security/efficiency trade-offs. All of our constructions make black-box use of an RC – one construction (Figure 4.8) requires the underlying RC to be RID-secure (Definition 37), and the rest only require correctness properties (Definitions 35 and 36) on the underlying RC to hold.

Two-Party Channels. In Chapter 6 we introduce another abstraction for two-party communication, namely *two-party channels* (2PC). This primitive is tailored to its use in instantiating Sender Keys and also supports immediate decryption. By contrast to `Init` in RC and ARC, we allow a given party to initialise channels with several parties over time via `InitCh` after initialising their own states with `Init`. In addition, our `Send` and `Receive` algorithms for 2PC do not take associated data as input, and we specifically assume ordinals are pairs of integers, which we call *channel epochs* that behave like the epoch/index pairs of Alwen et al. [ACD19]. We do not provide an explicit instantiation for 2PC here but we note that `InitCh` can be instantiated as above via an AKE and PKI, and `Send` and `Receive` via, e.g., the Double Ratchet protocol [ACD19] as done in implementations of Sender Keys in practice.

7.1.2 Group Communication

Authentication. Throughout Chapter 5, like our DAKE K-Waay, our protocols directly assume the existence of a PKI on which parties can upload keys. Unlike for DAKE for which we captured static long-term keys, we assume that parties can update their keys over time. Note that our protocols do not require a two-party AKE protocol: in practice in MLS, parties upload so-called *key bundles* that parties use as bootstrapping key material when a party is added to a group. In Chapter 6, authentication is instead abstracted away in the underlying two-party channels. Thus, Sender Keys and Sender Keys+ as written in Chapter 6 do not directly invoke a PKI, although two-party channels in practice are of course instantiated with PKI and AKE.

Continuous Group Key Agreement. In Chapter 5, we study continuous group key agreement (CGKA) and a generalisation that we introduce, namely *administrated CGKA* (A-CGKA). In CGKA, a dynamic group agrees on a sequence of secret keys over time (that can be used to build group messaging proper [ACDT21a] but may be useful in other contexts). Our formulation of CGKA, following the current MLS standard, is in the propose-and-commit paradigm, where group changes and key updates are first proposed (`Prop`) before being combined and committed (`Commit`) by some group member, after which group members process the corre-

sponding commit message (Proc). We discuss and compare variants of CGKA in the literature in Section 5.2.1.

Extending plain CGKA, A-CGKA provides first-class support for the addition, removal and updating of the keying material of *group administrators*. In Section 5.3, we formalise and prove secure two A-CGKA constructions, namely *individual admin signatures*, or IAS, and *dynamic group signature*, or DGS.

Group Messenger. In Chapter 6, we model Sender Keys and our enhanced protocol Sender Keys+ as instances of a *group messenger*, a primitive which we introduce. Unlike CGKA, a group messenger captures the sending and receiving of application messages (via Send and Recv). Our syntax here does not capture input associated data, and only Recv outputs an ordinal, which we simply assume is an epoch/index pair. Epochs here increase whenever the group membership changes (add/rem) or a group member performs a key update operation (upd). Group operations are captured via dedicated algorithms Exec and Proc: for simplicity we do not work in the propose-and-commit paradigm but it is straightforward to capture this in a modified group messenger primitive.

7.2 Discussion and Future Work

7.2.1 Composability

Throughout this thesis, we have worked with game-based security definitions, which do not immediately or necessarily provide security guarantees when the corresponding primitives are used in a broader context and under concurrent composition. Note our DAKE indistinguishability notion (Definition 26) considers a single real/random challenge query (TEST in our syntax); multi-challenge security follows with q multiplicative tightness loss for q challenge queries via a standard argument. Moreover, results of Brzuska et al. [BFWW11] indicate that Bellare-Rogaway-style key indistinguishability notions provide security guarantees for a broad class of protocols which assume that key exchange has taken place at initialisation time. As many two-party ratcheting and messaging primitives assume this already, including our RC primitive, they benefit from these composition results. Since these results only consider two-party key exchange, a natural direction is to devise similar ones for game-based CGKA. Previous work on both two-party [BFG⁺22a, CJSV22] and group messaging [AJM22, HKP⁺21, AHKM22, ACJM20] has alternatively considered simulation-based security notions, mostly in the *universal composability* framework [Can01] (the work of Jost et al. [JMM19b] in a variant of the *constructive cryptography* framework being an exception here).

7.2.2 Model Limitations

Whether the security notions are game-based or simulation-based, one must be careful interpreting the security claims of a given work. As we note below in more detail, there are

Chapter 7. Conclusion

attack vectors, such as randomness manipulation [BRV20], that we do not always capture in our security notions in this thesis.

Our notion of state exposure is not very fine-grained: in Chapter 3, we allow individual session states and long-term keys to be exposed separately, but otherwise we assume the entire secret state of a party is given to the adversary at once. Consequently the effects of partial leakage or state exposure are only captured through these means: bridging this gap further, for example considering leakage resilience in messaging, is of interest.

Our security notions, like all work we are aware of on messaging, further assume where relevant that code is faithfully and honestly executed at all times which may not always hold in practice.

7.2.3 Unification and Verification

Due to the complexity of messaging primitives, different works tend to introduce their own primitive syntax and security notions that can differ enough to render comparison across works complex and time-consuming. For pedagogical purposes and to enable easier formal comparison between different works, it is thus of interest to overcome this complexity. The work of Jost et al. [JMM19b] provided some recourse for two-party messaging: they show components of protocols in the literature [PR18, JS18, ACD19, DV19, JMM19a] can be expressed with some minor modifications in their composable framework. In addition, some authors have also compared related work in their particular security model, including the game-based models of Durak and Vaudenay [DV18] and Cremers and Zhao [CZ22].

In addition, the complexity of messaging makes it error-prone and difficult to verify. It is not always sufficient to reduce the complexity by considering a simplified security definition either, since attack vectors can and have previously been overlooked when taking this approach [AJM22]. Formal analysis and verification is therefore a natural direction to pursue, and has successfully been applied in messaging [KBB17, WPBB23]. However, there is still a gap in the literature for messaging protocols as much of the previous work considers symbolic/Dolev-Yao-style modelling rather than full computational security. Nonetheless, results like the analysis of Signal’s PQXDH key exchange protocol [BJKS23b] in CryptoVerif [Bla07] are promising and indicate that formal computational analysis of different messaging protocols is in reach.

7.2.4 Deniable Post-Quantum X3DH (Chapter 3)

In Chapter 3, we proved the key indistinguishability of our protocol K-Waay in a model that considers state exposures like Hashimoto et al.’s [HKKP22] but is nonetheless weaker (our protocol is provably secure under a notion similar to that of Brendel et al. [BFG⁺22b]). This is mainly since K-Waay only uses ephemeral split-KEM keys. As noted by Brendel et al. [BFG⁺20], however, it seems much more difficult to construct split-KEM secure under

several encapsulation/decapsulation queries, which we leave as important future work.

An interesting line of research would be to try to build other unforgeable IND-1BatchCCA split-KEMs that are more efficient (mostly in key and ciphertext size). One obvious direction would be to work over structured lattices [LS15, LPR10, SSTX09]. Indeed, Ring/Module-LWE with hints (similar to our Extended-LWE problem) have already been analysed from a theoretical point of view [BJRW21, MKMS22]. We also believe that our techniques can also be applied in the ring setting. However, for security purposes one needs to take a ring dimension d to be at least linear in the security parameter λ which becomes problematic when proving deniability. Indeed, the leaked hint is informally the product of secret keys of both parties. Thus, in the ring setting the hint would be at least a single polynomial, which contains $d = O(\lambda)$ coefficients. We predict that this would result with much larger reduction loss than what we have now. However, the concrete analysis is left as future work.

On the more practical side, it would be informative to benchmark our protocol and others in a real-life scenario or something close to it, and to implement other ring signatures schemes to have a more complete comparison. In light of the recent deployment of the PQXDH protocol [KS23], it would be prudent also to benchmark this and compare it with our protocols, replacing our choice of Kyber-512 with Kyber-1024 to be consistent with PQXDH which uses the latter KEM (note that the non-standard primitive used for the deniable PQ protocols should nonetheless be more of a bottleneck).

One could also try to build one-time ring signatures that are both efficient and provably secure (possibly in the QROM). In turn, these could possibly be used to build efficient ephemeral split-KEMs. For instance, Scafuro and Zhang [SZ21] designed an efficient linkable one-time ring signatures from hash functions alone and proved the security of the scheme in the ROM. It would be of interest to understand whether split-KEMs can be built out of such a construction or a variant, and/or to prove security in the QROM. Different parameter sets for K-Waay to achieve higher levels of security could also be provided and benchmarked.

7.2.5 Active Attack Detection in Messaging (Chapter 4)

Firstly, it is of interest to evaluate our protocols/transformations and determine more practically the overhead that active attack detection incurs. We believe that our most practical scheme is the epoch-based s-RID-secure one (Section 4.6.3) since the overhead can reasonably be expected to be a single hash and a relatively small amount of indices in practice. By applying a comparable optimisation for r-RID in future work and with some benchmarking and evaluation, it may be deemed viable in some deployments to achieve r-RID security. Note also here that our lower bound for r-RID security considers a worst-case execution where one party is continually sending without receiving any messages from their counterpart.

A natural direction is to explore active attack detection in *group messaging*. In a first step, one could generalise the RECOVER/RID definitions to the group setting and in particular capture

Chapter 7. Conclusion

any additional properties that may only arise in the group setting. In a second step, as done in this thesis for the two-party case, it would be prudent to design practical schemes and explore different trade-offs between security and efficiency. One could leverage properties of a given group messaging scheme along the way – for example, it may be prudent to take advantage of the assumption of in-order delivery provided by CGKA control messages in MLS to avoid the inherent complexity of out-of-order delivery on the messaging layer. Rather than modelling pairwise two-party out-of-band channels, assuming the existence of a ‘global’ out-of-band channel may simplify protocol design and make more sense practically to capture something like a public ledger.

7.2.6 Group Administration (Chapter 5)

We consider possible extensions of A-CGKA as a primitive and corresponding construction ideas. We note that these extensions may provide stronger security guarantees, or additional functionality, at reduced cost if the number of admins is small.

Admins Beyond CGKA. CGKA is not a suitable formalism for some group messaging protocols used in practice like pairwise channels and Sender Keys (the latter by Signal and WhatsApp [BCG23b]). In these protocols, each user is associated with their own key or keying material rather than a common group secret. Nevertheless, an IAS-like protocol can be easily adapted to this setting. For Sender Keys, admins could replace their keying material at a low cost (a signature attesting to their new signing key) for PCS authentication guarantees. We leave it as useful future work to formalise group administration beyond CGKA.

Telegram, although not end-to-end encrypted, offers fine-grained administration features like message filtering and delays. Some of these could be conceivably implemented cryptographically, e.g., by entrusting admins to process messages or through NIZKs.

Private Admins. In some applications, it may be desirable to hide the set of admins from (non-admin) users within a group (or between themselves). DGS could achieve some notion of administrative privacy if the underlying admin CGKA provides privacy guarantees. IAS could be modified to achieve anonymity guarantees using ring signatures [RST01]. However, there is overhead with ring signatures over regular signatures, at a minimum to parse the anonymity set.

In MLS’ TreeKEM protocol, proposals are constant-sized, but commits are variable, which leaks information about the contents of the commit even if it is encrypted. Thus, padding is required at a minimum for privacy. In the MLS standard, ciphertexts (`PrivateMessage`) leak the group ID, epoch and message content type (proposal or commit) in plaintext¹, which need to be hidden for additional privacy. The work of Hashimoto et al. [HKP22] tackles this issue by hiding this information and further protecting against some additional leakage. In particular, their CGKA compiler allows group members to anonymously authenticate themselves, taking

¹<https://www.rfc-editor.org/rfc/rfc9420.html#name-encoding-and-decoding-a-pri>

advantage of the CGKA secret, like DGS, to this end. In practice, additional attack vectors like timing and traffic analysis preclude privacy also, which are considered by some messaging systems [TGL⁺17, CSM⁺20] that do not, however, provide FS and PCS; it remains open to, e.g., adapt CGKA to defend against these attack vectors. As discussed in the introduction, the Signal Private Group System [CPZ20] hides the group membership from non-members (although not all metadata); the mechanism could be extended for admins but adapting the technique to (A)-CGKA is also open. Recently, some steps towards preserving anonymity in messaging even under state exposure have been made [DHRR22, BRT23].

External Admins. Our A-CGKA constructions assume the admins comprise a subset of all group members, i.e., $G^* \subseteq G$. Some applications may be better suited for *external* administration. For example, an online platform may wish to control the set of conversation participants to ensure they are subscribers but nevertheless ensure they are provided confidentiality. External admins who then attempt to add users that group members do not trust can be detected on the protocol level, rather than the less well-defined application level as previously done.

Given that the underlying CGKA allows for external commits, it is straightforward to administrate IAS and DGS-based groups externally. Namely, the admin who is approving the change can inspect proposals and make commit messages for the corresponding parties. However, TreeKEM and its variants are not ideal for this since the committer is the party who derives new group secrets and can thus violate confidentiality. One way around this issue is to essentially write a wrapper around each CGKA algorithm which declares that some CGKA group members are not actually in the group. Here, the wrapper would also force admins to delete group secrets as soon as they derive them, and would not consider admins as part of the group; the solution is however clearly vulnerable to corrupted admins.

One conceptually simple solution is to allow commit messages from regular users which play the role of proposals which have to be “committed” (e.g. signed) by admins. Additional machinery like the use of NIZKs is however required in the malicious setting to enable admins to verify that such commits are well-formed.

Hierarchical Admins. In messaging apps like Telegram and WhatsApp, the group creator has stronger capabilities than other admins. For instance, the group creator can never be removed by another admin. Extending this concept, one can conceive a hierarchy of administrators of several levels, e.g., of the form $G^{**} \subseteq G^* \subseteq G$, where G^{**} are super-administrators. Extending IAS, one can imagine using signatures that attest to other signatures in a chain-of-trust fashion. DGS can be extended by considering many CGKAs where the $(i + 1)$ -th CGKA must sign commit messages for the i -th CGKA for each $i \geq 1$. Attribute-based admins would enable greater flexibility in the access structure.

Muting Admins. It is possible to provide some cryptographic guarantees to the process of muting conversation participants. Although we do not explicitly consider group messaging proper, we sketch how such a solution would look. One solution entails a DGS-like construction in which members must sign messages using a common signature key spk derived from

Chapter 7. Conclusion

a secondary CGKA; honest group members would then process application messages if and only if they are signed using the common signature key (i.e., only from the set of unmuted users). Then, muted members will be able to filter messages from other muted users (since they could still be informed of the state of spk over time), but they will not be able to sign their own messages. Mechanisms that enable a central server to filter messages while maintaining privacy [HS20] can also be integrated into the encryption layer over A-CGKA. Nevertheless, we note that in messaging services where the identity of the group members is known, muted members can generally bypass a ban by sending individual messages to all group members using two-party messages. At an application level, muting group members is a functionality supported by both Signal and Telegram.

Threshold Admins. One issue with our A-CGKA constructions is that security breaks down if a single admin is compromised. To improve the robustness of the protocol, a protocol can enforce that some $k > 1$ admins must attest to a particular commit before it may be processed, which can be achieved using threshold cryptography [Sho00].

Decentralised Admins. To allow for network decentralisation, it is straightforward in theory for a given messaging group G to simply execute a state machine replication protocol [CL⁺99] to order commit messages and require that users reliably broadcast [BT85] all proposal messages. Given that group members who are expected to execute the protocol on, e.g., mobile devices, may not be available often, thus leading to liveness (and possibly unintended safety) violations in protocol execution, a natural solution is to entrust administrators to provide messages to users. These admins could indeed execute consensus.

7.2.7 Sender Keys (Chapter 6)

Our security model, both for two-party channels and our Group Messenger primitive that captures Sender Keys and Sender Keys+, could be extended to encompass attack vectors like randomness manipulation, successful message injections and insider threats not explicitly captured in our modelling.

Investigating the practical behaviour of Sender Keys would provide valuable insights for improved modelling and the identification of potential vulnerabilities. In particular, Sender Keys is commonly supplemented by additional mechanisms not considered in our study, such as support for multiple devices and encrypted cloud backups (the latter considered in isolation recently by Davies et al. [DFG⁺23]) that increase the attack surface. Conversely, features of Signal such as sealed sender (by increasing privacy) and delivery receipts (by refreshing keys) can also improve security. Benchmarking both the baseline and extended Sender Keys protocols would also contribute to assessing their practicality.

Additionally, it is important to address the challenges that arise when total order is violated, and to design a protocol that avoids the drawbacks associated with decentralised continuous group key agreement (DCGKA) such as the need for multi-round communication [WKHB21].

Towards a more concurrency-friendly Sender Keys protocol, an important direction is the design of a mechanism for resolving ties in control messages that are sent concurrently. The way that this is achieved for group membership in Signal's private group system [CPZ20], as discussed in Section 5.4.2 in the context of group administration, is by the central server maintaining the canonical set of group members (bypassing the need for group members to agree between themselves). It could be of interest therefore to model such a system composed with Sender Keys, since the existing modelling of the private group system does not itself capture messaging proper as written (even given it is modelled in a variant of the universal composability framework [Can01]).

A Appendices

A.1 QROM Preliminaries

In this section, we assume the random oracles output values in $\{0, 1\}^n$ for some integer n . We first recall the notion of extractable random oracle simulator introduced by Don et al. [DFMS22] and the corresponding properties as presented by Huguenin-Dumittan and Vaudenay [HV22], and a useful lemma. We refer the reader to the original paper for more details.

Definition 50 (Extractable RO [DFMS22]). An extractable RO-simulator is a tuple (S, Ext) , where S is a compressed RO efficiently simulatable and Ext is the extractor, such that the following properties hold.

1. If the extractor is never called, the simulator is indistinguishable from a (standard) RO.
2. Any two subsequent independent queries to S commute.
3. Any two subsequent independent queries to Ext commute.
4. Any two subsequent independent queries to Ext and S $8\sqrt{2/2^n}$ -almost commute.
5. Querying classically the simulator S on the same value multiples times in a row has the same effect on the state of S as making *one* of these queries.
6. Let $x^e \leftarrow \text{Ext}(t)$ for some t , and $t' \leftarrow S(x^e)$ be two subsequent classical queries. Then,

$$\Pr[t \neq t' \wedge x^e \neq \perp] \leq 2/2^n .$$

7. Let $t \leftarrow S(x)$ for some x and $x^e \leftarrow \text{Ext}(t)$ be two subsequent classical queries. Then,

$$\Pr[\hat{x} = \perp] \leq 2/2^n .$$

<p style="margin: 0;">COLL(\mathcal{A})</p> <hr style="margin: 2px 0;"/> <p style="margin: 0;">1: $(x_1, t_1), \dots, (x_m, t_m) \leftarrow \mathcal{A}^S$</p> <p style="margin: 0;">2: for $i \in \{1, \dots, m\}$: $t'_i \leftarrow S(x_i)$</p> <p style="margin: 0;">3: for $i \in \{1, \dots, m\}$: $x_i^e \leftarrow \text{Ext}(t_i)$</p> <p style="margin: 0;">4: if $\exists i : x_i^e \neq x_i$ and $t_i = t'_i$:</p> <p style="margin: 0;">5: return 1</p> <p style="margin: 0;">6: return 0</p>

Figure A.1: Collision game for Property 8 Definition 50.

8. Let COLL be the game defined in Figure A.1. Then, for any \mathcal{A} we have

$$\Pr[\text{COLL}(\mathcal{A}) \Rightarrow 1] \leq \frac{40e^2(q+m+1)^3 + 2}{2^n},$$

where q is the number of queries \mathcal{A} makes to S and m is the number of tuples output by \mathcal{A} in the game.

Lemma 28 (Early Extraction). Let Γ be a game where an adversary runs on some input, queries a quantum RO H , and outputs two values t and o . Then, the game applies some deterministic function on o to obtain a value x and queries $h \leftarrow H(x)$. The game returns 1 if $h = t$. Now let Γ' be the same as Γ , except the extractor is called on t right after it is returned by \mathcal{A} , and the game returns 1 if $h = t$ and $x = x^*$, where x^* is the value extracted. Then,

$$\Pr[\Gamma \Rightarrow 1] - \Pr[\Gamma' \Rightarrow 1] \leq \frac{2}{2^n} + 8\sqrt{2/2^n} + \frac{40e^2(q_H+2)^3 + 2}{2^n}.$$

Proof. This follows from Corollary 4.7 in Don et al. and the fact that if $h = t$, where $h = H(x)$, then $\Pr[x^* = \perp] \leq \frac{2}{2^n}$. □

We also recall the algorithm one way to hiding lemma [HHK17].

Lemma 29 (AOW2H [HHK17]). Let \mathcal{A} be a quantum adversary making at most q_H queries to the QRO $H : \{0, 1\}^\ell \mapsto \{0, 1\}^n$ and outputting 0 or 1. Then, for any algorithm F that does not use H

$$\begin{aligned} & \left| \Pr[\mathcal{A}^H(x) \Rightarrow 1 : \sigma^* \xleftarrow{\$} \{0, 1\}^\ell; x \leftarrow F(\sigma^*, H(\sigma^*))] \right. \\ & \left. - \Pr[\mathcal{A}^H(x) \Rightarrow 1 : (\sigma^*, K) \xleftarrow{\$} \{0, 1\}^{\ell+n}; x \leftarrow F(\sigma^*, K)] \right| \\ & \leq 2q_H \sqrt{\Pr \left[\sigma^* \leftarrow E^{\mathcal{A}, H}(x) : \begin{array}{l} (\sigma^*, K) \xleftarrow{\$} \{0, 1\}^{\ell+n}; \\ x \leftarrow F(\sigma^*, K) \end{array} \right]}. \end{aligned}$$

where E is an algorithm that runs \mathcal{A} , measures the input register of a random query made to H , and outputs the result.

$S_{i^*}^{H, \mathcal{A}}(\Theta)$	$H'(x)$
1: $(j, b) \xleftarrow{\$} (\{0, \dots, q_H - 1\} \setminus \{i^*\} \times \{0, 1\}) \cup \{(q_H, 0)\}$	1: if $q = i^*$:
2: $q \leftarrow 1$	2: return Θ
3: $(x, z) \xleftarrow{\$} \mathcal{A}^{H'}$ and $x' \leftarrow$ measure \mathcal{A} 's $j + 1$ -th query input register	3: if $q < j + b + 1$:
4: return (x, z)	4: return $H(x)$
	5: else
	6: if $x = x'$: return Θ
	7: else : return $H(x)$

Figure A.2: Algorithm S for Lemma 30.

Finally, we recall the measure-and-reprogram lemma of Jiang et al. [JMZ23].

Lemma 30 (Lemma 3.1 [JMZ23]). Let $H : \{0, 1\}^m \mapsto \{0, 1\}^n$ be a quantum random oracle and \mathcal{A}^H be a quantum algorithm that makes q quantum queries to H and outputs (x, z) , where x and z are classical. Furthermore, we assume the i^* -th query of \mathcal{A} to H is classical and equal to x , for some $i^* \in [q_H]$. In addition, let $V(x, y, z)$ be some predicate s.t. $V(x, y, z) = 1$ implies that y was output on \mathcal{A} 's i^* -th query to H . Then, there exists an algorithm S_{i^*} (see Figure A.2), that takes some $\Theta \in \{0, 1\}^n$ as input and is such that

$$\Pr[V(x, H(x), z) = 1 : (x, z) \leftarrow \mathcal{A}^H] \leq 2(2q_H + 1)^2 \Pr[V(x, \Theta, z) = 1 : (x, z) \leftarrow S_{i^*}^{\mathcal{A}}(\Theta)] + \frac{8q_H^2}{2^n}$$

where the probabilities are taken over the randomness of the algorithms, the random oracle H and the sampling of Θ at random.

Informally, the previous lemma states that if some adversary \mathcal{A}^H can satisfy a predicate with probability p , one can build another algorithm $S^{\mathcal{A}}$ that does not query H on the i^* -th query (but uses its input instead) but that can satisfy with probability $\approx \frac{p}{q_H^2}$.

A.2 Proof of Theorem 5 (Chapter 3)

A.2.1 Proof in the QROM

We proceed with a sequence of games that is detailed in Figure A.3. The proof uses the extractable RO-simulator of Don et al. [DFMS22] (see Definition 50).

Game Γ_0 : This is the UNF-1KCA game with $\text{sKEM} := T_{\text{CH}}(\text{sKEM}_0)$ written explicitly. In addition, the RO used to compute the tag corresponding to ct (i.e. $t = H_1(\text{pk}_A, \text{pk}_B, \text{ct}, K_B)$) is different from the one used to compute the tag for ct' (i.e. $t_c = H_2(\text{pk}_A, \text{pk}_B, \text{ct}', \mathcal{K}_A)$). Note that since $(\text{pk}, \text{ct}) \neq (\text{pk}_A, \text{ct}')$ for the adversary to win, both oracles can be separated in this way.

Appendix A. Appendices

Game Γ_1 : The game is the same as the previous one, except we use the simulated RO for H_2 , and we use the extractor on t' (the tag output by the adversary) at the end. Note that this does not change anything to the probability of success of the game.

Game Γ_2 : Now the game outputs 0 if the values extracted are different than (pk_A, pk_B, ct', K_A) . For the game to return 1, t_c must be equal to t' , so let's assume it is the case. Hence, Γ^2 and Γ^1 differ only if $S.Ext(t_c) \neq (pk_A, pk_B, ct', K_A)$ and $H_2(pk_A, pk_B, ct', K_A) = t_c$. By Lemma 28, this happens with probability at most $\epsilon := \frac{2}{2^n} + 8\sqrt{2/2^n} + \frac{40e^2(q_{H'}+2)^3+2}{2^n}$. Hence, we have

$$\Pr[\Gamma_1] - \Pr[\Gamma_2] \leq \epsilon .$$

Game Υ_1 : We see that if an adversary \mathcal{A} wins Γ_2 , one can build an adversary \mathcal{B} that wins the game Υ_1 defined in Figure A.3. The reduction works simply by \mathcal{B} running \mathcal{A} , simulating H_2 with the simulated RO, and running the extractor on t' at the end. Therefore, we have

$$\Pr[\Gamma_2] \leq \Pr[\Upsilon_1] .$$

In addition, note that one can consider oracles H and H_1 as *one* oracle $H^* := H_1 \otimes H$ with images in $\{0, 1\}^{2n}$ that can be accessed $q_H + q_{H'}$ times by the adversary.

Game Υ_2 : We change the game such that (t, K) are picked at random and the oracle used is now \hat{H} instead of $H^* := H_1 \otimes H$. Now, let's consider a game \mathcal{C} that runs Γ^1 and outputs $(x = (pk, pk_B, ct, K_B), z = ((t, K), K_A, K'_A))$. In addition, let $V(x, y, z) := \mathbb{1}_{z_1=y} \wedge \mathbb{1}_{z_2 = z_3}$. Clearly, we have that

$$\Pr[\Upsilon_1] \leq \Pr[V(x, H^*(x), z) : (x, z) \xleftarrow{\$} \mathcal{C}^{H^*}]$$

as V is satisfied iff $K_A = K'_A$. Also, note that the condition $z_1 = y$ in the predicate is always satisfied by the definition of z_1 itself. Therefore, one can apply Lemma 30 with i^* equal to the query to H^* made by the game (i.e. $(t, K) \leftarrow H^*(pk, pk_B, ct, K_B)$) and we get

$$\Pr[\Upsilon_1] \leq 2(2(q_H + q_{H'}) + 1)^2 \Pr \left[V(x, (t, K), z) = 1 : (x, z) \leftarrow S_{i^*}^{\mathcal{A}}((t, K)) \right] + \frac{8(q_H + q_{H'})^2}{2^{2n}}$$

where (t, K) is sampled at random and S_{i^*} is the algorithm shown in Figure A.2. By inspection, one can see that if the output of S_{i^*} satisfies the predicate V then Υ_2 would output 1. Therefore, we have

$$\Pr[\Upsilon_1] \leq 2(2(q_H + q_{H'}) + 1)^2 \Pr \left[V(x, (t, K), z) = 1 : (x, z) \leftarrow S_{i^*}^{\mathcal{A}}((t, K)) \right] + \frac{8(q_H + q_{H'})^2}{2^{2n}} \leq \Pr[\Upsilon_2].$$

Finally, one can see that if \mathcal{A} wins Υ_2 , one can build an adversary \mathcal{B} s.t. \mathcal{B} wins the decaps-CPA game against $sKEM_0$. That is, the first phase of \mathcal{B} runs the first phase of \mathcal{A} and outputs the same public key pk . Then, in the second phase, \mathcal{B} runs $\mathcal{A}^{\hat{H}}$ with its own input (pk_A, pk_B, ct)

and random tag and key (t, K) . In addition, note that \mathcal{B} can perfectly simulate \hat{H} . Finally, \mathcal{B} outputs the same as the adversary \mathcal{A} . If $K_A = K'_A$ then \mathcal{B} wins the decaps-OW-CPA game. Hence, we have that

$$\Pr[Y_2] \leq \Pr[\text{decaps-CPA}_{\text{sKEM}_0}(\mathcal{B}) \Rightarrow 1].$$

Collecting the probabilities concludes the proof. \square

$\Gamma_0(\mathcal{A})$	$\Gamma_1(\mathcal{A})$	$\Gamma_2(\mathcal{A})$
1: $\text{pk}_A, \text{sk}_A \stackrel{\$}{\leftarrow} \text{KeyGenA}(1^\lambda)$	1: $\text{pk}_A, \text{sk}_A \stackrel{\$}{\leftarrow} \text{KeyGenA}(1^\lambda)$	1: $\text{pk}_A, \text{sk}_A \stackrel{\$}{\leftarrow} \text{KeyGenA}(1^\lambda)$
2: $\text{pk}_B, \text{sk}_B \stackrel{\$}{\leftarrow} \text{KeyGenB}(1^\lambda)$	2: $\text{pk}_B, \text{sk}_B \stackrel{\$}{\leftarrow} \text{KeyGenB}(1^\lambda)$	2: $\text{pk}_B, \text{sk}_B \stackrel{\$}{\leftarrow} \text{KeyGenB}(1^\lambda)$
3: $\text{pk}, \text{st} \stackrel{\$}{\leftarrow} \mathcal{A}^{H, H_1, H_2}(\text{pk}_A, \text{pk}_B)$	3: $\text{pk}, \text{st} \stackrel{\$}{\leftarrow} \mathcal{A}^{H, H_1, H_2}(\text{pk}_A, \text{pk}_B)$	3: $\text{pk}, \text{st} \stackrel{\$}{\leftarrow} \mathcal{A}^{H, H_1, H_2}(\text{pk}_A, \text{pk}_B)$
4: $(K_B, \text{ct}) \stackrel{\$}{\leftarrow} \text{Encaps}(\text{pk}, \text{sk}_B)$	4: $(K_B, \text{ct}) \stackrel{\$}{\leftarrow} \text{Encaps}(\text{pk}, \text{sk}_B)$	4: $(K_B, \text{ct}) \stackrel{\$}{\leftarrow} \text{Encaps}(\text{pk}, \text{sk}_B)$
5: $t \leftarrow H_1(\text{pk}, \text{pk}_B, \text{ct}, K_B)$	5: $t \leftarrow H_1(\text{pk}, \text{pk}_B, \text{ct}, K_B)$	5: $t \leftarrow H_1(\text{pk}, \text{pk}_B, \text{ct}, K_B)$
6: $K \leftarrow H(\text{pk}, \text{pk}_B, \text{ct}, K_B)$	6: $K \leftarrow H(\text{pk}, \text{pk}_B, \text{ct}, K_B)$	6: $K \leftarrow H(\text{pk}, \text{pk}_B, \text{ct}, K_B)$
7: $(\text{ct}', t') \stackrel{\$}{\leftarrow} \mathcal{A}^{H, H_1, H_2}(\text{pk}_A,$	7: $(\text{ct}', t') \stackrel{\$}{\leftarrow} \mathcal{A}^{H, H_1, H_2}(\text{pk}_A,$	7: $(\text{ct}', t') \stackrel{\$}{\leftarrow} \mathcal{A}^{H, H_1, H_2}(\text{pk}_A,$
8: $\text{pk}_B, (\text{ct}, t), K, \text{st})$	8: $\text{pk}_B, (\text{ct}, t), K, \text{st})$	8: $\text{pk}_B, (\text{ct}, t), K, \text{st})$
9: if $(\text{pk}_A, \text{ct}') = (\text{pk}, \text{ct})$: return 0	9: if $(\text{pk}_A, \text{ct}') = (\text{pk}, \text{ct})$: return 0	9: $(\text{pk}_0^*, \text{pk}_1^*, \text{ct}^*, K_A^*) \leftarrow \text{S.Ext}(t')$
10: $K_A \leftarrow \text{Decaps}(\text{pk}_B, \text{sk}_A, \text{ct}')$	10: $K_A \leftarrow \text{Decaps}(\text{pk}_B, \text{sk}_A, \text{ct}')$	10: if $(\text{pk}_A, \text{ct}') = (\text{pk}, \text{ct})$: return 0
11: $t_c \leftarrow H_2(\text{pk}_A, \text{pk}_B, \text{ct}', K_A)$	11: $t_c \leftarrow H_2(\text{pk}_A, \text{pk}_B, \text{ct}', K_A)$	11: $K_A \leftarrow \text{Decaps}(\text{pk}_B, \text{sk}_A, \text{ct}')$
12: if $t_c \neq t'$: return 0	12: if $t_c \neq t'$: return 0	12: $t_c \leftarrow H_2(\text{pk}_A, \text{pk}_B, \text{ct}', K_A)$
13: return 1	13: $(\text{pk}_1^*, \text{pk}_2^*, \text{ct}^*, K_A^*) \leftarrow \text{S.Ext}(t')$	13: if $t_c \neq t'$: return 0
	14: return 1	14: if $(\text{pk}_1^*, \text{pk}_2^*, \text{ct}^*, K_A^*)$ $\neq (\text{pk}_A, \text{pk}_B, \text{ct}', K_A)$:
		15: return 0
		16: return 1
$\Upsilon_1(\mathcal{A})$	$\Upsilon_2(\mathcal{A})$	$\hat{H}(x)$
1: $\text{pk}_A, \text{sk}_A \stackrel{\$}{\leftarrow} \text{KeyGenA}(1^\lambda)$	1: $(j, b) \stackrel{\$}{\leftarrow} (\{0, \dots, q_H - 1\} \times \{0, 1\}) \cup \{(q_H, 0)\}$	1: $q \leftarrow q + 1$
2: $\text{pk}_B, \text{sk}_B \stackrel{\$}{\leftarrow} \text{KeyGenB}(1^\lambda)$	2: $x' \leftarrow \text{measure } \mathcal{A}'\text{'s } j + 1\text{-th query input register}$	2: if $q < j + b + 1$:
3: $\text{pk}, \text{st} \stackrel{\$}{\leftarrow} \mathcal{A}^{H, H_1}(\text{pk}_A, \text{pk}_B)$	3: $q \leftarrow 0$	3: return $H^*(x)$
4: $(K_B, \text{ct}) \stackrel{\$}{\leftarrow} \text{Encaps}(\text{pk}, \text{sk}_B)$	4: $\text{pk}_A, \text{sk}_A \stackrel{\$}{\leftarrow} \text{KeyGenA}(1^\lambda)$	4: else
5: $(t, K) \leftarrow H^*(\text{pk}, \text{pk}_B, \text{ct}, K_B)$	5: $\text{pk}_B, \text{sk}_B \stackrel{\$}{\leftarrow} \text{KeyGenB}(1^\lambda)$	5: if $x = x'$:
6: $\text{in} \leftarrow (\text{pk}_A, \text{pk}_B, (\text{ct}, t), K, \text{st})$	6: $(K_B, \text{ct}) \stackrel{\$}{\leftarrow} \text{Encaps}(\text{pk}_A, \text{sk}_B)$	6: return (t, K)
7: $K'_A, \text{ct}' \stackrel{\$}{\leftarrow} \mathcal{A}^{H, H_1}(\text{in})$	7: $\text{pk}, \text{st} \stackrel{\$}{\leftarrow} \mathcal{A}^{\hat{H}}(\text{pk}_A, \text{pk}_B)$	7: else :
8: $K_A \leftarrow \text{Decaps}(\text{pk}_B, \text{sk}_A, \text{ct}')$	8: $(K_B, \text{ct}) \stackrel{\$}{\leftarrow} \text{Encaps}(\text{pk}, \text{sk}_B)$	8: return $H^*(x)$
9: if $(\text{pk}_A, \text{ct}') = (\text{pk}, \text{ct})$ or $K_A \neq K'_A$:	9: $(t, K) \stackrel{\$}{\leftarrow} \{0, 1\}^{2n}$	
10: return 0	10: $K'_A, \text{ct}' \stackrel{\$}{\leftarrow} \mathcal{A}^{\hat{H}}(\text{pk}_A, \text{pk}_B, (\text{ct}, t), K, \text{st})$	
11: return 1	11: $K_A \leftarrow \text{Decaps}(\text{pk}_B, \text{sk}_A, \text{ct}')$	
	12: if $K_A \neq K'_A$:	
	13: return 0	
	14: return 1	

Figure A.3: Sequence of games for the proof of Theorem 5. H^* is defined as $H_1 \otimes H_1$.

Appendix A. Appendices

A.2.2 Proof in the ROM

The proof follows a similar approach as the one in the QROM.

Game Γ_0 : This is the same as the UNF-1KCA game with $\text{sKEM} = T_{\text{CH}}(\text{sKEM}_0)$, except we assume there is no collision on H' . Thus, Γ^0 is the same as UNF-1KCA except with probability at most $\frac{q_{H'}^2}{2^n}$.

Game Γ_1 : In this game, we return 0 if \mathcal{A} did not query $H'(\text{pk}_A, \text{pk}_B, \text{ct}', K_A)$. As we can assume $(\text{pk}, \text{ct}) \neq (\text{pk}_A, \text{ct}')$, this changes the probability of \mathcal{A} winning only if \mathcal{A} outputs $t' = H'(\text{pk}_A, \text{pk}_B, \text{ct}', K_A)$ without having made the oracle query. Since the query was not made, one can actually lazy sample the value of $H'(\text{pk}_A, \text{pk}_B, \text{ct}', K_A)$ after \mathcal{A} returns t' , and the probability both values are equal is $\frac{1}{2^n}$. Hence,

$$\Pr[\Gamma_0] - \Pr[\Gamma_1] \leq \frac{1}{2^n}.$$

Game Υ_1 : If Γ_1 outputs 1, it means \mathcal{A} outputs (ct, t') s.t. $((\text{pk}_A, \text{pk}_B, \text{ct}', K_A), t')$ is in the list of queries made by the \mathcal{A} . Hence, if that happens, one can find ct' and K_A s.t. $\text{Decaps}(\text{pk}_B, \text{sk}_A, \text{ct}') = K_A$ by running $(\text{ct}', t') \xleftarrow{\$} \mathcal{A}$ and looking for t' in the list of queries (note that we assume there is no collision). Therefore, it means one can build an adversary that wins the game Υ_1 in Fig. A.3, and we have

$$\Pr[\Gamma_1] \leq \Pr[\Upsilon_1].$$

Game Υ_2 : We modify the game s.t. the tag t and the key given to the adversary are picked uniformly at random as shown in Figure A.3. Both games are indistinguishable unless \mathcal{A} queries $(\text{pk}, \text{pk}_B, \text{ct}, K_B)$ to H or H' . Then, an adversary \mathcal{B} playing Υ_2 can perfectly simulate \mathcal{A} 's view in Υ_1 if it guesses correctly which query it is going to be and if such a query is going to happen. Overall, \mathcal{B} can make a correct guess with probability $\frac{1}{q_{H'} + q_H + 1}$. If that happens though, one can build an OW-CPA adversary \mathcal{B} against sKEM_0 that runs \mathcal{A} and picks a random query made by \mathcal{A} to H or H' . Hence, we have

$$\Pr[\Upsilon_1] \leq (q_{H'} + q_H + 1) \Pr[\Upsilon_2].$$

Finally, one can see that Υ_2 is the same as the decaps-CPA for sKEM_0 if we omit the random values K and t and the more restrictive winning condition $(\text{pk}_A, \text{ct}') \neq (\text{pk}, \text{ct})$. Hence, one can build an adversary \mathcal{C} such that

$$\Pr[\Upsilon_2] \leq \Pr[\text{decaps-CPA}_{\text{sKEM}_0}(\mathcal{C}) \Rightarrow 1].$$

$\mathcal{O}(\mathcal{L}_{H'}, \{(\text{pk}_i, (\text{ct}_i, t_i))\}_{i=1}^d)$
<pre> 1: for $i \in [d]$: 2: $K'_i \leftarrow \text{Decaps}(\text{pk}_i, \text{sk}_A, \text{ct}_i)$ 3: if $((\text{pk}_i, \text{ct}_i, K'_i), t_i) \notin \mathcal{L}_{H'}$: return 0 4: return 1 </pre>

Figure A.4: Oracle \mathcal{O} used in the proof of Theorem 6.

□

A.3 Proof of Theorem 6 (Chapter 3)

A.3.1 Proof in the ROM

Proof. The idea of the proof is very similar to the IND-qCCA proof of the T_{CH} transform by Huguenin-Dumittan et al. [HV22] and is the following. Either all tags in the decapsulation query are valid and thus they are the form $H'(\text{pk}_A, \text{pk}_i, \text{ct}_i, K'_i)$, or the oracle returns \perp . Then, if they are valid, with high probability the adversary queried $(\text{pk}_A, \text{pk}_i, \text{ct}_i, K'_i)$ to H' and thus K'_i can be recovered from the list of queries to the RO, i.e. the decapsulation oracle can be simulated without the knowledge of sk_A . In other words, the only information leaked by a query to the decapsulation oracle is whether *all* tags are valid or not, i.e. 1 bit of information, which is not sufficient to break the OW-CPA game. We prove this formally with a sequence of hybrid games.

Game Γ_0 : This is the IND-1BatchCCA game with $\text{sKEM} = T_{\text{CH}}(\text{sKEM}_0)$.

Game Γ_1 : We modify the previous game s.t. we abort if the adversary finds any collision when querying H' . We have that

$$\Pr[\Gamma_0] - \Pr[\Gamma_1] \leq \frac{q_H'^2}{2^n}$$

where q_H' is the number of queries the adversary makes to H' .

Game Γ_2 : We modify the game s.t. it aborts if $\text{BatchDEC}(\{(\text{pk}_i, (\text{ct}_i, t_i))\}_{i=1}^d)$ does not return \perp but one of the tags t_i was not obtained through an adversary's query to H' . The probability that some tag t_i is valid but $H'(\text{pk}_A, \text{pk}_i, \text{ct}_i, K'_i)$ (with $(\text{pk}_A, \text{pk}_i, \text{ct}_i, t_i) \neq (\text{pk}_A, \text{pk}_B, \text{ct}^*, t_i^*)$) was

Appendix A. Appendices

not queried by the adversary is $\frac{1}{2^n}$. Hence, overall we have

$$\Pr[\Gamma_1] - \Pr[\Gamma_2] \leq \frac{d}{2^n}$$

Game Γ_3 : We now change the game as follows. We record all queries to H' of the form $(pk_A, \cdot, \cdot, \cdot)$ made by the adversary in a list $\mathcal{L}_{H'} = \{((pk_j, ct_j, K_j), h_j)\}_{j=1}^{q_{H'}}$ s.t. $H'(pk_A, pk_j, ct_j, K_j) = h_j$ for all $j \in [q_{H'}]$. Then, the BatchDEC oracle is modified as follows. If some tag t_i is s.t. for all $K \in \mathcal{K}$ $((pk_i, ct_i, K), t_i) \notin \mathcal{L}_{H'}$ then \perp is returned. Then, $\mathcal{O}(\mathcal{L}_{H'}, \{(pk_i, (ct_i, t_i))\}_{i=1}^d) \rightarrow r$ is queried, where \mathcal{O} is defined in Figure A.4. If $r = 0$ BatchDEC outputs \perp , otherwise it outputs $H(pk_A, pk_i, ct_i, K_i)$ for all $i \in [d]$, where K_i is s.t. $((pk_i, ct_i, K_i), t_i) \in \mathcal{L}_{H'}$. Note that all these modifications are only syntactical as \mathcal{O} outputs 1 iff for all $i \in [d]$, K_i is (the unique) value in $\mathcal{L}_{H'}$ s.t. $H'(pk_A, pk_i, ct_i, K_i := \text{Decaps}(pk_i, sk_A)) = t_i$. Hence, we have

$$\Pr[\Gamma_2] = \Pr[\Gamma_3].$$

Game Γ_4 : We replace the challenge tag t^* and the real key K_0 by random values. This change can only be noticed if the adversary or the BatchDEC oracle queries $H(pk_A, pk_B, ct^*, K^*)$ or $H'(pk_A, pk_B, ct^*, K^*)$ at some point in the game. Let QUERY be this event. We show that if QUERY occurs, then one can break the OW-CPA security of sKEM_0 with high probability. The reduction works as follows. The OW-CPA adversary \mathcal{B} receives a challenge ciphertext ct^* and public keys pk_A, pk_B from its own challenger. Next, it samples random values K, t^* and passes all these to the IND-1BatchCCA adversary \mathcal{A} . Then, \mathcal{B} can simulate everything in BatchDEC (except the oracle call to \mathcal{O}) by recording \mathcal{A} 's queries to H' . In order to simulate \mathcal{O} , \mathcal{B} samples a bit r at random instead, which succeeds with probability $\frac{1}{2}$. Finally, it samples at random a query made by \mathcal{A} to H or H' or a query made to H by itself, and it outputs the key K that was part of this query. Overall, the simulation is correct with probability $\frac{1}{2}$ and if QUERY occurs \mathcal{B} recovers K^* with probability $\frac{1}{q_H + q_{H'} + d}$. Hence,

$$\Pr[\Gamma_3] - \Pr[\Gamma_4] \leq \Pr[\text{QUERY}] \leq 2(q_H + q_{H'} + d) \text{Adv}_{\text{sKEM}}^{\text{ow-cpa}}(\mathcal{A}).$$

Finally, we see that the adversary's view is independent of b in Γ_4 , therefore $\Pr[\Gamma_4] = \frac{1}{2}$. This concludes the proof. \square

A.3.2 Proof in the QROM

Proof. As in the proof of Theorem 5, we use the extractable RO-simulator by Don et al. [DFMS22] and we proceed with a sequence of hybrid games. We note the proof is nearly identical to the QROM IND-qCCA proof of T_{CH} [HV22] and we refer the reader to the latter for

a detailed explanation of the game transitions.

Game Γ_0 : This is the IND-1BatchCCA game with $\text{sKEM} = \text{T}_{\text{CH}}(\text{sKEM}_0)$. We also assume that the adversary only makes queries of the form $(\text{pk}_A, \cdot, \cdot, \cdot)$ to the oracles. This has no consequence on the winning probability of the adversary as other type of queries are independent of the key.

Game Γ_1 : We modify the BatchDEC oracle s.t. it returns \perp whenever the list of $(\text{pk}_i, \text{ct}_i, t_i)$ in the query contains $(\text{pk}_i, \text{ct}_i) = (\text{pk}_B, \text{ct}^*)$ (and thus $t_i \neq t^*$). This change has no impact except if $\text{Decaps}(\text{pk}_B, \text{sk}_A, \text{ct}^*) \neq K_0$, where K_0 is the challenge real key. In turn, this would imply that ct^* is an incorrect ciphertext. Hence,

$$\Pr[\Gamma_0] - \Pr[\Gamma_1] \leq \delta .$$

Game Γ_2 : Now, we split the random oracle H' into two oracles H'_0 and H'_1 s.t.

$$H'(\text{pk}_A, \text{pk}, \text{ct}, K) := \begin{cases} H'_0(K), & \text{if } (\text{pk}, \text{ct}) = (\text{pk}_B, \text{ct}^*) \\ H'_1(\text{pk}, \text{ct}, K), & \text{otherwise} \end{cases}$$

and we give the adversary access to H'_0, H'_1 instead of H' . We also switch to the RO simulator instead of using H'_1 . In addition, at the end of the game, the challenger calls the extractor on all tags t_i queried as part of the call to the BatchDEC oracle to obtain extracted values $(\text{pk}_i^e, \text{ct}_i^e, K_i^e)$, $i \in [d]$. Note that H'_0 is never called as part of a BatchDEC query due to the modification in the previous game. These changes have no impact on the success of the game and thus

$$\Pr[\Gamma_1] = \Pr[\Gamma_2] .$$

Game Γ_3 : We abort whenever the decapsulation oracle does not return \perp but the extracted values $(\text{pk}_i^e, \text{ct}_i^e, K_i^e)$ are not equal to \perp or $(\text{pk}_i, \text{ct}_i, K_i')$, where $K_i' = \text{Decaps}(\text{pk}_i, \text{sk}_A, \text{ct}_i)$. By Property 8 of the extractable oracle, we have

$$\Pr[\Gamma_2] - \Pr[\Gamma_3] \leq \frac{40e^2(q_{H'} + d + 1)^3 + 2}{2^n} .$$

Game Γ_4 : We move the extraction to the BatchDEC oracle, right after the corresponding tag verification. By Property 4 of the extractable oracle, we have

$$\Pr[\Gamma_3] - \Pr[\Gamma_4] \leq 8d(d + q_{H'})\sqrt{2/2^n} .$$

Appendix A. Appendices

Game Γ_5 : We modify the BatchDEC oracle s.t. we abort if all tag checks pass, i.e. $H'(\text{pk}_i, \text{ct}_i, K'_i) = t_i, \forall i \in [d]$ but some extracted value is equal to \perp , i.e. $(\text{pk}_i^e, \text{ct}_i^e, K_i^e) = \perp$ for some $i \in [d]$. By Property 7 of the extractable oracle we have

$$\Pr[\Gamma_4] - \Pr[\Gamma_5] \leq d \frac{2}{2^n}.$$

Game Γ_6 : We modify the BatchDEC oracle s.t. the queries to H' made for the tag verification are made after the corresponding extraction. By Property 8 of the extractable oracle we have

$$\Pr[\Gamma_5] - \Pr[\Gamma_6] \leq 8d\sqrt{2/2^n}.$$

Game Γ_7 : We modify the previous game as follows. Let r be a bit set to 1 iff for all $i \in [d]$ $(\text{pk}_i^e, \text{ct}_i^e) = (\text{pk}_i, \text{ct}_i)$ and $\text{Decaps}(\text{pk}_i^e, \text{sk}_A, \text{ct}_i^e) = K_i^e$. Then, we change BatchDEC s.t. it returns \perp if $r = 0$, otherwise it returns $H(\text{pk}_A, \text{pk}_i, \text{ct}_i, K_i^e)$ for all $i \in [d]$. In addition, the tag verification is now skipped. We argue this affects only negligibly the advantage of the adversary compared to the previous game:

- If BatchDEC returns $H(\text{pk}_A, \text{pk}_i, \text{ct}_i, K'_i), i \in [d]$ in Γ_6 , then by the previous changes we know that $(\text{pk}_i^e, \text{ct}_i^e, K_i^e) = (\text{pk}_i, \text{ct}_i, K'_i)$ for all $i \in [d]$, therefore BatchDEC returns $H(\text{pk}_A, \text{pk}_i, \text{ct}_i, K'_i), i \in [d]$ in Γ_7 as well.
- If BatchDEC returns $H(\text{pk}_A, \text{pk}_i, \text{ct}_i, K_i^e), i \in [d]$ in Γ_7 , we know that $(\text{pk}_i^e, \text{ct}_i^e, K_i^e) = (\text{pk}_i, \text{ct}_i, K'_i)$. In addition, for each $i \in [d]$, $t_i = H(\text{pk}_i^e, \text{ct}_i^e, K_i^e)$ with probability $1 - \frac{2}{2^n}$ by Property 6 of the extractable oracle. Therefore, the tag verification would pass in Γ_6 with high probability and BatchDEC would return the same values in that game as well.

Overall, we have

$$\Pr[\Gamma_6] - \Pr[\Gamma_7] \leq d \frac{2}{2^n}.$$

Game Γ_8 : Now we move all d queries to H' made in BatchDEC to the end of the game. By Property 8 of the extractable oracle, we have

$$\Pr[\Gamma_7] - \Pr[\Gamma_8] \leq 8dq_{H'}\sqrt{2/2^n}.$$

Note that we can now forget about the queries to H' we just moved to the end of the game.

Game Γ_9 : We replace the real key K_0 and the challenge tag t^* by random values. We have

$\Pr[\Gamma_9] = \frac{1}{2}$. Applying the OW2H lemma on $H \otimes H'_0$, we get

$$\Pr[\Gamma_8] - \Pr[\Gamma_9] \leq 2(q_{H'} + q_H + d)\sqrt{\Pr[\Upsilon]}$$

where Υ is the same as Γ_9 , except the challenger measures a random query made to $H \otimes H'_0$ and outputs 1 iff the query contains K^* , where K^* is the key encapsulated in ct^* . We can build an OW-CPA adversary \mathcal{B} against sKEM_0 that wins with high probability when Υ outputs 1. The reduction works nearly as in the ROM proof: \mathcal{B} receives a challenge ciphertext ct^* and two public keys pk_A, pk_B , then it samples t^* and K^* at random and passes all these values to \mathcal{A} . Then, \mathcal{B} can perfectly simulate BatchDEC as in Γ_9 , except for the bit r that it can guess correctly with probability $\frac{1}{2}$. Finally, \mathcal{B} measures a random query that was made to H or H' in the execution and outputs the corresponding value K . Overall, we have

$$\Pr[\Upsilon] \leq 2\text{Adv}_{\text{sKEM}}^{\text{ow-cpa}}(\mathcal{A})$$

which concludes the proof. □

A.4 Tables for Sender Keys (Chapter 6)

We provide here two tables (Table A.1 and Table A.2) that serve as a reference for the variables that we use in Chapter 6.

Variable	Description
General	
G	sender keys group
γ	state
ID	user
ME	caller
m	application message
T	control message
(e, i)	message epoch counter
H	random oracle
Keys	
SK	sender key
spk	public signature key
ssk	secret signature key
ck	symmetric chain key
mk	message key
τk	MAC key
Game Dictionaries	
M	ciphertexts
CH	challenged ciphertexts
T	control messages
E	message epochs
I	message indices
SM	skipped messages

Table A.1: Summary of variables used throughout Chapter 6 (part I).

Variable	Description
Protocol	
add	add command
crt	create command
rem	remove command
upd	update command
acc	acceptance bit
welcome	welcome ciphertext
ep	current epoch
i_{ck}	current index of user's chain key
i_{ME}	current epoch-specific index of caller's chain key
last-kc	last chain key
$\max-i_{ck}$	maximum index of chain key
kc	chain key counter
no-SK	boolean indicator of whether sender key exists
rs	self-sampled randomness (for updates)
r	randomness sent (for updates)
\vec{C}	ciphertext vector
M	message
τ	MAC tag
σ	signature
C	ciphertext
γ_{2pc}	two-party channel state
C_{2pc}	two-party channel ciphertext
Upd-Ind	view of update initiator

Table A.2: Summary of variables used throughout throughout Chapter 6 (part II).

Bibliography

- [AAB⁺21] Joël Alwen, Benedikt Auerbach, Mirza Ahad Baig, Miguel Cueto Noval, Karen Klein, Guillermo Pascual-Perez, Krzysztof Pietrzak, and Michael Walter. Grafting key trees: Efficient key management for overlapping groups. In Kobbi Nissim and Brent Waters, editors, *TCC 2021: 19th Theory of Cryptography Conference, Part III*, volume 13044 of *Lecture Notes in Computer Science*, pages 222–253, Raleigh, NC, USA, November 8–11, 2021. Springer, Heidelberg, Germany.
- [AAN⁺22a] Joël Alwen, Benedikt Auerbach, Miguel Cueto Noval, Karen Klein, Guillermo Pascual-Perez, and Krzysztof Pietrzak. DeCAF: Decentralizable continuous group key agreement with fast healing. *Cryptology ePrint Archive*, Report 2022/559, 2022. <https://eprint.iacr.org/2022/559>.
- [AAN⁺22b] Joël Alwen, Benedikt Auerbach, Miguel Cueto Noval, Karen Klein, Guillermo Pascual-Perez, Krzysztof Pietrzak, and Michael Walter. CoCoA: Concurrent continuous group key agreement. In Orr Dunkelman and Stefan Dziembowski, editors, *Advances in Cryptology – EUROCRYPT 2022, Part II*, volume 13276 of *Lecture Notes in Computer Science*, pages 815–844, Trondheim, Norway, May 30 – June 3, 2022. Springer, Heidelberg, Germany.
- [ABH⁺21] Joël Alwen, Bruno Blanchet, Eduard Hauck, Eike Kiltz, Benjamin Lipp, and Doreen Riepel. Analysing the HPKE standard. In Anne Canteaut and François-Xavier Standaert, editors, *Advances in Cryptology – EUROCRYPT 2021, Part I*, volume 12696 of *Lecture Notes in Computer Science*, pages 87–116, Zagreb, Croatia, October 17–21, 2021. Springer, Heidelberg, Germany.
- [ACD19] Joël Alwen, Sandro Coretti, and Yevgeniy Dodis. The double ratchet: Security notions, proofs, and modularization for the Signal protocol. In Yuval Ishai and Vincent Rijmen, editors, *Advances in Cryptology – EUROCRYPT 2019, Part I*, volume 11476 of *Lecture Notes in Computer Science*, pages 129–158, Darmstadt, Germany, May 19–23, 2019. Springer, Heidelberg, Germany.
- [ACD]23] Martin R. Albrecht, Sofia Celi, Benjamin Dowling, and Daniel Jones. Practically-exploitable Cryptographic Vulnerabilities in Matrix. In *2023 IEEE Symposium on Security and Privacy*, 2023.

Bibliography

- [ACDT19] Joël Alwen, Sandro Coretti, Yevgeniy Dodis, and Yiannis Tselekounis. Security analysis and improvements for the IETF MLS standard for group messaging. Cryptology ePrint Archive, Report 2019/1189, 2019. <https://eprint.iacr.org/2019/1189>.
- [ACDT20] Joël Alwen, Sandro Coretti, Yevgeniy Dodis, and Yiannis Tselekounis. Security analysis and improvements for the IETF MLS standard for group messaging. In Daniele Micciancio and Thomas Ristenpart, editors, *Advances in Cryptology – CRYPTO 2020, Part I*, volume 12170 of *Lecture Notes in Computer Science*, pages 248–277, Santa Barbara, CA, USA, August 17–21, 2020. Springer, Heidelberg, Germany.
- [ACDT21a] Joël Alwen, Sandro Coretti, Yevgeniy Dodis, and Yiannis Tselekounis. Modular design of secure group messaging protocols and the security of MLS. In Giovanni Vigna and Elaine Shi, editors, *ACM CCS 2021: 28th Conference on Computer and Communications Security*, pages 1463–1483, Virtual Event, Republic of Korea, November 15–19, 2021. ACM Press.
- [ACDT21b] Joël Alwen, Sandro Coretti, Yevgeniy Dodis, and Yiannis Tselekounis. Modular design of secure group messaging protocols and the security of MLS. Cryptology ePrint Archive, Report 2021/1083, 2021. <https://eprint.iacr.org/2021/1083>.
- [ACJM20] Joël Alwen, Sandro Coretti, Daniel Jost, and Marta Mularczyk. Continuous group key agreement with active security. In Rafael Pass and Krzysztof Pietrzak, editors, *TCC 2020: 18th Theory of Cryptography Conference, Part II*, volume 12551 of *Lecture Notes in Computer Science*, pages 261–290, Durham, NC, USA, November 16–19, 2020. Springer, Heidelberg, Germany.
- [ACPS09] Benny Applebaum, David Cash, Chris Peikert, and Amit Sahai. Fast cryptographic primitives and circular-secure encryption based on hard learning problems. In Shai Halevi, editor, *Advances in Cryptology – CRYPTO 2009*, volume 5677 of *Lecture Notes in Computer Science*, pages 595–618, Santa Barbara, CA, USA, August 16–20, 2009. Springer, Heidelberg, Germany.
- [ADJ24] Martin R. Albrecht, Benjamin Dowling, and Daniel Jones. Device-Oriented Group Messaging: A Formal Cryptographic Analysis of Matrix’ Core. In *2024 IEEE Symposium on Security and Privacy (to appear)*, 2024.
- [AHKM22] Joël Alwen, Dominik Hartmann, Eike Kiltz, and Marta Mularczyk. Server-aided continuous group key agreement. In Heng Yin, Angelos Stavrou, Cas Cremers, and Elaine Shi, editors, *ACM CCS 2022: 29th Conference on Computer and Communications Security*, pages 69–82, Los Angeles, CA, USA, November 7–11, 2022. ACM Press.

- [AJM20] Joël Alwen, Daniel Jost, and Marta Mularczyk. On the insider security of MLS. Cryptology ePrint Archive, Report 2020/1327, 2020. <https://eprint.iacr.org/2020/1327>.
- [AJM22] Joël Alwen, Daniel Jost, and Marta Mularczyk. On the insider security of MLS. In Yevgeniy Dodis and Thomas Shrimpton, editors, *Advances in Cryptology – CRYPTO 2022, Part II*, volume 13508 of *Lecture Notes in Computer Science*, pages 34–68, Santa Barbara, CA, USA, August 15–18, 2022. Springer, Heidelberg, Germany.
- [AMPS22] Martin R. Albrecht, Lenka Mareková, Kenneth G. Paterson, and Igors Stepanovs. Four attacks and a proof for Telegram. In *2022 IEEE Symposium on Security and Privacy*, pages 87–106, San Francisco, CA, USA, May 22–26, 2022. IEEE Computer Society Press.
- [AMT23] Joël Alwen, Marta Mularczyk, and Yiannis Tselekounis. Fork-resilient continuous group key agreement. In Helena Handschuh and Anna Lysyanskaya, editors, *Advances in Cryptology – CRYPTO 2023, Part IV*, volume 14084 of *Lecture Notes in Computer Science*, pages 396–429, Santa Barbara, CA, USA, August 20–24, 2023. Springer, Heidelberg, Germany.
- [ANPPP23] Benedikt Auerbach, Miguel Cueto Noval, Guillermo Pascual-Perez, and Krzysztof Pietrzak. On the cost of post-compromise security in concurrent continuous group-key agreement. In *TCC 2023: 21st Theory of Cryptography Conference, Part I*, *Lecture Notes in Computer Science*. Springer, Heidelberg, Germany, November 2023.
- [ANS23] ANSSI. ANSSI views on the Post-Quantum Cryptography transition (2023 follow up). https://www.ssi.gouv.fr/uploads/2023/09/follow_up_position_paper_on_post_quantum_cryptography.pdf, 2023. Accessed: 28-09-2023.
- [AP12] Jacob Alperin-Sheriff and Chris Peikert. Circular and KDM security for identity-based encryption. In Marc Fischlin, Johannes Buchmann, and Mark Manulis, editors, *PKC 2012: 15th International Conference on Theory and Practice of Public Key Cryptography*, volume 7293 of *Lecture Notes in Computer Science*, pages 334–352, Darmstadt, Germany, May 21–23, 2012. Springer, Heidelberg, Germany.
- [APS15] Martin R. Albrecht, Rachel Player, and Sam Scott. On the concrete hardness of learning with errors. Cryptology ePrint Archive, Report 2015/046, 2015. <https://eprint.iacr.org/2015/046>.
- [Bal21] David Balbás. On Secure Administrators for Group Messaging Protocols, 2021. MSc Thesis.

Bibliography

- [BBL⁺22] Olivier Blazy, Ioana Boureanu, Pascal Lafourcade, Cristina Onete, and Léo Robert. How fast do you heal? A taxonomy for post-compromise security in secure-channel establishment. *Cryptology ePrint Archive*, Report 2022/1090, 2022. <https://eprint.iacr.org/2022/1090>.
- [BBLW22] Richard Barnes, Karthikeyan Bhargavan, Benjamin Lipp, and Christopher A. Wood. Hybrid Public Key Encryption. Technical Report 9180, February 2022.
- [BBR18] Karthikeyan Bhargavan, Richard Barnes, and Eric Rescorla. TreeKEM: Asynchronous Decentralized Key Management for Large Dynamic Groups A protocol proposal for Messaging Layer Security (MLS). Research report, Inria Paris, May 2018.
- [BBR⁺23] Richard Barnes, Benjamin Beurdouche, Raphael Robert, Jon Millican, Emad Omara, and Katriel Cohn-Gordon. The Messaging Layer Security (MLS) Protocol. RFC 9420, July 2023.
- [BCC⁺23a] Khashayar Barooti, Daniel Collins, Simone Colombo, Loïs Huguenin-Dumittan, and Serge Vaudenay. On active attack detection in messaging with immediate decryption. In Helena Handschuh and Anna Lysyanskaya, editors, *Advances in Cryptology – CRYPTO 2023, Part IV*, volume 14084 of *Lecture Notes in Computer Science*, pages 362–395, Santa Barbara, CA, USA, August 20–24, 2023. Springer, Heidelberg, Germany.
- [BCC⁺23b] Khashayar Barooti, Daniel Collins, Simone Colombo, Loïs Huguenin-Dumittan, and Serge Vaudenay. On active attack detection in messaging with immediate decryption. *Cryptology ePrint Archive*, Report 2023/880, 2023. <https://eprint.iacr.org/2023/880>.
- [BCD⁺16] Joppe W. Bos, Craig Costello, Léo Ducas, Ilya Mironov, Michael Naehrig, Valeria Nikolaenko, Ananth Raghunathan, and Douglas Stebila. Frodo: Take off the ring! Practical, quantum-secure key exchange from LWE. In Edgar R. Weippl, Stefan Katzenbeisser, Christopher Kruegel, Andrew C. Myers, and Shai Halevi, editors, *ACM CCS 2016: 23rd Conference on Computer and Communications Security*, pages 1006–1018, Vienna, Austria, October 24–28, 2016. ACM Press.
- [BCG22] David Balbás, Daniel Collins, and Phillip Gajland. Analysis and Improvements of the Sender Keys Protocol for Group Messaging. *XVII Reunión española sobre criptología y seguridad de la información. RECSI 2022*, 265:25, 2022.
- [BCG23a] David Balbás, Daniel Collins, and Phillip Gajland. Whatsupp with sender keys? analysis, improvements and security proofs. *Cryptology ePrint Archive*, Report 2023/1385, 2023. <https://eprint.iacr.org/2023/1385>.
- [BCG23b] David Balbás, Daniel Collins, and Phillip Gajland. WhatsUpp with Sender Keys? Analysis, improvements and security proofs. In *Advances in Cryptology –*

-
- ASIACRYPT 2023, Part I*, Lecture Notes in Computer Science. Springer-Verlag, December 7–11, 2023.
- [BCK22] Chris Brzuska, Eric Cornelissen, and Konrad Kohbrok. Security analysis of the MLS key derivation. In *2022 IEEE Symposium on Security and Privacy*, pages 2535–2553, San Francisco, CA, USA, May 22–26, 2022. IEEE Computer Society Press.
- [BCLZL22] Renas Bacho, Daniel Collins, Chen-Da Liu-Zhang, and Julian Loss. Network-agnostic security comes for (almost) free in DKG and MPC. Cryptology ePrint Archive, Report 2022/1369, 2022. <https://eprint.iacr.org/2022/1369>.
- [BCLZL23] Renas Bacho, Daniel Collins, Chen-Da Liu-Zhang, and Julian Loss. Network-agnostic security comes (almost) for free in DKG and MPC. In Helena Handschuh and Anna Lysyanskaya, editors, *Advances in Cryptology – CRYPTO 2023, Part I*, volume 14081 of *Lecture Notes in Computer Science*, pages 71–106, Santa Barbara, CA, USA, August 20–24, 2023. Springer, Heidelberg, Germany.
- [BCM23a] Subhadeep Banik, Daniel Collins, and Willi Meier. Near collision attack against grain V1. In Mehdi Tibouchi and Xiaofeng Wang, editors, *ACNS 23: 21st International Conference on Applied Cryptography and Network Security, Part I*, volume 13905 of *Lecture Notes in Computer Science*, pages 178–207, Kyoto, Japan, June 19–22, 2023. Springer, Heidelberg, Germany.
- [BCM23b] Subhadeep Banik, Daniel Collins, and Willi Meier. Near collision attack against grain v1. Cryptology ePrint Archive, Paper 2023/884, 2023. <https://eprint.iacr.org/2023/884>.
- [BCV22] David Balbás, Daniel Collins, and Serge Vaudenay. Cryptographic administration for secure group messaging. Cryptology ePrint Archive, Report 2022/1411, 2022. <https://eprint.iacr.org/2022/1411>.
- [BCV23] David Balbás, Daniel Collins, and Serge Vaudenay. Cryptographic administration for secure group messaging. In *USENIX Security 2023: 32nd USENIX Security Symposium*, pages 1253–1270, Anaheim, CA, USA, August 2023. USENIX Association.
- [BDF⁺11] Dan Boneh, Özgür Dagdelen, Marc Fischlin, Anja Lehmann, Christian Schaffner, and Mark Zhandry. Random oracles in a quantum world. In Dong Hoon Lee and Xiaoyun Wang, editors, *Advances in Cryptology – ASIACRYPT 2011*, volume 7073 of *Lecture Notes in Computer Science*, pages 41–69, Seoul, South Korea, December 4–8, 2011. Springer, Heidelberg, Germany.
- [BDG⁺22] Alexander Bienstock, Yevgeniy Dodis, Sanjam Garg, Garrison Grogan, Mohammad Hajiabadi, and Paul Rösler. On the worst-case inefficiency of CGKA. In

Bibliography

- Eike Kiltz and Vinod Vaikuntanathan, editors, *TCC 2022: 20th Theory of Cryptography Conference, Part II*, volume 13748 of *Lecture Notes in Computer Science*, pages 213–243, Chicago, IL, USA, November 7–10, 2022. Springer, Heidelberg, Germany.
- [BDK⁺18] Joppe Bos, Léo Ducas, Eike Kiltz, Tancrede Lepoint, Vadim Lyubashevsky, John M Schanck, Peter Schwabe, Gregor Seiler, and Damien Stehlé. Crystalskyber: a cca-secure module-lattice-based kem. In *2018 IEEE European Symposium on Security and Privacy (EuroS&P)*, pages 353–367. IEEE, 2018.
- [BDR20] Alexander Bienstock, Yevgeniy Dodis, and Paul Rösler. On the price of concurrency in group ratcheting protocols. In Rafael Pass and Krzysztof Pietrzak, editors, *TCC 2020: 18th Theory of Cryptography Conference, Part II*, volume 12551 of *Lecture Notes in Computer Science*, pages 198–228, Durham, NC, USA, November 16–19, 2020. Springer, Heidelberg, Germany.
- [Bel06] Mihir Bellare. New proofs for NMAC and HMAC: Security without collision-resistance. In Cynthia Dwork, editor, *Advances in Cryptology – CRYPTO 2006*, volume 4117 of *Lecture Notes in Computer Science*, pages 602–619, Santa Barbara, CA, USA, August 20–24, 2006. Springer, Heidelberg, Germany.
- [Ber06] Daniel J. Bernstein. Curve25519: New Diffie-Hellman speed records. In Moti Yung, Yevgeniy Dodis, Aggelos Kiayias, and Tal Malkin, editors, *PKC 2006: 9th International Conference on Theory and Practice of Public Key Cryptography*, volume 3958 of *Lecture Notes in Computer Science*, pages 207–228, New York, NY, USA, April 24–26, 2006. Springer, Heidelberg, Germany.
- [Beu20] Ward Beullens. Sigma protocols for MQ, PKP and SIS, and Fishy signature schemes. In Anne Canteaut and Yuval Ishai, editors, *Advances in Cryptology – EUROCRYPT 2020, Part III*, volume 12107 of *Lecture Notes in Computer Science*, pages 183–211, Zagreb, Croatia, May 10–14, 2020. Springer, Heidelberg, Germany.
- [BFG⁺20] Jacqueline Brendel, Marc Fischlin, Felix Günther, Christian Janson, and Douglas Stebila. Towards post-quantum security for Signal’s X3DH handshake. In Orr Dunkelman, Michael J. Jacobson Jr., and Colin O’Flynn, editors, *SAC 2020: 27th Annual International Workshop on Selected Areas in Cryptography*, volume 12804 of *Lecture Notes in Computer Science*, pages 404–430, Halifax, NS, Canada (Virtual Event), October 21–23, 2020. Springer, Heidelberg, Germany.
- [BFG⁺22a] Alexander Bienstock, Jaiden Fairuze, Sanjam Garg, Pratyay Mukherjee, and Srinivasan Raghuraman. A more complete analysis of the Signal double ratchet algorithm. In Yevgeniy Dodis and Thomas Shrimpton, editors, *Advances in Cryptology – CRYPTO 2022, Part I*, volume 13507 of *Lecture Notes in Computer Science*, pages 784–813, Santa Barbara, CA, USA, August 15–18, 2022. Springer, Heidelberg, Germany.

- [BFG⁺22b] Jacqueline Brendel, Rune Fiedler, Felix Günther, Christian Janson, and Douglas Stebila. Post-quantum asynchronous deniable key exchange and the Signal handshake. In Goichiro Hanaoka, Junji Shikata, and Yohei Watanabe, editors, *PKC 2022: 25th International Conference on Theory and Practice of Public Key Cryptography, Part II*, volume 13178 of *Lecture Notes in Computer Science*, pages 3–34, Virtual Event, March 8–11, 2022. Springer, Heidelberg, Germany.
- [BFWW11] Christina Brzuska, Marc Fischlin, Bogdan Warinschi, and Stephen C. Williams. Composability of Bellare-Rogaway key exchange protocols. In Yan Chen, George Danezis, and Vitaly Shmatikov, editors, *ACM CCS 2011: 18th Conference on Computer and Communications Security*, pages 51–62, Chicago, Illinois, USA, October 17–21, 2011. ACM Press.
- [BG21] Colin Boyd and Kai Gellert. A modern view on forward security. *The Computer Journal*, 64(4):639–652, 2021.
- [BGB04] Nikita Borisov, Ian Goldberg, and Eric Brewer. Off-the-record communication, or, why not to use pgp. In *Proceedings of the 2004 ACM workshop on Privacy in the electronic society*, pages 77–84, 2004.
- [BHJ⁺15] Christoph Bader, Dennis Hofheinz, Tibor Jager, Eike Kiltz, and Yong Li. Tightly-secure authenticated key exchange. In Yevgeniy Dodis and Jesper Buus Nielsen, editors, *TCC 2015: 12th Theory of Cryptography Conference, Part I*, volume 9014 of *Lecture Notes in Computer Science*, pages 629–658, Warsaw, Poland, March 23–25, 2015. Springer, Heidelberg, Germany.
- [BJKS23a] Karthikeyan Bhargavan, Charlie Jacomme, Franziskus Kiefer, and Rolfe Schmidt. An analysis of signal’s pqxdh, 2023. <https://cryspen.com/post/pqxdh/> Accessed: 23.10.23.
- [BJKS23b] Karthikeyan Bhargavan, Charlie Jacomme, Franziskus Kiefer, and Rolfe Schmidt. An analysis of signal’s pqxdh, 2023.
- [BJRW21] Katharina Boudgoust, Corentin Jeudy, Adeline Roux-Langlois, and Weiqiang Wen. On the hardness of module-LWE with binary secret. In Kenneth G. Paterson, editor, *Topics in Cryptology – CT-RSA 2021*, volume 12704 of *Lecture Notes in Computer Science*, pages 503–526, Virtual Event, May 17–20, 2021. Springer, Heidelberg, Germany.
- [Bla07] Bruno Blanchet. Cryptoverif: Computationally sound mechanized prover for cryptographic protocols. In *Dagstuhl seminar “Formal Protocol Verification Applied*, volume 117, page 156, 2007.
- [BLP⁺13] Zvika Brakerski, Adeline Langlois, Chris Peikert, Oded Regev, and Damien Stehlé. Classical hardness of learning with errors. In Dan Boneh, Tim Roughgarden, and Joan Feigenbaum, editors, *45th Annual ACM Symposium on Theory of Computing*, pages 575–584, Palo Alto, CA, USA, June 1–4, 2013. ACM Press.

Bibliography

- [BLR04] Boaz Barak, Yehuda Lindell, and Tal Rabin. Protocol initialization for the framework of universal composability. *Cryptology ePrint Archive*, Report 2004/006, 2004. <https://eprint.iacr.org/2004/006>.
- [BM99] Mihir Bellare and Sara K. Miner. A forward-secure digital signature scheme. In Michael J. Wiener, editor, *Advances in Cryptology – CRYPTO’99*, volume 1666 of *Lecture Notes in Computer Science*, pages 431–448, Santa Barbara, CA, USA, August 15–19, 1999. Springer, Heidelberg, Germany.
- [BR93] Mihir Bellare and Phillip Rogaway. Random oracles are practical: A paradigm for designing efficient protocols. In Dorothy E. Denning, Raymond Pyle, Ravi Ganesan, Ravi S. Sandhu, and Victoria Ashby, editors, *ACM CCS 93: 1st Conference on Computer and Communications Security*, pages 62–73, Fairfax, Virginia, USA, November 3–5, 1993. ACM Press.
- [BR94] Mihir Bellare and Phillip Rogaway. Entity authentication and key distribution. In Douglas R. Stinson, editor, *Advances in Cryptology – CRYPTO’93*, volume 773 of *Lecture Notes in Computer Science*, pages 232–249, Santa Barbara, CA, USA, August 22–26, 1994. Springer, Heidelberg, Germany.
- [BRT23] Alexander Bienstock, Paul Rösler, and Yi Tang. Asmesh: Anonymous and secure messaging in mesh networks using stronger, anonymous double ratchet. *Cryptology ePrint Archive*, 2023.
- [BRV20] Fatih Balli, Paul Rösler, and Serge Vaudenay. Determining the core primitive for optimally secure ratcheting. In Shiho Moriai and Huaxiong Wang, editors, *Advances in Cryptology – ASIACRYPT 2020, Part III*, volume 12493 of *Lecture Notes in Computer Science*, pages 621–650, Daejeon, South Korea, December 7–11, 2020. Springer, Heidelberg, Germany.
- [BSI23] BSI. Cryptographic Mechanisms: Recommendations and Key Lengths. <https://www.bsi.bund.de/SharedDocs/Downloads/EN/BSI/Publications/TechGuidelines/TG02102/BSI-TR-02102-1.html>, 2023. Accessed: 28-09-2023.
- [BSJ⁺15] Richard Barnes, Bruce Schneier, Cullen Fluffy Jennings, Ted Hardie, Brian Trammell, Christian Huitema, and Daniel Borkmann. Confidentiality in the Face of Pervasive Surveillance: A Threat Model and Problem Statement. RFC 7624, August 2015.
- [BSJ⁺17] Mihir Bellare, Asha Camper Singh, Joseph Jaeger, Maya Nyayapati, and Igors Stepanovs. Ratcheted encryption and key exchange: The security of messaging. In Jonathan Katz and Hovav Shacham, editors, *Advances in Cryptology – CRYPTO 2017, Part III*, volume 10403 of *Lecture Notes in Computer Science*, pages 619–650, Santa Barbara, CA, USA, August 20–24, 2017. Springer, Heidelberg, Germany.

- [BSSW02] Dirk Balfanz, Diana K. Smetters, Paul Stewart, and H. Chi Wong. Talking to strangers: Authentication in ad-hoc wireless networks. In *ISOC Network and Distributed System Security Symposium – NDSS 2002*, San Diego, CA, USA, February 6–8, 2002. The Internet Society.
- [BT85] Gabriel Bracha and Sam Toueg. Asynchronous consensus and broadcast protocols. *Journal of the ACM (JACM)*, 32(4):824–840, 1985.
- [Can01] Ran Canetti. Universally composable security: A new paradigm for cryptographic protocols. In *42nd Annual Symposium on Foundations of Computer Science*, pages 136–145, Las Vegas, NV, USA, October 14–17, 2001. IEEE Computer Society Press.
- [CCBR22a] Andrea Caforio, Daniel Collins, Subhadeep Banik, and Francesco Regazzoni. A small GIFT-COFB: Lightweight bit-serial architectures. In Lejla Batina and Joan Daemen, editors, *AFRICACRYPT 22: 13th International Conference on Cryptology in Africa*, volume 2022 of *Lecture Notes in Computer Science*, pages 53–77, Fes, Morocco, July 18–20, 2022. Springer Nature Switzerland.
- [CCBR22b] Andrea Caforio, Daniel Collins, Subhadeep Banik, and Francesco Regazzoni. A small GIFT-COFB: Lightweight bit-serial architectures. *Cryptology ePrint Archive*, Report 2022/955, 2022. <https://eprint.iacr.org/2022/955>.
- [CCD⁺20] Katriel Cohn-Gordon, Cas Cremers, Benjamin Dowling, Luke Garratt, and Douglas Stebila. A formal security analysis of the Signal messaging protocol. *Journal of Cryptology*, 33(4):1914–1983, October 2020.
- [CCG16] Katriel Cohn-Gordon, Cas J. F. Cremers, and Luke Garratt. On post-compromise security. In Michael Hicks and Boris Köpf, editors, *CSF 2016: IEEE 29th Computer Security Foundations Symposium*, pages 164–178, Lisbon, Portugal, June 27–1, 2016. IEEE Computer Society Press.
- [CCG⁺18] Katriel Cohn-Gordon, Cas Cremers, Luke Garratt, Jon Millican, and Kevin Milner. On ends-to-ends encryption: Asynchronous group messaging with strong security guarantees. In David Lie, Mohammad Mannan, Michael Backes, and XiaoFeng Wang, editors, *ACM CCS 2018: 25th Conference on Computer and Communications Security*, pages 1802–1819, Toronto, ON, Canada, October 15–19, 2018. ACM Press.
- [CCG⁺19] Katriel Cohn-Gordon, Cas Cremers, Kristian Gjøsteen, Håkon Jacobsen, and Tibor Jager. Highly efficient key exchange protocols with optimal tightness. In Alexandra Boldyreva and Daniele Micciancio, editors, *Advances in Cryptology – CRYPTO 2019, Part III*, volume 11694 of *Lecture Notes in Computer Science*, pages 767–797, Santa Barbara, CA, USA, August 18–22, 2019. Springer, Heidelberg, Germany.

Bibliography

- [CCGB21a] Andrea Caforio, Daniel Collins, Ognjen Glamočanin, and Subhadeep Banik. Improving first-order threshold implementations of SKINNY. In *Progress in Cryptology–INDOCRYPT 2021: 22nd International Conference on Cryptology in India, Jaipur, India, December 12–15, 2021, Proceedings 22*, pages 246–267. Springer, 2021.
- [CCGB21b] Andrea Caforio, Daniel Collins, Ognjen Glamocanin, and Subhadeep Banik. Improving first-order threshold implementations of SKINNY. Cryptology ePrint Archive, Report 2021/1425, 2021. <https://eprint.iacr.org/2021/1425>.
- [CCHD23] Daniel Collins, Simone Colombo, and Loïs Huguenin-Dumittan. Real world deniability in messaging. Cryptology ePrint Archive, Paper 2023/403, 2023. <https://eprint.iacr.org/2023/403>.
- [CD23] Wouter Castryck and Thomas Decru. An efficient key recovery attack on sidh. In *Annual International Conference on the Theory and Applications of Cryptographic Techniques*, pages 423–447. Springer, 2023.
- [CDv⁺03] Dwaine E. Clarke, Srinivas Devadas, Marten van Dijk, Blaise Gassend, and G. Edward Suh. Incremental multiset hash functions and their application to memory integrity checking. In Chi-Sung Laih, editor, *Advances in Cryptology – ASIACRYPT 2003*, volume 2894 of *Lecture Notes in Computer Science*, pages 188–207, Taipei, Taiwan, November 30 – December 4, 2003. Springer, Heidelberg, Germany.
- [CDV21] Andrea Caforio, F. Betül Durak, and Serge Vaudenay. Beyond security and efficiency: On-demand ratcheting with security awareness. In Juan Garay, editor, *PKC 2021: 24th International Conference on Theory and Practice of Public Key Cryptography, Part II*, volume 12711 of *Lecture Notes in Computer Science*, pages 649–677, Virtual Event, May 10–13, 2021. Springer, Heidelberg, Germany.
- [CEST22] Kelong Cong, Karim Eldefrawy, Nigel P. Smart, and Ben Terner. The key lattice framework for concurrent group messaging. Cryptology ePrint Archive, Report 2022/1531, 2022. <https://eprint.iacr.org/2022/1531>.
- [CF11] Cas Cremers and Michele Feltz. One-round strongly secure key exchange with perfect forward secrecy and deniability. Cryptology ePrint Archive, Report 2011/300, 2011. <https://eprint.iacr.org/2011/300>.
- [CGG⁺22] Pierre Civit, Seth Gilbert, Vincent Gramoli, Rachid Guerraoui, Jovan Komatovic, Zarko Milosevic, and Adi Serendinschi. Crime and punishment in distributed byzantine decision tasks. *Cryptology ePrint Archive*, 2022.
- [CGK⁺20a] Daniel Collins, Rachid Guerraoui, Jovan Komatovic, Petr Kuznetsov, Matteo Monti, Matej Pavlovic, Yvonne-Anne Pignolet, Dragos-Adrian Seredinschi, Andrei Tonkikh, and Athanasios Xytkis. Online payments by merely broadcasting

- messages. In *2020 50th Annual IEEE/IFIP International Conference on Dependable Systems and Networks (DSN)*, pages 26–38. IEEE, 2020.
- [CGK⁺20b] Daniel Collins, Rachid Guerraoui, Jovan Komatovic, Matteo Monti, Athanasios Xygkis, Matej Pavlovic, Petr Kuznetsov, Yvonne-Anne Pignolet, Dragos-Adrian Seredinschi, and Andrei Tonkikh. Online payments by merely broadcasting messages (extended version). *arXiv preprint arXiv:2004.13184*, 2020.
- [CHDN⁺24a] Daniel Collins, Loïs Huguenin-Dumittan, Ngoc Khanh Nguyen, Nicolas Rolin, and Serge Vaudenay. K-Waay: Fast and deniable post-quantum X3DH without ring signatures. In *USENIX Security 2024: 33rd USENIX Security Symposium*. USENIX Association, August 2024.
- [CHDN⁺24b] Daniel Collins, Loïs Huguenin-Dumittan, Ngoc Khanh Nguyen, Nicolas Rolin, and Serge Vaudenay. K-waay: Fast and deniable post-quantum x3dh without ring signatures. Cryptology ePrint Archive, Paper 2024/120, 2024. <https://eprint.iacr.org/2024/120>.
- [CHK21] Cas Cremers, Britta Hale, and Konrad Kohbrok. The complexities of healing in secure group messaging: Why cross-group effects matter. In Michael Bailey and Rachel Greenstadt, editors, *USENIX Security 2021: 30th USENIX Security Symposium*, pages 1847–1864. USENIX Association, August 11–13, 2021.
- [CJSV22] Ran Canetti, Palak Jain, Marika Swanberg, and Mayank Varia. Universally composable end-to-end secure messaging. In Yevgeniy Dodis and Thomas Shrimpton, editors, *Advances in Cryptology – CRYPTO 2022, Part II*, volume 13508 of *Lecture Notes in Computer Science*, pages 3–33, Santa Barbara, CA, USA, August 15–18, 2022. Springer, Heidelberg, Germany.
- [CK02] Ran Canetti and Hugo Krawczyk. Security analysis of IKE’s signature-based key-exchange protocol. In Moti Yung, editor, *Advances in Cryptology – CRYPTO 2002*, volume 2442 of *Lecture Notes in Computer Science*, pages 143–161, Santa Barbara, CA, USA, August 18–22, 2002. Springer, Heidelberg, Germany. <https://eprint.iacr.org/2002/120/>.
- [CL⁺99] Miguel Castro, Barbara Liskov, et al. Practical byzantine fault tolerance. In *OSDI*, volume 99, pages 173–186, 1999.
- [CLM⁺18] Wouter Castryck, Tanja Lange, Chloe Martindale, Lorenz Panny, and Joost Renes. CSIDH: An efficient post-quantum commutative group action. In Thomas Peyrin and Steven Galbraith, editors, *Advances in Cryptology – ASIACRYPT 2018, Part III*, volume 11274 of *Lecture Notes in Computer Science*, pages 395–427, Brisbane, Queensland, Australia, December 2–6, 2018. Springer, Heidelberg, Germany.

Bibliography

- [CN11] Yuanmi Chen and Phong Q. Nguyen. BKZ 2.0: Better lattice security estimates. In Dong Hoon Lee and Xiaoyun Wang, editors, *Advances in Cryptology – ASIACRYPT 2011*, volume 7073 of *Lecture Notes in Computer Science*, pages 1–20, Seoul, South Korea, December 4–8, 2011. Springer, Heidelberg, Germany.
- [CPZ20] Melissa Chase, Trevor Perrin, and Greg Zaverucha. The Signal private group system and anonymous credentials supporting efficient verifiable encryption. In Jay Ligatti, Xinming Ou, Jonathan Katz, and Giovanni Vigna, editors, *ACM CCS 2020: 27th Conference on Computer and Communications Security*, pages 1445–1459, Virtual Event, USA, November 9–13, 2020. ACM Press.
- [CSM⁺20] Raymond Cheng, William Scott, Elisaweta Masserova, Irene Zhang, Vipul Goyal, Thomas Anderson, Arvind Krishnamurthy, and Bryan Parno. Talek: Private group messaging with hidden access patterns. In *Annual Computer Security Applications Conference*, pages 84–99, 2020.
- [CZ22] Cas Cremers and Mang Zhao. Provably post-quantum secure messaging with strong compromise resilience and immediate decryption. *Cryptology ePrint Archive*, Report 2022/1481, 2022. <https://eprint.iacr.org/2022/1481>.
- [CZ24] Cas Cremers and Mang Zhao. Secure messaging with strong compromise resilience, temporal privacy, and immediate decryption. In *IEEE S&P*, 2024.
- [DDF21] Julien Devigne, Céline Duguey, and Pierre-Alain Fouque. MLS group messaging: How zero-knowledge can secure updates. In Elisa Bertino, Haya Shulman, and Michael Waidner, editors, *ESORICS 2021: 26th European Symposium on Research in Computer Security, Part II*, volume 12973 of *Lecture Notes in Computer Science*, pages 587–607, Darmstadt, Germany, October 4–8, 2021. Springer, Heidelberg, Germany.
- [DDGR20] Dana Dachman-Soled, Léo Ducas, Huijing Gong, and Mélissa Rossi. LWE with side information: Attacks and concrete security estimation. In Daniele Micciancio and Thomas Ristenpart, editors, *Advances in Cryptology – CRYPTO 2020, Part II*, volume 12171 of *Lecture Notes in Computer Science*, pages 329–358, Santa Barbara, CA, USA, August 17–21, 2020. Springer, Heidelberg, Germany.
- [DFG⁺23] Gareth T. Davies, Sebastian H. Faller, Kai Gellert, Tobias Handirk, Julia Hesse, Máté Horváth, and Tibor Jager. Security analysis of the WhatsApp end-to-end encrypted backup protocol. In Helena Handschuh and Anna Lysyanskaya, editors, *Advances in Cryptology – CRYPTO 2023, Part IV*, volume 14084 of *Lecture Notes in Computer Science*, pages 330–361, Santa Barbara, CA, USA, August 20–24, 2023. Springer, Heidelberg, Germany.
- [DFMS22] Jelle Don, Serge Fehr, Christian Majenz, and Christian Schaffner. Online-extractability in the quantum random-oracle model. In Orr Dunkelman and Stefan Dziembowski, editors, *Advances in Cryptology – EUROCRYPT 2022, Part III*,

- volume 13277 of *Lecture Notes in Computer Science*, pages 677–706, Trondheim, Norway, May 30 – June 3, 2022. Springer, Heidelberg, Germany.
- [DG22] Samuel Dobson and Steven D. Galbraith. Post-quantum signal key agreement from SIDH. In Jung Hee Cheon and Thomas Johansson, editors, *Post-Quantum Cryptography - 13th International Workshop, PQCrypto 2022, Virtual Event, September 28-30, 2022, Proceedings*, volume 13512 of *Lecture Notes in Computer Science*, pages 422–450. Springer, 2022.
- [DGK06] Mario Di Raimondo, Rosario Gennaro, and Hugo Krawczyk. Deniable authentication and key exchange. In Ari Juels, Rebecca N. Wright, and Sabrina De Capitani di Vimercati, editors, *ACM CCS 2006: 13th Conference on Computer and Communications Security*, pages 400–409, Alexandria, Virginia, USA, October 30 – November 3, 2006. ACM Press.
- [DGP22] Benjamin Dowling, Felix Günther, and Alexandre Poirrier. Continuous authentication in secure messaging. In *ESORICS*, 2022.
- [DH76] Whitfield Diffie and Martin E. Hellman. New directions in cryptography. *IEEE Transactions on Information Theory*, 22(6):644–654, 1976.
- [DH20] Benjamin Dowling and Britta Hale. There can be no compromise: The necessity of ratcheted authentication in secure messaging. Cryptology ePrint Archive, Report 2020/541, 2020. <https://eprint.iacr.org/2020/541>.
- [DH21] Benjamin Dowling and Britta Hale. Secure Messaging Authentication against Active Man-in-the-Middle Attacks. In *EuroS&P*, 2021.
- [DHRR22] Benjamin Dowling, Eduard Hauck, Doreen Riepel, and Paul Rösler. Strongly anonymous ratcheted key exchange. In Shweta Agrawal and Dongdai Lin, editors, *Advances in Cryptology – ASIACRYPT 2022, Part III*, volume 13793 of *Lecture Notes in Computer Science*, pages 119–150, Taipei, Taiwan, December 5–9, 2022. Springer, Heidelberg, Germany.
- [DKL⁺18] Léo Ducas, Eike Kiltz, Tancrede Lepoint, Vadim Lyubashevsky, Peter Schwabe, Gregor Seiler, and Damien Stehlé. CRYSTALS-Dilithium: A lattice-based digital signature scheme. *IACR Transactions on Cryptographic Hardware and Embedded Systems*, 2018(1):238–268, 2018. <https://tches.iacr.org/index.php/TCHES/article/view/839>.
- [DNS04] Cynthia Dwork, Moni Naor, and Amit Sahai. Concurrent zero-knowledge. *Journal of the ACM (JACM)*, 51(6):851–898, 2004.
- [DSGHK23] Dana Dachman-Soled, Huijing Gong, Tom Hanson, and Hunter Kippen. Revisiting security estimation for lwe with hints from a geometric perspective. In Helena Handschuh and Anna Lysyanskaya, editors, *Advances in Cryptology – CRYPTO 2023*, pages 748–781, Cham, 2023. Springer Nature Switzerland.

Bibliography

- [DV18] F. Betül Durak and Serge Vaudenay. Bidirectional asynchronous ratcheted key agreement with linear complexity. Cryptology ePrint Archive, Report 2018/889, 2018. <https://eprint.iacr.org/2018/889>.
- [DV19] F. Betül Durak and Serge Vaudenay. Bidirectional asynchronous ratcheted key agreement with linear complexity. In Nuttapong Attrapadung and Takeshi Yagi, editors, *IWSEC 19: 14th International Workshop on Security, Advances in Information and Computer Security*, volume 11689 of *Lecture Notes in Computer Science*, pages 343–362, Tokyo, Japan, August 28–30, 2019. Springer, Heidelberg, Germany.
- [EM19] Ksenia Ermoshina and Francesca Musiani. “standardising by running code”: the signal protocol and de facto standardisation in end-to-end encrypted messaging. *Internet Histories*, 3(3-4):343–363, 2019.
- [ESZ22] Muhammed F. Esgin, Ron Steinfeld, and Raymond K. Zhao. MatRiCT⁺: More efficient post-quantum private blockchain payments. In *2022 IEEE Symposium on Security and Privacy*, pages 1281–1298, San Francisco, CA, USA, May 22–26, 2022. IEEE Computer Society Press.
- [fIS23] BSI German Federal Office for Information Security. Bsi tr-01102-1, 2023. <https://www.bsi.bund.de/SharedDocs/Downloads/EN/BSI/Publications/TechGuidelines/TG02102/BSI-TR-02102-1.html>.
- [FMB⁺16] Tilman Frosch, Christian Mainka, Christoph Bader, Florian Bergsma, Jörg Schwenk, and Thorsten Holz. How secure is textsecure? In *2016 IEEE European Symposium on Security and Privacy (EuroS&P)*, pages 457–472. IEEE, 2016.
- [Gal21] Tarek Galal. yowsup, Code Repository, 2021. <https://github.com/tgalal/yowsup>.
- [GHP18] Federico Giacon, Felix Heuer, and Bertram Poettering. KEM combiners. In Michel Abdalla and Ricardo Dahab, editors, *PKC 2018: 21st International Conference on Theory and Practice of Public Key Cryptography, Part I*, volume 10769 of *Lecture Notes in Computer Science*, pages 190–218, Rio de Janeiro, Brazil, March 25–29, 2018. Springer, Heidelberg, Germany.
- [Gie23] Hauke Gierow. Wire welcomes the publication of Messaging Layer Security as RFC 9420. <https://wire.com/en/blog/wire-welcomes-the-publication-of-messaging-layer-security-as-rfc-9420/>, 2023. Accessed: 28-09-2023.
- [GM82] Shafi Goldwasser and Silvio Micali. Probabilistic encryption and how to play mental poker keeping secret all partial information. In *14th Annual ACM Symposium on Theory of Computing*, pages 365–377, San Francisco, CA, USA, May 5–7, 1982. ACM Press.

- [GPA19] Lachlan J. Gunn, Ricardo Vieitez Parra, and N. Asokan. Circumventing cryptographic deniability with remote attestation. *Proceedings on Privacy Enhancing Technologies*, 2019(3):350–369, July 2019.
- [Gua13] The Guardian. The NSA files. <https://www.theguardian.com/us-news/the-nsa-files>, 2013. Accessed: 28-09-2023.
- [Gün90] Christoph G. Günther. An identity-based key-exchange protocol. In Jean-Jacques Quisquater and Joos Vandewalle, editors, *Advances in Cryptology – EUROCRYPT’89*, volume 434 of *Lecture Notes in Computer Science*, pages 29–37, Houthalen, Belgium, April 10–13, 1990. Springer, Heidelberg, Germany.
- [HHK17] Dennis Hofheinz, Kathrin Hövelmanns, and Eike Kiltz. A modular analysis of the Fujisaki-Okamoto transformation. In Yael Kalai and Leonid Reyzin, editors, *TCC 2017: 15th Theory of Cryptography Conference, Part I*, volume 10677 of *Lecture Notes in Computer Science*, pages 341–371, Baltimore, MD, USA, November 12–15, 2017. Springer, Heidelberg, Germany.
- [HKD07] Andreas Haeberlen, Petr Kouznetsov, and Peter Druschel. Peerreview: Practical accountability for distributed systems. *SIGOPS*, 41(6):175–188, 2007.
- [HKKP21] Keitaro Hashimoto, Shuichi Katsumata, Kris Kwiatkowski, and Thomas Prest. An efficient and generic construction for Signal’s handshake (X3DH): Post-quantum, state leakage secure, and deniable. In Juan Garay, editor, *PKC 2021: 24th International Conference on Theory and Practice of Public Key Cryptography, Part II*, volume 12711 of *Lecture Notes in Computer Science*, pages 410–440, Virtual Event, May 10–13, 2021. Springer, Heidelberg, Germany.
- [HKKP22] Keitaro Hashimoto, Shuichi Katsumata, Kris Kwiatkowski, and Thomas Prest. An efficient and generic construction for Signal’s handshake (X3DH): Post-quantum, state leakage secure, and deniable. *Journal of Cryptology*, 35(3):17, July 2022.
- [HKP⁺21] Keitaro Hashimoto, Shuichi Katsumata, Eamonn Postlethwaite, Thomas Prest, and Bas Westerbaan. A concrete treatment of efficient continuous group key agreement via multi-recipient PKEs. In Giovanni Vigna and Elaine Shi, editors, *ACM CCS 2021: 28th Conference on Computer and Communications Security*, pages 1441–1462, Virtual Event, Republic of Korea, November 15–19, 2021. ACM Press.
- [HKP22] Keitaro Hashimoto, Shuichi Katsumata, and Thomas Prest. How to hide Meta-Data in MLS-like secure group messaging: Simple, modular, and post-quantum. In Heng Yin, Angelos Stavrou, Cas Cremers, and Elaine Shi, editors, *ACM CCS 2022: 29th Conference on Computer and Communications Security*, pages 1399–1412, Los Angeles, CA, USA, November 7–11, 2022. ACM Press.

Bibliography

- [Hog23] Giles Hogben. An important step towards secure and interoperable messaging. <https://security.googleblog.com/2023/07/an-important-step-towards-secure-and.html>, 2023. Accessed: 28-09-2023.
- [HS20] Martha Norberg Hovd and Martijn Stam. Vetted encryption. In Karthikeyan Bhargavan, Elisabeth Oswald, and Manoj Prabhakaran, editors, *Progress in Cryptology - INDOCRYPT 2020: 21st International Conference in Cryptology in India*, volume 12578 of *Lecture Notes in Computer Science*, pages 488–507, Bangalore, India, December 13–16, 2020. Springer, Heidelberg, Germany.
- [HV22] Loïs Huguenin-Dumittan and Serge Vaudenay. On IND-qCCA security in the ROM and its applications - CPA security is sufficient for TLS 1.3. In Orr Dunkelman and Stefan Dziembowski, editors, *Advances in Cryptology – EUROCRYPT 2022, Part III*, volume 13277 of *Lecture Notes in Computer Science*, pages 613–642, Trondheim, Norway, May 30 – June 3, 2022. Springer, Heidelberg, Germany.
- [IETF23] (IETF) Internet Engineering Task Force. Messaging layer security, mailing list, 2023. <https://mailarchive.ietf.org/arch/browse/mls/>.
- [IOZ14] Yuval Ishai, Rafail Ostrovsky, and Vassilis Zikas. Secure multi-party computation with identifiable abort. In Juan A. Garay and Rosario Gennaro, editors, *Advances in Cryptology – CRYPTO 2014, Part II*, volume 8617 of *Lecture Notes in Computer Science*, pages 369–386, Santa Barbara, CA, USA, August 17–21, 2014. Springer, Heidelberg, Germany.
- [Jef20] Kee Jefferys. Session Protocol: Technical implementation details. <https://getsession.org/blog/session-protocol-technical-information> (accessed July 4th 2023), 2020.
- [JLN19] Riko Jacob, Kasper Green Larsen, and Jesper Buus Nielsen. Lower bounds for oblivious data structures. In Timothy M. Chan, editor, *30th Annual ACM-SIAM Symposium on Discrete Algorithms*, pages 2439–2447, San Diego, CA, USA, January 6–9, 2019. ACM-SIAM.
- [JMM19a] Daniel Jost, Ueli Maurer, and Marta Mularczyk. Efficient ratcheting: Almost-optimal guarantees for secure messaging. In Yuval Ishai and Vincent Rijmen, editors, *Advances in Cryptology – EUROCRYPT 2019, Part I*, volume 11476 of *Lecture Notes in Computer Science*, pages 159–188, Darmstadt, Germany, May 19–23, 2019. Springer, Heidelberg, Germany.
- [JMM19b] Daniel Jost, Ueli Maurer, and Marta Mularczyk. A unified and composable take on ratcheting. In Dennis Hofheinz and Alon Rosen, editors, *TCC 2019: 17th Theory of Cryptography Conference, Part II*, volume 11892 of *Lecture Notes in Computer Science*, pages 180–210, Nuremberg, Germany, December 1–5, 2019. Springer, Heidelberg, Germany.

- [JMZ23] Haodong Jiang, Zhi Ma, and Zhenfeng Zhang. Post-quantum security of key encapsulation mechanism against cca attacks with a single decapsulation query. In *ASIACRYPT 2023*. Springer-Verlag, 2023.
- [JS18] Joseph Jaeger and Igors Stepanovs. Optimal channel security against fine-grained state compromise: The safety of messaging. In Hovav Shacham and Alexandra Boldyreva, editors, *Advances in Cryptology – CRYPTO 2018, Part I*, volume 10991 of *Lecture Notes in Computer Science*, pages 33–62, Santa Barbara, CA, USA, August 19–23, 2018. Springer, Heidelberg, Germany.
- [KBB17] Nadim Kobeissi, Karthikeyan Bhargavan, and Bruno Blanchet. Automated verification for secure messaging protocols and their implementations: A symbolic and computational approach. In *2017 IEEE European symposium on security and privacy (EuroS&P)*, pages 435–450. IEEE, 2017.
- [KGP23] Kien Tuong Truong Kenneth G. Paterson, Matteo Scarlata. Three Lessons From Threema: Analysis of a Secure Messenger. In *2023 USENIX Security Symposium*, 2023.
- [KP05] Caroline Kudla and Kenneth G. Paterson. Modular security proofs for key agreement protocols. In Bimal K. Roy, editor, *Advances in Cryptology – ASIACRYPT 2005*, volume 3788 of *Lecture Notes in Computer Science*, pages 549–565, Chennai, India, December 4–8, 2005. Springer, Heidelberg, Germany.
- [KPPW⁺21] Karen Klein, Guillermo Pascual-Perez, Michael Walter, Chethan Kamath, Margarita Capretto, Miguel Cueto, Iliia Markov, Michelle Yeo, Joël Alwen, and Krzysztof Pietrzak. Keep the dirt: Tainted TreeKEM, adaptively and actively secure continuous group key agreement. In *2021 IEEE Symposium on Security and Privacy*, pages 268–284, San Francisco, CA, USA, May 24–27, 2021. IEEE Computer Society Press.
- [KPRR23] Eike Kiltz, Jiaxin Pan, Doreen Riepel, and Magnus Ringerud. Multi-user cdh problems and the concrete security of naxos and hmqv. In *Cryptographers’ Track at the RSA Conference*, pages 645–671. Springer, 2023.
- [Kra05] Hugo Krawczyk. HMQV: A high-performance secure Diffie-Hellman protocol. In Victor Shoup, editor, *Advances in Cryptology – CRYPTO 2005*, volume 3621 of *Lecture Notes in Computer Science*, pages 546–566, Santa Barbara, CA, USA, August 14–18, 2005. Springer, Heidelberg, Germany.
- [KS05] Jonathan Katz and Ji Sun Shin. Modeling insider attacks on group key-exchange protocols. In Vijayalakshmi Atluri, Catherine Meadows, and Ari Juels, editors, *ACM CCS 2005: 12th Conference on Computer and Communications Security*, pages 180–189, Alexandria, Virginia, USA, November 7–11, 2005. ACM Press.

Bibliography

- [KS23] Ehren Kret and Rolfe Schmidt. The pqxdh key agreement protocol, 2023. <https://signal.org/docs/specifications/pqxdh/pqxdh.pdf>.
- [KT11] Ralf Küsters and Max Tuengerthal. Composition theorems without pre-established session identifiers. In Yan Chen, George Danezis, and Vitaly Shmatikov, editors, *ACM CCS 2011: 18th Conference on Computer and Communications Security*, pages 41–50, Chicago, Illinois, USA, October 17–21, 2011. ACM Press.
- [Laa16] Thijs Laarhoven. Search problems in cryptography: from fingerprinting to lattice sieving, February 2016. Proefschrift.
- [LAZ19] Xingye Lu, Man Ho Au, and Zhenfei Zhang. Raptor: A practical lattice-based (linkable) ring signature. In Robert H. Deng, Valérie Gauthier-Umaña, Martín Ochoa, and Moti Yung, editors, *ACNS 19: 17th International Conference on Applied Cryptography and Network Security*, volume 11464 of *Lecture Notes in Computer Science*, pages 110–130, Bogota, Colombia, June 5–7, 2019. Springer, Heidelberg, Germany.
- [LLM07] Brian A. LaMacchia, Kristin Lauter, and Anton Mityagin. Stronger security of authenticated key exchange. In Willy Susilo, Joseph K. Liu, and Yi Mu, editors, *ProvSec 2007: 1st International Conference on Provable Security*, volume 4784 of *Lecture Notes in Computer Science*, pages 1–16, Wollongong, Australia, November 1–2, 2007. Springer, Heidelberg, Germany.
- [LN18] Kasper Green Larsen and Jesper Buus Nielsen. Yes, there is an oblivious RAM lower bound! In Hovav Shacham and Alexandra Boldyreva, editors, *Advances in Cryptology – CRYPTO 2018, Part II*, volume 10992 of *Lecture Notes in Computer Science*, pages 523–542, Santa Barbara, CA, USA, August 19–23, 2018. Springer, Heidelberg, Germany.
- [LN22] Vadim Lyubashevsky and Ngoc Khanh Nguyen. BLOOM: Bimodal lattice one-out-of-many proofs and applications. In Shweta Agrawal and Dongdai Lin, editors, *Advances in Cryptology – ASIACRYPT 2022, Part IV*, volume 13794 of *Lecture Notes in Computer Science*, pages 95–125, Taipei, Taiwan, December 5–9, 2022. Springer, Heidelberg, Germany.
- [LPR10] Vadim Lyubashevsky, Chris Peikert, and Oded Regev. On ideal lattices and learning with errors over rings. In Henri Gilbert, editor, *Advances in Cryptology – EUROCRYPT 2010*, volume 6110 of *Lecture Notes in Computer Science*, pages 1–23, French Riviera, May 30 – June 3, 2010. Springer, Heidelberg, Germany.
- [LS15] Adeline Langlois and Damien Stehlé. Worst-case to average-case reductions for module lattices. *Des. Codes Cryptogr.*, 75(3):565–599, 2015.

- [Lun18] Joshua Lund. Technology preview: Sealed sender for signal. <https://signal.org/blog/sealed-sender/>, 2018. Last visited on 13-09-2023.
- [M⁺16] Moxie Marlinspike et al. Signal protocol repository, 2016.
- [Mar14] Moxie Marlinspike. Private Group Messaging. <https://signal.org/blog/private-groups/> (accessed Sep 5th 2023), 2014.
- [Mar17] Moxie Marlinspike. Safety number updates. <https://signal.org/blog/verified-safety-number-updates/>, 2017. Accessed: 22-05-2022.
- [MKMS22] Jose Maria Bermudo Mera, Angshuman Karmakar, Tilen Marc, and Azam Soleimani. Efficient lattice-based inner-product functional encryption. In Goichiro Hanaoka, Junji Shikata, and Yohei Watanabe, editors, *PKC 2022: 25th International Conference on Theory and Practice of Public Key Cryptography, Part II*, volume 13178 of *Lecture Notes in Computer Science*, pages 163–193, Virtual Event, March 8–11, 2022. Springer, Heidelberg, Germany.
- [MM11] Daniele Micciancio and Petros Mol. Pseudorandom knapsacks and the sample complexity of LWE search-to-decision reductions. In Phillip Rogaway, editor, *Advances in Cryptology – CRYPTO 2011*, volume 6841 of *Lecture Notes in Computer Science*, pages 465–484, Santa Barbara, CA, USA, August 14–18, 2011. Springer, Heidelberg, Germany.
- [MMM02] Tal Malkin, Daniele Micciancio, and Sara K. Miner. Efficient generic forward-secure signatures with an unbounded number of time periods. In Lars R. Knudsen, editor, *Advances in Cryptology – EUROCRYPT 2002*, volume 2332 of *Lecture Notes in Computer Science*, pages 400–417, Amsterdam, The Netherlands, April 28 – May 2, 2002. Springer, Heidelberg, Germany.
- [MP16a] Moxie Marlinspike and Trevor Perrin. The double ratchet algorithm, 2016.
- [MP16b] Moxie Marlinspike and Trevor Perrin. The x3dh key agreement protocol. *Open Whisper Systems*, 283:10, 2016.
- [NRS20] Moni Naor, Lior Rotem, and Gil Segev. Out-of-band authenticated group key exchange: From strong authentication to immediate key delivery. In Yael Tauman Kalai, Adam D. Smith, and Daniel Wichs, editors, *ITC 2020: 1st Conference on Information-Theoretic Cryptography*, pages 9:1–9:25, Boston, MA, USA, June 17–19, 2020. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik.
- [Pei14] Chris Peikert. Lattice cryptography for the internet. In Michele Mosca, editor, *Post-Quantum Cryptography - 6th International Workshop, PQCrypto 2014*, pages 197–219, Waterloo, Ontario, Canada, October 1–3, 2014. Springer, Heidelberg, Germany.

Bibliography

- [PM16] Trevor Perrin and Moxie Marlinspike. The double ratchet algorithm. *GitHub wiki*, 2016.
- [PP22] Jeroen Pijnenburg and Bertram Poettering. On secure ratcheting with immediate decryption. In Shweta Agrawal and Dongdai Lin, editors, *Advances in Cryptology – ASIACRYPT 2022, Part III*, volume 13793 of *Lecture Notes in Computer Science*, pages 89–118, Taipei, Taiwan, December 5–9, 2022. Springer, Heidelberg, Germany.
- [PR18] Bertram Poettering and Paul Rösler. Towards bidirectional ratcheted key exchange. In Hovav Shacham and Alexandra Boldyreva, editors, *Advances in Cryptology – CRYPTO 2018, Part I*, volume 10991 of *Lecture Notes in Computer Science*, pages 3–32, Santa Barbara, CA, USA, August 19–23, 2018. Springer, Heidelberg, Germany.
- [PRSS21] Bertram Poettering, Paul Rösler, Jörg Schwenk, and Douglas Stebila. SoK: Game-based security models for group key exchange. In Kenneth G. Paterson, editor, *Topics in Cryptology – CT-RSA 2021*, volume 12704 of *Lecture Notes in Computer Science*, pages 148–176, Virtual Event, May 17–20, 2021. Springer, Heidelberg, Germany.
- [PV06] Sylvain Pasini and Serge Vaudenay. An optimal non-interactive message authentication protocol. In David Pointcheval, editor, *Topics in Cryptology – CT-RSA 2006*, volume 3860 of *Lecture Notes in Computer Science*, pages 280–294, San Jose, CA, USA, February 13–17, 2006. Springer, Heidelberg, Germany.
- [Reg05] Oded Regev. On lattices, learning with errors, random linear codes, and cryptography. In Harold N. Gabow and Ronald Fagin, editors, *37th Annual ACM Symposium on Theory of Computing*, pages 84–93, Baltimore, MA, USA, May 22–24, 2005. ACM Press.
- [Res18] Eric Rescorla. The Transport Layer Security (TLS) Protocol Version 1.3. RFC 8446, August 2018.
- [RMA⁺23] Nathan Reitering, Nathan Malkin, Omer Akgul, Michelle L Mazurek, and Ian Miers. Is cryptographic deniability sufficient? non-expert perceptions of deniability in secure messaging. In *2023 IEEE Symposium on Security and Privacy (SP)*, pages 274–292. IEEE, 2023.
- [RMS18] Paul Rösler, Christian Mainka, and Jörg Schwenk. More is less: On the end-to-end security of group chats in signal, whatsapp, and threema. In *2018 IEEE European Symposium on Security and Privacy (EuroS&P)*, pages 415–429, London, UK, 2018. IEEE.
- [Rog02] Phillip Rogaway. Authenticated-encryption with associated-data. In Vijayalakshmi Atluri, editor, *ACM CCS 2002: 9th Conference on Computer and Communi-*

-
- cations Security*, pages 98–107, Washington, DC, USA, November 18–22, 2002. ACM Press.
- [RST01] Ronald L. Rivest, Adi Shamir, and Yael Tauman. How to leak a secret. In Colin Boyd, editor, *Advances in Cryptology – ASIACRYPT 2001*, volume 2248 of *Lecture Notes in Computer Science*, pages 552–565, Gold Coast, Australia, December 9–13, 2001. Springer, Heidelberg, Germany.
- [SE94] Claus-Peter Schnorr and M. Euchner. Lattice basis reduction: Improved practical algorithms and solving subset sum problems. *Math. Program.*, 66:181–199, 1994.
- [Sha48] Claude Elwood Shannon. A mathematical theory of communication. *The Bell system technical journal*, 27(3):379–423, 1948.
- [Sho94] Peter W. Shor. Polynomial time algorithms for discrete logarithms and factoring on a quantum computer. In Leonard M. Adleman and Ming-Deh A. Huang, editors, *Algorithmic Number Theory, First International Symposium, ANTS-I, Ithaca, NY, USA, May 6-9, 1994, Proceedings*, volume 877 of *Lecture Notes in Computer Science*, page 289. Springer, 1994.
- [Sho00] Victor Shoup. Practical threshold signatures. In Bart Preneel, editor, *Advances in Cryptology – EUROCRYPT 2000*, volume 1807 of *Lecture Notes in Computer Science*, pages 207–220, Bruges, Belgium, May 14–18, 2000. Springer, Heidelberg, Germany.
- [SRCM⁺22] John Scott-Railton, Elies Campo, Bill Marczak, Bahr Abdul Razzak, Siena Anstis, Gözde Böcü, Salvatore Solimano, and Ron Deibert. CatalanGate: Extensive Mercenary Spyware Operation against Catalans Using Pegasus and Candiru. <https://citizenlab.ca/2022/04/catalangate-extensive-mercenary-spyware-operation-against-catalans-using-pegasus-candiru/>, 2022. Accessed: 22-05-2022.
- [SSTX09] Damien Stehlé, Ron Steinfeld, Keisuke Tanaka, and Keita Xagawa. Efficient public key encryption based on ideal lattices. In Mitsuru Matsui, editor, *Advances in Cryptology – ASIACRYPT 2009*, volume 5912 of *Lecture Notes in Computer Science*, pages 617–635, Tokyo, Japan, December 6–10, 2009. Springer, Heidelberg, Germany.
- [SSW20] Peter Schwabe, Douglas Stebila, and Thom Wiggers. Post-quantum TLS without handshake signatures. In Jay Ligatti, Xinming Ou, Jonathan Katz, and Giovanni Vigna, editors, *ACM CCS 2020: 27th Conference on Computer and Communications Security*, pages 1461–1480, Virtual Event, USA, November 9–13, 2020. ACM Press.

Bibliography

- [Sup22] Signal Support. Twilio Incident: What Signal Users Need to Know. <https://support.signal.org/hc/en-us/articles/4850133017242>, 2022. Accessed: 03-10-2022.
- [SZ21] Alessandra Scafuro and Bihan Zhang. One-time traceable ring signatures. In Elisa Bertino, Haya Shulman, and Michael Waidner, editors, *ESORICS 2021: 26th European Symposium on Research in Computer Security, Part II*, volume 12973 of *Lecture Notes in Computer Science*, pages 481–500, Darmstadt, Germany, October 4–8, 2021. Springer, Heidelberg, Germany.
- [Tea21] NTRU Prime Risk-Management Team. Risks of lattice KEMs. <https://ntruprime.cr.yt.to/latticerisks-20211031.pdf>, 2021. Accessed: 28-09-2023.
- [Tel] Telegram. Group Chats on Telegram. <https://telegram.org/tour/groups>.
- [TGL⁺17] Nirvan Tyagi, Yossi Gilad, Derek Leung, Matei Zaharia, and Nikolai Zeldovich. Stadium: A distributed metadata-private messaging system. In *Proceedings of the 26th Symposium on Operating Systems Principles*, pages 423–440, 2017.
- [UG15] Nik Unger and Ian Goldberg. Deniable key exchanges for secure messaging. In Indrajit Ray, Ninghui Li, and Christopher Kruegel, editors, *ACM CCS 2015: 22nd Conference on Computer and Communications Security*, pages 1211–1223, Denver, CO, USA, October 12–16, 2015. ACM Press.
- [UG18] Nik Unger and Ian Goldberg. Improved strongly deniable authenticated key exchanges for secure messaging. *Proceedings on Privacy Enhancing Technologies*, 2018(1):21–66, January 2018.
- [US] US National Security Agency. Announcing the commercial national security algorithm suite 2.0. https://media.defense.gov/2022/Sep/07/2003071834/-1/-1/0/CSA_CNSA_2.0_ALGORITHMS_.PDF.
- [VGIK20] Nihal Vatandas, Rosario Gennaro, Bertrand Ithurburn, and Hugo Krawczyk. On the cryptographic deniability of the Signal protocol. In Mauro Conti, Jianying Zhou, Emiliano Casalicchio, and Angelo Spognardi, editors, *ACNS 20: 18th International Conference on Applied Cryptography and Network Security, Part II*, volume 12147 of *Lecture Notes in Computer Science*, pages 188–209, Rome, Italy, October 19–22, 2020. Springer, Heidelberg, Germany.
- [Wei19] Matthew A Weidner. Group messaging for secure asynchronous collaboration, 2019.
- [Wha20] WhatsApp. WhatsApp Encryption Overview Technical white paper, v.3, oct 2020. <https://www.whatsapp.com/security/WhatsApp-Security-Whitepaper.pdf>.
- [Wha23] WhatsApp. How to create and invite into a group. https://faq.whatsapp.com/3242937609289432/?cms_platform=web, 2023. Accessed: 29-09-2023.

- [WKHB21] Matthew Weidner, Martin Kleppmann, Daniel Hugenothe, and Alastair R. Beresford. Key agreement for decentralized secure group messaging with strong security guarantees. In Giovanni Vigna and Elaine Shi, editors, *ACM CCS 2021: 28th Conference on Computer and Communications Security*, pages 2024–2045, Virtual Event, Republic of Korea, November 15–19, 2021. ACM Press.
- [WPBB23] Théophile Wallez, Jonathan Protzenko, Benjamin Beurdouche, and Karthikeyan Bhargavan. Treessync: Authenticated group management for messaging layer security. In Joseph A. Calandrino and Carmela Troncoso, editors, *32nd USENIX Security Symposium, USENIX Security 2023, Anaheim, CA, USA, August 9-11, 2023*, pages 1217–1233. USENIX Association, 2023.
- [YEL⁺21] Tsz Hon Yuen, Muhammed F. Esgin, Joseph K. Liu, Man Ho Au, and Zhimin Ding. DualRing: Generic construction of ring signatures with efficient instantiations. In Tal Malkin and Chris Peikert, editors, *Advances in Cryptology – CRYPTO 2021, Part I*, volume 12825 of *Lecture Notes in Computer Science*, pages 251–281, Virtual Event, August 16–20, 2021. Springer, Heidelberg, Germany.
- [YGS23] Tarun Kumar Yadav, Devashish Gosain, and Kent Seamons. Cryptographic deniability: A multi-perspective study of user perceptions and expectations. In *32nd USENIX Security Symposium (USENIX Security 23)*, pages 3637–3654, 2023.
- [YVCC23] Hailun Yan, Serge Vaudenay, Daniel Collins, and Andrea Caforio. Optimal symmetric ratcheting for secure communication. *The Computer Journal*, 66(4):987–1016, 2023.
- [Zhe97] Yuliang Zheng. Digital signcryption or how to achieve $\text{cost}(\text{signature} \ \& \ \text{encryption}) \ll \text{cost}(\text{signature}) + \text{cost}(\text{encryption})$. In Burton S. Kaliski Jr., editor, *Advances in Cryptology – CRYPTO'97*, volume 1294 of *Lecture Notes in Computer Science*, pages 165–179, Santa Barbara, CA, USA, August 17–21, 1997. Springer, Heidelberg, Germany.

Daniel Collins | Curriculum Vitae

Website: <https://dcol.me> | Nationality: Australian | Email: danielpatcollins@gmail.com

Education

PhD (Cryptography) (EPFL, Lausanne, Switzerland) Sep. 2019-Apr. 2024
- Advisor: Prof. Serge Vaudenay, LASEC (lasec.epfl.ch)

B.Sc. (Advanced Mathematics) (Honours) (University of Sydney, Australia) 2015-2018
Mathematics & Computer Science. Result: Honours Class I and the University Medal. Thesis: Byzantine fault tolerant boardroom elections without synchrony. Advisor: AProf. Vincent Gramoli

Research Visits

- FAU, Erlangen-Nuremberg, Germany: Lab of Paul Rösler. Jul. 2023
- CISPA, Saarbrücken, Germany: Lab of Julian Loss. Sept. 2022, Dec. 2022
- MPI for Security and Privacy, Bochum, Germany: Lab of Giulio Malavolta. Aug. 2022
- IMDEA Software Institute, Madrid, Spain: Lab of Dario Fiore. Apr. 2022

Academic Publications

- **K-Waay: Fast and Deniable Post-Quantum X3DH without Ring Signatures.** *Daniel Collins, Loïs Huguenin-Dumittan, Ngoc Khanh Nguyen, Nicolas Rolin, Serge Vaudenay.* USENIX Security 2024. Preprint: <https://eprint.iacr.org/2024/120>

- **WhatsApp with Sender Keys? Analysis, Improvements and Security Proofs.** *David Balbás, Daniel Collins, Phillip Gajland.* ASIACRYPT 2023. Preprint: <https://eprint.iacr.org/2023/1385>. Poster at USENIX Security 2023. Preliminary Version: Best Paper Runner-Up at RECSI 2022.

- **Network-Agnostic Security Comes (Almost) for Free in DKG and MPC.** *Renas Bacho, Daniel Collins, Chen-Da Liu-Zhang, Julian Loss.* CRYPTO 2023. Preprint: eprint.iacr.org/2022/1369

- **On Active Attack Detection in Messaging with Immediate Decryption.** *Khashayar Barooti, Daniel Collins, Simone Colombo, Loïs Huguenin-Dumittan, Serge Vaudenay.* CRYPTO 2023. Preprint: eprint.iacr.org/2023/880

- **Cryptographic Administration for Secure Group Messaging.** *David Balbás, Daniel Collins, Serge Vaudenay.* USENIX Security 2023. Preprint: eprint.iacr.org/2022/1411

- **Near Collision Attack against Grain v1.** *Subhadeep Banik, Daniel Collins, Willi Meier.* ACNS 2023. Preprint: eprint.iacr.org/2023/884

- **A Small GIFT-COFB: Lightweight Bit-Serial Architectures.** *Andrea Caforio, Daniel Collins, Subhadeep Banik, Francesco Regazzoni.* AFRICACRYPT 2022. Preprint: eprint.iacr.org/2022/955

- **Optimal Symmetric Ratcheting for Secure Communication.** *Hailun Yan, Serge Vaudenay, Daniel Collins, Andrea Caforio.* The Computer Journal (2022).

- **Improving First-Order Threshold Implementations of SKINNY.** *Andrea Caforio, Daniel Collins, Ognjen Glamocanin, Subhadeep Banik.* INDOCRYPT 2021. Preprint: eprint.iacr.org/2021/1425

- **Online Payments by Merely Broadcasting Messages.** *Daniel Collins, Rachid Guerraoui, Jovan Komatovic, Petr Kuznetsov, Matteo Monti, Matej Pavlovic, Yvonne-Anne Pignolet, Dragos-Adrian Seredinschi, Andrei Tonkikh, Athanasios Xyghis.* DSN 2020 (Best Paper Runner-Up). Preprint: arxiv.org/abs/2004.13184

- **Anonymity Preserving Byzantine Vector Consensus.** *Christian Cachin, Daniel Collins, Tyler Crain, Vincent Gramoli.* ESORICS 2020.

Selected Presentations

- **Cryptographic Administration for Secure Group Messaging.** Swiss Crypto Day 2023 (8 September 2023).
- **On Active Attack Detection in Messaging with Immediate Decryption.** CRYPTO 2023; NYU Crypto Reading Group (16 August 2023).
- **Network-Agnostic Security Comes for Free in DKG and MPC.** CRYPTO 2023.
- **Near Collision Attack against Grain v1.** ACNS 2023.
- **Real World Deniability in Messaging.** *Joint work with Simone Colombo and Loïs Huguenin-Dumittan.* RWC 2023; FAU Erlangen-Nuremberg (26 June 2023). Preprint: eprint.iacr.org/2023/403
- **K-Waay: Fast and Deniable Post-Quantum X3DH without Ring Signatures.** RWPQC 2023.
- **Secure Messaging: Past, Present and Future.** IMDEA Software Institute (March 29 2022).
- **Anonymity Preserving Byzantine Vector Consensus.** ESORICS 2020 (online).

Academic Service

Program committee: AFRICACRYPT 2024

External reviewer: CRYPTO 2024, EUROCRYPT 2024, CANS 2023, CCS 2023, EUROCRYPT 2023, ASIACCS 2023, PKC 2022, CANS 2022, ASIACCS 2022, ANTS XV 2022, ASIACRYPT 2021, ACISP 2021, CANS 2020, SRDS 2018

Master thesis supervision at EPFL (overseen by Prof. Serge Vaudenay):

- Post-Quantum X3DH (*Nicolas Rolin*, Spring 2022; co-supervised with Loïs Huguenin-Dumittan)
- On Secure Administrators for Group Messaging Protocols (*David Balbás*, August 2021)

Master project supervision at EPFL: Andris Suter-Dörig (with Simone Colombo), David Dervishi, Olivier Becker, Oliver Tran Si An, Nathan Duchesne (with Loïs Huguenin-Dumittan), Sina Schaeffler, Alain Gautschi, Nicolas Rolin, Tanguy Rocher (with Andrea Caforio), Maxence Courtet, Dejan Kovac.

Teaching assistant experience:

- | | |
|-----------------------------------------------------------------------|------------------|
| - Advanced information, computation, communication I (1st year, EPFL) | Autumn 2022-2023 |
| - Advanced cryptography (master, EPFL) | Spring 2021-2022 |
| - Cryptography and security (master, EPFL) | Autumn 2021-2022 |
| - Student seminar: security protocols and applications (master, EPFL) | Spring 2020-2021 |
| - Concurrent algorithms (master, EPFL) | Autumn 2020-2021 |
| - Theory of computation (2nd year, EPFL) | Spring 2019-2020 |
| - Computer Science Project (2nd year, USyd) | 2018 |
| - Introduction to Artificial Intelligence (2nd year, USyd) | 2018 |
| - Distributed Systems (3rd year, USyd) | 2018, 2019 |

Awards

- EDIC Fellowship, EPFL 2019
- Semester 1 Supervisor: Prof. Rachid Guerraoui
- Semester 2 Supervisor: Prof. Serge Vaudenay
- CommBank Cyber Prize for Exceptional Academic Performance in Cyber Security 2018
- University of Sydney Academic Merit Prize 2018
- School of Information Technologies Senior High Honour Roll 2016
- Farrand Science Scholarship 2015
- HSC All-rounder Achievers List member (90+ in all subjects) 2014

Other Employment

- Research assistant** (University of Sydney) Dec. 2018-Aug. 2019
 - First co-author to a paper (Anonymity Preserving Byzantine Consensus) at ESORICS'2020.
 - Contributor to the Datalog compiler Soufflé; web development
- Summer researcher** (University of Sydney) Nov. 2017-Feb. 2018
 - Surveyed some verifiable secret sharing schemes; implementation in golang
- Major development project** (GeoInteractive; University of Sydney) Jul. 2017-Nov. 2017
 - Developed a computer vision driven back-end paired with a React/Node web application leveraging AWS for feature detection and classification of pipe imagery
- Programming intern** (FluidIntel Pty Ltd) Nov. 2016-Feb. 2017
 - Fully internationalised a Ruby on Rails web application and a C# application
- Private tutor** 2015-2018
 - Taught students from years 7-12 in math and English and first/second year undergraduate courses

References

Available upon request.