

Cybersecurity Solutions for Active Power Distribution Networks

THÈSE N° 7484 (2017)

PRÉSENTÉE LE 23 FÉVRIER 2017

À LA FACULTÉ INFORMATIQUE ET COMMUNICATIONS
LABORATOIRE POUR LES COMMUNICATIONS INFORMATIQUES ET LEURS APPLICATIONS 2
PROGRAMME DOCTORAL EN INFORMATIQUE ET COMMUNICATIONS

ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

POUR L'OBTENTION DU GRADE DE DOCTEUR ÈS SCIENCES

PAR

Teklemariam Tsegay TESFAY

acceptée sur proposition du jury:

Prof. P. Thiran, président du jury
Prof. J.-Y. Le Boudec, directeur de thèse
Prof. E. Poll, rapporteur
Prof. G. Manimaran, rapporteur
Prof. M. Paolone, rapporteur



ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

Suisse
2017

Be at war with your vices,
at peace with your neighbours,
and let every new year find you a better man.
— Benjamin Franklin

To my loving parents...

Acknowledgements

First and foremost, I would like to thank my supervisor, Professor Jean-Yves Le Boudec, for accepting me, first as an intern and later as a PhD student. Working with him has been such a special privilege. His continued guidance, moral support and, at times, a necessary push led to the successful completion of my PhD. I will never truly be able to express my sincere gratitude for the father-like support he afforded me. I appreciate his ability to put up with my seemingly unstoppable habit of doing things at the 11th hour. I will always remain grateful for the opportunity and the experience.

I would also like to express my gratitude to Prof. Mario Paolone whose feedback throughout the PhD was crucial. I am also grateful to him for granting me the opportunity to participate in the C-DAX project. I am thankful to Professor Jean-Pierre Hubaux, Dr. Philippe Oechslin and Dr. Ola Svensson for their collaboration on parts of my PhD work.

I would like to thank my thesis committee members: Prof. Erik Poll, Prof. Manimaran Govindarasu, Prof. Mario Paolone for accepting to evaluate my thesis, giving me feedback, and making my defense such a happy ending, and Prof. Patrick Thiran for presiding the committee. Special thanks go to my colleagues at the LCA2 lab: Miroslav, Nadia, Maaz, Sergio, Roman, Wajeb, Cong and Elena. Miroslav was my officemate for five years and has been a great friend. In spite of his absolute disregard for political correctness, he managed to maximally exploit my gaffe-prone remarks to make me look like I was the bad guy in the office. After these many years, I still cannot say whether he was a good or a bad influence. Nadia is the European sister I never had. I thank her for being such a good friend. Maaz reminds me of the down to earth brilliant Indian classmates I had at IIT Bombay. His positive attitude and cheerful personality are qualities I will always remember. I am grateful to Sergio for being so kind to assume my responsibilities at EPFL while I was away visiting my family in the US. I am also thankful to Ramin Khalili for the guidance and friendship during my internship.

I thank my friends at the DESL lab for answering so many of my silly questions related to power systems. I thank Marco for reminding me that Ethiopia does not always defeat Italy by consistently beating me at the Lausanne 20km races.

I am very grateful for all the support and friendly atmosphere from our secretaries, Patricia, Holly, Angela, Danielle, and our system administrators, Marc-André and Yves.

Thanks go to my good Ethiopian friends in Switzerland; especially to Roba, Yohannes, MeKonen, Yonas, Walta, Kedija and Tadiwos. If my friends at EPFL blamed me for missing some of the group outings, I surely used them as an excuse by telling them I had an appointment with my Ethiopian friends. Many thanks to the Ethiopian community in Switzerland for making me

Acknowledgements

feel at home. To my Italian friends Simone and Pietro, many thanks for the friendship and the trips to beautiful Italy.

I am greatly indebted to Tekea, Roman, Seifu, Birhti and Mengistu for their gracious support during baby Nathan's arrival. I cannot thank them enough. Many thanks to Dr. Gebreyohannes for all the guidance since my childhood. His achievements have inspired me to never set a limit on what I can achieve. I thank Kibrom for being such a great uncle. I appreciate his continued care. Yaacob (a.k.a James), is an amazing friend and will always have my unlimited respect.

I am so thankful to God for blessing me with the most awesome siblings Fisseha (a.k.a, Fish), Roman, Genet, Haileselassie, Abrehaley and Berhane. Your confidence in me has always encouraged me to try harder when things get tougher. I know you already know how much precious you all are to me. I love you to no end. To my amazing brother, Fisseha, if there is any one I cannot survive without talking to for more than a few weeks, it is you. Not only have you been the best elder brother in the world since my early childhood, but also you have become my best friend. I have always looked up to you and you are the best role model anyone can wish for. To my loving parents, you will never know the amount of love and respect I have for you. I am grateful for the moral values you instilled in me. I would not have made it this far without your guidance and love. I hope I have made you proud.

I am boundlessly thankful to Tsega for all the burden taking care of Nathan all alone in the US. You are the epitome of what it means to be strong, kind and positive. I still have to learn how you manage to remain friendly and cheerful with everyone at all times. To my handsome son, Nathan, thank you for transforming me from an idealist to a realist. You have grown up to be such a charming boy. My love for you knows no bounds.

Lausanne, 15 December 2016

Teklemariam Tsegay Tesfay

Abstract

An active distribution network (ADN) is an electrical-power distribution network that implements a real-time monitoring and control of the electrical resources and the grid. Effective monitoring and control in an ADN is realised by deploying a large number of sensing and actuating intelligent electronic devices (IEDs) and a reliable two-way communication infrastructure that facilitates the transfer of measurement data, as well as control and protection signals. The reliance of ADN operations on a large number of electronic devices and on pervasive communication networks poses an unprecedented challenge in protecting the system against cyber-attacks emanating from outsiders and insiders. Identifying these different challenges and commissioning appropriate security solutions to counter them is of utmost importance for the realization of the full potential of a smart grid that seamlessly integrates distributed generation, such as renewable energy sources, at the distribution level.

As a first step towards achieving this goal, we perform a thorough threat analysis of a typical ADN automation system. We identify all potential threats against field devices, the communication infrastructure and servers at control centers. We also propose a check-list of security solutions and best practices that guarantee a distribution network's resilient operation in the presence of malicious attackers, natural disasters, and other unintended failures that could potentially lead to islanding.

For the next step, we focus on investigating the security aspects of Multi-Protocol Label Switching - Transport Profile (MPLS-TP), a technology that is mainly used for long-distance communication between control centers and between control centers and substations. Our findings show that an MPLS-TP implementation in Cisco IOS has serious security vulnerabilities in two of its protocols, bidirectional forwarding detection (BFD) and protection state coordination (PSC). These two protocols control protection-switching features in MPLS-TP. In our test-bed, we demonstrate that an attacker who has physical access to the network can exploit the vulnerabilities in the protocols in order to inject forged BFD or PSC messages that will lead to disruption of application data communication.

Third, we consider source-authentication problem for multicast communication of synchrophasor data in grid monitoring systems (GMS). Given resource constrained multicast sources, ensuring source authentication without violating the stringent real-time requirement of GMS is a challenging problem. In our effort to identify a suitable multicast authentication schemes, we set out by making an extensive review of existing authentication schemes and identifying a set of schemes that satisfy some desirable properties for GMS. The identified schemes are ECDSA, TV-HORS and Incomplete-key-set. The comparison metrics are

Abstract

computation, communication and key management overheads. The relatively low message sending rate of PMUs in GMS results in some idle CPU time. This fact enables us to implement an ECDSA variant that uses pre-computed tokens to sign messages. This tweak in ECDSA's implementation significantly improves the computation overhead of ECDSA, making it the preferred scheme for GMS. This finding is contrary to the generally accepted view that public key cryptography is inapplicable for real-time applications.

Finally, we study a planning problem that arises when a utility wants to roll out a software patch that requires rebooting to all PMUs in a grid while maintaining full system observability. We assume a PMU placement with enough redundancy to enable a utility to apply the patch to a subset of PMUs at a time and maintain system observability with the remaining ones. The problem we address is how to find a partitioning of the set of the deployed PMUs into as few subsets as possible such that all the PMUs in one subset can be patched in one round while all the PMUs in the other subsets provide full observability of the system. We show that the problem is NP-complete in the general case. We have provided a binary integer linear programming formulation of the problem. We have also proved that finding an optimal solution to the problem is equivalent to maximizing a submodular set function and have proposed an efficient heuristic algorithm that finds an approximate solution by using a greedy approach. Furthermore, we have identified a special case of the problem where the grid has a radial structure and have provided a polynomial-time algorithm that finds an optimal patching plan that requires only two rounds to patch the PMUs.

Key words: Active distribution network, phasor measurement unit, smart grid, cybersecurity, multicast authentication, key management, real-time application, performance evaluation, patching plan, MPLS-TP.

Résumé

Un réseau de distribution actif (active distribution network (ADN)) est un réseau de distribution d'énergie électrique qui implémente en temps réel un suivi et un contrôle des ressources électriques et du réseau. Le suivi et le contrôle efficace d'un ADN sont réalisés grâce au déploiement d'un grand nombre de senseurs, de dispositifs électroniques intelligents (intelligent electronic devices (IEDs)) et d'une infrastructure de communication à double sens fiable facilitant les transferts de mesures de données ainsi que le contrôle et la protection des signaux. Une bonne performance des opérations de l'ADN sur un grand nombre de dispositifs électroniques et sur des réseaux de communication complexes pose un challenge sans précédent pour la protection du système contre des cyber-attaques venant de l'extérieur comme de l'intérieur. Identifier ces différentes attaques et apporter des solutions de sécurité appropriées pour contrer ces attaques est d'une grande importance afin de réaliser pleinement le potentiel des réseaux smart grid, qui sont en plein développement à cause de la pénétration des énergies renouvelables.

Dans une première étape, nous avons analysé les menaces d'un système automatisé typique d'un ADN. Nous avons identifié toutes les menaces potentielles sur les appareils de terrain, les infrastructures de communication ainsi que les serveurs des centres de contrôle. Nous proposons aussi une liste de vérification des solutions de sécurité et des meilleurs pratiques permettant de garantir les opérations du réseau de distribution résilientes en présence d'attaques malicieuses, de catastrophes naturelles et autres défaillances inattendues qui pourraient potentiellement causer le fonctionnement en mode îloté.

L'étape suivante consiste à analyser les différents aspects de sécurité du Multi-Protocol Label Switching - Transport Profile (MPLS-TP), une technologie très souvent utilisée pour la communication à longue distance entre centres de contrôle et sous-stations. Nous montrons que l'implémentation de MPLS-TP par Cisco IOS présentent de sérieuses failles de sécurité dans deux de ces protocoles, bidirectional forwarding detection (BFD) et protection state coordination (PSC). Ces deux protocoles contrôlent le basculement sur des chemins de secours. Nous démontrons expérimentalement qu'un attaquant qui a un accès physique au réseau peut exploiter cette vulnérabilité du protocole pour injecter des messages BFD et PSC forgés, ce qui pourrait conduire à une perte de communication totale entre certains sites, malgré la redondance du réseau.

Troisièmement, nous nous penchons sur le problème d'authentification de source pour la communication multicast des synchroniseurs de données dans un système de surveillance. Assurer l'authentification de source sans déroger aux exigences strictes du temps réel connais-

Résumé

sant les ressources soumises aux contraintes des sources multicast est un challenge. Dans nos efforts pour identifier un schéma d'authentification multicast convenable, nous commençons par analyser les schémas d'authentification existants et identifions un groupe de schémas satisfaisant certaines propriétés désirables. Les schémas identifiés sont ECDSA, TV-HORS et Incomplete-key-set. La métrique de comparaison est la complexité de calcul, de communication et de gestion des clés. Le taux d'envoi de messages relativement bas des PMUs induit des temps d'inactivité du CPU. Nous avons exploité cette observation pour implémenter une variante de ECDSA qui utilise des jetons pré-calculés pour signer les messages. Cette astuce dans l'implémentation d'ECDSA améliore de façon significative le calcul d'ECDSA et en fait le système de choix. Ce résultat va à l'encontre de l'opinion généralement admise sur la cryptographie à clé publique et son impossibilité à être utilisée dans des applications en temps réel.

Enfin, nous avons étudié un problème de planification de mise à jour, qui se présente quand on veut déployer un correctif logiciel demandant le redémarrage de toutes les PMUs dans un réseau tout en conservant l'observabilité du réseau dans son intégralité pendant le déploiement. Nous supposons qu'il existe une redondance suffisante dans la couverture par PMUs pour qu'on puisse appliquer le correctif à un sous-ensemble de PMUs et maintenir l'observation du système grâce aux PMUs restantes. Le problème que nous avons résolu est comment trouver un partitionnement du groupe des PMUs déployées en sous-groupes possibles de façon à ce que les PMUs d'un sous-groupe puissent être mises à jour tandis que les PMUs des autres sous-groupes permettent une observation globale du système. Nous montrons que ce problème est NP-complet dans le cas général. Nous montrons aussi qu'il y a un algorithme polynomial qui trouve un plan de mise à jour de logiciel avec seulement deux étapes quand le réseau est un arbre.

Mots clefs : Réseau de distribution actif, synchrophaseur, réseau intelligent, cyber-sécurité, authentification multicast, gestion des clés, application en temps réel, évaluation de performance, plan de mise à jour, MPLS-TP.

List of Abbreviations

AAA	authentication, authorisation, accounting
ADN	Active distribution network
BFD	Bidirectional forwarding detection
CA	Certificate authority
CIP	Critical Infrastructure Protection
CPS	cyber-physical system
CPU	Central processing unit
DMZ	Demilitarized Zone
DNO	Distribution network operator
DoS	Denial of service
DSLAM	Digital subscriber line access multiplexer
DTLS	Datagram transport layer security
ECDSA	Elliptic curve digital signature algorithm
FERC	Federal Energy Regulatory Commission
FIPS	Federal Information Processing Standard
GMS	Grid monitoring systems
GOOSE	Generic Object Oriented Substation Event
GPS	Global positioning system
IC	Island controller
ICS	Industrial control system
ICT	Information and communication technologies
IEC	The International Electrotechnical Commission
IED	Intelligent electronic device
IP	Internet protocol
IPS	Intrusion prevention systems
LAN	Local area networks
LDAP	Lightweight Directory Access Protocol
LER	Label edge router
LIDS	Log-based intrusion detection system
LSP	Label switched path
LSR	Label-switching router
MAC	Message authentication code

List of Abbreviations

MPLS	Multiprotocol label switching
MPLS-TP	Multiprotocol label switching - transport profile
NERC	North American Electric Reliability Corporation
NIDS	Network-based intrusion detection system
NIST	National Institute of Standards and Technology
OAM	Operations administration and Maintenance
PDC	Phasor data concentrator
PMU	Phasor measurement unit
PSC	Protection state coordination
RFC	Request for comments
RSA	Rivest-Shamir-Adleman
RSVP	Resource reservation protocol
SCADA	Supervisory control and data acquisition
SE	State estimator
SHDSL	Single-pair high-speed digital subscriber line
SMV	Sampled Measured Values
TCP	Transmission control protocol
TESLA	Timed efficient stream loss-tolerant authentication
TLS	Transport layer security
TPM	Trusted platform module
TSO	Transmission System Operator
TV-HORS	Time valid hash to obtain random subsets
UDP	User datagram protocol
VPN	Virtual private network
WAMS	Wide area measurement system
WAN	Wide area network

Contents

Acknowledgements	i
Abstract (English, French)	iii
List of Abbreviations	vii
1 Introduction	1
1.1 Motivation	1
1.1.1 Challenges in Securing Smart Grid Automation Systems	2
1.1.2 Known Cyber Attacks on Smart Grid Automation Systems	3
1.2 Dissertation Outline	4
1.3 Contributions	5
2 State of the Art	7
3 Cyber-secure Communication Architecture for Active Power Distribution Networks	11
3.1 Introduction	11
3.2 Related Work	14
3.3 Threat Analysis	15
3.3.1 Unauthorized Access	15
3.3.2 Man-in-the-Middle Attacks	15
3.3.3 Rogue Device Installation	16
3.3.4 Denial of Service (DoS) Attacks	16
3.3.5 Malicious Software Patching	17
3.4 Security Solutions	17
3.4.1 Centralized User Authentication	18
3.4.2 End-to-End Secure Delivery of Messages	18
3.4.3 Scalable Key Management	19
3.4.4 Secure Software Patching	19
3.4.5 Tamper-resistant Credential Protection	20
3.4.6 Event Logging and Intrusion Detection	20
3.5 Secure Bootstrapping of a Field Device	21
3.5.1 Device Installation During Normal Operations	22
3.5.2 Device Installation During Emergency Conditions	24

Contents

3.5.3	Back Synchronization of an Islanded communication zone	26
3.5.4	Securing Legacy Devices	26
3.6	The EPFL-Campus Smart Grid Pilot	27
3.6.1	Security Architecture	27
3.6.2	Communication Architecture	27
3.6.3	Lessons Learnt	30
3.7	Conclusion	31
4	Security Vulnerabilities of the Cisco IOS Implementation of the MPLS Transport Profile	33
4.1	Introduction	33
4.2	MPLS-TP Protocol Overview	35
4.2.1	Bidirectional Forwarding Detection (BFD)	35
4.2.2	Protection State Coordination (PSC)	36
4.3	Testbed Description	37
4.4	Attacks on the MPLS-TP Protocol	38
4.4.1	Attackers Capabilities	38
4.4.2	Spoofing Attacks on BFD	38
4.4.3	Spoofing Attack on PSC Messages	41
4.5	Discussion and Countermeasures	43
4.6	Conclusion	44
5	Experimental Comparison of Multicast Authentication for Grid Monitoring Systems	45
5.1	Introduction	45
5.2	Authentication Mechanisms for IP Multicast	47
5.2.1	Asymmetric cryptography based schemes	48
5.2.2	One-time signature (OTS) schemes	48
5.2.3	Message authentication code (MAC) based schemes	49
5.2.4	Delayed key disclosure schemes	49
5.2.5	Signature amortization schemes	50
5.3	Candidate multicast authentication schemes for wide area monitoring systems	50
5.3.1	Elliptic Curve Digital Signature Algorithm (ECDSA)	51
5.3.2	Time Valid Hash to Obtain Random Subsets (TV-HORS)	52
5.3.3	Incomplete-key-set	53
5.4	System setup and evaluation methodology	55
5.4.1	EPFL-Campus Smart Grid Monitoring System	55
5.4.2	Comparison Metrics	56
5.5	Performance evaluation and comparisons	57
5.5.1	Implementation and Parameter Settings	57
5.5.2	Performance results and comparison	60
5.5.3	Support for addition and revocation	63
5.5.4	Impact of the scale of WAMS	63
5.6	Conclusion	64

6 Optimal Software Patching Plan for PMUs	65
6.1 Introduction	65
6.2 PMU Patching Problem	67
6.2.1 State Estimation and Assumptions	67
6.2.2 Observability Rules	68
6.2.3 System Model and Problem Definition	68
6.3 The Sensor Patching Problem (SPP)	69
6.3.1 Set Theoretic Formulation and NP-completeness Proof of SPP	70
6.3.2 BILP Formulation of SPP	72
6.4 The Case of Radial Structured Networks	73
6.5 Approximation Algorithm for Mesh Grid Structure	76
6.5.1 A Greedy Approximation Algorithm	76
6.5.2 Formulation as Submodular Maximization	77
6.6 Simulation Results and Comparisons	80
6.7 Conclusion	82
7 Conclusions	85
Bibliography	96
Publications	97
Curriculum Vitae	99

1 Introduction

1.1 Motivation

Modern life is so intimately dependent on electric energy that a major power grid blackout would amount to billions in economic loss and a massive disruption of critical infrastructures that provide health, transport, communication and other crucial public services. Therefore, it is important to ensure the reliable and stable operation of the electric grid. For this reason, nations are investing heavily to revamp their ageing electric-power infrastructure and transform it into a *smart grid*. One such example is the American Recovery and Reinvestment Act (ARRA) [1] that was enacted in 2009 to invest \$4.5 billion, matched by private funding of \$8 billion, to modernize the US power grid infrastructure.

Smart grid is a generic term encompassing many aspects of the modernization of the electric-power system. The most accommodating definition of smart grid is that it is a blending of traditional electric-power infrastructure with information and communication technology (ICT) infrastructures to realise reliable and efficient electric energy management and use. The ICT infrastructure facilitates the implementation of real-time monitoring and control (automation) systems for the power grid. The automation system enables the power grid to seamlessly integrate distributed intermittent renewable energy sources, such as solar photovoltaic and wind. A power grid automation system requires a large number of sensing devices that continuously measure the state of key components of the grid and actuating devices that receive control commands in response to perceived disturbances that might push the grid outside of its normal operating conditions [2]. Moreover, it requires a high-speed and reliable two-way communication infrastructure to facilitate a real-time transfer of the measurement and control data.

In spite of the tremendous benefits the ICT infrastructure brings to smart grids by enabling the real-time assessment of system conditions and taking corrective measures when necessary, it is also a source of grave concern. The large number of sensing and actuating field devices, IT systems at control centers, as well as all the devices in the communication network, are potential sources of vulnerabilities; thus increasing the grid's susceptibility to cyber attacks.

1.1.1 Challenges in Securing Smart Grid Automation Systems

Like any large and complex IT network, the smart grid automation system provides a vast attack surface that can be exploited by an attacker. Below, we provide some of the main reasons why securing a smart grid automation systems is challenging.

- *Legacy systems:* As much as a smart grid introduces new hardware and software components, it also leverages existing assets that include heterogenous legacy automation devices. However, many of these legacy devices were not designed with cyber-security features and run proprietary software that no longer has maintenance support. Therefore, no matter how secure the new components are, the legacy devices can serve as open gateways for an attacker to gain access to more components of the grid's automation system.
- *Rogue devices:* The complex smart grid automation system deploys diverse types of hardware that come from different international vendors. This affords a unique opportunity for an attacker to compromise the supply chain and then to pre-install malicious code or hardware into a device prior to shipment to a target location and later use it as a backdoor [3, 4].
- *Physical exposure:* Unlike traditional IT systems where devices have some level of physical protection, smart grid field devices are deployed in remote physically exposed locations. Besides, the communication network spans over a large unprotected geographic area. Hence, an attack has unhampered access to physically tamper the cyber-enabled devices or networks. Successfully tampering one such device enables an attacker to exploit the trust relationship this device has with its communicating partners (including those in the control center) and to launch further attacks and to compromise more devices in the automation system.
- *Life span of devices:* Devices in a smart grid are expected to last for more than a decade. However, some devices come with built-in cryptographic systems that might not be secure for this long. Security solutions for such devices should be designed such that they are proactively updatable to adapt to long-term evolutions in security threats [5]. Furthermore, preparing a catalogue of all software and hardware in these devices is important. Keeping track of all software patches and replacing or disabling obsolete hardware components during their long lifetime can be error prone.
- *Resource constraints:* Many smart grid applications have stringent latency constraints on the data they receive from sensing field devices that have limited computational power. Finding an appropriate authentication (and encryption) scheme that secures data from the sensors, without compromising on the security level and still not violating the latency requirement, is a challenge.
- *Human factor:* This refers to all those human aspects and conditions that an attacker could take advantage of to successfully achieve its malicious objectives. The human

factor is arguably the weakest link in the security of a smart grid. No matter how strong the deployed security solutions are, employees tend to become lax and find work-arounds to security measures put in place. Unless employees receive continuous security training and awareness-raising actions, an attacker can use social-engineering techniques to steal an employee's authentication credentials in order to use them to gain access to the grid's automation system [6]. Another aspect of the human factor is an insider attack; attacks from a malicious insider who has intimate knowledge of the deployed defence mechanisms and has privileged access to the network are difficult to detect [3,5].

1.1.2 Known Cyber Attacks on Smart Grid Automation Systems

The frequency of cyber attacks involving smart grid automation systems has been growing recently. Below, we describe only a few of the recently reported attacks.

Aurora Vulnerability: This vulnerability affects systems that control rotating machinery such as turbines and diesel generators. The vulnerability was demonstrated in a controlled environment by researchers at the Idaho National Laboratory in 2007. The researchers showed that an intentional out-of-synch closing of an open circuit breaker induces a high electrical torque that puts stress on the mechanical components of rotating equipment in generators [7]. To exploit the vulnerability, an attacker needs to have an in-depth knowledge of the target power system and gain either physical or electronic access to a protective relay that initiates a circuit breaker to open/close [8].

Stuxnet: Stuxnet is a computer worm that targets the Siemens SIMATIC WinCC SCADA system. The worm exploits several zero-day vulnerabilities in Windows OS and is believed to have been spread via USB drives. In 2010, Stuxnet was able to take over the PLCs controlling the centrifuges at Iran's nuclear facilities and disrupt the centrifuges speed. It has been shown that, with some modifications, Stuxnet could be tailored as a platform for attacking smart grid SCADA systems [9].

Slammer: Slammer is a malware that targeted the Davis-Besse nuclear power plant in Ohio in 2003 and took off its safety monitoring system for nearly five hours. The breach did not pose a safety hazard because the plant was under maintenance and not in operation at the time. The malware entered the Davis-Besse plant by first infecting an unsecured network of one of Davis-Besse contractors and then by following a T1 line connecting that network to Davis-Besse's corporate network. It was later found that this T1 line was just one of the many connections that bypassed the plant's firewall [10].

BlackEnergy: Blackenergy is a Malware that targets the human-machine interface ("HMI") software of industrial control systems. It is believed that the first ever hacker-caused power-outage in Ukraine on December 23, 2015 was caused by BlackEnergy. The Ukraine attack lasted for several hours and affected up to 225,000 customers in three different distribution networks. The adversary is believed to have gained access using spear-phishing by sending an e-mail with a BlackEnergy Malware attachment to a specific individual within the organization. The Malware was then used to steal a legitimate user's virtual private network (VPN) credentials.

The adversary used the credentials to gain remote access to distribution network's SCADA network and to control the human-machine interface (HMI). Once they gained access to the network, they installed custom firmware on serial-to-Ethernet devices at substations to take them off-line, they used KillDisk to delete the master boot records of hacked systems and they waged a denial-of-service attack on the power companies' telephone systems [7, 11].

Havex: Havex is a malware that uses Remote Access Trojan (RAT) to infiltrate and modify legitimate software in ICS and SCADA systems. It has targeted a number of European companies that develop industrial applications and appliances. The Malware is distributed through three possible means: (1) using spam e-mail, (2) using a watering-hole attack that compromises an intermediary target (the ICS vendor site) in order to gain access to the actual targets and (3) by using trojanized installers planted on compromised vendor sites [12].

Cyberspies: The US smart grid had been penetrated by espionage agents who are suspected of having inserted rogue code (BlackEnergy) in software that controls electrical turbines [13].

1.2 Dissertation Outline

Through the involvement in the Nano-Tera program (<http://www.nano-tera.ch>), we were tasked to design and implement a secure communication infrastructure for the smart grid pilot on the EPFL campus (<http://smartgrid.epfl.ch>). In Chapter 3, we present the thorough threat analysis we carried out and the comprehensive security framework we propose for a typical active distribution network. We also describe the network architecture and security solutions we have deployed for the EPFL smart grid network.

Because of MPLS-TP's support for bounded delay and guarantee for high degree of network availability, it is identified as one of the suitable technologies for long-distance communication in smart grid. As part of the LCA2 laboratory's collaborative work with ABB, we cover MPLS Transport Profile (MPLS-TP) security in Chapter 4. In this work, we specifically focus on the security of two Operations, Administration, and Maintenance (OAM) protocols that facilitate protection switching. Through our literature review, we find there is lack of a unified approach that addresses security issues for these protocols. Our testbed based study of MPLS-TP's implementation in Cisco IOS confirms that there is no support for source authentication for BFD and PSC messages. As a result, we demonstrate spoofing attacks on these protocols and show the disruptive effect of such attacks on MPLS-TP's proper operation.

IP Multicast is the preferred communication paradigm for phasor data in grid monitoring systems (GMS) because of its efficiency for one-to-many communication and because it minimizes setting changes in already deployed PMUs and PDCs when new receivers are added. In spite of multicast benefits, designing a multicast source authentication scheme for time-critical systems such as GMS is a challenging problem especially when the devices are resource constrained. In Chapter 5, we deal with identifying an appropriate source authentication for multicast communication of phasor measurement data in GMS. We present a detailed literature review of existing multicast authentication schemes and perform an experimental

comparison among a selected set of schemes on the EPFL-campus smart grid network. We show that making some changes in the implementation of ECDSA by exploiting the sending rate of PMUs significantly improves ECDSA performance compared to the other schemes, which is contrary to the popular belief that public key cryptography is not applicable for real-time applications.

In Chapter 6 we address the software patch planning problem. The problem arises when a utility wants to roll out a software patch to all PMUs deployed in a smart grid by patching only a subset of PMUs at a time while maintaining full grid observability using the PMUs that are not being patched. We use set theoretic formulation to model the problem as an instance of a sensor patching problem and prove that it is NP-complete. For the special case where the active configuration of a power grid is a tree, we show there is a polynomial-time algorithm that finds an optimal software patching plan. We also present results from a heuristic algorithm we propose for the general case.

Finally, Chapter 7 concludes by providing a summary of our findings.

1.3 Contributions

A list of the main contributions of this thesis are provided below.

- We perform a thorough cyber-threat analysis of a smart grid network. We consider threats that come from malicious outsiders and insiders. In our analysis we consider all possible assets that can be exploited by an attacker.
- We propose a check-list of security solutions and best practices for an ADN to counter the identified threats. The solutions entail mechanisms to prevent an attacker from exploiting emergency situations that cause an islanded communication zone to install a rogue devices that could be used as a backdoor later on. We have also built a secure communication infrastructure for the EPFL-campus smart grid pilot using our proposed security solutions for ADN as guidelines.
- We built a testbed to evaluate the security of MPLS-TP's OAM protocols. We demonstrate that MPLS-TP implementation in Cisco IOS lacks support for source authentication for BFD and PSC messages. We exploit this vulnerability to launch spoofing attacks against both protocols and we demonstrate the devastating consequences of such attacks.
- Through a qualitative comparison of existing multicast authentication schemes, we identified a set of schemes that satisfy some of the critical requirements for grid monitoring systems (GMS). We implement the identified schemes and experimentally compare their performance at the EPFL-campus smart-grid pilot.
- We have shown that implementing ECDSA in such a way that it uses pre-generated tokens for signature generation significantly improves its performance making it the

scheme of choice for GMS.

- We model a software patch planning problem for PMUs in a smart grid as an instance of a sensor patching problem and prove that it is NP-complete. We also formulate the problem as a binary integer linear programming (BILP) problem and used an ILP solver to find (sub)optimal solutions for small network sizes. Moreover, we have proved that finding an optimal solution to the problem is equivalent to maximizing a submodular set function and proposed a heuristic algorithm to find an approximate solution to the problem and compare the approximate results with those obtained from the ILP solver.
- For the special case when the grid is a tree, we have provided a polynomial-time algorithm that computes an optimal software patching plan that patches all the PMUs in only two rounds.

2 State of the Art

The increasing number of cyber attacks on critical infrastructures has enabled cybersecurity for smart grid automation systems to take centerstage. The awareness of the vulnerabilities of the grid and the dangers that cyber attacks pose on such a critical infrastructure have led governments, standard bodies and the research community to develop standards, guidelines and tools that are necessary for a cyber-secure smart grid.

In the United States, the Federal Energy Regulatory Commission (FERC) is given the authority to approve mandatory cybersecurity standards for the bulk power system (transmission). FERC, in turn, has designated the North American Electric Reliability Corporation (NERC) as the organization responsible for the development of reliability standards. Among NERC's approved reliability standards are the Critical Infrastructure Protection (CIP) standards. The NERC CIP standards address critical cyber-asset identification and their physical security, cybersecurity, and management and control of the electronic security perimeters. The standards also prescribe the recovery plans that must be put in place for the identified critical assets. These standards are mandatory only for utilities in the United States and Canada. Utilities that violate these mandatory standards can be fined up to \$1 million per day by FERC [14, 15].

The Energy Independence and Security Act (EISA) of 2007 [16] directed the National Institute of Standards and Technology (NIST) to coordinate the development of a framework including protocols and model standards that are necessary for a safe and secure smart grid. Under this obligation, NIST published "NISTIR 7628: Guidelines for Cyber Security in the Smart Grid" [17]. The guidelines propose methods for assessing risks in the smart grid, and then identifies and applies appropriate security requirements for mitigating these risks. The guidelines are presented as a non-mandatory framework for utilities to use in developing effective cybersecurity strategies.

The International Electrotechnical Commission (IEC) TC57/WG15 developed the IEC 62351 [18] series of security standards for the high- and low-voltage power-system communication protocols defined by IEC TC 57, specifically the IEC 60870-5 series, the IEC 60870-6 series, the IEC 61850 series, the IEC 61970 series, and the IEC 61968 series [19]. The primary focus

of this standardization is to provide end-to-end security. The IEC 62351-6 standard, for example, specifies security measures for protecting communications between intelligent electronic devices (IEDs) in substations. The standard specifically suggests RSA be used to authenticate IEC 61850 Generic Object Oriented Substation Event (GOOSE) / Sampled Measured Values (SMV) messages that have a 4ms response time. The performance evaluation done for this standard [19] showed that both software and hardware solutions could not satisfy the performance requirements of the applications. Therefore, the working group is currently looking at new approaches that will likely use symmetric-key based schemes. The ISA 99 working group, in collaboration with IEC TC65/WG10, is working on the IEC 62443 series of cybersecurity standards for industrial automation control systems (IACS). The IEC 62443-2-4 series deals with security requirements for vendors of IACS, which includes smart grid devices [20].

Other standards and technical reports include IEEE 1686 [21] on substation intelligent electronic devices (IEDs) cybersecurity capabilities; NIST Special Publication 800-82 [22], Guide to Industrial Control Systems (ICS) Security; IEEE PC37.240 [23], standard for cybersecurity requirements for substation automation, protection, and control systems; IEEE P1711 [24], trial-use standard for a cryptographic protocol for cybersecurity of substation serial links; IEC TR 62210 [25], Power system control and associated communications - Data and communication security; NIST Special Publication 1108R2 [26], NIST Framework and Roadmap for Smart Grid Interoperability Standards. In addition to the above standards, almost all countries have their own smart grid security-related standards, guidelines and regulatory documents.

Although the different security standards and guidelines are crucial in providing a general blueprint that can serve as a starting point, they are not comprehensive enough to provide a complete solution to all potential security threats. For example, they deal only with what are considered critical assets in the grid and fall short of providing an exhaustive list of all assets that need to be protected [7]. Unlike enterprise IT systems that provide more protection to the important components (central servers) than the client nodes, a power automation system needs to provide equal importance to protecting both critical and non-critical assets. Otherwise, an attacker can exploit the unprotected non-critical assets to gain access to the critical ones. Therefore, in addition to following the standards and guidelines they deem fit for their needs, utilities also need to apply tailor-made security solutions pertinent to their specific environment.

In addition to government agencies and standard bodies, the academic research community has also made significant contributions towards smart grid security. The authors in [3] and [5] treat a smart grid as a cyber-physical system where cyber attacks can cause disruptions that transcend the cyber infrastructure and affect the physical power infrastructure. The Aurora vulnerability, demonstrated by researches at the Idaho National Lab [27], was significant in showing the true cyber-physical nature of a smart grid, i.e., malicious instructions, issued to a protection relay in order to open and close a circuit breaker such that it creates an out-of-phase synchronization of the generator to the grid, cause physical damage to the rotating parts of

the generator.

The pioneering work by Liu et al. in [28] on false-data injection attacks on state estimation shows that an adversary with the knowledge of the power-system model can corrupt a selected set of measurements to introduce arbitrary errors into certain state variables while bypassing existing bad-data detection techniques. This research served as a precursor for several research works that are related to data-injection attacks [29–33].

Bhatti and Humphreys in [34] describe how they coerced a 65-meter yacht off its course, by a kilometer, using their GPS spoofing device. The attack was conducted such that the signals from the spoofing device gradually overrode those from the satellites and fed the receiver false coordinates. [35]. Gong and Li in [35] describe a similar spoofing attack against PMUs in grid monitoring systems (GMS) that use GPS signals for time synchronization. There is also a volume of work that deals with anti-spoofing of GPS signals. The three main mechanisms for protecting against GPS spoofing are cryptography (GPS signal source authentication), signal-distortion detection, and direction-of-arrival sensing. There is not yet any single good solution that can effectively protect GPS receivers from this attack.

The few attacks we describe above demonstrate the importance of source authentication for control messages, measurement data and GPS signals. However, there is no one-size-fits-all authentication scheme that can be used in all cases. For example, the communication paradigm used (unicast vs multicast) dictates the kind of authentication schemes that can be used. Moreover, an application's latency requirement, the scale of the network and a device's computing power put additional constraints on our choice of schemes.

In this thesis, we give particular emphasis to identifying appropriate source authentication solutions for smart grid applications in active power-distribution networks — in Chapter 4 for unicast communication and in Chapter 5 for multicast communication. Note that source authentication is only a small part of a comprehensive defence-in-depth-based security framework that utilities needs to deploy to secure their distribution network. In Chapter 3, we propose different security solutions and best practices for such networks and we implement them in the EPFL-campus smart grid pilot.

3 Cyber-secure Communication Architecture for Active Power Distribution Networks

3.1 Introduction

Conventional power distribution networks are passive and are characterised by unidirectional power flows with a minimum level of centralised monitoring and control strategies. However, the large-scale penetration of embedded distributed energy resources and the introduction of energy storage at the distribution premises is paving way for the emergence of active distribution networks (ADNs). An active distribution network is a distribution network with local energy generation, storage capabilities and bidirectional power flow; it requires more sophisticated active monitoring and control strategies. An active distribution network is divided into a subset of loosely-coupled autonomous regional controllers that can perform monitoring and control actions for their geographical subnetwork [36]. Under normal circumstances, each subnetwork is connected to the main power grid and each autonomous controller is able to cooperate with peer controllers when necessary. Inter-domain communication among autonomous controllers is necessary for detecting unexpected power system failures and other anomalous conditions in adjacent regions or in the main grid.

In most extreme cases, when a controller detects a widespread disturbance or power failure, the active distribution subnetwork within the controller's domain can automatically isolate itself from the grid and continue to operate as an island. The power demand within the island is then supplied by the local energy generation and storage until the island back-synchronises with the grid when the faults are resolved [37]. During this *islanding* process, power flow control and voltage and frequency regulations are carried out by the autonomous island controller (IC) in coordination with sensing and actuating devices deployed within the island.

Figure 3.1 illustrates the cyber-physical nature of a typical active distribution network where the sensing and control cyber infrastructure is superimposed on the physical power system infrastructure to facilitate the sophisticated automation operations (monitoring, control and protection) of the distribution network. A sophisticated automation system at the distribution level requires deployment of a large number of electronic data-acquisition and actuating field devices, which are nonexistent today [2]. Moreover, a high-speed and reliable two-way

Chapter 3. Cyber-secure Communication Architecture for Active Power Distribution Networks




Figure 3.1 – An active distribution network where the sensing and control cyber infrastructure is superimposed on the physical power system infrastructure (adopted from [38]). Different possible islanding configurations are shown such that an island can be a superset of islands depending on where the fault occurs.

communication infrastructure is required to facilitate a real-time transfer of sensor data and control signals.

The increasing reliance of distribution network operations on pervasive electronic automation devices and on communication networks poses an unprecedented challenge in protecting the system against cyber incidents. Cyber incidents can be intentional or unintentional. Unintentional cyber incidents can occur due to natural disasters, system failures or human errors, whereas intentional cyber incidents occur due to deliberate attacks from outsiders or insiders.

An attacker has a wide range of options to compromise a distribution automation. For example, many of the electronic automation (sensing and actuating) devices are field-deployed in remote locations where there is little protection against intruders. Moreover, the com-

munication infrastructure for an active distribution network spans a large geographic area. Hence some of the communication cables are likely to pass through physically insecure locations, thus providing an attacker physical access to the network. Furthermore, grid operators are increasingly adopting IP-based communication standards and commercial off-the-shelf hardware and software in their networks for interoperability and for cost reduction reasons. Such standards and products are well studied by attackers and are known to be vulnerable to network attacks such as IP spoofing and denial of service (DoS) attacks.

Given such a range of vulnerability points, a malicious attacker can launch sophisticated attacks to cause maximum damage on the distribution network. An attacker can, for example, launch a coordinated cyber-physical attack by first physically destroying a critical component of the grid (e.g., one of the distributed generators) and simultaneously (or with very little time difference) attack the communication infrastructure that transfers information about the status of the critical component. This way, the operator will not know about the state of the damaged component and thus will not take any corrective actions. With no corrective actions taken, such an attack can have a cascading effect, causing a blackout. Although not due to a malicious attack, the North-East American blackout of 2003 was caused mainly because of lack of system-state awareness by an operator.

Although both insiders and outsiders can attack a distribution automation system, insider attacks are more dangerous than outsider attacks mainly because an insider has better access privileges and has better information about internal-procedures and potential weak spots in the automation system [5]. In general, protecting a system against insider attacks is very difficult. However, implementing automated security tools and techniques to detect and identify suspicious activities from insiders can minimise the level of damage.

The main contribution of this chapter is to thoroughly assess insider and outsider security threats against a power distribution automation system and propose a check-list of security solutions and best practices to counter such threats. The proposed solution guarantees secure operations even when a sub-domain of the distribution network operates in an islanded mode by preventing outsider attackers and malicious insiders from installing a rogue field device by exploiting the emergency situation.

The rest of the chapter is structured as follows. In the following section we identify possible cyber-security threats in a typical active distribution network. In Section 3.4 we discuss security solutions and best practices that should be implemented to counter the identified security threats. In Section 3.5 we detail a secure device installation mechanism that guarantees only authorised field engineers can install field devices from accredited device manufacturers. We also devise an extension to the scheme that can be used to securely install field devices during an emergency situation when communication with a user authentication facility is not available from the installation location.

3.2 Related Work

Smart Grid security has recently received a lot of attention both from the research community and standardisation bodies. The NISTIR 7628 [17], “Guidelines for Cyber Security in the Smart Grid” standard provides a comprehensive set of guidelines for designing cyber-security mechanisms or systems for the smart grid. The standard proposes methods for assessing risks in the smart grid, and then identifies and applies appropriate security requirements to mitigate these risks. NIST has also released a draft on Cyber Security Framework for critical infrastructure [39], which was available for review in 2013. This draft follows a risk-based approach to secure critical infrastructures, as opposed to the process-based approach proposed by Langner in [40]. The latter approach stresses that maximising *security capability* is a prerequisite for security assurance of a critical infrastructure. The IEC 62351 standard series [18], developed by WG15 of IEC TC57, defines security mechanisms to protect communication protocols for substation systems, in particular, IEC 60870 and IEC 61850. The primary focus of this standardisation is to provide end-to-end security. The Critical Infrastructure Protection (CIP) set of standards [14] developed by the North American Electric Reliability Corporation (NERC) aims at introducing compliance requirements to enforce baseline cyber-security efforts throughout the bulk power system (transmission).

A large number of publications have also addressed smart grid security as a research problem. Research works in [3, 5, 41, 42] define smart grid as a cyber-physical system (CPS) and identify unique security challenges and issues encountered in such systems that are not prevalent in traditional IT security. They also discuss security solutions to address these unique challenges. [43] proposes a layered security framework for protecting power grid automation systems against cyberattacks. The security framework satisfies the desired performance in terms of modularity, scalability, extendibility, and manageability and protects the smart grid against attacks from either Internet or internal network via integrating security agents, security switches and security managements. Metke et al. in [44] propose a security solution for smart grid utilising PKI along with trusted computing. The paper suggests automation tools be used to ease management of the different PKI components such as registration authorities (RA), certificate authorities (CA). A comprehensive survey of smart grid security requirements and possible vulnerabilities and potential cyberattacks is provided in [45] and [46]. They also discuss existing security solutions to counter cyberattacks on the smart grid.

In spite of the rich set of publications and standardisation on smart grid security, no work has, to our best of knowledge, addressed security challenges associated with an ADN’s islanded operation in the presence of a malicious insider. In addition to proposing state of the art security solutions to the well known security issues in an ADN automation system, we also propose a scheme that prevents outsider attackers and malicious insiders from installing a rogue field device by exploiting the emergency situation during islanding.

3.3 Threat Analysis

An appropriate security architecture for an active distribution network can be determined only after a thorough threat analysis of the network architecture, information flow and security of each of the infrastructure's components. Cyberattacks can happen anywhere in a distribution automation system including at field devices (sensing and actuating devices), communication infrastructure (routers, switches etc) and at the control and monitoring centre.

Although different techniques can be used to launch cyberattacks on any of these components, the ultimate goal of an attacker is either to initiate erroneous control actions or to prevent or delay required control actions, thereby disrupting the proper operations of the physical power system. Erroneous control actions can happen either due to compromised sensor data fed to the control centre or due to a malicious injection or modification of the control signal. Likewise, an inability to send timely control signals can happen either due to absence of timely sensor data or due to control signals being maliciously dropped or delayed in the network. In the following, we discuss different possible attack vectors that can be exploited by an attacker to realise the stated goals.

3.3.1 Unauthorized Access

Although most field devices are usually located in a relatively secure location, physical access by an adversary cannot be completely ruled out. Even if devices are physically inaccessible, an adversary can still manage to gain access to a device through the network unless there is a secure perimeter that prevents unauthorised access to the communication infrastructure.

An adversary who gains local or remote access to a field device can reconfigure it such that it behaves in an undesirable way. An adversary can, for example, configure a metering device, such as a PMU, to stream incorrect phasor data so that the controller will have incorrect situational awareness about the system. Moreover, an adversary can misconfigure an actuating device to perform inaccurate actions in response to commands from a controller.

3.3.2 Man-in-the-Middle Attacks

An adversary who intrudes in the communication channel of a distribution network can launch a man-in-the-middle attack by selectively dropping or modifying sensor data (control signals) sent from a field device (controller), thus compromising the availability and/or integrity of message exchanges. A replay attack is another form of the man-in-the-middle attack: an attacker sniffing the communication channel can copy measurement data or control commands and forward them later on. Replay attacks can have catastrophic consequences especially when applied to control signals.

Note that man-in-the-middle attacks on measurement data are effective mainly if the attack is persistent. This is because the system is a dynamic system, i.e., measurement data are

Chapter 3. Cyber-secure Communication Architecture for Active Power Distribution Networks

continuously refreshed by a new set of measurements. Thus the effect of a single man-in-the-middle attack is negligible, especially for synchrophasor measurements that are refreshed several times per second. On the contrary, a single attack on control signals can be catastrophic. For example, a control signal that turns off a switchgear that protects a high-voltage circuit can throw an entire city into a blackout.

3.3.3 Rogue Device Installation

A metering field device, such as a PMU, comprises sensors that sample analogue signals from the power system and a computing component that converts the sampled analogue signals to digital data. An attacker who has physical access to a metering device can tamper with the analogue signals (voltage and/or current waveforms) and provides these wrong signals to the computing part of the field device. Similar attacks also apply to actuators. An attacker can replace an actuator with a rogue one that incorrectly acknowledges it has performed a certain control action, whereas in reality it has not.

Implementing cryptographic solutions that ensure device authentication before any meaningful communication starts can prevent an attacker from installing a field device. However, attacks that involve physical tampering of only the analogue component of field devices are difficult to prevent. The best that can be done to prevent such attacks is to harden the physical protection of the devices. Bad-data detection techniques at the control centre can be employed to filter out bad measurements from rogue sensors. However, it has been shown that existing bad-data detection (BDD) techniques do not always detect all bad measurements. Liu et al. [28] have shown that an intelligent adversary with knowledge of the power system model can corrupt a carefully selected set of sensor data to introduce arbitrary errors in the estimates of certain state variables without triggering an alarm from the BDD. A wrong state estimator output can, for example, falsely indicate a significant voltage drop (surge) in a bus, triggering the utility to inject more (less) reactive power to the bus, which may in turn have a catastrophic effect on the stable operation of the grid [47].

3.3.4 Denial of Service (DoS) Attacks

An attacker who manages to gain access to the communication infrastructure, either remotely or locally, can launch a denial-of-service (DoS) attack by flooding a critical link with bogus traffic or by saturating the computing resources of a critical network device such as a router or metering field device. Such an attack causes real-time measurement data from field devices to be delayed or at worst dropped. As a result, a DNO will not have a complete view of the distribution network's status, leading to incorrect decision making. Likewise, the attack can also delay or drop critical control signals from a controller.

3.3.5 Malicious Software Patching

Smart grid devices, such as PMUs, run software and firmware that need to be updated in order to patch bugs, to fix security vulnerabilities or to add new features for better usability or performance. Unless necessary authentication and integrity checks are performed during update, an attacker can use deceptive methods to install a malicious code (a malware) that masquerades as a legitimate software update. What is worse, a malicious insider (field engineer) can deliberately install compromised software update to field devices.

A malicious code (malware) can be used by an attacker to perform any kind of malicious activities. For example, it can be implemented as a “logic bomb” such that it runs in parallel to the legitimate code and sets off a malicious function when a specified condition is met. Stuxnet [48] is one such example of a sophisticated logic bomb believed to be designed to attack Iran’s nuclear facilities by specifically targeting Programmable Logic Controllers (PLCs) made by Siemens.

3.4 Security Solutions

The cyber threats discussed in the previous section are by no means exhaustive, but they serve to illustrate risks to help us develop a secure distribution network. The first step towards securing a distribution network is to separate the automation network from the enterprise network of a DNO and to maintain a secure perimeter around the automation network. A security perimeter is achieved by using a security gateway (a perimeter firewall) that provides a protective barrier from incoming (outgoing) traffic to (from) the automation network. Moreover, internal firewalls should also be used to provide more specific protection to certain parts of the automation network. All firewalls should be deployed with tightly configured rule bases such that the default policy is to “deny everything”, and then open up only what is needed (maintain a white list). Figure 3.2 depicts a logical positioning of firewalls in a typical distribution automation network.

Maintaining a secure perimeter and deploying firewalls is not sufficient to secure a distribution automation network for two reasons. First, security perimeters can fail, either due to misconfiguration or due to inherent weaknesses in the defence mechanism of the firewall. Second, a distribution network spans a large geographic area. Hence, it is impractical to define the perimeter as an attacker has a large attack space to physically connect to the distribution network and launch the attack from within the network.

Therefore, it is desirable to design a security framework that prevents attacks that emanate both from within the distribution network and from external networks. To address the security threats discussed in the previous section, we propose a set of security solutions and best practices discussed below.

Chapter 3. Cyber-secure Communication Architecture for Active Power Distribution Networks

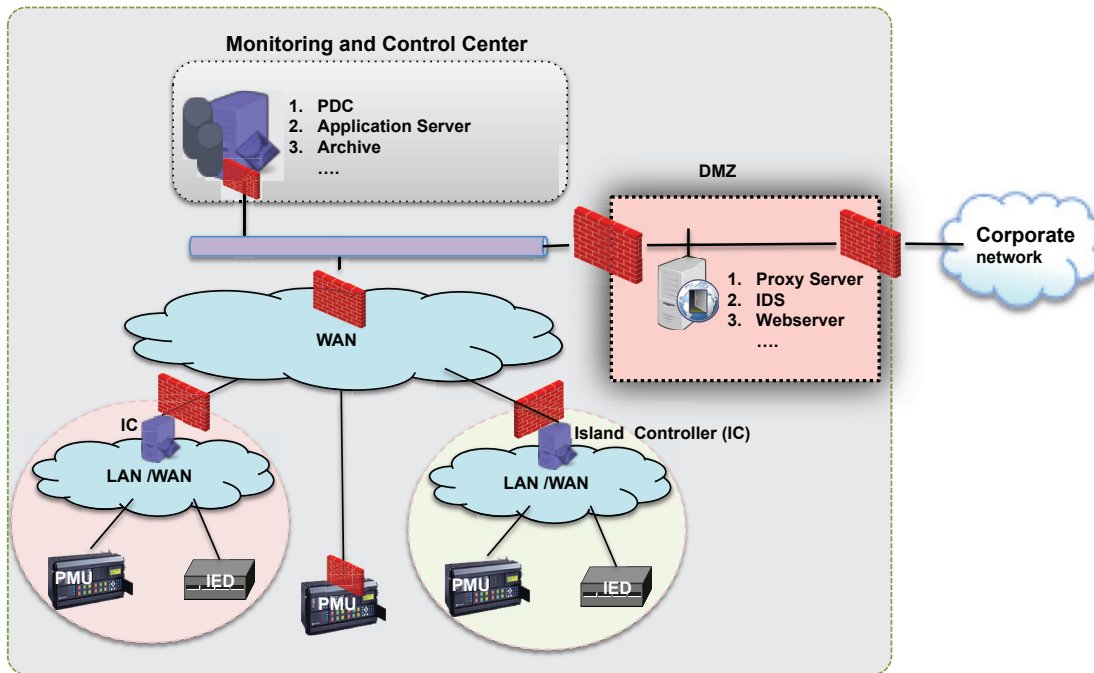


Figure 3.2 – Logical positioning of firewalls in a distribution automation network.

3.4.1 Centralized User Authentication

Access to all devices and services should be limited only to authorised personnel. Each person authorised to access a device or a service has to have a separate user account and a secure password. All user accounts are centrally managed in a central authentication, authorisation, and accounting (AAA) server. All standard security policies such as role-based access control, putting a limit on the number of unsuccessful access attempts, specifying password strength rules, etc should be enforced.

Creating and managing user accounts in a central server reduces the burden of creating and managing several accounts in each device for every authorised employee. A user's account can also be blocked from a single location when necessary. An employee's account can be blocked when he is no longer responsible for the tasks he was initially assigned to, when he leaves his job or when he is suspected as malicious based on a postmortem analysis of activity logs.

3.4.2 End-to-End Secure Delivery of Messages

Guaranteeing end-to-end security for message exchanges is essential for preventing man-in-the-middle attacks and for detecting messages from rogue devices. End-to-end security encompasses guaranteeing the confidentiality, integrity, source authenticity and freshness of measurements, control signals and other important message exchanges at all layers. Although confidentiality is not a critical requirement for measurement and control messages, a distribution network operator (DNO) may want to protect its sensor data's confidentiality in case such

data contains information sensitive to the market that could be exploited by competitors.

Time-stamping, which is already part of existing SCADA communication protocols, is used to guarantee message freshness. For protocols that do not support time-stamping, sequence numbers can be used as an alternative. A systematic use of IPsec, (D)TLS or other standard protocols can guarantee message source authenticity, integrity and confidentiality.

3.4.3 Scalable Key Management

Secure end-to-end communication depends on the existence of a secret key shared between communicating parties. Manual provisioning of such keys and updating them when necessary in a smart grid network, where there is a large number of communicating devices, can be unsafe and cumbersome. Therefore, it is crucial to design a secure and scalable key management scheme to generate, distribute and update the shared cryptographic keys. NISTIR 7628 [17], the foundation document for the architecture of the US Smart Grid, mentions key management as one of the most important research areas in smart grid security.

There is a general consensus in the smart grid research community that Public Key Infrastructure (PKI) is a viable solution as a key management scheme [44, 49]. For distribution automation systems, a DNO should support its own PKI architecture and be responsible for its devices' certificate management. Each communicating device in the distribution network is issued a digital certificate during installation by the DNO's certificate authority (CA). The exact procedure of how a DNO's certificate authority issues a certificate to a device is described in Section 3.5.

Once devices are issued digital certificates, they authenticate each other's identities using standard protocols such as Transport Layer Security (TLS). Following the authentication phase, the communicating parties use a key agreement protocol such as Diffie-Hellman to derive a session key that is used to secure messages exchanged during the TLS session.

A device requires the public key of the DNO's certificate authority (trust anchor) to verify the other party's certificate. Therefore, devices have to store the root CA's public key in a secure location where an adversary cannot delete or modify it. Protecting such sensitive information using file system permissions can be bypassed. An alternative and more efficient solution to protecting sensitive information such as cryptographic keys is to use tamper-proof, special-purpose hardware tokens such as the Trusted Platform Module (TPM).

3.4.4 Secure Software Patching

Attacks that exploit software patches in order to inject malicious code (malware) can be thwarted by requiring a device to validate the authenticity and integrity of any software prior to installation. A DNO has to have its own *approval body* that approves and signs software patches from device manufacturers or third party developers. Whenever a device in the DNO's

Chapter 3. Cyber-secure Communication Architecture for Active Power Distribution Networks

network installs a software patch, it has to first verify that the patch is signed by a DNO's approval body.

3.4.5 Tamper-resistant Credential Protection

Most field devices are deployed in remote geographic locations exposed to unauthorised physical access. Therefore, it is important to provide protection against unauthorised modification and disclosure of sensitive information, such as digital certificates and cryptographic keys, in these devices. An efficient solution to provide the required level of protection for keying materials within field devices is to use a FIPS140-validated tamper-resistant, special-purpose cryptographic module, such as Trusted Platform Module (TPM). A TPM is a secure crypto-processor that offers functionalities for secure generation and storage of cryptographic keys [50]. In addition to serving as tamper-proof storage to sensitive data like cryptographic keys and digital certificates, [44] discusses additional security benefits of using TPM for smart grid devices. Some of the benefits include secure software upgrade, high assurance booting, dynamic attestation of running software and device attestation.

3.4.6 Event Logging and Intrusion Detection

Even after the above security solutions are put in place, there can still be security incidents. Incidents could happen because an attacker installs a malware by exploiting zero-day vulnerabilities, which are inevitable in software. Incidents could also happen because of a field engineer's negligence to follow a DNO's security policy that prohibit the usage of removable media, such as USB, without a proper check for malware prior to use. Besides, disgruntled insiders can abuse their privileges to perform malicious operations.

To minimise the risks that result from such incidents, a DNO should implement automated intrusion-detection techniques to monitor events that occur in the network and to analyse them for signs of suspicious activities that violate the DNO's security policies and acceptable practices.

One type of intrusion detection is log-based intrusion detection system (LIDS) [51]. LIDS uses log data from network devices to detect suspicious activities in a device. This intrusion detection requires each device in the network to implement a secure logging mechanism that maintains a record of system events and user activities in the device. Log data must record noteworthy events such as user activity, program execution status, device configuration change, etc. Each log entry for an event must also contain detailed information about the event including identity of the user, time of the event, type of the event, etc.

LIDS should be implemented both at a device level and at a network level. For the network-level detection, devices send duplicates of their log entries to a centralised logging server. A postmortem analysis of the log files (at individual devices and at a central logging server) is used to reconstruct events and detect intrusions. The intrusion detection system can, for

example, identify insiders engaged in suspicious activities and flag them as malicious.

Another type of intrusion detection is called network-based intrusion detection system (NIDS) [51]. NIDS monitors traffic directed towards critical components of the network to detect suspicious traffic patterns such as denial of service (DoS) attacks. The best location for a NIDS is to deploy it in the same location where a firewall is deployed. In general, distribution automation network traffic is more or less predictable and follows regular traffic patterns, compared to network traffic in enterprise systems. Therefore, a network-based intrusion detection for such systems can be very effective in detecting intrusions.

Note that intrusion detection should be combined with automated intrusion prevention systems (IPS) that send an alarm when intrusions are detected and are capable of taking automated prevention measures, such as resetting the connection and blocking traffic from offending IP address where such actions do not have catastrophic consequences on the grid's operations. Moreover, the operator must have proper incident response and disaster recovery procedures in place to be able to rapidly recover from any emergency (including a cyberattack) and to mitigate damage caused by such incidents.

3.5 Secure Bootstrapping of a Field Device

This section focuses on secure initialization and certification of a newly installed field device before it starts any meaningful communication. This initial stage of securely bootstrapping a field device is a precursor for the effective implementation of the end-to-end security and secure software patching solutions described in Section 3.4.

A secure device-installation scheme should guarantee that the device comes from one of the trusted manufacturers and that the installation is carried out by an authorised field engineer. In other words, the scheme should prevent a malicious outsider or an insider (field engineer), who is suspected as malicious after postmortem log data analysis, from installing a rogue field device. The installation scheme described below assumes that each field device comes with a certificate pre-provisioned by an accredited manufacturer's certificate authority. Furthermore, we assume that the DNO's controllers, certificate authority and Device Registry (described below) know the public keys of all accredited manufacturers whose devices are installed in the DNO's network.

Our installation scheme puts full trust on an authorised field engineer to initialise a field device by securely loading the public key of the DNO's certificate authority and configuring some parameters such as disabling unnecessary ports and changing insecure default settings. An alternative to this would be for a DNO to have a safe central location where all field devices are received and securely initialised with the DNO's certificates and a field engineer is merely responsible for plugging the device into the network and setting some parameters. We choose the first option because we assume that a DNO might not always have pre-initialised devices that are readily available for use during emergency conditions. Thus we want to make it

Chapter 3. Cyber-secure Communication Architecture for Active Power Distribution Networks

possible for a field engineer to be able to take uninitialised field devices (for example, borrow them from a neighbouring DNO or buy them from the closest vendor available) and securely install these devices to the network whenever required.

3.5.1 Device Installation During Normal Operations

In this subsection we describe the set of procedures required to securely install a field device in a distribution network when communication is possible from the installation location to the DNO's network management centre. The network management centre comprises among other components the AAA server, the DNO's certificate authority and the Device Registry, as depicted in Figure 3.3.

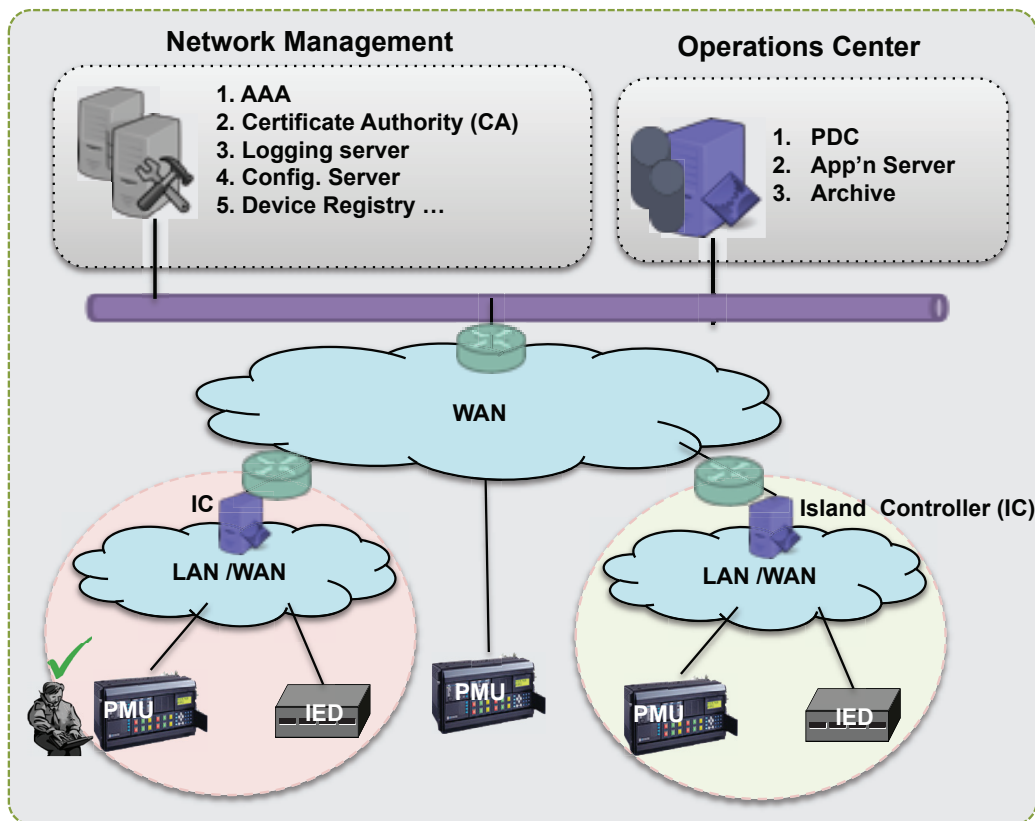


Figure 3.3 – An active distribution network's communication infrastructure and a network management module that facilitates secure communication.

A successful secure installation of a field device entails execution of the following three steps before the device participates in any communicating session.

- A field engineer is authenticated by the central AAA server and obtains an authorisation token for installing the device into the network.
- An authorised field engineer registers the device as a member of the distribution network

3.5. Secure Bootstrapping of a Field Device

in a central database called Device Registry. This database contains a list of all devices in the network and a metadata of each device.

- The device is issued a certificate by the DNO's certificate authority. A certificate is issued only after the CA verifies that the device has a valid certificate from an accredited manufacturer and that the device is registered at the Device Registry by an authorised field engineer.

User authorisation for installing a device can be accomplished by utilising any token/ticket-based standard authentication protocols such as Security Assertion Markup Language (SAML) or Kerberos. In this case we will use SAML to describe how the installation proceeds.

To install a device, an engineer performs the required initial configurations on the device and plugs it into the network. He then authenticates himself to the AAA server and is issued a SAML assertion (SAML security token) by the server. A SAML security token is an XML file that specifies whom it is issued to, what privileges the token holder has (registering a device as a member of the network). The token also contains information about its lifetime (validity period) and a digital signature signed by the token issuer (AAA server) in order to guarantee its integrity.

Once an engineer receives the security token, he initiates the device registration process. The registration proceeds only if the Device Registry verifies that the device comes from a trusted manufacturer and the engineer has the privilege of registering it. The Device Registry verifies the authenticity of the device by using the certificate issued by its manufacturer. The certificate is also used to initiate a secure session with the server. The engineer then sends the device's metadata along with the SAML security token to the Device Registry over the secure channel.

After a successful verification of the token's validity, the Device Registry assigns a unique ID to the device and creates a new entry for the device's metadata in its database. Note that a successful verification of the token guarantees the Device Registry that the engineer is trusted by the AAA server. The Device Registry then confirms a successful completion of the registration by sending back the unique ID to the device.

Upon receiving the unique device ID, the device again authenticates itself to the DNO's certificate authority (CA) and initiates a secure session by using the certificate issued by its manufacturer. A certificate request is then sent to the CA over the secure channel. The CA checks if there is an entry in the Device Registry database corresponding to the device ID that is received as part of the certificate request. If such an entry exists, the CA is convinced that the authenticated device requesting for a certificate is registered by a trusted field engineer. Therefore, the CA signs a new certificate and sends it back to the requesting device.

Now that the device has a certificate issued by the DNO's CA, it can authenticate itself to any communicating partner in the distribution network and initiate secure communication with them using standard protocols such as TLS or IPsec.

3.5.2 Device Installation During Emergency Conditions

When an island controller (IC) detects a widespread disturbance or power failure in the grid, the active distribution subnetwork within the controller's domain can automatically isolate itself from the grid and continue to operate as an island for an extended duration of time. It is possible that portions of the grid's communication infrastructure beyond the island's perimeter could be rendered unreachable as a result of the disturbance that caused the islanding. A subnetwork of a distribution communication infrastructure can also be isolated (islanded) due to a communication breakdown, irrespective of a power system failure. During such emergency situations, a DNO might want to replace some failed field devices within the communication-islanded region. However, if the DNO's network management centre is unreachable from the island, the device installation procedure described above cannot be applied.

Therefore, it is important to design a secure device installation scheme to prevent an attacker from exploiting the emergency situation in order to install a rogue device in the island. In the following we discuss an out-of-band challenge-response-based user-authentication scheme to securely install a device within an island. The scheme utilises the island controller (IC) to serve as a proxy for the security operations required during device installation. For this we assume each island controller knows the public key of the AAA server and the public key of the CA's of all accredited manufacturers whose devices are installed in the network. Furthermore, we assume that each IC is sufficiently secure to be delegated as a subordinate certificate authority for issuing temporary certificates to devices installed within the island during the emergency situation.

With these assumptions, the installation of a device in an islanded network proceeds as follows. The engineer first configures the device and plugs it into the network. Then the device uses the manufacturer issued certificate to authenticate itself and to setup a secure session with the island controller (IC). The device's metadata is then sent to the IC over the secure channel. Before locally registering the device's metadata, the IC replies with a random challenge (nonce) to prove that an authorised engineer is registering the device.

Assuming there exists an out-of-band means of communication (for example, a mobile network) from the island to the network management centre, the engineer authenticates himself to the AAA server using his mobile phone and requests the server for an authorisation token by forwarding the random challenge. Depending on which privileges the engineer has, he receives a signature of the random challenge signed by the AAA server. This signature is sent to the controller as a proof that the engineer is trusted by the AAA server to register a device. The controller then verifies the signature and accepts the device as part of the network by registering its metadata until communication with the network management centre is restored.

If, for some reason, the engineer in the island has lost his password or is unable to login to the AAA server, he can still install the device with the help of any other engineer who is in




Figure 3.4 – Islanding - where a portion of an active distribution’s communication network is cut off from the rest of the main grid’s communication network. A DNO securely installs new devices in the island in the presence of malicious outsiders or suspected insiders who would like to utilize the emergency situation to install rogue devices.

a location where he can communicate both to the network management centre and to the island. The only purpose of the engineer in the island is to forward the random challenge to the second engineer and receive the signature from him to use it in order to finish the registration

Chapter 3. Cyber-secure Communication Architecture for Active Power Distribution Networks

of the device (Figure 3.4). This way, the engineer in the island serves as a delegate to the authenticated engineer for registering the device. Note that the delegation is accomplished without revealing the authenticated user's password to the delegated engineer.

After the device is successfully registered, the island controller issues it with a new certificate. The device uses this certificate to authenticate and to securely communicate with other devices in the island. Other devices can verify the authenticity of the certificate by building a chain of trust starting from the device's certificate up to the root CA (trust anchor) of the DNO. Note that the signing key of the island controller is certified by the root CA and the public key of the root CA is preloaded to every device during installation.

The above description considers a single island controller per island. However, an island can be a superset of multiple islands with each member island having its own island controller. In such a situation, the different island controllers need to run a decision protocol among them to select a "master" controller which will be responsible for the tasks described above.

3.5.3 Back Synchronization of an Islanded communication zone

When the fault that caused islanded communication zone is cleared, the islanded zone synchronises back to the main communication infrastructure. The devices that are installed during an islanded communication are not recognised by the central Device Registry and do not yet have a certificate issued by the root certificate authority. The devices can still continue to communicate using the certificate issued to them by the island controller. However, building a chain of trust to verify such certificates can be complicated during another islanding incident. For example, assume a "master" controller issued a certificate to a device during a previous islanding. Furthermore, assume the device is now in another island that does not contain the previous "master" controller. If the device wants to securely communicate with another partner within the current island, the communicating partner will not be able to build the chain of trust for the device's certificate. To ease this complexity, we propose that each device be re-certified by the root CA, once the connection with the network management centre is restored. The re-certification can be automated as follows. First the IC forwards the temporarily stored metadata of these devices to the Device Registry over a secure channel. The Device Registry creates a new entry for each of these devices in its database. Following this, each such device auto-requests the CA for a certificate. The CA, upon successful verification of the existence of an entry for requesting the device in the Device Registry's database, issues a new certificate to it.

3.5.4 Securing Legacy Devices

The distribution automation network will contain not only new advanced field devices but also legacy devices, which do not have enough computational power or memory space to perform security functionalities. Communication with such legacy devices should be secured by

installing a modern security device, also known as bump-in-the-wire (BITW) device, adjacent to them [43]. The BITW device is issued a digital certificate from the CA on behalf of the legacy device. All security operations on data sent from and received by the legacy device are performed in the BITW device. Note that data transfer between the legacy device and the BITW is not protected.

3.6 The EPFL-Campus Smart Grid Pilot

The threat analysis and the security solutions presented in the previous sections served us as guidelines while building a secure communication infrastructure for the smart grid pilot on the EPFL campus. The smart grid pilot is deployed to monitor and control the medium-voltage electrical grid of the EPFL campus. The grid is a typical example of an active distribution network (ADN) in that it incorporates distributed power-generation (photo-voltaic systems and fuel cells) and energy storage and has a variable demand load.

3.6.1 Security Architecture

As shown in Figure 3.5 The monitoring and control system deploys a total of seven phasor measurement units (PMUs) to measure the state of the grid at different medium-voltage transformers within the campus. The PMUs use timing signals from the Global Positioning System Satellite (GPS) for synchronization. All the PMUs stream the time-synchronized phasor measurements to the Phasor Data Concentrator (PDC) every $20ms$. The PDC correlates phasor data from all the PMUs with equal time stamps and feeds these correlated data to a state estimator (SE), which is deployed within the same machine as the PDC. The SE uses the correlated measurement data to compute the estimated state of the grid in real time.

3.6.2 Communication Architecture

For security and for robustness reasons, the communication network for the smart grid pilot is built on a dedicated infrastructure. We have re-used existing twisted pair cables, originally installed for telephony. Since the twisted pairs are too long to support Ethernet-based communication, instead we use single-pair high-speed digital subscriber line (SHDSL) technology. Therefore, a PMU is connected to ZyXEL SHDSL line terminal (modem) using a short Ethernet cable and the SHDSL modem forwards the data over the long twisted pair cable to a digital subscriber line access multiplexer (DSLAM) router at a central location. The DSLAM serves as a concentration point for all traffic from all the PMUs and forwards it to the PDC over an optical cable.

We put different security mechanisms into place to ensure that the ICT infrastructure of the EPFL smart grid pilot is resilient to insider and outsider cyber-attacks. By deploying these security mechanisms, we aim to achieve the following three main security goals:

Chapter 3. Cyber-secure Communication Architecture for Active Power Distribution Networks




Figure 3.5 – A security architecture of the EPFL-campus smart grid pilot.

- Secure perimeter
- Secure end-to-end delivery of message
- Centralized access control

Secure perimeter: The first security measure we took towards ensuring a secure perimeter is that we build a dedicated communication network that is physically separate from the EPFL-campus public network and from the Internet. There is no direct communication between a device in the smart grid and another one from an external network. A proxy server in the Demilitarized Zone (DMZ) serves as an intermediary node to terminate and forward any valid communication between a device in the ADN and another one in the outside world.

The proxy server also functions as a software repository. It fetches software patches from the legitimate sources on the Internet and makes them available to the devices in the ADN. This guarantees that devices in the ADN don't connect to the Internet directly whenever there is a software patch available for them.

The external firewall (Cisco RV325) and the internal firewall (Juniper SRX100) serve as a two-stage protective barrier. They block all IP addresses and port numbers except those that are explicitly allowed by the system administrator. The external firewall filters traffic between the devices in the DMZ and those on the external network and the Internet firewall filters traffic between devices in the ADN and the DMZ.

Secure end-to-end delivery of messages: We use our own certificate authority (CA) to issue certificates to all devices in the ADN. The devices use their certificates for mutual authentication and to set up a secure communication channel to exchange confidential information (e.g. symmetric session keys).

We implement message source authentication to guarantee end-to-end security for the phasor data communication. As will be discussed, in detail, in 5.4.1, PMUs use multicast to communicate their phasor data with the different receivers within the ADN, namely the PDC and the proxy server in the DMZ. We will discuss in 5.1 why multicast communication paradigm is preferred. To guarantee end-to-end security (message origin authenticity) for the multicast data, we implement ECDSA that uses pre-generated tokens for signature generation.

End-to-end security is also guaranteed for unicast communication between devices in the ADN and those devices in the EPFL-campus network (refer to Figure 3.5). The proxy server being one of the multicast receivers, it forwards the received multicast data from the PMUs to the SCADA server as well as to the Data Historian over two separate secure DTLS channels. Moreover, the SE sends its output to the proxy server over a secure DTLS channel. The proxy server, in turn, uses another DTLS channel to forward this data to a web server outside of the ADN domain. The web server displays a live stream of the SE output to the public.

Centralized access control: Access to all devices within the ADN (PMUs, DSLAM router, Fire-

Chapter 3. Cyber-secure Communication Architecture for Active Power Distribution Networks

walls (Cisco and Juniper routers), SHDSL modems, and servers) is limited to only authorized personnel. The separate user accounts per personnel that are used to access these devices (except the SHDSL modems) are managed in a central openLDAP server that is setup on the same machine where the CA is setup. The certificate authority (CA) issues a certificate to the openLDAP server. This certificate is used by client nodes to verify the authenticity of the LDAP server. The certificate is also used to create a TLS session with the client during user authentication so that the user password is sent over a secure channel.

The DSLAM router and the Firewalls already support LDAP-based user authentication. Thus, we only needed to configure them with the appropriate LDAP server address. More work was need to enable LDAP-based authentication for the PMUs. The PMUs run an NI Linux Real-Time OS, which is a variant of the Angstrom distribution. This OS did not have some of the required client side packages required for LDAP-based authentication. Specifically, the client packages NSS_LDAP and PAM_LDAP where missing from the NI repository. Therefore, we had to cross-compile these packages from the source code and install them in the PMUs. Only then could we configure the PMU to support LDAP-based authentication. The SHDSL modems had no built-in support for LDAP-based authentication. Besides, it was not possible to access the firmware to be able to install cross-compiled modules required for LDAP-based authentication. Therefore, our only means to login to these modems is using the local user account.

The default password of the local accounts in all the devices are changed to strong passwords and these passwords are known to only designated network administrators. The accounts are used only when authentication via the LDAP is not possible.

3.6.3 Lessons Learnt

The most important lesson we learnt from our experience in securing the EPFL smart grid communication network is that it is very difficult to cover all security aspects even for such a small network. We learn that completely separating the operational network from the Internet all the time can be difficult. For example, there was a time when our local software repository was not fully deployed during the evolution of the ADN network. During this time, we had to bypass the DMZ and directly connect the PMUs to the Internet to patch some software bugs (e.g., the Heartbleed bug). Events like this, even if they are done for a very short time, are the exact events that a persistent attacker can exploit to gain access to the ADN network. However, it is not uncommon for network administrators to take temporary lax measures to fix critical issues if such issues could not be fixed when the tight security measures are in place.

We have also gained some insight into how challenging it is to have a heterogeneous set of devices. The fact that we needed to cross-compile some packages for the PMUs and that the SHDSL modems don't support LDAP-based authentication demonstrates how difficult it can be for major utilities that deploy a large number of heterogenous devices from different vendors.

The bottomline is that it is impossible to cover all possible security loopholes. It is inevitable that a motivated attacker will find a means to breach an ADN's network either due to overlooked vulnerabilities or as a result of a personnel's violation of security measures put in place. Therefore, a utility also needs to put proactive incidence response mechanisms to minimize risks associated with successful cyberattacks. Moreover, a utility needs to do a continuous revision of its security policies and adherence to them in order to account for unforeseen vulnerabilities and to strengthen weaker security links.

3.7 Conclusion

A smart grid's communication infrastructure is key to enabling a utility to collect and analyse data about current operating conditions of the grid and issue control signals as required. However, the critical nature of power grid makes its communication infrastructure a suitable target for cyberattacks. Therefore, implementing a comprehensive cybersecurity solution is necessary. We analysed different cybersecurity threats in a typical active distribution network and proposed security solutions and best practices to counter such threats. Our solution entails secure bootstrapping of field devices such that only an authorised personnel is able to install such devices and no malicious insider or outsider is able to install rogue field devices. We have also used our proposed solutions as a guideline to build a proof-of-concept secure communication architecture for EPFL-campus smart grid network.

4 Security Vulnerabilities of the Cisco IOS Implementation of the MPLS Transport Profile

4.1 Introduction

The MPLS Transport Profile (MPLS-TP) is an extension of MPLS standards that is compatible with already deployed IP/MPLS. In addition to adopting the quality of service (QoS) mechanisms like bounded delay defined within MPLS, MPLS-TP defines path-based, in-band Operations, Administration, and Maintenance (OAM) protection mechanisms. MPLS-TP OAM ensures high degree of network availability by providing tools needed to monitor and manage the network and to facilitate protection switching [52]. Two OAM protocols defined by MPLS-TP OAM are bidirectional forwarding detection (BFD) and protection state coordination (PSC). While BFD is responsible for detecting Label Switched Path (LSP) data plane failures, PSC handles protection switching.

MPLS-TP is identified as a promising Packet Switched Network technology for smart grid networks [53, 54]. MPLS-TP is suitable for long-distance communication between substations and control centers or between control centers. Since MPLS-TP can operate on non-IP Layer 2 Ethernet networks, it is suitable to transport smart grid data, like the time-critical IEC 61850 Generic Object Oriented Substation Event (GOOSE) and Sampled Measured Values (SMV) messages from substations to control centers.

Securing the communication infrastructure in smart grid networks is challenging. One of the reasons why it is challenging is because the communication network spans a large physically unprotected geographic area. This makes the network prone to man-in-the-middle attacks by leveraging physical sabotage and tampering of unprotected devices such as routers. As stated above MPLS-TP in the context of smart grid is used for long distance communication. Hence an MPLS-TP network is also prone to similar attacks.

Our literature review of different MPLS-TP related standards and RFCs reveals that cybersecurity is not given due attention. We find that the different RFCs [55–59] that marginally address security issues provide fragmented and incomplete pieces of information which makes it difficult to draw exactly which solutions to use and where. RFC 5085 [59] for example proposes

Chapter 4. Security Vulnerabilities of the Cisco IOS Implementation of the MPLS Transport Profile

that IPsec be used to secure the MPLS-TP data plane. However, as stated above, there are some smart grid applications that are not IP based. Hence IPsec is not applicable for such systems. This confusion in the security solutions for MPLS-TP led us to believe that many vendors may not provide proper security for their MPLS-TP implementation. In this chapter we present our results on the security analysis of MPLS-TP implementation in Cisco IOS. We decide to study this specific implementation because Cisco is a leading network equipment manufacturer and is heavily investing in smart grid.

From our experimental analysis we found that spoofing attacks are possible against the OAM (BFD and PSC) protocols with devastating effects. While performing the attacks, we assume an attacker gains physical access to an MPLS-TP network cables. In a smart grid network that employs MPLS-TP, the fiber-to-Ethernet converters or optical terminators are usually located at remote substations which are usually unmanned. Besides, pole mounted optical repeaters along transmission lines are common. Therefore, an attacker can easily gain physical access to such facilities to perform his cyber attack.

Besides the easy access to the physical network, the kind of equipment an attacker needs to launch the spoofing attacks are quite cheap. For cases where there are only optical cables in the network, NetFPGA-10G card [60] with SFP+ modules can be used to connect an attacker's device to the network. For those networks that have optical terminators or fiber-to-Ethernet converters that interface routers to optical cables, an attacker needs only an Ethernet switch and a laptop. In our study, we consider networks with fiber-to-Ethernet converters. However, our attacks equally apply to the case where there are only optical fibers in networks.

We studied a total of three spoofing attacks. Two of the attacks apply to the BFD protocol and one applies to PSC. One of the attacks on BFD requires physical access to the network at two locations. The second attack on BFD and the attack on PSC both require physical access at only one location. The spoofing attack on BFD that requires physical access at one location falsely convinces label-edge routers (LERs) that there is no protection LSP to switch to when a fault occurs in a working LSP. The second attack that requires access at two locations hides the presence of a link failure and stops the LER from taking a corrective measure by switching to a protection LSP. The spoofing attack on PSC disables both the working as well as the protection LSPs using forged PSC packets that emulate network operator issued commands to lock out the LSPs. The forged packets are inserted only from one location. This shows that the attack is worse than a physical sabotage that cuts a wire. This is because while cutting a wire affects only the LSP that makes use of the link, our spoofing attack locks out both the working and the protection LSPs. In a way, this is contrary to the increased availability objective smart grids intend to achieve, i.e., by introducing a communicating infrastructure, the power grid becomes more vulnerable to more devastating attacks than a physical sabotage.

The rest of the chapter is organized as follows: Section 4.2 briefly discusses MPLS-TP features and the OAM protocols. In Section 4.3 the testbed used to perform the spoofing attacks is introduced. Section 4.4 describes how the attacks are carried out and what the consequences

of the attacks are. In Section 4.5 we discuss countermeasures that can be applied to thwart the described attacks. Section 4.6 concludes the chapter.

4.2 MPLS-TP Protocol Overview

In the context of MPLS, a label-switched path (LSP) is a unidirectional point-to-point connection between two routers. The routers at the end of the connection are called label-edge routers (LERs) and the intermediate nodes that are capable of label switching are called label-switching routers (LSRs). In contrast, LSPs in MPLS-TP are bidirectional. They are comprised of two unidirectional MPLS LSPs, one in each direction. The unidirectional LSPs are congruent and traverse the same path in both directions and are managed as a single entity. This MPLS-TP feature is called co-routed bidirectional.

MPLS-TP sets the LSPs either using a static control plane via a network management system (NMS) or a dynamic control plane, which is optional [61]. Network operators, usually prefer the static control plane because it does not involve interactions with protocols such as OSPF-TEE, RSVP-TE, BGP etc; hence it is considered more secure than the dynamic control plane. For our experiment, we use static control plane to set up the LSPs.

One of the most important enhancement of MPLS-TP over MPLS is the use of OAM-triggered network recovery [62]. Network recovery is the ability of a network to recover traffic delivery following failure or degradation, of network resources. Although MPLS-TP supports different recovery types [57], we consider only protection switching. Protection switching enables fast repair from a fault condition. It can also be triggered by operator-issued commands. It is a fully allocated survivability mechanism, meaning that the route and resources of the protection (backup) LSP are reserved for a selected working LSP [56]. Protection switching is facilitated by BFD and PSC protocols. BFD fast detects failures (within 12ms in Cisco's implementation) and PSC coordinates the switching from a working LSP to a protection LSP. The overall time required to detect a failure and switch to a protection LSP is less than 50ms.

4.2.1 Bidirectional Forwarding Detection (BFD)

BFD [63] is an OAM protocol that provides fast fault detection in an LSP. A separate BFD session is created between the LERs for each of the two bidirectional (the working and protection) LSPs. Once a BFD session is established, an LER sends BFD control packets every 3.33ms as per the RFC6372 [57] recommendation. The Cisco implementation we study uses a 4ms inter BFD control packet interval. BFD control packets are switched in-band in the same LSP with the data packets. If an LER observes that three BFD packets in a row are missing, it declares the BFD session is down. It informs the remote LER about the down state of the unidirectional LSP by sending a "Session Down" control packet via the outgoing unidirectional LSP. On receiving the "Session Down" packet, the remote LER is expected to trigger protection switching.

Chapter 4. Security Vulnerabilities of the Cisco IOS Implementation of the MPLS Transport Profile

Ver	Diag	Sta	P	F	C	A	D	M	Detect Mult	Length
My Discriminator										
Your Discriminator										
Desired Min TX Interval										
Required Min RX Interval										
Required Min Echo RX Interval										
Auth Type		Auth Length			Authentication Data					

Figure 4.1 – Format of a BFD Control Packet.

Figure 4.1 describes a BFD control packet format. Out of the many fields shown in the figure, we manipulate only two of them to launch the attacks described in Section 4.4.2. One is the five-bit diagnostics (*Diag*) field that is used to report failure conditions to a remote LER. Out of the 32 possible values, we consider only two of them - 0 for when there is no fault to report and 1 for “Control Detection Timer Expired”. The “Control Detection Time Expired” is used when an LER misses three BFD control packets in a row. The second field we use is the two-bit state (*Sta*) field. The *Sta* field describes the state of the BFD session. *Sta* is set to zero (“Administrative Down”) when a network operator issues a “Down” command and *Sta* is set to one if the session is “Down” for nonadministrative reason. *Sta* is set to two (“Init”) when an LER brings a BFD session up and the remote LER has not replied yet and *Sta* is set to three (“Up”) when the session is up and running.

4.2.2 Protection State Coordination (PSC)

A PSC is an in-band, data-plane-driven signalling protocol that coordinates protection switching. Since MPLS-TP LSPs are bidirectional, PSC ensures that both ends of LSP are switched to the protection LSP (even when the fault is unidirectional). PSC is a single-phased coordination protocol in that the initiating LER performs the protection switchover to the protection LSP and informs the remote LER to do the same. On receiving the message, the remote LER completes the switchover [56].

The PSC control logic maintains information about the current state of the two LSPs on each LER. The state that is maintained contains information about what the current states of the LSPs are, what caused the current state and whether the current state is related to a remote or local condition. There are a total of six possible states supported by the control logic. The PSC control logic in an LER calculates the next state, determines what actions need to be taken (location actions at the LER or messages to be sent to the remote LER), based on the highest priority request received from three potential request sources. One input comes from a local request logic. The local request logic itself receives requests from the OAM (e.g. BFD triggered),

Ver	Request	Prot. Type	Revertive	Reserved ₁	FPath	Path
TLV Length			Reserved ₂			
Optional TLVs						

Figure 4.2 – Format of a PSC Control Packet.

network operator commands, local control plane (if present) and local timers. It then selects the request with the highest priority and passes it on to the PSC control logic. The second input for the PSC control logic comes from request messages received from the remote LER. These messages indicate the status of the LSPs from the viewpoint of the remote LER. The third input comes from the current state of the PSC Control logic.

Out of all the requests, the network-operator commands have the highest priority [56]. There are three possible commands in this category - "Clear", "Lockout of Protection" and "Forced Switch". We make use of these requests in the spoofing attack on the PSC messages described in Section 4.4.3. The attack manipulates the 4-bit request (*Request*) field in an PSC control packet and the 8-bit fault path (*FPath*) field (see Figure 4.2). The *Request* field reflects the current state at the LER that is sending the PSC packet, i.e., it is an expression of the source LER's wish for what the next state should be and the *FPath* field indicates which LSP (working or protection) is identified to be in a fault condition or affected by an administrative command. A *Request* field set to 14 ("Lockout of Protection") indicates an operator has issued command to prevent protection switching.

4.3 Testbed Description

Figure 4.3 describes the network topology we use while evaluating MPLS-TP's security. There are a total of eight routers and each router runs Cisco Cloud Service router 1000V (CSR1000V) IOS, that implements MPLS-TP [64]. Two physical machines, each with eight-core processors and 16GB RAM are used to host the eight routers as virtual machines (four virtual routers per physical machine). A virtual router is assigned one core processor and 3GB of RAM. We make use of the 60-day evaluation license that Cisco provides to all features of the CSR 1000V image at a throughput of 50Mbps.

The testbed is setup as a one-to-one (1:1) protection switching. We statically configure a working and a protection LSP between routers R1 and R8. RSVP is used to reserve network resources. The path R1-R2-R3-R4-R8 forms the working LSP and R5-R6-R7-R8 forms the protection LSP. The virtual links between the virtual routers and the physical links between the physical machines are full duplex and have a 1Gbps capacity. Following MPLS-TP's recommendation, we use co-routed bidirectional LSPs. We assigned 25Mbps of bandwidth for each link in both directions. BFD sessions with a 4ms message interval and a 12ms detection interval are configured for both the working and protection LSPs [65].

4.4 Attacks on the MPLS-TP Protocol

In this section, we discuss three spoofing attacks that a malicious attacker uses to temporarily or permanently disrupt MPLS-TP's normal operations. Two of the attacks target BFD messages and one targets PSC messages.

4.4.1 Attackers Capabilities

The attacker is assumed to be capable having physical access to the MPLS network at most at two locations and is able to insert his own Ethernet switch between a router and an Ethernet-to-Optical convert. We also assume the attacker has a laptop configured with a spoofed IP address and has enough processing power to continuously inject forged packets.

4.4.2 Spoofing Attacks on BFD

Scenario I - Removing Protection from a Target LSP

The goal of this attack is to first force the label edge routers to switch to the protection LSP and then make them continue to believe that they have no longer a backup LSP to switch to in case a fault occurs in the current LSP.

To demonstrate the attack, we need access to the network at only one location in the working LSP. We insert our Ethernet switch at location *a* in Figure 4.4 and connected our laptop to the switch.

After this setup, we proceed with a reconnaissance stage to learn about the MAC addresses of *R2* and *R3*, the TTL value of in the BFD packets, and the BFD session number of the target working LSP between *R1* and *R8*. We use packet sniffing to gather these pieces of information

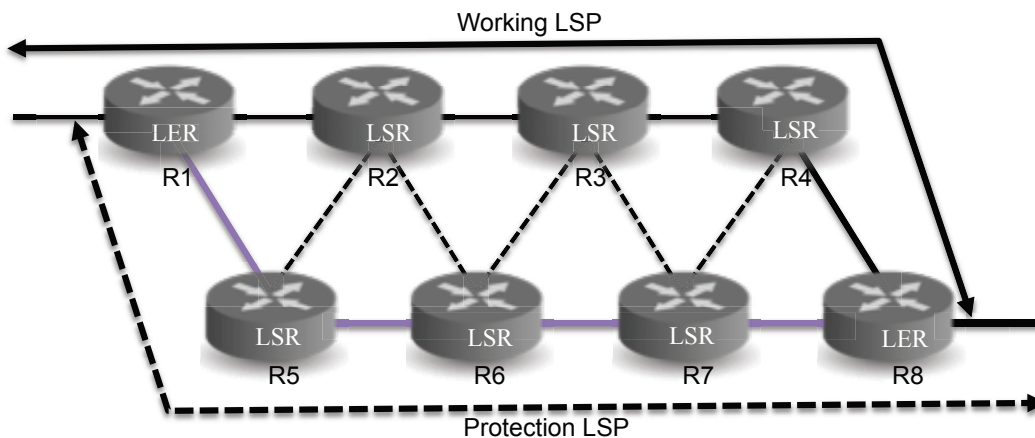


Figure 4.3 – An eight router MPLS-TP testbed in a 1:1 Protection setup.

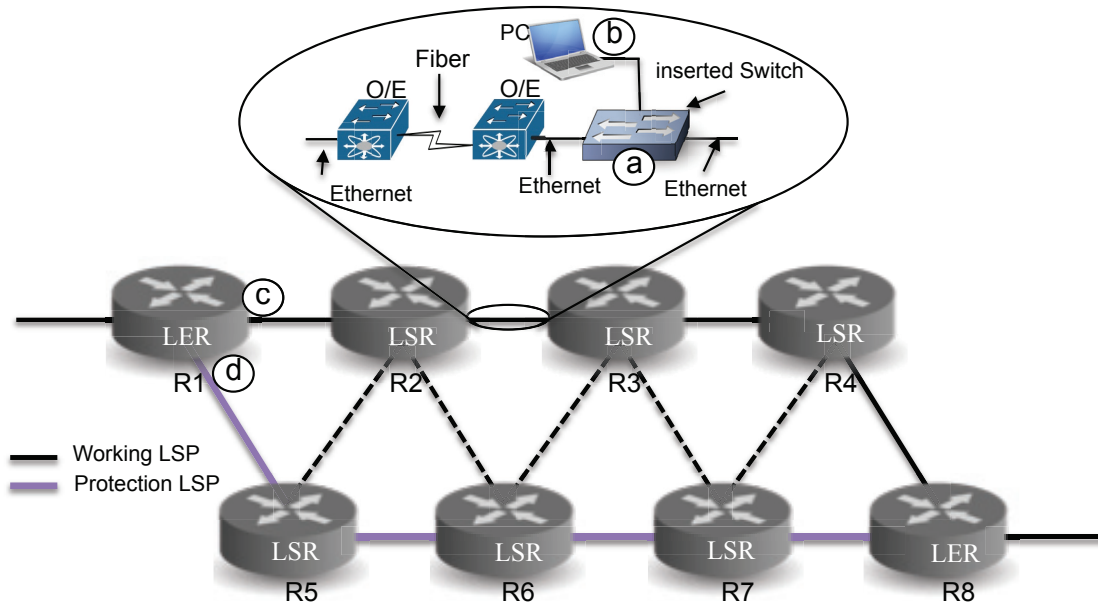


Figure 4.4 – Spoofing Attack on BFD: Removing protection from a target LSP.

from BFD control packets.

After the reconnaissance stage, we use the Scapy [66] tool to create forged BFD control packets in our laptop using the information we gathered. The forged BFD control packets created by setting the *Diagnostic (Diag)* field to “Control Detection Time Expired” and the *Sta* to “Down”. These packets are then injected to the network through the switch. Figure 4.5 shows a Wireshark capture of a forged BFD packet.

```

Time      DeltaTime  Source          Destination      Protocol  Length  Info
0.02701200003132000  VMware_8a:e6:5d  VMware_18:c0:0f  BFD Control  50  Diag: Control Detection Time Expired, State: Down, Flags: 0x28
[+] Frame 12: 50 bytes on wire (400 bits), 50 bytes captured (400 bits) on interface 0
[+] Ethernet II, Src: VMware_8a:e6:5d (00:0c:29:8a:e6:5d), Dst: VMware_18:c0:0f (00:0c:29:18:c0:0f)
[+] MultiProtocol Label Switching Header, Label: 90, Exp: 6, S: 0, TTL: 255
[+] MultiProtocol Label Switching Header, Label: 13 (Generic Associated Channel Label (GAL)), Exp: 6, S: 1, TTL: 1
[+] Generic Associated Channel Header
[+] BFD control message
    001. .... = Protocol version: 1
    ...0 0001 = Diagnostic Code: Control Detection Time Expired (0x01)
    01. .... = Session State: down (0x01)
[+] Message Flags: 0x28 (P, C)
    Detect Time Multiplier: 3 (= 12 ms Detection time)
    Message Length: 24 bytes
    My Discriminator: 0x00000002
    Your Discriminator: 0x00000002
    desired Min TX Interval: 4 ms (4000 us)
    Required Min RX Interval: 4 ms (4000 us)
    Required Min Echo Interval: 0 ms (0 us)
    
```

Figure 4.5 – A Wireshark capture of a spoofed BFD packet to remove protection from a target LSP.

We inserted packet sniffers at locations *c* and *d* of Figure 4.4 to observe the effect of the forged BFD packets. Few seconds after *R8* receives and processes the forged packets, we observe three PSC packets with a “Signal Fail (SF)” message at location *d* coming via the protection LSP towards *R1*. These three packets were sent by *R8* as a trigger for protection switching because it interprets the forged BFD packets as valid packets from *R1* informing it about the

Chapter 4. Security Vulnerabilities of the Cisco IOS Implementation of the MPLS Transport Profile

failure ("Down") state of the working LSP from *R1* to *R8*. In addition to the PSC packets, we also observe BFD control packets at location *c* coming via the working LSP with the *Sta* field set to "Down" (this observation confirms what was discussed in Section 4.2.1). After receiving the PSC and BFD packets from *R8*, we observe *R1* finalizes the packet switching by sending 3 PSC packets with a "Signal Fail (*SF*)" message towards *R8* via the protection switching. At this stage, the original protection LSP becomes the current working LSP.

To make the attack persistent, we kept injecting the forged packets on the original working path towards *R8*. These packets convince *R8* that the older working LSP is still down and cannot be relied up on as a backup LSP. This effectively removes protection from the current working LSP. Therefore, in case there is failure or quality degradation in the working LSP, the LERs don't have any backup LSP to switch to.

The reason why we were able to make the attack persistent is because *R8* did not have a means to detect the abnormal rate of BFD packets it received, the normal rate being 1 BFD packet per 4ms. Moreover, even though *R8* was receiving a mix of "Up" and "Down" BFD messages, it only reacted to the "Down" messages and did not raise any alarm.

Scenario II - Disabling LSP Fault Detection

The goal of this attack is to hide the presence of a link failure along the working LSP by letting the level edge routers receive all expected BFD messages for the LSP. In a way, this is the opposite of the attack we discussed in Section 4.4.2.

As shown in Figure 4.6, the attack requires physical access to insert a malicious Ethernet switches at two locations along the working LSP path; one at location *b* and another at location *d*.

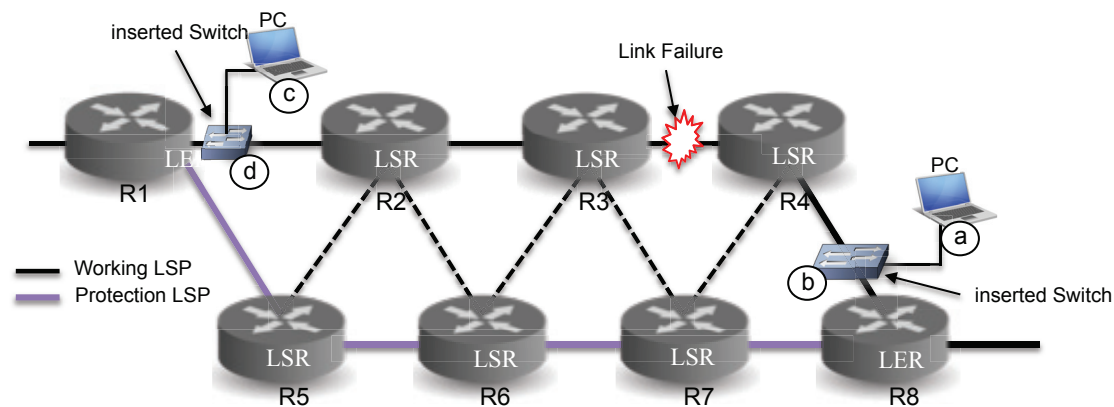


Figure 4.6 – Spoofing Attack BFD: LERs fail to detect a link failure along a working LSP.

After inserting the switches, connecting the laptops to the network via the switches and doing the necessary reconnaissance work, we create forged BFD packets with the *Sta* field set to

“Up” and simultaneously inject these packets to the network. The forged BFD packets injected by laptop *a* are destined to *R8* and appear to have come from *R1*, and those injected by laptop *c* are destined to *R1* and appear to have come from *R8*.

The next step of the attack is to physically sabotage one of the links in the working LSP path. In our experiment, we broke the one between *R3* and *R4*. Under normal conditions, this link failure would have stopped delivery of BFD messages for the working LSP to the LERs and would have triggered protection switching. By continuously injecting the forged BFD messages at the two locations, we fooled the two LERs to think that there is no link failure in the LSP. The side effect of this attack is that any application data that was being routed through the working LSP will be dropped at the broken link (technically, at the routers to which the broken link is incident). We verified this by sending ping packets from *R1* to *R8* and sniffing for traffic in the working LSP using laptop *c*. We were able to capture the ping requests but we could not see any ping replies coming from *R8*. Besides, we also look at *R1* for any replies that may have come through the protection LSP and there were none.

4.4.3 Spoofing Attack on PSC Messages

The goal of the attack is to cause a complete shut down of both the working and the protection LSP. As shown in Figure 4.7, an attacker needs to physically access the network at only one location. This attack is easier than the attacks on BFD messages because it takes only two forged PSC packets to get the desired result.




Figure 4.7 – Network setup for spoofing attacks on PSC control messages.

Chapter 4. Security Vulnerabilities of the Cisco IOS Implementation of the MPLS Transport Profile

Time	DeltaTime	Source	Destination	Protocol	Length	Info
0.000000	0.000000	Vmware_38:80:fa	Vmware_c7:be:aa	PSC	60	FS(1,1)
[+] Frame 1: 60 bytes on wire (480 bits), 60 bytes captured (480 bits) on interface 0						
[+] Ethernet II, Src: Vmware_38:80:fa (00:50:56:38:80:fa), Dst: Vmware_c7:be:aa (00:0c:29:c7:be:aa)						
[+] MultiProtocol Label Switching Header, Label: 800, Exp: 6, S: 0, TTL: 254						
[+] MultiProtocol Label Switching Header, Label: 13 (Generic Associated Channel Label (GAL)), Exp: 6, S: 1, TTL: 1						
[+] Generic Associated Channel Header						
[+] PSC						
01.. = Version: 1						
..11 00.. = Request: Forced Switch (12)						
.... ..10 = Protection type: bidirectional switching using a selector bridge (2)						
1... = R: revertive mode (1)						
Fault Path: working (1)						
Data Path: protection is in use (1)						
TLV Length: 0						

Figure 4.8 – A wireshark capture of a spoofed PSC shutdown command to a working LSP.

To carry out the attack, we insert an Ethernet switch at location *a* in Figure 4.7 and connect our laptop to the network via the switch. We sniffed the traffic and learnt about the MPLS label and the MAC address of *R6* and *R7*. We use this information to build two forged PSC packets that emulate a “shutdown” command from an operator.

To accomplish this, we forged two PSC packets one for the working LSP and the other for the protection LSP by manipulating the *Request* and the *Fault Path* fields. To disable the working LSP, the *Request* field was set to “Forced switch (12)” and the *Fault Path* field to “Working (1)” (Figure 4.8). For the protection LSP, the *Request* field is set to “Lockout of protection (14)” and the *Fault Path* field to “Protection (0)” (Figure 4.9). The two forged packets are then sent to *R8* through *R7*.

By looking at incoming and outgoing packets through the two interfaces of *R8* (location *c*), we see that the router sent 3 PSC packets towards *R1* with “Forced switch” in the *Request* field and “Working” in the *Fault Path* field. These three packets were sent out to lock out the working LSP and trigger protection switching. This is because on receiving the forged PSC packets targeting the working LSP, it assumed it was issued by a network operator at *R1* to force protection switching. We also observed other three PSC packets from *R8* to *R1* with “Lockout of protection (14)” in the *Request* field and “Protection” in the *Fault Path* field. These three packets were sent by *R8* to inform *R1* that it was locking out the protection LSP in response to the forged PSC packets that targeted the protection LSP. This is because the forged packets were understood by *R8* as network operator commands issued at *R1* to lockout the protection LSP.

The final result was both the working and LSPs locked out and a protection switching triggered but could not materialize. Hence, there was a complete shutdown of both LSPs. Since the lockdown of the LSPs was assumed to be caused by network operator issued commands, another network operator command needs to be explicitly issued to bring the two LSPs up to work.

Time	DeltaTime	Source	Destination	Protocol	Length	Info
0.0000000000000000		Vmware_38:80:fa	Vmware_c7:be:aa	PSC	60	LO(0,0)
[Frame 1: 60 bytes on wire (480 bits), 60 bytes captured (480 bits) on interface 0]						
[Ethernet II, Src: Vmware_38:80:fa (00:50:56:38:80:fa), Dst: Vmware_c7:be:aa (00:0c:29:c7:be:aa)]						
[MultiProtocol Label Switching Header, Label: 800, Exp: 6, S: 0, TTL: 254]						
[MultiProtocol Label Switching Header, Label: 13 (Generic Associated Channel Label (GAL)), Exp: 6, S: 1, TTL: 1]						
[Generic Associated Channel Header]						
.... 0000 = Channel version: 0						
Channel Type: Protection State Coordination Protocol (PSC) (0x0024)						
[PSC]						
01.. = Version: 1						
..11 10.. = Request: Lockout of protection (14)						
.... ..10 = Protection type: bidirectional switching using a selector bridge (2)						
1... = R: revertive mode (1)						
Fault Path: protection (0)						
Data Path: protection is not in use (0)						
TLV Length: 0						

Figure 4.9 – A wireshark capture of a spoofed PSC shutdown command to a protection LSP.

4.5 Discussion and Countermeasures

All the three attacks we demonstrated in Section 4.4 were caused by forged BFD and PSC messages injected by a malicious device that managed to masquerade as an authentic source of the messages. None of the attacks would be possible if the message receivers (the label-edge routers) had a means to verify the authenticity of the messages. Therefore, the attacks could be avoided by implementing a proper message authentication mechanism that enables a receiver to verify the identity of the message source and that the message was not tampered with while in transit.

RFC5880 [63] proposes an optional authentication field BFD control packets that can be used to carry the necessary information to allow the receiving LER to determine the validity of the received packet. However, the RFC does not suggest any specific authentication scheme to use nor does it say anything on how the keys required for the authentication are to be exchanged. Our experiment reveals that the Cisco IOS MPLS-TP does not provide any form of message authentication.

In our literature review, we could not find any RFC or standard that discusses authentication support for PSC packets. One possible solution is to propose a built-in security that uses one of the optional TLV fields as an authentication field, similar to the one proposed for BFD control packets in RFC5880.

The other means to provide authentication to BFD and PSC messages is to add an authentication layer that is not built-in to the packet format. One such solution is to use IPsec. RFC5085 [59] suggest that IPsec be used to protect packets in MPLS/GMPLS networks using MPLS-in-IP encapsulation. IPsec provides end-to-end security (including authentication) for IP packets. IPsec can be configured either in Tunnel mode or in Transport mode. In the tunnel mode, which is the default mode, the entire IP packet is protected (encrypted and authenticated) by IPsec. Therefore, if IP is supported in the an MPLS-TP core networks, IPsec can be used to secure BFD and PSC packets. Other standard solutions such as D(TLS) can also provide the required security. However, there are several cases in the smart grid context where an MPLS-TP core network does not necessarily support IP. For example, MPLS-TP networks

Chapter 4. Security Vulnerabilities of the Cisco IOS Implementation of the MPLS Transport Profile

that transport the IEC 61850 Generic Object Oriented Substation Event (GOOSE) and Sampled Measured Values (SMV) messages in substation automation systems don't usually use IP [67].

For cases where MPLS-TP does not support IP, a hop-by-hop security (e.g., MACsec) is an alternative solution. MACsec requires every pair of neighbouring nodes in an LSP has a shared secret key to be used for message authentication at the MAC layer. MACsec is not supported in today's Cisco routers. One key issue with MACsec is that it suffers from a single point of failure, i.e., if an attacker manages to compromise one of the MACsec sessions between any two neighbouring routers along an LSP, he will be able to launch the attacks discussed in the previous section. Therefore, it is important that routers save the cryptographic materials for MACsec in a secure location. Using a tamper resistant hardware module such as Trusted Platform Module (TPM) is one solution. Besides, proper access control mechanisms to prevent unauthorized access to sensitive data within the device should be put into place. Note that in addition to preventing spoofing attacks, MACsec also prevents DoS attacks by making sure that spoofed packets are detected at the end of the link where the packets are injected to the network. Therefore, among the different available options, we believe MACsec is the preferred solution to secure BFD and PSC packets in MPLS-TP networks.

4.6 Conclusion

MPLS-TP is envisioned to be a promising Packet Switched Network technology for smart grid networks. Our literature review of the different standards and RFC related to MPLS-TP showed that security is only marginally addressed for MPLS-TP's OAM protocols. Our experimental study of MPLS-TP implementation in Cisco IOS confirms that more needs to be done to secure MPLS-TP's OAM protocols. Our findings show that lack of support for message origin authentication in BFD and PSC protocols leads to an attacker being able to remove protection support, hide the presence of link failure and to completely shutdown both working and protection LSPs. All these attacks are of grave concern for a smart grid network that uses MPLS-TP to transport critical data used for grid monitoring, protection and control. Therefore, it is important the RFC's be more directive in suggesting built-in authentication and other security solutions for MPLS-TP protocols. In the case when built-in security is not possible, other standard security solutions that provide end-to-end security like (D)TLS or those that provide a hop-by-hop security, like MACsec, should be used.

5 Experimental Comparison of Multicast Authentication for Grid Monitoring Systems

5.1 Introduction

The smart grid, a superimposition of cyber infrastructure on a physical power system infrastructure, is envisioned to provide a reliable and efficient power supply with a smooth integration of renewables. Smart grid is a generic term that comprises different systems. Advanced metering systems, demand response management systems, substation automation systems, and wide area monitoring systems (WAMS) are a few of several systems that define a smart grid.

The cyber infrastructure in a smart grid facilitates two-way communication of sensing (metering) data and control signals among field devices and control centres. The field devices and the communication infrastructure usually span a large unprotected geographic area. One challenge for such a system is protecting against cyber attacks; in particular, guaranteeing message source authenticity to data consumers is difficult. Different systems in a smart grid use different communication paradigms, have different real-time requirements and the devices they use have different levels of resource constraints. Hence, there is no a one-size-fits-all security solution that works for all systems.

In this paper, we focus on identifying the best source authentication scheme for phasor data communication in wide area monitoring systems (WAMS). WAMS use high-resolution phasor data from several phasor measurement units (PMUs) to provide real-time information about a power grid's state and can be used to trigger corrective actions to maintain reliability. The North American Synchronphasor Initiative (NASPI) was founded to facilitate the deployment and use of synchronphasor technology for grid reliability and efficiency [68]. Although a few years ago there were only a few hundred PMUs deployed across the North American power grid and elsewhere, their adoption has exponentially increased due to their perceived benefits and their reduced cost [14]. PMUs provide absolute time-synchronized data at a high data rate compared to supervisory control and data acquisition (SCADA) systems, typically from 30 to 60 samples/second. This enables power grid operators to have real-time situational awareness of their grid, which in turn enables them to implement fast response to unstable conditions

Chapter 5. Experimental Comparison of Multicast Authentication for Grid Monitoring Systems

observed in the grid. Depending on the nature of the different control applications that use WAMS [14], the overall delay budget for synchrophasor data ranges from 4 to 20 ms [69]. Most of this budget is consumed by the communication and computation excluding security related operations. Therefore, the additional delay due to security is preferred to be in the order of sub-milliseconds.

IP multicast is envisioned to be a preferred communication paradigm for PMU to Phasor Data Concentrator (PDC) streaming [70, 71] because it is efficient for one-to-many communication in that it relieves a PMU from sending several copies of the same packet destined to multiple PDCs. Besides, since a multicast group address is used as a destination, new receivers can be added to an already operational WAMS seamlessly without any setting changes to other PMUs or other PDCs in the group. Multiple receivers are used for different reasons. A common reason is to support redundant PDCs for reliability. Another common reason is to have the SCADA system and archive servers receive the PMU data for supervision, fault detection and post mortem analysis. In some cases, a utility shares synchrophasor data with other neighboring organizations so that all utilities have a common understanding of the state of the entire grid, which allows them to better respond to detected conditions across the grid.

In spite of its benefits, multicast also comes with its own security challenges. More specifically, designing a multicast source authentication scheme for time-critical systems such as WAMS is a challenging problem [72, 73]. As a result, this problem is extensively studied by the research community [74]. Guaranteeing source authentication (thereby message integrity) is crucial for WAMS because any tampering of the synchrophasor data while in transit or injection of bogus data by an attacker leads to wrong real-time situational awareness of the grid; which in turn can lead to issuing wrong corrective measures with catastrophic consequences.

A trivial approach to providing multicast source authentication is to use a shared key (group key) scheme that uses message authentication codes (MACs). Several studies have proposed this as fast authentication mechanism for different smart grid applications [75, 76]. Although such a scheme is computationally fast and provides group authentication, it does not give any protection against an untrusted receiver since such a receiver can impersonate the source using the shared key. Group authentication can be considered sufficient for homogenous substation automation systems as we can assume that if one of the receivers in such a system is compromised, other receivers are also likely to be compromised. In this paper we consider WAMS, which are heterogeneous systems compared to substation automation systems. In WAMS, receivers are not necessarily colocated and may not have the same level of security. Therefore, group key based authentication is not viable in our framework because an attacker needs to compromise only one receiver or source to compromise the whole network.

An efficient multicast authentication requires a source of asymmetry in the authentication information. In other words, receivers should be able to verify the authentication information, but should not be able to generate valid authentication information [73, 74]. Different schemes use different sources of asymmetry. Some schemes use as a source of asymmetry the difference

in the number of symmetric keys that sources and receivers know [77]; others use time [78] and yet others use the computational intractability of the cryptographic primitives used to generate the keys (e.g., one-wayness of a function, collusion resistance of hash functions, factoring difficulty, discrete log problem) [79–81].

In this paper, we evaluate the different multicast authentication schemes that use asymmetry in the authentication information and identify the best candidate for WAMS. The set of metrics we use to evaluate the performance of these schemes are computation overhead, communication overhead and key management (key generation, distribution and storage) overhead. From the literature review, the short-list we identify for further evaluation are two variants of elliptic curve digital signature algorithm (ECDSA) [80], “time valid hash to obtain random subsets” (TV-HORS) [79] and three variants of Incomplete-key-set [77]. An experimental comparison of the short-list is then made in an operational wide area monitoring system that deploys the National Instruments CompactRIO 9068 based PMUs and phasor data concentrators (PDCs) to monitor a medium-voltage distribution network on the EPFL campus [82]. To the best of our knowledge, we are the first to perform an experimental comparison of different authentication schemes using actual PMUs deployed on an operational WAMS.

From our experiment, we find that even though the Incomplete-key-set variants use only symmetric key operations, their high computation and communication overheads make them impractical for WAMS based real-time applications. The ECDSA with no pre-computed tokens has low communication and key management overheads; however it has high computation overhead due to a slow key generation at resource-constrained PMUs. Therefore, for all practical purposes it requires hardware support in PMUs. The ECDSA variant which uses pre-computed tokens for fast signature generation has small computation and communication overheads which make it an ideal candidate for WAMS. TV-HORS also has low computation and communication overheads, but it has a large key management overhead as it requires frequent distribution of a large public key that needs to be reliably delivered to each receiver within a specified time window.

The rest of the paper is organized as follows. In Section 5.2 we present the state of the art. In Section 5.3, we provide an in-depth discussion of the short-listed schemes. We describe the wide area monitoring system for the EPFL active distribution network which we use as our testbed for the experimental comparison of the schemes in Section 5.4. We present the experimental results and comparison of the schemes in Section 6.6. Finally, in Section 6.7 we conclude the paper.

5.2 Authentication Mechanisms for IP Multicast

In this section, we cover the state of the art for multicast authentication. We also identify which source authentication schemes are more feasible for phasor data communication in wide area monitoring systems (WAMS).

Chapter 5. Experimental Comparison of Multicast Authentication for Grid Monitoring Systems

5.2.1 Asymmetric cryptography based schemes

Authentication schemes in this category include all schemes that are based on digital signatures, such as RSA and ECDSA [83]. Sources use their private keys to sign messages and receivers use the source's public key to verify received message source authenticity. These schemes are scalable in that they require a single small-size public/private key pair for every multicast source. However, directly applying these schemes for most real-time (e.g., smart grid) applications is a challenge because of their expensive computation overhead. The IEC standardization body in its IEC 62351-6 [18] standard suggests that RSA be used to authenticate IEC 61850 Generic Object Oriented Substation Event (GOOSE) / Sampled Measured Values (SMV) messages that have a 4ms response time. However, resource constrained intelligent electronic devices (IEDs) in substations are generally incapable of computing and verifying a digital signature using the RSA algorithm within the required response time. Yavuz in [84] proposed a fast RSA based scheme by exploiting an existing structure in command and control messages. Such a scheme, though efficient, is not applicable for WAMS because the structure assumed in [84] is not present in PMU measurements. Hohlbaum et al. [19] show that, with today's IED's hardware, the software implementation of digital signatures would not meet the real-time requirements of GOOSE/SMV messages. They also show the FPGA implementation of RSA signature with a key length of 1024 bits is not feasible for systems that have less than 4ms response time requirement. However, an RSA implementation on hardware like ASIC platforms and specialized crypto-chips are shown to be feasible solutions.

The cost of specialized hardware are expected to be affordable in the future that we can imagine digital signature solutions be preferred solutions in future smart grid devices. Therefore, we consider digital signature based solutions as one of the candidates for multicast authentication. More specifically, we choose ECDSA as the preferred candidate among digital signature schemes to be included in the short-list, as it has a shorter public/private key length and signature size compared to RSA for a similar security level.

5.2.2 One-time signature (OTS) schemes

One-time signature were first proposed by Lamport [85] and by Rabin [86]. Subsequent works on OTS [79, 87–89] improved the signature length and computation overhead required for signing and verification. Law et al. in [90] provide a simulation-validated mathematical analysis of the different OTS schemes and identify TV-HORS [79] as the favourable authentication scheme for real-time applications in terms of providing a balanced computation and communication efficiencies relative to security level. In a different context from WAMS, Lu et. al in [91] compare by simulation TV-HORS with RSA when applied for multicast authentication in substation automation systems. Their results show that TV-HORS performs better than RSA, in terms of computation cost. From our literature review and from works that did theoretical and simulated comparison of OTS systems, TV-HORS is shown to be the preferred scheme among OTS schemes. Therefore, TV-HORS is included in our short-list of candidate schemes

for further evaluation.

5.2.3 Message authentication code (MAC) based schemes

MAC based schemes use a shared symmetric key between a sender and a receiver to generate a cryptographically secure authentication tag for a given message. The simplest scheme in this category uses a group key shared among the multicast source and all the receivers. For example, a multicast extension to IPsec (RFC 5374) uses group keys to provide message authenticity and confidentiality. Secure distribution of the key to the multicast group members is handled by the group domain of interpretation protocol (GDOI, RFC 6407). The IEC 61850-90-5 [75] standard specifies the multicast extension of IPsec to secure synchrophasor data. Zhang and Gunter [76] also propose using IPsec for securing multicast data in substation automation and show the stringent latency constraints (less than 4ms) can be satisfied with their solution. The problem with all group key based solutions is they do not provide protection against a malicious receiver, i.e., any receiver that has the shared key can impersonate a legitimate source.

Another variant of the symmetric key based solution uses a secret-information asymmetry to cope with the impersonation problem stated above. Canetti et al. [77] propose such a scalable scheme suitable for systems with a large number of multicast receivers. In this scheme, the source knows a set of secret keys to authenticate a multicast message and each receiver knows only a subset of these keys that enable it only to verify the authenticity of received messages without being able to generate valid authentication information for messages [74]. The source attaches MACs computed using all its keys to the messages and each receiver uses its subset of keys to verify the authenticity of the received message. We refer to this scheme as the *Incomplete-key-set* scheme [91].

As the *Incomplete-key-set* scheme uses only fast MAC computations and does not require buffering before authentication, we include this scheme in the short-list of candidate schemes for further evaluation. In Section 5.3.3, we provide a more detailed description of the scheme.

5.2.4 Delayed key disclosure schemes

Like the schemes in 5.2.3, schemes in this category use a keyed-hash message authentication code (HMAC) for source and message authentication. The main difference between the two categories is the source of asymmetry, i.e. delayed key disclosure based schemes use time as a source of asymmetry. The source computes the HMAC of a message by using a symmetric key that only it knows. The receiver buffers the message until it receives the authentication key from the source. The source then discloses the key in its subsequent messages. Timed efficient stream loss-tolerant authentication (TESLA) [78] and its variants [92, 93] are examples of this scheme. To minimize the effect of packet losses, TESLA employs a chain of authentication keys linked to each other by a pseudo random function. Each key in the key chain is the image

Chapter 5. Experimental Comparison of Multicast Authentication for Grid Monitoring Systems

of the next key under the pseudo random function.

Delayed key disclosure schemes have low computation overhead (only one MAC function) and low communication overhead. The drawback with these schemes is they need to buffer messages, which makes them inapplicable for real-time smart grid applications like WAMS. Thus, we do not include schemes from this category in our short-list.

5.2.5 Signature amortization schemes

Signature amortization refers to using a single signature for authenticating a group of multicast packets, thereby spreading (amortizing) the signature verification cost across this group of packets [94]. A receiver has to assemble all the packets in the group before verifying their collective signature. As the introduced delay due to buffering makes them inapplicable for real-time applications, we do not consider schemes in this category for further evaluations.

Table 5.1 provides a summary of the different authentication schemes with respect to some desirable properties for WAMS. We have selected these desirable properties that are applicable for WAMS from those identified in [73] and [84]. A perfect scheme would be one that performs well in all the identified properties. As can be seen from the table none of the schemes satisfy that requirement. The subset of schemes we have chosen for further evaluation are those that satisfy the first three properties.

Table 5.1 – Summary of different multicast authentication schemes with respect to different desirable properties for WAMS.

	PKC		OTS	MAC based		Delayed disclosure	Amortized
	RSA	ECDSA	TV-HORS	Group key	IKS	TESLA	RSA based
Immediate authentication (no buffering)	Yes	Yes	Yes	Yes	Yes	No	No
Provides asymmetry	Yes	Yes	Yes	No	Yes	Yes	Yes
Robust to data packet loss	Yes	Yes	Yes	Yes	Yes	Partial	Partial
Scalable for large systems	Yes	Yes	Moderate	Yes	No	Yes	Yes
Free from time-bounded security	Yes	Yes	No	Yes	Yes	No	Yes
Low computation overhead	No	No	Yes	Yes	No	Yes	No
Low communication overhead	Yes	Yes	Yes	Yes	No	Yes	Yes
Low key storage at source	Yes	Yes	No	Yes	No	Moderate	Yes
Low key storage at receiver	Yes	Yes	No	Yes	No	Yes	Yes

IKS: Independent-key-set; PKC: Public key cryptography

5.3 Candidate multicast authentication schemes for wide area monitoring systems

In this section, we give a description of the three multicast authentication schemes that we identified in Section 5.2 as candidates for wide area monitoring systems.

5.3.1 Elliptic Curve Digital Signature Algorithm (ECDSA)

The elliptic curve digital signature algorithm (ECDSA) is a public-key authentication scheme whose security is based on the computational intractability of the Elliptic Curve Discrete Logarithm Problem (ECDLP) [80]. ECDSA provides the same level of security as other digital signatures, such as RSA, but with a smaller key size. Smaller keys enable ECDSA to have a faster computation time. For this reason, ECDSA is the digital signature scheme of choice for new applications: for example, Bitcoin relies on ECDSA for its security.

Below, we provide a brief description of the steps required to set up an ECDSA based multicast authentication system. More specifically we describe the domain parameter setup, key pair generation, signature generation and signature verification.

Domain parameters setup

The public/private key pairs used by ECDSA are generated with respect to a particular set of domain parameters (p, a, b, G, n) , where p is the prime modulus, a and b are coefficients of the elliptic curve, G is a group generator of prime order n . For better security, the elliptic curve should be chosen from a small set of elliptic curves referenced as NIST Recommended Elliptic Curves in FIPS publication 186 [83].

Key pair generation

Once the domain parameters are chosen, public/private key pair is generated as follows:

- (a) Private key is a random integer $d \in [1, n - 1]$.
- (b) Public key $Q = dG$ is a point on the elliptic curve.

Signature generation

Given a hash function h and a sender's key pair (d, Q) , a message m is signed as follows:

- (a) Select random $k \in [1, n - 1]$.
- (b) Compute $(x_1, y_1) = kG$.
- (c) $r = x_1 \bmod n$. If $r = 0$, go back to step *a*.
- (d) Compute $s = k^{-1}(h(m) + rd) \bmod n$.
If $s = 0$, go back to step *a*.
- (e) The signature for message m is the pair (s, r) .

Chapter 5. Experimental Comparison of Multicast Authentication for Grid Monitoring Systems

Signature verification

Given a sender's public key Q , the authenticity of a received message m is verified as follows:

- (a) Compute $(x_2, y_2) = s^{-1}(h(m)G + rQ)$.
- (b) Verification succeeds if $x_2 \equiv r \pmod{n}$ and $r, s \in [1, n - 1]$.

An interesting feature of ECDSA is that signature generation is faster than signature verification. This is a desirable feature for applications like WAMS because message sources (PMUs) are more resource constrained than message receivers (PDCs). Even with such asymmetry, signature generation is still expensive. A typical approach to achieve fast signature generation is to pre-compute r and k 's modular inverse k^{-1} before the message is known [95]. By pre-computing \aleph of these tokens offline, we later use them to sign \aleph messages as they appear at a minimum cost. In this paper, we evaluate the performance of ECDSA signature generation with and without pre-computed tokens.

5.3.2 Time Valid Hash to Obtain Random Subsets (TV-HORS)

TV-HORS [79] is an extension of hash to obtain random subsets (HORS) [88] authentication scheme. TV-HORS inherits HORS's advantages of fast message signing and verification. TV-HORS achieves small signature size and faster computational efficiency by signing only part of the hash of the message and by using a time-bounded signatures to prevent signature forgery. The signature period (a.k.a., epoch) is the maximum possible duration a signature can be exposed before it is verified. This duration has to be short enough so that an attacker cannot get a partial-hash collision of the signed message within that time duration.

One drawback of TV-HORS is the need for a periodic exchange of a large public key. TV-HORS uses two approaches to decrease the public key refresh rate: (1) It reuses its private key to sign multiple messages within a given epoch, i.e., it functions as a multiple-time instead of a one-time signature scheme. (2) It uses multiple key pairs linked together by using one-way hash chains, as shown in Figure 5.1, to authenticate a large number of streaming packets without needing to redistribute a new public key at the end of every epoch.

Though the "multiple timed-ness" feature improves the public key refresh rate, it also has security ramifications. It exposes more elements in the private key with every signed message. Thus, it provides an attacker with more opportunities to forge a message using the released private key elements.

The security level L for TV-HORS is expressed as a function of three parameters: the maximum number of messages that can be signed by a private key within an epoch ν , the number of elements in a private key N and the number of elements in a signature t . As shown in [79], $L = t \log_2(N/\nu t)$. The security level L is a security parameter such that an adversary has to compute 2^L hash computations on average to obtain a valid signature for a new message.

5.3. Candidate multicast authentication schemes for wide area monitoring systems

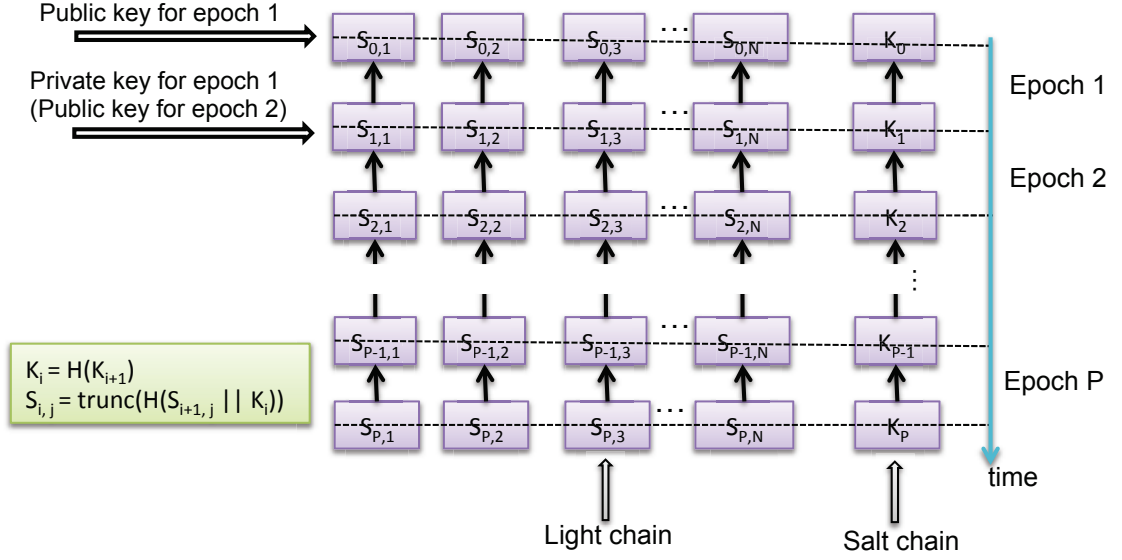


Figure 5.1 – TV-HORS key pairs linked using one-way hash chains. At epoch j , the light chain $s(j, _)$ and the salt k_j form the active private key. This private key can sign upto ν message within that epoch. A session has a total of P epochs. A key chain is refreshed at the end of epoch P .

Hence the TV-HORS parameters N, t, ν should be chosen such that the above formula satisfies a required security level L .

5.3.3 Incomplete-key-set

The basic idea behind the Incomplete-key-set scheme is the sender appends to each multicast message multiple MACs computed by using different symmetric keys. The asymmetry between senders and receivers is provided by the fact that the source knows more secret keys than each receiver.

Below we present three variants of this scheme that apply for two different scenarios.

Incomplete-key-set for a small number of receivers per group

In WAMS where the number of receivers is small (in the order of tens), implementing a variant that we refer to as *perfectly-secure Incomplete-key-set* is sufficient. For a multicast group of R receivers and any number of sources, this scheme uses a total of R primary secret keys $\kappa = \{k_1, \dots, k_R\}$ from which R secondary secret keys $\kappa_s = \{f(s, k_1), \dots, f(s, k_R)\}$ are generated and assigned to each source s , where $f(\cdot)$ is a pseudo-random function. Each receiver r is assigned a distinct primary key k_r from the set κ . The source authenticates a message m by computing R MACs using its R secondary secrets and concatenates all the MACs with the message. Each receiver r computes the secondary key of s that corresponds to its primary key k_r and verifies the authenticity of the message by verifying the MAC that was computed using this secondary key. However, it is not a scalable solution since the communication overhead (size of the

Chapter 5. Experimental Comparison of Multicast Authentication for Grid Monitoring Systems

MACs) grows linearly with the number of receivers.

Incomplete-key-set for a large number of receivers per group

In a system where there are a large number of multicast receivers, Canetti et al. [77] proposed a scheme that we will refer to as the *basic Incomplete-key-set* scheme. This addresses the scalability issue associated with the variant introduced above. This scheme uses a set of $l < R$ primary keys $\kappa = \{k_1, \dots, k_l\}$ from which a set of l secondary keys $\kappa_s = \{f(s, k_1), \dots, f(s, k_l)\}$ are assigned to each multicast source s . Each receiver r is assigned a set κ_r of primary keys such that $\kappa_r \subset \kappa$. When sender s wants to multicast message m , it computes l MACs using the secondary keys in κ_s and sends the message m , along with the l MACs. On receiving a message from sender s , receiver r computes the secondary keys of s with the primary keys in κ_r . It then verifies all the MACs that were computed using these secondary keys. If any of these MACs is incorrect, then r rejects the message.

The basic Incomplete-key-set scheme is susceptible to collusion attacks. A group of fraudulent receivers can collude among each other such that for each receiver j in the fraudulent group, $\cup \kappa_j$ can completely cover the key subset κ_u of a given receiver u with a certain probability.




Figure 5.2 – Key distribution for the Incomplete-key-set authentication scheme.

Let a multicast group have a maximum number of w corrupt users and let q be the probability that κ_u for any receiver u is completely covered by the subsets held by the coalition members. The authors in [77] show that the number of primary keys l is given by $l = e(w + 1) \ln(1/q)$. Each receiver r obtains a subset κ_r of primary keys such that $|\kappa_r| = e \ln(1/q)$.

Depending on the values of the system parameter w and q , the number of keys l can be large

thus the communication overhead can be large. The authors in [77] propose a *communication-efficient* variant of the basic scheme that uses MACs with a single bit as output so that the authentication information is reduced to only l bits. For such a setting, the number of MAC computations are four times that of the basic scheme, i.e., the total number of primary keys l and $|\kappa_r|$ are four times that of the basic scheme.

5.4 System setup and evaluation methodology

In this section, we describe the active power distribution network that we used as a testbed to perform our experiment to compare the three multicast authentication schemes introduced in the previous section. We also introduce the performance metrics we use to evaluate the schemes.

5.4.1 EPFL-Campus Smart Grid Monitoring System

We carry out the experimental comparison of the authentication schemes on the smart grid infrastructure deployed at EPFL to monitor the power distribution network of the campus.




Figure 5.3 – The EPFL smart grid infrastructure with 7 PMUs as multicast sources and 3 PDCs as multicast receivers.

Figure 5.3 depicts the map of the EPFL campus smart grid infrastructure. The smart grid infrastructure deploys PMUs at different locations on the campus. The PMUs measure synchrophasor data at the different locations at a rate of 50 samples/second, encapsulate the data according to the IEEE C37.118.2-2011 standard [96] and multicast it over UDP to aggregation points called phasor data concentrators (PDCs). Each synchrophasor measurement from a PMU is 74 bytes long. A PDC time-aligns the measurements from the different PMUs and feeds

Chapter 5. Experimental Comparison of Multicast Authentication for Grid Monitoring Systems

the time-aligned synchrophasor data to a real-time state estimator that is co-located with each PDC. The output of the real-time state estimator enables us to determine the most likely state of the grid. Our monitoring infrastructure of the smart grid pilot on the EPFL campus has a total of 7 PMUs and 3 PDCs in total. A more complete description of the smart grid infrastructure can be found in [82].

With no security (authentication or encryption) deployed, the overall latency between the time the synchrophasor data is sent from the PMU to the time the state estimator output is computed has a mean value of 17 ms. This relatively low latency in computing the state of the grid enables us to have a real-time grid monitoring system, which in turn enables us to implement real-time corrective measures when the state estimator output indicates a deviation from the grid's stable state. Any tampering of the synchrophasor data by an attacker, while in transit from the PMUs to the PDCs, leads to a wrong state estimator output; which in turn can lead to issuing wrong corrective measures with catastrophic consequences - thus the need for message authentication.

5.4.2 Comparison Metrics

The set of metrics we use to compare the performance of the multicast authentication schemes are *computation overhead* per message, *communication overhead* per message and *key management overhead*. Computation overhead refers to the processing time required to generate an authentication code (signature) at the sender and to verify the authenticity of the message at the receiver. Some of the schemes we evaluate have asymmetric computation overhead for authentication and verification. An authentication scheme is considered efficient for a real-time application if the sum of the authentication and the verification time is small. Communication overhead as a metric refers to the length of the authentication data that a scheme generates per message. This metric is important especially in systems where the network bandwidth is a constraint. The third metric, key management overhead, is the cost associated with the generation, distribution and storage of the key material. The key generation overhead is the CPU time required by a PMU to generate the keys. The distribution overhead is the bandwidth required to distribute the key material to the communicating partners. The storage overhead is the amount of memory required to store the key materials.

An ideal authentication scheme for WAMS is one that has low overhead in all the metrics. However, finding a scheme that satisfies all such requirements is difficult. WAMS are real-time applications. Thus, a small computation overhead is considered a critical requirement. In contrast, utilities are likely to have dedicated state-of-the-art communication infrastructure for their synchrophasor data communication. Therefore, low communication overhead can be considered a soft requirement. The key management overhead, however, is a combination of both computation and communication overheads. Thus, a low key management overhead is also a critical requirement.

It is important to mention here that the three schemes are immune to packet losses if the

packets contain application data (not key materials). For these reasons, we don't make any comparison among the schemes based on resistance to loss of packets containing application data. In contrast, packet losses during key distribution may affect the performance of a scheme and is discussed in Section 5.5.2.

5.5 Performance evaluation and comparisons

5.5.1 Implementation and Parameter Settings

The multicast sources at the EPFL smart grid pilot are National Instrument's CompactRIO 9068 based PMUs with a 667 MHz dual-core ARM Cortex-A9 processor, 512 MB DDR3 memory and 1 GB nonvolatile storage running NI Linux Real-Time OS. Likewise, each receiver is a PC with an Intel 2.8 GHz Core *i7* processor and a 4GB RAM running Ubuntu 12.04 with Linux 3.2. The source and receiver are implemented in *C* and use OpenSSL [97] open source tool kit to implement the authentication schemes. We use SHA-256 whenever we need a hash output for any of the schemes.

Threat model

The attacker is assumed to have an indepth knowledge of the the power system model so that he can launch an attack similar to the one proposed by Liu et al in [28] by corrupting measurement data from a selected set of PMUs to stealthily introduce arbitrary errors in the state estimator's output of certain state variables without triggering an alarm from a bad data detection algorithm. The first ever cyber-attack on three Ukrainian regional electric power distribution companies that caused a widespread power-outage in Ukraine on December 23, 2015 demonstrates the practical feasibility of mounting such an attack successfully [98]. Moreover, we assume that an attacker has continuous remote or physical access to the communication network of the WAMS from which he can intercept and capture measurement data from the selected PMUs. We also assume the attacker has access to a cloud computing resource that is equivalent to the computing capacity of a few thousand PCs. The attacker uses the computing resources to recover the secret (private) keys used to authenticate the synchrophasor messages in real time and uses them to authenticate forged messages and send them to the receivers as if they were sent from the legitimate PMUs whose keys are compromised. Since the PMUs refresh their keys periodically, the attacker can use a compromised key only until it is refreshed. Hence, the attacker needs to continuously follow the key refresh by the PMUs and re-do the key retrieval from captured messages after every refresh.

Security level and key refresh rate

The different authentication schemes have different parameters whose values affect the schemes' performance and security level. In order to make a fair comparison of the schemes,

Chapter 5. Experimental Comparison of Multicast Authentication for Grid Monitoring Systems

we set their parameters so that they all have equivalent security levels. According to [99], an ECDSA in a subgroup of m -bit size has an equivalent security level with a symmetric key based scheme of $m/2$ bits key-length. The security level of a symmetric key-based scheme is equal to the key length. As stated in Section 5.3.2, the security level for TV-HORS is defined by $L=t \log_2(N/vt)$.

Message authentication in WAMS is a short-term issue, i.e., it is enough to guarantee that the signing key is hard to break between the signing time and the signature delivering time [100]. Therefore, in our implementation, we use short-term keys by putting a bound on the life time of these keys.

As shown in [79], it takes 16×10^3 workstations to break TV-HORS with $L=54$ in 6 days. Eberle et al. in [101] show it takes 3.01×10^7 machines equipped with ECC-processor to work together for about 24 hours to break an 112-bit ECC key ($L=56$) and 1.02×10^{15} machines to break a 160-bit ECC key ($L=80$). In our experiment we considered two security levels: an intermediate security level $L=56$ and a stronger, future proof security level $L=80$. Based on the above data, we believe that a security level of $L=56$ is strong enough in the presence of an attacker with a computing capacity stated above if the keys are refreshed with in a few tens of seconds or even minutes. We have considered $L=80$, to see how the schemes compare when an attacker is likely to have more powerful computing capability in the future as cloud computing resources become more affordable.

For the intermediate security level, we generate the ECDSA key pairs from the elliptic curve domain *secp112r2* - a SECG curve over a 112-bit prime field. ECDSA keys generated from this curve have a security level $L=56$. For the Incomplete-key-set variants, we use a symmetric key-length of 56-bits. We set the TV-HORS parameters ($N=1024$, $t=13$, $v=4$), which give us $L=56$. For the stronger security level $L=80$, we use the 160-bit elliptic curve *secp160r2* for ECDSA, a symmetric key-length of 80-bits for the Incomplete-key-set and the parameters ($N=1024$, $t=16$, $v=2$) for TV-HORS. From the contour lines in Figure 5.4, we see that there are a range of values for v and t for a fixed value N to achieve a required security level L . A contour line in the $v-t$ plane show all the possible (v, t) pairs (only integer pairs) that give a value on the L axis that has the same color as the contour line. We took two representative set of values for t and v (one for $L=56$ and another for $L=80$) to conduct our experiment.

For all the schemes, we use a session duration $T_s=20$ sec. The message sending rate of the PMUs in our WAMS is $\lambda=50$ msg/sec, where each message is 74 bytes long synchrophasor data. Therefore, the PMUs stream 1000 messages during one session. We assume the key material for the entire session for all the schemes are pre-generated. For TV-HORS, the key-chain length (number of epochs P) is given by $P=T_s * \lambda/v$. Therefore, for the case where $L=56$, the number of epochs $P=250$ and for the case $L=80$, $P=500$. Note that a larger P value means a larger key generation and storage overhead. It also means the average verification time increases at the PDC.

The public keys for ECDSA and for TV-HORS and the symmetric keys for the Incomplete-




Figure 5.4 – TV-HORS security level (L) as a function of v and t for a fixed $N=1024$.

key-set that are used during session i are pre-generated and distributed during session $i - 1$. Similarly, for the ECDSA with pre-computed tokens, all the tokens required for the entire session i are locally pre-computed by each PMU during session $i - 1$. The public keys for TV-HORS and ECDSA are multicast to all receivers in an authenticated manner. For the Incomplete-key-set the keys are distributed from the key server to PMUs and PDC using a secure unicast channel. In our implementation, the public keys are distributed only once. However, to guarantee a reliable delivery of the keys, we suggest implementing the progressive public key distribution (PPKD) scheme proposed in [100]. Note that the relative difference in the key management overhead between ECDSA and TV-HORS remains the same even when the reliable key distribution scheme is implemented.

Following the proposals in [79], we use 48-bit light-chain elements and 80-bit salt-chain elements for TV-HORS. These parameters along with the t value affect the signature length. For the perfectly-secure Incomplete-key-set we assume a total number of receivers equal to 50. For the basic and the communication-efficient variants of the Incomplete-key-set, we set the system parameters $w = 10$ and $q = 10^{-4}$.

Chapter 5. Experimental Comparison of Multicast Authentication for Grid Monitoring Systems

5.5.2 Performance results and comparison

In Tables 5.2 and 5.3, we present experimental results for the performance of the candidate authentication schemes. The results show how the performance of the schemes vary depending on the values of corresponding parameters for each scheme. Below, we analyse the results for the schemes and draw conclusions on which scheme provides a better security versus performance tradeoff for WAMS.

Table 5.2 – Key management overhead of different multicast authentication schemes.

Scheme	Key management overhead per session (20 sec)							
	key generation time at PMU (ms)		key distribution overhead at PMU (bytes)		Key storage overhead at PMU (bytes)		key storage overhead at PDC per PMU (bytes)	
	L=56	L=80	L=56	L=80	L=56	L=80	L=56	L=80
ECDSA without precomputed tokens	3.367	5.335	29	41	14	20	29	41
ECDSA with precomputed tokens	3'340.367	5'447.335	29	41	28'014	40'020	29	41
TV-HORS	523.439	1'047.332	6'154	6'154	1'538'500	3'077'000	6'154	6'154
Basic Incomplete-key-set	0	0	1'932	2'760	1'932	2'760	175	250
Comm. efficient Incomplete-key-set	0	0	7'728	11'040	7'728	11'040	700	1'000
Perfectly-secure Incomplete-key-set	0	0	350	500	350	500	7	10

Table 5.3 – Performance comparison of multicast authentication schemes using per message computation and communication overheads.

Scheme	Computation overhead per synchrophasor message						Communication overhead (bytes) per synchrophasor message	
	Auth. time (ms)		Verif. time (ms)		Total (ms)		L=56	L=80
	L=56	L=80	L=56	L=80	L=56	L=80		
ECDSA without precomputed tokens	3.431	5.563	0.223	0.327	3.654	5.890	34	48
ECDSA with precomputed tokens	0.104	0.111	0.223	0.331	0.327	0.442	34	48
TV-HORS	0.014	0.014	0.110	0.217	0.124	0.231	88	106
Basic Incomplete-key-set	4.559	4.589	0.068	0.069	4.627	4.658	1'932	2'760
Comm. efficient Incomplete-key-set	18.151	18.361	0.172	0.181	18.323	18.542	138	138
Perfectly-secure Incomplete-key-set	0.848	0.853	0.018	0.019	0.866	0.872	350	500

Incomplete-key-set variants

Even though these schemes use only MAC computations, the large number of such computations introduces large computation and communication overheads per message that they are inapplicable for WAMS. Besides, the Incomplete-key-set requires a key server, which is a single point of failure, whereas EDSA and TV-HORS don't use one. Furthermore, key update for the Incomplete-key-set involves setting up a unicast encrypted channel between the key server and each of the sources and receivers, while EDSA and TV-HORS require only an authenticated multicast delivery of public keys. Therefore, given the large number of sources (and receivers)

in WAMS, the Incomplete-key-set schemes is inefficient from the key server's point of view.

ECDSA variants

The ECDSA without pre-computed tokens scheme performs best in all metrics except in the computation overhead per message. The computation overhead for both security levels is high, which makes it unsuited for WAMS applications that have strict real-time requirement. Adding a cryptographic accelerator hardware to PMUs is one way to speed up signature generation.

Implementing ECDSA with pre-computed tokens significantly improves the computation overhead per message. The pre-computation of the tokens also introduces a non-negligible key-generation overhead (we consider token-generation part of the key generation overhead). However, the tokens for session i are generated during session $i-1$. Hence a token-generation times in Table 5.2 for both security levels during a 20 second long session is within the computational capability of the kind of PMUs deployed in our smart grid. Besides, there is no significant change in the signing overhead between $L=56$ and $L=80$. The small increase in the overall computation overhead can be mitigated by deploying more powerful PDCs or by implementing an optimized ECDSA verification (which we have not implemented). Therefore, the sub-millisecond computation overheads and low communication overheads of ECDSA with pre-computed tokens for both security levels make it an ideal scheme for WAMS applications with real-time requirements for the foreseeable future. This finding is contrary to the generally accepted view that public key cryptography is inapplicable for real-time applications.

TV-HORS

TV-HORS has the lowest computation overhead and relatively low communication overhead per message. The only drawback of TV-HORS is that it requires frequently refreshing the public/private key pair and sending a large public key message to all receivers. WAMS are normally characterized by a large number of PMUs. Unless a proper randomization of key distribution is implemented, a large public key (≈ 6 kbytes) per PMU can cause periodic burst synchronization of packets that can have significant effect on the network bandwidth that could lead to synchrophasor packet losses. The burst of packets from each PMU can also have a non-negligible computation overhead on the receivers if the number of PMUs is in the order of hundreds or thousands. This effect is magnified if the public key has to be sent multiple times to guarantee reliable delivery.

Lu et al. in [91] identify two potential threats in TV-HORS when applied to substation automation systems (SAS) - *delay compression attack* and *key depletion attack*. The sending rate in WAMS is much slower than that of SAS - typically 50 msgs/sec; whereas a typical rate for SMV messages in SAS is 4800 msgs/sec. In our implementation a signing-key update occurs at the end of every epoch. An epoch duration of 80 ms for $L=56$ or 40ms for $L=80$ is long enough for any synchrophasor message to be verified within this time period. In fact, the

Chapter 5. Experimental Comparison of Multicast Authentication for Grid Monitoring Systems

overall end-to-end delay for phasor messages in our smart grid is less than 4 ms. Therefore, the *delay compression attack* is not an issue for WAMS. Moreover, TV-HORS replenishes its key-chain at the end of the last epoch. The time required to generate the whole key-chain for $P=500$ is only 1.047 sec (Table 5.2). Given the relatively lower message sending rate of PMUs, pre-generating the key-chain during the 20 sec duration of session $i - 1$ for session i is within the computational capacity of the PMUs we used in our experiment. Hence, the *key depletion attack* (key generation speed being slower than the key consumption speed) can also be ignored as an issue in WAMS. Finally, the comparison between RSA and TV-HORS in [91] is unfair since the chosen security levels for the two schemes are not the same.

From the above observations, we can conclude ECDSA with pre-computed tokens is the preferred scheme for WAMS applications. In spite of TV-HORS' desirable low computation overhead, it has inherent drawbacks due its hard-deadline requirement to deliver a large public key to receivers within a short duration. Each private key in a TV-HORS key chain has a time window during which it can be used to sign messages. These messages must be verified by the receiver during this assigned time window or else the message is discarded by the receiver. The private key cannot be used to sign messages sent after its time window expires. By the end of the P^{th} epoch, the last private key in the key-chain will be used to sign the v^{th} message of that epoch. Beyond that epoch, the multicast source has to use a new key-chain to sign new messages. However, if the public key for this new key chain is not successfully communicated to the PDCs, they will not be able to verify the messages signed using the private keys from the new key-chain. In our experiment, TV-HORS has only 20 sec to reliably deliver a large public key that is required for the next 20 sec session. As explained above, this 20 sec duration is a hard-deadline since the old key-chain cannot be used to sign more than the number of messages transmitted in 20 sec.

In contrast, ECDSA has a time window of 20 sec to deliver a relatively small public key for the next session. Besides, the 20 sec session duration for ECDSA is a conservative value. Hence, ECDSA could continue to use its old public/private key pair until the next public key is reliably delivered even beyond the 20 sec time window. The only means to extend the life time of the private/public key-chain for TV-HORS to increase P , which in turn introduces key generation, storage and verification overheads.

The two security levels we consider in our experiment are relatively high if we assume an attacker with low computational capabilities. Therefore, utilities who want to protect their WAMS against such an attacker may be willing to consider security levels less than 56. From the results in Table 5.3 we see that when the security level is decreased, the improvement in ECDSA's signing and verification times are much more than the other two schemes'. Hence, for lower security levels, ECDSA with pre-computed tokens is still the preferred scheme for such systems since it will still have lower overheads in all the other metrics.

5.5.3 Support for addition and revocation

All the three schemes support dynamic addition (revocation) of senders and receivers to (from) a multicast group. In all the three schemes, we assume there is a multicast group controller similar to the one described in [76] that is responsible for granting and revoking group membership to PMUs and PDCs and for announcing the addition and revocation of members to the already existing members.

In all schemes addition/revocation of a receiver (PDC) does not cause any change in any of the existing group members. However, addition/revocation of a new source (PMU) to a group introduces some changes to existing PDCs. The group controller has to inform all PDCs (receivers) about the identity of the new PMU. Once informed about the new member, the PDCs will be able to receive the key material (public key for ECDSA and TV-HORS) from the new PMU that they can use to verify messages they will subsequently receive from it. For the Independent-key-set, the key server has to send the secondary key set κ_s to the new PMU s . Performance wise, addition of a new PMU increases the aggregate verification time at the PDC. This increase per every additional PMU is proportional to the verification time in Table 5.3.

Revocation of a PMU involves a controller informing all PDCs about the identity of the revoked PMU and each PDC removing the identity (thus the corresponding authentication key) of the revoked PMU from their list of authentic sources. Performance wise, revocation of a PMU decreases the aggregate verification times at the PDCs. Again, the decrease in the aggregate value per every revoked PMU is proportional to the verification time in Table 5.3.

5.5.4 Impact of the scale of WAMS

The aggregate verification time as well as the key storage requirement at the PDC is proportional to the total number of PMUs in a multicast group. Therefore, the aggregate time that a PDC spends processing (verifying the authenticity, decapsulating and aggregating) synchrophasor messages can be large if the number of PMUs in a group is very large. The IEEE C37.244 Guide for Phasor Data Concentrator Requirements for Power System Protection, Control, and Monitoring [102] specifies a PDC uses a “wait timer” to wait for all messages to arrive from all PMUs before generating the aggregate data and passing it on to the state estimator. The value of the “wait timer” is user defined. Messages from all PMUs should be verified and aggregated before the timer expires. Therefore, a utility needs to determine the computational capacity of the PDC they deploy such that the aggregate processing time for all PMUs is within this limit. Our results in Table 5.3 for the verification time can be used to find the total number of PMUs that a PDC can support. Gomez-Exposito et al. in [103] propose a hierarchical multilevel state estimation framework to avoid using a single powerful central PDC that deals with aggregating synchrophasor data from a large number of PMUs. In such a paradigm, PDCs at the lowest level deal with only a small set of PMUs that are geographically closer to it and the PDCs at higher levels correlate pre-filtered data from PDCs in lower levels and possibly from other PMUs that are close to them. This way, multicast groups will have a

Chapter 5. Experimental Comparison of Multicast Authentication for Grid Monitoring Systems

manageable number of PMUs. The PDCs in the lower levels will be multicast sources in the multicast group for which the higher level PDCs are receivers. Hence, PDCs in the lowest level and in the intermediate levels can be both a receiver in one multicast group and a source in another multicast group.

5.6 Conclusion

In this paper, we have evaluated the performance of available multicast authentication schemes for WAMS. Contrary to the generally accepted notion that public key cryptography is impractical for real-time applications due to its high computation cost, we have shown that an ECDSA implementation that utilizes short-term keys and pre-computed tokens for signature generation provides the required performance for WAMS based real-time applications. TV-HORS is also widely treated as the scheme of choice for real-time applications in smart grid. Our findings show that even though TV-HORS has very low computation overhead even compared to ECDSA with pre-computed tokens, its potential drawbacks due to its hard-deadline requirement to reliably distribute a large public key makes it less preferable than ECDSA.

6 Optimal Software Patching Plan for PMUs

6.1 Introduction

The information and communications technology (ICT) infrastructure in a smart grid network consists of a large number of heterogeneous field devices and servers running a variety of software systems. Utilities deploy state of the art cybersecurity solutions to fend off attacks against the ICT infrastructure. However, no matter how strong the deployed security solutions are, the fact remains that there is no fool proof solution that provides absolute security against all possible attack vectors. There will always be unknown vulnerabilities in the software or hardware that an attacker will discover and exploit through time to compromise one or more of the devices.

Therefore, in addition to deploying state-of-the-art security solutions and taking reactive measures like incident response whenever there is a cyber attack, it is important for utilities to take pro-active measures, such as putting a software patch management in place. Deploying an efficient patch management process for industrial control systems (ICSs) has been addressed in [104–106]. A software patch management system guarantees that patches are applied to all devices running the vulnerable software. It is important that software patches that fix vulnerabilities are rolled out uniformly to all devices as soon as they are available. That is because if a patch is not applied on time and the vulnerability is of public knowledge to an attacker, the attacker will compromise one of the devices by exploiting the vulnerability. Once an attacker gets access to one such device, he can maintain access to the device by privilege escalation even after the patch is applied later on. By maintaining access to the device, the attacker can exploit the trust relationship the device has with other communicating partners in order to launch further attacks and compromise more devices in the network.

In light of the need to roll out software patches fast to all devices, we study the problem of software patching for phasor measurement units (PMUs) in smart grids. PMUs measure time-synchronized, high-resolution phasor data from several locations of the grid and stream this data to a central location called phasor data concentrators (PDC). The PDC time-aligns the measurements from the different PMUs and feeds the time-aligned synchrophasor data to

a real-time state estimator.

Since a PMU placed in a particular bus measures the bus's voltage phasors as well as the current phasors of all the branches incident to the bus, Kirchhoff's laws make it possible for the PMU to indirectly measure the voltage phasors of all incident buses. Therefore, the total number of PMUs required for full system observability is less than the total number of buses in the network. Finding an optimal PMU placement that minimizes the number of PMUs that provide full system observability is a widely studied problem. Research done to address this problem can be broadly categorized into two groups [107]: (1) deterministic approaches that formulate the problem as an ILP problem satisfying some constraints [107–112] (2) meta-heuristic algorithms [113–117].

While deciding on an optimal placement of PMUs to a grid, a utility normally adds a contingency constraint that ensures that the placement provides full observability even when any one of the PMUs fails or is offline for maintenance purposes. Adding more PMUs than the minimum number required for observability also increases measurement redundancy, which improves a state estimator's accuracy as well as its ability to detect bad data [118]. A PMU placement that provides enough measurement redundancy also enables a utility to roll out a software patch to all PMUs by patching a subset of the PMUs at a time while maintaining system observability at all times. In a large-scale power system that deploys a large number of PMUs, applying the patch to one or only a few PMUs at a time is infeasible. The main challenge we address in this chapter is, therefore, a patching plan that minimizes the number of rounds required to patch all the PMUs without losing full observability of the grid during the entire time. Stated otherwise, our goal is to find a partitioning of the set of the deployed PMUs into as few subsets as possible such that all the PMUs in one subset can be patched at a time while all the PMUs in the other subsets provide full observability of the system.

The main contributions of this chapter are:

- We formulate the PMU patching problem as a sensor patching problem and show that the problem of finding an optimal sensor patching plan is NP-complete.
- For the case when a power grid has a radial structure (is a tree), we show the minimum number of rounds required to patch all deployed PMUs is equal to two. We also provide a polynomial-time algorithm that finds the optimal patching plan.
- For mesh grids (non-radial structured grids), we formulate the sensor patching plan problem as a binary integer linear programming (BILP) problem and used a branch-and-bound based ILP solver to compute a patching plan for different bus systems. For grids that are too large to be solved by the ILP-solver, we propose a greedy heuristic algorithm to compute an approximate solution. Moreover, we have proved that finding an optimal solution to the problem is equivalent to maximizing a submodular set function.
-

Although we study the problem as a planning problem for offline time of PMUs caused by software patching, it can be generalized to any scheduled maintenance work that affects all PMUs and requires a PMU to go offline for some time.

The rest of this chapter is organized as follows. In Section 6.2 we state assumptions, introduce the system model and define the PMU patching problem. In Section 6.3, we formally define the PMU patching problem as an instance of a sensor patching problem using set theoretic approach. We also show it is NP-complete. The BILP formulation of the problem using the asymmetric representatives method is also introduced in this section. In Section 6.4 we introduce a polynomial-time algorithm that finds an optimal two round patching plan on a tree and prove its correctness. Section 6.5 discusses the heuristic algorithm for the general case networks. Results from the heuristic approach and the ILP solver are presented and compared in Section 6.6. Section 6.7 provides concluding remarks and future directions.

6.2 PMU Patching Problem

In this section, we briefly describe our assumptions on state estimation and system observability. We also introduce the system model and define the PMU patching problem.

6.2.1 State Estimation and Assumptions

The static estimation of a power system state is defined as determining the phase-to-ground voltage phasors at all the system buses through analysis of measurements collected from different locations of the grid [119]. The state estimator uses the set of measurements along with the power system model as an input to compute the most likely state of the grid at a given time. The set of measurements may come from conventional P-Q measurement devices that measure real and reactive nodal power injection and real and reactive line power flows or from phasor measurement units (PMUs) that directly measure nodal voltage magnitudes and phase angles and branch current magnitudes and phase angles.

The measurement model for system estimators is defined by [120]

$$\mathbf{z} = \mathbf{h}(\mathbf{x}) + \mathbf{v} \quad (6.1)$$

where $\mathbf{z} = (z_1, z_2, \dots, z_m)^T$ is an m -dimensional measurement vector; $\mathbf{x} = (x_1, x_2, \dots, x_n)^T$ is an n -dimensional state vector (phase-to-ground voltage phasors at all the system buses); $\mathbf{v} = (v_1, v_1, \dots, v_m)^T$ is an m -dimensional random measurement error vector. The measurement errors are assumed to be independent, zero-mean Gaussian variables with known covariance matrix \mathbf{W} . \mathbf{W} is a diagonal matrix with values σ_i^2 , where σ_i is the standard deviation of the error associated with measurement i . $\mathbf{h}(\mathbf{x}) = (h_1(\mathbf{x}), h_2(\mathbf{x}), \dots, h_m(\mathbf{x}))^T$ is a vector of power flow functions relating error free measurements to the state variables.

Although state estimation using AC power flow model is more accurate, it can be computation-

ally expensive and may not always converge to a solution. Therefore, power system engineers use DC power flow model which is a simplification, and linearization of an AC power flow model. In the DC power flow model, the measurement model is represented by the following linear regression model [28],

$$\mathbf{z} = \mathbf{H}\mathbf{x} + \mathbf{v} \quad (6.2)$$

where \mathbf{H} is an $m \times n$ matrix that reflects the configuration of the system.

Given the imperfect set of measurements \mathbf{z} , the purpose of a state estimator is to determine an optimal estimate $\hat{\mathbf{x}}$ of the system state that best fits the measurement model.

In this Chapter, we consider only measurements from PMUs. We assume that a PMU placed in a bus has enough number of channels to measure the bus' voltage phasor as well as the phasor currents of all lines incident to the bus.

6.2.2 Observability Rules

System observability depends on the connectivity among the buses as well as the location where measurement devices are placed. A power system is fully observable if all its buses are observable. A bus is said to be observable if the bus' state (voltage phasors) can be estimated from the set of available measurements. As stated above, we focus on system observability using measurements from PMUs. A bus is observable if a PMU is placed at the bus or if any of its neighboring buses have a PMU placed at them [107, 109, 121]. This condition implies that the system is fully observable if and only if the matrix \mathbf{H} introduced in Equation 6.2 has full rank. There are other observability rules that exploit the presence of zero-injection buses that we don't consider in this chapter and leave for future work.

6.2.3 System Model and Problem Definition

We model a power system as an undirected graph on the set of vertices $B = \{1, 2, \dots, n\}$ that represent the buses. We define the set $P = \{1, 2, \dots, m\}$ as the set of PMUs deployed in the power system and $\beta: P \rightarrow B$ the mapping such that $\beta(j) = b$ when PMU j is placed at bus b . From hereon, we use *PMU bus* to refer to a bus where a PMU is placed at.

During the time a utility rolls out a software patch, a PMU in a grid is in one of the following three states:

- State (1): unpatched and streaming phasor measurement,
- State (2): being patched and offline,
- State (3): patched and streaming phasor measurement.

6.3. The Sensor Patching Problem (SPP)

We assume that a state estimator receives and processes measurements from PMUs that are in state (1) as well as those in state (3) to compute the system state during the patching time window. Further, we assume that no PMU goes offline due to failure during the time a software patch is being rolled out.

A PMU patching problem is stated as finding a partitioning of deployed PMUs into as few disjoint groups as possible such that all the PMUs in one group can be transformed from State (1) through State (2) to State (3) in one round while the PMUs in all the other subsets provide full system observability during that round. Once such a partition of the PMU set is found, the patch is applied to all PMUs in as many rounds as there are subsets in the partition.

Note that a feasible patching plan exists if and only if every bus is observed by at least two PMUs. Indeed, if each bus is observed by at least two PMUs, there exists a patching plan that patches one PMU at a time and all the buses that are observed by this PMU will still be observed during that round by the remaining PMU(s). Such a patching plan requires as many rounds as there are deployed PMUs. Conversely, if we have full observability during all patching rounds, it means each bus has at least one PMU that is not being patched at any given round. Since this PMU has to be patched in one of the rounds, the bus must have at least one other PMU which makes it observable during that round. Hence the bus is observed by at least two PMUs (See Figure 6.1 for an example).

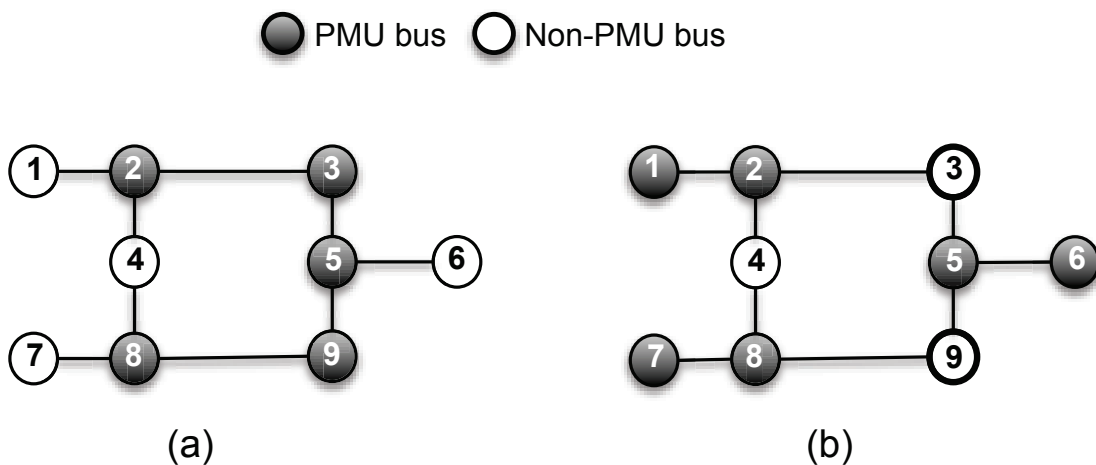


Figure 6.1 – A feasible patching plan exists if and only if every bus is observed by at least two PMUs. (a) no feasible patching plan (b) A feasible patching plan exists.

6.3 The Sensor Patching Problem (SPP)

In this section we give a set theoretic formulation of our problem, which we call the sensor patching problem (SPP). Further, we prove that it is NP-complete and give an ILP formulation.

6.3.1 Set Theoretic Formulation and NP-completeness Proof of SPP

Let $B = \{1, 2, 3, \dots, n\}$ is a finite set of sites to be observed and $P = \{1, 2, 3, \dots, m\}$ is a finite set of sensors that observe the sites. Further, let $\Gamma : B \rightarrow 2^P$ be a mapping such that $\Gamma(b)$ is the set of sensors in P that observe site $b \in B$. In our PMU patching problem, $\Gamma(b)$ is the set of PMUs placed in bus b , if there is one, and in any of the buses that are adjacent to b .

Definition 1. Given a non-empty finite set P , a k -tuple $\{c_1, c_2, c_3, \dots, c_k\}$ partitions P if:

- $c_i \neq \emptyset, \forall i \in \{1, 2, \dots, k\}$.
- $\cup_{i=1}^k c_i = P$.
- $c_i \cap c_j = \emptyset$, for $1 \leq i < j \leq k$.

A feasible sensor patching plan is a partition $\{c_1, c_2, c_3, \dots, c_k\}$ of the set P such that the following observability condition is satisfied:

$$|\Gamma(b) \setminus c_i| \geq 1, \forall b \in B, \text{ and } i = 1, 2, \dots, k \quad (6.3)$$

Each subset c_i in the family of subsets that partition P defines the set of sensors that are patched at round i . A given sensor placement P has a feasible patching plan if and only if $|\Gamma(b)| \geq 2, \forall b \in B$, i.e., each site is observed by at least two sensors.

The sensor patching problem (SPP) is finding a sensor patching plan that minimizes k . Below, we show that the decision problem version of the SPP is NP-complete.

SPP Decision problem:

- Instance: Finite sets B and P , a mapping $\Gamma : B \rightarrow 2^P$ and an integer $k \geq 2$.
- Question: Is there a partitioning of the set P into at most k disjoint subsets $\{c_1, c_2, c_3, \dots, c_k\}$ such that the observability condition in Eq. 6.3 is satisfied?

Theorem 1. *The decision version of SPP is NP-complete.*

Proof. The first step of the proof is to show that SPP is in NP. Given a nondeterministically selected partition of P into k disjoint subsets, we can determine if the partition satisfies the observability condition in Eq. 6.3 in polynomial time. Hence SPP is in NP.

The second step of our proof is to select a known NP-complete problem and construct a polynomial-time transformation that maps any instance of the NP-complete problem to an SPP problem. For our proof, we choose the hypergraph coloring problem (HCP), which is NP-complete.

6.3. The Sensor Patching Problem (SPP)

A hypergraph is denoted by $H = (V, E)$, where V is a finite set of vertices and E is a set of hyperedges whose elements are subsets $e \subseteq V$ such that $\cup_{e \in E} = V$. Given a hypergraph $H = (V, E)$ and an integer $k \geq 2$, a k -coloring of a hypergraph H is an allocation of colors to the vertices such that:

- A vertex has just one color.
- We use k colors to color all the vertices.
- No hyperedge with a cardinality more than one has all its vertices of the same color, i.e., no such hyperedge is monochromatic.

Any feasible coloring of a hypergraph using k colors induces a partition of the set of vertices V in k color classes: $\{c_1, c_2, c_3, \dots, c_k\}$ such that for $e \in E, |e| \geq 2$ then $e \not\subseteq c_i, \forall i \in \{1, 2, 3, \dots, k\}$ [122].

HCP Decision problem:

- Instance: Hypergraph $H = (V, E)$, an integer $k \geq 2$.
- Question: Is there a partitioning of the set of vertices V into at most k classes $\{c_1, c_2, c_3, \dots, c_k\}$ such that $\forall e \in E, |e| \geq 2, e \not\subseteq c_i, \forall i \in \{1, 2, 3, \dots, k\}$?

Having introduced HCP, let's now look at how to transform an instance of an HCP to an instance of SPP in polynomial time. Given an instance $HCP(V', E', k)$ where $|e'| \geq 2, \forall e' \in E'$, we construct an instance $SPP(P', B', \Gamma, k)$, where $P' \leftarrow V', B' \leftarrow E', k \leftarrow k, \Gamma \leftarrow Id_E$ such that $\Gamma : e' \rightarrow e', \forall e' \in E'$. This transformation from HCP to SPP is a polynomial-time (trivial) transformation.

Assume we have an oracle that solves any given SPP decision problem. The oracle outputs "yes" to the instance $SPP(P', B', \Gamma, k)$ if and only if there exists a partition of P' to k subsets $\{c_1, c_2, c_3, \dots, c_k\}$ such that $(\Gamma(b') \setminus c_i) \neq \emptyset, \forall b' \in B', \forall i \in \{1, 2, 3, \dots, k\}$. Because of the mapping stated above, this is also the same as saying the oracle outputs "yes" if and only if $e' \not\subseteq c_i, \forall e' \in E', \forall i \in \{1, 2, 3, \dots, k\}$, which is the same as the "yes" output if there is a solution to the HCP decision problem. Therefore, if we can transform HCP to SPP and solve it, it means SPP is at least as hard as HCP. Hence, SPP is NP-complete.

By showing that the SPP is as hard as HCP, it also follows that even if we were told the set of sensors in a given instance of SPP could be patched in only two rounds, there is no efficient algorithm that can find any reasonable approximation for the number of rounds.

□

6.3.2 BILP Formulation of SPP

Now that we have shown SPP is NP-complete, we formulate it as a binary integer linear programming (BILP) minimization problem and use a BILP solver to find optimal solutions for small size networks and sub-optimal solutions for large network sizes.

To formulate SPP as a BILP problem, we use the *representatives* method introduced in [123]. As stated above, our goal is to find the minimum number of subsets $\{c_1, c_2, \dots, c_k\}$ that partition set P such that for any $b \in B$ the sensors in $\Gamma(b)$ cannot all be assigned to the same subset. The representatives formulation, as its name indicates, chooses one element from each of the partitioning subsets as a representative element to the subset (to all the elements in the subset). Therefore, each element in P can be in one of two states: either it represents the subset it is an element of or there exists another element that represents its subset. To describe this, we use an $m \times m$ matrix r of binary variables where $m = |P|$ is the number of sensors and the variables are defined by:

$$r_{i,j} = \begin{cases} 1 & \text{if element } i \text{ represents element } j, \\ 0, & \text{otherwise} \end{cases} \quad (6.4)$$

Variable $r_{i,j}$ can be 1 only if elements i and j are in the same subset. By definition the representative elements are the elements i with $r_{i,i} = 1$. If $r_{i,i} = 1$, the row $r_{i,-}$ is an indicator vector of one of the subsets that partition the set P .

A BILP formulation of SPP is given as follows:

$$\min \quad \sum_{i=1}^m r_{i,i} \quad (6.5)$$

$$\text{s.t.} \quad \sum_{i=1}^m r_{i,j} = 1, \forall j \in \{1, 2, \dots, m\} \quad (6.6)$$

$$\sum_{j \in \Gamma(b)} r_{i,j} < |\Gamma(b)| r_{i,i}, \quad \forall b \in B, \forall i \in \{1, 2, \dots, m\} \quad (6.7)$$

$$r_{i,j} \in \{0, 1\}, \forall i, j \in \{1, 2, \dots, m\} \quad (6.8)$$

Claim 1. *A solution to the BILP problem 6.5 - 6.8 is an optimal solution to the SPP.*

Proof. Constraint (6.6) guarantees each sensor has only one representative. Since each subset has only one representative sensor, this constraint is equivalent to saying each sensor is assigned to only one subset. This means two things: first, it means no two subsets can have a common element; second, the union of the subsets is P . Therefore, the subsets are feasible

partitions of set P . Constraint (6.7) makes sure that the sensors in the set $\Gamma(b)$ cannot all choose the same representative sensor i and requires that $r_{i,i} = 1$ if sensor i is chosen as representative to one of the sensors in $\Gamma(b)$. This constraint guarantees that every bus has at least two of the sensors that observe it assigned to different subsets, i.e., the observability condition is satisfied. Constraint (6.8) states the variables are binary.

All the constraints represent the constraints for an SPP. Since the objective function (6.5) minimizes the number of representative sensors, which is the same as minimizing the number of subsets that partition P , the solution to the BILP problem is an optimal solution to the optimization version of SPP. \square

In Section 6.6, we solve the above BILP problem using the LP solver package *lpsolve* [124] for different bus systems.

6.4 The Case of Radial Structured Networks

In section 6.3, we have seen that the general case PMU placement is NP-complete. Therefore, the problem is in general solved using a heuristic approach. However, there is an important case (when the grid has a radial structure) where the problem can be optimally solved in polynomial time. The special case is of interest to us because the active configuration of many power distribution networks has a radial (tree) structure.

Theorem 2.

1. *Given a system model as stated in Section 6.2 where the graph is a tree and a PMU placement P that has a feasible patching plan ($\forall b \in B, |\Gamma(b)| \geq 2$), the minimum number of rounds required to patch all PMUs is equal to 2.*
2. *An optimal patching plan is given by Algorithm 1; its complexity is $O(|B|^2)$.*

In the description of the algorithm, we phrase the problem as a two-coloring problem and say that two PMUs have the same color if they are allocated to the same round. We say c_0 [resp. c_1] is the set of PMUs that are assigned to the first [resp. second] round i.e., colored in, say, red [resp. blue].

Chapter 6. Optimal Software Patching Plan for PMUs

Algorithm 1 Find a 2-round patching plan on a tree

Inputs: P, B, Γ, β

Output: c_0, c_1

Steps

1. Select one bus $\rho \in B$ and call it the root of the tree.
2. For each $j \in P$, color j according to its distance from the root $d(\beta(j), \rho)$ and build the color classes c_0 and c_1 as follows:

$$\forall i \in \{0, 1\}, c_i = \cup\{j : i = d(\beta(j), \rho) \bmod 2\} \quad (6.9)$$

3. **While** $\exists b \in B$ that violates the condition:

$$\forall i \in \{0, 1\}, |\Gamma(b) \setminus c_i| \geq 1 \quad (6.10)$$

- (a) Select b with the maximum $d(b, \rho)$ (breaking ties arbitrarily).
- (b) Select a PMU bus u that is a child of b and let T_u denote the sub-tree rooted at u .
- (c) Update the color assignment of the PMUs by flipping the color of each PMU placed in a bus in T_u .

4. **End while**
-

Figure 6.2 shows how the algorithm progresses on a 13-bus power system that deploys 10 PMUs.

To prove that Algorithm 1 is correct, we need to verify two properties:

- The algorithm is well-defined, i.e., at Step 3b, vertex $b \in B$ always has a child that is a PMU bus.
- The algorithm terminates.

Lemma 1. *At the beginning of each iteration, the selected vertex b that violates the condition in Eq. (6.10) is not a PMU bus (and hence one of its children is a PMU bus). Moreover, after the iteration, the updated coloring causes vertex b to satisfy the condition and no vertex that satisfies the condition by the previous coloring violates the condition as a result of the updated coloring.*

Proof. We claim that initially all the vertices that violate the condition in Eq. (6.10) are not PMU buses. This is true by the design of the initial coloring. By definition, all PMUs in $\Gamma(b)$ except the PMU placed in b (if there is one) are assigned the same color, which is different

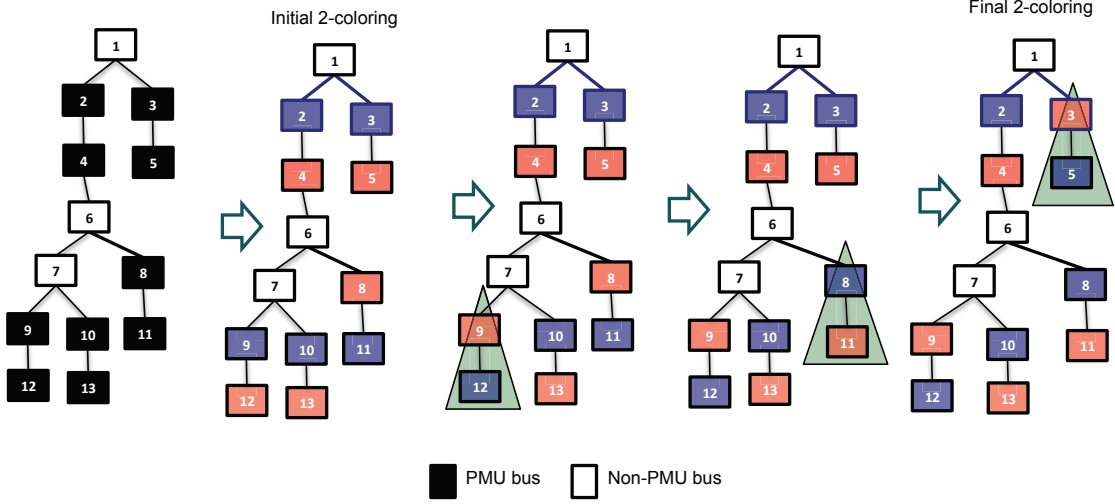


Figure 6.2 – A polynomial-time algorithm to find two disjoint subsets of a set of 10 deployed PMUs in a 13-bus system such that one subset of PMUs can provide full observability while the other subset of PMUs is being patched.

from the color of the PMU in b , in the initial coloring. Therefore, $|\Gamma(b)| \geq 2$ implies that b cannot be a PMU bus if it violates the condition.

Now consider an iteration where the coloring is changed. Let b be the selected violating vertex. As we will show that our algorithm does not introduce any new violating vertex and in our initial coloring we have shown b is not a PMU bus, $|\Gamma(b)| \geq 2$ implies that at least one child of b must be a PMU bus. So Step 3b is well defined. Let $u \in B$ be the selected child of b that is a PMU bus. Recall that T_u denotes the subtree rooted at u . Now, as b was selected to be the violating vertex farthest from the root, no vertex in T_u violates the condition by the initial coloring. Moreover, any vertex in T_u is observed by a PMU placed in T_u because b is not a PMU bus. Hence, flipping the colors of the PMUs placed in T_u does not introduce any new vertices that violate the condition. In other words, no newly violating vertices are introduced by our operation.

Now let us show that b is no longer violating the condition at the end of the iteration. Let c_0, c_1 and c'_0, c'_1 denote the coloring before and after the iteration, respectively. We have $|\Gamma(b)| \geq 2$ and since the condition was violated initially, we have for some $i \in \{0, 1\}$

$$|\Gamma(b) \setminus c_i| = 0 \text{ and } |\Gamma(b) \setminus c_{i \oplus 1}| \geq 2 \quad (6.11)$$

Here \oplus denotes addition modular 2. So after flipping the color of PMU bus u , we have

$$|\Gamma(b) \setminus c'_i| = 1 \text{ and } |\Gamma(b) \setminus c'_{i \oplus 1}| \geq 1 \quad (6.12)$$

and hence b satisfies the condition.

In the initial coloring, the maximum possible number of violating vertices is $|B|$. Since each iteration in our algorithm fixes only one violating vertex, all violating vertices are fixed in a maximum of $|B|$ iterations. Each iteration runs in linear time because the maximum possible number of vertices in any subtree T_u is $|B|$. Hence the complexity of the algorithm to obtain an optimal coloring is $O(|B|^2)$. The final coloring partitions the set of PMUs into two disjoint color classes. Consequently, all the PMUs can be patched in only two rounds by patching PMUs in one color class in the first round and those in the other color class in the second round. \square

6.5 Approximation Algorithm for Mesh Grid Structure

It is common to model NP-complete problems as ILP problems and use ILP solvers to find optimal or suboptimal solutions for relative small size of input. However, ILP solvers tend to be too slow to find even a suboptimal solution as the input size grows. The alternative is to use heuristic algorithms that find approximate solutions much faster than ILP solvers. For this reason, we propose a heuristic algorithm that finds an approximate solution to the SPP, which we have already shown to be NP-complete.

6.5.1 A Greedy Approximation Algorithm

Before going to the details of the heuristic algorithm, let's first define observability set o_j as the set of buses that are observed by PMU j . Given the set of buses in the grid B and the set of PMUs P , o_j is defined as follows:

$$o_j = \{b : b \in B, j \in \Gamma(b)\} \quad (6.13)$$

The collection $\mathcal{O} = \{o_j : j \in P\}$ is a set of all the observability sets of the deployed PMUs.

The heuristic algorithm we propose follows a greedy approach that maximizes the set of PMUs that are patched at each round while still maintaining full system observability. Given a set of unpatched PMUs P , finding the maximum number of PMUs to patch is equivalent to finding the minimum number of PMUs that provide full system observability, which is exactly the same as solving the minimum set cover (MSC) problem over a universe B and a collection of subsets \mathcal{O} . The set of PMUs to patch is, therefore, the set that contains the PMUs that are not in the MSC solution. Once these PMUs are patched, they will resume streaming for the rest of the time. Therefore, the observability condition for the set of buses that are in the observability sets of these PMUs will always be satisfied for the remaining patch rounds. Hence, before we select the next set of PMUs to patch, we perform the following preprocessing:

- Remove all the observability sets of all the already patched PMUs from \mathcal{O} .

- Remove all the buses in the observability sets of the already patched PMUs from the universe B .
- Remove all the buses in the observability sets of the already patched PMUs from the observability set of the yet unpatched PMUs.

After the pre-processing, we proceed with the same greedy approach (solving the MSC problem) for the updated universe B and the updated collection \mathcal{O} . We repeat this process until $B = \emptyset$ (until all buses are observed by the already patched PMUs). At this stage, if there are still any PMUs that are not yet patched ($\mathcal{O} \neq \emptyset$), we patch all such PMUs at once in the final round. Algorithm 2 shows the pseudocode for our heuristic algorithm. The algorithm outputs a collection $\mathcal{C} = \{c_1, c_2, \dots, c_k\}$, where $c_i \subset P$ is the set of PMUs patched in round i and k is the total number of rounds required to patch all the PMUs.

Since the MSC problem is itself NP-complete, we use the most commonly used greedy heuristic to solve it. The greedy heuristic for MSC chooses the subset that maximizes the number of new elements in the universe B that are not yet covered by the already selected subsets.

6.5.2 Formulation as Submodular Maximization

Here we show that the SPP can be formulated as the maximization of a submodular set function. It is known that a greedy heuristic guarantees a reasonably good approximation to the optimal solution for problems that are submodular and the SPP is in that category. This may be used as a justification to why the greedy algorithm proposed above can be expected to perform well.

Given the set of PMU's P and $m = |P|$, let's define the collection Ψ as:

$$\Psi = \{\psi : \psi \subseteq P, \cup_{j \in (P \setminus \psi)} o_j = B\} \quad (6.14)$$

In other words, an element in Ψ is a set of PMUs that can be taken offline and full system observability can still be maintained. From this, it follows that if $\mu \in \Psi$ and $\mu' \subset \mu$, then $\mu' \in \Psi$.

Consider a non-negative submodular set function $f : 2^\Psi \rightarrow \mathbb{R}_+$ on Ψ that assigns a non-negative number to every subset of the set Ψ .

Claim 2. *A collection $\mathcal{C} \subset \Psi$ that maximizes the following set function,*

$$f(\mathcal{C}) = Q \cdot |\cup_{c \in \mathcal{C}} c| - |\mathcal{C}|, \text{ where } Q > m \text{ is a constant.} \quad (6.15)$$

is an optimal solution to the SPP.

Proof. Let $\mathcal{C}^* = \{c_1, c_2, \dots, c_{k^*}\}$ be an optimal solution to the SPP. Passing \mathcal{C}^* as an input to f ,

Algorithm 2 Partiton P into minimum patchable subsets using a greedy heuristic

```
1: Input:  $\mathcal{O}, B$ 
2: Output:  $\mathcal{C} := \{c_1, c_2, \dots, c_k\}$ 
3:  $round = 1$ 
4: while  $B \neq \emptyset$  do
5:    $\sigma := \text{FindMSC}(\mathcal{O}, B)$ 
6:    $c_{round} := \{j : o_j \notin \sigma\}$ 
7:    $\mathcal{C} := \mathcal{C} \cup \{c_{round}\}$ 
8:    $\mathcal{O} := \mathcal{O} \setminus \{o_j : j \in c_{round}\}$ 
9:    $B := B \setminus \{o_j : j \in c_{round}\}$ 
10:  for  $o_u \in \mathcal{O}$  do
11:     $o_u := o_u \setminus \{o_j : j \in c_{round}\}$ 
12:  end for
13:   $round++$ 
14: end while
15: if  $\mathcal{O} \neq \emptyset$  then
16:    $c_{round} := \{j : o_j \in \mathcal{O}\}$ 
17:    $\mathcal{C} := \mathcal{C} \cup \{c_{round}\}$ 
18: end if

19: procedure  $\text{FINDMSC}(\mathcal{O}, B)$ 
20:    $B' := \emptyset$ 
21:    $\sigma := \emptyset$ 
22:   while  $B' \neq B$  do
23:      $MaxCount := 0$ 
24:      $idx = 0$ 
25:     for all  $o_j \in \mathcal{O}$  do
26:       if  $|o_j \setminus B'| > MaxCount$  then
27:          $MaxCount := |o_j \setminus B'|$ 
28:          $idx := j$ 
29:       end if
30:     end for
31:      $B' := B' \cup o_{idx}$ 
32:      $\sigma := \sigma \cup \{o_{idx}\}$ 
33:   end while
34:   Return  $\sigma$ 
35: end procedure
```

we get.

$$f(\mathcal{C}^*) = Q \cdot |\cup_{c \in \mathcal{C}^*} c| - |\mathcal{C}^*| = -k^* + Q \cdot m \quad (6.16)$$

We want to show that $f(\mathcal{C}) < -k^* + Q \cdot m$ for any input $\mathcal{C} = \{c_1, c_2, \dots, c_k\}$, where $k > k^*$.

Given a collection $\mathcal{C} = \{c_1, c_2, \dots, c_k\}$ for some $k \geq 1$,

$$f(\mathcal{C}) = -k + Q \cdot |\cup_{c \in \mathcal{C}} c| \quad (6.17)$$

Let $m' = |\cup_{c \in \mathcal{C}} c|$

$$\begin{aligned} f(\mathcal{C}) &= -k + Q \cdot m' \leq -k + Q \cdot (m - 1) \\ f(\mathcal{C}) &\leq -k + Q \cdot (m - 1) \end{aligned} \quad (6.18)$$

Since $Q > m$ and $k^* < k \leq m$, it is easy to show that

$$-k + Q \cdot (m - 1) < -k^* + Q \cdot m \quad (6.19)$$

Therefore,

$$\begin{aligned} f(\mathcal{C}) &\leq -k + Q \cdot (m - 1) < -k^* + Q \cdot m \\ f(\mathcal{C}) &< -k^* + Q \cdot m, \quad \forall \mathcal{C} \text{ where } |\mathcal{C}| > k^* \end{aligned} \quad (6.20)$$

This means,

$$\max_{\mathcal{C} \in \Psi} f(\mathcal{C}) = -k^* + Q \cdot m \quad (6.21)$$

which is the same as the optimal solution for SPP. \square

Claim 3. *The set function f in Eq. 6.15 is submodular.*

Proof. Function f is submodular if for all subsets $Y \subset X \subset \Psi$ and all $\mu \in \Psi \setminus X$,

$$f(Y \cup \{\mu\}) - f(Y) \geq f(X \cup \{\mu\}) - f(X) \quad (6.22)$$

We want to see if this holds true for f given in Eq. 6.15,

$$\begin{aligned} &Q \cdot (|\cup_{y \in Y} y \cup \mu| - |Y \cup \{\mu\}|) - Q \cdot (|\cup_{y \in Y} y|) + |Y| \stackrel{?}{\geq} \\ &Q \cdot (|\cup_{x \in X} x \cup \mu| - |X \cup \{\mu\}|) - Q \cdot (|\cup_{x \in X} x|) - |X| \end{aligned} \quad (6.23)$$

Using the substitutes $\mathbb{Y} = \cup_{y \in Y} y$ and $\mathbb{X} = \cup_{x \in X} x$, we get

$$\begin{aligned} Q \cdot |\mathbb{Y} \cup \mu| - |\mathbb{Y}| - 1 - Q \cdot |\mathbb{Y}| + |\mathbb{Y}| &\stackrel{?}{\geq} \\ Q \cdot |\mathbb{X} \cup \mu| - |\mathbb{X}| - 1 - Q \cdot |\mathbb{X}| + |\mathbb{X}| &\end{aligned} \quad (6.24)$$

$$|\mathbb{Y} \cup \mu| - |\mathbb{Y}| \stackrel{?}{\geq} |\mathbb{X} \cup \mu| - |\mathbb{X}| \quad (6.25)$$

$$|\mathbb{Y} \cup \mu| + |\mathbb{X}| \stackrel{?}{\geq} |\mathbb{X} \cup \mu| + |\mathbb{Y}| \quad (6.26)$$

$$|(\mathbb{Y} \cup \mu) \cup \mathbb{X}| + |(\mathbb{Y} \cup \mu) \cap \mathbb{X}| \stackrel{?}{\geq} |\mathbb{X} \cup \mu| + |\mathbb{Y}| \quad (6.27)$$

$$|\mu \cup \mathbb{X}| + |\mathbb{Y} \cup (\mu \cap \mathbb{X})| \stackrel{?}{\geq} |\mathbb{X} \cup \mu| + |\mathbb{Y}| \quad (6.28)$$

$$|\mathbb{Y} \cup (\mu \cap \mathbb{X})| \stackrel{?}{\geq} |\mathbb{Y}| \quad (6.29)$$

If $\mu \cap \mathbb{X} = \emptyset$, the right hand side of Eq. 6.29 is the same as the left hand side. Otherwise, the right hand side is strictly greater than the left hand side. Hence, f is submodular. \square

6.6 Simulation Results and Comparisons

We compare the performance of our heuristic algorithm to results obtained from an ILP solver for different feeder bus systems. In order to make the comparison, we first find an optimal PMU placement that has a feasible patching plan by solving the following BILP minimization problem:

$$\min \quad \sum_{j=1}^n p_j \quad (6.30)$$

$$\sum_{j=1}^n a_{i,j} \cdot p_j \geq 2 \quad (6.31)$$

$$p_j \in \{0, 1\}, j \in \{1, 2, \dots, n\} \quad (6.32)$$

where

$$a_{i,j} = \begin{cases} 1 & \text{if bus } i = j \text{ or } i \text{ is incident to } j \\ 0, & \text{otherwise} \end{cases} \quad (6.33)$$

The set $\{p_i : i \in B\}$ is a set of binary variables such that $p_i = 1$ if a PMU is placed in bus i and $p_i = 0$ otherwise. Solving the above BILP problem returns a placement with the fewest number of PMUs such that each bus is observed by at least two PMUs. That means it is guaranteed that such a placement has a feasible patching plan.

We use the open source ILP solver *lpsover* [124] to solve the above BILP as well as the BILP

6.6. Simulation Results and Comparisons

Table 6.1 – Performance comparison of ILP solver and a greedy algorithm for PMUs’ patching plan.

Bus system	PMU Placement		Patching Plan from ILP Solver			Greedy Patching Plan		
	#PMUs	PMU buses	Patchable groups	#R	Exec. time (sec.)	Patchable groups	#R	Exec. time (sec.)
14-Bus	9	[1, 2, 4, 6, 7, 8, 9, 11, 13]	[1, 4, 7, 11, 13] [2, 6, 8, 9]	2	0.014	[2, 8, 11, 13] [1, 6, 7, 9] [4]	3	0.004
30-Bus	21	[1, 2, 3, 5, 6, 8, 9, 10, 11, 12, 13, 15, 16, 18, 19, 21, 24, 25, 26, 27, 29]	[1, 2, 6, 9, 13, 15, 16, 19, 21, 25, 29] [3, 5, 8, 10, 11, 12, 18, 24, 26, 27]	2	0.127	[3, 5, 8, 11, 13, 16, 19, 21, 24, 26, 29] [1, 2, 6, 9, 10, 15, 18, 25, 27] [12]	3	0.005
39-Bus	28	[2, 3, 6, 8, 9, 10, 11, 13, 14, 16, 17, 19, 20, 22, 23, 25, 26, 29, 30, 31, 32, 33, 34, 35, 36, 37, 38, 39]	[2, 6, 9, 11, 14, 17, 19, 20, 22, 23, 25, 29, 32] [3, 8, 10, 13, 16, 26, 30, 31, 33, 34, 35, 36, 37, 38, 39]	2	0.399	[8, 11, 14, 17, 30, 31, 32, 33, 34, 35, 36, 37, 38, 39] [3, 6, 9, 10, 13, 19, 20, 22, 23, 25, 29] [2, 16, 26]	3	0.012
57-Bus	33	[1, 2, 4, 6, 9, 11, 12, 15, 19, 20, 22, 24, 25, 26, 28, 29, 30, 32, 33, 34, 36, 37, 38, 39, 41, 44, 46, 47, 50, 51, 53, 54, 56]	[1, 6, 19, 22, 32, 36, 39, 41, 44, 51] [4, 9, 11, 15, 20, 24, 25, 28, 34, 37, 47, 50, 53, 56] [2, 12, 26, 29, 30, 33, 38, 46, 54]	3	340	[2, 6, 12, 19, 22, 28, 30, 33, 34, 39, 41, 44, 47, 51, 54] [1, 4, 9, 11, 20, 25, 26, 32, 37, 46, 50, 53, 56] [15, 24, 29, 36, 38]	3	0.016
118-Bus†	68	[1, 2, 5, 6, 9, 10, 11, 12, 15, 17, 19, 21, 22, 24, 25, 26, 27, 28, 29, 32, 34, 35, 37, 40, 41, 44, 45, 46, 49, 50, 51, 52, 54, 56, 59, 62, 64, 65, 66, 68, 70, 71, 73, 75, 76, 77, 78, 80, 83, 85, 86, 87, 89, 90, 92, 94, 96, 100, 101, 105, 106, 108, 110, 111, 112, 114, 116, 117]	[2, 5, 10, 12, 22, 24, 27, 28, 32, 34, 37, 41, 45, 49, 52, 56, 62, 64, 73, 75, 77, 80, 85, 87, 90, 94, 101, 105, 110, 116] [1, 6, 9, 11, 17, 21, 25, 29, 35, 40, 44, 46, 50, 51, 54, 59, 65, 66, 68, 70, 71, 76, 78, 83, 86, 89, 92, 96, 100, 106, 108, 111, 112, 114, 117] [15, 19, 26]	3	500	[2, 6, 10, 15, 19, 22, 26, 29, 35, 41, 44, 46, 54, 56, 65, 66, 73, 76, 78, 83, 87, 90, 96, 101, 106, 108, 111, 112, 114, 116, 117] [1, 9, 11, 12, 21, 27, 28, 32, 34, 40, 45, 50, 52, 62, 64, 71, 75, 77, 80, 86, 89, 94, 105, 110] [5, 17, 25, 37, 49, 51, 59, 68, 70, 85, 92, 100] [24]	4	0.125
189-Bus†	160	[1, 4, 5, 6, 7, 8, 9, 10, 12, 13, 14, 15, 16, 17, 18, 19, 20, 21, 24, 26, 27, 28, 29, 31, 32, 33, 34, 35, 36, 37, 38, 39, 40, 41, 42, 43, 44, 45, 47, 48, 49, 50, 51, 52, 53, 54, 55, 58, 60, 61, 62, 63, 64, 65, 66, 67, 68, 69, 70, 71, 72, 73, 74, 76, 77, 78, 79, 80, 81, 82, 83, 84, 85, 86, 87, 88, 89, 90, 91, 94, 95, 96, 97, 98, 99, 100, 101, 102, 104, 105, 106, 107, 108, 110, 111, 112, 113, 114, 115, 117, 118, 120, 121, 122, 123, 124, 126, 127, 128, 129, 130, 131, 134, 135, 136, 137, 138, 139, 141, 142, 143, 144, 145, 147, 150, 151, 152, 153, 154, 155, 156, 157, 159, 160, 162, 163, 164, 165, 167, 168, 169, 170, 171, 172, 173, 174, 175, 176, 178, 179, 180, 181, 182, 183, 184, 185, 186, 187, 188, 189]	[1, 9, 13, 14, 15, 17, 24, 26, 31, 39, 40, 42, 45, 49, 60, 62, 66, 69, 71, 72, 78, 81, 82, 84, 88, 91, 94, 98, 100, 101, 104, 106, 108, 110, 114, 115, 117, 120, 126, 127, 130, 134, 138, 141, 144, 145, 147, 156, 159, 162, 168, 169, 174, 175, 176, 178, 179, 180, 181, 184, 186, 187] [5, 8, 10, 12, 16, 18, 19, 20, 21, 27, 28, 29, 32, 33, 34, 36, 37, 38, 41, 44, 47, 48, 50, 51, 52, 53, 54, 58, 61, 63, 64, 67, 68, 70, 73, 76, 77, 79, 83, 85, 86, 89, 96, 97, 102, 105, 111, 112, 113, 122, 123, 128, 129, 131, 136, 137, 139, 142, 143, 150, 151, 154, 155, 157, 160, 164, 165, 167, 172, 173, 182, 183, 185, 189] [4, 6, 7, 35, 43, 55, 65, 74, 80, 87, 90, 95, 99, 107, 118, 121, 124, 135, 152, 153, 163, 170, 171, 188]	3	500	[4, 6, 8, 9, 10, 13, 14, 16, 18, 19, 20, 21, 27, 28, 29, 32, 33, 34, 36, 37, 38, 42, 43, 45, 47, 48, 50, 51, 52, 53, 54, 55, 58, 61, 63, 65, 67, 68, 70, 74, 77, 79, 80, 81, 83, 85, 87, 89, 91, 95, 96, 98, 99, 100, 101, 102, 106, 107, 108, 111, 112, 114, 115, 117, 121, 122, 124, 128, 129, 131, 135, 136, 138, 142, 144, 145, 147, 150, 151, 154, 157, 160, 164, 165, 168, 169, 171, 172, 174, 175, 176, 178, 179, 180, 181, 182, 183, 186, 187, 189] [5, 7, 12, 15, 17, 31, 35, 39, 40, 41, 44, 49, 60, 62, 64, 66, 69, 73, 76, 78, 82, 84, 86, 88, 90, 94, 97, 104, 105, 113, 118, 120, 126, 127, 130, 134, 139, 141, 152, 153, 156, 159, 162, 163, 170, 173, 184, 185, 188] [1, 24, 26, 71, 72, 110, 123, 137, 143, 155, 167]	3	0.438
4941-Bus‡	3468	No enough space to display	ILP solver did not converge	NA	NA	No enough space to display	4	7716

#R:= number of rounds

†The ILP results for 118-bus and for 189-bus systems are sub-optimal. The simulation was stopped after 500 seconds.

‡The ILP solver does not have enough memory space to solve the 4941-bus system.

formulation for the SPP introduced in Section 6.3.2. We use a laptop with an Intel 2.8 GHz Core *i7* processor and 8GB RAM running Ubuntu 12.04 with Linux 3.2 for the simulations.

Our results in Table 6.1 show the optimal PMU placement and the number of rounds obtained both from the ILP solver and from the heuristic algorithm for different bus systems. The 4941-bus system represents the power transmission grid covering much of the western states in the United States as presented in [125]. The system has 4941 buses and 6594 branches. The data set for the bus system was obtained from [126]. The 189-bus system¹ represents Iceland's transmission network. It has 189 buses and 206 branches. All the other bus systems used in the simulation are the standard IEEE bus systems².

The simulation results show that the ILP performs better than the greedy algorithm in terms of finding fewer (optimal) number of rounds for small size networks. However, the execution time for the ILP solver quickly increases as the network size increases. The execution times for the 118-bus and the 189-bus systems are too large that we had to stop the executions after 500 seconds forcing it to return sub-optimal solutions. Similarly, the memory requirement for the 4941-bus system is too large that the ILP solver could not solve it even sub-optimally using the machine we used for the simulation. Although the greedy approach does not find the optimal patching plan even for the small size networks, it finds a sub-optimal solution much faster. It also solves the 4941-bus system and finds a total number of rounds equal to only 4 within 7991 seconds. One can only imagine how slow an ILP solver can be to solve this problem even if the machine had enough memory size.

It is important to remember that a patching plan obtained using either of the methods can be re-used only if the network setting remains static. If there is change either in the PMU placement or in the connectivity among the buses, a utility needs to re-compute the patching plan for the new setting.

6.7 Conclusion

We have studied the PMU patching problem that arises when a utility wants to maintain system observability while applying software patches to PMUs. We have used set theoretic formulation to model the problem as an instance of sensor patching problem, which we have shown to be NP-complete. We have proved that finding an optimal solution to the problem is equivalent to maximizing a submodular set function and proposed a heuristic algorithm that finds a sub-optimal solution. We have also formulated the problem as a BILP problem and solved it using an ILP solver. A comparison of the performance of the ILP solver and the greedy heuristic is also presented. Moreover, we have shown an interesting case, when the power grid has a radial structure, for which we have devised a polynomial-time algorithm that finds an optimal patching plan that requires only two rounds.

¹<http://www.maths.ed.ac.uk/optenergy/NetworkData/>

²<http://www2.ee.washington.edu/research/pstca/>

While studying the PMU patching problem, we have made simplifying assumptions in that we did not consider observability rules that exploit the presence of zero-injection buses as well as conventional P-Q measurements. As future works, we plan to study the effect of these observability rules on the PMU patching problem. We also plan to consider PMUs with limited channel capacity that measure only the nodal voltage and current phasors. Taking this assumption on a PMU's capacity is also interesting for the well-studied PMU placement problem.

7 Conclusions

In this thesis we studied cybersecurity issues and counter measures in active power distribution networks. The presence of long-lived heterogeneous devices with diversified computing power, support for applications with different latency requirements and a complex unprotected information and communication infrastructure demands for fine-grained security solutions that are tailored to the needs of specific applications in the grid.

We started by performing a threat analysis of a typical active distribution network. We cover threats that could emanate from malicious insiders and outsiders. We then proposed cybersecurity solutions and best practices to counter the threats. Our solutions provide protection against malicious agents who might try to exploit an emergency situation that creates an islanded communication zone to install rogue devices. We also build a secure communication network for the EPFL-campus smart grid pilot using the proposed solutions as guidelines.


Furthermore, we studied the security aspects of two OAM protocols for MPLS-TP, a technology that is envisioned to be used for long-distance inter-domain communication in smart grid. We focus on BFD and PSC protocols, two protocols that are responsible for monitoring the state of the network and for facilitating protection switching when a fault occurs. Following a literature review that shows lack of a unified security guidelines for these protocols, we built a testbed to study if one of the major network device manufacturers provides appropriate security solutions for these protocols. Our findings revealed there is no support for source authentication in both protocols. This allowed us to carry out several spoofing attacks with severe consequences on the network's availability.

In order to identify a suitable multicast source authentication scheme for grid monitoring systems, we studied existing schemes and experimentally compared a selected set of schemes on the EPFL-campus smart grid pilot. Our findings show that an ECDSA implementation that uses pre-generated tokens for signature generation performs better than other candidates. This contradicts the widely held belief that asymmetric cryptography is too computationally expensive to be applicable for real time applications. Two factors played a role in reaching at this conclusion: (1) the not-so-high message sending rate of a PMU provides enough CPU

Chapter 7. Conclusions

time that can be used to pre-generate tokens, (2) a key length that guarantees an intermediate security level is sufficient because authentication for real-time applications is a short-term issue.

The final part of the thesis dealt with a software patch planning problem for PMUs in a smart grid. Given a software patch that requires rebooting, we wanted to find a partitioning of the set of the deployed PMUs into as few subsets as possible such that all the PMUs in one subset can be patched at a time while all the PMUs in the other subsets provide full observability of the grid. We have modelled the problem as an instance of a sensor patching problem and have proved it to be NP-complete. Further, we have formulated the problem as a binary integer linear programming (BILP) problem and used an ILP solver to find a patching plan for relatively small-size networks. Moreover, we have proved that finding an optimal solution to the problem is equivalent to maximizing a submodular set function and we obtained an approximate solution using a heuristic algorithm based on a greedy approach. The results show that whereas the ILP solver does not converge for large-size networks, our heuristic algorithm finds a plan takes only small number of rounds even for very large networks (e.g., only 4 rounds for a 4941 bus system that deploys 3468 PMUs). For a special case of the problem where the grid is a tree, we have provided a polynomial-time algorithm that finds an optimal plan that patches all the PMUs in only two rounds. One way to extend this work is to consider different sets of constraints on the observability rules and to take measurements from non-PMU devices into consideration.

- 
- [1] US Congress, “America recovery and reinvestment act,” Feb. 2009.
- [2] J. Hull, H. Khurana, T. Markham, and K. Staggs, “Staying in control: Cybersecurity and the modern electric grid,” *Power and Energy Magazine, IEEE*, vol. 10, no. 1, pp. 41–48, 2012.
- [3] Y. Mo, T.-H. Kim, K. Brancik, D. Dickinson, H. Lee, A. Perrig, and B. Sinopoli, “Cyber-physical security of a smart grid infrastructure,” *Proceedings of the IEEE*, vol. 100, no. 1, pp. 195–209, 2012.
- [4] S. Fuloria and R. Anderson, “Towards a security architecture for substations,” in *Innovative Smart Grid Technologies (ISGT Europe), 2011 2nd IEEE PES International Conference and Exhibition on*, Dec 2011, pp. 1–6.
- [5] S. Sridhar, A. Hahn, and M. Govindarasu, “Cyber-physical system security for the electric power grid,” *Proceedings of the IEEE*, vol. 100, no. 1, pp. 210–224, 2012.
- [6] ENISA, “Smart Grid Security: Annex II. Security aspects of the smart grid,” 2012.
- [7] E. Smith, S. Corzine, D. Racey, P. Dunne, C. Hassett, and J. Weiss, “Going beyond cybersecurity compliance: What power and utility companies really need to consider,” *IEEE Power and Energy Magazine*, vol. 14, no. 5, pp. 48–56, Sept 2016.
- [8] D. Salmon, M. Zeller, A. Guzman, V. Mynam, and M. Donolo, “Mitigating the aurora vulnerability with existing technology,” in *36th annual western protection relay conference*, 2009.
- [9] S. Karnouskos, “Stuxnet worm impact on industrial cyber-physical system security,” in *IECON 2011 - 37th Annual Conference on IEEE Industrial Electronics Society*, Nov 2011, pp. 4490–4494.
- [10] K. Poulsen, “Slammer worm crashed Ohio nuke plant network,” 2003. [Online]. Available: <http://www.securityfocus.com/news/6767>
- [11] DHS, “Alert (ICS-ALERT-14-281-01E) Ongoing Sophisticated Malware Campaign Compromising ICS (Update E),” 2014. [Online]. Available: <https://ics-cert.us-cert.gov/alerts/ICS-ALERT-14-281-01B>

Bibliography

- [12] F-Secure, "Havex Hunts For ICS/SCADA Systems," 2014. [Online]. Available: <https://www.f-secure.com/weblog/archives/00002718.html>
- [13] S. Gorman, "Electricity Grid in U.S. Penetrated By Spies," 2009. [Online]. Available: <http://www.wsj.com/articles/SB123914805204099085>
- [14] North American Electric Reliability Corporation, "Reliability Standards for the Bulk Electric Systems of North America, NERC CIP standards CIP001 - CIP 011," July 2013.
- [15] K. C. Budka, J. G. Deshpande, and M. Thottan, *Communication Networks for Smart Grids: Making Smart Grid Real*. Springer Publishing Company, Incorporated, 2014.
- [16] United States, *Energy Independence and Security Act (EISA) of 2007*. US Government Printing Office, 2007.
- [17] NIST , "Guidelines for Smart Grid Cyber Security," http://www.nist.gov/smartgrid/upload/nistir-7628_total.pdf, Sept 2010.
- [18] IEC TC57, "IEC 62351 - Power systems management and associated information exchange - Data and communications security," Geneva, Switzerland, 2013.
- [19] F Hohlbaum, M. Braendle, and F Alvarez, "Cyber Security Practical considerations for implementing IEC 62351," ABB, Switzerland, 2010.
- [20] International Society for Automation, "ISA99, Industrial Automation and Control Systems Security." [Online]. Available: <https://www.isa.org/isa99/>
- [21] IEEE, "1686-2013 - IEEE Standard for Intelligent Electronic Devices Cyber Security Capabilities," 2013. [Online]. Available: <https://standards.ieee.org/findstds/standard/1686-2013.html>
- [22] NIST , "NIST Special Publication 800-82 Revision 2: Guide to Industrial Control Systems (ICS) Security," May 2015. [Online]. Available: http://csrc.nist.gov/publications/drafts/800-82r2/sp800_82_r2_second_draft.pdf
- [23] IEEE, "C37.240-2014 - IEEE Standard Cybersecurity Requirements for Substation Automation, Protection, and Control Systems," 2014. [Online]. Available: <https://standards.ieee.org/findstds/standard/C37.240-2014.html>
- [24] —, "1711-2010 - IEEE Trial-Use Standard for a Cryptographic Protocol for Cyber Security of Substation Serial Links," 2014. [Online]. Available: <https://standards.ieee.org/findstds/standard/1711-2010.html>
- [25] IEC TC57, "IEC TR 62210:2003 - Power system control and associated communications - Data and communication security," 2003.

-
- [26] NIST, “NIST Special Publication 1108R2 - NIST Framework and Roadmap for Smart Grid Interoperability Standards, Release 2.0,” Feb. 2012. [Online]. Available: https://www.nist.gov/sites/default/files/documents/smartgrid/NIST_Framework_Release_2-0_corr.pdf
- [27] CNN, “Sources: Staged cyber attack reveals vulnerability in power grid,” Sept.. 2007. [Online]. Available: <http://edition.cnn.com/2007/US/09/26/power.at.risk/>
- [28] Y. Liu, P. Ning, and M. K. Reiter, “False data injection attacks against state estimation in electric power grids,” *ACM Trans. Inf. Syst. Secur.*, vol. 14, no. 1, pp. 13:1–13:33, Jun. 2011.
- [29] O. Kosut, L. Jia, R. J. Thomas, and L. Tong, “Malicious data attacks on the smart grid,” *IEEE Transactions on Smart Grid*, vol. 2, no. 4, pp. 645–658, Dec 2011.
- [30] F. Pasqualetti, F. Dörfler, and F. Bullo, “Attack detection and identification in cyber-physical systems,” *IEEE Transactions on Automatic Control*, vol. 58, no. 11, pp. 2715–2729, Nov 2013.
- [31] G. Dan and H. Sandberg, “Stealth attacks and protection schemes for state estimators in power systems,” in *Smart Grid Communications (SmartGridComm), 2010 First IEEE International Conference on*, Oct 2010, pp. 214–219.
- [32] H. Sandberg, A. Teixeira, and K. H. Johansson, “On security indices for state estimators in power networks,” in *First Workshop on Secure Control Systems (SCS), Stockholm, 2010*, 2010.
- [33] L. Xie, Y. Mo, and B. Sinopoli, “False data injection attacks in electricity markets,” in *Smart Grid Communications (SmartGridComm), 2010 First IEEE International Conference on*, Oct 2010, pp. 226–231.
- [34] J. Bhatti and T. Humphreys, “Hostile control of ships via false gps signals: Demonstration and detection,” *submitted to Navigation, in review*, 2015.
- [35] S. Gong, Z. Zhang, M. Trinkle, A. D. Dimitrovski, and H. Li, “Gps spoofing based time stamp attack on real time wide area monitoring in smart grid,” in *Smart Grid Communications (SmartGridComm), 2012 IEEE Third International Conference on*, Nov 2012, pp. 300–305.
- [36] Q. Yang, J. Barria, and T. Green, “Communication infrastructures for distributed control of power distribution networks,” *IEEE Transactions on Industrial Informatics*, vol. 7, no. 2, pp. 316–327, 2011.
- [37] A. Borghetti, C. Nucci, M. Paolone, G. Ciappi, and A. Solari, “Synchronized phasors monitoring during the islanding maneuver of an active distribution network,” *IEEE Transactions on Smart Grid*, vol. 2, no. 1, pp. 82–91, March 2011.

Bibliography

- [38] H. Laaksonen and K. Kauhaniemi, "Synchronized re-connection of island operated lv microgrid back to utility grid," in *Innovative Smart Grid Technologies Conference Europe (ISGT Europe), 2010 IEEE PES*, 2010, pp. 1–8.
- [39] NIST, "Discussion Draft of the Preliminary Cybersecurity Framework," http://www.nist.gov/itl/upload/discussion-draft_preliminary-cybersecurity-framework-082813.pdf, Aug. 2013.
- [40] Ralph Langner, "The RIPE Framework: A Process-Driven Approach towards Effective and Sustainable Industrial Control System Security," <http://www.langner.com/en/wp-content/uploads/2013/09/The-RIPE-Framework.pdf>, Sept. 2013.
- [41] F. Aloula, A.-A. A. R., R. Al-Dalkya, M. Al-Mardinia, and W. El-Hajjb, "Smart grid security: Threats, vulnerabilities and solutions," *Intelnational Journal of Smart Grid and Clean Energy*, vol. 1, no. 1, 2012.
- [42] E. Wang, Y. Ye, X. Xu, S. Yiu, L. C. K. Hui, and K. Chow, "Security issues and challenges for cyber physical system," in *Green Computing and Communications (GreenCom), 2010 IEEE/ACM Int'l Conference on Cyber, Physical and Social Computing (CPSCom)*, 2010, pp. 733–738.
- [43] D. Wei, Y. Lu, M. Jafari, P. Skare, and K. Rohde, "Protecting smart grid automation systems against cyberattacks," *IEEE Transactions on Smart Grid*, vol. 2, no. 4, pp. 782–795, 2011.
- [44] A. Metke and R. Ekl, "Security technology for smart grid networks," *IEEE Transactions on Smart Grid*, vol. 1, no. 1, pp. 99–107, june 2010.
- [45] W. Wang and Z. Lu, "Survey cyber security in the smart grid: Survey and challenges," *Comput. Netw.*, vol. 57, no. 5, pp. 1344–1371, Apr 2013.
- [46] Y. Yan, Y. Qian, H. Sharif, and D. Tipper, "A survey on cyber security for smart grid communications," *IEEE Communications Surveys Tutorials*, vol. 14, no. 4, pp. 998–1010, 2012.
- [47] C.-C. Liu, A. Stefanov, J. Hong, and P. Panciatici, "Intruders in the grid," *IEEE Power and Energy Magazine*, vol. 10, no. 1, pp. 58–66, 2012.
- [48] T. Chen and S. Abu-Nimeh, "Lessons from stuxnet," *Computer*, vol. 44, no. 4, pp. 91–93, 2011.
- [49] T. Baumeister, "Adapting pki for the smart grid," in *Smart Grid Communications (Smart-GridComm), 2011 IEEE International Conference on*, Oct. 2011, pp. 249–254.
- [50] D. Grawrock, *Dynamics of a Trusted Platform: A Building Block Approach*, 1st ed. Intel Press, 2009.
- [51] Sunil Gupta, "Logging and Monitoring to Detect Network Intrusions and Compliance Violations in the Environment," SANS Institute Reading Room, Jul. 2012.

- [52] Cisco, "Understanding MPLS-TP and Its Benefits," 2009. [Online]. Available: http://www.cisco.com/en/US/technologies/tk436/tk428/white_paper_c11-562013.pdf
- [53] CIGRE, "Line and System Protection using Digital Circuit and Packet Communications," 2012. [Online]. Available: <http://www.e-cigre.org/Order/select.asp?ID=15831>
- [54] A. Durai and V. Varakantam, "Building Smart Grid Core Networks," IEEE Smart Grid Newsletter, October 2012.
- [55] L. Fang, B. Niven-Jenkins, S. Mansfield, and R. Graveman, "MPLS Transport Profile (MPLS-TP) Security Framework," IETF, RFC 6941 (Informational), Internet Engineering Task Force, Apr. 2013. [Online]. Available: <http://www.ietf.org/rfc/rfc6941.txt>
- [56] Y. Weingarten, S. Bryant, E. Osborne, N. Sprecher, and A. Fulignoli, "MPLS Transport Profile (MPLS-TP) Linear Protection," IETF, RFC 6378 (Proposed Standard), Internet Engineering Task Force, Oct. 2011, updated by RFCs 7214, 7271, 7324. [Online]. Available: <http://www.ietf.org/rfc/rfc6378.txt>
- [57] N. Sprecher and A. Farrel, "MPLS Transport Profile (MPLS-TP) Survivability Framework," IETF, RFC 6372 (Informational), Internet Engineering Task Force, Sep. 2011. [Online]. Available: <http://www.ietf.org/rfc/rfc6372.txt>
- [58] L. Fang, "Security Framework for MPLS and GMPLS Networks," IETF, RFC 5920 (Informational), Internet Engineering Task Force, Jul. 2010. [Online]. Available: <http://www.ietf.org/rfc/rfc5920.txt>
- [59] T. Nadeau and C. Pignataro, "Pseudowire Virtual Circuit Connectivity Verification (VCCV): A Control Channel for Pseudowires," IETF, RFC 5085 (Proposed Standard), Internet Engineering Task Force, Dec. 2007, updated by RFC 5586. [Online]. Available: <http://www.ietf.org/rfc/rfc5085.txt>
- [60] J. W. Lockwood, N. McKeown, G. Watson, G. Gibb, P. Hartke, J. Naous, R. Raghuraman, and J. Luo, "NetFPGA—An Open Platform for Gigabit-Rate Network Switching and Routing," in *Microelectronic Systems Education, 2007. MSE '07. IEEE International Conference on*, June 2007, pp. 160–161.
- [61] L. Andersson, L. Berger, L. Fang, N. Bitar, and E. Gray, "MPLS Transport Profile (MPLS-TP) Control Plane Framework," IETF, RFC 6373 (Informational), Internet Engineering Task Force, Sep. 2011. [Online]. Available: <http://www.ietf.org/rfc/rfc6373.txt>
- [62] Juniper Networks, "MPLS Transport Profile (MPLS-TP): A Set of Enhancements to the Rich MPLS Toolkit," 2011. [Online]. Available: <http://opti500.cian-erc.org/opti500/pdf/sm/mpls-tp%20Juniper.pdf>
- [63] D. Katz and D. Ward, "Bidirectional Forwarding Detection (BFD)," IETF, RFC 5880 (Proposed Standard), Internet Engineering Task Force, Jun. 2010. [Online]. Available: <http://www.ietf.org/rfc/rfc5880.txt>

Bibliography

- [64] Cisco Systems, Inc., “Cisco CSR 1000V Series Cloud Services Router Overview,” <http://www.cisco.com/c/en/us/td/docs/routers/csr1000/software/configuration/csr1000Vswcfg/csroverview.html>, Jul. 2012.
- [65] —, “IP Routing BFD Configuration Guide, Cisco IOS XE Release 3S,” http://www.cisco.com/c/en/us/td/docs/ios-xml/ios/iproute_bfd/configuration/xe-3s/irb-xe-3s-book.html, 2013.
- [66] Philippe Biondi, “Scapy Project,” <http://www.secdev.org/projects/scapy/>, Feb. 2011.
- [67] IEC TC/SC 57, “IEC 61850 Communication networks and systems for power utility automation - Part 8-1: Specific communication service mapping (SCSM) - Mappings to MMS and to ISO/IEC 8802-3,” Jun. 2011.
- [68] NASPI. [Online]. Available: <https://www.naspi.org/home>
- [69] “IEEE Standard Communication Delivery Time Performance Requirements for Electric Power Substation Automation,” *IEEE Std 1646-2004*, pp. 1–24, 2005.
- [70] M. Seewald, “Building an architecture based on IP-Multicast for large phasor measurement unit (PMU) networks,” in *Innovative Smart Grid Technologies (ISGT), 2013 IEEE PES*, Feb 2013.
- [71] P. Myrda, J. Taft, and P. Donner, “Recommended Approach to a NASPInet Architecture,” in *System Science (HICSS), 2012 45th Hawaii International Conference on*, Jan 2012.
- [72] D. Boneh, G. Durfee, and M. K. Franklin, “Lower bounds for multicast message authentication,” in *Proceedings of the International Conference on the Theory and Application of Cryptographic Techniques: Advances in Cryptology*, ser. EUROCRYPT ’01. London, UK, UK: Springer-Verlag, 2001, pp. 437–452. [Online]. Available: <http://dl.acm.org/citation.cfm?id=647086.715684>
- [73] M. Luk, A. Perrig, and B. Whillock, “Seven Cardinal Properties of Sensor Network Broadcast Authentication,” in *Proceedings of the Fourth ACM Workshop on Security of Ad Hoc and Sensor Networks*, ser. SASN ’06. New York, NY, USA: ACM, 2006.
- [74] Y. Challal, H. Bettahar, and A. Bouabdallah, “A taxonomy of multicast data origin authentication: Issues and solutions,” *Communications Surveys Tutorials, IEEE*, vol. 6, no. 3, Third 2004.
- [75] IEC TC57, “Communication networks and systems for power utility automation - Part 90-5: Use of IEC 61850 to transmit synchrophasor information according to IEEE C37.118,” 2012.
- [76] J. Zhang and C. Gunter, “Application-aware secure multicast for power grid communications,” in *Smart Grid Communications (SmartGridComm), 2010 First IEEE International Conference on*, Oct 2010.

-
- [77] R. Canetti, J. Garay, G. Itkis, D. Micciancio, M. Naor, and B. Pinkas, "Multicast security: a taxonomy and some efficient constructions," in *INFOCOM '99. Eighteenth Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings. IEEE*, vol. 2, Mar 1999, pp. 708–716 vol.2.
- [78] A. Perrig, R. Canetti, J. Tygar, and D. Song, "Efficient authentication and signing of multicast streams over lossy channels," in *Security and Privacy, 2000. S P 2000. Proceedings. 2000 IEEE Symposium on*, 2000.
- [79] Q. Wang, H. Khurana, Y. Huang, and K. Nahrstedt, "Time Valid One-Time Signature for Time-Critical Multicast Data Authentication," in *INFOCOM 2009, IEEE*, April 2009, pp. 1233–1241.
- [80] S. Vanstone, "Responses to NIST's Proposal," *Communications of the ACM*, 35, American National Standards Institute, July 1992.
- [81] R. L. Rivest, A. Shamir, and L. Adleman, "A Method for Obtaining Digital Signatures and Public-key Cryptosystems," *Commun. ACM*, vol. 21, no. 2, Feb. 1978.
- [82] M. Pignati, M. Popovic, S. Barreto, R. Cherkaoui, G. Dario Flores, J.-Y. Le Boudec, M. Mohiuddin, M. Paolone, P. Romano, S. Sarri, T. Tesfay, D.-C. Tomozei, and L. Zanni, "Real-time state estimation of the EPFL-campus medium-voltage grid by using PMUs," in *Innovative Smart Grid Technologies Conference, 2015 IEEE Power Energy Society*, Feb 2015.
- [83] NIST, "Federal Information Processing Standards Publication 186-4, Digital Signature Standard (DSS)," <http://nvlpubs.nist.gov/nistpubs/FIPS/NIST.FIPS.186-4.pdf>, 2013.
- [84] A. A. Yavuz, "An efficient real-time broadcast authentication scheme for command and control messages," *Trans. Info. For. Sec.*, vol. 9, no. 10, Oct. 2014.
- [85] L. Lamport, "Constructing Digital Signatures from a One Way Function," Technical Report, SRI-CSL-98, SRI Intl. Computer Science Laboratory, October 1979.
- [86] M. O. Rabin, "Digitized Signatures and Public-key Functions as Intractable as Factorization," Cambridge, MA, USA, Tech. Rep., 1979.
- [87] A. Perrig, "The BiBa One-time Signature and Broadcast Authentication Protocol," in *Proceedings of the 8th ACM Conference on Computer and Communications Security*, ser. CCS '01. NY, USA: ACM, 2001.
- [88] L. Reyzin and N. Reyzin, "Better than BiBa: Short One-time Signatures with Fast Signing and Verifying," in *In Seventh Australasian Conference on Information Security and Privacy (ACISP 2002)*, 2002.
- [89] Q. Li and G. Cao, "Multicast Authentication in the Smart Grid With One-Time Signature," *Smart Grid, IEEE Transactions on*, Dec 2011.

Bibliography

- [90] Y. W. Law, Z. Gong, T. Luo, S. Marusic, and M. Palaniswami, "Comparative Study of Multicast Authentication Schemes with Application to Wide-area Measurement System," in *Proceedings of the 8th ACM SIGSAC Symposium on Information, Computer and Communications Security*, ser. ASIA CCS '13. New York, NY, USA: ACM, 2013.
- [91] X. Lu, W. Wang, and J. Ma, "Authentication and Integrity in the Smart Grid: An Empirical Study in Substation Automation Systems," *International Journal of Distributed Sensor Networks*, 2012.
- [92] A. Perrig, R. Szewczyk, J. D. Tygar, V. Wen, and D. E. Culler, "SPINS: Security Protocols for Sensor Networks," *Wirel. Netw.*, Sep 2002.
- [93] D. Liu and P. Ning, "Multilevel μ -TESLA: Broadcast Authentication for Distributed Sensor Networks," *ACM Trans. Embed. Comput. Syst.*, vol. 3, no. 4, Nov. 2004.
- [94] C. Tartary, H. Wang, and S. Ling, "Authentication of digital streams," *Information Theory, IEEE Transactions on*, vol. 57, no. 9, Sept 2011.
- [95] D. Naccache, D. M'Raihi, S. Vaudenay, and D. Rphaeli, *Advances in Cryptology - EUROCRYPT'94: Workshop on the Theory and Application of Cryptographic Techniques Perugia, Italy, May, 1994 Proceedings*. Berlin, Heidelberg: Springer Berlin Heidelberg, 1995, ch. Can D.S.A. be improved? Complexity trade-offs with the digital signature standard.
- [96] "IEEE Standard for Synchrophasor Data Transfer for Power Systems," *IEEE Std C37.118.2-2011*, Dec 2011.
- [97] OpenSSL. [Online]. Available: <http://www.openssl.org>
- [98] G. Liang, S. R. Weller, J. Zhao, F. Luo, and Z. Y. Dong, "The 2015 ukraine blackout: Implications for false data injection attacks," *IEEE Transactions on Power Systems*, vol. PP, no. 99, pp. 1–1, 2016.
- [99] N. Smart, S. Babbage, D. Catalano, C. Cid, B. de Weger, O. Dunkelman, C. Christian Gehrman, L. Granboulan, T. Guneyssu, and M. Ward, "ECRYPT II Yearly Report on Algorithms and Keysizes," European Network of Excellence in Cryptology (ECRYPT II), Sep. 2012.
- [100] R. Wang, W. Du, X. Liu, and P. Ning, "ShortPK: A Short-term Public Key Scheme for Broadcast Authentication in Sensor Networks," *ACM Trans. Sen. Netw.*, vol. 6, no. 1, Jan 2010.
- [101] H. Eberle, N. Gura, S. C. Shantz, V. Gupta, L. Rarick, and S. Sundaram, "A public-key cryptographic processor for RSA and ECC," in *Application-Specific Systems, Architectures and Processors, 2004. Proceedings. 15th IEEE International Conference on*, Sept 2004.
- [102] "IEEE Guide for Phasor Data Concentrator Requirements for Power System Protection, Control, and Monitoring," *IEEE Std C37.244-2013*, May 2013.

- [103] A. Gomez-Exposito, A. Abur, A. de la Villa Jaen, and C. Gomez-Quiles, "A multilevel state estimation paradigm for smart grids," *Proceedings of the IEEE*, June 2011.
- [104] S. Tom, D. Christiansen, and D. Berrett, "Recommended practice for patch management of control systems," *DHS control system security program (CSSP) Recommended Practice*, 2008.
- [105] K. Knorr, *Securing Critical Infrastructures and Critical Control Systems: Approaches for Threat Protection*. IGI Global, 2013, ch. Patching our Critical Infrastructure - Towards an Efficient Patch and Update Management for Industrial Control Systems.
- [106] M. Souppaya and K. Scarfone, "Guide to enterprise patch management technologies," *NIST Special Publication*, vol. 800, p. 40, 2013.
- [107] E. Abiri, F. Rashidi, and T. Niknam, "An optimal pmu placement method for power system observability under various contingencies," *International Transactions on Electrical Energy Systems*, vol. 25, no. 4, pp. 589–606, 2015.
- [108] T. L. Baldwin, L. Mili, M. B. Boisen, and R. Adapa, "Power system observability with minimal phasor measurement placement," *IEEE Transactions on Power Systems*, vol. 8, May 1993.
- [109] F. Aminifar, A. Khodaei, M. Fotuhi-Firuzabad, and M. Shahidehpour, "Contingency-constrained pmu placement in power networks," *IEEE Transactions on Power Systems*, vol. 25, Feb 2010.
- [110] J. G. Philip and T. Jain, "Optimal placement of pmus for power system observability with increased redundancy," in *2015 Conference on Power, Control, Communication and Computational Technologies for Sustainable Growth (PCCCTSG)*, Dec 2015.
- [111] M. Esmaili, K. Gharani, and H. A. Shayanfar, "Redundant observability pmu placement in the presence of flow measurements considering contingencies," *IEEE Transactions on Power Systems*, Nov 2013.
- [112] A. Enshae, R. A. Hooshmand, and F. H. Fesharaki, "A new method for optimal placement of phasor measurement units to maintain full network observability under various contingencies," *Electric Power Systems Research*, vol. 89, 2012.
- [113] T. Kerdchuen and W. Ongsakul, "Optimal pmu placement by stochastic simulated annealing for power system state estimation," *GMSARN International Journal*, vol. 2, 2008.
- [114] R. F. Nuqui and A. G. Phadke, "Phasor measurement unit placement techniques for complete and incomplete observability," *IEEE Transactions on Power Delivery*, vol. 20, Oct 2005.
- [115] D. Gyllstrom, E. Rosensweig, and J. Kurose, "On the impact of pmu placement on observability and cross-validation," in *Proceedings of the 3rd International Conference*

Bibliography

- on Future Energy Systems: Where Energy, Computing and Communication Meet*, ser. e-Energy '12. New York, NY, USA: ACM, 2012.
- [116] J. Peng, Y. Sun, and H. Wang, "Optimal pmu placement for full network observability using tabu search algorithm," *International Journal of Electrical Power and Energy Systems*, vol. 28, 2006.
- [117] E. Abiri and F. Rashidi, "Optimal phasor measurement units placement to maintain network observability using a novel binary particle swarm optimization and fuzzy system," *J. Intell. Fuzzy Syst.*, vol. 28, Jan. 2015.
- [118] N. Xia, H. Gooi, S. Chen, and M. Wang, "Redundancy based pmu placement in state estimation," *Sustainable Energy, Grids and Networks*, vol. 2, pp. 23 – 31, 2015.
- [119] F. C. Schweppe and J. Wildes, "Power system static-state estimation, part i: Exact model," *IEEE Transactions on Power Apparatus and Systems*, vol. PAS-89, Jan 1970.
- [120] A. Monticelli and A. Garcia, "Reliable bad data processing for real-time state estimation," *Power Apparatus and Systems, IEEE Transactions on*, vol. PAS-102, no. 5, pp. 1126 –1139, May 1983.
- [121] A. Mahari and H. Seyedi, "Optimal pmu placement for power system observability using bica, considering measurement redundancy," *Electric Power Systems Research*, vol. 103, 2013.
- [122] A. Bretto, *Hypergraph Colorings*. Heidelberg: Springer International Publishing, 2013, pp. 43–56.
- [123] M. Campêlo, R. Corrêa, and Y. Frota, "Cliques, holes and the vertex coloring polytope," *Information Processing Letters*, vol. 89, no. 4, 2004.
- [124] Michel Berkelaar. [Online]. Available: <http://web.mit.edu/lpsolve/doc/>
- [125] D. J. Watts and S. H. Strogatz, "Collective dynamics of 'small-world' networks," *Nature*, vol. 393, no. 6684, pp. 440–442, 1998.
- [126] M. Newman, "Western states power grid data set," *Retrieved on October 12, 2016*. [Online]. Available: <http://www-personal.umich.edu/~mejn/netdata/power.zip>

Publications

Published

1. Teklemariam Tsegay Tesfay, J.-P. Hubaux, J.-Y. Le Boudec, and P. Oechslin, "Cyber-secure communication architecture for active power distribution networks," in Proceedings of the 29th Annual ACM Symposium on Applied Computing, ser. SAC '14. New York, NY, USA: ACM, 2014.
2. W. K. Chai, N. Wang, K. V. Katsaros, G. Kamel, G. Pavlou, S. Melis, M. Hoefling, B. Vieira, P. Romano, S. Sarri, T. T. Tesfay, B. Yang, F. Heimgaertner, M. Pignati, M. Paolone, M. Menth, E. Poll, M. Mampaey, H. H. I. Bontius, and C. Develder, "An information-centric communication infrastructure for real-time state estimation of active distribution networks," IEEE Transactions on Smart Grid, vol. 6, no. 4, July 2015.
3. U. Jayasinghe, S. Barreto, M. Popovic, T. T. Tesfay, and J.-Y. Le Boudec, "Security vulnerabilities of the Cisco IOS implementation of the MPLS transport profile," in Proceedings of the 2nd Workshop on Smart Energy Grid Security, Ser. SEGS '14. New York, NY, USA: ACM, 2014.
4. M. Pignati, M. Popovic, S. Barreto, R. Cherkaoui, G. D. Flores, J.-Y. Le Boudec, M. Mohiuddin, M. Paolone, P. Romano, S. Sarri, T. Tesfay, D. C. Tomozei, and L. Zanni, "Real-time state estimation of the EPFL-campus medium-voltage grid by using PMUs," in Innovative Smart Grid Technologies Conference (ISGT), 2015 IEEE Power Energy Society, Feb 2015.
5. M. Hoefling, F. Heimgaertner, D. Fuchs, M. Menth, P. Romano, T. Tesfay, M. Paolone, J. Adolph, and V. Gronas, "Integration of IEEE C37.118 and publish/subscribe communication," in 2015 IEEE International Conference on Communications (ICC), June 2015.
6. Teklemariam Tsegay Tesfay and Jean-Yves Le Boudec, "Experimental Comparison of Multicast Authentication for Wide Area Monitoring Systems", IEEE Transactions on Smart Grid, *Status: Accepted*, 2017. [Online]. Available: <http://infoscience.epfl.ch/record/216923>

Under Review

1. Teklemariam Tsegay Tesfay and Jean-Yves Le Boudec and Ola Svensson, "Optimal Software Patching Plan for PMUs," Submitted to IEEE Transactions on Smart Grid, *Status: under review*, 2017. Available: <https://infoscience.epfl.ch/record/222733/>



Teklemariam Tesfay |

EPFL, LCA2 INF 011, Station 14 – Lausanne, CH-1015, Switzerland
+41 78 635 83 80 • +41 21 693 12 66 • tech.tesfay@gmail.com
<http://people.epfl.ch/tech.tesfay>

*"Amateurs hack systems, professionals hack people." –
Bruce Schneier*

Education

École Polytechnique Fédérale de Lausanne

PhD Candidate

Supervisor: Prof. Jean-Yves Le Boudec

Thesis: Cybersecurity solutions for active power distribution networks

Switzerland

2011–Present

Indian Institute of Technology, Bombay

M.Tech in Computer Science and Engineering

Supervisor: Prof. S. Sudarshan

Thesis: Designing flash-aware indexing algorithms

CGPA: **9.53/10.0**

India

2007–2009

Mekelle Institute of Technology

BSc in Computer Science and Engineering

Supervisor: Prof. M.K. Chandra

Thesis: Handwritten signature recognition and verification using neural networks

CGPA: **4.0/4.0**

Ethiopia

2002–2007

Professional Experience

École Polytechnique Fédérale de Lausanne

Research Assistant

Pursuing my PhD thesis under the supervision of Prof. Jean-Yves Le Boudec. My main research interests are various aspects of smart-grid security with a focus on designing access control and key management schemes for sensing field devices in an active power distribution network. This includes designing and implementation of centralized authentication mechanisms to devices, identifying efficient multicast authentication schemes for synchrophasor data communication and devising algorithms for an optimal software patching plan for phasor measurement units in grid monitoring systems.

Switzerland

2011–Present

Mekelle Institute of Technology

Lecturer

I have taught Design and analysis of algorithms, Artificial intelligence and System modeling and simulation courses to 3rd year and 4th year undergraduate computer science and engineering students at Mekelle Institute of Technology, Ethiopia.

Ethiopia

2010–2011

École Polytechnique Fédérale de Lausanne

Intern

During my one-year internship at EPFL, I studied energy savings and capacity gain in 4G cellular networks with a careful deployment of low power micro-basestations along with macro-basestations. We proposed a multi-class product form queuing model to determine the traffic capacity of cellular networks while taking both the physical and traffic layer specifications of the network into consideration.

Switzerland

2009–2010

Other Professional Activities

Securing EPFL's smart grid communication network – <http://smartgrid.epfl.ch/>: 2013–2016

Successfully put into operation mechanisms to secure the cyber assets that supports PMU-based smart-grid monitoring system. **Teaching Assistant for the following courses at EPFL, Switzerland:**

TCP/IP Networking (MSc level): 2012–2015

Designed several lab work with hands-on exercises on socket programming, TCP congestion control, tunnelling, network security and BGP routing protocol.

Smart Grid technologies (MSc level): 2015

Designed labs with hands-on exercise on cyber-attacks on smart grid communication and security using DTLS.

Performance Evaluation of Computer and Communication Systems (MSc level): 2013–2015

Responsible for the lab problems that involved performance patterns (bottlenecks, congestion collapse), model fitting and forecasting, discrete event simulation and queuing theory.

Information Sciences (BSc level): 2014

Responsible for the lab problems that involved the mathematical foundations for different security protocols such as RSA and Diffie-Hellman key exchange protocol.

Peer Reviews

Reviewer: Elsevier Sustainable Energy, Grids and Networks Journal, since 2014

Reviewer: IEEE Transactions on Industrial Informatics Journal, since 2015

Trainings and Workshops

- Zurich Information Security and Privacy Center (ZISC), ETHZ, Zurich, Switzerland (2014)
- EES-UETP workshop on Cyber-Physical System Security of the Power Grid, KTH, Sweden (2013)
- Physical layer of LTE mobile networks, Vodafone at Dresden University of Technology, Germany (2010)
- MySQL server administration, Sun Micro Systems, Bombay, India (2009)
- Bandwidth management in optical networks, Cisco Systems, Bombay, India (2009)

Skills

Programming languages and tools: C, C++, Python, Perl, Shell, Metasploit, Nmap, Wireshark, Scapy

Operating Systems: Linux (Debian, RHEL, Kali Linux for security-assessments), Mac OS, Windows

Honors and Awards

- Teaching Assistant Award, EPFL, Switzerland (2015)
- Fellowship, EPFL (2011-2017)
- Excellence scholarship for M.Tech at IIT, Bombay, India (2007)
- Gold medalist graduate from Mekelle Institute of Technology, Ethiopia (2007)
- Gold medalist graduate from Mekelle Institute of Technology, Ethiopia (2007)

Languages

English: Fluent

French: Limited working proficiency (B1 Level)

Tigrigna: Native speaker

Amharic: Fluent

