

Routing for Wireless Sensor Networks: From Collection to Event-Triggered Applications

THÈSE N° 8879 (2018)

PRÉSENTÉE LE 21 SEPTEMBRE 2018

À LA FACULTÉ INFORMATIQUE ET COMMUNICATIONS

LABORATOIRE DE MODÉLISATION SYSTÉMIQUE

PROGRAMME DOCTORAL EN INFORMATIQUE ET COMMUNICATIONS

ÉCOLE POLYTECHNIQUE FÉDÉRALE DE LAUSANNE

POUR L'OBTENTION DU GRADE DE DOCTEUR ÈS SCIENCES

PAR

Daniel Camilo ROJAS QUIRÓS

acceptée sur proposition du jury:

Dr R. Bouluc, président du jury
Prof. J.-D. Decotignie, directeur de thèse
Prof. J.-C. Grégoire, rapporteur
Dr A. El-Hoiydi, rapporteur
Prof. R. Guerraoui, rapporteur



ÉCOLE POLYTECHNIQUE
FÉDÉRALE DE LAUSANNE

Suisse
2018

To my wife,
my parents
and my brother.

Acknowledgements

I have wanted to do a Ph.D. since I was 15 years old. Today, 15 years later, I look back feeling that I do not have enough pages to thank all the people that have been part of this journey. This complex path required transitioning through multiple geographies (Costa Rica -> Switzerland, and other countries), domains of knowledge (Physics -> Electrical Engineering -> Communication Systems -> Computer Science) and languages (Spanish -> English -> French, Romanian and German). This project has been a success because I have been surrounded by good friends. Some contributed by providing the right advice at the right time (probably not knowing that it would lead me here), by being constantly available to listen to my concerns, or simply by joining me in celebration. My heartfelt gratitude goes to all of you.

Jean-Do, thank you for believing in me. You supported my crazy ideas and guided me to transform them in this thesis. I treasure your example of commitment, integrity and patience as a guidance in life. I feel I could not have wished for a better Ph.D. supervisor and role model. During more than four years you have invested extensive time and energy in helping me develop the skills that I will carry for the rest of my professional career. Beyond my research, I appreciate that you have always been available to provide valuable advice on my personal and professional life.

Elena, you have been an extraordinary friend and colleague since your first day at CSEM. Thank you for saving my rear end countless times, including (but not limited to) matters of: writing, science, career, housing, family and life. You walked me during my early steps in research, and kept coaching me until today. You have been an invaluable source of advice.

I would like to extend my sincere thanks to John Farserotu and Philippe Dallemagne for always being available for providing advice, and supporting my research and career. You have constantly gone out of your way to help me discover and develop my potential.

I would like to thank my colleagues at CSEM and Division M for giving me a great place to work during my Ph.D.. You transformed my daily routine in a multicultural learning adventure. Special thanks for many interesting discussions go to: Engin, Marc, Marine, Amina, Camilo, Ernesto, Alain Serge, Solange, Leila, Nathalie and Rafa.

Thanks to the colleagues from Division T for a zillion laughs during climbing sessions, lake picnics and dinners. I am especially grateful to Francesca, Dara, Nemanja, Maxime, Kausik and Ivan.

Colleagues from Sector 141, it has been a privilege to be surrounded by people that are both excellent in their domain of expertise, and well rounded persons. You filled my Ph.D. days with interesting technical discussions and joyful conversations about every conceivable topic (from

Acknowledgements

religion in Madagascar and the royal line of succession in Thailand, to Colombian gastronomy). I was lucky to work side-by-side with such a diverse and open-minded group where no topic can be too weird.

Special thanks to Prof. Jean-Charles Grégoire, Dr. Amre El-Hoiydi and Prof. Rachid Guerraoui for being part of my jury, thoroughly revising the manuscript, and providing detailed feedback. Thanks to Dr. Ronan Boulic for presiding the jury. I especially appreciate the efforts of the jury for traveling to Lausanne, from as far as Canada, to evaluate my work.

I would like to thank CSEM and the NanoTera program for contributing extensively to the financing of my Ph.D.. I had the privilege of performing my research as part of project WiseSkin, which provided me with a deep sense of purpose and the prospect of a tangible positive impact in the world. Moreover, I was fortunate to share this exciting journey with an excellent multidisciplinary team including: Dr. Farserotu, Prof. Decotignie, Dr. Daskalaki, Prof. Enz, Dr. Kopta, Prof. Dr. Koch, Dr. Huang, Prof. Lacour and Dr. Michaud.

Big thanks to the reviewers of my papers, who provided valuable feedback and suggested interesting research directions. I am grateful to the Computer Engineering and Networks Laboratory at the ETH Zürich for maintaining the Flocklab testbed, which I have used extensively. I also appreciate the efforts that Carlo Boano and his team at TU Graz have undergone to organize the EWSN Dependability Competition. I had the privilege to participate in this event and I found inspiration in the passion of Carlo's team for pushing our domain forward.

I have been very fortunate to count on multiple supervisors through my scientific path, whose patience and dedication have shaped my career. Special thanks go to: Arturo Ramírez, Giovanni Ansalonni, David Atienza, Jose Araya Pochet, Jorge Romero, Emilio de Biasi, Flávio Garcia, Miguel Algueró and Jordi Sacristan.

I was fortunate to navigate the ups and downs of the academic life with other Ph.D. students and postdocs: Eleonora, Josua, Vladimir, Ana, Allen, Elisa von Brücke, Janina, Johannes, Esteban (el Rookie), Monica (la triple M), Jan, Arttu, Oana, Igor, Bogdan, Bojan and Afsoon. Thanks for filling these years with some of the best memories of my life.

I would like to dedicate a special shout-out to the colleagues at the PV-LAB and the IC-LAB in Microcity: you fully integrated me in the social life of your labs. You became the cornerstone of my integration in Neuchâtel through countless dinners, BBQs, parties and coffee breaks. I was also lucky to count with friends beyond EPFL who constantly gave me a healthy kick out of my research world: Sarah, Oscar, Natalia, Alberto, Marcelo, Chiara, Marina and Patrick.

I could not have remained several years in Switzerland without maintaining also a life in Costa Rica. This has only been possible due to the efforts of my dear friends: Yohel, Diana, Diego and Henry. Thank you for always pulling me out of my routine and for your unconditional friendships that have passed the test of time and distance.

I would like to thank the Stoenescus (Geo, Gaby, Ina, Dan and Livi) for integrating me in your family and for turning Romania into my second home. You have put immense time and effort into taking care of Corina and me, and supporting our crazy plans. Just when I thought that I could not be more blessed with the family that I got in Costa Rica, you guys appeared and doubled my reasons to be grateful with life.

I might not have lasted beyond the first week in Switzerland without the support of the Beretta

family (Julie, Paule, Jen and Pierre). They accompanied me through the difficulties of my settling down in Europe. Their kindness will remain a life example for me.

Having a home to miss is a privilege. I would like to thank my aunts, uncles and cousins in Costa Rica for filling my years in the other side of the Atlantic with countless laughs and adventures, for always keeping me in their thoughts, and for making every visit memorable.

Papi, gracias por sus esfuerzos para cuidarnos, convertirnos en personas de bien y enseñarnos a ser felices. Gracias por despertar una pasión por la historia, la política, la geografía y otras ciencias sociales, la cual con el tiempo se convirtió en mi motivación por viajar y descubrir la vida. Gracias por enseñarme a buscar y disfrutar lo bueno en cada persona.

Rubén, gracias por ser el mejor hermano que uno podría desear. Usted ha llenado mi vida de recuerdos alegres, desde jugar Play, hasta bucear en Thailandia. Gracias por ser mi primer compañero para descubrir el mundo, por sus esfuerzos para vencer la distancia, por cuidarme desde Costa Rica y por siempre apoyar mi loqueras.

Mami, gracias por sacar adelante la familia: la posibilidad de soñar en la vida se la debemos a usted. Gracias por tantos años de trabajo duro y sacrificios que construyeron la vida de ensueño que tenemos ahora. Gracias por enseñarnos a no tener miedo, a luchar, a ser humildes y agradecidos. Gracias por apoyar con todas sus fuerzas mi pasión por la ciencia, por darme su confianza desde pequeño y por un amor tan fuerte que no teme dejarme volar. Gracias por todos sus esfuerzos para superar las distancias.

Having my dear wife Corina next to me during these years has been essential to accomplish this Ph.D.. Thank you for filling my days with love, humor and excitement about our future. Your unconditional support and trust gave me the impulse to take risks and dream with the stars. You shared with me, without hesitation, the load during my most difficult moments. I could not have been able to dream with a better wife and I look forward to spending the rest of our lives together. Today, 15 years later, we begin a new adventure.

Camilo Rojas

Neuchâtel, 2018

Abstract

Wireless Sensor Networks (WSNs) are collections of sensing devices using wireless communication to exchange data. In the past decades, steep advancements in the areas of microelectronics and communication systems have driven an explosive growth in the deployment of WSNs. Novel WSN applications have penetrated multiple areas, from monitoring the structural stability of historic buildings, to tracking animals in order to understand their behavior, or monitoring humans' health.

The need to convey data from increasingly complex applications in a reliable and cost-effective manner translates into stringent performance requirements for the underlying WSNs. In the frame of this thesis, we have focused on developing routing protocols for multi-hop WSNs, that significantly improve their reliability, energy consumption and latency. Acknowledging the need for application-specific trade-offs, we have split our contribution into two parts. Part 1 focuses on collection protocols, catering to applications with high reliability and energy efficiency constraints, while the protocols developed in part 2 are subject to an additional bounded latency constraint.

The two mechanisms introduced in the first part, WiseNE and Rep, enable the use of composite metrics, and thus significantly improve the link estimation accuracy and transmission reliability, at an energy expense far lower than the one achieved in previous proposals. The novel beaconing scheme WiseNE enables the energy-efficient addition of the RSSI (Received Signal Strength Indication) and LQI (Link Quality Indication) metrics to the link quality estimate by decoupling the sampling and exploration periods of each mote. This decoupling allows the use of the Trickle Algorithm, a key driver of protocols' energy efficiency, in conjunction with composite metrics. WiseNE has been applied to the Triangle Metric and validated in an online deployment. The section continues by introducing Rep, a novel sampling mechanism that leverages the packet repetitions already present in low-power preamble-sampling MAC protocols in order to improve the WSN energy consumption by one order of magnitude. WiseNE, Rep and the novel PRSSI (Penalized RSSI, a combination of PRR and RSSI) composite metric have been validated in a real smart city deployment.

Part 2 introduces two mechanisms that were developed in the frame of the WiseSkin project (an initiative aimed at designing highly sensitive artificial skin for human limb prostheses), and are generally applicable to the domain of cyber-physical systems. It starts with Glossy-W, a protocol that leverages the superior energy-latency trade-off of flooding schemes based on concurrent transmissions. Glossy-W ensures the stringent synchronization requirements

Acknowledgements

necessary for robust flooding, irrespective of the number of motes simultaneously reporting an event. Part 2 also introduces SCS (Synchronized Channel Sampling), a novel mechanism capable of reducing the power required for periodic polling, while maintaining the event detection reliability, and enhancing the network coexistence. The testbed experiments performed show that SCS manages to reduce the energy consumption of the state-of-the-art protocol Back-to-Back Robust Flooding by over one third, while maintaining an equivalent reliability, and remaining compatible with simultaneous event detection. SCS' benefits can be extended to the entire family of state-of-the-art protocols relying on concurrent transmissions.

Keywords: Wireless Sensor Networks, WSN, Internet of Things, Routing, Medium Access Control, Link Estimation, Composite Metrics, Concurrent Transmissions, Urban Outdoor Deployments, Cyber-Physical Systems.

Résumé

Les réseaux de capteurs sans fil (WSN) sont constitués d'une myriade de nœuds qui échangent de l'information à l'aide de liaisons sans fil. Au cours des décennies précédentes, de nombreuses avancées technologiques ont permis le développement et la commercialisation d'émetteurs-récepteurs, de capteurs et de microprocesseurs à basse consommation. Grâce à ces progrès, les applications exploitant les WSN se sont multipliées dans de nombreux domaines, depuis la surveillance des structures de bâtiments anciens, en passant par la localisation et le suivi du comportement des animaux, jusqu'à la mesure de paramètres physiologiques du corps humain.

Le besoin d'échanger de l'information de manière fiable et à moindre coût au sein d'applications de plus en plus complexes se traduit par des exigences élevées en terme de performance à l'égard des WSN. Dans le cadre de cette thèse, nous nous sommes penchés sur les protocoles de routage pour les WSN à saut multiples, afin d'en augmenter la fiabilité et d'en diminuer la consommation énergétique et le temps de réponse. Reconnaisant la nécessité de proposer des compromis spécifiques aux applications, nous avons divisé notre travail en deux parties. La première partie se concentre sur les protocoles de collecte de données, destinés à des applications présentant de fortes contraintes de consommation énergétique et de fiabilité. La deuxième partie porte sur les protocoles prenant aussi en compte le besoin d'offrir des temps de réponse bornés.

Les deux mécanismes introduits dans la première partie, WiseNE et Rep, utilisent des métriques composites, augmentant ainsi notablement la précision de l'estimation de la qualité du lien et la fiabilité des transmissions, tout en abaissant drastiquement la consommation d'énergie par rapport aux solutions proposées dans le passé. En dissociant les périodes d'échantillonnage et d'exploration de chaque nœud du réseau, le nouveau principe de balisage utilisé dans WiseNE permet de combiner efficacement d'un point de vue énergétique les métriques RSSI (indicateur de force du signal reçu) et LQI (indicateur de qualité du lien) à la qualité estimée. Ce découplage permet d'associer l'algorithme Trickle, qui est un élément clé de la basse consommation des protocoles, à des métriques composites. WiseNE a été appliqué à la «Triangle Metric» et validé lors d'un déploiement en ligne. Cette partie introduit ensuite un nouveau mécanisme d'échantillonnage appelé Rep. Celui-ci s'appuie sur la répétition de paquets déjà présente dans les protocoles MAC basse consommation basés sur l'échantillonnage de préambules, dans le but de diminuer d'un ordre de grandeur la consommation énergétique du WSN. WiseNE, Rep et la nouvelle métrique composite «Penalized RSSI» (PRSSI), combinant

Acknowledgements

PRR et RSSI, ont été validés dans le cadre d'un déploiement dans une ville intelligente. La deuxième partie du travail introduit deux résultats issus du projet WiseSkin qui vise à concevoir une peau artificielle hautement sensible pour les prothèses humaines, généralisables au domaine des systèmes cyber-physiques. Le premier résultat, Glossy-W, est un protocole qui améliore le compromis déjà favorable entre consommation et latence offert par les techniques d'inondation basées sur les transmissions simultanées. Glossy-W remplit les critères de synchronisation requis pour obtenir une inondation robuste, indépendamment du nombre de nœuds devant simultanément notifier un événement. Nous proposons un mécanisme supplémentaire appelé échantillonnage de canaux synchronisés (SCS) qui peut réduire la consommation de la scrutation périodique tout en propageant de manière fiable la détection d'un événement et en améliorant la cohabitation des réseaux. Les expériences réalisées montrent que SCS réduit la consommation énergétique de plus d'un tiers par rapport au protocole de référence «Back-to-Back Robust Flooding», tout en offrant une fiabilité comparable et en supportant la notification d'événements simultanés. Les avantages offerts par SCS sont valables pour toute la famille de protocoles actuels basés sur le mécanisme de transmission synchrone.

Mots-clés: réseaux de capteurs sans fil, WSN, internet des objets, routage, contrôle d'accès au medium, estimation du lien, métriques composites, transmissions simultanées, déploiements urbains en plein air, systèmes cyber-physiques.

Contents

Acknowledgements	i
Abstract	v
Résumé	vii
List of figures	xiii
List of tables	xvii
Introduction	1
1.1 Wireless Sensor Networks	1
1.1.1 Definition, Development and Applications	1
1.1.2 Key Performance Indicators	1
1.1.3 Multi-hop topologies	2
1.2 Communication Protocols	3
1.2.1 WSN Architecture and Protocol Layers	3
1.2.2 Routing Techniques	5
1.3 Problem Statement	6
1.4 Contributions	6
1.5 Thesis Structure	8
Part I	11
2 Increasing WSN Reliability by Enabling the Use of Composite Metrics	17
2.1 Introduction	17
2.2 Problem Statement	18
2.3 State-of-the-art Acquisition of Link Quality Information	19
2.4 Proposed Protocol - WiseNE	21
2.4.1 Design Objectives	21
2.4.2 Design	23
2.4.3 Composite Metric	26
2.5 Experimental Results	27
2.5.1 Methodology	27
2.5.2 Evaluation	28

Contents

2.5.3	Difficulties Encountered	33
2.6	Conclusion	34
3	Leveraging MAC Preambles for an Efficient Link Estimation	35
3.1	Introduction	35
3.2	Problem Statement	36
3.3	Related Work and State-of-the-art	36
3.4	Proposed Protocol - Rep	38
3.4.1	Key Features	38
3.4.2	Implementation	39
3.4.3	Advantages	40
3.4.4	Other Issues	41
3.5	Experimental Results	42
3.5.1	Energy Considerations	42
3.5.2	Performance Analysis	47
3.5.3	Consequences of Reducing the Sampling Granularity	58
3.6	Applications	59
3.7	Encountered Difficulties	61
3.8	Conclusions	62
4	Deploying Collection Protocols in Real World Applications – The Fly Project	63
4.1	Introduction	63
4.2	Project Fly	64
4.2.1	Project Description	64
4.2.2	Implications for the WSN Protocols	65
4.3	Problem Statement	67
4.4	State-of-the-art Studies in Real Deployments	67
4.5	Design of the WiseFly Collection Protocol	69
4.5.1	Proposed Composite Metric - PRSSI	70
4.5.2	Enabling PRSSI through WiseNE	71
4.5.3	Link Quality Sampling Based on Rep	71
4.6	Experimental Setup	72
4.7	Results	72
4.7.1	Indoors Pre-validation of PRSSI	73
4.7.2	Benchmarking PRSSI to ETX in Project Fly in an Urban Environment	76
4.7.3	Fine-grain (Rep) Sampling vs. Coarse-grain	80
4.7.4	Combined Performance Analysis of the Deployment – Rep and WiseNE-enabled RSSI	82
4.8	Practical Lessons	87
4.9	Conclusions	91

Part II	93
5 Leveraging Constructive Interference for Improving the Energy-Latency Tradeoff	95
5.1 Introduction	95
5.2 Time-bounded Event Driven WSNs – the WiseSkin Project	97
5.2.1 WiseSkin - Project Description	97
5.2.2 Performance Requirements	98
5.2.3 Two-mode Operation: Low and High-Traffic	99
5.3 Problem Statement	100
5.4 State-of-the-Art Solutions	100
5.4.1 Avoiding Concurrent Transmissions	101
5.4.2 Leveraging Concurrent Transmissions	102
5.5 Evaluation: Capture Effect vs. Constructive Interference	105
5.6 Proposed Protocol – Glossy-W	107
5.6.1 Using Glossy-W in a Two-mode Operation	109
5.7 Evaluation	110
5.7.1 Modeling Avoiding vs. Leveraging Synchronous Transmissions	110
5.7.2 Experiments on Avoiding vs. Leveraging Synchronous Transmissions	117
5.7.3 Comparing Glossy-W with Floods over an Ideal Collision-Avoidance MAC	119
5.8 Conclusion	120
6 Improving the Energy-Latency Tradeoff in Event-driven Protocols based on Synchronous Transmissions	121
6.1 Introduction	121
6.2 Problem Statement	122
6.3 State-of-the-art	123
6.3.1 Back-to-Back Robust Flooding (B2B)	123
6.3.2 Other Related Work	123
6.4 Proposed Protocol - Synchronous Channel Sampling (SCS)	125
6.4.1 Synchronous Channel Sampling	126
6.4.2 Improving Glossy-like Protocols	128
6.4.3 Wake-up from Multiple Sources	129
6.5 Experimental Results	131
6.5.1 Impact of the External Interference and the Mitigation via Verifications in Multiple Channels	132
6.5.2 Dependability Benchmark	135
6.5.3 Characterization of SCS	137
6.5.4 Wake-up from Multiple Sources	140
6.6 Optimizations	142
6.7 Conclusion	143
7 Conclusion	145

Contents

Bibliography	161
Curriculum Vitae and Publication List	163

List of Figures

1.1	Single-hop versus multi-hop communication in Wireless Sensor Networks. . .	3
1.2	Open Systems Interconnection reference model and data link layer architecture. . .	4
1.3	Flooding in Wireless Sensor Networks.	5
1.4	The three reception regions for indoor conditions vs. LQI metric.	14
2.1	Timing diagram of probing beacons for extracting link quality.	22
2.2	Working principle of a collection protocol based on WiseNE.	24
2.3	Timing Diagram of the Neighborhood Exploration	25
2.4	The assessment of the link quality improves through the use of a window average and it is consistent in <i>indoor</i> and <i>outdoor</i> conditions.	29
2.5	Future PRR vs. the current value of PRR (a), \overline{LQI} (b), \overline{SNR} (c) and TM (d) for <i>indoor</i> conditions.	31
2.6	Future PRR vs. the current value of PRR (a), \overline{LQI} (b), \overline{SNR} (c) and TM (d) for <i>outdoor</i> conditions.	32
2.7	Progression of the prediction accuracy as a function of the time horizon.	33
3.1	Comparison between coarse-grain (top) and fine-grain (bottom) link estimation. 38	
3.2	Rep's working principle.	40
3.3	Timing diagram of the link estimation operation, between motes M1 and M2, using multiple packets (Coarse-grain sampling).	44
3.4	Timing diagram of the link estimation operation, between motes M1 and M2, using multiple repetitions of a single packet with Rep (fine grained sampling).	45
3.5	Power savings of using Rep instead of using CG to obtain a number of samples, as a function of the number of samples and the network density.	46
3.6	Location of the experiment in an office hallway.	47
3.7	Variance of LQI (top) and RSSI (bottom) during 13 days.	48
3.8	Future PRR vs. the current value of \overline{LQI} with negligible wireless interference and variable fading.	50
3.9	Future PRR vs. the current value of \overline{LQI} to study the effect of wireless interference. 52	
3.10	Future PRR vs. the current value of \overline{RSSI} to study the effect of wireless interference. 53	
3.11	Regression from Future PRR vs. the current value of \overline{LQI} datasets using a sigmoid model.	55
3.12	Regression from Future PRR vs. the current value of \overline{RSSI} datasets using a sigmoid model.	55

List of Figures

3.13	Future PRR vs. the current value of \overline{LQI} to study the effect of variable fading. . .	57
3.14	Future PRR vs. the current value of \overline{RSSI} to study the effect of variable fading. .	58
3.15	Future PRR vs. the current value of \overline{LQI} to study the effect of simultaneous wireless interference and variable fading.	59
4.1	Sample of a typical routing tree provided by the PRSSI metric with 25 motes. . .	73
4.2	Placement of the motes in the csem-tb deployment.	74
4.3	PRR prediction using PRSSI (RSSI in mW units).	75
4.4	PRR prediction using PRSSI (RSSI in mW units). Corresponds to Fig. 4.3 augmented in x-axis.	75
4.5	PRR prediction using PRSSI (RSSI in dBm units).	77
4.6	PRR prediction using ETX.	78
4.7	PRR prediction using ETX. Corresponds to Fig. 4.6 augmented in x-axis. . . .	78
4.8	PRR prediction using RSSI from Rep.	79
4.9	PRR prediction using CG RSSI.	79
4.10	Relation between \overline{RSSI} measuring using 20 samples obtained through Rep and CG. The line $x = y$ is shown in green for reference.	81
4.11	Number of samples per \overline{RSSI}_{Rep} in Fig. 4.10, between \overline{RSSI} measuring using 20 samples obtained through Rep and CG.	81
4.12	Packet received at the sink (type Routing Statistics) vs. time. The graph covers 30 days in project Fly.	83
4.13	Routing tree in project Fly before the loop.	83
4.14	Routing tree in deployment project Fly after the loop.	84
4.15	Packet Reception Ration per mote in deployment project Fly.	84
4.16	Time evolution of the number of successfully sent unicast packets (ack received) per hour.	85
4.17	Component of the radio Duty Cycle spent with the radio on listening or reception state (%) vs. time.	86
4.18	Component of the radio Duty Cycle spent with the radio on transmission state (%) vs. time.	86
4.19	Time evolution of the number of unicast packets sent with the TxMore optimization in WiseMAC vs. time.	87
4.20	Network topology during the sending of audiograms 1.	88
4.21	Average end-to-end latency for each mote.	88
4.22	Average hop count for each mote.	89
5.1	Schema of WiseSkin.	97
5.2	Timing diagram of Glossy relaying a flood through a 2-hop WSN, where every mote transmits only once ($N_{gl} = 1$).	102
5.3	Location of the testbed in an office hallway.	105
5.4	Measurements of the Event Detection Ratio using two mechanisms for concurrent transmissions.	106
5.5	Timing diagram of Glossy-W relaying a flood through a 2-hop WSN.	108

5.6	Timing diagram of an event-triggered WSN alternating between two modes: low-traffic mode and high-traffic mode.	109
5.7	Timing diagram of ContikiMAC relaying a flood through a 2-hop WSN using Netflood.	112
5.8	Duty Cycle (D) as a function of the Latency (L) for a 2-hop network. Glossy-W displays a lower duty-cycle in the entire range of latency values.	117
5.9	Measured Latency (L) as a function of the Wake-Up period (T_w) in a 2-hop network.	118
6.1	Schema of a glossy-like protocol operating in an event-triggered application.	122
6.2	Schema of a flood in B2B.	124
6.3	Timing diagram of SCS in a 3-hop network in the absence of events to be notified, and waking up the network upon detection of an event.	127
6.4	Schema of SCS enhancing a glossy-like protocol, for operating in an event-triggered application.	128
6.5	Schema of an event-triggered system with distributed event detection without SCS and with SCS.	130
6.6	Schematic of the Flocklab testbed.	131
6.7	SCS' sensitivity to interference from jammers - power, as a function of the number of CS verifications (N_{ch}).	134
6.8	SCS' sensitivity to interference from jammers - false wake-ups, as a function of the number of CS verifications (N_{ch}).	134
6.9	Evaluation of the mean number of hops per node in B2B.	136
6.10	Improvement of the mean power per node brought by SCS on B2B. The missing columns correspond to the jammer motes or non-existing motes.	137
6.11	Event Latency per node for B2B and SCS.	138
6.12	Impact of SCS on B2B's mean power consumption as a function of the mean event inter-arrival time.	139
6.13	Impact of SCS on B2B's mean power consumption as a function of the mean event inter-arrival time (augmented view).	139
6.14	Comparison between the mean power consumption of the SCS-enhanced protocol for different values of the channel sampling threshold.	140
6.15	Evolution of the mean power consumption with and without SCS, with the number of motes that simultaneously detect events.	141

List of Tables

3.1	Model Parameters	43
3.2	Experimental Conditions	49
3.3	Results of the regression using the sigmoid in Eqs. (3.9) (LQI) and (3.10) (RSSI).	56
5.1	Protocol Parameters	115

Introduction

1.1 Wireless Sensor Networks

1.1.1 Definition, Development and Applications

Wireless Sensor Networks (WSN) is a term used to describe collections of spatially distributed and autonomous electronic devices (motes) equipped with sensing, on-board processing, storage and wireless communication capabilities [1]. They are typically used to monitor the physical parameters of an environment and then transmit the relevant data wirelessly to a sink, which forwards it to remote devices with superior processing and decision making capabilities. Some applications, including the ones studied in the second part of our dissertation also require the motes to have actuation capabilities, in order to execute the commands given by the sink.

The accelerated research advances in the fields of microelectronics and communication systems have made WSNs increasingly cost-effective, opening the door to a myriad of use-cases. Nowadays, wireless sensor networks' applications range from the monitoring and control of industrial processes, to environmental monitoring and the support of health care applications through wearable and even implanted miniature WSNs.

While the advancements in related technologies have resulted in an ubiquitous use of WSNs, the desire to employ them in increasingly challenging applications and environments (e.g., extensive and remote areas in the Arctics, human organs, etc.) is imposing more stringent performance criteria, in terms of network size (large WSNs can comprise thousands of motes), sensitivity, transmission reliability, energy autonomy, density, speed, miniaturization, etc. These requirements generate both individual costs and often trade-offs (e.g., between a high reliability and a low energy consumption), driving research efforts, such as the ones in optimizing the information traffic through routing protocols that we are dedicating this dissertation to.

1.1.2 Key Performance Indicators

Depending on the characteristics of the application, the underlying WSN can be subject to a multitude of specific performance requirements. Nonetheless, most WSN use-cases pose a

series of common requirements:

- **Reliability:** A key challenge of WSNs, their reliability consists in the motes' capacity to sense and convey the sensed data to the destination mote, by making sure meaningful packets are not lost on their way. WSN reliability is often classified into Hop-by-Hop or End-to-End reliability [2]. In applications where data transmissions represent a dominant energy carrier, high levels of reliability often enter a trade-off with the WSNs' energy efficiency.
- **Energy Efficiency:** As WSNs are often deployed in environments difficult to access, be they remote areas [3] or sensitive human tissue, as well as composed of a high number of motes, their ability to operate over a long time (even multiple years) without the need of human intervention is a key requirement. At the same time, the ongoing quest for miniaturization poses a further limitation on the battery power available to each mote. This translates into the need to design systems for minimizing the energy use of each mote, especially the communication one, a dominant component of the energy consumption [4].
- **Bounded Latency:** Many applications (e.g., reporting alarms or controlling an actuator in a closed-loop) require the network to convey the information within a given deadline. This creates the need for the network to operate with bounded latencies, which creates an additional trade-off in terms of an efficient energy use [5].
- **Self-organization:** Since motes are typically not aware of their neighbors when the WSN is deployed [6], and at the same time some of them might change positions or fail during operation, the capacity of the network to adapt the topology becomes paramount.
- **Scalability:** With applications requiring increasingly more spatially detailed sensing, the ability to scale the WSN to a high number of motes is growing in importance.
- **Cost Efficiency:** Many performance requirements often boil down to making sure the WSN is built and operated with the minimum financial effort possible. Cost efficiency can be achieved through multiple levers, from a low-cost manufacturing of robust devices, to the minimization of user intervention once the network is operational.

Additional requirements include a low network complexity (both in terms of hardware and software), the ability to co-exist with multiple other networks, high robustness and high levels of security.

1.1.3 Multi-hop topologies

WSN covering large geographic surfaces often cannot rely on every mote communicating directly to the sink, in what is known as a star topology [1]. This limitation appears both

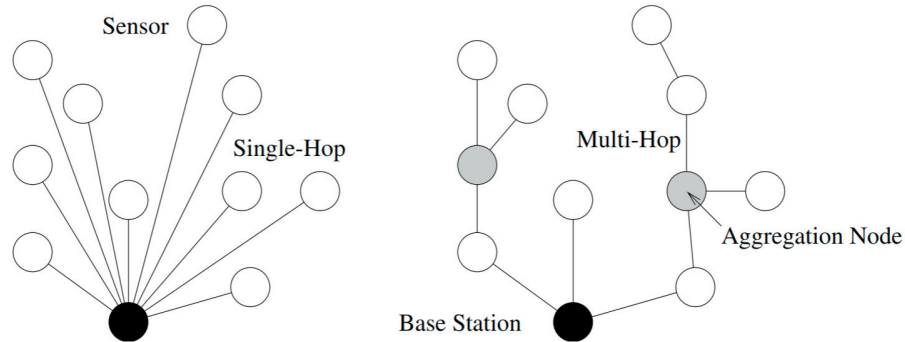


Figure 1.1 – Single-hop versus multi-hop communication in Wireless Sensor Networks (figure reproduced from [1]).

because the distance between the motes and the sink can be larger than the communication range of the mote's transceiver, and due to the fact that the propagation path loss is approximately proportional to the square of the distance, thus making long distance transmissions significantly less energy efficient [7]. In these situations, very frequent especially in outdoor deployments, the WSN relies on a multi-hop topology, in which multiple motes collaborate to send their information to the sink through consecutive transmissions. In this topology, motes do not only need to communicate their own data, but also help relay other motes' information. Figure 1.1 shows a graphical depiction of a star (left) and multi-hop (right) topology .

Finding the optimal path from each mote to the sink in a way that balances the need for a reliable, fast and energy efficient transmission is however a non-trivial task that has motivated extensive research in the area of routing protocols, including the current dissertation.

1.2 Communication Protocols

The reliable and resource efficient transmission of information between the multiple motes of a multi-hop WSN depends on the definition of adequate communication and network protocols. In order to reduce the complexity of the protocol design process, and allow for their modular development, the international research community generally relies on a hierarchical layering of the functions that need to be performed by the protocols. The Open Systems Interconnection (OSI) Reference Model [8] [9] is one of the most broadly used structure, consisting of 7 layers: Application, Presentation, Session, Transport, Network, Data Link and Physical.

1.2.1 WSN Architecture and Protocol Layers

WSNs are typically concerned with the bottom four layers of the OSI model, namely the Transport, Network, Data Link and Physical layers, plus the application layer. Figure 1.2 shows a broad decomposition of the layers in the OSI Model.

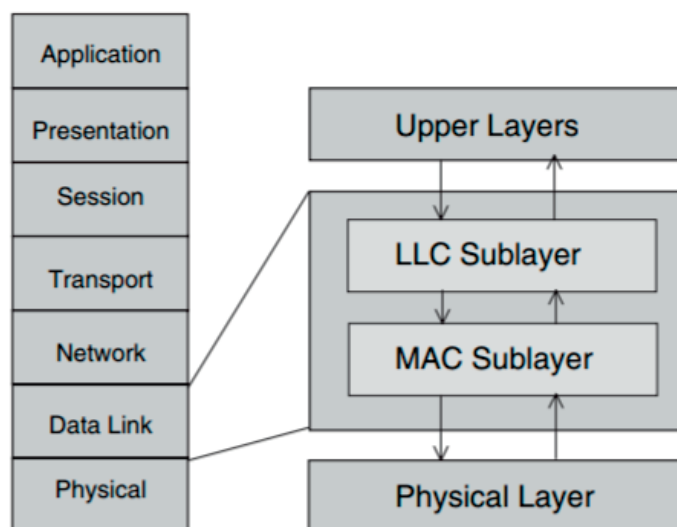


Figure 1.2 – Open Systems Interconnection reference model and data link layer architecture (figure reproduced from [10]).

While the Network layer is served by Routing protocols, the Data Link layer's functions are split into the MAC (Medium Access Control) sublayer and the LLC (Logical Link Control) sublayer. At the time, significant research effort is being invested into the development of MAC and Routing protocols.

- **The Data Link Layer - MAC Protocols:** This family of protocols deals with the way in which nodes using the same physical space for transmitting their electromagnetic signals will be sharing their access medium. MAC Protocols have to fulfill a wide set of performance requirements, ranging from a high throughput and reliability, low energy consumption and access delay, scalability, robustness and fairness [10] [11]. MAC protocols typically rely on three types of addresses: 1) Unicast - the data to be transmitted by one node is directed towards another specific node, 2) Multicast/Broadcast - the data to be transmitted by one node is directed at a predefined set of nodes and 3) Anycast - the data to be transmitted by one mote is only sent to the first neighbor that answers with an acknowledgement [12].
- **The Network Layer - Routing Protocols:** As multi-hop topologies require nodes to collaborate in the transfer of data (both their own and the one coming from other nodes), routing protocols are required in order to identify the optimal succession of nodes to be used for the transmission, namely the routing path. These protocols need to generate reliable, stable and resource (energy, time) efficient paths [10]. In order to minimize the operations performed and save resources, nodes typically collect transmission relevant information about their neighbors (e.g., the quality of the link to the respective node) in routing tables, which they periodically update in order to cater for the changes in the channel dynamics.

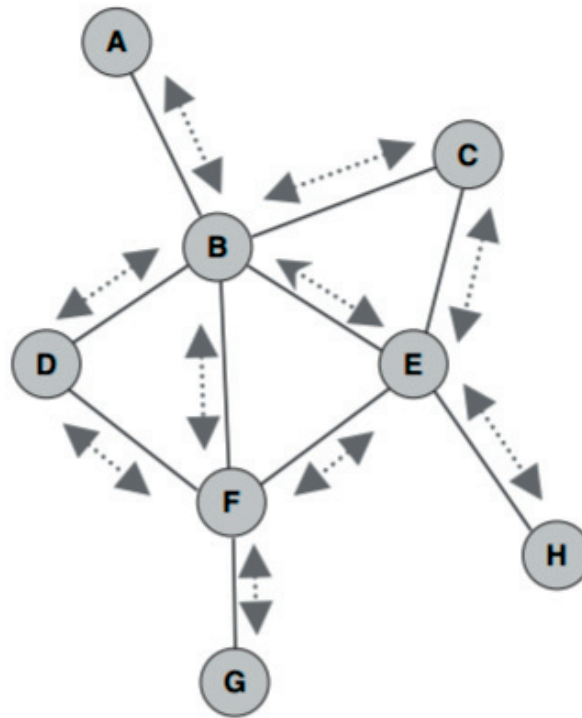


Figure 1.3 – Flooding in Wireless Sensor Networks (figure reproduced from [10]).

1.2.2 Routing Techniques

WSNs rely on a multitude of routing techniques. Two broad categories of techniques that are broadly used and will be further studied in the frame of this dissertation are i) Collection Tree Protocols and ii) Flood-based protocols.

- **Collection Tree Protocols:** define the routes that the packets follow to a given destination through a multi-hop network. They enable a mote to communicate beyond its transmission range by using the neighbors as relays. The collection operates on top of the MAC layer, which provides the service of exchanging packets with the neighbors and offers information about them (e.g., quality of the links). Collection tree protocols will be the focus of the first part of this dissertation.
- **Flood-based Protocols:** instead of investing in maintaining information about each node's neighbors, flood-based protocols rely on every neighbor forwarding its packets (Fig. 1.3). While flooding can increase the likelihood of data being reliably received due to exploiting every possible route, they can also cause indefinite loops [10], high overheads, and lower reliability in sparsely connected networks. This might necessitate retransmissions, typically not available to the broadcast primitive. The second part of this dissertation will detail the use of floods and the respective Glossy-like protocols.

1.3 Problem Statement

The need to reliably convey WSN data in deployments with increasingly challenging demands, be they a long-term energy autonomy in order to avoid expensive and intrusive interventions, or the real-time transmission of time-critical information, has motivated extensive research into the design of communication protocols.

Routing protocols in particular represent a challenging, yet also promising research space, given their responsibility of identifying the most robust data paths across extensive regions, subject to a myriad of perturbations. Our dissertation is thus focused on the development of mechanisms that manage to significantly improve the performance of state-of-the-art routing protocols. Acknowledging the emergence of requirement clusters in the multitude of applications, we have decided to address the issues of two different application families.

The first part of our thesis will be addressing the trade-off between the high reliability and the long energy autonomy expected from multi-hop WSNs. While composite metrics have been establishing themselves as enablers of significant reliability increases, their use has been limited to offline deployments, given their significant energy overheads. We will proceed to propose a series of systems that will help unlock composite metrics' potential to be used in real-life online deployments, by significantly reducing their energy consumptions.

In the second part, we will also address the energy-latency tradeoff faced by WSNs tasked with the real-time transmission of sporadic events. In order to reliably detect events, protocols catering for such applications typically rely on energy-expensive network polling mechanisms. Moreover, when an event does happen, the network has to deal with potentially disruptive packet collisions, given its detection by multiple neighboring nodes in high density WSNs. We will thus aim at designing solutions that significantly reduce the overhead of the polling mechanism, while being robust to simultaneous event detections.

Beyond these typical WSN trade-offs, this work will also address a number of additional challenges encountered by WSN deployments, including the effects of interferences, and the WSNs' ability to co-exist with other networks. Last, but not least, it will aim at leveraging the opportunity of having tested and implemented the systems in real-life projects in the direction of enriching the currently limited practical knowledge related to real life WSN deployments.

1.4 Contributions

The first part of the thesis has introduced two main contributions, WiseNE, and Rep. WiseNE is a novel beaconing mechanism that manages to increase the reliability of collection tree protocols, by enabling an energy efficient use of the RSSI (Received Signal Strength Indication), a key component of many composite metrics. By decoupling the sampling and exploration periods of the protocol, WiseNE manages to maintain the Trickle Algorithm active, thus generating a significantly lower overhead. This development also allowed us to perform the

first online validation of the Triangle Metric (composed of PRR, RSSI and LQI).

This overhead is further reduced (by one order of magnitude) with the introduction of a novel sampling mechanism, Rep. This scheme achieves the significant drop in energy consumption and traffic, and additionally increase the LQE (Link Quality Estimation) speed by leveraging the pre-existing repetitions in low-power preamble-sampling MAC layers, as opposed to waiting for additional beacons. Given the high frequency variations pattern of the dominant source of interference present in indoor office deployments (WiFi), the fine-grain sampling performed by Rep over 1 second achieves the same accuracy as the SOTA coarse-grain sampling over 2.7 minutes, while consuming one order of magnitude less energy.

The two mechanisms were also validated in the frame of a real deployment in an outdoors urban environment, as part of the Fly project, a WSN that we have developed and deployed for the monitoring of public utilities in a European Smart City initiative. The experiments conducted proved that WiseNE and Rep continue to deliver their promised improvements in environments subject to significant perturbations, notably multipath fading variations as a result of nodes being installed at ground level and large distances (in the order of hundreds of meters) without line-of-sight. The deployment has also introduced a newly developed composite metric, PRSSI (a combination between PRR and RSSI). Given the scarcity of such deployments, mostly due to the prohibitive cost of monitoring the network, we have decided to also document some of the most important practical lessons learned, ranging from the need to install a data collection mechanism in order to facilitate a non-intrusive experimentation, to the value of maintaining a close collaboration with the real beneficiary to insure the optimal utilization of the WSN.

The second part of the thesis has focused on the development of mechanisms also able to cater to very low latency constraints. This additional requirement was inspired by the need to develop a WSN for the WiseSkin project. Our particular role in this interdisciplinary project consisted in the development of a WSN able to convey the tactile information through the artificial skin of a prosthetic arm. As the ability to create a seamless tactile experience is one of the main criteria for patients' accepting the prosthesis, our WSN had to achieve very low latencies, on top of high reliabilities and an energy consumption compatible with a prolonged use. We have thus investigated the problem of distributed wake-up of an event-triggered WSN with tight latency, reliability and energy consumption constraints.

In this context, we have developed Glossy-W, a modified version of Glossy tailored to reduce the energy consumption for a given latency, while catering for multi-node simultaneous event detection. Glossy-W incorporates two main contributions. The first contribution consisted in the validation of the fact that leveraging concurrent transmissions as opposed to avoiding them can actually result in better energy performances for a given latency, despite the contrary expectation. Additionally, we show that, even when removing the per-hop contention by using an Ideal Collision Avoidance MAC, Glossy-W maintains a better energy-latency trade-off for the dense WSNs under consideration. The second contribution consisted in validating

that, among protocols leveraging concurrent transmissions, the ones relying on constructive interference are preferable to the ones based on the capture effect, given their superior robustness to simultaneous event detection, a key constraint of high density WSNs. Glossy-W is thus the first protocol able to leverage floods that rely on constructive interference in a simultaneous event detection scenario. While inspired by WiseSkin, Glossy-W has a broad range of applicability, especially as an increasing number of WSN applications require a dense network, consistently fast responsiveness and high energy autonomy.

Last, but not least, we have introduced a novel primitive: Synchronized Channel Sampling (SCS). SCS consists of consecutive measurements of the signal strength in different channels, performed synchronously by all the motes in the network. Its main contribution in improving the state-of-the-art protocols, such as B2B, relies on its providing a more power efficient mechanism for polling the network for events, compared to the traditional approach of sending floods. The improved power performance results from SCS' substituting the floods that B2B used for performing the network polling with more power efficient Synchronous Channel Sampling rounds. This proved particularly useful in low event frequency scenarios, where the network energy consumption is driven by the polling mechanism. The testbed experiments performed demonstrate that the use of SCS reduces the power consumption of the top performing protocol (from the EWSN17 competition) by up to 40%, while keeping an equivalent reliability performance.

1.5 Thesis Structure

While the entirety of our work has focused on the development of high-performing WSN routing protocols, the variability in application requirements has motivated us to organize the thesis into two parts.

The first part will focus on applications subject to high reliability and energy autonomy constraints. Chapter 1 will introduce WiseNE, a novel beaconing mechanism that enhances collection tree protocols' reliability, while maintaining a high energy efficiency. This chapter will also validate the Triangle Metric in an online deployment. Chapter 2 will introduce Rep, a sampling mechanism that will manage to reduce the energy required for the extraction of the RSSI and LQI metric by one order of magnitude, by exploiting the pre-existent repetitions at the MAC layer. The 3rd, and final chapter in Part 1 will take the opportunity to validate and deepen our understanding of composite metrics, as well as the two contributions of the previous chapters, WiseNE and Rep in a real-life smart city deployment called WiseFly.

The second part of the thesis will transition to routing protocols catering for applications subject to an additional low latency requirement, and simultaneous event detection such as WiseSkin, a WSN aimed at conveying the tactile stimuli perceived by the artificial skin of a human prosthesis in real time. The first chapter of this part (4th in the overall thesis organization) will introduce Glossy-W, a protocol that enhances Glossy's ability to optimize the energy – latency trade-off, while being compatible with simultaneous event detections

in multiple motes. Finally, Chapter 5 will present Synchronized Channel Sampling (SCS), a novel mechanism that reduces the WSN energy consumption by 33.3 - 40%, while remaining compatible with simultaneous event detection.

Part I

Introduction

In the first part, this dissertation focuses on the design of routing protocols for collection applications, where high standards of reliability and energy consumption are paramount.

In the last years, the quest for high reliabilities has been promoting the use of composite metrics as promising alternatives to simple ones for the Link Quality Estimation (LQE) performed by collection protocols. However, while composite metrics' superiority in terms of reliability is broadly acknowledged, protocols aiming at using them end up generating significant energy overheads, due to the high number of samplings required by RSSI (Received Signal Strength Indication) and LQI (Link Quality Indication), key components of best performing composite metrics. This results in applications of composite metrics being limited to offline environments, or testbeds with a steady energy supply. In this first part of our dissertation we will introduce two mechanisms, WiseNE and Rep, that manage to enable the use of composite metrics, while creating a significantly lower overhead than the state-of-the art protocols, thus allowing their use in online applications under strict energy efficiency constraints.

The novel beaconing scheme WiseNE described in Chapter 1 manages to extract sufficient samples for the RSSI and LQI components of composite metrics, by decoupling the sampling and exploration period of each mote. This allows the protocol to keep leveraging the Trickle Algorithm for the network exploration phase, and thus continue benefiting from its energy savings. WiseNE has been used to validate the Triangle Metric, a composite metric comprising RSSI, PRR (Packet Reception Ratio) and LQI in an online setting, proving to significantly enhance the WSN's reliability, while requiring far less energy for the addition of RSSI and LQI, than state-of-the art applications. The use of WiseNE also translates into a sharp improvement in the LQE speed, allowing for faster routing decisions.

Moreover, the novel sampling mechanism introduced in Chapter 2 (Rep), will manage to further reduce the energy consumption required by the RSSI sampling operation by one order of magnitude. Rep's working principle relies on leveraging the packet repetitions already present in low-power preamble-sampling medium access control protocols in order to improve the WSN's energy consumption. Under the typical interference sources of an office environment (e.g., WiFi), the fine-grain sampling mechanism (over 1s) that Rep relies upon has proven

to result in equivalent reliabilities as the typical coarse-grain sampling (over 4 min), while providing a faster, and more energy efficient sampling.

In the frame of Chapter 3, we have taken the opportunity to validate WiseNE, Rep and the composite metric PRSSI (Penalized Receive Signal Strength Indication) in a real smart city deployment. This endeavor has proven that the two mechanisms (WiseNE and Rep), as well as PRSSI continue to achieve their stated contributions under challenging environments (e.g., lack of line-of-sight, dominant multipath fading). Moreover, given the scarcity of real-world deployments of collection protocols available, the chapter also adds value by documenting multiple practical considerations relevant to outdoors urban deployments, than can hopefully prove useful to our field's community.

Background

Collection Protocols in WSNs

A collection protocol defines the routes that the packets follow to a given destination through a multi-hop network. It enables a mote to communicate beyond its transmission range by using the neighbors as relays. The collection operates on top of the MAC layer, which provides the service of exchanging packets with the neighbors and offers information about them (e.g., the quality of the links).

The network can operate with one or several sinks, which are motes that can forward the collected data beyond the WSN (e.g., through a cellular gateway). The routing aims to autonomously organize the network in destination-oriented acyclic graphs (DODAGs) that provide every mote with a preferred route to reach each sink, acting as the roots of the DODAGs (e.g. RPL [13])¹. A common topology is the one of a routing tree with the sink at the root, resulting from keeping a single DODAG with a single sink (e.g., CTP [14]). In order to create and maintain a topology, the core mechanism of the routing protocol must perform the following three operations:

- **Exploration:** probing the neighborhood of a mote in order to gather information about other motes within range. This is necessary in order to identify the potential parents or children in the DODAGs, and to gather link quality data.
- **Link and Route Estimation:** quantifying the quality of the links to the potential parents and their associated routes to the sink(s). This uses the link estimation techniques described in Sec. 1.2.2.
- **Parent Selection:** identifying the suitable parents (connected with high quality and stable links, not prone to generating routing loops, etc.) and choosing the best candidate.

¹The DAG structure used by RPL enables each mote to associate to multiple parents in the DAG. Therefore, it is a more general topology compared to the collection tree, where motes have a single parent. Through this thesis, we include RPL in our study of tree collection protocols because the analysis is agnostic to the number of parents.

These operations follow a best-effort approach to define a topology for the network, which: is composed by the links with the best available quality, remains stable in the future, and has the shortest routes possible, thus minimizing the energy and number of transmissions necessary for reaching the sink. Nevertheless, the conditions for improving the topology depend on the dynamic environment of the network, where links might break, parents might disappear and interference patterns are usually variable, thus changing the solutions that provide an acceptable performance. The result is that the routing protocol must constantly repeat this process in order to maintain a high reliability, at the expense of increasing the control traffic and the energy consumption. This reliability-energy trade-off is managed by state-of-the-art routing protocols by adapting the control traffic to the dynamics of the network environment by means of an Adaptive Beaconing mechanism, which relies on a trickle timer to optimize the message transmission frequency to the network conditions [15] [14]. In a nutshell, the frequency is quickly increased (i.e., exponentially) whenever an inconsistent network management information is received (e.g., a loop or a lost parent) for faster recovery from a potential failure, and progressively decreased (i.e., linearly) in case the network environment is stable, in order to reduce the control traffic and save energy.

Link Quality Estimation

Radio links are affected by multiple factors with the potential to significantly degrade their quality, from external interference from Wi-Fi routers or Bluetooth devices, to broken links due to obstacles between the antennas, or degraded RF channels due to harsh weather conditions or changes in the multipath fading pattern. In WSNs, these effects are particularly significant, given the limited resources to ensure robust links in the context of low-power and narrow-band radios. The quality of the link thus often fluctuates over time [16] [17] and space [18] [19] [20], and is often asymmetric with respect to the link direction [21] [20].

The ability to consistently estimate link quality, in the Link Quality Estimation (LE) phase [16] is thus key to enabling the motes to identify and use the links with the best potential for a high Packet Reception Ratio (PRR) in the future transmissions, a key indicator of the network's performance in terms of reliability. This prediction can be based on a series of present metrics. Current applications normally use one single metric as input data, a limitation we will be addressing in this chapter.

Figure 1.4 shows a classical categorization of the future expected Packet Reception Data (PRR), as a function of the present Link Quality Information (LQI) metric. In the frame of this study, we will be relying on this categorization of the PRR, into three quality bands, as defined by [22]. The categorization is characterized by a specific Packet Reception Ratio (PRR) and three dynamic bounds (changing over time):

- Connected: links display a good quality, stability and symmetry. They can be classified as good ($\text{PRR} > 75\%$) and very good ($\text{PRR} = 100\%$).

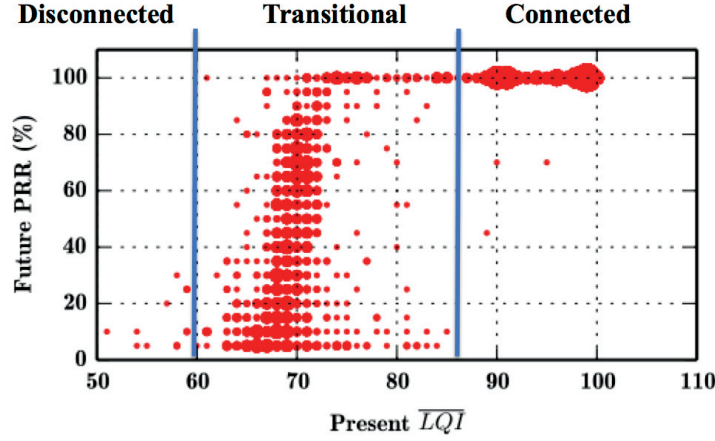


Figure 1.4 – The three reception regions for indoor conditions vs. LQI metric. These scatter plots provide insight about the prediction capabilities of LQI. The area of each dot is proportional to the number of data points with the same coordinates.

- **Transitional:** links are characterized by persistent packet losses (in long-term assessments), unstable, not correlated with distance, and often asymmetric. They are also called intermediate quality links ($75\% \geq \text{PRR} > 35\%$).
- **Disconnected:** links have poor quality and are not suitable for communication. They are often referred to as bad links ($\text{PRR} \leq 35\%$).

LE uses a statistical characterization in order to describe the behavior of the links, an approach that facilitates the use of measured empirical data with a random component, to estimate the value of a parameter. However, the identification of input metrics that can accurately represent the link quality (LQ) and its trends, and could thus be a reliable predictor of the future PRR is a research problem that remains open [23]. The perturbing factors that make the prediction of the future PRR necessary, such as Wi-Fi routers or Bluetooth devices, the broken links due to obstacles between the antennas, or the degraded RF channels due to harsh weather conditions, also represent a difficulty in our ability to estimate it.

This has motivated the search for metrics that accurately represent the status of links, with different qualities and their temporal dynamics, for over a decade (see [23] for a comprehensive review of these efforts). Some notable examples are:

- **(Past) Packet Reception Ratio (PRR):** represents the ratio of the number of successfully received packets to the number of transmitted packets in the past.
- **Expected Transmission Count (ETX):** is the inverse of the product of the PRR of the forward link and the PRR of the backward link [24]. ETX takes the link asymmetry property into account.
- **Link Quality Indicator (LQI):** is a measure of the correlation between the received symbol and the symbol to which it is mapped after the soft decoding stage in the radio [16]. LQI

is sometimes referred to as Chip Correlation Indicator (CCI).²

- Received Signal Strength Indicator (RSSI): is the signal strength of the received packet. This value is provided through a register by most radio transceivers³. In the absence of transmissions, the register gives the noise floor.
- Signal-to-Noise Ratio (SNR): is computed as the difference (in dBm) of the RSSI, without noise, and the noise floor.

For the purpose of this thesis, we will use a broad definition of the LE, that allows us to assess different metrics. We will refer to LE as the evaluation of a mathematical or logical expression (i.e., metric) within an observation window (a number of packets or seconds), used to characterize the quality of a link, as defined by Baccour *et al* in [23].

Composite Metrics

Despite the metrics in the previous section having enabled various practical collection applications, it is generally accepted that a single metric is only capable of assessing one particular link property, thus providing an incomplete characterization of the link [25]. For example, a low RSSI or LQI is an accurate indicator of bad quality links, while high values may indicate both good or bad quality links [22]. Conversely, PRR or ETX display a better accuracy with good quality links, compared to the one in the case of bad quality links. Therefore, extensive research efforts focus on composite metrics, addressing the aforementioned limitations by combining the individual strengths of different metrics in order to provide accurate and holistic characterizations. Notable composite metrics include Four-Bit [26], F-LQE [25], Triangle Metric [22] and the scheme proposed by Rondinone *et al* in [27].

An important challenge in the domain of composite metrics is the ability to combine metrics with different physical natures, such as RSSI (in dBm) and PRR (a percentage). A common approach is to normalize the values, as done in [27] or [25]

More recent work proposes the use of Fuzzy logic (F-LQE) [25] or Machine Learning schemes [28] [29] for finding metric combinations that more accurately estimate the quality of the links. The latter techniques rely on training data to identify optimal metrics to describe the behavior of the links. However, such techniques are currently not feasible due to their high requirements in terms of computational power, not yet accessible to the vast majority of WSNs.

Typically, the research in link estimation focuses on gathering the LQ information in the motes where it is immediately available, normally for offline analysis. This is the case of the Triangle Metric and other composite metrics such as [27] and [30].

²The cc2420, a transceiver used extensively through this thesis, measures the LQI on the first eight symbols of the received packet and is represented with a score between 50 and 110 (the higher the index, the better the quality).

³The CC2420 returns the average measured on the first eight symbols

2 Increasing WSN Reliability by Enabling the Use of Composite Metrics

2.1 Introduction

Wireless Sensor Networks (WSNs) use collection protocols in order to operate in areas that extend beyond the transmission range of a single mote, or that are difficult to access, such as glaciers [3], orchards [31], multi-store buildings [32] or even cruise ships [33]. Oftentimes, WSNs deployed in these less accessible areas are in charge of critical applications, such as passenger tracking for evacuation in case of sea accidents, or fire alarms. The protocols' reliability, namely their capacity to convey all data packets generated by the motes towards collection points, is thus a critical requirement in such applications. Moreover, the nature of these areas makes higher levels of WSN energy autonomy particularly desirable.

In terms of reliability, high performing collection protocols depend on an accurate performance of a link quality estimation (LE) [23], a piece of information that enables the routing protocols to choose the most reliable and efficient routes. However, the identification of metrics that can accurately reflect link quality is still an open research problem [23] [29] [28]. Latest research has shown that composite metrics have a significant potential to improve the delivery rate of WSNs, while at the same time reducing their power consumption. For example, the Triangle Metric (TM), a promising composite metric has the capacity to increase the accuracy of the state-of-the-art metric ETX for detecting good links from 82.3 to 98%. Nevertheless, SOTA routing protocols have not been leveraging this potential in collection applications, mostly due to two main types of implementation challenges.

First of all, the increased noise sensitivity of composite metrics requires the collection of a sufficiently high number of samples of the metric in each mote, in order to ensure a meaningful estimation of the link quality, and a fair comparison between the different links. This would however hamper the protocol's ability to maintain its low energy consumption, as the Trickle Algorithm that the energy savings of the SOTA protocols rely upon limits their ability to control the number of samples performed by each mote. SOTA protocols attempting the use of composite metrics, such as F-LQE [25] thus decide to disable the Trickle Algorithm in order to ensure sufficient samples, but end up with high energy consumptions.

WiseNE, the novel beaconing mechanism designed in the frame of this study, aims at solving this issue, by decoupling the sampling and exploration periods of each mote. This allows the mote to benefit from the energy savings of the Trickle Algorithm (driven by its exploration period) in the extraction of RSSI and LQI as well, while achieving the desired number of samples in the observation phase. WiseNE thus becomes the first algorithm enabling the use of composite metrics, while avoiding prohibitive increases in the energy consumption.

Second, the meaningful use of composite metrics is also enhanced by the protocol's ability to operate online, in order to be able to isolate the quality of the uplink and return the link quality information to the motes that need it for routing decisions. At the moment, however, SOTA link estimation schemes such as TM [22] or Foresee (4C) [29] rely on parallel infrastructures that evaluate the link quality offline. WiseNE is thus also the first protocol designed to validate the Triangle Metric in an online environment. The experiments performed confirm the metric's potential to combine the strengths of the Signal-to-Noise Ratio, the Link Quality Indication and the Packet Reception Rate, for accurately predicting the quality of the links, both inside an office and in outdoor environments. The study concludes that WiseNE can enable novel composite metrics and proposes an optimization method aimed at significantly reducing WiseNE's overhead.

In terms of structure, we will start this chapter by providing our Problem Statement, as articulated in Section 2.2. Section 2.3 will show an analysis of the limitations of state-of-the-art protocols in terms of using composite metrics. We will then proceed to introducing WiseNE in Section 2.4, describing both the mechanism which allows it to solve the aforementioned challenges, and its design. Section 2.5 will show the experimental setup and results that help substantiate our claim that the Link Quality Estimate obtained with the help of WiseNE presents an improvement in terms of accuracy from 82.3% to 98%. Lastly, the main conclusions will be reiterated in Section 2.6.

The work presented in this chapter was published in WiMob 2016 [34].

2.2 Problem Statement

There is a significant gap between the advancements in the domain of link estimation and their use in practical WSN applications. While multiple composite metrics have proven to outperform the individual metrics used traditionally, none of the collection protocols with widespread use leverages them, and they are rarely characterized in practical WSN deployments for data collection.

For example, ContikiRPL [35], the Contiki implementation of the IPv6 Routing Protocol for Low Power and Lossy Networks (RPL) [13] relies on ETX as a routing metric. Contiki-Collect is the default routing protocol of Contiki's rime stack [36] and it also uses the ETX. The Collection Tree Protocol (CTP) [14], which is the default routing protocol of TinyOS [37], does rely on ETX and on a composite metric, the Four-Bit. Nevertheless, the use of other metrics besides

2.3. State-of-the-art Acquisition of Link Quality Information

the confirmation of the packet reception in Four-Bit is limited to aiding the selection of a mote to be evicted when the neighbor's table is over its full capacity, by means of the "white bit" [26][14], instead of leveraging the composition of multiple metrics as a primary input for selecting the preferred parent, as is the intention of modern composite metrics such as TM or F-LQE.

One of the key barriers for the use of composite metrics in popular collection protocols is the lack of energy-efficient mechanisms for recovering metrics that are available on the receiver's side of a transmission. This occurs because routing protocols are optimized to obtain an accurate enough LE with few packet transmissions, thus enabling an acceptable balance between reactivity and low energy consumption.

A fundamental problem with ignoring the metrics available on the receiver side is that they characterize the quality of the link in the direction that a mote will use to forward the application data (uplink), which can be significantly different from the quality in the opposite direction (downlink) [20]. At the same time, several of the metrics on the receiver side are among the most promising for accurate LE, such as the RSSI, LQI and SNR.

Using these metrics in practical collection applications poses three main challenges: i) the information must be returned to the sender mote fast enough for it to still be representative of the link condition, ii) these metrics are subject to noise and require several samples to provide an accurate estimation, and iii) the estimation should entail a minimum transmission overhead in order to improve the energy efficiency and bandwidth utilization. Addressing these challenges is the focus of section 2.4.

2.3 State-of-the-art Acquisition of Link Quality Information

State-of-the-art collection protocols (e.g., CTP, RPL and ContikiCollect) broadcast beacons to probe the links to the neighbors and sample the link quality. The beacons include a list of neighboring motes and the respective qualities of the links leading to them. This allows the beacons to also serve as a mechanism for disseminating topology information and discovering new neighbors.

However, the beaconing represents a significant energy and traffic overhead, which protocols aim to reduce with the help of a Trickle Algorithm [15]. This scheme adapts the beaconing interval to the perceived stability of the network, thus progressively increasing it if the topology is working well, but quickly decreasing it upon detection of inconsistencies (e.g., loops), or of significant changes in the topology. The overhead reduction has been demonstrated by Gnawali et al. in [38], reporting that applying Trickle to a periodic beaconing protocol decreases the number of beacons sent per mote by 63.7% (from 550 to 200), after 5 hours. Moreover, the enhanced reactivity resulting from a low beaconing period enables the prompt bootstrapping, reparation or adaptation of the network.

However, this quest for energy efficiency through adaptive beaconing also means that the protocol generates an insufficient number of samples to ensure a meaningful accuracy of composite metrics. For example, an accurate estimation of intermediate links using LQI requires averaging over a large number of samples (about 40 up to 120) [39], but the quickest that a mote will obtain 16 samples is 1.16 hours (after an interval reset), and one extra sample every hour (for the default limits of the beaconing interval in CTP: 64ms - 1h). Sampling becomes even slower if the beaconing interval has reached the maximum value (1h), thus it would require 40 hours to obtain 40 samples. The long delays render the use of LQI for link estimation impractical, since motes would use this information for promptly adapting the topology to minimize disconnection times. This creates a clear trade-off between the energy efficiency of the protocol, reliant on the Trickle Algorithm, and its ability to enhance its reliability with the help of composite metrics.

Baccour et al. have validated the composite metric F-LQE by plugging it in the Collection Tree Protocol (CTP). However, they have adapted CTP by deactivating Trickle and keeping a constant beaconing interval. This permitted the acquisition of a sufficient number of samples required as input for their composite metric. Nevertheless, this setup is not feasible for a real application of a collection protocol, due to the high levels of energy consumption incurred by not using Trickle. Their work however focused on demonstrating the reliability of the metric, leaving aside its possible impact on the power consumption. The collection protocols MultihopLQI [40] and Arbutus + DUCHY [41] also rely on a constant beaconing interval in order to obtain sufficient LQI and RSSI samples.

Liu and Cerpa in [28] have validated the composite metric TALENT by plugging it in CTP. Their work has also recognized as a main limitation the fact that the resulting solution only functions in high data rate scenarios, due to the need of quickly gathering RSSI and LQI samples.

Another method often employed for validating composite metrics relies on the construction of a parallel infrastructure for the acquisition of data and the computation of the composite metric. Real experiments with composite metrics are thus often designed to provide results that are evaluated offline [42] [22] [29] [27]. This means that the raw link quality data is collected in each mote's memory, or using a secondary infrastructure (e.g., a cabled testbed), and the composite metrics are calculated and compared later on by a remote computer. This circumvents the problems of reliably recovering the link quality in a sink mote through a multihop network, and of returning the link quality metrics available upon reception to the sender mote. Additionally, the secondary infrastructure renders the adaptation of the beaconing interval in order to save energy unnecessary, thus facilitating the collection of many link quality samples.

Some state-of-the-art routing protocols (e.g., CTP, RPL and ContikiCollect) also rely on the application traffic for extracting link quality information. The application traffic normally consists of unicast packets in the uplink, which can provide bidirectional information about the link quality, thanks to the acknowledgement mechanism (e.g., the number of transmissions

required for sending a packet). Nevertheless, the usefulness of the application traffic is strongly dependent on the specific scenario, since it might be very scarce or dominated by sporadic surges, thus playing a minor role in updating the routing topology. The previous LE schemes share the limitation of extracting the metrics on received packets, thus evaluating the downlink direction, while requiring an assessment of the uplink for conveying traffic towards the sink.

A number of additional studies have touched upon the issue of integrating composite metrics in LE. Amongst them, the study reported in [43] aims at developing a mechanism tailored for collecting the link quality information necessary for using composite metrics, in the uplink direction. It presents LQR, an alternative to LE for identifying the best link available, with the minimum possible overhead packets, by probing the neighbors using broadcast packets and recovering the link quality information using unicasts. Instead of attempting to quantify the quality of the links, LQR compares several physical layer metrics and outputs a relative ranking among the links to the neighbors within transmission range of a mote. LQR proposes using fewer packets for ranking instead of several packets for estimation. Therefore, it is proposed for applications that require one-hop bursts, instead of long term routes. These applications are typically served by opportunistic routing protocols, such as ORW [44], and do not aim to perform long-term link estimation. While they provide a ranking at any given moment, they cannot be used to build routes, as they cannot compound the quantification of the link quality estimate along the route.

2.4 Proposed Protocol - WiseNE

In order to enable the use of composite metrics in collection protocols, while avoiding prohibitive energy expenses, and at the same time allowing for online network deployments, we have developed a mechanism called WiseNE.

2.4.1 Design Objectives

Enabling composite metrics in collection protocols demands a series of properties from WiseNE: i) returning the link quality that is available on the receiver side to the transmitter, the mote needing it for taking routing decisions, ii) providing a mote with multiple metrics simultaneously, as required by several of the most promising composite metrics available and iii) enabling the rapid extraction of link quality metrics that require several samples.

Our use of the neighborhood exploration (NE) to extract link quality is motivated as follows:

- Collection protocols already perform periodic NE to discover new motes and update the LE.
- The link quality extracted during the NE can be used for predicting the behavior of the links.

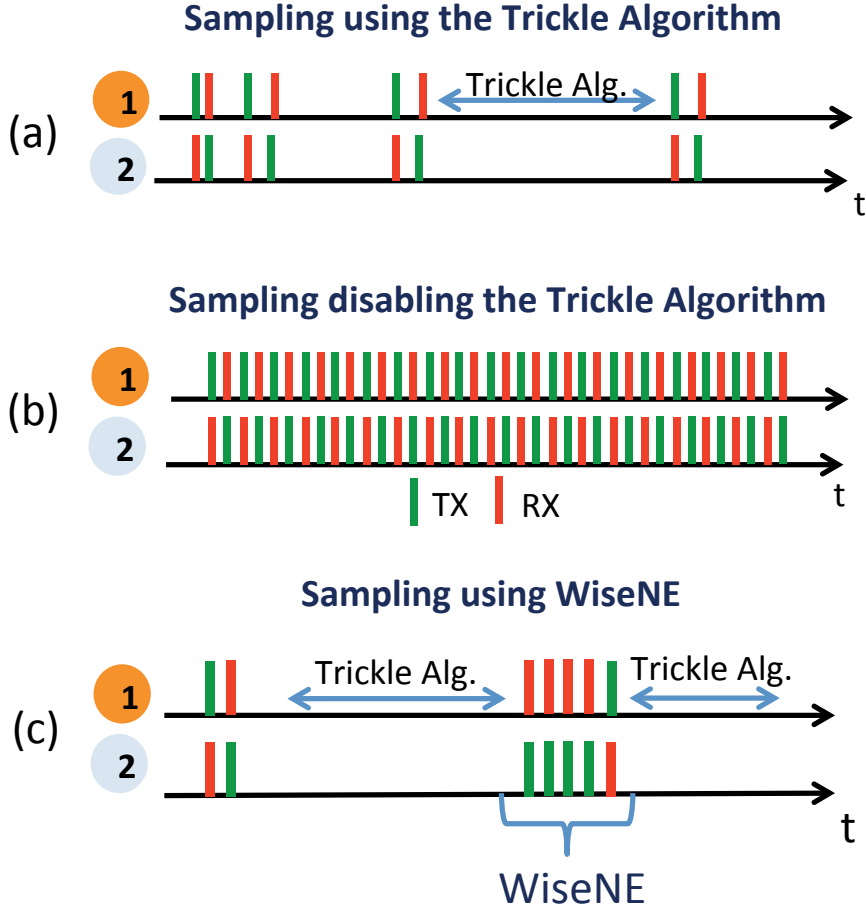


Figure 2.1 – Timing diagram of probing beacons for extracting link quality: (a) using the Trickle Algorithm (i.e. progressive increase of the interval between beacons if the network is stable), (b) deactivating Trickle (i.e. keeping the interval constant in a low value) and (c) using WiseNE (i.e. reducing the timing upon request and use Trickle otherwise).

- The energy and bandwidth overhead required for extracting the link quality can be reduced by one order of magnitude by using the packet repetitions available in low-power preamble sampling MAC protocols, such as WiseMAC [45] or ContikiMAC [46]. We elaborate on this solution in Chapter 3.

WiseNE enables the use of composite metrics by providing large numbers of link quality samples in a timely manner, while keeping the Trickle algorithm, which is the main driver of energy efficiency.

Before going into details about the operation of WiseNE, it is useful to remember the traditional approach for obtaining link quality samples (used by CTP and RPL, depicted in Fig. 2.1(a)): motes constantly broadcast beacons to evaluate the quality of the links, and the interval between consecutive beacons is increased progressively using the Trickle algorithm (upon detection of significant perturbations in the network, the interval is reset to a small value). This approach brings the shortcoming that collecting 16 samples will require 1.16 hours (after an

interval reset), and one extra sample every hour (as explained in Sec. 2.3). This long duration renders unfeasible the timely collection of 20 or more samples, as required for composite metrics.

An approach of several authors (detailed in Sec. 2.3) for accelerating the sampling is to keep the beaconing period constant (depicted in Fig. 2.1(b)). This approach provides the large number of samples required by composite metrics in a timely manner (e.g. 20 samples in less than a minute [25]). Nevertheless, it essentially deactivates Trickle, resulting in prohibitive energy expenses.

WiseNE's approach is depicted in Fig. 2.1(c), which shows the timing diagram of the beacon transmissions of a mote using Trickle. Therefore, at the beginning of the timeline the beaconing interval is long (the network has been stable since some time ago, but this is not a requirement for WiseNE). Later on (at the middle of the timeline), the mote needs to calculate a composite metric, therefore it triggers WiseNE, which consists in quickly broadcasting a series of beacons as required by the composite metric, promptly receiving the link quality estimates from the neighbours using unicasts, and immediately returning to the long beaconing period as required by Trickle.

An important advantage of WiseNE is highlighted in the phrase above: "promptly receiving the link quality estimates from the neighbours using unicasts", thus we will elaborate on recovery mechanisms. We have established that a mote that needs link information to choose a parent will transmit multiple broadcasts. Nevertheless, the most popular inputs for composite metrics are typically available upon reception (e.g. RSSI, LQI, PRR, etc.). Therefore, there must be a mechanism that enables the sender to recover the link quality data in a timely and energy-efficient manner. In the case shown in Fig. 2.1(a) (corresponding to CTP and RPL), the recovery is done by appending the link quality to the broadcasts of the neighbours (explained in Sec. 2.3), and a quick reply can be forced by resetting the Trickle interval of the neighbours (using a flag in the beacon header, such as the "Pull bit" in CTP). This approach brings an important energy penalty, which is avoided in WiseNE by using dedicated unicasts as recovery mechanism.

The end result is that the mote was able to calculate the composite metric, while Trickle was operating, thus maintaining an energy-efficient operation.

2.4.2 Design

The working principle of WiseNE is described in Figure 2.2, where it is depicted as part of a simple collection application. Mote N_6 explores the neighborhood by broadcasting a known number of Beacon Advertisements (B_A) (step 1). The neighboring motes will receive all or some of the B_A , and register the PRR and average RSSI, LQI and SNR (\overline{RSSI} , \overline{LQI} , \overline{SNR} , respectively). If the neighbor is a sink, or has a route to the sink, it will answer the beacons with a unicast acknowledged Beacon Reply (B_R), which contains an estimation of the link

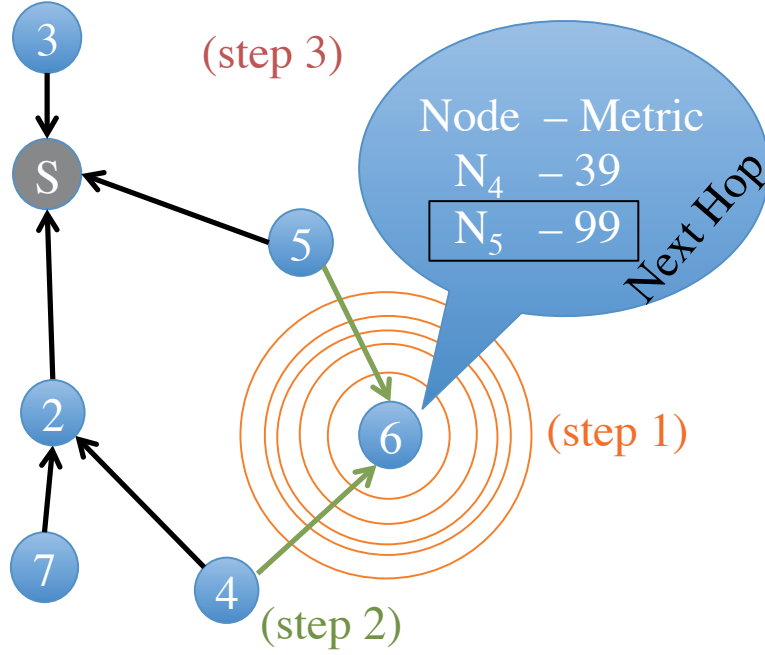


Figure 2.2 – Working principle of a collection protocol based on WiseNE.

metric (based on the B_A series) and its own route metric to the sink (step 2). The new mote will select the route with the lowest cost to the sink among the ones provided by the neighbors (step 3).

The following list details the WiseNE mechanism, and Figure 2.3 presents its timing diagram.

1. Advertising Mote (N_A) is looking for a route. It broadcasts n beacon advertisements with a period T_{BA} . The default values $n=20$ and $T_{BA} = 1s$, are chosen to match the evaluation parameters of the TM publication and ensure comparability. Each B_A has a sequence number s ($s = 1 \dots n$).
2. The neighbor N_R starts counting the B_A received. Only the sink or neighbors with a valid route to the sink will engage in this process.
3. Upon reception (RX) of the first B_A , N_R starts a time counter: $T_{BRI} = T_{BA} * (n - s)$ (Beacon Reception Interval, $s = 1$ for the first B_A sent). The purpose of T_{BRI} is to create a window where N_R expects to receive all remaining B_A . T_{BRI} is a function of s to account for the loss of the first beacons.
4. When T_{BRI} elapses, N_R has received a certain number (m) of B_A out of the n sent. This provides a value for the PRR of the link to the new mote:

$$PRR = \frac{m}{n} \quad (2.1)$$

This will also provide a value for the \overline{LQI} and \overline{SNR} (obtained by subtracting the RSSI of the beacon reception from the RSSI of the noise floor) of the B_A received, used to



5. When T_{BRRW} elapses, N_R unicasts a B_R to N_A . This packet will contain the metric of the link estimated by N_R and the metric of the route to the sink from N_R .
6. After sending the last B_A , the N_A will await the reception of BR during the interval T_{BRRI} (Beacon Reply Reception Interval). The interval T_{BRRI} is defined according to equation 2.2 to ensure that all the B_R were received. The interval T_{SB} (Safety Bound, default: $T_{SB} = 1s$) accounts for delays of the B_R due to re-transmissions.

When a B_R is received, the metric provided by N_R is compared to the best neighbor so far. If the metric is better, it will be declared the best neighbor. This change will not trigger a selection of the parent.

- Each B_A includes a field with the value n . This provides the option of doing WiseNE on demand. For example, N_A could report an $n = 1$, to send a single beacon for routine topology dissemination, but later use an $n = 20$ to quickly obtain an up-to-date estimate of \overline{LOI} .

¹The maximum value of 2s is dimensioned to avoid congestion with a MAC sampling period of 256ms and 8 neighbors. It must be dimensioned proportionally for other sampling periods.

1. Obtaining link quality of the uplink, decoupled from the downlink: This feature aims to reduce the probability of a mote accepting a parent that has a link with good downlink quality, but that it is unreachable for the uplink traffic. This problem has already been reported in RPL [47] and the subsequent re-routing results in delays, traffic overhead and, potentially, packet losses.
2. Enabling the collection of link quality metrics of different natures: WiseNE enables the collection of link quality metrics that are typically available upon reception, such as the uplink PRR, RSSI, SNR and LQI.
3. Supporting multiple samples: The extracted link quality metrics are the averages of the samples taken from the B_A series, which makes the LE process more robust to random variations in link quality.

2.4.3 Composite Metric

The Triangle Metric

We have chosen to use TM as a metric in WiseNE in order to illustrate the versatility of WiseNE in terms of enabling the use of composite metrics. This is due to the fact that TM poses a series of requirements that make its implementation particularly challenging, such as: i) simultaneous measurements of three metrics (SNR, LQI and PRR), ii) averaging over multiple samples and iii) the need for the sender to recover the link quality as soon as possible, and with a high reliability.

Calculation

The stream of B_A allows a neighbor to gather link quality metrics. In the current implementation, this mechanism gathers RSSI, LQI and SNR from each B_A , plus an estimation of the PRR by comparing the number of expected vs. received B_A . These metrics can be combined in different ways in order to characterize the link. This work draws on the TM presented in [22], which uses the LQI, SNR and PRR of a link as inputs in order to predict the PRR in the near future. It thus manages to leverage the individual strengths of the input metrics to provide a holistic characterization of the link. Moreover, its results are promising in both static and mobile networks. The TM is calculated through equations (2.4-2.6). Equations 2.4 and 2.5 define "window averages", which are averages over the total number of expected packets, instead of the number of received packets. This technique is suggested in [22] to merge the PRR with the SNR and LQI, respectively. Parameters m and n are defined in section 2.4.2.

The SNR of a packet was calculated according to equation 2.3, using the RSSI (S_{dBm}) and the Noise Floor (N_{dBm} , obtained from sampling the RSSI once after reception). This methodology

has been previously used in similar studies, such as [22], [43] and [48].

$$SNR = S_{dBm} - N_{dBm} \quad (2.3)$$

$$\overline{SNR}_w = \frac{\sum_{k=1}^m SNR_k}{n} \quad (2.4)$$

$$\overline{LQI}_w = \frac{\sum_{k=1}^m LQI_k}{n} \quad (2.5)$$

$$TM = \sqrt{\overline{SNR}_w^2 + \overline{LQI}_w^2} \quad (2.6)$$

2.5 Experimental Results

2.5.1 Methodology

The experiments presented were performed in order to validate the accuracy of the link quality extracted through WiseNE, for predicting the PRR. They followed the methodology established in [22] to study the correlation between present values of LE metrics and the future PRR values. The target system consisted of two motes, labeled N_S (sink) and N_A . This simple system was chosen to avoid collisions with traffic from other motes that might perturb the LE. Mote N_A performs an NE every 5 min, which consists of sending $n = 20$ Beacon Advertisements with $T_{BA} = 1s$. The $n = 20$ value was chosen because it represents an acceptable trade-off between minimizing the transmission overhead and having a less than 10% granularity in the PRR metric. As a first assessment, it is important to avoid a large granularity in order to better understand the prediction accuracy. The $T_{BA} = 1s$ is kept the same as in the reference experiments from [22].

Next, N_S sends a B_R with a maximum of 20 retransmissions. This value has been overestimated to make the reception of the B_R highly likely, even if the link has a bad quality. An alternative approach suggested in [25] is to send it with a higher transmission power. The traffic generated by WiseNE represents the control traffic.

Mote N_A also sends to N_S a unicast application packet every 5s, which represents the application traffic. This is sent without retransmissions, in order to study the correlation between the PRR measured during the NE and the PRR of the application traffic. Each application packet contains the link quality recovered by N_A from the last NE. This scheme is representative of the control and application traffic that would simultaneously appear in a deployment.

The application packets obtained at the sink provide the raw data used for analysis. Each NE extracts the information that N_A would use to perform an assessment of the current state of the link during an observation window ($w_o = 20s$). The application traffic after each NE is

used to study the future behavior of the link. It is analyzed during a prediction window of 100s ($w_p = 100$ s). The size of the windows is selected to allow comparisons with the reference publication of TM.

The experiment was repeated in several environments. The label *indoor* refers to an office hallway with metallic walls (distance between the two motes = 10m, transmission power = -25 dBm). The label *outdoor* refers to an urban environment (separation = 30m, transmission power = 0 dBm). The distances were selected close to the transitional region of the links, i.e. the limit where data traffic is received by N_S in order to study the link subject to temporal changes in reliability. In both variants, the model of the motes is TelosB (Crossbow, 2.4 GHz), operating on channel 26 (2.48 GHz), which is known to display low interference due to 802.11b and Bluetooth networks [49]. WiseNE was implemented in Contiki v2.7 [50] and uses ContikiMAC with the default parameters (sampling period = 125 ms), unless otherwise noted.

The *indoor* experiment lasted 13 days (April 19 - May 2, 2016), and the *outdoor* experiment lasted 13 days (May 11 - May 24, 2016). There are 1205 and 1339 points, for the indoor and outdoor cases, respectively. The relative low number of points is explained by periods when the motes would not receive any packet due to link degradation.

2.5.2 Evaluation

Figure 2.4 displays the scatter plots of \overline{LQI} vs. \overline{SNR} (using the arithmetic average, displayed in the left plot) and \overline{LQI}_w vs. \overline{SNR}_w (using the window average explained in Section 2.4.3, displayed in the plot on the right), for the *indoor* (top) and *outdoor* (bottom) deployments. Each point summarizes the information extracted during a single NE. The color of each point signals the link quality, classified according to the PRR interval (defined in p. 13): black (bad), red (intermediate), blue (good), green (very good). We proceed without normalizing LQI (range of values: 0 - 110) and SNR (typical range of values: 0 - 40), even though their ranges differ significantly, in order to follow the original definition of TM's authors in [22] and enable comparability with their results. Nevertheless, we believe that future work with TM should use a normalized version.

In the right side plots, the TM is equal to the Cartesian distance from each point to the origin, as defined in equation 2.6. The results show that the use of the window average in the plots of the right side improves the classification accuracy of the PRR of each point based on the TM. They also confirm the observation made in [22], namely that the PRR carries important link information that is not captured by the \overline{SNR} and \overline{LQI} , and that using a window average to include the PRR information provides a more accurate estimation of the link quality. In the context of a routing protocol, the figure evidences that a mote requiring the selection of a routing parent can use WiseNE to classify and compare the current quality of the links to the possible parents.

The figures also show that \overline{SNR}_w and \overline{LQI}_w intervals associated with each PRR category are

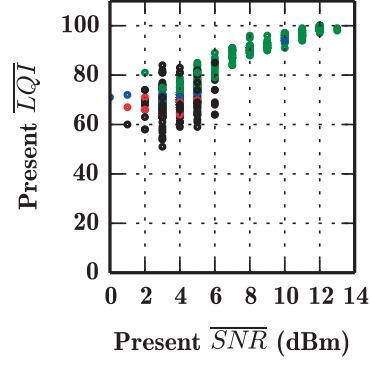
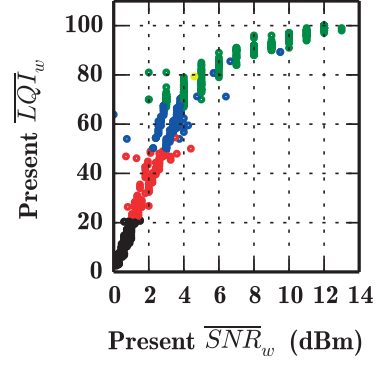
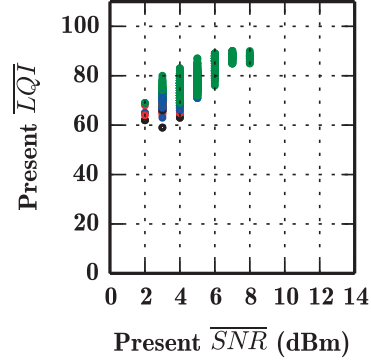
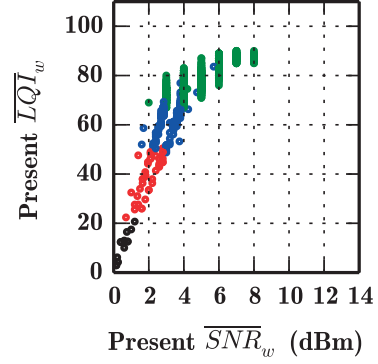

(a) *indoor*, arithmetic average

(b) *indoor*, window average

(c) *outdoor*, arithmetic average

(d) *outdoor*, window average

Figure 2.4 – The assessment of the link quality improves through the use of a window average and it is consistent in *indoor* and *outdoor* conditions. The color of each point signals the link quality (PRR interval from the NE): black (bad), red (intermediate), blue (good), green (very good).

approximately the same for *indoor* and *outdoor* conditions. In general, they are smaller than the values reported in [22]. These differences are attributed to the use of observation windows with different sizes, but no further comparison is possible given that the size of the observation window is not reported in [22] for the analogous figure.

The \overline{SNR} , \overline{LQI} , \overline{SNR}_w and \overline{LQI}_w values obtained from WiseNE have a granularity of 1 unit. The reason is that the input information is averaged in the micro-controller before sending the B_R and represented as integer values. This process differs from the one used in [22] in that Boano et al. log the information from each sample in each mote and perform the averaging operations offline with the possibility of using floating point values. Nevertheless, the results from [43], where a similar system is compared with and without the use of floating-point values, suggest that the use of floating-point only provides a minor benefit in the comparison of the quality of multiple links when PRR, LQI and SNR are used as inputs. The granularity of the current and future PRR in this study is determined solely by the size of the observation and prediction windows, respectively.

Figure 2.5 displays the future PRR (calculated over w_p) vs. the current value of several metrics (calculated over w_o). w_p starts immediately after w_o . It represents the prediction capabilities of several metrics that a mote can obtain from WiseNE. The area of each dot is proportional to the number of data points with the same coordinates. Figure 2.5a shows a wide range of future PRR values associated with very high current PRR ($> 90\%$). This behavior is due to periods when the link maintains a good quality and it is stable (PRR = 100% despite external effects) and periods when the quality degrades due to minor environmental changes, such as changes in the multipath fading pattern due to the movement of people [39]. The lower current PRR values ($< 80\%$) are associated with future PRR values generally lower than 35%, which supports the use of PRR to predict the bad links. These observations agree with the results from [22].

Figure 2.5b shows that the SNR values over 5 dBm are associated with future PRR values of 100%, which shows that the SNR can be used to identify very good links and predict their behavior. The lower SNR are associated with the full range of future PRR values, therefore they do not provide a useful insight to classify or predict the link. An analogous observation can be made from Figure 2.5c, which shows that LQI values over 85 are highly correlated with future PRR values of 100%.

However, the LQI scatter plot (Figure 2.5c) shows a smoother decay in the transitional region compared to the SNR plot (Figure 2.5b). Previous studies have performed this comparison and obtained similar results [16] [22]. This supports the use of LQI for predicting intermediate and good links, instead of SNR. Which metric is better is still an open discussion [23].

Previous studies have shown that the variation of the LQI increases significantly in the transitional region [16], and that LQI can predict the behavior of intermediate links only if it is averaged over large observation windows (40-120 packets) [39]. This suggests that it is possible to reduce the width of the transitional region (i.e., the variation of LQI measurements) if more

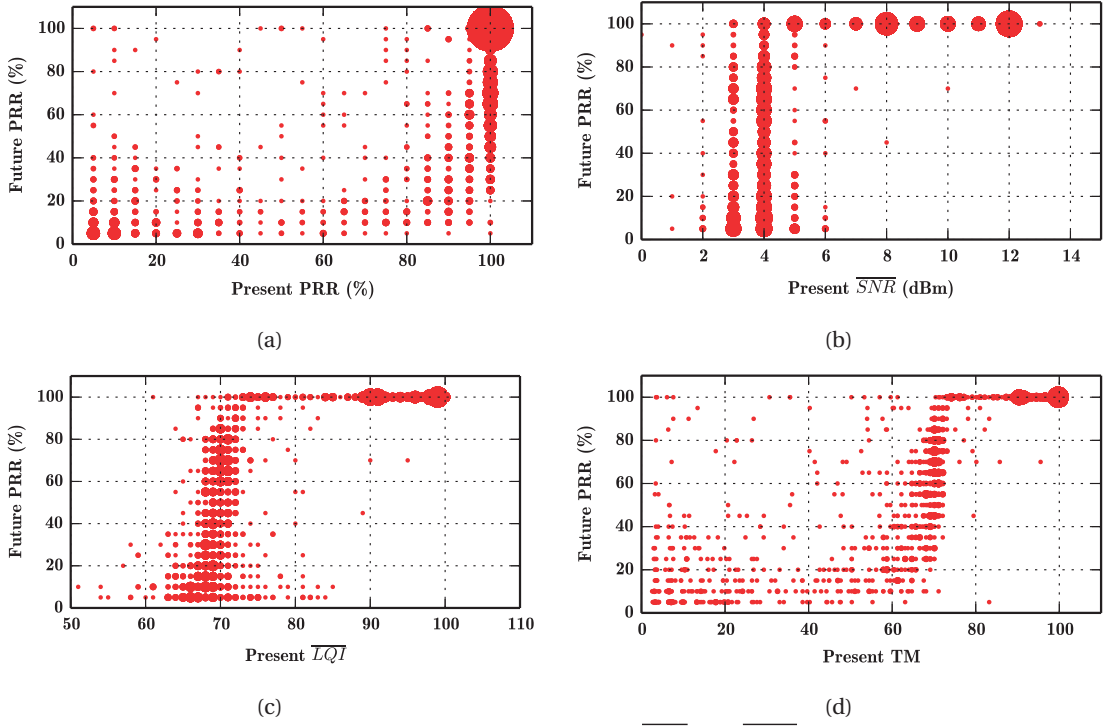


Figure 2.5 – Future PRR vs. the current value of PRR (a), \overline{LQI} (b), \overline{SNR} (c) and TM (d) for *indoor* conditions. These scatter plots provide insight about the prediction capabilities of each metric. The area of each dot is proportional to the number of data points with the same coordinates.

samples are used in w_o . This behavior is further studied in Chapter 3.

Figure 2.5d displays the future PRR vs. present TM. The low PRR variation for $TM \geq 75$ suggests that the TM can identify and predict the behavior of very good quality links with a high accuracy. The PRR variation is higher for $TM \leq 60$ with $PRR \leq 35\%$, which indicates that the TM can identify and predict the behavior of bad links with a lower accuracy than it does with very good links. The plot also shows the transitional region for $60 < TM < 75$, where $20\% \leq PRR \leq 100\%$.

The high variation of the PRR and the low accuracy in predicting the link quality are typical for the transitional region [18]. Nevertheless, the steep decline of the scatter plot in the transitional region makes it difficult to use it for predicting links with intermediate and very good qualities. Given that the SNR and LQI are known to have a significant variance in the transitional region [16], and that they provide the prediction capabilities of the TM for good and intermediate links, it can be expected that the prediction accuracy for these categories will improve if a longer w_o is used.

The prediction accuracy of the TM for bad and very good links can be quantified with the following bounds: $TM > 75$ for a very good link and $TM < 60$ for bad links. This results in 75% of the bad links (from 398 points) and 98% (from 343 points) of the very good links correctly predicted. The high accuracy for predicting good links surpasses the accuracy of the other

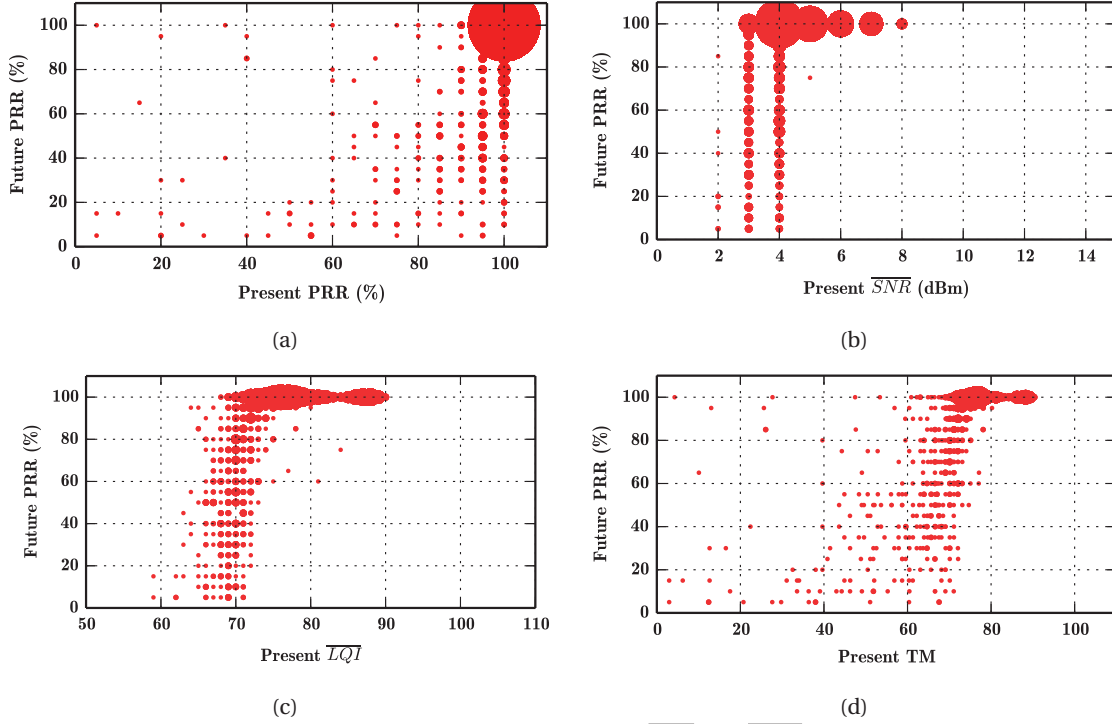


Figure 2.6 – Future PRR vs. the current value of PRR (a), \overline{LQI} (b), \overline{SNR} (c) and TM (d) for *outdoor* conditions. The weaknesses and strengths of each metric to predict the PRR agree with the results from Figure 2.5. The area of each dot is proportional to the number of data points with the same coordinates.

metrics analyzed for the same task: PRR 82.3%, SNR 87.4% and LQI 90.3%. Therefore, TM brings an improvement of 7.7%-15.7% over state-of-the-art metrics. Notably, the highest improvement (15.7%) is with respect to PRR, which is the input for the most commonly used metric, ETX. These observations agree with the report from Boano et al. [22], where TM was found to improve PRR in the prediction of very good links by 18%, while also exceeding the performance of SNR and LQI, with accuracies 94% and 89%, respectively.

These results confirm that the TM is able to merge the information in a way that leads to a more robust estimation compared to the input metrics taken individually. Moreover, the results suggest that it is possible to use the TM for predicting the behavior of the links in the near future and it could be used by a mote to compare the links to the neighbors and select the one that promises the highest PRR.

Figure 2.6 displays the future PRR vs. the current value of several metrics for outdoor (urban environment) conditions. The individual figures display the same prediction properties observed for the indoor case (Figure 2.5). Moreover, the TM intervals that define very good and bad links are found to be approximately the same for both environments.

Figure 2.7 displays the accuracy of the PRR prediction for the very good (green trace) and bad (black trace) categories as a function of the temporal horizon (time between the end of w_o and the beginning of w_p). The plot shows that the prediction that a link will be very good holds a

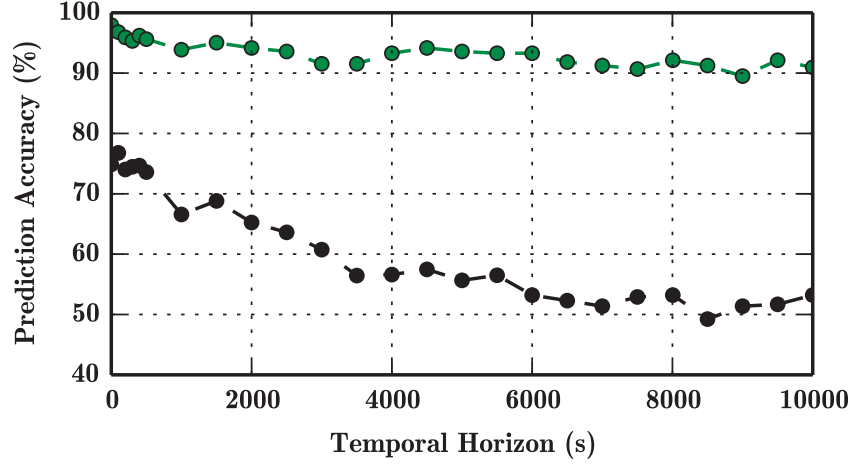


Figure 2.7 – Progression of the prediction accuracy as a function of the time horizon. Each trace represents the accuracy of the TM to predict the behavior of a type of link: bad (black trace) or very good (green trace).

high accuracy ($\sim 90\%$), even 10000s after the prediction. This demonstrates the capacity of the TM to correctly predict that a link will be very good and stable in the future. This feature is very important for collection protocols because the selection of a link that is unstable implies that packets will likely be lost in the future, and the protocol will need to redefine some of the routes, a very expensive operation in terms of energy and bandwidth [14] [43].

In the case of the bad link category, the accuracy of the prediction decreases significantly over time, and achieves a value close to 50% at 10000s after the prediction. This means that a TM that predicts that the link will behave as a bad link in 10000s has a prediction accuracy of 50%, and that link will likely behave as if it belonged to a different category. This agrees with previous studies showing that bad links have PRR values that vary significantly over time [39].

2.5.3 Difficulties Encountered

The experimental data employed was collected using the Collect application in ContikiOS, which by default uses the serial port of a sink mote as the sole data outlet. Using this popular public code base led us to rely on application packets to convey the data to the sink (scheme described in Section 2.5.1). The first steps of this study could have been improved in terms of simplicity, by connecting all the motes to a computer using the USB port. Thus, would have enabled the output of experimental data directly from each mote by using the serial port. This alternative setup was however only evident to us at a late stage of the experiment, upon studying the data collection techniques in public testbeds, such as Flocklab [51] and Indriya [52].

Another approach that could have facilitated the first stages of the study would have been validating the advantages of the proposed mechanism in publicly available offline traces, such as the dataset from the Life Under Your Feet (LYUF) project [53]. This dataset provides

the information collected over 225 days by WSNs deployed in forests in order to study soil properties. This approach accelerates the early stages of studies by providing insights into its performance, at a minimum effort, as shown by [54] and [30].

The experiments presented in this chapter would have also benefited from an easy extension aimed at studying the effect of the interference with wireless networks. The extension could have consisted in repeating the experiment in multiple channels, as some channels are known to be very susceptible to WiFi and Bluetooth sources, while others are known to be immune, as exemplified by [49] and [54].

2.6 Conclusion

The latest research in composite metrics has shown promising potential to improve the delivery rate and the power consumption of multihop WSNs. Nevertheless, providing the link quality information to the nodes that need it for taking routing decisions is an important challenge that must be overcome in order to integrate these metrics in collection protocols and leverage them in practical WSN applications.

We have addressed this challenge by proposing WiseNE, a novel neighborhood exploration mechanism that promptly provides inputs for composite metrics in the uplink, supports multiple samples, and is at the same time compatible with the energy-saving mechanism Trickle. The benefits that WiseNE can provide to collection applications were validated by performing, to the best of the authors' knowledge, the first online verification of the prediction capabilities of the Triangle Metric, as well as the first outdoor evaluation.

The results show that WiseNE can provide the inputs required for benefiting from the high prediction power of the Triangle Metric in an online setting. The results also confirm the potential of the TM to combine the strengths of the SNR, LQI and PRR for predicting the quality of the links in office and outdoor environments. Besides enabling TM, WiseNE can reasonably be expected to be compatible with other composite metrics as well.

3 Leveraging MAC Preambles for an Efficient Link Estimation

3.1 Introduction

Collection protocols' ability to identify and make use of the most reliable routes available relies on the computation of a Link Quality Estimation (LQE). The LQE thus has a fundamental impact on the network performance, affecting the design of higher-layers of the protocol stack [23], and constituting a central piece of state-of-the-art collection protocols, such as RPL [13], Contiki-Collect [35] and CTP [14].

As shown in the previous chapter, making use of physical metrics, such as SNR (Signal-to-Noise Ratio) and LQI (Link Quality Indication), on top of the traditional PRR (Packet Reception Ratio), in the computation of the LQE results in significant accuracy, and thus WSN reliability increases. However, the accuracy improvement potential of RSSI (the input for SNR) and LQI, and thus composite metrics relies on their access to multiple samples. The protocol presented in the previous chapter, WiseNE has managed to enable the use of multiple samplings in a way that remains compatible with the Trickle algorithm, thus avoiding the otherwise expected prohibitive increase in the WSN's energy consumption.

However, state-of-the-art algorithms, including WiseNE, still operate by extracting a single sample of the link state with each individual packet. Thus, increasing the number of samples in order to mitigate the interference and obtain higher accuracies of the LQE would necessitate a proportional increase in the number of packets. This translates in an increase of the energy overhead, as well as the traffic and the duration of the LQE operation [23] [43] [22]. Moreover, this accuracy-energy consumption trade-off creates additional challenges in mobile WSNs, since the high variability of the medium imposes more pressing constraints on the duration of link estimation [22].

In this chapter, we propose a mechanism to overcome this trade-off, which we have named Rep. Rep manages to obtain the high number of samples necessary for a reliable computation of the RSSI and LQI, while lowering the number of packets transmitted, and thus the energy consumption. The novelty of this sampling mechanism lies in its ability to leverage the already

existing packet repetitions in the preamble sampling medium access control (MAC) protocols in order to extract the link quality. The experimental evidence provided in this chapter shows that Rep can reduce the traffic, energy consumption and duration of a link estimation by one order of magnitude, without compromising the metric's accuracy.

The chapter is structured as follows: Section 3.2 briefly describes the problem at hand, followed by Section 3.3 which positions the chapter's contributions with respect to previous work, and elaborates on the limitations of current link estimation schemes. Section 3.4 presents the design and implementation of the proposed sampling mechanism, Rep, including a summary of its advantages and an explanation of its possible limitations. Section 3.5 compares Rep with the state-of-the-art on the basis of both an analytical model (in subsection 3.5.1) and a real testbed evaluation (in subsection 3.5.2). Section 3.6 highlights other domains that can benefit from using Rep, followed by Section 3.7 which elaborates on the difficulties faced during this study, and Section 3.8 which concludes the chapter.

The work presented in this chapter has been accepted for publication in WiMob 2018 [55].

3.2 Problem Statement

This chapter addresses the trade-off between the high number of samples required for LQE in order to leverage RSSI for improving the protocol's accuracy, and the increased energy consumption, delays and overhead caused by increasing this number of samples.

State-of-the-art protocols, including WiseNE, operate by extracting a single sample of the link state per packet. Increasing the number of samples in order to enable the use of RSSI in the calculation of the LQE would thus translate in a high number of packets. However, probing the channel with multiple packet transmissions in order to extract the relevant information for LQE also increases the energy overhead, traffic and the duration of the LQE operation. These factors thus become a limitation to the number of samples obtainable, hindering the use of promising information sources. Moreover, this trade-off creates additional challenges in mobile WSNs, since the high variability of the medium imposes more pressing constraints on the duration of link estimation [22].

3.3 Related Work and State-of-the-art

The problem of accurately and efficiently estimating link quality in WSNs is considered an essential [23], yet open research challenge [25]. For example, several studies suggest that the Link Quality Indication (LQI) metric is a promising LQE [39] [56] [57]. Nevertheless, LQI needs to be averaged over a high number of samples (40 up to 120) in order to provide a result accurate enough to be used by routing protocols [39] [22]. More specifically, each node in a WSN operating with this metric would need to transmit 40 to 120 packets for a single estimation, and to repeat this process periodically, generating a prohibitive energy and traffic

overhead. Additionally, the delay caused by obtaining this information hinders the taking of critical decisions in the upper layers, such as bootstrapping and repairing the topology, potentially resulting in packet losses.

Common approaches to improving the LQE attempt at either increasing the accuracy of metrics that can be obtained from the radio hardware or from the reception statistics with a low overhead, or at combining multiple such metrics into composite metrics ([23] presents a comprehensive review of these efforts). Nevertheless, the information required by many of these metrics is often not available in commonly used WSN protocols. The most cited reasons include the prohibitive number of samples necessary for an accurate calculation of the RSSI metric [39], the fact that some metrics' inputs are only available upon reception, thus of no use to the sender [23], or the fact that the different inputs necessary for the computation of the composite metric are not simultaneously available [25].

A previous approach at solving these limitations is F-LQE [25], a link estimator using fuzzy-logic to combine multiple link quality metrics. Its authors also proposed several design changes that must be implemented in Collection Tree Protocol in order to extract the required input information. F-LQE necessitates many samples of RSSI and LQI, which the authors solve with frequent beaconing (every 1s). The resulting high energy consumption limits the practical applications of F-LQE.

Additionally, [43] propose LQR, a novel paradigm for coping with unreliable links. Instead of performing a link quality estimation which would require multiple packets, LQR only performs a ranking (i.e., identifying the best link available at that given moment), thus needing less packets. LQR operates by probing all the neighbors of a mote with multiple broadcasts, receiving their replies with the LQE, and comparing these inputs in order to create a ranking. Therefore, it significantly improves the efficiency in terms of the number of sample packets, at the expense of not quantifying link quality. The authors also address the problem of designing a tailored mechanism for extracting the required input information. However, by only performing the ranking and not the estimate, LQR cannot support collection tree protocols, perhaps the most widespread enabler of WSNs. The reason is that these protocols require quantifying and comparing routes, which is an operation that necessitates the estimation of the links along each route.

The shared concern for the energy and traffic cost of link quality estimation, as evidenced by multiple efforts to minimize the transmission overhead [43] [22], has also stimulated the research on fine-granularity link estimators. An example of such an effort is presented in [39], which studies the estimation accuracy of LQI and RSSI with sampling periods of 10ms (the lowest period experiments found in the literature). To the best of our knowledge, the studies that analyze the accuracy of link estimation with a sampling period inferior to 100ms, do not develop a technique to integrate such fast sampling in MAC protocols. The proposed mechanism, Rep will build upon these efforts and propose a concrete mechanism to integrate these link estimators in low power MAC protocols and use them in practical applications.

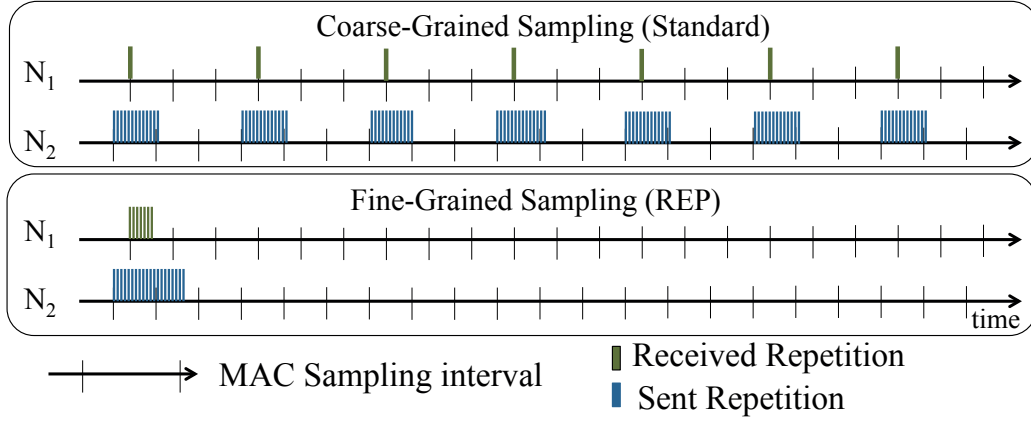


Figure 3.1 – Comparison between coarse-grain (top) and fine-grain (bottom) link estimation. The fine-grain link estimation enabled by Rep results in significantly less delay and overhead than the standard coarse-grain approach.

Additionally, state-of-the-art protocols tend to isolate the experimental system from external disruptions such as variations in the multipath fading pattern and wireless interference. In the frame of this study, we have chosen to keep the system subject to stress in order to analyze the effect of each perturbing factor. This approach is not commonly reported in link estimation literature.

3.4 Proposed Protocol - Rep

3.4.1 Key Features

Rep extracts useful link quality statistics, such as the average LQI and RSSI, from a single packet in low-power preamble sampling MAC (LP-MAC) protocols (ex. B-MAC [48], WiseMAC [45] and ContikiMAC [46]) (Fig. 3.1, bottom). This is achieved by leveraging a resource that is commonly available in state-of-the-art protocol stacks, yet so far under-used, namely the fact that LP-MACs transmit the exact same packet several times (hereafter referred to as packet repetitions) in order to wake the intended receiver(s) up.

At this point, it is useful to understand better the role of the packet repetitions. We will use the example of B-MAC [58], which is an LP-MAC protocol used by default in TinyOS' stack. B-MAC uses radio duty cycling to significantly decrease the time that the radio remains active, which is the main driver of energy consumption in a mote. This technique is also called low power listening, and results on activating the radio periodically during a short interval, named channel sampling. On each wake up, the mote turns on the radio and checks for activity. In case it detects activity in the channel, the mote remains awake until the incoming packet is received, afterwards it returns to sleep. A transmitter will add a preamble as long as the channel sampling period in front of each transmitted packet. This ensures that the intended receiver(s) will wake up before the actual transmission of that packet.

In more recent LP-MACs, such as a WiseMAC or ContikiMAC, the preamble consists of multiple

repetitions of the same packet. Both protocols use a preamble with a duration equal to the channel sampling period for broadcast transmissions, which wakes up all the motes within range to receive a packet repetition. In case of unicast transmissions, the phase-lock optimization (proposed in WiseMAC [59]) enables reducing the preamble length, by learning the offset between the wake-up schedules of a mote and its neighbours (from previous transmissions).

Estimating the quality of the link from each of these repetitions enables a reduction in the number of packets required for link estimation, and consequently the scaling down of the energy consumption. Moreover, it significantly reduces the duration of the link estimation. This study analyses the impact of sampling with Rep over the accuracy of the link estimation, as compared to the state-of-the-art coarse-grain approach.

Throughout this chapter, the term *packet* is used exclusively in the sense of the OSI model, namely the data transmission unit at the Network layer, while the transmission unit at the MAC layer is named *frame*. In the case of LP-MAC protocols, each packet is encapsulated in a frame and transmitted in multiple instances, that we refer to as repetitions.

3.4.2 Implementation

The repetitions from unicast, anycast and broadcast packets can be used for link estimation. This study explores the use of broadcasts, mostly given the fact that they rely on a large number of repetitions, and the fact that they are ubiquitous in routing protocols for discovering new nodes and routes.

The integration of Rep in LP-MACs is demonstrated through ContikiMAC-R, described in our previous study [60]. ContikiMAC-R is based in ContikiMAC, modified as follows:

- The sender uses a MAC flag ($F_{LE} = \text{true}$, LE comes from Link Estimation) to identify a Rep broadcast, i.e., it will provide the link quality to the receiver. With Rep, the broadcast transmission is prolonged by a time w_m (see Fig. 3.2). The extension corresponds to the time needed to send m repetitions, and it is necessary to prevent a mote from receiving less than m repetitions as a consequence of waking up at the end of the series. Receiving less repetitions can hinder the accuracy of the resulting LE; thus the extension has the role of decoupling the resulting estimation from the wake-up offset.
- Upon reception of a broadcast repetition, the radio keeps listening during a time window w_m . If the link has a high quality, it will receive at least m repetitions.
- When $F_{LE} = \text{false}$, ContikiMAC-R operates as ContikiMAC. Hence, a broadcast transmission will consist of enough repetitions to fill a sampling period (T_w) (plus one extra repetition), and only one repetition will be received (if successful).

The following events illustrate its operation with $m = 10$ (Fig. 3.2):

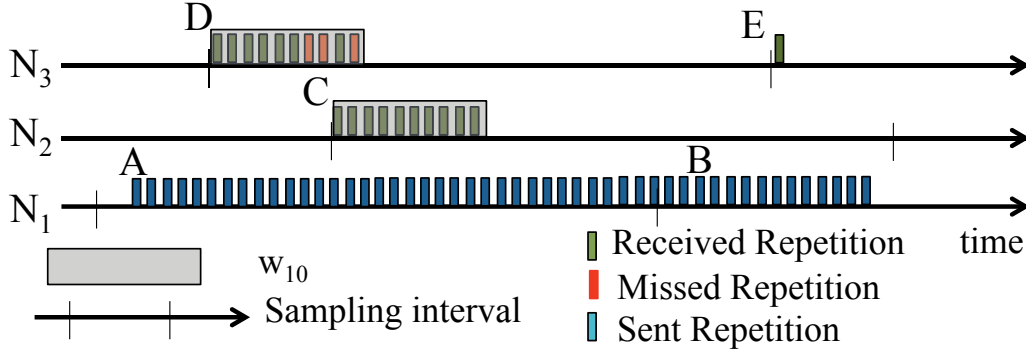


Figure 3.2 – Rep's working principle: N_2 and N_3 overhear several MAC repetitions from the packet broadcasted by N_1 , this samples can be used for fine-grain link estimation.

- (A) Node N_1 broadcasts a beacon by sending repetitions to fill a sampling period ($F_{LE} = \text{true}$).
- (B) Since $F_{LE} = \text{true}$, N_1 extends the broadcast time by w_{10} .
- (C) N_2 wakes up and receives one of the repetitions. Since it is a broadcast and $F_{LE} = \text{true}$, it extends its listening interval to w_{10} . During the window, it receives ten out of ten repetitions (PRR=100% for this link, PRR is the Packet Reception Ratio).
- (D) N_3 receives seven out of ten repetitions (PRR=70%).
- (E) N_3 wakes up and receives a repetition from the packet already received in (D) for link estimation, hence N_3 does not extend the listening interval, as there is no guarantee of receiving at most m packets. Otherwise, reusing the packet might provide an erroneous link estimation if there are less than 10 repetitions left for reception.

In case the receiver misses the repetition "in the air" during the channel sampling, it will miss the entire packet and not be able to extract the link quality. The probability of this event happening is the same when Rep is enabled or disabled, as the number of channel samplings of each node able to receive a particular broadcast packet stays constant.

3.4.3 Advantages

By leveraging packet repetitions to extract the link quality, Rep uses a sampling period in the order of milliseconds, resulting in a fine-grain (FG) sampling (Fig. 3.1, bottom). This is faster than traditional link estimation schemes that use a sampling period in the order of seconds (Fig. 3.1, top), hereafter referred to as coarse-grain (CG) sampling. This feature provides, by design, the following advantages:

- **Reduction of the duration of link estimation by one order of magnitude:** Rep avoids the delay for extracting multiple link quality samples. This delay is due to contention for accessing the channel, and depends on the network density, thus ranging from 100ms

to seconds. The result of using Rep is a significant increase in the speed of the link estimation (section 3.5.2 further quantifies this improvement). Moreover, this speed increase holds for different routing protocols, irrespective of the transmission primitive they use for link estimation (i.e., broadcast, unicast or anycast). This happens as the core saving in terms on speed comes from circumventing the delay in accessing the channel, which is determined by the sampling period (T_w), and the subsequent backoff delay.

Moreover, the presented fine-grain approach can be used in collection protocols that take advantage of the sporadic high quality intervals of unstable links. These schemes necessitate a mechanism to quickly identify a good link in order to use it before it degrades (e.g., [43]). Rep can improve the speed and accuracy of the required link characterization.

- **Reduction of the overhead (traffic) of link estimation by one order of magnitude:** Rep requires a single packet to extract multiple link quality samples, compared to CG schemes which only extract one sample per packet (even though a packet will anyway be composed of multiple repetitions). Therefore, Rep results in a significant overhead reduction in terms of traffic.

Another relevant factor are the possible changes in the distribution of packet collisions due to the reduction of the sampling timescale. Even though this demands a specific analysis, which is out of the focus of this chapter, we highlight that the reduction of the traffic with respect to the CG case should decrease the probability of collisions.

- **Increase in the number of samples:** The fast probing enables the collection of a volume of samples otherwise prohibitive for LQEs based on coarse-grain sampling. For example, Rep can collect 120 samples in less than a second, as compared to the duration in the order of minutes required by coarse-grain schemes (consider a MAC sampling frequency of 2 Hz, details explained in Sec. 3.5.2). The difference is even more significant if we consider the channel contention due to multiple nodes within range, all required to transmit for link estimation purposes.
- **Combined use of Rep and CG:** We note that both FG and CG sampling approaches can be used simultaneously by extracting the link quality from multiple repetitions of multiple packets. Therefore, in situations where the fine-grain sampling hinders the accuracy of a particular LQE, the information obtained can still be used to complement the coarse-grain scheme. This approach provides a more complete description of the link state, and we expect it to improve the accuracy compared to the sole use of coarse-grain information.

3.4.4 Other Issues

The influence of fine-grain link dynamics over Rep data must also be considered, as links of intermediate quality are known to have a bursty behaviour [23]. This means that they

switch between extreme qualities (good or bad) in the order of seconds [49]. Therefore, the link metrics behave differently as the interval between samples decreases (the granularity becomes finer). For example, the PRR tends to assume extreme values (closer to 0% or 100%), as the granularity becomes finer. Hence, it is not a good indicator of intermediate links at this time-scale [23]. Conversely, several studies suggest that the LQI is able to accurately characterize intermediate links even when the channel is probed during time-scales in the order of 10 ms [39].

3.5 Experimental Results

3.5.1 Energy Considerations

The following analysis is meant to validate the energy reduction provided by the LE operation with Rep, compared to CG. We focus on the time that the motes need to keep the radio active (i.e., transmitting or listening), since it is a widely-used hardware-agnostic metric for energy consumption [61].

This energy analysis concerns solely the LE packets, since the Rep mechanism is bypassed for the rest of the traffic by using a flag (explained in Sec. 3.4.2).

The analysis is based on the low-power preamble-sampling MAC protocol ContikiMAC v3.0, because it is the default MAC protocol of ContikiOS [50], the *de facto* standard operating system in the WSN domain. Moreover, we ignore the energy consumed by the microcontroller, which is a common assumption because the consumption of the transceiver is the dominant driver of the overall energy performance (e.g., [62]).

A more precise quantification of the network energy savings depends on the frequency of the LE operation and the use of the collected LQ information, both determined by the routing protocol. While we will leave a more detailed, scenario specific characterization and quantification of energy savings out of the scope of this chapter, it remains a valuable next step in our study.

We compare the energy that a mote requires for a LE operation using CG and Rep, by modeling each of the approaches. The model parameters are summarized in Table 3.1 and detailed in the following sections. We focus on the case where beacons are sent using broadcasts, because this is the main type of beacons used by state-of-the-art protocols. In this analysis, we ignore the processing delays in the microcontroller, which are small compared to the transition times between states of the radio.

Using Rep to extract RSSI and LQI means over many samples enables the reduction of the number of broadcast packets, which should diminish the overall energy consumption of the link estimation operation. Equation 3.8 details the percentage of energy savings provided by Rep (S), where E_{CG} and E_{Rep} represent the energy required for obtaining w_m samples using

Table 3.1 – Model Parameters

Parameter	Value	Meaning
n	1-9	Mean nb. of neighbors
T_d	1.54ms	Transmission duration of a repetition
T_D	2.2ms	Radio-on time for transmitting a repetition
T_{gap}	0.5ms	Time with radio-off between consecutive repetitions
T_{oh}^{tx}	0.872ms	Overhead time for sending a packet
T_{oh}^{rx}	1.27ms	Mean overhead time for receiving a packet
w_m	1 - 120	Nb. of samples
T_w	125ms	ContikiMAC Sampling period
w_{tw}	50	Nb. of repetitions to fill a T_w

CG and Rep, respectively.

$$S = 100 * (E_{CG} - E_{Rep}) / E_{CG} \quad (3.1)$$

First, we address the CG scheme (Fig. 3.3) in Eq. 3.2, and describe the total energy (E_{CG}) as the addition of the energy consumed in transmissions by a single mote (E_{CG}^{tx}) and the reception of the n mote(s) within range (E_{CG}^{rx}). The parameter n measures the density of the network.

$$E_{CG} = E_{CG}^{tx} + E_{CG}^{rx} \quad (3.2)$$

E_{CG}^{tx} (detailed in Eq. 3.3) represents the energy spent for transmitting w_m packets. Before the transmission, ContikiMAC performs two channel samplings, each with a duration of T_{cs} , thus creating a per-packet transmission overhead of $T_{oh}^{tx} = 2 * T_{cs}$. Assuming the channel is found available, the mote will proceed with the transmission of the w_{tw} repetitions that compose a broadcast, each of them requiring a radio-on time equal to T_D (including the time to turn the radio on and off). P^{tx} is the power consumption of the radio transceiver during transmission state.

The values of the previous parameters have been measured using Cooja [63] (the state-of-the-art simulator for WSNs, with cycle-accurate emulators for the microcontroller and radio) for ContikiMAC (version Contiki 3.0 vanilla, transceiver cc2420) and they are shown in Table 3.1.

The energy used for radio transmission decreases significantly, as a result of the reduction in the number of transmitted LE beacons. Even though Rep requires a small energy investment for extending the transmission of packet repetitions by w_m (see Sec. 3.4.2), this can be considered negligible in comparison to the reduction in the number of transmitted beacons, since each beacon is composed by tens to hundreds of packet repetitions (w_{tw} , depending on the settings of the MAC protocol).

$$E_{CG}^{tx} = P^{tx} * w_m * (T_{oh}^{tx} + (w_{tw} * T_D)) \quad (3.3)$$

E_{CG}^{rx} (detailed in Eq. 3.4) shows the energy spent by n motes within range, to receive one

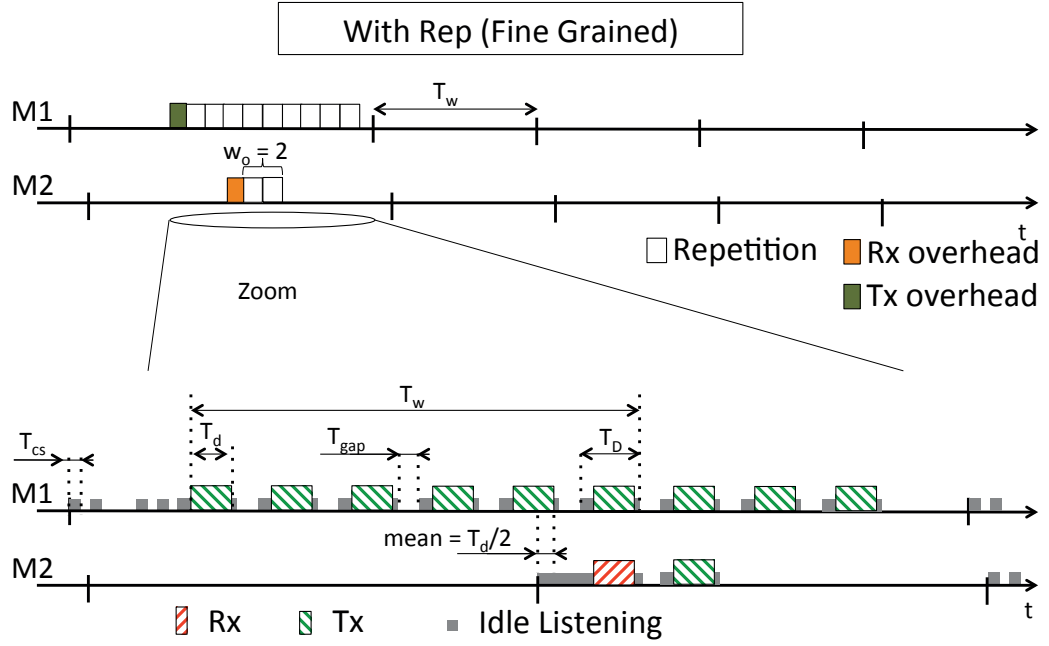


Figure 3.4 – Timing diagram of the link estimation operation, between nodes M1 and M2, using multiple repetitions of a single packet with Rep (fine grained sampling).

only once, due to the transmission of a single packet.

$$E_{Rep}^{tx} = P^{tx} * (T_{oh}^{tx} + ((w_{tw} + w_m) * T_D)) \quad (3.7)$$

In the case of the energy spent in idle listening and reception (E_{Rep}^{rx} , eq. 3.8), we consider the reception of w_m repetitions coming from a single packet, each of them requiring a radio-on time T_D (explained in the CG case). In the case of Rep, the reception overhead T_{oh}^{rx} is only taken into account once, due to the reception of a single packet.

$$E_{Rep}^{rx} = P^{rx} * (T_{oh}^{rx} + (T_D * w_m)) * n \quad (3.8)$$

The energy used for radio listening in Rep decreases when compared to the CG case, due to the reduction of wake-up probes (each requiring a time T_{oh}^{rx}), which are required for signaling the reception of a packet, and have a duration similar to the one of a packet in ContikiMAC. Therefore, Rep decreases the listening energy by reducing the number of wake-up probes to a single one (from a single packet), instead of w_m (from w_m packets, compare eqs. 3.4 and 3.8). Rep does not change the number of repetitions received, remaining w_m in both Rep and CG schemes.

Considering that low power transceivers typically have similar power consumptions during the transmission, reception and listening states, we use the expression $P^{tx} = P^{rx}$ and cancel these factors out of eq. 3.8.

Figure 3.5 shows the energy savings (S in 3.1) obtained by using Rep instead of CG. Note that

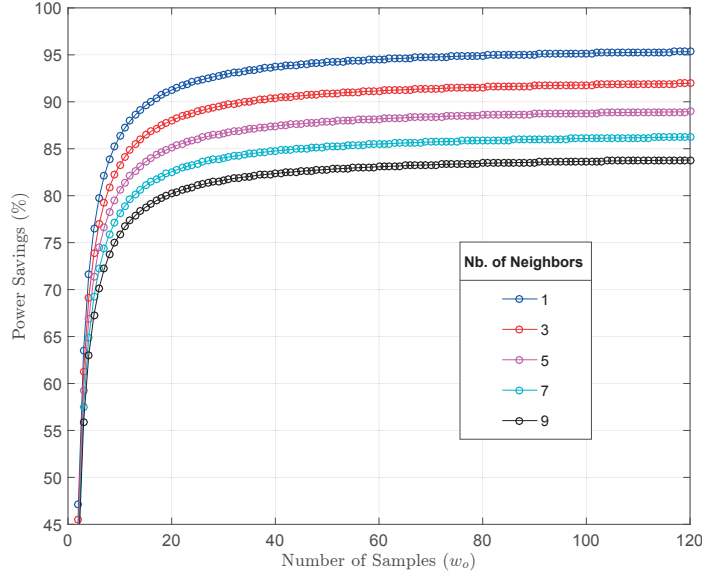


Figure 3.5 – Power savings of using Rep instead of using CG to obtain a number of samples, as a function of the number of samples and the network density (from the analytic model in Eq. 3.1). Note that the bottom of the y-axis is located at 45% to visualize more clearly the plateau region of the curves.

the bottom of the y-axis is located at 45%, in order to visualize more clearly the plateau region of the plots. We consider a number of neighbors in the range of 1 to 9, which corresponds to the spectrum between a sparse network and dense one. The limit of 9 is chosen close to the capacity of the neighbor table in CTP (10 neighbors) [64].

The savings in the figure display an elbow around $w_m = 18$ samples and they reach a plateau beyond this point. The plateau indicates that using Rep instead of CG brings an energy reduction of the link estimation operation of one order of magnitude (83 - 95%).

Figure 3.5 also shows that the savings decrease as the density of the network increases, thus reaching the value 83% for a $n = 9$ neighbors. We can obtain a deeper insight of this property by analyzing the limits of the Savings equation (eq. 3.1) with extreme density ($\lim_{n \rightarrow \infty} S$) and in the plateau of the curve ($\lim_{w_o \rightarrow \infty} S$), simultaneously. The conditions reduce S to $100 * T_{oh}^{rx} / (T_{oh}^{rx} + T_D)$. This represents a theoretical lower bound for the savings, which evaluated according to the values in Table 3.1 is 37%, a value that remains significant.

The energy consumption results were not experimentally validated, as a fair comparison would have required the simultaneous execution of the evaluations with Rep and CG, in order to have both schemes testing the same link quality conditions. We leave this further validation for future work.

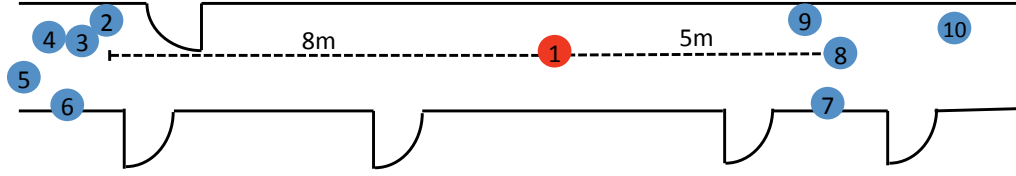


Figure 3.6 – Location of the experiment in an office hallway. N_1 periodically broadcasts beacons that are used by N_i ($i=2..10$) to estimate the link quality.

3.5.2 Performance Analysis

Overview

By design, Rep improves the speed and reduces the traffic of the link estimation primitive. Thus, we have performed four experiments aimed at comparing the accuracy of LQEs when using input information provided by Rep with the one provided by the state-of-the-art sampling scheme (detailed later in this section). The objective is to understand Rep's capacities and limitations under environmental disturbances commonly found in WSN deployments.

Setup

We have created a testbed for studying the link estimation accuracy inspired by RadiaLE [42]. RadiaLE is a framework for designing and assessing link quality estimators for WSNs, consisting of a star network, where nodes N_i ($i = 2..n$) are placed at different distances around a central node N_1 , as shown in Fig. 3.6.

Distance and direction are fundamental factors affecting link quality, therefore the links between N_1 and the surrounding nodes will have different qualities. We chose a particular direction for each node and determined the appropriate distance empirically, prior to the experiments, aiming to be as close as possible to the transitional region (quantified by means of the broadcast packets PRR).

The transitional region describes the operation of links with intermediate quality and it is generally characterized by a high-variance in reception rates [65]. We aim at setting the links in this range because, in practice, some of them will deviate from this operation point, thus also providing information about the regions below (low quality) and above (high quality) the transitional region. We study the performance of Rep under the most challenging conditions, by focusing on transitional links, as opposed to low- and high-quality links, which are more stable and require significantly less samples for an accurate characterization [39] [26] [43].

Our testbed consists of $n = 10$ TelosB nodes (Crossbow) and was placed on the ceiling of a hallway in an office building (Fig. 3.6). Most of the hallway walls are metallic.

We use the following traffic pattern: N_1 is the only node that transmits, the others are listeners during the entire experiment. N_1 sends a broadcast every 2s and the other nodes extract the link quality from the MAC repetitions of each packet (*RSSI* and *LQI*).

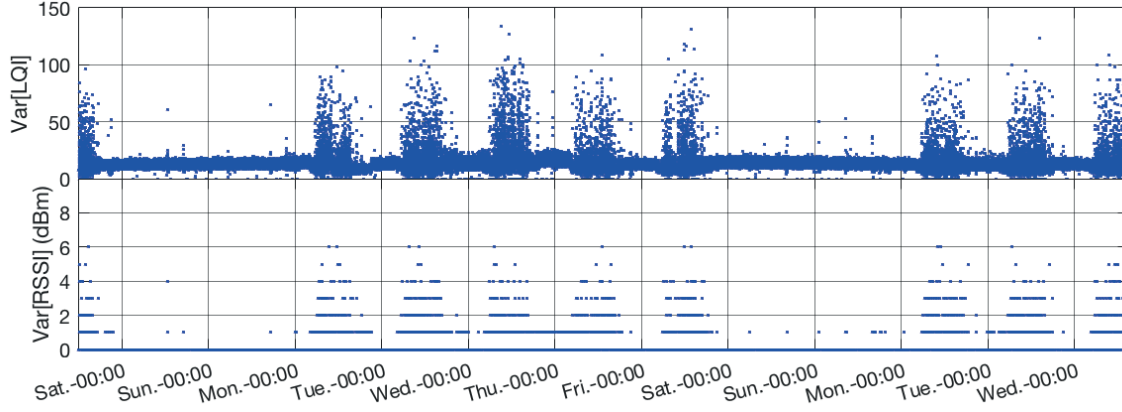


Figure 3.7 – Variance of LQI (top) and RSSI (bottom) during 13 days ($w_m=120$ repetitions). The system is sensitive to the traffic of people through the hallway as can be seen during the working hours (Channel 11).

The focus of this study on *RSSI* and *LQI*, over *PRR*, has a threefold motivation: i) the domain shows a significant and growing interest in exploiting these physical information sources for link estimation [25] [66], ii) there is significant experimental evidence that both metrics preserve a correlation with the link quality at small sampling periods ($\sim 10\text{ms}$, similar to the reception period of MAC repetitions) [39], and iii) both metrics are widely available in low-power transceivers. We do not elaborate on *PRR* measurements obtained through REP, because *PRR* has the tendency to acquire extreme values (e.g., 0-100%) as the period of the sampling becomes shorter. We have observed this in our experiments and it has been reported by other authors in [39] [67].

All the nodes are running Contiki OS v3.0 [50] and contikiMAC-R (sampling period 2Hz).

Independent Variables

WSNs are subject to multiple environmental disruptions with different time and intensity patterns. Therefore, we have designed experiments to study the effects of common disruptions by analysing the following variables:

- **Wireless Interference:** all the experiments were repeated in two channels: 11 (2.405 GHz) and 26 (2.480 GHz) of the TelosB radio. The former is susceptible to interference from Wi-Fi and Bluetooth networks, while the latter is out of the interference spectrum. This technique for studying the effects of interference has been previously used in RadiaLE [42].
- **Variable Fading:** the movement of people causes changes in the multipath fading pattern, which can increase or reduce the RSSI and LQI. This can be seen in the increase of the variance of RSSI and LQI during working hours, in Figure 3.7 [68].
- **Number of Repetitions:** the maximum number of repetitions that a node can receive from a packet (length of the observation window) determines the accuracy of link

Table 3.2 – Experimental Conditions

Experiment \ Conditions	Wireless Interference	Variable Fading
1: Stable Channel	✗	✗
2: Effect of Wireless Interference	✓	✗
3: Effect of Variable Fading	✗	✓
4: Wireless Interference + Variable Fading	✓	✓

estimation. All experiments were evaluated using broadcasts ($F_{LE} = \text{true}$) and repeated with $w_m=20$ and $w_m=80$ repetitions. Moreover, the link quality from the first repetition was registered separately in order to evaluate the link estimation accuracy when $F_{LE} = \text{false}$. We select the size of the windows motivated by the promising LE accuracy reported in [22] (for $w_m=20$) and [39] for (for $w_m=80$).

Table 3.2 summarizes the conditions of several experiments, designed to understand the effect of the previous variables on the accuracy of Rep. We have analysed each combination of conditions during at least 2h, resulting in 3600 broadcasts (sampling opportunities).

Dependent Variable

We have performed four experiments aimed at comparing the accuracy of LQEs based on a state-of-the-art sampling scheme, and based on Rep. We have measured the link estimation accuracy by evaluating the PRR in the near future (PRR_f) and studying its correlation to the link estimation metrics, for each received packet. This methodology has been borrowed from the LE study in [22].

PRR_f was calculated offline over a window of 100 packets that are expected after each reception, as this scheme emulates the application traffic. This methodology is inspired by previous work in the link estimation domain [22].

Control System (Coarse-Grain Sampling)

The experiments compare the accuracy of link estimation based on the following schemes:

- Coarse-grain (Standard Method): This is the traditional scheme for link estimation. The information is obtained by each receiver (N_i , $i=2\dots10$) by considering the link quality from the first repetition received from each packet, and then averaging the values obtained from the packets received during a time interval when m packets are expected (m packets result in \overline{RSSI} and \overline{LQI}).
- Fine-grain (Rep): A receiver node considers the link quality from the repetitions received

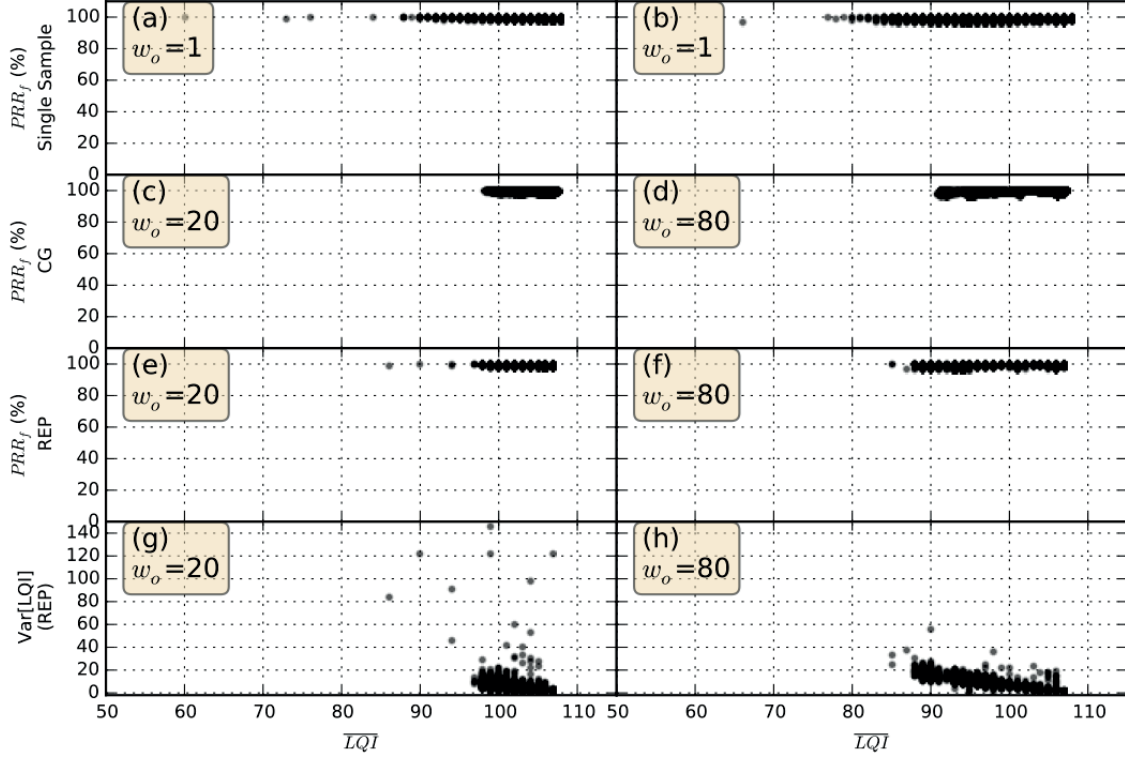


Figure 3.8 – Future PRR vs. the current value of \overline{LQI} with negligible wireless interference and variable fading (Experiment 1). The plots obtained using CG sampling (c and d) and Rep (e and f) show both schemes result in equivalent link estimations (same trend and comparable dispersion). \overline{LQI} is calculated with 20 (left column) and 80 samples (right column). Plots a and b use a single sample. The variance of the \overline{LQI} obtained with Rep is shown in plots g and h. The black dots are semitransparent, therefore the darkness of the tone is proportional to the number of overlapped data points.

during a time interval when m repetitions are expected, as part of a single packet (a window w_m provides \overline{RSSI} and \overline{LQI}).

Given that ContikiMAC uses a standard sampling period (512ms), the transmission of 80 packets for a CG link estimation requires at least 40.96s ($= 80 \times 0.512s$, assuming there is no contention to access the channel). In our experiment, the packets are transmitted with a 2s period, which is closer to values used in practical applications, therefore a single CG link estimation requires 160s ($= 80 \times 2s$). With Rep, a single link estimation takes less than 1.024s ($= 2 \times 0.512s$), as it requires less than two sampling periods: the sampling period required for sending a broadcast, plus the w_m window (e.g., the train of transmitted repetitions from M1 exceeds T_w in Fig. 3.4).

Upon reception of a packet, ContikiMAC-R provides the link quality data from the first repetition (CG) and the aggregated data from m repetitions (FG). We use the latter for studying the Rep link estimation, and the former in two different ways: for studying the link estimation with a single sample (for example, Fig. 3.9 row 1) and for studying the CG estimation by averaging it over multiple packets (for example, Fig. 3.9 row 2).

Results

Experiment 1: Stable Channel

This experiment represents the baseline scenario, since it is performed without additional interference or variable fading (details about the conditions can be found in Table 3.2). The plots only display points above the transitional region, due to the good conditions of the channel, that narrow down the transitional region [65].

Figure 3.8 displays several scatter plots showing the dependency of PRR_f on the \overline{LQI} , for $m = 20$ repetitions (left column) and for $m = 80$ repetitions (right column). Each row in the figure displays the results of a different sampling method: single-sample (first row, plots a and b), CG (second row, plots c and d) and Rep (third row, plots e and f). The fourth row displays the LQI variance vs. \overline{LQI} for Rep (plots g and h).

Figure 3.8 shows that the links have a $PRR_f > 96\%$ for $\overline{LQI} > 98$, which translates into a good quality. Given that the channel is stable, it can be expected that they will keep a high quality in the future, hence their future reliability can be predicted with a high accuracy.

A comparison between the plots that use a single sample in Fig. 3.8 row 1 and their counterparts that take into account several samples (rows 2 and 3), shows that the multiple samples reduce the dispersion in the x-axis, thus providing a more accurate \overline{LQI} .

The CG plots (row 2) exhibit negligible differences when compared to Rep plots (row 3), for each w_m value, thus showing that the link estimation accuracy that can be achieved with both sampling methods is equivalent when the channel is stable. This result is expected since, in the absence of channel perturbations, the accuracy should be independent of the sampling time-scale.

The results for RSSI of this experiment are omitted because they bring the same observations already presented for LQI.

Experiment 2: Effect of Wireless Interference

This experiment studies the effect of interference from wireless networks over the estimation accuracy of link quality metrics obtained from Rep. Fig. 3.9 is structured analogously to Fig. 3.8.

Plots in Fig. 3.9 obtained from averaging multiple samples (rows 2 and 3) display 4 separate data point islands. It was not possible to obtain data to connect them as there were no links available in this operation region. We overlap a regression curve (displayed in red), which is detailed in Fig. 3.11 and Sec. 3.5.2. The discretization in the x-axis for single-sample and Rep plots is noticeable, the reason is that \overline{LQI} is an integer because it is calculated in the node's microcontroller, while the \overline{LQI} from CG plots (row 2) is a floating point value because the average is calculated offline.

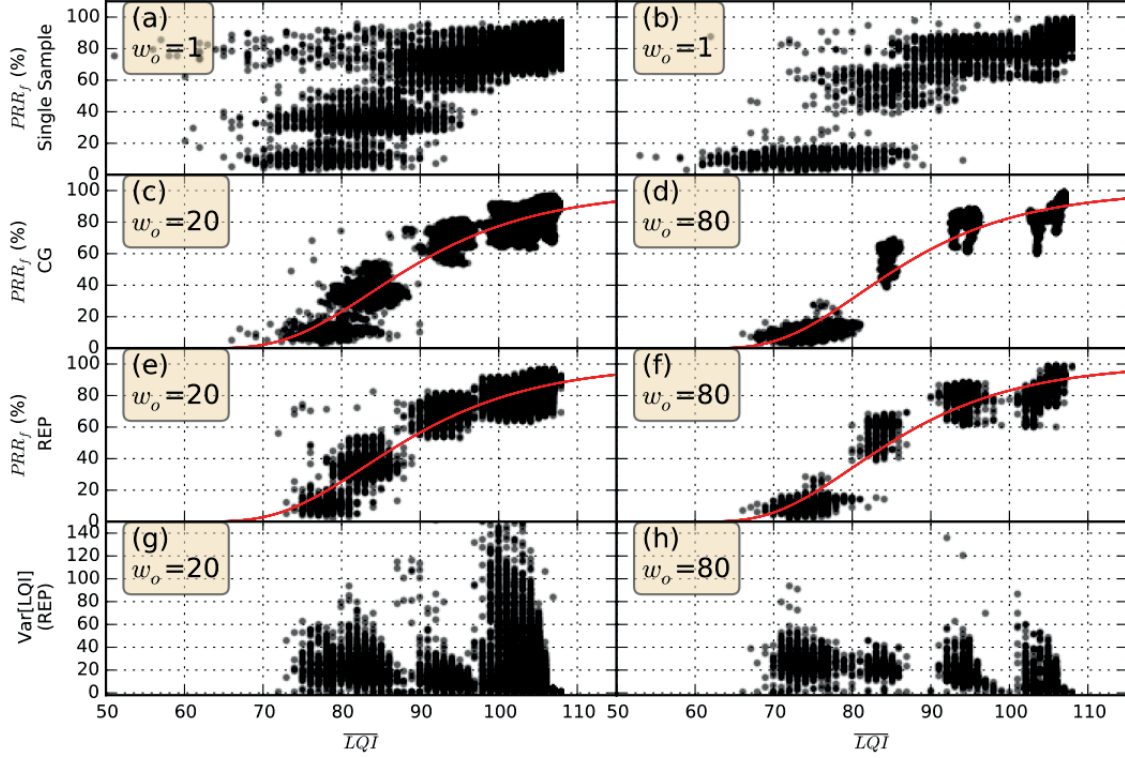


Figure 3.9 – Future PRR vs. the current value of \overline{LQI} to study the effect of wireless interference (Experiment 2). The plots from CG (c and d) and Rep (e and f) show both schemes result in equivalent link estimations (same trend and comparable dispersion). For further interpretation details of the plot refer to the caption of Figure 3.8. The regression curves (in red) are compared in Fig. 3.11.

Each subfigure can be compared to its analog from Fig. 3.8, which shows that wireless interference increases the variance on both axes. We note that CG (row 2) and Rep (row 3) plots in Fig. 3.9, tend to have a sigmoid shape as described in [39] [69], and display a clear transition region at $80 < \overline{LQI} < 90$. These features come from the variance introduced by wireless interference (compare to Fig. 3.8). This effect has been reported in previous studies [23].

The comparison between CG and Rep plots for each w_m value does not show a significant difference in terms of trend of data and its dispersion. A comparison of the regression curves shows that the corresponding CG and Rep plots highlight similar underlying sigmoids. This behavior supports our assertion that the accuracy of Rep is equivalent to the one of the standard method. However, while each point of the former required transmitting 20 and 80 packets, Rep only required a single packet.

The sigmoid trend in the results validates the RadiaLE-like approach to study the quality of the links, since each link only provides information from a particular operation region. Thus, using several links gives insight about multiple portions of the characteristic curve.

The comparison between the plots in Fig. 3.9 obtained from a single sample and the ones obtained from multiple samples, shows that averaging multiple samples significantly reduces

the dispersion in the x-axis, thus improving the accuracy of the \overline{LQI} calculation.

Figure 3.10 is analogous to Fig. 3.9, but displays the dependency of the future PRR to the RSSI. The regression curves are compared in Fig. 3.11 and explained further in Sec. 3.5.2. The transitional region in Fig. 3.10 (first and second row) is well defined between -93 and -90 dBm and it is significantly steeper than the transitional region for LQI in Fig. 3.9. These features agree with previous work reporting that the average \overline{LQI} has a stronger correlation with PRR when compared to \overline{RSSI} , in the transitional region [39] [56] [57].

Additionally, Fig. 3.10 displays a lower dispersion in the y-axis for high RSSI values than its LQI counterparts in Fig. 3.9. This result is aligned with previous work that reports a smaller variance of RSSI compared to LQI for a given number of samples [39]. This explains the negligible reduction of $\text{Var}[\text{RSSI}]$ when it is calculated with 80 samples instead of 20 (Fig. 3.10 g and h). The wireless interference increases the variance of PRR_f , as it can be noted by comparing Figures 3.8 and 3.9 (second and third row). This observation is valid both for LQI and RSSI.

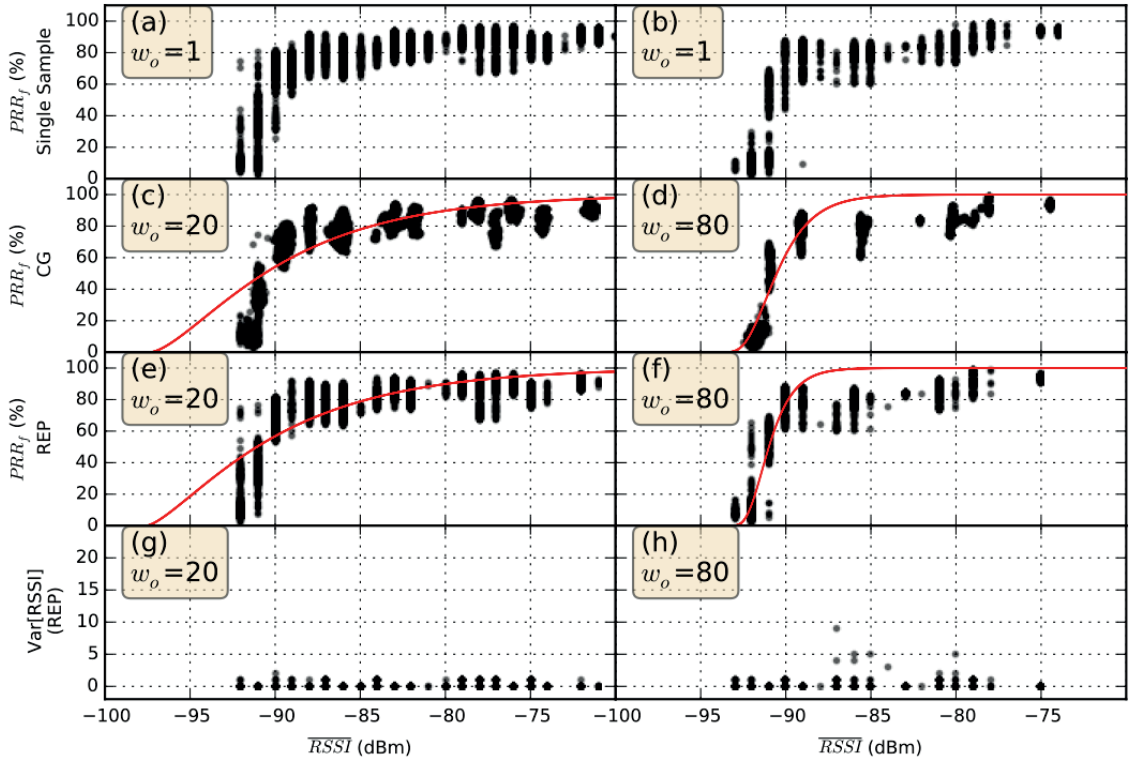


Figure 3.10 – Future PRR vs. the current value of \overline{RSSI} to study the effect of wireless interference (Experiment 2). The plots from CG (c and d) and Rep (e and f) show both schemes result in equivalent link estimations (same trend and comparable dispersion). For further interpretation details of the plot refer to the caption of Fig. 3.8. The regression curves (in red) are compared in Fig. 3.12.

The data displayed in Figs. 3.9 and 3.10 was analyzed through a non-linear regression, in order to better study the link estimation capabilities of the Rep and CG methods. The fitting curves chosen are the sigmoids in Eq. (3.9) (LQI) and Eq. (3.10) (RSSI). Eq. (3.10) (RSSI) has

been proposed as a theoretical model for the PRR of a low-power wireless link as a function of RSSI [65]. We extend the use of this model for the LQI, based on the reports of multiple authors that show a sigmoid-like behavior in LQI similar to RSSI's [39] [70]. The position of the transition zone of the sigmoid (x_t in the equations) has been chosen to maximize the r-squared coefficient (i.e., minimize the regression error).

The non-linear regression does not intend to prove the goodness of fit of the measured data points to a theoretical model. Instead, the objective is to show that the two link estimation mechanisms, Rep and CG, result in similar best fit curves, thus stressing the equivalent accuracy of both methods. The technique of fitting PRR vs. LQI or RSSI curves is widely used for predicting the behavior of a link from LQI or RSSI measurements through sigmoid models (e.g. [39] and [65]), or piece-wise linear models (e.g., in [70] and [71]).

$$PRR_f = 100 * (1 - (0.5 * e^{-(\overline{LQI} - x_t)/b}))^a, x_t = 65 \quad (3.9)$$

$$PRR_f = 100 * (1 - (0.5 * e^{-(\overline{RSSI} - x_t)/b}))^a, x_t = -93 \quad (3.10)$$

Table 3.3 summarizes the coefficients obtained from the regressions, and Figs. 3.11 and 3.12 display the regression curves. The Rep and CG curves are very similar for a given w_m , thus supporting the equivalent accuracy of both methods. Moreover, the distinct w_m values bring more significant differences between the curves compared to the use of a given data extraction method (Rep vs. CG). These observations hold for both LQI and RSSI.

RSSI curves display more significant differences between w_m values, which are clearly visible in Fig. 3.12. This can be explained considering the lower influence of external interference during the experiment with w_{20} compared to w_{80} , thus resulting in less datapoints in the lower part of the curve in Experiment 1. This difference is noticeable in Figs. 3.10 (c) and (d) (w_{80}), which display the lower tail of the curve extending to -93 dBm, not present in (e) and (f) (w_{20}).

We do not extend the regression study to Experiments 1, 3 and 4, because their experimental conditions distance the results from the sigmoid model. In the case of experiment 1, the lack of interference only provides data for the upper region of the curve associated with good links. Experiments 3 and 4 are subject to significant disturbances, due to variable fading, that are not part of the model in Eqs. 3.9 and 3.10.

We do not include a comparison between the results of the experiments $w_o = 20$ vs. 80 samples, because the measurements were not done at the same time, under the same perturbations over the network. We leave this analysis to be done in the future in a setting designed for controllable perturbations. Nevertheless, we believe that the results presented are a good starting point for validating the controllable perturbations in a future experiment. These observations are valid for all the results presented in this section (3.5.2).

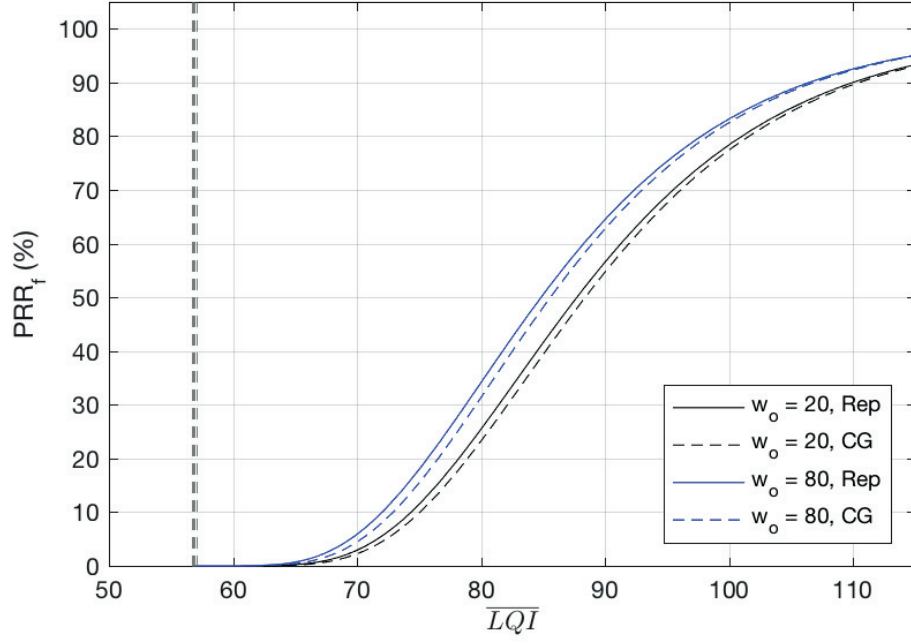


Figure 3.11 – Regression from Future PRR vs. the current value of \overline{LQI} datasets (Fig. 3.9) using a sigmoid model (Eq. 3.9). The similarity of the resulting curves for Rep and CG stress the equivalent accuracy of both data extraction schemes.

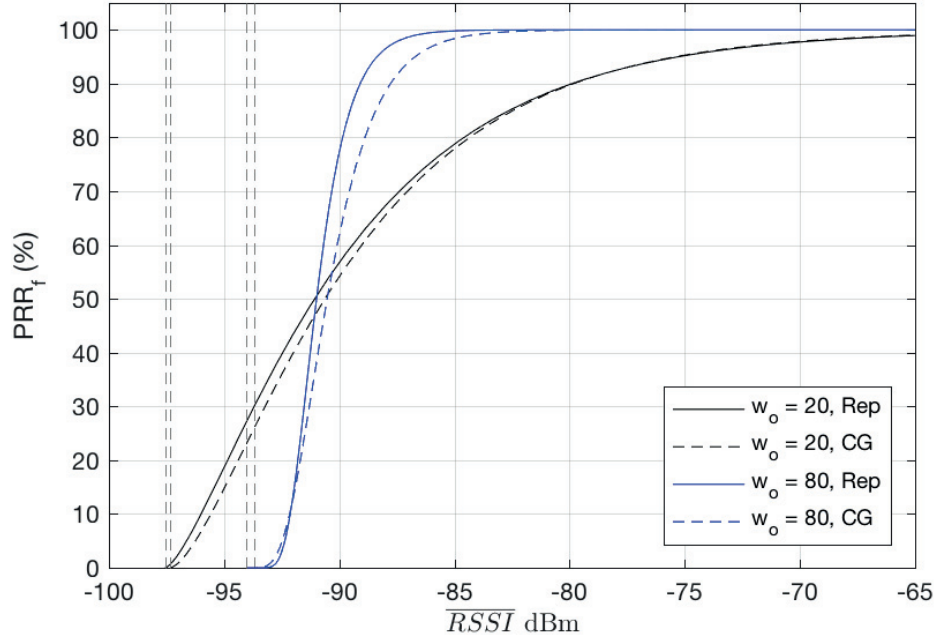


Figure 3.12 – Regression from Future PRR vs. the current value of \overline{RSSI} datasets (Fig. 3.10) using a sigmoid model (Eq. 3.10). The similarity of the resulting curves for Rep and CG stress the equivalent accuracy of both data extraction schemes.

Chapter 3. Leveraging MAC Preambles for an Efficient Link Estimation

Table 3.3 – Results of the regression using the sigmoid in Eqs. (3.9) (LQI) and (3.10) (RSSI).

Experiment	coef. a	coef. b	R-squared
LQI, $m = 20$, Rep	8.832	12.02	0.8682
LQI, $m = 20$, CG	9.536	11.91	0.8715
LQI, $m = 80$, Rep	7.1704	11.73	0.9175
LQI, $m = 80$, CG	7.9544	11.5	0.9167
RSSI, $m = 20$, Rep	1.4832	6.592	0.6491
RSSI, $m = 20$, CG	1.6488	6.3	0.6638
RSSI, $m = 80$, Rep	9.568	1.021	0.7881
RSSI, $m = 80$, CG	6.8832	1.494	0.7884

Experiment 3: Effect of Variable Fading

Figure 3.13 displays the effect of changes under multipath fading. This scenario is especially challenging because the link degradation is caused by human actions, such as the opening of a door or the crossing of a hallway, and therefore cannot be predicted. This explains the high dispersion in both x and y-axes in all the plots.

The CG plot for w_{20} (Fig. 3.13.c) shows a transitional region with a large LQI dispersion (in the x-axis) between $\overline{LQI} = 75$ and 97. This result agrees with previous work that shows that multi-path fading affects the size of the transitional region [65].

Notably, values of 100% future PRR are achievable, which is not observed in the plots with wireless interference (e.g., Fig. 3.9). This behavior comes from the fact that the links normally have a good quality, unless the multipath fading pattern changes adversely.

The comparison between CG and Rep plots in Fig. 3.13 shows that, in spite of the unpredictability of the channel distortions, both sampling schemes produce scatter plots that follow similar trends.

Figure 3.13 displays several green arrows that highlight a "leg" (an anomaly visualized as an approximately vertical section) in the plots. This feature is caused by the loss of a significant number of consecutive packets, which progressively affects the future PRR of multiple points as the sliding window advances through the anomaly. Therefore, we interpret the leg shape as a result of accounting for the same anomaly from the previous points several times.

Moreover, the vertical distribution of a leg proves a lack of correlation between the variables in the scatter-plot. Thus the legs are the consequence of events that cannot be predicted (i.e., anomalies, e.g., human-triggered events), that cause changes in the fading pattern and result in a prolonged link disconnection. This also explains why the legs are not present in experiments performed outside the office's working hours (experiments 1 and 2).

The anomalies affect both sampling schemes in the same way (rows 2 and 3). The same effect is present in the RSSI plots (Fig. 3.14). We explain the presence of more legs for w_{20} (left column) compared to w_{80} to an exceptional increase in the traffic of people due to a company activity during the former experiment.

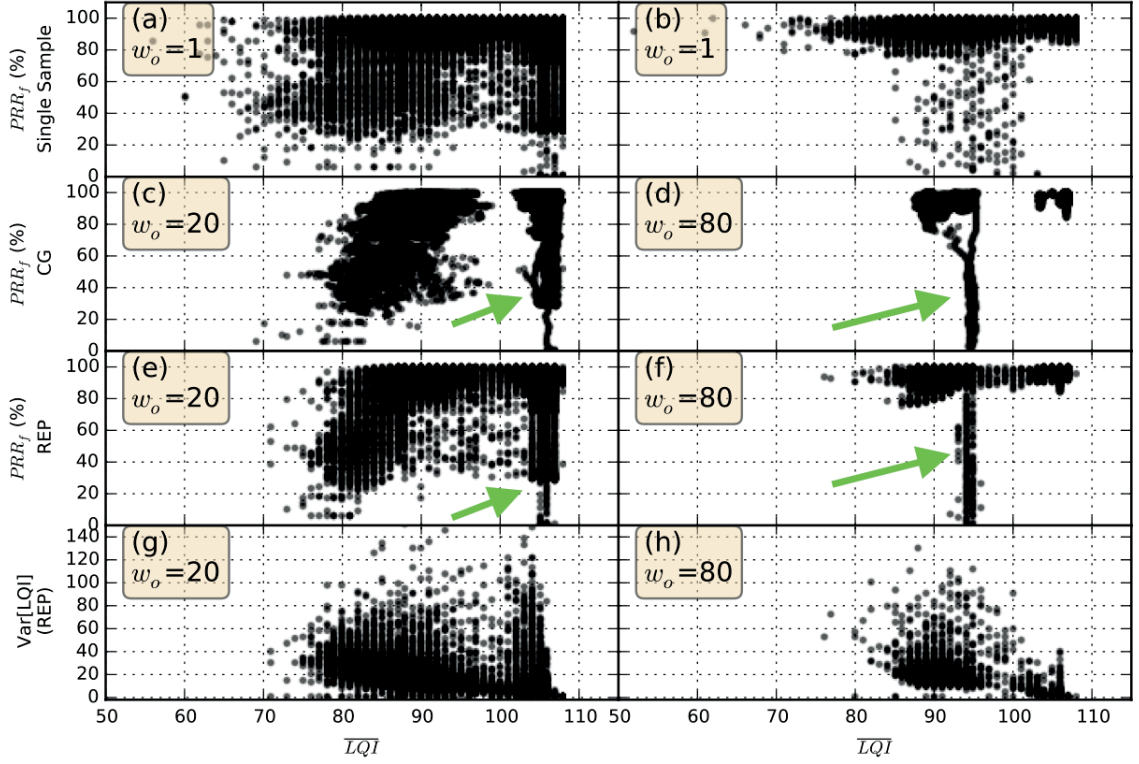


Figure 3.13 – Future PRR vs. the current value of \overline{LQI} to study the effect of variable fading (Experiment 3). The plots obtained using CG sampling (c and d) and Rep (e and f) show both schemes result in equivalent link estimations (same trend and comparable dispersion. For details about the interpretation of each plot refer to the caption of Fig. 3.8.

Experiment 4: Wireless Interference + Variable Fading

Figure 3.15 displays the results with both wireless and human interference, hence it is the worst case scenario under consideration. The PRR is always under 90% and, for high LQI, it ranges between 50% and 90%. The absence of PRR values above 90% is attributed to the presence of wireless interference, since this feature is only observed in the experiments that share this condition (for example, experiment 2).

The plots do not display legs as prominent as the ones seen in Fig. 3.13. We attribute this to two factors: i) the human events during this experiment caused shorter interruptions of the links, thus making the legs smaller on the y-axis and ii) the dispersion due to the wireless interference hides the short legs.

In this difficult scenario, the plots from Rep and CG still follow similar trends. This observation holds for the entire range of LQI values analysed. Therefore, in the worst scenario under scrutiny, the results show a remarkably similar link estimation accuracy between the two methods (Rep is as good as CG). There is also no clear advantage in using 80 repetitions (Figures 3.15.d and f) over using 20 repetitions (Figures 3.15.c and e). The results for RSSI of this experiment are omitted because they bring the same observations already presented for LQI (Fig. 3.15).

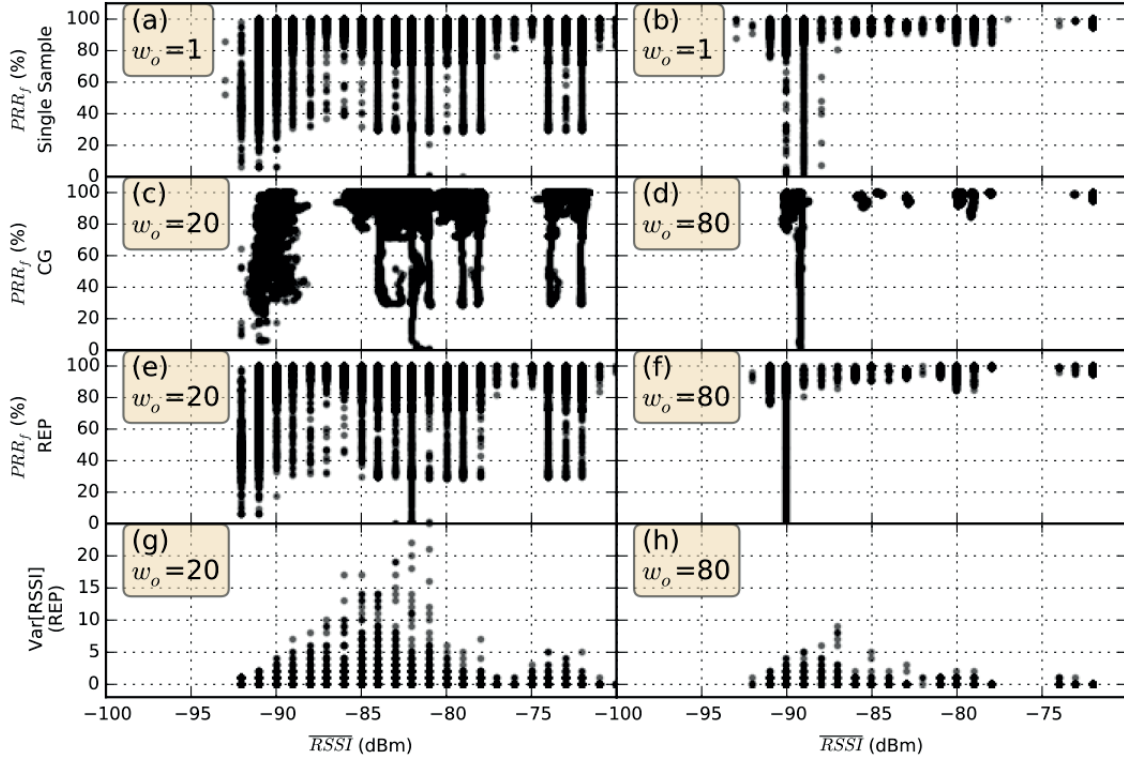


Figure 3.14 – Future PRR vs. the current value of \overline{RSSI} to study the effect of variable fading (Experiment 3). The plots from CG (c and d) and Rep (e and f) show both schemes result in the same trend and comparable dispersion. For details about the interpretation of each plot refer to the caption of Figure 3.8.

3.5.3 Consequences of Reducing the Sampling Granularity

The equivalence between the LE obtained from CG or Rep will depend on the temporal dynamics of the perturbations over the links (in the case of non-mobile networks). We have identified three cases where the two methods should yield similar results. First, there are no perturbations, therefore the LE will be independent of the sampling period. Second, the changes of the perturbations occur progressively and over a time-scale that is significantly bigger than the one of CG sampling, so that Rep (being faster) will still capture the same behavior of a link. Third, the changes follow a pattern, with a period small enough to be fully captured with the short duration of Rep.

We will elaborate on the third case to interpret the results obtained in section 3.5. In an office environment, the dominant source of interference creating the need for multiple samples for a reliable RSSI and LQI measurement is represented by the use of WiFi [72]. Potential other interferences such as Bluetooth and Zigbee are both far less ubiquitous, and microwave ovens tend to have a sporadic use, mostly around lunchtime.

Studies characterizing WiFi variations over time, such as [49] have found that WiFi typically displays a radiation pattern composed of spikes with a period of 24.67ms, each with a width

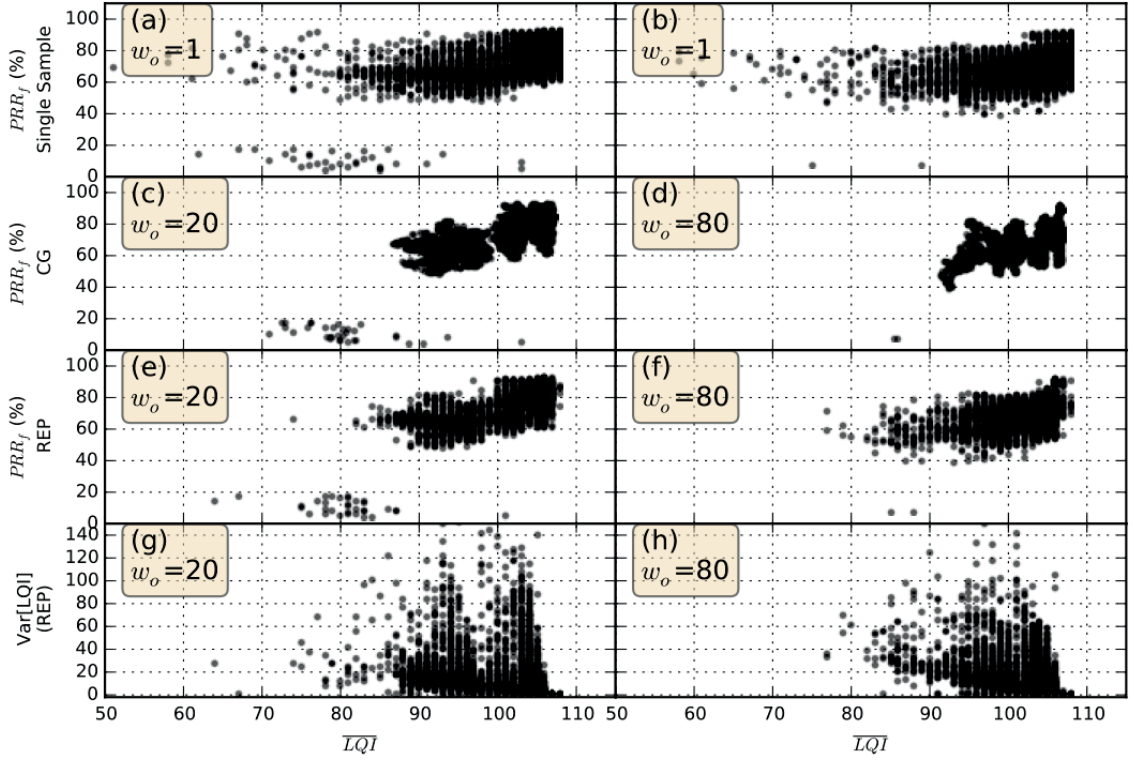


Figure 3.15 – Future PRR vs. the current value of \overline{LQI} to study the effect of simultaneous wireless interference and variable fading (Experiment 4). The plots from CG (c and d) and Rep (e and f) show both schemes result in the same trend and comparable dispersion. For details about the interpretation of each plot refer to the caption of Figure 3.8.

of 8ms. Moreover, when constructing empirical models, studies based on JamLab [72], the state-of-the-art mechanism for generating typical interference patterns have also considered that, when downloading a large file, with 97% probability the clear channel period is smaller than 1ms. The overlap of multiple WiFi users, typical for an office environment, creates a continuous pattern during the working hours, while also potentially increasing the frequency of the spikes.

Therefore, the dominant interference sources present a periodic pattern with temporal dynamics that should be fully captured by Rep. This behavior explains our ability to achieve equivalent noise mitigation effects with our fine-grain sampling (of 2.725ms per sample in the Contiki 3.0 vanilla code), as the ones of the traditional coarse-grain sampling in the scale of minutes.

3.6 Applications

This section details benefits that several domains can perceive from the reduction in energy, duration and traffic of the channel sampling enabled by Rep. The following points detail the benefits in the context of data collection:

- Increasing the lifetime of collection networks: WSNs used for collecting data over regions beyond the transmission range of a single node typically use broadcast beacons to find possible parents and to estimate the quality of the links. Thus, link estimation is a basic primitive to build a routing tree and convey the data to a gathering point [23], but it also represents an energy overhead. Moreover, it can be difficult or expensive to change the battery of the nodes, hence maximizing the lifetime translates into lower operational costs. A collection network can use Rep to reduce the number of beacons required for an accurate link estimation and consequently increase the battery lifetime.
- Improving the reliability of multi-hop WSNs: the nodes use a best-effort approach to define the parent in a routing topology based on link quality information that can be noisy and outdated [14]. The reason is that the number of link quality samples is limited by the energy overhead and delay required for probing the channel. This introduces a catch-22 dilemma: nodes must select parents with high quality links in order to ensure a reliable operation, while at the same time the process requires a long delay, limiting the network's ability to adapt quickly to perturbations, and hindering its capacity to reliably transfer application data. This pain-point is common to static [14], mobile [22] and opportunistic WSNs [43]. With Rep, the nodes can perform a quick and accurate link estimation, thus overcoming this trade-off and improving the reliability of the network.
- Enhancing the dataset for machine learning applications in WSNs: recent technologies aim at using machine learning techniques to improve the selection of the routing topology [25] [29] [28]. These efforts rely on coarse-grain link quality that is typically collected in WSNs. Therefore, we propose to complement it with the fine-grain information that Rep can provide. This approach can offer additional information, while reducing the time required to build the machine learning datasets.
- Improving the recovery and bootstrapping of WSNs: previous efforts have been aimed at improving the recovery of a WSN upon disconnection of several nodes and the set-up upon start-up (i.e., bootstrapping). This process determines the delay until the network is operational [14], and the reliability of the initial configurations [73]. Rep can improve the reliability of the initial topologies by providing accurate link quality when no historic data is available, while preserving the speed of bootstrapping schemes that rely on a single packet transmission.

The fast and energy-efficient link estimation provided by Rep can be used in applications beyond data collection applications, for example:

- Improving the accuracy of RSSI-based wireless localization: an approach for wireless localization is to leverage the dependency of the RSSI on the transmission distance. Thus, multiple nodes with known locations can register several RSSI samples from a node in order to determine its position via triangulation [74]. When multiple nodes must be localized, the shared access to the channel becomes a limitation to collecting enough RSSI samples for obtaining an accurate position. In this context, Rep can increase the number of RSSI samples while reducing the load in the channel.

- Improving the accuracy of Radio Tomographic Imaging: this technique uses a WSN installed in the periphery of a target area to estimate the locations of people and objects within it. Its working principle relies on detecting the obstruction of multiple radio links by analysing changes in their transmission (TX) statistics [75]. As shown in [68], the fine time granularity of Rep is an enabler of the LQI variance for detecting disruptions in the links due to the movement of people, thus providing a new source of information to complement the RSSI variance that is typically used in the domain.

3.7 Encountered Difficulties

Among the challenges faced during the study of fine-grain link estimation techniques, data collection and analysis proved to be the two most difficult ones. The data collection represented a challenge due to a steep transition between bad and good quality links, producing a narrow transitional zone. Therefore, during initial experiments we were only capable of capturing the behavior of the good and bad links, the second ones being especially difficult to characterize due to their low packet reception rate. We were only able to address these challenges by manually adjusting the distance and orientation of each mote to approach its link to the transition region. The state-of-the-art approach, introduced by RadiaLE, is to increase the diversity in link conditions by using a large number of motes (48 units) arranged with different orientations and in circles with diverse radius (from the transmitter mote in the center). Nevertheless, we were limited by the number of available motes (10 units), and it is not possible to use remote testbeds, such as Flocklab [51], for this type of experiment because the motes cannot be re-arranged.

The data analysis proved to be challenging, the main reason being that we were not able to find a statistical test commonly used by the WSN community to demonstrate that the data coming from two sampling mechanisms (Rep and CG) belong to similar distributions or display the same accuracy, within a given confidence interval. In other words, our hypothesis that the data obtained with CG (e.g., Fig. 3.9 (c)) and Rep (e.g., Fig. 3.9 (e)) come from similar probability distributions would have been further supported by using a statistical test, such as a multivariate t-student. Nevertheless, the best approach found to compare our datasets was to rely on linear regression.

If we had to reproduce the same experiments in this chapter, the data collection could be significantly improved by printing through the serial port the link quality obtained from every sample in the observation window (w_m), as opposed to printing aggregated metrics for the entire w_m . The advantage of the first method is that the size of the observation window can be adjusted offline, thus enabling the study of the effect of many window sizes, and significantly increasing the data analysis capabilities. For example, for every 80 samples collected with $w_m = 80$, it is possible to study the effects of any $0 \leq w_m \leq 80$. In order to carry a live data collection strategy, it is indispensable to connect all the motes to a computer via USB cables, otherwise the motes' memory would severely limit the length of the experiment.

Another foreseeable improvement for the presented experiments is to use reproducible interference conditions in order to enable better comparability between experiments by ensuring that the conditions are as close as possible. In this spirit, predefined interference patterns can be generated using Jamlab [72] in several motes or WiFi transmissions in Raspberry-Pi computers [61]. This technique requires the minimization of external interference sources, which can be achieved by using an interference free channel (e.g., ch26 in the transceiver cc2420), and running the experiments during the nights and weekends, when the external WiFi or Bluetooth interference is minimal.

3.8 Conclusions

This chapter has elaborated on Rep, a new sampling mechanism that leverages existing resources of LP-MAC protocols in order to reduce the duration and overhead of link estimation schemes (in terms of traffic) by one order of magnitude, with significant improvements in terms of speed and energy autonomy. The experimental evidence provided has proven that the link estimation accuracy, when using Rep, remains equivalent to the state-of-the-art methods' in a wide range of challenging scenarios. They have also shown, that, under the perturbation pattern of an indoors office deployment, a fast and energy efficient fine-grain sampling can result in measurements as accurate as those given by more time and energy expensive coarse-grain sampling.

4 Deploying Collection Protocols in Real World Applications – The Fly Project

4.1 Introduction

Real world deployments of collection protocols represent unique opportunities to test the strengths and limitations of their design features under challenging requirements. Nevertheless, the difficulty in deploying and maintaining experimental real-world WSNs [76] [77] [78] results in the overwhelming majority of collection protocols being analyzed either in testbed applications or in controlled indoor environments. Until now, we have only been able to identify a handful of relatively small-scale WSNs tested “in the wild”.

The WiseFly project thus represented a unique opportunity for validating and enriching the findings of the previous two chapters in the context of a real-world deployment. The project consisted in the development of a collection protocol (WiseFly) that incorporates the main findings of the previous two chapters (WiseNE, composite metrics and Rep), and its subsequent deployment for the monitoring of a public utility system in the frame of a Smart City initiative in Europe.

The application required the deployment of a multi-hop collection protocol in an urban setting, a challenge because of the dynamic conditions of operation, such as weather, vehicle and pedestrian traffic, hardware failures (due to natural factors or vandalism [79]), and multiple interference sources. The sensitivity of the data flows being monitored, as well as their direct use by operators for taking intervention decisions made the application particularly demanding in terms of reliability. Moreover, the low accessibility of the extensive network deployed gave further weight to its energy efficiency requirements. At the same time, several motes had to be able to record audiograms in a coordinated manner, which necessitated an on-demand network synchronization with an error of under 1 ms, as well as the capacity to cater for significant traffic surges.

As expected, real-life deployments are often subject to more numerous sources of variability, mostly due to the open environment, subject to multiple sources of noise, such as vehicle and pedestrian traffic. A common issue is the lack of Line-of-Sight (LOS) between motes being

Chapter 4. Deploying Collection Protocols in Real World Applications – The Fly Project

installed by the user company, often resulting in much weaker, reflection-based communications between the nodes. Moreover, some parameters taken for granted by most protocols, e.g., the LQI are not made available by all transceivers.

This challenging environment has allowed us to validate the strengths, limitations and trade-offs displayed by our previous findings, in particular WiseNE - a neighborhood exploration mechanism designed for providing the inputs required by composite metrics (presented in Chapter 1), and Rep - an energy efficient sampling mechanism extracting the link quality information from the packet repetitions at the MAC layer (presented in Chapter 3). Our experiments have shown that composite metrics such as PRSSI (Penalized Received Signal Strength Indication) continue to bring an improvement to simple ones even in environments subject to high variability. Moreover, by using the Rep beaconing mechanism, we have managed to reduce the energy consumption of the RSSI estimation by one order of magnitude, as compared to the coarse-grain mechanism, while not compromising on its accuracy. The deployment has also led to the development of a promising composite metric, PRSSI, a combination between PRR (Packet Reception Ratio) and RSSI (Received Signal Strength Indication).

At the same time, the project provided us with the opportunity to acquire a series of practical lessons related to the specificities of outdoor urban deployments, that, due to their relative scarcity in the field, we have decided to document and disseminate.

In the following section, we will start by providing a more detailed description of the Fly project, and the way its specificities have translated into requirements from our routing protocols (Section 4.2). We will proceed to describing a series of previous approaches to optimizing the reliability-energy tradeoff in outdoor urban deployments, highlighting the state-of-the art mechanisms, and the way our protocols aim at building upon and improving them (Sections 4.3 and 4.4). Section 4.5 will elaborate on the integration of the two mechanisms (WiseNE and Rep), and the design of the composite metric, PRSSI, in the specific WiseFly context, followed by Sections 4.6 and 4.7 which will briefly describe the experimental setup, and then elaborate on the experimental results. Last, but not least, Section 4.8 will summarize a series of practical lessons, followed by the conclusions in Section 4.9.

The work presented in this chapter is a result of joint work with Damien Pigué, who provided valuable feedback on the design of WiseFly protocol, supervised its implementation and oversaw the development of Project Fly.

4.2 Project Fly

4.2.1 Project Description

Project Fly¹ seeks to upgrade the capabilities of an European company to remotely monitor its utility distribution systems (e.g., gas, water or electricity) by means of WSNs. Typically,

¹"Project Fly" is a pseudonym used to protect the identity of the industrial client.

checking the status of the system is slow and expensive, as it requires a group of operators to visit multiple network points scattered across a city and perform diagnostic measurements or retrieve logged data. Moreover, many points of the network might be underground or in transmission towers, thus becoming cumbersome to access.

This difficulties become a critical problem in case the network experiences a failure, since the operators must promptly detect the presence of the fault and locate its cause, in order to minimize the downtime of the service and the collateral damage.

A common solution is the use of measurement devices installed across the distribution system that are controlled centrally through the cellular network or wired Internet infrastructure, and can perform on-demand remote evaluations. Such systems reduce the human effort required to access the network, as well as the response time in the case of failures. However, the sensor density is significantly limited by the fees that must be paid to the cellular operators for accessing the network, the reduced battery-life due to the hefty energy consumption of cellular communications, and the lack of access to communication infrastructures in remote or underground areas.

A WSN provides an alternative communication infrastructure composed by low cost motes that do not require spectrum-usage fees. The economical advantage thus allows for a significant increase in the density of remote sensing points, thus minimizing the impact of failures.

4.2.2 Implications for the WSN Protocols

Area Coverage. The WSN used in this application must be capable to collect the data from multiple sensing units dispersed across a city and convey it to a central office. This requires the installation of a sink mote connected through a gateway to the company office. The WSN conveys the data towards one or more sinks, which then stores it in a database via the gateway. The need to cover a large area (tens of square kilometers in surface), beyond the typical range of a single mote (tens to hundreds of meters) [80] translates into a multi-hop WSN being necessary in order to relay the data through multiple intermediary motes. This is commonly addressed with a collection protocol able to build routes that enable each of the motes to communicate with the sink.

Multiple traffic schemes. The monitoring of the distribution system imposes several operations on the sensing points, each generating WSN traffic with a particular volume, temporal characteristics and final destination. The operations are the following:

1. **Keep-alive reports:** every mote in the network must send a keep-alive message every 6h, to confirm its connectivity to the operator. A sudden absence of keep-alive messages can be interpreted as a lack of connectivity of a mote or of an entire section of the network.
2. **Alarms:** a mote can notify the operator of an emergency condition by sending an alarm packet. Alarms are triggered by events such as failures in the distribution infrastructure,

or a low battery. They represent sporadic traffic with a volume in the order of tens of packets, and they are repeated until acknowledged.

3. **Audiogram streams:** the distribution system can record audiograms upon demand, which enables the operators to precisely narrow down the location of infrastructure failures in the distribution system. Audiograms are large sets of data (typically 4 KB), and thus need to be fragmented in about 60-70 small packets (payload of 64 Bytes). Moreover, audiograms are sent by several motes simultaneously, thus representing sporadic and temporary traffic increases, significantly over the traffic baseline.
4. **Commands:** the operator must be able to communicate from the sink to the motes, in order to trigger actions that control the network, such as calibrating the sensors, or requesting audiograms. This generates traffic which traverses the network in a direction opposite to the collection traffic, as opposed to the one generated by the previous three operations.

Reliability. The traffic generated by the four operations mentioned above has a very sensitive nature, since a loss of packets can hide an alarm or compromise the quality of the service delivery.

The long term operation in an urban environment will result in the WSN undergoing numerous perturbations. For example, harsh weather conditions, such as snow covering the antennas and rain degrading the quality of the wireless links, or vandalism, can cause instabilities in the collecting topology due to the sudden removal of a key mote that relays a large portion of the network traffic. Another important factor that is particular to urban deployments is the temporary unavailability of links due to blockage by pedestrian or vehicle traffic.

These perturbations are exacerbated by the fact that the motes may not always have LOS between the antennas, and will thus need to communicate via reflections of the wireless waves. Relying solely on reflections significantly degrades the quality of the link, as it reduces the power of the signal upon reception, and increases the sensitivity to changes in the multipath fading pattern. The multipath fading pattern results from the superposition upon reception of many reflected components of the transmission. It can be easily altered by the movement of elements reflecting the signal, such as cars or pedestrians.

In project Fly, the lack of LOS between neighboring motes is the norm, rather than the exception. This is also caused by limitations in the infrastructure available to attach the motes, which are determined by the institutions involved in the project. The installation is limited to fixed stations at ground level and certain lamp posts used for public illumination. Additionally, the thick rock walls in European old cities make it impossible for the signals to traverse buildings, thus further limiting the availability of direct paths between motes. Additionally, the personnel that plans and executes the installation showed the tendency to stretch the distance between motes above their nominal transmission range, resulting in motes getting close to a transitional behavior (details on transitional links can be found in p.13).

Battery life. Changing the batteries of the motes implies the deployment of a team to visit each mote, which can be time consuming and even a dangerous operation for motes located in zones with difficult/restricted access.

Additional features. The application also required the network to be able to synchronize on demand, as well as display high levels of data security.

4.3 Problem Statement

Urban deployments represent challenging environments for WSNs, due to high levels of multipath fading and unpredictable events, ranging from weather, to perturbations caused by human activity or even vandalism. The limited positions in the city infrastructure available for the client to install the motes, generated a deployment where motes seldom have LOS and communicate using reflections. This condition exacerbates the sensitivity of the link quality to dynamic environments, due to multipath fading.

This effect, corroborated with the high cost associated with larger scale outdoors deployments limits the availability of experiments under these conditions. Up to date, we have been able to find very few such studies. Moreover, the studies found were showing very high variability in terms of reliability [78] [81] [77].

A limitation in terms of performance is also driven by the difficulty in creating protocols that can show superior reliability without causing prohibitive energy consumptions. While the use of composite metrics has proven superior in indoor environments, it is broadly believed that RSSI, a key component of high accuracy composite metric, would be additionally challenged by the sensitivity to multipath fading.

Our deployment thus aims at proposing and evaluating a mechanism that can enable the use of higher reliability LQE metrics, in our case RSSI, and maintain low energy consumptions, in spite of the challenging multipath fading conditions. This will be based on the findings of the previous two chapters - WiseNE and Rep, whose superior performance will be validated in a highly dynamic environment.

4.4 State-of-the-art Studies in Real Deployments

There are few reports of WSN deployments in outdoor environments subject to multipath fading. Moreover, as reported by multiple authors [82], the few real world deployments available have tended to be challenging due to the variability of outdoors conditions and significant perturbations due to multipath fading.

CTP is the default collection protocol in TinyOS and it has been extensively used in the literature as a benchmark. CTP uses the ETX as a routing metric. Despite numerous reports of CTP providing a reliability over 99% in simulations or indoor deployments, there are few

experiments in urban environments prone to multipath fading.

Istomin *et al* [83] performed a comparative evaluation in a smart city scenario of several of the most used routing protocols (including TinyRPL, ContikiRPL) that rely on the ETX metric, showing that the reliability of both protocols decreases from 90%, with low multipath fading, to 25%, as multipath fading increases. The authors conclude that one of the main causes of packet loss are the topology reconstructions, normally caused by fluctuations in the ETX due to multipath fading. The performance statistics were also obtained via simulations, due to the difficulty in accessing the motes, once they had been deployed. The study thus does not factor in the energy consumption, as the motes are expected to have access to the electrical infrastructure of street-lamps, even once deployed outdoors. This analysis highlights the importance of using accurate composite metrics in order to minimize the need for topology corrections.

PRR-based metrics tend to provide lower levels of accuracy. A common appreciation [25][28] is that PRR-based metrics are less suitable for such deployments given its difficulty in distinguishing between links that are good, but close to or in the transitional region, from links that are robust and very far from it.

RSSI is difficult to rely upon under high levels of multipath fading. A common concern in studies attempting to use RSSI in order to improve WSN reliability is that its high sensitivity to multipath fading would make it very difficult to use in environments with high levels of multipath fading. This observation was also expressed in [39].

Alternative protocols. A relevant alternative for enabling urban deployments could also be opportunistic protocols, such as ORPL [84], Orinoco [85] and ORW [44]. Opportunistic protocols are designed to have each mote forwarding the traffic to any available neighbor that represents a progression towards the sink, as opposed to having each mote predefining a favorite parent. Their use is appealing because they circumvent the long-term link estimation by leveraging the intervals when a link to a neighbor displays a high quality. Their development was motivated by applications with very high levels of variability (e.g., WSNs embedded in wearable devices). Nevertheless, they rely each data packet using anycast primitives on every hop, which enables targeting more than one neighbor, at the expense of a significant energy overhead compared to the use of unicast. Moreover, the use of anycast results in significant duplicated traffic, which further increases the energy overhead. This design-driven overhead thus represents a prohibitive, and in our case unnecessary cost.

A promising approach to cope with the effects of multipath fading is also the exploitation of channel diversity. Motes in protocols such as Oppcast [86] and IPv6 Time Slotted Channel Hopping (6TiSCH) [87] send their packets through several channels to a parent, expecting to increase the probabilities of finding a channel that is unaffected by multipath fading at that given time. We do not consider Time Slotted Channel Hopping and other protocols that require a tight and permanent network synchronization to ensure a coordinated use of the channels, the reason being the significant overhead required to maintain the synchronization

through the network, which pays-off when the traffic is periodic or predictable, but is not justified with the sporadic traffic expected in our application. Even though our deployment can use network synchronization, this is performed only on-demand and it is valid for a predefined time interval, thus it cannot be used to support coordinated channel hopping.

Additionally, the energy-efficient fine-grain link quality sampling brought by Rep (presented in chapter 3) and used in this chapter, can also be used to improve existing protocols that require a fast and reliable link estimation in multiple channels, such as [43].

Link Quality Ranking (LQR) [43] is a mechanism proposed as a complement for opportunistic protocols in scenarios where unreliable links are ubiquitous. LQR aims at quickly identifying the best link available at a given time, by probing them with multiple broadcasts and forwarding a short data stream through the top ranked link. The use of LQR is however not well suited for our deployment, as we do not need to perform a link estimation for each packet (we can leverage longer periods of channel stability). The sparsity of the typical application traffic would thus bring an unnecessary control overhead for ranking the quality of each packet on each hop.

A family of solutions that is gaining popularity for urban deployments is Low-Power Wide Area Networks (LPWAN), where LoRaWan and SigFox are amongst the most popular ones [88]. LPWANs enable long range communications at a low bit rate and, typically, rely on unlicensed spectrum, thus not requiring paying usage fees. Even though this solution could support the sending of sporadic alarms, it is not feasible for Project Fly because the low bit rate generates a prohibitive delay for sending data-heavy audiograms.

4.5 Design of the WiseFly Collection Protocol

The WiseFly protocol makes use of PRSSI, a composite metric based on the state-of-the-art PRR (Packet Reception Rate, commonly used to approximate $ETX=1/PRR$ [23]) and RSSI (Received Signal Strength Indication). WiseFly incorporates the two main contributions of the previous two chapters, validating their performance enhancement potential in terms of link estimation accuracy and energy efficiency in the presence of high multipath fading levels specific to urban outdoors deployments.

With the help of WiseNE - the neighborhood exploration mechanism proposed in chapter 1 - WiseFly manages to provide sufficient samples for the meaningful use of RSSI (a metric whose addition to PRR has a significant reliability enhancing potential, but is, at the same time highly sensitive to multipath fading), while keeping the Trickle Algorithm, thus generating a significantly lower overhead.

At the same time, with the help of Rep, the protocol manages to avoid significant overhead, by leveraging the preexisting repetitions in the MAC layer. Our results show that, under the multipath fading patterns found in the urban outdoors environment of a European city, Rep's

fine-grain sampling over 1 second results in an equivalent multipath fading mitigation performance to a coarse-grain sampling over 10 minutes, while reducing the energy consumption by an order of magnitude.

4.5.1 Proposed Composite Metric - PRSSI

While the previous chapters have mainly focused on single link estimations, as required by node-to-node communication, this larger scale deployment will require the use of full fledged collection protocols and the respective routing metrics. A routing metric is the input that a mote uses in order to select a route to the sink, and is obtained on the basis of multiple link metrics.

A routing metric defines a trade-off between the path length (number of hops) and the quality of individual links that compose the path [25]. For example, prioritizing on selecting high quality links might result in paths with a large number of hops, which might seem less energy efficient, but more reliable than a shorter path that includes a single low quality link. We aim at maximizing the link quality, motivated two-fold: i) the priority requirement for the routing protocol is the reliability (as defined in Sec. 4.2.2), and ii) there is no guarantee that a shorter path will result in energy savings, since low quality links increase packet retransmissions and can compromise the stability of the topology.

Our search for a metric that enables a mote to identify the link with the highest quality, has landed on RSSI. This decision is motivated by 1) extensive reports proving that RSSI can provide an accurate measure of whether a link is of very good quality or not, 2) the fact that RSSI tends to be a core building block of multiple other metrics e.g., (four-bit [24], Rondinone [27] and TM [22]), and 3) the fact that it can be easily obtained from the transceiver. We have decided not to use TM in the Fly Project, because the radio transceiver (CC1125) selected by the client does not provide a wide range of LQI values (we measured from 0 to 7), compared to 0 to 120 used in the original TM design (with the cc2420).

While RSSI has proven to be a meaningful enhancer of state-of-the art single metrics (including by our experiments in Chapter 1), it is highly susceptible to multipath fading [89], which translates in the need to acquire multiple samples in order to obtain a reliable measurement. The combination of RSSI and PRR emerges as a desirable choice, given their synergies in terms of prediction power. This creates PRSSI, a composite metric which provides a holistic description of link quality.

We define the PRSSI of a link $i \in path$ as shown in Eq. 4.1, which will later be validated in simulations and indoor deployments before being installed in our outdoors application. We use the mean RSSI (\overline{RSSI} in dBm) as the indicator of the quality of the link (the higher the better [23] [39]) and combine it with the PRR by penalizing the \overline{RSSI} proportionally to the number of packets lost, hence subtracting 10dBm for every 10% in the Packet Error Rate ($PER = 100 - PRR$). It was not possible to perform a parametric study on the penalty value in

the real deployment, due to the restricted access to the network. Therefore, we chose the value 10dBm to strongly penalize the loss of even a single packet, and it corresponds to the width of the transitional region observed in the deployment (see Fig. 4.8). Eq. 4.1 uses a division to represent the cost of the link (the lower, the better).

$$PRSSI_i^{dBm} = 1/(\overline{RSSI} - PER * 0.1dBm) \quad (4.1)$$

We follow the composition method proposed in [25], as shown in Eq. 4.2, since the use of the addition operation requires less computation resources from the microcontrollers used in the motes, as opposed to multiplications (e.g., the composition method used in [27]). We convert the value $PRSSI_i^{dBm}$ from dBm to mW ($PRSSI_i^{mW}$), in order to accentuate small differences in RSSI values. The resulting value is expressed in units normalized with the factor 10^{11} .

$$PRSSI/RM_{path} = \sum_{i \in path} PRSSI_i^{mW} \quad (4.2)$$

4.5.2 Enabling PRSSI through WiseNE

WiseNE is the Neighborhood Exploration mechanism proposed in Chapter 1, that provides an accurate use of RSSI, and thus also of RSSI-based composite metrics such as PRSSI in link quality estimates (by providing the uplink RSSI and PRR), while avoiding the previously punitive increases in energy consumption. By decoupling the observation and exploration periods of each mote, WiseNE allows the motes to continue benefiting from the energy performance of the Trickle Algorithm (driven by its exploration period), while achieving a sufficient number of samples in the observation phase for a reliable measurement of the multipath fading-sensitive RSSI.

WiseNE also handles the bootstrapping and adaptation of the topology. The protocol can deal with a rapid addition of motes, the reason behind its proactive philosophy. The use of WiseNE also relies on the assumption that the network is not congested. This depends on the traffic and density of the network, but has not proven to be a challenge in our experiments, despite the high traffic levels observed during the streaming of audiograms.

4.5.3 Link Quality Sampling Based on Rep

Presented in Chapter 2, Rep is a novel channel sampling mechanism that manages to extract the link quality information from the existing packet repetitions in preamble sampling medium access control (MAC) protocols. Rep extracts useful link quality statistics, such as the average LQI and RSSI, from a single packet in low-power preamble sampling MAC (LP-MAC) protocols (e.g., BMAC [48], WiseMAC [45] and ContikiMAC [46]) (Fig. 3.1, bottom). The mechanism leverages a resource that is commonly available in state-of-the-art protocol stacks, yet so far underused, namely the fact that LP-MACs transmit the exact same packet several times

(hereafter referred to as packet repetitions) in order to wake the intended receiver(s) up.

The experimental evidence provided in Chapter 2 shows that Rep can reduce the traffic and duration of a link estimation by one order of magnitude, while at the same time keeping an equivalent accuracy to the state-of-the-art. The experiments provided in the current chapter show that Rep manages to maintain its multipath fading mitigation effect and enable a reliable RSSI estimation under highly variable environments such as our outdoors deployment. Rep's ability to use fine-grain sampling over 1 second and still provide a multipath fading mitigation effect on RSSI equivalent to a coarse-grain sampling over 10 minutes also relies on the short term variation in the multipath fading pattern of our urban outdoor environment.

4.6 Experimental Setup

WiseFly has been tested in the frame of Project Fly, a Smart City initiative run in an European city of approx. 20'000 inhabitants, with the objective of optimizing the management of one of its public utility systems. We have analyzed the network over the period of one month in early 2018. The deployment had a single sink (M1000) and 15 motes.

In terms of protocols, we have used WiseFly together with WiseMAC (sampling period 1s). The *more bit mechanism* in WiseMAC provided support for the traffic increases due to audiograms. More bit is an energy optimization mechanism for streaming packets. It is triggered by a sender mote to signal to a neighbor that it must remain listening after receiving a packet, because a stream of frames will follow suit [90].

In terms of hardware, we have used a tailor-made mote based on the microcontroller MSP430F5438A and radio CC1125 (868 MHz, narrowband, FSK modulation). The modules have an external antenna, type standard omnidirectional 868 MHz. The operating frequency was helpful in terms of handling long distances (bitrate = 50 Kbps, sensitivity = -100 dBm. The transmission power is 14 dBm, unless otherwise noted.

However, before deploying the protocol in the beneficiary's network, we have first validated the newly proposed composite metric, PRSSI, indoors.

4.7 Results

The experimental results have validated two main contributions: 1) The WiseNE-enabled PRSSI results in a higher link estimation accuracy than the state of the art ETX, even in a noisy environment previously believed to hinder the use of RSSI and 2) Under the multipath fading pattern of our urban outdoor deployment, the fine-grain sampling mechanism (over 1s) proposed by Rep results in an equivalent multipath fading mitigation effect as a typical coarse-grain sampling over 10 minutes, while consuming one order of magnitude less energy.

The latest implemented WiseFly protocol based on a combination of WiseNE (indeed applied

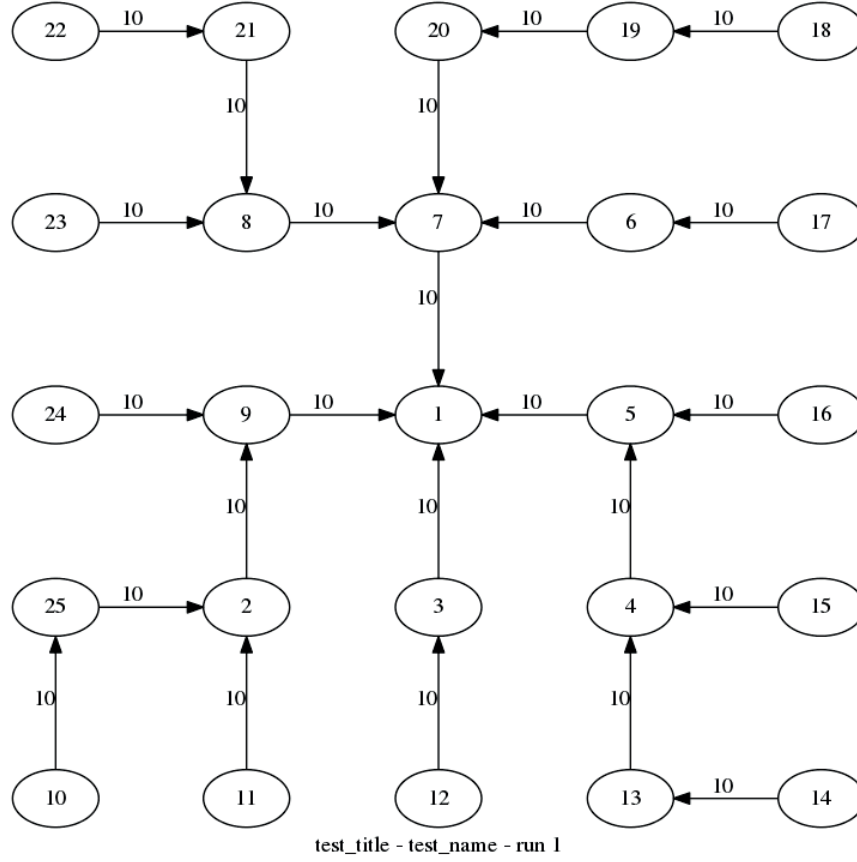


Figure 4.1 – Sample of a typical routing tree provided by the PRSSI metric with 25 motes.

to the simple RSSI metric) and Rep has been functioning at a reliability of 98% over the testing period. While this performance is very encouraging, its generalization to any urban outdoor environment and in-place comparison with reference collection protocols will require further testing, not yet possible under the current operation of our client deployment.

4.7.1 Indoors Pre-validation of PRSSI

Before deciding to deploy PRSSI in the real-life deployment of our beneficiary, we have performed a pre-validation, both in a simulated and in an indoors deployment. The testing of WiseFly² using the metric PRSSI/RM in OMNet++ simulations [91] consistently provided routing trees where each mote chose one of its closest neighbors. A representative example is shown in Fig. 4.1. This result was sustained for multiple topology sizes (25, 50, 75 and 100), with 32 experiment repetitions each. The simulated topologies were grids with the sink in the center and subject to traffic with a period of 300s from each mote (25 bytes payload). The channel mode used was "simple path loss", without any stochastic component or external interference. Therefore, all the simulations resulted in PRR=100%.

²This version did not use the Rep optimization, since we preferred to first test each mechanism separately.

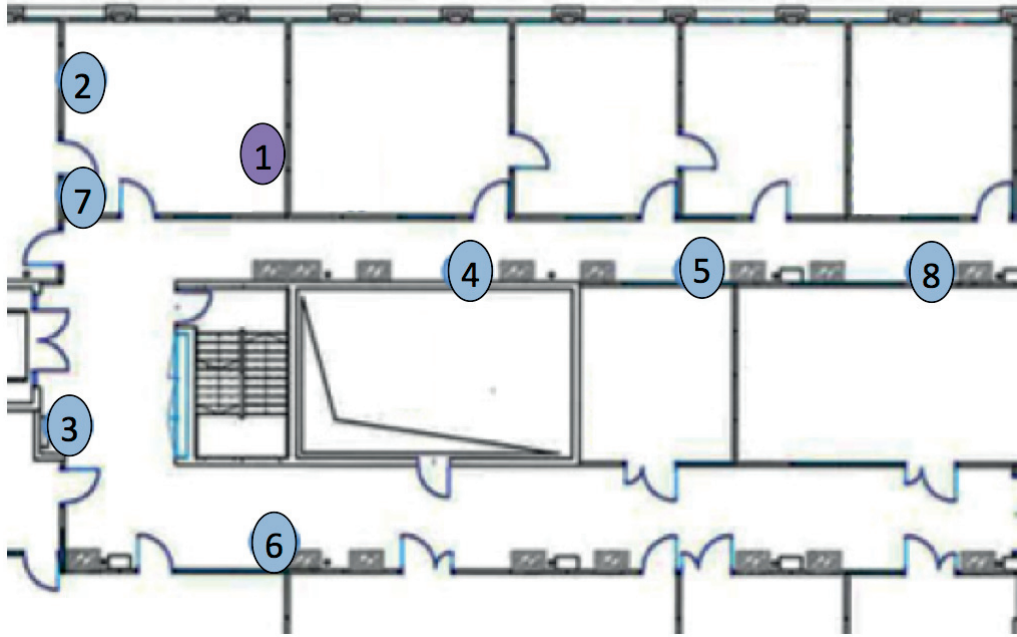


Figure 4.2 – Placement of the motes in the csem-tb deployment. Mote 1 (in purple) is the sink, the rest of the network requires traversing multiple hops to reach it.

Instead of continuing the simulations with a more challenging simulation channel, we have decided to repeat the experiments in a real, yet indoors deployment. Our rationale was that a real experiment would be more relevant for metric design, since the evaluation is highly dependent on the characteristics of the channel, which can only be partially captured in simulations. Therefore, we tested WiseFly, the same version used in simulations, in a local testbed (hereafter called csem-tb), shown in Fig. 4.2. Csem-tb was composed of 8 motes, whose hardware is detailed in Sec. 4.6. The motes were placed in an office hallway with metallic walls, at varying distances between 4 and 20 m. Each mote was sending packets with a period of 30s (payload 57 bytes) at a transmission power of 10 dB.

The csem-tb deployment revealed that the routing protocol strikes a packet error rate of 0.29% after the transmission of 15'789 packets (46 packets lost). Moreover, the motes find stable routes with the spatially closest neighbor, which tend to be long, but have proven to be reliable.

Both the simulations and the testbed (csem-tb) results show that PRSSI/RM is successful in providing high reliability routes, thus validating the value of adding signal strength information to the link quality estimation data. Moreover, the preferred routes tend to be long (in number of hops), as a consequence of having each mote preferring its closest neighbor in the parent selection.

Unlike the outdoors deployment, in the case of the testbed, we could make sure that each mote was always counting on at least one neighbor within range and with LOS. This is an ubiquitous assumption in both simulations and public testbeds in the literature (representative examples include [42], [61] and [51]). However, while this was a specification for the installation of the motes, it could not always be ensured, being out of our control. This is a realistic feature of

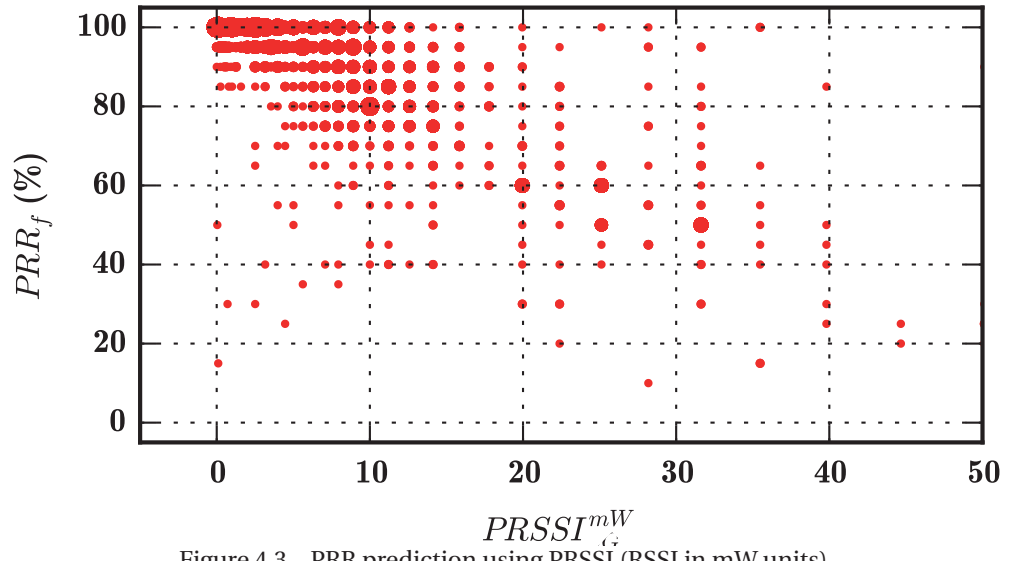


Figure 4.3 – PRR prediction using PRSSI (RSSI in mW units).

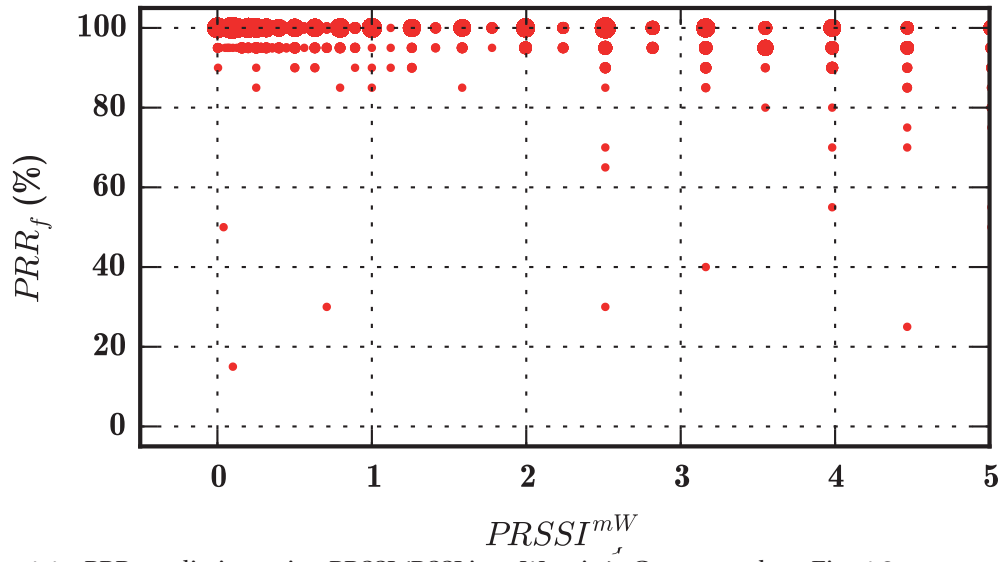


Figure 4.4 – PRR prediction using PRSSI (RSSI in mW units). Corresponds to Fig. 4.3 augmented in x-axis.

outdoors deployments, which only comes to underline the significant difficulty of catering for outdoor WSN deployments.

4.7.2 Benchmarking PRSSI to ETX in Project Fly in an Urban Environment

This experiment was designed to benchmark PRSSI's ability to predict the reliability of the links in a deployment subject to high levels of multipath fading, against the one of the non-composite metric ETX (de-facto metric for state-of-the-art protocols). The data for the analysis of both metrics has been collected using WiseNE, in order to ensure the availability of the RSSI inputs. Hereafter all the results reported correspond to an outdoor network deployed in a urban setting as described in Sec. 4.2.

We use the metrics extracted from a window of 10min (20 packets, coarse-grain as explained in Chapter 3) and from the first packet of the window (20 repetitions, using Rep), to predict the PRR of the next window of 10min (20 packets, PRR_f or future PRR). The PRR is obtained using CG, and the RSSI using Rep. The Rep and CG input data for the figures was registered simultaneously during the same WiseNE rounds, in order to enable a fair side-to-side comparison. Each plot contains 4'697 points, resulting from the same number of WiseNE rounds. The rounds were performed by 9 motes during 4 days (521 rounds/mote). During this validation we were only able to use 9 of the 15 motes, as the remaining 6 were undergoing technical maintenance at the time the experiment was scheduled.

Fig. 4.3 displays the scatter plot of PRR_f vs. the current $PRSSI$. The radius of each point is proportional to the square root of the number of overlapped points, i.e., with the same x-y coordinates. We used the square root to prevent large points from occluding smaller ones.

If $PRSSI^{mW}$ has a value lower than 5, we can predict with a 99.5% accuracy that the Future PRR will be >90%. This can be seen more clearly in the zoom of the values near zero of $PRSSI^{mW}$ (Fig. 4.4). If $PRSSI^{mW}$ is lower than 0.1, we can predict with an accuracy of 98.3% that the Future PRR will be 100%. Therefore, $PRSSI^{mW}$ is able to identify good and very good links with a high accuracy.

The values of $2 < PRSSI^{mW}$ indicate that the link will present a bad reliability, under 60%.

The range of $PRSSI^{mW}$ values is significant. The good values can be arbitrarily close to 0, therefore the cost spans 11 orders of magnitude. The reason is the conversion of RSSI from dBm to mW in the calculation of the formula.

Figure 4.5 shows the PRSSI with the RSSI factored in dBm units ($PRSSI^{dBm}$) instead of converting it to mW (as in Fig. 4.3).

The figure displays similar prediction capabilities as $PRSSI^{mW}$. A disadvantage with respect to $PRSSI^{mW}$ is that the transitional region extends over a wide interval, while the "good behavior" region is compressed. This increases the resolution in the transitional region and

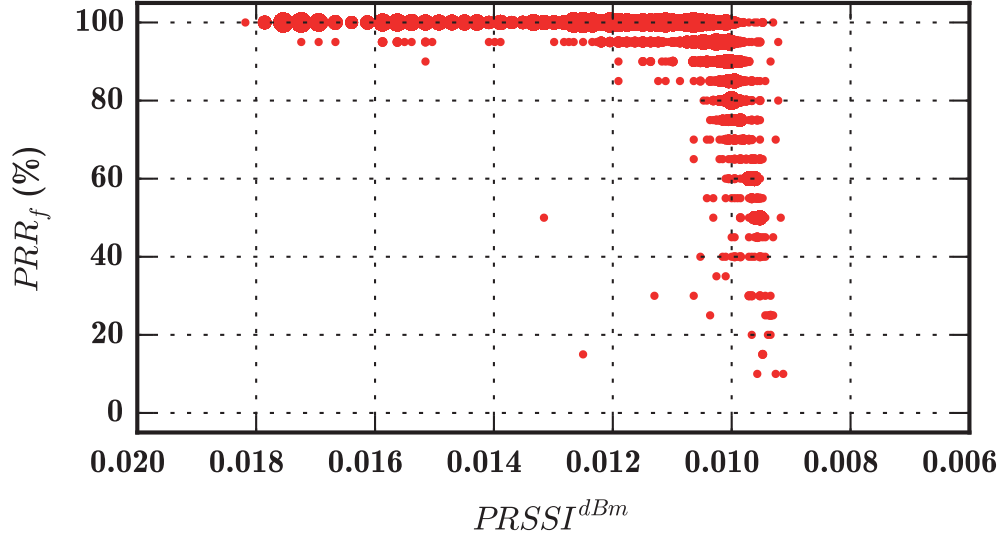


Figure 4.5 – PRR prediction using PRSSI (RSSI in dBm units).

compresses, without compromising the accuracy, the "good behavior" region.

A disadvantage of $PRSSI^{mW}$ is that the exponential behavior can bring instabilities when the addition operation is used for composition. This can cause significant variations in the cost of a route due to normal RSSI variations in our deployment, thus pushing the motes to overreact and unnecessarily search for a better parent.

Fig. 4.6, and its augmented view in Fig. 4.7, show the prediction capabilities of ETX, to be compared with Fig. 4.3. The links predicted to have the highest reliability ($ETX = 1$) present a wide variance of PRR_f values, ranging between 80-100%.

This shows that ETX alone is displaying a lower accuracy than PRSSI in identifying the good quality links (96.6% vs. 99.5%). The reason is that ETX cannot distinguish between a link that displays a good behavior, but is close to the transitional region and thus prone to degradation, and a good link that is stable. Adding the signal strength information allows for this, more granular distinction and thus improved prediction.

Studying in more depth the 5'261 points with $ETX < 1.11$ (i.e., $PRR > 90\%$), which are the ones expected to keep the maximum reliability in the future, we found 181 points that are worse than predicted ($PRR_f \leq 90\%$), of which 97.8% (177 points) present an $RSSI \leq -80dBm$. Leveraging the RSSI information would have thus reduced the ETX misclassification by 97.8%.

This result is in line with our hypothesis that using a composite metric that adds the signal strength information to the past reception statistics (respectively the RSSI to the ETX and / or PRR) significantly increases the link estimation accuracy, and consequently network reliability, in WSN deployments subject to high levels of multipath fading.

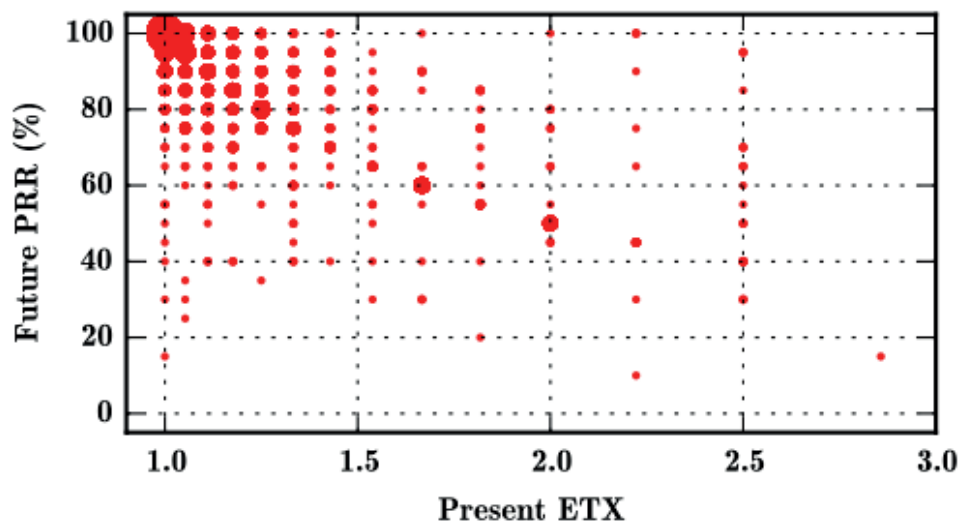


Figure 4.6 – PRR prediction using ETX.

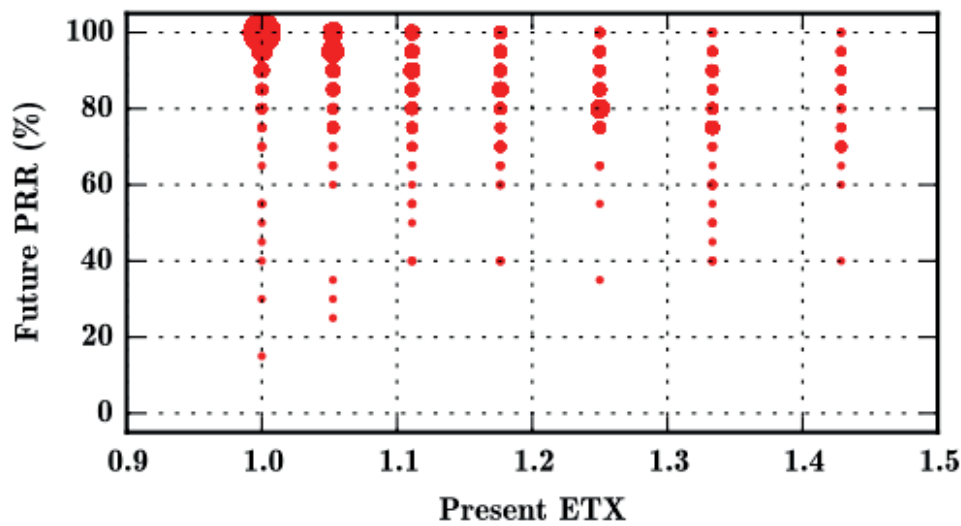


Figure 4.7 – PRR prediction using ETX. Corresponds to Fig. 4.6 augmented in x-axis.

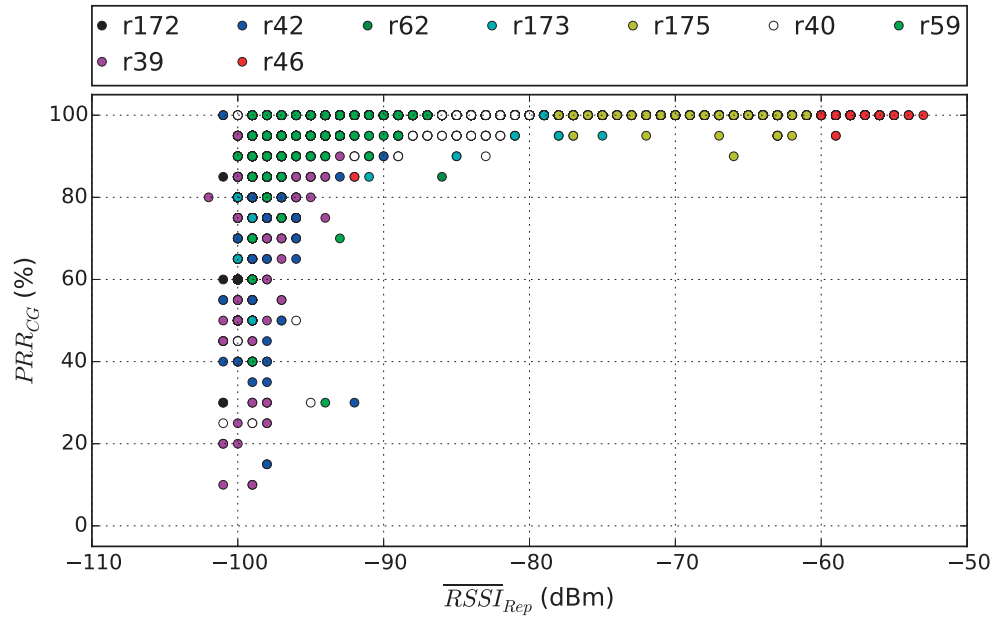


Figure 4.8 – PRR prediction using RSSI from Rep.

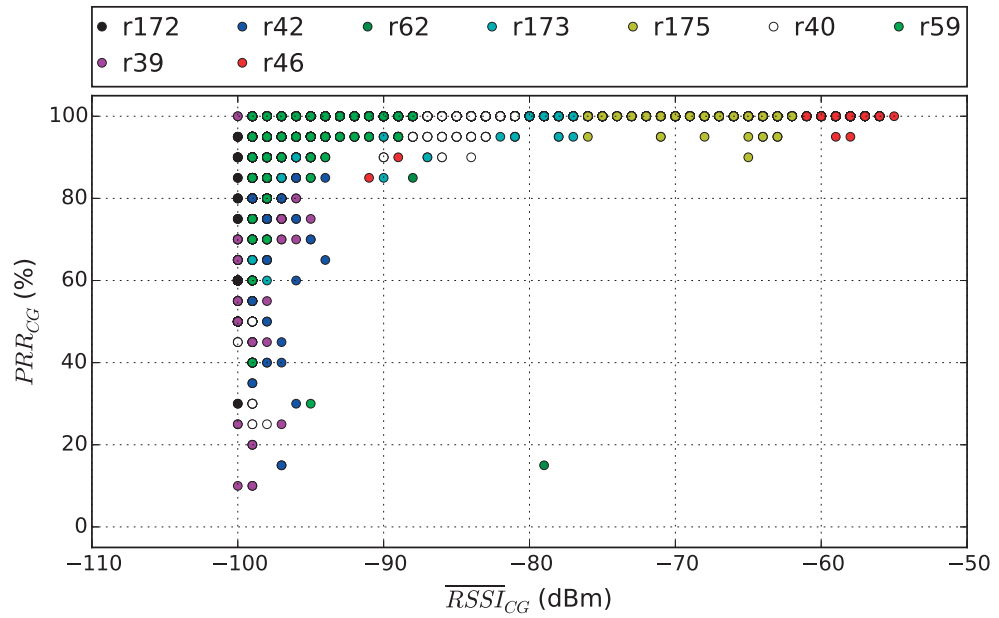


Figure 4.9 – PRR prediction using CG RSSI.

4.7.3 Fine-grain (Rep) Sampling vs. Coarse-grain

By using the Rep beaconing mechanism, we manage to drop the energy consumption of the link estimation operation by an order of magnitude, as compared to the coarse-grain mechanism, while not compromising on its ability to enable the use of RSSI. This confirms our hypothesis that fine-grain sampling has an equivalent capacity to mitigate the sensitivity to multipath fading of RSSI under the specific multipath fading pattern of our outdoors urban deployment, being relatively unaffected by the RSSI variability and link instability.

Figures 4.8 and 4.9 display the prediction capabilities of RSSI, collected by Rep and CG. Both curves present a future PRR in the y-axis over 90% for RSSI values over -95 dBm. The transitional region is visible for lower RSSI values, spanning PRR values from 5% to 100% and a width close to 10 dBm. No significant differences can be seen between both figures.

The RSSI information provided by Rep does not introduce a bias with respect to the one obtained by coarse-grain techniques. This can be observed in Fig. 4.10, which displays the mean RSSI obtained via CG (\overline{RSSI}_{CG}), as a function of the mean RSSI obtained via Rep (\overline{RSSI}_{Rep}). The blue line in the figure ($x = y$) represents the ideal scenario, where every measure done with Rep is identical to the one taken with CG. The line is within the standard deviation of all the measured points.

The mechanism explaining this equivalent result of the two sampling frequencies (fine and coarse-grain) is similar to the one detailed in Chapter 3. While there is no significant WiFi interference with RSSI (as the 868MHz band is basically interference-free [92]), the multipath fading phenomenon can be reasonably expected to show a similar, sub-second spikes pattern.

As motes in the deployment lack LOS among each other, this translates in the communication being carried over through multiple reflections, subject to multipath fading. Therefore, the quality of the link will be sensitive to elements such as pedestrian and vehicle traffic along transmission distances typically higher than 100m.

To date, there are very few studies characterizing deployments without of LOS in real environments at 868MHz, mostly given the difficulty in recovering the LQ information without a secondary testbed given the high energy consumptions. Most studies thus look at 2.4GHz, or in indoors environments. However, existing studies confirm the existence of gray areas in 868MHz [93], as a result of multipath reflections, and the fact that motes tend to show large variations in performance [92].

While we could not identify studies characterizing the variations over time of multipath fading in outdoor deployments, we find that a reasonable explanation is that the overlap of multiple and frequent perturbing factors (e.g., vehicle traffic) results in the multipath fading experiencing spikes patterns with high frequency, and thus periods lower than 1s.

Since Rep and CG were tested simultaneously to have a side-to-side comparison, the energy consumption in the network is not representative of either of them and cannot be used to

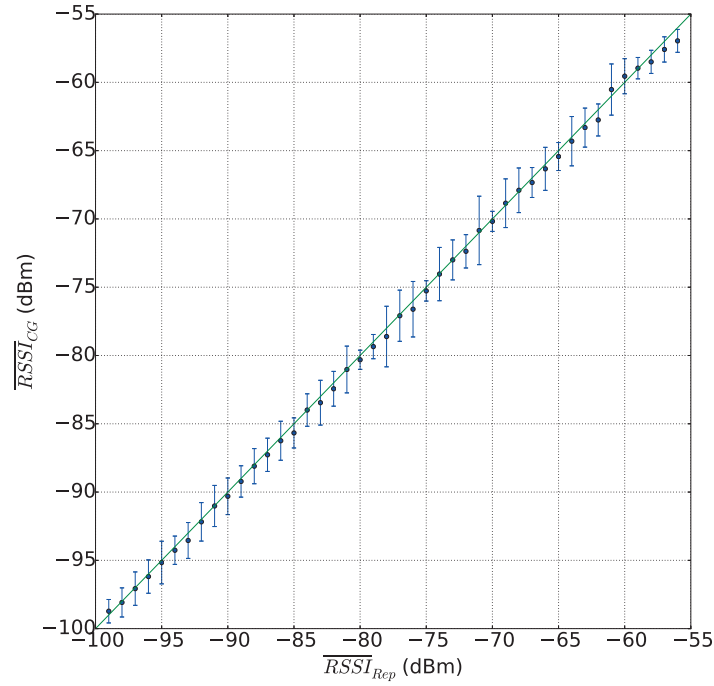


Figure 4.10 – Relation between \overline{RSSI} measuring using 20 samples obtained through Rep and CG. The line $x = y$ is shown in green for reference.

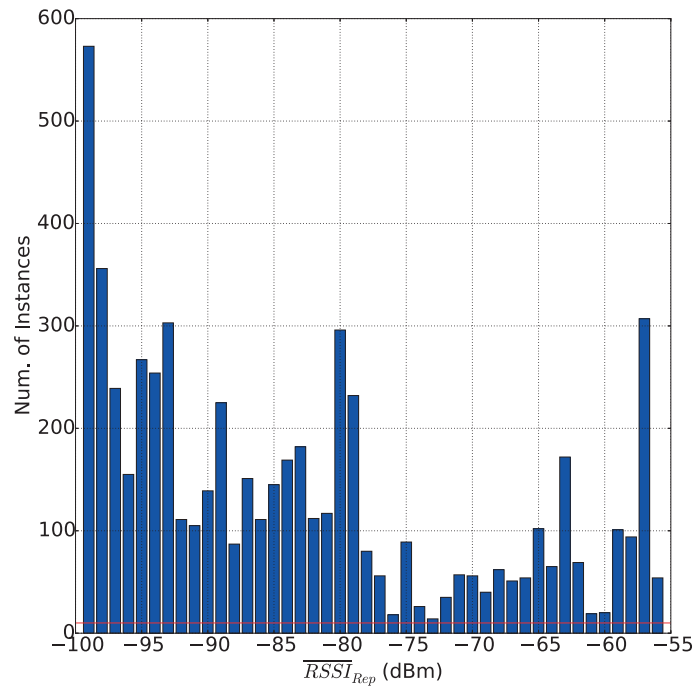


Figure 4.11 – Number of samples per \overline{RSSI}_{Rep} in Fig. 4.10, between \overline{RSSI} measuring using 20 samples obtained through Rep and CG.

compare energy performance. Therefore, we rely on the results provided in Chapter 3 for quantifying the energy savings of Rep. This means that Rep allows us to use one order of magnitude less energy in order to provide 20 samples of RSSI, compared with the case where the samples are obtained using CG.

4.7.4 Combined Performance Analysis of the Deployment – Rep and WiseNE-enabled RSSI

By combining WiseNE based on the (indeed single) RSSI metric, and REP, the system has proven a remarkable performance in terms of reliability, namely a mean PRR of 98% (as compared to the broad range of reliability measurements declared by typical outdoor deployments) and an energy consumption one order of magnitude lower than the course-grain reliant alternative.

In the previous sections and chapters, we have benchmarked the individual mechanisms that compose WiseFly (WiseNE and Rep) against the state-of-the-art. Nevertheless, testing the deployment of Project Fly against a state-of-the-art routing protocol in order to compare its global performance was not possible. Even though it is technologically feasible, our access to the network for testing was limited by it being currently used by our industrial customer. Benchmarking by using a third-party testbed was also not a feasible option because of the lack of public testbeds with a representative setting for Project Fly (i.e., in an outdoor urban environment). Neither was it possible to identify in the existing literature a representative proxy, with comparable requirements, installed in a similar environment and using the same frequency band. Therefore, we leave the benchmarking of WiseFly as a next step, but pursue to present its performance and validate the fact that it satisfies the requirements of the Fly project, as presented in Sec. 4.2.

Figure 4.12 displays the packets with management statistics received per mote during 30 days (period of this type of traffic: 12h). Mote 50 (M50) displays more packet losses than the others. This is due to the lack of neighbors with high-quality links.

Figure 4.12 shows that the entire network misses one packet at a similar time (it can be seen in the middle of the time interval). The cause of this loss is a full-network loop.

The loop begins when M62, which relays all the traffic from the network to the sink (Fig. 4.13), changes the parent to M44, thus relaying the entire traffic of the network through the loop. The loop is successfully detected and corrected (Fig. 4.14). The packet losses shown in Fig. 4.12 (1 per mote) are expected, since the tested version of WiseFly does not have a buffer to resend packets trapped in a loop after correction. The root cause of the loop is being further investigated, but the most likely reason is an increase in the RSSI, and the delay on back-propagating topology updates.

Figure 4.15 shows that the Packet Reception Ratio (PRR) is over 98% in all the motes except M50. These results include the losses due to the loop. At the same time, the total PRR in the

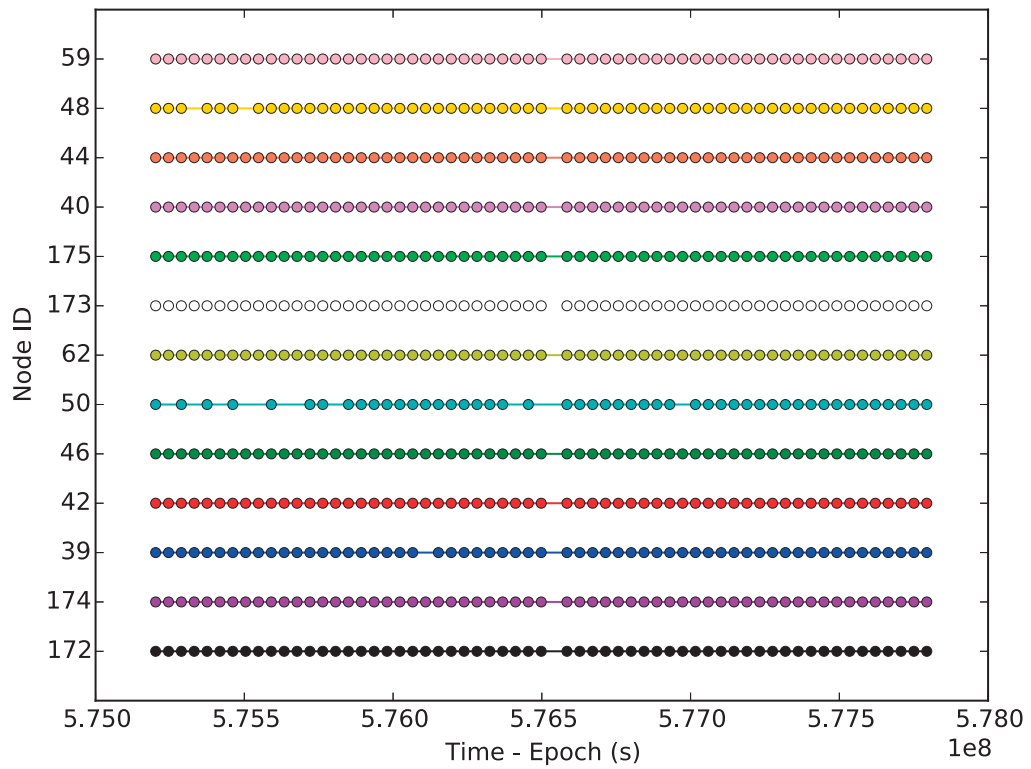


Figure 4.12 – Packet received at the sink (type Routing Statistics) vs. time. The graph covers 30 days in project Fly.

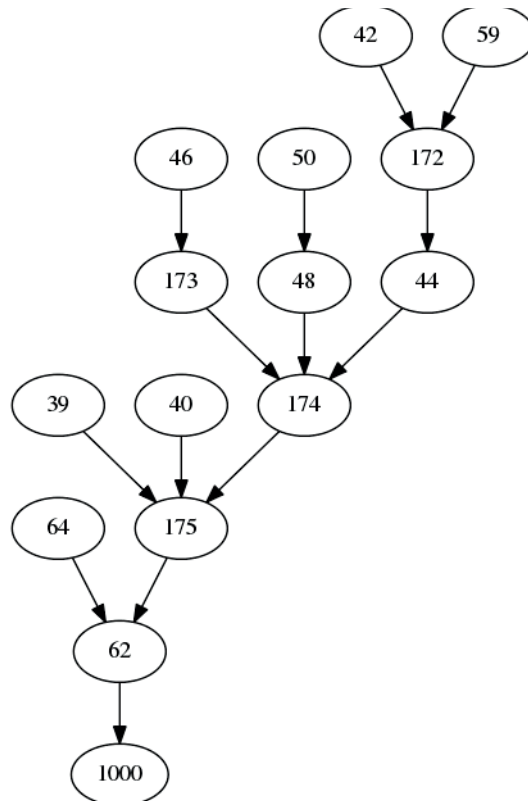


Figure 4.13 – Routing tree in project Fly before the loop.

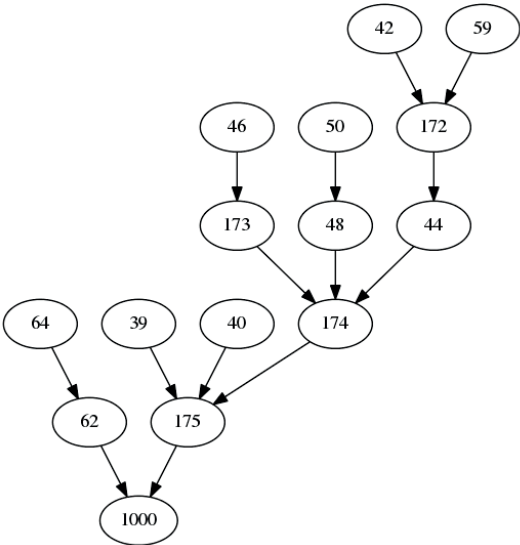


Figure 4.14 – Routing tree in deployment project Fly after the loop.

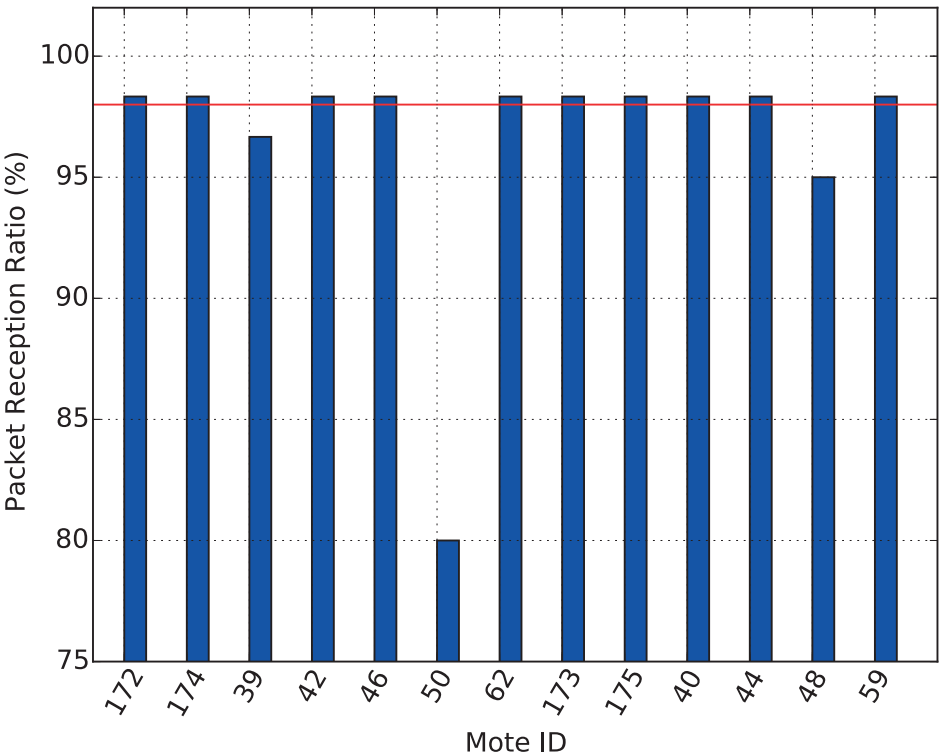


Figure 4.15 – Packet Reception Ration per mote in deployment project Fly.

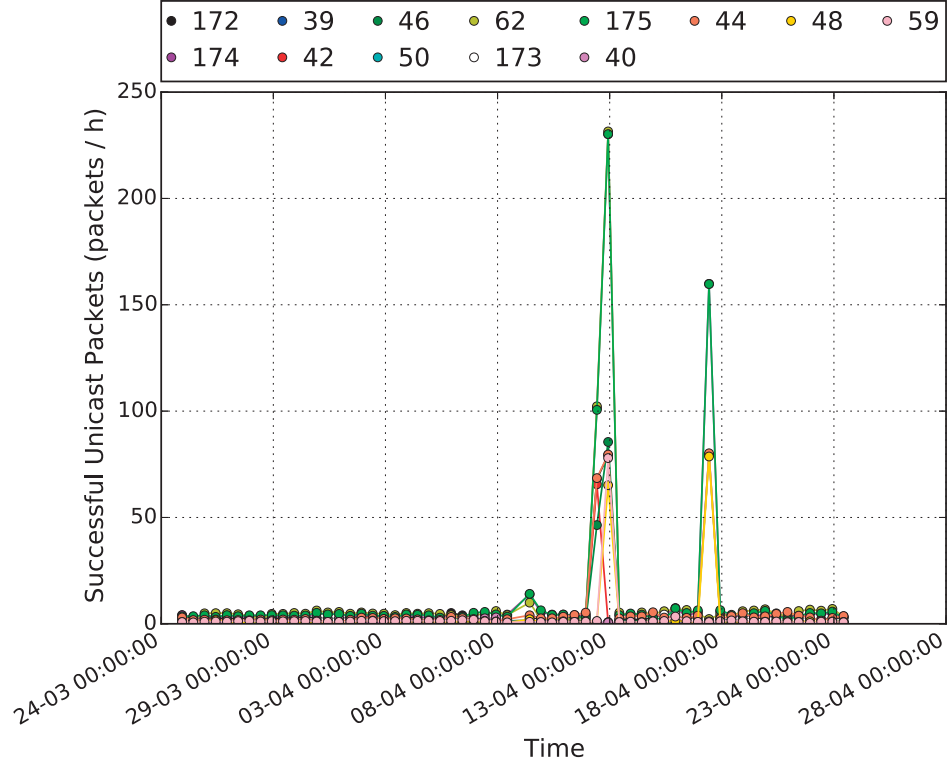


Figure 4.16 – Time evolution of the number of successfully sent unicast packets (ack received) per hour.

deployment is 96.54% (std dev = 23.70), which includes the packet losses due to the routing loop.

Figure 4.16 shows three sudden and temporary increases in the number of successful unicast packets (i.e., acknowledged). The first traffic increase (the smallest of the three) corresponds to the loop correction. The second and third increases correspond to the streaming of segmented audiograms.

The three traffic increases correspond to spikes in the radio duty cycle on listening or reception state (Fig. 4.17) and on transmission state (Fig. 4.18). The three increases also match corresponding increases in the number of packets using the "more bit" optimization in WiseMAC for increasing the streaming capacity (Fig. 4.19).

The PER is very low during the entire observation period, including the transmission of both audiograms and the energy consumed by a mote during the transmission of the audiograms shows an increase with the number of descendants.

A further analysis of the sending of audiograms reveals that there are two instances: audiogram 1 was received at the sink on April 12th, 2018, at 10:00 AM, and audiograms 2 on April 17th, 2018 at 3:00 AM. The network maintains the topology shown in Fig. 4.20 during the sending of audiograms 1. The PER in the MAC layer during the sending of each audiograms is under 0.2%, defined as the number of missing acknowledgments divided by the number of transmitted

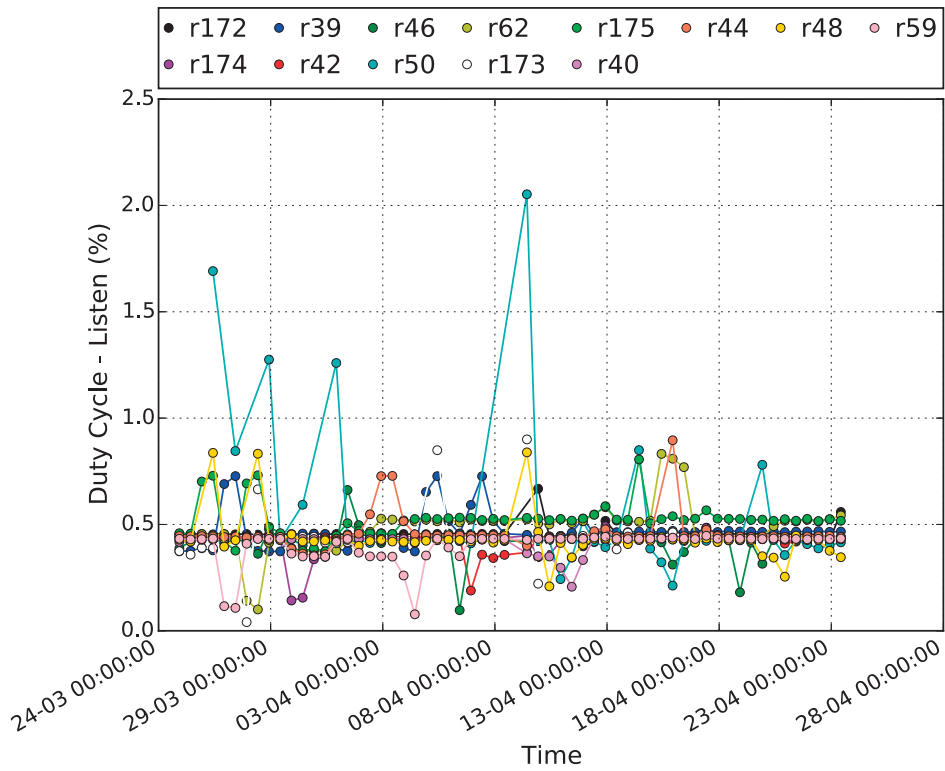


Figure 4.17 – Component of the radio Duty Cycle spent with the radio on listening or reception state (%) vs. time.

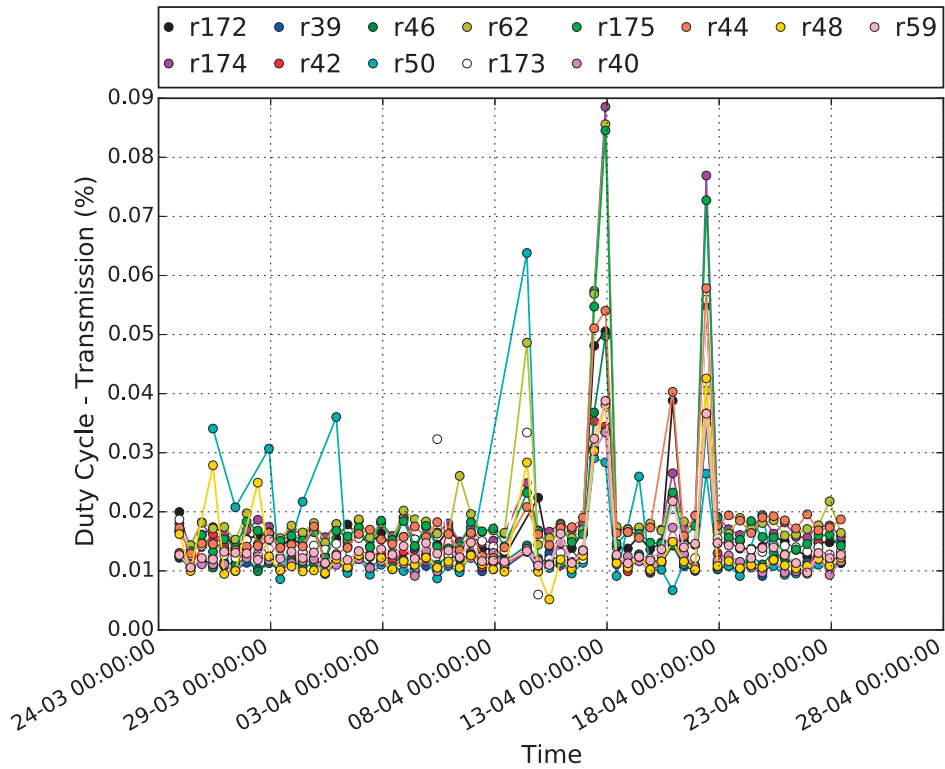


Figure 4.18 – Component of the radio Duty Cycle spent with the radio on transmission state (%) vs. time.

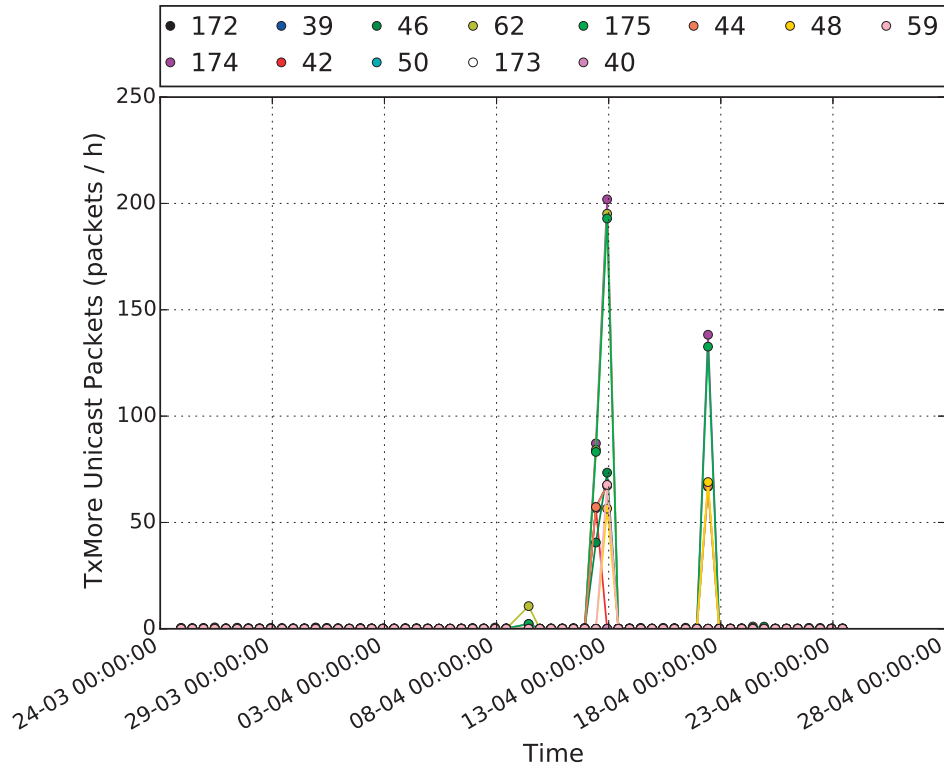


Figure 4.19 – Time evolution of the number of unicast packets sent with the TxMore optimization in WiseMAC vs. time.

frames, where each frame has an acknowledgment.

Figs. 4.21 shows the average and standard deviation of the latency of the motes, which displays the same pattern as the hop count (Fig. 4.22), as expected. Moreover, the latencies display maximum values well below 5s, which is satisfactory for the target of 20s maximum latency for delivering alarm messages. On average, traversing a hop in the network brings a latency of 0.5s, which is expected from the use of a sampling period in the MAC layer of 1.024s, and supports the assumption that the network is not congested. These results suggest that future deployments can cover significantly larger areas with deeper paths extending to 40 hops while still operating within the 20s latency target, as long as the network maintains a traffic pattern and density that enables a non-congested operation.

4.8 Practical Lessons

This deployment has revealed that, while simulations and the work on local testbeds are very useful for pre-validating results, real-world deployments often result in significantly different conditions. Moreover, the level of variability in the network inputs is far beyond what is typically used in simulators.

We could summarize our practical observations in a series of emerging "rules":

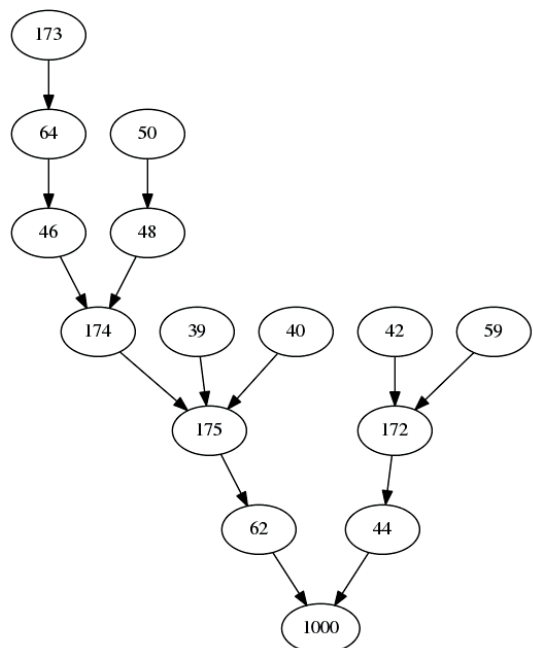


Figure 4.20 – Network topology during the sending of audiograms 1.

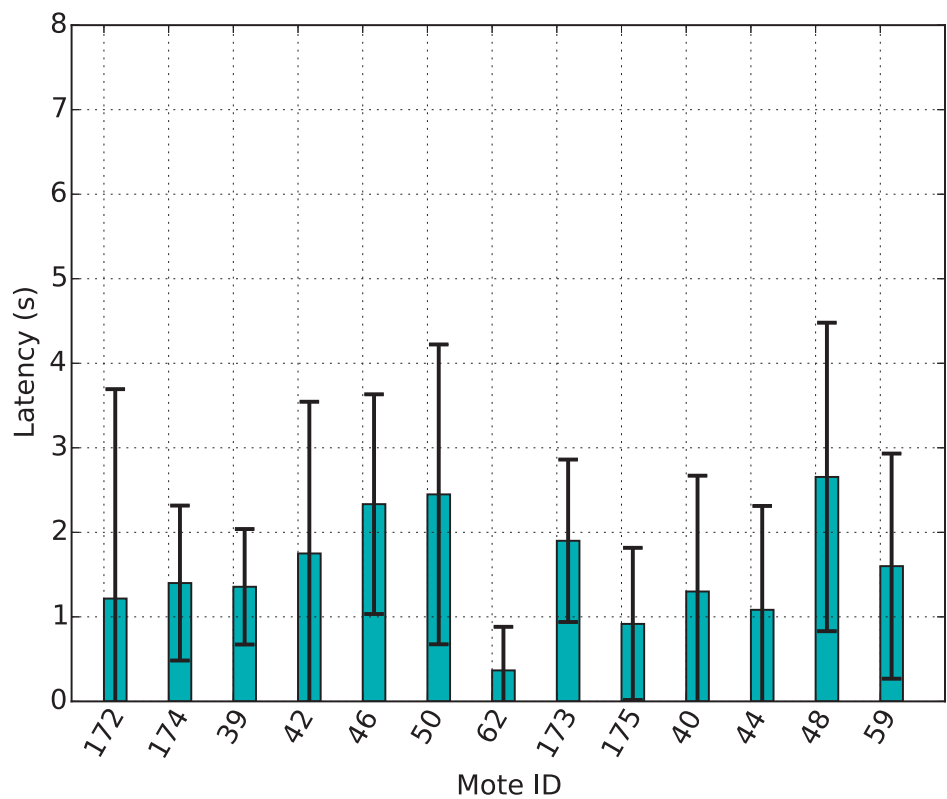


Figure 4.21 – Average end-to-end latency for each mote.

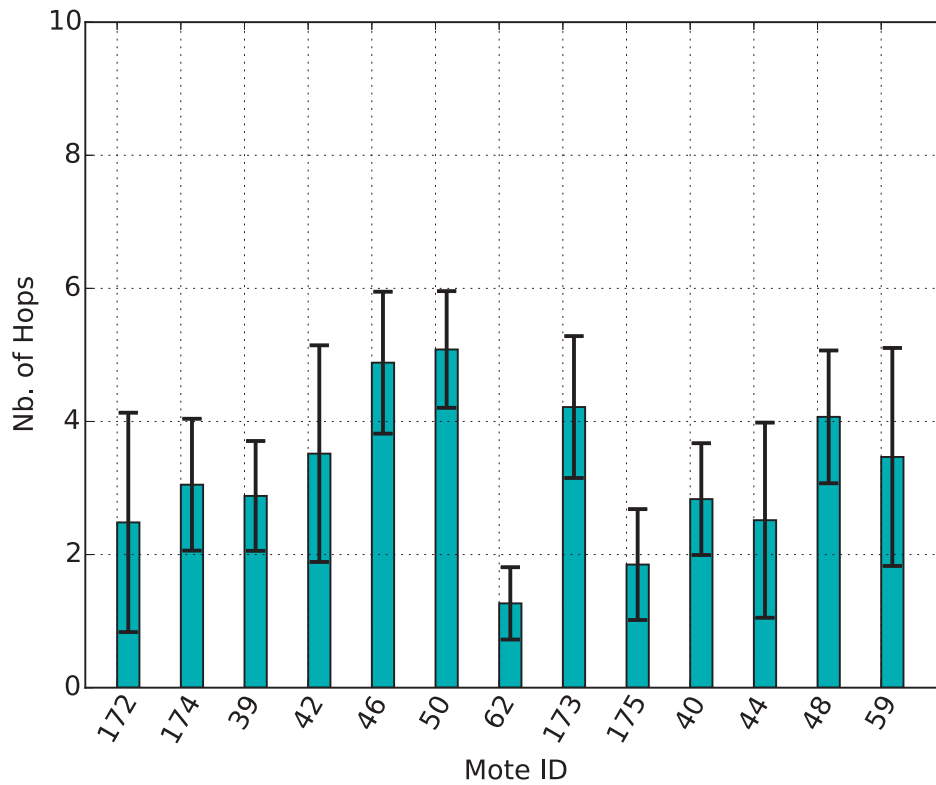


Figure 4.22 – Average hop count for each mote.

Foster a close collaboration with the beneficiary doing the deployment.

In real-life deployments, your involvement in the installation and / or operation of the network is more limited than in testbeds. This is due to multiple reasons, ranging from the specificities of the contractual relation with the beneficiary, to on-site confidentiality constraints. This necessitates a strong collaboration with the beneficiary, in order to both avoid unwanted network behavior and benefit from the data collected for scientific purposes.

In terms of testing, in the absence of a secondary infrastructure for collecting the performance statistics from the motes (as done in remote testbeds), the only feasible recovery mechanism is the collection topology. Logging is also a possibility, but it does not scale over time. Therefore, there is limited capacity to convey the statistics, as they are likely to be infrequent and/or to be deactivated by default.

The real network operator might not manually activate the performance statistics after reflash-ing of a subset or of the entire network, or after replacing a broken mote or adding new motes. Moreover, employees with no previous experience in the IoT or WSN domain might need support to store the data in a secure manner and learn the value of accumulating it. However, in one of our cases we have noticed beneficiaries only keeping the last month worth of data and discarding the rest.

Moreover, the network might end up being mobile, despite the application requiring a static

Chapter 4. Deploying Collection Protocols in Real World Applications – The Fly Project

one. Even though we plan for a static network, all the motes were flashed in a lab, stored in a box (thus becoming a super dense network) and then distributed progressively during the installation. The routing protocol must thus cope with this scenario. Some solutions that need to be communicated to the installation team include: resetting the network after installation, freezing the routing until the installation is completed or relying on the adaptation mechanisms of the protocols.

Last but not least, as the installation teams will likely stretch the links to the limit, it is reasonable to assume they will at first be transitional. It is also reasonable to assume that the installation team might not always be able to ensure a LOS between motes. In practice, the LOS might be even intermittent.

Consider the limitations of not having a secondary infrastructure.

Relying solely on network statistics for tracking the performance of the network requires a careful pre-design. In our experience, it will be necessary to change the content of the statistics packets often, as the required information to debug issues or evaluate the operation of a particular mechanism becomes more clear. Upgrading the statistics collected through the network implies changes in the packet format and in the structure of the database, which is significantly more difficult than the serial port used in academic testbeds.

Moreover, several variables that increase with time and reach large values (e.g., consumed energy requires 4 bytes) could be reported as the change since the last packet. This practice is recommendable, given that the reliability issues in the network can compromise the accuracy of the performance statistics.

Additionally, the statistics should be sent with a large period in order to minimize the overhead. This is also a challenge compared to testbeds, which do not have any period limitations. This should be factored in during the design of the statistics, since the long periods can hinder the characterization of target phenomena, and the troubleshooting in case of issues.

In any case, the custom motes should always have a USB port. This greatly simplifies the testing in the lab before the deployment, since it makes it possible to reuse the resources made available by the community for creating local deployments. Moreover, it simplifies on-site debugging.

The parameters that might have to be set remotely to tune the operation of the protocol in the field will require previous creation of set/get commands.

Other practical findings include the facts that:

1. The reliability of the routing protocol can only be as good as the available routes. Ergo, if the links are of a bad quality, the protocol will have a performance upper bound.
2. Despite the claims of the state-of-the-art in LQE, the LQI is not always available in

commercial transceivers.

3. It is worth planning for being able to capture data to work offline, thus interfering as little as possible with the actual deployment.

4.9 Conclusions

The possibility to develop and analyze a real world deployment for a month has given us the opportunity to validate two key learnings in environments of high variability: the composite metrics' (PRSSI) superiority in terms of link estimation accuracy, and the fine-grained sampling power to enable RSSI, even in environments subject to high levels of multipath fading.

Our experiments have shown that composite metrics such as PRSSI continue to bring an improvement to simple ones even in environments subject to high variability. This result was not obvious given the expectation that RSSI, a key component of composite metrics, but also a metric with a high sensitivity to noise, would be challenged by the exposure to the high multipath fading levels. Our results indicate that performing multiple samplings with the use of WiseNE and Rep manages to ensure a meaningful use of the RSSI metric. This was proven by comparing the reliability performance of a collection protocol based on the composite metric PRSSI to one based on the state of the art simple metric, ETX. While the PRSSI metric is subject to high variability, and needs further development, its composite nature results in a higher accuracy in predicting Rep's link behavior.

Moreover, by using Rep beaconing mechanism, we have managed to reduce the energy consumption of RSSI extraction operation by an order of magnitude, as compared to the coarse-grain mechanism, while not compromising its accuracy. This confirms our hypothesis that in our urban outdoor environment, performing a more energy economic fine-grain sampling within 1 second has an equivalent effect in terms of mitigating the multipath fading effects on the signal strength (RSSI), as a coarse-grain sampling over the period of 10 minutes.

By combining WiseNE based on the RSSI metric, and REP, the system has demonstrated a remarkable performance in terms of link quality estimation, namely a reliability of 98%, and an energy consumption one order of magnitude lower than the alternative coarse-grained approach. In interpreting this result we should of course consider the effects of the particular traffic and multipath fading patterns of our deployment. While generalizing this to any urban outdoor deployment requires further testing not yet possible in the frame of our real-life client deployment, we believe it opens a valuable area of exploration.

Beyond the scientifically relevant findings of this deployment, the experience has also resulted in extensive practical learnings. Due to the scarcity of such deployments, we found it meaningful to mention some of them, in the hope they will prove useful to the day-to-day work of our community.

Part II

Introduction

The second part of this thesis focuses on applications that are subject to stringent latency constraints, in addition to high reliability and energy autonomy, while at the same time requiring the feature of simultaneous event detection from multiple motes.

We focus on the particularly challenging application represented by event-triggered systems, where WSNs need to convey events that are not periodic, but require a time bounded response, as well as high levels of reliability and energy efficiency. Such systems apply to a broad variety of tasks, from automatic fire suppression [94], to landslide monitoring [95], or sniper shooter localization [96]. Motes in event-triggered systems wake-up periodically to detect transmissions that signal an event, but in many cases (such as the previous applications) most of the time there will not be events to report.

To cater for these applications, we have developed Gloosy-W, a protocol that leverages Glossy's superior energy - latency trade-off, while at the same time being compatible with simultaneous event detections. In addition to leveraging concurrent transmissions, Glossy-W uses floods to generate a unique temporal reference by designating one mote as a timekeeper in order to ensure reliable concurrent transmissions. This second part of the thesis thus provides the analytical and experimental evidence that motivates the use of flooding protocols based on constructive interference. Moreover, it introduces Synchronized Channel Sampling (SCS), a novel mechanism that reduces the energy consumption by 33.3 - 40 percent, while remaining compatible with simultaneous event detections.

5 Leveraging Constructive Interference for Improving the Energy-Latency Tradeoff

5.1 Introduction

This second part of our work focuses on applications whose performance relies on a time-bounded reaction, thus being subject to stringent latency requirements, on top of the high reliability and energy efficiency ones we have focused on in the first part. Common examples of such applications occur in the area of cyber-physical systems, that need to reliably provide a real-time response.

Our work was motivated by the requirements of WiseSkin [97], an interdisciplinary project aimed at the development of artificial skin for human prostheses, able to help the wearer regain the sense of touch. For the prosthesis to not be rejected by the patient, the WSN embedded in the skin needs to be able to relay the tactile stimulus fast enough to enable a timely actuation and create a seamless user experience. Moreover, the prosthesis needs to be able to rationalize its energy consumption, in a way that it does not necessitate frequent charging, another key determinant of patients' acceptability of the device.

In terms of system requirements, this results in a latency-energy consumption trade-off. Moreover, the density of the WSN embedded in the artificial skin, while paramount for its performance, results in a very high likelihood of simultaneous event detection, as a tactile stimulus is likely to affect an area with multiple sensing nodes. The protocols able to serve cyber-physical systems, such as WiseSkin thus often have to deal with two issues, the latency-energy consumption tradeoff and the robustness to simultaneous event detection.

To save energy, WSNs exposed to variable traffic are generally designed to operate in two modes: a low-traffic mode and a high-traffic mode. The more often a network wakes up to check for the occurrence of an event, the lower the latency. However, frequent wake-ups result in high energy consumption. Our work has thus focused on designing a low-traffic mode that operates within latency constraints, while minimizing the energy consumed, and not being impacted by simultaneous event detections.

Chapter 5. Leveraging Constructive Interference for Improving the Energy-Latency Tradeoff

The mobility of the motes in WiseSkin has motivated us to investigate the space of flooding algorithms, as opposed to the more classical tree-routing ones explored in the first three chapters. Flood-based protocols generally outperform routing ones in terms of latency and reliability, because they do not pose the risk of routes being unavailable at the time of transmission due to the change of position of the motes, a trigger of significant overheads. Flood-based algorithms, not necessitating to select and maintain routes, are thus well tailored to cater for our constraints.

In this context, we offer two key contributions. We show that, while simultaneous transmissions would be expected to result in high energy consumption for a given latency, leveraging them can actually result in better energy performance. The second contribution consists in showing that, among protocols leveraging concurrent transmissions, the ones relying on constructive interference are preferable to the ones based on the capture effect, given their superior robustness to simultaneous event detection, a key constraint of high density WSNs. Our experiments in a real testbed show a striking difference in reliability: 100% (constructive interference) vs. 0% (capture effect), for 6 or more neighboring motes.

These findings have been incorporated in the development of Glossy-W, a novel flood-based protocol tailored to cater for multi-node simultaneous event detection, while reducing the collision avoidance MACs' energy consumption for a given latency by 30%, for the dense WSNs under consideration. Glossy-W is the first protocol able to leverage floods that rely on constructive interference in a simultaneous event detection scenario. The protocol will be embedded in the WiseSkin pilot project to be tested by human patients. While inspired by WiseSkin, Glossy-W has a broad range of applicability, especially as an increasing number of WSN applications require a dense network, consistently fast responsiveness and high energy autonomy.

In terms of structure, this chapter will start by presenting the WiseSkin project, together with its requirements and their significance for the underlying routing protocols (Section 5.2). Section 5.3 will highlight the key challenges in fulfilling the system's stringent latency, reliability and energy efficiency requirements, followed by Section 5.4 which will present the existing state-of-the-art attempts at addressing these challenges. Section 5.6 will introduce the Glossy-W protocol, together with its key design features, namely the leveraging of concurrent transmissions and of the constructive interference. Section 5.7 will briefly describe the experimental setup constructed to validate the desirability of the two main design choices and highlight the results of these experiments, as well as the ones of benchmarking Glossy-W against an ideal Collision-Avoidance Medium Access Control. Finally, the main conclusions will be summarized in Section 5.8.

The work presented in this chapter was published in RTCSA 2017 [98].

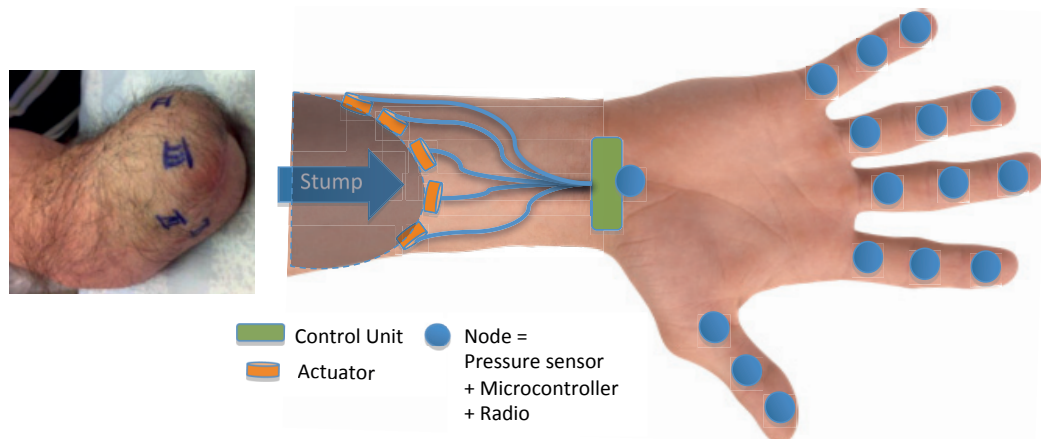


Figure 5.1 – Schema of WiseSkin. A multi-hop WSN embedded in a polymer is capable of sensing the tactile stimulus and sending it to a sink. The sink transfers the stimuli to the user by mechanical stimulation of the corresponding nerves (source of the stump image: Lund University).

5.2 Time-bounded Event Driven WSNs – the WiseSkin Project

This section elaborates on the requirements and challenges of time-bounded event-driven WSNs. It relies on the WiseSkin project as an exemplification of these challenges, otherwise broadly generalizable to WSNs requiring consistently fast responsiveness, high energy autonomy and simultaneous event detection.

5.2.1 WiseSkin - Project Description

Aimed at developing artificial skin capable of providing patients with a seamless tactile sensation, WiseSkin is a multi-disciplinary project composed of 5 research lines: materials, communication system, radio, tactile sensors and feedback to the nervous system. This endeavor is jointly pursued by multiple R&D and academic partners¹, as well as medical collaborators involved in the system's clinical testing. The WiseSkin project thus proposes the application of WSNs to the retrieval of tactile sensation in the artificial skin of human prostheses [97] (Fig. 5.1).

One of the main use cases for WiseSkin is the enhancement of commercially available prosthesis by adding tactile capabilities. This enables the user to "feel" the pressure applied to distinct regions of the skin and regulate the actuation accordingly. To exemplify the use, we can think of a patient who has undergone an arm amputation and is currently using a myoelectric prosthesis i.e., controlled through nervous impulses. Without tactile feedback, the user cannot regulate the pressure he or she applies on a given object. Therefore, a common task like grabbing an egg, can result in either dropping or smashing it.

The WSN substitutes the portion of the peripheral nervous system that conveys tactile infor-

¹ The five research lines are shared between the Swiss Center for Electronics and Microtechnology (CSEM), École Polytechnique Fédérale de Lausanne and the Bern University of Applied Sciences.

Chapter 5. Leveraging Constructive Interference for Improving the Energy-Latency Tradeoff

mation and consists of miniature motes embedded in a polymer, that sense the pressure over the skin and transmit the data through a wireless multihop network to a sink, which then transfers the stimuli to the nervous system. The network is considered multihop because the small form factor of the antenna and the absorption of the polymer surrounding the motes are expected to limit the communication range.

In the context of WiseSkin and its WSN, we consider an event as an increase in the pressure measured by a node, above a predefined threshold. The threshold should be calibrated to minimize the propagation of trivial information (e.g. the hand is resting), while it should not suppress important developments (e.g. tickling). The definition of the threshold could be a highly complex task, which is out of the scope of this thesis.

We rely on wireless communications, as opposed to a wired solution, for four main reasons: i) wires can break due to constant deformation of the skin, ii) the increase in the quantity of wires with the density of the WSN would hinder scalability, iii) off-the-shelf radios for enabling this type of WSN are easily available, and iv) the resulting wireless solution can be reused in systems where each mote has a battery or a power scavenging source.

The need to provide a seamless natural experience subjects the WSN to stringent performance constraints, presented in more detail in Sec. 5.2.2. The first iteration of WiseSkin, based on a small 5 motes WSN, has been successfully tested with human patients [99]. The next step consists of improvements in the sensing resolution by progressively increasing the network density.

5.2.2 Performance Requirements

We use WiseSkin as a case study to illustrate the challenges faced by time-bounded event-triggered systems. WiseSkin's goal of conveying the tactile stimuli in a manner that the user perceives as natural generates a series of strict requirements for the embedded WSN, notably:

Adaptability: WiseSkin's traffic is characterized by long periods of low activity, interrupted by sudden traffic increases. In the absence of tactile events, the application traffic of the network is minimal. Conversely, the stimulation of the skin triggers events from multiple motes in the same region and at a similar time. The traffic volume will further increase with the density of the WSN, as we improve sensing resolution in future iterations. Therefore, the network must cope with the different traffic schemes.

Responsiveness: the user must feel the sense of touch without a perceivable lag. The average reaction time to touch in a human is approximately 155ms [100]. Therefore, we seek a worst-case latency for the WSN of under 100ms, a target already very challenging for typical low-power WSNs.

Energy Efficiency: we aim at minimizing the power consumption of WiseSkin in order to extend the battery life of the prosthesis. Nevertheless, this goal creates a trade-off with the

responsiveness of the system, as reducing the delay requires the entire network to frequently wake up and be ready to relay an event. Whenever the skin is not stimulated, this energy investment would not be useful, but would significantly increase the power consumption.

Reliability: the WSN must minimize packet losses in order to avoid missing potentially important tactile events. The reliability of the network becomes critical in case data-fusion techniques are used to reduce the traffic. We therefore aim to ensure a reliability $>99\%$. This target can be further improved through diversity techniques, if required by an application (we elaborate on frequency diversity in chapter 6).

Mobility: the flexible nature of WiseSkin implies that the motes will change their relative position in the network during the operation, for example, when moving the fingers. Nevertheless, we do not expect that the movement will change the mean density or diameter of the network, since the motion of the motes is constrained to a volume that is approximately constant, defined by the range of movement of the human hand. A mobile network that keeps the density constant can be frequently found in domains such as body area networks [101] [102] or robot swarms [103].

In order to enable a low-power operation, we resort to the state-of-the-art solution of using duty-cycling mechanisms [48]. These schemes seek to minimize the time that a mote radio spends active by periodically listening to the channel, in case there is something to be received.

5.2.3 Two-mode Operation: Low and High-Traffic

Event-driven WSN applications, which experience sporadic events resulting in traffic surges are generally approached with two-mode operation protocols (e.g., [104], SAML [105], eLWB [106] and Crystal [54]), thus we continue with this practice along this chapter. This is mainly driven by the sharp distinction between the two naturally emerging types of behaviors in the network:

1. Low-traffic: assumed when the network is waiting for an event and there is only control traffic. The system must be ready to communicate the detection of an event within the deadline and with high reliability. As the network is expected to spend most of its lifetime in this mode, minimizing its power consumption becomes crucial.
2. High-traffic: assumed after detecting an event, therefore it must be capable of conveying a high-traffic volume, despite the consequent energy increase. Several motes (maybe even the entire network) will periodically transmit data subject to latency bounds and high reliability.

We elaborate on each of the modes and the transition from low to high-traffic mode (i.e., network wake-up) in section 5.6. The transition from high to low-traffic is left out of the scope of this chapter because it does not impact the latency.

5.3 Problem Statement

In order to save energy and cater to extreme traffic variations, a popular design for WSNs consists in operating in two modes: low-traffic and high-traffic. In our context of event triggered systems with tight latency bounds, the transition from the low-traffic to the high-traffic mode requires the mote that detected an event to quickly propagate the information to a coordinator mote, in order to trigger the wake-up. This generates two key challenges:

- Steep increases in traffic volume: Multiple motes can detect an event closely in time, thus attempting to report it simultaneously. This may result in failed transmissions due to collisions, or in an increased latency due to transmissions' rescheduling (e.g., with a contention algorithm).
- Energy vs. latency trade-off: A mote cannot immediately receive a frame when the sender is ready to transmit it. This meeting delay comes from the periodic wake-up introduced by the duty cycling, and it can be reduced if the motes wake up more often, at the expense of a higher power consumption due to idle listening (see eqs. 5.5 and 5.14 later in the chapter). Moreover, a message is also subject to the meeting delay on each hop of its path towards the sink. This creates a proportional dependency between the average and worst-case event latency and the network diameter.

Our work has thus focused on designing a low-traffic protocol that enables the propagation of event notification under stringent latency constraints, while minimizing the energy consumed, and not being impacted by simultaneous event detections.

5.4 State-of-the-Art Solutions

In this section, we will review the main state-of-the-art approaches for enabling the transition from low to high-traffic mode, while catering for latency targets and energy limitations.

We start by analyzing the routing protocol. A classical approach is to use a protocol that defines the routes a message will follow through multiple hops towards the sink in order to trigger the transition to high-traffic mode, such as RPL [13] or CTP [14]. This approach has the advantage of providing spatial reuse, by enabling different sections of the network to convey different events simultaneously. Nevertheless, we de-prioritize the spatial reuse advantage in favor of an approach relying on flooding, because of the relative mobility of the motes (e.g., the bending of the fingers in WiseSkin). The change in position of the motes can result in routes not being available upon detection of an event, thus bringing significant latency penalties. Flooding, on the other hand, is robust to mobility.

Flooding can rely on different types of distributed Medium Access Control (MAC) protocols. We differentiate between two fundamental approaches for addressing the presented challenges for the low-traffic mode: **avoiding** and **leveraging** concurrent transmissions.

The first common approach consists of designing a MAC that seeks to avoid concurrent transmissions by granting each mote with an unique and fair access to the channel. We will elaborate on this avoiding concurrent transmissions approach in Sec. 5.4.1.

The second approach consists of using concurrent transmissions. Such protocols rely on physical layer effects to receive one or several simultaneous transmissions. We elaborate on this leveraging concurrent transmissions approach in Sec. 5.4.2.

5.4.1 Avoiding Concurrent Transmissions

The classical approach in the design of MAC protocols provides exclusive channel access to one mote at a time. This can be achieved through two main techniques: synchronous or asynchronous MACs [62]. The third category of hybrid protocols, represents a combination of the previous two approaches.

Synchronous MACs rely on synchronization among the motes in order to organize transmission slots. Typically, the slots are intended to provide each mote with exclusive channel access, but it is also possible to open contention slots for operations such as bootstrapping. Some example of such protocols include [62] SMAC, TMAC and TSCH. The main advantage of this approach is that slots enable the motes to avoid delays due to contention and backoff. Nevertheless, the coordination of slots between motes requires synchronization and results in a consequent energy overhead. The use of Synchronous MACs is better tailored to periodic traffic, leveraging the energy investment required to maintain the slot structure.

Another advantage of synchronized MACs is that slots could be organized back-to-back along a route in order to minimize the delay due to contention and backoff, thus reducing the propagation latency of a message from a mote that detected an event to a defined destination. However, these approaches require the synchronization of sleep phases with respect to some source or destination. Thus, they do not provide low-latency communication from arbitrary locations in the network. Furthermore, they place constraints on the sleep cycles and phase of nodes in the network which are often modified by the application for other purposes, such as sensing [107]. The approaches we will present in this chapter do not require phase synchronization and are application independent.

In order to account for protocols where there is no per-hop contention in our benchmarking of Glossy-W, we perform a comparison with an ideal Collision Avoidance MAC that requires synchronization, detailed in Section 5.7.3.

Asynchronous MACs do not require synchronization among the motes, thus reducing the energy required in absence of traffic. They typically rely on carrier sensing operations on each mote, which enables them to periodically check if there is a packet to be received, or to confirm that the channel is free before transmitting. In the case of transmission, if the channel is found busy, the mote will try again in the future; otherwise it will send a train of frames to signal the reception to the neighbors. Examples include WiseMAC [45], ContikiMAC [46] and

Chapter 5. Leveraging Constructive Interference for Improving the Energy-Latency Tradeoff

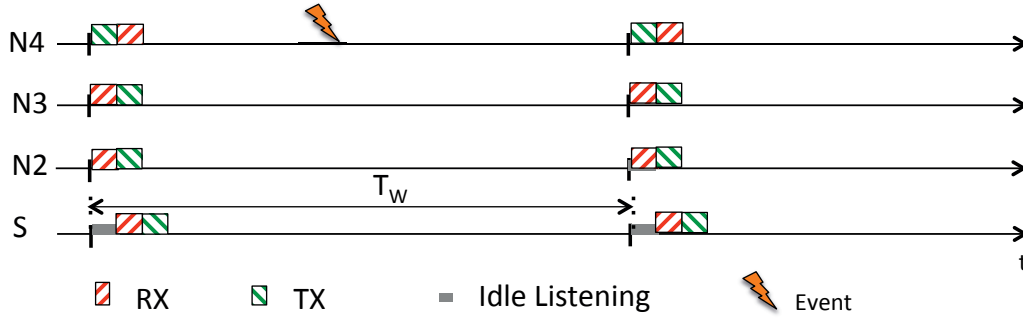


Figure 5.2 – Timing diagram of Glossy relaying a flood through a 2-hop WSN, where every mote transmits only once ($N_{gl} = 1$).

X-MAC [108].

Asynchronous MACs are better suited for sporadic traffic, but their latency has a random component due to the contention delay. This brings an unpredictable penalty to the latency, that increases with the traffic volume.

5.4.2 Leveraging Concurrent Transmissions

The family of protocols that leverage concurrent transmissions enable the motes to successfully receive one or several packets at the same time from multiple neighbors. This family can exploit two different principles for reception: capture effect and constructive interference. We elaborate both approaches in the following sections.

We highlight that we do not consider spread spectrum modulation techniques for enabling concurrent transmissions, such as the one used in WiFi. The reason is that the transceivers in WSNs have a low hardware complexity, which rarely provides this resource.

Enabling Concurrent Transmissions by using Constructive Interference

Constructive interference occurs upon reception of multiple transmissions with the same content that are temporally aligned, in a way that enables the receiver to successfully and reliably decode them. For example, the time alignment must lay within $0.5\mu s$ in the case of IEEE 802.15.4 symbols at 250 kbps [109]. This superposition increases the power upon reception (as compared to the individual components), thus improving the transmission's reliability [110].

Glossy [109] has demonstrated the feasibility of using *constructive interference* in WSNs, despite the challenging synchronization required. It is able to keep a tight synchronization through multiple hops by creating floods in the network that are relayed with a deterministic in-mote processing delay.

Figure 5.2 shows the working principle of Glossy in a 2-hop synchronized network with 4 motes, where N2 and N3 are within range of both N4 and S. The network wakes up simultaneously

with a period T_w to receive a flood from N4, which serves both to notify potential events in the payload and to keep the network synchronized. Every mote will retransmit an incoming packet after a predefined delay timed from the reception. Therefore, N2 and N3 will forward the flood simultaneously and the tight synchronization of their transmissions will make it possible for S to receive the packet, despite its coming from two concurrent transmissions.

Glossy can be configured to have every mote relaying each flood N_{gl} times. In order to simplify the presentation of our analysis, we consider along this chapter $N_{gl} = 1$ (as in Fig. 5.2), but our conclusions are independent of the value of N_{gl} . The duration of the wake-up interval is configured before the deployment of the network to cover an entire topology.

Some examples of glossy-like protocols that rely on constructive interference are Sparkle [111], RedFixHop [112] and B2B [113]).

The high reliability of glossy-like protocols stems from the following characteristics [113]:

- Short-term time diversity: each mote transmits the same frame several times during a flood. Each repetition thus represents an additional reception opportunity for each neighbor.
- Spatial diversity: the flood propagates the event through all the existing routes in the network. This maximizes the probability of a successful propagation through sections of the network connected with low quality links.
- Frequency diversity: each transmission of a mote during a flood can take place in a different channel, which is facilitated by network synchronization. This increases the probability of successful propagation in the face of frequency selective interference. This property was not exploited in the original version of Glossy, but it recently started gaining popularity (e.g., in [113] [61]).

We highlight that there are two types of synchronization involved in Glossy, which are key to the operation of the protocol and for the new mechanisms that will be proposed in the remaining of the thesis. Let us call them “Tight” and “Loose” synchronization, which are signaled in Fig. 5.2.

The first one enables all the motes in the network to wake-up at similar times to relay a flood and requires an error bound under 0.5ms. This condition is obtained by periodically sending floods from a timekeeper mote (e.g. mote N4 in Fig. 5.2) and it is readily available in glossy-like protocols.

The second one is the tight synchronization used within a glossy-like flood that enables the constructive interference. The second synchronization requires a much more challenging error bound of 0.25us (in the case of a 250 kbps transmission), which is obtained by using the first reception of a flood as the time reference to trigger subsequent transmissions within the same flood (e.g. the time alignment between packets of the same flood shown in Fig. 5.2).

Chapter 5. Leveraging Constructive Interference for Improving the Energy-Latency Tradeoff

Although promising, Glossy is known to have two limitations that must be factored in before considering its use in practical applications. The First is that the probability of successful reception of a flood (i.e. reliability of constructive interference) decreases as the number of traversed hops increases, thus limiting the diameter of a target deployment. The reason is that each relay of the flood introduces a small error of stochastic nature in the tight synchronization, which is accumulated across multiple hops ([109] presents a detailed analysis). The Second limitation is that, in the vanilla version of Glossy, all the motes of the network collaborate in propagating a flood, which has the same data content, without leveraging the spatial reuse of the spectrum (i.e. sending different packets through different routes).

Enabling Concurrent Transmissions by using Capture Effect

The few SOTA algorithms that attempt to make use of concurrent transmissions for distributed event detection leverage the capture effect (e.g., eLWB [106] and Crystal [54]). This is due to the fact that they rely on contention slots where multiple motes trigger Glossy floods concurrently upon detection of events. But since the floods from different sources will not have the tight synchronization required for constructive interference, their propagation solely relies on the capture effect.

The *capture effect* enables a receiver to successfully decode a packet in case its signal strength is at least 3 dB stronger than the sum of the received signals from all other concurrent transmissions [114]. An example of a protocol that relies on this principle is Crystal [54], which enables a mote to signal that it has detected an event and perform event identification (share its mote ID and event-related information). The mote uses a Glossy flood during a contention slot to disseminate this information through the network. Capture effect becomes necessary when multiple motes detect events closely on time, and trigger their floods during the same contention slot. Since the packet content of each flood will be different and they lack the tight synchronization (described in Sec. 5.4.2), they cannot propagate through constructive interference. Crystal operates on the assumption that at least one of the floods will eventually arrive at the sink, which will notify the sender using an end-to-end acknowledgment disseminated with a Glossy flood. On the next contention slot, the mote that received a confirmation will no longer transmit and the rest of motes that detected events will compete again.

Nevertheless, using the capture effect brings a negative impact on the reliability [115]. This is addressed in Crystal by continuing the retransmissions of the flood from a mote (up to a predefined limit) in absence of an end-to-end acknowledgement. Another approach, is proposed by the Chaos protocol, where the dissemination of the flood is inhibited in certain motes to reduce the number of simultaneous transmitters, specially in dense areas, thus mitigating the impact on the reliability.

In the next section, we further explore the impact on the reliability of using capture effect and constructive interference, and their dependency on the density of the network.

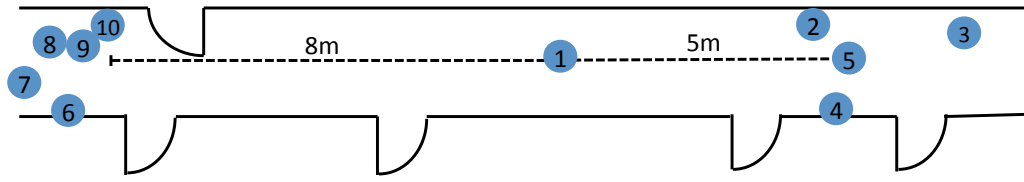


Figure 5.3 – Location of the testbed in an office hallway.

5.5 Evaluation: Capture Effect vs. Constructive Interference

We perform an experimental evaluation of the two strategies presented in Sec. 5.4 for managing concurrent transmissions: capture effect and constructive interference.

Testbed. The experiments are performed with a testbed of 10 motes located at the level of the ceiling of an office hallway (Fig. 5.3). The motes N2-N5 and N6-N10 are placed at opposing extremes of the hallway. N1 is in the center of the configuration, with a distance of 9-10m to the rest of the motes. They have line-of-sight access to each other. The motes are TelosB model (Crossbow, 2.4 GHz), operate on channel 26 (2.48 GHz) and are under little influence of 802.11b and Bluetooth networks [49]. Changes in the multipath fading pattern due to the movement of people and doors can also be expected.

The setup aims to reproduce the conditions of the EWSN dependability competition 2016 [116], an event which has become an increasingly popular opportunity for characterizing event-driven WSN protocols. We highlight that the evaluation of the protocol in this setup is also representative of the case of WiseSkin (Sec. 5.2), even though the testbed represents a version of the WSN which is scaled-up in terms of distance compared to the artificial skin application. The representativeness is driven by the fact that both systems display similar network densities and numbers of motes. Moreover, the preliminary versions of the WiseSkin protocol must be tested in large scale testbeds due to the difficulty of accessing the motes for re-flashing, once they are inside the skin. The setup does not consider the mobility of the motes, since we expect the impact to be low considering that we use flooding protocols and the network density remains constant. Moreover, the density of the testbed is designed to emulate a low-density scenario for WiseSkin: 4 neighbors on average and a network diameter of 2 hops. The reason is that this corresponds with the existing prototype of WiseSkin and, in hindsight, we expect the floods that rely on avoiding concurrent transmissions to perform better in low density scenarios, as opposed to high density scenarios where they will face increased contention.

In our setup, the 9 motes at the extremities of the hallway (N2-10) simultaneously detect an event with a 1s period and start a flood concurrently. The sink (N1) listens to concurrent transmissions. We note that the roles of the motes are different than the ones used in Sec. 5.7.2. The reason is that we require similar distances between N2-10 and N1, in order to discard the effect of different distances over the variation of the signal power upon reception. This allows us to control the effectiveness of the capture effect by changing the transmission power of the motes.

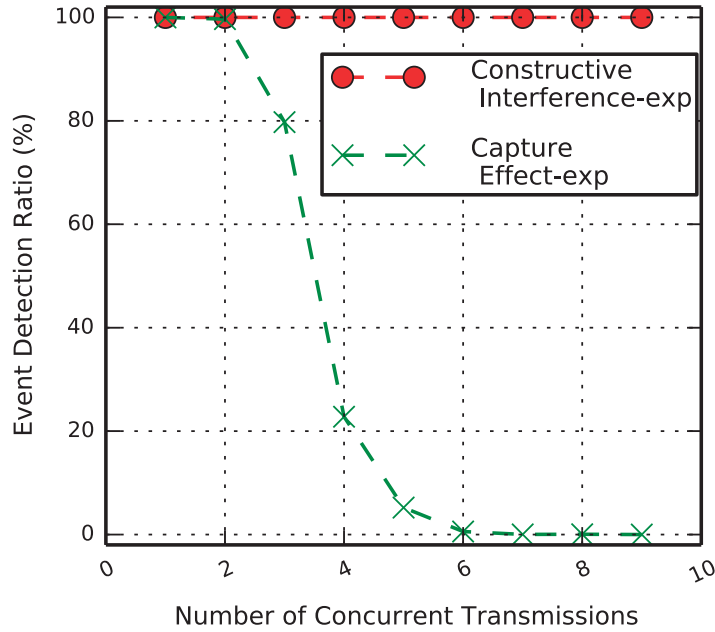


Figure 5.4 – Measurements of the Event Detection Ratio using two mechanisms for concurrent transmissions. The constructive interference is a more reliable scheme than the capture effect.

Additionally, we perform the analysis over a 1 hop network in order to focus on the core mechanisms for the reception of concurrent transmissions. The content of all the transmitted packets in both experiments and all motes is the same.

We repeat the experiment changing the number of motes transmitting simultaneously (from 1 to 9). Each repetition of the experiment lasts 1h (~3'600 events). The unused motes remain inactive.

The performance metric is the Event Detection Ratio i.e., the percentage of events successfully reported to N1. A successful report requires the reception of at least one of the floods sent by N2-10 triggered by an event.

The testing of the capture effect uses a glossy flood from N1 to trigger the event. The flood is sent with maximum power to ensure the reception by the entire network (this was previously tested with a 100% success rate). N2-10 schedules the flood simultaneously using software timers in the microcontroller with a ms accuracy. This setup ensures that the transmissions will be concurrent, but it makes the tight synchronization required for constructive interference very unlikely. Therefore, a reception can be attributed to the capture effect. This technique for experimentally studying the capture effect has previously been used in [110] and [117]. The flood from N2-10 is sent with a random power to improve the probabilities of reception due to capture effect. We note that there are 28 possible power values in a range between -25 and 0

dBm. N1 is within range even with the lowest power.

The testing of the constructive interference also starts with a glossy flood from N1, sent at maximum power. Its reception by N2-N10 triggers simultaneous floods. The transmissions are scheduled to ensure a $< 0.5 \mu\text{s}$ alignment to generate constructive interference. We use the flood from N1 to periodically synchronize the network with both schemes.

Fig. 5.4 shows that the Event Detection Ratio, when relying on the capture effect, significantly decreases with more than 2 concurrent transmissions and it is close to zero with more than 4. Conversely, the use of constructive interference yields a 100% Event Detection Ratio for the entire range of concurrent transmissions (2-9). Previous studies comparing the reliability of floods enabled by constructive interference and the capture effect support our experience [110].

These results show that event-triggered systems should use a scheme based on constructive interference, instead of a scheme based on the capture effect. This design would enable a high reliability when communicating the detection of an event to the sink, which scales with the density of the network. Extended listening is required to accommodate the flood from the sink and the potential extension in case one (or several) motes detect an event.

5.6 Proposed Protocol – Glossy-W

In the frame of this chapter, we are proposing an enhanced version of Glossy that, unlike Crystal or eLWB (as explained in the previous section), relies on constructive interference instead of the capture effect, for triggering the transition from the low to the high-traffic mode. Additionally, in Sec. 5.7.2 we demonstrate analytically and experimentally that leveraging concurrent transmissions using constructive interference delivers a better energy-latency tradeoff, namely a lower energy consumption for a given latency, than avoiding concurrent transmissions. Glossy-W is thus compatible with simultaneous event-detection, a common occurrence in high density networks such as WiseSkin.

First, it is useful to remember that Glossy relies on constructive interference to enable the reception of a packet with a high probability, despite simultaneous transmissions with identical content (as explained in Sec. 5.4.2). This effect depends on a tight temporal alignment between transmissions of neighboring motes, which is made possible by triggering each forward transmission of a flood after a deterministic processing time. This chain reaction is started by a single transmission from a mote.

In the scenario where multiple motes detect events at similar times, they will trigger contending floods. Therefore, the resulting transmissions would not have the temporal alignment required for constructive interference and would impact the reliability. Moreover, the contending floods would be propagated using the capture effect, meaning that their reliability would be degraded as the density of the network increases (see Sec. 5.4.2).

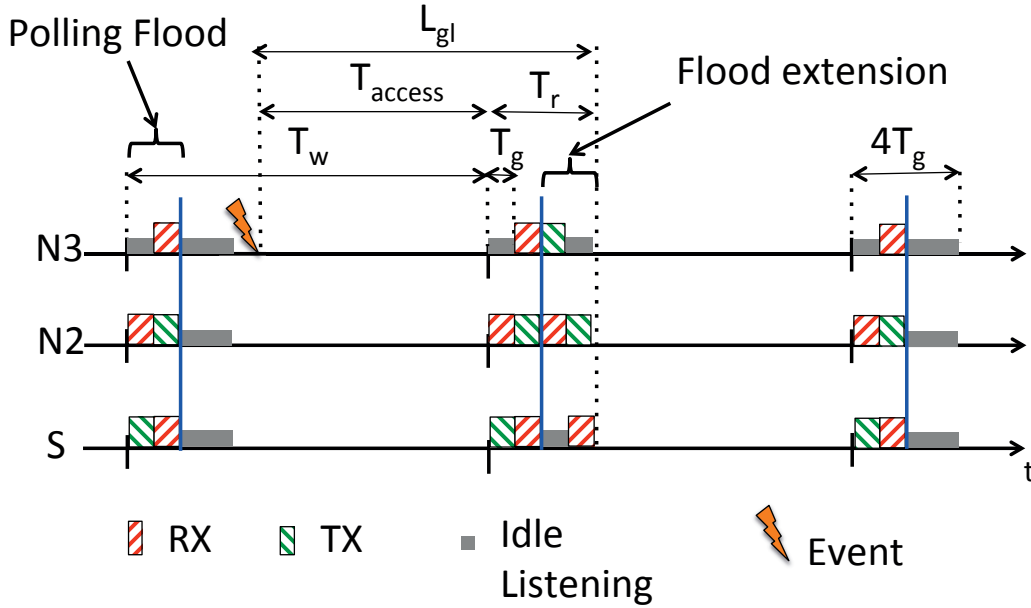


Figure 5.5 – Timing diagram of Glossy-W relaying a flood through a 2-hop WSN.

The key idea behind Glossy-W is to decouple the operations of event signaling and identification. The signaling consists on reporting to the rest of the network that an event has been detected, while the identification reports information specific to the event (e.g. the ID of the source mote or the magnitude/duration of the sensor reading that triggered it). The focus on signaling the detection of events, allows Glossy-W to use floods with identical content seeking to leverage constructive interference. Therefore, we constrain the number of events that can be reported during low-traffic mode to a single one: "something happened", which may trigger a mode change, for instance, to high-traffic. This approach defers the event identification to take place after the transition (explained in Sec. 5.6.1).

Glossy-W leverages constructive interference by periodically sending a probe flood from a host mote. After traversing the entire network, motes that detected an event can trigger an extension, which will back-propagate to the host, signaling the detection of the event. The operating principle is depicted in Fig. 5.5 for a 2-hop network, where the motes wake up every T_w for an interval $4 * T_g$ (the dimensioning of this interval will be explained later in this section). The host mote S will always initiate a flood (as Glossy would do) at the beginning of the interval. The purpose of this flood is poll the other motes for an event.

All the motes are aware of the moment when the polling flood concludes, this temporal horizon is shown as a blue line in Fig. 5.5. This is achieved with the synchronization, and the fact that all the motes know the duration of the propagation of the flood through the entire network in advance. This last factor is calibrated previous to the deployment, as done by Glossy [109].

In case no event have been detected, each mote will only relay the flood once (this parameter is decided during the configuration before deployment, it has the value of one in the example of the figure), resulting in an early termination (before the blue line depicted in Fig. 5.5).

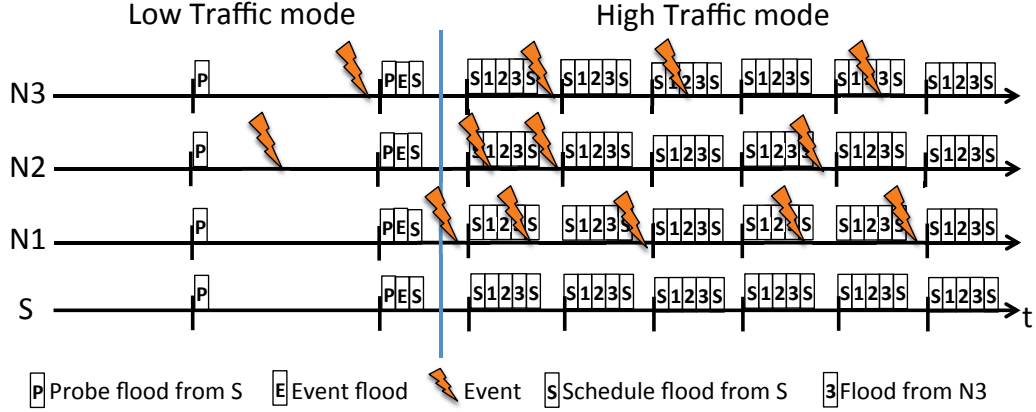


Figure 5.6 – Timing diagram of an event-triggered WSN alternating between two modes: low-traffic mode (= no tactile events) and high-traffic mode (= events are present). The idle listening intervals are not shown.

In case a mote has detected an event (e.g., N3 in the second wake-up round), it will extend the polling flood by transmitting twice the pre-configured number of times (two in the example). The result is that the flood will continue beyond the blue line and will reach back to S^2 . In case additional motes report an event at the same time, their mechanism of extending the flood is a natural continuation of the polling flood, thus it will continue the propagation with constructive interference.

The interval $4 * T_g$ in the previous example is designed to fit two glossy floods, each dimensioned to traverse the diameter of the network. A single flood requires an interval $2 * T_g$, for a two hop network and a single transmission per mote ($N_{gl} = 1$). The periodic polling flood from S also serves to keep the network synchronized.

5.6.1 Using Glossy-W in a Two-mode Operation

In this section we detail the operation of Glossy-W in a two-mode protocol designed to cater to extreme changes in the traffic, as described in Sec. 5.2.3. The intention of this section is to explain how Glossy-W fits in the low-traffic mode and the transition to high-traffic mode. We highlight that the study of the high-traffic mode is out of the scope of the chapter, since a state-of-the-art protocol tailored for periodic transitions can be plugged after a consistent transition of the network.

Low-traffic mode

Fig. 5.6 shows a 4-mote event-triggered WSN operating with Glossy-W. It is a zoomed-out version of Fig. 5.5. The system starts in low traffic mode, where the motes wake up synchronously and relay a polling flood (slot "P") started by mote S. In case a mote detects an event, it can

²A mote that detected an event must start the extension of the flood with a delay, since the beginning of the last reception, equal to T_g (even number of hops from S, e.g. N3 in Fig. 5.5) or $2 * T_g$ (odd number of hops, e.g. N2). This maintains the tight synchronization required for the flood extension, a principle introduced in [113].

Chapter 5. Leveraging Constructive Interference for Improving the Energy-Latency Tradeoff

extend the probe flood to notify the mote S (slot "E").

Upon notification of an event, the network will know that the sink will immediately answer with a new wake-up schedule to reorganize the network and transition to high-traffic mode.

High-traffic mode

Fig. 5.6 depicts the operation of high traffic mode. The goal of this mode is to reliably transfer the periodic traffic generated by each mote with a bounded worst-case event latency smaller than 100ms.

The figure shows the operation with the Low Power Wireless Bus (LWB) [118]. This protocol keeps a tight synchronization and periodically provides each mote an exclusive slot to send its data with a Glossy flood. In Fig. 5.6 these slots are represented with intervals that enclose the mote id. The sink periodically propagates a global communication schedule that determines when a mote is allowed to initiate a flood. These slots are marked with "s" in the figure.

The schedule of LWB is customizable to the requirements of the application's high-traffic. Nevertheless, the small latency requires each mote to wake up frequently and increases energy consumption. Given that the network spends a relatively small portion of its lifetime in this mode, this energy investment is not expected to significantly impact battery life.

LWB also brings a trade-off between the maximum number of motes in the network and the latency, as the period of the schedule imposed by the latency can fit a finite number of slots. This limitation can be addressed by several techniques depending on the specificities of the application, such as data aggregation. The performance evaluation of LWB under high and periodic traffic is out of scope, but can be found in [118].

After an interval without event notifications, the sink sends a new wake-up schedule to transition to the low traffic mode.

5.7 Evaluation

5.7.1 Modeling Avoiding vs. Leveraging Synchronous Transmissions

Modeling the Two Collision Approaches

For the purpose of this study, we fix a target average latency (L) and estimate the duty cycle (D) of each flooding approach.

The duty cycle of a mote is defined as the fraction of an interval T during which the radio is active (either transmitting, listening or receiving). D is a common metric to compare the power of WSN protocols, since the active state of the radio is typically the dominant factor in

the consumption.

$$D = t_{radio-on}/T \quad (5.1)$$

The latency is measured starting from the occurrence of an event and until it is reported in the top layer of the sink. Therefore, it includes the delay between the occurrence of the event and the start of the TX by the mote(s) that detected it, the rendezvous delays on each hop, and the duration of the propagation.

Additionally, we consider that each mote only relays a flood once. Table 5.1 details the parameters of our study.

We assume in our study that the network uses a single channel. Even though a multi-channel protocol is desirable because it increases the robustness against frequency-selective interference [61], this would create an unfair comparison, since protocols with a tight synchronization (e.g., Glossy-W) can easily include channel-hopping, while incurring a marginal energy overhead. Conversely, asynchronous protocols (e.g., ContikiMAC) would require significant design changes and an increase in the energy consumption to include it.

We also do not consider simultaneous event detection. We expect this scenario to favor the avoiding-collisions approach, since additional motes detecting transmissions concurrently will generate congestion and collisions, thus hampering the latency of the protocol. This scenario does not make a difference for the leveraging synchronous approach, since Glossy-W is agnostic to the number of motes detecting an event simultaneously. We highlight that we are concerned with notifying the network that an event has been detected, in order to enable the transition to the high-performance mode. Therefore, we do not consider the propagation of any additional event related information, such as the ID of the source, since this can be handled by the high-performance mode.

Avoiding Concurrent Transmissions

The objective of this section is to derive an analytical expression of the mean Duty Cycle (D_{cm}) as a function of the mean Latency (L_{cm}), for flooding based on ContikiMAC, triggered by an event (Fig. 5.7). We consider the flooding scheme netflood provided in Contiki (version Apr-18-2016), which is a best effort approach to send a single packet to all nodes in the network. Netflood operates by forwarding a broadcast transmission, a single time, by each mote. The header of the packets include the destination and packet id, which enables the motes to limit duplicate transmissions, by not forwarding packets with an already-seen packet id for a given destination. The header also includes a time-to-live counter, which is decreased on each forward and serves to limit forwarding loops, by not forwarding packets with a time-to-live equal to zero.

In ContikiMAC, each broadcast consists of multiple repetitions of the same frame, barely

Chapter 5. Leveraging Constructive Interference for Improving the Energy-Latency Tradeoff

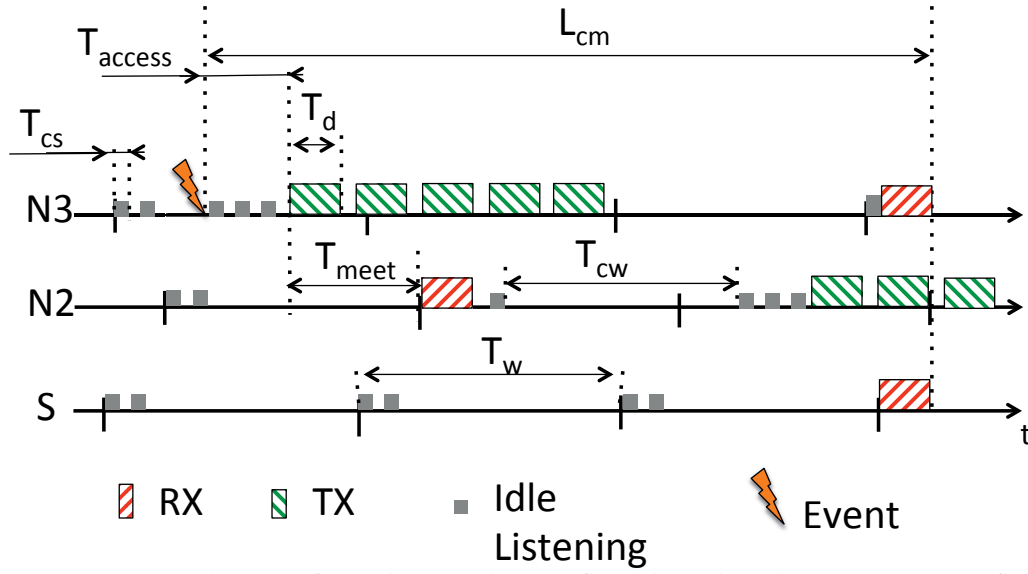


Figure 5.7 – Timing diagram of ContikiMAC relaying a flood through a 2-hop WSN using Netflood.

exceeding the duration of a sampling interval T_w , which ensures that all the neighbors within range will have an opportunity to receive the packet. This design is typical of low-power preamble sampling MAC protocols, such as WiseMAC [59]. Upon reception of a packet to forward, a mote will perform two channel samplings to check that the channel is available, without waiting for its next wake-up (defined by the T_w period in fig. 5.7). If the channel is found free, the mote will begin the transmission. Otherwise, the mote will wait a random delay proportional to T_w , using an exponential back-off before trying again.

In order to simplify our modeling, we do not consider collisions, assuming a low initial traffic.

Other authors have developed models in agreement with ours, but adapted to their respective systems [119].

Latency

We can model the mean flood latency in the absence of collisions by adding the following components (Eq. (5.2)):

$$L_{cm} = h * T_{fix} + h * T_{meet} + T_{cw} * (h - 1) \quad (5.2)$$

$h * T_{fix}$ accounts for the fixed delays of the flood over h hops before reaching the sink. T_{fix} considers the delay in processing the TX request, two channel samplings to determine if the channel is clear, and the duration of the reception.

$h * T_{meet}$ is the delay between a mote receiving a frame and the next hop waking up to sample the channel, over h hops. The wake-up can happen anytime during T_w with uniform probability, thus the average delay introduced by this offset is $T_{meet} = T_w/2$.

$T_{cw} * (h - 1)$ represents the delays due to contention over h hops, assuming that the first broadcast finds the channel free. The contention forces each mote to wait for a random interval before attempting to transmit, in case it detects that the channel is busy. The default contention mechanism³ reschedules the transmission randomly within the window $T_{cw} \in [T_w, T_w + (2^{n_{tx}} * T_w)[$, where n_{tx} is the number of retransmission attempts. A mote that intends to forward a packet often finds the channel busy because the transmission from the previous hop has not yet finished. We build on this observation to propose an optimistic scenario: a mote only finds the channel busy because of the ongoing precedent TX, thus each intermediate hop only finds the channel busy once ($n_{tx} = 0$). This scenario is reasonable given that the minimum T_{cw} is T_w , which is enough to ensure that the precedent TX will be over by the next transmission attempt. Moreover, we recall that we are not considering simultaneous event detection on this analysis, thus eliminating the possibility of contention due the transmission of another mote that detected the event. This simplifies the expression for the contention window $T_{cw} \in [T_w, 2 * T_w[$ and provides the average value for $T_{cw} = 3 * T_w / 2$.

Typically, models of the preamble sampling MAC protocols consider the guard time (T_{guard}) to compensate for clock drift. This term is important for modeling unicast traffic when the motes learn the wake-up offset of their neighbors (also known as phase-lock optimization). Nevertheless, this term is not considered in the analysis, as it has very little impact on the broadcast transmissions on which the floods rely.

Taking into consideration the previous observations, Eq. (5.2) can be reduced to Eq. (5.3).

$$\begin{aligned} L_{cm} &\approx h * T_{fix} + h * T_w / 2 + (h - 1) * 3 * T_w / 2 \\ L_{cm} &\approx h * T_{fix} + (4h - 3) * T_w / 2 \end{aligned} \quad (5.3)$$

Power

The duty-cycle of the floods in the avoiding collisions approach (Eq. (5.4)) is composed of three contributions: channel sampling (D_{cs}), flood transmission (D_{evtx}) and flood reception (D_{evrx}).

$$D_{cm} = D_{cs} + D_{evrx} + D_{evtx} \quad (5.4)$$

D_{cs} is determined by the proportion of time spent to sample the channel twice with respect to the sampling period.

$$D_{cs} = 2 * T_{cs} / T_w \quad (5.5)$$

D_{evtx} is the proportion between the time that each mote requires to transmit the flood and

³The default contention algorithm in ContikiMAC is imposed by the CSMA module of Contiki's Rime communication stack [120]

Chapter 5. Leveraging Constructive Interference for Improving the Energy-Latency Tradeoff

the mean event-arrival period (T_{ev}). In asynchronous duty-cycled MAC protocols, the flood is relayed by doing a broadcast transmission, that consists of repeatedly transmitting the frame during an entire sampling period (T_w), which ensures that all the motes within range will wake up to receive a repetition of the frame. In ContikiMAC, there is a pause between the transmission of successive frame repetitions when the radio is turned off. To consider the time that the radio is off during the transmissions, we define the parameter $0 < \alpha < 1$. α is the proportion of the time that the radio is in active (listening, transmitting or receiving) state with respect to the time between the end of the transmission of two consecutive frames.

$$D_{evtx} = \alpha * T_w / T_{ev} \quad (5.6)$$

D_{evrx} is the mean proportion of the time that each mote requires to receive a flood and the mean event-arrival period (T_{ev}). Each mote detects a repetition of the frame during channel sampling, and remains with the radio in listening state to receive the next one (frame duration T_d). Thus, on average the mote will keep the radio on for an interval $0.5 * T_d$ in order to detect the transmission, and an additional T_d for actually receiving the frame. This pattern is repeated for each event (every T_{ev}). We consider the mote that initiates the flood to eventually receive a packet when a neighbor forwards it. D_{evrx} is also proportional to the average number of neighbors (n) that receive the transmission.

$$D_{evrx} = 1.5 * n * T_d / T_{ev} \quad (5.7)$$

D_{evrx} can be expected to be significantly smaller than D_{evtx} (Eq. 5.6), given that $T_w \gg T_d$ on duty cycled MACs. Moreover, as T_w assumes smaller values, D_{cs} (Eq. 5.5) becomes significantly larger than D_{evrx} (which remains constant). Therefore, we can drop the term D_{evrx} from Eq. 5.4 without a significant loss in accuracy, since it will only have a small contribution for every value of T_w . We highlight that the model presented in this section corresponds to the low density WSN used as the current prototype of WiseSkin and represented by our testbed, which means that n has low values (e.g., under 5), and this will not make the term D_{evrx} become significant with respect to other components of the duty cycle. Under the same argument, we ignore the listening during the gap time between repetitions (T_g).

From the previous observations, we express D_{cm} as a function of T_w :

$$D_{cm} = 2 * T_{cs} / T_w + \alpha * T_w / T_{ev} \quad (5.8)$$

We use Eqs. (5.8) and (5.3) to express D_{cm} as a function of L_{cm} :

$$\begin{aligned} D_{cm} &\approx T_{cs}(4h-3)/((L_{cm}-h*T_{fix}) \\ &+ 2 * \alpha * (L_{cm}-h*T_{fix})/((4h-3)*T_{ev}) \end{aligned} \quad (5.9)$$

Table 5.1 – Protocol Parameters

Parameter	Value	Meaning
n	10	Number of motes
T_{ev}	5-10s (rand)	Event period
T_d	1.54ms	Frame duration ContikiMAC
T_{fix}	12.5ms	Fixed delays ContikiMAC
T_g	0.5ms	Frame duration Glossy-W
T_{cs}	0.436ms	Channel sampling
T_w	8-1000ms	Channel sampling period
α	1.54/2.7	sleep-between TX (Sec.5.6)
T_{sync}	1s	Synchronization period
T	30s	Duty cycle period (Eq.5.1)
h	2	Number of hops

Leveraging Synchronous Transmissions

The objective of this section is to derive an analytical expression of the mean Duty Cycle (D_{gl}) as a function of the mean Latency (L_{gl}) for Glossy-W relaying a flood triggered by an event (Fig. 5.5). In Glossy-W, a flood is formed by forwarding a broadcast TX immediately after reception by each mote. The forward TX is performed by tightly synchronized concurrent transmitters, thus enabling constructive interference upon reception. The network wakes up simultaneously with a period (T_w) and remains active for an interval long enough to transmit the flood. This interval is proportional to the diameter of the network.

We consider a that Glossy-W is based on the code from Glossy (version Nov-14-2012), with the parameters unchanged unless otherwise noted.

Latency

The average latency for transmitting a flood with Glossy-W can be expressed according to Eq. (5.10).

$$L_{gl} = T_{access} + T_r \quad (5.10)$$

T_{access} considers the waiting time between the detection of an event by a mote and the next wake-up of the network to initiate the flood. The wake-up can occur anytime during the period T_w with equal probability, thus the mean access time is $T_w/2$.

T_r considers the probing and event relaying delay across h hops. It accounts for the reception, processing and propagation delays. Glossy-W has packets with minimum protocol overhead and we only consider a payload of 1 byte (Table 5.1), thus reducing the duration of the reception to < 1 ms. The processing and the propagation (at the speed of light) also imply durations $<< 1$ ms. Therefore, we neglect the contribution T_r .

Chapter 5. Leveraging Constructive Interference for Improving the Energy-Latency Tradeoff

The previous observations simplify L_{gl} to Eq. 5.11.

$$L_{gl} \approx T_w/2 \quad (5.11)$$

Power

The average per-mote duty cycle is expressed in Eq. (5.12). The equation comprises the contributions of two periodic floods: i) synchronization packets from the sink (D_{sync}) and ii) probe packets from the sink to enable event reports from other nodes (D_{probe}).

$$D_{gs} = D_{sync} + D_{probe} \quad (5.12)$$

D_{sync} considers the impact of flooding a synchronization packet every T_{sync} on the duty cycle. This requires the network to keep the radio-on during an interval $T_g * h$ (Eq. (5.13)).

$$D_{sync} = T_g * h / T_{sync} \quad (5.13)$$

D_{probe} accounts for the duty cycle required for probing. The entire network will wake up periodically to receive a probing flood from the sink (which requires a radio-on time $T_g * h$). In case an event needs to be relayed, one or several motes will use the flood from sink to trigger another flood that will eventually reach the sink (which requires an additional radio-on time $T_g * h$). The motes keep the radio on during an interval $2 * T_g * h$ every T_w , independently of the presence of an event (Eq. (5.14)).

$$D_{probe} = 2 * T_g * h / T_w \quad (5.14)$$

Considering the previous observations, we express D_{gs} as a function of T_w :

$$D_{gs} \approx T_g * h / T_{sync} + 2 * T_g * h / T_w \quad (5.15)$$

We use Eqs. (5.11) and (5.15) to express D_{gs} as a function of L_{gs} :

$$D_{gs} \approx T_g * h / T_{sync} + T_g * h / L_{gl} \quad (5.16)$$

The probing floods (Eq. 5.14) can be used to synchronize the network, thus eliminating the contribution of Eq. 5.13. Nevertheless, we keep the two mechanisms decoupled because certain synchronization algorithms require packets longer than the probing ones.

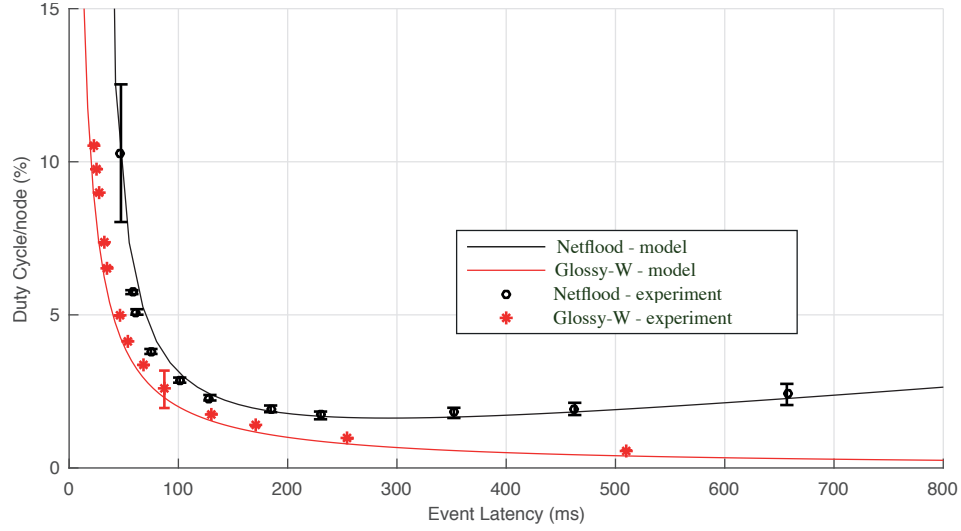


Figure 5.8 – Duty Cycle (D) as a function of the Latency (L) for a 2-hop network. Glossy-W displays a lower duty-cycle in the entire range of latency values.

5.7.2 Experiments on Avoiding vs. Leveraging Synchronous Transmissions

The objective of this set of experiments is the comparison of the energy consumption required for a target latency of two approaches: avoiding and leveraging concurrent transmissions (presented in Sec. 5.6.1). We also validate the calculations presented in Sec. 5.7.1.

Testbed. We use the testbed already described in Sec. 5.5. We measure the duty cycle using the energest module of Contiki, which measures the time each mote spends on each state of the radio with a μs accuracy.

Protocols. The scheme to avoid concurrent transmissions uses ContikiMAC to relay the flood, while the approach for leveraging concurrent transmissions uses Glossy-W. The parameters used for the modeling in Sec. 5.7.1 are kept consistent for these experiments.

ContikiMAC belongs to Contiki version Apr-18-2016 and Glossy-W is based on Glossy's version Nov-14-2012. Both protocols are used with the default parameters, unless otherwise noted. The floods in ContikiMAC are generated with Rime's netflood module. We note that the packets in ContikiMAC are larger than in Glossy-W because of protocol overhead (e.g., sender id).

Metrics. We monitor the duty cycle (D), Event Detection Ratio and latency. The former is measured since the occurrence of the event until it is reported in the top layer of N2's stack.

Parameters. We obtain the values of the Table 5.1 directly using the Cooja emulator [121] and the settings in the source code (the values do not depend on the emulation scenario). These parameters enable us to plot the model predictions as solid lines that overlap with the measured values in Fig. 5.8. Cooja emulates the behavior of the motes flashed with a given binary file with cycle accuracy. The error bars in the figures denote the standard deviation.

The value T_g is optimized for the diameter of the testbed as per the indications in [109]. We

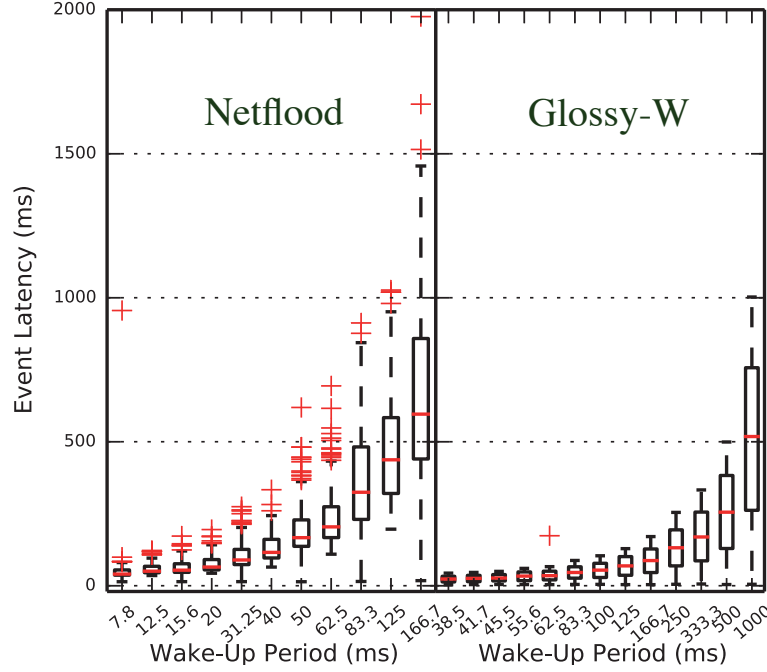


Figure 5.9 – Measured Latency (L) as a function of the Wake-Up period (T_w) in a 2-hop network. The lack of outliers for Glossy-W shows that the worst-case latency is highly predictable and determined by the wake-up period. The red band is the median. The bottom and top of the box represent the first and third quartiles. The whiskers denote the maximum and minimum values, excluding outliers (red crosses).

test multiple values and choose the minimum that does not impact the reliability of the floods.

In our setup, mote N6 periodically detects an event and starts a flood to notify N2 (the sink). The event arrives at a random interval between 5-10s. We set the transmission power level of all the motes to -25 dBm, which ensures that the floods will traverse 2-hops. This matches the measured network diameter during the experiments.

We repeat the experiment changing the wake-up period (T_w) of the protocol, thus operating on different trade-offs between energy consumption and latency. Each repetition of the experiment lasts 1h (~480 events).

We compare the duty-cycle required to achieve a given latency (Fig. 5.8, Eqs. 5.9 and 5.16). This shows that Glossy-W is more energy efficient (i.e., has a lower duty-cycle) than ContikiMAC with Netflood for the entire range of latencies, by at least 30%. We explain this behavior by noting that the motes using ContikiMAC must contend for accessing the channel on each hop of the flood, thus the latency is proportional to the network's diameter ($T_{cw} * (h - 1)$ in Eq. 5.2) and higher than Glossy-W's, for a given wake-up period (Eqs. 5.3 and 5.11). Therefore, ContikiMAC compensates with a lower T_w (i.e., higher energy consumption) for striking a given latency.

Given that ContikiMAC's delay due to contention grows with the network diameter, we expect this result to hold for larger topologies.

Moreover, the distribution of latency values as a function of the Wake-Up period (T_w) (Fig. 5.9) shows that the worst-case latency of Glossy-W is highly predictable, compared to ContikiMAC with Netflood. For Glossy-W, the ticks associated with the maximum and the minimum (excluding outliers) closely match the range of expected values ($L \in [0, T_w]$), and the median is close to the average of the distribution ($T_w/2$). This evidences a uniform distribution with clear bounds determined by the T_w , as predicted by the model (Eq. 5.11). Conversely, the distribution of values for ContikiMAC presents many outliers, which stems from the contention window producing an unbounded latency. The outliers are defined as the points over $1.5 * \text{third quartile}$ or under $1.5 * \text{first quartile}$.

We define the reliability as the percentage of events detected at the sink from all the occurred events. This metric is over 99% for both protocols in all cases.

In general, the experimental results match closely the model predictions for both protocols.

5.7.3 Comparing Glossy-W with Floods over an Ideal Collision-Avoidance MAC

The comparison presented here may seem unfair as the floods using ContikiMAC are handicapped with respect to Glossy-W, due to a significant per hop delay bigger than $T_w/2$. The delay is introduced because flooding in ContikiMAC requires each mote to forward a packet by transmitting during an entire wake-up period (T_w). This means that the neighboring mote(s) will have to wait until the end of a transmission to relay the packet. To further check the superiority of the Glossy-W approach over a collision avoidance approach, let us compare it with the same flooding scheme on top of an ideal collision avoidance MAC. Such a MAC adopts a “magic” schedule that avoids collisions. Let us look at the makespan of such a schedule. In the case of a linear network, the makespan would be the depth of the network (h). In case of a binary tree, the makespan is at least (assuming that nodes at 2 hops do not interfere) $(2h + 1) * T_s$. This is due to the fact that at each hop, 2 nodes receive a broadcast from the previous hop and have to relay it one after the other. The transmission time is thus at minimum $2 * T_s$ per hop. This corresponds to the latency of Glossy-W in the worst case. Other types of tree topologies would yield longer latency.

With respect to energy, the cost of a flood is proportional to the radio-on time per hop. In the case of the ideal MAC, the minimum energy would be proportional to $2 * T_s$, which becomes bigger for trees with a higher degree. In the case of Glossy-W, the radio-on time per hop $2 * T_s$ is independent of the degree of the tree.

This shows that Glossy-W is certainly superior to a collision avoidance approach, with the exception of the special case of line topologies.

5.8 Conclusion

We have studied the problem of distributed wake-up of an event-triggered WSN with tight energy and latency constraints. The wake-up is a key component of state-of-the-art protocols that rely on multiple modes of operation tailored to specific traffic characteristics (e.g., low vs. high).

In this context, we have offered two key contributions. Our experiments have shown that leveraging simultaneous transmissions instead of avoiding them can actually result in better energy performance. The key differentiator stems from the fact that avoiding concurrent transmissions requires contention for access to the channel on every hop, thus necessitating a higher duty cycle in order to strike a latency. Moreover, leveraging concurrent transmissions also provides a highly predictable worst-case latency. The result has been shown analytically and validated experimentally in the frame of a real testbed measurement campaign, leading to at least a 30% reduction in the energy consumption for a given latency. We have also performed a comparison with an ideal collision avoidance MAC, which confirms the superiority of our approach. This scheme can be extended to a multitude of other applications in the CPS domain.

The second contribution consisted in showing that, among protocols leveraging simultaneous transmissions, the ones relying on constructive interference are preferable to the ones based on the capture effect, given their superior robustness to simultaneous event detection, a key constraint of high density WSNs. This happens because the reliability of the capture effect does not scale with the density of the network. The result has been validated in a real testbed experiment leading to 100% (constructive interference) vs. 0% (capture effect) reliability for 6 or more neighboring nodes.

The findings have been incorporated in the development of Glossy-W, a modified version of Glossy that is tailored to achieve the best available latency-energy efficiency trade-off, while catering for multi-node simultaneous event detection. Glossy-W is the first protocol able to leverage floods that rely on constructive interference in a simultaneous event detection scenario. It will be embedded in the WiseSkin pilot project to be tested by human patients. While inspired by WiseSkin, Glossy-W has a broad range of applicability, especially as an increasing number of WSN applications require a dense network, consistently fast responsiveness and high energy autonomy.

6 Improving the Energy-Latency Trade-off in Event-driven Protocols based on Synchronous Transmissions

6.1 Introduction

Protocols relying on floods based on synchronous transmissions, hereafter referred to as *glossy-like* protocols [109] have been systematically displaying the best performance in WSN event-triggered systems (details are provided in Sec. 6.3 and an example is shown in Fig. 6.1). These protocols have consistently won the top places at the latest series of the EWSN dependability competition, an independent benchmark of state-of-the-art protocols for event-triggered systems [61] [122].

However, the fact that glossy-like protocols require a node to use the same flooding primitive for both the network wake-up and the event notification (propagating the event information upon detection of an event) results in high levels of energy consumption. A prompt notification requires the entire network to frequently probe the channel in order to receive the flood in case an event has been detected. As the flooding is tailored for high-reliability packet transmission, using it for probing the channel consumes more energy than necessary, this operation becoming the dominant driver of energy consumption in the case of sparse events. Moreover, excessive flooding leads to co-existence issues with neighboring networks, a potential obstacle for practical event-triggered applications.

This study proposes Synchronized Channel Sampling (SCS), a polling mechanism that has been tailored to trigger a reliable and energy-efficient network wake-up. SCS is capable of decoupling the wake-up from the event notification, thus reducing the power required for periodic polling, while reliably waking the network up upon detection of an event. SCS is compatible with state-of-the-art protocols that rely on synchronous transmissions, its benefits being thus extendable to Crystal [54], Glossy [109], eLWB [123], Chaos [115] and Back-to-Back Robust Flooding (B2B) [113], among others. The testbed experiments performed show that SCS manages to reduce the energy consumption of B2B (the winner protocol of the 2017 dependability competition) by at least 33.3% and up to 40%, while maintaining equivalent reliability.

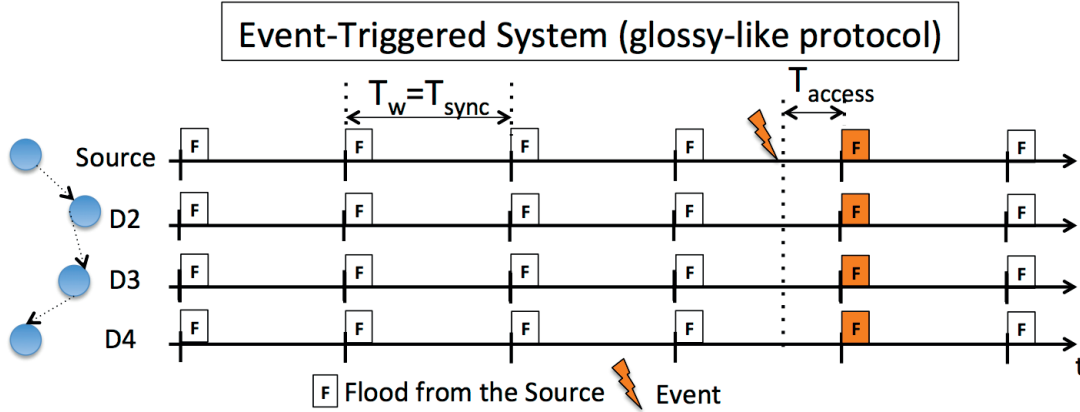


Figure 6.1 – Schema of a glossy-like protocol operating in an event-triggered application (for the case of a single event detector, i.e. the source). The source triggers periodic floods that keep the system synchronized and report events. An example of the structure of a flood (F slot) is shown in Fig. 6.2 for the case of the B2B protocol.

In the following pages, we will start by describing the state-of-the-art protocols that rely on simultaneous transmissions, as well as their key performance levers and limitations (Sec. 6.2). Section 6.3 will discuss other related work. We will then proceed to introducing the Synchronized Channel Sampling (SCS) mechanism, describing its benefits, its design and functionality (Section 6.4). Section 6.5 will present the experimental tests performed in order to validate SCS' performance, as well as the quantitative results and sensitivity analysis. Section 6.6 presents potential further optimizations, before drawing the main conclusions in Section 6.7.

The work presented in this chapter was published in WFCs 2018 [124].

6.2 Problem Statement

Flooding in glossy-like protocols constitutes a highly reliable primitive for spreading information throughout the network, with minimum propagation delays (as explained in Chapter 5), a highly desirable property for event-triggered applications.

Flooding's performance in terms of both energy consumption and network coexistence can however be significantly improved by exploiting the fact that event-triggered systems do not always have events to report. In many applications, the events are extremely rare, but once they appear, they require a short end-to-end latency. In the event-triggered application from Fig. 6.1, the source could need to initiate several flooding cycles per second in order to rapidly communicate an event that might appear with large inter-arrival times. Therefore, replacing polling in absence of events with a more energy-efficient mechanism can significantly reduce the WSN's power consumption, this translating into a longer battery life and lower maintenance costs in practical deployments.

Moreover, as the application requires a shorter end-to-end latency, the frequency of flooding

must increase proportionally in order to reduce the access delay (T_{access} , also known as dead-time) shown in Fig. 6.1. This results in high radio activity, contributing to the saturation of the radio spectrum, and hampering the operation of neighboring networks. This issue becomes more acute when several glossy-like networks, eventually serving different applications from multiple owners, share the same space.

This co-existence issue limits the use of glossy-like networks in practical applications. A mechanism that reduces the transmissions when no events are present would thus improve the general adoption of this promising protocol family.

6.3 State-of-the-art

6.3.1 Back-to-Back Robust Flooding (B2B)

Back-to-Back Robust Flooding (B2B) is a novel Glossy-like protocol based on periodic floods enabled by synchronous transmissions [113]. Its core strength relies on a mechanism for providing the tight synchronization required for constructive interference (as explained in Sec. 6.3), which increases the reliability of the floods. Traditional Glossy-like protocols rely on the simultaneous reception of a packet by multiple motes in order to trigger the synchronous sending of the subsequent repetition in a flood. Hence precious reception energy is used for timing the transmissions. B2B's main advantage is the fact that it transmits packets based on timeouts triggered from the first reception, thus eliminating the necessity for additional receptions in order to generate the remaining packets. This means that the time the radio would have used for receiving packets only aimed at aligning the transmissions can actually be used for successive transmissions, thus reducing latency and energy consumption.

The main criteria for considering B2B in this chapter is its diverse set of mechanisms aimed at ensuring a high reliability (including frequency diversity), as well as its proven performance superiority, as validated by independent evaluations [122]. Nevertheless, SCS finds general applicability beyond B2B in the entire family of Glossy-like protocols (details about the requirements in Sec. 6.4.1).

6.3.2 Other Related Work

The use of channel sampling (CS) for reducing the idle power consumption in WSNs is a technique frequently employed in asynchronous low power listening medium access control protocols, such as B-MAC [58] and WiseMAC [59]. These protocols use CS to evaluate the availability of the medium before sending packages, as well as to detect any transmission to be received. SCS, the protocol we will introduce in the frame of this chapter also uses CS to detect a transmission; however it adapts the technique to effectively replace flooding in protocols based on synchronous transmissions.

A protocol that relies on CS must properly calibrate the RSSI threshold used to determine the

Chapter 6. Improving the Energy-Latency Tradeoff in Event-driven Protocols based on Synchronous Transmissions

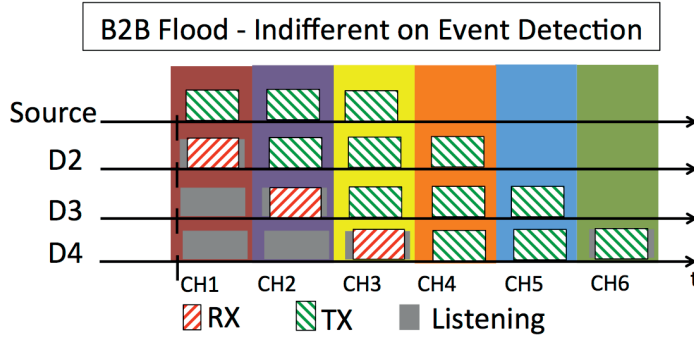


Figure 6.2 – Schema of a flood in B2B. The floods are being sent periodically, even when no events have been detected. The round presented in this figure is the content of an F slot in Figure 6.1.

channel availability, given its significant impact on performance. SCS can leverage previous research in CS, such as [125] and [126], in terms of using strategies aimed at a more accurate adaptation of the threshold value to the wireless environment of the network.

Before B2B, several protocols have approached event-triggered systems by using simultaneous transmissions, among them: Glossy [109], Sparkle [111] and RedFixHop [112]. These examples rely on periodic floods in order to propagate the events in the network, regardless of whether events are actually present. Therefore, the points of improvement presented in Sec. 6.2 are equally valid, and SCS can also increase their energy efficiency. Furthermore, the EWSN 2018 competition showed that the family of Glossy-like protocols keeps growing, as the majority of participants relied on this technology, thus further stressing the importance of widely compatible mechanisms like SCS.

A recent approach for adapting a protocol based on synchronous transmissions to event-triggered applications is Event-based Low Power Wireless Bus (eLWB) [106]. This study focuses on Low Power Wireless Bus (LWB) [127], which is tailored to periodic traffic, and proposes a wake-up mechanism that manages to reduce the power consumption in case no traffic is present, thus, adapting LWB to be used in event-driven systems. SCS shares the spirit of proposing a more energy efficient solution for waking the network up. Nevertheless, the core mechanisms of both proposals differ significantly, as eLWB reduces the flooding period in order to detect a wake-up signal, while SCS seeks to avoid them.

The Crystal protocol [54] is another approach to minimizing the transmissions in event-triggered networks during long periods without events, while exploiting synchronous transmissions to react quickly and reliably upon detection of a trigger. Despite the notable traffic reductions achieved by Crystal, we chose B2B as a benchmark, as its higher energy efficiency provided us with a more stringent reference. Nevertheless, SCS can also be used to improve the energy consumption of Crystal by providing a network polling mechanism that is more energy-efficient than Crystal's flooding.

Figure 6.2 displays the schema of a B2B flood in a 3-hop network. In the first slot, all the nodes start listening for a short time before the source initiates a transmission. D2 is the only node within range of the source, thus becoming the only one that receives the packet. In the

next slot, all the motes will change the channel and the source and D2 will retransmit the packet. In the following slots, the motes will repeat the previous actions until each of them has transmitted the flood a predefined number of times (three in this example). This number must be calibrated proportionally to the diameter of the network in order to ensure a successful propagation to all the motes.

In the publicly available version of B2B, the size of the packets is 8 bytes and the data rate is 250 kbps.

6.4 Proposed Protocol - Synchronous Channel Sampling (SCS)

The key idea behind our proposal for addressing the challenges described in Sec. 6.2 is to decouple the event signaling from the event identification operations.

Before elaborating further, we present a small reminder of the concept of wake-up and its relation with events signaling and identification. Low Power WSNs typically maintain a dormant state with the radio transceiver deactivated, enabling significant reductions on its energy consumption. In order to engage in sending or receiving packets, each mote wakes-up (i.e. turns to an active state the transceiver) periodically for a limited time interval. In the glossy-networks under consideration, a timekeeper mote maintains the network synchronized by triggering a flood every few seconds. This coarse synchronization enables wake-up of all the motes to occur almost at the same time, in case they have to trigger or relay a flood. In applications where events are scarce, the wake-up interval is of special importance because in most wake-ups, nothing will happen as no event will have taken place (thus increasing the idle listening), but it will represent the driver of energy consumption in the network. This justifies our focus on minimizing its duration.

In state-of-the-art glossy-like protocols, motes trigger a flood on the next wake-up instant upon detection of an event. Using the protocol Crystal as an example, this flood contains the mote ID and event-specific information, enabling the mote to report to the rest of the network that an event has been detected (signaling) and also the ID and event information (identification). In case other motes also detected events closely in time, they will aim to trigger a flood during the same wake-up instant, thus creating concurrent floods with different packet contents. Even though the contending floods cannot be relayed using constructive interference because the packets are not identical, they propagate via capture effect. This enables to signal to the other nodes that an event has been detected, but to identify only one event. Subsequent contending rounds are necessary to identify the remaining events in contention.

By decoupling event signaling from event identification into separate operations, we focus on minimizing the energy consumption required to signal to the other motes that an event was detected. If the event identification is also desired, the event signaling can trigger a transition of the network to an operating mode where each mote has an exclusive slot to propagate event

Chapter 6. Improving the Energy-Latency Tradeoff in Event-driven Protocols based on Synchronous Transmissions

information (e.g. the high traffic mode presented previously in Sec. 5.2.3). In this chapter we assume that handling the allocation of exclusive slots is performed using a state-of-the-art protocol, but optimizing this stage is not within the scope of this thesis.

The duration of the transmission to signal a detection is shorter than in the case of event identification, since it does not include event information, hence enabling a reduction in the wake up duration. Considering that the target applications use aperiodic traffic, we are thus addressing the dominant power consumption factor.

The core component of our proposal is the Synchronized Channel Sampling (SCS) scheme, presented in Section 6.4.1, a scheme that is performed synchronously in multiple channels of the network, reliably waking up the system in case of an event (Sec. 6.4.2). SCS rounds can replace the systematic probing using floods in protocols that rely on synchronous transmissions, thus providing a more energy efficient alternative. SCS has been accepted for the EWSN 2018 dependability competition [128].

After SCS wakes the network up, it can transition to a high-capacity operating mode to convey event-related data. For example, it can switch to a slotted protocol such as eLWB or Crystal. SCS can thus greatly improve the performance of any WSN application catering to traffic that is not periodic, while requiring a bounded latency, by providing an energy efficient and reliable wake-up mechanism that decreases the power consumption in absence of events.

6.4.1 Synchronous Channel Sampling

We have designed Synchronized Channel Sampling (SCS) as a wake-up mechanism. SCS listens during short intervals and only wakes the network up in case an event is detected. Therefore, we manage to minimize the idle listening required for SCS, while ensuring a high event-detection reliability.

Figure 6.3 displays the timing diagram of an SCS round in a 3-hop synchronized WSN performing two basic actions: **checking** the network, in the absence of any events (top) and **waking up** the network upon detection of an event (bottom, propagating an event detected notification). We assume that the network is synchronized with an error smaller than $500 \mu\text{s}$, which is easily achievable in protocols that rely on synchronous transmissions (e.g., [54]). The protocol thus performs the following steps:

Checking. SCS checks the network by using several Channel Samplings (CS). A CS consists of measuring the Received Signal Strength Indication (RSSI) in the channel, in order to determine if there is a transmission to be received. Periodic CS enables the motes to keep the radio inactive most of the time. It is typically used in asynchronous low-power listening medium access control protocols, as a way of minimizing the energy consumption (e.g., B-MAC [58] and WiseMAC [59]).

Fig. 6.3 (top) displays each mote performing three synchronized CS in channel 1 with period

6.4. Proposed Protocol - Synchronous Channel Sampling (SCS)

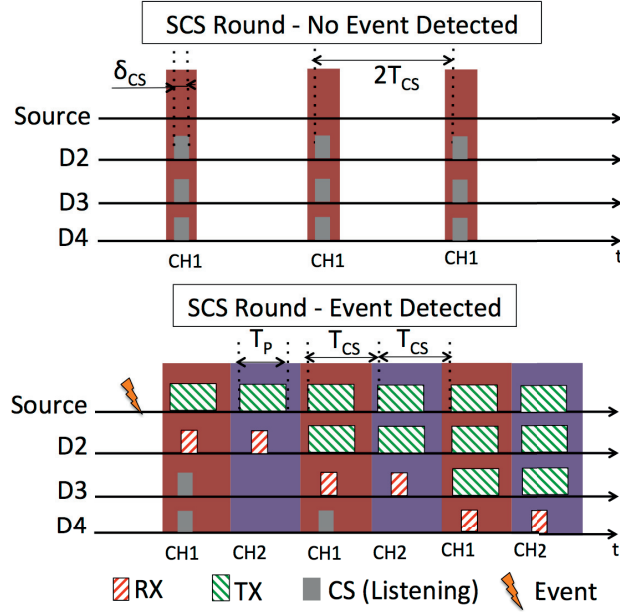


Figure 6.3 – Timing diagram of SCS in a 3-hop network in the absence of events to be notified (top), and waking up the network upon detection of an event (bottom).

$2 * T_{CS}$, where $T_{CS} = 1\text{ms}$. Each channel sampling has a duration $\delta_{CS} = 280\mu\text{s}$, which is 36.4% of the radio-on time for sending a packet in B2B - vanilla (767 μs at 250kbps and 4 Bytes payload).

Waking Up. Upon detection of an event, the source transmits a tone to wake the neighbors up. The synchronization enables starting of the transmission shortly before CS in the network, in order to ensure a reliable detection. The motes that wake up will relay the tone in order to propagate the wake-up indication. Since CS relies on RSSI measurements, the payload of the tone is irrelevant. Moreover, the tones do not require a tight temporal alignment for a correct detection, as opposed to primitives that depend on constructive interference or capture effects (refer to Sec. 6.3).

Fig. 6.3 (bottom) is a timeline of the motes relaying a wake-up indication from the source through the network. Upon detection of the tone during CS in channel 1, a mote repeats CS in channel 2 to minimize the false wake-ups due to interference. The total number of channel verifications per hop is defined as N_{ch} . For illustration purposes we use two verifications ($N_{ch} = 2$), but additional ones can be used.

The actions for checking and waking up are repeated three times in this example. This number corresponds to the diameter of the network in number of hops, and it is the minimum required to propagate the wake-up signal to all the motes.

The active time of the radio for an SCS check is smaller than the one for a flood (compare Fig. 6.3 - Top with 6.2). Therefore, it is reasonable to expect that SCS is a more energy efficient mechanism.

Eq. 6.1 expresses the duration of an SCS round (T_{SCS}), which depends linearly on the diameter

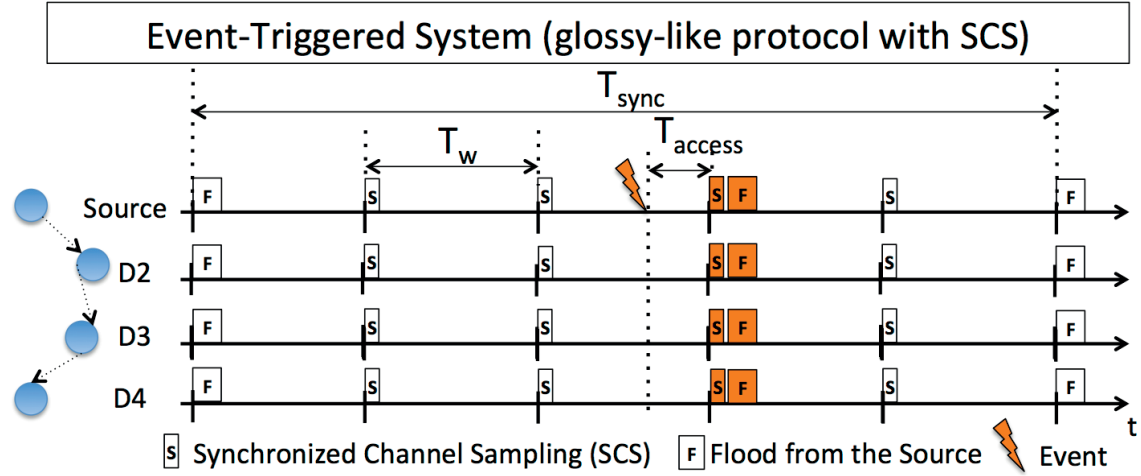


Figure 6.4 – Schema of SCS enhancing the protocol from Fig. 6.1, for operating in an event-triggered application. The enhanced protocol leverages the fact that events are not always detected, thus using SCS rounds instead of probing floods (compare with Fig. 6.1), which are more energy efficient. An example of the structure of a flood (F slot) is shown in Fig. 6.2 for the case of B2B protocol, while the structure of SCS (S slot) is shown in Fig. 6.3.

of the network (h), the number of verifications (N_{ch}), and the spacing between consecutive verifications ($T_{cs} = 1ms$). In the example in Fig. 6.3 (bottom), the duration of the wake-up is $T_{SCS} = 1ms * 3 * 2 = 6ms$. The impact of the latency compared to a B2B round is discussed later in Section 6.5.2. We have selected the value $T_{cs} = 1ms$ because this is the minimum granularity of the timers provided in Contiki OS, thus it could be further reduced by developing sub-ms timer methods, although the synchronization error of the network will define a lower bound.

$$T_{SCS} = T_{cs} * h * N_{ch} \quad (6.1)$$

We highlight that there are two types of synchronization involved in SCS. The first one enables all the motes in the network to wake-up at similar times to relay a flood. This synchronization requires an error bound under 0.5ms, it is obtained by periodically sending floods from a timekeeper mote and it is readily available in glossy-like protocols (e.g. the time alignment shown in Figs. 6.1, 6.4 and 6.2). The second one is the tight synchronization used within a glossy-like flood that enables the constructive interference. The second synchronization requires a much more challenging error bound of 0.25us, which is obtained by using the first reception of a flood as the time reference to trigger subsequent transmissions within the same flood (e.g. the time alignment shown in Fig. 6.3). In the case of SCS, we rely on the first type of synchronization, which has more relaxed requirements.

6.4.2 Improving Glossy-like Protocols

The main idea behind SCS is to use this mechanism for polling rather than flooding when there is no additional information to propagate. Figure 6.4 shows the working principle of SCS improving a glossy-like protocol, and can be compared to Fig. 6.1. The following points

describe the operation of SCS within the target protocol:

1. The source periodically initiates flooding for synchronization (displayed as the slot F, detailed in Fig. 6.2). Note that the period (T_{sync}) is significantly higher than the flooding period (T_w). The synchronization floods are also used to convey events detected during the last T_w .
2. All floods, except for the ones required for synchronization, are substituted by SCS rounds (displayed as the slot S, detailed in Fig. 6.3).
3. In case the source detects an event, it will send a wake-up transmission that will be propagated throughout the network. This is followed by an additional flood from the source in order to propagate the event. Note that the wake-up transmission alone could be used to signal the event, but it would not be able to relay event specific information.
4. To further increase reliability, the flood that propagates the event can be repeated N_{retx} times. The flood repetitions take places instead of the following SCS, or by using the synchronization flood. These repetitions are optional, thus they are not shown in Fig. 6.4.

After a successful wake-up, the network is in a consistent state and can thus transition to a different state-of-the-art protocol, in order to convey the traffic required by the specific application. For example, the transition can lead to sending a single B2B flood or structuring a slotted high-performance protocol such as eLWB. This strategy enables the network to transition between operation modes and cater for different traffic patterns requiring a bounded latency.

6.4.3 Wake-up from Multiple Sources

In previous sections we have considered the case of a single source network, where only one mote can detect events. In this section we introduce the distributed wake-up as the capacity of a network to wake the network up triggered by multiple motes, potentially at the same time. This scenario is common in dense networks where multiple motes can detect the same event.

SCS seamlessly supports applications that require distributed wake-up, which derives from the use of CS to detect the wake-up tone, a reliable mechanism that does not necessitate a precise time alignment between simultaneous transmissions. Furthermore, the robustness of the mechanism is expected to improve with the number of simultaneous transmissions, given the increasing levels of energy in the channel.

Figure 6.5 (bottom) displays a 3-hop WSN in an event-driven system that requires any node in the network to be able to detect events. Mote S1 assumes the task of keeping the network synchronized by periodic flooding. The motes will perform regular SCS rounds (Figure 6.5

Chapter 6. Improving the Energy-Latency Tradeoff in Event-driven Protocols based on Synchronous Transmissions

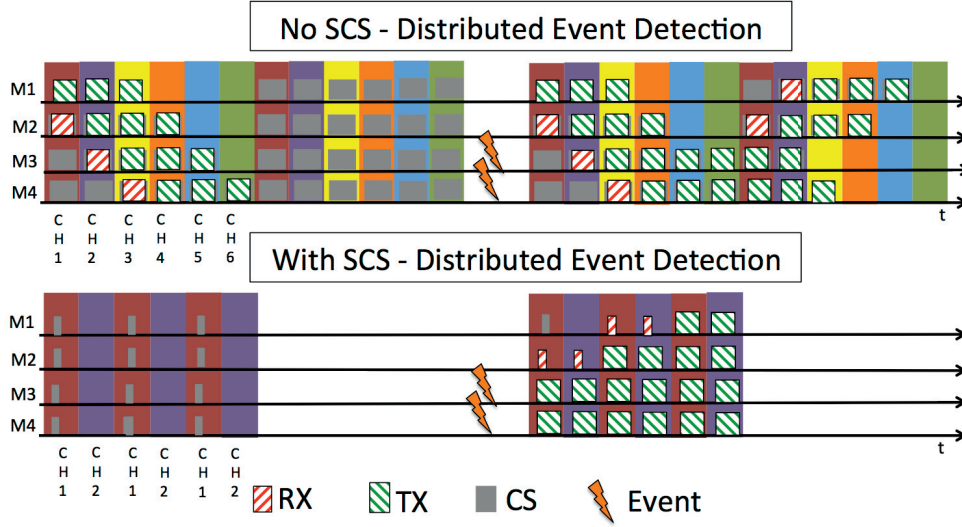


Figure 6.5 – Schema of an event-triggered system with distributed event detection without SCS (top) and with SCS (bottom). The former requires significant modifications to support this application, which increase its power consumption. SCS has built-in support and its wake-up power consumption is independent of the number of simultaneous events.

(bottom - left)). In case several motes detect an event, they can wake the network up by following the procedure described in Sec. 6.4.1 and depicted in Fig. 6.5 (bottom - right). The figure displays motes S3 and S4 detecting two events simultaneously and waking the network up during the same SCS round.

The simplicity of this wake-up mechanism is in stark contrast with the requirements for enabling distributed event detection in protocols that rely solely on the capture effect, or on constructive interference, for enabling simultaneous transmissions (contrast top vs. bottom in Fig. 6.5). As detailed in the previous chapter, Glossy-like protocols typically have two options for triggering a network wake-up from multiple motes:

- Capture effect: multiple motes can trigger concurrent floods and rely on the capture effect for its reception (explained in Sec. 5.4.2). Nevertheless, the reliability of the capture effect scales poorly with the density of the network [109].
- Glossy flood extension: the synchronizing mote can send regular polling floods in order to enable other motes to signal an event, by extending the flood. Moreover, this mechanism necessitates the motes to keep the radio active for the time needed to convey two floods, even when no events are present. This is the approach of Glossy-W explained in the previous chapter in sec. 5.6.

Neither the wake-up with Glossy flood extensions, nor SCS, will enable the motes that detected events to identify themselves. Upon wake-up, the network must transition to an operation mode to enable the sources that detected the event to identify themselves and reserve a slot to disseminate event specific information, or request bandwidth for data stream dissemination.



Figure 6.6 – Schematic of Flocklab. The Source (green) detects the events and reports them to the destinations (gray). The red motes are jammers (source of the floor plan image: ETH-Zurich).

For example, after the SCS wake-up, the network can transition to a round of Chaos [115], which is an efficient all-to-all data sharing primitive for protocols that rely on synchronous transmissions. Chaos would enable the motes that detected events to request a slot. Later on, the network can transition to a round of the protocol Low power Wireless Bus (LWB) [127], which would allocate slots to the motes to propagate event information using glossy-like floods.

6.5 Experimental Results

The quantitative experiments conducted have been aimed at evaluating the performance improvements obtained by B2B from incorporating SCS (Sec. 6.5.2) and studying the evolution of the performance for different event inter-arrival times (Sec. 6.5.3).

Metrics. We consider three key performance metrics: (i) *power* - the mean power consumed by a mote during the experiment; (ii) *latency* - the time measured between the occurrence of an event and the moment it is reported in the top layer of the destination mote, thus taking into account the delay until the initial mote can start the transmission (i.e., end-to-end latency); and (iii) *reliability* - the percentage of events that are correctly received in the destination mote. We compute the latency and reliability based on the serial output of the nodes, and measure the power from sampling the current draw of the mote at a frequency of 28 kHz [51] using Flocklab's power profiler.

Testbed. All of the experiments have been performed in the Flocklab testbed [51], using 27 Tmote nodes deployed in a university building (Fig. 6.6). The motes are subject to interference from Wi-Fi and Bluetooth networks, as well as from the movement of people inside the building, particularly during office hours.

Interference. The experiments in this section have been performed under minimum interference conditions. This means that they have been performed without jammer motes and at night (19:00 - 6:00) or during weekends, when the influence of Bluetooth and WiFi networks is

Chapter 6. Improving the Energy-Latency Tradeoff in Event-driven Protocols based on Synchronous Transmissions

minimal. A sensitivity analysis under the effects of interference can be found in Sec. 6.5.1.

Bootstrapping. B2B's procedure for bootstrapping the network remains identical whether SCS is used or not. The initial 20s of every test are considered for bootstrapping, thus they are not taken into consideration for the mean power calculation.

Traffic pattern. We consider a scenario with a single-source and multiple-destinations, emulating the EWSN competition during the editions 2016 and 2017. Without SCS, B2B leverages the periodic ($T_{sync} = T_w = 200\text{ms}$, Fig. 6.1) probing floods from the source for synchronizing the network. When using SCS, we choose the synchronization period $T_{sync} = 1\text{s}$ (Fig. 6.4), which can be as high as 20s, while still meeting SCS' clock skew requirements (Sec. 6.4.1) (e.g., [54]). We selected this configuration in order to analyze a case where SCS is in disadvantage in terms of energy consumption.

The setup used is similar to the one employed in the dependability competition at EWSN 2016 and 2017 [122]. Therefore, mote D1 (the source) periodically detects an event and starts a flood in order to notify all the motes in the network. The transmission power level of all the motes is set at 0 dBm and the CS threshold is -77 dBm (the derivation of this value results from Fig. 6.14 and it is discussed in Sec. 6.5.3).

Unless otherwise noted, the event is generated at a random interval of [2s, 4s] ($\overline{T_e} = 3\text{s}$), $T_w = 200\text{ms}$ was used in all protocols, the duration of each experiment is 5min.

The default configuration of B2B for testing in Flocklab, as provided by its authors, assumes a network diameter of 6 hops. We continue to use this assumption throughout all experiments.

6.5.1 Impact of the External Interference and the Mitigation via Verifications in Multiple Channels

This sensitivity study aims to characterize the impact on SCS' power consumption due to interference and the number of SCS verifications.

It has been extensively documented that the channel sampling primitive is susceptible to false wakeups caused by radiofrequency activity external to the network detected as activity in the channel [129] [130]. The origin of the external transmissions can be due to co-existing networks, malicious jammers, microwave ovens, WiFi or Bluetooth devices, among others.

Due to its reliance on the channel sampling primitive, SCS inherits a susceptibility to false wake-ups. The direct consequence of a false wake-up is an increase in the energy consumption due to the extended active time that the radios spend sending the wake-up probe, the additional SCS verifications triggered in the neighbors and subsequent listening time for propagating a Glossy-like flood. Moreover, a false wake-up in any mote will result in a full neighbor wake-up, thus underlining the relevance of a mechanism that is robust against interference. An additional effect of false wake-ups is the increase in the spectrum occupation, due to the

propagation of the wake-up tones. In the design presented in Sec. 6.4.2, a false wake-up has no negative effects on the reliability and redundancy, the reason being that the network wake-up is contained in a slot that is isolated from the future ones.

In order to generate a false wake-up, the frequency of the interference source(s) must match all the channels used in SCS, and have an intensity strong enough to surpass the CS threshold. Therefore, the effect of interference strongly depends on the spatial and temporal distribution of the sources, and on their intensity.

We aim to ground our evaluation on standard tools and experimental conditions used by the WSN community. We have created a challenging environment for SCS by selecting the interference pattern included in JamLab that maximizes the occupancy of the channel: emulating the download of a heavy file transfer in Ubuntu (pattern J1WIFI3 in [72]). We continue using the Flocklab testbed, as in the previous section, with seven jammer motes (displayed in red in Fig. 6.6). The jammer motes operated during the entire experiment and only in a given channel (J10 in ch26, J19 in ch25, J4 in ch26, J32 in ch25, J28 in ch15, J13 in ch15, J3 in ch25), which matches the channels used during SCS, thus increasing the probabilities of matching the channel sequence in SCS and further challenging the robustness against false wake-ups. The interfering motes are calibrated to keep the emission power constant during the experiment. All the experiments have been performed during the night and weekends in order to minimize the influence of 802.11b and Bluetooth networks, as well as the effects of the movement of people inside the building, with the aim of improving the reproducibility of the interference pattern. The setup emulates similar experiments that study the effects of interference over WSN protocols, such as [131], [54] [132].

Figs. 6.7 and 6.8 show the mean power consumption and the percentage of SCS rounds (in the entire network) that result in a false wake-up, respectively. The labels in the figures specify the number of jammers and their transmission power, while the x-axis shows the number of CS verifications (N_{ch}). Each bar results from an individual 5 min experiment and 23'060 data points (SCS rounds). Note that the base of Fig. 6.7 is 1 mW, instead of zero.

Fig. 6.7 shows B2B's mean power consumption as a reference (2.5 mW, matching the values previously presented in Fig. 6.12). The adjacent column displays the worst-case power consumption of B2B when using SCS ($N_{ch} = 2$), obtained by forcing a false wake-up in every mote and SCS round, which exceeds the value of the B2B reference by 96%. This represents an upper bound on the effects of jammer motes, independently of their characteristics. We were not able to perform worst-case experiments with other number of channels due to time limitations in the testbed, but we chose $N_{ch} = 2$, which can serve as reference for the results in other sections of this chapter.

The "No jammers" displays the lowest power consumption in the series, as a consequence of a false wake-up rate under 5.6% for all the N_{ch} values. Even though this percentage is relatively low, it shows that there is still a minor effect of external interference sources over the network, which could be attributed to control traffic in WiFi routers, or to the operation of

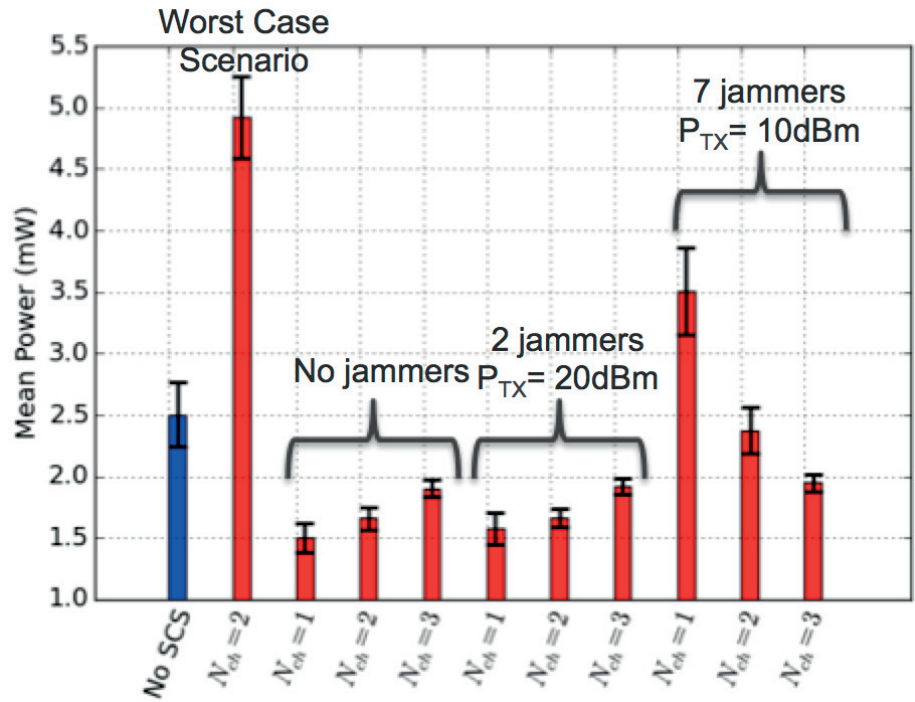


Figure 6.7 – SCS’ sensitivity to interference from jammers - power, as a function of the number of CS verifications (N_{ch}).

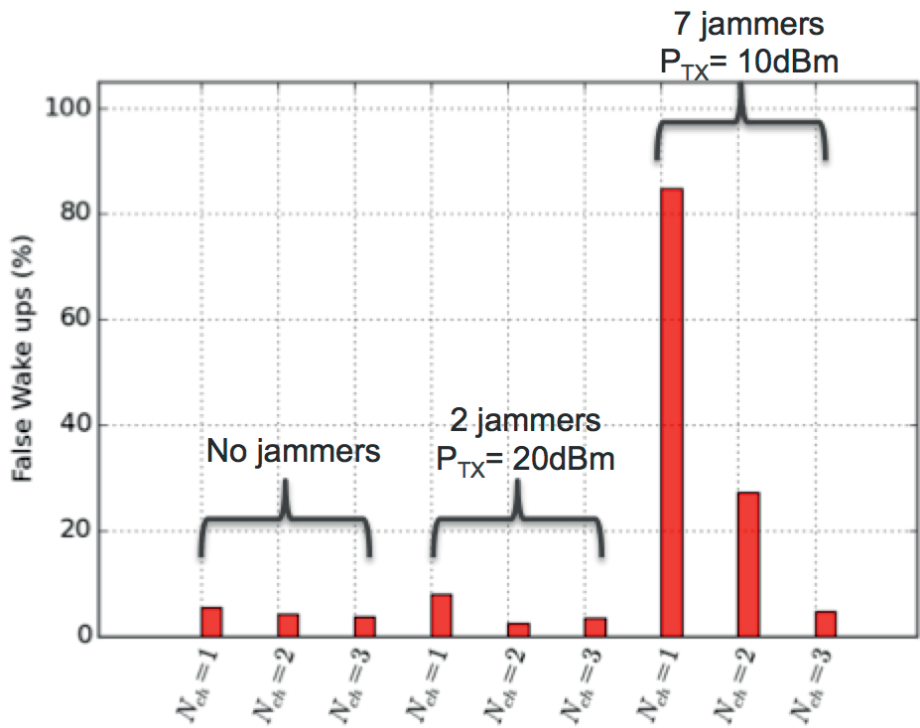


Figure 6.8 – SCS’ sensitivity to interference from jammers - false wake-ups, as a function of the number of CS verifications (N_{ch}).

HVAC engines.

The introduction of two jammers (J13 and J33 in the testbed shown in Fig. 6.6) with a high transmission power (20 dBm) in opposite extremes of the topology does not cause an increase in the number of false wake-ups and, consequently in the energy consumption of the network, when compared to the no jammers case. The reason behind this is that the jammers are too far apart, and their emission power is too low to exert their interference simultaneously and over the threshold on any of the network nodes.

Incrementing the number of jammers to seven increases the power consumption for $N_{ch} = 1$ by 126%, with respect to the case with 2 jammers, even if the TX power is smaller (10 vs. 20 dBm). In the case of $N_{ch} = 2$, the higher number of jammers result in an increase by 38.2%, which is significantly lower, but still notable. Fig. 6.8 shows that the power increases can be traced to higher proportions of false wake-ups, which is explained by noting that a more uniform distribution of the jammers increases the probabilities of a node suffering a simultaneous influence in the channels used for SCS. These results suggest a strong dependency of SCS on the spatial distribution of the interference sources, which is as important as their intensity. This behavior is a consequence of SCS' verification system that requires both an above-threshold RSSI and activity in several channels.

Further increasing the number of channel verifications to $N_{ch} = 3$, results, for both interference patterns, in the same energy consumption (1.95 mW) as in the no-jammers case, despite the fact that the jammers are using the three channels checked by SCS. The rationale behind this relies on the fact that the probability of finding several channels occupied during the same short time interval decreases with the number of CS verifications. The effective mitigation of the interference validates SCS' approach of using multiple channels for verification. However, the increase in the number of verifications will result in a longer SCS, which can increase the latency of the event propagation. We elaborate on the impact over the latency in Sec. 6.5.2.

SCS' power consumption dependency on the jamming characteristics means that under a variable interference pattern, SCS will automatically adjust its energy consumption. Therefore, an intense activity of jammers during a defined time interval might temporarily increase the energy consumption, which should go back to the original levels as soon as the interference passes.

This experiment suggests that $N_{ch} = 2$ strikes a good trade-off between effectively mitigating false wake-ups in a typical office environment, and the power consumption in the absence of interference. Therefore, we use this value for the detailed performance analysis of SCS in the remaining of the chapter.

6.5.2 Dependability Benchmark

The objective of this experiment is evaluating the improvement in the performance of B2B when using SCS. SCS is set to two successful CS verifications in order to validate a wake-up

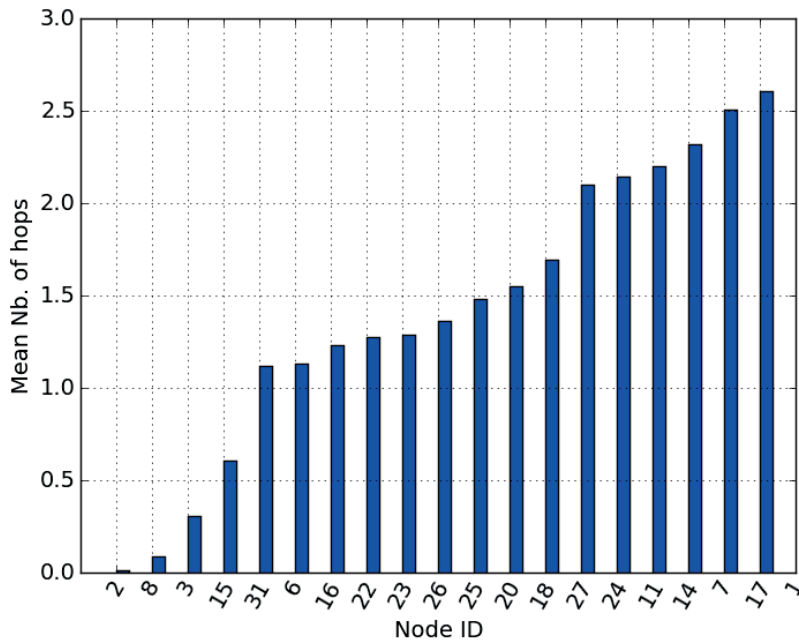


Figure 6.9 – Evaluation of the mean number of hops per node in B2B.

signal.

The mean number of hops (Fig. 6.9, without SCS) traversed by the event floods remains equivalent with and without SCS, with values between 1 (neighbors of the source) and 3.5 (motes further away such as D7 and D17). This suggests that the diameter of the network assumed for the experiments (i.e., six hops) is sufficient to cover the entire network with a flood. It also includes a safety margin, in case the routes available at the instant of the flood propagation are longer than usual.

In this setup, D7 detects the events, which are generated at an interval of [2s, 4s] (randomly and uniformly distributed). The performance of the network is monitored in order to convey the notification to the source (S1), which requires multiple hops (at least 4). In this experiment, we do not include additional interference over the one already present in the testbed, e.g., from WiFi, Bluetooth, and so forth.

The results yield that both protocols display equivalent reliability and 100% of the events were correctly received in all the motes (corresponding graph is not shown). The results are based on a population of 300 events and remain consistent throughout the three repetitions of the experiment.

The mean power consumption of each mote during the experiment (Fig. 6.10) is at least 33.3% lower when using SCS. This metric is measured by an external module on each mote.

The distribution of event latencies per mote (Fig. 6.11) shows that the mean value through the network without SCS is 124.1 ms, which corresponds to the contribution of the access latency

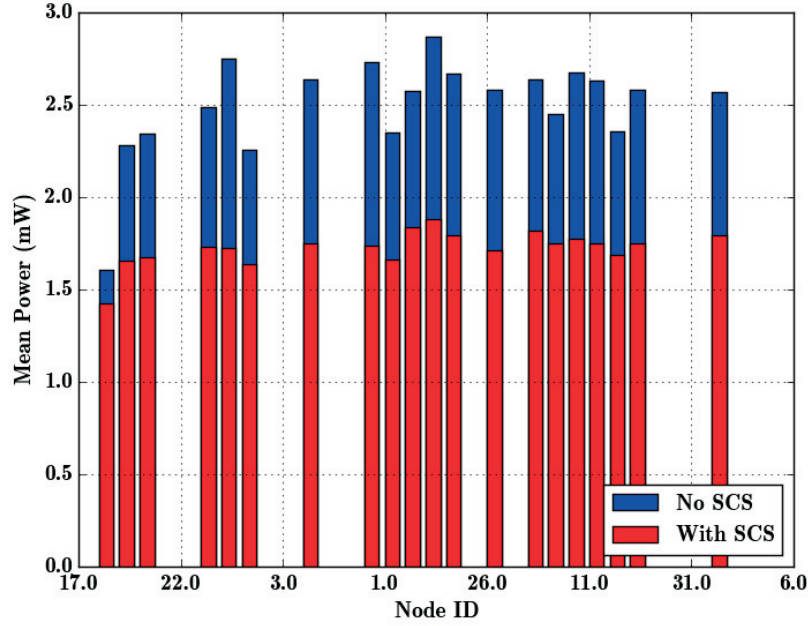


Figure 6.10 – Improvement of the mean power per node brought by SCS on B2B. The missing columns correspond to the jammer nodes or non-existing nodes.

($T_W/2 = 100ms$ with $T_W = 200ms$) plus the duration of the B2B flood to propagate through the network (mean value 10.3 ms, measured from the testbed gpio traces). The remaining contribution to the latency budget ($124.1ms - 10.3ms - 100ms = 13.7ms$) belongs to guard times used in B2B that could be optimized.

Fig. 6.11 also shows that the mean value through the network with SCS (136.3 ms, right plot) is slightly higher than without SCS (124.1 ms, left plot). The difference of 12.2 ms is due to the latency overhead of SCS, which is provided by Eq. 6.1: $T_{SCS} = 1ms * N_{ch} * hops = 1ms * 2 * 6 = 12ms$, with our settings. Therefore, SCS represents a mean impact in the latency of 9.7% ($= 100 * T_{SCS} / (T_W / 2) = 100 * 12ms / 124.1ms$). This impact becomes less relevant as the sampling period (T_w) increases, i.e., for larger target latencies.

This experiment has been repeated three times with and without SCS, with consistent results.

6.5.3 Characterization of SCS

The objective of this experiment is to study the evolution of the performance improvements brought by SCS over B2B for different event inter-arrival times and channel sampling thresholds.

Figure 6.12 shows the evolution of the mean power consumption among the nodes in the network with (red curve) and without (blue curve) SCS, as the mean event inter-arrival time (\bar{T}_e) increases (i.e., the events become more scarce). Figure 6.13 is an augmented view of the

Chapter 6. Improving the Energy-Latency Tradeoff in Event-driven Protocols based on Synchronous Transmissions

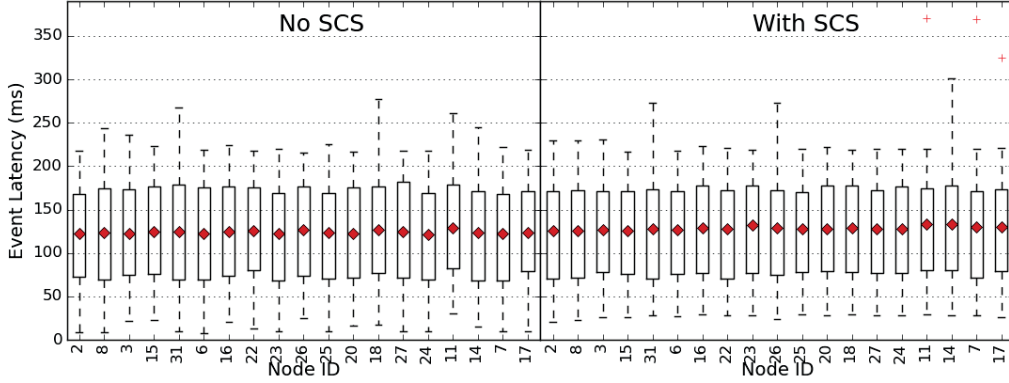


Figure 6.11 – Event Latency per node for B2B (left) and SCS (right). The mean latencies are the diamonds. The bottom and top of the box represent the first and third quartiles. The whiskers denote the maximum and minimum values, excluding outliers (red crosses).

Figure 6.12 for low \bar{T}_e values. \bar{T}_e is the mean value of the event inter-arrival time T_e , which is selected from the interval $[\bar{T}_e \pm \bar{T}_e/3]$ with a random uniform distribution.

Each point belongs to a separate experiment, whose duration was adjusted from 5min to 20min, as required to receive a statistically significant number of events: 100 - 200 events for $\bar{T}_e > 1$ s and 600 - 1500 events for $\bar{T}_e < 1$ s. The points are joined by straight lines that represent a linear interpolation, and the error bars represent the standard deviation. We perform the analysis for $T_w=200$ ms. This experiment is performed with the background interference present in the testbed and no influence from the jammers.

The plot shows that the power consumption of B2B does not change significantly with \bar{T}_e , and small variations of the mean are within the standard deviation. This behavior is expected because B2B has the source sending periodic floods regardless of the detection of an event.

Conversely, the power consumption of B2B, when using SCS, decreases as events become more scarce. The decrease in the curve is strong for small event intervals and becomes asymptotic towards larger ones. This feature arises from the fact that the source does not send any floods when events are not present, beyond the required ones for keeping the synchronization. The change of behavior in the curve is marked by an elbow at $\bar{T}_e = 3$ s, which is one order of magnitude superior to T_w . The elbow marks the value of \bar{T}_e after which the energy savings due to more scarce events become less significant.

The crossing point between both curves is $\bar{T}_e = 500$ ms and marks the \bar{T}_e value above which using SCS is a more energy efficient alternative. Nevertheless, in the operation point $\bar{T}_e = 750$ ms, the events arrive almost in every T_w , thus the system finds itself at the limit where the assumption of event-driven traffic is no longer valid, turning the use of a protocol tailored for specific traffic potentially more beneficial.

The mean energy of B2B is 33.3% lower when using SCS, at the elbow of the curve ($\bar{T}_e = 3$ s). This energy saving further improves to 40% at the highest \bar{T}_e value analyzed over the

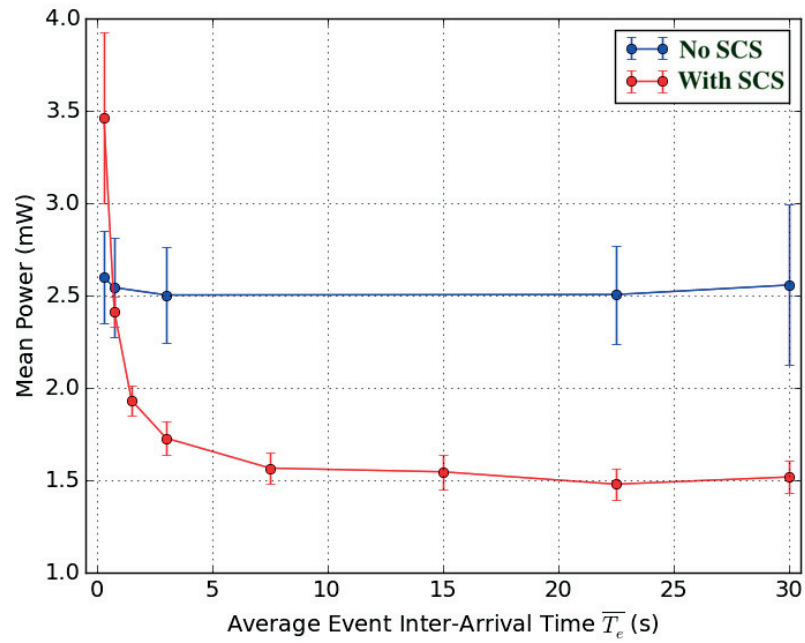


Figure 6.12 – Impact of SCS on B2B’s mean power consumption as a function of the mean event inter-arrival time.

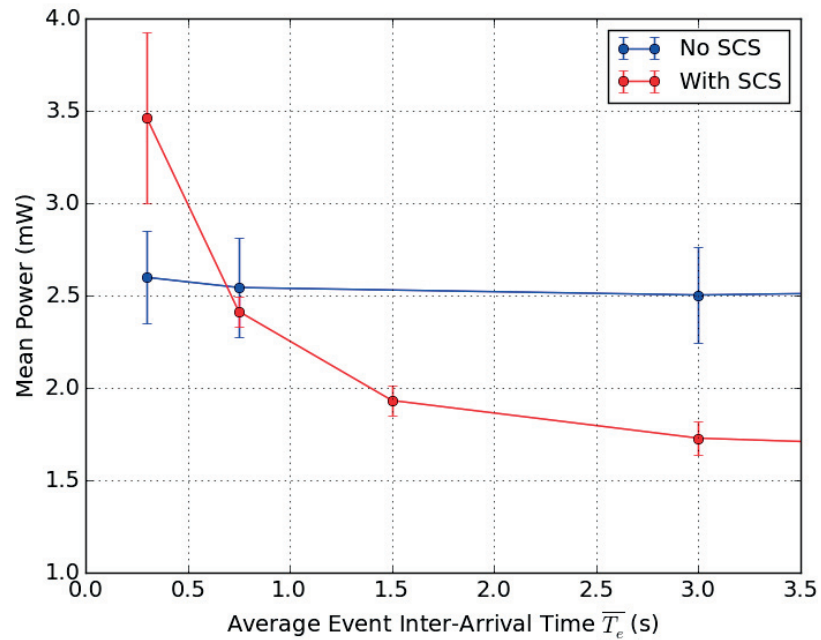


Figure 6.13 – Impact of SCS on B2B’s mean power consumption as a function of the mean event inter-arrival time (augmented view of Fig. 6.12).

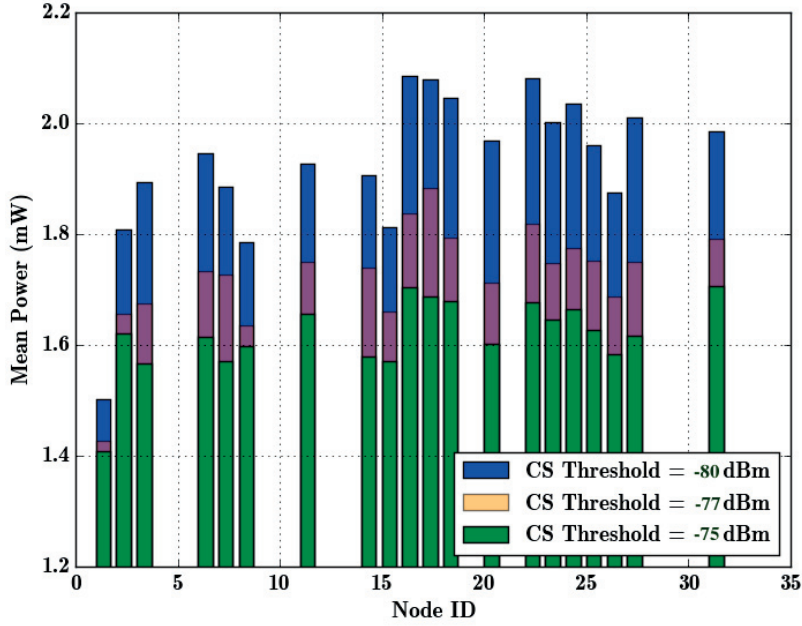


Figure 6.14 – Comparison between the mean power consumption of the SCS-enhanced protocol for different values of the channel sampling threshold.

asymptote.

For each value of \bar{T}_e , the observations made in Sec. 6.5.2 about the latency and reliability were reconfirmed. Moreover, both metrics remain equivalent for both protocols.

Figure 6.14 shows a comparison between the mean power consumption for different values of the channel sampling threshold, when using SCS over B2B. The orange bar uses a CS threshold of -77 dBm (the data is the same as in Fig. 6.10), while the other bars correspond to -75 dBm (green) and -80 dBm (blue). The power consumption decreases with higher values of the CS threshold, because this reduces the instances of false wake-ups. Nevertheless, the value of -80 dBm impacts the reliability because the wake-up tone is occasionally not detected. In all the other cases the latency and reliability remain equivalent to B2B.

6.5.4 Wake-up from Multiple Sources

The objective of this experiment is to evaluate the energy consumption and the reliability of SCS in terms of waking the network up due to events that are reported simultaneously by multiple motes (refer to Sec. 6.4.3). We compare the performance to the case of B2B, modified for supporting distributed event detection using flood extensions as described in Sec. 6.4.3.

The reliability of the network was monitored when one or multiple (5, 8 and 12) motes detect an event and wake the network up during the same SCS-round (S slot in Fig. 6.4). We quantify the proportion of successful wake-ups on all the other motes in the network.

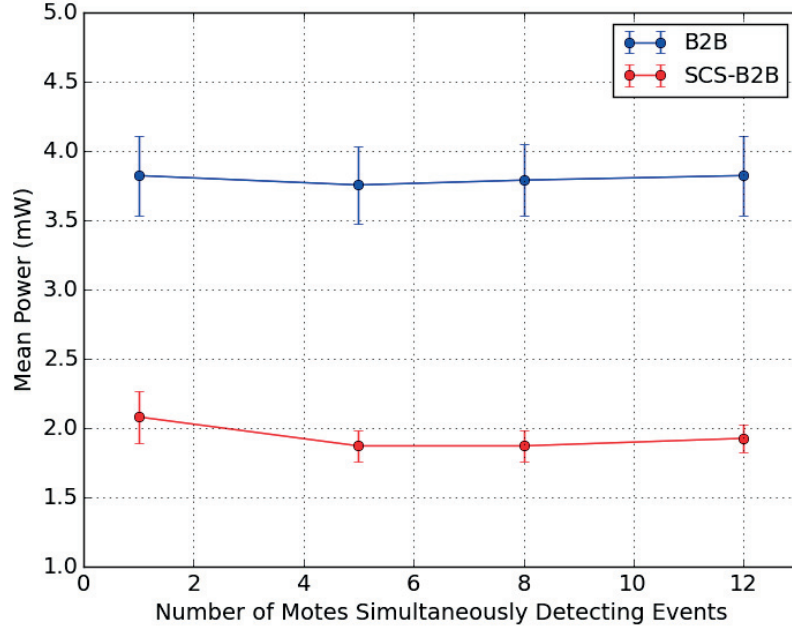


Figure 6.15 – Evolution of the mean power consumption with and without SCS, with the number of motes that simultaneously detect events.

We highlight that this experiment only aims to quantify the reliability of the wake-up operation. Therefore, the motes do not trigger a flood as a response to the wake-up (orange F slot in Fig. 6.4). Provided that the wake-up is reliable, it can be conceived that the network can reorganize itself with another protocol in order to convey the data from each mote.

The results in Fig. 6.15 show that the mean energy consumption of the case using SCS is on average 47% lower than B2B's, while maintaining a 100% wake-up reliability. These results are valid irrespective of the number of motes simultaneously waking the network up.

They are also consistent with the fact that SCS only relies on energy measurements in the channel in order to wake the network up, which does not require the receipt of coherent symbols. Moreover, it can be expected that the robustness of the system will increase with the number of motes simultaneously reporting an event, as it will increase the energy in the channel, thus facilitating its detection.

The energy consumption of each mote with B2B in this experiment is approximately twice the baseline results in Fig. 6.10. This relation is expected as B2B requires the radios to remain active approximately two times longer in the case of distributed event detection, than it does in the case of a single mote. Doubling the active time enables the network to be ready for conveying two floods: a probing from the source and a (possible) flood extension triggered by the mote(s) in case of event(s).

6.6 Optimizations

SCS can be configured to further improve its robustness against false wake-ups, its power consumption or its resilience against malicious attacks. The following points lay the foundations of six schemes:

- **Single wake-up tone:** reduces the number of tones sent upon wake-up to one, by any mote in the network. Fig. 6.12 shows that upon detection of a wake-up signal, a mote transmits the wake-up tone for the remainder of the SCS round. Nevertheless, after the first transmissions in the required sequence of channels, the additional ones provide redundancy and could be suppressed by turning the radio off in order to reduce the power consumption. We expect the impact of this redesign to be a significant reduction in the energy consumption of the worst case scenario, compared to the reduction in a no jammers scenario. The reason behind the impact's dependency on the interference is that the presented optimization only has an effect when the network wakes-up.
- **Network diameter reduction for SCS:** the number of CS performed during an SCS round is calibrated to the diameter of the network, but it can be further reduced to optimize the power consumption. The rationale behind this improvement is that the diameter of the network required to propagate a Glossy-like flood is larger than the one required to propagate an SCS wake up, the reason being that the flood's diameter must factor in the effect of interference, which increases the mean diameter by rendering some of the links unavailable. Conversely, the diameter required for an SCS wake-up is not affected by external interference, due to its reliance on the measurement of noise through the CS mechanism, thus resulting in smaller diameters. For example, the version of SCS analyzed in Sec. 6.12 could be modified to operate with a wake-up diameter of four instead of six hops, while maintaining the flood diameter of six hops. This change would result in significant reductions in the energy consumption, both in absence of external interference and in the worst-case scenario. Moreover, this modification could be used together with the optimization presented in the previous point.
- **Selective CS suppression:** reduces the power consumption by avoiding the CSs that are usually not necessary for relaying wake-up tones. This is prevalent for the motes that are several hops away from the source in the centralized event detection scenario. Figure 6.3 shows that the first CS in D3 is never used, irrespective of the presence of an event; thus suppressing it can reduce the energy consumption without affecting the reliability. The gains through this optimization increase as the motes are further away from the source (in number of hops). Figure 6.3 shows that D4, which is three hops away, can suppress the first two CS and make energy savings twice as large as D3.
- **Dynamic CS threshold:** adapts the RSSI level used during the CS to differentiate between the wake-up tone and the noise baseline. This tuning aims at increasing the

robustness against false wake-ups, and avoiding their associated energy waste.

The adaptation can be performed by each mote individually. Each CS provides a sample of the noise baseline in the environment. The mote will determine afterwards if the sample belongs to a false wake-up, in case the B2B flood is not received after the SCS round, or a true wake-up, if the flood arrives. The mote can use this ground truth to define the new CS threshold to be as high as possible, while being smaller than the typical RSSI associated with the true wake-ups. Similar techniques have been proposed previously in the context of channel sampling MAC protocols, such as [129] and [130].

- **Blacklisting:** avoids channels subject to high interference. Different combinations of SCS channels can be defined by the mote that handles the synchronization and propagated to the network for testing during a time interval. The false network wake-ups resulting from each channel subset can be monitored by the synchronizing mote in order to add in the the black list the channels subject to interference. This feature would also increase the resilience of the network against a malicious denial-of-sleep attack, where external jammers can induce constant false wake-ups on the network through blasting interference in the right combination of channels.

6.7 Conclusion

This chapter has introduced Synchronized Channel Sampling, a mechanism that manages to improve the energy efficiency of event-triggered WSNs by up to 40%, while maintaining an equivalent reliability performance. Its strategy relies on using periodic and synchronous channel sampling in the motes, thus enabling them to detect a wake-up tone emitted in case of an event, without needing to rely on floods. This allows the network to reliably transition towards a consistent state in order to cater to the resulting traffic.

In particular, this study demonstrates SCS' impact on improving the energy performance of B2B, one of the state-of-the-art protocols for event-triggered WSNs. The improved power performance results from SCS' substituting the floods that B2B used for performing the network wake-up with more power efficient synchronous channel sampling rounds. This proves particularly useful in low event frequency scenarios, where the network energy consumption is driven by the wake-up mechanism.

The testbed experiments show that SCS reduces the power consumption of B2B by at least 33.3% and up to 40% in set-ups with a single source. Moreover, SCS seamlessly supports applications where multiple motes can detect events, displaying power consumptions 47% lower than B2B. All evaluations confirm that SCS leads to reliability levels equivalent to B2B's.

While we have chosen B2B as a benchmark protocol due to its state-of-the art performance in terms of reliability, the benefits of SCS can easily be extrapolated to other state-of-the-art protocols for event-triggered systems based on Wireless Sensor Networks, such as eLWB, Glossy or RedFixHop. Moreover, the reduction of transmissions enabled by SCS brings the

Chapter 6. Improving the Energy-Latency Tradeoff in Event-driven Protocols based on Synchronous Transmissions

family of protocols based on synchronous transmissions closer to practical applications, by improving their co-existence with other networks.

Our next steps will focus on extending the evaluation of SCS under increasingly challenging interference scenarios and, in parallel, studying the effectiveness of the optimizations described in Sec. 6.6 for increasing the robustness against interference. We will also extend the study to different target latencies, particularly towards lower values, by characterizing the performance as a function of the probing frequency (T_w).

Given its contribution to enhancing network coexistence, our forthcoming efforts will also be oriented towards the first deployments of this type of protocols with industrial customers.

7 Conclusion

This dissertation has addressed the development of routing protocols designed to meet the specific requirements of multi-hop wireless sensor networks. The protocols thus aimed at providing an improved performance along the three key WSN performance indicators, namely the high reliability in the data transmission, the low energy consumption necessary for a prolonged autonomous operation, and the low and predictable latency needed for a seamless communication. Depending on the underlying applications, the mechanisms also had to display a high tolerance to traffic surges, as well as the ability to self-organize and be readily scalable. While in the first part of the thesis we have focused on collection protocols aimed at applications subject to high reliability requirements and energy efficiency constraints, the second part has dealt with Glossy-like protocols dedicated to applications also requiring very low latencies.

The first part of the thesis has introduced two main contributions, WiseNE, and Rep. Described in Chapter 1, WiseNE is a novel beaconing mechanism that manages to increase the reliability of collection tree protocols, by enabling an energy efficient use of composite metrics. While the gain in accuracy from the use of composite metrics was already considered an established fact in the field, previous protocols attempting at using them relied on de-activating the Trickle Algorithm in order to obtain sufficient samples of the noise-susceptible RSSI (Received Signal Strength Indication) and LQI (Link Quality Indication), key components on composite metrics. This generated a very high increase in the energy consumption. By decoupling the sampling and exploration periods of the protocol, WiseNE manages to maintain the Trickle Algorithm active, thus generating a significantly lower overhead. This development also allowed us to perform the first online validation of the Triangle Metric (composed of PRR, RSSI and LQI).

This overhead is further reduced (by one order of magnitude) with the introduction of Rep, in Chapter 2. Rep manages to achieve this significant drop in energy consumption and traffic, and additionally increase the LQE (Link Quality Estimation) speed by leveraging the pre-existing repetitions in low-power preamble-sampling MAC layers, as opposed to waiting for additional beacons. Given the high frequency variations pattern of the dominant source of interference present in indoor office deployments (WiFi), the fine-grain sampling performed

Chapter 7. Conclusion

by Rep over 1 second achieves the same accuracy as the SOTA coarse-grain sampling over 2.7 minutes, while consuming one order of magnitude less energy.

In Chapter 3, the two mechanisms were also validated in the frame of a real deployment in an outdoors urban environment, as part of the Fly project, a WSN that we have developed and deployed for the monitoring of public utilities in a European Smart City initiative. The experiments conducted proved that WiseNE and Rep continue to deliver their promised improvements in environments subject to significant perturbations, including variations due to the multipath fading, as a result of nodes being installed at ground level and large distances (in the order of hundreds of meters) without line-of-sight. The deployment has also introduced a newly developed composite metric, PRSSI (a combination between PRR and RSSI). The currently installed protocol (a combination between WiseNE based on the RSSI single metric, and Rep) has consistently been operating at a very high reliability (98%) over the 1 month it has been under surveillance, a useful observation given the broad range of reliabilities declared by typical outdoor deployments. Given the scarcity of such deployments, mostly due to the prohibitive cost of monitoring the network, we have decided to also document some of the most important practical learnings, ranging from the need to install a data collection mechanism in order to facilitate a non-intrusive experimentation, to the value of maintaining a close collaboration with the real beneficiary to insure the optimal utilization of the WSN.

Future efforts will focus on a more granular characterization of the scenarios when Rep is to be used on demand, in order to achieve an even further reduction of the energy overhead. Moreover, a more detailed study of the perturbation factors present in outdoors deployments would help further optimize the sampling mechanism. This can be achieved in the second stage of the Fly project. As the applicability of neither WiseNE, nor Rep relies on our use case, we also aim to further explore their potential for low-energy indoor localization systems, the monitoring of mobile networks, Radio Tomographic Imaging, as well as experiment with other composite metrics. The introduction of Rep in unicast and anycast communication primitives, as well as Channel hopping schemes is also worth exploring. Moreover, an interesting venue for further optimizing the reliability and energy performance of the two mechanisms can rely on using the link quality extracted by the two schemes as an additional information input for link estimation schemes based on machine learning.

The second part of the thesis has focused on the development of mechanisms also able to cater to very low latency constraints. This additional requirement was inspired by the need to develop a WSN for the WiseSkin project. Our particular role in this interdisciplinary project consisted in the development of a WSN able to convey the tactile information through the artificial skin of a prosthetic arm. As the ability to create a seamless tactile experience is one of the main criteria for patients' accepting the prosthesis, our WSN had to achieve very low latencies, on top of high reliabilities and an energy consumption compatible with a prolonged use. We have thus investigated the problem of distributed wake-up of an event-triggered WSN with tight latency, reliability and energy consumption constraints.

The first chapter of the second part (Chapter 4 in the overall thesis organization) has introduced Glossy-W, a modified version of Glossy tailored to reducing the energy consumption for a given latency, while catering for multi-node simultaneous event detection. Glossy-W incorporates two main contributions. The first contribution consisted in the validation of the fact that leveraging concurrent transmissions can actually result in better energy performances for a given latency, even though the initial expectation would be an increase in the energy consumptions. The key differentiator stems from the fact that avoiding concurrent transmissions requires contention for access to the channel on every hop, thus necessitating a higher duty cycle in order to strike a latency. Additionally, we show that even removing the per-hop contention by using an Ideal Collision Avoidance MAC, Glossy-W maintains a better energy-latency trade-off for the dense WSNs under consideration. The second contribution consisted in validating that, among protocols leveraging concurrent transmissions, the ones relying on constructive interference are preferable to the ones based on the capture effect, given their superior robustness to simultaneous event detection, a key constraint of high density WSNs.

Glossy-W is thus the first protocol able to leverage floods that rely on constructive interference in a simultaneous event detection scenario. Glossy-W will be embedded in the WiseSkin pilot project to be tested by human patients. While inspired by WiseSkin, Glossy-W has a broad range of applicability, especially as an increasing number of WSN applications require a dense network, consistently fast responsiveness and high energy autonomy. The approach of "leveraging concurrent transmissions" is poised to support novel applications in real-time cyber-physical systems based on WSNs. Nevertheless, extensive work is still required to improve its co-existence with other networks and scalability, before widespread practical applications and standardization become a reality.

The second chapter of the second part (Chapter 5) has introduced a novel primitive: Synchronized Channel Sampling (SCS). SCS consists of consecutive measurements of the signal strength in different channels performed synchronously by all the motes in the network. Its main contribution in improving the state-of-the-art protocols, such as B2B, relies on its providing a more power efficient mechanism for polling the network for events, compared to the traditional approach of sending floods. The improved power performance results from SCS' substituting the floods that B2B used for performing the network polling with more power efficient Synchronous Channel Sampling rounds. This proved particularly useful in low event frequency scenarios, where the network energy consumption is driven by the polling mechanism. The testbed experiments performed demonstrate that the use of SCS reduces the power consumption of the top performing protocol (from the EWSN17 competition) by up to 40%, while keeping an equivalent reliability performance.

Our next steps will focus on extending the evaluation of SCS under increasingly challenging interference scenarios and, in parallel, studying the effectiveness of the optimizations described in Sec. 6.6 for increasing the robustness against interference. We will also extend the study to different target latencies, particularly towards lower values, by characterizing the performance

Chapter 7. Conclusion

as a function of the probing frequency (T_w). Given its contribution to enhancing network coexistence, our forthcoming efforts will also be oriented towards the first deployments of this family for industrial customers. Both SCS and Glossy-W are to be additionally validated in the next round of WiseSkin, in the frame of clinical trials with human patients.

Bibliography

- [1] W. Dargie and C. Poellabauer, *Fundamentals of wireless sensor networks: theory and practice*. John Wiley & Sons, 2010.
- [2] M. A. Mahmood, W. K. G. Seah, and I. Welch, "Reliability in wireless sensor networks: A survey and challenges ahead," *Computer Networks*, vol. 79, pp. 166–187, 2015.
- [3] I. Martin, T. O'Farrell, R. Aspey, S. Edwards, T. James, P. Loskot, T. Murray, I. Rutt, N. Selmes, and T. Baugé, "A high-resolution sensor network for monitoring glacier dynamics," *IEEE Sensors Journal*, vol. 14, no. 11, pp. 3926–3931, Nov 2014.
- [4] N. Zaman, T. Low, and M. Yasin, "Enhancing energy efficiency of wireless sensor network through the design of energy efficient routing protocol," vol. 2016, pp. 1–16, 01 2016.
- [5] A. H. Lee, M. H. Jing, and C. Y. Kao, "Lmac: An energy-latency trade-off mac protocol for wireless sensor networks," in *International Symposium on Computer Science and its Applications*, Oct 2008, pp. 233–238.
- [6] K. Sohrabi, J. Gao, V. Ailawadhi, and G. Pottie, "A self organizing wireless sensor network," in *Proceedings of the Annual Allerton Conference on Communication Control and Computing*, vol. 37. The University; 1998, 1999, pp. 1201–1210.
- [7] A. El-Hoiydi, "Energy efficient medium access control for wireless sensor networks," p. 177, 2005.
- [8] D. E. Comer, *Internetworking con TCP/IP*. Pearson Italia Spa, 2006, vol. 1.
- [9] M. H. Weik, *open systems architecture*. Boston, MA: Springer US, 2001, pp. 1145–1145. [Online]. Available: https://doi.org/10.1007/1-4020-0613-6_12822
- [10] K. Sohraby, D. Minoli, and T. Znati, *Wireless sensor networks: technology, protocols, and applications*. John Wiley & Sons, 2007.
- [11] A. A. A. Alkhatib and G. S. Baicher, "Wireless sensor network architecture," 2012.
- [12] Omniseu.com. (2008) Unicast, multicast and anycast - types of communication in ipv6. [Online]. Available: <http://www.omniseu.com/tcpip/ipv6/unicast-multicast-anycast-types-of-network-communication-in-ipv6.php>

Bibliography

- [13] T. Winter, P. Thubert, and A. Brandt, "Rpl: Ipv6 routing protocol for low power and lossy networks," *IETF Request for Comments 6550*.
- [14] O. Gnawali, R. Fonseca, K. Jamieson, D. Moss, and P. Levis, "Collection Tree Protocol," in *Proceedings of the 7th ACM Conference on Embedded Networked Sensor Systems*, ser. SenSys '09. New York, NY, USA: ACM, 2009, pp. 1–14. [Online]. Available: <http://doi.acm.org/10.1145/1644038.1644040>
- [15] P. Levis, N. Patel, D. Culler, and S. Shenker, "Trickle: A Self-regulating Algorithm for Code Propagation and Maintenance in Wireless Sensor Networks," in *Proceedings of the 1st Conference on Symposium on Networked Systems Design and Implementation - Volume 1*, ser. NSDI'04. Berkeley, CA, USA: USENIX Association, 2004, pp. 2–2. [Online]. Available: <http://dl.acm.org/citation.cfm?id=1251175.1251177>
- [16] K. Srinivasan, P. Dutta, A. Tavakoli, and P. Levis, "An Empirical Study of Low-power Wireless," *ACM Trans. Sen. Netw.*, vol. 6, no. 2, pp. 16:1–16:49, Mar. 2010. [Online]. Available: <http://doi.acm.org/10.1145/1689239.1689246>
- [17] A. Cerpa, J. L. Wong, M. Potkonjak, and D. Estrin, "Temporal properties of low power wireless links: modeling and implications on multi-hop routing," in *Proceedings of the 6th ACM international symposium on Mobile ad hoc networking and computing*. ACM, 2005, pp. 414–425.
- [18] A. Cerpa, N. Busek, and D. Estrin, "Scale: A tool for simple connectivity assessment in lossy environments," *Center for Embedded Network Sensing*, 2003.
- [19] J. Zhao and R. Govindan, "Understanding packet delivery performance in dense wireless sensor networks," in *Proceedings of the 1st international conference on Embedded networked sensor systems*. ACM, 2003, pp. 1–13.
- [20] G. Zhou, T. He, S. Krishnamurthy, and J. A. Stankovic, "Models and solutions for radio irregularity in wireless sensor networks," *ACM Transactions on Sensor Networks (TOSN)*, vol. 2, no. 2, pp. 221–262, 2006.
- [21] A. Cerpa, J. L. Wong, L. Kuang, M. Potkonjak, and D. Estrin, "Statistical model of lossy links in wireless sensor networks," in *Information Processing in Sensor Networks, 2005. IPSN 2005. Fourth International Symposium on*. IEEE, 2005, pp. 81–88.
- [22] C. A. Boano, M. A. Zuniga, T. Voigt, A. Willig, and K. Romer, "The Triangle Metric: Fast Link Quality Estimation for Mobile Wireless Sensor Networks," in *2010 Proceedings of 19th International Conference on Computer Communications and Networks (ICCCN)*, Aug. 2010, pp. 1–7.
- [23] Baccour, Nouha and Koubâa, Anis and Mottola, Luca and Zúñiga, Marco Antonio and Youssef, Habib and Boano, Carlo Alberto and Alves, Mário, "Radio Link Quality Estimation in Wireless Sensor Networks: A Survey," *ACM Trans.*

- Sen. Netw.*, vol. 8, no. 4, pp. 34:1–34:33, Sep. 2012. [Online]. Available: <http://doi.acm.org/10.1145/2240116.2240123>
- [24] D. S. J. De Couto, D. Aguayo, J. Bicket, and R. Morris, “A High-throughput Path Metric for Multi-hop Wireless Routing,” *Wirel. Netw.*, vol. 11, no. 4, pp. 419–434, Jul. 2005. [Online]. Available: <http://dx.doi.org/10.1007/s11276-005-1766-z>
- [25] N. Baccour, A. Koubâa, H. Youssef, and M. Alves, “Reliable link quality estimation in low-power wireless networks and its impact on tree-routing,” *Ad Hoc Networks*, vol. 27, pp. 1–25, Apr. 2015. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S1570870514002637>
- [26] R. Fonseca, O. Gnawali, K. Jamieson, and P. Levis, “Four-bit wireless link estimation.” in *HotNets*, 2007.
- [27] M. Rondinone, J. Ansari, J. Riihijärvi, and P. Mähönen, “Designing a Reliable and Stable Link Quality Metric for Wireless Sensor Networks,” in *Proceedings of the Workshop on Real-world Wireless Sensor Networks*, ser. REALWSN ’08. New York, NY, USA: ACM, 2008, pp. 6–10. [Online]. Available: <http://doi.acm.org/10.1145/1435473.1435476>
- [28] T. Liu and A. E. Cerpa, “TALENT: Temporal Adaptive Link Estimator with No Training,” in *Proceedings of the 10th ACM Conference on Embedded Network Sensor Systems*, ser. SenSys ’12. New York, NY, USA: ACM, 2012, pp. 253–266. [Online]. Available: <http://doi.acm.org/10.1145/2426656.2426682>
- [29] —, “Temporal adaptive link quality prediction with online learning,” *ACM Transactions on Sensor Networks (TOSN)*, vol. 10, no. 3, p. 46, 2014. [Online]. Available: <http://dl.acm.org/citation.cfm?id=2594766>
- [30] L. Mottola, G. P. Picco, M. Ceriotti, c. Gună, and A. L. Murphy, “Not all wireless sensor networks are created equal: A comparative study on tunnels,” *ACM Trans. Sen. Netw.*, vol. 7, no. 2, pp. 15:1–15:33, Sep. 2010. [Online]. Available: <http://doi.acm.org/10.1145/1824766.1824771>
- [31] T. Watteyne, A. L. Diedrichs, K. Brun-Laguna, J. E. Chaar, D. Dujovne, J. C. Taffernaberry, and G. Mercado, “PEACH: Predicting Frost Events in Peach Orchards Using IoT Technology,” *EAI Endorsed Transactions on the Internet of Things*, Jun. 2016. [Online]. Available: <https://hal.inria.fr/hal-01312685/document>
- [32] M. Ceriotti, L. Mottola, G. P. Picco, A. L. Murphy, S. Guna, M. Corra, M. Pozzi, D. Zonta, and P. Zanon, “Monitoring heritage buildings with wireless sensor networks: The torre aquila deployment,” in *Proceedings of the 2009 International Conference on Information Processing in Sensor Networks*, ser. IPSN ’09. Washington, DC, USA: IEEE Computer Society, 2009.
- [33] A. Kounoudes, S. Hirdaris, D. Owen, A. Kalis, and Z. Siokouros, “Lynceus - people localisation for safe ship evacuation during emergency,” 11 2012.

- [34] C. Rojas, D. Piguet, and J. D. Decotignie, "Enabling composite metrics in collection protocols for wsns," in *2016 IEEE 12th International Conference on Wireless and Mobile Computing, Networking and Communications (WiMob)*, Oct 2016, pp. 1–8.
- [35] N. Tsiftes, J. Eriksson, and A. Dunkels, "Low-power wireless ipv6 routing with contikirpl," in *Proceedings of the 9th ACM/IEEE International Conference on Information Processing in Sensor Networks*. ACM, 2010, pp. 406–407.
- [36] A. Dunkels, F. Österlind, and Z. He, "An adaptive communication architecture for wireless sensor networks," in *Proceedings of the 5th international conference on Embedded networked sensor systems*. ACM, 2007, pp. 335–349.
- [37] P. Levis, S. Madden, J. Polastre, R. Szewczyk, K. Whitehouse, A. Woo, D. Gay, J. Hill, M. Welsh, E. Brewer *et al.*, "Tinyos: An operating system for sensor networks," in *Ambient intelligence*. Springer, 2005, pp. 115–148.
- [38] O. Gnawali, R. Fonseca, K. Jamieson, M. Kazandjieva, D. Moss, and P. Levis, "Ctp: An efficient, robust, and reliable collection tree protocol for wireless sensor networks," *ACM Trans. Sen. Netw.*, vol. 10, no. 1, pp. 16:1–16:49, Dec. 2013. [Online]. Available: <http://doi.acm.org/10.1145/2529988>
- [39] K. Levis, "Rssi is under appreciated," in *Proceedings of the Third Workshop on Embedded Networked Sensors, Cambridge, MA, USA*, vol. 3031, 2006, p. 239242.
- [40] T. Gilman, P. Levis, P. Buonadonna, A. Woo, and J. Polastre, "The multihop lqi protocol," in <http://www.tinyos.net/tinyos-2.x/tos/lib/net/lqi>, 2007.
- [41] D. Puccinelli and M. Haenggi, "Reliable data delivery in large-scale low-power sensor networks," *ACM Trans. Sen. Netw.*, vol. 6, no. 4, pp. 28:1–28:41, Jul. 2010. [Online]. Available: <http://doi.acm.org/10.1145/1777406.1777407>
- [42] N. Baccour, A. Koubía, M. Ben Jamía, D. do Rosário, H. Youssef, M. Alves, and L. B. Becker, "Radiale: A framework for designing and assessing link quality estimators in wireless sensor networks," *Ad Hoc Netw.*, vol. 9, no. 7, pp. 1165–1185, Sep. 2011. [Online]. Available: <http://dx.doi.org/10.1016/j.adhoc.2011.01.006>
- [43] M. Zuniga, I. Irzynska, J.-H. Hauer, T. Voigt, C. A. Boano, and K. Roemer, "Link quality ranking: Getting the best out of unreliable links," in *Distributed Computing in Sensor Systems and Workshops (DCOSS), 2011 International Conference on*. IEEE, 2011, pp. 1–8.
- [44] E. Ghadimi, O. Landsiedel, P. Soldati, S. Duquennoy, and M. Johansson, "Opportunistic routing in low duty-cycle wireless sensor networks," *ACM Trans. Sen. Netw.*, vol. 10, no. 4, pp. 67:1–67:39, Jun. 2014. [Online]. Available: <http://doi.acm.org/10.1145/2533686>
- [45] A. El-Hoiydi and J.-D. Decotignie, "WiseMAC: an ultra low power MAC protocol for the downlink of infrastructure wireless sensor networks," in *Ninth International Symposium*

- on Computers and Communications, 2004. Proceedings. ISCC 2004*, vol. 1, Jun. 2004, pp. 244–251 Vol.1.
- [46] A. Dunkels, “The contikimac radio duty cycling protocol,” 2011.
 - [47] T. Clausen, U. Herberg, and M. Philipp, “A critical evaluation of the IPv6 Routing Protocol for Low Power and Lossy Networks (RPL),” in *2011 IEEE 7th International Conference on Wireless and Mobile Computing, Networking and Communications (WiMob)*, Oct. 2011, pp. 365–372.
 - [48] J. Polastre, J. Hill, and D. Culler, “Versatile low power media access for wireless sensor networks,” in *Proceedings of the 2Nd International Conference on Embedded Networked Sensor Systems*, ser. SenSys ’04. New York, NY, USA: ACM, 2004, pp. 95–107. [Online]. Available: <http://doi.acm.org/10.1145/1031495.1031508>
 - [49] K. Srinivasan, M. A. Kazandjieva, S. Agarwal, and P. Levis, “The β -factor: Measuring wireless link burstiness,” in *Proceedings of the 6th ACM Conference on Embedded Network Sensor Systems*, ser. SenSys ’08. New York, NY, USA: ACM, 2008, pp. 29–42. [Online]. Available: <http://doi.acm.org/10.1145/1460412.1460416>
 - [50] A. Dunkels, B. Grönvall, and T. Voigt, “Contiki-a lightweight and flexible operating system for tiny networked sensors,” in *Local Computer Networks, 2004. 29th Annual IEEE International Conference on*. IEEE, 2004, pp. 455–462.
 - [51] R. Lim, F. Ferrari, M. Zimmerling, C. Walser, P. Sommer, and J. Beutel, “Flocklab: A testbed for distributed, synchronized tracing and profiling of wireless embedded systems,” in *Proceedings of the 12th International Conference on Information Processing in Sensor Networks*, ser. IPSN ’13. New York, NY, USA: ACM, 2013, pp. 153–166.
 - [52] M. Doddavenkatappa, M. C. Chan, and A. L. Ananda, “Indriya: A low-cost, 3d wireless sensor network testbed,” in *Testbeds and Research Infrastructure. Development of Networks and Communities*, T. Korakis, H. Li, P. Tran-Gia, and H.-S. Park, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2012, pp. 302–316.
 - [53] K. Szlavecz, A. Terzis, S. Ozer, R. Musaloiu-Elefteri, J. Cogan, S. Small, R. C. Burns, J. Gray, and A. S. Szalay, “Life under your feet: An end-to-end soil ecology sensor network, database, web server, and analysis service,” *CoRR*, vol. abs/cs/0701170, 2007. [Online]. Available: <http://arxiv.org/abs/cs/0701170>
 - [54] T. Istomin, A. L. Murphy, G. P. Picco, and U. Raza, “Data prediction + synchronous transmissions = ultra-low power wireless sensor networks,” in *Proceedings of the 14th ACM Conference on Embedded Network Sensor Systems CD-ROM*, ser. SenSys ’16. New York, NY, USA: ACM, 2016, pp. 83–95.
 - [55] C. Rojas and J. D. Decotignie, “Leveraging mac preambles for an efficient link estimation,” in *2018 IEEE 14th International Conference on Wireless and Mobile Computing, Networking and Communications (WiMob)*, Oct 2018.

Bibliography

- [56] J. Polastre, R. Szewczyk, and D. Culler, "Telos: Enabling ultra-low power wireless research," in *Proceedings of the 4th International Symposium on Information Processing in Sensor Networks*, ser. IPSN '05. Piscataway, NJ, USA: IEEE Press, 2005. [Online]. Available: <http://dl.acm.org/citation.cfm?id=1147685.1147744>
- [57] L. Tang, K.-C. Wang, Y. Huang, and F. Gu, "Channel characterization and link quality assessment of ieee 802.15. 4-compliant radio for factory environments," *IEEE Transactions on Industrial Informatics*, vol. 3, no. 2, pp. 99–110, 2007.
- [58] J. Polastre, J. Hill, and D. Culler, "Versatile low power media access for wireless sensor networks," in *Proceedings of the 2nd international conference on Embedded networked sensor systems*. ACM, 2004, pp. 95–107.
- [59] A. El-Hoiydi and J.-D. Decotignie, "Wisemac: an ultra low power mac protocol for the downlink of infrastructure wireless sensor networks," in *Computers and Communications, 2004. Proceedings. ISCC 2004. Ninth International Symposium on*, vol. 1. IEEE, 2004, pp. 244–251.
- [60] C. Rojas, D. Piguet, and J.-D. Decotignie, "Poster: Single Packet Link Estimation," in *Proceedings of the 2016 International Conference on Embedded Wireless Systems and Networks*, ser. EWSN '16. USA: Junction Publishing, 2016, pp. 263–264. [Online]. Available: <http://dl.acm.org/citation.cfm?id=2893711.2893764>
- [61] M. Schuß, C. A. Boano, M. Weber, and K. Römer, "A competition to push the dependability of low-power wireless protocols to the edge," in *Proceedings of the 2017 International Conference on Embedded Wireless Systems and Networks*, ser. EWSN '17. USA: Junction Publishing, 2017, pp. 54–65. [Online]. Available: <http://dl.acm.org/citation.cfm?id=3108009.3108018>
- [62] K. Langendoen and A. Meier, "Analyzing mac protocols for low data-rate applications," *ACM Trans. Sen. Netw.*, vol. 7, no. 1, pp. 10:1–10:34, Aug. 2010. [Online]. Available: <http://doi.acm.org/10.1145/1806895.1806905>
- [63] F. Osterlind, A. Dunkels, J. Eriksson, N. Finne, and T. Voigt, "Cross-level sensor network simulation with cooja," vol. 0, pp. 641–648, 11 2006.
- [64] U. Colesanti and S. Santini, "The collection tree protocol for the castalia wireless sensor networks simulator," Department of Computer Science, ETH Zurich, Zurich, Switzerland, Technical Report 729, Jun. 2011.
- [65] M. Zuniga and B. Krishnamachari, "Analyzing the transitional region in low power wireless links," in *Sensor and Ad Hoc Communications and Networks, 2004. IEEE SECON 2004. 2004 First Annual IEEE Communications Society Conference on*. IEEE, 2004, pp. 517–526.
- [66] T. Liu and A. Cerpa, "Temporal adaptive link quality prediction with online learning," *TOSN*, vol. 10, pp. 46:1–46:41, 2014.

- [67] N. Baccour, D. Puccinelli, T. Voigt, A. Koubaa, C. Noda, H. Fotouhi, M. Alves, H. Youssef, M. A. Zuniga, C. A. Boano, and K. Römer, *Overview of Link Quality Estimation*. Heidelberg: Springer International Publishing, 2013, pp. 65–86. [Online]. Available: http://dx.doi.org/10.1007/978-3-319-00774-8_3
- [68] C. Rojas and J.-D. Decotignie, “Poster abstract: Enabling a new resource for wsn radio tomographic imaging: Lqi in transitional links,” in *Proceedings of the 14th ACM Conference on Embedded Network Sensor Systems*, ser. SenSys ’16. New York, NY, USA: ACM, to be published, 2016.
- [69] D. LaI, A. Manjeshwar, F. Herrmann, E. Uysal-Biyikoglu, and A. Keshavarzian, “Measurement and characterization of link quality metrics in energy constrained wireless sensor networks,” in *Global Telecommunications Conference, 2003. GLOBECOM’03. IEEE*, vol. 1. IEEE, 2003, pp. 446–452.
- [70] B.-r. Chen, K.-K. Muniswamy-Reddy, and M. Welsh, “Ad-hoc multicast routing on resource-limited sensor nodes,” in *Proceedings of the 2Nd International Workshop on Multi-hop Ad Hoc Networks: From Theory to Reality*, ser. REALMAN ’06. New York, NY, USA: ACM, 2006, pp. 87–94. [Online]. Available: <http://doi.acm.org/10.1145/1132983.1132998>
- [71] C. Gomez, A. Boix, and J. Paradells, “Impact of lqi-based routing metrics on the performance of a one-to-one routing protocol for ieee 802.15.4 multihop networks,” *EURASIP J. Wirel. Commun. Netw.*, vol. 2010, pp. 6:1–6:20, Feb. 2010. [Online]. Available: <http://dx.doi.org/10.1155/2010/205407>
- [72] C. A. Boano, T. Voigt, C. Noda, K. Römer, and M. Zúñiga, “Jamlab: Augmenting sensornet testbeds with realistic and controlled interference generation,” in *Information Processing in Sensor Networks (IPSN), 2011 10th International Conference on*. IEEE, 2011, pp. 175–186.
- [73] K. Iwanicki and M. van Steen, “On hierarchical routing in wireless sensor networks,” in *Proceedings of the 2009 International Conference on Information Processing in Sensor Networks*, ser. IPSN ’09. Washington, DC, USA: IEEE Computer Society, 2009, pp. 133–144. [Online]. Available: <http://dl.acm.org/citation.cfm?id=1602165.1602179>
- [74] G. Mao, B. Fidan, and B. D. Anderson, “Wireless sensor network localization techniques,” *Computer Networks*, vol. 51, no. 10, pp. 2529 – 2553, 2007. [Online]. Available: <http://www.sciencedirect.com/science/article/pii/S1389128606003227>
- [75] J. Wilson and N. Patwari, “See-through walls: Motion tracking using variance-based radio tomography networks,” *IEEE Transactions on Mobile Computing*, vol. 10, no. 5, pp. 612–621, 2011.
- [76] K. Langendoen, A. Baggio, and O. Visser, “Murphy loves potatoes: Experiences from a pilot sensor network deployment in precision agriculture,” in *Proceedings of the*

Bibliography

- 20th International Conference on Parallel and Distributed Processing*, ser. IPDPS'06. Washington, DC, USA: IEEE Computer Society, 2006, pp. 174–174. [Online]. Available: <http://dl.acm.org/citation.cfm?id=1898953.1899108>
- [77] G. Barrenetxea, F. Ingelrest, G. Schaefer, and M. Vetterli, “The hitchhiker’s guide to successful wireless sensor network deployments,” in *Proceedings of the 6th ACM Conference on Embedded Network Sensor Systems*, ser. SenSys '08. New York, NY, USA: ACM, 2008, pp. 43–56. [Online]. Available: <http://doi.acm.org/10.1145/1460412.1460418>
- [78] M. Breza, I. Tomic, and J. McCann, “Failures from the environment, a report on the first failsafe workshop,” *SIGCOMM Comput. Commun. Rev.*, vol. 48, no. 2, pp. 40–45, May 2018. [Online]. Available: <http://doi.acm.org/10.1145/3213232.3213238>
- [79] R. Hartung, U. Kulau, B. Gernert, S. Rottmann, and L. Wolf, “On the experiences with testbeds and applications in precision farming,” in *Proceedings of the First ACM International Workshop on the Engineering of Reliable, Robust, and Secure Embedded Wireless Sensing Systems*, ser. FAILSAFE'17. New York, NY, USA: ACM, 2017, pp. 54–61. [Online]. Available: <http://doi.acm.org/10.1145/3143337.3143338>
- [80] R. P. Narayanan, T. Veedu Sarath, and V. Veetil Vineeth, “Survey on motes used in wireless sensor networks: Performance and parametric analysis,” vol. 8, pp. 67–76, 04 2016.
- [81] J. Beutel, K. Römer, M. Ringwald, and M. Woehrle, *Deployment Techniques for Sensor Networks*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2009, pp. 219–248. [Online]. Available: https://doi.org/10.1007/978-3-642-01341-6_9
- [82] P. A. John, R. Agren, Y. Chen, C. Rohner, and E. C. H. Ngai, “868 mhz wireless sensor network - A study,” *CoRR*, vol. abs/1609.00475, 2016. [Online]. Available: <http://arxiv.org/abs/1609.00475>
- [83] T. Istomin, C. Kiraly, and G. P. Picco, “Is rpl ready for actuation? a comparative evaluation in a smart city scenario,” in *Wireless Sensor Networks*, T. Abdelzaher, N. Pereira, and E. Tovar, Eds. Cham: Springer International Publishing, 2015, pp. 291–299.
- [84] S. Duquennoy, O. Landsiedel, and T. Voigt, “Let the tree bloom: Scalable opportunistic routing with orpl,” in *Proceedings of the 11th ACM Conference on Embedded Networked Sensor Systems*, ser. SenSys '13. New York, NY, USA: ACM, 2013, pp. 2:1–2:14. [Online]. Available: <http://doi.acm.org/10.1145/2517351.2517369>
- [85] C. Haas, J. Wilke, and V. Stöhr, “Realistic simulation of energy consumption in wireless sensor networks,” in *Proceedings of the 9th European Conference on Wireless Sensor Networks*, ser. EWSN'12. Berlin, Heidelberg: Springer-Verlag, 2012, pp. 82–97. [Online]. Available: http://dx.doi.org/10.1007/978-3-642-28169-3_6
- [86] M. Mohammad, X. Guo, and M. C. Chan, “Oppcast: Exploiting spatial and channel diversity for robust data collection in urban environments,” in *Proceedings of the*

- 15th International Conference on Information Processing in Sensor Networks*, ser. IPSN '16. Piscataway, NJ, USA: IEEE Press, 2016, pp. 19:1–19:12. [Online]. Available: <http://dl.acm.org/citation.cfm?id=2959355.2959374>
- [87] D. Dujovne, T. Watteyne, X. Vilajosana, and P. Thubert, “6tisch: deterministic ip-enabled industrial internet (of things),” *IEEE Communications Magazine*, vol. 52, no. 12, pp. 36–41, December 2014.
- [88] F. Adelantado, X. Vilajosana, P. Tuset-Peiro, B. Martinez, J. Melia-Segui, and T. Watteyne, “Understanding the limits of lorawan,” *IEEE Communications Magazine*, vol. 55, no. 9, pp. 34–40, 2017.
- [89] A. Anglès-Vázquez, X. Vilajosana-Guillèn, J. López-Vicario, A. Morell-Pérez, P. Tuset-Peiró, and I. Vilajosana-Guillèn, “Generic empiric propagation model for low power wireless networks operating at the 868 mhz band in smart cities,” *IET Microwaves, Antennas Propagation*, vol. 8, no. 14, pp. 1143–1153, 2014.
- [90] A. El-Hoiydi and J.-D. Decotignie, “Wisemac: An ultra low power mac protocol for multi-hop wireless sensor networks,” in *Algorithmic Aspects of Wireless Sensor Networks*, S. E. Nikolettseas and J. D. P. Rolim, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2004, pp. 18–31.
- [91] A. Varga and R. Hornig, “An overview of the omnet++ simulation environment,” in *Proceedings of the 1st International Conference on Simulation Tools and Techniques for Communications, Networks and Systems & Workshops*, ser. Simutools '08. ICST, Brussels, Belgium, Belgium: ICST (Institute for Computer Sciences, Social-Informatics and Telecommunications Engineering), 2008, pp. 60:1–60:10. [Online]. Available: <http://dl.acm.org/citation.cfm?id=1416222.1416290>
- [92] M. Woehrle, M. Bor, and K. Langendoen, “868 mhz: a noiseless environment, but no free lunch for protocol design,” in *9th int. conf. on Networked Sensing Systems*, ser. INSS, jun 2012, pp. 1–8. [Online]. Available: <http://www.es.ewi.tudelft.nl/papers/2012-Woehrle-868MHz.pdf>
- [93] N. Reijers, G. Halkes, and K. Langendoen, “Link layer measurements in sensor networks,” in *Mobile Ad-hoc and Sensor Systems, 2004 IEEE International Conference on*. IEEE, 2004, pp. 224–234.
- [94] Y. Sabri and N. Kamoun, “Forest fire detection and localization with wireless sensor networks,” in *Revised Selected Papers of the First International Conference on Networked Systems - Volume 7853*, ser. NETYS 2013. New York, NY, USA: Springer-Verlag New York, Inc., 2013, pp. 321–325.
- [95] K. Georgieva, K. Smarsly, M. König, and K. Law, “An autonomous landslide monitoring system based on wireless sensor networks,” in *Computing in Civil Engineering (2012)*, 2012, pp. 145–152.

Bibliography

- [96] G. Simon, M. Maróti, A. Lédeczi, G. Balogh, B. Kusy, A. Nádas, G. Pap, J. Sallai, and K. Frampton, "Sensor network-based countersniper system," in *Proceedings of the 2Nd International Conference on Embedded Networked Sensor Systems*, ser. SenSys '04. New York, NY, USA: ACM, 2004, pp. 1–12.
- [97] J. Farserotu, J.-D. Decotignie, J. Baborowski, P.-N. Volpe, C. Quirós, V. Kopta, C. Enz, S. Lacour, H. Michaud, R. Martuzzi *et al.*, "Tactile prosthetics in wiseskin," in *Proceedings of the 2015 Design, Automation & Test in Europe Conf. & Exhibition*, 2015, pp. 1695–1697.
- [98] C. Rojas and J. D. Decotignie, "Artificial skin for human prostheses, enabled through wireless sensor networks," in *2017 IEEE 23rd International Conference on Embedded and Real-Time Computing Systems and Applications (RTCSA)*, Aug 2017, pp. 1–8.
- [99] "Wiseskin project," <http://www.nano-tera.ch/projects/353.php>, accessed: 2010-04-30.
- [100] E. S. Robinson, "Work of the integrated organism." 1934.
- [101] W. W. V. Srinivasan and K.-C. Chua, "Trade-offs between mobility and density for coverage in wireless sensor networks," in *Proceedings of the 13th Annual ACM International Conference on Mobile Computing and Networking*, ser. MobiCom '07. New York, NY, USA: ACM, 2007, pp. 39–50. [Online]. Available: <http://doi.acm.org/10.1145/1287853.1287860>
- [102] S. Carruthers and V. King, "Connectivity of wireless sensor networks with constant density," in *Ad-Hoc, Mobile, and Wireless Networks*, I. Nikolaidis, M. Barbeau, and E. Kranakis, Eds. Berlin, Heidelberg: Springer Berlin Heidelberg, 2004, pp. 149–157.
- [103] J. D. Bjerknes and A. F. T. Winfield, *On Fault Tolerance and Scalability of Swarm Robotic Systems*. Berlin, Heidelberg: Springer Berlin Heidelberg, 2013, pp. 431–444. [Online]. Available: https://doi.org/10.1007/978-3-642-32723-0_31
- [104] J. Rousselot and J.-D. Decotignie, "When ultra low power meets high performance: The wisemac high availability protocol," in *Proceedings of the 8th ACM Conference on Embedded Networked Sensor Systems*, ser. SenSys '10. New York, NY, USA: ACM, 2010, pp. 441–442. [Online]. Available: <http://doi.acm.org/10.1145/1869983.1870064>
- [105] M. Sha, R. Dor, G. Hackmann, C. Lu, T.-S. Kim, and T. Park, "Self-adapting mac layer for wireless sensor networks," in *Proceedings of the 2013 IEEE 34th Real-Time Systems Symposium*, ser. RTSS '13. Washington, DC, USA: IEEE Computer Society, 2013, pp. 192–201. [Online]. Available: <http://dx.doi.org/10.1109/RTSS.2013.27>
- [106] F. Sutton, R. Da Forno, D. Gschwend, T. Gsell, R. Lim, J. Beutel, and L. Thiele, "The design of a responsive and energy-efficient event-triggered wireless sensing system," in *Proceedings of the 2017 International Conference on Embedded Wireless Systems and Networks*. USA: Junction Publishing, 2017, pp. 144–155.

-
- [107] J. Lu and K. Whitehouse, *Flash flooding: Exploiting the capture effect for rapid flooding in wireless sensor networks*. IEEE, 2009.
- [108] M. Buettner, G. V. Yee, E. Anderson, and R. Han, “X-mac: A short preamble mac protocol for duty-cycled wireless sensor networks,” in *Proceedings of the 4th International Conference on Embedded Networked Sensor Systems*, ser. SenSys ’06. New York, NY, USA: ACM, 2006, pp. 307–320. [Online]. Available: <http://doi.acm.org/10.1145/1182807.1182838>
- [109] F. Ferrari, M. Zimmerling, L. Thiele, and O. Saukh, “Efficient network flooding and time synchronization with glossy,” in *Proceedings of the Int. Conf. on Information Processing in Sensor Networks*, 2011, pp. 73–84.
- [110] M. König and R. Wattenhofer, “Maintaining Constructive Interference Using Well-Synchronized Sensor Nodes,” in *12th Annual International Conference on Distributed Computing in Sensor Systems (DCOSS)*, Washington, D.C., USA, May 2016.
- [111] D. Yuan, M. Riecker, and M. Hollick, “Making ‘glossy’ networks sparkle: Exploiting concurrent transmissions for energy efficient, reliable, ultra-low latency communication in wireless control networks,” in *European Conference on Wireless Sensor Networks*. Springer, 2014, pp. 133–149.
- [112] A. Escobar, J. Garcia-Jimenez, F. J. Cruz, J. Klaue, A. Corona, and D. Tati, “Competition: Redfloxhop with channel hopping,” in *Proceedings of the 2017 International Conference on Embedded Wireless Systems and Networks*. USA: Junction Publishing, 2017, pp. 264–265.
- [113] R. Lim, R. D. Forno, F. Sutton, and L. Thiele, “Competition: Robust flooding using back-to-back synchronous transmissions with channel-hopping,” in *Proceedings of the 14th International Conference on Embedded Wireless Systems and Networks (EWSN 2017)*, Uppsala, Sweden, Feb 2017.
- [114] P. Dutta, S. Dawson-Haggerty, Y. Chen, C.-J. M. Liang, and A. Terzis, “Design and evaluation of a versatile and efficient receiver-initiated link layer for low-power wireless,” in *Proceedings of the 8th ACM Conference on Embedded Networked Sensor Systems*, ser. SenSys ’10. New York, NY, USA: ACM, 2010, pp. 1–14.
- [115] O. Landsiedel, F. Ferrari, and M. Zimmerling, “Chaos: Versatile and efficient all-to-all data sharing and in-network processing at scale,” in *Proceedings of the 11th ACM Conference on Embedded Networked Sensor Systems*, ser. SenSys ’13. New York, NY, USA: ACM, 2013, pp. 1:1–1:14. [Online]. Available: <http://doi.acm.org/10.1145/2517351.2517358>
- [116] “Ewsn’16 dependability competition,” <http://ewsn2016.tugraz.at/cms/index.php>
- [117] D. Son, B. Krishnamachari, and J. Heidemann, “Experimental study of concurrent transmission in wireless sensor networks,” in *Proceedings of the 4th International*

Bibliography

- Conference on Embedded Networked Sensor Systems*, ser. SenSys '06. New York, NY, USA: ACM, 2006, pp. 237–250. [Online]. Available: <http://doi.acm.org/10.1145/1182807.1182831>
- [118] F. Ferrari, M. Zimmerling, L. Mottola, and L. Thiele, “Low-power wireless bus,” in *Proceedings of the 10th ACM Conference on Embedded Network Sensor Systems*, ser. SenSys '12. New York, NY, USA: ACM, 2012, pp. 1–14. [Online]. Available: <http://doi.acm.org/10.1145/2426656.2426658>
- [119] K. Langendoen and A. Meier, “Analyzing mac protocols for low data-rate applications,” *ACM Trans. Sen. Netw.*, vol. 7, no. 1, pp. 10:1–10:34, Aug. 2010. [Online]. Available: <http://doi.acm.org/10.1145/1806895.1806905>
- [120] A. Dunkels, F. Österlind, and Z. He, “An adaptive communication architecture for wireless sensor networks,” in *Proceedings of the 5th International Conference on Embedded Networked Sensor Systems*, ser. SenSys '07. New York, NY, USA: ACM, 2007, pp. 335–349. [Online]. Available: <http://doi.acm.org/10.1145/1322263.1322295>
- [121] F. Osterlind, A. Dunkels, J. Eriksson, N. Finne, and T. Voigt, “Cross-level sensor network simulation with cooja,” in *Local computer networks, proceedings 2006 31st IEEE conference on*. IEEE, 2006, pp. 641–648.
- [122] “Ewsn'17 dependability competition,” <http://www.ewsn2017.org/dependability-competition.html>, accessed: 2017-09-26.
- [123] F. Sutton, R. D. Forno, D. Gschwend, T. Gsell, R. Lim *et al.*, “The design of a responsive and energy-efficient event-triggered wireless sensing system,” in *Proceedings of EWSN'17*, 2017, pp. 144–155.
- [124] C. Rojas and J.-D. Decotignie, “Synchronous transmissions + channel sampling = energy efficient event-triggered wireless sensing systems,” in *Proceedings of the 2018 IEEE International Workshop on Factory Communication Systems*, ser. WFCS '18, 2018.
- [125] T. Sparber, C. A. Boano, S. Kanhere, and K. Römer, “Mitigating radio interference in large IoT networks through dynamic CCA adjustment,” *Open Journal of Internet Of Things (OJIOT)*, vol. 3, no. 1, pp. 103–113, Aug. 2017, proceedings of the 1st International Workshop on Very Large Internet of Things (VLIoT). Munich, Germany.
- [126] M. Sha, G. Hackmann, and C. Lu, “Energy-efficient low power listening for wireless sensor networks in noisy environments,” in *Proceedings of the 12th International Conference on Information Processing in Sensor Networks*, ser. IPSN '13. New York, NY, USA: ACM, 2013, pp. 277–288.
- [127] F. Ferrari, M. Zimmerling, L. Mottola, and L. Thiele, “Low-power wireless bus,” in *Proceedings of the 10th ACM Conference on Embedded Network Sensor Systems*. ACM, 2012, pp. 1–14.

-
- [128] C. Rojas and J.-D. Decotignie, "Competition: Synchronous transmissions + channel sampling = energy efficient event-triggered wireless sensing systems," in *Proceedings of the 2018 International Conference on Embedded Wireless Systems and Networks*, ser. EWSN '18. USA: Junction Publishing, 2018.
- [129] R. Marfievici, P. Corbalán, D. Rojas, A. McGibney, S. Rea, and D. Pesch, "Tales from the c130 horror room: A wireless sensor network story in a data center," in *Proceedings of the First ACM International Workshop on the Engineering of Reliable, Robust, and Secure Embedded Wireless Sensing Systems*, ser. FAILSAFE'17. New York, NY, USA: ACM, 2017, pp. 24–31. [Online]. Available: <http://doi.acm.org/10.1145/3143337.3143343>
- [130] M. Sha, G. Hackmann, and C. Lu, "Energy-efficient low power listening for wireless sensor networks in noisy environments," in *Proceedings of the 12th International Conference on Information Processing in Sensor Networks*, ser. IPSN '13. New York, NY, USA: ACM, 2013, pp. 277–288. [Online]. Available: <http://doi.acm.org/10.1145/2461381.2461415>
- [131] P. Zhang, O. Landsiedel, and O. Theel, "Mor: Multichannel opportunistic routing for wireless sensor networks," in *Proceedings of the 2017 International Conference on Embedded Wireless Systems and Networks*. USA: Junction Publishing, 2017, pp. 36–47.
- [132] C. A. Boano, K. Römer, M. A. Z. Zamalloa, and T. Voigt, "Jam-x: Wireless agreement under interference," *CoRR*, vol. abs/1202.4865, 2012.

Camilo Rojas

30 years old, Costa Rican.

📍 Neuchâtel, Switzerland 📞 +41 787 67 08 81 ✉ camilo.rojas@epfl.ch 📖 [Google Scholar Profile](#)

Education

2014-2018	École Polytechnique Fédérale de Lausanne (EPFL) Ph.D. in Computer and Communication Sciences "Routing for Wireless Sensor Networks: From Collection to Event-Triggered Applications" Sup. Jean-Dominique Decotignie	Lausanne, Switzerland
2011-2013	École Polytechnique Fédérale de Lausanne (EPFL) M.Sc. in Communication Systems (Grade 5,45/6) Minor in Management, Technology and Entrepreneurship Awarded Swiss Confederation Excellence Scholarship	Lausanne, Switzerland
2005-2010	University of Costa Rica Honors B.Sc. Physics (Grade 9,52/10) Honors B.Sc. Electrical Engineering (Grade 9,41/10) Awarded Academic Excellence Scholarship	San José, Costa Rica

Summer internships with full scholarship at: Microelectronics Institute of Barcelona (2010), Balseiro Institute of Argentina (2009), Materials Science Institute of Madrid (2009) and Synchrotron Light Laboratory of Brazil (2008).

Experience

2014-2018	CSEM - Industrial Ph.D. Internet-of-Things R&D. Main projects: <ul style="list-style-type: none">• Skin: A smart artificial skin for humans to restore the sense of touch. Designed and implemented the communication stack of an artificial nervous system consisting of a wireless sensor network. The system is currently being tested with human patients (video).• Rock: A long-term landslide monitoring system. Developed and deployed a mission-critical routing protocol for the dedicated wireless sensor network. The system is currently in use by a client to monitor landslides in the Swiss Alps.• Fox: A smart hydrant for the Internet-of-Things. Designed and implemented a high-reliability routing protocol for a network of connected hydrants. The system is being used by an industrial client in Switzerland for > 1year.	Neuchâtel, Switzerland
2013 (6 months)	Logitech – Software Engineering Intern Integrated the communication stack of multiple wireless peripheral devices to enable novel Human-Computer Interaction schemes. Identified new use cases and lead user validation (Master Thesis).	Lausanne, Switzerland
2011 (6 months)	Intel - Component Design Engineer Owned module development for low power System-on-a-Chip products. Co-designed the camera and display digital interfaces of microprocessors for smartphones.	San José, Costa Rica & Texas, U.S.A.

Supervision of students/junior researchers

Jeremy Weber (EPFL) – Master Thesis, 2016
Title: Comparison of Routing Protocols for IoT.

Ryoma Matsubara (Osaka University) – Research Internship, 2017
Title: Improving the downlink routing in WSNs.

Other

Awarded the Swiss VentureLab sponsorship for the **SmartGem** start-up: An intelligent jewel that enhances personal security, through a discretely embedded electronic device that activates various defense mechanisms in case of danger (validated the business model in Bangalore, India).

Languages: Spanish (Native), English, French (Fluent), Portuguese, German and Romanian (Conversational).

Volunteering Activities: Sea turtle anti-poaching vigilance, planetarium presentation design, etc.

163

Interests: Geopolitics, backpacking, scuba diving, soccer and snowboarding.

Main Publications (Full List)

Synchronous Transmissions + Channel Sampling = Energy Efficient Event-Triggered Wireless Sensing Systems.
C. Rojas and J-D. Decotignie. Presented at IEEE WFCs 2018, Imperia, Italy. **Best Paper Award.**

Leveraging MAC Preambles for an Efficient Link Estimation.
C. Rojas and J-D. Decotignie. Accepted at IEEE WiMob 2018, Lymassol, Cyprus.

Artificial Skin for Human Prostheses, Enabled Through Wireless Sensor Networks.
C. Rojas and J-D. Decotignie. Presented at IEEE RTCSA 2017, Hsinchu, Taiwan.

Enabling Composite Metrics in Collection Protocols for WSNs.
C. Rojas, D. Piguet, J-D. Decotignie. Presented at IEEE WiMob 2016, NYC, U.S.

Tactile prosthetics in WiseSkin.
J. Farserotu, J.D. Decotignie, J. Baborowski, P.N. Volpe, **C.R. Quirós**, V. Kopta, C. Enz, S. Lacour, H. Michaud, R. Martuzzi and V. Koch. In *Proceedings of the 2015 Design, Automation & Test in Europe Conference & Exhibition* (pp. 1695-1697). EDA Consortium.

Competition: Synchronous Transmissions + Channel Sampling = Energy Efficient Event-Triggered Wireless Sensing Systems. **C. Rojas** and J-D. Decotignie. Presented at ACM EWSN 2018, Madrid, Spain.

Poster: Enabling a New Resource for WSN Radio Tomographic Imaging: LQI in Transitional Links.
C. Rojas and J.-D. Decotignie. Presented at ACM SenSys 2016, Stanford, U.S.A.

Poster: Single packet link estimation.
C. Rojas, D. Piguet, J-D. Decotignie. Presented at ACM EWSN 2016, Graz, Austria.

