THESIS FOR THE DEGREE OF LICENTIATE OF ENGINEERING

Synchronous Protocols for Low-Power Wireless

Towards Reliable and Low-Latency Autonomous Networking
for the Internet of Things and Wireless Sensor Networks

BESHR AL NAHAS

Division of Networks and Systems
Department of Computer Science and Engineering
CHALMERS UNIVERSITY OF TECHNOLOGY
Gothenburg, Sweden 2017

Synchronous Protocols for Low-Power Wireless
Towards Reliable and Low-Latency Autonomous Networking for the Internet of Things and Wireless Sensor Networks
Beshr Al Nahas

Technical Report 165L
ISSN 1652-876X
Department of Computer Science and Engineering
Division of Networks and Systems
Distributed Computing and Systems Group

Chalmers University of Technology
SE-412 96 Gothenburg
Telephone +46 31 772 1000

Author e-mail: beshr@chalmers.se

Printed by Chalmers Reproservice
Gothenburg, Sweden 2017

Synchronous Protocols for Low-Power Wireless
Towards Reliable and Low-Latency Autonomous Networking for the Internet of Things and Wireless Sensor Networks

Beshr Al Nahas
Division of Networks and Systems, Chalmers University of Technology

Abstract

With the emergence of the Internet of Things, autonomous vehicles and the Industry 4.0, the need for reliable yet dynamic connectivity solutions is arising. Many of these applications build their operation on distributed consensus. For example, networked cooperative robots and UAVs agree on manoeuvres to execute, and industrial control systems agree on set-points for actuators. Many applications are mission- and safety-critical, too. Failures could cost lives or incur economic losses.

Any wireless network connecting safety-critical devices must be reliable, and often energy-efficient, as many devices are battery powered and we expect them to last for years. It shall be self-forming and self-fixing as well, to allow for reliable autonomous operation; as many applications cannot afford to stop and wait for external configuration. In this context, synchronised communication has emerged as a prime option for low-power critical applications. Solutions such as Chaos or Time Slotted Channel Hopping (TSCH) have demonstrated end-to-end reliability upwards of 99.99%.

In this thesis, we design and implement protocols to support highly reliable and low latency communication in low-power wireless settings. First, we present a standard-based solution that integrates with the 6TiSCH stack (IPv6 over TSCH) without the need of static scheduling or schedule negotiation. Second, we identify key challenges when it comes to implementing the 6TiSCH stack, and demonstrate how these challenges can be addressed. Then, we take a step beyond the standards and focus on synchronous network flooding such as Glossy and Chaos. We show how to enhance them by adding time-slotting and frequency diversity to achieve high reliability and low latency under interference. Finally, we design and realise a network stack that combines and extends ideas from TSCH and synchronous transmissions to achieve highly reliable data delivery with a loss rate lower than 10^{-5} and achieve network-wide consensus with a radio duty cycle of 0.5%. On top of this robust kernel, we enable two- and three-phase commit protocols to provide network-wide consensus.

We implement our protocols, evaluate them on public testbeds of sensor nodes equipped with IEEE 802.15.4 compatible radios and compare to state-of-the-art protocols. We contribute the source code of our main protocols to the community as a step towards enabling ubiquitous connectivity in the context of the Internet of Things.

Keywords Dependability, Industrial Internet of Things, IoT, WSN, Wireless Networks, Sensing, Distributed Computing.

Acknowledgements

I would like to express my deepest gratitude to my mentor and supervisor Olaf Landsiedel. His professional and personal support is invaluable to my development as a researcher. His patience and mentoring techniques of soft-pushing, motivational talks and the rare hard-pushes were essential to the completion of this work. I would like to thank my friend, collaborator, mentor and advisor Simon Duquennoy for his continuous support and his valuable input. Many people think that it is rare to find supervisors that get involved with their students works so closely. I feel extremely lucky to work with Olaf and Simon. They are both intellectual researchers, excellent writers and coders and great mentors.

I am also indebted to my co-supervisor Professor Philippas Tsigas for the valuable discussions and feedback. It gives me great pleasure as well in acknowledging the guidance and the high-level feedback of my examiner Professor Andrei Sabelfeld.

I thank my collaborators Thomas, Atis and Goerge for their amazing work and intellectual feedback. I would like to thank all the great people in the devision of Networks and Systems for providing a friendly work atmosphere and engaging discussions; especially, Babis, Valentin, Iosif, Yiannis, Thomas P., Boel, Elena, Wissam, Aras, Ivan, Tomas, Elad, Magnus, Marina and Ali. I thank my office mates Aljoscha, Thomas and Nasser for the interesting discussions and nice company.

I extend my acknowledgments to the fantastic administration staff with special thanks to Marianne, Eva and Rebecca for magically providing help with all the practical matters from paperwork to office furniture. Special thanks go to Peter H and Lasse for giving the “extra” technical care.

Finally, I owe it all to my parents Randa and Bakri who always grant me unlimited support, love, faith and wisdom. Last but not least, I thank Fawz, my wife, for her love, support and enduring my tamper and irregular work schedule at deadlines.

Beshr Al Nahas, Gothenburg, June 2017

List of Papers

Included Papers

- ▷ Simon Duquennoy, **Beshr Al Nahas**, Olaf Landsiedel, Thomas Watteyne.
  Orchestra: Robust Mesh Networks Through Autonomously Scheduled TSCH,
  Proceedings of the Conference on Embedded Networked Sensor Systems (ACM SenSys), 2015.

- ▷ Simon Duquennoy, Atis Elsts, **Beshr Al Nahas**, George Oikonomou.
  TSCH and 6TiSCH for Contiki: Challenges, Design and Evaluation,
  Proceedings of the Conference Distributed Computing in Sensor Systems (DCOSS), 2017.

- ▷ **Beshr Al Nahas**, Olaf Landsiedel.
  Competition: Towards Low-Latency, Low-Power Wireless Networking under Interference,
  Proceedings of the International Conference on Embedded Wireless Systems and Networks (EWSN), 2016.
  This is an extended abstract of our approach with which we participated in the EWSN Dependability Competition 2016 where we scored the third place.

- ▷ **Beshr Al Nahas**, Olaf Landsiedel.
  Competition: Towards Low-Power Wireless Networking that Survives Interference with Minimal Latency,
  Proceedings of the International Conference on Embedded Wireless Systems and Networks (EWSN), 2017.
  This is an extended abstract of our approach with which we participated in the EWSN Dependability Competition 2017 where we scored the third place again.

- ▷ **Beshr Al Nahas**, Simon Duquennoy, Olaf Landsiedel.
  Network-wide Consensus in Low-power Wireless Networks,
  under submission, 2017.

Not Included Papers

- ▷ **Beshr Al Nahas**, Simon Duquennoy, Venkatraman Iyer, Thiemo Voigt.
  Low-Power Listening Goes Multi-Channel,
  Proceedings of the Conference Distributed Computing in Sensor Systems (DCOSS), 2014.

- ▷ Liam McNamara, **Beshr Al Nahas**, Simon Deuqennoy, Joakim Eriksson, Thiemo Voigt.
  Demo Abstract: SicsthSense - Dispersing the Cloud,
  Proceedings of the International Conference on Embedded Wireless Systems and Networks (EWSN), 2014.

- ▷ Domenico De Guglielmo, **Beshr Al Nahas**, Simon Deuqennoy, Thiemo Voigt, Giuseppe Anastasi.
  Analysis and experimental evaluation of IEEE 802.15.4e TSCH CSMA-CA Algorithm,
  IEEE Transactions on Vehicular Technology (TVT), 2016.

Table of Contents

List of Figures

List of Tables

1
Introduction

Today, connected objects are everywhere. We have become so used to being always connected that we feel nervous when we are not [19]. Our laptops and cell phones are always connected to the Internet, and even our homes are connected, too: Alarm systems, security cameras and the smart grid which feeds our homes with electricity and updates the utility company with our consumption and the grid status. Industrial giants like Ericsson and Cisco predict a growing connectivity and project 29 billions devices to be connected by 2022 [2]. If this connectivity trend lives up to the predictions, a variety of appliances will be connected either to each others only or to the Internet –in what is called the Internet of Things (IoT)– in order to enable remote control, automatic actions and smart behaviour.

Not only small appliances get connected but bigger ones too. Car manufacturers, for example, are racing for increasingly complex connectivity in cars. A car equipped with online maps, media, weather and traffic services is already a decade old. With the recent hype of autonomous cars, however, arises the need for new connectivity methods to enable car-to-car and car-to-infrastructure communication in order to support a safer and a smoother mobility experience. Apart from everyday objects, industrial actors aim to enhance the automation of their factories with sensor networks and connections to cloud services to, for example, predict failures and trigger maintenance procedures automatically, as envisioned by Industry 4.0 [8].

In this thesis, we develop, implement and evaluate network protocols as a step towards enabling ubiquitous connectivity in the context of the Internet of Things. The remainder of this chapter is organised as following: First, I overview the topic in §1.1. Then, I present the research questions the thesis tackles in §1.2 and outline the roadmap of our approach in §1.3. In section §1.4, I summarise our work that appears in the append papers §2, §3, §4, §5 §6, which represent the main body of our contributions. Last, I reflect on the papers in §1.5, and conclude in §1.6.

1.1 Overview

Context. With the emergence of the Internet of Things, autonomous vehicles (*e.g.*, drones, cars) and the Industry 4.0, the need for reliable yet dynamic connectivity solutions is arising. Additionally, many of these applications build their operation on consensus. For example, networked cooperative robots and UAVs agree on manoeuvres to execute [1], and industrial control systems agree on set-points for actuators. Many applications are mission- and safety-critical, too. Failures could cost lives or incur economic losses. A building management system, for example, cannot afford to lose emergent sensor readings signalling a fire alarm. Similarly, autonomous cars crossing an intersection cannot afford a disagreement on which car to go first.

Connectivity Requirements. These applications require a connectivity solution that is *highly reliable:* with an end-to-end reliability suiting everyday applications needs and mission-critical applications, as data loss is undesirable and could be disastrous in critical scenarios; *self forming:* forms a mesh on its own without relying on external components or manual configuration; *self fixing:* copes with the links dynamics and nodes failures; *low power:* as many IoT and WSN devices have a limited energy source such as a battery or energy harvester, and the radio is one of the most energy consuming component in such small devices, we are interested in minimising its energy use; *low latency:* provides a timely information delivery; and suits target applications *traffic requirements* such as periodic traffic or random traffic patterns. The desired connectivity solution shall support rapid *network-wide consensus* as well to enable reliable distributed agreement. In summary, we desire a communication protocol that is (a) highly reliable, (b) autonomous (self forming and self fixing), (c) low power, (d) low latency and (e) suitable to the applications data traffic patterns, and (f) provides network-wide consensus.

Challenge. It is challenging to realise a solution that combines these properties. For example, classic highly reliable solutions use time scheduled medium access (TDMA) approaches with a central scheduler; as in WirelessHART [7] and ISA100.11a [11]. Such solutions lack a timely response to link dynamics and depend on careful network planning; thus, lack the network dynamicity needed in the context of IoT. Other solutions utilise asynchronous low-power listening (LPL). For example, ContikiMAC [3] running over RPL [20] offers a highly flexible alternative but fails to deliver extremely highly reliable communication [5]. The robustness of the TDMA communication approaches pushed the community to standardise it under the relatively new mode of the ZigBee MAC protocol called Time Synchronised Channel Hopping (TSCH) [10]. However, the derivation of the network's schedule is out the scope of this standard.

Standard-based Approach. To this end, we develop a simple yet efficient approach to enable autonomous scheduling for TSCH networks to support the IoT connectivity, which we call Orchestra [5]. We integrate TSCH with the IETF's low-power IPv6 stack to offer a standard-based solution that achieves the connectivity properties we desire. We evaluate Orchestra in testbeds and show how it outperforms state-of-the-art asynchronous solutions while bringing the dynamic network forming and self-fixing to TSCH. In this context, the IETF Working Group 6TiSCH is currently standardising the mechanisms to use TSCH in low-power IPv6 scenarios. We identify a number of challenges when it comes to implementing the 6TiSCH stack, and show how these challenges can be addressed with practical solutions for locking, queuing, scheduling and other aspects.

Synchronous Transmissions Approach. Taking a step beyond the standardised solutions, recent state-of-the-art protocols build on flooding-based synchronous transmissions to offer a more efficient yet simple approach as pioneered by Ferrari *et al.* in Glossy [6]. Glossy capitalises on the capture-effect[1] and non-destructive interference[2] in wireless communications and realises network synchronisation and data sharing by letting nodes synchronously transmit the same data packet. Landsiedel *et al.* builds on Glossy and adds in-network processing to achieve low-latency network-wide data aggregation and collection in Chaos [12]. We get inspirations from Chaos and Glossy as communication paradigms. To achieve higher reliability and lower latency, we enhance the synchronous transmissions with frequency diversity mechanisms. We put this system to test in the EWSN dependability competition [17] in 2016 and 2017 and score the third position twice.

Network-wide Consensus. While highly reliable data delivery is required as discussed, specific applications; such as, cooperative UAVs, need to reliably *agree* on certain actions, *e.g.*, manoeuvres to take. A simple acknowledgment is not sufficient for such consensus, as (a) the receiver that sent the acknowledgment is not sure that it is received, and (b) it is not sure whether the rest

[1] *Capture-effect:* A receiving radio is able to recover one of the many colliding packets under specific conditions related to the used technology. For example, in 802.15.4, the radio can recover a packet of the many different colliding packets if they are synchronised, and the stronger packet signal strength is 2-3 dB above the noise floor. We refer the interested reader to [12] for in-depth evaluation.

[2] *Non-destructive interference:* If the colliding packets are tightly synchronised and have the same contents, then it is highly probable that they do not destruct each other; thus, enabling the receiving radio to recover the contents with a high probability. We refer the interested reader to [6] for in-depth evaluation.

of the nodes agree and that they are aware of its agreement. This requires protocols to guarantee *distributed consensus*. Distributed consensus protocols is a mature field of research in a wired context, but has received little attention in low-power wireless settings. We combine and extend ideas from Chaos and TSCH, and we introduce A^2: Agreement in the Air, which builds a new synchronous transmission kernel that achieves high reliability and low latency through utilising multiple channels in parallel. On top of this robust kernel, we enable two- and three-phase commit protocols (2PC and 3PC) [9, 18] to provide network-wide consensus. In addition, we address the consistent group membership problem and build reliable primitives for nodes to join and leave the network. We implement and evaluate A^2 on testbeds and show how it outperforms state-of-the-art synchronous transmissions solutions.

1.2 Research Questions

This thesis focuses on the following questions:
How do we:
RQ1 *build a highly reliable low-power and low-latency wireless network?*
RQ2 *build an autonomous, self-fixing network?*
RQ3 *achieve network-wide consensus in the above settings?*
The chapters §2, §3, §4, §5 target the first and second questions, while §6 targets the third.

1.3 Roadmap

We use experimental computer science methods in our research. We design and implement protocols targeting real systems, and we evaluate in both simulations and testbeds.

We start by using TSCH since TDMA approaches are known to improve on reliability, latency and energy. However, the main obstacle of using a TDMA style MAC layer is that it requires complex solutions to drive dynamic transmission schedules. We integrate TSCH with the IPv6 stack and use topology information from RPL to create the communication schedules on the fly as we explain in §2.

From this work, we identify challenges and learn important lessons regarding the design and implementation of a synchronised protocol (TSCH) and integrating it in the IPv6 stack to realise 6TiSCH; all of which we discuss in §3.

We build on this knowledge, and come up with a hybrid solution combining ideas from TSCH, synchronous transmissions approaches inspired by Glossy and in-network processing inspired by Chaos. Equipped with this knowledge, we participated in the EWSN Dependability Competition in 2016 and 2017 [17]. We summarise our approaches in §4 and §5, respectively.

Last, to achieve the goals of low-power, low-latency, high reliability, autonomous operation and network-wide consensus in one system, we develop a new MAC layer inspired by ideas from TSCH and Chaos that integrates the knowledge we gained so far. On top of this MAC, we develop a full stack with the focus on network-wide consensus. We present the system in §6.

1.4 Summary of appended papers

In this section, I will give summaries of the papers that constitute the main body of the thesis.

1.4.1 Orchestra: Robust Mesh Networks Through Autonomously Scheduled TSCH

This paper addresses the challenge of bringing TSCH (Time Slotted Channel Hopping MAC) to such dynamic networks. We focus on low-power IPv6 and RPL networks, and introduce Orchestra. In Orchestra, nodes autonomously compute their own local schedules. They maintain multiple schedules, each allocated to a particular traffic plane (application, routing, MAC), and updated automatically as the topology evolves. Orchestra (re)computes local schedules without signalling overhead, and does not require any central or distributed scheduler. Instead, it relies on the existing network stack information to maintain the schedules. The key idea is to provision a set of slots for different traffic planes, and to define the slots in such a way that they can be automatically installed/removed as the RPL topology evolves. This scheme allows Orchestra to build non-deterministic networks while exploiting the robustness of TSCH.

We implement Orchestra in Contiki and demonstrate the practicality of Orchestra and quantify its benefits through extensive evaluation in simulation and two testbeds utilising two hardware platforms. Orchestra reduces, or even eliminates, network contention. In long running experiments of up to 72 hours we show that Orchestra achieves end-to-end delivery ratios of over 99.99%. Compared to RPL in asynchronous low-power listening networks, Orchestra improves reliability by two orders of magnitude with a loss rate of 10^{-4} vs. 10^{-2}, while keeping a good latency-energy balance with twice the energy budget: 1.4% duty cycle for 0.5 second latency vs. 0.8% duty cycle for ContikiMAC.

1.4.2 TSCH and 6TiSCH for Contiki: Challenges, Design and Evaluation

The IETF Working Group 6TiSCH is currently standardising the mechanisms to use TSCH in low-power IPv6 scenarios. This paper identifies a number of challenges when it comes to implementing the 6TiSCH stack. It shows how these challenges can be addressed with practical solutions for locking, queuing, scheduling and other aspects as done in our 6TiSCH implementation for Contiki.
With this implementation as an enabler, we present an experimental validation and comparison with state-of-the-art low power listening (LPL) MAC protocols and CSMA, in terms of reliability, latency and energy. We conduct fine-grained energy profiling, showing the impact of link-layer security on packet transmission.
Through a series of testbed experiments, we find that: (a) Synchronisation in large networks is possible at high accuracy: 97% of the time under 160 μs for a low cost: 0.3% duty cycle in a network of 340 nodes; (b) TSCH, when running dedicated slots, outperforms LPL in all key metrics: reliability, latency, duty cycle; (c) At a micro-level, TSCH and LPL spend about the same amount of energy for receptions, but TSCH has an edge with a factor 3 on transmissions. Link-layer security comes at a low overhead.

1.4.3 Solutions to the EWSN Dependability Competition in 2016 and 2017

We begin with an overview of the competition scenario and evaluation metrics. Then, we proceed with summarising our solutions of the challenge.
The EWSN Dependability Competition [17] reconstructs a scenario for monitoring discrete events under the presence of controlled radio interference. The events are sporadic with a known bound of inter-event interval and a known jitter range, while the radio interference patterns are undisclosed. The evaluation metrics for the competing solutions are:

- *Reliability [%]*: the percentage of events correctly reported at the sink;
- *Latency [milliseconds]*: the delay until an event is detected at the sink, and
- *Energy [Joules]*: the amount of energy collectively consumed by the nodes.

The competing solutions are ranked based on the three metrics, where a higher reliability, lower latency and lower energy consumption are the goals. The competition was held in 2016 and 2017 colocated with the EWSN conference in Graz, Austria and Uppsala, Sweden, respectively. In 2016, the competition featured a dense network setup with 15-20 nodes in 150 m^2 with 3-4 hops, while in 2017 the network was sparser; covering an area of 350 m^2 with 3-7

Year	Place	Nodes [#]	Jammers per node	Area [m^2]	Hops [#]
2016	Graz	15-20	3-5	150	3-4
2017	Uppsala	15-20	3-5	350	3-7

Table 1.1: EWSN Dependability Competition Setup. *We notice the difference in density in the two occasions. We report a range for the number of nodes, as the organisers do not reveal the exact number.*

hops, as summarised in Table 1.4.3. The final evaluation is mainly based on a long run of 45 minutes that starts with low interference but increases and varies with time. In 2017, a second evaluation was added with a short run of 5 minutes that keeps the interference levels high for the duration of the run. We refer the interested reader to Schuß *et al.* [17] for details. We participated in both events and scored the third place in both occasions. The remainder of this section introduces our solutions in both occasions.

A. Competition: Towards Low-Latency, Low-Power Wireless Networking under Interference, 2016.

Chaos [12] performs well under normal operating conditions but not under interference as it does not implement frequency diversity. In this paper, we present a robust version of Chaos under both short- and long-term interference where we extend Chaos with channel hopping, utilise multiple channels in parallel, and employ local and global blacklisting of channels. In addition, we adapt Chaos to the specific application and traffic requirements of the competition settings.

Actual solution. The competition abstract is submitted ahead of the actual competition; thus, the actual solution differs. On the competition day, we decided to not employ blacklisting, and we used a simple random transmission strategy: After the first valid reception from the sink, nodes randomly transmit or wait to receive. They stop the communication round after a fixed number of slots (16) or if they receive a packet that contains an acknowledgement flag signalling that the sink received the data.

Results. We scored the third position in the competition. The actual challenge was to find a competitive balance between the latency and energy usage while maintaining a high end-to-end reliability. Our solution moderately balances these three factors, where we achieve 95.49% reliability for an energy budget

of 699J slightly higher than the best solution. The latency is not bad either although not too competitive with an average of 129 ms compared to 14 ms and 59 ms for the first and second best solutions, respectively[3].
The wining solution has (a) extremely optimised slot length which helps in both latency and energy consumption, and (b) used custom RF channels that occur between the standard 802.15.4 channels which turned out to be interference-free. On the other hand, we have not optimised our implementation to shorten the slot length as we did not have time and we did not think of using non-standard channels as such a usage was not clearly permitted.

B. Competition: Towards Low-Power Wireless Networking that Survives Interference with Minimal Latency, 2017.

Low-power wireless networking needs to survive interference in order to accommodate the requirements of serious applications of Internet of Things. Synchronous transmission techniques like Glossy and Chaos perform well under normal operating conditions. However, their data delivery latency suffers under interference even when extended to use channel hopping.
In this paper, we present a synchronous flooding protocol that incorporates channel hopping to survive interference and to decrease latency. We employ VHT [16] to keep the nodes synchronised even after long periods of no communication with the source node. We keep the design simple to limit the overhead of the protocol and because we believe simplicity helps minimising the number of surprising bugs.

Actual solution. We use simple channel hopping with a fixed sequence to keep the code complexity low. We employ back-to-back transmissions, where every node repeats the packet for a fixed number of times (16) and on a different channel every slot. We optimise the code to compress the slot length. In addition, we provide a smart channel-scan operation to ease node synchronisation in the initial startup phase. A joining node scans the channels and listens on the channel with the least detected noise.

Results. We scored the third position in the competition, again. Our solution is a big improvement compared to the previous one, where we achieve 100% reliability in the main evaluation. Our energy-latency balance is good too with a latency of 87 ms for 1117 J of energy compared to 67 ms for 877 J and 56 ms for 1733 J for the first and second solutions, respectively[4]. The main issues

[3] The results are available at: `http://www.iti.tugraz.at/EWSN2016/cms/index.php?id=49`

[4] The results are available at: `http://www.carloalbertoboano.com/documents/Awards_EWSN2017_DC1.zip`

appeared in the short evaluation with the intense interference, where we lost many events due to bugs in our implementation as we discuss later.

The winning solution used synchronous back-to-back transmissions with channel hopping in a similar manner to our approach. They differed mainly by: (a) using 8 channels only, and (b) limiting the number of transmissions in case of lack of new events. In contrast, we were conservative on optimising the parameters of our protocol; such as, the number of used channels and the number of slots. We used the 16 available channels and 16 slots every round as we considered the possibility for the interference to cover all channels at one slot. In reality, however, there was a small number of jammers per nodes and using 8 channels would have been sufficient. Second, we had bugs in the implementation that lead the joining nodes to listen on the most interfered channel; thus, causing them to take a longer time to join and to lose the first few events, which is exactly the opposite of our intention.

1.4.4 Network-wide Consensus in Low-power Wireless Networks

This paper addresses low-latency and reliable consensus in low-power wireless networks. We argue that new approaches to synchronous transmissions, such as Glossy and Chaos, combined with slotted architecture, are key enablers for such protocols. We present A^2: Agreement in the Air, a system that brings distributed consensus to low-power multi-hop networks. A^2 introduces Synchrotron, a synchronous transmissions kernel that builds a robust mesh by exploiting the capture effect, frequency hopping with parallel channels, and link-layer security. A^2 builds on top of this reliable base layer and enables the two- and three-phase commit protocols, as well as network services such as group membership, hopping sequence distribution and re-keying.

We evaluate A^2 on four public testbeds with different deployment densities and sizes. Our extensive experimental evaluation shows that A^2 (a) is highly reliable, achieving zero losses over millions of points; (b) achieves low power and low latency, *e.g.*, A^2 requires only 475 ms to complete a two-phase commit over 180 nodes with a duty cycle of 0.5% for 1-minute intervals; and (c) enables network-wide agreement, with different consistency/liveness tradeoffs, *e.g.*, when adding controlled failures, we show that two-phase commit ensures transaction consistency in A^2 while three-phase commit provides liveness at the expense of inconsistency under specific failure scenarios.

1.5 Reflections

In this thesis we investigate orthogonal approaches to build highly reliable, low power and low latency autonomous wireless networks. In this section, we dissect, analyse and reflect on our approaches.

Our first approach, Orchestra, builds on top of the TSCH MAC layer and utilises RPL to build and maintain the mesh network, *c.f.*, §2 and §3. TSCH provides a reliable communication substrate and RPL provides the autonomous operation with self forming and self fixing abilities. Orchestra integrates TSCH in the IPv6 stack, extracts topology information from RPL and uses a rules-based approach to create communication schedules that fit the application traffic pattern, while saves power by turning the radio off when no communication is taking place. Orchestra's integration with IP shifts this approach toward the general purpose applications and supports IoT connectivity out-of-the-box at the cost of code complexity.

The second approach builds on synchronous flooding and in-network processing methods, and extends them with time slotting and channel hopping, *c.f.*, §4, §5 and §6. Synchronous flooding gives a reliable communication substrate that is autonomous as well since flooding is a greedy strategy that does not need to maintain routes. In-network processing enhances the low latency of the communication, as it aggregates data while flooding. Frequency diversity and the utilisation of channels in parallel enhances the reliability and decreases the latency of our communication. A^2 combines time slotting together with in-network processing over a parallel channel flooding paradigm to enable the realisation of the reliable consensus services.

The evaluation results present in the appended papers §2 and §3, shows that Orchestra has a higher latency and energy consumption than A^2 although A^2 uses a greedy flooding approach while Orchestra uses directed routing. The reasons behind this are: (a) RPL's overhead of signalling for maintaining the routing tree, while A^2 does not employ signalling since it is flooding-based; (b) TSCH beaconing for maintaining synchronisation, while A^2 utilises VHT [16] and builds on a periodic traffic pattern that maintains the synchronisation as well; (c) RPL's unicast for routing data from each node towards the sink, while A^2 provides all-to-all data sharing with in-network aggregation to limit traffic; and (d) Orchestra is evaluated with an application emulating a random traffic pattern with a fixed inter-packet interval and jitter, while A^2 is evaluated with periodic traffic.

Although A^2 is very efficient, it has key limitations: (a) point-to-point communications are costly as it builds on all-to-all primitives; (b) supporting general random traffic patterns is costly since it builds on periodic operation; and (c)

supporting A^2 on different radio technologies is not a straight forward task as A^2 depends heavily on the capture-effect.
In a nutshell, I do not argue that A^2 is better than Orchestra, or vice-versa. Instead, I argue that for a generic IoT connectivity style, Orchestra is an efficient solution, while for industrial settings and scenarios that require lower latency or rapid network-wide consensus, then A^2 is recommended. It is possible, however, to carefully investigate the introduction of some of A^2 approaches inside Orchestra, such as, in-network aggregation and the consensus primitives.
Finally, it is important to point out key limitations of our work, which we plan to target in our future work: (a) the level of autonomy is limited: The root node or the network coordinator shall be defined statically as we do not have a leader election mechanism; (b) we do not evaluate the connectivity in mobile scenarios. That's said, we realise that having these two features is essential for satisfying the needs of real-world fully autonomous systems; (c) we do not optimise the initial synchronisation of the network; (d) we do not use consensus protocols designed for lossy networks; instead, we use two- and three-phase commit protocols which assume a lossless network; and (e) we evaluate our protocols on platforms equipped with 802.15.4 compatible radios only.

1.6 Conclusions

This thesis introduces network protocols for building highly reliable, low power and low latency autonomous networks. Our solutions run without any central scheduling entity nor schedule negotiation, and provide high reliability communication. We introduce two approaches: (a) a standardised stack, where we integrate TSCH with RPL in the low power IPv6 stack; and (b) a channel hopping synchronous transmissions based stack with in-network processing. We also discuss and address the main challenges that lie in providing a flexible and efficient TSCH and 6TiSCH implementation. Finally, we present a protocol stack that addresses low-latency and reliable consensus in low-power wireless networks as well. We implement these protocols, evaluate them on testbeds and compare their performance to relevant state-of-the-art protocols.

References

[1] Jude Allred, Ahmad Bilal Hasan, Saroch Panichsakul, William Pisano, Peter Gray, Jyh Huang, Richard Han, Dale Lawrence, and Kamran Mohseni. "SensorFlock: An Airborne Wireless Sensor Network of Micro-air Vehicles". In: *Proceedings of the Conference on Embedded Networked Sensor Systems (ACM SenSys)*. 2007.

[2] *Ericsson Mobility Report - On the pulse of the Networked Society*. Tech. rep. Ericsson, Stockholm, Sweden, Nov. 2016. URL: `https://www.ericsson.com/mobility-report`.

[3] Adam Dunkels. *The ContikiMAC Radio Duty Cycling Protocol*. Tech. rep. T2011:13. Swedish Institute of Computer Science, 2011.

[4] Simon Duquennoy, Atis Elsts, Beshr Al Nahas, and George Oikonomou. "TSCH and 6TiSCH for Contiki: Challenges, Design and Evaluation". In: *Proceedings of the Conference Distributed Computing in Sensor Systems (DCOSS)*. 2017.

[5] Simon Duquennoy, Beshr Al Nahas, Olaf Landsiedel, and Thomas Watteyne. "Orchestra: Robust Mesh Networks Through Autonomously Scheduled TSCH". In: *Proceedings of the Conference on Embedded Networked Sensor Systems (ACM SenSys)*. 2015.

[6] Federico Ferrari, Marco Zimmerling, Lothar Thiele, and Olga Saukh. "Efficient Network Flooding and Time Synchronization with Glossy". In: *Proceedings of the Conference on Information Processing in Sensor Networks (ACM/IEEE IPSN)*. 2011.

[7] HART Communication Foundation. *WirelessHART Specification 75: TDMA Data-Link Layer*. HCF_SPEC-75. HART Communication Foundation, 2008.

[8] Alasdair Gilchrist. *Industry 4.0: The Industrial Internet of Things*. 1st. Berkely, CA, USA: Apress, 2016. Chap. 13. ISBN: 1484220463, 9781484220467.

[9] Jim Gray. "Notes on Data Base Operating Systems". In: *Operating Systems, An Advanced Course*. 1978.

[10] IEEE. *802.15.4-2015: IEEE Standard for Local and metropolitan area networks–Part 15.4: Low-Rate Wireless Personal Area Networks (LR-WPANs) 1: MAC sublayer*. 22 April 2016.

[11] ISA. *ISA-100.11a-2011 – Wireless Systems for Industrial Automation: Process Control and Related Applications*. ISA, 2011.

[12] Olaf Landsiedel, Federico Ferrari, and Marco Zimmerling. "Chaos: Versatile and Efficient All-to-All Data Sharing and In-Network Processing at Scale". In: *Proceedings of the Conference on Embedded Networked Sensor Systems (ACM SenSys)*. 2013.

[13] Beshr Al Nahas, Simon Duquennoy, and Olaf Landsiedel. "Network-wide Consensus in Low-power Wireless Networks". In: *under submission*. 2017.

[14] Beshr Al Nahas and Olaf Landsiedel. "Competition: Towards Low-Latency, Low-Power Wireless Networking under Interference". In: *Pro-*

ceedings of the International Conference on Embedded Wireless Systems and Networks (EWSN). 2016.

[15] Beshr Al Nahas and Olaf Landsiedel. “Competition: Towards Low-Power Wireless Networking that Survives Interference with Minimal Latency”. In: *Proceedings of the International Conference on Embedded Wireless Systems and Networks (EWSN)*. 2017.

[16] Thomas Schmid, Prabal Dutta, and Mani B. Srivastava. “High-resolution, Low-power Time Synchronization an Oxymoron No More”. In: *Proceedings of the Conference on Information Processing in Sensor Networks (ACM/IEEE IPSN)*. 2010.

[17] Markus Schuß, Carlo Alberto Boano, Manuel Weber, and Kay Uwe Römer. “A Competition to Push the Dependability of Low-Power Wireless Protocols to the Edge”. In: *Proceedings of the International Conference on Embedded Wireless Systems and Networks (EWSN)*. 2017.

[18] D. Skeen and M. Stonebraker. “A Formal Model of Crash Recovery in a Distributed System”. In: *IEEE Transactions on Software Engineering* SE-9.3 (1983).

[19] Piercarlo Valdesolo. “Scientists Study Nomophobia — Fear of Being without a Mobile Phone”. In: (Oct. 2015).

[20] T. Winter, P. Thubert, A. Brandt, J. Hui, R. Kelsey, P. Levis, K. Pister, R. Struik, JP Vasseur, and R. Alexander. *RPL: IPv6 Routing Protocol for Low power and Lossy Networks*. RFC 6550. IETF, 2012.