Universidad de la República
Facultad de Ingeniería

Improving the Performance of Wireless Sensor Networks using Directional Antennas

Tesis presentada a la Facultad de Ingeniería de la
Universidad de la República por

Javier Schandy

en cumplimiento parcial de los requerimientos
para la obtención del título de
Doctor en Ingeniería Eléctrica.

Directores de Tesis
Dr. Thiemo Voigt . . . . Uppsala University
Dr. Leonardo Steinfeld . . . . Universidad de la República

Tribunal
Dr. Carlo Alberto Boano . . . . Graz University of Technology
Dr. Xavier Vilajosana . . . . Universitat Oberta de Catalunya
Dr. George Oikonomou . . . . University of Bristol

Director Académico
Dr. Leonardo Steinfeld . . . . Universidad de la República

Montevideo
Thursday 27th August, 2020

UNIVERSIDAD DE LA REPÚBLICA

ACTA DE DEFENSA

TESIS DE DOCTORADO

Fecha: Viernes 31 de julio de 2020 .-

Lugar: Montevideo, Facultad de Ingeniería – Universidad de la República.-

Plan de Estudio: Doctorado en Ingeniería Eléctrica.-

Aspirante: **Javier Andrés Schandy Wood**.-

Documento de Identidad: 4.551.376-7

Director/es de Tesis: Dr. Thiemo Voigt(DT); Dr. Leonardo Steinfeld (coDT).-

Tribunal: Dr. Carlo Alberto Boano (Graz University of Technology, Austria);

Dr. George Oikonomou (University of Bristol, UK);
Dr. Xavier Vilajosana (Universitat Oberta de Catalunya, España);
Dr. Thiemo Voigt (Uppsala University);
Dr. Leonardo Steinfeld (IIE).

Los miembros del Tribunal hacen constar que en el día de la fecha el **Sr. Ing. Javier Schandy** ha sido

APROBADO en la defensa de su **Tesis de Doctorado** titulada: **"Improving Performance of Wireless**

Sensor Networks using Directional Antennas"

La resolución del Tribunal se fundamenta en los puntos detallados a continuación:

The thesis investigates how to improve the performance of wireless sensor networks (WSN) using electronically switched directional antennas. Although the benefits of the latter are well-known, especially the higher gain and the interference minimization, only a few works have analyzed in depth how to design solutions that can improve the performance of WSN in which nodes are equipped with directional antennas. The thesis fills this gap and proposes several optimizations and concepts both on a hardware and on a protocol level that are shown to effectively improve network performance. The PhD candidate has addressed a diverse and complete set of concepts in the telecommunication space including antenna design and characterization, network simulation, neighbour discovery and medium access.

The thesis also includes original and important findings and technical contributions in complementary areas of knowledge. From a scientific perspective, the development of DANDi and DirMAC are the deepest contributions, as they go beyond the optimization of existing concepts and protocol schemes. Each proposed concept has been thoroughly evaluated through both simulations and field experiments, which increases the depth of the thesis' contributions.

The thesis document is well structured and well written, and is presented as an integrated and coherent whole. Specifically, the structure of the thesis breaks down the contributions into manageable and easy-to-read units, clearly summarizing their contributions and limitations.

It has an appropriate balance between theoretical background, state of the art and proper evaluation and discussion about the obtained results. The thesis also succeeds in conceptually integrating the three technical contributions into a common flow, despite these were not merged and evaluated together.

The candidate conducted himself professionally throughout the examination process. The oral presentation was excellent and he answered all questions with honesty.
The candidate demonstrated clear understanding of the research area under investigation. We are satisfied that the work documented in the thesis is the candidate's own work, with some elements of the work undertaken collaboratively with other researchers.

This thesis constitutes a clear advancement of the state-of-the-art, and it is clear to the committee that the candidate masters the topic. The thesis, oral presentation and answer to questions demonstrate the candidate's ability to undertake independent research.

Para que conste,

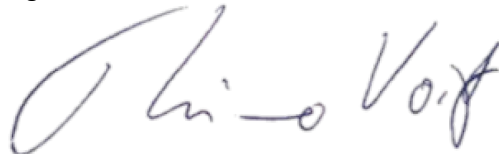*Firmas originales*

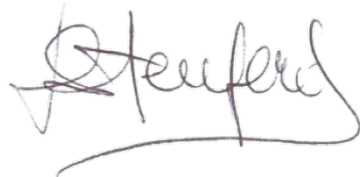**Dr. Carlo Alberto Boano**

Dr. George Oikonomou

Dr. Xavier Vilajosana

Dr. Thiemo Voigt

Dr. Leonardo Steinfeld

This page was intentionally left blank.

Acknowledgments

I would like to thank my thesis advisors Thiemo Voigt and Leonardo Steinfeld for guiding me through my Ph.D. studies and for their constant support and patience. I would also like to especially thank Thiemo for receiving me exceptionally at Uppsala for a research stay, where he was always available to discuss ideas and guide me. I learned a lot from him and his team, and I'm very grateful for the opportunity to see firsthand how they work and do research. I'm also thankful to Dr. Carlo Alberto Boano, Dr. Xavier Vilajosana, and Dr. George Oikonomou for generously accepting the invitation to be a part of the committee.

I would like to express my deepest gratitude to my partners at Focus, Mauricio Gonzalez and Agustin Rodriguez, for their patience, and for giving me the time and space to finish the Ph.D. Last but not least, I would like to especially thank my girlfriend Virginia, my friends, and my family for their constant support and for encouraging me to pursue this adventure.

This page was intentionally left blank.

Abstract

Over the last decades, lots of new applications have emerged thanks to the availability of small devices capable of wireless communications that form Wireless Sensor Networks (WSNs). These devices allow sensing, processing, and communication of multiple physical variables while keeping a low power consumption. During the last years, most of the research efforts were spent on the development and optimization of wireless communication protocols, aiming to maximize the reliability of the network while achieving the lowest possible power consumption.

In this thesis, we study how to improve the performance of these WSNs by using directional antennas. Directional antennas can provide a higher gain and reduce the interference with other nodes by concentrating the radiated power in a certain direction.

We present the different kinds of directional antennas available for WSNs, and we select the 6-element SPIDA antenna as a case of study. We present an electromagnetic model of this antenna, and we incorporate it into the COOJA network simulator. We report the first complete characterization of this antenna, including the radiation pattern and S_{11} parameters. The characterization shows that the antenna has a maximum gain of 6.8 dBi, a Half-Power Beamwidth (HPBW) of 113° and a module of S_{11} parameter of -7.5 dB at the central frequency ($f_c = 2.4525$ GHz). We also present a novel way to optimize the antenna without changing its design by isolating multiple director elements. We show that with this technique, the performance of the antenna can be improved in terms of maximum gain, narrower HPBW, and a lower module of the S_{11} parameter without making any changes in the antenna itself.

We evaluate the impact of supporting directional communications in the different layers of the network stack. We analyze the different challenges that arise and propose optimizations to overcome them in order to take advantage of the benefits of directional communication.

We present an analysis of the state-of-the-art in neighbor discovery protocols for WSNs with directional antennas, and we propose, implement end evaluate two novel fully directional protocols: Q-SAND and DANDi. We compare both of them with SAND, a fully directional neighbor discovery protocol. DANDi is a fully directional asynchronous and dynamic neighbor discovery protocol where the contention resolution relies on a collision detection mechanism. To the best of our knowledge, DANDi is the fastest neighbor discovery protocol for WSN with directional antennas, with the additional advantage of being able to discover every reliable communication link in a network without requiring any prior information

of the network topology.

We combine the directional neighbor discovery protocol with MAC and routing optimizations in order fully take advantage of the benefits of using directional antennas. We focus on convergecast, a typical data collection application where every node sends packets periodically to a sink node. We present DirMAC, a novel MAC protocol that fully supports directional communication, together with four different heuristics to optimize the performance of the protocols. One of these heuristics has the added major benefit of being completely distributed and with no need for offline processing. Our evaluation shows that optimizations at both the MAC and routing layers are needed in order to reap the benefits of using directional antennas for convergecast. Our results show that the performance of the network can be greatly improved in terms of packet delivery rate, energy consumption, and energy per received packet, and that we obtain the largest performance improvements in networks with dense traffic. Simulations with different node densities show that when using directional antennas the PDR increases up to 29%, while energy consumption and energy per received packet decreases by up to 55% and 46% respectively. Experiments with real nodes validate these results showing a significant performance increase when using directional antennas in our scenarios, with a reduction in the RDC and EPRP of 25% and 15% respectively, while maintaining a PDR of 100%.

Resumen

Durante las últimas décadas, la disponibilidad de pequeños dispositivos con comunicación inalámbrica ha permitido el desarrollo de muchas nuevas aplicaciones. Estos dispositivos forman Redes de Sensores Inalámbricos (RSI, o WSN por sus siglas en inglés) que permiten sensar, procesar y comunicar datos provenientes de variables físicas, mientras que mantienen un bajo consumo energético. En los últimos años, la mayor parte de los esfuerzos de la comunidad científica estuvieron concentrados en el desarrollo y optimización de los protocolos de comunicación inalámbricos, buscando maximizar la confiabilidad de la red y minimizar el consumo energético.

En esta tesis estudiamos cómo mejorar el rendimiento de las RSI usando antenas direccionales. Las antenas direccionales pueden proporcionar una mayor ganancia y reducir la interferencia con otros nodos al concentrar la potencia radiada en una cierta dirección.

Comenzamos presentando los distintos tipos de antenas direccionales disponibles para las RSI, y seleccionamos la antena SPIDA de 6 elementos como caso de estudio. Luego presentamos un modelo electromagnético de la antena, que incorporamos al simulador de red COOJA. Construimos un primer prototipo con el que realizamos la primera caracterización completa de ésta antena, incluyendo el patrón de radiación y el parámetro S_{11}. La caracterización muestra que la antena tiene una ganancia máxima de 6,8 dBi, un ancho de haz a mitad de potencia (HPBW por sus siglas en inglés) de 113° y un módulo del parámetro S_{11} de -7,5 dB en la frecuencia central ($f_c = 2{,}4525$ GHz). También mostramos una forma innovadora de optimizar la antena sin cambiar su diseño utilizando varios elementos directores al mismo tiempo. Mostramos que con esta técnica se puede mejorar el rendimiento de la antena en términos de ganancia máxima, ancho de haz a mitad de potencia, y módulo del parámetro S_{11}.

Luego evaluamos el impacto de usar comunicaciones direccionales en las diferentes capas del *stack* de red. Analizamos los diferentes desafíos que surgen y proponemos optimizaciones para sortearlos.

Presentamos un análisis del estado del arte en protocolos de descubrimiento de vecinos en RSI con antenas direccionales, y proponemos, implementamos y evaluamos dos protocolos direccionales: Q-SAND y DANDi. DANDi es un protocolo de descubrimiento de vecinos direccional, asíncrono y dinámico, donde la contienda por el canal se resuelve con un mecanismo basado en la detección de colisiones. Hasta donde sabemos, DANDi es el protocolo de descubrimiento de vecinos más

rápido para RSI con antenas direccionales, con la ventaja adicional de que permite descubrir todos los enlaces de comunicación confiables de una red sin requerir ningún conocimiento previo de la topología.

Luego combinamos los protocolos de descubrimiento de vecinos con optimizaciones en las capas de ruteo y acceso al medio para construir una aplicación de recolección de datos, donde cada nodo envía paquetes periódicamente a un nodo centralizador. Presentamos DirMAC, un protocolo de acceso al medio innovador que soporta comunicaciones direccionales, junto con cuatro heurísticas que permiten optimizar el rendimiento de los protocolos (una de ellas con la ventaja adicional que es totalmente distribuida). Los resultados muestran que usar antenas direccionales en este tipo de aplicaciones permite mejorar sustancialmente el rendimiento de la red, mostrando las mayores mejoras en redes con alto tráfico. Las simulaciones con diferentes densidades de nodos muestran que al usar antenas direccionales se puede aumentar el ratio de entrega de paquetes en hasta 29%, mientras que el consumo energético y la energía por paquete recibido bajan en hasta 55% y 46% respectivamente. Los experimentos en nodos reales validan estos resultados, mostrando una reducción en el consumo energético y en la energía por paquete recibido de 25% y 15% respectivamente, mientras que mantienen un ratio de entrega de paquetes de 100%.

List of Papers

This thesis is based on the following papers, which are referred to in the text by their Roman numerals.

Journal Papers:

I. **Schandy, J.**, Olofsson, S., Gammarano, N., Steinfeld, L., & Voigt, T. (2020). Improving Sensor Network Convergecast Performance with Directional Antennas: A Real World Evaluation. Under Review by ACM Transactions on Sensor Networks.

II. Gammarano, N., **Schandy, J.**, & Steinfeld, L. (2020). Reducing neighbor discovery time in sensor networks with directional antennas using dynamic contention resolution. Design Automation for Embedded Systems, 1-25.

III. **Schandy, J.**, Steinfeld, L., Rodríguez, B., González, J. P., & Silveira, F. (2019). Enhancing Parasitic Interference Directional Antennas with Multiple Director Elements. Wireless Communications and Mobile Computing, 2019.

Conference Papers:

IV. **Schandy, J.**, Olofsson, S., Steinfeld, L., & Voigt, T. (2019, February). Improving Sensor Network Convergecast Performance with Directional Antennas. In Proceedings of the 2019 International Conference on Embedded Wireless Systems and Networks (pp. 13-22). Junction Publishing.

V. Gammarano, N., **Schandy, J.**, & Steinfeld, L. (2018, November). DANDi: Dynamic Asynchronous Neighbor Discovery Protocol for Directional Antennas. In 2018 VIII Brazilian Symposium on Computing Systems Engineering (SBESC) (pp. 16-23). IEEE. **Best Paper Award**.

VI. Gammarano, N., **Schandy, J.**, & Steinfeld, L. (2018). Q-SAND: A Quick Neighbor Discovery Protocol for Wireless Networks with Sectored Antennas. Ninth Argentine Symposium and Conference on Embedded Systems (CASE) (pp. 19-24). IEEE.

VII. Rodríguez, B., González, J. P., Steinfeld, L., **Schandy, J.**, & Silveira, F. (2018). Antenna Characterization without Using Anechoic Chambers or TEM Cells. In Proceedings of the 10th Latin America Networking Conference (pp. 97-101). ACM.

VIII. Rodríguez, B., **Schandy, J.**, González, J. P., Steinfeld, L., & Silveira, F. (2017). Fabrication and characterization of a directional SPIDA antenna for wireless sensor networks. In URUCON, 2017 IEEE (pp. 1-4). IEEE.

Preface

One of the most important things I learned during my Ph.D. studies is the importance of collaboration. In today's world, research advances quickly and we not only have to find good solutions to the problems, but we need to do it in a timely manner. I'm convinced that discussing ideas with other colleagues encourages greater creativity, and the collaboration with researchers from different areas of expertise allows us to cover a wider set of topics and leads to more complete evaluations. This is why in this chapter I want to acknowledge all the researchers that actively collaborated with my research and without whom I couldn't have achieved my goals.

The topics covered in this thesis were also part of a research project entitled "Empowering Wireless Sensor Networks with Directional Antennas applied to Agriculture", funded by the Uruguayan National Research and Innovation Agency (ANII) through their Fondo Maria Viñas grant FMV_1_2014_1_104872. I want to specially acknowledge the researchers who participated in this project: Leonardo Steinfeld, Thiemo Voigt, Benigno Rodriguez, Juan Pablo Gonzalez, Nicolas Gammarano, Santiago Radi, Ignacio Brugnoli, and Fernando Silveira.

The contributions related to the characterization and optimization of the SPIDA Antenna presented in this thesis were partially developed as the final project of a postgraduate course "Advanced Topics in Wireless Systems" that I took as a part of my Ph.D. studies, with the collaboration of Benigno Rodriguez and Juan Pablo Gonzalez. The ten prototypes of the SPIDA antenna were fabricated with the help of Ignacio Brugnoli.

The initial implementation of SAND in Contiki-OS and the contributions related to the Q-SAND neighbor discovery protocol were proposed by the undergraduate students Nicolas Gammarano, Joaqun Facal, and Alex Gurevich as the final project of a wireless sensor networks course supervised by Leonardo Steinfeld and me. Gammarano later started his master's degree continuing his work in neighbor discovery in networks with directional antennas. The original design and implementation of DANDi were developed by Gammarano in the context of his master's thesis, which were only operational in simulation. We then worked together in the modifications needed for the protocol to work in a real environment with interference and lossy links, and in the evaluation of the protocol, both in simulations and with real nodes. The design and evaluation of DANDi will also be included in Gammarano's master's thesis.

Finally, I want to acknowledge Simon Olofsson and Ambuj Varshney for their

contribution to the initial idea and implementation of the medium access protocols discussed in this thesis.

Report on the doctoral Thesis

Title: Improving Performance of Wireless Sensor Networks using   Directional Antennas
PhD Candidate: Javier Schandy
Examiner (Revisor): Prof. Xavier Vilajosana
Date: 7th of July 2020

Summary

The thesis is divided in 3 main topics organized in 6 chapters. Technical achievements by the candidate are presented in chapter 3, 4 and 5. Chapter 2 describes the tools and software/hardware ecosystem used by the candidate to conduct the proposed research. Chapters 1 and 6 introduce and conclude the work.

The central topic addressed by the thesis is the exploration of different techniques to improve the performance of low power wireless nodes by using directional antennas. The research conducted along the thesis addresses different technical aspects including the creation, optimization and characterization of an electronically switched directional antenna, studies mechanisms to improve neighbour discovery in the context of wireless mesh sensor networks and develops novel medium access mechanisms to exploit directional antennas and optimize the network operation to respond to the needs of convergecast traffic.

Chapter 3 specifically focus in the electronically switched directional antenna study. In it the candidate implements and characterises a SPIDA antenna. From the characterization it implements a module for the Cooja simulator in order to take advantage of simulation tools to further explore the topic.

With Cooja an extensive set of simulations are used to determine optimal use of the antenna director elements, exploring different configurations and combinations.

Chapter 4. presents a detailed analysis of the state of the art in protocols for Neighbour Discovery in networks with directional antennas and proposes and evaluates two novel fully directional neighbour discovery protocols taking advantage of sectored antennas. These protocols are Q-SAND and DANDi.

Q-SAND is evaluated through simulation based on the SAND protocol implementation comparing it to SAND considering the neighbor discovery time.
DANDi is both evaluated using simulation and practical experimentation. For it, discovery time, discovered links, and robustness to real-world conditions like interference and lossy links are evaluated.

Chapter 5 takes a cross-layer perspective to exploit the possible benefits of directional antennas. On of the main contributions of this chapter is a MAC layer design and evaluation, referred to as DirMAC. DirMAC is the first protocol exploiting ESD antennas to handle medium access. A second contribution presented in the chapter is a novel technique to optimize antenna pair selection among wireless nodes. This also includes a parent selection protocol (i.e routing) that aims to minimize the number of selected parents to establish a link to. Since the selection aims to favour well connected links the PDR is shown to increase in scenarios under network interference.

Evaluation

The thesis aims to improve the performance of low power wireless nodes by using directional antennas. The rationale behind this approach is that directional antennas can provide a higher gain and reduce the interference with other nodes by concentrating the radiated power in a certain direction. This approach has not been well studied in the literature and therefore the proposed work can be considered a scientifically sound and valuable research effort.

The PhD candidate has addressed a diverse and complete set of concepts in the telecommunication space including antenna design and characterization, network simulation, neighbour discovery and medium access. The work combines theoretical, practical and simulation approaches covering demonstrating ability to find the proper methodology to drive his research.

The thesis also includes original and important findings and technical contributions in complementary areas of knowledge. The characterization of the SPIDA antenna as well as the extensive set of simulations to determine the proper use of director elements are basis for the future research in the area from which new researcher can develop upon. The PhD candidate have also explored key questions in the area of networking once electronically configurable antennas are in place. A key and fundamental question is how the network is formed and how mesh topologies are constructed given the new options that such an antenna setting bring. In this line novel and pioneering work has been done in order to speed up neighbour discovery when nodes need to search for neighbours using directional antennas. The Q-SAND and the DANDI protocols constitute clear advances in the state of the art as backed by relevant publications. Finally, the candidate has when through a cross-layer a medium access control protocol and routing design exploiting the possibility to dynamically focus the antenna. The results demonstrate that exploiting antenna configurability can lead to significant improvements in the operation of mesh networks.

The thesis document is well structured, synthetic and well written. It has an appropriate balance between theoretical background, state of the art and proper evaluation and discussion about the obtained results.

As a concluding remark, this thesis successfully combines deep scientific insights with solid theoretical underpinning and an acute awareness of engineering issues. I thus warmly recommend that Javier Schandy be awarded the doctoral degree after successful presentation of his achievements.

Prof. Xavier Vilajosana

Universitat
Oberta
de Catalunya

Institut für
Technische Informatik

Institutsleiter:
Univ.-Prof.Dipl.-Inform.Dr.sc ETH Kay Römer
Stellvertretender Institutsleiter:
Ao.Univ.-Prof. Dr.techn. Eugen Brenner

Inffeldgasse 16/1
A-8010 Graz
AUSTRIA

Ass.Prof. Dr. Carlo Alberto Boano
Tel.: +43 (0) 316 873-6413
Fax.: +43 (0) 316 873-6903
E-Mail: cboano@tugraz.at

DVR: 008 1833   UID: ATU 574 77 929

Graz University of Technology

Dr. Leonardo Steinfeld
Universidad de la República
Instituto de Ingenieria Electrica, Facultad de Ingenieria
Julio Herrera y Reissig 565
11300 Montevideo - Uruguay

Graz, 14. July 2020

Evaluation report for the PhD thesis "Improving Performance of Wireless Sensor Networks using Directional Antennas" authored by Javier Schandy

The thesis investigates how to improve the performance of wireless sensor networks (WSN) using electronically switched directional antennas. Although the benefits of the latter are well-known, especially the higher gain and the interference minimization, only a few works have analyzed in depth how to design solutions that can improve the performance of WSN in which nodes are equipped with directional antennas. The thesis fills this gap and proposes several optimizations and concepts both on a hardware and on a protocol level that are shown to effectively improve network performance.

Thesis summary. The thesis includes 6 chapters on at total of 97 pages. Besides three chapters presenting the main technical contributions of the thesis, two introductory chapters (containing the motivation of the work, a summary of the contributions, an outline, as well as background information on the employed tools) and a conclusive chapter (summarizing the contributions, discussing their limitations, as well as possible future work) are included.

Chapter 3 presents the modeling, fabrication, and characterization of a low-cost six-element SPIDA antenna designed for the 2.4 GHz ISM band. Simulation and real-world measurements show that the performance of the fabricated antennas is comparable or better than that of similar prototypes. The antenna model has been plugged into the COOJA simulator: this allows to simulate the benefits of directional antennas directly on a full-fledged and popular WSN simulator. Another contribution presented in this chapter is an optimized antenna design by using multiple director elements at the same time, which is shown to improve the antenna gain and directivity in both simulation and field experiments. Simulation and measurements are in good agreement with each other and show that the proposed approach outperforms the state of the art.

Chapter 4 focuses on neighbor discovery schemes for WSN making use of directional antennas. After analyzing existing approaches and highlighting their limitations, two directional protocols for neighbor discovery are proposed. Q-SAND is an extension of SAND (Sectored-Antenna Neighbor Discovery) that can, under certain assumptions, find the combination of antenna sectors exhibiting the highest signal strength between node pairs several times faster than the original SAND version. DANDi is a novel discovery protocol that exploits collision detection to improve contention resolution without requiring any prior information of the network topology. Both protocols have been evaluated both in simulation and in real-world office settings, showing that they can outperform SAND.

Chapter 5 presents DirMAC, an extension of the ContikiMAC radio duty cycling protocol that supports directional communications and four heuristic mechanisms to optimize its performance; some of these involving cross-layer optimizations with the routing layer. The protocol is designed for convergecast applications and a thorough evaluation shows that it can offer a comparable or better performance compared to the use of omnidirectional antennas, especially when using higher transmission rates. The key contribution here consists in experimentally showing the performance trade-offs when making use of directional antennas in convergecast applications. Among others, the results are very valuable as a guideline to WSN application designers.

Thesis assessment. The thesis presents original work and important findings that have resulted in three journal articles and five conference papers at regarded IEEE/ACM venues; some of which also received an award. Among the three contributions, from a scientific perspective, the development of DANDi and DirMAC are the deepest ones, as they go beyond the optimization of existing concepts and protocol schemes. Each proposed concept has been thoroughly evaluated through both simulations and field experiments, which increases the depth of the thesis' contributions. Even though some implementation details could have been described in more depth and the proposed concepts have not been integrated into a single system, the three main contributions are conceptually solid and clearly advance the state of the art in the field.

Thesis write-up. The thesis is well-written and there are only a few remaining typos and formatting issues that can be easily fixed before the submission of the final version (this reviewer will communicate these directly to the candidate). The thesis exhibits a good balance between theoretical and practical concepts. Specifically, the structure of the thesis breaks down the contributions into manageable and easy-to-read units, clearly summarizing their contributions and limitations. The thesis also succeeds in conceptually integrating the three technical contributions into a common flow, despite these were not merged and evaluated together.

Overall evaluation. Based on the above assessment, it is clear that the problems tackled in the thesis are both important and timely. The thesis successfully combines novel scientific insights with thorough experimental work. Therefore, I recommend the acceptance of the submitted thesis and warmly recommend that Javier Schandy is awarded the doctoral degree in electrical engineering.

Yours sincerely,

Ass.Prof. Dr. Carlo Alberto Boano

Contents

Contents

Chapter 1

Introduction

In this chapter, we introduce the work presented in this thesis. We discuss the motivation of our work, we highlight the main contributions and we present a brief summary of the organization of this thesis.

1.1   Motivation

Over the last decades, lots of new applications have emerged thanks to the availability of small devices capable of wireless communications. This allowed sensing, processing, and communication of multiple physical variables or interacting with the physical world while keeping a low power consumption. These devices are not only expected to be of low cost and small size, but also to reach years of autonomy with small batteries, forming Wireless Sensor Networks (WSNs) with low operational and maintenance costs.

The IEEE 802.15.4 [1] protocol is nowadays one of the most widely accepted standards in the 2.4 GHz ISM band for WSNs [44, 56, 58]. During the last decade, most of the research efforts were spent on the development and optimization of wireless communication protocols. The main goals were to try to maximize the reliability of the network while achieving the lowest possible power consumption [19, 20, 27, 50].

However, the design and optimization of antennas for WSN nodes have not been studied thoroughly, even though it could have helped to achieve a lower power consumption and better efficiency. Most of the commercial nodes come with integrated omnidirectional antennas that radiate the energy in every direction, resulting in a sub-optimal behavior for certain applications. Making the antenna directional could provide a better gain without increasing the overall radiated power, or it may extend the battery lifetime since the output power can be reduced. Another advantage of improving the antenna is that it may reduce the interference with other nodes by concentrating the radiated power in a certain direction, and thus reducing the congestion that is known to be an issue in multi-hop WSNs [8]. Also, directional antennas have been proposed as an alternative to increasing the security in WSN because of their immunity to eavesdroppers or jammers placed

outside their narrow radiation region and because of their feature of determining the exact position of a sender node using the signal's angle of arrival [9].

Directional antennas have proven significant benefits for several applications like high-throughput bulk forwarding [59] or radio tomographic imaging [64], but the benefits of using them for data collection applications like convergecast has been questioned stating they provide limited benefits that can only be achieved under specific conditions [55].

These results motivated us to make our own evaluation and try to answer the question: can we improve the performance of WSNs by using directional antennas? This question is very difficult to answer because WSNs have countless applications and each one can have very different requirements [41]. It also depends on which directional antenna we choose to use, how we use it, and how we integrate it into our application. The selection and optimization of the network protocols are other key aspects to consider when we try to answer this question, and that will be carefully analyzed in the following chapters of this thesis.

The first decision we have to make when working with directional antennas is which antennas we are going to use. This decision is based on many factors such as the cost, the size, the ease of fabrication, the maximum gain, and the Half-Power Beamwidth (HPBW). In this work, we analyze different kinds of antennas available for WSN and we select the SPIDA antenna as a case of study.

One of the first problems that arises when using directional antennas is how to discover all the possible communication links between neighboring nodes. When using omnidirectional antennas, one single broadcast message is enough to query neighbor nodes, and the ones in range may reply including their network address. But with directional antennas, there may be many sector-to-sector (S2S) links between two nodes, making the neighbor discovery mechanism an interesting challenge. This mechanism must guarantee that every S2S link is discovered in the shortest time possible.

Another problem we face is that radio communication is often the most power-consuming activity in WSN, and the energy consumption is dominated by idle listening [2, 51]. But the existing RDC mechanisms do not work outright with directional antennas as they would have to periodically listen for incoming packets in every possible antenna direction, significantly increasing the power consumption.

The S2S links used for the communication between a pair of nodes should also be carefully selected. Minimizing the number of antenna directions in use by the nodes could prevent unnecessary transmissions and listenings in directions with no neighboring nodes, while maximizing the received signal strength could be used to improve the link quality. The routing protocol should also be carefully analyzed, as minimizing the number of links that each node establishes could also impact on the power consumption.

1.2 Main Contributions

The main contribution of this thesis is having enabled to improve the performance of a WSN application by using electronically switched directional antennas. We

show that in order to reap the benefits of using this kind of antennas, we need to jointly optimize the neighbor discovery, medium access and routing protocols.

The particular core contributions of this thesis are described as follows.

1. **Simulation Models and Antenna Fabrication and Characterization**
   We fabricate and perform a complete characterization of a SPIDA antenna, including the radiation pattern and S_{11} parameter. We also fabricate ten working prototypes for experimental evaluation. We propose and implement a model for the COOJA network simulator that allows the use of the SPIDA antenna in simulated WSN nodes, calculating the antenna gain based on the angle between the transmitting and receiving nodes (Papers [IV, VII, VIII]).

2. **Optimization of the SPIDA antenna**
   We propose a novel technique to optimize the SPIDA antenna by using multiple director elements at the same time. We present simulations and real-world measurements analyzing different antenna configurations and show that we can improve the performance in terms of maximum gain, narrower HPBW, and a lower module of the S_{11} parameter without changing the geometry of the antenna (Paper [III]).

3. **Novel Fully Directional Neighbor Discovery Protocols**
   We propose, implement, and evaluate two novel fully directional neighbor discovery protocols for WSNs with directional antennas: Q-SAND and DANDi. To the best of our knowledge, DANDi is the fastest neighbor discovery protocol in the state-of-the-art for WSN with directional antennas, with the additional advantage of being able to discover every reliable communication link in a network without requiring any prior information of the network topology (Papers [II, V, VI]). This contribution is the result of a collaborative work with Nicolas Gammarano. The original design and implementation of DANDi were developed by Gammarano in the context of his master's thesis, which were only operational in simulation. We then worked together in the modifications needed for the protocol to work in a real environment with interference and lossy links, and in the evaluation of the protocol, both in simulations and with real nodes.

4. **Improve Convergecast Performance with Directional Antennas**
   We present a novel approach to jointly optimize the neighbor discovery, medium access, and routing protocols to support directional communication in a convergecast application. Results show that in dense networks, these protocols can significantly improve the performance of the network by increasing the packet delivery rate and decreasing the radio duty cycle compared to networks using omnidirectional antennas (Papers [I, IV]).

3

1.3 Thesis Organization

The rest of the thesis is divided in five chapters. In Chapter 2 we present the hardware platform, network stack, and tools we used for our experimental evaluation in the different chapters of this thesis.

In Chapter 3, we study the different kind of directional antennas available for WSNs. We select the SPIDA antenna as a case of study. We fully characterize it and build ten working prototypes. After the characterization, we propose and evaluate optimizations to the antenna without changing its geometry.

We also evaluate the impact of supporting directional communications in the different layers of the network stack. We analyze the different challenges that arise and propose optimizations to overcome them in order to take advantage of the benefits of directional communication.

In Chapter 4, we analyze the neighbor discovery mechanism, and we propose and evaluate two novel fully-directional protocols: Q-SAND and DANDi. To the best of our knowledge, DANDi is the fastest neighbor discovery protocol in the state-of-the-art for WSN with directional antennas, with the additional advantage of being able to discover every reliable communication link in a network without requiring any prior information of the network topology.

In Chapter 5, we present a novel approach to jointly optimize the neighbor discovery, medium access and routing protocols to support directional communication in a convergecast application. Our experimental results, both in simulation and with real nodes, show that when the traffic is dense, networks with directional antennas can significantly outperform networks with omnidirectional ones in terms of packet delivery rate, energy consumption, and energy per received packet.

Finally, the conclusions of each chapter are summarized in Chapter 6. In this chapter, we also present a general discussion of the research problem and future directions of work.

Chapter 2

Tools for Experimental Evaluation

In this chapter, we present the hardware platform, network stack, and tools we use for our experimental evaluations in the different chapters of this thesis.

2.1 Hardware Platform

One of the most widely used hardware platforms for WSNs is the TMote-Sky [36] presented in Figure 2.1. The TMote-Sky is an open hardware platform, originally manufactured by Moteiv but now offered by many suppliers. It is a wireless node that features an 8 MHz MSP430F1611 [26] low-power microcontroller, a CC2420 [6] radio transceiver compliant with the 802.15.4 standard, an external flash of 1 MB and several sensors. We take this network node as a reference and use it in the following chapters to evaluate different network protocols for several reasons. The first and most practical reason is that we have several units in our laboratory. The second reason is that it can be simulated (including the wireless communications) with the COOJA network simulator [42], which allows us to rapidly deploy networks with different topologies to test different network protocols. The TMote-Sky also has an external SMA connector and available GPIO pins that allow the connection of external antennas. Last but not least, it is one of the most widely used platforms for WSN evaluations, and it is used by most of the related work we will analyze in the following chapters.

Figure 2.1   TMote-Sky wireless node.

2.2   Operating System

We choose to program the TMote-Sky nodes using Contiki OS [13]. Contiki OS is an open-source, event-driven operating system oriented to WSN applications using constrained hardware like the TMote-Sky node. This operating system manages the hardware resources and includes different libraries such as network stacks and a file system. Contiki's scheduler manages sleep modes, powering down the microprocessor when there is neither processing needed nor events scheduled in the event queue.

2.3   Network Stack

We also work with Contiki's implementation of one of the most widely used network stacks for WSNs, shown in Table 2.1. Almost every layer is standardized by the Internet Engineering Task Force (IETF), except the physical and media access control layers which are standardized by the IEEE in the 802.15.4 standard.

Table 2.1   Communication protocol stack.

Layer	Protocol	Standard
Transport	UDP	IETF RFC 768
Network	IPv6 / RPL	IETF RFC 6550
Adaptation	6LoWPAN	IETF RFC 6282
Data link	IEEE 802.15.4 MAC (CSMA)	IEEE 802.15.4
Radio Duty Cycling	ContikiMAC	-
Physical	IEEE 802.15.4 PHY	IEEE 802.15.4

RPL (IPv6 Routing Protocol for LLNs) [66] is a proactive routing protocol based on a Destination-Oriented Directed Acyclic Graph (DODAG), that uses a distance metric based on link quality indicators. The network nodes exchange ICMPv6 control packets to find and propagate the routes in the network, creating the DODAG. RPL uses three types of ICMPv6 control packets: DODAG Information Solicitation (DIS), DODAG Information Object (DIO), and Destination Advertisement Object (DAO). The former two messages are sent as link-layer broadcasts. DIS messages are requests to join a DODAG, whereas DIO messages contain information about the DODAG of the sender. The DIO messages inform the distance to the root node using the DODAG metric. Based on this information, each node selects its preferred parent that will route its packets to the root. The frequency of the DIO messages is controlled with the trickle algorithm [34], which is based on a timer with a duration that is doubled each time it is fired, sending fewer messages per unit of time when the network is stable.

6LoWPAN [38] is an adaptation layer protocol that allows the transport of IPv6 packets over 802.15.4 links. It is in charge of the compression of IPv6 and the upper layer headers and the fragmentation and reassembly of IPv6 packets.

ContikiMAC [12] is a radio duty cycling (RDC) protocol based on the Low Power Listening (LPL) mechanism, that uses periodic wake-ups to listen for packet

transmissions from neighbor nodes. This enables the radio transceiver to achieve duty cycles below 1% [32], making a sensible reduction of the energy consumption in comparison with the systems that do not manage the RDC layer. If during a periodic wake-up a transmission is detected, the radio transceiver is kept active in order to receive the packet. After the packet is successfully received, a link-layer acknowledgment is sent. For the transmission of packets, the phase lock mechanism is used. Every time a node sends a packet to a neighbor, it records the time at which the neighbor replied with an acknowledgment, learning its phase. Using the phase information, the sender does not start to send its strobes until the neighbor is awake. If the sender misses the awake period of the neighbor, the sender will simply continue to send strobes for an entire wake-up period. Broadcast transmissions do not receive link-layer acknowledgments and are intended to be received by every neighbor node, so the sender needs to repeatedly send the packet during a full wake-up period.

2.4 Evaluation of Power Consumption

For the evaluation of the power consumption, we use the Energest module that is implemented in Contiki OS. Energest is a software-based on-line energy estimation mechanism, that measures the accumulated time the node is in different states such as running CPU code (CPU), Low Power Mode (LPM), attending Interrupt Requests (IRQ), transmitting network packets (Tx) or receiving network packets (Rx). The mechanism runs directly on the nodes and provides real-time estimates of the current energy consumption.

LPM state is activated when the node goes to LPM but is not deactivated in the interrupt requests handling, so the actual time the node is in LPM is $LPM - IRQ$. CPU state is activated whenever the node is active, including radio transmissions, so the time the CPU is active without using the radio transceiver is $CPU - Tx - Rx$.

Taking this into account, we can calculate the average power as stated in Equation 2.1.

$$P_{avg}^{state} = DC_{state} \times I_{state} \times V = \frac{t_{state}}{t_{total}} \times I_{state} \times V \qquad (2.1)$$

where I_{state} is current consumption at a certain state and V is the operating voltage. The current consumption of each state is obtained from the TMote-Sky datasheet [36], and the operating voltage used in the rest of this work is $3V$. Note that for Tx and Rx states, CPU current consumption is taken into account.

2.5 COOJA Network Simulator

Cooja [42] is a Java-based network simulator capable of emulating several platforms with radio communication. In particular, it supports the TMote-Sky platform we use for our evaluations. It has a graphical user interface that allows the visualization of the network topology, the packets exchanged between the nodes, the

serial output, and a timeline with the radio events such as packet transmissions, receptions, and collisions. It can also run in a headless mode, with which we can automate and run simulations in several scenarios. By default, it only supports omnidirectional antennas with fixed radiation patterns.

Although some related work [55] uses the Castalia simulator [3] to evaluate the performance of the network protocols, this simulator has the major drawback of not modeling the hardware layers (e.g. radio hardware ACK). For this reason, and due to the ease of integration with Contiki OS, we choose to work with the COOJA network simulator.

Chapter 3

Electronically Switched Directional Antennas

In this chapter, we present a brief introduction to the directional antennas available for WSNs. We select the SPIDA antenna as a case of study due to the advantage of having a low cost and an easy fabrication process, as well as a very small size that makes it very convenient for large scale deployments. In Section 3.2 we show the fabrication process of a SPIDA antenna. We build a first prototype with a fixed direction that we fully characterize both in simulation and with field measurements with very good correspondence. We then show the construction of our second prototype with full functionality that we replicate in ten units. We continue by presenting a model to represent this kind of antennas in COOJA, one of the most popular network simulators for WSNs. In Section 3.3 we present novel ideas to improve the SPIDA antenna gain and directivity without changing its geometry, by using multiple director elements simultaneously. Finally in Section 3.4 we conclude the chapter and present the final remarks.

3.1 Introduction

One of the prominent ways to optimize an antenna is by using dynamic beam-forming employing an electronically switched directional antenna. This technique allows us to increase the antenna gain in some directions selected dynamically for each transmission [31, 57]. One example of this kind of antennas is the Electronically Steerable Passive Array Radiators (ESPAR) [30, 35, 52], which consist of a single feed element surrounded by a ring of reactively loaded parasitic elements. By electrically controlling the loading reactances, directional beams and nulls can be formed and steered throughout the azimuth of the antenna.

Another kind of dynamic beamforming antennas widely used in communications systems is the phase-shifting antenna, but their use of heavy signal processing techniques makes them inadequate for WSN nodes with constrained hardware. Nevertheless, Selavo et al. propose PHASER [53], a phase-shifting directional antenna prototype for WSNs. They use a WSN low-power radio chip, an RF sig-

nal processing chip, and two monopole antennas that are $\lambda/2$ apart. Directional communication is achieved by splitting the output signal from the radio chip and programmatically controlling the phase of each signal as it transmitted to each antenna. In this way, constructive and destructive interference patterns are generated. The main problem with these antennas is that they require custom hardware to manage the signal processing efficiently and cannot be used with regular sensor nodes.

The SPIDA (*Swedish Institute of Computer Science* Parasitic Interference Directional Antenna) [40, 43, 47] is another kind of dynamic beamforming antenna designed specifically for WSN nodes in the 2.4 GHz ISM band. It consists of a central active element surrounded by parasitic elements. The central element is a quarter-wavelength whip antenna and the parasitic elements can be switched between grounded and isolated: when grounded, they work as reflectors of the radiated power and when isolated, they act as directors of the radiated power. The SPIDA has six parasitic elements, so by using electronically controlled switches the shape and direction of the antenna's main lobe can be modified. The parasitic elements can be controlled individually, allowing multiple configurations. This antenna has the advantage of having a low cost and an easy fabrication process, as well as a very small size that makes it very convenient for large scale deployments. The dynamic beamforming features of the SPIDA antenna are a promising alternative to optimize WSNs that has motivated researchers to perform several works in this area [21, 29, 37, 60].

Comparing SPIDA with ESPAR antennas, the first are simpler and cheaper to fabricate, resulting in quicker and more affordable deployments. SPIDA antennas are also more suitable for WSN than phase-shifting antennas because they do not need any additional hardware and can be attached to any sensor node with six output pins available. For these reasons, we select the SPIDA antenna as our case-of-study to analyze, characterize, and optimize for WSN applications.

The main contributions of this chapter are: i) the fabrication and complete characterization of the reference design of the SPIDA antenna [40], including the radiation pattern and the S_{11} parameter, the last one missing in the bibliography; ii) a model for the COOJA network simulator that allows to use the SPIDA antenna in simulated WSN nodes, calculating the antenna gain based on the angle between the transmitting and receiving nodes; iii) the optimization of a SPIDA antenna by using multiple director elements. We present simulations and analyze the performance of the different configurations in terms of maximum gain in the main direction, HPBW, and S_{11} parameter. We also present real-world measurements of the three director elements configuration (identified as an equivalent to the reference configuration but with an improved performance) to validate the simulation results.

3.2 Original SPIDA Antenna

3.2.1 Design

The original antenna proposed by Nilsson [40] has six parasitic elements, thus the legs are separated 60° forming a hexagon. Figure 3.1 shows a model we designed in CST Studio Suite® following the description of the reference antenna. The overall size is such that it can be fitted in a cylinder with a radius of 52 mm and a height of 60 mm.
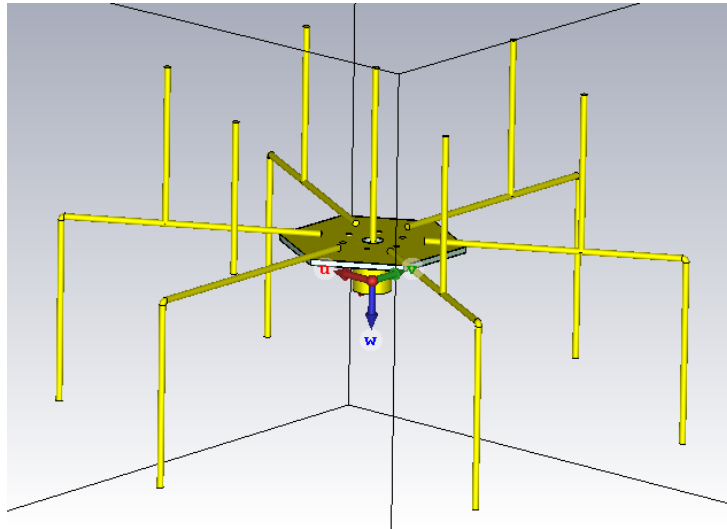

Figure 3.1   SPIDA Antenna Model designed in CST Studio Suite®.

The antenna geometry is shown in Figure 3.2. The antenna is based on a planar structure in a symmetrical arrangement, and a central vertical active element of 29.2 mm length and 1 mm width. Each of the six structures attached to the central hexagon is formed by a leg (that resembles the leg of a spider) with a vertical parasitic element. Each parasitic element is 27 mm length and 1 mm width, and can act either as a director, if it is left isolated (i.e. not connected), or as a reflector if it is connected to ground. Each connection can be controlled electronically by an RF-switch that allows a microcontroller to manage dynamically the configuration of the antenna. The gain of the SPIDA antenna is between 4 and 7 dBi in the principal direction with a HPBW of 130° [37, 40]. The HPBW is the angle for which the maximum gain decays 3 dB. However, the S_{11} parameter is missing in the corresponding reports and is very important to quantify the impedance matching and the power transfer. Thus, in this chapter a complete characterization is included, consisting of simulations and measurements.

3.2.2 Electromagnetic Simulation

The CST Studio Suite® simulator was used to assess the performance of the antenna. CST is a powerful electromagnetic simulator that allows us to obtain the
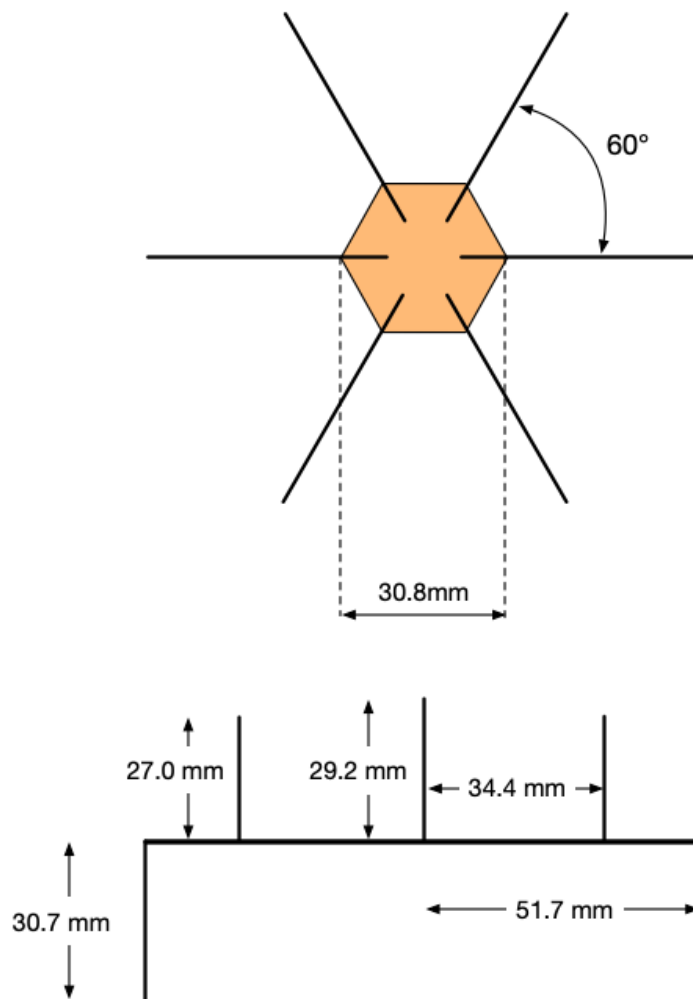
Figure 3.2   SPIDA Antenna geometry.

radiation pattern and the S_{11} parameter of the target antenna. The model used to simulate the antenna is shown in Figure 3.1 and was designed following the original design carefully (e.g., the vias used to connect both faces of the hexagon were also considered in the model).

Figure 3.3 presents the Smith Diagram obtained by simulation for frequencies between 2 GHz and 3 GHz. In this figure, the S_{11} parameter is represented, which allows us to observe how the input impedance varies in the range from 2 GHz to 3 GHz. At the frequency of 2.4525 GHz (the central frequency of the IEEE 802.15.4 band), the simulations show a module of S_{11} parameter of -3.7 dB. Figure 3.4 shows the radiation pattern obtained by simulation for the same central frequency. This figure shows that the main lobe of this antenna is in the direction of the director element, and the lowest gain is in the opposite direction. The shades of red indicate higher gains while the shades of green represent lower gains.
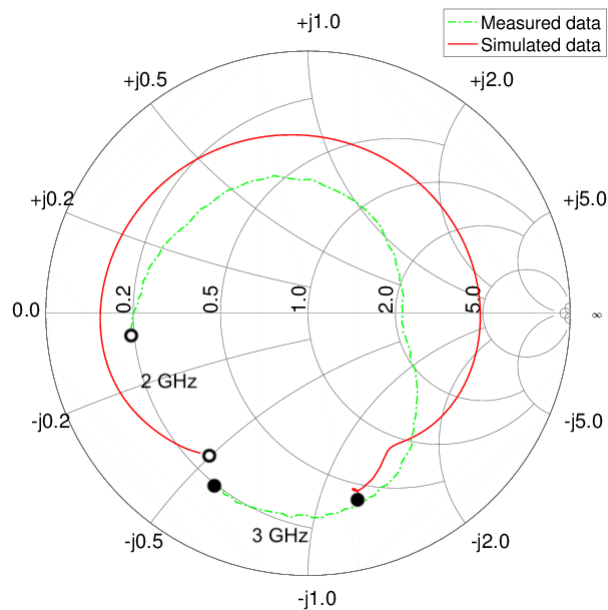
Figure 3.3   Simulated and measured Smith Diagram for the SPIDA antenna.

Figures 3.5a and 3.5b complement this information showing the gain in the horizontal and elevation planes. The elevation plane shows that the maximum gain is 20° above the direction of the leg of the antenna. The horizontal plane shows that a maximum gain of 6 dB is found in the direction of the director element, and the front-to-back ratio (FTBR) is 21 dB. This figure also shows a HPBW of 128°.
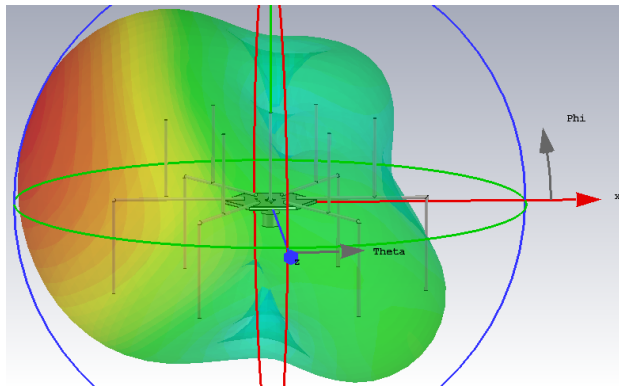

Figure 3.4   Simulated 3D Radiation Pattern for the SPIDA antenna.

3.2.3   Fabrication

To characterize the antennas with real experiments, we fabricated a SPIDA antenna following the original design. Simulations show that the antenna is symmetrical in every direction, so we choose to characterize the antenna in one direction and then extend the results.
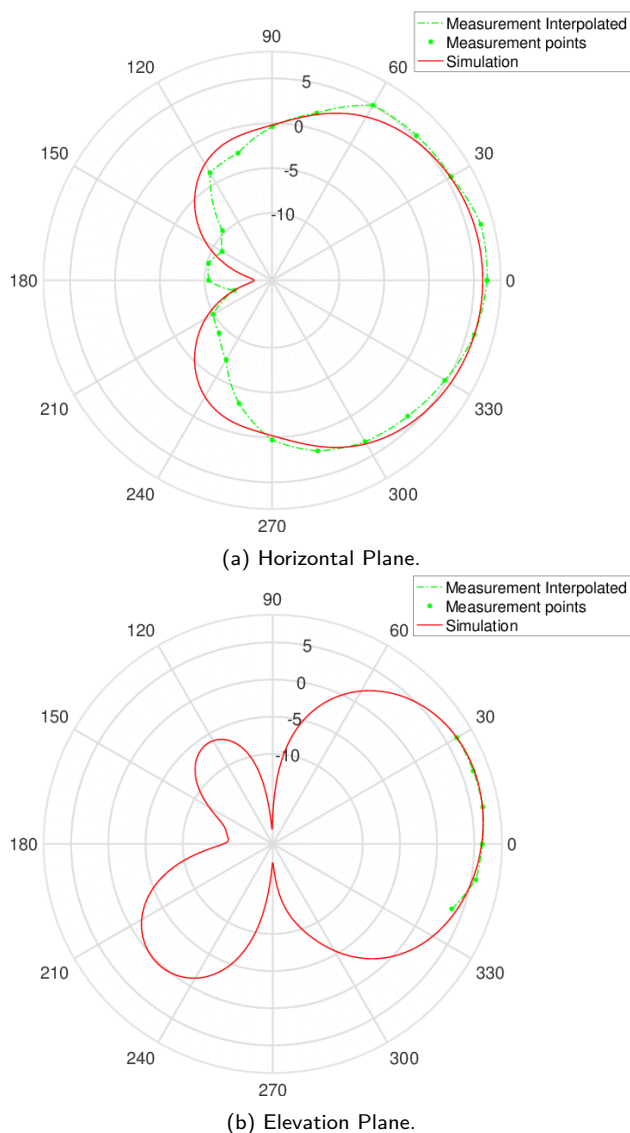
(a) Horizontal Plane.

(b) Elevation Plane.

Figure 3.5   Simulated and Measured Gain in H and E planes for the SPIDA antenna.

We built a first prototype with a fixed direction that is only used for characterization purposes, and a second one with full functionality that we replicate in ten units to test the protocols designed in the following chapters.

First prototype with one fixed director element

Our first prototype of the SPIDA antenna for electromagnetic characterization is shown in Figure 3.6. This antenna has a parasitic element defined as a director that is isolated with silicone and other five parasitic elements defined as reflectors that are welded to its corresponding legs to ground (which are connected to the hexagon and through it to the cable shield). This structure defines a static configuration where the radiated power is concentrated in a fixed direction.
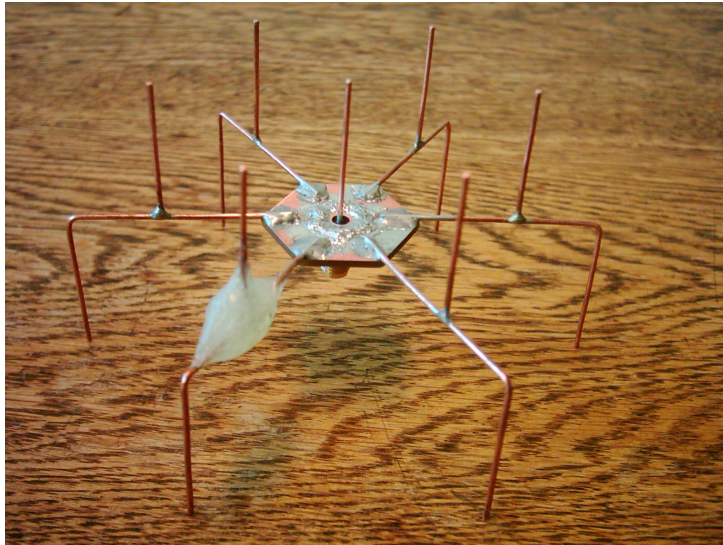
Figure 3.6 First prototype of SPIDA antenna with a fixed director element.

The elements are made of copper wire of 1 mm^2 section (the dielectric shield was removed). A central PCB hexagon was used to fix the legs and connect them to ground. This hexagon was made using a two-layer 1.6 mm FR4 PCB board of 35 μm of copper thickness. Both copper layers of the hexagon were connected using vias of 1 mm^2 of section. An SMA connector was welded to the lower copper layer of the hexagon to feed the antenna through it. The central and active element of the antenna was connected through the SMA connector to the central wire of the coaxial cable used to feed the antenna.

The hexagon was designed using CadSoft Eagle PCB Design Software and fabricated with LPKF ProtoMat S63 circuit board plotter. The circuit board plotter features a resolution of 0.5 μm and an accuracy of $\pm$ 0.02 mm, allowing a very precise fabrication. This equipment enables the production of identical hexagons for the fabrication of the antennas.

Second prototype with full functionality

Our second prototype extended the previous one by adding the RF switches to dynamically control the configuration of the parasitic elements using the schematics shown in Figure 3.7. We fabricated ten units as the one shown in Figure 3.8, and we used TMote-Sky nodes to control the direction of transmission. The cost of the materials for each antenna is below USD 10, including the RF switches, copper wire, and the SMA connector.

We used the same ADG902 [11] switches as in the original design [40] that provide a typical off-isolation of around 30 dB and a typical insertion loss of about 1.4 dB at 2.4 GHz. The supply and ground pins from the switches were connected to the corresponding supply and ground pins from the nodes. The control pin from each switch was connected to the different GPIO pins of the nodes, while the $RF1$ pin was connected to the parasitic element, and the $RF2$ pin was connected
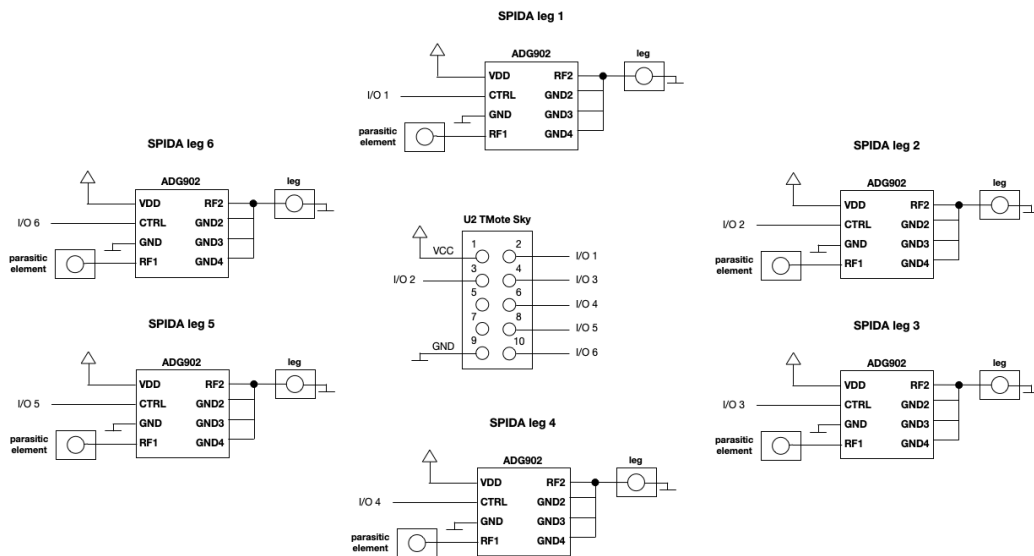
Figure 3.7   Schematics for the second prototype of SPIDA antenna with full functionality.

to ground. This way, the node can switch between isolation and grounding of each parasitic element by changing the voltage in the respective GPIO pins.
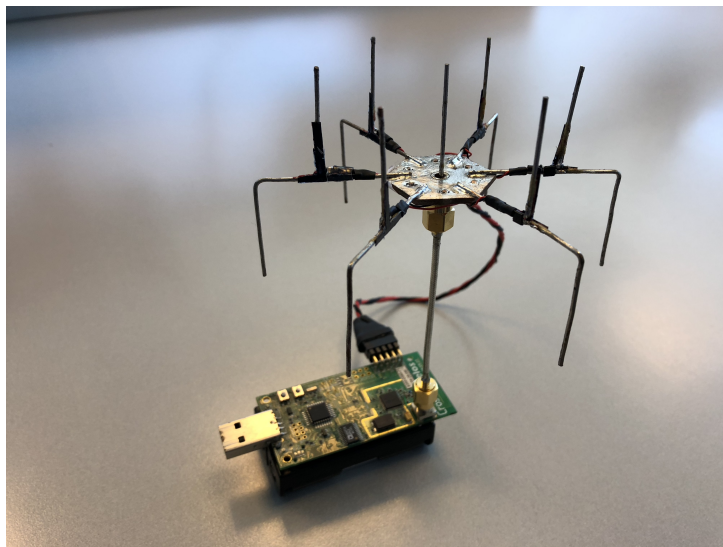

Figure 3.8   Second prototype of SPIDA antenna with full functionality.

3.2.4   Characterization

A vectorial network analyzer (Rohde & Schwarz ZVB 8 Vector Network Analyzer, 300 kHz - 8 GHz), an RF generator (Agilent, E4438C, 250 kHz - 3 GHz, ESG Vector Signal Generator) and a spectrum analyzer (Agilent Technologies, EXA Signal Analyzer, N9010 A, 9 kHz - 7 GHz) were used for the characterization process.

We fully characterized our first prototype described in Section 3.2.3 by measuring the S_{11} parameter and the radiation pattern in the H and E planes. The obtained measurements can be respectively seen in Figures 3.3 and 3.5. In these figures, the measurements and the simulations are superimposed showing good correspondence. The E plane was only measured in five points near 0°, as this is the area of interest for the WSN applications that we are considering, where all nodes are placed almost in the same horizontal plane.

A maximum gain of 6.8 dBi according to measurements and 6 dBi according to the simulations was obtained. This gain is higher than the other reported maximum gains we could find for similar antennas, that resulted in 5.1 dBi [52] and 4.3 dBi [40]. According to the measurements, the HPBW for this antenna is 113°(129° according to simulations), and the module of the S_{11} parameter is -7.5 dB (-3.73 dB according to the simulations). As the characterized prototype did not include the RF switches, it is expected that the maximum gain of the second prototype is 1.4 dB lower due to the attenuation of the switches.

These results quantitatively show that these antennas are appropriate for the intended application, opening the path to our further work.

3.2.5 Modeling for COOJA Simulator

Existing WSN simulators lack support for ESD antennas such as SPIDA. ESD antennas are also difficult to test in real testbeds as they require time-consuming manual set-up, and have to be manually manufactured as they are not commercially available. This difficulty to test ESD antennas in large networks hinders the development of WSN protocols that leverage directional communication. By default, the COOJA simulator assumes an omnidirectional antenna behavior. We introduce support for the SPIDA antenna in the COOJA simulator for the experiments conducted in our work. We also describe the antenna model we introduce in the COOJA simulator. Instead of using a fixed antenna range, we modify the COOJA simulator to use the radiation pattern of the SPIDA antenna, calculating the antenna gain based on the angle between the transmitting and receiving nodes.

The total received signal strength (RSS) for a transmission between two nodes A and B is calculated using the log-normal path loss model [23]:

$$RSS = P_{tx} + P_{PL}(d_0) - 10 \cdot K \cdot \log_{10}\left(\frac{d_{AB}}{d_0}\right) + G_{sp}(\theta_{AB}) + G_{sp}(\theta_{BA}) \quad (3.1)$$

where P_{tx} is the transmission power, $P_{PL}(d_0)$ is the path loss at reference distance d_0, K is the path loss exponent, d_{AB} is the distance between nodes A and B, $\theta_{AB} \in [0, 2\pi)$ is the angle between the active antenna direction of nodes A and B and finally $\theta_{BA} \in [0, 2\pi)$ is the angle between the active antenna direction of nodes B and A. Figure 3.9 depicts these angles for a pair of nodes A, B.

For successful transmissions, the RSS has to be above the receiver's sensitivity. For the CC2420 radio used by the TMote-Sky nodes, this value is -90 dBm for the worst case [25]. Based on typical outdoor radio environments, the reference distance d_0 is set to two meters, the reference path loss P_{PL} to -52 dBm, and
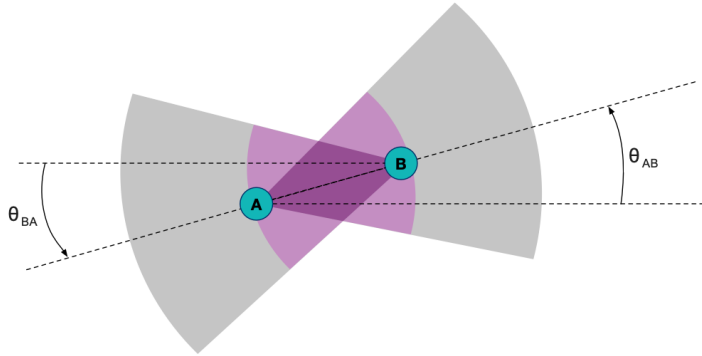
Figure 3.9   Visual representation of the antenna model in COOJA. θ_{AB} is the angle between the active antenna direction of nodes A and B and θ_{BA} is the angle between the active antenna direction of nodes B and A.

the path loss exponent K to 2.5. We also implement the capture effect in the simulator, where a packet can be received correctly despite interference, if its RSS is at least 3 dB stronger than the sum of the received signals from all other nodes, and if the strongest signal arrives within 160 μs of the first weaker signal [33, 59]. This time corresponds to the air time of the IEEE 802.15.4 time synchronization header.

To obtain the antenna gain $G_{sp}(\theta_{AB})$ as a function of the angle θ_{AB} between the active antenna direction of the transmitter and receiver nodes, we use the data from the simulations presented in Section 3.2.4.

We approximate $G_{sp}(\theta)$ as a sum of sines of sixth-degree fitted to the simulation data with an R-squared value of 0.9995. Figure 3.10 shows this fitted gain compared to the reference internal omnidirectional antenna of the TMote-Sky [45]. The SPIDA antenna has a 3 dB higher gain at 0º increasing the transmission range, and 12 dB lower gain at 180º, producing much less interference on neighboring nodes.

We can easily replace the radiation pattern of the SPIDA antenna for the one of any other antenna by changing the formula of the antenna gain. Particularly, we can use the simulation data of any of the configurations that will be presented in Section 3.3 to model them in COOJA.

Because of how the COOJA simulator works, all the simulated nodes will have the same orientation, although this is only relevant for the Q-SAND protocol that will be presented in Section 4.3.

3.3   SPIDA Antenna Optimization

In Section 3.1 we present related work showing the use of the SPIDA antenna for WSNs. While it has been used for several applications [21, 29, 37, 60], every work is based on a six-element SPIDA antenna with only one as a director. In this section, we propose and study the use of multiple parasitic elements as directors, as a way of improving the antenna without changing its geometry. We show we
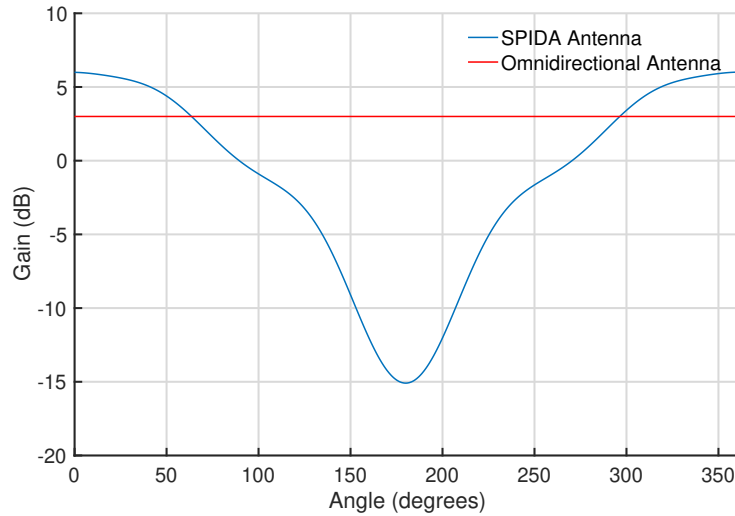
Figure 3.10 SPIDA antenna gain compared to reference omnidirectional antenna. SPIDA antenna has 3 dB higher gain at 0° and thus increases the transmission range, and 12 dB lower gain at 180° producing a much lesser interference with neighboring nodes.

can improve the performance in terms of maximum gain, narrower HPBW, and a lower module of the S_{11} parameter without making any changes to the antenna itself.

3.3.1 Design

In this section, different configurations of director elements are considered. We study how the adoption of multiple director elements affects the complete shape of the radiation pattern, the input impedance matching, FTBR, and gain in the different directions.

Table 3.1 lists the eight different configurations considered, and Figure 3.11 depicts them graphically, where the direction of maximum gain is aligned with the horizontal axis (0°) for experiments 1, 2, 3, 4 and 5. For experiments 6 and 7, the direction of maximum gain is also aligned with the horizontal axis but equally high for 0°and 180°, and experiment 8 is omnidirectional. All these configurations were simulated to assess their performance. Section 3.3.2 describes the simulation results in detail. Among them, the most promising configuration (using three director elements) was measured and analyzed more deeply in Section 3.3.4.

3.3.2 Electromagnetic Simulation

We modify the antenna model presented in Section 3.2.2 to isolate or ground the different parasitic elements, to obtain the configurations from Table 3.1. We then use CST Studio Suite® to obtain the radiation pattern and the S_{11} parameter.
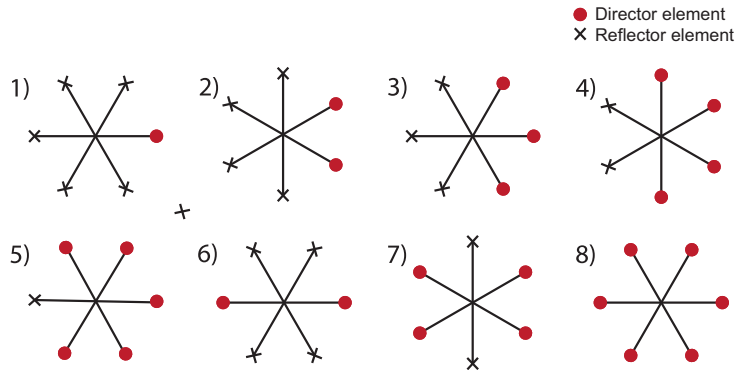
Figure 3.11 Eight different experiments, showing the location of the director and reflector elements.

Table 3.1 List of eight different configurations simulated with the location of the director element(s).

| #	Description	Location
1	One director	0°
2	Two consecutive directors	30° and 330°
3	Three consecutive directors	0°, 60° and 300°
4	Four consecutive directors	30°, 90°, 270° and 330°
5	Five consecutive directors	0°, 60°, 120°, 240° and 300°
6	Two opposed directors	0° and 180°
7	Four opposed directors	30°, 150°, 210° and 330°
8	Six directors (all)	-

One director (conf. #1)

This configuration, in which a single element acts as director, corresponds to the original configuration previously reported in Section 3.2.2 and is the reference design for comparison.

Two consecutive directors (conf. #2)

This configuration presents a radiation pattern with good directivity as shown in Figure 3.12a. The HPBW is 87°, significantly narrower than the 129° from the reference design. The maximum gain is 7.70 dBi, 1.72 dB higher than the configuration with one director but the FTBR is 9.00 dB, 12.0 dB lower than the reference design. Figure 3.12b shows the S_{11} parameter over the Smith Diagram when it varies from 2 GHz to 3 GHz. The S_{11} parameter in the central frequency 2.4525 GHz is -7.73 dB, about 4.00 dB lower than the reference design. According to these simulation results, this configuration outperforms the original one in maximum gain, presenting a narrower beamwidth which could be favorable in many scenarios, and having a better input impedance matching.

Another interesting characteristic of this configuration is that the main lobe
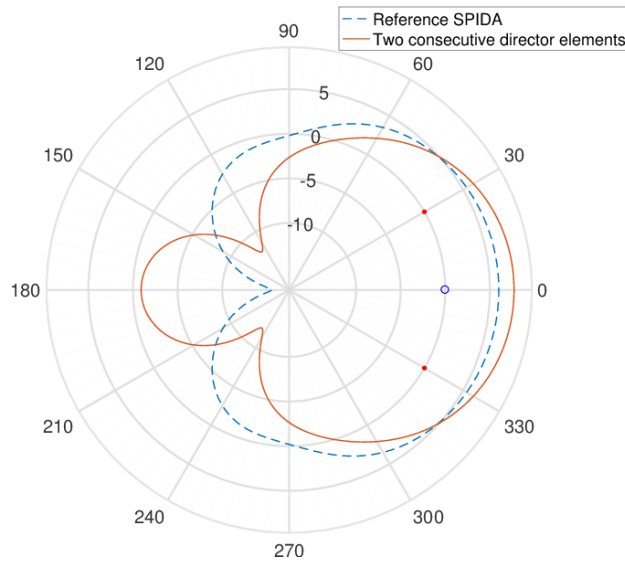
(a) Polar radiation pattern in H plane.
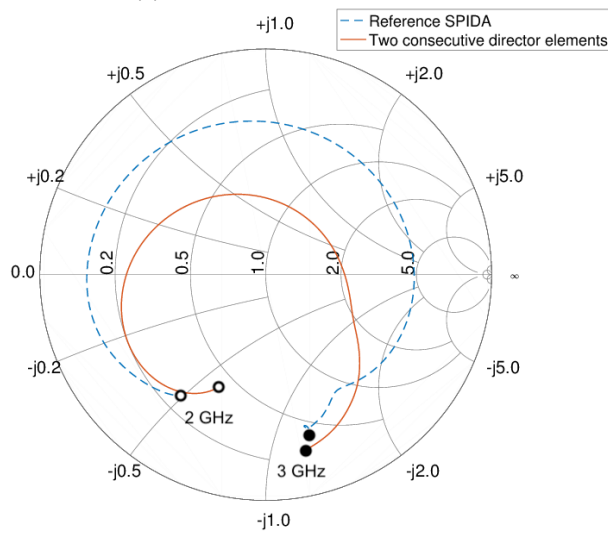
(b) S_{11} parameter.

Figure 3.12   Two consecutive director elements (conf. #2).

direction is in the middle of the two directors, so with a six-element antenna and using two consecutive directors it would be possible to direct the main beam in 12 different directions.

Three consecutive directors (conf. #3)

The radiation pattern in the H plane of this configuration is shown in Figure 3.13a. It can be observed that this configuration has an even better directivity than for one and two consecutive directors, with a HPBW of 76° compared to the 129° of the original design. The maximum gain is also better than the corresponding one for one and two directors, achieving 8.35 dBi, 2.37 dB higher than the reference

design. This configuration presents a back lobe resulting in an FTBR of 13 dB, lower than in the reference design but better than in the configuration with two directors. Figure 3.13b shows the radiation pattern in E plane; for the sake of brevity E plane is shown for this configuration only, since the other configurations have very similar characteristics.

The S_{11} parameter is shown in Figure 3.13c, and the module at 2.4525 GHz is -11.25 dB , outperforming the previous configurations. The main lobe direction is aligned with the center director element, allowing to transmit directionally in any of the original six directions, but with these improved characteristics.

So far, these simulation results show that by using two and three director elements, we can focus the main beam to 12 different directions with better performance in terms of gain, directivity, and impedance matching than the original configuration using only one director. An aspect to consider in these configurations is that the FTBR is lower than the reference design, resulting in the radiation of more energy in the opposite direction.

Four and five consecutive directors (confs. #4 and #5)

The remaining configurations using consecutive directors, that is four and five (configuration #4 and #5 respectively), do not show improvements -in terms of maximum gain- over the two previously analyzed configurations, so these results are not plotted for the sake of brevity.

Two and four opposed directors (confs. #6 and #7)

Configurations using opposing directors can be used to radiate power simultaneously in two opposite directions, to minimize the broadcast transmissions in linear deployments. The configuration #7 uses four opposed director elements and the resulting radiation pattern in the H plane is shown in Figure 3.14a. We observe that the power is radiated in two opposite directions, each one with a gain of 7.40 dBi (1.42 dB higher than the reference design). The S_{11} parameter is shown in Figure 3.14b presenting a value of -12.2 dB at 2.4525 GHz, about 8.47 dB lower than the reference design. The HPBW is 58°, being much narrower than the reference design. The results obtained for two opposed directors (conf. #6) can be observed in Table 3.2 and do not present any improvement regarding conf. #7.

Six directors (conf. #8)

The configuration with six director elements is a special case, as we obtain an almost omnidirectional pattern with a gain of 3.9 dBi (1.8 dBd), which shows an improvement in the radiation characteristics with respect to a dipole, being in this way a very good option as an omnidirectional antenna. The use of a combination of directional antennas together with omnidirectional antennas has been proposed in some communication protocols [46, 49, 68], where the omnidirectional antenna is adopted for broadcast messages.
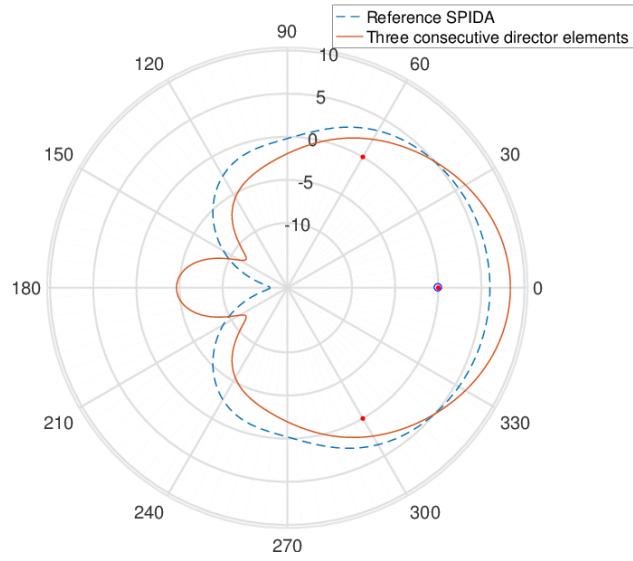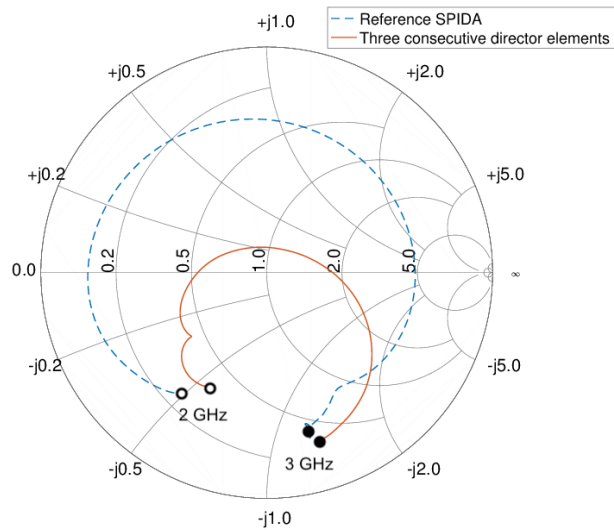
(a) Polar radiation pattern in H plane.
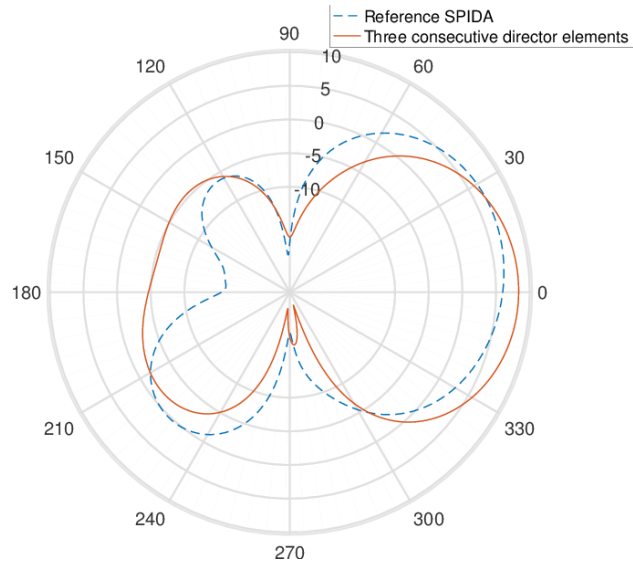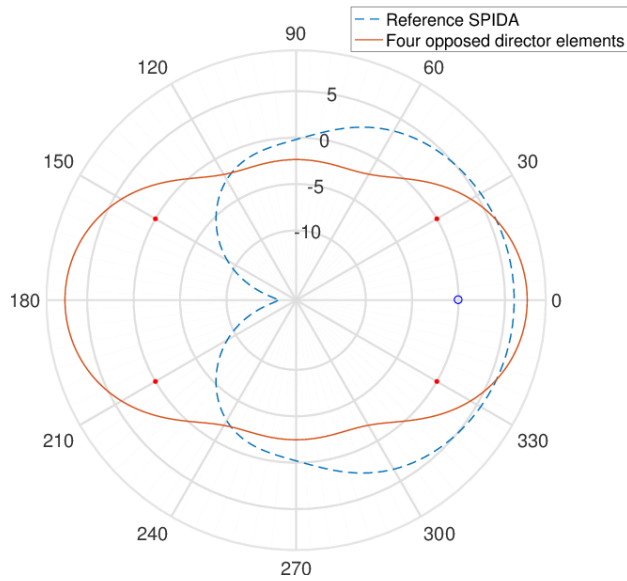

(b) Polar radiation pattern in E plane.

(c) S_{11} parameter.

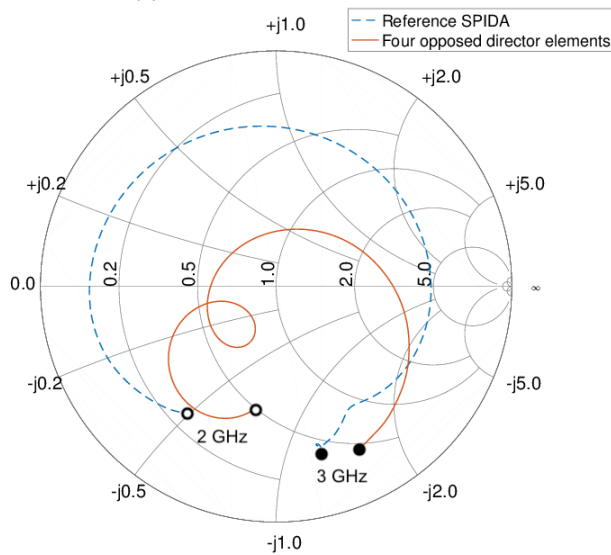Figure 3.13   Three consecutive director elements (conf. #3).

(a) Polar radiation pattern in H plane.

(b) S_{11} parameter.

Figure 3.14   Four opposed director elements (conf. #7).

General Analysis

Table 3.2 summarizes the main performance parameters of the simulations for the different configurations of the SPIDA antenna. All configurations using consecutive director elements outperform the reference design in terms of input impedance matching. The configurations with two, three, and five consecutive directors (confs. #2, #3, and #5) present narrower beamwidth, while the configuration with four directors (conf. #4) has a wider beamwidth. In terms of maximum gain, confs. #2, #3, and #4 outperform the reference design. Considering the FTBR, the reference antenna performs similar to the one with four consecutive directors but

better than the configurations with two and three consecutive director elements.

Considering the configurations with opposed directors, confs. #6 and #7 outperform the reference design in terms of input impedance matching and beam directivity in the desired directions, but only conf. #7 has a higher maximum gain than the reference design. These configurations have the potential to direct RF power to opposite directions with the benefit of enhanced performance, which would be beneficial in many applications with linear topologies.

Simulations also show that this antenna can be used as an omnidirectional one, with a better gain than a dipole.

The choice of the best configuration for each application would depend on network design aspects, such as medium access protocols requirements and sensor nodes arrangement. This is why it is so important to be able to change these configurations dynamically without changing the geometry of the antenna. Each deployment could have a configuration phase where the nodes choose the best antenna configurations to communicate with each other, or nodes could also change the configuration of the antenna dynamically depending on some input parameters (position of the neighboring nodes, interference, packet delivery rate, etc.). Moreover, the increased RF power delivered to the air due to better impedance matching (achieved in all the cases) represents an important improvement to the energy efficiency of the system.

Measured and simulated values of the S_{11} parameter for the most interesting cases are shown in Figures 3.12b, 3.13c, 3.14b and 3.16. Considering a Voltage Standing Wave Ratio lower than two (module of S_{11} in dB lower than -9.54 dB), the configuration with three director elements presents an impedance bandwidth of 310 MHz, from 2180 MHz to 2490 MHz. For the band used in the IEEE 802.15.4 standard (from 2400 MHz to 2483.5 MHz), the module of S_{11} in dB for this antenna configuration is always lower than -9.7 dB, satisfying the selected criterion. During this research, it was observed that the gain, directivity and impedance bandwidth of the reference antenna could be modified (and improved) by changing its geometry, in particular changing the distance between the parasitic and the active elements. But it is important to have in mind that the number of director elements in use in this antenna is selected dynamically, and any geometrical optimization of this antenna has a different impact in each of the possible configurations.

Table 3.2   Simulation results.

#	Description	Half Power Beamwidth (3 dB)	Maximum Gain	$|S_{11}|$ at 2.4525 GHz	FTBR
1	One director	129°	5.98 dBi	-3.73 dB	21 dB
2	Two consecutive directors	87°	7.70 dBi	-7.73 dB	9 dB
3	Three consecutive directors	76°	8.35 dBi	-11.25 dB	13 dB
4	Four consecutive directors	137°	6.42 dBi	-11.94 dB	25 dB
5	Five consecutive directors	103°	5.65 dBi	-8.39 dB	-
6	Two opposed directors	70°	5.91 dBi	-8.70 dB	-
7	Four opposed directors	58°	7.40 dBi	-12.20 dB	-
8	Six directors (all)	Omni	3.90 dBi	-5.81 dB	-

3.3.3 Fabrication

We used the same procedure and materials described in Section 3.2.3 to build a new prototype with three director elements, shown in Figure 3.15. In this prototype, three parasitic elements are defined as directors (glued with silicone), and the other three parasitic elements are defined as reflectors (connected to ground).
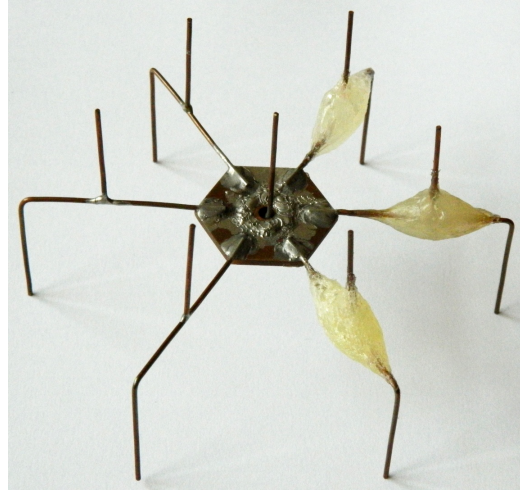

Figure 3.15   SPIDA antenna with three fixed consecutive director elements.

3.3.4 Characterization

For the characterization, we used the prototype previously described in Section 3.3.3 and the same procedure described in Section 3.2.4. We measured the S_{11} parameter and the radiation pattern in the H and E planes for the case with three director elements. The obtained measurements can be respectively seen in Figures 3.16 and 3.17. In these figures, the measurements and the simulations are superimposed showing good correspondence. For the same reason as in the original prototype, the E plane was only measured in five points near 0°.

According to the measurements for the SPIDA antenna with three director elements, a maximum gain of 8.20 dBi is obtained (8.35 dBi according to the simulations), 1.40 dB higher than the gain for the single director antenna taken as reference.

The measured HPBW for this antenna is 59° (76° according to the simulations) against 113° for the reference antenna (129° according to simulations). The measured module of the S_{11} parameter is -9.8 dB (-11.25 dB according to the simulations) against -7.5 dB for the reference antenna (-3.73 dB according to the simulations). All these results show a very important improvement compared with the single director SPIDA antenna, which would justify the use of three director elements instead of only one for almost every application.
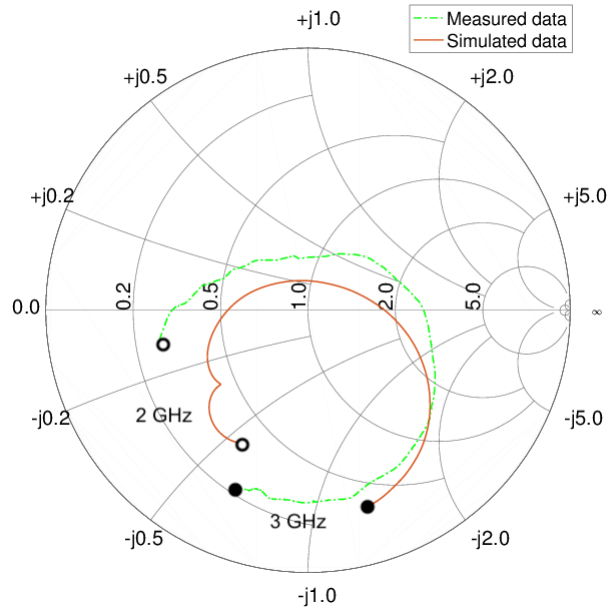
Figure 3.16 S_{11} parameter for the fabricated SPIDA antenna with three consecutive director elements.

3.4 Conclusions

A 6 element SPIDA antenna for the 2.4 GHz ISM band is modeled, fabricated, and fully characterized in Section 3.2. As far as we know, it is the first complete characterization reported for this kind of antenna including the S_{11} parameter. We present a simple and cheap fabrication process that leads to a directional antenna with very good performance. The measurement and simulation results are in good agreement. The obtained performance is comparable and even better than the best reported for similar antennas. The antenna shows a maximum gain of 6.8 dBi, an HPBW of 113° and a module of S_{11} parameter of -7.5 dB at the central frequency ($f_c = 2.4525$ GHz). We present a model for the COOJA network simulator that allows us to use the SPIDA antenna in simulated WSN nodes, calculating the antenna gain based on the angle between the transmitting and receiving nodes.

We show how to optimize the antenna without changing its geometry by using multiple director elements in Section 3.3. To the best of our knowledge, there is not any published report proposing the use of multiple director elements for this kind of antennas, which is very convenient because it allows to improve the performance in terms of maximum gain, narrower HPBW, and a lower module of the S_{11} parameter without making any changes to the antenna itself.

All configurations using consecutive director elements outperform the reference design in terms of input impedance matching. In terms of maximum gain, confs. #2, #3, and #4 outperform the reference design. Considering the FTBR, the reference antenna performs similar to the one with four consecutive directors but better than the configurations with two and three consecutive director elements.

For the configuration with three director elements, an improved radiation pat-
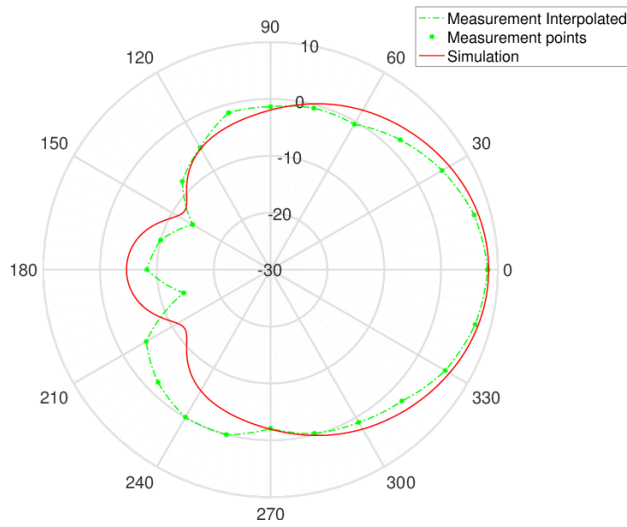
(a) Polar radiation pattern in H plane.
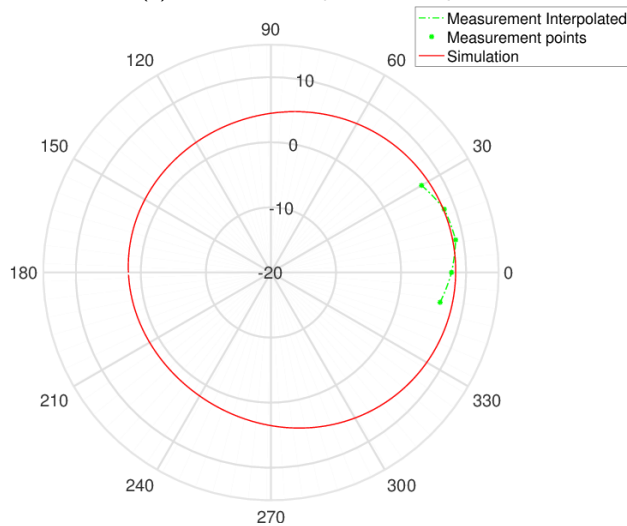
(b) Polar radiation pattern in E plane.

Figure 3.17   Fabricated SPIDA antenna with three consecutive director elements.

tern and input impedance matching is obtained both in simulation and measurements. It presents a maximum gain of 8.2 dBi (against 6.8 dBi for the original design), and a module of the S_{11} parameter of -9.8 dB at the central frequency (against -7.5 dB for the original design). In the configuration with two consecutive director elements, the main lobe direction is in the middle of the two directors, so with a six-element antenna it would be possible to direct the main beam in 12 different directions. The configurations with opposed director elements show the potential to direct RF power to opposite directions with the benefit of enhanced performance. Finally, the simulations also show that this antenna can be used as an omnidirectional one, with a better gain than a dipole.

By considering multiple director elements the flexibility in the beam orientation can be duplicated, having twelve beam directions instead of only six. The

possibility to change the configuration of the antenna dynamically makes the optimizations suitable for a wide variety of applications and network topologies.

In this chapter, we have deeply analyzed the SPIDA antenna, from its modeling to its fabrication, characterization, and optimization. In the following chapters, we will study how to use this antenna for WSN applications. In particular, the COOJA model described in Section 3.2.5 and the fabricated prototypes described in Section 3.2.3 are used to test different protocols in simulation and field measurements.

This page was intentionally left blank.

Chapter 4

Neighbor Discovery in Networks with Directional Antennas

In this chapter, we analyze neighbor discovery mechanisms for WSNs that use directional antennas. We start with a brief introduction to the challenges of finding neighbor nodes when using directional antennas, and present related work with several protocols that use different approaches to achieve this. We deepen the analysis in the SAND protocol [18, 39] in Section 4.2, as it is a state-of-the-art fully directional neighbor discovery protocol.

In Sections 4.3 and 4.4 we propose, implement end evaluate Q-SAND and DANDi, two novel fully directional protocols for neighbor discovery. Q-SAND is an improved version of SAND that is capable of finding the combination of sectors with the highest RSSI between every pair of neighbor nodes in a network up to four times faster than SAND. DANDi is a novel fully directional asynchronous and dynamic neighbor discovery protocol based on collision detection, that does not require any prior information of the network topology. Simulations and experimental results show that it can perform up to 2.6 times better than SAND in terms of neighbor discovery time, while finding every reliable link.

Finally in Section 4.5 we conclude the chapter and present the final remarks.

4.1   Introduction

One of the first problems that arise when using directional antennas is how to discover all the possible communication links with neighbor nodes. When using omnidirectional antennas, one single broadcast message is enough to query neighbor nodes, and the ones in range may reply including their network address. But with ESD antennas, there may be (and generally there are) more than one sector-to-sector (S2S) links between two nodes. Besides, the transmitter must know when its neighbor is pointing in a certain direction.

The main goal of a neighbor discovery protocol is to enable every node in a network to gather information about its respective neighboring nodes. If we consider nodes with sectored antennas, the essential purpose is to ensure the discovery

of all reliable S2S links between the nodes in the network. Considering K-sectored antennas, there are K^2 possible sector combinations between two neighbor nodes: K possible sectors of one node and K possible sectors of the other.

Many algorithms have been proposed to perform neighbor discovery when using directional antennas. Some neighbor discovery algorithms are based on a probabilistic approach. One of the simplest is to divide the time in timeslots and to send with probability p_t and to receive with probability $1 - p_t$. Vasudevan et al. [61] deduce the probability p_t that maximizes the number of nodes discovered in time. This optimal probability depends on the beamwidth of the directional antennas and the number of nodes in the network. The disadvantage of a fully probabilistic approach is that it will take much more time to discover the last links.

Other neighbor discovery algorithms rely on omnidirectional antennas to assist the neighbor discovery process [62], or assume that only the sink node is equipped with a directional antenna [54]. The most common approach is to use directional transmissions and omnidirectional receptions, eliminating the need of knowing in which direction the neighbor's antenna is pointing. The main drawback of this approach is that neighbors that would only be reached with both directional transmissions and receptions would not be discovered, and that the spatial reuse decreases.

Other algorithms [17,67] rely on previous time synchronization among nodes in the network. This time synchronization is hard to achieve before establishing the network. Xiong et al. [67] propose SBA, an algorithm where each node transmits with probability p_t to the sector S_i and receives with probability $1 - p_t$ to the opposite of sector S_i. The synchronization is needed for every node in the network to know where to point at every moment.

There are fully directional (directional transmissions and receptions) algorithms that do not require previous time synchronization among the nodes [5, 18,39,63]. The first two algorithms [5,63] use scan sequences based on each node's unique ID, guaranteeing that all the sector combinations among all nodes are tested. In both works, the shortcoming is that the collisions are not considered, so the discovery is not guaranteed for dense networks, or even for networks with dense areas. In the SAND (Sectored-Antenna Neighbor Discovery) protocol [18] (extended in [39]), a serialized mechanism where one node is discovering its neighbors at a time guarantees that all the sector combinations among all the nodes are tested. In all the algorithms mentioned above, the respective parameters are chosen using the number of nodes in the network or the estimated number of neighbors using the network density and the coverage area of the antenna. So some knowledge of the network topology is needed in advance.

These drawbacks motivated us to create DANDi, a novel fully directional asynchronous and dynamic neighbor discovery protocol based on collision detection, that does not require any prior information of the network topology. We also propose Q-SAND, an optimization of the SAND protocol that reduces the discovery time by only finding the strongest combination of sectors between a pair of neighbor nodes (the ones with the highest RSSI).

The main contributions of this chapter are: i) a detailed analysis of the state of

the art in protocols for Neighbor Discovery in networks with directional antennas and ii) the proposal, implementation, and evaluation of two novel fully directional neighbor discovery protocol for wireless networks with sectored antennas: Q-SAND and DANDi.

We evaluate Q-SAND only in simulation because it is an optimization of the SAND protocol, and our main objective is to compare its performance with SAND in terms of neighbor discovery time. To evaluate DANDi we make both simulations and experiments with real nodes to evaluate every aspect of the protocol, including discovery time, discovered links, and robustness to real-world conditions like interference and lossy links. We also evaluate SAND with field experiments in the same conditions as DANDi to have a reference to compare to.

4.2 SAND

This section describes SAND, a neighbor discovery protocol for sensor nodes with sectored-antennas initially proposed by Felemban et al. [18], and then extended by Murawski et al. [39]. This protocol is based on the idea of serializing the neighbor discovery process sequentially in time. This means that there is always a single node discovering neighbors at a time. To achieve this, a token is used to identify the node that is discovering neighbors.

4.2.1 Protocol Description

The main idea of SAND is to serialize the neighbor discovery process among all nodes. This serialization is based on a token-passing approach, where only the node that holds the token (Token-Holder or TH) is allowed to discover neighbors. Once a node has the token, it holds the attention of its neighbors using the Hone-In mechanism, discovers its neighbors using the Hello-Reply mechanism, and passes the token using the Token-Passing and Token-Releasing mechanisms.

Summarizing, the protocol is formed by the three steps introduced before:

1. Hone-In Mechanism

2. Hello-Reply Mechanism

3. Token-Passing Mechanism or Token Releasing Mechanism

Initialization

All nodes except the first TH node start in Fast Scan mode. In this mode, the nodes scan for activity in the network, switching its active sector every t_{switch}. It is important to note that initially the nodes are not synchronized.

Hone-In Mechanism

As mentioned before, the Hone-In Mechanism is necessary to synchronize the TH node with all its neighbors. The TH sends periodically every $t_{Hone\text{-}In}$ a number h of Hone-In messages per each sector.

The message carries the information of the number of messages remaining (including the current message) before starting the next step: the Hello-Reply Mechanism. So the first Hone-In message will contain hK, the second $hK - 1$, until the last one, that will contain 1 (in total, hK Hone-In messages will be sent). This mechanism is shown in Figure 4.1 for a 4-sectored antenna and 3 nodes.

In turn, when receiving a Hone-In message containing the number m, the neighbor nodes (that are in Fast Scan mode) will set a timer that will expire in $mt_{Hone\text{-}In}$ and cease switching sectors. In this way, a local synchronization between the TH node and its neighbors is achieved, as the timer of all its neighbors will expire (ideally) at the same time. The total duration of the Hone-In Mechanism is $T_{HI} = hKt_{Hone\text{-}In}$.

In order to guarantee that all the neighbor nodes have the opportunity to receive messages in all its sectors, the switching time t_{switch} must be greater than $t_{Hone\text{-}In}$.

To minimize the duration of the Hone-In Mechanism, the number of messages h must be also minimized. Nevertheless, h must be large enough to ensure that at least 1 message could be received for each one of the K^2 sector combinations, so h must satisfy $h > K \frac{t_{switch}}{t_{Hone\text{-}In}}$.

Hello-Reply Mechanism

Once the TH node has sent the hK Hone-In messages, the Hello-Reply Mechanism starts. On the side of the neighbor nodes, this mechanism starts when the timer set during the Hone-In Mechanism expires.

The idea is to explore the K^2 sector combinations in a predetermined way. To accomplish this, the TH node sends Hello messages, to which the neighbor nodes that receive it respond with a Reply message containing its ID, its current sector and a bit indicating whether it has already done neighbor discovery or not (this is for the TH node to determine to which neighbor to pass the token after finishing the neighbor discovery process). When receiving a Reply message, the TH node adds to its neighbor discovery table the corresponding neighbor with its ID, current sector (both TH and neighbor current sectors), the neighbor discovery bit mentioned above and a link quality indicator, such as the RSSI.

Since many neighbors can respond to the same Hello message, there can be a collision between the Reply messages, causing those neighbors not to be discovered for that sector combination. To reduce the number of collisions, the time after a Hello message is divided into N_{slots} time slots of duration t_{slots}. The neighbor nodes choose a time slot randomly to send their Reply message. To reduce the number of collisions even more, the TH repeats the Hello message N_{rounds} times per each sector combination, and adds to each Hello message the ID of every node that has already been discovered for the current sector combination (this type
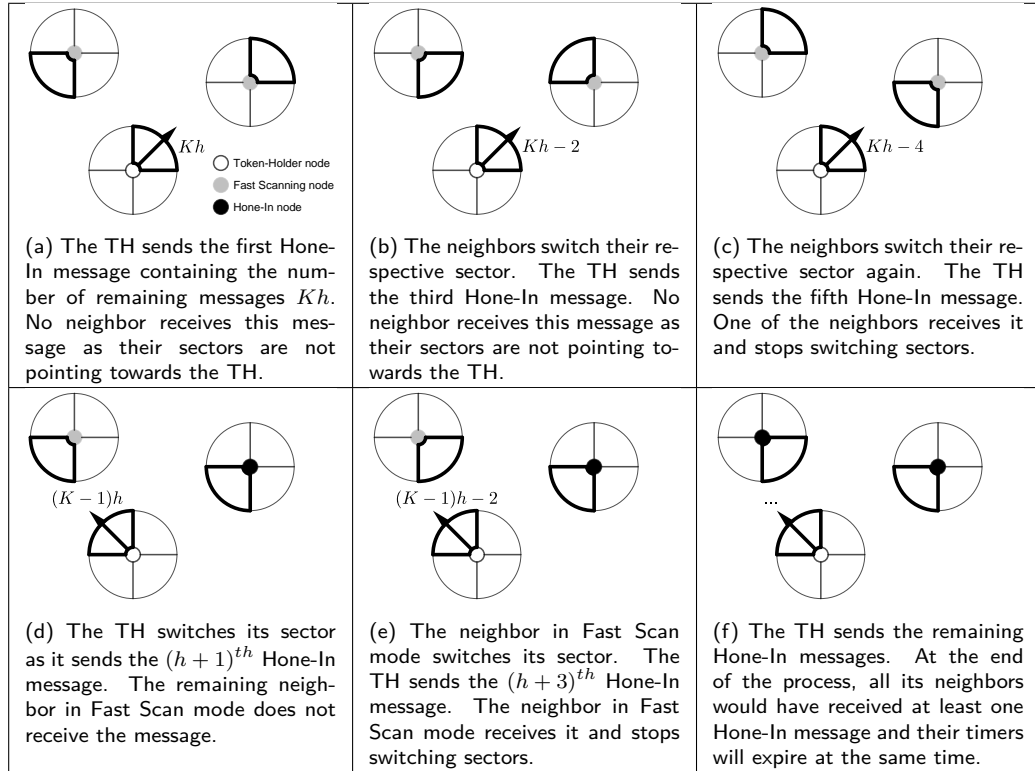
Figure 4.1 Hone-In process for $K = 4$ and $t_{switch} = 2t_{Hone-In}$. In this case, h must be greater than $K\frac{t_{switch}}{t_{Hone-In}} = 4 \times 2 = 8$. At the beginning of the mechanism, all the neighbor nodes are in Fast Scan mode. At the end of the mechanism, all the neighbor nodes have received at least one Hone-In message from the TH and their respective active sectors are pointing towards it.

of message is called Long-Hello message). The nodes that recognize their ID in the Long-Hello message know they have already been discovered, so they do not respond with a Reply message.

At the end of the mechanism, the TH node will proceed to pass the token. The neighbor nodes will activate the sector in which the Hone-In message was initially received. The total duration of the Hello-Reply Mechanism is $T_{HR} = K^2 N_{rounds} N_{slots} t_{slots}$.

Token Passing and Token Releasing Mechanisms

Once the Hello-Reply Mechanism is completed, the TH node goes over its neighbor table, searching for the first node that has not performed neighbor discovery yet (this is checked with the bit mentioned in Section 4.2.1). If such a node is found, the TH node proceeds to pass the token to it. Otherwise, the TH node proceeds to pass the token to the node from which it originally received the token.

The token passage between nodes then forms a tree, whose root is the node that started the neighbor discovery process (the first TH node). To pass the token, two different cases need to be distinguished.

In case the TH node has just finished discovering neighbors, all its neighbors will point to a specific sector, directing their antenna to the TH node. The TH must send GoToFastScan messages telling its neighbors to start Fast Scan mode (otherwise the neighbor nodes would stay with the same active sector). One single GoToFastScan message per sector ensures that all the neighbors will receive at least one of such messages. Along with the GoToFastScan message, the ID of the node that will receive the token next is also sent. This is done for the next TH to know that it should not go to Fast Scan mode as it will receive the token. Then, the TH node sends the token and waits for an acknowledgment. When the acknowledgment is received, the older TH node goes to Fast Scan mode. The total duration of the Token-Passing Mechanism is $T_{TP} = (K - 1) t_{GoToFastScan} + t_{TokenAck}$, where $t_{TokenAck}$ is the time it takes to send the token and receive the acknowledgment.

In case the TH node had already discovered neighbors (it is not the first time it has the token), all its neighbors will be in Fast Scan mode. It is necessary for the TH to make a reduced version of the Hone-In Mechanism (we call it Mini-Hone-In), just for the node that will receive the token next. After enough Mini-Hone-In messages to guarantee that the other node is listening are sent, the TH node sends the token and waits for an acknowledgment. When the acknowledgment is received, the older TH node goes to Fast Scan mode. The total duration of the Token Releasing Mechanism is $T_{TR} = (h - 1) t_{Hone\text{-}In} + t_{TokenAck}$.

In the following paragraphs we will find an analytical expression for the total time taken by the protocol. After the protocol ends, all n nodes in the network have discovered their neighbors forming a tree graph.

The tree graph is formed by the token passage and has n vertexes (because each node is represented by a vertex) and $n-1$ edges (because it is a tree). So there have been n Hone-In, Hello-Reply and Token-Passing mechanisms independently of the network topology, and $2 (n - 1)$ Token Passing and Token Releasing mechanisms (together). There have been then $2 (n - 1) - n = n - 2$ Token Releasing mechanisms, independently of the network topology.

$$T_{SAND} = n (T_{HI} + T_{HR} + T_{TP}) + (n - 2) T_{TR}. \tag{4.1}$$

The neighbor discovery time is independent of the network topology, and is increasingly linear with the number of nodes of the network.

4.2.2 Implementation

We implement the SAND protocol using Contiki OS, and adopting the network stack described in Section 2.3. We use preprocessor macros to define all the protocol parameters: t_{switch}, $t_{Hone\text{-}In}$, h, N_{slots}, t_{slots}, N_{rounds} and $t_{GoToFastScan}$ so they can be easily modified. We define several types of broadcast messages according to SAND: Hone-In message, Hello message, Reply message, GoToFastScan message, Mini-Hone-In message, Token message, and Acknowledgment message. These messages are used by the nodes to communicate using link-local addresses.

When the protocol is initialized, the first TH is the one with ID = 1. In practice, this node usually is the border router of the network. Once the TH finishes discovering neighbors, it checks its neighbor table and passes the token to the neighbor with the lowest ID that has not discovered neighbors yet. Our implementation is based on a state machine, whose diagram is presented in Figure 4.2. In this figure, each row represent each one of the three steps that form the protocol (Hone-In Mechanism, Hello-Reply Mechanism, Token Passing, and Token Releasing Mechanisms), and the two columns represent whether the node is acting as the TH (left) or as the neighbor node being discovered (right).
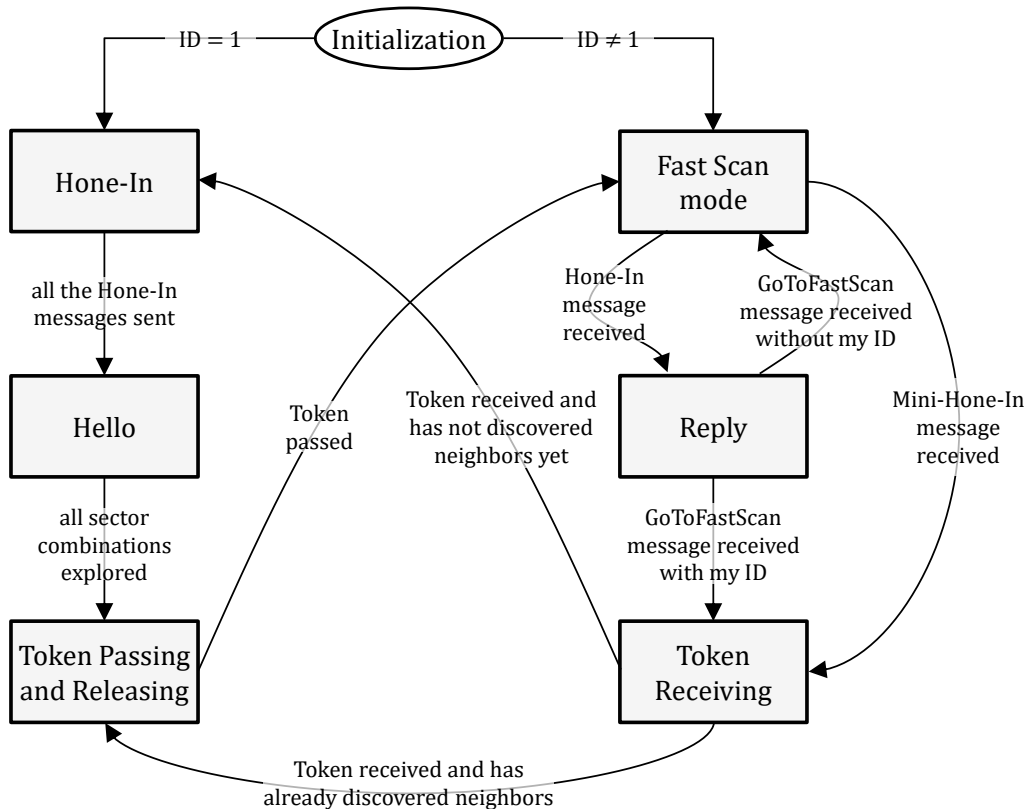

Figure 4.2   State diagram of SAND.

The neighbor discovery table is saved in each node's RAM. This table contains the node's sector, the neighbor ID, the neighbor sector, the RSSI, and a bit indicating whether the neighbor has already been a TH or not. This bit is used by the TH to know to which node to pass the token, as described in Section 4.2. When the protocol ends and all the nodes have completed their respective neighbor discovery table, these bits should be all set.

One simplification of this implementation of the protocol is that the neighbor discovery tables are not passed along with the token, so each node of the network only has access to its own neighbor discovery table. This simplification reduces the size of the messages exchanged, as well as the memory burden of each node, but does not allow to gather the neighbors' information in a centralized location.

However, as each node does not receive the other nodes' tables, it has to update the neighbor discovery bit somehow. We solve this problem by passing the IDs of all the nodes that have already discovered neighbors along with the token when it is passed upwards in the tree. When the token is passed downwards in the tree, the ID of the TH is passed along with the token, for the new TH to know which neighbor to pass the token when all its neighbors have completed the neighbor discovery process.

4.3   Q-SAND

One of the drawbacks of the SAND protocol is that since the neighbor discovery time is increasingly linear with the number of nodes in the network, large networks need large times to complete. We propose an optimization to the SAND protocol to reduce the neighbor discovery time we call Q-SAND (for Quick SAND).

4.3.1   Protocol Description

The main idea of Q-SAND is to reduce the discovery time by only finding the strongest combination of sectors between a pair of neighbor nodes (the ones with the highest RSSI). Reducing the number of sector combinations that the protocol tries to discover reduces the neighbor discovery time proportionally.

We assume that the links are reciprocal [60] and that the strongest sector combination between two nodes is achieved when they point their sectors to each other. This last assumption is only valid for line of sight and open environments. A mathematical way to express this is that a node has a maximum link quality indicator with a neighbor when their active sector is the one that contains the line that connects both nodes. Examples are shown in Figure 4.3 for an even K and in Figure 4.4 for an odd K.
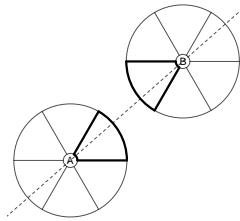

Figure 4.3   Best sector combination between two nodes A and B for K even ($K = 6$).

With this consideration, it is clear that if the sector of node A that contains the line that connects both nodes is $S_A \in [0, ..., K - 1]$, then the sector of node B that contains the same line is:

- $S_B = \left(S_A + \frac{K}{2}\right) \bmod K$ if K is even, and

- $S_B = \left(S_A + \frac{K-1}{2}\right) \bmod K$ or $\left(S_A + \frac{K+1}{2}\right) \bmod K$ if K is odd.
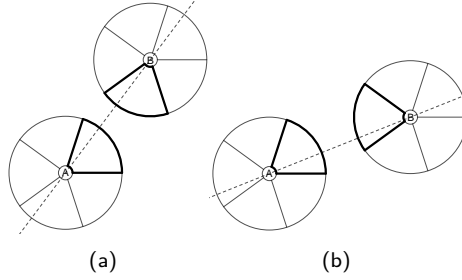
Figure 4.4 Best sector combination between two nodes A and B for K odd ($K = 5$).

If K is even, we only need to try K sector combinations to find the best one (and $2K$ sector combinations if K is odd), compared to the K^2 sector combinations that are needed without this modification. The main drawback is that we would only find the best sector combination, so this improvement is not applicable if we would like to obtain all the discovered sector combinations between two nodes and not just the best one.

An extra consideration compared to SAND is that the nodes need to be oriented in the same manner, for example, all nodes with sector 0 pointing to the north, so this improvement is not applicable if the nodes are free to rotate. The advantage is that the new duration of the Hello-Reply Mechanism (which is the mechanism that commonly takes longer) is now

$$
\begin{cases}
KN_{rounds}N_{slots}t_{slots} & = \quad \frac{T_{HR}}{K} \quad \text{if } K \text{ is even} \\[2ex]
2KN_{rounds}N_{slots}t_{slots} & = \quad 2\frac{T_{HR}}{K} \quad \text{if } K \text{ is odd}
\end{cases}
\tag{4.2}
$$

Then, with Q-SAND, we have that the total time T_{Q-SAND} the protocol takes is

$$
\begin{cases}
n\left(T_{HI} + \frac{T_{HR}}{K} + T_{TP}\right) + (n-2)\,T_{TR} & \text{if } K \text{ is even} \\[2ex]
n\left(T_{HI} + 2\frac{T_{HR}}{K} + T_{TP}\right) + (n-2)\,T_{TR} & \text{if } K \text{ is odd}
\end{cases}
\tag{4.3}
$$

The neighbor discovery time in Q-SAND compared to SAND is greatly reduced when $T_{HR} \gg (T_{HI} + T_{TP} + T_{TR})\,K$ (up to K times if K is even and up to $\frac{K}{2}$ times if K is odd), and little reduced when $T_{HR} \ll T_{HI} + T_{HR} + T_{TP}$.

4.3.2 Implementation

We modify the implementation of SAND described in Section 4.2.2 to follow the protocol description presented in Section 4.3.1. The Hone-In, Token-Passing, and Token-Releasing mechanisms remain unchanged. We introduce our modifications in the Hello-Reply mechanism. If the TH is sending Hello messages with the sector S_A, the neighbor nodes always reply using the sector $S_B = (S_A + 3) \bmod 6$, reducing K times the time of the Hello-Reply mechanism.

4.3.3 Evaluation in Simulation

In this section, we evaluate the Q-SAND protocol in simulation and compare its performance with SAND. We simulate the implementations described in Sections 4.2.2 and 4.3.2 with COOJA for different network topologies and densities. The parameters for both SAND and Q-SAND protocols used in the simulations are shown in Table 4.1.

Table 4.1   Parameters of SAND and Q-SAND protocols used in the simulations.

Protocol mechanism	Parameter	Value
Fast Scan mode	t_{switch}	$31.25\,\mathrm{ms}$
Hone-In Mechanism	$t_{Hone\text{-}In}$	$15.625\,\mathrm{ms}$
	h	12
Hello-Reply Mechanism	N_{slots}	5
	t_{slots}	$15.625\,\mathrm{ms}$
	N_{rounds}	5
Token-Passing Mechanism	$t_{GoToFastScan}$	$15.625\,\mathrm{ms}$

To verify the results of our theoretical analysis, we simulate different random networks with the same number of nodes but different topologies. We obtain the exact same neighbor discovery time for networks with the same number of nodes, confirming the independence of the neighbor discovery time with respect to the network topology for both SAND and Q-SAND protocols.

We then simulate a random network with 16 nodes for both protocols. Figure 4.5 shows the simulated network. The colored node is the border router of the network with ID = 1, and is the node that starts with the discovery process. The tree formed by the token passage is shown with continuous lines. The links between nodes are shown both in continuous and dotted lines.

The resulting neighbor discovery tables are the same for both protocols: the exact same best combination of sectors was discovered for each pair of nodes. Table 4.2 shows the neighbor discovery table of the node with ID 1. As we state

Table 4.2   Neighbor Discovery Table of Node 1.

Node 1 sector	Neighbor ID	Neighbor sector	RSSI
1	7	4	-42
5	9	2	-39
0	16	3	-44

in Section 4.3.1, if the best sector combination between a node and its neighbor is achieved for the sector S_1 of the first node, then it is also achieved for the sector $\left(S_1 + \frac{K}{2}\right) \bmod K$ of its neighbor.

We verify that the best sector combination $(sector A, sector B)$ between every node was achieved by one of the following combinations: $(0, 3)$, $(1, 4)$, $(2, 5)$, $(3, 0)$, $(4, 1)$ or $(5, 2)$.
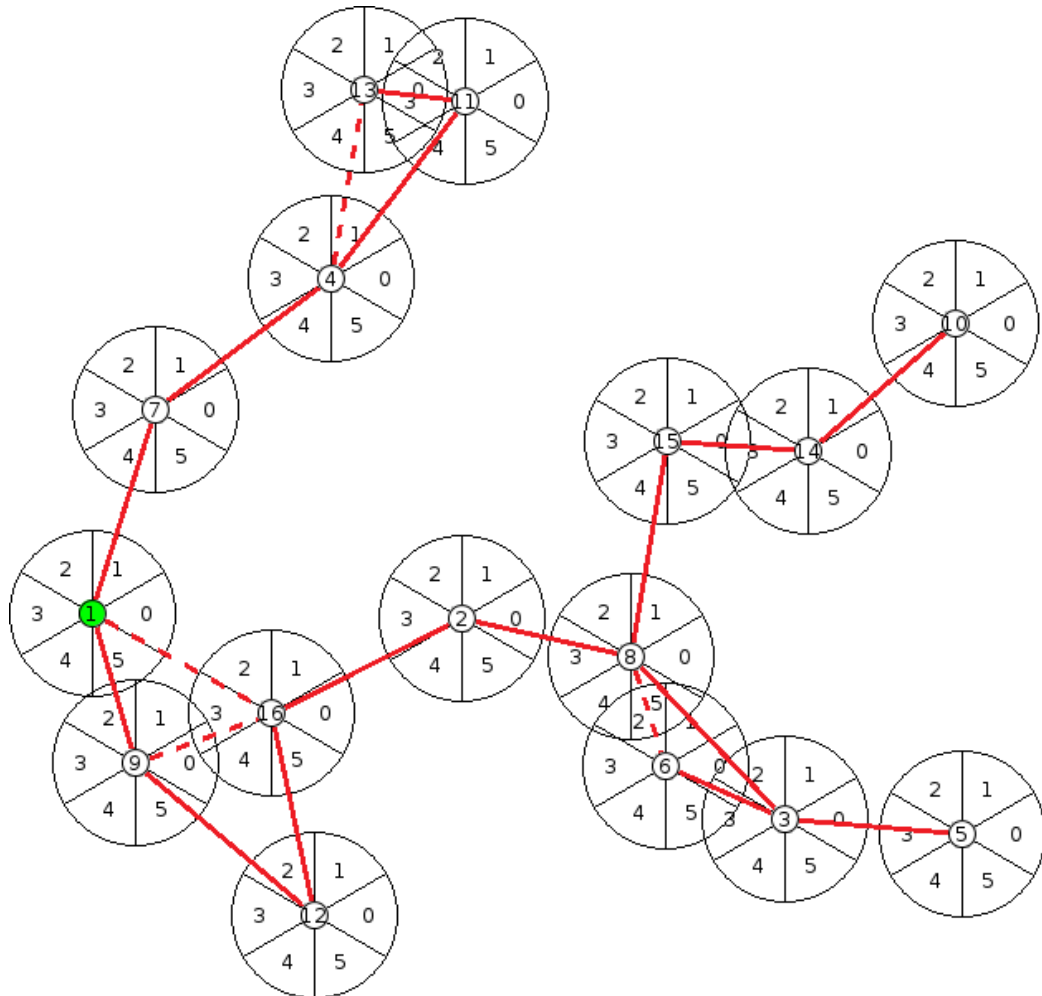
Figure 4.5   Simulated network with 16 nodes.

In addition, we simulate SAND and Q-SAND protocols for different random networks with up to 16 nodes. In every case, the best sector combinations discovered by both protocols has the same RSSI, showing that Q-SAND effectively finds the best sector combinations. Figure 4.6 shows the total time taken by each protocol, containing both theoretical results calculated through equations 4.1 and 4.3 and simulation results. We can see that the simulation results match very well the theoretical equations for both protocols. The total time taken is increasingly linear with the number of nodes in the network for both protocols, taking $15.5\,\text{s}$ per node for SAND and $3.78\,\text{s}$ for Q-SAND. With the set of parameters used, Q-SAND is more than 4 times faster than SAND in discovering neighbors.

4.4   DANDi

In this section, we present the Dynamic Asynchronous Neighbor Discovery protocol for Directional Antennas (DANDi), a novel protocol based on collision detection.
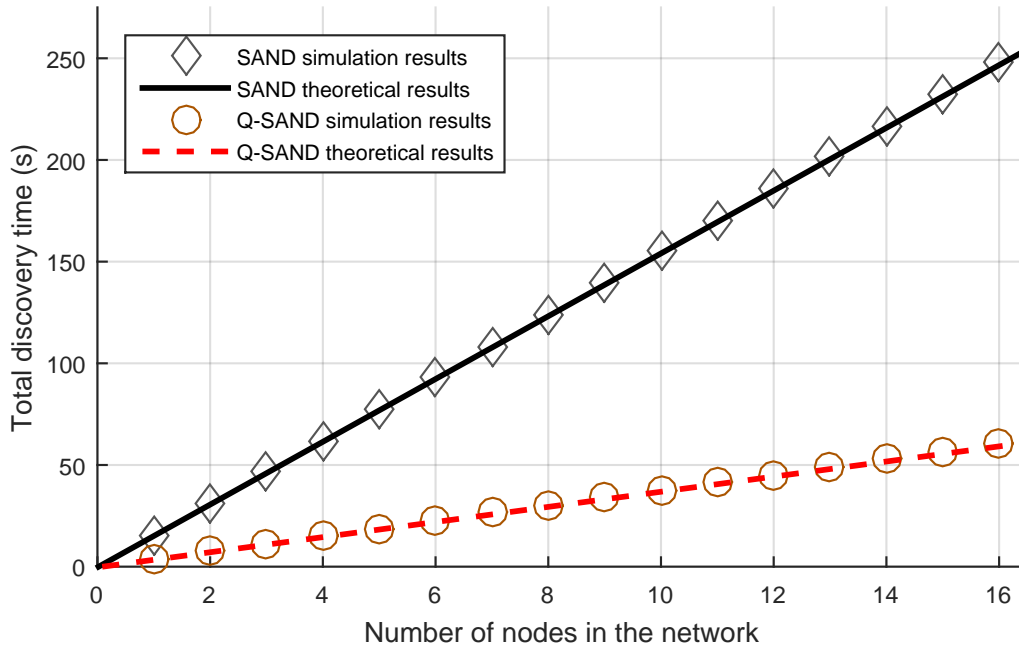
Figure 4.6   Neighbor discovery time vs. number of nodes in the network.

DANDi does not need omnidirectional antennas nor time synchronization; it relies on the serialization of the discovery process (like SAND) and collision detection. DANDi is dynamic, so no network or topology information is needed in advance.

4.4.1   Protocol Description

The main idea of the DANDi protocol is that one single discoverer node (DN) at a time discovers all reliable S2S links with its neighbor nodes (NN) by transmitting probe messages repeatedly, while the remainder network nodes listen for incoming messages in each sector sequentially.

The DN probes one sector at a time using the *probe-reply with dynamic contention resolution mechanism* described below. Once the DN finishes discovering nodes in one sector, it continues with the next one. When the DN ends discovering neighbors in its K sectors, it passes the DN role to a NN through the *token-passing mechanism*.

Figure 4.7 shows a state diagram representing both roles (DN and NN) and the main processes. In this case, the node ID is used to select the initial role, so the first DN is the node with ID 1. The discovery process ends when the DN role is taken again by the node that started as DN, and all of its neighbors have already discovered their own neighbors. The above mechanisms are thoroughly explained in this section.
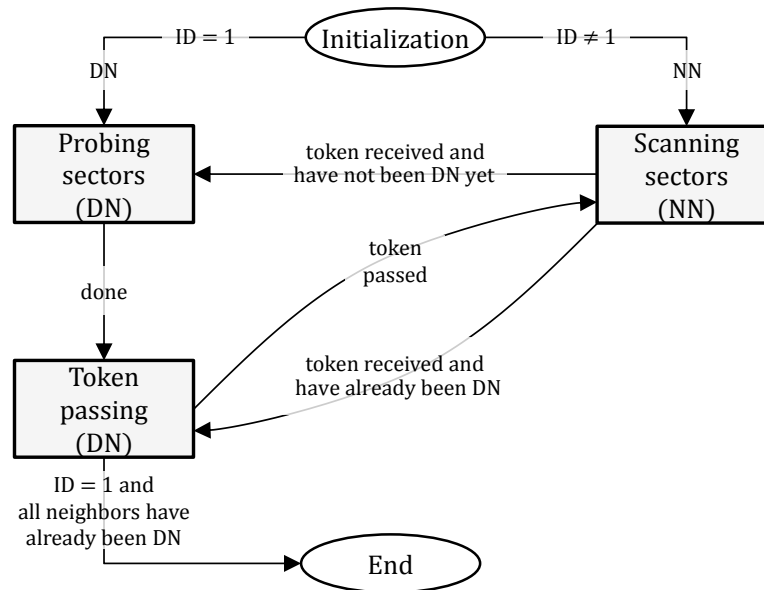
Figure 4.7 State diagram: node roles and main process.

Probe-reply with dynamic contention resolution

The messages exchanged between DN and NN are organized in rounds. The probe message sent by the DN delimits the start of a round. A *round* is composed by a *probe slot* and one or many *reply slots*. A *probe message* is sent at the beginning of the *probe slot*. After the probe slot follows a reply round composed of one or more reply slots. The number of reply slots is specified in the leading probe message. The probe message also includes other information, such as the list of nodes discovered in the last round. Figure 4.8 shows a first round with one reply slot followed by a second round with many reply slots.
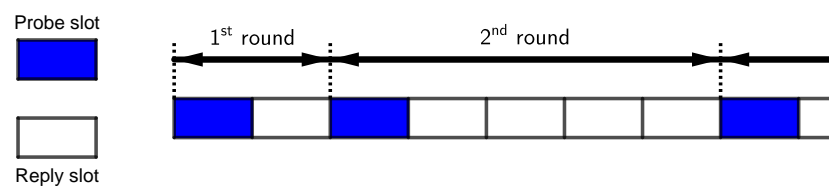

Figure 4.8 Slots and rounds.

The probe-reply with dynamic contention resolution algorithm involves the DN and all the NN in the network. Depending on its actual role, a node executes the DN part of the algorithm (*probing sectors*) or the NN part (*scanning sectors*).

Figure 4.9 shows the probing sectors algorithm of the DN for discovering neighbor nodes sequentially starting at sector 0. The probing of a sector starts by selecting an antenna sector and locking it at that position until it discovers all the neighbors at that sector. The DN initializes the number of reply slots R_{slots} equal to an initial predefined value (e.g. equal to one). Then it sends the probe message of the first round and waits for a reply in each slot during the time t_{slot}. In each
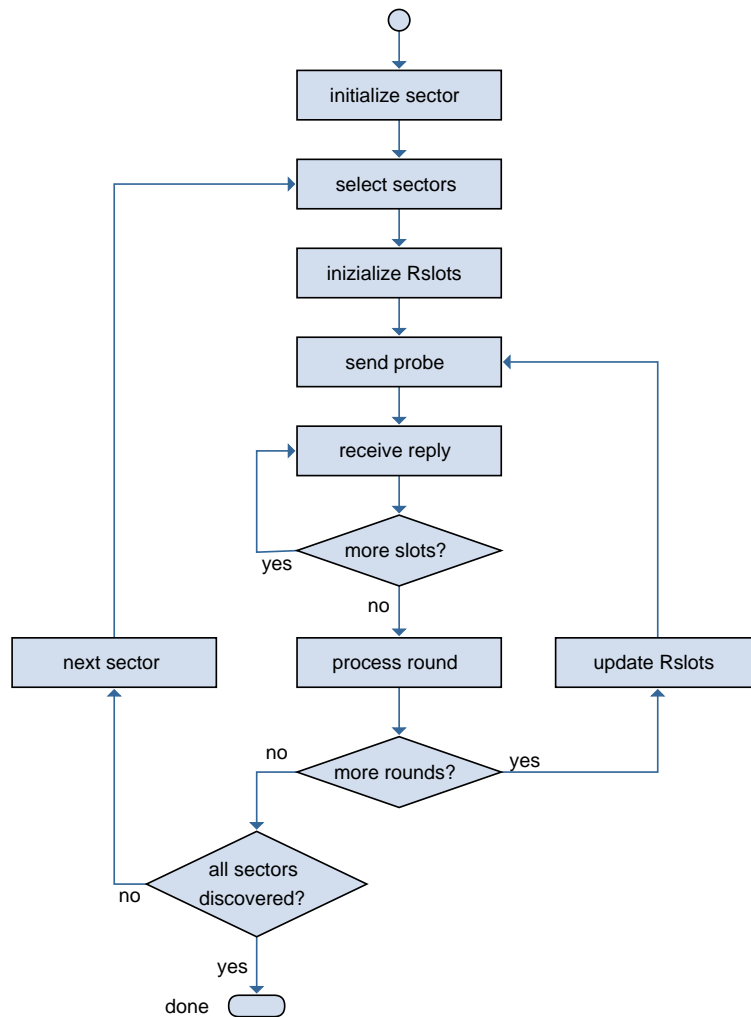
Figure 4.9 Probing sectors algorithm.

reply slot, there are four possible situations shown in Figure 4.10 and summarized below:

(a) No answer.

(b) A single node replies.

(c) Many nodes reply, and one frame prevails over the others due to the capture effect.

(d) Many nodes reply, and results in a collision.

The first case implicates that no neighbor can receive in that sector at that time. From the point of view of the DN, the second and third cases are indistinct, and the DN needs to give the opportunity to any node that had a suppressed

message due to the capture effect, to send its reply again. The capture effect occurs when a node receives two or more messages almost at the same time, and it can demodulate the message with higher signal strength when some conditions are met [65]. For the last case, in which the DN detects a collision, all the replies in that slot are lost. Summarizing, more rounds are needed to ensure that all NN are discovered correctly for every case except the first one.

(a) No answer.

(b) A single node replies.

(c) Many nodes reply: capture effect.
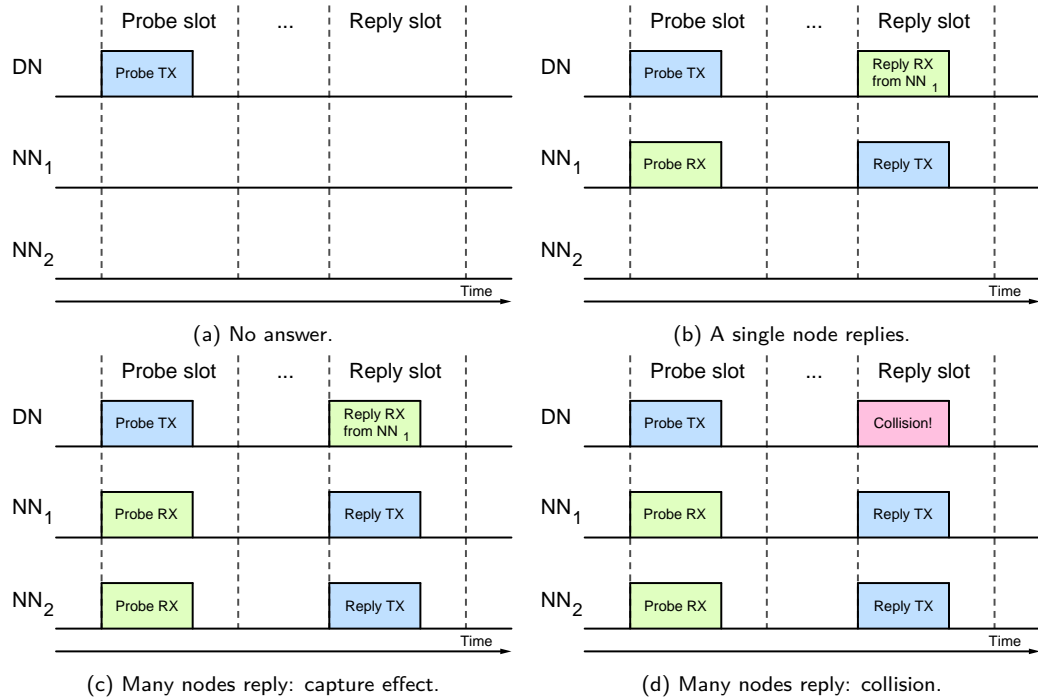
(d) Many nodes reply: collision.

Figure 4.10   Possible results for a reply slot.

When the reply round completes, the result is processed to determine whether a new round is initiated or not. This process consists in analyzing the following cases:

1. No answer in any slot.

2. A reply or a collision happens in any slot.

The second case needs a new round to resolve either the collision/s or give the opportunity to any potential suppressed message to be sent again by the corresponding node. In any of these situations, the DN starts a new round. In this case, the DN selects the number of reply slots for the next round R_{slots} based on the results. The DN could increment the number of slots to speed up the contention resolution if there is a collision in any slot. If there are no collisions, the number of reply slots could be kept unchanged or even reduced. In this work we adopted a simple yet effective solution that is *exponential back-off* [1]. The DN acknowledges the replies in the probe message of the following round, by including the NN identification in the list of discovered nodes.

In the first case where there is no answer in any slot, the DN can switch to the next sector. However, if there is a round with more than one reply slot for contention resolution, there will be no probe slots for a certain time (determined below) since the nodes are resolving the contention. To ensure that the missing slots will not lead a neighbor to miss a probe, the missing probe slots must be recovered. To achieve this, the DN must explore a certain number of rounds with a single reply slot (N_{probe}, determined in Section 4.4.1), and receive no answer in any of them in order to change the sector.

Only after this process is completed, can the DN continue with the next sector. When the DN finishes discovering all the sectors, it passes the token to a NN that has not discovered neighbors yet.
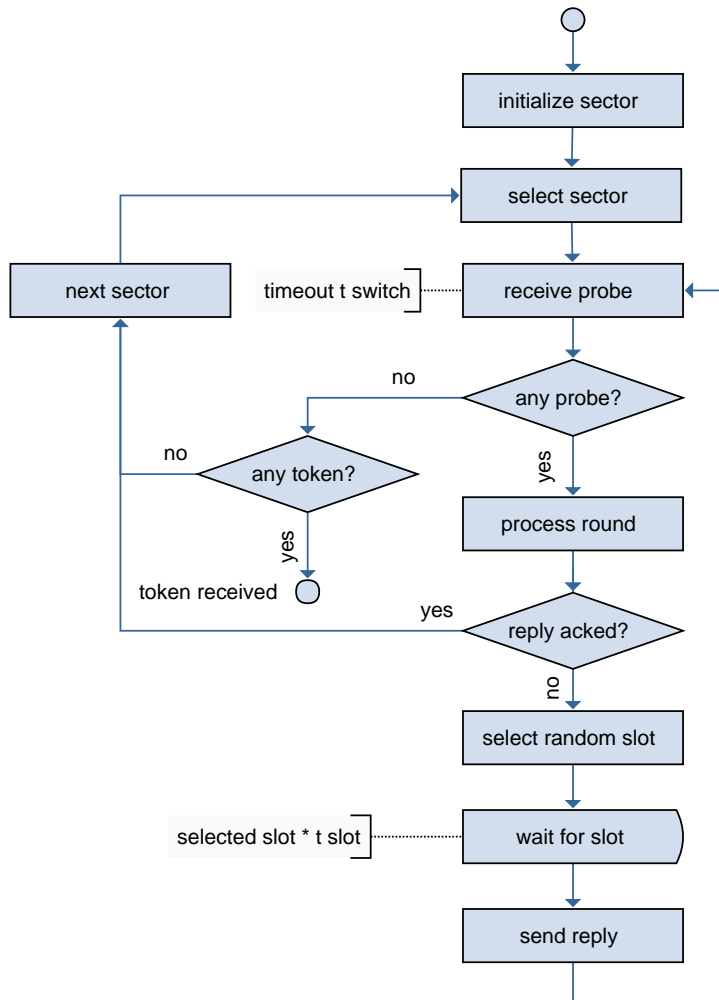

Figure 4.11 Scanning sectors algorithm.

Figure 4.11 shows the scanning sectors algorithm of the NN, where the NN selects one sector and listens for incoming messages during a certain period t_{switch}.

If the NN does not receive a message, it continues with the next sector. If it receives a probe message from the DN, the NN processes the message. It verifies in the probe message whether it is in the list of discovered nodes or not. If it is not in the list, the NN randomly chooses a slot and sends a reply message.

If the NN is in the list, it means that the DN received the previous reply correctly and so the NN continues scanning sectors. The NN also continues with the next sector if it had already received a probe sent by this DN from the same S2S link. This is why probe messages need to include the sector of the DN.

If the NN receives a token message instead of a probe at any time, it stops scanning sectors and assumes the DN role.

Token-passing mechanism

After a DN finishes discovering neighbors in all of its sectors, it passes the discoverer role to a NN that has not discovered neighbors yet. To achieve this, a token identifies the node with the DN role, and a special message is used to pass the token to another node. In the case that all the neighbors of a DN have already discovered their neighbors, the DN passes the token to the node from which it received the token initially. Since the NN are scanning sectors, the DN locks its antenna in the direction of the node that will receive the token message and sends probe messages repeatedly until it can guarantee that the neighbor node received it. For the NN to be able to receive at least one probe per sector, the restriction is the same as when probing a sector. After having sent the probe messages, the DN sends the token and waits to receive an acknowledgment.

In this way the token-passing process between nodes gradually forms a directed rooted tree, where the node that initiates the discovery process is the tree root, the remainder vertexes are the nodes that have already been in the DN role, and the edges represent the token passing relation.

Figure 4.12 depicts the neighbor discovery performed by each DN, the token passing between nodes, and the formed tree.

Probe, reply and switching time restrictions

The probe slot duration has some restrictions. Considering that probe and reply messages are at most of the maximum frame length, then the duration of these messages is limited by a given time τ_{frame}. The period for sending probe messages is t_{probe}. A guard time t_{guard} between slots may be necessary for turning around and switching from transmitting to receiving or vice-versa. Then, the probe period must satisfy

$$t_{probe} > 2\left(\tau_{frame} + t_{guard}\right). \tag{4.4}$$

Figure 4.13a shows the minimum probe period required to allocate a probe or a reply and the actual time used by the DN probe and a NN reply.

NN asynchronously switch their active sector every t_{switch}, scanning for incoming messages. Note that t_{switch} is the period for switching the active sector, but while switching, there is a time τ_{switch} which accounts for the hardware switching time, where the node cannot receive anything.
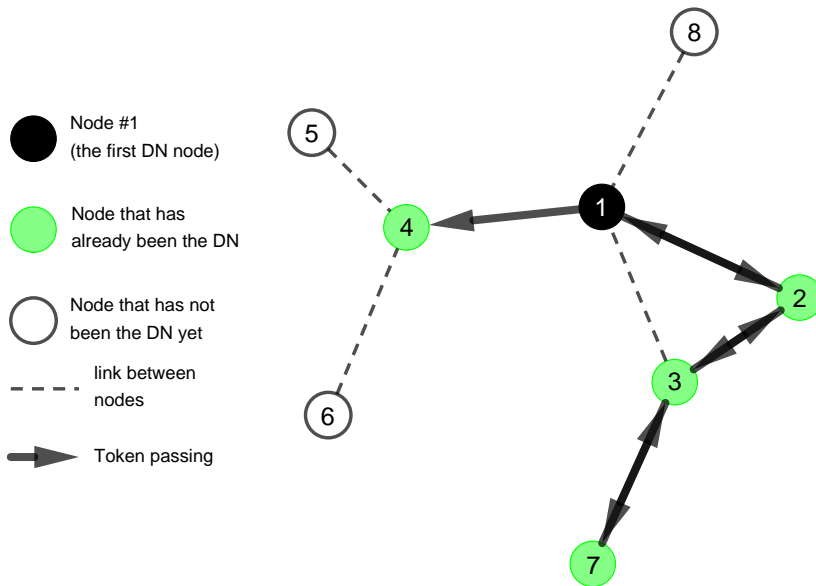
Figure 4.12 Partial tree started by node #1. Nodes #2, #3 and #7 have already been the DN, while node #4 is the current DN. Nodes #5, #6 and #8 have not been the DN yet.

(a) Restriction on t_{probe}.

(b) Restriction on t_{switch}.

(c) Restriction on N_{probe}.

Figure 4.13 Time restrictions.

To ensure that the DN neighbor nodes are able to receive at least one probe message, the following condition must be satisfied:

$$t_{switch} > t_{probe} + \tau_{frame} + \tau_{switch}. \tag{4.5}$$

Figure 4.13b depicts the time involved in the probing process and shows the minimum t_{switch}.

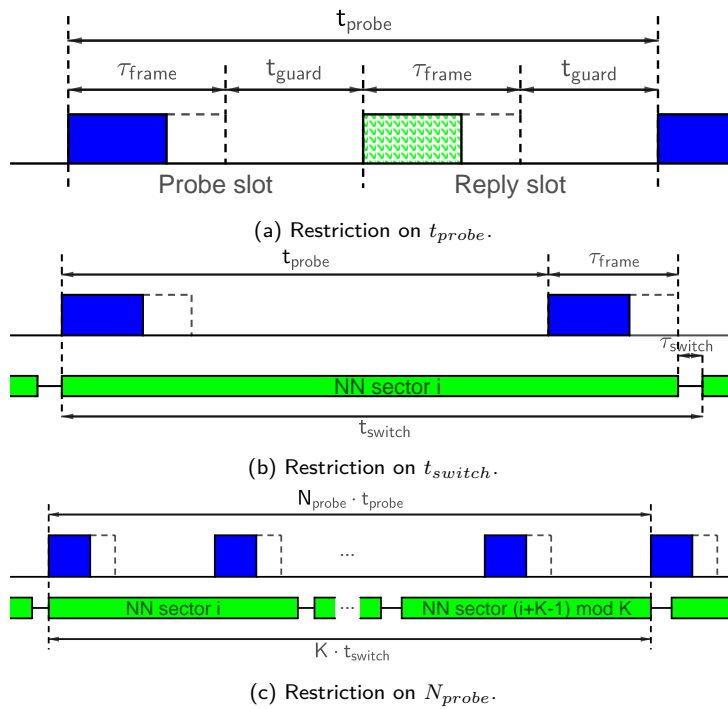In order to send enough probe messages so that every NN has the opportunity to receive at least one of them, the total time sending probes at one sector must be larger than the time of one "full turn" scan of the neighbors:

$$N_{probe} \cdot t_{probe} > K \cdot t_{switch}, \tag{4.6}$$

where N_{probe} is the number of probes sent per sector. This relationship can be observed in Figure 4.13c.

Hereinafter, for the sake of simplicity, we define a generic time slot for sending probe and reply messages of duration $t_{probe} = t_{switch}/2$, and long enough to satisfy the above restrictions. Note that during the time t_{switch} there is enough time for a probe slot and a reply slot.

Finally, to minimize the duration of the overall neighbor discovery process, the number of messages N_{probe} must be also minimized. Then considering Eqs. (4.5) and (4.6), we obtain

$$N_{probe} > K \left(1 + \frac{\tau_{frame} + \tau_{switch}}{t_{probe}} \right). \tag{4.7}$$

The discovery process forms a directed rooted tree, whose edges represent the passage of the token between nodes. The network has n nodes (vertexes), and since it is a tree it has $n-1$ edges. At the end of the protocol, there have been $2(n-1)$ token passing (two tokens passing per edge: one upwards in the tree and one downwards). Taking this into consideration, the total time the protocol takes to complete is the sum of the time taken to discover neighbors by each network node through the *probe-reply mechanism* and the time taken to pass the token through the *token-passing mechanism* (there are $2(n-1)$ in total):

$$T_{DANDi} = \sum_{i=1}^{n} \{T_{PR}(i)\} + 2(n-1)T_{TP}, \tag{4.8}$$

where $T_{PR}(i)$ and T_{TP} are the times taken by the probe-reply mechanism of node i and by the token-passing mechanism respectively.

The time taken by the probe-reply mechanism of node i can be expressed as:

$$T_{PR}(i) = t_{slot} \sum_{j=0}^{K-1} \left\{ \sum_{k=1}^{R_{ij}} \{1 + s_{ijk}\} \right\}, \tag{4.9}$$

where R_{ij} is the number of rounds needed to complete the discovery process for node i for its active sector j, and s_{ijk} is the number of reply slots that has the k^{th} round, when node i is the DN and its active sector is j.

$T_{PR}(i)$ is probabilistic, given that the random choice of the reply slot by the NN impacts on the number of suppressed messages or collisions, and then in the number of rounds needed to resolve the contention. It could happen that two NN

choose the same reply slot for a lot of consecutive rounds and as a result, a lot of rounds would be needed. In this case the time $T_{PR}(i)$ would be greater than if the nodes had chosen different reply slots earlier. So the time $T_{PR}(i)$ will vary with i, expecting it to be greater for nodes with more neighbors. Even more, the time of probing sectors will vary from sector to sector for a given DN, being greater for sectors in which the DN has more neighbors.

In case that a DN has at most one neighbor per sector, there will be no collisions and $T_{PR}(i)$ will be the minimum possible time of a probe-reply mechanism duration. In this case $T_{PR}(i)$ is deterministic and it is given by the following equation:

$$T_{PR}(i) = 2 \cdot t_{slot} \cdot K \cdot N_{probe}. \tag{4.10}$$

The time taken by the token-passing mechanism can be expressed in the following way:

$$T_{TP} = (N_{TP} - 1) t_{probe} + t_{token-ack}, \tag{4.11}$$

where N_{TP} is the number of messages sent from the current DN to the NN, and $t_{token-ack}$ is the time that takes to send the token and receive an acknowledgment.

4.4.2 Considerations for Lossy Networks

Our first implementation of the DANDi protocol was extensively tested and assessed using the COOJA network simulator. The simulation setup used a radio medium with no packet losses and without interference.

A real-world wireless network faces some practical issues such as frame losses (due to interference, fading, etc.) that must be taken into account in the protocol's design and implementation for it to work properly.

Our original design in a network with real nodes performs poorly. In some cases, not all the sectors or nodes are discovered, and in other cases, the protocol does not even finish because the token messages are lost. Considering these aspects, the protocol's design must include mechanisms to cope with lossy links to guarantee: i) the discovery of all reliable S2S links and ii) the passage of the token.

We analyze the effect of losing the different kinds of messages in the original protocol, and we later present modifications to deal with these issues.

Probe

Losing a probe message may lead to the failure to discover some S2S links. If a probe is lost, any NN that would have received the probe will continue scanning sectors. Depending on the moment the probe is lost, NNs may receive subsequent probes or fail to discover that particular S2S link if the DN changes its active sector. To reduce the probability of losing a probe message, the number of probe messages sent without an answer before changing the active sector (N_{probe}) can be increased. This would also increase the discovery time per sector.

Reply

In case a NN's reply message is lost, the DN will not send the ACK message (sent in the probe message of the following round), and the NN will send again a new reply. Consequently, there is no need to introduce modifications. The effect is that more rounds may be needed to discover every S2S link.

Token

The loss of a token message would lead to stopping the discovering process, leaving the protocol in an inconsistent state, so it is fundamental to make the token passage reliable. The token message is then acknowledged. A timer is used to retransmit the token in case the DN does not receive the ACK message.

Ack

If the ACK message corresponding to a token message is lost, the DN will continue to send token messages until the corresponding ACK is received (even if the NN received the token correctly). If the previous DN fails to receive the ACK but receives any message that indicates that the target node received the token correctly (a probe message sent by the new DN or a reply message of any node), it assumes that the token message was correctly delivered and takes the NN role.

4.4.3 Implementation

We implement DANDi in Contiki OS for the Tmote Sky and use the default 6LoW-PAN protocol stack described in Section 2.3. To validate DANDi, we simulate a network of 16 Tmote Sky nodes with 6-sectored antennas in the COOJA network simulator using the model presented in Section 3.2.5. The radio medium does not consider packet losses or interference from external sources (e.g. WiFi, Bluetooth).

The experiments with real sensor nodes are performed in an office environment with moderate WiFi interference. We deploy a network of 10 Tmote Sky nodes with 6-sectored SPIDA antennas, described in Section 3.2.3. The output power is set near the minimum value supported by the radio hardware ($-25\,\mathrm{dBm}$) in order to enable testing the network in a reduced space. In this configuration, the maximum communication range is around six meters. Ten nodes were distributed over an area of 20 by 8 meters as shown in Figure 4.14.

We compare the performance of DANDi with our own implementation of SAND, developed following the description presented in Section 4.2. We defined the following types of messages according to the protocol: i) probe, ii) reply, iii) token, and iv) acknowledgment. These messages are sent using link-local addresses. Probe messages are *broadcast* (since their destination is any NN), while reply messages are *unicast*, addressed to the DN.

The parameters of DANDi used in the simulations are shown in Table 4.3. In this implementation, if a NN chooses the first reply slot, it replies immediately after having received the probe message. This is done to reduce the discovery
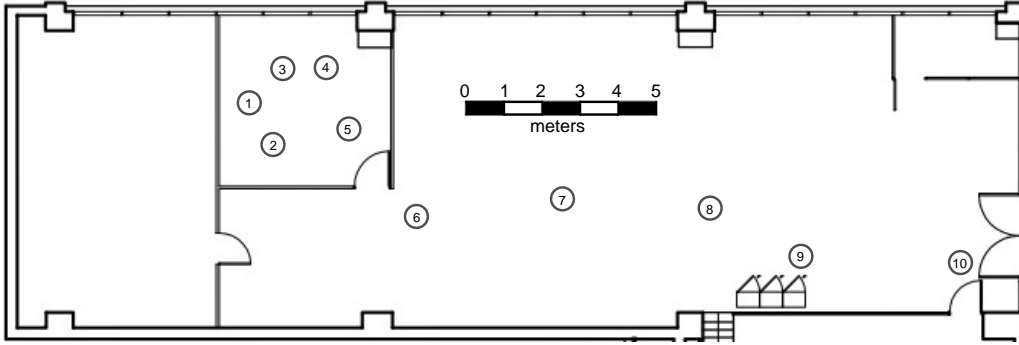
Figure 4.14   Network deployed in an office environment.

Table 4.3   Parameters of the DANDi protocol used in the simulations.

DANDi	
Parameter	Value
t_{switch}	$62.5\,\mathrm{ms}$
t_{slot}	$31.25\,\mathrm{ms}$
N_{probe}	13

time, as well as to keep the implementation simple. As a consequence, the probe slot is very short (negligible) compared to the reply slots, so the time taken by the probe-reply mechanism of node i can be simplified as:

$$T_{PR}(i) = t_{slot} \sum_{j=0}^{K-1} \left\{ \sum_{k=1}^{R_{ij}} \{s_{ijk}\} \right\}, \tag{4.12}$$

and the minimum possible time (in case of no collisions whatsoever) becomes:

$$T_{PR}(i) = t_{slot} \cdot K \cdot N_{probe}. \tag{4.13}$$

We still have to choose how the number of slots of each round s_{r_i} is selected depending on the number of slots $s_{r_{i-1}}$, the number of collisions $c_{r_{i-1}}$ and the number of discovered neighbors $N_{r_{i-1}}$ in the past round. A mathematical way to express this is through a function: $s_{r_i} = f\left(s_{r_{i-1}}, c_{r_{i-1}}, N_{r_{i-1}}\right)$. The function we used in the simulations is the following, based on exponential back-off as mentioned in Section 4.4.1:

$$s_{r_i} = \begin{cases} 1 & \text{if } i = 0 \text{ or } c_{r_{i-1}} = 0 \\ 2s_{r_{i-1}} & \text{otherwise.} \end{cases} \tag{4.14}$$

This function strongly impacts in the time required by the protocol to complete. The optimization of the function used to determine the number of slots of a given round is out of the scope of this work.

4.4.4   Collision Detection

One of the key features of DANDi that enables the dynamic contention resolution is the collision detection mechanism. The DN must be able to detect if more

than one node replies in the same slot, generating a collision. A receiver-side collision detection can be implemented using Cyclic Redundancy Check (CRC) or Clear Channel Assessment-based (CCA) techniques [10], adopted in some other protocols [28]. We chose to implement the collision detection based on received signal strength measured by the radio receiver as the RSSI.

To determine if a collision occurs, the DN checks the RSSI level during the reply slots. If the RSSI is above a certain threshold for a certain time, it means one or many NN may have sent a reply message. If the DN does not receive a reply message it assumes a collision occurred.

The RSSI level is polled repeatedly during the reply slot. In our implementation, we were able to obtain a valid RSSI measurement every $200\,\mu s$. A reply packet has 72 bytes of data and takes approximately $2.46\,ms$ to be transmitted, so we poll the RSSI level around 12 times during the transmission of a reply packet. We set our threshold to seven consecutive checks were the RSSI is above a certain power level to determine if a packet is being transmitted. Seven consecutive checks correspond to $1.4\,ms$. This time restriction is very important to distinguish between packet collisions and interference. WiFi signals can make the RSSI level increase significantly for short periods, even though the nodes are working in the channels with lower overlap with WiFi. In Section 4.4.4, we show that there are no false-positive measurements due to WiFi interference if we consider that a packet transmission should be at least $1.4\,ms$ long.

The power level above which we say there is a packet transmission must be properly selected. On the one hand, a low threshold close to the radio sensitivity (e.g. $-95\,dBm$ for the CC2420 radio of the Tmote Sky) could increase the number of false-positive (FP) detections (a collision is detected and there was none) due to a high noise level (e.g. Wi-Fi interference). On the other hand, a relatively high threshold could increase the number of false-negative (FN) detections (a collision occurred and it is not detected), resulting in missing detecting a real collision just because the received signal is low.

A threshold of $-88\,dB$ was selected based on empirical experiments, detailed in Section 4.4.4. We used the criteria of selecting the lower threshold that resulted in $0\,\%$ of FP, even though we know that the protocol could fail to discover links with an RSSI lower than $-88\,dB$ due to FNs in the collision detection mechanism. The S2S links that we would fail to discover would be the weaker ones, and probably would not be used for further communication to avoid packet losses even if they were discovered.

Collision Detection Evaluation

We performed several experiments to assess the performance of the collision detection mechanism, both in simulation and with real sensor nodes.

We performed two simulated experiments as a first step in verifying our approach. In experiment #1, we placed one DN periodically sending probe messages and one NN that replied to each probe message. This experiment corresponds to a scenario without collisions and it aims to obtain an estimation of the FP rate. In experiment #2, we placed a second NN replying the probe messages in the

same slot as the other NN to force collisions. This experiment aims to obtain an estimation of the FN rate. Table 4.4 shows the results of the simulation of these experiments.

Experiment	Probe messages sent	Reply messages received	Collisions detected
#1	10 000	10 000	0
#2	10 000	924	9076

Table 4.4   Collision detection simulation results.

Results from experiment #1 show that we have no FPs, as no collisions were detected. In experiment #2, even though there is only one slot for the NNs to reply, the DN can receive some packets due to the capture effect. The remaining cases are collisions correctly detected so we can conclude that we have no FNs. These results show that the logic of the implementation works fine, but to assess the real performance we have to test it with real sensor nodes.

We repeat the experiments performed in simulation with real sensor nodes in an office environment (as described in Section 4.4.3) adding a third experiment. For experiment #3 we placed a single DN sending probe messages, without any NN. The main objective of this experiment was to find out if WiFi interference was detected as packet collisions. Results from these three experiments are shown in Table 4.5.

Experiment	Probe messages sent	Reply messages received	Collisions detected
#1	10 000	9561	125
#2	10 000	3016	6822
#3	10 000	0	0

Table 4.5   Collision detection experimental results.

For experiment #1 (scenario without collisions), we can see that we have detected 125 collisions when we were expecting none. We can also observe that 9561 reply messages were received by the DN. The remaining cases are probe messages that did not receive the corresponding reply message nor were detected as collisions.

In these kinds of networks, depending on the link quality, some packets are expected to be lost. We analyzed the sequence number of each packet and found that 229 reply messages were lost. When a reply packet is lost, it is possible that it was sent by the NN but could not be received correctly by the DN due to interference. If this is the case, the DN would find that the RSSI exceeds the threshold long enough to mistake it with the collision of two or more reply messages. Then, the 125 detected collisions correspond to some of the 229 reply messages sent by the NN but not received correctly by the DN due to interference.

For experiment #2 (scenario forcing collisions), we can observe that there are 162 probe messages that did not receive a reply nor were detected as a collision ($10\,000 - 3016 - 6822 = 162$). Analyzing the sequence number of the packets, we can observe that from these probe messages, 125 were not received by either of

the NN. So the remaining 37 probe messages that were replied by the NN, but not received nor detected as a collision by the DN are the candidates for FNs. We say candidates because it is also possible that the NN sent the replies but the DN did not receive any of them. But if we take a worst-case scenario, the FN ratio would be:

$$FN \leq \frac{37}{6822 + 37} = 0.54\,\%.$$

For experiment #3 we obtained no detections of collisions for 10 000 probe messages transmitted. Based on the results of the experiments, we can conclude that our implementation is robust to interference from the environment.

4.4.5   Evaluation in Simulation

This section presents the performance results of the DANDi protocol and a comparison with the state-of-the-art protocol SAND. We first present simulation results and in Section 4.4.5 we present experimental results with real sensor nodes. To assess the protocol effectiveness, we simulate DANDi for many different network topologies, and we verify that every link is discovered. We observe that the total time taken by the DANDi protocol varies with the network topology, where denser networks take more time to discover neighbors because there are more links to be discovered, and lesser dense networks take less time to discover. To investigate different aspects, we choose to deepen the analysis on a particular topology where we have different areas with different densities.
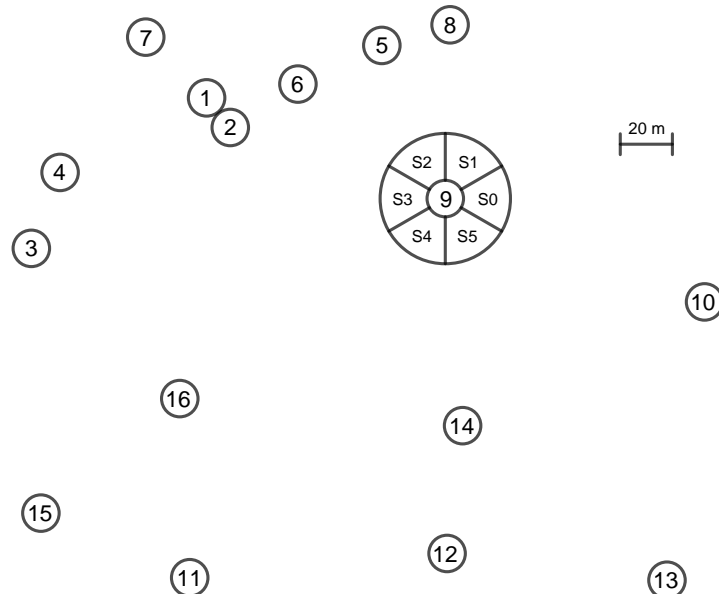

Figure 4.15   Simulated network with 16 nodes. The sectors of all nodes are oriented as shown for node #9.

Performance evaluation

In the first place, we simulate the 16-node network shown in Figure 4.15 with 25 different simulation seeds. In every case, there were 666 S2S links discovered in total.

The total time taken by the protocol was 1 min 31 s averaging the 25 simulations. The simulation that took the least time took 1 min 15 s, and the one that took the longest took 1 min 48 s.

Table 4.6 Number of S2S links per sector of each node of the network.

Node	S_0	S_1	S_2	S_3	S_4	S_5	Total
#1	20	18	12	16	15	15	96
#2	16	16	16	15	14	10	87
#3	12	12	6	3	2	7	42
#4	15	11	7	5	5	10	53
#5	7	5	12	13	16	11	64
#6	9	9	15	14	14	16	77
#7	16	8	0	6	13	17	60
#8	0	2	6	13	13	6	40
#9	1	6	12	8	2	3	32
#10	0	0	1	4	3	0	8
#11	1	3	5	4	2	0	15
#12	5	5	4	3	0	1	18
#13	0	1	3	5	1	0	10
#14	6	4	2	3	4	6	25
#15	6	7	3	0	1	3	20
#16	1	2	4	4	5	3	19

Table 4.6 shows the number of S2S links discovered for each sector of each node in the network. Note that a node can discover more than one S2S link per neighbor. At most, a node can discover K^2 (in our case 36) S2S links per neighbor. As expected from the position of the nodes, we identify two unevenly dense areas in the network that generate different number of S2S links per node: a very dense area near nodes #1 to #8 (with 40 to 87 S2S links per node) and a lesser dense area near nodes #10 to #16 (with 8 to 25 S2S links per node). Node #9 is between both areas with 32 S2S. Another expected result is that the nodes that are on the edges of the network have no S2S links when pointing away from the network (e.g. the case of node #10 for sectors S_0, S_1, and S_5).

Table 4.7 shows the average time taken by the probe-reply mechanism of each sector of each node in the network. We can see that for a sector with more S2S links to discover, it takes longer for the probe-reply mechanism to complete. This result is expected, since more S2S links in a sector generate more collisions which induce more rounds and thus more time. We can also see that the minimum time taken by the probe-reply mechanism per sector is 0.406 s. This corresponds with Equation (4.13), since $T_{PR}(i)/K = t_{slot} N_{probe} = (31.25\,\text{ms}) \times 13 = 406.25\,\text{ms}$.

Table 4.7 Average time taken by the probe-reply mechanism of each sector of each node in the network.

Node	S_0 (s)	S_1 (s)	S_2 (s)	S_3 (s)	S_4 (s)	S_5 (s)	Total (s)
#1	1.839	1.550	1.083	1.783	1.519	1.261	9.034
#2	2.543	1.340	1.415	1.438	1.421	1.201	9.358
#3	1.215	1.020	0.550	0.406	0.406	1.103	4.700
#4	1.596	1.096	0.611	0.406	0.406	0.863	4.979
#5	0.610	0.406	1.100	1.106	1.359	1.423	6.004
#6	0.888	0.785	1.299	1.759	1.414	1.968	8.111
#7	1.623	0.955	0.406	0.661	1.139	1.835	6.619
#8	0.406	0.406	0.486	1.278	1.311	0.713	4.600
#9	0.406	0.489	1.286	1.029	0.406	0.488	4.104
#10	0.406	0.406	0.406	0.504	0.481	0.406	2.610
#11	0.406	0.476	0.476	0.406	0.406	0.406	2.578
#12	0.421	0.458	0.406	0.492	0.406	0.406	2.590
#13	0.406	0.406	0.478	0.561	0.406	0.406	2.664
#14	0.609	0.428	0.406	0.439	0.406	0.490	2.778
#15	0.551	0.833	0.464	0.406	0.406	0.406	3.066
#16	0.406	0.406	0.525	0.538	0.626	0.450	2.951

The average of the total values shown in Tables 4.6 and 4.7 (last column in both tables) are graphically presented in Figures 4.16 and 4.17 respectively. Figure 4.17 also includes the minimum and the maximum time taken by all simulations.
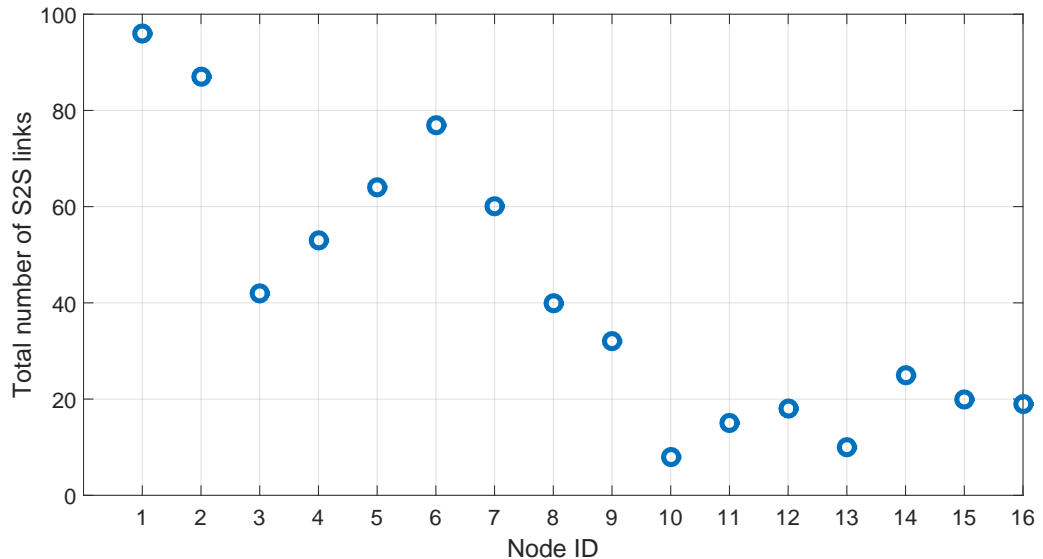

Figure 4.16 Total number of S2S links for each node of the network.

Figures 4.16 and 4.17 show similar tendencies, confirming that the nodes that take longer to discover its neighbors are the ones that have more S2S links. We
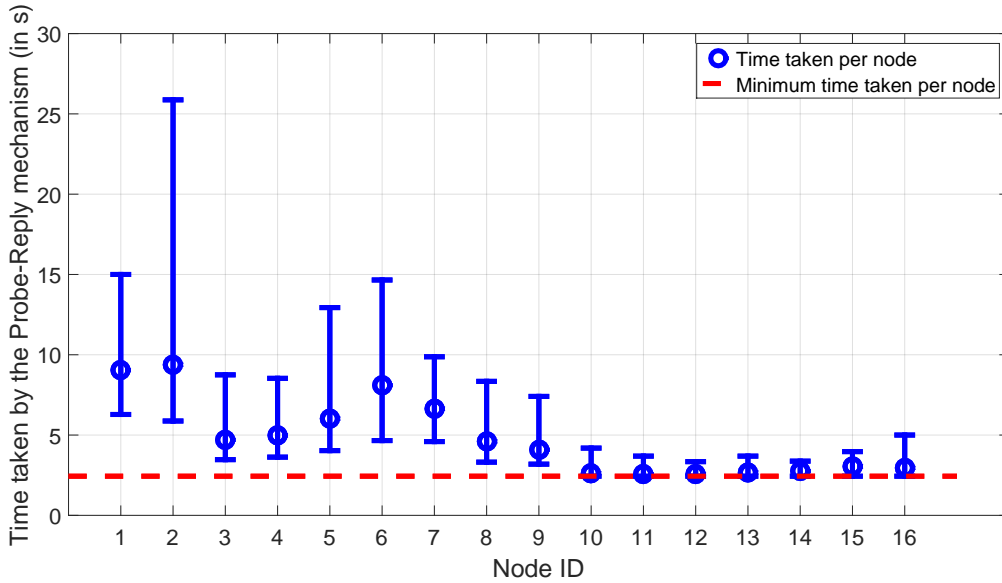
Figure 4.17 Minimum, average and maximum time taken by the probe-reply mechanism for each node of the network.

also see in Figure 4.17 the minimum time taken per node (in case there are no collisions between reply messages), calculated using Equation (4.13): $T_{PR}(i) = t_{slot}KN_{probe} = (31.25\,\text{ms}) \times 6 \times 13 = 2.4375\,\text{s}$. This result shows that nodes #10 to #16 take a time very close to the minimum.

Comparison with SAND

We also implemented and simulated SAND protocol for the same networks to compare it with DANDi. For SAND, s and r are fixed and must be configured in advance before running the protocol. The parameters of both DANDi and SAND protocols used in the simulations are shown in Table 4.8. For the comparison to be fair, we used the same protocol parameters and the same network topologies. Table 4.8 shows that the DANDi protocol is simpler than SAND in that it has fewer parameters.

Table 4.8 Parameters of DANDi and SAND protocols used in the simulations.

DANDi		SAND	
Parameter	Value	Parameter	Value
t_{switch}	$62.5\,\text{ms}$	t_{switch}	$62.5\,\text{ms}$
$t_{pretoken}$	$31.25\,\text{ms}$	$t_{Hone\text{-}In}$	$31.25\,\text{ms}$
N_{probe}	13	h	12
dynamic s and r		s	$s(N_{max}, 99\,\%)$
		r	$r(N_{max}, 99\,\%)$
t_{slot}	$31.25\,\text{ms}$	t_{slot}	$31.25\,\text{ms}$
- (not used)		$t_{GoToFastScan}$	$31.25\,\text{ms}$

For SAND, the parameters s and r were optimized taking into account the maximum number of neighbors per sector N_{max} for each simulated network and a probability of discovering all of them of $p = 99\%$. The authors of SAND proposed an algorithm that given N_{max} and p, allows to obtain s and r minimizing the protocol duration [18,39]. We computed that algorithm to make a fair comparison between DANDi and SAND.

In the first place, we simulate the network shown in Figure 4.15 with SAND for $s = 5$ slots and $r = 4$ rounds. If s or r were smaller, SAND would take less time to finish, but it would fail to discover all the 666 S2S links with a probability greater than 1%. The total time taken by SAND was 6 min 46 s. This time is approximately 4.46 times higher than the average time taken by DANDi. In the case of SAND, once the parameters s and r are fixed, the time the protocol takes is deterministic, and all the nodes in the network take the same amount of time. Even more, each sector of every node in the network takes the same amount of time to discover its neighbors.

It is clear from the results that for a random network with a different number of neighbors in each active sector of each network node, DANDi outperforms SAND regarding network discovery time. This is an expected result because, on one hand, if we choose for SAND a small number of slots s and rounds r (suitable for sectors with very few neighbors), the sectors that have more neighbors will very likely fail to discover all of them. On the other hand, if we choose a number of slots s and rounds r large enough (suitable for sectors with the largest number of neighbors in the network), we can make the probability of discovering all of those neighbors arbitrarily high, but the time taken by the protocol would be unnecessarily high for the sectors with few neighbors. In the case of DANDi, r and s are dynamic, in such a way that the time taken to discover neighbors on a sector is according to the neighbors of that sector. The only time overhead in DANDi is due to the added probe messages to recover probe gaps in case of contention (round with a number of reply slots greater than one).

Besides reducing considerably the neighbor discovery time, DANDi does not require a previous knowledge of the network. Note that we compared DANDi to the better-optimized version of SAND ($s = 5$ and $r = 4$), which was chosen considering the network topology. Since we want to discover all S2S links, even for the sectors with more neighbors in the network, it is not sufficient to simply estimate the number of neighbors based on the network density as if the nodes were evenly distributed, but we need the topology information. In a real deployment, it is not very likely to have such information beforehand. Additionally, if the network changes (i.e. nodes are added, moved, or removed), no changes are required for the DANDi protocol. However, for SAND protocol, a change in the parameters s and r might be needed to re-optimize SAND, and all the nodes in the network would have to be reprogrammed with the new parameters.

We proceeded to compare DANDi with SAND for the scenario of a network with no collisions whatsoever: every sector must have at most one neighbor. We created the network depicted in Figure 4.18, where each node is in the range of the consecutive nodes only. As there are no collisions, the better optimized version

of SAND is with $s = 1$ and $r = 1$. We used those parameters of SAND for this network simulation.

Figure 4.18   Simulated network with 16 nodes and no collisions.

For both protocols, the 30 S2S links were discovered. We obtained the same results (milliseconds of difference) for different simulation seeds, both for DANDi and SAND. This indicates that in the case of a network with no collisions, DANDi's duration is deterministic. DANDi took 51.26 s to discover all nodes, while SAND took 63.29 s. We can see that even in the case where DANDi performs "worse", it outperforms SAND in terms of discovery time taking 19 % less.

Using Equation (4.8), (4.11) and (4.13), and the parameters of DANDi from Table 4.8 (assuming that $t_{token\text{-}ack} \approx 0$ or negligible), we obtain $T_{DANDi} = 50.25$ s, confirming the theoretical equations deduced.

These results show that the DANDi protocol implementation works as expected. This analysis was repeated for many different topologies confirming the results presented in this section, but we do not include them for the sake of brevity. We proceeded then to test it with real sensor nodes.

4.4.6   Evaluation with Real Nodes

We deployed the network shown in Figure 4.14 composed of ten TMote-Sky nodes equipped with 6-sectored antennas, without any particular orientation. These nodes use only one director element at a time, as in the original design. We performed the experiments running the DANDi protocol and repeated them six times. The total time required by DANDi to perform the neighbor discovery in the different experiments was 134 s, 126 s, 124 s, 113 s, 121 s and 109 s, giving an average time of 121 s.

As expected from the position of the nodes, we can observe that there are two unevenly dense areas in the network: a very dense area near nodes #1 to #6 and a lesser dense area near nodes #7 to #10.

Table 4.9 shows that nodes in the first area discover more S2S links in average (between 124 and 171) and, correspondingly take more time (between 9.805 s and 21.820 s). On the other hand, nodes in the second area discover lesser S2S links (between 46 and 79), and correspondingly take lesser time (between 3.547 s and 10.852 s). This shows how the protocol dynamically adapts to the network topology, taking more time in denser areas with more S2S links to discover, and speeding up in areas with less nodes and less links to discover.

Comparison with SAND.

For the same network topology, we run SAND protocol with different number of slots and rounds, to compare it with DANDi. We run SAND $(s = 2, r = 2)$, $(s = 3, r = 3)$ and $(s = 5, r = 5)$, where s is the number of slots and r is the number

Node	S2S links	Time (s)
#1	125	10.117
#2	137	10.984
#3	132	9.805
#4	158	12.789
#5	171	14.898
#6	124	21.820
#7	75	7.602
#8	79	10.852
#9	75	5.273
#10	46	3.547
Total	1122	107.687

Table 4.9   Average S2S links discovered and average time taken by the probe-reply mechanism for each node using the DANDi protocol.

of rounds. Table 4.10 shows the total number of S2S links discovered in the network (per node and total) for the different experiments.

Node	DANDi	SAND		
		$(s = 2, r = 2)$	$(s = 3, r = 3)$	$(s = 5, r = 5)$
#1	125	96	114	129
#2	137	119	150	160
#3	132	105	126	130
#4	158	135	156	163
#5	171	138	152	178
#6	124	74	115	146
#7	75	78	106	110
#8	79	63	79	81
#9	75	69	71	79
#10	46	38	45	52
Total	1122	915	1114	1228

Table 4.10   Total number of S2S links discovered in the network per node for the different experiments.

Table 4.11 shows the time taken by the probe-reply mechanism of each node. To compare the total time taken by each protocol, besides the time taken by the probe-reply mechanism, we have to consider the time required for the token passing present in both protocols. For SAND we also have to consider the extra time taken by the mechanism needed to locally synchronize the nodes before beginning the probe-reply mechanism. These total times are shown in Table 4.12.

Table 4.11 also shows that for the DANDi protocol, the time taken by each node to perform the discovery is different and depends on the number of neighbors each node has, showing its dynamic property. For SAND, the time taken by each

Node	DANDi (s)	SAND		
		$(s=2, r=2)$ (s)	$(s=3, r=3)$ (s)	$(s=5, r=5)$ (s)
1	10.117	4.500	10.125	28.125
2	10.984	4.500	10.125	28.125
3	9.805	4.500	10.125	28.125
4	12.789	4.500	10.125	28.125
5	14.898	4.500	10.125	28.125
6	21.820	4.500	10.125	28.125
7	7.602	4.500	10.125	28.125
8	10.852	4.500	10.125	28.125
9	5.273	4.500	10.125	28.125
10	3.547	4.500	10.125	28.125
Total	107.687	45.000	101.250	281.250

Table 4.11 Probe-reply time taken by each node of the network for the different experiments.

DANDi (s)	SAND		
	$(s=2, r=2)$ (s)	$(s=3, r=3)$ (s)	$(s=5, r=5)$ (s)
121	81	137	318

Table 4.12 Total time taken for the different experiments.

node to perform the discovery is fixed and depends on the number of rounds and slots. To ensure that a node in a dense area discovers all the S2S links, we have to force nodes in lesser dense areas to take more time to discover their neighbors.

We also see that while SAND$[s=2, r=2]$ takes less time than DANDi to finish, it fails to discover some S2S links in the denser areas. SAND$[s=3, r=3]$ takes more time than DANDi to complete, while it discovers almost the same number of S2S links than DANDi. If we compare DANDi with SAND$[s=5, r=5]$, we can see that it is 2.6 times faster but it discovers 10 % less S2S links. However, the S2S links that are not discovered by DANDi are the ones with RSSI less than the collision detection threshold (-88 dB). Those links are not considered as reliable links.

4.5 Conclusions

In this chapter, we present a detailed analysis of the state-of-the-art in neighbor discovery protocols for wireless networks with directional antennas, and we propose, implement and evaluate two novel fully directional protocols: Q-SAND and DANDi. We compare both of them with SAND, a state-of-the-art fully directional neighbor discovery protocol.

Q-SAND, an improved version of SAND that is capable of finding the combination of sectors with the highest RSSI between every pair of neighbor nodes in a network was proposed, implemented, and tested through extensive simulations. Simulation results show that Q-SAND can discover this combination of sectors up

to four times faster than SAND. Although it is faster, it has some drawbacks: i) it only works in networks where line of sight can be assumed between neighbor nodes and ii) it does not discover every S2S links in a network and, but just the ones that have the highest RSSI, and iii) it assumes that every node in the network has the same orientation. This could be suitable for many applications, but not for others as we show in Chapter 5.

DANDi, a fully directional asynchronous and dynamic neighbor discovery protocol for WSNs is proposed, implemented, and tested through extensive simulations and experiments with real nodes. The contention resolution mechanism on which DANDi is based relies on a collision detection mechanism that is also implemented and tested in simulations and experiments, achieving an FP rate of $0\,\%$ and an FN rate of $0.54\,\%$. The results of the simulations show that all the reliable S2S links in the network are discovered independently of the network topology and that the sectors that have fewer neighbors take less time to discover than the sectors that have more neighbors. The comparison with SAND shows that DANDi takes from $19\,\%$ to $78\,\%$ less time to discover the network, without having to set any parameter depending on the network topology. The experimental results confirm that DANDi can effectively discover every reliable S2S link in a given network with a better performance than SAND. While DANDi may fail to discover the weaker S2S links below the collision detection threshold, SAND may fail to discover some random S2S links if the number of slots and rounds is not high enough ($\text{SAND}[s = 2, r = 2]$ and $\text{SAND}[s = 3, r = 3]$). This could lead SAND to miss reliable S2S links. If the selected number of slots and rounds is high enough ($\text{SAND}[s = 5, r = 5]$), the time taken by the protocol is 2.6 times higher than the time taken by DANDi.

To the best of our knowledge, this makes DANDi the fastest neighbor discovery protocol in the state-of-the-art for WSN with directional antennas, with the additional advantage of being able to discover every reliable communication link in a network without requiring any prior information of the network topology.

In this chapter, we assume that the network topology and the channel conditions are static, and under these circumstances neighbor discovery needs to be performed only once at the beginning. If the channel conditions or the network topology were dynamic, we would have to repeat neighbor discovery periodically, or in an event-driven fashion.

In the following chapter, we will use the implementation of the SAND neighbor discovery protocol described in Section 4.2 to build a convergecast application. We use SAND instead of DANDi because at the time we performed the experiments we did not have a complete implementation and evaluation of DANDi. In this chapter we have deeply analyzed the neighbor discovery mechanism, and in the next one we focus on the MAC, routing and application layers.

This page was intentionally left blank.

Chapter 5

Convergecast Application in Networks with Directional Antennas

In this chapter, we introduce novel cross-layer optimizations in order to fully exploit the benefits of using directional antennas in convergecast applications. We start by presenting related work in the use of ESD antennas for different WSN applications in Section 5.2. In Section 5.3 we present the design of our modified MAC and routing protocols to support directional communication. We evaluate these protocols in Section 5.4, both in simulation and with real nodes, under different configurations and a large number of different network topologies. Our experimental results, both in simulation and with real nodes, show when the traffic is dense, networks with directional antennas can significantly outperform networks with omnidirectional ones in terms of packet delivery rate, energy consumption, and energy per received packet. Finally, in Section 5.5 we conclude the chapter.

5.1 Introduction

While there is consensus that directional antennas in WSN can provide several benefits for specific applications like high-throughput bulk forwarding [59] or radio tomographic imaging [64], the benefits of using them for convergecast applications has been questioned with diverse results. While some research shows that ESD antennas can reduce channel contention and the radio duty cycle [37, 60], other research shows that they provide limited benefits that can be only achieved under specific conditions [55].

In this chapter, we target convergecast applications on top of UDP/IPv6/RPL/ 6LoWPAN/IEEE 802.15.4 protocols. Convergecast is a typical data collection application where every node sends packets periodically to the sink node. We present a novel approach that jointly optimizes the neighbor discovery, medium access, and routing protocols to support directional communication in this kind of applications. The antenna pair selection is optimized (in terms of energy efficiency) by minimizing the number of antenna directions in use by the nodes, as this prevents unnecessary transmissions and listenings in directions with no neighboring nodes.

But from a communication perspective, the antenna pairs with better received signal strength (RSS) should be used to improve the link quality. We compare both techniques by analyzing their impact on the average Radio Duty-Cycle (RDC), Packet Delivery Rate (PDR), and Energy per Received Packet (EPRP) of the network. Finally, we combine these ideas with routing protocol optimization by minimizing the number of links that each node establishes, and modify Contiki's RPL implementation to use this information.

To evaluate these protocols, we work both in simulation and with real nodes. We analyze the different protocols through over 500 simulations with different network topologies and transmission rates, achieving up to 24% increase in PDR, 55% decrease in RDC and 46% decrease in EPRP compared to the omnidirectional reference case. To work with real nodes, we use the hardware described in Section 3.2.3 and evaluate the protocols in an office environment, performing experiments with different configurations, and confirming the benefits of using directional antennas.

The main contributions of this chapter are (i) DirMAC, the first MAC protocol for ESD antennas in WSNs, (ii) a novel technique to optimize antenna pair selection between nodes, including the proposal of a routing protocol optimization method that minimizes the number of links that each node establishes, (iii) results showing that these protocols can significantly increase the PDR and decrease the RDC compared to networks using omnidirectional antennas.

5.2   Related Work

Varshney et al. [60] show that using directional transmissions and receptions together considerably reduces channel contention by exploiting the capture effect and allowing simultaneous communication flows between multiple nodes. They do not target convergecast but predefined point-to-point transmission scenarios.

Wei et al. [64] show how directional antennas improve the localization accuracy of a radio tomographic imaging system based on WSN nodes. Großwindhager et al. [24] also show how to use directional antennas to improve the performance of an UWB-based indoor localization system. Varshney et al. [59] propose a high-throughput bulk transfer protocol that leverages ESD antennas, where data is transmitted over disjoint paths.

Mottola et al. [37] study the impact of introducing ESD antennas in WSN. They show that they can increase the performance of WSN by reducing the radio on time per delivered packet and increasing the packet delivery rate. But they also observe that, with increased network densities, parents found beyond the omnidirectional range are more likely to be affected by collisions. This is a consequence of using directionality only for packet transmissions and relying on omnidirectional communications for packet reception, MAC and routing protocols.

Another recent work regarding the use of ESD antennas in WSN is presented by Tarter et al. [55]. They make a quantitative analysis of the performance of WSN protocols in a convergecast application using directional transmissions and omnidirectional receptions, and conclude that directional antennas provide limited

benefits that can only be leveraged under specific conditions. They also state that when using ESD antennas there is an increased likelihood of hidden terminals. However, in this work nodes have different ranges for transmission and reception as the latter is done omnidirectionally, and thus doing CSMA-CA (the de-facto MAC protocol in WSN) is inappropriate as the hidden terminal problem is exacerbated when the channel is checked in a shorter range than the transmission range. Also, a fair analysis of the power consumption cannot be performed as they do not use any RDC mechanism in the MAC layer. Finally, they simulate the WSN with the Castalia simulator which does not model the hardware layers (e.g. radio hardware ACK) nor the capture effect, which is key to reducing the interference.

In this chapter, we try to overcome the limitations of these works in order to make a fair analysis of a convergecast application using directional antennas. The protocols we design use directional antennas both for transmission and reception in every layer of the stack. We incorporate directional antennas to the ContikiMAC protocol to make a proper analysis of the power consumption. We first evaluate the protocols in simulation using COOJA with the radio model presented in Section 3.2.5 that considers the capture effect, and then in a real environment with real nodes.

5.3 Design

In this section, we describe the optimization of the different protocols in the network stack to support directional communications and improve network performance. We work with the default network stack of the Contiki OS based on the IEEE 802.15.4 standard (UDP, IPv6, RPL, 6LoWPAN, ContikiMAC) described in Section 2.3 and make the necessary modifications to fully support directional communication.

5.3.1 Neighbor Discovery

In order to implement directional communication, we use the implementation of the SAND protocol described in Section 4.2.2, where a token passes through the network multiple times, and the token holder sends out probing messages in each direction, with neighbors simultaneously collecting RSS values for all receiving directions. Through this procedure, a table of RSS values for all the sector pairs is created. The neighbors determine which direction the token-holder should use (by determining which direction yields the highest RSS) for future communication, and pass this information back to the token-holder, which then passes the token on to the next reachable node that has not discovered its neighbors yet. The table of RSS values are stored by each node, and can be passed along with the token to be processed centrally. If that is the case, the sink node will have the tables of RSS values of the entire network, which can be used for offline optimizations of the MAC and routing protocol.

5.3.2 MAC

Radio communication is often the most power-consuming activity in WSN, and the energy consumption is dominated by idle listening [2, 51]. The WSN research community has put in significant effort over the past decade to reduce this with different RDC mechanisms [4, 12, 16]. However, the support for directional communication is still missing. The existing RDC mechanisms do not work outright with directional antennas as they simply do not take multiple possible antenna directions into consideration. ContikiMAC [12] is one such RDC mechanism, that keeps the radio off for almost 99% of the time and yet allows seamless communication with other sensor nodes. It is the default RDC mechanism included with Contiki.
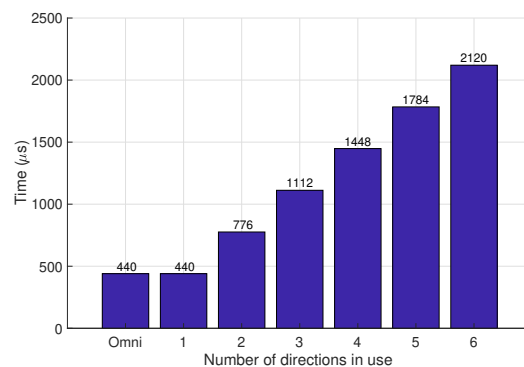
Figure 5.1 Time required to listen to the channel for the different number of directions in use. This shows how expensive it is to perform multiple CCA checks each time a node wakes up.

(a) Clear Channel Assessment

(b) Unicast transmissions

(c) Broadcast transmissions
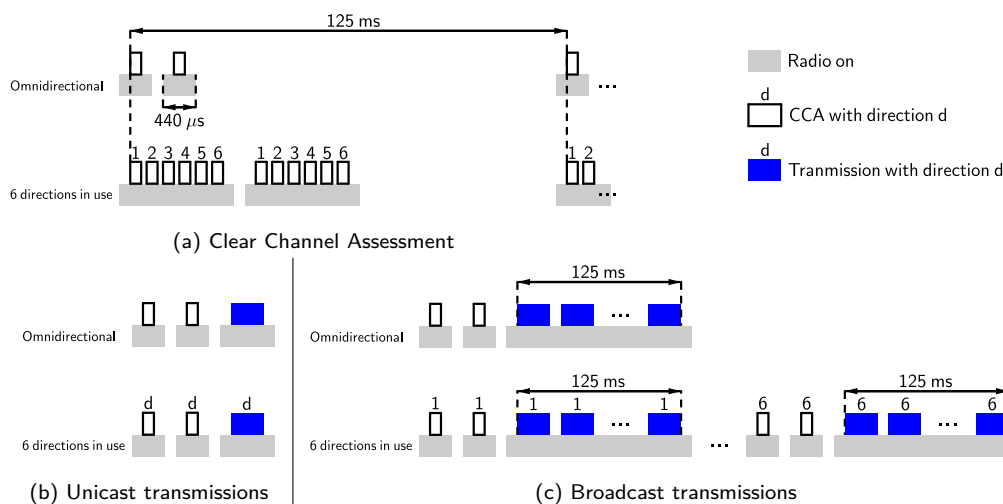
Figure 5.2 Transmissions and Clear Channel Assessments using omnidirectional antennas and ESD antennas with 6 directions in use.

We choose to base our design on an asynchronous RDC mechanism like Con-

tikiMAC over some new synchronous ones (like TSCH [14]) because in the latter interference is handled using time-slotting and channel hopping mechanisms, so using ESD antennas may just increase the transmission range (which may be equivalent to using a higher output power, for example).

In the remainder of this section, we review ContikiMAC and introduce DirMAC, a modified MAC layer protocol that uses directional antennas for the transmission and reception of packets.

Packet reception. For packet reception, the ContikiMAC protocol periodically checks the wireless channel for any ongoing transmission by performing Clear Channel Assessment (CCA) checks. These CCA checks are by default performed at a frequency of 8 Hz, corresponding to a 125 ms period. When working with ESD antennas, the simplest solution is to perform CCA checks omnidirectionally [37,55] with the main drawback that packets from nodes beyond the omnidirectional range may not be received. If the nodes listen to the channel in one given direction they might observe different channel conditions based on the antenna configuration, being deaf to transmissions in some directions while being able to receive them in others.

To support directional reception we should perform CCA checks in every possible antenna direction. Figure 5.1 shows how the listening time increases proportionally with the number of directions in use. Figure 5.2a shows the periodic channel checks when six directions are in use. In this situation, the overall energy consumption drastically increases as the radio is kept on for a longer period compared to omnidirectional antennas, and it is known that for convergecast applications the energy consumption of the CCA checks has a great impact on the overall energy consumption [51].

The proposed *DirMAC protocol* checks for incoming communications (performing CCA checks) only in the antenna directions that are in use, i.e. directions where neighbor nodes have previously been discovered. This reduces radio-on time. For example, if we have a node with a six-sectored antenna with all the neighbors aligned to one direction, this node will only have to perform CCA checks in that direction, having a radio-on time equivalent to nodes with omnidirectional antennas.

Packet transmission. For the transmission of unicast packets, ContikiMAC uses ideas from the CSMA-CA protocol, where the channel is checked twice before transmitting a packet. If the channel is clear, the node transmits the packet, and if it is busy, it waits for a random back-off time and tries again. In related work CCA checks are performed omnidirectionally [37, 55]. Figure 5.3a shows how in these scenarios, the hidden terminal problem is exacerbated, and thus collisions increase.

Figure 5.2b shows how the *DirMAC protocol* supports the directional transmission of unicast packets by sending or forwarding each packet to the receiver node using the appropriate antenna direction. The phase information from the neighbors is obtained in the same way as in ContikiMAC, so the only change needed to support directional communication is to select the correct antenna direction before transmitting the packet. Then the channel will be checked in that direction before

transmitting, mitigating the hidden terminal problem that Tarter et al. [55] consider a major drawback of using ESD antennas in WSN, as shown in Figure 5.3b. When the parent transmits, the receiver node will be listening to the channel and scanning through its sectors and thus it will receive the packet correctly.
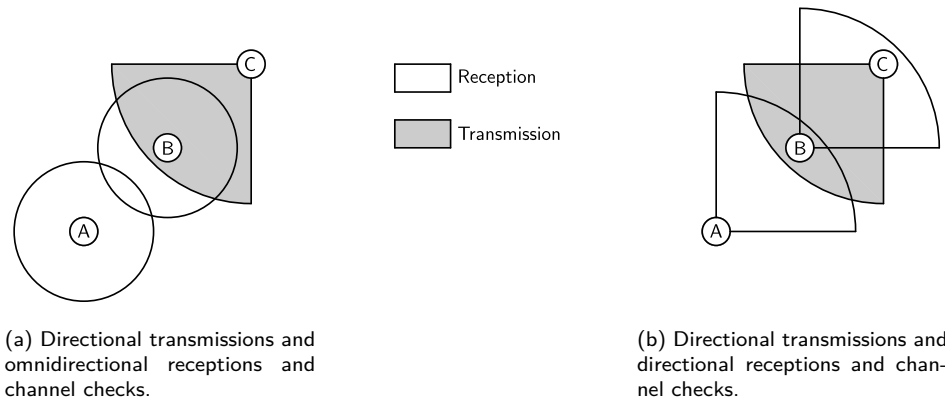

(a) Directional transmissions and omnidirectional receptions and channel checks.

(b) Directional transmissions and directional receptions and channel checks.

Figure 5.3   In both scenarios, node C is transmitting to node B and node A is performing channel checks. In the scenario shown in Figure 5.3a, node A will fail to hear node C's transmission, so it will transmit resulting in a collision (hidden terminal problem). Using DirMAC will result in the scenario shown in Figure 5.3b, where node A will hear node C's transmission and will not transmit, avoiding the collision.

For the transmission of broadcast packets, in ContikiMAC each node checks if the channel is clear by using the same procedure as with unicast packets, and then transmits the packet during a whole listening period (called strobe time), to ensure that every node has the possibility to hear it.

Figure 5.2c shows how the *DirMAC protocol* supports the transmission of directional broadcast messages by sending each packet during a whole listening period in every direction in use. This means that for a node with N directions in use, each broadcast packet is going to be transmitted during $N \times strobe\_time$ seconds (N times longer than with omnidirectional antennas). Before transmission, the nodes check that the channel is clear in the selected direction.

Both for unicast and broadcast transmissions, multiple directions can serve the same neighbor with different link qualities, so we could optimize the selection of the different directions that the nodes use. The number of directions in use by each node has a great impact on the radio-on time and thus on overall power consumption. There is also a trade-off between network performance and the number of directions in use by the nodes. The more directions we use, the stronger the links can be, but power consumption increases.

The relationship between the network performance and antenna directions in use is complex to model, especially as the network size increases. In Section 5.3.3, we tackle this problem heuristically.

5.3.3 Antenna Pair Selection

In this section we propose four different heuristics to optimize the antenna pair selection between neighboring nodes: BestSS_Dist, MinDir_Cent, MinDirOpt_Cent, and BestSSOpt_Dist. The objective is to find the mechanism of antenna pair selection that achieves the best trade-off between the highest overall PDR in the network and the lowest overall RDC and EPRP. The MinDir_Cent and MinDirOpt_Cent antenna pair selection mechanisms are centralized and require offline processing, while the BestSS_Dist and BestSSOpt_Dist are distributed and can be performed online. A distributed approach has many benefits that enable the use of these protocols in real networks. First of all, these heuristics can be repeated periodically to deal with changing channel conditions or movement of nodes. Also, having to process routing or neighbor discovery information offline and transmitting this information back to the nodes is impractical in some deployments.

BestSS_Dist Heuristic. The BestSS_Dist heuristic determines the choice of the direction to transmit to every neighbor node by choosing the direction which maximizes the signal strength at the receiver. This is a completely distributed method of choosing directions, and no global information is used to minimize the total number of directions in use.

MinDir_Cent Heuristic. MinDir_Cent is a centralized heuristic to find the antenna combinations that enable communication between neighboring nodes while minimizing the number of directions in use by each node. This heuristic cannot be implemented locally as some communication links may only be established when both transmitter and receiver are using some given directions.

In our implementation, computations are performed offline using the RSS tables collected during Neighbor Discovery. The heuristic produces a list of possible links and the directions that each node has to use. For each node, any neighbor that can only be reached through one combination of antenna directions is designated as a bad neighbor and avoided. Bad neighbors are avoided because they force the heuristic to use a single available combination of antenna directions for communication between the node and the given neighbor, and hence limit the freedom to minimize the total number of directions in use.

Optimization begins by placing all the nodes in the set S_n of non-optimized nodes. The heuristic then iterates through the nodes in the set S_n one-by-one in a random order and evaluates whether there is a single direction in which the node can reach all of its neighbors, except bad neighbors, based on the table of RSS values and given neighbors' directions in use. If such a direction exists, the node is removed from the set S_n. Otherwise, the heuristic goes on to check the next node. After it has gone through all the nodes, the heuristic goes through the remaining nodes in S_n one-by-one and evaluates whether there exist combinations of two directions in which the nodes can reach their respective neighbors. Each time the heuristic has gone through the set S_n, it goes through the set again but allows for one more direction to be used by the nodes.

71

After following these steps, the set S_n is empty. The combinations of directions in use chosen for all the nodes are saved as a possible solution. All nodes are then put back into the set S_n, and their choice of directions in use kept except for one randomly chosen node where the choice is reset. The heuristic then iterates through the set S_n again in the same manner as before. When the set S_n is empty again, the new combinations of directions in use are compared to the previous solution. If they are the same, the optimization is complete. Otherwise, the heuristic saves the new solution and repeats the optimization.

MinDirOpt_Cent Heuristic. In the MinDirOpt_Cent heuristic, ideas from MinDir_Cent are combined with offline routing optimization which allows us to decrease the number of links in the network. The offline routing optimization yields *a priori* knowledge of the links that have to be established between parents and children. The number of directions in use by every node is minimized only considering these links.

The heuristic is centralized, performed offline, and has two outputs: the choice of the parent for each node, and which directions to use in each node in order to minimize the total number of directions in use in the network. The choice of node parents is used by RPL as the next hop to the sink, overriding the default metrics.

The heuristic starts by using a metric similar to RPL's hop count. It places all nodes that can reach the sink in Tier 1. All nodes not in Tier 1, but within reach of a Tier 1 node, are placed in Tier 2, and so on until all nodes have been placed in a tier in which they can reach a node in the tier below. The choice of parent is made by selecting the node in the lower tier with the strongest link, i.e. the highest RSS. This tier-based system minimizes the number of hops to the sink. After all the nodes have chosen their parent, the heuristic uses the same procedure as in MinDir_Cent to find the smallest number of directions each node needs to use to reach its parent and all of its children.

BestSSOpt_Dist Heuristic. In BestSSOpt_Dist, we combine ideas from BestSS_Dist and MinDirOpt_Cent heuristics to reduce the number of directions in use in a distributed way. BestSSOpt_Dist starts with a neighbor discovery phase where the nodes collect the RSS tables, and store the direction that maximizes the signal strength with each neighbor. After neighbor discovery is complete, we use Contiki's default RPL implementation to select the parent and children of each node in a distributed way using the ETX metric. During this phase, the nodes communicate selecting the antenna pairs that maximize the signal strength between each other, obtained during neighbor discovery. When the routing phase is complete, the nodes process the RSS tables and discard the directions used only to communicate with neighbors that are neither parent nor children. After they are discarded, the DirMAC protocol stops listening and transmitting packets in these directions. This is a completely distributed approach to choose the directions that maximize the signal strength while reducing the number of directions in use.

A summary of the four heuristics is presented in Table 5.1.

Table 5.1   Summary of the different heuristics.

Heuristic	Type	Optimization	Remarks
BestSS_Dist	Distributed	Maximize signal strength.	Each node selects the strongest link to communicate to each neighbor node.
MinDir_Cent	Centralized	Minimize directions in use.	Algorithm minimizes the overall number of directions in use using RSS tables from every node.
MinDirOpt_Cent	Centralized	Minimize links and directions in use.	Combines MinDir_Cent heuristic with routing information (minimizing hops to the sink) to minimize directions in use.
BestSSOpt_Dist	Distributed	Maximize signal strength and minimize links in use.	Nodes only communicate with RPL parent and children, selecting the strongest links to communicate with each one.

5.4   Evaluation

We first evaluate the performance of the protocols proposed in Section 5.3 with over 500 simulations, and then with experiments on real nodes with real antennas. With the simulations we can test many different network topologies and node densities, and get initial insights on the performance of the different protocols. The experiments with real nodes allow us to confirm these results and test them in real-world conditions.

The purpose is to assess the benefits of using directional antennas in a WSN convergecast application, and compare the results with omnidirectional protocols. Therefore, we obtain the average PDR, RDC, and EPRP of all nodes in the network. We select these magnitudes as representative of the performance of the network protocols. Due to the additional energy consumption required to perform CCA checks, it is not evident that our heuristics for directional antennas will reduce the energy consumption compared to convergecast with omnidirectional antennas. In this kind of applications we want to minimize the energy consumption (and thus the RDC), while maximizing the PDR to collect as much data as possible from the sensors.

In every experiment, the sink is assumed to be connected to a main power

Table 5.2   Values used in the different simulations.

	Values
Protocol	Omnidir, BestSS_Dist, MinDir_Cent, MinDirOpt_Cent, BestSSOpt_Dist
Network density	10, 15, 20, 25, 30 nodes/$10.000\,\text{m}^2$
Transmission rate	2, 5, 10, 20 pkt/minute

supply, i.e. it does not suffer from the same power constraints as the other nodes. Therefore, it does not need to keep its radio off to minimize idle listening. This allows it to continuously scan all directions in use for incoming transmissions. However, the continuous scanning diverts CPU time from packet processing, which means the sink does not have time to process incoming packets before they are over-written in the sink's memory by packets received later. Hence, a delay is added, forcing the sink's antenna direction switching period to 110 μs.

5.4.1   Experimental Setup for Simulations

We evaluate a convergecast scenario, where every node sends packets periodically to the sink node with a fixed rate. The transmission rate determines the application throughput of the network and it is varied according to the values in Table 5.2. The table also shows the analyzed protocols and the different network densities. We use the Contiki implementation of a data collection protocol based on Collect (similar to Collection Tree Protocol [22]) on top of our modified and optimized protocols. The payload of the packets contain a sequence number, the nodes' duty cycles (estimated with the Energest module) and the RPL parent, among others.

We run over 500 experiments in a noise-free environment with networks of 10 to 30 nodes pseudo-randomly distributed over a bounded area. We follow Varshney et al. [59] by placing the first node ($i = 0$) completely randomly and then placing the subsequent $(i + 1)^{\text{th}}$ node at a distance d_i from the i^{th} node, where d_i follows a Gaussian distribution with a mean three quarters of the omnidirectional communication range and a standard deviation of half this range. We work with up to 30 nodes because with larger networks the COOJA simulator takes too much time to complete the simulations. We enforce a minimum distance between all nodes, as well as a maximum distance to enable communication with omnidirectional antennas. We bound our environment to 100×100 meters to be able to simulate high-density networks with a limited number of nodes. Five different network topologies are generated for each network density. Nodes run the SAND neighbor discovery protocol until it completes, and then they run for four minutes before starting the collect application, giving the routing protocol time to stabilize. Subsequently, the nodes periodically send data to the sink. We repeat each experiment and vary the transmission rate, using the different rates listed in Table 5.2. Each experiment runs for 30 simulated minutes, giving the RDC and PDR time to stabilize. We repeat the simulations for the five protocols listed in Table 5.2. For the experiments with omnidirectional antennas, we consider a constant antenna

(a) Transmission rate of 2 pkt/minute.

(b) Transmission rate of 5 pkt/minute.

(c) Transmission rate of 10 pkt/minute.

(d) Transmission rate of 20 pkt/minute.
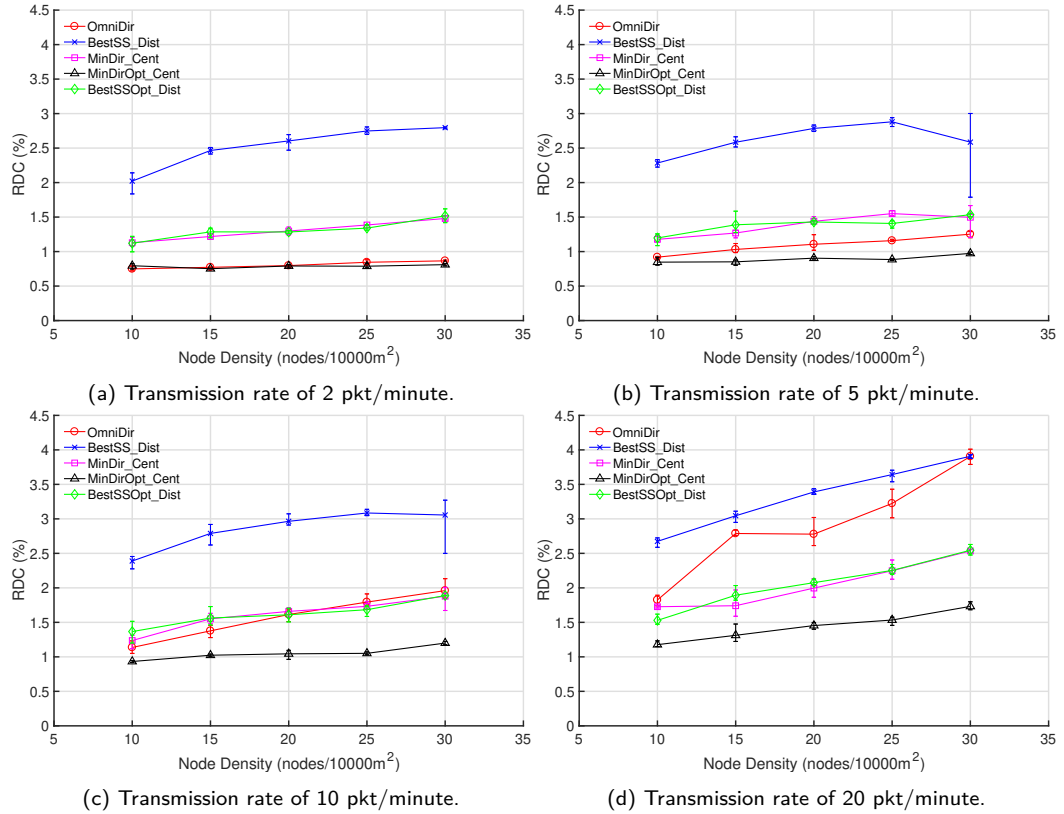
Figure 5.4   Average radio duty cycles for the different transmission rates. Most directional heuristics perform better than OmniDir for the higher transmission rates, while they perform similar or worse for the lower transmission rates.

gain of 3 dB, and we use the default implementation of ContikiMAC and RPL.

5.4.2   Results of Simulations

Energy Consumption. In the first set of experiments, we evaluate the potential energy savings possible with directional antennas. As a proxy for energy consumption we use the RDC as other researchers have also done [37].

Figure 5.4 shows the nodes' average RDC for different transmission rates, where for each point we plot the maximum, minimum, and average over the different network topologies. Figure 5.4a shows that when the traffic is light, both OmniDir and MinDirOpt_Cent achieve a duty cycle below 1%. The result for OmniDir is in line with previous results for RPL [15]. The other heuristics for directional antennas, however, have much higher duty cycles due to the overhead of additional CCA checks. Using directional antennas and the MinDirOpt_Cent heuristic, the radio duty cycle remains low, below 2% even when the traffic load is high (Figure 5.4d) while it drastically increases with omnidirectional antennas. Figure 5.4d also shows that directional antennas using the BestSS_Dist heuristic have a higher energy consumption than using omnidirectional antennas even in high traffic load scenarios. One of the reasons for this high energy consumption is the high number
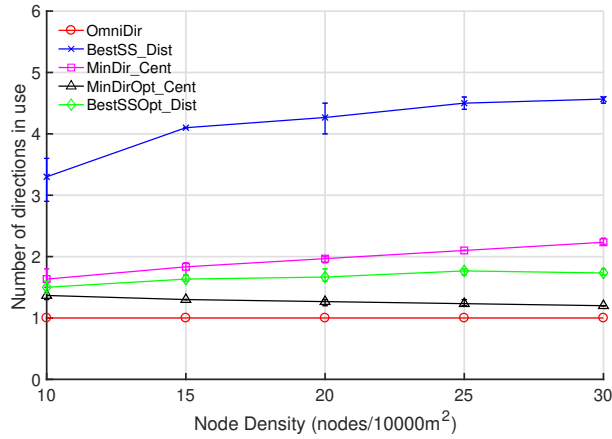
Figure 5.5   Average number of directions in use by each heuristic, for each network density. MinDirOpt_Cent and BestSSOpt_Dist heuristics perform very similar to the omnidirectional reference case, while MinDir_Cent and BestSS_Dist use more directions in average.

of average directions in use compared to the other heuristics, as shown in Figure 5.5. Note that for each direction in use, a CCA check needs to be performed. Although the number of hops to the sink decreases by one hop on average when using directional antennas, this is not enough to compensate for the additional CCA checks that force BestSS_Dist to keep the radio on for a longer time.

Figures 5.4a and 5.4b show scenarios with lower transmission rates and hence lighter traffic. The figures show that using the MinDir_Cent and BestSSOpt_Dist heuristics also results in a higher RDC than using omnidirectional antennas. Also in these scenarios, the RDC is dominated by idle listening and thus the reduction in the number of directions in use shown in Figure 5.5 is not enough to outperform omnidirectional communication. The MinDirOpt_Cent heuristic, however, has lower energy consumption than OmniDir. The reason for this is that MinDirOpt_Cent uses in average only slightly more directions than one. Hence its reception RDC is only slightly higher than OmniDir's. MinDirOpt_Cent's transmission RDC, however, is smaller than OmniDir's due to reduced interference from neighboring nodes achieved with directional antennas. For a network with 30 nodes and a transmission rate of 2 packets per minute, the reception RDC of MinDirOpt_Cent is higher than the RDC of OmniDir (0.76% against 0.75%), but its transmission RDC is significantly lower (0.053% against 0.12%).

Figures 5.4c and 5.4d show that as we increase the transmission rate and hence the traffic density, also the MinDir_Cent and BestSSOpt_Dist heuristics outperform OmniDir. As we increase the transmission rate, transmissions dominate the power consumption and since directional antennas decrease the interference from neighboring nodes, the overall RDC of the MinDir_Cent and BestSSOpt_Dist heuristics decreases. This effect is more noticeable with the MinDirOpt_Cent heuristic, where the decrease of the transmission RDC due to the decreased interference is combined with the decrease of the reception RDC due to the reduction of the number of directions in use.

For a network with 30 nodes and a transmission rate of 20 packets per minute,

(a) Transmission rate of 2 pkt/minute.

(b) Transmission rate of 5 pkt/minute.

(c) Transmission rate of 10 pkt/minute.

(d) Transmission rate of 20 pkt/minute.
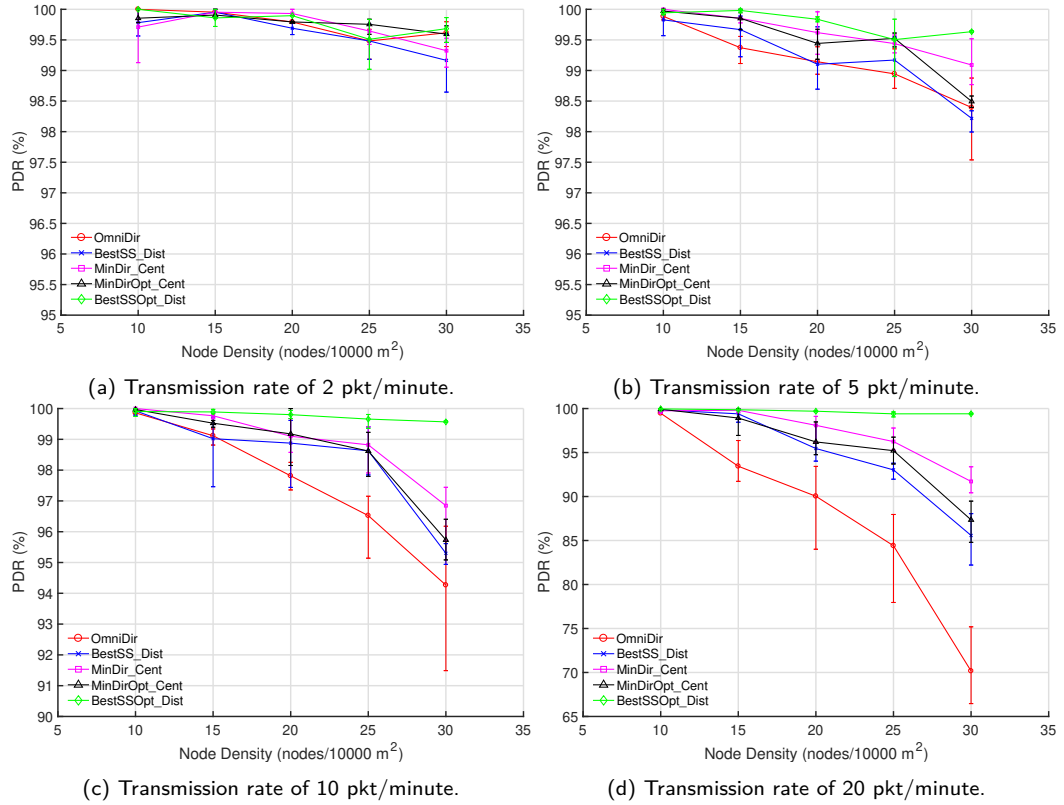
Figure 5.6   Average packet delivery rates for the different transmission rates. Directional heuristics perform similar to OmniDir for the lower transmission rates, and significantly better for the higher transmission rates.

the RDC of MinDirOpt_Cent is 55% lower than in the omnidirectional case (50% lower in reception and 65% lower in transmission), while the RDC of BestSSOpt_Dist is 35% lower than the reference omnidirectional case (37% lower in reception and 32% lower in transmission). The decrease in the reception RDC can be explained by the effect of overhearing [2], where nodes drain energy by receiving irrelevant packets destined to other nodes, and this effect is exacerbated in dense networks. Nodes using directional antennas receive fewer packets destined to other nodes as the radio energy is concentrated in a single direction.

Packet Delivery Rate. The PDR is an important metric for convergecast, as it shows the percentage of data packets that the sensor nodes can effectively deliver to the sink. Figure 5.6 shows the average PDR for the different transmission rates. Figures 5.6a and 5.6b present the results when the traffic transmission rates are low and hence the traffic is light. In these scenarios, all the heuristics perform similar to the omnidirectional case and achieve a PDR of above 99% for transmission rates of two packets per minute and above 98% for transmission rates of five packets per minute.

When the transmission rate increases, as shown in Figures 5.6c and 5.6d, all heuristics for directional antennas achieve a higher PDR than OmniDir. The reason is the reduction of the interference achieved with the directional antennas.

(a) Transmission rate of 2 pkt/minute.

(b) Transmission rate of 5 pkt/minute.

(c) Transmission rate of 10 pkt/minute.

(d) Transmission rate of 20 pkt/minute.
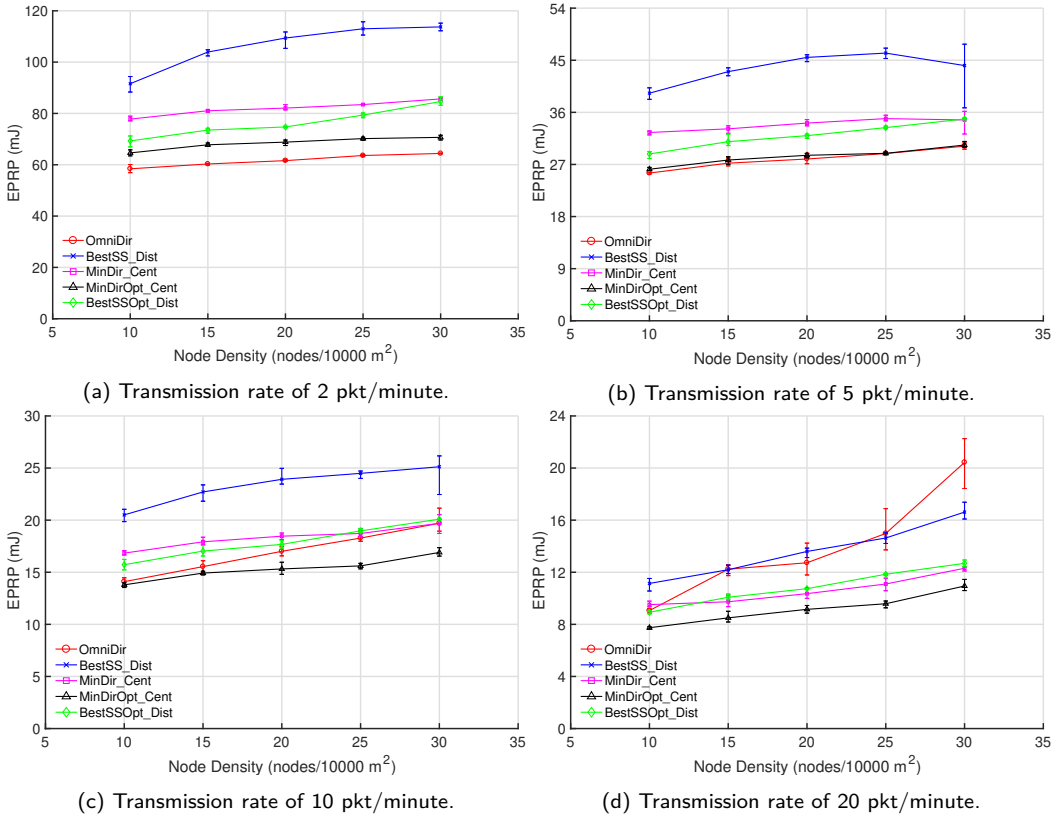
Figure 5.7   Average energy per received packet for the different transmission rates. Most directional heuristics perform better than OmniDir for the higher transmission rates, while they perform similar or worse for the lower transmission rates.

In the most dense scenario with 30 nodes and a transmission rate of 20 packets per minute, MinDirOpt_Cent's PDR is above 85% (up to 24% higher than OmniDir's), while BestSSOpt_Dist's PDR is above 99% (up to 32% higher than OmniDir's).

Energy per received packet. The energy per received packet (EPRP) combines PDR and RDC and is a good metric to assess the performance of convergecast applications. It is calculated as the total energy consumed by each node, divided by the number of packets received by the sink from that node. We average these values over every node in the network. Figure 5.7 shows the EPRP for the different transmission rates. This energy is calculated as the total energy consumed by each node, divided by the number of packets received by the sink.

Figures 5.7a and 5.7b show that in scenarios with light traffic, the performance difference between MinDirOpt_Cent and OmniDir is low, while BestSSOpt_Dist and MinDir_Cent have a slightly higher consumption. When traffic is dense, however, the MinDirOpt_Cent heuristic consumes up to 46% less energy per received packet than the omnidirectional reference case, and BestSSOpt_Dist consumes up to 38% less EPRP than OmniDir, demonstrating the overall benefit of using directional antennas with cross-layer optimizations.

Table 5.3   Values used in the different experiments with real nodes.

	Values
Protocol	Omnidir, BestSS_Dist, BestSSOpt_Dist
Network density	10 nodes/10.000 m^2
Transmission rate	2, 5, 10, 20, 60 pkt/minute

5.4.3   Experimental Setup for Real Nodes

We evaluate the same convergecast application as in the simulations, where information is sent to the sink with a fixed rate, and we vary the transmission rate and protocols according to the values in Table 5.3. We evaluate the distributed heuristics BestSS_Dist and BestSSOpt_Dist instead of the centralized ones (MinDir_Cent and MinDirOpt_Cent) because we could achieve similar performance results in simulation and they are better suited to implement in a real network. As we could not increase the node density due to the limited availability of antennas, we added a new higher transmission rate of 60 packets per minute to the evaluation, in order to consider a dense traffic scenario.

We deploy a network of 10 Tmote-Sky nodes with the 6-sectored SPIDA antennas described in Section 3.2.3, without any particular orientation. These nodes use only one director element at a time, as in the original design. We perform the experiments in an office environment with moderate WiFi interference and the nodes are randomly distributed over a bounded area and run for 20 minutes. As the network in this evaluation has 10 nodes, the routing protocol stabilizes in less than one minute.

5.4.4   Results of Experiments with Real Nodes

Before analyzing the performance of the different heuristics, we validate the experiments with real nodes by comparing them with simulations under similar scenarios. To achieve this, we run simulations according to the parameters in Table 5.3 with the same distribution of nodes as in the experiments with real nodes.

Results comparing simulations and experiments with real nodes are shown in Figures 5.8, 5.9 and 5.10. Figures 5.8a and 5.8b show that in both cases the RDC tends to grow linearly with the transmission rate, and that BestSS_Dist performs significantly worse than OmniDir and BestSSOpt_Dist. Figures 5.9a and 5.9b show that PDR is above 96% for every experiment, while the directional heuristics perform better than OmniDir in every scenario. Finally, Figures 5.10a and 5.10b show that both in simulation and experiments with real nodes, the EPRP is inversely proportional to the transmission rate. These results validate our simulation methodology.

Energy Consumption. Figure 5.8b shows the average RDC of the nodes for the different scenarios. It shows that BestSSOpt_Dist always performs similar or better than OmniDir for every transmission rate, while BestSS_Dist has a significantly higher power consumption. For a transmission rate of 60 packets per minute, the RDC of BestSSOpt_Dist is 25% lower than in the omnidirectional case (16% lower

79

(a) Simulation.

(b) Real nodes.
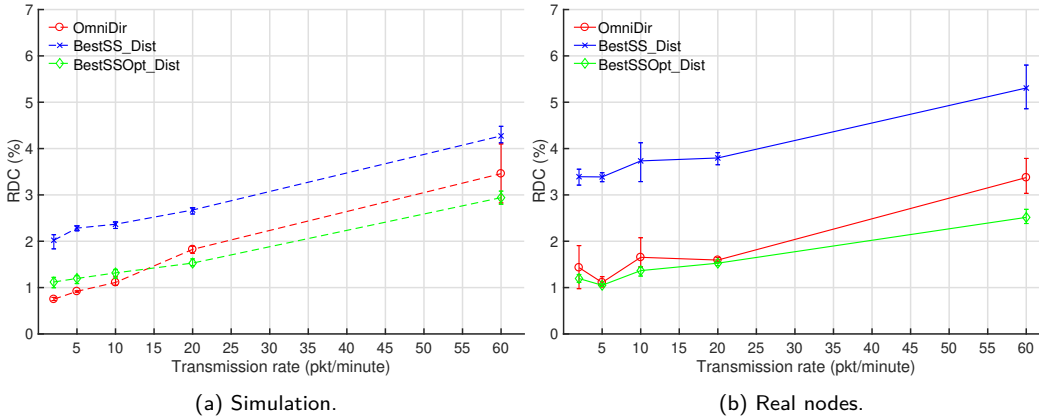
Figure 5.8   Radio Duty Cycle for a density of 10 nodes/10.000 m^2.

in reception and 43% lower in transmission). The reduction in the reception RDC can also be explained by the effect of overhearing in Omnidir heuristic, while the reduction in the transmission RDC is due to reduced interference from neighboring nodes achieved by using directional antennas, meaning fewer transmissions.

Packet Delivery Rate. Our results from Figure 5.9b show that for a network of 10 nodes, the PDR is above 96% for every heuristic and every transmission rate. Nevertheless, the figure shows that for OmniDir the PDR drops in some experiments due to interference, while BestSSOpt_Dist has a PDR of almost 100% in every scenario. This shows that the robustness against interference, enabled both by the capture effect and the reduced interference between nodes due to the directional radiation pattern, is a key feature of using directional antennas in WSN.

Energy per received packet. Figure 5.10b shows that the EPRP decreases as the transmission rate increases. This result is expected because there are very few packet losses for every scenario, so as we increase the transmission rate, the overall energy consumed by the nodes is divided by a greater number of received packets. Even though the RDC due to transmission and reception of packets increases for higher transmission rates, the power consumption due to idle listening is still constant and dominates the overall power consumption. This figure shows that the EPRP for BestSSOpt_Dist is always below the EPRP for OmniDir, which confirms all the previous results. For the higher transmission rates, BestSSOpt_Dist consumes 15% less EPRP than OmniDir.

5.5   Conclusions

The benefits of using directional antennas for convergecast in WSN are not clear in the existing literature. To the best of our knowledge, we are the first to jointly optimize the neighbor discovery, medium access and routing protocols to support directional communication in WSN in a faithful simulation scenario while improving the performance of the network for a convergecast application.

(a) Simulation.

(b) Real nodes.
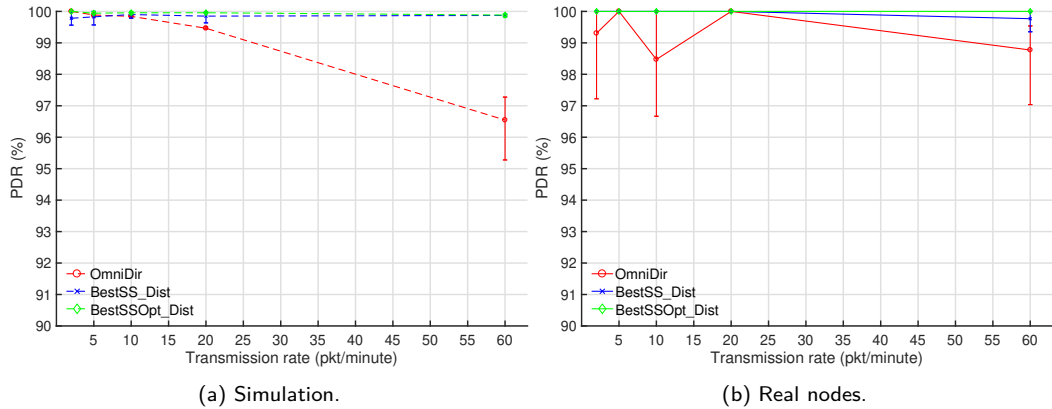
Figure 5.9   Packet Delivery Rate for a density of 10 nodes/10.000 m^2.

(a) Simulation.

(b) Real nodes.
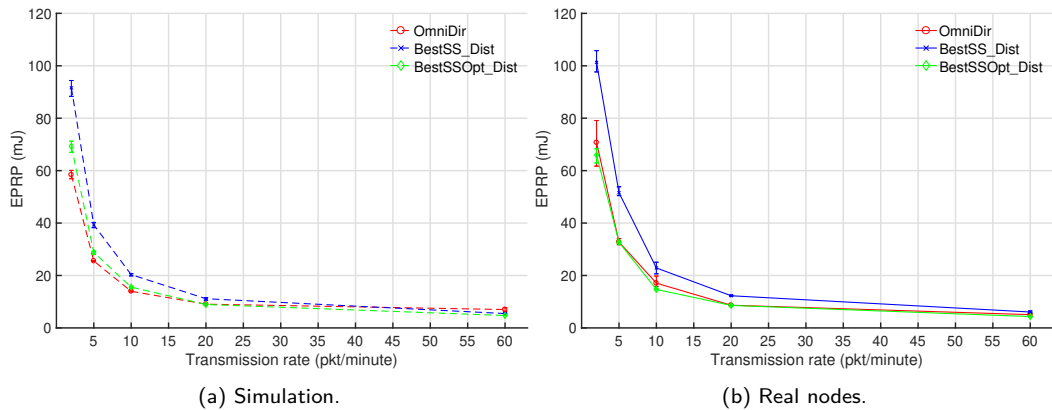
Figure 5.10   Energy per Received Packet for a density of 10 nodes/10.000 m^2.

We design and implement DirMAC, a new MAC protocol that fully supports directional communication. We also propose four different heuristics to optimize the performance of the protocols. Our evaluation shows that optimizations at both the MAC and routing layers are needed in order to reap the benefits of using directional antennas for convergecast. The best results are obtained when we reduce the number of antenna directions used by the nodes, and reduce the number of links that each node establishes (MinDirOpt_Cent and BestSSOpt_Dist heuristics), the latter having the added major benefit of being completely distributed and with no need for offline processing. We evaluate the performance of these protocols in simulation and with real nodes under different application scenarios. Our results show that the performance of the network can be greatly improved and we obtain the largest performance improvements in networks with dense traffic. Simulations with different node densities show that when using directional antennas the PDR increases up to 29%, while energy consumption and energy per received packet decreases by up to 55% and 46% respectively. Experiments with real nodes validate these results showing a significant performance increase when using directional antennas in our scenarios, with a reduction in the RDC and EPRP of 25% and 15% respectively, while maintaining a PDR of 100%.

This page was intentionally left blank.

Chapter 6

Conclusions and Future Work

In this chapter, we conclude the work and we present a general discussion of the research problem addressed in this thesis, before discussing future directions of work.

6.1 Conclusions

In this thesis, we first addressed the modeling, fabrication, and characterization of a six-element SPIDA antenna, presented in Chapter 3. As far as we know, it is the first complete characterization reported of this kind of antenna including the S_{11} parameter. The characterization shows that the antenna has a maximum gain of 6.8 dBi, an HPBW of 113° and a module of S_{11} parameter of -7.5 dB at the central frequency ($f_c = 2.4525$ GHz). We also presented a model of this antenna for the COOJA network simulator, based on our previous characterization. This model enables the simulation of large networks, where every node is equipped with a SPIDA antenna, allowing the evaluation of the network protocols presented in the following chapters. Even though we implemented a model for a six-element SPIDA antenna, it can easily be modified to change the number of sectors and the resulting radiation pattern.

In the same chapter, we presented a novel way to optimize the antenna without changing its geometry by using multiple director elements. We show that by using multiple director elements we can improve the performance in terms of maximum gain, narrower HPBW, and a lower module of the S11 parameter without making any changes in the antenna itself. The configuration with three consecutive director elements presents a maximum gain of 8.2 dBi (against 6.8 dBi for the original design), and a module of the S_{11} parameter of -9.8 dB at the central frequency (against -7.5 dB for the original design). The configuration with two consecutive director elements shows an increased performance, with the additional feature that the main lobe direction is in the middle of the two directors, so with a six-element antenna it would be possible to direct the main beam in 12 different directions. The configurations with opposed director elements show the potential to direct RF power to opposite directions with the benefit of enhanced performance.

The possibility to change the configuration of the antenna dynamically makes the optimizations suitable for a wide variety of applications and network topologies.

In this thesis we also studied the state-of-the-art in neighbor discovery protocols for WSNs with directional antennas, presented in Chapter 4. We also propose, implement end evaluate two novel fully directional protocols: Q-SAND and DANDi. We compare both of them with SAND, a state-of-the-art fully directional neighbor discovery protocol.

Q-SAND is an improved version of SAND that is capable of finding the combination of sectors with the highest RSSI between every pair of neighbor nodes in a network. Results show that Q-SAND can discover these combinations of sectors up to four times faster than SAND. Although it is faster, it has two drawbacks. The first one is that it assumes that the best combination of sectors between two nodes is with their antennas pointing at each other, which is only valid in networks where there is line of sight between neighbor nodes. The second one is that it does not discover every S2S link in a network, but just the ones that have the highest signal strength.

DANDi is a fully directional asynchronous and dynamic neighbor discovery protocol for WSNs. In this protocol, the contention resolution relies on a collision detection mechanism that is also implemented and tested in simulations and experiments with real nodes, achieving an FP rate of 0% and an FN rate of 0.54%. The results show that every reliable S2S link in the network is discovered independently of the network topology and that the sectors that have few neighbors take less time to discover than the sectors that have more neighbors. The comparison with SAND shows that DANDi takes from 19% to 78% less time to discover the network, without having to set any parameter for the network topology. One drawback of DANDi is that it may fail to discover the weaker S2S links below the collision detection threshold. In comparison, SAND may fail to discover some random S2S links if the number of slots and rounds is not high enough, which could lead to miss reliable S2S links. This makes DANDi the fastest neighbor discovery protocol in the state-of-the-art for WSN with directional antennas, with the additional advantage of being able to discover every reliable communication link in a network without requiring any prior information of the network topology.

We also studied the benefits of using directional antennas for convergecast in WSN, that are presented in Chapter 5. To the best of our knowledge, we are the first to jointly optimize the neighbor discovery, medium access and routing protocols to support directional communication in WSN while improving the performance of the network for a convergecast application.

We present DirMAC, a novel MAC protocol that fully supports directional communication together with four different heuristics to optimize the performance of the protocols. Our evaluation shows that optimizations at both the MAC and routing layers are needed in order to reap the benefits of using directional antennas for convergecast. The best results are obtained when we reduce the number of antenna directions used by the nodes and reduce the number of links that each node establishes (MinDirOpt_Cent and BestSSOpt_Dist heuristics), the latter having the added major benefit of being completely distributed and with no need for

offline processing. Our results show that the performance of the network can be greatly improved and we obtain the largest performance improvements in networks with dense traffic. Simulations with different node densities show that when using directional antennas the PDR increases up to 29%, while energy consumption and energy per received packet decreases by up to 55% and 46% respectively. Experiments with real nodes validate these results showing a significant performance increase when using directional antennas in our scenarios, with a reduction in the RDC and EPRP of 25% and 15% respectively, while maintaining a PDR of 100%.

These contributions enable the integration of electronically switched directional antennas such as the SPIDA to WSN applications, while improving the overall performance of the network. We present optimizations to the antenna itself, and also design and implement cross-layer optimizations that can be used in any WSN deployment.

6.2 Discussion

In this thesis, we have analyzed how to improve the performance of WSNs by using directional antennas. In Chapter 3, we have presented the characterization, fabrication, and optimization of a simple ESD antenna for WSN: the SPIDA antenna. We have verified that it can provide a higher gain than the off-the-shelf omnidirectional antennas embedded in sensor nodes, and also an improved directivity.

The most direct consequence of an increased gain is that nodes can achieve a higher signal strength at the receiver without increasing the radiated power. But, does this have a significant impact on the performance of a WSN? An increased gain would allow nodes to extend their communication range, which can be very beneficial for many applications. Extending the range can allow the creation of new communication paths that can reduce the number of hops between two distant nodes, or it can simply allow the communication between two nodes that were previously out of range. But in dense networks where the nodes generally have several communication paths to reach a certain node, the stronger links are chosen for communication and this extended range does not have a significant impact. We know from our own experience and related work [48] that an increased signal strength does not always translate into a better PDR. When we have weak links where the RSSI is near the sensitivity of the radio, a higher gain can result in an increased PDR. But if we have strong links, an increased RSSI will probably not affect the PDR. Therefore, the benefit of the higher gain would be more noticeable in networks forced to used weak links, for example, if we have a network topology where the nodes are separated by a distance close to the maximum communication range.

Another consequence of increasing the gain is that we could lower the output power to reduce the power consumption while maintaining the RSSI. The first thing to note here is that the power consumption of the radio is not directly proportional to the output power. A reduction from 0 dB to -3 dB in the output power results in a cutback of just 10% in the transmission power consumption [7]. We also know from our experience [51] that the power consumption in WSNs using

RDC mechanisms such as ContikiMAC is dominated by idle listening, and though reducing the output power would indeed reduce the power consumption, its impact on the overall power consumption would be negligible.

But besides the increased gain, the directional antennas provide a narrower beamwidth and the possibility to switch between directions. Concentrating the radiated power in a certain direction reduces the interference between neighbor nodes, which results in fewer collisions and thus a decreased transmission duty cycle that is noticeable in networks with dense traffic. The increased gain and reduced interference enables simultaneous communication paths that are possible thanks to the capture effect. This reduced interference also generates a decrease in the reception RDC, explained by the effect of overhearing, where nodes drain energy by receiving irrelevant packets destined to other nodes. Nodes using directional antennas receive fewer packets destined to other nodes as the radio energy is concentrated in a single direction. This effect is exacerbated in very dense networks. In Chapter 5 we show that these effects combined can generate a significant increase in the performance of the network, resulting in an increased packet delivery rate, and a decreased energy consumption and energy per received packet.

So going back to the initial question: can we improve the performance of WSNs by using directional antennas? The answer would be yes, but not always. WSNs have countless possible applications with very different requirements. While some may demand low latency, others may require high reliability, high throughput, increased battery autonomy, or any combination of the above. Over the years there has been a significant effort to improve the network protocols and tune its parameters to fit numerous applications, and it is clear that there is no one-fits-all solution. And directional antennas are not the exception. Even though they may not be a good fit for several applications, they have shown to provide significant benefits for several other applications like bulk data transmission [59], radio tomographic imaging [64] or data collection applications over dense networks with dense traffic.

6.3   Future Work

In this thesis, we presented contributions to optimize a SPIDA antenna, improve the neighbor discovery mechanism, and improve the performance of a data collection application by performing cross-layer optimizations. Nevertheless, not all of these contributions were combined. When we analyze the performance of a data collection application, we use the SAND protocol for neighbor discovery instead of DANDi, that showed to find all the reliable S2S links in a reduced time. For this analysis, we also did not use three consecutive director elements in the SPIDA antennas, that showed to provide a higher gain and narrower beamwidth. The main reason why we did not combine some of these contributions is that they were not developed in the same order as they are presented in this thesis. Future work should include evaluating a similar data collection application, but including the DANDi neighbor discovery protocol and the antenna optimizations.

In Chapter 4 we compare the DANDi protocol with SAND, but we do not

compare it with Q-SAND. This is because Q-SAND is an optimization of SAND that only discovers the links with the highest RSSI, and also assumes line of sight between nodes. DANDi and SAND are both general, work for any environment, and aim to discover every S2S link. This being said, it would be interesting to design and implement a "Quick" version DANDi that only discovers the links with the highest RSSI, and then compare it with Q-SAND. This could be achieved by logging the RSSI of the received probes, and avoiding to reply to newer ones with lower RSSI, and thus reducing the contention and neighbor discovery time.

We have worked with the ContikiMAC RDC mechanism based on low power listening, which was the most widely used for a long time. Now the new generation of the Contiki OS uses a timed slotted mechanism called TSCH [14], where the communication between nodes is scheduled in time slots. Scheduled communication could simplify some of the problems that arise when using directional antennas with protocols based on low power listening. After establishing the neighbors' table using a neighbor discovery mechanism like the ones analyzed in Chapter 4, the nodes could leverage the scheduled nature of TSCH to achieve synchronization and point their antennas between the nodes assigned to a certain time slot. As we analyzed before in this chapter, one of the most important benefits of using directional antennas is to reduce contention and interference with neighbor nodes. In TSCH this problem is to some extent handled with time-slotting and channel hopping mechanisms, so using directional antennas may just increase the transmission range. However, the use of directional antennas could still be beneficial in very dense networks, or in scenarios where multiple TSCH networks are deployed, as they could help to reduce inter-network interference. In any case, studying this in more detail could be interesting.

We consider a six-element antenna for our evaluations as the SPIDA antenna is the reference directional antenna for WSN, also used in most of the related work. If an antenna with more or fewer elements is used, the analysis can be repeated by changing the antenna model in COOJA and building new antennas, but the general design of the protocols will still be valid.

Another assumption we make is that nodes do not move and that channel conditions are static. Under these circumstances, neighbor discovery and antenna selection heuristics need to be performed only once at the beginning. If channel conditions are dynamic or the nodes move, we would have to repeat neighbor discovery and antenna selection heuristics frequently, which is only possible with the distributed heuristics (BestSS_Dist and BestSSOpt_Dist) and dynamic neighbor discovery protocols like DANDi. Future work could include analyzing the performance of the network under changing channel conditions, or with nodes in movement.

In Chapter 5 we choose to analyze a convergecast application where every node sends packets periodically to a sink node. If we wanted to establish communications between any pair of nodes, some of the techniques we used to reduce the number of links in use would not be adequate, and we would have to think of other optimizations. Another aspect that would be interesting to analyze in future work is how would the protocols presented in this chapter behave in larger networks.

Chapter 6.   Conclusions and Future Work

Even though the design does not have any limitation in terms of network size or density, there are some hardware constraints that could affect the results like the amount of memory needed to store the neighbor tables, although this is a common problem for any neighbor discovery protocol.

Bibliography

[1] IEEE Standard for Low-Rate Wireless Networks. *IEEE Std 802.15.4-2015 (Revision of IEEE Std 802.15.4-2011)*, pages 1–709, April 2016.

[2] A. Bachir, M. Dohler, T. Watteyne, and K. K. Leung. MAC essentials for wireless sensor networks. *IEEE Communications Surveys & Tutorials*, 12(2):222–248, 2010.

[3] A. Boulis. Castalia 3.2, user's manual. *Australia: National ICT Australia Ltd*, 2011.

[4] M. Buettner, G. V. Yee, E. Anderson, and R. Han. X-MAC: a short preamble MAC protocol for duty-cycled wireless sensor networks. In *SenSys '06: Proceedings of the 4th international conference on Embedded networked sensor systems*, Boulder, Colorado, USA, 2006.

[5] L. Chen, Y. Li, and A. V. Vasilakos. On oblivious neighbor discovery in distributed wireless networks with directional antennas: Theoretical foundation and algorithm design. *IEEE/ACM Transactions on Networking*, 25(4):1982–1993, Aug 2017.

[6] Chipcon AS. *CC2400 datasheet - 2.4 GHz Low-Power RF Transceiver (Rev. 1.5)*, Mar. 2006. http://focus.ti.com/lit/ds/symlink/cc2400.pdf - Last visited January 2009.

[7] Chipcon AS. *CC2420 datasheet - 2.4 GHz IEEE 802.15.4 / ZigBee-Ready RF Transceiver (Rev. B)*, Mar. 2007.

[8] O. Chughtai, N. Badruddin, M. Rehan, and A. Khan. Congestion detection and alleviation in multihop wireless sensor networks. *Wireless Communications and Mobile Computing*, 2017, 2017.

[9] D.-I. Curiac. Wireless sensor network security enhancement using directional antennas: State of the art and research challenges. *Sensors*, 16(4):488, 2016.

[10] M. Demirbas, O. Soysal, and M. Hussain. A singlehop collaborative feedback primitive for wireless sensor networks. In *INFOCOM 2008. The 27th Conference on Computer Communications. IEEE*, pages 2047–2055. IEEE, 2008.

Bibliography

[11] A. Devices. ADG902 RF switch, Product manual. *Accessed: Jun*, 6:2018, 2005.

[12] A. Dunkels. The ContikiMAC Radio Duty Cycling Protocol. Technical Report T2011:13, Swedish Institute of Computer Science, Dec. 2011.

[13] A. Dunkels, B. Gronvall, and T. Voigt. Contiki - a lightweight and flexible operating system for tiny networked sensors. In *29th Annual IEEE International Conference on Local Computer Networks*, pages 455–462, Nov 2004.

[14] S. Duquennoy, A. Elsts, A. Nahas, and G. Oikonomou. TSCH and 6TiSCH for Contiki: Challenges, Design and Evaluation. In *Proceedings of the International Conference on Distributed Computing in Sensor Systems (IEEE DCOSS 2015)*, 2017.

[15] S. Duquennoy, O. Landsiedel, and T. Voigt. Let the Tree Bloom: Scalable Opportunistic Routing with ORPL. In *Proceedings of the International Conference on Embedded Networked Sensor Systems (ACM SenSys 2013)*, Rome, Italy, Nov. 2013.

[16] P. Dutta, S. Dawson-Haggerty, Y. Chen, C.-J. M. Liang, and A. Terzis. A-MAC: A Versatile and Efficient Receiver-initiated Link Layer for Low-power Wireless. *ACM Trans. Sen. Netw.*, 8(4):30:1–30:29, Sept. 2012.

[17] B. El Khamlichi, D. H. N. Nguyen, J. El Abbadi, N. W. Rowe, and S. Kumar. Collision-aware neighbor discovery with directional antennas. In *2018 International Conference on Computing, Networking and Communications (ICNC)*, pages 220–225, March 2018.

[18] E. Felemban, R. Murawski, E. Ekici, S. Park, K. Lee, J. Park, and Z. Hameed. SAND: Sectored-Antenna Neighbor Discovery Protocol for Wireless Networks. In *2010 7th Annual IEEE Communications Society Conference on Sensor, Mesh and Ad Hoc Communications and Networks (SECON)*, pages 1–9, June 2010.

[19] 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*, pages 1–14, 2012.

[20] F. Ferrari, M. Zimmerling, L. Thiele, and O. Saukh. Efficient network flooding and time synchronization with glossy. In *Proceedings of the 10th ACM/IEEE International Conference on Information Processing in Sensor Networks*, pages 73–84. IEEE, 2011.

[21] B. Geletu, L. Mottola, T. Voigt, and F. Österlind. Poster abstract: Modeling an electronically switchable directional antenna for low-power wireless networks. In *IPSN*, 2011.

[22] 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*, pages 1–14. ACM, 2009.

[23] A. Goldsmith. *Wireless Communications*. Cambridge University Press, New York, NY, USA, 2005.

[24] B. Großwindhager, M. Rath, J. Kulmer, M. S. Bakr, C. A. Boano, K. Witrisal, and K. Römer. SALMA: UWB-based single-anchor localization system using multipath assistance. In *Proceedings of the 16th ACM Conference on Embedded Networked Sensor Systems*, pages 132–144, 2018.

[25] T. Instruments. CC2420 datasheet. *Reference SWRS041B*, 2007.

[26] T. Instruments. MSP430F1611. *Revised Datasheet, May*, 2009.

[27] T. Istomin, M. Trobinger, A. L. Murphy, and G. P. Picco. Interference-resilient ultra-low power aperiodic data collection. In *2018 17th ACM/IEEE International Conference on Information Processing in Sensor Networks (IPSN)*, pages 84–95. IEEE, 2018.

[28] X. Ji, Y. He, J. Wang, W. Dong, X. Wu, and Y. Liu. Walking down the STAIRS: Efficient collision resolution for wireless sensor networks. In *IEEE INFOCOM 2014 - IEEE Conference on Computer Communications*, pages 961–969, Apr. 2014.

[29] Y. Jiang, H. Zhang, B. Zhao, and S. Rangarajan. Optimizing multicast delay with switched beamforming in wireless networks. In *2011 IEEE International Conference on Communications (ICC)*, pages 1–6. IEEE, 2011.

[30] A. Kalis, A. G. Kanatas, and C. B. Papadias. *Parasitic Antenna Arrays for Wireless MIMO Systems*. Springer, 2014.

[31] V. V. Khairnar, B. V. Kadam, C. Ramesha, and L. J. Gudino. A reconfigurable parasitic antenna with continuous beam scanning capability in H-plane. *AEU-International Journal of Electronics and Communications*, 88:78–86, 2018.

[32] J. G. Ko, N. Tsiftes, A. Dunkels, and A. Terzis. Pragmatic low-power interoperability: ContikiMAC vs TinyOS LPL. In *2012 9th Annual IEEE Communications Society Conference on Sensor, Mesh and Ad Hoc Communications and Networks (SECON)*, pages 94–96. IEEE, 2012.

[33] 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*, page 1. ACM, 2013.

[34] P. Levis, T. Clausen, J. Hui, O. Gnawali, and J. Ko. The trickle algorithm. *Internet Engineering Task Force, RFC6206*, 2011.

[35] J. Lu, D. Ireland, and R. Schlub. Dielectric embedded ESPAR (DE-ESPAR) antenna array for wireless communications. *IEEE transactions on antennas and propagation*, 53(8):2437–2443, 2005.

[36] Moteiv Corporation. *Tmote Sky Datasheet*, June 2006. Rev. 1.0.2.

[37] L. Mottola, T. Voigt, and G. Picco. Electronically-switched directional antennas for wireless sensor networks: A full-stack evaluation. In *Int. Conf. on Sensor and Ad-Hoc Communication and Networks (SECON)*, New Orleans, USA, June 2013.

[38] G. Mulligan, N. Kushalnagar, and G. Montenegro. IPv6 over IEEE 802.15.4 BOF (6lowplan). Web page. Visited 2005-02-21.

[39] R. Murawski, E. Felemban, E. Ekici, S. Park, S. Yoo, K. Lee, J. Park, and Z. Hameed Mir. Neighbor discovery in wireless networks with sectored antennas. *Ad Hoc Networks*, 10(1):1 – 18, 2012.

[40] M. Nilsson. Directional antennas for wireless sensor networks. In *Proc. 9th Scandinavian Workshop on Wireless Adhoc Networks (Adhoc'09)*, 2009.

[41] F. J. Oppermann, C. A. Boano, and K. Römer. A decade of wireless sensing applications: Survey and taxonomy. In *The Art of Wireless Sensor Networks*, pages 11–50. Springer, 2014.

[42] F. Osterlind, A. Dunkels, J. Eriksson, N. Finne, and T. Voigt. Cross-Level Sensor Network Simulation with COOJA. In *Proceedings. 2006 31st IEEE Conference on Local Computer Networks*, pages 641–648, Nov 2006.

[43] E. Öström, L. Mottola, and T. Voigt. Evaluation of an electronically switched directional antenna for real-world low-power wireless networks. In *International Workshop on Real-world Wireless Sensor Networks*, pages 113–125. Springer, 2010.

[44] U. Pešović and P. Planinšič. Error probability model for IEEE 802.15. 4 wireless communications in the presence of co-channel interference. *Physical Communication*, 25:43–53, 2017.

[45] B. Raman, K. Chebrolu, N. Madabhushi, D. Y. Gokhale, P. K. Valiveti, and D. Jain. Implications of link range and (in) stability on sensor network architecture. In *Proceedings of the 1st international workshop on Wireless network testbeds, experimental evaluation & characterization*, pages 65–72. ACM, 2006.

[46] R. Ramanathan, J. Redi, C. Santivanez, D. Wiggins, and S. Polit. Ad hoc networking with directional antennas: a complete system solution. *Selected Areas in Communications, IEEE Journal on*, 23(3):496–506, 2005.

[47] B. Rodríguez, J. Schandy, J. P. González, L. Steinfeld, and F. Silveira. Fabrication and characterization of a directional SPIDA antenna for wireless sensor networks. In *URUCON, 2017 IEEE*, pages 1–4. IEEE, 2017.

[48] D. Rojas and J. Barrett. Experimental analysis of a wireless sensor network in a multi-chamber metal environment. In *European Wireless 2016; 22th European Wireless Conference*, pages 1–6. VDE, 2016.

[49] R. A. Santosa, B.-S. Lee, C. K. Yeo, and T. M. Lim. Distributed neighbor discovery in ad hoc networks using directional antennas. In *Computer and Information Technology, 2006. CIT'06. The Sixth IEEE International Conference on*, pages 97–97. IEEE, 2006.

[50] C. Sarkar, R. V. Prasad, R. T. Rajan, and K. Langendoen. Sleeping beauty: Efficient communication for node scheduling. In *2016 IEEE 13th International Conference on Mobile Ad Hoc and Sensor Systems (MASS)*, pages 56–64. IEEE, 2016.

[51] J. Schandy, L. Steinfeld, and F. Silveira. Average power consumption breakdown of Wireless Sensor Network nodes using IPv6 over LLNs. In *Distributed Computing in Sensor Systems (DCOSS), 2015 International Conference on*, pages 242–247. IEEE, 2015.

[52] R. Schlub, J. Lu, and T. Ohira. Seven-element ground skirt monopole ESPAR antenna design from a genetic algorithm and the finite element method. *IEEE Transactions on Antennas and Propagation*, 51(11):3033–3039, 2003.

[53] L. Selavo and O. Chipara. Directional Antenna Platform for Low Power Wireless Networks. In *EWSN*, pages 258–259, 2017.

[54] G. D. S. Sidibe, A. Surier, R. Bidaud, G. Delisle, N. Hakem, M. Servajean, B. Rmili, G. Chalhoub, and M. Misson. Use of a Switched Beam Antenna in a Star Wireless Sensor Network for Data Collection: Neighbor Discovery Problem. In *2019 IEEE International Conference on Wireless for Space and Extreme Environments (WiSEE)*, pages 21–26, Oct 2019.

[55] G. Tarter, L. Mottola, and G. P. Picco. Directional antennas for convergecast in wireless sensor networks: Are they a good idea? In *Mobile Ad Hoc and Sensor Systems (MASS), 2016 IEEE 13th International Conference on*, pages 172–182. IEEE, 2016.

[56] S. Tennina, O. Gaddour, A. Koubâa, F. Royo, M. Alves, and M. Abid. Z-Monitor: A protocol analyzer for IEEE 802.15. 4-based low-power wireless networks. *Computer Networks*, 95:77–96, 2016.

[57] N. Tiwari and T. R. Rao. A switched beam antenna array with butler matrix network using substrate integrated waveguide technology for 60 GHz wireless communications. *AEU-International Journal of Electronics and Communications*, 70(6):850–856, 2016.

[58] E. Toscano and L. L. Bello. Comparative assessments of IEEE 802.15. 4/ZigBee and 6LoWPAN for low-power industrial WSNs in realistic scenarios. In *2012 9th IEEE International Workshop on Factory Communication Systems*, pages 115–124. IEEE, 2012.

Bibliography

[59] A. Varshney, L. Mottola, M. Carlsson, and T. Voigt. Directional transmissions and receptions for high-throughput bulk forwarding in wireless sensor networks. In *Proceedings of the 13th ACM Conference on Embedded Networked Sensor Systems*, pages 351–364. ACM, 2015.

[60] A. Varshney, T. Voigt, and L. Mottola. *Using Directional Transmissions and Receptions to Reduce Contention in Wireless Sensor Networks*, pages 205–213. Springer International Publishing, Cham, 2014.

[61] S. Vasudevan, J. Kurose, and D. Towsley. On neighbor discovery in wireless networks with directional antennas. In *INFOCOM 2005. 24th Annual Joint Conference of the IEEE Computer and Communications Societies. Proceedings IEEE*, volume 4, pages 2502–2512. IEEE, 2005.

[62] J. Wang, W. Peng, and S. Liu. Neighbor discovery algorithm in wireless local area networks using multi-beam directional antennas. *Journal of Physics: Conference Series*, 910:012067, oct 2017.

[63] Y. Wang, S. Mao, and T. S. Rappaport. On directional neighbor discovery in mmwave networks. In *2017 IEEE 37th International Conference on Distributed Computing Systems (ICDCS)*, pages 1704–1713, June 2017.

[64] B. Wei, A. Varshney, N. Patwari, W. Hu, T. Voigt, and C. T. Chou. dRTI: Directional Radio Tomographic Imaging. In *Proceedings of the 14th International Conference on Information Processing in Sensor Networks*, IPSN '15, pages 166–177, New York, NY, USA, 2015. ACM.

[65] K. Whitehouse, A. Woo, F. Jiang, J. Polastre, and D. Culler. Exploiting the capture effect for collision detection and recovery. In *Embedded Networked Sensors, 2005. EmNetS-II. The Second IEEE Workshop on*, pages 45–52. IEEE, 2005.

[66] T. Winter, P. Thubert, A. Brandt, J. Hui, R. Kelsey, P. Levis, K. Pister, R. Struik, J. Vasseur, and R. Alexander. RPL: IPv6 Routing Protocol for Low-Power and Lossy Networks. RFC 6550, RFC Editor, March 2012.

[67] W. Xiong, B. Liu, and L. Gui. Neighbor discovery with directional antennas in mobile ad-hoc networks. In *2011 IEEE Global Telecommunications Conference - GLOBECOM 2011*, pages 1–5, Dec 2011.

[68] S. Zhang and A. Datta. A directional-antenna based MAC protocol for wireless sensor networks. In *International Conference on Computational Science and Its Applications*, pages 686–695. Springer, 2005.

List of Tables

List of Tables

List of Figures

List of Figures

This is the last page.
Compiled on Thursday 27th August, 2020.
`http://iie.fing.edu.uy/`