



Part I.

Introduction

The year 2008 marked an important turning point in human history. For the first time, more humans were living in urban than in rural areas.¹ Today, cities are growing and transforming at an ever-increasing rate. The challenges on social, political, technical, and administrative levels are manifold. They span from waste management, logistics, and energy all the way to education, poverty, and the environment.

Citizens, the economy, administration, and other city stakeholders are in desperate need for tools that allow them to handle their old responsibilities but, more importantly, adapt to new ones. The term *smart city* [65] describes the technological perspective of a city meeting this demand. An important aspect of a smart city is the ability to collect data about itself, preferably in real time. This data enables city stakeholders to make better and more transparent decisions, reach their goals faster, and be more eco-friendly in the process. This leads to the inherent need for a city data layer. Wireless sensor networks, combined with other technologies, can provide this data layer. Their scalability both from a cost and network management perspective make them the perfect fit for large-scale citywide deployments.

Wireless sensor networks have been an active research field, since the smart dust idea was published by Kahn et al. in 1999 [77]. The possibility of small wireless transceivers lead to the idea of very small (< 1 mm), very cheap sensor nodes that connect wirelessly and scale into millions. While this vision hasn't become a reality yet, sensor networks have been used in many applications. Examples include the tracking of animals [139, 160, 10, 135, 5], people, and vehicles; the monitoring of data centers [93] or vineyards [17, 16]; and the development of wearable wireless sensor networks [99].

Wireless sensor networks have also been used heavily in city deployments monitoring the environment or buildings [20, 11, 22, 57, 67, 106, 59]. But even though there is vast potential for wireless sensor networks applications, some challenges remain unsolved, preventing a full breakthrough.

The main challenges that hinder large urban sensor networks are the following:

Deployment: Sensor networks are supposed to scale into millions of nodes. Today, sensor networks require a large amount of manual labor for both monitoring and deploying the sensor nodes. Scaling networks into hundred or even thousands of nodes is only possible with a large investment both in labor and money. And even though hardware cost and size are decreasing, they have not yet reached a level suitable for commercial large-scale deployment.

Standards and generic solutions: Sensor networks are designed to be application-specific; thus, most solutions presented today will only work under specific

¹ <http://www.unfpa.org/pds/urbanization.htm>

circumstances. Different layers of the network stack are often merged to optimize parameters, such as energy consumption, message loss, etc. In order to enable a true Internet of Things (IoT) we need general standards that are interoperable with other networks. Some work has been done over the past years to standardize IPv6 (RFC 4944) [105] for wireless sensor networks, but more work needs to be done in the future.

Energy: Energy remains the most vital challenge. Since sensor nodes are still more expensive and harder to deploy than envisioned, it is important that a given node remains active as long as possible. In the past, it was shown that two approaches are possible to enable energy-efficiency, either by deploying more energy-efficient algorithms or by enabling the nodes to harvest environmental energy (e.g., by using solar panels).

Deployment is an application-independent challenge, even though it might be harder for some applications (e.g., remote monitoring). As cheaper hardware will be available more commercial applications are built. This will increase the need for maintenance, and the industry will focus on tools for maintenance. In turn, deployment costs will go down. Industry or industry-driven research is, therefore, better equipped to deal with deployment challenges. The same is true for *standards and generic solutions*. As the research converges against well-known algorithms, standardization will follow. Indeed, with the implementation of IPv6 [105, 142] on wireless sensor nodes and the standardization of RPL as routing protocol by the Internet Engineering Task Force (IETF) [156], we have reached a point where parts of the network stack are already standardized. Other parts will probably follow soon.

The challenge addressed in this thesis is *energy*. We want to design wireless sensor networks that are energy-aware. The contribution of this thesis is threefold:

kTC: Wireless communication is responsible for most of the energy consumed in sensor nodes. Energy can, for example, be saved by better medium access control (MAC) mechanisms and better routing or data collection mechanisms. Both has been done extensively in the past. Topology control (i.e., the reduction of transmission range without disconnecting the network) has received mostly theoretical attention. Hence, we introduce kTC, a practical topology control approach. kTC creates planar, symmetric topologies, which are degree bounded and contain the minimum spanning tree of the underlying connectivity graph. kTC was compared against three state-of-the-art algorithms in simulation. Using idealistic connection models, kTC was able to perform better or as good as the best state-of-the-art algorithm. kTC is

vastly superior to all state-of-the-art algorithms, if the model is modified to be more realistic. These results indicate practicality and robustness in real-world networks, which was further evaluated by implementing kTC on real hardware. Using TUD μ Net, an open testbed, and different load scenarios, we are able to show up to 16% in total energy saved. This is in addition to already energy-efficient MAC and routing protocols.

Solar-aware distributed flow (SDF): Batteries are the most limiting factor in wireless sensor networks and algorithms, such as kTC help to increase lifetime. To be really energy-independent methods such as energy harvesting must be deployed. If energy harvesting is used, another challenge arises. Nodes might harvest more energy than they consume. This energy can either be wasted or it can be used to generate a higher quality of service. The goal is to reach the so-called energy neutral consumption rate, where no energy is wasted. In urban applications, the quality of service is the data coverage of a given area. The coverage is increased if the sampling rate of each node is increased. This implies that in order to reach the energy neutral consumption rate, the sampling rate of each node should be maximized. SDF will control the sampling rate and message flow, such that the sampling rate of each node is maximized. To the best of our knowledge, SDF is the first fully distributed approach. Performing extensive simulations using different routing protocols, we can report an average energy utilization of almost 100%. At the same time, the average sampling rate of SDF is always within 90% of the theoretical optimum. The testbed implementation confirmed both the feasibility, as well as energy utilization.

da_sense: In order to understand the requirements of wireless sensor networks in urban environments (urban sensor networks), we have developed da_sense. The da_sense platform integrates three sensor sources: (i) inductive loops counting cars, (ii) a wireless sensor network deployed on trams, and (iii) a participatory sensor application to measure noise pollution.² In da_sense we tackle challenges, such as data delivery in mobile networks and incentives for participatory sensing. We also concentrate on integrating different sensor sources and gather real-world data. The platform is available at <http://www.da-sense.de>.

Figure 1 offers a classification of the contributions and how they are related to other parts of the communication stack. The figure also illustrates how da_sense

² <http://www.youtube.com/watch?v=JimyRNYLZKo>

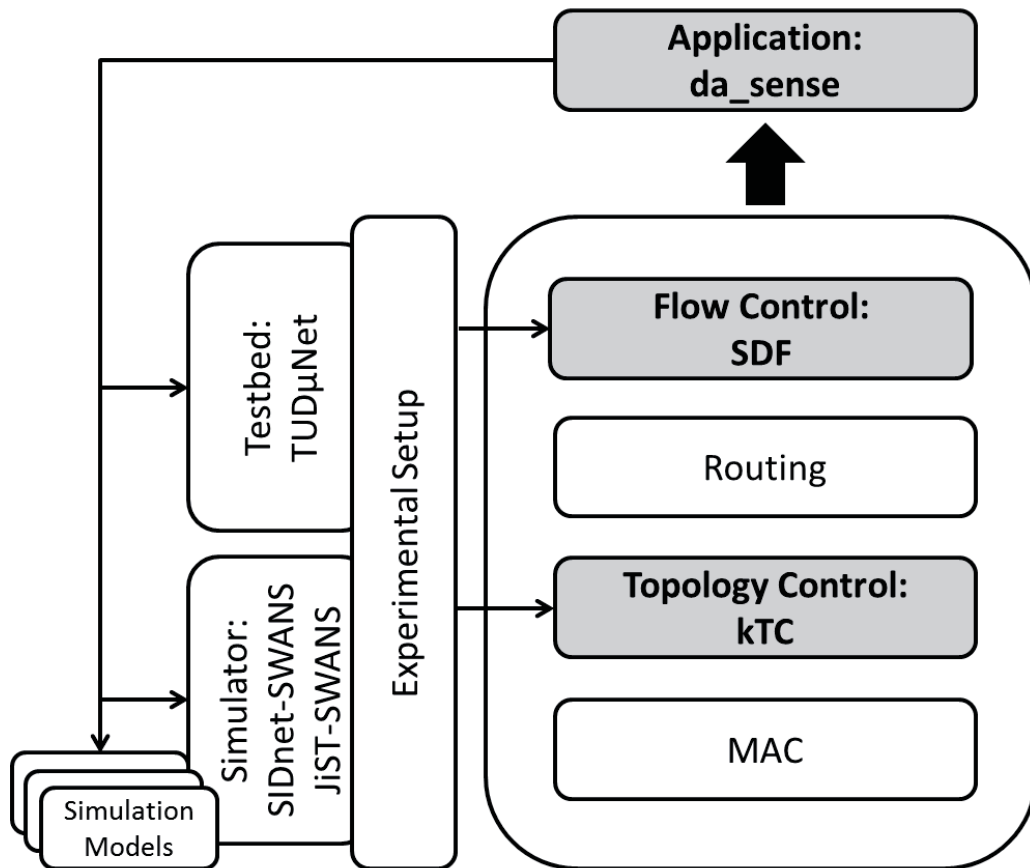


Figure 1.: Classification of the contributions

is related to the experimental setup, consisting of testbed and simulation, which is used to evaluate both approaches.

The remainder of the thesis is organized as follows. Part II will introduce the related work. Afterward, Part III introduces kTC and the respective evaluation results. The SDF algorithm, as well as the evaluation and testbed implementation are described in Part IV. Part V introduces the da_sense platform and all data sources available today. Part VI concludes the thesis and discusses future work.

1 Publications

This chapter is meant to give a short summary of my publications related to this thesis.

kTC was first published at the IEEE International Conference on Computer Communication and Networks (ICCCN) in 2012 [129]. The work that led to the development of kTC was published at IEEE Globecom [86] and did investigate the idea of using graph patterns to balance peer-to-peer topologies. To understand graphs (e.g., wireless topologies), I was also involved in the development of GTNA, a framework for graph-theoretic network analysis [118].

Solar-aware distributed flow control was published at the 2011 IEEE Local Computer Networks conference [125].

da_sense has a number of publications related to different aspects of the platform. In [124] we are investigating data quality in mobile sensing and introduce a frequency-based calibration scheme. Noisemap has also incorporated different non-monetary incentive schemes, which are used to gather more data. The incentive schemes were evaluated in [126]. The wireless sensor network deployed on trams and the future integration in the TUD μ Net testbed was described in [54].

In addition, I have been heavily involved in investigating the use of wireless networks for resilient first responder communication [13, 82, 108, 109, 110] and have some work on artificial intelligence [127].