On the Benefits and Challenges of Crowd-Sourced Network Performance Measurements for IoT Scenarios

Lars Møller Mikkelsen¹ · Tatiana Kozlova Madsen¹ · Hans-Peter Schwefel^{1,2}

Published online: 3 October 2019
© The Author(s) 2019

Abstract

Systems within IoT domains such as ITS, Smart City, Smart Grid and other, often rely on real-time information and communication. These types of systems often include geographically distributed nodes which are connected via cellular or other wireless networks. This means great variability and uncertainty in network connection performance, effectively increasing the expected minimum system response time. Having information about network connection performance means that it is possible to predict the performance of the system in terms of sensor access delay or application response time. We obtain the performance information, in terms of signal strength and transport layer round trip time, using crowd sourcing and consumer devices which causes the measurements to be heterogeneously distributed. From these measurements we want to create a network performance map but in areas with sparse measurements the reliability of the map values will be low. To solve this problem we include neighboring measurements and evaluate the impact of doing so. We show that generally there is a benefit from including neighboring measurements, and that transport layer round trip times are less sensitive to bias when increasing the size of the extended area to include measurements from.

Keywords Crowd-sourcing · Cellular network performance · Sparse measurements

1 Introduction

Internet of Things (IoT) is rapidly being developed and starting to being deployed. More and more IoT devices, systems and services are emerging [8]. IoT systems rely on information, and especially information about the world they operate in, such as temperature,

✉ Tatiana Kozlova Madsen
tatiana@es.aau.dk

Lars Møller Mikkelsen
lmm@es.aau.dk

Hans-Peter Schwefel
hps@es.aau.dk

¹ Aalborg University (AAU), Aalborg, Denmark

² GridData, Anger, Germany

air quality, number of users, device state and other. This information can both be used as historical or live information, i.e. from a database with previous values or directly from the sensor as the source of the information. In both cases it is often not enough just to have the information, but also awareness of the quality of the information is needed. One measure of quality could be freshness of the information, or knowledge of what freshness to expect from future updates of the information, i.e. prediction based on historical information.

To be able to predict the freshness of information it is necessary to obtain measurements of end-to-end network delay, which can be obtained cost effectively using crowd sourcing. From the measurements we will create a map of network performance, which represents the geo-dependent cellular network performance by the mean value of different performance metrics; in this paper, we use transport-layer round-trip times and signal strength. In [3] it is shown how a network performance map can be created and used for optimizing TCP based data transfer. To create the map we will divide the geographical area into cells (not to be confused with radio cells) in which we will aggregate the measurements in the mean value. There are cells with only few measurements where the mean estimator is showing a high variance, so it is not 'trustworthy' for further prediction use. In order to reduce the variability, this paper investigates an approach to include neighboring measurements. This increases the number of measurements and therefore reduces variability, but on the other hand may introduce bias, as these are sampled from different locations. This trade-off is analyzed in this paper.

To understand the problem we first have to look at what measurements we include in the map, how we obtain them, and what they will be used for.

In an end-to-end connection in a IoT system between a sensor and an application on a user device, there are at least two wireless links; the sensor connection, and the connection of the user device running the application. In the connection between sensor and application there will also be several wired links, but we assume that the wireless links will by far have the greatest impact on the end-to-end connection performance. The wireless connection of the sensor is typically achieved by using low-power technologies [9], as sensors often are fixed in location and only need to transmit low amounts of data. The wireless connection of the user device will typically be a cellular connection to achieve high mobility and ubiquitous high speed network coverage, but also to support a wide range of usages. We choose to focus on obtaining information about how the cellular wireless link influences the end-to-end connection performance. Subsequent we will denote this as the connection performance.

To get information about the actual connection performance we will apply an active measurement methodology, meaning we will generate measurement traffic and not just utilize already present traffic. Generally there are two approaches that connection performance can be acquired in; dedicated measurements or crowd sourcing. When applying dedicated measurements a very accurate picture is obtained of exactly what is measured, but the measurements can be costly in terms of time and measurement equipment [7]. From the crowd sourcing measurements more unknowns are included in the results which must be handled in post processing [1], while the costs of performing the measurements are low [2]. In this work we use crowd sourcing to measure the connection performance from the end user devices, realized using the NetMap system [6].

The measurements are influenced from factors such as signal disturbances and interferences, network load, device load, different device and antenna characteristics, different networks and network technologies, etc., all of which cause measurement values to vary. This has been explored and documented in works such as [12] that shows that movement highly influences the measured connection performance, while [5] studies the impact of

cellular connections content access in general. Furthermore, [10] show that signal strength and higher layer metrics not necessarily are highly correlated, underlining the need for measurement of both type of metrics. The varying measurement distribution is due to the layout of roads and buildings, and how users move and where they spend more or less time. This means that one area can have many measurements, and the immediately neighboring area can have few or no measurements. This is investigate in [11] where a bandwidth map is created from measurements performed only while driving on roads. Furthermore, [4] evaluate the influence of the hidden state of the network, i.e. other factors than location, on network performance is evaluated.

In this work we focus on how to handle the varying measurement density, by evaluating the mean value and the impact on the mean when enhancing sparse measurement cells with measurements from neighboring cells.

The rest of this paper is organized as follows; in Sect. 2 the measurement method and the measurements are introduced. Section 3 describes the approach we apply for evaluation of the measurements and presents the evaluation results. Section 4 concludes on the results and gives an outlook to future work.

2 Crowd-Sourced Network Performance Measurements

We will base our evaluation on two measurement sets, one obtained in an urban area, and one in a rural area. In this section we first describe the measurements and how they are obtained, followed by a description of the processing approach and evaluation of the results.

As mentioned in the introduction we have decided to focus on the performance of the cellular connection, because typically this is the link with the highest influence on the end to end connection performance as seen from the user perspective. The measurements will be performed by consumer smartphones using the crowd sourcing measurement system NetMap.

2.1 Measurement Collection Software and Metrics

The measurements were collected using the NetMap system [6] on Android devices, and are all performed using a 3G connection. NetMap is a system developed for performing crowd based network performance measurements. Users install an app that periodically performs measurements of various QoS metrics on the cellular data connection, and automatically submits the results to the back end, along with a wide range of additional context information about the device at the time of measurement.

For our purpose the QoS metrics we measure are packet round-trip time (RTT) and signal strength. The measurement system distinguishes between the RTT of TCP packets and of UDP packets. Both times are measured while actively exchanging data between the mobile device and a measurement server. The signal strength values are the result of a passive measurement that does not include sending any data, but are measured while the connection is active measuring RTT. Consequently each individual sample contains either a TCP RTT or a UDP RTT value along with a signal strength value, a timestamp, and longitude and latitude of the location.

RTT is recorded as the time it takes to send a data packet from the client to a server and to send a data packet from the server to the client, where the server replies as fast a

possible. This is done with 20 request/reply sequences both for TCP and UDP. For TCP, the connection is initialized before the measurement starts. For both TCP and UDP the client waits for the reply to the previous request to arrive before sending the next request, and the size of the data is 20 bytes.

As soon as the sampling process has been started in the NetMap app, measurements are done periodically until the process is manually stopped. The schedule of measurements is configured such that in each round first TCP RTT is measured together with signal strength, and then, after a delay of 2 s, measurements of UDP RTT and signal strength are performed. The process then remains idle for a uniform random time of 0–10 s before the next round starts.

2.2 Measurement Setting

We collected measurements in two different settings: rural and urban. In the urban setting we collected measurements in an area confined to a few streets with a mix of residential buildings and shops, while walking in a normal pace. The measurements were performed using 2 similar devices (LG G4c (LG-H525N) and LG G3 (LG-D855)), both connected to the same network, using 3G UMTS 2100 MHz as connection technology.

In the rural setting we collected measurements on a 13km stretch of road going through a rural area. The measurements were collected both while driving and walking. The measurements were performed using 2 identical devices (Motorola Nexus 6), both connected to the same network, using 3G UMTS 900 and 2100 MHz as connection technology. In both scenarios the measurements were collected during several days, but only between 8 in the morning and 8 in the evening.

2.3 Measurements

Table 1 lists the amount of measurements collected in the two settings.

In Figs. 1 and 2 it can be seen where the measurements were collected in the two scenarios.

In Figs. 3 and 4 the measurement value distributions of measurements from the two settings can be seen.

3 Evaluation of Measurements

In this section we describe how we create a cellular network performance map and how we employ a simple interpolation approach to handle sparse measurement cells. Furthermore, we describe our approach to evaluate the impact on the map values when interpolating measurements to sparse measurement cells.

Table 1 Number of measurements per type and per setting

Measurement type	Rural	Urban
TCP RTT	1586	475
UDP RTT	1580	477
Signal Strength	3166	952

Fig. 1 Locations of urban measurements map on Google Maps

3.1 Measurement Processing Approach

Here we describe how we generate a network performance map, and from that highlight the problem we investigate. We will refer to the area where we have measurements within as the full geographical area. The full geographical area is divided into square non-overlapping cells (not to be confused with radio cells), and from the measurements within each cell we calculate the sample mean and a confidence interval of the sample mean. The chosen size of the cells, and thereby the resolution of the map, will depend on requirements of the use case of the map, which we will not specify further, but we will evaluate cell sizes between 20m and 65m. Depending on the cell size and measurement density, we will have cells with statistically sufficient measurements and cells with less. We select 30 measurements as a minimum to ensure a good statistical basis for the sample mean. In cells with fewer than 30 measurements we interpolate neighboring measurements by including them in the sample mean calculation. Adding more measurements will reduce the variance of the mean estimator, but on the other hand may add bias as the additional measurements may be subject to different environment influences than the measurements in the initial cell. We will analyze this impact in the following.

Fig. 2 Locations of rural measurements on Google Maps

Fig. 3 Measurement distributions of full urban dataset

Fig. 4 Measurement distributions of full rural dataset

3.2 Impact Evaluation Approach

Here we describe how we evaluate the impact of including neighboring measurements when having a sparse measurement cell. The goal is to evaluate the impact as a function of the distance to the included measurements. We assume a given cell size, which however may vary depending on the performance map and the application using it. Practically we will evaluate the impact at different cell sizes.

We start with selecting a cell as the initial cell containing a minimum of 30 measurements. We denote the measurements within the initial cell as m_S . We calculate the sample mean of measurements in this cell and denote this as the ground truth (GT). Furthermore, we calculate the 95% confidence interval of the GT (CI_{GT}) based on the measurements. CI_{GT} will be the basis of the further evaluation. In our evaluation we will for simplicity define cells as circles, where the radius of the circle is D . We will evaluate cells with diameter between 20m and 65m.

To simulate a cell with sparse measurements we sample m_S to get $n=20$ subsets, S_n (see Fig. 5). Each of these subsets will contain 5 measurements randomly selected from m_S . We calculate the mean of each of these subsets, which we will call μ_n .

We now start to include neighbor measurements in the subsets from outside the initial cell. We do this by extending the initial cell by defining a radius R , where $R \geq D$ and include the measurements placed outside the initial cell with radius D , and inside the extended cell with radius R . We call this set m_R . We include measurements by combining the full m_R with each of S_n subsets, which gives us new subsets $\hat{S}_n$. For each of the new subsets we calculate the mean $\hat{\mu}_n$.

Now we evaluate for each $\hat{\mu}_n$ if it is similar to GT, by checking if it is within CI_{GT} . If inside we give it the indicator value of 1 and 0 if outside. By averaging the indicator values over all subsets we get the average similarity as a value between 0 and 1. We repeat this for several values of R giving the similarity between the enhanced subset means $\hat{\mu}_n$ and GT.

Fig. 5 Representation of sampling approach

Following are listed the relevant parameters for the evaluation:

- D : Initial cell radius. In the range of 20m to 65m
- R : Extended cell radius to include further measurements within. In the range of D to 180m (300m for rural)
- Measurement types evaluated: TCP RTT, UDP RTT, and Signal Strength
- n : $n = 20$ subset samples from the initial cell
- Subset sample size: 5 randomly selected measurements as a subset of m_s

Example of initial cell with sparse measurements: In Fig. 6 we see an example of the processing of UDP RTT measurements from urban setting in an initial cell for $D=20$ m, where R is in the range of 20m to 180m. In the left plot we see the individual measurement values from the initial cell (m_s). In the right plot we see the evolution of the subset means ($\hat{\mu}_n$) from $R=20$ m up to $R=180$ m. At $R=20$ m no measurements outside the initial cell are included in the subsets while as R increases more and more measurements are included in the subsets from the extended cell. As R increases and we include more and more measurements (m_R) in the subsets, and in effect the means become less and less spread out. In this example we have relatively high spread in the measurements in the initial cell, which leads to spread in subset means for low R . We can see that for low R some of the means are inside CI_{GT} , and some are outside. In this initial cell from $R=75$ m onwards for the investigated range until $R=180$ m the subset means happen to all stay within CI_{GT} , but for another initial cell it could happen that they are outside CI_{GT} .

3.3 Impact Evaluation

Now we will analyze the impact of including neighboring measurements in sparsely populated measurement cells. We do this based on the output of our evaluation approach, described in previous section. Figures 7, 8, 9, 10, 11 and 12 show the evaluation output. We have evaluated results for all integer values of D between $D = 20$ m and $D = 65$ m, but to make the plots easier to read we only show results from a subset of D values. The conclusions do however still hold as the graphs evolve gradually from low to high D values. Note the graphs shows the similarity indicator averaged over several initial cell.

Fig. 6 Example of an initial cell in urban setting with $D=20\text{m}$, and the processing of UDP RTT measurements

3.3.1 Rural and Urban Signal Strength Evaluation

In Figures 7 and 8 we see the similarity between GT mean and the means of sparse measurements sets, when including measurements from m_R for Signal Strength measured in rural and urban settings. For both measurements in the rural and urban setting we see the similarity graphs increase when we include neighboring measurements. For rural setting the graphs rise to a maximum value between 75 and 85%, and around $R = 165\text{ m}$ they start decreasing again. For urban setting the graphs rise quickly to between 50 and 80%, and starts to decrease immediately after the initial increase. This drop continues to around $R = 140\text{ m}$ after which the graphs even out. Besides the decrease behavior after the initial increase another difference between measurements from rural and urban setting is the max similarity values that the graphs rise to initially. For urban setting the max level seem to be dependent on the D value and starting similarity.

3.3.2 Signal Strength Measurements Impact Considerations

From the Signal Strength measurements there is a clear indication of the initial benefit when including neighboring measurements in sparse measurement sets. But as we increase the size of the cell where we include measurements from the similarity decreases, with the

Fig. 7 Subset means similarity to GT of Signal Strength measurements in rural setting

Fig. 8 Subset means similarity to GT of Signal Strength measurements in urban setting

Fig. 9 Subset means similarity to GT of TCP RTT measurements in urban setting

Fig. 10 Subset means similarity to GT of UDP RTT measurements in urban setting

Fig. 11 Subset means similarity to GT of TCP RTT measurements in rural setting

Fig. 12 Subset means similarity to GT of UDP RTT measurements in rural setting

decrease being faster for urban setting than for rural. This indicates that the wireless signals change much faster in urban setting than in rural setting. This makes sense as in urban setting there are more obstacles to signal propagation, i.e. turning around a corner of a building will give you a significantly different signal path. So the maximum distance from which we should include neighboring measurements depends on the area type.

3.3.3 Urban TCP and UDP RTT Evaluation

In Figures 9 and 10 we see the similarity between GT mean and the means of sparse measurements sets when including neighboring measurements, for both TCP and UDP RTT measured in the urban setting. We see that for both TCP and UDP the similarity increases as soon as we start including neighboring measurements from m_R . For TCP RTT the maximum similarity values are between 75 and 85% while for UDP RTT the maximum values are between 60 and 75%. Furthermore, the similarity graph for the smallest value of D for UDP RTT experience a drop after the initial increase before rising to the maximum value.

3.3.4 Rural TCP and UDP RTT Evaluation

In Figures 11 and 12 we see the similarity between GT mean and the means of sparse measurements sets when including neighboring measurements, for both TCP and UDP RTT measured in the rural setting. Again here for both TCP and UDP RTT the similarity graphs increase as soon as measurements are included from m_R . For TCP RTT the maximum similarity for small values of D is at 100%, while for bigger values of D the maximum similarity values are around 80%. For UDP RTT the maximum similarity for small values of D is also at 100%, while for bigger values of D the maximum similarity is between 90 and 95%. Again here for $D = 20$ m, as it was the case for UDP RTT in urban setting, the similarity shows an initial drop, before rising to 100%.

3.3.5 TCP and UDP RTT Measurements Impact Considerations

For UDP and TCP RTT there is also a clear benefit of including neighboring measurements in sparse measurement sets. But what is different here from the Signal Strength measurements is that the benefit does not seem to disappear when including measurements from further and further away, or at least not within the distance in the available data. This means that if we obtain statistically sufficient measurements by including measurements within the first 20–40m outside the initial cell, then there is no additional benefit in similarity by looking further away, or at least not within 180 m and 300 m for urban and rural setting respectively. This is because for large R the measurements will always introduce bias as it will pull the mean in the limit to the mean value over the whole space.

In both rural and urban setting for TCP RTT the maximum similarity is around 80%, with an exception of small initial cell size for rural setting, which evens out at 100%. For UDP RTT in rural setting the maximum similarity value is higher than for urban setting. These observations for TCP and UDP RTT can be explained by looking at the distribution plots of the urban and rural setting measurements in Figs. 3 and 4. For rural setting the TCP and UDP RTT both have more narrow distributions than for urban setting measurements, why the maximum similarities are higher for rural than for urban setting. Furthermore, the TCP RTT distributions for urban and rural settings are more similar than the

UDP RTT distributions, why the similarity graphs look more alike for TCP RTT than for UDP RTT.

3.4 Recommended Size of R

Based on the evaluation in the previous paragraphs we can now make some recommendations of how far away from the sparsely measurement populated cell, or initial cell, to include measurements from.

For Signal Strength measurements in urban setting we can recommend to only go 20–40 m outside the cell as after that the benefit is reduced. In rural setting the distance is greater going up to around 100m outside the cell, as further away we see a small decrease in the benefit.

For TCP and UDP RTT measurements the recommendation is not as strict because we do not see a decrease in similarity values within the range of distance values that we investigated in the experiments. We would advice to not increase the cell further when statistically sufficient measurements have been obtained due to the introduction of bias.

4 Conclusion

In the previous sections we evaluated the impact on the mean estimate from cells with sparse measurements when including neighboring measurements, by comparing to GT mean of the cell. We did this by evaluating the similarity between means of subsets with GT mean. Increase in the similarity indicates that subset means are improved on average, i.e. more of the subset means are inside CI_{GT} . Generally when including neighboring measurements the similarity increases. This seem true for both Signal Strength and TCP and UDP RTT, and both in rural and urban settings. But limits to the distance to the included measurements vary depending on measurement metric and setting.

From this we can conclude that we can enrich the sparse measurement sets without compromising the accuracy of the mean estimate. For RTT, which is a transport layer performance metric, there seems not to be any significant impact of including neighboring measurements up to 180 m or 300 m for urban and rural setting respectively. For Signal Strength however, there seems to be a limit to how far away we can include measurements from, where the measurement setting is the defining factor.

Outlook In this paper we focused on the distance outside the initial cell, while not considering the size of the initial cell. This is however also an interesting topic to explore, as different use cases will have different requirements to map resolution. So we would like to investigate if it makes sense to make small cells in the map, or if there is a minimum limit.

In this paper we looked at the cell by cell impact of including neighboring measurements, but looking further ahead we would like to investigate the impact on a network performance map. This could be in terms of impact on detail level and coverage level.

Furthermore, we would also like to look at the actual use case of the network performance map, i.e. using it to attach quality metrics to transfered information in IoT.

Acknowledgements The work is supported by BIG IoT (Bridging the Interoperability Gap of the Internet of Things) project funded by the European Union Horizon 2020 research and innovation programme under Grant Agreement No 688038.

Open Access This article is distributed under the terms of the Creative Commons Attribution 4.0 International License (<http://creativecommons.org/licenses/by/4.0/>), which permits unrestricted use, distribution, and reproduction in any medium, provided you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons license, and indicate if changes were made.

References

1. Faggiani, A., Gregori, E., Lenzini, L., Luconi, V., & Vecchio, A. (2014). Smartphone-based crowd-sourcing for network monitoring: Opportunities, challenges, and a case study. *IEEE Communications Magazine*, 52(1), 106–113. <https://doi.org/10.1109/MCOM.2014.6710071>.
2. Howe, J. (2006). The rise of crowdsourcing. *Wired Magazine*, 14, 1–4.
3. Højgaard-Hansen, K., Madsen, T., & Schwefel, H. P. (2012). *Reducing communication overhead by scheduling TCP transfers on mobile devices using wireless network performance maps* (pp. 1–8). Piscataway: IEEE Press.
4. Højgaard-Hansen, K., Madsen, T., & Schwefel, H. P. (2015). *Addressing the influence of hidden state on wireless network optimizations using performance maps* (pp. 177–189). Berlin: Springer. https://doi.org/10.1007/978-3-319-23126-6_16.
5. Kaup, F., Michelinakis, F., Bui, N., Widmer, J., Wac, K., & Hausheer, D. (2016). Assessing the implications of cellular network performance on mobile content access. *IEEE Transactions on Network and Service Management*, 13(2), 168–180. <https://doi.org/10.1109/TNSM.2016.2544402>.
6. Mikkelsen, L., Thomsen, S., Pedersen, M., & Madsen, T. (2014). Netmap—creating a map of application layer qos metrics of mobile networks using crowd sourcing. In S. Balandin, S. Andreev, & Y. Koucheryavy (Eds.), *Internet of things, smart spaces, and next generation networks and systems* (Vol. 8638, pp. 544–555). Lecture notes in computer science Berlin: Springer. https://doi.org/10.1007/978-3-319-10353-2_50.
7. Misra, A., Gooze, A., Watkins, K., Asad, M., & Le Dantec, C. (2014). Crowdsourcing and its application to transportation data collection and management. *Transportation Research Record: Journal of the Transportation Research Board*, 2414, 1–8.
8. Porkodi, R., & Bhuvaneswari, V. (2014). The internet of things (IoT) applications and communication enabling technology standards: An overview. In *2014 International conference on intelligent computing applications* (pp. 324–329). <https://doi.org/10.1109/ICICA.2014.73>
9. Rakshit, S. M., Hempel, M., Banerjee, S., & Sharif, H. (2016). Towards a modular iot network model: Low power radio domain use-case in linear-topology wireless sensor network. In *2016 International conference on selected topics in mobile wireless networking (MoWNeT)* (pp. 1–7). <https://doi.org/10.1109/MoWNeT.2016.7496627>
10. Sonntag, S., Schulte, L., & Manner, J. (2013). Mobile network measurements—it's not all about signal strength. In *2013 IEEE Wireless Communications and Networking Conference (WCNC)* (pp. 4624–4629). <https://doi.org/10.1109/WCNC.2013.6555324>
11. Yao, J., Kanhere, S. S., & Hassan, M. (2012). Improving qos in high-speed mobility using bandwidth maps. *IEEE Transactions on Mobile Computing*, 11(4), 603–617. <https://doi.org/10.1109/TMC.2011.97>.
12. Beuran, R., Tan, Y., & Shinoda, Y. (2011). Challenges of using wireless network testbeds: A case study on ORBIT. In *2011 Proceedings ACM workshop wireless networks testbeds experimental evaluation and characterization (WiNTECH)* (pp. 11–18).

Publisher's Note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Lars Møller Mikkelsen received the B.Sc., M.Sc. and Ph.D. in electronics engineering all from Aalborg University, Denmark, in 2011, 2013 and 2017 respectively. He has worked on smart city and smart transportation related European projects MOBiNET and BIG IoT where he has developed various IoT based solutions for enhancing the transportation experience for travelers. His research interests include IoT, context information, smart transportation, mobile network performance, M2M service discovery and distributed systems.

Tatiana Kozlova Madsen received Ph.D. in mathematics from Lomonosov Moscow State University, Russia in 2000. Since 2001 she is working at Department of Electronic Systems, Aalborg University, Denmark where she currently holds a position of an Associate Professor in Wireless Networking. Her research interests lie in the wide area of wireless networking, including Quality of Service and performance optimization of converging networks, communication protocols for IoT systems and mesh networks and methods and tools for performance evaluation of communication systems. She has been involved in a number of national and international projects developing network architectures, network protocols and solutions for Intelligent Transportation Systems; Smart Grids; automotive and train industries.

Hans-Peter Schwefel is Professor for Communication Networks at Aalborg University, Denmark, and managing director of the startup company GridData (www.griddata.eu). In parallel to his position at Aalborg University, he has previously been Scientific Director of the Research Center for Telecommunications (Forschungszentrum Telekommunikation Wien—FTW) in Vienna, Austria in 2008–2016. His research focuses on IP-based communication networks and their applications for critical infrastructures with main interest in performance and dependability aspects. Before he joined Aalborg University, he was a project manager at Siemens Information and Communication Mobile, supervising research projects and responsible for the development of technical concepts for next generation mobile networks. He obtained his doctoral degree in the area of IP traffic and performance modeling from the Technical University in Munich, Germany. For his research activities, he also spent extended periods of time at the University of Connecticut, USA, at AT&T Labs, Middletown, USA, and at University of Florence and CNR-ISTI in Pisa, Italy.