e&i Elektrotechnik und Informationstechnik manuscript No.
(will be inserted by the editor)

Dependability for the Internet of Things

From Dependable Networking in Harsh Environments to a Holistic View on Dependability

Carlo Alberto Boano · Kay Römer · Roderick Bloem · Klaus Witrisal · Marcel Baunach · Martin Horn

Received: date / Accepted: date

Abstract Internet of Things (IoT) applications in several domains such as surveillance of civil infrastructure, smart grids, and smart healthcare are of utmost importance for our society and require *dependable* performance. Guaranteeing that application-specific dependability requirements are met is however still an open research challenge. The IoT indeed exposes highly resource-constrained computing devices to harsh environmental conditions (e.g., heat, mechanical shock, electromagnetic radiation) and physical attacks. Unfortunately, traditional methods to withstand these threats heavily rely on redundancy, a concept that is incompatible with the resource constraints of common IoT devices.
In this article, we illustrate our efforts in providing methods and tools to predict, guarantee, and raise the level of dependability of the IoT. We first outline our contributions in the area of dependable wireless networking and describe a cost-effective solution allowing to guarantee that IoT applications meet specific performance requirements despite the challenging interaction of low-power wireless networks with their surrounding environment.
We then argue that dependable networking alone is insufficient to guarantee the correct operation of a complex IoT system, and outline how we join different scientific disciplines in a long-term endeavor and work towards a coherent view of dependability.

Keywords Dependability · Environmental Impact · Internet of Things · Performance Guarantees

Carlo Alberto Boano, Kay Römer, and Marcel Baunach
Graz University of Technology, Austria
Institute for Technical Informatics
E-mail: {cboano,roemer,baunach}@tugraz.at

Roderick Bloem
Graz University of Technology, Austria
Institute of Applied Inform. Processing and Communications
E-mail: roderick.bloem@iaik.tugraz.at

Klaus Witrisal
Graz University of Technology, Austria
Signal Processing and Speech Communication Laboratory
E-mail: witrisal@tugraz.at

Martin Horn
Graz University of Technology, Austria
Institute of Automation and Control
E-mail: martin.horn@tugraz.at

1 Introduction

The Internet of Things (IoT) is a key enabling technology for applications with high societal relevance and impact. Application domains include smart cities, making the life in dense urban environments more comfortable; smart cars, increasing driving safety and comfort; smart factories, controlling and optimizing production processes; smart grids, improving the efficiency of production, distribution, and consumption of energy; as well as smart buildings, maximizing the comfort of its inhabitants and reducing energy consumption.

These attractive applications represent a long-term investment and are only feasible if the underlying IoT technology does not fail. Any failure in meeting application-specific requirements and in conveying information about the state of things and places in a reliable, timely, and energy-efficient manner may result in high costs, insufficient user satisfaction, and physical damage to people or things.

The key challenge is that in the traditional Internet powerful servers are sheltered in air-conditioned data centers directly connected to power plants, whereas the Internet of Things requires highly resource-constrained computers to be embedded into adverse environments. Smart objects are indeed often directly exposed to heat, humidity, mechanical shock, electromagnetic radiation, and physical attacks, and some of these *hostile environmental conditions* can drastically affect the overall system performance. Electromagnetic radiation and ambient temperature, for example, have

a profound impact on the achievable performance of low-power wireless sensor networks used to sense and convey information about the state of things [5]. For example, radio interference from co-located Wi-Fi and Bluetooth networks may cause a significant message loss, which in turn leads to an increase in end-to-end latency and energy consumption [6]. Variations in temperature over time and space affect the operation of electrical and electronic components and can have a significant impact on clock drift, battery capacity and discharge, as well as on the efficiency of low-power radios [9].

Adverse environmental conditions have typically a strong impact on the reliability and energy-efficiency of IoT communication and are not only hard to predict for a given deployment site, but they may also largely vary from one deployment site to another, thus hindering the scalable deployment of IoT applications. As a result, the development of IoT solutions is severely constrained and often limited to non-critical monitoring applications. Typical methods to avoid failures and to increase fault-tolerance rely heavily on redundancy, which however collides with the resource constraints of common IoT devices and with the space/cost requirements for additional backup hardware. For this reason, withstanding the severe threats posed by adverse environmental effects with scarce resources remains an open challenge.

In this article, we illustrate our efforts in tackling this challenge and providing methods and tools to predict, guarantee, and ultimately raise the level of dependability of the IoT. In Sect. 2, we describe a systematic framework that enables the development of dependable IoT applications by taking into account the challenging interaction of IoT platforms and protocols with the surrounding environment, and show how it can be employed to mitigate the impact of temperature variations on low-power wireless networks.

We then describe in Sect. 3 additional dependability threats for typical IoT systems such as *complexity* (e.g, scaling bugs) and *physical attacks*, arguing that different scientific disciplines should join forces and work towards a coherent view of dependability. We provide examples of open research challenges (e.g., how to systematically ascertain that independently developed smart things interact correctly and robustly) and delineate our research road-map for the systematic construction of a dependable IoT that is resilient against failures and attacks. We finally conclude this article in Sect. 4 with a summary of our contributions.

2 Dependable IoT Networking despite Harsh Environmental Conditions

In order to allow IoT applications to withstand adverse environmental conditions, we have studied methods to increase the reliability and energy-efficiency of networks deployed in harsh settings. In the context of the RELYonIT project[1], we have created a systematic framework that enables the development of dependable IoT applications by taking into account the challenging interaction of IoT platforms and communication protocols with the surrounding environment [8]. To this end, we followed the methodology illustrated in Fig. 1 and devised parametrizable *environmental models* that capture how environmental properties (e.g., temperature and radio interference) vary over time and *platform models* that capture how these environmental properties affect the operation of a hardware (HW) platform. A specification language allows a user to specify *dependability requirements* for a given application that drive the automatic selection and parametrization of *environment-aware communication protocols* such that performance requirements can be met for a given environment and HW platform (or the infeasibility of these requirements is detected). If environmental properties change at runtime, the framework automatically adapts protocol parameters to reflect the new environmental model.

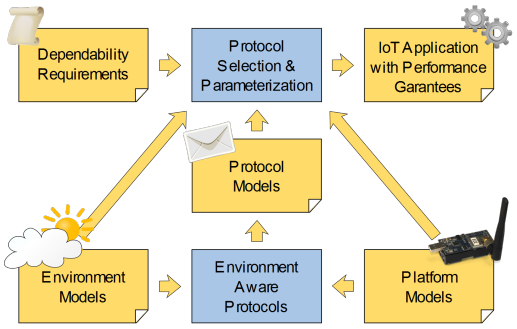

Fig. 1 Methodology employed to build a systematic framework enabling the development of dependable IoT applications despite the interaction of IoT HW platforms with their surrounding environment.

We illustrate next an example of how we employed this framework to mitigate the impact of temperature variations on low-power wireless networks deployed outdoors [16].

Mitigating the impact of temperature variations on low-power wireless networks. Real-world deployments of IoT systems have shown how low-power wireless sensors deployed outdoors often experience high on-board temperature fluctuations, especially if they are placed inside IR-transparent enclosures exposed to direct sun radiation [9]. These large temperature variations can have a severe impact on the operation of carrier sense multiple access (CSMA) protocols, because they can reduce the effectiveness of clear channel assessment (CCA) and compromise the ability of a sensor node to avoid collisions and to successfully wake-up from low-power mode [7]. At high temperatures, indeed, the efficiency of low-power radios reduces significantly: as a result, the signal strength between two wireless sensor nodes A and B decreases by up to 10 dB when the on-board tem-

[1] http://www.relyonit.eu/

perature of both nodes increases from 5 °C to 55 °C [9]. State-of-the-art CSMA protocols typically compare the received signal strength s_r to a static CCA threshold ζ to determine if a node should remain awake to receive a packet or if it should return to sleep mode. As illustrated in Fig. 2, the decrease in signal strength induced by an increase in temperature may lead to a situation in which s_r decreases and becomes lower than ζ. If this happens, the receiver node remains constantly in sleep mode, causing the disruption of the wireless link [7].

In order to allow IoT applications to meet their performance requirements despite temperature variations, we employ the framework shown in Fig. 1 to find a suitable configuration of ζ. To this end, we derive three models:

1. An *environmental model* capturing the relevant aspects of the environment. Such model can be simply based on on-board temperature ranges recorded on the sensor nodes at specific times of the day. If we sub-divide each day into intervals of equal length and run a data collection application prior deployment capturing the on-board temperature variations over several days, we can obtain a model that outputs the minimum/maximum expected temperature as a function of the time of the day.
2. A *platform model* mapping environmental parameters to variables that are relevant for the operation of IoT hardware: such models would capture the relationship between the on-board temperature of sender and receiver nodes and the attenuation of the received signal strength for the HW in use. Denoting PL as the path loss between a transmitter-receiver pair, P_t as the transmission power, $P_\mathrm{r} = P_t - PL$ as the received power, and P_n as the noise floor at the receiver, the impact of temperature on the signal-to-noise ratio (SNR) can be described as:

$$\begin{aligned} SNR &= (P_\mathrm{t} - \alpha\Delta T_\mathrm{t}) - (PL + \beta\Delta T_\mathrm{r}) \\ &\quad -(P_\mathrm{n} - \gamma\Delta T_\mathrm{r} + 10\log_{10}(1 + \tfrac{\Delta T_\mathrm{r}}{T_\mathrm{r}})) \\ &= (P_\mathrm{r} - \alpha\Delta T_\mathrm{t} - \beta\Delta T_\mathrm{r}) \\ &\quad -(P_\mathrm{n} - \gamma\Delta T_\mathrm{r} + 10\log_{10}(1 + \tfrac{\Delta T_\mathrm{r}}{T_\mathrm{r}})) \end{aligned} \quad (1)$$

   where the constants α, β, and γ with units dB/K denote respectively the effect on transmitted power, received power, and on the noise floor, the values T_t and T_r represent the reference temperature of transmitter and receiver; whilst ΔT_t and ΔT_r capture the difference of current temperature with respect to T_t and T_r [9].
3. A *protocol model* describing how the operations of the employed protocol are affected by temperature changes. In our case, the packet reception rate (PRR) in IoT CSMA-based protocols such as ContikiMAC [10] can be estimated by analyzing how the signal strength s_r with which the packet is received relates to the selected CCA threshold ζ and to the transitional phase of the radio response. If $s_\mathrm{r} \geq \zeta$, the node infers that an ongoing transmission is present and remains awake to receive the packet; if $s_\mathrm{r} < \zeta$, the node infers that there is no ongoing transmission and returns to sleep mode without receiving the packet. The PRR further decreases in the transitional region according to a sigmoid curve f in which the probability p of receiving a packet is $p = (1 - f(P_\mathrm{r} - P_\mathrm{n}))^b$ with P_r being the received signal strength, P_n the sensitivity threshold of the radio, and b the number of bits in the packet [20]. Denoting $s_{0.99}$ and $s_{0.01}$ as the signal strength that leads to a delivery rate of 0.99 and 0.01, respectively, we can define three reception regions, as shown in Fig. 3: a connected region, where the received signal strength is above $s_{0.99}$; a disconnected region, where the signal strength is below $s_{0.01}$, and a transitional region of length τ dB, where the delivery rate drops monotonically between 1 and 0.
   Because of the dependency between signal strength and temperature, a variation in the on-board temperature at the receiver or at the transmitter will cause the receiver to measure a signal strength $s_\mathrm{r}^{'} = (s_\mathrm{r} - \alpha\Delta T_\mathrm{t} - \beta\Delta T_\mathrm{r})$ (see Eq. 1), i.e., an increase (decrease) in T_t and/or T_r will attenuate (strengthen) s_r into $s_\mathrm{r}^{'}$. Fig. 3 shows an example in which the received signal strength s_r decreases (i.e., is shifted to the left) due to an increase of temperature in both transmitter ($\alpha\Delta T_\mathrm{t}$ component) and receiver ($\beta\Delta T_\mathrm{r}$ component). To predict if a change in the on-board temperature affects packet reception, we need to verify if

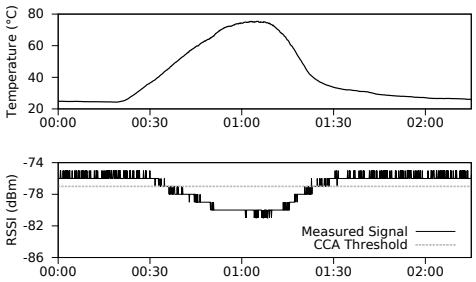

Fig. 2 High temperatures affect low-power radios and decrease the received signal strength. This can cause the received signal strength s_r to drop below the CCA threshold ζ (dotted line) as a result of temperature increase that can compromise link connectivity [7].

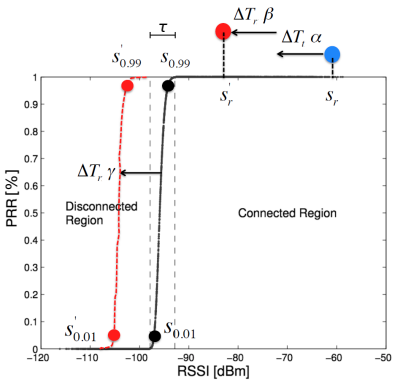

Fig. 3 Temperature impact on CSMA-based IoT MAC protocols [16].

$s_{\mathrm{r}}^{'} < \zeta$. If this is the case, no packet will be received, as the device will return to sleep mode after having assumed no ongoing transmission. Similarly, if the onboard temperature of the receiver changes, also the position of the sigmoid curve may change. To predict how $s_{0.01}$ and $s_{0.99}$ would change in relation to temperature variation we use Eq. 1 to derive $s'_{0.99} = s_{0.99} - \Delta T_{\mathrm{r}}\gamma$ and $s'_{0.01} = s_{0.01} - \Delta T_{\mathrm{t}}\gamma$.
Finally, we can estimate the packet delivery rate PRR' given a specific ζ value for each link i in the network as:

$$PRR' = \begin{cases} 1, & \text{if } \max\{\zeta, s'_{0.99}\} < s'_{\mathrm{r}} \\ p, & \text{if } s'_{0.01} \leq \zeta < s'_{\mathrm{r}} \leq s'_{0.99} \\ 0, & \text{otherwise} \end{cases} \quad (2)$$

and use this model to estimate the worst case delivery rate given a specific temperature variation/range.

These three models allow predictions of the performance of a low-power wireless network in presence of temperature variations given specific configurations of ζ. To automatically identify a (near-)optimal configuration that satisfies the dependability requirements of an IoT application, the framework employs mathematical optimization. Given *user-defined dependability requirements* as input, the framework outputs optimal parameters for the employed communication protocol. In our case, the framework exposes one configuration parameter, ζ, and provides a number of metrics that can be used to define goals or constraints. The primary metric is the worst-case PRR of a link that can be accepted by the application of interest. Another metric is minimal energy consumption to ensure a long system lifetime. These requirements lead to an optimization problem such as:

$$\begin{aligned} &\text{Maximize } \mathrm{CCA}([\zeta]) \\ &\text{Subject to } \mathrm{PRR}([\zeta]) \geq 0.85 \text{ with probability } 1.00 \\ &\qquad\quad \mathrm{PRR}([\zeta]) \geq 0.95 \text{ with probability } 0.9 \end{aligned} \quad (3)$$

where ζ should be maximized while not violating the constraints on PRR (in this example at least 85% at all times and at least 95% in at least 90% of the cases). Maximizing CCA ensures low energy-consumption, as lower CCA values lead to a higher number of false wake-ups and an increased radio usage. A real-world evaluation has shown that the framework is indeed able to pick the most optimal CCA threshold given a set of dependability requirements and correctly parametrize IoT protocols such that the impact of harsh temperature variations can be predicted and minimized [16].

In our research, we developed further models and employed the described framework to configure other protocol parameters and increase the resilience of low-power networks to surrounding radio interference and optimize duty cycling selection in order to minimize energy expenditure. Further details can be found in [8].

3 Beyond Dependable IoT Networking

The framework presented in Sect. 2 enables resilient low-power wireless communication despite harsh environmental conditions. However, dependability is the combination of several attributes that allow a user to put trust into and rely on a system [3]. Such attributes are *reliability* (i.e., continuity of correct, accurate, and timely service), *availability* (i.e., readiness for correct service), *safety* (i.e., absence of catastrophic consequences on users and environments); *confidentiality* (absence of unauthorized disclosure of information); and *integrity* (i.e., absence of improper system alteration).

In the IoT context there are several dependability threats that can affect those attributes. In the previous sections, we already discussed how *harsh environmental conditions* represent a major threat for the reliability and availability of a system due to their impact on communication performance. Similarly, two other common threats to all IoT systems are:

- *Physical and remote attacks.* IoT devices are deployed everywhere, and attackers can not only mount attacks remotely via network interfaces, but also physically (e.g., by performing dynamic fault induction [13], or by collecting information through side-channels [15]).
- *Complexity.* The IoT is a complex system (of systems) where many devices with continuously updated software and services cooperate using a dynamically changing communication network and where the number of devices is not known in advance. These properties make designed IoT applications prone to design and implementation flaws, as well as scaling bugs.

As dependability is the combination of several attributes, it is sufficient that one of them is poorly addressed to affect the overall system performance. An analogy is the most fragile link of a chain (the one that is most likely to break) compromising the usability of the whole system when splitting apart. To build a dependable IoT that is resilient against failures and attacks, it is hence not sufficient to consider only dependable networking in harsh environments as discussed in Sect. 2, but to address all relevant functions and threats of the IoT in an interdisciplinary fashion. We describe next our efforts in this regard by providing an overview of our joint research activities in the areas of dependable localization and communication (Sect. 3.1), dependable embedded computing (Sect. 3.2), composition of smart objects (Sect. 3.3), and dependable networked control (Sect. 3.4). These efforts are carried out in the DependableThings[2] project funded by Graz University of Technology.

[2] `http://dependablethings.tugraz.at/`

3.1 Dependable Wireless Localization and Communication

Wireless technologies suffer from physical and man-made impairments (e.g., multi-path propagation and interference from competing transmissions, as well as from the effect of temperature variations and other environmental properties): this impairs the accuracy, latency, loss, and energy consumption of wireless services. A key challenge is therefore to offer statistical guarantees on the reliability and availability of correct wireless localization and communication by automatically adapting system parameters, using models of the transceiver hardware and the environment.

In our research, we employ location-resolved models of the environment [19] to obtain robustness and scalability, as well as adaptive radio front-ends (i.e., tunable filters and antennas) to support low-power operation [2], in contrast to power-hungry software defined radios that are normally considered for flexible (cognitive) radios. We then map the problem to a model-predictive control system to gain control over the dependability requirements. The control loop includes the adaptable radio front-ends, physical-layer signal processing for environment modeling/mapping, robust wireless communication and localization, and communication protocols for distributed control of the radio transceivers.

3.2 Dependable Embedded Computing

The IoT requires all kinds of connected computing devices to execute software dependably: operations have to be completed within guaranteed response times, functions must be immune to environmental perturbation or attacks, and secret information must not be revealed via physical side-channels or communication interfaces. While security, reactivity, and dynamic modularity are still considered independently in today's embedded system design, this separation has unacceptable implications on the resilience, versatility, and longevity of IoT devices. We hence investigate an integrated approach that spans across all system layers.

Our research focuses on co-designed hardware and software that jointly handle the inherent complexity of both physical attack scenarios and modular real-time applications, accounting for the additional support for dynamically changing software and service composition. We tackle the co-design of processor architectures and operating systems as a base for dependable applications and services: in contrast to established concepts for comparable embedded systems [14], we focus on the solid anchoring of security, real-time, and modularity features throughout the entire stack.

3.3 Dependable Composition

In the IoT, smart things collaborate to provide services. They do so by using protocols that are vital for the safe and secure operation of the system, but that are not always documented well and rarely implemented correctly. Indeed, bugs such as Heartbleed [11] have achieved notoriety because of their deleterious effects, and studies show that implementations of important protocols like TLS are rife with bugs [4].

Our research activities in this domain aim to develop methods and tools to automatically find bugs in communication protocols as efficiently as possible. We focus on observing and experimenting with the behavior of implementations under the assumption that a full specification is not available. Towards this goal, we use language modeling techniques to learn models of protocol implementations [1]. This model can be used for different purposes: we can compare learned models and flag differences as suspicious behavior; we can use a model for model-based fuzzing of the implementation; we can formally verify the learned model; or we can use the model for runtime verification and enforcement to detect attacks that use the protocol in untypical ways.

3.4 Dependable Networked Control

As communication between smart items is prone to errors and likely to be corrupted by unpredictable distortions and losses, the stability and performance of the respective feedback loops have to be robust with respect to these phenomena which are inherent to the IoT. The fact that conventional control theories are based on ideal assumptions such as non-delayed actuation and sensing and perfect synchronization motivates the need for innovative methods for the design of dependable systems

Towards this goal, we propose an information-theoretic approach to networked control building on our earlier results on the characterization of information processing in deterministic input-output systems using information loss as a key parameter [12]. We focus on a recovery of disturbed or lost interconnecting signals based on information theoretic principles [18]. We plan to adapt methods from interactive real-time multimedia communication to automatic control, and aim to investigate the potential of these methods to increase the stability margins. In contrast to related theoretic works [17], our approach is intended for practical applications and constrained CPUs.

4 Conclusions

Providing methods and concepts to guarantee that IoT applications meet specific dependability requirements is a hot research topic. In this article, we have illustrated our efforts in the area of low-power wireless networking and described a framework that helps guaranteeing that specific performance requirements can be met despite the impact of the surrounding environment. We have further analyzed common dependability threats for IoT systems and outlined our

research road-map for the systematic construction of a dependable IoT that is resilient against failures and attacks.

Acknowledgements This work was performed within the LEAD-Project "Dependable Internet of Things in Adverse Environments", funded by Graz University of Technology, and within the EU FP7 project RELYonIT, funded by the European Commission.

References

1. Aichernig, B., Bloem, R., Pernkopf, F., Röck, F., Schrank, T., Tappler, M.: Poster: Learning models of a network protocol using neural network language models. IEEE Security & Privacy (2016)
2. Aigner, R.: Tunable filters? reality check foreseeable trends in system architecture for tunable RF filters. IEEE Microwave Magazine **16**(7) (2015)
3. Avižienis, A., Laprie, J.C., Randell, B., Landwehr, C.: Basic concepts and taxonomy of dependable and secure computing. IEEE Transactions on Dependable and Secure Computing **1**(1) (2004)
4. Beurdouche, B., Bhargavan, K., Delignat-Lavaud, A., Fournet, C., Kohlweiss, M., Pironti, A., Strub, P., Zinzindohoue, J.: A messy state of the union: Taming the composite state machines of TLS. IEEE Security & Privacy (2015)
5. Boano, C.A.: Dependable wireless sensor networks. Ph.D. thesis, Graz University of Technology, Graz, Austria (2014)
6. Boano, C.A., Römer, K.: External radio interference. In: Radio Link Quality Estimation in Low-Power Wireless Networks, SpringerBriefs in Electrical and Computer Engineering (2013)
7. Boano, C.A., Römer, K., Tsiftes, N.: Mitigating the adverse effects of temperature on low-power wireless protocols. In: Proc. of the 11^{th} IEEE MASS Conf. (2014)
8. Boano, C.A., Römer, K., et al.: RELYonIT: Publishable Summary Report. Tech. rep. (2015)
9. Boano, C.A., Wennerström, H., et al.: Hot Packets: A systematic evaluation of the effect of temperature on low power wireless transceivers. In: Proc. of the 5^{th} ExtremeCom Conf. (2013)
10. Dunkels, A.: The ContikiMAC radio duty cycling protocol. Tech. Rep. T2011:13, Swedish Institute of Computer Science (2011)
11. Durumeric, Z., Kasten, J., Adrian, D., Halderman, J., Bailey, M., Li, F., Weaver, N., Amann, J., Beekman, J., Payer, M., Paxson, V.: The Matter of Heartbleed. In: Proc. of the ACM IMC Conf. (2014)
12. Geiger, B.: Information loss in deterministic systems. Ph.D. thesis, Graz University of Technology, Graz, Austria (2014)
13. Joye, M., Tunstall, M. (eds.): Fault Analysis in Cryptography. Information Security and Cryptography. Springer (2012)
14. Malenko, M., Baunach, M.: Real-time and security requirements for the internet of things operating systems. In: Proc. of the ECHTZEIT Conf. (2016)
15. Mangard, S., Oswald, E., Popp, T.: Power analysis attacks - revealing the secrets of smart cards. Springer (2007)
16. Oppermann, F.J., Boano, C.A., Zúñiga, M.A., Römer, K.: Automatic protocol configuration for dependable internet of things applications. In: Proc. of the 10^{th} IEEE SenseApp Workshop (2015)
17. S. Mastellone, C.T.A., Dorato, P.: Stochastic Control over Finite Capacity Channels: Causality and Feedback. In: Proc. of the ECC2009 Conf. (2009)
18. Tang, P.L., de Silva, C.: Compensation for transmission delays in an ethernet-based control network using variable-horizon predictive control. IEEE Trans. on Control Systems Techn. **14**(4) (2006)
19. Witrisal, K., Meissner, P., et al.: High-accuracy localization for assisted living: 5G systems will turn multipath channels from foe to friend. IEEE Signal Processing Magazine **33**(2) (2016)
20. Zúñiga, M.A., Krishnamachari, B.: Analyzing the transitional region in low-power wireless links. In: Proc. of the 1^{st} IEEE SECON Conf. (2004)