

A MoDeST Evaluation of the Interplay between Energy Consumption and Clock Drift in ZigBee

CHRISTIAN GROSS, HOLGER HERMANNNS AND REZA PULUNGAN

Abstract. Wireless embedded sensor networks are predicted to provide attractive application possibilities in industry as well as at home. IEEE 802.15.4 and ZigBee are proposed as standards for such networks with a particular focus on pairing reliability with energy efficiency, while sacrificing high data rates.

IEEE 802.15.4 is configurable in many aspects, including the synchronicity of the communication, and the periodicity in which battery-powered sensors need to wake up to communicate. This paper develops a formal behavioral model for the energy implications of these options. The model is modularly specified using the language MODEST, which has an operational semantics mapping on stochastic timed automata. The latter are simulated using a variant of discrete-event simulation implemented in the tool MÖBIUS. We obtain estimated energy consumptions of a number of possible communication scenarios in accordance with the standards, and derive conclusions about the energy-optimal configuration of such networks. As a specific fine point, we investigate the effects of drifting clocks on the energy behavior of various application scenarios.

Key words and Phrases: Sensor networks, formal modelling, distributed coordination, power-aware design, clock drift.

INTRODUCTION

Quantitative analysis of ad hoc and wireless networks has in the past been concentrating on scalability and routing questions [6]. The predominantly applied techniques are based on simulation using enhanced tools such as GloMoSim or NS-2 [24, 5], Omnet [22] or commercially available simulation tools, such as Opnet.

2000 Mathematics Subject Classification:

Another approach is based on instruction-level simulation of the actual microcontroller codes [23, 21].

The credibility of simulation results obtained using the above enhanced modelling tools seems not to be free of doubts. Such studies generally suffer from (1) unclarities of how simulation models are obtained from the modelling language, (2) the sheer number of parameters with non-obvious effects adjustable by the user, some of them having hidden effects on the simulation outcomes, (3) excessive simulation times needed to simulate the ensemble of many protocol stacks and states, and (4) the impossibility to validate simulation results through reproducible real-life experiments.

As a matter of fact, some recent articles have criticized the extremely poor quality and reproducibility of simulation-based experimentations with ad hoc or wireless networks [13, 6]. Indeed, this questions the validity of simulation-based predictions for this area as a whole.

The work presented in this paper is among the few which attempt to attack the above mentioned principal problems. Other loosely similar approaches include [17, 12]. While we still rely on discrete-event simulation as our analytic workhorse, we proceed in a drastically different way. The main difference is that (1) we use a language with a strictly formal semantics, which is equipped with well-established abstraction techniques. Consequently, the underlying stochastic model for simulation is well-defined and the obtained simulation results are trustworthy. (2) We expose all assumptions explicitly, since they are part of the formal system specification. (3) We do not model entire protocol stacks, but work with well-justified abstractions of lower layer effects. We consider worst-case scenarios, if no other information is available. But (4) we do not yet provide real-life experiments to back up our simulation results.

The model is modularly specified using the language MODEST, which has an operational semantics mapping on stochastic timed automata. The latter are simulated using a variant of discrete-event simulation implemented in the tool MÖBIUS. In this paper, that approach is applied to the IEEE 802.15.4 and ZigBee standards. This is a protocol family dedicated to low-bandwidth sensor networks operating on battery. IEEE 802.15.4 is configurable in many aspects, including the synchronicity of the communication, and the periodicity in which battery-powered sensors need to wake up to communicate. The particular configuration chosen has obvious – and non-obvious – implications on the lifetime of battery-powered devices. An obvious rule of thumb is, for instance, that battery-operated devices can survive longer timespans if they need to wake up less often. This paper investigates the non-obvious rules.

We obtain estimated energy consumptions of a number of possible communication scenarios in accordance with the standards, and derive conclusions about the energy-optimal configuration of such networks. In particular, we investigate the effects of time-slotted and unslotted medium access techniques, and their interplay with drifting clocks. Our observations allow us to establish rules of the following kind: (1) Unslotted CSMA/CA is more favorable w.r.t. energy saving

than slotted CSMA/CA. (2) If devices using GTSSs and CSMA/CA coexist, those operating in GTSSs expend considerably less energy. (3) Small clock drifts can have far over-proportional effects on energy consumption, but with only minor adverse effects on battery lifetimes. For sure, power consumption does not only depend on clock drift and synchronization policy. It also depends on many other factors, such as link quality and other environmental conditions. For our studies, these are assumed constant, since we see no way to include them in our studies without losing focus. Further, we do not include comparison with experimental measurements. The reason is that (1) controlling physical clock drifts is virtually impossible with available hardware, and (2) real-life experiments would take several months or years to show measurably distinct effects on battery lifetimes. The main contribution of this paper is that it pioneers a model-based analysis of the interplay of clock drift and energy cost in sensor networks. The paper is organized as follows: Section provides a brief introduction to ZigBee and IEEE 802.15.4. In Section , the modelling formalism we employ is described together with the tool chain supporting it. Section describes the modelling of ZigBee and IEEE 802.15.4 in MODEST. We also clarify the modelling assumptions we make and discuss in detail a particular model representative. In Section , we describe and discuss the simulation experiments and their results. Section concludes the paper.

PROBLEM DOMAIN

This section provides a general introduction to ZigBee and IEEE 802.15.4, focussed on the characteristics and features that are important for the scope of the paper. For more detailed information, the interested reader is invited to consult the standard documents [2, 1].

ZigBee

ZigBee is a wireless communications standard for low-cost and low-power consumption networks, based on the IEEE 802.15.4 standard (see Section). The standard was developed by ZigBee Alliance, which is a cooperation of a number of industrial enterprises established in 2002. The ZigBee Specification Version 1.0 [2] was ratified in December 2004, and is publicly available since June 2005.

ZigBee-compliant products will typically be embedded in consumer electronics, home and building automation, industrial controls, PC peripherals, medical sensor applications, toys and games [2]. The strict emphasis on low power consumption allows implementation of wireless communications with devices operating for several years without needing battery replacement.

The architecture of ZigBee follows the layering principle. Figure 1 depicts its simplified stack architecture [2]. Both physical and medium access control (MAC) layers are defined in IEEE 802.15.4 standard. The ZigBee standard, on the other hand, provides the definition of the Network layer and the framework for the Application layer.




Figure 1: Simplified ZigBee Stack Architecture

Beside data delivery between layers (Network Message Broker–NMB), the Network layer provides mechanisms for joining and leaving networks, together with starting a new network and address assignment in ZigBee coordinator (Network Management–NM); applying security to data frames (Network Security Management–NSM); and routing-related activities, including neighbors discovery, creating and maintaining routing information between devices and routing the data frames (Routing Management–RM).

In the Application layer, Application Support sublayer is responsible for maintaining tables for binding, namely matches between two devices based on their services and needs, and forwarding messages between these devices. The ZigBee Device Object defines the role of the devices within the network, binding activities, establishing secure relationships between devices, discovering devices on the network and determining which services they provide.

ZigBee supports star, tree and mesh network topologies. In these topologies, a device called ZigBee coordinator controls the network. In the star topology, all devices communicate with the coordinator, while in tree and mesh topologies, communication can be mediated by ZigBee routers. Peer-to-peer communication is allowed in mesh topology.

IEEE 802.15.4

IEEE 802.15.4 is an open standard for ultra-low complexity, cost, power consumption and low data rate wireless connectivity among inexpensive devices in wireless personal area networks (WPAN) [1]. The standard specifies the physical and the medium access control (MAC) layers.

Devices participating in WPAN can be distinguished into full-function devices (FFD) and reduced-function devices (RFD). Communication between FFDs and RFDs is possible, but an RFD can only communicate with an FFD. An FFD may become a personal area network (PAN) coordinator, a coordinator, or a device. There are two topologies where a WPAN can operate: star and peer-to-peer. In

the star topology a PAN coordinator controls all communications between devices. In the peer-to-peer topology, any device can communicate with any other device as long as they are in the range of each others.

The physical layer specifies mechanisms for (de)activating the radio transceiver, detecting the channel's energy level, indicating the link quality of received packets, performing Clear Channel Assessment (CCA) for CSMA/CA, selecting channel frequency and data transmission and reception. The standard maintains that compliant devices shall operate on 2.4 GHz, 915 MHz and 868 MHz bands. The 2.4 GHz band, which is available worldwide, consists of 16 channels, each with data rate 250 kbps. The 915 MHz band is available in the North America and has 10 channels, each with data rate 40 kbps. The 868 MHz band has only one channel with data rate 20 kbps and is available in Europe. In all three bands, the radio transmissions use the direct sequence spread spectrum coding with chip rate 2000, 600 and 300 kcps, respectively. The 2.4 GHz band uses the Orthogonal Quadrature Phase-Shift Key modulation, while the others use the Binary Phase-Shift Key. As a result, the symbol rates of the three bands are 62.5, 40 and 20 ksymbols/s, respectively.

Beside managing access to the physical radio channel and providing a reliable link between peer MAC entities, the MAC layer is responsible for generating network beacons (for coordinators), synchronizing to the beacons, supporting PAN (dis)association, supporting device security, and employing the channel access mechanism.

The functional characteristic of low-rate PAN can be distinguished into beacon-enabled and nonbeacon-enabled networks. The simplest manner of operation is nonbeacon-enabled, where the network operates by using the (unslotted) Carrier Sense Multiple Access with Collision Avoidance (CSMA/CA). CSMA/CA is a multiple access protocol similar to CSMA/CD. As opposed to CSMA/CD, which operates by listening to the channel while sending in order to detect collisions, CSMA/CA tries to avoid collisions by listening to the channel for a predetermined amount of time prior to transmissions. If a collision anyhow occurs, it can be determined by the absence of acknowledgements from the receiver. CSMA/CA is more suitable for wireless communication since listening to the medium is energy-expensive, and the spatial setup might not allow to witness collisions anyhow.

In beacon-enabled mode, on the other hand, the coordinator periodically emits beacon signals, which provide a frame of reference for a time-slotted access to the medium. More precisely, a so-called superframe structure, which is defined by the coordinator, is used. Inside of the superframe structure, communication can be carried out with a guaranteed time slot (GTS) mechanism or with slotted CSMA/CA mechanisms. Devices using GTS and those using CSMA/CA may coexist, as we will explain below.

The structure of a superframe is depicted in Figure 2. The coordinator broadcasts network beacons regularly to all devices. These beacons mark the beginnings and the ends of superframes. The beacons can be used for synchronization purpose, to identify the PAN, to describe the structure of the superframe, as well as

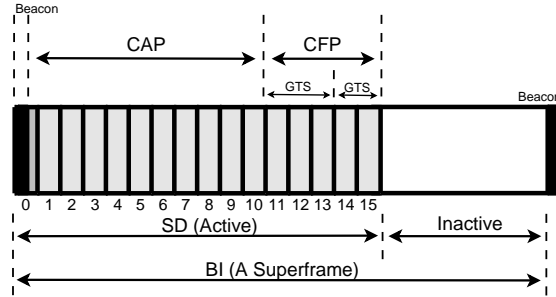


Figure 2: The Superframe Structure

to announce the GTS allocations.

The coordinator can divide a superframe into an active and an inactive portion. The active portion is further divided into 16 equally sized slots. The coordinator may decide to allocate up to 7 of these slots as GTSs. These GTSs form the Contention-Free Period (CFP) and must appear at the end of the active portions. The rest of the slots forms the Contention-Access Period (CAP), in which devices compete using a slotted CSMA/CA mechanism.

The length of a superframe (Beacon Interval–BI) is determined by the coordinator by varying the value of Beacon Order (BO) which influences the length exponentially as follows

$$BI = \text{aBaseSuperframeDuration} \times 2^{\text{BO}}, 0 \leq \text{BO} \leq 14,$$

where $\text{aBaseSuperframeDuration} = 960$ symbols. The duration of active portion (Superframe Duration–SD) is set by varying the value of Superframe Order (SO) and follows

$$SD = \text{aBaseSuperframeDuration} \times 2^{\text{SO}}, 0 \leq \text{SO} \leq \text{BO} \leq 14.$$

A device wishing to transmit using unslotted CSMA/CA mechanism, begins by setting $\text{NB} = 0$ and $\text{BE} = \text{macMinBE}$, where NB is the number of performed backoff so far, and BE is the backoff exponent. According to the standard, the value of macMinBE is in the range $\{0 \dots 3\}$ with default 3. The device then randomly selects a number r from the interval $[0, 2^{\text{BE}} - 1]$ and waits for r backoff periods before performing a clear channel assessment (CCA). If the channel is idle, the device transmits immediately.

If the channel is not idle then another backoff period must expire before a new CCA is performed. For this, the value of NB is increased by one, and similarly with BE as long as it does not exceed $\text{aMaxBE} = 5$. Thus, the upper bound of the interval, from which the random number is selected, grows exponentially every time the CCA finds the channel busy. However, this does not go on indefinitely. The number of backoff (NB) is limited by macMaxCSMABackoff , a value in the range

$\{0 \dots 5\}$ with default value 4. Once NB reaches this limit, the channel access attempt is considered failed, and the device has to start the transmission anew.

In a beacon-enabled network, slotted CSMA/CA can be employed during the CAP. The backoff period boundaries of each device in the PAN are then aligned with the superframe slot boundaries, and a device can only start transmitting on the boundary of a backoff period. A backoff period, in both slotted and unslotted mode, is the same as the required time to transmit `aUnitBackoffPeriod` = 8 symbols.

Beside the two variables used in the unslotted version, in the slotted version a transmitting device begins by setting $CW = 2$, where the CW is the contention window length. It then tries to locate the backoff period boundary, and proceeds to randomly select some number r from the interval $[0, 2^{BE} - 1]$, and waits for r backoff periods, similar to the unslotted mechanism. The differences are that on finding the channel idle on the first CCA, a further CCA is commenced before the actual transmission is started, and that the CCAs are performed on the boundaries of backoff periods. Every time the CCA confirms that the channel is idle, CW is decreased by one. During these CCAs, a busy channel will trigger backoff.

MODEST AND SUPPORTING TOOLS

The *Modelling and Description Language for Stochastic and Timed Systems* (MODEST) [8] is a specification formalism for stochastic real-time systems. The language is rooted in classical process algebra, *i.e.* the specification of models is compositional. Basic activities are expressed with atomic actions, more complex behavior with constructs for sequential composition, nondeterministic choice, parallel composition with CSP-style synchronization, looping and exception handling. A special construct exists to describe probabilistic choice. Clocks, variables and random variables are used to describe stochastic real-time aspects.¹

All constructs and language concepts have a pleasant syntax, inspired by Promela, LOTOS, FSP, and Java. We will see examples later. MODEST is equipped with a structural operational semantics mapping on stochastic timed automata (STA). The MODEST semantics is described in full detail in [8]. MODEST allows one to describe a very large spectrum of models, including: ordinary labelled transition systems, timed automata, probabilistic automata, stochastic automata [15], Markov decision processes, and various combinations thereof (*cf.* [8]). Remarkably, the language is designed in a way that all these models correspond to syntactic subsets of the language, and can thus be identified while parsing a MODEST specification.

MODEST takes a *single-formalism, multi-solution* approach. This advocates to have a single specification that addresses various aspects of the system under consideration. This is contrary to the more common approach to construct different models to describe different aspects of a system and then analyze these models.

¹MODEST also supports for modelling time variation, especially time non-determinism, but this feature is not used in this case study.

Generally, no guarantee of consistency between these models can be given, be it for lack of a rigorous semantics or because a proper relation between the different model classes is not known. Thus, the validity of results w.r.t. the original system under study is often questionable.

In order to facilitate the analysis of the different models, tool support is indispensable. The MODEST Tool Environment (MOTOR) is a software tool that implements the MODEST semantics and is the central vehicle in the multi-solution analysis of MODEST models. The fundamental idea behind MOTOR is to simplify specifying MODEST models (*e.g.*, by providing a macro-preprocessor), and to translate or adapt the models in a way such that the actual analysis work can be carried out by third-party state-of-the-art tools, such as PRISM [20], UPPAAL [7], or CADP [18].

MOTOR is designed to facilitate easy access to all language features of MODEST, and thus to allow easy extraction of all imaginable model classes from a specification. The design allows straightforward extensibility of the tool. To realize this, MOTOR provides two abstract programming interfaces, one on the level of the abstract syntactic representation of the MODEST specification, and another one which provides backend-developers with access means to the STA defined by the MODEST semantics.

The by now most mature backend of MOTOR is provided by a link to the MÖBIUS performability evaluation environment. MÖBIUS has been developed independently from MODEST and MOTOR at the University of Illinois at Urbana-Champaign [14]. MÖBIUS is designed as an integrated tool environment for the analysis of performability and dependability models. It allows specification of models in different formalisms, based on, for instance, Petri net-like formalisms or Markovian process algebra. The tool provides efficient discrete-event simulation capabilities and numerical solvers, such as Markov chain solvers. From a user perspective, the MOTOR/MÖBIUS tandem enables one to perform simulation of MODEST models, and to gather performability results.

Simulation-based analysis covers the largest language fragment of MODEST: the only concept that cannot be supported by simulation is nondeterminism, in particular of delay durations and non-deterministic choice between actions. We exclude the former by assuming maximal-progress with respect to delays. We do not restrict action nondeterminism, since it is a convenient modelling instrument. However, no mechanisms, like a well-specified-check [16], is implemented yet to ensure the validity of the simulation statistics when action nondeterminism is present.

MOTOR and its connection to MÖBIUS is mature and has been tested in a number of non-trivial case studies. In [19], it has been used for reliability analysis of the upcoming European Train Control System (ETCS) standard. In [9], it has been applied to the analysis of an innovative plug-and-play communication protocol, which has led to a patent application of our industrial partner. In [10], MOTOR has been used for the optimization of production schedules, in combination with timed automata-based schedule synthesis with UPPAAL.

MODELLING

In this section, we describe the modelling of ZigBee and IEEE 802.15.4 in MODEST. First, we provide an overview of the models. At this point, we clarify the parts of the standard being modelled, as well as the parts we abstract from. Subsequently, we describe the framework model, taking the slotted CSMA/CA version as representative. In the end of the section, the modelling of energy consumption and clock precision is explained.

Overview

We model ZigBee/IEEE 802.15.4-compliant personal area networks in a star topology. Each network consists of a single PAN coordinator and a number of stations or devices. We assume that the PAN coordinator has continuous power supply, while the stations do not. A station can be either an FFD or an RFD, attached to a sensor. Periodically, a station communicates with the PAN coordinator, either to transmit its gathered sensor data or to receive instructions, but a station cannot communicate with other stations.

Two separate models, beacon and nonbeacon-enabled PAN models, are developed. The simpler, nonbeacon-enabled model is parameterized by BI, but this is not used to signal beacons. Instead, BI is used to indicate the arrival of messages to each station from its sensor. It is assumed that each station, in both models, always has message to send: almost one slot-length of data every beacon interval.

In the beacon-enabled PAN model, some of the stations use the CAP for CSMA/CA communications, and some use GTS in the CFP, the detailed scenarios considered are described later. The model is parameterized by BI, the beacon interval. We set the superframes to have the same duration as BI, hence there is no inactive period. Every BI-equivalent time units, the PAN coordinator broadcasts a beacon, and all stations must be ready to receive it. Henceforth, the coordinator is ready to receive transmissions from the stations. The stations which are not assigned any GTSs compete with other similar stations to send their messages by using slotted CSMA/CA, while the stations with assigned GTSs wait for their turns.

To save energy, a station goes to sleep mode whenever it has a chance to. For instance, a station sleeps whenever it has no messages to transmit, or while it must wait for its GTS turn or when it is performing backoff. However, a station must always wake up before its turn to transmit or before the beacon is transmitted by the PAN coordinator.

We abstain from modelling the initialization of the PANs. Hence, the selection of the name, the (dis)association of stations into the PAN, and the address assignment of stations by the coordinator are not modelled. Instead, we concentrate on the typical operations of the PAN, when the stations are transmitting messages. The physical layer of IEEE 802.15.4 is not modelled either. For instance, we assume that there is no significant propagation delay and no channel

selection procedure. Nevertheless, some physical layer constants that affect the timing of the communication, for instance the duration of CCAs, are taken into consideration. Furthermore, we model PAN in 2.4 GHz band, which means that all devices transmit at 62.5 ksymbols/s. The medium is assumed uniform, in the sense that all communicating parties have complete knowledge of it.

Models

The model of the PAN in beacon-enabled mode is described in this section, focusing on slotted CSMA/CA, which is more complex and interesting than the GTS case. Complete details are available as a tutorial-style modelling guide [4]. The model consists of two distinct process definitions: `coordinator()` and `station()`, modelling the behaviors of a PAN coordinator and a station, respectively. In all experiments, we set `macMaxCSMABackoff` (the maximum backoff attempts before declaring a channel access failure) to 5 and `macMinBE` (the minimum value of the backoff exponent) to 2.

The System The model of the overall system is depicted in MODEST model 1. The system consists of 11 process instances, one coordinator and ten stations, run in parallel. This is achieved by using the parallel composition construct `par{}`. Pro-

MODEST **model 1** The complete system

```

01 par{
02  :: coordinator()
03  :: relabel {...} by {...} station(1)
  .. ...
12  :: relabel {...} by {...} station(10)
13 }
```

cesses inside of a parallel composition construct run concurrently and synchronize on their common actions, if existing. The relabelling operator `relabel { } by { } p()` relabels the actions in the first set by the actions in the second in a particular instance of process `p()`. This allows multiple instantiations of process definition `station()`.

The Coordinator A simple model of the PAN coordinator is shown in MODEST model 2, especially to highlight some of the clock manipulation features of MODEST. The coordinator has two clock variables: `btimer` modelling the time progress between beacons, and clock `c` modelling the transmission time of a beacon. Clock variables increase linearly with time and can only be reset to zero. The coordinator process begins by immediately sending a beacon (action `sendb_start`). All beacons are of length $52 \mu\text{s}$ (lines 05 & 09), namely the duration of the smallest possible beacon. Action `sendb_end` signals the end of the beacon's transmission. From then on, the coordinator waits until clock `btimer` is equal to the value of `binterval` (line

```

MODEST model 2 The coordinator
01 process coordinator() {
02   clock btimer, c;
03
04   sendb_start {= bintheair=true =} ;
05   when(c==52) sendb_end {= bintheair=false =} ;
06   do{
07     ::when(btimer==binterval)
08       sendb_start {= c=0, btimer=0, bintheair=true =} ;
09       when(c==52) sendb_end {= bintheair=false =}
10   }
11 }

```

07). At this point, a beacon interval has expired and the coordinator broadcasts a new beacon. The coordinator proceeds thus continuously, broadcasting beacons every time a beacon interval expires.

The process makes use of two global variables, *i.e.* variables accessible to all processes in the model. The variables are `binterval`, of type integer, which corresponds to the time it takes for a BI in the standard, namely the length of the time interval between two beacons; and `bintheair`, which is a boolean variable used to indicate to the whole system that a beacon is being transmitted. Delimiters `{= =}` wrap a set of variable assignments. Such sets of assignments are executed atomically at the time instant of the action preceding them.

The Stations The model of the station is shown in MODEST model 3 and 4. In the beginning, a station waits until `bintheair` is true, indicating the beginning of the transmission of a new beacon. At the same time, it resets its main clock and sets `ttosend` (the duration of the remaining data to send) equal to constant `sendingtime = binterval/16 - 194 μs`, almost a slot duration in the used BI. Once the beacon finishes, the station aligns its backoff boundary with the superframe slot boundary by waiting until a multiple of the backoff period (20 symbols, hence 320 μs in 2.4 GHz band) expires since the beginning of the previous beacon (line 20). The station then performs backoff and attempts to transmit. This is repetitively done as long as there is still something to send and enough time to do so (line 17), otherwise (line 15) the station just waits for the next beacon.

The maximum length of messages submittable to the physical layer is 133 bytes, which takes 4256 μs to transmit in 2.4 GHz band. Thus the remaining data is split accordingly (line 19) into `attosend`, the actual time required to send the partitioned message. In line 23, more lines of code are needed to model the random selection of variable `r` than the code indicates. For readability we shorten the code and put `r=//Uniform(0,2^BE-1),//` as a placeholder.

As shown in MODEST model 4, once the duration of the backoff delay is determined, the backoff is performed only if there is still enough time to complete

```

MODEST model 3 The station (Part 1)


---


01 process station(int id) {
.. ...
10 do{
11  ::when(bintheair) beacon_received
12    {= mainclock=0, ttosend=sendingtime =} ;
13    when(!bintheair) start_waiting;
14    do{
15      ::when(ttosend==0 || !enoughtime) do_nothing
16        {= enoughtime=true =} ; break
17      ::when(ttosend>0 && enoughtime) start_csmaca
18        {= NB=0, CW=2, BE=2,
19          attosend=(ttosend>=4256)?4256:ttosend =} ;
20      when(mainclock%320==0)
21        do{
22          ::choose_random
23            {= c=0, r=//Uniform(0,2^BE-1),//
24              backofftime=r*320 =} ;
.. ...

```

the backoff together with a CCA before the contention-access period ends (line 28). CCA detection time, denoted by `ccatime` in the code, is equal to 8 symbols period, namely $128 \mu\text{s}$ in 2.4 GHz band. If there is not enough time, the station stops and waits until the next beacon comes (line 26). A CCA is carried out immediately after the backoff finishes (line 29 & 31). There are three possible outcomes of the CCA: a busy channel, an idle channel with enough time to send the current portion of the message and an idle channel but not enough time. The outcome of a CCA is determined by the value of global variable `sending`. Variable `sending` indicates the number of stations transmitting at the current moment. A station about to transmit increases this variable and decreases it again once the transmission finishes.

A busy channel triggers another backoff if the number of backoff so far does not exceed the maximum allowed – `maxbackoff` – (line 33 & 35), otherwise the whole CSMA/CA procedure must be restarted to transmit the current portion of the message (line 37). When the channel is idle but there is not enough time to complete the transmission (line 51) the station escapes the CSMA/CA procedure and waits for the next beacon for further attempts.

In a beacon-enabled PAN, two CCAs are required after backoffs before the transmission of the message. Hence, when the channel is idle and there is enough time to complete the transmission (line 39), the transmission is only commenced if $CW = 0$ (line 42), namely when two consecutive CCAs find the channel idle. The transmission, which starts at the next backoff boundary, is announced to all other devices by increasing the global variable `sending`, and it takes the amount of time to send the portion of the message, namely `attosend`.

Once the portion of the message is transmitted, the station decreases the

```

MODEST model 4 The station (Part 2)
.. ...
25 alt{
26 ::when(mainclock>=CAP-backofftime-ccatime)
27   {= enoughtime=false =} ; break
28 ::when(mainclock<CAP-backofftime-ccatime)
29   when(c==backofftime) {= c=0 =} ;
30   do{
31     ::when(c==ccatime)
32       alt{
33         ::when(sending>0)
34           alt{
35             ::when(NB<=maxbackoff) channel_busy
36               {= CW=2, NB=NB+1, BE=(BE<6)?BE+1:6 =} ; break
37             ::when(NB>maxbackoff) {= restart=true =} ; break
38           }
39         ::when(sending==0 && mainclock<CAP-attosend-960)
40           count_down_CW {= CW=CW-1 =} ;
41           alt{
42             ::when(CW==0) wait_for_boundary ;
43             when (c==320) send_message_start
44               {= sending+=1, c=0 =} ;
45             when (c==attosend) send_message_end
46               {= ttosend-=attosend, sending-=1,
47                 restart=true =} ; break
48             ::when(CW>0) wait_for_boundary ;
49             when(c==320) {= c=0 =}
50           }
51         ::when(sending==0 && mainclock>=CAP-attosend-960)
52           {= enoughtime=false, restart=true =} ; break
53       }
54   } ; alt{
55     ::when(restart) {= restart=false =} ; break
56     ::when(!restart)
57     } } } } } }

```

global variable `send` and updates the remaining portions of the message to transmit by changing `ttosend` (line 46). Henceforth, the station restarts the CSMA/CA procedure to transmit the remaining message, without waiting for a new beacon.

Energy Consumption

In modelling the energy consumption of a station, we use the technical specification of CC2420 [3]. CC2420 is a 2.4 GHz ZigBee/IEEE 802.15.4-compliant radio frequency (RF) transceiver produced by Chipcon AS. Figure 3 summarizes the specification relevant for this paper. During its operation, the transceiver can be in four modes: *shutdown*, *idle*, *transmitting* or *receiving* modes. The rate of energy




Figure 3: Energy Modes of CC2420

consumption while occupying these modes [11] are 0.144, 712, 30672 and 35280 μW , respectively. A small amount of energy is required during the shutdown mode to power clock and to witness power-ups.

Furthermore, the transitions from mode *s* to *i*, from mode *i* to *t*, and from mode *i* to *r* do not happen instantaneously, but take 970, 194 and 194 μs to complete, respectively. During these transitions [11], the transceiver is considered to be still in the original mode, while consuming energy at the level of the destination mode. This reflects that the transceiver requires time (and power) to turn on its transmitting and receiving devices. All other transitions, namely those with dotted edges in the figure, take no time to complete. In the models, all stations are assumed to use transceivers with such energy consumption.

When the PAN coordinator broadcasts a beacon, all stations must be in receiving mode. That means, some time beforehand, they must have woken up and proceeded to the idle and then to the receiving mode. If a station is assigned GTSSs, it immediately changes to shutdown mode upon the completion of the beacon. However, the station must already be in transmitting mode, when its assigned GTSSs begin. To anticipate this, the station must leave the shutdown mode 1164 (namely $970+194$) μs before the GTSSs. After finishing the transmission in the GTSSs, the station moves to idle mode. If there is still ‘enough’ time before the next beacon, it changes further to shutdown mode.

A station without assigned GTSSs must compete with similar stations by using slotted or unslotted CSMA/CA to gain access to the medium. Immediately after the end of a beacon, the station enters the CSMA/CA procedure. Depending on

the duration of the backoff delay, the station may transition to idle or shutdown mode. If the backoff delay is long enough, the station may sleep and wake up 1164 μ s before it must perform a CCA, and be ready in receiving mode. If two consecutive CCAs result in idle channel, the station changes to transmitting mode and sends the message portion. This is performed continuously until the whole message is sent. Afterwards, the station may move to shutdown mode if there is enough time to do so before the next beacon.

To incorporate the energy consumption to the PAN models, each instant of station changes variables `time_in_s_mode`, `time_in_i_mode`, `time_in_t_mode` and `time_in_r_mode` every time it spends some portions of time in shutdown, idle, transmitting, and respectively receiving mode. The framework models are annotated accordingly and the four variables are turned into four global variable arrays, which each instant of station accesses through its `id`. The amount of energy consumed by a station during a beacon interval is the sum of amount of time it spends in each mode weighted by the rate of the energy consumption of the mode.

Referring back to MODEST model 3 and 4, the stations start being in receiving mode in lines 11, 29 and 49. They are starting in transmitting mode in line 44. They start in idle mode in lines 13, 31, and 37. In lines 16, 23, 27, 36 46 and 52, they have to make decisions whether to enter idle or shutdown mode based on the progress of their main clocks so far. If there is enough time to sleep and then follow the wake up procedure, which takes time, they enter shutdown mode. Otherwise, they just enter idle mode.

Clock Precision

Clock precision is important to ensure the correct functioning of ZigBee/IEEE 802.15.4-compliant PANs. In the case of PANs which provide CFPs, it is crucial that stations with assigned GTSs transmit exactly in their allocated period of time. Similarly with PANs employing the slotted CSMA/CA mechanism, the ability to correctly determine the backoff boundaries, which requires a precise clock, is necessary to avoid collisions.

Quartz-based clocks are generally used for electronic components in industry. Such clocks suffer from inaccuracy due to aging and temperature variations. Usually, the manufacturers guarantee a certain upper bound inaccuracy for their clock products. Assuming this upper bound to be constant is actually not realistic. However, we do so in our models. Hence, a clock with a guaranteed accuracy may deviate from the real time within a given time interval based on the accuracy and exhibit at different times a different deviation from the real time.

Clock inaccuracy is usually expressed in ‘parts per million’ (ppm), namely the maximum deviation is some certain time units within a million time units. In the models, we assume that a clock with a guaranteed accuracy of p ppm may show at time t , a clock value within the interval

$$[t - p \cdot 10^{-6} \cdot t, t + p \cdot 10^{-6} \cdot t].$$

A station, however, does not know the exact value of the deviation. Nevertheless, it must be able to precisely observe the deadlines, for instance the arrival of a beacon. Therefore, it counts on the maximum inaccuracy and be ready for the beacon even before the actual time. It is also assumed that the stations synchronize their clocks to the PAN coordinator's, while receiving the beacons.

We model the effect of clock inaccuracy to energy consumption in the following way. Assume that a station must wait for W time units to be ready for some event. The clock of the station has inaccuracy p ppm. The actual waiting time W' for the station is

$$W' = W - W \cdot p \cdot 10^{-6} + W \cdot p' \cdot 10^{-6}, \quad (1)$$

where $W \cdot p \cdot 10^{-6}$ is the maximum deviation of the clock given the clock inaccuracy and $W \cdot p' \cdot 10^{-6}$ is the actual deviation. The actual inaccuracy p' is a value in the interval $[-p, p]$. The models can be parameterized by the actual deviation. In full generality, a clock's inaccuracy may be time-dependent, in which case the above formula involves integration.

SIMULATIONS

In this section, we describe the simulations of the MODEST models presented in the previous section. The simulation was done with discrete-event simulator of MÖBIUS, and we only present an excerpt of several thousand simulation runs we performed. First, the experimental setups are described. The result of the experiments is presented afterwards, followed by its analysis and discussions.

Experiments

In all experiments, the system is a personal area network, which consists of a single PAN coordinator and 10 stations with star topology. The network uses 2.4 GHz band and all durations appearing in the standard which are defined by amount of symbols are adapted accordingly. The networks do not use any inactive periods, thus the beacon interval is always the same as the superframe duration (BI=SD).

The experiments cover $BO = 0, \dots, 10$, which means a beacon interval ranges from 15360 μs to 15728640 μs . Similarly the length of a slot ranges from 960 μs to 983040 μs . In all experiments, a station always has about a slot-length message to transmit during the duration between two beacons. A station transmitting in GTSS can transmit the whole message continuously. On the other hand, a station using slotted or unslotted CSMA/CA can only transmit 133 bytes at a time, which takes 4256 μs to transmit.

The simulation time of each experiment takes the duration of 10 beacon intervals. For instance, if $BO = 7$, a beacon interval takes 1966080 μs . Therefore, the simulation is run for around 20 seconds (19660800 μs). Simulations are carried out to estimate the mean values of some measures of interest, such as the time a station spends in shutdown mode. The simulation is repeated until the mean

values of all measures of interest converge with relative confidence interval 0.1 and confidence level 95%. MÖBIUS allows users to adjust these settings as desired.

The experiments are divided into two groups: (P) experiments for models with perfect clocks and (D) experiments for models with drifting clocks. There are 3 series of experiments in experiment group P. In series 1, 7 stations use GTSS, while 3 stations use slotted CSMA/CA. This is due to the restriction that there can only be at most 7 slots used for CFP in beacon-enabled networks. In series 2 and 3, 10 stations use slotted and unslotted CSMA/CA, respectively. In experiment group D there are 12 series of experiments, namely the three series in group P repeated for clock inaccuracies of 40, 20, 10 and 5 ppm.

As mentioned earlier, we consider worst-case scenarios whenever appropriate. For clock inaccuracy, the worst scenarios occur when the clock of the PAN coordinator is progressing as slow as possible within its inaccuracy bounds, while the clock of a station is in its fastest possible progress. Referring to Equation (1), the clock of a station is fastest when $p' = -p$. Then the station, which needs to wake up some time before the actual deadlines, actually wakes up even earlier because its clock is too fast. We proved this situation to be worst case by analytical means, which is backed up by simulation. The experiments reported in group D are worst scenario experiments.

All simulations are conducted on a PC with a Pentium 4 3.0 GHz processor with 1 GB RAM running Linux 2.6.17-2.686. The CPU time required ranges from the briefest around 5.2 seconds for experiment group P series 1 with $B0 = 0$ precise clock to the longest around 15429 seconds (around 4.25 hours) for experiment group D series 3 with $B0 = 10$ and clock inaccuracy 40 ppm.

Results

The graphs in Figure 4 summarize the result of experiment group P, where clocks are perfect. The first graph in the figure shows the percentage of the time spent in transmit, receive and idle modes for the 7 stations assigned with GTSS in experiment series 1 (we denote this as series 1G). While the portion of time spent in transmit mode is constant for all $B0$'s, the portions in receive and idle modes diminish as the beacon interval gets larger, and become less than 1% after $B0 = 5$. This is because the actual time spent in receive and idle modes remains constant for all beacon intervals, thus their percentages decrease as the beacon intervals get longer. They are constant because a station only goes to receive mode while receiving beacons and to idle mode prior to the arrivals of beacons and before the transmissions of the messages.

The second graph in the figure depicts the same percentage distribution for the 3 stations using slotted CSMA/CA in experiment series 1 (denoted as series 1C). The portion of time spent in the transmit mode remains almost constant and around 0.25% higher throughout than the previous case. The increase can be explained by the fact that there are more transitions from idle to transmit mode compared to the previous case, because of the restriction on the length of message

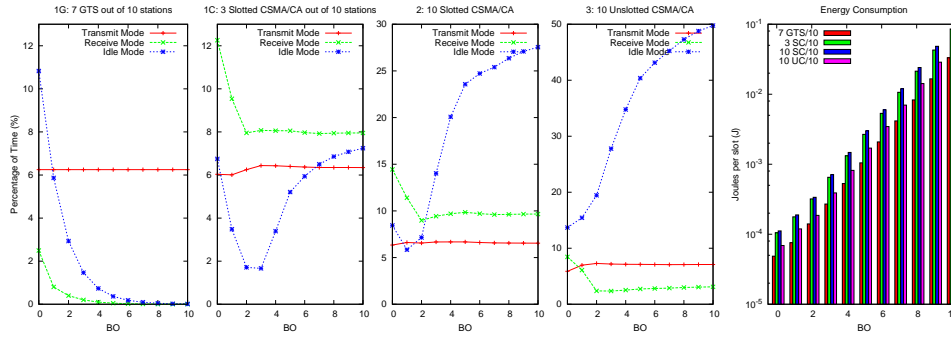


Figure 4: P: Percentage of time per energy mode and energy consumption for perfect clocks




Figure 5: D: Percentage of time per energy mode and energy consumptions for drifting clocks (40 ppm)

portions. When the beacon intervals are too brief, some of the stations may find not enough time to completely transmit its apportioned message due to the overhead of the CSMA/CA procedure, hence the small perturbations in the percentage of transmit mode when $BO < 5$. The percentage of time spent in receive mode is also almost constant at around 8%, which is higher than in series 1G. The percentage increase in receive mode is due to the CCAs, and since the number of CCAs to be performed is proportional to the length of the message, the percentage is steady for most BO values. The time spent in the idle mode, on the other hand, increases with the length of beacon intervals. The longer the beacon interval, the longer a slot-length message. However, since the length of message portion transmittable is restricted, the more often slotted CSMA/CA procedure is repeated to send the whole message. Hence the station spends more percentage of time in idle mode.

The third and fourth graphs in the figure show similar percentage distribution for experiment series 2 and 3, namely 10 stations using slotted and unslotted CSMA/CA, respectively. Compared to series 1C, the percentage of time spent in

receive mode is higher in series 2. This is due to the fact that there are 10 stations for 15 CAP slots in series 2 compared to 3 stations for 8 CAP slots in series 1C. Thus stations perform more CCAs and spend more percentage of time in receive mode. Similar reason applies to the percentage of idle mode. In series 3, however, the percentage of time spent in receive mode is lower than in series 1C and 2 because the stations perform only one CCA in every CSMA/CA cycle. Between themselves, the stations in experiment series 3 spend more time in idle mode compared to those in series 2. This can be explained by the absence of beacons in unslotted CSMA/CA. Once a station finishes transmitting the whole message, it does not have any estimation when the next message will be available and thus must stay in idle mode. The station only enters the shutdown mode when it is performing backoff.

The rightmost graph in Figure 4 depicts the amount of energy consumed by a station during a beacon interval time for each value of B_0 and for each of the formerly described series. Note that the y-axis of the graph is in logarithmic scale. The largest energy consumption is around 0.096946940587 Joules per beacon interval when $B_0 = 10$. We can observe that for all values of B_0 , a station using GTSS consumes the least amount of energy. A station using the unslotted CSMA/CA consumes more energy than that using GTSS but less than the energy consumption of a station using slotted CSMA/CA. A question arises as to why the energy consumption of a station in series 1C is lower than that in series 2, considering the fact that the former has more CAP-slots per station. The reason is in series 1C, a station spends more time in shutdown mode, namely during the CFP period. This can also be observed from the second graph in the figure. It is also worth noting that the differences between the energy consumptions amongst the series are significant, even though they are not apparent due to the logarithmic scale. For instance, a station in series 2 consumes almost thrice the energy of a station in series 1G and almost twice of a station in series 3.

Figure 5 depicts the corresponding graphs for the experiment group D, in which clocks are now drifting with inaccuracy 40 ppm. Recall that we consider the worst case of the coordinator driven by a slowest possible clock, and the stations being fastest. The graphs show similar trends as the graphs in Figure 4, but with some differences. These differences are not apparent, due to the scale of the graphs. There is actually a slight increase of percentage of time spent in receiving and idle modes, decrease of percentage of time spent in shutdown mode and no changes in transmitting mode, which will be emphasized in the next figure.

Figure 6 shows the influence of worst-case clock drifts on the energy consumptions of stations in series 1G, 2 and 3, respectively. Consistently, clock drifts increase the amount of energy consumed for all experiments. The graphs in the figure show the increase in energy expended relative to the amount of energy needed in the precise clock setting. The first graph depicts the deviations for a station in series 1G with clock inaccuracies 40, 20, 10 and 5 ppm for $B_0 = 0, 2, 4, 6, 8, 10$. We observe a linear correlation between the clock inaccuracies and the deviations in the graph. For instance, for most of B_0 values, a station with clock inaccuracy

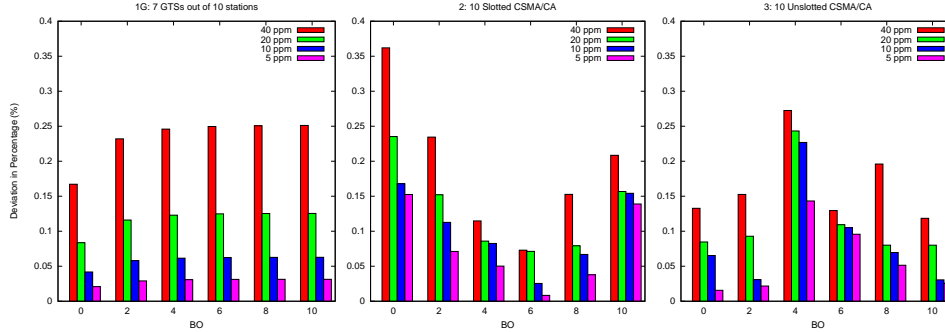


Figure 6: Increase of energy consumptions when clocks drift (40, 20, 10 and 5 ppm) relative to perfect clocks

40 and 20 ppm spends approximately 0.25% and respectively 0.125% more energy than a station with precise clock. The second and third graphs in Figure 6 depicts the deviations for a station in series 2 and 3, respectively. We abstain from explaining the precise patterns exhibited by these plots for increasing BO. However, the effects of the level of clock inaccuracies are obvious. In both graphs, higher clock inaccuracy tends to cause higher energy consumption level.

Overall, a deviation of 0.25% or 0.125% may seem small. If extrapolated, it amounts to about half a day per year decrease in lifetime of a battery. However, a clock drift of 40 ppm is – percentwise – only 0.004%, and we thus conclude that such a clock drift may be amplified by a factor of 62.5 when looking at the excess in energy consumption caused.

CONCLUSIONS

This paper has advocated a rigorous approach to modelling and simulation of energy-aware networked embedded systems. In this approach, MODEST, a language with a strictly formal semantics is used to specify the system under study. Since all assumptions are exposed explicitly in the specification, the underlying stochastic model is well-defined and the obtained simulation results are trustworthy. This is different from many other simulation-based results lately published in the networking literature. The simulation is carried out via the tools MOTOR and MÖBIUS. MOTOR is available as source code from <http://www.purl.org/net/motor> under the GPL license. MÖBIUS is freely available for educational and research purposes from <http://www.mobius.uiuc.edu/>. MOTOR can be installed as an add-on package to the MÖBIUS installation.

We applied the proposed approach to the modelling of crucial medium access mechanisms in IEEE 802.15.4, with a particular emphasis on energy specific configuration relevant to ZigBee operation. When investigating the effects of time-

slotted (*i.e.*, beacon-enabled) and unslotted medium access techniques, we observed that slotted CSMA/CA is clearly inferior to unslotted CSMA/CA w.r.t. energy efficiency. In slotted mode, devices using GTSs use considerably less energy than those using CSMA/CA.

Our studies of drifting clocks reveal some interesting observations. Most of them are consistent with what one might expect – but we did not see them published elsewhere. However interesting, they only show minor quantitative effects. The IEEE standard limits the allowed clock drift to 40 ppm, and this appears to be caused by correctness concerns. (One can show that with 64 ppm the clock boundaries could drift apart to an extent that the protocol could desynchronize and thus malfunction). In the worst case we estimate that a clock drift of 40 ppm induces a shortage in battery life of about half a day per year. While this amplifies the clock drift by a factor of about 60, we think it is not particularly significant, leading us to finally negate the question in the title of this paper.

Acknowledgement. This work is supported by the German Research Council (DFG) as part of the Transregional Collaborative Research Center 'Automatic Verification and Analysis of Complex Systems' (SFB/TR 14 AVACS). See <http://www.avacs.org> for more information.

REFERENCES

- [1] *IEEE 802.15.4: Wireless Medium Access Control (MAC) and Physical Layer (PHY) Specifications for Low-Rate Wireless Personal Area Networks*. IEEE Press, New York, 2003.
- [2] *ZigBee Specification Version 1.0*. ZigBee Alliance, 2004.
http://www.zigbee.org/en/spec_download/download_request.asp.
- [3] *CC2420 Product Information*. Chipcon AS, 2005.
http://www.chipcon.com/index.cfm?kat_id=2&subkat_id=12&dok_id=115.
- [4] *MoDeST Tutorial: development of a complex model in MoDeST*. Dependable Systems and Software, Saarland University, Germany, 2006.
<http://depend.cs.uni-sb.de/modesttutorial/index.html>.
- [5] The network simulator – ns-2 website, 2007.
<http://www.isi.edu/nsnam/ns/>.
- [6] T. R. Andel and A. Yasinac. On the credibility of Manet simulations. *IEEE Computer*, 39(7):48–54, 2006.
- [7] J. Bengtsson, K. G. Larsen, F. Larsson, P. Pettersson, and W. Yi. UPPAAL – a Tool Suite for Automatic Verification of Real-Time Systems. In *Workshop on Verification and Control of Hybrid Systems III*, pages 232–243. Springer, 1995.
- [8] H. C. Bohnenkamp, P. R. D’Argenio, H. Hermanns, and J.-P. Katoen. MoDeST: A compositional modeling formalism for hard and softly timed systems. *IEEE Transactions on Software Engineering*, 32(10):812–830, 2006.
- [9] H. C. Bohnenkamp, J. Gortler, J. Guidi, and J.-P. Katoen. Are you still there? - A lightweight algorithm to monitor node presence in self-configuring networks. In *2005 International Conference on Dependable Systems and Networks, DSN’05*, pages 704–709. IEEE Computer Society, 2005.

- [10] H. C. Bohnenkamp, H. Hermanns, R. Klaren, A. Mader, and Y. S. Usenko. Synthesis and stochastic assessment of schedules for lacquer production. In *1st International Conference on Quantitative Evaluation of Systems, QEST'04*, pages 28–37. IEEE Computer Society, 2004.
- [11] B. Bougard, F. Catthoor, D. C. Daly, A. Chandrakasan, and W. Dehaene. Energy efficiency of the IEEE 802.15.4 standard in dense wireless microsensor networks: Modeling and improvement perspectives. In *8th Conference on Design, Automation and Test in Europe, DATE'05*, pages 196–201. IEEE Computer Society, 2005.
- [12] M. Cadilhac, T. Héroult, R. Lassaigne, S. Peyronnet, and S. Tixeuil. Evaluating complex mac protocols for sensor networks with apmc. *Electron. Notes Theor. Comput. Sci.*, 185:33–46, 2007.
- [13] D. Cavin, Y. Sasson, and A. Schiper. On the accuracy of MANET simulators. In *2nd ACM International Workshop on Principles of Mobile Computing*, pages 38–43. ACM Press, 2002.
- [14] D. Daly, D. D. Deavours, J. M. Doyle, P. G. Webster, and W. H. Sanders. Möbius: An extensible tool for performance and dependability modeling. In *11th International Conference on Computer Performance Evaluation: Modelling Techniques and Tools, TOOLS'00*, pages 332–336. Springer, 2000.
- [15] P. R. D'Argenio and J.-P. Katoen. A theory of stochastic systems part I: Stochastic Automata. *Information and Computation*, 203(1):1–38, 2005.
- [16] D. D. Deavours and W. H. Sanders. An efficient well-specified check. In *8th International Workshop on Petri Nets and Performance Models*, pages 124–133. IEEE Computer Society, 1999.
- [17] M. Fruth. Probabilistic model checking of contention resolution in the IEEE 802.15.4 low-rate wireless personal area network protocol. In *2nd International Symposium on Leveraging Applications of Formal Methods, Verification and Validation, ISOLA'06*, 2006.
- [18] H. Garavel, F. Lang, and R. Mateescu. An overview of CADP 2001. *European Association for Software Science and Technology Newsletter*, 4:13–24, 2001.
- [19] H. Hermanns, D. N. Jansen, and Y. S. Usenko. From StoCharts to MoDeST: a comparative reliability analysis of train radio communications. In *5th International Workshop on Software and Performance, WOSP'05*, pages 13–23. ACM Press, 2005.
- [20] M. Z. Kwiatkowska, G. Norman, and D. Parker. PRISM: Probabilistic symbolic model checker. In *12th International Conference on Computer Performance Evaluation: Modelling Techniques and Tools, TOOLS'02*, pages 200–204. Springer, 2002.

- [21] O. Landsiedel, K. Wehrle, and S. Gotz. Accurate prediction of power consumption in sensor networks. In *2nd IEEE workshop on Embedded Networked Sensors, EmNets'05*, pages 37–44. IEEE Computer Society, 2005.
- [22] G. Pongor. OMNeT: Objective modular network testbed. In *International Workshop on Modeling, Analysis, and Simulation On Computer and Telecommunication Systems*, pages 323–326. Society for Computer Simulation, 1993.
- [23] B. L. Titzer, D. K. Lee, and J. Palsberg. Avrora: scalable sensor network simulation with precise timing. In *4th International Symposium on Information Processing in Sensor Networks, IPSN'05*, page 67. IEEE Press, 2005.
- [24] X. Zeng, R. Bagrodia, and M. Gerla. GloMoSim: A library for parallel simulation of large-scale wireless networks. In *12th Workshop on Parallel and Distributed Simulation*, pages 154–161, 1998.

CHRISTIAN GROSS: comlet Verteilte Systeme GmbH, Amerikastraße 21,
D-66482, Zweibrücken, Germany.
E-mail: christian.gross@comlet.de

HOLGER HERMANN: Department of Computer Science, Saarland University
D-66123, Saarbrücken, Germany
E-mail: hermanns@cs.uni-sb.de

REZA PULUNGAN: Department of Computer Science, Saarland University
D-66123, Saarbrücken, Germany
E-mail: pulungan@cs.uni-sb.de