



Exploring IoT Trickle-Based Dissemination Using Timed Model-Checking and Symbolic Execution

Boutheina Bannour¹(✉), Arnault Lapitre¹, and Pascale Le Gall²

¹ CEA LIST, Gif-sur-Yvette, France

{[boutheina.bannour](mailto:boutheina.bannour@cea.fr),[arnault.lapitre](mailto:arnault.lapitre@cea.fr)}@cea.fr

² MICS Laboratory, University of Paris-Saclay, Gif-sur-Yvette, France
pascale.legall@centralesupelec.fr

Abstract. We focus on studying an IoT algorithm called Trickle using a formal model-based approach. The algorithm has an essential role in traffic regulation across distributed networks of wireless sensors which are part of IoT. The algorithm allows efficient dissemination of information such as critical applicative data, firmware upgrades or security fixes. In this paper, we develop timed asynchronous computational models for Trickle. We show how reachability properties can be assessed on such models using a combination of model-checking and symbolic execution implemented by the tools UPPAAL and DIVERSITY, respectively. Our experiments produce promising results on highlighting updated or outdated nodes situations during dissemination.

1 Introduction

Context. Sensors networks (WSN) play an essential role in the uptake of the Internet of Things (IoT) as they allow direct connection between the physical environment and the digital systems. They come with a reduced economical cost, and they can easily be deployed in inaccessible areas. WSN involve constrained-energy devices (sensors) which operate over long periods. The information dissemination across these networks is often subject to constraints to reduce the communication cost, with the objective not to exhaust the batteries of such nodes that are in general neither rechargeable nor replaceable after the deployment. Gossip paradigm has been recognized as being efficient in practice to control the communications of each node, roughly speaking: (i) every node try quickly to transmit new data, and (ii) in case of redundant data reception, the node reduces the transmissions frequency over time. The algorithm Trickle [23, 25, 27] is one of the most known: It comes as a standard library in TinyOS [24] and Contiki [16], two well-known firmware Operating Systems (OS) for WSN. The algorithm is involved in recently standardized WSN protocols namely the Multicast Protocol for Low Power and Lossy Networks (MPL) [17] and the IPv6 Routing Protocol for Low Power and Lossy Networks (RPL) [1]. There are many others like FireFly Gossip (FiGo) [8], Energy Efficient Gossiping (E-Gossip) [22], Multi Randomized Gossip-Consensus-based Sync (Multi RGCS) [29], and new ones continue to be proposed given the economic interest of WSN.

Case Study: Trickle Dissemination. The goal of Trickle is to reach a stable global state of the network where all the nodes have the same up-to-date information. Each node applies a set of rules to control its transmissions as follows [27]:

- each node maintains a current interval I , a counter c and a broadcasting time t in interval $[I/2, I[$,
- global parameters to all nodes are k the redundancy constant, I_{min} (resp. I_{max}) the smallest (resp. largest) interval,
- each node applies the following rules:
 1. at the start of a new interval, the timer and counter c are reset and t is randomly set to a value in $[I/2, I[$,
 2. if a received message is consistent with the information the node holds, the counter c is incremented,
 3. when the timer reaches t and $c < k$, a message carrying the node information is transmitted to neighbours in broadcast,
 4. when the timer expires at I , the interval length is increased by setting I to $\min(2 \cdot I, I_{max})$ and a new interval starts,
 5. when a received message is inconsistent with the node information, then I is set to I_{min} , and a new interval starts. Otherwise, nothing happens.

Trickle uses “polite gossip” to exchange information with its network neighbours. It breaks time into dynamically adjustable intervals, and at a random point in each interval, it considers broadcasting the information it holds. If Trickle has already heard several other nodes gossip the same information in this interval, it politely stays quiet: “repeating what someone else has said is rude” [25].

Motivation and Related Work. As it may show from Trickle, popular gossip protocols are sophisticated and complex on the nature of applied control on node transmissions. For that reason, most of the validation effort of this class of protocols relies mainly on testbeds or simulations, and at a less extend on analytical methods, more exhaustive, yet they lack automation: among the latter, we can cite those developed for Trickle in [7, 12, 19, 28, 31]. On the other hand, formal methods and in particular model-checking [11] come with a high degree of automation. They consider computational models which, by its nature, delimit the perimeter of the analysis and the kind of properties to be verified. The survey [9] overviews many relevant works on applying formal methods on WSN protocols, including gossip-based protocols. The following papers [15, 20, 32–34]) have successfully applied model-checking on gossip protocols, among which [15, 33, 34] concern Trickle. The early work [15] proposed formal models for WSN based on classic process algebras. The work developed a simplified model of Trickle for illustration purposes, then translated into a network of Timed Automata (TA) [2] supported by the model-checker UPPAAL. In their model, Trickle intervals are not adjustable and are of fixed length, which restricts the coverage of performed analyses. More recently, authors in [33, 34] have proposed model-checking techniques for WSN, and Trickle has been proposed for illustration as well. The work [34] provides formal semantics for a subset of NesC programs used to built TinyOS [24] applications; such semantics have been implemented in the model-checking framework PAT [26] which supports TA too. The work [33] focuses

on combining probabilistic model checking provided by the tool PRISM [21] with automated debugging algorithms in order to find pathological typologies which cause some failure: typically in case of Trickle, it is about finding the topology pattern that prevents recent information from being spread. In both works [33,34], Trickle models have not been given. In this paper, we are interested in the distributed nature of WSN nodes [31], that is they are dephased by some duration because often they do not share a common zero as classically in distributed systems. Gossip protocols, Trickle included, do not have any assumption on nodes that should be synchronized. The control they provide on transmission instants is usually implemented by introducing some time randomness to benefit from desynchronized nodes. Besides, inline with such assumption, we consider asynchronous communications which can take time to deliver data, all previous works have not taken into account such desynchronized hypothesis.

Contribution and Paper Outline. We provide formal models for Trickle as Extended Timed Automata (XTA) in which data can be used to define computations, random updates, constraints on clocks and communication actions. A network of XTA can then be designed to form the overall Trickle topology of nodes. The network is endowed with operational semantics in which XTA communicate their data via unbounded queues. In this paper, we are interested in highlighting some situations of updated or outdated nodes using reachability properties under the desynchronized assumption. To assess such properties on XTA, which are more expressive than classical TA because they introduce data, we propose to combine model-checking [11] with symbolic execution [18]. The latter virtually executes models (or code) for symbolic input parameters rather than concrete values. Each execution path is associated with logical constraints on those input parameters computed at each execution step, the so-called Path Conditions (PC). PCs are a compact representation of classes of actual values for input parameters for executing those paths; they can be solved using SMT solvers such as CVC4 [6] or Z3 [13]. Symbolic execution can be applied on timed models (e.g., [4]).

Yet enumerating all execution paths to check a property of interest is combinatorial, this is an identified problem of the technique that we have experimented as well on some early work on Trickle [5,30]. The idea is then to apply model-checking first on XTA in which data is numeric and random updates are restricted to some values, yet clock constraints are handled as usual by zone-based abstraction implemented in UPPAAL using the efficient Difference Bounded Matrices data structure (DBM) [14]. In case the property is verified, the corresponding sequence of transitions of a solution is used to guide symbolic execution. We experiment with the combined techniques in UPPAAL and in the symbolic execution tool DIVERSITY [3]. The rest of the paper is organized as follows. Section 2 introduces the formal model of the network of XTA and its operational semantics. Section 3 proposes a Trickle model designed as a network of XTA. Section 4 introduces our approach of combining model-checking and symbolic execution to assess reachability properties and evaluates it on UPPAAL-DIVERSITY connection. Section 5 concludes the paper.

2 Network of Extended Timed Automata

The model of Timed automata [2] is a well-established formalism for modelling the timing behaviour of systems. This section defines syntax and semantics of an extension of timed automata introducing data updates, communication actions with data transmission, and non-trivial data-dependent time constraints. Those appear in functional specification of systems, as in the specification of Trickle in the introduction.

Data Domain. We use a universal data domain D to abstract all values of time variables, called clocks as usual, and other data variables. Data variables can be of any type, whereas clocks are typed in a time domain $T \subseteq D$ which is isomorphic to \mathbb{Q}_+ , the set of positive rational numbers.

Data Valuations. For a set of data variables V , a data valuation is a type-preserving mapping $v : V \rightarrow D$. We canonically extend data valuation to usual arithmetical expressions defined over V , i.e., $v(e)$ is the value of e for the valuation v . We denote by D^V the set of all such valuations.

Data Formulae. The set $\mathcal{F}(V)$ of data formula f over V is either: an atomic formula of the form *true*, *false*, $e_1 = e_2$, $e_1 < e_2$, and $e_1 \succ e_2$ with $< \in \{<, \leq\}$ and $\succ \in \{>, \geq\}$; or built over those using usual connectives: conjunction (\wedge), disjunction (\vee), and negation (\neg).

Sequential Data Updates. We consider sequential updates defined as follows:

$$u ::= \text{skip} \mid x := e_1 \mid e_1 <: x <: e_2 \mid u_1; u_2 \mid \text{if}(f)\text{then}\{u_1\}\text{else}\{u_2\} \mid \text{repeat}(n)\{u_1\}$$

skip is the null update; $x := e_1$ assigns the variable x with a new value denoted by e_1 ; $e_1 <: x <: e_2$ assigns x with a new random value bounded from below by the value denoted by e_1 , the value of x is also bounded from above by the value denoted by e_2 . Moreover, updates can be built using usual control primitives: sequence ($;$), condition (*if ... then ... else ...*) or counted-loop (*repeat(n)*) allowing the repetition of the enclosed update n times.

The set of update functions $\mathcal{U}(V)$ is defined by functions $\llbracket u \rrbracket$ from D^V to $2^{(D^V)}$. The set of valuations $\llbracket u \rrbracket(v) \in 2^{(D^V)}$ is defined on the form of u as follows¹:

$$\llbracket \text{skip} \rrbracket(v) = \{v\} \tag{1}$$

$$\llbracket x := e_1 \rrbracket(v) = \{v[x \rightarrow v(e_1)]\} \tag{2}$$

$$\llbracket e_1 <: x <: e_2 \rrbracket(v) = \{v[x \rightarrow y] \mid v(e_1) < y < v(e_2)\} \tag{3}$$

$$\llbracket u_1; u_2 \rrbracket(v) = \{v' \mid \exists v'' \in u_1(v), v' \in u_2(v'')\} \tag{4}$$

$$\llbracket \text{if}(f)\text{then}\{u_1\}\text{else}\{u_2\} \rrbracket(v) = u_1(v) \quad v(f) \tag{5}$$

$$\llbracket \text{if}(f)\text{then}\{u_1\}\text{else}\{u_2\} \rrbracket(v) = u_2(v) \quad \neg v(f) \tag{6}$$

$$\llbracket \text{repeat}(n)\{u_1\} \rrbracket(v) = \llbracket u_1; \text{repeat}(n-1)\{u_1\} \rrbracket(v) \quad n > 0 \tag{7}$$

$$\llbracket \text{repeat}(0)\{u_1\} \rrbracket(v) = \{v\} \tag{8}$$

¹ Given a function $h : A \rightarrow B$, a subset $X \subset A$, the function $h' = h[x \rightarrow y, x \in X]$ is defined as follows: $h'(z) = y$ if $z \in X$ otherwise $h'(z) = h(z)$. In case X is a singleton of the form $\{x\}$, we denote $h' = h[x \rightarrow y]$ in short.

Clock Formulae. Given a set of clocks Cl disjointed from V ($Cl \cap V = \emptyset$), a clock valuation is a mapping $w : Cl \rightarrow T$. With previous notation, the set $\mathcal{G}(Cl, V)$ of clock formulas g over Cl and V is either an atomic formulas of the form *true*, *false*, $clk \prec e$ or $clk \succ e$ where e is an expression over V typed in time domain T ; or a conjunction of those. The set of clock invariants $\mathcal{I}(Cl)$ is defined by conjunctions of formulas of the form $clk \prec e$. We define a universal valuation $v \oplus w : V \cup Cl \rightarrow D$ as the resulting valuation which coincide with v and w on V and Cl respectively. This valuation can be canonically extended to formulas as usual.

Communication Actions with Data. Given a set of interaction points, often called ports, P , the set of communication actions $\mathcal{C}(P, V)$ contains two kind elements: output actions of the form $p!e$ which denotes an emission on some port p of a piece of data corresponding to the current valuation of e ; or input actions of the form $p?x$ which denotes a reception of a piece of data that is stored in the variable x . Moreover, we consider the special action ϵ which denotes the absence of communication action. The valuations of actions of the form $p!e$, $p?x$, and ϵ , are defined by $v(p!e) = p!v(e)$, $v(p?x) = p?v(x)$, and $v(\epsilon) = \epsilon$ respectively.

Extended Timed Automaton. An extended timed automaton (XTA in short) is a tuple $(L, l_0, V, Cl, P, Tr, Inv)$ where L is a finite set of locations, $l_0 \in L$ is the initial location, V is a set of variables, Cl is a set of clocks, P is a set of ports, $Tr \subseteq L \times \mathcal{F}(V) \times \mathcal{G}(Cl, V) \times (\mathcal{C}(P, V) \cup \{\epsilon\}) \times \mathcal{U}(V) \times 2^{Cl} \times L$ is a set of transitions, and $Inv : L \rightarrow \mathcal{I}(Cl)$ is a state invariant mapping.

For a transition $tr = (l, f, g, ca, u, R, l') \in Tr$, l and l' are respectively the source and target location of tr ; f and g are respectively the data guard and time guard of tr , i.e., enabling conditions on data variables and clocks; ca is the communication action of tr ; u is an update function through which data variables are updated when tr is executed; and $R \subseteq Cl$ is the set of clocks to be reset.

The semantics of an XTA is a labeled transition system where states s are triples (l, v, w) where l is a location, v and w are data and clock valuations respectively. The transition relation is defined as follows:

- delay transition: $(l, v, w) \xrightarrow{d} (l, v, w')$ where for all $clk \in Cl$, $w'(clk) = w(clk) + d$ with some $d \in T$ such that $v \oplus w' \models Inv(l)$
- action transition: $(l, v, w) \xrightarrow{a} (l', v', w')$ if and only if there exists a transition $(l, f, g, ca, u, R, l') \in Tr$:
 - $v \oplus w \models f \wedge g$,
 - $a = v(ca)$,
 - $v' \in \llbracket u \rrbracket(v)$,
 - $w' = w[clk \rightarrow 0, clk \in R]$.

In the following, we introduce a network of XTA which exchange data using broadcast communication.

A Network of Extended Timed Automata. A network of XTA denoted by $\mathcal{N} = ((\mathcal{A}_i)_{i \in \{1, \dots, n\}}, \mathcal{K})$ is defined as follows:

- $(\mathcal{A}_i)_{i \in \{1, \dots, n\}}$ a family of XTA $\mathcal{A}_i = (L_i, l_0^i, V_i, Cl_i, P_i, Tr_i, Inv_i)$ which do not share variables and clocks, i.e., for all $i, j \leq n$ we have that $V_i \cap V_j = \emptyset$, $Cl_i \cap Cl_j = \emptyset$, and $P_i \cap P_j = \emptyset$,
- a total function $\mathcal{K} : P_{\mathcal{N}} \rightarrow 2^{P_{\mathcal{N}}}$ specifying connections between ports.

The set of all ports of \mathcal{N} is denoted by $P_{\mathcal{N}} = \bigcup_{i \leq n} P_i$. Besides, the set of all variables (resp. clock variables) of \mathcal{N} is denoted by $V_{\mathcal{N}} = \bigcup_{i \leq n} V_i$ (resp. $Cl_{\mathcal{N}} = \bigcup_{i \leq n} Cl_i$).

We make the hypothesis that latent data issued (potentially by different sources) and targeting some internal port in delivered on that port in the order they were sent, i.e., we will implement this using queues with a policy of first-in-first-out (fifo). In the following, we denote by the function $q : P_{\mathcal{N}} \rightarrow D^*$ the pending data in the network as being the content of queues associated with receiving ports.

The semantics of a network of XTA is a labeled transition system in which: states are tuples of the form $S = ((l_1, v_1, w_1), \dots, (l_n, v_n, w_n), q)$ with initial states S_0 verify for all $i \leq n$, $l_i = l_0^i$; and transitions are defined as follows:

- delay transition:

$$((l_1, v_1, w_1), \dots, (l_n, v_n, w_n), q) \xrightarrow{d} ((l_1, v_1, w'_1), \dots, (l_n, v_n, w'_n), q)$$

iff for all $i \leq n$ there exists $(l_i, v_i, w_i) \xrightarrow{d} (l_i, v_i, w'_i)$

- internal output transition:

$$((l_1, v_1, w_1), \dots, (l_i, v_i, w_i), \dots, (l_n, v_n, w_n), q) \xrightarrow{a} ((l_1, v_1, w_1), \dots, (l'_i, v'_i, w'_i), \dots, (l_n, v_n, w_n), q')$$

iff there exists $(l_i, v_i, w_i) \xrightarrow{a} (l'_i, v'_i, w'_i)$ with $a = p!m$, q' is such that for all port p_1 either $p_1 \in \mathcal{K}(p)$ then $q'(p_1) = q(p_1).m$ otherwise $q'(p_1) = q(p_1)$,

- internal input transition:

$$((l_1, v_1, w_1), \dots, (l_i, v_i, w_i), \dots, (l_n, v_n, w_n), q) \xrightarrow{a} ((l_1, v_1, w_1), \dots, (l'_i, v'_i, w'_i), \dots, (l_n, v_n, w_n), q')$$

iff there exists $(l_i, v_i, w_i) \xrightarrow{a} (l'_i, v'_i, w'_i)$ with $a = p?m$, $q(p)$ is not empty and is of the form $q(p) = m.q_1$, and q' is such that $q'(p) = q_1$ and for all $p_1 \neq p$ we have $q'(p_1) = q(p_1)$,

- silent or external action transition:

$$((l_1, v_1, w_1), \dots, (l_i, v_i, w_i), \dots, (l_n, v_n, w_n), q) \xrightarrow{a} ((l_1, v_1, w_1), \dots, (l'_i, v'_i, w'_i), \dots, (l_n, v_n, w_n), q)$$

iff there exists $(l_i, v_i, w_i) \xrightarrow{a} (l'_i, v'_i, w'_i)$ with $a = \epsilon$ or $a = p!m$ (resp. $a = p?m$) such that $\mathcal{K}(p) = \emptyset$ (resp. for all p_1 , $p \notin \mathcal{K}(p_1)$).

In a nutshell, the above definition shows that time advances in the same way for all clocks of XTA forming the network. Besides, an internal emission $p!m$ on a port p has the effect of filling all the fifo associated to the ports of $\mathcal{K}(p)$ and an internal reception $p?m$ on port p consumes the first message stored in its fifo.

When a silent action or an external action (reception or an emission) occurs on a port which is not connected to other XTA ports, it is executed with no effect on fifo queues, since it is assumed to be connected to some implicit environment.

A run of the network is derived from the labelled transition system as a path starting in an S_0 and alternating delay transitions and action transitions. The property we are interested in is the reachability of states S . A state S is reachable iff there exists a run in which S occurs. In practice, such states are those which satisfy some user-specified formula $\phi = f \wedge g$ on data and clocks.

Let us denote by $\text{tr-seq}(r)$ the sequence of (syntactic) transitions in $\bigcup_{i \leq n} Tr_i$ covered by a run r . We recall that such sequence intertwines transitions of different automata composing the network based on induced fifo-communications causalities discussed above.

Two runs r_1 and r_2 are said to be *coverage equivalent* if and only if $\text{tr-seq}(r_1) = \text{tr-seq}(r_2)$, i.e., they cover the same (syntactic) transitions sequence.

The equivalence classes characterized by this relation guide the symbolic execution to search for all runs of a given class (as a symbolic path together with its path condition). On the other hand, model-checking will be used to compute some representative runs of the class that satisfy a reachability property ϕ . In practice, from the latter, we extract the transition sequence that guides the exploration performed by the symbolic execution.

3 Trickle Models

Trickle Node Behavior. We propose the XTA $(L, Init, V, Cl, P, Tr, Inv)$ which specifies Trickle behaviour of each node in the network, the automaton is depicted in Fig. 1. The XTA has 5 locations.

$L = \{Init, Listen_1, Listen_2, Check_1, Check_2\}$ in which $Init$ is the initial location. The clocks set is a singleton $Cl = \{clk\}$ containing one clock used to implement the Trickle timer. The set of variables contains 5 variables $V = \{I, t, c, myv, rcv\}$: the former three variables I , t and c are Trickle variables which respectively represent the value of the current interval, the instant of transmission and the counter value (whereas I_{min} , I_{max} , k are Trickle constants). Without loss of generality, in this automaton, Trickle is used to maintain consistency of version number across the network, the variable myv stores the most recent version the node holds and rcv is used to store received versions from neighbors.

The automaton has 10 transitions composing the set Tr , that we will overview next together with meaning associated with state invariants defined by mapping Inv . A node can be started at any time; this is captured by state invariant $Inv(Init) : true$ in location l_0 , which means that any duration can elapse in this location. The transition $Init \rightarrow Listen_1$ is fired to start the Trickle behaviour. It sets I to I_{min} , assigns counter c and clk with 0, and finally chooses a the transmission time $I/2 \leq t < I$ within the second-half of the interval current interval I . The state invariant $Inv(Listen_1) : clk \leq t$ constrain time elapsing to be bounded by t . When clk reaches t , there two possible behaviors: either the transmission

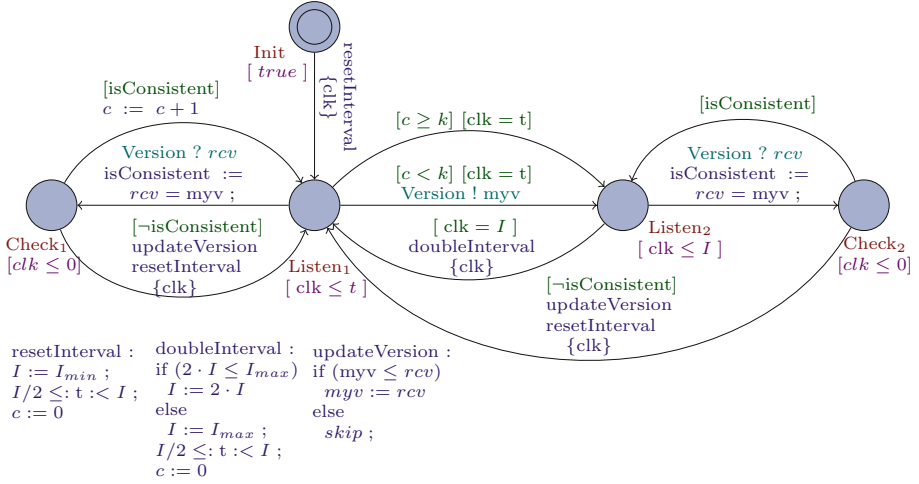


Fig. 1. An extended timed automaton (XTA) of a Trickle node behavior

occurs given $c < k$ is fulfilled (horizontal transition $Listen_1 \rightarrow Listen_2$ with action $Version?rcv$), otherwise the transmission is suppressed (curved transition $Listen_1 \rightarrow Listen_2$ with action ϵ). A first reception handling is defined by transition $Listen_1 \rightarrow Check_1$. Once started, the automaton satisfies the input enable less property: in every state ($Listen_1$ or $Listen_2$), it is possible to receive every input (action $Version?rcv$). Each time, a version is received, it is compared to the current version of the node: in case of consistency (same version $rcv = myv$), the counter c is incremented ($c := c + 1$), we recall that the latter counts redundant versions; otherwise (case of inconsistency) a new interval is started, and the node updated its version if it is older.

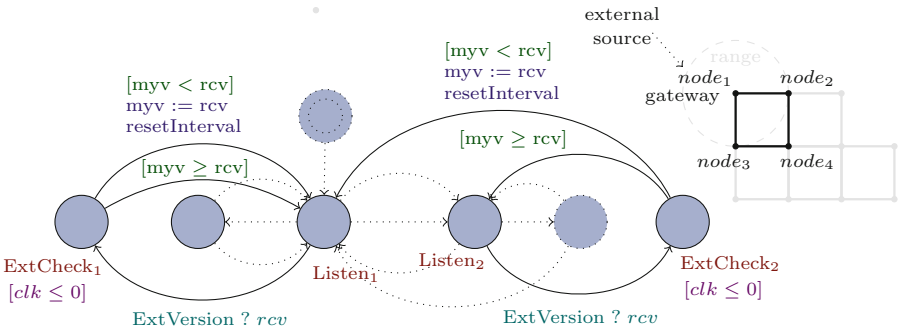


Fig. 2. Setting a network of XTA for a Trickle dissemination

State invariant $Inv(Listen_2) : clk \leq I$ constrains time elapsing to be at most of the value denoted by current interval length I . Subsequently, a new interval is

started by doubling I (until I_{max}) and a new t is chosen as previous (transition $Listen_2 \rightarrow Listen_1$). Similar reception handling as in location $Listen_1$ is defined.

Network Modelling. A transmission can be associated with several receivers, which are exactly those situated within the broadcast range of the node, i.e., they are its neighbours. A typical Trickle topology contains some gateway node $node_1$ in Fig. 2, which can receive versions to be disseminated across the network from an external source. A network of XTA one per node can be naturally designed for such topology. Among those, XTA of gateway nodes are extended with extra transitions which allows the reception of new versions from the external source (transitions $Listen_1 \rightarrow ExtCheck_1$ and $Listen_2 \rightarrow ExtCheck_2$ with input $ExtVersion?myv$). The connections between XTA ports, defined by function \mathcal{K} , are inferred from topology connections: e.g., $\mathcal{K}(node_1.Version) = \{node_2.Version, node_3.Version\}$ for the four-nodes topology depicted by the bidirectional graph in Fig. 2. Note that ports of gateway nodes are implicitly connected to external source and can receive any value.

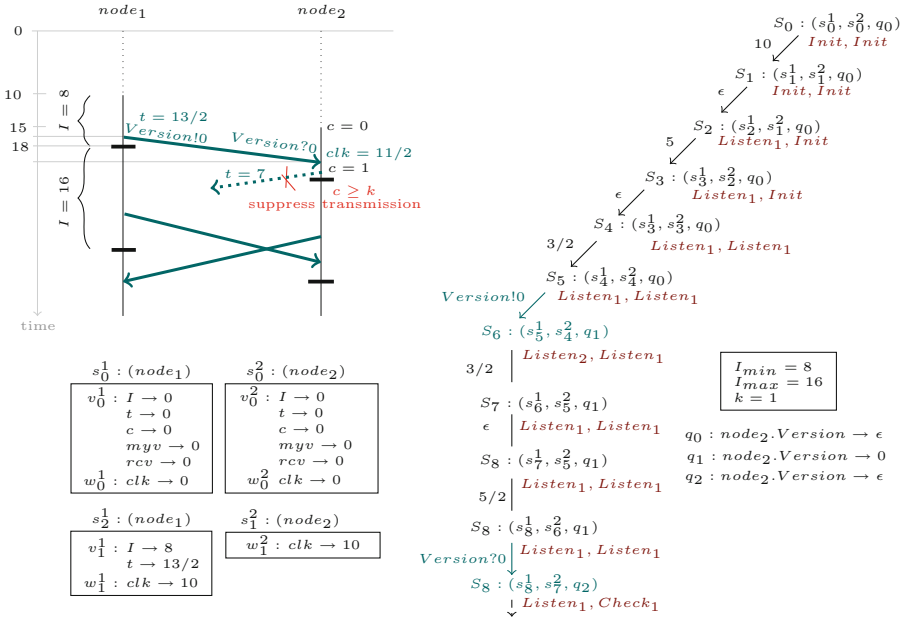


Fig. 3. A run of a two-nodes network of XTA.

Illustration of Network Runs. Figure 3 depicts a simple sequence diagram together with a run of a simple two-nodes network of XTA. The run shows that both nodes exchange a version of value 0, that they both initially hold (see data valuations in initial state S_0). As the redundancy constant k is set to 1, the receiver node gets its counter saturated, i.e., c reaches k . Therefore, it suppresses its transmission. This is a typical trickle behaviour which reduces the number of transmissions (gossip) when the neighbourhood is up-to-date.

4 Exploring Trickle with UPPAAL and DIVERSITY

UPPAAL Model. We have created a model in UPPAAL, which corresponds to the network of XTA presented in Sect. 3. The UPPAAL model has a very similar structure in terms of states and transitions. Locations *Check*₁ and *Check*₂ have been declared as *Committed* (marked with a “C”) which means that time cannot elapse in this location as intended in the original model. Also, such locations have a higher priority to be taken than non-committed ones. This reduces interleaving between automata if the latter are not executing on their turn transitions from committed locations. To implement asynchronous communication actions, we have created c-like functions in UPPAAL which implement fifo operations on queues. Unlike XTA which uses unbounded queues, those are arrays of fixed parameter size `QUEUE_SIZE`.

In UPPAAL, clocks are only compared to integer expressions, and clock guards are essentially conjunctions. This does not allow the specification of guards of the form $clk \leq t$ where t can take any random value in the dense interval $[I/2, I[$. In XTA, those values are (isomorphic to) positive rationals (\mathbb{Q}_+). In Fig. 4, we propose a UPPAAL pattern so that values assigned with t are positive integers (\mathbb{Z}_+). This is compliant with the nature of clock constraints supported by UPPAAL, i.e., clocks in guards are bounded by integer expressions. On the other hand, UPPAAL provides a *select* statement $s : [L, U]$ on transition which selects a random value for an integer variable s within a specified integer interval. Interval bounds L and U are necessarily constants. It is equivalent to an update $L \leq s \leq U$ in XTA. The pattern allows the selection of at most N values for t within the current interval $[I/2, I[$: first an integer s is selected in the interval $[0, N - 1]$, s is then used to assign t with a value denoted by the expression $e = (N + s) \cdot I / (2 \cdot N)$, however t is an integer variable, so t will be assigned exactly by the greatest integer less than or equal to the valuation of e . For instance, for $I = 8$ (resp. for $I = 16$) and $N = 4$, t can be assigned with the four integers 4, 5, 6 and 7 (resp. 8, 10, 12 and 14). The variable t can have an infinite number of possible values within the second half of I ; the pattern allows exploring with UPPAAL just a few (at most N integers, we experiment with small values). But since we are interested in reachability, if a solution exists for those, the verification concludes.

Model Exploration in UPPAAL. The tool uses the notion of *zone* to represent the set of valuations of clocks symbolically. A zone is defined by the conjunction of difference constraints of the form $clk \prec s$ or $clk - clk' \prec s$ where s is an integer. The simulation graph in UPPAAL is composed of nodes of the form $((l_1, v_1), \dots, (l_n, v_n), Z, v)$ where l_i is the location (resp. v_i is the data valuation)




Fig. 4. UPPAAL transitions.

for the i^{th} automaton, Z is a zone over clocks of the n involved automata, and v is the valuation of global or shared data variables. An example of such node is $((Listen_1, v_0^1), (Init, v_0^2), node_1.clk \leq 4, v_0)$ with notation of Fig. 3, v_0 associates Trickle constants to their values, it sets queues to empty at start. The exploration of the graph uses inclusion on zones which checks whether a zone of a successor node in the graph is already covered by some zones of previously explored nodes. In case of inclusion, data valuations must coincide in order to prune the search. This helps master the search when infinite cycles exist. Intuitively, a typical cycle is when all nodes reach I_{max} , share the same version and the content of the queues coincides, then same behaviours will start over again. To enable detect this situation: (i) we choose small values for I_{max} and (ii) we stop increment the counter c once it reaches k (see transition $Check_1 \rightarrow Listen_1$ in Fig. 4). Note that *saturation* is an extra clock of the node that will be discussed later; obviously, this has no effect on Trickle behaviour since the decision to suppress transmission depends only on reaching exactly k . In fact, after exchanging k or many more redundant versions is similar concerning subsequent behaviours. Otherwise, counter c will be assigned differently depending on the number of received versions. In which case, matching data valuations fails despite zone inclusion, and the cycle never exit. UPPAAL provides classical search strategies Depth-First Search (DFS) and Breadth-First Search (BFS), as well as Random Depth-First Search (RDFS). As we consider reachability, so one solution is wanted, DFS or RDFS are typically the most efficient option according to the tool documentation. When applying a DFS (or RDFS), inclusion on zones is of practical use as is it avoids getting lost in an infinite cycle.

Reachability Properties in UPPAAL. The tool supports a subset of Computation Tree Logic (CTL) [10]. As we are interested in reachability, we propose to use CTL formulae of the form $E\Diamond\phi$ where $\phi = f \wedge g$ is a formula on data variables and clocks. The satisfaction of such formula is defined on the tree with root an S_0 extracted from the labelled transition system of the network (see Sect. 2). The operator E quantifies over paths (or runs) of such a tree: it checks if there exists a path (with root S_0) in the tree satisfying the sub formula $\Diamond\phi$, the latter on the other hand is satisfied by that path if a state S satisfying ϕ occurs in the path. A simple property is $E\Diamond(node_2.c \geq k \wedge node_2.clk = node_2.I)$. The property is satisfied by the run discussed in Fig. 3, in which the node $node_2$ suppresses its transmission. Let us discuss the following two properties expressed for the four-nodes topology of Fig. 2.

Updated: $E\Diamond((node_1.myv = NEW) \wedge (node_2.myv = NEW) \wedge (node_3.myv = NEW) \wedge (node_4.myv = NEW))$

Outdated: $E\Diamond((node_4.isStarted) \wedge (node_4.myv = OLD) \wedge (node_2.myv = NEW) \wedge (node_2.c \geq k) \wedge (node_3.myv = NEW) \wedge (node_3.c \geq k) \wedge (node_2.saturation \geq D) \wedge (node_3.saturation \geq D))$

The first property states that it is possible that all the nodes are updated. The second property states that there exists a node which is still outdated (holds an old version) while its neighbours are all updated, yet they have suppressed

their transmissions. We use an extra clock *saturation* per node which is reset when the counter c reaches k (see Fig. 4). The clock measures the delay elapsed since then, and the formula requires that such delay is bounded by a parameter $D > 0$. This situation is not desirable, especially after having observed numerous exchanges of messages in the networks.




Fig. 5. Symbolic execution of a DIVERSITY transition.

DIVERSITY Models and Symbolic Execution. DIVERSITY [3] tools provides symbolic execution for state-based models (e.g., [4]) involving data expressions, data and clock guards. Input parameters or fresh symbols (they are used only once) can substitute uninitialized variables, reception variables used in communication with the external environment or any variable using a dedicated explicit *newfresh* statement. The latter is of the form *newfresh*(x); it associates the variable x with a new fresh symbol. We have developed a DIVERSITY model of the XTA network for Trickle, as the UPPAAL model it has a similar structure in terms of states and transitions. DIVERSITY provides communication over unbounded fifo queues. We declare a single queue per node as each node owns only one port for internal communications. Those are initially empty, and their size is automatically adjusted by the tool as communication actions are evaluated. Figure 5 depicts a DIVERSITY transition ($Init \rightarrow Listen_1$). It suggests a pattern in DIVERSITY which allows to assign the variable t with a random value within current interval $[I/2, I]$: *newfresh*(t) associates with t with a new fresh symbol, the latter is constrained by subsequent *guard statement* $I/2 \leq t < I$. It is possible in DIVERSITY to declare I_{min} , I_{max} , I and t to be typed as a positive rational numbers (\mathbb{Q}_+) so as to be compatible with clock clk . DIVERSITY computes the so-called symbolic tree in which nodes are called *execution contexts* ec : they store pieces of information about the execution including the current location, about the transition which allows reaching the context, and importantly about Path Conditions and Substitutions of data variables, clocks, queue places, ... by arithmetical expressions over input parameters. Figure 5 depicts the symbolic execution step of the previous transition (being of $node_1$) from ec_1 . The context ec_2 is reached, the transmission variable t is substitutes by a new symbol $node_1.t_1$ (substituted by t_0 in initial context ec_0), t_1 is constrained by the PC sub-formula $((I_{min}/2) \leq node_1.t_1 < I_{min})$. Note that sub-formula

($z_3 \leq node1.t_1$) shows that time elapsing in ec_2 (denoted by duration symbol z_3) is bounded by the value denoted by $node1.t_1$. This condition corresponds to the evaluation of clock invariant $clk \leq t$ in location $Init$ for $node_1$ ($node_1.clk$ is substituted by z_3 in ec_2). Symbolic execution techniques characterize all intended runs. DIVERSITY provides different classical search strategies which can be used to unfold the symbolic tree from the initial context ec_0 up to criteria on tree size (depth, width or number of nodes). Naturally, this results in a huge tree. DIVERSITY provides heuristic search [3] guided by a user-specified sequence of transitions, possibly non-consecutive as it is difficult in general to guess strict sequencing when it comes to automata network. We use UPPAAL to find such a sequence corresponding to some runs satisfying a user-specified property. Let us now overview the connection between both tools.

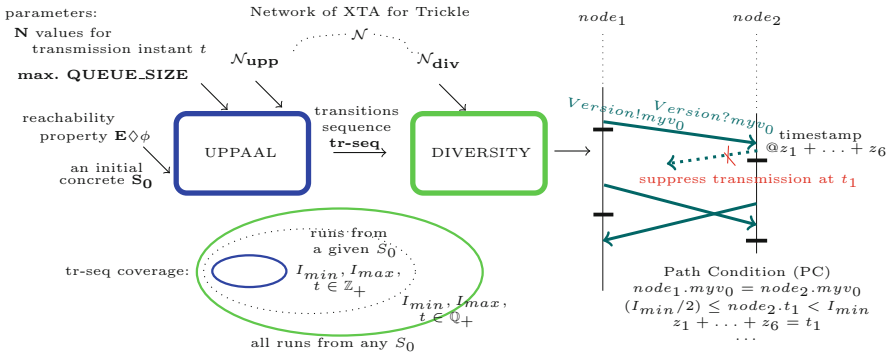


Fig. 6. Workflow UPPAAL-DIVERSITY. (Color figure online)

Tools Workflow and Experiments. The workflow is given in Fig. 6. UPPAAL takes as input a user-specified property $E\Diamond\phi$. The aim is to check the property from an initial state S_0 . Parameters on the number N of transmission instants t to be selected within the Trickle interval have to be given together with a bound $QUEUE_SIZE$ on communication queues. The property is assessed on the UPPAAL model N_{supp} . In case of the property is verified, i.e., there exists at least a run from S_0 which satisfies ϕ (such runs are schematically depicted by a blue ellipse), then the sequence of covering transitions $tr-seq$ is derived from the UPPAAL solution. The idea now is to compute the equivalence class composed by all runs from any states S_0 which are covered by the sequence $tr-seq$. For this, a symbolic exploration guided by the sequence is conducted on the DIVERSITY model N_{div} , a symbolic path together with its Path Condition (PC) is obtained then. The path is naturally feasible (satisfiability of its PC is assessed with SMT solvers). It identifies all runs covered by $tr-seq$ with dense domain for I_{min} , I_{max} and t as being positive rationals (those are schematically depicted by a green ellipse). Since the path often represents pairwise communication actions from different automata, it can be depicted in a natural manner

as a Sequence Diagram (SD). Figure 7 depicts an SD which highlights an outdated node situation, that of $node_4$. It has been computed by DIVERSITY for the four-nodes grid topology given in Fig. 2. The situation is atypical: neighbors of $node_4$, that is $node_2$ and $node_3$ are first updated by gateway $node_1$ with a new version (green messages), they hence reset their interval to I_{min} ; right away, $node_4$ gets them to reset their interval again by transmitting its old version (blue messages), their transmissions are postponed; this somehow gives $node_1$ time to retransmit the new version and saturate them (orange messages); therefore they suppress their transmissions for $node_4$. This unfortunate circumstance for $node_4$ can be prevented by increasing the redundancy constant so that its neighbors, $node_2$ and $node_3$, can still transmit, yet this comes at a cost, the number of messages increases for the entire nodes lifetime. Or $node_4$ can still be updated later because it together with $node_1$ will get their intervals doubled, and their transmission instants are at least dephased of I_{min} of those of $node_2$ and $node_3$ which leaves them time to update $node_4$. This is in favour of using Trickle, even if a node is in such outdated situation, it will not remain for a long time, thanks to the dynamic interval adjustment which avoids flood the network with messages. We have experimented with more nodes in the grid (up to 9), Table 1 reports on those. *Updated* and *Outdated* properties were successfully checked, which is satisfactory given the non-trivial kind of interactions in case of outdated nodes.

Table 1. The nodes graph is bidirectional in the form of a grid topology (See Fig. 2), 7(2) (resp. 7(3)) denotes that $node_7$ is two blocks (resp. three blocks) from the gateway $node_1$. Results measured on an Intel Core i7-7920HQ processor with RAM 32 GB, the redundancy constant k was set to 1 the minimum transmission case in Trickle. DIVERSITY time includes PC check with CVC4. UPPAAL concluded in few more trials inline with the exponential growth in running time.

Nodes	Updated			Outdated		
	UPPAAL	DIVERSITY	Messages	UPPAAL	DIVERSITY	Messages
3	1 ms	1 s 542 ms	7	12 ms	4 s 156 ms	8
4	3 ms	2 s 731 ms	16	984 ms	9 s 796 ms	12
5	4 ms	9 s 395 ms	12	168 ms	11 s 828 ms	12
6	7 ms	32 s 750 ms	21	7 s 621 ms	1 m 28 s 953 ms	23
7(2)	6 ms	29 s 453 ms	29	5 s 772 ms	2 m 5 s 375 ms	26
7(3)	5 ms	49 s 375 ms	29	3 s 227 ms	1 m 42 s 718 ms	24
8	8 ms	59 s 640 ms	41	8 m 9 s 503 ms	9 m 36 s 375 ms	42
9	9 ms	2 m 58 s 408 ms	64	9 m 19 s 22 ms	15 m 36 s 11 ms	62




Fig. 7. Outdated node situation - Sequence Diagram generated by DIVERSITY. (Color figure online)

5 Conclusion

We have developed first models for desynchronized Trickle network. Those models use data to express adjustable transmission intervals and abstract transmitted information. To assess reachability properties, we combine model checking and symbolic execution. If the property is verified on the model in which data is concrete, we derive from the returned solution a sequence of transitions that guides the symbolic execution to compute the corresponding symbolic path. The latter is a compact representation of the equivalence class of behaviours for the transitions coverage and can be depicted in the user-friendly format of sequence diagrams to enable their understanding. For future work, we plan to extract from the model-checking solution more information about structural coverage of the mini-language of updates on transitions to refine the equivalence classes by symbolic execution. We also plan to experiment with other gossip protocols. We believe that our approach can be of practical use to highlight for those non-trivial nodes interactions and give hints on their benefit for transmissions control.

Acknowledgment. This work was financially supported by European commission through ECSEL-JU 2018 program under grant agreement No. 826276.

References

1. RPL: IPv6 routing protocol for low-power and lossy networks, request for comments: 6550. Technical report, Cooper Power Systems and Cisco Systems and Stanford University, March 2012 (2012)
2. Alur, R., Dill, D.: A theory of timed automata. *J. Theoret. Comput. Sci.* **126**, 183–235 (1994)
3. Arnaud, M., Bannour, B., Lapitre, A.: An illustrative use case of the DIVERSITY platform based on UML interaction scenarios. *Electr. Notes Theoret. Comput. Sci.* **320**, 21–34 (2016)
4. Bannour, B., Escobedo, J.P., Gaston, C., Le Gall, P.: Off-line test case generation for timed symbolic model-based conformance testing. In: Nielsen, B., Weise, C. (eds.) *ICTSS 2012*. LNCS, vol. 7641, pp. 119–135. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-34691-0_10
5. Bannour, B., Lapitre, A.: Heuristic-aided symbolic simulation for trickle-based wireless sensors networks configuration. In: *International Workshop on RAPIDO@HiPEAC*. ACM (2020)
6. Barrett, C., et al.: CVC4. In: Gopalakrishnan, G., Qadeer, S. (eds.) *CAV 2011*. LNCS, vol. 6806, pp. 171–177. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-22110-1_14
7. Becker, M., Kuladinithi, K., Görg, C.: Modelling and simulating the Trickle algorithm. In: Pentikousis, K., Aguiar, R., Sargento, S., Agüero, R. (eds.) *MONAMI 2011*. LNICST, vol. 97, pp. 135–144. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-30422-4_10
8. Breza, M.J., McCann, J.A.: Lessons in implementing bio-inspired algorithms on wireless sensor networks. In: *International Conference on NASA/ESA*. IEEE (2008)

9. Chen, Z., Zhang, D., Zhu, R., Ma, Y., Yin, P., Xie, F.: A review of automated formal verification of ad hoc routing protocols for wireless sensor networks. CoRR, abs/1305.7410 (2013)
10. Clarke, E.M., Emerson, E.A.: Design and synthesis of synchronization skeletons using branching time temporal logic. In: Grumberg, O., Veith, H. (eds.) 25 Years of Model Checking. LNCS, vol. 5000, pp. 196–215. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-69850-0_12
11. Clarke, E.M., Grumberg, O., Peled, D.A.: Model Checking. MIT Press, Cambridge (2001)
12. Coladon, T., Vucinic, M., Tourancheau, B.: Multiple redundancy constants with trickle. In: PIMRC. IEEE (2015)
13. de Moura, L., Bjørner, N.: Z3: an efficient SMT solver. In: Ramakrishnan, C.R., Rehof, J. (eds.) TACAS 2008. LNCS, vol. 4963, pp. 337–340. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-78800-3_24
14. Dill, D.L.: Timing assumptions and verification of finite-state concurrent systems. In: Sifakis, J. (ed.) CAV 1989. LNCS, vol. 407, pp. 197–212. Springer, Heidelberg (1990). https://doi.org/10.1007/3-540-52148-8_17
15. Dong, J.S., Sun, J., Sun, J., Taguchi, K., Zhang, X.: Specifying and verifying sensor networks: an experiment of formal methods. In: Liu, S., Maibaum, T., Araki, K. (eds.) ICFEM 2008. LNCS, vol. 5256, pp. 318–337. Springer, Heidelberg (2008). https://doi.org/10.1007/978-3-540-88194-0_20
16. Dunkels, A., Gronvall, B., Voigt, T.: Contiki - a lightweight and flexible operating system for tiny networked sensors. In: IEEE ICLCN (2004)
17. Hui, J., Kelsey, R.: Multicast protocol for low-power and lossy networks, request for comments: 7731. Technical report, Silicon Labs (2016)
18. King, J.C.: Symbolic execution and program testing. Commun. ACM **19**, 385–394 (1976)
19. Kermajani, H.R., Gomez, C., Arshad, M.H.: Modeling the message count of the Trickle algorithm in a steady-state, static wireless sensor network. IEEE Commun. Lett. **16**, 1960–1963 (2012)
20. Kwiatkowska, M.Z., Norman, G., Parker, D.: Analysis of a gossip protocol in PRISM. SIGMETRICS Perform. Eval. Rev. **36**, 17–22 (2008)
21. Kwiatkowska, M., Norman, G., Parker, D.: PRISM 4.0: verification of probabilistic real-time systems. In: Gopalakrishnan, G., Qadeer, S. (eds.) CAV 2011. LNCS, vol. 6806, pp. 585–591. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-22110-1_47
22. Lee, B., Song, H.K., Suh, Y., Oh, K.H., Youn, H.Y.: Energy-efficient gossiping protocol of WSN with realtime streaming data. In: International Conference on DASC (2014)
23. Levis, P., Clausen, T., Hui, J., Gnawali, O., Ko, J.: The Trickle algorithm, request for comments: 6206. Technical report, March 2011 (2011)
24. Levis, P., et al.: TinyOS: an operating system for sensor networks. In: Weber, W., Rabaey, J.M., Aarts, E. (eds.) Ambient Intelligence. Springer, Heidelberg (2005)
25. Levis, P., Patel, N., Culler, D., Shenker, S.: Trickle: a self-regulating algorithm for code propagation and maintenance in wireless sensor networks. In: International Symposium on NSDI. USENIX Association (2004)
26. Liu, Y., Sun, J., Dong, J.S.: PAT 3: an extensible architecture for building multi-domain model checkers. In: Dohi, T., Cukic, B. (eds.) International Conference on ISSRE. IEEE (2011)

27. Meyfroyt, T., Borst, S.C., Boxma, O.J., Denteneer, D.: On the scalability and message count of Trickle-based broadcasting schemes. *Queueing Syst.* **81**, 203–230 (2015)
28. Meyfroyt, T.M.M.: An analytic evaluation of the Trickle algorithm: towards efficient, fair, fast and reliable data dissemination. In: *WoWMoM*. IEEE (2015)
29. Nan, X., Fei, M., Yang, T.: Randomized and efficient time synchronization in dynamic wireless sensor networks: a gossip-consensus-based approach. *Complexity* **2018**, 1–16 (2018)
30. Nguyen, N.M.T., Bannour, B., Lapitre, A., Le Gall, P.: Behavioral models and scenario selection for testing IoT Trickle-based lossy multicast networks. In: *International Workshop on VVIoT@ICST*. IEEE (2019)
31. Vucinic, M., Król, M., Jonglez, B., Coladon, T., Tourancheau, B.: Trickle-D: high fairness and low transmission load with dynamic redundancy. *IEEE IoT J.* **4**, 1477–1488 (2017)
32. Webster, M., Breza, M., Dixon, C., Fisher, M., McCann, J.A.: Formal verification of synchronisation, gossip and environmental effects for wireless sensor networks. In: *ECEASST* (2018)
33. Woehrle, M., Bakhshi, R., Mousavi, M.R.: Mechanized extraction of topology anti-patterns in wireless networks. In: Derrick, J., Gnesi, S., Latella, D., Treharne, H. (eds.) *IFM 2012*. LNCS, vol. 7321, pp. 158–173. Springer, Heidelberg (2012). https://doi.org/10.1007/978-3-642-30729-4_12
34. Zheng, M., Sun, J., Liu, Y., Dong, J.S., Gu, Yu.: Towards a model checker for NesC and wireless sensor networks. In: Qin, S., Qiu, Z. (eds.) *ICFEM 2011*. LNCS, vol. 6991, pp. 372–387. Springer, Heidelberg (2011). https://doi.org/10.1007/978-3-642-24559-6_26