On the Verification of Timed Ad Hoc Networks

Parosh Aziz Abdulla[1], Giorgio Delzanno[2], Othmane Rezine[1], Arnaud Sangnier[3*], and Riccardo Traverso[2]

[1] Uppsala University   [2] University of Genova   [3] Liafa-University Paris 7

Abstract. We study decidability and undecidability results for parameterized verification of a formal model of timed Ad Hoc network protocols. The communication topology is represented by a graph and the behavior of each node is represented by a timed automaton communicating with its neighbors via broadcast messages. We consider verification problems formulated in terms of reachability, starting from initial configurations of arbitrary size, of a configuration that contain at least one occurrence of a node in a certain state. We study the problem for dense and discrete time and compare the results with those obtained for (fully connected) networks of timed automata.

1 Introduction

In recent years there has been an increasing interest in analysis and verification methods for ad hoc networks, see e.g. [7,8,11,14,15,16]. Ad hoc networks consist of collections of wireless hosts that communicate with their neighbors by sending broadcast messages. In this context, protocols are supposed to work independently from a specific configuration of the network. Indeed, discovery protocols are often applied in order to identify the vicinity of a given node. In the AHN model proposed in [7] undirected graphs are used to represent a network in which each node executes an instance of a fixed (untimed) interaction protocol. Since individual nodes are not aware of the network topology, it is natural to consider here verification problems in which the size and shape of the initial configuration is not fixed a priori. For an untimed model of ad hoc network protocols, decidability and undecidability of this kind of parameterized verification problems are studied in [7]. We observe however that protocols for ad hoc networks are often based on time-sensitive conditions like time-outs and time-stamps added to flooded data. Parameterized verification of timed automata has been studied for fully connected networks with rendez-vous communication (Timed Networks), e.g., in [3,4,5,6].

A natural combination of the AHN model in [7] and of Timed Networks is obtained by adding a connectivity graph to a network of timed automata communicating via broadcast messages, we call the resulting model Timed Ad hoc Networks (TAHNs). For a fixed initial configuration, TAHNs can be specified in model checkers like Uppaal [17] by using a shared global matrix to specify the

---

communication topology. This idea has been used to verify safety properties of the LMAC protocol for initial configurations with a small number of nodes [11].

Following [7,8], in this paper we study decidability and undecidability properties of the local state reachability problem parametric on the initial configuration of a TAHN, i.e., the problem of checking the existence of an initial configuration that can evolve using continuous and discrete steps into a configuration exposing a given local state – usually representing an error.

Compared to the positive results obtained for Timed Networks where the control state reachability problem is decidable for processes with a single clock, for the same type of processes the problem becomes undecidable in a very simple class of topologies in which nodes are connected so as to form stars with diameter five. The undecidability result can be ported to the more general class of graphs with bounded path (for some bound $N \geq 5$ on the length – number of nodes – of paths). In the untimed case local state reachability is decidable for bounded path topologies [7]. Furthermore, the problem turns out to be undecidable in the class of cliques of arbitrary order (that contains graphs with arbitrarily long paths) in which each timed automaton has at least two clocks. Decidability holds for special topologies like stars with diameter three and cliques of arbitrary order assuming that the timed automaton in each node has a single clock (as in Timed Networks).

For discrete time, we show that the local state reachability problem becomes decidable for processes with any number of clocks in the class of graphs with bounded path. The same holds for cliques of arbitrary order as in the case of dense time.

Related Work Decidability issues for untimed models of ad hoc Networks have been considered in [7,8]. Abstraction techniques for untimed selective broadcast protocols has been considered in [14]. Model checking for timed automata has been applied to verify protocols for ad hoc networks with a fixed number of nodes in [11]. Models with a discrete global clock and a lazy exploration of configurations of fixed size has been considered in [16]. Formal specification languages for timed models of ad hoc networks have been proposed, e.g., in [13]. In contrast to these works, we consider computatibility issues for verification of timed ad hoc networks with parametric initial configurations. Decidability of some cases is proved by resorting to an extension of Timed Networks with transfers. In the untimed case the combination of rendez-vous and transfers is considered in Datanets, a model in which processes have data taken from an ordered domain [12].

Outline In the next section we give some preliminaries. The models of Timed Ad Hoc Networks and Time Networks are introduced in Sections 3–4. In Section 5, we introduce the undecidability results for three different topologies. We give decidability for different topologies under the dense time semantic in Section 6 and under the discrete time semantics in Section 7. Finally, we give some conclusions in Section 8.
Detailed proofs are given in [2].

2 Preliminaries

We use $\mathbb{N}$ and $\mathbb{R}^{\geq 0}$ for the set of natural numbers and set of non-negative real numbers respectively. For sets A and B, we use $f : A \mapsto B$ to denote that f is a total function that maps A to B. For $a \in A$ and $b \in B$, we use $f[a \hookleftarrow b]$ to denote the function f' defined as follows: $f'(a) = b$ and $f'(a') = f(a')$ for all $a' \neq a$. We use $[A \mapsto B]$ to denote the set of all total functions from A to B.

3 Timed Ad Hoc Networks

An ad hoc network consists of a graph where the nodes represent processes that run a common predefined protocol. This protocol is defined by a communicating timed automaton. The values of the clocks of the automata inside the processes are incremented continuously all at the same rate. In addition, a process may perform a discrete transition. The latter is either a local transition or the result of a communication event. In a *local* transition, the process changes its local state without interacting with the other processes. Communication is performed through *selective broadcast*, where a process sends a broadcast message to the network. The effect of a broadcast is local to the vicinity of the sender, i.e., only the neighbors of the sending process are able to receive the message. The connectivity of the nodes is reflected by the edges of the graph. Furthermore, transitions are conditioned by values of the clocks of the process, and may reset the values of some clocks.

We assume that each process operates on a set X of clocks. A *guard* is a Boolean combination of predicates of the form $k \lhd x$ for $k \in \mathbb{N}$, $\lhd \in \{=, <, \leq, >, \geq\}$, and $x \in X$. A *reset* R is a subset of X. We will use guards to impose conditions on the clocks of processes that participate in transitions, and use resets to identify the clocks that will be reset during the transition. A *clock valuation* is a mapping $\mathtt{X} : X \mapsto \mathbb{R}^{\geq 0}$. For a guard g and a clock valuation $\mathtt{X}$, we write $\mathtt{X} \models g$ to indicate the validity of the formula we get by replacing each clock x in g by its value $\mathtt{X}(x)$. Also, we will assume a finite alphabet Σ. The alphabet induces a set of events, where an event is of one of three forms: (i) *empty event* τ that represents a local move; (ii) *broadcast event* $!!a$, with $a \in \Sigma$, that represents broadcasting the message a; or (ii) *receive event* $??a$ with $a \in \Sigma$, that represents receiving the message a (that has been broadcast by another process).

Formally, a *Timed Ad Hoc Network* (*TAHN* for short) is defined by a pair $\mathcal{T} = (G, P)$. The first component $G = (V, E)$ is graph where V is a finite set of vertices and $E \subseteq V \times V$ is a set of edges. The second component P, called the *protocol*, is a pair $(Q, \mathcal{R})$ where Q is a finite set of *states*, and $\mathcal{R}$ is a finite set of *rules*. Intuitively, the graph G defines the topology of $\mathcal{T}$ where the set V represents the nodes, and E defines the connectivity of the nodes. The vertices belonging to an edge are called the *endpoints* of the edge. For an edge $(u, v) \in E$, we often use the notation $u \sim v$ and say that the vertices u and v are adjacent to each other. Furthermore, P defines the protocol that runs inside the nodes, where Q is the set of *local states* of each node, while $\mathcal{R}$ is a set of *rules* describing

the behavior of each node. A rule $\rho \in \mathcal{R}$ is of the form $\left(q, g \xrightarrow{e} R, q'\right)$ where $q, q' \in Q$, g is a guard, e is an event, and R is a reset. We use $\texttt{CLOCKS}(P)$ to denote that number of clocks inside each process.

Configurations A configuration γ is a pair $(\mathcal{Q}, \mathcal{X})$, where $\mathcal{Q} : V \mapsto Q$ and $\mathcal{X} : V \mapsto [X \to \mathbb{R}^{\geq 0}]$, i.e., the configuration assigns to each node a local state and assigns to each clock in the node a value (in $\mathbb{R}^{\geq 0}$).

Transition Relation For configurations $\gamma = (\mathcal{Q}, \mathcal{X})$ and $\gamma' = (\mathcal{Q}', \mathcal{X}')$, we write $\gamma \Longrightarrow_\mathcal{T} \gamma'$ to denote that one of the following conditions is satisfied:

- *Local:* There exists a rule $\rho = \left(q, g \xrightarrow{\tau} R, q'\right)$ and a vertex $v \in V$ such that $\mathcal{Q}(v) = q$, $\mathcal{X}(v) \models g$, $\mathcal{Q}' = \mathcal{Q}\left[v \hookleftarrow q'\right]$, and $\mathcal{X}' = \mathcal{X}\left[v \hookleftarrow \texttt{X}\right]$ where $\texttt{X}(x) = 0$ if $x \in R$ and $\texttt{X}(x) = \mathcal{X}(v)(x)$ otherwise.
- *Broadcast:* There exists a rule $\left(q, g \xrightarrow{!!a} R, q'\right)$ and a vertex $v \in V$ such that $\mathcal{Q}(v) = q$, $\mathcal{X}(v) \models g$, $\mathcal{Q}'(v) = q'$, $\mathcal{X}'(v)(x) = 0$ if $x \in R$, and $\mathcal{X}'(v)(x) = \mathcal{X}(v)(x)$ otherwise. Furthermore, for each $v_1 \in V - \{v\}$, one of the following conditions is satisfied:
    - $v_1 \sim v$ and there is a rule of the form $\left(q_1, g_1 \xrightarrow{??a} R_1, q_1'\right)$ such that $\mathcal{Q}(v_1) = q_1$, $\mathcal{X}(v_1) \models g_1$, $\mathcal{Q}'(v_1) = q_1'$, $\mathcal{X}'(v_1)(x) = 0$ if $x \in R_1$, and $\mathcal{X}'(v_1)(x) = \mathcal{X}(v_1)(x)$ otherwise.
    - $\mathcal{Q}'(v_1) = \mathcal{Q}(v_1)$, $\mathcal{X}'(v_1) = \mathcal{X}(v_1)$, and either $v_1 \not\sim v$ or there is no rule of the form $\left(q_1, g_1 \xrightarrow{??a} R_1, q_1'\right)$ for any g_1, R_1, q_1' with $\mathcal{Q}(v_1) = q_1$ and $\mathcal{X}(v_1) \models g_1$.
- *Time:* There is a $\delta \in \mathbb{R}^{\geq 0}$ such that $\mathcal{Q}(v') = \mathcal{Q}(v)$ and $\mathcal{X}'(v)(x) = \mathcal{X}(v)(x) + \delta$ for all $v \in V$ and $x \in X$.

Topology A *topology Top* restricts the shape of the underlying graph $G = (V, E)$ in a TAHN $\mathcal{T} = (P, G)$. We write $G \in Top$ to indicate that G satisfies *Top*. Below, we give examples of some topologies.

- We denote by $\texttt{GRAPH}$ the topology consisting of all finite graphs.
- The star topology of depth ℓ (with $\ell \geq 0$), denoted $\texttt{STAR}(\ell)$, characterizes graphs G for which there is a partitioning $\{v_0\} \cup V_1 \cup \cdots \cup V_\ell$ of V such that (i) $v_0 \sim v_1$ for all $v_1 \in V_1$, (ii) for each $1 \leq i < \ell$ and $v_i \in V_i$ there is exactly one $v_{i+1} \in V_{i+1}$ with $v_i \sim v_{i+1}$, (iii) for each $1 < i \leq \ell$ and $v_i \in V_i$ there is exactly one $v_{i-1} \in V_{i-1}$ with $v_i \sim v_{i-1}$, and (iv) no other nodes are adjacent to each other. In other words, in a star topology of dimension ℓ, there is a central node v_0 and an arbitrary number of *rays*. A ray consists of a sequence of ℓ nodes, starting from v_0 followed by some $v_1 \in V_1$, and then by some $v_2 \in V_2$, etc. We call v_0 the *root*, call the nodes in $V_1, \ldots, V_{\ell-1}$ the *internal nodes*, and call the nodes in V_ℓ the *leaves* of G.
- The bounded path topology of bound ℓ (with $\ell \geq 0$), denoted $\texttt{BOUNDED}(\ell)$, characterizes graphs G for which the length of the maximal simple path in

G is bounded by ℓ. This means that there does not exist a finite sequence of vertices $(v_i)_{1 \le i \le m}$ satisfying the following conditions (1) $m > \ell$, (2) $v_i \ne v_j$ for all i, j in $\{1, \ldots, m\}$ such that $i \ne j$ and (3) $v_i \sim v_{i+1}$ for all i such that $1 \le i < m - 1$.

  – The set of cliques, denoted CLIQUE characterizes graphs G where $v_1 \sim v_2$ for all $v_1, v_2 \in V$ with $v_1 \ne v_2$.

Reachability We assume a distinguished local state $q^{init} \in Q$. A configuration γ^{init} is said to be *initial* if it is of the form $\left(\mathcal{Q}_Q^{init}, \mathcal{X}^{init}\right)$ where $\mathcal{Q}_Q^{init}(v) = q^{init}$, and $\mathcal{Q}_Q^{init}(v)(x) = 0$ for all $v \in V$ and $x \in X$. In other words, all the processes are in their initial local states and all the clocks have value 0. A *computation* π of $\mathcal{T}$ is a sequence $\gamma_0 \Longrightarrow_\mathcal{T} \gamma_1 \Longrightarrow_\mathcal{T} \cdots \Longrightarrow_\mathcal{T} \gamma_n$ where γ_0 is an initial configuration. In such a case, we say that γ_n is *reachable* in $\mathcal{T}$. For a local state $q \in Q$, we say that q is *reachable* in $\mathcal{T}$ if there is a configuration $\gamma = (\mathcal{Q}, \mathcal{X})$ such that γ is reachable in $\mathcal{T}$ and $\mathcal{Q}(v) = q$ for some $v \in V$.

The *(local state) reachability problem* for a topology *Top* and a number K, denoted TAHN−Reach(*Top*, K), is defined as follows:

> Given a protocol P with CLOCKS$(P) = K$ and a local state $q \in Q$, is there a TAHN $\mathcal{T} = (P, G)$ such that $G \in$ *Top* and q is reachable in $\mathcal{T}$.

The following results concerning untimed Ad Hoc Networks have been proved.

Theorem 1 ([7]). TAHN−Reach(GRAPH, 0) *is undecidable. For each $K \ge 0$, the problem* TAHN−Reach(BOUNDED(K), 0) *is decidable.*

4 Timed Networks

In this section, we recall the model of *Timed Networks* (*TN* for short) [6]. In a similar manner to TAHNs, a TN contains an arbitrary number of identical *timed processes* that operate on a finite number of local real-valued clocks. However, there are three main differences between TNs and TAHNs. First, a TN contains a distinguished *controller* that is a finite-state automaton without any clocks[1]. Second, each process in a TN may communicate with all other processes and hence it is not meaningful to describe topologies in the case of TNs. Finally, communication takes place through rendez-vous between fixed sets of processes rather than broadcast messages. In a similar manner to TAHNs, the values of all clocks in a TN are incremented continuously at the same rate. In addition, the TN can change its configuration according to a finite number of *rules*. Each rule describes a set of transitions in which the controller and a fixed number of processes synchronize and simultaneously change their states. A rule may be conditioned on the local state of the controller, together with the local states and clock values of the processes. If the conditions for a rule are satisfied, then a transition may be performed where the controller and each participating process

---

[1] This is the model defined in [6]. Adding clocks to the controller does not affect the decidability results.

changes its state. During a transition, a process may reset some of its clocks to 0.

We assume a finite set X of clocks and define guards and resets in a similar manner to TAHNs. A *family of timed networks* (*timed network* for short) $\mathcal{N}$ is a pair $(Q, \mathcal{R})$, where Q is a finite set of *states*, partitioned into a set Q^{ctrl} of *controller states*, and a set Q^{proc} of *process states*; and $\mathcal{R}$ is a finite set of *rules* where each rule is of the form

$$[q_0 \to q_0'] \, [q_1; g_1 \to R_1; q_1'] \, \cdots \, [q_n; g_n \to R_n; q_n']$$

such that $q_0, q_0' \in Q^{ctrl}$, for all $i : 1 \leq i \leq n$ we have: $q_i, q_i' \in Q^{proc}$, g_i is a guard, and R_i is a reset. Intuitively, the set Q^{ctrl} represents the states of the controller and the set Q^{proc} represents the states of the processes. A rule of the above form describes a set of transitions of the network. The rule is enabled if the state of the controller is q_0 and if there are n processes with states $q_1, \cdots, q_n$ whose clock values satisfy the corresponding guards. The rule is executed by simultaneously changing the state of the controller to q_0' and the states of the n processes to $q_1', \cdots, q_n'$, and resetting the clocks belonging to the sets $R_1, \ldots, R_n$.

Configurations A configuration γ of a timed network $(Q, \mathcal{R})$ is a tuple of the form $(I, q, \mathcal{Q}, \mathcal{X})$, where I is a finite *index set*, $q \in Q^{ctrl}$, $\mathcal{Q} : I \to Q^{proc}$, and $\mathcal{X} : I \to X \to \mathbb{R}^{\geq 0}$. Intuitively, the configuration γ refers to the controller whose state is q, and to $|I|$ processes, whose states are defined by $\mathcal{Q}$. The clock values of the processes are defined by $\mathcal{X}$. More precisely, for $i \in I$ and $x \in X$, $\mathcal{X}(i)(x)$ gives the value of clock x in the process with index i. We use $|\gamma|$ to denote the number of processes in γ, i.e., $|\gamma| = |I|$.

Transition Relation The timed network $\mathcal{N}$ above induces a transition relation $\longrightarrow_{\mathcal{N}}$ on the set of configurations. The relation $\longrightarrow_{\mathcal{N}}$ is the union of a *discrete* transition relation $\longrightarrow_D$, representing transitions induced by the rules, and a *timed* transition relation $\longrightarrow_T$ which represents passage of time.

The discrete relation $\longrightarrow_D$ is the union $\bigcup_{r \in \mathcal{R}} \longrightarrow_r$, where $\longrightarrow_r$ represents a transition performed according to rule r. Let r be a rule of the form described in the above definition of timed networks. Consider two configurations $\gamma = (I, q, \mathcal{Q}, \mathcal{X})$ and $\gamma' = (I, q', \mathcal{Q}', \mathcal{X}')$. We use $\gamma \longrightarrow_r \gamma'$ to denote that there is an injection $h : \{1, \ldots, n\} \to I$ such that for each $i : 1 \leq i \leq n$ we have:

1. $q = q_0$, $\mathcal{Q}(h(i)) = q_i$, and $\mathcal{X}(h(i)) \models g_i$. That is, the rule r is enabled.
2. $q' = q_0'$, and $\mathcal{Q}'(h(i)) = q_i'$. The states are changed according to r.
3. If $x \in R_i$ then $\mathcal{X}'(h(i))(x) = 0$, while if $x \notin R_i$ then $\mathcal{X}'(h(i))(x) = \mathcal{X}(h(i))(x)$. In other words, a clock is reset to 0 if it occurs in the corresponding set R_i. Otherwise its value remains unchanged.
4. $\mathcal{Q}'(j) = \mathcal{Q}(j)$ and $\mathcal{X}_k'(j) = \mathcal{X}_k(j)$, for $j \in I \setminus \mathrm{range}(h)$, i.e., the process states and the clock values of the non-participating processes remain unchanged.

A *timed transition* is of the form $\gamma \longrightarrow_{T=\delta} \gamma'$ where $\gamma = (I, q, \mathcal{Q}, \mathcal{X})$, $\delta \in \mathbb{R}^{\geq 0}$, $\gamma' = (I, q, \mathcal{Q}, \mathcal{X}')$, $\mathcal{X}'(j)(x) = \mathcal{X}(j)(x) + \delta$ for all $j \in I$ and $x \in X$. We use $\gamma \longrightarrow_T \gamma'$ to denote that $\gamma \longrightarrow_{T=\delta} \gamma'$ for some $\delta \in \mathbb{R}^{\geq 0}$.

We define $\longrightarrow_{\mathcal{N}}$ to be $\longrightarrow_D \cup \longrightarrow_T$ and use $\stackrel{*}{\longrightarrow}_{\mathcal{N}}$ to denote the reflexive transitive closure of $\longrightarrow_{\mathcal{N}}$. Notice that if $\gamma \longrightarrow_{\mathcal{N}} \gamma'$ then the index sets of γ and γ' are identical and therefore $|\gamma| = |\gamma'|$. For a configuration γ and a controller state q, we use $\gamma \stackrel{*}{\longrightarrow}_{\mathcal{N}} q$ to denote that there is a configuration γ' of the form $(I', q, \mathcal{Q}', \mathcal{X}')$ such that $\gamma \stackrel{*}{\longrightarrow}_{\mathcal{N}} \gamma'$.

Given $\gamma_0 \longrightarrow_{\mathcal{N}} \gamma_1 \longrightarrow_{\mathcal{N}} \gamma_2 \ldots \longrightarrow_{\mathcal{N}} \gamma_n$, we say that $\gamma_0, \ldots, \gamma_n$ is a computation of $\mathcal{N}$.

Reachability. We assume a distinguished initial controller state $q_{ctrl}^{init} \in Q^{ctrl}$ and a distinguished initial process state $q_{proc}^{init} \in Q^{proc}$. A configuration $\gamma^{init} = (I, q, \mathcal{Q}, \mathcal{X})$ is said to be *initial* if $q = q_{ctrl}^{init}$, $\mathcal{Q}(i) = q_{proc}^{init}$, and $\mathcal{X}(i)(x) = 0$ for each $i \in I$ and $x \in X$. This means that an execution of a timed network starts from a configuration where the controller and all the processes are in their initial states, and the clock values are all equal to 0. Notice that there is an infinite number of initial configurations, namely one for each index set I. Concepts such as that of computations and reachability are extended from TAHNs to TNs in the obvious way.

The *(controller state) reachability problem* $\mathtt{TN-Reach}(K)$ is defined by a timed network $(Q, \mathcal{R})$ with K clocks (in each process), and a controller state q. The task is to check whether q is reachable or not. In [3], the following result is shown.

Theorem 2. $\mathtt{TN-Reach}(2)$ *is undecidable.*

5 Undecidability with Dense Time

In this section, we show undecidability of the reachability problem for three classes of TAHNs, namely (i) those with *star* topologies of depth 2 (one root and several rays with two nodes) and with a single clock in each node; (ii) those with *clique* topologies provided that each node has two clocks; and (iii) those with *bounded path* topologies if the length of simple paths of the underlying graph is at least 5 and with a single clock in each node. In the first two cases, the undecidability result is shown through a reduction from $\mathtt{TN-Reach}(2)$ (that is undecidable by Theorem 2). The main idea of the proofs is to show that we can simulate rendez-vous (the communication model of TNs) by broadcast (the communication model of TAHNs) In each case we will simulate a TN $\mathcal{N}$ with two clocks per process by a TAHN $\mathcal{T}$. We will refer to the clocks inside a process of $\mathcal{N}$ as x_1 and x_2 respectively. For each state q in $\mathcal{N}$, we will have a corresponding state $\mathcal{T}(q)$ in $\mathcal{T}$. Furthermore, we will have a number of auxiliary states in $\mathcal{T}$ that we need to perform the simulation. The third undecidability result (for bounded path topologies) is shown by simulating the star case.

Two-Star Topologies We show that the reachability problem for the star topology is undecidable even when the rays are restricted to have depth 2 and the nodes are restricted to have a single clock.

Theorem 3. $\mathtt{TAHN-Reach}(\mathtt{STAR}(2),1)$ *is undecidable.*

Given a TN $\mathcal{N} = (Q, \mathcal{R})$ and a controller state q in $\mathcal{N}$, we define a TAHN $\mathcal{T} = (P, G)$ such that $G \in \mathtt{STAR}(2)$ and $\mathtt{CLOCKS}(P) = 1$, together with a local state $\mathcal{T}(q)$, such that q is reachable in $\mathcal{T}$ iff q is reachable in $\mathcal{N}$. The root of $\mathcal{T}$ plays the role of the controller in $\mathcal{N}$. Furthermore, each ray in $\mathcal{T}$ plays the role of one process in $\mathcal{N}$. The local state of a process in $\mathcal{T}$ is stored in the internal node of the corresponding ray. Furthermore, the two clocks x_1, x_2 of a process are represented by the clock of the internal node resp. the clock of the leaf of the ray. For technical reasons, we require that $\mathcal{T}$ has at least three rays. In case $\mathcal{N}$ has fewer than three processes, the additional rays will not simulate any processes, and remain passive (except during the initialization phase; see below). The simulation consists of two phases.

Initialization Recall that the nodes of a TAHN are identical in the sense that they execute the same (predefined) protocol. This means that the nodes are not *a priori* aware of their positions inside the network. The purpose of the initialization phase is to identify the nodes that play the roles of the controller and those that play the roles of the different processes. First, a node may broadcast a message where it requests to become the node that simulates the controller in $\mathcal{N}$. In order to succeed, P requires that it should receive acknowledgements from at least three other processes.

Notice that only the root of $\mathcal{T}$ can be successful since it is the only node that is connected to more than two other nodes (the internal nodes are connected to two other nodes while the leafs are connected to only one other node).

Once the root has become the controller, it will make the internal nodes aware of their positions. It does that by sending a broadcast message. Due to the star topology, this message is received only by the internal nodes. A node receiving this broadcast message will initiate a "local protocol" inside its ray as follows: (i) It changes local state to reflect that it now knows that it is indeed an internal node. (ii) It makes the leaf of the ray aware of its position by broadcasting a message. Such a message is received only by the leaf of the ray and by the root (the latter simply ignores the message). (iii) The leaf broadcasts an acknowledgment (that can only be received by the internal node of the ray). (iv) The internal node changes state when it receives the acknowledgement and declares itself ready for the next step. Notice that the internal node and the leaf may choose to ignore performing steps (ii) or (iv). In such a case we say that the ray has "failed", otherwise we declare the ray to be "successful". In the last step of the initialization, the root will send one more broadcast where the following takes places: (i) It changes local state to $\mathcal{T}(q_{ctrl}^{init})$ which means that it is now simulating the initial controller state. (ii) It checks that its clock is equal to 0 which means that the initialization phase has been performed instantaneously. (iii) The internal nodes of the successful rays will change state to $\mathcal{T}(q_{proc}^{init})$. The rest of the nodes will remain passive throughout the rest of the simulation. Now all the nodes are ready: the root of $\mathcal{T}$ is in the initial state of the controller of $\mathcal{N}$; the internal nodes of the successful rays are in the initial states of the processes of $\mathcal{N}$, and all clocks have values equal to 0.

Simulating Discrete Transitions Below, we show how $\mathcal{T}$ simulates a rule of the form $[q_0 \to q'_0]\ [q_1; g_1 \to R_1; q'_1]\ \cdots\ [q_n; g_n \to R_n; q'_n]$. The root of $\mathcal{T}$ is in the state $\mathcal{T}(q_0)$. First, the root resets its clock to 0 (this is done so that it can later make sure that the simulation of the rule has not taken time). The simulation consists of different phases, where in each phase the root tries to identify a ray that can play the role of process k for $1 \leq k \leq n$. To find the first ray, it sends a broadcast message. An (internal) node that receives the broadcast and whose local state is q_1 may either decide to ignore the message or to try to become the node that simulates the first process in the rule. In the latter case it will enter a temporary state from which it initiates a sub-protocol whose goal is to confirm its status as the simulator of the first process. In doing so, the node has guessed (perhaps wrongly) that its clocks satisfy the values specified by the guard. If the node has guessed wrongly it will eventually be excluded from the rest of the simulation (will remain passive in the rest of the simulation). At the end of this phase, exactly one node will be chosen among the ones that have correctly guessed that their clocks satisfy g_1. The successful node will be the one that plays the role of the first process. The sub-protocol proceeds as follows: (i) The internal node checks whether the value of its clock satisfies the guard g_1. Recall that each node contains one clock. Since the guard g_1 only compares the clocks x_1, x_2 with constants, the conditions of g_1 can be tested on each of x_1 and x_2 separately. If the clock of the node does not satisfy g_1 (which means that x_1 does not satisfy g_1), the node will remain passive from now on (it has made the wrong guess). Otherwise, the node resets its clock if R_1 contains x_1, and then broadcasts a message (such a message is received by the leaf of the ray); (ii) The leaf checks whether the value of its clock satisfies the guard g_1 (i.e., if x_2 satisfies g_1); if *yes* it resets its clock if x_2 is included in R_1, and then broadcasts an acknowledgement. (iii) Upon receiving the above acknowledgement, the internal node declares itself ready for the next step by broadcasting an acknowledgement itself. At the same time, it moves to new local state and waits for a last acknowledgement from the root (described below) after which it will move to local state $\mathcal{T}(q'_1)$. (iv) When the root receives the acknowledgement it sends a broadcast declaring that it has successfully found a ray to simulate the first process. All the nodes in temporary states will now enter local states from which they remain passive. To prevent multiple nodes to play the role of the first process, the root enters en error state if it happens to receive acknowledgements from several internal nodes. The root now proceeds to identify the ray to simulate the second process. This continues until all n processes have been identified. Then the root makes one final move where the following events take place: (i) It moves its local state to $\mathcal{T}(q'_0)$ (ii) It sends a final boraodcast where the node ready for simulating the i^{th} process will now move to $\mathcal{T}(q'_i)$ for all $i : 1 \leq i \leq n$ (notice that there is at most one such node for each i). (iii) It checks that its clock is equal to 0 (the simulation of the rule has not taken any time).

Simulating Timed Transitions This is done in a straightforward manner by letting time pass in $\mathcal{T}$ by the same amount as it has done in $\mathcal{N}$.

Cliques We show that the reachability problem for the clique topology is undecidable if the nodes have two clocks.

Theorem 4. $\text{TAHN}-\text{Reach}(\text{CLIQUE}, 2)$ *is undecidable.*

We will build a protocol P with $\text{CLOCKS}(P) = 2$ which will simulate $\mathcal{N}$ on the clique topology. In a similar manner to the case of star topologies, the simulation consists of two phases.

Initialization Phase The purpose of the initialization phase is to choose a node that will simulate the controller. This choice is done non-deterministically through a protocol that is initialized by a broadcast message. Notice that this protocol exists in all the nodes since they run the same pre-defined protocol. The first node which will perform the broadcast will become the controller (from now on we refer to this node as the *controller node*). When the controller node performs the above broadcast it moves to the state $\mathcal{T}(q_{ctrl}^{init})$, while all the other nodes will move to $\mathcal{T}(q_{proc}^{init})$.

Simulating Discrete Transitions Below, we show how a rule of the form of the previous sub-section is simulated. In a similar manner to the case of stars, the controller node first resets its clock to 0. The simulation again consists of different phases, where in each phase the controller node tries to identify a ray that can play the role of process i for $1 \leq i \leq n$. To find the first process it sends a broadcast. A node that receives the broadcast, whose local state is q_1, and whose clocks (x_1 and x_2) satisfy the guard g_1, may decide to ignore the message or to try to become the node that simulates the first process in the rule. In the latter case, the node declares itself ready for the next step by broadcasting an acknowledgement. At the same time, it moves to a new local state and waits for a last acknowledgement from the controller node (described below) after which it will move to local state $\mathcal{T}(q_1')$. To prevent multiple nodes to play the role of the first process, the controller node enters en error state if it happens to receive acknowledgements from several nodes. The controller node now proceeds to identify the node to simulate the second process. This continues until all n processes have been identified. Then the controller node performs the same three steps as the ones in the final phase of the simulation described above for stars.

Bounded Path Topologies Using the result of Theorem 3 we can show that the reachability problem can be extended to bounded path topologies. The result uses a reduction to the two-star case, thus we need to consider topologies in which the (number of vertices) simple paths can have 5 vertices in order to be able to rebuild stars with rays of depth 2.

For such a reduction, we need a preliminary protocol that discovers a two-star topology in an arbitrary graph in paths are allowed to have (at least) five nodes. The discovery protocol first selects root, internal and leaf candidates and then verify that they are connected in the desired way by sending all other nodes in their vicinities to a special null state. Combining the discovery protocol and the undecidability results for two-star topology we obtain the following theorem.

Theorem 5. $\mathtt{TAHN-Reach}(\mathtt{BOUNDED}(5), 1)$ *is undecidable.*

6 Decidability with Dense Time

In the previous section we showed that $\mathtt{TAHN-Reach}(\mathtt{STAR}(2), 1)$ is undecidable. In this section we consider two other classes of topologies for which reachability becomes decidable when nodes have a single clock, namely the class $\mathtt{STAR}(1)$ and $\mathtt{CLIQUE}$. A convenient way to prove these results is to resort to an extension of Timed Networks with transfers for which control state reachability is still decidable. When executed together with a rendez-vous, a transfer action from q to q' forces all processes in state q to move to state q' (as in transfer arcs for Petri nets). We give next a more detailed definition of the extended TN model.

Timed Networks with Transfers (TNT) A rule of a *timed network with transfer* (TNT) $\mathcal{N}$ combines rendez-vous synchronization and transfer actions. Namely, a TNT rule has the following form:

$$[q_0 \to q_0'] \ [q_1; g_1 \to R_1; q_1'] \ \cdots \ [q_n; g_n \to R_n; q_n']$$
$$[p_1; g_1' \to_t R_1'; p_1'] \cdots [p_\ell; g_\ell' \to_t R_\ell'; p_\ell']$$

where actions with $\to$ denote rendez-vous communication, whereas actions with $\to_t$ denote transfers. For $i : 1 \le i \le \ell$, $g_i \to_t R_i$ is a guarded transfer where g_i is a guard, R_i is a reset, $p_i, p_i' \in Q^{proc}$. We assume that there are no $i : 1 \le i \le \ell$ and $j : 1 \le j \le n$ such that $p_i = q_j$; and assume that $p_i \ne p_j$ if $i \ne j$. A TNT rule is enabled if the state of the controller is q_0 and if there are n processes with states $q_1, \cdots, q_n$ whose clock values satisfy the corresponding guards. The rule is executed by simultaneously changing the state of the controller to q_0' and together with the states of selected processes $p_1, \cdots, p_\ell$. Furthermore, each process in state p_i for $i : 1 \le i \le \ell$ whose clock value satisfies the guard g_i' moves to p_i' while its clocks are reset according to R_i'. Note that for a rule to be enabled there is no requirement on the presence or absence of processes in states $p_1, \cdots, p_\ell$.

The (controller state) reachability problem $\mathtt{TNT-Reach}(K)$ for processes with K clocks is defined by replacing the TN transition relation with that for TNT. The following result then holds.

Theorem 6. $\mathtt{TNT-Reach}(1)$ *is decidable.*

Sketch of proof. In [6], it is proved that $\mathtt{TN-Reach}(1)$ is decidable. The proof is based on the general results in [1] based on a well quasi orderings of TN configurations under which the transition relation is monotonic. Monotonicity of the transition relation of a TNT still holds for the same ordering used in [6]. Furthermore, the algorithm used for computing predecessors can be extended in order to cope with transfer action. This extension is similar to that used for Data nets [12] or, for processes without clocks, to that used for Petri nets with transfer arcs. We prove next that in TAHNs with one clock restricted to the clique topology, the reachability problem is decidable.

Theorem 7. TAHN−Reach(CLIQUE, 1) *is decidable.*

Proof. We reduce TAHN−Reach(CLIQUE, 1) to TNT−Reach(1). The reduction works as follows. Since in a clique graph all nodes are connected to each other, a broadcast message sent by a node is always received by all other nodes. In other words working with a clique is like working with a multiset of processes as in a TNT. Broadcast communication can then be simulated by using TNT rules in which the sender perfoms an individual step and reception of a message is simulated by transfer actions, one for each state in which the message can be received. Furthermore, we can insert a special rule to synchronize the controller with the local state we want to reach in the TAHN. Local state reachability corresponds then to control state reachability in the resulting TNT. The claim then follows from Theorem 6.

A similar positive result can be obtained for TAHN with 1 clock restricted to the star topology of depth 1. The main difference compared to the previous result is that in a star of depth 1 we have to distinguish the root (the central node) from the leaves. When the root performs a broadcast, it is transmitted to all the leaf nodes, but when a leaf performs one, only the root can receive it.

Theorem 8. TAHN−Reach(STAR(1), 1) *is decidable.*

Proof. We reduce TAHN−Reach(STAR(1), 1) to TNT−Reach(1). Let $\mathcal{T}$ be a TAHN with star topology of depth 1. We construct a TNT N that simulates $\mathcal{T}$ as follows. Initially all TNT processes are in state q^i. We first define an initialization step in which the controller non-deterministically selects one of the processes in state q^i and elects it as the root of the simulated TAHN. All the remaining processes in state q^i are then used to simulate the leaves of the star. After this step, the simulation of each TAHN rule r is split in two TNT rules: one for the root process and one for a leaf process. A local rule is simulated by a rendez-vous (with no transfer) for the root process and by a rendez-vous for a leaf process. A broadcast rule executed by the root node is simulated with a TNT transfer action involving the root process and all the leaf processes, whereas a broadcast executed by a leaf is translated into a rendez-vous between a leaf and the root process only. As for cliques, we can add rules for the controller to observe when the root of a leaf process has reached a given local state. The claim then follows from Theorem 6.

7 Decidability with Discrete Time

In this section we study the reachability problem for Discrete Time Ad Hoc Networks (DTAHN). In this model clocks range over the natural numbers instead of the reals. Let μ' be the maximum constant used in the protocol rules and let $\mu = \mu' + 1$. When using discrete time, it is enough to restrict the valuation of clocks to the finite range $I = [0, \mu]$. This follows from the fact that guards of the form $x > c$ remain enabled when time passes once the clock associated to

variable x reaches a value greater or equal to μ; while guards of the form $x \leq c$ remain disabled. Therefore, beyond μ we need not distinguish between different values for the same clock (see e.g. [3]). For every DTAHN $\mathcal{T}$, we can then define a finite-state process that describes the behavior of a node. We use $\mathcal{C}_{\mu,K}$ to denote configurations over undirected graphs with labels in $Q \times I^K$, where Q is the set of local states of a process, I is the interval $[0, \mu]$, and K is the number of clocks in each process.

The transition relation $\rightharpoonup_{\mathcal{T}}$ is obtained from that of TAHN by assuming that the valuation of clock variables is a function $\mathcal{X} : V \mapsto [X \rightarrow I]$ and by replacing the time step by the discrete time step defined as follows.

Discrete Time Step For configurations $\gamma = (\mathcal{Q}, \mathcal{X})$ and $\gamma' = (\mathcal{Q}', \mathcal{X}')$, we write $\gamma \rightharpoonup_{\mathcal{T}} \gamma'$ if, for all $v \in V$ and $x \in X$, the following conditions are satisfied:

- $\mathcal{Q}(\gamma') = \mathcal{Q}(\gamma)$,
- $\mathcal{X}'(v)(x) = \mathcal{X}(v)(x) + 1$, if $\mathcal{X}(v)(x) < \mu$
- $\mathcal{X}'(v)(x) = \mathcal{X}(v)(x) = \mu$, otherwise.

For a topology class *Top* and $K \geq 0$, the control state reachability problem $\texttt{DTAHN} - \texttt{Reach}(\textit{Top}, K)$ is the natural reformulation of the one defined for TAHN.

We show next that reachability is decidable when restricting the topology to the class of bounded path graphs $\texttt{BOUNDED}(N)$ for any fixed $N > 1$. The decision procedure is obtained by resorting to the theory of well-structured transition systems [1]. The procedure is based on a symbolic backward exploration algorithm in which we use *constraints* to finitely represent sets of configurations of *variable size* taken from the class $\texttt{BOUNDED}(N)$ (the configurations may potentially belong to different TAHNs).

Ordering We first introduce the following ordering between configurations of variable size. Given configurations $\gamma = (\mathcal{Q}, \mathcal{X})$ defined over $G = (V, E)$ and $\gamma' = (\mathcal{Q}', \mathcal{X}')$ defined over $G' = (V', E')$, $\gamma \preceq \gamma'$ iff there exists an injective function $h : V \mapsto V'$ such that:

- $\forall u, u' \in V$, $(u, u') \in E$ if and only if $(h(u), h(u')) \in E'$;
- $\forall u \in V$, $\mathcal{Q}(u) = \mathcal{Q}'(h(u))$ and $\mathcal{X}(u) = \mathcal{X}'(h(u))$.

An upward closed set U satisfies the property that $U = \{\gamma' \mid \gamma \preceq \gamma', \ \gamma \in U\}$. The following property then holds.

Proposition 1. *If U is an upward closed set of configurations (of variable size), then $Pre(U) = \{\gamma \mid \gamma \rightharpoonup_{\mathcal{T}} \gamma', \ \gamma' \in U\}$ is still upward closed wrt. $\preceq$.*

Proof. In [7] it has been proven, that for untimed AHN, selective broadcast is monotone w.r.t. $\preceq$. We observe that DTAHN restricted to configurations in $\mathcal{C}_{\mu,K}$ can be viewed as untimed AHN extended with a time step transition. Thus, we just have to show monotonicity wrt. a time step $\gamma \rightharpoonup_{\mathcal{T}} \gamma'$. Since time can always proceed, for every $\beta \succeq \gamma$ there is a configuration β' such that $\beta \rightharpoonup_{\mathcal{T}} \beta'$ with a time step. Now since h is label preserving both time steps will have the same effect on nodes identified by h and thus $\gamma' \preceq \beta'$.

Constraints Assume a given process definition P. A constraint is a tuple $\Phi = (G, \mathcal{Q}, \mathcal{X})$, where $G = (V, E)$ is a graph in `BOUNDED`(N), and $\mathcal{Q}$ and $\mathcal{X}$ are defined on the set of vertices V. We call γ_Φ the associated configuration $(\mathcal{Q}, \mathcal{X})$ over the graph G. The denotation of a constraint Φ is defined as the infinite set of configurations $[\![\Phi]\!] = \{\gamma' \mid \gamma_\Phi \preceq \gamma'\}$ with variable size and topology. The following proposition then holds.

Proposition 2. *Given a constraint Φ, there exists an algorithm to compute a finite set of constraints whose denotation corresponds to $Pre([\![\Phi]\!])$.*

Proof. We can extend the symbolic predecessor operator defined for untimed AHN in [7] by adding some new conditions and steps. In order to correctly obtain predecessors of γ, rules of the form $\left(q, g \xrightarrow{e} R, q'\right)$ must be applied only to nodes in state q' which satisfy the post-condition R. When R is not empty, i.e. some clocks have been reset, we do not know their values before the step. In this case it is necessary to add a predecessor to the set $Pre(\gamma)$ for every possible combination of $i \in I$ for those clocks – keeping the old values of the node from configuration γ otherwise.

As an addition to regular protocol rules, we also have to consider time steps. Thus, for every configuration γ that does not have 0-valued clocks, we have to add to its set of predecessors every configuration γ' such that γ' can make a time step into γ.

Theorem 9. `DTAHN`$-$`Reach`$(`BOUNDED`(N), K)$ *is decidable for any $N \geq 1$ and $K \geq 0$.*

Proof. From propositions 1 and 2, we can apply the general results in [1] to define a backward search algorithm working on upward closed sets of extended configurations represented by their finite basis. The finite set of constraints (G, Q, X) in which G is a single node v, $Q(v) = q$, and $X(v) \in I$ can be used as finite representation of all configurations containing the control state q. Furthermore, for a fixed $N \geq 1$ the induced subgraph relation is a well-quasi-ordering on the class of N-bounded path graphs [10]. This implies that for any sequence of constraints $\Phi_0 \Phi_1 \ldots$ there exists i and j with $i < j$ s.t. $\Phi_i \preceq \Phi_j$. This property ensures the termination of the symbolic backward reachability algorithm.

8 Conclusions

We have studied local state reachability for Timed Ad Hoc Networks in different classes of topologies by taking the number of clocks of each node as a parameter. In order to refine our analysis, it could be interesting to investigate restricted form of guards, e.g., open and closed specifications as in [5], to obtain decidable fragments of the full model. Furthermore, similarly to the untimed case [9], it would be interesting to evaluate the impact of failures and delays in the decidability of local reachability for TAHNs. From a more practical point of view, by exploiting symbolic representations of timed regions it could be interesting to

apply backward reachability analysis in order to extract preconditions (in form of minimal representation of topologies) for violations of safety conditions in protocol models like those considered in [11]. Finally, finding abstractions for topology- and time-dependent broadcast protocols that could be used to reason on arbitrary configurations is another interesting direction for future work.

References

1. P. Abdulla, K. Cerans, B. Jonsson, and Y. Tsay. General decidability theorems for infinite-state systems. In *LICS'96*, pages 313–321. IEEE Computer Society, 1996.
2. P. Abdulla, G. Delzanno, O. Rezine, A. Sangnier, and R. Traverso. On the verification of timed ad hoc networks. Disi-11-05, Univ. of Genova, July 2011.
3. P. Abdulla, J. Deneux, and P. Mahata. Multi-clock timed networks. In *LICS'04*, pages 345–354. IEEE Computer Society, 2004.
4. P. Abdulla and A. Nylén. Timed petri nets and bqos. In *ICATPN'01*, volume 2075 of *LNCS*, pages 53–70. Springer, 2001.
5. P. A. Abdulla, J. Deneux, and P. Mahata. Closed, open, and robust timed networks. *ENTCS*, 138(3):117–151, 2005.
6. P. A. Abdulla and B. Jonsson. Model checking of systems with many identical timed processes. *TCS*, 290(1):241–264, 2003.
7. G. Delzanno, A. Sangnier, and G. Zavattaro. Parameterized verification of ad hoc networks. In *CONCUR'10*, volume 6269 of *LNCS*. Springer, 2010.
8. G. Delzanno, A. Sangnier, and G. Zavattaro. On the power of cliques in the parameterized verification of ad hoc networks. In *FoSSaCS'11*, volume 6604 of *LNCS*, pages 441–455. Springer, 2011.
9. G. Delzanno, A. Sangnier, and G. Zavattaro. Verification of ad hoc networks with node and communication failures. Disi-11-06, Univ. of Genova, July 2011.
10. G. Ding. Subgraphs and well quasi ordering. *J. of Graph Theory*, 16(5):489–502, 1992.
11. A. Fehnker, L. van Hoesel, and A. Mader. Modelling and verification of the LMAC protocol for wireless sensor networks. In *IFM'07*, volume 4591 of *LNCS*, pages 253–272. Springer, 2007.
12. R. Lazic, T. Newcomb, J. Ouaknine, A. Roscoe, and J. Worrell. Nets with tokens which carry data. *Fund. Inf.*, 88(3):251–274, 2008.
13. M. Merro, F. F. Ballardin, and E. Sibilio. A timed calculus for wireless systems. In *FSEN '09*, volume 5961 of *LNCS*, pages 228–243. Springer, 2010.
14. S. Nanz, F. Nielson, and H. Nielson. Static analysis of topology-dependent broadcast networks. *Inf. Comput.*, 208(2):117–139, 2010.
15. M. Saksena, O. Wibling, and B. Jonsson. Graph grammar modeling and verification of Ad Hoc Routing Protocols. In *TACAS'08*, volume 4963 of *LNCS*, pages 18–32. Springer, 2008.
16. A. Singh, C. R. Ramakrishnan, and S. A. Smolka. Query-based model checking of ad hoc network protocols. In *CONCUR'09*, volume 5710 of *LNCS*, pages 603–619. Springer, 2009.
17. The uppaal web page. http://www.uppaal.com/.