Modelling Real-time Systems with Bigraphs

Maram Albalwe
University of Glasgow
Glasgow, UK
University of Tabuk
Tabuk, Saudi Arabia
m.albalwe.1@research.gla.ac.uk

Blair Archibald Michele Sevegnani
University of Glasgow
Glasgow, UK
{blair.archibald, michele.sevegnani}@glasgow.ac.uk

Bigraphical Reactive Systems (BRSs) are a graph-rewriting formalism describing systems evolving in two dimensions: spatially, *e.g.* a person in a room, and non-spatially, *e.g.* mobile phones communicating regardless of location. Despite use in domains including communication protocols, agent programming, biology, and security, there is no support for real-time systems. We extend BRSs to support real-time systems by using a modelling approach that employs multiple perspectives to represent digital clocks. We use Action BRSs, a recent extension of BRSs, where the resulting transition system is a Markov Decision Process (MDP). This allows a natural representation of the choices in each system state: to either allow time to pass or perform a specific action. We implement our proposed approach using the BigraphER toolkit, and demonstrate the effectiveness using multiple examples including the timed aspects of the Routing Protocol for Low-Power and Lossy Networks (RPL).

1 Introduction

Bigraphs are a computational model where systems are described based on two types of relationships: spatially via *nesting* (and parallel adjacency) and non-local linking through hyperlinks regardless of locations. Like standard graphs, bigraphs have an equivalent diagrammatic and algebraic representation. Bigraphical Reactive Systems (BRSs) equip bigraphs with a set of reaction rules that specify how a system evolves over time, *i.e.* reaction rules substitute sub-bigraphs with other bigraphs.

The standard theory of bigraphs has been extended to model a wider range of systems *e.g.* stochastic bigraphs [13] assign rates to reaction rules, bigraphs with sharing [23] allow intersecting locations, directed bigraphs [9] associate directions to links, conditional bigraphs [3] add conditions to rules, and probabilistic and action bigraphs [4] support non-determinism and probabilistic behaviour. Utilising these extensions, bigraphs have been successfully used in the literature to model a variety systems including mixed-reality games [7], cloud systems [20], self-adaptive fog systems [21], sensor network infrastructure [24], and rational agents [5].

One extension that has not been explored is real-time bigraphs that can model systems exhibiting non-deterministic behaviour that is controlled by time constraints. While the non-deterministic aspects of real-time systems could be modelled by action bigraphical reactive systems (ABRSs) in which the underlying Markov Decision Process (MDP) semantics allows for a transition choice at each state, there is currently no notion of *clocks constraints* in the theory, *e.g.* specifying that after 3 time units have passed, a given action *must* occur.

We propose a modelling strategy to encode timed systems within ABRSs through the well-known MDP-based digital clock approximation. We introduce a clock, as a new bigraph entity, for each timed entity in the system and we place these in a separate region so all the clocks can be manipulated without polluting the actual system model. This allows one reaction rule to advance all clocks at once, and this

Submitted to:
GCM 2024

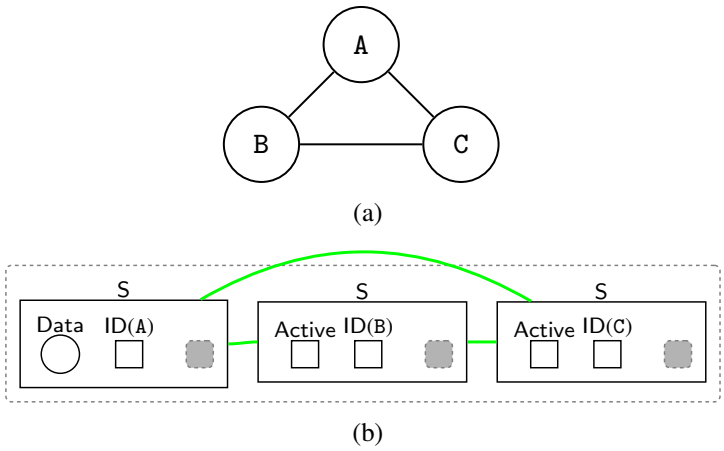

Figure 1: (a) Example network topology with three sensors; (b) corresponding bigraph with data to be transmitted by sensor A and the status of each sensor (*e.g.* Active) (b).

reduces the state space overhead of the approach. We mirror the real-time passage by introducing a global clock that controls the system execution time, *i.e.* system terminates when we reach the time limit of this clock, and we explicitly model the nondeterminism feature of real-time systems, *i.e.* at each state there is a choice between taking an action or allowing time to pass when the constraints are not yet satisfied.

We make the following research contributions:

- We propose a modelling strategy to express clocks within action bigraphs.
- We define a set of reaction rules to model clock constraints in real-time systems through a digital clocks approximation.
- We illustrate how to use our approach in practice by providing a model of some timed aspects of a real-world network protocol: the Routing Protocol for Low-Power and Lossy Networks (RPL).
- We implement this approach in the BigraphER [22] toolkit to show this is not just a theoretical contribution, but a practical one.

Outline. We give an informal description of bigraphs and BRSs in Section 2. Section 3 provides a description of non-deterministic models and shows how to formalise digital clocks within ABRSs. We illustrate our approach by providing a BRS implementation for a simple scenario modelled as a (probabilistic) timed automaton and a fragment of the RPL protocol in Section 4. Section 5 gives a brief literature review, and we conclude in Section 6.

2 Background

2.1 Bigraphs

Introduced by Milner [18], Bigraphs provide a powerful diagrammatic representation for modelling systems that evolve in both spatial and non-spatial dimensions. Bigraphs represent the structure of a system through two relations over its entities: a *place graph* that encodes their spatial arrangement, *e.g.* a person in a room, and a *link graph* that specifies, through hyper-edges, non-spatial relations, *e.g.* communication capabilities between network devices. We give an informal description of bigraphs using the example

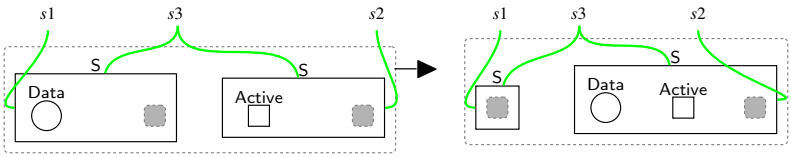

Figure 2: Reaction rule `send_Data` applies when the receiver is active.

in Figure 1. A comprehensive formal description can be found elsewhere [18]. Bigraphs have both a diagrammatic representation and *equivalent* algebraic notation. To provide an intuitive presentation we use the diagrammatic notion throughout this paper.

Figure 1a shows a simple network topology consisting of three sensors connected to each other by communication links. Data may be exchanged between two sensors when the receiver is active. A corresponding bigraph representation is in Figure 1b. Sensors are represented as *entities* of type S that have *nested* identifiers (also just a different type of entity) ID(A), ID(B), and ID(C). The sender contains an entity Data that can be sent to other sensors when they are in the Active mode. We sometimes draw different entity types by using different shapes or colours *e.g.* we have used squares for identifiers and a circles for Data. Entities can be *atomic*, *e.g.* Data, meaning they have no children, or contain any number of other entities. We allow entities to be parameterised to represent families of entities, *i.e.* ID(x) specifies a new entity for every possible value of x (which may be string, integer, or float typed). Sites—filled dashed rectangles—represent unspecified bigraphs: an arbitrary bigraph, including the empty bigraph, might exist there. Bigraphs may consist of more than one region—clear dashed rectangles—which represent adjacent parts of the system. For the link graph, entities have a fixed *arity* that represents the number of *links* they must have (the green edges). Links are hyperlinks that may connect 1-to-n. Open links—a link that has a name—indicates a link that may connect elsewhere, *i.e.* to currently unspecified entities. Links may also be closed, a 1-to-0 hyperedge, which is shown by an orthogonal line at the end of the link.

Open names, sites, and regions allow model composition: regions of one bigraph can be placed in the sites of another and like-names joined. This forms the basis of the rewriting theory.

2.2 Bigraphical Reactive Systems (BRSs)

A bigraph represents a system at a single point in time. To model dynamic behaviour, Bigraphical Reactive Systems (BRSs) equip bigraphs with a set of *reaction rules* specifying how the system may evolve. Reaction rules take the form $L \longrightarrow\!\!\!\!\blacktriangleright R$ meaning that when the rule is applicable to a state S (state S *matches* bigraph L) the system evolves by replacing an occurrence of bigraph L in S with bigraph R. BRSs form a transition system that has an initial state (*bigraph*), and the transition from one state (*bigraph*) to another is defined by generating all possible rewrites (*reactions*). For example, the rule in Figure 2 is applicable where there is a sensor that has Data to send to another Active sensor in its range.

We control the execution of a set of reaction rules via *priority classes*. That is $\{\mathtt{r}_1, \mathtt{r}_2\} < \{\mathtt{r}_3\}$ means rules $\mathtt{r}_1$ and $\mathtt{r}_2$ can be applied only if it is not possible to apply $\mathtt{r}_3$. Similarly to entities, rules can be parameterised to allow multiple rule applications over a predefined set of values.

3 Modelling Time with Action Bigraphs

The big idea is that by utilising Action Bigraphs [4] we can encode timed systems. This is based on a well-known approximation of (probabilistic) timed automata to Markov Decision Processes [11]. Intuitively, we extend the usual action semantics to allow two forms of action: *discrete actions* that encode system events (which may only be enabled at some time), and *time* actions that progress time.

Action Bigraphical Reactive Systems (ABRSs) allow a choice of probability distributions at each rewriting step by assigning *weights* to the reaction rules, and by giving action labels to specific sets of reaction rules. An action is enabled/possible whenever there is a reaction rule from the set which is enabled. The resulting transition system is a Markov Decision Processes (MDP) [6, 12] which can be verified using off-the-shelf model checkers such as PRISM [14]. An MDP is a tuple $(S, s_0, A, Step)$ where S is a set of states with a predefined initial state $s_0 \in S$. For each state, we can choose an (enabled) action $a \in A$ which gives an associated probability distribution over future states as specified by $Step$.

To show how ABRSs model non-determinism behaviour of real-time systems where underlying transition system is an MDP, we use the bigraph example shown in Figure 1b to include a choice of probability distributions. As an example, we can permit sensor `A` to send `Data` to either sensor by replacing `send_Data` rule (Figure 2) with two rules each represents a non-deterministic action through assigning the related rule with a weight (Figure 3). For example, with $weight = 0.7$ sensor `A` sends `Data` to sensor `B` or to `C` with $weight = 0.3$. Weights are then normalised to probabilities depending on which rule(s) are applicable in a state. We use *bigraphs* throughout this paper to mean *Action Bigraphical Reactive Systems (ABRSs)*. Using BigraphER [22], an open-source toolkit for working with bigraphs and it is the only tool supporting ABRSs, we export the corresponding MDPs transition system to be formally checked.

Importantly, action bigraphs still respect any rule priorities. This means an action will fire with any high priority rule before firing with those of lower priorities. We use this feature in our encoding of clock invariants.

A Probabilistic Timed Automata (PTA) [2] introduces support for clocks. A PTA is defined as a tuple in the form (L, l_0, E, A, C, I) where L is a set of locations (*i.e.* states) with $l_0 \in L$ a start location. E is a set of edges that represents the possible transitions between locations while A is a set of (discrete) actions that may occur in the system. C is a finite set of clocks and I is a set of clock invariants associated to each location. While these invariants could be flexible, practically they are typically used to force a transition from a location, *e.g.* an invariant $x \leq 2$ forces a location move at time 2 if we have not already left the location (alongside transition conditions these usually encode a "move within a maximum of 2 time units" semantics). Transitions are associated to actions and probability distributions, with the addition of *clock constraints* (*e.g.* $x > 3$) and *clock resets* (*e.g.* $x := 0$).

In the *digital clocks approximation* we give the semantics of PTAs as a Markov Decision Process where states (including initial) are all the locations where clock invariants are met, and we form a single action set consisting of both actions manipulating time and user-specified discrete actions as follows:

- a *discrete* action is enabled if all clock constraints (given in E) associated with the transition are satisfied;
- a *time* action is available while invariants associated to $l \in L$ are satisfied as time elapses. That is, we cannot progress time if something must happen beforehand.

Since we use action bigraphs, we can associate discrete transitions with a probability distribution. If required we can, for example, draw with uniform probability to recover semantics for a non-probabilistic timed automata.

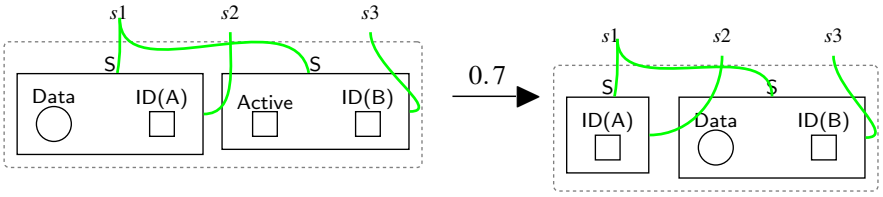

(a) $Sensor_A$ sends to $Sensor_B$ with weight 0.7

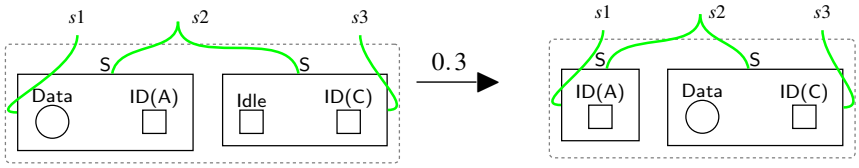

(b) $Sensor_A$ sends to $Sensor_C$ with weight 0.3

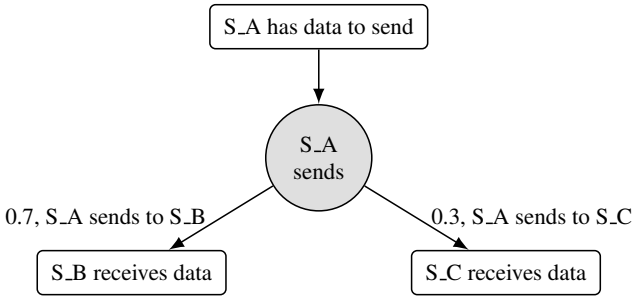

(c) Non-determinism model with ABRSs

Figure 3: Example Action Bigraphical Reactive System: probabilistic reaction rules using weights.

3.1 Digital Clocks in Bigraphs

Standard BRSs evolve to a new state whenever there is a match of a reaction rule[1]. In this sense they are non-deterministic but not explicitly so, *i.e.* we cannot label particular actions without ABRSs. Models of real-time systems need to both find enabled matches, and meet any requirements introduced by clock constraints. We propose a modelling technique that encodes digital clocks as *entities* to model real-time systems. Although PTAs can work with real-valued clocks, for the digital clocks approximation we fix clock values to non-negative integers and bound the total runtime of the system (to ensure a finite action set which means the models can be analysed through existing MDP reachability techniques). As for the standard digital clocks approximation, our approach does not support strict inequalities and comparisons between clocks. As we have action bigraphs, these clocks can live within the bigraph model itself, and a separate type of semantics is not needed (unlike for probabilistic bigraphs for example). We show our approach by an example, and the main idea is in Figure 4. We put all clocks into their own parallel region which we call the *clocks perspective*. A global clock is represented by an entity family $\mathsf{GC}(n)$—*i.e.* there is one entity for each $n \in \mathbb{M}$ for some max time $\mathbb{M}$—and is used to model *wall clock time*. Wall-time cannot be reset. Local clocks are entity families $\mathsf{LC}(n)$ where $0 \leq n \leq \mathbb{M}$ for each timed entity. For clarity of modelling, we place these in $\mathsf{LocalClocks}$, although this is not strictly required. We identify a set of *timed* entities and expand their arity by 1 to link them to a specific local clock. Using this, we can *identify* a clock through the linked entity, *i.e.* clocks do not need specific identifiers of their own. Not all

[1] Subject to any requirements about rule priorities and conditions.

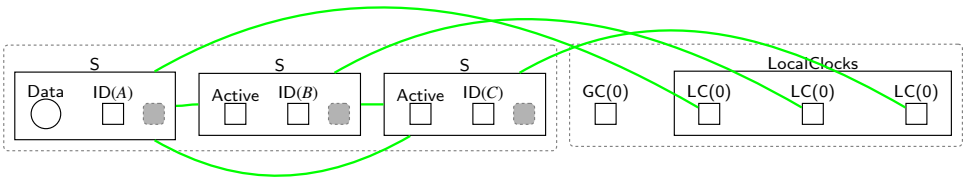

Figure 4: Example from Figure 1b extended with clocks perspective.

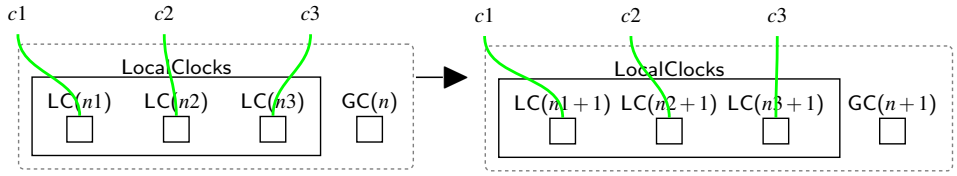

Figure 5: Reaction rule `clock_advance`$(n1,n2,n3,n)$ for a three timed entity system.

entities in the system will own a clock, *e.g.* Data is not timed. We assume all clocks are initialised to 0 in the initial state.

The approach to place clocks within their own region is a design choice, and because of the flexibility of bigraphs, other choices would be possible. For example, we could *nest* local clocks within particular timed entities instead of linking to them. We chose the extra region as it does not clutter any existing model, *i.e.* we can convert an existing model to a timed representation without significant changes (other than arity) within the existing regions.

3.2 Reaction Rules for Digital Clocks

In our digital clocks approximation we have two main types of action: *discrete* actions which are user-defined and system specific, and *time* actions that deal with the passage of time. As actions in ABRSs are modelled as sets of reaction rules, the main challenge here is defining appropriate reactions for each type of action.

For timing actions, the main rule is `clock_advance(x,y,...)` shown in Fig. 5. This rule advances all local and global clocks simultaneously. All clocks advance at the same speed (one unit per application). This is a parameterised rule, which, like paramterised entities, defines a family of rules *i.e.* one for each possible value of the parameters. The *tick* action consists of the set of all possible `clock_advance(x,y,...)` rules for parameters drawn from $x \in \mathbb{M}, y \in \mathbb{M}, \ldots$, for some max time $\mathbb{M}$. Due to the parameters only one instance of `clock_advance` will ever be enabled meaning this action always advances time with probability one.

For *discrete* actions we extend existing system rules whenever a time constraint must be met. We take rules that have a standard (untimed) definition of a bigraphs rule ($L \longrightarrow\!\!\!\!\blacktriangleright R$) and, like with the clocks perspective, utilise parallel regions to link timed entities with their related clocks as follows:

$$\mathsf{L}_c \| \mathsf{LC}(n)_c \longrightarrow\!\!\!\!\blacktriangleright \mathsf{R}_c \| \mathsf{LC}'(n)_c$$

Here the notation $\|$ is the algebraic operator placing L and the new clock LC in *different* regions (as seen graphically by the dashed lines). The subscript c means there is (at least) a link c that connects to other c linked entities. This is the new link we are using to identify specific clocks.

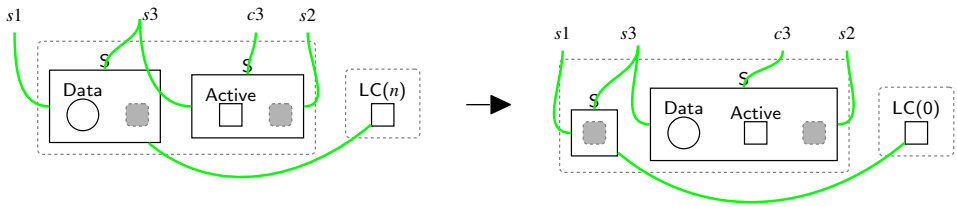

Figure 6: Reaction rule(s) `sending_data(n)`. Time constraints are encoded by setting $n \in \{2,3,\ldots,\mathbb{M}\}$.

Importantly, this change also makes existing rules into parameterised rules (over parameter n). This is used later to encode timing constraints, *e.g.* we chose a set of parameters that define at what (local) *time* a particular rule can be applied. For rules that are already parameterised, we can simply add an additional parameter for the clock.

As an example we extend our rule in Figure 2, that allows sensors to send data, to only fire when the receiver is active, *and* at least two time units have passed *e.g.* $\mathsf{LC}(n) \geq 2$. As we are modelling an inequality we cannot do this with a single reaction rule and instead we provide a new parameterised rule `sending_data(n)`, shown in Figure 6, that explicitly links the sending sensor to its local clock. The rule is only valid for $n \in \{2,3,\ldots,\mathbb{M}\}$ so this is the family of rules we generate. In this case, we include a clock reset on the right-hand-side of the rule, and clocks may only be reset during the application of a rule.

Note that `clock_advance` rule possesses the same priority as the corresponding rules whose applications are triggered by clock constraints. This allows time to pass until the clock valuation satisfies the first invariant. In a such state, the non-deterministic choices are applicable *i.e.* the system can take a discrete action or permit the time to pass if it is not the last valid clock valuation in which an action must occur.

We must also encode any location-based clock invariants, *e.g.* locations that are only valid for $x \leq 2$. In practice this is used to force a transition to occur rather than allowing an indefinite wait within a particular location, and is almost always *true* (allowing indefinite waits) or a $\leq$ constraint. To encode these invaraints we make use of priority classes in the ABRS. For a state invariant, *e.g.* $t \leq x \leq n$, we add any outgoing transition rules `r(m)` such that $\{$`clock_advance(...)`,`r(t)`,`r(t+1)`,...,`r(n-1)`$\} <$ $\{$`r(n)`$\}$. This means `r(n)` applies *before* any clock updates and *forces* a state transition before the invariant is broken.

4 Examples and Implementation

We implement our approach in BigraphER which also generates the MDPs (that are the semantics for our digital clocks representation). The exported MDPs can be formally checked using standard (probabilstic) model checkers, *e.g.* PRISM, allowing us to benefit from existing model checking algorithms. Modelling a timed system with our approach follows these general steps:

- Define sets of clock valuations/invariants according to the system requirements.
- Define a new clock perspective and add clock entities equal to the number of the timed entities required, and one global clock. Increase the arity of timed entities by 1, and link to the local clock. All clocks are initialised to 0.

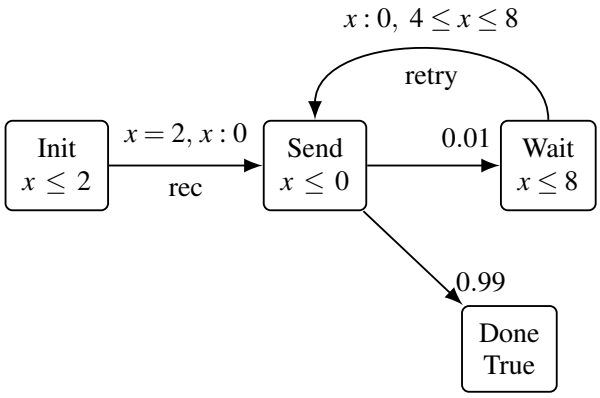

Figure 7: A probabilistic timed automata example model for a simple sending data process from [15].

- For a given $\mathbb{M}$ (max time), and number of timed entities, generate the set of:

$$\texttt{clock\_advance}(\{0,\ldots,\mathbb{M}\},\{0,\ldots,\mathbb{M}\},\ldots)$$

  rules
- Extend system rules by adding in the clock perspective when required to restrict their application based on the time constraints. Parameters for these rules must meet the clock invariants. Any state-based invariants are encoded using priorities.

We illustrate our approach by encoding two different examples: a probabilistic timed automata example, and the timed behaviour of sending a DIS message in Routing Protocol for Low-Power and Lossy Networks (RPL) [1]. The BigraphER models for both examples are in the appendix.

4.1 PTA example

We apply our approach to the probabilistic timed automata example shown in Figure 7 and recreated from [15]. This simple communication system attempts to send data based on the clock X constraints. Here we do not model the actual sends, only the state machine the system goes though. This is somewhat unnatural for bigraphs, where we are more concerned with global system updates (possibly of multiple agents) than encoding a particular state machine for an entity, but is used to illustrate that we can recover existing PTA semantics.

The system starts in the initial state `Init`. When the time has elapsed exactly 2 time units (as forced by the transition constraint *and* the clock invariant on the state) the system moves to the `Send` state. The clock invariant $n \leq 0$ forces data to be sent immediately. With probability 0.99 the system reaches its final state, `Done`, and with 0.01 probability it fails and moves to a `Wait` state. The system waits at least 4 time units and at most 8 time units (forced by the invaraints) before moving back to the `Send` state to retry. The clock is reset with each transition.

We start modelling this PTA example by defining the required time constraints. For each state we define the valid clock valuations as follow. `init_transition` rule (Figure 8a) applies over $n \in \{0,1,2\}$ for `Init` state, and the rules that are associated to `Send` state (Figure 8b and Figure 8c) fire immediately upon receiving data *i.e.* at $n = 0$. While `wait_transition` rule (Figure 8d) is applicable over $n \in \{4,5,6,7,8\}$. For all parameters except 2 for `Init` and 8 for `Wait` states, the rules application should be in the same priority class as `clock_advance(...)` rule allowing the non-deterministic behaviour:

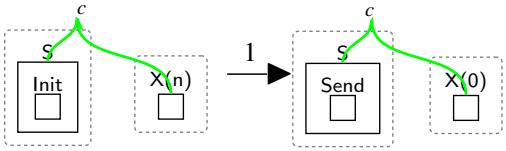

(a) Reaction rule `init_transition(n)`, $n \in \{0,1,2\}$

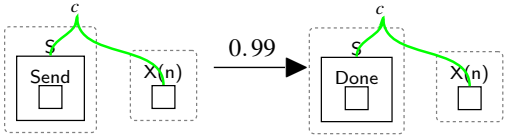

(b) Reaction rule `send_transition_success(n)`, $n = 0$

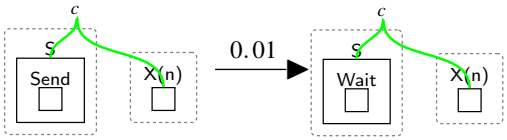

(c) Reaction rule `send_transition_fail(n)`, $n = 0$

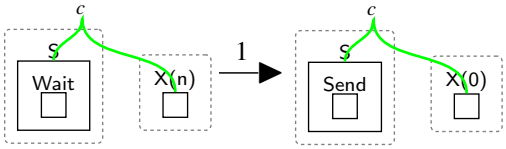

(d) Reaction rule `wait_transition(n)`, $n \in \{4,5,6,7,8\}$

Figure 8: Reaction Rules for the PTA example.

time passes or an data sends. However, as the invariants force leaving if more than $\mathbb{M}$ units elapse, `init_transition(2)` `send_transition(0)`, and `wait_transition(8)` are in a higher priority class. We encode a single rule for each system transition, and show how the probabilities are encoded using rule weights[2], *e.g.* to allow `Send` state to move to either `Wait` or `Done`. In Figure 9, we show the clocks bigraph model conforms to the probabilistic timed automata example by giving the transition system. For space, we only provide the first few transitions. Here the system moves from the initial state `Init` to the `Send` state. When the system in `Init` state and the clock constraints are satisfied, there exist two possibilities: the time passes or an action occurs. Once the clock becomes $X(2)$, the system **has to** move to `Send`.

4.2 Routing Protocol for Low-Power and Lossy Networks (RPL) example

RPL is a distance vector routing protocol that finds the optimal path to send a packet in a wireless sensor network (WSN). We give a brief description of RPL, and the full standard can be found in [1]. RPL uses four control messages to build a route path, called a *a Destination Oriented Directed Acyclic Graph (DODAG)*, from each network node to a predefined gateway (Figure 10).

Briefly, the initial construction of the *DODAG* proceeds as follows

- A node *periodically* sends a DODAG Information Solicitation message (DIS) to find a nearby DODAG.

[2]In this case, as there is only one agent moving state, the weights are equivalent to their probabilities since there will never be additional rule application matches to account for.

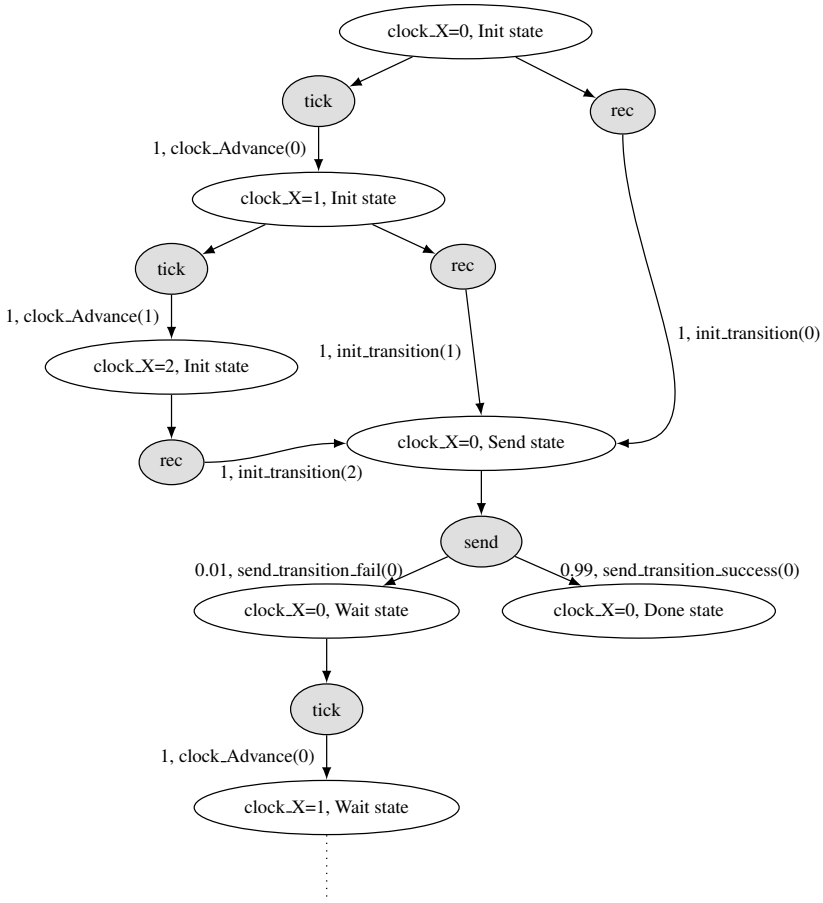

Figure 9: Resulting MDP (partial) for the probabilistic timed automata shown in Figure 7.

- Upon receiving a DIS, a node (or gateway), that is already a part of the DODAG, replies with a DODAG Information Object message (DIO) containing information including the node's rank. A node's rank represents the distance from the gateway and used to avoid/detect loops in a routing path.
- When the DIO is received by the joining node it:
  - Sets the sender of the DIO as a *preferred* parent.
  - Joins the DODAG on a rank larger than its parent's rank, which is given by an *Objective Function*. A simple objective function is simply to add one to the rank.
  - Sends a Destination Advertisement Object message (DAO) to its parent to update the routing tables. For nodes other than the gateway, DAO messages are passed further up the routing tree.
- Eventually, the parent sends a Destination Advertisement Object Acknowledgement message (DAO-ACK) to the joining node to let it know the join was successful.

The newly joined node may now respond to DIS messages allowing more nodes to join the *DODAG*.

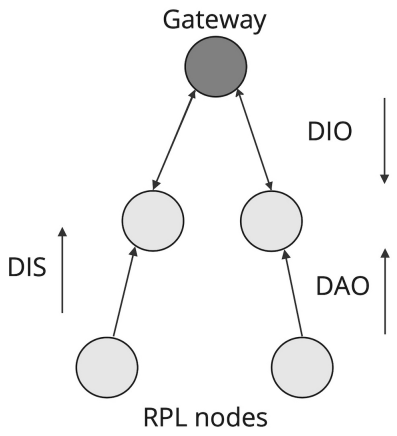

Figure 10: Main RPL control messages. DIS is sent first.

4.2.1 Modelling DIS Messages

We provide a real-time bigraph model for the periodic sending of DIS control messages. Using the same network topology in Figure 1a, our model uses three regions: one is the physical deployment of the sensors PhySensor entity (defining which sensors are in range), the second consists of Sensors modelling the internals of the sensors, *i.e.* what messages they are handling and current states, and the third is the clocks perspective. A physical sensor is connected to the sensor internal information, and to the local clock, using a link.

PhyLink entities represent the physical link status between two sensors. Nested entities, *e.g.* Ready, show the status of the link. In this case Ready indicates that a sensor is allowed to send or receive a DIS message. A DIS entity, representing the DIS control message, is added to sensors that are not part of *DODAG* yet, and Connected entity shows that a sensor is already a part of the DODAG. We abstract unrelated components, including the identifiers, using sites. The *clocks perspective* is as described previously.

We model the DIS sending process to a nearby *DODAG* sensor using two separate rules that are `sendDIS_success` (Figure 11a) and `sendDIS_fail` (Figure 11b). A sensor can send a DIS successfully to any connected sensor in the range when its clock satisfies *e.g.* 1 or 2 time units with weight *e.g.* 0.9; or with weight *e.g.* 0.1 the sending DIS fails. The sensor's clock is reset with each send. Once the gateway or any other joined sensor in the range receives a DIS successfully from a sensor, that sensor stops sending DIS as the receiver will reply with a DIO and the DODAG constructing process starts.

4.3 Model Analysis

To check the feasibility of our proposed modelling technique, we export the bigraph models into MDPs automatically using BigraphER. We can then verify them against required properties using PRISM by expressing properties in the Probabilistic Temporal Logic (PCTL) [10]. To help writing properties, we utilise *bigraph patterns* that allow states to be labelled based on matches [7], *e.g.* label any state that has a Data entity. Since clocks are just part of the model, we can also label clock values, *e.g.* $clock\_LC_0$ for LC(0). Interestingly, the clocks used in the properties are those in the model, *i.e.* they are just bigraph entities, rather than clock variables you would generate if you, for example, expressed the PTA directly in PRISM. This gives flexibility, *i.e.* we can write predicates over many different types of clock matches.

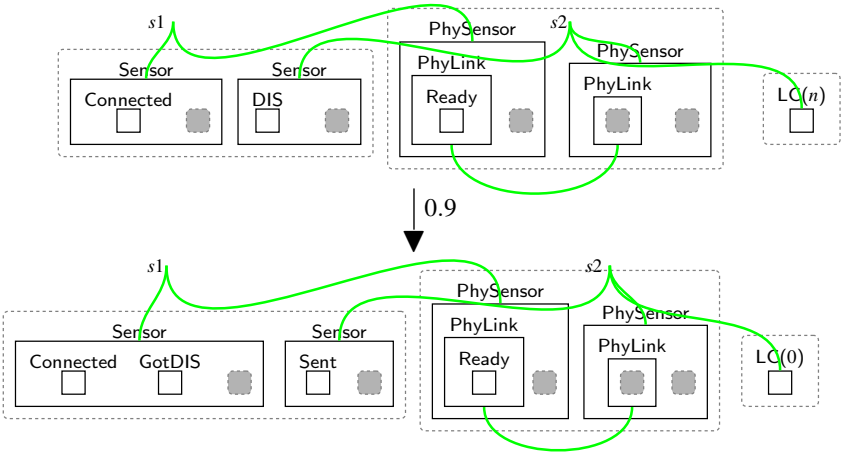

(a) Reaction rule `sendDIS_success(`n`)`, $n \in \{1,2,\ldots,\mathbb{M}\}$

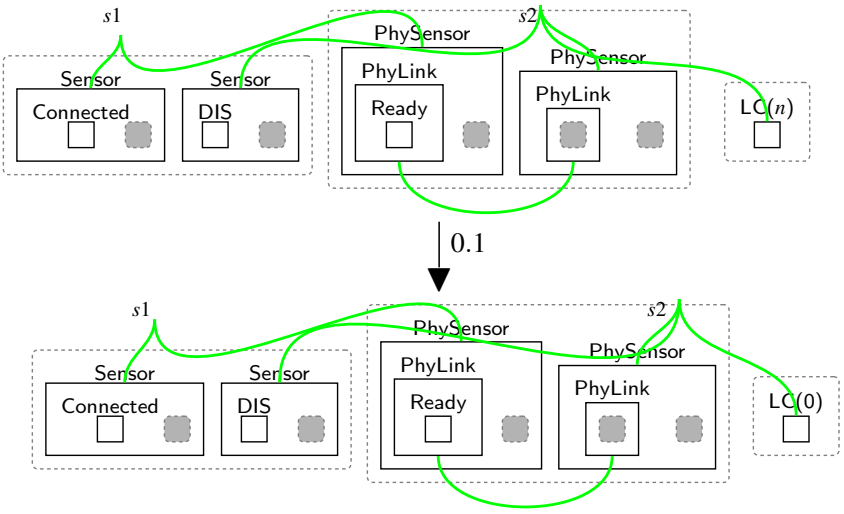

(b) Reaction rule `sendDIS_fail(`n`)`, $n \in \{1,2,\ldots,\mathbb{M}\}$

Figure 11: Sending DIS reaction rules.

We make use of the following PCTL quantifiers to specify our properties: **A** (for all) and **E** (there exists), and path formulae, **F** (eventually), **X** (next), **U** (until) and $\wedge$ (and). We also use probability quantifier **P** to check the probability of reaching a specific state. For example, $\mathbf{P} \leq 0.5\,[\,\mathbf{F}\,stateA\,]$ means: eventually there is at most 0.5 probability that the system reaches a state satisfying *stateA*, we label a state as *stateA* using a bigraph pattern.

We conduct model checking on several interesting properties. For example, given the Section 4.1, we can perform probabilistic reachability properties such as "Is there at least a 0.99 percent probability that the system moves to `Done` state successfully", which holds for this system.

$$\mathbf{P} \geq 0.99\,[\,\mathbf{F}\,(\texttt{In\_Done\_state})\,] \tag{1}$$

We can also utilise clock predicates to check properties relating to time *e.g.* "eventually there is at least a 99% chance that the system moves to the Done state and resets clock X", which is true for our system. Importantly, these are based on our clock entities, not time-bounded properties within PRISM which operate on their own time units.

$$\mathbf{P} \geq 0.99\,[\,\mathbf{F}\,(\,\texttt{In\_Done\_state} \wedge \texttt{clock\_X}_0)\,] \tag{2}$$

We can check clock invariants are satisfied, *e.g.* that the system is not be in Wait state after 8 time units have elapsed:

$$\mathbf{A}\,[\,\neg\,(\mathbf{F}\,(\texttt{In\_Wait\_state}) \wedge (\texttt{clock\_X}_9)\,)\,] \tag{3}$$

As expected this is true in all cases.

Similarly, for RPL we can check that a DIS message is periodically sent *i.e.* at $\texttt{LC}(1), \texttt{LC}(2), \ldots,$ until it is received. That is the property ensures the prohibition of intermediate states so that a time interval must be followed immediately by the transmission of a DIS message.

$$\mathbf{E}\,[\,\mathbf{F}\ \texttt{clock\_LC}(\texttt{n}) \wedge (\mathbf{X}\ \texttt{DIS\_Sent})\ \mathbf{U}\ (\texttt{DIS\_received})\,] \quad \text{where } 1 \leq n \leq \mathbb{M} \tag{4}$$

5 Related work

Many modelling formalisms have been extended to real-time systems by introducing clocks. Timed CCS [26] extends Milner's CCS by introducing the time notion to create concurrency models for real-time systems that are interpreted as non-deterministic behaviour. Timed CCS introduces another variable to record time delays *e.g.* before a message arrives. Similar to our work, timed CCS considers the positive real number including zero as the time domain for convenience, but the model can deal with another numerical domain *e.g.* the natural numbers. Timed Petri nets [28] extends Petri nets to allow time triggered transitions. It uses tokens to allow transitions to start and then they are placed into the output when the firing process terminates. The authors use random variables with continuous or discrete probability distribution functions for the firing time. Timed π-Calculus [19] extends the π-Calculus with continuous time to describe and reason about concurrent Cyber-Physical Systems and real-time systems. An executable operational semantics of π-Calculus is developed in Logic Programming to model concurrency.

Graph Transformation Systems (GTS) has been extended into Probabilistic Timed Graph Transformation Systems (PTGTSs) which enable modelling and analysing structure dynamic and timed and probabilistic behaviour of embedded systems [17]. The clock is a typed node that is contained in a graph to identify the nodes that are used for time measurement only. In this work, PTGTSs is formally mapped into Probabilistic Timed Automata (PTA) where the Probabilistic Timed Structure (PTS) of the mapped PTGTSs is equal to the PTS of the resulting PTA, hence they both satisfy the same set of PTCTL properties. The obtained PTA can be checked using PRISM. However, the mapping process does not consider three aspects of the PTA. First, since the PTA does not consider valuations in the labelling function, constraints are ignored in the mapping process so the constraint should be true for any such atomic proposition. It also considers the PTGTS that does not show timelocks during its execution instead of finding and removing them when mapping PTGTS into PTA. Finally, the GTS state space that is constructed for PTGTS is up to isomorphism as preserved clock nodes are respected which may affect the

state space finiteness of the resulting PTA. In our work, we bound the clock valuations which ensures finite transition systems. Additionally, in case that a system frequently resets its local clocks, the employment of the global clock guarantees the finiteness of the transition system. To prevent timelocks and so obtain a proper transition between reachable states, we assign the maximum states invariant valuation to the rules that are in a higher priority.

Bigraphs have been applied to model some timing aspects, for example, they have been used to model cloud systems that involve time constraints [27]. The cloud model includes aspects such as power consumption, mobile location, and latency. They perform analysis by extending the BigMc model checker to associate the reaction rule with the time cost of tasks at the evaluation step for cloud systems. BigrTimo [25] combines rTiMO process algebra and bigraphs to model the location and connectivity of components of structure-aware mobile systems. BigrTimo uses real-time constraints to control actions by showing the waiting time for communication.

Another work explicitly encodes clocks as entities within bigraphs to model and reason about cloud applications [16], and shares some similarities with our approach. It adds a set of clocks and a set of clocks constraints that are associated with nodes. It then utilises two different types of rules: 1. a set of reaction rules to advance all clocks, all clocks are advanced at the same speed; 2. instantaneous rules[3] that are executed only when there is a match and time constraints of one or many nodes are satisfied. These rules can also update the clock constraints and resets the clocks. When a new time constraint is satisfied, another instantaneous rule may apply subsequently. The work confines the use of clocks to entities where a clock is nested in an entity. This results in encoding multiple reaction rules to advance clocks which may cause unnecessary state-space explosion. It also does not provide a semantic definition that reflects the wall-time. Unlike our approach, it does not consider non-deterministic behaviour that real-time systems often have. The work uses Real-Time Maude language [8] and its TCTL model checker to implement and analyse the approach. In contrast, we encode timed aspects as action bigraphs resulting in an MDP transition system that explicitly models the non-deterministic behaviour of timed systems, that can be formally checked by different model checking tools. State-space explosion is reduced here by employing a strategy that models all clocks as a separate region allowing us to use only one reaction rule to advance all clocks simultaneously. We imitate the wall-time by utilising the global clock entity.

6 Conclusion

Bigraphical Reactive Systems (BRSs) have been successfully used to model a wide range of systems but no explicit support for real-time systems as they do not explicitly support clocks. We overcome this limitation by introducing a modelling technique that uses the digital clocks approximation of (probabilistic) timed automata to encode timing aspects. This approach relies on action bigraphs, which have as a semantics Markov Decision Processes, meaning we can express both probabilistic and timing behaviour within models. Using BigraphER, we encode the proposed strategy and verify the MDP corresponding to each model using PCTL model checking. Using two examples, we show our approach supports multiple clocks and the transition system obtained via rewriting faithfully encodes the digital-clock approximation of the behaviour of a real-time system.

Our approach suffers from the same limitations as the digital-clock approximation *i.e.* currently we do not support diagonal clocks and strict inequalities. We mitigate state space explosion by adopting the following two strategies. First, we bound clock valuations thus we obtain finite transition systems. Second, we allow the base time unit to be specified in each model, effectively allowing clocks to advance

[3] Silent rules that do not appear in an output transition system

by multiple ticks in one step. Another limitation of our approach is that we assume clocks are always synchronised, *i.e.* they all progress at the same speed, which might not always be the case in scenarios like wireless sensor networks and IoT.

In future, we will develop syntactic support for clock constraints in the BigraphER language to generate real-time ABRS models like the ones considered in the paper. This will ensure, for example, that our extended set of reaction rules is correct by construction. We also aim to extend our approach to also support diagonal clocks and strict inequalities.

Acknowledgement

This work is supported by the UK EPSRC projects CHEDDAR (EP/X040518/1) and CHEDDAR Uplift (EP/Y037421/1), and an Amazon Research Award on Automated Reasoning.

References

[1] Roger Alexander et al. *RPL: IPv6 Routing Protocol for Low-Power and Lossy Networks*. RFC 6550. Mar. 2012. DOI: 10.17487/RFC6550.

[2] Rajeev Alur and David L. Dill. "A Theory of Timed Automata". In: *Theor. Comput. Sci.* 126.2 (1994), pp. 183–235. DOI: 10.1016/0304-3975(94)90010-8.

[3] Blair Archibald, Muffy Calder, and Michele Sevegnani. "Conditional Bigraphs". In: *Graph Transformation - 13th International Conference, ICGT 2020, Held as Part of STAF, Norway, June 25-26, 2020, Proceedings*. Ed. by Fabio Gadducci and Timo Kehrer. Vol. 12150. Lecture Notes in Computer Science. Springer, pp. 3–19. DOI: 10.1007/978-3-030-51372-6_1.

[4] Blair Archibald, Muffy Calder, and Michele Sevegnani. "Probabilistic Bigraphs". In: *Form. Asp. Comput.* 34.2 (2022). ISSN: 0934-5043. DOI: 10.1145/3545180.

[5] Blair Archibald et al. "Modelling and verifying BDI agents with bigraphs". In: *Sci. Comput. Program.* 215 (2022), p. 102760. DOI: 10.1016/J.SCICO.2021.102760.

[6] Richard Bellman. "A Markovian decision process". In: *Journal of mathematics and mechanics* (1957), pp. 679–684.

[7] Steve Benford et al. "On Lions, Impala, and Bigraphs: Modelling Interactions in Physical/Virtual Spaces". In: *ACM Trans. Comput. Hum. Interact.* 23.2 (2016), 9:1–9:56. DOI: 10.1145/2882784.

[8] Manuel Clavel et al., eds. *All About Maude - A High-Performance Logical Framework, How to Specify, Program and Verify Systems in Rewriting Logic*. Vol. 4350. Lecture Notes in Computer Science. Springer, 2007. ISBN: 978-3-540-71940-3. DOI: 10.1007/978-3-540-71999-1.

[9] Davide Grohmann and Marino Miculan. "Directed Bigraphs". In: *Proceedings of the 23rd Conference on the Mathematical Foundations of Programming Semantics, MFPS, LA, USA, April 11-14, 2007*. Ed. by Marcelo Fiore. Vol. 173. Electronic Notes in Theoretical Computer Science. Elsevier, pp. 121–137. DOI: 10.1016/J.ENTCS.2007.02.031.

[10] Hans Hansson and Bengt Jonsson. "A Logic for Reasoning about Time and Reliability". In: *Formal Aspects Comput.* 6.5 (1994), pp. 512–535. DOI: 10.1007/BF01211866.

[11] Thomas A. Henzinger, Zohar Manna, and Amir Pnueli. "What Good Are Digital Clocks?" In: *Automata, Languages and Programming, 19th International Colloquium, ICALP92, Vienna, Austria, July 13-17, 1992, Proceedings*. Ed. by Werner Kuich. Vol. 623. Lecture Notes in Computer Science. Springer, pp. 545–558.

[12] Ronald A Howard. “Dynamic programming and markov processes.” In: (1960).

[13] Jean Krivine, Robin Milner, and Angelo Troina. “Stochastic Bigraphs”. In: *Proceedings of the 24th Conference on the Mathematical Foundations of Programming Semantics, MFPS, PA, USA, May 22-25, 2008*. Ed. by Andrej Bauer and Michael W. Mislove. Vol. 218. Electronic Notes in Theoretical Computer Science. Elsevier, pp. 73–96. DOI: 10.1016/J.ENTCS.2008.10.006.

[14] Marta Kwiatkowska, Gethin Norman, and David Parker. “PRISM 4.0: Verification of Probabilistic Real-time Systems”. In: *International conference on computer aided verification*. Springer. 2011, pp. 585–591.

[15] Marta Z. Kwiatkowska et al. “Performance analysis of probabilistic timed automata using digital clocks”. In: *Formal Methods Syst. Des.* 29.1 (2006), pp. 33–78. DOI: 10.1007/S10703-006-0005-2.

[16] Fateh Latreche and Faiza Belala. “Timed CTL checking of time critical cloud applications using timed bigraphs”. In: *Int. J. Crit. Comput. Based Syst.* 9.4 (2019), pp. 379–406. DOI: 10.1504/IJCCBS.2019.106818.

[17] Maria Maximova, Holger Giese, and Christian Krause. “Probabilistic timed graph transformation systems”. In: *J. Log. Algebraic Methods Program.* 101 (2018), pp. 110–131. DOI: 10.1016/J.JLAMP.2018.09.003.

[18] Robin Milner. *The Space and Motion of Communicating Agents*. Cambridge University Press, 2009. ISBN: 978-0-521-73833-0.

[19] Neda Saeedloei and Gopal Gupta. “Timed π-Calculus”. In: *Trustworthy Global Computing - 8th International Symposium, TGC Argentina, August 30-31, 2013, Revised Selected Papers*. Ed. by Martín Abadi and Alberto Lluch-Lafuente. Vol. 8358. Lecture Notes in Computer Science. Springer, pp. 119–135. DOI: 10.1007/978-3-319-05119-2_8.

[20] Hamza Sahli, Faiza Belala, and Chafia Bouanaka. “A BRS-Based Approach to Model and Verify Cloud Systems Elasticity”. In: *1st International Conference on Cloud Forward: From Distributed to Complete Computing, October 6-8, 2015, Italy*. Ed. by Keith G. Jeffery and Dimosthenis Kyriazis. Vol. 68. Procedia Computer Science. Elsevier, pp. 29–41. DOI: 10.1016/J.PROCS.2015.09.221.

[21] Hamza Sahli, Thomas Ledoux, and Éric Rutten. “Modeling Self-adaptive Fog Systems Using Bigraphs”. In: *Software Engineering and Formal Methods - SEFM 2019 Collocated Workshops: CoSim-CPS, ASYDE, CIFMA, and FOCLASA, Norway, September 16-20, Revised Selected Papers*. Ed. by Javier Cámara and Martin Steffen. Vol. 12226. Lecture Notes in Computer Science. Springer, pp. 252–268. DOI: 10.1007/978-3-030-57506-9_19.

[22] Michele Sevegnani and Muffy Calder. “BigraphER: Rewriting and Analysis Engine for Bigraphs”. In: *Computer Aided Verification - 28th International Conference, CAV, ON, Canada, July 17-23, 2016, Proceedings, Part II*. Ed. by Swarat Chaudhuri and Azadeh Farzan. Vol. 9780. Lecture Notes in Computer Science. Springer, pp. 494–501. DOI: 10.1007/978-3-319-41540-6_27.

[23] Michele Sevegnani and Muffy Calder. “Bigraphs with sharing”. In: *Theor. Comput. Sci.* 577 (2015), pp. 43–73. DOI: 10.1016/J.TCS.2015.02.011.

[24] Michele Sevegnani et al. “Modelling and Verification of Large-Scale Sensor Network Infrastructures”. In: *23rd International Conference on Engineering of Complex Computer Systems, ICECCS, Australia, December 12-14, 2018*. IEEE Computer Society, pp. 71–81. DOI: 10.1109/ICECCS2018.2018.00016.

[25] Wanling Xie, Huibiao Zhu, and Qiwen Xu. "BigrTiMo-A Process Algebra for Structure-Aware Mobile Systems". In: *22nd International Conference on Engineering of Complex Computer Systems, ICECCS, Fukuoka, Japan, November 5-8, 2017*. IEEE Computer Society, pp. 50–59. DOI: 10.1109/ICECCS.2017.13.

[26] Wang Yi. "CCS + Time = An Interleaving Model for Real Time Systems". In: *Automata, Languages and Programming, 18th International Colloquium, ICALP91, Spain, July 8-12, 1991, Proceedings*. Ed. by Javier Leach Albert, Burkhard Monien, and Mario Rodríguez-Artalejo. Vol. 510. Lecture Notes in Computer Science. Springer, pp. 217–228. DOI: 10.1007/3-540-54233-7_136.

[27] Lian Yu et al. "Modeling and Analysis of Mobile Cloud Computing Based on Bigraph Theory". In: *2nd IEEE International Conference on Mobile Cloud Computing, Services, and Engineering, MobileCloud United Kingdom, April 8-11, 2014*. IEEE Computer Society, pp. 67–76. DOI: 10.1109/MOBILECLOUD.2014.11.

[28] Wlodek M Zuberek. "Timed Petri nets definitions, properties, and applications". In: *Microelectronics Reliability* 31.4 (1991), pp. 627–644.

Appendix A BigraphER Implementation Of Section 4 Examples

```
################  PTA Example #################

atomic fun ctrl X(n) = 1 ;
ctrl S = 1;
atomic ctrl Init = 0;
atomic ctrl Send = 0;
atomic ctrl Wait = 0;
atomic ctrl Done = 0;

################  Reaction Rules #################

##Init
fun react init_transition(n) =
  S{c}.Init || X(n){c}
  -[1]->
  S{c}.Send || X(0){c};

## Send
fun react send_transition_success(n) =
  S{c}.Send || X(n){c}
  -[0.99]->
  S{c}.Done || X(n){c};

fun react send_transition_fail(n) =
  S{c}.Send || X(n){c}
  -[0.01]->
  S{c}.Wait || X(n){c};


# Wait
fun react wait_transition(n) =
  S{c}.Wait || X(n){c}
  -[1]->
  S{c}.Send || X(0){c};

#Done
react done_done =
  S{c}.Done -[1]-> S{c}.Done;

################ Clock Advance Rule #################
fun react clock_advance(n) =
  X(n){c}
  -[1]->
  X(n + 1){c};

################ Predicates #################
fun big clock_X(n) = X(n){c};
big in_Init_state = S{c}.Init;
big in_Send_state = S{c}.Send;
big in_Wait_state = S{c}.Wait;
big in_Done_state = S{c}.Done;

################ Initial State #################
big example_PTA = /c (S{c}.Init || X(0){c});

 begin abrs
  int n = {0,1,2,3,4,5,6,7,8};
  int maxInitT = {2};
  int init_Sending_Time = {0,1};
  int maxSendT = 0;
  int maxWaitT = 8;
  int wait_Sending_Time={4,5,6,7};
```

```
63
64   init example_PTA;
65
66   rules = [
67     # Higher in the list => higher priority
68     {done_done,
69      init_transition(maxInitT),
70      send_transition_fail(maxSendT),
71      send_transition_success(maxSendT),
72      wait_transition(maxWaitT)},
73
74     {clock_advance(n),
75      wait_transition(wait_Sending_Time),
76      init_transition(init_Sending_Time)}
77   ];
78
79   actions = [
80     send={send_transition_success,send_transition_fail},
81     retry={wait_transition},
82     rec={init_transition},
83     deadlock={done_done},
84     tick={clock_advance}
85   ];
86
87   preds = {
88     in_Init_state,
89     in_Send_state,
90     in_Wait_state,
91     in_Done_state,
92     clock_X(n)
93   };
94 end
```

```
################  DIS_RPL Example #################
ctrl Physical = 0;
ctrl PhySensor = 1;
ctrl PhyLink = 1;
ctrl Sensors = 0;
ctrl Sensor = 1;
atomic fun ctrl ID(i) = 0;
atomic ctrl Ready = 0;
atomic ctrl Connected = 0;
ctrl LocalClocks = 0;
atomic fun ctrl LC(n1) = 1;
atomic fun ctrl GC(n) = 0;
atomic ctrl Sent=0;
atomic ctrl DIS=0;
atomic ctrl GotDIS=0;

###################################################################

fun react clock_advance( n1,n2, n) = LocalClocks.(  LC(n1){s1} | LC(n2){s2}  ) | GC(n)
     -[1]->
           LocalClocks.(LC(n1 + 1){s1} | LC(n2 + 1 ){s2}  ) | GC(n + 1)  ;

###################################################################

fun react sensor_sendDIS_Success(n1) =

  ( Sensor{s1}.(Connected | id) | Sensor{s2}.(DIS | id))

  || /x ( PhySensor{s1}.(PhyLink{x}.Ready | id) | PhySensor{s2}.(PhyLink{x}.id | id) )

  || ( LC(n1){s2} )

  -[0.9]->

 ( Sensor{s1}.(Connected | GotDIS | id) | Sensor{s2}.( Sent | id))

  || /x ( PhySensor{s1}.(PhyLink{x}.Ready | id) | PhySensor{s2}.(PhyLink{x}.id | id) )

  || ( LC(0){s2} );


fun react sensor_sendDIS_Fail(n1) =

  ( Sensor{s1}.(Connected | id) | Sensor{s2}.(DIS | id))

  || /x ( PhySensor{s1}.(PhyLink{x}.Ready | id) | PhySensor{s2}.(PhyLink{x}.id | id) )

  || ( LC(n1){s2} )

  -[0.1]->

  ( Sensor{s1}.(Connected | id) | Sensor{s2}.(DIS | id))

  || /x ( PhySensor{s1}.(PhyLink{x}.Ready | id) | PhySensor{s2}.(PhyLink{x}.id | id) )

  || ( LC(0){s2} );

###################################################################

#Initial State

big example_RPL =
     /s1/s2/s3/l1/l2/l3
  (
```

```
    Physical.(
      PhySensor{s1}.(PhyLink{l1}.Ready | PhyLink{l2}.Ready)
    | PhySensor{s2}.(PhyLink{l1}.Ready | PhyLink{l3}.Ready)
    | PhySensor{s3}.(PhyLink {l2}.Ready | PhyLink{l3}.Ready )
  )
  ||
    Sensors.(
      Sensor{s1}.( ID("A") | Connected )
    | Sensor{s2}.( ID("B") | DIS )
    | Sensor{s3}.( ID("C") | DIS )
  )
 ||
  ( GC(0) | LocalClocks.( LC(0){s2} | LC(0){s3}  ) )
 );

#####################################################################


fun big sensor_send_DIS(i) = Sensor{s2}.( ID(i) | Sent | id);

fun big sensor_received_DIS(i) = Sensor{s2}.( ID(i) | GotDIS | id);

fun big clockB_LC(forPredicate) = Sensor{s2}.(ID("B") | id )
                               || LC(forPredicate){s2} ;

fun big clockC_LC(forPredicate) = Sensor{s3}.(ID("C") | id )
                               || LC(forPredicate){s3};


begin abrs
  int n1={0,1};
  int n2={0,1};
  int n={0,1};
  int forPredicate={0,1,2};
  int sensor_Sending_Time={1,2};

  string i = {"A", "B", "C"};

#################################################################

init example_RPL;

rules = [
         { clock_advance(n1,n2, n),
           sensor_sendDIS_Success(sensor_Sending_Time),
           sensor_sendDIS_Fail(sensor_Sending_Time)}
        ];

actions=[
         tick={clock_advance},
         sendDIS = {sensor_sendDIS_Success, sensor_sendDIS_Fail}
    ];
preds = { clockB_LC(forPredicate),
          clockC_LC(forPredicate),
          sensor_send_DIS(i),
          sensor_received_DIS(i)};
  end
```