



Checking conformance for time-constrained scenario-based specifications [☆]



S. Akshay ^{a,*}, Paul Gastin ^b, Madhavan Mukund ^c, K. Narayan Kumar ^c

^a Indian Institute of Technology Bombay, India

^b LSV, ENS Cachan, INRIA, CNRS, France

^c Chennai Mathematical Institute, Chennai, India

ARTICLE INFO

Article history:

Received 4 March 2014

Received in revised form 1 February 2015

Accepted 18 March 2015

Available online 24 March 2015

Communicated by P. Aziz Abdulla

Keywords:

MSC graphs

Timed automata

Model checking

ABSTRACT

We consider the problem of model checking message-passing systems with real-time requirements. As behavioral specifications, we use message sequence charts (MSCs) annotated with timing constraints. Our system model is a network of communicating finite state machines with local clocks, whose global behavior can be regarded as a timed automaton. Our goal is to verify that all timed behaviors exhibited by the system conform to the timing constraints imposed by the specification. In general, this corresponds to checking inclusion for timed languages, which is an undecidable problem even for timed regular languages. However, we show that we can translate regular collections of time-constrained MSCs into a special class of event-clock automata that can be determined and complemented, thus permitting an algorithmic solution to the model checking/conformance problem.

© 2015 Elsevier B.V. All rights reserved.

1. Introduction

In a distributed system, several agents interact to generate a global behavior. This interaction is usually specified in terms of scenarios, using message sequence charts (MSCs) [23]. Protocol specifications typically include timing requirements for messages and descriptions of how to recover from timeouts, so a natural and useful extension to MSCs is to add timing constraints between pairs of events, yielding time-constrained MSCs (TCMSCs) [8,1].

Infinite collections of MSCs are typically described using message sequence graphs (MSGs) [23,9]. An MSG, a finite directed graph with nodes labeled by MSCs, is the most basic form of a High-level Message Sequence Chart (HMSC) [26]. Any path through the graph generates a new MSC by concatenating the MSCs seen along the path. Thus, the set of all paths through an MSG generates a possibly infinite collection of MSCs. In this article, we generalize MSGs to time-constrained MSGs (TCMSGs), where nodes are labeled by TCMSCs and edges may have additional time constraints between nodes. Thus, TCMSGs generate infinite collections of time-constrained scenarios, i.e., TCMSCs. This forms our basic model of specification.

A natural system model in this setting is a timed message-passing automaton (timed MPA), a set of communicating finite-state machines equipped with clocks that are used to guard transitions, as in timed automata [11]. Just as timed words are used to describe the runs of timed automata, the interactions exhibited by timed MPAs can be described using

[☆] Supported by CNRS, LIA InForMel and DST-INSPIRE faculty award [IFA12-MA-17].

* Correspondence to: Dept. of Computer Science and Engineering, IIT Bombay, Powai, Mumbai, 400076, India. Tel.: +91 2225767711.

E-mail addresses: akshayss@cse.iitb.ac.in (S. Akshay), Paul.Gastin@lsv.ens-cachan.fr (P. Gastin), madhavan@cmi.ac.in (M. Mukund), kumar@cmi.ac.in (K. Narayan Kumar).

timed MSCs—MSCs in which each event is assigned an explicit timestamp. However, the global state space of a timed MPA in fact defines a timed automaton over a distributed alphabet and in this paper we focus on this simplified global view of timed message-passing systems, though our results go through smoothly for the distributed system model as well. Thus, our main interest in this paper lies in considering a distributed specification (formalized using TCMSGs) and comparing it against a global timed implementation.

Our aim is to check if all timed MSCs accepted by a timed MPA conform to the time constraints given by a TCMSG specification. This problem can naturally be seen as comprising of two parts. The first asks if for a given timed MPA A and TCMSG G , every timed execution exhibited by A is in the specification. Indeed, this is the standard model-checking question for timed MPAs. The second part, the coverage problem, asks if every TCMSG generated by a given TCMSG can be witnessed by some timed execution of the TMAPA. To make the problem tractable, we focus on *locally synchronized* TCMSGs—those for which the underlying behavior is guaranteed to be regular [22].

In general, the model checking problem above corresponds to checking inclusion for timed languages, which is known to be undecidable even for timed regular languages [6]. Fortunately, it turns out that timing constraints in a TCMSG correspond to a very restricted use of clocks. This allows us to associate with each TCMSG an extended event clock automaton that accepts all timed executions that are consistent with the timing constraints of the TCMSG. We prove that these extended event clock automata can be determinized and complemented (as in the case of ordinary event clock automata [7]), yielding an algorithmic solution to our model checking problem.

Turning to the coverage problem, we observe this cannot be directly reduced to a timed inclusion problem. The timed inclusion problem in this direction would ask if there is a witnessing execution of the timed MPA for every timed linearization of a TCMSG generated by the TCMSG. But an implementation (timed MPA) having strictly better time bounds than the specification might have a witnessing execution for every TCMSG generated by the TCMSG, even if it does not satisfy every timed linearization of the TCMSG. Such an implementation should be considered as a valid one and this is precisely what our definition of the coverage problem achieves. For solving this problem, we need an additional assumption on the specification. We assume that the locally synchronized TCMSG has a special form that every process on the TCMSG labeling any node has some event. Now, we use the same extended event clock automaton as above accepting all timed executions that are consistent with the TCMSG. Then, using a product construction, we can recover the set of paths of the TCMSG which have some valid execution in the timed MPA, thus solving the coverage problem.

Related work. We have used TCMSGs as the basic model for specifying high level distributed and timed systems. However, there are other formalisms which also tackle time and concurrency issues in systems. In Petri nets [30] *tokens* are positioned in *places* and a transition fires by consuming tokens and creates new ones, in general in other places. Thus, transitions that consume different tokens, can fire independently. Many timed extensions of Petri nets have been considered, for instance, time Petri nets [12], timed Petri nets [29]. Unfoldings of Petri nets provide a way to model the partial order behavior of these systems and by lifting these unfoldings to the timed extensions, they provide a timed partial order semantics [17]. For more discussion on this refer to [16]. However, these unfoldings are seldom graphically representable in a compact manner unlike MSCs (and their timed extensions). Further, unfoldings in Petri nets correspond to “branching time” whereas MSCs express “linear time” behavior.

Other models dealing with time and concurrency include networks of timed automata [6] and products of timed automata [20]. Again in [13], unfolding techniques were applied to study such networks of timed automata. However, these models do not allow communication via explicit message passing which is one of the main features of the timed MPA and TCMSGs that we have introduced.

The formal semantics and analysis of timing in MSCs has been addressed earlier in [8,10,15,24]. In [8] and [10], only single timed MSCs or high-level timed MSCs were considered, while in [24] one of the first models of timed MPAs was introduced. However, the latter do not consider MSCs as a semantics of their automata but rather look at restricted channel architectures (e.g., one-channel systems) to transfer decidability of reachability problems from the untimed to the timed setting. The automaton model in [15] links the two approaches by considering a similar automaton model with semantics in terms of timed MSCs. But they tackle only a specific matching problem for which they propose a practical solution using the tool UPPAAL. More recently, in [4] the authors have considered TCMSGs under restrictions that are weaker than being locally-synchronized. Though this allows modeling more general non-regular languages of TCMSGs, they only tackle the emptiness problem and do not address more complicated issues of consistency or conformance as we do.

In [19], the authors develop a specification theory that combines notions of specifications and implementations and provides constructs for checking consistency etc., in the setting of sequential real-timed systems. However, they define the implementation as another specification and relate the two using a notion of refinement defined as an alternating (timed) simulation relation. In our setting, the implementation and specification are different objects to begin with and we relate them at the level of behaviors rather than systems. Thus, checking consistency corresponds to checking inclusion of timed behaviors which is often a harder problem than defining a simulation. In addition, we consider timed and distributed systems, where concurrency plays a major role and gives rise to several additional challenges.

Preliminary versions of some of the results were presented as extended abstracts in [3,5]. Here, we establish a generic framework that combines those results as well as completes and generalizes the proofs and techniques.

Structure of the paper. The paper is organized as follows. We begin with some preliminaries where we introduce (timed) MSCs, MSGs and the timed automata formalisms. In the subsequent section, we discuss the conformance problem in detail.




Fig. 1. Different views of a system with two users and a server.

In Section 4 we introduce MSC event clock automata and show that they can be determinized and complemented. The next section has the first main technical result: translating locally synchronized TCMSCs to finite state MSC event clock automata, which yields a solution to the model-checking problem in Section 6. In Section 7, we provide a partial solution to the reverse problem of checking coverage and finally conclude in Section 8 with a short discussion.

2. Preliminaries

Message sequence charts. A message sequence chart (MSC) describes the messages exchanged between a set $Proc$ of processes in a distributed system. The first diagram in Fig. 1 is an MSC involving two users and a server. Each process evolves vertically along a lifeline. Messages are shown by arrows between the lifelines of the sender and receiver.

Each message consists of two events, send and receive, and is labeled using a finite set of message labels \mathcal{M} . For instance, the events u_1 and a_1 are the send and receive events of a message labeled req from process p (User1) to process q (Server). Each (ordered) pair of processes p and q is connected by a dedicated fifo channel (p, q) —for example, in Fig. 1, the messages sent at s_1 and s_2 are on channel (r, q) and the second message cannot be received before the first one. Note that the channels (p, q) and (q, p) are distinct under this definition.

Since processes are locally sequential, the set of events E_p along a process p is linearly ordered by a relation denoted \leq_{pp} . In addition, for each message sent along a channel (p, q) , the send and receive events of the message are related by an ordering relation \leq_{pq} . Thus, for example, $a_1 \leq_{qq} a_5$ and $a_3 \leq_{qp} u_2$. Together, the local linear orders \leq_{pp} and the message orders \leq_{pq} generate a partial order \leq over the set of events—for instance, $u_3 \leq s_3$.

Finally, we label each event using a finite alphabet Act of communication actions. We write $p!q(m)$ to denote the action where p sends message m to q and $p?q(m)$ to denote the action where p receives message m from q . We abbreviate by $p!q$ and $p?q$ the set of all actions of the form $p!q(m)$ and $p?q(m)$, respectively, over all possible choices of m .

Overall, an MSC can then be captured as a labeled partial order $M = (E, \leq, \lambda)$ where $\lambda : E \rightarrow Act$ associates each event with its corresponding action. A cut is a subset of events that is downward closed: $c \subseteq E$ is a cut if $\downarrow c = c$, where $\downarrow c = \{e \in E \mid \exists e' \in c. e \leq e'\}$.

Like any partial order, an MSC can be reconstructed up to isomorphism from its linearizations, i.e., words over Act that extend \leq . In fact, the fifo condition on channels ensures that a single linearization suffices to reconstruct an MSC. In this way, an MSC M corresponds to a set $\text{lin}(M)$ of words over Act and a set of MSCs defines the word language $\bigcup_{M \in \mathcal{L}} \text{lin}(M)$. We say that a set \mathcal{L} of MSCs is regular if its associated word language is regular.

Time-constrained message sequence charts. A time-constrained MSC (TCMSC) is an MSC annotated with time intervals between pairs of events. We restrict timing constraints to pairs of distinct events on the same process and to the matching send and receive events across messages. Intervals have rational endpoints and may be open or closed at either end.

For example, in the second diagram in Fig. 1, the constraint $[0, 3]$ between a_3 and a_4 bounds the time that the Server waits for a User to confirm a grant. On the other hand, the constraint $[0, 1]$ between a_3 and u_2 bounds the time taken to deliver this particular message.

A TCMSC over Act is a pair $\mathfrak{M} = (M, \tau)$, where $M = (E, \leq, \lambda)$ is an MSC over Act and τ is a partial map from $E \times E$ to the set of intervals such that $(e, e') \in \text{dom}(\tau)$ implies that $e \neq e'$ and either $e \leq_{pp} e'$ or $e \leq_{pq} e'$ for some processes p and q .

Timed message sequence charts. A timed MSC (TMSC) describes a concrete timed behavior in the MSC setting. In a TMSC, we assign events timestamps that are consistent with the underlying partial order. Thus, a TMSC over Act is a pair $T = (M, t)$ where $M = (E, \leq, \lambda)$ is an MSC over Act and $t : E \rightarrow \mathbb{R}_{\geq 0}$ is a function such that if $e \leq e'$ then $t(e) \leq t(e')$ for all $e, e' \in E$.

For instance, consider the TMSC in the third diagram of Fig. 1. The message sent at a_3 is received instantaneously while the message sent at s_2 is received 3 time units later.

A timed word over Act is a sequence $(a_1, t_1)(a_2, t_2) \dots (a_n, t_n)$ where $a_1 a_2 \dots a_n$ is a word over Act and $t_1 \leq t_2 \leq \dots \leq t_n$ is a nondecreasing sequence over $\mathbb{R}_{\geq 0}$. The set of timed words over Act is denoted TW_{Act} . A timed linearization of a TMSC is thus a timed word in TW_{Act} . We let $t\text{-lin}(T)$ denote the set of timed linearizations of TMSC T . A single TMSC may admit more than one timed linearization if concurrent events on different processes have the same timestamp. As for untimed MSCs, under the fifo assumption for channels, a timed MSC can be reconstructed from any one of its timed linearizations.




Fig. 2. A TCMSC, with a TCMSC and a TMSMC that it generates.

With this definition, TCMSCs can be considered as abstractions of TMSCs and timed words. For instance, we will say that the TMSMC in Fig. 1 realizes the TCMSC in the same figure since each interval constraint between events in the TCMSC is satisfied by the time-stamps of the corresponding events in the TMSMC. In this way, a TCMSC \mathfrak{M} defines a family of TMSCs—the set of all TMSCs that realize \mathfrak{M} , which we denote $\mathcal{L}_{time}(\mathfrak{M})$. We also consider the set $\mathcal{L}_{tw}(\mathfrak{M}) = \bigcup_{T \in \mathcal{L}_{time}(\mathfrak{M})} \text{t-lin}(T)$ of timed words that realize \mathfrak{M} .

Message sequence graphs. A message sequence graph (MSG) is a directed graph in which nodes are labeled by MSCs. We begin with a graph $G = (V, \rightarrow, v_{in}, V_F)$ with nodes V , initial node $v_{in} \in V$, final nodes $V_F \subseteq V$ and edge relation \rightarrow . An MSG is a structure $\mathfrak{G} = (G, \mathcal{L}^M, \Phi)$ where \mathcal{L}^M is a set of MSCs and $\Phi: V \rightarrow \mathcal{L}^M$ associates an MSC with each node. A path in G is a sequence of nodes $v_0 v_1 \dots v_n$ where each adjacent pair of states is related by \rightarrow . An accepting path is one that starts in v_{in} and ends in some node of V_F .

A path $\pi = v_0 v_1 \dots v_n$ in G defines an MSC $\Phi(v_0 v_1 \dots v_n) = \Phi(v_0) \circ \Phi(v_1) \circ \dots \circ \Phi(v_n)$, where \circ denotes MSC concatenation. When we concatenate two MSCs $M_1 = (E^1, \leq^1, \lambda_1)$ and $M_2 = (E^2, \leq^2, \lambda_2)$ we attach the lifelines in M_2 below those of M_1 to obtain an MSC $M_1 \circ M_2 = (E^1 \cup E^2, \leq, \lambda)$ where λ combines λ_1 and λ_2 and \leq is generated by $\leq^1 \cup \leq^2 \cup \{(e_1, e_2) \mid \exists p. e_1 \in E_p^1, e_2 \in E_p^2\}$.

More formally, for a path $\pi = v_0 v_1 \dots v_n$ we define the MSC $\Phi(\pi)$ as follows. First, we use \leq to denote the prefix relation and write $\rho \leq \pi$ to denote that path ρ is a prefix of path π . Now, for each vertex v , let $\Phi(v)$ be the MSC $M_v = (E^v, \leq^v, \lambda_v)$. We assume that the events are disjoint across the MSCs M_v . We then define $\Phi(\pi) = M_\pi = (E^\pi, \leq^\pi, \lambda_\pi)$, where,

- $E^\pi = \bigcup_{\rho v \leq \pi} E_v \times \{\rho v\}$
- For each $\rho v \leq \pi$, $\lambda_\pi((e, \rho v)) = \lambda_v(e)$.
- \leq^π is defined as the reflexive transitive closure of $\bigcup_{p, q \in \text{Proc}} <_{pq}^\pi$, where
 - $(e, \rho v) <_{pp}^\pi (e', \rho' v')$ for some $p \in \text{Proc}$ if $e \in E_p^v$, $e' \in E_p^{v'}$ and either $\rho v \leq \rho' v'$ or $(\rho v = \rho' v'$ and $e <_{pp}^v e')$.
 - $(e, \rho v) <_{pq}^\pi (e', \rho' v')$ for some processes $p \neq q$, if $\rho v = \rho' v'$ and $e <_{pq}^v e'$.

We associate with an MSG \mathfrak{G} a language of MSCs $\mathcal{L}(\mathfrak{G}) = \{\Phi(\pi) \mid \pi \text{ is an accepting path in } G\}$. In general, it is undecidable to determine whether $\mathcal{L}(\mathfrak{G})$ is regular [22]. This is because processes move asynchronously along the MSC traced out by accepting paths and there is no bound, in general on this asynchrony. However, there is a sufficient structural condition to guarantee regularity [9,27].

Given an MSC M , we construct its communication graph $CG(M)$ as follows: the vertices are the processes and we have a directed edge (p, q) if M contains a message from p to q . An MSC M is said to be *connected* if the non-isolated vertices in $CG(M)$ form a single strongly connected component. An MSG \mathfrak{G} is said to be *locally synchronized* if for every loop π in \mathfrak{G} , the MSC $\Phi(\pi)$ is connected. Intuitively, this means that every message sent in a loop is implicitly acknowledged, because if p sends a message, there is a path in the communication graph back to p . This ensures that all channels are *universally bounded*—there is a uniform bound B such that across all linearizations, no channel ever has more than B pending messages. Thus, if \mathfrak{G} is locally synchronized, $\mathcal{L}(\mathfrak{G})$ is a regular set of MSCs.

Time-constrained message sequence graphs. We generalize MSGs to the timed setting in a natural way. In a *time-constrained MSG (TCMSG)*, states are labeled by TCMSCs rather than MSCs. In addition, we also permit process-wise timing constraints along the edges of the graph. A constraint for process p along an edge $v \rightarrow v'$ specifies a constraint between the final p -event of $\Phi(v)$ and the initial p -event of $\Phi(v')$, provided p actively participates in both these nodes. If p does not participate in either of these nodes, the constraint is ignored. Formally, a TCMSC is a tuple $\mathfrak{G} = (G, \mathcal{L}^{TC}, \Phi, \text{EdgeC})$ where $G = (V, \rightarrow, v_{in}, V_F)$ is a graph as before, $\Phi: V \rightarrow \mathcal{L}^{TC}$ labels each node with a TCMSC from a set \mathcal{L}^{TC} and EdgeC associates a tuple of constraints with each edge—for convenience, we assume that any edge constraint not explicitly specified corresponds to the trivial constraint $(-\infty, \infty)$.

Each accepting path in a TCMMSG defines a TCMSC. Given a path $v_0v_1 \cdots v_n$, we concatenate the TCMSCs $\Phi(v_0), \Phi(v_1), \dots, \Phi(v_n)$ and insert the additional constraints specified by *EdgeC*. We define $\mathcal{L}_{TC}(\mathfrak{G})$ to be the set of all TCMSCs over *Act* generated by accepting paths in G . We also let $\mathcal{L}_{time}(\mathfrak{G}) = \bigcup_{\mathfrak{M} \in \mathcal{L}_{TC}(\mathfrak{G})} \mathcal{L}_{time}(\mathfrak{M})$ and $\mathcal{L}_{tw}(\mathfrak{G}) = \bigcup_{\mathfrak{M} \in \mathcal{L}_{TC}(\mathfrak{G})} \mathcal{L}_{tw}(\mathfrak{M})$. Fig. 2 shows a TCMMSG, a TCMSC that it generates and a realizing TMSC.

Just as for MSGs, a TCMMSG \mathfrak{G} is said to be *locally synchronized* if for every loop π in \mathfrak{G} , the MSC $\Phi(\pi)$ is connected.

Timed automata. We can formulate many types of machine models for timed MSCs. One natural choice is a message-passing automaton (MPA) equipped with (local) clocks. In a timed MPA, we have one component for each process p , which is a finite state automaton over actions of the form $p!q(m)$ and $p?q(m)$. Each component also has local clocks that can be used to guard transitions. The global state space defines a timed automaton over *Act*.

A timed automaton over an alphabet Σ is a tuple $\mathcal{A} = (Q, \Delta, q_{in}, F, Z)$ where Q is a finite set of states, $q_{in} \in Q$ is the initial state, $F \subseteq Q$ are the final states and Z is a set of clocks that take values over $\mathbb{R}_{\geq 0}$. Each transition in Δ is of the form $q \xrightarrow{\varphi, a, X} q'$ where $q, q' \in Q$, $a \in \Sigma$, $X \subseteq Z$ and φ is a boolean combination of clock constraints of the form $x \text{ op } c$ where $x \in Z$, $c \in \mathbb{Q}_{\geq 0}$ and $\text{op} \in \{\leq, <, >, \geq\}$. This transition is enabled if the current values of all clocks satisfy the guard φ . On taking this transition, the clocks in X are reset to 0. As is standard, time elapses between transitions, transitions occur instantaneously and such an automaton accepts timed words from TW_{Σ} . More details can be found in [6,11].

For our purposes, we only need the following results about timed automata.

- Given timed automata \mathcal{A}_1 and \mathcal{A}_2 , we can construct a timed automaton \mathcal{A}_{12} such that $L(\mathcal{A}_{12}) = L(\mathcal{A}_1) \cap L(\mathcal{A}_2)$.
- Checking if the language of a timed automaton is empty is decidable.

A *timed MPA* is defined as a timed automata over *Act* whose languages can be interpreted as (timed) linearizations of timed MSCs. A timed word in TW_{Act} corresponds to a linearization of a timed MSC provided the timed word is well-formed and complete. A word w over *Act* is *well-formed* if for each channel (p, q) , in every prefix v of w , the sequence of messages received by q from p in v is a prefix of the messages sent from p to q in v . A well-formed word w is *complete* if $\#_{p!q}(w) = \#_{q?p}(w)$ for each matching pair of send–receive actions, where $\#_X(u)$ counts the number of occurrences in u of $X \subseteq Act$. Finally, a well-formed word w is *B-bounded* if, in every prefix v of w , $\#_{p!q}(v) - \#_{q?p}(v) \leq B$ for each channel (p, q) . Correspondingly, a timed word is said to be well-formed (complete, B-bounded) if its projection onto *Act* is well-formed (complete, B-bounded). Well-formedness captures the intuition that any receive action has an earlier matching sending action. Completeness guarantees that all pending messages have been received. B-boundedness promises that no channel ever has more than B messages.

Note that we could have alternatively defined timed MPA over a distributed state space and their semantics directly over timed MSCs, as done in [2,1] (and originally in [25] in the untimed setting), instead of defining it via timed linearizations. However, to preserve notational clarity and for ease of presentation, we prefer to adopt the global state space approach in this paper.

3. The problem statement

In this paper, we are interested in comparing the distributed timed behavior of a TCMMSG with the global timed behavior of a timed MPA. As detailed in the previous section, the behavior of a timed MPA can be described as a timed automaton over the global alphabet of actions *Act*. Thus, given a TCMMSG \mathfrak{G} and a timed automaton \mathcal{A} over *Act*, we address the question of checking whether the implementation \mathcal{A} *conforms* to the specification \mathfrak{G} . This breaks up as two natural problems.

The model checking problem. Given a timed automaton \mathcal{A} over *Act* and a TCMMSG specification \mathfrak{G} , the model checking problem is to check that every timed word accepted by \mathcal{A} realizes some TCMSC in $\mathcal{L}_{TC}(\mathfrak{G})$. Since \mathcal{A} may accept timed words that are not well-formed or not complete, this implicitly includes checking that \mathcal{A} accepts only well-formed and complete timed words in TW_{Act} .

From this, it is clear that the model checking problem corresponds to checking whether $L(\mathcal{A}) \subseteq \mathcal{L}_{tw}(\mathfrak{G})$. To make the problem tractable, we restrict our attention to locally synchronized TCMMSGs, so that $\mathcal{L}_{tw}(\mathfrak{G})$ is a timed regular language. However, checking inclusion is undecidable even for timed regular languages [6], so this restriction is not sufficient to solve the model-checking problem.

The coverage problem. In the reverse direction, given a TCMMSG \mathfrak{G} and a timed automaton \mathcal{A} over *Act*, we ask if every $\mathfrak{M} \in \mathcal{L}_{TC}(\mathfrak{G})$ is witnessed by the implementation \mathcal{A} . Note that this is not the same as checking if $\mathcal{L}_{tw}(\mathfrak{G}) \subseteq \mathcal{L}_{tw}(\mathcal{A})$. Indeed, the timed inclusion problem would correspond to asking if every *timed linearization* of every $\mathfrak{M} \in \mathcal{L}_{TC}(\mathfrak{G})$ is witnessed by the implementation \mathcal{A} . This is rather strong, as the implementation is required to witness every possible way in which every TCMSC specification (generated by the TCMMSG) can be implemented. Indeed, the implementation can have strictly better time bounds than the specification and therefore be a valid implementation without witnessing all timed linearizations of the MSG. When we interpret TCMMSGs as incomplete positive specifications, a natural verification problem is to ask if the implementation witnesses every TCMSC generated by the TCMMSG (rather than every possible run of every TCMSC). This is the second problem that we address in this paper.

Let \mathcal{G} be a TCMSC and a timed automaton \mathcal{A} over Act . The *coverage problem* for \mathcal{G} and \mathcal{A} is to determine whether for each TCMSC $\mathfrak{M} \in \mathcal{L}_{TC}(\mathcal{G})$, there is a $w \in \mathcal{L}_{tw}(\mathfrak{M})$ such that $w \in \mathcal{L}_{tw}(\mathcal{A})$.

In the untimed case, the corresponding problem of scenario matching considered in [28,21], asks whether $\mathcal{L}_{MSC}(\mathcal{G}) \subseteq \mathcal{L}_{MSC}(\mathcal{A})$ where \mathcal{G} is an MSG and \mathcal{A} is an MPA. In the timed case, as mentioned above we cannot reduce coverage to language inclusion of timed MSCs. A TCMSC \mathfrak{M} represents an infinite family of TCMSCs, each of which realizes \mathfrak{M} . However, the implementation need not, in general, permit all these realizations. In other words, checking inclusion would correspond to checking if for each $\mathfrak{M} \in \mathcal{L}_{TC}(\mathcal{G})$, and each $w \in \mathcal{L}_{tw}(\mathfrak{M})$, it is the case that $w \in \mathcal{L}_{tw}(\mathcal{A})$. Indeed, this does not directly provide an answer to the coverage problem we have defined above.

Another plausible approach is to treat this as a timed game between Spoiler, who picks a path in the TCMSC \mathcal{G} , and Duplicator, who picks $w \in \mathcal{L}_{tw}(\mathcal{A}) \cap \mathcal{L}_{tw}(\mathfrak{M})$. At each step, Spoiler adds a node to the path in \mathcal{G} . Duplicator has to match this move by extending the current timed word so that it stays in $\mathcal{L}_{tw}(\mathcal{A})$ and also realizes the TCMSC described by the extended path. However, a winning strategy in this game would have the following property: if two paths π_1 and π_2 have a common prefix π , then w generated by Duplicator for the prefix π must be the same for the plays in which Spoiler generates π_1 and π_2 . This is not what we want since it may happen that π_1 and π_2 are realized by different plays that do not match on π and this solution will not be obtained by any winning strategy. In other words, the game-theoretic formulation introduces too strict a correlation between the timed words realizing different paths through the TCMSC.

These observations suggest that traditional approaches for scenario matching in the untimed case do not generalize to the coverage problem in the timed case.

Our strategy. Our approach to tackle both the above problems follows from the following basic observation. We observe that in locally synchronized TCMSCs, clocks are in used in a very particular way and do not in fact require the power of timed automata. We exploit this observation by showing that TCMSCs correspond to a strictly more restrictive model of timed automata which are closed under complementation.

In the next section, we introduce our restricted machine model for timed MSCs called MSC event clock automata. It turns out that $\mathcal{L}_{tw}(\mathcal{G})$ can be recognized by MSC event clock automata as demonstrated in Section 5. This yields a solution to our model checking problem in Section 6. Finally in Section 7, we use the theorem of Section 5 to obtain a solution for the coverage problem with an additional restriction on the specification.

4. An extended event clock automaton—the MSC-ECA

We now define MSC event clock automata or MSC-ECA. These will be used to capture exactly the guards that occur in the TCMSCs that we have defined. We denote an MSC-ECA over Act by $\mathcal{C} = (Q, Act, \delta, q_0, F)$, with states Q , initial state $q_0 \in Q$ and final states $F \subseteq Q$. A transition in δ is of the form (q, φ, a, q') , which we also write as $q \xrightarrow{\varphi, a} q'$, where $q, q' \in Q$, $a \in Act$ and φ is a conjunction of event clock guards. There are two types of event clock guards. First, we may check that the time elapsed between the k th-previous p -event and the current p -event is in some interval I . This guard is denoted $Y_p^k \in I$ (where I is an interval, as used in TCMSC timing constraints). In particular, we may check the time elapsed between the previous and the current p -events with $Y_p^1 \in I$. Second, we may check at a receive event the time elapsed since the matching write event. This guard is denoted $Msg^{-1} \in I$. We interpret these guards over timed words. Let $\sigma = (a_1, t_1) \cdots (a_n, t_n) \in TW_{Act}$. We define below when σ satisfies a guard φ at some position $1 \leq j \leq n$, denoted $\sigma, j \models \varphi$.

- (D1) $\sigma, j \models Y_p^k \in I$ if $a_j \in Act_p$ and there exists $1 \leq i < j$ such that $a_i \in Act_p$, $|\{\ell \mid i \leq \ell < j \wedge a_\ell \in Act_p\}| = k$ and $t_j - t_i \in I$. That is, the time elapsed between the k th-previous p -action a_i in σ and this action a_j is in the interval I .
- (D2) $\sigma, j \models Msg^{-1} \in I$ if a_j is a receive action and the time elapsed since the occurrence of its matching send action a_i is in the interval I . Formally, if there exists $p, q \in Proc$, $1 \leq i < j$ such that $a_i \in p!q$, $a_j \in q?p$, $|\{a_k \mid 1 \leq k \leq i, a_k \in p!q\}| = |\{a_k \mid 1 \leq k \leq j, a_k \in q?p\}|$ and $t_j - t_i \in I$ (recall that we write $a_k \in p!q$ and $a_k \in q?p$ to mean $a_k = p!q(m)$ and $a_k = q?p(m)$ for some $m \in \mathcal{M}$, respectively).

In both these definitions, note that action a_i is uniquely defined, i.e., there is at most one position i that matches a given position j with respect to a given event clock guard.

Now, we define runs of the MSC-ECA \mathcal{C} over timed words. For a timed word $\sigma = (a_1, t_1) \cdots (a_n, t_n)$, we say there is a run of \mathcal{C} from q to q' on σ , denoted $q \xrightarrow{\sigma} q'$ in \mathcal{C} , if there exists a sequence of transitions $q = q_0 \xrightarrow{\varphi_1, a_1} \cdots \xrightarrow{\varphi_n, a_n} q_n$ such that for all j , $1 \leq j \leq n$, $\sigma, j \models \varphi_j$. The timed word σ is said to be accepted if it has a run from the initial state to some final state in F . We denote by $\mathcal{L}_{tw}(\mathcal{C})$ the set of timed words accepted by the MSC-ECA \mathcal{C} . Notice that, time words in $\mathcal{L}_{tw}(\mathcal{C})$ need not be well-formed. An MSC-ECA is said to be *finite* if it has finitely many states.

4.1. Determinization and complementation of MSC-ECA

We now prove that MSC-ECA can be determinized and complemented, which is crucial for solving the model checking problem. We obtain this by constructing a deterministic and complete version of any given MSC-ECA. Intuitively, this works as for classical ECA's [7] and the main reason is that there are no explicit clocks. Since the reset of an event clock only

depends on the timed word being read and not on the path followed in the automaton, we can use the subset construction. However, instead of trying to encode the extended guards arising from the MSC structure in an MSC-ECA into a classical ECA, which seems rather difficult, we directly prove in this section that MSC-ECA can be determinized and complemented. Note that this is unlike the case of finite timed automata, which allow arbitrary clock resets and in general are not closed under complementation [6].

More precisely, let $\mathcal{C} = (Q, Act, \delta, q_0, F)$ be a finite MSC-ECA. The set of states of the universal automaton \mathcal{C}^{univ} is 2^Q . For a set $X \subseteq Q$ and an action a , we let $T(X, a)$ denote the set of transitions in δ having action a and a source state in X . Then, for some $T' \subseteq T(X, a) = T$, we denote by $\text{target}(T')$ the set of target states of transitions in T' and we define

$$\varphi(T', T) = \bigwedge_{t=(q, \varphi_t, a, q') \in T'} \varphi_t \wedge \bigwedge_{t=(q, \varphi_t, a, q') \in T \setminus T'} \neg \varphi_t.$$

We denote the set of transitions of \mathcal{C}^{univ} by Δ , where we say that $X \xrightarrow{\varphi, a} X' \in \Delta$ if there exists $T' \subseteq T = T(X, a)$ such that $\varphi = \varphi(T', T)$ and $X' = \text{target}(T')$.

Note that, once we have fixed X, a and the set T' , the transition is uniquely defined. Also for $X = \emptyset$, we have $T(X, a) = \emptyset$ and the only possible transition is $\emptyset \xrightarrow{\text{true}, a} \emptyset$. As before, a run on a timed word $\sigma = (a_1, t_1) \cdots (a_n, t_n)$ is defined as a sequence of transitions $X_0 \xrightarrow{\varphi_1, a_1} X_1 \cdots \xrightarrow{\varphi_n, a_n} X_n$ such that $\sigma, j \models \varphi_j$ for all j . The crucial property of \mathcal{C}^{univ} is that it is deterministic and complete (and finite, if \mathcal{C} is).

Lemma 1. *Given any timed word $\sigma = (a_1, t_1) \cdots (a_n, t_n) \in \text{TW}_{Act}$, there exists a unique run $X_0 \xrightarrow{\varphi_1, a_1} X_1 \xrightarrow{\varphi_2, a_2} \cdots X_{n-1} \xrightarrow{\varphi_n, a_n} X_n$ of \mathcal{C}^{univ} on σ starting from $X_0 = \{q_0\}$. Moreover, $X_n = \{q \in Q \mid q_0 \xrightarrow{\sigma} q \text{ in } \mathcal{C}\}$.*

Proof. Given $\sigma = (a_1, t_1) \cdots (a_n, t_n)$ and $X_0 = \{q_0\}$, for $j \in \{1, \dots, n\}$, we define inductively $T_j = T(X_{j-1}, a_j)$, $T'_j = \{(q, \varphi, a_j, q') \in T_j \mid \sigma, j \models \varphi\}$ and $X_j = \text{target}(T'_j)$. Observe that $X_{j-1} \xrightarrow{\varphi(T'_j, T_j), a_j} X_j$ is a transition of \mathcal{C}^{univ} . Also by definition of T'_j , for all $T' \subseteq T_j$, we have

$$\sigma, j \models \varphi(T', T_j) \text{ if and only if } T' = T'_j \quad (1)$$

Using the “if” part above, we obtain that $X_0 \xrightarrow{\varphi(T'_1, T_1), a_1} X_1 \cdots \xrightarrow{\varphi(T'_n, T_n), a_n} X_n$ is a run of \mathcal{C}^{univ} on σ . Conversely, we show by induction that this is the unique run of \mathcal{C}^{univ} on σ starting from $X_0 = \{q_0\}$. Let $X_0 \xrightarrow{\varphi_1, a_1} X'_1 \xrightarrow{\varphi_2, a_2} \cdots X'_n$ be any such run. Suppose $X'_{j-1} = X_{j-1}$. Then we show that $\varphi_j = \varphi(T'_j, T_j)$ and $X'_j = X_j$. By definition of a run, we have $X_{j-1} \xrightarrow{\varphi_j, a_j} X'_j$ and $\sigma, j \models \varphi_j$. But by the definition of a transition, there exists $T' \subseteq T(X_{j-1}, a_j) = T_j$ such that $\varphi_j = \varphi(T', T_j)$ and $X'_j = \text{target}(T')$. Thus, $\sigma, j \models \varphi(T', T_j)$ which by Eq. (1) implies that $T' = T'_j$. Thus, we conclude $\varphi_j = \varphi(T'_j, T_j)$ and $X'_j = \text{target}(T') = \text{target}(T'_j) = X_j$.

For the second statement, we prove both inclusions. First, if $q_0 \xrightarrow{\sigma} q$ in \mathcal{C} , let $q_0 \xrightarrow{\varphi_1, a_1} \cdots \xrightarrow{\varphi_n, a_n} q_n = q$ with $\sigma, j \models \varphi_j$ for $1 \leq j \leq n$. Using the notations above we show that for all $j \in \{0, \dots, n\}$, $q_j \in X_j$. Clearly $q_0 \in X_0$. Assume $q_{j-1} \in X_{j-1}$. Then we have $(q_{j-1}, \varphi_j, a_j, q_j) \in T'_j$. Hence we conclude that $q_j \in X_j = \text{target}(T'_j)$.

Conversely, for all j and for all $q_j \in X_j$ we show that $q_0 \xrightarrow{(a_1, t_1) \cdots (a_j, t_j)} q_j$ is a run of \mathcal{C} . The proof is by induction on j . $j = 0$ is obvious. Assume $j > 0$ and let $q_j \in X_j$. Then there exists $(q_{j-1}, \varphi_j, a_j, q_j) \in T'_j$, i.e., $q_{j-1} \in X_{j-1}$ and $\sigma, j \models \varphi_j$. By the induction hypothesis we have $q_0 \xrightarrow{(a_1, t_1) \cdots (a_{j-1}, t_{j-1})} q_{j-1}$. This implies that $q_0 \xrightarrow{(a_1, t_1) \cdots (a_{j-1}, t_{j-1})} q_{j-1} \xrightarrow{a_j, t_j} q_j$ is a run of \mathcal{C} . \square

By suitably choosing the final states, \mathcal{C}^{univ} will accept either the same language as \mathcal{C} or its complement. Let $F_1 = \{X \in 2^Q \mid F \cap X \neq \emptyset\}$ and $F_2 = 2^Q \setminus F_1$. Define $\mathcal{C}_i^{univ} = (2^Q, Act, \Delta, \{q_0\}, F_i)$ for $i = \{1, 2\}$. From Lemma 1 we obtain:

Corollary 2. *We have $\mathcal{L}_{tw}(\mathcal{C}_1^{univ}) = \mathcal{L}_{tw}(\mathcal{C})$ and $\mathcal{L}_{tw}(\mathcal{C}_2^{univ}) = \text{TW}_{Act} \setminus \mathcal{L}_{tw}(\mathcal{C})$.*

4.2. From MSC-ECA to TA

Not every MSC-ECA can be translated into an equivalent (classical) timed automaton. The problem comes from the event guards $\text{Msg}^{-1} \in I$, which may require infinitely many clocks if channels are unbounded. Fortunately, thanks to the locally synchronized assumption on TCMSCs, we are only interested in bounded channels. Let $B > 0$. We show below how to construct a timed automaton \mathcal{B}_C^B from an MSC-ECA $\mathcal{C} = (Q, Act, \delta, q_0, F)$ such that \mathcal{B}_C^B and \mathcal{C} are equivalent, in the sense formalized below, when restricted to B -bounded channels.

Let $K = \max\{k \mid Y_p^k \in I \text{ occurs in some guard in } \delta\}$. A state of \mathcal{B}_C^B is either a dead state denoted \perp or a tuple $\mathfrak{s} = (s, \bar{b}, \bar{n}, \bar{\alpha}, \bar{\beta})$ where $s \in Q$, $\bar{b} = (b_p)_{p \in \text{Proc}} \in \{0, 1\}^{\text{Proc}}$ ($b_p = 1$ if we have already seen at least K p -events), $\bar{n} = (n_p)_{p \in \text{Proc}} \in$

$\{0, \dots, K-1\}^{Proc}$ (n_p is the number of p -events already seen modulo K), $\bar{\alpha} = (\alpha_{p,q})_{p,q \in Proc} \in \{0, \dots, B\}^{Proc^2}$ ($\alpha_{p,q}$ is the number of $q?p$ events modulo $B+1$), $\bar{\beta} = (\beta_{p,q})_{p,q \in Proc} \in \{0, \dots, B\}^{Proc^2}$ ($\beta_{p,q}$ is the number of $p!q$ events modulo $B+1$). The set of all states is denoted Q' and the initial state is $s_0 = (s_0, (0), (0), (0), (0))$. The set of clocks is $Y \cup Z$ where $Y = \{y_p^i \mid p \in Proc, 0 \leq i < K\}$ and $Z = \{z_{p,q}^i \mid p, q \in Proc, 0 \leq i \leq B\}$. We will reset clock y_p^i when executing the i th p -event mod K . Also, $z_{p,q}^i$ will be reset when executing the i th $p!q$ event mod $B+1$.

We say that channel (p, q) is *empty* if $\alpha_{p,q} = \beta_{p,q}$ and *full* if $\beta_{p,q} = \alpha_{p,q} + B \bmod (B+1)$. The set of transitions $\delta_{\mathcal{B}_C^B}$ is defined as follows: Assume $s \xrightarrow{\varphi, a} s'$ in \mathcal{C} with $a \in Act_p$. Then, we have three types of transitions in \mathcal{B}_C^B .

(Tr1) $(s, \bar{b}, \bar{n}, \bar{\alpha}, \bar{\beta}) \xrightarrow{\text{true}, a, \emptyset} \perp$ is in \mathcal{B}_C^B if either $a \in p!q$ and channel (p, q) is full (the bound was exceeded), or $a \in p?q$ and channel (p, q) is empty.

(Tr2) $(s, \bar{b}, \bar{n}, \bar{\alpha}, \bar{\beta}) \xrightarrow{\varphi', a, R} (s', \bar{b}', \bar{n}', \bar{\alpha}', \bar{\beta}')$ is in \mathcal{B}_C^B if we are not in the above case and the following conditions hold:

1. $b'_r = b_r$ for $r \neq p$ and $b'_p = \begin{cases} 1 & \text{if } n_p = K-1 \\ b_p & \text{otherwise.} \end{cases}$
2. $n'_r = n_r$ for $r \neq p$ and $n'_p = (n_p + 1) \bmod K$.
3. if $a \in p!q$, then $\beta'_{p,q} = (\beta_{p,q} + 1) \bmod (B+1)$ and $\beta'_{p',q'} = \beta_{p',q'}$ for $(p', q') \neq (p, q)$. Also $\bar{\alpha}' = \bar{\alpha}$, $R = \{y_p^{n'_p}, z_{p,q}^{\beta'_{p,q}}\}$ and φ' is φ where $Y_p^k \in I$ is replaced with

$$\begin{cases} \text{false} & \text{if } b_p = 0 \text{ and } k > n_p \\ y_p^{(K+n'_p-k) \bmod K} \in I & \text{otherwise} \end{cases}$$

4. if $a \in p?q$, then $\alpha'_{q,p} = \alpha_{q,p} + 1 \bmod (B+1)$ and $\alpha'_{q',p'} = \alpha_{q',p'}$ for $(q', p') \neq (q, p)$. Also $\bar{\beta}' = \bar{\beta}$, $R = \{y_p^{n'_p}\}$ and φ' is φ where $Y_p^k \in I$ is replaced as above and $\text{Msg}^{-1} \in I$ is replaced with $z_{q,p}^{\alpha'_{q,p}} \in I$.

(Tr3) $\perp \xrightarrow{\text{true}, a, \emptyset} \perp$ is in \mathcal{B}_C^B for all $a \in Act$.

In the following, we call a timed word w *weakly well-formed* (wwf) if for each channel (p, q) , in every prefix v of w , we have $\#_{q?p}(v) \leq \#_{p!q}(v)$. This weak form does not require the sequence of received messages to be a prefix of the sequence of the sent messages—it only demands that at any point, the number of messages received does not exceed the number of messages sent. Let $\text{TW}_{Act}^{B, \text{wwf}}$ denote the set of timed words $\sigma \in \text{TW}_{Act}$ which are both wwf and B -bounded.

We can immediately observe some invariant properties that are maintained by the above transitions. Let $s_0 \xrightarrow{\varphi_1, a_1, R_1} \dots \xrightarrow{\varphi_m, a_m, R_m} s_m$ for $m \geq 0$ be a path in \mathcal{B}_C^B from the initial state s_0 to some state $s_m \neq \perp$. Then, for $s_m = (s_m, \bar{b}, \bar{n}, \bar{\alpha}, \bar{\beta})$,

1. $b_p = 1$ if $|\{\ell \mid 1 \leq \ell \leq m \wedge a_\ell \in Act_p\}| \geq K$ and $b_p = 0$ otherwise.
2. $n_p = |\{\ell \mid 1 \leq \ell \leq m \wedge a_\ell \in Act_p\}| \bmod K$
3. $\alpha_{p,q} = |\{\ell \mid 1 \leq \ell \leq m \wedge a_\ell \in q?p\}| \bmod (B+1)$
4. $\beta_{p,q} = |\{\ell \mid 1 \leq \ell \leq m \wedge a_\ell \in p!q\}| \bmod (B+1)$

On the other hand, suppose $s_0 \xrightarrow{\varphi_1, a_1, R_1} \dots \xrightarrow{\varphi_m, a_m, R_m} s_m$ for $m \geq 0$ is a path in \mathcal{B}_C^B from the initial state s_0 to $s_m = \perp$. Then, either σ is not wwf or it exceeds the bound B for some channel.

We can define different notions of acceptance (i.e., final states) on \mathcal{B}_C^B constructed from \mathcal{C} to derive the results below.

Proposition 3. Let $\mathcal{C} = (Q, Act, \delta, q_0, F)$ and $\mathcal{B}_C^B = (Q', Act, \delta_{\mathcal{B}_C^B}, s_0, F', (Y \cup Z))$ be as above, with some set F' of final states.

1. With final states $F' = \{(s, \bar{b}, \bar{n}, \bar{\alpha}, \bar{\beta}) \mid s \in F\}$ the timed automaton \mathcal{B}_C^B accepts the language $\mathcal{L}_{\text{tw}}(\mathcal{C}) \cap \text{TW}_{Act}^{B, \text{wwf}}$.
2. If \mathcal{C} is complete (i.e., it has a run on every timed word over Act) then with final states $F' = \{\perp\}$ the timed automaton \mathcal{B}_C^B accepts the complement of $\text{TW}_{Act}^{B, \text{wwf}}$ in TW_{Act} .

Further if \mathcal{C} is finite so is \mathcal{B}_C^B .

Proof. 1. Let $\sigma = (a_1, t_1) \dots (a_m, t_m)$ be a wwf and B -bounded timed word. Consider a path $\pi = s_0 \xrightarrow{\varphi_1, a_1} s_1 \xrightarrow{\varphi_2, a_2} \dots \xrightarrow{\varphi_m, a_m} s_m$ of \mathcal{C} . We can build inductively a path $\pi' = s_0 \xrightarrow{\varphi'_1, a_1, R_1} s_1 \xrightarrow{\varphi'_2, a_2, R_2} \dots \xrightarrow{\varphi'_m, a_m, R_m} s_m$ of \mathcal{B}_C^B starting from its initial state s_0 and using (Tr2) only. Then, $\mathcal{L}_{\text{tw}}(\mathcal{C}) \cap \text{TW}_{Act}^{B, \text{wwf}} \subseteq \mathcal{L}_{\text{tw}}(\mathcal{B}_C^B)$ follows immediately from the following claim.

Claim 1. If σ has a run through π in \mathcal{C} (i.e., $\sigma, i \models \varphi_i$ for all $i \in \{1, \dots, m\}$) then σ has a run through π' in \mathcal{B}_C^B .

We define inductively the valuation sequence for the run through π' : ν_0 is the valuation mapping all clocks to 0, and $\nu_i = (\nu_{i-1} + t_i - t_{i-1})[R_i \rightarrow 0]$ for $1 \leq i \leq m$. To establish [Claim 1](#) we show for all $i \in \{1, \dots, m\}$ that $(\nu_{i-1} + t_i - t_{i-1}) \models \varphi'_i$. There are three cases:

1. φ'_i contains $z_{p,q}^k \in I$ where $k = \alpha'_{p,q} = |\{a_\ell \mid 1 \leq \ell \leq i \wedge a_\ell \in q?p\}| \bmod (B+1)$. From the definition of the transition, φ_i must contain $\text{Msg}^{-1} \in I$. Since, $\sigma, i \models \varphi_i$ we have $t_i - t_j \in I$, where j is the index of the matching send: $a_j = p!q(m), a_i = q?p(m), 1 \leq j \leq i$ and $|\{a_\ell \mid 1 \leq \ell \leq j \wedge a_\ell \in p!q\}| = |\{a_\ell \mid 1 \leq \ell \leq i \wedge a_\ell \in q?p\}|$. Thus, $k = |\{a_\ell \mid 1 \leq \ell \leq j \wedge a_\ell \in p!q\}| \bmod (B+1)$. Using the invariant at state s_j , we get $z_{p,q}^k \in R_j$. Using the invariant at s_i , we can replace $\text{Msg}^{-1} \in I$ by $z_{p,q}^k \in I$ in φ'_i . Moreover, $z_{p,q}^k \notin R_\ell$ for $j < \ell \leq i$ —otherwise, the number of events labeled $p!q$ between j and ℓ would be B more than the number of events labeled $q?p$ between j and i (and therefore ℓ). This implies that the channel was full and so, at ℓ , transition (*Tr1*) is enabled which means that transition (*Tr2*) cannot be fired, which contradicts the transition at i . Now, $z_{p,q}^k \in R_j$ implies $\nu_j(z_{p,q}^k) = 0$, and $z_{p,q}^k \notin R_\ell$ for $j < \ell \leq i$ implies that $(\nu_{i-1} + t_i - t_{i-1})(z_{p,q}^k) = \nu_j(z_{p,q}^k) + t_i - t_j = t_i - t_j$. So, we have $(\nu_{i-1} + t_i - t_{i-1}) \models (z_{p,q}^k \in I)$.
2. We will show that φ' cannot contain false. If φ' contains false, then $b_p = 0$ and there exists $Y_p^k \in I$ in φ such that $k > n_p$. But $b_p = 0$ implies that K events have not been seen and so we are trying to relate two events that are k apart when we have not seen k events on p . This contradicts the fact that $\sigma, i \models Y_p^k \in I$, so this cannot happen.
3. φ'_i contains $y_p^\ell \in I$. Then $\ell = (K + n'_p - k) \bmod K$ and $Y_p^k \in I$ is in φ_i . Consider the event j such that the number of p -events between j and i is k . Such an event exists since either $n'_p > k$ or $b_p = 1$ (which means that $K > k$ many p -events have been seen). But if $k < n'_p$, then $\ell = n'_p - k$ and so the value of n_p -component at j is ℓ . If $k > n'_p$, then at j , we have $\ell = K + n'_p - k < K$. Thus in both cases, y_p^ℓ was reset at j and not reset again between j and i . Again, $\sigma, i \models \varphi_i$ implies that $t_i - t_j \in I$ and so $(\nu_{i-1} + t_i - t_{i-1})(y_p^\ell) = \nu_j(y_p^\ell) + t_i - t_j = t_i - t_j \in I$. Thus, $(\nu_{i-1} + t_i - t_{i-1}) \models y_p^\ell \in I$.

For the converse inclusion, we start with a path of \mathcal{B}_C^B starting from its initial state s_0 and which does not reach \perp :

$\pi' = s_0 \xrightarrow{\varphi'_1, a_1, R_1} s_1 \xrightarrow{\varphi'_2, a_2, R_2} \dots \xrightarrow{\varphi'_m, a_m, R_m} s_m$. Since we did not reach \perp , the timed word $\sigma = (a_1, t_1) \dots (a_m, t_m)$ must be wwf and B -bounded. Moreover, transitions in π' comes from (*Tr2*) only and we can recover a corresponding path $\pi = s_0 \xrightarrow{\varphi_1, a_1} s_1 \xrightarrow{\varphi_2, a_2} \dots \xrightarrow{\varphi_m, a_m} s_m$ in \mathcal{C} . Again, we can prove that if σ has a run through π' in \mathcal{B}_C^B then σ has a run through π in \mathcal{C} . This follows by considering each case for the guards and observing that if $\nu_{i-1} + t_i - t_{i-1} \models \varphi'_i$ then $\sigma, i \models \varphi_i$ in each case.

2. We have already noted that if a timed word σ has a run through a path of \mathcal{B}_C^B reaching the dead state \perp then σ is either not wwf or not B -bounded. Conversely, assume that σ is either not wwf or not B -bounded and let σ' be the greatest prefix of σ which is both wwf and B -bounded. Since \mathcal{C} is complete, the timed word σ' has a run through a path π of \mathcal{C} . As above we deduce that σ' has a run through a corresponding path π' of \mathcal{B}_C^B . The next letter of σ will violate either the B -bound or the wwf condition. Hence the run reaches \perp with this next letter and loops on \perp until the end of σ . \square

5. From a locally synchronized TCMMSG to a finite MSC-ECA

The main result of this section is that locally synchronized TCMMSGs define timed regular languages.

Theorem 4. *If $\mathfrak{G} = (G, \mathcal{L}^{TC}, \Phi, \text{EdgeC})$ is a locally synchronized TCMMSG, then there exists a finite MSC-ECA \mathcal{C} , such that $\mathcal{L}_{\text{tw}}(\mathcal{C}) = \mathcal{L}_{\text{tw}}(\mathfrak{G})$.*

In the untimed case, the corresponding result has been stated and proved in different ways [9,18,27,14]. We describe a different proof that is more suitable for the timed version. It is split in three main steps that are described in the following sections.

5.1. TCMMSG to an infinite MSC-ECA

In this section, from a TCMMSG, we construct an MSC-ECA (with infinitely many states) which accepts exactly the same set of timed linearizations. We start with a definition and a remark. For an MSC $M = (E, \leq, \lambda)$ over Act , recall that a cut c of M over Act to be a subset of the events E which is closed under the partial order \leq . That is, $e \in c, e' \leq e$ implies that $e' \in c$. In what follows, we will use this definition in the setting of MSCs generated by a path π , namely M^π . We also recall that any event of E^π is of the form $(e, \rho u)$ where $\rho u \leq \pi$ and $e \in E^u$. Indeed, keeping the prefix of the path along with the event uniquely identifies the event's occurrence in the path.

For a fixed TCMMSG $\mathfrak{G} = (G, \mathcal{L}^{TC}, \Phi, \text{EdgeC})$, where $G = (V, \rightarrow, \nu_{in}, V_F)$, we define the infinite MSC-ECA denoted $\mathcal{C}_{\mathfrak{G}}$, which we sometimes call the *global system* of \mathfrak{G} . A state of $\mathcal{C}_{\mathfrak{G}}$ is a pair $\bar{s} = (\pi, C)$ where

- π is a path in G .
- $C \subseteq E^\pi$ is a cut of M^π

Now, an event (e, ρ) is said to have been *executed* in \bar{s} if $(e, \rho) \in C$. The event is said to be *enabled* in \bar{s} if it has not been executed, i.e., $(e, \rho) \notin C$, and all the events below it (in the partial order) have been executed, i.e., for all $(e', \rho') \in E^\pi$ with $(e', \rho') <^\pi (e, \rho)$, we have $(e', \rho') \in C$.

A state $\bar{s} = (\pi, C)$ is *initial* if π is any path in G from an initial state to a final state and C is empty. It is *final* if $C = E^\pi$. We denote the set of all states of this global system by $Q_{\mathcal{G}}$.

Next, let \mathcal{I} denote the set of all intervals appearing as constraints in the TCMMSG. For any node u in the TCMMSG, we note that $\Phi(u) = M_u = (E^u, \leq^u, \lambda_u, \tau^u)$ is a TCMSC from \mathcal{L}^{TC} . This definition lifts as before to paths π of $\mathcal{C}_{\mathcal{G}}$ as well.

Now, the transitions can be defined by saying that at any state we execute an enabled event. We have $\bar{s} = (\pi, C) \xrightarrow{\varphi, a} \bar{s}' = (\pi, C')$ if there exists an event $(e, \rho u)$ enabled in \bar{s} such that $\lambda^u(e) = a$ and

- $C' = C \uplus \{(e, \rho u)\}$ (where \uplus stands for disjoint union)
- the guard φ checks all local and edge constraints that are matched here,

$$\varphi = \left(\bigwedge_{e' \in E^u, I \in \mathcal{I} | \tau^u(e', e) = I} \varphi(u, e', e, I) \right) \wedge \varphi^{edge} \text{ where,} \quad (2)$$

$$\varphi(u, e', e, I) = \begin{cases} \text{Msg}^{-1} \in I & \text{if } \exists p, q, p \neq q \text{ s.t. } e' <_{qp}^u e \\ Y_p^k \in I & \text{if } e, e' \in E_p^u, |\{e'' \in E_p^u \mid e' <_{pp}^u e'' <_{pp}^u e\}| = k \end{cases} \quad (3)$$

$$\text{and } \varphi^{edge} = \begin{cases} Y_p^1 \in I & \text{if } \rho = \rho' u', \text{ and for some } p \in \text{Proc}, \text{ we have} \\ & \text{EdgeC}((u', u), p) = I \text{ and } e = \min(E_p^u) \\ \text{true} & \text{otherwise} \end{cases} \quad (4)$$

Note that, in the transition above, the event $(e, \rho u)$ which is enabled in \bar{s} becomes an executed event of \bar{s}' . Thus we can say that the transition $\bar{s} \xrightarrow{\varphi, a} \bar{s}'$ *executes* the event $(e, \rho u)$.

As before, a run of $\mathcal{C}_{\mathcal{G}}$ on a timed word $\sigma = (a_1, t_1) \cdots (a_n, t_n)$ is a sequence of transitions $\bar{s}_0 \xrightarrow{\varphi_1, a_1} \cdots \xrightarrow{\varphi_n, a_n} \bar{s}_n$ such that for each $j \in \{1, \dots, n\}$, $\sigma, j \models \varphi_j$. Again a run is accepting if it starts at an initial state and ends in a final state. We say a timed word σ belongs to $\mathcal{L}_{tw}(\mathcal{C}_{\mathcal{G}})$, if there is an accepting run on σ .

In what follows, we often refer to runs of the global system, i.e., the MSC-ECA $\mathcal{C}_{\mathcal{G}}$ as *global runs*. Also since we have fixed a TCMMSG $\mathcal{G} = (G, \mathcal{L}^{TC}, \Phi, \text{EdgeC})$ throughout this section, we often just write *the global system*, when referring to $\mathcal{C}_{\mathcal{G}}$.

Lemma 5. *We have the following relation between the timed languages of $\mathcal{C}_{\mathcal{G}}$ and \mathcal{G} : $\mathcal{L}_{tw}(\mathcal{C}_{\mathcal{G}}) = \mathcal{L}_{tw}(\mathcal{G}) = \{\sigma \mid \sigma \text{ is a timed linearization of some TMSC } T \text{ over Act, such that } T \text{ realizes some } \mathfrak{M} \in \mathcal{L}_{TC}(\mathcal{G})\}$.*

Proof. (\subseteq) Let $\sigma = (a_1, t_1) \cdots (a_n, t_n) \in \mathcal{L}_{tw}(\mathcal{C}_{\mathcal{G}})$. Then there exists an accepting run

$$\bar{s}_0 \xrightarrow{\varphi_1, a_1} \cdots \xrightarrow{\varphi_n, a_n} \bar{s}_n$$

where for each $i \in \{1, \dots, n\}$, $\sigma, i \models \varphi_i$ and $\bar{s}_{i-1} = (\pi_{i-1}, C_{i-1}) \xrightarrow{\varphi_i, a_i} (\pi_i, C_i) = \bar{s}_i$ executes some enabled event (e_i, ρ_i) .

First, $\pi_0 = \cdots = \pi_n = \pi$ (say). Then, as $\bar{s}_0 = (\pi, C_0)$ is an initial state of the global system, π is a path from the initial vertex v_{in} to some final one in G . Therefore $\mathfrak{M}^\pi = (M^\pi, \tau^\pi) \in \mathcal{L}_{TC}(\mathcal{G})$. Now, for each $i \in \{1, \dots, n\}$, $C_i = C_{i-1} \uplus \{(e_i, \rho_i)\}$ is a cut of M^π . Moreover, $C_n = E^\pi$ since \bar{s}_n is final. From this we get that $(e_1, \rho_1) \cdots (e_n, \rho_n)$ is a linearization of M^π and $\lambda^\pi(e_i, \rho_i) = a_i$ for all $i \in \{1, \dots, n\}$. Now consider the TMSC $T = (M^\pi, t)$ where we define t by $t((e_i, \rho_i)) = t_i$. Thus, $(a_1, t_1) \cdots (a_n, t_n)$ is a timed linearization of T since $i < j$ implies $t(e_i, \rho_i) = t_i \leq t_j = t(e_j, \rho_j)$.

We are done if we show that T realizes \mathfrak{M}^π . That is, for all $((e_i, \rho_i), (e_j, \rho_j)) \in \text{dom}(\tau^\pi)$, we want to show that $|t(e_j, \rho_j) - t(e_i, \rho_i)| = t_j - t_i \in \tau^\pi((e_i, \rho_i), (e_j, \rho_j))$. We have two cases to handle:

- If $\rho_i = \rho_j = \rho v$ (say) then $\tau^\pi((e_i, \rho v), (e_j, \rho v)) = \tau^v(e_i, e_j) = I$. Then, first $\varphi(v, e_i, e_j, I)$ is in φ_j . Indeed, if $\tau^v(e_i, e_j) = I$ then one of the two following cases hold:
 - Either $e_i, e_j \in E_p^v$ for some $p \in \text{Proc}$. Then, $|\{e_\ell \in E_p^v \mid e_i <_{pp}^v e_\ell <_{pp}^v e_j\}| = k$ for some $k \in \mathbb{N}_{>0}$. Thus $\varphi(v, e_i, e_j, I) = Y_p^k \in I$. At state j , $\sigma, j \models \varphi_j$ implies $\sigma, j \models \varphi(v, e_i, e_j, I)$ which implies that $\sigma, j \models Y_p^k \in I$. Now, $e_\ell \in E_p^v$ such that $e_i <_{pp}^v e_\ell <_{pp}^v e_j$ if and only if $i \leq \ell < j$ such that $a_\ell \in \text{Act}_p$. Thus, by Definition (D1), we conclude that $t_j - t_i \in I$.
 - Or $e_i <_{qp}^v e_j$ for some $p, q \in \text{Proc}, p \neq q$. Then, $\varphi(v, e_i, e_j, I) = \text{Msg}^{-1} \in I$. Again, we have $\sigma, j \models \text{Msg}^{-1} \in I$. Now, $e_i <_{qp}^v e_j$ implies that $\lambda^\pi(e_j, \rho_j) = a_j = p?q(m)$ for some $m \in \mathcal{M}$ and $\lambda^\pi(e_i, \rho_i) = a_i = q?p(m)$ is its matching send. Thus, $|\{a_\ell \mid 1 \leq \ell \leq i, a_\ell \in q?p\}| = |\{a_\ell \mid 1 \leq \ell \leq j, a_\ell \in p?q\}|$. Now, by Definition (D2) we conclude that $t_j - t_i \in I$.
- Otherwise, $\rho_j = \rho_i v, \rho_i = \rho v'$ for some $\rho, e_i = \max(E_p^v)$ and $e_j = \min(E_p^v)$ for some $p \in \text{Proc}$, then $\tau^\pi((e_i, \rho_i), (e_j, \rho_j)) = \text{EdgeC}((v', v), p) = I$. Then at stage j , we have $\varphi_j^{edge} = (Y_p^1 \in I)$. Indeed, $a_i = \lambda^{v'}(e_i)$ is the last p -action before $a_j = \lambda^v(e_j)$ in σ . Thus, by Definition (D1), $t_j - t_i \in I$ and so we are done.

(\supseteq) Suppose $\mathfrak{M} \in \mathcal{L}_{TC}(\mathfrak{G})$, then there exists a path $\pi = v_1 \cdots v_m$ in G such that v_1 is an initial vertex and v_m is a final vertex and $\mathfrak{M} = \mathfrak{M}^\pi = (M^\pi, \tau^\pi)$. Now, suppose $\sigma = (a_1, t_1) \cdots (a_n, t_n)$ is a timed linearization of $T = (M^\pi, t)$ and T realizes \mathfrak{M} . Then first we observe that $a_1 \cdots a_n \in \text{lin}(M^\pi)$ and so there is $(e_1, \rho_1) \cdots (e_n, \rho_n)$ a linearization of the events of M^π where for each i $\rho_i \preceq \pi$, $\lambda^\pi(e_i, \rho_i) = a_i$.

Then we can construct the run of the global system on this timed word. First, we define $C_i = \{(e_1, \rho_1) \cdots (e_i, \rho_i)\}$ and $\bar{s}_i = (\pi, C_i)$ for all $i \in \{0, \dots, n\}$, where $C_0 = \emptyset$. Then observe that (e_i, ρ_i) is enabled in \bar{s}_{i-1} . Thus there exists a transition $\bar{s}_{i-1} \xrightarrow{\varphi_i, a_i} \bar{s}_i$ that executes event (e_i, ρ_i) . We show that $\sigma, i \models \varphi_i$ where φ_i is defined by the transition. Again there are two cases:

- Either φ_i contains an edge constraint, i.e., $\varphi_i^{\text{edge}} = (Y_p^1 \in I)$ for some $p \in \text{Proc}$. In this case, by Condition 4, $\rho_i = \rho' v' v$, e_i is the first p -event in M^v and for some $j < i$, we have $\rho_j = \rho' v'$, e_j is the last p -event on M^v and $\text{EdgeC}((v', v), p) = I$. First, this implies that a_i is the next p -action with respect to a_j in σ . Also, $\tau^\pi((e_j, \rho_j), (e_i, \rho_i)) = I$ and since T realizes \mathfrak{M} , we have $t(e_i, \rho_i) - t(e_j, \rho_j) \in I$ which implies that $t_i - t_j \in I$. By Definition (D1), we conclude that $\sigma, i \models \varphi_i^{\text{edge}}$.
- Or there is a local constraint of a node of the form $\varphi(u, e_j, e_i, I)$, where $\rho_i = \rho u = \rho_j$ for some $j < i$ and $\tau^u(e_j, e_i) = I$. Again since T realizes \mathfrak{M} , $t(e_i, \rho u) - t(e_j, \rho u) = t_i - t_j \in I$. By Condition 3, if the constraint is of the form $\text{Msg}^{-1} \in I$, then $e_j \prec_{qp}^u e_i$ and otherwise the constraint is of the form $Y_p^k \in I$ where the number of p -events between e_j and e_i is k . Using Definition (D2) in the former case and Definition (D1) in the latter case, we conclude that $\sigma, i \models \varphi(u, e_j, e_i, I)$.

Thus, we have shown that $\sigma, i \models \varphi_i$ for all $i \in \{1, \dots, n\}$ and therefore $\bar{s}_0 \xrightarrow{\varphi_1, a_1} \cdots \xrightarrow{\varphi_n, a_n} \bar{s}_n$ is a run of the global system on σ . Finally, the run ends in a final state of the global system, since σ is a full linearization of M^π . Thus, our proof is complete. \square

5.2. Removing unexecuted nodes

We want to simulate the global run of a TCMMSG in a finite way. So, instead of maintaining the whole path along the run, we want to maintain only the relevant portions, i.e., the nodes in which there is at least an event that has occurred.

For segments of nodes in the path that have not seen any event yet, we replace them by a special gap symbol $\#$. Thus, having a $\#$ symbol between two nodes denotes that some (nonempty) sequence of nodes must be inserted here later.

In fact, the insertion must satisfy two conditions: (1) when we insert a node it must not conflict with the events that have already occurred in later nodes and (2) finally, after all insertions, we do obtain a path in the graph. The latter is done by checking that when we fill a gap the corresponding bordering nodes have an edge in the graph.

This construction is formalized next. However, note that this construction is still infinite since we might still have unboundedly many completed nodes, i.e., nodes in which all events have been seen. In the next section, we describe how to perform a sequence of reductions to throw away such completed nodes from the current path. However, we have to be careful that the two conditions, in the infinite case above, are still maintained.

We start by observing that the cut C that we keep in a state in the simulation in the previous section is global. Thus, if we want to remove some nodes we would need to maintain the cut C locally within each node. To do this we break up each state $(u_1 \cdots u_n, C)$ into $(u_1, c_1) \cdots (u_n, c_n)$. Formally, we define the map Φ which we call *stratification* as follows:

$$\Phi((u_1 \cdots u_n, C)) = (u_1, c_1) \cdots (u_n, c_n)$$

where each $c_i \subseteq E^{u_i}$ is defined by $c_i = \{e \in E^{u_i} \mid (e, u_1 \cdots u_i) \in C\}$. Notice that each c_i is a cut of E^{u_i} . Φ is in fact a bijection since we also have the inverse map given by $C = \{(e, u_1 \cdots u_i) \in E^{u_1 \cdots u_n} \mid e \in c_i\}$.

We define an *extended node* to be a pair (u, c) where $u \in V$ and $c \subseteq E^u$ is a cut of E^u . As before, c contains the events that have been executed in node u . For simplicity, we extend the set of vertices V with two dummy vertices $\triangleright, \triangleleft$ and add edges from \triangleright to the initial vertex v_{in} and from every final vertex $v \in V_F$ to \triangleleft . We also set $E^\triangleright = \emptyset = E^\triangleleft$ so that for $u \in \{\triangleright, \triangleleft\}$, the only extended node is (u, \emptyset) . We will also maintain that for any extended node (u, c) , if $u \notin \{\triangleright, \triangleleft\}$, then $c \neq \emptyset$. The set of all extended nodes is denoted ExtNodes and we let $\Gamma = \text{ExtNodes} \cup \{\#\}$.

Now, we construct our new automaton $C_{\mathfrak{G}}^\#$. A state α of $C_{\mathfrak{G}}^\#$ is an element of Γ^* . The initial state is $\alpha_0 = (\triangleright, \emptyset)\#(\triangleleft, \emptyset)$. Now, we lift the notion of events to *extended events* of a state in this new automaton. An *extended event* of $\alpha \in \Gamma^*$ is a pair $(e, \alpha_1(u, c))$ where $e \in E^u$ and $\alpha_1(u, c) \preceq \alpha$. We say that the extended event $(e, \alpha_1(u, c))$ is

- *executed* in α if $e \in c$ and
- *enabled* in α if the following hold:
 - (E1) it has not been executed, i.e., $e \notin c$,
 - (E2) all events within the node which are below it (in the partial order) have been executed, i.e., for all $e' \in E^u$ with $e' \prec^u e$, we have $e' \in c$
 - (E3) and if e belongs to process p , then all p -events on any node occurring before this node in α have been executed, i.e., if $e \in E_p^u$ then for all $\alpha'_1(u', c') \preceq \alpha_1$, we have $E_p^{u'} \subseteq c'$.

An extended node (u, c) is said to be *completed* if $c = E^u$. Note that $(\triangleright, \emptyset)$ and $(\triangleleft, \emptyset)$ are completed by default. A state α is *final* if it is a sequence of completed nodes.

We will need some notations to describe the set of processes that participate in node, path or a state. First, for a node $u \in V$, $OProc(u) = \{p \in Proc \mid E_p^u \neq \emptyset\}$ denotes the set of processes that participate (*occur*) in u . This is extended to V^* as a morphism. Also, with $OProc(u, c) = OProc(u)$ and $OProc(\#) = \emptyset$ it extends to Γ^* . In addition, for $\beta \in \Gamma^*$, $EProc(\beta)$ denoting the set of *executed* events in β , is given by the morphism defined by $EProc((u, c)) = \{p \in Proc \mid E_p^u \cap c \neq \emptyset\}$, $EProc(\#) = \emptyset$.

Now, the transitions can be defined by saying that at any state we can choose to execute an enabled event or add a new (extended) node to the state and then we must execute an enabled event on the new node. In fact, we always add a node by inserting it in a $\#$.

Let us now define the *node insertion* operation which tells us how a node is inserted in a gap. Formally, this is defined as a macro $\alpha_1 \# \alpha_2 \xrightarrow{u} \alpha'_1(u, \emptyset) \alpha'_2$ which is said to hold if

- (11) for every process that participates in u , there is no executed event in the segment α_2 on that process, i.e., $OProc(u) \cap EProc(\alpha_2) = \emptyset$.
- (12) $\alpha'_1 \in \{\alpha_1, \alpha_1 \#\}$ and if $\alpha'_1 = \alpha_1$ then $\alpha_1 = \alpha''_1(v, c)$ and $v \rightarrow u$ in G .
- (13) $\alpha'_2 \in \{\alpha_2, \#\alpha_2\}$ and if $\alpha'_2 = \alpha_2$ then $\alpha_2 = (v, c) \alpha'_2$ and $u \rightarrow v$ in G .

Now, using this macro we can define the transition relation as follows. Formally, $\alpha \xrightarrow{\varphi, a} \alpha'$ is a transition in $\mathcal{C}_{\mathcal{G}}^{\#}$ if there exists $\beta = \beta_1(u, c) \beta_2$ and an extended event $(e, \beta_1(u, c))$ enabled in β such that

- one of the two following conditions hold:
 - (i) either $\beta = \beta_1(u, c) \beta_2 = \alpha$, i.e., the enabled event is already present in the current state,
 - (ii) or $\alpha = \alpha_1 \# \alpha_2 \xrightarrow{u} \beta_1(u, \emptyset) \beta_2 = \beta$. Hence, $c = \emptyset$, $\beta_1 \in \{\alpha_1, \alpha_1 \#\}$ and $\beta_2 \in \{\alpha_2, \#\alpha_2\}$
- and all the below conditions hold:
 - (T1) $a = \lambda^u(e)$
 - (T2) the guard φ must check all local and edge constraints, i.e.,

$$\varphi = \left(\bigwedge_{e' \in E^u, I \in \mathcal{I} \mid \tau^u(e', e) = I} \varphi(u, e', e, I) \right) \wedge \varphi^{edge} \text{ where,} \quad (5)$$

$$\varphi(u, e', e, I) = \begin{cases} \text{Msg}^{-1} \in I & \text{if } \exists p, q, p \neq q \text{ s.t. } e' <_{qp}^u e \\ Y_p^k \in I & \text{if } e, e' \in E_p^u, |\{e'' \in E_p^u \mid e' \leq_{pp}^u e'' <_{pp}^u e\}| = k \end{cases} \quad (6)$$

$$\text{and } \varphi^{edge} = \begin{cases} Y_p^1 \in I & \text{if } \beta_1 = \beta'_1(u', c') \text{ and for some } p \in Proc, \text{ we have} \\ & \text{EdgeC}((u', u), p) = I \text{ and } e = \min(E_p^u) \\ \text{true} & \text{otherwise} \end{cases} \quad (7)$$

- (T3) $\alpha' = \beta_1(u, c') \beta_2$, where $c' = c \uplus \{e\}$.

Observe as in the case of the automaton $\mathcal{C}_{\mathcal{G}}$, once the state and the enabled event which is to be executed are fixed, the transition that is taken and indeed the state reached after the transition are uniquely determined.

We can also observe that reachable states of this system satisfy some nice properties. To capture this we define the notion of a *valid state* of $\mathcal{C}_{\mathcal{G}}^{\#}$.

A state α of $\mathcal{C}_{\mathcal{G}}^{\#}$ is said to be *valid* if

- (V1) Every $\#$ symbol in α is surrounded by nodes from $ExtNodes$, i.e., there are no successive $\#$'s in α . Also α starts with $(\triangleright, \emptyset)$ and ends with $(\triangleleft, \emptyset)$.
- (V2) For any two consecutive extended nodes in α , there exists an edge between the nodes in G , i.e., for all $\alpha_1(u, c) \alpha_2(u', c') \preceq \alpha$, we have $u \rightarrow u'$ in G .
- (V3) Executed events in α are *downward closed*. By this we mean that the following two conditions are satisfied:
 - (a) For all $\alpha_1(u, c) \preceq \alpha$, if $e \in c$ and $e' \leq^u e$ then $e' \in c$.
 - (b) For all $\alpha_1(u, c) \alpha_2(u', c') \preceq \alpha$, if $e \in E_p^u, e' \in E_p^{u'}$ for some $p \in Proc$, then $e' \in c' \implies e \in c$.

Proposition 6. Every state of $\mathcal{C}_{\mathcal{G}}^{\#}$ reachable from the initial state is valid.

Proof. First note that the initial state is valid. Now, suppose α is valid and $\alpha \xrightarrow{\varphi, a} \alpha'$ we want to show that α' is valid as well. The first two properties follow from the node-insertion definition. The third follows from the definition of an enabled event. \square

We may note however that the converse is not true in general, i.e., a valid state need not always be reachable.

Lemma 7. $\mathcal{L}_{tw}(\mathcal{C}_{\mathfrak{G}}^{\#}) = \mathcal{L}_{tw}(\mathcal{C}_{\mathfrak{G}})$

Proof. We consider a morphism $\Psi : ExtNodes^* \rightarrow \Gamma^*$ defined by $(u, \emptyset) \mapsto \#$ and $(u, c) \mapsto (u, c)$ if $c \neq \emptyset$. We also define a reduction operation which acts on Γ^* and reduces consecutive multiple occurrences of $\#$ into a single $\#$. Formally, it is a rewrite operation where the rule is $\alpha_1 \# \# \alpha_2 \xrightarrow{red_{\#}} \alpha_1 \# \alpha_2$. Then, for a state $\alpha \in \Gamma^*$, we denote by $Red_{\#}(\alpha)$ the state that we reach by a maximal sequence of repeated applications of this rule. We denote by Υ , the function that, given a state $\tilde{s} \in Q_{\mathfrak{G}}$ of the global system $\mathcal{C}_{\mathfrak{G}}$, assigns the state of $\mathcal{C}_{\mathfrak{G}}^{\#}$ obtained as $\beta = (\triangleright, \emptyset) Red_{\#}(\Psi(\Phi(\tilde{s}))) (\triangleleft, \emptyset)$ where Φ is the stratification function defined earlier.

Now using the above definitions, we can relate accepting paths of the global semantics (i.e., of $\mathcal{C}_{\mathfrak{G}}$) and accepting paths of the automaton $\mathcal{C}_{\mathfrak{G}}^{\#}$. The equality of the languages $\mathcal{L}_{tw}(\mathcal{C}_{\mathfrak{G}}^{\#}) = \mathcal{L}_{tw}(\mathcal{C}_{\mathfrak{G}})$ follows immediately.

(\Leftarrow) Consider any global path of \mathfrak{G} , i.e.,

$$\tilde{s}_0 \xrightarrow{\varphi_1, a_1} \tilde{s}_1 \cdots \tilde{s}_{n-1} \xrightarrow{\varphi_n, a_n} \tilde{s}_n$$

where each $\tilde{s}_i = (\pi, C_i)$. For all $i \in \{0, \dots, n\}$, let $\beta_i = \Upsilon(\tilde{s}_i)$. We will show that

$$\beta_0 \xrightarrow{\varphi_1, a_1} \beta_1 \cdots \beta_{n-1} \xrightarrow{\varphi_n, a_n} \beta_n$$

is a path of $\mathcal{C}_{\mathfrak{G}}^{\#}$.

Since $\tilde{s}_0 = (\pi, C_0)$ is initial, we have $C_0 = \emptyset$ which implies that $\beta_0 = (\triangleright, \emptyset) \# (\triangleleft, \emptyset)$ which is the initial state of $\mathcal{C}_{\mathfrak{G}}^{\#}$. Fix $1 \leq i \leq n$ and let $\Phi(\tilde{s}_{i-1}) = (u_1, c_1) \cdots (u_m, c_m)$ where $\pi = u_1 \cdots u_m$. Now, the transition $\tilde{s}_{i-1} \xrightarrow{\varphi_i, a_i} \tilde{s}_i$ executes some event $(e, u_1 \cdots u_j)$ which is enabled in \tilde{s}_{i-1} . Then, $\tilde{s}_i = (\pi, C_i)$ with $C_i = C_{i-1} \uplus \{(e, u_1 \cdots u_j)\}$. There are two cases to consider:

- Either $c_j \neq \emptyset$. In this case, we observe that $\beta_{i-1} = \Upsilon(\tilde{s}_{i-1}) = \alpha_1(u_j, c_j)\alpha_2$ where we can write

$$\begin{aligned} \alpha_1 &= (\triangleright, \emptyset) Red_{\#}(\Psi((u_1, c_1) \cdots (u_{j-1}, c_{j-1}))) \\ \alpha_2 &= Red_{\#}(\Psi((u_{j+1}, c_{j+1}) \cdots (u_m, c_m))) (\triangleleft, \emptyset). \end{aligned}$$

Then, we observe that $(e, u_1 \cdots u_j)$ is enabled in \tilde{s}_{i-1} implies that $(e, \alpha_1(u_j, c_j))$ is enabled in β_{i-1} . Thus, there exists a transition of $\mathcal{C}_{\mathfrak{G}}^{\#}$ which executes this event, namely $\beta_{i-1} \xrightarrow{\varphi'_i, a_i} \alpha_1(u_j, c'_j)\alpha_2$ where $c'_j = c_j \uplus \{e\}$. From Conditions (2)–(4) and (5)–(7), we deduce that $\varphi' = \varphi$. Then, $\Phi(\tilde{s}_i) = (u_1, c_1) \cdots (u_j, c'_j) \cdots (u_m, c_m)$ by definition of C_i and so $\Upsilon(\tilde{s}_i) = \alpha_1(u_j, c'_j)\alpha_2 = \beta_i$.

- Or $c_j = \emptyset$. That is, the event being executed is on a node that is not present in β_{i-1} . Then, there was a gap in β_{i-1} instead and we can write $\beta_{i-1} = \alpha_1 \# \alpha_2$ where $\alpha_1 \# = (\triangleright, \emptyset) Red_{\#}(\Psi((u_1, c_1) \cdots (u_j, c_j)))$. Now if $c_{j-1} = \emptyset$, then we let $\beta' = \alpha_1 \#$ and else $\beta' = \alpha_1$. Similarly if $c_{j+1} = \emptyset$, then we let $\beta'' = \# \alpha_2$ and $\beta'' = \alpha_2$ otherwise. Then we can observe that $(e, \beta'(u_j, \emptyset))$ is enabled in $\beta'(u_j, \emptyset)\beta''$. Also, we have $\alpha_1 \# \alpha_2 \xrightarrow{u_j} \beta'(u_j, \emptyset)\beta''$ since Conditions (11), (12) and (13) hold. Indeed the latter two conditions follow from above, and if $\beta' = \alpha_1$ or $\beta'' = \alpha_2$, the presence of the edge in (12), (13) follows from the fact that the corresponding nodes are consecutive in π which is a path through G . Also if Condition (11) is violated this would contradict the downward-closed property of the cut C_i .

Thus there exists a transition in $\mathcal{C}_{\mathfrak{G}}^{\#}$, $\beta_{i-1} \xrightarrow{\varphi'_i, a_i} \beta_i = \beta'(u_j, c'_j)\beta''$ where $c'_j = \{e\}$. As above, we can conclude that $\varphi' = \varphi$. Now, $\Phi(\tilde{s}_i) = (u_1, c_1) \cdots (u_j, c'_j) \cdots (u_m, c_m)$ where $c'_j \neq \emptyset$ and so $\Upsilon(\tilde{s}_i) = \beta'(u_j, c'_j)\beta'' = \beta_i$.

Finally, since \tilde{s}_n is a final state of $\mathcal{C}_{\mathfrak{G}}$, $\beta_n = \Upsilon(\tilde{s}_n)$ is a final state as well as it is a sequence of completed nodes. This completes the proof in one direction.

(\Rightarrow) For the converse consider an accepting path in $\mathcal{C}_{\mathfrak{G}}^{\#}$,

$$\alpha_0 \xrightarrow{\varphi_1, a_1} \alpha_1 \cdots \alpha_{n-1} \xrightarrow{\varphi_n, a_n} \alpha_n$$

where each $\alpha_i \in \Gamma^*$.

Then, α_n is final if it is a sequence of completed nodes, which we write as $(\triangleright, \emptyset)(u_1, c_1) \cdots (u_m, c_m)(\triangleleft, \emptyset)$. Then we claim that $\pi = u_1 \cdots u_m$ is a path in G from an initial state to a final state. This follows since this state is reachable and therefore valid and so Property (V2) holds (and from the definition of $\triangleright, \triangleleft$). Then, we will construct the global run inductively maintaining the invariant $\Upsilon(\tilde{s}_i) = \alpha_i$ for all $i \in \{0, \dots, n\}$.

At $i = 0$, $\tilde{s}_0 = (\pi, C_0) = (\pi, \emptyset)$ and $\Upsilon(\tilde{s}_0) = (\triangleright, \emptyset) \# (\triangleleft, \emptyset) = \alpha_0$. Suppose we have defined till $\tilde{s}_{i-1} = (\pi, C_{i-1})$ such that $\Upsilon(\tilde{s}_{i-1}) = \alpha_{i-1}$, with $\Phi(\tilde{s}_{i-1}) = (u_1, c_1) \cdots (u_m, c_m)$. Consider $\alpha_{i-1} \xrightarrow{\varphi_i, a_i} \alpha_i$. Then again we have two cases:

- either the transition executes the event $(e, \beta'_1(u_j, c_j))$ which is enabled in $\alpha_{i-1} = \beta'_1(u_j, c_j)\beta'_2 = \beta'$ where we let $\beta'_1 = (\triangleright, \emptyset) Red_{\#}(\Psi((u_1, c_1) \cdots (u_{j-1}, c_{j-1})))$ and $\beta'_2 = Red_{\#}(\Psi((u_{j+1}, c_{j+1}) \cdots (u_m, c_m))) (\triangleleft, \emptyset)$.
- Or the transition inserts a node and then executes an enabled event, i.e., $\alpha_{i-1} = \beta_1 \# \beta_2$ and $\beta_1 \# \beta_2 \xrightarrow{u_j} \beta'_1(u, \emptyset)\beta'_2 = \beta'$ and $(e, \beta'_1(u, \emptyset))$ is enabled in β' . Then $\beta'_1 \in \{\beta_1, \beta_1 \#\}$ and $\beta'_2 \in \{\beta_2, \#\beta_2\}$. In π consider the first occurrence of u , say u_j , which has no executed event in \tilde{s}_{i-1} , i.e., $C_{i-1} \cap (E^{u_1 \cdots u_j} \setminus E^{u_1 \cdots u_{j-1}}) = \emptyset$. Thus, in this case, $c_j = \emptyset$.

Now, in both of the above cases, we claim that $(e, u_1 \cdots u_j)$ is enabled in \tilde{s}_{i-1} . Suppose not, choose a maximal event $(e', u_1 \cdots u_{j'})$ which was not executed in \tilde{s}_{i-1} , such that $(e', u_1 \cdots u_{j'}) <^\pi (e, u_1 \cdots u_j)$. This implies $j' \leq j$ and in fact, we have $j' < j$ since otherwise $e' <^{u_j} e$ which contradicts enabledness of $(e, \beta'_1(u_j, c_j))$ in β' . Thus, e' belongs to the same process as e . But then, there can't be any executed event in node $u_{j'}$, since if there was, the node would occur in α_{i-1} and so would contradict the fact that $(e, \beta'_1(u_j, c_j))$ is enabled in β' by violating Condition (E3). Now, if there was no executed event it would have been replaced by $\#$ in α_{i-1} . But then since we are simulating an accepting run of $C_{\mathcal{G}}^\#$, at some later transition, node $u_{j'}$ will be inserted in this $\#$. At that stage, we would violate Condition (I1) for node insertion since the process has seen an event, namely e to the right. Thus, we have a contradiction.

Once again, the existence of the enabled event immediately implies that there exists a transition that executes it in $C_{\mathcal{G}}$, namely $\tilde{s}_{i-1} \xrightarrow{\varphi_i, a_i} \tilde{s}_i$ such that $C_i = C_{i-1} \uplus \{(e, u_1 \cdots u_j)\}$. Then we can also observe that $\Phi(\tilde{s}_i) = (u_1, c_1) \cdots (u_j, c'_j) \cdots (u_m, c_m)$ and $c'_j = c_j \uplus \{e\}$. Thus, we conclude that $\Upsilon(\tilde{s}_i) = \alpha_i$. \square

In fact, we can strengthen the above lemma slightly without much change in the proof. If we restrict the above automaton to states that are both reachable and co-reachable even then the result holds. It turns out that this property of co-reachability is easy to capture in the automaton. Formally, we call a state α *completable* if whenever $\alpha = \alpha_1(u, c)\#(v, c')\alpha_2$, there is $\beta \in V^+$ such that $u\beta v$ is a path in G and $OProc(\beta) \cap EProc((v, c')\alpha_2) = \emptyset$.

Corollary 8. Consider the timed automaton obtained from $C_{\mathcal{G}}^\#$ by restricting to valid and completable states. Then, the timed language of this automaton is $\mathcal{L}_{tw}(C_{\mathcal{G}}^\#)$.

5.3. Removing completed nodes

As we mentioned earlier, from a state α we would like to obtain a finite abstraction of α , such that

1. the set of events left to be executed are the same,
2. if $\alpha = \alpha_1\#\alpha_2$ where $\alpha_2 \in \Gamma^*$, then we want to preserve the information about the processes in $EProc(\alpha_2)$ so that if some nodes in α_2 are deleted we still know which processes must not be inserted in this gap.
We accomplish this by enlarging the alphabet of nodes and $\#$ symbol with subsets of processes $P \subseteq Proc$. The idea is that this set P keeps track of the processes that are not allowed to participate in a node inserted on the left.
3. we preserve (do not throw away) the nodes around a $\#$ occurrence in α and also nodes that start an edge constraint which needs to be verified later.

Formally, the set of states of our new automaton $C_{\mathcal{G}}^{fin}$ will be a finite subset of Π^* where $\Pi = \Gamma \cup 2^{Proc}$. Then, in our definition of the morphisms earlier we need to add $OProc(P) = P$, $EProc(P) = P$. Now, we define the reduction as a rewrite operation $\alpha \xrightarrow{redn} \alpha'$. There are two rewrite rules:

- (R1) The first says that if two process sets are together they can be merged, i.e., $\alpha_1 P P' \alpha_2 \xrightarrow{redn} \alpha_1 (P \cup P') \alpha_2$.
- (R2) Now, we define the rule that removes a completed extended-node (v, c) and replaces it by the set of processes participating in v , i.e., we have $\alpha_1(v, c)\alpha_2 \xrightarrow{redn} \alpha_1 OProc(v)\alpha_2$ if the following hold:
- (C2.1) $v \in V$, $\varepsilon \neq \alpha_1 \notin \Pi^*\#$, $\varepsilon \neq \alpha_2 \notin \#\Pi^*$ i.e., the node v is not next to a gap or at the beginning or the end.
- (C2.2) $c = E^v$, i.e., all events in the node have been completed,
- (C2.3) and one of the two following cases hold:
- (i) either $\alpha_2 \in (v', c')\Pi^*$ and then for each $p \in Proc$ we must have either $E_p^v = \emptyset$ or $E_p^{v'} = \emptyset$ or $(c' \cap E_p^{v'}) \neq \emptyset$. In other words, if the first symbol of α_2 is an extended node (v', c') and there is an event in both E_p^v and $E_p^{v'}$, then some event in $E_p^{v'}$ has occurred and so, the edge constraint has indeed been checked,
 - (ii) or $\alpha_2 \in 2^{Proc}\Pi^*$ in which case there is no unchecked edge constraint.

Remark 1. We can observe that, in some sense, the negation of Rule (R2) is an invariant of the reduction operation. More precisely, let $\alpha = \alpha_1(u, c)\alpha_2$ be such that we cannot apply Rule (R2) to remove node (u, c) (given by its occurrence $\alpha_1(u, c) \leq \alpha$) and suppose $\alpha \xrightarrow{redn} \alpha'$. This, of course, implies that (u, c) (or rather, this occurrence of (u, c)) is present in α' as well. Then, we can easily check that we cannot apply Rule (R2) to remove this node in α' either.

Lemma 9. The rewrite system defined by the operation \xrightarrow{redn} is confluent.

Proof. Indeed it is easy to see that if the reduction rules apply on non-adjacent segments in a path, then they can be executed in any order. For instance, for $\beta \neq \varepsilon$, if we have $\alpha(u, c)\beta P P' \gamma \xrightarrow{redn} \alpha P'' \beta P P' \gamma$ where $P'' = OProc(u)$ and $\alpha(u, c)\beta P P' \gamma \xrightarrow{redn} \alpha(u, c)\beta(P \cup P')\gamma$, then of course $\alpha P'' \beta P P' \xrightarrow{redn} \alpha P'' \beta(P \cup P')\gamma$ and $\alpha(u, c)\beta(P \cup P') \xrightarrow{redn} \alpha P'' \beta(P \cup P')\gamma$. The interesting case is when two reduction rules apply on adjacent segments. Again, we may consider

several subcases. If one of the reductions is by applying Rule (R1), then it is easy to handle since, in some sense, this rule does not depend on the context (i.e., the surrounding nodes/symbols). We now explicitly illustrate the subcase when we have two applications of Rule (R2) on adjacent nodes, i.e., let

- $\alpha(u, c)(u', c')\beta \xrightarrow{\text{redn}} \alpha(u, c)P'\beta$ where $P' = OProc(u')$ and
- $\alpha(u, c)(u', c')\beta \xrightarrow{\text{redn}} \alpha P(u', c')\beta$ where $P = OProc(u)$.

Then, from the first reduction we get $c' = E^{u'}$, $\varepsilon \neq \beta \notin \Pi^*$ and Condition (C2.3) holds with $\alpha_2 = \beta$. Using these and observing that $\alpha P \notin \Pi^*\#$, we can conclude that the first reduction is applicable after the second, i.e., $\alpha P(u', c')\beta \xrightarrow{\text{redn}} \alpha P P'\beta$. From the second reduction we have $c = E^u$ and $\varepsilon \neq \alpha \notin \Pi^*\#$. Now from these and the fact that Condition (C2.3)(ii) holds, we can conclude that the second reduction is applicable after the first, i.e., $\alpha(u, c)P'\beta \xrightarrow{\text{redn}} \alpha P P'\beta$. \square

Using the above lemma we can conclude that, from any state α after any maximal sequence of reductions, we reach the same state which we denote by $Red(\alpha)$. Note that if $\alpha \xrightarrow{\text{redn}} \alpha'$, then $EProc(\alpha) = EProc(\alpha')$ and therefore, $EProc(\alpha) = EProc(Red(\alpha))$. In fact, from confluence, we derive some useful properties of the reduction operation,

(P1) $Red(\alpha_1 \# \alpha_2) = Red(\alpha_1) \# Red(\alpha_2)$.

(P2) $Red(\alpha_1 \alpha_2) = Red(Red(\alpha_1) \alpha_2) = Red(\alpha_1 Red(\alpha_2)) = Red(Red(\alpha_1) Red(\alpha_2))$

(P3) Let $\alpha = \alpha_1(u, c)\alpha_2$ be such that this (u, c) (given by its occurrence $\alpha_1(u, c)$) cannot be reduced in α , i.e., Rule (R2) cannot be applied. Then $Red(\alpha) = \gamma_1(u, c)\gamma_2$ where $\gamma_1(u, c) = Red(\alpha_1(u, c))$ and $(u, c)\gamma_2 = Red((u, c)\alpha_2)$.

Proof. The first two properties are self-evident. For the third, using Remark 1 we deduce that (u, c) is not deleted during the reductions. Now, let $\gamma_1(u, c) = Red(\alpha_1(u, c))$ and $(u, c)\gamma_2 = Red((u, c)\alpha_2)$. Then applying Property (P2) twice on α , we obtain $Red(\alpha_1(u, c)\alpha_2) = Red(Red(\alpha_1(u, c))\alpha_2) = Red(\gamma_1(u, c)\alpha_2) = Red(\gamma_1 Red((u, c)\alpha_2)) = Red(\gamma_1(u, c)\gamma_2)$. Now since $\gamma_1(u, c)$ and $(u, c)\gamma_2$ are already in reduced form and (u, c) cannot be deleted in $Red(\alpha)$, we obtain $Red(\gamma_1(u, c)\gamma_2) = \gamma_1(u, c)\gamma_2$. \square

The set of final states of $\mathcal{C}_{\mathfrak{G}}^{\text{fin}}$ are all states of the form $(\triangleright, \emptyset)P(\triangleleft, \emptyset)$ where $P \subseteq Proc$.

In the definition of a transition of $\mathcal{C}_{\mathfrak{G}}^{\text{fin}}$ we replace the final condition (T3) with the following condition:

(T3') $\alpha' = Red(\beta_1(u, c')\beta_2)$ where $c' = c \uplus \{e\}$.

Observe indeed that the node deleted by the above transition, need not necessarily be the one that had an event executed (but instead be one of its neighboring nodes). Now, if we maintain the rest of the definition of a transition of $\mathcal{C}_{\mathfrak{G}}^{\text{fin}}$ to be the same as a transition of $\mathcal{C}_{\mathfrak{G}}^{\#}$, we can prove that $\mathcal{C}_{\mathfrak{G}}^{\text{fin}}$ is a finite MSC-ECA which accepts the same timed language as $\mathcal{C}_{\mathfrak{G}}^{\#}$. We can also observe that in all reachable states of $\mathcal{C}_{\mathfrak{G}}^{\text{fin}}$, Properties (V1), (V2) and (V3) continue to hold with the enlarged alphabet Γ . In other words, for any state β of $\mathcal{C}_{\mathfrak{G}}^{\text{fin}}$, it is still the case that (1) no two #'s can occur consecutively in β as they would be either separated by an extended node or a set of processes (if a node was removed by application of Rule (R2)); (2) for any two consecutive extended nodes in β , there is an edge between the respective nodes in G and (3) executed events in β are downward closed (as defined in (V3)). Indeed the latter two properties hold as they are not affected by removal of nodes.

Lemma 10. *If \mathfrak{G} is locally synchronized, then $\mathcal{C}_{\mathfrak{G}}^{\text{fin}}$ as defined above is a finite MSC-ECA.*

Proof. We show that if \mathfrak{G} is locally synchronized, then the number of states of $\mathcal{C}_{\mathfrak{G}}^{\text{fin}}$ is finite. For this, it is enough to show that the length of each reachable, completable state of $\mathcal{C}_{\mathfrak{G}}^{\text{fin}}$ is bounded. Note that by definition in every state in every extended node there is at least one executed event. We begin with some properties about a loop in a state which follow from the locally synchronized assumption.

Claim 2. *Let $\alpha(u, c)\beta(u, c')\gamma$ be a valid completable state of $\mathcal{C}_{\mathfrak{G}}^{\text{fin}}$. If $(u, c)\beta$ is not completely executed or if $\#$ occurs in β , then we have $EProc((u, c')\gamma) \subsetneq EProc((u, c)\beta(u, c')\gamma)$.*

Proof. First, since $\alpha(u, c)\beta(u, c')\gamma$ is completable, for each occurrence of $\#$ in β , there exists $u_1 \cdots u_n \in V^*$ in G such that if we replace the $\#$ by $(u_1, \emptyset) \cdots (u_n, \emptyset)$, then we obtain a path β' such that $\alpha(u, c)\beta'(u, c')\gamma$ is a valid state.

Now, we can write $(u, c)\beta' = \beta_1(v, c'')\beta_2$ with $c'' \subsetneq E^v$. This follows, since either there is a $\#$ in β , and so for any node (v, c'') on the path inserted we have $c'' = \emptyset$, or else $\beta' = \beta$ and by assumption $(u, c)\beta$ is not completely executed; so there exists some node v such that $c'' \subsetneq E^v$.

Now, let $e' \in (E^v \setminus c')$ such that $e \in E_p^v$ for some p' and also, let $e \in c$ such that $e \in E_p^u$ for some $p \in Proc$. Consider the path $\widehat{\beta}'$ in G , obtained by restricting β' to its first component. Now, as \mathfrak{G} is locally synchronized, in the communication graph of $M^{u\widehat{\beta}'}$ there exists a path from p' to p . Then let this path be $p' = p_0 \rightarrow p_1 \rightarrow \dots \rightarrow p_n = p$ for some $n \geq 1$. We call a process q good if there is an executed event and an unexecuted event on q in $(u, c)\beta'$. If q is good, then $q \in (EProc((u, c)\beta') \setminus EProc((u, c')\gamma))$. We will now show that there is some good process $q \in \{p_0, \dots, p_n\}$.

Suppose, $p_n = p$ has an unexecuted event in $(u, c)\beta'$ then it is good and we are done. Otherwise, p must have completed its events in $(u, c)\beta'$ and so it must have received a message from p_{n-1} . Therefore, p_{n-1} has also taken part in $(u, c)\beta'$ since it must have sent the message that was received by p_n . Now if p_{n-1} has another event in $(u, c)\beta'$ which is unexecuted, then it is good and again we are done. Otherwise, we repeat this argument till we reach an executed event in $p_0 = p'$. But this implies that p' is good and so we are done. \square

Claim 3. *If $\alpha(u, c)\beta(u, c')\gamma$ is a valid state such that $(u, c)\beta(u, c')$ is completely executed and β has no #, then $\alpha = \alpha'\#$.*

Proof. Since $(u, c)\beta(u, c')$ is completely executed, the first occurrence of node u , i.e., (u, c) would have been deleted unless $\alpha = \alpha'\#$ or $\beta = \#\beta'$. But since β does not contain # the latter case is not possible and so we are done. \square

From the above claim we can conclude that after every two occurrences of node u in a path, there must exist a # or the segment is not completely executed. Then, along with Claim 2 this implies that we can bound the number of occurrences of a node u in a path by $2|Proc|$. From which we can conclude that we have a bound of $(2|Proc|)|V|$ on the number of extended nodes in a path. But we know that each # or $P \subseteq Proc$ must have a node $u \in V$ next to it on the left, so we can conclude that the length of the path is $\mathcal{O}(|Proc||V|)$. Thus $\mathcal{C}_{\mathfrak{G}}^{fin}$ is finite. \square

Now, we will show that the timed language accepted by $\mathcal{C}_{\mathfrak{G}}^{fin}$ is the same as the timed language accepted by $\mathcal{C}_{\mathfrak{G}}^{\#}$. We will accomplish this by defining a bisimulation relation \sim between the states of the abstract automata $\mathcal{C}_{\mathfrak{G}}^{\#}$ and $\mathcal{C}_{\mathfrak{G}}^{fin}$, i.e., a binary relation \sim between states of $\mathcal{C}_{\mathfrak{G}}^{\#}$ and $\mathcal{C}_{\mathfrak{G}}^{fin}$ such that:

- if $\alpha \xrightarrow{\varphi, a} \alpha'$ is a transition of $\mathcal{C}_{\mathfrak{G}}^{\#}$ and $\alpha \sim \beta$ for some β a state of $\mathcal{C}_{\mathfrak{G}}^{fin}$, then there exists a state β' of $\mathcal{C}_{\mathfrak{G}}^{fin}$ and a transition $\beta \xrightarrow{\varphi, a} \beta'$ of $\mathcal{C}_{\mathfrak{G}}^{fin}$ such that $\alpha' \sim \beta'$.
- conversely, if $\beta \xrightarrow{\varphi, a} \beta'$ is a transition of $\mathcal{C}_{\mathfrak{G}}^{fin}$ and $\alpha \sim \beta$, then there exists a state α' of $\mathcal{C}_{\mathfrak{G}}^{\#}$ and a transition $\alpha \xrightarrow{\varphi, a} \alpha'$ of $\mathcal{C}_{\mathfrak{G}}^{\#}$ such that $\alpha' \sim \beta'$.

From this, we will be able to conclude that their accepting paths (and therefore their timed languages) coincide. We define the relation \sim between states of $\mathcal{C}_{\mathfrak{G}}^{\#}$ and $\mathcal{C}_{\mathfrak{G}}^{fin}$:

$$\alpha \sim \beta \text{ if } \beta = Red(\alpha) \quad (8)$$

Now, we have the lemma,

Lemma 11. \sim is a bisimulation on abstract automata $\mathcal{C}_{\mathfrak{G}}^{\#}$ and $\mathcal{C}_{\mathfrak{G}}^{fin}$.

Proof. Let α be a state of $\mathcal{C}_{\mathfrak{G}}^{\#}$ and β a state of $\mathcal{C}_{\mathfrak{G}}^{fin}$ such that $\alpha \sim \beta$, i.e., $\beta = Red(\alpha)$.

(\implies) In one direction we start from a move $\alpha \xrightarrow{\varphi, a} \alpha'$ in $\mathcal{C}_{\mathfrak{G}}^{\#}$ and show that there is a move $\beta \xrightarrow{\varphi, a} \beta'$ in $\mathcal{C}_{\mathfrak{G}}^{fin}$, where $\beta' = Red(\alpha')$. There are two broad cases to consider depending on whether the transition in $\mathcal{C}_{\mathfrak{G}}^{\#}$ extends the path or not.

- Suppose the path is extended. Then, we have $\alpha = \alpha_1\#\alpha_2 \xrightarrow{u} \alpha'_1(u, \emptyset)\alpha'_2 = \alpha''$ where $\alpha'_1 \in \{\alpha_1, \alpha_1\#\}$ and $\alpha'_2 \in \{\alpha_2, \#\alpha_2\}$. Also, there exists an extended event $(e, \alpha'_1(u, \emptyset))$ enabled in α'' such that $\alpha' = \alpha'_1(u, c')\alpha'_2$ where $c' = \{e\}$. Then, we observe that
 1. we can write $\beta = \beta_1\#\beta_2$ where $\beta_1 = Red(\alpha_1)$ and $\beta_2 = Red(\alpha_2)$. This follows by Property (P1).
 2. we have $\beta_1\#\beta_2 \xrightarrow{u} \beta'_1(u, \emptyset)\beta'_2 = \beta''$ where $\beta'_1 \in \{\beta_1, \beta_1\#\}$ and $\beta'_2 \in \{\beta_2, \#\beta_2\}$. Further $\beta'_1 = \beta_1$ if and only if $\alpha'_1 = \alpha_1$ and $\beta'_2 = \beta_2$ if and only if $\alpha'_2 = \alpha_2$. The existence of this node insertion move follows from the node insertion in $\mathcal{C}_{\mathfrak{G}}^{\#}$ above since we have $OProc(u) \cap EProc(\alpha_2) = \emptyset$, which implies that $OProc(u) \cap EProc(\beta_2) = \emptyset$ (since $\beta_2 = Red(\alpha_2)$). Notice that we also have for $i \in \{1, 2\}$, $\beta'_i = Red(\alpha'_i)$ since $\beta_i = Red(\alpha_i)$.
 3. $(e, \beta'_1(u, \emptyset))$ is enabled in β'' . Indeed, Conditions (E1), (E2) hold since they hold for $(e, \alpha'_1(u, \emptyset))$. And if there exists $(\hat{e}, \hat{\beta}(\hat{u}, \hat{c}))$ such that e, \hat{e} are on the same process, $\hat{\beta}(\hat{u}, \hat{c}) \leq \beta'_1 = Red(\alpha'_1)$ and $\hat{e} \notin \hat{c}$, then $\hat{\beta}'(\hat{u}, \hat{c}) \leq \alpha'_1$ for some $\hat{\beta}'$. This contradicts the fact that $(e, \alpha'_1(u, \emptyset))$ is enabled in α'' . Therefore Condition (E3) holds as well.

Then by definition of a transition, we have $\beta \xrightarrow{\varphi, a} \beta' = Red(\beta'_1(u, \{e\})\beta'_2)$ which executes this enabled event in $\mathcal{C}_{\mathfrak{G}}^{fin}$.

Now, we show that the same guard is used, i.e., $\varphi' = \varphi$. For this, observe that $\varphi = \varphi^{edge}$ and $\varphi' = \varphi'^{edge}$ since there are no local-constraints. Now $\varphi^{edge} = (Y_p^1 \in I)$ for some $p \in Proc$ if and only if $e = \min(E_p^u)$, $\alpha'_1 = \alpha_1 = \alpha_1''(u', c'')$, $EdgeC((u', u), p) = I$. But now, the node u' cannot be removed during the reduction of α since it is next to a #, so we have $\beta'_1 = \beta_1 = \beta_1''(u', c'')$ which implies that we have the constraint $\varphi'^{edge} = (Y_p^1 \in I)$.

Finally, we will be done with this case if we show that $Red(\alpha') = \beta'$. We have $\beta' = Red(\beta'_1(u, c')\beta'_2) = Red(Red(\alpha'_1)(u, c')Red(\alpha'_2))$. But by Property (P2) this is equal to $Red(\alpha'_1(u, c')\alpha'_2) = Red(\alpha')$ and so we are done.

- Else, it was not extended then there exists an enabled event $(e, \alpha_1(u, c))$ in α which is executed in the transition $\alpha \xrightarrow{\varphi, a} \alpha'$, where $\alpha = \alpha_1(u, c)\alpha_2$, $\alpha' = \alpha_1(u, c')\alpha_2$ with $c' = c \uplus \{e\}$ and φ is defined by Eq. (T2). Then (u, c) is not completely executed and so it cannot be reduced in α . Thus by Property (P3), $\beta = Red(\alpha) = \gamma_1(u, c)\gamma_2$, where $\gamma_1(u, c) = Red(\alpha_1(u, c))$ and $(u, c)\gamma_2 = Red((u, c)\alpha_2)$. Now, $(e, \gamma_1(u, c))$ is enabled in β , since $(e, \alpha_1(u, c))$ was enabled in α , and Conditions (E1), (E2) and Condition (E3) follow as in the previous case. That is, if there exists $(\hat{e}, \hat{\beta}(\hat{u}, \hat{c}))$ such that $\hat{\beta}(\hat{u}, \hat{c}) \leq \gamma_1$, then $\hat{\beta}'(\hat{u}, \hat{c}) \leq \alpha_1$ for some $\hat{\beta}'$.

Thus, there exists a transition $\beta \xrightarrow{\varphi, a} \beta'$ that executes $(e, \gamma_1(u, c))$ in $C_{\mathfrak{G}}^{fin}$. Again we check that $\varphi' = \varphi$. This follows as in the previous case except that we also need to check local constraints in φ' . But as the guards are local to the node (u, c) which is not deleted in β , this follows directly from the definition.

It remains to show that $Red(\alpha') = \beta'$. Since $\alpha' = \alpha_1(u, c')\alpha_2$ is such that $c \subseteq c' \subseteq E^u$, we have $Red(\alpha') = Red(\alpha_1(u, c')\alpha_2) = Red(\gamma_1(u, c')\gamma_2) = \beta'$. This follows because, firstly, every reduction that can be performed on α can be performed on α' (since $\alpha = \alpha_1(u, c)\alpha_2$ and $\alpha' = \alpha_1(u, c')\alpha_2$ and $c' \supseteq c$). Then, by Property (P2) we can perform reductions in any order, so we can choose to compute $Red(\alpha')$, by first performing the same sequence of reductions as was used to compute $Red(\alpha)$. Now, since $Red(\alpha) = (\gamma_1(u, c)\gamma_2)$, by applying the same sequence on α' we obtain $(\gamma_1(u, c')\gamma_2)$. But indeed, this may not be a maximal sequence of reductions since in $(\gamma_1(u, c')\gamma_2)$ it may be possible to remove further nodes by reduction (due to events in $(c' \setminus c)$). Thus, we conclude that $Red(\alpha') = Red(\gamma_1(u, c')\gamma_2)$, which by the definition of a transition (Condition (T3')) is equal to β' .

(\Leftarrow) For the other direction, the result follows by observing that the enabled event that gets executed in the infinite system $C_{\mathfrak{G}}^{\#}$ is obtained from the corresponding event in the finite system $C_{\mathfrak{G}}^{fin}$. More formally, we assume that $\beta \xrightarrow{\varphi, a} \beta'$ is a transition in $C_{\mathfrak{G}}^{\#}$ and show that there is a transition $\alpha \xrightarrow{\varphi, a} \alpha'$ in $C_{\mathfrak{G}}^{\#}$.

Let the transition in $C_{\mathfrak{G}}^{fin}$ execute the event $(e, \beta_1(u, c))$ enabled in $\beta = \beta_1(u, c)\beta_2$. Indeed there is another case where the executed event is not in β and so we need to perform a node insertion before we obtain the enabled event. But as this case follows by the same arguments (and indeed, is simpler due to presence of #), we only consider the first case.

Let $\alpha_1(u, c')$ be the least prefix of α such that $e \notin c'$. Then (u, c') is not removed by the reduction operation. Since $\beta = Red(\alpha)$ and $(e, \beta_1(u, c))$ is enabled in β , we deduce from (E3) that $c' = c$ and $Red(\alpha_1(u, c)) = \beta_1(u, c)$. Now we claim that $(e, \alpha_1(u, c))$ is enabled in α . Conditions (E1), (E2) hold since they hold for $(e, \beta_1(u, c))$. Suppose Condition (E3) did not hold, then for $p \in Proc$ such that $e \in E_p^u$, there exists an event $(e', \hat{\alpha}_1(v, c'))$ with $e' \in (E_p^v \setminus c')$ and $\hat{\alpha}_1(v, c') \leq \alpha_1$. Again, $Red(\hat{\alpha}_1(v, c')) = \hat{\beta}_1(v, c') < \beta_1$ (since (v, c') cannot be removed by reductions). But then $e' \in (E_p^v \setminus c')$ is a contradiction of Condition (E3) on $(e, \beta_1(u, c))$. Thus all the conditions hold and $(e, \alpha_1(u, c))$ is enabled in α .

Thus, we can conclude that there is a transition that executes $(e, \alpha_1(u, c))$ in $C_{\mathfrak{G}}^{\#}$, i.e., $\alpha \xrightarrow{\varphi, a} \alpha'$. The fact that $\varphi' = \varphi$ and $\beta' = Red(\alpha')$ follows exactly as in the previous direction so we are done. \square

Corollary 12. $\mathcal{L}_{tw}(C_{\mathfrak{G}}^{fin}) = \mathcal{L}_{tw}(C_{\mathfrak{G}}^{\#})$

Proof. From the above bisimulation at the symbolic level of paths, we deduce easily that the timed language of $C_{\mathfrak{G}}^{\#}$ is equal to the timed language of $C_{\mathfrak{G}}^{fin}$. \square

Proof of Theorem 4. Given a locally synchronized TCMSC \mathfrak{G} , consider the finite MSC-ECA $C_{\mathfrak{G}}^{fin}$. Then, by using the above corollary, Lemma 7 and Lemma 5, we conclude that $\mathcal{L}_{tw}(C_{\mathfrak{G}}^{fin}) = \mathcal{L}_{tw}(\mathfrak{G})$. \square

6. Solving the model checking problem

Now, we are in a position to solve the model checking problem.

Theorem 13. For a locally synchronized TCMSC \mathfrak{G} and a timed automaton \mathcal{A} , the model checking problem $\mathcal{L}_{tw}(\mathcal{A}) \subseteq \mathcal{L}_{tw}(\mathfrak{G})$ is decidable, i.e., it is decidable to check if for all timed words σ generated by \mathcal{A} there exists some \mathfrak{M} specified by \mathfrak{G} such that σ is a linearization of a TMSCT which realizes \mathfrak{M} .

Proof. We have to prove that $\mathcal{L}_{tw}(\mathcal{A}) \cap (TW_{Act} \setminus \mathcal{L}_{tw}(\mathfrak{G})) = \emptyset$. By Theorem 4 we can construct an MSC-ECA \mathcal{C} such that $\mathcal{L}_{tw}(\mathcal{C}) = \mathcal{L}_{tw}(\mathfrak{G})$. Using the complementation construction of Section 4.1 we can build a deterministic and complete MSC-ECA $\mathcal{C}' = C_{\mathfrak{G}}^{univ}$ such that by Corollary 2 we have $\mathcal{L}_{tw}(\mathcal{C}') = TW_{Act} \setminus \mathcal{L}_{tw}(\mathcal{C}) = TW_{Act} \setminus \mathcal{L}_{tw}(\mathfrak{G})$.

Since \mathfrak{G} is locally synchronized, there is a bound $B > 0$ such that each timed word $\sigma \in \mathcal{L}_{tw}(\mathfrak{G})$ is wwff and B -bounded: $\mathcal{L}_{tw}(\mathfrak{G}) \subseteq TW_{Act}^{B,wwf}$. Consider the timed automaton $\mathcal{B}_{C'}^B$ associated with C' and the bound B by the construction of Section 4.2. For final states of $\mathcal{B}_{C'}^B$, we choose $F' \cup F''$ as defined in Proposition 3. We get $\mathcal{L}_{tw}(\mathcal{B}_{C'}^B) = (TW_{Act} \setminus TW_{Act}^{B,wwf}) \cup (\mathcal{L}_{tw}(C') \cap TW_{Act}^{B,wwf}) = (TW_{Act} \setminus TW_{Act}^{B,wwf}) \cup (TW_{Act}^{B,wwf} \setminus \mathcal{L}_{tw}(\mathfrak{G}))$. Using $\mathcal{L}_{tw}(\mathfrak{G}) \subseteq TW_{Act}^{B,wwf}$ we deduce $\mathcal{L}_{tw}(\mathcal{B}_{C'}^B) = TW_{Act} \setminus \mathcal{L}_{tw}(\mathfrak{G})$.

Hence, the model checking problem is reduced to checking emptiness of the intersection of two timed automata, \mathcal{A} and $\mathcal{B}_{C'}^B$, which is indeed decidable. \square

7. Solving the coverage problem

Let us fix a TCMSC $\mathfrak{G} = (G, \mathcal{L}^{TC}, \Phi, EdgeC)$, where $G = (V, \rightarrow, v_{in}, V_F)$ is a graph. Let \mathcal{A} be a timed automaton over Act . We recall that the *coverage problem* for \mathfrak{G} and \mathcal{A} is to determine whether for each TCMSC $\mathfrak{M} \in \mathcal{L}_{TC}(\mathfrak{G})$, there exists $w \in \mathcal{L}_{tw}(\mathcal{A})$ such that $w \in \mathcal{L}_{tw}(\mathfrak{M})$.

Our strategy for the solution is as follows. Note that every TCMSC $\mathfrak{M} \in \mathcal{L}_{TC}(\mathfrak{G})$ is defined by some path in G . Moreover if two paths define the same TCMSC then either both or neither are witnessed by \mathcal{A} . We record the set of paths in G that can be witnessed by \mathcal{A} by synchronizing \mathcal{A} with $\mathcal{C}_{\mathfrak{G}}^{fin}$. Comparing this set to the set of all paths in G , we obtain a solution to the coverage problem.

For recording a path, our strategy is to *emit* the sequence of nodes visited by the path. However, if a process, say p , does not participate in a node u but does participate in the next node v in the path, then by this strategy, we may emit v before u . Thus, we additionally need to handle the out of order emission of node labels. The problem is that, instead of a single node u , we could have a loop (which is still locally synchronized) in which p does not participate. In this case, it becomes very hard to recover the actual path traversed from the sequence of nodes emitted.

One way to get around this problem is by introducing a structural restriction on the TCMSC forbidding such behavior. We propose a natural restriction that handles this in the following section.

Event-saturated TCMSCs. A locally synchronized TCMSC \mathfrak{G} is said to be *event-saturated* if in every node of \mathfrak{G} there is an event present on each process.

Now we see why coverage is easier to establish for event-saturated TCMSCs. Intuitively, every move $\alpha \xrightarrow{\varphi, a} \alpha'$ in $\mathcal{C}_{\mathfrak{G}}^{\#}$ or $\mathcal{C}_{\mathfrak{G}}^{fin}$ between reachable and completable states, either executes an event in α or extends the current path by exactly one node. Formally,

Proposition 14. *Let \mathfrak{G} be an event-saturated TCMSC and $\mathcal{C}_{\mathfrak{G}}^{\#}$ and $\mathcal{C}_{\mathfrak{G}}^{fin}$ be the associated MSC-ECA as defined in Sections 5.2 and 5.3, respectively. Then, all reachable and completable states of $\mathcal{C}_{\mathfrak{G}}^{\#}$ and $\mathcal{C}_{\mathfrak{G}}^{fin}$ are of the form $(\triangleright, \emptyset)\alpha\#(\triangleleft, \emptyset)$ or $(\triangleright, \emptyset)\alpha(\triangleleft, \emptyset)$ where $\alpha \in (ExtNodes)^*$. Further, if $\alpha = (u_1, c_1) \dots (u_m, c_m)$, then $u_1 \dots u_m$ is a path in G .*

Proof. For any node u , we have $OProc(u) = Proc$. Thus in any node insertion move $\alpha_1\#\alpha_2 \xrightarrow{u} \alpha'_1(u, \emptyset)\alpha'_2$, by Condition (11) we infer that $EProc(\alpha_2) = \emptyset$ which implies that $\alpha_2 = (\triangleleft, \emptyset)$. In addition, from the fact that the state is completable we obtain $\alpha'_1 = \alpha_1$ and $\alpha'_2 = \#(\triangleleft, \emptyset)$ or $\alpha'_2 = (\triangleleft, \emptyset)$. Thus, from this and by Conditions (12), (13) it follows that the node insertion extends the path with a single node. Thus, any move either executes an event in the current path or it is a node insertion which extends the path with a single node. Finally, when a reduction is applied in $\mathcal{C}_{\mathfrak{G}}^{fin}$, it always removes the leftmost node (which is not the endpoint $(\triangleright, \emptyset)$) in the current path. This follows from the definition of the reduction rules. Hence, we conclude that any completable state reached defines a path in G and the proposition follows. \square

Remark 2. As observed earlier, each run of $\mathcal{C}_{\mathfrak{G}}^{\#}$ or $\mathcal{C}_{\mathfrak{G}}^{fin}$ defines a path through G . In this case, as each process occurs in each node of G , we can further infer that each process visits all the nodes in the path traced out by the run in the same order.

Coverage for event-saturated TCMSCs. Recall that our proof strategy is to record the paths that \mathcal{A} can follow in \mathfrak{G} by constructing a product of \mathfrak{G} and \mathcal{A} . We enlarge the communication actions in Act to include the set of nodes in \mathfrak{G} . Then, we build a product of \mathcal{A} and the timed automaton obtained from $\mathcal{C}_{\mathfrak{G}}^{fin}$ thus synchronizing the runs of \mathcal{A} with the runs of $\mathcal{C}_{\mathfrak{G}}^{fin}$. The language of the resulting timed automaton would be the set of all runs of \mathcal{A} that are consistent with some run of \mathfrak{G} . Now, in this automaton, using our enlarged alphabet, we emit the nodes seen along these runs.

Finally, we use the region construction [6] to obtain an untimed regular language over $(Act \cup V)$, where V is the set of nodes of G in TCMSC \mathfrak{G} . This language projected onto the alphabet V precisely describes the set of all paths in \mathfrak{G} that are covered by some run of \mathcal{A} .

The product construction. Formally, given an event-saturated TCMSC \mathfrak{G} , we first apply Theorem 4 to construct a finite MSC-ECA $\mathcal{C} = \mathcal{C}_{\mathfrak{G}}^{fin}$ such that $\mathcal{L}_{tw}(\mathcal{C}) = \mathcal{L}_{tw}(\mathfrak{G})$. Now, since \mathfrak{G} is locally synchronized, there is a bound $B > 0$ such that each timed word $\sigma \in \mathcal{L}_{tw}(\mathfrak{G})$ is wwff and B -bounded: $\mathcal{L}_{tw}(\mathfrak{G}) \subseteq TW_{Act}^{B,wwf}$. Thus by the construction in Section 4.2, we obtain

the (finite) timed automaton \mathcal{B}_C^B associated with \mathcal{C} and bound B such that, by [Proposition 3](#) and the above, we have $\mathcal{L}_{tw}(\mathfrak{G}) = \mathcal{L}_{tw}(\mathcal{C}) = \mathcal{L}_{tw}(\mathcal{B}_C^B)$.

Now, by taking the product of \mathcal{A} and \mathcal{B}_C^B , we obtain the timed automaton $\mathcal{B}_{\mathcal{A},\mathfrak{G}}^{prod}$ accepting the intersection of the timed languages $\mathcal{L}_{tw}(\mathcal{A})$ and $\mathcal{L}_{tw}(\mathfrak{G})$. Let (r, s) denote a state of this product automaton $\mathcal{B}_{\mathcal{A},\mathfrak{G}}^{prod}$ where r is a state of \mathcal{A} and s is a state of \mathcal{B}_C^B . Without loss of generality, we can assume that the set of clocks used by \mathcal{A} , denoted $\mathcal{Z}_{\mathcal{A}}$, and \mathcal{B}_C^B , denoted $\mathcal{Z}_{\mathfrak{G}}$ are disjoint. Then, the set of clocks of $\mathcal{B}_{\mathcal{A},\mathfrak{G}}^{prod}$ is the union of the set of clocks of \mathcal{A} and \mathcal{B}_C^B .

A transition is of the form $(r, s) \xrightarrow{\varphi, a, R} (r', s')$ where

- $\varphi = \varphi_1 \wedge \varphi_2$ and $R = R_1 \cup R_2$, for some $\varphi_1 \in \text{Form}(\mathcal{Z}_{\mathcal{A}})$, $\varphi_2 \in \text{Form}(\mathcal{Z}_{\mathfrak{G}})$
- $(r, \varphi_1, a, R_1, r')$ is a transition of \mathcal{A}
- $(s, \varphi_2, a, R_2, s')$ is a transition of \mathcal{B}_C^B .

From the above it follows that,

Lemma 15. $\mathcal{L}_{tw}(\mathcal{B}_{\mathcal{A},\mathfrak{G}}^{prod}) = \{\sigma \mid \sigma \in \mathcal{L}_{tw}(\mathcal{A}) \text{ and } \sigma \in \mathcal{L}_{tw}(\mathfrak{M}) \text{ for some } \mathfrak{M} \in \mathcal{L}_{TC}(\mathfrak{G})\}$.

Now, we modify the construction by considering the alphabet $\Sigma = Act \times (V \cup \diamond)$ where V is the set of nodes of \mathfrak{G} and \diamond is an extra symbol. The set of states and clocks are the same as before. We redefine the transitions as follows: $(r, s) \xrightarrow{\varphi, (a,b), R} (r', s')$ is a transition if $(r, s) \xrightarrow{\varphi, a, R} (r', s')$ is a transition of the product automaton and

$$b = \begin{cases} u & \text{if } (s, \varphi_2, a, R_2, s') \text{ and } s' \text{ extends } s \text{ by the node } u \in V \\ \diamond & \text{otherwise} \end{cases}$$

The soundness of the above definition follows from [Remark 2](#), [Proposition 14](#). Let us call this new timed automaton $\mathcal{B}_{\mathcal{A},\mathfrak{G}}^{cov}$. Then,

Lemma 16. $\mathcal{L}_{tw}(\mathcal{B}_{\mathcal{A},\mathfrak{G}}^{cov}) = \{\sigma \in \text{TW}_{Act \times (V \cup \diamond)} \mid \sigma = ((a_1, b_1), t_1) \dots ((a_n, b_n), t_n) \text{ such that } (a_1, t_1) \dots (a_n, t_n) \in \mathcal{L}_{tw}(\mathcal{A}) \text{ is a timed linearization of a TMS } T \text{ that realizes the TCMSC } \mathfrak{M}_{b_{i_1} \dots b_{i_m}} \in \mathcal{L}_{TC}(\mathfrak{G}) \text{ where } b_{i_1} \dots b_{i_m} \text{ is the projection of } b_1 \dots b_n \text{ onto } V\}$.

Proof. We define $\mathcal{L}_{tw}^1(\mathcal{B}_{\mathcal{A},\mathfrak{G}}^{cov}) = \{(a_1, t_1) \dots (a_n, t_n) \in \text{TW}_{Act} \mid \text{there exists } \sigma = \sigma_1 \dots \sigma_n \in \mathcal{L}_{tw}(\mathcal{B}_{\mathcal{A},\mathfrak{G}}^{cov}) \text{ such that } \forall i \in \{1, \dots, n\}, \sigma_i = ((a_i, b_i), t_i)\}$. Then this language coincides with the language of $\mathcal{B}_{\mathcal{A},\mathfrak{G}}^{prod}$ defined above, i.e., we have $\mathcal{L}_{tw}^1(\mathcal{B}_{\mathcal{A},\mathfrak{G}}^{cov}) = \mathcal{L}_{tw}(\mathcal{B}_{\mathcal{A},\mathfrak{G}}^{prod})$. Thus, by [Lemma 15](#), $(a_1, t_1) \dots (a_n, t_n) \in \mathcal{L}_{tw}(\mathcal{B}_{\mathcal{A},\mathfrak{G}}^{prod})$ if and only if it is a timed linearization of some TMS that realizes a TCMSC $\mathfrak{M} \in \mathcal{L}_{TC}(\mathfrak{G})$.

Now, consider $\mathcal{L}_{tw}^2(\mathcal{B}_{\mathcal{A},\mathfrak{G}}^{cov}) = \{b_1 \dots b_n \in (V \cup \{\diamond\})^* \mid \text{there exists } \sigma = (\sigma_1 \dots \sigma_n) \in \mathcal{L}_{tw}(\mathcal{B}_{\mathcal{A},\mathfrak{G}}^{cov}) \text{ such that } \forall i \in \{1, \dots, n\}, \sigma_i = ((a_i, b_i), t_i)\}$. Then, from [Proposition 14](#) and the translation in [Section 4.2](#), it follows that $\mathcal{L}_{tw}^2(\mathcal{B}_{\mathcal{A},\mathfrak{G}}^{cov})$ lists out the nodes of G in the order in which they are traversed by $\mathcal{B}_{\mathcal{A},\mathfrak{G}}^{prod}$. Therefore by projecting $\mathcal{L}_{tw}^2(\mathcal{B}_{\mathcal{A},\mathfrak{G}}^{cov})$ onto V (i.e., by erasing \diamond), we obtain the actual path through G that corresponds to the run of $\mathcal{B}_{\mathcal{A},\mathfrak{G}}^{cov}$. In other words, the path $b_{i_1} \dots b_{i_m}$ obtained by projecting $b_1 \dots b_n$ to V is exactly the path that generates the TCMSC $\mathfrak{M} \in \mathcal{L}_{TC}(\mathfrak{G})$, which completes the proof. \square

We define $\mathcal{L}_V^2(\mathcal{B}_{\mathcal{A},\mathfrak{G}}^{cov}) = \{u_1 \dots u_n \in V^* \mid \text{there exists } w \in \text{Untime}(\mathcal{L}_{tw}(\mathcal{B}_{\mathcal{A},\mathfrak{G}}^{cov})) \text{ such that } w \text{ projected on its second component and projected onto the alphabet } V \text{ gives } u_1 \dots u_n\}$.

To check coverage, we just need to verify that the *node language* of G , $\mathcal{L}_V(\mathfrak{G}) = \{u_0 u_1 \dots u_n \in V^* \mid u_0 \rightarrow u_1 \rightarrow \dots \rightarrow u_n \text{ is a run}\}$, is included in $\mathcal{L}_V^2(\mathcal{B}_{\mathcal{A},\mathfrak{G}}^{cov})$. This would imply that for every path π through \mathfrak{G} , the TCMSC \mathfrak{M}_π has some timed linearization which is in $\mathcal{L}_{tw}(\mathcal{A})$.

Lemma 17. *If $\mathcal{L}_V(\mathfrak{G}) \subseteq \mathcal{L}_V^2(\mathcal{B}_{\mathcal{A},\mathfrak{G}}^{cov})$, then for all $\mathfrak{M} \in \mathcal{L}_{TC}(\mathfrak{G})$, there exists $\sigma \in \mathcal{L}_{tw}(\mathcal{A})$ such that $\sigma \in \mathcal{L}_{tw}(\mathfrak{M})$.*

Proof. $\mathfrak{M} \in \mathcal{L}_{TC}(\mathfrak{G})$ implies that $\mathfrak{M} = \mathfrak{M}_\pi$ for some path $\pi \in \mathcal{L}_V(\mathfrak{G})$. But then $\pi \in \mathcal{L}_V^2(\mathcal{B}_{\mathcal{A},\mathfrak{G}}^{cov})$ which implies by definition that there exists $\sigma \in \mathcal{L}_{tw}(\mathcal{B}_{\mathcal{A},\mathfrak{G}}^{cov})$ such that $\text{Untime}(\sigma)$ projected onto V is π . But then by [Lemma 16](#), we are done since σ projected onto the first component gives us the witness. \square

The converse of the above lemma may not hold in general, since some paths in \mathfrak{G} may define TCMSCs that cannot be realized, because of self-contradictory timing constraints. However, it is easy to exclude such paths. We start with the trivial automaton \mathcal{A}_U that recognizes Act^* , which can be regarded as a degenerate timed automaton with no timing constraints. To \mathcal{A}_U , we apply the same construction as we have done for \mathcal{A} . The resulting timed automaton $\mathcal{B}_{\mathcal{A}_U,\mathfrak{G}}^{cov}$ will mark out all

paths π through \mathfrak{G} for which \mathfrak{M}_π can be realized by some TMSG. Hence, checking that every realizable path is witnessed by \mathcal{A} amounts to checking that $\mathcal{L}_V^2(\mathcal{B}_{\mathcal{A}_U, \mathfrak{G}}^{\text{cov}})$ is included in $\mathcal{L}_V^2(\mathcal{B}_{\mathcal{A}, \mathfrak{G}}^{\text{cov}})$.

Lemma 18. $\mathcal{L}_V^2(\mathcal{B}_{\mathcal{A}_U, \mathfrak{G}}^{\text{cov}}) \subseteq \mathcal{L}_V^2(\mathcal{B}_{\mathcal{A}, \mathfrak{G}}^{\text{cov}})$ if and only if for all $\mathfrak{M} \in \mathcal{L}_{\text{TC}}(\mathfrak{G})$ there exists $\sigma \in \mathcal{L}_{\text{tw}}(\mathcal{A})$ such that $\sigma \in \mathcal{L}_{\text{tw}}(\mathfrak{M})$.

Since both $\mathcal{L}_V^2(\mathcal{B}_{\mathcal{A}_U, \mathfrak{G}}^{\text{cov}})$ and $\mathcal{L}_V^2(\mathcal{B}_{\mathcal{A}, \mathfrak{G}}^{\text{cov}})$ are regular languages, the result now follows. That is,

Theorem 19. For an event-saturated TMSG \mathfrak{G} and a timed automaton \mathcal{A} over Act, the coverage problem for \mathfrak{G} and \mathcal{A} is decidable, i.e., it is decidable to check if for all $\mathfrak{M} \in \mathcal{L}_{\text{TC}}(\mathfrak{G})$, there exists $\sigma \in \mathcal{L}_{\text{tw}}(\mathcal{A})$ such that $\sigma \in \mathcal{L}_{\text{tw}}(\mathfrak{M})$.

8. Conclusion

Given a locally synchronized TMSG and a timed MPA, we have shown that the model checking problem is decidable. That is, we can check if every timed execution of the timed MPA is witnessed by some TMSG generated by the TMSG. In the reverse direction, if the TMSG is in addition event-saturated, then we can prove that the coverage problem is also decidable. That is, we can check if every TMSG generated by the TMSG is witnessed by some timed execution of the timed MPA. We have argued why the above restrictions make sense and demonstrated why our proof techniques are likely to fail in more general settings. Together, this provides a new framework for defining and solving the problem of conformance for time-constrained distributed specifications.

References

- [1] S. Akshay, Specification and verification for distributed and timed systems, PhD thesis, Laboratoire Spécification et Vérification, ENS Cachan, France, July 2010.
- [2] S. Akshay, B. Bollig, P. Gastin, Automata and logics for timed message sequence charts, in: Proc. of FSTTCS, in: LNCS, vol. 4855, Springer, 2007, pp. 290–302.
- [3] S. Akshay, P. Gastin, M. Mukund, K. Narayan Kumar, Model checking time-constrained scenario-based specifications, in: Proc. of FSTTCS, in: LIPIcs, vol. 8, Schloss Dagstuhl – Leibniz-Zentrum fuer Informatik, 2010, pp. 204–215.
- [4] S. Akshay, B. Genest, L. Hélouët, S. Yang, Symbolically bounding the drift in time-constrained MSC graphs, in: Proc. of ICTAC, in: LNCS, vol. 7521, Springer, 2012, pp. 1–15.
- [5] S. Akshay, M. Mukund, K. Narayan Kumar, Checking coverage for infinite collections of timed scenarios, in: Proc. of CONCUR, in: LNCS, vol. 4703, Springer, 2007, pp. 181–196.
- [6] R. Alur, D.L. Dill, A theory of timed automata, Theoret. Comput. Sci. 126 (2) (1994) 183–235.
- [7] R. Alur, L. Fix, T.A. Henzinger, Event-clock automata: a determinizable class of timed automata, Theoret. Comput. Sci. 211 (1–2) (1999) 253–273.
- [8] R. Alur, G. Holzmann, D. Peled, An analyser for message sequence charts, in: Proc. of TACAS, in: LNCS, vol. 1055, Springer, 1996, pp. 35–48.
- [9] R. Alur, M. Yannakakis, Model checking of message sequence charts, in: Proc. of CONCUR, in: LNCS, vol. 1664, Springer, 1999, pp. 114–129.
- [10] H. Ben-Abdallah, S. Leue, Timing constraints in message sequence chart specifications, in: Proc. of FORTE, in: IFIP Conference Proceedings, vol. 107, Chapman & Hall, 1997, pp. 91–106.
- [11] J. Bengtsson, W. Yi, Timed automata: semantics, algorithms and tools, in: Lectures on Concurrency and Petri Nets, in: LNCS, vol. 3098, Springer-Verlag, 2003, pp. 87–124.
- [12] B. Berthomieu, M. Diaz, Modeling and verification of time dependent systems using time Petri nets, IEEE Trans. Softw. Eng. 17 (3) (1991) 259–273.
- [13] P. Bouyer, S. Haddad, P.-A. Reynier, Timed unfoldings for networks of timed automata, in: Proc. of ATVA, in: LNCS, vol. 4218, Springer, 2006, pp. 292–306.
- [14] J. Chakraborty, D. D'Souza, K. Narayan Kumar, Analysing message sequence graph specifications, in: Proc. of ISOla (1), in: LNCS, vol. 6415, Springer, 2010, pp. 549–563.
- [15] P. Chandrasekaran, M. Mukund, Matching scenarios with timing constraints, in: Proc. of FORMATS, in: LNCS, vol. 4202, Springer, 2006, pp. 98–112.
- [16] Th. Chatain, Déplages symboliques de réseaux de Petri de haut niveau et application à la supervision des systèmes répartis, Thèse de doctorat, Université Rennes 1, Rennes, France, November 2006.
- [17] Th. Chatain, C. Jard, Time supervision of concurrent systems using symbolic unfoldings of time Petri nets, in: Proc. of FORMATS, in: LNCS, vol. 3829, Springer, 2005, pp. 196–210.
- [18] M. Clerbout, M. Latteux, Semi-commutations, Inform. and Comput. 73 (1) (1987) 59–74.
- [19] A. David, K.G. Larsen, A. Legay, U. Nyman, A. Wasowski, Timed i/o automata: a complete specification theory for real-time systems, in: Proc. of HSCC, ACM, 2010, pp. 91–100.
- [20] D. D'Souza, A logical study of distributed timed automata, PhD thesis, BITS, Pilani, 2000.
- [21] D. D'Souza, M. Mukund, Checking consistency of SDL+MSC specifications, in: Proc. of SPIN, in: LNCS, vol. 2648, Springer, 2003, pp. 151–165.
- [22] J.G. Henriksen, M. Mukund, K. Narayan Kumar, M. Sohoni, P.S. Thiagarajan, A theory of regular MSC languages, Inform. and Comput. 202 (1) (2005) 1–38.
- [23] ITU-TS Recommendation Z.120: Message Sequence Chart 1999 (MSC99), 1999.
- [24] P. Krcal, W. Yi, Communicating timed automata: the more synchronous, the more difficult to verify, in: Proc. of CAV, in: LNCS, vol. 4144, Springer, 2006, pp. 243–257.
- [25] D. Kuske, Regular sets of infinite message sequence charts, Inform. and Comput. 187 (2003) 80–109.
- [26] S. Mauw, M.A. Reniers, High-level message sequence charts, in: Proc. of SDL Forum, Elsevier, 1997, pp. 291–306.
- [27] A. Muscholl, D. Peled, Message sequence graphs and decision problems on Mazurkiewicz traces, in: Proc. of MFCS, in: LNCS, vol. 1672, Springer, 1999, pp. 81–91.
- [28] A. Muscholl, D. Peled, Z. Su, Deciding properties for message sequence charts, in: Proc. of FoSSaCS, in: LNCS, vol. 1378, Springer, 1998, pp. 226–242.
- [29] C. Ramchandani, Analysis of asynchronous concurrent systems by timed Petri nets, PhD thesis, Massachusetts Institute of Technology, 1974.
- [30] W. Reisig, Petri Nets: An Introduction, Monographs in Theoretical Computer Science. An EATCS Series, vol. 4, Springer, 1985.