

Meeting Deadlines Together

Laura Bocchi¹, Julien Lange², and Nobuko Yoshida²

1 University of Kent

2 Imperial College London

Abstract

This paper studies safety, progress, and non-zeno properties of Communicating Timed Automata (CTAs), which are timed automata (TA) extended with unbounded communication channels, and presents a procedure to *build* timed global specifications from systems of CTAs. We define safety and progress properties for CTAs by extending properties studied in communicating finite-state machines to the timed setting. We then study non-zenoness for CTAs; our aim is to prevent scenarios in which the participants have to execute an infinite number of actions in a finite amount of time. We propose sound and decidable conditions for these properties, and demonstrate the practicality of our approach with an implementation and experimental evaluations of our theory.

1998 ACM Subject Classification F.3.1 Specifying and Verifying and Reasoning about Program

Keywords and phrases timed automata, multiparty session types, global specification

Digital Object Identifier 10.4230/LIPIcs.xxx.yyy.p

1 Introduction

Meeting deadlines is part of our everyday life; this is also the case for distributed software systems that have real-time constraints, such as e-business and financial systems, where exchanges of agreements and data transmissions need to be completed within specified time-frames. Guaranteeing that a single entity will finish its assigned task within an upcoming deadline is a crucial requirement that is generally difficult to attain. It is even harder to ensure that several, distributed, and interdependent entities will work *together* in a timely fashion to meet each other's deadlines. To model such real-time distributed behaviours, communicating timed automata (CTAs) [17] have been introduced as an extension of communicating finite-state machines (CFSMs) [9] with time constraints. A system of CTAs consists of several automata that exchange messages through unbounded FIFO channels and must comply with time constraints on emission/reception of messages. These two features (unbounded channels and time) make CTAs difficult to verify, e.g., reachability is undecidable in general [12].

This paper tackles the following two shortcomings of the current state-of-the-art of CTAs. First, to the best of our knowledge, safety and progress properties, such as absence of deadlocks and unspecified reception (type) errors, which are standard in the literature on CFSMs [10], and essential for distributed systems, have not been studied in the context of CTAs. Moreover, customary properties for TAs such as time-divergence [2] and non-zenoness [7, 21] (preventing that some participant's only possible way forward is by firing actions at increasingly short intervals of time) have not been investigated for CTAs.

Second, while global specifications such as message sequent charts (MSC) and choreographies [8, 16] are useful to model protocols from a global viewpoint, there has not been any work to *build* global specifications from CTAs. The top-down approach [6] alone, which requires a preexisting global specification, is not satisfactory in agile development life-cycles [23], in refinement and reverse-engineering of existing systems, or to compose real-time distributed components, possibly dynamically (see [14, 18, 19]).



© Laura Bocchi, Julien Lange, and Nobuko Yoshida;
licensed under Creative Commons License CC-BY


Conference title on which this volume is based on.

Editors: Billy Editor and Bill Editors; pp. 1–33



Leibniz International Proceedings in Informatics

Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl Publishing, Germany



■ **Figure 1** Scheduled Task Protocol (System S_{ST})

This work introduces classical properties of CFSMs and TAs to the world of CTAs, and investigates the interplay between asynchronous communications through unbounded channels and time constraints. We define the classes of CTAs that enjoy four properties – *safety*, *progress*, *non-zenoness*, and *eventual reception* – and give a sound decision procedure for checking whether a system of CTAs belongs to these classes. This procedure does not rely on any other information than the CTAs themselves. Interestingly, a property of CFSMs called multiparty compatibility (MC) [14], which characterises a sound and complete correspondence with multiparty session types in the untimed setting [16], soundly characterises safe CTAs and offers a basis for decidable decision procedures for progress and non-zenoness in the timed setting. We give: (i) a sound characterisation for progress by checking the satisfiability of first order logic formulae (thus verifiable by generic SMT solvers), and (ii) a sound characterisation of non-zenoness by using a synchronous execution of CTAs. Eventual reception follows from (i) and (ii). In addition, we present an algorithm to build a timed global type [6] from CTAs, whose traces are equivalent to the original system. Thus, if a system validates some of the properties discussed above, then the CTAs obtained by projecting its timed global type onto its participants will preserve these properties.

The system S_{ST} in Fig. 1 (Scheduled Task Protocol) will be used to illustrate our approach throughout the paper. S_{ST} consists of three participants (or machines): a user U , a worker W , and an aggregator A , who exchange messages through *unbounded* FIFO buffers. Each machine is equipped with one or more clocks, initially set to 0 and possibly reset during the protocol. Time elapses at the same pace for all clocks, which is a standard assumption [17]. The protocol is as follows: U sends a task to W , W progressively sends intermediary data to A , and finally A sends the aggregated result to U . The time constraints are:

- U must send a task to W within one time unit, reset its clock x , and expects to receive the result within 15 time units.
- W must consume U 's *task* message at time 1, reset its clocks y and y' , and repeatedly send *data* to A , waiting less than 1 time unit between each emission (modelled by the constraint and reset on y). The overall iteration cannot last more than 10 time units (modelled by the constraint on y' , which is not reset in the loop). When W has finished, it must send a notification *stop* to A .
- A must read intermediary data every 1 time unit, reset each time its clock z , and send the overall result to U within 5 time units after receiving *stop*.

This example, albeit small, models a complex interaction where each machine has its own, interdependent, deadlines; e.g., U relies on the other machines' deadlines to receive the final result within 20 time units. Note that the channel between W and A is *unbounded*: W can send to A an arbitrary number of messages before A receives them, cf. $WA!data(y < 2 \wedge y' < 10, y := 0)$.

Contribution and synopsis In the rest of the paper, we give several conditions that guarantee that no participant misses its deadlines, that every message sent is eventually received *on*

time, and that no participant is forced to perform actions infinitely fast, i.e., forced into a zeno behaviour. In § 2 we recall basic definitions on CTAs. In § 3 we extend the standard safety properties of CFSMs to the timed setting, and show that multiparty compatibility (MC) is a sound condition for safety (Theorem 6). MC CTAs still allow undesirable scenarios when, e.g., (1) the system gets stuck because of unmeetable deadlines, (2) the system's only possibility to meet its deadlines is through zeno behaviours, or (3) sent messages are never received. We give sound and decidable conditions to rule out (1) in § 4 (Theorem 13) and (2-3) in § 5 (Theorem 17 and Theorem 19). In § 6, we discuss the applications of our theory and its implementation. The work in [6] studies a correspondence between timed local types (projected from timed global types) and CTAs, focusing on type-checking timed π -calculus processes. The present work studies CTAs directly, i.e., without relying on a priori global knowledge of the system, and gives more general conditions for safety, progress, and non-zenoness. None of the previous works [14, 18, 19] on building global specifications from local ones caters for time constraints. Unlike existing work on the properties of CTAs (e.g., reachability) our results do not set limitations to channel size or to network topologies [12, 17]. We discuss related work further in § 7. The proofs, additional material, and the implementation are available online [3].

2 Communicating Timed Automata

We introduce communicating timed automata (CTA) following definitions from [14, 17]. Fix a finite set \mathcal{P} of *participants* (ranged over by $\mathbf{p}, \mathbf{q}, \mathbf{r}, \mathbf{s}$, etc.). Let \mathbb{A} be a finite alphabet of messages ranged over by a, b , etc. The set of finite words on \mathbb{A} is denoted by \mathbb{A}^* , ww' is the concatenation of w and w' , and ε is the empty word (overloaded on any alphabet). The set of *channels* is $C \stackrel{\text{def}}{=} \{\mathbf{pq} \mid \mathbf{p}, \mathbf{q} \in \mathcal{P} \text{ and } \mathbf{p} \neq \mathbf{q}\}$. Given a (finite) set of *clocks* \mathcal{X} (ranged over by x, y , etc.), the set of *actions* (ranged over by ℓ) is $Act_{\mathcal{X}} \stackrel{\text{def}}{=} C \times \{!, ?\} \times \mathbb{A} \times \Phi(\mathcal{X}) \times 2^{\mathcal{X}}$, and the set of *guards* (ranged over by g) $\Phi(\mathcal{X})$ is

$$g ::= \text{true} \mid x \leq c \mid c \leq x \mid \neg g \mid g_1 \wedge g_2$$

where c ranges over constants in $\mathbb{Q}_{\geq 0}$, and from which we derive the usual abbreviations. We write $\text{fc}(g)$ for the set of clocks in g and $\mathbf{sr}!a(g, \lambda)$ or $\mathbf{sr}?a(g, \lambda)$ for an element of $Act_{\mathcal{X}}$. Action $\mathbf{sr}!a(g, \lambda)$ says that \mathbf{s} sends a message a to \mathbf{r} , provided that guard g is satisfied, and resets the clocks in $\lambda \subseteq \mathcal{X}$; the dual receiving action is $\mathbf{sr}?a(g, \lambda)$. Given $\ell = \mathbf{sr}!a(g, \lambda)$ or $\ell = \mathbf{sr}?a(g, \lambda)$, we define: $\text{msg}(\ell) = a$, $\text{guard}(\ell) = g$, and $\text{reset}(\ell) = \lambda$. We define the subject of an action: $\text{subj}(\mathbf{pr}!a(g, \lambda)) = \text{subj}(\mathbf{sp}?a(g, \lambda)) \stackrel{\text{def}}{=} \mathbf{p}$.

A *communicating timed automaton*, or *machine*, is a finite transition system given by a tuple $M = (Q, q_0, \mathcal{X}, \delta)$ where Q is a finite set of *states*, $q_0 \in Q$ is the initial state, \mathcal{X} is a set of clocks, and $\delta \subseteq Q \times Act_{\mathcal{X}} \times Q$ is a set of *transitions*. We write $q \xrightarrow{\ell} q'$ when $(q, \ell, q') \in \delta$.

A machine $M = (Q, q_0, \mathcal{X}, \delta)$ is *deterministic* if for all states $q \in Q$ and all actions $\ell, \ell' \in Act_{\mathcal{X}}$, if $(q, \ell, q'), (q, \ell', q'') \in \delta$ and $\text{msg}(\ell) = \text{msg}(\ell')$, then $q' = q''$ and $\ell = \ell'$. A state $q \in Q$ is: *final* if it has no outgoing transitions; *sending* (resp. *receiving*) if it is not final and each of its outgoing transitions is of the form $\mathbf{sr}!a(g, \lambda)$ (resp. $\mathbf{sr}?a(g, \lambda)$); and *mixed* if it is neither final, sending, nor receiving. We say that q is *directed* if it contains only sending/receiving actions to/from the same participant. Hereafter, we only consider deterministic machines, whose states are directed and not mixed. These assumptions, adapted from [14], ensure that a machine corresponds to a syntactic local session type [16]. We discuss how to lift some of these restrictions in § 7.

A *timed communicating system* consists of a finite set of machines and a set of queues (one

4 Meeting Deadlines Together

for each channel) used for asynchronous message passing. Given a valuation $\nu : \mathcal{X} \rightarrow \mathbb{R}_{\geq 0}$ of the clocks in \mathcal{X} , $\nu \models g$ denotes that the guard g is satisfied by ν and $\lambda(\nu)$ denotes a valuation where all clocks in λ are set to 0 (reset) and clocks not in λ keep their values in ν .

► **Definition 1** (Timed communicating system). A timed communicating system (or *system*), is a tuple $S = (M_p)_{p \in \mathcal{P}}$ where each $M_p = (Q_p, q_{0p}, \mathcal{X}_p, \delta_p)$ is a CTA and for all $p \neq q \in \mathcal{P}$: $\mathcal{X}_p \cap \mathcal{X}_q = \emptyset$. A *configuration* of S is a triple $s = (\vec{q}; \vec{w}; \nu)$ where: $\vec{q} = (q_p)_{p \in \mathcal{P}}$ is the *control state* and $q_p \in Q_p$ is the *local state* of machine M_p ; $\vec{w} = (w_{pq})_{pq \in C}$ with $w_{pq} \in \mathbb{A}^*$ is a vector of queues; $\nu : \bigcup_{p \in \mathcal{P}} \mathcal{X}_p \rightarrow \mathbb{R}_{\geq 0}$ is a clock valuation. The *initial configuration* of S is $s_0 = (\vec{q}_0; \vec{\varepsilon}; \nu_0)$ with $\vec{q}_0 = (q_{0p})_{p \in \mathcal{P}}$, $\vec{\varepsilon}$ being the vector of empty queues, and $\nu_0(x) = 0$ for each clock $x \in \bigcup_{p \in \mathcal{P}} \mathcal{X}_p$. ◊

Hereafter, we fix a machine $M_p = (Q_p, q_{0p}, \mathcal{X}_p, \delta_p)$ for each participant $p \in \mathcal{P}$ (assuming that $\forall p \in \mathcal{P} : (q, \ell, q') \in \delta_p \implies \text{subj}(\ell) = p$), and let $S = (M_p)_{p \in \mathcal{P}}$ be the corresponding system. We write \mathcal{X} for $\bigcup_{p \in \mathcal{P}} \mathcal{X}_p$ and $\nu + t$ for the valuation mapping each $x \in \mathcal{X}$ to $\nu(x) + t$. The definition below is from [17, Definition 1], omitting internal transitions.

► **Definition 2** (Reachable configuration). Configuration $s' = (\vec{q}'; \vec{w}'; \nu')$ is *reachable* from configuration $s = (\vec{q}; \vec{w}; \nu)$ by *firing the transition* α , written $s \xrightarrow{\alpha} s'$ (or $s \rightarrow s'$ when the label is immaterial), if either:

1. $(q_s, \text{sr}!a(g, \lambda), q'_s) \in \delta_s$ and (a) $q'_p = q_p$ for all $p \neq s$; (b) $w'_{sr} = w_{sr}a$ and $w'_{pq} = w_{pq}$ for all $pq \neq sr$; (c) $\nu' = \lambda(\nu)$; (d) $\alpha = \text{sr}!a(g, \lambda)$, and $\nu \models g$;
2. $(q_r, \text{sr}?a(g, \lambda), q'_r) \in \delta_r$ and (a) $q'_p = q_p$ for all $p \neq r$; (b) $w_{sr} = aw'_{sr}$ and $w'_{pq} = w_{pq}$ for all $pq \neq sr$; (c) $\nu' = \lambda(\nu)$; (d) $\alpha = \text{sr}?a(g, \lambda)$ and $\nu \models g$; or
3. $\alpha = t \in \mathbb{R}_{\geq 0}$, $\nu' = \nu + t$, $w'_{pq} = w_{pq}$ for all $pq \in C$, and $q'_p = q_p$ for all $p \in \mathcal{P}$.

We let ρ range over sequences of labels $\alpha_1 \cdots \alpha_k$ and write \rightarrow^* for the reflexive transitive closure of \rightarrow . The *reachability set* of S is $RS(S) \stackrel{\text{def}}{=} \{s \mid s_0 \rightarrow^* s\}$. ◊

Condition (1) allows a machine s to put a message a on queue sr , if the time constraints in g are satisfied by ν ; dually, (2) allows r to consume a message from the queue, if g is satisfied; and (3) models the elapsing of time (or a delay).

3 Safety in CTAs

This section defines safe CTAs and gives a sufficient condition for safety, called multiparty compatibility (MC) [14], in the timed setting. Here, we present a new approach based on *synchronous transition systems* (STS); the STS is also useful for defining progress and non-zeno properties in § 4.

Let n range over vectors of local states; and e range over events, which are elements of the set $C \times \mathbb{A} \times \Phi(\mathcal{X}) \times 2^{\mathcal{X}} \times \Phi(\mathcal{X}) \times 2^{\mathcal{X}}$, and write $(s \rightarrow r : a; g_s, \lambda_s; g_r, \lambda_r)$ for the event in which s sends message a to r , with s (resp. r) having guard g_s (resp. g_r) and resets λ_s (resp. λ_r). We introduce the synchronous transition system of S , following [19].

► **Definition 3** (Synchronous transition system). The *synchronous transition system* of S , written $STS(S)$, is a tuple $(N, n_0, \hookrightarrow, E)$ such that:

- \hookrightarrow is the relation defined as $n \xrightarrow{e} n'$ with $e = (s \rightarrow r : a; g_s, \lambda_s; g_r, \lambda_r)$ iff $n = \vec{q}, n' = \vec{q}', q_s \xrightarrow{\text{sr}!a(g_s, \lambda_s)} q'_s, q_r \xrightarrow{\text{sr}?a(g_r, \lambda_r)} q'_r$, and $\forall p \in \mathcal{P} \setminus \{s, r\} : q_p = q'_p$ (write \hookrightarrow when e is unimportant and \hookrightarrow^* for the reflexive and transitive closure of \hookrightarrow);
- $n_0 = \vec{q}_0$ is the initial node; $N = \{n \mid n_0 \hookrightarrow^* n\}$ is the (finite) set of nodes; and $E = \{e \mid \exists n, n' \in N \text{ and } n \xrightarrow{e} n'\}$ is the set of events.

We write $n_1 \xrightarrow{e_1 \cdots e_k} n_{k+1}$, when, for some $n_2, \dots, n_k \in N$, $n_1 \xrightarrow{e_1} n_2 \cdots n_k \xrightarrow{e_k} n_{k+1}$. Let φ range over (possibly empty) sequences of events $e_1 \cdots e_k$, and ε denote the empty sequence. \diamond


The *STS* of the Scheduled Task Protocol (S_{ST}) is given in Fig. 2; essentially, it models all the synchronous executions of S_{ST} . In the following, we fix $STS(S) = (N, n_0, \leftrightarrow, E)$.

Given $e = (\mathbf{s} \rightarrow \mathbf{r} : a; g_s, \lambda_s; g_r, \lambda_r)$, we define $\text{sid}(e) \stackrel{\text{def}}{=} \mathbf{s}$, $\text{rid}(e) \stackrel{\text{def}}{=} \mathbf{r}$, and $\text{id}(e) \stackrel{\text{def}}{=} \{\mathbf{s}, \mathbf{r}\}$. The projection of e on \mathbf{p} (written $e|_{\mathbf{p}}$) is given by: $(\mathbf{s} \rightarrow \mathbf{r} : a; g_s, \lambda_s; g_r, \lambda_r)|_{\mathbf{s}} = \mathbf{sr}!a(g_s, \lambda_s)$; $(\mathbf{s} \rightarrow \mathbf{r} : a; g_s, \lambda_s; g_r, \lambda_r)|_{\mathbf{r}} = \mathbf{sr}?a(g_r, \lambda_r)$; and $(\mathbf{s} \rightarrow \mathbf{r} : a; g_s, \lambda_s; g_r, \lambda_r)|_{\mathbf{p}} = \varepsilon$, if $\mathbf{p} \notin \{\mathbf{s}, \mathbf{r}\}$. We extend $\varphi|_{\mathbf{p}}$ to sequences of events and, given $n \in N$, define $\text{ids}(n) \stackrel{\text{def}}{=} \bigcup \{\text{id}(e) \mid n \xrightarrow{*} \cdot \xrightarrow{e}\}$.

► **Definition 4** (Multiparty compatibility (MC)). System S is *multiparty compatible* if for all $\mathbf{p} \in \mathcal{P}$, for all $q \in Q_{\mathbf{p}}$, and for all $n = \vec{q} \in N$, if $q_{\mathbf{p}} = q$, then

1. if q is a sending state, then $\forall (q, \ell, q') \in \delta_{\mathbf{p}} : \exists \varphi, \exists e \in E : n \xrightarrow{\varphi} \ell \wedge e|_{\mathbf{p}} = \ell \wedge \varphi|_{\mathbf{p}} = \varepsilon$;
2. if q is a receiving state, then $\exists (q, \ell, q') \in \delta_{\mathbf{p}} : \exists \varphi, \exists e \in E : n \xrightarrow{\varphi} \ell \wedge e|_{\mathbf{p}} = \ell \wedge \varphi|_{\mathbf{p}} = \varepsilon$. \diamond

Intuitively, condition (1) ensures that for every sending state, all messages that can be sent can also be received, while (2) guarantees that, for every receiving state, at least one transition will be eventually fireable, i.e., an *expected* message will eventually be received. System S_{ST} , in Fig. 1, is multiparty compatible.



■ **Figure 2** *STS* for Scheduled Task, cf. Fig. 1

Observe that $STS(S)$ and MC do not address time constraints. In fact, $STS(S)$ might include interactions forbidden by time constraints. These can be ruled out at a later stage when analysing time properties in § 4. We deliberately kept communication and time properties separated, so that we can provide simpler and modular definitions in § 7. Crucially, MC guarantees that any *asynchronous* execution can be mapped to a path in $STS(S)$, i.e., it can be simulated by $STS(S)$.

We recall two types of errors from the CFSM model, which are ruled out by MC also in the timed setting. Let $s = (\vec{q}; \vec{w}; \nu)$ be a configuration of a system S ; s is a **deadlock configuration** [10, Def. 12] if $\vec{w} = \vec{\varepsilon}$, there is $\mathbf{r} \in \mathcal{P}$ such that $q_{\mathbf{r}}$ is a receiving state, and for every $\mathbf{p} \in \mathcal{P}$, $q_{\mathbf{p}}$ is a receiving or final state, i.e., all machines are blocked waiting for messages; and s is an **orphan message configuration** if all $q_{\mathbf{p}} \in \vec{q}$ are final but $\vec{w} \neq \vec{\varepsilon}$, i.e., there is at least a non-empty buffer and all the machines are in a final state.

► **Definition 5** (Safe system). S is *safe* if for all $s \in RS(S)$, s is not a deadlock, nor an orphan message configuration. \diamond

► **Theorem 6** (Safety). *If S is multiparty compatible, then it is safe.*

The proof follows from the fact that (i) MC guarantees safety in CFSMs [14] and (ii) time constraints imply that a subset of the configurations reachable in the untimed setting are reachable in the timed setting (modulo clock valuations). Thus, if there is a deadlock or an orphan message configuration in the timed setting, there is one in the untimed setting, which contradicts the results in [14].

The projection $STS(S)|_{\mathbf{p}}$ of a synchronous transition system $STS(S)$ on a machine \mathbf{p} is given by substituting each event $e \in E$ with its projection $e|_{\mathbf{p}}$, then minimising the automaton w.r.t. language equivalence. For example, the projections of $STS(S)$ onto \mathbf{U} , \mathbf{W} , and \mathbf{A} are isomorphic to the system S_{ST} in Fig. 1. Below \sim denotes the standard timed bisimulation [15].

► **Theorem 7** (Equivalence). *If $S = (M_p)_{p \in \mathcal{P}}$ is MC then $S \sim (STS(S)|_p)_{p \in \mathcal{P}}$.*

Theorem 7 says that the behaviour of the original system is preserved by $STS(S)$, this result is crucial to be able to construct a global specification that is equivalent to a system of CTAs. It follows from the fact that, (i) if the system is MC, then all the machine’s behaviour is preserved except for the receive actions that are never executed; and (ii) since we assume that the machines are deterministic w.r.t. messages, the projections of $STS(S)$ also preserve all required transitions.

4 Progress with Time Constraints

This section introduces a *progress* property for CTAs, ensuring that no communication mismatch prevents the progress of the overall system (cf. § 4.1). In § 4.2, we give a sufficient condition to guarantee progress in CTAs (cf. Theorem 13).

4.1 Progress Properties

We identify several types of errors, inspired by their counterparts in the (untimed) CFSM model, which may arise in timed communicating systems. Let $s = (\vec{q}; \vec{w}; \nu) \in RS(S)$; s is an **unsuccessful reception configuration** if there exists $r \in \mathcal{P}$ such that q_r is a receiving state, and for all $(q_r, sr?a(g, \lambda), q'_r) \in \delta_r$ either (i) $w_{sr} \neq \varepsilon$ and $w_{sr} \notin a\mathbb{A}^*$ or (ii) $\forall t \in \mathbb{R}_{\geq 0} : \nu + t \not\models g$ (i.e., r cannot receive messages from any of its queues, as they either contain an unexpected message or none of the transition guards will ever be satisfied); and s is an **unfeasible configuration** if there exists $s \in \mathcal{P}$ such that q_s is a sending state, and $(q_s, sr!a(g, \lambda), q'_s) \in \delta_s$ implies that $\forall t \in \mathbb{R}_{\geq 0} : \nu + t \not\models g$ (i.e., s is unable to send a message because none of its guards will ever be satisfied).

► **Definition 8** (Progress). S satisfies the *progress* property if for all $s \in RS(S)$, s is not a deadlock, an orphan message, an unsuccessful reception, nor an unfeasible configuration. ◊

Observe that the original semantics of CTAs in [17] and in Def. 2 do not allow us to identify unsuccessful reception or unfeasible configurations. From Def. 2, a system may take a time transition which *permanently* prevents a machine from firing further actions. Below, we adjust the semantics of CTAs and give examples of “undesirable” scenarios it prevents.

► **Definition 9** (Reachable configuration (2)). $s \xrightarrow{\alpha} s'$ is defined as Def. 2, replacing (3) with:

3. $\alpha = t \in \mathbb{R}_{>0}$, $\nu' = \nu + t$, $\forall pq \in C : w'_{pq} = w_{pq}$, and $\forall p \in \mathcal{P} : q'_p = q_p$ and
 - a. q_p sending $\implies \exists (q_p, \ell, q''_p) \in \delta_p : \exists t' \in \mathbb{R}_{\geq 0} : \nu' + t' \models \text{guard}(\ell)$
 - b. $\forall (q_p, sp?a(g, \lambda), q''_p) \in \delta_p : (w_{sp} \in a\mathbb{A}^* \implies \exists t' \in \mathbb{R}_{\geq 0} : \nu' + t' \models g)$

Unless stated otherwise, we only consider this semantics hereafter. ◊

Condition (3a) handles the case of machines waiting to perform send actions, and (3b) handles receive transitions, as illustrated by the examples below:



Consider configuration $((q_0, q_2); \vec{\varepsilon}; \nu_0)$ in which s must send a message within 3 time units. Condition (3a) prevents a time transition with delay $t = 3$. Indeed, with a clock valuation $\nu_0 + 3$, none of the action of s from q_0 can be fired. Consider now configuration $((q_1, q_2); \vec{w}; \nu)$ with $w_{sr} = a$ and $\nu(x) = \nu(y) = 3.5$. Condition (3b) rules out a time transition with $t = 1$. Indeed, even if r has a transition whose guard will be enabled after time $\nu(y) + 1 = 4.5$, i.e.,

$(q_2, \text{sr}^?b(y = 5), q_3)$, this transition cannot be fired due to the content of queue $w_{\text{sr}} \notin b\mathbb{A}^*$; on the other hand transition $(q_2, \text{sr}^?a(y = 4), q_3)$ is no longer fireable, due to its time constraint.

4.2 A Sound Characterisation of Progress

Roadmap We give a sound condition that guarantees progress in the presence of time constraints. The main property, *interaction-enabling* (IE) in Def. 12, essentially checks that future actions are possible. IE guarantees that: (1) whatever the past, each machine that is in a sending state is eventually able to fire one of its transitions and (2) for every message that is sent, there exists a (future) time where this message can be received. IE relies on checking whether an action ℓ is *progress enabling* (Def. 11) which ensures that, for all possible past clock valuations, there exists a future time where the guard of ℓ is satisfied.

In the rest of this section, we give (i) a procedure for understanding the past of a configuration, based on a graph modelling the causal dependencies between previously executed actions; and (ii) a procedure to check that, for any reachable configuration, there is always a future time where an available action can be fired.

Understanding the past We check that S has progress by analysing paths, i.e., sequences of events, in $STS(S)$. Since $STS(S)$ gives an over-approximation of the causal dependencies between actions, we will construct a graph of the actual dependencies of the underlying actions of a path. We compute the underlying actions of a path via the function:

$$\text{nodes}(e_1 \cdots e_k) \stackrel{\text{def}}{=} e_1 \downarrow_{\text{sid}(e_1)} \cdot e_1 \downarrow_{\text{rid}(e_1)} \cdots e_k \downarrow_{\text{sid}(e_k)} \cdot e_k \downarrow_{\text{rid}(e_k)} \quad (k \geq 0)$$

Remarkably, given a path φ and two actions ℓ_i and ℓ_j in $\text{nodes}(\varphi)$, $i < j$ does not imply that there is a causal dependency between ℓ_i and ℓ_j . For instance, in

$$\text{nodes}(\varphi) = \text{sr}!a(x < 10, \emptyset) \cdot \text{sr}^?a(10 \leq y, \emptyset) \cdot \text{sp}!a(x < 10, \emptyset) \cdot \text{sp}^?a(10 \leq z, \emptyset)$$

the two receive actions $\text{sr}^?a(10 \leq y, \emptyset)$ and $\text{sp}^?a(10 \leq z, \emptyset)$ may not always be executed in that order, since they are executed by two different participants.

The graph of dependencies of an action ℓ_k in a sequence of actions $\ell_1 \cdots \ell_k$ (Def. 10 below) gives an abstraction of all actions on which ℓ_k depends. This is done by taking into account two kinds of dependencies: output/input dependencies between matching send and receive actions, and local dependencies within a single machine.


► **Definition 10** (Graph of Dependencies). Let $\text{dep}(\varepsilon; \ell) \stackrel{\text{def}}{=} \emptyset$ and

$$\text{dep}(\rho \cdot \ell_1; \ell_2) \stackrel{\text{def}}{=} \begin{cases} \{(\ell_1, \ell_2)\} \cup \text{dep}(\rho; \ell_i)_{i=1,2} & \text{if } \ell_1 = \text{sr}!a(g_1, \lambda_1), \ell_2 = \text{sr}^?a(g_2, \lambda_2) \\ \{(\ell_1, \ell_2)\} \cup \text{dep}(\rho; \ell_1) & \text{if } \text{subj}(\ell_1) = \text{subj}(\ell_2) \\ \text{dep}(\rho; \ell_2) & \text{otherwise} \end{cases}$$

The graph of dependencies of $\rho = \ell_1 \cdots \ell_k$ ($k > 0$), written $\text{DG}(\rho)$, is the graph (D, A) s.t. $A = \text{dep}(\ell_1 \cdots \ell_{k-1}; \ell_k) \setminus \{(\ell_i, \ell_k) \mid 1 \leq i < k\}$ and $D = \{\ell_r \mid \exists (\ell_i, \ell_j) \in A \wedge r \in \{i, j\}\}$.¹ ◊

$\text{DG}(\ell_1 \cdots \ell_k)$ is a graph whose nodes form a subset of $\{\ell_1, \dots, \ell_{k-1}\}$ and whose edges model causal dependencies between actions (computed backwards starting from ℓ_k). In Fig. 3 (in solid black), we give the graph of dependencies of $\text{WA}^?data(z = 1, \{z\})$ in the sequence ρ_{ST} , corresponding to an execution of the Scheduled Task Protocol.

¹ For the sake of presentation, we write ℓ_i for the node (i, ℓ_i) in D where ℓ_i is an action in ρ and i is its position in ρ . This guarantees that each element in ρ is assigned a unique node in D .



■ **Figure 3** Graph of dependencies $\text{DG}(\rho_{ST})$, in solid black, cf. Scheduled Task Protocol (Fig. 1)

$$\begin{aligned} \text{idX}(\rho) &\stackrel{\text{def}}{=} \{i \mid \ell_i \in D\} & W_x^i(\rho) &\stackrel{\text{def}}{=} \begin{cases} v_i - v_j & \text{if } 0 \leq j = \max\{j < i \mid \ell_j \in D \wedge x \in \text{reset}(\ell_j)\} \\ v_i & \text{otherwise} \end{cases} \\ \text{allpast}(\rho) &\stackrel{\text{def}}{=} \bigwedge_{i \in \text{idX}(\rho)} \text{absolute}_\rho(\ell_i) \\ \text{elapse}(\rho) &\stackrel{\text{def}}{=} \bigwedge_{(\ell_i, \ell_j) \in A} v_i \leq v_j & \text{absolute}_\rho(\ell_i) &\stackrel{\text{def}}{=} \text{guard}(\ell_i) \{x \mapsto W_x^i(\rho)\}_{x \in \mathcal{X}} \end{aligned}$$

■ **Figure 4** Functions on graphs of dependencies, where $\text{DG}(\rho) = (D, A)$

Given a graph of dependencies $\text{DG}(\rho)$, we define several functions that allow us to construct predicates modelling the past. The definitions of these functions are given in Fig. 4, where we fix $\text{DG}(\rho) = (D, A)$. Below, we illustrate how they behave using $\text{DG}(\rho_{ST})$ in Fig. 3. First, we transform the guard of an action ℓ_i such that its solutions are the possible *absolute* times (i.e., from the initial configuration of the system) in which one may execute ℓ_i (taking into account the last reset of each clock in ρ). In our example, we have:

$$\text{absolute}_{\rho_{ST}}(\ell_5) = v_5 - v_3 < 2 \wedge v_5 - v_2 < 10 \quad \text{with } \ell_5 = \text{WA!data}(y < 2 \wedge y' < 10, \{y\})$$

Observe that clock y (resp. y') is replaced by the difference between variable v_5 and variable v_3 (resp. v_2) corresponding to the *latest* step where y (resp. y') was reset. Unifying, e.g., y and y' into v_5 models the fact that time elapses at the same pace for all clocks. Next, we aggregate the information in $\text{DG}(\rho)$, by (i) recording the indices of all the actions on which ℓ_k depends ($\text{idX}(\rho)$); (ii) taking the conjunction of all constraints in absolute time ($\text{allpast}(\rho)$); and (iii) recording the fact that time never decreases between two causally dependent actions ($\text{elapse}(\rho)$). Taking the dependencies for ρ_{ST} in Fig. 3, we have:

$$\begin{aligned} \text{allpast}(\rho_{ST}) &= v_1 < 1 \wedge v_2 = 1 \wedge (v_3 - v_2 < 2 \wedge v_3 - v_2 < 10) \wedge v_4 = 3 \wedge (v_5 - v_3 < 2 \wedge v_5 - v_2 < 10) \\ \text{elapse}(\rho_{ST}) &= v_1 \leq v_2 \wedge v_2 \leq v_3 \wedge v_3 \leq v_4 \wedge v_3 \leq v_5 & \text{idX}(\rho_{ST}) &= \{1, 2, 3, 4, 5\} \end{aligned}$$

Predicting the future We now give the main definition of this section, allowing to check whether the past implies that there exists a satisfiable future. We use the functions defined above to check whether a given event in $\text{STS}(S)$ can indeed meet its time constraints.

► **Definition 11** (Progress enabling (PE)). A pair (n, e) is *progress enabling* (PE) for $\mathbf{p} \in \text{id}(e)$ if for *all* paths φ such that $n_0 \xrightarrow{\varphi} n$, letting:

$$\rho = \begin{cases} \text{nodes}(\varphi \cdot e) & \text{if } \mathbf{p} = \text{rid}(e) \\ \text{nodes}(\varphi) \cdot e|_{\text{sid}(e)} & \text{otherwise} \end{cases}$$

and $k = |\rho|$, $\ell_k = e|_{\mathbf{p}}$, $\vec{v} = \{v_i \mid i \in \text{idX}(\rho)\}$; the following holds

$$\forall \vec{v} \exists v_k : \text{allpast}(\rho) \wedge \text{elapse}(\rho) \implies \text{absolute}_\rho(\ell_k) \wedge \bigwedge_{v_i \in \vec{v}} v_i \leq v_k$$

A pair (n, φ) is *recursively progress enabling* (RPE) for $P \subseteq \mathcal{P}$ if $\varphi = \varepsilon$ and $P = \emptyset$; or if (n, e) is PE for $\text{sid}(e)$ and for $\text{rid}(e)$ and (n', φ') is RPE for $P \setminus \text{id}(e)$ with $\varphi = e \cdot \varphi'$ and $n \xrightarrow{e} n'$. \diamond

Given a node n and an event e in $\text{STS}(S)$, and a participant p , the above definition ensures that for all possible past clock valuations, there exists a *future* time where participant p has the possibility to execute action $e|_p$. For instance, the pair $((U_1, W_1, A_0), (W \rightarrow A : \text{data}; y < 1 \wedge y' < 10, \{y\}; z = 1, \{z\}))$ is PE for A , notably because the following holds:

$$\forall v_1 \dots v_5 \exists v_6 : \text{allpast}(\rho_{\text{ST}}) \wedge \text{elapse}(\rho_{\text{ST}}) \implies (v_6 - v_4) = 1 \wedge v_1 \leq v_6 \dots v_5 \leq v_6$$

Below, Def. 11 is used in $\text{STS}(S)$ to ensure progress of the overall system.


► **Definition 12** (Interaction enabling (IE)). A node $n \in N$ is interaction enabling (IE) if either (i) it is final or (ii) the following conditions hold:

1. There is $e \in E$ and φ such that $n \xrightarrow{e, \varphi}$ and $(n, e \cdot \varphi)$ is RPE for $\text{ids}(n)$;
2. For all $e \in E$ such that $n \xrightarrow{e} n'$, (n, e) is PE for $\text{rid}(e)$, and n' is IE.

A system S is *interaction enabling* (IE) if n_0 is IE. \diamond

Def. 12 recursively checks the nodes of $\text{STS}(S)$ (starting from n_0) and for each ensures that: (1) there is at least one path, involving all the participants still active at node n , that is RPE, i.e., where each guard along that path is satisfied for any past; (2) each receive action is PE and its successor is IE (note that a send action is always a dependency of its receive action). Condition (1) ensures that no sender will be left in a configuration where it cannot send any message, due to time constraints being unsatisfiable; condition (2) ensures that a receive action is always feasible given that its corresponding send action was executed.

Examples (1) The first example shows how resets affect the satisfiability of guards.




The system above is IE, notably, because the following holds:

$$\forall v_1 v_2 v_3 v_4 v_5 \exists v_6 : v_1 = 3 \wedge v_2 \leq 3 \wedge v_3 = 5 \wedge v_4 - v_2 \leq 2 \wedge v_5 = 7 \wedge v_1 \leq \dots \leq v_5 \implies v_6 - v_4 \leq 2 \wedge v_1 \leq v_6 \dots v_5 \leq v_6$$


Notice that the resets of clock y (recorded by subtracting v_2 and v_4 in the formula above) allow r to receive message c before *absolute* time 7. If we modified the example by removing the second reset of y in machine r , then the system would not be IE because message c would be expected before *absolute* time 5, while c can only be sent at time 7. In fact, the RHS of the implication above would become: $v_6 - v_2 \leq 2 \wedge v_1 \leq v_6 \dots v_5 \leq v_6$.

(2) The second example shows a system of three machines, which violates IE (Def. 12).



If participant s does not send b before time 2, then message c (sent by p), will never be received. This system is *not* IE because there is no path from n_0 that is RPE for $\{s, r, p\}$. The only transition that is PE from n_0 is the loop on n_0 (which does not involve p).

(3) The third example shows that IE captures a “global” notion of progress (i.e., *all* participants must be able to proceed). Consider the system of four machines below:



this system is *not* IE. Indeed, although there is one RPE path outgoing node n_1 (machines s_2 and r_2 can continue interacting), there is no path that is RPE for *all* participants $\{s_1, r_1, s_2, r_2\}$. Observe that s_1 is stuck in n_1 , as the transition with label $s_1r_1!b(x < 2, \{x\})$ can never be fired by s_1 , i.e., $\forall v_0 \exists v_1 : v_0 > 2 \implies v_0 \leq v_1 < 2$ does not hold.

► **Theorem 13** (Progress). *Suppose S is multiparty compatible (Def. 4) and interaction enabling (Def. 12). (1) Then S satisfies the progress property. (2) For all $s = (\vec{q}; \vec{w}; \nu) \in RS(S)$, if there is $p \in \mathcal{P}$ such that q_p is not final, then there is s' such that $s \rightarrow s'$.*

► **Theorem 14** (Decidability). *Interaction enabling (Def. 12) is decidable.*

The decidability of Def. 12 relies on the fact that the logic used in Def. 11 forms a subset of the Presburger arithmetic, which is decidable; and that it is enough to check *finite* paths in $STS(S)$. The complexity of the decision procedure is mostly affected by the enumeration of paths in $STS(S)$ (which can be reduced via partial order reduction techniques) and the satisfiability of Presburger formulae (which can be relegated to an SMT solver).


5 Non-Zenoness and Eventual Reception in CTAs

In the presence of time constraints, one needs to make sure that some participant’s only possible way forward is not by firing actions at increasingly short intervals of time, i.e., by zeno behaviours. This is a common requirement in real-time systems [2], and it is justified by the assumption that “*any physical process, no matter how fast, cannot be infinitely fast*” [21].

In order to identify zeno behaviours in our systems, we assume without loss of generality that there is a special clock $\hat{x} \in \mathcal{X}$ which is *never* reset, i.e., for all $p \in \mathcal{P}$ and all $(q, \ell, q') \in \delta_p : \hat{x} \notin \text{reset}(\ell)$. Hence, \hat{x} keeps the absolute time since the beginning of the interactions. Let $s = (\vec{q}; \vec{w}; \nu)$ be a configuration of a system S , s is a **zeno configuration** if there exists $t \in \mathbb{R}_{\geq 0}$ such that for all $s' = (\vec{q}'; \vec{w}'; \nu')$, $s \rightarrow^* s'$ implies $\nu'(\hat{x}) < t$ and $s' \rightarrow s''$, for some s'' .

► **Definition 15** (Non-zeno system). S is *non-zeno* (NZ) if $\forall s \in RS(S)$, s is not a zeno configuration. \diamond

The following example shows that a zeno configuration may still occur in systems that are multiparty compatible and interaction enabling.



The system above (*ignoring the dashed transitions*) satisfies MC and IE, e.g., $\forall v_0 \exists v_1 : v_0 < 3 \implies v_1 \geq 3 \wedge v_0 \leq v_1$, but is *not* NZ. Because of the upper bound $x < 3$ and the fact that x is not reset in the loop, machine s has to produce an *infinite* number of (send) actions in

a *finite* amount of time (3 time units). A dramatic consequence of this zeno behaviour is that machine r will never be able to consume any message a due to the fact that constraint $y \geq 3$ will never be satisfied (cf. Def. 9). This system violates *eventual reception*, a property which guarantees that every message that is sent is eventually received. Formally, a system S satisfies *eventual reception* (ER) if for all $s = (\vec{q}; \vec{w}; \nu) \in RS(S)$, if $w_{sr} \in a\mathbb{A}^*$, then $s \xrightarrow{*} \text{sr}^?a(g, \lambda)_s$.

The system above (*considering the dashed transitions*) is NZ and satisfies ER: the dashed transitions offer an ‘escape’ from zeno-only behaviours where time can elapse and thus allow machine r to consume any messages that were sent. Observe that in general NZ alone is not sufficient to guarantee ER. However, ER is guaranteed for systems which validate all the condition presented in this paper, see Theorem 19 below.

The example also shows a fundamental difference between CTAs and models with synchronous communications, such as Networks of Timed Automata (NTAs) [2]. The work in [7] shows that it is sufficient that one machine in each loop of an NTA satisfies non-zenoness for the whole system to be non-zeno. This is not generally true for CTAs. In the example above (*ignoring the dashed transitions*), time cannot diverge despite the machine on the right being non-zeno.

Checking non-zenoness Now we give a condition on $STS(S)$ that, together with MC, guarantees non-zenoness. A walk in $STS(S)$ is an alternating sequence $n_1 \cdot e_1 \cdot n_2 \cdots e_{k-1} \cdot n_k$ such that $n_i \xrightarrow{e_i} n_{i+1}$ for all $1 \leq i < k$. We let ω range over walks in $STS(S)$. A walk is *elementary* if $(n_i \cdot e_i) \neq (n_j \cdot e_j)$ for all $1 \leq i \neq j < k$. A (elementary) cycle in $STS(S)$ is a (elementary) walk $n_1 \cdot e_1 \cdot n_2 \cdots e_{k-1} \cdot n_k$ such that $n_1 = n_k$.

Given guard g and clock x , we say that g is an *upper bound* for x , written g is *UB* for x , if there is a sub-term $x \leq c$ in g (not under a negation) or a sub-term $c \leq x$ under a negation. We say that g is *strictly positive*, written g is *SP*, if for all clocks $x \in \text{fc}(g)$ and for all sub-terms in g of the form $x \leq c$ (not under negation) or $c \leq x$ (under negation), $c \in \mathbb{Q}_{>0}$.

► **Definition 16** (Cycle enabling (CE)). System S is cycle enabling (CE) if for each elementary cycle ω in $STS(S)$, and for each clock x such that there is $(\mathbf{s} \rightarrow \mathbf{r} : a; g_s, \lambda_s; g_r, \lambda_r)$ in ω and g_s is *UB* for x , the following holds, either

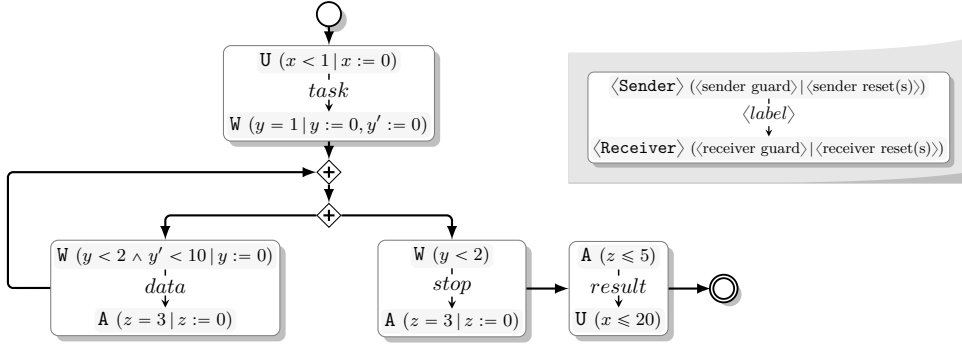
1. there are (i) $(\mathbf{p} \rightarrow \mathbf{q} : b; g_p, \lambda_p; g_q, \lambda_q)$ in ω s.t. $x \in \lambda_p \cup \lambda_q$, and (ii) $(\mathbf{p}' \rightarrow \mathbf{q}' : b'; g_{p'}, \lambda_{p'}; g_{q'}, \lambda_{q'})$ in ω s.t. $g_{p'}$ is *SP*; or
2. for each $(n_i \cdot e \cdot n_{i+1})$ in ω , there is $n' \neq n_i \in N$ and $e' \neq e \in E$ such that $\text{id}(e) = \text{id}(e')$, $n_i \xrightarrow{e'} n'$, and (n_i, e') is PE for $\text{sid}(e')$ ◊

Condition (1) adapts *structural non-zenoness* from [22] to CTAs by requiring that: (i) each x is reset in ω , and (ii) it is possible to let some time elapse at each iteration. Condition (2) requires that the “escape” event e' , leading to a different node n' , can *always* be taken. Our running example satisfies CE (Def. 16); $STS(S_{\text{ST}})$ has one (elementary) cycle in which two clocks have an upper bound: clock y satisfies (1) since it is reset and the guards have upper bounds strictly greater than 0 in the cycle; clock y' satisfies (2) since there is an escape event, $e' = (\mathbb{W} \rightarrow \mathbb{A} : \text{stop}; y < 1, \emptyset; z = 1, \{z\})$, which is PE for \mathbb{W} .

► **Theorem 17** (Non-zenoness). *If S is MC and CE, then S is non-zeno.*

► **Theorem 18** (Decidability). *Cycle enabling (Def. 16) is decidable.*

► **Theorem 19** (Eventual reception). *If S is MC, IE, and CE, then S satisfies ER.*



■ **Figure 5** Timed choreography for the Scheduled Task Protocol (S_{ST})

6 Applications and Implementation

Constructing global specifications Our theory can be easily applied and integrated with other works, to construct sound (i.e., satisfying safety, progress, and non-zenoness) timed global specifications, such as (syntactic) multiparty session types [6, 16], or graphical choreographies [8, 13, 19]. Thanks to Theorem 7, we can build on the algorithm in [14] to construct (syntactic) timed global types from CTAs. In Appendix [3], we give the formal definitions of the adaptation of the algorithm in [14]. Given an MC system S our algorithm generates a timed global type [6] equivalent to the original system S (i.e., its projections are timed bisimilar to those of S). This implies that if S is IE (resp. CE) then the constructed timed global type will also enjoy progress (resp. non-zenoness). Similarly, building on the algorithm in [19], we obtain a *graphical* representation reminiscent of BPMN Choreographies, see [8, 19]. When applied to the Scheduled Task Protocol, the algorithm adapted from [19] produces the choreography in Fig. 5; giving a much clearer specification for S_{ST} .

Implementation To assess the applicability and cost of our theory, we have integrated our theory into the tool first introduced in [19], which builds graphical choreographies from CFSMs. Our tool [3] (implemented in Haskell and using Z3) takes as input a textual representation of CTAs on which each condition (MC, IE, and CE) is checked for, and produces an equivalent choreography (such as the one in Fig. 5). The results of our experiments (executed on a Intel i7 computer, with 16GB of RAM) are below; where $|\mathcal{P}|$ is the number of machines, and $|N|$ (resp. $|\hookrightarrow|$) is the number of nodes (resp. transitions) in $STS(S)$.

S	$ \mathcal{P} $	$ N $	$ \hookrightarrow $	MC	IE	CE	s	$ \mathcal{P} $	$ N $	$ \hookrightarrow $	MC	IE	CE	s	
Running Example	3	4	4	✓	✓	✓	0.41	×4	12	256	1024	✓	✓	✓	28.49
Bargain	3	5	5	✓	✓	✓	0.44	×2	6	25	50	✓	✓	✓	12.30
Temp. calculation [6]	3	6	6	✓	✓	✓	0.45	×2	6	36	72	✓	✓	✓	9.24
Word Count [20]	3	6	6	✓	✓	✓	0.41	×2	6	36	72	✓	✓	✓	8.63
ATM (Template) [11]	3	9	8	✓	✓	✓	0.36	×3	9	729	1944	✓	✓	✓	94.01
ATM (Instance) [11]	3	9	8	✓	✓	✓	0.53	×3	9	729	1944	✓	✓	✓	96.09
Consumer-Producer [11]	2	1	1	✓	✓	✓	0.16	×5	10	1	5	✓	✓	✓	43.19
Fischers Mutual Excl. [5]	2	4	3	✓	✓	✓	0.21	×4	8	256	768	✓	✓	✓	3.19

Most of the protocols are taken from the literature and all are checked within a minute on average. For the sake of space, we have used small examples throughout the paper, however our benchmarks include bigger protocols (up-to 12 machines), which have comparable size with those we encountered through our collaboration with Cognizant [19, 23]. Since the size of the STS is the most critical parameter for scalability, we have tested systems consisting of the parallel composition of several instances of a protocol. For instance, *Running Example* ×4 is the parallel composition of four instances of S_{ST} , cf. Fig. 1.

7 Conclusions and Related Work

Our results are summarised in the table below. Multiparty compatibility (MC) gives (i) an equivalence between an MC system and a system consisting of the projections of its *STS*; and (ii) a sufficient condition for *safety*. MC and interaction enabling (IE) form a sufficient condition for *progress*; while MC and cycle enabling (CE) form a sufficient condition for *non-zenoness* (NZ). Together, MC, IE, and CE ensure safety, progress, NZ, and *eventual reception* (ER).

Property	$S \sim (STS(S) _p)_{pEP}$	Safety	Progress	Non-Zeno	ER
MC (Def. 4)	✓	✓	✗	✗	✗
MC+IE (Def. 12)	✓	✓	✓	✗	✗
MC+CE (Def. 16)	✓	✓	✗	✓	✗
MC+IE+CE	✓	✓	✓	✓	✓

Multiparty session types The work in [6] studies a typing system for a timed π -calculus using timed global types. A class of CTAs which are safe and have progress is given in [6] via projection of (well-formed) timed global types onto timed local types (which correspond to deterministic, non-mixed, and directed CTAs). Well-formedness yields conditions on CTAs that are more restrictive than the ones given in this paper. For instance, the system in Fig. 1, which is safe and enjoys progress, is ruled out by the conditions in [6]. In addition, this paper gives sufficient conditions for CTAs to belong to the class of safe CTAs with progress, which was left as an open problem in [6]. The construction of timed global types from either local types or CTAs is not addressed in [6]. Recently, [4] introduced a compliance and sub-typing relation for *binary* timed session types *without queues* (synchronous communication semantics). The existing works for constructing global specifications from local specifications [14, 18, 19] only apply to *untimed* models. Our conditions (IE and CE) are given independently of the definition of MC. The use of a more general notion of MC, as the one given in [19], would allow us to lift the assumptions that the machines are directed and have no mixed states (cf. § 2). Hence, we could capture more general timed choreographies.

Reachability and decidability When extending NTAs [2] with unbounded channels, reachability is no longer decidable in general [17]. Existing work tackles undecidability by restricting the network topologies [12, 17] or the channel size [1]. We give general (w.r.t. topology and channel size) decidable conditions ensuring that a configuration violating safety, progress, or NZ will not be reached. Observe that the scenario in Fig. 1 would be ruled out in [17] (its topology is not a polyforest) and in [1] (w_{WA} is unbounded). Our conditions are based, instead, on the conversation structures, which also enable the construction of global specifications.

Non-zeno conditions In § 5 we set the conditions for time divergence, by ruling out specifications in which the only way forward is a zeno behaviour. This condition is called time progress in [2] and it is built-in in the definition of runs of a TA. Several conditions have been proposed to ensure absence of non-zeno behaviours in TAs: some, e.g., [21], do not allow any zeno execution, and some, e.g., [7], and this work (cf. Def. 15), ensure that there is always a non-zeno way forward. The condition in [7] can be checked with a simple form of reachability analysis which introduced the notion of ‘escape’ from a zeno loop, which we also use. [7, 21] consider Networks of TAs (NTAs), which do not feature asynchrony nor unbounded channels.

Acknowledgements We would like to thank the ZDLC team at Cognizant for their stimulating conversations and Dominic Orchard for some (very useful) Haskell tips. This


work is partially supported by UK EPSRC projects EP/K034413/1, EP/K011715/1, and EP/L00058X/1; and by EU 7FP project under grant agreement 612985 (UPSCALE).

References

- 1 S. Akshay, Paul Gastin, Madhavan Mukund, and K. Narayan Kumar. Model checking time-constrained scenario-based specifications. In *FSTTCS*, volume 8 of *LIPICs*, pages 204–215, 2010.
- 2 Rajeev Alur and David L. Dill. A theory of timed automata. *TCS*, 126:183–235, 1994.
- 3 Webpage of this paper, 2015. <http://www.doc.ic.ac.uk/~jlange/cta/>.
- 4 Massimo Bartoletti, Tiziana Cimoli, Maurizio Murgia, Alessandro Sebastian Podda, and Livio Pompianu. Compliance and subtyping in timed session types. In *FORTE*, volume 9039 of *LNCS*, pages 161–177. Springer, 2015.
- 5 Johan Bengtsson et al. Uppaal - a tool suite for automatic verification of real-time systems. In *Hybrid Systems III*, volume 1066 of *LNCS*, pages 232–243. Springer, 1996.
- 6 Laura Bocchi, Weizhen Yang, and Nobuko Yoshida. Timed multiparty session types. In *CONCUR*, volume 8704 of *LNCS*, pages 419–434. Springer, 2014.
- 7 Howard Bowman and Rodolfo Gómez. How to stop time stopping. *FAC*, 18(4):459–493, 2006.
- 8 BPMN 2.0 Choreography, 2012. <http://en.bpmn-community.org/tutorials/34/>.
- 9 Daniel Brand and Pitro Zafropulo. On communicating finite-state machines. *JACM*, 30(2):323–342, 1983.
- 10 Gérard Cécé and Alain Finkel. Verification of programs with half-duplex communication. *ISC*, 202(2):166–190, 2005.
- 11 Prakash Chandrasekaran and Madhavan Mukund. Matching scenarios with timing constraints. In *FORMATS*, volume 4202 of *LNCS*, pages 98–112. Springer, 2006.
- 12 Lorenzo Clemente, Frédéric Herbreteau, Amelie Stainer, and Grégoire Sutre. Reachability of communicating timed processes. In *FOSSACS*, volume 7794 of *LNCS*, pages 81–96. Springer, 2013.
- 13 Pierre-Malo Deniérou and Nobuko Yoshida. Multiparty session types meet communicating automata. In *ESOP*, volume 7211 of *LNCS*, pages 194–213. Springer, 2012.
- 14 Pierre-Malo Deniérou and Nobuko Yoshida. Multiparty compatibility in communicating automata: Characterisation and synthesis of global session types. In *ICALP (2)*, volume 7966 of *LNCS*, pages 174–186. Springer, 2013.
- 15 Uno Holmer, Kim Guldstrand Larsen, and Wang Yi. Deciding properties of regular real time processes. In *CAV*, volume 575 of *LNCS*, pages 443–453. Springer, 1991.
- 16 Kohei Honda, Nobuko Yoshida, and Marco Carbone. Multiparty asynchronous session types. In *POPL*, pages 273–284. ACM, 2008.
- 17 Pavel Krcál and Wang Yi. Communicating timed automata: The more synchronous, the more difficult to verify. In *CAV*, volume 4144 of *LNCS*, pages 249–262, 2006.
- 18 Julien Lange and Emilio Tuosto. Synthesising Choreographies from Local Session Types. In *CONCUR*, volume 7454 of *LNCS*, pages 225–239. Springer, 2012.
- 19 Julien Lange, Emilio Tuosto, and Nobuko Yoshida. From communicating machines to graphical choreographies. In *POPL*, pages 221–232, 2015.
- 20 Rumyana Neykova, Laura Bocchi, and Nobuko Yoshida. Timed runtime monitoring for multiparty conversations. In *BEAT*, volume 162 of *EPTCS*, pages 19–26, 2014.
- 21 Stavros Tripakis. Verifying progress in timed systems. In *Formal Methods for Real-Time and Probabilistic Systems*, volume 1601 of *LNCS*, pages 299–314. Springer, 1999.
- 22 Stavros Tripakis, Sergio Yovine, and Ahmed Bouajjani. Checking timed büchi automata emptiness efficiently. *Formal Methods in System Design*, 26(3):267–292, 2005.
- 23 Zero Deviation Lifecycle. <http://www.zd1c.co>.

A Additional Examples for Interaction Enabling

This example illustrates a non-trivial set of constraints in cycle. The system below, where we let $k \in \mathbb{N}_{>1}$, be a constant, is not IE.



This system does not have progress, e.g., when $\nu(x) = 2k + 1$, machine **s** will get stuck the state highlighted in black. It is ruled out by IE since we have to check n in $STS(S)$ that correspond to the configuration where both machines are in their highlighted state is IE. This amount to checking that the pair (n, e) is \forall -PE for **s**, with $e = (\mathbf{s} \rightarrow \mathbf{r} : b; (\dots), \emptyset; (\dots), \emptyset)$. However, the path $n_0 \xrightarrow{a} n \xrightarrow{b} n' \xrightarrow{c} n$ violates the condition, i.e.,

$$\begin{aligned}
 & \forall v_1, v_2, v_3 \exists v_4 : \\
 & v_1 < 1 \wedge \\
 & (v_2 = 2 \vee v_2 = 4 \vee \dots \vee v_2 = 2k) \wedge \\
 & (v_3 = 3 \vee v_3 = 5 \vee \dots \vee v_3 = 2k + 1) \wedge \\
 & v_1 \leq v_2 \wedge v_2 \leq v_3 \\
 & \implies (v_4 = 2 \vee v_4 = 4 \vee \dots \vee v_4 = 2k) \wedge v_1 \leq v_4 \wedge v_2 \leq v_4 \wedge v_3 \leq v_4
 \end{aligned}$$

does not hold. In addition, observe that this system is not CE (due to the upper bound on x , which is never reset).

B Additional Definitions for Section 5

We give the formal definition of upper and lower bounds, used in § 5.

$$g \text{ is UB for } x \iff \begin{cases} g = x \leq c, \\ g = \neg g' \text{ and not } g' \text{ is UB for } x \\ g = g_1 \wedge g_2 \text{ and } \exists i \in \{1, 2\} : g_i \text{ is UB for } x \end{cases}$$

$$g \text{ is SP} \iff \begin{cases} g = x \leq c \text{ with } c \in \mathbb{Q}_{>0}, \\ g = c \leq x, \\ g = \neg g' \text{ and not } g' \text{ is SP} \\ g = g_1 \wedge g_2 \text{ and } \forall i \in \{1, 2\} : g_i \text{ is SP} \\ g = \text{true} \end{cases}$$


C Additional Examples for Zeno Systems

In this section we give additional examples to illustrate the notions of zeno configuration and zeno system.

C.1 Progress vs non-zenoness

First we show that progress and non-zenoness are orthogonal properties. The example in § 5 shows that a MC and IE system (hence enjoying progress) may or may not be zeno. The

example below shows that a system *not* enjoying progress may be non-zero.




Consider the configuration $s = ((q_1, q'_1); \vec{\varepsilon}; \nu)$. s is an unsuccessful reception configuration since s cannot receive b at time $x = 4$ (b will be sent only at a time $y > 4$) hence $s \rightarrow^* s'$ only for an empty step with $s = s'$. Taking an arbitrary t , the (non-zero configuration) condition

for all $s' = (q'; \vec{w}'; \nu')$, $s \rightarrow^* s'$ implies $\nu'(\hat{x}) < t$ and $s' \rightarrow s''$, for some s''

does not hold since $s = s' \not\rightarrow$, hence s is a non-zero configuration. Similarly, one can show that all the other reachable configurations of the system above are non-zero.

C.2 Progress and non-zenoness

The example in this paragraph shows the application of the definition of non-zero configuration to a configuration that: (1) has a zeno loop, (2) *and* all the other non-zero way forward are not viable (stuck).




Consider the configuration $s = ((q_1, q'_1); \vec{\varepsilon}; \nu)$ with $\nu(x) = 1$. Since $\nu(x) = 1$ machine s can no longer send c (having constraint $x < 1$), hence $s \rightarrow^* s'$ only for s' of the form $((q_1, q'_1); \vec{w}'; \nu')$ for some \vec{w}' and $\nu'(x) < 3$. Similarly s' can always and only move to configurations having the same form. We observe that s (and s') is a zeno configuration as $\nu'(\hat{x}) = t$ can never be reached for $t = 3$.

Note that s above is not an unsuccessful reception configuration: constraint $y = 3$ of r is satisfiable at some point in the future, although this point will never be reached due to zeno behaviour.

C.3 Non-zero configurations and systems


The example in this paragraph illustrates how the definition of non-zero configuration applies to the definition of non-zero system. In particular, we show how the definitions capture the intended property: some zeno-behaviour may be possible but there is always a non-zero way forward.



Configuration $s_0 = (q_0, q'_0; \vec{\varepsilon}; \nu_0)$ is non-zero as there exists a future configuration, of the form $(q_2, q'_2; \vec{w}'; \nu)$ which allows time to diverge (note that final states allow time

divergence). On the other hand, configurations of the form $(q_1, q'_1; \vec{w}'; \nu')$ are obviously zeno configurations. The overall system does *not* satisfy Def. 15 (non-zenoness) since it can reach some zeno configuration (i.e., with no possible non-zeno way forward).

If we modify the example above by *moving the outgoing non-zeno path* to the state with the outgoing zeno loop, we obtain a non-zeno system. For instance, in the example below configurations of the form $(q_1, q'_1; \vec{w}'; \nu')$ have a possible non-zeno way forward. The system below is non-zeno.



D Construction of Timed Global Types

In this section, we give a new approach to construct timed global types from CTAs; it is adapted from the construction algorithm in [14] while using projection functions similar to the ones given in [19].

A timed global type is generated from the following grammar:

$$G ::= \mathbf{s} \rightarrow \mathbf{r} : \{a_i \langle \mathbf{A}_i \rangle . G_i\}_{i \in I} \mid \mu \mathbf{t} . G \mid \mathbf{t} \mid \mathbf{end} \quad \text{where } \mathbf{A}_i ::= g_{\mathbf{s}i}, \lambda_{\mathbf{s}i}; g_{\mathbf{r}i}, \lambda_{\mathbf{r}i}$$

Type $\mathbf{s} \rightarrow \mathbf{r} : \{a_i \langle \mathbf{A}_i \rangle . G_i\}_{i \in I}$ models an interaction: \mathbf{s} chooses a branch $i \in I$, and sends a message a_i to \mathbf{r} , the interaction then continues as specified by G_i . Each branch is annotated with a time assertion $\mathbf{A}_i = g_{\mathbf{s}i}, \lambda_{\mathbf{s}i}; g_{\mathbf{r}i}, \lambda_{\mathbf{r}i}$, where $g_{\mathbf{s}i}$ and $\lambda_{\mathbf{s}i}$ are the clock constraint and reset for the send action, and $g_{\mathbf{r}i}$ and $\lambda_{\mathbf{r}i}$ are for the receive action. We write $\mathbf{s} \rightarrow \mathbf{r} : a \langle \mathbf{A} \rangle . G$ for interactions with one branch. Type $\mu \mathbf{t} . G$ is a recursive type, \mathbf{t} is a type variable, and \mathbf{end} is for termination.

The main transformation from CTAs to timed global types concerns the removal of branches in $STS(S)$ that correspond to concurrent executions of the system (i.e., interleaving). This transformation is realised by Definition 20 below, relying on the following auxiliary function which separates a set \hat{E} of events according to the (two) participants they refer to:

$$\text{indep}(\hat{E}) \stackrel{\text{def}}{=} \left\{ E' \subseteq \hat{E} \mid \forall e, e' \in E' : \text{id}(e) = \text{id}(e') \wedge \forall e \in \hat{E} \setminus E' \forall e' \in E' : \text{id}(e) \neq \text{id}(e') \right\}$$

We define $n^\bullet \stackrel{\text{def}}{=} \{e \mid \exists n' \in N : n \xrightarrow{e} n'\}$.

► **Definition 20 (Reduction).** Let $STS(S) = (N, n_0, \leftrightarrow, E)$. The transition system $RTS(S) = (N', n_0, \mapsto, E)$ is a *reduction of $STS(S)$* , if the following conditions holds: (1) $n_0 \in N'$, $N' \subseteq N$, and $\mapsto \subseteq \leftrightarrow$; (2) $\forall n \in N : n^\circ = \emptyset \iff n^\bullet = \emptyset$; (3) for all cycle $n_1 \cdot e_1 \cdots n_{k-1} \cdot e_{k-1} \cdot n_1$ in $RTS(S)$: $\forall 1 \leq i < k : e \in n_i^\bullet \implies \exists 1 \leq j < k : e = e_j$; (4) $\forall n \in N' : n^\circ \in \text{indep}(n^\bullet)$ where $n^\circ \stackrel{\text{def}}{=} \{e \mid \exists n' \in N' : n \xrightarrow{e} n'\}$. \diamond

Condition (1) ensures that the $RTS(S)$ is a sub-graph of $STS(S)$; condition (2) says that if a node n has at least one successor in $STS(S)$, then n has also a successor in $RTS(S)$; condition (3) is adapted from partial order reduction techniques and ensures that no transition is forever disabled in a cycle of $RTS(S)$; and condition (4) selects only the outgoing transitions of node which involve the same (two) participants.

For instance, in Section 4.2, we gave the $STS(S)$ (right) of the system consisting of machines \mathbf{s}_1 , \mathbf{r}_1 , \mathbf{s}_2 , and \mathbf{r}_2 (left). The reduction of $STS(S)$, i.e., $RTS(S)$, is the sub-graph of $STS(S)$ consisting of edges:

- $(n_0, (\mathbf{s}_1 \rightarrow \mathbf{r}_1 : a; x > 2, \emptyset; \text{true}, \emptyset), n_1)$
- $(n_1, (\mathbf{s}_2 \rightarrow \mathbf{r}_2 : c; \text{true}, \emptyset; \text{true}, \emptyset), n_2)$
- $(n_2, (\mathbf{s}_1 \rightarrow \mathbf{r}_1 : b; x < 2, \emptyset; \text{true}, \emptyset), n_3)$
- $(n_3, (\mathbf{s}_2 \rightarrow \mathbf{r}_2 : d; \text{true}, \emptyset; \text{true}, \emptyset), n_0)$

Observe that the reduction preserves all the events of $STS(S)$ (see below).

The lemma below is useful to show the completeness of the transformation as well as the equivalence between the original system and its projections; it follows from Definitions 3 and 20, and the definition of $\text{indep}(_)$.

► **Lemma 21.** *The following holds:*

1. For all S , there exists a reduction $RTS(S)$ of $STS(S)$.
2. For all $\mathbf{p} \in \mathcal{P}$, $STS(S)|_{\mathbf{p}} = RTS(S)|_{\mathbf{p}}$
3. Let $RTS(S) = (N, n_0, \mapsto, E)$, then $\forall n \in N, \forall e, e' \in E : (n \xrightarrow{e} \wedge n \xrightarrow{e'} \implies \text{sid}(e) = \text{sid}(e') \wedge \text{rid}(e) = \text{rid}(e'))$.

Proof. (1) and (3) follow since we assumed that the machines are directed and non-mixed, and once we observe that for all $n \in N$ and for all $e_1, e_2 \in E$, with $\text{id}(e_1) \cap \text{id}(e_2) = \emptyset$, such that $n \xrightarrow{e_i} n_i$, there is $n' \in N$ such that $n_i \xrightarrow{e_j} n'$, with $i \neq j \in \{1, 2\}$ (two independent event commute).

For (2), we use the fact that independent event commute (i.e, we do not forget event in a sequence) and the fact that no event is “left out” in cycles. ◀

We now give an algorithm to construct a syntactic timed global type from $RTS(S)$.

► **Definition 22 (Construction of timed global types).** Suppose S is multiparty compatible. Let $RTS(S) = (N, n_0, \mapsto, E)$, then the function $G(n, N')$ is defined as follows:

1. if $n^\circ = \emptyset$, then return **end**;
2. if $n \in N'$, then return \mathbf{t}_n ;
3. if $n \notin N'$, then, letting $n^\circ = \{e_i = (\mathbf{s} \rightarrow \mathbf{r} : a_i; g_{\mathbf{s}i}, \lambda_{\mathbf{s}i}; g_{\mathbf{r}i}, \lambda_{\mathbf{r}i}) \mid 1 \leq i \leq k\}$, return $\mu \mathbf{t}_n. \mathbf{s} \rightarrow \mathbf{r} : \{a_i \langle g_{\mathbf{s}i}, \lambda_{\mathbf{s}i}; g_{\mathbf{r}i}, \lambda_{\mathbf{r}i} \rangle. G(n_i, N' \cup \{n\})\}$, with $n \xrightarrow{e_i} n_i$.

Finally, erase any unnecessary $\mu \mathbf{t}$ if $\mathbf{t} \notin G(n, N')$. We write $G(S) \stackrel{\text{def}}{=} G(n_0, \emptyset)$ for the timed global type obtained from a reduction $RTS(S)$ of S . ◊

Synchronous transition systems and timed global types can be *projected* onto participants so to obtain CTAs. We define a generic projection function from a transition system (whose transitions are labelled by events e) onto machines.

► **Definition 23 (Projection).** The projection of a transition system $T = (N, n_0, \leftrightarrow)$ on participant \mathbf{p} , written $T|_{\mathbf{p}}$, is the (minimised w.r.t. language equivalence) automaton $(Q, q_0, \mathcal{X}, \delta)$ where $Q = N$, $q_0 = n_0$, $\delta \subseteq Q \times \text{Act}_{\mathcal{X}} \cup \{\varepsilon\} \times Q$ is s.t. $(n_1, e|_{\mathbf{p}}, n_2) \in \delta \iff n_1 \xrightarrow{e} n_2$, and \mathcal{X} is the set of clocks in δ . ◊

The projection of a timed global type G , written $G|_{\mathbf{p}}$, is defined by translating G into a transition system, then applying Def. 23.

Hereafter, we assume that bound (recursion) variables are pairwise distinct and that there is no free variable in G .

► **Definition 24 (TS of G).** The transition system of G is obtained as follows. Let $TG(G) \stackrel{\text{def}}{=} (N, n_0, \leftrightarrow)$, with $TG(G, n_0, \circ) = (N, \leftrightarrow)$; where \circ is the empty map and $TG(G, n, \Gamma)$ is defined as follows:²

² We assume that each node n_0, n, n' and n_i is freshly created.

1. if $G = \mathbf{s} \rightarrow \mathbf{r} : \{a_i \langle g_{si}, \lambda_{si}; g_{ri}, \lambda_{ri} \rangle . G_i\}_{i \in I}$ then return

$$(\{n\} \cup \cup_{i \in I} N_i, \cup_{i \in I} \{(n, (\mathbf{s} \rightarrow \mathbf{r} : a_i; g_{si}, \lambda_{si}; g_{ri}, \lambda_{ri}), n_i)\} \cup \hookrightarrow_i)$$

where $(N_i, \hookrightarrow_i) = TG(G_i, n_i, \Gamma)$ for $i \in I$.

2. if $G = \mu \mathbf{t} . G'$ then return $(\{n\} \cup N', \{(n, \varepsilon, n')\} \cup \hookrightarrow')$ where $(N', \hookrightarrow') = TG(G', n', \Gamma \uplus \{\mathbf{t} \mapsto n\})$.
3. if $G = \mathbf{end}$, then return $(\{n\}, \emptyset)$.
4. if $G = \mathbf{t}$, then return $(\{n\}, \{(n, \varepsilon, \Gamma(\mathbf{t}))\})$.

◇

In the Scheduled Task Protocol, our construction algorithm yields the following timed global type:

$$\begin{aligned} G_{\text{ST}} = & \mathbf{U} \rightarrow \mathbf{W} : \text{task} \langle x < 1, x := 0; y = 1, y := 0, y' := 0 \rangle . \\ & \mu \mathbf{t} . \mathbf{W} \rightarrow \mathbf{A} : \{ \text{data} \langle y < 2 \wedge y' < 10, y := 0; z = 3, z := 0 \rangle . \mathbf{t}, \\ & \text{stop} \langle y < 1; z = 3, z := 0 \rangle . \mathbf{A} \rightarrow \mathbf{U} : \text{result} \langle z \leq 5; x \leq 20 \rangle \} \end{aligned}$$

The projection of a timed global type G , written $G \downarrow_{\mathbf{p}}$, is given by translating G into a transition system, where each transitions is *projected* onto \mathbf{p} (cf. projection of an event e). As in the case of projection of a *STS*, the projections of G_{ST} onto \mathbf{U} , \mathbf{W} , and \mathbf{A} are isomorphic to the system S_{ST} in Fig. 1.

► **Theorem 25** (Equivalence 2). *If $S = (M_{\mathbf{p}})_{\mathbf{p} \in \mathcal{P}}$ is MC then $S \sim (G(S) \downarrow_{\mathbf{p}})_{\mathbf{p} \in \mathcal{P}}$.*

E Extended Related Work

E.1 Multiparty session types

The conditions required in [6] are more restrictive than the ones required in this paper. For instance, the system in Fig. 1, which is safe and enjoys progress, is ruled out by the conditions in [6]. This is due, in particular, to a condition called *infinite satisfiability* [6] which requires that, in each loop, either all clocks are reset or there is no upper bound on any clock. The example in Fig. 1 is not infinite satisfiable since z is never reset and there is an upper bound in constraint $z < 10$. We relaxed infinite satisfiability by using a notion of ‘escape’ along the lines of [7].

E.2 Non-zeno conditions

The condition given in [21], called Strong Non-Zenoness (SNZ), is syntactic and requires that time elapses of some $c \geq 1$ in all loop cycles. This condition disallows zeno executions altogether. We allow *possible* zeno behaviour as long as there is a non-zeno way to continue the computation.

E.3 Reachability and decidability

Although complex (NL-complete with one clock [33], NP-hard [33] and PSPACE-complete [30] with two clocks, and PSPACE-hard with three or more clocks [28]) reachability is a decidable problem for Timed Automata and NTAs [2]. These models are well-supported by model-checking tools such as UPPAAL [5] and KRONOS [25]. Several works tackle this complexity by using abstraction [29], approximations [31], local time semantics [24], or considering limited scenarios (with no recursion) [26]. Unfortunately, when one extends NTAs with unbounded

data structures [32] or unbounded channels [12] (as in the case of CTAs) decidability no longer holds in the general case. In the case of CTAs, existing work [12, 17] tackles undecidability by limiting the network topologies. For CTAs with dense time and urgent semantics (i.e., internal actions are disabled if a message can be received) decidability is limited to topologies in the form $(M_1, M_2, w_{1,2})$ while all other topologies have the power of Turing machines [17]. With a non-urgent semantics, reachability is decidable if the topology is a polyforest (i.e., it is a direct acyclic graph with no undirected cycles) [12]. Asynchronous communication is considered in [1] using Message Sequence Charts (MSCs) annotated with time constraints and where participants communicate through FIFO channels. For decidability of verification, the topology of the participants involved in each loop should be a single strongly connected component, which implies that each message is implicitly acknowledged, hence channels have an upper bound. Unlike [1, 12, 17], our verification framework does not require limitations of the topology or buffer size. Note, for example, that the scenario in Fig. 1 is a not polyforest, and it does not satisfy the condition on the topology posed by [1] due to the iterative one-way messaging from W to A (i.e., channel w_{WA} is unbounded since). Reachability is also decidable when restricting the unbounded data structures [32] or, in the untimed scenario, the channels (i.e., using unordered ‘bag’ channels instead of FIFO channels) [27]. Our conditions are based, instead, on the conversation structures, which also enables the construction of global specifications.

Additional References for the Appendix

- 24 J. Bengtsson, B. Jonsson, J. Lilius, and W. Yi. Partial order reductions for timed systems. In *CONCUR*, volume 1466 of *LNCS*, pages 485–500. Springer, 1998.
- 25 M. Bozga, C. Daws, O. Maler, A. Olivero, S. Tripakis, and S. Yovine. Kronos: A model-checking tool for real-time systems. In *FTRTFT*, volume 1486 of *LNCS*, pages 298–302. Springer, 1998.
- 26 M. Cambroner, G. Díaz, V. Valero, and E. Martínez. Validation and verification of web services choreographies by using timed automata. *J. Log. Algebr. Program.*, 80(1):25–49, 2011.
- 27 L. Clemente, F. Herbreteau, and G. Sutre. Decidable topologies for communicating automata with FIFO and bag channels. In *CONCUR*, pages 281–296, 2014.
- 28 C. Courcoubetis and M. Yannakakis. Minimum and maximum delay problems in real-time systems. *Formal Methods in System Design*, 1(4):385–415, 1992.
- 29 C. Daws and S. Tripakis. Model checking of real-time reachability properties using abstractions. In *TACAS*, LNCS, pages 313–329. Springer, 1998.
- 30 J. Fearnley and M. Jurdzinski. Reachability in two-clock timed automata is pspace-complete. *CoRR*, abs/1302.3109, 2013.
- 31 F. Herbreteau, D. Kini, B. Srivathsan, and I. Walukiewicz. Using non-convex approximations for efficient analysis of timed automata. In *FSTTCS*, volume 13 of *LIPICs*, pages 78–89. Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, 2011.
- 32 R. Lanotte, A. Maggiolo-Schettini, and A. Troina. Reachability results for timed automata with unbounded data structures. *Acta Inf.*, 47(5-6):279–311, 2010.
- 33 F. Laroussinie, N. Markey, and P. Schnoebelen. Model checking timed automata with one or two clocks. In *CONCUR*, volume 3170 of *LNCS*, pages 387–401. Springer, 2004.

F Proofs

F.1 Correspondence between S and $STS(S)$

In this section, we introduce a formal correspondence between the reachable configurations of S and its synchronous transition system ($STS(S)$). This correspondence does not rely on time constraints because: (i) $STS(S)$ disregards time, (ii) the machines are deterministic, and (iii) the reachable configurations of a system without its time constraints form a superset of the original system (modulo clock valuations); cf. Prop. F.1. Hence, we will first in this subsection ignore time constraints.

► **Definition 26** (Untimed systems). The untimed system of a system $S = (M_p)_{p \in \mathcal{P}}$, where each $M_p = (Q_p, q_{0p}, \mathcal{X}_p, \delta_p)$, written $UT(S)$, is the system S where every transition $(q_p, \ell, q'_p) \in \delta_p$ in every machine M_p is replaced by $(q_p, \hat{\ell}, q'_p)$ where $\hat{\ell}$ is the same as ℓ except that the guard $\text{guard}(\hat{\ell}) = \text{true}$ and $\text{reset}(\hat{\ell}) = \emptyset$. \diamond

Since an untimed system does not have any time constraints, we can safely ignore clock valuations, i.e., we write $(\vec{q}; \vec{w})$ for a configuration $(\vec{q}; \vec{w}; \nu) \in RS(UT(S))$. Similarly, we write $\text{sr}!a$ (resp. $\text{sr}?a$) instead of $\text{sr}!a(\text{true}, \emptyset)$ (resp. $\text{sr}?a(\text{true}, \emptyset)$), and $\mathbf{s} \rightarrow \mathbf{r} : a$ for $(\mathbf{s} \rightarrow \mathbf{r} : a; \text{true}, \emptyset; \text{true}, \emptyset)$. We formalise the relationship between $UT(S)$ and S in the proposition below.

► **Proposition F.1.**

$$(\vec{q}; \vec{w}; \nu) \in RS(S) \implies (\vec{q}; \vec{w}) \in RS(UT(S))$$

We now give safety properties guaranteed by multiparty compatible (untimed) systems. First, we define the projection of a sequence of events φ onto channel sr , written $\varphi|_{\text{sr}}$, as follows:

$$\varphi|_{\text{sr}} = \begin{cases} a \cdot \varphi'|_{\text{sr}} & \text{if } \varphi = \mathbf{s} \rightarrow \mathbf{r} : a \cdot \varphi' \\ \varphi'|_{\text{sr}} & \text{if } \varphi = \mathbf{s}' \rightarrow \mathbf{r}' : a \cdot \varphi' \text{ and } \mathbf{s}'\mathbf{r}' \neq \text{sr} \\ \varepsilon & \text{otherwise} \end{cases}$$

Def. 27 below introduces a correspondence between a (untimed) configuration $(\vec{q}; \vec{w})$ and a sequence of events φ linking two nodes n and n' in $STS(S)$. Below given $n = (q_p)_{p \in \mathcal{P}}$, we write $n[\mathbf{p}]$ for q_p .

► **Definition 27** (Correspondence with $STS(S)$). Given a system S , $STS(S) = (N, n_0, \leftrightarrow, E)$, $s = (\vec{q}; \vec{w}) \in RS(S)$ and, $n, n' \in N$, we write $s \dashv\vdash (n, \varphi, n')$ iff

$$\forall \text{sr} \in C : \exists n_1, n_2 \in N : \quad n \xrightarrow{\varphi_1} n_1 \xrightarrow{\varphi_2} n_2 \xrightarrow{\varphi_3} n' \text{ with } \varphi = \varphi_1 \cdot \varphi_2 \cdot \varphi_3 \\ n_1[\mathbf{r}] = q_{\mathbf{r}} \wedge n_2[\mathbf{s}] = q_{\mathbf{s}} \wedge \varphi_2|_{\text{sr}} = w_{\text{sr}}$$

\diamond

Later, we will often consider only *minimal* correspondences, as defined in Def. 28 below.

► **Definition 28** (Minimality). Given a system S , $STS(S) = (N, n_0, \leftrightarrow, E)$, $s = (\vec{q}; \vec{w}) \in RS(S)$ and, $n, n' \in N$, we say that $s \dashv\vdash (n, \varphi, n')$ is *minimal* if

$$\forall n'' \in N : n \xrightarrow{\varphi_1} n'' \xrightarrow{\varphi_2} n' \wedge \varphi = \varphi_1 \cdot \varphi_2 \implies \neg (s \dashv\vdash (n, \varphi_1, n'') \vee s \dashv\vdash (n'', \varphi_2, n'))$$

\diamond

► **Lemma 29.** *Let S be multiparty compatible and $STS(S) = (N, n_0, \leftrightarrow, E)$. For all $s = (\vec{q}; \vec{w}) \in RS(S)$, there exist φ and $n, n' \in N$ such that $s \multimap (n, \varphi, n')$.*

Proof. We show this result by induction on the length of the execution reaching s (from s_0). If the execution has length 0, then $s = s_0$ and we have $\vec{q} = n_0$ and $w_{\mathbf{sr}} = \varepsilon$ for all $\mathbf{sr} \in C$. The result follows trivially from the definition of $STS(S)$, with $s \multimap (n_0, \varepsilon, n_0)$.

Assume the result holds for any s reachable with n steps and let us show that the result still holds with an $n + 1$ step execution. Hence, we have $s \multimap (n, \varphi, n')$ (minimal, without loss of generality) and $s \xrightarrow{\ell} s'$, and we have to show that $s' \multimap (n'', \varphi', n''')$, for some n'', n''' and φ' .

Case 1. Assume $\ell = \mathbf{sr}?a$. Then, by the semantics of CFSMs, we must have $w_{\mathbf{sr}} = a w'_{\mathbf{sr}}$ and $(q_{\mathbf{r}}, \mathbf{sr}?a, q'_{\mathbf{r}}) \in \delta_{\mathbf{r}}$. While, by Def. 27, we must have

$$\exists n_1, n_2 \in N : n \xrightarrow{\varphi_1} n_1 \xrightarrow{\varphi_2} n_2 \xrightarrow{\mathbf{s} \rightarrow \mathbf{r} a} n_3 \xrightarrow{\varphi_3} n_4 \xrightarrow{\varphi_4} n'$$

with $\varphi = \varphi_1 \cdot \varphi_2 \cdot \mathbf{s} \rightarrow \mathbf{r} : a \cdot \varphi_3 \cdot \varphi_4$, $n_1[\mathbf{r}] = q_{\mathbf{r}}$, $n_4[\mathbf{s}] = q_{\mathbf{s}}$, and $\varphi_2 \cdot \mathbf{s} \rightarrow \mathbf{r} : a \cdot \varphi_3 \downarrow_{\mathbf{sr}} = w_{\mathbf{sr}}$. We can safely assume that $\varphi_2 \downarrow_{\mathbf{s}} = \varepsilon$, otherwise there would be a previous occurrence of $\mathbf{s} \rightarrow \mathbf{r} : a$ in φ .

We have to show that $s' \multimap (n, \varphi, n')$, i.e., the correspondence is preserved. Let n'' be the first node on the walk following φ such that $n''[\mathbf{r}] = q_{\mathbf{r}}$, i.e., we have

$$n \xrightarrow{\varphi_5} n'' \xrightarrow{\varphi_6} n_2 \xrightarrow{\mathbf{s} \rightarrow \mathbf{r} a} n_3 \xrightarrow{\varphi_7} n' \quad \text{with } \varphi = \varphi_5 \cdot \varphi_6 \cdot \mathbf{s} \rightarrow \mathbf{r} : a \cdot \varphi_7$$

Hence, by Def. 27 and since only machine \mathbf{r} makes a move, we have to show that

$$\forall \mathbf{p} \neq \mathbf{r} \in \mathcal{P} : \varphi_6 \downarrow_{\mathbf{pr}} = \varepsilon \wedge \varphi_6 \downarrow_{\mathbf{rp}} = \varepsilon$$

By contradiction, if there is \mathbf{rp} or \mathbf{pr} such that $\varphi_6 \downarrow_{\mathbf{rp}} \neq \varepsilon$ or $\varphi_6 \downarrow_{\mathbf{pr}} \neq \varepsilon$, then there must be an event involving \mathbf{r} in φ_6 . Let e be the first event in φ_6 involving \mathbf{r} .

(a) If $e = \mathbf{r} \rightarrow \mathbf{p} : b$, then we obtain a contradiction with the assumption that the machines have no mixed states.

(b) If $e = \mathbf{p} \rightarrow \mathbf{r} : b$, then, by directedness, we must have $\mathbf{p} = \mathbf{s}$. Assuming $b \neq a$, if b is in the queue $w_{\mathbf{sr}}$, then it is at the top of the queue $w_{\mathbf{sr}}$, which contradicts the fact that we considered the first such element (see assumption wrt $\varphi_2 \downarrow_{\mathbf{sr}} = \varepsilon$ above).

Case 2. Assume $\ell = \mathbf{sr}!a$. By Def. 27, we must have

$$\exists n_1, n_2 \in N : n \xrightarrow{\varphi_1} n_1 \xrightarrow{\varphi_2} n_2 \xrightarrow{\varphi_3} n'$$

with $\varphi = \varphi_1 \cdot \varphi_2 \cdot \varphi_3$, $n_1[\mathbf{r}] = q_{\mathbf{r}}$, and $n_2[\mathbf{s}] = q_{\mathbf{s}}$ (with $q_{\mathbf{s}}$ a sending state).

By multiparty compatibility, we have $n_2 \xrightarrow{\varphi' \cdot \mathbf{s} \rightarrow \mathbf{r} a} n'$ with $\varphi' \downarrow_{\mathbf{s}} = \varepsilon$.


(a) If $\varphi_3 = \varphi'_3 \cdot \mathbf{s} \rightarrow \mathbf{r} : a \cdot \varphi''_3$, with $\varphi'_3 \downarrow_{\mathbf{s}} = \varepsilon$, then we have the result immediately since must have $w_{\mathbf{ps}} = w_{\mathbf{rp}} = \varepsilon$ for all $\mathbf{p} \neq \mathbf{s} \in \mathcal{P}$ (since \mathbf{s} is not involved in φ'_3).

(b) If $\varphi_3 \neq \varphi'_3 \cdot \mathbf{s} \rightarrow \mathbf{r} : a \cdot \varphi''_3$, with $\varphi'_3 \downarrow_{\mathbf{s}} = \varepsilon$, then there are two sub-cases, either

1. $\varphi_3 \downarrow_{\mathbf{s}} = \varepsilon$, i.e., \mathbf{s} is not involved in φ_3 , or
2. $\varphi_3 = \varphi'_3 \cdot \mathbf{s} \rightarrow \mathbf{r} : b \cdot \varphi''_3$, with $\varphi'_3 \downarrow_{\mathbf{s}} = \varepsilon$.

For (1), we have the result immediately by multiparty compatibility, i.e., we must have $n' \xrightarrow{\varphi'' \cdot \mathbf{s} \rightarrow \mathbf{r} a} n''$ (with $\varphi'' \downarrow_{\mathbf{s}} = \varepsilon$). Hence, we have $s' \multimap (n, \varphi \cdot \varphi'' \cdot \mathbf{s} \rightarrow \mathbf{r} : a, n'')$.

In the case (2), we have the following situation (by multiparty compatibility):



with $\mathbf{s} \rightarrow \mathbf{r} : a \in \varphi_4$ and $\mathbf{s} \rightarrow \mathbf{r} : b \in \varphi_3$.

Hence, we only have to show that $s \multimap (n, \varphi_1 \cdot \varphi_2 \cdot \varphi_4, n'')$. Then we can return to the previous case. Let $\varphi' = \varphi_1 \cdot \varphi_2 \cdot \varphi_4$. We show that for all $\mathbf{pq} \in C : \exists n_3, n_4 \in N$ such that, letting $\varphi' = \varphi'_1 \cdot \varphi'_2 \cdot \varphi'_3$,

$$n \xrightarrow{\varphi'_1} n'_1 \xrightarrow{\varphi'_2} n'_2 \xrightarrow{\varphi'_3} n'' \text{ with } n'_1[\mathbf{r}] = q_q, n'_2[\mathbf{s}] = q_p, \text{ and } w_{\mathbf{pq}} = \varphi'_2|_{\mathbf{pq}}$$

Without loss of generality, assume that φ'_2 is the shortest possible.

First, we note that if $\varphi'_1 \cdot \varphi'_2$ is a prefix of $\varphi_1 \cdot \varphi_2$, then the result holds by the assumption that $s \multimap (n, \varphi, n')$.

Otherwise, there is a suffix of $\varphi'_1 \cdot \varphi'_2$ which is a prefix of φ_3 (there is an overlapping), call it $\hat{\varphi}$. We have to show that $\varphi_4 = \varphi_5 \cdot \varphi_6$ with $\varphi_5|_{\mathbf{pq}} = \hat{\varphi}|_{\mathbf{pq}}$. Since φ_2 is the shortest possible and there is an overlapping, we must have $w_{\mathbf{pq}} = \rho \cdot \rho'$ with $\rho' = \hat{\varphi}|_{\mathbf{pq}}$. If the queue is empty, we have the result immediately. Otherwise, it implies that machine \mathbf{p} is a sending state in node n_2 , which implies, by multiparty compatibility, that its behaviour must be the same in both branches, hence we obtain the result. \blacktriangleleft

► Lemma 30. *Let S be a multiparty compatible system and $STS(S) = (N, n_0, \multimap, E)$. For all $s = (\vec{q}; \vec{w}) \in RS(S)$, such that $s \multimap (n, \varphi, n')$, there is s' stable such that $s \rightarrow^* s' = (\vec{q}'; \vec{w}')$ and $\vec{q}' = n'$.*

Proof. By induction on the length of φ , assuming that $s \multimap (n, \varphi, n')$ is minimal. If $\varphi = \varepsilon$, then s is stable, hence we have the result with $s \multimap (n, \varepsilon, n)$ and $s = s'$.

Assume the result holds for path of length k , $\varphi = \mathbf{s} \rightarrow \mathbf{r} : a \cdot \varphi'$ (with $|\varphi| = k$), and that $n \xrightarrow{\mathbf{s} \rightarrow \mathbf{r} : a} n''$. There are four cases:

1. If $q_{\mathbf{s}} = n[\mathbf{s}]$ and $q_{\mathbf{r}} = n[\mathbf{r}]$, then we have $s \xrightarrow{\mathbf{sr}!a} s'$ and $s' \multimap (n'', \varphi', n')$. We have the result by induction hypothesis.
2. If $q_{\mathbf{s}} \neq n[\mathbf{s}]$ and $q_{\mathbf{r}} \neq n[\mathbf{r}]$, then we have a contradiction with the fact that $s \multimap (n, \varphi, n')$ is minimal (i.e., φ' is a sufficient path).
3. If $q_{\mathbf{s}} = n[\mathbf{s}]$ and $q_{\mathbf{r}} \neq n[\mathbf{r}]$, then, by Def. 27, the content of the queue $w_{\mathbf{sr}}$ cannot rely on the $\mathbf{s} \rightarrow \mathbf{r} : a$ (since $q_{\mathbf{r}} \neq n[\mathbf{r}]$); hence we have a contradiction with the fact that $s \multimap (n, \varphi, n')$ is minimal, as above.
4. If $q_{\mathbf{s}} \neq n[\mathbf{s}]$ and $q_{\mathbf{r}} = n[\mathbf{r}]$, then we have $w_{\mathbf{sr}} = a w'_{\mathbf{sr}}$. Thus, we have $s \xrightarrow{\mathbf{sr}^?a} s'$ and $s' \multimap (n'', \varphi', n')$. We have the result by induction hypothesis. \blacktriangleleft

► Lemma 31. *Let S be a multiparty compatible system. For all $s = (\vec{q}; \vec{w}) \in RS(S)$ and $\mathbf{r} \in \mathcal{P}$, if $(q_{\mathbf{r}}, \mathbf{sr}^?b, q'_{\mathbf{r}}) \in \delta_{\mathbf{r}}$ and $w_{\mathbf{sr}} = \varepsilon$ then $s \rightarrow^* \xrightarrow{\mathbf{sr}!a}$ and $(q_{\mathbf{r}}, \mathbf{sr}^?a, q''_{\mathbf{r}}) \in \delta_{\mathbf{r}}$.*

Proof. Follows directly from Lemma 30. If $s \multimap (n, \varphi, n')$ and $\mathbf{s} \rightarrow \mathbf{r} : a \in \varphi$, then we have the result since $s \rightarrow^* (\vec{n}; \vec{\varepsilon})$. Otherwise, then the result follows directly by MC, once $(\vec{n}; \vec{\varepsilon})$ has been reached. \blacktriangleleft

► **Lemma 32** (General input availability). *Assume $S = (M_p)_{p \in \mathcal{P}}$ is multiparty compatible, then for all $s \in RS(\text{UT}(S))$, if $s \xrightarrow{\text{sr}!a} s_1 \xrightarrow{\rho} s'$, with $\text{sr}?a \notin \rho$, then $s' \xrightarrow{\rho'} s_2 \xrightarrow{\text{sr}?a}$.*

Proof. Follows directly from Lemmas 29 and 30. ◀

► **Lemma 33.** *Let S be a multiparty compatible system. Then for all $s \in RS(S)$, s is neither a deadlock nor an orphan message configuration.*

Proof. Absence of deadlock configurations. By contradiction, let $s = (\vec{q}; \vec{w}; \nu) \in RS(S)$ be a *deadlock* configuration. Then, by definition of deadlock, we have $\vec{w} = \vec{\varepsilon}$, hence $s \rightarrow (\vec{q}, \varepsilon, \vec{q})$, and by multiparty compatibility there is $\vec{q} \xrightarrow{(s \rightarrow r : a : g_s, \lambda_s; g_r, \lambda_r)} n'$, which leads to a contradiction with the fact that each machine is either in a final or receiving state.

Absence of orphan message configurations. By contradiction, let $s = (\vec{q}; \vec{w}; \nu) \in RS(S)$ be an *orphan message* configuration. Then, by definition of orphan message configurations, all q_p are final and $\vec{w} \neq \vec{\varepsilon}$, which contradicts Lemma 29, since we cannot find a path φ such that $\varepsilon \neq w_{\text{sr}} = \varphi|_{\text{sr}}$. ◀

F.2 Equivalence between systems and projections

In this section, we study an equivalence between S and $(STS(S)|_p)_{p \in \mathcal{P}}$.

We first define a simulation relation between two machines M and M' .

► **Definition 34** (Simulation \lesssim). Let $M_1 = (Q_1, q_{01}, \mathcal{X}_1, \delta_1)$ and $M_2 = (Q_2, q_{02}, \mathcal{X}_2, \delta_2)$. We write $M_1 \lesssim M_2$ iff $q_{01} \lesssim q_{02}$ where, given $q_1 \in Q_1$ and $q_2 \in Q_2$, $q_1 \lesssim q_2$ iff

1. $q_1 \xrightarrow{\ell} q'_1$ implies that there is q'_2 such that $q_2 \xrightarrow{\ell} q'_2$ and $q'_1 \lesssim q'_2$.
2. $q_2 \xrightarrow{\text{sr}!a(g, \lambda)} q'_2$ implies that there is q'_1 such that $q_1 \xrightarrow{\text{sr}!a(g, \lambda)} q'_1$ and $q'_1 \lesssim q'_2$.
3. q_1 is final implies that q_2 is final. ◊

► **Definition 35** (Bisimulation \sim). We write $S_1 \sim S_2$ iff

1. $S_1 \xrightarrow{\ell} S'_1$ implies that there is S'_2 such that $S_2 \xrightarrow{\ell} S'_2$ and $S'_1 \sim S'_2$.
2. $S_2 \xrightarrow{\ell} S'_2$ implies that there is S'_1 such that $S_1 \xrightarrow{\ell} S'_1$ and $S'_1 \sim S'_2$. ◊

► **Lemma 36.** *Let $S = (M_p)_{p \in \mathcal{P}}$ be a multiparty compatible system. Then for all $p \in \mathcal{P}$: $STS(S)|_p \lesssim M_p$.*

Proof. Straightforward from the definition of multiparty compatibility (Def. 4), the definition of projections (Def. 23), and the fact that each machine is deterministic. ◀

► **Theorem 37.** *If S is multiparty compatible, then $\hat{S} = (STS(S)|_p)_{p \in \mathcal{P}}$ is also multiparty compatible.*

Proof. We have to show that multiparty compatibility holds for \hat{S} wrt $STS(\hat{S})$. By Lemma 38, we have that $STS(S) \sim STS(\hat{S})$; hence checking for multiparty compatibility on either transition system is equivalent (if a node n validates the conditions, then any bisimilar node n' does as well).

A state \hat{q} in a machine obtained by minimising the projection of $STS(\hat{S})$, maps to a set of nodes in $STS(\hat{S})$, which in turn map to bisimilar nodes in $STS(S)$. Since S is multiparty compatible, by Lemma 36, there must be a state q in the original machine such that $q \lesssim \hat{q}$ for which the conditions in Def. 4 hold. Since $q \lesssim \hat{q}$ and $STS(S) \sim STS(\hat{S})$, these conditions must also hold for \hat{q} . ◀

Below, we overload \sim for the standard bisimulation between two transition systems.

► **Lemma 38.** *Let S be a multiparty compatible system, and $\hat{S} = (STS(S)|_{\mathcal{P}})_{\mathcal{P} \in \mathcal{P}}$. Then $STS(S) \sim STS(\hat{S})$.*

Proof. Let $STS(S) = (N, n_0, \hookrightarrow, E)$, write $\hat{S} \sim n$ ($n \in N$) if

- $\hat{S} \xrightarrow{\mathbf{sr}!a, \mathbf{sr}?a} \hat{S}'$ implies that $n \xrightarrow{\mathbf{s}\text{-}\mathbf{r}a} n'$ and $\hat{S}' \sim n'$, and
- $n \xrightarrow{\mathbf{s}\text{-}\mathbf{r}a} n'$ implies that $\hat{S} \xrightarrow{\mathbf{sr}!a, \mathbf{sr}?a} \hat{S}'$ and $\hat{S}' \sim n'$.

By definition of $STS(\hat{S})$, we have to show that $\hat{S} \sim n_0$. Note that after projection and minimisation, each state $q_{\mathbf{p}}$ of a machine $M_{\mathbf{p}}$ in \hat{S} corresponds to a set of nodes from $STS(S)$; hence, we write, e.g., $n \in q_{\mathbf{p}}$.

(1) First, we show the following:

$$\forall \mathbf{p} \in \mathcal{P} : n \in q_{\mathbf{p}} \implies (s \xrightarrow{\mathbf{sr}!a, \mathbf{sr}?a} \iff n \xrightarrow{\mathbf{s}\text{-}\mathbf{r}a}) \quad (1)$$

for $(\vec{q}; \vec{\varepsilon}) = s \in RS(\hat{S})$ such that $\vec{q}_0 \hookrightarrow \vec{q}$ and $n \in N$.

(\implies) If $q_{\mathbf{s}} \xrightarrow{\mathbf{sr}!a}$ and $q_{\mathbf{r}} \xrightarrow{\mathbf{sr}?a}$, and $\forall \mathbf{p} \in \mathcal{P} : n \in q_{\mathbf{p}}$, then we must have $n \xrightarrow{\mathbf{s}\text{-}\mathbf{r}a}$ by definition of \hookrightarrow .

(\impliedby) If $n \xrightarrow{\mathbf{s}\text{-}\mathbf{r}a}$ and $\forall \mathbf{p} \in \mathcal{P} : n \in q_{\mathbf{p}}$, then we must also have $q_{\mathbf{s}} \xrightarrow{\mathbf{sr}!a}$ and $q_{\mathbf{r}} \xrightarrow{\mathbf{sr}?a}$, by definition of \hookrightarrow .

(2) We show that $\hat{S} \sim n_0$ by induction on the length of a (synchronous) execution, as per definition of \hookrightarrow , showing that the property in (1) is preserved.

Base case. If the execution has length 0, then we consider s_0 and n_0 . By definition of projection and minimisation, we have $\forall \mathbf{p} \in \mathcal{P} : n_0 \in q_{\mathbf{p}0}$.

Inductive case. Assume the property holds for s (stable) and n (i.e., $\forall \mathbf{p} \in \mathcal{P} : n \in q_{\mathbf{p}}$), $s \xrightarrow{\mathbf{sr}!a, \mathbf{sr}?a} s'$ and $n \xrightarrow{\mathbf{s}\text{-}\mathbf{r}a} n'$. We have to show that $\forall \mathbf{p} \in \mathcal{P} : n' \in q'_{\mathbf{p}}$.

- For \mathbf{s} (resp. \mathbf{r}), it must be the case that for each n' that is reachable by firing $\xrightarrow{\mathbf{sr}!a}$ (resp. $\xrightarrow{\mathbf{sr}?a}$), $n' \in q_{\mathbf{s}}$ (resp. $n' \in q_{\mathbf{r}}$); by definition of minimisation.
- For $\mathbf{p} \in \mathcal{P} \setminus \{\mathbf{s}, \mathbf{r}\}$, by definition of minimisation, each state $q_{\mathbf{p}}$ must contain all the nodes n' that are ε -reachable. In particular, n' is ε -reachable by definition of projection. ◀

► **Lemma 39.** *Let S be a multiparty compatible system. Then $(STS(S)|_{\mathcal{P}})_{\mathcal{P} \in \mathcal{P}} \sim S$.*

Proof. Let $s \in RS(S)$, $\hat{S} = (STS(S)|_{\mathcal{P}})_{\mathcal{P} \in \mathcal{P}}$, $\hat{s} \in RS(\hat{S})$.

By Lemma 36, we know that $\hat{S} \xrightarrow{\ell} \hat{S}'$ implies that $S \xrightarrow{\ell} S'$; and $S \xrightarrow{\mathbf{sr}!a(g, \lambda)} S'$ implies that $\hat{S} \xrightarrow{\mathbf{sr}!a(g, \lambda)} \hat{S}'$. Hence the only difficulty is to show that if $S \xrightarrow{\mathbf{sr}?a(g, \lambda)} S'$ then $\hat{S} \xrightarrow{\mathbf{sr}?a(g, \lambda)} \hat{S}'$.

Abusing notation slightly, we show that $s = (\vec{q}; \vec{w}; \nu) \sim (\vec{q}'; \vec{w}'; \nu) = \hat{s}$, with $\vec{w} = \vec{w}'$, by induction on the length of the execution.

Initial configuration. If $q_{\mathbf{p}} = q_{0\mathbf{p}}$ and $q'_{\mathbf{p}} = q'_{0\mathbf{p}}$, for all $\mathbf{p} \in \mathcal{P}$ and $w_{\mathbf{sr}} = \varepsilon$ for all $\mathbf{sr} \in C$, then we have that $s \xrightarrow{\ell} s_1$ iff $\hat{s} \xrightarrow{\ell} s_2$ since ℓ must be of the form $\mathbf{sr}!a(g, \lambda)$ (due to the queues being empty). We maintain $\vec{w} = \vec{w}'$, trivially, hence we can apply the induction hypothesis on $s_1 \sim s_2$.

Send action. If $s \xrightarrow{\mathbf{sr}!a(g, \lambda)} s_1$ then we must have $\hat{s} \xrightarrow{\mathbf{sr}!a(g, \lambda)} s_2$; and vice-versa if $\hat{s} \xrightarrow{\mathbf{sr}!a(g, \lambda)} s_2$, then $s \xrightarrow{\mathbf{sr}!a(g, \lambda)} s_1$, by Lemma 36. In both cases, we maintain $\vec{w} = \vec{w}'$, trivially, hence we can apply the induction hypothesis on $s_1 \sim s_2$.

Receive action (1). If $\hat{s} \xrightarrow{\mathbf{sr}?a(g, \lambda)} s_2$, then $s \xrightarrow{\mathbf{sr}?a(g, \lambda)} s_1$, by Lemma 36. Hence, we maintain $\vec{w} = \vec{w}'$, trivially, hence we can apply the induction hypothesis on $s_1 \sim s_2$.

Receive action (2). By contradiction, assume $s \xrightarrow{\mathbf{sr}?a(g, \lambda)} s_1$ and $\neg(\hat{s} \xrightarrow{\mathbf{sr}?a(g, \lambda)})$. By Lemma 36, machine \mathbf{r} is not in a terminal state in neither system. By directedness, machine \mathbf{r} in \hat{S} must also expect a message ($\neq a$) from machine \mathbf{s} . However, since a is on top of the queue, machine \mathbf{r} is therefore permanently stuck in state $q_{\mathbf{p}}$. This contradicts the fact that \hat{S}

is multiparty compatible (Theorem 37) and that multiparty compatibility implies that every message is eventually received, cf. Lemma 32. \blacktriangleleft

► **Theorem 25** (Equivalence 2). *If $S = (M_p)_{p \in \mathcal{P}}$ is MC then $S \sim (G(S))_{p \in \mathcal{P}}$.*

Proof. Follows from Lemma 39. \blacktriangleleft

► **Theorem 6** (Safety). *If S is multiparty compatible, then it is safe.*

Proof. Follows from the fact that (i) MC guarantees safety in CFSMs and (ii) timing constraints imply that only a subset of the configurations reachable in the untimed setting are reachable in the timed setting (cf. Prop. F.1). Hence, if there is a deadlock or an orphan message configuration in the timed setting, there is also one in the untimed setting, which contradicts Lemma 33. \blacktriangleleft

In the following, given a system of CTAs S , $STS(S) = (N, n_0, \hookrightarrow, E)$, and $s = (\vec{q}; \vec{w}; \nu) \in RS(S)$, we write $s \multimap (n, \varphi, n')$ when $(\vec{q}; \vec{w}) \multimap (n, \varphi, n')$. This is a sound notation, cf. Lemma 29 and Proposition F.1.

F.3 Progress properties

► **Lemma 40.** *Let S be a MC and IE system and $STS(S) = (N, n_0, \hookrightarrow, E)$. For all $s = (\vec{q}; \vec{w}; \nu) \in RS(S)$ with $s \xrightarrow{\ell}$, and $\text{subj}(\ell) = \mathbf{p}$,*

- (1) $\exists n \in N : s \multimap (n_0, \varphi_0 \cdot e \cdot \varphi_1, n)$ with $n_0 \xrightarrow{\varphi_0} n_1 \xrightarrow{e \cdot \varphi_1} n$, $e|_{\mathbf{p}} = \ell$ and $n_1[\mathbf{p}] = q_{\mathbf{p}}$, and,
- (2) letting, $\rho = \text{nodes}(\varphi_0 \cdot e)$ if $\mathbf{p} = \text{rid}(e)$, $\rho = \text{nodes}(\varphi_0) \cdot e|_{\text{sid}(e)}$ otherwise, and $k = |\rho|$, $\ell = \ell_k = e|_{\mathbf{p}}$, $\vec{v} = \{v_i \mid i \in \text{idx}(\rho)\}$, the following holds:

$$\exists \vec{v} \exists v_k : \text{allpast}(\rho) \wedge \text{elapse}(\rho) \wedge \text{absolute}_{\rho}(\ell_k) \wedge \bigwedge_{v_i \in \vec{v}} v_i \leq v_k \wedge \bigwedge_{x \in \mathcal{X}_{\mathbf{p}}} W_x^k(\rho) = \nu(x)$$

Proof. (1) follows straightforwardly from Lemma 29. We show (2) by induction on the length h of an execution $s \xrightarrow{\alpha_1 \dots \alpha_h} s$. First observe that $s \xrightarrow{\ell}$ implies $\nu \models \text{guard}(\ell)$.

Base Case. If $h = 0$ then we have $s = s_0$, $\ell_k = \ell = \text{pr!}a(g, \lambda)$, and $(q_{0\mathbf{p}}, \ell, q'_{\mathbf{p}}) \in \delta_{\mathbf{p}}$. We can take $s_0 \multimap (n_0, \varphi_0 \cdot e, n)$ (since all queues are empty) and let $n_0 \xrightarrow{\varphi_0} n_1 \xrightarrow{e} n$ and $\rho = \text{nodes}(\varphi_0) \cdot e|_{\mathbf{p}}$, $|\rho| = k$. Since ℓ is the first action fired, it does not depend on any other transitions, hence we have $\text{DG}(\rho) = (\emptyset, \emptyset)$ which implies that

- $\text{idx}(\rho) = \emptyset$
- $\text{allpast}(\rho) \iff \text{true}$
- $\text{elapse}(\rho) \iff \text{true}$
- $\text{absolute}_{\rho}(\ell_k) = \text{guard}(\ell) \{x \mapsto v_k\}_{x \in \mathcal{X}}$
- $\bigwedge_{v_i \in \vec{v}} v_i \leq v_k \iff \text{true}$
- $\bigwedge_{x \in \mathcal{X}_{\mathbf{p}}} W_x^k(\rho) = \nu(x) \iff \bigwedge_{x \in \mathcal{X}_{\mathbf{p}}} v_k = 0$

We have the result by the hypothesis: $\nu_0 = \nu \models \text{guard}(\ell)$.

Inductive Case. Assume the result hold for execution of length less than h and let us show that result holds for h .

(1) Assume that $\ell = \text{pr!}a(g, \lambda)$. Without loss of generality, let us assume that the last action executed by \mathbf{p} is $\ell' = \ell_{k-1} = e'|_{\mathbf{p}}$ $s \multimap (n_0, \varphi_0 \cdot e' \cdot e \cdot \varphi_1, n)$ and let $n_0 \xrightarrow{\varphi_0 \cdot e'} n_1 \xrightarrow{e} n$. If there are other actions between the last actions of \mathbf{p} before ℓ , they are ignored the construction of the graph of dependencies. Pose $\rho' = \text{nodes}(\varphi_0 \cdot e')$, with $|\rho'| = k - 1$ and assume that ν'

is the clock valuation at the time where ℓ' was executed. By induction hypothesis, we have that

$$\exists \vec{v}' \exists v_{k-1} : \text{allpast}(\rho') \wedge \text{elapse}(\rho') \wedge \text{absolute}_{\rho'}(\ell_{k-1}) \wedge \bigwedge_{v_i \in \vec{v}'} v_i \leq v_{k-1} \wedge \bigwedge_{x \in \mathcal{X}_p} W_x^{k-1}(\rho') = \nu'(x)$$

is true, with $\vec{v}' = \{v_i \mid i \in \text{idx}(\rho')\}$.

Now, let $\nu(\hat{x}) - \nu'(\hat{x}) = t$, and $\rho = \text{nodes}(\varphi_0 \cdot e') \cdot e \downarrow_p$, with $|\rho| = k$ and $\vec{v} = \{v_i \mid i \in \text{idx}(\rho)\}$ we have, by definition of $\text{DG}(_)$ and its associated functions:

1. $\text{idx}(\rho) = \text{idx}(\rho') \cup \{k-1\}$
2. $\text{allpast}(\rho) \iff \text{allpast}(\rho') \wedge \text{absolute}_{\rho'}(\ell_{k-1})$
3. $\text{elapse}(\rho) \iff \text{elapse}(\rho') \wedge v_i \leq v_{k-1}$ for any $i \in \text{idx}(\rho')$ such that there is no $(_, \ell_i) \in A$.
4. $\bigwedge_{v_i \in \vec{v}} v_i \leq v_k \iff \bigwedge_{v_i \in \vec{v}'} v_i \leq v_{k-1} \wedge v_{k-1} \leq v_k$

Hence the first three components of the formulae holds by induction hypothesis (remember that v_{k-1} does not appear in $\text{allpast}(\rho')$ nor $\text{elapse}(\rho')$). We have left to show that

$$\text{absolute}_{\rho}(\ell_k) \stackrel{\text{def}}{=} \text{guard}(\ell) \{x \mapsto W_x^k(\rho)\}_{x \in \mathcal{X}}$$

holds assuming that

$$\bigwedge_{x \in \mathcal{X}_p} W_x^k(\rho) = \nu(x)$$

The assumption that $\nu \models \text{guard}(\ell)$ implies that, letting $M(x) = \nu(\hat{x}) - \nu(x)$ (i.e., the difference between the absolute time and the current valuation of clock x):

$$\text{guard}(\ell) \{x \mapsto \nu(\hat{x}) - M(x)\}_{x \in \mathcal{X}} \text{ is true}$$

We have, for each $x \in \mathcal{X}_p$

- $W_x^k(\rho) = v_k - v_{k-1} = t$ if $x \in \text{reset}(\ell_{k-1})$,
- $W_x^k(\rho) = W_x^k(\rho') \{v_{k+t}/v_{k-1}\} = \nu'(x) + t$ otherwise.

thus we have the result since the only further restriction on v_k is $v_{k-1} \leq v_k$.

(2) The case were $\ell = \text{pr}?a(g, \lambda)$ is similar but for the fact that ℓ may have up to *two* dependencies. However, these dependencies must come from different participants and must therefore be independent (in $\text{DG}(_)$), hence we can reason as in the case above. \blacktriangleleft

► **Lemma 41.** *Let S be a MC and IE system; for all $(\vec{q}; \vec{w}; \nu) = s \in \text{RS}(S)$ and for all $\mathbf{p} \in \mathcal{P}$:*

- *if $q_{\mathbf{p}}$ is sending, then $\exists (q_{\mathbf{p}}, \ell, q'_{\mathbf{p}}) \in \delta_{\mathbf{p}}$ such that $\exists t \in \mathbb{R}_{\geq 0} : \nu + t \models \text{guard}(\ell)$,*
- *if $q_{\mathbf{p}}$ is receiving, then $((q_{\mathbf{p}}, \text{sp}?a(g, \lambda), q'_{\mathbf{p}}) \in \delta_{\mathbf{p}} \wedge w_{\text{sp}} \in aA^*) \implies \exists t \in \mathbb{R}_{\geq 0} : \nu + t \models g$.*

Proof. We show this result by induction on the length of an execution from s_0

Base Case. If the execution has length 0, then $s = s_0 = (\vec{q}; \vec{e}; \nu_0)$ and $\vec{q} = n_0$ (by definition of $\text{STS}(S)$). Since all the queues are empty ($\vec{w} = \vec{e}$), we only have to show the first item of the result (the second holds trivially). Hence, we must show that for all $\mathbf{p} \in \mathcal{P}$ if $q_{0\mathbf{p}}$ is sending, then $\exists (q_{0\mathbf{p}}, \ell, q'_{\mathbf{p}}) \in \delta_{\mathbf{p}}$ such that $\exists t \in \mathbb{R}_{\geq 0} : \nu + t \models \text{guard}(\ell)$. We show this by contradiction: assume there is $\mathbf{p} \in \mathcal{P}$ such that for all $(q_{0\mathbf{p}}, \ell, q'_{\mathbf{p}}) \in \delta_{\mathbf{p}}$ and $\nu_0 + t \not\models \text{guard}(\ell)$, for all $t \in \mathbb{R}_{\geq 0}$.

By definition of IE, there must be $e \in E$ and φ with $n_0 \xrightarrow{e:\varphi}$ such that $(n_0, e \cdot \varphi)$ is RPE for $\text{ids}(n_0)$ Note that $\text{ids}(n_0) = \mathcal{P}$, by MC. Hence, there should be a path $e \cdot \varphi = e_1 \cdots e_i \cdots e_k$ such that $\text{sid}(e_i) = \mathbf{p}$ and, assuming n is the node from which e_i is fired, (n, e_i) is PE for \mathbf{p} .

Since $e_i \downarrow_{\mathbf{p}}$ is a sending action and it is the first action for \mathbf{p} , we have

$$\text{DG}(\text{nodes}(e_1 \cdots e_{i-1}) \cdot e_i \downarrow_{\mathbf{p}}) = (\emptyset, \emptyset)$$

Thus checking PE for this pair amounts to checking whether the following holds:

$$\exists v : \text{true} \implies \text{guard}(e_i \downarrow_p) \{x \mapsto v\}_{x \in \mathcal{X}}$$

which contradicts our hypothesis that no action of \mathbf{p} was eventually fireable.

Inductive Case. Let us assume that result holds for $s = (\vec{q}; \vec{w}; \nu) \in RS(S)$ and let us show that it holds for $s' = (\vec{q}'; \vec{w}'; \nu')$ with $s \xrightarrow{\alpha} s'$. There are five cases, depending on the type of action that is α and, if α is a send or receive action, on the type of state reached by $\text{subj}(\alpha)$.

1. $\alpha = \text{sp}^?a(g, \lambda)$ and q'_p is a sending state. Then we have $\nu' = \nu$, $q_q = q'_q$ for all $q \in \mathcal{P} \setminus \{\mathbf{p}\}$. We have to show that

$$\exists (q'_p, \ell', q''_p) \in \delta_p \text{ such that } \exists t \in \mathbb{R}_{\geq 0} : \nu + t \models \text{guard}(\ell')$$

Take $s \multimap (n, \varphi, _)$ such that, assuming that $e' \downarrow_p = \ell'$,

$$\varphi = \varphi_1 \cdot (\mathbf{s} \rightarrow \mathbf{p} : a; g_s, \lambda_s; g_p, \lambda_p) = e \cdot \varphi_2 \cdot (\mathbf{p} \rightarrow _ : a'; g'_s, _ ; _) = e' \cdot \varphi_3$$

and

$$n \xrightarrow{\varphi_1} n_1 \xrightarrow{(\mathbf{s} \rightarrow \mathbf{p} : a; g_s, \lambda_s; g_r, \lambda_r)} n_2 \xrightarrow{\varphi_2} n_3 \xrightarrow{(\mathbf{p} \rightarrow _ : a'; g'_s, _ ; _)} n_4 \xrightarrow{\varphi_3} \quad n_1[\mathbf{p}] = q_p \text{ and } n_3[\mathbf{p}] = q'_p$$

this is possible since (i) $w_{\text{sp}} \in a\mathbb{A}^*$ and (ii) q'_p is a sending, thus there must be continuation of that path so that \mathbf{p} does fire a sending action.

By Lemma 40, we know that there is path φ' from n_0 to n_1 for which the pair (n_1, e) is PE for \mathbf{p} . In addition, by IE, there must be a path φ'' from n_2 that must be RPE for $\text{ids}(n_2)$ (note that $\mathbf{p} \in \text{ids}(n_2)$). Assume that φ above is a suffix of this path. For φ'' to be RPE, one must check, notably, that the pair (n_3, e') is PE for \mathbf{p} . Since, by assumption, e' is the first event involving \mathbf{p} since e and \mathbf{p} is the sender; then the dependency graph is for e' is the same as the one for e with the addition of the guard corresponding to $e \downarrow_p$, and this guard is satisfiable since the corresponding action has just been executed and $\nu' = \nu$. Hence, the LHS of the implication is satisfiable in the check for PE of e' , thus there must also be a (future) time where $e' \downarrow_p$ is fireable; and we have the result.

2. $\alpha = \text{sp}^?a(g, \lambda)$ and q'_p is a receiving state. We have $\nu' = \nu$, $q_q = q'_q$ for all $q \in \mathcal{P} \setminus \{\mathbf{p}\}$. We have to show that

$$\text{if } (q'_p, \text{qp}^?b(g', \lambda'), q''_p) \in \delta_p \text{ and } w_{\text{qp}} \in b\mathbb{A}^* \text{ then } \exists t \in \mathbb{R}_{\geq 0} : \nu + t \models g'$$

Take $s \multimap (n, \varphi, _)$ such that

$$\varphi = \varphi_1 \cdot (\mathbf{s} \rightarrow \mathbf{p} : a; g_s, \lambda_s; g_r, \lambda_r) = e \cdot \varphi_2 \cdot (\mathbf{q} \rightarrow \mathbf{p} : b; _ , _ ; g'_s) = e' \cdot \varphi_3$$

and

$$n \xrightarrow{\varphi_1} n_1 \xrightarrow{(\mathbf{s} \rightarrow \mathbf{p} : a; g_s, \lambda_s; g_r, \lambda_r)} n_2 \xrightarrow{\varphi_2} n_3 \xrightarrow{(\mathbf{q} \rightarrow \mathbf{p} : b; _ , _ ; g'_s)} n_4 \xrightarrow{\varphi_3} \quad n_1[\mathbf{p}] = q_p \text{ and } n_3[\mathbf{p}] = q'_p$$

We reason similarly as above (cf. (1)), however in this case, the dependency graph for the check of the pair $(n_3, e' \downarrow_p)$ must also include the action $e' \downarrow_q$. However, since $e' \downarrow_q$ was indeed executed (there is a corresponding message in the queue) and there is no dependence between $e' \downarrow_q$ and $e \downarrow_p$ (if $\mathbf{q} = \mathbf{s}$, then it contradicts the fact that the first message in the queue is a); the LHS of the implication is again satisfiable for the current clock valuation (Lemma 40). Therefore, there must be a future time when g' is satisfied.

3. $\alpha = \text{pr!}a(g, \lambda)$ and q'_p is a sending state. In this case, we have to show that
 - a. $\exists(q'_p, \ell', q''_p) \in \delta_p$ such that $\exists t \in \mathbb{R}_{\geq 0} : \nu + t \models \text{guard}(\ell')$, and
 - b. if $(q'_r, \text{pr?}a(g', \lambda'), q''_r) \in \delta_r$ and $w_{\text{pr}} \in a\mathbb{A}^*$ then $\exists t \in \mathbb{R}_{\geq 0} : \nu + t \models g'$
 We can show (3a) following the same reasoning as (1), and show (3b) following the same reason as (2).
4. $\alpha = \text{pr!}a(g, \lambda)$ and q'_p is a receiving state. We have to show that
 - a. if $(q'_p, \text{qp?}b(g', \lambda'), q''_p) \in \delta_p$ and $w_{\text{qp}} \in b\mathbb{A}^*$ then $\exists t \in \mathbb{R}_{\geq 0} : \nu + t \models g'$, and
 - b. if $(q'_r, \text{pr?}a(g'', \lambda''), q''_r) \in \delta_r$ and $w_{\text{pr}} \in a\mathbb{A}^*$ then $\exists t \in \mathbb{R}_{\geq 0} : \nu + t \models g''$
 for which we use the same reasoning as the case above.
5. $\alpha = t$. Follows immediately from the semantics of CTAs (cf. Def. 9). ◀

► **Lemma 42.** *Let S be a MC and IE system. For all $s = (\vec{q}; \vec{w}; \nu) \in RS(S)$ and $\mathbf{r} \in \mathcal{P}$, if $(q_r, \text{sr?}b(g', \lambda'), q'_r) \in \delta_r$ and $w_{\text{sr}} = \varepsilon$ then $s \rightarrow^* \xrightarrow{\text{sr!}a}$ and $(q_r, \text{sr?}a(g', \lambda'), q''_r) \in \delta_r$.*

Proof. The result follows from its untimed counterpart (Lemma 31) and the fact that IE requires each node to allow an RPE path which involves all the active participant, include \mathbf{r} . By MC, if \mathbf{r} is active at from a node n , then \mathbf{r} is involved in all branching outgoing n . ◀

► **Theorem 13 (Progress).** *Suppose S is multiparty compatible (Def. 4) and interaction enabling (Def. 12). (1) Then S satisfies the progress property. (2) For all $s = (\vec{q}; \vec{w}; \nu) \in RS(S)$, if there is $\mathbf{p} \in \mathcal{P}$ such that q_p is not final, then there is s' such that $s \rightarrow s'$.*

Proof. We show (1), (2) is a direct corollary of (1).

- Absence of *deadlock* and *orphan message* follows directly from Theorem 6.
- Absence of *unfeasible configuration* follows from Lemma 41.
- Absence of *unsuccessful reception configuration* follows from Lemmas 33, 41 and 42. In particular, Lemma 33 guarantees that only *expected* messages can arrive in a queue; while Lemma 41 guarantees there is always a future time where this message can be received. Finally, concerning the special case of unsuccessful reception configuration where $s = (\vec{q}; \vec{w}; \nu)$ and there exists $\mathbf{r} \in \mathcal{P}$ such that q_r is a receiving state, $(q_r, \text{sr?}a(g, \lambda), q'_r) \in \delta_r$, $w_{\text{sr}} = \varepsilon$, and $\forall t \in \mathbb{R}_{\geq 0} : \nu + t \not\models g$ Lemma 42 guarantees that if a participant is in a receiving state, there is a future time where it will receive a message. Hence Lemma 41 eventually applies and this type of configuration is also ruled out. ◀

Given two walks ω_1 and ω_2 in $STS(S)$ such that the last node of ω_1 is the same as the first node of ω_2 , we write $\omega_1 \cdot \omega_2$ for the walk corresponding to the concatenation of ω_1 and ω_2 .

► **Definition 43 (Non-redundant walk).** A walk ω in $STS(S)$ is *non-redundant* if $\omega \neq \omega_1 \cdot \omega_2 \cdot \omega_3$, with $\omega_2 \neq \varepsilon$. ◊

► **Lemma 44.** *Given $STS(S) = (N, n_0, \leftrightarrow, E)$, there is a finite number of non-redundant walks in $STS(S)$.*

Proof. Straightforward since there is a finite number of elementary cycle starting from a given node (also finite) and a finite number of permutation of these cycles. ◀

► **Lemma 45.** *Let $\varphi = e_1 \cdots e_k$ ($0 \leq k$) and $\rho = \text{nodes}(\varphi)$, $A_1 = \text{dep}(\rho; \ell)$ and $A_2 = \text{dep}(\rho \cdot \rho; \ell)$. $\forall (\ell_i, \ell_j) \in A_1 \implies \exists (\ell_i, \ell_j) \in A_2$.*

Proof. Straightforward since the dependencies are *at least* duplicated, e.g., $e_i \downarrow_{\mathbf{p}}$ is a (possibly indirect) dependency of $e_{i+k} \downarrow_{\mathbf{p}}$ (where $k = |\rho|$) since $\text{subj}(e_i \downarrow_{\mathbf{p}}) = \text{subj}(e_{i+k} \downarrow_{\mathbf{p}})$. ◀

► **Theorem 14** (Decidability). *Interaction enabling (Def. 12) is decidable.*

Proof. We first note that $STS(S)$ is a finite transition system and that our logic is a subset of the (decidable) Presburger arithmetic (once each constant c has been multiplied by a common multiple of all the constants in the guards, as in [2]). Next, we show that it is enough to consider *finite* path in $STS(S)$ by showing that if a par is PE for a non-redundant path $\omega_1 \cdot \omega_2 \cdot \omega_3$ path, then it is also PE for $\omega_1 \cdot \omega_2 \cdot \omega_2 \cdot \omega_3$.

Assume that a pair (n, e) is PE for $\mathbf{p} \in \text{id}(e)$ on a *given path* $\varphi = \varphi_1 \cdot \varphi_2 \cdot \varphi_3$. Take $\mathbf{p} = \text{sid}(e)$ (for simplicity) and

$$n_0 \xrightarrow{\varphi_1} n_1 \xrightarrow{\varphi_2} n_1 \xrightarrow{\varphi_3} n \quad \text{and} \quad \rho = \text{nodes}(\varphi \cdot e) \quad \text{and} \quad |\rho| = k$$

If the pair is PE then the following holds, with $\ell_k = e \downarrow_{\mathbf{p}}$, $\vec{v} = \{v_i \mid i \in \text{idx}(\rho)\}$

$$\forall \vec{v} \exists v_k : \text{allpast}(\rho) \wedge \text{elapse}(\rho) \implies \text{absolute}_{\rho}(\ell_k) \wedge \bigwedge_{v_i \in \vec{v}} v_i \leq v_k$$

Assume $|\text{nodes}(\varphi_1)| = h_1$ $|\text{nodes}(\varphi_2)| = h_2$, let us write σ for the variable substitution that maps

- each v_i s.t. $1 \leq i \leq h_1 + h_2$ to itself
- each v_i $h_1 + h_2 \leq i \leq k$ to v_{i+h_2} .

Now consider the path $\varphi' = \varphi_1 \cdot \varphi_2 \cdot \varphi_2 \cdot \varphi_3$ and $\rho' = \text{nodes}(\varphi' \cdot e)$. Take $\vec{v} = \{v_i \mid i \in \text{idx}(\rho')\}$. By Lemma 45 and logical weakening, we have

$$\text{allpast}(\rho') \implies \text{allpast}(\rho)\sigma \quad \text{and} \quad \text{elapse}(\rho') \implies \text{elapse}(\rho)\sigma$$

In addition, we have $\text{absolute}_{\rho}(\ell_k)\sigma \iff \text{absolute}_{\rho}(\ell_{k+h_2})$.

Finally,

$$\left(\text{elapse}(\rho') \wedge \bigwedge_{v_i \in \vec{v}} v_i \leq v_k \right) \sigma \implies \bigwedge_{v_i \in \vec{v}} v_i \leq v_{k+h_2}$$

since each variable that is added by the duplication is either (i) related to its predecessor (which was already present) or (ii) is added by a dependency on the duplication of the cycle, which means that it is dependent on a variable in the cycle, with was also required to be less or equal than than v_{k+h_2} . ◀

F.4 Non-zeno properties

In this appendix we give the proof of the main theorems of § 5: Theorem 17, Theorem 18, and Theorem 19, after a few auxiliary lemmas. Lemma 46 states that if a system S has an infinite execution then its STS has an elementary cycle. Lemma 46 also sets a correspondence between states of the infinite execution of S and nodes of the cycle in $STS(S)$. Reversely, Lemma 47 states that if $STS(S)$ has no cycles then S will stop executing non-time actions at some point (either because the system is stuck or because each machine is in a final state). By Lemma 48 these ('final' or stuck) configurations are non-zeno.

► **Lemma 46.** *Let S be a multiparty compatible system and $s \in RS(S)$. If for all s' such that $s \rightarrow^* s'$ there exists s'' such that $s' \xrightarrow{\ell} s''$ with ℓ not being a time action, then there exist $s_i \in RS(S)$, $n_i, n'_i \in STS(S)$, and φ_i such that $s \rightarrow^* s_i$, $s_i \multimap (n_i, \varphi, n'_i)$ and n_i is in an elementary cycle in $STS(S)$.*

Proof. By hypothesis, the set $RS(s) = \{s_i \mid s \rightarrow^* s_i\}$ is infinite. By multiparty compatibility of S we can build $STS(S)$. By Lemma 29, for each $s_i \in RS(s)$ there exist $n_i, n'_i \in STS(S)$ and φ_i such that $s_i \multimap (n_i, \varphi_i, n'_i)$. Since $RS(s)$ is infinite, but the set of nodes of $STS(S)$ is finite, then there exist $s_i, s_j \in RS(s)$ such that $s_i \multimap (n_i, \varphi_i, n'_i)$, $s_j \multimap (n_j, \varphi_j, n'_j)$ and $n_i = n_j$. By $n_i = n_j$, the walk from n_i to n_j is a cycle. The lemma follows by the fact that any cycle includes an elementary cycle. ◀

► **Lemma 47.** *Let S be a multiparty compatible system. If $STS(S)$ has no cycles then for all $s \in RS(S)$ there exists s' such that $s \rightarrow^* s'$ and s' cannot make any action which is not a time action.*

Proof. Follows similar arguments than the proof of Lemma 46, by observing the correspondence between the number of possible walks in $STS(S)$, which in this case is finite, and the corresponding configurations of S . ◀

► **Lemma 48.** *Let S be a multiparty compatible system. If $s \in RS(S)$ and for all s' such that $s \rightarrow^* s'$ there exists s'' such that $s' \xrightarrow{\ell} s''$ where ℓ can only be a time action then s is not a zeno configuration.*

Proof. There are two cases in which s' can only make time actions: (1) s' is a final state, and (2) s' is a configuration violating progress. In case (1) s' is non-zeno as a final state always allows time to diverge. Case (2) cannot occur as this would imply that s' could reach a state with no further possible transitions, which contradicts the hypothesis. ◀

► **Theorem 17 (Non-zenoness).** *If S is MC and CE, then S is non-zeno.*

Proof. We proceed by contradiction, assuming that both of the following conditions hold:

1. S is MC and CE,
2. there exists $s \in RS(S)$ which is a zeno configuration.

Since s is zeno then for all s' such that $s \rightarrow^* s'$ there exists s'' such that $s' \xrightarrow{\ell} s''$. Fix one s' . By Lemma 48 we can fix an action ℓ that is not a time action. Since S is MC then we can apply Lemma 46, yielding that there exist n, n', φ such that $s' \multimap (n, \varphi, n')$ and n is in an elementary cycle ω . Let $n \xrightarrow{e} n''$ for some $e = (\mathbf{s} \rightarrow \mathbf{r} : a; g_{\mathbf{s}}, \lambda_{\mathbf{s}}; g_{\mathbf{r}}, \lambda_{\mathbf{r}})$. Since S is CE, then by Def. 16 we can create a partition of $\text{fc}(g_{\mathbf{s}})$, which consists of the three sets of the following form:

- $C_1 = \{x \mid x \in \text{fc}(g_{\mathbf{s}}) \wedge g_{\mathbf{s}} \text{ is not } UB \text{ for } x\}$
- $C_2 = \{x \mid x \in \text{fc}(g_{\mathbf{s}}) \wedge g_{\mathbf{s}} \text{ is } UB \text{ for } x \wedge$
 $\exists (\mathbf{p} \rightarrow \mathbf{q} : b; g_{\mathbf{p}}, \lambda_{\mathbf{p}}; g_{\mathbf{q}}, \lambda_{\mathbf{q}}) \text{ in } \omega \text{ s.t. } x \in \{\lambda_{\mathbf{p}} \cup \lambda_{\mathbf{q}}\} \wedge$
 $\exists (\mathbf{p}' \rightarrow \mathbf{q}' : c; g'_{\mathbf{p}}, \lambda'_{\mathbf{p}}; g'_{\mathbf{q}}, \lambda'_{\mathbf{q}}) \text{ in } \omega \text{ s.t. } g'_{\mathbf{p}} \text{ is } SP \}$
- $C_3 = \{x \mid x \in \text{fc}(g_{\mathbf{s}}) \wedge g_{\mathbf{s}} \text{ is } UB \text{ for } x \wedge \text{ there is } n''' \neq n \text{ and } e_1 \neq e$
 $\text{s.t. } \text{id}(e) \neq \text{id}(e_1), n \xrightarrow{e_1} n''' \text{ and } (n, e_1) \text{ is } \forall\text{-PE for } \text{sid}(e_1)\}$

Note that more than one partition of $\text{fc}(g_{\mathbf{s}})$ may exist (since e.g., a clock can be both reset and not have an upper bound). However, as all partition are of the form above, we consider one without loss of generality.

We next proceed by induction on the number of elementary cycles in $STS(S)$.

(Base case - ω includes one elementary cycle) We consider two possible, and exhaustive scenarios: $C_3 \neq \emptyset$ and $C_3 = \emptyset$.

If $C_3 \neq \emptyset$ then the (one) cycle ω has an ‘escape’ to n''' . Since by hypothesis (base case of induction) there is just one elementary cycle in ω this implies that n''' is not in another cycle in $STS(S)$. Since n''' is not in a cycle, then by Lemma 47 it follows that there is no infinite execution from s' . Therefore s' is non zero, which contradicts condition (2).

If $C_3 = \emptyset$ then we proceed by induction on the number of nodes in ω that have at least one clock with upper bound in the sender’s constraint (i.e., nodes in ω that have the second set in the partition, of the form of C_2 , non empty).

- In the base case there is only one node that has upper bounded clocks. For simplicity of notation and without loss of generality, we assume the only node with clocks in C_2 is n (and recall that n has an outgoing event $e = (\mathbf{s} \rightarrow \mathbf{r} : a; g_s, \lambda_s; g_r, \lambda_r)$). Since C_2 is a finite set, there exists a constant c which is the smallest upper bound for the clocks in the guards g_s of e . Also, note that by condition (1-ii) of Def. 16 c is strictly greater than 0 (i.e., $c \in \mathbb{Q}_{>0}$). We know, by definition of C_2 , that each $x \in C_2$ is reset at least once in ω . Since all other nodes in ω do not have upper bounds (hypothesis of this base case) then at every execution of the cycle, when reaching s' , it is possible to let time elapse at least c time units. Note that this delay is allowed by: (i) the constraints on clocks in C_1 (as they have no upper bound), (ii) the constraints in C_2 as they have upper bound greater than or equal to c , and (iii) by sender’s constraints the other nodes in ω which have no upper bound (by hypothesis of this base case). Since $c > 0$ time units can elapse at each cycle, and the execution is infinite by hypothesis, then after $\frac{t - \nu_{s'}(\hat{x})}{c}$ cycles $\nu_{s'}(\hat{x}) \geq t$ which contradicts the condition (2) by which s' is zero.
- In the inductive case, we consider a cycle ω' which is as ω except that the upper-binding constraints of n have been removed from n in ω' . We assume by induction that s' in ω' is non zero and there are m nodes in ω' with upper binding (sender’s) constraints. Since s' is non zero in ω' then either: the time elapsed in each iteration in ω' is infinite or has an upper bound $c' > 0$ (i.e., some time elapsed in each iteration, and it is not progressively decreasing). In either case, at each iteration in ω' at least c' time units can elapse. Now, considering the upper binding constraints of n in ω , which again we assume to have an upper bound greater than or equal to c , we have that at least $\min(c, c')$ time units can elapse at each iteration in ω . Since both c and c' are strictly greater than 0 then $\min(c, c') > 0$ hence after $\frac{t - \nu_{s'}(\hat{x})}{\min(c, c')}$ cycles $\nu_{s'}(\hat{x}) \geq t$ which, again, contradicts the condition (2).

(Inductive case - ω includes more than one elementary cycles) If $C_3 = \emptyset$ then we can proceed exactly as for the base case (i.e., the case for only one cycle).

If $C_3 \neq \emptyset$, we assume by induction that (being S MC and CE) $STS(S)$ has m loops which do not have any zero configuration. We extend $STS(S)$ with one more loop. The topology of this extension is irrelevant to this proof. By the case analysis we know that $C_3 \neq \emptyset$ hence C_3 includes at least one clock that has an upper bound in s' . Recall that n is a node such that $s' \multimap (n, \varphi, n')$: since $STS(S)$ is CE there exists an ‘escape’ from n yielding out of the elementary cycle ω . Namely, $n \xrightarrow{e_1} n'''$ for some e_1 and n''' .

If n''' is not included in any walk into or yielding to an elementary cycle then Lemma 47 yields that there is no infinite execution from s' . Therefore s' is non zero, which contradicts condition (2).

If n''' is included in an elementary cycle ω'' then we have $\omega'' \neq \omega'$ since $n \neq n'''$ (although ω'' and ω' could have some nodes in common). The condition in C_3 , which corresponds to condition 1 in Def. 16, and requiring that (n, e_1) is \forall -PE for $\text{sid}(e_1)$, yields that it is always possible for S to make a step $s \xrightarrow{e_1} s''$ such that $s' \multimap (n''', \varphi', n''')$ (for some $\varphi', n''')$. By induction ω'' does not have zero configurations, hence it is possible to follow ω'' an arbitrary

number of times until $\nu_s(\hat{x}) \geq t$ which, again, contradicts condition (2). ◀

► **Theorem 18** (Decidability). *Cycle enabling (Def. 16) is decidable.*

Proof. Observe that the checks CE performs in each state are either (a) decidable syntactic checks on the form of the guards, or (b) checks for PE, which are decidable by Lemma 14. Observe also that the checks are made on the states all elementary cycles in $STS(S)$, which are a finite number. Decidability follows from the fact that CE performs decidable checks on a finite number of states. ◀

► **Theorem 19** (Eventual reception). *If S is MC, IE, and CE, then S satisfies ER.*

Proof. Take a configuration $s = (\vec{q}; \vec{w}; \nu) \in RS(S)$ such that $w_{sr} \in a\mathbb{A}^*$. We must have $s \multimap (n, \varphi_1 \cdot (\mathbf{s} \rightarrow \mathbf{r} : a; g_s, \lambda_s; g_r, \lambda_r) \cdot \varphi_2, n)$. Assume we have the following situation:

$$n \xrightarrow{\varphi_1} n_1 \xrightarrow{(\mathbf{s} \rightarrow \mathbf{r} : a; g_s, \lambda_s; g_r, \lambda_r)} n_2 \xrightarrow{\varphi_2} n' \quad \text{with } n_1[\mathbf{r}] = q'_r \neq q_r \text{ and } n_2[\mathbf{s}] = q_s$$

i.e., \mathbf{s} has just put the message a into queue w_{sr} and \mathbf{r} is not (yet) able to receive it. We first show that there is an execution for \mathbf{r} to reach a (local) state where it can receive a . This follows simply for IE and RPE since there must be a satisfiable path involves all participant (and by MC if a is sent in several branches, it is always received).

Once we are in a configuration $s = (\vec{q}; \vec{w}; \nu)$ such that $w_{sr} \in a\mathbb{A}^*$ and $(q_r, \mathbf{sr}?a(g, \lambda), q''_r) \in \delta_r$, then we know by Theorem 13 that there is a future time where g is satisfiable. In addition, this time is indeed reachable since s is not a zeno configuration (Theorem 19). ◀