

Kent Academic Repository

Full text document (pdf)

Citation for published version

Murgia, Maurizio (2018) On Urgency in Asynchronous Timed Session Types. In: Electronic Proceedings in Theoretical Computer Science. Proceedings 11th Interaction and Concurrency Experience. Open Publishing Association pp. 85-94.

DOI

<https://doi.org/10.4204/EPTCS.279>

Link to record in KAR

<https://kar.kent.ac.uk/69683/>

Document Version

Author's Accepted Manuscript

Copyright & reuse

Content in the Kent Academic Repository is made available for research purposes. Unless otherwise stated all content is protected by copyright and in the absence of an open licence (eg Creative Commons), permissions for further reuse of content should be sought from the publisher, author or other copyright holder.

Versions of research

The version in the Kent Academic Repository may differ from the final published version.

Users are advised to check <http://kar.kent.ac.uk> for the status of the paper. **Users should always cite the published version of record.**

Enquiries

For any further enquiries regarding the licence status of this document, please contact:

researchsupport@kent.ac.uk

If you believe this document infringes copyright then please contact the KAR admin team with the take-down information provided at <http://kar.kent.ac.uk/contact.html>

On Urgency in Asynchronous Timed Session Types*

Maurizio Murgia

School of Computing
University of Kent.
Canterbury, Uk

M.Murgia@kent.ac.uk

We study an urgent semantics of asynchronous timed session types, where input actions happen as soon as possible. We show that with this semantics we can recover to the timed setting an appealing property of untimed session types: namely, deadlock-freedom is preserved when passing from synchronous to asynchronous communication.

1 Introduction

Session types are abstractions of communication protocols [14], used to statically or dynamically check that distributed programs interact correctly. The original binary synchronous theory has subsequently been extended in several directions: explicit support for multiparty protocols and choreographies [11], asynchronous communication through FIFO buffers [11], time [8, 4], and others [10].

In this paper, we start an investigation on the relationships between synchronous and asynchronous session types in the timed binary setting. A related study has been performed in the untimed setting [6], where, among other things, it has been proved that deadlock freedom in session types interacting with synchronous communication is preserved if messages are buffered. As reasoning about synchronous systems is easier (synchronous progress is decidable, asynchronous one is not), concurrent applications can be designed and verified with synchronous communication in mind, and then run on top of real-world asynchronous mediums (e.g., TCP) while preserving correctness. We refer to this practice as *design synchronous/deploy asynchronous* methodology.

In the timed setting, as noted in [4], this property is lost, at least with the asynchronous semantics of [8]. In this work, we propose an alternative semantics of asynchronous session types, that forbids delays when reading actions are possible, similar to an urgent semantics of Communicating Timed Automata [3]. As noted in [3], this semantics, that we call *input urgent* asynchronous semantics, better captures common reading primitives of programming languages/APIs, that return as soon as a message is available. Our semantics makes therefore session types abstract models of *programs*. The semantics of [8], instead, being more general (allows more behaviour), seems preferable for modelling *protocols* at a higher level of abstraction.

The main contribution of this paper is that the preservation result of [6] can be lifted to the timed setting, when using input urgent semantics. As timed synchronous progress is decidable [4], and, as discussed above, input urgent semantics models programs, our result paves the way for the application of the design synchronous/deploy asynchronous methodology to time-sensitive distributed software.

*This work has been partially sponsored by EPSRC EP/N035372/1.

2 Synchronous timed session types

We now introduce timed session types (TST), their synchronous semantics and the associated notion of progress. The material of this section is taken from [4], with minor variations. The clock based model of time is borrowed from Timed Automata [1].

Preliminaries. Let \mathbf{A} be a set of *actions*, ranged over by $\mathbf{a}, \mathbf{b}, \dots$. We denote with $\mathbf{A}^!$ the set $\{!a \mid a \in \mathbf{A}\}$ of *output actions*, with $\mathbf{A}^?$ the set $\{?a \mid a \in \mathbf{A}\}$ of *input actions*, and with $\mathbf{L} = \mathbf{A}^! \cup \mathbf{A}^?$ the set of *branch labels*, ranged over by ℓ, ℓ', \dots . We use δ, δ', \dots to range over the set $\mathbb{R}_{\geq 0}$ of not-negative real numbers, and d, d', \dots to range over the set of natural numbers \mathbb{N} . Let \mathbf{C} be a set of *clocks*, variables in $\mathbb{R}_{\geq 0}$, ranged over by t, t', \dots . We use $R, T, \dots \subseteq \mathbf{C}$ to range over sets of clocks. The syntax of guards (ranged over by g, g', \dots) is:

$$g ::= \text{true} \mid \neg g \mid g \wedge g \mid t \circ d \mid t - t' \circ d. \quad \text{where } \circ \in \{<, \leq, =, \geq, >\}$$

We give meaning to guards in terms of *clock valuations*, namely functions of type $\mathbf{C} \rightarrow \mathbb{R}_{\geq 0}$ which associate each clock with its value. We denote with $\mathbb{V} = \mathbf{C} \rightarrow \mathbb{R}_{\geq 0}$ the set of clock valuations (ranged over by \mathbf{v}, η, \dots), and with \mathbf{v}_0 the valuation mapping each clock to zero. We use $\mathcal{H}, \mathcal{H}', \dots$ to range over sets of clock valuations. We write $\mathbf{v} + \delta$ for the valuation which increases \mathbf{v} by δ , i.e., $(\mathbf{v} + \delta)(t) = \mathbf{v}(t) + \delta$. For a set $R \subseteq \mathbf{C}$, we write $\mathbf{v}[R]$ for the *reset* of the clocks in R , i.e.,

$$\mathbf{v}[R](t) = \begin{cases} 0 & \text{if } t \in R \\ \mathbf{v}(t) & \text{otherwise} \end{cases}$$

Definition 1 (Semantics of guards). Let g be a guard. We define the set of clock valuations $\llbracket g \rrbracket$ inductively as follows, where $\circ \in \{<, \leq, =, \geq, >\}$:

$$\begin{aligned} \llbracket \text{true} \rrbracket &= \mathbb{V} & \llbracket \neg g \rrbracket &= \mathbb{V} \setminus \llbracket g \rrbracket & \llbracket g_1 \wedge g_2 \rrbracket &= \llbracket g_1 \rrbracket \cap \llbracket g_2 \rrbracket \\ \llbracket t \circ d \rrbracket &= \{\mathbf{v} \mid \mathbf{v}(t) \circ d\} & \llbracket t - t' \circ d \rrbracket &= \{\mathbf{v} \mid \mathbf{v}(t) - \mathbf{v}(t') \circ d\} \end{aligned}$$

Definition 2 (Past). Let \mathcal{H} be a set of clock valuations. We define $\downarrow \mathcal{H}$ (the past of \mathcal{H}) as follows:

$$\downarrow \mathcal{H} = \{\mathbf{v} \mid \exists \delta \geq 0 : \mathbf{v} + \delta \in \mathcal{H}\}$$

To model messages in transit we use queues. Queues are terms of the following grammar:

$$\rho ::= \emptyset \mid \mathbf{a}; \rho$$

We use ρ, σ to range over queues and we omit trailing occurrences of \emptyset . We write $|\rho|$ for the number of messages in the queue ρ (we omit the straightforward definition).

Syntax. A TST p models the behaviour of a single participant involved in an interaction. Roughly, in an internal choice $\bigoplus_i !\mathbf{a}_i \{g_i, R_i\} . p_i$ a participant has to perform one of the outputs $!\mathbf{a}_i$ in a time window where g_i is true. Conversely, in an external choice $\sum_i ?\mathbf{a}_i \{g_i, R_i\} . q_i$ the participant is available to receive each message \mathbf{a}_i in any instant within the time window defined by g_i .

Definition 3. Timed session types p, q, \dots are terms of the following grammar:

$$p ::= \mathbf{1} \mid \bigoplus_{i \in I} !\mathbf{a}_i \{g_i, R_i\} . p_i \mid \sum_{i \in I} ?\mathbf{a}_i \{g_i, R_i\} . p_i \mid \text{rec } X . p \mid X$$

where (i) $I \neq \emptyset$ and finite, (ii) actions in internal/external choices are pairwise distinct, (iii) recursion is guarded. We omit true guards, empty resets, and trailing occurrences of $\mathbf{1}$.

$$\begin{array}{c}
(\text{!a}\{g, R\}.p \oplus p', \emptyset, \mathbf{v}) \xrightarrow{\tau}_s (p, \mathbf{a}, \mathbf{v}[R]) \quad \text{if } \mathbf{v} \in \llbracket g \rrbracket \quad [\oplus] \\
(p, \mathbf{a}, \mathbf{v}) \xrightarrow{\text{!a}}_s (p, \emptyset, \mathbf{v}) \quad [\text{!}] \\
(\text{?a}\{g, R\}.p + p', \emptyset, \mathbf{v}) \xrightarrow{\text{?a}}_s (p, \emptyset, \mathbf{v}[R]) \quad \text{if } \mathbf{v} \in \llbracket g \rrbracket \quad [?] \\
(p, \emptyset, \mathbf{v}) \xrightarrow{\delta}_s (p, \emptyset, \mathbf{v} + \delta) \quad \text{if } \mathbf{v} + \delta \in \text{rdy}(p) \quad [\text{DEL}] \\
\frac{(p\{\text{rec}X.p/X\}, \rho, \mathbf{v}) \xrightarrow{\alpha}_s (p', \rho', \mathbf{v}')}{(\text{rec}X.p, \rho, \mathbf{v}) \xrightarrow{\alpha}_s (p', \rho', \mathbf{v}')} \quad [\text{REC}] \\
\frac{(p, \rho, \mathbf{v}) \xrightarrow{\tau}_s (p', \rho', \mathbf{v}')}{(p, \rho, \mathbf{v}) \mid (q, \sigma, \eta) \xrightarrow{\tau}_s (p', \rho', \mathbf{v}') \mid (q, \sigma, \eta)} \quad [\text{S-}\oplus] \\
\frac{(p, \rho, \mathbf{v}) \xrightarrow{\delta}_s (p, \rho, \mathbf{v}') \quad (q, \sigma, \eta) \xrightarrow{\delta}_s (q, \sigma, \eta')}{(p, \rho, \mathbf{v}) \mid (q, \sigma, \eta) \xrightarrow{\delta}_s (p, \rho, \mathbf{v}') \mid (q, \sigma, \eta')} \quad [\text{S-DEL}] \\
\frac{(p, \rho, \mathbf{v}) \xrightarrow{\text{!a}}_s (p', \rho', \mathbf{v}') \quad (q, \sigma, \eta) \xrightarrow{\text{?a}}_s (q', \sigma', \eta')}{(p, \rho, \mathbf{v}) \mid (q, \sigma, \eta) \xrightarrow{\tau}_s (p', \rho', \mathbf{v}') \mid (q', \sigma', \eta')} \quad [\text{S-}\tau] \\
\text{rdy}(\oplus \text{!a}_i\{g_i, R_i\}.p_i) = \downarrow \cup \llbracket g_i \rrbracket \quad \text{rdy}(\sum \dots) = \text{rdy}(\mathbf{1}) = \mathbb{V} \\
\text{rdy}(\text{rec}X.p) = \text{rdy}(p\{\text{rec}X.p/X\})
\end{array}$$

Figure 1: Semantics of synchronous timed session types (symmetric rules omitted).

Synchronous semantics. Semantics of TSTs is given in terms of a timed labelled transition relation between configurations (defined below). Labels (ranged over by α, α', \dots) are either silent actions τ , delays $\delta > 0$, or branch labels. Labels δ model elapse of time. Branch labels and τ model discrete actions, and take no time.

Definition 4. (Configurations) A configuration is a term of the form $(p, \rho, \mathbf{v}) \mid (q, \sigma, \eta)$. A configuration is synchronous if $|\rho|, |\sigma| \leq 1$.

We are now ready to define the synchronous semantics of TSTs. Unlike [4], that uses committed choices, we use queues in synchronous configurations to simplify the comparison of synchronous and asynchronous semantics. The two semantics are equivalent.

Definition 5. (Synchronous semantics of TSTs) The semantics of TSTs is defined as the smallest labelled relation between synchronous configurations closed under the rules in Figure 1. As usual, we denote with \rightarrow_s^* the reflexive and transitive closure of the relation \rightarrow_s .

Rule $[\oplus]$ allows to commit to the branch \mathbf{a} of an internal choice, when the corresponding guard is satisfied in the clock valuation \mathbf{v} and the queue is empty. This results in the configuration $(p, \mathbf{a}, \mathbf{v})$ which can only fire !a ($[\text{!}]$). Rule $[?]$ allows an external choice to fire any of its enabled input actions. Note that the queue must be empty. This is what makes the semantics synchronous. Indeed, without the emptiness requirement, we would have obtained a 1-bounded asynchronous semantics. Rule $[\text{DEL}]$ allows time to pass; this is always possible for external choices and success term $\mathbf{1}$, while for an internal choice we require, through the function rdy , that some guard remains satisfiable (now or in the future). Note that also here we require empty queues: this guarantees that messages are read at the same time of writing, i.e. communication is synchronous. The other rules are almost standard.

Example 1. Let $p = !a \oplus !b\{t \geq 2\}$ and $q = ?b\{t \geq 5\}$. p internally chooses whether to send a at any time or b after a delay of at least 2 time units. q instead waits for a b message after a delay of 5 time units. Three possible reductions of the composition of p with q are the following:

$$(p, \emptyset, \nu_0) \mid (q, \emptyset, \eta_0) \xrightarrow{7}_s \xrightarrow{\tau}_s (\mathbf{1}, b, \nu_0 + 7) \mid (q, \emptyset, \eta_0 + 7) \xrightarrow{\tau}_s (\mathbf{1}, \emptyset, \nu_0 + 7) \mid (\mathbf{1}, \emptyset, \eta_0 + 7) \quad (1)$$

$$(p, \emptyset, \nu_0) \mid (q, \emptyset, \eta_0) \xrightarrow{\delta}_s \xrightarrow{\tau}_s (\mathbf{1}, a, \nu_0 + \delta) \mid (q, \emptyset, \eta_0 + \delta) \quad (2)$$

$$(p, \emptyset, \nu_0) \mid (q, \emptyset, \eta_0) \xrightarrow{3}_s \xrightarrow{\tau}_s (\mathbf{1}, b, \nu_0 + 3) \mid (q, \emptyset, \eta_0 + 3) \quad (3)$$

The computation in (1) reaches success. In (2), p commits to the choice $!a$ after some delay δ ; at this point, time cannot pass, and no synchronisation is possible. In (3), p commits to $!b$ after 3 time units; here, the rightmost endpoint would offer $?b$, — but not in the time chosen by the leftmost endpoint.

Synchronous progress. We recall the progress based notion of compliance of [4], that we refer here as synchronous compliance. TSTs p and q are synchronous compliant when their composition never reaches a deadlock state.

Definition 6 (Synchronous compliance). We say that $(p, \rho, \nu) \mid (q, \sigma, \eta)$ is success whenever $p = \mathbf{1} = q$ and $\rho = \emptyset = \sigma$. We say that $(p, \rho, \nu) \mid (q, \sigma, \eta)$ is s-stuck whenever $(p, \rho, \nu) \mid (q, \sigma, \eta) \not\xrightarrow{\tau}_s$ and there is no δ such that $(p, \rho, \nu) \mid (q, \sigma, \eta) \xrightarrow{\delta}_s \xrightarrow{\tau}_s$. We say that $(p, \rho, \nu) \mid (q, \sigma, \eta)$ is s-deadlock whenever (i) $(p, \rho, \nu) \mid (q, \sigma, \eta)$ not success, and (ii) $(p, \rho, \nu) \mid (q, \sigma, \eta)$ is s-stuck. We then write $(p, \nu) \bowtie_s (q, \eta)$ whenever:

$$(p, \emptyset, \nu) \mid (q, \emptyset, \eta) \rightarrow_s^* (p', \rho, \nu') \mid (q', \sigma, \eta') \quad \text{implies} \quad (p', \rho, \nu') \mid (q', \sigma, \eta') \text{ not s-deadlock}$$

We say that p and q are synchronous compliant whenever $(p, \nu_0) \bowtie_s (q, \eta_0)$ (in short, $p \bowtie_s q$).

Example 2. Let $p = ?a\{t \leq 3\}.!b\{t \leq 3\}$. We have that p is compliant with $q = !a\{t \leq 2\}.?b\{t \leq 3\}$, but it is not compliant with $q' = !a\{t \leq 4\}.?b\{t \leq 4\}$.

3 Input urgent asynchronous timed session types

In this section we introduce input urgent asynchronous semantics, the associated notion of progress, and we show some relationships with the synchronous semantics.

Input urgent asynchronous semantics. We now introduce the input urgent semantics of TSTs. Note that here we use configurations (Definition 4), i.e. queues can be unbounded.

Definition 7. (Input urgent asynchronous semantics of TSTs) The input urgent asynchronous semantics of TSTs is defined as the smallest labelled relation between configuration closed under the rules in Figure 2. As usual, we denote with \rightarrow_a^* the reflexive and transitive closure of the relation \rightarrow_a .

Rule $[\oplus]$ allows to append the message a to the queue, when the corresponding guard is satisfied in the clock valuation ν . Rule $[\!]$ just says that the message in the head of the queue can be consumed by the communication partner. Rule $[\?]$ allows an external choice to fire any of its enabled input actions. Rule $[\text{DEL}]$ allows time to pass; this is always possible for external choices and success term, while for an internal choice we require, through the function rdy , that some guard remains satisfiable. Rule $[\text{S-DEL}]$

$$\begin{array}{c}
(\mathbf{!a}\{g, R\}.p \oplus p', \rho, \mathbf{v}) \xrightarrow{\tau}_a (p, \rho; \mathbf{a}, \mathbf{v}[R]) \quad \text{if } \mathbf{v} \in \llbracket g \rrbracket \quad [\oplus] \\
(p, \mathbf{a}; \rho, \mathbf{v}) \xrightarrow{\mathbf{!a}}_a (p, \rho, \mathbf{v}) \quad [!] \\
(\mathbf{?a}\{g, R\}.p + p', \rho, \mathbf{v}) \xrightarrow{\mathbf{?a}}_a (p, \rho, \mathbf{v}[R]) \quad \text{if } \mathbf{v} \in \llbracket g \rrbracket \quad [?] \\
(p, \rho, \mathbf{v}) \xrightarrow{\delta}_a (p, \rho, \mathbf{v} + \delta) \quad \text{if } \mathbf{v} + \delta \in \text{rdy}(p) \quad [\text{DEL}] \\
\frac{(p\{p/x\}, \rho, \mathbf{v}) \xrightarrow{\alpha}_a (p', \rho', \mathbf{v}')}{(\text{rec } X.p, \rho, \mathbf{v}) \xrightarrow{\alpha}_a (p', \rho', \mathbf{v}')} \quad [\text{REC}] \\
\frac{(p, \rho, \mathbf{v}) \xrightarrow{\tau}_a (p', \rho', \mathbf{v}')}{(p, \rho, \mathbf{v}) \mid (q, \sigma, \eta) \xrightarrow{\tau}_a (p', \rho', \mathbf{v}') \mid (q, \sigma, \eta)} \quad [\text{S-}\oplus] \\
\frac{(p, \rho, \mathbf{v}) \xrightarrow{\delta}_a (p, \rho, \mathbf{v}') \quad (q, \sigma, \eta) \xrightarrow{\delta}_a (q, \sigma, \eta') \quad \forall \delta' < \delta : (p, \rho, \mathbf{v}) \mid (q, \sigma, \eta) \text{ not } \delta' \text{-sync}}{(p, \rho, \mathbf{v}) \mid (q, \sigma, \eta) \xrightarrow{\delta}_a (p, \rho, \mathbf{v}') \mid (q, \sigma, \eta')} \quad [\text{S-DEL}] \\
\frac{(p, \rho, \mathbf{v}) \xrightarrow{\mathbf{!a}}_a (p', \rho', \mathbf{v}') \quad (q, \sigma, \eta) \xrightarrow{\mathbf{?a}}_a (q', \sigma', \eta')}{(p, \rho, \mathbf{v}) \mid (q, \sigma, \eta) \xrightarrow{\tau}_a (p', \rho', \mathbf{v}') \mid (q', \sigma', \eta')} \quad [\text{S-}\tau] \\
(p, \rho, \mathbf{v}) \mid (q, \sigma, \eta) \delta \text{-sync} \iff \exists \mathbf{a} : \begin{cases} (p, \rho, \mathbf{v} + \delta) \xrightarrow{\mathbf{!a}}_a \wedge (q, \sigma, \eta + \delta) \xrightarrow{\mathbf{?a}}_a & \text{or} \\ (p, \rho, \mathbf{v} + \delta) \xrightarrow{\mathbf{?a}}_a \wedge (q, \sigma, \eta + \delta) \xrightarrow{\mathbf{!a}}_a \end{cases}
\end{array}$$

Figure 2: Urgent semantics of asynchronous timed session types (symmetric rules omitted).

allows time to pass for composite systems. While the first two premises are standard, the third one is what makes the semantics urgent: we require, through the predicate δ -sync, that elapsing of time does not prevent nor delay any possible communication. The other rules are almost standard.

Example 3. Let $p = \mathbf{!a}\{t \leq 2\}. \mathbf{!b}\{t \leq 3\}$, $q = \mathbf{?a}\{t \geq 4\}. \mathbf{?b}\{t \geq 5\}$, a possible execution of the system is:

$$(p, \emptyset, \mathbf{v}_0) \mid (q, \emptyset, \eta_0) \xrightarrow{\tau}_a \xrightarrow{\tau}_a (\mathbf{1}, \mathbf{a}; \mathbf{b}, \mathbf{v}_0) \mid (q, \emptyset, \eta_0) \xrightarrow{4}_a \xrightarrow{\tau}_a \xrightarrow{1}_a \xrightarrow{\tau}_a$$

Where the for τ actions represent, respectively, an output of \mathbf{a} , an output of \mathbf{b} , an input of \mathbf{a} , and an input of \mathbf{b} . Note that urgency prevents transitions $(\mathbf{1}, \mathbf{a}; \mathbf{b}, \mathbf{v}_0) \mid (q, \emptyset, \eta_0) \xrightarrow{\delta}_a$ if $\delta > 4$. Let $q' = \mathbf{?a}\{t > 4\}. \mathbf{?b}\{t \geq 5\}$, i.e. q' is like q but the constraint $t \geq 4$ is substituted with $t > 4$. Message \mathbf{a} cannot be consumed anymore:

$$(p, \emptyset, \mathbf{v}_0) \mid (q, \emptyset, \eta_0) \xrightarrow{\tau}_a \xrightarrow{\tau}_a \xrightarrow{4}_a (\mathbf{1}, \mathbf{a}; \mathbf{b}, \mathbf{v}_0 + 4) \mid (q, \emptyset, \eta_0 + 4)$$

Configuration $(\mathbf{1}, \mathbf{a}; \mathbf{b}, \mathbf{v}_0 + 4) \mid (q, \emptyset, \eta_0 + 4)$ cannot read \mathbf{a} (because $\mathbf{v}_0 + 4(t) = 4$ and $4 \not> 4$), and cannot delay, because for any δ there is a $\delta' < \delta$ such that $(\mathbf{1}, \mathbf{a}; \mathbf{b}, \mathbf{v}_0 + 4) \mid (q, \emptyset, \eta_0 + 4)$ is δ' -sync. Problems like that are well-known when dealing with urgency [9], therefore it is usually assumed that there is a first instant in which an urgent action becomes enabled. In our setting, this assumption corresponds to forbid guards in the form $x > n$. We do not make this assumption here just because our result does not rely on it.

Synchrony vs asynchrony. We remark some differences between the semantics in Figures 1 and 2:

- Synchronous semantics is defined on synchronous configurations, namely buffers are 1-bounded; asynchronous semantics is defined on configurations, where buffers are unbounded.
- With synchronous semantics, by rule [DEL] of Figure 1, time can pass only if all buffer are empty. With asynchronous semantics, time may pass even if buffers are not empty. However, by rule [S-DEL] of Figure 2, not empty buffers still constrain time passing: message consumption is never delayed.

Example 4. Let p and q be as in Example 3. Intuitively, their composition should not succeed with synchronous semantics, as p writes a strictly earlier than when q is going to read it. A possible complete execution is:

$$(p, \emptyset, v_0) \mid (q, \emptyset, \eta_0) \xrightarrow{\tau}_s (p, a, v_0) \mid (q, \emptyset, \eta_0)$$

Where $(p, a, v_0) \mid (q, \emptyset, \eta_0)$ is s -deadlock: it is not success, it cannot perform actions, and it cannot delay (one buffer is not empty).

Below, we sketch a proof of some relations between synchronous and asynchronous semantics. Namely, asynchronous semantics simulates the synchronous one. Furthermore, when queues are empty, asynchronous delays are mimicked by the synchronous semantics.

Lemma 1. Let $(p, \rho, v) \mid (q, \sigma, \eta)$ be a synchronous configuration. Then:

$$(p, \rho, v) \mid (q, \sigma, \eta) \xrightarrow{\alpha}_s (p', \rho', v') \mid (q', \sigma', \eta') \implies (p, \rho, v) \mid (q, \sigma, \eta) \xrightarrow{\alpha}_a (p', \rho', v') \mid (q', \sigma', \eta')$$

Furthermore:

$$(p, \emptyset, v) \mid (q, \emptyset, \eta) \xrightarrow{\delta}_a (p', \rho', v') \mid (q', \sigma', \eta') \implies (p, \emptyset, v) \mid (q, \emptyset, \eta) \xrightarrow{\delta}_s (p', \rho', v') \mid (q', \sigma', \eta')$$

Proof. First note the following facts (can be easily proved by rule induction):

$$\forall (p, \rho, v), (p', \rho', v') : (p, \rho, v) \xrightarrow{\alpha}_s (p', \rho', v') \implies (p, \rho, v) \xrightarrow{\alpha}_a (p', \rho', v') \quad (4)$$

$$\forall (p, \rho, v) : (p, \rho, v) \xrightarrow{\delta}_s \implies \rho = \emptyset \quad (5)$$

$$\forall p, v, p', v' : (p, \emptyset, v) \xrightarrow{\delta}_a (p', \emptyset, v') \implies (p, \emptyset, v) \xrightarrow{\delta}_s (p', \emptyset, v') \quad (6)$$

Back to the main statement, the first part can be proved by cases on the rule used in the derivation of $(p, \rho, v) \mid (q, \sigma, \eta) \xrightarrow{\alpha}_s (p', \rho', v') \mid (q', \sigma', \eta')$. We only show the more complicated case, namely rule [S-DEL]. Suppose:

$$\frac{(p, \rho, v) \xrightarrow{\delta}_s (p, \rho, v') \quad (q, \sigma, \eta) \xrightarrow{\delta}_s (q, \sigma, \eta')}{(p, \rho, v) \mid (q, \sigma, \eta) \xrightarrow{\delta}_s (p, \rho, v') \mid (q, \sigma, \eta')}$$

By Equation (5): $\sigma = \emptyset = \rho$. Therefore, by an inspection of the rules in Figure 2, we can conclude that, for all δ' and for all a , both $(p, \rho, v) \not\xrightarrow{\delta'}_a$ and $(q, \sigma, \eta) \not\xrightarrow{\delta'}_a$. So, $(p, \rho, v) \mid (q, \sigma, \eta)$ not δ' -sync for any δ' . Therefore, thanks to Equation (4), we can use rule [S-DEL]:

$$\frac{(p, \rho, v) \xrightarrow{\delta}_a (p, \rho, v') \quad (q, \sigma, \eta) \xrightarrow{\delta}_a (q, \sigma, \eta') \quad \forall \delta' < \delta : (p, \rho, v) \mid (q, \sigma, \eta) \text{ not } \delta' \text{-sync}}{(p, \rho, v) \mid (q, \sigma, \eta) \xrightarrow{\delta}_a (p, \rho, v') \mid (q, \sigma, \eta')}$$

For the furthermore case, the only possibility is:

$$\frac{(p, \emptyset, v) \xrightarrow{\delta}_a (p, \emptyset, v') \quad (q, \emptyset, \eta) \xrightarrow{\delta}_a (q, \emptyset, \eta') \quad \forall \delta' < \delta : (p, \emptyset, v) \mid (q, \emptyset, \eta) \text{ not } \delta' - \text{sync}}{(p, \emptyset, v) \mid (q, \emptyset, \eta) \xrightarrow{\delta}_a (p, \emptyset, v') \mid (q, \emptyset, \eta')}_{\text{[S-DEL]}}$$

By Equation (6):

$$\frac{(p, \emptyset, v) \xrightarrow{\delta}_s (p, \emptyset, v') \quad (q, \emptyset, \eta) \xrightarrow{\delta}_s (q, \emptyset, \eta')}{(p, \emptyset, v) \mid (q, \emptyset, \eta) \xrightarrow{\delta}_s (p, \emptyset, v') \mid (q, \emptyset, \eta')}_{\text{[S-DEL]}}$$

□

Asynchronous progress. We extend the notion of progress to the asynchronous setting. It differs from Definition 6 only in that it uses the asynchronous semantics.

Definition 8 (Asynchronous compliance). *We say that $(p, \rho, v) \mid (q, \sigma, \eta)$ is a-stuck whenever $(p, \rho, v) \mid (q, \sigma, \eta) \not\xrightarrow{\tau}_a$ and there is no δ such that $(p, \rho, v) \mid (q, \sigma, \eta) \xrightarrow{\delta}_s \xrightarrow{\tau}_a$. We say that $(p, \rho, v) \mid (q, \sigma, \eta)$ is a-deadlock whenever (i) $(p, \rho, v) \mid (q, \sigma, \eta)$ not success, and (ii) $(p, \rho, v) \mid (q, \sigma, \eta)$ is a-stuck.. We then write $(p, v) \bowtie_a (q, \eta)$ whenever:*

$$(p, \emptyset, v) \mid (q, \emptyset, \eta) \rightarrow_a^* (p', \rho, v') \mid (q', \sigma, \eta') \quad \text{implies} \quad (p', \rho, v') \mid (q', \sigma, \eta') \text{ not a-deadlock}$$

We say that p and q are asynchronous compliant whenever $(p, v_0) \bowtie_a (q, \eta_0)$ (in short, $p \bowtie_a q$).

4 Results

In this section we sketch a proof of the main result of the paper (Theorem 1), namely that synchronous progress implies asynchronous progress. The proof is quite standard: we introduce a property (being the composition of *r-compliant* TSTs, Definition 9) that is enjoyed by $(p, \emptyset, v_0) \mid (q, \emptyset, \eta_0)$, provided $p \bowtie_a q$. We then show that r-compliance is preserved by transitions (Proposition 1) and that configurations of r-compliant TSTs are not a-deadlock. (Proposition 2).

R-compliance is defined below. It is based on the notion of *remainder*, that, given a configuration and a queue, returns the configuration obtained after consuming the given queue immediately (without delays). Note that the remainder is a partial operation, and it is not defined if the queue cannot be consumed, or some delay is required. Then, r-compliance requires that:

- Queues can be consumed immediately.
- The resulting configuration is composed by synchronous compliant TSTs.

Definition 9. *We define the remainder of (p, ρ, v) and queue σ , in symbols $(p, \rho, v) - \sigma$, inductively as follows:*

$$\begin{aligned} (p, \rho, v) - \emptyset &= (p, \rho, v) \\ (p, \rho, v) - a; \sigma &= (p', \rho', v') - \sigma \quad \text{if } (p, \rho, v) \xrightarrow{?a}_a (p', \rho', v') \end{aligned}$$

We say that (p, ρ, v) is r-compliant with (q, σ, η) (in symbols $(p, \rho, v) \bowtie (q, \sigma, \eta)$) if, for some p', v', q', η' :

$$(p, \rho, v) - \sigma = (p', \rho, v') \wedge (q, \sigma, \eta) - \rho = (q', \sigma, \eta') \wedge (p', v') \bowtie_s (q', \eta')$$

The following auxiliary lemma says that the r-compliant configurations, under asynchronous semantics, never allow delays unless both the queues are empty.

Lemma 2. *If $(p, \rho, \nu) \boxtimes (q, \sigma, \eta)$ and $(p, \rho, \nu) \mid (q, \sigma, \eta) \xrightarrow{\delta}_a$, then $\rho = \emptyset = \sigma$.*

Proof. Suppose $(p, \rho, \nu) \mid (q, \sigma, \eta) \xrightarrow{\delta}_a$. First note that the only applicable rule is [S-DEL]. We have to show $\rho = \emptyset = \sigma$. Suppose, by contradiction, this is not the case, and assume that, say, $\rho = \mathbf{a}; \rho''$ for some \mathbf{a}, ρ'' . By rule [⊕], $(p, \rho, \nu) \xrightarrow{!a}_a$. Since $(p, \rho, \nu) \boxtimes (q, \sigma, \eta)$, it follows that $(q, \sigma, \eta) - \mathbf{a}; \rho''$ is defined. Then, by Definition 9, it must be $(q, \sigma, \eta) \xrightarrow{?a}_a$. But then $(p, \rho, \nu) \mid (q, \sigma, \eta)$ is 0-sync, and so rule [S-DEL] does not apply: contradiction. \square

The following proposition states that r-compliance is preserved by asynchronous transitions.

Proposition 1. *Let $(p, \rho, \nu) \boxtimes (q, \sigma, \eta)$ and $(p, \rho, \nu) \mid (q, \sigma, \eta) \xrightarrow{\alpha}_a (p', \rho', \nu') \mid (q', \sigma', \eta')$. Then:*

$$(p', \rho', \nu') \boxtimes (q', \sigma', \eta')$$

Proof. Since $(p, \rho, \nu) \boxtimes (q, \sigma, \eta)$, there exist p'', ν'', q'', η'' such that:

$$(p, \rho, \nu) - \sigma = (p'', \rho, \nu'') \wedge (q, \sigma, \eta) - \rho = (q'', \sigma, \eta'') \wedge (p'', \nu'') \boxtimes_s (q'', \eta'')$$

We proceed by cases on the rule used.

- [S-⊕]. It must be $(q, \sigma, \eta) = (q', \sigma', \eta')$, $\rho' = \rho; \mathbf{a}$ for some \mathbf{a} , and $(p, \rho, \nu) \xrightarrow{\tau}_a (p', \rho', \nu')$. By an inspection of rules in Figure 2, we can conclude that $(p, \rho, \nu) \not\xrightarrow{?b}_a$ for all \mathbf{b} . Then $\sigma = \emptyset$: otherwise, $(p, \rho, \nu) - \sigma$ would be undefined. Therefore, $(p, \rho, \nu) = (p'', \rho, \nu'')$ and $(p', \rho', \nu') - \sigma = (p', \rho', \nu')$. By a simple induction on the length of ρ , we can conclude $(q, \sigma, \eta) - \rho; \mathbf{a} = (q'', \sigma, \eta'') - \mathbf{a}$. We have to show that $(q'', \sigma, \eta'') - \mathbf{a} = (q''', \sigma, \eta''')$ for some q''', η''' such that $(p', \nu') \boxtimes_s (q''', \eta''')$. Note that, since $(p'', \nu'') \boxtimes_s (q'', \eta'')$ and (p'', ν'') writes \mathbf{a} , it must be (by lemmas A.2 and 3.6 of [4], modulo minor notational differences) $(q'', \sigma, \eta'') \xrightarrow{?a}_s (q''', \sigma, \eta''')$, with $(p', \nu') \boxtimes_s (q''', \eta''')$. By Equation (4) in the proof of Lemma 1, $(q'', \sigma, \eta'') \xrightarrow{?a}_a (q''', \sigma, \eta''')$. Therefore, $(q'', \sigma, \eta'') - \mathbf{a} = (q''', \sigma, \eta''')$ with $(p'', \nu'') \boxtimes_s (q'', \eta'')$, and we are done.
- [S-DEL]. By Lemma 2, it follows $\rho = \emptyset = \sigma$, and by rule [S-DEL] $\rho' = \emptyset = \sigma'$ as well. By Definition 9, $(p, \nu) \boxtimes_s (q, \eta)$. By Lemma 1 and Definition 6, it follows $(p', \nu') \boxtimes_s (q', \eta')$, and so, since ρ' and σ' are both empty, $(p', \rho', \nu') \boxtimes (q', \sigma', \eta')$.
- [S-τ]. It must be $(p, \rho, \nu) \xrightarrow{!a}_a (p', \rho', \nu')$ and $(q, \sigma, \eta) \xrightarrow{?a}_a (q', \sigma', \eta')$. By a simple induction on the rules in Figure 2, we can conclude $\rho = \mathbf{a}; \rho'$, $\sigma = \sigma'$, $p = p'$ (up to unfolding of recursion), and $\nu = \nu'$. Since $(q, \sigma, \eta) \xrightarrow{?a}_a (q', \sigma', \eta')$, by Definition 9 it follows that $(q, \sigma, \eta) - \mathbf{a}; \rho' = (q', \sigma, \eta') - \rho'$. Clearly, up to unfolding of recursion, $(p, \rho, \nu) - \sigma = (p', \rho, \nu) - \sigma'$. Therefore, since $(p, \rho, \nu) \boxtimes (q, \sigma, \eta)$ by assumption, also $(p', \rho', \nu') \boxtimes (q', \sigma', \eta')$. \square

The following proposition states a-deadlock freedom of r-compliant configurations.

Proposition 2. *If $(p, \rho, \nu) \boxtimes (q, \sigma, \eta)$, then $(p, \rho, \nu) \mid (q, \sigma, \eta)$ is not a-deadlock.*

Proof. We have two cases:

- $\rho = \emptyset \wedge \sigma = \emptyset$. Then, $(p, \rho, \nu) - \sigma = (p, \rho, \nu)$ and $(q, \sigma, \eta) - \rho = (q, \sigma, \eta)$, with $(p, \nu) \bowtie_s (q, \eta)$. Therefore, $(p, \rho, \nu) \mid (q, \sigma, \eta)$ is not s-deadlock. By Definition 6, if $p = \mathbf{1} = q$, then $(p, \rho, \nu) \mid (q, \sigma, \eta)$ is success and therefore not a-deadlock. If it is not the case that $p = \mathbf{1} = q$, by Definition 6 there is δ such that $(p, \rho, \nu) \mid (q, \sigma, \eta) \xrightarrow{\delta}_s \xrightarrow{\tau}_s$ (we omit the simpler case where the τ move is performed without delay). Then, by Lemma 1, $(p, \rho, \nu) \mid (q, \sigma, \eta) \xrightarrow{\delta}_a \xrightarrow{\tau}_a$. Therefore $(p, \rho, \nu) \mid (q, \sigma, \eta)$ is not a-deadlock.
- $\rho \neq \emptyset \vee \sigma \neq \emptyset$. We show only the case $\rho \neq \emptyset$. The other is similar. It must be $\rho = \mathbf{a}; \rho'$ for some \mathbf{a} and ρ' . Since $(q, \sigma, \eta) - \rho$ is defined, it must be $(q, \sigma, \eta) \xrightarrow{?a}$. Therefore, by rule [S- τ], $(p, \rho, \nu) \mid (q, \sigma, \eta) \xrightarrow{\tau}_a$, and so $(p, \rho, \nu) \mid (q, \sigma, \eta)$ is not a-deadlock. □

The main result follows.

Theorem 1. *If $p \bowtie_s q$ then $p \bowtie_a q$.*

Proof. Let $p \bowtie_s q$, and assume $(p, \emptyset, \nu_0) \mid (q, \emptyset, \eta_0) \rightarrow_a^* (p', \rho', \nu') \mid (q', \rho', \eta')$. We have to show $(p', \rho', \nu') \mid (q', \rho', \eta')$ is not a-deadlock. First note that, since queues are empty and $p \bowtie_s q$, it holds that $(p, \emptyset, \nu_0) \bowtie (q, \emptyset, \eta_0)$. By Proposition 1 and a simple induction on the length of the reduction, we can derive $(p', \rho', \nu') \bowtie (q', \rho', \eta')$. By Proposition 2, $(p', \rho', \nu') \mid (q', \rho', \eta')$ is not a-deadlock. □

5 Conclusions and related work

Following [4], we pursued a line of research aimed at lifting key properties of session types to the timed setting. We have shown that the interesting property of preservation of untimed synchronous progress when passing to asynchronous semantics (discovered in [6]), can be recovered using a certain urgent asynchronous semantics, that closely models realistic programming primitives.

Timed session types have been introduced in [8], in the multiparty asynchronous version, where they have been used to statically type check a timed π -calculus. Their theory has subsequently been extended to dynamic verification [13]. [4] introduced the binary and synchronous theory, subsequently applied in a contract-oriented middleware [5] and the companion verification tool-chain [2]. [7] studies progress in the context of Communicating Timed Automata [12]. Several works study urgency in timed systems, here we mention [9]. As far as we know, urgency in the context of asynchronous communication has only been studied in [3], where the idea of modelling input primitives with input urgency originates.

References

- [1] Rajeev Alur & David L. Dill (1994): *A Theory of Timed Automata*. *Theor. Comput. Sci.* 126(2), pp. 183–235, doi:10.1016/0304-3975(94)90010-8.
- [2] Nicola Atzei & Massimo Bartoletti (2016): *Developing Honest Java Programs with Diogenes*. In: *FORTE 2016, LNCS 9688*, Springer, pp. 52–61, doi:10.1007/978-3-319-39570-8_4.
- [3] Massimo Bartoletti, Laura Bocchi & Maurizio Murgia (2018): *Progress-preserving Refinements of CTA*. In: *CONCUR 2018*. To appear.
- [4] Massimo Bartoletti, Tiziana Cimoli & Maurizio Murgia (2017): *Timed Session Types*. *Logical Methods in Computer Science* 13(4), doi:10.23638/LMCS-13(4:25)2017.
- [5] Massimo Bartoletti, Tiziana Cimoli, Maurizio Murgia, Alessandro Sebastian Podda & Livio Pompianu (2015): *A contract-oriented middleware*. In: *FACS, LNCS 9539*, Springer, pp. 86–104, doi:10.1007/978-3-319-28934-2_5.
- [6] Massimo Bartoletti, Alceste Scalas & Roberto Zunino (2014): *A Semantic Deconstruction of Session Types*. In: *Proc. CONCUR, LNCS 8704*, Springer, pp. 402–418, doi:10.1007/978-3-662-44584-6_28.
- [7] Laura Bocchi, Julien Lange & Nobuko Yoshida (2015): *Meeting Deadlines Together*. In: *CONCUR, LIPIcs 42*, Schloss Dagstuhl - Leibniz-Zentrum fuer Informatik, pp. 283–296, doi:10.4230/LIPIcs.CONCUR.2015.283.
- [8] Laura Bocchi, Weizhen Yang & Nobuko Yoshida (2014): *Timed Multiparty Session Types*. In: *CONCUR, LNCS 8704*, Springer, pp. 419–434, doi:10.1007/978-3-662-44584-6_29.
- [9] Sébastien Bornot, Joseph Sifakis & Stavros Tripakis (1997): *Modeling Urgency in Timed Systems*. In: *COMPOS, LNCS 1536*, Springer, pp. 103–129, doi:10.1007/3-540-49213-5_5.
- [10] Mariangiola Dezani-Ciancaglini & Ugo de'Liguoro (2009): *Sessions and Session Types: An Overview*. In: *WS-FM, LNCS 6194*, Springer, pp. 1–28, doi:10.1007/978-3-642-14458-5_1.
- [11] Kohei Honda, Nobuko Yoshida & Marco Carbone (2016): *Multiparty Asynchronous Session Types*. *J. ACM* 63(1), pp. 9:1–9:67, doi:10.1145/2827695.
- [12] Pavel Krcál & Wang Yi (2006): *Communicating Timed Automata: The More Synchronous, the More Difficult to Verify*. In: *CAV, LNCS 4144*, Springer, pp. 249–262, doi:10.1007/11817963_24.
- [13] Rumyana Neykova, Laura Bocchi & Nobuko Yoshida (2017): *Timed runtime monitoring for multiparty conversations*. *Formal Asp. Comput.* 29(5), pp. 877–910, doi:10.1007/s00165-017-0420-8.
- [14] Kaku Takeuchi, Kohei Honda & Makoto Kubo (1994): *An Interaction-based Language and its Typing System*. In: *PARLE, LNCS 817*, Springer, pp. 398–413, doi:10.1007/3-540-58184-7_118.