

Metric Semantics for True Concurrent Real Time

CHRISTEL BAIER^a, JOOST-PIETER KATOEN^b AND DIEGO LAELLA^c

^a*Fakultät für Mathematik und Informatik, Universität Mannheim
Seminargebäude A5, D-68159 Mannheim, Germany*

^b*Lehrstuhl für Informatik VII, Friedrich-Alexander-Universität Erlangen-Nürnberg
Martensstrasse 3, D-91058 Erlangen, Germany*

^c*CNUCE Istituto del CNR, Via Santa Maria 36, I-56100 Pisa, Italy*

Abstract. This paper investigates the use of a complete metric space framework for providing denotational semantics to a real-time process algebra. The study is carried out in a non-interleaving setting and is based on a timed extension of Langerak’s bundle event structures, a variant of Winskel’s event structures. The distance function is based on the amount of time to which event structures do ‘agree’. We show that this intuitive notion of distance is a pseudo metric (but not a metric) on the set of timed event structures. A generalisation to equivalence classes of timed event structures in which we abstract from event names and non-executable events (events that can never appear) is shown to be a complete ultra-metric space. We show that the resulting metric semantics is an abstraction of an existing cpo-based denotational and a related operational semantics for the considered language.

1 Introduction

In this paper we consider a metric denotational semantics for an algebraic specification language that besides concurrency, synchronisation, and non-determinism, encompasses the notion of real-time. The language that we consider is a real-time extension of a process algebra based on the standardised specification language LOTOS [7]. As semantic domain we take a timed extension (defined in [10]) of Langerak’s bundle event structures, a variant of Winskel’s event structures that has been shown to adequately deal with the operators of LOTOS (in particular, parallel composition and disruption) [11]. The suitability of this timed truly concurrent model for modelling time-critical systems is addressed in [10]. The metric approach of this paper can also be applied to timed variants of other brands of event structures, like prime and stable event structures.

The basic idea of this paper is to consider behaviours of event structures up to a certain time. This is in fact a continuous version of the idea in [12] to consider (untimed) event structures up to a certain depth (i.e. length of a causal chain). The distance function is based on the amount of time to which event structures do ‘agree’. We show that this intuitive notion of distance is a pseudo metric (but not a metric) on TES, the set of timed event structures. As a first step towards obtaining a metric (rather than a pseudo metric), we consider TES modulo an isomorphism \simeq_{iso} that abstracts from event names and from non-executable events, events that can never appear. Secondly, we refine this notion

towards finitely approximable timed event structures module \simeq_{iso} and show that this model is a complete ultra-metric space. The resulting domain is used as a semantic domain for time-guarded processes. A process is time-guarded if it cannot generate instantaneous recursive process instantiations. We show that the proposed metric semantics is an abstraction of the cpo-based semantics of [10].

2 A real-time process algebra

We assume a given set of observable actions Obs and an *invisible action* τ ; $\tau \notin \text{Obs}$. The action \surd indicates the *successful termination* action of a process; $\surd \notin \text{Obs}$ and $\surd \neq \tau$. In addition, let $\text{Act} = \text{Obs} \cup \{\tau, \surd\}$, $a \in \text{Obs} \cup \{\tau\}$, $I \subseteq \mathbb{R}^+ \cup \{\infty\}$, $t \in \mathbb{R}^+ \cup \{\infty\}$, $A \subseteq \text{Obs}$, $\lambda : \text{Act} \rightarrow \text{Act}$ with $\lambda(\tau) = \tau$ and $\lambda(\surd) = \surd$, and Var a set of process variables with $x \in \text{Var}$. The set of expressions defined in the following is denoted Expr .

$$P ::= \mathbf{0} \mid \mathbf{1} \mid a_I . P \mid P + P \mid P ; P \mid P [> P \mid P \parallel_A P \mid P \setminus A \mid P[\lambda] \mid P \triangleright_t P \mid x.$$

$+$, $\setminus A$, and $[\lambda]$ are the usual process algebra operators choice, abstraction and relabelling, respectively.

- $\mathbf{1}$ represents the successful termination process; it can only perform action \surd and then becomes $\mathbf{0}$, the process that cannot perform any action.
- $a_I . P$ denotes the prefix of a and P where a is allowed (but not forced) to occur at $t \in I$.
- $P ; Q$ denotes the sequential composition of P and Q ; the control is passed to Q by the termination of P as indicated by the occurrence of \surd .
- $P [> Q$ denotes the disruption of P by Q ; i.e. P may at any point of its execution be disrupted by Q , unless P has terminated.
- $P \parallel_A Q$ denotes the parallel composition of P and Q ; P and Q execute actions not in A independently from each other, while actions in A (and successful termination actions) must be performed by both processes simultaneously.
- $P \triangleright_t Q$ initially behaves like P , but if P does not perform an action before time t (since its enabling) then a timeout occurs and control is passed to Q .

Using these operators a timed interrupt, for instance, can easily be modelled: the process $P [> (\mathbf{0} \triangleright_t Q)$ specifies that P is disrupted by Q at time t , unless P has terminated before. Various case studies have proven that the timed operators like $a_I . P$ and $P \triangleright_t Q$ are convenient to specify practical real-time systems [1, 18].

Process variables are considered in the context of a set of process definitions of the form $x := P$, where P might contain occurrences of x or of other process variables. For process variable x let $\text{decl}(x)$ denote the body of x , i.e. $\text{decl}(x) = P$ for $x := P$. A *process* is a pair $\langle \text{decl}, P \rangle$ consisting of a declaration $\text{decl} : \text{Var} \rightarrow \text{Expr}$ and an expression $P \in \text{Expr}$. PA denotes the set of all processes.

3 Timed event structures

Event structures consist of *events* labelled with actions (an event modelling the occurrence of its action), together with relations of causality and conflict between events. We take Langerak's (extended bundle) event structures [11] and equip this with timing information. Event structures incorporate a *conflict* relation (denoted \rightsquigarrow) that—as opposed to what is common in other types of event structures—is not required to be symmetric and a *bundle* relation (denoted \mapsto) for modelling causality.

The meaning of $e \rightsquigarrow e'$ is that (i) if e' occurs it disables the occurrence of e , and (ii) if e and e' both occur in a single system run then e precedes e' . $e \rightsquigarrow e'$ and $e' \rightsquigarrow e$ is equivalent with $e \# e'$, the usual symmetric conflict in event structures. The reason for adopting \rightsquigarrow rather than $\#$ is to model the disrupt operator $[>]$ adequately.

Causality is represented by the bundle relation. For set X of events and an event e , $X \mapsto e$ means that if e happens in a system run, some event in X must have happened before. X is called the *bundle set* and we use \mapsto to denote the set of bundles of an event structure. The reason for not having a binary causality relation between events (as in prime event structures [16]) is to model parallel composition $\parallel_{\mathcal{A}}$ in a less complex way.

Time is added to event structures in the following way [10]. Relative delays between events are attached to bundles, and delays relative to the start are attached to events. Delays determine when an event may happen, they do not specify that an event should happen at a particular time. For the latter purpose we use *urgent* events; an urgent event should happen as soon as it is enabled.

Definition 1. A *timed event structure* (tes) \mathcal{E} is a tuple $(E, \rightsquigarrow, \mapsto, l, \mathcal{A}, \mathcal{R}, \mathcal{U})$ with E , a set of *events*, $\rightsquigarrow \subseteq E \times E$, the (irreflexive) *conflict* relation, $\mapsto \subseteq \mathcal{P}(E) \times E$, the *bundle* relation, $l : E \rightarrow \text{Act}$, the *labelling* function, $\mathcal{A} : E \rightarrow \mathcal{P}(\mathbb{R}^+ \cup \{\infty\})$, the *event delay* function, $\mathcal{R} : \mapsto \rightarrow \mathcal{P}(\mathbb{R}^+ \cup \{\infty\})$, the *bundle delay* function, and $\mathcal{U} \subseteq E$, the set of *urgent* events, such that:

(P1) $(X \times X) \setminus \text{Id}_E \subseteq \rightsquigarrow$ for any bundle set X

and for all $e \in \mathcal{U}$:

(P2) $\forall e' \in E, X \subseteq E' : ((e' \rightsquigarrow e \vee e \rightsquigarrow e') \wedge X \mapsto e) \Rightarrow (X \mapsto e' \vee X \rightsquigarrow e')$

(P3) $\exists t \in \mathbb{R}^+ : (\mathcal{A}(e) \in \{\emptyset, \{t\}\}) \vee (\exists X : X \mapsto e \wedge \mathcal{R}(X, e) \in \{\emptyset, \{t\}\})$.

Here, $\mathcal{P}(\cdot)$ denotes the power-set function, $X \rightsquigarrow e'$ denotes $(\forall e'' \in X : e'' \rightsquigarrow e')$ and Id_E denotes the identity relation on set E . Note that $\emptyset \rightsquigarrow e'$ for all e' .

Event structures are depicted as follows. Events are denoted as dots; near the dot the action label is given. If no confusion arises we often use action labels rather than event identities to denote events. $e \rightsquigarrow e'$ is indicated by a dotted arrow from e to e' ; if also $e' \rightsquigarrow e$, then a dotted line is drawn instead. A bundle $X \mapsto e$ is indicated by an arrow to which each event in X is connected via a line. Bundle and event delays are depicted near to a bundle and event, respectively. Urgent events are denoted by open dots, other events by closed dots. A bundle

$X \mapsto e$ with $\mathcal{R}(X, e) = I$ is denoted by $X \xrightarrow{I} e$. Delays $[t, \infty)$ are simply denoted by t ; delays $[0, \infty)$ are usually omitted. Figure 1(a) shows an example tes with e.g. $\{a\} \xrightarrow{[0,7]} b$ and $\{a\} \xrightarrow{[0,5]} c$, $b \rightsquigarrow \tau$ and $\tau \rightsquigarrow c$. $\mathcal{U} = \{\tau\}$ and \mathcal{A} is 0 for all events.




Fig. 1. A tes and a non-tes

The concept of a system run for tes's is captured by the notion of a *timed event trace*. For σ a sequence of distinct events let the set of events enabled in \mathcal{E} after σ be defined as¹

$$\text{en}^{\mathcal{E}}(\sigma) \triangleq \{e \in E \mid e \notin \sigma \wedge (\forall e_i \in \sigma : e \not\rightsquigarrow e_i) \wedge (\forall X \mapsto e : X \cap \sigma \neq \emptyset)\}.$$

The time instants at which an enabled event in \mathcal{E} after $\sigma = (e_1, t_1) \dots (e_n, t_n)$ could potentially happen is determined as

$$\text{time}_{\sigma}^{\mathcal{E}}(e) \triangleq \mathcal{A}(e) \cap \bigcap_{e_i \rightsquigarrow e} [t_i, \infty) \cap \bigcap_{X \xrightarrow{I} e, e_i \in X} t_i + I.$$

Definition 2. $\sigma = (e_1, t_1) \dots (e_n, t_n)$ with $e_i \in E$ (all events being pairwise distinct) and $t_i \in \mathbb{R}^+$, is a *timed event trace* of $\mathcal{E} \in \text{TES}$ iff for all $0 < i \leq n$:

1. $e_j \rightsquigarrow e_i \Rightarrow (j < i \wedge t_j \leq t_i)$ for all $0 < j \leq n$
2. $X \xrightarrow{I} e_i \Rightarrow (\exists j : X \cap \{e_1, \dots, e_{i-1}\} = \{e_j\} \wedge t_i \in t_j + I)$ for all $X \subseteq E$
3. $t_i \in \mathcal{A}(e_i)$
4. $(e_i \rightsquigarrow e \vee e \rightsquigarrow e_i) \Rightarrow t_i \leq \min(\text{time}_{e_1 \dots e_{i-1}}^{\mathcal{E}}(e))$ for $e \in \mathcal{U} \cup \text{en}^{\mathcal{E}}(e_1 \dots e_{i-1})$.

The set of timed event traces of \mathcal{E} is denoted by $\text{Traces}(\mathcal{E})$.

By convention we use $\min \emptyset = \infty$. The last constraint takes care of the fact that urgent events may prevent the events that they disable (or by which they are disabled) to occur after a certain time. That is, event e_i can occur at time t_i provided there is no enabled urgent event e that disables e_i (or that is disabled by e_i) and that (if it occurs) must occur before t_i .

For example, for the following sequences of timed events the conditions are given under which they are timed event traces of Figure 1(a):

$$\begin{aligned} &(a, t_a)(c, t_c)(b, t_b) \text{ if } 0 \leq t_a \leq t_c \leq t_b \wedge t_b \leq t_a + 4 \wedge t_c \leq t_a + 4 \\ &(a, t_a)(\tau, t_\tau)(d, t_d) \text{ if } 0 \leq t_a \leq t_\tau \leq t_d \wedge t_\tau = t_a + 4. \end{aligned}$$

¹ Often the set of events of a sequence is identified with the sequence itself.

Note that Figure 1(a) models a typical timeout scenario: if after the occurrence of a neither b nor c happen within 4 time units, then a timeout (τ) is forced to occur. It τ would not be urgent, the conditions for t_a and t_b in the first case would be $t_b \leq t_a + 7$ and $t_c \leq t_a + 5$, since τ is not forced to occur and time does not resolve the choice.

4 Operators for timed event structures

In this section we present some operators on timed event structures that are needed to define a compositional semantics for PA. They are basically adopted from [9, 10]. We start with some basic notions. Let **Events** be a set such that for all actions $a \in \text{Act}$ there is an event $e_a \in \text{Events}$, and (i) if $e \in \text{Events}$ then $(e, *)$, $(*, e) \in \text{Events}$, and (ii) if $e, e' \in \text{Events}$ then $(e, e') \in \text{Events}$. Let TES denote the set of tes's \mathcal{E} with $E \subseteq \text{Events}$. Let $\text{init}(\mathcal{E})$ be the set of initial events of \mathcal{E} and $\text{exit}(\mathcal{E})$ its set of successful termination events, i.e. $\text{init}(\mathcal{E}) \triangleq \{e \in E \mid \neg(\exists X \subseteq E : X \mapsto e)\}$ and $\text{exit}(\mathcal{E}) \triangleq \{e \in E \mid l(e) = \sqrt{}\}$.

In the rest of this section let $\mathcal{E} \in \text{TES}$ and $\mathcal{E}_1 = (E_1, \rightsquigarrow_1, \mapsto_1, l_1, \mathcal{A}_1, \mathcal{R}_1, \mathcal{U}_1)$, $\mathcal{E}_2 = (E_2, \rightsquigarrow_2, \mapsto_2, l_2, \mathcal{A}_2, \mathcal{R}_2, \mathcal{U}_2)$ such that w.l.o.g. $E_1 \cap E_2 = \emptyset$. Let $\hat{\tau}$ denote the urgent variant of τ .

Definition 3. For $a \in \text{Obs} \cup \{\tau, \hat{\tau}\}$ and $I \subseteq [0, \infty)$, $a_I \cdot \mathcal{E}_1 \triangleq (E_1 \cup \{e_a\}, \rightsquigarrow_1, \mapsto, l_1 \cup \{(e_a, a)\}, \mathcal{A}, \mathcal{R}, \mathcal{U})$ where for $e \notin E_1$

- $\mapsto = \mapsto_1 \cup (\{\{e\}\} \times E_1)$
- $\mathcal{A} = \{(e, I)\} \cup (E_1 \times \{[0, \infty)\})$
- $\mathcal{R} = \mathcal{R}_1 \cup \{(\{e\}, e'), \mathcal{A}_1(e) \mid e' \in E_1\}$
- $\mathcal{U} = \text{if } a = \hat{\tau} \text{ then } \mathcal{U}_1 \cup \{e\} \text{ else } \mathcal{U}_1.$

$\hat{\tau}_I \cdot \mathcal{E}$ denotes the prefixing of τ_I and \mathcal{E} where e is declared to be urgent. The possibility $\hat{\tau}_I \cdot \mathcal{E}$ is used to define the semantics of the timeout operator \triangleright in a concise way. Notice that for $\hat{\tau}_I \cdot \mathcal{E}$ set I must be either empty or equal to $[t, t]$ for some t in order to guarantee axiom (P3).

Definition 4. $\mathcal{E}_1 + \mathcal{E}_2 \triangleq (E_1 \cup E_2, \rightsquigarrow, \mapsto_1 \cup \mapsto_2, l_1 \cup l_2, \mathcal{A}_1 \cup \mathcal{A}_2, \mathcal{R}_1 \cup \mathcal{R}_2, \mathcal{U}_1 \cup \mathcal{U}_2)$ where $\rightsquigarrow = \rightsquigarrow_1 \cup \rightsquigarrow_2 \cup (\text{init}(\mathcal{E}_1) \times \text{init}(\mathcal{E}_2)) \cup (\text{init}(\mathcal{E}_2) \times \text{init}(\mathcal{E}_1))$.

Definition 5. Let $A \subseteq \text{Obs}$. Then $\mathcal{E} \setminus A \triangleq (E, \rightsquigarrow, \mapsto, l', \mathcal{A}, \mathcal{R}, \mathcal{U})$ where $(l(e) \in A \Rightarrow l'(e) = \tau) \wedge (l(e) \notin A \Rightarrow l'(e) = l(e))$.


Definition 6. For $\lambda : \text{Act} \rightarrow \text{Act}$ with $\lambda(\tau) = \tau$ and $\lambda(\sqrt{}) = \sqrt{}$ let $\mathcal{E}[\lambda] \triangleq (E, \rightsquigarrow, \mapsto, \lambda \circ l, \mathcal{A}, \mathcal{R}, \mathcal{U})$.

Definition 7. $\mathcal{E}_1 ; \mathcal{E}_2 \triangleq (E_1 \cup E_2, \rightsquigarrow, \mapsto, l, \mathcal{A}, \mathcal{R}, \mathcal{U}_1 \cup \mathcal{U}_2)$ where

- $\rightsquigarrow = \rightsquigarrow_1 \cup \rightsquigarrow_2 \cup (\text{exit}(\mathcal{E}_1) \times \text{exit}(\mathcal{E}_1)) \setminus \text{Id}_{E_1}$
- $\mapsto = \mapsto_1 \cup \mapsto_2 \cup (\{\text{exit}(\mathcal{E}_1)\} \times E_2)$
- $l = ((l_1 \cup l_2) \setminus (\text{exit}(\mathcal{E}_1) \times \{\sqrt{}\})) \cup (\text{exit}(\mathcal{E}_1) \times \{\tau\})$
- $\mathcal{A} = \mathcal{A}_1 \cup (E_2 \times \{[0, \infty)\})$

$$- \mathcal{R} = \mathcal{R}_1 \cup \mathcal{R}_2 \cup \{ ((\text{exit}(\mathcal{E}_1), e), \mathcal{A}_2(e)) \mid e \in E_2 \}.$$

As an example of how $\mathcal{E}_1 ; \mathcal{E}_2$ is computed consider:



Notice that the delay of d in \mathcal{E}_2 now becomes relative to the termination of \mathcal{E}_1 .

Definition 8. $\mathcal{E}_1 [> \mathcal{E}_2 \triangleq (E_1 \cup E_2, \rightsquigarrow, \mapsto_1 \cup \mapsto_2, l_1 \cup l_2, \mathcal{A}_1 \cup \mathcal{A}_2, \mathcal{R}_1 \cup \mathcal{R}_2, \mathcal{U}_1 \cup \mathcal{U}_2)$ where $\rightsquigarrow = \rightsquigarrow_1 \cup \rightsquigarrow_2 \cup (E_1 \times \text{init}(\mathcal{E}_2)) \cup (\text{init}(\mathcal{E}_2) \times \text{exit}(\mathcal{E}_1))$.

The events of $\mathcal{E}_1 \parallel_A \mathcal{E}_2$ are constructed in the following way: an event e of E_i ($i=1, 2$) that does not need to synchronise is paired with the auxiliary symbol $*$, and an event which is labelled with \checkmark or with an action in A is paired with all events (if any) in the other tes that are equally labelled. Two events are put in conflict if any of their components are in conflict, or if different events have a common component different from $*$ (such events appear if two or more events in one tes synchronise with the same event in the other tes). A bundle is introduced iff when we take the projection on the component \mathcal{E}_i of the bundle-set we obtain a bundle in \mapsto_i . Let for $A \subseteq \text{Obs}$, $E_i^s \triangleq \{ e \in E_i \mid l_i(e) \in A \cup \{ \checkmark \} \}$ be the set of synchronising events and $E_i^f \triangleq E_i \setminus E_i^s$ the set of ‘free’ events.

Definition 9. Let $A \subseteq \text{Obs}$. Then $\mathcal{E}_1 \parallel_A \mathcal{E}_2 \triangleq (E, \rightsquigarrow, \mapsto, l, \mathcal{A}, \mathcal{R}, \mathcal{U})$ where

- $E = (E_1^f \times \{ * \}) \cup (\{ * \} \times E_2^f) \cup \{ (e_1, e_2) \in E_1^s \times E_2^s \mid l_1(e_1) = l_2(e_2) \}$
- $(e_1, e_2) \rightsquigarrow (e'_1, e'_2)$ iff
 - $(e_1 \rightsquigarrow_1 e'_1) \vee (e_2 \rightsquigarrow_2 e'_2)$ or
 - $(e_1 = e'_1 \neq * \wedge e_2 \neq e'_2) \vee (e_2 = e'_2 \neq * \wedge e_1 \neq e'_1)$
- $X \mapsto (e_1, e_2)$ iff
 - $(\exists X_1 : X_1 \mapsto_1 e_1 \wedge X = \{ (e, e') \in E \mid e \in X_1 \})$ or
 - $(\exists X_2 : X_2 \mapsto_2 e_2 \wedge X = \{ (e, e') \in E \mid e' \in X_2 \})$
- $l(e_1, e_2) = \text{if } e_1 = * \text{ then } l_2(e_2) \text{ else } l_1(e_1)$
- $\mathcal{A}(e_1, e_2) = \mathcal{A}_1(e_1) \cap \mathcal{A}_2(e_2)$ with $\mathcal{A}_i(*) = [0, \infty)$.
- $\mathcal{R}(X, (e_1, e_2)) = \bigcap_{X_1 \in S_1} \mathcal{R}_1(X_1, e_1) \cap \bigcap_{X_2 \in S_2} \mathcal{R}_2(X_2, e_2)$ with
 - $S_1 = \{ X_1 \subseteq E_1 \mid X_1 \mapsto_1 e_1 \wedge X = \{ (e, e') \in E \mid e \in X_1 \} \}$ and
 - $S_2 = \{ X_2 \subseteq E_2 \mid X_2 \mapsto_2 e_2 \wedge X = \{ (e, e') \in E \mid e' \in X_2 \} \}$
- $(e_1, e_2) \in \mathcal{U}$ iff $e_1 \in \mathcal{U}_1 \vee e_2 \in \mathcal{U}_2$ with $* \notin \mathcal{U}_i$.

Parallel composition is illustrated by the following example where the left-hand tes is composed with the empty tes:



For $\mathcal{E}_1 \triangleright_t \mathcal{E}_2$ a new urgent event e with delay $[t, t]$ is introduced that models the expiration of the timer. Since either the timer expires or \mathcal{E}_1 performs an initial event before (or at) t , e is put in mutual conflict with all initial events of \mathcal{E}_1 .

Definition 10. For $t \in [0, \infty)$ let $\mathcal{E}_1 \triangleright_t \mathcal{E}_2 \triangleq \mathcal{E}_1 + \hat{\tau}_{[t, t]} \cdot \mathcal{E}_2$.

By straightforward proof one can establish that TES is closed under the operators $a_I \cdot, +, \setminus A, [\lambda], ;, \triangleright, \parallel_A,$ and \triangleright_t .

5 A metric denotational semantics

The approach. We only give a brief account of our approach; see [2] for a full treatment, and [15, 5, 6] for more information on the use of metrics for denotational semantics. The semantic domain S for PA is equipped with a set Op' of operators that reflect the operators Op of Expr. For any fixed declaration $decl$, the function $P \mapsto \mathcal{M}(\langle decl, P \rangle)$ is a homomorphism from $(Expr, Op)$ to (S, Op') such that the meaning of process variable x is given by $decl(x)$. Function \mathcal{M} satisfies these conditions iff, for any fixed declaration $decl$, the function $P \mapsto \mathcal{M}(\langle decl, P \rangle)$ is a fixed point of $F_{decl} : [Expr \rightarrow S] \rightarrow [Expr \rightarrow S]$, defined (in our case) by:

$$\begin{aligned}
 F_{decl}(\phi)(\mathbf{0}) &\triangleq \mathbf{0}' \\
 F_{decl}(\phi)(\mathbf{1}) &\triangleq \mathbf{1}' \\
 F_{decl}(\phi)(x) &\triangleq \phi(decl(x)) \\
 F_{decl}(\phi)(op\ P) &\triangleq op'\ F_{decl}(\phi)(P) \quad \text{for } op \in \{a_I \cdot, \setminus A, [\lambda]\} \\
 F_{decl}(\phi)(P\ op\ Q) &\triangleq F_{decl}(\phi)(P)\ op'\ F_{decl}(\phi)(Q) \quad \text{for } op \in \{+, \parallel_A, ;, \triangleright, \triangleright_t\}.
 \end{aligned}$$

The semantics of PA is now obtained by $\mathcal{M}(\langle decl, P \rangle) = \phi_{decl}(P)$, where $\phi_{decl} : Expr \rightarrow S$ is the unique fixed point of F_{decl} . By Banach's fixpoint theorem, F_{decl} has a unique fixed point, provided that F_{decl} is contracting with respect to a distance function \tilde{d} where $([Expr \rightarrow S], \tilde{d})$ is a complete metric space (cms). \tilde{d} is obtained from the cms $\langle S, d \rangle$ where

$$\tilde{d}(\phi_1, \phi_2) = \sup \{ d(\phi_1(P), \phi_2(P)) \mid P \in Expr \}. \tag{1}$$

F_{decl} is contracting on $([Expr \rightarrow S], \tilde{d})$ if its constituents $;, \triangleright_t$ and so on, are non-distance increasing on $\langle S, d \rangle$ and contracting in certain arguments.

Time truncation. The basis of our distance function d is time truncation. The minimal time at which e can occur in \mathcal{E} is defined by $mintime_{\mathcal{E}}(e) \triangleq \inf \{ t \in \mathbb{R}^+ \mid \exists \sigma \in Traces(\mathcal{E}) : (e, t) \in \sigma \}$, where $\inf \emptyset \triangleq \infty$. For $t \in \mathbb{R}^+$ and $X \subseteq E$ let $X \upharpoonright t \triangleq \{ e \in X \mid mintime_{\mathcal{E}}(e) < t \}$ and $X \upharpoonright \omega \triangleq \bigcup_{t \geq 0} X \upharpoonright t$. Event e is called *executable* iff $e \in E \upharpoonright \omega$, i.e. if $mintime_{\mathcal{E}}(e) < \infty$.


Definition 11. The *truncation* of \mathcal{E} up to $t \in \mathbb{R}^+$ is defined by $\mathcal{E} \upharpoonright t \triangleq (E \upharpoonright t, \rightsquigarrow_t, \mapsto_t, l_t, A_t, \mathcal{R}_t, \mathcal{U}_t)$ where $l_t(e) = l(e)$, $A_t(e) = A(e) \cap [0, t)$, $\mathcal{U}_t = \mathcal{U} \upharpoonright t$, and

$$- \rightsquigarrow_t = \rightsquigarrow \cap (E \upharpoonright t \times E \upharpoonright t)$$

- $X \mapsto_t e$ iff there exists $Y \mapsto e$ with $Y \upharpoonright t = X$
- $\mathcal{R}_t(X, e) = \bigcap \{ \mathcal{R}(Y, e) \mid Y \mapsto e, Y \upharpoonright t = X \}$.

It is not difficult to check that for $\mathcal{E} \in \text{TES}$ we have $\mathcal{E} \upharpoonright t \in \text{TES}$, for all $t \in \mathbb{R}^+$.

Example 1. Time truncation is illustrated by the following figure.



Events e_b, e_f and e_g are eliminated since the minimal time at which they can occur, time 11, 8 and 6, respectively, is at least 6. Note that $\{e_a\} \xrightarrow{[1,9]}_t e_c$ for $t=6$, since $\{e_a, e_b\} \xrightarrow{1} e_c$ and $\{e_a\} \xrightarrow{[0,9]} e_c$ and $[1, \infty) \cap [0, 9] = [1, 9]$.

Lemma 12. $\text{Traces}(\mathcal{E}) = \bigcup_{t \geq 0} \text{Traces}(\mathcal{E} \upharpoonright t)$.

A complete ultra-metric space. The idea is to use time truncation as a basis for defining a distance d on TES. In particular, the distance between two tes's will be determined by the maximum amount of time they "agree", that is:

$$d(\mathcal{E}_1, \mathcal{E}_2) = \inf \{ 2^{-t} \mid \mathcal{E}_1 \upharpoonright t = \mathcal{E}_2 \upharpoonright t \}. \tag{2}$$

Remark that $\mathcal{E} \upharpoonright 0$ is the empty tes, so each pair of tes's agrees at least up to time 0. Although this basic idea is rather intuitive, it is, unfortunately, too naive. The problem is that some distinct tes's cannot be distinguished according to d . This means that d is a pseudo-metric rather than a metric. For instance, the tes consisting of a single event e with an empty bundle pointing to e is indistinguishable from the empty tes, since their time truncations are all empty. That is, according to (2) their distance is 0. The problem with this example is that tes's may contain events that can never appear. Such events can, for instance, appear in the semantics for expressions that contain circular causal dependencies, like in $a . b . \mathbf{0} \parallel_{\{a,b\}} b . a . \mathbf{0}$, or timing constraints that avoid certain actions from happening, like in $a_2 . \mathbf{0} \triangleright_1 b . \mathbf{0}$ where a will never happen. (Such events can be removed by applying the transformations exposed in [11, 9] that preserve timed event traces.)

A solution to this problem is to impose an equivalence relation, \simeq say, on TES, while aiming at $d(\mathcal{E}_1, \mathcal{E}_2) = 0 \Leftrightarrow \mathcal{E}_1 \simeq \mathcal{E}_2$. Stated in other words, where \mathbf{d} is the equivalent of d on TES/\simeq and \mathbf{E}_i denotes the equivalence class of \mathcal{E}_i under \simeq , we aim at $\mathbf{d}(\mathbf{E}_1, \mathbf{E}_2) = 0 \Leftrightarrow \mathbf{E}_1 = \mathbf{E}_2$. In order to obtain \simeq , the example suggests to abstract from events that can never be executed:

Definition 13. The *normal form* of \mathcal{E} , denoted $NF(\mathcal{E})$, is defined as $NF(\mathcal{E}) \triangleq (E \upharpoonright \omega, \rightsquigarrow_\omega, \mapsto_\omega, l_\omega, \mathcal{A}_\omega, \mathcal{R}_\omega, \mathcal{U}_\omega)$ where $l_\omega(e) = l(e)$, $\mathcal{A}_\omega(e) = \mathcal{A}(e)$, $\mathcal{U}_\omega = \mathcal{U} \cap (E \upharpoonright \omega)$, $\rightsquigarrow_\omega = \rightsquigarrow \cap (E \upharpoonright \omega \times E \upharpoonright \omega)$ and

- $X \mapsto_{\omega} e$ iff $X \upharpoonright t \mapsto e$ is a bundle in $\mathcal{E} \upharpoonright t$ for all $t > \text{mintime}_{\mathcal{E}}(e)$
- $\mathcal{R}_{\omega}(X, e) = \bigcup_{t > \text{mintime}_{\mathcal{E}}(e)} \{ \mathcal{R}(X \upharpoonright t, e) \mid X \upharpoonright t \mapsto e \text{ is a bundle of } \mathcal{E} \upharpoonright t \}$.

For $\mathcal{E} \in \text{TES}$ it follows by straightforward verification that $NF(\mathcal{E}) \in \text{TES}$.

Lemma 14. $\text{Traces}(NF(\mathcal{E})) = \text{Traces}(\mathcal{E})$.

As in the untimed case [12] the metric approach also allows to abstract from the names of the events, i.e. to deal with isomorphism classes of tes's. The names of the events are only needed for technical reasons but they are meaningless for the semantics of a PA-process. The advantage of abstraction from event names is that the definitions of operators like $+$, $[\>$, and so on, become less awkward.

Definition 15. $\mathcal{E}_i = (E_i, \sim_i, \mapsto_i, l_i, \mathcal{A}_i, \mathcal{R}_i, \mathcal{U}_i)$ for $i=1, 2$ are *isomorphic* if there exists a bijection $f : E_1 \upharpoonright \omega \longrightarrow E_2 \upharpoonright \omega$ such that $l_2 \circ f = l_1$, $\mathcal{A}_2 \circ f = \mathcal{A}_1$ and

1. $e_1 \sim_1 e_2$ iff $f(e_1) \sim_2 f(e_2)$ for all $e_1, e_2 \in E_1 \upharpoonright \omega$,
2. $X \xrightarrow{I}_1 e$ iff $f(X) \xrightarrow{I}_2 f(e)$ for all $e \in E_1 \upharpoonright \omega, X \subseteq E_1 \upharpoonright \omega$, and
3. $e \in \mathcal{U}_1 \upharpoonright \omega$ iff $f(e) \in \mathcal{U}_2$.

Let $\mathcal{E}_1 \simeq_{iso} \mathcal{E}_2$ iff there exists an isomorphism from \mathcal{E}_1 to \mathcal{E}_2 . Note that $\mathcal{E} \simeq_{iso} NF(\mathcal{E})$.

For $\mathcal{E} \in \text{TES}$ let $\mathbf{E}_{\mathcal{E}}$ denote the equivalence class of \mathcal{E} under \simeq_{iso} . For $\mathbf{E} \in \text{TES}/\simeq_{iso}$ let $\mathbf{E} \upharpoonright t \triangleq \mathbf{E}_{\mathcal{E} \upharpoonright t}$, where \mathcal{E} is a representative of \mathbf{E} . The distance between equivalence classes (under \simeq_{iso}) of tes's is given by:

$$\mathbf{d}(\mathbf{E}_1, \mathbf{E}_2) \triangleq \inf \{ 2^{-t} \mid \mathbf{E}_1 \upharpoonright t = \mathbf{E}_2 \upharpoonright t \}. \quad (3)$$

Remark that $\mathbf{d}(\mathbf{E}, \mathbf{E} \upharpoonright t) \leq 2^{-t}$ for all $t \geq 0$.

Example 2. Let $\mathcal{E}_i = (E_i, \emptyset, \mapsto_i, E_i \times \{a\}, \mathcal{A}_i, \mathcal{R}_i, \emptyset)$, for $i=1, 2$ where

- $E_1 = \{ (k, j) \mid j \geq 1, 0 < k \leq j \}$ and $E_2 = E_1 \cup \{ (k, 0) \mid k \geq 1 \}$
- $\{ (k, j) \} \mapsto_i (k+1, j)$ for $0 < k < j$ and $\{ (k, 0) \} \mapsto_2 (k+1, 0)$ for $k \geq 1$
- $\mathcal{A}_i(k, j) = [k, k]$ for all $(k, j) \in E_i$, and
- $\mathcal{R}_i(\{ (k, j) \}, (k+1, j)) = [1, 1]$.

Then, $\mathcal{E}_1 \not\simeq_{iso} \mathcal{E}_2$ while $\mathcal{E}_1 \upharpoonright t \simeq_{iso} \mathcal{E}_2 \upharpoonright t$ for all $t \geq 0$. If we now would define \mathbf{d} as suggested in (3) on TES/\simeq_{iso} then $\mathbf{d}(\mathbf{E}_1, \mathbf{E}_2) = 0$, although \mathcal{E}_1 and \mathcal{E}_2 are not isomorphic, thus yielding a pseudo-metric.

The problem with this example is that both tes's allow an infinite number of events to occur in a finite amount of time. This is avoided by considering finitely approximable tes's, a timed analogon of approximable event structures [12].

Definition 16. \mathcal{E} is called *finitely approximable* iff $E \upharpoonright t$ is finite for all $t \in \mathbb{R}^+$.

Let $\text{TES}_{fn}/\simeq_{iso}$ denote the isomorphism classes of finitely approximable tes's. The main result that we need in order to define a metric semantics for PA is:

Theorem 17. $\langle \text{TES}_{\text{fin}}/\simeq_{\text{iso}}, \mathbf{d} \rangle$ is a complete ultra-metric space.

A metric semantics for PA. We now give a metric semantics for (a subset of) PA based on equivalence classes (under \simeq_{iso}) of tes's. The main difference with the standard (untimed) case is the notion of ‘guardedness’ which ensures the well-definedness of recursive programs. While in the untimed case [12] guardedness requires that each process instantiation is preceded by an action we use a notion of *time guardedness* (like in timed CSP [17]) which requires that a recursive process instantiation can only happen after a positive amount of time. Let functions $\sqrt{\min} : \text{Expr} \rightarrow [0, \infty)$ and $\text{tg} : \text{Expr} \rightarrow [0, \infty)$ be defined as in Table 1. For declaration decl let $\text{tg}(\text{decl}) \triangleq \inf \{ \text{tg}(\text{decl}(x)) \mid x \in \text{Var} \}$. decl is called *time-guarded* iff $\text{tg}(\text{decl}) > 0$. We give a metric semantics to TGPA,

P	$\sqrt{\min}(P)$	$\text{tg}(P)$
$\mathbf{0}$	∞	∞
$\mathbf{1}$	0	∞
x	0	0
$a_I . P$	$\inf(I) + \sqrt{\min}(P)$	$\inf(I) + \text{tg}(P)$
$P[\lambda], P \setminus A$	$\sqrt{\min}(P)$	$\text{tg}(P)$
$P + Q, P [> Q$	$\min \{ \sqrt{\min}(P), \sqrt{\min}(Q) \}$	$\min \{ \text{tg}(P), \text{tg}(Q) \}$
$P \parallel_A Q$	$\max \{ \sqrt{\min}(P), \sqrt{\min}(Q) \}$	$\min \{ \text{tg}(P), \text{tg}(Q) \}$
$P ; Q$	$\sqrt{\min}(P) + \sqrt{\min}(Q)$	$\min \{ \text{tg}(P), \sqrt{\min}(P) + \text{tg}(Q) \}$
$P \triangleright_t Q$	$\min \{ \sqrt{\min}(P), t + \sqrt{\min}(Q) \}$	$\min \{ \text{tg}(P), t + \text{tg}(Q) \}$

Table 1. Auxiliary functions $\sqrt{\min}$ and tg .

the set of *time-guarded processes*, i.e. the set of pairs $\langle \text{decl}, P \rangle$ where decl is a time-guarded declaration and P an expression. For the definition of the meaning function $\mathcal{M}_{\text{cms}} : \text{TGPA} \rightarrow \text{TES}_{\text{fin}}/\simeq_{\text{iso}}$ we lift the semantic operators of Section 4 to operators on $\text{TES}_{\text{fin}}/\simeq_{\text{iso}}$. Given that all operators defined in Section 4 preserve \simeq_{iso} and finitely approximability (as can be shown by straightforward proof) we may define for $\mathbf{E}, \mathbf{F} \in \text{TES}_{\text{fin}}/\simeq_{\text{iso}}$:

$$\begin{aligned} op \mathbf{E} &\triangleq \mathbf{E}_{op} \mathcal{E} \quad \text{for } op \in \{ a_I ., \setminus A, [\lambda] \} \text{ and} \\ \mathbf{E} op \mathbf{F} &\triangleq \mathbf{E}_{\mathcal{E} op \mathcal{F}} \quad \text{for } op \in \{ +, ;, \parallel_A, [>, \triangleright_t \} \end{aligned}$$

where \mathcal{E}, \mathcal{F} are representatives of \mathbf{E} and \mathbf{F} , respectively. Let \mathbf{E}_0 be the equivalence class of the empty tes and \mathbf{E}_1 the equivalence class of the tes

$$\mathcal{E}_1 \triangleq (\{ e \}, \emptyset, \emptyset, \{ (e, \sqrt{}) \}, \{ (e, [0, \infty)) \}, \emptyset, \emptyset).$$

Together with these semantic operators, $\text{TES}_{\text{fin}}/\simeq_{\text{iso}}$ constitutes a PA-algebra.

Lemma 18. For $\mathbf{E}, \mathbf{E}', \mathbf{F}, \mathbf{F}' \in \text{TES}_{\text{fin}}/\simeq_{\text{iso}}$ we have

$$1. \mathbf{d}(a_I . \mathbf{E}, a_I . \mathbf{E}') = 2^{-\inf(I)} \cdot \mathbf{d}(\mathbf{E}, \mathbf{E}')$$

2. $\mathbf{d}(\mathbf{E} \text{ op } \mathbf{F}, \mathbf{E}' \text{ op } \mathbf{F}') \leq \max \{ \mathbf{d}(\mathbf{E}, \mathbf{E}'), \mathbf{d}(\mathbf{F}, \mathbf{F}') \}$ for $\text{op} \in \{ +, \parallel_A, [>] \}$
3. $\mathbf{d}(\text{op } \mathbf{E}, \text{op } \mathbf{E}') \leq \mathbf{d}(\mathbf{E}, \mathbf{E}')$ for $\text{op} \in \{ \setminus A, [\lambda] \}$
4. $\mathbf{d}(\mathbf{E}; \mathbf{F}, \mathbf{E}'; \mathbf{F}') \leq \max \{ \mathbf{d}(\mathbf{E}, \mathbf{E}'), 2^{-\sqrt{\min}(\mathbf{E})} \cdot \mathbf{d}(\mathbf{F}, \mathbf{F}') \}$
5. $\mathbf{d}(\mathbf{E} \triangleright_t \mathbf{F}, \mathbf{E}' \triangleright_t \mathbf{F}') \leq \max \{ \mathbf{d}(\mathbf{E}, \mathbf{E}'), 2^{-t} \cdot \mathbf{d}(\mathbf{F}, \mathbf{F}') \}$.

where $\sqrt{\min}(\mathbf{E}\mathcal{E}) \triangleq \inf \{ \text{mintime}_{\mathcal{E}}(e) \mid e \in E, l(e) = \sqrt{\cdot} \}$.

Lemma 19. For each decl and homomorphisms $\phi_1, \phi_2 : \text{Expr} \longrightarrow \text{TES}_{\text{fin}}/\simeq_{\text{iso}}$:

$$\tilde{d}(F_{\text{decl}}(\phi_1), F_{\text{decl}}(\phi_2)) \leq 2^{-\text{tg}(\text{decl})} \cdot \tilde{d}(\phi_1, \phi_2).$$

where \tilde{d} is defined in equation (1).

This result says that F_{decl} is contracting with contraction coefficient $2^{-\text{tg}(\text{decl})}$ provided that decl is time-guarded. Thus, for time-guarded declaration decl , F_{decl} has a unique fixed point, say ϕ_{decl} . The metric semantics $\mathcal{M}_{\text{cms}} : \text{TGPA} \longrightarrow \text{TES}_{\text{fin}}/\simeq_{\text{iso}}$ is now defined by $\mathcal{M}_{\text{cms}}(\langle \text{decl}, P \rangle) \triangleq \phi_{\text{decl}}(P)$.

6 Concluding remarks

Relation with untimed case. This paper defines a metric semantics \mathcal{M}_{cms} for an expressive real-time process algebra PA that contains delay and timeout operators. The distance d measures the amount of time in which two processes coincide, i.e., $d(\mathcal{M}_{\text{cms}}(P_1), \mathcal{M}_{\text{cms}}(P_2)) \leq 2^{-t}$ iff P_1 and P_2 have the same behaviour upto time t . This notion of distance is a timed analogon of the distance proposed in [12] which is based on the number of steps processes coincide.

Consistency. [10] defines a cpo-based denotational semantics \mathcal{M}_{cpo} and an (event-based) operational semantics for PA such that $\text{Traces} \circ \mathcal{M}_{\text{cpo}}$ are precisely the timed event traces that are generated operationally. The formal relationship between our cpo and metric semantics is as follows. Let TES_{fin} be the set of timed event structures that are finitely approximable. For time-guarded $\langle \text{decl}, P \rangle$ it follows that $\mathcal{M}_{\text{cpo}}(\langle \text{decl}, P \rangle)$ is finitely approximable. Function $f : \text{TES}_{\text{fin}} \longrightarrow \text{TES}_{\text{fin}}/\simeq_{\text{iso}}$ with $f(\mathcal{E}) \triangleq \mathbf{E}_{\mathcal{E}}$ is a homomorphism between the PA-algebras TES_{fin} and $\text{TES}_{\text{fin}}/\simeq_{\text{iso}}$. Then, according to the results of [3], we obtain for any time-guarded process $\langle \text{decl}, P \rangle$: $f(\mathcal{M}_{\text{cpo}}(\langle \text{decl}, P \rangle)) = \mathcal{M}_{\text{cms}}(\langle \text{decl}, P \rangle)$. This entails that the presented metric semantics is significantly more abstract than the cpo-based, and consequently, the operational semantics of TGPA.

Related work. Several real-time extensions of process algebras have been proposed in the literature; for an overview see [14]. Usually, timed process algebras are provided with an operational semantics in the style of Plotkin that is based on some notion of (timed) transition system. Notably exceptions are the works on timed CSP by Reed & Roscoe [17] who use a metric denotational semantics based on timed refusals, and real-time LOTOS by Bryans, Davies & Schneider [8] who use a (non-standard) fixed point semantics based on an advanced form of timed refusals in order to deal with divergence. Both works provide an interleaving semantics. In the non-interleaving setting, related work has been done by

Murphy [13] on interval event structures in which events have a duration. Murphy uses time truncation—in a similar way as we do in the metric semantics—as a basis for obtaining limiting infinite objects using ideal completions.

References

1. A.F. Ates, M. Bilgic, S. Saito and B. Sarikaya. Using timed CSP for specification verification and analysis of multi-media synchronization. *IEEE J. on Sel. Areas in Comm.*, **14**(1):126–137, 1996.
2. C. Baier and M.E. Majster-Cederbaum. Denotational semantics in the cpo and metric approach. *Th. Comp. Sci.*, **135**:171–220, 1994.
3. C. Baier and M.E. Majster-Cederbaum. How to interpret consistency and establish consistency results for semantics of concurrent programming languages. *Fund. Inf.*, **29**:225–256, 1997.
4. C. Baier and M.E. Majster-Cederbaum. Metric semantics from partial order semantics. *Acta Inf.*, **34**:701–735, 1997.
5. J.W. de Bakker and J.I. Zucker. Processes and the denotational semantics of concurrency. *Inf. and Contr.*, **54**(1/2):70–120, 1982.
6. J.W. de Bakker and E.P. de Vink. *Control Flow Semantics*. MIT Press, 1996.
7. T. Bolognesi and E. Brinksma. Introduction to the ISO specification language LOTOS. *Comp. Netw. & ISDN Syst.*, **14**:25–59, 1987.
8. J. Davies, J.W. Bryans and S.A. Schneider. Real-time LOTOS and timed observations. In *Formal Description Techniques VIII*. Chapman & Hall, 1995.
9. J-P. Katoen. *Quantitative and Qualitative Extensions of Event Structures*. PhD thesis, University of Twente, 1996.
10. J-P. Katoen, D. Latella, R. Langerak and E. Brinksma. On specifying real-time systems in a causality-based setting. In *Formal Techniques in Real-Time and Fault-Tolerant Systems*, LNCS 1135, pages 385–405. Springer-Verlag, 1996.
11. R. Langerak. Bundle event structures: a non-interleaving semantics for LOTOS. In *Formal Description Techniques V*, pages 331–346. North-Holland, 1993.
12. R. Loogen and U. Goltz. Modelling nondeterministic concurrent processes with event structures. *Fund. Inf.*, **14**(1):39–74, 1991.
13. D. Murphy. Time and duration in noninterleaving concurrency. *Fund. Inf.*, **19**:403–416, 1993.
14. X. Nicollin and J. Sifakis. An overview and synthesis on timed process algebras. In *Real-Time: Theory in Practice*, LNCS 600, pages 526–548. Springer-Verlag, 1992.
15. M. Nivat. Infinite words, infinite trees, infinite computations. In *Foundations of Computer Science III*, Mathematical Centre Tracts **109**, pages 3–52, 1979.
16. M. Nielsen, G.D. Plotkin and G. Winskel. Petri nets, event structures and domains, part 1. *Th. Comp. Sc.*, **13**(1):85–108, 1981.
17. G.M. Reed and A.W. Roscoe. A timed model for Communicating Sequential Processes. *Th. Comp. Sc.*, **58**:249–261, 1988.
18. J.J. Žic. Time-constrained buffer specifications in CSP+T and timed CSP. *ACM Transactions on Programming Languages and Systems*, **16**(6):1661–1674, 1994.