

Revisiting Timed Specification Theories: A Linear-Time Perspective

Chris Chilton, Marta Kwiatkowska, and Xu Wang

Department of Computer Science, University of Oxford, UK

Abstract. We consider the setting of component-based design for real-time systems with critical timing constraints. Based on our earlier work, we propose a compositional specification theory for timed automata with I/O distinction, which supports substitutive refinement. Our theory provides the operations of parallel composition for composing components at run-time, logical conjunction/disjunction for independent development, and quotient for incremental synthesis. The key novelty of our timed theory lies in a weakest congruence preserving safety as well as bounded liveness properties. We show that the congruence can be characterised by two linear-time semantics, *timed-traces* and *timed-strategies*, the latter of which is derived from a game-based interpretation of timed interaction.

1 Introduction

Component-based design methodologies can be encapsulated in the form of compositional specification theories, which allow the mixing of specifications and implementations, admit substitutive refinement to facilitate reuse, and provide a rich collection of operators. Previously [1], we developed a linear-time specification theory for reasoning about untimed components that interact by synchronisation of input and output (I/O) actions, inspired by interface automata [2]. Models can be specified operationally by means of transition systems augmented by an inconsistency predicate on states, or declaratively using traces. The theory admits non-determinism, a refinement preorder based on traces, and the operations of parallel composition, conjunction and quotient. The refinement is strictly weaker than alternating simulation and is actually the weakest pre-congruence preserving inconsistent states. This implies that our refinement is *substitutive*, meaning component A refines component B iff A can replace B in any environmental context without introducing additional errors.

In this paper we target component-based development for real-time systems with critical timing constraints. We formulate a timed extension of the linear-time specification theory of [1], by allowing for both operational descriptions of components, as well as declarative specifications based on traces. Our operational models are based on a variant of timed automata with I/O distinction (although we do not insist on input-enabledness, cf [3]), augmented by two special states: \perp for safety and bounded-liveness errors, and \top for timestep. Trace-based declarative specifications are shown to be a suitable semantic domain for the operational models. In addition to timed-trace semantics, we present timed-strategy semantics, which coincides with the former but relates our work closer

to the timed-game frameworks used by [4] and [5]. The *substitutive refinement* of our framework gives rise to the weakest congruence preserving \perp , and is shown to coincide across all our formalisms.

Amongst notable works in the literature, we briefly mention a theory of timed interfaces [5] and a theory of timed specifications [4]. Timed interface theory contributes a framework based on timed games to formalise notions such as interfaces and compatibility, and also provides a parallel composition operator. However, the work cannot be considered a specification theory as it does not deal with the notion of refinement for component substitution or the operations of conjunction, disjunction and quotient. In this respect, [4] provides a complete theory; however, the refinement is a timed version of the alternating simulation originally defined for interface automata [2]. Consequently, it is too strong for determining when a component can be safely substituted with another (cf the example in Figure 3).

Outline. In Section 2 we introduce timed I/O automata, their semantic mapping to timed I/O transition systems, and supply the operational definitions for the operations of parallel composition, conjunction, disjunction and quotient. In Section 3 we use the timed-game framework to introduce timed-strategy semantics, which we relate to the operational framework. Similarly in Section 4, we present timed-trace semantics and relate these to the operational definitions. Section 5 discusses related work, and finally Section 6 concludes.

2 Formal Framework

In this section we introduce timed I/O automata, timed I/O transition systems and a semantic mapping from the former to the latter. Timed I/O automata are compact representations of timed I/O transition systems. Our theory will be developed using timed I/O transition systems, which are endowed with a richer repertoire of semantic machinery.

2.1 Timed I/O Automata

Clock constraints. Given a set X of real-valued clock variables, a *clock constraint* over X , $cc : CC(X)$, is a boolean combination of atomic constraints of the form $x \bowtie d$ and $x - y \bowtie d$ where $x, y \in X$, $\bowtie \in \{\leq, <, =, >, \geq\}$, and $d \in \mathbb{N}$.

A *clock valuation* over X is a map t that assigns to each clock variable x in X a real value from $\mathbb{R}^{\geq 0}$. We say t satisfies cc , written $t \in cc$, if cc evaluates to true under valuation t . $t + d$ denotes the valuation derived from t by increasing the assigned value on each clock variable by $d \in \mathbb{R}^{\geq 0}$ time units. $t[rs \mapsto 0]$ denotes the valuation obtained from t by resetting the clock variables in rs to 0. Sometimes we use 0 for the clock valuation that maps all clock variables to 0.

Definition 1. A timed I/O automaton (*TIOA*) is a tuple $(C, I, O, L, l^0, AT, Inv, coInv)$, where:

- $C \subseteq X$ is a finite set of clock variables
- $A (= I \cup O)$ is a finite alphabet, where I and O are disjoint sets of input actions and output actions respectively
- L is a finite set of locations
- $l^0 \in L$ is the initial location
- $AT \subseteq L \times CC(C) \times A \times 2^C \times L$ is a set of action transitions
- $Inv : L \rightarrow CC(C)$ and $coInv : L \rightarrow CC(C)$ assign invariants and co-invariants to states, each of which is a downward-closed clock constraint.

We use l, l', l_i to range over L and use $l \xrightarrow{g, a, rs} l'$ as a shorthand for $(l, g, a, rs, l') \in AT$. $g : CC(C)$ is the enabling guard of the transition, $a \in A$ the action, and rs the subset of clock variables to be reset.

Our TIOAs are similar to existing variants of timed automata with input/output distinction, except for the introduction of co-invariants and non-insistence on input-enabledness. While invariants specify the bounds beyond which time may not progress, co-invariants specify the bounds beyond which the system will *time-out* and enter error states. Our TIOAs can be used to describe both the assumptions made by the component on the inputs, together with the guarantees provided by the component on the outputs. Such assumptions and guarantees can be time constrained: guards on output transitions express safety timing guarantees, while guards on input transitions express safety timing assumptions; invariants (urgency) express liveness timing guarantees on outputs while co-invariants (time-out) express liveness timing assumptions on inputs.

When components interact together, we check whether the guarantees they provide meet the assumptions they make on each other. If not, there are two types of errors:

- An input arrives in a state and at a time when it is not expected (i.e. not satisfying the guards on the input transitions). This is a *safety error*.
- An input does not arrive in a state within a time bound (specified by a co-invariant) as expected. This is a *bounded-liveness error*.

Example. Figure 1 depicts TIOAs representing a job scheduler together with a printer controller. The invariant at location A of the scheduler forces a bounded-liveness guarantee on outputs in that location. As time must be allowed to progress beyond $t = 100$, the *start* action must be fired within the range $0 \leq t \leq 100$. After *start* has been fired, the clock x is reset to 0 and the scheduler waits (possibly indefinitely) for the job to *finish*. If the job does finish, the scheduler is only willing for this to take place between $5 \leq t \leq 8$ after the job started (safety assumption), otherwise an unexpected input error will be thrown.

The controller waits for the job to *start*, after which it will wait exactly 1 time unit before issuing *print* (forced by the invariant $y \leq 1$ on state 2 and the guard $y = 1$). The controller now requires the printer to indicate the job is *printed* within 10 time units of being sent to the printer, otherwise a time-out error on inputs will occur (co-invariant $y \leq 10$ in state 3 as liveness assumption). After the job has finished printing, the controller must indicate to the scheduler that the job has *finished* within 5 time units.




Fig. 1. Job scheduler and printer controller.

2.2 Timed Actions and Words

In this section we introduce some notation relating to timed actions and timed words that will be of use to us in later sections.

Timed actions. For a set of input actions I and a set of output actions O , define $tA = I \uplus O \uplus \mathbb{R}^{>0}$ to be the set of *timed actions*, $tI = I \uplus \mathbb{R}^{>0}$ to be the set of *timed inputs*, and $tO = O \uplus \mathbb{R}^{>0}$ to be the set of *timed outputs*. We use symbols like α, β , etc. to range over tA .

Timed words. A *timed word* (ranged over by w, w', w_i etc.) is a finite mixed sequence of positive real numbers ($\mathbb{R}^{>0}$) and visible actions such that *no two numbers are adjacent to one another*. For instance, $\langle 0.33, a, 1.41, b, c, 3.1415 \rangle$ is a timed word denoting the observation that action a occurs at 0.33 time units, then another 1.41 time units lapse before the simultaneous occurrence of b and c , which is followed by 3.1415 time units of no event occurrence. The empty word is denoted by ϵ .

Operations on timed words. We use $\text{last}(w)$ to denote the last element in the sequence w , and $l(w)$ to indicate the length, which is obtained as the sum of all the reals in w . Concatenation of timed words w and w' is obtained by appending w' onto the end of w and coalescing adjacent reals (summing them). For instance, $\langle a, 1.41 \rangle \frown \langle 0.33, b, 3.1415 \rangle = \langle a, (1.41 + 0.33), b, 3.1415 \rangle = \langle a, 1.74, b, 3.1415 \rangle$. Prefix/extension are defined as usual by concatenation, and we use \leq for the prefix partial order. We write $w \upharpoonright tA_0$ for the projection of w onto timed alphabet tA_0 , which is defined by removing from w all actions not inside tA_0 and coalescing adjacent reals.

2.3 Semantics as Timed I/O Transition Systems

The semantics of TIOAs are given as timed I/O transition systems, which are a special class of infinite labelled transition systems.

Definition 2. A timed I/O transition system (TIOTS) is a tuple $\mathcal{P} = \langle I, O, S, s^0, \rightarrow \rangle$, where: I and O are the input and output actions respectively, S is a set of states, s^0 is the designated initial state, and $\rightarrow \subseteq S \times I \uplus O \uplus \mathbb{R}^{>0} \times S$ is the action and time-labelled transition system.

The states of the TIOTS for a TIOA capture the configurations of the automaton, i.e. its location and clock valuation. Therefore, each state of the TIOTS is a pair drawn from $L \times \mathbb{R}^C$, which we refer to as the set of *plain states*, denoted P . In addition, we introduce two special states \perp and \top , which are required for the semantic mapping of disabled inputs/outputs, invariants and co-invariants.

\perp is called the *inconsistent state*, representing safety and bounded-liveness errors. \top is the so-called *timestop state*, representing the *magic moment* from which no error can occur.¹

An intuitive way to understand \top and \perp is from an input/output game perspective. The component controls output and delay while the environment controls input. \perp is the losing state for the environment. So a disabled input at a state p is equated to an input transition from p to \perp . \top is the losing state for the component. So a disabled output/delay at p is equated to an output/delay transition from p to \top . Thus we can have two semantics-preserving transformations on TIOTSs.

The \perp -completion of a TIOTS \mathcal{P} , denoted \mathcal{P}^\perp , adds an a -labelled transition from p to \perp for every $p \in P_{\mathcal{P}}$ and $a \in I$ s.t. a is not enabled at p . \perp -completion will make a TIOTS *input-receptive*, i.e. input-enabled at all states. The \top -completion of a TIOTS \mathcal{P} , denoted \mathcal{P}^\top , adds an α -labelled transition from p to \top for every $p \in P_{\mathcal{P}}$ and $\alpha \in tO$ s.t. α is not enabled at p .

Furthermore, for technical convenience (e.g. ease of defining time additivity), the definition of TIOTSs requires that 1) \top is a *quiescent state*, i.e. a state in which the set of outgoing transitions are all self-loops, one for each $d \in \mathbb{R}^{>0}$, and 2) \perp is a *chaotic state*, i.e. a state in which the set of outgoing transitions are all self-loops, one for each $\alpha \in tA$. The set of all possible states is denoted $S = P \uplus \{\perp, \top\}$. We use p, p', p_i to range over P while s, s', s_i range over S .

The transition relation \rightarrow of the TIOTS is derived from the execution semantics of the TIOA.

Definition 3. Let \mathcal{P} be a TIOA. The semantic mapping of \mathcal{P} is a TIOTS $\langle I, O, S, s^0, \rightarrow \rangle$, where:

- $S = (L \times \mathbb{R}^C) \uplus \{\perp, \top\}$
- $s^0 = \top$ providing $0 \notin \text{Inv}(l^0)$, $s^0 = \perp$ providing $0 \in \text{Inv}(l^0) \wedge \neg \text{coInv}(l^0)$ and $s^0 = (l^0, 0)$ providing $0 \in \text{Inv}(l^0) \wedge \text{coInv}(l^0)$,
- \rightarrow is the smallest relation satisfying:
 1. If $l \xrightarrow{g, a, rs} l'$, $t' = t[rs \mapsto 0]$, $t \in \text{Inv}(l) \wedge \text{coInv}(l) \wedge g$, then:
 - (a) plain action: $(l, t) \xrightarrow{a} (l', t')$ providing $t' \in \text{Inv}(l') \wedge \text{coInv}(l')$
 - (b) error action: $(l, t) \xrightarrow{a} \perp$ providing $t' \in \text{Inv}(l') \wedge \neg \text{coInv}(l')$
 - (c) magic action: $(l, t) \xrightarrow{a} \top$ providing $t' \in \neg \text{Inv}(l')$.

¹ For instance, a location with *true* as co-invariant and *false* as invariant is mapped to \top , while a location with *true* as invariant and *false* as co-invariant is mapped to \perp . A location with *false* for both invariant and co-invariant is mapped to \top since invariants have priority over co-invariants according to our semantics; whereas a location with $x \leq 0$ as invariant and *true* as co-invariant is mapped to a plain state.

2. plain delay: $(l, t) \xrightarrow{d} (l, t + d)$ if $t, t + d \in \text{Inv}(l) \wedge \text{coInv}(l)$
3. time-out delay: $(l, t) \xrightarrow{d} \perp$ if $t \in \text{Inv}(l) \wedge \text{coInv}(l)$, $t + d \notin \text{coInv}(l)$ and $\exists 0 < \delta \leq d : t + \delta \in \text{Inv}(l) \wedge \neg \text{coInv}(l)$.

Note that our semantics tries to minimise the use of transitions leading to \top/\perp states. Thus there are no delay transitions leading to \top . This creates implicit timesteps, which we capture using the concept of *semi-timestep* (i.e. semi- \top). We say a plain state p is a *semi- \top* iff 1) all output transitions enabled in p or any of its time-passing successors lead to the \top state, and 2) there exists $d \in \mathbb{R}^{>0}$ s.t. $p \xrightarrow{d} \top$ or d is not enabled in p . Thus a semi- \top is a state in which it is impossible for the component to avoid the timestep without suitable inputs from the environment.

TIOTS terminology. A TIOTS is *time additive* providing $p \xrightarrow{d_1+d_2} s'$ iff $p \xrightarrow{d_1} s$ and $s \xrightarrow{d_2} s'$ for some s . In the sequel of this paper we only consider TIOTSs that are time-additive.

We say a TIOTS is *deterministic* iff there is no ambiguous transition in the TIOTS, i.e. $s \xrightarrow{\alpha} s' \wedge s \xrightarrow{\alpha} s''$ implies $s' = s''$.

Given a TIOTS \mathcal{P} , a timed word can be derived from a finite execution of \mathcal{P} by extracting the labels in each transition and coalescing adjacent reals. The timed words derived from such executions are called *traces* of \mathcal{P} . We use tt, tt', tt_i to range over the set of traces and use $s^0 \xRightarrow{tt} s$ to denote a finite execution that produces trace tt and leads to s .

2.4 Operational Specification Theory

In this section we develop a compositional specification theory for TIOTSs based on the operations of parallel composition \parallel , conjunction \wedge , disjunction \vee and quotient $\%$. The operators are defined via transition rules that are a variant on synchronised product.

Parallel composition yields a TIOTS that represents the combined effect of its operands interacting with one another. The remaining operations must be explained with respect to a refinement relation, which corresponds to safe-substitutivity in our theory. A TIOTS is a refinement of another if it will work in any environment that the original worked in without introducing safety or bounded-liveness errors. Conjunction yields the coarsest TIOTS that is a refinement of its operands, while disjunction yields the finest TIOTS that is refined by both of its operands. The operators are thus equivalent to the join and meet operations on TIOTSs². Quotient is the adjoint of parallel composition, meaning that $\mathcal{P}_0 \% \mathcal{P}_1$ is the coarsest TIOTS such that $(\mathcal{P}_0 \% \mathcal{P}_1) \parallel \mathcal{P}_1$ is a refinement of \mathcal{P}_0 .

Let $\mathcal{P}_i = \langle I_i, O_i, S_i, s_i^0, \rightarrow_i \rangle$ for $i \in \{0, 1\}$ be two TIOTSs that are both \perp and \top -completed, satisfying (wlog) $S_0 \cap S_1 = \{\perp, \top\}$. The composition of \mathcal{P}_0

² As we write $A \sqsubseteq B$ to mean A is refined by B , our operators \wedge and \vee are reversed in comparison to the standard symbols for meet and join.

Table 1. State representations under composition operators.

\parallel	\top	p_0	\perp
\top	\top	\top	\top
p_1	\top	$p_0 \times p_1$	\perp
\perp	\top	\perp	\perp

\wedge	\top	p_0	\perp
\top	\top	\top	\top
p_1	\top	$p_0 \times p_1$	p_1
\perp	\top	p_0	\perp

\vee	\top	p_0	\perp
\top	\top	p_0	\perp
p_1	p_1	$p_0 \times p_1$	\perp
\perp	\perp	\perp	\perp

$\%$	\top	p_0	\perp
\top	\perp	\perp	\perp
p_1	\top	$p_0 \times p_1$	\perp
\perp	\top	\top	\perp

and \mathcal{P}_1 under the operation $\otimes \in \{\parallel, \wedge, \vee, \%\}$, written $\mathcal{P}_0 \otimes \mathcal{P}_1$, is only defined when certain *composability* restrictions are imposed on the alphabets of the TIOTSs. $\mathcal{P}_0 \parallel \mathcal{P}_1$ is only defined when the output sets of \mathcal{P}_0 and \mathcal{P}_1 are disjoint, because an output should be controlled by at most one component. Conjunction and disjunction are defined only when the TIOTSs have *identical alphabets* (i.e. $O_0 = O_1$ and $I_0 = I_1$). This restriction can be relaxed at the expense of more cumbersome notation, which is why we focus on the simpler case in this paper. For the quotient, we require that the alphabet of \mathcal{P}_0 *dominates* that of \mathcal{P}_1 (i.e. $A_1 \subseteq A_0$ and $O_1 \subseteq O_0$), in addition to \mathcal{P}_1 being a deterministic TIOTS. As quotient is a synthesis operator, it is difficult to give a definition using just *state-local* transition rules, since quotient needs global information of the transition systems. This is why we insist on \mathcal{P}_1 being deterministic³.

Definition 4. Let \mathcal{P}_0 and \mathcal{P}_1 be TIOTSs composable under $\otimes \in \{\parallel, \wedge, \vee, \%\}$. Then $\mathcal{P}_0 \otimes \mathcal{P}_1 = \langle I, O, S, s^0, \rightarrow \rangle$ is the TIOTS where:

- If $\otimes = \parallel$, then $I = (I_0 \cup I_1) \setminus O$ and $O = O_0 \cup O_1$
- If $\otimes \in \{\wedge, \vee\}$, then $I = I_0 = I_1$ and $O = O_0 = O_1$
- If $\otimes = \%$, then $I = I_0 \cup O_1$ and $O = O_0 \setminus O_1$
- $S = P_0 \times P_1 \uplus P_0 \uplus P_1 \uplus \{\top, \perp\}$
- $s^0 = s_0^0 \otimes s_1^0$
- \rightarrow is the smallest relation containing $\rightarrow_0 \cup \rightarrow_1$, and satisfying the rules:

$$\frac{p_0 \xrightarrow{\alpha}_0 s'_0 \quad p_1 \xrightarrow{\alpha}_1 s'_1}{p_0 \otimes p_1 \xrightarrow{\alpha} s'_0 \otimes s'_1} \quad \frac{p_0 \xrightarrow{a}_0 s'_0 \quad a \notin A_1}{p_0 \otimes p_1 \xrightarrow{a} s'_0 \otimes p_1} \quad \frac{p_1 \xrightarrow{a}_0 s'_1 \quad a \notin A_0}{p_0 \otimes p_1 \xrightarrow{a} p_0 \otimes s'_1}$$

We adopt the notation of $s_0 \otimes s_1$ for states, where the associated interpretation is supplied in Table 1. Furthermore, given two plain states $p_i = (l_i, t_i)$ for $i \in \{0, 1\}$, we define $p_0 \times p_1 = ((l_0, l_1), t_0 \uplus t_1)$.

Table 1 tells us how states should be combined under the composition operators. From the environment's point of view, \top refines plain states, which in turn refines \perp . For parallel, a state is magic if one component state is magic, and a state is error if one component is error while the other is not magic. For conjunction, encountering error in one component implies the component can be discarded and the rest of the composition behaves like the other component. The conjunction table follows the intuition of the join operation on the refinement

³ Technically speaking, the problem lies in that state quotient operator is right-distributive but not left-distributive over state disjunction (cf Table 1).

preorder. Similarly for disjunction. Quotient is the adjoint of parallel composition. If the second component state does not refine the first, the quotient will try to rescue the refinement by producing \top (so that its composition with the second will refine the first). If the second component state does refine the first, the quotient will produce the least refined value so that its composition with the second will not break the refinement.

An *environment* for a TIOTS \mathcal{P} is any TIOTS \mathcal{Q} such that the alphabet of \mathcal{Q} is *complementary* to that of \mathcal{P} , meaning $I_{\mathcal{P}} = O_{\mathcal{Q}}$ and $O_{\mathcal{P}} = I_{\mathcal{Q}}$. Refinement in our framework corresponds to contextual substitutability, in which the context is an arbitrary environment.

Definition 5. Let \mathcal{P}_{imp} and \mathcal{P}_{spec} be TIOTSs with identical alphabets. \mathcal{P}_{imp} refines \mathcal{P}_{spec} , denoted $\mathcal{P}_{spec} \sqsubseteq \mathcal{P}_{imp}$, iff for all environments \mathcal{Q} , $\mathcal{P}_{spec} \parallel \mathcal{Q}$ is \perp -free implies $\mathcal{P}_{imp} \parallel \mathcal{Q}$ is \perp -free. We say \mathcal{P}_{imp} and \mathcal{P}_{spec} are substitutively equivalent, i.e. $\mathcal{P}_{spec} \simeq \mathcal{P}_{imp}$, iff $\mathcal{P}_{imp} \sqsubseteq \mathcal{P}_{spec}$ and $\mathcal{P}_{spec} \sqsubseteq \mathcal{P}_{imp}$.

It is obvious that \simeq induces an equivalence on TIOTSs and no equivalence that preserves the \perp state can be weaker than \simeq . In the sequel we will give two concrete characterisations of \simeq and show that \simeq is also a congruence w.r.t. the parallel composition, conjunction, disjunction and quotient operators.

The operational definition of quotient requires that \mathcal{P}_1 is determinised, which can be accomplished by a modified subset construction procedure on $(\mathcal{P}_1^\perp)^\top$. If the current state subset S_0 contains \perp , it reduces S_0 to \perp ; if $\perp \notin S_0 \neq \{\top\}$, it reduces S_0 by removing any potential \top in S_0 . As expected, the determinisation of \mathcal{P} , denoted \mathcal{P}^D , is substitutively equivalent to \mathcal{P} .

Proposition 1. Any TIOTS is substitutively equivalent to a deterministic TIOTS.

Equipped with determinisation, quotient is a fully defined operator on any pair of TIOTSs. Furthermore, we can give an alternative (although substitutively equivalent) formulation of quotient as the derived operator $(\mathcal{P}_0^\top \parallel \mathcal{P}_1)^\neg$, where \neg is a mirroring operation that first determinises its argument, then interchanges the input and output sets, as well as the \top and \perp states.

Example. Figure 2 shows the parallel composition of the job scheduler with the printer controller. In the transition from $B4$ to $A1$, the guard combines the effects of the constraints on the clocks x and y . As *finish* is an output of the controller, it can be fired at a time when the scheduler is not expecting it, meaning that a safety error will occur. This is indicated by the transition to \perp when the guard constraint $5 \leq x \leq 8$ is not satisfied.

3 Timed I/O Game

Our specification theory can be understood from a game theoretical point of view. It is an input-output game between a *component* and an *environment* that

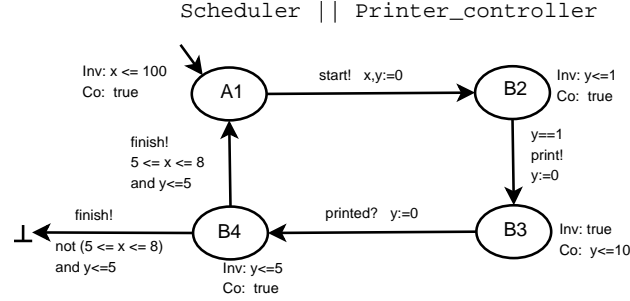


Fig. 2. Parallel composition of the job scheduler and printer controller.

uses a *coin* to break ties. The specification of a component (in the form of a TIOA or TIOTS) is built to encode the set of strategies possible for the component in the game (just like an NFA encodes a set of words).

- Given two TIOTSs \mathcal{P} and \mathcal{Q} with identical alphabets, we say \mathcal{P} is a partial unfolding [6] of \mathcal{Q} if there exists a function f from $S_{\mathcal{P}}$ to $S_{\mathcal{Q}}$ s.t. 1) f maps \top to \top , \perp to \perp , and plain states to plain states, 2) $f(s_{\mathcal{P}}^0) = s_{\mathcal{Q}}^0$, and 3) $p \xrightarrow{\alpha}_{\mathcal{P}} s \Rightarrow f(p) \xrightarrow{\alpha}_{\mathcal{Q}} f(s)$.
- We say an acyclic TIOTS is a *tree* if 1) there does not exist a pair of transitions in the form of $p \xrightarrow{a} p''$ and $p' \xrightarrow{d} p''$, 2) $p \xrightarrow{a} p'' \wedge p' \xrightarrow{b} p''$ implies $p = p'$ and $a = b$ and 3) $p \xrightarrow{d} p'' \wedge p' \xrightarrow{d} p''$ implies $p = p'$.
- We say an acyclic TIOTS is a *simple path* if 1) $p \xrightarrow{a} s' \wedge p \xrightarrow{\alpha} s''$ implies $s' = s''$ and $a = \alpha$ and 2) $p \xrightarrow{d} s' \wedge p \xrightarrow{d} s''$ implies $s' = s''$.
- We say a simple path \mathcal{L} is a *run* of \mathcal{P} if \mathcal{L} is a partial unfolding of \mathcal{P} .

Strategies. A *strategy* \mathcal{G} is a deterministic tree TIOTS s.t. each plain state in \mathcal{G} is ready to accept all possible inputs by the environment, but allows a single move (delay or output) by the component, i.e. $eb_{\mathcal{G}}(p) = I \uplus mv_{\mathcal{G}}(p)$ s.t. $mv_{\mathcal{G}}(p) = \{a\}$ for some $a \in O$ or $\{\} \subset mv_{\mathcal{G}}(p) \subseteq \mathbb{R}^{>0}$, where $eb_{\mathcal{G}}(p)$ denotes the set of enabled timed actions in state p of LTS \mathcal{G} , and $mv_{\mathcal{G}}(p)$ denotes the unique component move allowed by \mathcal{G} at p .

A TIOTS \mathcal{P} *contains* a strategy \mathcal{G} if \mathcal{G} is a partial unfolding of $(\mathcal{P}^{\perp})^{\top}$. The set of strategies⁴ contained in \mathcal{P} is denoted $stg(\mathcal{P})$. Since it makes little sense to distinguish strategies that are isomorphic, we will freely use strategies to refer to their isomorphism classes and write $\mathcal{G} = \mathcal{G}'$ to mean \mathcal{G} and \mathcal{G}' are isomorphic.

Let us give some examples in Figure 3. For the sake of simplicity we use two untimed transition systems \mathcal{P} and \mathcal{Q} , which have identical alphabets $I = \{e, f\}$ and $O = \{a, b, c\}$, to illustrate the idea of strategies. The transition systems use solid lines while strategies use dotted lines. Plain states are unmarked while the \top and \perp states are marked by \top and \perp resp.⁵ We show four strategies of \mathcal{P} and

⁴ In this paper we use a set of strategies (say Π) to mean a set of strategies with identical alphabets

⁵ To simplify drawing, multiple copies of \top and \perp are allowed but the self-loops on them are omitted.




Fig. 3. Strategy example.

two strategies of \mathcal{Q} on the right hand side of \mathcal{P} and \mathcal{Q} resp. in Figure 3. (They are not the complete sets of strategies for \mathcal{P} and \mathcal{Q} .) Note that the strategies 3 and 4 own their existence to the \top completion.

Comparing strategies. When the game is played, the component tries to avoid reaching \top while the environment tries to avoid reaching \perp . Different strategies in $stg(\mathcal{P})$ vary in their effectiveness to achieve the objective. Such effectiveness can be compared if two strategies closely resemble each other: we say \mathcal{G} and \mathcal{G}' are *affine* if $s_{\mathcal{G}}^0 \xrightarrow{tt} p$ and $s_{\mathcal{G}'}^0 \xrightarrow{tt} p'$ implies $mv_{\mathcal{G}}(p) = mv_{\mathcal{G}'}(p')$. Intuitively, it means \mathcal{G} and \mathcal{G}' propose the same move at the ‘same’ states. For instance, the strategies 1, 3 and A in Figure 3 are pairwise affine and so are the strategies 2, 4 and B .

Given two affine strategies \mathcal{G} and \mathcal{G}' , we say \mathcal{G} is *more aggressive* than \mathcal{G}' , denoted $\mathcal{G} \preceq \mathcal{G}'$, if 1) $s_{\mathcal{G}'}^0 \xrightarrow{tt} \perp$ implies there is a prefix tt_0 of tt s.t. $s_{\mathcal{G}}^0 \xrightarrow{tt_0} \perp$ and 2) $s_{\mathcal{G}}^0 \xrightarrow{tt} \top$ implies there is a prefix tt_0 of tt s.t. $s_{\mathcal{G}'}^0 \xrightarrow{tt_0} \top$. Intuitively, it means \mathcal{G} can reach \perp faster but \top slower than \mathcal{G}' . \preceq forms a partial order over $stg(\mathcal{P})$, or more generally, over any set of strategies with identical alphabets. For instance, strategy A is more aggressive than 1 and 3, while strategy B is more aggressive than 2 and 4.

When the game is played, the component \mathcal{P} prefers to use the maximally aggressive strategies in $stg(\mathcal{P})$ ⁶. Thus two components that differ only in non-maximally aggressive strategies should be equated. We define the *strategy semantics* of component \mathcal{P} to be $[\mathcal{P}]_s = \{\mathcal{G}' \mid \exists \mathcal{G} \in stg(\mathcal{P}) : \mathcal{G} \preceq \mathcal{G}'\}$, i.e. the upward-closure of $stg(\mathcal{P})$ w.r.t. \preceq .

Game rules. When a component strategy \mathcal{G} is played against an environment strategy \mathcal{G}' , at each game state (i.e. a product state $p_{\mathcal{G}} \times p_{\mathcal{G}'}$) \mathcal{G} and \mathcal{G}' each propose a move (i.e. $mv_{\mathcal{G}}(p_{\mathcal{G}})$ and $mv_{\mathcal{G}'}(p_{\mathcal{G}'})$). If one of them is a delay and the other is an action, the action will prevail. If both propose delay moves (i.e. $mv_{\mathcal{G}}(p_{\mathcal{G}}), mv_{\mathcal{G}'}(p_{\mathcal{G}'}) \subseteq \mathbb{R}^{\geq 0}$), the smaller one (w.r.t. set containment) will prevail.

⁶ This is because our semantics is designed to preserve \perp rather than \top .

Since a delay move proposed at a strategy state is the maximal set of possible delays enabled at that state, the next move proposed at the new state after firing the set must be an action move (due to time additivity). Thus a play cannot have two consecutive delay moves.

If, however, both propose action moves, there will be a tie, which will be resolved by tossing the coin. For uniformity's sake, the coin can be treated as a special component. A strategy of the coin is a function h from tA^* to $\{0, 1\}$. We denote the set of all possible coin strategies as H .

A play of the game can be formalised as a composition of three strategies, one each from the component, environment and coin, denoted $\mathcal{G}_{\mathcal{P}} \parallel_h \mathcal{G}_{\mathcal{Q}}$. At a current game state $p_{\mathcal{P}} \times p_{\mathcal{Q}}$, if the prevailing action is α and we have $p_{\mathcal{P}} \xrightarrow{\alpha} s'_{\mathcal{P}}$ and $p_{\mathcal{Q}} \xrightarrow{\alpha} s'_{\mathcal{Q}}$, then the next game state is $s_{\mathcal{P}} \parallel s_{\mathcal{Q}}$. The play will stop when it reaches either \top or \perp . The composition will produce a simple path \mathcal{L} that is a run of $\mathcal{P} \parallel \mathcal{Q}$. Since $\mathcal{P} \parallel \mathcal{Q}$ gives rise to a *closed system* (i.e. the input alphabet is empty), a run of $\mathcal{P} \parallel \mathcal{Q}$ is a strategy of $\mathcal{P} \parallel \mathcal{Q}$.

This is crucial since it reveals that strategy composition of \mathcal{P} and \mathcal{Q} is closely related to their parallel composition: $stg(\mathcal{P} \parallel \mathcal{Q}) = \{\mathcal{G}_{\mathcal{P}} \parallel_h \mathcal{G}_{\mathcal{Q}} \mid \mathcal{G}_{\mathcal{P}} \in stg(\mathcal{P}), \mathcal{G}_{\mathcal{Q}} \in stg(\mathcal{Q}) \text{ and } h \in H\}$.

Parallel composition. Strategy composition, like component (parallel) composition, can be generalised to any pair of components \mathcal{P} and \mathcal{Q} with *composable alphabets*. That is, $O_{\mathcal{P}} \cap O_{\mathcal{Q}} = \{\}$. For such \mathcal{P} and \mathcal{Q} , $\mathcal{G}_{\mathcal{P}} \parallel_h \mathcal{G}_{\mathcal{Q}}$ gives rise to a tree rather than simple path TIOTS. That is, at each game state $p_{\mathcal{P}} \times p_{\mathcal{Q}}$, besides firing the prevailing $\alpha \in tO_{\mathcal{P}} \cup tO_{\mathcal{Q}}$, we need also to fire 1) all the synchronised inputs, i.e. $e \in I_{\mathcal{P}} \cap I_{\mathcal{Q}}$, and reach the new game state $s_{\mathcal{P}} \parallel s_{\mathcal{Q}}$ (assuming $p_{\mathcal{P}} \xrightarrow{e} s_{\mathcal{P}}$ and $p_{\mathcal{Q}} \xrightarrow{e} s_{\mathcal{Q}}$) and 2) all the independent inputs, i.e. $e \in (I_{\mathcal{P}} \cup I_{\mathcal{Q}}) \setminus (A_{\mathcal{P}} \cap A_{\mathcal{Q}})$, and reach the new game state $s_{\mathcal{P}} \times p_{\mathcal{Q}}$ or $p_{\mathcal{P}} \times s_{\mathcal{Q}}$. It is easy to verify that $\mathcal{G}_{\mathcal{P}} \parallel_h \mathcal{G}_{\mathcal{Q}}$ is a strategy of $\mathcal{P} \parallel \mathcal{Q}$.

Conjunction/disjunction. Besides strategy composition, *strategy conjunction* ($\&$) and *strategy disjunction* ($+$) are also definable. They are binary operators defined only on pairs of affine strategies. We define $\mathcal{G} \& \mathcal{G}' = \mathcal{G} \wedge \mathcal{G}'$ and $\mathcal{G} + \mathcal{G}' = \mathcal{G} \vee \mathcal{G}'$. Note that if \mathcal{G} and \mathcal{G}' are not affine, $\mathcal{G} \wedge \mathcal{G}'$ and $\mathcal{G} \vee \mathcal{G}'$ do not necessarily produce a strategy. For instance the disjunction of the strategies 1 and 2 in Figure 3 will produce a transition system that stops to output after the a transition.

Refinement. Strategy semantics induce an equivalence on TIOTSs. That is, \mathcal{P} and \mathcal{Q} are *strategy equivalent* iff $[\mathcal{P}]_s = [\mathcal{Q}]_s$. However, strategy equivalence is too fine for the purpose of *substitutive refinement* (cf Definition 5). For instance, transition systems \mathcal{P} and \mathcal{Q} in Figure 3 are substitutively equivalent, but are not strategy equivalent, because 1, 2, 3 and 4 are strategies of \mathcal{Q} (due to upward-closure w.r.t. \preceq), but A and B are not strategies of \mathcal{P} .

However, we demonstrate that *substitutive equivalence is reducible to strategy equivalence* providing we perform *disjunction closure* on strategies.

Lemma 1. *Given a pair of affine component strategies \mathcal{G}_0 and \mathcal{G}_1 , $\mathcal{G}_0 \parallel_h \mathcal{G}$ and $\mathcal{G}_1 \parallel_h \mathcal{G}$ are \perp -free for some environment strategy \mathcal{G} and $h \in H$ iff $\mathcal{G}_0 + \mathcal{G}_1 \parallel_h \mathcal{G}$ is \perp -free.*

We say Π^+ is a *disjunction closure* of Π iff it is the least superset of Π s.t. $\mathcal{G} + \mathcal{G}' \in \Pi^+$ for all pairs of affine strategies $\mathcal{G}, \mathcal{G}' \in \Pi^+$. It is easy to see the disjunction closure operation preserves the upward-closedness of strategy sets.

Theorem 1. *Given TIOTSs \mathcal{P} and \mathcal{Q} , $\mathcal{P} \sqsubseteq \mathcal{Q}$ iff $[\mathcal{Q}]_s^+ \subseteq [\mathcal{P}]_s^+$.*

For instance, the disjunction of strategies 1 and 3 produces A , while the disjunction of strategies 2 and 4 produces B . Thus $[\mathcal{P}]_s^+ = [\mathcal{Q}]_s^+$,

Relating operational composition to strategies. The operations of parallel composition, conjunction and disjunction defined on the operational models of TIOTSs (Section 2.4) can be characterised by simple operations on strategies in the game-based setting.

Lemma 2. *For \parallel -composable TIOTSs \mathcal{P} and \mathcal{Q} , $[\mathcal{P} \parallel \mathcal{Q}]_s^+ = \{\mathcal{G}_{\mathcal{P} \parallel \mathcal{Q}} \mid \exists \mathcal{G}_{\mathcal{P}} \in [\mathcal{P}]_s^+, \mathcal{G}_{\mathcal{Q}} \in [\mathcal{Q}]_s^+, h \in H : \mathcal{G}_{\mathcal{P}} \parallel_h \mathcal{G}_{\mathcal{Q}} \preceq \mathcal{G}_{\mathcal{P} \parallel \mathcal{Q}}\}$.*

Lemma 3. *For \vee -composable TIOTSs \mathcal{P} and \mathcal{Q} , $[\mathcal{P} \vee \mathcal{Q}]_s^+ = ([\mathcal{P}]_s^+ \cup [\mathcal{Q}]_s^+)^+$.*

Lemma 4. *For \wedge -composable TIOTSs \mathcal{P} and \mathcal{Q} , $[\mathcal{P} \wedge \mathcal{Q}]_s^+ = [\mathcal{P}]_s^+ \cap [\mathcal{Q}]_s^+$.*

Lemma 5. *For $\%$ -composable TIOTSs \mathcal{P} and \mathcal{Q} , $[\mathcal{P} \% \mathcal{Q}]_s^+ = \{\mathcal{G}_{\mathcal{P} \% \mathcal{Q}} \mid \forall \mathcal{G}_{\mathcal{Q}} \in [\mathcal{Q}]_s^+, h \in H : \mathcal{G}_{\mathcal{P} \% \mathcal{Q}} \parallel_h \mathcal{G}_{\mathcal{Q}} \in [\mathcal{P}]_s^+\}$.*

Thus conjunction and disjunction are the join and meet operations and quotient produces the coarsest TIOTS s.t. $(\mathcal{P}_0 \% \mathcal{P}_1) \parallel \mathcal{P}_1$ is a refinement of \mathcal{P}_0 .

Lemma 6. *For any TIOTS \mathcal{P} , $[\mathcal{P}^\neg]_s^+ = \{\mathcal{G}_{\mathcal{P}^\neg} \mid \forall \mathcal{G}_{\mathcal{P}} \in [\mathcal{P}]_s^+, h \in H : \mathcal{G}_{\mathcal{P}^\neg} \parallel_h \mathcal{G}_{\mathcal{P}} \text{ is } \perp\text{-free}\}$.*

Theorem 2. *\simeq is a congruence w.r.t. \parallel , \vee , \wedge and $\%$ subject to composability.*

Summary. Strategy semantics has given us a weakest \perp -preserving congruence (i.e. $[\mathcal{P}]_s^+$) for timed specification theories based on operators for (parallel) composition, conjunction, disjunction and quotient. Strategy semantics captures nicely the game-theoretical nature as well as the operational intuition of the specification theories. However, in a more declarative manner, the equivalence can also be characterised by timed traces, as we see in the next section.

4 Declarative Specification Theory

In this section, we develop a compositional specification theory based on timed traces. We introduce the concept of a timed-trace structure, which is an abstract

representation for a timed component. The timed-trace structure contains essential information about the component, for checking whether it can be substituted with another in a safety and liveness preserving manner.

Given any TIOTS $\mathcal{P} = \langle I, O, S, s^0, \rightarrow \rangle$, we can extract three sets of traces from $(\mathcal{P}^\perp)^\top$: TP (plain traces) is a set of timed traces leading to plain states, TE (error traces) a set of timed traces leading to \perp and TM (magic traces) a set of timed traces leading to \top . The three sets contain sufficient but not necessary information for our substitutive refinement, which is designed to preserve \perp rather than \top . For instance, adding any trace $tt \in TE$ to TP should not change the semantics of the component; similarly it is true for removing any trace $tt \in TP$ from TM . Based on a slight abstraction of the three sets we can define a *triple-trace structure* as the semantics of \mathcal{P} .

Definition 6 (Triple-trace structure). $\mathcal{TT}(\mathcal{P}) := (I, O, TT, TR, TE)$, where $TT := TE \cup TP \cup TM$ is the set of all traces and $TR := TE \cup TP$ the set of realisable traces.

Obviously, TE is extension-closed. TT is non-empty and prefix-closed. TR is prefix-closed and *fully branching*⁷ w.r.t. TT (i.e. $tt \hat{\cdot} \langle \alpha \rangle \in TT$ for all $tt \in TR$ and $\alpha \in tA$). $TT \setminus TR$ is time-extension closed (i.e. $tt \in X \Rightarrow tt \hat{\cdot} \langle d \rangle \in X$) and any pair of traces from $TT \setminus TR$ that are related by extension are related by time-extension.

From hereon let \mathcal{P}_0 and \mathcal{P}_1 be two TIOTSs with triple trace structures $\mathcal{TT}(\mathcal{P}_i) := (I_i, O_i, TT_i, TR_i, TE_i)$ for $i \in \{0, 1\}$. Define $\bar{i} = 1 - i$.

The substitutive refinement relation \sqsubseteq in Section 2.4 can equally be characterised by means of trace containment. Consequently, $\mathcal{TT}(\mathcal{P}_0)$ can be regarded as providing an alternative encoding of the set $[\mathcal{P}_0]_s^+$ of strategies.

Theorem 3. $\mathcal{P}_0 \sqsubseteq \mathcal{P}_1$ iff $TT_1 \subseteq TT_0$, $TR_1 \subseteq TR_0$ and $TE_1 \subseteq TE_0$.

We are now ready to define the timed-trace structure semantics for the operators of our specification theory. Intuitively, the timed-trace semantics mimic the synchronised product of the operational definitions in Section 2.4. An important fact utilised in formulating these operations on traces is that for any trace $tt \in tA^*$ and TIOTS \mathcal{P} , either tt is a trace of \mathcal{P} or there is some prefix tt_0 of tt s.t. tt_0 is an error or magic trace of \mathcal{P} .

Parallel composition. The idea behind parallel composition is that the projection of any trace in the composition onto the alphabet of one of the components should be a trace of that component.

Proposition 2. If \mathcal{P}_0 and \mathcal{P}_1 are \parallel -composable, then $\mathcal{TT}(\mathcal{P}_0 \parallel \mathcal{P}_1) = (I, O, TT, TR, TE)$ where $I = (I_0 \cup I_1) \setminus O$, $O = O_0 \cup O_1$ and the trace sets are given by:

- $TE = \{tt \mid tt \upharpoonright tA_i \in TE_i \wedge tt \upharpoonright tA_{\bar{i}} \in TR_{\bar{i}}\} \cdot tA^*$
- $TR = TE \uplus \{tt \mid tt \upharpoonright tA_i \in (TR_i \setminus TE_i) \wedge tt \upharpoonright tA_{\bar{i}} \in (TR_{\bar{i}} \setminus TE_{\bar{i}})\}$

⁷ This is due to \top/\perp -completion.

$$- TT = TR \uplus \{tt \mid tt \upharpoonright tA_i \in (TT_i \setminus TR_i) \wedge tt_0 < tt \upharpoonright tA_{\bar{i}} \Rightarrow tt_0 \in (TR_{\bar{i}} \setminus TE_{\bar{i}})\} \cdot \mathbb{R}^{\geq 0}.$$

The above says tt is an error trace if the projection of tt on one component is an error trace while the projection of tt on the other component is not a magic trace. tt is a realisable trace if tt is either an error trace or a plain trace. tt is a plain trace if the projection of tt on both components are plain traces. Finally, tt is a magic trace if its projection on one component is a magic trace, while the projection of all strict prefixes of tt on the other component is a plain trace.

Disjunction. From any composite state in the disjunction of two components, the composition should only be willing to accept inputs that are accepted by both components, but should accept the union of outputs. After witnessing an output enabled by only one of the components, the disjunction should behave like that component. Because of the way that \perp and \top work in Table 1, this loosely corresponds to taking the union of the traces from the respective components.

Proposition 3. *If \mathcal{P}_0 and \mathcal{P}_1 are \vee -composable, then $\mathcal{TT}(\mathcal{P}_0 \vee \mathcal{P}_1) = (I, O, TR_0 \cup TR_1 \cup TM, TR_0 \cup TR_1, TE_0 \cup TE_1)$, where $I = I_0 = I_1$, $O = O_0 = O_1$ and $TM = \{tt \mid tt \in (TT_i \setminus TR_i) \wedge \exists tt_0 \leq tt : tt_0 \in (TT_{\bar{i}} \setminus TR_{\bar{i}})\} \cdot \mathbb{R}^{\geq 0}$.*

Essentially, tt is a magic trace if it is a magic trace on one component while one of its prefixes is a magic trace on the other component. The realisable and error traces are simply the union of the corresponding traces on \mathcal{P}_0 and \mathcal{P}_1 .

Conjunction. Similarly to disjunction, from any composite state in the conjunction of two components, the composition should only be willing to accept outputs that are accepted by both components, and should accept the union of inputs, until a stage when one of the component's input assumptions has been violated, after which it should behave like the other component. Because of the way that both \perp and \top work in Table 1, this essentially corresponds to taking the intersection of the traces from the respective components.

Proposition 4. *If \mathcal{P}_0 and \mathcal{P}_1 are \wedge -composable, then $\mathcal{TT}(\mathcal{P}_0 \wedge \mathcal{P}_1) = (I, O, (TR_0 \cap TR_1) \cup TM, TR_0 \cap TR_1, TE_0 \cap TE_1)$, where $I = I_0 = I_1$, $O = O_0 = O_1$ and $TM = \{tt \mid tt \in (TT_i \setminus TR_i) \wedge tt_0 < tt \Rightarrow tt_0 \in TR_{\bar{i}}\} \cdot \mathbb{R}^{\geq 0}$.*

A trace tt is a magic trace if it is a magic trace on one of the components, and all strict prefixes of the trace are realisable by the other component. The realisable and error traces are simply the intersection of the corresponding traces on \mathcal{P}_0 and \mathcal{P}_1 .

Quotient. Quotient ensures its composition with the second component is a refinement of the first. Given the synchronised running of \mathcal{P}_0 and \mathcal{P}_1 , if \mathcal{P}_0 is in a more refined state than \mathcal{P}_1 , the quotient will try to rescue the refinement by taking \top as its state (so that its composition with \mathcal{P}_1 's state will refine \mathcal{P}_0 's). If \mathcal{P}_0 is in a less or equally refined state than \mathcal{P}_1 's, the quotient will take the worst possible state without breaking the refinement.

Proposition 5. *If \mathcal{P}_0 dominates \mathcal{P}_1 , then $\mathcal{TT}(\mathcal{P}_0 \% \mathcal{P}_1) = (I, O, TT, TR, TE)$, where $I = I_0 \cup O_1$, $O = O_0 \setminus O_1$, and the trace sets satisfy:*

- $TE = \{tt \mid (tt \in TE_0 \wedge tt_0 < tt \Rightarrow tt_0 \upharpoonright tA_1 \notin TE_1) \vee (tt \upharpoonright tA_1 \in (TT_1 \setminus TR_1) \wedge tt_0 < tt \Rightarrow tt_0 \notin TT_0 \setminus TR_0)\} \cdot tA^*$
- $TR = TE \uplus \{tt \mid tt \in (TR_0 \setminus TE_0) \wedge tt \upharpoonright tA_1 \in (TR_1 \setminus TE_1)\}$
- $TT = TR \uplus \{tt \mid (tt \in (TT_0 \setminus TR_0) \wedge tt_0 \leq tt \Rightarrow tt_0 \upharpoonright tA_1 \in TR_1) \vee (tt \upharpoonright tA_1 \in TE_1 \wedge tt_0 \leq tt \Rightarrow tt_0 \notin TE_0)\}$.

The above says tt is an error trace if either 1) tt is an error trace in \mathcal{P}_0 , but the projection of any strict prefix of tt on \mathcal{P}_1 is not an error trace, or 2) the projection of tt on \mathcal{P}_1 is a magic trace, but no strict prefix of tt is a magic trace in \mathcal{P}_0 . tt is a magic trace if either 1) tt is a magic trace in \mathcal{P}_0 , but the projection of any prefix of tt is not a magic trace in \mathcal{P}_1 , or 2) the projection of tt on \mathcal{P}_1 is an error trace, but no prefix of tt is an error trace in \mathcal{P}_0 .

Mirroring of triple trace structures is straightforward: $\mathcal{TT}(\mathcal{P}_0)^\neg = (O_0, I_0, TT_0, TT_0 \setminus TE_0, TT_0 \setminus TR_0)$. This is because dealing with traces means we have implicit determinism, so we can skip the determinisation step. Consequently, quotient can also be defined as the derived operator $(\mathcal{TT}(\mathcal{P}_0)^\neg \parallel \mathcal{TT}(\mathcal{P}_1))^\neg$.

5 Comparison with Related Works

Based on linear-time, our timed theory owes much to the pioneering work of trace theories in asynchronous circuit verification, such as Dill's trace theory [7]. Our mirror operator is essentially a timed extension of the mirror operator from asynchronous circuit verification. The definition of quotient based on mirroring (for the untimed case) was first presented by Verhoeff as his Factorisation Theorem [8].

Our work is also deeply influenced by the work of [5] on timed games, with some modifications. Firstly, a TIOTS is regarded as a set of component strategies, rather than a timed game graph. We adopt most of the game rules in [5], except that, due to our requirement that proposed delay moves are maximal delays allowed by a strategy, a play cannot have consecutive delay moves. This enables us to avoid the complexity of time-blocking strategies and blame assignment, but does not ensure non-Zenoness⁸. Secondly, we do not use timestop/semi-timestop to model time errors (i.e. bounded-liveness errors). Rather, we introduce the explicit inconsistent state \perp to model both time and immediate (i.e. safety) errors. Timestop is used to model the magic state, which can simplify the definition of parallel, conjunction and quotient and enables us to avoid the complexity of having two transition relations and well-formedness of timed interfaces.

Last but not least, our work is related to [4], as both devise a complete timed specification theory. The major differences lie in the use of timed alternating simulation as refinement in [4], while ours is linear-time. An advantage of

⁸ Zeno behaviours (infinite action moves within finite time) in a play are not regarded as abnormal behaviours in our semantics.

our work is that refinement is the weakest congruence preserving inconsistency, while beneficial in [4] is the algorithmic efficiency of branching-time simulation checking. Moreover, [4] has fully implemented the timed-game algorithms.

We briefly mention other related works, which include timed modal transition systems [9,10], the timed I/O model [3] and embedded systems [11,12].

6 Conclusions

We have formulated a rich compositional specification theory for components with real-time constraints based on a linear-time notion of substitutive refinement. The operators of hiding and renaming can also be defined, according to our past experiences [13]. We believe that our theory can be reformulated as a timed extension of Dill's trace theory [7]. Future work will include an investigation of realisability and assume-guarantee reasoning.

Acknowledgments. The authors are supported by EU FP7 project CONNECT and ERC Advanced Grant VERIWARE.

References

1. Chen, T., Chilton, C., Jonsson, B., Kwiatkowska, M.: A Compositional Specification Theory for Component Behaviours. In Seidl, H., ed.: Programming Languages and Systems, Proc. 21st European Symposium on Programming (ESOP'12). Volume 7211 of Lecture Notes in Computer Science., Springer-Verlag (2012) 148–168
2. de Alfaro, L., Henzinger, T.A.: Interface automata. SIGSOFT Softw. Eng. Notes **26** (2001) 109–120
3. Kaynar, D.K., Lynch, N.A., Segala, R., Vaandrager, F.W.: Timed i/o automata: A mathematical framework for modeling and analyzing real-time systems. In: RTSS. (2003)
4. David, A., Larsen, K.G., Legay, A., Nyman, U., Wasowski, A.: Timed I/O automata: a complete specification theory for real-time systems. In: Proc. 13th ACM International Conference on Hybrid systems: computation and control. HSCC '10, ACM (2010) 91–100
5. de Alfaro, L., Henzinger, T.A., Stoelinga, M.: Timed interfaces. In Sangiovanni-Vincentelli, A., Sifakis, J., eds.: Embedded Software. Volume 2491 of LNCS. Springer-Verlag (2002) 108–122
6. Wang, X.: Maximal Confluent Processes. In: Proc. of PETRI NETS 2012. Volume 7347 of Lecture Notes in Computer Science., Springer-Verlag (2012)
7. Dill, D.L.: Trace theory for automatic hierarchical verification of speed-independent circuits. ACM distinguished dissertations. MIT Press (1989)
8. Verhoeff, T.: A Theory of Delay-Insensitive Systems. PhD thesis, Dept. of Math. and C.S., Eindhoven Univ. of Technology (1994)
9. Bertrand, N., Pinchinat, S., Raclet, J.B.: Refinement and consistency of timed modal specifications. In: LATA. (2009) 152–163
10. Cerans, K., Godskesen, J.C., Larsen, K.G.: Timed modal specification - theory and tools. In: CAV. (1993) 253–267

11. Thiele, L., Wandeler, E., Stoimenov, N.: Real-time interfaces for composing real-time systems. In: EMSOFT. (2006)
12. I. Lee, J.Y.T.L., Song, S.: Handbook of Real-Time and Embedded Systems. Chapman (2007)
13. Wang, X., Kwiatkowska, M.Z.: On process-algebraic verification of asynchronous circuits. Fundam. Inform. **80** (2007) 283–310