

Conjunction on Processes: Full–Abstraction via Ready–Tree Semantics

Gerald Lüttgen^{1*} and Walter Vogler²

¹ Department of Computer Science, University of York, York YO10 5DD, U.K.
luettgen@cs.york.ac.uk

² Institut für Informatik, Universität Augsburg, D-86135 Augsburg, Germany
vogler@informatik.uni-augsburg.de

Abstract. A key problem in mixing operational (e.g., process–algebraic) and declarative (e.g., logical) styles of specification is how to deal with inconsistencies arising when composing processes under conjunction. This paper introduces a conjunction operator on labelled transition systems capturing the basic intuition of “ a and $b = false$ ”, and considers a naive preorder that demands that an inconsistent specification can only be refined by an inconsistent implementation.

The main body of the paper is concerned with characterising the largest precongruence contained in the naive preorder. This characterisation will be based on a novel semantics called ready–tree semantics, which refines ready traces but is coarser than ready simulation. It is proved that the induced ready–tree preorder is compositional and fully–abstract, and that the conjunction operator indeed reflects conjunction.

The paper’s results provide a foundation for, and an important step towards a unified framework that allows one to freely mix operators from process algebras and linear–time temporal logics.

Keywords. Labelled transition system, conjunction, consistency preorder, ready–tree semantics, ready–tree preorder, full abstraction.

1 Introduction

Process algebra [2] and *temporal logic* [13] are two popular approaches to formally specifying and reasoning about reactive systems. The process–algebraic paradigm is founded on notions of *refinement*, where one typically formulates a system specification and its implementation in the same notation and then proves using *compositional reasoning* that the latter refines the former. The underlying semantics is often given operationally, and refinement relations are formalised as precongruences. In contrast, the temporal–logic paradigm is based on the use of temporal logics to formulate specifications abstractly, with implementations being denoted in an operational notation. One then verifies a system by establishing that it is a model of its specification.

* Research support was partially provided by the NSF under grant CCR–9988489.

Recently, two papers have been published aimed at marrying process algebras and temporal logics [5, 6]. While the first paper introduces a semantic framework based on Büchi automata, the second paper considers labelled transition systems augmented with an “unimplementability predicate”. This predicate captures *inconsistencies* arising when composing processes conjunctively; e.g., the composition $a \wedge b$ is contradictory since a run of a process cannot begin with both actions a and b . Moreover, the frameworks in [5, 6] are equipped with a refinement preorder based on De Nicola and Hennessy’s must-testing preorder [11]. However, the obtained results are unsatisfactory: the refinement preorder in [5] is not a precongruence, while the \wedge -operator in [6] is not *conjunction* with respect to the studied precongruence.

This paper solves the deficiencies of [5, 6] within a simple setting of labelled transition systems in which a state represents either external (non-deterministic) or internal (disjunctive) choice. Moreover, states that are vacuously *true* or *false* are tagged accordingly. The tagging of *false* states, or inconsistent states, is given by an inductive *inconsistency predicate* that is defined very similar but subtly different to the unimplementability predicate of [6]. We then equip our setting with two operators: the conjunction operator \wedge is in essence a synchronous composition on observable actions and an interleaving product on the unobservable action τ , but additionally captures inconsistencies; the disjunction operator \vee simply resembles the process-algebraic operator of internal choice.

Our variant of labelled transition systems gives rise to a naive refinement preorder \sqsubseteq_F requiring that an inconsistent specification cannot be refined except by an inconsistent implementation. We characterise the *consistency preorder*, i.e. the largest precongruence contained in \sqsubseteq_F when conjunctively closing under all contexts. To do so, we define a novel semantics, called *ready-tree semantics* which is — at least when disallowing divergent behaviour — finer than both must-testing semantics [11] and ready-trace semantics [7], but coarser than ready simulation [3]. The resulting *ready-tree preorder* \sqsubseteq is not only compositional for \wedge and \vee and fully-abstract with respect to \sqsubseteq_F , but also possesses several other desired properties. In particular, we prove that \wedge (\vee) is indeed conjunction (disjunction) relative to \sqsubseteq , and that \wedge and \vee satisfy the expected boolean laws, such as the distributivity laws.

Our results are a significant first step towards the goal of developing a uniform calculus in which one can freely mix process-algebraic and temporal-logic operators. This will give engineers powerful tools to model system components at different levels of abstraction and to impose logical constraints on the execution behaviour of components. The proposed ready-tree preorder will allow engineers to step-wise and component-wise refine systems by trading off logical content for operational content.

Organisation. The next section presents our setting of labelled transition systems augmented with *true* and *false* predicates, together with a conjunction and a disjunction operator. Sec. 3 defines the novel ready-tree semantics, addresses expressiveness issues of several ready-tree variants and introduces the ready-tree preorder. Our compositionality and full-abstraction results are proved in

Sec. 4. Finally, Sec. 5 discusses our results in light of related work, while Sec. 6 presents our conclusions and suggests directions for future research.

2 Labelled Transition Systems & Conjunction

This section first introduces our process-algebraic setting and particularly conjunctive composition informally, discusses semantic choices and their implications, and finally gives a formal account of our framework.

Motivation. Our setting models processes as labelled transition systems, which may be composed conjunctively and disjunctively. As usual in process algebra, transition labels are actions taken from some alphabet $\mathcal{A} = \{a, b, \dots\}$. When an action a is offered by the environment and the process under consideration is in a state having one or more outgoing a -transitions, the process must choose and perform one of them. If there is no outgoing a -transition, then the process stays in its state, at least in classical process-algebraic frameworks where the composition between a process and its environment is modelled using some parallel operator. However, in a conjunctive setting we wish to mark the composed state between process and environment as inconsistent, if the environment offers an action that the process cannot perform, or vice versa. Hence, taking ordinary synchronous composition as operator for conjunction is insufficient.




Fig. 1. Basic intuition behind conjunctive composition.

We illustrate this intuition behind our conjunction operator \wedge and its implications by the example labelled transition systems of Fig. 1. First, consider the processes p , q and r . Process p and q specify that exactly action a and respectively action b is offered initially. Similarly, process r specifies that a and b are offered initially. From this perspective, $p \wedge q$ as well as $p \wedge r$ are *inconsistent* and should be tagged as such. Formally, our labelled transition systems will be augmented by an *inconsistency predicate* F , so that $p \wedge q$, $p \wedge r \in F$ in our example. We also refer to inconsistent states as *false-states*.

Now consider the conjunction $p' \wedge q'$ shown on the right in Fig. 1. Since both conjuncts require action a to be performed, $p' \wedge q'$ should have an a -transition. From the preceding discussion, this transition should lead to a *false-state*. No sensible process can meet these requirements of being able to perform a and being inconsistent afterwards. Thus, our inconsistency predicate will propagate backwards to the conjunction itself, as indicated in Fig. 1.




Fig. 2. Backward propagation of inconsistencies.

Fig. 2 shows more intricate examples of backward propagation. The inconsistency of the target state of the a -transition of the process on the left propagates backwards to its source state. This is the case although the source state is able to offer a transition leading to a consistent state. However, that transition can only be taken if the environment offers action b . The process is forced into the inconsistency when the environment offers action a .

The situation is different for the process in the middle, which has an additional a -transition leading to a consistent state. Here, the process is consistent, as it can choose to execute this new a -transition and thus avoid to enter a *false*-state. In fact, this choice can be viewed as a disjunction between the two a -branches. As an aside, note that in [6] the design decision was to consider a process already as inconsistent if some a -derivative is. While there might be an intuitive justification for that, it led to a setting where the implied conjunction operator did not reflect conjunction for the studied refinement preorder, i.e., where Thm. 21(1) did not hold.

Disjunction can be made explicit by using the classical *internal-choice* operator. This operator may as usual be expressed by employing the special, unobservable action $\tau \notin \mathcal{A}$ as shown on the right in Fig. 2. Hence, we may identify the internal-choice operator with the disjunction operator \vee desired in our setting. Moreover, a disjunction $p \vee q$ is inconsistent if both p and q are *false*-states. In particular, the process on the right in Fig. 2 will represent $\text{false} \vee q$ in our approach, with q from Fig. 1, which clearly should be consistent.

Formalisation. For notational convenience we denote $\mathcal{A} \cup \{\tau\}$ by \mathcal{A}_τ and use α, β, \dots as representatives of \mathcal{A}_τ . We start off by defining our notion of *labelled transition systems* (LTS). The LTSs considered here are augmented with a *false*-predicate F on states, as discussed above, and dually with a *true*-predicate T . A state in F represents inconsistent, empty behaviour, while a state in T represents completely underspecified, arbitrary behaviour.

Formally, an LTS is a quadruple $\langle P, \longrightarrow, T, F \rangle$, where P is the set of *processes* (states), $\longrightarrow \subseteq P \times \mathcal{A}_\tau \times P$ is the *transition relation*, and $T \subseteq P$ and $F \subseteq P$ are the *true*- and *false*-predicates, respectively. We write $p \xrightarrow{\alpha} p'$ instead of $\langle p, \alpha, p' \rangle \in \longrightarrow$, $p \xrightarrow{\alpha}$ instead of $\exists p' \in P. p \xrightarrow{\alpha} p'$, and $p \longrightarrow$ instead of $\exists p' \in P. p \xrightarrow{\alpha} p'$. When $p \xrightarrow{\alpha} p'$, we say that process p can perform an α -step to p' , and we call p' an α -derivative. We also require an LTS to satisfy the following τ -purity condition: $p \xrightarrow{\tau}$ implies $\nexists a \in \mathcal{A}. p \xrightarrow{a}$, for all $p \in P$. Hence, each process represents either an external or internal (disjunctive) choice

between its outgoing transitions. This restriction turns out to be technically convenient, and we leave exploring the consequences of lifting it for future work. The LTSs of interest to us need to satisfy four further properties, as stated in the following formal definition:

Definition 1 (Logical LTS). An LTS $\langle P, \longrightarrow, T, F \rangle$ is a *logical LTS* if it satisfies the following conditions:

1. $T \cap F = \emptyset$
2. $T \subseteq \{p \mid p \not\rightarrow\}$
3. $F \subseteq P$ such that $p \in F$ whenever $\exists \alpha \in \mathcal{I}(p) \forall p' \in P. p \xrightarrow{\alpha} p' \implies p' \in F$
4. p cannot stabilise $\implies p \in F$

Naturally, we require that a process cannot be tagged *true* and *false* at the same time. As a *true*-process specifies arbitrary, full behaviour, any behaviour made explicit by outgoing transitions is already implicitly included; hence, any outgoing transitions may simply be cut off. The third condition formalises the backwards propagation of inconsistencies as discussed in the motivation section above; here, $\mathcal{I}(p)$ stands for the set $\{\alpha \in \mathcal{A}_\tau \mid p \xrightarrow{\alpha}\}$ of initial actions of process p , to which we also refer as *ready set*.

The fourth condition relates to *divergence*, i.e., infinite sequences of τ -transitions. In many semantic frameworks, e.g. [11, 7], divergence is considered catastrophic, while in our setting catastrophic behaviour is inconsistent behaviour. We view divergence only as catastrophic if a process cannot *stabilise*, i.e., if it cannot get out of an infinite, internal computation. While this is intuitive, there is also a technical reason to which we come back shortly.

To formalise our notion of stabilisation, we first introduce a weak transition relation $\Longrightarrow \subseteq P \times (\mathcal{A}_\tau \cup \{\epsilon\}) \times P$ which is defined by (1) $p \xRightarrow{\epsilon} p'$ if $p \equiv p' \notin F$, where \equiv denotes syntactic equality, or if $p \notin F$ and $p \xrightarrow{\tau} p'' \xRightarrow{\epsilon} p'$ for some p'' , and (2) $p \xRightarrow{a} p'$ if $p \notin F$ and $p \xrightarrow{a} p'' \xRightarrow{\epsilon} p'$ for some p'' . Our definition of a weak transition is slightly unusual: a weak transition cannot pass through *false*-states since these cannot occur in computations, and it does not abstract from τ -transitions preceding a visible transition. However, we only will use weak visible transitions from *stable* states, i.e., states with no outgoing τ -transition. Finally, we can now formalise stabilisation: a process p *can stabilise* if $p \xRightarrow{\epsilon} p'$ for some stable p' .

Note that both Conds. (3) and (4) are inductively defined conditions. We refer to them as *fixed point conditions of F for LTS*. For convenience, we will often write LTS instead of logical LTS in the sequel. Moreover, whenever we mention a process p without giving a respective LTS explicitly, we assume implicitly that such an LTS $\langle P, \longrightarrow, T, F \rangle$ is given. We let tt (ff) stand for the *true* (*false*) process, which is the only process of an LTS with $tt \in T$ ($ff \in F$).

Operators. Our conjunction operator is essentially a synchronous composition for visible transitions and an asynchronous composition for τ -transitions. However, we need to take care of the T - and F -predicates.

Definition 2 (Conjunction Operator). The conjunction of two logical LTSs $\langle P, \longrightarrow_P, T_P, F_P \rangle$, $\langle Q, \longrightarrow_Q, T_Q, F_Q \rangle$ is the LTS $\langle P \wedge Q, \longrightarrow_{P \wedge Q}, T_{P \wedge Q}, F_{P \wedge Q} \rangle$ defined by:

- $P \wedge Q =_{\text{df}} \{p \wedge q \mid p \in P, q \in Q\}$
- $\longrightarrow_{P \wedge Q}$ is determined by the following operational rules:

$$\begin{array}{lll}
p \xrightarrow{\tau}_P p' & \Longrightarrow & p \wedge q \xrightarrow{\tau}_{P \wedge Q} p' \wedge q \\
q \xrightarrow{\tau}_Q q' & \Longrightarrow & p \wedge q \xrightarrow{\tau}_{P \wedge Q} p \wedge q' \\
p \xrightarrow{a}_P p', q \xrightarrow{a}_Q q' & \Longrightarrow & p \wedge q \xrightarrow{a}_{P \wedge Q} p' \wedge q' \\
p \in T_P, q \xrightarrow{\alpha}_Q q' & \Longrightarrow & p \wedge q \xrightarrow{\alpha}_{P \wedge Q} p \wedge q' \\
q \in T_Q, p \xrightarrow{\alpha}_P p' & \Longrightarrow & p \wedge q \xrightarrow{\alpha}_{P \wedge Q} p' \wedge q
\end{array}$$

- $p \wedge q \in T_{P \wedge Q}$ if and only if $p \in T_P$ and $q \in T_Q$
- $p \wedge q \in F_{P \wedge Q}$ if at least one of the following conditions hold:
 1. $p \in F_P$ or $q \in F_Q$
 2. $p \notin T_P$ and $q \notin T_Q$ and $p \wedge q \not\xrightarrow{\tau}_{P \wedge Q}$ and $\mathcal{I}(P) \neq \mathcal{I}(Q)$
 3. $\exists \alpha \in \mathcal{I}(p \wedge q) \forall p' \wedge q'. p \wedge q \xrightarrow{\alpha}_{P \wedge Q} p' \wedge q' \Longrightarrow p' \wedge q' \in F_{P \wedge Q}$
 4. $p \wedge q$ cannot stabilise

Note that the treatment of *true*-processes when defining $\longrightarrow_{P \wedge Q}$ and $T_{P \wedge Q}$ reflects our intuition that these processes allow arbitrary behaviour. We are left with explaining Conds. (1)–(4). Firstly, a conjunction is inconsistent if any conjunct is. Conds. (2) and (3) reflect our intuition of inconsistency and, respectively, backward propagation stated in the motivation section above. Cond. (4) is added to enforce Def. 1(4). We refer to Conds. (3) and (4) as *fixed point conditions of F for \wedge* .

It is easy to check that conjunction is well-defined, i.e., that the conjunctive composition of two logical LTSs satisfies the four conditions of Def. 1. For Def. 1(1) in particular, note that $p \wedge q \in T_{P \wedge Q}$ does not satisfy any of the four conditions for $F_{P \wedge Q}$.




Fig. 3. Counter-example demonstrating non-associativity.

We may now demonstrate why we have treated non-escapable divergence as catastrophic in our setting. This is because, otherwise, our conjunction operator

would not be associative as demonstrated by the example depicted in Fig. 3. If the conjunction is computed from the left, the result is the first conjunct. Computed from the right, the result is the same but with both processes being in F . Hence, in the first case, the divergence hides the inconsistency. Since this is not really plausible and associativity of conjunction is clearly desirable, we need some restriction for divergence; it turns out that restricting divergence to escapable divergence, i.e., potential stabilisation, is sufficient for our purposes.

Definition 3 (Disjunction Operator). The disjunction of two logical LTSs $\langle P, \longrightarrow_P, T_P, F_P \rangle$ and $\langle Q, \longrightarrow_Q, T_Q, F_Q \rangle$ satisfying (w.l.o.g.) $P \cap Q = \emptyset$, is the logical LTS $\langle P \vee Q, \longrightarrow_{P \vee Q}, T_{P \vee Q}, F_{P \vee Q} \rangle$ defined by:

- $P \vee Q =_{\text{df}} \{p \vee q \mid p \in P, q \in Q\} \cup P \cup Q$
- $\longrightarrow_{P \vee Q}$ is determined by the following operational rules:

$$\begin{array}{lll}
\text{always} & \Longrightarrow & p \vee q \xrightarrow{\tau}_{P \vee Q} p \\
\text{always} & \Longrightarrow & p \vee q \xrightarrow{\tau}_{P \vee Q} q \\
p \xrightarrow{\alpha}_P p' & \Longrightarrow & p \xrightarrow{\alpha}_{P \vee Q} p' \\
q \xrightarrow{\alpha}_Q q' & \Longrightarrow & q \xrightarrow{\alpha}_{P \vee Q} q'
\end{array}$$

- $p \vee q \notin T_{P \vee Q}$ always
- $p \vee q \in F_{P \vee Q}$ if and only if $p \in F_P$ and $q \in F_Q$

The definition of disjunction, which reflects internal choice, is quite straightforward and well-defined. Only the definition of $T_{P \vee Q}$ for $p \vee q$ is unusual, as one would expect to simply have $p \vee q \in T$ whenever p or q is in T . However, then Cond. (2) of Def. 1 would be violated. Our alternative definition respects this condition and is semantically equivalent. In the sequel we leave out indices of relations and predicates whenever the context is clear.

Refinement preorder. As a basis for our semantical considerations we now define a naive refinement preorder stating that an inconsistent specification cannot be implemented except by an inconsistent implementation.

Definition 4 (Naive Consistency Preorder). The *naive consistency preorder* \sqsubseteq_F on processes is defined by $p \sqsubseteq_F q$ if $p \in F \implies q \in F$.

One main objective of this paper is to identify the corresponding fully-abstract preorder with respect to conjunction and disjunction, which is contained in \sqsubseteq_F . Our approach follows the testing idea of De Nicola and Hennessy [11], for which we define a testing relation \sqsubseteq as usual. Note that a process and an observer need to be composed not simply synchronously but conjunctively. This is because we want the observer to be sensitive to inconsistencies, so that $p \sqsubseteq q$ if each “conjunctive observer” that sees an inconsistency in p also sees one in q .

Definition 5 (Consistency Testing Preorder). The *consistency testing preorder* \sqsubseteq on processes is defined as the conjunctive closure of the naive consistency preorder under all processes (observers), i.e., $p \sqsubseteq q$ if $\forall o. p \wedge o \sqsubseteq_F q \wedge o$.

To characterise the full-abstract precongruence contained in \sqsubseteq_F we will introduce a new semantics, called *ready-tree semantics*, and an associated preorder, the *ready-tree preorder*, which is compositional for conjunction and disjunction and which coincides with \sqsubseteq .




Fig. 4. Example processes.

Example. As an illustration for our approach, consider process *spec* in Fig. 4. For $\mathcal{A} = \{a, b, c\}$, *spec* specifies that action *c* can only occur after action *a*. In the light of the above discussions, an implementation of this intuition should offer initially either just *a*, or *a* and *b*, or just *b*, so that *spec* is an internal choice between three states. Moreover, after an action *a*, nothing more is specified; after an action *b*, the same is required as initially. While our specification of this simple behaviour may look quite complex, we may imagine that process *spec* is generated automatically from a temporal-logic formula. Fig. 4 also shows process *impl* which repeats sequence *abc*, and *spec* \wedge *impl*. It will turn out that *spec* \sqsubseteq *impl*, as we will show in Sec. 4.

3 Ready-Tree Semantics

A first guess for achieving a compositional semantics reflecting consistency testing is to use a kind of *ready-trace semantics* [7]. Such a semantics would refine trace semantics by checking the initial action set of every stable state along each trace. However, this is not sufficient when dealing with inconsistencies, since inconsistencies propagate backwards along traces as explained in Sec. 2. It turns out that a set of tree-like observations is needed, which leads to a novel denotational-style semantics which we call *ready-tree semantics*.

Observation trees & ready trees. A tree-like observation can itself be seen as a deterministic LTS with empty *F*-predicate, reflecting that observers are internally consistent.

Definition 6 (Observation Tree). An *observation tree* is a LTS $\langle V, \rightarrow, T, \emptyset \rangle$ satisfying the following properties:

1. $\langle V, \rightarrow \rangle$ is a tree whose root is referred to as v_0
2. $\forall v \in V. v$ stable
3. $\forall v \in V, a \in \mathcal{I}(v) \exists v' \in V. v \xrightarrow{a} v'$

We often denote such an observation tree by its root v_0 .

Next we define the observations of a process p , called *ready trees*. Note that p can only be observed at its stable states.

Definition 7 (Ready Tree). An observation tree v_0 is a *ready tree* of p if there is a labelling $h : V \rightarrow P$ satisfying the following conditions:

1. $\forall v \in V. h(v)$ stable and $h(v) \notin F$
2. $p \xRightarrow{\varepsilon} h(v_0)$
3. $\forall v \in V, a \in \mathcal{A}. v \xrightarrow{a} v'$ implies (a) $h(v') = h(v) \in T$ or (b) $h(v) \xRightarrow{a} h(v')$
4. $\forall v \in V. (v \notin T \text{ and } h(v) \notin T) \text{ implies } \mathcal{I}(v) = \mathcal{I}(h(v))$

Intuitively, nodes v in a ready tree represent stable states $h(v)$ of p (cf. Cond. (1), first part) and transitions represent computations containing exactly one observable action (cf. Cond. (3)(b)). Since computations do not contain *false*-states, no represented state is in F (cf. Cond. (1), second part). Since p might not be stable, the root v_0 of a ready tree represents a stable state reachable from p by some internal computation (cf. Cond. (2)). If the state $h(v)$ represented by node v is in T , the subtree of v is arbitrary since $h(v)$ is considered to be completely underspecified (cf. Conds. (3)(a) and (4)). In case $h(v) \notin T$, one distinguishes two cases: (i) if $v \notin T$, then v and $h(v)$ must have the same initial actions, i.e., the same *ready set*; (ii) if $v \in T$, the observation stops at this node and nothing is required in Conds. (3) and (4).

In the following, we write $\text{RT}(p)$ for the set of all ready trees of p , $\text{fRT}(p)$ for the set of all ready trees of p that have finite depth (*finite-depth* ready trees), and $\text{cRT}(p)$ for the set of ready trees $\langle V, \rightarrow, T, \emptyset \rangle$ where $T = \emptyset$ (*complete* ready trees). Note that a complete ready tree is called *complete* as it never stops its task of observing; hence, complete ready trees are often infinite in practice. Moreover, *false*-states may be characterised as follows:

Lemma 8. $\text{RT}(p) = \emptyset$ if and only if $p \in F$.

Proof. Direction “ \Leftarrow ” follows immediately from Def. 7(2) and the definition of $\xRightarrow{\varepsilon}$. For Direction “ \Rightarrow ” we know by $p \notin F$ and Def. 1(4) of the existence of some p' such that $p \xRightarrow{\varepsilon} p' \not\rightarrow$. Hence, $tt \in \text{RT}(p)$ by $h(tt) =_{\text{df}} p'$. \square




Fig. 5. Some ready trees of *spec*.

We illustrate our concept of ready trees by returning to our example of Fig. 4. Some of the ready trees of process *spec* are shown in Fig. 5. In the first ready tree, the observation stops after the third *b*. In the second tree, we see that we can observe an arbitrary tree after *a*, since the respective state of *spec* is in *T*. An arbitrary tree can also consist of just the root, as shown for the right-most *a* in the third tree; this tree is also complete. Process *impl* in Fig. 4 has only one complete ready tree which is an infinite path repeating sequence *abc*; this is also a ready tree of *spec*.

Ready-tree preorder & expressiveness. Our ready-tree semantics suggests the following refinement preorder:

Definition 9 (Ready-Tree Preorder). The *ready-tree preorder* \sqsubseteq on processes is defined as reverse ready-tree inclusion, i.e., $p \sqsubseteq q$ if $\text{RT}(q) \subseteq \text{RT}(p)$.

This preorder will turn out to be the desired fully-abstract preorder contained in the naive consistency preorder. We first show that \sqsubseteq could just as well be formulated on the basis of complete ready trees and, for finitely branching LTS, of finite-depth ready trees. A crucial notion for our results is the following:

Definition 10 (Ready-Tree Prefix). Ready tree v_0 is *prefix* of ready tree w_0 , in signs $v_0 \leq w_0$, if there exists an injective mapping $\rho : V \hookrightarrow W$ such that:

1. $\rho(v_0) = w_0$
2. $v \xrightarrow{a} v' \implies \rho(v) \xrightarrow{a} \rho(v')$
3. $\rho(v) \xrightarrow{a} w' \implies v \in T \text{ or } (\exists v'. v \xrightarrow{a} v' \text{ and } \rho(v') = w')$
4. $\rho(v) \in T \implies v \in T$

Intuitively, one observation is a prefix of another if it stops observing earlier. Recall that a *true*-node indicates that observation stops (cf. Cond. (3)). Intuitively, we obtain a prefix of w_0 by cutting all transitions from some nodes (and adding them to *T*), while cutting just some transitions of a node is not allowed. It is easy to see that our definition of $\text{RT}(p)$ is closed under prefix:

Lemma 11. $(v_0 \leq w_0 \text{ and } w_0 \in \text{RT}(p)) \text{ implies } v_0 \in \text{RT}(p)$.

Proof. Let $w_0 \in \text{RT}(p)$ due to h and $v_0 \leq w_0$ with injection ρ . We define $h' : V \longrightarrow P$ such that $v \longmapsto h(\rho(v))$ and check that v_0 is a ready tree of p :

1. $h'(v) = h(\rho(v))$ is stable and not in *F* by Def. 7(1) for w_0 .
2. $p \xRightarrow{\varepsilon} h(w_0) = h(\rho(v_0)) = h'(v_0)$ by Def. 7(2) for w_0 .
3. $v \xrightarrow{a} v' \implies \rho(v) \xrightarrow{a} \rho(v') \implies h(\rho(v')) = h(\rho(v)) \in T \text{ or } h(\rho(v)) \xrightarrow{a} h(\rho(v')) \implies h'(v') = h(\rho(v')) \in T \text{ or } h'(v) \xrightarrow{a} h'(v')$.
4. Assume $v \notin T$ and $h'(v) \notin T$. Then, $\rho(v) \notin T$ by Def. 10(4) as well as $h(\rho(v)) \notin T$ by the definition of h' . This implies $\mathcal{I}(\rho(v)) = \mathcal{I}(h(\rho(v))) = \mathcal{I}(h'(v))$. Furthermore, $\mathcal{I}(v) = \mathcal{I}(\rho(v))$ by Defs. 10(2) and (3) and since $v \notin T$. Hence, $\mathcal{I}(v) = \mathcal{I}(h'(v))$. \square

Note that, since we do not want to distinguish isomorphic observation trees, we may always assume, without loss of generality, that the embedding ρ in Def. 10 is the identity, i.e., that the node set V of the prefix is a subset of W .

Lemma 12. $\{v_0 \mid \exists w_0 \in cRT(p). v_0 \leq w_0\} = RT(p)$.

Proof. Inclusion “ \subseteq ” is an application of Lemma 11; note that $cRT(p) \subseteq RT(p)$ by definition.

For proving the reverse inclusion “ \supseteq ”, let $v_0 \in RT(p)$ due to h . We construct a suitable w_0 such that the respective injection is the identity, by successively extending the T -nodes of v_0 . Let v_0 be the 0-extension of v_0 . Given the k -extension of v_0 we construct the $(k+1)$ -extension as follows:

For each $v \in T$ with $h(v) \in T$, remove v from T . For each $v \in T$ with $h(v) \notin T$, and every $a \in \mathcal{I}(h(v))$, choose some p' with $h(v) \xrightarrow{a} p' \notin F$ and p' stable. Such a p' exists since $h(v) \notin F$ (due to Def. 7(1)) implies, by the first fixed point condition of F for LTS (Def. 1(3)), the existence of a $p'' \notin F$ with $h(v) \xrightarrow{a} p''$. Moreover, by the second fixed point condition of F for LTS (Def. 1(4)), p'' can stabilise, i.e., there is a stable $p' \notin F$ such that $p'' \xrightarrow{\varepsilon} p'$. Now, choose a fresh node v_a and add $v \xrightarrow{a} v_a$ into the tree, with $h(v_a) = p'$ and $v_a \in T$. Remove v from T .

Note that the $(k+1)$ -extension is indeed a ready tree for p by construction. Finally, let w_0 be the component-wise union of all k -extensions with T set to the empty set. This yields a complete ready tree, i.e., $w_0 \in cRT(p)$; note in particular that our construction ensures that $h(v) \notin T \implies \mathcal{I}(v) = \mathcal{I}(h(v))$. \square

As a consequence of Lemma 12 we obtain the following corollary:

Corollary 13.

1. $RT(p)$ is uniquely determined by $cRT(p)$, and vice versa.
2. $RT(p) \subseteq RT(q) \iff cRT(p) \subseteq cRT(q)$
3. $fRT(p) = \{v_0 \text{ of finite depth} \mid \exists w_0 \in cRT(p). v_0 \leq w_0\}$

Before stating the next lemma we introduce the following definitions that allow us to approximate ready trees:

Definition 14 (k -Ready Tree). A k -tree $\langle V, \longrightarrow, T, \emptyset \rangle$, where $k \in \mathbb{N}_0$, is an observation tree where all nodes have depth at most k , and T is the set of all nodes of depth k . A k -ready tree of p is a ready tree of p that is also a k -tree. Moreover, $k\text{-RT}(p) =_{\text{df}} \{v_0 \in RT(p) \mid v_0 \text{ is a } k\text{-tree}\}$.

Intuitively, k -trees represent observations of k steps.

Definition 15 (Limit). Let \mathbf{v} be an infinite sequence $(v_k)_{k \in \mathbb{N}}$ where $v_k \in k\text{-RT}(p)$ and $v_k \leq v_{k+1}$, with the identity as injection, for all $k \in \mathbb{N}$. Then, $\lim \mathbf{v}$ is the component-wise union of all v_k with T set to empty; $\lim \mathbf{v}$ is called a *limit* of p .

Note that a node of some v_k in such a sequence is not in T_{k+1} , whence nodes in T are successively pushed out. In the limit, we may thus set T to empty. Moreover, if $v_k = v_{k+1} = v_{k+2} = \dots$ for some k , then the limit is v_k ; this happens exactly when v_k is complete. According to the following definition, we base the notion of finite branching on the weak transition relation $\xRightarrow{\alpha}$.

Definition 16 (Finite Branching). A process p is finite branching if, for every p' reachable from p , there are only finitely many $\langle \alpha, p'' \rangle$ with $p' \xRightarrow{\alpha} p''$.

For finite-branching processes p , $cRT(p)$ is characterised by the limits of p .

Lemma 17. *If p is finite branching, $cRT(p)$ equals the set of all limits of p .*

Proof. For proving inclusion “ \subseteq ”, let $w_0 \in cRT(p)$. We define a sequence $\mathbf{v} = (v_k)_{k \in \mathbb{N}}$ as follows: v_k consists of all nodes of w_0 of depth at most k , and the arcs between them. Moreover, T is the set of all nodes at depth k . Hence, v_k is a k -tree, and $v_k \leq w_0$ with the identity as injection. By Lemma 11, each v_k is in $RT(p)$. Obviously, $v_k \leq v_{k+1}$ for all $k \in \mathbb{N}$ and $w_0 = \lim(\mathbf{v})$.

For proving inclusion “ \supseteq ”, let $\mathbf{v} = (v_k)_{k \in \mathbb{N}}$ with $w_0 \in \lim(\mathbf{v})$. Hence, for each v_k we have at least one h_k such that $v_k \in RT(p)$ due to h_k . To show $w_0 \in cRT(p)$ we have to find a labelling $g : W \rightarrow P$ so that the definition of cRT is satisfied. This h will be assembled from the h_k ’s by an application of König’s Lemma.

We construct a graph with vertices $\langle v_k, h \rangle$ such that $v_k \in k\text{-}RT(p)$ due to h . Note again that there may be several such h . The edges of our graph are given by $\langle v_k, h \rangle \rightarrow \langle v_{k+1}, h' \rangle$ if $h = h'|_{V_k}$. Since p is finite branching we have only finitely many $\langle v_k, h \rangle$ for each k . Adding a root vertex r that is connected to all $\langle v_0, h \rangle$, we therefore obtain an infinite, finitely branching tree.

According to König’s Lemma, there exists an infinite path $\langle v_0, h_0 \rangle \rightarrow \langle v_1, h_1 \rangle \rightarrow \dots$. We now set $g =_{\text{df}} \bigcup_{k \in \mathbb{N}} h_k$. That g satisfies the conditions of the definition of $cRT(p)$ is obvious for Conds. (1–3); for Cond. (4), observe that each node w of w_0 has some depth k , hence it is in v_{k+1} and not in T_{k+1} , and we can use that v_{k+1} satisfies Cond. (4). \square

Note that the premise “ p is finite branching” is only needed for direction “ \supseteq ” in the above lemma. We may now obtain the following corollary of Cor. 13(3) and of Lemma 17, which is key to proving compositionality and full abstraction of our ready-tree preorder in the next section.

Corollary 18.

1. $cRT(p) \subseteq cRT(q) \implies fRT(p) \subseteq fRT(q)$, *always*.
2. $cRT(p) \subseteq cRT(q) \iff fRT(p) \subseteq fRT(q)$, *if p is finite branching*.

4 Compositionality & Full Abstraction

This section establishes our full-abstraction result of the ready-tree preorder \sqsubseteq with respect to the consistency-testing preorder \sqsubseteq , and proves that \wedge and \vee are indeed conjunction and, respectively, disjunction for \sqsubseteq . To do so, we first show that \wedge and \vee correspond to intersection and union on the semantic level, respectively. While the correspondence for \vee holds for ready trees in general, the correspondence for \wedge only holds for *complete* ready trees.

Theorem 19 (Set-Theoretic Interpretation of \wedge and \vee).

1. $cRT(p \wedge q) = cRT(p) \cap cRT(q)$
2. $RT(p \vee q) = RT(p) \cup RT(q)$




Fig. 6. Necessity of considering complete ready trees for conjunction.

Fig. 6 illustrates that Thm. 19(1) is invalid when considering all ready trees instead of complete read trees. The two processes displayed on the left and in the middle have the ready tree displayed on the right in common. However, the conjunction of the two processes is false and has no ready trees. Intuitively, the shown common ready tree formalises an observation that finished too early to encounter the inconsistency.

Proof. We first establish statement (1) of Thm. 19. For proving direction “ \subseteq ”, take $v_0 \in cRT(p \wedge q)$ due to h . We define $h'(v) =_{\text{df}} p'$ if $h(v) = p' \wedge q'$ for some q' , for all $v \in V$, and check the conditions of the definition of cRT :

1. If p' is instable (resp., in F), then $p' \wedge q'$ is instable (resp., in F), by our definition of \wedge (Def. (2)).
2. $p \wedge q \xRightarrow{\varepsilon} p' \wedge q'$ implies $p \xRightarrow{\varepsilon} p'$.
3. Let $v \xrightarrow{a} v'$. If $h(v') = h(v) = p' \wedge q' \in T$, then $h'(v) = h'(v') = p' \in T$, by Def. (2). If $h(v) = p' \wedge q' \xrightarrow{a} p'' \wedge q'' = h(v')$, then either $h'(v) = p' = p'' = h'(v') \in T$, or $h'(v) = p' \xrightarrow{a} p'' = h'(v')$; note that we can avoid F -processes along $p' \xrightarrow{a} p''$ since we can do so along $p' \wedge q' \xrightarrow{a} p'' \wedge q''$, and that $h'(v') = p'$ is stable since $h(v) = p' \wedge q'$ is stable.
4. Let $h'(v) \notin T$, and recall that $v \notin T$ always. Then, $h(v) = p' \wedge q' \notin T$ by Def. (2). Hence, $\mathcal{I}(v) = \mathcal{I}(h(v))$ by Def 7(4). Since q' is stable we have $\mathcal{I}(h(v)) = \mathcal{I}(h'(v))$, or $h'(v) \in T$.

Note that this direction of the theorem is also valid for RT in place of cRT .

For proving direction “ \supseteq ”, take $v_0 \in cRT(p) \cap cRT(q)$ due to h_1 and h_2 , respectively. Define $h(v) =_{\text{df}} h_1(v) \wedge h_2(v)$. We check the four conditions of the definition of cRT , starting with Cond. (4):

Let $v \notin T$ and $h(v) \notin T$. Without loss of generality, $h_1(v) \notin T$ according to our definition of \wedge (Def. (2)). If $h_2(v) \in T$, we have $\mathcal{I}(h(v)) = \mathcal{I}(h_1(v)) = \mathcal{I}(v)$ by Cond. (4) for h_1 . If $h_2(v) \notin T$, then $\mathcal{I}(v) = \mathcal{I}(h_1(v)) = \mathcal{I}(h_2(v))$ by Cond. (4) for h_1 and h_2 ; hence, $\mathcal{I}(v) = \mathcal{I}(h(v))$ according to the definition of $p \wedge q$.

The remaining conditions, Conds. (1)–(3), are proved together. Note that, since $h_1(v)$ and $h_2(v)$ are stable, we have that $h(v)$ is stable, too. We will prove simultaneously that a number of processes are not in F . To do so, we will collect a number of processes in a list \mathcal{F} and argue that the complement $\overline{\mathcal{F}}$ meets the conditions for F in \wedge (Def. 2). Then we know that the least set $F_{p \wedge q}$ satisfying these conditions is contained in $\overline{\mathcal{F}}$, whence no process on our list is in F .

We now simply show that the processes on the list do not satisfy any of the four conditions, using for the fixed point conditions (Defs. 2(3) and (4)) that no process in \mathcal{F} is in $\overline{\mathcal{F}}$.

Our list \mathcal{F} firstly contains all processes $p' \wedge q'$, so that p' (q') is a process along the derivation $p \xRightarrow{\varepsilon} h_1(v_0)$ ($q \xRightarrow{\varepsilon} h_2(v_0)$) according to Def. 7(2). Analogously, we treat p' on the subderivation $p'' \xRightarrow{\varepsilon} h_1(v')$ contained in $h_1(v) \xRightarrow{a} h_1(v')$ and q' on the subderivation $q'' \xRightarrow{\varepsilon} h_2(v')$ contained in $h_2(v) \xRightarrow{a} h_2(v')$ according to Def. 7(3), if both derivations exist. Finally, if $h_1(v) \xRightarrow{a} h_1(v')$ ($h_2(v) \xRightarrow{a} h_2(v')$) exists and $h_2(v') = h_2(v) \in T$ ($h_1(v') = h_1(v) \in T$), we combine each such p' (q') with $h_2(v)$ ($h_1(v)$).

We next show that $\overline{\mathcal{F}}$ is consistent with our constraints on F (Defs. 2(1)–(4)):

1. If $p' \in F$ or $q' \in F$, then $p' \wedge q'$ is not on the list, i.e., $p' \wedge q'$ is in $\overline{\mathcal{F}}$. In other words, if $p' \wedge q'$ is on the list, then $p' \notin F$ and $q' \notin F$ such that the first constraint on F is satisfied.
2. Assume $p' \wedge q'$ is on the list, and $p' \notin T$ and $q' \notin T$ and $p' \wedge q'$ stable. The last condition implies $p' \equiv h_1(v)$ and $q' \equiv h_2(v)$ for some v . Since $v \notin T$ by completeness of v_0 , we get $\mathcal{I}(h_1(v)) = \mathcal{I}(v) = \mathcal{I}(h_2(v))$ by Def. 7(4).
3. Assume $p' \wedge q'$ is on the list \mathcal{F} . If $p' \wedge q' \xrightarrow{\tau}$, then (without loss of generality) $p' \xrightarrow{\tau} p''$ for some p'' on the same derivation as p' . Hence, $p' \wedge q' \xrightarrow{\tau} p'' \wedge q'$ which is also on our list \mathcal{F} .

If $p' \wedge q' \not\xrightarrow{\tau}$, then $p' \equiv h_1(v)$ and $q' \equiv h_2(v)$ for some v . Let $a \in \mathcal{I}(h_1(v) \wedge h_2(v))$ and distinguish the following cases:

- $h_1(v) \in T$, i.e., $h_1(v) \equiv tt$: By Lemma 23, $h_1(v) \wedge h_2(v) \cong h_2(v)$. We must have $a \in \mathcal{I}(h_2(v))$, whence $h_2(v) \notin T$. Since $v \notin T$ by completeness of v_0 , we have $\mathcal{I}(v) = \mathcal{I}(h_2(v))$ by Def. 7(4). Thus, $v \xrightarrow{a} v'$ for some v' , and $h_2(v) \xrightarrow{a} q'' \xRightarrow{\varepsilon} h_2(v')$ by Def. 7(3). This implies $h_1(v) \wedge h_2(v) \xrightarrow{a} h_1(v) \wedge q''$ which is on our list \mathcal{F} .
 - $h_2(v) \in T$, i.e., $h_2(v) \equiv tt$: This case is analogous to the first one.
 - $h_1(v) \notin T$ and $h_2(v) \notin T$: We must have $a \in \mathcal{I}(h_1(v))$ and $a \in \mathcal{I}(h_2(v))$. As above, $v \xrightarrow{a} v'$ for some v' ; we conclude $h_1(v) \wedge h_2(v) \xrightarrow{a} p'' \wedge q''$ which is on our list \mathcal{F} .
4. Suppose $p' \wedge q'$ is on our list \mathcal{F} . Then, either $p' \wedge q' \not\xrightarrow{\tau}$, or $p' \wedge q'$ lies along the way to $h_1(v') \wedge h_2(v') \not\xrightarrow{\tau}$, using only processes on list \mathcal{F} .

This concludes the proof of Cond. (1). The validity of Cond. (2) is now immediate: by the above, $p \wedge q \xRightarrow{\varepsilon} h_1(v_0) \wedge h_2(v_0)$ since $p \xRightarrow{\varepsilon} h_1(v_0)$ and $q \xRightarrow{\varepsilon} h_2(v_0)$. To show the validity of Cond. (3), we consider $v \xrightarrow{a} v'$ and distinguish four cases, as suggested by Def. 7(3). We only show one case here as the others are equally easy: if $h_1(v) \xrightarrow{a} p' \xRightarrow{\varepsilon} h_1(v')$ and $h_2(v) \xrightarrow{a} q' \xRightarrow{\varepsilon} h_2(v')$, then $h_1(v) \wedge h_2(v) \xrightarrow{a} p' \wedge q' \xRightarrow{\varepsilon} h_1(v') \wedge h_2(v')$ using only processes not in F . This finishes the proof of statement (1) of Thm. 19.

The proof of statement (2) of Thm. 19 is much easier. For inclusion “ \supseteq ”, let $v_0 \in \text{RT}(p) \cup \text{RT}(q)$, whence (without loss of generality) $v_0 \in \text{RT}(p)$ due to h . Now, it is straightforward to check that $v_0 \in \text{RT}(p \vee q)$ due to h . We only note that Cond (2) of Def. 7 follows from $p \vee q \xrightarrow{\tau} p \xRightarrow{\varepsilon} h(v_0)$, due to $p \notin F$ by Lemma 8. For the reverse inclusion “ \subseteq ”, let $v_0 \in \text{RT}(p \vee q)$ due to h . By Def. 7(2), $p \vee q \xRightarrow{\varepsilon} h(v_0)$ and (without loss of generality) $p \vee q \xrightarrow{\tau} p \xRightarrow{\varepsilon} h(v_0)$. Obviously, $v_0 \in \text{RT}(p)$ by h . Note that, by Def. 7(2), all $h(v)$ are reachable from $h(v_0)$, whence h maps into P . \square

Compositionality of our conjunction and disjunction operators for \preceq is now an immediate consequence of Thm. 19

Theorem 20 (Compositionality).

1. $p \preceq q \implies p \wedge r \preceq q \wedge r$
2. $p \preceq q \implies p \vee r \preceq q \vee r$

Proof. The compositionality of \wedge follows from the following implication chain: $\text{RT}(p) \subseteq \text{RT}(q) \implies$ (by Cor. 13(2)) $\text{cRT}(p) \subseteq \text{cRT}(q) \implies \text{cRT}(p) \cap \text{cRT}(r) \subseteq \text{cRT}(q) \cap \text{cRT}(r) \implies$ (by Thm. 19(1)) $\text{cRT}(p \wedge r) \subseteq \text{cRT}(q \wedge r) \implies$ (by Cor. 13(2)) $\text{RT}(p \wedge r) \subseteq \text{RT}(q \wedge r)$. The compositionality of \vee can be proved analogously by referring to Thm. 19(2) instead of Thm. 19(1). \square

Thm. 19 also allows us to prove that \wedge and \vee really behave as conjunction and disjunction with respect to our refinement relation.

Theorem 21 (\wedge is And & \vee is Or).

1. $p \wedge q \preceq r \iff p \preceq r \text{ and } q \preceq r$
2. $r \preceq p \vee q \iff r \preceq p \text{ and } r \preceq q$

Proof. Part (1) follows from the following equivalences: $\text{RT}(p \wedge q) \subseteq \text{RT}(r) \iff$ (by Cor. 13(2)) $\text{cRT}(p \wedge q) \subseteq \text{cRT}(r) \iff$ (by Thm. 19(1)) $\text{cRT}(p) \cap \text{cRT}(q) \subseteq \text{cRT}(r) \iff \text{cRT}(p) \subseteq \text{cRT}(r) \text{ and } \text{cRT}(q) \subseteq \text{cRT}(r) \iff$ (by Cor. 13(2)) $\text{RT}(p) \subseteq \text{RT}(r) \text{ and } \text{RT}(q) \subseteq \text{RT}(r)$. Again, for Part (2) we get a similar but simpler proof by referring to Thm. 19(2) instead of Thm. 19(1). \square

In order to see that ready trees are indeed fully-abstract with respect to our naive consistency preorder, it now suffices to prove that \preceq coincides with our consistency testing preorder. This means that \preceq is *the* adequate preorder in our setting of logical LTS with conjunction and disjunction.

Theorem 22 (Full Abstraction). $\sqsubseteq = \sqsubset$

For the proof of this theorem, the following technical lemma will be convenient.

Lemma 23. *The LTSs of $tt \wedge p$ and p are isomorphic, in signs $tt \wedge p \cong p$.*

Proof. It is easy to check that $tt \wedge p' \mapsto p'$ is an isomorphism. Note that with $tt \wedge p' \in F$ iff $p' \in F$, the fixed point conditions of F are satisfied. \square

We may now proof Thm. 22.

Proof. We first prove the easier direction “ \supseteq ”. If $p \sqsubset q$ then, for all o , we have $p \wedge o \sqsubset q \wedge o$ by Thm. 20. If $p \wedge o \in F$, then $\text{RT}(p \wedge o) = \emptyset$ by Lemma 8. Hence, $\text{RT}(q \wedge o) = \emptyset$, i.e., $q \wedge o \in F$ by Lemma 8 again. Thus, $p \sqsubseteq q$.

For proving the reverse inclusion, let $p \sqsubseteq q$ and $v_0 \in \text{RT}(q)$ due to h . Note that v_0 is a process itself. We show that $q \wedge v_0 \notin F$. To do so, we use a fixed point argument similar to the one in the proof of Thm. 19. Here, our list firstly includes processes $q' \wedge v_0$, with q' on the derivation $q \xRightarrow{\varepsilon} h(v_0)$ (cf. Def. 7(3)), as well as processes $q' \wedge v'$, with q' along the derivations $q'' \xRightarrow{\varepsilon} h(v')$ that emerge due to $h(v) \xrightarrow{\alpha} q'' \xRightarrow{\varepsilon} h(v')$ in Def. 7(3). Furthermore, we include all processes $h(v') \wedge v'$ on our list, whenever $h(v') = h(v) \in T$ in Def. 7(3). Finally, we include all $q' \wedge v'$ such that $q' \notin F$, $v' \in T$ and q' is a process in the same LTS as q .

We now check that the complement of this list is a fixed point, i.e., satisfies Conds. (1)–(4) of Def. 2. Let $q' \wedge v'$ be on our list:

1. According to the definition of observation trees (Def. 6), $v' \notin F$. Moreover, $q' \notin F$ by the definition of $\xRightarrow{\varepsilon}$, or $q' \in T$ or ‘immediate’, depending on why $q' \wedge v'$ has been included in the list.
2. If $q' \notin T$ and $v' \notin T$ and $q' \wedge v' \not\rightarrow$, then $q' \equiv h(v')$, as these are the only stable processes on our list. Hence, we may apply Def. 7(4) for v' to obtain $\mathcal{I}(q') = \mathcal{I}(v')$.
3. Let $\alpha \in \mathcal{I}(q' \wedge v')$. If $\alpha = \tau$, then $q' \xrightarrow{\tau} q''$ on the respective derivation according to the definition of our list. Then, $q' \wedge v' \xrightarrow{\tau} q'' \wedge v'$ which is on our list as well. If $\alpha \neq \tau$, then q' is stable, whence $q' \equiv h(v')$ again or $v' \in T$. We proceed by a case distinction:
 - $q' \in T$: Hence, $v' \xrightarrow{\alpha} v''$ for some v'' and $v' \notin T$. Since $q' \not\rightarrow$, we are in the case of Def. 7(3a), i.e., $h(v'') = h(v') = q'$. Moreover, $q' \wedge v' \xrightarrow{\alpha} q' \wedge v''$ which is on our list.
 - $v' \in T$: Hence, $q' \xrightarrow{\alpha} q''$ for some $q'' \notin F$. Then, $q' \wedge v' \xrightarrow{\alpha} q'' \wedge v'$ which is on our list.
 - $q' \notin T$ and $v' \notin T$: Hence, $v' \xrightarrow{\alpha} v''$ for some v'' . Since $h(v') = q' \notin T$ we know, by Def. 7(3), that $q' \xrightarrow{\alpha} q'' \xRightarrow{\varepsilon} h(v'')$. Thus, $q' \wedge v' \xrightarrow{\alpha} q'' \wedge v''$ which is on our list.
4. If $q' \wedge v'$ is on the list due to $q' \xRightarrow{\varepsilon} h(v')$, then $q' \wedge v' \xRightarrow{\varepsilon} h(v') \wedge v'$ along processes that are on our list, and $h(v') \wedge v'$ is stable since $h(v')$ is stable and v' is trivially stable.

If $q' \wedge v'$ is on the list due to $v' \in T$, then q' can stabilise. Thus, $q' \wedge v'$ can stabilise in an isomorphic way using only processes on the list.

Thus, we have established $q \wedge v_0 \notin F$. This implies by $p \sqsubseteq q$ that $p \wedge v_0 \notin F$. To show $v_0 \in \text{RT}(p)$, we will construct a respective labelling g according to depth. Since $p \wedge v_0 \notin F$, process $p \wedge v_0$ can stabilise by Def. 1(4) with $p \wedge v_0 \xRightarrow{\varepsilon} p' \wedge v_0$ for some stable p' with $p \xRightarrow{\varepsilon} p'$. We define $g(v_0) =_{\text{df}} p'$, so that Cond. (2) of Def. 7 is satisfied, and Cond. (1) of Def. 7 holds for v_0 . Note that $g(v_0) \wedge v_0 \equiv p' \wedge v_0 \notin F$ by the definition of $\xRightarrow{\varepsilon}$.

Assume that g is defined up to depth k such that Conds. (1), (3) and (4) of Def. 7 hold for all v with depth less than k , and that Cond. (1) is satisfied for depth k as well. Moreover, assume that $g(v) \wedge v \notin F$ whenever $g(v)$ is defined. These assumptions are our induction hypothesis, which we have just checked for $k = 0$. For each v at depth k we proceed as follows. If $v \in T$, then Conds. (3) and (4) are vacuously true; since v has no children at depth $k + 1$, we are done. Thus, let $v \notin T$ and distinguish the following cases:

- $g(v) \in T$: For all v' with $v \xrightarrow{a} v'$, define $g(v') =_{\text{df}} g(v)$. Thus, Conds. (3) and (4) are satisfied for v , and $g(v')$ is stable and not in F . Since $g(v) \in T$ we have, by a consequence of Lemma 23, that $g(v) \wedge v$ and v are isomorphic. As $v \xrightarrow{a} v' \notin F$, this implies $g(v') \wedge v' \equiv g(v) \wedge v' \notin F$.
- $g(v) \notin T$: We first show Cond. (4) whose premise is now true. As $g(v) \wedge v \notin F$, the facts $v \notin T$, $g(v) \notin T$ and $g(v) \wedge v$ stable imply $\mathcal{I}(v) = \mathcal{I}(g(v))$. Next, we show Cond. (3), i.e., how to extend g to all v' with $v \xrightarrow{a} v'$ such that $g(v) \xRightarrow{a} g(v')$. Consider some $v \xrightarrow{a} v'$. Since $\mathcal{I}(v) = \mathcal{I}(g(v))$ we know $a \in \mathcal{I}(g(v) \wedge v)$ and, since $g(v) \wedge v \notin F$, there is some p' such that $g(v) \wedge v \xrightarrow{a} p' \wedge v' \notin F$. But $p' \wedge v' \notin F$ implies that $p' \wedge v'$ can stabilise according to Def. 1(4), with some derivation $p' \wedge v' \xRightarrow{\varepsilon} p'' \wedge v' \xrightarrow{\varepsilon}$. In particular, p'' is stable and not in F . In addition, $g(v) \xrightarrow{a} p' \xRightarrow{\varepsilon} p''$. Note that all processes along the derivation $p' \xRightarrow{\varepsilon} p''$ are not in F since, otherwise, some process along $p' \wedge v' \xRightarrow{\varepsilon} p'' \wedge v'$ would be in F . Finally, we now define $g(v') =_{\text{df}} p''$.

When applying this construction to all v' with $v \xrightarrow{a} v'$, Cond. (3) is satisfied for v . Furthermore, Cond. (1) is satisfied for all v' , and $g(v') \wedge v' \notin F$.

Treating all v at level k in the above way, we extend g to depth $k + 1$ such that the induction hypothesis now also holds for $k + 1$. With this induction, we can thus define g for each v in the observation tree. Hence, $v_0 \in \text{RT}(p)$ due to g . \square

The following proposition states the validity of several boolean properties desired of conjunction and disjunction operators. Here, $=$ denotes the kernel of our consistency testing preorder (ready–tree preorder).

Proposition 24 (Properties of \wedge and \vee).

<i>Commutativity:</i>	$p \wedge q = q \wedge p$	$p \vee q = q \vee p$
<i>Associativity:</i>	$(p \wedge q) \wedge r = p \wedge (q \wedge r)$	$(p \vee q) \vee r = p \vee (q \vee r)$
<i>Idempotence:</i>	$p \wedge p = p$	$p \vee p = p$
<i>False:</i>	$p \wedge \text{ff} = \text{ff}$	$p \vee \text{ff} = p$
<i>True:</i>	$p \wedge \text{tt} = p$	$p \vee \text{tt} = \text{tt}$
<i>Distributivity:</i>	$p \wedge (q \vee r) = (p \wedge q) \vee (p \wedge r)$	$p \vee (q \wedge r) = (p \vee q) \wedge (p \vee r)$

Proof. All the above properties are straightforward since \wedge (\vee) on processes corresponds to \cap (\cup) on complete ready trees by Thm. 19, and since complete ready trees and ready trees induce the same preorder by Cor. 13(2). Moreover, $\text{cRT}(\text{ff})$ is the empty set of complete observation trees, while $\text{cRT}(tt)$ is the set of all complete observation trees. \square

Proposition 25 (Relating \wedge , \vee to \sqsubseteq).

1. $p \wedge q = q \iff p \sqsubseteq q$
2. $p \vee q = p \iff p \sqsubseteq q$

Proof. Again, the statements follow directly from Thm. 19 and Cor. 13(2). \square

We conclude this section by briefly returning to the illustrative processes *spec* and *impl* of Fig. 4. We have already remarked that the only complete ready tree of the latter is also one of the former. Hence, by Thm. 22, *impl* is indeed a refinement of *spec* according to our ready-tree preorder. Considering the conjunction of these processes, also shown in Fig. 4, it might be easier to see this using Prop. 25(1).

5 Related Work

Traditionally, process-algebraic and temporal-logic formalisms are not mixed but co-exist side by side. Indeed, the process-algebra school often uses synchronous composition and internal choice to model conjunction and disjunction, respectively. The compositionality of classic process-algebraic refinement preorders, such as failures semantics [4] and must-testing [11], enables component-based reasoning. However, inconsistencies in specifications are not captured so that, e.g., the conjunctive composition of *a* and *b* is identified with deadlock rather than *ff*. In contrast, the temporal-logic school distinguishes between deadlock and *ff* but does not support component-based refinement.

Much research on mixing operational and logical styles of specification avoids dealing with inconsistencies by translating one style into the other. On the one hand, operational content may be translated into logic formulas, as is implicitly done in Lamport's TLA [10] or in the work of Graf and Sifakis [8]. In these approaches, logical implication serves as refinement relation. On the other hand, logical content may be translated into operational content. This is the case in automata-theoretic work, such as Kurshan's work on ω -automata [9], which includes synchronous and asynchronous composition operators and uses maximal trace inclusion as refinement relation. However, both logical implication and trace inclusion are insensitive to deadlock and thus do not support component-based reasoning.

A seminal approach to compositional refinement relations in a mixed setting was proposed by Olderog in [12], where process-algebraic constructs are combined with trace formulas expressed in a predicate logic. In this approach,

trace formulas can serve as processes, but not vice versa. Thus, freely mixing operational and logical styles is not supported and, in particular, conjunction cannot be applied to processes. For his setting, Olderog develops a denotational semantics that is a slight variation of standard failures semantics. Remarkably, an inconsistent formula is given a semantics that is not an element of the appropriate domain, as stated on pp. 172–173 of [12].

Recently, a more general approach to combining process–algebraic and temporal–logic approaches was proposed in two papers by Cleaveland and Lüttgen [5, 6], which adopt a variant of De Nicola and Hennessy’s must–testing preorder [11] as refinement preorder. However, Cleaveland and Lüttgen have not successfully solved the challenge of defining a semantics that is both deadlock–sensitive and compositional, and in which the conjunction operator and the refinement relation are compatible in the sense of Prop. 25(1). Our work solves this problem in the basic setting of logical LTS. Key for the solution is our new understanding of inconsistency, which is reflected by the fact that we consider processes a and $a + b$ as inconsistent, whereas they were treated as consistent in [6]. Observe that also in failure semantics and must–testing, a and $a + b$ are inconsistent in the sense that they do not have a common implementation.




Fig. 7. Backward propagation of inconsistency.

In addition, our backward propagation of inconsistency, as formalised in Def. 1(3), is in line with traditional semantics, as illustrated in Fig. 7. The first conjunct would be a specification of the second conjunct with respect to failures semantics and must–testing, whence their conjunction, also shown in Fig. 7, should be consistent. In fact, the conjunction equals the second process in our ready–tree semantics.

Comparing Ready–Tree Semantics to Other Semantics

To the best of our knowledge, ready–tree semantics is novel and has not been studied in the literature before. We thus briefly compare it to three popular semantics, namely *ready–trace semantics*, *failures semantics* and *ready simulation* [7]. Since our treatment of divergence is different from the one of failures semantics, we restrict our discussion to τ –free processes.

A *ready trace* [1] of a process is a sequence of actions that it can perform and where, at the beginning of the trace, between any two actions and at the end, the ready set of the process reached at the respective stage is inserted. Such a




Fig. 8. Ready-tree semantics is strictly finer than ready-trace semantics.

ready trace can be understood as a particular type of ready tree that consists only of a single path and includes additional transitions representing the ready sets. These additional transitions ensure that each state on the path has, for each action in its ready set, exactly one transition that either belongs to the path or ends in a *true*-state. For example, the first ready tree in Fig. 5 in Sec. 3 represents the ready trace $\{a, b\}b\{b\}b\{a, b\}$. Consequently, the ready traces of a process can be read off from its ready trees, and ready-tree inclusion implies ready-trace inclusion. The reverse implication does not hold as demonstrated by the two left-most processes in Fig. 8. These possess the same ready traces; however, the observation tree on the right-hand side is a ready tree of the first, but not of the second process.

The *failures semantics* of a process is the set of its refusal pairs. Such a pair consists of a trace followed by a refusal set, i.e., a set of actions that the process reached by the trace cannot perform. Such a refusal pair can be read off from the respective ready trace by deleting all its ready sets and adding a set of actions having an empty intersection with the last ready set on the trace. Thus, ready-tree semantics is finer than failures semantics.




Fig. 9. The ready-tree preorder is strictly coarser than ready simulation.

A process q *ready-simulates* some process p if there exists a simulation relation from p to q such that related states have identical ready sets. When tracing a ready tree of p , it is easy to see that such a simulation translates this ready tree to the same ready tree for q . Thus, the ready-tree preorder is coarser than

ready simulation. Fig. 9 shows that it is indeed *strictly* coarser. Both processes displayed have the same ready trees, all of which are paths. However, the second process cannot even simulate the first one.

6 Conclusions & Future Work

This paper introduced a new semantics, the *ready-tree semantics*, that lies between ready-trace semantics and ready simulation. Our framework was one of τ -pure LTSs, with distinguished *true*- and *false*-states, and is equipped with *conjunction* and *disjunction* operators. Key for defining the conjunction operator was the careful, inductive formalisation of an inconsistency predicate. The implied ready-tree preorder proved to be compositional and fully-abstract with respect to a naive preorder that allows inconsistent specifications to be refined only by inconsistent implementations. Standard laws of boolean algebra hold as expected, due to the fact that conjunction and disjunction on LTSs correspond to intersection and union on ready trees, respectively.

Consequently, this paper solves the problems of defining conjunction which are reported in closely related work [5, 6], albeit in a simpler setting that does not consider process-algebraic operators but only conjunction and disjunction. However, it is the simplicity of our setting that brought the subtleties of defining a fully-abstract semantics in the presence of conjunction to light, and which offered a way forward in addressing the challenge of defining “logical” process calculi, i.e., process calculi that allow one to freely mix process-algebraic and temporal-logic operators [6].

Future work shall extend our results to richer frameworks. Firstly, we plan to lift our requirement of τ -purity on LTS and extend our framework by standard process-algebraic operators such as parallel composition, hiding and recursion. In particular hiding is likely to prove challenging due to its transformation of observable infinite behaviour into divergent behaviour. Secondly, our framework shall be extended semantically from LTS to Büchi LTS [5] so that one may express liveness and fairness properties, and syntactically to linear-time temporal-logic formulas [6]. Last, but not least, we wish to explore tool support.

Acknowledgements. We thank Rance Cleaveland for many fruitful discussions and particularly for suggesting the use of an inconsistency predicate.

References

- [1] J.C.M. Baeten, J.A. Bergstra, and J.W. Klop. Ready-trace semantics for concrete process algebra with the priority operator. *Computer J.*, 30(6):498–506, 1987.
- [2] J.A. Bergstra, A. Ponse, and S.A. Smolka. *Handbook of Process Algebra*. Elsevier Science, 2001.
- [3] B. Bloom, S. Istrail, and A.R. Meyer. Bisimulation can’t be traced. *J. ACM*, 42(1):232–268, 1995.
- [4] S.D. Brookes, C.A.R. Hoare, and A.W. Roscoe. A theory of communicating sequential processes. *J. ACM*, 31(3):560–599, 1984.

- [5] R. Cleaveland and G. Lüttgen. A semantic theory for heterogeneous system design. In *FSTTCS 2000*, vol. 1974 of *LNCS*, pp. 312–324. Springer-Verlag, 2000.
- [6] R. Cleaveland and G. Lüttgen. A logical process calculus. In *EXPRESS 2002*, vol. 68,2 of *ENTCS*. Elsevier Science, 2002.
- [7] R. van Glabbeek. The linear time – branching time spectrum II. In *CONCUR '93*, vol. 715 of *LNCS*, pp. 66–81. Springer-Verlag, 1993.
- [8] S. Graf and J. Sifakis. A logic for the description of non-deterministic programs and their properties. *Information and Control*, 68(1–3):254–270, 1986.
- [9] R.P. Kurshan. *Computer-Aided Verification of Coordinating Processes: The Automata-Theoretic Approach*. Princeton Univ. Press, 1994.
- [10] L. Lamport. The temporal logic of actions. *TOPLAS*, 16(3):872–923, 1994.
- [11] R. De Nicola and M.C.B. Hennessy. Testing equivalences for processes. *TCS*, 34:83–133, 1983.
- [12] E.R. Olderog. *Nets, Terms and Formulas*. Cambridge Tracts in Theoretical Computer Science 23. Cambridge Univ. Press, 1991.
- [13] A. Pnueli. The temporal logic of programs. In *FOCS '77*, pp. 46–57. IEEE Computer Society Press, 1977.