



HAL
open science

Constructive characterisations of the must-preorder for asynchrony

Giovanni Bernardi, Ilaria Castellani, Paul Laforgue, Léo Stefanescu

► **To cite this version:**

Giovanni Bernardi, Ilaria Castellani, Paul Laforgue, Léo Stefanescu. Constructive characterisations of the must-preorder for asynchrony. 2025. hal-04642776v2

HAL Id: hal-04642776

<https://hal.science/hal-04642776v2>

Preprint submitted on 17 Jan 2025

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Constructive characterisations of the must-preorder for asynchrony

Giovanni Bernardi¹[0009–0008–3653–3040], Ilaria Castellani²[0000–0001–9820–0892], Paul Laforgue^{1,3}[0009–0009–6688–4850], and Léo Stefanescu⁴[0000–0002–4719–2922]*

¹ Université Paris Cité, CNRS, IRIF, F-75013, Paris, France

² INRIA, Université Côte d’Azur, France

³ Nomadic Labs, Paris, France

⁴ University of Cambridge, United Kingdom

Abstract. De Nicola and Hennessy’s MUST-preorder is a liveness preserving refinement which states that a server q refines a server p if all clients satisfied by p are also satisfied by q . Owing to the universal quantification over clients, this definition does not yield a practical proof method, and alternative characterisations are necessary to reason over it. Finding these characterisations for asynchronous semantics, i.e. where outputs are non-blocking, has thus far proven to be a challenge, usually tackled via ad-hoc definitions.

We show that the standard characterisations of the MUST-preorder carry over as they stand to asynchronous communication, if servers are enhanced to act as forwarders, *i.e.* they can input any message as long as they store it back into the shared buffer. Our development is constructive, is completely mechanised in Coq, and is independent of any calculus: our results pertain to Selinger output-buffered agents with feedback. This is a class of Labelled Transition Systems that captures programs that communicate via a shared unordered buffer, as in asynchronous CCS or the asynchronous π -calculus.

We show that the standard coinductive characterisation lets us prove in Coq that concrete programs are related by the MUST-preorder.

Finally, our proofs show that Brouwer’s bar induction principle is a useful technique to reason on liveness preserving program transformations.

1 Introduction

Code refactoring is a routine task to develop or update software, and it requires methods to ensure that a program p can be safely replaced by a program q . One way to address this issue is via refinement relations, *i.e.* preorders. For programming languages, the most well-known one is Morris *extensional* preorder [75, pag. 50], defined by letting $p \leq q$ if for all contexts C , whenever $C[p]$ reduces to a normal form N , then $C[q]$ also reduces to N .

* Work done at MPI-SWS, Germany.

1. INTRODUCTION

Comparing servers This paper studies a version of Morris preorder for nondeterministic asynchronous client-server systems. In this setting it is natural to reformulate the preorder by replacing reduction to normal forms (*i.e.* termination) with a suitable *liveness* property. Let $p \parallel r$ denote a *client-server system*, that is a parallel composition in which the identities of the server p and the client r are distinguished, and whose computations have the form $p \parallel r = p_0 \parallel r_0 \longrightarrow p_1 \parallel r_1 \longrightarrow p_2 \parallel r_2 \longrightarrow \dots$, where each step represents either an internal computation of one of the two components, or an interaction between them. Interactions correspond to handshakes, where two components ready to perform matching input/output actions advance together. We express liveness by saying that p *must pass* r , denoted $p \text{ MUST } r$, if in every maximal computation of $p \parallel r$ there exists a state $p_i \parallel r_i$ such that $\text{GOOD}(r_i)$, where GOOD is a decidable predicate indicating that the client has reached a successful state. Servers are then compared according to their capacity to satisfy clients, *i.e.* via contexts of the form $[-] \parallel r$ and the predicate MUST . Morris preorder then becomes the MUST -preorder by De Nicola and Hennessy [43]: $p \sqsubseteq_{\text{MUST}} q$ when $\forall r. p \text{ MUST } r$ implies $q \text{ MUST } r$.

Advantages The MUST -preorder is by definition liveness preserving, because $p \text{ MUST } r$ literally means that “in every execution something good must happen (on the client side)”. Results on $\sqsubseteq_{\text{MUST}}$ thus shed light on liveness-preserving program transformations.

The MUST -preorder is independent of any particular calculus, as its definition requires simply (1) a reduction semantics for the parallel composition $p \parallel r$, and (2) a predicate GOOD over programs. Hence $\sqsubseteq_{\text{MUST}}$ may relate servers written in different languages. For instance, servers written in OCAML may be compared to servers written in JAVA according to clients written in PYTHON, because all these languages use the same basic protocols for communication.

Drawback The definition of the MUST -preorder is *contextual*: proving $p \sqsubseteq_{\text{MUST}} q$ requires analysing an *infinite* amount of clients, and so the definition of the preorder does not entail an effective proof method. A solution to this problem is to define an *alternative (semantic) characterisation* of the preorder $\sqsubseteq_{\text{MUST}}$, *i.e.* a preorder \preceq_{alt} that coincides with $\sqsubseteq_{\text{MUST}}$ and does away with the universal quantification over clients (*i.e.* contexts). In *synchronous* settings, *i.e.* when both input and output actions are blocking, such alternative characterisations have been thoroughly investigated, typically via a behavioural approach based on labelled transition systems.

Labelled transition systems A program p is associated with a labelled transition system (LTS) representing its behaviour, which we denote by $\text{LTS}(p)$. Figure 1 presents two instances of LTSs, where transitions are labelled by input actions such as str , output actions such as $\overline{\text{str}}$, or the internal action τ (not featured in Figure 1), while dotted nodes represent successful states, *i.e.* those satisfying the predicate GOOD . There, the server p_0 is ready to input either a string or a float. It is the environment that, by offering an output of either type, will make p move to either p_1 or p_2 . The client r_0 , on the other hand, is ready




Fig. 1. The behaviours of a server p_0 and of a client r_0 .

to either output a string, or input an integer. The input int makes the client move to the successful state r_2 , while the output $\overline{\text{str}}$ makes the client move to the state r_1 , where it can still perform the input int to reach the successful state r_3 . In an asynchronous setting, output transitions enjoy a commutativity property on which we will return later. Programs p are usually associated with their behaviours $LTS(p)$ via inference rules that we omit in the main body of the paper, as they are standard.

Alternative preorders for synchrony Program behaviours, *i.e.* LTSs, are used to define the alternative preorders for $\sqsubseteq_{\text{MUST}}$ following one of two different approaches: MUST-sets or acceptance sets.

Both approaches were originally proposed for Milner's Calculus of Communicating Systems (CCS) [74], where communication is *synchronous*. The first alternative preorder, which we denote by \preceq_{MS} , was put forth by De Nicola [43], and it compares server behaviours according to their MUST-sets, *i.e.* the sets of actions that they may perform after doing a given sequence of actions. The second alternative preorder, which we denote by \preceq_{AS} , was put forth by Hennessy [55], and it compares the acceptance sets of servers, *i.e.* how servers can be moved out of their potentially deadlocked states, namely, states from which the servers cannot evolve autonomously. Both these preorders characterise $\sqsubseteq_{\text{MUST}}$ in the following sense:

$$\forall p, q \in \text{CCS}. p \sqsubseteq_{\text{MUST}} q \text{ iff } LTS(p) \preceq_{\text{MS}} LTS(q) \quad (1)$$

$$\forall p, q \in \text{CCS}. p \sqsubseteq_{\text{MUST}} q \text{ iff } LTS(p) \preceq_{\text{AS}} LTS(q) \quad (2)$$

While these alternative preorders do away with the universal quantification over clients, they are not practical to use directly, as they still universally quantify over (finite) traces of actions. A more practical approach [2] is to use a coinductively defined preorder \preceq_{co} based on \preceq_{AS} [2,72,16]. This preorder has two advantages: first, its definition quantifies universally only on single actions; second, it allows the user to use standard coinductive methods, as found in the literature on bisimulation. In the case where the LTS is image-finite, such as for CCS and most process calculi, the coinductive preorder is sound and complete:

$$\forall p, q \in \text{CCS}. p \sqsubseteq_{\text{MUST}} q \text{ iff } LTS(p) \preceq_{\text{co}} LTS(q) \quad (3)$$

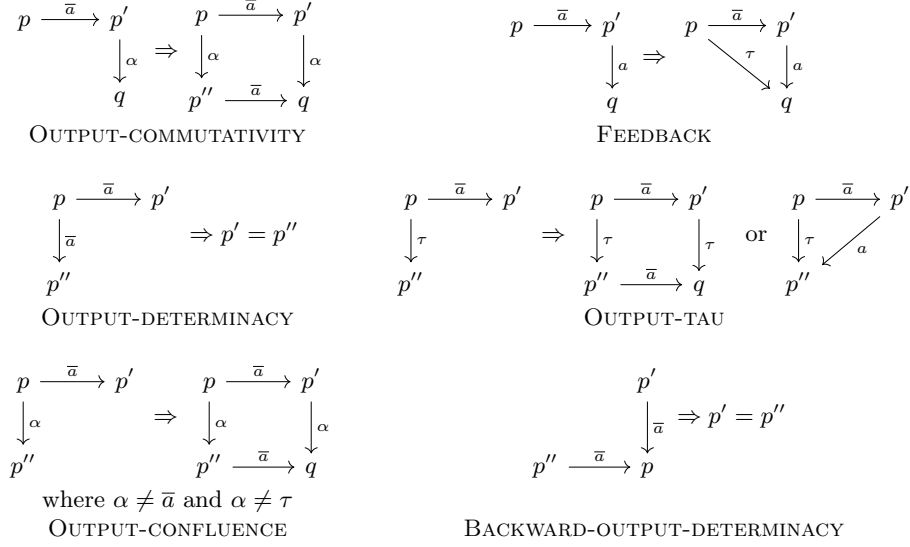


Fig. 2. First-order axioms for output-buffered agents with feedback as given by Selinger [90], extended with the BACKWARD-OUTPUT-DETERMINACY axiom.

Asynchrony In distributed systems, communication is inherently asynchronous. For instance, the standard TCP transmission on the Internet is asynchronous. Actor languages like ELIXIR and ERLANG implement asynchrony via mailboxes, and both PYTHON and JAVASCRIPT offer developers the constructs ASYNC/WAIT, to return promises (of results) or wait for them. In this paper we model asynchrony via *output-buffered agents with feedback*, as introduced by Selinger [90]. These are LTSs obeying the axioms in Figure 2, where a denotes an input action, \bar{a} denotes an output action, τ denotes the internal action, and α ranges over all these actions. For instance, the OUTPUT-COMMUTATIVITY axiom states that an output \bar{a} can always be postponed: if \bar{a} is followed by any action α , it can commute with it. In other words, outputs are non-blocking, as illustrated by the LTS for r_0 in Figure 1. We defer a more detailed discussion of these axioms to Section 2.

Technical difficulties The practical importance of asynchrony motivates a specific study of $\sqsubseteq_{\text{MUST}}$. Efforts in this direction have been made, all of which focussed on process calculi [38,23,94,57], while the axioms in Figure 2 apply to LTSs. Note that these axioms impose conditions only over outputs, and this asymmetric treatment of inputs and outputs substantially complicates the proofs of completeness and soundness of the alternative characterisations of $\sqsubseteq_{\text{MUST}}$. To underline the subtleties due to asynchrony, we note that the completeness result for asynchronous CCS given by Castellani and Hennessy in [38], and subsequently extended to the π -calculus by Hennessy [57], is false (see Appendix I).

Contributions and paper structure.

Our main contributions may be summarised as follows (where for each of them, we indicate where it is presented in the paper):

- The first behavioural characterisations of the MUST-preorder (Theorem 1, Theorem 3) that are calculus independent, in that both our definitions and our proofs work directly on LTSs. Contrary to all the previous works on the topic, we show that the *standard* alternative preorders characterise the MUST-preorder also in Selinger asynchronous setting. To this end, it suffices to enrich the server semantics with *forwarding*, i.e. ensure that servers are ready to input any message, as long as they store it back in a global shared buffer. This idea, although we use it here in a slightly different form, was pioneered by Honda et al. [64]. In this paper we propose a construction that works on any LTS (Lemma 2) and we show the following counterparts of Equations (1), (2), and (3) where FDB denotes the LTSs of output-buffered agents with feedback, and FW is the function that enhances them with forwarding:

$$\forall p, q \in \text{FDB}. p \sqsubseteq_{\text{MUST}} q \text{ iff } \text{FW}(p) \preceq_{\text{MS}} \text{FW}(q) \quad (\text{a})$$

$$\forall p, q \in \text{FDB}. p \sqsubseteq_{\text{MUST}} q \text{ iff } \text{FW}(p) \preceq_{\text{AS}} \text{FW}(q) \quad (\text{b})$$

$$\forall p, q \in \text{FDB}. p \sqsubseteq_{\text{MUST}} q \text{ iff } \text{FW}(p) \preceq_{\text{co}} \text{FW}(q) \quad (\text{c})$$

Quite surprisingly, the alternative preorders \preceq_{AS} , \preceq_{MS} and \preceq_{co} need not be changed. We present these results in Section 3. We use the coinductive preorder \preceq_{co} to prove the correctness of a form of code hoisting (6).

- The first characterisations of the MUST-preorder that fully exploit asynchrony, *i.e.* disregard irrelevant (that is, non-causal) orders of visible actions in traces (Corollary 4).
- The first constructive account of the MUST-preorder. We show that if the MUST and termination predicates are defined *intensionally* (in the sense of Brede and Herbelin [28]), then $\sqsubseteq_{\text{MUST}}$ can be characterised constructively. The original definitions of MUST and termination given by De Nicola [43], though, are *extensional*. We show how to use Brouwer bar induction principle to prove that the two approaches are logically equivalent (Corollary 2). Since Rahli et al. [81] have shown bar induction to be compatible with constructive type theory, we argue that our development is entirely constructive.
- The first mechanisation of the theory of MUST-preorder in a fully nondeterministic setting, which consists of around 8000 lines of Coq. In Appendix J we gather the Coq versions of all the definitions and the results presented in the main body of the paper.

In Section 6, we discuss the impact of the above contributions, as well as related and future work. In Section 2, we recall the necessary background definitions and illustrate them with a few examples.

```

Class Sts (A: Type) := MkSts {
  sts_step: A → A → Prop;
  sts_stable: A → Prop; }.

Inductive ExtAct (A: Type) :=      Inductive Act (A: Type) :=
| ActIn (a: A) | ActOut (a: A).   | ActExt (ext: ExtAct A) | τ.

Class Label (L: Type) :=          Class Lts (A L : Type) {Label L} :=
MkLabel {                          MkLts {
  label_eqdec: EqDecision L;       lts_step: A → Act L → A → Prop;
  label_countable: Countable L;    lts_outputs: A → finite_set L;
                                   lts_performs: A → (Act L) → Prop; }.

```

Fig. 3. Highlights of our Sts and Lts typeclasses.

2 Preliminaries

We model individual programs such as servers p and clients r as LTSs obeying Selinger axioms, while client-server systems $p \parallel r$ are modelled as state transition systems with a reduction semantics. We now formally define this two-level semantics.

Labelled transition systems A *labelled transition system* (LTS) is a triple $\mathcal{L} = \langle A, L, \longrightarrow \rangle$ where A is the set of states, L is the set of labels and $\longrightarrow \subseteq A \times L \times A$ is the transition relation. When modelling programs as LTSs, we use transition labels to represent program actions. The set of labels in Selinger LTSs has the same structure as the set of actions in Milner’s calculus CCS: one assumes a set of names \mathcal{N} , denoting input actions and ranged over by a, b, c , a complementary set of conames $\overline{\mathcal{N}}$, denoting output actions and ranged over by $\overline{a}, \overline{b}, \overline{c}$, and an *invisible* action τ , representing internal computation. The set of all actions, ranged over by α, β, γ , is given by $\text{Act}_\tau \stackrel{\text{def}}{=} \mathcal{N} \uplus \overline{\mathcal{N}} \uplus \{\tau\}$. We use μ, μ' to range over the set of visible actions $\mathcal{N} \uplus \overline{\mathcal{N}}$, and we extend the complementation function $\bar{\cdot}$ to this set by letting $\overline{\overline{a}} \stackrel{\text{def}}{=} a$. In the following, we will always assume $L = \text{Act}_\tau$. Once the LTS is fixed, we write $p \xrightarrow{\alpha} p'$ to mean that $(p, \alpha, p') \in \longrightarrow$ and $p \xrightarrow{\alpha}$ to mean $\exists p'. p \xrightarrow{\alpha} p'$.

We use \mathcal{L} to range over LTSs. To reason simultaneously on different LTSs, we will use the symbols \mathcal{L}_A and \mathcal{L}_B to denote respectively the LTSs $\langle A, L, \longrightarrow_A \rangle$ and $\langle B, L, \longrightarrow_B \rangle$.

In our mechanisation LTSs are borne out by the typeclass `Lts` in Figure 3. The states of the LTS have type A , labels have type L , and `lts_step` is the characteristic function of the transition relation, which we assume to be decidable. We let $O(p) = \{\overline{a} \in \overline{\mathcal{N}} \mid p \xrightarrow{\overline{a}}\}$ and $I(p) = \{a \in \mathcal{N} \mid p \xrightarrow{a}\}$ be respectively the set of outputs and the set of inputs of state p . We assume that the set $O(p)$

$$\begin{array}{ccc}
\frac{[S\text{-SRV}]}{p \xrightarrow{\tau} p'} & \frac{[S\text{-CLT}]}{r \xrightarrow{\tau} r'} & \frac{[S\text{-COM}]}{p \xrightarrow{\mu} p' \quad r \xrightarrow{\bar{\mu}} r'} \\
\hline
p \parallel r \longrightarrow p' \parallel r & p \parallel r \longrightarrow p \parallel r' & p \parallel r \longrightarrow p' \parallel r'
\end{array}$$

Fig. 4. The STS of server-client systems.

is finite for any p . In our mechanisation, the set $O(p)$ is rendered by the function `lts_outputs`, and we shall also use a function `lts_performs` that lets us decide whether a state can perform a transition labelled by a given action.

Client-server systems A *client-server* system (or *system*, for short) is a pair $p \parallel r$ in which p is deemed to be the server of client r . In general, every system $p \parallel r$ is the root of a *state transition system* (STS), $\langle S, \longrightarrow \rangle$, where S is the set of states and \longrightarrow is the reduction relation. For the sake of simplicity⁵ we derive the reduction relation from the LTS semantics of servers and clients as specified by the rules in Figure 4. In our mechanisation (Figure 3), `sts_step` is the characteristic function of the reduction relation \longrightarrow , and `sts_stable` is the function that states whether a state can reduce or not. Both functions are assumed decidable.

Definition 1 (Computation). *Given an STS $\langle S, \longrightarrow \rangle$ and a state $s_0 \in S$, a computation of s_0 is a finite or infinite reduction sequence starting from s_0 . A computation is maximal if either it cannot be extended or it is infinite.*

To formally define the MUST-preorder, we assume a decidable predicate `GOOD` over clients. A computation $p \parallel r = p_0 \parallel r_0 \longrightarrow p_1 \parallel r_1 \longrightarrow p_2 \parallel r_2 \longrightarrow \dots$ is *successful* if there exists a state $p_i \parallel r_i$ such that `GOOD(r_i)`. We assume the predicate `GOOD` to be *invariant under outputs*:

$$\text{If } r \xrightarrow{\bar{a}} r' \text{ then } \text{GOOD}(r) \iff \text{GOOD}(r') \quad (4)$$

All the previous works on asynchronous calculi implicitly make this assumption, since they rely on ad-hoc actions such as ω or \checkmark to signal success and they treat them as outputs. In Appendix H we show that this assumption holds for the language ACCS (the asynchronous variant of CCS) extended with the process `1`, which is used as a syntactic means to denote `GOOD` states. Moreover, when considering an equivalence on programs \simeq that is compatible with transitions, in the sense of Figure 5, we assume the predicate `GOOD` to be preserved also by this equivalence. These assumptions are met by the frameworks in [38,23,57].

Definition 2 (Client satisfaction). *We write $p \text{ MUST } r$ if every maximal computation of $p \parallel r$ is successful.*

⁵ In general the reduction semantics and the LTS of a calculus are defined independently, and connected via the Harmony lemma ([85], Lemma 1.4.15 page 51). We have a mechanised proof of it.

2. PRELIMINARIES

Definition 3 (must-preorder). *We let $p \sqsubseteq_{\text{MUST}} q$ whenever for every client r we have that $p \text{ MUST } r$ implies $q \text{ MUST } r$.*

Example 1. Consider the system $p_0 \parallel r_0$, where p_0 and r_0 are the server and client given in Figure 1. The unique maximal computation of this system is $p_0 \parallel r_0 \longrightarrow p_1 \parallel r_1 \longrightarrow p_3 \parallel r_3$. This computation is successful since it leads the client to the GOOD state r_3 . Hence, client r_0 is satisfied by server p_0 . Since OUTPUT-COMMUTATIVITY implies an absence of causality between the output $\overline{\text{str}}$ and the input int in the client, it is the order between the input str and the output $\overline{\text{int}}$ in the server that guides the order of client-server interactions. \square

A closer look at Selinger axioms Let us now discuss the axioms in Figure 2. The OUTPUT-COMMUTATIVITY axiom expresses the non-blocking behaviour of outputs: an output cannot be a cause of any subsequent transition, since it can also be executed after it, leading to the same resulting state. Hence, outputs are concurrent with any subsequent transition. The FEEDBACK axiom says that an output followed by a complementary input can also synchronise with it to produce a τ -transition. These first two axioms specify properties of outputs that are followed by another transition. Instead, the following three axioms, OUTPUT-CONFLUENCE, OUTPUT-DETERMINACY and OUTPUT-TAU, specify properties of outputs that are co-initial with another transition⁶. The OUTPUT-DETERMINACY and OUTPUT-TAU axioms apply to the case where the co-initial transition is an identical output or a τ -transition respectively, while the OUTPUT-CONFLUENCE axiom applies to the other cases. When taken in conjunction, these three axioms state that outputs cannot be in conflict with any co-initial transition, except when this is a τ -transition: in this case, the OUTPUT-TAU axiom allows for a confluent nondeterminism between the τ -transition on one side and the output followed by the complementary input on the other side.

We now explain the novel BACKWARD-OUTPUT-DETERMINACY axiom. It is the dual of OUTPUT-DETERMINACY, as it states that also backward transitions with identical outputs lead to the same state. The intuition is that if two programs arrive at the same state by removing the same message from the mailbox, then they must coincide. This axiom need not be assumed in [90] because it can be derived from Selinger axioms when modelling a calculus like ACCS equipped with a parallel composition operator \parallel (see Lemma 53 in Appendix H). We use the BACKWARD-OUTPUT-DETERMINACY axiom only to prove a technical property of clients (Lemma 25) that is used to prove our completeness result.

Calculi A number of asynchronous calculi [64,24,38,61,77,86] have an LTS that enjoys the axioms in Figure 2, at least up to some structural equivalence \equiv . The reason is that these calculi syntactically enforce outputs to have no continuation, *i.e.* outputs can only be composed in parallel with other processes.⁷ For

⁶ Two transitions are co-initial if they stem from the same state.

⁷ In the calculus TACCS (a variant of ACCS tailored for testing semantics) of [38] there is a construct of asynchronous output prefix, but its behaviour is to spawn the corresponding atom in parallel with the continuation, so it does not act as a prefix.

example, Selinger [90] shows that the axioms of Figure 2 hold for the LTS of the calculus ACCS (the asynchronous variant of CCS⁸) modulo bisimulation, and in Appendix H (Lemma 55) we prove that they hold also for the LTS of ACCS modulo \equiv :

Lemma 1. *We have that $\langle \text{ACCS}_{\equiv}, L, \longrightarrow_{\equiv} \rangle \in \text{FDB}$.*

To streamline reasoning modulo some (structural) equivalence we introduce the typeclass **LtsEq**, whose instances are LTSs equipped with an equivalence \simeq that satisfies the property in Figure 5. Defining output-buffered agents with feedback using **LtsEq** does not entail any loss of generality, because the equivalence \simeq can be instantiated using the identity over the states A . Further details can be found in Appendix H.1.

When convenient we denote LTSs using the following minimal syntax for ACCS:

$$p, q, r ::= \bar{a} \mid g \mid p \parallel p \mid \text{rec}x.p \mid x \quad g ::= 0 \mid a.p \mid \tau.p \mid g + g \quad (5)$$

as well as its standard LTS⁹ whose properties we discuss in detail in Appendix H. This is exactly the syntax used in [90,23], without the operators of restriction and relabelling. Here the syntactic category g defines *guards*, *i.e.* the terms that may be used as arguments for the $+$ operator. As in most process calculi, 0 denotes the terminated process that cannot do any action. Note that, apart from 0, only input-prefixed and τ -prefixed terms are allowed as guards, and that the output prefix operator is replaced by *atoms* \bar{a} . In fact, this syntax is completely justified by Selinger axioms, which, as we argued above, specify that outputs cannot cause any other action, nor be in conflict with it.

$$\begin{array}{ccc} p & & p \xrightarrow{\alpha} p' \\ \vdots & \Rightarrow & \vdots \\ \simeq & & \simeq \\ \vdots & & \vdots \\ q \xrightarrow{\alpha} q' & & q \xrightarrow{\alpha} q' \end{array}$$

Fig. 5. Axiom stating that equivalence \simeq is compatible with a transition relation.

Definition 4 (Transition sequence). *Given an LTS $\langle A, L, \longrightarrow \rangle$ and a state $p_0 \in A$, a transition sequence of p_0 is a finite or infinite sequence of the form $p_0 \alpha_1 p_1 \alpha_2 p_2 \cdots$ with $p_i \in A$ and $\alpha_i \in L$, and such that, for every $n \geq 1$ such that p_n is in the sequence we have $p_{n-1} \xrightarrow{\alpha_n} p_n$.*

⁸ The syntax of ACCS, which is closely inspired by that of the asynchronous π -calculus with input- and τ -guarded choice [4,5], is given in Equation (5) and discussed later.


⁹ Where the recursion rule is replaced by the one usually adopted for testing semantics, which introduces a τ -transition before each unfolding.

2. PRELIMINARIES

If a transition sequence is made only of τ -transitions, it is called a *computation* by abuse of notation, the idea being that usually τ -steps are related to reduction steps via the Harmony lemma (see footnote on page 7).

We give now an example that illustrates the use of the testing machinery in our asynchronous setting. This is also a counter-example to the completeness of the alternative preorder proposed in [38], as discussed in detail in Appendix I.

Example 2. Let $\Omega = \text{rec}x.\tau.x$ and $Pierre = b.(\tau.\Omega + c.\bar{d})$. The LTS of $Pierre$ is as follows:



$Pierre$ models a citizen confronted with an unpopular pension reform. To begin with, $Pierre$ can only do the input b , which models his getting aware of the brute-force imposition of the reform by the government. After performing the input, $Pierre$ reaches the state $\tau.\Omega + c.\bar{d}$, where he behaves in a nondeterministic manner. He can internally choose not to trust the government for any positive change, in which case he will diverge, refusing any further interaction. But this need not happen: in case the government offers the action \bar{c} , which models a positive change in political decision, $Pierre$ can decide to accept this change, and then he expresses his agreement with the output \bar{d} , which stands for “done”. \square

Example 3. We prove now the inequality $Pierre \sqsubseteq_{\text{MUST}} 0$ by leveraging the possibility of divergence of $Pierre$ after the input b . Fix an r such that $Pierre \text{MUST } r$. Note that, since 0 is the terminated process, the condition for the server 0 to satisfy r is that r reaches by itself a successful state in each of its maximal computations. We distinguish two cases, according to whether $r \xrightarrow{\bar{b}}$ or $r \not\xrightarrow{\bar{b}}$.

i) Let $r \xrightarrow{\bar{b}} r'$ for some r' . Consider the maximal computation $Pierre \parallel r \longrightarrow \tau.\Omega + c.\bar{d} \parallel r' \longrightarrow \Omega \parallel r' \longrightarrow \dots$ in which $Pierre$ diverges and r does not move after the first output. Since $Pierre \text{MUST } r$, either $\text{GOOD}(r)$ or $\text{GOOD}(r')$. In case $\text{GOOD}(r')$, by Equation (4) we get also $\text{GOOD}(r)$. Hence $0 \text{MUST } r$.

ii) Let $r \not\xrightarrow{\bar{b}}$. Suppose $r = r_0 \xrightarrow{\tau} r_1 \xrightarrow{\tau} r_2 \xrightarrow{\tau} \dots$ is a maximal computation of r . Then $Pierre \parallel r$ has a maximal computation $Pierre \parallel r_0 \longrightarrow Pierre \parallel r_1 \longrightarrow Pierre \parallel r_2 \longrightarrow \dots$. As $Pierre \text{MUST } r$, there must exist an $i \in \mathbb{N}$ such that $\text{GOOD}(r_i)$. Hence $0 \text{MUST } r$. \square

The argument in Example 3 can directly use Definition 3 because it is very simple to reason on the process 0 . The issues brought about by the contextuality of Definition 3, though, hinder showing general properties of $\sqsubseteq_{\text{MUST}}$. Consider the following form of code hoisting:

$$\tau.(\bar{a} \parallel \bar{b}) + \tau.(\bar{a} \parallel \bar{c}) \sqsubseteq_{\text{MUST}} \bar{a} \parallel (\tau.\bar{b} + \tau.\bar{c}) \quad (6)$$

If we see the above nondeterministic sums as representing the two branches of a conditional statement, this refinement corresponds to hoisting the shared action \bar{a} before the conditional statement, a common compiler optimisation. Proving Equation (6) via the contextual definition of $\sqsubseteq_{\text{MUST}}$ is cumbersome. This motivates the study of alternative characterisations for $\sqsubseteq_{\text{MUST}}$, and in the rest of the paper we present several preorders that fit the purpose, in particular the coinductive preorder \preceq_{co} , which we will use to establish Equation (6) in Section 3.3.

We conclude this section by recalling auxiliary and rather standard notions: given an LTS $\langle A, L, \longrightarrow \rangle$, the weak transition relation $p \xRightarrow{s} p'$, where $s \in \text{Act}^*$, is defined via the rules

$$\begin{aligned} [\text{wt-refl}] \quad & p \xRightarrow{\varepsilon} p \\ [\text{wt-tau}] \quad & p \xRightarrow{s} q \text{ if } p \xrightarrow{\tau} p' \text{ and } p' \xRightarrow{s} q \\ [\text{wt-mu}] \quad & p \xRightarrow{\mu.s} q \text{ if } p \xrightarrow{\mu} p' \text{ and } p' \xRightarrow{s} q \end{aligned}$$

We write $p \xRightarrow{s}$ to mean $\exists p'. p \xRightarrow{s} p'$.

We write $p \downarrow$ and say that p *converges* if every computation of p is finite, and we lift the convergence predicate to finite traces by letting the relation $\Downarrow \subseteq A \times \text{Act}^*$ be the least one that satisfies the following rules

$$\begin{aligned} [\text{cnv-epsilon}] \quad & p \Downarrow \varepsilon \text{ if } p \downarrow, \\ [\text{cnv-mu}] \quad & p \Downarrow \mu.s \text{ if } p \downarrow \text{ and } p \xRightarrow{\mu} p' \text{ implies } p' \Downarrow s. \end{aligned}$$

To understand the next section, one should keep in mind that all the predicates defined above have an implicit parameter: the LTS of programs. By changing this parameter, we may change the meaning of the predicates. For instance, letting Ω be the ACCS process $\text{rec}x.\tau.x$, in the standard LTS $\langle \text{ACCS}, \longrightarrow, \text{Act}_\tau \rangle$ we have $\Omega \xrightarrow{\tau} \Omega$ and $\neg(\Omega \downarrow)$, while in the LTS $\langle \text{ACCS}, \emptyset, \text{Act}_\tau \rangle$ we have $\Omega \xrightarrow{\tau} \Omega$ and thus $\Omega \downarrow$. In other words, the *same* predicates can be applied to different LTSs, and since the alternative characterisations of $\sqsubseteq_{\text{MUST}}$ are defined using such predicates, they can relate different LTSs.

3 Preorders based on acceptance sets

We first recall the definition of the standard alternative preorder \preceq_{AS} , and show how to use it to characterise $\sqsubseteq_{\text{MUST}}$ in our asynchronous setting. We also present a new characterisation that disregards the order of non-causally related actions. We then explain the tools we use to prove these characterisations, and in particular their soundness. This section ends with the coinductive version \preceq_{co} of \preceq_{AS} , which we use to prove the hoisting refinement (6).

3.1 Trace-based characterisations

The *ready set* of a program p is defined as $R(p) = I(p) \cup O(p)$, and it contains all the *visible* actions that p can immediately perform. If a program p is stable, *i.e.* it cannot perform any τ -transition, we say that it is a *potential deadlock*. In general,

3. PREORDERS BASED ON ACCEPTANCE SETS

the ready set of a potential deadlock p shows how to make p move to a different state, possibly one that can perform further computation: if $R(p) = \emptyset$ then there is no way to make p move on, while if $R(p)$ contains some action, then p is a state waiting for the environment to interact with it. Indeed, potential deadlocks are called *waiting states* in [64]. In particular, in an asynchronous setting the outputs of a potential deadlock p show how it can unlock the inputs of a client, which in turn may lead the client to a novel state that can make p move, possibly to a state that can perform further computation. A standard manner to capture all the ways out of the potential deadlocks that a program p encounters after executing a trace s is its *acceptance set*: $\mathcal{A}(p, s, \longrightarrow) = \{R(p') \mid p \xrightarrow{s} p' \xrightarrow{\tau} \}$.

In our presentation we indicate explicitly the third parameter of \mathcal{A} , *i.e.* the transition relation of the LTS at hand, because when necessary we will manipulate this parameter. For any two LTSs $\mathcal{L}_A, \mathcal{L}_B$ and servers $p \in A, q \in B$, we write $\mathcal{A}(p, s, \longrightarrow_A) \ll \mathcal{A}(q, s, \longrightarrow_B)$ if for every $R \in \mathcal{A}(q, s, \longrightarrow_B)$ there exists $\widehat{R} \in \mathcal{A}(p, s, \longrightarrow_A)$ such that $\widehat{R} \subseteq R$. We can now recall the definition of the behavioural preorder à la Hennessy, \preceq_{AS} , which is based on acceptance sets [55].

Definition 5. *We write*

- $p \preceq_{\text{cnv}} q$ whenever $\forall s \in \text{Act}^*. p \Downarrow_A s$ implies $q \Downarrow_B s$,
- $p \preceq_{\text{acc}} q$ whenever $\forall s \in \text{Act}^*. p \Downarrow_A s$ implies $\mathcal{A}(p, s, \longrightarrow_A) \ll \mathcal{A}(q, s, \longrightarrow_B)$,
- $p \preceq_{AS} q$ whenever $p \preceq_{\text{cnv}} q$ and $p \preceq_{\text{acc}} q$.

In the synchronous setting, the behavioural preorder \preceq_{AS} is closely related to the denotational semantics based on Acceptance Trees proposed by Hennessy in [54,55]. There the predicates need not be annotated with the LTS that they are used on, because those works treat a unique LTS. Castellani and Hennessy [38] show in their Example 4 that the condition on acceptance sets, *i.e.* \preceq_{acc} , is too demanding in an asynchronous setting.

Letting $p = a.0$ and $q = 0$, they show that $p \sqsubseteq_{\text{MUST}} q$ but $p \not\preceq_{AS} q$, because $\mathcal{A}(p, \epsilon) = \{\{a\}\}$ and $\mathcal{A}(q, \epsilon) = \{\emptyset\}$, and corresponding to the ready set $\emptyset \in \mathcal{A}(q, \epsilon)$ there is no ready set $\widehat{R} \in \mathcal{A}(p, s)$ such that $\widehat{R} \subseteq \emptyset$. Intuitively this is the case because acceptance sets treat inputs and outputs similarly, while in an asynchronous setting only outputs can be tested.

Nevertheless \preceq_{AS} characterises $\sqsubseteq_{\text{MUST}}$, if servers are enhanced as with forwarding. We now introduce this concept.

Forwarders We say that an LTS \mathcal{L} is of output-buffered agents *with forwarding*, for short is FWD, if it satisfies all the axioms in Figure 2 except FEEDBACK, and also the two following axioms:

$$\begin{array}{ccc}
 \begin{array}{c}
 \begin{array}{ccc}
 p & \begin{array}{c} \xrightarrow{a} \\ \xleftarrow{\bar{a}} \end{array} & p' \\
 \text{INPUT-BOOMERANG} & &
 \end{array} \\
 \end{array} &
 \begin{array}{c}
 \begin{array}{ccc}
 p & \xrightarrow{\bar{a}} & p' \\
 & & \downarrow_a \\
 & & q \\
 \text{FWD-FEEDBACK} & &
 \end{array} \\
 \Rightarrow p \xrightarrow{\tau} q \text{ or } p = q & & (7)
 \end{array}
 \end{array}$$

The INPUT-BOOMERANG axiom states a kind of input-enabledness property, which is however more specific as it stipulates that the target state of the input should loop back to the source state via a complementary output. This is the essence of the behaviour of a forwarder, whose role is simply to pass on a message and then get back to its original state. The FWD-FEEDBACK axiom is a weak form of Selinger’s FEEDBACK axiom, which is better understood in conjunction with the INPUT-BOOMERANG axiom: if the transition sequence $p \xrightarrow{\bar{a}} p' \xrightarrow{a} q$ in the FWD-FEEDBACK axiom is taken to be the transition sequence $p' \xrightarrow{\bar{a}} p \xrightarrow{a} p'$ in the INPUT-BOOMERANG axiom, then we see that it must be $q = p$ in the FWD-FEEDBACK axiom. Moreover, no τ action is issued when moving from p to q , since no synchronisation occurs in this case: the message is just passed on.

To prove that \preceq_{AS} is sound and complete with respect to \preceq_{MUST} :

1. we define a function $FW : FDB \rightarrow FWD$ that lifts any LTS $\mathcal{L} \in FDB$ into a suitable LTS $FW(\mathcal{L}) \in FWD$, and
2. we check the predicates \Downarrow and $\mathcal{A}(-, -, -)$ over the LTS $FW(\mathcal{L})$.

Let MO denote the set of all finite multisets of output actions, for instance we have $\emptyset, \{\bar{a}\}, \{\bar{a}, \bar{a}\}, \{\bar{a}, \bar{b}, \bar{a}, \bar{b}\} \in MO$. We let M, N, \dots range over MO . The symbol M stands for *mailbox*. We denote with \uplus the multiset union.

Definition 6. Let $FW(\mathcal{L}) = \langle A \times MO, L, \rightarrow_{fw} \rangle$ for every $\mathcal{L} = \langle A, L, \rightarrow \rangle$, where the states in $FW(\mathcal{L})$ are pairs denoted $p \triangleright M$, such that $p \in A$ and $M \in MO$, and the transition relation \rightarrow_{fw} is defined via the rules in Figure 6.

$$\begin{array}{c}
\text{[L-PROC]} \frac{p \xrightarrow{\alpha} p'}{p \triangleright M \xrightarrow{\alpha}_{fw} p' \triangleright M} \qquad \text{[L-COMM]} \frac{p \xrightarrow{a} p'}{p \triangleright (\{\bar{a}\} \uplus M) \xrightarrow{\tau}_{fw} p' \triangleright M} \\
\text{[L-MOUT]} \frac{}{p \triangleright (\{\bar{a}\} \uplus M) \xrightarrow{\bar{a}}_{fw} p \triangleright M} \qquad \text{[L-MINP]} \frac{}{p \triangleright M \xrightarrow{a}_{fw} p \triangleright (\{\bar{a}\} \uplus M)}
\end{array}$$

Fig. 6. Lifting of an LTS to an LTS with forwarding.


Let us briefly comment on Definition 6 and the rules in Figure 6. The pair $p \triangleright M$ is a kind of asymmetric parallel composition between a process p and a mailbox M . Rule [L-PROC] says that the process can evolve independently of the mailbox. Rule [L-COMM] says that an input in the process can synchronise with a complementary output in the mailbox. Rules [L-MINP] and [L-MOUT] express the essence of the forwarding behaviour: the pair $p \triangleright M$ may input any message from the environment and store it into the mailbox M (Rule [L-MINP]); dually, the pair $p \triangleright M$ may output any message in the mailbox M towards the environment (Rule [L-MOUT]). Note that in both cases, the interaction occurs between the mailbox M and the environment, without any participation of the process p .

3. PREORDERS BASED ON ACCEPTANCE SETS

Example 4. If a calculus is fixed, then the function FW may have a simpler definition. For instance Castellani and Hennessy [38] define it in their calculus TACCS by letting $\xrightarrow{\alpha}_{\text{fw}}$ be the least relation over TACCS such that (1) for every $\alpha \in \text{Act}_\tau$. $\xrightarrow{\alpha} \subseteq \xrightarrow{\alpha}_{\text{fw}}$, and (2) for every $a \in \mathcal{N}$. $p \xrightarrow{a}_{\text{fw}} p \parallel \bar{a}$. \square

The transition relation $\xrightarrow{\text{fw}}$ is reminiscent of the one introduced in Definition 8 by Honda and Tokoro in [64]. The construction given in our Definition 6, though, does not yield the LTS of Honda and Tokoro, as $\xrightarrow{\text{fw}}$ adds the forwarding capabilities to the states only at the top-level, instead of descending structurally into terms. As a consequence, in the LTS of [64] $a.0 + 0 \xrightarrow{b} \bar{b}$, while $a.0 + 0 \triangleright M \not\xrightarrow{b}_{\text{fw}} M \uplus \{\bar{b}\}$.

Example 5. As the set \mathcal{N} is countable, every process p that belongs to the LTS $\langle \text{ACCS} \times \text{MO}, \text{Act}_\tau, \xrightarrow{\text{fw}} \rangle$ is infinitely-branching: for every mailbox M we have $p \triangleright M \xrightarrow{a_i}_{\text{fw}} p \triangleright (\{\bar{a}_i\} \uplus M)$ for every $a_i \in \mathcal{N}$. This is illustrated by the following picture, where for simplicity we omit the subscript fw under the arrows.



\square

The intuition behind Definition 6 is that, when a client interacts with a server asynchronously, the client can send any message it likes, regardless of the inputs that the server can actually perform. In fact, asynchronous clients behave as if the server was saturated with *forwarders*, namely processes of the form $a.\bar{a}$, for any $a \in \mathcal{N}$.

The function FW enjoys two crucial properties: it lifts any LTS of output-buffered agents with feedback to an LTS with forwarding, and the lifting preserves the MUST predicate. We can thus reason on $\sqsubseteq_{\text{MUST}}$ using LTSs in FWD.

Lemma 2. For every LTS $\mathcal{L} \in \text{FDB}$, $\text{FW}(\mathcal{L}) \in \text{FWD}$.

Proof. See Appendix C. \square

Lemma 3. For every $\mathcal{L}_A, \mathcal{L}_B, \mathcal{L}_C \in \text{FDB}$, $p \in A$, $q \in B$, $r \in C$,

1. $p \text{ MUST } r$ if and only if $\text{FW}(p) \text{ MUST } r$,
2. $p \sqsubseteq_{\text{MUST}} q$ if and only if $\text{FW}(p) \sqsubseteq_{\text{MUST}} \text{FW}(q)$.

We now simplify the definition of acceptance sets to reason on LTSs that are in FWD: for any LTS $\mathcal{L} = \langle A, \text{Act}, \xrightarrow{\text{fw}} \rangle \in \text{FWD}$ and program $p \in A$ we let $\mathcal{A}_{\text{fw}}(p, s, \xrightarrow{\text{fw}}) = \{O(p') \mid p \xrightarrow{s} p' \xrightarrow{\tau} \}$. This definition suffices to characterise $\sqsubseteq_{\text{MUST}}$ because in each LTS that is FWD every state performs every input,

3. PREORDERS BASED ON ACCEPTANCE SETS

thus comparing inputs has no impact on the preorder \preceq_{acc} of Definition 5. More formally, for every $\mathcal{L}_A, \mathcal{L}_B \in \text{FWD}$ and every $p \in A$ and $q \in B$, we let

$$p \preceq_{\text{acc}}^{\text{fw}} q \text{ iff } \forall s \in \text{Act}^*. p \Downarrow s \text{ implies } \mathcal{A}_{\text{fw}}(p, s, \rightarrow_A) \ll \mathcal{A}_{\text{fw}}(q, s, \rightarrow_B)$$

We have the following logical equivalence.

Lemma 4. *Let $\mathcal{L}_A, \mathcal{L}_B \in \text{FWD}$. For every $p \in A, q \in B, p \preceq_{\text{acc}} q$ if and only if $p \preceq_{\text{acc}}^{\text{fw}} q$.*

Proof. The *only if* implication is trivial, so we discuss the *if* one. Suppose that $p \preceq_{\text{acc}}^{\text{fw}} q$ and that for some s we have that $R \in \mathcal{A}(q, s, \rightarrow_B)$. Let X be the possibly empty subset of R that contains only output actions. Since \mathcal{L}_B is FWD we know by definition that $R = X \cup \mathcal{N}$. By definition $X \in \mathcal{A}_{\text{fw}}(q, s, \rightarrow_B)$, and thus by hypothesis there exists a set of output actions $Y \in \mathcal{A}_{\text{fw}}(p, s, \rightarrow_A)$ such that $Y \subseteq X$. It follows that the set $Y \cup \mathcal{N} \in \mathcal{A}(p, s, \rightarrow_A)$, and trivially $Y \cup \mathcal{N} \subseteq X \cup \mathcal{N} = R$. \square

In view of the second point of Lemma 3, to prove completeness it suffices to show that \preceq_{AS} includes $\sqsubseteq_{\text{MUST}}$ over LTSs with forwarding. This is indeed true:

Lemma 5. *For every $\mathcal{L}_A, \mathcal{L}_B \in \text{FWD}$ and $p \in A, q \in B$, if $p \sqsubseteq_{\text{MUST}} q$ then we have $p \preceq_{\text{AS}} q$.*

By a slight abuse of notation, given an LTS $\mathcal{L} = \langle A, L, \rightarrow \rangle$ and a state $p \in A$, we denote with $\text{FW}(p)$ the LTS rooted at $p \triangleright \emptyset$ in $\text{FW}(\mathcal{L})$.

Theorem 1. *For every $\mathcal{L}_A, \mathcal{L}_B \in \text{FDB}$ and $p \in A, q \in B$,*

$$p \sqsubseteq_{\text{MUST}} q \text{ iff } \text{FW}(p) \preceq_{\text{AS}} \text{FW}(q).$$

This theorem is the linchpin of this paper, as all other results presented in this paper are corollaries. To begin with, instantiating this theorem to a calculus which can be given an LTS that satisfies the axioms for output-buffered agents with feedback such as ACCS (Lemma 1), the core join-calculus [48] or KLAIM [42], we get a characterisation of the MUST-preorder essentially for free. In our Coq development, we instantiate it with ACCS:

Corollary 1. *For every $p, q \in \text{ACCS}_{\equiv}, p \sqsubseteq_{\text{MUST}} q$ iff $\text{FW}(p) \preceq_{\text{AS}} \text{FW}(q)$.*

Another application of Theorem 1 is a novel behavioural characterisation of the MUST-preorder, which fully exploits asynchrony, *i.e.* disregards irrelevant non-causal orders of visible actions in traces. For space reasons, we defer the discussion of this result to Appendix G.

So far, we have seen the more direct applications of Theorem 1. Before we explain two other applications, namely the coinductive characterisation of the MUST-preorder and the relation of $\sqsubseteq_{\text{MUST}}$ with the failure refinement, we outline the proof of Theorem 1, and, in particular, the technical tools we used.

3.2 Proof of Theorem 1

The full proof of Theorem 1, which is given in the Appendix, as well as in the Coq development, comprises two parts: Appendix D deals with completeness, where the main aim is to show Lemma 5, and Appendix E deals with soundness. Here we outline the main tools we use to prove the soundness of the \preceq_{AS} preorder: *Bar-induction*, which allows us to relate the standard definition of MUST with an inductive one that is more practical to use, especially in a constructive setting, and the *LTS of sets*, which is derived from the LTS of the processes.

Bar induction: from extensional to intensional definitions We present the inductive characterisations of \downarrow and MUST in any state transition system (STS) $\langle S, \rightarrow \rangle$ that is countably branching. In practice, this condition is satisfied by most concrete LTS of programming languages, which usually contain countably many terms; this is the case for ACCS and for the asynchronous π -calculus.

Following the terminology of [28] we introduce extensional and intensional predicates associated to any decidable predicate $Q : S \rightarrow \mathbb{B}$ over an STS $\langle S, \rightarrow \rangle$, where \mathbb{B} denotes the set of booleans.

Definition 7. *The extensional predicate $\text{ext}_Q(s)$ is defined, for $s \in S$, as*

$$\forall \eta \text{ maximal execution of } S. \eta_0 = s \text{ implies } \exists n \in \mathbb{N}, Q(\eta_n)$$

The intensional predicate int_Q is the inductive predicate (least fixpoint) defined by the following rules:

$$[\text{AXIOM}] \frac{Q(s)}{\text{int}_Q(s)} \quad [\text{IND-RULE}] \frac{s \rightarrow \quad \forall s'. s \rightarrow s' \text{ implies } \text{int}_Q(s')}{\text{int}_Q(s)}$$

For instance, by letting

$$Q_1(p) \iff p \dashv \dashv \quad Q_2(p, r) \iff \text{GOOD}(r)$$

we have by definition that

$$p \downarrow \iff \text{ext}_{Q_1}(p) \quad p \text{ MUST } r \iff \text{ext}_{Q_2}(p, r) \quad (\text{ext-preds})$$

that is the standard definitions of \downarrow and MUST are extensional. Our aim now is to prove that they coincide with their intensional counterparts. The reader not familiar with this terminology may find in Appendix B an informal and hopefully intuitive explanation. Since we will use the intensional predicates in the rest of the paper a little syntactic sugar is in order, let

$$p \downarrow_i \iff \text{int}_{Q_1}(p) \quad p \text{ MUST}_i r \iff \text{int}_{Q_2}(p, r) \quad (\text{int-preds})$$

The proofs of soundness, *i.e.* that the inductively defined predicates imply the extensional ones, are by rule induction:

Lemma 6. For $p \in S$,

- (a) $p \downarrow_i$ implies $p \downarrow$,
- (b) for every r . $p \text{ MUST}_i r$ implies $p \text{ MUST } r$.

The proofs of completeness are more delicate. To the best of our knowledge, the ones about CCS [39,12] proceed by induction on the greatest number of steps necessary to arrive at termination or at a successful state. Since the STS of $\langle \text{CCS}, \xrightarrow{\tau} \rangle$ is finite branching, König's lemma guarantees that such a bound exists. This technique does not work on infinite-branching STSs, for example the one of CCS with infinite sums [14]. If we reason in classical logic, we can prove completeness without König's lemma and also over infinite-branching STSs via a proof *ad absurdum*: suppose $p \downarrow$. If $\neg(p \downarrow_i)$ no finite derivation tree exists to prove $p \downarrow_i$, and then we construct an infinite sequence of τ moves starting with p , thus $\neg(p \downarrow)$. Since we strive to be constructive we replace reasoning *ad absurdum* with a constructive axiom: (decidable) *bar induction*. In the rest of this section we discuss this axiom, and adapt it to our client-server setting. This requires a little terminology.

Bar induction The axiom we want to use is traditionally stated using natural numbers. We use the standard notations \mathbb{N}^* for finite sequences of natural numbers, \mathbb{N}^ω for infinite sequences, and $\mathbb{N}^\infty = \mathbb{N}^* \cup \mathbb{N}^\omega$ for finite or infinite sequences. Remark that, in constructive logics, given $u \in \mathbb{N}^\infty$, we cannot do a case analysis on whether u is finite or infinite. The set \mathbb{N}^∞ equipped with the prefix order can be seen as a *tree*, denoted $T_{\mathbb{N}}$, in the sense of set theory: a tree is an ordered set (A, \leq) such that, for each $a \in A$, the set $\{b \mid b < a\}$ is well-ordered by $<$. A *path* in a tree A is a maximal element in A . In the tree \mathbb{N}^∞ , each node has ω children, and the paths are exactly the infinite sequences \mathbb{N}^ω .

A predicate $P \subseteq \mathbb{N}^*$ over finite words is a *bar* if every infinite sequence of natural numbers has a finite prefix in P . Note that a bar defines a subtree of $T_{\mathbb{N}}$ *extensionally*, because it defines each path of the tree, as a path $u \in \mathbb{N}^\omega$ is in the tree if and only if there exists a finite prefix which is in the bar P .

A predicate $Q \subseteq \mathbb{N}^*$ is *hereditary* if

$$\forall w \in \mathbb{N}^*, \text{ if } \forall n \in \mathbb{N}, w \cdot n \in Q \text{ then } w \in Q.$$

Bar induction states that the extensional predicate associated to a bar implies its *intensional* counterpart: a predicate $P_{int} \subseteq \mathbb{N}^*$ which contains Q and which is hereditary.

Axiom 8 (Decidable bar induction over \mathbb{N}) Given two predicates P_{int}, Q over \mathbb{N}^* , such that:

1. for all $\pi \in \mathbb{N}^\omega$, there exists $n \in \mathbb{N}$ such that $(\pi_1, \dots, \pi_n) \in Q$;
2. for all $w \in \mathbb{N}^*$, it is decidable whether $Q(w)$ or $\neg Q(w)$;
3. for all $w \in \mathbb{N}^*$, $Q(w) \Rightarrow P_{int}(w)$;
4. P_{int} is hereditary;

then P_{int} holds over the empty word: $P_{int}(\varepsilon)$.

Bar induction is a generalisation of the fan theorem, i.e. the constructive version of König's lemma [47, pag. 56], and states that any extensionally well-founded tree T can be turned into an inductively-defined tree t that realises T [28,68].

Our mechanisation of bar induction principle is formulated as a Proposition that is proved using classical reasoning, since it is not provable directly in the type theory of Coq. Unfortunately, while bar induction is a constructive principle, mainstream proof assistants do not support it yet, which is why on the one hand we had to postulate it as a proof principle while on the other hand we proved it in classical logic using the Excluded Middle axiom. This principle though has a computational content, Spector bar recursion¹⁰, which, currently, cannot be used in mainstream proof assistants such as Coq. Developing a type theory with a principle of bar induction is recent and ongoing work [50,81].

Encoding states The version of bar induction we just outlined is not directly suitable for our purposes, as we need to reason about sequences of reductions rather than sequences of natural numbers. The solution is to encode STS states by natural numbers. This leads to the following issue: the nodes of the tree $T_{\mathbb{N}}$ have a fixed arity, namely \mathbb{N} , while processes have variably many reducts, including zero if they are stable. To deal with this glitch, it suffices to assume that there exists the following family of surjections:

$$F(p) : \mathbb{N} \rightarrow \{q \mid p \rightarrow q\} \tag{8}$$

where a surjection is defined as follows.

Definition 9. *A map $f : A \rightarrow B$ is a surjection if it has a section $g : B \rightarrow A$, that is, $f \circ g = \text{Id}_B$.*

This definition implies the usual one which states the existence of an antecedent $x \in A$ for any $y \in B$, and it is equivalent to it if we assume the Axiom of Choice.

Using this map F as a decoding function, any sequence of natural numbers corresponds to a path in the STS. Its subjectivity means that all paths of the LTS can be represented as such a sequence. This correspondence allows us to transport bar induction from sequences of natural numbers to executions of processes.

Note that such a family of surjections F exists for ACCS processes, and generally to most programming languages, because the set Act_τ is countable, and so are processes. This leads to the following version of bar induction where words and sequences are replaced by finite and infinite executions.

Proposition 1 (Decidable bar induction over an STS). *Let $\langle S, \rightarrow \rangle$ be an STS such that a surjection as in (8) exists. Given two predicates Q, P_{int} over finite executions, if*

¹⁰ https://en.wikipedia.org/wiki/Bar_recursion

1. for all infinite execution η , there exists $n \in \mathbb{N}$ such that $(\eta_1, \dots, \eta_n) \in Q$;
2. for all finite execution ζ , $Q(\zeta)$ or $\neg Q(\zeta)$ is decidable;
3. for all finite execution ζ , $Q(\zeta) \Rightarrow P_{int}(\zeta)$;
4. P_{int} is hereditary, as defined above except that $\zeta \cdot q$ is a partial operation defined when ζ is empty or its last state is p and $p \rightarrow q$;

then P_{int} holds over the empty execution: that is $P_{int}(\varepsilon)$ holds.

The last gap towards a useful principle is the requirement that every state in our STS has an outgoing transition. This condition is necessary to ensure the existence of the surjection in Equation (8). To ensure this requirement given any countably-branching STS, we enrich it by adding a *sink* state, which (a) is only reachable from stable states of the original STS, and (b) loops. This is a typical technique, see for instance [71, pag. 17].

Definition 10. Define $Sink(S, \rightarrow) := \langle S \cup \{\top\}, \rightarrow^\top \rangle$, where \rightarrow^\top is defined inductively as follows:

$$p \rightarrow q \implies p \rightarrow^\top q \quad p \dashrightarrow \implies p \rightarrow^\top \top \quad \top \rightarrow^\top \top$$

A maximal execution of $Sink(S, \rightarrow)$ is always infinite, and it corresponds (in classical logic) to either an infinite execution of S or a maximal execution of S followed by infinitely many \top . We finally prove the converse of Lemma 6.

Proposition 2. Given a countably branching STS $\langle S, \rightarrow \rangle$, and a decidable predicate Q on S , we have that, for all $s \in S$, $\text{ext}_Q(s)$ implies $\text{int}_Q(s)$.

Now we thus obtain completeness of the intensional predicates.

Corollary 2. For every $p \in A$ we have

1. $p \downarrow$ if and only if $p \downarrow_i$, and
2. for every r we have that $p \text{MUST} r$ if and only if $p \text{MUST}_i r$.

Proof. Direct consequence of Proposition 2, and Equation (ext-preds) and Equation (int-preds) above.

As we have outlined why Corollary 2 is true, from now on we use \downarrow_i and MUST_i instead of \downarrow and MUST . In Appendix B.2 we prove the properties of these predicates, that we use in the rest of the paper.

The LTS of sets Recall that soundness of \preceq_{AS} means that $\preceq_{AS} \subseteq \overline{\preceq}_{\text{MUST}}$. The naïve reasoning does not work. Fix two servers p and q such that $p \preceq_{AS} q$. We need to prove that for every client r , if $p \text{MUST}_i r$ then $p \text{MUST}_i r$. Rule induction on the predicate $p \text{MUST}_i r$ fails, as demonstrated by the following example.

Example 6. Consider the two servers $p = \tau.(\bar{a} \parallel \bar{b}) + \tau.(\bar{a} \parallel \bar{c})$ and $q = \bar{a} \parallel (\tau.\bar{b} + \tau.\bar{c})$ of Equation (6). Fix a client r such that $p \text{MUST}_i r$. Rule induction yields the following inductive hypothesis:

3. PREORDERS BASED ON ACCEPTANCE SETS

$$\forall p', q'. p \parallel r \xrightarrow{\tau} p' \parallel r' \wedge p' \preceq_{\text{AS}} q' \Rightarrow q' \text{MUST}_i r'.$$

In the proof of $q \text{MUST}_i r$ we have to consider the case where there is a communication between q and r such that, for instance, $q \xrightarrow{\bar{a}} \tau.\bar{b} + \tau.\bar{c}$ and $r \xrightarrow{a} r'$. In that case, we need to show that $\tau.\bar{b} + \tau.\bar{c} \text{MUST}_i r'$. Ideally, we would like to use the inductive hypothesis. This requires us to exhibit a p' such that $p \parallel r \xrightarrow{\tau} p' \parallel r'$ and $p' \preceq_{\text{acc}} \tau.\bar{b} + \tau.\bar{c}$. However, note that there is no way to derive $p \parallel r \xrightarrow{\tau} p' \parallel r'$, because $p \xrightarrow{\bar{a}}$. The inductive hypothesis thus cannot be applied, and the naïve proof does not go through. \square

This example suggests that defining an auxiliary predicate MUST_{aux} in some sense equivalent to MUST_i , but that uses explicitly *weak* outputs of servers, should be enough to prove that \preceq_{AS} is sound with respect to $\sqsubseteq_{\text{MUST}}$. Unfortunately, though, there is an additional nuisance to tackle: server nondeterminism.

Example 7. Assume that we defined the predicate MUST_i using weak transitions on the server side. Recall the argument put forward in the previous example. The inductive hypothesis now becomes the following:

For every p', q', μ such that $p \xrightarrow{\mu} p'$ and $r \xrightarrow{\mu} r'$, $p' \preceq_{\text{AS}} q'$ implies $q' \text{MUST}_i r'$.

To use the inductive hypothesis we have to choose a p' such that $p \xrightarrow{\bar{a}} p'$ and $p' \preceq_{\text{AS}} \tau.\bar{b} + \tau.\bar{c}$. This is still not enough for the entire proof to go through, because (modulo further τ -moves) the particular p' we pick has to be related also to either \bar{b} or \bar{c} . It is not possible to find such a p' , because the two possible candidates are either \bar{b} or \bar{c} ; neither of which can satisfy $p' \preceq_{\text{AS}} \tau.\bar{b} + \tau.\bar{c}$, as the right-hand side has not committed to a branch yet.

If instead of a single state p in the novel definition of MUST_i we used a set of states and a suitable transition relation, the choice of either \bar{b} or \bar{c} would be suitably delayed. It suffices for instance to have the following states and transitions: $\{p\} \xrightarrow{\bar{a}} \{\bar{b}, \bar{c}\}$. \square

Now that we have motivated the main intuitions behind the definition of our novel auxiliary predicate MUST_{aux} , we proceed with the formal definitions.

Definition 11 (LTS of sets). Let $\mathcal{P}^+(Z)$ be the set of non-empty parts of Z . For any LTS $\langle A, L, \longrightarrow \rangle$, $X \in \mathcal{P}^+(A)$ and $\alpha \in L$, we define the sets

$$D(\alpha, X) = \{p' \mid \exists p \in X. p \xrightarrow{\alpha} p'\}, \quad WD(\alpha, X) = \{p' \mid \exists p \in X. p \xrightarrow{\alpha} p'\}.$$

We construct the LTS $\langle \mathcal{P}^+(A), \text{Act}_\tau, \longrightarrow \rangle$ by letting $X \xrightarrow{\alpha} D(\alpha, X)$ whenever $D(\alpha, X) \neq \emptyset$. Similarly, we have $X \xrightarrow{\alpha} WD(\alpha, X)$ whenever $WD(\alpha, X) \neq \emptyset$.

Intuitively, this definition lifts the standard notion of state derivative to sets of states. This construction is standard [41,19,20] and goes back to the determination of nondeterministic automata.

Let MUST_{aux} be defined via the rules in Figure 7. This predicate lets us reason on MUST_i via sets of servers, in the following sense:

<p>[MSET-NOW]</p> $\frac{\text{GOOD}(r)}{X \text{ MUST}_{\text{aux}} r}$	<p>[MSET-STEP]</p> $\frac{\neg \text{GOOD}(r) \quad \forall X'. X \xrightarrow{\tau} X' \text{ implies } X' \text{ MUST}_{\text{aux}} r \quad \forall p \in X. p \parallel r \xrightarrow{\tau} \quad \forall r'. r \xrightarrow{\tau} r' \text{ implies } X \text{ MUST}_{\text{aux}} r'}{\forall X', \mu \in \text{Act}^*. X \xrightarrow{\bar{\mu}} X' \text{ and } r \xrightarrow{\mu} r' \text{ imply } X' \text{ MUST}_{\text{aux}} r'}$ $X \text{ MUST}_{\text{aux}} r$
--	--

Fig. 7. Rules to define inductively the predicate MUST_{aux} .

Lemma 7. *For every LTSs $\mathcal{L}_A, \mathcal{L}_B$ and every set of servers $X \in \mathcal{P}^+(A)$, we have that $X \text{ MUST}_{\text{aux}} r$ if and only if for every $p \in X$. $p \text{ MUST}_i r$.*

The LTS of sets has two important applications in this paper: first, it is used to define the MUST_{aux} relation, on which we rely to prove the soundness of the characterisation (see Appendix E). Additionally, it is used in the definition of the coinductive characterisation, which is the topic of the next section.

3.3 The action-based coinductive characterisation

We conclude this section by introducing a characterisation of the MUST -preorder that is more practical than \preceq_{AS} , as it allows one to use the usual coinductive techniques. In addition, in the asynchronous case where processes are enhanced with forwarding, being able to use the coinductive proof method allows us to deal easily with the additional transitions due to forwarding. As a demonstration, we use this preorder to prove the code hoisting refinement shown in (6).

First, we recall the definition of this alternative preorder, which, like the other ones, is the same as in the synchronous case [2,72,16].

Definition 12 (Coinductive preorder). *For all image-finite LTSs $\mathcal{L}_A, \mathcal{L}_B$ and all $X \in \mathcal{P}^+(A), q \in B$, we let the coinductive preorder \preceq_{co} be defined as the greatest relation such that whenever $X \preceq_{\text{co}} q$, the following requirements hold:*

1. $X \downarrow$ implies $q \downarrow$,
2. For each q' such that $q \xrightarrow{\tau} q'$, we have that $X \preceq_{\text{co}} q'$,
3. $X \downarrow$ and $q \xrightarrow{\tau}$ imply that there exist $p \in X$ and $p' \in A$ such that $p \Longrightarrow p' \xrightarrow{\tau}$ and $R(p') \subseteq R(q)$,
4. For any $\mu \in \text{Act}$, if $X \Downarrow \mu$, then for every X' and q' such that $X \xrightarrow{\mu} X'$ and $q \xrightarrow{\mu} q'$, we have that $X' \preceq_{\text{co}} q'$.

This preorder characterises $\sqsubseteq_{\text{MUST}}$ when the set X of servers is a singleton.

Theorem 2. *For every image-finite LTS $\mathcal{L}_A, \mathcal{L}_B \in \text{FDB}$, every $p \in A$ and $q \in B$, we have that $p \sqsubseteq_{\text{MUST}} q$ if and only if $\{\text{FW}(p)\} \preceq_{\text{co}} \text{FW}(q)$.*

4. PREORDERS BASED ON MUST-SETS AND FAILURE REFINEMENT

The idea of the proof is to establish that the coinductive preorder characterises a version of the MUST-preorder that also has a set of servers on its LHS, and is defined as follows:

$$X \sqsubseteq_{\text{MUST}}^{\text{Set}} q \iff \forall t. (\forall p \in X, p \text{ MUST}_i t) \text{ implies } q \text{ MUST}_i t.$$

Observe that Definition 12 is based on single actions, instead of traces like \preceq_{AS} , thus it gives us a practical proof method. To make this point, we now prove the code hoisting refinement (6). According to Theorem 2, it suffices to prove:

$$\{\tau.(\bar{a} \parallel \bar{b}) + \tau.(\bar{a} \parallel \bar{c})\} \preceq_{\text{co}} \bar{a} \parallel (\tau.\bar{b} + \tau.\bar{c}) \quad (9)$$

As for proofs by induction, when using coinduction it is helpful to prove a more general statement, which yields a useful *coinductive hypothesis*. This vocabulary corresponds to the proof theoretic point of view of coinduction, which matches how Coq implements coinductive proofs using `cofix`. In practice, the prover can use the coinductive hypothesis *after* the predicate defined coinductively has been unfolded at least once. In the set-theoretic setting used in Definition 12, this corresponds to choosing a relation R that is closed under the operations given in the definition. This is borne out by Lemma 47 in Appendix F.

4 Preorders based on must-sets and failure refinement

We now establish the second *standard* characterisation of the MUST-preorder, defined using MUST-sets, again thanks to Theorem 1. As an application, we relate the *failure refinement* used by the CSP community to the MUST-preorder.

We begin by defining formally the \preceq_{MS} preorder, and we relate it to the MUST-preorder. For every $X \subseteq_{\text{fin}} \text{Act}$, that is for every finite set of visible actions, with a slight abuse of notation we write $p \text{ MUST } X$ whenever $p \xrightarrow{\varepsilon} p'$ implies that $p' \xrightarrow{\mu}$ for some $\mu \in X$, and we say that X is a MUST-set of p . Let $(p \text{ after } s, \longrightarrow) = \{p' \mid p \xrightarrow{s} p'\}$. For every $\mathcal{L}_A, \mathcal{L}_B$ and $p \in A, q \in B$, let $p \preceq_{\text{M}} q$ whenever $\forall s \in \text{Act}^*$ we have that $p \Downarrow s$ implies that $(\forall X \subseteq_{\text{fin}} \text{Act}$ if $(p \text{ after } s, \longrightarrow_A) \text{ MUST } X$ then $(q \text{ after } s, \longrightarrow_B) \text{ MUST } X$).

Definition 13. For all $\mathcal{L}_A, \mathcal{L}_B \in \text{FDB}$ and servers $p \in A$ and $q \in B$, we let $p \preceq_{\text{MS}} q$ whenever $p \preceq_{\text{cnv}} q \wedge p \preceq_{\text{M}} q$.

Lemma 8. Let $\mathcal{L}_A, \mathcal{L}_B \in \text{FDB}$. For all $p \in A$ and $q \in B$ such that $\text{FW}(p) \preceq_{\text{cnv}} \text{FW}(q)$, we have that $\text{FW}(p) \preceq_{\text{M}} \text{FW}(q)$ if and only if $\text{FW}(p) \preceq_{\text{acc}}^{\text{fw}} \text{FW}(q)$.

As a direct consequence, we obtain the following result.

Theorem 3. Let $\mathcal{L}_A, \mathcal{L}_B \in \text{FDB}$. For all $p \in A$ and $q \in B$, we have that $p \sqsubseteq_{\text{MUST}} q$ if and only if $\text{FW}(p) \preceq_{\text{MS}} \text{FW}(q)$.

Failure refinement MUST-sets have been used mainly by De Nicola and collaborators, for instance in [44,23], and are closely related to the failure refinement proposed in [33] by Hoare, Brookes and Roscoe for TCSP (the process algebra based on Hoare’s language CSP [63,31]). Following [33], a *failure* of a process p is a pair (s, X) such that $p \xrightarrow{s} p'$ and $p' \not\xrightarrow{\mu}$ for all $\mu \in X$. Then, failure refinement is defined by letting $p \leq_{\text{fail}} q$ whenever the failures of q are also failures of p . This refinement was designed to give a denotational semantics to processes, and mechanisations in Isabelle/HOL have been developed to ensure that the refinement is well defined [91,11]. Both Hennessy [55, pag. 260] and [37] highlight that the failure model can be justified operationally via the MUST testing equivalence: it is folklore dating back to [43, Section 4] that failure equivalence and \approx_{MUST} coincide. Thanks to Theorem 3 we conclude that in fact $\sqsubseteq_{\text{MUST}}$ coincides with the *failure divergence refinement* [92], that is, the intersection of \leq_{fail} and \preceq_{cnv} .

Corollary 3. *Let $\mathcal{L}_A, \mathcal{L}_B \in \text{FDB}$. For every $p \in A$ and $q \in B$, we have that $p \sqsubseteq_{\text{MUST}} q$ if and only if $\text{FW}(p) \preceq_{\text{cnv}} \text{FW}(q)$ and $\text{FW}(p) \leq_{\text{fail}} \text{FW}(q)$.*

Thanks to Corollary 4, we obtain $\preceq_{\text{MS}}^{\text{NF}}$, the analogous of Definition 16 based on normal forms, and we prove that it characterises $\sqsubseteq_{\text{MUST}}$.

5 Related work

Here we discuss in detail the works more closely related to the results of this paper. Further discussion on related work may be found in Appendix A.

The first investigation on the MUST-preorder in an asynchronous setting was put forth by [38]. While their very clear examples shed light on the preorder, their alternative preorder (Definition 6 in that paper) is more complicated than necessary: it uses the standard LTS of TACCS, its lifting to an LTS with forwarding, and two somewhat ad-hoc notions: a predicate \sim^I and a condition on multisets of inputs. Moreover that preorder is not complete because of a glitch in the treatment of divergence. The details of the counter-example we found to that completeness result are given in Appendix I.

In [57] Hennessy outlines how to adapt the approach of [38] to a typed asynchronous π -calculus. While the LTS with forwarding is replaced by a Context LTS, the predicates to define the alternative preorder are essentially the same as those used in the preceding work with Castellani. Acceptance sets are given in Definition 3.19 there, and the predicate \rightsquigarrow is denoted \searrow , while the generalised acceptance sets of [38] are given in Definition 3.20. Owing to the glitch in the completeness of [38], it is not clear that Theorem 3.28 of [57] is correct either.

Also the authors of [23] consider the MUST-preorder in ACCS. There is a major difference between their approach and ours. When studying theories for asynchronous programs, one can either

- (1) keep the definitions used for synchronous programs, and enhance the LTS with forwarders; or
- (2) adapt the definitions, and keep the standard LTS.

6. CONCLUSION

In the first case, the complexity is moved into the LTS, which becomes infinite-branching and infinite-state. In the second case, the complexity is moved into the definitions used to reason on the LTS (i.e. in the meta-language), and in particular in the definition of the alternative preorder, which deviates from the standard one. The authors of [23] follow the second approach. This essentially explains why they employ the standard LTS of CCS and to tackle asynchrony they reason on traces via (1) a preorder \preceq (Table 2 of that paper) that defines on *input* actions the phenomena due to asynchrony; and (2) a rather technical operation on traces, namely $s \ominus s' = (\{s\}_i \setminus \{s'\}_i) \setminus (\{s\}_o \setminus \{s'\}_o)$. We favour instead the approach in (1), for it helps achieve a modular mechanisation.

The authors of [45] give yet another account of the MUST-preorder. Even though non-blocking outputs can be written in their calculus, they use a left-merge operator that allows writing *blocking* outputs. The contexts that they use to prove the completeness of their alternative preorder use such blocking outputs, consequently their arguments need not tackle the asymmetric treatment of input and output actions. This explains why they can use smoothly a standard LTS, while [38] and [23] have to resort to more complicated structures.

Theorem 5.3 of the PhD thesis by [94] states an alternative characterisation of the MUST-preorder, but it is given with no proof. The alternative preorder given in Definition 5.8 of that thesis turns out to be a mix of the ones by [38] and [23]. In particular, the definition of the alternative preorder relies on the LTS with forwarding, there denoted \rightarrow_A (Point 1. in Definition 5.1 defines exactly the input transitions that forward messages into the global buffer). The condition that compares convergence of processes is the same as in [38], while server actions are compared using MUST-sets, and not acceptance sets. In fact, Definition 5.7 there is titled “acceptance sets” but it actually defines MUST-sets.

6 Conclusion

In this paper we have shown that the standard characterisations of the MUST-preorder by De Nicola and Hennessy [43,55] are sound and complete also in an asynchronous setting, provided servers are enhanced with the forwarding ability. Lemma 2 shows that this lifting is always possible. We have also shown that the standard coinductive characterisation carries over to the asynchronous setting. Our results are supported by the first mechanisation of the MUST-preorder, and increase proof (i.e. code) factorisation and reusability since the alternative preorders do not need to be changed when shifting between synchronous and asynchronous semantics: it is enough to parameterise the proofs on the set of non-blocking actions. Corollary 3 states that MUST-preorder and failure refinement essentially coincide. This might spur further interest in the mechanisations of failure refinement, carried out so far in Isabelle/HOL [91,11], possibly opening up opportunities of joint efforts for automated checking.

Proof method for MUST-preorder Theorems 1, 2 and 3 endow researchers in programming languages for message-passing software with a proof method for \sqsubset_{MUST} ,

namely: to define for their calculi an LTS that enjoys the axioms of output-buffered agents with feedback. An example of this approach is Corollary 1.

Live programs have barred trees We argued that a proof of $p \text{MUST} r$ is a proof of liveness (of the client). This paper is thus de facto an example that proving liveness amounts to prove that a computational tree has a bar (identified by the predicate GOOD), and hence bar induction is a natural way to reason constructively on liveness-preserving manipulations on programs. While this fact seems to be by and large unexploited by the PL community, we believe that it may be of interest to practitioners reasoning on liveness properties in theorem provers in particular, and to the PL community at large.

Mechanisation Boreale and Gadducci [21] remark that the MUST-preorder lacks a tractable proof method. In contrast, we argue that our contributions, in particular the coinductive characterisation (Theorem 2), being fully mechanised in Coq, let practitioners pursue non-trivial results about testing preorders for real-world programming languages. To make this point, we have proved a form of code-hoisting using this characterisation. Our mechanisation lowers the barrier to entry for researchers versed into theorem provers and wishing to use testing preorders; adds to the toolkit of Coq users an alternative to the well-known (and already mechanised) bisimulation equivalence [78]; and provides a starting point for researchers willing to study testing preorders and analogous refinements within type theory. Researchers working on testing preorders may benefit from it, as there are analogies between reasoning techniques for MAY, MUST, COMPLIANCE, SHOULD, and FAIR testing. For instance Baldan et al. show with pen and paper that a technique similar to forwarding works to characterise the MAY-preorder [8].

Future work Thanks to Theorems 1, 2 and 3 we can now set out to (1) devise an axiomatisation of $\sqsubseteq_{\text{MUST}}$ for asynchronous calculi, as done in [59,22,55,56] for synchronous ones; (2) study for which asynchronous calculi $\sqsubseteq_{\text{MUST}}$ is a pre-congruence; (3) machine-check semantic models of subtyping for session types [16]; (4) study the decidability of $\sqsubseteq_{\text{MUST}}$.

More in general, given the practical relevance of asynchronous communication, it seems crucial not only to adapt the large body of theory for synchronous communication to the asynchronous setting but also to resort to machine supported reasoning to do it. This paper is meant to be a step in this direction.

Data availability The mechanised proofs have been archived on Zenodo[13].

Acknowledgments. We thank the anonymous reviewers for their helpful comments, and Roberto Amadio for useful remarks. The first author would like to acknowledge his “N+1 honoris causa” Pierre-Evariste Dagand for the systematic and useful feedback, Guillaume Geoffroy for suggesting the use of bar induction, and Hugo Herbelin for discussions on the topic. This work has received funding from the European Research Council (ERC) under the European Union’s Horizon 2020 research and innovation programme (grant agreement No. 101003349).

References

1. Aceto, L., Achilleos, A., Francalanza, A., Ingólfssdóttir, A., Lehtinen, K.: Adventures in monitorability: from branching to linear time and back again. *Proc. ACM Program. Lang.* **3**(POPL), 52:1–52:29 (2019). <https://doi.org/10.1145/3290365>, <https://doi.org/10.1145/3290365>
2. Aceto, L., Hennessy, M.: Termination, Deadlock, and Divergence. *J. ACM* **39**(1), 147–187 (1992). <https://doi.org/10.1145/147508.147527>, <https://doi.org/10.1145/147508.147527>
3. Affeldt, R., Kobayashi, N.: A Coq Library for Verification of Concurrent Programs. In: Schürmann, C. (ed.) *Proceedings of the Fourth International Workshop on Logical Frameworks and Meta-Languages, LFM@IJCAR 2004*. Cork, Ireland, July 5, 2004. *Electronic Notes in Theoretical Computer Science*, vol. 199, pp. 17–32. Elsevier (2004). <https://doi.org/10.1016/J.ENTCS.2007.11.010>, <https://doi.org/10.1016/j.entcs.2007.11.010>
4. Amadio, R.M., Castellani, I., Sangiorgi, D.: On Bisimulations for the Asynchronous pi-calculus. In: Montanari, U., Sassone, V. (eds.) *Proceedings CONCUR 96*, Pisa. *Lecture Notes in Computer Science*, vol. 1119, pp. 147–162. Springer Verlag (1996)
5. Amadio, R.M., Castellani, I., Sangiorgi, D.: On Bisimulations for the Asynchronous pi-calculus. *Theoretical Computer Science* **195**, 291–324 (1998)
6. Athalye, A.: *CoqIOA: A Formalization of IO Automata in the Coq Proof Assistant*. Master’s thesis, Massachusetts Institute of Technology (Jun 2017)
7. Aubert, C., Varacca, D.: Processes against tests: On defining contextual equivalences. *J. Log. Algebraic Methods Program.* **129**, 100799 (2022). <https://doi.org/10.1016/J.JLAMP.2022.100799>, <https://doi.org/10.1016/j.jlamp.2022.100799>
8. Baldan, P., Bonchi, F., Gadducci, F., Monreale, G.V.: Asynchronous Traces and Open Petri Nets. In: Bodei, C., Ferrari, G., Priami, C. (eds.) *Programming Languages with Applications to Biology and Security - Essays Dedicated to Pierpaolo Degano on the Occasion of His 65th Birthday*. *Lecture Notes in Computer Science*, vol. 9465, pp. 86–102. Springer (2015). https://doi.org/10.1007/978-3-319-25527-9_8, https://doi.org/10.1007/978-3-319-25527-9_8
9. Barbanera, F., de’Liguoro, U.: Two notions of sub-behaviour for session-based client/server systems. In: Kutsia, T., Schreiner, W., Fernández, M. (eds.) *Proceedings of the 12th International ACM SIGPLAN Conference on Principles and Practice of Declarative Programming*, July 26–28, 2010, Hagenberg, Austria. pp. 155–164. ACM (2010). <https://doi.org/10.1145/1836089.1836109>, <https://doi.org/10.1145/1836089.1836109>
10. Barendregt, H., Manzonetto, G.: A Lambda Calculus Satellite. *College Publications* (2022), <https://www.collegepublications.co.uk/logic/mlf/?00035>, chapter 11
11. Baxter, J., Ribeiro, P., Cavalcanti, A.: Sound reasoning in tock-CSP. *Acta Informatica* **59**(1), 125–162 (2022). <https://doi.org/10.1007/s00236-020-00394-3>, <https://doi.org/10.1007/s00236-020-00394-3>
12. Bernardi, G.: *Behavioural equivalences for Web services*. Ph.D. thesis, Trinity College Dublin (2013), <http://www.tara.tcd.ie/handle/2262/77595>
13. Bernardi, G., Castellani, I., Laforgue, P., Stefanescu, L.: Artifact for constructive characterisations of the must-preorder for asynchrony (jan 2025). <https://doi.org/10.5281/zenodo.14617145>, [https://doi.org/10.5281/zenodo.14617145](https://doi.org/10.5281/zenodo.10.5281/zenodo.14617145)

14. Bernardi, G., Hennessy, M.: Mutually Testing Processes. *Log. Methods Comput. Sci.* **11**(2) (2015). [https://doi.org/10.2168/LMCS-11\(2:1\)2015](https://doi.org/10.2168/LMCS-11(2:1)2015), [https://doi.org/10.2168/LMCS-11\(2:1\)2015](https://doi.org/10.2168/LMCS-11(2:1)2015)
15. Bernardi, G., Hennessy, M.: Using higher-order contracts to model session types. *Log. Methods Comput. Sci.* **12**(2) (2016). [https://doi.org/10.2168/LMCS-12\(2:10\)2016](https://doi.org/10.2168/LMCS-12(2:10)2016), [https://doi.org/10.2168/LMCS-12\(2:10\)2016](https://doi.org/10.2168/LMCS-12(2:10)2016)
16. Bernardi, G.T., Hennessy, M.: Modelling session types using contracts. *Math. Struct. Comput. Sci.* **26**(3), 510–560 (2016). <https://doi.org/10.1017/S0960129514000243>, <https://doi.org/10.1017/S0960129514000243>
17. Berry, G., Boudol, G.: The Chemical Abstract Machine. *Theor. Comput. Sci.* **96**(1), 217–248 (1992). [https://doi.org/10.1016/0304-3975\(92\)90185-I](https://doi.org/10.1016/0304-3975(92)90185-I), [https://doi.org/10.1016/0304-3975\(92\)90185-I](https://doi.org/10.1016/0304-3975(92)90185-I)
18. Bizjak, A., Birkedal, L., Miculan, M.: A model of countable nondeterminism in guarded type theory. In: Dowek, G. (ed.) *Rewriting and Typed Lambda Calculi*. pp. 108–123. Springer International Publishing, Cham (2014)
19. Bonchi, F., Caltais, G., Pous, D., Silva, A.: Brzozowski’s and Up-To Algorithms for Must Testing. In: Shan, C. (ed.) *Programming Languages and Systems - 11th Asian Symposium, APLAS 2013, Melbourne, VIC, Australia, December 9-11, 2013*. Proceedings. Lecture Notes in Computer Science, vol. 8301, pp. 1–16. Springer (2013). https://doi.org/10.1007/978-3-319-03542-0_1, https://doi.org/10.1007/978-3-319-03542-0_1
20. Bonchi, F., Sokolova, A., Vignudelli, V.: The Theory of Traces for Systems with Nondeterminism, Probability, and Termination. *Log. Methods Comput. Sci.* **18**(2) (2022). [https://doi.org/10.46298/LMCS-18\(2:21\)2022](https://doi.org/10.46298/LMCS-18(2:21)2022), [https://doi.org/10.46298/lmcs-18\(2:21\)2022](https://doi.org/10.46298/lmcs-18(2:21)2022)
21. Boreale, M., Gadducci, F.: Processes as formal power series: A coinductive approach to denotational semantics. *Theor. Comput. Sci.* (2006). <https://doi.org/10.1016/j.tcs.2006.05.030>, <https://doi.org/10.1016/j.tcs.2006.05.030>
22. Boreale, M., Nicola, R.D.: Testing Equivalence for Mobile Processes. *Inf. Comput.* **120**(2), 279–303 (1995). <https://doi.org/10.1006/inco.1995.1114>, <https://doi.org/10.1006/inco.1995.1114>
23. Boreale, M., Nicola, R.D., Pugliese, R.: Trace and Testing Equivalence on Asynchronous Processes. *Inf. Comput.* **172**(2), 139–164 (2002). <https://doi.org/10.1006/inco.2001.3080>, <https://doi.org/10.1006/inco.2001.3080>
24. Boudol, G.: Asynchrony and the Pi-calculus. Research Report RR-1702, INRIA (1992), <https://hal.inria.fr/inria-00076939>
25. Boudol, G.: The pi-Calculus in Direct Style. *High. Order Symb. Comput.* **11**(2), 177–208 (1998). <https://doi.org/10.1023/A:1010064516533>, <https://doi.org/10.1023/A:1010064516533>
26. Boudol, G., Lavatelli, C.: Full Abstraction for Lambda Calculus with Resources and Convergence Testing. In: Kirchner, H. (ed.) *Trees in Algebra and Programming - CAAP’96, 21st International Colloquium, Linköping, Sweden, April, 22-24, 1996*, Proceedings. Lecture Notes in Computer Science, vol. 1059, pp. 302–316. Springer (1996). https://doi.org/10.1007/3-540-61064-2_45, https://doi.org/10.1007/3-540-61064-2_45
27. Bravetti, M., Lange, J., Zavattaro, G.: Fair Refinement for Asynchronous Session Types. In: FOSSACS (2021)
28. Brede, N., Herbelin, H.: On the logical structure of choice and bar induction principles. In: 36th Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2021, Rome, Italy, June 29 - July 2, 2021. pp. 1–13. IEEE

6. CONCLUSION

- (2021). <https://doi.org/10.1109/LICS52264.2021.9470523>, <https://doi.org/10.1109/LICS52264.2021.9470523>
29. Breuvert, F., Manzonetto, G., Polonsky, A., Ruoppolo, D.: New Results on Morris's Observational Theory: The Benefits of Separating the Inseparable. In: Kesner, D., Pientka, B. (eds.) 1st International Conference on Formal Structures for Computation and Deduction, FSCD 2016, June 22-26, 2016, Porto, Portugal. LIPIcs, vol. 52, pp. 15:1–15:18. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2016). <https://doi.org/10.4230/LIPIcs.FSCD.2016.15>, <https://doi.org/10.4230/LIPIcs.FSCD.2016.15>
 30. Breuvert, F., Manzonetto, G., Ruoppolo, D.: Relational Graph Models at Work. *Log. Methods Comput. Sci.* **14**(3) (2018). [https://doi.org/10.23638/LMCS-14\(3:2\)2018](https://doi.org/10.23638/LMCS-14(3:2)2018), [https://doi.org/10.23638/LMCS-14\(3:2\)2018](https://doi.org/10.23638/LMCS-14(3:2)2018)
 31. Brookes, S.D.: On the Relationship of CCS and CSP. In: Díaz, J. (ed.) Automata, Languages and Programming, 10th Colloquium, Barcelona, Spain, July 18-22, 1983, Proceedings. *Lecture Notes in Computer Science*, vol. 154, pp. 83–96. Springer (1983). <https://doi.org/10.1007/BFb0036899>, <https://doi.org/10.1007/BFb0036899>
 32. Brookes, S.D.: Deconstructing CCS and CSP Asynchronous Communication, Fairness, and Full Abstraction (2016), <http://www.cs.cmu.edu/afs/cs/Web/People/brookes/papers/mfps16paper.pdf>
 33. Brookes, S.D., Hoare, C.A.R., Roscoe, A.W.: A Theory of Communicating Sequential Processes. *J. ACM* **31**(3), 560–599 (1984). <https://doi.org/10.1145/828.833>, <https://doi.org/10.1145/828.833>
 34. Caruana, C.: Compositional Reasoning about Actor Based Systems. Master thesis, University of Malta (2019)
 35. Castagna, G., Dezani-Ciancaglini, M., Giachino, E., Padovani, L.: Foundations of session types. In: Porto, A., López-Fraguas, F.J. (eds.) Proceedings of the 11th International ACM SIGPLAN Conference on Principles and Practice of Declarative Programming, September 7-9, 2009, Coimbra, Portugal. pp. 219–230. ACM (2009). <https://doi.org/10.1145/1599410.1599437>, <https://doi.org/10.1145/1599410.1599437>
 36. Castagna, G., Gesbert, N., Padovani, L.: A theory of contracts for Web services. *ACM Trans. Program. Lang. Syst.* **31**(5), 19:1–19:61 (2009). <https://doi.org/10.1145/1538917.1538920>, <https://doi.org/10.1145/1538917.1538920>
 37. Castellan, S., Clairambault, P., Winskel, G.: The Mays and Musts of Concurrent Strategies. In: Palmigiano, A., Sadrzadeh, M. (eds.) Samson Abramsky on Logic and Structure in Computer Science and Beyond. pp. 327–361. Springer International Publishing, Cham (2023). https://doi.org/10.1007/978-3-031-24117-8_9, https://doi.org/10.1007/978-3-031-24117-8_9
 38. Castellani, I., Hennessy, M.: Testing Theories for Asynchronous Languages. In: Arvind, V., Ramanujam, R. (eds.) Foundations of Software Technology and Theoretical Computer Science, 18th Conference, Chennai, India, December 17-19, 1998, Proceedings. *Lecture Notes in Computer Science*, vol. 1530, pp. 90–101. Springer (1998). https://doi.org/10.1007/978-3-540-49382-2_9, https://doi.org/10.1007/978-3-540-49382-2_9
 39. Cerone, A., Hennessy, M.: Process Behaviour: Formulae vs. Tests. Tech. rep., Trinity College Dublin, School of Computer Science and Statistics (2010)
 40. Cleaveland, R., Zwarico, A.E.: A Theory of Testing for Real-Time. In: LICS (1991)
 41. Cleaveland, R., Hennessy, M.: Testing Equivalence as a Bisimulation Equivalence. In: Sifakis, J. (ed.) Automatic Verification Methods for Finite State Systems, International Workshop, Grenoble, France, June 12-14, 1989, Proceedings. *Lecture*

- Notes in Computer Science, vol. 407, pp. 11–23. Springer (1989). https://doi.org/10.1007/3-540-52148-8_2, https://doi.org/10.1007/3-540-52148-8_2
42. De Nicola, R., Ferrari, G., Pugliese, R.: Klaim: a kernel language for agents interaction and mobility. *IEEE Transactions on Software Engineering* **24**(5), 315–330 (1998). <https://doi.org/10.1109/32.685256>
 43. De Nicola, R., Hennessy, M.: Testing Equivalences for Processes. *Theor. Comput. Sci.* **34**, 83–133 (1984). [https://doi.org/10.1016/0304-3975\(84\)90113-0](https://doi.org/10.1016/0304-3975(84)90113-0), [https://doi.org/10.1016/0304-3975\(84\)90113-0](https://doi.org/10.1016/0304-3975(84)90113-0)
 44. De Nicola, R., Melgratti, H.C.: Multiparty testing preorders. *Log. Methods Comput. Sci.* **19**(1) (2023). [https://doi.org/10.46298/lmcs-19\(1:1\)2023](https://doi.org/10.46298/lmcs-19(1:1)2023), [https://doi.org/10.46298/lmcs-19\(1:1\)2023](https://doi.org/10.46298/lmcs-19(1:1)2023)
 45. De Nicola, R., Pugliese, R.: Linda-based applicative and imperative process algebras. *Theor. Comput. Sci.* **238**(1-2), 389–437 (2000). [https://doi.org/10.1016/S0304-3975\(99\)00339-4](https://doi.org/10.1016/S0304-3975(99)00339-4), [https://doi.org/10.1016/S0304-3975\(99\)00339-4](https://doi.org/10.1016/S0304-3975(99)00339-4)
 46. Dreyer, D., Neis, G., Birkedal, L.: The impact of higher-order state and control effects on local relational reasoning. *J. Funct. Program.* **22**(4-5), 477–528 (2012). <https://doi.org/10.1017/S095679681200024X>, <https://doi.org/10.1017/S095679681200024X>
 47. Dummett, M.: *Elements of Intuitionism*. Oxford logic guides, Clarendon Press (2000), <https://books.google.fr/books?id=JVFzknbGBVAC>
 48. Fournet, C., Gonthier, G.: The join calculus: A language for distributed mobile programming. In: Barthe, G., Dybjer, P., Pinto, L., Saraiva, J. (eds.) *Applied Semantics*. pp. 268–332. Springer Berlin Heidelberg, Berlin, Heidelberg (2002)
 49. Francalanza, A.: A theory of monitors. *Inf. Comput.* **281**, 104704 (2021). <https://doi.org/10.1016/j.ic.2021.104704>, <https://doi.org/10.1016/j.ic.2021.104704>
 50. Fridlender, D.: An Interpretation of the Fan Theorem in Type Theory. In: *TYPES* (1998), https://doi.org/10.1007/3-540-48167-2_7
 51. Frumin, D., Krebbers, R., Birkedal, L.: ReLoC: A Mechanised Relational Logic for Fine-Grained Concurrency. In: Dawar, A., Grädel, E. (eds.) *Proceedings of the 33rd Annual ACM/IEEE Symposium on Logic in Computer Science, LICS 2018, Oxford, UK, July 09-12, 2018*. pp. 442–451. ACM (2018). <https://doi.org/10.1145/3209108.3209174>, <https://doi.org/10.1145/3209108.3209174>
 52. Gay, S.J., Hole, M.: Subtyping for session types in the pi calculus. *Acta Informatica* **42**(2-3), 191–225 (2005). <https://doi.org/10.1007/s00236-005-0177-z>, <https://doi.org/10.1007/s00236-005-0177-z>
 53. van Glabbeek, R.: Just Testing. In: Kupferman, O., Sobocinski, P. (eds.) *Foundations of Software Science and Computation Structures - 26th International Conference, FoSSaCS 2023, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2023, Paris, France, April 22-27, 2023, Proceedings*. *Lecture Notes in Computer Science*, vol. 13992, pp. 498–519. Springer (2023). https://doi.org/10.1007/978-3-031-30829-1_24, https://doi.org/10.1007/978-3-031-30829-1_24
 54. Hennessy, M.: Acceptance Trees. *J. ACM* **32**(4), 896–928 (1985). <https://doi.org/10.1145/4221.4249>, <https://doi.org/10.1145/4221.4249>
 55. Hennessy, M.: *Algebraic theory of processes*. MIT Press series in the foundations of computing, MIT Press (1988)
 56. Hennessy, M.: A fully abstract denotational semantics for the pi-calculus. *Theor. Comput. Sci.* **278**(1-2), 53–89 (2002). [https://doi.org/10.1016/S0304-3975\(00\)00331-5](https://doi.org/10.1016/S0304-3975(00)00331-5), [https://doi.org/10.1016/S0304-3975\(00\)00331-5](https://doi.org/10.1016/S0304-3975(00)00331-5)

6. CONCLUSION

57. Hennessy, M.: The security pi-calculus and non-interference. *J. Log. Algebraic Methods Program.* (2005). <https://doi.org/10.1016/j.jlap.2004.01.003>, <https://doi.org/10.1016/j.jlap.2004.01.003>
58. Hennessy, M.: *A distributed Pi-calculus*. Cambridge University Press (2007)
59. Hennessy, M., Ingólfssdóttir, A.: Communicating Processes with Value-passing and Assignments. *Formal Aspects Comput.* **5**(5), 432–466 (1993). <https://doi.org/10.1007/BF01212486>, <https://doi.org/10.1007/BF01212486>
60. Hennessy, M., Plotkin, G.D.: A Term Model for CCS. In: Dembinski, P. (ed.) *Mathematical Foundations of Computer Science 1980 (MFCS'80)*, Proceedings of the 9th Symposium, Rydzyna, Poland, September 1-5, 1980. *Lecture Notes in Computer Science*, vol. 88, pp. 261–274. Springer (1980). <https://doi.org/10.1007/BFb0022510>, <https://doi.org/10.1007/BFb0022510>
61. Hennessy, M., Riely, J.: Information flow vs. resource access in the asynchronous pi-calculus. *ACM Trans. Program. Lang. Syst.* **24**(5), 566–591 (2002). <https://doi.org/10.1145/570886.570890>, <https://doi.org/10.1145/570886.570890>
62. Hirschhoff, D., Jaber, G., Prebet, E.: Deciding Contextual Equivalence of ν -Calculus with Effectful Contexts. In: Kupferman, O., Sobocinski, P. (eds.) *Foundations of Software Science and Computation Structures - 26th International Conference, FoSSaCS 2023, Held as Part of the European Joint Conferences on Theory and Practice of Software, ETAPS 2023, Paris, France, April 22-27, 2023, Proceedings*. *Lecture Notes in Computer Science*, vol. 13992, pp. 24–45. Springer (2023). https://doi.org/10.1007/978-3-031-30829-1_2, https://doi.org/10.1007/978-3-031-30829-1_2
63. Hoare, C.A.R.: *Communicating Sequential Processes (Reprint)*. *Commun. ACM* (1983)
64. Honda, K., Tokoro, M.: An Object Calculus for Asynchronous Communication. In: America, P. (ed.) *ECOOP'91 European Conference on Object-Oriented Programming*, Geneva, Switzerland, July 15-19, 1991, Proceedings. *Lecture Notes in Computer Science*, vol. 512, pp. 133–147. Springer (1991). <https://doi.org/10.1007/BFb0057019>, <https://doi.org/10.1007/BFb0057019>
65. Honda, K., Yoshida, N., Carbone, M.: Multiparty Asynchronous Session Types. In: Necula, G.C., Wadler, P. (eds.) *POPL*. pp. 273–284. ACM Press, New York (2008)
66. Honda, K., Yoshida, N., Carbone, M.: Multiparty Asynchronous Session Types. *Journal of ACM* **63**(1), 9:1–9:67 (2016)
67. Intrigila, B., Manzonetto, G., Polonsky, A.: Degrees of extensionality in the theory of Böhm trees and Sallé's conjecture. *Log. Methods Comput. Sci.* **15**(1) (2019). [https://doi.org/10.23638/LMCS-15\(1:6\)2019](https://doi.org/10.23638/LMCS-15(1:6)2019), [https://doi.org/10.23638/LMCS-15\(1:6\)2019](https://doi.org/10.23638/LMCS-15(1:6)2019)
68. Kleene, S.C., Vesley, R.E.: *The Foundations of Intuitionistic Mathematics: Especially in Relation to Recursive Functions*. *Studies in logic and the foundations of mathematics*, North-Holland Publishing Company (1965), <https://books.google.fr/books?id=2EHVxQEACAAJ>
69. Koutavas, V., Tzevelekos, N.: Fully Abstract Normal Form Bisimulation for Call-by-Value PCF. In: *LICS* (2023)
70. Krebbers, R., Timany, A., Birkedal, L.: Interactive proofs in higher-order concurrent separation logic. In: Castagna, G., Gordon, A.D. (eds.) *Proceedings of the 44th ACM SIGPLAN Symposium on Principles of Programming Languages, POPL 2017, Paris, France, January 18-20, 2017*. pp. 205–217. ACM (2017). <https://doi.org/10.1145/3009837.3009855>, <https://doi.org/10.1145/3009837.3009855>

-
71. Lamport, L.: Specifying Systems, The TLA+ Language and Tools for Hardware and Software Engineers. Addison-Wesley (2002), <http://research.microsoft.com/users/lamport/tla/book.html>
72. Laneve, C., Padovani, L.: The *Must* Preorder Revisited. In: Caires, L., Vasconcelos, V.T. (eds.) CONCUR 2007 - Concurrency Theory, 18th International Conference, CONCUR 2007, Lisbon, Portugal, September 3-8, 2007, Proceedings. Lecture Notes in Computer Science, vol. 4703, pp. 212–225. Springer (2007). https://doi.org/10.1007/978-3-540-74407-8_15, https://doi.org/10.1007/978-3-540-74407-8_15
73. Milner, R.: Functions as Processes. In: Paterson, M. (ed.) Automata, Languages and Programming, 17th International Colloquium, ICALP90, Warwick University, England, UK, July 16-20, 1990, Proceedings. Lecture Notes in Computer Science, vol. 443, pp. 167–180. Springer (1990). <https://doi.org/10.1007/BFb0032030>, <https://doi.org/10.1007/BFb0032030>
74. Milner, R.: Communicating and Mobile Systems - the Pi-Calculus. Cambridge University Press (1999)
75. Morris, J.H.: Lambda-calculus models of programming languages. Ph.D. thesis, Massachusetts Institute of Technology (1969), <https://dspace.mit.edu/handle/1721.1/64850>
76. Nicola, R.D., Hennessy, M.: CCS without tau's. In: Ehrig, H., Kowalski, R.A., Levi, G., Montanari, U. (eds.) TAPSOFT'87: Proceedings of the International Joint Conference on Theory and Practice of Software Development, Pisa, Italy, March 23-27, 1987, Volume 1: Advanced Seminar on Foundations of Innovative Software Development I and Colloquium on Trees in Algebra and Programming (CAAP'87). Lecture Notes in Computer Science, vol. 249, pp. 138–152. Springer (1987). https://doi.org/10.1007/3-540-17660-8_53, https://doi.org/10.1007/3-540-17660-8_53
77. Palamidessi, C.: Comparing the Expressive Power of the Synchronous and Asynchronous *pi*-calculi. Mathematical Structures in Computer Science **13**(5), 685–719 (2003). <https://doi.org/10.1017/S0960129503004043>
78. Pous, D.: Coinduction All the Way Up. In: Grohe, M., Koskinen, E., Shankar, N. (eds.) Proceedings of the 31st Annual ACM/IEEE Symposium on Logic in Computer Science, LICS '16, New York, NY, USA, July 5-8, 2016. pp. 307–316. ACM (2016). <https://doi.org/10.1145/2933575.2934564>, <https://doi.org/10.1145/2933575.2934564>
79. Prebet, E.: Functions and References in the Pi-Calculus: Full Abstraction and Proof Techniques. In: Bojanczyk, M., Merelli, E., Woodruff, D.P. (eds.) 49th International Colloquium on Automata, Languages, and Programming, ICALP 2022, July 4-8, 2022, Paris, France. LIPIcs, vol. 229, pp. 130:1–130:19. Schloss Dagstuhl - Leibniz-Zentrum für Informatik (2022). <https://doi.org/10.4230/LIPIcs.ICALP.2022.130>, <https://doi.org/10.4230/LIPIcs.ICALP.2022.130>
80. Pugliese, R.: Semantic Theories for Asynchronous Languages. Ph.D. thesis, Università di Roma "La Sapienza" (1996)
81. Rahli, V., Bickford, M., Cohen, L., Constable, R.L.: Bar Induction is Compatible with Constructive Type Theory. J. ACM **66**(2), 13:1–13:35 (2019). <https://doi.org/10.1145/3305261>, <https://doi.org/10.1145/3305261>
82. Ravara, A., Resende, P., Vasconcelos, V.T.: An Algebra of Behavioural Types. Inf. Comput. **212**, 64–91 (2012)
83. Rensink, A., Vogler, W.: Fair testing. Inf. Comput. **205**(2), 125–198 (2007). <https://doi.org/10.1016/j.ic.2006.06.002>, <https://doi.org/10.1016/j.ic.2006.06.002>

6. CONCLUSION

84. Sabel, D., Schmidt-Schauß, M.: A call-by-need lambda calculus with locally bottom-avoiding choice: context lemma and correctness of transformations. *Math. Struct. Comput. Sci.* **18**(3), 501–553 (2008). <https://doi.org/10.1017/S0960129508006774>, <https://doi.org/10.1017/S0960129508006774>
85. Sangiorgi, D.: *Introduction to Bisimulation and Coinduction*. Cambridge University Press (2011)
86. Sangiorgi, D.: Asynchronous pi-calculus at Work: The Call-by-Need Strategy. In: Alvim, M.S., Chatzikokolakis, K., Olarte, C., Valencia, F. (eds.) *The Art of Modelling Computational Systems: A Journey from Logic and Concurrency to Security and Privacy - Essays Dedicated to Catuscia Palamidessi on the Occasion of Her 60th Birthday*. *Lecture Notes in Computer Science*, vol. 11760, pp. 33–49. Springer (2019). https://doi.org/10.1007/978-3-030-31175-9_3, https://doi.org/10.1007/978-3-030-31175-9_3
87. Sangiorgi, D., Walker, D.: *The Pi-Calculus - a Theory of Mobile Processes*. Cambridge University Press (2001)
88. Schmidt-Schauß, M., Sabel, D.: Correctly Implementing Synchronous Message Passing in the Pi-Calculus By Concurrent Haskell’s MVars. In: Dardha, O., Rot, J. (eds.) *Proceedings Combined 27th International Workshop on Expressiveness in Concurrency and 17th Workshop on Structural Operational Semantics, EXPRESS/SOS 2020, and 17th Workshop on Structural Operational Semantics Online*, 31 August 2020. *EPTCS*, vol. 322, pp. 88–105 (2020). <https://doi.org/10.4204/EPTCS.322.8>, <https://doi.org/10.4204/EPTCS.322.8>
89. Schmidt-Schauß, M., Sabel, D., Dallmeyer, N.: Sequential and Parallel Improvements in a Concurrent Functional Programming Language. In: Sabel, D., Thiemann, P. (eds.) *Proceedings of the 20th International Symposium on Principles and Practice of Declarative Programming, PPDP 2018, Frankfurt am Main, Germany, September 03-05, 2018*. pp. 20:1–20:13. ACM (2018). <https://doi.org/10.1145/3236950.3236952>, <https://doi.org/10.1145/3236950.3236952>
90. Selinger, P.: First-Order Axioms for Asynchrony. In: Mazurkiewicz, A.W., Winkowski, J. (eds.) *CONCUR ’97: Concurrency Theory, 8th International Conference, Warsaw, Poland, July 1-4, 1997, Proceedings*. *Lecture Notes in Computer Science*, vol. 1243, pp. 376–390. Springer (1997). https://doi.org/10.1007/3-540-63141-0_26, https://doi.org/10.1007/3-540-63141-0_26
91. Taha, S., Ye, L., Wolff, B.: HOL-CSP Version 2.0. *Archive of Formal Proofs* (April 2019)
92. Taha, S., Wolff, B., Ye, L.: The hol-csp refinement toolkit. *Archive of Formal Proofs* (November 2020), https://isa-afp.org/entries/CSP_RefTK.html, Formal proof development
93. Tanti, E., Francalanza, A.: *Towards Sound Refactoring in Erlang* (2015), <https://api.semanticscholar.org/CorpusID:63046364>
94. Thati, P.: *A Theory of Testing for Asynchronous Concurrent Systems*. Ph.D. thesis, University of Illinois Urbana-Champaign, USA (2003), <https://hdl.handle.net/2142/81630>
95. Turon, A.J., Thamsborg, J., Ahmed, A., Birkedal, L., Dreyer, D.: Logical Relations for Fine-Grained Concurrency. In: Giacobazzi, R., Cousot, R. (eds.) *The 40th Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL ’13, Rome, Italy - January 23 - 25, 2013*. pp. 343–356. ACM (2013). <https://doi.org/10.1145/2429069.2429111>, <https://doi.org/10.1145/2429069.2429111>

A Further related works

Contextual preorders in functional languages Morris preorder is actively studied in the pure λ -calculus [29,30,67,10], λ -calculus with references [79,62], in PCF [69] as well as in languages supporting shared memory concurrency [95], and mutable references [46]. The more sophisticated the languages, the more intricate and larger the proofs. The need for mechanisation became thus apparent, in particular to prove that complex logical relations defined in the framework Iris (implemented in Coq) are sound, *i.e.* included in the preorder [70,51]. The paper [7] provides a framework to study contextual equivalences in the setting of process calculi. It is worth noting, though, that as argued in [25, Section 3], Morris equivalence coincides with MAY-equivalence, at least if the operational semantics at hand enjoys the Church-Rosser property. In fact [26] define Morris preorder literally as a testing one, via tests for convergence. The studies of the MUST-preorder in process calculi can thus be seen as providing proof methods to adapt Morris equivalence to nondeterministic settings, and using contexts that are really external observers. To sum up, one may say that Morris equivalence coincides with MAY-equivalence when nondeterminism is confluent and all states are viewed as accepting states, while it coincides with MUST equivalence in the presence of true nondeterminism and when only successful states are viewed as accepting states.

In the setting of nondeterministic and possibly concurrent applicative programming languages [89,88,18], also a contextual preorder based on may and must-termination has been studied [84,18]. Our preorder \preceq_{cnv} is essentially a generalisation of the must-termination preorder of [84] to traces of visible actions.

Theories for synchronous semantics Both [72] and [36] employed LTSs as a model of contracts for web-services (*i.e.* WSCL), and the MUST-preorder as refinement for contracts. The idea is that a search engine asked to look for a service described by a contract p_1 can actually return a service that has a contract p_2 , provided that $p_1 \sqsubseteq_{\text{MUST}} p_2$.

The MUST-preorder for *clients* proposed by [14] has partly informed the theory of monitors by [1], in particular the study of preorders for monitors by [49]. Our results concern LTSs that are more general than those of monitors, and thus our code could provide the basis to mechanise the results of [1].

The first subtyping relation for binary session types was presented in [52] using a syntax-oriented definition. The semantic model of that subtyping is a refinement very similar to the MUST-preorder. The idea is to treat types as CCS terms, assign them an LTS [35,9,82,16], and use the resulting testing preorders as semantic models of the subtyping. In the setting of coinductively defined higher-order session types, the correspondence is implicitly addressed in [35]. In the setting of recursive higher-order session types, it is given by Theorem 4.10 of [15].

A. FURTHER RELATED WORKS

Models of asynchrony While synchronous (binary) communication requires the simultaneous occurrence of a send and a receive action, asynchronous communication allows a delay between a send action and the corresponding receive action. Different models of asynchrony exist, depending on which medium is assumed for storing messages in transit. In this paper, following the early work on the asynchronous π -calculus [64,24,4], we assume the medium to be an unbounded unordered mailbox, shared by all processes. Thus, no process needs to wait to send a message, namely the send action is non-blocking. This model of communication is best captured via the output-buffered agents with feedback of [90]. The early style LTS of the asynchronous π -calculus is a concrete example of this kind of LTSs. A similar global unordered mailbox is used also in Chapter 5 of [94], by [32], which relies explicitly on a mutable global state, and by [77], which manipulates it via two functions *get* and *set*.

More deterministic models of asynchrony are obtained assigning a data structure to every channel. For example [65,66] use an even more deterministic model in which each ordered pair of processes is assigned a dedicated channel, equipped with an *ordered* queue. Hence, messages along such channels are received in the same order in which they were sent. This model is used for asynchronous session calculi, and mimics the communication mode of the TCP/IP protocol. The obvious research question here is how to adapt our results to the different communication mechanisms and different classes of LTSs. For instance, both [93] and [34] define LTSs for ERLANG. We will study whether at least one of these LTSs is an instance of output-buffered agents with feedback. If this is not the case, we will first try to adapt our results to ERLANG LTSs.

MAY-preorder MAY testing and the MAY-preorder, have been widely studied in asynchronous settings. The first characterisation for ACCS appeared in [38] and relies on comparing traces and asynchronous traces of servers. Shortly after [23] presented a characterisation based on operation on traces. A third characterisation appeared in [8], where the saturated LTS \longrightarrow_s is essentially our \longrightarrow_{fw} . That characterisation supports our claim that results about synchronous semantics are true also for asynchronous ones, modulo forwarding. Compositionality of trace inclusion, i.e. the alternative characterisation of the MAY-preorder, has been partly investigated in Coq by [6] in the setting of IO-automata. The MAY-preorder has also been studied in the setting of actor languages by [34,93].

Fairness Van Glabbeek [53] argues that by amending the semantics of parallel composition (i.e. the scheduler) different notions of fairness can be embedded in the MUST-preorder. We would like to investigate which notion of fairness makes the MUST-preorder coincide with the FAIR-preorder of [83].

Mutable state Prebet [79] has recently shown an encoding of the asynchronous π -calculus into a λ -calculus with references, which captures Morris equivalence via a bisimulation. This renders vividly the intuition that output-buffered agents manipulate a shared common state. We therefore see our work also as an analysis of the MUST-preorder for a language in which programs manipulate a global




Fig. 8. The state transition system of client-server system.

mutable store. Since the store is what contains output messages, and our formal development shows that only outputs are observable, our results suggest that characterisations of testing preorders for impure programming languages should predicate over the content of the mutable store, *i.e.* the values written by programs. Another account of π -calculus synchronisation via a functional programming language is provided in [88], that explains how to use Haskell M-VARS to implement π -calculus message passing.

B Bar-Induction

Section B.1 is an informal introduction to the intuitions behind bar induction, together with a comment on the admissibility of this principle in Coq. In Section B.2 we gather the properties of the predicates \Downarrow and MUST that we use throughout our technical development.

B.1 A visual introduction

We explain the difference between *extensional* definitions of predicates and *intensional* ones, by discussing how the two different approaches make us reason on computational trees.

Suppose that we have a client-server system $p \parallel r$ and that we want to prove either $p \text{MUST} r$ or $p \text{MUST}_i r$. For both proofs, what matters is the state transition system (STS) of $p \parallel r$, *i.e.* the computation steps performed by the client-server system at issue. In fact it is customary to treat this STS as a computational tree, as done for instance in the proofs of [55, Lemma 4.4.12] and [39, Theorem 2.3.3]. In the rest of this subsection we discuss the tree depicted in Figure 8. It contains three maximal computations, the middle one being infinite. In the figures of this subsection, the states in which the client is successful (*i.e.* in the predicate GOOD) contain the symbol \checkmark .




Fig. 9. Extensional approach: finding successful prefixes in every maximal path of the computational tree.

The extensional approach To prove $p \text{ MUST } r$, the extensional definition of MUST requires checking that every maximal path in the tree in Figure 8 starts with a finite prefix that leads to a successful state. The proof that $p \text{ MUST } r$ amounts to looking for a suitable prefix maximal path by maximal path, via a loop whose iterations are suggested in Figure 9. There at every iteration a different maximal path (highlighted by dashed arrows) is checked, and each time a successful prefix is found (indicated by a red arrow), the loop moves on to the next maximal path. Once a maximal path is explored, it remains dashed, to denote that there a succesful prefix has been found. The first iteration looks for a successful prefix in the left-most maximal path, while the last iteration looks for a successful prefix in the right-most path. In the current example the loop terminates because the tree in Figure 8 has conveniently a finite number of maximal paths, but in general the mathematical reasoning has to deal with an infinite amount of maximal path. An archetypal example is the tree in Figure 11: it has countably many maximal paths, each one starting with a successful prefix.

The intensional approach Consider now the predicate MUST_i - which is defined intensionally - and a proof that $p \text{ MUST}_i r$. The base case of MUST_i ensures that all the nodes that contain a successful client (i.e. that satisfies the predicate Q_2 , defined on line 553 of the submission) are in MUST_i . Pictorially, this is the step from (1) to (2) in Figure 10, where the nodes in MUST_i are drawn using dashed borders, and the freshly added ones are drawn in red. Once the base case is established, the inductive rule of MUST_i ensures that any node that inevitably goes to nodes that are in MUST_i , is also in the predicate MUST_i . This leads to the step from (2) to (3) and then from (3) to (4). Note that the argument is concise, for in the tree the depth at which successful states can be found is finite. In general though is may not be the case. The tree in Figure 11 is again the archetypal example: every maximal path there contains a finite prefix that leads to a successful state, but there is no upper bound on the length on those prefixes.




Fig. 10. Intensional approach: visiting the tree bottom-up, starting from the bar.

Do extensional and intensional predicates coincide? Extensional and intensional definitions make us reason on computational trees in strikingly different fashions: *extensionally* we reason maximal path by maximal path, while *intensionally* we reason bottom-up, starting from the nodes in a predicate that bars the tree.¹¹ It is natural to ask whether reasoning in these different manners ultimately leads to the same outcomes. In our setting this amounts to proving that the predicates MUST and MUST_i are logically equivalent, and similarly for the convergence predicates \downarrow and \downarrow_i . The proof that $p \text{MUST}_i r$ implies $p \text{MUST} r$ is - obviously - by induction on the derivation of $p \text{MUST}_i r$. Proving that the extensional predicates imply the intensional ones is, on the other hand, delicate, because we may have to deal with unbounded structures. The tree in Figure 11 is once more the archetypal example: it has countably many maximal paths, and there is no upper bound on the depth at which successful states (i.e. nodes in the *bar*) are found.

In classical logic one can prove that $p \text{MUST} r$ implies $p \text{MUST}_i r$ by contradiction. As we wish to avoid this reasoning principle, the only tool we have is the axiom of Bar induction, which states exactly that under suitable hypotheses, extensionally defined predicates imply their intensionally defined counter-parts.

¹¹ Whence the name *bar*-induction.




Fig. 11. An infinite branching computational tree where the *bar* ✓ is at unbounded depth.

Admissibility. To show that the principle is admissible, we prove that it follows from the Classical Epsilon (CE) axiom of the Coq standard library. In short, CE gives a choice function ϵ such that if p is a proof of $\exists x : A, Px$, then $\epsilon(p)$ is an element of A such that $P(\epsilon(p))$ holds. It implies Excluded Middle, and thus classical reasoning, because $A \vee \neg A$ is equivalent to $\exists b : \text{bool}, (b = \text{true} \wedge A) \vee (b = \text{false} \wedge \neg A)$. Since CE is guaranteed by the Coq developers to be admissible, our statement of bar induction is also admissible.

B.2 Properties of inductively defined predicates

Convergence along traces is preserved by the strong transitions \longrightarrow .

Lemma 9. *In every LTS, for every $p, p' \in C$ and $s \in \text{Act}^*$ the following facts are true,*

1. *if $p \Downarrow s$ and $p \xrightarrow{\tau} p'$ then $p' \Downarrow s$,*
2. *for every $\mu \in \text{Act}$. $p \Downarrow \mu.s$ and $p \xrightarrow{\mu} p'$ imply $p \Downarrow s$.*

Lemma 10. *For every $s \in \text{Act}^*$ and $p \in \text{ACCS}$, if $p \Downarrow s$ then $|\{q \mid p \xrightarrow{s} q\}| \in \mathbb{N}$.*

The hypothesis of convergence in Lemma 10 is necessary. This is witnessed by the process $p = \text{rec}x.(x \parallel \bar{a})$, which realises an ever lasting addition of a message to the mailbox:

$$p \xrightarrow{\tau} p \parallel \bar{a} \xrightarrow{\tau} p \parallel \bar{a} \parallel \bar{a} \xrightarrow{\tau} p \parallel \bar{a} \parallel \bar{a} \parallel \bar{a} \xrightarrow{\tau} \dots$$

In more general languages also image-finiteness may fail. An example is given on page 267 of [60].

The predicate MUST_i is preserved by atoms freely changing their locations in systems. This is coherent with the intuition that the mailbox is a global and shared one. For instance the systems $a.0 \parallel \bar{d} \parallel d.1$ and $a.0 \parallel d.1 \parallel \bar{d}$, which in the mechanisation are respectively

$(\text{pr_par } (\text{pr_input } a \text{ pr_nil}) (\text{pr_out } d), \text{pr_input } d \text{ pr_succes})$

and

$(\text{pr_input } a \text{ pr_nil}, \text{pr_par } (\text{pr_input } a \text{ pr_succes}) (\text{pr_out } d))$

have the same mailbox, namely \bar{d} .

The predicate MUST_i enjoys three useful properties: it ensures convergence of servers interacting with clients that are not in a good state; it is preserved by internal computation with servers; and it is preserved also by interactions with unhappy clients. The arguments to show these facts are by rule induction on the hypothesis $p \text{MUST}_i r$. The last fact is a consequence of a crucial property of MUST_i , namely Lemma 14.

Lemma 11. *Let $\mathcal{L}_A \in \text{FWD}$ and $\mathcal{L}_B \in \text{FDB}$. For every $p \in A$, $r \in B$ we have that $p \text{MUST}_i r$ implies that $p \downarrow_i$ or $\text{GOOD}(r)$.*

Lemma 12. *Let $\mathcal{L}_A \in \text{FWD}$ and $\mathcal{L}_B \in \text{FDB}$. For every $p, p' \in A$, $r \in B$ we have that $p \text{MUST}_i r$ and $p \xrightarrow{\tau} p'$ imply $q \text{MUST}_i r$.*

Lemma 13. *For every $\mathcal{L}_B \in \text{OBA}$, $r \in B$ and name $a \in \mathcal{N}$ such that $p \xrightarrow{\bar{a}} p'$ then $\text{GOOD}(p)$ iff $\text{GOOD}(p')$.*

Proof. This is a property of Good, more specifically `good_preserved_by_lts_output` and `good_preserved_by_lts_output_converse`.

The next lemma states that when reasoning on MUST_i , outputs can be freely moved from the client to the server side of systems, if servers have the forwarding ability. Its proof uses *all* the axioms for output-buffered agents with feedback, and the lemma itself is used in the proof of the main result on acceptance bsets, namely Lemma 19.

Lemma 14 (Output swap). *Let $\mathcal{L}_A \in \text{FWD}$ and $\mathcal{L}_B \in \text{FDB}$. For every $p_1, p_2 \in A$, every $r_1, r_2 \in B$ and name $a \in \mathcal{N}$ such that $p_1 \xrightarrow{\bar{a}} p_2$ and $r_1 \xrightarrow{\bar{a}} r_2$, if $p_1 \text{MUST}_i r_2$ then $p_2 \text{MUST}_i r_1$.*

Proof. We proceed by induction on $p_1 \text{MUST}_i r_2$. In the base case $p_1 \text{MUST}_i r_2$ is derived using the rule [AXIOM] and thus $\text{GOOD}(r_2)$. Lemma 13 implies that $\text{GOOD}(r_1)$, and so we prove $p_2 \text{MUST}_i r_1$ using rule [AXIOM]. We are done with the base case.

In the inductive case, the hypothesis $p_2 \text{MUST}_i r_1$ has been derived via an rule [IND-RULE], and we therefore know the following facts:

B. BAR-INDUCTION

1. $p_1 \parallel r_2 \xrightarrow{\tau} \hat{p} \parallel \hat{r}$, and
2. For every p', r' such that $p_1 \parallel r_2 \xrightarrow{\tau} p' \parallel r'$ we have that $p' \text{ MUST}_i r'$.

We prove $p_2 \text{ MUST}_i r_1$ by applying rule [IND-RULE]. In turn this requires us to show that

- (i) $p_2 \parallel r_1 \xrightarrow{\tau}$, and that
- (ii) for each p' and r' such that $p_2 \parallel r_1 \xrightarrow{\tau} p' \parallel r'$, we have $p' \text{ MUST}_i r'$.

We prove (i). The argument starts with a case analysis on how the transition (1) has been derived. There are the following three cases:

- [S-Srv] a τ -transition performed by the server such that $p_1 \xrightarrow{\tau} \hat{p}$ and that $\hat{r} = r_2$, or
- [S-Clt] a τ -transition performed by the client such that $r_2 \xrightarrow{\tau} \hat{r}$ and that $\hat{p} = p_1$, or
- [S-com] an interaction between the server p_1 and the client r_2 such that $p_1 \xrightarrow{\mu} \hat{p}$ and that $r_2 \xrightarrow{\bar{\mu}} \hat{r}$.

In case [S-Srv] we use the OUTPUT-TAU axiom together with the transitions $p_1 \xrightarrow{\bar{a}} p_2$ and $p_1 \xrightarrow{\tau} \hat{p}$ to obtain that either:

- there exists a p_3 such that $p_2 \xrightarrow{\tau} p_3$ and $\hat{p} \xrightarrow{\bar{a}} p_3$, or
- $p_2 \xrightarrow{a} p_3$.

In the first case $p_2 \xrightarrow{\tau} p_3$ let us construct the transition $p_2 \parallel r_1 \xrightarrow{\tau} p_3 \parallel r_1$ as required. In the second case recall that by hypothesis $r_1 \xrightarrow{\bar{a}} r_2$, and thus the transition $p_2 \xrightarrow{a} \hat{p}$ and rule [S-COM] let us construct the desired reduction $p_2 \parallel r_1 \xrightarrow{\tau} \hat{p} \parallel r_2$.

In case [S-CLT] we use the OUTPUT-COMMUTATIVITY axiom together with the transitions $r_1 \xrightarrow{\bar{a}} r_2 \xrightarrow{\tau} \hat{r}$ to obtain a r_3 such that $r_1 \xrightarrow{\tau} r_3 \xrightarrow{\bar{a}} \hat{r}$ and it follows that there exists the silent move $p_2 \parallel r_1 \xrightarrow{\tau} p_2 \parallel r_3$.

In case [S-COM] we have that $p_1 \xrightarrow{\mu} \hat{p}$ and $r_2 \xrightarrow{\bar{\mu}} \hat{r}$. We distinguish whether $\mu = \bar{a}$ or not. If $\mu = \bar{a}$ then observe that $r_1 \xrightarrow{\bar{a}} r_2 \xrightarrow{a} \hat{r}$. Since by hypothesis $r_1, r_2 \in B$ and $\mathcal{L}_B \in \text{FDB}$ we apply FEEDBACK axiom to these transitions and obtain $r_1 \xrightarrow{\tau} \hat{r}$. An application of [S-COM] let us construct the desired transition $p_2 \parallel r_1 \xrightarrow{\tau} p_2 \parallel \hat{r}$.

If $\mu \neq \bar{a}$ we apply the OUTPUT-CONFLUENCE axiom to the transitions $p_1 \xrightarrow{\bar{a}} p_2$ and $p_1 \xrightarrow{\mu} \hat{p}$ to obtain a p_3 such that $p_2 \xrightarrow{\mu} p_3$ and $\hat{p} \xrightarrow{\bar{a}} p_3$. We then apply the OUTPUT-COMMUTATIVITY axiom to obtain $r_1 \xrightarrow{\bar{\mu}} r_3 \xrightarrow{\bar{a}} \hat{r}$ for some r_3 . Finally, we have the desired $p_2 \parallel r_1 \xrightarrow{\tau} \hat{p} \parallel r_3$ thanks to the existence of an interaction between p_2 and r_1 that follows from $p_2 \xrightarrow{\mu} p_3$ and $r_1 \xrightarrow{\bar{\mu}} r_3$. This concludes the proof of (i).

We now tackle (ii). First of all, note that the inductive hypothesis states the following fact,

For every p', r', p_0 and r_0 , such that $p_1 \parallel r_2 \xrightarrow{\tau} p' \parallel r'$, $p' \xrightarrow{\bar{a}} p_0$ and $r_0 \xrightarrow{\bar{a}} r'$ then $p_0 \text{ MUST}_i r_0$.

Fix a transition

$$p_2 \parallel r_1 \xrightarrow{\tau} p' \parallel r',$$

we must show $p' \text{ MUST}_i r'$. We proceed by case analysis on the rule used to derive the transition at issue, and the cases are as follows,

- (a) a τ -transition performed by the server such that $p_2 \xrightarrow{\tau} p'$ and that $r' = r_1$,
or
- (b) a τ -transition performed by the client such that $r_1 \xrightarrow{\tau} r'$ and that $p' = p_2$,
or
- (c) an interaction between the server p_2 and the client r_1 such that $p_2 \xrightarrow{\mu} p'$ and that $r_1 \xrightarrow{\bar{\mu}} r'$.

In case (a) we have $p_2 \xrightarrow{\tau} p'$ and $r' = r_1$ and hence we must show $p' \text{ MUST}_i r_1$.

We apply the OUTPUT-COMMUTATIVITY axiom to the transitions $p_1 \xrightarrow{\bar{a}} p_2 \xrightarrow{\tau} p'$ to obtain a p_3 such that $p_1 \xrightarrow{\tau} p_3 \xrightarrow{\bar{a}} p'$. We apply the inductive hypothesis with $p' = p_3$, $r' = r_2$, $p_0 = p'$ and $r_0 = r_1$ and obtain $p_2 \text{ MUST}_i r_1$ as required.

In case (b) we have $r_1 \xrightarrow{\tau} r'$ and $p' = p_2$, we therefore must show $p_2 \text{ MUST}_i r'$.

We apply the OUTPUT-TAU axiom to the transitions $r_1 \xrightarrow{\tau} r'$ and $r_1 \xrightarrow{\bar{a}} r_2$ to obtain that

- either there exists a \hat{r} such that $r_2 \xrightarrow{\tau} \hat{r}$ and $r' \xrightarrow{\bar{a}} \hat{r}$,
- or $r_2 \xrightarrow{a} r'$.

In the first case we apply the inductive hypothesis with $p' = p_1$, $r' = \hat{r}$, $p_0 = p_2$ and $r_0 = r'$ and obtain $p_2 \text{ MUST}_i r'$ as required. In the second case, the transitions $p_1 \xrightarrow{\bar{a}} p_2$ and $r_2 \xrightarrow{a} r'$ and rule [S-COM] let us prove $p_1 \parallel r_2 \xrightarrow{\tau} p_2 \parallel r'$. We apply part (2) to obtain $p_2 \text{ MUST}_i r'$ as required.

We now consider the case (c) in which $p_2 \xrightarrow{\mu} p'$ and $r_1 \xrightarrow{\bar{\mu}} r'$. We must show $p' \text{ MUST}_i r'$ and to do so we distinguish whether $\mu = a$ or not.

If $\mu = a$ then we apply the OUTPUT-DETERMINACY axiom to the transitions $r_1 \xrightarrow{\bar{a}} r_2$ and $r_1 \xrightarrow{\bar{\mu}} r'$ to obtain that $r_2 = r'$. Since by hypothesis $p_1, p_2 \in A$ and $\mathcal{L}_A \in \text{FWD}$ we apply the FWD-FEEDBACK axiom to the transitions $p_1 \xrightarrow{\bar{a}} p_2 \xrightarrow{a} p'$ to prove that either $p_1 \xrightarrow{\tau} p'$ or $p_1 = p'$ must hold. If $p_1 \xrightarrow{\tau} p'$ then we have that $p_1 \parallel r_2 \xrightarrow{\tau} p' \parallel r_2$. The property in (2) ensures that $p' \text{ MUST}_i r_2$ and from $r_2 = r'$ we have that the required $p' \text{ MUST}_i r'$ holds too. If $p_1 = p'$ then $p' \text{ MUST}_i r_2$ is a direct consequence of the hypothesis $p_1 \text{ MUST}_i r_2$.

If $\mu \neq a$ then we are allowed to apply the OUTPUT-CONFLUENCE axiom to the transitions $r_1 \xrightarrow{\bar{a}} r_2$ and $r_1 \xrightarrow{\bar{\mu}} r'$ to obtain a \hat{r} such that $r_2 \xrightarrow{\bar{\mu}} \hat{r}$ and $r' \xrightarrow{\bar{a}} \hat{r}$. An application of the OUTPUT-COMMUTATIVITY axiom to the transitions $p_1 \xrightarrow{\bar{a}} p_2 \xrightarrow{\mu} p'$ provides us with a \hat{p} such that $p_1 \xrightarrow{\mu} \hat{p} \xrightarrow{\bar{a}} p'$. We now apply the inductive hypothesis with $p' = \hat{p}$, $r' = \hat{r}$, $p_0 = p'$ and $r_0 = r'$ and obtain $p_2 \text{ MUST}_i r'$ as required. This concludes the proof of (ii), and therefore of the lemma.

Lemma 15. *Let $\mathcal{L}_A \in \text{FWD}$ and $\mathcal{L}_B \in \text{FDB}$. For every $p, p' \in A$, $r, r' \in B$ and every action $\mu \in \text{Act}$ such that $p \xrightarrow{\mu} p'$ and $r \xrightarrow{\bar{\mu}} r'$ we have that $p \text{MUST}_i r$ and $\neg \text{GOOD}(r)$ implies $p' \text{MUST}_i r'$.*

Proof. By case analysis on the hypothesis that $p \text{MUST}_i r$.

C Forwarders

The intuition behind forwarders, quoting [64], is that “any message can come into the configuration, regardless of the forms of inner receptors. [...] As the experimenter is not synchronously interacting with the configuration [...], he may send any message as he likes.”

In this appendix we give the technical results to ensure that the function FW builds an LTS that satisfies the axioms of the class **LtsEq**.

For any LTS $\mathcal{L} = \langle A, L, \longrightarrow \rangle$ of output-buffered agents we assume a function $\text{mbox} : A \rightarrow MO$

- (i) $\bar{a} \in O(p)$ if and only if $\bar{a} \in \text{mbox}(p)$, and
- (ii) for every p' , if $p \xrightarrow{\bar{a}} p'$ then $\text{mbox}(p) = \{\bar{a}\} \uplus \text{mbox}(p')$.

Note that by definition $\text{mbox}(p)$ is a finite multiset.

Definition 14. *For any LTS $\mathcal{L} = \langle C, L, \longrightarrow \rangle \in \text{FDB}$, we define the function $\text{strip} : C \rightarrow C$ by induction on $\text{mbox}(p)$ as follows: if $\text{mbox}(p) = \emptyset$ then $\text{strip}(p) = p$, while if $\exists \bar{a} \in \text{mbox}(p)$ and $p \xrightarrow{\bar{a}} p'$ then $\text{strip}(p) = \text{strip}(p')$. Note that $\text{strip}(p)$ is well-defined thanks to the **OUTPUT-DETERMINACY** and the **OUTPUT-COMMUTATIVITY** axioms.*

We now wish to show that $\text{FW}(\mathcal{L}) \in \text{FWD}$ for any LTS \mathcal{L} of output-buffered agents with feedback. Owing to the structure of our typeclasses, we have first to construct an equivalence \doteq over $\text{FW}(\mathcal{L})$ that is compatible with the transition relation, *i.e.* satisfies the axiom in Figure 5. We do this in the obvious manner, *i.e.* by combining the equivalence \simeq over the states of \mathcal{L} with an equivalence over mailboxes.

Definition 15. *For any LTS $\mathcal{L} \in \text{FDB}$, two states $p \triangleright M$ and $q \triangleright N$ of $\text{FW}(\mathcal{L})$ are equivalent, denoted $p \triangleright M \doteq q \triangleright N$, if $\text{strip}(p) \simeq \text{strip}(q)$ and $M \uplus \text{mbox}(p) = N \uplus \text{mbox}(q)$.*

Lemma 16. *For every \mathcal{L}_A and every $p \triangleright M, q \triangleright N \in A \times MO$, and every $\alpha \in L$, if $p \triangleright M (\doteq \cdot \xrightarrow{\alpha}_{\text{fw}}) q \triangleright N$ then $p \triangleright M (\xrightarrow{\alpha}_{\text{fw}} \cdot \doteq) q' \triangleright N'$.*

Lemma 17. *For every $\mathcal{L}_A \in \text{FDB}$ and every $p, q \in A$, $M \in MO$, if $p \simeq q$ then $p \triangleright M \doteq q \triangleright M$.*

Proof. This follows from the fact that if $p \simeq q$ then $\text{strip}(p) \simeq \text{strip}(q)$ and $\text{mbox}(p) = \text{mbox}(q)$.

Lemma 2. *For every LTS $\mathcal{L} \in \text{FDB}$, $\text{FW}(\mathcal{L}) \in \text{FWD}$.*

Proof. We must show that, given an LTS $\mathcal{L} = \langle C, L, \longrightarrow \rangle \in \text{FDB}$, we have that $\text{FW}(\mathcal{L}) \in \text{FWD}$. To do so, we need to show that $\text{FW}(\mathcal{L})$ obeys to the axioms given in Equation (7), namely INPUT-BOOMERANG and FWD-FEEDBACK. We first show that $\text{FW}(\mathcal{L})$ obeys to the INPUT-BOOMERANG axiom.

We pick a process $p \in C$, a mailbox $M \in \text{MO}$ and a name $a \in \mathcal{N}$. The axiom INPUT-BOOMERANG requires us to exhibit a process $p' \in A$ and a mailbox $M' \in \text{MO}$ such that the following transitions hold.

$$p \triangleright M \xrightarrow{a}_{\text{fw}} p' \triangleright M' \xrightarrow{\bar{a}}_{\text{fw}} p \triangleright M$$

We choose $p' = p$ and $M' = (\{\bar{a}\} \uplus M)$. An application of the rule [L-MINP] and then the rule [L-MOUT] from Figure 6 allows us to derive the required sequence of transitions as shown below.

$$p \triangleright M \xrightarrow{a}_{\text{fw}} p \triangleright (\{\bar{a}\} \uplus M) \xrightarrow{\bar{a}}_{\text{fw}} p \triangleright M$$

We now show that $\text{FW}(\mathcal{L})$ obeys to the FWD-FEEDBACK axiom. To begin we pick three processes $p_1, p_2, p_3 \in C$, three mailboxes $M_1, M_2, M_3 \in \text{MO}$ and a name $a \in \mathcal{N}$ such that:

$$p_1 \triangleright M_1 \xrightarrow{\bar{a}}_{\text{fw}} p_2 \triangleright M_2 \xrightarrow{a}_{\text{fw}} p_3 \triangleright M_3$$

We need to show that either:

1. $p_1 \triangleright M_1 \xrightarrow{\tau}_{\text{fw}} p_3 \triangleright M_3$, or
2. $p_1 \triangleright M_1 \doteq p_3 \triangleright M_3$

We proceed by case analysis on the last rule used to derive the transition $p_1 \triangleright M_1 \xrightarrow{\bar{a}}_{\text{fw}} p_2 \triangleright M_2$. This transition can either be derived by the rule [L-MOUT] or the rule [L-PROC].

We first consider the case where the transition has been derived using the rule [L-MOUT]. We then have that $p_1 = p_2$ and $M_1 = (\{\bar{a}\} \uplus M_2)$. We continue by case analysis on the last rule used to derive the transition $p_2 \triangleright M_2 \xrightarrow{a}_{\text{fw}} p_3 \triangleright M_3$. If this transition has been derived using the rule [L-MINP] then it must be the case that $p_2 = p_3$ and that $(\{\bar{a}\} \uplus M_2) = M_3$. This lets us conclude by the following equality to show that $p_1 \triangleright M_1 \doteq p_3 \triangleright M_3$.

$$p_1 \triangleright M_1 = p_2 \triangleright (\{\bar{a}\} \uplus M_2) = p_3 \triangleright M_3$$

Otherwise, this transition has been derived using the rule [L-PROC], which implies that $p_2 \xrightarrow{a} p_3$ together with $M_2 = M_3$. An application of the rule [L-COMM] ensures the following transition and allows us to conclude this case with $p_1 \triangleright M_1 \xrightarrow{\tau}_{\text{fw}} p_3 \triangleright M_3$.

$$p_1 \triangleright M_1 = p_2 \triangleright (\{\bar{a}\} \uplus M_2) \xrightarrow{\tau} p_3 \triangleright M_2 = p_3 \triangleright M_3$$

We now consider the case where the transition $p_1 \triangleright M_1 \xrightarrow{\bar{a}}_{\text{fw}} p_2 \triangleright M_2$ has been derived using the rule [L-PROC] such that $p_1 \xrightarrow{\bar{a}} p_2$ and $M_1 = M_2$.

Again, we continue by case analysis on the last rule used to derive the transition $p_2 \triangleright M_2 \xrightarrow{a}_{\text{fw}} p_3 \triangleright M_3$. If this transition has been derived using the rule [L-MINP] then it must be the case that $p_2 = p_3$ and $(\{\bar{a}\} \uplus M_2) = M_3$. Also, note that, as $\mathcal{L} \in \text{FDB}$, the transition $p_1 \xrightarrow{\bar{a}} p_2$ implies $\text{mbox}(p_1) = \text{mbox}(p_2) \uplus \{\bar{a}\}$. In order to prove $p_1 \triangleright M_1 \doteq p_3 \triangleright M_3$, it suffices to show the following:

- (a) $\text{strip}(p_1) \simeq \text{strip}(p_3)$, and
- (b) $M_1 \uplus \text{mbox}(p_1) = M_3 \uplus \text{mbox}(p_3)$

We show that $\text{strip}(p_1) \simeq \text{strip}(p_3)$ by definition of $\text{strip}()$ together with the transition $p_1 \xrightarrow{\bar{a}} p_2$ and the equality $p_2 = p_3$.

The following ensures that $M_1 \uplus \text{mbox}(p_1) = M_3 \uplus \text{mbox}(p_3)$.

$$\begin{aligned}
 & M_1 \uplus \text{mbox}(p_1) \\
 = & M_1 \uplus \text{mbox}(p_2) \uplus \{\bar{a}\} && \text{from } \text{mbox}(p_1) = \text{mbox}(p_2) \uplus \{\bar{a}\} \\
 = & M_2 \uplus \text{mbox}(p_3) \uplus \{\bar{a}\} && \text{from } M_1 = M_2, p_2 = p_3 \\
 = & (M_2 \uplus \{\bar{a}\}) \uplus \text{mbox}(p_3) \\
 = & M_3 \uplus \text{mbox}(p_3) && \text{from } M_3 = M_2 \uplus \{\bar{a}\}
 \end{aligned}$$

If the transition $p_2 \triangleright M_2 \xrightarrow{a}_{\text{fw}} p_3 \triangleright M_3$ has been derived using the rule [L-PROC] then it must be the case that $p_2 \xrightarrow{a} p_3$ and $M_2 = M_3$. As $\mathcal{L} \in \text{FDB}$, we are able to call the axiom FEEDBACK together with the transitions $p_1 \xrightarrow{\bar{a}} p_2$ and $p_2 \xrightarrow{a} p_3$ to obtain a process p'_3 such that $p_1 \xrightarrow{\tau} p'_3$ and $p'_3 \simeq p_3$. An application of Lemma 17 and rule [L-PROC] allows us to conclude that $p_1 \triangleright M_1 \xrightarrow{\tau}_{\text{fw}} \cdot \doteq p_3 \triangleright M_3$. \square

D Completeness

This section is devoted to the proof that the alternative preorder given in Definition 5 includes the MUST-preorder. First we present a general outline of the main technical results to obtain the proof we are after. Afterwards, in Subsection (D) we discuss in detail on all the technicalities.

Proofs of completeness of characterisations of contextual preorders usually require using, as the name suggests, syntactic contexts. Our calculus-independent setting, though, does not allow us to define them. Instead we phrase our arguments using two functions $tc : \text{Act}^* \rightarrow C$, and $ta : \text{Act}^* \times \mathcal{P}(\mathcal{N}) \rightarrow C$ where $\langle C, L, \longrightarrow \rangle$ is some LTS of FDB. In Table 1 we gather all the *properties* of tc and ta that are sufficient to give our arguments. The properties (1) - (5) must hold for both tc and $ta(\varepsilon, -)$ for every set of names O , the properties (c1) - (c2) must hold for tc , and (t1) - (t4) must hold for ta .

We use the function tc to test the convergence of servers, and the function ta to test the acceptance sets of servers.

A natural question is whether such tc and ta can actually exist. The answer depends on the LTS at hand. In Appendix H.2, and in particular Figure 16, we

$\forall s \in \text{Act}^*, \forall a \in \mathcal{N}$,

- (1) $\neg \text{GOOD}(f(s))$
- (2) $\forall \mu \in \text{Act}, f(\mu.s) \xrightarrow{\bar{\mu}} f(s)$
- (3) $f(\bar{a}.s) \xrightarrow{\tau}$
- (4) $\forall r \in C, f(\bar{a}.s) \xrightarrow{\tau} r$ implies $\text{GOOD}(r)$
- (5) $\forall r \in C, \mu \in \text{Act}, f(\bar{a}.s) \xrightarrow{\mu} r$ implies $\mu = a$ and $r = f(s)$

$\forall E \subseteq \mathcal{N}$,

- (t1) $ta(\varepsilon, E) \xrightarrow{\tau}$
- (t2) $\forall a \in \mathcal{N}, ta(\varepsilon, E) \xrightarrow{\bar{a}}$
- (t3) $\forall a \in \mathcal{N}, ta(\varepsilon, E) \xrightarrow{a}$ if and only if $a \in E$
- (t4) $\forall \mu \in \text{Act}, r \in C, ta(\varepsilon, E) \xrightarrow{\mu} r$ implies $\text{GOOD}(r)$

- (c1) $\forall \mu \in \text{Act}, tc(\varepsilon) \xrightarrow{\mu}$
- (c2) $\exists r, tc(\varepsilon) \xrightarrow{\tau} r$
- (c3) $\forall r, tc(\varepsilon) \xrightarrow{\tau} r$ implies $\text{GOOD}(r)$

Table 1. Properties of the functions that generate clients.

define these functions for the standard LTS of ACCS, and it should be obvious how to adapt those definitions to the asynchronous π -calculus [57].

In short, our proofs show that \preceq_{AS} is complete with respect to $\sqsubseteq_{\text{MUST}}$ in any LTS of output-buffered agents with feedback wherein the functions tc and ta enjoying the properties in Table 1 can be defined.

First, converging along a finite trace s is logically equivalent to passing the client $tc(s)$. In other words, there exists a bijection between the proofs (i.e. finite derivation trees of $p \text{MUST}_i tc(s)$) and the ones of $p \Downarrow s$. We first give the proposition, and then discuss the auxiliary lemmas to prove it.

Proposition 3. *For every $\mathcal{L}_A \in \text{FWD}$, $p \in A$, and $s \in \text{Act}^*$ we have that $p \text{MUST}_i tc(s)$ if and only if $p \Downarrow s$.*

The *if* implication is Lemma 31 and the *only if* implication is Lemma 28. The hypothesis that $\mathcal{L}_A \in \text{FWD}$, i.e. the use of forwarders, is necessary to show that convergence implies passing a client, as shown by the next example.

Example 8. Consider a server p in an LTS $\mathcal{L} \in \text{FDB}$ whose behaviour amounts to the following transitions: $p \xrightarrow{b} \Omega \xrightarrow{\tau} \Omega \xrightarrow{\tau} \dots$. Note that this entails that \mathcal{L} does not *not* enjoy the axioms of forwarders.

Now let $s = a.b$. Since $p \Downarrow$ and $p \xrightarrow{a}$ we know that $p \Downarrow a.b$. On the other hand Table 1(2) implies that the client $tc(s)$ performs the transitions $tc(s) \xrightarrow{\bar{a}} tc(b) \xrightarrow{\bar{b}} tc(\varepsilon)$. Thanks to the OUTPUT-COMMUTATIVITY axiom we

D. COMPLETENESS

obtain $tc(s) \xrightarrow{\bar{b}} \xrightarrow{\bar{a}} tc(\varepsilon)$. Table 1(1) implies that the states reached by the client are unsuccessful, and so by zipping the traces performed by p and by $tc(s)$ we build a maximal computation of $p \parallel tc(s)$ that is unsuccessful, and thus $p \text{ MUST}_i tc(s)$. \square

This example explains why in spite of Lemma 3 output-buffered agents with feedback do not suffice to use the standard characterisations of the MUST-preorder.

We move on to the more involved technical results, *i.e.* the next two lemmas, that we use to reason on acceptance sets of servers.

Lemma 18. *Let $\mathcal{L}_A \in \text{FWD}$. For every $p \in A$, $s \in \text{Act}^*$, and every $L, E \subseteq \mathcal{N}$, if $\bar{L} \in \mathcal{A}_{\text{fw}}(p, s)$ then $p \text{ MUST}_i ta(s, E \setminus L)$.*

Lemma 19. *Let $\mathcal{L}_A \in \text{FWD}$. For every $p \in A$, $s \in \text{Act}^*$, and every finite set $O \subseteq \mathcal{N}$, if $p \Downarrow s$ then either*

- (i) $p \text{ MUST}_i ta(s, \bigcup \overline{\mathcal{A}_{\text{fw}}(p, s)} \setminus O)$, or
- (ii) there exists $\hat{O} \in \mathcal{A}_{\text{fw}}(p, s)$ such that $\hat{O} \subseteq O$.

We can now show that the alternative preorder \preceq_{AS} includes \preceq_{MUST} when used over LTSs of forwarders.

Lemma 20. *For every $\mathcal{L}_A, \mathcal{L}_B \in \text{FWD}$ and servers $p \in A, q \in B$, if $p \preceq_{\text{MUST}} q$ then $p \preceq_{\text{AS}} q$.*

Proof. Let $p \preceq_{\text{MUST}} q$. To prove that $p \preceq_{\text{cnv}} q$, suppose $p \Downarrow s$ for some trace s . Proposition 3 implies $p \text{ MUST}_i tc(s)$, and so by hypothesis $q \text{ MUST}_i tc(s)$. Proposition 3 ensures that $q \Downarrow s$.

We now show that $p \preceq_{\text{acc}} q$. Thanks to Lemma 4, it is enough to prove that $p \preceq_{\text{acc}}^{\text{fw}} q$. So, we show that for every trace $s \in \text{Act}^*$, if $p \Downarrow s$ then $\mathcal{A}_{\text{fw}}(p, s) \ll \mathcal{A}_{\text{fw}}(q, s)$. Fix an $O \in \mathcal{A}_{\text{fw}}(q, s)$. We have to exhibit a set $\hat{O} \in \mathcal{A}_{\text{fw}}(p, s)$ such that $\hat{O} \subseteq O$.

By definition $O \in \mathcal{A}_{\text{fw}}(q, s)$ means that for some q' we have $q \xrightarrow{s} q' \xrightarrow{\tau}$ and $O(q') = O$. Let $E = \bigcup \mathcal{A}_{\text{fw}}(p, s)$ and $X = E \setminus O$. The hypothesis that $p \Downarrow s$, and the construction of the set X let us apply Lemma 19, which implies that either

- (a) $p \text{ MUST}_i ta(s, \bar{X})$, or
- (b) there exists a $\hat{O} \in \mathcal{A}_{\text{fw}}(p, s)$ such that $\hat{O} \subseteq O(q')$.

Since (b) is exactly what we are after, to conclude the argument it suffices to prove that (a) is false. This follows from Lemma 18, which proves $q \text{ MUST}_i ta(s, \bar{X})$, and the hypothesis $p \preceq_{\text{MUST}} q$, which ensures $p \text{ MUST}_i ta(s, \bar{X})$.

The fact that the MUST-preorder can be captured via the function $\text{FW}(-)$ and \preceq_{AS} is a direct consequence of Lemma 3 and Lemma 20.

Proposition 4 (Completeness). *For every $\mathcal{L}_A, \mathcal{L}_B \in \text{FDB}$ and servers $p \in A, q \in B$, if $p \preceq_{\text{MUST}} q$ then $\text{FW}(p) \preceq_{\text{AS}} \text{FW}(q)$.*

We now gather all the technical auxiliary lemmas and then discuss the proofs of the main ones.

By assumption, outputs preserve the predicate GOOD. For stable clients, they also preserve the negation of this predicate.

Lemma 21. *For all $r, r' \in A$ and trace $s \in \overline{\mathcal{N}}^*$, $r \xrightarrow{\tau} \neg\text{GOOD}(r)$ and $r \xRightarrow{s} r'$ implies $\neg\text{GOOD}(r')$.*

D.1 Testing convergence

We start with preliminary facts, in particular two lemmas that follow from the properties in Table 1.

A process p *converges along* a trace s if for every p' reached by p performing any prefix of s , the process p' converges.

Lemma 22. *For every $\langle A, L, \longrightarrow \rangle$, $p \in A$, and $s \in \text{Act}^*$, $p \Downarrow s$ if and only if $p \xRightarrow{s'} p'$ implies $p' \Downarrow$ for every s' prefix of s .*

Traces of output actions impact neither the stability of servers, nor their input actions.

Lemma 23. *For every \mathcal{L}_A , every $p, p' \in A$ and every trace $s \in \overline{\mathcal{N}}^*$,*

1. $p \xrightarrow{\tau} \neg\text{GOOD}(p)$ and $p \xRightarrow{s} p'$ implies $p' \xrightarrow{\tau} \neg\text{GOOD}(p')$.
2. $p \xrightarrow{\tau} \neg\text{GOOD}(p)$ and $p \xRightarrow{s} p'$ implies $I(p) = I(p')$.

Lemma 24. *For every $s \in \text{Act}^*$, $tc(s) \xrightarrow{\tau} \neg\text{GOOD}(tc(s))$.*

The BACKWARD-OUTPUT-DETERMINACY axiom is used in the proof of the next lemma.

Lemma 25. *For every $s \in \text{Act}^*$, if $tc(s) \xrightarrow{\mu} r$ then either*

- (a) $\text{GOOD}(r)$, or
- (b) $s = s_1.\bar{\mu}.s_2$ for some $s_1 \in \mathcal{N}^*$ and $s_2 \in \text{Act}^*$ such that $r \simeq tc(s_1.s_2)$.

Lemma 26. *For every $s \in \text{Act}^*$, if $tc(s) \xrightarrow{\tau} r$ then either:*

- (a) $\text{GOOD}(r)$, or
- (b) there exist b, s_1, s_2 and s_3 with $s_1.b.s_2 \in \mathcal{N}^*$ such that $s = s_1.b.s_2.\bar{b}.s_3$ and $r \simeq tc(s_1.s_2.s_3)$.

Lemma 27. *Let $\mathcal{L}_A \in \text{FWD}$. For every server $p, p' \in A$, trace $s \in \text{Act}^*$ and action $\mu \in \text{Act}$ such that $p \xRightarrow{\mu} p'$ we have that $p \text{MUST}_i tc(\mu.s)$ implies $p' \text{MUST}_i tc(s)$.*

Proof. By rule induction on the reduction $p \xRightarrow{\mu} p'$ together with Lemma 12 and Lemma 15.

Lemma 28. *Let $\mathcal{L}_A \in \text{FWD}$. For every server $p \in A$, trace $s \in \text{Act}^*$ we have that $p \text{MUST}_i tc(s)$ implies $p \Downarrow s$.*

D. COMPLETENESS

Proof. We proceed by induction on the trace s . In the base case s is ε . Table 1(1) states that $\neg\text{GOOD}(tc(\varepsilon))$ and we apply Lemma 11 to obtain $p \downarrow_i$, and thus $p \Downarrow \varepsilon$. In the inductive case s is $\mu.s'$ for some $\mu \in \text{Act}$ and $s' \in \text{Act}^*$. We must show the following properties,

1. $p \downarrow_i$, and
2. for every p' such that $p \xrightarrow{\mu} p'$, $p' \Downarrow s'$.

We prove the first property as we did in the base case, and we apply Lemma 27 to prove the second property.

Lemma 29. *Let $\mathcal{L}_A \in \text{FWD}$. For every $p \in A$, $s_1 \in \mathcal{N}^*$ and $s_3 \in \text{Act}^*$ we have that*

1. for every $\mu \in \text{Act}$, if $p \Downarrow s_1.\mu.s_3$ and $p \xrightarrow{\mu} q$ then $q \Downarrow s_1.s_3$,
2. for every $a.s_2 \in \mathcal{N}^*$ if $p \Downarrow s_1.a.s_2.\bar{a}.s_3$ then $p \Downarrow s_1.s_2.s_3$.

Lemma 30. *For every LTS \mathcal{L}_A and every $p \in A$, $p \downarrow_i$ implies $p \text{MUST}_i tc(\varepsilon)$.*

Proof. Rule induction on the derivation of $p \downarrow_i$.

Lemma 31. *For every $\mathcal{L}_A \in \text{FWD}$, every $p \in A$, and $s \in \text{Act}^*$, if $p \Downarrow s$ then $p \text{MUST}_i tc(s)$.*

Proof. The hypothesis $p \Downarrow s$ ensures $p \downarrow_i$. We show that $p \text{MUST}_i tc(s)$ reasoning by complete induction on the length of the trace s . The base case is Lemma 30 and here we discuss the inductive case, i.e. when $\text{len}(s) = n + 1$ for some $n \in \mathbb{N}$.

We proceed by rule induction on $p \downarrow_i$. In the base case $p \xrightarrow{\tau}$, and the reduction at hand is due to either a τ transition in $tc(s)$, or a communication between p and $tc(s)$.

In the first case $tc(s) \xrightarrow{\tau} r$, and so Lemma 26 ensures that one of the following conditions holds,

1. $\text{GOOD}(r)$, or
2. there exist $a \in \mathcal{N}$, s_1, s_2 and s_3 with $s = s_1.a.s_2.\bar{a}.s_3$ and $r \simeq tc(s_1.s_2.s_3)$.

If $\text{GOOD}(r)$ then we conclude via rule [AXIOM]; otherwise Lemma 29(2) and the hypothesis that $p \Downarrow s$ imply $p \Downarrow s_1.s_2.s_3$, thus prove $p \text{MUST}_i r$ via the inductive hypothesis of the complete induction on s .

We now consider the case when the transition is due to a communication, i.e. $p \xrightarrow{\mu} p'$ and $tc(s) \xrightarrow{\bar{\mu}} r$. Lemma 25 tells us that either $\text{GOOD}(r)$ or there exist s_1 and s_2 such that $s = s_1.\mu.s_2$ and $r \simeq tc(s_1.s_2)$. In the first case we conclude via rule [AXIOM]. In the second case we apply Lemma 29(1) to prove $p' \Downarrow s_1.s_2$, and thus $p' \text{MUST}_i r$ follows from the inductive hypothesis of the complete induction. In the inductive case of the rule induction on $p \downarrow_i$, we know that $p \xrightarrow{\tau} p'$ for some process p' . We reason again by case analysis on how the reduction we fixed has been derived, i.e. either via a τ transition in $tc(s)$, or via a communication between p and $tc(s)$, or via a τ transition in p . In the first two cases we reason as we did for the base case of the rule induction. In the third case $p \Downarrow s$ and $p \xrightarrow{\tau} p'$ imply $p' \Downarrow s$, we thus obtain $p' \text{MUST}_i tc(s)$ thanks to the inductive hypothesis of the rule induction which we can apply because the tree to derive $p' \downarrow_i$ is smaller than the tree to derive that $p \downarrow_i$.

D.2 Testing acceptance sets

In this section we present the properties of the function $ta(-, -)$ that are sufficient to obtain completeness. To begin with, $ta(-, -)$ function enjoys a form of monotonicity with respect to its second argument.

Lemma 32. *Let $\mathcal{L}_A \in \text{FDB}$. For every $p \in A$, trace $s \in \text{Act}^*$, and sets of outputs O_1, O_2 , if $p \text{ MUST}_i ta(s, O_1)$ and $O_1 \subseteq O_2$ then $p \text{ MUST}_i ta(s, O_2)$.*

Proof. Induction on the derivation of $p \text{ MUST}_i ta(s, O_1)$.

Let OBA denote the set of LTS of output-buffered agents. Note that any $\mathcal{L} \in \text{OBA}$ need not enjoy the FEEDBACK axiom.

Lemma 33. *Let $\mathcal{L}_A \in \text{OBA}$, and $\mathcal{L}_B \in \text{OBA}$. For every $p \in A$, trace $s \in \text{Act}^*$, set of outputs O and name $a \in \mathcal{N}$, such that*

- (i) $p \downarrow_i$ and,
- (ii) For every $p' \in A$, $p \xrightarrow{\bar{a}} p'$ implies $p' \text{ MUST}_i ta(s, \bar{O})$,

we have that $p \text{ MUST}_i ta(\bar{a}.s, \bar{O})$.

Proof. We proceed by induction on the hypothesis $p \downarrow_i$.

(Base case: p is stable) We prove $p \text{ MUST}_i ta(\bar{a}.s, \bar{O})$ by applying rule [IND-RULE]. Since Table 1(3) implies that $p \parallel ta(\bar{a}.s, \bar{O}) \xrightarrow{\tau} p' \parallel r$, all we need to prove is the following fact,

$$\forall p' \in A, r \in B. \text{ if } p \parallel ta(\bar{a}.s, \bar{O}) \xrightarrow{\tau} p' \parallel r \text{ then } p' \text{ MUST}_i r. \quad (\star)$$

Fix a transition $p \parallel ta(\bar{a}.s, \bar{O}) \xrightarrow{\tau} p' \parallel r$. As p is stable, this transition can either be due to:

1. a τ -transition performed by the client such that $ta(\bar{a}.s, \bar{O}) \xrightarrow{\tau} r$, or
2. an interaction between the server p and the client $ta(\bar{a}.s, \bar{O})$.

In the first case Table 1(4) implies $\text{GOOD}(r)$, and hence we obtain $p' \text{ MUST}_i r$ via rule [AXIOM]. In the second case there exists an action μ such that

$$p \xrightarrow{\mu} p' \quad \text{and} \quad ta(\bar{a}.s, \bar{O}) \xrightarrow{\bar{\mu}} r$$

Table 1(5) implies μ is \bar{a} and $r = ta(s, \bar{O})$. We then have $p \xrightarrow{\bar{a}} p'$ and thus the reduction $p \xrightarrow{\bar{a}} p'$, which allows us to apply the hypothesis (ii) and obtain $p' \text{ MUST}_i r$ as required.

D. COMPLETENESS

(*Inductive case: $p \xrightarrow{\tau} p'$ implies p'*) The argument is similar to one for the base case, except that we must also tackle the case when the transition $p \parallel ta(\bar{a}.s, \bar{O}) \xrightarrow{\tau} p' \parallel r$ is due to a τ action performed by p , i.e. $p \xrightarrow{\tau} p'$ and $r = ta(\bar{a}.s, \bar{O})$. The inductive hypothesis tells us the following fact:

For every p_1 and a , such that $p \xrightarrow{\tau} p_1$, for every p_2 , if $p_1 \xrightarrow{\bar{a}} p_2$ then
 $p_2 \text{ MUST}_i ta(s, O)$.

To apply the inductive hypothesis we have to show that for every p_2 such that $p' \xrightarrow{\bar{a}} p_2$ we have that $p_2 \text{ MUST}_i ta(s, \bar{O})$. This is a consequence of the hypothesis (ii) together with the reduction $p \xrightarrow{\tau} p' \xrightarrow{\bar{a}} p_2$, and thus concludes the proof.

Lemma 34. *Let $\mathcal{L}_A \in \text{FDB}$. For every $p \in A$ and set of outputs O , if p is stable then either*

- (a) $p \text{ MUST}_i ta(\varepsilon, \overline{O(p) \setminus O})$, or
- (b) $O(p) \subseteq O$.

Proof. We distinguish whether $O(p) \setminus O$ is empty or not. In the first case, $O(p) \setminus O = \emptyset$ implies $O(p) \subseteq O$, and we are done.

In the second case, there exists $\bar{a} \in O(p)$ such that $\bar{a} \notin O$. Note also that Table 1(1) ensures that $\neg \text{GOOD}(ta(\varepsilon, \overline{O(p) \setminus O}))$, and thus we construct a derivation of $p \text{ MUST}_i ta(\varepsilon, \overline{O(p) \setminus O})$ by applying the rule [IND-RULE]. This requires us to show the following facts,

1. $p \parallel ta(s, \overline{O(p) \setminus O}) \longrightarrow$, and
2. for each p', r such that $p \parallel ta(s, \overline{O(p) \setminus O}) \xrightarrow{\tau} p' \parallel r$, $p' \text{ MUST}_i r$ holds.

To prove (1), we show that an interaction between the server p and the test $ta(s, \overline{O(p) \setminus O})$ exists. As $\bar{a} \in O(p)$, we have that $p \xrightarrow{\bar{a}}$. Then $\bar{a} \in O(p) \setminus O$ together with (3) ensure that $ta(s, \overline{O(p) \setminus O}) \xrightarrow{\bar{a}}$. An application of the rule [S-COM] gives us the required transition $p \parallel ta(s, \overline{O(p) \setminus O}) \longrightarrow$.

To show (2), fix a silent transition $p \parallel ta(s, \overline{O(p) \setminus O}) \xrightarrow{\tau} p' \parallel r$. We proceed by case analysis on the rule used to derive the transition under scrutiny. Recall that the server p is stable by hypothesis, and that $ta(s, \overline{O(p) \setminus O})$ is stable thanks to Table 1(1). This means that the silent transition must have been derived via rule [S-COM]. Furthermore, Table 1(2) implies that the test $ta(s, \overline{O(p) \setminus O})$ does not perform any output. As a consequence, if there is an interaction it must be because the test performs an input and becomes r . Table 1(4) implies that $\text{GOOD}(r)$, and hence we obtain the required $p' \text{ MUST}_i r$ applying rule [AXIOM].

Lemma 18 Let $\mathcal{L}_A \in \text{FWD}$. For every $p \in A$, $s \in \text{Act}^*$, and every $L, E \subseteq \mathcal{N}$, if $\bar{L} \in \mathcal{A}_{\text{fw}}(p, s)$ then $p \text{ MUST}_i ta(s, E \setminus L)$.

Proof. By hypothesis there exists a set $\bar{L} \in \mathcal{A}_{\text{fw}}(p, s)$, i.e. for some p' we have $p \xrightarrow{s} p' \xrightarrow{\tau} \bar{L}$ and $O(p') = \bar{L}$. We have to show that $p \text{ MUST}_i ta(s, E \setminus L)$, i.e. $p \text{ MUST}_i ta(s, E \setminus L)$ implies \perp . For convenience, let $X = E \setminus L$.

We proceed by induction on the derivation of the weak transitions $p \xRightarrow{s} p'$. In the base case the derivation consists in an application of rule [WT-REFL], which implies that $p = p'$ and $s = \varepsilon$. We show that there exists no derivation of judgement $p \text{ MUST}_i ta(s, X)$. By definition, $\neg \text{GOOD}(ta(s, X))$ and thus no tree that ends with [AXIOM] can have $p \text{ MUST}_i ta(s, X)$ as conclusion. The hypotheses ensure that p is stable, and $ta(\varepsilon, X)$ is stable by definition. The set of inputs of $ta(\varepsilon, X)$ is X , which prevents an interaction between p and $ta(s, X)$, i.e. an application of rule [S-COM]. This proves that $p \parallel ta(s, X)$ is stable, thus a side condition of [IND-RULE] is false, and the rule cannot be employed to prove $p \text{ MUST}_i ta(s, X)$.

In the inductive cases $p \xRightarrow{s} p'$ is derived using either:

- (i) rule [WT-TAU] such that $p \xrightarrow{\tau} \hat{p} \xRightarrow{s} p'$, or
- (ii) rule [WT-MU] such that $p \xrightarrow{\mu} \hat{p} \xRightarrow{t} p'$, with $s = \mu.t$.

In the first case, applying the inductive hypothesis requires us to show $\hat{p} \text{ MUST}_i ta(s, X)$, which is true since $p \text{ MUST}_i ta(s, X)$ is preserved by the τ -transitions performed by the server.

In the second case, applying the inductive hypothesis requires us to show $\hat{p} \text{ MUST}_i ta(t, X)$. Table 1(2) implies that $ta(\mu.t, X) \xrightarrow{\mu} ta(t, X)$. Then we derive via [S-COM] the transition $p \parallel ta(\mu.t, E) \xrightarrow{\tau} \hat{p} \parallel ta(t, E)$. Since $p \text{ MUST}_i ta(s, X)$ is preserved by the interactions occurring between the server and the client, which implies $\hat{p} \text{ MUST}_i ta(t, X)$ as required.

Lemma 19 Let $\mathcal{L}_A \in \text{FWD}$. For every $p \in A, s \in \text{Act}^*$, and every finite set $O \subseteq \overline{\mathcal{N}}$, if $p \Downarrow s$ then either

- (i) $p \text{ MUST}_i ta(s, \bigcup \overline{\mathcal{A}_{\text{fw}}(p, s)} \setminus O)$, or
- (ii) there exists $\hat{O} \in \mathcal{A}_{\text{fw}}(p, s)$ such that $\hat{O} \subseteq O$.

Proof. We proceed by induction on the trace s .

(Base case, $s = \varepsilon$) The hypothesis $p \Downarrow \varepsilon$ implies $p \Downarrow_i$ and we continue by induction on the derivation of $p \Downarrow_i$.¹² In the base case $p \Downarrow_i$ was proven using rule [AXIOM], and hence $p \xrightarrow{\tau}$. We apply Lemma 34 to obtain either:

- (i) $p \text{ MUST}_i ta(\varepsilon, \overline{O(p)} \setminus O)$, or
- (ii) $O(p) \subseteq O$.

In case (i) we are done. In case (ii), as p is stable we have $\{p' \mid p \xRightarrow{\varepsilon} p' \xrightarrow{\tau}\} = \{p\}$ and thus $\mathcal{A}_{\text{fw}}(p, \varepsilon) = \{O(p)\}$ and we conclude by letting $\hat{O} = O(p)$.

In the inductive case $p \Downarrow_i$ was proven using rule [IND-RULE]. We know that $p \xrightarrow{\tau}$, and the inductive hypothesis states that for any p' such that $p \xrightarrow{\tau} p'$, either:

- (a) $p' \text{ MUST}_i ta(\varepsilon, \overline{O(p')} \setminus O)$, or

¹² Recall that the definition of \Downarrow_i is in Equation (int-preds)

D. COMPLETENESS

(b) there exists $\widehat{O} \in \mathcal{A}_{\text{fw}}(p', s)$ such that $\widehat{O} \subseteq O$.

It follows that either

- (\forall) for each $p' \in \{p' \mid p \xrightarrow{\tau} p'\}$, $p' \text{ MUST}_i ta(\varepsilon, \bigcup \overline{\mathcal{A}_{\text{fw}}(p', s)} \setminus O)$, or
- (\exists) there exists a $p' \in \{p' \mid p \xrightarrow{\tau} p'\}$ and a $\widehat{O} \in \mathcal{A}_{\text{fw}}(p', \varepsilon)$ such that $\widehat{O} \subseteq O$,

We discuss the two cases. If (\exists) the argument is straightforward: we pick the existing p' such that $p \xrightarrow{\tau} p'$. The definition of $\mathcal{A}_{\text{fw}}(-, -)$ ensures that and show that $\mathcal{A}_{\text{fw}}(p', \varepsilon) \subseteq \mathcal{A}_{\text{fw}}(p, \varepsilon)$, and thus we conclude by choosing \widehat{O} .

Case (\forall) requires more work. We are going to show that $p \text{ MUST}_i ta(\varepsilon, \bigcup \overline{\mathcal{A}_{\text{fw}}(p, s)} \setminus O)$ holds. To do so we apply the rule [IND-RULE] and we need to show the following facts,

- (a) $p \Vdash ta(\varepsilon, \bigcup \overline{\mathcal{A}_{\text{fw}}(p, s)} \setminus O) \xrightarrow{\tau}$, and
- (b) for each $p' \Vdash r'$ such that $p \Vdash ta(\varepsilon, \bigcup \overline{\mathcal{A}_{\text{fw}}(p, s)} \setminus O) \xrightarrow{\tau} p' \Vdash r'$, we have $p' \text{ MUST}_i r'$.

The first requirement follows from the fact that p is not stable. To show the second requirement we proceed by case analysis on the transition $p \Vdash ta(\varepsilon, \bigcup \overline{\mathcal{A}_{\text{fw}}(p, s)} \setminus O) \xrightarrow{\tau} p' \Vdash r'$. As $ta(\varepsilon, \overline{O(p)} \setminus O)$ is stable by (1), it can either be due to:

1. a τ -transition performed by the server p such that $p \xrightarrow{\tau} p'$, or
2. an interaction between the server p and the client $ta(\varepsilon, \bigcup \overline{\mathcal{A}_{\text{fw}}(p, s)} \setminus O)$.

In the first case we apply the first part of the inductive hypothesis to prove that $p' \text{ MUST}_i ta(\varepsilon, \bigcup \overline{\mathcal{A}_{\text{fw}}(p', s)} \setminus O)$, and we conclude via Lemma 32 to get the required

$$p' \text{ MUST}_i ta(\varepsilon, \bigcup \overline{\mathcal{A}_{\text{fw}}(p, s)} \setminus O).$$

In the second case, there exists a $\mu \in \text{Act}$ such that

$$p \xrightarrow{\mu} p' \text{ and } ta(\varepsilon, \bigcup \overline{\mathcal{A}_{\text{fw}}(p, s)} \setminus O) \xrightarrow{\bar{\mu}} r$$

Thanks to Table 1(4) we apply rule [AXIOM] to prove that $p' \text{ MUST}_i r$ and we are done with the base case of the main induction on the trace s .

(*Inductive case, $s = \mu.s'$*) By induction on the set $\{p' \mid p \xrightarrow{\mu} p'\}$ and an application of the inductive hypothesis we know that either:

- (i) there exists $p' \in \{p' \mid p \xrightarrow{\mu} p'\}$ and $\widehat{O} \in \mathcal{A}_{\text{fw}}(p', s')$ such that $\widehat{O} \subseteq O$, or
- (ii) for each $p' \in \{p' \mid p \xrightarrow{\mu} p'\}$ we have that $p' \text{ MUST}_i ta(s', \bigcup \overline{\mathcal{A}_{\text{fw}}(p', s')})$.

In the first case, the inclusion $\mathcal{A}_{\text{fw}}(p', s') \subseteq \mathcal{A}_{\text{fw}}(p, \mu.s')$ and $\widehat{O} \in \mathcal{A}_{\text{fw}}(p', s')$ imply $\widehat{O} \in \mathcal{A}_{\text{fw}}(p, s)$ and we are done.

In the second case, we show $p \text{ MUST}_i ta(s, \bigcup \overline{\mathcal{A}_{\text{fw}}(p, s)})$ by case analysis on the action μ , which can be either an input or an output.

- If μ is an input, $\mu = a$ for some $a \in \mathcal{N}$. An application of the axiom of forwarders gives us a p' such that $p \xrightarrow{a} p' \xrightarrow{\bar{a}} p$. An application of Table 1(2) gives us the following transition,

$$ta(a.s', \bigcup \overline{\mathcal{A}_{\text{fw}}(p, a.s') \setminus O}) \xrightarrow{\bar{a}} ta(s', \bigcup \overline{\mathcal{A}_{\text{fw}}(p, a.s') \setminus O})$$

By an application of Lemma 14 it is enough to show

$$p' \text{ MUST}_i ta(s', \bigcup \overline{\mathcal{A}_{\text{fw}}(p, a.s') \setminus O})$$

to obtain the required $p \text{ MUST}_i ta(a.s', \bigcup \overline{\mathcal{A}_{\text{fw}}(p, a.s') \setminus O})$.

- If μ is an output, $\mu = \bar{a}$ for some $a \in \mathcal{N}$ and we must show that

$$p \text{ MUST}_i ta(\bar{a}.s', \bigcup \overline{\mathcal{A}_{\text{fw}}(p, \bar{a}.s') \setminus O}).$$

We apply Lemma 33 together with (ii) to obtain $p \text{ MUST}_i ta(\bar{a}.s', \bigcup \overline{\mathcal{A}_{\text{fw}}(p', s') \setminus O})$. Again, Lemma 32 together with the inclusion $\mathcal{A}_{\text{fw}}(p', s') \subseteq \mathcal{A}_{\text{fw}}(p, \bar{a}.s')$ ensures the required $p \text{ MUST}_i ta(s, \bigcup \overline{\mathcal{A}_{\text{fw}}(p, s) \setminus O})$.

E Soundness

In this section we prove the converse of Proposition 4, i.e. that \preceq_{AS} is included in $\overline{\preceq}_{\text{MUST}}$. We remark immediately that a naïve reasoning does not work. Fix two servers p and q such that $p \preceq_{\text{AS}} q$.

To lift the predicates \preceq_{cnv} and \preceq_{acc} to sets of servers, we let $\mathcal{A}_{\text{fw}}(X, s) = \{O \mid \exists p \in X. O \in \mathcal{A}_{\text{fw}}(p, s)\}$, and for every finite $X \in \mathcal{P}^+(A)$, we write $X \downarrow$ to mean $\forall p \in X. p \downarrow$, we write $X \downarrow s$ to mean $\forall p \in X. p \downarrow s$, and let

- $X \preceq_{\text{cnv}}^{\text{set}} q$ to mean $\forall s \in \text{Act}^*$, if $X \downarrow s$ then $q \downarrow s$,
- $X \preceq_{\text{acc}}^{\text{set}} q$ to mean $\forall s \in \text{Act}^*$, $X \downarrow s$ implies $\mathcal{A}_{\text{fw}}(X, s) \ll \mathcal{A}_{\text{fw}}(q, s)$,
- $X \preceq_{\text{AS}}^{\text{set}} q$ to mean $X \preceq_{\text{cnv}}^{\text{set}} q \wedge X \preceq_{\text{acc}}^{\text{set}} q$.

These definitions imply immediately the following equivalences, $\{p\} \preceq_{\text{acc}}^{\text{set}} q \iff p \preceq_{\text{cnv}} q$, $\{p\} \preceq_{\text{cnv}}^{\text{set}} q \iff p \preceq_{\text{acc}} q$ and thereby the following lemma.

Lemma 35. *For every LTS $\mathcal{L}_A, \mathcal{L}_B, p \in A, q \in B$, $p \preceq_{\text{AS}} q$ if and only if $\{p\} \preceq_{\text{AS}}^{\text{set}} q$.*

The preorder $\preceq_{\text{AS}}^{\text{set}}$ is preserved by τ -transitions on its right-hand side, and by visible transitions on both sides. We reason separately on the two auxiliary preorders $\preceq_{\text{cnv}}^{\text{set}}$ and $\preceq_{\text{acc}}^{\text{set}}$. We need one further notion.

Lemma 36. *Let $\mathcal{L}_A, \mathcal{L}_B \in \text{FWD}$. For every set $X \in \mathcal{P}^+(A)$, and $q \in B$, such that $X \preceq_{\text{cnv}}^{\text{set}} q$,*

1. $q \xrightarrow{\tau} q'$ implies $X \preceq_{\text{cnv}}^{\text{set}} q'$,

2. $X \downarrow_i$, $X \xrightarrow{\mu} X'$ and $q \xrightarrow{\mu} q'$ imply $X' \preceq_{\text{cnv}}^{\text{set}} q'$.

Lemma 37. *Let $\mathcal{L}_A, \mathcal{L}_B \in \text{FWD}$. For every $X, X' \in \mathcal{P}^+(A)$ and $q \in B$, such that $X \preceq_{\text{acc}}^{\text{set}} q$, then*

1. $q \xrightarrow{\tau} q'$ implies $X \preceq_{\text{acc}}^{\text{set}} q'$,
2. if $X \downarrow_i$ then for every $\mu \in \text{Act}$, every q' and X' such that $q \xrightarrow{\mu} q'$ and $X \xrightarrow{\mu} X'$ we have $X' \preceq_{\text{acc}}^{\text{set}} q'$.

The main technical work for the proof of soundness is carried out by the next lemma.

Lemma 38. *Let $\mathcal{L}_A, \mathcal{L}_B \in \text{FWD}$ and $\mathcal{L}_C \in \text{FDB}$. For every set of servers $X \in \mathcal{P}^+(A)$, server $q \in B$ and client $r \in C$, if $X \text{ MUST}_{\text{aux}} r$ and $X \preceq_{\text{AS}}^{\text{set}} q$ then $q \text{ MUST}_i r$.*

Proposition 5 (Soundness). *For every $\mathcal{L}_A, \mathcal{L}_B \in \text{FDB}$ and servers $p \in A, q \in B$, if $\text{FW}(p) \preceq_{\text{AS}} \text{FW}(q)$ then $p \sqsubseteq_{\text{MUST}} q$.*

Proof. Lemma 3 ensures that the result follows if we prove that $\text{FW}(p) \sqsubseteq_{\text{MUST}} \text{FW}(q)$. Fix a client r such that $\text{FW}(p) \text{ MUST}_i r$. Lemma 38 implies the required $\text{FW}(q) \text{ MUST}_i r$, if we show that

- (i) $\{\text{FW}(p)\} \text{ MUST}_{\text{aux}} r$, and that
- (ii) $\{\text{FW}(p)\} \preceq_{\text{AS}}^{\text{set}} \text{FW}(q)$.

The first fact follows from the assumption that $\text{FW}(p) \text{ MUST}_i r$ and Lemma 7 applied to the singleton $\{\text{FW}(p)\}$. The second fact follows from the hypothesis that $\text{FW}(p) \preceq_{\text{AS}} \text{FW}(q)$ and Lemma 35.

E.1 Technical results to prove soundness

We now discuss the proofs of the main technical results behind Proposition 5. The predicate MUST_{aux} is monotonically decreasing with respect to its first argument, and it enjoys properties analogous to the ones of MUST_i that have been shown in Lemma 11 and Lemma 12.

Lemma 39. *For every LTS $\mathcal{L}_A, \mathcal{L}_B$ and every set $X_1 \subseteq X_2 \subseteq A$, client $r \in B$, if $X_2 \text{ MUST}_{\text{aux}} r$ then $X_1 \text{ MUST}_{\text{aux}} r$.*

Lemma 40. *Let $\mathcal{L}_A \in \text{FWD}$ and $\mathcal{L}_B \in \text{FDB}$. For every set $X \in \mathcal{P}^+(A)$, client $r \in B$, if $\neg \text{GOOD}(r)$ and $X \text{ MUST}_{\text{aux}} r$ then $X \downarrow_i$.*

Lemma 41. *For every $\mathcal{L}_A, \mathcal{L}_B$, every set $X_1, X_2 \in \mathcal{P}^+(A)$, and client $r \in B$, if $X_1 \text{ MUST}_{\text{aux}} r$ and $X_1 \xrightarrow{\varepsilon} X_2$ then $X_2 \text{ MUST}_{\text{aux}} r$.*

Lemma 42. *For every LTS $\mathcal{L}_A, \mathcal{L}_B$ and every $X \in \mathcal{P}^+(A)$ and $r \in B$, if $X \text{ MUST}_{\text{aux}} r$ then for every X' such that*

- (a) *If $X \xrightarrow{\varepsilon} X'$ then $X' \text{ MUST}_{\text{aux}} r$,*

(b) For any $\mu \in \text{Act}$ and client r' , if $X \xrightarrow{\mu} X'$, $r \xrightarrow{\bar{\mu}} r'$ and $\neg\text{GOOD}(r)$, then $X' \text{MUST}_{\text{aux}} r'$.

Lemma 43. Given two LTS \mathcal{L}_A and \mathcal{L}_B then for every $X \in \mathcal{P}^+(A)$ and $r \in B$, if for each $p \in X$ we have that $p \text{MUST}_i r$, then $X \text{MUST}_{\text{aux}} r$.

Lemma 44. Let $\mathcal{L}_A, \mathcal{L}_B \in \text{FWD}$ and $\mathcal{L}_C \in \text{FDB}$. For every $X \in \mathcal{P}^+(A)$ and $q \in B$ such that $X \preceq_{\text{AS}}^{\text{set}} q$, for every $r \in C$ if $\neg\text{GOOD}(r)$ and $X \text{MUST}_{\text{aux}} r$ then $q \parallel r \xrightarrow{\tau}$.

Proof. If either $q \xrightarrow{\tau}$ or $r \xrightarrow{\tau}$ then we prove that $q \parallel r$ performs a τ -transition vis [S-SRV] or [S-CLT], so suppose that both q and r are stable. Since q is stable we know that

$$\mathcal{A}_{\text{fw}}(q, \varepsilon) = \{O(q)\}$$

The hypotheses $\neg\text{GOOD}(r)$ and $X \text{MUST}_{\text{aux}} r$ together with Lemma 40 imply $X \downarrow_i$ and thus $X \downarrow \varepsilon$. The hypothesis $X \preceq_{\text{acc}}^{\text{set}} q$ with $s = \varepsilon$, gives us a p' such that $p \xrightarrow{\varepsilon} p' \xrightarrow{\tau}$ and $O(p') \subseteq O(q)$. By definition there exists the weak silent trace $X \xRightarrow{} X'$ for some set X' such that $\{p'\} \subseteq X'$. The hypothesis $X \text{MUST}_{\text{aux}} r$ together with Lemma 41 and Lemma 39 ensure that $\{p'\} \text{MUST}_{\text{aux}} r$.

As $\neg\text{GOOD}(r)$, $\{p'\} \text{MUST}_{\text{aux}} r$ must have been derived using rule [IND-RULE] which implies that $p' \parallel r \xrightarrow{\tau}$. As both r is stable by assumption, and p' is stable by definition, this τ -transition must have been derived using [S-COM], and so $p' \xrightarrow{\mu}$ and $r \xrightarrow{\bar{\mu}}$ for some $\mu \in \text{Act}$. Now we distinguish whether μ is an input or an output. In the first case μ is an input. Since $\mathcal{L}_B \in \text{FWD}$ we use the INPUT-BOOMERANG axiom to prove $q \xrightarrow{\mu}$, and thus $q \parallel r \xrightarrow{\tau}$ via rule [S-COM]. In the second case μ is an output, and so the inclusion $O(p') \subseteq O(q)$ implies that $q \xrightarrow{\mu}$, and so we conclude again applying rule [S-COM].

Lemma 45. Let $\mathcal{L}_A, \mathcal{L}_B \in \text{FWD}$. For every $X \in \mathcal{P}^+(A)$ and $q, q' \in B$, such that $X \preceq_{\text{AS}}^{\text{set}} q$, then for every $\mu \in \text{Act}$, if $X \downarrow \mu$ and $q \xrightarrow{\mu} q'$ then $X \xrightarrow{\mu}$.

Proof. Then, from $X \preceq_{\text{cnv}}^{\text{set}} q$ and $X \downarrow \mu$ we have that $q \downarrow \mu$ and thus $q' \downarrow_i$. As q' converges, there must exist q'' such that

$$q \xrightarrow{\mu} q' \xrightarrow{\varepsilon} q'' \xrightarrow{\tau}$$

and so $\mathcal{A}_{\text{fw}}(q, \mu, \rightarrow_B) \neq \emptyset$. An application of the hypothesis $X \preceq_{\text{acc}}^{\text{set}} q$ implies that there exists a set $\hat{O} \in \mathcal{A}_{\text{fw}}(X, \mu, \rightarrow_A)$, and thus there exist two servers $p' \in X$ and p'' such that $p' \xrightarrow{\mu} p'' \xrightarrow{\tau}$. Since $p' \in X$ it follows that $X \xrightarrow{\mu}$.

Lemma 36 Let $\mathcal{L}_A, \mathcal{L}_B \in \text{FWD}$. For every set $X \in \mathcal{P}^+(A)$, and $q \in B$, such that $X \preceq_{\text{cnv}}^{\text{set}} q$ then

1. $q \xrightarrow{\tau} q'$ implies $X \preceq_{\text{cnv}}^{\text{set}} q'$,
2. if $X \downarrow_i$ and $q \xrightarrow{\mu} q'$ then for every set $X \xrightarrow{\mu} X'$ we have that $X' \preceq_{\text{cnv}}^{\text{set}} q'$.

Proof. We first prove part (1). Let us fix a trace s such that $X \Downarrow s$. We must show $q' \Downarrow s$. An application of the hypothesis $X \preceq_{\text{cnv}}^{\text{set}} q$ ensures $q \Downarrow s$. From the transition $q \xrightarrow{\tau} q'$ and the fact that convergence is preserved by the τ -transitions we have that $q' \Downarrow s$ as required.

We now prove part (2). Fix a trace s such that $X' \Downarrow s$. Since $q \xrightarrow{\mu} q'$, the required $q' \Downarrow s$ follows from $q \Downarrow \mu.s$. Thanks to the hypothesis $X \preceq_{\text{cnv}}^{\text{set}} q$ it suffices to show that $X \Downarrow \mu.s'$, *i.e.* that

$$\forall p \in X.p \Downarrow \mu.s'$$

Fix a server $p \in X$. We must show that

1. $p \downarrow_i$ and that
2. for any p' such that $p \xrightarrow{\mu} p'$ we have $p' \Downarrow s$.

The first requirement follows from the hypothesis $X \downarrow_i$. The second requirement follows from the transition $p \xrightarrow{\mu} p'$, from the assumption $X' \Downarrow s$, and the hypothesis that $X \xrightarrow{\mu} X'$, which ensures that $p' \in X'$ and thus by definition of $X' \Downarrow s$ that $p' \Downarrow s$.

Lemma 37 Let $\mathcal{L}_A, \mathcal{L}_B \in \text{FWD}$. For every $X, X' \in \mathcal{P}^+(A)$ and $q \in B$, such that $X \preceq_{\text{acc}}^{\text{set}} q$, then

1. $q \xrightarrow{\tau} q'$ implies $X \preceq_{\text{acc}}^{\text{set}} q'$,
2. for every $\mu \in \text{Act}$, if $X \downarrow_i$, then for every $q \xrightarrow{\mu} q'$ and set $X \xrightarrow{\mu} X'$ we have $X' \preceq_{\text{acc}}^{\text{set}} q'$.

Proof. To prove part (1) fix a trace $s \in \text{Act}^*$ such that $X \Downarrow s$. We have to explain why $\mathcal{A}_{\text{fw}}(X, s) \ll \mathcal{A}_{\text{fw}}(q', s)$. By unfolding the definitions, this amounts to showing that

$$\forall O \in \mathcal{A}_{\text{fw}}(q', s). \exists p_{\text{attaboy}} \in X. \exists \widehat{O} \in \mathcal{A}_{\text{fw}}(p_{\text{attaboy}}, s). \widehat{O} \subseteq O \quad (\star)$$

Fix a set $O \in \mathcal{A}_{\text{fw}}(q', s)$. By definition there exists some q'' such that $q' \xrightarrow{s} q'' \xrightarrow{\tau} q$, and that $O = O(q'')$. The definition of $\mathcal{A}_{\text{fw}}(-, -)$ and the silent move $q \xrightarrow{\tau} q'$ ensures that $O \in \mathcal{A}_{\text{fw}}(q, s)$. The hypothesis $X \preceq_{\text{acc}}^{\text{set}} q$ and that $X \Downarrow s$ now imply that $\mathcal{A}_{\text{fw}}(X, s) \ll \mathcal{A}_{\text{fw}}(q, s)$, which together with $O \in \mathcal{A}_{\text{fw}}(q, s)$ implies exactly Equation (\star) .

We now prove part (2). To show $X' \preceq_{\text{acc}}^{\text{set}} q'$ fix a trace $s \in \text{Act}^*$ such that $X' \Downarrow s$.

We have to explain why $\mathcal{A}_{\text{fw}}(X, s) \ll \mathcal{A}_{\text{fw}}(q', s)$. By unfolding the definitions we obtain our aim,

$$\forall O \in \mathcal{A}_{\text{fw}}(q', s). \exists p_{\text{attaboy}} \in X'. \exists \widehat{O} \in \mathcal{A}_{\text{fw}}(p_{\text{attaboy}}, s). \widehat{O} \subseteq O \quad (\star\star)$$

To begin with, we prove that $X \Downarrow \mu.s$. Since $X \xrightarrow{\mu} X'$ we know that $X \xrightarrow{\mu} X'$. This, together with $X \downarrow_i$ and $X' \Downarrow s$ implies the convergence property we are after.

Now fix a set $O \in \mathcal{A}_{\text{fw}}(q', s)$. Thanks to the transition $q \xrightarrow{\mu} q'$, we know that $O \in \mathcal{A}_{\text{fw}}(q, \mu.s)$. The hypothesis $X \preceq_{\text{acc}}^{\text{set}} q$ together with $X \Downarrow \mu.s$ implies that there exists a server $p_{\text{attaboy}} \in X$ such that there exists an $\widehat{O} \in \mathcal{A}_{\text{fw}}(p_{\text{attaboy}}, \mu.s)$. This means that $p_{\text{attaboy}} \xrightarrow{\mu} p'_{\text{attaboy}}$ and that $\widehat{O} \in \mathcal{A}_{\text{fw}}(p'_{\text{attaboy}}, s)$. Since $X \xrightarrow{\mu} X'$ we know that $p'_{\text{attaboy}} \in X'$ and this concludes the argument.

Lemma 46. *For every $\mathcal{L}_A \in \text{FWD}$, $\mathcal{L}_B \in \text{FDB}$, every set of processes $X \in \mathcal{P}^+(A)$, every $r \in B$, and every $\mu \in \text{Act}$, if $X \text{ MUST}_{\text{aux}} r$, $\neg \text{GOOD}(r)$ and $r \xrightarrow{\mu}$ then $X \Downarrow \bar{\mu}$.*

Lemma 38 Let $\mathcal{L}_A, \mathcal{L}_B \in \text{FWD}$ and $\mathcal{L}_C \in \text{FDB}$. For every set of processes $X \in \mathcal{P}^+(A)$, server $q \in B$ and client $r \in C$, if $X \text{ MUST}_{\text{aux}} r$ and $X \preceq_{\text{AS}}^{\text{set}} q$ then $q \text{ MUST}_i r$.

Proof. We proceed by induction on the derivation of $X \text{ MUST}_{\text{aux}} r$. In the base case, $\text{GOOD}(r)$ so we trivially derive $q \text{ MUST}_i r$. In the inductive case the proof of the hypothesis $X \text{ MUST}_{\text{aux}} r$ terminates with an application of [MSET-STEP]. Since $\neg \text{GOOD}(r)$, we show the result applying [IND-RULE]. This requires us to prove that

- (1) $q \parallel r \xrightarrow{\tau}$, and that
- (2) for all q', r' such that $q \parallel r \xrightarrow{\tau} q' \parallel r'$ we have $q' \text{ MUST}_i r'$.

The first fact is a consequence of Lemma 44, which we can apply because $\neg \text{GOOD}(r)$ and thanks to the hypothesis $X \preceq_{\text{acc}}^{\text{set}} q$ and $X \text{ MUST}_{\text{aux}} r$. To prove the second fact, fix a transition $q \parallel r \xrightarrow{\tau} q' \parallel r'$. We have to explain why the following properties are true,

- (a) for every q' . $q \xrightarrow{\tau} q'$ implies $q' \text{ MUST}_i r$,
- (b) for every r . $r \xrightarrow{\tau} r'$ implies $q \text{ MUST}_i r'$,
- (c) for every q', r' and $\mu \in \text{Act}$, $q \xrightarrow{\mu} q'$ and $r \xrightarrow{\bar{\mu}} r'$ imply $q' \text{ MUST}_i r'$.

First, note that $X \text{ MUST}_{\text{aux}} r$, $\neg \text{GOOD}(r)$, and Lemma 40 imply $X \Downarrow$. Second, the inductive hypotheses state that for every r' , non-empty set X' , and q the following facts hold,

- (i) $X \xrightarrow{\tau} X'$ and $X' \preceq_{\text{AS}}^{\text{set}} q$ implies $q \text{ MUST}_i r$,
- (ii) $r \xrightarrow{\tau} r'$ and $X \preceq_{\text{AS}}^{\text{set}} q$ implies $q \text{ MUST}_i r'$,
- (iii) for every and $\mu \in \text{Act}$, $X \xrightarrow{\bar{\mu}} X'$ and $r \xrightarrow{\mu} r'$, and $X' \preceq_{\text{AS}}^{\text{set}} q$ implies $q \text{ MUST}_i r'$.

To prove (a) we use $X \Downarrow$ and the hypothesis $X \preceq_{\text{cnv}}^{\text{set}} q$ to obtain $q \Downarrow_i$. A rule induction on $q \Downarrow_i$ now suffices: in the base case (a) is trivially true and in the inductive case (a) follows from Lemma 36(1) and Lemma 37(1), and the inductive hypothesis.

The requirement (b) follows directly from the hypothesis $X \preceq_{\text{AS}}^{\text{set}} q$ and part (ii) of the inductive hypothesis.

To see why (c) holds, fix an action μ such that $q \xrightarrow{\mu} q'$ and $r \xrightarrow{\bar{\mu}} r'$. Since $\neg\text{GOOD}(r)$ Lemma 46 implies that $X \Downarrow \mu$, and so Lemma 45 proves that $X \xrightarrow{\mu}$. In turn this implies that there exists a X' such that $X \xrightarrow{\mu} X'$, and thus Lemma 36(2) and Lemma 37(2) prove that $X' \preceq_{\text{AS}}^{\text{set}} q'$ holds, and (iii) ensures the result, *i.e.* that $q' \text{MUST}_i r'$.

F Technical details about coinduction

Here we outline the lemma that lets us prove Equation (9).

Lemma 47. *For every mailbox M and every set X of processes, we have*

$$\{\tau.(\bar{a} \parallel \bar{b}) + \tau.(\bar{a} \parallel \bar{c}) \triangleright M\} \cup X \preceq_{\text{co}} \bar{a} \parallel (\tau.\bar{b} + \tau.\bar{c}) \triangleright M. \quad (\text{CH})$$

Proof (sketch). Let M be a mailbox and X be a set of processes. We focus on the last point of Definition 12, for the other cases are similar. Consider an arbitrary reduction of the RHS of Equation (CH):

$$\bar{a} \parallel (\tau.\bar{b} + \tau.\bar{c}) \triangleright M \xrightarrow{\mu} q \quad (10)$$

We proceed by case analysis.

1. The cases that correspond to a reduction of the mailbox M are dealt with directly using the coinductive hypothesis (CH), since the mailbox is quantified universally in (CH). In more detail, consider the case where the reduction (10) is of the form:

$$\bar{a} \parallel (\tau.\bar{b} + \tau.\bar{c}) \triangleright M \xrightarrow{b} \bar{a} \parallel (\tau.\bar{b} + \tau.\bar{c}) \triangleright \{\bar{b}\} \uplus M$$

Let X' be $WD(b, \{\tau.(\bar{a} \parallel \bar{b}) + \tau.(\bar{a} \parallel \bar{c}) \triangleright M\} \uplus X)$, that is, the unique set of processes such that

$$\{\tau.(\bar{a} \parallel \bar{b}) + \tau.(\bar{a} \parallel \bar{c}) \triangleright M\} \cup X \xrightarrow{b} X'$$

It is easy to prove that $\{\tau.(\bar{a} \parallel \bar{b}) + \tau.(\bar{a} \parallel \bar{c}) \triangleright \{\bar{b}\} \uplus M\} \in X'$ and thus we conclude by applying the coinductive hypothesis (CH).

2. If (10) corresponds to a transition of the process, it must be that $\mu = \bar{a}$ and $q = \tau.\bar{b} + \tau.\bar{c} \triangleright M$. In that case, the set X' of processes reached from $\{\tau.(\bar{a} \parallel \bar{b}) + \tau.(\bar{a} \parallel \bar{c}) \triangleright M\} \uplus X$ while outputting a contains \bar{b} and \bar{c} , so that $X' \preceq_{\text{co}} q$ follows from the general fact that $\{p, q\} \preceq_{\text{MUST}}^{\text{Set}} \tau.p + \tau.q$. \square

G Traces in normal form

Let MI be the set of multisets of input actions, and let $\text{nf} : \text{Act}^* \rightarrow (MI \times MO)^*$ be the function

$$\text{nf}(s) = (I_0, M_0), (I_1, M_2), \dots, (I_n, M_n)$$

$$\begin{aligned}
\text{nf}(\varepsilon) &= \varepsilon \\
\text{nf}(s) &= \text{nf}'(s, \emptyset, \emptyset) \\
\text{nf}'(\varepsilon, I, M) &= (I, M) \\
\text{nf}'(a.s, I, M) &= \text{nf}'(s, \{a\} \uplus I, M) \\
\text{nf}'(\bar{a}.s, I, M) &= \begin{cases} (I, \{\bar{a}\} \uplus M), \text{nf}'(s, I, \emptyset) & \text{if } s = b.s' \text{ for some } b, s' \\ \text{nf}'(s, I, \{\bar{a}\} \uplus M) & \text{otherwise} \end{cases}
\end{aligned}$$

Fig. 12. Definition of the trace normalisation function nf

which is defined inductively in Figure 12. The intuition is that given a trace s , the function nf forgets the order of actions in sequences of consecutive inputs, and in sequences of consecutive outputs, thereby transforming them in multisets. On the other hand nf preserves the causal order in the trace, in the sense that the execution of all inputs in I_k is necessary to execute any of the outputs in M_k , as well as the actions in the multisets at index h with $k < h$. For instance:

$$\text{nf}(\text{cabddae}\bar{f}e) = (\{c, a\}, \{\bar{b}, \bar{d}, \bar{d}\}), (\{a\}, \{\bar{e}, \bar{f}\}), (\{e\}, \emptyset)$$

Let σ range over the set $(MI \times MO)^*$. We say that σ is a trace in *normal form*, and we write $p \xrightarrow{\sigma} q$, whenever there exists $s \in \text{Act}^*$ such that $p \xrightarrow{s} q$ and $\text{nf}(s) = \sigma$.

Definition 16. We lift in the obvious way the predicates \preceq_{cnv} , \preceq_{acc} , and $\preceq_{\text{acc}}^{\text{fw}}$ to traces in normal forms. For every $\mathcal{L}_A, \mathcal{L}_B$ and $p \in A, q \in B$, let

- $p \preceq_{\text{asyn}}^{\text{cnv}} q$ to mean $\forall \sigma \in (MI \times MO)^*. p \Downarrow \sigma$ implies $q \Downarrow \sigma$,
- $p \preceq_{\text{asyn}}^{\text{acc}} q$ to mean $\forall \sigma \in (MI \times MO)^*. p \Downarrow \sigma$ implies $\mathcal{A}_{\text{fw}}(p, \sigma) \ll \mathcal{A}_{\text{fw}}(q, \sigma)$,
- $p \preceq_{\text{AS}}^{\text{NF}} q$ whenever $q \preceq_{\text{asyn}}^{\text{cnv}} p \wedge p \preceq_{\text{asyn}}^{\text{acc}} q$.

In LTSs with forwarding, the transition relation \longrightarrow is *input-receptive* (Axiom (IB4), Table 2 of [90]), and thus Lemma 48 shows that \longrightarrow enjoys a restricted version of INPUT-COMMUTATIVITY, and that so does its weak version. Sequences of input actions $s \in \mathcal{N}^*$ enjoy a form of diamond property in \Longrightarrow . The crucial fact pertains consecutive input actions. Recall that \simeq denotes any equivalence that satisfies the compatibility property depicted in Figure 5.

Lemma 48. For every $\mathcal{L} = \langle A, \text{Act}, \longrightarrow \rangle \in \text{FWD}$, every $p, q \in A$ and $a, b \in \mathcal{N}$, if $p \xrightarrow{a.b} q$ then $p \xrightarrow{b.a} \cdot \simeq q$.

Lemma 48, together with an induction on traces, allows us to prove that nf preserves convergence and acceptance sets.

Lemma 49. For every $\mathcal{L} = \langle A, \text{Act}, \longrightarrow \rangle \in \text{FWD}$, every $p \in A$ and $s \in \text{Act}^*$ we have that

$$\begin{array}{ll}
 \text{[INPUT]} & \frac{}{a.p \xrightarrow{a} p} \quad \text{[TAU]} \quad \frac{}{\tau.p \xrightarrow{\tau} p} \\
 \text{[MB-OUT]} & \frac{}{\bar{a} \xrightarrow{\bar{a}} 0} \quad \text{[UNF]} \quad \frac{}{\text{rec}x.p \xrightarrow{\tau} p[\text{rec}x.p/x]} \\
 \text{[SUM-L]} & \frac{p \xrightarrow{\alpha} p'}{p + q \xrightarrow{\alpha} p'} \quad \text{[SUM-R]} \quad \frac{q \xrightarrow{\alpha} q'}{p + q \xrightarrow{\alpha} q'} \\
 \text{[PAR-L]} & \frac{p \xrightarrow{\alpha} p'}{p \parallel q \xrightarrow{\alpha} p' \parallel q} \quad \text{[PAR-R]} \quad \frac{q \xrightarrow{\alpha} q'}{p \parallel q \xrightarrow{\alpha} p \parallel q'} \\
 \text{[COM]} & \frac{p \xrightarrow{\mu} p' \quad q \xrightarrow{\bar{\mu}} q'}{p \parallel p' \xrightarrow{\tau} q \parallel q'}
 \end{array}$$

Fig. 13. The LTS of processes. The meta-variables are $a \in \mathcal{N}$, $\mu \in \text{Act}$, and $\alpha \in \text{Act}_\tau$.

1. $p \xRightarrow{s} q$ iff $p \xRightarrow{\text{nf}(s)} \cdot \simeq q$, and if the first trace does not pass through a successful state then the normal form does not either,
2. $p \Downarrow \text{nf}(s)$ iff $p \Downarrow s$,
3. $\mathcal{A}_{\text{fw}}(p, s) = \mathcal{A}_{\text{fw}}(p, \text{nf}(s))$.

We thereby obtain two other characterisations of the contextual preorder $\sqsubseteq_{\text{MUST}}$: Theorem 1 and Lemma 49 ensure that the preorders $\sqsubseteq_{\text{MUST}}$ and $\preceq_{\text{AS}}^{\text{NF}}$ coincide.

Corollary 4. For every $\mathcal{L}_A, \mathcal{L}_B \in \text{FDB}$, every $p \in A$ and $q \in B$,

$$p \sqsubseteq_{\text{MUST}} q \quad \text{iff} \quad \text{FW}(p) \preceq_{\text{AS}}^{\text{NF}} \text{FW}(q).$$

H Asynchronous CCS

Here we recall the syntax and the LTS of asynchronous CCS, or **ACCS** for short, a version of CCS where outputs have no continuation and sum is restricted to input- and τ -guards.

$$p, q, r ::= \bar{a} \mid g \mid p \parallel p \mid \text{rec}x.p \mid x, \quad g ::= 0 \mid 1 \mid a.p \mid \tau.p \mid g + g$$

This calculus, which is inspired by the variant of the asynchronous π -calculus considered by [4,5] for their study of asynchronous bisimulation, was first investigated by [90], and subsequently resumed by other authors such as [23]. Different asynchronous variants of CCS were studied in the same frame of time by [80], whose calculus included output prefixing and operators from ACP, and by [38], whose calculus TACCS included asynchronous output prefixing and featured

two forms of choice, internal and external, in line with previous work on testing semantics [76].

The syntax of terms is given in Equation (5). As usual, $\text{rec}x.p$ binds the variable x in p , and we use standard notions of free variables, open and closed terms. Processes, ranged over by p, q, r, \dots are *closed* terms. The operational semantics of processes is given by the LTS $\langle \text{ACCS}, \text{Act}_\tau, \longrightarrow \rangle$ specified by the rules in Figure 13.

The prefix $a.p$ represents a *blocked* process, which waits to perform the input a , *i.e.* to interact with the atom \bar{a} , and then becomes p ; and atoms \bar{a}, \bar{b}, \dots represent output messages. We will discuss in detail the role played by atoms in the calculus, but we first overview the rest of the syntax. We include 1 to syntactically denote successful states. The prefix $\tau.p$ represents a process that does one step of internal computation and then becomes p . The sum $g_1 + g_2$ is a process that can behave as g_1 or g_2 , but not both. Thus, for example $\tau.p + \tau.q$ models an **if ...then ...else**, while $a.p + b.q$ models a **match ...with**. Note that the sum operator is only defined on *guards*, namely it can only take as summands $0, 1$ or input-prefixed and τ -prefixed processes. While the restriction to guarded sums is a standard one, widely adopted in process calculi, the restriction to input and τ guards is specific to asynchronous calculi. We will come back to this point after discussing atoms and mailboxes. Parallel composition $p \parallel q$ runs p and q concurrently, allowing them also to interact with each other, thanks to rule [COM]. For example

$$b.a.0 \parallel b.c.0 \parallel (\bar{a} \parallel \bar{b} \parallel \bar{c}) \quad (11)$$

represents a system in which two concurrent processes, namely $b.a.0$ and $b.c.0$, are both ready to consume the message \bar{b} from a third process, namely $\bar{a} \parallel \bar{b} \parallel \bar{c}$. This last process is a parallel product of atoms, and it is not guarded, hence it is best viewed as an unordered mailbox shared by *all* the processes running in parallel with it. For instance in (11) the terms $b.a.0$ and $b.c.0$ share the mailbox $\bar{a} \parallel \bar{b} \parallel \bar{c}$. Then, depending on which process consumes \bar{b} , the overall process will evolve to either $b.c.0 \parallel \bar{c}$ or $b.a.0 \parallel \bar{a}$, which are both stuck.¹³

Concerning the sum construct, we follow previous work on asynchronous calculi ([4,5,87,23]) and only allow input-prefixed or τ -prefixed terms as summands. The reason for forbidding atoms in sums is that the nondeterministic sum is essentially a synchronising operator: the choice is solved by executing an action in one of the summands and *simultaneously* discarding all the other summands. Then, if an atom were allowed to be a summand, this atom could be discarded by performing an action in another branch of the choice. This would mean that a

¹³ The global shared mailbox that we treat is reminiscent but less general than the chemical “soup” of [17]. In that context the components of the soup are not just atoms, but whole parallel components: in fact, the chemical soup allows parallel components to come close in order to react with each other, exactly as the structural congruence of [73], which indeed was inspired by the Chemical Abstract Machine.

process would have the ability to withdraw a message from the mailbox without consuming it, thus contradicting the intuition that the mailbox is a shared entity which is out of the control of any given process, and with which processes can only interact by feeding a message into it or by consuming a message from it. In other words, this restriction on the sum operator ensures that atoms indeed represent messages in a global mailbox. For further details see the discussion on page 191 of [87].

A structural induction on the syntax ensures that processes perform only a finite number of outputs:

Lemma 50. *For every $p \in \text{ACCS}$. $|O(p)| \in \mathbb{N}$.*

Together with Lemma 53, this means that at any point of every execution the global mailbox contains a finite number of messages. Since the LTS is image-finite under any visible action, a consequence of Lemma 50 is that the number of reducts of a program is finite.

Lemma 51. *For every $p \in \text{ACCS}$. $|\{p' \in \text{ACCS} \mid p \xrightarrow{\tau} p'\}| \in \mathbb{N}$.*

Proof. Structural induction on p . The only non-trivial case is if $p = p_1 \parallel p_2$. In this case the result is a consequence of the inductive hypothesis, of Lemma 50 and of the following fact: $p \xrightarrow{\tau} q$ iff

1. $p_1 \xrightarrow{\tau} p'_1$ and $q = p'_1 \parallel p_2$,
2. $p_2 \xrightarrow{\tau} p'_2$ and $q = p_1 \parallel p'_2$,
3. $p_1 \xrightarrow{a} p'_1$ and $p_2 \xrightarrow{\bar{a}} p'_2$ and $q = p'_1 \parallel p'_2$,
4. $p_1 \xrightarrow{\bar{a}} p'_1$ and $p_2 \xrightarrow{a} p'_2$ and $q = p'_1 \parallel p'_2$.

In the third case the number of possible output actions \bar{a} is finite thanks to Lemma 50, and so is the number of reducts p'_1 and p'_2 , so the set of term $p'_1 \parallel p'_2$ is decidable. The same argument works for the fourth case.

Thanks to Lemma 50 and Lemma 51, Lemma 51 holds also for the LTS modulo structural congruence, i.e. $\langle \text{ACCS}_{\equiv}, \xrightarrow{\tau}_{\equiv}, \text{Act}_{\tau} \rangle$.

H.1 Structural equivalence and its properties

To manipulate the syntax of processes we use a standard structural congruence denoted \equiv , stating that ACCS is a commutative monoid with identity 0 with respect to both sum and parallel composition.

A first fact is the following one.

Lemma 52. *For every $\mu \in \bar{\mathcal{N}}$ and $\alpha \in \text{Act}_{\tau}$, if $p \xrightarrow{\mu, \alpha} q$ then $p \xrightarrow{\alpha, \mu} \cdot \equiv q$.*

As sum and parallel composition are commutative monoids, we use the notation

$$\begin{aligned} \Sigma\{g_0, g_1, \dots, g_n\} & \text{ to denote } g_0 + g_1 + \dots + g_n \\ \Pi\{p_0, p_1, \dots, p_n\} & \text{ to denote } p_0 \parallel p_1 \parallel \dots \parallel p_n \end{aligned}$$

```

Class LtsEq (A L : Type) ` {Lts A L} := {
  eq_rel : A → A → Prop;

  eq_rel_refl p : eq_rel p p;
  eq_symm p q : eq_rel p q → eq_rel q p;
  eq_trans p q r :
  eq_rel p q → eq_rel q r → eq_rel p r;

  eq_spec p q (α : Act L) :
  (∃ p', (eq_rel p p') ∧ p' →{α} q)
  →
  (∃ q', p →{α} q' ∧ (eq_rel q' q))
}.

```

Fig. 14. A typeclass for LTSs where a structural congruence exists over states.

This notation is useful to treat the global shared mailbox. In particular, if $\{\mu_0, \mu_1, \dots, \mu_n\}$ is a multiset of output actions, then the syntax $\Pi\{\mu_0, \mu_1, \dots, \mu_n\}$ represents the shared mailbox that contains the messages μ_i ; for instance we have $\Pi\{\bar{a}, \bar{a}, \bar{c}\} = \boxed{\bar{a} \parallel \bar{a} \parallel \bar{c}}$. We use the colour — to highlight the content of the mailbox. Intuitively a shared mailbox contains the messages that are ready to be read, i.e. the outputs that are immediately available (i.e. not guarded by any prefix operation). For example in

$$\boxed{\bar{c}} \parallel a.(\bar{b} \parallel c.d.1) \parallel \boxed{\bar{d}} \parallel \tau.\bar{e}$$

the mailbox is $\boxed{\bar{c} \parallel \bar{d}}$. The global mailbox that we denote with — is exactly the buffer B in the *configurations* of [94], and reminiscent of the ω used by [27]. The difference is that ω represents an unbounded *ordered* queue, while our mailbox is an unbounded *unordered* buffer.

As for the relation between output actions in the LTS and the global mailbox, an output \bar{a} can take place if and only if the message \bar{a} appears in the mailbox:

Lemma 53. *For every $p \in ACCS$,*

1. *for every $a \in \mathcal{N}$. $p \xrightarrow{\bar{a}} p'$ implies $p \equiv p' \parallel \boxed{\bar{a}}$,*
2. *there exists p' such that $p \equiv p' \parallel \boxed{\Pi M}$, and p' performs no output action.*

This lemma and Lemma 52 essentially hold, because, as already pointed out in Section 2, the syntax enforces outputs to have no continuation.

The following lemma states a fundamental fact ([58, Lemma 2.13], [74, Proposition 5.2], [87, Lemma 1.4.15]). Its proof is so tedious that even the references we have given only sketch it. In this paper we follow the masters example, and give merely a sketch. However, we have a complete machine-checked proof.

[S-SZERO]	$p + \mathbf{0} \equiv p$	
[S-SCOM]	$p + q \equiv q + p$	
[S-SASS]	$(p + q) + r \equiv p + (q + r)$	
[S-PZERO]	$p \parallel \mathbf{0} \equiv p$	
[S-PCOM]	$p \parallel q \equiv q \parallel p$	
[S-PASS]	$(p \parallel q) \parallel r \equiv p \parallel (q \parallel r)$	
[S-REFL]	$p \equiv p$	
[S-SYMM]	$p \equiv q$	<i>if $q \equiv p$</i>
[S-TRANS]	$p \equiv q$	<i>if $p \equiv p'$ and $p' \equiv q$</i>
[S-PREFIX]	$\alpha.p \equiv \alpha.q$	<i>if $p \equiv q$</i>
[S-SUM]	$p + q \equiv p' + q$	<i>if $p \equiv p'$</i>
[S-PPAR]	$p \parallel q \equiv p' \parallel q$	<i>if $p \equiv p'$</i>

Fig. 15. Rules to define structural congruence on ACCS.

Lemma 54. *For every $p, q \in \text{ACCS}$ and $\alpha \in \text{Act}_\tau$, $p \equiv \cdot \xrightarrow{\alpha} q$ implies $p \xrightarrow{\alpha} \cdot \equiv q$.*

Proof (Proof sketch). We need to show that if there exists a process p' such that $p \equiv p'$ and $p' \xrightarrow{\alpha} q$ then there exists a process q' such that $p \xrightarrow{\alpha} p'$ and $p' \equiv q$. The proof is by induction on the derivation $p \equiv p'$.

We illustrate one case with the rule [S-TRANS]. The hypotheses tell us that there exists \hat{p} such that $p \equiv \hat{p}$ and $\hat{p} \equiv p'$, that $p' \xrightarrow{\alpha} q$, and the inductive hypotheses that

- (a) for all q' s.t. $\hat{p} \xrightarrow{\alpha} q'$ implies that there exists a \hat{q} such that $p \xrightarrow{\alpha} \hat{q}$ and $\hat{q} \equiv q'$
- (b) for all q' s.t. $p' \xrightarrow{\alpha} q'$ implies that there exists a \hat{q} such that $\hat{p} \xrightarrow{\alpha} \hat{q}$ and $\hat{q} \equiv q'$

By combining part (b) and $p' \xrightarrow{\alpha} q$ we obtain a \hat{q}_1 such that $\hat{p} \xrightarrow{\alpha} \hat{q}_1$ and $\hat{q}_1 \equiv q$. Using part (a) together with $\hat{p}_1 \xrightarrow{\alpha} \hat{q}_1$ we have that there exists a \hat{q}_2 such that $p \xrightarrow{\alpha} \hat{q}_2$ and $\hat{q}_2 \equiv \hat{q}_1$. We then have that $p \xrightarrow{\alpha} \hat{q}_2$ and it remains to show that $\hat{q}_2 \equiv q$. We use the transitivity property of the structural congruence relation to show that $\hat{q}_2 \equiv \hat{q}_1$ and $\hat{q}_1 \equiv q$ imply $\hat{q}_2 \equiv q$ as required and we are done with this case.

Time is a finite resource. The one spent to machine check Lemma 54 would have been best invested into bibliographical research. Months after having implemented the lemma we realised that [3] already had an analogous result for a mechanisation of the π -calculus. Lemma 54 is crucial to prove the Harmony Lemma, which states that τ -transitions coincide with the standard reduction

relation of ACCS. This is out of the scope of our discussion, and we point the interested reader to Lemma 1.4.15 of [87], and to the list of problems presented on the web-page of THE CONCURRENT CALCULI FORMALISATION BENCHMARK.¹⁴

We give a corollary that is useful to prove Lemma 55.

Corollary 5. *For every $p, q \in \text{ACCS}$, and $\alpha \in \text{Act}_\tau$. $p \equiv q$ implies that $p \xrightarrow{\alpha} \cdot \equiv r$ if and only if $q \xrightarrow{\alpha} \cdot \equiv r$.*

Proof. Since $q \equiv p \xrightarrow{\alpha} p' \equiv r$ Lemma 54 implies $q \xrightarrow{\alpha} \cdot \equiv p'$, thus $q \xrightarrow{\alpha} \cdot \equiv r$ by transitivity of \equiv . The other implication follows from the same argument and the symmetry of \equiv .

A consequence of Lemma 53 is that the LTS $\langle \text{ACCS}_{\equiv}, \xrightarrow{\equiv}, \text{Act}_\tau \rangle$ enjoys the axioms in Figure 2, and thus it is FDB. [90, Theorem 4.3] proves it reasoning modulo bisimilarity, while we reason modulo structural equivalence.

Lemma 55. *For every $p \in \text{ACCS}$, and $a \in \mathcal{N}$ the following properties are true,*

- for every $\alpha \in \text{Act}_\tau$. $p \xrightarrow{\bar{a}} \alpha p_3$ implies $p \xrightarrow{\alpha} \bar{a} \cdot \equiv p_3$;
- for every $\alpha \in \text{Act}_\tau$. $\alpha \notin \{\tau, \bar{a}\}$. $p \xrightarrow{\bar{a}} p'$ and $p \xrightarrow{\alpha} p''$ imply that $p'' \xrightarrow{\bar{a}} q$ and $p' \xrightarrow{\alpha} q$ for some q ;
- $p \xrightarrow{\bar{a}} p'$ and $p \xrightarrow{\bar{a}} p''$ imply $p' \equiv p''$;
- $p \xrightarrow{\bar{a}} p' \xrightarrow{a} q$ implies $p \xrightarrow{\tau} \cdot \equiv q$;
- $p \xrightarrow{\bar{a}} p'$ and $p \xrightarrow{\tau} p''$ imply that $p' \xrightarrow{\tau} q$ and $p'' \xrightarrow{\bar{a}} q$; or that $p' \xrightarrow{\tau} p''$.
- for every p' if there exists a \hat{p} such that $p \xrightarrow{\bar{a}} \hat{p}$ and $p' \xrightarrow{\bar{a}} \hat{p}$ then $p \equiv p'$

Proof. To show FEEDBACK we begin via Lemma 53 which proves $p \equiv p' \parallel \boxed{\bar{a}}$. We derive $p' \parallel \boxed{\bar{a}} \xrightarrow{\tau} q'$ and apply Corollary 5 to obtain $p \xrightarrow{\tau} \cdot \equiv q'$.

We prove OUTPUT-TAU. The hypothesis and Lemma 53 imply that $p \equiv p' \parallel \boxed{a}$. Since $p \xrightarrow{\tau} p''$ it must be the case that $p' \xrightarrow{\tau} \hat{p}$ for some \hat{p} , and $p'' = \hat{p} \parallel \boxed{a}$. Let $q = \hat{p}$. We have that $p'' \xrightarrow{a} \hat{p} \parallel 0 \equiv q$.

Processes that enjoy OUTPUT-TAU are called *non-preemptive* in [40, Definition 10].

Each time a process p reduces to a *stable* process p' , it does so by consuming at least part of the mailbox, for instance a multiset of outputs N , thereby arriving in a state q whose inputs cannot interact with what remains of the mailbox, i.e. $M \setminus N$, where M is the original mailbox.

Lemma 56. *For every $M \in \text{MO}$, $p, p' \in \text{ACCS}$, if $p \parallel \boxed{IM} \xrightarrow{\varepsilon} p' \xrightarrow{\tau}$ then there exist an $N \subseteq M$ and some $q \in \text{ACCS}$ such that $p \xrightarrow{\bar{N}} q \xrightarrow{\tau}$, $O(q) \subseteq O(p')$, and $I(q) \# (M \setminus N)$.*

¹⁴ <https://concurrentbenchmark.github.io/>

Proof. By induction on the derivation of $p \parallel \boxed{IIM} \xRightarrow{\varepsilon} p'$. In the base case this is due to [WT-REFL], which ensures that

$$p \parallel \boxed{IIM} = p',$$

from which we obtain $p \parallel \boxed{IIM} \xrightarrow{\tau}$. This ensures that $\overline{I(p)} \# M$. We pick as q and N respectively $p \parallel \boxed{IIM}$ and \emptyset as $p \parallel \boxed{IIM} \xRightarrow{\emptyset} p \parallel \boxed{IIM}$ by reflexivity, and $O(p \parallel \boxed{IIM}) = O(p')$.

In the inductive case the derivation ends with an application of [WT-TAU] and

$$\frac{\begin{array}{c} \vdots \\ p \parallel \boxed{IIM} \xrightarrow{\tau} p' \quad \overline{p' \xRightarrow{\varepsilon} p'} \end{array}}{p \parallel \boxed{IIM} \xRightarrow{\varepsilon} p'}$$

We continue by case analysis on the rule used to infer the transition $p \parallel \boxed{IIM} \xrightarrow{\tau} p'$. As by definition $\boxed{IIM} \xrightarrow{\tau}$, the rule is either [PAR-L], i.e. a τ -transition performed by p , or [COM], i.e. an interaction between p and \boxed{IIM} .

Rule [PAR-L]: In this case $p \xrightarrow{\tau} p''$ for some p'' , thus $p'' \parallel \boxed{IIM} \xRightarrow{\varepsilon} p'$ and the result follows from the inductive hypothesis.

Rule [COM]: The hypothesis of the rule ensure that $p \xrightarrow{a} p''$ and $\boxed{IIM} \xrightarrow{\bar{a}} p'$, and as the process \boxed{IIM} does not perform any input, it must be the case that $a \in \mathcal{N}$, that $\bar{a} \in M$, and that $q \equiv \boxed{II(M \setminus \{\bar{a}\})}$.¹⁵ Note that $p' \equiv p'' \parallel \boxed{II(M \setminus \{\bar{a}\})}$.

The inductive hypothesis ensures that for some $N' \subseteq M \setminus \{\bar{a}\}$ and some $q_3 \in \text{ACCS}$ we have

- (a) $p' \xRightarrow{N'} q_3 \xrightarrow{\tau}$,
- (b) $O(q_3) \subseteq O(p')$, and
- (c) $\overline{I(q_3)} \# ((M \setminus \{\bar{a}\}) \setminus N')$

We conclude by letting $q = q_3$, and $N = \{a\} \uplus N'$. The trace $p \xrightarrow{a} p' \xRightarrow{N'} q_3$ proves that $p \xRightarrow{\{a\} \uplus N'} q_3$, moreover we already know that q_3 is stable. The set inclusion $O(q_3) \subseteq O(p')$ follows from b, and lastly $\overline{I(q)} \# (M \setminus (\{\bar{a}\} \uplus N'))$ is a consequence of $\overline{I(q_3)} \# ((M \setminus \{\bar{a}\}) \setminus N')$ and of $(M \setminus \{\bar{a}\}) \setminus N' = (I \setminus (\{\bar{a}\} \uplus N'))$.

We define the predicate GOOD,

$$\begin{aligned} \text{GOOD}(1) \\ \text{GOOD}(p \parallel q) \text{ if } \text{GOOD}(p) \text{ or } \text{GOOD}(q) \\ \text{GOOD}(p + q) \text{ if } \text{GOOD}(p) \text{ or } \text{GOOD}(q) \end{aligned}$$

This predicate is preserved by structural congruence.

¹⁵ In terms of LTS with mailboxes, $p' = (M \setminus \{\bar{a}\})$.

Lemma 57. *For every $p, q \in ACCS$. $p \equiv q$ and $\text{GOOD}(p)$ imply $\text{GOOD}(q)$.*

Lemma 58. *For every $p, q \in ACCS$. $p \equiv q$ and $p \downarrow_i$ imply $q \downarrow_i$.*

Lemma 59. *For every $p, q \in ACCS$ and $s \in \text{Act}^*$, we have that $p \equiv q$ and $\text{FW}(p) \downarrow s$ imply $\text{FW}(q) \downarrow s$.*

Lemma 60. *For every $p, r, r' \in ACCS$. $r \equiv r'$ and $p \text{MUST}_i r$ then $p \text{MUST}_i r'$.*

Lemma 61. *For every $p, q, r \in ACCS$. $p \equiv q$ and $p \text{MUST}_i r$ then $q \text{MUST}_i r$.*

A typical technique to reason on the LTS of concurrent processes, and so also of client-server systems, is trace zipping: if $p \xrightarrow{s} p'$ and $q \xrightarrow{\bar{s}} q'$, an induction on s ensures that $p \parallel q \Longrightarrow p' \parallel q'$. Zipping together different LTS is slightly more delicate: we can zip weak transitions $\xrightarrow{s}_{\text{fw}}$ together with the co-transitions $\xrightarrow{\bar{s}}$, but possibly moving inside equivalence classes of \equiv instead of performing actual transitions in \longrightarrow .

Lemma 62 (Zipping). *For every $p, q \in ACCS$*

1. *for every $\mu \in \text{Act}$. if $p \xrightarrow{\mu}_{\text{fw}} p'$ and $q \xrightarrow{\bar{\mu}} q'$ then $p \parallel q \xrightarrow{\tau} p' \parallel q'$ or $p \parallel q \equiv p' \parallel q'$;*
2. *for every $s \in \text{Act}^*$. if $p \xrightarrow{s}_{\text{fw}} p'$ and $q \xrightarrow{\bar{s}} q'$ then $p \parallel q \xrightarrow{\varepsilon} \cdot \equiv p' \parallel q'$.*

Obviously, for every $p, q \in A$ and output $a \in \mathcal{N}$ we have

$$p \xrightarrow{\tau}_{\text{fw}} q \text{ if and only if } p \xrightarrow{\tau} q \quad (12)$$

$$p \xrightarrow{\varepsilon}_{\text{fw}} q \text{ if and only if } p \xrightarrow{\varepsilon} q \quad (13)$$

$$p \xrightarrow{\bar{a}}_{\text{fw}} q \text{ if and only if } p \xrightarrow{\bar{a}} q \quad (14)$$

together with the expected properties of finiteness, the first one amounting to the finiteness of the global mailbox in any state:

$$|\{\bar{a} \in \bar{\mathcal{N}} \mid p \xrightarrow{\bar{a}}_{\text{fw}} q\}| \in \mathbb{N} \quad (15)$$

$$\text{For every } \mu \in \text{Act}. |\{q \mid p \xrightarrow{\mu}_{\text{fw}} q\}| \in \mathbb{N} \quad (16)$$

$$|\{q \in A \mid p \xrightarrow{\tau}_{\text{fw}} q\}| \in \mathbb{N} \quad (17)$$

H.2 Client generators and their properties

This subsection is devoted to the study of the semantic properties of the clients produced by the function g . In general these are the properties sufficient to obtain our completeness result.

Lemma 63. *For every $p \in ACCS$ and $s \in \text{Act}^*$, if $p \xrightarrow{\tau}$ then $g(p, s) \xrightarrow{\tau}$.*

$$\begin{aligned}
 g(\varepsilon, r) &= r \\
 g(a.s, r) &= \boxed{\bar{a}} \parallel g(s, r) \\
 g(\bar{a}.s, r) &= a.g(s, r) + \tau.1
 \end{aligned}$$

$$c(s) = g(s, \tau.1) \tag{18}$$

$$ta(s, L) = g(s, h(L)) \text{ where } h(L) = \Pi\{\mu.1 \mid \mu \in L\} \tag{19}$$

Fig. 16. Functions to generate clients.

Proof. By induction on the sequence s . In the base case $s = \varepsilon$. The test generated by g is p , which reduces by hypothesis, and so does $g(\varepsilon, p)$. In the inductive case $s = \alpha.s_2$, and we proceed by case-analysis on α . If α is an output then $g(\alpha.s_2, p) = \bar{\alpha}.(g(s_2, p)) + \tau.1$ which reduces to 1 using the transition rule [SUM-R]. If μ is an input then $g(\alpha.s_2, p) = \boxed{\bar{\alpha}} \parallel g(s_2, p)$ which reduces using the transition rule [PAR-R] and the inductive hypothesis, which ensures that $g(s_2, p)$ reduces.

Lemma 64. For every p and s , if $\neg\text{GOOD}(p)$ then for every s , $\neg\text{GOOD}(g(s, p))$.

Proof. The argument is essentially the same of Lemma 63

Lemma 65. For every $s \in \text{Act}^*$, if $g(s, q) \xrightarrow{\mu} o$ then either

- (a) $q \xrightarrow{\mu} q'$, $s \in \mathcal{N}^*$, and $o = \boxed{\Pi\bar{s}} \parallel q'$, or
- (b) $s = s_1.\boxed{\bar{\mu}}.s_2$ for some $s_1 \in \mathcal{N}^*$ and $s_2 \in \text{Act}^*$, and $o \equiv \boxed{\Pi\bar{s}_1} \parallel g(s_2, q)$, and
 - (i) $\mu \in \mathcal{N}$ implies $g(s, q) \equiv \boxed{\Pi\bar{s}_1} \parallel (\tau.1 + \mu.g(s_2, q))$,
 - (ii) $\mu \in \boxed{\mathcal{N}}$ implies $g(s, q) \equiv \boxed{\Pi\bar{s}_1 \parallel \mu} \parallel (\tau.1 + \mu.g(s_2, q))$.

Proof. The proof is by induction on s .

In the case, $s = \varepsilon$, and hence by definition $g(s, q) = q$. The hypotheses $g(s, q) \xrightarrow{\mu} o$ implies $q \xrightarrow{\mu} o$, and $o \equiv 0 \parallel o \equiv \Pi\varepsilon \parallel o$.

In the inductive case, $s = \nu.s'$. We have two cases, depending on whether ν is an output action or an input action.

Suppose ν is an output. In this case $g(s, q) = \tau.1 + \bar{\nu}.g(s', q)$. The hypothesis $g(s, q) \xrightarrow{\mu} o$ ensures that $\boxed{\bar{\nu}} = \mu$, thus μ is an input action. By letting $s_1 = \varepsilon$ and $s_2 = s'$ we obtain the required

$$g(s, q) = \tau.1 + \mu.g(s_2, q) \equiv \boxed{\Pi\bar{s}_1} \parallel (\tau.1 + \mu.g(s_2, q))$$

and $o \equiv \boxed{\Pi\bar{s}_1} \parallel g(s_2, q)$.

Now suppose that ν is an input action. By definition

$$g(s, q) = \boxed{\bar{\nu}} \parallel g(s', q) \tag{20}$$

and the inductive hypothesis ensures that either

- (1) $q \xrightarrow{\mu} q'$, $s' \in \mathcal{N}^*$, and $o' = \Pi \overline{s'} \parallel q'$, or
 (2) $s' = s'_1 \cdot \overline{\mu} \cdot s'_2$, for some $s'_1 \in \mathcal{N}^*$ and $s_2 \in \text{Act}^*$, and

$$o' \equiv \Pi \overline{s'_1} \parallel g(s'_2, q) \quad (21)$$

and

$$\mu \in \mathcal{N} \text{ implies } g(s', q) \equiv \Pi \overline{s'_1} \parallel (\tau.1 + \mu.g(s'_2, q)) \quad (22)$$

$$\mu \in \overline{\mathcal{N}} \text{ implies } g(s', q) \equiv \Pi \overline{s'_1} \parallel \mu \parallel (\tau.1 + \mu.g(s'_2, q)) \quad (23)$$

The action μ is either an input or an output, and we organise the proof accordingly.

Suppose μ is an input. Since $\overline{\nu}$ is an output, the transition $g(s, q) \xrightarrow{\mu} o$ must be due to a transition $g(s', q) \xrightarrow{\mu} o'$, thus Equation (20) implies

$$o = \overline{\nu} \parallel o' \quad (24)$$

In case 1, then $s' \in \mathcal{N}^*$ and $\nu \in \mathcal{N}$ ensure $s \in \mathcal{N}^*$ and the equality $o \equiv \Pi \overline{s} \parallel q'$ follows from $o' = \Pi \overline{s'} \parallel q'$ and Equation (24).

In case 2, let $s_1 = \nu.s'_1$, $s_2 = s'_2$. Since ν is an input we have $s_1 \in \mathcal{N}^*$. The equalities $s = \nu.s'$ and $s' = s'_1 \cdot \overline{\mu} \cdot s'_2$ imply that $s = s_1 \cdot \overline{\mu} \cdot s_2$. The required $o \equiv \Pi \overline{s_1} \parallel g(s'_2, q)$ follows from $o' \equiv \Pi \overline{s'_1} \parallel g(s'_2, q)$ and Equation (24).

Now we proceed as follows,

$$\begin{aligned} g(s, q) &= \overline{\nu} \parallel g(s', q) && \text{By Equation (20)} \\ &\equiv \overline{\nu} \parallel (\Pi \overline{s'_1} \parallel (\tau.1 + \mu.g(s'_2, q))) && \text{By Equation (22)} \\ &\equiv (\overline{\nu} \parallel \Pi \overline{s'_1}) \parallel (\tau.1 + \mu.g(s'_2, q)) && \text{Associativity} \\ &\equiv \Pi \overline{s_1} \parallel (\tau.1 + \mu.g(s'_2, q)) && \text{Because } s_1 = \nu.s'_1 \\ &\equiv \Pi \overline{s_1} \parallel (\tau.1 + \mu.g(s_2, q)) && \text{Because } s_2 = s'_2 \end{aligned}$$

Now suppose that μ is an output. Then either $\overline{\nu} = \mu$ or $\overline{\nu} \neq \mu$.

In the first case we let $s_1 = \varepsilon$ and $s_2 = s'$. Equation (20) and $\overline{\nu} = \mu$ imply $g(s, q) = \mu \parallel g(s_2, q)$ from which we obtain the required $g(s, q) \equiv \Pi \overline{s_1} \parallel \mu \parallel g(s_2, q)$, and $o \equiv \Pi \overline{s_1} \parallel g(s_2, q)$.

If $\overline{\nu} \neq \mu$ then Equation (20) ensures that the transition $g(s, q) \xrightarrow{\mu} o$ must be due to $g(s', q) \xrightarrow{\mu} o'$ and Equation (24) holds. We use the inductive hypothesis.

If 1 is true, we proceed as already discussed. In case 2 holds, let $s_1 = \nu.s'_1$, we have that $s_1 \in \mathcal{N}^*$. Let $s_2 = s'_2$, now we have that

$$\begin{aligned}
 g(s, q) &= \bar{\nu} \parallel g(s', q) && \text{By Equation (20)} \\
 &\equiv \bar{\nu} \parallel (\Pi \bar{s}'_1 \parallel \mu \parallel g(s'_2, q)) && \text{By Equation (23)} \\
 &\equiv (\bar{\nu} \parallel \Pi \bar{s}'_1) \parallel \mu \parallel g(s'_2, q) && \text{Associativity} \\
 &= \Pi \bar{s}_1 \parallel \mu \parallel g(s'_2, q) && \text{Because } s_1 = \nu.s'_1 \\
 &= \Pi \bar{s}_1 \parallel \mu \parallel g(s_2, q) && \text{Because } s_2 = s'_2
 \end{aligned}$$

Lemma 66. *For every $s \in \text{Act}^*$, if $g(s, q) \xrightarrow{\mu} p$ then either:*

- (a) *there exists q' such that $q \xrightarrow{\mu} q'$, $s \in \mathcal{N}^*$ with $p \equiv \Pi \bar{s} \parallel q'$, or*
- (b) *$s = s_1.\bar{\mu}.s_2$ for some $s_1 \in \mathcal{N}^*$ and $s_2 \in \text{Act}^*$ with $p \equiv g(s_1.s_2, q)$.*

Proof. The proof is by induction over the sequence s .

In the base case $s = \varepsilon$ and we have $g(\varepsilon, q) = q$. We show part (a) and choose $q' = p$. We have $g(\varepsilon, q) = q \xrightarrow{\mu} p$, $p \equiv \Pi \bar{\varepsilon} \parallel q'$ and $\varepsilon \in \mathcal{N}^*$ as required.

In the inductive case $s = \nu.s'$. We proceed by case-analysis on ν . If ν is an input, then $g(\nu.s', q) = \bar{\nu} \parallel g(s', q)$. The hypothesis $g(\nu.s', q) \xrightarrow{\mu} p$ implies that either:

- (i) $\bar{\nu} \xrightarrow{\mu} 0$ with $p = 0 \parallel g(s', q)$ and $\bar{\nu} = \mu$, or
- (ii) $g(s', q) \xrightarrow{\mu} \hat{p}$ with $p = \bar{\nu} \parallel \hat{p}$.

In the first case we show part (b). We choose $s_1 = \varepsilon$, $s_2 = s'$. We have $s = \mu.s' = \bar{\nu}.s'$, $p = 0 \parallel g(s', q) \equiv g(\varepsilon.s', q)$ and $\varepsilon \in \mathcal{N}^*$ as required.

In the second case the inductive hypothesis tells us that either:

- (H-a) *there exists q' such that $q \xrightarrow{\mu} q'$, $s' \in \mathcal{N}^*$ with $\hat{p} \equiv \Pi \bar{s}' \parallel q'$, or*
- (H-b) *$s = s_1.\bar{\mu}.s_2$ for some $s_1 \in \mathcal{N}^*$ and $s_2 \in \text{Act}^*$ with $\hat{p} \equiv g(s_1.s_2, q)$.*

part (a) or part (b) is true.

If part (a) is true then $s' \in \mathcal{N}^*$ and there exists q'' such that $q \xrightarrow{\mu} q''$ with $\hat{p} \equiv \Pi \bar{s}' \parallel q''$. We prove part (a). We choose $q' = q''$ and $s = \nu.s'$. We have $p \equiv \bar{\nu} \parallel \hat{p} \equiv \bar{\nu} \parallel \Pi \bar{s}' \parallel q'' \equiv \Pi \bar{\nu.s}' \parallel q''$ and $\nu.s' \in \mathcal{N}^*$ as required.

If part (b) is true then $s' = s'_1.\bar{\mu}.s'_2$ for some $s'_1 \in \mathcal{N}^*$ and $s'_2 \in \text{Act}^*$ with $\hat{p} \equiv g(s'_1.s'_2, q)$. We prove part (a). We choose $s_1 = \nu.s'_1$ and $s_2 = s'_2$. We have $p \equiv \bar{\nu} \parallel \hat{p} \equiv \bar{\nu} \parallel g(s'_1.s'_2, q) \equiv g(\nu.s'_1.s'_2, q)$ as required.

If ν is an output, then $g(\nu.s', q) = \bar{\nu}.(g(s', q)) + \tau.1$. We prove part (b) and choose $s_1 = \varepsilon$, $s_2 = s'$. The hypothesis $g(\nu.s', q) \xrightarrow{\mu} p$ implies that $\mu = \bar{\nu}$ and $p \equiv g(s', q) \equiv g(\varepsilon.s', q)$ as required.

Lemma 67. *For every $s \in \text{Act}^*$, if $g(s, q) \xrightarrow{\tau} o$ then either*

- (a) *GOOD(o), or*

- (b) $s \in \mathcal{N}^*$, $q \xrightarrow{\tau} q'$, and $o = \Pi \bar{s} \parallel q'$, or
 (c) $s \in \mathcal{N}^*$, $q \xrightarrow{\nu} q'$, and $s = s_1.\nu.s_2$, and $o \equiv \Pi \overline{s_1.s_2} \parallel q'$, or
 (d) $o \equiv g(s_1.s_2.s_3, q)$ where $s = s_1.\mu.s_2.\bar{\mu}.s_3$ with $s_1.\mu.s_2 \in \mathcal{N}^*$.

Proof. By induction on the structure of s .

In the base case $s = \varepsilon$. We prove b. Trivially $s \in \mathcal{N}^*$, and by definition $g(\varepsilon, q) = q$, the hypothesis implies therefore that $q \xrightarrow{\tau} o$. The q' we are after is o itself, for $o \equiv \mathbf{0} \parallel o = \Pi \bar{s} \parallel o$.

In the inductive case $s = \nu.s'$. We proceed by case analysis on whether $\nu \in \overline{\mathcal{N}}$ or $\nu \in \mathcal{N}$.

If ν is an output, by definition $g(s, q) = \tau.1 + \bar{\nu}.gens'q$. Since $\bar{\nu}.g(s', q) \xrightarrow{\tau}$, the silent move $g(s, q) \xrightarrow{\tau} o$ is due to rule [SUM-L], thus $o = 1 + \bar{\nu}.g(s', q)$, and thus GOOD(o). We have proven a.

Suppose now that ν is an input, by definition

$$g(s, q) = \bar{\nu} \parallel g(s', q) \quad (25)$$

The silent move $g(s, q) \xrightarrow{\tau} o$ must have been derived via the rule [COM], or the rule [PAR-R].

If [COM] was employed we know that

$$\frac{\frac{\boxed{\bar{\nu}}}{\bar{\nu} \xrightarrow{\tau} \mathbf{0}} \quad \frac{\vdots}{g(s', q) \xrightarrow{\nu} o'}}{\bar{\nu} \parallel g(s', q) \xrightarrow{\tau} \mathbf{0} \parallel o'}}$$

and thus $o \equiv o'$. Since ν is an input and $g(s', q) \xrightarrow{\nu} o'$, Lemma 65 ensures that either

- (1) $q \xrightarrow{\nu} q'$, $s' \in \mathcal{N}^*$, and $o' = \Pi \bar{s}' \parallel q'$, or
 (2) $s' = s'_1.\boxed{\bar{\nu}}.s'_2$ for some $s'_1 \in \mathcal{N}^*$ and $s'_2 \in \text{Act}^*$, and

$$o' \equiv \boxed{\Pi \bar{s}'_1} \parallel g(s'_2, q) \quad (26)$$

$$g(s, q) \equiv \boxed{\Pi \bar{s}'_1} \parallel (\tau.1 + \nu.g(s'_2, q)) \quad (27)$$

In case 1 we prove part (c). Since ν is an input, $s' \in \mathcal{N}^*$ ensures that $s \in \mathcal{N}^*$. By letting $s_1 = \varepsilon$ and $s_2 = s'$ we obtain $s = s_1.\nu.s_2$. We have to explain why $o \equiv \boxed{\Pi \overline{s_1.s_2}} \parallel q'$. This follows from the definitions of s_1 and s_2 , from $o \equiv \mathbf{0} \parallel o'$ and from $o' \equiv \boxed{\Pi \bar{s}'_1} \parallel q'$.

In case 2 we prove part (d). Let $s_1 = \varepsilon$, $s_2 = s'_1$ and $s_3 = s'_2$.

$$\begin{array}{ll} s' &= \nu s'_1.\bar{\nu}.s'_2 && \text{By inductive hypothesis} \\ \nu.s' &= \nu.s'_1.\bar{\nu}.s'_2 && \\ s &= \nu.s'_1.\bar{\nu}.s'_2 && \text{Because } s = \nu.s' \\ s &= s_1.\nu.s_2.\bar{\nu}.s_3 && \text{By definition} \end{array}$$

and $s_1.\nu.s_2 \in \mathcal{N}^*$ as required. Moreover $o' = \boxed{\Pi s'_1} \parallel g(s'_2, q) = \boxed{\Pi s_2} \parallel g(s_3, q) = g(s_2.s_3, q) = g(s_1.s_2.s_3, q)$ as required. This concludes the argument due to an application of [COM].

If [PAR-R] was employed we know that

$$\begin{array}{c} \vdots \\ \hline g(s', q) \xrightarrow{\tau} o' \\ \hline \bar{\nu} \parallel g(s', q) \xrightarrow{\tau} \bar{\nu} \parallel o' \end{array}$$

thus $g(s', q) \xrightarrow{\tau} o'$ and

$$o = \boxed{\bar{\nu}} \parallel o' \quad (28)$$

Since s' is smaller than s , thanks to $g(s', q) \xrightarrow{\tau} o'$ we apply the inductive hypothesis to obtain either

- (i) $\text{GOOD}(o')$, or
- (ii) $s' \in \mathcal{N}^*$, $q \xrightarrow{\tau} q'$, and $o' = \boxed{\Pi s'}$ $\parallel q'$, or
- (iii) $s' \in \mathcal{N}^*$, $q \xrightarrow{\mu} q'$, and $s' = s'_1.\mu.s'_2$, or
- (iv) $o' \equiv g(s'_1.s'_2.s'_3, q)$ where $s' = s'_1.\mu.s'_2.\bar{\mu}.s'_3$ with $s'_1.\mu.s'_2 \in \mathcal{N}^*$,

If i then Equation (28) implies a. If ii then $s = \nu.s'$ and the assumption that ν is input imply that $s \in \mathcal{N}^*$. Equation (28) and $o' = \boxed{\Pi s'}$ $\parallel q'$ imply that $o = \boxed{\Pi s}$ $\parallel q'$. We have proven b. If iii we prove b, because $s' \in \mathcal{N}^*$ ensures $s \in \mathcal{N}^*$ and $s' = s'_1.\mu.s'_2$ let use prove $s = s_1.\nu.s_2$ by letting $s_1 = \nu.s'_1$ and $s_2 = s'_2$.

If iv we prove d. We have $o' \equiv g(s'_1.s'_2.s'_3, q)$ and $s' = s'_1.\lambda.s'_2.\bar{\lambda}.s'_3$ with $s'_1.\lambda.s'_2 \in \mathcal{N}^*$.

Let $s_1 = \nu.s'_1$, $s_2 = s'_2$ and $s_3 = s'_3$. Since $s'_1.\mu.s'_2 \in \mathcal{N}^*$, we have $s_1.\mu.s_2 \in \mathcal{N}^*$. We also have

$$\begin{array}{ll} s' &= s'_1.\mu.s'_2.\bar{\mu}.s'_3 && \text{By inductive hypothesis} \\ \nu.s' &= \nu.s'_1.\mu.s'_2.\bar{\mu}.s'_3 && \\ s &= \nu.s'_1.\mu.s'_2.\bar{\mu}.s'_3 && \text{Because } s = \nu.s' \\ s &= s_1.\mu.s_2.\bar{\mu}.s_3 && \text{By definition} \end{array}$$

It remains to prove that $o \equiv g(s_1.s_2.s_3, q)$. This is a consequence of Equation (28), of $o' \equiv g(s'_1.s'_2.s'_3, q)$, and of the definitions of s_1, s_2 , and s_3 .

Lemma 68. *For every $s \in \text{Act}^*$, and process q such that or $q \xrightarrow{\tau} q'$ implies $\text{GOOD}(q')$, and for every $\mu \in \mathcal{N}.q \xrightarrow{\mu} q'$ implies $\text{GOOD}(q')$, if $g(s, q) = o_0 \xrightarrow{\tau} o_1 \xrightarrow{\tau} o_2 \xrightarrow{\tau} \dots o_n \xrightarrow{\tau}$ and $n > 0$ then $\text{GOOD}(o_i)$ for some $i \in [1, n]$.*

Proof. Lemma 67 implies that one of the following is true,

- (a) $\text{GOOD}(o_1)$, or
- (b) $s \in \mathcal{N}^*$, $q \xrightarrow{\tau} q'$, and $o_1 = \boxed{\Pi s}$ $\parallel q'$, or

- (c) $s \in \mathcal{N}^*$, $q \xrightarrow{\mu} q'$, and $s = s_1.\mu.s_2$, and $o_1 \equiv II \boxed{s_1.s_2} \parallel q'$, or
 (d) $o_1 \equiv g(s_1.s_2.s_3, q)$ where $s = s_1.\mu.s_2.\bar{\mu}.s_3$ with $s_1.\mu.s_2 \in \mathcal{N}$.

If a we are done. If b or c then $\text{GOOD}(q')$, and thus $\text{GOOD}(o_1)$.

In the base case $\text{len}(s) = 0$, thus d is false. It follows that $\text{GOOD}(o_1)$.

In the inductive case $s = \nu.s'$. We have to discuss only the case in which d is true. The inductive hypothesis ensures that

For every $s' \in \bigcup_{i=0}^{\text{len}(s)-1}$, if $g(s', q) \xrightarrow{\tau} o'_1 \xrightarrow{\tau} o'_2 \xrightarrow{\tau} \dots o'_m \xrightarrow{\tau}$ and $m > 0$ then $\text{GOOD}(o'_j)$ for some j .

Note that $o_1 \xrightarrow{\tau}$ so the reduction sequence $o_1 \Longrightarrow o_n$ cannot be empty, thus $m > 0$. This and $\text{len}(s_1.s_2.s_3) < \text{len}(s)$ let us apply the inductive hypothesis to state that

$$g(s_1.s_2.s_3, q) \xrightarrow{\tau} o'_1 \xrightarrow{\tau} o'_2 \xrightarrow{\tau} \dots o'_m \xrightarrow{\tau} \text{ implies } o'_j \text{ for some } j.$$

We conclude the argument via Lemma 54 and because \equiv preserves success.

Lemma 69. *For every $s \in \text{Act}^*$ and process q , if $g(s, q) \xrightarrow{\tau}$ then*

1. $s \in \mathcal{N}^*$,
2. $q \xrightarrow{\tau}$,
3. $I(q) \cap \bar{s} = \emptyset$,
4. $R(g(s, q)) = \bar{s} \cup R(q)$.

Proof. By induction on s . In the base case $\varepsilon \in \mathcal{N}^*$, and $g(\varepsilon, q) = q$, thus $q \xrightarrow{\tau}$. The last two points follow from this equality and from ε containing no actions.

In the inductive case $s = \mu.s'$. The hypothesis $g(\mu.s', q) \xrightarrow{\tau}$ and the definition of g imply that $g(\mu.s', q) = \bar{\mu} \parallel g(s', q)$, thus $\mu \in \mathcal{N}$. The inductive hypothesis ensures that

1. $s' \in \mathcal{N}^*$,
2. $q \xrightarrow{\tau}$,
3. for every $I(q) \cap \bar{s}' = \emptyset$,
4. for every $R(g(s', q)) = \bar{s}' \cup R(q)$

Since $\bar{\mu} \parallel g(s', q) \xrightarrow{\tau}$ rule [COM] cannot be applied, thus $q \not\xrightarrow{\mu}$, and so $I(q) \cap \bar{s} = \emptyset$. From $R(g(s', q)) = \bar{s}' \cup R(q)$ we obtain $R(g(s, q)) = \bar{s} \cup R(q)$.

Lemma 70. *For every $\mu \in \text{Act}$, s and p , $g(\mu.s, p) \xrightarrow{\bar{\mu}} g(s, p)$.*

Proof. We proceed by case-analysis on μ . If μ is an input then $g(\mu.s, p) = \bar{\mu} \parallel g(s, p)$. We have $\boxed{\bar{\mu}} \parallel g(s, p) \xrightarrow{\bar{\mu}} 0 \parallel g(s, p) \equiv g(s, p)$ as required. If μ is an output then $g(\mu.s, p) = \bar{\mu}.g(s, p) + \tau.1$. We have $\bar{\mu}.g(s, p) + \tau.1 \xrightarrow{\bar{\mu}} g(s, p)$ as required.

Lemma 71. *For every $s \in \text{Act}^*$, and $q \in \text{ACCS}$. $c(s) \xrightarrow{\tau}_{\text{fw}} q$ either*

I. COUNTER-EXAMPLE TO EXISTING COMPLETENESS RESULT

- (a) $\text{GOOD}(q)$, or
(b) there exist b, s_1, s_2 and s_3 with $s_1.b.s_2 \in \mathcal{N}^*$ such that $s = s_1.b.s_2.\bar{b}.s_3$ and $q \equiv c(s_1.s_2.s_3)$.

Proof. The proof is by induction on s .

In the base case $s = \varepsilon$, $c(\varepsilon) = \tau.1$ and then $q = 1$. We prove a with $\text{GOOD}(1)$.

In the inductive case $s = \mu.s'$. We proceed by case-analysis over μ .

If μ is an input then $c(\mu.s') = \bar{\mu} \parallel c(s')$. We continue by case-analysis over the reduction $\bar{\mu} \parallel c(s') \xrightarrow{\tau} q$. It is either due to:

- (i) a communication between $\bar{\mu}$ and $c(s')$ such that $\bar{\mu} \xrightarrow{\mu} 0$ and $c(s') \xrightarrow{\mu} q'$ with $q = 0 \parallel q'$, or
(ii) a reduction of $c(s')$ such that $c(s') \xrightarrow{\tau} q'$ with $q = \bar{\mu} \parallel q'$.

If i is true then Lemma 66 tells us that there exist s'_1 and s'_2 such that $s' = s'_1.\bar{\mu}.s'_2$ and $q' \equiv c(s'_1.s'_2)$ with $s'_1 \in \mathcal{N}^*$. We prove (b). We choose $b = \mu, s_1 = \varepsilon, s_2 = s'_1, s_3 = s'_2$. We show the first requirement by $s = \mu.s' = \mu.s'_1.\bar{\mu}.s'_2 = \varepsilon.\mu.s'_1.\bar{\mu}.s'_2 = s_1.b.s_2.\bar{b}.s_3$. The second requirement is $q = 0 \parallel q' \equiv c(s'_1.s'_2) = c(\varepsilon.s'_1.s'_2) = c(s_1.s_2.s_3)$.

We now consider the case (ii). The inductive hypothesis tells us that either:

1. $\text{GOOD}(q')$, or
2. there exist ι, s'_1, s'_2 and s'_3 with $s'_1.\iota.s'_2 \in \mathcal{N}^*$ such that $s' = s'_1.\iota.s'_2.\bar{\iota}.s'_3$ and $q' \equiv c(s'_1.s'_2.s'_3)$.

If (1) is true then we prove a with $q = \bar{\mu} \parallel q'$ and $\bar{\mu} \parallel q'$ and $\text{GOOD}(q')$. If (2) is true then we prove (b). We choose $b = \iota, s_1 = \mu.s'_1, s_2 = s'_2, s_3 = s'_3$. We show the first requirement with $s = \mu.s' = \mu.s'_1.\iota.s'_2.\bar{\iota}.s'_3 = s_1.b.s_2.\bar{b}.s_3$. The second requirement is $q = \bar{\mu} \parallel q' \equiv \bar{\mu} \parallel c(s'_1.s'_2.s'_3) = c(\mu.s'_1.s'_2.s'_3) = c(s_1.s_2.s_3)$.

If μ is an output then $c(\mu.s') = \bar{\mu}.c(s') + \tau.1$. The hypothesis $c(\mu.s') \xrightarrow{\tau} q$ implies $q = 1$. We prove (a) with $\text{GOOD}(1)$.

I Counter-example to existing completeness result

In this section we recall the definition of the alternative preorder \ll_{ch} by [38], and show that it is not complete with respect to $\bar{\sim}_{\text{MUST}}$, i.e. $\bar{\sim}_{\text{MUST}} \not\subseteq \ll_{\text{ch}}$. We start with some auxiliary definitions.

The predicate $\overset{I}{\rightsquigarrow}$ is defined by the following two rules:

- $p \overset{I}{\rightsquigarrow} p$ if $p \xrightarrow{\tau}$ and $I(p) \cap I = \emptyset$,
- $p \overset{I \uplus \{a\}}{\rightsquigarrow} p''$ if $p \xrightarrow{a} p'$ and $p' \overset{I}{\rightsquigarrow} p''$

The *generalised acceptance set* of a process p after a trace s with respect to a multiset of input actions I is defined by

$$\mathcal{GA}(p, s, I) = \{O(p'') \mid p \xrightarrow{s}_{\text{fw}} p' \overset{I}{\rightsquigarrow} p''\}$$

I. COUNTER-EXAMPLE TO EXISTING COMPLETENESS RESULT

The set of *input multisets* of a process p after a trace s is defined by

$$IM(p, s) = \{\{a_1, \dots, a_n\} \mid a_i \in \mathcal{N}, p \xrightarrow{s} \text{fw} \xrightarrow{a_1} \dots \xrightarrow{a_n}\}$$

The convergence predicate over traces performed by forwarders is denoted \Downarrow_a , and defined as \Downarrow , but over the LTS given in Example 4.

The preorder \ll_{ch} is now defined as follows:

Definition 17 (Alternative preorder \ll_{ch} [38]). Let $p \ll_{\text{ch}} q$ if for every $s \in \text{Act}^*$. $p \Downarrow_a s$ implies

1. $q \Downarrow_a s$,
2. for every $R \in \mathcal{A}(q, s)$ and every $I \in IM(p, s)$ such that $I \cap R = \emptyset$ there exists some $O \in \mathcal{GA}(p, s, I)$ such that $O \setminus \bar{I} \subseteq R$.

We illustrate the three auxiliary definitions using the process $Pierre = b.(\tau.\Omega + c.\bar{d})$ introduced in Example 3. We may infer that

$$Pierre \overset{\{b,c\}}{\rightsquigarrow} \boxed{\bar{d}} \tag{29}$$

thanks to the following derivation tree

$$\frac{\frac{\frac{\boxed{\bar{d}} \overset{\emptyset}{\rightsquigarrow} \boxed{\bar{d}}}{\boxed{\bar{d}} \xrightarrow{\tau} \text{ and } I(\boxed{\bar{d}}) \cap \emptyset = \emptyset}}{\tau.\Omega + c.\bar{d} \overset{\{c\}}{\rightsquigarrow} \boxed{\bar{d}}} \quad \tau.\Omega + c.\bar{d} \xrightarrow{c} \boxed{\bar{d}}}{Pierre \overset{\{b,c\}}{\rightsquigarrow} \boxed{\bar{d}}} \quad Pierre \xrightarrow{b} \tau.\Omega + c.\bar{d}}$$

Let us now consider the generalised acceptance set of $Pierre$ after the trace ε with respect to the multiset $\{b, c\}$. We prove that

$$\mathcal{GA}(Pierre, \varepsilon, \{b, c\}) = \{\{\bar{d}\}\} \tag{30}$$

By definition $\mathcal{GA}(Pierre, \varepsilon, \{b, c\}) = \{O(p'') \mid Pierre \xrightarrow{\varepsilon} \text{fw} p' \overset{\{b,c\}}{\rightsquigarrow} p''\}$. Since $Pierre \xrightarrow{\tau}$, we have

$$\mathcal{GA}(Pierre, \varepsilon, \{b, c\}) = \{O(p'') \mid Pierre \overset{\{b,c\}}{\rightsquigarrow} p''\} \tag{31}$$

Then, thanks to Equation (29) we get $O(\bar{d}) = \{\bar{d}\} \in \mathcal{GA}(Pierre, \varepsilon, \{b, c\})$. We show now that $\{\bar{d}\}$ is the only element of this acceptance set. By (31) above, it is enough to show that $Pierre \overset{\{b,c\}}{\rightsquigarrow} p''$ implies $p'' = \boxed{\bar{d}}$. Observe that

1. $I(Pierre) \cap \{b, c\} \neq \emptyset$,
2. $Pierre \xrightarrow{a} p'$ implies $a = b$, and
3. There are two different states p' such that $Pierre \xrightarrow{b} p'$, but the only one that can do the input c is $p' = \tau.\Omega + c.\bar{d}$.

This implies that the only way to infer $Pierre \stackrel{\{\{b,c\}\}}{\rightsquigarrow} p''$ is via the derivation tree that proves Equation (29) above. Thus $p'' = \bar{d}$.

Counterexample 18 *The alternative preorder \ll_{ch} is not complete for $\sqsubseteq_{\text{MUST}}$, namely $p \sqsubseteq_{\text{MUST}} q$ does not imply $p \ll_{\text{ch}} q$.*

Proof. The cornerstone of the proof is the process $Pierre = b.(\tau.\Omega + c.\bar{d})$ discussed above. In Example 3 we have shown that $Pierre \sqsubseteq_{\text{MUST}} 0$. Here we show that $Pierre \not\ll_{\text{ch}} 0$, because the pair $(Pierre, 0)$ does not satisfy Condition 2 of Definition 17.

Since $Pierre \xrightarrow{\tau} \Omega$, we know that $Pierre \downarrow$, and thus by definition $Pierre \Downarrow_a \varepsilon$. We also have by definition $\mathcal{A}(0, \varepsilon) = \{\emptyset\}$, and $IM(Pierre, \varepsilon) = \{\emptyset, \{b\}, \{b, c\}\}$.

Let us check Condition 2 of Definition 17 for $p = Pierre$ and $q = 0$. Since there is a unique $R \in \mathcal{A}(0, \varepsilon)$, which is \emptyset , and $I \cap \emptyset = \emptyset$ for any I , we only have to check that for every $I \in IM(Pierre, \varepsilon)$ there exists some $O \in \mathcal{GA}(Pierre, \varepsilon, I)$ such that $O \setminus \bar{I} \subseteq \emptyset$.

Let $I = \{b, c\}$. By Equation (30) it must be $O = \{\bar{d}\}$. Since $\{\bar{d}\} \setminus \bar{I} = \{\bar{d}\} \setminus \{\{b, c\}\} = \{\bar{d}\} \not\subseteq \emptyset$, the condition is not satisfied. Thus $Pierre \not\ll_{\text{ch}} 0$.

J Highlights of the Coq mechanisation

J.1 Preliminaries

We begin this section recalling the definition of MUST, which is given in Definition 2. It is noteworthy that the mechanised definition, i.e. `must_extensional`, depends on the typeclass `Sts` (Figure J.1), and *not* the type class `Lts`. This lays bare what stated in Section 1: to define MUST a reduction semantics (i.e. a state transition system), and a predicate `GOOD` over clients suffice.

State Transition Systems The typeclass for state transition systems (`Sts`) is defined as follows, where `A` is the set of states of the `Sts`. It included a notion of stability which is axiomatized and decidable.

```

1 Class Sts (A: Type) := {
2   sts_step: A → A → Prop;
3   sts_state_eqdec: EqDecision A;
4   sts_step_decidable: RelDecision sts_step;
5
6   sts_stable: A → Prop;
7   sts_stable_decidable p : Decision (sts_stable p);
8   sts_stable_spec1 p : ¬ sts_stable p → { q | sts_step p q };
9   sts_stable_spec2 p : { q | sts_step p q } → ¬ sts_stable p;
10  }.

```

Maximal computations A computation is maximal if it is infinite or if its last state is stable. Given a state s , the type `max_exec_from s` contains all the maximal traces that start from s . Note the use of a coinductive type to allow for infinite executions.

```

1 Context `{Sts A}.
2
3 CoInductive max_exec_from: A -> Type :=
4   | MExStop s (Hstable: sts_stable s) : max_exec_from s
5   | MExStep s s' (Hstep: sts_step s s') (η: max_exec_from s') :
6     max_exec_from s.

```

J.2 The must-preorder

Client satisfaction The predicate `GOOD` is defined as any predicate over the states of an LTS that satisfies certain properties: it is preserved by structural congruence, by outputs in both directions (if $p \xrightarrow{a} p'$ then $\text{GOOD}(p) \Leftrightarrow \text{GOOD}(p')$).

It is defined as a typeclass indexed over the type of states and labels, because we expect a practitioner to reason on a single canonical notion of "good" at a time.

```

1
2 Class Good (A L : Type) `{Lts A L, ! LtsEq A L} := {
3   good : A -> Prop;
4   good_preserved_by_eq p q : good p -> p ≡ q -> good q;
5   good_preserved_by_lts_output p q a :
6     p →[ActOut a] q -> good p -> good q;
7   good_preserved_by_lts_output_converse p q a :
8     p →[ActOut a] q -> good q -> good p
9 }.
10

```

Must testing Definition 2: We write $p \text{ MUST } r$ if every maximal computation of $p \parallel r$ is successful.

Given an integer n and a maximal execution η , the function `mex_take_from n` applied to η returns `None` if η is shorter than n and `Some p`, where p is a finite execution corresponding to the first n steps of η .

Then, we define the extensional version of $p \text{ MUST } e$ by stating that, for all maximal executions η starting from (p, e) , there exists an integer n such that the n -th element of η is good. The n th element is obtained by taking the last element of the finite prefix of length n computed using the function above.

```

1 Context `{good : B -> Prop}.
2
3 Fixpoint mex_take_from (n: nat) {x} (η: max_exec_from x) :
4   option (finexec_from x) :=
5   match n with

```

```

6   | 0 => Some $ FExSingl x
7   | S n => match  $\eta$  with
8       | MExStep x Hstable => None
9       | MExStep x x' Hstep  $\eta'$  =>
10          let p' := mex_take_from n  $\eta'$  in
11             ( $\lambda$  p', FExStep x x' (bool_decide_pack _ Hstep) p') <$> p'
12          end
13   end.
14
15 Definition must_extensional (p : A) (e : B) : Prop :=
16   forall  $\eta$  : max_exec_from (p, e), exists n fex,
17     mex_take_from n  $\eta$  = Some fex /\ good (fex_from_last fex).2.

```

The preorder Definition 3 is mechanised in a straightforward way:

```

1 Definition pre_extensional (p : A) (q : R) : Prop :=
2   forall (r : B), must_extensional p r -> must_extensional q r.
3
4 Notation "p  $\sqsubseteq_e$  q" := (pre_extensional p q).

```

J.3 Behavioural characterizations

Labeled Transition Systems An LTS is a typeclass indexed by the type of states and the type of labels. The type of labels must be equipped with decidable equality and be countable, as enforced by the `Label` typeclass. An action $a : \text{Act } L$ is either an internal action τ or an external action: an input or an output of a label in L .

```

1 Class Label (L: Type) := {
2   label_eqdec: EqDecision L;
3   label_countable: Countable L;
4 }.
5
6 Inductive Act (A: Type) := ActExt ( $\mu$ : ExtAct A) |  $\tau$ .
7
8 Class Lts (A L : Type) `{Label L} := {
9   lts_step: A  $\rightarrow$  Act L  $\rightarrow$  A  $\rightarrow$  Prop;
10  lts_state_eqdec: EqDecision A;
11
12  lts_step_decidable a  $\alpha$  b : Decision (lts_step a  $\alpha$  b);
13
14  lts_outputs : A  $\rightarrow$  gset L;
15  lts_outputs_spec1 p1 x p2 :
16    lts_step p1 (ActExt (ActOut x)) p2  $\rightarrow$  x  $\in$  lts_outputs p1;
17  lts_outputs_spec2 p1 x :
18    x  $\in$  lts_outputs p1  $\rightarrow$  {p2 | lts_step p1 (ActExt (ActOut x)) p2};
19
20  lts_stable: A  $\rightarrow$  Act L  $\rightarrow$  Prop;
21  lts_stable_decidable p  $\alpha$  : Decision (lts_stable p  $\alpha$ );

```

```

22   lts_stable_spec1 p α : ¬ lts_stable p α → { q | lts_step p α q };
23   lts_stable_spec2 p α : { q | lts_step p α q } → ¬ lts_stable p α;
24   }.
25
26   Notation "p → q"      := (lts_step p τ q).
27   Notation "p →{ α } q" := (lts_step p α q).
28   Notation "p →[ α ] q" := (lts_step p (ActExt μ) q).

```

An LTS L is cast into an STS by taking only the τ -transitions, as formalised by the following instance, which says that A can be equipped with an STS structure when, together with some labels L , A is equipped with a LTS structure.

```

1   Program Instance sts_of_lts `{Label L} (M: Lts A L): Sts A :=
2     { |
3       sts_step p q := sts_step p τ q;
4       sts_stable s := lts_stable s τ;
5     | }.

```

Weak transitions Let $\Longrightarrow \subseteq A \times \text{Act}^* \times A$ denote the least relation such that:

[wt-refl] $p \xrightarrow{\varepsilon} p'$,
[wt-tau] $p \xrightarrow{s} q$ if $p \xrightarrow{\tau} p'$, and $p' \xrightarrow{s} q$
[wt-mu] $p \xrightarrow{\mu.s} q$ if $p \xrightarrow{\mu} p'$ and $p' \xrightarrow{s} q$.

```

1   Definition trace L := list (ExtAct L).
2
3   Inductive wt : A → trace L → A → Prop :=
4     | wt_nil p : wt p [] p
5     | wt_tau s p q t (l : p → q) (w : wt q s t) : wt p s t
6     | wt_act μ s p q t (l : p →[μ] q) (w : wt q s t) : wt p (μ :: s) t.
7
8   Notation "p ⇒[s] q" := (wt p s q).

```

Product of LTS The characteristic function of the transition relation of the LTS resulting from the parallel composition of two LTS. States of the parallel product of L_1 and L_2 are pairs $(a, b) \in L_1 \times L_2$. The first two cases correspond to unsynchronized steps from either LTS, and the third case corresponds to the LTS taking steps with dual actions. The predicate `act_match l1 l2` states that the two actions are visible and are dual of each other.

```

1   Inductive parallel_step `{M1: Lts A L, M2: Lts B L} :
2     A * B → Act L → A * B → Prop :=
3     | ParLeft l a1 a2 b: a1 -[l]→ a2 → parallel_step (a1, b) l (a2, b)
4     | ParRight l a b1 b2: b1 -[l]→ b2 → parallel_step (a, b1) l (a, b2)
5     | ParSync l1 l2 a1 a2 b1 b2:
6       act_match l1 l2 → a1 -[l1]→ a2 → b1 -[l2]→ b2 →
7       parallel_step (a1, b1) τ (a2, b2)
8     .

```





Fig. 17. Typeclasses to formalise LTSs.

J.4 Typeclasses for LTS

Our basic typeclasses are $\text{Lts} \geq \text{LtsEq} \geq \text{LtsOba}$, where LtsOba is a super-class of both LtsObaFb and LtsObaFw . The class LtsObaFb represents output-buffered agents with feedback, while LtsObaFw represents output-buffered agents with forwarding.

```

1  Class LtsOba (A L : Type) ` {Lts A L, !LtsEq A L} :=
2    MkOba {
3      lts_oba_output_commutativity {p q r a α} :
4        p → [ActOut a] q → q → {α} r →
5        ∃ t, p → {α} t ∧ t → ≡ [ActOut a] r ;
6      lts_oba_output_confluence {p q1 q2 a μ} :
7        μ ≠ ActOut a → p → [ActOut a] q1 → p → [μ] q2 →
8        ∃ r, q1 → [μ] r ∧ q2 → ≡ [ActOut a] r ;
9      lts_oba_output_tau {p q1 q2 a} :
10     p → [ActOut a] q1 → p → q2 →
11     (∃ t, q1 → t ∧ q2 → ≡ [ActOut a] t) ∨ q1 → ≡ [ActIn a] q2 ;
12     lts_oba_output_deter {p1 p2 p3 a} :
13     p1 → [ActOut a] p2 → p1 → [ActOut a] p3 → p2 ≡ p3 ;
14     lts_oba_output_deter_inv {p1 p2 q1 q2} a :
15     p1 → [ActOut a] q1 → p2 → [ActOut a] q2 → q1 ≡ q2 → p1 ≡ p2;
16     (* Multiset of outputs *)
17     lts_oba_mo p : gmultiset L;
18     lts_oba_mo_spec1 p a : a ∈ lts_oba_mo p <-> a ∈ lts_outputs p;
19     lts_oba_mo_spec2 p a q :
20     p → [ActOut a] q → lts_oba_mo p = {[+ a +]} ∪ lts_oba_mo q;
21   }.
22
23  Class LtsObaFb (A L : Type) ` {LtsOba A L} :=
24    MkLtsObaFb {
25     lts_oba_fb_feedback {p1 p2 p3 a} :
26     p1 → [ActOut a] p2 → p2 → [ActIn a] p3 → p1 → ≡ p3
27   }.
28
29  Class LtsObaFw (A L : Type) ` {LtsOba A L} :=
30    MkLtsObaFw {

```

```

31   lts_oba_fw_forward p1 a :
32     ∃ p2, p1 →[ActIn a] p2 ∧ p2 →≡[ActOut a] p1;
33   lts_oba_fw_feedback {p1 p2 p3 a} :
34     p1 →[ActOut a] p2 → p2 →[ActIn a] p3 → p1 →≡ p3 ∨ p1 ≡ p3;
35   }.

```

Termination We write $p \downarrow$ and say that p *converges* if every sequence of τ -transitions performed by p is finite. This is expressed extensionally by the property that all maximal computations starting from p contain a *stable* process, meaning that it is finite.

```

1  Definition terminate (p : A) : Prop :=
2    forall η : max_exec_from p, exists n fex,
3    mex_take_from n η = Some fex /\ lts_stable (fex_from_last fex) τ.

```

Convergence along a trace To define the behavioural characterisation of the preorder, we first define $\Downarrow \subseteq A \times \text{Act}^*$ as the least relation such that,

[cnv-epsilon] $p \Downarrow \varepsilon$ if $p \downarrow$,
[cnv-mu] $p \Downarrow \mu.s$ if $p \downarrow$ and for each p' , $p \xRightarrow{\mu} p'$ implies $p' \Downarrow s$.

This corresponds to the following inductive predicate in Coq:

```

1  Inductive cnv : A -> trace L -> Prop :=
2    | cnv_ext_nil p : terminate p -> cnv p []
3    | cnv_ext_act p μ s :
4      terminate p -> (forall q, p ==>{μ} q -> cnv q s) -> cnv p (μ :: s).
5
6  Notation "p ↓ s" := (cnv p s).

```

J.5 Forwarders

We define a mailbox MO as a multiset of names.

```

1  Definition mb (L : Type) `{Label L} := gmultiset L.

```

Definition 6 and Figure 6 *Lifting of a transition relation to transitions of forwarders*.

```

1  Inductive lts_fw_step {A L : Type} `{Lts A L} :
2    A * mb L -> Act L -> A * mb L -> Prop :=
3    | lts_fw_p p q m α :
4      lts_step p α q -> lts_fw_step (p ▷ m) α (q ▷ m)
5    | lts_fw_out_mb m p a :
6      lts_fw_step (p ▷ {[+ a +]} ⊔ m) (ActExt $ ActOut a) (p ▷ m)
7    | lts_fw_inp_mb m p a :
8      lts_fw_step (p ▷ m) (ActExt $ ActIn a) (p ▷ {[+ a +]} ⊔ m)
9    | lts_fw_com m p a q :
10     lts_step p (ActExt $ ActIn a) q ->
11     lts_fw_step (p ▷ {[+ a +]} ⊔ m) τ (q ▷ m).

```

Definition 14 and Definition 15 *For any LTS \mathcal{L} , two states of $\text{FW}(\mathcal{L})$ are equivalent, denoted $p \triangleright M \doteq q \triangleright N$, if $\text{strip}(p) \simeq \text{strip}(q)$ and $M \uplus \text{mbox}(p) = N \uplus \text{mbox}(q)$.*

```

1 Inductive strip `{Lts A L} : A -> gmultiset L -> A -> Prop :=
2 | strip_nil p : p ~>{\emptyset} p
3 | strip_step p1 p2 p3 a m :
4   p1 ->[ActOut a] p2 -> p2 ~>{m} p3 -> p1 ~>{\{[+ a +]\} \uplus m} p3
5
6 where "p ~>{ m } q" := (strip p m q).
7
8 Definition fw_eq `{LtsOba A L} (p : A * mb L) (q : A * mb L) :=
9   forall (p' q' : A),
10     p.1 ~>{lts_oba_mo p.1} p' ->
11     q.1 ~>{lts_oba_mo q.1} q' ->
12     p' \simeq q' /\ lts_oba_mo p.1 \uplus p.2 = lts_oba_mo q.1 \uplus q.2.
13
14 Infix "\doteq" := fw_eq (at level 70).

```

Lemma 16 *For every \mathcal{L}_A and every $p \triangleright M, q \triangleright N \in A \times MO$, and every $\alpha \in L$, if $p \triangleright M (\doteq \cdot \xrightarrow{\alpha}_{\text{fw}}) q \triangleright N$ then $p \triangleright M (\xrightarrow{\alpha}_{\text{fw}} \cdot \doteq) q' \triangleright N'$.*

```

1 Lemma lts_fw_eq_spec `{LtsObaFB A L} p q t mp mq mt \alpha :
2   p \triangleright mp \doteq t \triangleright mt -> (t \triangleright mt) ->{\alpha} (q \triangleright mq) -> p \triangleright mp ->{\doteq{\alpha}} q \triangleright mq.

```

Lemma 2. *For every LTS $\mathcal{L} \in \text{FDB}$, $\text{FW}(\mathcal{L}) \in \text{FWD}$.*

```

1 Program Instance LtsMB0baFW `{LtsObaFB A L} : LtsObaFW (A * mb L) L.

```

Lemma 3 *For every $\mathcal{L}_A, \mathcal{L}_B \in \text{FDB}$, $p \in A, r \in B$, $p \text{MUST}_i r$ if and only if $\text{FW}(p) \text{MUST}_i r$.*

```

1 Lemma must_iff_must_fw
2   {@LtsObaFB A L IL LA LOA V, @LtsObaFB B L IL LB LOB W,
3    !FiniteLts A L, !Good B L good }
4   (p : A) (e : B) : must p e \leftrightarrow must (p, \emptyset) e.

```

J.6 The Acceptance Set Characterisation

The behavioural characterisation with acceptance sets (Definition 5) is formalised as follows. Note that `lts_outputs`, used in the second part of the definition, is part of the definition of an `Lts`, and produces the finite set of outputs that a process can immediately produce.

```

1 Definition bhv_pre_cond1 `{Lts A L, Lts B L} (p : A) (q : B) :=
2   forall s, p \Downarrow s -> q \Downarrow s.
3
4 Notation "p \preceq_1 q" := (bhv_pre_cond1 p q) (at level 70).
5
6 Definition bhv_pre_cond2 `{Lts A L, Lts B L} (p : A) (q : B) :=
7   forall s q', p \Downarrow s -> q \implies [s] q' -> q' \dashv\rightarrow ->

```

```

8      $\exists p', p \implies [s] p' \wedge p' \not\rightarrow \wedge \text{lts\_outputs } p' \subseteq \text{lts\_outputs } q'$ .
9
10    Notation " $p \preceq_2 q$ " := (bhv_pre_cond2 p q) (at level 70).
11
12    Definition bhv_pre `{@Lts A L HL, @Lts B L HL} (p : A) (q : B) :=
13      p  $\preceq_1$  q  $\wedge$  p  $\preceq_2$  q.
14
15    Notation " $p \preceq q$ " := (bhv_pre p q) (at level 70).

```

Given an LTS that satisfies the right conditions, MUST-equivalence coincides with the behavioural characterisation above on the LTS of forwarders (Theorem 1).

```

1  Section correctness.
2  Context `{LtsObaFB A L, LtsObaFB R L, LtsObaFB B L}.
3  Context `{!FiniteLts A L, !FiniteLts B L, !FiniteLts R L, !Good B L}.
4  (* The LTS can express the tests required for completeness *)
5  Context `{!gen_spec_conv gen_conv, !gen_spec_acc gen_acc}.
6
7  Theorem equivalence_bhv_acc_ctx (p : A) (q : R) :
8    p  $\sqsubseteq_e$  q  $\leftrightarrow$  (p,  $\emptyset$ )  $\preceq$  (q,  $\emptyset$ ).
9  End correctness.

```

J.7 The Must Set characterisation

The behavioural characterisation with must sets (Definition 13) is formalised as follows.

```

1  Definition MUST `{Lts A L} (p : A) (G : gset (ExtAct L)) :=
2    forall p', p  $\implies$  p'  $\rightarrow$  exists  $\mu$  p0,  $\mu \in G \wedge p' \implies \{\mu\} p0$ .
3
4  Definition MUST__s `{FiniteLts A L} (ps : gset A) (G : gset (ExtAct L)) :=
5    forall p, p  $\in$  ps  $\rightarrow$  MUST p G.
6
7  Definition AFTER `{FiniteLts A L} (p : A) (s : trace L) (hcnv : p  $\Downarrow$  s) :=
8    wt_set p s hcnv.
9
10 Definition bhv_pre_ms_cond2
11   `{@FiniteLts A L HL LtsA, @FiniteLts B L HL LtsB} (p : A) (q : B) :=
12   forall s h1 h2 G, MUST__s (AFTER p s h1) G  $\rightarrow$  MUST__s (AFTER q s h2) G.
13
14 Notation " $p \preceq_2 q$ " := (bhv_pre_ms_cond2 p q) (at level 70).
15
16 Definition bhv_pre_ms `{@FiniteLts A L HL LtsA, @FiniteLts B L HL LtsB}
17   (p : A) (q : B) := p  $\preceq_1$  q  $\wedge$  p  $\preceq_2$  q.
18
19 Notation " $p \preceq q$ " := (bhv_pre_ms p q).

```

Lemma 8 Let $\mathcal{L}_A, \mathcal{L}_B \in \text{FDB}$. For every $p \in A$ and $q \in B$ such that $\text{FW}(p) \preceq_{\text{cnv}} \text{FW}(q)$, we have that $\text{FW}(p) \preceq_M \text{FW}(q)$ if and only if $\text{FW}(p) \preceq_{\text{acc}}^{\text{fw}} \text{FW}(q)$.

```

1 Context `{@LtsObaFB A L LL LtsA LtsEqA LtsObaA}.
2 Context `{@LtsObaFB B L LL LtsR LtsEqR LtsObaR}.
3
4 Lemma equivalence_bhv_acc_mst2 (p : A) (q : B) :
5   (p, ∅) ≼1 (q, ∅) → (p, ∅) ≼2 (q, ∅) ↔ (p, ∅) ≼2 (q, ∅).

```

Given an LTS that satisfies the right conditions, MUST-equivalence coincides with the behavioural characterisation above on the LTS of forwarders (Theorem 3).

```

1 Section correctness.
2   Context `{LtsObaFB A L, LtsObaFB R L, LtsObaFB B L}.
3   Context `{!FiniteLts A L, !FiniteLts B L, !FiniteLts R L, !Good B L}.
4   (* The LTS can express the tests required for completeness. *)
5   Context `{!gen_spec_conv gen_conv, !gen_spec_acc gen_acc}.
6
7   Theorem equivalence_bhv_mst_ctx (p : A) (q : R) :
8     p ⊑e q ↔ (p, ∅) ≼ (q, ∅).
9 End correctness.

```

J.8 From extensional to intensional definitions

Proposition 2 *Given a countably branching STS $\langle S, \rightarrow \rangle$, and a decidable predicate Q on S , for all $s \in S$, $\text{ext}_Q(s)$ implies $\text{int}_Q(s)$.*

```

1 Context `{Hsts: Sts A, @CountableSts A Hsts}.
2 Context `{@Bar A Hsts}.
3
4 Theorem extensional_implies_intensional x:
5   extensional_pred x → intensional_pred x.

```

Corollary 2 *For every $p \in A$,*

1. $p \downarrow$ if and only if $p \downarrow_i$,
2. for every r we have that $p \text{MUST } r$ if and only if $p \text{MUST}_i r$.

```

1 Context `{Label L}.
2 Context `{!Lts A L, !FiniteLts A L}.
3
4 Lemma terminate_extensional_iff_terminate (p : A) :
5   terminate_extensional p ↔ terminate p.
6
7 Inductive must_sts `{Sts (A * B), good : B → Prop} (p : A) (e : B) :
8   Prop :=
9   | m_sts_now : good e → must_sts p e
10  | m_sts_step
11    (nh : ¬ good e)
12    (nst : ¬ sts_stable (p, e))
13    (l : forall p' e', sts_step (p, e) (p', e') → must_sts p' e')
14  : must_sts p e

```

```

15 .
16
17 Lemma must_extensional_iff_must_sts
18   `{good : B -> Prop, good_decidable : forall (e : B), Decision (good e)}
19   `{Lts A L, !Lts B L, !LtsEq B L, !Good B L good,
20     !FiniteLts A L, !FiniteLts B L} (p : A) (e : B) :
21     must_extensional p e <-> must_sts p e.

```

Equivalence between the inductive definitions of MUST defined using Sts and MUST defined using Lts.

```

1 Inductive must `{Lts A L, !Lts B L, !LtsEq B L, !Good B L good}
2   (p : A) (e : B) : Prop :=
3   | m_now : good e -> must p e
4   | m_step
5     (nh : ¬ good e)
6     (ex : ∃ t, parallel_step (p, e) τ t)
7     (pt : forall p', p -> p' -> must p' e)
8     (et : forall e', e -> e' -> must p e')
9     (com : forall p' e' μ, e ->[μ] e' -> p ->[co μ] p' -> must p' e')
10    : must p e
11 .
12
13 Lemma must_sts_iff_must `{Lts A L, !Lts B L, !LtsEq B L, !Good B L good}
14   (p : A) (e : B) : must_sts p e <-> must p e.

```

J.9 Completeness

Properties of the functions that generate clients (Table 1).

```

1 Class gen_spec {A L : Type} `{Lts A L, !LtsEq A L, !Good A L good}
2   (gen : list (ExtAct L) -> A) := {
3     gen_spec_ungood : forall s, ¬ good (gen s) ;
4     gen_spec_mu_lts_co μ s : gen (μ :: s) ->[co μ] gen s ;
5     gen_spec_out_lts_tau_ex a s : ∃ e', gen (ActOut a :: s) -> e' ;
6     gen_spec_out_lts_tau_good a s e : gen (ActOut a :: s) -> e -> good e ;
7     gen_spec_out_lts_mu_uniq {e a μ s} :
8     gen (ActOut a :: s) ->[μ] e -> e = gen s /\ μ = ActIn a ;
9   }.
10
11 Class gen_spec_conv {A L : Type} `{Lts A L, !LtsEq A L, !Good A L good}
12   (gen_conv : list (ExtAct L) -> A) := {
13     gen_conv_spec_gen_spec : gen_spec gen_conv ;
14     gen_spec_conv_nil_stable_mu μ : gen_conv [] ->[μ] ;
15     gen_spec_conv_nil_lts_tau_ex : ∃ e', gen_conv [] -> e' ;
16     gen_spec_conv_nil_lts_tau_good e : gen_conv [] -> e -> good e ;
17   }.
18
19 Class gen_spec_acc {A : Type} `{Lts A L, !LtsEq A L, !Good A L good}
20   (gen_acc : gset L -> list (ExtAct L) -> A) := {

```

```

21   gen_acc_spec_gen_spec 0 : gen_spec (gen_acc 0);
22   gen_spec_acc_nil_stable_tau 0 : gen_acc 0 [] ↗;
23   gen_spec_acc_nil_stable_out 0 a : gen_acc 0 [] ↗[ActOut a];
24   gen_spec_acc_nil_mu_inv 0 a e : gen_acc 0 [] →[ActIn a] e -> a ∈ 0;
25   gen_spec_acc_nil_mem_lts_inp 0 a :
26     a ∈ 0 -> ∃ r, gen_acc 0 [] →[ActIn a] r;
27   gen_spec_acc_nil_lts_inp_good μ e' 0 :
28     gen_acc 0 [] →[μ] e' -> good e';
29 }.
```

Proposition 3 For every $\mathcal{L}_A \in \text{FWD}$, $p \in A$, and $s \in \text{Act}^*$ we have that $p \text{MUST}_i tc(s)$ if and only if $p \Downarrow s$.

```

1 Lemma must_iff_cnv
2   `{@LtsObaFW A L IL LA LOA V, @LtsObaFB B L IL LB LOB W,
3     !Good B L good, !gen_spec_conv gen_conv} (p : A) s :
4     must p (gen_conv s) <-> p ↓ s.
5 Proof. split; [eapply cnv_if_must | eapply must_if_cnv]; eauto. Qed.
```

Lemma 14 Let $\mathcal{L}_A \in \text{FWD}$ and $\mathcal{L}_B \in \text{FDB}$. For every $p_1, p_2 \in A$, every $r_1, r_2 \in B$ and name $a \in \mathcal{N}$ such that $p_1 \xrightarrow{\bar{a}} p_2$ and $r_1 \xrightarrow{\bar{a}} r_2$, if $p_1 \text{MUST}_i r_2$ then $p_2 \text{MUST}_i r_1$.

```

1 Lemma must_output_swap_l_fw
2   `{@LtsObaFW A L IL LA LOA V, @LtsObaFB B L IL LB LOB W, !Good B L good}
3   (p1 p2 : A) (e1 e2 : B) (a : L) :
4   p1 →[ActOut a] p2 -> e1 →[ActOut a] e2 -> must p1 e2 -> must p2 e1.
```

Lemma 18. Let $\mathcal{L}_A \in \text{FWD}$. For every $p \in A$, $s \in \text{Act}^*$, and every $L, E \subseteq \mathcal{N}$, if $\bar{L} \in \mathcal{A}_{\text{fw}}(p, s)$ then $p \text{MUST}_i ta(s, E \setminus L)$.

```

1 Lemma not_must_gen_a_without_required_output
2   `{@LtsObaFW A L IL LA LOA V, @LtsObaFB B L IL LB LOB W,
3     !Good B L good, !gen_spec_acc gen_acc} (q q' : A) s 0 :
4   q ⇒[s] q' -> q' ↗ -> ¬ must q (gen_acc (0 \ lts_outputs q') s).
```

Lemma 19 Let $\mathcal{L}_A \in \text{FWD}$. For every $p \in A, s \in \text{Act}^*$, and every finite set $O \subseteq \mathcal{N}$, if $p \Downarrow s$ then either

- (i) $p \text{MUST}_i ta(s, \overline{\bigcup \mathcal{A}_{\text{fw}}(p, s) \setminus O})$, or
- (ii) there exists $\hat{O} \in \mathcal{A}_{\text{fw}}(p, s)$ such that $\hat{O} \subseteq O$.

```

1 Lemma must_gen_a_with_s
2   `{@LtsObaFW A L IL LA LOA V, @LtsObaFB B L IL LB LOB W,
3     !FiniteLts A L, !Good B L good, !gen_spec_acc gen_acc}
4   s (p : A) (hcnv : p ↓ s) 0 :
5   (exists p', p ⇒[s] p' ∧ lts_stable p' τ ∧ lts_outputs p' ⊆ 0)
6   ∨ must p (gen_acc (oas p s hcnv \ 0) s).
```

Lemma 20. For every $\mathcal{L}_A, \mathcal{L}_B \in \text{FWD}$ and servers $p \in A, q \in B$, if $p \overset{\sqsubseteq}{\text{MUST}} q$ then $p \preceq_{\text{AS}} q$.

```

1 Lemma completeness_fw
2   `{@LtsObaFW A L IL LA LOA V, @LtsObaFB B L IL LB LOB W,
3     @LtsObaFW C L IL LC LOC VC, !FiniteLts A L, !FiniteLts C L,
4     !FiniteLts B L, !Good B L good,
5     !gen_spec_conv gen_conv, !gen_spec_acc gen_acc}
6   (p : A) (q : C) : p  $\sqsubseteq$  q  $\rightarrow$  p  $\preceq$  q.

```

Proposition 4. For every $\mathcal{L}_A, \mathcal{L}_B \in \text{FDB}$ and servers $p \in A, q \in B$, if $p \preceq_{\text{MUST}} q$ then $\text{FW}(p) \preceq_{\text{AS}} \text{FW}(q)$.

```

1 Lemma completeness
2   `{@LtsObaFB A L IL LA LOA V, @LtsObaFB B L IL LB LOB W,
3     @LtsObaFB C L IL LC LOC VC,
4     !FiniteLts A L, !FiniteLts B L, !FiniteLts C L, !Good C L good,
5     !gen_spec_conv gen_conv, !gen_spec_acc gen_acc}
6   (p : A) (q : B) : p  $\sqsubseteq$  q  $\rightarrow$  p  $\triangleright \emptyset \preceq$  q  $\triangleright \emptyset$ .

```

J.10 Soundness

Figure 7. Rules to define inductively the predicate MUST_{aux} .

```

1 Inductive mustx
2   `{Lts A L, !FiniteLts A L, !Lts B L, !LtsEq B L, !Good B L good}
3   (ps : gset A) (e : B) : Prop :=
4 | mx_now (hh : good e) : mustx ps e
5 | mx_step
6   (nh :  $\neg$  good e)
7   (ex : forall (p : A), p  $\in$  ps  $\rightarrow$   $\exists$  t, parallel_step (p, e)  $\tau$  t)
8   (pt : forall ps',
9     lts_tau_set_from_pset_spec1 ps ps'  $\rightarrow$  ps'  $\neq$   $\emptyset$   $\rightarrow$ 
10    mustx ps' e)
11  (et : forall (e' : B), e  $\rightarrow$  e'  $\rightarrow$  mustx ps e')
12  (com : forall (e' : B)  $\mu$  (ps' : gset A),
13    lts_step e (ActExt  $\mu$ ) e'  $\rightarrow$ 
14    wt_set_from_pset_spec1 ps [co  $\mu$ ] ps'  $\rightarrow$  ps'  $\neq$   $\emptyset$   $\rightarrow$ 
15    mustx ps' e')
16  : mustx ps e.

```

Lemma 7 For every LTS $\mathcal{L}_A, \mathcal{L}_B$ and every $X \in \mathcal{P}^+(A)$, we have that $X \text{MUST}_{\text{aux}} r$ if and only if for every $p \in X$. $p \text{MUST}_i r$.

```

1 Lemma must_set_iff_must_for_all
2   `{Lts A L, !FiniteLts A L, !Lts B L, !LtsEq B L, !Good B L good}
3   (X : gset A) (e : B) : X  $\neq$   $\emptyset$   $\rightarrow$ 
4   (forall p, p  $\in$  X  $\rightarrow$  must p e)  $\leftrightarrow$  mustx X e.

```

Lifting of the predicates \preceq_{cnv} and \preceq_{acc} to sets of servers.

```

1 Definition bhv_pre_cond1_xx `{FiniteLts P L, FiniteLts Q L}
2   (ps : gset P) (q : Q) := forall s, (forall p, p  $\in$  ps  $\rightarrow$  p  $\Downarrow$  s)  $\rightarrow$  q  $\Downarrow$  s.

```



```

3
4 Notation "ps ≲x1 q" := (bhv_pre_cond1__x ps q) (at level 70).
5
6 Definition bhv_pre_cond2__x
7   `{@FiniteLts P L HL LtsP, @FiniteLts Q L HL LtsQ}
8   (ps : gset P) (q : Q) :=
9   forall s q', q ==>[s] q' -> q' → ->
10    (forall p, p ∈ ps -> p ↓ s) ->
11    exists p, p ∈ ps ∧ exists p',
12    p ==>[s] p' ∧ p' → ∧ lts_outputs p' ⊆ lts_outputs q'.
13
14 Notation "ps ≲x2 q" := (bhv_pre_cond2__x ps q) (at level 70).
15
16 Notation "ps ≲x q" := (bhv_pre_cond1__x ps q ∧ bhv_pre_cond2__x ps q)
17   (at level 70).

```

Lemma 35. For every LTS $\mathcal{L}_A, \mathcal{L}_B$ and servers $p \in A, q \in B$, $p \preceq_{AS} q$ if and only if $\{p\} \preceq_{AS}^{\text{set}} q$.

```

1 Lemma alt_set_singleton_iff
2   `{@FiniteLts P L HL LtsP, @FiniteLts Q L HL LtsQ}
3   (p : P) (q : Q) : p ≲ q <-> {[ p ]} ≲x q.

```

Lemma 36. Let $\mathcal{L}_A, \mathcal{L}_B \in \text{FWD}$. For every set $X \in \mathcal{P}^+(A)$, and $q \in B$, such that $X \preceq_{\text{cnv}}^{\text{set}} q$ then

1. $q \xrightarrow{\tau} q'$ implies $X \preceq_{\text{cnv}}^{\text{set}} q'$,
2. $X \downarrow_i, X \xrightarrow{\mu} X'$ and $q \xrightarrow{\mu} q'$ imply $X' \preceq_{\text{cnv}}^{\text{set}} q'$.

Lemma 37 Let $\mathcal{L}_A, \mathcal{L}_B \in \text{FWD}$. For every $X, X' \in \mathcal{P}^+(A)$ and $q \in B$, such that $X \preceq_{\text{acc}}^{\text{set}} q$, then

1. $q \xrightarrow{\tau} q'$ implies $X \preceq_{\text{acc}}^{\text{set}} q'$,
2. for every $\mu \in \text{Act}$, if $X \downarrow_i$, then for every $q \xrightarrow{\mu} q'$ and set $X \xrightarrow{\mu} X'$ we have $X' \preceq_{\text{acc}}^{\text{set}} q'$.

```

1 Lemma bhvx_preserved_by_tau
2   `{@FiniteLts P L HL LtsP, @FiniteLts Q L HL LtsQ}
3   (ps : gset P) (q q' : Q) : q → q' -> ps ≲x q -> ps ≲x q'.
4
5 Lemma bhvx_preserved_by_mu
6   `{@FiniteLts P L HL LtsP, @FiniteLts Q L HL LtsQ}
7   (ps0 : gset P) (q : Q) μ ps1 q'
8   (htp : forall p, p ∈ ps0 -> terminate p) :
9   q →[μ] q' -> wt_set_from_pset_spec ps0 [μ] ps1 ->
10  ps0 ≲x q -> ps1 ≲x q'.

```

Lemma Let $\mathcal{L}_A, \mathcal{L}_B \in \text{FWD}$ and $\mathcal{L}_C \in \text{FDB}$. For every $X \in \mathcal{P}^+(A)$ and $q \in B$ such that $X \preceq_{AS}^{\text{set}} q$, for every $r \in C$ if $\neg \text{GOOD}(r)$ and $X \text{MUST}_{\text{aux}} r$ then $q \Vdash r \xrightarrow{\tau}$.

```

1 Lemma stability_nbhvleqtwo
2   `{@LtsObaFW P L Lbl LtsP LtsEqP LtsObaP,
3     @LtsObaFW Q L Lbl LtsQ LtsEqQ LtsObaQ,
4     !FiniteLts P L, !FiniteLts Q L, !Lts B L, !LtsEq B L, !Good B L good}
5   (X : gset P) (q : Q) e :
6   ¬ good e → mustx X e → X ≲x2 q → exists t, (q, e) →{τ} t.

```

Lemma *Let $\mathcal{L}_A, \mathcal{L}_B \in \text{FWD}$. For every $X \in \mathcal{P}^+(A)$ and $q, q' \in B$, such that $X \preceq_{\text{AS}}^{\text{set}} q$, then for every $\mu \in \text{Act}$, if $X \Downarrow \mu$ and $q \xrightarrow{\mu} q'$ then $X \xrightarrow{\mu}$.*

```

1 Lemma bhvx_mu_ex `{@FiniteLts P L HL LtsP, @FiniteLts Q L HL LtsQ}
2   (ps : gset P) (q q' : Q) μ
3   : ps ≲x q → (forall p, p ∈ ps → p ↓ [μ]) →
4     q →[μ] q' → exists p', wt_set_from_pset_spec1 ps [μ] {[ p' ]}.

```

Lemma *For every $\mathcal{L}_A \in \text{FWD}$, $\mathcal{L}_B \in \text{FDB}$, every set of processes $X \in \mathcal{P}^+(A)$, every $r \in B$, and every $\mu \in \text{Act}$, if $X \text{MUST}_{\text{aux}} r$, $\neg \text{GOOD}(r)$ and $r \xrightarrow{\mu}$ then $X \Downarrow \bar{\mu}$.*

```

1 Lemma ungood_acnv_mu `{LtsOba A L, !FiniteLts A L, !Lts B L, !LtsEq B L,
2   !Good B L good} ps e e' : mustx ps e →
3   forall μ p, p ∈ ps → e →[co μ] e' → ¬ good e → p ↓ [μ].

```

Lemma 38. *Let $\mathcal{L}_A, \mathcal{L}_B \in \text{FWD}$ and $\mathcal{L}_C \in \text{FDB}$. For every set of processes $X \in \mathcal{P}^+(A)$, server $q \in B$ and client $r \in C$, if $X \text{MUST}_{\text{aux}} r$ and $X \preceq_{\text{AS}}^{\text{set}} q$ then $q \text{MUST}_i r$.*

```

1 Lemma soundnessx `{
2   @LtsObaFW A L Lbl LtsA LtsEqA LtsObaA,
3   @LtsObaFW C L Lbl LtsC LtsEqC LtsObaC,
4   @LtsObaFB B L Lbl LtsB LtsEqB LtsObaB,
5   !FiniteLts A L, !FiniteLts C L, !FiniteLts B L, !Good B L good}
6   (ps : gset A) (e : B) : mustx ps e →
7   forall (q : C), ps ≲x q → must q e.

```

Proposition 5. *For every $\mathcal{L}_A, \mathcal{L}_B \in \text{FDB}$ and servers $p \in A, q \in B$, if $\text{FW}(p) \preceq_{\text{AS}} \text{FW}(q)$ then $p \sqsubseteq_{\text{MUST}} q$.*

```

1 Lemma soundness
2   `{@LtsObaFB A L IL LA LOA V, @LtsObaFB C L IL LC LOC T,
3     @LtsObaFB B L IL LB LOB W,
4     !FiniteLts A L, !FiniteLts C L, !FiniteLts B L, !Good B L good }
5   (p : A) (q : C) : p ▷ ∅ ≲ q ▷ ∅ → p ⊆ q.

```

Corollary 6. *Let $\mathcal{L}_A, \mathcal{L}_B \in \text{FDB}$. For every $p \in A$ and $q \in B$, we have that $p \sqsubseteq_{\text{MUST}} q$ if and only if $p \leq_{\text{fail}} q$.*

```

1
2 Section failure.
3

```

```

4  Definition Failure `{FiniteLts A L} (p : A)
5    (s : trace L) (G : gset (ExtAct L)) :=
6    p ↓ s -> exists p', p ==>[s] p' /\
7      forall μ, μ ∈ G -> ¬ exists p0, p' ==>[μ] p0.
8
9  Definition fail_pre_ms_cond2
10   `{@FiniteLts A L HL LtsA, @FiniteLts B L HL LtsB}
11   (p : A) (q : B) := forall s G, Failure q s G -> Failure p s G.
12
13  Definition fail_pre_ms
14   `{@FiniteLts A L HL LtsA, @FiniteLts B L HL LtsB} (p : A) (q : B) :=
15   p ≼1 q /\ fail_pre_ms_cond2 p q.
16
17  Context `{LL : Label L}.
18  Context `{LtsA : !Lts A L, !FiniteLts A L}.
19  Context `{LtsR : !Lts R L, !FiniteLts R L}.
20
21  Context `{@LtsObaFB A L LL LtsA LtsEqA LtsObaA}.
22  Context `{@LtsObaFB R L LL LtsR LtsEqR LtsObaR}.
23
24  Theorem equivalence_pre_failure_must_set (p : A) (q : R) :
25   (p ▷ ∅) ≼ (q ▷ ∅) <-> (p ▷ ∅) < (q ▷ ∅).
26
27  End failure.

```

J.11 Coinductive definition

The coinductive preorder is defined in Coq using a coinductive predicate as follows:

```

1  CoInductive copre `{@FiniteLts A L HL LtsP, @FiniteLts B L HL LtsQ}
2    (ps : gset A) (q : B) : Prop := {
3    c_tau q' : q → q' -> copre ps q'
4    ; c_now : (forall p, p ∈ ps -> p ↓i) -> q ⇏ ->
5      exists p p', p ∈ ps /\ p ==> p' /\ p' ⇏ /\ lts_outputs p' ⊆ lts_outputs q
6    ; c_step : forall μ q' ps', (forall p, p ∈ ps -> p ↓ [μ]) ->
7      q → [μ] q' -> wt_set_from_pset_spec ps [μ] ps' -> copre ps' q'
8    ; c_cnv : (forall p, p ∈ ps -> p ↓i) -> q ↓i
9    }.

```

The soundness and completeness of the preorder is expressed in Theorem 2, which is formalised as follows in our development.

```

1  Theorem eqx `{@FiniteLts A L HL LtsP, @FiniteLts B L HL LtsQ} (X : gset A) (q : B) :
2    X ≼x q <-> X ≤ q.

```

J.12 The preorder on traces in normal forms

A trace in normal form is simply a list of pairs of multisets, representing inputs and outputs:

1 **Definition** ntrace L `{Label L} : Type := list (gmultiset L * gmultiset L).

Given such a trace in normal form, we can linearize it to a trace by choosing an arbitrary order for the elements of the multisets. This is achieved using the `elements` that maps a multiset to a list that contains its elements in an arbitrary order.

```
1 Fixpoint linearize `{Label L} (nt : ntrace L) : trace L :=
2   match nt with
3     | [] => []
4     | (mi, mo) :: nt' =>
5       let inputs := map ActIn (elements mi) in
6       let outputs := map ActOut (elements mo) in
7       inputs ++ outputs ++ linearize nt'
8   end.
```

Using this operation, we can define the preorder as follows:

```
1 Definition bhv_lin_pre_cond1 `{Lts P L, Lts Q L} (p : P) (q : Q) :=
2   forall s, p ↓ linearize s -> q ↓ linearize s.
```

```
3
4 Notation "p ≈1 q" := (bhv_lin_pre_cond1 p q) (at level 70).
```

```
5
6 Definition bhv_lin_pre_cond2 `{@Lts P L HL, @Lts Q L HL} (p : P) (q : Q) :=
7   forall nt q',
8     p ↓ linearize nt -> q ⇒ [linearize nt] q' -> q' ↗ ->
9     ∃ p', p ⇒ [linearize nt] p' ∧ p' ↗ ∧ lts_outputs p' ⊆ lts_outputs q'.
```

```
10
11 Notation "p ≈2 q" := (bhv_lin_pre_cond2 p q) (at level 70).
```

```
12
13 Definition bhv_lin_pre `{@Lts P L HL, @Lts Q L HL} (p : P) (q : Q) := p ≈1 q ∧ p ≈2 q.
```

```
14
15 Notation "p ≈ q" := (bhv_lin_pre p q) (at level 70).
```

Finally, the theorem stating that this preorder characterises the MUST-preorder is formulated as follows in our development:

```
1 Lemma asyn_iff_bhv
2   `{@LtsObaFW P L IL LA LOA V, @LtsObaFW Q L IL LB LOB W,
3    !FiniteLts Q L, !FiniteLts P L} :
4   forall (p : P) (q : Q), p ≈ q <-> p ≡ q.
```

K Mapping of results from the paper to the Coq code

Paper	Coq File	Coq name
Figure 2	TransitionSystems.v	Class LtsOba
Figure 3	TransitionSystems.v	Class Sts, ExtAct, Act, Label, Lts
Definition 2	Equivalence.v	must_extensional
Definition 3	Equivalence.v	pre_extensional
Equation (5)	ACCSInstance.v	proc
Figure 5	TransitionSystems.v	LtsEq
Definition 2	MustEx.v	must_extensional
Definition 4	TransitionSystems.v	max_exec_from
$p \xrightarrow{s} p'$	TransitionSystems.v	wt
$p \downarrow$	Equivalence.v	terminate_extensional
$p \Downarrow s$	TransitionSystems.v	cnv
Lemma 22	TransitionSystems.v	cnv_iff_prefix_terminate
Lemma 23	TransitionSystems.v	stable_tau_preserved_by_wt_output, stable_tau_input_preserved_by_wt_output
Lemma 21	Must.v	ungood_preserved_by_wt_output
Equation (7)	TransitionSystems.v	Class LtsObaFW
Definition 5	Must.v	bhv_pre
Figure 6	TransitionSystems.v	lts_fw_step
Definition 6	TransitionSystems.v	MbLts
Definition 14	TransitionSystems.v	strip
Definition 15	TransitionSystems.v	fw_eq
Lemma 16	TransitionSystems.v	lts_fw_eq_spec
Lemma 2	TransitionSystems.v	Instance LtsMBObaFW
Lemma 3	Lift.v	must_iff_must_fw
Lemma 3	Lift.v	lift_fw_ctx_pre
Theorem 1	Equivalence.v	equivalence_bhv_acc_ctx
Definition 13	Must.v	bhv_pre_ms
Lemma 8	Must.v	equivalence_bhv_acc_mst
Theorem 3	Must.v	equivalence_bhv_mst_ctx
Lemma 1	ACCSInstance.v	ACCS_ltsObaFB
Corollary 1	ACCSInstance.v	bhv_iff_ctx_ACCS
Proposition 2	Bar.v	extensional_implies_intensional
$p \downarrow_i$	TransitionSystems.v	terminate
$p \text{ MUST}_i q$	Must.v	must_sts
Corollary 2	Equivalence.v	terminate_extensional_iff_terminate
Table 1	Completeness.v	Class gen_spec, gen_spec_conv, gen_spec_acc
Proposition 3	Completeness.v	must_iff_cnv
Lemma 14	Lift.v	must_output_swap_l_fw
Lemma 18	Completeness.v	not_must_gen_a_without_required_output
Lemma 19	Completeness.v	must_gen_a_with_s
Lemma 20	Completeness.v	completeness_fw
Proposition 4	Completeness.v	completeness
Lemma 7	Soundness.v	must_set_iff_must_for_all
Figure 7	Soundness.v	mustx
$X \xrightarrow{\text{set}}_{\text{cnv}} q$ and $X \xrightarrow{\text{set}}_{\text{acc}} q$	Soundness.v	bhv_pre_cond1 _x and bhv_pre_cond2 _x
Lemma 35	Soundness.v	must_set_iff_must
Lemma 36, Lemma 37	Soundness.v	bhvx_preserved_by_tau, bhvx_preserved_by_mu
Lemma 38	Soundness.v	soundnessx
Proposition 5	Soundness.v	soundness
Corollary 4	Normalisation.v	asyn_iff_bhv
Theorem 2	Coin.v	eq_ctx