

Full-Abstraction for Must Testing Preorders (Extended Abstract)

Giovanni Bernardi¹(✉) and Adrian Francalanza²

¹ IRIF, Université Paris-Diderot, Paris, France
gio@irif.fr

² University of Malta, Msida, Malta
adrian.francalanza@um.edu.mt

Abstract. The client **must** preorder relates tests (clients) instead of processes (servers). The existing characterisation of this preorder is unsatisfactory for it relies on the notion of *usable* clients which, in turn, are defined using an existential quantification over the servers that ensure client satisfaction. In this paper we characterise the set of usable clients for finite-branching LTSs, and give a sound and complete decision procedure for it. We also provide a novel coinductive characterisation of the client preorder, which we use to argue that the preorder is decidable, thus positively answering the question opened in [3,6].

1 Introduction

The standard testing theory of De Nicola–Hennessy [12,15] has recently been employed to provide theoretical foundations for web-services [9,25] (where processes denote servers). To better fit that setting, in [6] this theory has been enriched with preorders for clients (tests) and peers (where both interacting parties mutually satisfy one another). Client preorders also tie testing theory with session type theory, as is outlined in [2]: they are instrumental in defining semantic models of the Gay & Hole subtyping [14] for first-order session types [3, Theorem 6.3.4] and [5, Theorem 5.2].

The testing preorders for clients and peers are *contextual* preorders, defined by comparing the capacity of either being satisfied by servers or the capacity of peers to mutually satisfy one another. This paper focuses on the client preorder due to the **must** testing relation [12,15]: a client r_2 is better than a client r_1 , denoted $r_1 \sqsubset_{\text{ct}} r_2$, whenever *every* server p that **must** pass r_1 also **must** pass r_2 . Although this definition is easy to understand, it suffers from the endemic universal quantification over contexts (servers) and, by itself, does not give any effective proof method to determine pairs in the preorder. To solve this problem, contextual preorders usually come equipped with *behavioural characterisations* that avoid universal context quantification thereby facilitating reasoning. In [6] the authors develop such characterisations for the client and the peer **must** preorders; these preorders are however *not* fully-abstract, for they are defined modulo *usable* clients, *i.e.*, clients that are satisfied by *some* server.




Fig. 1. LTS depictions of the behaviours described in Eq. (1)

Usability is a pivotal notion that appears frequently in the literature of process calculi and web-service foundations, *cf.* viability in [18,26] and controllability in [8,24], and has already been studied, albeit for restricted or different settings, in [6,7,18,25,26]. In general though, the characterisation of usability is problematic, for solving it requires finding the conditions under which one can either (a) construct a server p that satisfies a given client, or (b) show that every p does *not* satisfy a given client. Whereas proving (b) is complicated by the universal quantification over *all* servers, the proof of (a) is complicated by the non-deterministic behaviour of clients. In particular, the approach in (a) is complicated because client usability is *not* compositional. For instance consider the following clients, whose behaviours are depicted in Fig. 1:

$$r_1 = c.(a.1 + b.0) \quad \text{and} \quad r_2 = c.(a.0 + b.1) \tag{1}$$

where 1 denotes satisfaction (success). Both clients are usable, since r_1 is satisfied by the server $\bar{c}.\bar{a}.0$, and r_2 is satisfied by server $\bar{c}.\bar{b}.0$. However, their composition $r_1 + r_2$ is *not* a usable client, *i.e.*, $p \text{ must } r_1 + r_2$ for every p ; intuitively, this is because r_1 and r_2 impose opposite constraints on the processes that pass one or the other (*e.g.*, $\bar{c}.\bar{a}.0 + \bar{c}.\bar{b}.0$ does not satisfy $r_1 + r_2$). A compositional analysis is even more unwieldy for recursive tests. For instance, the client $\mu x.(c.(a.1 + b.x) + c.(a.0 + b.1))$ is not usable because of the non-determinism analogous to $r_1 + r_2$, and the unsuccessful computations along the infinite trace of interactions $(c.b)^*$; this argument works because infinite unsuccessful computations are catastrophic *wrt.* **must** testing.

This paper presents a sound and complete characterisation for usable clients with finite-branching LTSs. Through the results of [6] — in particular, the equivalence of usability for clients and peers stated on [6, p. 11] — our characterisation directly yields a fully-abstract characterisation for the **must** preorder for clients and peers. We go a step further and use this characterisation to develop a novel *coinductive* and fully-abstract characterisation of \sqsubseteq_{ct} , which we find easier to use than the one of [6] when proving inequalities involving recursive clients. This coinductive characterisation turns out to be informed by our study on usability, and differs from related coinductive characterisations for the server preorder [18,25] in a number of respects. Finally, our inductive definition for

usable clients also provides deeper insights into the original client preorder of [6]: we show that limiting contexts to servers offering only *finite* interactions preserves the discriminating power of the original preorder. Our contributions are:

- a fully-abstract characterisation of usable clients, Theorem 2;
- a coinductive, fully-abstract characterisation of the client preorder \sqsubseteq_{ct} , Theorem 5;
- a contextual preorder $\sqsubseteq_{\text{ct}}^f$ that is equivalent to \sqsubseteq_{ct} but relies only on non-recursive contexts Theorem 6;
- decidability results for usable clients and the client preorder, Theorem 7.

The solutions devised here addressing client usability are directly relevant to controllability issues in service-oriented architectures [21,30]. Our techniques may also be extended beyond this remit. The ever growing sizes of test suites, together with the ubiquitous reliance on testing for the increasing quality-assurance requirements in software systems, has directed the attention to non-deterministic (or *flaky*) tests. Such tests arise frequently in practice and their impact on software development has been the subject of various studies [19,20,22]. By some measures, $\approx 4.56\%$ of test failures of the TAP (Test Anything Protocol) system at Google are caused by flaky tests [19]. We believe that our concepts, models and procedures can be extended to such testing methodologies to analyse detrimental non-deterministic behaviour arising in test suites, thereby reducing the gap between empirical practices and theory.

Structure of the paper: Sect. 2 outlines the preliminaries for client must testing. Section 3 tackles client usability and gives a fully-abstract definition for it. Section 4 uses this result to give a coinductive characterisation for client preorders. In Sect. 5 we present expressiveness results for servers with finite interactions together with decidability results for client usability and the client testing preorder. Section 6 concludes.

2 Preliminaries

Let $a, b, c, \dots \in \text{Act}$ be a set of actions, and let τ, \checkmark be two distinct actions *not* in Act ; the first denotes *internal* unobservable activity whereas the second is used to *report success* of an experiment. To emphasise their distinctness, we use $\alpha \in \text{Act}_\tau$ to denote $\text{Act} \cup \{\tau\}$, and similarly for $\lambda \in \text{Act}_{\checkmark}$. We assume Act has an involution function, with \bar{a} being the complement to a .

A *labelled transition system*, LTS, consists of a triple $\langle \text{Proc}, \text{Act}_{\tau\checkmark}, \longrightarrow \rangle$, where Proc is a set of processes and $\longrightarrow \subseteq (\text{Proc} \times \text{Act}_{\tau\checkmark} \times \text{Proc})$ is a transition relation between processes decorated with labels drawn from the set $\text{Act}_{\tau\checkmark}$; we write $p \xrightarrow{\lambda} q$ in lieu of $(p, \lambda, q) \in \longrightarrow$. An LTS is *finite-branching* if for all $p \in \text{Proc}$ and for all $\lambda \in \text{Act}_{\tau\checkmark}$, the set $\{q \mid p \xrightarrow{\lambda} q\}$ is finite. For $s \in (\text{Act}_{\checkmark})^*$ we also have the standard weak transitions, $p \xRightarrow{s} q$, defined by *ignoring* the occurrences of τ s.

Syntax	$p, q, r, o \in \text{CCS}^{\mu} ::= 0 \mid 1 \mid \alpha.p \mid p + q \mid \mu x.p \mid x$		
Semantics			
$\frac{}{1 \xrightarrow{\checkmark} 0}$ (A-OK)	$\frac{}{\alpha.p \xrightarrow{\alpha} p}$ (A-PRE)	$\frac{}{\mu x.p \xrightarrow{\tau} p\{\mu x.p/x\}}$ (A-UNFOLD)	
$\frac{p \xrightarrow{\lambda} p'}{p + q \xrightarrow{\lambda} p'}$ (R-EXT-L)	$\frac{q \xrightarrow{\lambda} q'}{p + q \xrightarrow{\lambda} q'}$ (R-EXT-R)		
Contract Composition Semantics			
$\frac{p \xrightarrow{\lambda} p'}{p \parallel r \xrightarrow{\lambda} p' \parallel r}$ (P-SRV)	$\frac{r \xrightarrow{\lambda} r'}{p \parallel r \xrightarrow{\lambda} p \parallel r'}$ (P-CLI)	$\frac{p \xrightarrow{a} p' \quad r \xrightarrow{\bar{a}} r'}{p \parallel r \xrightarrow{\tau} p' \parallel r'}$ (P-SYN)	

Fig. 2. Syntax and Semantics of recursive CCS^{μ} with 1.

We limit ourselves to finite-branching LTSs. Whenever sufficient, we describe such LTSs using a version of CCS with recursion [23] and augmented with a *success* operator, denoted as 1. The syntax of this language is depicted in Fig. 2 and assumes a denumerable set of variables $x, y, z \dots \in \text{Var}$. For finite I , we use the notation $\sum_{i \in I} p_i$ to denote the *resp.* sequence of summations $p_1 + \dots + p_n$ where $I = 1..n$. Similarly, when I is a non-empty set, we define $\bigoplus_{i \in I} p_i = \sum_{i \in I} \tau.p_i$ to represent process *internal* choice. The transition relation $p \xrightarrow{\lambda} q$ between terms of the language is the least one determined by the (standard) rules in Fig. 2. As usual, $\mu x.p$ binds x in p and we identify terms up to alpha conversion of bound variables. The operation $p\{\mu x.p/x\}$ denotes the unfolding of the recursive process $\mu x.p$, by substituting the term $\mu x.p$ for the free occurrences of the variable x in p .

To model the interactions taking place between the server and the client contracts, we use the standard binary composition of contracts, $p \parallel r$, whose operational semantics is given in Fig. 2. A *computation* consists of sequence of τ actions of the form

$$p \parallel r = p_0 \parallel r_0 \xrightarrow{\tau} p_1 \parallel r_1 \xrightarrow{\tau} \dots \xrightarrow{\tau} p_k \parallel r_k \xrightarrow{\tau} \dots \quad (2)$$

It is *maximal* if it is infinite, or whenever $p_n \parallel r_n$ is the last state then $p_n \parallel r_n \not\xrightarrow{\tau}$. We say (2) is *client-successful* if there exists some $k \geq 0$ such that $r_k \xrightarrow{\checkmark}$.

Definition 1 (Client Testing preorder [6]). We write p *must* r if every maximal computation from $p \parallel r$ is *client-successful*, and write $r_1 \sqsubseteq_{\text{ct}} r_2$ if, for every p , p *must* r_1 implies p *must* r_2 . ■

Although intuitive, the universal quantification on servers in Definition 1 complicates reasoning about \sqsubseteq_{ct} . One way of surmounting this is by defining alternative characterisations for \sqsubseteq_{ct} of Definition 1, that come equipped with practical proof methods.

2.1 Characterising the Client Preorder

In [6, Definition 3.10, p. 9], an alternative characterisation for the preorder \sqsubseteq_{ct} is given and proven to be sound and complete. We recall this characterisation, restating the *resp.* notation. The alternative characterisation relies on *unsuccessful* traces: $r \xrightarrow{s} r'$ means that r may weakly perform the trace of external actions s reaching state r' *without* passing through *any* successful state; in particular neither r nor r' are successful. Formally, $r \xrightarrow{s} r'$ is the least relation satisfying (a) $r \not\check{\rightarrow}$ implies $r \xrightarrow{\varepsilon} r$, and (b) if $r'' \xrightarrow{s} r'$ and $r \not\check{\rightarrow}$ then (i) $r \xrightarrow{a} r''$ implies $r \xrightarrow{as} r'$, and (ii) $r \xrightarrow{\tau} r''$ implies $r \xrightarrow{s} r'$. The *unsuccessful* acceptance set of r after s , are defined as

$$\text{Acc}_{\not\check{\rightarrow}}(r, s) = \{ S(r') \mid r \xrightarrow{s} r' \not\check{\rightarrow} \} \quad (3)$$

where $S(r) = \{ a \in \text{Act} \mid r \xrightarrow{a} \}$ denotes the strong actions of r . Intuitively, for the client r , the set $\text{Acc}_{\not\check{\rightarrow}}(r, s)$ records all the actions that lead r out of *potentially deadlocked* (i.e. stable) states that it reaches performing *unsuccessfully* the trace s . It turns out that these abstractions are fundamental to characterise must-testing preorders and also compliance preorders [3, 6, 25]. In the sequel, we shall also use $r \xrightarrow{\alpha} r'$ whenever $r \xrightarrow{\alpha} r'$, $r \not\check{\rightarrow}$ and $r' \not\check{\rightarrow}$ hold.

Example 1. For client $r_3 = \tau.(1 + \tau.0)$ we have $\text{Acc}_{\not\check{\rightarrow}}(r_3, \varepsilon) = \emptyset$, but for $r'_3 = r_3 + \tau.0$ we have $\text{Acc}_{\not\check{\rightarrow}}(r'_3, \varepsilon) = \{ \emptyset \}$. We also have $\text{Acc}_{\not\check{\rightarrow}}(r''_3, \varepsilon) = \emptyset$ for $r''_3 = r_3 + \mu x.x$. ■

Note that, whenever $\text{Acc}_{\not\check{\rightarrow}}(r, s) = \emptyset$, then any sequence of moves with trace s from r to a *stable* reduct r' must pass through a successful state, for otherwise we would have $S(r') \in \text{Acc}_{\not\check{\rightarrow}}(r, s)$ for some r' .

Definition 2 (Usable Clients). $\mathcal{U} = \{ r \mid \text{there exists } p.p \text{ must } r \}$. ■

Example 2. Recall clients r_1 and r_2 from (1) in Sect. 1. We show that despite being individually usable, the sum of these clients is not: $p \not\text{m}\mu\text{st } r_1 + r_2$ for every p . Fix a process p . If p does not offer an interaction on \bar{c} , then, plainly, $p \text{ m}\mu\text{st } r_1 + r_2$. Suppose that $p \xrightarrow{\bar{c}} p'$; to prove $p \not\text{m}\mu\text{st } r_1 + r_2$, it suffices to show that there exists a client r reached by $r_1 + r_2$ by performing action c (i.e., $r \in \{ a.1 + b.0, a.0 + b.1 \}$) such that $p' \not\text{m}\mu\text{st } r$. Indeed, for $r = a.1 + b.0$, if $p' \text{ must } r$ implies p' has to interact on a and *not* on b , but then such a p' does not satisfy the derivative $r = a.0 + b.1$, i.e., $p' \not\text{m}\mu\text{st } r$ (because the composition $p' \parallel r$ is stable but *not* client-successful). Using a symmetric argument we deduce that if $p' \text{ must } a.0 + b.1$ then $p' \not\text{m}\mu\text{st } a.1 + b.0$, and thus no process p exists that satisfies $r_1 + r_2$; note that the argument above crucially exploits the external non-determinism of $r_1 + r_2$. The client $\mu x.(c.(a.1 + b.x) + c.b.1)$ from Sect. 1 is unusable for similar reasons, the analysis being more involved due to infinite computations. ■

We let $(r \text{ after } \not\llcorner s) = \{r' \mid r \xrightarrow{s} \not\llcorner r'\}$, and call the set $(r \text{ after } \not\llcorner s)$ the *residuals* of r after the *unsuccessful* trace s . We extend the notion of usability and say that r is *usable* along an unsuccessful trace s whenever $r \text{ usbl}_{\not\llcorner} s$, which is the least predicate satisfying the conditions (a) $r \text{ usbl}_{\not\llcorner} \varepsilon$ whenever $r \in \mathcal{U}$, and (b) $r \text{ usbl}_{\not\llcorner} as$ whenever (i) $r \in \mathcal{U}$ and (ii) if $r \xrightarrow{a} \not\llcorner$ then $\bigoplus(r \text{ after } \not\llcorner a) \text{ usbl}_{\not\llcorner} s$. If $r \text{ usbl}_{\not\llcorner} s$, any state reachable from r by performing any unsuccessful subsequence of s is usable [6]. Finally, let $ua_{\text{ctt}}(r, s) = \{a \in \text{Act} \mid r \xrightarrow{sa} \not\llcorner \text{ implies } r \text{ usbl}_{\not\llcorner} sa\}$ denote all the usable actions for a client r after the unsuccessful trace s .

Definition 3 (Semantic client-preorder). *Let $r_1 \lesssim_{\text{ctt}} r_2$ if, for every $s \in \text{Act}^*$ such that $r_1 \text{ usbl}_{\not\llcorner} s$, we have (i) $r_2 \text{ usbl}_{\not\llcorner} s$, (ii) for every $B \in \text{Acc}_{\not\llcorner}(r_2, s)$ there exists a $A \in \text{Acc}_{\not\llcorner}(r_1, s)$ such that $A \cap ua_{\text{ctt}}(r_1, s) \subseteq B$, (iii) $r_2 \xrightarrow{s} \not\llcorner$ implies $r_1 \xrightarrow{s} \not\llcorner$.* ■

Theorem 1. *In any finite branching LTS, $r_1 \sqsubseteq_{\text{ctt}} r_2$ if and only if $r_1 \lesssim_{\text{ctt}} r_2$.*

Proof. Follows from [6, Theorem 3.13] and König’s Infinity Lemma.

Definition 3 enjoys a few pleasing properties and, through Theorem 1, sheds light on behavioural properties of clients related by \sqsubseteq_{ctt} . Concretely, it shares a similar structure to well-studied characterisations of the (standard) must-testing preorder of [12, 15], where process convergence is replaced by client usability, and traces and acceptance sets are replaced by their unsuccessful counterparts (modulo usable actions). Unfortunately, Definition 3 has a major drawback: it is parametric wrt. the set of usable clients \mathcal{U} (Definition 2), which relies on an existential quantifications over servers. As a result, the definition is *not* fully-abstract, and this makes it hard to use as proof technique and to ground decision procedures for \sqsubseteq_{ctt} on it.

3 Characterising Usability

We use the behavioural predicates of Sect. 2.1, together with the new predicate in Definition 4, to formulate the characterising properties of the set of usable clients \mathcal{U} (Proposition 1). We use these predicates to construct a set \mathcal{U}_{bhv} that coincides with \mathcal{U} (Theorem 2); this gives us an inductive proof method for determining usability.

Definition 4. *We write $r \Downarrow_{\checkmark}$ whenever for every infinite sequence of internal moves $r \xrightarrow{\tau} r_1 \xrightarrow{\tau} r_2 \xrightarrow{\tau} \dots$, there exists a state r_i such that $r_1 \checkmark$.* ■

Recalling Eq. (3), let $\text{Acc}_{\not\llcorner}(r) = \text{Acc}_{\not\llcorner}(r, \varepsilon)$. Proposition 1 crystallises the characteristic properties of usable clients, providing a blue print for our alternative definition Definition 5. Instead of giving a direct proof of this proposition, we obtain it indirectly as consequence of our other results.

Proposition 1. *For every $r \in \text{Proc}$, $r \in \mathcal{U}$ if and only if*

1. $r \Downarrow_{\surd}$, and
 2. if $A \in \text{Acc}_{\surd}(r)$, then there exists $a \in A$. ($r \xrightarrow{a}_{\surd}$ implies $\bigoplus(r \text{ after}_{\surd} a) \in \mathcal{U}$).
-

The proposition above states that a client r is usable if and only if, for every potentially deadlocked state r' reached via silent moves by r , there exists an action a that leads r' out of the potential deadlock, i.e., into another state r'' where r'' is certainly usable.

Example 3. We use Proposition 1 to discuss the (non) usability of clients from previous example. Recall $r_3 = \tau.(1 + \tau.0)$, $r'_3 = r_3 + \tau.0$ and $r''_3 = r_3 + \mu x.x$ from Example 1. Since we have $r_3 \Downarrow_{\surd}$ and $\text{Acc}_{\surd}(r_3) = \emptyset$, r_3 satisfies both condition of Proposition 1, with the second one being trivially true. As a consequence r_3 is usable, and indeed 0 must r_3 . On the contrary, we have $\text{Acc}_{\surd}(r'_3) = \{\emptyset\}$, thus r'_3 violates Proposition 1(2) and thus r'_3 is unusable. Client r''_3 is unusable as well, but violates Proposition 1(1) instead. Conversely, client $r'''_3 = r_3 + \tau.(1 + \mu x.x)$ satisfies both conditions of Proposition 1, and it is usable. For instance, 0 must r'''_3 .

A more involved client is $r_1 + r_2$ from Example 2. There we proved that $r_1 + r_2 \notin \mathcal{U}$, and indeed $r_1 + r_2$ does not satisfy Proposition 1(2). This is true because $\text{Acc}_{\surd}(r_1 + r_2) = \{\{c\}\}$, and $r' \notin \mathcal{U}$, where

$$r' = \bigoplus((r_1 + r_2) \text{ after}_{\surd} c) = \tau.(a.1 + b.0) + \tau.(a.0 + b.1).$$

In turn, the reason why r' is not usable is that $\text{Acc}_{\surd}(r') = \{\{a, b\}\}$, and Proposition 1(2) requires us to consider every set in $\{\{a, b\}\}$ — we have only $\{a, b\}$ to consider — and show that for some action $a' \in \{a, b\}$, $\bigoplus(r' \text{ after}_{\surd} a') \in \mathcal{U}$. It turns out that neither action in $\{a, b\}$ satisfies this condition. For instance, in the case of action b , we have $\bigoplus(r' \text{ after}_{\surd} b) = \tau.1 + \tau.0$ and $\text{Acc}_{\surd}(\tau.1 + \tau.0) = \{\emptyset\}$, so $\bigoplus(r' \text{ after}_{\surd} b)$ violates Proposition 1(2) and as a result $\bigoplus(r' \text{ after}_{\surd} b) \notin \mathcal{U}$. The reasoning why action a is *not* a good candidate either is identical. ■

Definition 5. *Let $\mathcal{F} : \mathcal{P}(\text{Proc}) \longrightarrow \mathcal{P}(\text{Proc})$ be defined by letting $r \in \mathcal{F}(S)$ whenever*

1. $r \Downarrow_{\surd}$, and
2. if $A \in \text{Acc}_{\surd}(r)$, then there exists an $a \in A$. ($r \xrightarrow{a}_{\surd}$ implies $\bigoplus(r \text{ after}_{\surd} a) \in S$).

We let $\mathcal{U}_{\text{bhv}} = \mu x.\mathcal{F}(x)$, the least fix-point of \mathcal{F} . ■

The function \mathcal{F} is continuous over the CPO $\langle \mathcal{P}(\text{Proc}), \subseteq \rangle$, thus Kleene fixed point theorem [31, Theorem 5.11] ensures that $\mu x.\mathcal{F}(x)$ (the least fix-point of \mathcal{F}) exists and is equal to $\bigcup_{n=0}^{\infty} \mathcal{F}^n(\emptyset)$ where $\mathcal{F}^0(S) = S$ and $\mathcal{F}^{n+1}(S) = \mathcal{F}(\mathcal{F}^n(S))$.

The bulk of the soundness result follows as a corollary from the next lemma, which also lays bare the role of non-recursive servers in proving usability of clients.

Lemma 1. *For every $n \in \mathbb{N}$ and $r \in \text{Proc}$, $r \in \mathcal{F}^n(\emptyset)$ implies that there exists a non-recursive server p such that p must r .* \square

An inductive argument is used to prove that \mathcal{U}_{bhv} is complete wrt. \mathcal{U} , where we define the following measure over which to perform induction. We let $MC(r, p)$ denote the set of maximal computations of a composition $r \parallel p$ and, for every computation $c \in MC(r, p)$, we associate the number $\#\text{itr}(c)$ denoting the number of *interactions* that take place between the initial state of c , and the *first successful state* of the computation c ($\#\text{itr}(c) = \infty$ whenever c is unsuccessful). Let $\text{itr}(r, p) = \max\{\#\text{itr}(c) \mid c \in MC(r, p)\}$. For instance, if $r = \mu x.a.x + b.1$, we have $\text{itr}(r, \bar{a}.\bar{a}.\bar{b}.0) = 3$, but $\text{itr}(r, \mu x.\bar{a}.x + \bar{b}.0) = \infty$.

Lemma 2. *Let T be a tree with root v . If T is finite branching and it has a finite number of nodes, then the number of paths $v \longrightarrow \dots$ is finite.* \square

Lemma 3. *In a finite branching LTS, p must r implies the number $\text{itr}(r, p)$ is finite.*

Proof. If p must r , every $c \in MC(r, p)$ reaches a successful state after a finite number of reductions. Since the number of interactions is not more than the number of reductions:

$$\text{for every } c \in MC(r, p). \#\text{itr}(c) \in \mathbb{N} \quad (4)$$

A set of successful computations from $r \parallel p$, e.g., $MC(r, p)$, may also be seen as a *computation tree*, where common prefixes reach the same node in the tree. In general, such a tree may have infinite depth. Consider the computation tree T obtained by *truncating* all the maximal computations of $r \parallel p$ at their *first* successful state, and let $TMC(r, p)$ be the set of all the computations obtained this way. It follows that

$$\{\#\text{itr}(c) \mid c \in MC(r, p)\} = \{\#\text{itr}(c) \mid c \in TMC(r, p)\} \quad (5)$$

From $\text{itr}(r, p) = \max\{\#\text{itr}(c) \mid c \in MC(r, p)\}$, (4) and (5) we know that that $\text{itr}(r, p)$ is finite if the set $\{c \mid c \in TMC(r, p)\}$ is finite. This will follow from Lemma 2 if we prove that the tree T has a finite number of nodes. By the contrapositive of König's Lemma [16, 17], since every node in the tree T above is finitely branching, and there are no infinite paths, then T necessarily contains a *finite* number of nodes. By Lemma 2, $\{c \mid c \in TMC(r, p)\}$ must also be finite, and hence we can put a (finite) natural number $\text{itr}(r, p) \in \mathbb{N}$ as an upper bound on the number of interactions required to reach success. \square

If the LTS is not *image-finite* then Lemma 3 is false. To see why, consider the infinite branching client r and the server p depicted in Fig. 3. Since r engages in *finite* sequences of a actions which are unbounded in size, and the p offers any number of interactions on action \bar{a} , we have that p must r , but the set $MC(r, p)$ contains an infinite amount of computations, and the number $\text{itr}(r, p)$ is not finite. Dually, even if the LTS of a composition $r \parallel p$ is finite branching and finite state,




Fig. 3. Servers and clients to discuss the hypothesis in Lemma 3

it is necessary that p must r for $\text{itr}(r, p)$ to be finite. Lemma 3 lets us associate a rank to every usable client r , defined as $\text{rank}(r) = \min\{\text{itr}(r, p) \mid p \text{ must } r\}$. The well-ordering of \mathbb{N} ensures that $\text{rank}(r)$ is defined for every usable r . When defined, the rank of a client r gives us information about its usability,¹ where we can stratify \mathcal{U} as follows:

$$\mathcal{U} = \bigcup_{i \in \mathbb{N}} \mathcal{U}^i, \quad \text{where } \mathcal{U}^i = \{r \in \text{Proc} \mid \text{rank}(r) = i\} \tag{6}$$

Lemma 4. *For every $i \in \mathbb{N}$, $r \in \mathcal{U}^i$ implies $r \in \mathcal{F}(\mathcal{F}^j(\emptyset))$ for some $j \leq i$. \square*

We are now ready to prove the main result of this section.

Theorem 2 (Full-abstraction usability). *The sets \mathcal{U} and \mathcal{U}_{bhv} coincide.*

Proof. To show $\mathcal{U} \subseteq \mathcal{U}_{\text{bhv}}$, pick an $r \in \mathcal{U}$. By (6), $r \in \mathcal{U}^i$ for some $i \in \mathbb{N}$, and by Lemma 4 we obtain $r \in \mathcal{F}^j(\emptyset) \subseteq \mathcal{U}_{\text{bhv}}$ for some $j \in \mathbb{N}^+$. To show $\mathcal{U}_{\text{bhv}} \subseteq \mathcal{U}$, pick an $r \in \mathcal{U}_{\text{bhv}}$. Definition 5 ensures that $\mathcal{U}_{\text{bhv}} \subseteq \bigcup_{n=0}^{\infty} \mathcal{F}^n(\emptyset)$, thus $r \in \mathcal{F}^n(\emptyset)$ for some $n \in \mathbb{N}$. Lemma 1 implies that $r \in \mathcal{U}$. The reasoning applies to any $r \in \mathcal{U}_{\text{bhv}}$, thus $\mathcal{U}_{\text{bhv}} \subseteq \mathcal{U}$. \square

4 The Client Preorder Revisited

By combining the definition of \preceq_{ct} with \mathcal{U}_{bhv} of Definition 5, Theorem 2 yields a fully-abstract characterisation of the client preorder \preceq_{ct} . In general, however, this characterisation still requires us to consider an infinite number of (unsuccessful) traces to establish client inequality. In this section, we put forth a novel coinductive definition for the client preorder and exploit the finite-branching property of the LTS to show that this definition characterises the contextual preorder \preceq_{ct} , Theorem 5. We also argue that this new characterisation is easier to use in practice than Definition 3, a claim that is substantiated by showing how this coinductive preorder can be used to prove the second result in this section, namely that servers offering a *finite* amount of interactions are sufficient and necessary to distinguish clients, Theorem 6. Subsequently, in Theorem 7, we also show that the coinductive preorder is decidable for our client language.

¹ Function min is *not* defined for empty sets, thus $\text{rank}(r)$ is undefined whenever r is unusable.

Example 4. The use of \preceq_{ct} is hindered, in practice, by the universal quantification over traces in its definition. Consider, for instance, clients r_4 and r_5 ,

$$r_4 = a.1 + \mu y.(a.r_3'' + b.y + c.1) \quad \text{and} \quad r_5 = (\mu z.(b.z + c.1)) + d.1$$

where $r_3'' = (\tau.(1 + \tau.0)) + \mu x.x$ from Example 1. One way to prove $r_4 \sqsubseteq_{\text{ct}} r_5$ amounts in showing that $r_4 \preceq_{\text{ct}} r_5$, even though this task is far from obvious. Concretely, the definition of \preceq_{ct} requires us to show that for *every* trace $s \in \text{Act}^*$ where $r_4 \text{ usbl}_{\not\sim} s$ holds, clauses (i), (ii) and (iii) of Definition 3 also hold. In this case, there are an *infinite* number of such unsuccessful traces s to consider and, a priori, there is no clear way how to do this in finite time. Specifically, there are (unsuccessful) traces that r_4 *can* perform while remaining usable at every step, such as $s = b^n$, but also (unsuccessful) traces that r_4 *cannot* perform (which trivially imply $r_4 \text{ usbl}_{\not\sim} s$ according to the definition in Sect. 2.1), such as $s = d(b^n)$, $s = (db)^n$ or $s = (ac)^n$.

The definition of $r_4 \text{ usbl}_{\not\sim} s$ does however rule out a number of traces to consider, and Definition 5 helps us with this analysis. For instance, for $s = a$, we have $\neg(r_4 \text{ usbl}_{\not\sim} a)$ because $\bigoplus(r_4 \text{ after}_{\not\sim} a) = (\tau.1 + \tau.r_3'' + \tau.0 + \tau.\mu x.x)$ and, by using similar reasoning to that in Example 3 for r_3'' , we know that $\neg((r_4 \text{ after}_{\not\sim} a) \Downarrow_{\checkmark})$ which implies $\bigoplus(r_4 \text{ after}_{\not\sim} a) \notin \mathcal{U}_{\text{bhv}}$ and, by Theorem 2, we have $\bigoplus(r_4 \text{ after}_{\not\sim} a) \notin \mathcal{U}$. \square

To overcome the problems outlined in Example 4, we identify three properties of the preorder \sqsubseteq_{ct} , stated in Lemma 5, which partly motivate the conditions defining the transfer function \mathcal{G} in Definition 6. Conditions (ii) and (iii) are explained in greater detail as discussions to points (2) and (3c) of Definition 6 below.

Lemma 5. $r_1 \sqsubseteq_{\text{ct}} r_2$ implies (i) if $r_2 \xrightarrow{\tau}_{\not\sim} r_2'$ then $r_1 \sqsubseteq_{\text{ct}} r_2'$; (ii) if $r_2 \not\checkmark$ then $r_1 \not\checkmark$; (iii) if $r_2 \xrightarrow{a}_{\not\sim}$ then $(r_1 \xrightarrow{a}_{\not\sim})$ and $\bigoplus(r_1 \text{ after}_{\not\sim} a) \sqsubseteq_{\text{ct}} \bigoplus(r_2 \text{ after}_{\not\sim} a)$. \square

Definition 6. Let $\mathcal{G} : \mathcal{P}(\text{Proc} \times \text{Proc}) \rightarrow \mathcal{P}(\text{Proc} \times \text{Proc})$ be the function such that $(r_1, r_2) \in \mathcal{G}(R)$ whenever all the following conditions hold:

1. if $r_2 \xrightarrow{\tau}_{\not\sim} r_2'$ then $r_1 R r_2'$
2. if $r_2 \not\checkmark$ then $r_1 \not\checkmark$
3. if $r_1 \in \mathcal{U}_{\text{bhv}}$ then
 - (a) $r_2 \in \mathcal{U}_{\text{bhv}}$
 - (b) if $B \in \text{Acc}_{\not\sim}(r_2)$ then there exists an $A \in \text{Acc}_{\not\sim}(r_1)$ such that $A \cap ua_{\text{bhv}}(r_1) \subseteq B$
 - (c) if $r_2 \xrightarrow{a}_{\not\sim}$ then $(r_1 \xrightarrow{a}_{\not\sim})$ and $\bigoplus(r_1 \text{ after}_{\not\sim} a) R \bigoplus(r_2 \text{ after}_{\not\sim} a)$

where $ua_{\text{bhv}}(r) = \{a \mid r \xrightarrow{a}_{\not\sim} \text{ implies } \bigoplus(r \text{ after}_{\not\sim} a) \in \mathcal{U}_{\text{bhv}}\}$. Let $\preceq_{\text{ct}} = \nu x.\mathcal{G}(x)$ where $\nu x.\mathcal{G}(x)$ denotes the greatest fixpoint of \mathcal{G} . The function \mathcal{G} is monotone over the complete lattice $\langle \mathcal{P}(\text{Proc} \times \text{Proc}), \subseteq \rangle$ and thus $\nu x.\mathcal{G}(x)$ exists. \square

The definition of \mathcal{G} follows a similar structure to that of the resp. definitions that coinductively characterise the must preorder for servers [18, 25]. Definition 6, however, uses predicates for clients, i.e., unsuccessful traces and usability, in place of the predicates for servers, i.e., traces and convergence. Note, in particular, that we use the *fully-abstract* version of usability, \mathcal{U}_{bhv} , from Definition 5 and adapt the definition of usable actions accordingly, $ua_{\text{bhv}}(r)$. Another subtle but crucial difference in Definition 6 is condition (2). The next example elucidates why such a condition is necessary for \preceq_{ct} to be sound.

Counterexample 3. Let \mathcal{G}_{bad} be defined as \mathcal{G} in Definition 6, but without part (2). In this case, we prove that the pair of clients $(1, \tau.1)$ is contained in the greatest fixed point of \mathcal{G}_{bad} , and then proceed to show that this pair is not contained in \preceq_{ct} . Let $R = \{(1, \tau.1)\}$. It follows that $R \subseteq \mathcal{G}_{\text{bad}}(R)$ if all the conditions for \mathcal{G}_{bad} are satisfied: condition (1) is trivially true, condition (3a) is true because $0 \text{ must } 1$ and $0 \text{ must } \tau.1$, condition (3b) holds trivially because $\text{Acc}_{\neq}(\tau.1) = \emptyset$, whereas condition (3c) is satisfied because $\tau.1$ does not perform any strong actions. It therefore follows that $(1, \tau.1) \in \mu x. \mathcal{G}_{\text{bad}}(x)$. Contrarily, $1 \not\preceq_{\text{ct}} \tau.1$ because the divergent server τ^∞ distinguishes between the two clients: whereas $\tau^\infty \text{ must } 1$ since the client succeeds immediately, we have $\tau^\infty \text{ m\text{ust} } \tau.1$ because the composition $\tau.1 \parallel \tau^\infty$ has an infinite unsuccessful computation due to the divergence of τ^∞ . ■

A more fundamental difference between Definition 6 and the coinductive server preorders in [18, 25] is that, in Definition 6(3c), the relation R has to relate internal sums of derivative clients on *both* sides. Although non-standard, this condition is sufficient to compensate for the lack of compositionality of usable clients (see clients r_1 and r_2 (1) from Sect. 1). Using the standard weaker condition makes the preorder \preceq_{ct} unsound wrt. \preceq_{ct} , as we proceed to show in the next example.

Counterexample 4. Let \mathcal{G}_{bad} be defined as \mathcal{G} in Definition 6, but replacing the condition (3c) with the relaxed condition in (3bad) below, which requires each derivative r'_2 to be analysed in isolation. We show that the greatest fixpoint of \mathcal{G}_{bad} , $\preceq_{\text{ct}}^{\text{bad}}$, contains client pairs that are not in \preceq_{ct} .

$$\text{if } r_2 \xrightarrow{a} r'_2 \text{ then } (r_1 \xrightarrow{a} r'_1 \text{ and } \bigoplus (r_1 \text{ after } a) R r'_2) \quad (3\text{bad})$$

Consider the clients $r_6 = c.r'_6$ and $r_7 = (r_1 + r_2) + \tau.1$ where

$$r'_6 = \tau.r_6^a + \tau.r_6^b \quad r_6^a = a.0 + \tau.1 \quad r_6^b = b.0 + \tau.1$$

and r_1 and r_2 are the clients defined in (1) above. On the one hand, we have that $r_6 \not\preceq_{\text{ct}} r_7$, because $\bar{c}.0 \text{ must } r_6$ whereas $\bar{c}.0 \text{ m\text{ust} } r_7$. On the other hand, we now show that $r_6 \preceq_{\text{ct}}^{\text{bad}} r_7$. Focusing on condition Definition 6(3), we start by deducing that $r_6 \in \mathcal{U}_{\text{bhv}}$ (either directly using Definition 5 or indirectly through $\bar{c}.0 \text{ must } r_6$, recalling Theorem 2). Now, Definition 6(3a) is true because

0 must r_7 , thus r_7 is usable, and thanks to Theorem 2 we have $r_7 \in \mathcal{U}_{\text{bhv}}$. Also point (3b) is satisfied, because $\text{Acc}_{\mathcal{H}}(r_7) = \text{Acc}_{\mathcal{H}}(r_6) = \{\{a\}\}$.² To prove that the (relaxed) condition (3bad) holds, we have to show that

$$r_6^c \preceq_{\text{ct}}^{\text{bad}} a.1 + b.0 \quad \text{and} \quad r_6^c \preceq_{\text{ct}}^{\text{bad}} a.0 + b.1, \quad \text{with } r_6^c = r_6' + \tau.r_6^a + \tau.r_6^b \quad (7)$$

Let $r_7' = a.1 + b.0$. We only show the proof for the inequality $r_6^c \preceq_{\text{ct}}^{\text{bad}} r_7'$, since the proof for the other inequality is analogous. We focus again on conditions (3a), (3b), and (3bad). Condition (3a) is true because 0 must r_6^c , and thus $r_6^c \in \mathcal{U} = \mathcal{U}_{\text{bhv}}$, and because $r_7' \in \mathcal{U} = \mathcal{U}_{\text{bhv}}$ as well (e.g., $\bar{a}.0$ must r_7'). Condition (3b) holds because $\text{Acc}_{\mathcal{H}}(r_7') = \{\{c\}\}$ and $\text{Acc}_{\mathcal{H}}(r_6^c) = \{\{b\}, \{c\}\}$. Finally for (3bad) we only have to check the case for $r_7' \xrightarrow{b}_{\mathcal{H}} 0$, which requires us to show that $\tau.0 \preceq_{\text{ct}}^{\text{bad}} 0$; this latter check is routine. As a result, we have $r_6^c \preceq_{\text{ct}}^{\text{bad}} r_7'$. Since we can also show that $r_6^c \preceq_{\text{ct}}^{\text{bad}} a.0 + b.1$ holds, we obtain (7), and consequently $r_6 \preceq_{\text{ct}}^{\text{bad}} r_7$. \blacksquare

After our digression on Definition 6, we outline why \preceq_{ct} coincides with $\bar{\preceq}_{\text{ct}}$. A detailed proof can be found in the full version of this paper [4].

Lemma 6. *Whenever $r_1 \preceq_{\text{ct}} r_2$, for every $s \in \text{Act}^*$, $r_1 \text{ usbl}_{\mathcal{H}} s$ implies $r_2 \text{ usbl}_{\mathcal{H}} s$ and also that for every $B \in \text{Acc}_{\mathcal{H}}(r_2, s)$, there exists an set $A \in \text{Acc}_{\mathcal{H}}(r_1, s)$ such that $A \cap \text{ua}_{\text{ct}}(r_2, s) \subseteq B$; and that if $r_2 \xrightarrow{s}_{\mathcal{H}}$ then $r_1 \xrightarrow{s}_{\mathcal{H}}$. \square*

Theorem 5. *In any finite branching LTS $r_1 \bar{\preceq}_{\text{ct}} r_2$ if and only if $r_1 \preceq_{\text{ct}} r_2$.*

Proof. We have to show the set inclusions, $\bar{\preceq}_{\text{ct}} \subseteq \preceq_{\text{ct}}$ and $\preceq_{\text{ct}} \subseteq \bar{\preceq}_{\text{ct}}$. Lemma 5 and Theorem 1 imply that $\bar{\preceq}_{\text{ct}} \subseteq \mathcal{G}(\bar{\preceq}_{\text{ct}})$, and thus, by the Knaster-Tarski theorem, we obtain the first inclusion. The second set inclusion follows from Theorem 1 and Lemma 6. \square

Example 5. Recall clients $r_4 = a.1 + \mu y.(a.r_3'' + b.y + c.1)$ and $r_5 = (\mu z.(b.z + c.1)) + d.1$ from Example 4, used to argue that the alternative relation \preceq_{ct} is still a burdensome method for reasoning on $\bar{\preceq}_{\text{ct}}$. By contrast, We now contend that it is simpler to show $r_4 \bar{\preceq}_{\text{ct}} r_5$ by proving $r_4 \preceq_{\text{ct}} r_5$, thanks to Theorem 5 and the Knaster-Tarski theorem. By Definition 6, it suffices to provide a witness relation R such that $(r_4, r_5) \in R$ and $R \subseteq \mathcal{G}(R)$. Let $R = \{(r_4, r_5), (r_4', r_5')\}$ where $r_3'' = (\tau.(1 + \tau.0)) + \mu x.x$ from Example 1, $r_4' = \mu y.(a.r_3'' + b.y + c.1)$, and $r_5' = \mu z.(b.z + c.1)$. Checking that R satisfies the conditions in Definition 6 is routine work. To prove condition (3b), though, note that $\text{Acc}_{\mathcal{H}}(r_5) = \text{Acc}_{\mathcal{H}}(r_5') = \{\{b, c\}\}$ and that $\text{Acc}_{\mathcal{H}}(r_4) = \{\{a, b, c\}\}$. However $\text{ua}_{\text{bhv}}(r_4) = \{b, c\}$ and thus the required set inclusion $(\{a, b, c\} \cap \{b, c\}) \subseteq \{b, c\}$ holds. \blacksquare

The coinductive preorder of \preceq_{ct} may also be used to prove that two clients are *not* in the contextual preorder $\bar{\preceq}_{\text{ct}}$: by iteratively following the conditions of Definition 6 one can determine whether a relation including the pair of clients exists.

² The restriction of the left hand side of the inclusion of Definition 6(3b) by $\text{ua}_{\text{bhv}}(r_6)$ is superfluous.

This approach is useful when guessing a discriminating server is not straightforward; in failing to define a such relation R one obtains information on how to construct the discriminating server.

Example 6. Recall the clients r_6 and r_7 considered in Counterexample 4. By virtue of the full-abstraction result, we can show directly that $r_6 \not\sqsubseteq_{\text{ct}} r_7$ by following the requirements of Definition 6 and arguing that no relation exists that contains the pair (r_6, r_7) while satisfying the conditions of the coinductive preorder. Without loss of generality, pick a relation R such that $r_6 R r_7$: we have to show that $R \subseteq \mathcal{G}(R)$. Since $r_6 \in \mathcal{U}_{\text{bhv}}$, $r_7 \xrightarrow{c} \not\llcorner$ and $r_6 \xrightarrow{c} \llcorner$, Definition 6(3c) requires that we show that

$$r_6^c R \tau.r_7' + \tau.r_7'' \text{ where } r_6^c = \bigoplus (r_6 \text{ after } \not\llcorner c) \text{ and } (\tau.r_7' + \tau.r_7'') = \bigoplus (r_7 \text{ after } \not\llcorner c) \quad (8)$$

and r_6^c , r_7' and r_7'' are the clients defined earlier in Counterexample 4. Since we want to show that $R \not\subseteq \mathcal{G}(R)$, the condition Definition 6(3a) requires that, if $r_6^c \in \mathcal{U}_{\text{bhv}}$, then $(\tau.r_7' + \tau.r_7'') \in \mathcal{U}_{\text{bhv}}$. However, even though $r_6^c \in \mathcal{U}_{\text{bhv}}$, we have $(\tau.r_7' + \tau.r_7'') \notin \mathcal{U}_{\text{bhv}}$, violating Definition 6(3a) and thus showing that no such R satisfying both $(r_6, r_7) \in R$ and $R \subseteq \mathcal{G}(R)$ can exist. We highlight the fact that whereas (7) of Counterexample 4 resulted in $r_6 \not\lesssim_{\text{ct}}^{\text{bad}} r_7$, (8) is instrumental to conclude that $r_6 \not\lesssim_{\text{ct}} r_7$. Note also that the path along c leading to a violation of the requirements of Definition 6 is related to the discriminating server $\bar{c}.0$ used in Counterexample 4 to justify $r_6 \not\sqsubseteq_{\text{ct}} r_7$. ■

5 Expressiveness and Decidability

We show that servers with finite interactions suffice to preserve the discriminating power of the contextual preorder \sqsubseteq_{ct} in Definition 1, which has ramifications on standard verification techniques for the preorder, such as counter-example generation [11]. We also show that, for finite-state LTSs, the set of usable clients is decidable. Using standard techniques [27] we then argue that, in such cases, there exists a procedure to decide whether two finite-state clients are related by \sqsubseteq_{ct} .

5.1 On the Power of Finite Interactions

We employ the coinductive characterisation of the client preorder, Theorem 5, to prove an important property of the client preorder of Definition 1, namely that servers that only offer a *finite amount of interactions* to clients are necessary and sufficient to distinguish all the clients according to our touchstone preorder \sqsubseteq_{ct} of Definition 1. Let $\text{CCS}^f ::= 0 \mid 1 \mid \alpha.p \mid p + q \mid \tau^\infty$, and

$$\begin{aligned} \sqsubseteq_{\text{ct}}^f &= \{ (r_1, r_2) \mid \text{for every } p \in \text{CCS}^f. p \text{ must } r_1 \text{ implies } p \text{ must } r_2 \} \\ \mathcal{U}^f &= \{ r \mid \text{there exists } p \in \text{CCS}^f. p \text{ must } r \} \end{aligned}$$

In what follows, we find it convenient to use the definitions above: CCS^f excludes recursively-defined processes, but explicitly adds the divergent process τ^∞ because of its discriminating powers (see Counterexample 3). Accordingly, $\bar{\simeq}_{\text{ct}}^f$ and \mathcal{U}^f restrict the *resp.* sets to the syntactic class CCS^f .

Corollary 1. *The sets \mathcal{U} and \mathcal{U}^f coincide.*

Proof. The inclusion $\mathcal{U}^f \subseteq \mathcal{U}$ is immediate. Suppose that $r \in \mathcal{U}$. By Theorem 2 we have $r \in \mathcal{U}_{\text{bhv}}$. By Lemma 1, there exists a non-recursive $p \in \text{CCS}^f$ such that p must r , thus $r \in \mathcal{U}^f$ follows. \square

Theorem 6. *In any finite-branching LTS $r_1 \bar{\simeq}_{\text{ct}}^f r_2$ if and only if $r_1 \bar{\simeq}_{\text{ct}} r_2$.*

Proof. The inclusion $\bar{\simeq}_{\text{ct}} \subseteq \bar{\simeq}_{\text{ct}}^f$ follows immediately from the *resp.* definitions. On the other hand, Theorem 5 provides us with a proof technique for showing the inclusion $\bar{\simeq}_{\text{ct}}^f \subseteq \bar{\simeq}_{\text{ct}}$: if we show that $\bar{\simeq}_{\text{ct}}^f \subseteq \mathcal{G}(\bar{\simeq}_{\text{ct}}^f)$ then $\bar{\simeq}_{\text{ct}}^f \subseteq \preceq_{\text{ct}} = \bar{\simeq}_{\text{ct}}$. In view of the Knaster-Tarski theorem it suffices to show that $\bar{\simeq}_{\text{ct}}^f \subseteq \mathcal{G}(\bar{\simeq}_{\text{ct}}^f)$. In turn, this requires us to prove the three conditions stated in Definition 6. The argument for the first two conditions is virtually the same to that of Lemma 5. Similarly, the arguments for the third condition follow closely those used in Theorem 1 (albeit in a simpler setting of unsuccessful traces of length 1). The only new reasoning required is that servers that exists because of $r_1 \in \mathcal{U}$ also belong to CCS^f , which we know from Corollary 1. \square

An analogous result should also hold for the server-preorder, for the proofs of completeness in [6, Theorem 3.1] rely on clients that can be written in the language CCS^f .

5.2 Deciding the Client Preorder

Figure 4 describes the pseudo-code for the eponymous function $\text{isUsable}(r, \text{acm})$, which is meant to determine whether a client r is usable. It adheres closely to the conditions of Definition 5 for \mathcal{U}_{bhv} , using acm as an *accumulator* to keep track of all the terms that have already been explored. Thus, if an r is revisited, the algorithm rejects it on the basis that a loop of unsuccessful interactions (leading to an infinite sequence of unsuccessful interactions that makes the client unusable) is detected (lines 2–3). If not, the algorithm checks for the conditions in Definition 5 (lines 4–9). In particular, line 4 checks that infinite sequences of internal moves are always successful (using function convtick defined on lines 11–17) and that partially deadlocked clients reached through a finite number of unsuccessful internal moves, $\text{Acc}_{\not\sim}(r) \neq \emptyset$, contain at least one action that unblocks them to some other usable client (lines 7–8). This latter check employs the function $\text{existsUnblockAction}$ (defined on lines 19–26) which recursively calls isUsable to determine whether the client reached after an action is indeed usable. $\text{isUsable}(r, \text{acm})$ of Fig. 4 relies on the LTS of r being *finite-state* in order to guarantee termination via the state accumulation held in acm . This is indeed the case for our expository language CCS^μ of Fig. 2. Concretely, we define the

```

1  isUsable (r, acm) =
2    if r in acm
3    then false
4    else if convtick (∅, r)
5         then if Acc∕(r) == empty
6              then true
7              else BoolSet = map ( existsUnblockAction acm r ) Acc∕(r)
8                  conjunction BoolSet
9    else false
10 where
11   convtick(acm, r) =
12     if r in acm
13     then false
14     else if r  $\xrightarrow{\checkmark}$ 
15          then true
16          else BoolSet = map ( convtick (acm ∪ {r}) ) {r' | r  $\xrightarrow{\tau}$  r' }
17                  conjunction BoolSet
18 and
19   existsUnblockAction (acm, r, A) =
20     case A of
21       empty -> false
22       {a} ∪ A' ->
23         if r  $\xrightarrow{a}$ 
24            then if isUsable (⊕(r after∕ a), acm ∪ { r })
25                 then true else existsUnblockAction (r, A', acm)
26         else true

```

Fig. 4. An algorithm for deciding inclusion in the set \mathcal{U}

set of internal-sums for the derivatives that a client r reaches via all the finite traces $\in \text{Act}^*$, and show that this set is finite. Let

$$\text{sumsRdx}(r) = \{ \bigoplus (r \text{ after}_{\not\sim} s) \mid \text{for some } s \in \text{Act}^* \},$$

Lemma 7. *For every $r \in \text{CCS}^\mu$, the set $\text{sumsRdx}(r)$ is finite.* \square

Proof. Let $\text{Reach}_r = \{r' \mid r \xRightarrow{s} r' \text{ for some } s \in \text{Act}^*\}$ denote the set of reachable terms from client r , and $\text{PwrR}_r = \{ \bigoplus B \mid B \in \mathcal{P}(\text{Reach}_r) \}$ denote the elements of the powerset of Reach_r , expressed as internal summations of the elements of $\mathcal{P}(\text{Reach}_r)$. By definition, we have that $\text{sumsRdx}(r) \subseteq \text{PwrR}_r$. Hence, it suffices to prove that Reach_r is finite to show that PwrR_r is finite, from which the finiteness of $\text{sumsRdx}(r)$ follows. The proof of the finiteness of Reach_r is the same as that of Lemma 4.2.11 of [29] for the language serial-CCS, which is homologous to CCS^μ of Fig. 2 modulo the satisfaction construct 1. \square

Theorem 7. *For every $r \in \text{Proc}$ we have that*

- (i) $r \in \mathcal{U}$ iff $\text{isUsable}(r, \emptyset) = \text{true}$,
- (ii) $r \notin \mathcal{U}$ iff $\text{isUsable}(r, \emptyset) = \text{false}$.

Proof. For the *only-if* case of clause (i), we use Theorem 2 and show instead that $r \in \mathcal{U}_{\text{bhv}}$ implies $\text{isUsable}(r, \emptyset) = \text{true}$; we do so by numerical induction on $n \in \mathbb{N}^+$ where $r \in \mathcal{F}^n(\emptyset)$. For the *if* case, we dually show that $\text{isUsable}(r, \emptyset) = \text{true}$ implies $r \in \mathcal{U}_{\text{bhv}}$, by numerical induction on the *least* number $n \in \mathbb{N}^+$ of (recursive) calls to isUsable that yield the outcome true . We note that in either direction of clause (i), there is a direct correspondence between the respective inductive indices (e.g., for the base case $n = 1$, $r \in \mathcal{F}^1(\emptyset) = \mathcal{F}(\emptyset)$ implies that $r \Downarrow_{\checkmark}$ and that $\text{Acc}_{\checkmark}(r) = \emptyset$).

For the second clause (ii), the statements ($r \notin \mathcal{U}$ implies $\text{isUsable}(r, \emptyset) = \text{true}$) and ($\text{isUsable}(r, \emptyset) = \text{false}$ implies $r \notin \mathcal{U}$) contradict the first clause (i) which we just proved. The required result thus holds if we ensure that $\text{isUsable}(r, \emptyset)$ is defined for any $r \in \text{Proc}$. This follows from Lemma 7. \square

From Theorems 5, 7 and Lemma 7, we conclude that Definition 6 can be used to decide \preceq_{ct} for languages such as CCS^μ of Fig. 2. We can do this by adapting the algorithm of [27, Chapter 21.5], and proving that in our setting [27, Theorems 21.5.9 and 21.5.12] are true. In particular, using the terminology of [27] we have that $\text{reachable}_G(X)$ is finite, essentially because the *resp.* LTS is finite-state, and thus the decidability of \preceq_{ct} follows from Theorem 21.5.12.

6 Conclusion

We present a study that revolves around the notion of usability and preorders for clients (tests). Preorders for clients first appeared for compliance testing [2], and were subsequently investigated in [3, 6] for must testing [12] and extended to include peers. The characterisations given in [6] relied fundamentally on the set of usable terms \mathcal{U} which made them not fully-abstract and hard to automate. This provided the main impetus for our study. In general, recursion poses obstacles when characterising usable terms, but the very nature of **must** testing — which regards infinite unsuccessful computations as catastrophic — let us treat recursive terms in a finite manner (see Definition 5).

We focus on the client preorder, even though [6] presents preorders for both client and peers; note however that [6, Theorem 3.20] and Theorem 2 imply full-abstractness for the peer preorder as well. Our investigations and the *resp.* proofs for Theorem 2, Theorems 5 and 6 are conducted in terms of finitely-branching LTSs, which cover the semantics used by numerous other work describing client and server contracts [6, 8, 9, 18] — we only rely on an internal choice construct to economise on our presentation, but this can be replaced by tweaking the *resp.* definitions so as to work on sets of processes instead. As a consequence, the results obtained should also extend to arbitrary languages enjoying the finite-branching property. Theorem 7 relies on a stronger property, namely that the

language is finite-state. In [29], it is shown that this property is also enjoyed by larger CCS fragments, and we therefore expect our results to extend to these fragments as well.

6.1 Related Work

Client usability depends both on language expressiveness and on the notion of testing employed. Our comparison with the related work is organised accordingly.

Session types [14] do not contain unsuccessful termination, 0, restrict internal (resp. external) choices to contain only pair-wise distinct outputs (resp. inputs) and are, by definition, strongly convergent [25] (i.e., no infinite sequences of τ -transitions). E.g., $\tau.!a.1 + \tau.!b.?c.1$ corresponds to a session type in our language (modulo syntactic transformations such as those for internal choices), whereas $\tau.!a.0 + \tau.!b.?c.1$, $\tau.!a.1 + \tau.!a.?b.1$ and $?a.1 + ?a.!b.1$ do *not*. Since they are mostly deterministic — only internal choices on outputs are permitted — usability is relatively easy to characterise. In fact [7, Section 5] shows that every session type is usable wrt. compliance testing (even in the presence of higher-order communication) whereas, in [26, Theorem 4.3], *non-usable* session types are characterised wrt. fair testing. First-order session types are a subset of our language, and hence, Theorem 2 is enough to (positively) characterise usable session types wrt. must testing; we leave the axiomatisation of \mathcal{U} in this setting as future work.

Contracts [25] are usually formalised as (mild variants of) our language CCS^μ . In the case of must testing, the authors in [6, Theorem 6.9, Lemma 7.8(2)] characterise *non-usable* clients (and peers) for the sublanguage CCS^f as the terms that can be re-written into 0 via equational reasoning. Full-abstraction for usable clients wrt. compliance testing has been solved for *strongly convergent terms* in [25, Proposition 4.3] by giving a coinductive characterisation for viable (i.e., usable wrt. compliance) contracts. If we restrict our language to strongly convergent terms, that characterisation is neither sound nor complete wrt. must testing. It is unsound because clients such as $\mu x.a.x$ are viable but *not* usable. It is incomplete because of clients such as $r = 1 + \tau.0$; this client is usable wrt. must because, for arbitrary p , any computation of $p \parallel r$ is successful (since we have $r \xrightarrow{\checkmark}$ immediately). On the other hand, r is *not* viable wrt. compliance testing of [25] (where every server is strongly convergent), because for any server p we observe the computation starting with the reduction $p \parallel r \xrightarrow{\tau} p \parallel 0$, and once p stabilises to some p' , the final state $p' \parallel 0$ contains an unsuccessful client. This argument relies on subtle discrepancies in the definitions of the testing relations: in must testing it suffices for maximal computations to *pass through* a successful state, whereas in compliance testing the *final state* of the computation (if any) is required to be successful. This aspect impinges on the technical development: although our Definition 5(2) resembles [25, Definition 4.2], the two definitions have strikingly different meanings: we are forced to reason wrt. *unsuccessful* actions and *unsuccessful* acceptance sets whereas [25, Definition 4.2] is defined in terms of (standard) weak actions and acceptance sets (note that Definition 5(1) holds trivially in the strongly convergent setting of [25]). We note

also that our Definition 5 is inductive whereas [25, Definition 4.2] is coinductive. More importantly, our work lays bare the *non-compositionality* of usable terms and how it affects other notions that depend on it, such as Definition 6 (and consequently Theorem 5). We are unaware of any full-abstraction results for contract usability in the case of should-testing [8, 24, 28].

Future work: In the line of [10], we plan to show a logical characterisation of the client and peer preorder. We also intend to investigate coinductive characterisations for the peer preorder of [6] and subsequently implement decision procedures for the server, client, and peer preorders in CAAL [1]. Usability is not limited to tests. We expect it to extend naturally to runtime monitoring [13], where it can be used as a means of lowering runtime overhead by not instrumenting unusable monitors.

Acknowledgements. This research was supported by the COST Action STSMs IC1201-130216-067787 and IC1201-170214-038253. The first author was supported by the EU FP7 ADVENT project. The second author is partly supported by the RANNIS THEOFOMON project 163406-051. The authors acknowledge the Dagstuhl seminar 17051 and thank L. Aceto, M. Bravetti, A. Gorla, M. Hennessy, C. Spaccasassi and anonymous reviewers for their help and suggestions.

References

1. Andersen, J.R., Andersen, N., Enevoldsen, S., Hansen, M.M., Larsen, K.G., Olesen, S.R., Srba, J., Wortmann, J.K.: CAAL: concurrency workbench, Aalborg edition. In: Leucker, M., Rueda, C., Valencia, F.D. (eds.) ICTAC 2015. LNCS, vol. 9399, pp. 573–582. Springer, Cham (2015). doi:[10.1007/978-3-319-25150-9_33](https://doi.org/10.1007/978-3-319-25150-9_33)
2. Barbanera, F., de'Liguoro, F.: Two notions of sub-behaviour for session-based client/server systems. In: PPDP (2010)
3. Bernardi, G.: Behavioural equivalences for web services. Ph.D. thesis, TCD (2013)
4. Bernardi, G., Francalanza, A.: Full-abstraction for must testing preorders (extended abstract). <https://www.irif.fr/gio/papers/BFcoordination2017.pdf>
5. Bernardi, G., Hennessy, M.: Modelling session types using contracts. In: SAC (2012)
6. Bernardi, G., Hennessy, M.: Mutually testing processes. LMCS **11**(2), 1–23 (2015)
7. Bernardi, G., Hennessy, M.: Using higher-order contracts to model session types. LMCS **12**(2), 1–43 (2016)
8. Bravetti, M., Zavattaro, G.: A foundational theory of contracts for multi-party service composition. Fundam. Inf. **89**(4), 451–478 (2008)
9. Castagna, G., Gesbert, N., Padovani, L.: A theory of contracts for web services. ACM Trans. Program. Lang. Syst. **31**(5), 1–61 (2009)
10. Cerone, A., Hennessy, M.: Process behaviour: formulae vs. tests. In: EXPRESS (2010)
11. Clarke, E., Veith, H.: Counterexamples revisited: principles, algorithms, applications. In: Dershowitz, N. (ed.) Verification: Theory and Practice. LNCS, vol. 2772, pp. 208–224. Springer, Heidelberg (2003). doi:[10.1007/978-3-540-39910-0_9](https://doi.org/10.1007/978-3-540-39910-0_9)
12. De Nicola, R., Hennessy, M.: Testing equivalences for processes. TCS **34**(1–2), 83–93 (1984)

13. Francalanza, A.: A theory of monitors. In: Jacobs, B., Löding, C. (eds.) FoSSaCS 2016. LNCS, vol. 9634, pp. 145–161. Springer, Heidelberg (2016). doi:[10.1007/978-3-662-49630-5_9](https://doi.org/10.1007/978-3-662-49630-5_9)
14. Gay, S.J., Hole, M.: Subtyping for session types in the pi calculus. *Acta Inf.* **42**(2–3), 191–225 (2005)
15. Hennessy, M.: *Algebraic Theory of Processes*. MIT Press, Cambridge (1988)
16. Knuth, D.E.: *The Art of Computer Programming, Volume 1 (3rd Ed.): Fundamental Algorithms*. Addison Wesley Longman Publishing Co., Inc., Redwood City (1997)
17. König, D.: Über eine schlussweise aus dem endlichen ins unendliche. *Acta Litt. ac. sci. Szeged* **3**, 121–130 (1927)
18. Laneve, C., Padovani, L.: The Must preorder revisited. In: Caires, L., Vasconcelos, V.T. (eds.) CONCUR 2007. LNCS, vol. 4703, pp. 212–225. Springer, Heidelberg (2007). doi:[10.1007/978-3-540-74407-8_15](https://doi.org/10.1007/978-3-540-74407-8_15)
19. Luo, Q., Hariri, F., Eloussi, L., Marinov, D.: An empirical analysis of flaky tests. In: FSE (2014)
20. Marinescu, P., Hosek, P., Cadar, C.: Covrig: a framework for the analysis of code, test, and coverage evolution in real software. In: ISSTA (2014)
21. Martens, A.: Analyzing web service based business processes. In: Cerioli, M. (ed.) FASE 2005. LNCS, vol. 3442, pp. 19–33. Springer, Heidelberg (2005). doi:[10.1007/978-3-540-31984-9_3](https://doi.org/10.1007/978-3-540-31984-9_3)
22. Memon, A.M., Cohen, M.B.: Automated testing of GUI applications: models, tools, and controlling flakiness. In: ICSE (2013)
23. Milner, R.: *Communication and Concurrency*. Prentice-Hall, Upper Saddle River (1989)
24. Mooij, A.J., Stahl, C., Voorhoeve, M.: Relating fair testing and accordance for service replaceability. *J. Log. Algebr. Program.* **79**(3–5), 233–244 (2010)
25. Padovani, L.: Contract-based discovery of web services modulo simple orchestrators. *TCS* **411**(37), 3328–3347 (2010)
26. Padovani, L.: Fair subtyping for multi-party session types. *MSCS* **26**(3), 238–302 (2016)
27. Pierce, B.: *Types and Programming Languages*. MIT Press, Cambridge (2002)
28. Rensink, A., Vogler, W.: Fair testing. *Inf. Comput.* **205**(2), 125–198 (2007)
29. Spaccasassi, C.: *Language support for communicating transactions*. Ph.D. thesis, TCD, (2015)
30. Weinberg, D.: Efficient controllability analysis of open nets. In: Bruni, R., Wolf, K. (eds.) WS-FM 2008. LNCS, vol. 5387, pp. 224–239. Springer, Heidelberg (2009). doi:[10.1007/978-3-642-01364-5_14](https://doi.org/10.1007/978-3-642-01364-5_14)
31. Winskel, G.: *The Formal Semantics of Programming Languages: An Introduction*. MIT Press, Cambridge (1993)