Full-Abstraction for Must Testing Preorders

Giovanni Bernardi, Adrian Francalanza

► To cite this version:

Giovanni Bernardi, Adrian Francalanza. Full-Abstraction for Must Testing Preorders. 19th International Conference on Coordination Languages and Models (COORDINATION), Jun 2017, Neuchâtel, Switzerland. pp.237-255, 10.1007/978-3-319-59746-1_13 . hal-01657337

HAL Id: hal-01657337

<https://hal.inria.fr/hal-01657337>

Submitted on 6 Dec 2017

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

Full-abstraction for Must Testing Preorders (Extended Abstract)

Giovanni Bernardi and Adrian Francalanza

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

² University of Malta, Msida; 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 [6,3].

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 $r_1 \sqsubseteq_{\text{clt}} r_2$, 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.

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 [18,25,7,6,26]. 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

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

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 Figure 1:

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

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., $p \not\text{must } r_1 + r_2$ 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., $\bar{c}.\bar{a}.0 + \bar{b}.0$ 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, pag. 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 $\sqsubseteq_{\text{clt}}$, 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 $\sqsubseteq_{\text{clt}}$, Theorem 5;
- a contextual preorder $\sqsubseteq_{\text{clt}}^f$ that is equivalent to $\sqsubseteq_{\text{clt}}$ 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 [22,20,19]. 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: Section 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 Section 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 $\tau, \checkmark$ 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 $\lambda \in \text{Act}_{\tau\checkmark}$. 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 $s \in (\text{Act}_{\tau\checkmark})^*$ we also have the standard weak transitions, $p \xRightarrow{s} q$, defined by *ignoring* the occurrences of τ s.

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 Figure 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 Figure 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

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

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

is given in Figure 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 $p \text{ must } r$ if every maximal computation from $p \parallel r$ is client-successful, and write $r_1 \sqsubseteq_{\text{clt}} r_2$ if, for every p , $p \text{ must } r_1$ implies $p \text{ must } r_2$. ■

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

2.1 Characterising the client preorder

In [6, Def. 3.10, pg. 9], an alternative characterisation for the preorder $\sqsubseteq_{\text{clt}}$ is given and proven to be sound and complete. We recall this characterisation, restating the *resp.* notation. The alternative characterisation relies on *unsuccessful* traces: $r \xRightarrow{s}_{\not\sim} r'$ 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, $r \xRightarrow{s}_{\not\sim} r'$ is the least relation satisfying (a) $r \not\xrightarrow{\checkmark}$ implies $r \xRightarrow{\varepsilon}_{\not\sim} r$, and (b) if $r'' \xRightarrow{s}_{\not\sim} r'$ and $r \not\xrightarrow{\checkmark}$ then (i) $r \xrightarrow{a} r''$ implies $r \xRightarrow{as}_{\not\sim} r'$, and (ii) $r \xrightarrow{\tau} r''$ implies $r \xRightarrow{s}_{\not\sim} r'$. The *unsuccessful* acceptance set of r after s , are defined as

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

where $S(r) = \{a \in \text{Act} \mid r \xrightarrow{a}\}$ denotes the strong actions of r . Intuitively, for the client r , the set $\text{Acc}_{\not\sim}(r, s)$ 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 $r \xrightarrow{\alpha}_{\not\sim} r'$ whenever $r \xrightarrow{\alpha} r', r \not\rightarrow^{\checkmark}$ and $r' \not\rightarrow^{\checkmark}$ hold.

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

Note that, whenever $\text{Acc}_{\not\sim}(r, s) = \emptyset$, 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 $S(r') \in \text{Acc}_{\not\sim}(r, s)$ 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 Section 1. We show that despite being individually usable, the sum of these clients is not: $p \not\text{must } r_1 + r_2$ for *every* p . Fix a process p . If p does not offer an interaction on $\bar{c}$, then, plainly, $p \not\text{must } r_1 + r_2$. Suppose that $p \xrightarrow{\bar{c}} p'$; to prove $p \not\text{must } r_1 + r_2$, 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 $p' \not\text{must } r$. 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., $p' \not\text{must } r$ (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 $p' \not\text{must } a.1 + b.0$, 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 Section 1 is unusable for similar reasons, the analysis being more involved due to infinite computations. ■

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

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

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

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 $\sqsubseteq_{\text{clt}}$. 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 $\sqsubseteq_{\text{clt}}$ on it.

3 Characterising usability

We use the behavioural predicates of Section 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 $\mathcal{U}_{\text{bhv}}$ 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 $r_1 \xrightarrow{\checkmark}$. ■

Recalling Eq. (3), let $\text{Acc}_{\checkmark}(r) = \text{Acc}_{\checkmark}(r, \varepsilon)$. 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. $r \Downarrow_{\checkmark}$, and
2. if $A \in \text{Acc}_{\checkmark}(r)$, then there exists $a \in A$. ($r \xRightarrow{a}_{\checkmark}$ implies $\bigoplus (r \text{ after}_{\checkmark} a) \in \mathcal{U}$). □

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 $r_3 \Downarrow_{\checkmark}$ and $\text{Acc}_{\checkmark}(r_3) = \emptyset$, 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 $\text{Acc}_{\checkmark}(r'_3) = \{\emptyset\}$, 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 $\text{Acc}_{\checkmark}(r_1 + r_2) = \{\{c\}\}$, and $r' \notin \mathcal{U}$, where

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

In turn, the reason why r' is not usable is that $\text{Acc}_{\not\sim}(r') = \{\{a, b\}\}$, 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 $a' \in \{a, b\}$, $\bigoplus(r' \text{ after }_{\not\sim} a') \in \mathcal{U}$. It turns out that neither action in $\{a, b\}$ satisfies this condition. For instance, in the case of action b , we have $\bigoplus(r' \text{ after }_{\not\sim} b) = \tau.1 + \tau.0$ and $\text{Acc}_{\not\sim}(\tau.1 + \tau.0) = \{\emptyset\}$, so $\bigoplus(r' \text{ after }_{\not\sim} b)$ violates Proposition 1(2) and as a result $\bigoplus(r' \text{ after }_{\not\sim} b) \notin \mathcal{U}$. 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. $r \Downarrow_{\vee}$, and
2. if $A \in \text{Acc}_{\not\sim}(r)$, then there exists an $a \in A$. ($r \xRightarrow{a}_{\not\sim}$ implies $\bigoplus(r \text{ after }_{\not\sim} a) \in S$).

We let $\mathcal{U}_{\text{bvh}} = \mu x. \mathcal{F}(x)$, 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 $p \text{ must } r$. □

An inductive argument is used to prove that $\mathcal{U}_{\text{bvh}}$ 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 $\#itr(c)$ denoting the number of *interactions* that take place between the initial state of c , and the *first successful state* of the computation c ($\#itr(c) = \infty$ whenever c is unsuccessful). Let $itr(r, p) = \max\{\#itr(c) \mid c \in MC(r, p)\}$. For instance, if $r = \mu x. a.x + b.1$, we have $itr(r, \bar{a}.\bar{a}.\bar{b}.0) = 3$, but $itr(r, \mu x. \bar{a}.x + \bar{b}.0) = \infty$.

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. □

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

Proof. If $p \text{ must } r$, 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:

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

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

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

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

From $\text{itr}(r, p) = \max\{\#\text{itr}(c) \mid c \in MC(r, p)\}$, (4) and (5) we know 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 [17,16], 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 Figure 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, 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,³ where we can stratify $\mathcal{U}$ as follows:

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

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 $\mathcal{U}_{\text{bhv}}$ 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 $\preceq_{\text{clt}}$ with $\mathcal{U}_{\text{bhv}}$ of Definition 5, Theorem 2 yields a fully-abstract characterisation of the client preorder $\preceq_{\text{clt}}$. In general, however, this characterisation still requires us to consider an infinite number of (unsuccessful) traces to

³ Function min is *not* defined for empty sets, thus $\text{rank}(r)$ is undefined whenever r is unusable.

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 $\sqsubseteq_{\text{clt}}$, 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.

Example 4. The use of $\lesssim_{\text{clt}}$ 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 $r_4 \sqsubseteq_{\text{clt}} r_5$ amounts in showing that $r_4 \lesssim_{\text{clt}} r_5$, even though this task is far from obvious. Concretely, the definition of $\lesssim_{\text{clt}}$ 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 Section 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 $\bigoplus(r_4 \text{ after}_{\not\sim} a) \notin \mathcal{U}_{\text{bvh}}$ 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 $\sqsubseteq_{\text{clt}}$, 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. $r_1 \sqsubseteq_{\text{clt}} r_2$ implies (i) if $r_2 \xrightarrow{\tau}_{\not\sim} r_2'$ then $r_1 \sqsubseteq_{\text{clt}} r_2'$; (ii) if $r_2 \not\xrightarrow{\checkmark}$ then $r_1 \not\xrightarrow{\checkmark}$; (iii) if $r_2 \xrightarrow{a}_{\not\sim}$ then $(r_1 \xrightarrow{a}_{\not\sim})$ and $\bigoplus(r_1 \text{ after}_{\not\sim} a) \sqsubseteq_{\text{clt}} \bigoplus(r_2 \text{ after}_{\not\sim} a)$. $\square$

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 $r_2 \not\xrightarrow{\checkmark}$ then $r_1 \not\xrightarrow{\checkmark}$
3. if $r_1 \in \mathcal{U}_{\text{bvh}}$ then
 - (a) $r_2 \in \mathcal{U}_{\text{bvh}}$
 - (b) if $B \in \text{Acc}_{\not\sim}(r_2)$ then there exists an $A \in \text{Acc}_{\not\sim}(r_1)$ such that $A \cap \text{ua}_{\text{bvh}}(r_1) \subseteq B$
 - (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 $ua_{bhv}(r) = \{a \mid r \xRightarrow{a} \text{ implies } \bigoplus(r \text{ after } a) \in \mathcal{U}_{bhv}\}$. Let $\leq_{\text{clt}} = vx.\mathcal{G}(x)$ where $vx.\mathcal{G}(x)$ 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 $vx.\mathcal{G}(x)$ 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, $ua_{bhv}(r)$. Another subtle but crucial difference in Definition 6 is condition (2). The next example elucidates why such a condition is necessary for $\leq_{\text{clt}}$ to be sound.

Counterexample 3 Let $\mathcal{G}_{\text{bad}}$ 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 $\mathcal{G}_{\text{bad}}$, and then proceed to show that this pair is not contained in $\leq_{\text{clt}}$. Let $R = \{(1, \tau.1)\}$. It follows that $R \subseteq \mathcal{G}_{\text{bad}}(R)$ if all the conditions for $\mathcal{G}_{\text{bad}}$ are satisfied: condition (1) in is trivially true, condition (3a) is true because $0 \text{ must } 1$ and $0 \text{ must } \tau.1$, condition (3b) holds trivially because $\text{Acc}_{\text{clt}}(\tau.1) = \emptyset$, 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, $1 \not\leq_{\text{clt}} \tau.1$ because the divergent server τ^∞ distinguishes between the two clients: whereas $\tau^\infty \text{ must } 1$ since the client succeeds immediately, we have $\tau^\infty \not\text{ must } \tau.1$ because the composition $\tau.1 \parallel \tau^\infty$ has an infinite unsuccessful computation due to the divergence of τ^∞ . $\blacksquare$

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 Section 1). Using the standard weaker condition makes the preorder $\leq_{\text{clt}}$ unsound wrt. $\leq_{\text{clt}}$, as we proceed to show in the next example.

Counterexample 4 Let $\mathcal{G}_{\text{bad}}$ 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 $\mathcal{G}_{\text{bad}}$, $\leq_{\text{clt}}^{\text{bad}}$, contains client pairs that are not in $\leq_{\text{clt}}$.

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

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 $r_6 \not\leq_{\text{clt}} r_7$, because $\bar{c}.0 \text{ must } r_6$ whereas $\bar{c}.0 \not\text{ must } r_7$. On the other hand, we now show that $r_6 \leq_{\text{clt}}^{\text{bad}} r_7$. Focusing on condition Definition 6(3), we start by deducing that $r_6 \in \mathcal{U}_{bhv}$ (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 $0 \text{ must } r_7$, 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 $\text{Acc}_{\mathcal{U}}(r_7) = \text{Acc}_{\mathcal{U}}(r_6) = \{\{a\}\}$.⁴ To prove that the (relaxed) condition (3bad) holds, we have to show that

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

Let $r_7' = a.1 + b.0$. We only show the proof for the inequality $r_6^c \leq_{\text{clt}}^{\text{bad}} r_7'$, 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 $\text{Acc}_{\mathcal{U}}(r_7') = \{\{c\}\}$ and $\text{Acc}_{\mathcal{U}}(r_6^c) = \{\{b\}, \{c\}\}$. Finally for (3bad) we only have to check the case for $r_7' \xrightarrow{b} 0$, which requires us to show that $\tau.0 \leq_{\text{clt}}^{\text{bad}} 0$; this latter check is routine. As a result, we have $r_6^c \leq_{\text{clt}}^{\text{bad}} r_7'$. Since we can also show that $r_6^c \leq_{\text{clt}}^{\text{bad}} a.0 + b.1$ holds, we obtain (7), and consequently $r_6 \leq_{\text{clt}}^{\text{bad}} r_7$. ■

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

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

Theorem 5. In any finite branching LTS $r_1 \sqsubseteq_{\text{clt}} r_2$ if and only if $r_1 \leq_{\text{clt}} r_2$.

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

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 $\leq_{\text{clt}}$ is still a burdensome method for reasoning on $\sqsubseteq_{\text{clt}}$. By contrast, We now contend that it is simpler to show $r_4 \sqsubseteq_{\text{clt}} r_5$ by proving $r_4 \leq_{\text{clt}} r_5$, 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 $\text{Acc}_{\mathcal{U}}(r_5) = \text{Acc}_{\mathcal{U}}(r_5') = \{\{b, c\}\}$ and that $\text{Acc}_{\mathcal{U}}(r_4) = \{\{a, b, c\}\}$. 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. ■

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

⁴ The restriction of the left hand side of the inclusion of Definition 6(3b) by $\text{ua}_{\text{bhv}}(r_6)$ is superfluous.

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 $r_6 \not\sqsubseteq_{\text{clt}} r_7$ 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 $r_6 \in \mathcal{U}_{\text{bhv}}$, $r_7 \xrightarrow{c} \not\sqsubseteq$ and $r_6 \xRightarrow{c} \not\sqsubseteq$, Definition 6(3c) requires that we show that

$$r_6^c R \tau.r'_7 + \tau.r''_7 \text{ where } r_6^c = \bigoplus (r_6 \text{ after } \not\sqsubseteq c) \text{ and } (\tau.r'_7 + \tau.r''_7) = \bigoplus (r_7 \text{ after } \not\sqsubseteq c) \quad (8)$$

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 $r_6 \leq_{\text{clt}}^{\text{bad}} r_7$, (8) is instrumental to conclude that $r_6 \not\sqsubseteq_{\text{clt}} r_7$. 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 $r_6 \not\sqsubseteq_{\text{clt}} r_7$. ■

5 Expressiveness and Decidability

We show that servers with finite interactions suffice to preserve the discriminating power of the contextual preorder $\sqsubseteq_{\text{clt}}$ 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 $\sqsubseteq_{\text{clt}}$.

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 $\sqsubseteq_{\text{clt}}$ of Definition 1. Let $\text{CCS}^f ::= 0 \mid 1 \mid \alpha.p \mid p + q \mid \tau^\infty$, and

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

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, $\sqsubseteq_{\text{clt}}^f$ 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 $r_1 \sqsubseteq_{\text{clt}}^f r_2$ if and only if $r_1 \sqsubseteq_{\text{clt}} r_2$.*

Proof. The inclusion $\sqsubseteq_{\text{clt}} \subseteq \sqsubseteq_{\text{clt}}^f$ follows immediately from the resp. definitions. On the other hand, Theorem 5 provides us with a proof technique for showing the inclusion $\sqsubseteq_{\text{clt}}^f \subseteq \sqsubseteq_{\text{clt}}$: if we show that $\sqsubseteq_{\text{clt}}^f \subseteq \mathcal{G}(\sqsubseteq_{\text{clt}}^f)$ then $\sqsubseteq_{\text{clt}}^f \subseteq \preceq_{\text{clt}} = \sqsubseteq_{\text{clt}}$. In view of the Knaster-Tarski theorem it suffices to show that $\sqsubseteq_{\text{clt}}^f \subseteq \mathcal{G}(\sqsubseteq_{\text{clt}}^f)$. 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 exist 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 $\mathcal{U}_{\text{bhv}}$, 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, $\text{Acc}_{\mathcal{U}}(r) \neq \emptyset$, 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 Figure 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 Figure 2. Concretely, we define the 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

$$\text{sumsRdx}(r) = \{ \bigoplus (r \text{ after } s) \mid \text{for some } s \in \text{Act}^* \},$$

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

```

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 $\mathcal{U}$

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 Figure 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 $r \in \mathcal{U}$) 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 Theorem 5, Theorem 7 and Lemma 7, we conclude that Definition 6 can be used to decide $\sqsubseteq_{\text{clt}}$ for languages such as CCS^μ of Figure 2. We can do this by adapting the algorithm of [27, Chapter 21.5], and proving that in our setting [27, Theorem 21.5.9 and Theorem 21.5.12] are true. In particular, using the terminology of [27] we have that $\text{reachable}_{\mathcal{G}}(X)$ is finite, essentially because the *resp.* LTS is finite-state, and thus the decidability of $\sqsubseteq_{\text{clt}}$ 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-abstraction for the peer preorder as well. Our investigations and the *resp.* proofs for Theorem 2, Theorem 5 and Theorem 6 are conducted in terms of finitely-branching LTSs, which cover the semantics used by numerous other work describing client and server contracts [8,18,9,6] — 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^u . 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 rewritten 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 $r \xrightarrow{\tau}$ 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 [28,8,24].

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. J. R. Andersen, N. Andersen, S. Enevoldsen, M. M. Hansen, K. G. Larsen, S. R. Olesen, J. Srba, and J. Wortmann. CAAL: concurrency workbench, aalborg edition. In *ICTAC*, 2015.
2. F. Barbanera and F. de'Liguoro. Two notions of sub-behaviour for session-based client/server systems. In *PPDP*, 2010.
3. G. Bernardi. *Behavioural Equivalences for Web Services*. PhD thesis, TCD, 2013.
4. G. Bernardi and A. Francalanza. Full-abstraction for must testing preorders (extended abstract). Available from <https://www.irif.fr/~gio/papers/BFcoordination2017.pdf>.
5. G. Bernardi and M. Hennessy. Modelling session types using contracts. In *SAC*, 2012.
6. G. Bernardi and M. Hennessy. Mutually testing processes. *LMCS*, 11(2), 2015.
7. G. Bernardi and M. Hennessy. Using higher-order contracts to model session types. *LMCS*, 12(2), 2016.
8. M. Bravetti and G. Zavattaro. A foundational theory of contracts for multi-party service composition. *Fundam. Inf.*, 89(4), 2008.
9. G. Castagna, N. Gesbert, and L. Padovani. A theory of contracts for web services. *ACM Trans. Program. Lang. Syst.*, 31(5), 2009.
10. A. Cerone and M. Hennessy. Process behaviour: Formulae vs. tests. In *EXPRESS*, 2010.
11. E. M. Clarke and H. Veith. Counterexamples revisited: Principles, algorithms, applications. In *Verification: Theory and Practice*, 2003.
12. R. De Nicola and M. Hennessy. Testing equivalences for processes. *TCS*, 34(1–2), 1984.
13. A. Francalanza. A Theory of Monitors. In *FoSSaCS*, LNCS, 2016.
14. S. J. Gay and M. Hole. Subtyping for session types in the pi calculus. *Acta Inf.*, 42(2-3), 2005.
15. M. Hennessy. *Algebraic Theory of Processes*. 1988.
16. D. E. Knuth. *The Art of Computer Programming, Volume 1 (3rd Ed.): Fundamental Algorithms*. Addison Wesley Longman Publishing Co., Inc., 1997.
17. D. König. Über eine schlussweise aus dem endlichen ins unendliche. *Acta Litt. ac. sci. Szeged*, 3, 1927.
18. C. Laneve and L. Padovani. The must preorder revisited. In *CONCUR*, 2007.
19. Q. Luo, F. Hariri, L. Eloussi, and D. Marinov. An empirical analysis of flaky tests. In *FSE*, 2014.
20. P. Marinescu, P. Hosek, and C. Cadar. Covrig: A framework for the analysis of code, test, and coverage evolution in real software. In *ISSTA*, 2014.
21. A. Martens. Analyzing Web Service Based Business Processes. In *FASE*, 2005.
22. A. M. Memon and M. B. Cohen. Automated testing of gui applications: Models, tools, and controlling flakiness. In *ICSE*, 2013.
23. R. Milner. *Communication and Concurrency*. Prentice-Hall, 1989.
24. A. J. Mooij, C. Stahl, and M. Voorhoeve. Relating fair testing and accordance for service replaceability. *J. Log. Algebr. Program.*, 79(3-5), 2010.
25. L. Padovani. Contract-based discovery of web services modulo simple orchestrators. *TCS*, 411(37), 2010.
26. L. Padovani. Fair subtyping for multi-party session types. *MSCS*, 26(3), 2016.
27. B. Pierce. *Types and programming languages*. 2002.
28. A. Rensink and W. Vogler. Fair testing. *Information and Computation*, 205(2), 2007.
29. C. Spaccasassi. *Language Support for Communicating Transactions*. PhD thesis, TCD, 2015.
30. D. Weinberg. Efficient controllability analysis of open nets. In *WS-FM*, 2009.
31. G. Winskel. *The Formal Semantics of Programming Languages: An Introduction*. 1993.