

Bisimulations for Communicating Transactions^{*}

(Extended Abstract)

Vasileios Koutavas, Carlo Spaccasassi^{**}, and Matthew Hennessy

Trinity College Dublin

Abstract. We develop a theory of bisimulations for a simple language containing communicating transactions, obtained by dropping the isolation requirement of standard transactions. Such constructs have emerged as a useful programming abstraction for distributed systems.

In systems with communicating transactions actions are tentative, waiting for certain transactions to commit before they become permanent. Our theory captures this by making bisimulations history-dependent, in that actions performed by transactions need to be recorded. The main requirement on bisimulations is the systems being compared need to match up exactly in the permanent actions but only those.

The resulting theory is fully abstract with respect to a natural contextual equivalence and, as we show in examples, provides an effective verification technique for comparing systems with communicating transactions.

1 Introduction

Communicating transactions, obtained by dropping the isolation requirement of standard transactions, is a novel and powerful programming construct for distributed systems. For example, it can be used to simplify the programming of complex concurrent consensus scenarios, avoiding the use of locks and explicit error handling [16]. Variants of such constructs have been proposed as extensions to programming languages [5,7,11,16] and process calculi [2,1,3]. However, before they can be adopted in mainstream programming, significant research is needed in efficient implementation strategies [6,11,16], programming paradigms [1,7], and viable verification techniques [8]. The last concern is the topic of this paper.

Bisimulations [12] provide an elegant and effective proof technique for proving equivalences between processes in a variety of settings (e.g., [15,14]). They are essentially defender strategies in a game where a challenger attempts to discover a difference in the extensional behaviour of two processes, while the defender tries to refute these attempts [17]. For example, consider the standard CCS processes

$$P_1 = a.(b.\mathbf{0} + c.\mathbf{0}) \qquad Q_1 = a.b.\mathbf{0} + a.c.\mathbf{0} \qquad (1)$$

^{*} This research was supported by SFI project SFI 06 IN.1 1898.

^{**} Supported by MSR (MRL 2011-039).

which can perform the action a followed by either b or c , with a slight difference in when this choice is taken. The challenger chooses Q_1 to perform the action a , with residual $Q'_1 = b.\mathbf{0}$ to which the defender must respond with a matching a action from P_1 ; the only possibility is for P_1 to perform a with residual $P'_1 = b.\mathbf{0} + c.\mathbf{0}$. But now the challenger chooses P'_1 to perform the c action, to which the defender has no response. The defender loses the game. In fact there is no possible winning strategy for the defender in this game and thus no bisimulation containing the pair (P_1, Q_1) . Therefore P_1 and Q_1 are deemed to be behaviourally distinct.

However, the appropriate notion of the bisimulation game for a language with communicating transactions or similar constructs is *a priori* unclear. An objective criterion for a potential bisimulation game is relation with contextual equivalence. If the existence of a winning defender's strategy in a game over two systems implies that the systems are contextually equivalent then this game is a *sound* verification technique. If the absence of such a strategy implies that the two systems are contextually inequivalent then the game is a *complete* technique. The main result of this paper is a weak bisimulation theory for a simple language containing communicating transactions, $TCCS^m$, which provides a sound and complete proof methodology with respect to a natural contextual equivalence.

$TCCS^m$ is obtained from CCS [12] by essentially adding one construct $\llbracket P \triangleright_k Q \rrbracket$ for communicating transactions, and a new command co for committing them. Here k is the name of the transaction, P the body which is expected to be completed in its entirety or not at all, and Q is the alternative, to be executed in case the transaction is aborted. However it should be emphasised that if an abort occurs not only is Q launched but P and all its effects on the environment are rolled back. For example consider the systems with p fresh from R and S :

$$\begin{aligned} P_2 &= \nu p. \llbracket a.p.\text{co}.R \triangleright_k \mathbf{0} \rrbracket \parallel \llbracket b.\bar{p}.\text{co}.S \triangleright_l \mathbf{0} \rrbracket \\ Q_2 &= \llbracket a.b.\text{co}.(R \mid S) + b.a.\text{co}.(R \mid S) \triangleright_m \mathbf{0} \rrbracket \end{aligned} \quad (2)$$

Here P_2 consists of two independent transactions which co-operate by synchronising on a private channel p . If after this synchronisation the left-hand transaction aborts then the effect of the a action, which is a communication with the environment, must be rolled back. But the synchronisation on p must also be undone, and therefore, because of the all-or-nothing nature of transactions, the effects of the b action in the right-hand transaction will be rolled back. Indeed in our reduction semantics, given in Sect. 2, this right-hand transaction is also aborted. Because of the synchronisation on p , the destiny of both transactions is conjoined. For this reason we should be able to demonstrate that P_2 is behaviourally equivalent to the single transaction Q_2 .

The standard bisimulation game outlined above cannot be easily extended to $TCCS^m$. Many actions, such as a, b in (2) above, are tentative, in the sense that their effect can only be considered to be permanent when the transactions k and l commit, if ever. To underline this consider the processes:

$$P_3 = \llbracket a.(b.\text{co} + c.\mathbf{0}) \triangleright_k \mathbf{0} \rrbracket \quad Q_3 = \llbracket a.b.\text{co} \triangleright_l \mathbf{0} \rrbracket \quad (3)$$

Here P_3 commits only after performing the a, b . It would be unreasonable for the challenger, after the a action, to demand a response to c . Not only is this action

tentative on the completion of transaction k but also if c is performed then k will *never* commit; so the defender should be able to ignore this challenge.

Our approach is to play the bisimulation game on *configurations*, of the form $\mathcal{C} = (H \triangleright P)$ where P is a system and H a *history* of all the tentative actions taken so far in the game by P . These are of the form $k(a)$, where k is the name of a transaction which needs to commit before the action becomes a permanent a . When playing bisimulation moves, the histories of both systems being scrutinised must remain *consistent*, in that permanent actions in the respective histories must match exactly (for simplicity old permanent actions are not garbage collected). The crucial aspect of this new game is that when a system commits a transaction k , and only then, all tentative actions in its history dependent on k are made permanent. This consistency requirement then forces a response in which the corresponding actions match exactly.

For example consider the following variation on (1), using transactions:

$$P_4 = \llbracket a.(b.\text{co} + c.\text{co}) \triangleright_l \mathbf{0} \rrbracket \quad Q_4 = \llbracket a.b.\text{co} + a.c.\text{co} \triangleright_k \mathbf{0} \rrbracket \quad (4)$$

Replaying the game from (1), where the challenger first chooses a from Q_4 and then c from P_4 with the same responses, we reach the configurations

$$\mathcal{C}_4 = (k(a), k(c) \triangleright \llbracket \text{co} \triangleright_k \mathbf{0} \rrbracket) \quad \mathcal{D}_4 = (l(a), l(b) \triangleright \llbracket \text{co} \triangleright_l \mathbf{0} \rrbracket)$$

At this stage the two histories are still consistent as they contain no permanent actions. However, now there is no possible response when the challenger chooses the commit move $\mathcal{C}_4 \rightarrow \mathcal{C}'_4 = (a, c \triangleright \mathbf{0})$. This is a silent move from \mathcal{C}_4 in which transaction k commits, making the two actions in the history permanent. There are various ways in which \mathcal{D}_4 can try to respond but all lead to an inconsistent history. Thus, with our version of bisimulations P_4 and Q_4 are not bisimilar.

Note that such a successful attack by the challenger cannot be mounted for (3) above. After one round in the game we have the configurations

$$\mathcal{C}_3 = (k(a) \triangleright \llbracket b.\text{co} + c.\mathbf{0} \triangleright_k \mathbf{0} \rrbracket) \quad \mathcal{D}_3 = (l(a) \triangleright \llbracket b.\text{co} \triangleright_l \mathbf{0} \rrbracket)$$

and since \mathcal{D}_3 has no possible c actions the challenger might request a response to the action $\mathcal{C}_3 \rightarrow \mathcal{C}'_3 = (k(a), k(c) \triangleright \llbracket \mathbf{0} \triangleright_k \mathbf{0} \rrbracket)$. However the tentative $k(c)$ recorded in the history will never become permanent and thus the defender can successfully respond with any tentative action of \mathcal{D}_3 which will also never become permanent. To ensure that such responses can always be made, our bisimulations will allow *any* configuration to make the degenerate silent move $(H \triangleright P) \rightarrow (H, k(\star) \triangleright P)$, where \star is a reserved symbol. This defender move represents a weak idle move whose validity is postponed until after k commits. So \mathcal{C}_3 's move above can be matched by \mathcal{D}_3 playing this degenerate move followed by the abort of the transaction. Later we prove that P_3 and Q_3 are indeed bisimilar.

Using recursion we can write *restarting transactions* as $\text{rec}X.\llbracket P \blacktriangleright X \rrbracket$. Here $\llbracket P \blacktriangleright X \rrbracket$ is an uninitiated transaction which has yet to be allocated a name. These transactions can abort and re-run internal steps, thus branching differences in initial silent actions can be hidden from the challenger. Consider a compiler performing common subexpression elimination, transforming P_5 to Q_5 :

$$P_5 = \text{rec}X. \llbracket \tau.b.\text{co} + \tau.c.\text{co} \blacktriangleright X \rrbracket \quad Q_5 = \text{rec}X. \llbracket \tau.(b.\text{co} + c.\text{co}) \blacktriangleright X \rrbracket \quad (5)$$

As we will show, all moves of the challenger can be matched by the defender in the bisimulation game. The interesting scenario is when (after initiating the transactions) the challenger picks the right τ action from P_5 and the defender responds with the τ action from Q_5 . We then get the configurations:

$$\mathcal{C}_5 = (\varepsilon \triangleright \llbracket c.\text{co} \triangleright_k P_5 \rrbracket) \quad \mathcal{D}_5 = (\varepsilon \triangleright \llbracket b.\text{co} + c.\text{co} \triangleright_l Q_5 \rrbracket)$$

The challenger then picks the b action from \mathcal{D}_5 . The defender responds with a silent abort of \mathcal{C}_5 which will reinstate P_5 , re-initialise the transaction, and select the left τ action in \mathcal{P}_5 followed by the b action. This would lead to the configurations $\mathcal{C}'_5 = (k'(b) \triangleright \llbracket \text{co} \triangleright_{k'} P_5 \rrbracket)$ and $\mathcal{D}'_5 = (l(b) \triangleright \llbracket \text{co} \triangleright_l Q_5 \rrbracket)$. We will in fact prove that this optimisation is sound in our setting, although it is not sound in the case of P_4 , Q_4 , even if we used restarting transactions.

In the remainder of this extended abstract we explain the language $TCCS^m$ (Sect. 2), which is a simplification of that used in previous work [4] in that we do not consider nested transactions, simplifying technical development. Inspired by cJoin [1], when transactions co-operate by synchronising on an action they are virtually *merged* by acquiring the same name. Thus in P_2 from (2), when synchronisation occurs on the private channel p , the residual will be a term equivalent to the single transaction $\nu p. \llbracket \text{co}.(R | S) \triangleright_n \mathbf{0} | \mathbf{0} \rrbracket$.

This is followed by an exposition of our history dependent bisimulations (Sect. 3). As we have already stated, these bisimulations demand the appropriate matching of all actions, even those dependent on transactions which can never commit. We also give a variation, called *predictive bisimulations*, in which dependent actions need only be matched when the transaction on which they depend has some future possibility of committing (Sect. 4). We ultimately show that both equivalences coincide but the former is an easier proof technique while the latter is easier to prove sound. We then outline the proof of the main result of the paper, namely that these bisimulation equivalences coincide with contextual equivalence (Sect. 5).

2 The Language $TCCS^m$

The syntax for terms in the language is given in Fig. 1 where $a \in \text{Act}$ are actions, $\mu \in \text{Act} \uplus \{\tau\} \uplus \Omega$ are prefixes, and X ranges over a collection of *recursion variables*. Here $\omega \in \Omega$ are special actions which will be used to define contextual equivalence, while $(\bar{\tau}) : \text{Act} \rightarrow \text{Act}$ is a total bijection over Act , used in the standard manner to formalise CCS synchronisation between processes. The language contains all the standard constructs of CCS, with the result that CCS is a sub-language of $TCCS^m$. There are three extra operators, discussed in the Introduction. We assume the standard notion of free and bound occurrence of recursion variables, and only consider closed terms, those which contain no free occurrences of variables. We use the standard abbreviations associated

TCCS^m Syntax		
$P, Q, R ::= \sum \mu_i.P_i \mid P Q \mid \nu a.P \mid X \mid \mathbf{rec}X.P$ $\mid \llbracket P \triangleright_k Q \rrbracket \mid \mathbf{co}.P \mid \llbracket P \blacktriangleright Q \rrbracket$		
CCS Transitions		
CCSSUM $\frac{}{\sum \mu_i.P_i \xrightarrow{\mu_i}_\varepsilon P_i}$	CCSSYNC $\frac{P \xrightarrow{\alpha}_\varepsilon P' \quad Q \xrightarrow{\bar{\alpha}}_\varepsilon Q'}{P Q \xrightarrow{\tau}_\varepsilon P' Q'}$	CCSREC $\frac{}{\mu X.P \xrightarrow{\tau}_\varepsilon P[\mu X.P/X]}$
Transactional Transitions		
TRTAU $\frac{P \xrightarrow{\tau}_\varepsilon P'}{\llbracket P \triangleright_k Q \rrbracket \xrightarrow{\tau}_\varepsilon \llbracket P' \triangleright_k Q \rrbracket}$	TRSUM $\frac{}{\sum \mu_i.P_i \xrightarrow{k(a)}_{\varepsilon \mapsto k} \llbracket P_j \mid \mathbf{co} \triangleright_k \sum \mu_i.P_i \rrbracket} \mu_j = a$	
TRACT $\frac{P \xrightarrow{\alpha}_\varepsilon P'}{\llbracket P \triangleright_l Q \rrbracket \xrightarrow{k(a)}_{l \mapsto k} \llbracket P' \triangleright_k Q \rrbracket} k \# l$	TRSYNC $\frac{P \xrightarrow{k(a)}_{\sigma_1} P' \quad Q \xrightarrow{k(\bar{a})}_{\sigma_2} Q' \quad \sigma_1 = \tilde{l}_1 \mapsto k \quad \sigma_2 = \tilde{l}_2 \mapsto k}{P Q \xrightarrow{k(\tau)}_{(\tilde{l}_1, \tilde{l}_2) \mapsto k} P'\sigma_2 Q'\sigma_1}$	
Propagation Transitions		
RESTR $\frac{P \xrightarrow{\alpha}_\sigma P'}{\nu a.P \xrightarrow{\alpha}_\sigma \nu a.P'} a \notin \alpha$	PARL $\frac{P \xrightarrow{\alpha}_\sigma P'}{P Q \xrightarrow{\alpha}_\sigma P' Q\sigma} \mathit{range}(\sigma) \# Q$	

Fig. 1. Communication and internal transitions (omitting symmetric rules)

with CCS, and write $s \# s'$ when the transaction names of the syntax object s are fresh from those in s' ; $\mathit{fn}(s)$ denotes the transaction names in s . Note that unlike previous work [3,4] transaction names are never bound and we do not require that all transaction names used in a term are distinct. Thus, we allow terms of the form $\llbracket P_1 \triangleright_k P_2 \rrbracket \mid R \mid \llbracket Q_1 \triangleright_k Q_2 \rrbracket$. Here k should be looked upon as a *distributed* transaction whose behaviour will be approximately the same as the centralised $\llbracket P_1 \mid Q_1 \triangleright_k P_2 \mid Q_2 \rrbracket$. The use of these distributed transactions will simplify considerably the exposition of the reduction semantics.

Definition 2.1. *A closed term is called well-formed if in every occurrence of $\llbracket P \triangleright_k Q \rrbracket$, $\llbracket P \blacktriangleright Q \rrbracket$, and $\mathbf{rec}X.P$, the subterms P and Q do not contain named transactions of the form $\llbracket - \triangleright_k - \rrbracket$. We refer to well-formed terms as processes.*

Note that dormant transactions *can* appear within other transactions and under recursion but they will be activated only when they end up at top-level. In the sequel we only consider well-formed terms.

The reduction semantics of the language is given as a binary relation between processes $P \rightarrow Q$. However this is defined indirectly in terms of three auxiliary relations, which will also be used in the formulation of bisimulations:

$$P \rightarrow Q \quad \text{when} \quad P \xrightarrow{\tau}_\sigma Q \quad \text{or} \quad P \xrightarrow{\beta} Q \quad \text{or} \quad P \xrightarrow{k(\tau)}_\sigma Q \quad (6)$$

TRNEW <hr style="border: 1px solid black;"/> $\frac{[[P \blacktriangleright Q]] \xrightarrow{\text{new } k} [[P \triangleright_k Q]]}{\text{TRBCAST}}$ $\frac{P \xrightarrow{\beta} P' \quad Q \xrightarrow{\beta} Q' \quad \beta \in \{\text{co } k, \text{ab } k\}}{P Q \xrightarrow{\beta} P' Q'}$	TRAB <hr style="border: 1px solid black;"/> $[[P \triangleright_k Q]] \xrightarrow{\text{ab } k} Q$ $\frac{P \xrightarrow{\beta} P' \quad \beta \# Q}{P Q \xrightarrow{\beta} P' Q}$	TRCO $\frac{P \rightsquigarrow_{\text{co}} P'}{[[P \triangleright_k Q]] \xrightarrow{\text{co } k} P'}$ TRRESTR $\frac{P \xrightarrow{\beta} P'}{\text{va. } P \xrightarrow{\beta} \text{va. } P'}$
---	--	---

Fig. 2. Transactional reconfiguration transitions

The first, $P \xrightarrow{\tau}_\sigma Q$, is essentially synchronisation between pure CCS processes. The second, $P \xrightarrow{\beta} Q$ (Fig. 2), where β ranges over $\text{co } k$, $\text{ab } k$ and $\text{new } k$, encode the creation of new named transactions, and commit/abort broadcast transitions (TRCO , TRAB , TRBCAST) which eliminate distributed transactions. The notation $P \rightsquigarrow_{\text{co}} P'$ means the execution of a top-level co in P and the replacement of all other top-level commits with τ -prefixes. We now concentrate on the third, $P \xrightarrow{k(\tau)}_\sigma Q$, or more generally $P \xrightarrow{k(\mu)}_\sigma Q$ where $\mu \in \text{Act} \cup \{\tau\}$.

Action $P \xrightarrow{k(a)}_\sigma Q$ should be viewed as the synchronisation between P and some transaction named k in the environment which can perform the complementary \bar{a} . Because this transaction is external the freshness side conditions in the rules of Fig. 1 ask that the name k is fresh with respect to P . Also, the effect of this synchronisation is that the future behaviour of P , or at least any transactions involved in the execution of a , is dependent on the eventual committing of k . This dependency is implemented by σ , a substitution renaming the responsible transaction in P to k . The essential rule in the generation of these judgements is TRACT in Fig. 1. For example, this rule ensures that we can derive $[[a.P_1 \triangleright_{l_1} Q_1]] \xrightarrow{k(a)}_{l_1 \mapsto k} [[P_1 \triangleright_k Q_1]]$ for any fresh k . The substitution recorded in the action is propagated by PARL into contexts. Note that by TRSUM , even pure CCS processes with no transactions can perform a $k(a)$ action; e.g., $a.P \xrightarrow{k(a)}_{\varepsilon \mapsto k} [[P | \text{co } \triangleright_k a.P]]$. This embeds P into the k -transaction; the distributed part of the k -transaction surrounding P is always ready to commit (hence the introduction of co). Note that this is a *communication-driven embedding*, which reduces the nondeterminism of embedding of previous work [3,4], making semantics more concrete [16]. Embedding leads to a uniform treatment of CCS processes and transactions, and a simple reduction semantics.

The conjoining of transactions is implemented in TRSYNC . Using it we infer:

$$[[a.P_1 \triangleright_{l_1} Q_1]] \mid [[\bar{a}.P_2 \triangleright_{l_2} Q_2]] \xrightarrow{k(\tau)}_{(l_1, l_2) \mapsto k} [[P_1 \triangleright_k Q_1]] \mid [[P_2 \triangleright_k Q_2]]$$

for any fresh k . Here the previously independent transactions l_1, l_2 have been merged into the transaction k (recorded in the substitution $(l_1, l_2) \mapsto k$). Note that this new transaction is distributed, in that its activity is divided in two. In order for it to commit the rules in Fig. 2 ensure that *both* components commit.

Example 2.2. Consider the process P_2 defined in (2) in the Introduction. Two applications of TRACT followed by rule RESTR gives the reduction from (6) above

$$P_2 \mid \bar{a}.b \rightarrow^* \nu p. \llbracket p.\text{co}.R \triangleright_{k_1} \mathbf{0} \rrbracket \mid \llbracket \bar{p}.\text{co}.S \triangleright_{k_2} \mathbf{0} \rrbracket = P'_2$$

where k_1 and k_2 are fresh names. The synchronisation rule TRSYNC then gives

$$P'_2 \rightarrow \nu p. \llbracket \text{co}.R \triangleright_k \mathbf{0} \rrbracket \mid \llbracket \text{co}.S \triangleright_k \mathbf{0} \rrbracket$$

where k is an arbitrary fresh name. Here the residual is a single transaction named k , albeit distributed. For it to commit both components have to commit: using TRBCAST this leads to the process $R \mid S$. \square

The semantics has a number of properties: it preserves well-formedness, generates only fresh transaction names and is equivariant. The properties about transaction names are important because they give us the liberty to pick fresh enough transaction names in proofs. To state these properties we use *renamings*, ranged over by r , which are bijective substitutions of the form $\{l_1/k_1, \dots, l_n/k_n\}$. The range of a *fresh* renaming r_{fr} has names not appearing in the proof.

Lemma 2.3 (Names). *Suppose $P \xrightarrow{l(\mu)}_{\sigma} Q$. Then*

1. l is fresh to P , Q is well-formed, and the substitution σ has the form $\tilde{k} \mapsto l$;
2. (Equivariance) $P r_{\text{fr}} \xrightarrow{l'(\mu)}_{\sigma r_{\text{fr}}} Q r_{\text{fr}}$, where $r_{\text{fr}}(l) = l'$. \square

Based on this semantics we give a natural contextual equivalence, using standard formulations [15]. We write \Rightarrow for the reflexive transitive closure of \rightarrow .

Definition 2.4 (Barb). $P \Downarrow \omega$ ($\omega \in \Omega$) if $\exists Q, Q'$ such that $P \Rightarrow Q \xrightarrow{\omega}_{\epsilon} Q'$.

Definition 2.5 (Reduction Barbed Equivalence (\cong_{rbe})). (\cong_{rbe}) is the largest relation for which $P \cong_{\text{rbe}} Q$ when:

1. $P \Downarrow \omega$ iff $Q \Downarrow \omega$,
2. if $P \rightarrow P'$ then there exists Q' such that $Q \Rightarrow Q'$ and $P' \cong_{\text{rbe}} Q'$,
3. if $Q \rightarrow Q'$ then there exists P' such that $P \Rightarrow P'$ and $P' \cong_{\text{rbe}} Q'$,
4. $P \mid R \cong_{\text{rbe}} Q \mid R$ for any R with $R \nmid P, Q$.

Here we consider contexts with fresh transaction names to enforce that observer transactions are distinct from process transactions before communication occurs. If this was not the case then transaction names would be observable: $\llbracket P \triangleright_k Q \rrbracket$ would not be equivalent to $\llbracket P \triangleright_l Q \rrbracket$ because by introducing the context $\llbracket \mathbf{0} \triangleright_k \mathbf{0} \rrbracket$ the k -transaction can no longer commit but l still can. Thus, *all* transaction names are considered local here; the side condition $R \nmid P, Q$ enforces this without the syntactic overhead of a ν -binder for all transaction names.

To see why in the above definition we use barbs from a distinct Ω consider:

$$P = \llbracket a.\text{co} \triangleright_k \mathbf{0} \rrbracket \qquad Q = a.\mathbf{0} + \tau.\mathbf{0}$$

Intuitively, we would expect P to have exactly the same behaviour as Q , eventually executing the single action a , or failing with a τ step. But if we allowed the barb $\Downarrow a$ in Thm. 2.5 then they would not be equivalent because $P \Downarrow a$ and $Q \not\Downarrow a$.

Example 2.6. P_4 and Q_4 from (4) in the Introduction are indeed inequivalent.

Assume $P_4 \cong_{\text{rbe}} Q_4$. Take $C_1 = \bar{a}$. Then $P_4 | C_1 \cong_{\text{rbe}} Q_4 | C_1$. We have:

$$Q_4 | C_1 \xrightarrow{l(\tau)} \llbracket [b.\text{co} \triangleright_l \mathbf{0}] \rrbracket \llbracket [\text{co} \triangleright_l C_1] \rrbracket = Q'_4 \quad \text{thus} \quad Q_4 | C_1 \rightarrow Q'_4$$

Process Q'_4 should be equivalent to $P_4 | C_1$ or one of its successors:

1. $P_4 | C_1$. Let $C_2 = \bar{c}.\omega$; then $Q'_4 | C_2 \not\Downarrow \omega$, $P_4 | C_1 | C_2 \Downarrow \omega$. Thus $P_4 | C_1 \not\cong_{\text{rbe}} Q'_4$.
2. $P'_4 = \llbracket [b.\text{co} + c.\text{co} \triangleright_m \mathbf{0}] \rrbracket \llbracket [\mathbf{0} \triangleright_m C_1] \rrbracket$. Again, $P'_4 | C_2 \Downarrow \omega$, thus $P'_4 \not\cong_{\text{rbe}} Q'_4$.
3. C_1 (after an abort). $Q'_4 | \bar{b}.\omega \Downarrow \omega$ but $C_1 | \bar{b}.\omega \not\Downarrow \omega$, so $C_1 \not\cong_{\text{rbe}} Q'_4$.

Thus $Q'_4 \not\cong_{\text{rbe}} P_4 | C_1$ or any later state (contradiction), and $P_4 \not\cong_{\text{rbe}} Q_4$. \square

Note that the difference in the branching structure of P_4 and Q_4 is not observable by the may- and must-testing equivalences [3,4]. These equivalences are characterised by so-called *clean traces*, which are traces in which all tentative actions are committed. If bisimulations were developed using such clean traces, P_4 and Q_4 would also be identified in the resulting bisimulation theory but it would not correspond to a natural definition of (\cong_{rbe}) .

3 Bisimulations

As mentioned in the Introduction, our bisimulations will be over configurations $(H \triangleright P)$ with a process P and a history H of the tentative interactions of P with its environment. An element of such a history can be a $k(a)$, a , ab , $k(\star)$, or \star . A past tentative action $k(a)$ that has not been committed or aborted is recorded as is in the history. If the k -transaction that performed this action commits, the action becomes a ; if k aborts, it becomes ab . Histories also record the trivial actions $k(\star)$ which can be performed by any process. If k commits, $k(\star)$ -recordings become \star , which terminates the bisimulation game in favour of the attacker; if k aborts they become ab . For technical convenience in proofs, elements in a history are uniquely indexed and permanent actions are not garbage-collected.

Definition 3.1 (History). A history H is a partial function from objects i of a countable set I to the set $\{a, \star, k(a), k(\star), \text{ab} \mid a \in \text{Act}\}$.

We often write histories as *lists*, omitting the indices of their elements. History composition, written as H_1, H_2 , is defined when $\text{dom}(H_1) \cap \text{dom}(H_2) = \emptyset$. We also let \hat{a} and \hat{b} range over $\text{Act} \cup \{\star\}$ and $\hat{\mu}$ range over $\text{Act} \cup \Omega \cup \{\tau, \star\}$. To express the effect of commits and aborts to histories we define the following operations.

Definition 3.2. $H \setminus_{\text{co}} k$ and $H \setminus_{\text{ab}} k$ are the lifting to lists of the operations:

$$\begin{array}{lll} (i \mapsto k(\hat{a})) \setminus_{\text{co}} k & = (i \mapsto \hat{a}) & (i \mapsto k(\hat{a})) \setminus_{\text{ab}} k & = (i \mapsto \text{ab}) \\ (i \mapsto l(\hat{a})) \setminus_{\text{co}} k & = (i \mapsto l(\hat{a})) & (i \mapsto l(\hat{a})) \setminus_{\text{ab}} k & = (i \mapsto l(\hat{a})) & \text{when } k \# l \\ (i \mapsto \hat{a}) \setminus_{\text{co}} k & = (i \mapsto \hat{a}) & (i \mapsto \hat{a}) \setminus_{\text{ab}} k & = (i \mapsto \hat{a}) \\ (i \mapsto \text{ab}) \setminus_{\text{co}} k & = (i \mapsto \text{ab}) & (i \mapsto \text{ab}) \setminus_{\text{ab}} k & = (i \mapsto \text{ab}) \end{array}$$

For the reasons we explained in the Introduction, weak bisimulations for $TCCS^m$ require configurations to agree on the committed actions in their histories, and only those actions. Soundness of our technique will establish this as a sufficient requirement for contextual equivalence between processes.

Definition 3.3 (Consistency). H_1 and H_2 are consistent when they have the same domain and for all $i \in I$, $a \in \text{Act}$: $H_1(i) = a$ iff $H_2(i) = a$.

History consistency is one of the two main requirements for weakly bisimilar configurations; the other is to have the same barbs. Thus the weak bisimulation game for $TCCS^m$ will be over transitions with three simple labels: $\zeta ::= \tau \mid k \mid \omega$ annotating internal (τ), tentative synchronisation (k), and barb (ω) transitions.

Definition 3.4 (Bisimulation Transitions). $C \xrightarrow{\zeta} C'$ is derived by the rules:

$$\begin{array}{lll}
(H \triangleright P) \xrightarrow{\tau} (\sigma(H) \triangleright Q) & \text{if } P \xrightarrow{\tau}_{\sigma} Q & (\text{LTS}_{\tau}) \\
(H \triangleright P) \xrightarrow{\tau} (\sigma(H) \triangleright Q) & \text{if } P \xrightarrow{k(\tau)}_{\sigma} Q \text{ and } k \# H & (\text{LTS}_{k(\tau)}) \\
(H \triangleright P) \xrightarrow{\tau} (H \triangleright Q) & \text{if } P \xrightarrow{\text{new } k}_{\sigma} Q \text{ and } k \# H & (\text{LTS}_{\text{new}}) \\
(H \triangleright P) \xrightarrow{\tau} (H \setminus_{\text{co}} k \triangleright Q) & \text{if } P \xrightarrow{\text{co } k}_{\sigma} Q & (\text{LTS}_{\text{co}}) \\
(H \triangleright P) \xrightarrow{\tau} (H \setminus_{\text{ab}} k \triangleright Q) & \text{if } P \xrightarrow{\text{ab } k}_{\sigma} Q & (\text{LTS}_{\text{ab}}) \\
(H \triangleright P) \xrightarrow{k} (\sigma(H), k(a) \triangleright Q) & \text{if } P \xrightarrow{k(a)}_{\sigma} Q \text{ and } k \# H & (\text{LTS}_{k(a)}) \\
(H \triangleright P) \xrightarrow{k} (H, k(\star) \triangleright P) & \text{if } k \# H, P & (\text{LTS}_{\star}) \\
(H \triangleright P) \xrightarrow{\omega} (\sigma(H) \triangleright Q) & \text{if } P \xrightarrow{\omega}_{\sigma} Q & (\text{LTS}_{\omega})
\end{array}$$

We define $\xrightarrow{\zeta}$ to be $\xrightarrow{\tau}^*$ when $\zeta = \tau$, and $\xrightarrow{\tau} \xrightarrow{\zeta} \xrightarrow{\tau}$ otherwise.

The first five rules encode the $TCCS^m$ reduction semantics of (6) in Sect. 2, updating the history of the configurations accordingly. $\text{LTS}_{k(a)}$ encodes the synchronisation between a transaction in the process and its environment, yielding a fresh transaction k ; this tentative action is recorded in the history. LTS_{ω} encodes top-level barbs and LTS_{\star} records a trivial defender synchronisation move. Weak τ - and k -transitions can always be performed by the defender in the bisimulation game. Moreover, there are no top-level a -transitions because they can always be simulated by a $k(a)$ -transition followed by the commit of k .

Lemma 3.5. Suppose $P \xrightarrow{a}_{\varepsilon} Q$; then $P \xrightarrow{k(a)}_{\varepsilon \rightarrow k} P' \xrightarrow{\text{co } k}_{\sigma} Q$ and $(H \triangleright P) \xrightarrow{k} (H, k(a) \triangleright P') \xrightarrow{\tau} (H, a \triangleright Q)$ for some P' , any H , and any $k \# P, H$. \square

To keep histories and the bisimulation game finite in examples, the challenger of this bisimulation game performs all-but- \star transitions.

Definition 3.6. $C_1 \xrightarrow{\zeta} C_2$ is a challenger move if it is derived without using LTS_{\star} .

We now give the definition for a weak bisimulation over the above transitions.

Definition 3.7 (Weak Bisimulation). A binary relation \mathcal{R} over configurations is a weak bisimulation when for all $C_1 \mathcal{R} C_2$:

$$\begin{aligned}
 \mathcal{R}_3 &\stackrel{\text{def}}{=} \{ ((\varepsilon \triangleright P_3), (\varepsilon \triangleright Q_3)), ((k(a) \triangleright \llbracket b.\text{co} + c.\mathbf{0} \triangleright_k \mathbf{0} \rrbracket), (k(a) \triangleright \llbracket b.\text{co} \triangleright_k \mathbf{0} \rrbracket)), \\
 &\quad ((k(a), k(b) \triangleright \llbracket \text{co} \triangleright_k \mathbf{0} \rrbracket)), (k(a), k(b) \triangleright \llbracket \text{co} \triangleright_k \mathbf{0} \rrbracket)), ((a, b \triangleright \mathbf{0}), (a, b \triangleright \mathbf{0})), \\
 &\quad ((k(a), k(c) \triangleright \llbracket \mathbf{0} \triangleright_k \mathbf{0} \rrbracket)), (\text{ab}, \text{ab} \triangleright \mathbf{0}), ((\text{ab}, \dots \triangleright \mathbf{0}), (\text{ab}, \dots \triangleright \mathbf{0})) \mid \text{any } k \} \\
 \mathcal{R}_2 &\stackrel{\text{def}}{=} \{ ((\varepsilon \triangleright P_2), (\varepsilon \triangleright Q_2)), ((\text{ab}, \dots \triangleright \mathbf{0}), (\text{ab}, \dots \triangleright \mathbf{0})), \\
 &\quad (k_1(a) \triangleright \nu p. \llbracket p.\text{co}.R \triangleright_{k_1} \mathbf{0} \rrbracket \mid \llbracket b.p.\text{co}.S \triangleright_{k_2} \mathbf{0} \rrbracket), (k_1(a) \triangleright \llbracket b.\text{co}.(R \mid S) \triangleright_{k_1} \mathbf{0} \rrbracket)), \\
 &\quad (k_2(b) \triangleright \nu p. \llbracket a.p.\text{co}.R \triangleright_{k_1} \mathbf{0} \rrbracket \mid \llbracket p.\text{co}.S \triangleright_{k_2} \mathbf{0} \rrbracket), (k_2(b) \triangleright \llbracket a.\text{co}.(R \mid S) \triangleright_{k_2} \mathbf{0} \rrbracket)), \\
 &\quad ((\mathbf{k}_1(a), \mathbf{k}_2(b) \triangleright \nu p. \llbracket p.\text{co}.R \triangleright_{k_1} \mathbf{0} \rrbracket \mid \llbracket p.\text{co}.S \triangleright_{k_2} \mathbf{0} \rrbracket)), \\
 &\quad (k_2(a), k_2(b) \triangleright \llbracket \text{co}.(R \mid S) \triangleright_{k_2} \mathbf{0} \rrbracket)), \\
 &\quad ((\mathbf{k}_2(b), \mathbf{k}_1(a) \triangleright \nu p. \llbracket p.\text{co}.R \triangleright_{k_1} \mathbf{0} \rrbracket \mid \llbracket p.\text{co}.S \triangleright_{k_2} \mathbf{0} \rrbracket)), \\
 &\quad (k_1(b), k_1(a) \triangleright \llbracket \text{co}.(R \mid S) \triangleright_{k_1} \mathbf{0} \rrbracket)), \\
 &\quad ((\mathbf{k}(x), \mathbf{k}(y) \triangleright \nu p. \llbracket \text{co}.R \triangleright_k \mathbf{0} \rrbracket \mid \llbracket \text{co}.S \triangleright_k \mathbf{0} \rrbracket)), (k(x), k(y) \triangleright \llbracket \text{co}.(R \mid S) \triangleright_k \mathbf{0} \rrbracket)) \\
 &\quad ((x, y, H \triangleright \nu p. R \mid S), (x, y, H \triangleright R \mid S)) \\
 &\quad \mid \text{any } k, k_1, k_2, R, S, H \text{ and } (x, y) = (a, b) \text{ or } (b, a) \} \\
 \mathcal{R}_5 &\stackrel{\text{def}}{=} \{ ((H \triangleright P_5), (H \triangleright Q_5)), ((H, x \triangleright \mathbf{0}), (H, x \triangleright \mathbf{0})) \\
 &\quad ((H \triangleright \llbracket \tau.b.\text{co} + \tau.c.\text{co} \blacktriangleright P_5 \rrbracket)), (H \triangleright \llbracket \tau.(b.\text{co} + c.\text{co}) \blacktriangleright Q_5 \rrbracket)), \\
 &\quad ((H \triangleright \llbracket \tau.b.\text{co} + \tau.c.\text{co} \triangleright_k P_5 \rrbracket)), (H \triangleright \llbracket \tau.(b.\text{co} + c.\text{co}) \triangleright_l Q_5 \rrbracket)), \\
 &\quad ((H \triangleright \llbracket b.\text{co} \triangleright_k P_5 \rrbracket)), (H \triangleright \llbracket (b.\text{co} + c.\text{co}) \triangleright_l Q_5 \rrbracket)), \\
 &\quad ((H \triangleright \llbracket c.\text{co} \triangleright_k P_5 \rrbracket)), (H \triangleright \llbracket (b.\text{co} + c.\text{co}) \triangleright_l Q_5 \rrbracket)), \\
 &\quad ((H, k(x) \triangleright \llbracket \text{co} \triangleright_k P_5 \rrbracket)), (H, k(x) \triangleright \llbracket \text{co} \triangleright_k Q_5 \rrbracket)), \\
 &\quad \mid \text{any } k, l, H = (\text{ab}, \dots), \text{ and } x = a \text{ or } b \}
 \end{aligned}$$

Fig. 3. Relations used to prove the equivalences in Ex(s). 3.10 to 3.12.

1. $\text{hist}(\mathcal{C}_1)$ and $\text{hist}(\mathcal{C}_2)$ are consistent,
2. if $\mathcal{C}_1 \xrightarrow{\zeta} \mathcal{C}'_1$ is a challenger move and $\zeta \sharp \mathcal{C}_2$ then $\exists \mathcal{C}'_2: \mathcal{C}_2 \xrightarrow{\zeta} \mathcal{C}'_2$ and $\mathcal{C}'_1 \mathcal{R} \mathcal{C}'_2$,
3. the converse of the preceding condition.

Condition $\zeta \sharp \mathcal{C}_2$ guarantees that the choice of fresh transaction names in ζ does not hinder the transition from \mathcal{C}_2 . Weak bisimilarity (\approx) is the largest weak bisimulation, and extends to processes $P \approx Q$ if $(\emptyset \triangleright P) \approx (\emptyset \triangleright Q)$. Bisimulation transitions and weak bisimulations are unaffected by fresh renaming. Thus, the name selected in a challenger move is unimportant.

Lemma 3.8 (ζ -Equivariance). *If $\mathcal{C} \xrightarrow{\zeta} \mathcal{C}'$ then $\mathcal{C}r_{\text{fr}} \xrightarrow{\zeta r_{\text{fr}}} \mathcal{C}'r_{\text{fr}}$.* □

Lemma 3.9 (Equivariance of \approx). *If $\mathcal{C} \approx \mathcal{D}$ then $\mathcal{C} \approx \mathcal{D}r$.* □

We close this section by showing the equivalence of the processes in the introduction by proving them weakly bisimilar. The soundness of our bisimulation technique establishes a proof of contextual equivalence between these processes.

Example 3.10. Recall P_3 and Q_3 from (3) in the Introduction. We show that $P_3 \approx Q_3$; i.e., $(\emptyset \triangleright P_3) \approx (\emptyset \triangleright Q_3)$. It suffices to check that \mathcal{R}_3 in Fig. 3 is a weak bisimulation. Related histories in \mathcal{R}_3 are consistent. The interesting case is when $(k(a) \triangleright \llbracket b.\text{co} + c.\mathbf{0} \triangleright_k \mathbf{0} \rrbracket) \xrightarrow{k'} (k'(a), k'(c) \triangleright \llbracket \mathbf{0} \triangleright_{k'} \mathbf{0} \rrbracket)$. The defender responds:

$(k(a) \triangleright \llbracket b.\text{co} \triangleright_k \mathbf{0} \rrbracket) \xrightarrow{k'}_{\text{LTS}_\star} (k'(a), k'(\star) \triangleright \llbracket b.\text{co} \triangleright_k \mathbf{0} \rrbracket) \xrightarrow{\tau}_{\text{LTSab}} (\text{ab}, \text{ab} \triangleright \mathbf{0})$
 and get $(k'(a), k'(c) \triangleright \llbracket \mathbf{0} \triangleright_{k'} \mathbf{0} \rrbracket) \mathcal{R}_3 (\text{ab}, \text{ab} \triangleright \mathbf{0})$. The rest is trivial, thus the defender always wins, therefore \mathcal{R}_3 is a weak bisimulation and $P_3 \approx Q_3$. \square

Example 3.11. Let us now prove $P_2 \approx Q_2$ from (2) in the Introduction. For this proof we construct relation \mathcal{R}_2 in Fig. 3. It is easy to verify that all histories related in \mathcal{R}_2 are consistent and all challenger moves can be matched by the defender. Here it is noteworthy that the two tentative actions a and b are recorded in the left-hand history under different transaction names (k_1 and k_2 , respectively) until the synchronisation on p merges the two transactions; these history annotations are highlighted in bold typeface. \square

Example 3.12. In our final example proof we show that $P_5 \approx Q_5$ from (5) in the Introduction. Here we construct relation \mathcal{R}_5 in Fig. 3. In this construction, H is a history with zero or more aborted actions; we add this to our configurations because restarting transactions can nondeterministically abort and restart. The proof that \mathcal{R}_5 is a weak bisimulation is again by an easy inspection of the moves of the challenger. The important move is when from the pair $((H \triangleright \llbracket b.\text{co} \triangleright_k P_5 \rrbracket), (H \triangleright \llbracket (b.\text{co} + c.\text{co}) \triangleright_l Q_5 \rrbracket))$ the challenger picks the transition $(H \triangleright \llbracket (b.\text{co} + c.\text{co}) \triangleright_l Q_5 \rrbracket) \xrightarrow{l'}$ $(H, l'(c) \triangleright \llbracket \text{co} \triangleright_{l'} Q_5 \rrbracket)$ and the defender:

$$\begin{array}{ll} ((H \triangleright \llbracket b.\text{co} \triangleright_k P_5 \rrbracket) & \xrightarrow{\tau}_{(\text{LTSab})} (H \triangleright P_5) \\ \xrightarrow{\tau}_{(\text{LTS}_\tau)} (H \triangleright \llbracket \tau.b.\text{co} + \tau.c.\text{co} \blacktriangleright P_5 \rrbracket) & \xrightarrow{\tau}_{(\text{LTS}_{\text{new}})} (H \triangleright \llbracket \tau.b.\text{co} + \tau.c.\text{co} \triangleright_k P_5 \rrbracket) \\ \xrightarrow{\tau}_{(\text{LTS}_{k(\tau)})} (H \triangleright \llbracket c.\text{co} \triangleright_{k'} P_5 \rrbracket) & \xrightarrow{l'}_{(\text{LTS}_{k(a)})} (H, l'(c) \triangleright \llbracket \text{co} \triangleright_{l'} P_5 \rrbracket) \end{array}$$

and get $(H, l'(c) \triangleright \llbracket \text{co} \triangleright_{l'} Q_5 \rrbracket) \mathcal{R}_5 (H, l'(c) \triangleright \llbracket \text{co} \triangleright_{l'} P_5 \rrbracket)$. \square

4 Predictive Bisimulations

In the previous section we showed that weak bisimulations provide an effective verification technique of equivalence. However, it does not enable a direct soundness proof. The difficulty is in proving weak bisimulation *compositional* (i.e., a congruence). Here we define *predictive bisimulations*, an alternative notion of bisimulations that allows us to give an indirect proof of soundness of (\approx) . First we explain the problem with directly proving (\approx) compositional.

Failing Proof (Compositionality) We need to prove \mathcal{R} a weak bisimulation, when

$$(H_1 \triangleright P \mid R) \mathcal{R} (H_2 \triangleright Q \mid R) \quad (7)$$

for any context R and $(H_1 \triangleright P) \approx (H_2 \triangleright Q)$. Let $(H_1 \triangleright P) \xrightarrow{k} (H_1, k(a) \triangleright P')$ and R can perform $k(\bar{a})$ to become R' . We have $(H_1 \triangleright P \mid R) \xrightarrow{\tau} (H_1 \triangleright P' \mid R')$. We need to show that there exist H'_2 , Q'' , and R'' such that $(H_2 \triangleright Q \mid R) \xrightarrow{\tau} (H'_2 \triangleright Q'' \mid R'')$ and $(H_1 \triangleright P' \mid R') \mathcal{R} (H'_2 \triangleright Q'' \mid R'')$. Weak bisimilarity gives

$$(H_2 \triangleright Q) \xrightarrow{k} (H_2, k(\hat{b}) \triangleright Q') \text{ and } (H_1, k(a) \triangleright P') \approx (H_2, k(\hat{b}) \triangleright Q')$$

for some \hat{b} , Q' , and the new parts $k(a)$ and $k(\hat{b})$ of the histories are consistent. However, consistency does not restrict the values of \hat{b} ; it can be a name different than a , or even \star , provided that k does not commit in any extension of the bisimulation game. When $\hat{b} = a$ we can complete the proof by taking $H'_2 = (H_2, k(\hat{a}))$, $Q'' = Q'$ and $R'' = R'$, showing $(H_1, k(a) \triangleright P' \mid R') \mathcal{R} (H_2, k(a) \triangleright Q' \mid R')$ because it is of the same shape as (7). However, when $\hat{b} = c \neq a$ or $\hat{b} = \star$ it is unclear how to proceed in this direct proof. \square

We instead prove soundness by defining (\approx_{prd}) and showing:

$$P \approx Q \text{ implies } P \approx_{\text{prd}} Q \text{ implies } P \cong_{\text{rbe}} Q \quad (8)$$

To show the second implication we need to prove (\approx_{prd}) compositional. The intuition of why this is possible is because (\approx_{prd}) only takes into account those challenger transitions that have the possibility to be committed later in the bisimulation game. This allows us to define a stronger definition of consistency which avoids the problematic cases of the above failed direct proof. Strong consistency is a reflexive, symmetric, and transitive relation.

Definition 4.1 (Strong Consistency (\simeq)). $H_1 \simeq H_2$ when:

$$(H_1(i) = \hat{a} \text{ iff } H_2(i) = \hat{a}) \quad \text{and} \quad (\exists k. H_1(i) = k(\hat{a}) \text{ iff } \exists l. H_2(i) = l(\hat{a}))$$

In a predictive bisimulation game the challenger only performs transitions within transactions when those transactions can commit later in the game. Thus, here we differentiate between τ - and $k(\tau)$ -transitions. Moreover, we emphasise that a challenger $k(a)$ move has to be matched with an identical defender move. Thus predictive bisimulations are over the actions $\eta ::= \tau \mid k(\tau) \mid k(a) \mid \omega$.

Definition 4.2 (Pred. Bisim. Transitions). $\mathcal{C} \xrightarrow{\eta} \mathcal{C}'$ is derived by the rules:

$$\begin{array}{ll} (H \triangleright P) \xrightarrow{\tau} (H \triangleright Q) & \text{if } P \xrightarrow{\tau}_{\varepsilon} Q \quad (\text{LTS}'\tau) \\ (H \triangleright P) \xrightarrow{k(\tau)} (\sigma(H) \triangleright Q) & \text{if } P \xrightarrow{k(\tau)}_{\sigma} Q \text{ and } k \# H \quad (\text{LTS}'k(\tau)) \\ (H \triangleright P) \xrightarrow{\tau} (H \triangleright Q) & \text{if } P \xrightarrow{\text{new } k} Q \text{ and } k \# H \quad (\text{LTS}'\text{new}) \\ (H \triangleright P) \xrightarrow{\tau} (H \setminus_{\text{co}} k \triangleright Q) & \text{if } P \xrightarrow{\text{co}k} Q \quad (\text{LTS}'\text{co}) \\ (H \triangleright P) \xrightarrow{\tau} (H \setminus_{\text{ab}} k \triangleright Q) & \text{if } P \xrightarrow{\text{ab}k} Q \quad (\text{LTS}'\text{ab}) \\ (H \triangleright P) \xrightarrow{k(a)} (\sigma(H), k(a) \triangleright Q) & \text{if } P \xrightarrow{k(a)}_{\sigma} Q \text{ and } k \# H \quad (\text{LTS}'k(a)) \\ (H \triangleright P) \xrightarrow{\omega} (\sigma(H) \triangleright \mathbf{0}) & \text{if } P \xrightarrow{\omega}_{\sigma} Q \quad (\text{LTS}'\omega) \end{array}$$

We define $\xrightarrow{\eta}$ to be $\left((\xrightarrow{\tau}) \cup (\xrightarrow{k(\tau)}) \right)^*$ when $\eta \in \{\tau, k(\tau)\}$, and $\xrightarrow{\tau} \xrightarrow{\eta} \xrightarrow{\tau}$ otherwise.

In a predictive bisimulation game, defender moves are weak η -transitions, with no need for $k(\star)$ -transitions. Challenger moves are *commitable* transitions.

Definition 4.3 (Commitable Transition). $\mathcal{C} \xrightarrow{\eta} \mathcal{C}'$ with $\eta \in \{\tau, \omega\}$ is *commitable*; $\mathcal{C} \xrightarrow{k(\mu)} (H_1 \triangleright P_1)$ is *commitable* when there exists $(H_1 \triangleright P_1) \xrightarrow{\eta_1} \dots \xrightarrow{\eta_{(n+1)}} (H_n \triangleright P_n)$ such that for any a and $i \notin \text{dom}(H_1)$:

$$(H_1, (i \mapsto k(a)) \triangleright P_1) \xrightarrow{\eta_1} \dots \xrightarrow{\eta_{(n+1)}} (H_n, (i \mapsto a) \triangleright P_n)$$

Weak predictive bisimulations are thus defined as follows.

Definition 4.4 (Weak Predictive Bisimulation). *A binary relation \mathcal{R} is a weak predictive bisimulation when for all $\mathcal{C}_1 \mathcal{R} \mathcal{C}_2$:*

1. $\text{hist}(\mathcal{C}_1)$ and $\text{hist}(\mathcal{C}_2)$ are strongly consistent,
2. if $\mathcal{C}_1 \xrightarrow{\eta} \mathcal{C}'_1$ is a committable transition and $\text{ftn}(\eta) \# \mathcal{C}_2$ then $\exists \mathcal{C}'_2$ such that $\mathcal{C}_2 \xrightarrow{\eta} \mathcal{C}'_2$ and $\mathcal{C}'_1 \mathcal{R} \mathcal{C}'_2$, and its converse.

Weak predictive bisimilarity (\approx_{prd}) is the largest such bisimulation and (\approx_{prd}) extends to processes in the same way as (\approx). The first part of (8) follows by:

Theorem 4.5. *Let $\mathcal{C} \approx \mathcal{C}'$ with strongly consistent histories; then $\mathcal{C} \approx_{\text{prd}} \mathcal{C}'$.*

Proof. The proof of this proposition relies on showing that strong consistency is closed under committable transitions of weakly bisimilar configurations. \square

To prove the second part of (8) we need to show that (\approx_{prd}) is compositional.

Theorem 4.6. *If $P \approx_{\text{prd}} Q$ and $\text{ftn}(R) \# P, Q$ then $P \mid R \approx_{\text{prd}} Q \mid R$.*

Proof. This relies on de-/re-composition of actions; e.g., we need to decompose $(H \triangleright P \mid R) \xrightarrow{\eta} (H \triangleright P' \mid R')$ into the constituent sub-actions from P, R with appropriate histories. This is facilitated by strongly consistent histories. \square

5 Full Abstraction

Using (\approx) we can prove soundness of our original bisimulation game.

Theorem 5.1 (Soundness). *If $P \approx Q$ then $P \cong_{\text{rbe}} Q$.*

Proof. In view of Thm. 4.5 it is sufficient to prove the result for (\approx_{prd}). The major step in this proof is already established in Thm. 4.6. \square

We prove completeness for \mathcal{L}_{Act} by first translating any history H into a process $\langle H \rangle$. Then we show that the transitions of a configuration $(H \triangleright P)$ examined by bisimulations can be modelled by reductions of the process $\langle H \rangle \mid P$, when put in parallel with the appropriate contexts. The translation of H is the parallel composition of the translation of each element in H according to:

$$\begin{aligned} \langle i \mapsto k(a) \rangle &= \llbracket \text{co} \mid \omega_{ai}^{\text{commit}} \triangleright_k \omega_i^{\text{abort}} \rrbracket & \langle i \mapsto a \rangle &= \omega_{ai}^{\text{commit}} & \langle i \mapsto \text{ab} \rangle &= \omega_i^{\text{abort}} \\ \langle i \mapsto k(\star) \rangle &= \llbracket \text{co} \triangleright_k \omega_i^{\text{abort}} \rrbracket & \langle i \mapsto \star \rangle &= \mathbf{0} \end{aligned}$$

A tentative $k(a)$ -action, recorded in the history as $(i \mapsto k(a))$, corresponds to a particular move of the bisimulation game—say the i th move. This is translated to a k -transaction which is ready to commit. When k commits because *all* distributed parts of k commit, a unique “success” barb $\omega_{ai}^{\text{commit}}$ becomes observable, signalling that the i th move in the bisimulation game was a synchronisation on a which became permanent. In the history this is recorded as $(i \mapsto a)$. If the k -transaction aborts then a unique ω_i^{abort} barb signals the abortion of the i th

move, corresponding to $(i \mapsto \mathbf{ab})$ in the history. The translation of a defender's $(i \mapsto k(\star))$ move is similar, with the exception that this is a no-action that has no "success" barb associated with it. A key proposition is that reductions of translated configurations model silent bisimulation transitions.

Proposition 5.2. $(H_1 \triangleright P) \xrightarrow{\tau} (H_2 \triangleright Q)$ iff $\langle H_1 \rangle \mid P \rightarrow \langle H_2 \rangle \mid Q$.

Moreover, the i th tentative k -transition in the bisimulation game is modelled by the reduction induced by the context:

$$\langle k \rangle^i = \llbracket \mathbf{co} \mid (\sum_{a \in \text{Act}} \bar{a}.\omega_{ai}^{\text{commit}}) + \tau.\mathbf{0} + \omega_i^{\text{before}} \triangleright_k \omega_i^{\text{abort}} \rrbracket$$

When synchronising with a process, this context becomes $\langle i \mapsto k(a) \rangle$ (for any a), modelling a tentative $k(a)$ -transition. It may also spontaneously become $\langle i \mapsto k(\star) \rangle$, modelling a $k(\star)$ -transition. In any case it loses the weak barb ω_i^{before} .

Proposition 5.3. Let $H_2 = \sigma(H_1)$, $(i \mapsto k(\hat{a}))$ and $k' \# k$, H_1, P , and $\text{ftn}(H_1) \subseteq \text{ftn}(P)$; then $(H_1 \triangleright P) \xrightarrow{k} (H_2 \triangleright Q)$ iff $\langle H_1 \rangle \mid P \mid \langle k' \rangle^i \rightarrow \langle H_2 \rangle \mid Q \not\Downarrow \omega_i^{\text{before}}$. \square

Theorem 5.4 (Completeness). If $P, Q \in \mathcal{L}_{\text{Act}}$ and $P \cong_{\text{rbe}} Q$ then $P \approx Q$.

Proof. Using the above propositions we show \mathcal{X} is a weak bisimulation, where $\mathcal{X} = \{((H \triangleright P), (H' \triangleright Q)) \mid H, H' \text{ cons.}, P, Q \in \mathcal{L}_{\text{Act}}, \langle H \rangle \mid P \cong_{\text{rbe}} \langle H' \rangle \mid Q\}$. \square

6 Conclusions

We presented a weak bisimulation theory for $TCCS^m$, a simple language with communicating transactions. In $TCCS^m$, two transactions that communicate are conjoined by being renamed to the same name forming a *distributed* version of cJoin's merged transactions [1]. When a transaction communicates with a non-transactional process, the latter is embedded in the former in line with the semantics of previous work [3]. Compared to that semantics, embedding and merging in $TCCS^m$ is communication-driven limiting nondeterminism. For simplicity this language has only single-level transactions, although we believe that the results of this paper can be adapted to a language with nested transactions.

The bisimulation equivalence is sound and complete with respect to a natural contextual equivalence which equates transactions that differ only in uncommitable actions; these are destined to eventually be rolled back. We motivated this with examples in the Introduction which we verified in Sect. 3 using our technique. In related work about reversible calculi [10,9], contextual equivalence can discriminate between transactions $\llbracket b + a.b.\mathbf{co} \triangleright_k \mathbf{0} \rrbracket$ and $\llbracket a.b.\mathbf{co} \triangleright_k \mathbf{0} \rrbracket$ by virtue of the fact that the former contains an extra uncommitable b -action.

Our bisimulations provide an effective verification technique for the aforementioned contextual equivalence which can be applied to related programming languages where uncommitable actions have no effect [5,7,11,16]. They can also serve as a verification technique for other forms of contextual equivalences, such as may- and must-testing equivalences [3,4]. Other bisimulation methods for

reversible computation [2,8,9,10] may also be used in these settings, but they are more fine-grained, discriminating even between the above two processes. Forward-reverse and hereditary history-preserving bisimulations [13] differentiate between forward and reverse transitions, which would discriminate between the processes P_5 and Q_5 shown in the Introduction capturing a possible compiler optimisation. To our knowledge, the bisimulation technique presented here is the only one that can be used to prove the correctness of such compiler optimisations.

References

1. Bruni, R., Melgratti, H.C., Montanari, U.: Nested commits for mobile calculi: extending Join. In: Levy, J.-J., Mayr, E.W., Mitchell, J.C. (eds.) TCS 2004. IFIP, vol. 155, pp. 563–576. Springer, Boston (2004)
2. Danos, V., Krivine, J.: Transactions in RCCS. In: Abadi, M., de Alfaro, L. (eds.) CONCUR 2005. LNCS, vol. 3653, pp. 398–412. Springer, Heidelberg (2005)
3. de Vries, E., Koutavas, V., Hennessy, M.: Communicating transactions. In: Gastin, P., Laroussinie, F. (eds.) CONCUR 2010. LNCS, vol. 6269, pp. 569–583. Springer, Heidelberg (2010)
4. de Vries, E., Koutavas, V., Hennessy, M.: Liveness of communicating transactions (Extended abstract). In: Ueda, K. (ed.) APLAS 2010. LNCS, vol. 6461, pp. 392–407. Springer, Heidelberg (2010)
5. Donnelly, K., Fluet, M.: Transactional events. In: ICFP, pp. 124–135. ACM (2006)
6. Effinger-Dean, L., Kehrt, M., Grossman, D.: Transactional events for ML. In: ICFP, pp. 103–114. ACM (2008)
7. Field, J., Varela, C.A.: Transactors: a programming model for maintaining globally consistent distributed state in unreliable environments. In: POPL, pp. 195–208. ACM (2005)
8. Krivine, J.: A verification technique for reversible process algebra. In: Glück, R., Yokoyama, T. (eds.) RC 2012. LNCS, vol. 7581, pp. 204–217. Springer, Heidelberg (2013)
9. Lanese, I., Lienhardt, M., Mezzina, C.A., Schmitt, A., Stefani, J.-B.: Concurrent flexible reversibility. In: Felleisen, M., Gardner, P. (eds.) ESOP 2013. LNCS, vol. 7792, pp. 370–390. Springer, Heidelberg (2013)
10. Lanese, I., Mezzina, C.A., Stefani, J.-B.: Reversing higher-order pi. In: Gastin, P., Laroussinie, F. (eds.) CONCUR 2010. LNCS, vol. 6269, pp. 478–493. Springer, Heidelberg (2010)
11. Lesani, M., Palsberg, J.: Communicating memory transactions. In: PPOPP, pp. 157–168. ACM (2011)
12. Milner, R.: A Calculus of Communication Systems. LNCS, vol. 92. Springer, Heidelberg (1980)
13. Phillips, I., Ulidowski, I.: Reversibility and models for concurrency. ENTCS 192(1), 93–108 (2007) (SOS 2007)
14. Sangiorgi, D., Kobayashi, N., Sumii, E.: Environmental bisimulations for higher-order languages. In: LICS, pp. 293–302. IEEE Computer Society (2007)
15. Sangiorgi, D., Walker, D.: The π -calculus: a Theory of Mobile Processes. Cambridge University Press (2001)
16. Spaccasassi, C., Koutavas, V.: Towards efficient abstractions for concurrent consensus. In: McCarthy, J. (ed.) TFP 2013. LNCS, vol. 8322, pp. 76–90. Springer, Heidelberg (2014)
17. Stirling, C.: Playing games and proving properties of concurrent systems. J. Comput. Sci. Technol. 13(6), 482 (1998)