

EPTCS 190

Proceedings of the
**Combined 22th International Workshop on
Expressiveness in Concurrency
and 12th Workshop on
Structural Operational Semantics**

Madrid, Spain, 31st August 2015

Edited by: Silvia Crafa and Daniel E. Gebler

Published: 26th August 2015
DOI: 10.4204/EPTCS.190
ISSN: 2075-2180
Open Publishing Association

Table of Contents

Table of Contents	i
Preface	ii
Comparing Deadlock-Free Session Typed Processes	1
<i>Ornela Dardha and Jorge A. Pérez</i>	
On Compensation Primitives as Adaptable Processes	16
<i>Jovana Dedeić, Jovanka Pantović and Jorge A. Pérez</i>	
SOS rule formats for convex and abstract probabilistic bisimulations	31
<i>Pedro R. D’Argenio, Matias David Lee and Daniel Gebler</i>	
Analysing and Comparing Encodability Criteria	46
<i>Kirstin Peters and Rob van Glabbeek</i>	
Encoding CSP into CCS	61
<i>Meike Hatzel, Christoph Wagner, Kirstin Peters and Uwe Nestmann</i>	
Encoding the Factorisation Calculus	76
<i>Reuben N. S. Rowe</i>	

Preface

This volume contains the proceedings of the Combined 22nd International Workshop on Expressiveness in Concurrency and the 12th Workshop on Structural Operational Semantics (EXPRESS/SOS 2015) which was held on 31 August 2015 in Madrid, Spain, as an affiliated workshop of CONCUR 2015, the 26th International Conference on Concurrency Theory.

The EXPRESS workshops aim at bringing together researchers interested in the expressiveness of various formal systems and semantic notions, particularly in the field of concurrency. Their focus has traditionally been on the comparison between programming concepts (such as concurrent, functional, imperative, logic and object-oriented programming) and between mathematical models of computation (such as process algebras, Petri nets, event structures, modal logics, and rewrite systems) on the basis of their relative expressive power. The EXPRESS workshop series has run successfully since 1994 and over the years this focus has become broadly construed.

The SOS workshops aim at being a forum for researchers, students and practitioners interested in new developments, and directions for future investigation, in the field of structural operational semantics. One of the specific goals of the SOS workshop series is to establish synergies between the concurrency and programming language communities working on the theory and practice of SOS. Reports on applications of SOS to other fields are also most welcome, including: modelling and analysis of biological systems, security of computer systems programming, modelling and analysis of embedded systems, specification of middle-ware and coordination languages, programming language semantics and implementation, static analysis software and hardware verification, semantics for domain-specific languages and model-based engineering.

Since 2012, the EXPRESS and SOS communities have organized an annual combined EXPRESS/SOS workshop on the expressiveness of mathematical models of computation and the formal semantics of systems and programming concepts.

We received ten full paper submissions out of which the programme committee selected six for publication and presentation at the workshop. These proceedings contain these selected contributions. The workshop had an invited presentation:

Bisimulation techniques in probabilistic higher-order languages,

by Davide Sangiorgi (Universita' di Bologna)

We would like to thank the authors of the submitted papers, the invited speaker, the members of the programme committee, and their subreviewers for their contribution to both the meeting and this volume. We also thank the CONCUR 2015 organizing committee for hosting EXPRESS/SOS 2015. Finally, we would like to thank our EPTCS editor Rob van Glabbeek for publishing these proceedings and his help during the preparation.

Silvia Crafa and Daniel E. Gebler,
August 2015

Program Committee

Johannes Borgstrom (Uppsala University, Sweden)
Matteo Cimini (Indiana University, Bloomington, Indiana)
Silvia Crafa (University of Padova, Italy)
Pedro R. D'Argenio (University of Cordoba, Argentina)
Daniel Gebler (VU University Amsterdam, The Netherlands)

Thomas Given-Wilson (Inria, France)
Thomas T. Hildebrandt (IT University of Copenhagen, Denmark)
Daniel Hirschhoff (ENS Lyon, France)
Stefan Milius (University of Erlangen-Nurnberg, Germany)
Mohammad R. Mousavi (Halmstad University, Sweden)
Kirstin Peters (Technical University of Berlin, Germany)
Damien Pous (ENS Lyon, France)
Irek Ulidowski (University of Leister, United Kingdom)

Additional Reviewers

Massimo Bartoletti, Paul Brunet, Cinzia Di Giusto, Carla Ferreira, Daniele Gorla, Barry Jay, Johannes Aman Pohjola.

Comparing Deadlock-Free Session Typed Processes

Ornela Dardha

University of Glasgow, United Kingdom

Jorge A. Pérez

University of Groningen, The Netherlands

Besides respecting prescribed protocols, communication-centric systems should never “get stuck”. This requirement has been expressed by liveness properties such as progress or (dead)lock freedom. Several typing disciplines that ensure these properties for mobile processes have been proposed. Unfortunately, very little is known about the precise relationship between these disciplines—and the classes of typed processes they induce.

In this paper, we compare \mathcal{L} and \mathcal{H} , two classes of deadlock-free, session typed concurrent processes. The class \mathcal{L} stands out for its canonicity: it results naturally from interpretations of linear logic propositions as session types. The class \mathcal{H} , obtained by encoding session types into Kobayashi’s usage types, includes processes not typable in other type systems.

We show that \mathcal{L} is strictly included in \mathcal{H} . We also identify the precise condition under which \mathcal{L} and \mathcal{H} coincide. One key observation is that the *degree of sharing* between parallel processes determines a new expressiveness hierarchy for typed processes. We also provide a type-preserving rewriting procedure of processes in \mathcal{H} into processes in \mathcal{L} . This procedure suggests that, while effective, the degree of sharing is a rather subtle criteria for distinguishing typed processes.

1 Introduction

The goal of this work is to formally relate different type systems for the π -calculus. Our interest is in *session-based concurrency*, a type-based approach to communication correctness: dialogues between participants are structured into *sessions*, basic communication units; descriptions of interaction sequences are then abstracted as *session types* [12] which are checked against process specifications. We offer the first formal comparison between different type systems that enforce *(dead)lock freedom*, the liveness property that ensures session communications never “get stuck”. Our approach relates the classes of typed processes that such systems induce. To this end, we identify a property on the structure of typed parallel processes, the *degree of sharing*, which is key in distinguishing two salient classes of deadlock-free session processes, and in shedding light on their formal underpinnings.

In session-based concurrency, types enforce correct communications through different safety and liveness properties. Basic correctness properties are *communication safety* and *session fidelity*: while the former ensures absence of errors (e.g., communication mismatches), the latter ensures that well-typed processes respect the protocols prescribed by session types. Moreover, a central (liveness) property for safe processes is that they should never “get stuck”. This is the well-known *progress* property, which asserts that a well-typed term either is a final value or can further reduce [17]. In calculi for concurrency, this property has been formalized as *deadlock freedom* (“a process is deadlock-free if it can always reduce until it eventually terminates, unless the whole process diverges” [15]) or as *lock freedom* (“a process is lock free if it can always reduce until it eventually terminates, even if the whole process diverges” [13]). Notice that in the absence of divergent behaviors, deadlock and lock freedom coincide.

(Dead)lock freedom guarantees that all communications will eventually succeed, an appealing requirement for communicating processes. Several advanced type disciplines that ensure deadlock-free processes have been proposed (see, e.g., [2, 3, 5, 10, 13, 15, 16, 20]). Unfortunately, these disciplines consider different process languages and/or are based on rather different principles. As a result, very little

is known about how they relate to each other. This begs several research questions: What is the formal relationship between these type disciplines? What classes of deadlock-free processes do they induce?

In this paper, we tackle these open questions by comparing \mathcal{L} and \mathcal{H} , two salient classes of deadlock-free, session typed processes (Definition 4.2):

- \mathcal{L} contains all session processes that are well-typed according to the Curry-Howard correspondence of linear logic propositions as session types [2, 3, 21]. This suffices, because the type system derived from such a correspondence ensures communication safety, session fidelity, and deadlock freedom.
- \mathcal{H} contains all session processes that enjoy communication safety and session fidelity (as ensured by the type system of Vasconcelos [19]) and are (dead)lock-free by combining Kobayashi’s type system based on *usages* [13, 15] with Dardha et al.’s encodability result [8].

There are good reasons for considering \mathcal{L} and \mathcal{H} . On the one hand, due to its deep logical foundations, \mathcal{L} appears to us as the *canonic* class of deadlock-free session processes, upon which all other classes should be compared. Indeed, this class arguably offers the most principled yardstick for comparisons. On the other hand, \mathcal{H} integrates session type checking with the sophisticated usage discipline developed by Kobayashi for π -calculus processes. This indirect approach to deadlock freedom (first suggested in [14], later developed in [4, 7, 8]) is fairly general, as it may capture sessions with subtyping, polymorphism, and higher-order communication. Also, as informally shown in [4], \mathcal{H} strictly includes classes of typed processes induced by other type systems for deadlock freedom in sessions [5, 10, 16].

One key observation in our development is that \mathcal{H} corresponds to a *family* of classes of deadlock-free processes, denoted $\mathcal{H}_0, \mathcal{H}_1, \dots, \mathcal{H}_n$, which is defined by the *degree of sharing* between their parallel components. Intuitively, \mathcal{H}_0 is the subclass of \mathcal{H} with *independent parallel composition*: for all processes $P \mid Q \in \mathcal{H}_0$, subprocesses P and Q do not share any sessions. Then, \mathcal{H}_1 is the subclass of \mathcal{H} which contains \mathcal{H}_0 but admits also processes with parallel components that share at most one session. Then, \mathcal{H}_n contains deadlock-free session processes whose parallel components share at most n sessions.

Contributions. In this paper, we present three main contributions:

1. We show that the inclusion between the constituent classes of \mathcal{H} is *strict* (Theorem 4.4). We have:

$$\mathcal{H}_0 \subset \mathcal{H}_1 \subset \mathcal{H}_2 \subset \dots \subset \mathcal{H}_n \subset \mathcal{H}_{n+1} \quad (1)$$

Although not extremely surprising, the significance of this result lies in the fact that it talks about concurrency (via the degree of sharing) but implicitly also about the potential sequentiality of parallel processes. As such, processes in \mathcal{H}_k are necessarily “more parallel” than those in \mathcal{H}_{k+1} . Interestingly, the degree of sharing in $\mathcal{H}_0, \dots, \mathcal{H}_n$ can be defined in a very simple way, via a natural condition in the rule for parallel composition in Kobayashi’s type system for deadlock freedom.

2. We show that \mathcal{L} and \mathcal{H}_1 coincide (Theorem 4.6). That is, there are deadlock-free session processes that cannot be typed by systems derived from the Curry-Howard interpretation of session types [2, 3, 21], but that can be admitted by the (indirect) approach of [8]. This result is significant: it establishes the precise status of systems based on [3, 21] with respect to previous (non Curry-Howard) disciplines. Indeed, it formally confirms that linear logic interpretations of session types naturally induce the most basic form of concurrent cooperation (sharing of exactly one session), embodied as the principle of “composition plus hiding”, a distinguishing feature of such interpretations.
3. We define a rewriting procedure of processes in \mathcal{H} into \mathcal{L} (Definition 5.7). Intuitively, due to our previous observation and characterization of the degree of sharing in session typed processes, it is

quite natural to convert a process in \mathcal{K} into another, more parallel process in \mathcal{L} . In essence, the procedure replaces sequential prefixes with representative parallel components. The rewriting procedure satisfies type-preservation, and enjoys the compositionality and operational correspondence criteria as stated in [11] (cf. Theorems 5.8 and 5.10). These properties not only witness the significance of the rewriting procedure; they also confirm that the degree of sharing is a rather subtle criteria for formally distinguishing deadlock-free, session typed processes.

To the best of our knowledge, these contributions define the first formal comparison between fundamentally distinct type systems for deadlock freedom in session communications. Previous comparisons, such as the ones in [4] and [3, §6], are informal: they are based on representative “corner cases”, i.e., examples of deadlock-free session processes typable in one system but not in some other.

The paper is structured as follows. § 2 summarizes the session π -calculus and associated type system of [19]. In § 3 we present the two typed approaches to deadlock freedom for sessions. § 4 defines the classes \mathcal{L} and \mathcal{K} , formalizes the hierarchy (1), and shows that \mathcal{L} and \mathcal{K}_1 coincide. In § 5 we give the rewriting procedure of \mathcal{K}_n into \mathcal{L} and establish its properties. § 6 collects some concluding remarks. Due to space restrictions, details of proofs are omitted; they can be found online [9].

2 Session π -calculus

Following Vasconcelos [19], we introduce the session π -calculus and its associated type system which ensures communication safety and session fidelity. The syntax is given in Figure 1 (upper part). Let P, Q range over processes x, y over channels and v over values; for simplicity, the set of values coincides with that of channels. In examples, we often use \mathbf{n} to denote a terminated channel that cannot be further used.

Process $\bar{x}(v).P$ denotes the output of v along x , with continuation P . Dually, process $x(y).P$ denotes an input along x with continuation P , with y denoting a placeholder. Process $x \triangleleft l_j.P$ uses x to select l_j from a labelled choice process, being $x \triangleright \{l_i : P_i\}_{i \in I}$, so as to trigger P_j ; labels indexed by the finite set I are pairwise distinct. We also have the inactive process (denoted $\mathbf{0}$), the parallel composition of P and Q (denoted $P \mid Q$), and the (double) restriction operator, noted $(\nu xy)P$: the intention is that x and y denote *dual session endpoints* in P . We omit $\mathbf{0}$ whenever possible and write, e.g., $\bar{x}(\mathbf{n})$ instead of $\bar{x}(\mathbf{n}).\mathbf{0}$. Notions of bound/free variables in processes are standard; we write $\text{fn}(P)$ to denote the set of free names of P . Also, we write $P^{[v/z]}$ to denote the (capture-avoiding) substitution of free occurrences of z in P with v .

The operational semantics is given in terms of a reduction relation, noted $P \rightarrow Q$, and defined by the rules in Figure 1 (lower part). It relies on a standard notion of structural congruence, noted \equiv (see [19]). We write \rightarrow^* to denote the reflexive, transitive closure of \rightarrow . Observe that interaction involves prefixes with different channels (endpoints), and always occurs in the context of an outermost (double) restriction. Key rules are (R-COM) and (R-CASE), denoting the interaction of output/input prefixes and selection/branching constructs, respectively. Rules (R-PAR), (R-RES), and (R-STR) are standard.

The syntax of session types, ranged over T, S, \dots , is given by the following grammar.

$$T, S ::= \mathbf{end} \mid ?T.S \mid !T.S \mid \&\{l_i : S_i\}_{i \in I} \mid \oplus\{l_i : S_i\}_{i \in I}$$

Above, \mathbf{end} is the type of an endpoint with a terminated protocol. The type $?T.S$ is assigned to an endpoint that first receives a value of type T and then continues according to the protocol described by S . Dually, type $!T.S$ is assigned to an endpoint that first outputs a value of type T and then continues according to the protocol described by S . Type $\oplus\{l_i : S_i\}_{i \in I}$, an *internal choice*, generalizes output types; type $\&\{l_i : S_i\}_{i \in I}$, an *external choice*, generalizes input types. Notice that session types describe *sequences* of structured behaviors; they do not admit parallel composition operators.

$P, Q ::= \bar{x}(v).P$	(output)	$\mathbf{0}$	(inaction)
$x(y).P$	(input)	$P \mid Q$	(composition)
$x \triangleleft l_j.P$	(selection)	$(\nu xy)P$	(session restriction)
$x \triangleright \{l_i : P_i\}_{i \in I}$	(branching)		
$v ::= x$	(channel)		

(R-COM) $(\nu xy)(\bar{x}(v).P \mid y(z).Q) \rightarrow (\nu xy)(P \mid Q[v/z])$	(R-PAR) $P \rightarrow Q \implies P \mid R \rightarrow Q \mid R$
(R-CASE) $(\nu xy)(x \triangleleft l_j.P \mid y \triangleright \{l_i : P_i\}_{i \in I}) \rightarrow (\nu xy)(P \mid P_j) \quad j \in I$	(R-RES) $P \rightarrow Q \implies (\nu xy)P \rightarrow (\nu xy)Q$
(R-STR) $P \equiv P', P \rightarrow Q, Q' \equiv Q \implies P' \rightarrow Q'$	

Figure 1: Session π -calculus: syntax and semantics.

(T-NIL) $\frac{}{x : \mathbf{end} \vdash_{\text{ST}} \mathbf{0}}$	(T-PAR) $\frac{\Gamma_1 \vdash_{\text{ST}} P \quad \Gamma_2 \vdash_{\text{ST}} Q}{\Gamma_1 \circ \Gamma_2 \vdash_{\text{ST}} P \mid Q}$	(T-RES) $\frac{\Gamma, x : T, y : \bar{T} \vdash_{\text{ST}} P}{\Gamma \vdash_{\text{ST}} (\nu xy)P}$	(T-IN) $\frac{\Gamma, x : S, y : T \vdash_{\text{ST}} P}{\Gamma, x : ?T.S \vdash_{\text{ST}} x(y).P}$
(T-OUT) $\frac{\Gamma, x : S \vdash_{\text{ST}} P}{\Gamma, x : !T.S, y : T \vdash_{\text{ST}} \bar{x}(y).P}$	(T-BRCH) $\frac{\Gamma, x : S_i \vdash_{\text{ST}} P_i \quad \forall i \in I}{\Gamma, x : \&\{l_i : S_i\}_{i \in I} \vdash_{\text{ST}} x \triangleright \{l_i : P_i\}_{i \in I}}$	(T-SEL) $\frac{\Gamma, x : S_j \vdash_{\text{ST}} P \quad \exists j \in I}{\Gamma, x : \oplus\{l_i : S_i\}_{i \in I} \vdash_{\text{ST}} x \triangleleft l_j.P}$	

Figure 2: Typing rules for the π -calculus with sessions.

A central notion in session-based concurrency is *duality*, which relates session types offering opposite (i.e., complementary) behaviors. Duality stands at the basis of communication safety and session fidelity. Given a session type T , its dual type \bar{T} is defined as follows:

$$\frac{\bar{!T.S} \triangleq ?T.\bar{S}}{\oplus\{l_i : S_i\}_{i \in I} \triangleq \&\{l_i : \bar{S}_i\}_{i \in I}} \quad \frac{\bar{?T.S} \triangleq !T.\bar{S}}{\&\{l_i : S_i\}_{i \in I} \triangleq \oplus\{l_i : \bar{S}_i\}_{i \in I}} \quad \bar{\mathbf{end}} \triangleq \mathbf{end}$$

Typing contexts, ranged over by Γ, Γ' , are sets of typing assignments $x : T$. Given a context Γ and a process P , a session typing judgement is of the form $\Gamma \vdash_{\text{ST}} P$. Typing rules are given in Figure 2. Rule (T-NIL) states that $\mathbf{0}$ is well-typed under a terminated channel. Rule (T-PAR) types the parallel composition of two processes by composing their corresponding typing contexts using a splitting operator, noted \circ [19]. Rule (T-RES) types a restricted process by requiring that the two endpoints have dual types. Rules (T-IN) and (T-OUT) type the receiving and sending of a value over a channel x , respectively. Finally, rules (T-BRCH) and (T-SEL) are generalizations of input and output over a labelled set of processes.

The main guarantees of the type system are *communication safety* and *session fidelity*, i.e., typed processes respect their ascribed protocols, as represented by session types.

Theorem 2.1 (Type Preservation for Session Types). *If $\Gamma \vdash_{\text{ST}} P$ and $P \rightarrow Q$, then $\Gamma \vdash_{\text{ST}} Q$.*

The following notion of well-formed processes is key to single out meaningful typed processes.

Definition 2.2 (Well-Formedness for Sessions). *A process is well-formed if for any of its structural congruent processes of the form $(\nu \tilde{xy})(P \mid Q)$ the following hold.*

- *If P and Q are prefixed at the same variable, then the variable performs the same action (input or output, branching or selection).*

- If P is prefixed in x_i and Q is prefixed in y_i where $x_i y_i \in \tilde{x} \tilde{y}$, then $P \mid Q \rightarrow$.

It is important to notice that well-typedness of a process does not imply the process is well-formed. We have the following theorem:

Theorem 2.3 (Type Safety for Sessions [19]). *If $\vdash_{\text{ST}} P$ then P is well-formed.*

We present the main result of the session type system. The following theorem states that a well-typed closed process does not reduce to an ill-formed one. It follows immediately from Theorems 2.1 and 2.3.

Theorem 2.4 ([19]). *If $\vdash_{\text{ST}} P$ and $P \rightarrow^* Q$, then Q is well-formed.*

An important observation is that the session type system given above does not exclude *deadlocked processes*, i.e., processes which reach a “stuck state.” This is because the interleaving of communication prefixes in typed processes may create extra causal dependencies not described by session types. (This intuitive definition of deadlocked processes will be made precise below.) A particularly insidious class of deadlocks is due to cyclic interleaving of channels in processes. For example, consider a process such as $P \triangleq (\nu xy)(\nu wz)(\bar{x}\langle \mathbf{n} \rangle . \bar{w}\langle \mathbf{n} \rangle \mid z(t).y(s))$: it represents the implementation of two (simple) independent sessions, which get intertwined (blocked) due to the nesting induced by input and output prefixes. We have that $\mathbf{n} : \mathbf{end} \vdash_{\text{ST}} P$ even if P is unable to reduce. A deadlock-free variant of P would be, e.g., process $P' \triangleq (\nu xy)(\nu wz)(\bar{x}\langle \mathbf{n} \rangle . \bar{w}\langle \mathbf{n} \rangle \mid y(s).z(t))$, which also is typable in \vdash_{ST} .

We will say that a process is *deadlock-free* if any communication action that becomes active during execution is eventually consumed; that is, there is a corresponding co-action that eventually becomes available. Below we define deadlock freedom in the session π -calculus; we follow [13, 15] and consider *fair* reduction sequences [6]. For simplicity, we omit the symmetric cases for input and branching.

Definition 2.5 (Deadlock Freedom for Session π -Calculus). *A process P_0 is deadlock-free if for any fair reduction sequence $P_0 \rightarrow P_1 \rightarrow P_2 \rightarrow \dots$, we have that*

1. $P_i \equiv (\nu \tilde{x} \tilde{y})(\bar{x}\langle v \rangle . Q \mid R)$, for $i \geq 0$, implies that there exists $n \geq i$ such that $P_n \equiv (\nu x' y')(\bar{x}\langle v \rangle . Q \mid y(z).R_1 \mid R_2)$ and $P_{n+1} \equiv (\nu x' y')(Q \mid R_1[v/z] \mid R_2)$;
2. $P_i \equiv (\nu \tilde{x} \tilde{y})(x \triangleleft l_j . Q \mid R)$, for $i \geq 0$, implies that there exists $n \geq i$ such that $P_n \equiv (\nu x' y')(x \triangleleft l_j . Q \mid y \triangleright \{l_k : R_k\}_{k \in I \cup \{j\}} \mid S)$ and $P_{n+1} \equiv (\nu x' y')(Q \mid R_j \mid S)$.

3 Two Approaches to Deadlock Freedom

We introduce two approaches to deadlock-free, session typed processes. The first one, given in § 3.1, comes from interpretations of linear logic propositions as session types [1–3, 21]; the second approach, summarized in § 3.2, combines usage types for the standard π -calculus with encodings of session processes and types [8]. Based on these two approaches, in § 4 we will define the classes \mathcal{L} and \mathcal{H} .

3.1 Linear Logic Foundations of Session Types

The linear logic interpretation of session types was introduced by Caires and Pfenning [3], and developed by Wadler [21] and others. Initially proposed for intuitionistic linear logic, here we consider an interpretation based on classical linear logic with mix principles, following a recent presentation by Caires [1].

The syntax and semantics of processes are as in § 2 except for the following differences. First, we have the standard restriction construct $(\nu x)P$, which replaces the double restriction. Second, we have a so-called *forwarding process*, denoted $[x \leftrightarrow y]$, which intuitively “fuses” names x and y . Besides these

$$\begin{array}{c}
\text{(T-1)} \frac{}{\mathbf{0} \vdash_{\text{CH}} x : \bullet} \quad \text{(T-}\perp\text{)} \frac{P \vdash_{\text{CH}} \Delta}{P \vdash_{\text{CH}} x : \bullet, \Delta} \quad \text{(T-id)} \frac{}{[x \leftrightarrow y] \vdash_{\text{CH}} x : A, y : \bar{A}} \\
\text{(T-}\wp\text{)} \frac{P \vdash_{\text{CH}} \Delta, y : A, x : B}{x(y).P \vdash_{\text{CH}} \Delta, x : A \wp B} \quad \text{(T-}\otimes\text{)} \frac{P \vdash_{\text{CH}} \Delta, y : A \quad Q \vdash_{\text{CH}} \Delta', x : B}{\bar{x}(y).(P \mid Q) \vdash_{\text{CH}} \Delta, \Delta', x : A \otimes B} \quad \text{(T-cut)} \frac{P \vdash_{\text{CH}} \Delta, x : \bar{A} \quad Q \vdash_{\text{CH}} \Delta', x : A}{(vx)(P \mid Q) \vdash_{\text{CH}} \Delta, \Delta'} \\
\text{(T-}\oplus\text{)} \frac{P \vdash_{\text{CH}} \Delta, x : A_j \quad j \in I}{x \triangleleft l_j.P \vdash_{\text{CH}} \Delta, x : \oplus \{l_i : A_i\}_{i \in I}} \quad \text{(T-}\&\text{)} \frac{P_i \vdash_{\text{CH}} \Delta, x : A_i \quad \forall i \in I}{x \triangleright \{l_i : P_i\}_{i \in I} \vdash_{\text{CH}} \Delta, x : \& \{l_i : A_i\}_{i \in I}} \quad \text{(T-mix)} \frac{P \vdash_{\text{CH}} \Delta \quad Q \vdash_{\text{CH}} \Delta'}{P \mid Q \vdash_{\text{CH}} \Delta, \Delta'}
\end{array}$$

Figure 3: Typing rules for the π -calculus with C-types.

differences in syntax, we have also some minor modifications in reduction rules. Differences with respect to the language considered in § 2 are summarized in the following:

$$P, Q ::= (vx)P \quad (\text{channel restriction}) \quad | \quad [x \leftrightarrow y] \quad (\text{forwarding})$$

$$\begin{array}{ll}
\text{(R-CHCOM)} \quad \bar{x}(v).P \mid x(z).Q \rightarrow P \mid Q[v/z] & \text{(R-FWD)} \quad (vx)([x \leftrightarrow y] \mid P) \rightarrow P[y/x] \\
\text{(R-CHCASE)} \quad x \triangleleft l_j.P \mid x \triangleright \{l_i : P_i\}_{i \in I} \rightarrow P \mid P_j \quad j \in I & \text{(R-CHRES)} \quad P \rightarrow Q \implies (vx)P \rightarrow (vx)Q
\end{array}$$

Observe how interaction of input/output prefixes and selection/branching is no longer covered by an outermost restriction. As for the type system, we consider the so-called C-types which correspond to linear logic propositions. They are given by the following grammar:

$$A, B ::= \perp \mid \mathbf{1} \mid A \otimes B \mid A \wp B \mid \oplus \{l_i : A_i\}_{i \in I} \mid \& \{l_i : A_i\}_{i \in I}$$

Intuitively, \perp and $\mathbf{1}$ are used to type a terminated endpoint. Type $A \otimes B$ is associated to an endpoint that first outputs an object of type A and then behaves according to B . Dually, type $A \wp B$ is the type of an endpoint that first inputs an object of type A and then continues as B . The interpretation of $\oplus \{l_i : A_i\}_{i \in I}$ and $\& \{l_i : A_i\}_{i \in I}$ as select and branch behaviors follows as expected.

We define a full duality on C-types, which exactly corresponds to the negation operator of CLL $(\cdot)^\perp$. The *dual* of type A , denoted \bar{A} , is inductively defined as follows:

$$\begin{array}{lll}
\bar{\mathbf{1}} = \perp & \bar{\perp} = \mathbf{1} & \overline{\oplus \{l_i : A_i\}_{i \in I}} = \& \{l_i : \bar{A}_i\}_{i \in I} \\
\overline{A \otimes B} = \bar{A} \wp \bar{B} & \overline{A \wp B} = \bar{A} \otimes \bar{B} & \overline{\& \{l_i : A_i\}_{i \in I}} = \oplus \{l_i : \bar{A}_i\}_{i \in I}
\end{array}$$

Recall that $A \multimap B \triangleq \bar{A} \wp B$. As explained in [1], considering mix principles means admitting $\perp \multimap \mathbf{1}$ and $\mathbf{1} \multimap \perp$, and therefore $\perp = \mathbf{1}$. We write \bullet to denote either \perp or $\mathbf{1}$, and decree that $\bar{\bullet} = \bullet$.

Typing contexts, sets of typing assignments $x : A$, are ranged over Δ, Δ', \dots . The empty context is denoted ‘ \cdot ’. Typing judgments are then of the form $P \vdash_{\text{CH}} \Delta$. Figure 3 gives the typing rules associated to the linear logic interpretation. Salient points include the use of bound output $(vy)\bar{x}(y).P$, which is abbreviated as $\bar{x}(y)P$. Another highlight is the ‘‘composition plus hiding’’ principle implemented by rule (T-cut), which integrates parallel composition and restriction in a single rule. Indeed, there is no dedicated rule for restriction. Also, rule (T-mix) enables the typing of *independent parallel compositions*, i.e., the composition of two processes that do not share sessions.

We now collect main results for this type system; see [1, 3] for details. For any P , define *live*(P) if and only if $P \equiv (v\tilde{n})(\pi.Q \mid R)$, where π is an input, output, selection, or branching prefix.

$$\begin{array}{ll}
U ::= \text{?}_{\kappa}.U & \text{(used in input)} & \emptyset & \text{(not usable)} \\
& \text{!}_{\kappa}.U & \text{(used in output)} & (U_1 \mid U_2) \text{ (used in parallel)} \\
T ::= U[\tilde{T}] & \text{(channel types)} & \langle l : T \rangle_{i \in I} & \text{(variant type)}
\end{array}$$

Figure 4: Syntax of usage types for the π -calculus.

Theorem 3.1 (Type Preservation for C-Types). *If $P \vdash_{\text{CH}} \Delta$ and $P \rightarrow Q$ then $Q \vdash_{\text{CH}} \Delta$.*

Theorem 3.2 (Progress). *If $P \vdash_{\text{CH}} \cdot$ and $\text{live}(P)$ then $P \rightarrow Q$, for some Q .*

3.2 Deadlock Freedom by Encodability

As mentioned above, the second approach to deadlock-free session processes is *indirect*, in the sense that establishing deadlock freedom for session processes appeals to usage types for the π -calculus [13, 15], for which type systems enforcing deadlock freedom are well-established. Formally, this reduction exploits encodings of processes and types: a session process $\Gamma \vdash_{\text{ST}} P$ is encoded into a (standard) π -calculus process $\llbracket \Gamma \rrbracket_f \vdash_{\text{KB}}^n \llbracket P \rrbracket_f$. Next we introduce the syntax of standard π -calculus processes with variant values (§ 3.2.1), the discipline of usage types (§ 3.2.2), and the encodings of session processes and types into standard π -calculus processes and usage types, respectively (§ 3.2.3).

3.2.1 Processes

The syntax and semantics of the π -calculus with usage types build upon those in § 2. We require some modifications. First, the encoding of terms presented in § 3.2.3, requires polyadic communication. Rather than branching and selection constructs, the π -calculus that we consider here includes a *case* construct **case** v of $\{l_i \text{-} x_i \triangleright P_i\}_{i \in I}$ that uses *variant value* $l_j \text{-} v$. Moreover, we consider the standard channel restriction, rather than double restriction. These modifications are summarized below:

$$\begin{array}{ll}
P, Q ::= (vx)P & \text{(channel restriction)} & \mid \text{case } v \text{ of } \{l_i \text{-} x_i \triangleright P_i\}_{i \in I} & \text{(case)} \\
v ::= l_j \text{-} v & \text{(variant value)}
\end{array}$$

$$\begin{array}{ll}
(\text{R}\pi\text{-COM}) & \bar{x}(\tilde{v}).P \mid x(\tilde{z}).Q \rightarrow P \mid Q[\tilde{v}/\tilde{z}] \\
(\text{R}\pi\text{-RES}) & P \rightarrow Q \implies (vx)P \rightarrow (vx)Q \\
(\text{R}\pi\text{-CASE}) & \text{case } l_j \text{-} v \text{ of } \{l_i \text{-} x_i \triangleright P_i\}_{i \in I} \rightarrow P_j[v/x_i] \quad j \in I
\end{array}$$

The definition of deadlock-freedom for the π -calculus follows [13, 15]:

Definition 3.3 (Deadlock Freedom for Standard π -Calculus). *A process P_0 is deadlock-free under fair scheduling, if for any fair reduction sequence $P_0 \rightarrow P_1 \rightarrow P_2 \rightarrow \dots$ the following hold*

1. *if $P_i \equiv (v\tilde{x})(\bar{x}(\tilde{v}).Q \mid R)$ for $i \geq 0$, implies that there exists $n \geq i$ such that $P_n \equiv (v\tilde{x})(\bar{x}(\tilde{v}).Q \mid x(\tilde{z}).R_1 \mid R_2)$ and $P_{n+1} \equiv (v\tilde{x})(Q \mid R_1[\tilde{v}/\tilde{z}] \mid R_2)$;*
2. *if $P_i \equiv (v\tilde{x})(x(\tilde{z}).Q \mid R)$ for $i \geq 0$, implies that there exists $n \geq i$ such that $P_n \equiv (v\tilde{x})(x(\tilde{z}).Q \mid \bar{x}(\tilde{v}).R_1 \mid R_2)$ and $P_{n+1} \equiv (v\tilde{x})(Q[\tilde{v}/\tilde{z}] \mid R_1 \mid R_2)$.*

3.2.2 Usage Types

The syntax of usage types is defined in Figure 4. For simplicity, we let α range over input ? or output ! actions. The usage \emptyset describes a channel that cannot be used at all. We will often omit \emptyset , and so we

will write U instead of $U.\emptyset$. Usages $?_{\kappa}^o.U$ and $!_{\kappa}^o.U$ describe channels that can be used once for input and output, respectively and then used according to the continuation usage U . The *obligation* o and *capability* κ range over the set of natural numbers. The usage $U_1 \mid U_2$ describes a channel that is used according to U_1 by one process and U_2 by another processes in parallel.

Intuitively, obligations and capabilities describe inter-channel dependencies:

- An obligation of level n must be fulfilled by using only capabilities of level *less than* n . Said differently, an action of obligation n must be prefixed by actions of capabilities less than n .
- For an action with capability of level n , there must exist a co-action with obligation of level *less than or equal to* n .

Typing contexts are sets of typing assignments and are ranged over Γ, Γ' . A typing judgement is of the form $\Gamma \vdash_{\text{KB}}^n P$: the annotation n explicitly denotes the greatest *degree of sharing* admitted in parallel processes. Before commenting on the typing rules (given in Figure 5), we discuss some important auxiliary notions, extracted from [13, 15]. First, the composition operation on types (denoted \mid , and used in rules $\text{T}\pi\text{-}(\text{PAR})_n$ and $\text{T}\pi\text{-}(\text{OUT})$) is based on the composition of usages and is defined as follows:

$$\langle l_i : T_i \rangle_{i \in I} \mid \langle l_i : T_i \rangle_{i \in I} = \langle l_i : T_i \rangle_{i \in I} \quad U_1[\tilde{T}] \mid U_2[\tilde{T}] = (U_1 \mid U_2)[\tilde{T}]$$

The generalization of \mid to typing contexts, denoted $(\Gamma_1 \mid \Gamma_2)(x)$, is defined as expected. The unary operation \uparrow^t applied to a usage U lifts its obligation level *up to* t ; it is defined inductively as:

$$\uparrow^t \emptyset = \emptyset \quad \uparrow^t \alpha_{\kappa}^o.U = \alpha_{\kappa}^{\max(o,t)}.U \quad \uparrow^t (U_1 \mid U_2) = (\uparrow^t U_1 \mid \uparrow^t U_2)$$

The \uparrow^t is extends to types/typing contexts as expected. *Duality* on usage types simply exchanges $?$ and $!$:

$$\overline{\emptyset} = \emptyset \quad \overline{?_{\kappa}^o.U[\tilde{T}]} = !_{\kappa}^o.\overline{U}[\tilde{T}] \quad \overline{!_{\kappa}^o.U[\tilde{T}]} = ?_{\kappa}^o.\overline{U}[\tilde{T}]$$

Operator “ $;$ ” in $\Delta = x : [T]\alpha_{\kappa}^o ; \Gamma$, used in rules $(\text{T}\pi\text{-}(\text{IN}))$ and $(\text{T}\pi\text{-}(\text{OUT}))$, is such that the following hold:

$$\text{dom}(\Delta) = \{x\} \cup \text{dom}(\Gamma) \quad \Delta(x) = \begin{cases} \alpha_{\kappa}^o.U[\tilde{T}] & \text{if } \Gamma(x) = U[\tilde{T}] \\ \alpha_{\kappa}^o[\tilde{T}] & \text{if } x \notin \text{dom}(\Gamma) \end{cases} \quad \Delta(y) = \uparrow^{\kappa+1} \Gamma(y) \text{ if } y \neq x$$

The final required notion is that of a *reliable usage*. It builds upon the following definition:

Definition 3.4. *Let U be a usage. The input and output obligation levels (resp. capability levels) of U , written $\text{ob}_?(U)$ and $\text{ob}_!(U)$ (resp. $\text{cap}_?(U)$ and $\text{cap}_!(U)$), are defined as:*

$$\begin{aligned} \text{ob}_{\alpha}(\alpha_{\kappa}^o.U) &= o & \text{cap}_{\alpha}(\alpha_{\kappa}^o.U) &= \kappa \\ \text{ob}_{\alpha}(U_1 \mid U_2) &= \min(\text{ob}_{\alpha}(U_1), \text{ob}_{\alpha}(U_2)) & \text{cap}_{\alpha}(U_1 \mid U_2) &= \min(\text{cap}_{\alpha}(U_1), \text{cap}_{\alpha}(U_2)) \end{aligned}$$

The definition of reliable usages depends on a reduction relation on usages, noted $U \rightarrow U'$. Intuitively, $U \rightarrow U'$ means that if a channel of usage U is used for communication, then after the communication occurs, the channel should be used according to usage U' . Thus, e.g., $?_{\kappa}^o.U_1 \mid ?_{\kappa}^o.U_2$ reduces to $U_1 \mid U_2$.

Definition 3.5 (Reliability). *We write $\text{con}_{\alpha}(U)$ when $\text{ob}_{\alpha}(U) \leq \text{cap}_{\alpha}(U)$. We write $\text{con}(U)$ when $\text{con}_?(U)$ and $\text{con}_!(U)$ hold. Usage U is reliable, noted $\text{rel}(U)$, if $\text{con}(U')$ holds $\forall U'$ such that $U \rightarrow^* U'$.*

$$\begin{array}{c}
\text{(T}\pi\text{-NIL)} \\
\frac{}{x : \emptyset \vdash_{\text{KB}}^n \mathbf{0}} \\
\\
\text{(T}\pi\text{-RES)} \\
\frac{\Gamma, x : U[\tilde{T}] \vdash_{\text{KB}}^n P \quad \text{rel}(U)}{\Gamma \vdash_{\text{KB}}^n (\nu x)P} \\
\\
\text{(T}\pi\text{-PAR}_n) \\
\frac{\Gamma_1 \vdash_{\text{KB}}^n P \quad \Gamma_2 \vdash_{\text{KB}}^n Q \quad |\Gamma_1 \cap \Gamma_2| \leq n}{\Gamma_1 \mid \Gamma_2 \vdash_{\text{KB}}^n P \mid Q} \\
\\
\text{(T}\pi\text{-IN)} \\
\frac{\Gamma, \tilde{y} : \tilde{T} \vdash_{\text{KB}}^n P}{x : ?_{\text{K}}^0[\tilde{T}] ; \Gamma \vdash_{\text{KB}}^n x(\tilde{y}).P} \\
\\
\text{(T}\pi\text{-OUT)} \\
\frac{\Gamma_1 \vdash_{\text{KB}}^n \tilde{v} : \tilde{T} \quad \Gamma_2 \vdash_{\text{KB}}^n P}{x : !_{\text{K}}^0[\tilde{T}] ; (\Gamma_1 \mid \Gamma_2) \vdash_{\text{KB}}^n \bar{x}(\tilde{v}).P} \\
\\
\text{(T}\pi\text{-LVAL)} \\
\frac{\Gamma \vdash_{\text{KB}}^n v : T_j \quad \exists j \in I}{\Gamma \vdash_{\text{KB}}^n l_{j \rightarrow v} : \langle l_i : T_i \rangle_{i \in I}} \\
\\
\text{(T}\pi\text{-CASE)} \\
\frac{\Gamma_1 \vdash_{\text{KB}}^n v : \langle l_i : T_i \rangle_{i \in I} \quad \Gamma_2, x_i : T_i \vdash_{\text{KB}}^n P_i \quad \forall i \in I}{\Gamma_1, \Gamma_2 \vdash_{\text{KB}}^n \mathbf{case } v \mathbf{ of } \{ l_i \cdot x_i \triangleright P_i \}_{i \in I}}
\end{array}$$

Figure 5: Typing rules for the π -calculus with usage types with degree of sharing n .

Typing Rules. The typing rules for the standard π -calculus with usage types are given in Figure 5. The only difference with respect to the rules in Kobayashi’s systems [13, 15] is that we annotate typing judgements with the degree of sharing, explicitly stated in rule (T π -PAR $_n$)—see below. Rule (T π -NIL) states that the terminated process is typed under a terminated channel. Rule (T π -RES) states that process $(\nu x)P$ is well-typed if the usage for x is reliable (cf. Definition 3.5). Rules (T π -IN) and (T π -OUT) type input and output processes in a typing context where the “;” operator is used in order to increase the obligation level of the channels in continuation P . Rules (T π -LVAL) and (T π -CASE) type a choice: the first types a variant value with a variant type; the second types a case process using a variant value as its guard.

Given a degree of sharing n , rule (T π -PAR $_n$) states that the parallel composition of processes P and Q (typable under contexts Γ_1 and Γ_2 , respectively) is well-typed under the typing context $\Gamma_1 \mid \Gamma_2$ only if $|\Gamma_1 \cap \Gamma_2| \leq n$. This allows to simply characterize the “concurrent cooperation” between P and Q . As a consequence, if $P \vdash_{\text{KB}}^n$ then $P \vdash_{\text{KB}}^k$, for any $k \leq n$. Observe that the typing rule for parallel composition in [13, 15] is the same as (T π -PAR $_n$), except for condition $|\Gamma_1 \cap \Gamma_2| \leq n$, which is not specified.

The next theorems imply that well-typed processes by the type system in Figure 5 are deadlock-free.

Theorem 3.6 (Type Preservation for Usage Types). *If $\Gamma \vdash_{\text{KB}}^n P$ and $P \rightarrow Q$, then $\Gamma' \vdash_{\text{KB}}^n Q$ for some Γ' such that $\Gamma \rightarrow \Gamma'$.*

Theorem 3.7 (Deadlock Freedom). *If $\emptyset \vdash_{\text{KB}}^n P$ and either $P \equiv (\nu \tilde{x})(x(\tilde{z}).Q \mid R)$ or $P \equiv (\nu \tilde{x})(\bar{x}(\tilde{v}).Q \mid R)$, then $P \rightarrow Q$, for some Q .*

Corollary 3.8. *If $\emptyset \vdash_{\text{KB}}^n P$, then P is deadlock-free, in the sense of Definition 3.3.*

Theorem 3.2 (progress for the linear logic system) and Theorem 3.7 (deadlock freedom for the standard π -calculus) have a rather similar formulation: both properties state that processes can always reduce if they are well-typed (under the empty typing context) and have an appropriate structure (i.e., condition $\text{live}(P)$ in Theorem 3.2 and condition $P \equiv (\nu \tilde{x})(x(\tilde{z}).Q \mid R)$ or $P \equiv (\nu \tilde{x})(\bar{x}(\tilde{v}).Q \mid R)$ in Theorem 3.7).

3.2.3 Encodings of Processes and Types

Encoding of Processes. To relate classes of processes obtained by the different type systems given so far, we rewrite a session typed or C-typed process into a usage typed process by following a continuation-passing style: this allows us to mimic the structure of a session or C-type by sending its continuation as a payload over a channel. This idea, suggested in [14] and developed in [8], is recalled in Figure 6.

$$\begin{array}{ll}
\llbracket \bar{x}(v).P \rrbracket_f \triangleq (\nu c) \bar{f}_x \langle v, c \rangle . \llbracket P \rrbracket_{f, \{x \rightarrow c\}} & \llbracket x \triangleright \{l_i : P_i\}_{i \in I} \rrbracket_f \triangleq f_x(y). \mathbf{case} \ y \mathbf{of} \ \{l_{i-c} \triangleright \llbracket P_i \rrbracket_{f, \{x \rightarrow c\}}\}_{i \in I} \\
\llbracket x(y).P \rrbracket_f \triangleq f_x(y, c) . \llbracket P \rrbracket_{f, \{x \rightarrow c\}} & \llbracket (\nu xy)P \rrbracket_f \triangleq (\nu c) \llbracket P \rrbracket_{f, \{x, y \rightarrow c\}} \\
\llbracket x \triangleleft l_j . P \rrbracket_f \triangleq (\nu c) \bar{f}_x \langle l_j, c \rangle . \llbracket P \rrbracket_{f, \{x \rightarrow c\}} & \llbracket P \mid Q \rrbracket_f \triangleq \llbracket P \rrbracket_f \mid \llbracket Q \rrbracket_f
\end{array}$$

Figure 6: Encoding of session processes into π -calculus processes.

$$\begin{array}{ll}
\llbracket \mathbf{end} \rrbracket_{\text{su}} = \emptyset & \llbracket \mathbf{end} \rrbracket_{\text{c}} = \bullet \\
\llbracket ?T.S \rrbracket_{\text{su}} = ?_{\kappa}^{\circ} \llbracket T \rrbracket_{\text{su}}, \llbracket S \rrbracket_{\text{su}} & \llbracket ?T.S \rrbracket_{\text{c}} = \llbracket T \rrbracket_{\text{c}} \wp \llbracket S \rrbracket_{\text{c}} \\
\llbracket !T.S \rrbracket_{\text{su}} = !_{\kappa}^{\circ} \llbracket T \rrbracket_{\text{su}}, \llbracket \bar{S} \rrbracket_{\text{su}} & \llbracket !T.S \rrbracket_{\text{c}} = \llbracket \bar{T} \rrbracket_{\text{c}} \otimes \llbracket S \rrbracket_{\text{c}} \\
\llbracket \&\{l_i : S_i\}_{i \in I} \rrbracket_{\text{su}} = ?_{\kappa}^{\circ} \langle l_i : \llbracket S_i \rrbracket_{\text{su}} \rangle_{i \in I} & \llbracket \&\{l_i : S_i\}_{i \in I} \rrbracket_{\text{c}} = \&\{l_i : \llbracket S_i \rrbracket_{\text{c}}\}_{i \in I} \\
\llbracket \oplus\{l_i : S_i\}_{i \in I} \rrbracket_{\text{su}} = !_{\kappa}^{\circ} \langle l_i : \llbracket \bar{S}_i \rrbracket_{\text{su}} \rangle_{i \in I} & \llbracket \oplus\{l_i : S_i\}_{i \in I} \rrbracket_{\text{c}} = \oplus\{l_i : \llbracket S_i \rrbracket_{\text{c}}\}_{i \in I}
\end{array}$$

Figure 7: Encodings of session types into usage types (Left) and C-types (Right).

Encoding of Types. We formally relate session types and logic propositions to usage types by means of the encodings given in Figure 7. The former one, denoted as $\llbracket \cdot \rrbracket_{\text{su}}$, is taken from [8].

Definition 3.9. Let Γ be a session typing context. The encoding $\llbracket \cdot \rrbracket_f$ into usage typing context and $\llbracket \cdot \rrbracket_{\text{c}}$ into C-typing context is inductively defined as follows:

$$\llbracket \emptyset \rrbracket_f = \llbracket \emptyset \rrbracket_{\text{c}} \triangleq \emptyset \quad \llbracket \Gamma, x : T \rrbracket_f \triangleq \llbracket \Gamma \rrbracket_f, f_x : \llbracket T \rrbracket_{\text{su}} \quad \llbracket \Gamma, x : T \rrbracket_{\text{c}} \triangleq \llbracket \Gamma \rrbracket_{\text{c}}, x : \llbracket T \rrbracket_{\text{c}}$$

Lemma 3.10 (Duality and encoding of session types). Let T, S be finite session types.

Then: (i) $\bar{T} = S$ if and only if $\llbracket \bar{T} \rrbracket_{\text{c}} = \llbracket S \rrbracket_{\text{c}}$; (ii) $\bar{T} = S$ if and only if $\llbracket \bar{T} \rrbracket_{\text{su}} = \llbracket S \rrbracket_{\text{su}}$.

On Deadlock Freedom by Encoding. The next results relate deadlock freedom, typing and encoding.

Proposition 3.11. Let P be a deadlock-free session process, then $\llbracket P \rrbracket_f$ is a deadlock-free π -process.

Proof. Follows by the encoding of terms given in Figure 6, Definition 2.5 and Definition 3.3. \square

Next we recall an important result relating deadlock freedom and typing, by following [4].

Corollary 3.12. Let $\vdash_{\text{ST}} P$ be a session process. If $\vdash_{\text{KB}}^n \llbracket P \rrbracket_f$ is deadlock-free then P is deadlock-free.

4 A Hierarchy of Deadlock-Free Session Typed Processes

Preliminaries. To formally define the classes \mathcal{L} and \mathcal{H} , we require some auxiliary definitions. The following translation addresses minor syntactic differences between session typed processes (cf. § 2) and the processes typable in the linear logic interpretation of session types (cf. § 3.1). Such differences concern output actions and the restriction operator:

Definition 4.1. Let P be a session process. The translation $\{\!\{ \cdot \}\!\}$ is defined as

$$\{\!\{ \bar{x}(y).P \}\!\} = \bar{x}(z).([z \leftrightarrow y] \mid \{\!\{ P \}\!\}) \quad \{\!\{ (\nu xy)P \}\!\} = (\nu w)\{\!\{ P \}\!\}[w/x][w/y] \quad w \notin \text{fn}(P)$$

and as an homomorphism for the other process constructs.

Let $\llbracket \cdot \rrbracket_c$ denote the encoding of session types into linear logic propositions in Figure 7 (right). Recall that $\llbracket \cdot \rrbracket_f$ stands for the encoding of processes and $\llbracket \cdot \rrbracket_{su}$ for the encoding of types, both defined in [8], and given here in Figure 6 and Figure 7 (left), respectively. We may then formally define the languages under comparison as follows:

Definition 4.2 (Typed Languages). *The languages \mathcal{L} and \mathcal{K}_n ($n \geq 0$) are defined as follows:*

$$\begin{aligned} \mathcal{L} &= \{P \mid \exists \Gamma. (\Gamma \vdash_{ST} P \wedge \{\{P\}\} \vdash_{CH} \llbracket \Gamma \rrbracket_c)\} \\ \mathcal{K}_n &= \{P \mid \exists \Gamma, f. (\Gamma \vdash_{ST} P \wedge \llbracket \Gamma \rrbracket_f \vdash_{KB}^n \llbracket P \rrbracket_f)\} \end{aligned}$$

Main Results. Our first observation is that there are processes in \mathcal{K}_2 but not in \mathcal{K}_1 :

Lemma 4.3. $\mathcal{K}_1 \subset \mathcal{K}_2$.

Proof. \mathcal{K}_2 contains (deadlock-free) session processes not captured in \mathcal{K}_1 . A representative example is:

$$P_2 = (\nu a_1 b_1)(\nu a_2 b_2)(a_1(x). \overline{a_2}\langle x \rangle \mid \overline{b_1}\langle \mathbf{n} \rangle. b_2(z))$$

This process is not in \mathcal{K}_1 because it involves the composition of two parallel processes which share two sessions. As such, it is typable in \vdash_{KB}^n (with $n \geq 2$) but not in \vdash_{KB}^1 . \square

The previous result generalizes easily, so as to define a hierarchy of deadlock-free, session processes:

Theorem 4.4. *For all $n \geq 1$, we have that $\mathcal{K}_n \subset \mathcal{K}_{n+1}$.*

Proof. Immediate by considering one of the following processes, which generalize process P_2 in Lemma 4.3:

$$\begin{aligned} P_{n+1} &= (\nu a_1 b_1)(\nu a_2 b_2) \cdots (\nu a_{n+1} b_{n+1})(a_1(x). \overline{a_2}\langle x \rangle. \cdots . \overline{a_{n+1}}\langle y \rangle \mid \overline{b_1}\langle \mathbf{n} \rangle. b_2(z). \cdots . b_{n+1}(z)) \\ Q_{n+1} &= (\nu a_1 b_1)(\nu a_2 b_2) \cdots (\nu a_{n+1} b_{n+1})(a_1(x). \overline{a_2}\langle x \rangle. \cdots . a_{n+1}(y) \mid \overline{b_1}\langle \mathbf{n} \rangle. b_2(z). \cdots . \overline{b_{n+1}}\langle \mathbf{n} \rangle) \end{aligned}$$

To distinguish \mathcal{K}_{n+1} from \mathcal{K}_n , we consider P_{n+1} if $n+1$ is even and Q_{n+1} otherwise. \square

One main result of this paper is that \mathcal{L} and \mathcal{K}_1 coincide. Before stating this result, we make the following observations. The typing rules for processes in \mathcal{L} do not directly allow free output. However, free output is representable (and typable) by linear logic types by means of the transformation in Definition 4.1. Thus, considered processes are not syntactically equal. In \mathcal{L} there is cooperating composition (enabled by rule (T-cut) in Figure 3); independent composition can only be enabled by rule (T-mix). Arbitrary restriction is not allowed; only restriction of parallel processes.

The following property is key in our developments: it connects our encodings of (dual) session types into usage types with reliability (Definition 3.5), a central notion to the type system for deadlock freedom in Figure 5. Recall that, unlike usage types, there is no parallel composition operator at the level of session types.

Proposition 4.5. *Let T be a session type. Then $\text{rel}(\llbracket T \rrbracket_{su} \mid \llbracket \overline{T} \rrbracket_{su})$ holds.*

Proof (Sketch). By induction on the structure of session type T and the definitions of $\llbracket \cdot \rrbracket_{su}$ and predicate $\text{rel}(\cdot)$, using Lemma 3.10 (encodings of types preserve session type duality). See [9] for details. \square

We then have the following main result, whose proof is detailed in [9]:

Theorem 4.6. $\mathcal{L} = \mathcal{K}_1$.

Therefore, we have the following corollary, which attests that the class of deadlock-free session processes naturally induced by linear logic interpretations of session types is strictly included in the class induced by the indirect approach of Dardha et al. [8] (cf. § 3.2).

Corollary 4.7. $\mathcal{L} \subset \mathcal{K}_n, n > 1$.

The fact that (deadlock-free) processes such as P_2 (cf. Lemma 4.3) are not in \mathcal{L} is informally discussed in [3, §6]. However, [3] gives no formal comparisons with other classes of deadlock-free processes.

5 Rewriting \mathcal{K}_n into \mathcal{L}

The hierarchy of deadlock-free session processes established by Theorem 4.4 is *subtle* in the following sense: if $P \in \mathcal{K}_{k+1}$ but $P \notin \mathcal{K}_k$ (with $k \geq 1$) then we know that there is a subprocess of P that needs to be “adjusted” in order to “fit in” \mathcal{K}_k . More precisely, we know that such a subprocess of P must become more independent in order to be typable under the lesser degree of sharing k .

Here we propose a *rewriting procedure* that converts processes in \mathcal{K}_n into processes in \mathcal{K}_1 (that is, \mathcal{L} , by Theorem 4.6). The rewriting procedure follows a simple idea: given a parallel process as input, return as output a process in which one of the components is kept unchanged, but the other is replaced by parallel representatives of the sessions implemented in it. Such parallel representatives are formally defined as characteristic processes and catalyzers, introduced next. The rewriting procedure is type preserving and satisfies operational correspondence (cf. Theorems 5.8 and 5.10).

5.1 Preliminaries: Characteristic Processes and Catalyzers

Before presenting our rewriting procedure, let us first introduce some preliminary results.

Definition 5.1 (Characteristic Processes of a Session Type). *Let T be a session type (cf. § 2). Given a name x , the set of characteristic processes of T , denoted $\{\!\{T}\!\}^x$, is inductively defined as follows:*

$$\begin{aligned} \{\!\{\mathbf{end}\}\!\}^x &= \{P \mid P \vdash_{\text{CH}} x : \bullet\} \\ \{\!\{?T.S\}\!\}^x &= \{x(y).P \mid P \vdash_{\text{CH}} y : \llbracket T \rrbracket_c, x : \llbracket S \rrbracket_c\} \\ \{\!\{!T.S\}\!\}^x &= \{\bar{x}(y).(P \mid Q) \mid P \in \{\!\{\bar{T}\}\!\}^y \wedge Q \in \{\!\{S\}\!\}^x\} \\ \{\!\{\&\{l_i : S_i\}_{i \in I}\}\!\}^x &= \{x \triangleright \{l_i : P_i\}_{i \in I} \mid \forall i \in I. P_i \in \{\!\{S_i\}\!\}^x\} \\ \{\!\{\oplus\{l_i : S_i\}_{i \in I}\}\!\}^x &= \bigcup_{i \in I} \{x \triangleleft l_i.P_i \mid P_i \in \{\!\{S_i\}\!\}^x\} \end{aligned}$$

Definition 5.2 (Catalyzer). *Given a session typing context Γ , we define its associated catalyzer as a process context $\mathcal{C}_\Gamma[\cdot]$, as follows:*

$$\mathcal{C}_\emptyset[\cdot] = [\cdot] \quad \mathcal{C}_{\Gamma, x:T}[\cdot] = (\nu x)(\mathcal{C}_\Gamma[\cdot] \mid P) \quad \text{with } P \in \{\!\{\bar{T}\}\!\}^x$$

We record the fact that characteristic processes are well-typed in the system of § 3.1:

Lemma 5.3. *Let T be a session type. For all $P \in \{\!\{T}\!\}^x$, we have: $P \vdash_{\text{CH}} x : \llbracket T \rrbracket_c$*

We use $\{\!\{T\}\!\}^x \vdash_{\text{CH}} x : \llbracket T \rrbracket_c$ to denote the set of processes $P \in \{\!\{T\}\!\}^x$ such that $P \vdash_{\text{CH}} x : \llbracket T \rrbracket_c$.

Lemma 5.4 (Catalyzers Preserve Typability). *Let $\Gamma \vdash_{\text{ST}} P$ and $\Gamma' \subseteq \Gamma$. Then $\mathcal{C}_{\Gamma'}[P] \vdash_{\text{CH}} \llbracket \Gamma \rrbracket_c \setminus \llbracket \Gamma' \rrbracket_c$.*

Corollary 5.5. *Let $\Gamma \vdash_{\text{ST}} P$. Then $\mathcal{C}_\Gamma[P] \vdash_{\text{CH}} \emptyset$.*

5.2 Rewriting \mathcal{K}_n in \mathcal{L}

We start this section with some notations. First, in order to represent pseudo-non deterministic binary choices between two equally typed processes, we introduce the following:

Notation 5.6. Let P_1, P_2 be two processes such that $k \notin \text{fn}(P_1, P_2)$. We write $P_1 \parallel_k P_2$ to stand for the process $(\nu k)(k \triangleleft \text{inx}.\mathbf{0} \mid k \triangleright \{\text{inl} : P_1, \text{inr} : P_2\})$, where label inx stands for either inl or inr .

Clearly, since session execution is purely deterministic, notation $P_1 \parallel_k P_2$ denotes that either P_1 or P_2 will be executed (and that the actual deterministic choice is not relevant). It is worth adding that Caires has already developed the technical machinery required to include non deterministic behavior into the linear logic interpretation of session types; see [1]. Casting our rewriting procedure into the typed framework of [1], so as to consider actual non deterministic choices, is interesting future work.

We find it convenient to annotate bound names in processes with session types, and write $(\nu xy : T)P$ and $x(y : T).P$, for some session type T . When the reduction relation involves a left or right choice in a binary labelled choice, as in reductions due to pseudo-non deterministic choices (Notation 5.6), we sometimes annotate the reduction as \rightarrow^{inl} or \rightarrow^{inr} . We let \mathbb{C} denote a *process context*, i.e., a process with a hole. And finally, for a typing context Γ , we shall write $\{\Gamma\}$ to denote the process $\prod_{(w_i : T_i) \in \Gamma} \{\{T_i\}^{w_i}\}$. We are now ready to give the rewriting procedure from \mathcal{K}_n to \mathcal{L} .

Definition 5.7 (Rewriting \mathcal{K}_n into \mathcal{L}). Let $P \in \mathcal{K}_n$ such that $\Gamma \vdash_{\text{ST}} P$, for some Γ . The encoding $(\Gamma \vdash_{\text{ST}} P)$ is a process of \mathcal{L} inductively defined as follows:

$$\begin{aligned}
(x : \mathbf{end} \vdash_{\text{ST}} \mathbf{0}) &\triangleq \mathbf{0} \\
(\Gamma \vdash_{\text{ST}} \bar{x}\langle v \rangle.P') &\triangleq \bar{x}(z).([v \leftrightarrow z] \mid (\Gamma', x : S \vdash_{\text{ST}} P')) & \Gamma = \Gamma', x : !T.S, v : T \\
(\Gamma \vdash_{\text{ST}} x(y : T).P') &\triangleq x(y).(\Gamma', x : S, y : T \vdash_{\text{ST}} P') & \Gamma = \Gamma', x : ?T.S \\
(\Gamma \vdash_{\text{ST}} x \triangleleft l_j.P') &\triangleq x \triangleleft l_j.(\Gamma', x : S_j \vdash_{\text{ST}} P') & \Gamma = \Gamma', x : \oplus\{l_i : S_i\}_{i \in I} \\
(\Gamma \vdash_{\text{ST}} x \triangleright \{l_i : P_i\}_{i \in I}) &\triangleq x \triangleright \{l_i : (\Gamma', x : S_i \vdash_{\text{ST}} P_i)\}_{i \in I} & \Gamma = \Gamma', x : \&\{l_i : S_i\}_{i \in I} \\
(\Gamma \vdash_{\text{ST}} (\nu \tilde{x}\tilde{y} : \tilde{S})(P \mid Q)) &\triangleq \{\Gamma_2\} \mid \mathcal{C}_{\tilde{z}, \tilde{S}}[(\Gamma_1, \tilde{x} : \tilde{S} \vdash_{\text{ST}} P)[\tilde{z}/\tilde{x}]] & \Gamma = \Gamma_1 \circ \Gamma_2 \wedge \Gamma_1, \tilde{x} : \tilde{S} \vdash_{\text{ST}} P \\
& \parallel_k \{\Gamma_1\} \mid \mathcal{C}_{\tilde{z}, \tilde{V}}[(\Gamma_2, \tilde{y} : \tilde{V} \vdash_{\text{ST}} Q)[\tilde{z}/\tilde{y}]] & \Gamma_2, \tilde{y} : \tilde{V} \vdash_{\text{ST}} Q \wedge \forall_i = \bar{S}_i
\end{aligned}$$

We illustrate the procedure in [9]. Notice that the rewriting procedure given in Definition 5.7 satisfies the compositionality criteria given in [11]. In particular, it is easy to see that the rewriting of a composition of terms is defined in terms of the rewriting of the constituent subterms. Indeed, e.g., $(\Gamma_1 \circ \Gamma_2 \vdash_{\text{ST}} (\nu xy : S)(P \mid Q))$ depends on a context including both $(\Gamma_1, x : S \vdash_{\text{ST}} P)$ and $(\Gamma_2, y : \bar{S} \vdash_{\text{ST}} Q)$.

We present two important results about our rewriting procedure. First, we show it is type preserving:

Theorem 5.8 (Rewriting is Type Preserving). Let $(\Gamma \vdash_{\text{ST}} P) \in \mathcal{K}_n$. Then, $(\Gamma \vdash_{\text{ST}} P) \vdash_{\text{CH}} \llbracket \Gamma \rrbracket_{\mathbb{C}}$.

Notice that the inverse of the previous theorem is trivial by following the definition of typed encoding. Theorem 5.8 is meaningful, for it says that the type interface of a process (i.e., the set of sessions implemented in it) is not modified by the rewriting procedure. That is, the procedure modifies the process structure by closely following the causality relations described by (session) types. Notice that causality relations present in processes, but not described at the level of types, may be removed.

The rewriting procedure also satisfies an operational correspondence result. Let us write $\Gamma \vdash_{\text{ST}} P_1, P_2$ whenever both $\Gamma \vdash_{\text{ST}} P_1$ and $\Gamma \vdash_{\text{ST}} P_2$ hold. We have the following auxiliary definition:

Definition 5.9. Let P, P' be such that $\Gamma \vdash_{\text{ST}} P, P'$. Then, we write $P \doteq P'$ if and only if $P = \mathbb{C}[Q]$ and $P' = \mathbb{C}[Q']$, for some context \mathbb{C} , and there is Γ' such that $\Gamma' \vdash_{\text{ST}} Q, Q'$.

Theorem 5.10 (Operational Correspondence). Let $P \in \mathcal{K}_n$ such that $\Gamma \vdash_{\text{ST}} P$ for some Γ . Then we have:

- I) If $P \rightarrow P'$ then there exist Q, Q' s.t. (i) $(\Gamma \vdash_{\text{ST}} P) \rightarrow^{\text{inx}} \rightarrow^* \equiv Q$; (ii) $Q \doteq Q'$; (iii) $(\Gamma \vdash_{\text{ST}} P') \rightarrow^{\text{inx}} Q'$.
- II) If $(\Gamma \vdash_{\text{ST}} P) \rightarrow^{\text{inx}} \rightarrow^* \equiv Q$ then there exists P' s.t. $P \rightarrow P'$ and $Q \doteq (\Gamma \vdash_{\text{ST}} P')$.

6 Concluding Remarks

We have presented a formal comparison of fundamentally distinct type systems for deadlock-free, session typed processes. To the best of our knowledge, ours is the first work to establish precise relationships of this kind. Indeed, prior comparisons between type systems for deadlock freedom are informal, given in terms of representative examples typable in one type system but not in some other.

An immediate difficulty in giving a unified account of different typed frameworks for deadlock freedom is the variety of process languages, type structures, and typing rules that define each framework. Indeed, our comparisons involve: the framework of session processes put forward by Vasconcelos [19]; the interpretation of linear logic propositions as session types by Caires [1]; the π -calculus with usage types defined by Kobayashi in [13]. Finding some common ground for comparing these three frameworks is not trivial—several translations/transformations were required in our developments to account for numerous syntactic differences. We made an effort to follow the exact definitions in each framework. Overall, we believe that we managed to concentrate on essential semantic features of two salient classes of deadlock-free session processes, noted \mathcal{L} and \mathcal{K} .

Our main contribution is identifying the *degree of sharing* as a subtle, important issue that underlies both session typing and deadlock freedom. We propose a simple characterization of the degree of sharing: in essence, it arises via an explicit premise for the typing rule for parallel composition in the type system in [13]. The degree of sharing is shown to effectively induce a strict hierarchy of deadlock-free session processes in \mathcal{K} , as resulting from the approach of [8]. We showed that the most elementary (and non trivial) member of this hierarchy precisely corresponds to \mathcal{L} —arguably the most canonical class of session typed processes known to date. Furthermore, by exhibiting an intuitive rewriting procedure of processes in \mathcal{K} into processes in \mathcal{L} , we demonstrated that the degree of sharing is a subtle criteria for distinguishing deadlock-free processes. As such, even if our technical developments are technically simple, in our view they substantially clarify our understanding of type systems for liveness properties (such as deadlock freedom) in the context of π -calculus processes.

As future work, we would like to obtain *semantic characterizations* of the degree of sharing, in the form of, e.g., preorders on typed processes that distinguish when one process “is more parallel” than another. We plan also to extend our formal relationships to cover typing disciplines with *infinite behavior*. We notice that the approach of [8] extends to recursive behavior [7] and that infinite (yet non divergent) behavior has been incorporated into logic-based session types [18]. Finally, we plan to explore whether the rewriting procedure given in § 5 could be adapted into a *deadlock resolution* procedure.

Acknowledgements. We are grateful to Luís Caires, Simon J. Gay, and the anonymous reviewers for their valuable comments and suggestions. This work was partially supported by the EU COST Action IC1201 (Behavioural Types for Reliable Large-Scale Software Systems). Dardha is supported by the UK EPSRC project EP/K034413/1 (From Data Types to Session Types: A Basis for Concurrency and Distribution). Pérez is also affiliated to NOVA Laboratory for Computer Science and Informatics, Universidade Nova de Lisboa, Portugal.

References

- [1] Luís Caires (2014): *Types and Logic, Concurrency and Non-Determinism*. In *Essays for the Luca Cardelli Fest - Microsoft Research Technical Report MSR-TR-2014-104*. Available at <http://research.microsoft.com/apps/pubs/default.aspx?id=226237>.
- [2] Luís Caires & Frank Pfenning (2010): *Session Types as Intuitionistic Linear Propositions*. In: *Proc. of CONCUR 2010, LNCS 6269*, Springer, pp. 222–236, doi:10.1007/978-3-642-15375-4_16.
- [3] Luís Caires, Frank Pfenning & Bernardo Toninho (2014): *Linear Logic Propositions as Session Types*. *MSCS*, doi:10.1017/S0960129514000218.
- [4] Marco Carbone, Ornella Dardha & Fabrizio Montesi (2014): *Progress as Compositional Lock-Freedom*. In: *COORDINATION, LNCS 8459*, Springer, pp. 49–64, doi:10.1007/978-3-662-43376-8_4.
- [5] Marco Carbone & Søren Debois (2010): *A Graphical Approach to Progress for Structured Communication in Web Services*. In: *Proc. of ICE 2010, Amsterdam, The Netherlands, 10th of June 2010., EPTCS 38*, pp. 13–27, doi:10.4204/EPTCS.38.4.
- [6] Gerardo Costa & Colin Stirling (1987): *Weak and Strong Fairness in CCS*. *Inf. Comput.* 73(3), pp. 207–244, doi:10.1016/0890-5401(87)90013-7.
- [7] Ornella Dardha (2014): *Recursive Session Types Revisited*. In: *Proceedings Third Workshop on Behavioural Types, BEAT 2014, Rome, Italy, 1st September 2014., EPTCS 162*, pp. 27–34, doi:10.4204/EPTCS.162.4.
- [8] Ornella Dardha, Elena Giachino & Davide Sangiorgi (2012): *Session types revisited*. In: *Proc. of PPDP'12, ACM*, pp. 139–150, doi:10.1145/2370776.2370794.
- [9] Ornella Dardha & Jorge A. Pérez (2015): *Full version of this paper*. Technical Report. Available at <http://www.jorgeaperez.net>.
- [10] Mariangiola Dezani-Ciancaglini, Ugo de'Liguoro & Nobuko Yoshida (2008): *On Progress for Structured Communications*. In: *Trustworthy Global Computing, LNCS 4912*, Springer, pp. 257–275, doi:10.1007/978-3-540-78663-4_18.
- [11] Daniele Gorla (2010): *Towards a unified approach to encodability and separation results for process calculi*. *Inf. Comput.* 208(9), pp. 1031–1053, doi:10.1016/j.ic.2010.05.002.
- [12] Kohei Honda, Vasco Thudichum Vasconcelos & Makoto Kubo (1998): *Language Primitives and Type Discipline for Structured Communication-Based Programming*. In: *Proc. of ESOP'98, LNCS 1381*, Springer, pp. 122–138, doi:10.1007/BFb0053567.
- [13] Naoki Kobayashi (2002): *A Type System for Lock-Free Processes*. *Inf. Comput.* 177(2), pp. 122–159, doi:10.1006/inco.2002.3171.
- [14] Naoki Kobayashi (2003): *Type Systems for Concurrent Programs*. In: *Formal Methods at the Crossroads, LNCS 2757*, Springer, pp. 439–453, doi:10.1007/978-3-540-40007-3_26.
- [15] Naoki Kobayashi (2006): *A New Type System for Deadlock-Free Processes*. In: *Proc. of CONCUR 2006, LNCS 4137*, Springer, pp. 233–247, doi:10.1007/11817949_16.
- [16] Luca Padovani (2013): *From Lock Freedom to Progress Using Session Types*. In: *Proceedings of PLACES 2013, Rome, Italy, 23rd March 2013., EPTCS 137*, pp. 3–19, doi:10.4204/EPTCS.137.2.
- [17] Benjamin C. Pierce (2002): *Types and programming languages*. MIT Press, MA, USA.
- [18] Bernardo Toninho, Luís Caires & Frank Pfenning (2014): *Corecursion and Non-divergence in Session-Typed Processes*. In: *Proc. of TGC 2014, LNCS 8902*, Springer, pp. 159–175, doi:10.1007/978-3-662-45917-1_11.
- [19] Vasco T. Vasconcelos (2012): *Fundamentals of session types*. *Inf. Comput.* 217, pp. 52–70, doi:10.1016/j.ic.2012.05.002.
- [20] Hugo Torres Vieira & Vasco Thudichum Vasconcelos (2013): *Typing Progress in Communication-Centred Systems*. In: *COORDINATION, LNCS 7890*, Springer, pp. 236–250, doi:10.1007/978-3-642-38493-6_17.
- [21] Philip Wadler (2012): *Propositions as sessions*. In: *Proc. of ICFP'12*, pp. 273–286, doi:10.1145/2364527.2364568.

On Compensation Primitives as Adaptable Processes

Jovana Dedeić

University of Novi Sad, Serbia

Jovanka Pantović

University of Novi Sad, Serbia

Jorge A. Pérez

University of Groningen, The Netherlands

We compare mechanisms for *compensation handling* and *dynamic update* in calculi for concurrency. These mechanisms are increasingly relevant in the specification of reliable communicating systems. Compensations and updates are intuitively similar: both specify how the behavior of a concurrent system changes at runtime in response to an exceptional event. However, calculi with compensations and updates are technically quite different. We investigate the *relative expressiveness* of these calculi: we develop encodings of core process languages with compensations into a calculus of *adaptable processes* developed in prior work. Our encodings shed light on the (intricate) semantics of compensation handling and its key constructs. They also enable the transference of existing verification and reasoning techniques for adaptable processes to core languages with compensation handling.

1 Introduction

Many software applications are based on *long-running transactions* (LRTs). Frequently found in service-oriented systems [8], LRTs are computing activities which extend in time and may involve distributed, loosely coupled resources. These features sharply distinguish LRTs from usual (database) transactions. One particularly delicate aspect of LRTs management is handling (partial) failures: mechanisms for detecting failures and bringing the LRT back to a consistent state need to be explicitly programmed. As designing and certifying the correctness of such mechanisms is error prone, the last decade has seen the emergence of specialized constructs, such as *exceptions* and *compensations*, which offer direct programming support. Our focus is in the latter: as their name suggests, compensation mechanisms are meant to compensate the fact that an LRT has failed or has been aborted. Upon reception of an abortion or failure signal, compensation mechanisms are expected to install and activate alternative behaviors for recovering system consistency. Such a compensation behavior may be different from the LRT's initial behavior.

A variety of calculi for concurrency with constructs for compensation handling has been proposed (see, e.g., [1, 5, 8, 14]). Building upon the tradition and approach of mobile process calculi such as the π -calculus [16], they capture different forms of error recovery and offer reasoning techniques (e.g., behavioral equivalences) on communicating processes with compensation constructs. The relative expressive power of such proposals has also been studied [4, 5, 12, 13]. On a related but different vein, a calculus of *adaptable processes* has been put forward as a process calculus approach to specify the dynamic evolution of interacting systems [2]. It is intended as a way of overcoming the limitations that process calculi have for describing patterns of dynamic evolution. In this calculus, process behaviors may be enclosed by nested, transparent *locations*; actions of dynamic update are targeted to particular locations. This model allows us to represent a wide range of evolvability patterns for concurrent processes. The theory of adaptable processes includes expressiveness, decidability, and verification results [2, 3], as well as the integration with structured communications governed by session types [9, 10].

Adaptable processes specify forms of dynamic reconfiguration which are triggered by exceptional events, not necessarily catastrophic. For instance, an external request for upgrading a working component is an exceptional event which is hard to predict and entails a modification of the system's behavior. Still, it is certainly not an error or a failure. Thus, adaptation intuitively appears to us as a general phenomenon

which includes the (negative) exceptional events dealt by compensations. That is, it should be possible to represent failures and compensation activities as particular instances of the behaviors expressible in [2].

In this paper, we make this intuitive observation precise by encoding calculi with compensations into adaptable processes. Our motivation is twofold. First, given the diversity of linguistic constructs for compensations, understanding how they can be implemented as adaptable processes could shed new light in their formal underpinnings. Since adaptable processes have a simple semantics (based on higher-order process communication [17]), the envisaged encodings could suggest alternative semantics for existing formalisms. Second, given that adaptable processes have been developed in several directions, encodings of calculi with compensations into adaptable processes could enable the transference of, e.g., decidability results or type systems, from adaptable processes to calculi with compensations.

As source languages in our study, we systematically consider the different classes of calculi with compensations developed in [12], a work that offers a unified presentation for many calculi proposed in the literature. In particular, we consider processes with *static* and *dynamic* compensations, each of them with *preserving*, *discarding*, and *aborting* semantics. (All these semantics are illustrated next.) As such, we offer six different encodings into adaptable processes, each one equipped with appropriate operational correspondence results. The encodings are rather involved; in particular, representing preserving, discarding, and aborting semantics by means of the transparent locations in [2] proved to be quite challenging. In our view, the intricate character of our representations into adaptable processes is directly related to the intricate semantics of each of the forms of calculi with compensations.

This paper is structured as follows. § 2 illustrates primitives for adaptable processes and compensation handling; § 3 formally presents the corresponding calculi. In § 4 we define and prove correct encodings of processes with static compensations into adaptable processes. We consider aborting, preserving, and discarding semantics. § 5 describes encodings of processes with dynamic compensations. § 6 collects some concluding remarks. Due to space restrictions, omitted proofs can be found online [7].

2 Adaptable and Compensable Processes, By Example

We give an intuitive account of the calculus of *adaptable processes* (introduced by Bravetti et al. [2]) and of the core calculus with primitives for *compensation handling* (as presented by Lanese et al. [12, 13]).

Adaptable Processes. The calculus of *adaptable processes* was introduced in [2] as a variant of Milner’s CCS [15] (without restriction and relabeling), extended with the following two constructs, aimed at representing the dynamic reconfiguration (or update) of active communicating processes:

1. A *located process*, denoted $l[P]$, represents a process P which resides in a location called l . Locations are *transparent*: the behavior of $l[P]$ is the same as the behavior of P . Locations can also be arbitrarily *nested*, which allows to organize process descriptions into meaningful hierarchical structures.
2. An *update prefix* $l\{(X).Q\}$ —where X is a process variable that occurs zero or more times in Q —denotes an adaptation mechanism for processes located at location l .

This way, in the calculus of adaptable process the possibility of updating a (located) process behavior is given the same status as communication prefixes. Intuitively, an update prefix for location l is able to interact with a located process at l , updating its current behavior. This is captured by the reduction rule

$$C_1[l[P]] \mid C_2[l\{(X).Q\}.R] \rightarrow C_1[Q\{P/X\}] \mid C_2[R]$$

where C_1 and C_2 denote *contexts* which may describe, e.g., nested locations and parallel components. Therefore, the adaptation mechanism (embodied by $l\{(X).Q\}$) moves to the place where $l[P]$ resides (C_1

above) and exercises a dynamic update there, as represented by substitution $Q\{P/X\}$. As such, adaptation is a form of *higher-order process communication* [17]. Observe that Q may not contain X , so the current behavior at l (i.e., P) may get erased as a result of the update. Notice also that this form of adaptation is *subjective*: located processes are influenced by (unknown) update prefixes in their environment.

Compensable Processes. Our core process language with compensations is based on the calculus in [13] (a variant of the language in [12]). The languages in [12, 13] are appealing because they uniformly capture several different proposals for calculi with compensation handling. These calculi were introduced as extensions of the π -calculus [16] with primitives for *static* and *dynamic recovery*. However, in order to focus on the essentials of compensation handling primitives, in this presentation we consider a variant of the languages in [12, 13] without name mobility. There are three salient constructs:

1. *Transaction scopes* (or simply *transactions*), denoted $t[P, Q]$, where t is a name and P, Q are processes;
2. *Protected blocks*, denoted $\langle Q \rangle$, for some process Q ;
3. *Compensation updates*, denoted $\text{inst}[\lambda X.Q].P$, where P, Q are processes and X is a process variable that occurs zero or more times in Q .

While transactions and protected blocks define static recovery mechanisms, compensation updates are used to define dynamic recovery. We now gradually introduce these constructs and their main features.

Basic Intuitions. A transaction $t[P, Q]$ consists of a *default activity* P with a *compensation activity* Q . Transactions can be nested, so process P in $t[P, Q]$ may contain other transactions. Transactions can be aborted: intuitively, process $t[P, Q]$ behaves as P until an *error notification* (abortion signal) arrives along name t . Error notifications are simply output messages which can originate inside or outside the transaction. To illustrate the simplest manifestation of compensations, we have the following transitions:

$$t[P, Q] \mid \bar{t}.R \xrightarrow{\tau} Q \mid R \qquad t[\bar{t}.P_1 \mid P_2, Q] \mid R \xrightarrow{\tau} Q \mid R$$

While the transition in the left shows how a transaction t can be aborted by an external signal, the transition in the right illustrate abortion due to an internal signal. In both cases, abortion leads to discarding the default behavior of the transition, and the compensation activity is executed instead (Q in both cases).

Protected Blocks. The transitions above illustrate the different sources of abortion signals that lead to compensation behaviors. One key element in calculi with compensations primitives are *protected blocks*: as their name suggests, these constructs protect a process from abortion signals. Similarly as locations, protected blocks are transparent: Q and $\langle Q \rangle$ have the same behavior, but $\langle Q \rangle$ cannot be affected by abortion signals. Protected blocks are meant to prevent abortions after a compensation:

$$t_2[P_2, Q_2] \mid \bar{t}_2 \xrightarrow{\tau} \langle Q_2 \rangle$$

That is, the compensation behavior Q_2 will be immune to external errors thanks to protected blocks. Consider now process $t_1[t_2[P_2, Q_2] \mid \bar{t}_2.R_1, Q_1]$, which includes a transaction named t_2 which is *nested* inside t_1 . Although in previous examples the default behavior has been erased following an abortion signal, the semantics of compensations actually may partially preserve such behavior. This is realized by *extraction functions*, denoted $\text{extr}(\cdot)$. For the previous process, we have the following transition:

$$t_1[t_2[P_2, Q_2] \mid \bar{t}_2.R_1 \mid R_2, Q_1] \xrightarrow{\tau} t_1[\langle Q_2 \rangle \mid \text{extr}(P_2) \mid R_1, Q_1]$$

In case transaction t_2 is aborted, its compensation behavior Q_2 will be preserved. Moreover, part of the behavior of P_2 will be preserved as well: this is expressed by process $\text{extr}(P_2)$, which consists of at least all protected blocks in P_2 ; it may also contain some other processes, related to transactions (see next).

We consider *discarding*, *preserving*, and *aborting* variants for $\text{extr}(\cdot)$; they define three different semantics for compensations. Noted $\text{extr}_D(\cdot)$, $\text{extr}_P(\cdot)$, and $\text{extr}_A(\cdot)$, respectively, these functions concern mostly protected blocks and transactions. Given a process P , we would have:

- $\text{extr}_D(P)$ keeps only protected blocks in P . Other processes (including transactions) are discarded.
- $\text{extr}_P(P)$ keeps protected blocks and transactions at the top-level in P . Other processes are discarded.
- $\text{extr}_A(P)$ keeps protected blocks and nested transactions in P , including their respective compensation activities. Other processes are discarded.

As an example, consider the process $P = t[t_1[P_1, Q_1] \mid t_2[\langle P_2 \rangle, Q_2] \mid R \mid \langle P_3 \rangle, Q_5]$. We then have:

$$\begin{array}{l} \text{Discarding semantics: } \bar{t} \mid P \xrightarrow{\tau}_D \langle P_3 \rangle \mid \langle Q_5 \rangle \\ \text{Preserving semantics: } \bar{t} \mid P \xrightarrow{\tau}_P \langle P_3 \rangle \mid \langle Q_5 \rangle \mid t_1[P_1, Q_1] \mid t_2[\langle P_2 \rangle, Q_2] \\ \text{Aborting semantics: } \bar{t} \mid P \xrightarrow{\tau}_A \langle P_3 \rangle \mid \langle Q_5 \rangle \mid \langle P_2 \rangle \mid \langle Q_1 \rangle \mid \langle Q_2 \rangle \end{array}$$

Thus, the three different semantics implement different levels of protection. The discarding semantics only concerns the compensation activity for transaction t and the protected block $\langle P_3 \rangle$. The preserving semantics protects also the nested transactions t_1 and t_2 ; a process such as R , without an enclosing protected block, is discarded. Finally, the aborting semantics preserves all protected blocks and compensation activities in the default activity for t , including those in nested transactions, such as $\langle P_2 \rangle$.

Dynamic Compensations. Up to here we have considered transactions with *static compensations*: while the default behavior may change due to transaction abortion, the compensable behavior remains unchanged. Given a transaction $t[P, Q]$, using *compensation updates* one may specify in P an update for the compensation behavior Q . This is achieved by the operator $\text{inst}[\lambda X.Q].P$, where $\lambda X.Q$ is a function which represents the compensation update. As a simple example, consider the following transition:

$$t[\text{inst}[\lambda X.R].P_1 \mid P_2, Q] \xrightarrow{\tau} t[P_1 \mid P_2, R\{Q/X\}]$$

This way, $\text{inst}[\lambda X.R].P$ produces a new compensation behavior $R\{Q/X\}$ after an internal transition. As variable X may not occur in R , this step may fully discard the previous compensation activity Q .

3 The Calculi

We introduce adaptable processes (§ 3.1) and compensable processes (§ 3.2). To focus on their essentials, both calculi are defined as extensions of CCS [15] (no name passing involved). In both cases, we assume a countable set of names N , ranged over by a, b, l, t, \dots . As a convention, we use names l, l', \dots to denote locations (in adaptable processes) and names t, t', \dots to denote transactions (in compensable processes).

3.1 Adaptable Processes

The syntax of the calculus of *adaptable processes* is defined by *prefixes* π, π', \dots and *processes* P, Q, \dots :

$$\pi ::= a \mid \bar{a} \mid l\{(X).Q\} \quad P ::= l[P] \mid \mathbf{0} \mid \pi.P \mid !P \mid P \mid Q \mid (\nu a)P \mid X$$

We consider input and output prefixes (noted a and \bar{a} , respectively) and the *update prefix* $l\{(X).Q\}$, where Q may contain zero or more occurrences of *process variable* X . The syntax of processes includes *located processes* (noted $l[P]$ and intuitively motivated above) as well as usual CCS constructs for inaction, prefix

$$\begin{array}{c}
\text{(R-I/O)} \quad E \left[C[\bar{a}.P] \mid D[a.Q] \right] \rightarrow E \left[C[P] \mid D[Q] \right] \quad \text{(R-UPD)} \quad E \left[C[l[P]] \mid D[l\{(X).Q\}.R] \right] \rightarrow E \left[C[Q\{P/X\}] \mid D[R] \right] \\
\text{(R-PAR)} \quad \frac{P \rightarrow P'}{P \mid Q \rightarrow P' \mid Q} \quad \text{(R-RES)} \quad \frac{P \rightarrow P'}{(va)P \rightarrow (va)P'} \quad \text{(R-STR)} \quad \frac{P \equiv P' \quad P' \rightarrow Q' \quad Q' \equiv Q}{P \rightarrow Q}
\end{array}$$

Figure 1: Reduction semantics for adaptable processes.

(sequentiality), replication, parallel composition, and restriction. We omit $\mathbf{0}$ whenever possible; we write, e.g., $l\{(X).P\}$ instead of $l\{(X).P\}.\mathbf{0}$. Name a is bound in $(va)P$ and process variable X is bound in $l\{(X).Q\}$; given a process P , its sets of free and bound names/variables—denoted $\text{fn}(P)$, $\text{bn}(P)$, $\text{fv}(P)$, and $\text{bv}(P)$ —are as expected. We rely on expected notions of α -conversion (noted \equiv_α) and process substitution: $P\{Q/X\}$ denotes the process obtained by (capture avoiding) substitution of Q for X in P .

The semantics of adaptable processes is given by a reduction semantics, denoted \rightarrow , and defined as the smallest relation on processes induced by the rules in Figure 1. \rightarrow^* denotes the reflexive and transitive closure of \rightarrow . Reduction relies on *structural congruence*, denoted \equiv , and *contexts*, denoted C, D, E . We define \equiv as the smallest congruence on processes that satisfies the axioms:

$$\begin{array}{lll}
P \mid Q \equiv Q \mid P & P \mid (Q \mid R) \equiv (P \mid Q) \mid R & P \mid \mathbf{0} \equiv P \\
P \equiv Q \text{ if } P \equiv_\alpha Q & (va)\mathbf{0} \equiv \mathbf{0} & (va)(vb)P \equiv (vb)(va)P \\
(va)P \mid Q \equiv (va)(P \mid Q) \text{ if } a \notin \text{fn}(Q) & (va)l[P] \equiv l[(va)P] & !P \equiv P \mid !P
\end{array}$$

The syntax of monadic contexts (processes with a single *hole*, denoted $[\bullet]$) is defined as:

$$C ::= [\bullet] \mid C \mid P \mid l[C]$$

We write $C[P]$ to denote the process resulting from filling in all occurrences of $[\bullet]$ in context C with process P . We comment on rules in Figure 1. Rule (R-I/O) formalizes synchronization between process $\bar{a}.P$ and process $a.Q$ (enclosed in contexts C and D , respectively). Rule (R-UPD) formalizes the dynamic update/evolvability of a location l . The result of the synchronization between a located process $l[P]$ and an update prefix $l\{(X).Q\}$ is the process $Q\{P/X\}$. This resulting process stays in the same context as process $l[P]$. Rules (R-PAR), (R-RES), and (R-STR) are standard and/or self-explanatory.

3.2 Compensable Processes

The calculus of *compensable processes* extends CCS with constructs for transactions, protected blocks, and compensation updates:

$$\pi ::= a \mid \bar{a} \quad P, Q ::= \mathbf{0} \mid \pi.P \mid !P \mid (va)P \mid P \mid Q \mid t[P, Q] \mid \langle Q \rangle \mid X \mid \text{inst}[\lambda X.R].P$$

Prefixes π include input and output actions. Processes for inaction ($\mathbf{0}$), sequentiality ($\pi.P$), replication ($!P$), restriction ($(va)P$), and parallel composition ($P \mid Q$) are standard. We omit $\mathbf{0}$ whenever possible. Protected blocks $\langle Q \rangle$, transactions $t[P, Q]$, and compensation updates $\text{inst}[\lambda X.R].P$ have been already motivated. Error notifications are simply output messages; they can be internal (coming from the default activity) or external (coming from outside of the transaction). Name a is bound in $(va)P$ and variable X is bound in $\text{inst}[\lambda X.R]$; given a process P , its sets of free and bound names/variables—denoted $\text{fn}(P)$, $\text{bn}(P)$, $\text{fv}(P)$, and $\text{bv}(P)$ —are as expected. α -conversion (noted \equiv_α) and substitution $P\{Q/X\}$

$$\begin{array}{lll}
\text{extr}_D(t[P, Q]) = \mathbf{0} & \text{extr}_P(t[P, Q]) = t[P, Q] & \text{extr}_A(t[P, Q]) = \text{extr}_A(P) \mid \langle Q \rangle \\
\text{extr}(\langle P \rangle) = \langle P \rangle & \text{extr}(P \mid Q) = \text{extr}(P) \mid \text{extr}(Q) & \text{extr}((va)P) = (va)\text{extr}(P) \\
\text{extr}(!P) = \mathbf{0} & \text{extr}(\text{inst}[\lambda X.Q].P) = \mathbf{0} & \text{extr}(\pi.P) = \mathbf{0}
\end{array}$$

Figure 2: Extraction functions.

are also as expected. We assume that protected blocks and transactions do not appear behind prefixes; this is key to ensure encoding correctness. We shall say that the sub-calculus without compensation updates $\text{inst}[\lambda X.R].P$ is the calculus with *static compensations*; the full calculus will be referred to as the calculus with *dynamic compensations*. The following definitions apply uniformly to both.

Following [12, 13], the semantics of compensable processes is given in terms of a Labeled Transition System (LTS). Ranged over α, α' , the set of labels includes a, \bar{a}, τ , and $\lambda X.Q$. As in CCS, a denotes an input action, \bar{a} denotes an output action, and τ denotes synchronization (internal action). Label $\lambda X.Q$ is associated to compensation updates. Formally, we have three different LTSs, corresponding to processes under discarding, preserving, and aborting semantics. Therefore, for each $\kappa \in \{D, P, A\}$, we will have an extraction function $\text{extr}_\kappa(\cdot)$ and a transition relation $\xrightarrow{\alpha}_\kappa$. The different extraction functions are defined in Fig. 2; the rules of the LTSs are given in Fig. 3. As a convention, whenever a notion coincides for the three semantics, we shall avoid decorations D, P, and A. This way, e.g., by writing $\text{extr}(\langle P \rangle) = \langle P \rangle$ we mean that the extraction function for protected blocks is the same for all three semantics.

We comment on the rules in Fig. 3. Axioms (L-OUT) and (L-IN) execute output and input prefixes, respectively. Rule (L-REP) deals with replication, while rule (L-PAR) allows one parallel component to progress independently. Rule (L-RES) is the standard rule for restriction: it states that a transition of process P determines a transition of process $(va)P$, where label α provides that the restriction name a does not occur inside α . Rule (L-COMM) defines communication on a . Rule (L-SCOPE-OUT) allows the default activity P of a transaction to progress, provided that the performed action is not a compensation update and that there is no pending compensation update to be executed. The latter is ensured by condition $\text{noComp}(P)$, defined in [7]: the condition is true if and only if process P does not have compensation update which waits for execution. This means that a compensation update has priority over other transitions; that is, if process P in transaction $t[P, Q]$ has a compensation update at top-level then it will be performed before any change of the current state. Rule (L-RECOVER-OUT) allows an external process to abort a transaction via an output action \bar{t} . The resulting process contains two parts: the first part is obtained from the default activity P of the transaction via the appropriate extraction function; the second part corresponds to compensation Q which will be executed inside a protected block. Similarly, rule (L-RECOVER-IN) handles abortion when the error notification comes from the default activity P of the transaction. Rule (L-BLOCK) essentially specifies that protected blocks are transparent units. Observe that the actual semantics of protected blocks is defined via the extraction functions $\text{extr}(\cdot)$. The final two rules are peculiar of processes with dynamic compensations: while rule (L-INST) performs a compensation update, rule (L-SCOPE-CLOSE) updates the compensation of a transaction.

We find it convenient to define structural congruence (\equiv) and contexts also for compensable processes. We define \equiv as the smallest congruence on processes that includes \equiv_α and satisfies the axioms:

$$\begin{array}{lll}
P \mid Q \equiv Q \mid P & P \mid (Q \mid R) \equiv (P \mid Q) \mid R & P \mid \mathbf{0} \equiv P \\
(va)(vb)P \equiv (vb)(va)P & (va)P \mid Q \equiv (va)(P \mid Q) \text{ if } a \notin \text{fn}(Q) & (va)\mathbf{0} \equiv \mathbf{0} \\
\langle \langle P \rangle \rangle \equiv \langle P \rangle & \langle (va)P \rangle \equiv (va)\langle P \rangle & \langle \mathbf{0} \rangle \equiv \mathbf{0} \\
t[(va)P, Q] \equiv (va)t[P, Q] \text{ if } t \neq a, a \notin \text{fn}(Q) & & (va)\bar{a} \equiv \mathbf{0}
\end{array}$$

$$\begin{array}{c}
\begin{array}{c}
\text{(L-OUT)} \quad \text{(L-IN)} \\
\bar{a}.P \xrightarrow{\bar{a}} P \quad a.P \xrightarrow{a} P
\end{array}
\quad
\begin{array}{c}
\text{(L-REP)} \\
\frac{P \xrightarrow{\alpha} P'}{!P \xrightarrow{\alpha} P' \mid !P}
\end{array}
\quad
\begin{array}{c}
\text{(L-PAR)} \\
\frac{P \xrightarrow{\alpha} P'}{P \mid Q \xrightarrow{\alpha} P' \mid Q}
\end{array}
\quad
\begin{array}{c}
\text{(L-RES)} \\
\frac{P \xrightarrow{\alpha} P' \quad \alpha \neq a, \bar{a}}{(va)P \xrightarrow{\alpha} (va)P'}
\end{array}
\quad
\begin{array}{c}
\text{(L-COMM)} \\
\frac{P \xrightarrow{a} P' \quad Q \xrightarrow{\bar{a}} Q'}{P \mid Q \xrightarrow{\tau} P' \mid Q'}
\end{array}
\\
\\
\begin{array}{c}
\text{(L-SCOPE-OUT)} \\
\frac{P \xrightarrow{\alpha} P' \quad \alpha \neq \lambda X.Q \quad \text{noComp}(P)}{t[P, Q] \xrightarrow{\alpha} t[P', Q]}
\end{array}
\quad
\begin{array}{c}
\text{(L-RECOVER-OUT)} \\
\frac{\text{noComp}(P)}{t[P, Q] \xrightarrow{t} \text{extr}(P) \mid \langle Q \rangle}
\end{array}
\quad
\begin{array}{c}
\text{(L-RECOVER-IN)} \\
\frac{P \xrightarrow{\bar{t}} P' \quad \text{noComp}(P)}{t[P, Q] \xrightarrow{\tau} \text{extr}(P') \mid \langle Q \rangle}
\end{array}
\\
\\
\begin{array}{c}
\text{(L-BLOCK)} \\
\frac{P \xrightarrow{\alpha} P'}{\langle P \rangle \xrightarrow{\alpha} \langle P' \rangle}
\end{array}
\quad
\begin{array}{c}
\text{(L-INST)} \\
\text{inst}[\lambda X.Q].P \xrightarrow{\lambda X.Q} P
\end{array}
\quad
\begin{array}{c}
\text{(L-SCOPE-CLOSE)} \\
\frac{P \xrightarrow{\lambda X.R} P'}{t[P, Q] \xrightarrow{\tau} t[P', R\{Q/X\}]}
\end{array}
\end{array}$$

Figure 3: LTS for compensable processes. Symmetric variants of (L-PAR) and (L-COMM) are omitted.

An n -adic context $C[\bullet_1, \dots, \bullet_n]$ is obtained from a process by replacing n occurrences of $\mathbf{0}$, that are neither compensations nor in continuation of prefixes, with indexed holes $[\bullet_1], \dots, [\bullet_n]$. This way, for instance, the syntax of monadic contexts is defined as:

$$C ::= [\bullet] \mid \langle C \rangle \mid t[C, P] \mid P \mid C \mid C \mid P \mid (va)C.$$

We write $C[P]$ to denote the process resulting from filling in all occurrences of $[\bullet]$ in context C with process P . The following proposition is central to our operational correspondence statements.

Proposition 3.1. *Let P be a compensable process. If $P \xrightarrow{\tau} P'$ then one of the following holds:*

- $P \equiv E[C[\bar{a}.P_1] \mid D[a.P_2]]$ and $P' \equiv E[C[P_1] \mid D[P_2]]$,
- $P \equiv E[C[t[P_1, Q]] \mid D[\bar{t}.R]]$ and $P' \equiv E[C[\text{extr}(P_1) \mid \langle Q \rangle] \mid D[R]]$,
- $P \equiv C[t[D[\bar{t}.P_1], Q]]$ and $P' \equiv C[\text{extr}(D[P_1]) \mid \langle Q \rangle]$,
- $P \equiv E[t[C'[\text{inst}[\lambda X.R].P], Q]]$ and $P' \equiv E[t[C'[P], R\{Q/X\}]]$,

for some contexts C, C', D, E , processes P_1, P_2, Q, R , and names a, t .

4 Encoding Static Compensation Processes

Here we present encodings of processes with static compensations into adaptable processes. We consider discarding, preserving and aborting semantics. We adopt the following abbreviations for update prefixes:

- $t\{\dagger\}$ for the update prefix $t\{(Y).\mathbf{0}\}$ which “kills” location t , together with the process located at t ;
- $t\{P\}$ for the update prefix $t\{(Y).P\}$ (with $Y \notin \text{fv}(P)$) that replaces the current behavior at t with P ;
- $t\{\text{id}\}$ for the update prefix $t\{(X).X\}$ which deletes the location name t ;
- $t\{(X_1, X_2, \dots, X_n).R\}$ for the sequential composition of updates $t\{(X_1).t\{(X_2). \dots .t\{(X_n).R\}\}\}$.

Basic Intuitions. We describe some commonalities in the encodings we are about to present. Unsurprisingly, the main challenge to encodability is in representing transactions $t[P, Q]$ and protected blocks $\langle R \rangle$ as adaptable processes. Our strategy consists in representing P and Q independently, using located

processes. Since locations are transparent units of behavior, this suffices for encoding P . However, the encoding of Q cannot freely execute unless an abortion signal (an output action) is received. Very approximately, our encodings of protected blocks and transactions have the following structure:

$$\llbracket \langle R \rangle \rrbracket_{t,\rho} = p_{t,\rho} \llbracket \llbracket R \rrbracket_{\varepsilon} \rrbracket \quad (1)$$

$$\llbracket t[P, Q] \rrbracket_{\rho} = \underbrace{t \llbracket \llbracket P \rrbracket_{t,\rho} \rrbracket}_{(a)} \mid \underbrace{l_t.\pi_1 \cdots \pi_k.p_t \llbracket \llbracket Q \rrbracket_{t,\rho} \rrbracket}_{(b)} \mid \underbrace{t.\bar{l}_t.K}_{(c)} \quad (2)$$

In our encodings we use *paths*, finite sequences of names, denoted t_1, t_2, \dots, t_n . The empty path is denoted ε . Ranged over ρ , paths capture the hierarchical structure of nested transactions. Using paths, for each protected block, we maintain an association with the name of its enclosing transaction. As such, the encoding of a protected block associated to transaction t will be enclosed in a location p_t (see (1) above). There could be more than one occurrence of such locations, as the transaction's body may contain several protected blocks. The encoding of transactions, given in (2), consists of three parallel components:

- Component (a) is a location which contains the encoding of the default activity of the transaction; we retain the name of the transaction in the source process.
- Component (b) represents the compensation activity of the transaction. It is given as a located process at p_t , and is protected by a number of prefixes π_1, \dots, π_k including an input prefix l_t .
- Component (c) handles abortion signals. After synchronizing with an output on t , it synchronizes with the input on l_t in component (b). This releases a process K which “collects” all protected blocks in the encoding of P (which occur inside locations named p_t) but also the encoding of the compensation activity Q . This collection process may involve synchronizations with π_1, \dots, π_k in (b). Once all protected blocks have been collected, location t is destroyed.

This (very approximate) strategy is used in all of our encodings, with variations motivated by discarding, preserving, and aborting semantics. Knowing the number of protected blocks to be collected is crucial in this scheme. To this end, appropriate counting functions on the default activity P are defined.

The following remark defines some basic conditions on “reserved names” used in our encodings:

Remark 4.1. *Let t be a name, then we know that there are names l_t, k_t, p_t and m_t which are associated with the name t . Also, if $t_1 \neq t_2$ then $l_{t_1} \neq l_{t_2}, k_{t_1} \neq k_{t_2}, p_{t_1} \neq p_{t_2}$ and $m_{t_1} \neq m_{t_2}$.*

4.1 Discarding Semantics

Before presenting the encoding, we introduce some auxiliary functions. First, we introduce a function that counts the number of protected blocks in a process.

Definition 4.2 (Number of protected blocks). *Let P be a compensable process. The number of protected blocks in P , denoted by $\text{npb}_D(P)$, is defined as follows:*

$$\text{npb}_D(P) = \begin{cases} 1 & \text{if } P = \langle P_1 \rangle \\ \text{npb}_D(P_1) + \text{npb}_D(P_2) & \text{if } P = P_1 \mid P_2 \\ \text{npb}_D(P_1) & \text{if } P = (\nu a)P_1 \\ 0 & \text{otherwise.} \end{cases}$$

We shall define an encoding $D\llbracket \cdot \rrbracket_{\rho}$ of compensable processes into adaptable processes, where ρ is a path (a sequence of location names). The encoding of transactions requires an auxiliary encoding, denoted $D\| \cdot \|_{\rho}^n$, loosely related to component (b) in (2). In case of an abortion signal \bar{t} , $D\| \cdot \|_{\rho}^n$ defines a process that collects the encodings of the n protected blocks included in the default activity (which is to be found at ρ) as well as the encoding of the compensation activity. We define $D\| \cdot \|_{\rho}^n$ by induction on n :

Definition 4.3 (Auxiliary Encoding). *Let Q be a compensable process and let $\rho_0 = t, \rho$ be a path. Also, let $n \geq 0$. The process $\mathbb{D}\|Q\|_{\rho_0}^n$ is defined as follows:*

$$\begin{aligned} \mathbb{D}\|Q\|_{t,\rho}^0 &= l_t.\overline{m}_t.p_\rho [\mathbb{D}\|Q\|_\varepsilon] \mid m_t.\overline{k}_t.t\{\dagger\} \\ \mathbb{D}\|Q\|_{t,\rho}^n &= l_t.p_{t,\rho} \left\{ (X_1, \dots, X_n).z\{p_\rho[X_1] \mid \dots \mid p_\rho[X_n] \mid \overline{m}_t.p_\rho [\mathbb{D}\|Q\|_\varepsilon]\} \right\}.(z[\mathbf{0}] \mid m_t.\overline{k}_t.t\{\dagger\}) \quad [n > 0] \end{aligned}$$

(The definition of $\mathbb{D}\llbracket \cdot \rrbracket_\rho$ is given next.) Consider the encoding of $t[P, Q]$: if P contains n top-level protected blocks, then process $\mathbb{D}\llbracket t[P, Q] \rrbracket_\rho$ will include n successive update prefixes that will look for n protected blocks at location $p_{t,\rho}$ (the path points that they were enclosed with t) and move them to their parent location p_ρ . As these n dynamic updates leave these located processes at location t , an update on z is introduced to take them out of t once the n updates are executed.

We are now ready to introduce the encoding $\mathbb{D}\llbracket \cdot \rrbracket_\rho$. Recall that we adhere to Remark 4.1:

Definition 4.4 (Encoding Discarding Semantics). *Let P be a compensable process and let ρ be a path. The encoding $\mathbb{D}\llbracket \cdot \rrbracket_\rho$ of compensable processes into adaptable processes is defined as follows:*

$$\begin{aligned} \mathbb{D}\llbracket \langle P \rangle \rrbracket_\rho &= p_\rho [\mathbb{D}\llbracket P \rrbracket_\varepsilon] & \mathbb{D}\llbracket t[P, Q] \rrbracket_\rho &= t \left[\mathbb{D}\llbracket P \rrbracket_{t,\rho} \right] \mid \mathbb{D}\|Q\|_{t,\rho}^{\text{npb}_D(P)} \mid t.\overline{l}_t.k_t.\mathbf{0} & \mathbb{D}\llbracket \mathbf{0} \rrbracket_\rho &= \mathbf{0} \\ \mathbb{D}\llbracket P_1 \mid P_2 \rrbracket_\rho &= \mathbb{D}\llbracket P_1 \rrbracket_\rho \mid \mathbb{D}\llbracket P_2 \rrbracket_\rho & \mathbb{D}\llbracket \pi.P \rrbracket_\rho &= \pi.\mathbb{D}\llbracket P \rrbracket_\rho & \mathbb{D}\llbracket !P \rrbracket_\rho &= !\mathbb{D}\llbracket P \rrbracket_\rho & \mathbb{D}\llbracket (va)P \rrbracket_\rho &= (va)\mathbb{D}\llbracket P \rrbracket_\rho \end{aligned}$$

Key cases are encodings of protected blocks and transactions, as motivated earlier. Each protected block is associated with a location p indexed with the path to the protected block. A transaction is encoded as the composition of three processes. The leftmost component encodes the default activity P preserving the nested structure. In case of an abortion signal on t , the rightmost component will execute the middle component by sending message \overline{l}_t . As already explained, this second component will find all the top-level encodings of protected blocks of P , moving them to locations p_ρ together with the encoding of compensation activity Q . We may formalize these observations using the following lemma:

Lemma 4.5. *Let $t[P, Q]$ be a transaction with default activity P and compensation Q . Then we have:*

$$t \left[\mathbb{D}\llbracket P \rrbracket_{t,\rho} \right] \mid \mathbb{D}\|Q\|_{t,\rho}^{\text{npb}_D(P)} \mid \overline{l}_t.k_t \rightarrow^* \mathbb{D}\llbracket \text{extr}_D(P) \rrbracket_\rho \mid \mathbb{D}\llbracket \langle Q \rangle \rrbracket_\rho.$$

The following statement attests the operational correspondence for our encoding:

Theorem 4.6. *Let P be a compensable process and let ρ be an arbitrary path.*

- a) *If $P \xrightarrow{\tau}_D P'$ then $\mathbb{D}\llbracket P \rrbracket_\rho \rightarrow^* \mathbb{D}\llbracket P' \rrbracket_\rho$.*
- b) *If $\mathbb{D}\llbracket P \rrbracket_\rho \rightarrow Q$ then there is P' such that $P \xrightarrow{\tau}_D P'$ and $Q \rightarrow^* \mathbb{D}\llbracket P' \rrbracket_\rho$.*

We illustrate our encoding by means of an example:

Example 4.7. *Let $P_0 = t[R \mid \langle P \rangle, Q] \mid \overline{l}_t$ be a compensable process with $\text{npb}_D(R) = 0$. Then $P_0 \xrightarrow{\tau}_D \langle P \rangle \mid \langle Q \rangle$. By expanding Def. 4.4, we obtain (recall that we omit $\mathbf{0}$ whenever possible):*

$$\begin{aligned} \mathbb{D}\llbracket P_0 \rrbracket_\varepsilon &= t \left[\mathbb{D}\llbracket R \mid \langle P \rangle \rrbracket_{t,\varepsilon} \right] \mid \mathbb{D}\|Q\|_{t,\varepsilon}^1 \mid t.\overline{l}_t.k_t \mid \overline{l}_t \\ &= t \left[\mathbb{D}\llbracket R \rrbracket_{t,\varepsilon} \mid p_{t,\varepsilon} [\mathbb{D}\llbracket P \rrbracket_\varepsilon] \right] \mid l_t.p_{t,\varepsilon} \left\{ (X).z\{p_\varepsilon[X] \mid \overline{m}_t.p_\varepsilon [\mathbb{D}\|Q\|_\varepsilon]\} \right\}.(z[\mathbf{0}] \mid m_t.\overline{k}_t.t\{\dagger\}) \mid t.\overline{l}_t.k_t \mid \overline{l}_t \\ &\rightarrow^* t \left[\mathbb{D}\llbracket R \rrbracket_{t,\varepsilon} \mid z\{p_\varepsilon [\mathbb{D}\llbracket P \rrbracket_\varepsilon] \mid \overline{m}_t.p_\varepsilon [\mathbb{D}\|Q\|_\varepsilon]\} \right] \mid z[\mathbf{0}] \mid m_t.\overline{k}_t.t\{\dagger\} \mid k_t \rightarrow^* p_\varepsilon [\mathbb{D}\llbracket P \rrbracket_\varepsilon] \mid p_\varepsilon [\mathbb{D}\|Q\|_\varepsilon] \\ &= \mathbb{D}\llbracket \langle P \rangle \mid \langle Q \rangle \rrbracket_\varepsilon \end{aligned}$$

4.2 Preserving Semantics

The encoding of compensable processes with preserving semantics is as the previous encoding. In this case, since the extraction function keeps both protected blocks and top-level transactions (cf. Fig. 2), our auxiliary encoding, denoted $P \parallel \cdot \parallel_{\rho}^{n,m}$, has two parameters: n denotes protected blocks and m denotes top-level transactions. We count protected blocks using Def. 4.2; to count transactions we use the following:

Definition 4.8 (Number of transactions). *Let P be a compensable process. The number of transactions which occur in P , denoted $\text{nts}(P)$, is defined as follows:*

$$\text{nts}(P) = \begin{cases} \text{nts}(P_1) + 1 & \text{if } P = t[P_1, Q_1] \\ \text{nts}(P_1) + \text{nts}(P_2) & \text{if } P = P_1 \mid P_2 \\ \text{nts}(P_1) & \text{if } P = (\nu a)P_1 \\ 0 & \text{otherwise.} \end{cases}$$

The encoding of the transaction body P with location t that is nested in location β_{ρ} . Before giving the encoding $P \llbracket \cdot \rrbracket_{\rho}$, we define the auxiliary encoding $P \parallel \cdot \parallel_{\rho}^{n,m}$, where ρ is a path, n is the number of protected blocks, and m is the number of transactions in the default activity.

Definition 4.9 (Auxiliary Encoding). *Let Q be a compensable process and let $\rho_0 = t, \rho$ be a path. Also, let $n, m \geq 0$. The process $P \parallel Q \parallel_{\rho_0}^{n,m}$ is defined as follows:*

$$\begin{aligned} P \parallel Q \parallel_{t, \rho}^{0,0} &= l_t \cdot \bar{m}_t \cdot a \cdot p_{\rho} [P \llbracket Q \rrbracket_{\varepsilon}] \mid m_t \cdot \bar{k}_t \cdot t \{ \dagger \} \\ P \parallel Q \parallel_{t, \rho}^{1,0} &= l_t \cdot p_{t, \rho} \left\{ (X_1) \cdot z \{ a \cdot p_{\rho} [X_1] \mid \bar{m}_t \cdot p_{\rho} [P \llbracket Q \rrbracket_{\varepsilon}] \} \right\} \cdot (z[\mathbf{0}] \mid m_t \cdot \bar{k}_t \cdot t \{ \dagger \}) \\ P \parallel Q \parallel_{t, \rho}^{0,1} &= l_t \cdot \beta_{t, \rho} \left\{ (Y_1) \cdot z \{ a \cdot \beta_{\rho} [Y_1] \mid \bar{m}_t \cdot p_{\rho} [P \llbracket Q \rrbracket_{\varepsilon}] \} \right\} \cdot (z[\mathbf{0}] \mid m_t \cdot \bar{k}_t \cdot t \{ \dagger \}) \\ P \parallel Q \parallel_{t, \rho}^{n,m} &= l_t \cdot p_{t, \rho} \left\{ (X_1, \dots, X_n) \cdot \beta_{t, \rho} \left\{ (Y_1, \dots, Y_m) \cdot z \{ p_{\rho} [X_1] \mid \dots \mid p_{\rho} [X_n] \right. \right. \\ &\quad \left. \left. \mid a \cdot (\beta_{\rho} [Y_1] \mid \dots \mid \beta_{\rho} [Y_m]) \mid \bar{m}_t \cdot p_{\rho} [P \llbracket Q \rrbracket_{\varepsilon}] \} \right\} \right\} \cdot (z[\mathbf{0}] \mid m_t \cdot \bar{k}_t \cdot t \{ \dagger \}) \quad [n, m > 0] \end{aligned}$$

We may now define the encoding $P \llbracket \cdot \rrbracket_{\rho}$:

Definition 4.10 (Encoding Preserving). *Let P be a compensable process and let ρ be a path. The encoding $P \llbracket \cdot \rrbracket_{\rho}$ of compensable processes into adaptable processes is defined as*

$$P \llbracket \langle P \rangle \rrbracket_{\rho} = p_{\rho} [P \llbracket P \rrbracket_{\varepsilon}] \quad P \llbracket t[P, Q] \rrbracket_{\rho} = \beta_{\rho} \left[t [P \llbracket P \rrbracket_{t, \rho}] \mid P \parallel Q \parallel_{t, \rho}^{\text{npb}_P(P), \text{nts}(P)} \mid t \cdot \bar{l}_t \cdot k_t \cdot \bar{j} \mid j \cdot \beta_{\rho} \{ \text{id} \} \cdot \bar{a} \right]$$

and as a homomorphism for the other operators.

The following lemma formalizes the execution of the encoding:

Lemma 4.11. *Let $t[P, Q]$ be a transaction with default activity P and compensation Q . Then we have:*

$$\beta_{\rho} \left[t [P \llbracket P \rrbracket_{t, \rho}] \mid P \parallel Q \parallel_{t, \rho}^{\text{npb}_P(P), \text{nts}(P)} \mid \bar{l}_t \cdot k_t \cdot \bar{j} \mid j \cdot \beta_{\rho} \{ \text{id} \} \cdot \bar{a} \right] \rightarrow^* P \llbracket \text{extr}_P(P) \rrbracket_{\rho} \mid P \llbracket \langle Q \rangle \rrbracket_{\rho}.$$

We then have the following statement of operational correspondence:

Theorem 4.12. *Let P be a compensable process and let ρ an arbitrary path.*

- a) *If $P \xrightarrow{\tau}_P P'$ then $P \llbracket P \rrbracket_{\rho} \rightarrow^* P \llbracket P' \rrbracket_{\rho}$.*
- b) *If $P \llbracket P \rrbracket_{\rho} \rightarrow Q$ then there is P' such that $P \xrightarrow{\tau}_P P'$ and $Q \rightarrow^* P \llbracket P' \rrbracket_{\rho}$.*

Example 4.13. Let P_0 be a compensable process in Example 4.7 with $R = t_1[P_1, Q_1]$ and $\text{npb}_p(P_1) = \text{nts}(P_1) = 0$. In the preserving semantics we have: $P_0 \xrightarrow{\tau}_p t_1[P_1, Q_1] \mid \langle P \rangle \mid \langle Q \rangle$. By expanding Def. 4.10, we obtain:

$$\begin{aligned} \mathbb{P}[\![P_0]\!]_\varepsilon &= \beta_\varepsilon \left[t \left[\mathbb{P}[\![t_1[P_1, Q_1]]\!]_{t,\varepsilon} \mid \mathbb{P}[\![Q]\!]_{t,\varepsilon}^{1,1} \mid t.\bar{l}_t.k_t.\bar{j} \right] \mid j.\beta_\varepsilon\{\text{id}\}.\bar{a} \mid \bar{t} \right] \\ &= \beta_\varepsilon \left[t \left[\beta_{t,\varepsilon}[M] \mid j.\beta_{t,\varepsilon}\{\text{id}\}.\bar{a} \mid p_{t,\varepsilon}[\mathbb{P}[\![P]\!]_\varepsilon] \mid l_t.p_{t,\varepsilon} \left\{ (X_1).\beta_{t,\varepsilon} \left\{ (Y_1).z\{p_\varepsilon[X_1] \mid a.\beta_\varepsilon[Y_1] \right. \right. \right. \right. \\ &\quad \left. \left. \left. \mid \bar{m}_t.p_\varepsilon[\mathbb{P}[\![Q]\!]_\varepsilon] \right\} \right\} \cdot (z[\mathbf{0}] \mid m_t.\bar{k}_t.t\{\dagger\}) \mid t.\bar{l}_t.k_t.\bar{j} \right] \mid j.\beta_\varepsilon\{\text{id}\}.\bar{a} \mid \bar{t} \right] \\ &\rightarrow^* \beta_\varepsilon \left[t \left[z\{p_\varepsilon[\mathbb{P}[\![P]\!]_\varepsilon] \mid a.\beta_\varepsilon[M] \mid \bar{m}_t.p_\varepsilon[\mathbb{P}[\![Q]\!]_\varepsilon] \} \mid j.\beta_{t,\varepsilon}\{\text{id}\}.\bar{a} \mid z[\mathbf{0}] \mid m_t.\bar{k}_t.t\{\dagger\} \mid k_t.\bar{j} \right] \right. \\ &\quad \left. \mid j.\beta_\varepsilon\{\text{id}\}.\bar{a} \right] \\ &\rightarrow^* \beta_\varepsilon \left[t \left[\mathbf{0} \mid j.\beta_{t,\varepsilon}\{\text{id}\}.\bar{a} \mid p_\varepsilon[\mathbb{P}[\![P]\!]_\varepsilon] \mid a.\beta_\varepsilon[M] \mid p_\varepsilon[\mathbb{P}[\![Q]\!]_\varepsilon] \mid t\{\dagger\} \mid \bar{j} \right] \mid j.\beta_\varepsilon\{\text{id}\}.\bar{a} \right] \\ &\rightarrow^* \beta_\varepsilon [M] \mid p_\varepsilon[\mathbb{P}[\![P]\!]_\varepsilon] \mid p_\varepsilon[\mathbb{P}[\![Q]\!]_\varepsilon] \end{aligned}$$

where $M = t_1[\mathbb{P}[\![P_1]\!]_{t_1,t,\varepsilon}] \mid \mathbb{P}[\![Q_1]\!]_{t_1,t,\varepsilon}^{0,0} \mid t_1.\bar{l}_{t_1}.k_{t_1}.\bar{j}$.

4.3 Aborting Semantics

We now discuss the encoding of compensable processes with abortion semantics. While preserving the structure of the two encodings already presented, in this case the extraction function (cf. Fig. 2) add some complications. We need to modify the function that counts the number of protected blocks in a process; also, collecting encodings of (nested) protected blocks requires so-called *activation processes* which capture the hierarchical structure of nested transactions (cf. Def. 4.16).

Definition 4.14 (Number of protected blocks). *Let P be a compensable process. The number of protected blocks in P , denoted by $\text{npb}_A(P)$, is defined as follows:*

$$\text{npb}_A(P) = \begin{cases} 1 & \text{if } P = \langle P_1 \rangle \\ \text{npb}_A(P_1) + 1 & \text{if } P = t[P_1, Q_1] \\ \text{npb}_A(P_1) + \text{npb}_A(P_2) & \text{if } P = P_1 \mid P_2 \\ \text{npb}_A(P_1) & \text{if } P = (\nu a)P_1 \\ 0 & \text{otherwise.} \end{cases}$$

We now define the auxiliary encoding, denoted $A\|Q\|_\rho^n$. This process, as explained above, collects all encoded protected blocks of a process, in a case that an error notification is activated.

Definition 4.15 (Auxiliary Encoding). *Let Q be a compensable process and let $\rho_0 = t, \rho$ be a path. Also, let $n \geq 0$. The process $A\|Q\|_{\rho_0}^n$ is defined as follows:*

$$\begin{aligned} A\|Q\|_{t,\rho}^0 &= l_t.\bar{m}_t.p_\rho [A\|Q\|_\varepsilon] \mid m_t.\bar{k}_t.t\{\dagger\}.\Gamma_{t,\rho} \\ A\|Q\|_{t,\rho}^n &= l_t.p_{t,\rho} \left\{ (X_1, \dots, X_n).z\{p_\rho[X_1] \mid \dots \mid p_\rho[X_n] \mid \bar{m}_t.p_\rho [A\|Q\|_\varepsilon] \} \cdot (z[\mathbf{0}] \mid m_t.\bar{k}_t.t\{\dagger\}.\Gamma_{t,\rho}) \mid n > 0 \right\} \end{aligned}$$

where $\Gamma_{t,\rho} = \gamma_{t_1}\{(Z_1).\gamma_{t_1}[(\nu l_{t_1})(\nu k_{t_1})(Z_1 \mid l_{t_1}.\bar{k}_{t_1})]\} \cdot \dots \cdot \gamma_{t_n}\{(Z_n).\gamma_{t_n}[(\nu l_{t_n})(\nu k_{t_n})(Z_n \mid l_{t_n}.\bar{k}_{t_n})]\}.\gamma_{t_n}\{\dagger\}$.

To appropriately collect nested protected blocks, we define a so-called *activation process* that captures the hierarchical structure of nested transactions.

Definition 4.16 (Activation Process). *Let $St(P)$ denote the containment structure of compensable process P , i.e., the labeled tree (with root t) in which nodes are labeled with transaction names and sub-trees capture transaction nesting. The activation process for P , denoted $\mathcal{T}_t(P)$, is the sequential process obtained by a post-order search in $St(P)$ in which the visit to a node labeled c_i adds prefixes $\bar{l}_{c_i}.k_{c_i}$.*

This way, e.g., given $P = a[c[P_1, Q_2] \mid P_2, Q_1] \mid b[P_3 \mid d[P_4, Q_4] \mid e[P_5, Q_5], Q_3]$ we will have the activation process $\mathcal{T}_t(P) = \bar{l}_c.k_c.\bar{l}_a.k_a.\bar{l}_d.k_d.\bar{l}_e.k_e.\bar{l}_b.k_b.\bar{l}_t.k_t$.

Now we have all necessary definitions for introducing of the encoding $A[\cdot]_\rho$ of compensable processes into adaptable processes with respect to aborting semantics. Notice the use of activation processes in the encoding of transactions:

Definition 4.17 (Encoding Aborting). *Let P be a compensable process and let ρ be a path. The encoding $A[\cdot]_\rho$ of compensable processes into adaptable processes is defined as*

$$A[\langle P \rangle]_\rho = p_\rho[A[P]_\varepsilon] \quad A[t[P, Q]]_\rho = t[A[P]_{t,\rho}] \mid A[Q]_{t,\rho}^{\text{npb}_A(P)} \mid \gamma_t[t, \mathcal{T}_t(P)]$$

and as a homomorphism for the other operators.

With respect to previous encodings, the encoding for aborting semantics differs in the rightmost process of the encoding. In this case, the activation process $\mathcal{T}_t(P)$ searches the subtree of the transaction body to activate the middle components of all nested transactions inside t .

The following correctness statements follow the same ideas as in the two previous encodings. In the sequel, we write \approx to denote a (weak) behavioral equivalence that abstracts from internal transitions (due to the synchronizations added by the activation process).

Lemma 4.18. *Let $t[P, Q]$ be a transaction with default activity P and compensation Q . Then we have:*

$$t[A[P]_{t,\rho}] \mid A[Q]_{t,\rho}^{\text{npb}_A(P)} \mid \gamma_t[\mathcal{T}_t(P)] \rightarrow^* A[\text{extr}_A(P)]_\rho \mid A[\langle Q \rangle]_\rho \mid \Gamma_{t,\rho} \mid \gamma_t[\mathbf{0}].$$

Theorem 4.19. *Let P be a compensable process and let ρ be an arbitrary path.*

- a) *If $P \xrightarrow{\tau}_A P'$ then $A[P]_\rho \rightarrow^* A[P']_\rho$.*
- b) *If $A[P]_\rho \rightarrow Q$ then there is P' such that $P \xrightarrow{\tau}_A P'$ and $Q \rightarrow^* Q'$ and $Q' \approx A[P']_\rho$.*

5 Encoding Dynamic Compensation Processes

We discuss how to extend the previous encodings to account for compensation updates $\text{inst}[\lambda Y.R].P$. Due to space constraints, we only describe required extensions to previously given definitions/statements.

Discarding Semantics. We first have the following extension to Def. 4.2:

Definition 5.1 (Number of protected blocks). *Let P be a compensable process such that $P = \text{inst}[\lambda Y.R].P_1$. The number of protected blocks in P , denoted by $\text{npb}(P)$, is equal to $\text{npb}(P_1)$.*

The definition of the auxiliary encoding, given in Def. 4.3, is extended as follows:

Definition 5.2 (Auxiliary encoding). *Let Q be a compensable process and let $\rho_0 = t, \rho$ be a path. Also, let $n \geq 0$. The process $\mathbb{D}^n Q \mid \rho_0$ is defined inductively on n as follows:*

$$\begin{aligned} \mathbb{D}^0 Q \mid \rho &= l_t.\bar{m}_t.p_\rho[u[\bar{f}.\bar{g}]] \mid m_t.\bar{k}_t.t\{\dagger\} \mid v[u\{(Z).(Z \mid v_1[\mathbb{D}[Q]_\varepsilon] \mid f.v_1\{\text{id}\}.v\{\text{id}\}.g)\}] \\ \mathbb{D}^n Q \mid \rho &= l_t.p_{t,\rho} \left\{ (X_1, \dots, X_n).z\{p_\rho[X_1] \mid \dots \mid p_\rho[X_n] \mid \bar{m}_t.p_\rho[u[\bar{f}.\bar{g}]]\} \right\} \\ &\quad .(z[\mathbf{0}] \mid m_t.\bar{k}_t.t\{\dagger\}) \mid v[u\{(Z).(Z \mid v_1[\mathbb{D}[Q]_\varepsilon] \mid f.v_1\{\text{id}\}.v\{\text{id}\}.g)\}] \quad [n > 0] \end{aligned}$$

Based on the above modifications, the encoding of processes with dynamic compensations is obtained by extending Def. 4.4 with the following:

$$\begin{aligned} D[[Y]]_\rho &= Y \\ D[[\text{inst}[\lambda Y.R].P]]_{t,\rho} &= u[\mathbf{0}] \mid v_1\{(Y).\bar{g}.v\{(X).X \mid v[u\{(Z).(Z \mid \\ &\quad v_1[D[[R]]_\rho] \mid f.v_1\{\text{id}\}.v\{\text{id}\}.g\}\}\}\} \mid D[[P]]_{t,\rho}\}.\bar{f}.(v[\mathbf{0}] \mid v_1[\mathbf{0}]) \end{aligned}$$

We then have the following property:

Lemma 5.3. *Let $t[P, Q]$ be a transaction with default activity P and compensation Q . Then we have:*

$$t[D[[P]]_\rho] \mid \frac{d}{D} \parallel Q \parallel_{t,\rho}^{\text{npb}_D(P)} \mid \bar{l}_t.k_t \rightarrow^* D[[\text{extr}_D(P)]]_\rho \mid D[[\langle P \rangle]]_\rho$$

Lemma 5.4. *If R is a compensable process such that all free occurrences of process variable X in it are replaced with a process Q then the following encoding holds: $[[R\{Q/X\}]]_\rho = [[R]]_\rho\{[[Q]]_\rho/X\}$.*

Operational correspondence for the extended encoding follows from the following theorem:

Theorem 5.5. *Let P be a compensable process and let ρ be an arbitrary path.*

- If $P \xrightarrow{\tau}_D P'$ then there is an adaptable process P'' such that $D[[P]]_\rho \rightarrow^* P''$ and $P'' \approx D[[P']]_\rho$.*
- If $D[[P]]_\rho \rightarrow Q$ then there is P' such that $P \xrightarrow{\tau}_D P'$ and $Q \rightarrow^* D[[P']]_\rho$.*

Preserving Semantics. The function that counts the number of protected blocks in $\text{inst}[\lambda Y.R].P$ is the same as in Def. 5.1, while a function that counts the number of transactions is defined next.

Definition 5.6 (Number of transactions). *Let P be a compensable process such that $P = \text{inst}[\lambda Y.R].P_1$. The number of transactions which occur in P , denoted $\text{nts}(P)$, is equal to $\text{nts}(P_1)$.*

We have the following extension of Def. 4.9:

Definition 5.7 (Auxiliary encoding). *Let Q be a compensable process and let $\rho_0 = t, \rho$ be a path. Also, let $n, m \geq 0$. The process $\frac{d}{P} \parallel Q \parallel_{\rho_0}^{n,m}$ is defined as follows:*

$$\begin{aligned} \frac{d}{P} \parallel Q \parallel_{t,\rho}^{0,0} &= l_t.\bar{m}_t.a.p_\rho [u[\bar{f}.\bar{g}]] \mid m_t.\bar{k}_t.t\{\dagger\} \mid v[u\{(Z).(Z \mid v_1[P[[Q]]_\varepsilon] \mid f.v_1\{\text{id}\}.v\{\text{id}\}.g\}\})] \\ \frac{d}{P} \parallel Q \parallel_{t,\rho}^{1,0} &= l_t.p_{t,\rho} \left\{ (X_1).z\{a.p_\rho[X_1] \mid \bar{m}_t.p_\rho[u[\bar{f}.\bar{g}]]\} \right\} .(z[\mathbf{0}] \mid m_t.\bar{k}_t.t\{\dagger\}) \\ &\quad \mid v[u\{(Z).(Z \mid v_1[P[[Q]]_\varepsilon] \mid f.v_1\{\text{id}\}.v\{\text{id}\}.g\}\})] \\ \frac{d}{P} \parallel Q \parallel_{t,\rho}^{0,1} &= l_t.\beta_{t,\rho} \left\{ (Y_1).z\{a.\beta_\rho[Y_1] \mid \bar{m}_t.p_\rho[u[\bar{f}.\bar{g}]]\} \right\} .(z[\mathbf{0}] \mid m_t.\bar{k}_t.t\{\dagger\}) \\ &\quad \mid v[u\{(Z).(Z \mid v_1[P[[Q]]_\varepsilon] \mid f.v_1\{\text{id}\}.v\{\text{id}\}.g\}\})] \\ \frac{d}{P} \parallel Q \parallel_{t,\rho}^{n,m} &= l_t.p_{t,\rho} \left\{ (X_1, \dots, X_n).\beta_{t,\rho} \left\{ (Y_1, \dots, Y_m).z\{p_\rho[X_1] \mid p_\rho[X_2] \mid \dots \mid p_\rho[X_n] \right. \right. \\ &\quad \left. \left. \mid a.(\beta_\rho[Y_1] \mid \dots \mid \beta_\rho[Y_m]) \mid \bar{m}_t.p_\rho[u[\bar{f}.\bar{g}]]\} \right\} \right\} .(z[\mathbf{0}] \mid m_t.\bar{k}_t.t\{\dagger\}) \mid v[u\{(Z).(Z \mid v_1[P[[Q]]_\varepsilon] \\ &\quad \mid f.v_1\{\text{id}\}.v\{\text{id}\}.g\}\})] \quad [n, m > 0] \end{aligned}$$

We then have the following extended correctness statements:

Lemma 5.8. *Let $t[P, Q]$ be a transaction with default activity P and compensation Q . Then we have:*

$$\beta_\rho \left[t[P[[P]]_{t,\rho}] \mid \frac{d}{P} \parallel Q \parallel_{t,\rho}^{\text{npb}_P(P), \text{nts}(P)} \mid \bar{l}_t.k_t.\bar{j} \mid j.\beta_\rho\{\text{id}\}.\bar{a} \rightarrow^* P[[\text{extr}_P(P)]]_\rho \mid P[[\langle Q \rangle]]_\rho \right]$$

Theorem 5.9. *Let P be a compensable process and let ρ be an arbitrary path.*

- If $P \xrightarrow{\tau}_P P'$ then there is an adaptable process P'' such that $P[[P]]_\rho \rightarrow^* P''$ and $P'' \approx P[[P']]_\rho$.*
- If $P[[P]]_\rho \rightarrow Q$ then there is P' such that $P \xrightarrow{\tau}_P P'$ and $Q \rightarrow^* P[[P']]_\rho$.*

Aborting Semantics. The encoding of processes with dynamic compensations and aborting semantics is obtained by extending Def. 4.17 with the encodings of process variables and compensation updates, which are the same as in discarding and preserving semantics. The function that counts protected blocks in compensation updates $\text{npb}(\text{inst}[\lambda Y.R].P)$ is as in Def. 5.1. We require an extension to Def. 4.15:

Definition 5.10 (Auxiliary encoding). *Let Q be a compensable process and let $\rho_0 = t, \rho$ be a path. Also, let $n \geq 0$. The process $\mathbb{A}^d\|Q\|_{\rho_0}^n$ is defined as follows:*

$$\begin{aligned} \mathbb{A}^d\|Q\|_{t,\rho}^0 &= l_t.\overline{m}_t.p_\rho[u[\overline{f}.\overline{g}]] \mid m_t.\overline{k}_t.t\{\dagger\}.\Gamma_{t,\rho} \mid v[u\{(Z).(Z \mid v_1[\mathbb{A}[Q]_\varepsilon] \mid f.v_1\{\text{id}\}.v\{\text{id}\}.g)\}] \\ \mathbb{A}^d\|Q\|_{t,\rho}^n &= l_t.p_{t,\rho} \left\{ (X_1, \dots, X_n).z\{p_\rho[X_1] \mid p_\rho[X_2] \mid \dots \mid p_\rho[X_n] \mid \overline{m}_t.p_\rho[u[\overline{f}.\overline{g}]]\} \right\} \\ &\quad .(z[\mathbf{0}] \mid m_t.\overline{k}_t.t\{\dagger\}.\Gamma_{t,\rho}) \mid v[u\{(Z).(Z \mid v_1[\mathbb{A}[Q]_\varepsilon] \mid f.v_1\{\text{id}\}.v\{\text{id}\}.g)\}] \quad [n > 0] \end{aligned}$$

We then have the following extended correctness statements:

Lemma 5.11. *Let $t[P, Q]$ be a transaction with default activity P and compensation Q . We have:*

$$t[\mathbb{A}[P]_{t,\rho} \mid \mathbb{A}^d\|Q\|_{t,\rho}^{\text{npb}_A(P)} \mid \gamma_t[\mathcal{T}_t(P)] \rightarrow^* \mathbb{A}[\text{extr}_A(P)]_\rho \mid \mathbb{A}[\langle Q \rangle]_\rho \mid \Gamma_{t,\rho} \mid \gamma_t[\mathbf{0}].$$

Theorem 5.12. *Let P be a compensable process and let ρ be an arbitrary path.*

- a) *If $P \xrightarrow{\tau}_A P'$ then there is an adaptable process P'' such that $\mathbb{A}[P]_\rho \rightarrow^* P''$ and $P'' \approx \mathbb{A}[P']_\rho$.*
- b) *If $\mathbb{A}[P]_\rho \rightarrow Q$ then there is P' such that $P \xrightarrow{\tau}_A P'$ and $Q \rightarrow^* \mathbb{A}[P']_\rho$.*

6 Concluding Remarks

We have compared, from the point of view of relative expressiveness, two related and yet fundamentally different process models: the calculus of *compensable processes* (introduced in [12]) and the calculus of *adaptable processes* (introduced in [2]). We provided encodings of processes with static and dynamic compensations (under discarding, preserving, and aborting semantics) into adaptable processes. Our encodings not only are a non trivial application of process mobility as present in adaptable processes; they also shed light on the intricate semantics of compensable processes. As encoding criteria, we have considered compositionality and operational correspondence (up-to weak equivalences), as in [11]. It would be insightful to establish encoding correctness with respect to all the criteria in [11].

Our study opens several interesting avenues for future work. Having addressed the encodability of compensable processes into adaptable processes, we plan to consider the reverse direction, i.e., encodings of adaptable processes into compensable processes. We conjecture that an encoding of adaptable process into a language with static compensations does not exist: compensation updates $\text{inst}[\lambda X.Q].P$ seem essential to model an update prefix $l\{X\}.Q.P$ —the semantics of both constructs induces process substitutions. Still, even by considering a language with dynamic compensations, an encoding of adaptable processes is far from obvious, because the semantics of compensation updates dynamically modifies the behavior of the compensation activity, the inactive part of a transaction. Formalizing these (non) encodability claims is interesting future work. Another promising direction is to cast our encodability results into variants of adaptable and compensable processes with *session types*: a candidate for source language could be the typed calculus with *interactional exceptions* developed in [6]; as target language, we plan to consider extensions of adaptable processes with session types [9, 10].

Acknowledgements. We are grateful to the anonymous reviewers for their comments and suggestions. This research was partially supported by the EU COST Action IC1201. Pérez is also affiliated to NOVA Laboratory for Computer Science and Informatics, Universidade Nova de Lisboa, Portugal.

References

- [1] Laura Bocchi, Cosimo Laneve & Gianluigi Zavattaro (2003): *A Calculus for Long-Running Transactions*. In: *Proc. of FMOODS 2003*, LNCS 2884, Springer, pp. 124–138, doi:10.1007/978-3-540-39958-2_9.
- [2] Mario Bravetti, Cinzia Di Giusto, Jorge A. Pérez & Gianluigi Zavattaro (2012): *Adaptable processes*. *Logical Methods in Computer Science* 8(4), doi:10.2168/LMCS-8(4:13)2012.
- [3] Mario Bravetti, Cinzia Di Giusto, Jorge A. Pérez & Gianluigi Zavattaro (2012): *Towards the Verification of Adaptable Processes*. In: *ISoLA*, LNCS 7609, Springer, pp. 269–283, doi:10.1007/978-3-642-34026-0_20.
- [4] Mario Bravetti & Gianluigi Zavattaro (2009): *On the expressive power of process interruption and compensation*. *Mathematical Structures in Computer Science* 19(3), pp. 565–599, doi:10.1017/S0960129509007683.
- [5] Luís Caires, Carla Ferreira & Hugo Torres Vieira (2009): *A Process Calculus Analysis of Compensations*. In: *Proc. of TGC 2008*, LNCS 5474, Springer, pp. 87–103, doi:10.1007/978-3-642-00945-7_6.
- [6] Marco Carbone, Kohei Honda & Nobuko Yoshida (2008): *Structured Interactional Exceptions in Session Types*. In: *CONCUR 2008*, LNCS 5201, Springer, pp. 402–417, doi:10.1007/978-3-540-85361-9_32.
- [7] Jovana Dedeić, Jovanka Pantović & Jorge A. Pérez (2015): *Full version of this paper*. Technical Report. Available at <http://www.jorgeaperez.net>.
- [8] Carla Ferreira, Ivan Lanese, António Ravara, Hugo Torres Vieira & Gianluigi Zavattaro (2011): *Advanced Mechanisms for Service Combination and Transactions*. In: *Results of SENSORIA*, LNCS 6582, Springer, pp. 302–325, doi:10.1007/978-3-642-20401-2_14.
- [9] Cinzia Di Giusto & Jorge A. Pérez (2015): *Disciplined structured communications with disciplined runtime adaptation*. *Sci. Comput. Program.* 97, pp. 235–265, doi:10.1016/j.scico.2014.04.017.
- [10] Cinzia Di Giusto & Jorge A. Pérez (2015): *An Event-Based Approach to Runtime Adaptation in Communication-Centric Systems*. In: *Proc. of WS-FM 2014*, LNCS, Springer. To appear.
- [11] Daniele Gorla (2010): *Towards a unified approach to encodability and separation results for process calculi*. *Inf. Comput.* 208(9), pp. 1031–1053, doi:10.1016/j.ic.2010.05.002.
- [12] Ivan Lanese, Cátia Vaz & Carla Ferreira (2010): *On the Expressive Power of Primitives for Compensation Handling*. In: *Proc. of ESOP 2010*, LNCS 6012, Springer, pp. 366–386, doi:10.1007/978-3-642-11957-6_20.
- [13] Ivan Lanese & Gianluigi Zavattaro (2013): *Decidability Results for Dynamic Installation of Compensation Handlers*. In: *COORDINATION*, LNCS 7890, Springer, pp. 136–150, doi:10.1007/978-3-642-38493-6_10.
- [14] Cosimo Laneve & Gianluigi Zavattaro (2005): *Foundations of Web Transactions*. In: *Proc. of FOSSACS 2005*, LNCS 3441, Springer, pp. 282–298, doi:10.1007/978-3-540-31982-5_18.
- [15] Robin Milner (1989): *Communication and concurrency*. PHI Series in computer science, Prentice Hall.
- [16] Robin Milner, Joachim Parrow & David Walker (1992): *A Calculus of Mobile Processes, I*. *Inf. Comput.* 100(1), pp. 1–40, doi:10.1016/0890-5401(92)90008-4.
- [17] Davide Sangiorgi (1992): *Expressing Mobility in Process Algebras: First-Order and Higher Order Paradigms*. Ph.D. thesis, University of Edinburgh.

SOS rule formats for convex and abstract probabilistic bisimulations*

Pedro R. D’Argenio

Matias David Lee

FaMAF, Universidad Nacional de Córdoba – CONICET.
Ciudad Universitaria, X5000HUA – Córdoba, Argentina.

dargenio@famaf.unc.edu.ar

lee@famaf.unc.edu.ar

Daniel Gebler

Department of Computer Science, VU University Amsterdam,
De Boelelaan 1081a, NL-1081 HV Amsterdam, The Netherlands

e.d.gebler@vu.nl

Probabilistic transition system specifications (PTSSs) in the $nt\mu f\theta/nt\mu x\theta$ format provide structural operational semantics for Segala-type systems that exhibit both probabilistic and nondeterministic behavior and guarantee that bisimilarity is a congruence for all operator defined in such format. Starting from the $nt\mu f\theta/nt\mu x\theta$, we obtain restricted formats that guarantee that three coarser bisimulation equivalences are congruences. We focus on (i) Segala’s variant of bisimulation that considers combined transitions, which we call here *convex bisimulation*; (ii) the bisimulation equivalence resulting from considering Park & Milner’s bisimulation on the usual stripped probabilistic transition system (translated into a labelled transition system), which we call here *probability obliterated bisimulation*; and (iii) a *probability abstracted bisimulation*, which, like bisimulation, preserves the structure of the distributions but instead, it ignores the probability values. In addition, we compare these bisimulation equivalences and provide a logic characterization for each of them.

1 Introduction

Structural operational semantics (SOS for short) [24] is a powerful tool to provide semantics to programming languages. In SOS, process behavior is described using transition systems and the behavior of a composite process is given in terms of the behavior of its components. SOS has been formalized using an algebraic framework as *Transition Systems Specifications (TSS)* [6, 7, 14, 15, 23, etc.]. Basically, a TSS contains a signature, a set of actions or labels, and a set of rules. The signature defines the terms in the language. The set of actions represents all possible activities that a process (i.e., a term over the signature) can perform. The rules define how a process should behave (i.e., perform certain activities) in terms of the behavior of its subprocesses, that is, the rules define compositionally the transition system associated to each term of the language. A particular focus of these formalizations was to provide a meta-theory that ensures a diversity of semantic properties by simple inspection on the form of the rules. (See [1, 2, 23] for overviews.) One of such kind of properties is to ensure that a given equivalence relation is a congruence for all operators whose semantics is defined in a TSS whose rules complies to a particular format. These so called *congruence theorems* have been proved for a variety of equivalences in the non-probabilistic case [6, 14, 15, etc.].

*Supported by ANPCyT project PICT-2012-1823, SeCyT-UNC program 05/BP12 and their related projects, and EU 7FP grant agreement 295261 (MEALS).

The introduction of probabilistic process algebras motivated the need for a theory of structural operational semantics to define *probabilistic* transition systems. Few earlier results appeared in this direction [4, 5, 17, 18] presenting congruence theorems for Larsen & Skou bisimulation equivalence [19]. Most of these formats have complicated restrictions that extend to sets of rules due to the fact that they considered transitions labeled both with an action and a probability value. By using a more modern view of probabilistic transition systems (where the target of the transition is a probability distribution on states) we manage to obtain the most general format for bisimulation equivalence, which we called $nt\mu f\theta/nt\mu x\theta$, following the nomenclature of [14, 15].

Starting from the $nt\mu f\theta/nt\mu x\theta$ format, in this paper we define formats to guarantee that three coarser versions of bisimulation equivalence are congruences for all operator definable in the respective format. The first relation we focus on is Segala's variant of bisimulation that considers combined transitions, here called *convex bisimulation* [25]. The second relation we explore originates here and we call it *probability abstracted bisimulation*. Like bisimulation and unlike convex bisimulation, it preserves the structure of the distributions of each transition, but instead, it ignores the probability values. This relation preserves the fairness introduced by the probability distributions. Finally, we study the bisimulation equivalence resulting from considering Park & Milner's bisimulation [22] on the usual stripped probabilistic transition system (translated into a labeled transition system). Here we call it *probability obliterated bisimulation*. This is the usual way to abstract probabilities, but it has the drawback that it breaks the basic fairness provided by probabilistic choices.

Apart from presenting congruence theorems for all previously mentioned bisimulation equivalences, we briefly study alternative definitions of these bisimulations, compare them with each other, and provide logical characterizations, which are particularly new here for probability abstracted and probability obliterated bisimulation equivalences.

The paper is organized as follows. Sec. 2 recalls the type of algebraic structure and Sec. 3 provides the basic notions and results of probabilistic transition system specifications (PTSS). Sec. 4 presents the different bisimulation equivalences and a brief study of them, including their logical characterizations. The study of all the PTSS formats and the respective congruence theorems is given in Sec. 5. The paper concludes in Sec 6.

2 Preliminaries

Let $S = \{s, d\}$ be a set denoting two sorts. Elements of sort $s \in S$ are intended to represent states in the transition system, while elements of sort $d \in S$ will represent distributions over states. We let σ range over S . An S -sorted signature is a structure (F, ar) , where (i) F is a set of *function names*, and (ii) $\text{ar}: F \rightarrow (S^* \times S)$ is the *arity function*. The rank of $f \in F$ is the number of arguments of f , defined by $\text{rk}(f) = n$ if $\text{ar}(f) = \sigma_1 \dots \sigma_n \rightarrow \sigma$. (We write " $\sigma_1 \dots \sigma_n \rightarrow \sigma$ " instead of " $(\sigma_1 \dots \sigma_n, \sigma)$ " to highlight that function f maps to sort σ .) Function f is a *constant* if $\text{rk}(f) = 0$. To simplify the presentation we will write an S -sorted signature (F, ar) as a pair of disjoint signatures (Σ_s, Σ_d) where Σ_s is the set of operations that map to s and Σ_d is the set of operations that map to d . Let \mathcal{V} and \mathcal{V}_d be two infinite sets of S -sorted variables where $\mathcal{V}, \mathcal{V}_d, F$ are all mutually disjoint. We use x, y, z (with possible sub- or super-scripts) to range over \mathcal{V} , μ, ν to range over \mathcal{V}_d and ζ to range over $\mathcal{V} \cup \mathcal{V}_d$.

Definition 1. Let Σ_s and Σ_d be two signatures as before and let $V \subseteq \mathcal{V}$ and $D \subseteq \mathcal{V}_d$. We simultaneously define the sets of state terms $T(\Sigma_s, V, D)$ and distribution terms $T(\Sigma_d, V, D)$ as the smallest sets satisfying: (i) $V \subseteq T(\Sigma_s, V, D)$; (ii) $D \subseteq T(\Sigma_d, V, D)$; (iii) $f(\xi_1, \dots, \xi_{\text{rk}(f)}) \in T(\Sigma_\sigma, V, D)$, if $\text{ar}(f) = \sigma_1 \dots \sigma_n \rightarrow \sigma$ and $\xi_i \in T(\Sigma_{\sigma_i}, V, D)$.

We let $\mathbb{T}(\Sigma) = T(\Sigma_s, \mathcal{V}, \mathcal{V}_d) \cup T(\Sigma_d, \mathcal{V}, \mathcal{V}_d)$ denote the set of all *open terms* and distinguish the sets $\mathbb{T}(\Sigma_s) = T(\Sigma_s, \mathcal{V}, \mathcal{V}_d)$ of *open state terms* and $\mathbb{T}(\Sigma_d) = T(\Sigma_d, \mathcal{V}, \mathcal{V}_d)$ of *open distribution terms*. Similarly, we let $\mathbb{T}(\Sigma) = T(\Sigma_s, \emptyset, \emptyset) \cup T(\Sigma_d, \emptyset, \emptyset)$ denote the set of all *closed terms* and distinguish the sets $\mathbb{T}(\Sigma_s) = T(\Sigma_s, \emptyset, \emptyset)$ of *closed state terms* and $\mathbb{T}(\Sigma_d) = T(\Sigma_d, \emptyset, \emptyset)$ of *closed distribution terms*. We let t, t', t_1, \dots range over state terms, $\theta, \theta', \theta_1, \dots$ range over distribution terms, and ξ, ξ', ξ_1, \dots range over any kind of terms. With $\mathcal{V}(\xi) \subseteq \mathcal{V} \cup \mathcal{V}_d$ we denote the set of variables occurring in term ξ .

Let $\Delta(\mathbb{T}(\Sigma_s))$ denote the set of all (discrete) probability distributions on $\mathbb{T}(\Sigma_s)$. We let π range over $\Delta(\mathbb{T}(\Sigma_s))$. For each $t \in \mathbb{T}(\Sigma_s)$, let $\delta_t \in \Delta(\mathbb{T}(\Sigma_s))$ denote the *Dirac distribution*, i.e., $\delta_t(t) = 1$ and $\delta_t(t') = 0$ if t and t' are not syntactically equal. For $X \subseteq \mathbb{T}(\Sigma_s)$ we define $\pi(X) = \sum_{t \in X} \pi(t)$. The convex combination $\sum_{i \in I} p_i \pi_i$ of a family $\{\pi_i\}_{i \in I}$ of probability distributions with $p_i \in (0, 1]$ and $\sum_{i \in I} p_i = 1$ is defined by $(\sum_{i \in I} p_i \pi_i)(t) = \sum_{i \in I} (p_i \pi_i(t))$.

The type of signatures we consider has a particular construction. We start from a signature Σ_s of functions mapping into sort s and construct the signature Σ_d of functions mapping into d as follows. For each $f \in F_s$ we include a function symbol $\mathbf{f} \in F_d$ with $\text{ar}(\mathbf{f}) = d \dots d \rightarrow d$ and $\text{rk}(\mathbf{f}) = \text{rk}(f)$. We call \mathbf{f} the *probabilistic lifting* of f . (We use boldface fonts to indicate that a function in Σ_d is the probabilistic lifting of another in Σ_s .) Moreover Σ_d may include any of the following additional operators: (i) δ with arity $\text{ar}(\delta) = s \rightarrow d$, and (ii) $\bigoplus_{i \in I} [p_i]_{-}$ with I being a finite or countable infinite index set, $\sum_{i \in I} p_i = 1$, $p_i \in (0, 1]$ for all $i \in I$, and $\text{ar}(\bigoplus_{i \in I} [p_i]_{-}) = d^{|I|} \rightarrow d$. Notice that if I is countably infinite, $\bigoplus_{i \in I} [p_i]_{-}$ is an infinitary operator.

Operators δ and $\bigoplus_{i \in I} [p_i]_{-}$ are used to construct discrete probability functions of countable support: $\delta(t)$ is interpreted as a distribution that assigns probability 1 to the state term t and probability 0 to any other term t' (syntactically) different from t , and $\bigoplus_{i \in I} [p_i] \theta_i$ represents a distribution that weights with p_i the distribution represented by the term θ_i . Moreover, a probabilistically lifted operator \mathbf{f} is interpreted by properly lifting the probabilities of the operands to terms composed with the operator f .

Formally, the algebra associated with a probabilistically lifted signature $\Sigma = (\Sigma_s, \Sigma_d)$ is defined as follows. For sort s , it is the freely generated algebraic structure $\mathbb{T}(\Sigma_s)$. For sort d , it is defined by the carrier $\Delta(\mathbb{T}(\Sigma_s))$ and the following interpretation: $\llbracket \delta(t) \rrbracket = \delta_t$ for all $t \in \mathbb{T}(\Sigma_s)$, $\llbracket \bigoplus_{i \in I} [p_i] \theta_i \rrbracket = \sum_{i \in I} p_i \llbracket \theta_i \rrbracket$ for $\{\theta_i \mid i \in I\} \subseteq \mathbb{T}(\Sigma_d)$, $\llbracket \mathbf{f}(\theta_1, \dots, \theta_{\text{rk}(\mathbf{f})}) \rrbracket (f(\xi_1, \dots, \xi_{\text{rk}(\mathbf{f})})) = \prod_{\sigma_j = s} \llbracket \theta_j \rrbracket (\xi_j)$ if for all j s.t. $\sigma_j = d$, $\theta_j = \xi_j$, and $\llbracket \mathbf{f}(\theta_1, \dots, \theta_{\text{rk}(\mathbf{f})}) \rrbracket (f(\xi_1, \dots, \xi_{\text{rk}(\mathbf{f})})) = 0$ otherwise. Here it is assumed that $\prod \emptyset = 1$. Notice that in the semantics of a lifted function \mathbf{f} , the big product only considers the distributions related to the s -sorted positions in f , while the distribution terms corresponding to the d -sorted positions in f should match exactly to the parameters of f .

A *substitution* ρ is a map $\mathcal{V} \cup \mathcal{V}_d \rightarrow \mathbb{T}(\Sigma)$ such that $\rho(x) \in \mathbb{T}(\Sigma_s)$, for all $x \in \mathcal{V}$, and $\rho(\mu) \in \mathbb{T}(\Sigma_d)$, for all $\mu \in \mathcal{V}_d$. A substitution is closed if it maps each variable to a closed term. A substitution extends to a mapping from terms to terms as usual.

Finally, we remark a general property of distribution terms: let $f \in \Sigma_s$ with $\text{ar}(f) = \sigma_1 \dots \sigma_n \rightarrow s$, and let $\sigma_j = s$; then $\mathbf{f} \in \Sigma_d$ is distributive w.r.t. \oplus in the position j , i.e. $\llbracket \rho(\mathbf{f}(\dots, \xi_{j-1}, \bigoplus_{i \in I} [p_i] \theta_i, \xi_{j+1}, \dots)) \rrbracket = \llbracket \rho(\bigoplus_{i \in I} [p_i] \mathbf{f}(\dots, \xi_{j-1}, \theta_i, \xi_{j+1}, \dots)) \rrbracket$ for any closed substitution ρ . The proof follows from the definition of $\llbracket _ \rrbracket$. However, notice that \mathbf{f} does not distribute w.r.t. \oplus in a position k such that $\sigma_k = d$.

3 Probabilistic Transition System Specifications

A (probabilistic) transition relation prescribes which possible activity can be performed by a term in a signature. Such activity is described by the label of the action and a probability distribution on terms that indicates the probability to reach a particular new term. We will follow the probabilistic automata style

of probabilistic transitions [25] which are a generalization of the so-called reactive model [19].

Definition 2 (PTS). A probabilistic labeled transition system (PTS) is a triple $(T(\Sigma_s), A, \rightarrow)$, where $\Sigma = (\Sigma_s, \Sigma_d)$ is a probabilistically lifted signature, A is a countable set of actions, and $\rightarrow \subseteq T(\Sigma_s) \times A \times \Delta(T(\Sigma_s))$, is a transition relation. We write $t \xrightarrow{a} \pi$ for $(t, a, \pi) \in \rightarrow$.

Transition relations are usually defined by means of structured operational semantics in Plotkin's style [24]. For PTS, algebraic characterizations of this style were provided in [8, 9, 21] where the term *probabilistic transition system specification* was used and which we adopt in our paper.

Definition 3 (PTSS). A probabilistic transition system specification (PTSS) is a triple $P = (\Sigma, A, R)$ where Σ is a probabilistically lifted signature, A is a set of labels, and R is a set of rules of the form:

$$\frac{\{t_k \xrightarrow{a_k} \theta_k \mid k \in K\} \cup \{t_l \xrightarrow{b_l} \theta_l \mid l \in L\} \cup \{\theta_j(T_j) \bowtie_j q_j \mid j \in J\}}{t \xrightarrow{a} \theta}$$

where K, L, J are index sets, $t, t_k, t_l \in T(\Sigma_s)$, $a, a_k, b_l \in A$, $T_j \subseteq T(\Sigma_s)$, $\bowtie_j \in \{>, \geq, <, \leq\}$, $q_j \in [0, 1]$ and $\theta_j, \theta_k, \theta \in T(\Sigma_d)$.

Expressions of the form $t \xrightarrow{a} \theta$, $t \xrightarrow{a} \theta$, and $\theta(T) \bowtie p$ are called *positive literal*, *negative literal*, and *quantitative literal*, respectively. For any rule $r \in R$, literals above the line are called *premises*, notation $\text{prem}(r)$; the literal below the line is called *conclusion*, notation $\text{conc}(r)$. We denote with $\text{pprem}(r)$, $\text{nprem}(r)$, and $\text{qprem}(r)$ the sets of positive, negative, and quantitative premises of the rule r , respectively. In general, we allow the sets of positive, negative, and quantitative premises to be infinite.

Substitutions provide instances to the rules of a PTSS that, together with some appropriate machinery, allow us to define probabilistic transition relations. Given a substitution ρ , it extends to literals as follows: $\rho(t \xrightarrow{a} \theta) = \rho(t) \xrightarrow{a} \rho(\theta)$, $\rho(\theta(T) \bowtie p) = \rho(\theta)(\rho(T)) \bowtie p$ (where $\rho(T) = \{\rho(t) \mid t \in T\}$), and $\rho(t \xrightarrow{a} \theta) = \rho(t) \xrightarrow{a} \rho(\theta)$.

We say that r' is a (closed) instance of a rule r if there is a (closed) substitution ρ so that $r' = \rho(r)$. We say that ρ is a *proper substitution of r* if for all quantitative premises $\theta(T) \bowtie p$ of r and all $t \in T$, $\llbracket \rho(\theta) \rrbracket(\rho(t)) > 0$ holds. We use only this kind of substitution in the paper.

In the rest of the paper, we will deal with models as *symbolic* transition relations in the set $T(\Sigma_s) \times A \times T(\Sigma_d)$ rather than the *concrete* transition relations in $T(\Sigma_s) \times A \times \Delta(T(\Sigma_s))$ required by a PTS. Hence we will mostly refer with the term “transition relation” to the symbolic transition relation. In any case, a symbolic transition relation induces always a unique concrete transition relation by interpreting every target distribution term as the distribution it defines; that is, the symbolic transition $t \xrightarrow{a} \theta$ is interpreted as the concrete transition $t \xrightarrow{a} \llbracket \theta \rrbracket$. If the symbolic transition relation turns out to be a model of a PTSS P , we say that the induced concrete transition relation defines a PTS associated to P .

To define an appropriate notion of model we consider *3-valued models*. A 3-valued model partitions the set $T(\Sigma_s) \times A \times T(\Sigma_d)$ in three sets containing, respectively, the transition that are known to hold, that are known not to hold, and those whose validity is unknown. Thus, a 3-valued model can be presented as a pair $\langle \text{CT}, \text{PT} \rangle$ of transition relations $\text{CT}, \text{PT} \subseteq T(\Sigma_s) \times A \times T(\Sigma_d)$, with $\text{CT} \subseteq \text{PT}$, where CT is the set of transitions that *certainly* hold and PT is the set of transitions that *possibly* hold. So, transitions in $\text{PT} \setminus \text{CT}$ are those whose validity is unknown and transitions in $(T(\Sigma_s) \times A \times T(\Sigma_d)) \setminus \text{PT}$ are those that certainly do not hold. A 3-valued model $\langle \text{CT}, \text{PT} \rangle$ that is justifiably compatible with the proof system defined by a PTSS P is said to be *stable* for P . (See Def. 5.)

Before formally defining the notions of proof and 3-valued stable model we introduce some notation. Given a transition relation $\text{Tr} \subseteq T(\Sigma_s) \times A \times T(\Sigma_d)$, $t \xrightarrow{a} \theta$ holds in Tr , notation $\text{Tr} \models t \xrightarrow{a} \theta$, if $t \xrightarrow{a} \theta \in \text{Tr}$; $t \xrightarrow{a} \theta$ holds in Tr , notation $\text{Tr} \models t \xrightarrow{a} \theta$, if for all $\theta \in T(\Sigma_d)$, $t \xrightarrow{a} \theta \notin \text{Tr}$. A closed quantitative constraint

$\theta(T) \bowtie p$ holds in Tr , notation $\text{Tr} \models \theta(T) \bowtie p$, if $\llbracket \theta \rrbracket(T) \bowtie p$. Notice that the satisfaction of a quantitative constraint does not depend on the transition relation. We nonetheless use this last notation as it turns out to be convenient. Given a set of literals H , we write $\text{Tr} \models H$ if for all $\phi \in H$, $\text{Tr} \models \phi$.

Definition 4 (Proof). *Let $P = (\Sigma, A, R)$ be a PTSS. Let ψ be a positive literal and let H be a set of literals. A proof of a transition rule $\frac{H}{\psi}$ from P is a well-founded, upwardly branching tree where each node is a literal such that: (i) the root is ψ ; and (ii) if χ is a node and K is the set of nodes directly above χ , then one of the following conditions holds: (a) $K = \emptyset$ and $\chi \in H$, or (b) $\chi = (\theta(T) \bowtie p)$ is a closed quantitative literal such that $\llbracket \theta \rrbracket(T) \bowtie p$ holds, or (c) $\frac{K}{\chi}$ is a valid substitution instance of a rule from R .*

$\frac{H}{\psi}$ is provable from P , notation $P \vdash \frac{H}{\psi}$, if there exists a proof of $\frac{H}{\psi}$ from P .

Before, we said that a 3-valued stable model $\langle \text{CT}, \text{PT} \rangle$ for a PTSS P has to be *justifiably compatible* with the proof system defined by P . By “compatible” we mean that $\langle \text{CT}, \text{PT} \rangle$ has to be consistent with every provable rule. With “justifiable” we require that for each transition in CT and PT there is actually a proof that justifies it. More precisely, we require that (a) for every certain transition in CT there is a proof in P such that all negative hypotheses of the proof are known to hold (i.e. there is no possible transition in PT denying a negative hypothesis), and (b) for every possible transition in PT there is a proof in P such that all negative hypotheses possibly hold (i.e. there is no certain transition in CT denying a negative hypothesis). This is formally stated in the next definition.

Definition 5 (3-valued stable model). *Let $P = (\Sigma, A, R)$ be a PTSS. A tuple $\langle \text{CT}, \text{PT} \rangle$ with $\text{CT} \subseteq \text{PT} \subseteq T(\Sigma_s) \times A \times T(\Sigma_d)$ is a 3-valued stable model for P if for every closed positive literal ψ ,*

(a) $\psi \in \text{CT}$ iff there is a set N of closed negative literals such that $P \vdash \frac{N}{\psi}$ and $\text{PT} \models N$

(b) $\psi \in \text{PT}$ iff there is a set N of closed negative literals such that $P \vdash \frac{N}{\psi}$ and $\text{CT} \models N$.

The least 3-valued stable model of a PTSS can be constructed using induction [8, 11, 12].

Lemma 1. *Let P be a PTSS. For each ordinal α , define the pair $\langle \text{CT}_\alpha, \text{PT}_\alpha \rangle$ as follows:*

- $\text{CT}_0 = \emptyset$ and $\text{PT}_0 = T(\Sigma_s) \times A \times T(\Sigma_d)$.
- For every non-limit ordinal $\alpha > 0$, define:

$$\text{CT}_\alpha = \left\{ t \xrightarrow{a} \theta \mid \text{for some set } N \text{ of negative literals, } P \vdash \frac{N}{t \xrightarrow{a} \theta} \text{ and } \text{PT}_{\alpha-1} \models N \right\}$$

$$\text{PT}_\alpha = \left\{ t \xrightarrow{a} \theta \mid \text{for some set } N \text{ of negative literals, } P \vdash \frac{N}{t \xrightarrow{a} \theta} \text{ and } \text{CT}_{\alpha-1} \models N \right\}$$

- For every limit ordinal α , define $\text{CT}_\alpha = \bigcup_{\beta < \alpha} \text{CT}_\beta$ and $\text{PT}_\alpha = \bigcap_{\beta < \alpha} \text{PT}_\beta$.

Then: 1. if $\beta \leq \alpha$, $\text{CT}_\beta \subseteq \text{CT}_\alpha$ and $\text{PT}_\beta \supseteq \text{PT}_\alpha$, and 2. there is an ordinal λ such that $\text{CT}_\lambda = \text{CT}_{\lambda+1}$ and $\text{PT}_\lambda = \text{PT}_{\lambda+1}$. Moreover, $\langle \text{CT}_\lambda, \text{PT}_\lambda \rangle$ is the least 3-valued stable model for P .

PTSSs with least 3-valued stable model that are also a 2-valued model are particularly interesting, since this model is actually the only 3-valued stable model [7, 13]. A PTSS P is said to be *complete* if its least 3-valued stable model $\langle \text{CT}, \text{PT} \rangle$ satisfies that $\text{CT} = \text{PT}$ (i.e., the model is also 2-valued). We associate a probabilistic transition system to each complete PTSS.

Definition 6. *Let P be a complete PTSS and let $\langle \text{Tr}, \text{Tr} \rangle$ be its unique 3-valued stable model. We say that Tr is the transition relation associated to P . We also define the PTS associated to P as the unique PTS $(T(\Sigma_s), A, \rightarrow)$ such that $t \xrightarrow{a} \pi$ if and only if $t \xrightarrow{a} \theta \in \text{Tr}$ and $\llbracket \theta \rrbracket = \pi$ for some $\theta \in T(\Sigma_d)$.*

The different examples that we give in the rest of the papers are in terms of a basic probabilistic process algebra. We introduce it here, but address the reader to [8] for an example of a PTSS with richer operators. Signature Σ_s contains the constant 0, representing the stop process, for each action $a \in A$, a unary probabilistic prefix operators $a._$ with arity $\text{ar}(a) = d \rightarrow s$, and a binary operator $+$, the alternative composition or sum, with arity $\text{ar}(+) = ss \rightarrow s$, while Σ_d contains the respective lifted signature, δ , and all binary operators $\bigoplus_{i \in \{1,2\}} [p_i] \theta_i$ which we denote by \oplus_p . The semantics is defined with the usual rules:

$$\frac{}{a.\mu \xrightarrow{a} \mu} \quad \frac{x \xrightarrow{a} \mu}{x+y \xrightarrow{a} \mu} \quad \frac{y \xrightarrow{a} \mu}{x+y \xrightarrow{a} \mu}$$

4 Bisimulation relations

This work revolves around four different types of bisimulation relations: (i) the usual (*strong*) *bisimulation* [19] relation on probabilistic system, in which each probabilistic transition should be matched with a single probabilistic transition so that the distributions of both transitions agree on the probabilities of jumping into equivalent states; (ii) the *convex bisimulation* [25] relation, in which the matching is performed instead with a convex combination of transition relations; (iii) the *probability abstracted bisimulation*, in which the matching is performed by a single transition so that the distributions of both transitions agree on jumping to the same equivalent classes of states but not necessarily with the same probability value; and (iv) the *probabilistic obliterated bisimulation*, which represents the usual bisimulation [22] once the probabilistic transition system is abstracted into a traditional labeled transition system in the usual way.

To our knowledge, the probability abstracted bisimulation originates here. Its intention is to strictly preserve the probabilistic structure of a system without caring about the probability values. Thus, probability abstracted bisimulation is consistent with any bisimulation preserving quantitative properties that only tests for positive quantifications, rather than a particular value. Instead, this kind of properties are not preserved by the probabilistic obliterated bisimulation as it is shown below in this section.

In the following we introduce all these relations and discuss their relationship as well as alternative definitions. For the rest of the section we assume given a PTS $P = (\mathcal{T}(\Sigma_s), A, \rightarrow)$.

Given a relation $R \subseteq \mathcal{T}(\Sigma_s) \times \mathcal{T}(\Sigma_s)$, a set $Q \subseteq \mathcal{T}(\Sigma_s)$ is *R-closed* if for all $t \in Q$ and $t' \in \mathcal{T}(\Sigma_s)$, $t R t'$ implies $t' \in Q$ (i.e. $R(Q) \subseteq Q$). It is easy to verify that if two relations $R, R' \subseteq \mathcal{T}(\Sigma_s) \times \mathcal{T}(\Sigma_s)$ are such that $R' \subseteq R$, then if $Q \subseteq \mathcal{T}(\Sigma_s)$ is *R-closed*, it is also *R'-closed*.

Definition 7. A relation $R \subseteq \mathcal{T}(\Sigma_s) \times \mathcal{T}(\Sigma_s)$ is a *bisimulation* if it is symmetric and for all $t, t' \in \mathcal{T}(\Sigma_s)$, $a \in A$, and $\pi \in \Delta(\mathcal{T}(\Sigma_s))$, $t R t'$ and $t \xrightarrow{a} \pi$ imply that there exists $\pi' \in \Delta(\mathcal{T}(\Sigma_s))$ s.t. $t' \xrightarrow{a} \pi'$ and $\pi R \pi'$, where $\pi R \pi'$ if and only if for all *R-closed* $Q \subseteq \mathcal{T}(\Sigma_s)$, $\pi(Q) = \pi'(Q)$. The relation \sim , called *bisimilarity* or *bisimulation equivalence*, is defined as the smallest relation that includes all bisimulations.

A *combined transition* $t \xrightarrow{a}_c \pi$ is defined whenever there is a family $\{\pi_i\}_{i \in I} \subseteq \Delta(\mathcal{T}(\Sigma_s))$ and a family $\{p_i\}_{i \in I} \subseteq [0, 1]$ such that $t \xrightarrow{a} \pi_i$ for all $i \in I$, $\sum_{i \in I} p_i = 1$ and $\pi = \sum_{i \in I} p_i \pi_i$.

Definition 8. A relation $R \subseteq \mathcal{T}(\Sigma_s) \times \mathcal{T}(\Sigma_s)$ is a *convex bisimulation* if it is symmetric and for all $t, t' \in \mathcal{T}(\Sigma_s)$, $a \in A$, and $\pi \in \Delta(\mathcal{T}(\Sigma_s))$, $t R t'$ and $t \xrightarrow{a} \pi$ imply that there exists $\pi' \in \Delta(\mathcal{T}(\Sigma_s))$ s.t. $t' \xrightarrow{a}_c \pi'$ and $\pi R \pi'$. The relation \sim_c , called *convex bisimilarity* or *convex bisimulation equivalence*, is defined as the smallest relation that includes all convex bisimulations.

Definition 9. A relation $R \subseteq \mathcal{T}(\Sigma_s) \times \mathcal{T}(\Sigma_s)$ is a *probability abstracted bisimulation* if it is symmetric and for all $t, t' \in \mathcal{T}(\Sigma_s)$, $a \in A$, and $\pi \in \Delta(\mathcal{T}(\Sigma_s))$, $t R t'$ and $t \xrightarrow{a} \pi$ imply that there exists $\pi' \in \Delta(\mathcal{T}(\Sigma_s))$ s.t. $t' \xrightarrow{a} \pi'$

and for all R -closed $Q \subseteq T(\Sigma_s)$, $\pi(Q) > 0$ iff $\pi'(Q) > 0$. The relation \sim_a , called probability abstracted bisimilarity or probability abstracted bisimulation equivalence, is defined as the smallest relation that includes all probability abstracted bisimulations.

Notice that the transfer property in this last case follows the same structure as the bisimulation, only that it only requires that $\pi(Q) > 0$ iff $\pi'(Q) > 0$ for all R -closed, instead of $\pi(Q) = \pi'(Q)$.

Definition 10. A relation $R \subseteq T(\Sigma_s) \times T(\Sigma_s)$ is a probability obliterated bisimulation if it is symmetric and for all $t, t' \in T(\Sigma_s)$, $a \in A$, and $\pi \in \Delta(T(\Sigma_s))$, $t R t'$ and $t \xrightarrow{a} \pi$ imply that for all R -closed $Q \subseteq T(\Sigma_s)$ with $\pi(Q) > 0$, there exists $\pi' \in \Delta(T(\Sigma_s))$ s.t. $t' \xrightarrow{a} \pi'$ and $\pi'(Q) > 0$. The relation \sim_o , called probability obliterated bisimilarity or probability obliterated bisimulation equivalence, is defined as the smallest relation that includes all probability obliterated bisimulations.

Compare this last definition with Def. 9. While for probability abstracted bisimulation we require that there is a single matching transition $t' \xrightarrow{a} \pi'$ so that π' gives some positive probability to all R -closed sets exactly whenever π does, the definition of probability obliterated bisimulation permits to choose different matching transitions for each R -closed set that measures positively on π .

It is well known that \sim and \sim_c are equivalences relations and that they also are, respectively, a bisimulation relation and a convex bisimulation relation. The fact that \sim_o is also an equivalence relation and itself a probability obliterated bisimulation follows from Lemma 4 which state that it agrees with Park & Milner's bisimulation. The same properties can be proven for probability abstracted bisimulation:

Lemma 2. \sim_a is an equivalence relation and is itself a probability abstracted bisimulation.

Similarly to the bisimulation [3, Prop 3.4.4], the probability abstracted bisimulation has a characterization in terms of an *abstract weight function*. This alternative characterization is the one used in the proof of Theorem 4 and that is why we present it in this paper.

Given a relation $R \subseteq T(\Sigma_s) \times T(\Sigma_s)$, we define $\equiv_R^w \in \Delta(T(\Sigma_s)) \times \Delta(T(\Sigma_s))$ as follows. For all $\pi, \pi' \in \Delta(T(\Sigma_s))$, $\pi \equiv_R^w \pi'$ if there is an *abstract weight function* $w : (T(\Sigma_s) \times T(\Sigma_s)) \rightarrow [0, 1]$ s.t. for all $t, t' \in T(\Sigma_s)$, (i) $w(t, T(\Sigma_s)) > 0$ iff $\pi(t) > 0$, (ii) $w(T(\Sigma_s), t') > 0$ iff $\pi'(t') > 0$, and (iii) $w(t, t') > 0$ implies $t R t'$.

Lemma 3. For all $t, t' \in T(\Sigma_s)$, $t \sim_a t'$ if and only if there is a symmetric relation $R \subseteq T(\Sigma_s) \times T(\Sigma_s)$ with $t R t'$ such that for all $t_1, t_2 \in T(\Sigma_s)$, $a \in A$, and $\pi_1 \in \Delta(T(\Sigma_s))$, $t_1 R t_2$ and $t_1 \xrightarrow{a} \pi_1$ imply that there exists $\pi_2 \in \Delta(T(\Sigma_s))$ s.t. $t_2 \xrightarrow{a} \pi_2$ and $\pi_1 \equiv_R^w \pi_2$.

The next lemma shows that the probability obliterated bisimulation agrees with Park & Milner's bisimulation. Denote $t \rightsquigarrow t'$ iff there is π such that $t \xrightarrow{a} \pi$ and $\pi(t') > 0$. Notice that this notation precisely defines the usual abstraction of probabilistic transition systems into labeled transition systems in which all information regarding the probability distribution is lost except from the fact that one state can reach another state with positive probability after a transition.

Lemma 4. For all $t, t' \in T(\Sigma_s)$, $t \sim_o t'$ iff there is a symmetric relation $R \subseteq T(\Sigma_s) \times T(\Sigma_s)$ with $t R t'$ s.t. for all $t_1, t_2, t'_1 \in T(\Sigma_s)$ and $a \in A$, $t_1 R t_2$ and $t_1 \rightsquigarrow t'_1$ imply that there exists $t'_2 \in T(\Sigma_s)$ s.t. $t_2 \rightsquigarrow t'_2$ and $t'_1 R t'_2$.

Finally we state the relation among the different bisimulations

Lemma 5. The following inclusions hold and are proper: $\sim \subsetneq \sim_c \subsetneq \sim_o$ and $\sim \subsetneq \sim_a \subsetneq \sim_o$. Besides \sim_c and \sim_a are incomparable.

In fact the results can be proved stronger as we explain in the following. Any bisimulation relation is also a convex bisimulation, which follow from the fact that $t \xrightarrow{a} \pi$ implies $t \xrightarrow{a}_c \pi$. Any convex bisimulation is also a probability obliterated bisimulation since $t \xrightarrow{a}_c \pi$ with $\pi(Q) > 0$ implies that there is a

π' such that $t \xrightarrow{a} \pi'$ and $\pi'(Q) > 0$. Any bisimulation is also a probability abstracted bisimulation since $\pi R \pi'$ implies $\pi(Q) > 0$ iff $\pi'(Q) > 0$ for all R -closed Q . Finally, any probability abstracted bisimulation is also a probability obliterated bisimulation since, for a given π and R , the existence of a π' s.t. $t' \xrightarrow{a} \pi'$ and $\pi(Q) > 0$ iff $\pi'(Q) > 0$ for all R -closed Q , guarantees that, for all R -closed Q with $\pi(Q) > 0$ there is a π' s.t. $t' \xrightarrow{a} \pi'$ and $\pi'(Q) > 0$.

Notice that $\mathbf{t}_1 = a.(b.\mathbf{0}) + a.(c.\mathbf{0})$ and $\mathbf{t}_2 = \mathbf{t}_1 + a.(b.\mathbf{0} \oplus_{0.5} c.\mathbf{0})$ are convex bisimilar but not probability abstracted bisimilar. Besides, notice that $\mathbf{t}_3 = a.(b.\mathbf{0} \oplus_{0.5} c.\mathbf{0})$ and $\mathbf{t}_4 = a.(b.\mathbf{0} \oplus_{0.1} c.\mathbf{0})$ are probability abstracted bisimilar but not convex bisimilar. These examples not only show that \sim_c and \sim_a are incomparable, but also that all stated inclusions are proper.

In the rest of the section we present logical characterizations for the different bisimulation equivalences. This work has already been done for bisimulation [10, 16] and convex bisimulation [16]. We adopt here the two-level logic style of [10].

We define the logic \mathcal{L}_b as the set of all formulas with the following syntax:

$$\phi := \top \mid \langle a \rangle \psi \mid \langle a \rangle_c \psi \mid \bigwedge_{i \in I} \phi_i \mid \neg \phi \qquad \psi := [\phi]_p \mid \bigcap_{i \in I} \psi_i$$

where $a \in A$, $p \in [0, 1] \cap \mathbb{Q}$, and I is any index set. The logic \mathcal{L}_c contains all formulas of \mathcal{L}_b without the modality $\langle a \rangle_-$. The logic \mathcal{L}_a contains all formulas of \mathcal{L}_b without modalities $\langle a \rangle_c$ - and $[_]_p$ for all $p > 0$ (i.e. it only accepts $[_]_0$ among this type of modalities.) Finally, the logic \mathcal{L}_o contains all formulas of \mathcal{L}_a without $\bigcap_{i \in I}$.

The semantics of \mathcal{L}_b is defined with the satisfaction relation \models on a PTS $P = (\mathcal{T}(\Sigma_s), A, \rightarrow)$ as follows.

- | | |
|--|--|
| (i) $t \models \top$ for all $t \in \mathcal{T}(\Sigma_s)$
(ii) $t \models \langle a \rangle \psi$ if there is $t \xrightarrow{a} \pi$ s.t. $\pi \models \psi$
(iii) $t \models \langle a \rangle_c \psi$ if there is $t \xrightarrow{a}_c \pi$ s.t. $\pi \models \psi$
(iv) $t \models \bigwedge_{i \in I} \phi_i$ if $t \models \phi_i$ for all $i \in I$ | (v) $t \models \neg \phi$ if $t \not\models \phi$
(vi) $\pi \models [\phi]_p$ if $\pi(\{t \in \mathcal{T}(\Sigma_s) \mid t \models \phi\}) > p$
(vii) $\pi \models \bigcap_{i \in I} \psi_i$ if $\pi \models \psi_i$ for all $i \in I$ |
|--|--|

The semantics of the other logics is defined in the same way but restricted to the respective operators.

For $\chi \in \{b, c, a, o\}$, let $\mathcal{L}_\chi(t) = \{\phi \in \mathcal{L}_\chi \mid t \models \phi\}$, for all $t \in \mathcal{T}(\Sigma_s)$, and $\mathcal{L}_\chi(\pi) = \{\psi \in \mathcal{L}_\chi \mid \pi \models \psi\}$, for all $\pi \in \Delta(\mathcal{T}(\Sigma_s))$. We write $t_1 \sim_{\mathcal{L}_\chi} t_2$ iff $\mathcal{L}_\chi(t_1) = \mathcal{L}_\chi(t_2)$ and $\pi_1 \sim_{\mathcal{L}_\chi} \pi_2$ iff $\mathcal{L}_\chi(\pi_1) = \mathcal{L}_\chi(\pi_2)$. Then, we have the following characterization theorem.

Theorem 1. *For all $\chi \in \{b, c, a, o\}$ and for all $t_1, t_2 \in \mathcal{T}(\Sigma_s)$, $t_1 \sim_\chi t_2$ iff $t_1 \sim_{\mathcal{L}_\chi} t_2$ (where $\sim_b = \sim$).*

Let $\mathbf{t}_1, \mathbf{t}_2, \mathbf{t}_3$, and \mathbf{t}_4 be as before. Recall $\mathbf{t}_1 \sim_c \mathbf{t}_2$ and $\mathbf{t}_3 \sim_a \mathbf{t}_4$. Notice that $\langle a \rangle([\langle b \rangle \top]_{0.5} \sqcap [\langle c \rangle \top]_{0.5})$ distinguish \mathbf{t}_1 from \mathbf{t}_2 , while $\langle a \rangle_c([\langle b \rangle \top]_{0.5} \sqcap [\langle c \rangle \top]_{0.5})$ is satisfied by both \mathbf{t}_1 and \mathbf{t}_2 . That is why $\langle a \rangle_-$ is not an operator of \mathcal{L}_c . Notice $[\langle b \rangle \top]_{0.5}$ distinguishes the distribution $\llbracket b.\mathbf{0} \oplus_{0.5} c.\mathbf{0} \rrbracket$ from $\llbracket b.\mathbf{0} \oplus_{0.1} c.\mathbf{0} \rrbracket$, while $[\langle b \rangle \top]_0$ does not (but it does distinguish them from e.g. $\llbracket c.\mathbf{0} \rrbracket$). Thus $\langle a \rangle[\langle b \rangle \top]_{0.5}$ distinguishes \mathbf{t}_3 from \mathbf{t}_4 . That is why $[_]_p$ is not an operator of \mathcal{L}_a if $p > 0$. Finally, notice that $\langle a \rangle([\langle b \rangle \top]_0 \sqcap [\langle c \rangle \top]_0)$ distinguishes $\mathbf{t}_5 = a.(b.\mathbf{0} \oplus_{0.5} c.\mathbf{0})$ from $\mathbf{t}_6 = a.b.\mathbf{0} + a.c.\mathbf{0}$, and observe that $\mathbf{t}_5 \sim_o \mathbf{t}_6$. However, neither $\langle a \rangle[\langle b \rangle \top]_0$ nor $\langle a \rangle[\langle c \rangle \top]_0$ can distinguish them. That is why $\bigcap_{i \in I}$ is not an operator of \mathcal{L}_o .

5 Formats

In this section we introduce rule and specification formats that guarantee that each bisimulation equivalences discussed in the previous section is a congruence for every operator whose semantics is defined

within the respective rule of the specification format. In particular, the format $nt\mu f\theta/nt\mu x\theta$, which ensures that bisimulation equivalence is a congruence for all operator in such format, has been already introduced in [9] and finally revised in [8]. We present here its more general form.

The following definition is important to ensure a symmetric treatment of variables and terms within the format. Let $\{Y_l\}_{l \in L}$ be a family of sets of state term variables with the same cardinality. The l -th element of a tuple \vec{y} is denoted by $\vec{y}(l)$. For a set of tuples $T = \{\vec{y}_i \mid i \in I\}$ we denote the l -th projection by $\Pi_l(T) = \{\vec{y}_i(l) \mid i \in I\}$. Fix a set $\text{Diag}\{Y_l\}_{l \in L} \subseteq \prod_{l \in L} Y_l$ such that: (i) for all $l \in L$, $\Pi_l(\text{Diag}\{Y_l\}_{l \in L}) = Y_l$; and (ii) for all $\vec{y}, \vec{y}' \in \text{Diag}\{Y_l\}_{l \in L}$, $(\exists l \in L : \vec{y}(l) = \vec{y}'(l)) \Rightarrow \vec{y} = \vec{y}'$. Property (ii) ensures that different tuples $\vec{y}, \vec{y}' \in \text{Diag}\{Y_l\}_{l \in L}$ differ in all positions, and by property (i) every variable of every Y_l is used in (exactly) one $\vec{y} \in \text{Diag}\{Y_l\}_{l \in L}$. Diag stands for “diagonal”, following the intuition that each \vec{y} represents a coordinate in the space $\prod_{l \in L} Y_l$, so that $\text{Diag}\{Y_l\}_{l \in L}$ can be seen as the line that traverses the main diagonal of the space. Therefore, notice that, for $Y_l = \{y_l^0, y_l^1, y_l^2, \dots\}$, a possible definition for the set $\text{Diag}\{Y_l\}_{l \in L}$ is $\{(y_1^0, y_2^0, \dots, y_L^0), (y_1^1, y_2^1, \dots, y_L^1), (y_1^2, y_2^2, \dots, y_L^2), \dots\}$. In addition, we use the following notation: $t(\zeta_1, \dots, \zeta_n)$ denotes a term that only has variables in the set $\{\zeta_1, \dots, \zeta_n\}$, that is $\mathcal{V}(t(\zeta_1, \dots, \zeta_n)) \subseteq \{\zeta_1, \dots, \zeta_n\}$, and moreover, $t(\zeta'_1, \dots, \zeta'_n)$ denotes the same term as $t(\zeta_1, \dots, \zeta_n)$ in which each variable ζ_i has been replaced by ζ'_i .

Definition 11. Let $P = (\Sigma, A, R)$ be a PTSS. A rule $r \in R$ is in $nt\mu f\theta$ format if it has the following form

$$\frac{\bigcup_{m \in M} \{t_m(\vec{z}) \xrightarrow{a_m} \mu_m^{\vec{z}} \mid \vec{z} \in \mathcal{Z}\} \cup \bigcup_{n \in N} \{t_n(\vec{z}) \xrightarrow{b_n} \mid \vec{z} \in \mathcal{Z}\} \cup \{\theta_l(Y_l) \triangleright_{l,k} p_{l,k} \mid l \in L, k \in K_l\}}{f(\zeta_1, \dots, \zeta_{rk(f)}) \xrightarrow{a} \theta}$$

with $\triangleright_{l,k} \in \{>, \geq\}$ for all $l \in L$ and $k \in K_l$, and $\mathcal{Z} = \text{Diag}\{Y_l\}_{l \in L} \times \prod_{\zeta \in W} \{\zeta\}$, with $W \subseteq \mathcal{V} \cup \mathcal{V}_d \setminus \bigcup_{l \in L} Y_l$. In addition, it has to satisfy the following conditions:

1. Each set Y_l should be at least countably infinite, for all $l \in L$, and the cardinality of L should be strictly smaller than that of the Y_l 's.
2. All variables $\zeta_1, \dots, \zeta_{rk(f)}$ are different.
3. All variables $\mu_m^{\vec{z}}$, with $m \in M$ and $\vec{z} \in \mathcal{Z}$, are different and $\{\zeta_1, \dots, \zeta_{rk(f)}\} \cap \{\mu_m^{\vec{z}} \mid \vec{z} \in \mathcal{Z}, m \in M\} = \emptyset$.
4. For all $l \in L$, $Y_l \cap \{\zeta_1, \dots, \zeta_{rk(f)}\} = \emptyset$, and $Y_l \cap Y_{l'} = \emptyset$ for all $l' \in L$, $l \neq l'$.
5. For all $m \in M$, the set $\{\mu_m^{\vec{z}} \mid \vec{z} \in \mathcal{Z}\} \cap (\mathcal{V}(\theta) \cup (\bigcup_{l \in L} \mathcal{V}(\theta_l)) \cup W)$ is finite.
6. For all $l \in L$, the set $Y_l \cap (\mathcal{V}(\theta) \cup \bigcup_{l' \in L} \mathcal{V}(\theta_{l'}))$ is finite.

A rule $r \in R$ is in $nt\mu x\theta$ format if its form is like above but has a conclusion of the form $x \xrightarrow{a} \theta$ and, in addition, it satisfies the same conditions as above, except that whenever we write $\{\zeta_1, \dots, \zeta_{rk(f)}\}$, we should write $\{x\}$. P is in $nt\mu f\theta$ format if all its rules are in $nt\mu f\theta$ format. P is in $nt\mu f\theta/nt\mu x\theta$ format if all its rules are in either $nt\mu f\theta$ format or $nt\mu x\theta$ format.

The rationale behind each of the restrictions are discussed in [8] in depth (see also [9]). In the following we briefly summarize it. Variables $\zeta_1, \dots, \zeta_{rk(f)}$ in the source of the conclusion, all variables $\mu_m^{\vec{z}}$ in the target of the positive premises, and all variables in the sets Y_l , $l \in L$, as part of the measurable sets in the quantitative premises, are binding. That is why all of them are requested to be different, which is stated in conditions 2, 3, and 4. If Y_l is finite, quantitative premises will allow to count the minimum number of terms that gather certain probabilities. This goes against the spirit of bisimulation that measures equivalence classes of terms regardless of the size of them. Therefore Y_l needs to be infinite (condition 1). Condition 5 ensures that, for each $m \in M$ there are sufficiently many distribution variables in the set $\{\mu_m^{\vec{z}} \mid \vec{z} \in \mathcal{Z}\}$ to be freely instantiated. The use of a distribution variable in a quantitative premise may disclose part of the structural nature of the distribution term that substitutes such variable. Thus, for instance, if all variables $\mu_m^{\vec{z}}$ are used in different quantitative premises together with some lookahead, we

may restrict the syntactic form of the eventually substituted distribution terms, hence revealing unwanted differences. A similar situation arises with the use of variables in Y_l for all $l \in L$, hence condition 6. The precise understanding of conditions 5 and 6 requires a rather lengthy explanation that is beyond the scope of this paper. The reader is referred to [8, 9] for details.

All congruence theorems in this article apply only to PTSSs whose rules are *well-founded*. A rule r is *well-founded* if there is no infinite backward chain in the *dependency graph* $G_r = (V, E)$ of r defined by $V = \mathcal{V} \cup \mathcal{V}_d$ and $E = \{\langle \zeta, \mu \rangle \mid (t \xrightarrow{a} \mu) \in \text{pprem}(r), \zeta \in \mathcal{V}(t)\} \cup \{\langle \zeta, y \rangle \mid (\theta(Y) \triangleright p) \in \text{qpem}(r), \zeta \in \mathcal{V}(\theta), y \in Y\}$. A PTSS is called *well-founded* if all its rules are well-founded.

The full proof of the following theorem can be found in [8].

Theorem 2. *Let $P = (\Sigma, A, R)$ be a complete well-founded PTSS in $nt\mu f\theta/nt\mu x\theta$ format. Then, the bisimulation equivalence is a congruence for all operators defined in P .*

The $nt\mu f\theta/nt\mu x\theta$ format is still too general to preserve the other (weaker) bisimulation equivalences presented in Sec. 4. In the remainder of the section, we will discuss through appropriate examples how the $nt\mu f\theta/nt\mu x\theta$ format should be further restricted or modified so that the other bisimulation equivalences are congruences for the resulting restricted formats.

We focus first on convex bisimulation. For this consider the terms $\mathbf{t}_1 = a.(\mathbf{b}.0) + a.(\mathbf{c}.0)$ and $\mathbf{t}_2 = \mathbf{t}_1 + a.(\mathbf{b}.0 \oplus_{0.5} \mathbf{c}.0)$. Notice that $\mathbf{t}_1 \sim_c \mathbf{t}_2$. Consider a possible extension of our running example with a unary operator f with the following $nt\mu f\theta$ rule:

$$\frac{x \xrightarrow{a} \mu \quad \mu(Y) \geq 0.5 \quad \{y \xrightarrow{b} \mu_y \mid y \in Y\} \quad \mu(Y') \geq 0.5 \quad \{y' \xrightarrow{c} \mu_{y'} \mid y' \in Y'\}}{f(x) \xrightarrow{a} \mathbf{0}} \quad (1)$$

Since $\mathbf{t}_2 \xrightarrow{a} (\mathbf{b}.0 \oplus_{0.5} \mathbf{c}.0)$, $f(\mathbf{t}_2) \xrightarrow{a} \mathbf{0}$. However it is easy to see that $f(\mathbf{t}_1)$ cannot perform any transition. Therefore $f(\mathbf{t}_1) \not\sim_c f(\mathbf{t}_2)$.

The problem arises precisely because, in order to show that $\mathbf{t}_1 \sim_c \mathbf{t}_2$, transition $\mathbf{t}_2 \xrightarrow{a} (\mathbf{b}.0 \oplus_{0.5} \mathbf{c}.0)$ is matched with the appropriate convex combination of the transitions $\mathbf{t}_1 \xrightarrow{a} \mathbf{b}.0$ and $\mathbf{t}_1 \xrightarrow{a} \mathbf{c}.0$. Thus, we need that a quantitative premise guarantees that the test is produced on a convex combination of target distributions rather than on a single target distribution. An appropriate modification of such rule would be to replace it by a family of rules of the form

$$\frac{\{x \xrightarrow{a} \mu_n \mid n \in \mathbb{N}\} \quad \left(\bigoplus_{n \in \mathbb{N}} [p_n] \mu_n\right)(Y) \geq 0.5 \quad \{y \xrightarrow{b} \mu_y \mid y \in Y\} \quad \left(\bigoplus_{n \in \mathbb{N}} [p_n] \mu_n\right)(Y') \geq 0.5 \quad \{y' \xrightarrow{c} \mu_{y'} \mid y' \in Y'\}}{f(x) \xrightarrow{a} \mathbf{0}},$$

one for each $\{p_n\}_{n \in \mathbb{N}}$ such that $\sum_{n \in \mathbb{N}} p_n = 1$ and each $p_i \in [0, 1] \cap \mathbb{Q}$.

Consider now that the semantic of f is defined by the rule

$$\frac{x \xrightarrow{a} \mu}{f(x) \xrightarrow{a} a.\mu} \quad (2)$$

and notice that $f(\mathbf{t}_2) \xrightarrow{a} a.(\mathbf{b}.0 \oplus_{0.5} \mathbf{c}.0)$. However, the only two possible transitions for $f(\mathbf{t}_1)$ are $f(\mathbf{t}_1) \xrightarrow{a} a.\mathbf{b}.0$ and $f(\mathbf{t}_1) \xrightarrow{a} a.\mathbf{c}.0$, and there is no $p \in [0, 1]$ such that $a.\mathbf{b}.0 \oplus_p a.\mathbf{c}.0 \sim_c a.(\mathbf{b}.0 \oplus_{0.5} \mathbf{c}.0)$. For this reason, we will require that a target of a positive premise does not appear in a d -sorted position of a subterm in the target of the conclusion.

For the next example, we consider an additional unary d -sorted operator g and the following rules

$$\frac{x \xrightarrow{a} \mu \quad g(\mu) \xrightarrow{b} \mu'}{f(x) \xrightarrow{a} \mathbf{0}} \quad \frac{\mu(Y) > 0 \quad \{y \xrightarrow{b} \mu \mid y \in Y\} \quad \mu(Y') > 0 \quad \{y' \xrightarrow{c} \mu' \mid y' \in Y'\}}{g(\mu) \xrightarrow{b} \mathbf{0}} \quad (3)$$

Notice that $g(\mathbf{b.0} \oplus_{0.5} \mathbf{c.0}) \xrightarrow{b} \mathbf{0}$. Therefore $f(\mathbf{t}_2) \xrightarrow{a} \mathbf{0}$. However, neither $g(\mathbf{b.0})$ nor $g(\mathbf{c.0})$ can perform any transition, and as a consequence $f(\mathbf{t}_1)$ cannot perform any transition either. Hence, $f(\mathbf{t}_1) \not\prec_c f(\mathbf{t}_2)$. For this reason we will require that a target of a positive premise does not appear in the source of a positive or negative premise.

Suppose now that g is a binary s -sorted operator and consider the following rules

$$\frac{x \xrightarrow{a} \mu}{f(x) \xrightarrow{a} g(\mu, \mu)} \quad \frac{x_1 \xrightarrow{b} \mu_1 \quad x_2 \xrightarrow{c} \mu_2}{g(x_1, x_2) \xrightarrow{a} \mathbf{0}} \quad (4)$$

Notice that the only possible transitions for $f(\mathbf{t}_1)$ are $f(\mathbf{t}_1) \xrightarrow{a} g(\mathbf{b.0}, \mathbf{b.0})$ and $f(\mathbf{t}_1) \xrightarrow{a} g(\mathbf{c.0}, \mathbf{c.0})$. Moreover, notice that $g(\mathbf{b.0}, \mathbf{b.0}) \sim_c g(\mathbf{c.0}, \mathbf{c.0}) \sim_c \mathbf{0}$. However, $f(\mathbf{t}_2) \xrightarrow{a} g(\mathbf{b.0} \oplus_{0.5} \mathbf{c.0}, \mathbf{b.0} \oplus_{0.5} \mathbf{c.0})$, and it is not difficult to see that $g(\mathbf{b.0} \oplus_{0.5} \mathbf{c.0}, \mathbf{b.0} \oplus_{0.5} \mathbf{c.0}) \sim_c (a.0 \oplus_{0.25} \mathbf{0})$. Therefore, $f(\mathbf{t}_1) \not\prec_c f(\mathbf{t}_2)$. In this case, the problem seems to arise because the same distribution variable occurs in the target of the conclusion of the first rule in two different s -sorts positions of the target distribution term. However, the problem is not so general. Notice that if the target in the conclusion is replaced by the term $g(\mu, \mathbf{c.0}) \oplus_p g(\mathbf{b.0}, \mu)$ we would have $f(\mathbf{t}_1) \sim_c f(\mathbf{t}_2)$. The difference arises from the fact that in the interpretation of $g(\theta, \theta)$ the probability distribution $\llbracket \theta \rrbracket$ multiplies with itself. This is not the case in the interpretation of $g(\theta, \mathbf{c.0}) \oplus_p g(\mathbf{b.0}, \theta)$ where the two instances of $\llbracket \theta \rrbracket$ are summed up. Thus, we will actually request that the target of the conclusion is *linear* with respect to each distribution variable on a target of a positive premise.

Definition 12. A distribution term $\theta \in \mathbb{T}(\Sigma_d)$ is linear for a set $V \subseteq \mathcal{V}_d$ if (i) $\theta \in T(\Sigma_d) \cup \mathcal{V}_d \cup \{\delta(x) \mid x \in \mathcal{V}\}$. (ii) $\theta = \bigoplus_{i \in I} [p_i] \theta_i$ and θ_i is linear for V , for all $i \in I$, (iii) $\theta = f(\theta_1, \dots, \theta_n)$, for all $i \in I$, θ_i is linear for V , and $\mathcal{V}(\theta_i) \cap \mathcal{V}(\theta_j) \cap V = \emptyset$, for all $i, j \in \{1, \dots, n\}$ and $i \neq j$.

Definition 13. Let $P = (\Sigma, A, R)$ be a PTSS. A rule $r \in R$ is in convex $nt\mu f\theta$ format if has the form

$$\frac{\bigcup_{m \in M} \{t_m(\vec{z}) \xrightarrow{a_m} \mu_m^{\vec{z}} \mid \vec{z} \in \mathcal{Z}\} \quad \bigcup_{n \in N} \{t_n(\vec{z}) \xrightarrow{b_n} \mid \vec{z} \in \mathcal{Z}\} \quad \bigcup_{\tilde{m} \in \tilde{M}} \{t_{\tilde{m}}(\vec{z}_{\tilde{m}}) \xrightarrow{a_{\tilde{m}}} \mu_{\tilde{m}}^{\vec{z}_{\tilde{m}}} \mid i \in \mathbb{N}\} \quad \bigcup_{\tilde{m} \in \tilde{M}} \{(\bigoplus_{i \in \mathbb{N}} [p_i^{\tilde{m}}] \mu_i^{\tilde{m}})(Y_l) \triangleright_{l,k} p_{l,k} \mid l \in L_{\tilde{m}}, k \in K_l\}}}{f(\zeta_1, \dots, \zeta_{rk(f)}) \xrightarrow{a} \theta}$$

with $L = \bigcup_{\tilde{m} \in \tilde{M}} L_{\tilde{m}}$, $L_{\tilde{m}} \cap L_{\tilde{m}'} = \emptyset$ whenever $\tilde{m} \neq \tilde{m}'$, $\triangleright_{l,k} \in \{>, \geq\}$ for all $l \in L$ and $k \in K_l$, $\mathcal{Z} = \text{Diag}\{Y_l\}_{l \in L} \times \prod_{\zeta \in W} \{\zeta\}$, with $W \subseteq \mathcal{V} \cup \mathcal{V}_d \setminus \bigcup_{l \in L} Y_l$. In addition, it should also satisfy conditions 1 to 6 in Def. 11 and the following extra conditions:

7. For every $\tilde{m} \in \tilde{M}$, the family $\{p_i^{\tilde{m}}\}_{i \in \mathbb{N}} \subseteq [0, 1] \cap \mathbb{Q}$ and $\sum_{i \in \mathbb{N}} p_i^{\tilde{m}} = 1$
8. For every $\tilde{m} \in \tilde{M}$, there is exactly one $j \in \mathbb{N}$ such that $\mu_j^{\tilde{m}} = \mu_m^{\vec{z}}$ for some $m \in M$ and $\vec{z} \in \mathcal{Z}$, in which case also $t_{\tilde{m}}(\vec{z}_{\tilde{m}}) \xrightarrow{a_{\tilde{m}}} \mu_j^{\tilde{m}} = t_m(\vec{z}) \xrightarrow{a_m} \mu_m^{\vec{z}}$. Moreover $\{\mu_i^{\tilde{m}} \mid i \in \mathbb{N}\} \cap \{\mu_i^{\tilde{m}'} \mid i \in \mathbb{N}\} = \emptyset$ for all $\tilde{m} \neq \tilde{m}'$, and $\{\mu_i^{\tilde{m}} \mid i \in \mathbb{N}\} \cap \{\zeta_1, \dots, \zeta_{rk(f)}\} = \emptyset$.
9. No variable $\mu_m^{\vec{z}}$, with $m \in M$ and $\vec{z} \in \mathcal{Z}$, appears in the source of a premise (i.e. in the set W) or in a d -sorted position of a subterm in the target of the conclusion θ .
10. θ is linear for $\{\mu_m^{\vec{z}} \mid m \in M, \vec{z} \in \mathcal{Z}\}$.

A rule $r \in R$ is in convex $nt\mu x\theta$ format if its form is like above but has a conclusion of the form $x \xrightarrow{a} \theta$ and it satisfies the same conditions, except that whenever we write $\{\zeta_1, \dots, \zeta_{rk(f)}\}$, we should write $\{x\}$. A set of convex $nt\mu f\theta/nt\mu x\theta$ rules R is convex closed if for all $r \in R$, for any term $\bigoplus_{i \in \mathbb{N}} [p_i^{\tilde{m}}] \mu_i^{\tilde{m}}$ appearing in a quantitative premise of r and any family $\{q_i\}_{i \in \mathbb{N}} \subseteq [0, 1] \cap \mathbb{Q}$ such that $\sum_{i \in \mathbb{N}} q_i = 1$, then the rule r' obtained by replacing each occurrence of $\bigoplus_{i \in \mathbb{N}} [p_i^{\tilde{m}}] \mu_i^{\tilde{m}}$ in r by $\bigoplus_{i \in \mathbb{N}} [q_i] \mu_i^{\tilde{m}}$ is also in R . A PTSS $P = (\Sigma, A, R)$ is in convex $nt\mu f\theta/nt\mu x\theta$ format if all rules in R are in convex $nt\mu f\theta/nt\mu x\theta$ format and R is convex closed.

The problem indicated in rule (1) is attacked with the requirement of having sets $\{t_{\vec{m}}(\vec{z}_{\vec{m}}) \xrightarrow{a_{\vec{m}}} \mu_i^{\vec{m}} \mid i \in \mathbb{N}\}$ as positive premises with which the convex closures $\bigoplus_{i \in \mathbb{N}} [p_i^{\vec{m}}] \mu_i^{\vec{m}}$ can be constructed, plus the request that the set of rules is convex closed. Notice that condition 8 states that these sets of positive premises are only used to construct such distribution terms and are only linked to the “actual” positive premises in $\bigcup_{m \in M} \{t_m(\vec{z}) \xrightarrow{a_m} \mu_{\vec{m}}^{\vec{z}} \mid \vec{z} \in \mathcal{Z}\}$ through a single transition $t_{\vec{m}}(\vec{z}_{\vec{m}}) \xrightarrow{a_{\vec{m}}} \mu_j^{\vec{m}}$.

Rules like (2) and the left rule on (3) are excluded on condition 9 since no variable of a positive premise can be used in the source of a premise (excluding (3)) or in a d -sort position in the target of the conclusion (excluding (2)). Finally, rules like on the left of (4) are excluded by requesting that the target of the conclusion is linear (condition 10).

Now, we can state the congruence theorem for convex bisimulation equivalence.

Theorem 3. *Let P be a complete well-founded PTSS in convex $nt\mu f\theta/nt\mu x\theta$ format. Then, convex bisimulation equivalence is a congruence for all operators defined by P .*

We focus now on the probability abstracted bisimulation. Notice that the terms $\mathbf{t}_3 = a.(\mathbf{b.0} \oplus_{0.5} \mathbf{c.0})$ and $\mathbf{t}_4 = a.(\mathbf{b.0} \oplus_{0.1} \mathbf{c.0})$ are probability abstracted bisimilar, i.e., $\mathbf{t}_3 \sim_a \mathbf{t}_4$. Consider now the unary operator f whose semantics is defined with rule (1). It should not be difficult so see that $f(\mathbf{t}_3) \xrightarrow{a} \mathbf{0}$ while $f(\mathbf{t}_4)$ cannot perform any transition. Therefore $f(\mathbf{t}_3) \not\sim_a f(\mathbf{t}_4)$. The problem is a consequence of the fact that the quantitative premises are tested against non-zero values which may distinguish distributions with the same support set but mapping into different probability values. Thus, in order to preserve probability abstracted bisimulation equivalence, the only extra restriction that we ask to a rule in $nt\mu f\theta/nt\mu x\theta$ format is that none of its quantitative premises test against a value different from 0.

Definition 14. *A PTSS $P = \langle \Sigma, A, R \rangle$ is in probability abstracted $nt\mu f\theta/nt\mu x\theta$ format if it is in $nt\mu f\theta/nt\mu x\theta$ format and for every rule $r \in R$ and quantitative premise $\theta(Y) \geq p \in qpem(r)$, $p = 0$.*

The proof of the congruence theorem for probability abstracted bisimulation equivalence (Theorem 4 below) follows closely the lines of the proof of Theorem 2 as given in [8].

Theorem 4. *Let P be a complete well-founded PTSS in probability abstracted $nt\mu f\theta/nt\mu x\theta$ format. Then, the probability abstracted bisimulation equivalence is a congruence for all operators defined in P .*

Given the alternative definition of the probability obliterated bisimulation provided by Lemma 4, we will now consider simpler definitions for the quantitative premises for the rule format associated to this relation. Thus, we consider quantitative premises of the form $\theta(\{y\}) \geq p$ rather than $\theta(Y) \geq p$.

Taking \mathbf{t}_3 and \mathbf{t}_4 as before, we have that $\mathbf{t}_3 \sim_o \mathbf{t}_4$. The same example of the unary operator f , whose semantics is defined with a conveniently modify rule (1), shows that $f(\mathbf{t}_3) \not\sim_a f(\mathbf{t}_4)$ and hence the need that the quantitative premises can only be tested against 0.

Let $\mathbf{t}_5 = a.(\mathbf{b.0} \oplus_{0.5} \mathbf{c.0})$ and $\mathbf{t}_6 = a.\mathbf{b.0} + a.\mathbf{c.0}$, and observe that $\mathbf{t}_5 \sim_o \mathbf{t}_6$. Take rule (2) as the semantic definition for f . Notice that $f(\mathbf{t}_5) \xrightarrow{a} a.(\mathbf{b.0} \oplus_{0.5} \mathbf{c.0})$ is the only transition for $f(\mathbf{t}_5)$, while the only possible transitions for $f(\mathbf{t}_6)$ are $f(\mathbf{t}_6) \xrightarrow{a} a.\mathbf{b.0}$ and $f(\mathbf{t}_6) \xrightarrow{a} a.\mathbf{c.0}$. Since $a.\mathbf{b.0} \not\sim_o a.(\mathbf{b.0} \oplus_{0.5} \mathbf{c.0}) \not\sim_o a.\mathbf{c.0}$, $f(\mathbf{t}_5) \not\sim_o f(\mathbf{t}_6)$. Like for the convex bisimulation case, this shows that the target of a positive premise cannot appear in a d -sorted position of a subterm in the target of the conclusion.

Suppose now that the semantics of f is defined with the rule

$$\frac{x \xrightarrow{a} \mu \quad \mu(\{y_1\}) > 0 \quad \mu(\{y_2\}) > 0 \quad y_1 \xrightarrow{b} \mu_1 \quad y_2 \xrightarrow{c} \mu_2}{f(x) \xrightarrow{a} \mathbf{0}} \quad (5)$$

Notice that $f(\mathbf{t}_5) \not\sim_o f(\mathbf{t}_6)$ since $f(\mathbf{t}_5) \xrightarrow{a} \mathbf{0}$ while $f(\mathbf{t}_6)$ cannot perform any transition. This is due to the fact that, by allowing the same distribution variable μ to occur in different quantitative premises, we gain some knowledge of the structure of (the instance of) μ , in particular of its support set.

Consider now that f is defined with the left rule in (3) and g with an appropriate modification of the right rule in (3). Notice that $f(t_5) \xrightarrow{a} \mathbf{0}$ but $f(t_6)$ cannot perform any transition. Thus $f(t_5) \not\sim_o f(t_6)$. In this case, we are also gaining knowledge of the support set of μ , but this time through the rule associated to the operator g . Therefore we require that a target of a positive premise does not appear in the source of a positive or negative premise.

Consider now the rules

$$\frac{x \xrightarrow{a} \mu \quad \mu(\{y\}) > 0 \quad y \xrightarrow{b} \mu'}{f(x) \xrightarrow{a} g(\mu)} \quad \frac{x \xrightarrow{c} \mu}{g(x) \xrightarrow{c} \mathbf{0}} \quad (6)$$

Notice that the only transition for $f(t_5)$ is $f(t_5) \xrightarrow{a} g(\mathbf{b}.\mathbf{0} \oplus_{0.5} \mathbf{c}.\mathbf{0})$ and the only transition for $f(t_6)$ is $f(t_6) \xrightarrow{a} g(\mathbf{b}.\mathbf{0})$. Then $f(t_5) \xrightarrow{a} g(\mathbf{c}.\mathbf{0}) \xrightarrow{c} \mathbf{0}$ while $f(t_6) \xrightarrow{a} g(\mathbf{b}.\mathbf{0})$ is the only possible ‘‘obliterated’’ transition for $f(t_6)$. Then $f(t_5) \not\sim_o f(t_6)$. This is an alternative way of gaining information on the support set of a possible instance of μ in (6): on the one hand, by the quantitative premise on the first rule, we deduce that such instance has an element in the support set that performs a b -transition and, on the other hand, by having μ as an argument in the target of the conclusion, we may gather extra information from the same instance of μ through the rules for the semantics of the target of the conclusion (in this case, that μ has another element in the support set that performs a c -transition.) Therefore, we forbid that the target of a positive premise is both tested in a quantitative premise and used in the target of the conclusion.

Notice that the example in rules (4) also apply for probability obliterated bisimulation since $t_1 \sim_o t_2$ but $f(t_1) \not\sim_o f(t_2)$ with exactly the same explanation. Thus, we also request that the target of the conclusion is linear for all distribution variables on targets of positive premises.

Finally, consider a modification (4) where the left rule is instead

$$\frac{x \xrightarrow{a} \mu \quad g(\mu, \mu)(\{y\}) > 0 \quad y \xrightarrow{a} \mu'}{f(x) \xrightarrow{a} \mathbf{0}} \quad (7)$$

It should not be difficult to observe that $f(t_5) \xrightarrow{a} \mathbf{0}$ but $f(t_6)$ cannot perform any transition. Thus $f(t_5) \not\sim_o f(t_6)$. For this reason we also require that the quantitative premises only allow linear distribution terms.

Definition 15. Let $P = (\Sigma, A, R)$ be a well-founded PTSS. A rule $r \in R$ is in probability obliterated $nt\mu f\theta$ format if it has the form

$$\frac{\bigcup_{m \in M} \{t_m \xrightarrow{a_m} \mu_m\} \quad \cup \quad \bigcup_{n \in N} \{t_n \xrightarrow{b_n} \mu_n\} \quad \cup \quad \bigcup_{l \in L} \{\theta_l(\{y_l\}) > 0\}}{f(\zeta_1, \dots, \zeta_{rk(f)}) \xrightarrow{a} \theta}$$

where all variables $\zeta_1, \dots, \zeta_{rk(f)}$, μ_m , with $m \in M$, and y_l , with $l \in L$, are different and the following restrictions are satisfied:

1. For all $m \in M$, $\mathcal{V}(t_m) \cap \{\mu_{m'} \mid m' \in M\} = \emptyset$. Similarly, for all $n \in N$, $\mathcal{V}(t_n) \cap \{\mu_{m'} \mid m' \in M\} = \emptyset$.
2. For all $l \in L$, θ_l is linear for $\{\mu_{m'} \mid m' \in M\}$ and, moreover, for all $l, l' \in L$ with $l \neq l'$, $\mathcal{V}(\theta_l) \cap \mathcal{V}(\theta_{l'}) \cap \{\mu_m \mid m \in M\} = \emptyset$.
3. θ is linear for $\{\mu_{m'} \mid m' \in M\}$, $\mathcal{V}(\theta) \cap (\bigcup_{l \in L} \mathcal{V}(\theta_l)) \cap \{\mu_m \mid m \in M\} = \emptyset$, and no variable μ_m appear in a d -sorted position of a subterm of the target of the conclusion θ .

A rule is in probability obliterated $nt\mu x\theta$ format if its form is like above but has a conclusion of the form $x \xrightarrow{a} \theta$. P is in probability obliterated $nt\mu f\theta/nt\mu x\theta$ format if all its rules are in probability obliterated $nt\mu f\theta/nt\mu x\theta$ format.

Condition 1 limits the form to exclude rules like the one on the left of (3). Condition 2 requires that the distribution terms on the quantitative premises are linear (excluding (6)), and that they do not share distributions variables on the target of positive premises (excluding (5)). Finally, condition 3 request that the target of the conclusion is linear (excluding (4)) and does not have targets of positive premises on d -sorted positions (excluding (2)) nor if they are used in quantitative premises (excluding (6)).

Finally, we state the congruence theorem for probability obliterated bisimulation equivalence.

Theorem 5. *Let P be a complete well-founded PTSS in probability obliterated $nt\mu f\theta/nt\mu x\theta$ format. Then, probability obliterated bisimulation equivalence is a congruence for all operators defined by P .*

6 Conclusion and Future Work

In this article, we presented three new rule formats that preserve three different bisimulation equivalences coarser than Larsen & Skou's bisimulation. These formats are more restricted variants of the $nt\mu f\theta/nt\mu x\theta$ format and notably, all of them can be seen as generalizations of the non-probabilistic $ntyf\theta/ntyxt$ format [7, 14]. For completeness we mention two other similar results on PTSSs that fall out of Larsen & Skou's bisimulation. They are [20], that presents a format for rooted branching bisimulation, and [26], that presents a format for non-expansiveness of ϵ -bisimulations.

Prior to the congruence theorems, we presented the different bisimulation equivalences, compare them, and, in particular, we gave a logic characterization for each of them. The intention of presenting these logic characterizations is to use them as the basis for the proof of full abstraction theorems (see, e.g., [8, 9, 14, 15].) Full abstraction theorems are somewhat dual to the congruence theorems. An equivalence relation is fully abstract with respect to a particular format and an equivalence relation \equiv if it is the largest relation included in \equiv that is a congruence for all operators definable by any PTSS in that format. In particular we are interested when \equiv is the coarsest reasonable behavioral equivalence, namely, (possibilistic) trace equivalence. We are busy now on trying to prove this results for the formats presented here using the logic characterization as a means to construct the so called *testers*. As the current point of our investigation, we do not foresee major problems for all relations except for convex bisimulation equivalences, for which we may need to relax some of the conditions of the convex $nt\mu f\theta/nt\mu x\theta$ format.

References

- [1] L. Aceto, W. Fokkink & C. Verhoef (2001): *Conservative Extension in Structural Operational Semantics*. In: *Current Trends in Theoretical Computer Science*, pp. 504–524, doi:10.1142/9789812810403_0004.
- [2] L. Aceto, W. Fokkink & C. Verhoef (2001): *Structural operational semantics*. In: *Handbook of Process Algebra*, Elsevier, pp. 197–292, doi:10.1016/B978-044482830-9/50021-7.
- [3] C. Baier (1998): *On Algorithmics Verification Methods for Probabilistic Systems*. Habilitation thesis, University of Mannheim.
- [4] F. Bartels (2002): *GSOS for Probabilistic Transition Systems*. *Electr. Notes Theor. Comput. Sci.* 65(1), pp. 29–53, doi:10.1016/S1571-0661(04)80358-X.
- [5] F. Bartels (2004): *On Generalised Coinduction and Probabilistic Specification Formats*. Ph.D. thesis, VU University Amsterdam.
- [6] B. Bloom, S. Istrail & A.R. Meyer (1995): *Bisimulation can't be traced*. *J. ACM* 42, pp. 232–268, doi:10.1145/200836.200876.
- [7] R. Bol & J.F. Groote (1996): *The meaning of negative premises in transition system specifications*. *J. ACM* 43, pp. 863–914, doi:10.1145/234752.234756.

- [8] P.R. D'Argenio, D. Gebler & M.D. Lee (2015): *A general SOS theory for the specification of probabilistic transition systems*. Available at <http://www.cs.famaf.unc.edu.ar/~lee/publications/DArgenioGeblerLee.pdf>. Submitted.
- [9] P.R. D'Argenio & M.D. Lee (2012): *Probabilistic Transition System Specification: Congruence and Full Abstraction of Bisimulation*. In: *Proc. FoSSaCS'12, LNCS 7213*, Springer, pp. 452–466, doi:10.1007/978-3-642-28729-9_30.
- [10] P.R. D'Argenio, P. Sánchez Terraf & N. Wolovick (2012): *Bisimulations for non-deterministic labelled Markov processes*. *Mathematical Structures in Computer Science* 22(1), pp. 43–68, doi:10.1017/S0960129511000454.
- [11] W. Fokkink (2000): *Rooted Branching Bisimulation as a Congruence*. *J. Comput. Syst. Sci.* 60(1), pp. 13–37, doi:10.1006/jess.1999.1663.
- [12] W. Fokkink & C. Verhoef (1998): *A Conservative Look at Operational Semantics with Variable Binding*. *Inf. Comput.* 146(1), pp. 24–54, doi:10.1006/inco.1998.2729.
- [13] R.J. van Glabbeek (2004): *The meaning of negative premises in transition system specifications II*. *J. Log. Algebr. Program.* 60-61, pp. 229–258, doi:10.1016/j.jlap.2004.03.007.
- [14] J.F. Groote (1993): *Transition system specifications with negative premises*. *Theoretical Computer Science* 118(2), pp. 263–299, doi:10.1016/0304-3975(93)90111-6.
- [15] J.F. Groote & F.W. Vaandrager (1992): *Structured Operational Semantics and Bisimulation as a Congruence*. *Inf. Comput.* 100(2), pp. 202–260, doi:10.1016/0890-5401(92)90013-6.
- [16] H. Hermanns, A. Parma, R. Segala, B. Wachter & L. Zhang (2011): *Probabilistic Logical Characterization*. *Inf. Comput.* 209(2), pp. 154–172, doi:10.1016/j.ic.2010.11.024.
- [17] R. Lanotte & S. Tini (2005): *Probabilistic Congruence for Semistochastic Generative Processes*. In: *Foundations of Software Science and Computational Structures, 8th International Conference, FOSSACS 2005.*, pp. 63–78, doi:10.1007/978-3-540-31982-5_4.
- [18] R. Lanotte & S. Tini (2009): *Probabilistic bisimulation as a congruence*. *ACM Trans. Comput. Log.* 10(2), doi:10.1145/1462179.1462181.
- [19] K.G. Larsen & A. Skou (1991): *Bisimulation through Probabilistic Testing*. *Inf. Comput.* 94(1), pp. 1–28, doi:10.1016/0890-5401(91)90030-6.
- [20] M.D. Lee & E. de Vink (2015): *Rooted branching bisimulation as a congruence for probabilistic transition systems*. In: *Proc. QAPL'15, EPTCS*. To appear.
- [21] M.D. Lee, D. Gebler & P.R. D'Argenio (2012): *Tree Rules in Probabilistic Transition System Specifications with Negative and Quantitative Premises*. In: *Proc. EXPRESS/SOS'12, EPTCS 89*, pp. 115–130, doi:10.4204/EPTCS.89.9.
- [22] R. Milner (1989): *Communication and Concurrency*. Prentice-Hall.
- [23] M.R. Mousavi, M.A. Reniers & J.F. Groote (2007): *SOS formats and meta-theory: 20 years after*. *Theor. Comput. Sci.* 373(3), pp. 238–272, doi:10.1016/j.tcs.2006.12.019.
- [24] G. Plotkin (1981): *A structural approach to operational semantics*. Report DAIMI FN-19, Aarhus University, doi:10.1016/j.jlap.2004.05.001. Reprinted in *J. Log. Algebr. Program.*, 60-61:17-139, 2004.
- [25] R. Segala (1995): *Modeling and Verification of Randomized Distributed Real-Time Systems*. Ph.D. thesis, MIT.
- [26] S. Tini (2010): *Non-expansive ϵ -bisimulations for probabilistic processes*. *Theor. Comput. Sci.* 411(22-24), pp. 2202–2222, doi:10.1016/j.tcs.2010.01.027.

Analysing and Comparing Encodability Criteria

Kirstin Peters*
TU Dresden, Germany

Rob van Glabbeek
NICTA[†] Sydney, Australia
Computer Science and Engineering, UNSW, Sydney, Australia

Encodings or the proof of their absence are the main way to compare process calculi. To analyse the quality of encodings and to rule out trivial or meaningless encodings, they are augmented with quality criteria. There exists a bunch of different criteria and different variants of criteria in order to reason in different settings. This leads to incomparable results. Moreover it is not always clear whether the criteria used to obtain a result in a particular setting do indeed fit to this setting. We show how to formally reason about and compare encodability criteria by mapping them on requirements on a relation between source and target terms that is induced by the encoding function. In particular we analyse the common criteria *full abstraction*, *operational correspondence*, *divergence reflection*, *success sensitiveness*, and *respect of barbs*; e.g. we analyse the exact nature of the simulation relation (coupled simulation versus bisimulation) that is induced by different variants of operational correspondence. This way we reduce the problem of analysing or comparing encodability criteria to the better understood problem of comparing relations on processes.

1 Introduction

Encodings are used to compare process calculi and to reason about their expressive power. Encodability criteria are conditions that limit the existence of encodings. Their main purpose is to rule out trivial or meaningless encodings, but they can also be used to limit attention to encodings that are of special interest in a particular domain or for a particular purpose. These quality criteria are the main tool in *separation results*, saying that one calculus is not expressible in another one; here one has to show that no encoding meeting these criteria exists. To obtain stronger separation results, care has to be taken in selecting quality criteria that are not too restrictive. For *encodability results*, saying that one calculus is expressible in another one, all one needs is an encoding, together with criteria testifying for the quality of the encoding. Here it is important that the criteria are not too weak.

In the literature various different criteria and different variants of the same criteria are employed to achieve separation and encodability results [6, 13, 15, 17, 18, 14, 3, 19, 2, 8, 24, 7]. Some criteria, like full abstraction or operational correspondence, are used frequently. Other criteria are used to enforce a property of encodings that might only be necessary within a certain domain. For instance, the homomorphic translation of the parallel operator—in general a rather strict criterion—was used in [17] to show the absence of an encoding from the synchronous into the asynchronous π -calculus, because this requirement forbids for the introduction of global coordinators. Thus this criterion is useful when reasoning about the concurrent behaviour of processes, although it is in general too strict to reason about their interleaving behaviour. Unfortunately it is not always obvious or clear whether the criteria used to obtain a result in a particular setting do indeed fit to this setting. Indeed, as discussed in [24], the

*Supported by funding of the Excellence Initiative by the German Federal and State Governments (Institutional Strategy, measure ‘support the best’).

[†]NICTA is funded by the Australian Government through the Department of Communications and the Australian Research Council through the ICT Centre of Excellence Program.

homomorphic translation of the parallel operator forbids more than global coordinators, i.e., is too strict even in a concurrent setting.

The different purposes of encodability criteria lead to very different kinds of conditions that are usually hard to analyse and compare directly. In fact even widely used criteria—as full abstraction—seem not to be fully understood by the community, as the need for articles as [9, 20] shows. In contrast to that, relations on processes—such as simulations and bisimulations—are a very well studied and understood topic (see for example [5]). Moreover it is natural to describe the behaviour of terms, or compare them, modulo some equivalence relation. Also many encodability criteria, like operational correspondence, are obviously designed with a particular kind of relation between processes in mind. Therefore, in order to be able to formally reason about encodability criteria, to completely capture and describe their semantic effect, and to analyse side conditions of combinations of criteria, we map them on conditions on relations between source and target terms.

We consider the disjoint union $\mathcal{P}_S \uplus \mathcal{P}_T$ of the terms or processes from the source and target languages of an encoding. Then we describe the effect an encodability criterion \mathbf{C} has on the class of permitted encoding functions in terms of a relation $\mathcal{R}_{[\cdot]}$ that relates at least all source terms to their literal translations, i.e., contains the pair $(S, \llbracket S \rrbracket)$ for all source terms S . If the encodability criterion \mathbf{C} is defined w.r.t. some additional relations on the source or target languages, as it is the case for full abstraction and operational correspondence, we usually also include these relations in $\mathcal{R}_{[\cdot]}$. In order to completely capture the effect of a criterion \mathbf{C} we aim at iff-results of the form

$[\cdot]$ satisfies \mathbf{C} iff there exists a relation $\mathcal{R}_{[\cdot]}$ such that $\forall S. (S, \llbracket S \rrbracket) \in \mathcal{R}_{[\cdot]}$ and $P(\mathcal{R}_{[\cdot]})$,

where P is the condition that captures the effect of \mathbf{C} . For example, an encoding reflects divergence iff there exists a relation $\mathcal{R}_{[\cdot]}$ such that $\forall S. (S, \llbracket S \rrbracket) \in \mathcal{R}_{[\cdot]}$ and $\mathcal{R}_{[\cdot]}$ reflects divergence.

We illustrate this approach by applying it to some very common criteria. We start with divergence reflection in §3.1, because it is simple and well understood. Accordingly, we do not gain significant new insights, but it suits us very well to introduce our approach. In the same way success sensitiveness and respect of barbs are analysed. We then switch to the criteria full abstraction in §3.2 and operational correspondence in §3.3, which are possibly not completely understood yet. In particular, we show a connection between full abstraction and transitivity, and prove to which kinds of simulation relations common variants of operational correspondence are linked. In §4 we analyse the effects of combining the above criteria. Since we first map the criteria to conditions on relations between source and target terms, analysing their combined effect requires us to identify a suitable witness relation for the combined conditions. Combining divergence reflection and success sensitiveness is simple, as illustrated in §4.1. Combining these two criteria with operational correspondence (§4.2) is more elaborate. Finally we analyse the effect of combining full abstraction with operational correspondence in §4.3.

All claims in this paper have been proved using the interactive theorem prover Isabelle/HOL [16]. The Isabelle implementation of the theories is available in the ‘Archive of Formal Proofs’ at

http://afp.sourceforge.net/entries/Encodability_Process_Calculi.shtml.

2 Technical Preliminaries

We analyse criteria used to reason about the quality of encodings between process calculi. We do not force any limitations on the considered calculi. A *process calculus* is a language $\mathcal{L}_C = (\mathcal{P}_C, \mapsto_C)$ consisting of a set of terms \mathcal{P}_C —its *syntax*—and a relation on terms $\mapsto_C \subseteq \mathcal{P}_C \times \mathcal{P}_C$ —its *semantics*. The elements of \mathcal{P}_C are called *process terms* or shortly processes or terms.

Here we assume that the semantics of the language is provided as a so-called reduction semantics, because in the context of encodings the treatment of reductions is simpler—the consideration of labelled semantics and of criteria using labelled steps is left for further work. A *step* $P \mapsto_C P'$ is an element $(P, P') \in \mapsto_C$. Let \Longrightarrow_C denote the reflexive and transitive closure of \mapsto_C . We write $P \mapsto_C$ if $\exists P'. P \mapsto_C P'$ and $P \mapsto_C^\omega$ if P can do an infinite sequence of steps. A term P such that $P \mapsto_C^\omega$ is called *divergent*.

Languages can be augmented with (a set of) relations $\mathcal{R}_C \subseteq \mathcal{P}_C^2$ on their processes. If $\mathcal{R} \subseteq B^2$ is a relation and $B' \subseteq B$, then $\mathcal{R}|_{B'} = \{(x, y) \mid x, y \in B' \wedge (x, y) \in \mathcal{R}\}$ denotes the restriction of \mathcal{R} to the domain B' . A relation \mathcal{R} *preserves* some condition $P : B \rightarrow \mathbb{B}$ (with \mathbb{B} representing the Booleans) if whenever $(P, Q) \in \mathcal{R}$ and P satisfies P then Q satisfies P . A relation \mathcal{R} *reflects* P if whenever $(P, Q) \in \mathcal{R}$ and Q satisfies P then also P satisfies P . Finally \mathcal{R} *respects* a condition P if \mathcal{R} preserves and reflects it. We use $r(\cdot)$, $s(\cdot)$, and $t(\cdot)$ to denote the *reflexive*, *symmetric*, and *transitive closure* of a binary relation, respectively.

Relations on process terms are an important tool to reason about processes and languages. Of special interest are simulation relations; in particular bisimulations. \mathcal{R} is a bisimulation if any two related processes mutually simulate their respective sequences of steps, such that the derivatives are again related.

Definition 2.1 (Bisimulation) \mathcal{R} is a (weak reduction) bisimulation if for each $(P, Q) \in \mathcal{R}$:

- $P \Longrightarrow P'$ implies $\exists Q'. Q \Longrightarrow Q' \wedge (P', Q') \in \mathcal{R}$
- $Q \Longrightarrow Q'$ implies $\exists P'. P \Longrightarrow P' \wedge (P', Q') \in \mathcal{R}$

Two terms are bisimilar if there exists a bisimulation that relates them.

The definition of a *strong (reduction) bisimulation* is obtained by replacing all \Longrightarrow by \mapsto in the above definition, i.e., a strong bisimulation requires that a step has to be simulated by a single step. Coupled similarity is strictly weaker than bisimilarity. As pointed out in [21], in contrast to bisimilarity it allows for intermediate states in simulations: states that cannot be identified with states of the simulated term. Each symmetric coupled simulation is a bisimulation.

Definition 2.2 (Coupled Simulation) A relation \mathcal{R} is a (weak reduction) coupled simulation if both $(\exists Q'. Q \Longrightarrow Q' \wedge (P', Q') \in \mathcal{R})$ and $(\exists Q'. Q \Longrightarrow Q' \wedge (Q', P') \in \mathcal{R})$ whenever $(P, Q) \in \mathcal{R}$ and $P \Longrightarrow P'$. Two terms are coupled similar if they are related by a coupled simulation in both directions.

An *encoding* from $\mathcal{L}_S = (\mathcal{P}_S, \mapsto_S)$ into $\mathcal{L}_T = (\mathcal{P}_T, \mapsto_T)$ relates two process calculi. We call \mathcal{L}_S the *source* and \mathcal{L}_T the *target language*. Accordingly, terms of \mathcal{P}_S are *source terms* and of \mathcal{P}_T *target terms*. In the simplest case an encoding from \mathcal{L}_S into \mathcal{L}_T is an *encoding function* $\llbracket \cdot \rrbracket : \mathcal{P}_S \rightarrow \mathcal{P}_T$ from source terms into target terms. Sometimes an encoding is defined by several functions, such as the encoding function and the renaming policy used in the framework of [8]. Else we identify an encoding with its encoding function.

An *encodability criterion* is a predicate on encoding functions, used to reason about the quality of encodings. We analyse such criteria by mapping them on requirements on relations $\mathcal{R}_{\llbracket \cdot \rrbracket} \subseteq (\mathcal{P}_S \uplus \mathcal{P}_T)^2$ on the disjoint union of the source and target terms of the considered encodings $\llbracket \cdot \rrbracket$. To simplify the presentation we assume henceforth that $\mathcal{P}_S \cap \mathcal{P}_T = \emptyset$ and thus $\mathcal{P}_S \uplus \mathcal{P}_T = \mathcal{P}_S \cup \mathcal{P}_T$. The Isabelle proofs do not rely on such an assumption. We say that a condition $P : (\mathcal{P}_S \uplus \mathcal{P}_T) \rightarrow \mathbb{B}$ is *preserved* by an encoding if for all source terms S that satisfy P , the condition P also holds for $\llbracket S \rrbracket$. A condition is *reflected* by an encoding if whenever $\llbracket S \rrbracket$ satisfies it, then so does S . Finally an encoding *respects* a condition if it both preserves and reflects it.

3 Analysing Encodability Criteria

An encoding function $\llbracket \cdot \rrbracket : \mathcal{P}_S \rightarrow \mathcal{P}_T$ maps source terms on target terms. Thereby it induces a relation on the combined domain of source and target terms that relates source terms with their literal translations. We start with this relation, i.e., in the simplest case we map an encodability criterion to a requirement on a relation $\mathcal{R}_{\llbracket \cdot \rrbracket} \subseteq (\mathcal{P}_S \uplus \mathcal{P}_T)^2$ that contains at least the pairs $(S, \llbracket S \rrbracket)$ for all source terms $S \in \mathcal{P}_S$. If we consider a criterion that is defined w.r.t. some relations on the source or target, we will also include these relations in $\mathcal{R}_{\llbracket \cdot \rrbracket}$, possibly closing the latter under reflexivity, symmetry, and/or transitivity.

Alternatively, we could require that $\mathcal{R}_{\llbracket \cdot \rrbracket}$ relates source terms and their literal translations in both directions, meaning that $(S, \llbracket S \rrbracket) \in \mathcal{R}_{\llbracket \cdot \rrbracket}$ and $(\llbracket S \rrbracket, S) \in \mathcal{R}_{\llbracket \cdot \rrbracket}$ for all source terms $S \in \mathcal{P}_S$. However, this condition limits our analysis to properties that are respected. It does not allow us to reason about properties like divergence reflection, where some condition need only to be reflected but not necessarily be preserved, or vice versa. Accordingly we follow the first approach.

3.1 Divergence Reflection and Observables

We start with divergence reflection as defined in [8], because it is often easy to establish and well understood. An encoding reflects divergence if it does not introduce divergence, i.e., if all divergent translations result from divergent source terms.

Definition 3.1 (Divergence Reflection) *An encoding $\llbracket \cdot \rrbracket : \mathcal{P}_S \rightarrow \mathcal{P}_T$ reflects divergence if $\llbracket S \rrbracket \mapsto_T^\omega$ implies $S \mapsto_S^\omega$ for all source terms $S \in \mathcal{P}_S$.*

We can reformulate this criterion as follows: An encoding reflects divergence if it reflects the predicate $\lambda P. P \mapsto^\omega$. To analyse this criterion it suffices to consider the relation $\{(S, \llbracket S \rrbracket) \mid S \in \mathcal{P}_S\}$. It is obvious that an encoding reflects divergence iff $\{(S, \llbracket S \rrbracket) \mid S \in \mathcal{P}_S\}$ reflects divergence, i.e., reflects the predicate $\lambda P. P \mapsto^\omega$. In fact we can generalise this case. If an encodability criterion can be described by the preservation or reflection of a predicate, then an encoding satisfies this criterion iff $\{(S, \llbracket S \rrbracket) \mid S \in \mathcal{P}_S\}$ preserves or reflects this predicate. Of course direction “if” holds for any relation that contains at least the pairs $(S, \llbracket S \rrbracket)$. We use the relation $\{(S, \llbracket S \rrbracket) \mid S \in \mathcal{P}_S\}$ as a witness and it allows us to analyse the combination of different criteria later.

Lemma 3.2 (Preservation) *Let $P : (\mathcal{P}_S \uplus \mathcal{P}_T) \rightarrow \mathbb{B}$ be a predicate. An encoding preserves the predicate P iff $\exists \mathcal{R}_{\llbracket \cdot \rrbracket}. (\forall S. (S, \llbracket S \rrbracket) \in \mathcal{R}_{\llbracket \cdot \rrbracket}) \wedge \mathcal{R}_{\llbracket \cdot \rrbracket}$ preserves P .*

We obtain a similar result if we replace the unary predicate $P(\cdot)$ by the binary predicate $P(\cdot, \cdot)$ of type $(\mathcal{P}_S \uplus \mathcal{P}_T) \times \mathcal{T} \rightarrow \mathbb{B}$ for some arbitrary type \mathcal{T} to represent predicates with several parameters. Moreover we obtain the same result for either reflection or respect instead of preservation.

Accordingly an encoding reflects divergence, i.e., the predicate $\lambda P. P \mapsto^\omega$, iff there exists a relation $\mathcal{R}_{\llbracket \cdot \rrbracket}$ that relates at least each source term to its literal translation and reflects this predicate.

Lemma 3.3 (Divergence Reflection) *An encoding $\llbracket \cdot \rrbracket : \mathcal{P}_S \rightarrow \mathcal{P}_T$ from \mathcal{L}_S into \mathcal{L}_T reflects divergence iff $\exists \mathcal{R}_{\llbracket \cdot \rrbracket}. (\forall S. (S, \llbracket S \rrbracket) \in \mathcal{R}_{\llbracket \cdot \rrbracket}) \wedge \mathcal{R}_{\llbracket \cdot \rrbracket}$ reflects divergence.*

In a similar way we can deal with the criterion barb sensitiveness. A *barb* is a property of a process that is treated as an observable, and whose reachability should be respected by an encoding. We assume that \mathcal{B} is a set of barbs that contains at least all barbs of the source and the target language. Moreover we assume that each language \mathcal{L} specifies its own predicate $\cdot \downarrow_{\mathcal{L}} \cdot$ such that $P \downarrow_{\mathcal{L}} a$ returns true if $P \in \mathcal{P}_{\mathcal{L}}$

and P has the barb a in \mathcal{L} . If a barb a is not relevant or present in a language \mathcal{L} then $P \downarrow_{\mathcal{L}} a$ does not hold for any $P \in \mathcal{P}_{\mathcal{L}}$. We use $P \downarrow_{\mathcal{L}} a$ if P reaches the barb a in \mathcal{L} , i.e., $P \downarrow_{\mathcal{L}} a \triangleq \exists P'. P \Longrightarrow_{\mathcal{L}} P' \wedge P' \downarrow_{\mathcal{L}} a$.

An encoding weakly respects source term barbs iff it respects the predicate $\lambda P a. P \downarrow a$. This holds iff $\{(S, \llbracket S \rrbracket) \mid S \in \mathcal{P}_S\}$ respects this predicate, which in turn is the case iff there exists a relation $\mathcal{R}_{[\cdot]}$ that relates at least each source term to its literal translation and respects this predicate.

Lemma 3.4 (Barb Sensitiveness) *Assume \mathcal{L}_S and \mathcal{L}_T each define a predicate $\cdot \downarrow \cdot : \mathcal{P} \times \mathcal{B} \rightarrow \mathbb{B}$. $\llbracket \cdot \rrbracket : \mathcal{P}_S \rightarrow \mathcal{P}_T$ weakly respects barbs iff $\exists \mathcal{R}_{[\cdot]}. (\forall S. (S, \llbracket S \rrbracket) \in \mathcal{R}_{[\cdot]}) \wedge \mathcal{R}_{[\cdot]}$ weakly respects barbs.*

Again we obtain a similar result if we replace respect by preservation or reflection or if we consider the existence instead of the reachability of barbs.

However, only very few encodings directly preserve or reflect barbs. More often barbs are translated, as for example in the encodings between different variants of the π -calculus in [15, 23] or the two translations from CSP into variants of CCS with name passing in [10]. Since we do not fix the definition of $\cdot \downarrow_{\mathcal{L}} \cdot$, this can for instance be expressed by adapting this predicate in the target language.

In a similar way we can deal with the criterion success sensitiveness. This criterion was proposed by Gorla as part of his encodability framework [8]. An encoding is success sensitive if it respects reachability of a particular process \checkmark that represents successful termination, or some other form of success, and is added to the syntax of the source as well as the target language. We write $P \downarrow_{\checkmark}$ to denote the fact that P is successful—however this predicate might be defined in the particular source or target language. Reachability of success is then defined as $P \downarrow_{\checkmark} \triangleq \exists P'. P \Longrightarrow P' \wedge P' \downarrow_{\checkmark}$. An encoding is success sensitive if each source term and its translation answer the test for reachability of success in the same way.

Definition 3.5 (Success Sensitiveness) *Let \mathcal{L}_S and \mathcal{L}_T each define a predicate $\cdot \downarrow_{\checkmark} : \mathcal{P} \rightarrow \mathbb{B}$. An encoding $\llbracket \cdot \rrbracket : \mathcal{P}_S \rightarrow \mathcal{P}_T$ is success sensitive if, for all $S \in \mathcal{P}_S$, $S \downarrow_{\checkmark}$ iff $\llbracket S \rrbracket \downarrow_{\checkmark}$.*

Accordingly, an encoding is success sensitive iff it respects the predicate $\lambda P. P \downarrow_{\checkmark}$. This is the case iff $\{(S, \llbracket S \rrbracket) \mid S \in \mathcal{P}_S\}$ respects this predicate, which in turn is the case iff there exists a relation $\mathcal{R}_{[\cdot]}$ that relates at least each source term to its literal translation and respects this predicate.

Lemma 3.6 (Success Sensitiveness) *Assume \mathcal{L}_S and \mathcal{L}_T each define a predicate $\cdot \downarrow_{\checkmark} : \mathcal{P} \rightarrow \mathbb{B}$. An encoding $\llbracket \cdot \rrbracket : \mathcal{P}_S \rightarrow \mathcal{P}_T$ is success sensitive iff $\exists \mathcal{R}_{[\cdot]}. (\forall S. (S, \llbracket S \rrbracket) \in \mathcal{R}_{[\cdot]}) \wedge \mathcal{R}_{[\cdot]}$ respects $\lambda P. P \downarrow_{\checkmark}$.*

Success sensitiveness links source term behaviours to behaviours of target terms. If the source and the target language are very different, they can impose quite different kinds of behaviour that might be hard to compare directly. For example, observables in the π -calculus refer to the existence of unguarded input or output prefixes [11], whereas in the core of mobile ambients there are no in- or outputs but only ambients and action prefixes that describe the entering, leaving, and opening of an ambient [4]. Success sensitiveness allows to compare such languages by introducing a new kind of barb that can be understood in both calculi. If we want to compare two languages that are very similar, such as two variants of the same calculus, we can demand stricter encodability criteria and compare their barbs directly.

Next we concentrate on criteria that cannot be expressed simply by the preservation or reflection of some predicate.

3.2 Full Abstraction

Full abstraction was probably the first criterion that was widely used to reason about the quality of encodings [26, 12, 22]. This criterion is defined w.r.t. a relation $\mathcal{R}_S \subseteq \mathcal{P}_S^2$ on source terms and a relation $\mathcal{R}_T \subseteq \mathcal{P}_T^2$ on target terms. An encoding is fully abstract w.r.t. \mathcal{R}_S and \mathcal{R}_T if two source terms are related by \mathcal{R}_S iff their literal translations are related by \mathcal{R}_T .

Definition 3.7 (Full Abstraction) *An encoding $\llbracket \cdot \rrbracket : \mathcal{P}_S \rightarrow \mathcal{P}_T$ is fully abstract w.r.t. the relations $\mathcal{R}_S \subseteq \mathcal{P}_S^2$ and $\mathcal{R}_T \subseteq \mathcal{P}_T^2$ if, for all $S_1, S_2 \in \mathcal{P}_S$, $(S_1, S_2) \in \mathcal{R}_S$ iff $(\llbracket S_1 \rrbracket, \llbracket S_2 \rrbracket) \in \mathcal{R}_T$.*

There are a number of trivial full abstraction results, i.e., results that hold for all (or nearly all) encodings (see e.g. [9, 20]). In particular, for each encoding and each target term relation $\mathcal{R}_T \subseteq \mathcal{P}_T^2$ there exists a source term relation $\mathcal{R}_S \subseteq \mathcal{P}_S^2$, namely $\{(S_1, S_2) \mid (\llbracket S_1 \rrbracket, \llbracket S_2 \rrbracket) \in \mathcal{R}_T\}$, such that the encoding is fully abstract w.r.t. \mathcal{R}_S and \mathcal{R}_T . For each injective encoding and each source term relation $\mathcal{R}_S \subseteq \mathcal{P}_S^2$, there exists $\mathcal{R}_T \subseteq \mathcal{P}_T^2$, namely $\{(\llbracket S_1 \rrbracket, \llbracket S_2 \rrbracket) \mid (S_1, S_2) \in \mathcal{R}_S\}$, such that the encoding is fully abstract w.r.t. \mathcal{R}_S and \mathcal{R}_T . Accordingly we consider full abstraction w.r.t. fixed source and target term relations.

As suggested above, we map this criterion on a relation that relates at least each source term to its literal translation and includes the relations \mathcal{R}_S and \mathcal{R}_T . If we additionally add pairs of the form $(\llbracket S \rrbracket, S)$ for all $S \in \mathcal{P}_S$, we make an interesting observation. If we surround the pair $(S_1, S_2) \in \mathcal{R}_S$ by the pairs $(\llbracket S_1 \rrbracket, S_1)$ and $(S_2, \llbracket S_2 \rrbracket)$ and add transitivity we obtain the pair $(\llbracket S_1 \rrbracket, \llbracket S_2 \rrbracket)$. Similarly, from transitivity, $(S_1, \llbracket S_1 \rrbracket)$, $(\llbracket S_1 \rrbracket, \llbracket S_2 \rrbracket)$, and $(\llbracket S_2 \rrbracket, S_2)$ we obtain the pair (S_1, S_2) . Because of this, an encoding is fully abstract w.r.t. the preorders \mathcal{R}_S and \mathcal{R}_T iff there exists a transitive relation $\mathcal{R}_{\llbracket \cdot \rrbracket}$ that relates at least each source term to its literal translation in both directions, such that the restriction of $\mathcal{R}_{\llbracket \cdot \rrbracket}$ to source/target terms is $\mathcal{R}_S/\mathcal{R}_T$.

Lemma 3.8 (Full Abstraction) *$\llbracket \cdot \rrbracket : \mathcal{P}_S \rightarrow \mathcal{P}_T$ is fully abstract w.r.t. the preorders $\mathcal{R}_S \subseteq \mathcal{P}_S^2$ and $\mathcal{R}_T \subseteq \mathcal{P}_T^2$ iff $\exists \mathcal{R}_{\llbracket \cdot \rrbracket}. (\forall S. (S, \llbracket S \rrbracket), (\llbracket S \rrbracket, S) \in \mathcal{R}_{\llbracket \cdot \rrbracket}) \wedge \mathcal{R}_S = \mathcal{R}_{\llbracket \cdot \rrbracket} \upharpoonright_{\mathcal{P}_S} \wedge \mathcal{R}_T = \mathcal{R}_{\llbracket \cdot \rrbracket} \upharpoonright_{\mathcal{P}_T} \wedge \mathcal{R}_{\llbracket \cdot \rrbracket}$ is transitive.*

Thus an encoding is fully abstract w.r.t. \mathcal{R}_S and \mathcal{R}_T if the encoding function combines the relations \mathcal{R}_S and \mathcal{R}_T in a transitive way.

In order to allow combinations with criteria like divergence reflection, i.e., predicates that are not respected but preserved or reflected, we get rid of the requirement on the pairs $(\llbracket S \rrbracket, S)$. Therefore we consider the symmetric closure of $\mathcal{R}_{\llbracket \cdot \rrbracket}$. An encoding is fully abstract w.r.t. the equivalences \mathcal{R}_S and \mathcal{R}_T iff there exists a relation $\mathcal{R}_{\llbracket \cdot \rrbracket}$ that relates at least each source term to its literal translation, such that the restriction of the symmetric closure of $\mathcal{R}_{\llbracket \cdot \rrbracket}$ to source/target terms is $\mathcal{R}_S/\mathcal{R}_T$ and the symmetric closure of $\mathcal{R}_{\llbracket \cdot \rrbracket}$ is a preorder.

Lemma 3.9 (Full Abstraction) *An encoding $\llbracket \cdot \rrbracket : \mathcal{P}_S \rightarrow \mathcal{P}_T$ is fully abstract w.r.t. the equivalences $\mathcal{R}_S \subseteq \mathcal{P}_S^2$ and $\mathcal{R}_T \subseteq \mathcal{P}_T^2$ iff $\exists \mathcal{R}_{\llbracket \cdot \rrbracket}. (\forall S. (S, \llbracket S \rrbracket) \in \mathcal{R}_{\llbracket \cdot \rrbracket}) \wedge \mathcal{R}_S = s(\mathcal{R}_{\llbracket \cdot \rrbracket}) \upharpoonright_{\mathcal{P}_S} \wedge \mathcal{R}_T = s(\mathcal{R}_{\llbracket \cdot \rrbracket}) \upharpoonright_{\mathcal{P}_T} \wedge s(\mathcal{R}_{\llbracket \cdot \rrbracket})$ is a preorder.*

Since it is always possible to construct a relation that includes \mathcal{R}_S , \mathcal{R}_T , and pairs $(S, \llbracket S \rrbracket)$, the crucial requirement on the right-hand side is transitivity. A discussion of this criterion and references to earlier such discussions can be found in [22, 9].

3.3 Operational Correspondence

To strengthen full abstraction it is often combined with operational correspondence. This criterion requires that source terms and their translations ‘behave’ similar, by requiring that steps are preserved and reflected modulo some target term relation $\mathcal{R}_T \subseteq \mathcal{P}_T^2$. Intuitively an encoding is operational corresponding w.r.t. \mathcal{R}_T if each source term step is simulated by its translation, i.e., $\llbracket \cdot \rrbracket$ does not remove source behaviour (*completeness*), and each step of the target is part of the simulation of a source term step, i.e., $\llbracket \cdot \rrbracket$ does not introduce new behaviour (*soundness*). There are a number of different variants of this criterion. We consider three unlabelled variants [15, 8]. In particular the last variant, proposed in [8], was used for numerous encodability and separation results.

Definition 3.10 (Operational Correspondence) An encoding $\llbracket \cdot \rrbracket : \mathcal{P}_S \rightarrow \mathcal{P}_T$ is strongly operationally corresponding w.r.t. $\mathcal{R}_T \subseteq \mathcal{P}_T^2$ if it is:

Strongly Complete: $\forall S, S'. S \mapsto_S S'$ implies $(\exists T. \llbracket S \rrbracket \mapsto_T T \wedge (\llbracket S' \rrbracket, T) \in \mathcal{R}_T)$

Strongly Sound: $\forall S, T. \llbracket S \rrbracket \mapsto_T T$ implies $(\exists S'. S \mapsto_S S' \wedge (\llbracket S' \rrbracket, T) \in \mathcal{R}_T)$

$\llbracket \cdot \rrbracket : \mathcal{P}_S \rightarrow \mathcal{P}_T$ is operationally corresponding w.r.t. $\mathcal{R}_T \subseteq \mathcal{P}_T^2$ if it is:

Complete: $\forall S, S'. S \Longrightarrow_S S'$ implies $(\exists T. \llbracket S \rrbracket \Longrightarrow_T T \wedge (\llbracket S' \rrbracket, T) \in \mathcal{R}_T)$

Sound: $\forall S, T. \llbracket S \rrbracket \Longrightarrow_T T$ implies $(\exists S'. S \Longrightarrow_S S' \wedge (\llbracket S' \rrbracket, T) \in \mathcal{R}_T)$

$\llbracket \cdot \rrbracket : \mathcal{P}_S \rightarrow \mathcal{P}_T$ is weakly operationally corresponding w.r.t. $\mathcal{R}_T \subseteq \mathcal{P}_T^2$ if it is:

Complete: $\forall S, S'. S \Longrightarrow_S S'$ implies $(\exists T. \llbracket S \rrbracket \Longrightarrow_T T \wedge (\llbracket S' \rrbracket, T) \in \mathcal{R}_T)$

Weakly Sound: $\forall S, T. \llbracket S \rrbracket \Longrightarrow_T T$ implies $(\exists S', T'. S \Longrightarrow_S S' \wedge T \Longrightarrow_T T' \wedge (\llbracket S' \rrbracket, T') \in \mathcal{R}_T)$

Again this criterion is trivial if we do not fix the target term relation. Each encoding is operational corresponding w.r.t. the universal relation on target terms.

The formulation of operational correspondence (in all its variants) strongly reminds us of simulation relations on processes, such as bisimilarity. Obviously this criterion is designed in order to establish a simulation-like relation between source and target terms. We now determine the exact nature of this relation. The first two variants exactly describe strong and weak bisimilarity up to \mathcal{R}_T . More precisely, an encoding is operational corresponding w.r.t. a preorder \mathcal{R}_T that is a bisimulation iff there exists a preorder $\mathcal{R}_{\llbracket \cdot \rrbracket}$, such as $t(\text{r}(\{(S, \llbracket S \rrbracket) \mid S \in \mathcal{P}_S\} \cup \mathcal{R}_T))$, that is a bisimulation, relates at least all source terms to their literal translations, and such that $\mathcal{R}_T = \mathcal{R}_{\llbracket \cdot \rrbracket} \upharpoonright_{\mathcal{P}_T}$, and for all pairs $(S, T) \in \mathcal{R}_{\llbracket \cdot \rrbracket}$ it holds that $(\llbracket S \rrbracket, T) \in \mathcal{R}_T$. The last condition is necessary to ensure operational correspondence, and $\mathcal{R}_T = \mathcal{R}_{\llbracket \cdot \rrbracket} \upharpoonright_{\mathcal{P}_T}$ ensures that \mathcal{R}_T is a bisimulation if $\mathcal{R}_{\llbracket \cdot \rrbracket}$ is. Accordingly, operational correspondence ensures that source terms and their translations are bisimilar.

Lemma 3.11 (Operational Correspondence) An encoding $\llbracket \cdot \rrbracket : \mathcal{P}_S \rightarrow \mathcal{P}_T$ is operational corresponding w.r.t. a preorder $\mathcal{R}_T \subseteq \mathcal{P}_T^2$ that is a bisimulation iff $\exists \mathcal{R}_{\llbracket \cdot \rrbracket}. (\forall S. (S, \llbracket S \rrbracket) \in \mathcal{R}_{\llbracket \cdot \rrbracket}) \wedge \mathcal{R}_T = \mathcal{R}_{\llbracket \cdot \rrbracket} \upharpoonright_{\mathcal{P}_T} \wedge (\forall S, T. (S, T) \in \mathcal{R}_{\llbracket \cdot \rrbracket} \rightarrow (\llbracket S \rrbracket, T) \in \mathcal{R}_T) \wedge \mathcal{R}_{\llbracket \cdot \rrbracket}$ is a preorder and a bisimulation.

We obtain the same result if we replace operational correspondence by strong operational correspondence and bisimulation by strong bisimulation.

Weak and strong bisimilarity are often considered as *the* standard reference relations for calculi like the π -calculus. Thus the above result imposes an important property for the comparison of languages. If bisimilarity is the standard reference relation, i.e., if we usually do not record differences between terms that cannot be observed by bisimilarity, then an encoding that ensures that source terms and their translations are bisimilar strongly validates the claim that the target language is at least as expressive as the source language. Nonetheless, comparisons of different languages are very often considered only modulo weak operational correspondence and not operational correspondence. As discussed in [21, 10], relating source terms and their literal translations by a bisimulation does not allow for intermediate states, i.e., states that occur in simulations of source term steps and thus intuitively are in between two source term translations but are not related to source terms themselves. Intermediate states result from partial commitments. If a source term can evolve to one of three different derivatives, operational correspondence (in all variants) ensures that the translation has the same possible evolutions. But operational correspondence requires that the decision on which of the three possibilities is chosen is done in a single step. Weak operational correspondence allows for partial commitments, where a first step may rule out one possibility but not decide on one of the remaining two. Thus weak operational correspondence is much more flexible and allows to encode source term concepts that have no direct counterpart in the target.

Obtaining a result similar to Lemma 3.11 for weak operational correspondence is not that easy. Again this criterion is linked to a simulation condition on relations between source and target terms up to \mathcal{R}_T , but weak operational correspondence does not directly map to a well-known kind of simulation relation. It is linked to a simulation relation that is in between coupled similarity and bisimilarity. We call it correspondence similarity.

Definition 3.12 (Correspondence Simulation) *A relation \mathcal{R} is a (weak reduction) correspondence simulation if for each $(P, Q) \in \mathcal{R}$:*

- $P \Longrightarrow P'$ implies $\exists Q'. Q \Longrightarrow Q' \wedge (P', Q') \in \mathcal{R}$
- $Q \Longrightarrow Q'$ implies $\exists P'', Q''. P \Longrightarrow P'' \wedge Q' \Longrightarrow Q'' \wedge (P'', Q'') \in \mathcal{R}$

Two terms are correspondence similar if a correspondence simulation relates them.

Just as coupled similarity, correspondence similarity allows for intermediate states that result from partial commitments, but in contrast to coupled similarity these intermediate states are not necessarily covered in the relation. Correspondence similarity is obviously strictly weaker than bisimilarity, but it implies coupled similarity.

Lemma 3.13 *For each correspondence simulation \mathcal{R} there exists a coupled simulation \mathcal{R}' such that $\forall (P, Q) \in \mathcal{R}. (P, Q), (Q, P) \in \mathcal{R}'$.*

Correspondence simulation is linked to weak operational correspondence in the same way as bisimilarity is linked to operational correspondence.

Lemma 3.14 (Weak Operational Correspondence) $\llbracket \cdot \rrbracket : \mathcal{P}_S \rightarrow \mathcal{P}_T$ *is weakly operat. corresp. w.r.t. a preorder $\mathcal{R}_T \subseteq \mathcal{P}_T^2$ that is a correspondence simulation iff $\exists \mathcal{R}_{\llbracket \cdot \rrbracket}. (\forall S. (S, \llbracket S \rrbracket) \in \mathcal{R}_{\llbracket \cdot \rrbracket}) \wedge \mathcal{R}_T = \mathcal{R}_{\llbracket \cdot \rrbracket} \upharpoonright_{\mathcal{P}_T} \wedge (\forall S, T. (S, T) \in \mathcal{R}_{\llbracket \cdot \rrbracket} \rightarrow (\llbracket S \rrbracket, T) \in \mathcal{R}_T) \wedge \mathcal{R}_{\llbracket \cdot \rrbracket}$ is a preorder and a correspondence simulation.*

Accordingly, weak operational correspondence ensures that source terms and their literal translations are correspondence similar and thus coupled similar.

Correspondence similarity and coupled similarity are weaker than bisimilarity. Nevertheless, proving that a relation is a correspondence simulation and, even more, showing that a particular pair of terms is contained in a correspondence simulation, can be more difficult than it is in the case of bisimulation. Fortunately, encodings that satisfy only weak operational correspondence—and introduce partial commitments—often do so w.r.t. a variant of bisimilarity. As example consider the de-centralised encoding of [10]. It translates from CSP into asynchronous CCS with name passing and matching. [10] proves that this encoding is operational corresponding w.r.t. a target term preorder \mathcal{R}_T that is a weak reduction bisimulation. Thus, by Lemma 3.14, the encoding ensures that source terms and their literal translations are correspondence similar, and thus coupled similar.

4 Combining Encodability Criteria

As done in [8], often several different criteria are combined to ensure the quality of an encoding. Of course we have to ensure that the criteria we want to combine do not contradict each other and thus trivially rule out any kind of encoding. Moreover, the combination of criteria might lead to unexpected side effects, such that their combined effect on the quality of encodings is no longer obvious or clear. One major motivation of our desire to analyse encodability criteria is to be able to formally compare them and analyse side effects that result from their combinations.

In the previous section we derive iff-results linking a single criterion with the existence of a relation between source and target terms satisfying specific conditions. Of course we can trivially combine two

such results by considering two different source-target relations on the right-hand side. But this way side effects that results from the combination of the criteria remain hidden. Instead we want to combine the criteria into conditions of a single source-target relation. Therefore we need to find a witness, i.e., a relation that satisfies the conditions of both relations.

4.1 Divergence Reflection and Success Sensitiveness

The combinations of criteria defined on the pairs of $\mathcal{R}_{[\cdot]}$ —such as the preservation, reflection, or respect of some predicate—are easy to analyse. Obviously an encoding reflects divergence and respects success iff there exists a relation $\mathcal{R}_{[\cdot]}$ that relates at least each source term to its literal translation and both reflects divergence and respects success.

Lemma 4.1 *Assume \mathcal{L}_S and \mathcal{L}_T each define a predicate $\downarrow_{\checkmark}: \mathcal{P} \rightarrow \mathbb{B}$. $[\cdot]: \mathcal{P}_S \rightarrow \mathcal{P}_T$ reflects divergence and respects success iff $\exists \mathcal{R}_{[\cdot]}. (\forall S. (S, [S]) \in \mathcal{R}_{[\cdot]}) \wedge \mathcal{R}_{[\cdot]}$ reflects divergence and respects success.*

The \Leftarrow -direction of Lemma 4.1 is an immediate corollary of Lemmas 3.3 and 3.6. For the other direction we obtain from these lemmata two relations that satisfy the condition $C \triangleq \forall S. (S, [S]) \in \mathcal{R}_{[\cdot]}$ and of which one reflects divergence and the other respects success. We have to combine these two relations into a single relation that satisfies all three conditions. If the latter two conditions are defined on the pairs of the respective relations, this is always possible. The reason is that the condition C ensures that we can use $\{(S, [S]) \mid S \in \mathcal{P}_S\}$ as a witness for both relations and thus as a witness for their combined effect. More precisely, if there are two relations that both satisfy C and each satisfies a predicate about the pairs of the respective relation, then there exists a single relation, namely $\{(S, [S]) \mid S \in \mathcal{P}_S\}$, that satisfies all three conditions.

Lemma 4.2 *Let $\mathcal{R}_1, \mathcal{R}_2 \subseteq (\mathcal{P}_S \uplus \mathcal{P}_T)^2$ and the predicates P_1, P_2 be such that $\forall i \in \{1, 2\}. \forall S. (S, [S]) \in \mathcal{R}_i$ and $\forall i \in \{1, 2\}. \forall (P, Q) \in \mathcal{R}_i. P_i((P, Q))$. Then there exists a relation $\mathcal{R}_{[\cdot]} \subseteq (\mathcal{P}_S \uplus \mathcal{P}_T)^2$ such that $\forall S. (S, [S]) \in \mathcal{R}_{[\cdot]}$ and $\forall i \in \{1, 2\}. \forall (P, Q) \in \mathcal{R}_{[\cdot]}. P_i((P, Q))$.*

4.2 Adding Operational Correspondence

Gorla [8] combines five criteria to define ‘good’ encodings. Three of these—the ‘semantical’ ones—we considered in Section 3: weak operational correspondence (called ‘operational correspondence’ in [8]) w.r.t. a relation \mathcal{R}_T , success sensitiveness, and divergence reflection. Gorla assumes that, ‘for the sake of coherence’ as he claims, the relation \mathcal{R}_T never relates two process T_P and T_Q such that $T_P \downarrow_{\checkmark}$ and $T_Q \not\downarrow_{\checkmark}$, i.e., \mathcal{R}_T has to respect (reachability of) success. This allows us to find a witness relation to combine the effect of weak operational correspondence and success sensitiveness. Our iff-result for weak operational correspondence requires that this relation is a preorder, has to relate source terms with their literal translations, and satisfies $\mathcal{R}_T = \mathcal{R}_{[\cdot]} \upharpoonright_{\mathcal{P}_T}$. Because of that, a minimal witness is $t(r(\{(S, [S]) \mid S \in \mathcal{P}_S\} \cup \mathcal{R}_T))$. This witness also satisfies $\forall S, T. (S, T) \in \mathcal{R}_{[\cdot]} \rightarrow ([S], T) \in \mathcal{R}_T$. Without the condition that \mathcal{R}_T respects success—or another suitable assumption—we cannot ensure that $t(r(\{(S, [S]) \mid S \in \mathcal{P}_S\} \cup \mathcal{R}_T))$ respects success and thus we find no witness for the combination of the respective conditions.

Lemma 4.3 *Assume $\mathcal{L}_S, \mathcal{L}_T$ each define a predicate $\downarrow_{\checkmark}: \mathcal{P} \rightarrow \mathbb{B}$. An encoding $[\cdot]: \mathcal{P}_S \rightarrow \mathcal{P}_T$ is success sensitive and weakly operational corresponding w.r.t. a preorder $\mathcal{R}_T \subseteq \mathcal{P}_T^2$ that is a success respecting correspondence simulation iff $\exists \mathcal{R}_{[\cdot]}. (\forall S. (S, [S]) \in \mathcal{R}_{[\cdot]}) \wedge \mathcal{R}_T = \mathcal{R}_{[\cdot]} \upharpoonright_{\mathcal{P}_T} \wedge \mathcal{R}_{[\cdot]}$ respects success $\wedge (\forall S, T. (S, T) \in \mathcal{R}_{[\cdot]} \rightarrow ([S], T) \in \mathcal{R}_T) \wedge \mathcal{R}_{[\cdot]}$ is a preorder and a correspondence simulation.*

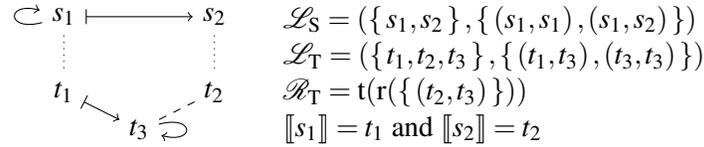


Figure 1: Encoding satisfying operational correspondence and divergence reflection

We obtain a similar result if we replace weak operational correspondence and correspondence simulation by operational correspondence and bisimulation. Similarly, we can replace weak operational correspondence, correspondence simulation, and the predicate \Downarrow_{\checkmark} in the definition of success sensitiveness by strong operational correspondence, strong bisimulation, and \downarrow_{\checkmark} . Also, in all variants of the above result we can add on the left-hand side that the encoding as well as \mathcal{R}_T (weakly) respect barbs iff we add on the right-hand side that $\mathcal{R}_{[\cdot]}$ (weakly) respects barbs.

As an example we can consider—once more—the de-centralised encoding from CSP into a variant of CCS of [10]. Additionally to operational correspondence w.r.t. a preorder \mathcal{R}_T that is a bisimulation, [10] proves that this encoding satisfies success sensitiveness, divergence reflection, barb sensitiveness—w.r.t. standard CSP barbs in the source and a notion of translated barbs in the target—and preservation of distributability (a criterion defined in [25]). \mathcal{R}_T respects success and weakly respects barbs (but does not reflect divergence). Thus the encoding ensures that source terms and their literal translations are correspondence similar and thus coupled similar w.r.t. a relation that respects success and weakly respects barbs.

Success sensitiveness significantly strengthens the requirements on a simulation relation like correspondence simulation or bisimulation. Thus the combined effect—a success respecting correspondence simulation—is stronger than the effects of both criteria considered in isolation—a correspondence simulation and a success respecting relation. Accordingly, the framework of Gorla in [8] ensures that (among other conditions) source terms and their literal translations are correspondence similar w.r.t. a success respecting relation and thus—to refer to a more established simulation relation—are coupled similar w.r.t. a success respecting relation.

In [8] there is no such condition that links \mathcal{R}_T and divergence reflection. Requiring that \mathcal{R}_T reflects divergence would e.g. exclude weak bisimulation. Since this relation is often referred to as *the* standard relation for calculi as the π -calculus, excluding it would be too strict a requirement. As a consequence, the criteria in [8] do not allow to combine the effects of weak operational correspondence and divergence reflection into a single relation, as done for success sensitiveness. Consider the following counterexample, visualised in Figure 1. Obviously the encoding—indicated by the dotted line—satisfies operational correspondence w.r.t. \mathcal{R}_T —indicated by the dashed line—and reflects divergence. But to relate s_1 and its literal translation $[[s_1]] = t_1$ by a correspondence simulation, we have to simulate the step $s_1 \mapsto_S s_2$. Therefore we need either the pair (s_2, t_3) —which can be obtained by including \mathcal{R}_T in $\mathcal{R}_{[\cdot]}$ —or the pair (s_2, t_1) , but in either case the respective source-target relation does not reflect divergence. Thus in general an encoding that satisfies the criteria of [8] induces a source-target relation that is a correspondence simulation that only partially reflects divergence.

Of course particular encodings might satisfy stronger requirements than enforced by the minimal setting in [8]. If the encoding is operational corresponding w.r.t. a relation that reflects divergence, we can combine the effects of these two criteria in one relation. Accordingly, if an encoding reflects divergence, respects success, and satisfies operational correspondence w.r.t. a preorder that is a success respecting and divergence reflecting bisimulation, we can combine the conditions of all three relations

as in the following lemma.

Lemma 4.4 *Assume \mathcal{L}_S and \mathcal{L}_T each define a predicate $\downarrow_{\mathcal{V}}: \mathcal{P} \rightarrow \mathbb{B}$. An encoding $\llbracket \cdot \rrbracket: \mathcal{P}_S \rightarrow \mathcal{P}_T$ reflects divergence, respects success, and is operational corresponding w.r.t. a preorder $\mathcal{R}_T \subseteq \mathcal{P}_T^2$ that is a success respecting and divergence reflecting bisimulation iff $\exists \mathcal{R}_{\llbracket \cdot \rrbracket}. (\forall S. (S, \llbracket S \rrbracket) \in \mathcal{R}_{\llbracket \cdot \rrbracket}) \wedge \mathcal{R}_T = \mathcal{R}_{\llbracket \cdot \rrbracket} \upharpoonright_{\mathcal{P}_T} \wedge (\forall S, T. (S, T) \in \mathcal{R}_{\llbracket \cdot \rrbracket} \rightarrow (\llbracket S \rrbracket, T) \in \mathcal{R}_T) \wedge \mathcal{R}_{\llbracket \cdot \rrbracket}$ reflects divergence, respects success, and is a preorder and a bisimulation.*

Again we obtain similar results for weak operational correspondence and correspondence simulation as well as for strong operational correspondence, strong bisimulation, and $\downarrow_{\mathcal{V}}$.

4.3 Full Abstraction and Operational Correspondence

Before the framework in [8] was proposed, often a combination of full abstraction and operational correspondence was used. For simplicity we switch to source-target relations that relate source terms and their literal translations in both directions and assume that \mathcal{R}_S and \mathcal{R}_T are equivalences in the following. Then a witness for the effect of operational correspondence is $t(r(\{(S, \llbracket S \rrbracket), (\llbracket S \rrbracket, S) \mid S \in \mathcal{P}_S\} \cup \mathcal{R}_T))$. Since this relation is transitive, it indeed suffices as witness to combine the effects of full abstraction and operational correspondence. The only obstacle left is that, to cover the effect of full abstraction, the source-target relation should also include \mathcal{R}_S . Fortunately we do not have to include \mathcal{R}_S by construction, because its inclusion is ensured by full abstraction and the inclusion of \mathcal{R}_T . For every encoding $\llbracket \cdot \rrbracket$ that is fully abstract w.r.t. \mathcal{R}_S and \mathcal{R}_T and for all transitive relations $\mathcal{R}_{\llbracket \cdot \rrbracket}$ that relate at least all source terms to their literal translations in both directions, $\mathcal{R}_{\llbracket \cdot \rrbracket}$ contains \mathcal{R}_S iff the restriction of $\mathcal{R}_{\llbracket \cdot \rrbracket}$ to encoded source terms contains the restriction of \mathcal{R}_T to encoded sources.

Lemma 4.5 *Let $\llbracket \cdot \rrbracket: \mathcal{P}_S \rightarrow \mathcal{P}_T$ be an encoding that is fully abstract w.r.t. $\mathcal{R}_S \subseteq \mathcal{P}_S^2$ and $\mathcal{R}_T \subseteq \mathcal{P}_T^2$ and let $\mathcal{R}_{\llbracket \cdot \rrbracket} \subseteq (\mathcal{P}_S \uplus \mathcal{P}_T)^2$ be transitive such that $\forall S. (S, \llbracket S \rrbracket), (\llbracket S \rrbracket, S) \in \mathcal{R}_{\llbracket \cdot \rrbracket}$. Then $\mathcal{R}_S = \mathcal{R}_{\llbracket \cdot \rrbracket} \upharpoonright_{\mathcal{P}_S}$ iff $\forall S_1, S_2. (\llbracket S_1 \rrbracket, \llbracket S_2 \rrbracket) \in \mathcal{R}_T \leftrightarrow (\llbracket S_1 \rrbracket, \llbracket S_2 \rrbracket) \in \mathcal{R}_{\llbracket \cdot \rrbracket}$.*

Because of that an encoding is fully abstract w.r.t. \mathcal{R}_S and \mathcal{R}_T and operational corresponding w.r.t. a bisimulation \mathcal{R}_T iff there exists a transitive bisimulation that relates source terms and their literal translations in both directions and contains \mathcal{R}_S and \mathcal{R}_T .

Lemma 4.6 *Let $\mathcal{R}_S \subseteq \mathcal{P}_S^2$ and $\mathcal{R}_T \subseteq \mathcal{P}_T^2$ be equivalences. An encoding $\llbracket \cdot \rrbracket: \mathcal{P}_S \rightarrow \mathcal{P}_T$ is fully abstract w.r.t. \mathcal{R}_S and \mathcal{R}_T and operational corresponding w.r.t. \mathcal{R}_T and \mathcal{R}_T is a bisimulation iff $\exists \mathcal{R}_{\llbracket \cdot \rrbracket}. (\forall S. (S, \llbracket S \rrbracket), (\llbracket S \rrbracket, S) \in \mathcal{R}_{\llbracket \cdot \rrbracket}) \wedge \mathcal{R}_S = \mathcal{R}_{\llbracket \cdot \rrbracket} \upharpoonright_{\mathcal{P}_S} \wedge \mathcal{R}_T = \mathcal{R}_{\llbracket \cdot \rrbracket} \upharpoonright_{\mathcal{P}_T} \wedge \mathcal{R}_{\llbracket \cdot \rrbracket}$ is a transitive bisimulation.*

So what do we gain by combining the two criteria, that we do not obtain from each of them in isolation? In comparison to our iff-result for operational correspondence we add only the condition that $\mathcal{R}_S = \mathcal{R}_{\llbracket \cdot \rrbracket} \upharpoonright_{\mathcal{P}_S}$. As a consequence \mathcal{R}_S has to be a bisimulation.

Full abstraction ensures that \mathcal{R}_S and \mathcal{R}_T have the same basic properties. For example, if we either consider surjective encodings ($\forall T. \exists S. T = \llbracket S \rrbracket$) or restrict \mathcal{R}_T to encoded source terms ($\{(T_1, T_2) \mid \exists S_1, S_2. T_1 = \llbracket S_1 \rrbracket \wedge T_2 = \llbracket S_2 \rrbracket \wedge (T_1, T_2) \in \mathcal{R}_T\}$), then \mathcal{R}_S is reflexive iff \mathcal{R}_T is reflexive, and similarly for symmetry and transitivity. But properties such as being a bisimulation are not respected by full abstraction on its own. As counterexample consider the fully abstract but not operational corresponding encoding in Figure 2. Here \mathcal{R}_S is a bisimulation but \mathcal{R}_T is not. By removing the arrow $t_2 \mapsto_T t_3$ from the target and adding it to the source $s_2 \mapsto_S s_3$, the encoding remains fully abstract and \mathcal{R}_T becomes a bisimulation but \mathcal{R}_S loses this property. Operational correspondence does not refer to a source relation \mathcal{R}_S and thus does not enforce any properties on this relation. But combining full abstraction with operational correspondence w.r.t. a bisimulation \mathcal{R}_T enforces \mathcal{R}_S to be a bisimulation.

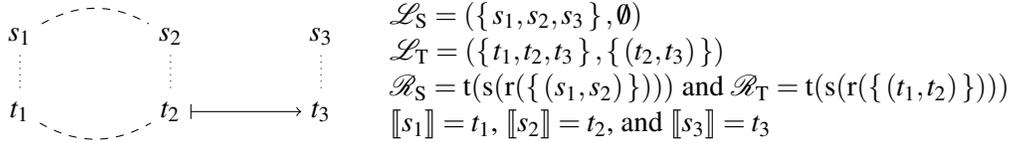


Figure 2: Encoding satisfying full abstraction but not operational correspondence

Lemma 4.7 *Let $\mathcal{R}_S \subseteq \mathcal{P}_S^2$ and $\mathcal{R}_T \subseteq \mathcal{P}_T^2$ be equivalences. If an encoding $\llbracket \cdot \rrbracket : \mathcal{P}_S \rightarrow \mathcal{P}_T$ is fully abstract w.r.t. \mathcal{R}_S and \mathcal{R}_T and operational corresponding w.r.t. \mathcal{R}_T and \mathcal{R}_T is a bisimulation then \mathcal{R}_S is a bisimulation.*

To conclude that \mathcal{R}_S is a bisimulation iff \mathcal{R}_T is a bisimulation, we have to get rid of pairs in \mathcal{R}_T that do not result from pairs of encoded source terms and their derivatives, because operational correspondence provides no information about such pairs. The simplest way to do so, is to assume a surjective encoding.

Lemma 4.8 *Let $\mathcal{R}_S \subseteq \mathcal{P}_S^2$ and $\mathcal{R}_T \subseteq \mathcal{P}_T^2$ be equivalences. If an encoding $\llbracket \cdot \rrbracket : \mathcal{P}_S \rightarrow \mathcal{P}_T$ is surjective ($\forall T. \exists S. T = \llbracket S \rrbracket$), fully abstract w.r.t. \mathcal{R}_S and \mathcal{R}_T , and operational corresponding w.r.t. \mathcal{R}_T then:*

$$\mathcal{R}_S \text{ is a bisimulation iff } \mathcal{R}_T \text{ is a bisimulation}$$

5 Conclusions

Within this paper we provide a number of results about different encodability criteria. In particular:

- We analyse divergence reflection, barb sensitiveness, success sensitiveness, full abstraction, and operational correspondence as well as several combinations of these criteria.
- We prove that different variants of operational correspondence correlate with different kinds of simulation relations from coupled similarity to strong bisimilarity.
- We define a new kind of simulation relation—correspondence similarity—that completely covers the effect of weak operational correspondence as proposed in [8].
- We relate the combination of success sensitiveness and operational correspondence w.r.t. a bisimulation with the existence of a success respecting bisimulation between source terms and their literal translations.
- We show that for surjective encodings the combination of full abstraction w.r.t. \mathcal{R}_S and \mathcal{R}_T and operational correspondence w.r.t. \mathcal{R}_T implies that \mathcal{R}_S is a bisimulation iff \mathcal{R}_T is a bisimulation.

In [7] a quality criterion for encodings was proposed that requires the translation $\llbracket S \rrbracket$ of a source term S to be related to S according to a behavioural equivalence or preorder defined on a domain of interpretation (such as labelled transition systems or reduction-based transition systems with barbs) that applies to both languages. This behavioural relation has to be chosen with care and should be meaningful for the application at hand. Possible choices include strong and weak barbed bisimilarity, barbed weak coupled simulation equivalence, or (in between) our new correspondence preorder. Iff-results—as the results above—relate these instances of the criterion discussed in [7] with other encodability criteria. In particular, by the results of Section 3, if an encoding satisfies the criterion of [7] w.r.t. (weak) bisimilarity, then it also satisfies operational correspondence w.r.t. (weak) bisimilarity.¹ Moreover, if an encoding satisfies the criterion of [7] w.r.t. correspondence similarity, then it also satisfies weak operational correspondence w.r.t. coupled similarity.²

¹And by the results of Section 4 the bisimulation may be required to (weakly) respect barbs at both sides of the implication.

²We may not conclude that it also satisfies weak operational correspondence w.r.t. correspondence similarity, at least not when also weakly respecting barbs. A counterexample can be found in Figure 3. There we have a weakly barb respecting

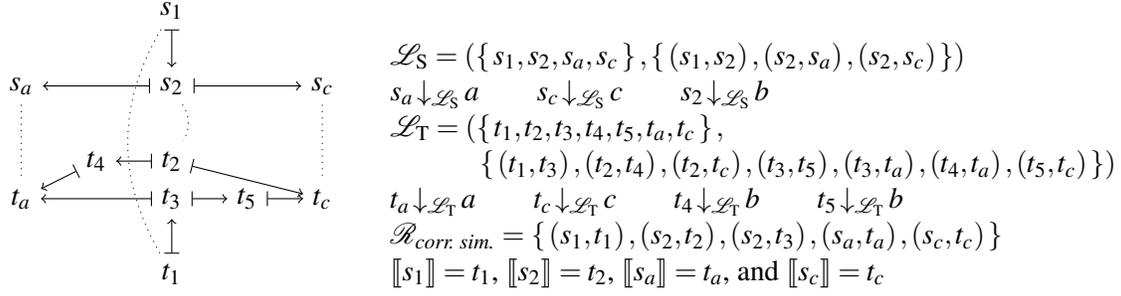


Figure 3: Encoding not satisfying weak operational correspondence w.r.t. to a correspondence simulation that weakly respects barbs, even though $\mathcal{R}_{corr.sim.}$ is a weakly barb respecting correspondence simulation relating each source term with its encoding

The above results may leave the impression that we try to replace common encodability criteria by conditions on relations between source terms and their translations. That is not the case. But we provide alternative ways to prove different criteria. An example of how the above results can be used to reason about the quality of an encoding are the two encodings of [10]. That paper analyses ways to encode the CSP synchronisation mechanism following an approach of [21]. The latter shows that a central encoding of a similar synchronisation mechanism ensures that source terms and their translations behave bisimilar, whereas a decentralised encoding only ensures coupled similarity. Proving coupled similarity can be more difficult than proving bisimilarity. Here our results allow to decrease the proof burden. With Lemma 3.14 we can conclude from weak operational correspondence w.r.t. a bisimulation that source terms and their literal translations are coupled similar, without having to deal with coupled similarity directly. This way we have to deal with the difficult partial commitments, which are introduced by the de-central implementation, only in operational correspondence and not when relating target terms.

In retrospective, mapping encodability criteria on requirements of a relation between source and target terms seems quite natural. Indeed the main challenge of the above presented iff-results was not in proving them but in finding the exact matches between variants of the considered criteria and requirements on the relation. As a consequence we had to define a new kind of simulation relation to capture the version of operational correspondence used in [8].

We do not claim that it is always simple to obtain iff-results as presented in this paper or that we provide a strategy to obtain such results. Instead we claim that proving such results formally captures the effect of a criterion on the quality of an encoding function and thus (1) helps us to understand a criterion, (2) allows to identify unexpectedly strict or weak criteria (3) allows to compare (sets of) criteria, and (4) allows to analyse the side effects that result from the combination of criteria. Analysing criteria this way is not necessarily straightforward. To illustrate this, consider the requirement on the preservation of the (degree of) distribution of a process (preservation of distributability). In the context of asynchronous distributed systems this requirement is very important.

Several attempts to capture it were proposed in the literature. At least for the π -calculus, the most correspondence simulation relating s_2 with t_2 and s_2 with t_3 , but, due to the asymmetric nature of correspondence simulations, there is no weakly barb respecting correspondence simulation relating t_2 and t_3 .³ By Lemma 3.13 the terms t_2 and t_3 are weakly barb respecting coupled similar, however.

³This example does not contradict the weakly barbed variant of Lemma 4.3, for its right-hand side does not hold. Namely, the condition $\forall S. (S, \llbracket S \rrbracket) \in \mathcal{R}_{\llbracket \cdot \rrbracket}$ forces $(s_1, t_1) \in \mathcal{R}_{\llbracket \cdot \rrbracket}$, and thus, since $\mathcal{R}_{\llbracket \cdot \rrbracket}$ is a correspondence simulation, also $(s_2, t) \in \mathcal{R}_{\llbracket \cdot \rrbracket}$ for some $t \in \{t_1, t_3, t_5, t_a, t_c\}$. As s_2 weakly respects barbs a and c , so must t , yielding $t \in \{t_1, t_3\}$. The requirement $(S, T) \in \mathcal{R}_{\llbracket \cdot \rrbracket} \rightarrow (\llbracket S \rrbracket, T) \in \mathcal{R}_T$ yields $(t_2, t) \in \mathcal{R}_{\llbracket \cdot \rrbracket}$, but there exists no correspondence simulation containing this pair.

prominent candidate is the homomorphic translation of the parallel operator ($\llbracket P \mid Q \rrbracket = \llbracket P \rrbracket \mid \llbracket Q \rrbracket$) as used in [17]. [25] shows that this criterion is too strict by providing an encoding that preserves distributability but does not translate the parallel operator homomorphically. Instead this encoding is compositional. Compositionality requires that the translation of an operator is the same for all occurrences of that operator in a term, i.e., it can be captured by a context. Compositionality is significantly weaker than the homomorphic translation of the parallel operator but also forbids the introduction of global coordinators. However, to ensure the preservation of distributability, this criterion is too weak. [25] claims to provide a suitable criterion for the degree of distributability, but without a formal way to reason about the effect of encodability criteria, there is no way to formally prove such a claim. Thus [25] can only provide arguments and illustrations. The inability to formally prove it was one of the original motivations for the present work. Unfortunately, analysing the three criteria compositionality, the homomorphic translation of the parallel operator, and the preservation of distributability is not an easy task.

Compositionality obviously implies some kind of congruence property on encoded source term contexts, but it is not obvious how to turn this observation into an iff-result. To map the homomorphic translation of the parallel operator on conditions on a relation between source and target terms, is even more difficult. This criterion clearly implies some strong properties on such a relation, but it is not clear which condition implies the homomorphic translation of the parallel operator. Because of that, we cannot completely capture the effect of this criterion on the quality of an encoding. This explains why this criterion was originally accepted as a criterion for the preservation of distributability. It is not easy to capture the cases for which it is too strict. The criterion for the preservation of distributability proposed in [25] can intuitively be understood as a concurrency respecting variant of operational correspondence. It not only requires that the source term behaviour is preserved and reflected, but also that the simulations of independent steps are independent. Thus analysing this criterion appears to require some kind of simulation relations that not only consider interleaving semantics. We leave the analysis of these criteria to further research.

All claims in this paper have been proved using the interactive theorem prover Isabelle/HOL. For this purpose, a rich theory of encodability criteria was implemented. Since we do not force any assumptions on process calculi except that they consist of a set of processes, i.e., a type \mathcal{P} , and a reduction relation, i.e., a relation of type $\mathcal{P}^2 = \mathcal{P} \times \mathcal{P}$, this theory can be used to formally reason in Isabelle about encodings for all kinds of source and target languages. A number of well-known process calculi including the π -calculus can for instance be represented in the Psi-calculi framework [1]. Thus there are Isabelle implementations of well-known process calculi that can directly be combined with our Isabelle implementation to formally reason about encodings between such calculi.

References

- [1] J. Bengtson, M. Johansson, J. Parrow & B. Victor (2009): *Psi-calculi: Mobile Processes, Nominal Data, and Logic*. In: *Proceedings of LICS, IEEE*, pp. 39–48, doi:10.1109/LICS.2009.20.
- [2] N. Busi, M. Gabbriellini & G. Zavattaro (2009): *On the expressive power of recursion, replication and iteration in process calculi*. *Mathematical Structures in Computer Science* 19(6), pp. 1191–1222, doi:10.1017/S096012950999017X.
- [3] D. Cacciagrano, F. Corradini, J. Aranda & F.D. Valencia (2008): *Linearity, Persistence and Testing Semantics in the Asynchronous Pi-Calculus*. *Electronic Notes in Theoretical Computer Science* 194(2), pp. 59–84, doi:10.1016/j.entcs.2007.11.006.
- [4] L. Cardelli & A. D. Gordon (2000): *Mobile ambients*. *Theoretical Computer Science* 240(1), pp. 177–213, doi:10.1016/S0304-3975(99)00231-5.

- [5] R. J. van Glabbeek (2001): *The Linear Time – Branching Time Spectrum I: The Semantics of Concrete, Sequential Processes*. *Handbook of Process Algebra*, pp. 3–99, doi:10.1016/B978-044482830-9/50019-9.
- [6] R. J. van Glabbeek (1994): *On the expressiveness of ACP (extended abstract)*. In: *Proceedings of ACP, Workshops in Computing*, Springer, pp. 188–217, doi:10.1007/978-1-4471-2120-6_8.
- [7] R.J. van Glabbeek (2012): *Musings on Encodings and Expressiveness*. In: *Proceedings of EXPRESS/SOS, EPTCS 89*, pp. 81–98, doi:10.4204/EPTCS.89.7.
- [8] D. Gorla (2010): *Towards a unified approach to encodability and separation results for process calculi*. *Information and Computation* 208(9), pp. 1031–1053, doi:10.1016/j.ic.2010.05.002.
- [9] D. Gorla & U. Nestmann (2014): *Full abstraction for expressiveness: history, myths and facts*. *Mathematical Structures in Computer Science*, pp. 1–16, doi:10.1017/S0960129514000279.
- [10] M. Hatzel, C. Wagner, K. Peters & U. Nestmann (2015): *Encoding CSP into CCS*. In: *Proceedings of EXPRESS/SOS, EPTCS 190*, pp. 61–75, doi:10.4204/EPTCS.190.5.
- [11] R. Milner, J. Parrow & D. Walker (1992): *A Calculus of Mobile Processes, Part I and II*. *Information and Computation* 100(1), pp. 1–77, doi:10.1016/0890-5401(92)90008-4.
- [12] J. C. Mitchell (1993): *On abstraction and the expressive power of programming languages*. *Science of Computer Programming* 21(2), pp. 141–163, doi:10.1016/0167-6423(93)90004-9.
- [13] U. Nestmann (2000): *What is a “Good” Encoding of Guarded Choice?* *Information and Computation* 156(1-2), pp. 287–319, doi:10.1006/inco.1999.2822.
- [14] U. Nestmann (2006): *Welcome to the Jungle: A Subjective Guide to Mobile Process Calculi*. In: *Proceedings of CONCUR, LNCS 4137*, Springer, pp. 52–63, doi:10.1007/11817949_4.
- [15] U. Nestmann & B. C. Pierce (2000): *Decoding Choice Encodings*. *Information and Computation* 163(1), pp. 1–59, doi:10.1006/inco.2000.2868.
- [16] T. Nipkow, L. C. Paulson & M. Wenzel (2002): *Isabelle/HOL: a proof assistant for higher-order logic*. 2283, Springer Science & Business Media, doi:10.1007/3-540-45949-9_1.
- [17] C. Palamidessi (2003): *Comparing The Expressive Power Of The Synchronous And Asynchronous Pi-Calculi*. *Mathematical Structures in Computer Science* 13(5), pp. 685–719, doi:10.1017/S0960129503004043.
- [18] C. Palamidessi, V. A. Saraswat, F. D. Valencia & B. Victor (2006): *On the Expressiveness of Linearity vs Persistence in the Asynchronous Pi-Calculus*. In: *Proceedings of LICS, IEEE Computer Society*, pp. 59–68, doi:10.1109/LICS.2006.39.
- [19] J. Parrow (2008): *Expressiveness of Process Algebras*. *Electronic Notes Theoretical Computer Science* 209, pp. 173–186, doi:10.1016/j.entcs.2008.04.011.
- [20] J. Parrow (2014): *General conditions for full abstraction*. *Mathematical Structures in Computer Science*, pp. 1–3, doi:10.1017/S0960129514000280.
- [21] J. Parrow & P. Sjödin (1992): *Multiway synchronization verified with coupled simulation*. In: *Proceedings of CONCUR, LNCS 630*, Springer, pp. 518–533, doi:10.1007/BFb0084813.
- [22] J. A. Perez (2009): *Higher-Order Concurrency: Expressiveness and Decidability Results*. Ph.d. thesis, University of Bologna.
- [23] K. Peters (2012): *Translational Expressiveness*. Ph.D. thesis, TU Berlin. Available at <http://opus.kobv.de/tuberlin/volltexte/2012/3749/>.
- [24] K. Peters & U. Nestmann (2012): *Is it a “Good” Encoding of Mixed Choice?* In: *Proceedings of FOSSACS, LNCS 7213*, Springer, pp. 210–224, doi:10.1007/978-3-642-28729-9_14.
- [25] K. Peters, U. Nestmann & U. Goltz (2013): *On Distributability in Process Calculi*. In: *Proceedings of ESOP, LNCS 7792*, Springer, pp. 310–329, doi:10.1007/978-3-642-37036-6_18.
- [26] J. G. Riecke (1991): *Fully abstract translations between functional languages*. In: *Proceedings of POPL, ACM*, pp. 245–254, doi:10.1145/99583.99617.

Encoding CSP into CCS^{*}

Meike Hatzel
TU Berlin

Christoph Wagner
TU Berlin

Kirstin Peters[†]
TU Dresden

Uwe Nestmann
TU Berlin

We study encodings from CSP into asynchronous CCS with name passing and matching, so in fact, the asynchronous π -calculus. By doing so, we discuss two different ways to map the multi-way synchronisation mechanism of CSP into the two-way synchronisation mechanism of CCS. Both encodings satisfy the criteria of Gorla except for compositionality, as both use an additional top-level context. Following the work of Parrow and Sjödin, the first encoding uses a centralised coordinator and establishes a variant of weak bisimilarity between source terms and their translations. The second encoding is decentralised, and thus more efficient, but ensures only a form of coupled similarity between source terms and their translations.

1 Introduction

In the context of a scientific meeting on Expressiveness in Concurrency and Structural Operational Semantics (SOS), likely very little needs to be said about the process algebras (or process calculi) CSP and CCS. Too many papers have been written since their advent in the 70's to be mentioned in our own paper; it is instructive, though, and recommended to appreciate Jos Baeten's historical overview [1], which also places CSP and CCS in the context of other process algebras like ACP and the many extensions by probabilities, time, mobility, etc. Here, we just select references that help to understand our motivation.

Differences. From the beginning, although CSP [8] and CCS [11] were intended to capture, describe and analyse reactive and interactive concurrent systems, they were designed following rather different philosophies. Tony Hoare described this nicely in his position paper [9] as follows: “A primary goal in the original design of CCS was to discover and codify a minimal set of basic primitive agents and operators . . . and a wide range of useful operators which have been studied subsequently are all definable in terms of CCS primitives.” and “CSP was more interested in this broader range of useful operators, independent of which of them might be selected as primitive.” So, at their heart, the two calculi use two different synchronisation mechanisms, one (CCS) using binary, i.e., two-way, handshake via matching actions and co-actions, the other (CSP) using multiway synchronisation governed by explicit synchronisation sets that are typically attached to parallel composition. Another difference is the focus on Structural Operational Semantics in CCS, and the definition of behavioural equivalences on top of this, while CSP emphasised a trace-based denotational model, enhanced with failures, and the question on how to design models such that they satisfy a given set of laws of equivalence.

Comparisons. From the early days, researchers were interested in more or less formal comparisons between CSP and CCS. This was carried out by both Hoare [9] and Milner [12] themselves, where they concentrate on the differences in the underlying design principles. But also other researchers joined the game, but with different analysis tools and comparison criteria.

For example, Brookes [3] contributed a deep study on the relation between the underlying abstract models, synchronisation trees for CCS and the failures model of CSP. Quite differently, Lanese and Montanari [10] used the power to transform graphs as a measure for the expressiveness of the two calculi.

^{*}Supported by the DFG via project “Synchronous and Asynchronous Interaction in Distributed Systems”.

[†]Supported by the German Federal and State Governments via the Excellence Initiative (Institutional Strategy).

Yet completely differently, Parrow and Sjödin [16, 21] tried to find an algorithm to implement—best in a fully distributed fashion—the multiway synchronisation operator of CSP (and its variant LOTOS [2]) using the supposedly simpler two-way synchronisation of CCS. They came up with two candidates—a reasonably simple centralised synchroniser, and a considerably less simple distributed synchroniser¹—and proved that the two are not weakly bisimilar, but rather coupled similar, which is only slightly weaker. Coupled simulation is a notion that Parrow and Sjödin invented for just this purpose, but it has proved afterwards to be often just the right tool when analysing the correctness of distribution- and divergence-sensitive encodings that involve partial commitments (whose only effect is to gradually perform internal choices) [15].

The probably most recent comparison between CSP and CCS was provided by van Glabbeek [5]. As an example for his general framework to analyse the relative expressive power of calculi, he studied the existence of syntactical translations from CSP into CCS, for which a common semantical domain is provided via labelled transition systems (LTS) derived from respective sets of SOS rules. The comparison is here carried out by checking whether a CSP term and its translation into CCS are distinguishable with respect to a number of equivalences defined on top of the LTS. The concrete results are: (1) there is a translation that is correct up to trace equivalence (and contains deadlocks), and (2) there is no translation that is correct up to weak bisimilarity equivalence that also takes divergence into account.

Contribution. Given van Glabbeek’s negative result, and given Parrow and Sjödin’s algorithm, we set out to check whether we can define a syntactical encoding from CSP into CCS—using Parrow and Sjödin’s ideas—that is correct up to coupled similarity.² We almost managed. In this paper, we report on our current results along these lines: (1) Our encoding target is an asynchronous variant of CCS, but enhanced with name-passing and matching, so it is in fact an asynchronous π -calculus; we kept mentioning CCS in the title of this paper, as it clearly emphasises the origin and motivation of this work. But, we could *not* do without name-passing. (2) We exhibit one encoding that is not distributability-preserving (so, it represents a centralised solution), but is correct up to weak bisimilarity and does not introduce divergence. This does not contradict van Glabbeek’s results, but suggests that van Glabbeek’s framework implies some form of distributability-preservation. (3) We exhibit another encoding that is distributability-preserving and divergence-reflecting, but is only correct up to coupled similarity.

Overview. We introduce the considered variants of CSP and CCS in § 2. There we also introduce the criteria—that are (variants of) the criteria in [6] and [20]—modulo which we prove the quality of the considered encodings. In § 3 we introduce the inner layer of our two encodings. It provides the main machinery to encode synchronisations of CSP. We complete this encoding with an outer layer that is either a centralised (§ 4) or a decentralised coordinator (§ 5). In § 6 we discuss the two encodings. Missing proofs and some additional informations can be found in [7].

2 Technical Preliminaries

A process calculus (\mathcal{P}, \mapsto) consists of a set \mathcal{P} of processes (syntax) and a reduction relation $\mapsto \subseteq \mathcal{P}^2$ (semantics). Let \mathcal{N} be the countably-infinite set of names. $\tau \notin \mathcal{N}$ denotes an internal unobservable action. We use a, b, x, \dots to range over names and P, Q, \dots to range over processes. We use $\alpha, \beta \dots$ to range over $\mathcal{N} \cup \{\tau\}$. \tilde{a} denotes a sequence of names. Let $\text{fn}(P)$ and $\text{bn}(P)$ denote the sets of free names and bound names occurring in P , respectively. Their definitions are completely standard. We use $\sigma, \sigma', \sigma_1, \dots$

¹Recently [4], a slight variant of the protocol behind this algorithm was used to implement the distributed compiler DLC for a substantial subset of LNT (successor of LOTOS New Technology) that yields reasonably efficient C code.

²The idea and a first draft of the encoding were developed by Nestmann and van Glabbeek during a stay at NICTA, Sydney.

to range over substitutions. A substitution is a mapping $[x_1/y_1, \dots, x_n/y_n]$ from names to names. The application of a substitution on a term $P[x_1/y_1, \dots, x_n/y_n]$ is defined as the result of simultaneously replacing all free occurrences of y_i by x_i for $i \in \{1, \dots, n\}$. For all names in $\mathcal{N} \setminus \{y_1, \dots, y_n\}$ the substitution behaves as the identity mapping. The relation \mapsto as defined in the semantics below defines the reduction steps processes can perform. We write $P \mapsto P'$ if $(P, P') \in \mapsto$ and call P' a *derivative* of P . Let \Longrightarrow denote the reflexive and transitive closure of \mapsto . P is *divergent* if it has an infinite sequence of steps $P \mapsto^\omega$. We use *barbs* or *observables* to distinguish between processes with different behaviours. We write $P \downarrow_\alpha$ if P has a barb α , where the predicate $\cdot \downarrow_\alpha$ can be defined differently for each calculus. Moreover P has a weak barb α , if P may reach a process with this barb, i.e., $P \Downarrow_\alpha \triangleq \exists P'. P \Longrightarrow P' \wedge P' \downarrow_\alpha$.

As source calculus we use the following variant of CSP [8].

Definition 1. The processes \mathcal{P}_{CSP} are given by

$$P ::= P \parallel_A P \quad | \quad \text{DIV} \quad | \quad \text{STOP} \quad | \quad P \sqcap P \quad | \quad P/b \quad | \quad f(P) \quad | \quad X \quad | \quad \mu X \cdot P \quad | \quad \sum_{i \in \mathcal{I}} a_i \rightarrow P$$

where $X \in \mathcal{X}$ is a process variable, $A \subseteq \mathcal{N}$, and \mathcal{I} is a finite index set.

$P \parallel_A Q$ is the parallel composition of P and Q , where P and Q can proceed independently except for actions $a \in A$, on which they have to synchronise. **DIV** describes *divergence*. **STOP** denotes *inaction*. **Internal choice** $P \sqcap Q$ reduces to either P or Q within a single internal step. **Concealment** P/b hides an action b and masks it as τ . **Renaming** $f(P)$ for some $f: \mathcal{N} \rightarrow \mathcal{N}$ extended by $f(\tau) = \tau$ behaves as P , where a is replaced by $f(a)$ for all $a \in \mathcal{N}$. **Recursion** $\mu X \cdot P$ describes a process behaving like P with every occurrence of X being replaced by $\mu X \cdot P$. **External choice** $\sum_{i \in \mathcal{I}} a_i \rightarrow P_i$ offers a selection of one of the *action prefixes* $a_i \rightarrow \cdot$ followed by the corresponding continuation P_i , so it may perform any a_i and then behave like P_i . Note that we enforce action prefixes to be syntactically part of an external choice construct. As usual, we use $M \sqcap N$ to denote binary external choice.

The CSP semantics is given by the following rules, using labelled steps $\xrightarrow{\alpha}$ to define \mapsto :

$\frac{E \xrightarrow{b} E'}{E/b \xrightarrow{\tau} E'/b}$	$\frac{E \xrightarrow{\alpha} E' \quad (\alpha \neq b)}{E/b \xrightarrow{\alpha} E'/b}$	$\frac{E \xrightarrow{\alpha} E'}{f(E) \xrightarrow{f(\alpha)} f(E')}$	$\frac{M_j \xrightarrow{a} M'_j \quad (j \in \mathcal{I})}{\sum_{i \in \mathcal{I}} M_i \xrightarrow{a} M'_i}$
$(a \rightarrow E) \xrightarrow{a} E$		$\mu X \cdot E \xrightarrow{\tau} E[\mu X \cdot E/X]$	
$\frac{E \xrightarrow{\alpha} E' \quad (\alpha \notin A)}{E \parallel_A F \xrightarrow{\alpha} E' \parallel_A F}$	$\frac{F \xrightarrow{\alpha} F' \quad (\alpha \notin A)}{E \parallel_A F \xrightarrow{\alpha} E \parallel_A F'}$	$\frac{E \xrightarrow{a} E' \quad F \xrightarrow{a} F' \quad (a \in A)}{E \parallel_A F \xrightarrow{a} E' \parallel_A F'}$	$\frac{P \xrightarrow{\tau} P'}{P \mapsto P'}$
$\text{DIV} \xrightarrow{\tau} \text{DIV}$		$E \sqcap F \xrightarrow{\tau} E$	$E \sqcap F \xrightarrow{\tau} F$

A barb of CSP is the possibility of a term, to perform an action, i.e., $P \downarrow_a \triangleq \exists P'. P \xrightarrow{a} P'$. Following the definition of distributability in [20] a CSP term P is distributable into P_1, \dots, P_n if P_1, \dots, P_n are unguarded subterms of P such that every action prefix in P occurs in exactly one of the P_1, \dots, P_n , where different but equally-named action prefixes are distinguished and unguarded occurrences of $\mu X \cdot P'$ may result in several copies of P' within the P_1, \dots, P_n .

As target calculus we use an asynchronous variant of CCS [11] with name-passing and matching.

Definition 2. The processes \mathcal{P}_{CCS} are given by

$$P ::= P \mid P \quad | \quad (\nu \tilde{c})P \quad | \quad * \underline{c}(\tilde{x}).P \quad | \quad \underline{c}(\tilde{x}).P \quad | \quad \bar{c}(\tilde{x}) \quad | \quad [c = z]P \quad | \quad \mathbf{0}$$

$P \mid Q$ is the parallel composition of P and Q , where P and Q can either proceed independently or synchronise on matching channels names. $(\nu\tilde{c})P$ restricts the visibility of actions using names in \tilde{c} to P . $\underline{c}(\tilde{x}).P$ denotes input on channel c . $\bar{c}(\tilde{x})$ is output on channel c . Since there is no continuation, we interpret this calculus as asynchronous. We use $*\underline{c}(\tilde{x}).P$ to denote *replicated input* on channel c with the continuation P . $[x = y]P$ is the matching operator, if $x = y$ then P is enabled. $\mathbf{0}$ denotes inaction.

The CCS semantics is given by following transition rules:

$$\boxed{\begin{array}{c} \frac{P \mapsto P'}{P \mid Q \mapsto P' \mid Q} \quad \frac{P \mapsto P'}{(\nu\tilde{c})P \mapsto (\nu\tilde{c})P'} \quad \frac{P \equiv P' \quad P' \mapsto Q' \quad Q' \equiv Q}{P \mapsto Q} \\ * \underline{c}(\tilde{x}).P \mid \bar{c}(\tilde{y}) \mapsto * \underline{c}(\tilde{x}).P \mid P[\tilde{y}/\tilde{x}] \quad \bar{c}(\tilde{y}) \mid \underline{c}(\tilde{x}).Q \mapsto P \mid Q[\tilde{y}/\tilde{x}] \end{array}}$$

where \equiv denotes structural congruence given by the rules: $P \mid \mathbf{0} \equiv P$, $P \mid Q \equiv Q \mid P$, $P \mid (Q \mid R) \equiv (P \mid Q) \mid R$, $(\nu\tilde{a})\mathbf{0} \equiv \mathbf{0}$, $P \mid (\nu\tilde{a})Q \equiv (\nu\tilde{a})(P \mid Q)$ if $\text{bn}(\tilde{a}) \notin \text{fn}(P)$, and $[x = x]P \equiv P$. As discussed in [20], a CCS term P is distributable into P_1, \dots, P_n if $P \equiv (\nu\tilde{x})(P_1 \mid \dots \mid P_n)$.

Simulation Relations. The semantics of a process is usually considered modulo some behavioural equivalence. For many calculi, *the* standard reference equivalence is some form of weak bisimilarity. In the context of encodings, the source and target language often differ in their relevant observables, i.e., barbs. In this case, it is advantageous to use a variant of reduction bisimilarity. With Gorla [6], we add a *success* operator \checkmark to the syntax of both CSP and CCS. Since \checkmark cannot be further reduced, the semantics is left unchanged in both cases. The test for the reachability of success is standard in both languages, i.e., $P \downarrow_{\checkmark} \triangleq \exists P'. P \equiv \checkmark \mid P'$. To obtain a non-trivial equivalence, we require that the bisimulation respects success and the reachability of barbs. We use the standard definition of barbs in CSP, i.e., action prefixes. Our encoding function will translate all source terms into closed terms, thus the standard definition of CCS barbs would not provide any information. Instead we use a notion of translated barb ($\cdot \downarrow_{\llbracket \cdot \rrbracket}$) that reflects how the encoding function translates source term barbs. Its definition is given in Section 3.

Definition 3 (Bisimulation). A relation $\mathcal{R} \subseteq \mathcal{P}^2$ is a (*success-sensitive, [translated-]barb-respecting, weak, reduction*) *bisimulation* if, whenever $(P, Q) \in \mathcal{R}$, then:

- $P \mapsto P'$ implies $\exists Q'. Q \Longrightarrow Q' \wedge (P', Q') \in \mathcal{R}$
- $Q \mapsto Q'$ implies $\exists P'. P \Longrightarrow P' \wedge (P', Q') \in \mathcal{R}$
- $P \downarrow_{\checkmark}$ iff $Q \downarrow_{\checkmark}$
- P and Q reach the same (translated) barbs, where we use $\cdot \downarrow_a$ for CSP and $\cdot \downarrow_{\llbracket a \rrbracket}$ for CCS

Two terms $P, Q \in \mathcal{P}$ are *bisimilar*, denoted as $P \dot{\approx} Q$, if there exists a bisimulation that relates P and Q .

We use the symbol $\dot{\approx}$ to denote either bisimilarity on our target language CCS or on the disjoint union of CSP and CCS that allows us to describe the relationship between source terms and their translations. In the same way we define a corresponding variant of coupled similarity.

Definition 4 (Coupled Simulation). A relation $\mathcal{R} \subseteq \mathcal{P}^2$ is a (*success-sensitive, [translated-]barb-respecting, weak, reduction*) *coupled simulation* if, whenever $(P, Q) \in \mathcal{R}$, then:

- $P \mapsto P'$ implies $\exists Q'. Q \Longrightarrow Q' \wedge (P', Q') \in \mathcal{R}$ and $\exists Q''. Q \Longrightarrow Q'' \wedge (Q'', P') \in \mathcal{R}$
- $P \downarrow_{\checkmark}$ iff $Q \downarrow_{\checkmark}$
- P and Q reach the same (translated) barbs, where we use $\cdot \downarrow_a$ for CSP and $\cdot \downarrow_{\llbracket a \rrbracket}$ for CCS

Two terms $P, Q \in \mathcal{P}$ are *coupled similar*, denoted as $P \dot{\approx}_{\text{cs}} Q$, if there exists a coupled simulation that relates P and Q in both directions.

Encodings and Quality Criteria. We consider two different translations from (the above-defined variant of) CSP into (the above-defined variant of) CCS with name passing and matching. In this context, we

refer to CSP terms as *source terms* \mathcal{P}_S and to CCS terms as *target terms* \mathcal{P}_T . Encodings often translate single source steps into a sequence or pomset of target steps. We call such a sequence or pomset a *simulation* of the corresponding source term step. Moreover, we assume for each encoding the existence of a so-called renaming policy φ , i.e., a mapping of names from the source into vectors of target term names.

To analyse the quality of encodings and to rule out trivial or meaningless encodings, Gorla [6] provide a general framework comprising five quality criteria, which have afterwards been used in many papers. In addition to our above-mentioned definition of process calculus, whough, Gorla requires the target calculus to be equipped with a notion of behavioural equivalence \asymp on target terms. Its purpose is to describe the ‘abstract’ behaviour of a target process, where ‘abstract’ refers to an observer at the source level. In [6], the equivalence \asymp is often defined as a barbed equivalence (cf. [13]) or can be derived directly from the reduction semantics, and it typically is a congruence, at least with respect to parallel composition. Bisimilarity and coupled similarity are such relations on CCS terms. The criteria are:

- (1) *Compositionality*: The translation of an operator op is the same for all occurrences of that operator in a term, i.e., it can be captured by a context \mathcal{C}_{op} such that $\text{enc}(\text{op}(x_1, \dots, x_n, S_1, \dots, S_m)) = \mathcal{C}_{\text{op}}^N(x_1, \dots, x_n, \text{enc}(S_1), \dots, \text{enc}(S_m))$ for $\text{fn}(S_1) \cup \dots \cup \text{fn}(S_m) = N$.
- (2) *Name Invariance*: The encoding does not depend on particular names, i.e., for every S and σ , it holds that $\text{enc}(\sigma(S)) \equiv \sigma'(\text{enc}(S))$ if σ is injective and $\text{enc}(\sigma(S)) \asymp \sigma'(\text{enc}(S))$ otherwise, where σ' is such that $\varphi(\sigma(n)) = \sigma'(\varphi(n))$ for every $n \in \mathcal{N}$.
- (3) *Operational Correspondence*: Every computation of a source term can be simulated by its translation, i.e., $S \Longrightarrow_S S'$ implies $\text{enc}(S) \Longrightarrow_T \asymp \text{enc}(S')$ (completeness), and every computation of a target term corresponds to some computation of the corresponding source term (soundness, compare to Section 5).
- (4) *Divergence Reflection*: The encoding does not introduce divergence, i.e., $\text{enc}(S) \dashv\vdash_T^\omega$ implies $S \dashv\vdash_S^\omega$.
- (5) *Success Sensitiveness*: A source term and its encoding answer the tests for success in exactly the same way, i.e., $S \Downarrow_\checkmark$ iff $\text{enc}(S) \Downarrow_\checkmark$.

Our encodings will satisfy all of these criteria except for compositionality, because both encodings consists of two layers. [20] shows that the above criteria do not ensure that an encoding preserves distribution and proposes an additional criterion for the preservation of distributability.

Definition 5 (Preservation of Distributability). An encoding $\text{enc}(\cdot)$ *preserves distributability* if for every S and for all terms S_1, \dots, S_n that are distributable within S there are some T_1, \dots, T_n that are distributable within $\text{enc}(S)$ such that $T_i \asymp \text{enc}(S_i)$ for all $1 \leq i \leq n$.

Here, because of the choice of the source and the target language, an encoding preserves distributability if for each sequence of distributable source term steps their simulations are pairwise distributable. In both languages two alternative steps of a term are in *conflict* with each other if—for CSP—they reduce the same action-prefix or—for CCS—they either reduce the same input using two outputs or they reduce the same output using two [replicated] inputs. Two alternative steps that are not in conflict are *distributable*.

3 Translating the CSP Synchronisation Mechanism

CSP and CCS—or the π -calculus—differ fundamentally in their communication and synchronisation mechanisms. In CSP there is only a single kind of action $c \rightarrow \cdot$, where c is a name. Synchronisation is implemented by the parallel operator $\cdot \|_A \cdot$ that in CSP is augmented with a set of names A containing the names that need to be synchronised at this point. By nesting parallel operators arbitrarily many actions on the same name can be synchronised. In CCS there are two different kinds of actions: inputs \underline{c} and

outputs \bar{c} . Again synchronisation is implemented by the parallel operator, but in CCS only a single input and a single matching output can ever be synchronised within one step.

To encode the CSP communication and synchronisation mechanisms in CCS with name passing we make use of a technique already used in [17, 19] to translate between different variants of the π -calculus. CSP actions are translated into action announcements augmented with a lock indicating, whether the respective action was already used in the simulation of a step. The other operators of CSP are then translated into handlers for these announcements and locks. The translation of sum combines several actions under the same lock and thus ensures that only one term of the sum can ever be used. The translation of the parallel operator combines announcements of actions that need to be synchronised into a single announcement under a fresh lock, whose value is determined by the combination of the respective underlying locks at its left and right side. Announcements of actions that do not need to be synchronised are simply forwarded. A second layer—containing either a centralised or a decentralised coordinator—then triggers and coordinates the simulation of source term steps.

Action announcements are of the form $\bar{a}\langle c, r, l, r' \rangle$: c is the translation of the source term action. r is used to trigger the computation of the Boolean value of l . The lock l evaluates to \top as long as the respective translated action was not successfully used in the simulation of a step. r' is used to guard the encoded continuation of the respective source term action. In the case of a successful simulation attempt involving this announcement, an output $\bar{r}'\langle \top \rangle$ allows to unguard the encoded source term continuation and ensures that all following evaluations of l return \perp . The message $\bar{r}'\langle \perp \rangle$ indicates an aborted simulation attempt and allows to restore l for later simulation attempts. Once a lock becomes \perp , all request for its computation return \perp .

Abbreviations. We introduce some abbreviations to simplify the presentation of the encodings. We use

$$[x \in A]P \triangleq \prod_{a \in A} [x = a]P$$

to test, whether an action belongs to the set of synchronised actions in the encoding of the parallel operator. As already done in [14, 15] we use Boolean-valued locks to ensure that every translation of an action is only used once to simulate a step. *Boolean locks* are channels on which only the Boolean values \top (true) or \perp (false) are transmitted. An unguarded output over a Boolean lock with value \top represents a positive instantiation of the respective lock, whereas an unguarded output sending \perp represents a negative instantiation. At the receiving end of such a channel, the Boolean value can be used to make a binary decision, which is done here within an *IF-construct*. This construct and according instantiations of locks are implemented as in [14, 15] using restriction and the order of transmitted values.

$$\bar{l}\langle \top \rangle \triangleq \underline{l}(t, f).\bar{t} \quad \bar{l}\langle \perp \rangle \triangleq \underline{l}(t, f).\bar{f}$$

$$\underline{l}(b).\text{IF } b \text{ THEN } P \text{ ELSE } Q \triangleq (\nu t, f)(\bar{l}\langle t, f \rangle \mid \underline{t}P \mid \underline{f}Q)$$

We observe that the Boolean values \top and \perp are realised by a pair of links without parameters. Both cases of the IF-construct operate as guard for its subterms P and Q . The renaming policy φ reserves the names t and f to implement the Boolean values \top and \perp .

The Algorithm. The encoding functions introduce some fresh names, that are reserved for special purposes. In Table 1 we list the reserved names \mathcal{R} and provide a hint on their purpose. Moreover we reserve the names $\{x_i \mid i \in \mathbb{N}\}$ and assume an injective mapping $\varphi' : \mathcal{X} \rightarrow \{x_i \mid i \in \mathbb{N}\}$ that maps process variables of CSP to distinct names. The renaming policy φ for our encodings is then a function that reserves the names in $\mathcal{R} \cup \{x_i \mid i \in \mathbb{N}\}$ and translates every source term name into three target term names. More precisely, choose $\varphi : \mathcal{N} \rightarrow \mathcal{N}^3$ such that:

1. No name is mapped onto a reserved name, i.e., $\varphi(n) \cap (\mathcal{R} \cup \{x_i \mid i \in \mathbb{N}\}) = \emptyset$ for all $n \in \mathcal{N}$.
2. No two different names are mapped to overlapping sets of names, i.e., $\varphi(n) \cap \varphi(m) = \emptyset$ for all $n, m \in \mathcal{N}$ with $n \neq m$.

reserved names	purpose
a, a'	announce the ability to perform an action
c, c _L , c _R , z	(translated) source term channel, channel from the left/right of a parallel operator
l, l _L , l _R	lock, lock from the left/right of a parallel operator
l'	re-instantiate a positive sum lock
r, r _L , r _R	request the computation of the value of a lock
r', r' _i , r' _L , r' _R	simulate a source term step and unguard the corresponding continuations
n	order left announcements for the same channel that need to be synchronised
s, s'	distribute right announcements that need to be synchronised
b	Boolean value (\perp or \top)
τ	fresh name used to announce τ -steps that result from concealment
once	used by the centralised encoding to avoid overlapping simulation attempts
m	fresh names used to encode internal choice
d	fresh names used to encode divergence
t, f	used to encode Boolean values

Table 1: Reserved Names.

We naturally extend the renaming policy to sets of names, i.e., $\varphi(X) \triangleq \{ \varphi(x) \mid x \in X \}$ if $X \subseteq \mathcal{N}$. Let $((x_1, \dots, x_n)).i \triangleq x_i$ denote the projection of a n -tuple to its i th element, if $1 \leq i \leq n$. Moreover $(X).i \triangleq \{ (x).i \mid x \in X \}$ for a set X of n -tuples and $1 \leq i \leq n$.

The inner part of our two encodings is presented in Figure 1. The most complex case is the translation of the parallel operator $\llbracket P \parallel_A Q \rrbracket$ that is based on the following four steps:

Step 1: Action announcements for channels $c \notin A$

In the case of actions on channels $c \notin A$ —that do not need to be synchronised here—the encoding of the parallel operator acts like a forwarder and transfers action announcements of both its subtrees further up in the parallel tree. Two different restrictions of the channel for action announcements a from the left side $\llbracket P \rrbracket$ and the right side $\llbracket Q \rrbracket$, allow to trace action announcements back to their origin as it is necessary in the following case. In the present case we use a' to bridge the action announcement over the restrictions on a .

Step 2: Action announcements for channels $c \in A$

Actions $c \in A$ need to be synchronised, i.e., can be performed only if both sides of the parallel operator cooperate on this action. Simulating this kind of synchronisation is the main purpose of the encoding of the parallel operator. The renaming policy φ translates each source term name into three target term names. The first target term name is used as reference to the original source term name and transferred in announcements. The other two names are used to simulate the synchronisation of the parallel operator in CSP. Announcements from the left are translated to outputs on the respective second name and announcements from the right to the respective third name. Restriction ensures that these outputs can only be computed by the current parallel operator encoding. The translations of the announcements into different outputs for different source term names allows us to treat announcements of different names concurrently using the term $\text{Synch}(c)$, where c is a source term name.

Step 3: The term $\text{Synch}(c)$

In $\text{Synch}(c)$ all announcements for the same source term name c from the left are ordered in order to combine each left and each right announcement on the same name. Several such announcements

$$\begin{aligned}
\llbracket P \rrbracket_A Q &\triangleq (\nu a', (\varphi(A)).2, (\varphi(A)).3) \left(\right. \\
&\quad (\nu a) \left(\llbracket P \rrbracket \mid *a(c, \tilde{x}). \left([c \in (\varphi(A)).1] \overline{(\varphi(c)).2}(\tilde{x}) \mid [c \notin (\varphi(A)).1] \overline{a'}(c, \tilde{x}) \right) \right) \\
&\quad (\nu a) \left(\llbracket Q \rrbracket \mid *a(c, \tilde{x}). \left([c \in (\varphi(A)).1] \overline{(\varphi(c)).3}(\tilde{x}) \mid [c \notin (\varphi(A)).1] \overline{a'}(c, \tilde{x}) \right) \right) \\
&\quad \mid \prod_{c \in A} \text{Synch}(c) \mid *a'(\tilde{x}).\overline{a}(\tilde{x}) \left. \right) \\
\text{Synch}(c) &\triangleq (\nu n) \left(\overline{n}(\langle (\varphi(c)).3 \rangle \right. \\
&\quad \mid *n(s) \left((\varphi(c)).2(r_L, l_L, r'_L). \left((\nu s') \left(\right. \right. \right. \\
&\quad \quad *s(r_R, l_R, r'_R). \left((\nu r, l, r') \left(\overline{a}(\langle (\varphi(c)).1, r, l, r' \rangle \mid \text{Sim}) \mid \overline{s'}(r_R, l_R, r'_R) \right) \right) \right) \\
&\quad \quad \mid (\nu s) \left(\overline{n}(s) \mid *s'(\tilde{x}).\overline{s}(\tilde{x}) \right) \left. \right) \left. \right) \\
\text{Sim} &\triangleq (\nu l') \left(\overline{l'} \mid *l'. \left(\underline{r}. \left(\overline{r_L} \mid \underline{l_L}(b). \left(\text{IF } b \text{ THEN } \left(\overline{r_R} \mid \underline{l_R}(b). \left(\text{IF } b \right. \right. \right. \right. \right. \right. \\
&\quad \quad \text{THEN } \left(\overline{l}(\top) \mid \underline{r}'(b). \left(\overline{r'_L}(b) \mid \overline{r'_R}(b) \mid \text{IF } b \text{ THEN } *r.\overline{l}(\perp) \text{ ELSE } \overline{l'} \right) \right) \\
&\quad \quad \text{ELSE } \left(\overline{l}(\perp) \mid \overline{r'_L}(\perp) \mid *r.\overline{l}(\perp) \right) \left. \right) \left. \right) \left. \right) \\
&\quad \text{ELSE } \left(\overline{l}(\perp) \mid *r.\overline{l}(\perp) \right) \left. \right) \left. \right) \\
\llbracket \sum_{i \in \mathcal{I}} c_i \rightarrow P_i \rrbracket &\triangleq (\nu r, l, r'_1, \dots, r'_n) \left(\underline{r}.\overline{l}(\top) \right. \\
&\quad \mid \prod_{i \in \mathcal{I}} \left(\overline{a}(\langle c_i \rangle.1, r, l, r'_i) \mid *r'_i(b). \text{IF } b \text{ THEN } (\llbracket P_i \rrbracket \mid *r.\overline{l}(\perp)) \text{ ELSE } \underline{r}.\overline{l}(\top) \right) \left. \right) \\
\llbracket (P) / z \rrbracket &\triangleq (\nu a') \left((\nu a, z) \left(\llbracket P \rrbracket \mid *a(c, \tilde{x}). \left([c = z] \overline{a'}(\tau, \tilde{x}) \mid [c \neq z] \overline{a'}(c, \tilde{x}) \right) \mid *a'(\tilde{x}).\overline{a}(\tilde{x}) \right) \right) \\
\llbracket f(P) \rrbracket &\triangleq (\nu a') \left((\nu a, z) \left(\llbracket P \rrbracket \mid *a(c, \tilde{x}). \left(\prod_{z/x \in f} [c = (\varphi(x)).1] \overline{a'}(\langle (\varphi(z)).1, \tilde{x} \rangle \right. \right. \right. \\
&\quad \quad \left. \left. \mid [c \notin \text{dom}(f)] \overline{a'}(c, \tilde{x}) \right) \mid *a'(\tilde{x}).\overline{a}(\tilde{x}) \right) \right) \\
\llbracket \text{DIV} \rrbracket &\triangleq (\nu d) \left(\overline{d} \mid *d.\overline{d} \right) \\
\llbracket \mu X \cdot P \rrbracket &\triangleq (\nu \varphi'(X)) \left(\overline{\varphi'(X)} \mid * \underline{\varphi'(X)}. \llbracket P \rrbracket \right) \\
\llbracket X \rrbracket &\triangleq \overline{\varphi'(X)} \\
\llbracket P \sqcap Q \rrbracket &\triangleq (\nu m) \left(\underline{m}. \llbracket P \rrbracket \mid \underline{m}. \llbracket Q \rrbracket \mid \overline{m} \right) \\
\llbracket \text{STOP} \rrbracket &\triangleq \mathbf{0} \\
\llbracket \checkmark \rrbracket &\triangleq \checkmark
\end{aligned}$$

where $\notin (\varphi(A)).1$ is short for $\in (\text{fn}(P) \cup \text{fn}(Q)) \setminus (\varphi(A)).1$, $\notin \text{dom}(f)$ is short for $\in \text{fn}(P) \setminus \text{dom}(f)$, and $\neq z$ is short for $\in \text{fn}(P) \setminus \{z\}$.

Figure 1: An encoding from CSP into CCS with value passing (inner part).

may result from underlying parallel operators, sums with similar summands, and junk left over from already simulated source term steps. For each left announcement a fresh instance of s is generated and restricted. The names s and s' are used to transfer right announcements to the respective next left announcement, where s' is used to bridge over the restriction on s . This way each right announcement will eventually be transferred to each left announcement on the same name. Note that this kind of forwarding is not done concurrently but in the source language a term $P\|_A Q$ also cannot perform two steps on the same name $c \in A$ concurrently. After combining a left and a right announcement on the same source term name a fresh set of auxiliary variables r, l, r' is generated and a corresponding announcement is transmitted. The term Sim reacts to requests regarding this announcement and is used to simulate a step on the synchronised action.

Step 4: The term Sim

If a request reaches Sim it starts questioning the left and the right side. First the left side is requested to compute the current value of the lock of the action. Only if \top is returned, the right side is requested to compute its lock as well. This avoids deadlocks that would result from blindly requesting the computation of locks in the decentralised encoding. If the locks of both sides are still valid the fresh lock l returns \top else \perp is returned. For each case Sim ensures that subsequently requests will obtain an answer by looping with \bar{l} or returning \perp to all requests, respectively. The messages $\bar{r}'_L(\perp)$ and $\bar{r}'_R(\perp)$ cause the respective underlying subterms on the left and the right side to do the same, whereas $\bar{r}'_L(\top)$ and $\bar{r}'_R(\top)$ cause the unguarding of encoded continuations as result of a successful simulation of a source term synchronisation step.

Basic Properties and Translated Observables. The protocol introduced by the encoding function in Figure 1 (and its outer parts introduced later) simulates a single source term step by a sequence of target term steps. Most of these steps are merely pre- and post-processing steps, i.e., they do not participate in decisions regarding the simulation of conflicting source term steps but only prepare and complete simulations. Accordingly we distinguish between *auxiliary steps*—that are pre- and post-processing steps—and *simulation steps*—that mark a point of no return by deciding which source term step is simulated. Note that the points of no return and thus the definition of auxiliary and simulation steps is different in the two variants of our encoding.

Auxiliary steps do not influence the choice of source terms steps that are simulated. Moreover they operate on restricted channels, i.e., are unobservable. Accordingly they do not change the state of the target term modulo the considered reference relations $\dot{\approx}$ and $\dot{\approx}_{\text{cs}}$. We introduce some auxiliary lemmata to support this claim.

The encoding $\llbracket \cdot \rrbracket$ translates source term barbs c into free announcements with $(\varphi(c)).l$ as first value and a lock l as third value that computes to \top . The two coordinators, i.e., outer encodings, we introduce later, restrict the free a -channel of $\llbracket \cdot \rrbracket$.

Definition 6 (Translated Barbs). Let $T \in \mathcal{P}_T$ such that $\exists S. \llbracket S \rrbracket \Longrightarrow_T T$, $\exists S. \llbracket S \rrbracket \Longrightarrow_T T$, or $\exists S. (\|S\|) \Longrightarrow_T T$. T has a translated barb c , denoted by $T \downarrow_{\llbracket c \rrbracket}$, if

- there is an unguarded output $\bar{a}((\varphi(c)).l, r, l, r')$ —on a free channel a in the case of $\llbracket \cdot \rrbracket$ or the outermost variant of a in the case of the later introduced encodings $\llbracket \cdot \rrbracket$ and (\cdot) —in T
- such an announcement was consumed to unguard an IF-construct testing l and this construct is still not resolved in T

such that all locks that are necessary to instantiate l are positively instantiated.

Analysing the encoding function in Figure 1 we observe that an encoded source term has a translated barb iff the corresponding source term has the corresponding source term barb.

$$\llbracket P \rrbracket \triangleq (\text{va, once})(\llbracket P \rrbracket \mid \overline{\text{once}} \mid * \text{once.a}(c, r, l, r').(\bar{r} \mid \perp(b).(\overline{\text{once}} \mid \text{IF } b \text{ THEN } \bar{r}'(\top))))$$

Figure 2: A **centralised** encoding from CSP into CCS with value passing.

Observation 7. For all $S \in \mathcal{P}_S$, it holds $S \downarrow_c$ iff $\llbracket S \rrbracket \downarrow_{\llbracket c \rrbracket}$.

All instances of success in the translation result from success in the source. More precisely the only way to obtain \checkmark in the translation is by $\llbracket \checkmark \rrbracket \triangleq \checkmark$.

Observation 8. For all $S \in \mathcal{P}_S$, it holds $S \downarrow_{\checkmark}$ iff $\llbracket S \rrbracket \downarrow_{\checkmark}$.

The encoding propagates announcements through the translated parallel structure. In the translation of parallel operators it combines all left and right announcements w.r.t. to the same channel name, if this channel needs to be synchronised. Therefore we copy announcements. We use locks carrying a Boolean value to indicate whether an announcement was already used to simulate a source term step. These locks carry \top in the beginning and are swapped to \perp as soon as the announcement was used. In each state there is at most one positive instantiation of each lock and as soon as a lock is instantiated negatively it never becomes positive again.

Lemma 9. Let $T \in \mathcal{P}_T$ such that $\exists S. \llbracket S \rrbracket \Longrightarrow_T T$. Then for each variant l of the names l, l_L, l_R

1. there is at most one positive instantiation of l in T ,
2. if there is a positive instantiation of l in T then there is no other instantiation of l in T ,
3. if there is a negative instantiation of l in T then no derivative of T contains a positive instantiation of l .

4 The Centralised Encoding

Figure 1 describes how to translate CSP actions into announcements augmented with locks and how the other operators are translated to either forward or combine these announcements and locks. With that $\llbracket \cdot \rrbracket$ provides the basic machinery of our encoding from CSP into CCS with name passing and matching. However it does not allow to simulate any source term step. Therefore we need a second (outer) layer that triggers and coordinates the simulation of source term steps. We consider two ways to implement this coordinator: a centralised and a decentralised coordinator. The centralised coordinator is depicted in Figure 2.

The channel `once` is used to ensure that simulation attempts of different source term steps cannot overlap each other. For each simulation attempt exactly one announcement is consumed. The coordinator then triggers the computation of the respective lock that was transmitted in the announcement. This request for the computation of the lock is propagated along the parallel structure induced by the translations of parallel operators until—in the leafs—encodings of sums are reached. There the request for the computation yields the transmission of the current value of the respective lock. While being transmitted back to the top of the tree, different locks that refer to synchronisation in the source terms are combined. If the computation of the lock results with \top at the top of the tree, the respective source term step is simulated. Else the encoding aborts the simulation attempt and restores the consumed informations about the values of the respective locks. In both cases a new instance of $\overline{\text{once}}$ allows to start the next simulation attempt. Accordingly only some post-processing steps can overlap with a new simulation attempt.

As we prove below, the points of no return in the centralised encoding can result from the consumption of action announcements by the outer encoding in Figure 2 if the corresponding lock computes to

\top . Moreover the encoding of internal choice and divergence introduces simulation steps, namely all steps on variants of the channels m , d , and $\varphi'(X)$. All remaining steps of the centralised encoding are auxiliary.

Definition 10 (Auxiliary and Simulation Steps). A step $T \mapsto_{\top} T'$ such that $\exists S \in \mathcal{P}_{\mathcal{S}}. \llbracket S \rrbracket \Longrightarrow_{\top} T$ is called a *simulation step*, denoted by $T \xrightarrow{\top} T'$, if $T \mapsto_{\top} T'$ is a step on the outermost channel a and the computation of the value of the received lock l will return \top or it is a step on a variant of m , d , or $\varphi'(X)$.

Else the step $T \mapsto_{\top} T'$ is called an *auxiliary step*, denoted by $T \dot{\mapsto} T'$.

Let \Longrightarrow denote the reflexive and transitive closure of $\dot{\mapsto}$ and let $\xrightarrow{\top} \triangleq \Longrightarrow \dot{\mapsto} \xrightarrow{\top} \Longrightarrow$. Auxiliary steps do not change the state modulo $\dot{\approx}$.

Lemma 11. $T \dot{\mapsto} T'$ implies $T \dot{\approx} T'$ for all target terms T, T' .

By distinguishing auxiliary and simulation steps, we can prove a condition stronger than operational correspondence, namely that each source term step is simulated by exactly one simulation step.

Lemma 12. For all S, S' , it holds $S \mapsto_{\mathcal{S}} S'$ iff $\exists T. \llbracket S \rrbracket \xrightarrow{\top} T \wedge \llbracket S' \rrbracket \dot{\approx} T$.

This direct correspondence between source term steps and the points of no return of their translation allows us to prove a variant of operational correspondence that is significantly stricter than the variant proposed in [6].

Definition 13 (Operational Correspondence).

An encoding $\text{enc}(\cdot) : \mathcal{P}_{\mathcal{S}} \rightarrow \mathcal{P}_{\top}$ is *operationally corresponding* w.r.t. $\dot{\approx} \subseteq \mathcal{P}_{\top}^2$ if it is:

Complete: $\forall S, S'. S \Longrightarrow_{\mathcal{S}} S'$ implies $\exists T. \llbracket S \rrbracket \Longrightarrow_{\top} T \wedge \llbracket S' \rrbracket \dot{\approx} T$

Sound: $\forall S, T. \llbracket S \rrbracket \Longrightarrow_{\top} T$ implies $\exists S'. S \Longrightarrow_{\mathcal{S}} S' \wedge \llbracket S' \rrbracket \dot{\approx} T$

The ‘if’-part of Lemma 12 implies operational completeness w.r.t. $\dot{\approx}$ and the ‘only-if’-part contains the main argument for operational soundness w.r.t. $\dot{\approx}$. Hence $\llbracket \cdot \rrbracket$ is operationally corresponding w.r.t. to $\dot{\approx}$.

Theorem 1. The encoding $\llbracket \cdot \rrbracket$ is operationally corresponding w.r.t. to $\dot{\approx}$.

To obtain divergence reflection we show that there is no infinite sequence of only auxiliary steps. Then divergence reflection follows from the combination of this fact and Lemma 12.

Theorem 2. The encoding $\llbracket \cdot \rrbracket$ reflects divergence.

The encoding function ensures that $\llbracket S \rrbracket$ has an unguarded occurrence of \checkmark iff S has such an unguarded occurrence. Operational correspondence ensures that S and $\llbracket S \rrbracket$ also answer the question for the reachability of \checkmark in the same way.

Theorem 3. The encoding $\llbracket \cdot \rrbracket$ is success sensitive.

In a similar way we can prove that a source term reaches a barb iff its translation reaches the respective translated barb.

Theorem 4. For all S, c , it holds $S \downarrow_c$ iff $\llbracket S \rrbracket \downarrow_{\llbracket c \rrbracket}$.

As proved in [18], Theorem 1, the fact that $\dot{\approx}$ is success sensitive and respects (translated) barbs, Theorem 3, and Theorem 4 imply that for all S it holds S and $\llbracket S \rrbracket$ are (success sensitive, (translated) barb respecting, weak, reduction) bisimilar, i.e., $S \dot{\approx} \llbracket S \rrbracket$. Bisimilarity is a strong relation between source terms and their translation. On the other hand, because of efficiency, distributability preserving encodings are more interesting. Because of once the encoding $\llbracket \cdot \rrbracket$ obviously does not preserve distributability. As discussed in [16] bisimulation often forbids distributed encodings. Instead they propose coupled simulation as a relation that still provides a strong connection between source terms and their translations but is more flexible. Following the approach in [16] we consider a decentralised coordinator next.

$$\langle P \rangle \triangleq (\nu a)(\llbracket P \rrbracket \mid *_a(c, r, l, r').(\bar{r} \mid \llbracket b \rrbracket . \text{IF } b \text{ THEN } \bar{r}' \langle T \rangle))$$

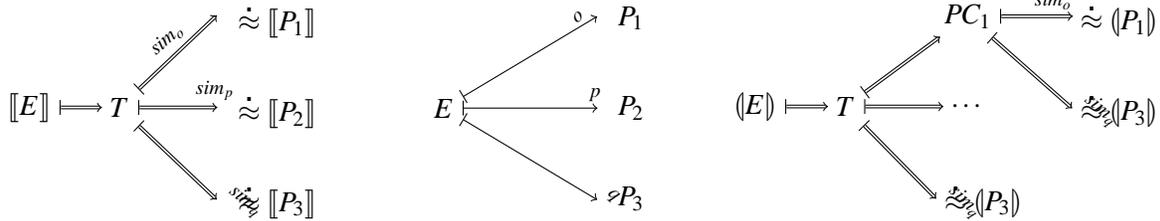
Figure 3: A **decentralised** encoding from CSP into CCS with value passing.

5 The Decentralised Encoding

Figure 3 presents a decentralised variant of the coordinator in Figure 2. The only difference between the centralised and the decentralised version of the coordinator is that the latter can request to check different locks concurrently. Technically $\llbracket \cdot \rrbracket$ and $\langle \cdot \rangle$ differ only by the use of once. As a consequence the steps of different simulation attempts can overlap and even (pre-processing) steps of simulations of conflicting source term steps can interleave to a certain degree. Because of this effect, $\langle \cdot \rangle$ does not satisfy the version of operational correspondence used above for $\llbracket \cdot \rrbracket$, but $\langle \cdot \rangle$ satisfies weak operational correspondence that was proposed in [6] as part of a set of quality criteria.

Since several announcements can be processed concurrently by the decentralised coordinator, here all consumptions of announcements are auxiliary steps. Instead the consumption of positive instantiations of locks can mark a point of no return. In contrast to $\llbracket \cdot \rrbracket$ not every point of no return in $\langle \cdot \rangle$ unambiguously marks a simulation of a single source term step, because in contrast to $\llbracket \cdot \rrbracket$ the encoding $\langle \cdot \rangle$ introduces *partial commitments* [17, 19].

Consider the example $E = (o \rightarrow P_1 \square p \rightarrow P_2) \parallel_{\{o, p\}} (o \rightarrow P_3 \square p \rightarrow P_4 \square q \rightarrow P_5)$.



In the example, two sides of a parallel operator have to synchronise on either action p , or action o , or action q happens without synchronisation. In the centralised encoding $\llbracket E \rrbracket$ the use of once ensures that different simulation attempts cannot overlap. Thus, only after finishing the simulation of a source term step, the simulation of another source term step can be invoked. As a consequence each state reachable from encoded source terms can unambiguously be mapped to a single state of the source term. This allows us to use a stronger version of operational correspondence and, thus, to prove that source terms and their translations are bisimilar. The corresponding 1-to-1 correspondence between source terms and their translations is visualised by the first two graphs above, where $T \approx \llbracket E \rrbracket$.

The decentralised encoding $\langle E \rangle$ introduces partial commitments. Assume the translation of a source term that offers several alternative ways to be reduced. Then some encodings—as our decentralised one—do not always decide on which of the source term steps should be simulated next. More precisely a partial commitment refers to a state reachable from the translation of a source term in that already some possible simulations of source term steps are ruled out, but there is still more than a single possibility left.

In the decentralised encoding announcements can be processed concurrently and parts of different simulation attempts can interleave. The only blocking part of the decentralised encoding are conflicting attempts to consume the same positive instantiation of a lock. In the presented example above there are two locks; one for each side of the parallel operator. The simulations of the step on o and p need both of

these locks, whereas to simulate the step on q only a positive instantiation of the right lock needs to be consumed. By consuming the positive instantiation of the left lock in an attempt to simulate the step on o , the simulation of the step on p is ruled out, but the simulation of the step on q is still possible. Since either the simulation of the step on o or the simulation of the step on q succeeds, the simulation of the step on p is not only blocked but ruled out. But the consumption of the instantiation of the left lock does not unambiguously decide between the remaining two simulations. The intermediate state that results from consuming the instantiation of the left lock and represents a partial commitment is visualised in the right graph above by the state PC_1 .

Partial commitments forbid a 1-to-1 mapping between the states of a source term and its translations by a bisimulation. But, as shown in [16], partial commitments do not forbid to relate source terms and their translations by coupled similarity.

Whether the consumption of a positive instantiation of a lock is an auxiliary step—does not change the state of the term modulo $\dot{\approx}$ —, is a partial commitment, or unambiguously marks a simulation of a single source term step depends on the surrounding term, i.e., cannot be determined without the context. For simplicity we consider all steps that reduce a positive instantiation of a lock as simulation steps. Also steps on variants of the channels m , d , and $\varphi'(X)$ are simulation steps, because they unambiguously mark a simulation of a single source term step. All remaining steps of the decentralised encoding are auxiliary.

Definition 14 (Auxiliary and Simulation Steps). A step $T \mapsto_T T'$ such that $\exists S \in \mathcal{P}_S. (\llbracket S \rrbracket \Longrightarrow_T T)$ is called a *simulation step*, denoted by $T \overset{\cdot}{\mapsto} T'$, if $T \mapsto T'$ reduces a positive instantiation of a lock or is a step on a variant of m , d , or $\varphi'(X)$.

Else the step $T \mapsto_T T'$ is called an *auxiliary step*, denoted by $T \overset{\circ}{\mapsto} T'$.

Again let $\overset{\cdot}{\mapsto}$ denote the reflexive and transitive closure of $\overset{\cdot}{\mapsto}$ and let $\overset{\cdot}{\Longrightarrow} \triangleq \overset{\cdot}{\mapsto} \overset{\cdot}{\mapsto} \overset{\cdot}{\mapsto}$. Since auxiliary steps do not introduce partial commitments, they do not change the state modulo $\dot{\approx}$. The proof of this lemma is very similar to the centralised case.

Lemma 15. $T \overset{\cdot}{\mapsto} T'$ implies $T \dot{\approx} T'$ for all target terms T, T' .

In contrast to the centralised encoding, the simulation of a source term step in the decentralised encoding can require more than a single simulation step and a single simulation step not unambiguously refers to the simulation of a particular source term step. The partial commitments described above forbid operational correspondence, but the weaker variant proposed in [6] is satisfied. We call this variant weak operational correspondence.

Definition 16 (Weak Operational Correspondence).

An encoding $\text{enc}(\cdot) : \mathcal{P}_S \rightarrow \mathcal{P}_T$ is *weakly operationally corresponding* w.r.t. $\dot{\approx}_{cs} \subseteq \mathcal{P}_T^2$ if it is:

Complete: $\forall S, S'. S \Longrightarrow_S S'$ implies $\exists T. (\llbracket S \rrbracket \Longrightarrow_T T \wedge (\llbracket S' \rrbracket \dot{\approx}_{cs} T)$

Weakly Sound: $\forall S, T. (\llbracket S \rrbracket \Longrightarrow_T T)$ implies $\exists S', T'. S \Longrightarrow_S S' \wedge T \Longrightarrow_T T' \wedge (\llbracket S' \rrbracket \dot{\approx}_{cs} T')$

The only difference to operational correspondence is the weaker variant of soundness that allows for T to be an intermediate state that does not need to be related to a source term directly. Instead there has to be a way from T to some T' such that T' is related to a source term.

Theorem 5. The encoding $(\llbracket \cdot \rrbracket)$ is weakly operational corresponding w.r.t. to $\dot{\approx}$.

As in the encoding $\llbracket \cdot \rrbracket$, there is no infinite sequence of only auxiliary steps in $(\llbracket S \rrbracket)$. Moreover each simulation of a source term requires only finitely many simulation steps (to consume the respective positive instantiations of locks). Thus $(\llbracket \cdot \rrbracket)$ reflects divergence.

Theorem 6. *The encoding (\cdot) reflects divergence.*

The encoding function ensures that (S) has an unguarded occurrence of \checkmark iff S has such an unguarded occurrence. Operational correspondence again ensures that S and (S) also answer the question for the reachability of \checkmark in the same way.

Theorem 7. *The encoding (\cdot) is success sensitive.*

Similarly, a source term reaches a barb iff its translation reaches the respective translated barb.

Theorem 8. *For all S, c , it holds $S \downarrow_c$ iff $(S) \downarrow_{\llbracket c \rrbracket}$.*

Weak operational correspondence does not suffice to establish a bisimulation between source terms and their translations. But, as proved in [18], Theorem 5, the fact that $\dot{\approx}$ is success sensitive and respects (translated) observables, Theorem 7, and Theorem 8 imply that $\forall S. S$ and $\llbracket S \rrbracket$ are (success sensitive, (translated) barbs respecting, weak, reduction) coupled similar, i.e., $S \dot{\approx}_{cs} \llbracket S \rrbracket$.

It remains to show, that (\cdot) indeed preserves distributability. Therefore we prove that all blocking parts of the encoding (\cdot) refer to simulations of conflicting source term steps.

Theorem 9. *The encoding (\cdot) preserves distributability.*

6 Conclusions

We introduced two encodings from CSP into asynchronous CCS with name passing and matching. As in [16] we had to encode the multiway synchronisation mechanism of CSP into binary communications and, similarly to [16], we did so first using a centralised controller that was then modified into a decentralised controller. By doing so we were able to transfer the observations of [16] to the present case:

1. The centralised solution allows to prove a stronger connection between source terms and their translations, namely by bisimilarity. Our decentralised solution does not relate source terms and their translations that strongly and we doubt that any decentralised solution can do so.
2. Nonetheless, decentralised solutions are possible as presented by the second encoding and they still relate source terms and their translations in an interesting way, namely by coupled similarity.

Thus as in [16] we observed a trade-off between *centralised* but *bisimilar* solutions on the one-hand side and *decentralised* but only *coupled similar* solutions on the other side.

More technically we showed here instead a trade-off between centralised but *operationally corresponding* solutions on the one-hand side and *weakly operationally corresponding* but decentralised solutions on the other side. The mutual connection between operational correspondence and bisimilarity as well as between weak operational correspondence and coupled similarity is proved in [18].

Both encodings make strict use of the renaming policy and translate into closed terms. Hence the criterion *name invariance* is trivially satisfied in both cases. Moreover we showed that both encodings are *success-sensitive*, *reflect divergence*, and even *respect barbs* w.r.t. to the standard source term (CSP) barbs and a notion of translated barbs on the target. The centralised encoding $\llbracket \cdot \rrbracket$ additionally satisfies a variant of *operational correspondence* that is stricter than the variant proposed in [6]. The decentralised encoding (\cdot) satisfies *weak operational correspondence* as proposed in [6] and *distributability preservation* as proposed in [20]. Thus both encodings satisfy all of the criteria proposed in [6] except for compositionality. However in both cases the inner part is obviously compositional and the outer part only adds a fixed context.

References

- [1] J. C. M. Baeten (2005): *A Brief History of Process Algebra*. *Theor. Comput. Sci.* 335(2-3), pp. 131–146, doi:10.1016/j.tcs.2004.07.036.
- [2] E. Brinksma (1985): *A tutorial on LOTOS*. In: *Proc. of PSTV*, pp. 171–194.
- [3] S. D. Brookes (1983): *On the Relationship of CCS and CSP*. In: *Proc. of ICALP, LNCS 154*, pp. 83–96, doi:10.1007/BFb0036899.
- [4] H. Evrard & F. Lang (2015): *Automatic Distributed Code Generation from Formal Models of Asynchronous Concurrent Processes*. In: *Proc. of PDP, IEEE*, pp. 459–466, doi:10.1109/PDP.2015.96.
- [5] R. van Glabbeek (2012): *Musings on Encodings and Expressiveness*. In: *Proc. of EXPRESS/SOS, EPTCS 89*, pp. 81–98, doi:10.4204/EPTCS.89.7.
- [6] D. Gorla (2010): *Towards a Unified Approach to Encodability and Separation Results for Process Calculi*. *Information and Computation* 208(9), pp. 1031–1053, doi:10.1016/j.ic.2010.05.002.
- [7] M. Hatzel, C. Wagner, K. Peters & U. Nestmann (2015): *Encoding CSP into CCS (Extended Version)*. Technical Report. Available at <http://arxiv.org/abs/1508.01127>.
- [8] C. A. R. Hoare (1978): *Communicating Sequential Processes*. *Communications of the ACM* 21(8), pp. 666–677, doi:10.1145/359576.359585.
- [9] C. A. R. Hoare (2006): *Why ever CSP?* *Electronic Notes in Theoretical Computer Science* 162(0), pp. 209–215, doi:10.1016/j.entcs.2006.01.031.
- [10] I. Lanese & U. Montanari (2006): *Hoare vs Milner: Comparing Synchronizations in a Graphical Framework With Mobility*. *Electronic Notes in Theoretical Computer Science* 154(2), pp. 55 – 72, doi:10.1016/j.entcs.2005.03.032.
- [11] R. Milner (1980): *A calculus of communicating systems*. Springer, doi:10.1007/3-540-10235-3.
- [12] R. Milner (1986): *Process Constructors and Interpretations (Invited Paper)*. In: *IFIP Congress*, pp. 507–514.
- [13] R. Milner & D. Sangiorgi (1992): *Barbed Bisimulation*. In: *Proc. of ICALP, LNCS 623*, pp. 685–695, doi:10.1007/3-540-55719-9_114.
- [14] U. Nestmann (1996): *On Determinacy and Nondeterminacy in Concurrent Programming*. Ph.D. thesis, Universität Erlangen-Nürnberg.
- [15] U. Nestmann & B. C. Pierce (2000): *Decoding Choice Encodings*. *Information and Computation* 163(1), pp. 1–59, doi:10.1006/inco.2000.2868.
- [16] J. Parrow & P. Sjödin (1992): *Multiway synchronization verified with coupled simulation*. In: *Proc. of CONCUR, LNCS 630*, Springer, pp. 518–533, doi:10.1007/bfb0084813.
- [17] K. Peters (2012): *Translational Expressiveness*. Ph.D. thesis, TU Berlin.
- [18] K. Peters & R. van Glabbeek (2015): *Analysing and Comparing Encodability Criteria*. In: *Proc. of EXPRESS/SOS, EPTCS 190*, pp. 46–60, doi:10.4204/EPTCS.190.4.
- [19] K. Peters & U. Nestmann (2012): *Is it a “Good” Encoding of Mixed Choice?* In: *Proc. of FoSSaCS, LNCS 7213*, pp. 210–224, doi:10.1007/978-3-642-28729-9_14.
- [20] K. Peters, U. Nestmann & U. Goltz (2013): *On Distributability in Process Calculi*. In: *Proc. of ESOP, LNCS 7792*, Springer, pp. 310–329, doi:10.1007/978-3-642-37036-6_18.
- [21] P. Sjödin (1991): *From LOTOS Specifications to Distributed Implementations*. Ph.D. thesis, Department of Computer Science, Uppsala University. Available as Report DoCS 91/31.

Encoding the Factorisation Calculus (Representing the Intensional in the Extensional)

Reuben N. S. Rowe

Department of Computer Science
University College London

r.rowe@ucl.ac.uk

Jay and Given-Wilson have recently introduced the Factorisation (or SF-) calculus as a minimal fundamental model of *intensional* computation. It is a combinatory calculus containing a special combinator, F, which is able to examine the internal structure of its first argument. The calculus is significant in that as well as being combinatorially complete it also exhibits the property of structural completeness, i.e. it is able to represent any function on terms definable using pattern matching on arbitrary normal forms. In particular, it admits a term that can decide the structural equality of any two arbitrary normal forms.

Since SF-calculus is combinatorially complete, it is clearly at least as powerful as the more familiar and paradigmatic Turing-powerful computational models of λ -calculus and Combinatory Logic. Its relationship to these models in the converse direction is less obvious, however. Jay and Given-Wilson have suggested that SF-calculus is strictly more powerful than the aforementioned models, but a detailed study of the connections between these models is yet to be undertaken.

This paper begins to bridge that gap by presenting a faithful encoding of the Factorisation Calculus into the λ -calculus preserving both reduction and strong normalisation. The existence of such an encoding is a new result. It also suggests that there is, in some sense, an equivalence between the former model and the latter. We discuss to what extent our result constitutes an equivalence by considering it in the context of some previously defined frameworks for comparing computational power and expressiveness.

1 Introduction

Mathematical models of computation are useful in studying the formal properties of programming practice. Indeed, the field of computing today arose partly out of the study of such abstract models: namely Turing Machines [33], the λ -calculus [8], and Combinatory Logic [10], which are consequently considered to be archetypal computational models. It is standard practice to qualify the abilities, or expressiveness, of a formal model of computation by demonstrating that it may *simulate* (and be simulated by) the operation of other formal models. This is the very essence of the notion of *Turing-completeness*, which encapsulates the intuition that a model may carry out any operation that is ‘effectively computable’. To construct such a simulation one must first give an injective mapping, showing how the terms of the source model may be represented by terms of the target. For example, this is the basis behind the process of Gödelization and the Church encoding of natural numbers [22]. Two basic properties are then required: that each atomic operational step of the source model is reflected by one or more steps of the target, and that a program of the target model *terminates* whenever the corresponding source program does. The former is a key ingredient of Landin’s influential work on comparing languages [25], while the latter is used as a criterion for comparing expressiveness by, e.g., Felleisen [12]. Formal definitions of encodings incorporating these properties, referred to as “faithful”, are already in use by the 90s, e.g. in [2], and are now common (see e.g. [15]).

Recently, Jay’s work on formal models of generic pattern matching [17] have led, in collaboration with Given-Wilson, to the formulation of the *Factorisation Calculus*. This is a combinatory calculus comprising two combinators: the S combinator, familiar from Combinatory Logic; and a new F combinator. The purpose of the latter is to enable arbitrary (head) normal terms to be *factorised*, that is split into their constituent parts, thereby allowing the examination of the internal structure of terms. This endows Factorisation Calculus with an interesting and powerful property: that of *structure completeness*. This means that any function on terms themselves definable by pattern matching over *arbitrary* normal forms is *representable*. Thus, Factorisation Calculus can be viewed as a minimal, fundamental model characterising not only the abstract notion of *pattern matching* but also *intensional computation*.

Jay and Given-Wilson show that Factorisation Calculus is Turing-complete, demonstrating a straightforward simulation of Combinatory Logic in their calculus. Moreover, due to its structural completeness, there is a term of Factorisation Calculus which can decide the structural equality of any two arbitrary normal forms. Conversely, factorisation and structural equality of normal forms *cannot* be so represented in λ -calculus and Combinatory Logic, thus these models are *not* structure complete. This hints at some sort of disparity in the expressivity of the two models. In their original and subsequent research [18, 14, 19], Jay and Given-Wilson speculate that the added expressive power may manifest itself in a non-existence result for simulations of Factorisation Calculus in λ -calculus, but this is not pursued in detail.

We show that there *does* exist a simulation of Factorisation Calculus within λ -calculus. The existence of such a simulation has not been demonstrated before, and this is the primary contribution of our paper. The simulation is made possible by a construction due to Berarducci and Böhm, which shows how to encode a certain class of term rewriting systems within λ -calculus. We show that Factorisation Calculus can be simulated by such a term rewriting system, whence the result follows. In the classical framework, our result signifies that Factorisation Calculus is no more powerful than λ -calculus. Thus there appears to be a mismatch between our result and the structure completeness property that the standard simulation-based notion of equivalence does not account for. To begin to try and resolve this, we consider some research in the literature which refines the concept of computational equivalence and discuss how our result relates to this.

Outline The rest of this paper is organised as follows. Section 2 recalls Jay and Given-Wilson’s Factorisation Calculus and its basic properties. Section 3 describes Berarducci and Böhm’s construction for encoding so-called canonical rewrite systems within the λ -calculus. In Section 4 we present our technical contribution: a simulation of the Factorisation Calculus in the λ -calculus via this construction. Section 5 then discusses, in light of our results, how the relative expressiveness of Factorisation Calculus and λ -calculus may be characterised. Section 6 concludes and remarks on areas for future work.

2 Factorisation Calculus

We begin by presenting Jay and Given-Wilson’s Factorisation Calculus itself, and review its principal properties. Factorisation Calculus, or more accurately SF-calculus¹, is a combinatory calculus whose terms are those of the free algebra over the two-element signature containing the combinators S and F, which each reduce upon being applied to three arguments. The former is the familiar combinator from Combinatory Logic [10] which applies its first and second arguments to duplicates of its third. The F combinator, on the other hand, introduces new capabilities in the form of *factorisation*: it is

¹We may say that any combinatory calculus that is structure complete is a factorisation calculus.

able to examine the *internal* structure of its first argument and process its second and third in different ways depending on whether that argument is *atomic* (i.e. itself a combinator) or *compound* (i.e. a partial application). To illustrate this, consider how the F combinator reduces in the following two instances:

$$FSMN \rightarrow M \qquad F(SX)MN \rightarrow NSX$$

Observe that when the first argument is atomic, it eliminates its third argument and returns its second. On the other hand, when the first argument is compound it eliminates its second argument and *factorises* the first into its left- and right-hand constituent components, passing these *separately* to its third argument.

Formally, SF-calculus is defined as follows.

Definition 2.1 (SF-calculus [18, § 4]). *The SF-calculus is a combinatory rewrite system over terms (ranged over by uppercase roman letters M, N , etc.) given by the following grammar:*

$$M, N ::= S \mid F \mid MN$$

Terms of the form S, F, SM, FM, SMN , or FMN (i.e. partially applied combinators) are called factorable forms. Reduction of terms is the smallest contextually closed binary relation \rightarrow_{SF} on terms (with the reflexive transitive closure denoted by \rightarrow_{SF}^) satisfying:*

$$\begin{array}{ll} SMNX \rightarrow_{SF} MX(NX) & \\ FOMN \rightarrow_{SF} M & \text{if } O \text{ is } S \text{ or } F \\ F(PQ)MN \rightarrow_{SF} NPQ & \text{if } PQ \text{ is a factorable form} \end{array}$$

Reduction of SF-calculus is confluent, and the K combinator of Combinatory Logic can be represented in SF-calculus by FF (also, indeed, by FS). Thus, there is a trivial encoding of Combinatory Logic in SF-calculus which preserves reduction and strong normalisation [14].

The behaviour of the F combinator gives SF-calculus an *intensional* quality: one may define higher order functions in SF-calculus which discriminate between functions whose *implementations* are different even when those functions are *extensionally* equal (i.e. produce identical outputs for identical inputs). For example, for any normal form X , the term $I_X \equiv S(FF)X$ implements the identity function (i.e. $I_X M \rightarrow_{SF}^* M$ for all M) and thus all such terms are extensionally equal. However, an SF-term T can be constructed which distinguishes them (by behaving as $T I_X \rightarrow_{SF}^* X$). Moreover, one can construct an SF-term that can decide the equality of any two *arbitrary* normal forms.

The intensional behaviour of SF-calculus is formally characterised by a property called *structure completeness*, which captures the notion that every *symbolic computation* (i.e. Turing-computable symbolic function) on normal forms is represented by some term.

Definition 2.2 (Structure Completeness [18, § 7-8]). *Let \mathcal{C} be a confluent combinatory calculus whose terms include variables, with reduction relation $\rightarrow_{\mathcal{C}}^*$. Define patterns to be the linear normal forms (i.e. containing no more than one occurrence of each variable), and matchable forms to be partially applied combinators.*

1. *A match $\{U/P\}$ of a pattern P against a term U may be defined to succeed with a substitution of terms for variables, or fail as follows (where ld denotes the identity function and \uplus the disjoint union of substitutions with match failure as an absorbing element):*

$$\begin{array}{lll} \{U/x\} = [U/x] & \{A/A\} = \text{ld} & \text{(if } A \text{ atomic)} \\ \{UV/PQ\} = \{U/P\} \uplus \{V/Q\} & & \text{(if } UV \text{ a compound)} \\ \{U/P\} = \text{fail} & & \text{(otherwise, if } U \text{ matchable)} \end{array}$$

2. A case is an equation of the form $P = M$ where P is a pattern and M an arbitrary term, which defines a symbolic function \mathcal{G} on terms by $\mathcal{G}(U) = \sigma(M)$ if $\{U/P\}$ succeeds with substitution σ , and $\mathcal{G}(U) = U$ if it fails.
3. A confluent combinatory calculus is structure complete if for every pattern P and term M , there is some term G such that $GU \rightarrow_{\mathcal{G}}^* \mathcal{G}(U)$ for every term U on which \mathcal{G} is defined; i.e. the symbolic function defined by every case is represented by some term.

Structure completeness subsumes combinatorial completeness since $\lambda x.M$ is given by the case $x = M$.

Theorem 2.3 ([18, Cor. 8.4]). *SF-calculus is structure complete.*

The F combinator itself represents a symbolic computation \mathcal{F} , namely that of factorisation:

$$\begin{array}{ll} \mathcal{F}(A, M, N) = M & \text{if } A \text{ is atomic} \\ \mathcal{F}(PQ, M, N) = NPQ & \text{if } PQ \text{ is compound} \end{array}$$

A significant (and arguably remarkable) fact is that \mathcal{F} cannot be represented in Combinatory Logic (for definitions of atomic and compound appropriate thereto).

Theorem 2.4 ([18, Thm. 3.2]). *Factorisation of SK-combinators is a symbolic computation that is not representable in Combinatory Logic.*

The equality predicate on normal forms also has no representation in Combinatory Logic. Thus, there exist (symbolic) functions, which are clearly ‘computable’ from an empirical point of view, that are not (directly) representable in Combinatory Logic (there is also a similar result for λ -calculus [4]). This result clearly points towards some form of added expressivity possessed by SF-calculus over the archetypal computational models. It is to this issue that we will return in Section 5.

3 Strongly Normalising Solutions of Equational Systems in λ -calculus

We now reiterate the interpretation result of Berarducci and Böhm [5], upon which our technical contribution rests. Essentially, this result says that systems of equations for a particular class of term algebras can be given solutions in the λ -calculus such that the representation of each atomic term of the algebra is *strongly normalising* thus having a *normal form*. Moreover, when the set of equations is interpreted as a rewrite system the encoding of terms preserves reduction and strong normalisation.

We assume the usual definitions of the λ -calculus without further explanation (readers may refer to [4] for details), with Λ denoting the set of lambda terms, \rightarrow_{β}^* denoting the (multi-step) β -reduction relation, and $=_{\beta}$ denoting β -equality (i.e. the equivalence relation on lambda terms induced by β -reduction). Furthermore, we also assume the familiar algebraic notion of the set $\text{Ter}(\Sigma)$ of (Σ -)terms over the signature (set of *function symbols*, each with an associated arity) Σ . We can then also consider the set $\Lambda(\Sigma)$ of *extended* lambda terms (i.e. lambda terms which may contain Σ -terms); notice that both $\text{Ter}(\Sigma)$ and Λ are (strict) subsets of $\Lambda(\Sigma)$.

Definition 3.1 (Canonical Systems of Equations). *Fix a signature Σ and let \mathcal{E} be a set of equations between terms $t \in \text{Ter}(\Sigma)$. We say that \mathcal{E} is canonical if Σ can be partitioned into two disjoint subsets Σ_0 and Σ_1 (i.e. $\Sigma = \Sigma_0 \cup \Sigma_1$) such that: each equation in \mathcal{E} is of the form $f(c(x_1, \dots, x_m), y_1, \dots, y_n) = t$ with $c \in \Sigma_0$ and $f \in \Sigma_1$ and where the variables $x_1, \dots, x_m, y_1, \dots, y_n$ are all distinct and form a superset of the variables in the term t ; and for each distinct pair $(c, f) \in \Sigma_0 \times \Sigma_1$ there is at most one equation in \mathcal{E} of this form. We say that \mathcal{E} is complete if for each distinct pair $(c, f) \in \Sigma_0 \times \Sigma_1$ there is exactly one such equation in \mathcal{E} .*

Canonical systems of equations, then, partition the signature into a set Σ_0 of (algebraic datatype) *constructors*, and Σ_1 of *programs* defined by pattern matching over the constructors on the first argument. Notice that any incomplete canonical system of equations can trivially be made complete by adding equations for the missing cases which simply project one of the function's arguments².

As an example of a canonical system of equations, we may observe that the usual recursive definition of addition over the datatype of (Peano) natural numbers is such a system:

$$\text{add}(\text{zero}, x) = x \quad \text{add}(\text{succ}(x), y) = \text{succ}(\text{add}(x, y))$$

We have a signature containing one function symbol `add`, one nullary constructor `zero`, and one unary constructor `succ`; moreover in this simple case, the equation system is already complete. In fact, every partial recursive function (on natural numbers) can be defined by a canonical system of equations [5, 6].

Given an equational system \mathcal{E} over a signature Σ , we can also take it to define a term rewriting system on $\text{Ter}(\Sigma)$ by reading each equation as a *rewrite rule*, i.e. $f_i(c_j(x_1, \dots, x_m), y_1, \dots, y_n) \rightarrow t$. We will write $\rightarrow_{\mathcal{E}}$ for the (one-step) reduction relation of the rewrite system defined by \mathcal{E} in this way (i.e. the smallest binary relation on terms satisfying the rewrite rules and closed under substitution and contexts), and $\rightarrow_{\mathcal{E}}^*$ for its reflexive, transitive closure (i.e. multi-step reduction). Ultimately, the aim is to interpret equational systems (and their associated rewrite systems) within λ -calculus.

Definition 3.2 (Interpretations). *A representation of the signature Σ is a function $\phi : \Sigma \rightarrow \Lambda$ from the function symbols of Σ to (closed) lambda terms, and induces a map $(\cdot)^\phi : \Lambda(\Sigma) \rightarrow \Lambda$ in the obvious way, namely by $x^\phi = x$, $(\lambda x.M)^\phi = \lambda x.M^\phi$, $(MN)^\phi = M^\phi N^\phi$, and for $f \in \Sigma$, $f(t_1, \dots, t_n)^\phi = \phi(f)t_1^\phi \dots t_n^\phi$. We say that a representation ϕ satisfies (or solves) \mathcal{E} if for each equation $t_1 = t_2$ (and corresponding rewrite rule $t_1 \rightarrow t_2$) in \mathcal{E} we have $t_1^\phi =_{\beta} t_2^\phi$ (and correspondingly also $t_1^\phi \rightarrow_{\beta}^* t_2^\phi$). When a representation ϕ satisfies \mathcal{E} , we say that ϕ is an interpretation (or a solution) of \mathcal{E} within λ -calculus.*

The following construction gives a special kind of representation for canonical systems of equations.

Definition 3.3 (Canonical Representations). *Let \mathcal{E} be a canonical system of equations that partitions the signature Σ into constructors $\Sigma_0 = \{c_1, \dots, c_r\}$ and programs $\Sigma_1 = \{f_1, \dots, f_k\}$. Without loss of generality we may assume that \mathcal{E} is complete, and so for each $1 \leq i \leq k$ and $1 \leq j \leq r$ let $b_{(i,j)}$ denote the term t such that $f_i(c_j(x_1, \dots, x_m), y_1, \dots, y_n) = t \in \mathcal{E}$.*

We will make use of the following notational abbreviations:

- Let $\langle t_1, \dots, t_n \rangle$ denote the Church n -tuple, i.e. $\lambda x.xt_1 \dots t_n$.
- Let Π_k^n (where $1 \leq k \leq n$) be the n -ary k^{th} projection function, i.e. $\lambda x_1 \dots x_n.x_k$.
- For $k \geq i > 1$, let $t_i, \dots, t_k, \dots, t_{i-1}$ denote the cyclic permutation of t_1, t_2, \dots, t_k beginning with t_i ; (in an abuse of notation we may also take $t_i, \dots, t_k, \dots, t_{i-1} = t_1, t_2, \dots, t_k$ when $i = 1$).

We now define two disjoint representations ϑ and ζ for constructors and programs respectively.

(Representation of Constructors) *For each $1 \leq i \leq r$, we define the representation of the constructor c_i as follows (where n is the arity of c_i):*

$$\vartheta(c_i) = \lambda x_1 \dots x_n f. f \Pi_i^r x_1 \dots x_n f$$

(Representation of Programs) *We choose k distinct fresh variables v_1, \dots, v_k not occurring in \mathcal{E} and fix a 'pre-representation', ψ , of Σ_1 defined by $\psi(f_i) = \langle v_i, \dots, v_k, \dots, v_{i-1} \rangle$ for each $1 \leq i \leq k$.*

²Alternatively, one might want to introduce a new nullary constructor (denoting an 'error' value) and add equations for the missing cases that simply return this value.

Using this representation, and the representation of constructors defined above, we then define $k \times r$ lambda terms $t_{(i,j)}$ ($1 \leq i \leq k$, $1 \leq j \leq r$), using the equations in \mathcal{E} as follows:

$$t_{(i,j)} = \lambda x_1 \dots x_m v_i \dots v_k \dots v_{i-1} y_1 \dots y_n. (b_{(i,j)}^\psi)^\vartheta$$

where $f_i(c_j(x_1, \dots, x_m), y_1, \dots, y_n) = b_{(i,j)} \in \mathcal{E}$ is the equation defining the behaviour of f_i when given a datum constructed using c_j as its first argument. We now define k terms, each one a Church r -tuple collating the bodies of all the cases for one of the programs in Σ_1 , as follows:

$$t_i = \langle t_{(i,1)}, \dots, t_{(i,r)} \rangle$$

(where $1 \leq i \leq k$). Each program is then represented by a Church k -tuple containing the collated representations of each program definition, beginning with its own. That is, ζ is defined by:

$$\zeta(f_i) = \langle t_i, \dots, t_k, \dots, t_{i-1} \rangle \quad (1 \leq i \leq k)$$

The representation $\phi = \vartheta \cup \zeta$ is called a canonical representation of Σ with respect to \mathcal{E} .

To gain some insight into the construction defined above, one can observe that it is related to an encoding of data attributed to Scott³ (and thus commonly referred to in the literature as the Scott encoding), which has subsequently been developed by others (e.g. [30, 27, 16, 31]). In the more familiar ‘standard’ encoding of functions, a fixed-point combinator is used to solve any recursion in the definition. This has the effect of making recursion *explicit*, and thus the representations of recursive functions have infinite expansions consisting of a ‘list’ of distinct instances of the function body, one for each recursive call that may be made. Applying the function to a datum then corresponds to a *fold* of the datum over this list, which discards the remaining infinity of recursive calls once the base case is reached. Therefore, as described by Böhm et al. [6, §3], in this scheme functions are ‘diverging objects which, when applied to data, may “incidentally” converge’. In encodings of the Scott variety, the recursive nature of functions is kept *implicit* and, while still triggered by application to a datum, only reproduced ‘on demand’. Hence we obtain finite objects which now ‘may “incidentally” diverge’ when applied to data⁴.

To explicate the particular encoding specified by Definition 3.3, we point out that the representation of a constructor is a (lambda) function that takes in the appropriate number of arguments (the *sub-data* of the datum that is subsequently constructed) and then waits to be given a function, which will be the program to be executed. Now, looking at how the constructor representation uses this function argument, we see that programs should expect to be given a projection function, followed by a number of sub-data, and then they are also passed *a copy of themselves*. It is this final element which is the key to Scott-type encodings, and allows recursion to be kept implicit. Looking now at the representation of programs we see that they are Church k -tuples containing an element for each program defined by \mathcal{E} (each of which is a Church r -tuple, where each element is a representation of one of the cases of that program’s definition). Thus the representation of each program contains the definition of *every* program defined by \mathcal{E} ; in particular it will contain the definition of each program which it may itself invoke. To illustrate in more detail how the encoding works, we can consider the general reduction sequence of a term representing the application of some program prog_i to some arguments, the first of which is a datum constructed as $c_j(d_1, \dots, d_m)$:

$$(\text{prog}_i(c_j(d_1, \dots, d_m)) \text{arg}_1 \dots \text{arg}_m)^\phi = \langle t_i, \dots, t_{i-1} \rangle(c_j(d_1, \dots, d_m))^\phi \text{arg}_1^\phi \dots \text{arg}_m^\phi$$

³The citation can be found in Curry, Hindley and Seldin [11, p. 504].

⁴This reversed form of the slogan is also due to Böhm et al., and illustrates the dual nature of the Scott and Church encodings.

$$\rightarrow_{\beta}^* \quad (\lambda x.x t_i \dots t_{i-1}) (\lambda f.f \Pi_j^r d_1 \dots d_m f) \arg_1^\phi \dots \arg_m^\phi \quad (1)$$

$$\rightarrow_{\beta} \quad (\lambda f.f \Pi_j^r d_1 \dots d_m f) t_i \dots t_{i-1} \arg_1^\phi \dots \arg_m^\phi \quad (2)$$

$$\rightarrow_{\beta} \quad t_i \Pi_j^r d_1 \dots d_m t_i t_{i+1} \dots t_{i-1} \arg_1^\phi \dots \arg_m^\phi \quad (3)$$

$$= \quad \langle t_{i,1}, \dots, t_{i,r} \rangle \Pi_j^r d_1 \dots d_m t_i t_{i+1} \dots t_{i-1} \arg_1^\phi \dots \arg_m^\phi \quad (4)$$

$$= \quad (\lambda x.x t_{i,1} \dots t_{i,r}) \Pi_j^r d_1 \dots d_m t_i t_{i+1} \dots t_{i-1} \arg_1^\phi \dots \arg_m^\phi \quad (5)$$

$$\rightarrow_{\beta} \quad \Pi_j^r t_{i,1} \dots t_{i,r} d_1 \dots d_m t_i t_{i+1} \dots t_{i-1} \arg_1^\phi \dots \arg_m^\phi \quad (6)$$

$$\rightarrow_{\beta} \quad t_{i,j} d_1 \dots d_m t_i t_{i+1} \dots t_{i-1} \arg_1^\phi \dots \arg_m^\phi \quad (7)$$

$$= \quad (\lambda x_1 \dots x_m v_i \dots v_k \dots v_{i-1} y_1 \dots y_n. (b_{(i,j)}^\psi)^\vartheta) d_1 \dots d_m t_i t_{i+1} \dots t_{i-1} \arg_1^\phi \dots \arg_m^\phi$$

$$\rightarrow_{\beta}^* \quad (b_{(i,j)} [d_1/x_1, \dots, d_m/x_m, \arg_1/y_1, \dots, \arg_n/y_n])^\phi \quad (8)$$

When the program is applied to a datum (Eq. (1)), its representation arranges to apply the datum first to the representations of each program beginning with its own, and then to the remainder of the arguments (Eq. (2)). Then, the particular structure of the datum will reduce the expression to pick out the appropriate case of the program definition to be executed, and apply it to the sub-data and the representations of each program, having duplicated the program being executed (Eqs. (3) to (7)). This then reduces to the representation of the appropriate substitution instance of the function body (Eq. (8)).

The result of Berarducci and Böhm says that a canonical representation gives an interpretation that also preserves strong normalisation.

Theorem 3.4 (Interpretation Theorem [5, Thm 3.4]). *Let Σ be a signature and \mathcal{E} a canonical set of equations for Σ ; then any canonical representation ϕ for Σ with respect to \mathcal{E} is an interpretation of \mathcal{E} within λ -calculus. In addition $(\cdot)^\phi$ preserves strong normalisation of closed terms.*

4 Encoding the Factorisation Calculus

In this section, we present our novel technical contribution: an encoding of SF-calculus in λ -calculus. We believe that this is the first such encoding presented in the literature. Our encoding is a faithful simulation; it preserves both the reduction behaviour of terms (thus also β -equality) and their termination behaviour (i.e. strong normalisation). In this section we shall make use of standard notation and results for term rewriting systems, details of which may be found in [24].

The key step to the encoding is to define a rewrite system behaviourally equivalent to SF-calculus that is also canonical, in the sense of Definition 3.1. It is then simply a matter of applying the construction of Berarducci and Böhm to obtain the encoding. Thus it is our translation of SF-calculus into this intermediate rewrite system that is the primary novelty of our contribution.

We are aiming to derive a set of rewrite rules that is *canonical* and so we must translate the schematic definition of Jay and Given-Wilson, as presented in Section 2, into one consisting of algebraic rewrite rules. There are two salient features of Definition 3.1 that we must take into account: that the rewrite rules must make a distinction between *programs* and *constructors*; and that the left-hand side of each rewrite rule must contain exactly one program symbol and one constructor. To obtain rewrite rules of the required form, we recast SF-calculus as a *curryfied, applicative* term rewriting system. That is, we first introduce an explicit program symbol `app` to denote application and use the symbols `S` and `F` solely as *constructors*. Secondly we stratify the combinators $C \in \{S, F\}$ into sets $\{C_0, C_1, C_2\}$ of constructors, each of which represent successive *partial* applications of their underlying combinator C . Although this

‘currying’ process is well-known from the world of functional programming, the reader may refer to [20, 3] for a formal definition of this process in the context of general term rewriting.

We may take the rewrite rules for the S combinator directly from the standard curried applicative formulation of Combinatory Logic (see e.g. [3]):

$$\text{app}(S_0, x) \rightarrow S_1(x) \quad \text{app}(S_1(x), y) \rightarrow S_2(x, y) \quad \text{app}(S_2(x, y), z) \rightarrow \text{app}(\text{app}(x, z), \text{app}(y, z))$$

The rules for producing the partial applications of the F combinator are similarly straightforward:

$$\text{app}(F_0, x) \rightarrow F_1(x) \quad \text{app}(F_1(x), y) \rightarrow F_2(x, y)$$

The rewrite rule for the full application of the F combinator is more tricky because we must find a way of implementing its two possible reductions. As in the original formulation of SF-calculus, since the choice of which reduction to make is determined by the structure of the first argument we should like to be able to use the pattern-matching capabilities inherent in the term rewriting discipline, e.g. by giving rewrite rules such as:

$$\text{app}(F_2(S_0, y), z) \rightarrow y \quad \text{app}(F_2(F_1(x), y), z) \rightarrow \text{app}(\text{app}(z, F_0), x)$$

However these rules are *not* canonical since they contain two occurrences of a constructor: they are pattern-matching ‘too deeply’. We can circumvent this by introducing an auxiliary *program* symbol f-reduce and then having the rewrite rule for the F₂ case of app delegate to this new program:

$$\text{app}(F_2(x, y), z) \rightarrow \text{f-reduce}(x, y, z)$$

Since f-reduce is an independent program symbol, and only needs to pattern match on its first argument to determine which result to compute, we may give canonical rewrite rules for it, such as the following:

$$\text{f-reduce}(S_0, y, z) \rightarrow y \quad \text{f-reduce}(F_1(x), y, z) \rightarrow \text{app}(\text{app}(z, F_0), x)$$

We now have all the components to be able to present a canonical rewrite system that faithfully implements SF-calculus.

Definition 4.1 (Curried Applicative SF-Calculus). *Let $\Sigma_{SF} = \Sigma_0 \cup \Sigma_1$ be the signature comprising the set $\Sigma_0 = \{S_0, S_1, S_2, F_0, F_1, F_2\}$ of constructors and the set $\Sigma_1 = \{\text{app}, \text{f-reduce}\}$ of programs. Curried Applicative SF-calculus is the term rewriting system $SF_{\text{@}}^{\mathcal{C}}$ defined by the rewrite rules given in Figure 1 over the signature Σ_{SF} . We denote its one-step and many-step reduction relations by $\rightarrow_{SF_{\text{@}}^{\mathcal{C}}}$ and $\rightarrow_{SF_{\text{@}}^{\mathcal{C}}}^*$.*

Notice that the rewrite rules of $SF_{\text{@}}^{\mathcal{C}}$ are canonical, in the sense of Definition 3.1, and that they are also *complete*. We also remark that $SF_{\text{@}}^{\mathcal{C}}$ is an *orthogonal* term rewriting system [24, Def. 2.1.1].

It is interesting to observe that our implementation of SF-calculus as a canonical applicative term rewriting system has involved an application of the *Visitor* design pattern⁵ [13]. When it comes to reducing a complete application of the F combinator, the computation must proceed based on the particular identity of some object (i.e. the first argument), but *without having any knowledge of that identity*. The solution is to apply the visitor pattern, which involves invoking a new ‘visit’ operation (that we call f-reduce) on the object, which in response executes the appropriate behaviour based on its *self-knowledge* of its own identity. We do not think it is entirely coincidental that the visitor pattern has arisen in our work: its connection with structural matching has already been noted [28], and investigating this connection further is an avenue for future research.

There is a straightforward translation from SF-calculus to $SF_{\text{@}}^{\mathcal{C}}$.

⁵In fact, the encoding that we are presenting in this paper arose as a direct result of considering how the Factorisation Calculus could be implemented in (Featherweight) Java.

$$\begin{array}{ll}
\text{app}(S_0, x) \rightarrow S_1(x) & \text{app}(F_0, x) \rightarrow F_1(x) \\
\text{app}(S_1(x), y) \rightarrow S_2(x, y) & \text{app}(F_1(x), y) \rightarrow F_2(x, y) \\
\text{app}(S_2(x, y), z) \rightarrow \text{app}(\text{app}(x, z), \text{app}(y, z)) & \text{app}(F_2(x, y), z) \rightarrow \text{f-reduce}(x, y, z) \\
\\
\text{f-reduce}(S_0, y, z) \rightarrow y & \text{f-reduce}(S_1(x), y, z) \rightarrow \text{app}(\text{app}(z, S_0), x) \\
\text{f-reduce}(F_0, y, z) \rightarrow y & \text{f-reduce}(F_1(x), y, z) \rightarrow \text{app}(\text{app}(z, F_0), x) \\
\\
\text{f-reduce}(S_2(p, q), y, z) \rightarrow \text{app}(\text{app}(z, \text{app}(S_0, p)), q) \\
\text{f-reduce}(F_2(p, q), y, z) \rightarrow \text{app}(\text{app}(z, \text{app}(F_0, p)), q)
\end{array}$$

Figure 1: A Complete Set of Canonical Rewrite Rules for Curried Applicative SF-calculus

Definition 4.2 (Translation of SF-calculus to $SF_{@}^{\mathcal{C}}$). *The translation $[[\cdot]]_{@}$ from SF-terms to $SF_{@}^{\mathcal{C}}$ -terms is defined by $[[S]]_{@} = S_0$, $[[F]]_{@} = F_0$, and $[[MN]]_{@} = \text{app}([M]_{@}, [N]_{@})$.*

We now show that $SF_{@}^{\mathcal{C}}$ faithfully implements SF-calculus.

Lemma 4.3 ($[[\cdot]]_{@}$ Preserves Reduction). *Let M and N be SF-terms; if $M \rightarrow_{SF}^* N$ then $[[M]]_{@} \rightarrow_{SF_{@}^{\mathcal{C}}}^* [[N]]_{@}$.*

Proof. It is sufficient to consider the basic reduction rules of SF-calculus. In the interests of clarity, we underline the redex that is contracted at each step. The case for S is straightforward:

$$\begin{aligned}
[[SMNX]]_{@} &= \text{app}(\text{app}(\text{app}(S_0, \underline{[[M]]_{@}}), [N]_{@}), [X]_{@}) \rightarrow_{SF_{@}^{\mathcal{C}}} \text{app}(\text{app}(S_1(\underline{[[M]]_{@}}), [N]_{@}), [X]_{@}) \\
&\rightarrow_{SF_{@}^{\mathcal{C}}} \text{app}(S_2(\underline{[[M]]_{@}}, [N]_{@}), [X]_{@}) \rightarrow_{SF_{@}^{\mathcal{C}}} \text{app}(\text{app}(\underline{[[M]]_{@}}, [X]_{@}), \text{app}([N]_{@}, [X]_{@})) \\
&= [[MX(NX)]]_{@}
\end{aligned}$$

The case for F with S the first argument (i.e. atomic) is as follows (the other atomic case is symmetric):

$$\begin{aligned}
[[FSMN]]_{@} &= \text{app}(\text{app}(\text{app}(F_0, S_0), \underline{[[M]]_{@}}), [N]_{@}) \rightarrow_{SF_{@}^{\mathcal{C}}} \text{app}(\text{app}(F_1(S_0), \underline{[[M]]_{@}}), [N]_{@}) \\
&\rightarrow_{SF_{@}^{\mathcal{C}}} \text{app}(F_2(S_0, \underline{[[M]]_{@}}), [N]_{@}) \rightarrow_{SF_{@}^{\mathcal{C}}} \text{f-reduce}(S_0, \underline{[[M]]_{@}}, [N]_{@}) \rightarrow_{SF_{@}^{\mathcal{C}}} [[M]]_{@}
\end{aligned}$$

When the first argument to F is a factorable form, we must further consider its structure. The case for when the first argument is SX (for some term X) is as follows:

$$\begin{aligned}
[[F(SX)MN]]_{@} &= \text{app}(\text{app}(\text{app}(F_0, \text{app}(S_0, \underline{[[X]]_{@}})), \underline{[[M]]_{@}}), [N]_{@}) \\
&\rightarrow_{SF_{@}^{\mathcal{C}}} \text{app}(\text{app}(\text{app}(F_0, S_1(\underline{[[X]]_{@}})), \underline{[[M]]_{@}}), [N]_{@}) \\
&\rightarrow_{SF_{@}^{\mathcal{C}}} \text{app}(\text{app}(F_1(S_1(\underline{[[X]]_{@}})), \underline{[[M]]_{@}}), [N]_{@}) \rightarrow_{SF_{@}^{\mathcal{C}}} \text{app}(F_2(S_1(\underline{[[X]]_{@}}), \underline{[[M]]_{@}}), [N]_{@}) \\
&\rightarrow_{SF_{@}^{\mathcal{C}}} \text{f-reduce}(S_1(\underline{[[X]]_{@}}), \underline{[[M]]_{@}}, [N]_{@}) \rightarrow_{SF_{@}^{\mathcal{C}}} \text{app}(\text{app}([N]_{@}, S_0), [X]_{@}) = [[NSX]]_{@}
\end{aligned}$$

Again, the case for when the first argument is FX (for some term X) is symmetric and can be obtained from the above sequence by replacing each occurrence of S_0 by F_0 and each occurrence of S_1 by F_1 .

The case for when the first argument is SXY (for some terms X and Y) is as follows:

$$[[F(SXY)MN]]_{@} = \text{app}(\text{app}(\text{app}(F_0, \text{app}(\text{app}(S_0, \underline{[[X]]_{@}}), \underline{[[Y]]_{@}})), \underline{[[M]]_{@}}), [N]_{@})$$

$$\begin{aligned}
& \rightarrow_{\text{SF}^{\mathcal{C}}_{\text{@}}} \text{app}(\text{app}(\text{app}(\text{F}_0, \text{app}(\text{S}_1(\llbracket X \rrbracket_{\text{@}}, \llbracket Y \rrbracket_{\text{@}})), \llbracket M \rrbracket_{\text{@}}), \llbracket N \rrbracket_{\text{@}})) \\
& \rightarrow_{\text{SF}^{\mathcal{C}}_{\text{@}}} \text{app}(\text{app}(\text{app}(\text{F}_0, \text{S}_2(\llbracket X \rrbracket_{\text{@}}, \llbracket Y \rrbracket_{\text{@}})), \llbracket M \rrbracket_{\text{@}}), \llbracket N \rrbracket_{\text{@}}) \\
& \rightarrow_{\text{SF}^{\mathcal{C}}_{\text{@}}} \text{app}(\text{app}(\text{F}_1(\text{S}_2(\llbracket X \rrbracket_{\text{@}}, \llbracket Y \rrbracket_{\text{@}})), \llbracket M \rrbracket_{\text{@}}), \llbracket N \rrbracket_{\text{@}}) \\
& \rightarrow_{\text{SF}^{\mathcal{C}}_{\text{@}}} \text{app}(\text{F}_2(\text{S}_2(\llbracket X \rrbracket_{\text{@}}, \llbracket Y \rrbracket_{\text{@}}), \llbracket M \rrbracket_{\text{@}}), \llbracket N \rrbracket_{\text{@}}) \\
& \rightarrow_{\text{SF}^{\mathcal{C}}_{\text{@}}} \underline{\text{f-reduce}(\text{S}_2(\llbracket X \rrbracket_{\text{@}}, \llbracket Y \rrbracket_{\text{@}}), \llbracket M \rrbracket_{\text{@}}, \llbracket N \rrbracket_{\text{@}})} \\
& \rightarrow_{\text{SF}^{\mathcal{C}}_{\text{@}}} \text{app}(\text{app}(\llbracket N \rrbracket_{\text{@}}, \text{app}(\text{S}_0, \llbracket X \rrbracket_{\text{@}})), \llbracket Y \rrbracket_{\text{@}}) = \llbracket N(\text{S}X)Y \rrbracket_{\text{@}}
\end{aligned}$$

Once more, the case for when the first argument is FXY (for some terms X and Y) is symmetric and can be obtained from the above sequence by replacing each occurrence of S_0 by F_0 , each occurrence of S_1 by F_1 , and each occurrence of S_2 by F_2 . \square

To show that $\llbracket \cdot \rrbracket_{\text{@}}$ preserves strong normalisation, we will rely on the notion of a *perpetual reduction sequence*. We recall the relevant definitions of perpetual reductions and their properties [21]. A term t is called an ∞ -term (also denoted $\infty(t)$) if it has an infinite reduction sequence. A reduction step $t \rightarrow s$ is called *perpetual* if $\infty(t)$ implies $\infty(s)$, that is it preserves divergence, and a reduction sequence $t_1 \rightarrow \dots \rightarrow t_n$ is perpetual if every step $t_i \rightarrow t_{i+1}$ is perpetual. Clearly, a perpetual reduction sequence $t \rightarrow^* t'$ also preserves divergence. A *redex* u is called perpetual if its contraction in every context yields a perpetual reduction step. Let $u \rightarrow t$ be a substitution instance of a rewrite rule r (so u is a redex and t its r -contraction), then call the subterms of u that are those substituted for the variables in r the *arguments* of u . Such an argument is said to be *erased* if it corresponds to a variable that does *not* occur in the right-hand side of r . It is the case that for orthogonal rewrite systems every redex whose erased arguments are strongly normalising and closed (i.e. containing no variables) is perpetual [21, Cor. 5.1].

We first prove a couple of auxiliary lemmas.

Lemma 4.4. *Let N be an SF-normal form, then $\llbracket N \rrbracket_{\text{@}}$ is strongly normalising.*

Proof. We characterise the normal forms as terms taking one of the following forms: S , F , SX , FY , SXY or FXY , in which each subterm is also a normal form. We then proceed by induction on the size of terms. The base cases, i.e. when N is either S or F are trivial since then $\llbracket N \rrbracket_{\text{@}}$ is itself a normal form. For the inductive cases, notice that the terms $S_1(\llbracket X \rrbracket_{\text{@}}, \llbracket Y \rrbracket_{\text{@}})$, $F_1(\llbracket X \rrbracket_{\text{@}}, \llbracket Y \rrbracket_{\text{@}})$, $S_2(\llbracket X \rrbracket_{\text{@}}, \llbracket Y \rrbracket_{\text{@}})$ and $F_2(\llbracket X \rrbracket_{\text{@}}, \llbracket Y \rrbracket_{\text{@}})$ are strongly normalising since they are head normal and by induction $\llbracket X \rrbracket_{\text{@}}$ and $\llbracket Y \rrbracket_{\text{@}}$ are strongly normalising as X and Y are by definition normal forms (smaller than N). It is then straightforward to show in each case that the (unique) reduction from $\llbracket N \rrbracket_{\text{@}}$ to its corresponding head normal form given above is perpetual since it does not erase any arguments, and $\text{SF}^{\mathcal{C}}_{\text{@}}$ is an orthogonal rewrite system. The result then follows. \square

Lemma 4.5. *Let X , M and N be SF-normal forms, and O an operator (i.e. either S or F) such that $OXMN \rightarrow_{\text{SF}} R$; then $\mathcal{C}[\llbracket OXMN \rrbracket_{\text{@}}] \rightarrow_{\text{SF}^{\mathcal{C}}_{\text{@}}}^* \mathcal{C}[\llbracket R \rrbracket_{\text{@}}]$ is a perpetual reduction sequence for any $\text{SF}^{\mathcal{C}}_{\text{@}}$ -term context \mathcal{C} .*

Proof. We show that there is a reduction sequence $\llbracket OXMN \rrbracket_{\text{@}} \rightarrow_{\text{SF}^{\mathcal{C}}_{\text{@}}}^* \llbracket R \rrbracket_{\text{@}}$ which contracts a perpetual redex at each step, from which the result immediately follows. In fact, the reduction sequences that witness this are exactly those that are used to show preservation of reduction. In each case notice that, in the reduction sequence demonstrated in the proof of Lemma 4.3, the only erased argument (when it exists) is in the final reduction step and in each case this argument is either $\llbracket M \rrbracket_{\text{@}}$ or $\llbracket N \rrbracket_{\text{@}}$ which by Lemma 4.4 is strongly normalising since M and N are normal forms (and also closed since we do not

consider SF-terms with variables). Thus, since $SF_{@}^{\mathcal{C}}$ is an orthogonal rewrite system, it follows that the redex contracted at each step is perpetual. \square

We can now prove the following result.

Lemma 4.6 ($[[\cdot]]_{@}$ Preserves Strong Normalisation). *Let M be a strongly normalising SF-term; then $[[M]]_{@}$ is strongly normalising.*

Proof. We use the same technique as used in the proof of [5, Thm. 3.4(2)], and proceed by (strong) induction on the length n of the longest reduction sequence from M to its normal form. When $n = 0$, then we have that M is a (SF-)normal form and thus $[[M]]_{@}$ is strongly normalising w.r.t $\rightarrow_{SF_{@}^{\mathcal{C}}}^*$ by Lemma 4.4. When $n > 0$ then M must contain at least one redex $OXYZ$. Consider an *innermost* redex, whose contractum is the term R . Thus $M = \mathcal{C}[OXYZ] \rightarrow_{SF} \mathcal{C}[R] = N$ (for some (SF-)term context \mathcal{C}) and X , Y and Z are normal forms (since the redex is innermost). Now, since M is strongly normalising so too is N , and the length of its longest reduction sequence must be strictly less than n (otherwise n would not be maximum). Thus by the inductive hypothesis $[[N]]_{@}$ is strongly normalising. Consider now the structure of $[[M]]_{@}$: we have $[[M]]_{@} = \mathcal{C}'[[[OXYZ]]_{@}]$ for some (SF $_{@}^{\mathcal{C}}$ -)term context \mathcal{C}' . Since X , Y and Z are normal forms, by Lemma 4.5 there is a perpetual reduction sequence from $[[M]]_{@} = \mathcal{C}'[[[OXYZ]]_{@}]$ to $[[N]]_{@} = \mathcal{C}'[[[R]]_{@}]$, i.e. one which preserves divergence. Therefore, since $[[N]]_{@}$ is strongly normalising so too is $[[M]]_{@}$ (if it were not, neither would $[[N]]_{@}$ be). \square

Using Berarducci and Böhm's construction, outlined in Section 3, we obtain an encoding of SF-calculus in the λ -calculus.

Definition 4.7 (Translation of SF-calculus to λ -calculus). *Fix a canonical representation ϕ_{SF} of Σ_{SF} w.r.t. the rewrite rules of $SF_{@}^{\mathcal{C}}$. The mapping $[[\cdot]]_{\lambda} = (\cdot)^{\phi_{SF}} \circ [[\cdot]]_{@}$ translates SF-calculus to λ -calculus.*

We leave it as an exercise to the reader to compute such a canonical representation ϕ_{SF} .

We now present our main result: that $[[\cdot]]_{\lambda}$ is a faithful encoding of SF-calculus in λ -calculus.

Theorem 4.8 (Faithful Encoding of SF-calculus in λ -calculus). *The translation $[[\cdot]]_{\lambda}$ of SF-calculus into λ -calculus preserves reduction and strong normalisation.*

Proof. The result follows directly from the fact that the two translations that are composed to obtain $[[\cdot]]_{\lambda}$, namely $[[\cdot]]_{@}$ and $(\cdot)^{\phi_{SF}}$, each satisfy both these properties. In the case of the former, we refer to Lemmas 4.3 and 4.6; for the latter to the results of Berarducci and Böhm, cf. Theorem 3.4 (note that all terms are closed, since we do not consider SF-calculus with variables). \square

5 Expressiveness of Factorisation: Discussion & Related Work

We now turn our attention to the question of the expressiveness of SF-calculus relative to λ -calculus and Combinatory Logic. On one hand our results, i.e. Theorem 4.8, along with those of Jay and Given-Wilson [18], show that SF-calculus and λ -calculus simulate the same executions. On the other, SF-calculus is structure complete whereas λ -calculus is not. Is this a contradiction and, if so, how may it be resolved? Notwithstanding the long tradition of using simulations to characterise computational equivalence, it has since been realised that a refinement of this notion is necessary to draw richer, more meaningful comparisons. While a complete and in-depth analysis is not within the scope of the current paper, by discussing our results with reference to some of this work we aim to draw some concrete conclusions about the how the expressiveness of structure completeness and SF-calculus may be characterised.

The ‘Standard’ Notion of Equivalence. From the earliest research into computability, two aspects of abstract notions of computation were identified as relevant to the idea of expressiveness. Firstly, one wants to compare the respective set of *functions* that each model computes. For example, in recursion theory it was shown that the set of primitive recursive functions (on natural numbers) is a strict subset of the recursive functions (see e.g. [32]). At the same time, there is a requirement to compare models that operate, at a fundamental level, in diverse domains. Turing Machines, λ -calculus and Combinatory Logic are, operationally, quite different ways to compute. It was shown however that via particular (and now canonical) representations of numbers, each model can ‘compute’ the same set of functions on natural numbers, namely the partial recursive functions. Conversely, Gödelization allows each of the computations in these models to be *simulated* by a partial recursive function [22].

This idea of simulation extends to the operation of the models themselves: each model may simulate the operational behaviour of the others. Such simulations also appear to abstract away the problem of *representation*: often we do want to compute functions of natural numbers, however more often we desire to compute functions over different domains; it is incumbent upon us to represent elements of the desired domain of discourse as terms of the computational model. The initial characterisation of the set of ‘computable’ functions was over the domain of natural numbers⁶, but what is the set of ‘computable’ functions over some other given domain? With simulations between models it seems that one may at least lay this question aside by observing that whatever can be represented in one model may then also be represented in the other, and therefore whatever functions they do compute it is the same in both cases. A stronger conclusion would be that the set of computable functions over arbitrary domains is isomorphic to the computable functions on natural numbers.

This simulation method has long since become the standard: to show two models are of equivalent computational power, demonstrate simulations of each in the other. One model of computation is only more powerful than another, then, if it is *not* possible to simulate the former in the latter. The approach has been cemented over the years, notably in Landin’s now seminal work [25]. As the search space of formal computation has been explored the notion of simulation has been adapted accordingly, and there are now a number of sophisticated simulations between all sorts of models, both sequential and concurrent (e.g. [1, 29]). In this tradition, Theorem 4.8 is a result showing that λ -calculus is computationally *as powerful as* SF-calculus.

Refining the Notion of Expressiveness. It was already noted over two decades ago that despite the broad applicability and application of the simulation method, it is not actually a fine-grained enough notion to provide a complete and universal characterisation of expressiveness. Felleisen observed that since the languages we wish to compare are (usually) Turing-complete, other methods (than simulation) must be found in order to verify claims of relative (in-)expressiveness [12]. He proposed a framework based on the concept, from logic, of *eliminability* of symbols from conservative extensions [23]. One logical system \mathcal{L} is a conservative extension of another system \mathcal{L}' if the expressions (formulae) and theorems of the latter are subsets of those of the former. A symbol of \mathcal{L} (which is not in \mathcal{L}') is eliminable if there is a homomorphism (i.e. a map preserving the syntactic structure) $\varphi : \text{Exp}(\mathcal{L}) \rightarrow \text{Exp}(\mathcal{L}')$ from the expressions of \mathcal{L} to the expressions of \mathcal{L}' , which acts as identity for expressions of \mathcal{L}' , such that an expression t is a theorem of \mathcal{L} if and only if $\varphi(t)$ is a theorem of \mathcal{L}' . Felleisen extends this to programming languages by analogy - formulae are (syntactically valid) programs and the theorems are the terminating programs. Then we may say that language \mathcal{L} is *more* expressive than language \mathcal{L}' when

⁶Turing’s work is different in this respect, since he deliberately embarked on a characterisation of computable functions over a different domain, namely that of strings of arbitrary symbols.

the former adds some *non-eliminable* syntactic construct, i.e. one which cannot be ‘translated away’. Thus the standard simulation approach is refined by imposing an extra criterion when one language is a superset of another: can the larger one be built from the smaller one using *macros*?

To place SF-calculus in this framework we can consider SKF-calculus, i.e. the extension of SF-calculus by including an additional atomic term: the familiar K combinator. Since SKF-calculus is a proper extension of Combinatory Logic, obtained by adding the F combinator, we can apply the expressiveness test of Felleisen. That is, we ask is F eliminable? The answer to this question is *no*; indeed this is guaranteed by Theorem 2.4. In this sense, Jay and Given-Wilson’s results *do* justify the claim that SF-calculus is more expressive than Combinatory Logic and λ -calculus.

More Abstract Notions of Computational Equivalence. The simulation method, and its refinement described above, are still firmly grounded in an *operational* view of computation, but recent work has sought to anchor formal comparisons of expressiveness in a more general, *abstract* basis. A notable contribution to this effort is the work of Boker and Dershowitz [7]. They abstract the notion of computational model as simply its *extension*, i.e. the set of functions over its inherent domain that it computes, and consider simulations (encodings) between them. They derive the remarkable result that combining the standard simulation approach with the natural containment of one extensionality within another leads to a paradox: some computational models can simulate models which are *strictly* more powerful in the sense that they have larger extensionalities (i.e. compute more functions). Thus, some representations *add more computational power*.

This result begs the question: do we consider simulations via such ‘active’ mappings to constitute an equivalence? One may suspect that the problem lies in allowing *injective* mappings between domains, and that imposing stricter conditions (e.g. bijections) would ensure ‘passiveness’. This is indeed the case, but adopting such restrictions is useless for most comparisons since the passive encodings are ones which are “almost identity”. We have no choice but to allow such encodings, although we do have the option of considering a hierarchy of equivalences based on the properties of the encodings used. Boker and Dershowitz define four increasingly stricter notions of (in-)equivalence based on, respectively: injective encodings (power equivalence), corresponding to the standard approach; injective encodings for which the images are computable in the simulating models (decent power equivalence); bijective encodings (bijective power equivalence); and bijections that are inverses of each other (isomorphism). Boker and Dershowitz also show that there are models (including Turing Machines and the recursive numeric functions) which cannot simulate any stronger models; they are *interpretation complete*.

We may also gain insight into the relative expressiveness of SF-calculus and λ -calculus using this framework. Our result shows that they are *power equivalent*. Theorem 2.4 shows that they cannot be bijectively power equivalent (nor, therefore, isomorphic) since that would imply that λ -calculus could distinguish arbitrary normal forms. Thus, in this sense they are not equivalent. We do not know if SF-calculus is *bijectionally stronger* than λ -calculus; this would require demonstrating a bijective encoding of the latter in the former. Also, we do not know if λ -calculus is decently power equivalent to SF-calculus; we consider it at least a possibility. We would also expect that, due to its intensional capabilities, SF-calculus is interpretation complete.

Related to the work of Boker and Dershowitz, is that of Cockett and Hofstra [9], and Longley [26] which are both concerned with category-theoretic descriptions of abstract computational models. In these frameworks model equivalence is interpreted by categorical isomorphism, and so akin to the strongest notion of equivalence considered by Boker and Dershowitz.

6 Conclusions & Future Work

In this paper, we have considered the relationship of the recently introduced SF-calculus to the ‘canonical’ computational model of λ -calculus and, so by extension, Combinatory Logic. We have demonstrated that SF-calculus can be faithfully encoded (i.e. simulated) in the λ -calculus by defining a behaviourally equivalent applicative term rewriting system and then interpreting this system in λ -calculus using a construction of Berarducci and Böhm. This result shows that SF-calculus and λ -calculus are of equivalent computational power, according to the classical interpretation of computational equivalence. We have also considered the relationship of SF-calculus to the λ -calculus using a more nuanced interpretation of equivalence, informed by research in the literature. Moreover, we hope to have exposed both SF-calculus and Berarducci and Böhm’s encoding to greater prominence. We feel that they are both subjects of great interest which deserve to be better known.

With respect to future work, there is still great scope for investigating the expressiveness of SF-calculus. There are the open questions we have highlighted regarding SF-calculus as it relates to the framework of Boker and Dershowitz. The categorical and denotational natures of SF-calculus also deserve exploration. Beyond this, the wider question of how best to qualify computational expressiveness still remains; we believe the study of SF-calculus can provide further insights on this. For example, the structural completeness property is already a new metric; Jay and Vergara have also considered a strengthening of the notion of decent power equivalence in which the simulation of each model in itself via the composition of the two encodings is computable in each model [19]. Quantitative questions regarding expressiveness also present themselves: e.g. our encoding of SF-calculus in λ -calculus leads to a large increase in the size of terms; are there lower bounds on the size of such increases which meaningfully quantify expressive power?

Acknowledgements We would like to thank Barry Jay, and also the anonymous reviewers for helpful comments in the preparation of the final version of this paper. This research was supported by EPSRC Grant EP/K040049/1.

References

- [1] M. Abadi & L. Cardelli (1996): *A Theory Of Objects*. Springer-Verlag New York, Inc., Secaucus, NJ, USA.
- [2] Martín Abadi, Luca Cardelli & Ramesh Viswanathan (1996): *An Interpretation of Objects and Object Types*. In: *Principles of Programming Languages (POPL’96)*, pp. 396–409, doi:10.1145/237721.237809.
- [3] S. van Bakel & M. Fernández (1997): *Normalization Results for Typeable Rewrite Systems*. *Information and Computation* 133(2), pp. 73–116, doi:10.1006/inco.1996.2617.
- [4] H. Barendregt (1981): *The Lambda Calculus, Its Syntax and Semantics*. North-Holland, Amsterdam.
- [5] Alessandro Berarducci & Corrado Böhm (1992): *A Self-Interpreter of Lambda Calculus Having a Normal Form*. In: *Computer Science Logic, 6th Workshop, CSL ’92*, pp. 85–99, doi:10.1007/3-540-56992-8_7.
- [6] Corrado Böhm, Adolfo Piperno & Stefano Guerrini (1994): *Lambda-Definition of Function(al)s by Normal Forms*. In: *ESOP’94, Edinburgh, U.K.*, pp. 135–149, doi:10.1007/3-540-57880-3_9.
- [7] Udi Boker & Nachum Dershowitz (2009): *The Influence of Domain Interpretations on Computational Models*. *Applied Mathematics and Computation* 215(4), pp. 1323–1339, doi:10.1016/j.amc.2009.04.063.
- [8] Alonzo Church (1936): *An Unsolvable Problem of Elementary Number Theory*. *American Journal of Mathematics* 58(2), pp. 345–363, doi:10.2307/2371045.

- [9] J.R.B. Cockett & Pieter J.W. Hofstra (2010): *Categorical Simulations*. *Journal of Pure and Applied Algebra* 214(10), pp. 1835 – 1853, doi:10.1016/j.jpaa.2009.12.028.
- [10] H. B. Curry & R. Feys (1958): *Combinatory Logic*. 1, North-Holland, Amsterdam.
- [11] H. B. Curry, J. R. Hindley & J. P. Seldin (1972): *Combinatory Logic*. 2, North-Holland, Amsterdam.
- [12] Matthias Felleisen (1991): *On the Expressive Power of Programming Languages*. *Sci. Comput. Program.* 17(1-3), pp. 35–75, doi:10.1016/0167-6423(91)90036-w.
- [13] Erich Gamma, Richard Helm, Ralph Johnson & John Vlissides (1995): *Design Patterns: Elements of Reusable Object-oriented Software*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA.
- [14] Thomas Given-Wilson (2014): *Expressiveness via Intensionality and Concurrency*. In: *Theoretical Aspects of Computing - ICTAC 2014, Bucharest, Romania*, pp. 206–223, doi:10.1007/978-3-319-10882-7_13.
- [15] Daniele Gorla (2010): *Towards a Unified Approach to Encodability and Separation Results for Process Calculi*. *Inf. Comput.* 208(9), pp. 1031–1053, doi:10.1016/j.ic.2010.05.002.
- [16] J.M. Jansen, P. W. M. Koopman & R. Plasmeijer (2006): *Efficient Interpretation by Transforming Data Types and Patterns to Functions*. In: *Trends in Functional Programming*, Nottingham, UK, pp. 73–90.
- [17] Barry Jay (2009): *Pattern Calculus - Computing with Functions and Structures*. Springer.
- [18] Barry Jay & Thomas Given-Wilson (2011): *A Combinatory Account of Internal Structure*. *J. Symb. Log.* 76(3), pp. 807–826, doi:10.2178/jsl/1309952521.
- [19] Barry Jay & Jose Vergara (2014): *Confusion in the Church-Turing Thesis*. Available at <http://arxiv.org/abs/1410.7103>. CoRR abs/1410.7103.
- [20] Stefan Kahrs (1995): *Confluence of Curried Term-Rewriting Systems*. *J. Symb. Comput.* 19(6), pp. 601–623, doi:10.1006/jscs.1995.1035.
- [21] Zurab Khasidashvili, Mizuhito Ogawa & Vincent van Oostrom (2001): *Perpetuality and Uniform Normalization in Orthogonal Rewrite Systems*. *Inf. Comput.* 164(1), pp. 118–151, doi:10.1006/inco.2000.2888.
- [22] S. C. Kleene (1936): *Lambda-definability and recursiveness*. *Duke Mathematical Journal* 2, pp. 340–353, doi:10.1215/S0012-7094-36-00227-2.
- [23] S.C. Kleene (1952): *Introduction to Metamathematics*. Bibliotheca Mathematica, Wolters-Noordhoff.
- [24] J. W. Klop (1992): *Term Rewriting Systems*. In S. Abramsky, Dov M. Gabbay & S. E. Maibaum, editors: *Handbook of Logic in Computer Science (Vol. 2)*, OUP, Inc., New York, NY, USA, pp. 1–116.
- [25] Peter J. Landin (1966): *The Next 700 Programming Languages*. *Commun. ACM* 9(3), pp. 157–166, doi:10.1145/365230.365257.
- [26] John Longley (2014): *Computability Structures, Simulations and Realizability*. *Mathematical Structures in Computer Science* 24(2), doi:10.1017/S0960129513000182.
- [27] Torben Æ. Mogensen (1992): *Efficient Self-Interpretation in Lambda Calculus*. *Journal of Functional Programming* 2, pp. 345–364, doi:10.1017/S0956796800000423.
- [28] Jens Palsberg & C. Barry Jay (1998): *The Essence of the Visitor Pattern*. In: *COMPSAC '98, August 19-21, 1998, Vienna, Austria*, pp. 9–15, doi:10.1109/COMPSAC.1998.716629.
- [29] D. Sangiorgi & D. Walker (2001): *PI-Calculus: A Theory of Mobile Processes*. CUP, New York, NY, USA.
- [30] J. Steensgaard-Madsen (1989): *Typed Representation of Objects by Functions*. *ACM Transactions on Programming Languages and Systems* 11(1), pp. 67–89, doi:10.1145/59287.77345.
- [31] Aaron Stump (2009): *Directly Reflective Meta-Programming*. *Higher-Order and Symbolic Computation* 22(2), pp. 115–144, doi:10.1007/s10990-007-9022-0.
- [32] G.J. Tourlakis (1984): *Computability*. Reston Publishing Company.
- [33] A. M. Turing (1937): *On Computable Numbers, with an Application to the Entscheidungsproblem*. *Proceedings of the London Mathematical Society* s2-42(1), pp. 230–265, doi:10.1112/plms/s2-42.1.230.