

**EPTCS 425**

Proceedings of the

**18th**

**Interaction and Concurrency Experience**

**Lille, France, 20th June 2025**

Edited by: Clément Aubert, Cinzia Di Giusto, Simon Fowler and Violet Ka I  
Pun

Published: 19th August 2025  
DOI: 10.4204/EPTCS.425  
ISSN: 2075-2180  
Open Publishing Association

## Table of Contents

Table of Contents .....	i
Preface .....	ii
<i>Clément Aubert, Cinzia Di Giusto, Simon Fowler and Violet Ka I Pun</i>	
<b>Invited Presentation:</b> On the Expressiveness of MPST .....	1
<i>Kirstin Peters</i>	
Modular Multiparty Sessions with Mixed Choice .....	2
<i>Franco Barbanera and Mariangiola Dezani-Ciancaglini</i>	
Reactive Semantics for User Interface Description Languages .....	21
<i>Basile Pesin, Celia Picard and Cyril Allignol</i>	
Bisimilarity and Simulatability of Processes Parameterized by Join Interactions .....	36
<i>Clemens Grabmayer and Maurizio Murgia</i>	
A Formalization of the Reversible Concurrent Calculus CCSKP in Beluga .....	55
<i>Gabriele Cecilia</i>	

# Preface

Clément Aubert

Augusta University, USA

Cinzia Di Giusto

Université Côte d'Azur, CNRS, I3S Sophia Antipolis, FR

Simon Fowler

University of Glasgow School of Computing Science, GB

Violet Ka I Pun

Western Norway University of Applied Sciences, NO

This volume contains the proceedings of ICE'25, the 18th Interaction and Concurrency Experience, which was held in Lille, France, as a satellite event of DisCoTec'25. The previous editions of ICE are now conveniently listed on the series website, which also links to ICE's dblp page. The conference ICE is co-located with since 2010, DisCoTec, also continued its efforts to collect and link to previous editions, sometimes using the services of the Wayback Machine.

The ICE workshop series features a distinguishing review and selection procedure: PC members are encouraged to interact with authors in a double anonymous environment. This year again, these interactions took place on the HotCRP platform which combines paper selection features with forum-like interactions. As in the past editions, the forum discussion during the review and selection phase of ICE'25 considerably improved the accuracy of the feedback from the reviewers and the quality of accepted papers<sup>1</sup> and offered the basis for lively discussion during the workshop. The time and effort spent on the interaction between reviewers and authors is rewarding for all parties. The discussions on HotCRP made it possible to resolve misunderstandings at an early stage of the review process, to discover and correct mistakes in key definitions, and to improve examples and presentation.

ICE was glad to receive 7 submissions, including one that was withdrawn and a short presentation that is not part of this proceedings. Each paper was reviewed by three PC members and about 75 comments were made in total, witnessing very lively discussions between PC members and authors, and during the discussion phase. We were proud to host one invited talk by Kirstin Peters. The abstract of her invited talk is included in this volume, together with the final versions of the research papers, which take into account the discussion at the workshop and after.

We would like to thank the 14 authors of all the submitted papers for their interest in the workshop. We thank Kirstin Peters for accepting our invitations to present her work. We are extremely grateful for the efforts of the PC members:

- Franco Barbanera (Dept. of Mathematics and Computer Science - University of Catania)
- Manel Barkallah (University of Namur)
- Matteo Cimini (University of Massachusetts Lowell)
- Farzaneh Derakhshan (Illinois Tech)
- Emanuele D'Ossualdo (University of Konstanz)
- **Luc Edixhoven (University of Southern Denmark) \***
- Lorenzo Gheri (University of Liverpool)
- Lucie Guillou (IRIF, Université Paris Cité)
- Ping Hou (University of Oxford)

---

<sup>1</sup>As well, we hope, as the quality of future resubmission in the case of papers being rejected.

- Andrew K. Hirsch (University at Buffalo, SUNY)
- Jonas Kastberg Hinrichsen (IT University of Copenhagen)
- Matthew Alan Le Brun (University of Glasgow)
- Andreia Mordido (LASIGE, University of Lisbon)
- Maurizio Murgia (Gran Sasso Science Institute)
- Jonah Pears (University of Kent)
- Felix Stutz (University of Luxembourg)
- Petra van den Bos (Formal Methods and Tools group (FMT) (University of Twente)
- Bas van den Heuvel (HKA Karlsruhe and University of Freiburg)

★ **The ICE 2025 Outstanding PC Member Award was awarded this year to Luc Edixhoven!**

Previously, were awarded:

- ICE 2024 Outstanding PC Member Award: *Bas van den Heuvel*
- ICE 2023 Outstanding PC Member Award: *Sergueï Lenglet*
- ICE 2022 Outstanding PC Member Award: *Duncan Paul Attard*
- ICE 2021 Outstanding PC Member Award: *Ivan Prokić*

We thank the DisCoTec'25 organizers, in particular the General Chair of the Organizing Committee, Simon Bliudze, for providing a welcoming, inclusive and lively environment for the preparation and staging of the event. We thank the editors of EPTCS for the publication of these post-proceedings. Cinzia and Clément are particularly thankful for the trust put in them by the ICE steering committee during their 4-years mandate, and wish Simon and Violet all the best in the organization of future editions—it has been a pleasure working with you all.



# On the Expressiveness of MPST

Kirstin Peters

University of Augsburg, Germany

kirstin.peters@uni-a.de

Multiparty session types (MPST) are a type discipline for enforcing the structured, deadlock-free communication of concurrent and message-passing programs. In this talk we will analyse the expressive power of MPST. In particular, we are interested in features that mark the difference expressive power of synchronous and asynchronous distributed languages. In the synchronous pi-calculus mixed choice is the main ingredient for its expressive power. Traditional MPST have in contrast usually a limited form of choice, in which alternative communication possibilities are offered by a single participant and selected by another. Accordingly, we extend MPST by a more general mixed choice construct.

**Biography** Kirstin Peters studied computer science at the university of Potsdam, did her PhD at the technical university of Berlin. After some stops at Uppsala University and the technical university of Dresden, she had a junior professorship at the technical university of Darmstadt. Since April 2022, she is a professor at the university of Augsburg.

# Modular Multiparty Sessions with Mixed Choice

Franco Barbanera \*

Dipartimento di Matematica e Informatica, Università di Catania, Catania, Italy

franco.barbanera@unict.it

Mariangiola Dezani-Ciancaglini

Dipartimento di Informatica, Università di Torino, Torino, Italy

dezani@di.unito.it

MultiParty Session Types (MPST) provide a useful framework for safe concurrent systems. *Mixed choice* (enabling a participant to play at the same time the roles of sender and receiver) increases the expressive power of MPST as well as the difficulty in controlling safety of communications. Such a control is more viable when modular systems are considered and the power of mixed choice fully exploited only inside loosely coupled modules. We carry over such idea in a type assignment approach to multiparty sessions. Typability for modular sessions entails Subject Reductions, Session Fidelity and Lock Freedom.

**Keywords:** Multiparty Sessions, Modular Systems, Global Types.

## 1 Introduction

MultiParty Session Types (MPST) offer a structured approach to the development and formal verification of concurrent and distributed systems [18, 19]. As in the vast majority of choreographic formalisms, two distinct but related views of concurrent systems are taken into account: (a) the *global view*, a formal specification via *global types* of the overall behaviour of a system; (b) the *local view*, namely a description, at different levels of abstraction, of the behaviours of the single components. A key issue in MPST, and choreographies in general, is the relation between these two views. Among others, we can refer to the notion of *projection*, used till recently in most of the MPST formalisms. Given a (well-formed) global type, the projection operator produces a tuple of local types – one for each component – which generalises binary session types [16, 17]. Such local types can be looked at as an abstraction of finer grained descriptions of processes. Another approach to the MPST global-local relationship is the one embodied in the so called *Simple MultiParty Sessions* (SMPS) formalisms. Such an approach was first introduced in [12] and [4], and further investigated in a bunch of papers, among which [3, 7, 5, 6, 1, 2]. Whereas the MPST approach typically considers two-layered local views – a layer of processes and a layer of local types – SMPS are based on single-layered local views, where only a fairly abstract notion of process is considered. In SMPS, which is the general setting of the present paper, systems of communicating processes are represented as *multiparty sessions*, i.e. parallel compositions of named processes (the *participants*). Then, by means of type systems, global types are inferred for such sessions. Typability is such to ensure relevant communication properties – typically Lock Freedom – for sessions. Besides the above mentioned ones, it is worth recalling that also other approaches have been investigated, like the one introduced in [32], where the global view is only implicitly considered.

---

\*Partially supported by Project “National Center for HPC, Big Data e Quantum Computing”, Programma M4C2, Investimento 1.3 – Next Generation EU; and by the PIAAno di inCEntivi per la Ricerca di Ateneo 2024-2026 UniCT (Linea di Intervento 1).

A common feature of all the above mentioned approaches, till lately, has been the use of communication models where, before any interaction, a process can clearly be identified as a sender or a receiver. The intrinsic potentiality of nondeterministically choosing among both inputs and outputs inside a single process interaction (usually referred to in the literature as “mixed choice”) has however recently intrigued session type researchers, both for the binary and the multiparty cases [10, 28, 29, 30, 31]. A thorough investigation of the expressivity of mixed choice in (synchronous) MPST formalisms has been carried on in [31]. For instance, mixed choice enables to implement protocols safely exploiting circular interactions, as shown in [31] through an example recalled in the present paper (see Example 2.4). This is not possible in usual MPST formalisms where typing [7] or, equivalently, the projectability condition on global types – as shown in [6] – consists essentially in checking the possibility of sequentialising the interactions in a protocol. Together with its expressive power, however, mixed choice brings subtly harmful features, as already exposed decades ago [15] in the setting of Communicating Finite State Machines [9] (an asynchronous formalism closely related to MPST). In the present paper we aim at exploiting such expressive power in a safe and controlled way using a SMPS setting. In order to do that we resort to the notion of *modularity*.

Modularity is a fairly general property of complex systems. Any complex system can be decomposed into smaller subsystems that are always going to be interdependent to some extent and independent to some other extent [33]. In fact, in many human activities, from business to biology, as well as to software engineering, modularisation offers a strategic approach enabling to cope with their complexity. Modularity in software engineering refers to the design approach that emphasises the separation of concerns: a complex software system is decomposed into smaller, loosely coupled modules, where coupling is the degree of interdependence between the modules. By means of project modularisation one manages to, among others, reduce complexity (breaking down a large system into smaller modules makes it more manageable and easier to understand [25]) as well as to improve testing and separation of concerns (SoC), a fundamental principle in software engineering. In particular, in modular programming, concerns are separated such that modules, performing logically coherent tasks, do interact through simple and manageable interfaces.

Our proposal is hence to restrict our attention to sessions corresponding to modularised systems. A type discipline is then proposed that profits, as in more rigid MPST formalisms, from a form of “sequentialisation” condition. Such a condition however, instead of being imposed on participants, is imposed on the modules forming a session, inside which the mixed choice can be freely used (at the cost of a thoroughly check, but limited inside the single modules, of all the possible interactions among participants). The inter-modules interactions are instead more controlled, so respecting the decoupling of modules characterising any sound decomposition of systems. It is then possible to prove the properties of Subject Reduction and Session Fidelity for typable modular sessions. Moreover, typability also entails the communication property of Lock Freedom. Typability is shown to be independent from the way a session can be modularised as well as from the order in which the typing of the modules is “sequentialised”. We propose, as working example, a modular extension of the above mentioned leader election example of [31].

*Overview.* In Section 2 we introduce the calculus of multiparty sessions with mixed choice as an extension of the SMPS calculus of [4] and [7]. Modularisable multiparty sessions are then formally presented in Section 3. Global types equipped with a coinductive LTS are defined in Section 4 in the style of [6]. Properties of typable modular sessions are proven in Section 5, namely Subject Reduction, Session Fidelity and Lock Freedom. In that section we also show that typability does not depend on how a session is modularised or on the particular modules considered during typing. A summing-up section, also discussing related and future works, concludes the paper.

## 2 Multiparty Sessions with Mixed Choice

We present now a SMPS synchronous calculus of multiparty sessions with mixed choice, inspired mainly by [31] and partially by [4]. We assume to have the following denumerable base sets: *messages* (ranged over by  $\lambda, \lambda', \dots$ ); *session participants* (ranged over by  $p, q, r, s, \dots$ ); *indexes* (ranged over by  $i, j, h, k, \dots$ ); *finite sets of indexes* (ranged over by  $I, J, H, K, \dots$ ). We refer to the denumerable set of participant names as  $\mathfrak{P}$ .

Processes, ranged over by  $P, Q, R, S, \dots$ , implement the behaviour of participants. In the following and in later definitions the symbol  $::=^{coind}$  does express that the productions have to be interpreted *coinductively* and that only *regular* terms are allowed. Then we can adopt in proofs the coinduction style advocated in [21] which, without any loss of formal rigour, promotes readability and conciseness.

**Definition 2.1 (Processes)** i) Action prefixes are defined by  $\pi ::= p?\lambda \mid p!\lambda$ .

ii) Processes are coinductively defined by

$$P ::=^{coind} \mathbf{0} \mid \sum_{i \in I} \pi_i.P_i$$

where  $I \neq \emptyset$  and finite, and  $\pi_l = q?\lambda_l$ ,  $\pi_j = q?\lambda_j$  (resp.  $\pi_l = q!\lambda_l$ ,  $\pi_j = q!\lambda_j$ ) imply  $\lambda_l \neq \lambda_j$ , for any  $l, j \in I$  such that  $l \neq j$ .

In the above definition,  $\sum_{i \in I} \pi_i.P_i$  stands, as usual, for the summand of processes  $\pi_i.P_i$ 's. A  $\sum_{i \in I} \pi_i.P_i$  process represents the nondeterministic choice of one of the actions  $\pi_i$ , after which the process continues as  $P_i$  with  $i \in I$ . As usual, we assume the summand of processes to be commutative and associative. A prefix  $\pi$  can be any input (i.e. of the form  $p?\lambda$ ) or output (i.e. of the form  $p!\lambda$ ) action. We use  $\mathbf{0}$  to denote the terminated process. For the sake of readability, we omit trailing  $\mathbf{0}$ 's in processes.

We define the participants of action prefixes by  $\text{prt}(p?\lambda) = \text{prt}(p!\lambda) = \{p\}$ . Moreover, we define the participants of processes by

$$\text{prt}(\mathbf{0}) = \emptyset \quad \text{prt}(\sum_{i \in I} \pi_i.P_i) = \bigcup_{i \in I} \text{prt}(\pi_i) \cup \bigcup_{i \in I} \text{prt}(P_i)$$

By the regularity condition and the finiteness of indexes,  $\text{prt}(P)$  is finite for any  $P$ .

Processes correspond to inductively defined terms of the calculus MCMP (Mixed Choice Multiparty Sessions) as defined in [31], where the  $\mu$ -operator is used to describe infinite behaviours. Our use of coinductively defined (possibly) infinite terms enables us to get simpler formalisations and proofs with respect to the use of  $\mu$ -terms. The latter entail just technicalities that can be dealt with as done in the literature on session types and MPST [19], where such terms are usually considered.

Multiparty sessions are parallel compositions of participant-process pairs, where all participants are different.

**Definition 2.2 (Multiparty sessions)** Multiparty sessions are defined by  $\mathbb{M} = p_1[P_1] \parallel \dots \parallel p_n[P_n]$  where  $p_j \neq p_l$  for  $1 \leq j, l \leq n$  and  $j \neq l$ .

We assume the standard structural congruence  $\equiv$  on multiparty sessions, stating that parallel composition is commutative and associative and has neutral elements  $p[\mathbf{0}]$  for any fresh  $p$ . Such a congruence can be inductively defined, as multiparty sessions are. We then write  $p[P] \in \mathbb{M}$  if  $\mathbb{M} \equiv p[P] \parallel \mathbb{M}'$  and  $P \neq \mathbf{0}$ . Moreover, we define the participants of multiparty sessions by  $\text{prt}(\mathbb{M}) = \{p \mid p[P] \in \mathbb{M}\}$ .

To define the *synchronous operational semantics* of sessions we use an LTS, whose transitions are decorated by *communication labels*, i.e. expressions of the shape  $p\lambda q$ . In the following we use  $\Lambda$  to range over communication labels.

**Notation:** We use  $\pi.P \dot{\vdash} Q$  as short for either  $\pi.P + Q$  or  $\pi.P$ . Such a notation enables us to present the

following LTS in a compact and yet formal way, since in our processes – as in [31] – we cannot have unprefixes  $\mathbf{0}$ 's as summands.

**Definition 2.3 (LTS for Multiparty Sessions)** *The labelled transition system (LTS) for multiparty sessions is the closure under structural congruence of the reduction specified by the unique axiom:*

$$[\text{COMM}] \frac{}{p[q!\lambda.P \dot{\vdash} P'] \parallel q[p?\lambda.Q \dot{\vdash} Q'] \parallel \mathbb{M} \xrightarrow{p\lambda q} p[P] \parallel q[Q] \parallel \mathbb{M}}$$

Rule [COMM] makes the communication possible: if participant  $p$  is enabled to send message  $\lambda$  to participant  $q$  which, in turn, is enabled to receive it, the message can be exchanged. This rule is non-deterministic in the choice of messages exchanged. The implementation issues raised by such operational semantics are similar to those for most of the calculi for concurrency and can be dealt with by resorting to suitable coordination protocols. Such issues are however outside the scope of the present paper.

Note that in the above semantics there is no difference between the behaviours of inputs and outputs, while usually a sender freely chooses among all its available messages. In actual communicating systems, messages would also carry values that are abstracted away in SMPS formalisms for the sake of simplicity. The present calculus (as well as its type system) could however be extended to messages with data.

We define *traces* as (possibly infinite) sequences of communication labels. Formally,

$$\sigma ::= \text{coind } \varepsilon \mid \Lambda \cdot \sigma$$

where  $\varepsilon$  is the empty sequence. When  $\sigma = \Lambda_1 \cdot \dots \cdot \Lambda_n$  ( $n \geq 0$ ) we write  $\mathbb{M} \xrightarrow{\sigma} \mathbb{M}'$  as short for

$$\mathbb{M} \xrightarrow{\Lambda_1} \mathbb{M}_1 \dots \xrightarrow{\Lambda_n} \mathbb{M}_n = \mathbb{M}'$$

We write  $\mathbb{M} \xrightarrow{\sigma}$  with the standard meaning. Moreover, we define the participants of labels by  $\text{prt}(p\lambda q) = \{p, q\}$  and the participants of traces – notation  $\text{prt}(\sigma)$  – as its obvious extension. We also denote by  $\mathcal{L}(\mathbb{M})$  the set of all labels the session  $\mathbb{M}$  can emit, i.e.  $\mathcal{L}(\mathbb{M}) = \{\Lambda \mid \mathbb{M} \xrightarrow{\Lambda}\}$ .

We present now, in our setting, the leader-election example used in [31] (inspired by [27], in turn inspired by [8]) to make evident the expressive power of mixed choice.

**Example 2.4 (Leader election [31])** Five participants ( $a, b, c, d, e$ ) interact with the aim of electing a leader. Each of them can send to the next participant, in a circular fashion, the message *leader* in order to ask it to become the leader. If such a communication succeeds, the sender terminates. Of course only two of this sort of communications can succeed. The protocol then allows two, among the remaining three participants, to be able to exchange the *leader* message. The receiver is hence considered as the *elect*ed leader and so it informs the station participant  $s$  which, in turn, provides to *delete* the participant which remained inactive during the previous interactions.

The above behaviour is implemented by the following session:


$$\begin{aligned} \mathbb{E} &= a[P_a] \parallel b[P_b] \parallel c[P_c] \parallel d[P_d] \parallel e[P_e] \parallel s[P_s] \\ \text{where } P_a &= e!leader \\ &\quad + b?leader.(c!leader + d?leader.s!elect) \\ &\quad + s?del \end{aligned}$$

and (with some abuse of notation)

$$P_s = \sum_{x \in \{a, b, c, d, e\}} (x?elect.\sum_{x \in \{a, b, c, d, e\}} x!del)$$

Processes of participants  $b, c, d$  and  $e$  are obtained out of  $P_a$  by applying the name substitution  $v = [a \mapsto b, b \mapsto c, c \mapsto d, d \mapsto e, e \mapsto a]$  as follows:

$$P_b = P_a v, P_c = P_b v, P_d = P_c v, P_e = P_d v$$



Session  $\mathbb{E}$  can be graphically represented by the diagram above on the left, where blue arrows represent the initial possible exchanges of the message *leader* and the red ones the potential further exchanges of such message. The following sequence of reductions is, for instance, the one leading to the election of  $e$ :

$$aleadere \cdot dleaderc \cdot cleadere \cdot eelects \cdot sdelb \quad \diamond$$

Lock Freedom is a relevant property of concurrent systems. We define it in our setting following [26]: roughly, there is always a continuation enabling a participant to communicate whenever it is willing to do so. Lock Freedom entails Deadlock Freedom, since it ensures progress for each participant.

**Definition 2.5 (Lock Freedom)** *A session  $\mathbb{M}$  is lock free if  $\mathbb{M} \xrightarrow{\sigma} \mathbb{M}'$  with  $\sigma$  finite and  $\mathbf{p} \in \text{prt}(\mathbb{M}')$  imply  $\mathbb{M}' \xrightarrow{\sigma' \cdot \Lambda}$  for some  $\sigma'$  and  $\Lambda$  such that  $\mathbf{p} \notin \text{prt}(\sigma')$  and  $\mathbf{p} \in \text{prt}(\Lambda)$ .*

The above definition corresponds to the notion of liveness used in [20] and [22] in a channel-based synchronous communication setting.

### 3 Modular Multiparty Sessions

“With great [expressive] power comes great responsibility” (Spider-Man’s Uncle Ben), since expressive power is often difficult to control and tame. Our aim is to provide a type system for multiparty sessions with mixed choice ensuring, like usual in SMPS, relevant communication properties, together with the guarantee that the overall behaviour of a session faithfully respects what the type assigned to the session, if any, describes. To do that, instead of restricting the expressive power of mixed choice, we decided to consider multiparty sessions that can be – as suggested by a well-known software engineering principle – modularised. A module, in our SMPS setting, is formalised in terms of a subsession inside which participants can freely interact by means of mixed choice. The communications among the modules are instead controlled by imposing them to be performed only by particular participants called “connectors”. The processes of the connectors (dubbed connecting processes) must satisfy the restriction that each choice involving a participant not belonging to the module must be between communications with only that participant. Definition 3.1, where  $\mathbf{P}$  is the set of module participants, formalises this condition.

**Definition 3.1 (Connecting processes)** *Given a set of participants  $\mathbf{P}$ , we say that a process  $P$  is  $\mathbf{P}$ -connecting if for any subprocess of  $P$ , say  $\Sigma_{i \in I} \pi_i.P_i$ , we have that  $\text{prt}(\pi_j) \notin \mathbf{P}$  for some  $j \in I$  implies  $\text{prt}(\pi_i) = \text{prt}(\pi_j)$  for all  $i \in I$ .*

A connector is hence a participant of a module (represented by the session  $\mathbb{M}$  in the definition below) whose process is a connecting process which can interact with the outside of the module.

**Definition 3.2 (Connectors)** *Let  $\mathbf{p}[P] \in \mathbb{M}$ . We say that the participant  $\mathbf{p}$  is a connector for  $\mathbb{M}$  if  $P$  is  $\text{prt}(\mathbb{M})$ -connecting and there is  $\mathbf{q} \in \text{prt}(P)$  such that  $\mathbf{q} \notin \text{prt}(\mathbb{M})$ .*

The notions of subsession and session partition are at the basis of that of modular session. We say that  $\mathbb{M}'$  is a *subsession* of  $\mathbb{M}$  and write  $\mathbb{M}' \subseteq \mathbb{M}$ , whenever  $\mathbf{p}[P] \in \mathbb{M}'$  implies  $\mathbf{p}[P] \in \mathbb{M}$ . A *partition* of a session  $\mathbb{M}$  is, as expected, a set of subsessions  $\{\mathbb{M}_h\}_{h \in H}$  of  $\mathbb{M}$  such that  $\text{prt}(\mathbb{M}) = \bigcup_{h \in H} \text{prt}(\mathbb{M}_h)$  and, for all  $h, k \in H$ ,  $\text{prt}(\mathbb{M}_h) \cap \text{prt}(\mathbb{M}_k) = \emptyset$ . Therefore,  $\mathbf{p}[P] \in \mathbb{M}$  implies that there is a unique  $k \in H$  such that  $\mathbf{p}[P] \in \mathbb{M}_k$ .

**Definition 3.3 ( $\mathcal{P}$ -partition)** *Let  $\{\mathbb{M}_h\}_{h \in H}$  be a partition of a session  $\mathbb{M}$ , and let  $\mathcal{P} = \{\mathbf{P}_k\}_{k \in K}$  be a partition of a finite subset of the set  $\mathfrak{P}$ . We say that  $\{\mathbb{M}_h\}_{h \in H}$  is a  $\mathcal{P}$ -partition of  $\mathbb{M}$  if  $H \subseteq K$  and  $\text{prt}(\mathbb{M}_h) \subseteq \mathbf{P}_h$  for all  $h \in H$ .*

It is not difficult to check that, given a session  $\mathbb{M}$  and a partition  $\mathcal{P} = \{\mathbf{P}_k\}_{k \in K}$  such that  $\text{prt}(\mathbb{M}) \subseteq \bigcup_{k \in K} \mathbf{P}_k$ , there is a unique  $\mathcal{P}$ -partition of  $\mathbb{M}$ .

A session is modularisable (with respect to a partition of participants) when it can be partitioned into subsessions that interact only by means of connectors.

**Definition 3.4 ( $\mathcal{P}$ -modularisation)** A  $\mathcal{P}$ -modularisation of  $\mathbb{M}$  is a  $\mathcal{P}$ -partition  $\{\mathbb{M}_h\}_{h \in H}$  of  $\mathbb{M}$  such that, for all  $h \in H$ , the following conditions hold

- i)  $\mathbf{p}[P] \in \mathbb{M}_h$  implies that either  $\mathbf{p}$  is a connector of  $\mathbb{M}_h$  or, for each  $\mathbf{q} \in \text{prt}(P)$ ,  $\mathbf{q} \in \text{prt}(\mathbb{M})$  implies  $\mathbf{q} \in \text{prt}(\mathbb{M}_h)$ ;
- ii) for each connector  $\mathbf{p}[P] \in \mathbb{M}_h$  and each  $\mathbf{q} \in \text{prt}(P) \setminus \text{prt}(\mathbb{M}_h)$ :  
if  $\mathbf{q} \in \text{prt}(\mathbb{M}_k)$ , then  $\mathbf{q}$  is a connector for  $\mathbb{M}_k$

In such a case, we say that  $\mathbb{M}$  is  $\mathcal{P}$ -modularisable.

It is worth noticing that we impose no limit on the number of connectors present in a module, as well as on the number of external connectors a connector can interact with (see the following example). Besides, a session  $\mathbb{M}$  is always  $\{\text{prt}(\mathbb{M})\}$ -modularisable. For example  $\mathbb{M} = \mathbf{p}[\mathbf{q}!\lambda + \mathbf{r}!\lambda']$  is  $\{\{\mathbf{p}\}\}$ -modularisable and its unique module does not contain any connector. Also, given a partition  $\mathcal{P}$  and a session  $\mathbb{M}$ , there exists a unique  $\mathcal{P}$ -modularisation of  $\mathbb{M}$ , if any.

**Example 3.5 (Modules with multiple connectors)** Let us consider  $\mathbb{M} = \mathbb{M}_1 \parallel \mathbb{M}_2$ , where  $\mathbb{M}_1 = \mathbf{u}[\mathbf{p}?\lambda.\mathbf{q}!\lambda + \mathbf{q}!\lambda.\mathbf{p}?\lambda] \parallel \mathbf{p}[\mathbf{u}!\lambda.(r!\lambda_1 + r?\lambda_2)] \parallel \mathbf{q}[\mathbf{u}?\lambda.(s!\lambda_1 + s?\lambda_2)]$  and  $\mathbb{M}_2 = \mathbf{v}[\mathbf{r}!\lambda.s?\lambda + \mathbf{s}?\lambda.r!\lambda] \parallel \mathbf{r}[\mathbf{v}?\lambda.(p?\lambda_1 + p!\lambda_2)] \parallel \mathbf{s}[\mathbf{v}!\lambda.(q?\lambda_1 + q!\lambda_2)]$ . This session is  $\{\{\mathbf{u}, \mathbf{p}, \mathbf{q}\}, \{\mathbf{v}, \mathbf{r}, \mathbf{s}\}\}$ -modularisable and it has multiple connectors. In fact,  $\mathbf{p}$  and  $\mathbf{q}$  are the connectors for the module  $\{\mathbf{u}, \mathbf{p}, \mathbf{q}\}$ , whereas  $\mathbf{r}$  and  $\mathbf{s}$  are the connectors for the module  $\{\mathbf{v}, \mathbf{r}, \mathbf{s}\}$ .  $\diamond$

In actual programming, refining a modularisation enables to enhance parameters like scalability, maintenance, reusability and many more. In the present setting, it also allows for simpler typings. Any modularisation refinement corresponds to a partition refinement.

**Definition 3.6 (Refinement)** A partition  $\mathcal{P}$  refines a partition  $\mathcal{P}'$  (notation  $\mathcal{P} \sqsubseteq \mathcal{P}'$ ) if  $\mathcal{P} = \{\mathbf{P}_h\}_{h \in H}$ ,  $\mathcal{P}' = \{\mathbf{P}'_k\}_{k \in K}$  and for all  $k \in K$  there is  $H_k \subseteq H$  such that  $\mathbf{P}'_k = \bigcup_{h \in H_k} \mathbf{P}_h$ .

As intuitively evident, it is possible to formally show that coarser partitions maintain modularisability.

**Lemma 3.7** If  $\mathbb{M}$  is  $\mathcal{P}$ -modularisable and  $\mathcal{P}$  refines  $\mathcal{P}'$ , then  $\mathbb{M}$  is  $\mathcal{P}'$ -modularisable.

*Proof.* It is not difficult to verify that if conditions (i) and (ii) of Definition 3.4 hold for  $\mathcal{P}$ , then they hold also for  $\mathcal{P}'$ .  $\square$

From the previous lemma, being  $\mathbb{M}$   $\{\text{prt}(\mathbb{M})\}$ -modularisable, it immediately follows that  $\mathbb{M}$  is  $\mathcal{P}$ -modularisable for all  $\mathcal{P}$  such that  $\text{prt}(\mathbb{M}) \subseteq \mathbf{P}$  for some  $\mathbf{P} \in \mathcal{P}$ .

We can build a minimal refined partition, with respect to  $\sqsubseteq$ , among those modularising a session. We start with  $\mathcal{P}_0 = \{\{\mathbf{p}\} \mid \mathbf{p} \in \text{prt}(\mathbb{M})\}$  and we iteratively build  $\mathcal{P}_{i+1}$  by replacing in  $\mathcal{P}_i$  the two sets  $\mathbf{P}, \mathbf{P}'$  with the unique set  $\mathbf{P} \cup \mathbf{P}'$  if there is  $\mathbf{p}[P] \in \mathbb{M}$  such that  $\mathbf{p} \in \mathbf{P}$ ,  $\mathbf{q} \in \mathbf{P}' \cap \text{prt}(P)$  and either  $P$  is not  $\mathbf{P}$ -connecting or  $\mathbf{q}$  is not a connector for the subsession of  $\mathbb{M}$  whose set of participants is  $\mathbf{P}'$ . It is possible to verify that  $\mathbb{M}$  is  $\mathcal{P}$ -modularisable, where  $\mathcal{P}$  is the fixed point of this procedure. We show now such a partition to be also the minimum among the partitions modularising a session.

**Lemma 3.8** The minimal partition modularising a session is unique, i.e. a minimum.

*Proof.* Assume toward a contradiction that there are two different minimal (w.r.t  $\sqsubseteq$ ) partitions  $\mathcal{P} = \{\mathbf{P}_h\}_{h \in H}$  and  $\mathcal{P}' = \{\mathbf{P}'_k\}_{k \in K}$  for modularising a session. This implies that there are  $h_1, h_2 \in H$  and  $k_0 \in K$  such that  $\mathbf{p} \in \mathbf{P}_{h_1}$ ,  $\mathbf{q} \in \mathbf{P}_{h_2}$  and  $\{\mathbf{p}, \mathbf{q}\} \subseteq \mathbf{P}'_{k_0}$ . Then also the partition obtained from  $\mathcal{P}'$  by replacing the set  $\mathbf{P}'_{k_0}$  with the two sets  $\mathbf{P}'_{k_0} \cap \mathbf{P}_{h_1}$  and  $\mathbf{P}'_{k_0} \cap \mathbf{P}_{h_2}$  modularises the same session. This is clearly a contradiction.  $\square$

It is crucial that  $\mathcal{P}$ -modularisation is preserved by reduction.

**Lemma 3.9** *Let  $\mathbb{M}$  be  $\mathcal{P}$ -modularisable and let  $\mathbb{M} \xrightarrow{\Delta} \mathbb{M}'$ . Then also  $\mathbb{M}'$  is  $\mathcal{P}$ -modularisable.*

*Proof.* Notice that conditions (i) and (ii) of Definition 3.4 are invariant by reduction. In fact a participant  $\mathbf{p}$  which is a connector in  $\mathbb{M}$  is either a connector in  $\mathbb{M}'$  too or it has a process whose participants all belong to the subsession containing  $\mathbf{p}$  in the  $\mathcal{P}$ -modulation of  $\mathbb{M}'$ . Therefore a  $\mathcal{P}$ -modularisation of  $\mathbb{M}$  is also a  $\mathcal{P}$ -modularisation of  $\mathbb{M}'$ .  $\square$

We can notice that, if  $\mathcal{P}$  is the minimal partition for modularising  $\mathbb{M}$  and  $\mathbb{M} \xrightarrow{\Delta} \mathbb{M}'$ , in general  $\mathcal{P}$  is not the minimal partition for modularising  $\mathbb{M}'$ . In fact  $\text{prt}(\mathbb{M}')$  can be a proper subset of  $\text{prt}(\mathbb{M})$  and a process in  $\mathbb{M}$  which is not a connector can reduce to a process which is a connector in  $\mathbb{M}'$ . For example the minimal partition of  $\mathbb{M} \equiv \mathbf{p}[\mathbf{q}!\lambda.r!\lambda + r?\lambda'.(\mathbf{q}!\lambda_1 + \mathbf{q}?\lambda_2)] \parallel \mathbf{q}[\mathbf{p}?\lambda + \mathbf{p}?\lambda_1 + \mathbf{p}!\lambda_2] \parallel \mathbf{r}[\mathbf{p}?\lambda + \mathbf{p}!\lambda']$  is  $\{\{\mathbf{p}, \mathbf{q}, \mathbf{r}\}\}$ , since  $\mathbf{p}$  is not a connector. But  $\mathbb{M} \xrightarrow{r!\lambda'p} \mathbf{p}[\mathbf{q}!\lambda_1 + \mathbf{q}?\lambda_2] \parallel \mathbf{q}[\mathbf{p}?\lambda + \mathbf{p}?\lambda_1 + \mathbf{p}!\lambda_2]$  and the minimal partition of this last session is  $\{\{\mathbf{p}\}, \{\mathbf{q}\}\}$ , since  $\mathbf{p}$  and  $\mathbf{q}$  are connectors and  $\mathbf{r}$  disappeared.

**Example 3.10 (Modular election)** We consider three “local” elections, all managed like in the Example 2.4. The names of the participants of the three local election are like the ones in the Example 2.4, but indexed with indexes in  $\{1, 2, 3\}$ . We also consider a further “global election” with participants  $\mathbf{w}_1, \mathbf{w}_2$  and  $\mathbf{w}_3$  and global station  $\mathbf{gs}$ . Such an election follows a protocol similar to that of the local elections (but simpler, since only three participants do compete for leadership). It can be seen as an election among the winners of the local elections. In the present example the “local leaders”, once they are elected, are informed whether they have been elected also “global leader” or not. The above sketched global behaviour is implemented by the following session  $\mathbb{E}^{gl}$  made of four subsessions: three local elections ( $\mathbb{E}_1, \mathbb{E}_2, \mathbb{E}_3$ ) and one global election  $\mathbb{G}$  among the “local” winners.

The processes of the participants, are fairly similar to the ones in the Example 2.4 to which a part is added implementing the communication to the local leaders of whether they are also global leader or not.

$$\mathbb{E}^{gl} = \mathbb{E}_1 \parallel \mathbb{E}_2 \parallel \mathbb{E}_3 \parallel \mathbb{G}$$

where, for  $1 \leq i \leq 3$ ,

$$\mathbb{E}_i = \mathbf{a}_i[\mathbf{P}_{\mathbf{a}_i}] \parallel \mathbf{b}_i[\mathbf{P}_{\mathbf{b}_i}] \parallel \mathbf{c}_i[\mathbf{P}_{\mathbf{c}_i}] \parallel \mathbf{d}_i[\mathbf{P}_{\mathbf{d}_i}] \parallel \mathbf{e}_i[\mathbf{P}_{\mathbf{e}_i}] \parallel \mathbf{s}_i[\mathbf{P}_{\mathbf{s}_i}]$$

$$\begin{aligned} \text{with } \mathbf{P}_{\mathbf{a}_i} = & \mathbf{e}_i!\mathit{leader} \\ & + \mathbf{b}_i?\mathit{leader}.\left(\mathbf{c}_i!\mathit{leader} + \mathbf{d}_i?\mathit{leader}.\mathbf{s}_i!\mathit{elect}.\left(\mathbf{s}_i?\mathit{gleader} + \mathbf{s}_i?\mathit{no}\right)\right) \\ & + \mathbf{s}_i?\mathit{del} \end{aligned}$$

$$\begin{aligned} \text{and } \mathbf{P}_{\mathbf{s}_i} = & \sum_{\mathbf{x} \in \{\mathbf{a}_i, \mathbf{b}_i, \mathbf{c}_i, \mathbf{d}_i, \mathbf{e}_i\}} \mathbf{x}?\mathit{elect}.\left(\mathbf{gs}?\mathit{gleader}.\mathbf{x}!\mathit{gleader}.\mathbf{Q}_i + \mathbf{gs}?\mathit{no}.\mathbf{x}!\mathit{no}.\mathbf{Q}_i\right) \\ & \text{with } \mathbf{Q}_i = \sum_{\mathbf{x} \in \{\mathbf{a}_i, \mathbf{b}_i, \mathbf{c}_i, \mathbf{d}_i, \mathbf{e}_i\}} \mathbf{x}!\mathit{del} \end{aligned}$$

and with processes of participants  $\mathbf{b}_i, \mathbf{c}_i, \mathbf{d}_i$  and  $\mathbf{e}_i$  obtained out of  $\mathbf{P}_{\mathbf{a}_i}$  by applying the name substitution

$$\mathbf{v}_i = [\mathbf{a}_i \mapsto \mathbf{b}_i, \mathbf{b}_i \mapsto \mathbf{c}_i, \mathbf{c}_i \mapsto \mathbf{d}_i, \mathbf{d}_i \mapsto \mathbf{e}_i, \mathbf{e}_i \mapsto \mathbf{a}_i]$$
 as in Example 2.4

and




Figure 1: The three local elections and the global one of Example 3.10.

$$\mathbb{G} = w_1[P_{w_1}] \parallel w_2[P_{w_2}] \parallel w_3[P_{w_3}] \parallel gs[P_{gs}]$$

with  $P_{w_i} = w_{i+2}!leader$   
 $+ w_{i+1}?leader.gs!gleader$   
 $+ gs?del$

and  $P_{gs} = \sum_{i \in \{1,2,3\}} w_i?gleader.s_i!gleader.s_{i+1}.no.s_{i+2}.no.(\sum_{i \in \{1,2,3\}} w_i!del)$   
 where all the indexes above have to be considered modulo 3, plus 1.

It is not difficult to check that  $\mathbb{E}^{gl}$  is  $\mathcal{P}$ -modularisable for  $\mathcal{P} = \{\text{prt}(\mathbb{E}_1), \text{prt}(\mathbb{E}_2), \text{prt}(\mathbb{E}_3), \text{prt}(\mathbb{G})\}$ , where  $s_1, s_2, s_3$  and  $gs$  are the connectors for, respectively, the modules  $\mathbb{E}_1, \mathbb{E}_2, \mathbb{E}_3$  and  $\mathbb{G}$ .  $\diamond$

## 4 A Type System for Modular Sessions

Global types are used to represent the overall behaviour of multiparty sessions. Here we use a notion of global type similar to the one in MPST but, following SMPS, we define them coinductively, as possibly infinite regular terms.

**Definition 4.1 (Global types)** Global types are coinductively defined by:

$$G ::=^{coind} \text{End} \mid \sum_{i \in I} \Lambda_i.G_i$$

where  $I \neq \emptyset$  and finite, and for any  $j, l \in I$  such that  $j \neq l$ ,  $\Lambda_i = p\lambda_i;q$  and  $\Lambda_j = p\lambda_j;q$  imply  $\lambda_j \neq \lambda_l$ .

We define the participants of global types by  $\text{prt}(\text{End}) = \emptyset$  and  $\text{prt}(\sum_{i \in I} \Lambda_i.G_i) = \bigcup_{i \in I} \text{prt}(\Lambda_i) \cup \bigcup_{i \in I} \text{prt}(G_i)$ . By the regularity condition,  $\text{prt}(G)$  is finite for any  $G$ . As usual, trailing  $\text{End}$ 's will be omitted.

As mentioned in the Introduction, in standard SMPS, typing rules take into account single interactions between pairs of participants. In our mixed-choice setting, that approach would allow to type only sessions whose independent parts were intrinsically sequential, so ruling out protocols like the leader election and the modular election. We hence consider a typing rule where all the interactions involving

the participants of a single module  $\widehat{\mathbb{M}}$  (so including also the communications of the connectors of  $\widehat{\mathbb{M}}$  with other modules) are taken into account. Such a set of interactions is formalised through the notion of coherent set of communication labels. A module  $\widehat{\mathbb{M}}$  is therefore indirectly represented in the rule in terms of its corresponding coherent set and referred to, when necessary, as the *witness* of such a set.

**Definition 4.2 (Coherent set of communication labels)** *A set of labels  $\{\Lambda_i\}_{i \in I}$  is  $\mathcal{P}$ -coherent for  $\mathbb{M}$  if  $\mathbb{M}$  is  $\mathcal{P}$ -modularisable and there exists an element  $\widehat{\mathbb{M}}$  of the (unique)  $\mathcal{P}$ -modularisation of  $\mathbb{M}$  such that*

$$\{\Lambda_i\}_{i \in I} = \{\Lambda \in \mathcal{L}(\mathbb{M}) \mid \text{prt}(\Lambda) \cap \text{prt}(\widehat{\mathbb{M}}) \neq \emptyset\}$$

*The  $\widehat{\mathbb{M}}$  above is called witness for the  $\mathcal{P}$ -coherence of  $\{\Lambda_i\}_{i \in I}$ .*

**Example 4.3 (Witness)** Let us consider  $\mathbb{M}'_1 = u[q!\lambda] \parallel p[r!\lambda_1 + r?\lambda_2] \parallel q[u?\lambda.(s!\lambda_1 + s?\lambda_2)]$  and  $\mathbb{M}_2 = v[r!\lambda.s?\lambda + s?\lambda.r!\lambda] \parallel r[v?\lambda.(p?\lambda_1 + p!\lambda_2)] \parallel s[v!\lambda.(q?\lambda_1 + q!\lambda_2)]$ . The session  $\mathbb{M}'_1 \parallel \mathbb{M}_2$  can be obtained by reducing the session of Example 3.5. It is still  $\{\{u, p, q\}, \{v, r, s\}\}$ -modularisable and  $\mathbb{M}'_1$  is the witness for the  $\{\{u, p, q\}, \{v, r, s\}\}$ -coherent set of labels  $\{u\lambda q, p\lambda_1 r, r\lambda_2 p\}$ .  $\diamond$

Here and in the following, the double line indicates that the rules are interpreted coinductively.

**Definition 4.4 (Type system)** *The type system  $\vdash^{\mathcal{P}}$  is defined by the following axiom and rule, where sessions are considered modulo structural congruence:*

$$\begin{array}{c} \text{[END]} \frac{}{\text{End} \vdash^{\mathcal{P}} \mathbf{p}[0]} \\ \text{[TCOMM]} \frac{\mathbb{M} \xrightarrow{\Lambda_i} \mathbb{M}_i \quad G_i \vdash^{\mathcal{P}} \mathbb{M}_i \quad \forall i \in I \neq \emptyset \quad \{\Lambda_i\}_{i \in I} \text{ is } \mathcal{P}\text{-coherent for } \mathbb{M} \quad \text{prt}(\sum_{i \in I} \Lambda_i.G_i) = \text{prt}(\mathbb{M})}{\sum_{i \in I} \Lambda_i.G_i \vdash^{\mathcal{P}} \mathbb{M}} \end{array}$$

It is not difficult to check that we can derive the global type  $u\lambda q.p\lambda u.G_1 + p\lambda u.u\lambda q.G_1$ , where  $G_1 = v\lambda r.s\lambda v.G_2 + s\lambda v.v\lambda r.G_2$ ,  $G_2 = p\lambda_1 r.G_3 + r\lambda_2 p.G_3$ , and  $G_3 = q\lambda_1 s + s\lambda_2 q$ , for the session of Example 3.5.

The condition “ $\{\Lambda_i\}_{i \in I}$  is  $\mathcal{P}$ -coherent for  $\mathbb{M}$ ” is essential to get Subject Reduction. In fact, by allowing any subset of  $\mathcal{L}(\mathbb{M})$  as  $\{\Lambda_i\}_{i \in I}$  in [TCOMM], we could derive  $p\lambda' r.p\lambda q \vdash^{\mathcal{P}} \mathbb{M}_0$  for

$$\mathbb{M}_0 \equiv p[q!\lambda + r!\lambda'.q!\lambda] \parallel q[p?\lambda] \parallel r[p?\lambda']$$

regardless of  $\mathcal{P}$ . However, we would also have  $\mathbb{M}_0 \xrightarrow{p\lambda q} r[p?\lambda']$ , with  $r[p?\lambda']$  untypable. The above example also shows that, in order to get Subject Reduction it is necessary that, at any moment, a connector can interact with one other connector only. Let us assume to relax Definition 3.1 as follows

*Given a set of participants  $\mathbf{P}$ , we say that a process  $P$  is  $\mathbf{P}$ -connecting if for any subprocess of  $P$ , say  $\sum_{i \in I} \pi_i.P_i$ , we have that  $\text{prt}(\pi_j) \notin \mathbf{P}$  for some  $j \in I$  implies  $\text{prt}(\pi_i) \notin \mathbf{P}$  for all  $i \in I$ .*

This would imply  $\mathbb{M}_0$  above to be  $\{\{p\}, \{q\}, \{r\}\}$ -modularisable and all the participants would turn out to be connectors in their respective modules. Hence  $\mathbb{M}_0$  would be  $\vdash^{\{\{p\}, \{q\}, \{r\}\}}$  typable, whereas  $r[p?\lambda']$  would not.

The condition “ $\text{prt}(\sum_{i \in I} \Lambda_i.G_i) = \text{prt}(\mathbb{M})$ ” is necessary to get Lock Freedom. For example without this condition we could derive  $G \vdash^{\{\{p\}, \{q\}, \{r\}\}} p[P] \parallel q[Q] \parallel r[s!\lambda]$  with  $P = q!\lambda.P$ ,  $Q = p?\lambda.Q$  and  $G = p\lambda q.G$ . For what concerns the condition  $I \neq \emptyset$ , let us consider one of our previous examples, namely  $\mathbb{M} = p[q!\lambda + r!\lambda']$ . Such a session can be uniquely modularised with itself as possible module and it is not lock free. In fact it is not typable, since its unique coherent set is empty.

It can be proved that  $\mathcal{P}$ -coherence of a label set for a session is preserved by reducing the session with a label not belonging to the  $\mathcal{P}$ -coherent set, see Lemma 5.6.

The type system is decidable, since processes and global types are regular, and there is only a finite number of partitions for the participants of a session. More interesting, typability of a session does depend on the choice neither of the  $\mathcal{P}$ -coherent sets nor of the partition  $\mathcal{P}$  (see, respectively, Theorems 5.11 and 5.12).

We define the semantics of global types via a coinductive formal system, as done first in [6]. Such a coinductive definition enables to take into account global types containing branches, where some communications can be indefinitely procrastinated, see Example 4.7. In order to do that, it is handy to associate to a global type the set of communication labels which might (not necessarily) decorate its transitions. We dub them capabilities of the global type.

**Definition 4.5 (Capabilities)** *Capabilities of global types are defined by:*

$$\text{cap}(\text{End}) = \emptyset \quad \text{cap}(\Sigma_{i \in I} \Lambda_i . G_i) = \{\Lambda_i\}_{i \in I} \cup \bigcup_{i \in I} \text{cap}(G_i)$$

**Definition 4.6 (LTS for global types)** *The labelled transition system (LTS) for global types is specified by the following axiom and rule:*

$$\begin{array}{c} \text{[E-COMM]} \frac{}{\Sigma_{i \in I} \Lambda_i . G_i \xrightarrow{\Lambda_j} G_j} \quad j \in I \\ \text{[I-COMM]} \frac{G_i \xrightarrow{\Lambda} G'_i \quad \Lambda \in \text{cap}(G_i) \quad \text{prt}(\Lambda) \cap \text{prt}(\Lambda_i) = \emptyset \quad \forall i \in I}{\Sigma_{i \in I} \Lambda_i . G_i \xrightarrow{\Lambda} \Sigma_{i \in I} \Lambda_i . G'_i} \end{array}$$

Axiom [E-COMM] formalises the fact that, in a session exposing the behaviour  $\Sigma_{i \in I} \Lambda_i . G_i$ , the communication labelled  $\Lambda_j$  for any  $j \in I$  can happen. When such a communication is actually performed, the resulting session will expose the behaviour  $G_j$ .

Rule [I-COMM] enables to describe independent and concurrent communications, even if global types apparently look like sequential descriptions of sessions' overall behaviours. In fact, behaviours involving participants ready to interact with each other uniformly in all branches of a global type, can do that if neither of them is involved in an interaction appearing at top level in the global type. The condition  $\Lambda \in \text{cap}(G_i)$  in Rule [I-COMM] is needed because such a rule is coinductive. In fact, without such a condition, we could get the following infinite derivation for  $G \xrightarrow{p\lambda q} G$  with  $G = r\lambda's.G$ :

$$\mathcal{D} = \frac{\mathcal{D}}{G \xrightarrow{p\lambda q} G} \text{[I-COMM]}$$

**Example 4.7 (Use of coinduction in Rule [I-COMM])** As shown in [6] the coinductive formulation of Rule [I-COMM] allows to get  $G \xrightarrow{p\lambda q} G'$ , where  $G = r\lambda_1s.G + r\lambda_2s.p\lambda q$  and  $G' = r\lambda_1s.G' + r\lambda_2s$ . The inductive definition of this rule does not allow the shown transition.  $\diamond$

**Example 4.8 (Typing the modular election session)** In Figure 2 we provide a typing for the modular session  $\mathbb{E}^{gl}$  of Example 3.10 in the type system  $\vdash^{\mathcal{P}}$  with

$$\mathcal{P} = \bigcup_{i \in \{1,2,3\}} \{\{a_i, b_i, c_i, d_i, e_i, s_i\}\} \cup \{\{w_1, w_2, w_3, gs\}\}$$

We use colours for representing reduced participants. In particular, a participant has the colour

- |  |   |  |
|--|---|--|
| $\square$ in its initial state;          | $\blacksquare$ after one interaction;   | $\blacksquare$ after two interactions; |
| $\blacksquare$ after three interactions; | $\blacksquare$ after four interactions; | $\blacksquare$ when terminated.        |

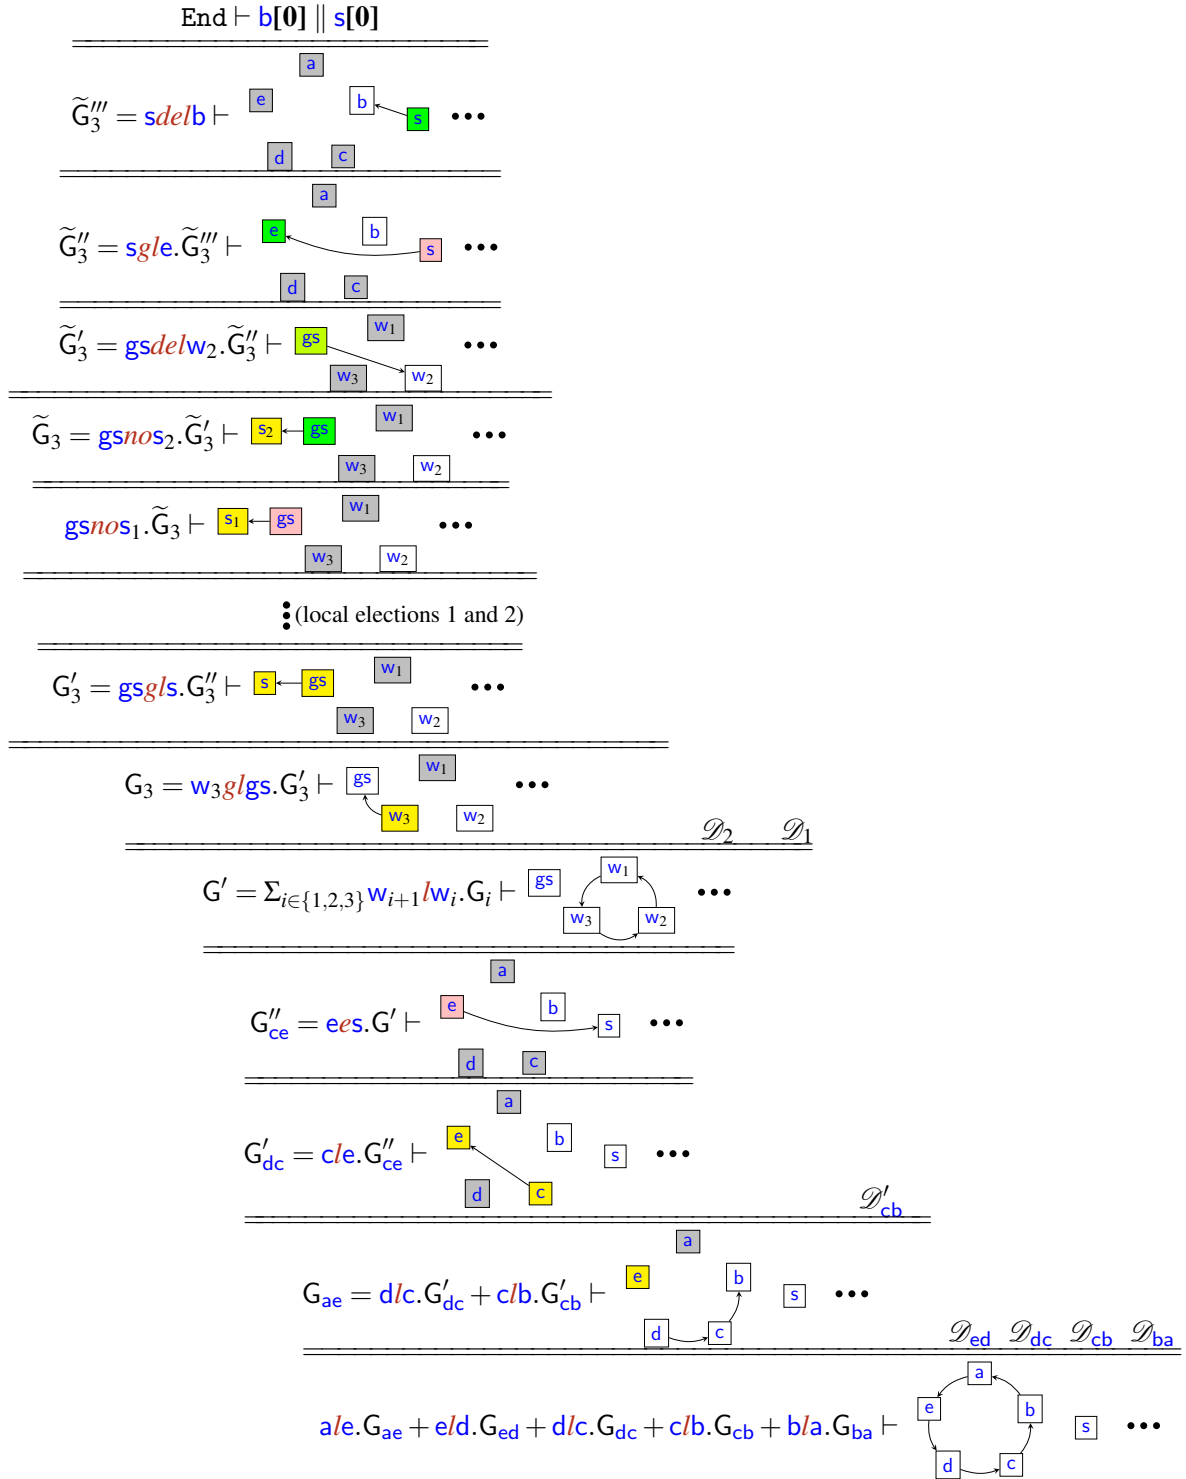


Figure 2: A type derivation for Example 3.10.

For instance:

$$\begin{aligned} \boxed{\text{gs}} &= \text{gs}[\sum_{i \in \{1,2,3\}} w_i ? \text{gleader}.s_i ! \text{gleader}.s_{i+1} ! \text{no}.s_{i+2} ! \text{no}.(\sum_{i \in \{1,2,3\}} w_i ! \text{del})] \\ \boxed{\text{gs}} &= \text{gs}[s_3 ! \text{gleader}.s_1 ! \text{no}.s_2 ! \text{no}.(\sum_{i \in \{1,2,3\}} w_i ! \text{del})] & \boxed{\text{gs}} &= \text{gs}[s_1 ! \text{no}.s_2 ! \text{no}.(\sum_{i \in \{1,2,3\}} w_i ! \text{del})] \\ \boxed{\text{gs}} &= \text{gs}[s_2 ! \text{no}.(\sum_{i \in \{1,2,3\}} w_i ! \text{del})] & \boxed{\text{gs}} &= \text{gs}[\sum_{i \in \{1,2,3\}} w_i ! \text{del}] & \boxed{\text{gs}} &= \text{gs}[\mathbf{0}] \end{aligned}$$

In Figure 2, arrows do connect pair of participants forming a redex. Moreover, in the conclusion of the rules we show only the participants of the module containing the coherent set of reductions and, in case, the “external” connectors. The rest of the session will be denoted by “•••”. Also, the figure shows only one branch of the typing derivation tree: the one concerning the global election of  $e_3$ .

For the sake of readability,  $l$  and  $gl$  are abbreviation for, respectively,  $leader$  and  $gleader$ . Moreover,  $a, b, c, d, e$  and  $s$  stand for  $a_3, b_3, c_3, d_3, e_3$  and  $s_3$ .  $\diamond$

## 5 Properties

A subsession of the shape  $p[q!\lambda.P \dot{\vdash} P'] \parallel q[p?\lambda.Q \dot{\vdash} Q']$  is called a *redex* and  $p[P] \parallel q[Q]$  is the *contractum* of the redex. In a transition labelled by  $p\lambda q$  both the redex and the contractum are uniquely determined.

**Lemma 5.1** *If  $\mathbb{M} \xrightarrow{p\lambda q} \mathbb{M}'$ , then there exists a unique redex  $p[q!\lambda.P \dot{\vdash} P'] \parallel q[p?\lambda.Q \dot{\vdash} Q']$  such that*

$$\mathbb{M} \equiv p[q!\lambda.P \dot{\vdash} P'] \parallel q[p?\lambda.Q \dot{\vdash} Q'] \parallel \mathbb{M}''$$

*and  $\mathbb{M}' \equiv p[P] \parallel q[Q] \parallel \mathbb{M}''$ .*

*Proof.* Immediate by the definition of session LTS.  $\square$

Rule [COMM] in Definition 2.3 entails an easy relation between the participants connected by reductions in a session.

**Lemma 5.2** *If  $\mathbb{M} \xrightarrow{\Lambda} \mathbb{M}'$ , then  $\text{prt}(\mathbb{M}) = \text{prt}(\Lambda) \cup \text{prt}(\mathbb{M}')$ .*

It is not difficult to check that the participants of a session and of its global type are the same.

**Lemma 5.3** *If  $G \vdash^{\mathcal{P}} \mathbb{M}$ , then  $\text{prt}(G) = \text{prt}(\mathbb{M})$ .*

The following technical lemma relating capabilities and possible reductions of a global type will be handy later on.

**Lemma 5.4** *If  $G \xrightarrow{\Lambda} G'$ , then  $\Lambda \in \text{cap}(G)$ .*

*Proof.* By cases on the applied axiom/rule justifying  $G \xrightarrow{\Lambda} G'$ . If this is [E-COMM], then  $G = \sum_{i \in I} \Lambda_i.G_i$  and  $\Lambda = \Lambda_j$  for some  $j \in I$  and  $\Lambda_j \in \text{cap}(\sum_{i \in I} \Lambda_i.G_i)$  by Definition 4.5.

Otherwise,  $G = \sum_{i \in I} \Lambda_i.G_i$  and  $G' = \sum_{i \in I} \Lambda_i.G'_i$  by Rule [I-COMM], where  $G_i \xrightarrow{\Lambda} G'_i$ ,  $\Lambda \in \text{cap}(G_i)$  and  $\text{prt}(\Lambda) \cap \text{prt}(\Lambda_i) = \emptyset$  for all  $i \in I$ . This implies  $\Lambda \in \text{cap}(G)$ , since  $\bigcup_{i \in I} \text{cap}(G_i) \subseteq \text{cap}(G)$  by Definition 4.5.  $\square$

For showing Subject Reduction it is crucial to ensure that the  $\mathcal{P}$ -coherence of a set of labels is preserved by reducing a label not belonging to this set, see Lemma 5.6 whose proof uses Lemma 5.5 below.

**Lemma 5.5** *Let  $\{\Lambda_i\}_{i \in I}$  be  $\mathcal{P}$ -coherent for  $\mathbb{M}$  and let  $\Lambda \in \mathcal{L}(\mathbb{M})$ . Moreover, let  $\Lambda \neq \Lambda_i$  for all  $i \in I$ . Then  $\text{prt}(\Lambda) \cap \text{prt}(\Lambda_i) = \emptyset$  for all  $i \in I$ .*

*Proof.* By definition of coherence (Definition 4.2), we have a subsession  $\widehat{M}$  of  $M$  witnessing the  $\mathcal{P}$ -coherence of  $\{\Lambda_i\}_{i \in I}$ . By contradiction, let us assume  $\text{prt}(\Lambda) \cap \text{prt}(\Lambda_j) \neq \emptyset$  for some  $j \in I$ . By definition of  $\mathcal{P}$ -modularisation, this implies that  $\text{prt}(\Lambda_j)$  contains a connector  $\mathbf{p}$  of  $\widehat{M}$ . Let  $\text{prt}(\Lambda_j) = \{\mathbf{p}, \mathbf{q}\}$  with  $\mathbf{q} \notin \text{prt}(\widehat{M})$  and  $\text{prt}(\Lambda) = \{\mathbf{q}, \mathbf{r}\}$ . Then  $\mathbf{q} \in \text{prt}(M)$  and the process  $Q$  of  $\mathbf{q}$  must have a choice between a communication with  $\mathbf{p}$  and a communication with  $\mathbf{r}$ . But this is impossible, since  $\mathbf{q}$  must be a connector for some subsession  $\widehat{M}'$  of  $M$  by condition (ii) of Definition 3.4 and then the process  $Q$  must be  $\text{prt}(\widehat{M}')$ -connecting by Definition 3.2.  $\square$

**Lemma 5.6 (Coherence preservation)** *Let  $\{\Lambda_i\}_{i \in I}$  be  $\mathcal{P}$ -coherent for  $M$  and let  $M \xrightarrow{\Lambda} M'$ . Then  $\Lambda \notin \{\Lambda_i\}_{i \in I}$  implies that  $\{\Lambda_i\}_{i \in I}$  is  $\mathcal{P}$ -coherent for  $M'$  as well.*

*Proof.* Let  $\widehat{M}$  be the subsession of  $M$  witnessing the  $\mathcal{P}$ -coherence of  $\{\Lambda_i\}_{i \in I}$  for  $M$ . From  $\Lambda \notin \{\Lambda_i\}_{i \in I}$  and Lemma 5.5 we get that  $\text{prt}(\Lambda) \cap \text{prt}(\Lambda_i) = \emptyset$  for all  $i \in I$ . This implies that the reduction  $\xrightarrow{\Lambda}$  cannot affect any reduction with label in  $\{\Lambda_i\}_{i \in I}$ . Hence  $\widehat{M}$  is a witness of the  $\mathcal{P}$ -coherence of  $\{\Lambda_i\}_{i \in I}$  also for  $M'$ .  $\square$

Notice how the conditions on connectors (Definitions 3.1 and 3.2) are crucial in getting the property that  $\text{prt}(\Lambda) \cap \text{prt}(\Lambda_i) = \emptyset$  for all  $i \in I$  in the above result of coherence preservation. If we allowed connectors to communicate with external partners having unrestricted processes we could consider the  $\{\{\mathbf{p}\}, \{\mathbf{r}, \mathbf{s}\}\}$ -modularisation of  $M = \mathbf{p}[r!\lambda] \parallel \mathbf{r}[p?\lambda + s?\lambda] \parallel \mathbf{s}[r!\lambda]$ . In such a case, the set  $\{\mathbf{p}\lambda\mathbf{r}\}$  would be  $\{\{\mathbf{p}\}, \{\mathbf{r}, \mathbf{s}\}\}$ -coherent with witness  $M' = \mathbf{p}[r!\lambda]$ . However, we would also have that  $M \xrightarrow{s\lambda r} M'$ , but  $\{\mathbf{p}\lambda\mathbf{r}\}$  would not be  $\{\{\mathbf{p}\}, \{\mathbf{r}, \mathbf{s}\}\}$ -coherent for  $M'$ , since  $\mathbf{p}\lambda\mathbf{r} \notin \mathcal{L}(M')$ .

**Theorem 5.7 (Subject Reduction)** *If  $G \vdash^{\mathcal{P}} M$  and  $M \xrightarrow{\Lambda} M'$ , then  $G' \vdash^{\mathcal{P}} M'$  and  $G \xrightarrow{\Lambda} G'$  for some  $G'$ .*

*Proof.* By coinduction on the derivation of  $G \vdash^{\mathcal{P}} M$ . Let  $\Lambda = \mathbf{p}\lambda\mathbf{q}$ . By Lemma 5.1, if  $M \xrightarrow{\Lambda} M'$ , then there exists a unique redex

$$\mathcal{R} = \mathbf{p}[\mathbf{q}!\lambda.P \dot{\vdash} P'] \parallel \mathbf{q}[\mathbf{p}?\lambda.Q \dot{\vdash} Q']$$

such that  $M \equiv \mathcal{R} \parallel M''$  and  $M' \equiv \mathbf{p}[P] \parallel \mathbf{q}[Q] \parallel M''$  for some  $M''$ . By the hypothesis that  $G \vdash^{\mathcal{P}} M$  we know that  $G$  is of the form  $\Sigma_{i \in I} \Lambda_i.G_i$  and the derivation ends by

$$\text{[TCOMM]} \frac{\frac{M \xrightarrow{\Lambda} M_i \quad G_i \vdash^{\mathcal{P}} M_i \quad \forall i \in I \neq \emptyset}{\{\Lambda_i\}_{i \in I} \text{ is } \mathcal{P}\text{-coherent for } M} \quad \text{prt}(\Sigma_{i \in I} \Lambda_i.G_i) = \text{prt}(M)}{\Sigma_{i \in I} \Lambda_i.G_i \vdash^{\mathcal{P}} M}$$

We proceed by distinguishing the two possible following cases.

*Case  $\mathbf{p}\lambda\mathbf{q} = \Lambda_j$  for some  $j \in I$ .* By the premises of the rule, we have  $M \xrightarrow{\Lambda} M_j$  and  $G_j \vdash^{\mathcal{P}} M_j$ , where  $M' = M_j$ . Moreover, it immediately follows that  $G \xrightarrow{\Lambda} G_j$  by Axiom [E-COMM].

*Case  $\mathbf{p}\lambda\mathbf{q} \neq \Lambda_i$  for all  $i \in I$ .* We have that, for all  $i \in I$ ,  $M_i \equiv \mathcal{R} \parallel M'_i$  for some  $M'_i$ . Hence we get that  $M_i \xrightarrow{\Lambda} \mathbf{p}[P] \parallel \mathbf{q}[Q] \parallel M'_i$  for all  $i \in I$ . Moreover, for all  $i \in I$ ,  $M'' \xrightarrow{\Lambda_i} M'_i$ . By the coinduction hypothesis on the premises of the rule, we have that, for all  $i \in I$ ,  $G'_i \vdash^{\mathcal{P}} \mathbf{p}[P] \parallel \mathbf{q}[Q] \parallel M'_i$  and  $G_i \xrightarrow{\Lambda} G'_i$  for some  $G'_i$ . Now, by Lemma 5.4, we get that  $\Lambda \in \text{cap}(G_i)$  for all  $i \in I$ , hence we have that  $G \xrightarrow{\Lambda} \Sigma_{i \in I} \Lambda_i.G'_i$  by Rule [I-COMM]. From  $M \xrightarrow{\Lambda} M' \equiv \mathbf{p}[P] \parallel \mathbf{q}[Q] \parallel M''$  and  $M'' \xrightarrow{\Lambda_i} M'_i$  we get  $M' \xrightarrow{\Lambda_i} \mathbf{p}[P] \parallel \mathbf{q}[Q] \parallel M'_i$  for all  $i \in I$ , which imply, by Lemma 5.2,

$$\text{prt}(M') = \bigcup_{i \in I} \text{prt}(\Lambda_i) \cup \bigcup_{i \in I} \text{prt}(\mathbf{p}[P] \parallel \mathbf{q}[Q] \parallel M'_i)$$

By Lemma 5.3,  $G'_i \vdash^{\mathcal{P}} \mathbf{p}[P] \parallel \mathbf{q}[Q] \parallel M'_i$  gives  $\text{prt}(G'_i) = \text{prt}(\mathbf{p}[P] \parallel \mathbf{q}[Q] \parallel M'_i)$  for all  $i \in I$ . Hence  $\text{prt}(\Sigma_{i \in I} \Lambda_i.G'_i) = \bigcup_{i \in I} \text{prt}(\Lambda_i) \cup \bigcup_{i \in I} \text{prt}(G'_i) = \text{prt}(M')$ . Moreover, from Lemma 5.6 and  $\mathbf{p}\lambda\mathbf{q} \neq \Lambda_i$  for all  $i \in I$ , it follows that  $\{\Lambda_i\}_{i \in I}$  is  $\mathcal{P}$ -coherent for  $M'$  as well. Therefore Rule [TCOMM] applies, namely

$$\frac{\begin{array}{c} M' \xrightarrow{\Lambda_i} \mathbf{p}[P] \parallel \mathbf{q}[Q] \parallel M'_i \quad G'_i \vdash^{\mathcal{P}} \mathbf{p}[P] \parallel \mathbf{q}[Q] \parallel M'_i \quad \forall i \in I \neq \emptyset \\ \{\Lambda_i\}_{i \in I} \text{ is } \mathcal{P}\text{-coherent for } M' \quad \text{prt}(\Sigma_{i \in I} \Lambda_i.G'_i) = \text{prt}(M') \end{array}}{\Sigma_{i \in I} \Lambda_i.G'_i \vdash^{\mathcal{P}} M'} \quad \square$$

**Theorem 5.8 (Session Fidelity)** *If  $G \vdash M$  and  $G \xrightarrow{\Lambda} G'$ , then  $M \xrightarrow{\Lambda} M'$  and  $G' \vdash M'$  for some  $M'$ .*

*Proof.* By coinduction on the derivation of  $G \xrightarrow{\Lambda} G'$ . We distinguish two cases according to the axiom/rule justifying  $G \xrightarrow{\Lambda} G'$ .

*Axiom [E-COMM]:* then  $G = \Sigma_{i \in I} \Lambda_i.G_i$ ,  $\Lambda = \Lambda_j$  and  $G' = G_j$  for some  $j \in I$ . Since  $G \neq \text{End}$ , the last rule in the derivation of  $G \vdash M$  must be [TCOMM], which implies that  $\Lambda = \mathbf{p}\lambda\mathbf{q}$  for some  $\mathbf{p}$ ,  $\lambda$  and  $\mathbf{q}$  such that

$$M \equiv \mathbf{p}[\mathbf{q}!\lambda.P \dot{\vdash} P'] \parallel \mathbf{q}[\mathbf{p}?\lambda.Q \dot{\vdash} Q'] \parallel M_0 \xrightarrow{\mathbf{p}\lambda\mathbf{q}} \mathbf{p}[P] \parallel \mathbf{q}[Q] \parallel M_0 \equiv M'$$

for some  $M_0$ , and  $G' \vdash^{\mathcal{P}} M'$ .

*Rule [I-COMM]:* then  $G = \Sigma_{i \in I} \Lambda_i.G_i$  and  $G' = \Sigma_{i \in I} \Lambda_i.G'_i$  with  $G_i \xrightarrow{\Lambda} G'_i$  and  $\Lambda \in \text{cap}(G_i)$  and  $\text{prt}(\Lambda) \cap \text{prt}(\Lambda_i) = \emptyset$  for all  $i \in I$ .

Since the last rule in the derivation of  $G \vdash M$  must be [TCOMM], it follows that

- $\{\Lambda_i\}_{i \in I}$  is  $\mathcal{P}$ -coherent for  $M$ ;
- $M \xrightarrow{\Lambda} M_i$  and  $G_i \vdash^{\mathcal{P}} M_i$ , for all  $i \in I \neq \emptyset$ ;
- $\text{prt}(\Sigma_{i \in I} \Lambda_i.G_i) = \text{prt}(M)$ .

By the coinduction hypothesis, we know that, for each  $i \in I$ , there exists  $M'_i$  such that

$$M_i \xrightarrow{\Lambda} M'_i \quad \text{and} \quad G'_i \vdash M'_i$$

Notice that, being the label  $\Lambda$  the same for all these reductions, by Lemma 5.1 there exists a unique redex

$$\mathbf{p}[\mathbf{q}!\lambda.P \dot{\vdash} P'] \parallel \mathbf{q}[\mathbf{p}?\lambda.Q \dot{\vdash} Q']$$

with contractum  $\mathbf{p}[P] \parallel \mathbf{q}[Q]$  in all the  $M_i$ , such that  $\Lambda = \mathbf{p}\lambda\mathbf{q}$ . On the other hand, since we know that  $\text{prt}(\Lambda) \cap \text{prt}(\Lambda_i) = \emptyset$  for all  $i \in I$ , it must be the case that  $\Lambda_i = r_i\lambda_i s_i$  and

$$M'_i \equiv r_i[R_i] \parallel s_i[S_i] \parallel M''_i$$

for some  $r_i, s_i, R_i, S_i, M''_i$  and for all  $i \in I$ . Hence, since  $M_i \xrightarrow{\mathbf{p}\lambda\mathbf{q}} M'_i$ , we have that, for each  $i \in I$ ,

$$M \equiv r_i[s_i!\lambda_i.R_i \dot{\vdash} R'_i] \parallel s_i[r?\lambda_i.S_i \dot{\vdash} S'_i] \parallel M''_i \xrightarrow{\mathbf{p}\lambda\mathbf{q}} r_i[s_i!\lambda_i.R_i \dot{\vdash} R'_i] \parallel s_i[r?\lambda_i.S_i \dot{\vdash} S'_i] \parallel M'''_i \equiv M'$$

for some  $R'_i, S'_i$  (if any),  $M'''_i$  and for all  $i \in I$ .

By Lemma 5.3,  $G'_i \vdash r_i[R_i] \parallel s_i[S_i] \parallel M''_i$  implies  $\text{prt}(G'_i) = \text{prt}(r_i[R_i] \parallel s_i[S_i] \parallel M''_i)$  and then  $\text{prt}(G') = \text{prt}(M')$ . Moreover, from  $\text{prt}(\Lambda) \cap \text{prt}(\Lambda_i) = \emptyset$  for all  $i \in I$ , we immediately get that  $\Lambda \notin \{\Lambda_i\}_{i \in I}$ . So, by Lemma 5.6 we get that  $\{\Lambda_i\}_{i \in I}$  is  $\mathcal{P}$ -coherent for  $M'$ . We conclude that there exists a derivation ending by the following application of Rule [TCOMM]

$$\text{[TCOMM]} \frac{\begin{array}{c} M' \xrightarrow{\Lambda} M'_i \quad G'_i \vdash M'_i \quad \forall i \in I \neq \emptyset \\ \{\Lambda_i\}_{i \in I} \text{ is } \mathcal{P}\text{-coherent for } M' \quad \text{prt}(G') = \text{prt}(M') \end{array}}{G' \vdash^{\mathcal{P}} M'} \quad \square$$

Toward establishing the property that typable sessions are lock free, we first prove the following lemma. In words, if  $\mathbf{p} \in \text{prt}(\mathbf{G})$ , then it must occur somewhere in its syntactic tree, hence there is a trace  $\sigma \cdot \Lambda$  out of  $\mathbf{G}$ , consisting just of external communications, which corresponds to a path in the tree ending by the first communication label  $\Lambda$  involving  $\mathbf{p}$ .

**Lemma 5.9** *If  $\mathbf{p} \in \text{prt}(\mathbf{G})$ , then there are  $\sigma$ ,  $\Lambda$  and  $\mathbf{G}'$  such that  $\mathbf{G} \xrightarrow{\sigma \cdot \Lambda} \mathbf{G}'$ ,  $\mathbf{p} \notin \text{prt}(\sigma)$  and  $\mathbf{p} \in \text{prt}(\Lambda)$ .*

*Proof.* The proof is by coinduction on  $\mathbf{G}$ . Since  $\mathbf{p} \in \text{prt}(\mathbf{G})$  we have that  $\mathbf{G} = \Sigma_{i \in I} \Lambda_i \cdot \mathbf{G}_i$ . Now, let us assume  $\mathbf{p} \in \bigcup_{i \in I} \text{prt}(\Lambda_i)$ . Without loss of generality, we can also assume that  $\mathbf{p} \in \text{prt}(\Lambda_j)$  for some  $j \in I$ . Then we immediately have that  $\mathbf{G} \xrightarrow{\Lambda_j} \mathbf{G}_j$  by Axiom [E-COMM], and the thesis trivially follows by taking  $\sigma = \varepsilon$ . Otherwise, since  $\mathbf{p} \in \text{prt}(\mathbf{G}) = \bigcup_{i \in I} \text{prt}(\Lambda_i) \cup \bigcup_{i \in I} \text{prt}(\mathbf{G}_i)$ , we have that  $\mathbf{p} \notin \bigcup_{i \in I} \text{prt}(\Lambda_i)$  implies  $\mathbf{p} \in \text{prt}(\mathbf{G}_j)$  for some  $j \in I$ . By the coinduction hypothesis, we have that there are a  $\sigma'$  and a  $\Lambda$  such that  $\mathbf{G}_j \xrightarrow{\sigma' \cdot \Lambda} \mathbf{G}'$ ,  $\mathbf{p} \notin \text{prt}(\sigma')$  and  $\mathbf{p} \in \text{prt}(\Lambda)$ . Then the thesis follows by setting  $\sigma = \Lambda_j \cdot \sigma'$ , since  $\mathbf{G} \xrightarrow{\Lambda_j} \mathbf{G}_j$  by Axiom [E-COMM] and  $\mathbf{G}_j \xrightarrow{\sigma' \cdot \Lambda} \mathbf{G}'$ .  $\square$

Observe that the last lemma is a sort of inverse implication w.r.t. Lemma 5.4, since it shows that the existence of a capability which is an actual communication of a global type  $\mathbf{G}$  follows by the fact that one of the involved participants is in  $\text{prt}(\mathbf{G})$ .

We are now in place to prove that typable sessions are lock free.

**Theorem 5.10 (Lock Freedom)** *If  $\mathbb{M}$  is typable, then  $\mathbb{M}$  is lock free.*

*Proof.* Let  $\mathbf{G} \vdash^{\mathcal{P}} \mathbb{M}$ . Following Definition 2.5, in order to prove Lock Freedom for  $\mathbb{M}$ , let  $\mathbb{M} \xrightarrow{\sigma} \mathbb{M}'$  for a finite  $\sigma$  and let  $\mathbf{p} \in \text{prt}(\mathbb{M}')$ . By Subject Reduction (Theorem 5.7) we get  $\mathbf{G}' \vdash^{\mathcal{P}} \mathbb{M}'$ . We can now recur to Lemma 5.3 and get  $\mathbf{p} \in \text{prt}(\mathbf{G}')$ . From the fact that  $\mathbf{p} \in \text{prt}(\mathbf{G}')$  and by Lemma 5.9 it follows that  $\mathbf{G}' \xrightarrow{\sigma' \cdot \Lambda} \mathbf{G}''$  for some  $\sigma'$  and  $\Lambda$  with  $\mathbf{p} \notin \text{prt}(\sigma')$  and  $\mathbf{p} \in \text{prt}(\Lambda)$ . Now the thesis follows by Session Fidelity (Theorem 5.8).  $\square$

We conclude this section by showing that typability of a session does depend on the choice neither of the  $\mathcal{P}$ -coherent sets nor of the partition  $\mathcal{P}$ .

**Theorem 5.11** *If  $\mathbb{M}$  is typable in  $\vdash^{\mathcal{P}}$  and  $\{\Lambda_i\}_{i \in I}$  is  $\mathcal{P}$ -coherent for  $\mathbb{M}$ , then there are  $\mathbf{G}_i$  for  $i \in I$  such that  $\Sigma_{i \in I} \Lambda_i \cdot \mathbf{G}_i \vdash^{\mathcal{P}} \mathbb{M}$ .*

*Proof.* The  $\mathcal{P}$ -coherence of  $\{\Lambda_i\}_{i \in I}$  for  $\mathbb{M}$  gives  $\mathbb{M} \xrightarrow{\Lambda_i} \mathbb{M}_i$ , which implies by Subject Reduction (Theorem 5.7)  $\mathbf{G}_i \vdash^{\mathcal{P}} \mathbb{M}_i$  for some  $\mathbf{G}_i$  and all  $i \in I$ . By Lemma 5.3  $\text{prt}(\mathbf{G}_i) = \text{prt}(\mathbb{M}_i)$  for all  $i \in I$ . From  $\mathbb{M} \xrightarrow{\Lambda_i} \mathbb{M}_i$  for all  $i \in I$  we get  $\text{prt}(\mathbb{M}) = \bigcup_{i \in I} \text{prt}(\Lambda_i) \cup \bigcup_{i \in I} \text{prt}(\mathbb{M}_i)$  by Lemma 5.2. By definition,  $\text{prt}(\Sigma_{i \in I} \Lambda_i \cdot \mathbf{G}_i) = \bigcup_{i \in I} \text{prt}(\Lambda_i) \cup \bigcup_{i \in I} \text{prt}(\mathbf{G}_i)$ . We conclude  $\text{prt}(\Sigma_{i \in I} \Lambda_i \cdot \mathbf{G}_i) = \text{prt}(\mathbb{M})$ , so we can derive  $\Sigma_{i \in I} \Lambda_i \cdot \mathbf{G}_i \vdash^{\mathcal{P}} \mathbb{M}$  using Rule [TCOMM].  $\square$

**Theorem 5.12** *If  $\mathbb{M}$  is typable in  $\vdash^{\mathcal{P}}$  and it is  $\mathcal{P}'$ -modularisable, then  $\mathbb{M}$  is typable in  $\vdash^{\mathcal{P}'}$  too.*

*Proof.* Since  $\mathbb{M}$  is typable in  $\vdash^{\mathcal{P}}$ , then  $\mathbb{M}$  is  $\mathcal{P}$ -modularisable by Definition 4.2. Let  $\mathcal{P} = \{\mathbf{P}_k\}_{k \in K}$  and  $\{\Lambda_h^k\}_{h \in H_k} = \{\Lambda \in \mathcal{L}(\mathbb{M}) \mid \text{prt}(\Lambda) \cap \mathbf{P}_k \neq \emptyset\}$ . By definition of partition we observe that  $\bigcup_{k \in K} \bigcup_{h \in H_k} \Lambda_h^k = \mathcal{L}(\mathbb{M})$ . By Definitions 4.2 and 3.4 for all  $k \in K$   $\{\Lambda_h^k\}_{h \in H_k}$  is  $\mathcal{P}$ -coherent for  $\mathbb{M}$ . By Lemma 5.11 for each  $k \in K$  there is a global type  $\mathbf{G}_k = \Sigma_{h \in H_k} \Lambda_h^k \cdot \widehat{\mathbf{G}}_h$  which can be assigned to  $\mathbb{M}$  in  $\vdash^{\mathcal{P}}$ , that is  $\Sigma_{h \in H_k} \Lambda_h^k \cdot \widehat{\mathbf{G}}_h \vdash^{\mathcal{P}} \mathbb{M}$ .

Let  $\{\Lambda_i\}_{i \in I}$  be  $\mathcal{P}'$ -coherent for  $\mathbb{M}$ . We proceed now by coinduction simultaneously on the derivations  $\Sigma_{h \in H_k} \Lambda_h^k \cdot \widehat{\mathbf{G}}_h \vdash^{\mathcal{P}} \mathbb{M}$  for all  $k \in K$ . From the above, for each  $i \in I$  there is  $k_i \in K$  and  $l_i \in H_{k_i}$  such that one

of the premises of the conclusion  $\Sigma_{h \in H_{k_i}} \Lambda_h^{k_i} \cdot \widehat{G}_h \vdash^{\mathcal{P}} \mathbb{M}$  is  $\widehat{G}_i \vdash^{\mathcal{P}} \mathbb{M}_i$  where  $\mathbb{M} \xrightarrow{\Lambda_i} \mathbb{M}_i$ . By coinduction we get  $G_i'' \vdash^{\mathcal{P}'} \mathbb{M}_i$  for some  $G_i''$  and for all  $i \in I$ . By Lemma 5.3  $\text{prt}(G_i'') = \text{prt}(\mathbb{M}_i)$  for all  $i \in I$ . We have  $\text{prt}(\Sigma_{i \in I} \Lambda_i \cdot G_i'') = \bigcup_{i \in I} \text{prt}(\Lambda_i) \cup \bigcup_{i \in I} \text{prt}(G_i'')$  and  $\text{prt}(\mathbb{M}) = \bigcup_{i \in I} \text{prt}(\Lambda_i) \cup \bigcup_{i \in I} \text{prt}(\mathbb{M}_i)$  by Lemma 5.2, so we get  $\text{prt}(\Sigma_{i \in I} \Lambda_i \cdot G_i'') = \text{prt}(\mathbb{M})$ . Then we can use Rule [TCOMM] to derive  $\Sigma_{i \in I} \Lambda_i \cdot G_i'' \vdash^{\mathcal{P}'} \mathbb{M}$ .  $\square$

## 6 Concluding Remarks, Related and Future Works

In the setting of the *message-passing* communication model, it is possible to envisage mechanisms of interaction where a process can be, at the very same time, both a potential sender and a potential receiver. In various frameworks for concurrent systems, like session types and communicating finite state machines, as well as the  $\pi$ -calculus, such mechanism is referred to as *mixed choice*. The flexibility and expressive power of mixed choice is understandably counterbalanced by a difficult control of the behaviour of systems. That was arguably the motivation that mostly restrained the session type community from pursuing a thorough investigation of this sort of interactions. A stimulus in that direction has been instead recently given by some papers like [10, 28, 29, 30, 31]. In particular, [10, 30] investigate mixed choice for binary session types, whereas [31] considers mixed choice in a MPST setting [18, 19] following the approach of [32] (global types are in fact not taken into account in [31]). Even if the main concern of [31] is the expressivity of multiparty calculi (according to the full range of possible restrictions of mixed choice), type systems assigning local types to processes are provided, where various predicates on contexts of local types are investigated. In [28, 29] binary sessions with timeout and mixed choice are enriched with a semantics guaranteeing Progress and a type system enjoying Subject Reduction.

Inspired by [31], we carry on an investigation on the use of mixed choice for synchronous communications in the setting of SMPS [12, 4]. In SMPS, global types are inferred for sessions, i.e. parallel compositions of named processes, the latter being an abstraction for both processes and local types usually considered in MPST. Our processes can use now mixed choice. Subject Reduction, Session Fidelity, as well as Lock Freedom are ensured for typable sessions. The most relevant aspect of our type system is that we look at sessions as implicitly composed by modules whose participants freely interact via unrestricted mixed choice, whereas the inter-module communications can be more easily controlled by allowing communications with only one participant. Such an approach does not discard a priori any session. In fact all sessions – even those developed without any specific modularisation in mind – can be modularised in a less or more refined way: from a single large module comprising the whole session, to a set of modules made by single participants. In the former case the typing is less effective, since the global type would result in a complete interaction tree, whereas in the latter case the typing coincides with the standard SMPS (with no mixed choice). In particular, our type system is conservative since, for processes without mixed choice, it coincides with the type system of [6], which is at present one of the most expressive. For mixed choice there is only the type system of [31], which is modularised by predicates on local types. An interesting property ensured by that type system is *safety*. Safety in [31] entails that the protocol for process interactions is such that, when a participant intends to perform an interaction, it nondeterministically chooses among all the participants that can interact with it. Then there is a nondeterministic decision concerning who has to play the role of the sender and who of the receiver. At that point, the sender performs an internal choice among the available outputs. Typing then guarantees that the possibility of interaction does not depend on the chosen output. We consider, instead, a simplified synchronisation protocol where there exists a nondeterministic choice among all the participants that can actually interact and all the possible communication interactions. As mentioned in Section 2, this approach makes “!” and “?” just two complementary synchronisation actions. It is however possible to

modify our type system in order to guarantee safety of sessions as defined in [31] by requiring that

$$p[q!\lambda.P \dot{\vdash} P'] \in \mathbb{M} \text{ and } q[p?\lambda'.Q \dot{\vdash} Q'] \in \mathbb{M} \text{ imply } p\lambda q \in \mathcal{L}(\mathbb{M})$$

The safety condition only ensures that an output finds the corresponding input if the receiver offers some inputs for the sender. An alternative condition is

$$p[q!\lambda.P \dot{\vdash} P'] \in \mathbb{M} \text{ implies } \mathbb{M} \xrightarrow{\sigma \cdot p\lambda q} \text{ for some } \sigma$$

With this last condition we would type less sessions, for example we would not type

$$p[q!\lambda + r!\lambda'] \parallel r[p?\lambda']$$

and the election example. The feature of this alternative condition is that the choice between outputs is internal, in agreement with an asynchronous implementation. In such a case it is worth remarking that a restriction of the subtyping relation used in [31] would be implicitly entailed by our typing.

As first pointed out in [14], the naive extension of the original type system [19] to sessions where input choices have different senders is unsound. In fact one can type sessions which reduce to untypable and stuck sessions. Suitable conditions ensuring a sound extension have been proposed both for the synchronous [14] and asynchronous [24, 11] communications. Notably, the type system proposed here does not have this problem.

We notice that modular sessions can be obtained by connecting independent sessions via gateways, according to the PaI approach to system composition. When composing several typable SMPS systems (through *compatible* interfaces) one gets a typable system [7]. Although the presence of mixed choice does not seem to be a major obstacle, it is unlikely a result like the one of [7] could be easily translated in the present context. In fact, as shown in Example 3.5 we allow the presence of multiple connectors per module. This possibility is actually a severe impediment to the safeness of PaI composition. In fact, let us take the following two typable sessions:

$$p[q!\lambda] \parallel q[p?\lambda] \quad r[s?\lambda] \parallel s[r!\lambda]$$

Taking all the participants as interfaces and considering that there exists a typable connection policy among them [7], the PaI composition would result in the following untypable  $\{\{p, q\}, \{r, s\}\}$ -modular session

$$p[r?\lambda.q!\lambda] \parallel q[p?\lambda.s!\lambda] \parallel r[s?\lambda.p!\lambda] \parallel s[q?\lambda.r!\lambda]$$

where all the processes begin with an input, so forming a deadlock. This problem was overcome in [13], in a setting using projections and without mixed choice, by means of a suitable extension of the syntax of global types.

It is worth noticing that mixed choice make PaI composition problematic even by allowing one connector only. For example, by composing using  $p$  and  $t$  the following typable sessions

$$\begin{aligned} & p[q!\lambda_1 + r?\lambda_2] \parallel q[p?\lambda_1 + s?\lambda_3] \parallel r[p!\lambda_2 + s!\lambda_4] \parallel s[q!\lambda_3 + r?\lambda_4] \\ & t[u?\lambda_1 + v!\lambda_2] \parallel u[t!\lambda_1 + w!\lambda_3] \parallel v[t?\lambda_2 + w?\lambda_4] \parallel w[u?\lambda_3 + v!\lambda_4] \end{aligned}$$

we get the session

$$\begin{aligned} & p[t?\lambda_1.q!\lambda_1 + r?\lambda_2.t!\lambda_2] \parallel q[p?\lambda_1 + s?\lambda_3] \parallel r[p!\lambda_2 + s!\lambda_4] \parallel s[q!\lambda_3 + r?\lambda_4] \parallel \\ & t[u?\lambda_1.p!\lambda_1 + p?\lambda_2.v!\lambda_2] \parallel u[t!\lambda_1 + w!\lambda_3] \parallel v[t?\lambda_2 + w?\lambda_4] \parallel w[u?\lambda_3 + v!\lambda_4] \end{aligned}$$

which reduces to the stuck session

$$p[t!\lambda_2] \parallel t[p!\lambda_1]$$

PaI composition provides also an intuitive justification for the shape of connectors in our modularisation. In fact, by composing systems via interfaces with unrestricted mixed choice, most of the communication properties, if any, are not preserved, as shown by the previous example.

As future work we plan to investigate PaI composition for sessions with mixed choice and asyn-

chronous communication, taking inspiration from [28, 29], where asynchronous communication for sessions with mixed choice is first modelled. We deem worth investigating the modular approach to session types also for standard MPST, as well as for the recent approaches to global types and projections devised in [24, 23].

**Acknowledgments** We wish to gratefully thank the anonymous reviewers for their thoughtful and helpful comments.

## References

- [1] Franco Barbanera, Viviana Bono & Mariangiola Dezani-Ciancaglini (2025): *Open compliance in multiparty sessions with partial typing*. *Journal of Logical and Algebraic Methods in Programming* 144, p. 101046, doi:10.1016/j.jlamp.2025.101046.
- [2] Franco Barbanera, Viviana Bono & Mariangiola Dezani-Ciancaglini (2025): *Partially typed multiparty sessions with internal delegation*. *Journal of Logical and Algebraic Methods in Programming* 142, p. 101018, doi:10.1016/J.JLAMP.2024.101018.
- [3] Franco Barbanera & Mariangiola Dezani-Ciancaglini (2023): *Partially typed multiparty sessions*. In Clément Aubert, Cinzia Di Giusto, Simon Fowler & Larisa Safina, editors: *ICE, EPTCS* 383, Open Publishing Association, pp. 15–34, doi:10.4204/EPTCS.383.2.
- [4] Franco Barbanera, Mariangiola Dezani-Ciancaglini & Ugo de’Liguoro (2022): *Open compliance in multiparty sessions*. In S. Lizeth Tapia Tarifa & José Proença, editors: *FACS, LNCS* 13712, Springer, pp. 222–243, doi:10.1007/978-3-031-20872-0\_13.
- [5] Franco Barbanera, Mariangiola Dezani-Ciancaglini & Ugo de’Liguoro (2024): *Partial typing for asynchronous multiparty sessions*. In Sandra Alves & Ian Mackie, editors: *DCM, EPTCS* 408, Open Publishing Association, pp. 1–20, doi:10.4204/EPTCS.408.1.
- [6] Franco Barbanera, Mariangiola Dezani-Ciancaglini & Ugo de’Liguoro (2024): *Un-projectable global types for multiparty sessions*. In Alessandro Bruni, Alberto Momigliano, Matteo Pradella & Matteo Rossi, editors: *PPDP*, ACM Press, pp. 15:1–15:13, doi:10.1145/3678232.3678245.
- [7] Franco Barbanera, Mariangiola Dezani-Ciancaglini, Lorenzo Gheri & Nobuko Yoshida (2023): *Multicompatibility for multiparty-session composition*. In Santiago Escobar & Vasco T. Vasconcelos, editors: *PPDP*, ACM Press, pp. 2:1–2:15, doi:10.1145/3610612.3610614.
- [8] Luc Bougé (1988): *On the existence of symmetric algorithms to find leaders in networks of communicating sequential processes*. *Acta Informaticae* 25(2), p. 179–201, doi:10.1007/BF00263584.
- [9] Daniel Brand & Pitro Zafiropulo (1983): *On communicating finite-state machines*. *Journal of ACM* 30(2), pp. 323–342, doi:10.1145/322374.322380.
- [10] Filipe Casal, Andreia Mordido & Vasco T. Vasconcelos (2022): *Mixed sessions*. *Theoretical Computer Science* 897, pp. 23–48, doi:10.1016/J.TCS.2021.08.005.
- [11] Iliaria Castellani, Mariangiola Dezani-Ciancaglini & Paola Giannini (2022): *Asynchronous Sessions with Input Races*. In Marco Carbone & Rumyana Neykova, editors: *PLACES, EPTCS* 356, Open Publishing Association, pp. 12–23, doi:10.4204/EPTCS.356.2.
- [12] Francesco Dagnino, Paola Giannini & Mariangiola Dezani-Ciancaglini (2023): *Deconfined global types for asynchronous sessions*. *Logical Methods in Computer Science* 19(1), pp. 1–41, doi:10.46298/lmcs-19(1:3)2023.
- [13] Lorenzo Gheri & Nobuko Yoshida (2023): *Hybrid multiparty session types: compositionality for protocol specification through endpoint projection*. *PACMPL* 7 (OOPSLA1), pp. 112–142, doi:10.1145/3586031.
- [14] Rob van Glabbeek, Peter Höfner & Ross Horne (2021): *Assuming just enough fairness to make session types complete for lock-freedom*. In Leonid Libkin, editor: *LICS, IEEE*, pp. 1–13, doi:10.1109/LICS52264.2021.9470531.

- [15] Mohamed G. Gouda, Eric G. Manning & Yao-Tin Yu (1984): *On the progress of communication between two machines*. *Information and Control* 63(3), pp. 200–2016, doi:10.1016/S0019-9958(84)80014-5.
- [16] Kohei Honda (1993): *Types for dyadic Interaction*. In Eike Best, editor: *CONCUR, LNCS 715*, Springer, pp. 509–523, doi:10.1007/3-540-57208-2\_35.
- [17] Kohei Honda, Vasco T. Vasconcelos & Makoto Kubo (1998): *Language primitives and type discipline for structured communication-based programming*. In Chris Hankin, editor: *ESOP, LNCS 1381*, Springer, pp. 122–138, doi:10.1007/BFb0053567.
- [18] Kohei Honda, Nobuko Yoshida & Marco Carbone (2008): *Multiparty asynchronous session types*. In George C. Necula & Philip Wadler, editors: *POPL, ACM Press*, pp. 273–284, doi:10.1145/1328897.1328472.
- [19] Kohei Honda, Nobuko Yoshida & Marco Carbone (2016): *Multiparty asynchronous session types*. *Journal of the ACM* 63(1), pp. 9:1–9:67, doi:10.1145/2827695.
- [20] Naoki Kobayashi & Davide Sangiorgi (2010): *A hybrid type system for lock-freedom of mobile processes*. *ACM Transactions on Programming Languages and Systems* 32(5), pp. 16:1–16:49, doi:10.1016/j.jlamp.2024.101018.
- [21] Dexter Kozen & Alexandra Silva (2017): *Practical coinduction*. *Mathematical Structures in Computer Science* 27(7), pp. 1132–1152, doi:10.1017/S0960129515000493.
- [22] Julien Lange, Nicholas Ng, Bernardo Toninho & Nobuko Yoshida (2017): *Fencing off go: liveness and safety for channel-based programming*. In Giuseppe Castagna & Andrew D. Gordon, editors: *POPL, ACM Press*, pp. 748–761, doi:10.1145/3009837.3009847.
- [23] Elaine Li, Felix Stutz, Thomas Wies & Damien Zufferey (2023): *Complete multiparty session type projection with automata*. In Constantin Enea & Akash Lal, editors: *CAV, LNCS 13966*, Springer, pp. 350–373, doi:10.1007/978-3-031-37709-9\_17.
- [24] Rupak Majumdar, Madhavan Mukund, Felix Stutz & Damien Zufferey (2021): *Generalising projection in asynchronous multiparty session types*. In Serge Haddad & Daniele Varacca, editors: *CONCUR, LIPIcs 203*, Schloss Dagstuhl - Leibniz-Zentrum für Informatik, pp. 35:1–35:24, doi:10.4230/LIPIcs.CONCUR.2021.35.
- [25] <https://www.ottia.com/en/post/project-modularization>.
- [26] Luca Padovani (2014): *Deadlock and lock freedom in the linear  $\pi$ -calculus*. In Thomas A. Henzinger & Dale Miller, editors: *CSL-LICS, ACM Press*, pp. 72:1–72:10, doi:10.1145/2603088.2603116.
- [27] Catuscia 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.
- [28] Jonah Pears, Laura Bocchi & Andy King (2023): *Safe Asynchronous Mixed-Choice for Timed Interactions*. In Sung-Shik Jongmans & Antónia Lopes, editors: *COORDINATION, LNCS 13908*, Springer, pp. 214–231, doi:10.1007/978-3-031-35361-1\_12.
- [29] Jonah Pears, Laura Bocchi, Maurizio Murgia & Andy King (2024): *Introducing TOAST: Safe Asynchronous Mixed-Choice For Timed Interactions*. *CoRR* abs/2401.11197, doi:10.48550/ARXIV.2401.11197.
- [30] Kirstin Peters & Nobuko Yoshida (2024): *Mixed choice in session types*. *Information and Computation* 298, p. 105164, doi:10.1016/J.IC.2024.105164.
- [31] Kirstin Peters & Nobuko Yoshida (2024): *Separation and encodability in mixed choice multiparty sessions*. In Pawel Sobocinski, Ugo Dal Lago & Javier Esparza, editors: *LICS, ACM Press*, pp. 62:1–62:15, doi:10.1145/3661814.3662085.
- [32] Alceste Scalas & Nobuko Yoshida (2019): *Less is more: multiparty session types revisited*. *PACMPL* 3 (POPL), pp. 30:1–30:29, doi:10.1145/3290343.
- [33] Herbert A. Simon (1991): *The architecture of complexity*. In: *Facets of Systems Science*, Springer, pp. 457–476, doi:10.1007/978-1-4899-0718-9\_31.

# Reactive Semantics for User Interface Description Languages

Basile Pesin

basile.pesin@enac.fr

Celia Picard

celia.picard@enac.fr

Cyril Allignol

cyril.allignol@enac.fr

Fédération ENAC ISAE-SUPAERO ONERA, Université de Toulouse, France

User Interface Description Languages (UIDLs) are high-level languages that facilitate the development of Human-Machine Interfaces, such as Graphical User Interface (GUI) applications. They usually provide first-class primitives to specify how the program reacts to an external event (user input, network message), and how data flows through the program. Although these domain-specific languages are now widely used to implement safety-critical GUIs, little work has been invested in their formalization and verification.

In this paper, we propose a denotational semantic model for a core reactive UIDL, Smalite, which we argue is expressive enough to encode constructs from more realistic languages. This preliminary work may be used as a stepping stone to produce a formally verified compiler for UIDLs.

## 1 Introduction and Context

With the democratization of interactive devices, the general interaction paradigm has changed. Users now expect to interact with their systems through tactile interactions or advanced interactive devices. This extends to critical systems too, notably aviation related ones. For instance, aircraft cockpits are now digital and tactile [13] and paper strips for air traffic controllers have been replaced by elaborate digital dashboards [10].

Using traditional programming languages to implement the interactive parts of systems implies the use of numerous callbacks. The resulting code quickly becomes spaghetti code and gets particularly difficult to analyze and maintain [16, 17]. The use of dedicated languages has shown to be beneficial in such circumstances [18]. Such languages are called User Interface Description Languages (UIDLs). These languages can efficiently describe both the appearance of the system (its scene graph) and the interactive behavior (mainly activation propagation). They are becoming more popular, including for the development of critical systems [2, 1]. However, these languages have not been formally defined, particularly not their semantics, which hinders any formal reasoning on these safety-critical programs. In this paper, we tackle this problem by proposing denotational semantics for a minimal declarative UIDL.

There has been a few other attempts to give a formal semantics to interactive languages or libraries such as React [14]. However, React is based on a functional language and as such, has very different mechanisms than other UIDL/interactive languages. Indeed, the behavior of a React program is described using small-step operational semantics which specify the order in which components are rendered. In this paper, we focus on a declarative language with denotational semantics where the order of updates is left implicit.

Previous work has paved the way towards a minimal common abstract syntax for Smala, a declarative UIDL [15]. A first formal semantic model based on bigraphs has been proposed [20], but in this work we explore an alternative expression of the semantics, which is simpler and should facilitate compiler


verification in the fashion of CompCert [12] or Vélus [4]. To give more confidence in those semantics, we also aim at proving the equivalence between the two semantics, as future work.

Hence, we propose a new denotational semantics for Smalite, a declarative UIDL which includes a minimal set of constructs from Smala. As a preliminary step to a compiler correctness proof, we mechanized the language and its semantics in the Rocq prover<sup>1</sup> [23], and implemented a prototype compiler that generates imperative code. In section 2, we give an informal presentation of our language through an example program implementing a simple Graphical User Interface (GUI) program. Then, in section 3, we describe our formalization of the reactive semantics of the language. Finally, in section 4, we discuss the future steps towards extending the language and building a formal proof of correctness for its compiler.

## 2 Specifying interactions with Smala

The language we propose closely resembles Smala [15], a UIDL used in safety-critical applications [2, 7]. In fig. 1, we present an example Smala program where two buttons control a counter by either decrementing it until it reaches 0 or setting it back to 3.

Smala programs are composed of named *processes*. The root process is a **Component** (line 1), usually named *root*, which contains child processes. The counter is implemented as a *property* *count* (line 2), which is declared with a type and an expression giving its initial value. Then, the program contains a **Spike** (line 3), which represents an event that may be triggered from inside or outside the program, and reacted to. This particular **Spike** represents the event “the counter has just reached 0”. Indeed, it is triggered by the *binding* on the same line, whose left-hand side is a condition checking that *count* equals 0. Conditions are checked only at reactions where the value of one of the properties involved changes; therefore, *zero* is triggered only at reactions where the value of *count* changes to 0.

The next process is a **Frame**, which is a special *graphical* component (line 6). For the purpose of our reactive semantics, a **Frame** is just syntactic sugar for a component with properties *title*, *width* and *height*, passed as parameters between parentheses, as well as a **Spike** *close*. On the other hand, our prototype compiler generates code which uses the **Frame** and its parameters to open a window. The compiler also binds user requests to close the window (e.g. by clicks on the  button) to triggering the *close Spike*. Therefore, the program can define custom reactions to user actions.

The first child of the **Frame** defines the **Font** used to display text in the interface (line 7). Since the rest of the program does not need to refer to this particular component, its name is unspecified (*\_*). Then, the frame contains three main components: two buttons, and a text label. Each button is implemented by a component **FillColor** (with parameters *red*, *green* and *blue*), which sets the color of its **Rectangle** child (with parameters *x*, *y*, *width* and *height*). Text labels are implemented by a **FillColor** which sets the color of its **Text** child (with parameters *text*, *x* and *y*).

**Rectangle** components have two predefined spikes, *press* and *release*, which are respectively triggered when the mouse is pressed/released while the mouse pointer is on top of the rectangle. The logic of each button is implemented as a reaction to these spikes: on *press*, the button is highlighted by increasing the *green* property of its **FillColor** (lines 14, 26). On *release*, the *green* component is set back to its default value (lines 15, 27). The graphical effect of these first two reactions are shown at right of fig. 1.

The *release Spike* of each button is bound to an action on the counter. Releasing *decr* decreases the counter (line 17) by setting its new value to its previous value (accessed with *last*) minus one.

---

<sup>1</sup>formerly known as Coq

```

1 Component root {
2   Int count 3;
3
4   Spike zero; (count == 0) -> zero;
5
6   Frame f ("ICE_2025", 300, 50) {
7     Font _ ("arial.ttf", 20) {
8       FillColor btn1 (150,150,150) {
9         Rectangle r (0,0,100,f.height) {
10          FillColor _ (0,0,0) {
11            Text _ ("decr", r.x + 30, 13) {}
12          };
13        };
14        r.press -> hg; hg: 255 =: green;
15        r.release -> dhg; dhg: 150 =: green;
16      };
17      btn1.r.release -> dec; dec: last count-1 =: count;
18      zero ->! btn1.r; (count > 0) -> btn1.r;
19
20      FillColor btn2 (150,150,150) {
21        Rectangle<d> r (btn1.r.width+10,0,100,f.height) {
22          FillColor _ (0,0,0) {
23            Text t ("restart", r.x + 20, 13) {}
24          };
25        };
26        r.press -> hg; hg: 255 =: green;
27        r.release -> dhg; dhg: 150 =: green;
28      };
29      btn2.r.release -> rst; rst: 3 =: count;
30      (count < 3) -> btn2.r; (count == 3) ->! btn2.r;
31
32      FillColor _ (255,255,255) {
33        Text t ("rem:␣" + str(count), btn2.r.x+btn2.r.width+10, 13) {
34          count -> chtext; chtext: "rem:␣" + str(count) =: text;
35        };
36      }
37    }
38  };
39
40  Exit e (0) { f.close -> trigger };
41 }

```

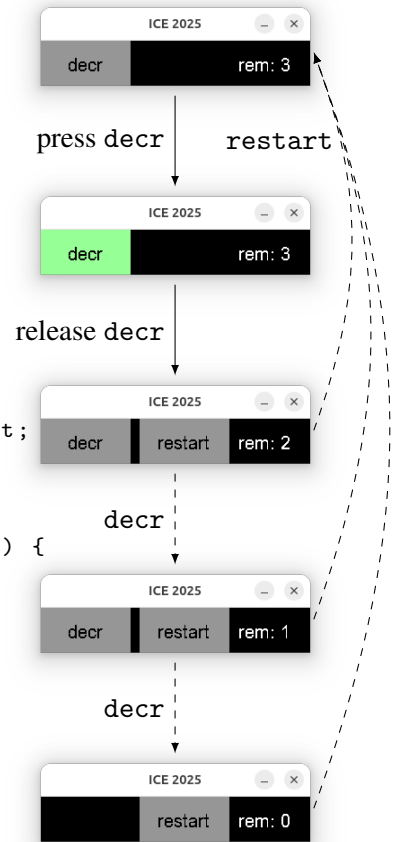


Figure 1: An example Smala program

Releasing restart sets the counter back to 3 (line 29). Buttons may be activated or deactivated depending on the value of the counter. If the counter equals 0, it is not possible anymore to decrease it, and therefore the `decr` button is deactivated (first binding `->! btn1.r` on line 18). As soon as the counter goes back above 0, it is reactivated. Conversely, the `restart` button is deactivated when the counter equals 3, and reactivated otherwise (line 30). Normally, the activation of a parent automatically activates all of its children. However, since the counter starts at 3, the `restart` button must be initially deactivated, which is specified by `<d>` on line 21.

The last child component of the `Frame` is a `Text` label displaying the value of `count`. It is updated every time the value of `count` changes (binding `count ->` on line 34).

$$\begin{array}{l}
p ::= ty \ x \ e \\
| \text{Spike } x \\
| x: e =: path \\
| x: lhs \rightarrow \langle ia \rangle rhs \\
| \text{Component} \langle ia \rangle x \{ p^* \} \\
\end{array}
\qquad
\begin{array}{l}
path ::= x \mid x.path \\
e ::= c \mid path \mid \text{last } path \mid \diamond e \mid e \oplus e \\
ia ::= a \mid d \\
lhs ::= T?(path) \mid A?(path) \mid D?(path) \mid C?(path) \mid (e)? \\
rhs ::= T!(path) \mid A!(path) \mid D!(path)
\end{array}$$

Figure 2: Abstract Syntax of Smalite

Concrete LHS	Concrete RHS	Abstract LHS	Abstract RHS	Event
path ->	-> path	T?(path)	T!(path)	Trigger path
path ->	N/A	C?(path)	N/A	Assign v path
(e) ->	N/A	(e)?	N/A	Assign v path, path ∈ free(e)
path ->	-> path	A?(path)	A!(path)	Activate path
path !->	->! path	D?(path)	D!(path)	Deactivate path

Figure 3: Correspondence between concrete syntax, abstract syntax and events

Finally, the program contains an `Exit` component, which takes an exit code as a parameter, and exposes a `Spike` trigger which halts the program. The `close Spike` of the `Frame` is bound to `trigger`, which allows the user to close the program by closing the window.

### 3 Formalizing Smalite’s semantics

In this paper, we define the formal semantics of Smalite, a more restricted form of Smala. Smalite is a reactive language: the execution of a program can be seen as a series of reactions to external events. We saw in the example of fig. 1 that the main type of external event is the triggering of a `Spike`, such as `close` (line 40) or `release` (line 15). Another possible external event not showcased in the example would be an outside modification of a property. For instance, an adaptive GUI program could listen for the resizing of a `Frame` which would be encoded as a change of its width or height property.

A reaction may then update the state of the process by *activating* or *deactivating* permanent processes, or *assigning* values to properties. A reaction may also *trigger* spikes which, as seen in the example, may have an effect on the view. Therefore, we generalize external events to also account for events triggered by the program during a reaction, as defined below.

$$ev ::= \text{Trigger } path \mid \text{Assign } v \ path \mid \text{Activate } path \mid \text{Deactivate } path$$

First, an event can be the triggering of a “transient” process, which does not retain its activation between cycles (`Spike` or `assign`). Second, it can be an assignment to a property. Last, it can be the activation or deactivation of a “permanent” process, which retains its activation between cycles (`Component` or `binding`).

#### 3.1 Abstract syntax of Smalite

We now detail the abstract syntax of Smalite as we formalize it. A minimalized version of Smala has been proposed in previous work [15]. It includes a small set of core elements that constitutes the heart

of the language, expressive enough to describe the full set of components of the whole language while minimizing the number of constructs to actually formally define and verify. The translation from Smala to this minimal set is done through a transpilation pass not described here. Our syntax is presented in fig. 2, and builds upon the one proposed previously [15].

A program is a process  $p$ . There are five different processes detailed hereafter. A process may be the declaration of a property with its type  $ty$ , name  $x$  and initial value  $e$ . There are two kinds of transient processes, which are triggered during a single cycle: spikes and assignments. The latter assigns the result of the evaluation of an expression  $e$  to a property specified by an absolute *path*. An expression is either a constant  $c$ , the *path* to access the value or the `last` value of a property, or a unary ( $\diamond$ ) or binary ( $\oplus$ ) operation. Types, constants and operations are those of the host language that the compiler targets (in the case of our prototype compiler, CompCert C).

Finally, there are two kinds of permanent processes, which retain their activation between cycles. Both kinds may be initially activated or deactivated following the *ia* flag. The first permanent process is the binding, denoted by an arrow  $\rightarrow$ . It binds the event specified by its left-hand-side to the event specified by its right-hand-side. An event detected by a left-hand-side may be the triggering (T?) of a transient process (`Spike`, `assign`), the activation (A?) or deactivation (D?) of a permanent process (binding, `Component`), the change of value of a property (C?), or a condition being true after one of its free variable changes ( $(e)?$ ). Right-hand-sides may trigger a transient process (T!) or activate (A!) or deactivate (D!) a permanent process. The second permanent process is the component, which contains a list of sub-processes.

This abstract syntax is slightly different from the concrete syntax used for the example. First, each process is explicitly named. Second, there is only one generic `Component`. It can then be instantiated to create all other components, particularly the graphical components (`Frame`, `Rectangle`, etc), as those do not play a special role in the semantics. Third, as we have seen, the left- and right-hand sides of bindings distinguish between triggering of transient processes, activation/deactivation of permanent processes, and change of property value. Figure 3 recapitulates the correspondence between concrete left- and right-hand-sides, as seen in the example, and their abstract counterpart. Last, processes refer to each other using their absolute paths, which are made up of a sequence of identifiers indicating the path from the root process to the process of interest.

In practice, our prototype compiler parses the source program of fig. 1 into an Abstract Syntax Tree (AST) represented as an inductive type in Rocq. Then, a sequence of functions recursively elaborates it into a second AST representing the more restricted syntax of fig. 2, in three passes. First, the elaborator fills-in missing names by generating globally unique identifiers. Then, it expands graphical components into generic `Component` by adding the predefined properties and spikes (e.g. `title`, `width`, `frame` for a `Frame`). Finally, it type-checks the program, makes relative paths absolute, differentiates transient and permanent processes in bindings, and builds the restricted AST.

## 3.2 Semantics of program initialization

We now describe the rules that specify the initialization of the system, presented in fig. 4. The first judgment,  $\vdash_{\text{init}}^{\text{activ}} p_r \Downarrow A$ , specifies the initial activation for a process  $p$  rooted at path  $r$ , where  $A$  is a set of paths of permanent processes. It is defined by the rules in fig. 4a. Non-permanent processes (property, spike, assignment) never appear in this set. Bindings and components appear in the set if-and-only-if they are declared with  $\langle a \rangle$ . For a component marked with  $\langle a \rangle$ , the sub-processes are initially activated according to the same rules.

The second judgment,  $E \vdash_{\text{init}}^{\text{env}} p_r \Downarrow E'$ , specifies how the initial property values are added to initial

$$\begin{array}{c}
\boxed{\vdash_{\text{init}}^{\text{activ}} p_r \Downarrow A} \\
\frac{\vdash_{\text{init}}^{\text{activ}} (ty \ x \ e)_{r_x} \Downarrow \emptyset}{\text{IAP}} \quad \frac{\vdash_{\text{init}}^{\text{activ}} (\text{Spike } x)_{r_x} \Downarrow \emptyset}{\text{IAS}} \quad \frac{\vdash_{\text{init}}^{\text{activ}} (x: e =: p)_{r_x} \Downarrow \emptyset}{\text{IAA}} \\
\frac{\vdash_{\text{init}}^{\text{activ}} (x: lhs \rightarrow \langle d \rangle rhs)_{r_x} \Downarrow \emptyset}{\text{IAB}_1} \quad \frac{\vdash_{\text{init}}^{\text{activ}} (x: lhs \rightarrow \langle a \rangle rhs)_{r_x} \Downarrow \{r_x.x\}}{\text{IAB}_2} \\
\frac{\vdash_{\text{init}}^{\text{activ}} (\text{Component} \langle d \rangle x \{ps\})_{r_x} \Downarrow \emptyset}{\text{IAC}_1} \quad \frac{\forall i. \vdash_{\text{init}}^{\text{activ}} (ps_i)_{r_x.x} \Downarrow A_i}{\vdash_{\text{init}}^{\text{activ}} (\text{Component} \langle a \rangle x \{ps\})_{r_x} \Downarrow (\bigcup_i A_i) \cup \{r_x.x\}} \text{IAC}_2
\end{array}$$

(a) Initialization of activation

$$\begin{array}{c}
\boxed{E \vdash_{\text{init}}^{\text{env}} p_r \Downarrow E'} \\
\frac{E \vdash_{\text{init}}^{\text{env}} (\text{Spike } x)_{r_x} \Downarrow E}{\text{IES}} \quad \frac{E \vdash_{\text{init}}^{\text{env}} (x: e =: p)_{r_x} \Downarrow E}{\text{IEA}} \quad \frac{E \vdash_{\text{init}}^{\text{env}} (x: lhs \rightarrow \langle \_ \rangle rhs)_{r_x} \Downarrow E}{\text{IEB}} \\
\frac{\emptyset, E \vdash_{\text{exp}} e \Downarrow v}{E \vdash_{\text{init}}^{\text{env}} (ty \ x \ e)_{r_x} \Downarrow E[r_x.x \mapsto v]} \text{IEP} \quad \frac{}{E \vdash_{\text{init}}^{\text{env}} (\text{Component} \langle \_ \rangle x \{\varepsilon\})_{r_x} \Downarrow E} \text{IEC}_1 \\
\frac{E \vdash_{\text{init}}^{\text{env}} p_{(r_x.x)} \Downarrow E' \quad E' \vdash_{\text{init}}^{\text{env}} (\text{Component} \langle \_ \rangle x \{ps\})_{r_x} \Downarrow E''}{E \vdash_{\text{init}}^{\text{env}} (\text{Component} \langle \_ \rangle x \{p; ps\})_{r_x} \Downarrow E''} \text{IEC}_2
\end{array}$$

(b) Initialization of environment

$$\begin{array}{c}
\boxed{\vdash_{\text{init}} p \Downarrow E, A} \\
\frac{\vdash_{\text{init}}^{\text{activ}} p_\varepsilon \Downarrow A \quad \emptyset \vdash_{\text{init}}^{\text{env}} p_\varepsilon \Downarrow E}{\vdash_{\text{init}} p \Downarrow E, A} \text{I}
\end{array}$$

(c) Program initialization

Figure 4: Initialization rules

environment  $E$  to form  $E'$ , where environments are represented by finite maps from property paths to values. It is defined by the rules in fig. 4b. As expected, spikes, assignments and bindings have no effect on the initial environment. For properties (rule IEP), the initialization expression is evaluated under the starting environment  $E$  and the resulting value is added to  $E$ . A value is either a numerical value from the C language (integer, floating-point number), a boolean value, or a string of characters. The rules for expression evaluation specifying  $E, E' \vdash_{\text{exp}} e \Downarrow v$  are unsurprising, and presented in fig. 5. The notation  $c^\#$  represents the semantic denotation of constant  $c$  (respectively  $\diamond^\#$  for unary operator  $\diamond$  and  $\oplus^\#$  for binary operator  $\oplus$ ). The current value of a variable is searched in  $E'$ , the updated environment (rule EV), while its **last** value is searched in  $E$ , which represents the environment at the end of the last step, and is empty at initialization (rule EL). The evaluation of operators may fail (e.g. in case of a division by zero), which is denoted on their return value by  $\lfloor v \rfloor$ .

Rules IEC<sub>1</sub> and IEC<sub>2</sub>, for components, highlight an interesting design choice. In our definitions, the

$$\begin{array}{c}
\boxed{E, E' \vdash_{\text{exp}} e \Downarrow v} \\
\frac{}{E, E' \vdash_{\text{exp}} c \Downarrow c^\#} \text{EC} \quad \frac{}{E, E' \vdash_{\text{exp}} x \Downarrow E'(x)} \text{EV} \quad \frac{}{E, E' \vdash_{\text{exp}} \text{last } x \Downarrow E(x)} \text{EL} \\
\frac{E, E' \vdash_{\text{exp}} e_1 \Downarrow v_1 \quad \diamond^\#(v_1) = \lfloor v \rfloor}{E, E' \vdash_{\text{exp}} \diamond e_1 \Downarrow v} \text{EU} \quad \frac{E, E' \vdash_{\text{exp}} e_1 \Downarrow v_1 \quad E, E' \vdash_{\text{exp}} e_2 \Downarrow v_2 \quad \oplus^\#(v_1, v_2) = \lfloor v \rfloor}{E, E' \vdash_{\text{exp}} e_1 \oplus e_2 \Downarrow v} \text{EB}
\end{array}$$

Figure 5: Expression evaluation rules

$$\begin{array}{c}
\boxed{E, A \vdash_{\text{react}} p(ev) \Downarrow E', A', T} \\
\frac{T = \{ev_0\} \cup \{ev \mid E, A, T, E', A' \vdash_{\text{prop}}^{\text{proc}} p \Downarrow ev\} \quad E, T, E', A' \vdash_{\text{safe}}^{\text{proc}} p_\varepsilon \quad E, A \vdash_{\text{update}} T \Downarrow E', A'}{E, A \vdash_{\text{react}} p(ev_0) \Downarrow E', A', (T \cap \{\text{Trigger path} \mid \text{Spike path} \in p\})} \text{R}
\end{array}$$

Figure 6: Main reaction rule

environment is updated sequentially, following the order the sub-processes are written in the program. One simpler definition would have been to update the same environment  $E$  in parallel for each sub-process  $ps_i$  into an environment  $E_i$ , and to take the union of these environments, but this would have prevented writing initial expressions that depend on the values of other properties, as seen for the button text position in fig. 1 (line 21). Conversely, a more expressive semantics would allow the initialization of a property to depend on the values of properties later in the program, as long as there is no cycle in the definitions. This would however require (1) a more complicated semantic judgment for initialization, (2) an analysis of the absence of cycle at the source level, and (3) a way to schedule property initialization in the compiled code. Although these are all feasible, we believe that the marginal expressivity gains are not sufficient to justify this additional complexity.

Finally, the two judgments described above are combined in rule I, presented in fig. 4c. It specifies the initialization of a Smalite program:  $\vdash_{\text{init}} p \Downarrow E, A$  indicates that the process  $p$ , rooted at the empty path  $\varepsilon$ , is initialized with initial environment  $E$  and initial activation  $A$ .

### 3.3 Semantics of reaction

To describe how a program reacts to an external event, let us start with the main reaction rule presented in fig. 6. The judgment  $E, A \vdash_{\text{react}} p(ev) \Downarrow E', A', T$  indicates that, with initial environment  $E$  and activation  $A$ , the process  $p$  reacts to an external event  $ev$  by updating its environment to  $E'$ , its activation to  $A'$ , and emitting a set  $T$  of external events. The relation between these parameters are specified by three premises, which we now detail.

#### 3.3.1 Event propagation

The first premise determines the set  $T$  of events that are emitted during the reaction step.  $T$  contains the external event that triggered the reaction,  $ev$ . It also contains any additional event propagated by the process: this is represented by judgment  $E, A, T, E', A' \vdash_{\text{prop}}^{\text{proc}} p \Downarrow ev$  which can be read as “when process  $p$

$$\begin{array}{c}
\boxed{E, T, E' \vdash_{\text{prop}}^{\text{lhs}} lhs} \\
\frac{\text{Trigger } x \in T}{E, T, E' \vdash_{\text{prop}}^{\text{lhs}} T?(x)} \text{PLT} \quad \frac{\text{Activate } x \in T}{E, T, E' \vdash_{\text{prop}}^{\text{lhs}} A?(x)} \text{PLA} \quad \frac{\text{Deactivate } x \in T}{E, T, E' \vdash_{\text{prop}}^{\text{lhs}} D?(x)} \text{PLD} \quad \frac{\text{Assign } v \ x \in T}{E, T, E' \vdash_{\text{prop}}^{\text{lhs}} C?(x)} \text{PLC} \\
\frac{x \in \text{free}(e) \quad \text{Assign } v \ x \in T \quad E, E' \vdash_{\text{exp}} e \Downarrow \text{true}}{E, T, E' \vdash_{\text{prop}}^{\text{lhs}} (e)?} \text{PLI} \\
\boxed{A, A' \vdash_{\text{prop}}^{\text{rhs}} rhs \Downarrow ev} \\
\frac{r_x \in A'}{A, A' \vdash_{\text{prop}}^{\text{rhs}} T!(r_x.x) \Downarrow \text{Trigger } r_x.x} \text{PRT} \quad \frac{r_x \in A' \quad r_x.x \notin A}{A, A' \vdash_{\text{prop}}^{\text{rhs}} A!(r_x.x) \Downarrow \text{Activate } r_x.x} \text{PRA} \\
\frac{r_x \in A' \quad r_x.x \in A}{A, A' \vdash_{\text{prop}}^{\text{rhs}} D!(r_x.x) \Downarrow \text{Deactivate } r_x.x} \text{PRD} \\
\boxed{E, A, T, E', A' \vdash_{\text{prop}}^{\text{proc}} p \Downarrow ev} \\
\frac{p[r_x.x] = [x: lhs \rightarrow \_ < \_ rhs] \quad r_x.x \in A' \quad E, T, E' \vdash_{\text{prop}}^{\text{lhs}} lhs \quad A, A' \vdash_{\text{prop}}^{\text{rhs}} rhs \Downarrow ev}{E, A, T, E', A' \vdash_{\text{prop}}^{\text{proc}} p \Downarrow ev} \text{PB} \\
\frac{p[r_x.x] = [x: e =: r_y.y] \quad r_x \in A' \quad r_y \in A' \quad \text{Trigger } r_x.x \in T \quad E, E' \vdash_{\text{exp}} e \Downarrow v}{E, A, T, E', A' \vdash_{\text{prop}}^{\text{proc}} p \Downarrow \text{Assign } v \ r_y.y} \text{PA}
\end{array}$$

Figure 7: Propagation rule for bindings and assignments

reacts to events  $T$ , it also emits event  $ev$ ". The rules defining this judgment are presented in figs. 7 and 8. These rules are not syntax-directed: instead, each one asserts the existence of a sub-process that may propagate events. We write the judgement  $p[\text{path}] = [p']$  for "there is a process  $p'$  rooted at  $\text{path}$  in  $p$ ".

The first set of rules, in fig. 7, describes how bindings propagate events. Rule PB, presented at the bottom, specifies that a binding only propagates events if it is activated. The third premise  $E, T, E' \vdash_{\text{prop}}^{\text{lhs}} lhs$  of the rule assumes that the left-hand side is activated by a matching event. In particular, if the left-hand side is a condition, it must evaluate to **true**. The final premise  $A, A' \vdash_{\text{prop}}^{\text{rhs}} rhs \Downarrow ev$  specifies which event  $ev$  the right-hand side emits. All rules specifying this judgment mandate that the parent of the process involved in the event is active. Activation (resp. deactivation) events are only emitted if the process was previously deactivated (resp. activated): in other words, "non-events" that do not modify the activation state are not propagated.

The last rule in fig. 7, covers the assignment. It states that an assignment produces an **Assign**  $v \ y$  event when (1) its parent is activated, (2) the parent of  $y$  is activated, (3) the assignment has been triggered by another event, and (4) the expression of the assignment evaluates to value  $v$ . Here, we made one interesting choice: the second premise implies that it is not possible to assign to a property whose parent is deactivated. This constraint was added because the behavior of programs where a property is assigned to when its parent is deactivated was not clear: should the assignment event be propagated during the

$$\begin{array}{c}
\boxed{\vdash_{\text{prop}}^{\text{activ}} p \Downarrow ev} \\
\frac{}{\vdash_{\text{prop}}^{\text{activ}} (x: lhs \rightarrow \langle a \rangle rhs)_{r_x} \Downarrow \text{Activate } r_x} \text{PAB} \\
\frac{}{\vdash_{\text{prop}}^{\text{activ}} (\text{Component} \langle a \rangle x \{ps\})_{r_x} \Downarrow \text{Activate } r_x.x} \text{PAC}_1 \quad \frac{\vdash_{\text{prop}}^{\text{activ}} (ps_i)_{r_x.x} \Downarrow ev}{\vdash_{\text{prop}}^{\text{activ}} (\text{Component} \langle a \rangle x \{ps\})_{r_x} \Downarrow ev} \text{PAC}_2 \\
\boxed{\vdash_{\text{prop}}^{\text{deactiv}} p \Downarrow ev} \\
\frac{}{\vdash_{\text{prop}}^{\text{deactiv}} (x: lhs \rightarrow \langle \_ \rangle rhs)_{r_x} \Downarrow \text{Deactivate } r_x.x} \text{PDB} \\
\frac{}{\vdash_{\text{prop}}^{\text{deactiv}} (\text{Component} \langle \_ \rangle x \{ps\})_{r_x} \Downarrow \text{Deactivate } r_x.x} \text{PDC}_1 \quad \frac{\vdash_{\text{prop}}^{\text{deactiv}} (ps_i)_{r_x.x} \Downarrow ev}{\vdash_{\text{prop}}^{\text{deactiv}} (\text{Component} \langle \_ \rangle x \{ps\})_{r_x} \Downarrow ev} \text{PDC}_2 \\
\boxed{E, A, T, E', A' \vdash_{\text{prop}}^{\text{proc}} p \Downarrow ev} \\
\frac{p[r_x.x] = [\text{Component} \langle a \rangle x \{ps\}] \quad \text{Activate } r_x.x \in T \quad \vdash_{\text{prop}}^{\text{activ}} (ps_i)_{r_x.x} \Downarrow \text{Activate } y}{E, A, T, E', A' \vdash_{\text{prop}}^{\text{proc}} p \Downarrow \text{Activate } y} \text{PC}_1 \\
\frac{p[r_x.x] = [\text{Component} \langle \_ \rangle x \{ps\}] \quad \text{Deactivate } r_x.x \in T \quad \vdash_{\text{prop}}^{\text{deactiv}} (ps_i)_{r_x.x} \Downarrow \text{Deactivate } y}{E, A, T, E', A' \vdash_{\text{prop}}^{\text{proc}} p \Downarrow \text{Deactivate } y} \text{PC}_2
\end{array}$$

Figure 8: Propagation rules for components

reaction, or when the parent becomes active again, or not at all? We chose to eliminate this question by forbidding this case entirely, as we believe it is not useful in real-world programs. In practice, our prototype compiler statically checks that this situation can never arise.

The last set of rules, in fig. 8, describes how components propagate events to their children. Rule PC<sub>1</sub> specifies the propagation of activation events, while rule PC<sub>2</sub> specifies the propagation of deactivation events. Each is specified by a judgment  $\vdash_{\text{prop}}^{(\text{de})\text{activ}} p \Downarrow ev$ , which follows the same logic as the judgment for initial activation. The only difference between these two sets of rules is that activation is only propagated to children for components/bindings marked with  $\langle a \rangle$ , while deactivation is always propagated.

### 3.3.2 Safe reactions

The propagation rule by itself is not enough to ensure the right semantics is given to propagation of events. Indeed, consider the pair of processes  $a: 0 = y; (x/y > 10) \rightarrow t$  and suppose that assignment  $a$  is triggered. According to our semantics, it generates an event `Assign 0 y`. In turn, this triggers the conditional binding, as the value of  $y$  has changed. However, the behavior of  $x/y$  is undefined because it divides by 0. This means that the predicate for expression evaluation does not apply, and therefore the premises of the rule for conditional left-hand side activation does not hold. In the end, our semantics say that this program evaluates fine, but does not trigger  $t$ . This is wrong: actually, if we compile and/or execute this program, it will crash. Therefore, this program should not admit a semantics at all. To

$$\begin{array}{c}
\boxed{E, T, E' \vdash_{\text{safe}}^{\text{lhs}} lhs} \\
\\
\frac{}{E, T, E' \vdash_{\text{safe}}^{\text{lhs}} x} \text{SLT} \quad \frac{}{E, T, E' \vdash_{\text{safe}}^{\text{lhs}} A?(x)} \text{SLA} \quad \frac{}{E, T, E' \vdash_{\text{safe}}^{\text{lhs}} D?(x)} \text{SLD} \quad \frac{}{E, T, E' \vdash_{\text{safe}}^{\text{lhs}} C?(x)} \text{SLC} \\
\\
\frac{\forall x \in \text{free}(e), \forall v, \text{Assign } v \ x \notin T}{E, T, E' \vdash_{\text{safe}}^{\text{lhs}} (e)?} \text{SLI}_1 \quad \frac{E, E' \vdash_{\text{exp}} e \Downarrow v \quad v \in \{\text{true}, \text{false}\}}{E, T, E' \vdash_{\text{safe}}^{\text{lhs}} (e)?} \text{SLI}_2 \\
\\
\boxed{E, T, E', A' \vdash_{\text{safe}}^{\text{proc}} p_r} \\
\\
\frac{}{E, T, E', A' \vdash_{\text{safe}}^{\text{proc}} (ty \ x \ e)_{r_x}} \text{SP} \quad \frac{}{E, T, E', A' \vdash_{\text{safe}}^{\text{proc}} (\text{Spike } x)_{r_x}} \text{SS} \quad \frac{\text{Trigger } r_x \cdot x \in T \implies E, E' \vdash_{\text{exp}} e \Downarrow v}{E, T, E', A' \vdash_{\text{safe}}^{\text{proc}} (x: e =: y)_{r_x}} \text{SA} \\
\\
\frac{r_x \cdot x \notin A'}{E, T, E', A' \vdash_{\text{safe}}^{\text{proc}} (x: lhs \rightarrow \_ > rhs)_{r_x}} \text{SB}_1 \quad \frac{E, T, E' \vdash_{\text{safe}}^{\text{lhs}} lhs}{E, T, E', A' \vdash_{\text{safe}}^{\text{proc}} x: lhs \rightarrow \_ > rhs} \text{SB}_2 \\
\\
\frac{r_x \cdot x \notin A'}{E, T, E', A' \vdash_{\text{safe}}^{\text{proc}} (\text{Component} \_ < \_ > x \{ps\})_{r_x}} \text{SC}_1 \quad \frac{\forall i, E, T, E', A' \vdash_{\text{safe}}^{\text{proc}} ps_{i_{r_x \cdot x}}}{E, T, E', A' \vdash_{\text{safe}}^{\text{proc}} (\text{Component} \_ < \_ > x \{ps\})_{r_x}} \text{SC}_2
\end{array}$$

Figure 9: Safe reaction rules

correct for these cases, we introduce an additional premise in fig. 6,  $E, T, E', A' \vdash_{\text{safe}}^{\text{proc}} p_r$ , which ensures that all expressions that need to be evaluated in the process  $p$  rooted at  $r$  evaluate without undefined behavior.

The definition of this judgment is presented in fig. 9. Execution of a process that does not propagate events (property, **Spike**) is always safe. An assignment is safe if, when triggered, the evaluation of its expression is safe (rule SA). Both bindings and components are safe if they are inactive (rules SB<sub>1</sub> and SC<sub>1</sub>). If an assignment is active, then its left-hand side must be safe: if the left-hand side is a condition, either there is no assignment to its free variables, in which case it does not need to be evaluated (rule SLI<sub>1</sub>), or it evaluates to a boolean value (rule SLI<sub>2</sub>). An active component is safe if all its children processes are safe (rule SC<sub>2</sub>).

### 3.3.3 State update

The two judgments described above specify what events are emitted during the reaction. It remains to specify how the state of the reactive system is updated by these events. This is the role of the judgment  $E, A \vdash_{\text{update}} T \Downarrow E', A'$ , which specifies that “when it receives events  $T$ , the state of a process updates from  $(E, A)$  to  $(E', A')$ ”. The unique rule specifying this judgment is presented in fig. 10.

Its first two premises specify the updated environment  $E'$ : for each path  $p$ , either there is an **Assign** event to  $p$ , in which case  $E'(x)$  is set to its value, or there is none, in which case  $E'(x)$  keeps the same value as  $E(x)$ . The first premise implies a strong constraint of Smalite programs, that was not explicit until now: two assignments to the same property with different values may not occur during the same reaction. This is a stronger constraint than was chosen in the original implementation of Smala (where the value of the “final” assignment was kept), but it simplifies the semantic model and facilitates the generation of simple and efficient imperative code. In practice, this semantic constraint is satisfied for

$$\boxed{E, A \vdash_{\text{update}} T \Downarrow E', A'}$$

$$\frac{
\begin{array}{l}
\forall p \ v, \text{Assign } v \ p \in T \implies E'(p) = \lfloor v \rfloor \quad \forall p, (\forall v, \text{Assign } v \ p \notin T) \implies E'(p) = E(p) \\
\forall p, \text{Activate } p \in T \implies p \in A' \quad \forall p, \text{Deactivate } p \in T \implies p \notin A' \\
\forall p, (\text{Activate } p \notin T \wedge \text{Deactivate } p \notin T) \implies (p \in A' \iff p \in A)
\end{array}
}{E, A \vdash_{\text{update}} T \Downarrow E', A'} \text{U}$$

Figure 10: State update rule

any program that respects a Reactive Static Single Assignment (RSSA) property: each reaction only contains one assignment to a given property. This property is checked by our prototype compiler.

The remaining premises specify the updated activation  $A'$ : an **Activate**  $p$  event adds  $p$  in  $A'$ , a **Deactivate**  $p$  event removes it. If neither type of events are emitted, then  $p$  keeps its activation status. These premises imply a second constraint: **Activate**  $p$  and **Deactivate**  $p$  events may not occur during the same reaction. Like the previous constraint, this choice simplifies both the semantics and compilation scheme, and is checked statically during compilation.

### 3.3.4 Discussion

Going back to the rule that composes them in fig. 6, we can now get a high-level view of how these three judgments interact to specify how a Smalite process reacts to an event. As discussed, the *prop* and *safe* rules specify which events are produced, while the *update* rule specifies how state is updated. These definitions appear to be mutually recursive. On the one hand, in *prop*, the updated state  $E', A'$  is used as argument to compute generated events. On the other hand, in *update*, the set of events  $T$  is used as an argument to compute the updated state  $E', A'$ . To make these semantics executable, one would most likely need to implement them as a pair of mutually-defined fixed-points.

## 4 Conclusion and Future work

In this paper, we have presented Smalite, a minimal language capable of encoding more fully-featured UIDLs such as Smala. The denotational semantics of the language, along with a prototype compiler, have been implemented in the Rocq prover: this is a first stepping stone towards a verification framework for UIDLs in general and Smala in particular. We propose here some leads for future work.

### 4.1 A formally verified compiler for Smalite

We have implemented a prototype compiler for Smalite programs that generates C code which can be linked against the SDL library [22] to define reactive GUI programs. This compiler generates executable code for Smalite programs by finding a static scheduling of instructions that implement events (setting the value of property, (un)setting activation flag), guarded by conditions that emulate the left-hand side of bindings. This involves three major passes between separate Intermediate Representations (IRs):

1. Explicitly compute the event propagation graph of the program, where vertices are instructions and edges are guarded by conditions that correspond to the left-hand side of bindings

2. Checking that this graph does not contain any cycle, and flattening it by transforming its transitive closure into direct associations between external events and lists of guarded instructions
3. For each external event, schedule the guarded instructions according to dependencies (e.g. an assignment to  $x$  must be processed before an instruction that depends on  $x$ ), and generate a function that implements the reaction to the event

Our prototype compiler targets the Obc IR of the Vélus compiler [5]. Obc is an imperative object-oriented language where each class has fields and methods. In our compiler, we use fields to store the values of properties and the activation state of permanent processes, and generate one method for each external event. Then, we reuse the Obc-to-Clight pass of Vélus to produce a Clight program. Clight is one of the frontend language of the CompCert verified compiler [12] on which Vélus, and therefore our prototype compiler, are based.

In the future, we hope to reuse the correctness proofs of Vélus and CompCert to build an end-to-end semantics preservation proof from the denotational semantics presented in this paper down to the assembly semantics provided by CompCert. The missing piece is, of course, a correctness proof that relates our denotational semantics for Smalite to the operational semantics of Obc [5, §4.1.2]. We expect this proof to be difficult, for three reasons. First, our denotational semantics model is clearly more abstract than the operational semantics used for Obc: the former asserts the existence of objects on which a set of relations hold, while the latter describes precisely how these objects are computed. Second, the correctness of our compilation scheme heavily depends on the well-formedness of the source program (absence of dependency loops, absence of contradictory assignments, etc.) which might be difficult to specify precisely and reason about. Last, the rule for reactions described in fig. 6 relies on a complex predicate to define the content of the set of events  $T$ . We are afraid that reasoning about this predicate in semantics preservation proofs will require proving implications in two directions (source-to-target and target-to-source), which might incur a significant amount of work for each compilation pass.

## 4.2 Static analysis, simplifications and optimizations

Our prototype compiler is, in some respects, very naïve, and future work will focus on improving it so that it generates more efficient imperative code, and accepts a larger set of source Smalite programs.

Indeed, compilation to imperative code with a fixed schedule places a restriction on which programs may be accepted: some programs with well-defined and deterministic semantics may be rejected during compilation. For instance, consider a program with only two bindings:  $x \rightarrow y$ ;  $y \rightarrow x$ . In theory, this program is not schedulable, as there is a dependency loop between the triggering of  $x$  and  $y$ . In practice however, the semantics of this program are clear: either  $x$  and  $y$  are both triggered, or they both are not. To generate imperative code, we need to somehow cut the dependency cycle, but it is not obvious how to do so in general.

Another, more complex loop is actually showcased in our example, on lines 29–30. Line 29 specifies that releasing `btn2` sets `count` back to 3. Line 30 specifies that changing `count` triggers a test that activates `btn2` if `count < 3`. How should these two bindings be scheduled? On the one hand, before checking the first binding, it must be determined whether or not `btn2` is active, so line 30 should be scheduled before line 29. On the other hand, line 29 sets `count`, so it should be scheduled before line 30 that listens to a change on `count`. This looks like a dependency loop, but on closer inspection the second scheduling is never useful: indeed, line 29 sets `count` to 3, therefore, `count < 3` will never be true on a cycle where this assignment is activated. A general way to filter this type of false dependency loops would be to “cut” chains of bindings that include conditions that will never evaluate

to true. To do so, we could use some type of static analysis such as abstract interpretation [6]. Moreover, statically simplifying conditions and cutting useless bindings would also make the generated code more efficient.

### 4.3 Extending Smalite

The example presented in section 2 could be simplified with a few higher-level constructs. First, we see a very similar code pattern of declaring an assignment and having a unique binding to that assignment at several places (lines 14, 15, 17, 29, ...). These could be simplified with a meta-process that encodes “binding to an assignment”. In Smala, this is implemented by assignment sequences, with syntax `press -> { 255 =: green }`. This feature could be added as part of a more general source language than the one we present in this paper, and compiled down to the simpler constructs of Smalite.

Another unwelcome repetition appears in the definitions of the buttons `btn1` and `btn2`, which are essentially identical bar a few parameter. The full Smala language allows the user to define parameterized components which can then be instantiated in more complex programs. It is not yet clear what would be the best way to treat such a feature in our formalization: either user-defined components could be inlined into a single Smalite program, or they could be compiled separately, making Smalite and the other intermediate representations more complex.

### 4.4 Graphical semantics and properties

The semantic model we propose in this paper describes how the internal state of a program is updated in reaction to an event. It does not specify how this internal state is related to the observable behavior of the program. In particular, all the components used in the example (`Frame`, `Rectangle`, etc) have graphical semantics: their activation, and the value of their property affects what is displayed on the screen. Furthermore, the events that are modeled by spikes (`close`, `released`, etc) correspond to user actions. The structure of graphical components could be formalized as a scene graph, while the behavior of interactions could be specified using low-level events (click on the mouse at specific coordinates, etc). Modeling these graphical and interactive semantics at the level of the Smala source language would facilitate two high-level goals.

First, mechanize a compilation correctness proof for the code that binds the reactive program to the system libraries that implement user interactions; for our prototype compiler, that would be SDL3. Writing this proof would first require axiomatizing the behavior of all the SDL3 functions in use: drawing functions would affect the scene graph, while event-listening functions would be related to low-level events.

Second, it would be possible to reason on graphical properties of the system. Indeed, many safety-critical interactive systems, such as airplane cockpit GUIs are bound by strict norms. For instance, the ED 143 [9] specifies the behavior of the Traffic Alert and Collision Avoidance System (TCAS), which prevents aircraft from crashing into each other by alerting the pilot of imminent collisions and requesting altitude change. In particular, it specifies how incoming aircraft should be displayed on the screen.

In [21], the authors present a technique to formally verify that this specification holds for a given implementation, by generating Weakest Preconditions (WPs) and verifying them using the Z3 SMT Solver [8]. We could ground these results by proving the correspondence of the WP generation algorithm with our semantics model, and using a tool such as SMTCoq [3] to transport the proof of correctness generated by Z3 into Rocq logic. This would give us a mechanized Rocq proof that these properties hold

for the source program and, by applying the compiler correctness proof mentioned in section 4.1, that they hold for the generated binary.

Another approach would be to reuse the work proposed in [20, 19], which proposes a denotational semantic model based on bigraphs for a UIDL very similar to ours. It then uses the PRISM model checker [11] to prove that the resulting bigraphical reactive system implies graphical properties of interest. To take advantage of these results, we would first need to prove that our own relational semantic model is equivalent to the bigraphical model.

## References

- [1] Govindaraj Anand et al. “Avionics display for Two-Seated aircraft using OpenGL”. In: *Artificial Intelligence and Evolutionary Computations in Engineering Systems: Computational Algorithm for AI Technology, Proceedings of ICAIECES 2020*. Springer, 2022, pp. 255–265. DOI: 10.1007/978-981-16-2674-6\_20.
- [2] Philippe Antoine and Stéphane Conversy. “Volta: The First All-Electric Conventional Helicopter”. In: *MEA 2017, More Electric Aircraft*. Bordeaux, France, Feb. 2017. URL: <https://enac.hal.science/hal-01609233>.
- [3] Michael Armand et al. “A Modular Integration of SAT/SMT Solvers to Coq through Proof Witnesses”. In: *Certified Programs and Proofs*. Vol. 7086. Berlin, Heidelberg: Springer Berlin Heidelberg, 2011, pp. 135–150. DOI: 10.1007/978-3-642-25379-9\_12.
- [4] Timothy Bourke, Basile Pesin, and Marc Pouzet. “Verified Compilation of Synchronous Dataflow with State Machines”. In: *ACM Trans. Embed. Comput. Syst.* 22 (5s Sept. 9, 2023), 137:1–137:26. DOI: 10.1145/3608102.
- [5] Léo Brun. “Mechanized Semantics and Verified Compilation for a Dataflow Synchronous Language with Reset”. PhD thesis. PSL Research University, June 2020. URL: <https://theses.hal.science/tel-03068862/>.
- [6] Patrick Cousot and Radhia Cousot. “Static determination of dynamic properties of programs”. In: *Proceedings of the Second International Symposium on Programming*. Dunod, Paris, France, 1976, pp. 106–130.
- [7] Mathieu Cousy et al. “AEON: Toward a concept of operation and tools for supporting engine-off navigation for ground operations”. In: *Toward Sustainable Aviation Summit (TSAS) 2022*. Toulouse, France, Oct. 2022, pp. 1–12. URL: <https://enac.hal.science/hal-03859908>.
- [8] Leonardo de Moura and Nikolaj Bjørner. “Z3: An Efficient SMT Solver”. In: *Tools and Algorithms for the Construction and Analysis of Systems*. Ed. by C. R. Ramakrishnan and Jakob Rehof. Berlin, Heidelberg: Springer, 2008, pp. 337–340. DOI: 10.1007/978-3-540-78800-3\_24.
- [9] EUROCAE. *ED 143 - Minimum Operational Performance Standards For Traffic Alert And Collision Avoidance System II (TCAS II)*. URL: <https://standards.globalspec.com/std/1609213/eurocae-ed-143>.
- [10] Stephan Huber, Johanna Gramlich, and Tobias Grundgeiger. “From Paper Flight Strips to Digital Strip Systems: Changes and Similarities in Air Traffic Control Work Practices”. In: *CSCW*. New York, NY, USA: Association for Computing Machinery, May 2020. DOI: 10.1145/3392833.

- [11] M. Kwiatkowska, G. Norman, and D. Parker. “PRISM 4.0: Verification of Probabilistic Real-time Systems”. In: *Proc. 23rd International Conference on Computer Aided Verification (CAV’11)*. Ed. by G. Gopalakrishnan and S. Qadeer. Vol. 6806. LNCS. Springer, 2011, pp. 585–591. DOI: 10.1007/978-3-642-22110-1\_47.
- [12] Xavier Leroy. “Formal Verification of a Realistic Compiler”. In: *Communications of the ACM* 52.7 (July 2009), pp. 107–115. DOI: 10.1145/1538788.1538814.
- [13] Yixiang Lim et al. “Avionics Human-Machine Interfaces and Interactions for Manned and Unmanned Aircraft”. In: *Progress in Aerospace Sciences* 102 (2018), pp. 1–46. DOI: 10.1016/j.paerosci.2018.05.002.
- [14] Magnus Madsen, Ondřej Lhoták, and Frank Tip. “A Semantics for the Essence of React”. In: *LIPICs, Volume 166, ECOOP 2020* 166 (2020). Ed. by Robert Hirschfeld and Tobias Pape, 12:1–12:26. DOI: 10.4230/LIPICs.ECOOP.2020.12.
- [15] Mathieu Magnaudet et al. “Djnn/Smala: A Conceptual Framework and a Language for Interaction-Oriented Programming”. In: *Proc. ACM Hum.-Comput. Interact.* 2 (EICS June 19, 2018), 12:1–12:27. DOI: 10.1145/3229094.
- [16] Ingo Maier, Tiark Rompf, and Martin Odersky. *Deprecating the observer pattern*. Tech. rep. EPFL-REPORT-148043. Lausanne: EPFL, 2010.
- [17] Alice Martin, Mathieu Magnaudet, and Stéphane Conversy. “Causette: user-controlled rearrangement of causal constructs in a code editor”. In: *Proceedings of the 30th IEEE/ACM International Conference on Program Comprehension*. 2022, pp. 241–252. DOI: 10.1145/3524610.3527885.
- [18] Brad A. Myers. “Separating application code from toolkits: Eliminating the spaghetti of callbacks”. In: *Proceedings of the 4th annual ACM symposium on User interface software and technology*. 1991, pp. 211–220. DOI: 10.1145/120782.120805.
- [19] Nicolas Nalpon. “Vers la vérification des langages de description d’interface utilisateur”. PhD thesis. INSA de Toulouse, Mar. 13, 2023. URL: <https://theses.hal.science/tel-04139613>.
- [20] Nicolas Nalpon, Cyril Allignol, and Celia Picard. “Towards a User Interface Description Language Based on Bigraphs”. In: *Theoretical Aspects of Computing – ICTAC 2022*. Ed. by Helmut Seidl, Zhiming Liu, and Corina S. Pasareanu. Vol. 13572. Cham: Springer International Publishing, 2022, pp. 360–368. DOI: 10.1007/978-3-031-17715-6\_23.
- [21] Daniel Prun and Pascal Béger. “Formal Verification of Graphical Properties of Interactive Systems”. In: *Proceedings of the ACM on Human-Computer Interaction* 6 (EICS June 14, 2022), pp. 1–30. DOI: 10.1145/3534521.
- [22] Simple DirectMedia Layer project. *SDL Wiki*. URL: <https://wiki.libsdl.org/SDL3/FrontPage>.
- [23] The Coq Development Team. *The Coq Proof Assistant*. Version 8.20. Zenodo, Sept. 4, 2024. DOI: 10.5281/zenodo.14542673.

# Bisimilarity and Simulatability of Processes Parameterized by Join Interactions

Clemens Grabmayer

GSSI  
L'Aquila, Italy  
clemens.grabmayer@gssi.it

Maurizio Murgia

GSSI  
L'Aquila, Italy  
maurizio.murgia@gssi.it

Departing from Larsen’s concept of parameterized bisimilarity of processes with respect to interaction with environments, we start an exploration of its natural weakening: bisimilarity of unrestricted join interactions with environments. Parameterized bisimilarity relates processes  $p$  and  $q$  with respect to an environment  $e$  if  $p$  and  $q$  behave bi-similarly while joining—respectively the same—transitions from  $e$ . The weakened variant relates processes  $p$  and  $q$  with respect to environment  $e$  if the join-interaction processes  $p \& e$  and  $q \& e$  of  $p$  and  $q$  with  $e$  are bisimilar. (Hereby join interactions  $r \& f$  facilitate a step with label  $a$  to  $r'$  &  $f'$  if and only if  $r$  and  $f$  permit  $a$ -steps to  $r'$  and  $f'$ , respectively.)

Join-interaction parameterized (ji-parameterized) bisimilarity coincides with parameterized bisimilarity for deterministic environments, but that it is a coarser equivalence in general. We explain how Larsen’s concept can be recovered from ji-parameterized bisimilarity by ‘determinizing’ interactions. We show that by adaptation to simulatability (simulation preorder) the same concept arises: parameterized simulatability coincides with ji-parameterized simulatability. For the discrimination preorder of (ji-)parameterized simulatability on environments we obtain the same result as Larsen did for parameterized bisimilarity. Also, we give a modal-logic characterization of (ji-)parameterized simulatability. Finally we gather open problems, and provide an outlook on our current related work.

## 1 Introduction

With the motivation of developing flexible formal methods for proving correctness of software programs incrementally, by showing compositional correctness under the formation of contexts, Larsen in [11, 12] introduced parameterized bisimilarity of processes as a helpful concept. It turned out, more recently, to be useful in an area with a similar motivation: contextual behavioural metrics (see work of Dal Lago and Murgia [6, 10]), which measure differences between programs as distances by means of pseudo-metrics. This is because parameterized bisimilarity provides natural examples of contextual behavioural metrics.

The idea underlying parameterized bisimilarity is that the behaviors of two processes are compared with respect to a third process that represents a common environment, in which both processes are placed, and with which both can interact separately. The environment is able to ‘consume’ a transition from a process by performing a transition with the same action label, after which both the process and the environment move to the target state of the interaction transition on their side, respectively. Such consumption interactions are intended to continue as long as possible. Yet in case that an environment state permits no transition with the same label as the current process state (this is the case, for example, if the environment or the considered process is in a deadlock state), the consumption process stops. Given this setup, processes  $p$  and  $q$  are called bisimilar with respect to an environment process  $e$  if  $p$  and  $q$  behave in a bisimilar way (fulfilling forth and back conditions as typical for bisimulations) for any pair of runs of *synchronous* consumption interactions of the environment with the two processes in which the environment takes the *same* transitions on its side.

It is distinctive for Larsen's concept of parameterized bisimilarity that the forth and back conditions of two processes  $p$  and  $q$ , and subsequently, of states reached via transitions from  $p$  and  $q$ , have to be verified *separately* for every run of the environment but while interacting *synchronously* with both processes. Indeed, comparisons of possible further interactions have to be carried out in synchronicity of the interactions, as long as the environment can interact with either of the processes. Thereby a mismatch is detected in the following situation: Suppose that by successful comparisons in a synchronous run derivative processes  $p'$  and  $q'$  as well as derivative environment  $e'$  are reached. Suppose further that  $e'$  permits, say, an  $a$ -transition that can be joined only with an  $a$ -transition from  $p'$ , but not from  $q'$  (in case  $q'$  does not permit  $a$ -transitions). Then it has been determined that  $p$  and  $q$  are not bisimilar with respect to  $e$ .

Parameterized bisimilarity thus compares the behavior of two processes with respect to *controlled* and *synchronous* interactions with an environment process. For determining whether two processes  $p$  and  $q$  are bisimilar with respect to an environment process  $e$  it is necessary to observe the consumption interactions of  $p$  with  $e$  and of  $q$  with  $e$  in a synchronous step-wise manner. It is not sufficient to be merely presented the completed processes that result from the interactions of  $p$  with  $e$ , and of  $q$  with  $e$ , respectively, and then to ask whether these results are bisimilar.

There are, however, conceivable practical situations, in which one lacks sufficient control over the environment process in order to perform, or merely to analyze, controlled and synchronous interactions with the considered processes. That is, situations in which a scientist has access only to the data of completed interactions of two processes with a given environment, but in which she lacks sufficient control over the environment in order to perform the two interactions synchronously in a step-by-step manner so that she can compare the behaviors that remain after each step.

Here we define, and start to investigate, the weaker concept of parameterized bisimilarity in which only the completed outcome processes of the possible interactions of two processes  $p$  and  $q$  with a given environment are compared as to whether they are bisimilar. For this purpose we stipulate that the consumption interaction takes place in the form of a 'join' operation ( $\&$ ) between each process and the environment, which produces transitions with the same action labels as the two interaction transitions. Indeed, only transitions with the same label from a process and the environment can be 'joined' to interact, and produce a resulting transition with again that same label. We call the concept of bisimilarity between the join interactions of each process with the environment 'join-interaction parameterized' (ji-parameterized) bisimilarity. Larsen briefly mentions this concept at the end of the article [12]. He calls it 'perhaps more immediate', but excludes it from further consideration on the basis that it lacks some distinctive properties that he was able to show for parameterized bisimilarity. (For more details, see the paragraph 'Larsen on ji-parameterized bisimilarity ...' in Section 5.) Although that is true, it remains the case that ji-parameterized bisimilarity has a much easier, and appealingly natural definition, and that it may be of practical use in cases in which parameterized bisimilarity cannot be used.

While these considerations may seem abstract, we got interested in studying ji-parameterized bisimilarity when we made the following concrete observations (many of which are explained here later):

- parameterized bisimilarity and ji-parameterized bisimilarity do not coincide, so we try to understand the conceptual difference between them, also by means of concrete examples (for an overview see Theorem 3.11);
- noticing that, for deterministic environments, parameterized bisimilarity and ji-parameterized bisimilarity coincide (see Proposition 3.10);
- recognizing that also parameterized bisimilarity can be formulated as bisimilarity of a special kind ( $\&\bullet$ ) of join interaction (see Definition 3.4 and Lemma 3.5);
- recognizing that simulation preorder adaptations of the two concepts of parameterized bisimilarity

and  $\text{ji}$ -parameterized bisimilarity do in fact coincide (see Proposition 3.10, (i));

- Larsens modal-logical characterization of parameterized bisimilarity for  $(\text{ji-})$ parameterized can be easily adapted to simulatability<sup>1</sup> (see Theorem 4.2 in Section 4);
- a natural modal-logical characterization also for  $\text{ji}$ -parameterized bisimilarity is worth studying, albeit we do not provide a solution to it in this work (see current work item (W1) in Section 5).

In Section 2 we summarize Larsen’s definition and main results on parameterized bisimilarity, and we define parameterized simulatability.<sup>1</sup> In Section 3 we define  $\text{ji}$ -parameterized bisimilarity and simulatability, and develop basic results about their relationships with parameterized bisimilarity and simulatability. We discover that Larsen’s theorem about the discrimination preorder induced by parameterized bisimilarity has an analogous version for the discrimination preorder that is induced by  $(\text{ji-})$ parameterized simulatability. Then in Section 4 we specialize Larsen’s modal-logical characterization of parameterized bisimilarity to  $(\text{ji-})$ parameterized simulatability. Finally in Section 5 we give a list that summarizes our results, we report about the literature and our ongoing related work, and we sketch further ideas and plans.

## 2 Preliminaries on Larsen’s parameterized bisimilarity

In this section we summarize definitions and results by Larsen in [11, 12] concerning parameterized bisimilarity, its induced discrimination preorder, and a modal-logical characterization for it. Additionally we define parameterized simulations, which relate to parameterized bisimulations in the same way as how simulations relate to bisimulations. We start with the basic concept of labeled transition system.

**Definition 2.1** (LTSs). A (simple) labeled transition system (LTS) is a triple  $\mathcal{T} = \langle \text{St}, A, \rightarrow \rangle$  that consists of a set  $\text{St}$  of states, a set  $A$  of actions, and a ternary transition relation  $\rightarrow \subseteq \text{St} \times A \times \text{St}$  that represents  $A$ -labeled transitions on the state set.

For LTSs we will use notation and terminology for basic properties as follows. For their stipulation, we let  $\mathcal{T} = \langle \text{St}, A, \rightarrow \rangle$  be an LTS. For  $\langle s, a, t \rangle \in \rightarrow$  we usually write  $s \xrightarrow{a} t$ , and say that “in state  $s$  there is a transition with label  $a$  (symbolizing an action called  $a$ ) to state  $t$ ”. In this case we also say that  $t$  is an  $a$ -derivative of  $s$ . For  $s \in \text{St}$  and  $a \in A$  we write  $s \xrightarrow{a}$  if there is an  $a$ -transition from  $s$  in  $\mathcal{T}$ , and  $s \not\xrightarrow{a}$  if there is no  $a$ -transition from  $s$  in  $\mathcal{T}$ .

We call an LTS  $\mathcal{T} = \langle \text{St}, A, \rightarrow \rangle$  *deterministic* (respectively *image-finite*) if  $|\{s' \mid s \xrightarrow{a} s'\}| \leq 1$  (and respectively if  $|\{s' \mid s \xrightarrow{a} s'\}| < \infty$ ) for all states  $s \in \text{St}$  and actions  $a \in A$ , that is, if every state of  $\mathcal{T}$  has at most one  $a$ -derivative (resp. has only finitely many  $a$ -derivatives), for all  $a \in A$ . We say that a state  $s \in \text{St}$  is *deterministic* (resp. *is image-finite*) if every state of  $\mathcal{T}$  that is reachable from  $s$  via a path of transitions has at most one  $a$ -derivative (resp. has only finitely many  $a$ -derivatives), for all  $a \in A$ .

Following well-known intuitions, bi-/simulations on such simple LTSs can be defined as follows.

**Definition 2.2** (bisimulation/bisimilar, simulation/simulated by). Let  $\mathcal{T} = \langle \text{St}, A, \rightarrow \rangle$  be an LTS.

- (i) A *bisimulation*  $B$  on  $\mathcal{T}$  is a non-empty binary relation  $B \subseteq \text{St} \times \text{St}$  with the following property: If  $s B t$  for  $s, t \in \text{St}$ , then the following two conditions hold:

$$\text{(forth)} \quad (\forall s' \in \text{St}) [s \xrightarrow{a} s' \implies (\exists t' \in \text{St}) [t \xrightarrow{a} t' \wedge s' B t']],$$

$$\text{(back)} \quad (\forall t' \in \text{St}) [t \xrightarrow{a} t' \implies (\exists s' \in \text{St}) [s \xrightarrow{a} s' \wedge s' B t']].$$

<sup>1</sup>We use ‘simulatability’ instead of ‘similarity’ for ‘simulation preorder’ for two reasons: to prevent the impression that a symmetrical relation were meant, and to avoid a possible confusion with ‘simulation equivalence’ that will also appear here.

For processes  $s, t \in \text{St}$ , we write  $s \sim t$  and say that  $s$  and  $t$  are *bisimilar* if there is a bisimulation  $B$  on  $\mathcal{T}$  such that  $s B t$ .

- (ii) A *simulation*  $B$  on  $\mathcal{T}$  is a non-empty binary relation  $S \subseteq \text{St} \times \text{St}$  with the following property: If  $s B t$  for  $s, t \in \text{St}$ , then the condition (forth) in (i) holds for  $B := S$  (but not necessarily the condition (back)). For processes  $s, t \in \text{St}$ , we write  $s \leq t$  and say that  $s$  can be simulated by  $t$ , and we write  $t \geq s$  and say that  $t$  can simulate  $s$ , if there is a simulation  $S$  on  $\mathcal{T}$  such that  $s S t$ .

Rather than defining simulations as weakened versions of bisimulations as above, bisimulations can also be defined from simulations, as follows. A relation  $B$  on an LTS  $\mathcal{T}$  is a bisimulation if and only if both  $B$  as its converse relation  $B^\smile := \{\langle t, s \rangle \mid s B t\}$  are simulations on  $\mathcal{T}$ .

For modeling processes whose behavior is studied according to how they interact with environments, both processes and environments are formalized as LTSs. However, in order to indicate their intended roles for occurring LTSs, we distinguish in notation, name, and in how they are referenced between *process LTSs*  $\mathcal{P} = \langle \text{Pr}, A, \rightarrow \rangle$ , whose states we call *processes*, and *environment LTSs*  $\mathcal{E} = \langle \text{Env}, A, \Rightarrow \rangle$ , whose states we call *environments*. We follow Larsen [11, 12] in this terminology and in most of the notation. Based on this distinction, Larsen defines parameterized bisimulation and bisimilarity as follows.

**Definition 2.3** (parameterized bisimulation (Larsen [11, 12])). Let  $\mathcal{P} = \langle \text{Pr}, A, \rightarrow \rangle$  be a process LTS, and let  $\mathcal{E} = \langle \text{Env}, A, \Rightarrow \rangle$  be an environment LTS. An  $\mathcal{E}$ -parameterized bisimulation  $\mathcal{B}$  on  $\mathcal{P}$  is an Env-indexed family  $\mathcal{B} = \{B_f\}_{f \in \text{Env}}$  of non-empty binary relations  $B_f \subseteq \text{Pr} \times \text{Pr}$  such that the following holds:<sup>2</sup> If  $p B_e q$  for  $e \in \text{Env}$ , then if  $e \xrightarrow{a} e'$  for  $a \in A$  the following conditions hold:

$$\begin{aligned} \text{(forth)} \quad & (\forall p' \in \text{Pr}) [p \xrightarrow{a} p' \implies (\exists q' \in \text{Pr}) [q \xrightarrow{a} q' \wedge p' B_{e'} q']], \\ \text{(back)} \quad & (\forall q' \in \text{Pr}) [q \xrightarrow{a} q' \implies (\exists p' \in \text{Pr}) [p \xrightarrow{a} p' \wedge p' B_{e'} q']]. \end{aligned}$$

For processes  $p, q \in \text{Pr}$ , and environments  $e \in \text{Env}$  we write  $p \sim_e q$  and say that  $p$  and  $q$  are *bisimilar with respect to  $e$*  if there is an  $\mathcal{E}$ -parameterized bisimulation  $\mathcal{B} = \{B_f\}_{f \in \text{Env}}$  such that  $p B_e q$ .

While simulation plays a crucial role in Larsen's main theorem on parameterized bisimulation, see Theorem 2.5 below, it is surprising that he did not also define the simulation version of this concept with only the forth condition from its progression conditions. For the reason that it can be linked directly to the simulation version of the concept of 'ji-parameterized bisimulation' that we will introduce in Section 3 (see Definition 3.3), we also define 'parameterized simulation' here.

**Definition 2.4** (parameterized simulation). Let  $\mathcal{P} = \langle \text{Pr}, A, \rightarrow \rangle$  be a process LTS, and let  $\mathcal{E} = \langle \text{Env}, A, \Rightarrow \rangle$  be an environment LTS. An  $\mathcal{E}$ -parameterized simulation  $\mathcal{S}$  on  $\mathcal{P}$  is an Env-indexed family  $\mathcal{S} = \{S_f\}_{f \in \text{Env}}$  of non-empty binary relations  $S_f \subseteq \text{Pr} \times \text{Pr}$  such that the following holds: If  $p S_e q$  holds for  $e \in \text{Env}$ , then for all  $a \in A$ , if  $e \xrightarrow{a} e'$  the condition (forth) in Def. 2.3 for  $B_{e'} := S_{e'}$  holds:

$$\text{(forth)} \quad (\forall p' \in \text{Pr}) [p \xrightarrow{a} p' \implies (\exists q' \in \text{Pr}) [q \xrightarrow{a} q' \wedge p' S_{e'} q]].$$

For processes  $p, q \in \text{Pr}$ , and environments  $e \in \text{Env}$  we write  $p \leq_e q$  and say that  $p$  can be simulated by  $q$  with respect to  $e$ , and  $q \geq_e p$  and say that  $q$  can simulate  $p$  with respect to  $e$ , if there is an  $\mathcal{E}$ -parameterized simulation  $\mathcal{S} = \{S_f\}_{f \in \text{Env}}$  such that  $p S_e q$ .

Also parameterized bisimulations can be defined from parameterized simulations: for process LTS  $\mathcal{P}$ , and environment LTS  $\mathcal{E}$ ,  $\mathcal{S} = \{S_f\}_{f \in \text{Env}}$  is an  $\mathcal{E}$ -parameterized bisimulation on  $\mathcal{P}$  if and only if  $\{S_f\}_{f \in \text{Env}}$ , and the family  $\{S_f^\smile\}_{f \in \text{Env}}$  of converse relations of  $S_f$  are  $\mathcal{E}$ -parameterized simulations.

<sup>2</sup>Note the occurrence of  $e'$  (instead of  $e$ ) in  $B_{e'}$  in both of the conditions (back) and (forth).

Larsen’s main result on parameterized bisimilarity concerns the *discrimination preorder*  $\sqsubseteq$  that orders environments according to their power of discriminating between processes. It is defined, for a given process LTS  $\mathcal{P} = \langle \text{Pr}, A, \rightarrow \rangle$  and a given environment LTS  $\mathcal{E} = \langle \text{Env}, A, \Rightarrow \rangle$  (so that, for all  $e \in \text{Env}$ , the relations  $\sim_e$  are then fixed as subsets of  $\text{Pr} \times \text{Pr}$ ), for all  $e, f \in \text{Env}$  by:

$$e \sqsubseteq f : \iff \approx_e \subseteq \approx_f \quad (\iff \sim_f \subseteq \sim_e). \quad (2.1)$$

Now Larsen’s result characterizes  $\sqsubseteq$  as coinciding with the simulation preorder  $\leq$  on environments. For the ‘completeness’ direction of this characterization to hold (“ $\Leftarrow$ ” in (2.2)), it is, however, necessary to assume that the underlying process LTS is, as Larsen formulates it, ‘sufficiently rich’ structurally. The weak natural assumption that he uses for the purpose of guaranteeing sufficient structural richness of any considered process LTS is that its set of processes is closed under action prefixing and finite summation (see Definition 3.2 in Section 3).

**Theorem 2.5** (Larsen [11, 12]). *The following logical equivalence holds, provided that the underlying process LTS is closed under action prefixing and finite summation, for all image-finite environments  $e, f$ :*

$$e \leq f \iff e \sqsubseteq f \quad (\iff \sim_f \subseteq \sim_e). \quad (2.2)$$

The implication “ $\Rightarrow$ ” holds for all (thus also for not necessarily image-finite) environments  $e$  and  $f$ .

Specifically for the direction “ $\Leftarrow$ ” in (2.2) Larsen provides an impressive, technical proof, which he found, as he writes, only after an intensive search that took several months.

We now turn to modal-logical characterizations of the relations of being able to be simulated by  $\leq$ , of bisimilarity  $\sim$ , and of parameterized bisimilarity  $\sim_e$ . For expressing properties of LTSs such as the existence of a transition with label  $a$  from a given state such that at the target state property  $\phi_0$  holds, modal formulas should include a diamond modality  $\langle a \rangle$  to build formulas like  $\langle a \rangle \phi_0$ . The set  $\mathcal{M}$  of simple modal formulas (and the set  $\mathcal{L}$  of positive formulas) are now defined with these diamond modalities and basic propositional connectives (resp. such connectives except negation) as constructors.

**Definition 2.6** (modal formulas). For given sets  $A$  of actions, we define the following classes of formulas:  $\mathcal{L}(A)$  of positive formulas, and  $\mathcal{M}(A)$  of (simple modal logic) formulas, via the following grammars:

$$\mathcal{L}(A) \quad \phi ::= \top \mid \phi \wedge \phi \mid \langle a \rangle \phi \quad (\text{where } a \in A), \quad (2.3)$$

$$\mathcal{M}(A) \quad \phi ::= \top \mid \neg \phi \mid \phi \wedge \phi \mid \langle a \rangle \phi \quad (\text{where } a \in A). \quad (2.4)$$

As above, we usually will keep the underlying set  $A$  of actions implicit, and write  $\mathcal{L}$  and  $\mathcal{M}$  for  $\mathcal{L}(A)$  and  $\mathcal{M}(A)$ , respectively.

**Definition 2.7** (satisfaction relation, sets of satisfied formulas). Let  $\mathcal{P} = \langle \text{Pr}, A, \rightarrow \rangle$  be a process LTS. The *satisfaction relation*  $\models \subseteq \text{Pr} \times \mathcal{M}$  on  $\mathcal{P}$  is defined by the following clauses:

$$\begin{aligned} p \models \top & : \iff p \in \text{Pr}, & p \models \phi_1 \wedge \phi_2 & : \iff p \models \phi_1 \text{ and } p \models \phi_2, \\ p \models \neg \phi_0 & : \iff p \not\models \phi_0, & p \models \langle a \rangle \phi_0 & : \iff \exists p' \in \text{Pr} (p \xrightarrow{a} p' \text{ and } p' \models \phi_0). \end{aligned}$$

by induction on the structure of formulas in  $\mathcal{M}$ . For all processes  $p \in \text{Pr}$  we define by:

$$\mathcal{L}(p) := \{ \phi \in \mathcal{L} \mid p \models \phi \}, \quad \mathcal{M}(p) := \{ \phi \in \mathcal{M} \mid p \models \phi \},$$

the set  $\mathcal{L}(p)$  of positive formulas in  $\mathcal{L}$  that are satisfied in  $p$ , and respectively, the set  $\mathcal{M}(p)$  of formulas in  $\mathcal{M}$  that are satisfied in  $p$ .

The classical characterization result via modal-logical formulas of the relations bisimilarity  $\sim$ , and ‘being able to be simulated by’  $\leq$  is the following well-known theorem by Hennessy and Milner.

**Theorem 2.8** (Hennessy, Milner [8]). *For all image-finite processes  $p$  and  $q$  the following statements hold:*

$$p \leq q \iff \mathcal{L}(p) \subseteq \mathcal{L}(q), \quad (2.5)$$

$$p \sim q \iff \mathcal{M}(p) = \mathcal{M}(q). \quad (2.6)$$

The implications “ $\Rightarrow$ ” hold for all (thus also for not necessarily image-finite) processes  $p$  and  $q$ .

For a modal-logical characterization of parameterized bisimilarity, the concept of negation closure of positive formulas will be needed. By departing slightly from Larsen’s exposition in [11, 12] we define it via a projection of general formulas to positive formulas.

**Definition 2.9** (positive-formula projection, negation closure). The *positive-formula projection* is the function  $|\cdot|_+ : \mathcal{M} \rightarrow \mathcal{L}$  that maps formulas  $\phi \in \mathcal{M}$  to positive formulas  $|\phi|_+ \in \mathcal{L}$ , and that is defined by induction on the structure of  $\phi$  via the following clauses, for all formulas  $\phi_0, \phi_1, \phi_2 \in \mathcal{M}$ :

$$|\top|_+ := \top, \quad |\neg\phi_0|_+ := |\phi_0|_+, \quad |\phi_1 \wedge \phi_2|_+ := |\phi_1|_+ \wedge |\phi_2|_+, \quad |\langle a \rangle \phi_0|_+ := \langle a \rangle |\phi_0|_+.$$

For every positive formula  $\phi \in \mathcal{L}$ , we define by  $\overline{\neg\phi} := \{\psi \in \mathcal{M} \mid |\psi|_+ = \phi\}$  the *negation closure of  $\phi$  in  $\mathcal{M}$* . For subclasses  $\mathcal{F} \subseteq \mathcal{L}$  of positive formulas, we define the *negation closure of  $\mathcal{F}$  in  $\mathcal{M}$*  by  $\overline{\neg\mathcal{F}} := \bigcup \{\overline{\neg\phi} \mid \phi \in \mathcal{F}\} = \{\psi \in \mathcal{M} \mid |\psi|_+ \in \mathcal{F}\}$ .

Larsen presents [11, 12] the following modal characterization theorem of parameterized bisimilarity, which he attributes to Colin Stirling. The characterization restricts consideration for possible discriminating formulas to those in the negation-closure of positive formulas that are satisfied by the environment.

**Theorem 2.10** (Stirling and Larsen, [11, 12]). *For all image-finite processes  $p, q$ , and environments  $e$ :*<sup>3</sup>

$$\begin{aligned} p \sim_e q &\stackrel{(\star)}{\iff} \mathcal{M}(p) \cap \overline{\neg\mathcal{L}(e)} = \mathcal{M}(q) \cap \overline{\neg\mathcal{L}(e)} \\ &\iff \forall \phi_0 \in \mathcal{L} [e \models \phi_0 \Rightarrow \forall \phi \in \overline{\neg\phi_0} (p \models \phi \Leftrightarrow q \models \phi)]. \end{aligned} \quad (2.7)$$

The implication “ $\Rightarrow$ ” in  $(\star)$  holds for all (thus also for not necessarily image-finite)  $p, q$ , and  $e$ .

### 3 Join-Interaction parameterized simulatability and bisimilarity

In this section we first define the weaker versions of parameterized bisimilarity and simulatability (simulation preorder) that are based on a definition of ‘join-interaction’ of LTSs (Definition 3.4): ji-parameterized simulatability and bisimilarity (Definition 3.3). Then we investigate the basic relationship between the new concepts and parameterized bisimilarity and simulatability (Theorem 3.11), and explain that also parameterized bisimilarity and simulatability can be viewed as bisimilarity and simulatability, resp., with respect to a special kind of join operation (Lemma 3.5). Finally we present a theorem (Theorem 3.14) that characterizes the discrimination preorder of (ji-)parameterized simulatability in analogy with Larsen’s characterization of the discrimination preorder of parameterized bisimilarity, see Theorem 2.5.

<sup>3</sup>The condition of being image-finite can be dropped for the environments  $e$ . This can be verified by means of a careful analysis of the proof in [11, 12] for this logical characterization.

### Definitions of ji-parameterized simulatability and bisimilarity

In order to prepare for the definition of ji-parameterized bisimilarity we define ‘join-interaction LTSs’ by using an operation of processes that Larsen calls ‘join’ [11, p.43,44]. Later we also need the subsequent stipulation of when a single LTS is closed under the ‘join’ operation.

**Definition 3.1** (join-interaction of LTSs). Let  $\mathcal{T}_1 = \langle \text{Pr}_1, A, \rightarrow_1 \rangle$  and  $\mathcal{T}_2 = \langle \text{Pr}_2, A, \rightarrow_2 \rangle$  two LTSs. By the *join-interaction of  $\mathcal{T}_1$  and  $\mathcal{T}_2$*  we mean the LTS  $\mathcal{T}_1 \& \mathcal{T}_2 = \langle \text{Pr}_1 \& \text{Pr}_2, A, \rightarrow \rangle$  where  $\rightarrow \subseteq (\text{Pr}_1 \& \text{Pr}_2) \times A \times (\text{Pr}_1 \& \text{Pr}_2)$  with  $\text{Pr}_1 \& \text{Pr}_2 := \{p_1 \& p_2 \mid p_1 \in \text{Pr}_1, p_2 \in \text{Pr}_2\}$  is defined via the transition system rule:

$$\frac{p_1 \xrightarrow{a}_1 p'_1 \quad p_2 \xrightarrow{a}_2 p'_2}{p_1 \& p_2 \xrightarrow{a} p'_1 \& p'_2}$$

The symbol “&” in processes  $p_1$  &  $p_2$  of  $\mathcal{T}_1$  &  $\mathcal{T}_2$  is to be understood as a term constructor that from any two processes  $q_1$  in  $\text{Pr}_1$  and  $q_2$  in  $\text{Pr}_2$  constructs a formal join-interaction process  $p_1 \& p_2$  in  $\text{Pr}_1 \& \text{Pr}_2$ .

**Definition 3.2** (closure of an LTS under action prefixing, sum, and join). Let  $\mathcal{T} = \langle \text{St}, A, \rightarrow \rangle$  be a labeled transition system. We say that  $\mathcal{T}$  is *closed under action prefixing*, resp. *under sum*, and resp. *under join* if for every states  $s, s_1, s_2 \in \text{St}$  there exists a state  $a.s \in \text{St}$  for all  $a \in A$ , resp. there exists a state  $s_1 + s_2 \in \text{St}$ , and resp. there exists a state  $s_1 \& s_2 \in \text{St}$  such that the respective transition rule below is satisfied:

$$\frac{}{a.s \xrightarrow{a} s} \quad \frac{s_i \xrightarrow{a} s'_i}{s_1 + s_2 \xrightarrow{a} s'_i} \text{ (where } i \in \{1, 2\}) \quad \frac{s_1 \xrightarrow{a} s'_1 \quad s_2 \xrightarrow{a} s'_2}{s_1 \& s_2 \xrightarrow{a} s'_1 \& s'_2}$$

We now proceed to defining join-interaction versions of parameterized simulatability, simulation equivalence, and bisimilarity as simulatability, simulation equivalence, and bisimilarity, respectively, of join-interactions between two processes and an environment.

**Definition 3.3.** Let  $\mathcal{P} = \langle \text{Pr}, A, \rightarrow \rangle$  be a process LTS, and let  $\mathcal{E} = \langle \text{Env}, A, \Rightarrow \rangle$  be an environment LTS.

For all environments  $e \in \text{Env}$ , we define three binary relations on  $\text{Pr}$ : *ji-parameterized simulatability*  $\leq_{\&e}$ , *ji-parameterized bisimilarity*  $\sim_{\&e}$ , and finally, *ji-parameterized simulation equivalence*  $(\leq \geq)_{\&e}$  where  $\leq_{\&e}, \sim_{\&e}, (\leq \geq)_{\&e} \subseteq \text{Pr} \times \text{Pr}$ , are defined by the following clauses, for all processes  $p, q \in \text{Pr}$ :

$$\begin{array}{l} (p \text{ can be simulated by } q \\ \text{with respect to join-interaction with } e) \end{array} \quad p \leq_{\&e} q : \iff p \& e \leq q \& e, \quad (3.1)$$

$$\begin{array}{l} (p \text{ is bisimilar to } q \\ \text{with respect to join-interaction with } e) \end{array} \quad p \sim_{\&e} q : \iff p \& e \sim q \& e, \quad (3.2)$$

$$\begin{array}{l} (p \text{ and } q \text{ are simulation equivalent} \\ \text{with respect to join-interaction with } e) \end{array} \quad p (\leq \geq)_{\&e} q : \iff p \leq_{\&e} q \wedge q \leq_{\&e} p, \quad (3.3)$$

where  $p \& e$  and  $q \& e$  on the right in (3.1) and in (3.2) are processes from the join-interaction LTS  $\mathcal{P} \& \mathcal{E}$ . By  $\geq_{\&e}$  we denote the converse of  $\leq_{\&e}$ , and express  $q \geq_{\&e} p$  verbally by saying that  $q$  can simulate  $p$  with respect to join-interaction with  $e$ .

### Relationship of ji-parameterized bisimilarity with parameterized bisimilarity

In order to recognize Larsen’s parameterized bisimilarity as bisimilarity with respect to a specific form of join-interaction, we introduce a ‘right-determinizing’ variant  $\&_{\bullet}$  of the join operation  $\&$ . For interactions of a process  $p$  with an environment  $e$  this operation yields the process  $p \&_{\bullet} e$  from which transitions are labeled by pairs  $\langle a, e' \rangle$  that result from joining an  $a$ -transition from  $p$  with an  $a$ -transition from  $e$  to target  $e'$ . In this way different environment steps that originally have the same action label are distinguished from

$\&\bullet$ -joins. Indeed, by making different targets of environment transitions visible as different transitions from  $p \&\bullet e$  for processes  $p$  and  $q$  and an environment  $e$ , a correspondence arises between bisimulations that link  $p \&\bullet e$  and  $q \&\bullet e$  and parameterized bisimulations that link  $p$  and  $q$  with respect to  $e$ .

**Definition 3.4** (right-determinizing join-interaction with environment LTSs). Let  $\mathcal{P} = \langle \text{Pr}, A, \rightarrow \rangle$  be a process LTS, and  $\mathcal{E} = \langle \text{Env}, A, \Rightarrow \rangle$  be an environment LTS. By the *right-determinizing join-interaction of  $\mathcal{P}$  and  $\mathcal{E}$*  we understand the LTS of the form  $\mathcal{P} \&\bullet \mathcal{E} = \langle \text{Pr} \&\bullet \text{Env}, A \times \text{Env}, \rightarrow \rangle$  with  $\text{Pr} \&\bullet \text{Env} := \{p \&\bullet e \mid p \in \text{Pr}, e \in \text{Env}\}$  and where  $\rightarrow \subseteq (\text{Pr} \&\bullet \text{Env}) \times (A \times \text{Env}) \times (\text{Pr} \&\bullet \text{Env})$  is defined by as transitions that are generated by the following rules:

$$\frac{p \xrightarrow{a} p' \quad e \xRightarrow{a} e'}{p \&\bullet e \xrightarrow{\langle a, e' \rangle} p' \&\bullet e'}$$

Hereby “ $\&\bullet$ ” in processes  $p \&\bullet e$  of  $\mathcal{P} \&\bullet \text{Env}$  has to be understood as a term constructor that from any process  $p \in \text{Pr}$  and environment  $e \in \text{Env}$  constructs a formal join-interaction process  $p \&\bullet e$  in  $\text{Pr} \&\bullet \text{Env}$ .

Now this variant “ $\&\bullet$ ” of the join operation “ $\&$ ” facilitates characterizations of parameterized simulatability and bisimilarity that are analogous in kind to the definitions of ji-parameterized simulatability and bisimilarity in Definition 3.3. As stated by logical equivalences in the following lemma, parameterized simulatability, and parameterized bisimilarity correspond to simulatability, and respectively to bisimilarity, of  $\&\bullet$ -interactions between two processes and an environment. From this we obtain inclusions of parameterized simulatability and bisimilarity in ji-parameterized simulatability and bisimilarity.

**Lemma 3.5.** *For all processes  $p$  and  $q$ , and environments  $e$  the following two chains of statements hold:*

$$p \leq_e q \iff (p \&\bullet e) \leq (q \&\bullet e) \quad p \sim_e q \iff (p \&\bullet e) \sim (q \&\bullet e) \quad (3.4)$$

$$\implies (p \& e) \leq (q \& e) \quad \implies (p \& e) \sim (q \& e) \quad (3.5)$$

$$\iff p \leq_{\&e} q, \quad \iff p \sim_{\&e} q,$$

where  $p \&\bullet e$  and  $q \&\bullet e$  are processes from the right-determinizing join-interaction LTS  $\mathcal{P} \&\bullet \mathcal{E}$ , and  $p \& e$  and  $q \& e$  are processes from the join-interaction LTS  $\mathcal{P} \& \mathcal{E}$ .

*Proof (Sketch).* We consider a process LTS  $\mathcal{P} = \langle \text{Pr}, A, \rightarrow \rangle$ , and an environment LTS  $\mathcal{E} = \langle \text{Env}, A, \Rightarrow \rangle$ .

We only argue for the chain of equivalences and implications on the right for  $\sim_e$ ,  $\sim$ , and  $\sim_{\&e}$ , since the chain of statements on the left for  $\leq_e$ ,  $\leq$ , and  $\leq_{\&e}$  can be demonstrated analogously.

We first consider statement (3.4). For showing “ $\implies$ ” it suffices to demonstrate that if  $\mathcal{B} = \{B_e\}_{e \in \text{Env}}$  is an  $\mathcal{E}$ -parameterized bisimulation on  $\mathcal{P}$ , then  $B := \{\langle p \&\bullet e, q \&\bullet e \rangle \mid p B_e q\}$  is a bisimulation on  $\mathcal{P} \&\bullet \mathcal{E}$ . For “ $\impliedby$ ” it suffices to show that if  $B$  is a bisimulation on  $\mathcal{P} \&\bullet \mathcal{E}$ , then  $\mathcal{B} = \{B_e\}_{e \in \text{Env}}$  with the defining clause  $B_e := \{\langle p, q \rangle \mid \langle p \&\bullet e, q \&\bullet e \rangle \in B\}$  for  $e \in \text{Env}$  is an  $\mathcal{E}$ -parameterized bisimulation on  $\mathcal{P}$ . Both auxiliary statements can be shown by using the conditions (forth) and (back) from the assumed (parameterized) bisimulation in order to demonstrate the conditions (forth) and (back) of the (parameterized) bisimulation in the conclusion of the implication.

The implication in (3.5) can be easily verified similarly: by showing that whenever  $B_\bullet$  is a bisimulation on  $\mathcal{P} \&\bullet \mathcal{E}$ , then  $B := \{\langle p \& e, q \& e \rangle \mid \langle p \&\bullet e, q \&\bullet e \rangle \in B_\bullet\}$  is a bisimulation on  $\mathcal{P} \& \mathcal{E}$ .  $\square$

In Figure 1 we illustrate the characterization (3.4) of parameterized bisimilarity  $\sim_e$  as bisimilarity of  $\&\bullet$ -interactions by an example. By using Lemma 3.5 we can now show that with respect to deterministic environments no difference arises between parameterized bisimilarity and ji-parameterized bisimilarity.

**Proposition 3.6.**  $\sim_e = \sim_{\&e}$  holds for all deterministic environments  $e$ .




Figure 1: Example that witnesses the correspondence (3.4) in Lemma 3.5: For the environment  $e := a.b + a$  and the processes  $p := a.b$  and  $q := e$ , it holds that  $p \approx_e q$  (indicated by the mismatches  $\times$  when building a parameterized bisimulation on the left), and also  $p \&_{\bullet} e \approx q \&_{\bullet} e$  (indicated by the mismatches  $\times$  when building a bisimulation on the right). Note that in contrast  $p \sim_{\&e} q$  holds  $p, q, e$ , see Fig. 2 later.

*Proof.* Let  $\mathcal{P} = \langle \text{Pr}, A, \rightarrow \rangle$  be a process LTS, and  $\mathcal{E} = \langle \text{Env}, A, \Rightarrow \rangle$  be an environment LTS. The Proposition follows from Lemma 3.5, once the converse implication “ $\Leftarrow$ ” in (3.5), which is the only implication that is missing there for  $\sim_e$  to coincide with  $\sim_{\&e}$ , is shown to hold for deterministic environments  $e$ :

$$e \text{ is deterministic} \implies [(p \&_{\bullet} e) \sim (q \&_{\bullet} e) \iff (p \& e) \sim (q \& e)]. \quad (3.6)$$

For this, it suffices to show, that whenever  $B$  is a bisimulation on  $\mathcal{P} \& \mathcal{E}$  in which all environments that occur in joins in pairs in  $B$  are deterministic, then  $B_{\bullet} := \{ \langle p \&_{\bullet} e, q \&_{\bullet} e \rangle \mid \langle p \& e, q \& e \rangle \in B \}$  is a bisimulation on  $\mathcal{P} \&_{\bullet} \mathcal{E}$ . The assumption that only deterministic environments occur in  $B$  is not too restrictive, because derivatives of deterministic environments are deterministic again.

We let  $B$  be a bisimulation on  $\mathcal{P} \& \mathcal{E}$  as described, and let  $B_{\bullet}$  be defined as above. We have to show that  $B_{\bullet}$  is a bisimulation on  $\mathcal{P} \&_{\bullet} \mathcal{E}$ .

For showing the condition (forth) for  $B_{\bullet}$  to be a bisimulation on  $\mathcal{P} \&_{\bullet} \mathcal{E}$ , we let  $\langle p \&_{\bullet} e, q \&_{\bullet} e \rangle \in B_{\bullet}$ , and a transition  $p \&_{\bullet} e \xrightarrow{l} r$  be arbitrary, where  $r \in \text{Pr} \&_{\bullet} \text{Env}$ ,  $l$  some label in  $A \times \text{Pr}$ . Due to operational semantics of  $\mathcal{P} \&_{\bullet} \mathcal{E}$ , this transition must actually be of the form  $p \&_{\bullet} e \xrightarrow{\langle a, e' \rangle} p' \&_{\bullet} e'$ , for  $p' \in \text{Pr}$  and  $e' \in \text{Env}$ . We have to show that there is  $s \in \text{Pr} \&_{\bullet} \text{Env}$  such that  $q \&_{\bullet} e \xrightarrow{\langle a, e' \rangle} s$  and  $\langle r, s \rangle = \langle p' \&_{\bullet} e', s \rangle \in B_{\bullet}$ .

From  $p \&_{\bullet} e \xrightarrow{\langle a, e' \rangle} p' \&_{\bullet} e'$  it follows that  $p \xrightarrow{a} p'$  and  $e \xrightarrow{a} e'$ . By the definition of  $\mathcal{P} \& \mathcal{E}$  it follows that there is also the transition  $p \& e \xrightarrow{a} p' \& e'$  in  $\mathcal{P} \& \mathcal{E}$ . From  $\langle p \&_{\bullet} e, q \&_{\bullet} e \rangle \in B_{\bullet}$  it follows by the definition of  $B_{\bullet}$  that  $\langle p \& e, q \& e \rangle \in B$ , and by the assumption on  $B$  also that  $e$  is deterministic. Then it follows from the condition (forth) of  $B$  as a bisimulation on  $\mathcal{P} \& \mathcal{E}$  that there is some  $s \in \text{Pr} \& \text{Env}$  such that  $q \& e \xrightarrow{a} s$  and  $\langle p' \& e', s \rangle \in B$ . By the definition of  $\mathcal{P} \& \mathcal{E}$  we find that  $s = q' \& e''$  and  $\langle p' \& e', q' \& e'' \rangle \in B$  for some  $q' \in \text{Pr}$  and  $e'' \in \text{Env}$  with  $q \xrightarrow{a} q'$  and  $e \xrightarrow{a} e''$ . But now from  $e \xrightarrow{a} e'$  and  $e \xrightarrow{a} e''$  we can conclude, because  $e$  is deterministic, that  $e'' = e'$ . From this we obtain  $\langle p' \& e', q' \& e' \rangle \in B$ , which entails  $\langle p' \&_{\bullet} e', q' \&_{\bullet} e' \rangle \in B_{\bullet}$ . From  $e \xrightarrow{a} e'$  and  $e \xrightarrow{a} e'$  we also obtain  $q \&_{\bullet} e \xrightarrow{\langle a, e' \rangle} q' \&_{\bullet} e'$ . Therefore we have found in  $s := q' \&_{\bullet} e'$  the desired  $s \in \text{Pr} \&_{\bullet} \text{Env}$  with  $q \&_{\bullet} e \xrightarrow{\langle a, e' \rangle} s$  and  $\langle r, s \rangle = \langle p' \&_{\bullet} e', s \rangle \in B_{\bullet}$ .

In this way we have established the condition (forth) for  $B_{\bullet}$  to be a bisimulation on  $\mathcal{P} \&_{\bullet} \mathcal{E}$ . Since the condition (back) can be verified analogously, we conclude that  $B_{\bullet}$  is indeed a bisimulation on  $\mathcal{P} \&_{\bullet} \mathcal{E}$ .

By having shown that  $B_{\bullet}$  is a bisimulation on  $\mathcal{P} \&_{\bullet} \mathcal{E}$  under the assumption that  $B$  is a bisimulation on  $\mathcal{P} \& \mathcal{E}$  in which only deterministic environments occur, we have established (3.6), from which the proposition follows from Lemma 3.5 as argued above.  $\square$




Figure 2: Example for witnessing  $\sim_e \neq \sim_{\&e}$ : For  $e := a.b + a$ ,  $p := a.b$ , and  $q := e$  it holds that  $p \not\sim_e q$  (indicated by the **mismatches**  $\times$  when building a parameterized bisimulation), but  $p \&e \sim q \&e$  (indicated by the **bisimulation links**) and hence  $p \sim_{\&e} q$ .

The proposition below clarifies which inclusions hold in general between  $\sim_e$ ,  $\sim_{\&e}$ , and  $(\leq\geq)_{\&e}$ .

**Proposition 3.7.** *The following set-theoretical relationships hold between parameterized bisimilarity  $\sim_e$ , ji-parameterized bisimilarity  $\sim_{\&e}$ , and ji-parameterized simulation equivalence  $(\leq\geq)_{\&e}$ :*

- (i)  $\sim_e \subseteq \sim_{\&e}$  for all environments  $e$ .
- (ii)  $\sim_e \neq \sim_{\&e}$  for some environments  $e$ , for which then  $\sim_e \subsetneq \sim_{\&e}$  holds due to (i).
- (iii)  $\sim_{\&e} \subseteq (\leq\geq)_{\&e}$  for all environments  $e$ .
- (iv)  $\sim_{\&e} \neq (\leq\geq)_{\&e}$  for some environments  $e$ , for which then  $\sim_{\&e} \subsetneq (\leq\geq)_{\&e}$  holds due to (iii).

The counterexample statements (ii) and (iv) hold under the proviso that environments are included among processes, they permit at least two actions, and are closed under action prefixing and sums. This can be weakened to merely require that  $a.b + a$  and  $a.b$  are contained among environments and processes.

*Proof.* Statement (i) follows directly from the chain of implications as guaranteed by Lemma 3.5.

A counterexample for (ii) is in Figure 2: We have that  $p \sim_{\&e} q$  holds due to  $p \&e = (a.b + a) \& (a.b) \simeq a.b + a \sim a.b + a + a + a \simeq (a.b + a) \& (a.b + a) = q \&e$ , where  $\simeq$  denotes being isomorphic, which shows  $p \&e \sim q \&e$ . However,  $p \not\sim_e q$  holds for the following reason: Suppose that  $p \sim_e q$  holds. Then due to  $e = a.b + a \xrightarrow{a} b$ , and the condition (back) of an underlying parameterized bisimilarity the transition  $q \xrightarrow{a} 0$  must be matched by the transition  $p \xrightarrow{a} b$  so that  $b \sim_b 0$  holds. However the latter is false, because  $b \not\sim_b 0$  holds, as the environment  $b$  and the process  $b$  can make a  $b$ -step, but  $0$  cannot. Statement (iii) follows from the fact that bisimilarity is symmetric and it is a simulation [20]. For (iv), let  $p$  and  $q$  be as in Figure 2, and let  $e = p$ . Then  $p (\leq\geq)_{\&e} q$  holds due to  $p \&e = (a.b) \& (a.b) \simeq a.b \leq\geq a.b + a \simeq (a.b + a) \& (a.b) = q \&e$ . However,  $p \not\sim_{\&e} q$ : indeed,  $q \&e \xrightarrow{a} b \&0 \sim 0$ , to which  $p \&e$  can only answer by reducing to  $b \&b \sim b$ , and clearly  $0 \not\sim b$ .  $\square$

### Parameterized simulatability coincides with ji-parameterized simulatability

While parameterized bisimilarity and ji-parameterized bisimilarity are two different relations in general by Proposition 3.7, (ii), it turns out that this does not hold for the corresponding two concepts of parameterized simulatability. For us it was surprising to find the proof of the first of the following two lemmas, which together show that parameterized simulatability and ji-parameterized simulatability coincide.

**Lemma 3.8.**  $\leq_{\&e} \subseteq \leq_e$  holds for all environments  $e$ .

*Proof.* We fix a process LTS  $\mathcal{P} = \langle \text{Pr}, A, \rightarrow \rangle$ , and an environment LTS  $\mathcal{E} = \langle \text{Env}, A, \Rightarrow \rangle$ .

As the crucial stepping stone, we show that  $\mathcal{S} = \{S_e\}_{e \in \text{Env}}$  as defined by, for all  $e \in \text{Env}$ :

$$S_e := \{ \langle p, q \rangle \in \text{Pr} \mid \exists e_2 \in \text{Env} [p \& e \leq q \& e_2] \} \subseteq \text{Pr} \times \text{Pr} \quad (3.7)$$

is an  $\mathcal{E}$ -parameterized simulation on  $\mathcal{P}$ . For this, we let  $e \in \text{Env}$ , and  $\langle p, q \rangle \in S_e$  be arbitrary. We assume that  $e \xrightarrow{a} e'$ , and  $p \xrightarrow{a} p'$  for some  $a \in A$ ,  $e' \in \text{Env}$ , and  $p' \in \text{Pr}$ . We have to show that there exists  $q' \in \text{Pr}$  with  $q \xrightarrow{a} q'$  such that  $\langle p', q' \rangle \in S_{e'}$ .

From  $e \xrightarrow{a} e'$  and  $p \xrightarrow{a} p'$  we find that  $p \& e \xrightarrow{a} p' \& e'$  holds. Due to  $\langle p, q \rangle \in S_e$  we can pick  $e_2 \in \text{Env}$  with  $p \& e \leq q \& e_2$ . It follows, by the forward-property (forth) of the (largest) simulation  $\leq$  applied to  $p \& e \leq q \& e_2$  and  $p \& e \xrightarrow{a} p' \& e'$ , and by the operational semantics of the join operation, that there are  $q' \in \text{Pr}$  and  $e'_2 \in \text{Env}$  such that  $q \& e_2 \xrightarrow{a} q' \& e'_2$ , as well as  $q \xrightarrow{a} q'$  and  $e_2 \xrightarrow{a} e'_2$  and with  $p' \& e' \leq q' \& e'_2$ . The latter shows that  $\langle p', q' \rangle \in S_{e'}$ , and thus we have found  $q \xrightarrow{a} q'$  such that  $\langle p', q' \rangle \in S_{e'}$ . In this way we have verified that  $\mathcal{S} = \{S_e\}_{e \in \text{Env}}$  as defined in (3.7) is an  $\mathcal{E}$ -parameterized simulation on  $\mathcal{P}$ .

For showing  $\leq_{\&e} \subseteq \leq_e$ , suppose now that  $p \leq_{\&e} q$  holds, for some  $p, q \in \text{Pr}$  and  $e \in \text{Env}$ . By the definition of  $\leq_{\&e}$ , this means that  $p \& e \leq q \& e$  holds. That, however, implies  $\langle p, q \rangle \in S_e$  due to (3.7). But since we have recognized  $\mathcal{S}$  as an  $\mathcal{E}$ -parameterized simulation, we conclude that  $p \leq_e q$  holds.  $\square$

**Lemma 3.9.**  $\leq_e \subseteq \leq_{\&e}$  holds for all environments  $e$ .

*Proof.* The inclusion as stated by the lemma follows from the chain of implications displayed on the left in Lemma 3.5, which as stated in its proof can be proved analogously as the implications on the right there. But since we dropped the argument there, we also provide the sketch of a direct proof here.

Let  $\mathcal{S} = \{S_e\}_{e \in \text{Env}}$  be an  $\mathcal{E}$ -parameterized simulation on a process LTS  $\mathcal{P} = \langle \text{Pr}, A, \rightarrow \rangle$  with respect to an environment LTS  $\mathcal{E} = \langle \text{Env}, A, \Rightarrow \rangle$ . Then it is easy to verify that:

$$S := \{ \langle p \& e, q \& e \rangle \mid p, q \in \text{Pr} \text{ and } e \in \text{Env} \text{ such that } p S_e q \} \subseteq \text{Pr} \& \text{Env}$$

is a simulation on  $\mathcal{P} \& \mathcal{E} = \langle \text{Pr} \& \text{Env}, A, \rightarrow \rangle$ . This statement implies that if  $p \leq_e q$ , then  $p \& e \leq q \& e$  follows, and hence  $p \leq_{\&e} q$ .  $\square$

**Proposition 3.10.** For all environments  $e$  the following two statements hold:

- (i)  $\leq_{\&e} = \leq_e$ .
- (ii)  $(\leq \geq)_{\&e} = (\leq \geq)_e$ .

*Proof.* The inclusions “ $\subseteq$ ” and “ $\supseteq$ ” that make up statement (i) are guaranteed by Lemma 3.8 and by Lemma 3.9, respectively. Then (ii) follows from (i) by:  $(\leq \geq)_{\&e} = \leq_{\&e} \cap \geq_{\&e} = \leq_e \cap \geq_e = (\leq \geq)_e$ .  $\square$

The theorem below collects results we have obtained about which inclusions hold in general between parameterized bisimilarity, ji-parameterized bisimilarity, and (ji-)parameterized simulation equivalence.

**Theorem 3.11.** The following set-theoretical relationships hold between parameterized bisimilarity, ji-parameterized bisimilarity, and (ji-)parameterized simulation equivalence, for environments  $e, f, g$ :

$$\begin{array}{llllllll} \sim_e & \subseteq & \sim_{\&e} & \subseteq & (\leq \geq)_e & = & (\leq \geq)_{\&e} & \text{(for all } e), \\ & \text{Prop. 3.7,(i)} & & \text{Prop. 3.7,(iii)} & & \text{Prop. 3.10,(ii)} & & \\ \sim_f & \subseteq & \sim_{\&f} & \subseteq & (\leq \geq)_f & = & (\leq \geq)_{\&f} & \text{(for some } f), \\ & \text{Prop. 3.7,(ii)} & & \text{Prop. 3.7,(iii)} & & \text{Prop. 3.10,(ii)} & & \\ \sim_g & \subseteq & \sim_{\&g} & \subseteq & (\leq \geq)_g & = & (\leq \geq)_{\&g} & \text{(for some } g), \\ & \text{Prop. 3.7,(i)} & & \text{Prop. 3.7,(iv)} & & \text{Prop. 3.10,(ii)} & & \end{array}$$

where the statements that guarantee the relationship in question are indicated.

### Discrimination preorder induced by (ji-)parameterized similarity

Larsen noted in [12, p.209–210]: “Due to the modal characterization [see Theorem 2.10] and the simple characterization of the discrimination ordering presented [see Theorem 2.5], we are confident that the notion of parameterized bisimulation equivalence proposed is indeed a natural one”. Indeed Larsen also explains that “the simulation ordering does not characterize the discrimination ordering generated by this alternative parameterized version [namely  $\sim_{\&e}$ ]”.

This is witnessed by the following proposition. Indeed it demonstrates that a characterization of the discrimination preorder induced by ji-parameterized bisimilarity  $\sim_{\&e}$  cannot, in analogy with Theorem 2.5 for  $\sim_e$ , be of the form  $e \leq f \iff \sim_{\&f} \subseteq \sim_{\&e}$ , for all environments  $e$  and  $f$ .

**Proposition 3.12.** *There are environments  $e$  and  $f$  such that:*

$$e \leq f \wedge \sim_{\&f} \not\subseteq \sim_{\&e} . \quad (3.8)$$

*Proof.* Let  $e = a.b$  and  $f = a.b + a$ . We have that  $e \leq f$ . Set  $p = e$  and  $q = f$ . We have that  $p \sim_{\&f} q$  but  $p \not\sim_{\&e} q$ , hence  $\sim_{\&f}$  is not contained in  $\sim_{\&e}$ .  $\square$

Unfortunately we have not yet found an appealing characterization of the discrimination preorder that is induced by  $\sim_{\&e}$ . We formulate this question together with a perhaps also interesting specialization as the open problems (P1) and (P2) in the conclusion.

However, and somewhat surprisingly, we do obtain characterizations analogous to Theorem 2.5 for the discrimination preorder on environments  $e$  with respect to (ji-)parameterized similarity  $\leq_{\&e}$  and  $\leq_e$ , and with respect to (ji-)parameterized simulation equivalence  $(\leq_{\geq})_{\&e}$  and  $(\leq_{\geq})_e$ .

Similar to the proviso for the ‘completeness’ direction “ $\Rightarrow$ ” in (2.2) of Theorem 2.5 our characterization of  $(\leq_{\geq})_{\&e}$  requires an assumption that guarantees that the structure of the underlying process LTS is sufficiently rich in relation to the environment LTS. While Larsen assumed closure under the formation of action prefixing and finite sums, we will assume the existence of a ‘universal’ process, and that the underlying process LTS contains the environment LTS, and is closed under the formation of joins. (The assumption of a universal process simplifies the proof, but can be dropped.)

Let  $\mathcal{P} = \langle \text{Pr}, A, \rightarrow \rangle$  be a process LTS. We say that a process  $u \in \text{Pr}$  is *universal* if for all  $a \in A$  there is a transition  $u \xrightarrow{a} u'$  with  $u' \in \text{Pr}$  in  $\mathcal{P}$  such that  $u' \sim u$  (consequently it permits all actions in transitions from any of its reachable states). Note that all universal processes in  $\mathcal{P}$  are bisimilar. If  $\mathcal{P}$  is additionally closed under the formation  $\&$ , then  $u \& p \sim p \& u \sim p$  holds for all  $p, u \in \text{Pr}$  where  $u$  is universal.

For proving our characterization below, Theorem 3.14 we will use the following lemma.

**Lemma 3.13.** *For all environments  $e, f$  it holds, provided that the environment LTS is closed under joins:*

$$e \leq f \& e \iff e \leq f .$$

*Proof.* For showing the direction “ $\Rightarrow$ ”, and the direction “ $\Leftarrow$ ” in the statement of the lemma, it suffices to prove that the relation  $\{ \langle e, f \rangle \mid e \leq f \& e \}$ , and respectively, that the relation  $\{ \langle e, f \& e \rangle \mid e \leq f \}$  is a simulation. Both of these statements can be verified in a straightforward manner.  $\square$

We now are in a position to show characterizations of the discrimination preorders of (ji-)parameterized simulatability and (ji-)parameterized simulation equivalence, as formulated by the theorem below.

**Theorem 3.14.** *The following logical equivalences hold, for all environments  $e$  and  $f$ , provided that: the underlying process LTS contains a universal process and the environment LTS, and additionally is closed under the formation of joins (but note that image-finiteness as in Theorem 2.5 is not required):*

$$e \leq f \iff \leq_{\&f} \subseteq \leq_{\&e} , \quad (3.9)$$

$$e \leq f \iff \geq_{\&f} \subseteq \geq_{\&e}, \quad (3.10)$$

$$e \leq f \iff (\leq\geq)_{\&f} \subseteq (\leq\geq)_{\&e}. \quad (3.11)$$

*Proof.* We let  $\mathcal{E} = \langle \text{Env}, A, \Rightarrow \rangle$  be an environment LTS, and we let  $\mathcal{P} = \langle \text{Pr}, A, \rightarrow \rangle$  be a process LTS that contains the universal process U.

We first note that (3.10) follows from, and is equivalent, to (3.9), because  $\geq_{\&f}$  is the converse relation of  $\leq_{\&e}$ . Furthermore (3.11) follows from (3.9) and (3.10), due to  $(\leq\geq)_{\&e} = \leq_{\&e} \cap \geq_{\&e}$ . Therefore it remains to prove (3.10). For that we show the two directions of 3.9, and proceed as follows:

“ $\Rightarrow$ ”: We show that  $e \leq f \implies \leq_f \subseteq \leq_e$ . Then (3.9) follows from Proposition 3.10, (i). So, let  $\{S_e\}_{e \in \text{Env}}$  be a  $\mathcal{E}$ -parameterized family of binary relations  $S_e \subseteq \text{Pr} \times \text{Pr}$  that is defined, for all  $e \in \text{Env}$  by:

$$p S_e q \iff \exists f \geq e [p \leq_f q]$$

It suffices to show that  $\{S_e\}_{e \in \text{Env}}$  is an  $\mathcal{E}$ -parameterized simulation. So, suppose  $p S_e q$ . Then  $p \leq_f q$  for some  $f \geq e$ . Suppose  $p \xrightarrow{a} p'$  and  $e \xrightarrow{a} e'$ . Then  $f \xrightarrow{a} f'$  for some  $f' \geq e'$ , and hence  $q \xrightarrow{a} q'$  for some  $q'$  such that  $p' \leq_{f'} q'$ . Then it follows that  $p' S_{e'} q'$  holds, as required.

“ $\Leftarrow$ ”: Assume  $\leq_{\&f} \subseteq \leq_{\&e}$ . Let  $u$  be a universal process in  $\mathcal{P}$ . Then  $u \& f \sim f$ . Moreover,  $f \leq f \& f$  (which is easy to show), and then  $u \leq_{\&f} f$ . By the assumption  $\leq_{\&f} \subseteq \leq_{\&e}$  we have that  $u \leq_{\&e} f$ . In other words:  $e \sim u \& e \leq f \& e$ . By Lemma 3.13 we get  $e \leq f$ , as required.  $\square$

## 4 Modal characterization of (ji-)parameterized simulatability

In this section we adapt the modal characterization of parameterized bisimilarity  $\sim_e$ , see Theorem 2.10, for (ji-)parameterized simulatability  $\leq_e$  and  $\leq_{\&e}$ . The crucial observation for our adaptation is Lemma 4.1 below which states that the set of positive formulas that a join interaction  $p_1 \& p_2$  satisfies is the intersection of the sets of positive formulas satisfied by the constituent processes  $p_1$  and  $p_2$ . Finally we explain why a similar line of argument is not possible in order to adapt the modal characterization for  $\sim_e$  to obtain one for ji-parameterized bisimilarity  $\sim_{\&e}$ . In doing so we provide some evidence for Larsen’s assessment, that (in view of that  $\sim_e \subsetneq \sim_{\&e}$  holds in general, see Proposition 3.7) “the modal characterization for  $\sim_e$  does not hold for  $\sim_{\&e}$ , and no other modal characterization seems immediate” [12, p.210]. However, we report about further work of ours on this issue in the final section (see (W1) in Section 5).

**Lemma 4.1.** *For all processes  $p$  and  $q$  of a process LTS  $\mathcal{P} = \langle \text{Pr}, A, \rightarrow \rangle$  it holds:*

$$\mathcal{L}(p \& q) = \mathcal{L}(p) \cap \mathcal{L}(q). \quad (4.1)$$

*Proof.* Statement (4.1) can be established by induction on the structure of positive modal formulas  $\phi$  according to their definition in grammar (2.4) of Definition 2.6.

The base case of (4.1) for  $\phi = \top$  is obviously true, because  $\top$  is satisfied for any process. It remains to establish the induction step for formulas of the forms  $\phi = \phi_1 \wedge \phi_2$  and  $\phi = \langle a \rangle \phi_0$ . Since in the first case the induction step is easy to demonstrate, we only treat the more interesting case of  $\phi = \langle a \rangle \phi_0$ . For this we argue as follows:

$$\begin{aligned} \langle a \rangle \phi_0 \in \mathcal{L}(p \& q) &\iff p \& q \models \langle a \rangle \phi_0 \\ &\iff (\exists p', q' \in \text{Pr}) [p \xrightarrow{a} p' \wedge q \xrightarrow{a} q' \wedge p' \& q' \models \phi_0] \\ &\iff (\exists p', q' \in \text{Pr}) [p \xrightarrow{a} p' \wedge q \xrightarrow{a} q' \wedge \phi_0 \in \mathcal{L}(p' \& q')] \end{aligned}$$

$$\begin{aligned}
&\stackrel{\text{IH}}{\iff} (\exists p', q' \in \text{Pr}) [p \xrightarrow{a} p' \wedge q \xrightarrow{a} q' \wedge \phi_0 \in \mathcal{L}(p') \cap \mathcal{L}(q')] \\
&\iff (\exists p' \in \text{Pr}) [p \xrightarrow{a} p' \wedge \phi_0 \in \mathcal{L}(p')] \wedge (\exists q' \in \text{Pr}) [q \xrightarrow{a} q' \wedge \phi_0 \in \mathcal{L}(q')] \\
&\iff \langle a \rangle \phi_0 \in \mathcal{L}(p) \wedge \langle a \rangle \phi_0 \in \mathcal{L}(q) \\
&\iff \langle a \rangle \phi_0 \in \mathcal{L}(p) \cap \mathcal{L}(q),
\end{aligned}$$

where we have marked by (IH) the logical equivalence in which the induction hypothesis is used.  $\square$

Based on this lemma, a modal characterization of  $\leq_{\&e}$  and  $\leq_e$  is now an easy consequence of the modal characterization of the simulation preorder  $\leq$  on processes, see (2.5) in Theorem 2.8. In this way we obtain, in analogy with Theorem 2.10, the following modal characterizations of (ji-)parameterized similarity and of (ji-)parameterized simulation equivalence with respect to positive formulas.

**Theorem 4.2.** *For all image-finite environments  $e$ , the following characterizations of  $\leq_{\&e}$  and  $\leq_e$ , as well as of  $(\leq_{\&e})_{\&e}$  and  $(\leq_e)_e$  hold for all image-finite processes  $p, q$ :*

$$p \leq_{\&e} q \quad (\iff p \leq_e q) \quad \iff \mathcal{L}(p) \cap \mathcal{L}(e) \subseteq \mathcal{L}(q) \cap \mathcal{L}(e), \quad (4.2)$$

$$p (\leq_{\&e})_{\&e} q \quad (\iff p (\leq_e)_e q) \quad \iff \mathcal{L}(p) \cap \mathcal{L}(e) = \mathcal{L}(q) \cap \mathcal{L}(e). \quad (4.3)$$

The implications “ $\Rightarrow$ ” in (4.2) and (4.3) hold also for not necessarily image-finite  $p, q$ , and  $e$ .

*Proof.* For (4.2) we argue as follows for all image-finite processes  $p$  and  $q$ , and environments  $e$ :

$$\begin{aligned}
p \leq_e q &\iff p \leq_{\&e} q && \text{(by Prop. 3.10, (i))} \\
&\iff p \&e \leq q \&e && \text{(by the definition of } \leq_{\&e} \text{)} \\
&\iff \mathcal{L}(p \&e) \subseteq \mathcal{L}(q \&e) && \text{(by (2.5) in Thm. 2.8)} \\
&\iff \mathcal{L}(p) \cap \mathcal{L}(e) \subseteq \mathcal{L}(q) \cap \mathcal{L}(e) && \text{(by using Lem. 4.1).}
\end{aligned}$$

The implication “ $\Rightarrow$ ” in the third equivalence statement holds also for not necessarily image-finite  $p, q$ , and  $e$  due to the Hennessy–Milner Theorem 2.8. Together with the fact that “ $\Rightarrow$ ” also holds for the other three equivalence statements above, this demonstrates that “ $\Rightarrow$ ” in (4.2) holds for all  $p, q$ , and  $e$ .

Statement (4.3) for the (ji-)parameterized simulation equivalences  $(\leq_{\&e})_{\&e}$  and  $(\leq_e)_e$  follows from (4.2) due to the definition of  $(\leq_{\&e})_{\&e}$  from  $\leq_{\&e}$  in (3.3), and of  $(\leq_e)_e$  from  $\leq_e$  in Definition 2.4.  $\square$

There is no obvious generalization of Lemma 4.1 that applies to all formulas of  $\mathcal{M}$ . In particular,  $\mathcal{M}(p_1 \& p_2) = \mathcal{M}(p_1) \cap \mathcal{M}(p_2)$  does not hold, because certainly “ $\subseteq$ ” is violated: in case that  $p_1$  and  $p_2$  are such that  $p_1 \xrightarrow{a}$  and  $p_2 \xrightarrow{a}$  holds for some  $a \in A$ , then  $\neg \langle a \rangle \top \in \mathcal{M}(p_1 \& p_2)$  due to  $(p_1 \& p_2) \xrightarrow{a}$ , but  $\neg \langle a \rangle \top \notin \mathcal{M}(p_1)$ , and hence  $\neg \langle a \rangle \top \notin \mathcal{M}(p_1) \cap \mathcal{M}(p_2)$ . Therefore the proof of the characterization above cannot be extended, at least not in an analogous manner, to obtain a modal characterization of ji-parameterized bisimilarity  $\sim_{\&e}$ . (But see the report about our current work (W1) in Section 5.)

Yet an interesting specialization of Lemma 4.1 concerns the specialized join operation  $\&_{\bullet}$  introduced in Definition 3.4:  $|\mathcal{M}(p \&_{\bullet} e)|_{\pi_A} = \mathcal{M}(p) \cap \overline{\mathcal{L}(e)}$  holds for all processes  $p$  and environments  $e$ , where  $|\cdot|_{\pi_A}$  projects modalities  $\langle \langle a, e \rangle \rangle$  in formulas to their action components  $\langle a \rangle$ . This observation can be used, together with the characterization of parameterized bisimilarity  $\sim_e$  via  $\&_{\bullet}$  in (3.4) of Lemma 3.5, to obtain, for Larsen’s characterization of  $\sim_e$  in Theorem 2.10, an alternative proof that is similar to the proof of Theorem 4.2 above.

## 5 Conclusion (summary, literature, current work, open problems, plans)

Here we first summarize our contributions in a list with references to statements in earlier sections. We then explain our path to the definition of  $\text{ji}$ -parameterized bisimilarity, and Larsen’s comments on the shortcomings of this concept. Furthermore we collect some references to the literature concerning work that has been done based on parameterized bisimilarity in the meantime. Subsequently we report about our current work on a modal characterization of (ji-)parameterized bisimilarity, and about generalizations of the modal characterizations here and by Stirling and Larsen. We also describe some open problems of which the solutions have evaded us thus far. Finally we mention our plan to investigate whether  $\text{ji}$ -parameterized bisimilarity can be used to refine Larsen’s results in his thesis [11] on a method to show program correctness under the formation of contexts.

**Contribution.** Below we provide a summary by listing the concepts that we have defined and the results we have obtained, together with references to the appertaining formal statements:

- (C1) We complemented Larsen’s parameterized bisimilarity  $\sim_e$  with respect to ‘synchronous’ interaction with environments  $e$  by also defining parameterized simulatability, the simulation preorder  $\leq_e$ , on processes with respect to ‘synchronous’ interaction with environment  $e$  (see Definition 2.4).
- (C2) We defined weaker versions  $\leq_{\&e}$  of  $\leq_e$  and  $\sim_{\&e}$  of  $\sim_e$  by relaxing the synchronicity condition of environment interaction for  $\leq_e$  and  $\sim_e$  to require only the existence of simulations, and respectively of bisimulations, between free join interactions ( $\&$ ) with environments  $e$  (see Definition 3.3).
- (C3) We showed that  $\leq_e$  and  $\sim_e$  can be characterized similarly to the definitions of  $\leq_{\&e}$ , and  $\sim_{\&e}$  via join interactions ( $\&$ ) as the existence of a simulation, and as bisimilarity, respectively, of free interactions with the specific form  $\&\bullet$  (see Definition 3.4) of join interactions that record targets of environment transitions in action labels (see Lemma 3.5).
- (C4) We established that  $\sim_e$  and  $\sim_{\&e}$  coincide for deterministic environments  $e$  (see Proposition 3.6).
- (C5) We settled the relationships between (ji-)parameterized bisimilarity  $\sim_e$  and  $\sim_{\&e}$ , and the (ji-)parameterized simulation equivalences  $(\leq\geq)_e$  and  $(\leq\geq)_{\&e}$ : for all environments  $\sim_e$  is contained in  $\sim_{\&e}$ , and furthermore  $\sim_{\&e}$  is contained in both of  $(\leq\geq)_e$  and  $(\leq\geq)_{\&e}$ , which coincide. The two inclusions in this chain are proper in general. (See Theorem 3.11).
- (C6) Larsen’s main technical result about  $\sim_e$  (see Theorem 2.5), that the discrimination preorder induced by  $\sim_e$  on environments coincides with the simulation preorder  $\leq$  on environments, does not hold analogously for the discrimination preorder induced by  $\sim_{\&e}$  (see Proposition 3.12). However, we showed that this coincidence with the simulation preorder  $\leq$  on environments *does* hold analogously for the discrimination preorders induced both by (ji-)parameterized similarity  $\leq_e = \leq_{\&e}$  and by (ji-)parameterized simulation equivalence  $(\leq\geq)_e = (\leq\geq)_{\&e}$  (see Theorem 3.14).
- (C7) We adapted Stirling and Larsen’s modal characterization of parameterized bisimilarity  $\sim_e$  (see Theorem 2.10) to obtain a modal characterization of (ji-)parameterized similarity  $\leq_e = \leq_{\&e}$  and also of (ji-)parameterized simulation equivalence  $(\leq\geq)_e = (\leq\geq)_{\&e}$  (see Theorem 4.2).

**Larsen on  $\text{ji}$ -parameterized bisimilarity, and our way to its definition.** We formulated  $\text{ji}$ -parameterized bisimilarity and  $\text{ji}$ -parameterized simulatability while reading Larsen’s article [12] from 1987, and trying to improve our intuitive understanding of parameterized bisimilarity. Afterwards we developed, in stages, the results that we report here. Only when diving deeper into the intricate proof of Larsen’s main result, the characterization of the discrimination preorder induced by parameterized bisimilarity  $\sim_e$  as simulatability of environments (see Theorem 2.5), did we find his remarks about an “alternative and perhaps more immediate parameterized version [of bisimulation equivalence]”. This passage appears on

page 210 in [12], at the end of Section 5 that is devoted to this central result. The version of bisimulation equivalence that Larsen sketches there coincides with  $\text{ji}$ -parameterized bisimilarity  $\sim_{\&e}$ .

Larsen refers to  $\text{ji}$ -parameterized bisimilarity in order to “give further support for the proposed parameterized version of bisimulation equivalence”, in addition to the following assessment: “Due to the modal characterization presented [...] and the simple characterization of the discrimination ordering presented [...], we are confident that the notion of parameterized bisimulation equivalence proposed is indeed a natural one.” As for the mentioned further evidence Larsen notes that  $\text{ji}$ -parameterized bisimilarity “lacks many of the properties presented in this paper”. Concretely he mentions three properties.

First, that “ $\sim_e$  is strictly included in  $\sim_{\&e}$  for all environments  $e$ ” (in general is meant [we use our notation for  $\sim_{\&e}$  here]), corresponding to Proposition 3.7, (i) and (ii). Second, that “thus the modal characterization for  $\sim_e$  does not hold for  $\sim_{\&e}$ , and no other modal characterization seems immediate.” This assessment stimulates us to work out (W1).

Finally third, Larsen writes that: “More important though is that the simulation ordering does not characterize the discrimination ordering generated by this alternative parameterized version[.]”, in contrast with his impressive and surprising main result in [12], Theorem 2.5 here, which shows that that is the case for parameterized bisimilarity  $\sim_e$ . For this observation Larsen uses a counterexample that is slightly different from the one we use for Proposition 3.12, the corresponding statement here.  $\text{Ji}$ -parameterized bisimilarity fits nicely in the recently proposed framework [1, 2]: the authors thereof advocate to make a clear distinction between processes and what they call tests. The composition of processes with tests give rise to instrumentations, which are to be understood as “compiled binaries ready to be executed” [2, page 6]. Environments in our setting can be seen as tests, and hence join is a notion of composition. In contrast,  $\sim_e$  lacks the composition operation (but still distinguishes processes and tests/environments). Below we formulate the question of a characterization of the discrimination preorder induced by  $\text{ji}$ -parameterized bisimilarity as the open problem (P1), and a specialization of this question as the open problem (P2).

**Literature on parameterized versions of bisimilarity.** Parameterized bisimilarity proved to be a very fruitful concept since its inception by Larsen in [11, 12]. His definition has been applied, specialized, and adapted in multiple ways in the meantime. Please see below for a few examples. But to the best of our knowledge this does not hold for the  $\text{ji}$ -parameterized concepts of simulatability and bisimilarity, apart from the passages in [12] that we cited and described above.

Parameterized bisimilarity in Larsen’s definition [11, 12] has later been called ‘relative bisimilarity’ and ‘relativized bisimilarity’ in [14] by Larsen and Milner, who used it also for the practical purpose of verifying the Alternating Bit Protocol [14]. As pointed out in [7], it was also the basis for ‘modal transition systems’ to which a large body of work has been devoted since, see for example [16, 15, 3, 9].

Environment parameterized bisimulations in the sense of Larsens’ definition or adapted and specialized variants of it have been used frequently, for example in [14, 13]. [19] introduces a notion of equivalence parameterized with respect to typing information, which, quoting from [19]: “can be seen as a disciplined instance of Larsen’s, in which one uses types to express constraints on the behaviors of the observers, rather than explicitly writing all their possible behavior”.

**Current Work.** We investigate modal-logical characterizations of ( $\text{ji}$ -)parameterized bisimilarity and simulatability, and of refinements of the modal characterizations already obtained by Larsen, and here.

(W1) We are working out a modal-logical characterization for  $\text{ji}$ -parameterized bisimilarity  $\sim_{\&e}$  that is based on a game characterization of  $\sim_{\&e}$ . However, our characterization will not just be of a simple form comparable to Theorem 2.10 and Theorem 4.2, for all environments  $e$ :

$$p \sim_{\&e} q \iff \mathcal{M}(p) \cap \mathcal{F}(e) = \mathcal{M}(q) \cap \mathcal{F}(e) \quad (\text{for all image-finite processes } p, q).$$

where  $\mathcal{F} \subseteq \mathcal{M}$  would be appropriately defined formulas with then  $\mathcal{F}(g) := \{\phi \in \mathcal{F} \mid g \models \phi\}$  defined for all environments  $g$ . It is nevertheless interesting to note that since  $\sim_{\&e} \subseteq \sim_e$  holds (due to  $\sim_e \subseteq \sim_{\&e}$  by Proposition 3.7, (i)), that whenever  $p \sim_{\&e} q$  holds, always  $p \sim_e q$  follows, and a formula  $\phi \in (\mathcal{M}(p) \cap \overline{\mathcal{L}(e)}) \Delta (\mathcal{M}(q) \cap \overline{\mathcal{L}(e)})$  (where  $\Delta$  denotes symmetric difference) that distinguishes  $p$  and  $q$  can always be found via Larsen’s characterization, Theorem 2.10.

(W2) The restriction to image-finite processes for the modal characterizations of simulatability  $\leq$  and bisimilarity  $\sim$  by Hennessy and Milner (Theorem 2.8) can be dropped by permitting infinitary formulas with infinite conjunctions. Indeed, Milner has described such an adaptation for infinitary formulas in [18].

We want to obtain similar extensions to not necessarily image-finite processes for Larsen’s characterization of  $\sim_e$  (Theorem 2.10) and our ones of  $\leq_e = \leq_{\&e}$  and  $(\leq \geq)_e = (\leq \geq)_{\&e}$  (Theorem 4.2).

**Open problems.** As problems to which (satisfactory) answers have evaded us so far, we want to mention:

- (P1) How can the discrimination order for  $\sim_{\&e}$  be characterized? Note that a similar characterization in terms of simulatability  $\leq$  as for the discrimination order of  $\sim_e$  in Thm. 2.5 by Larsen, and for  $\leq_{\&e}$  and  $(\leq \geq)_{\&e}$  in Theorem. 3.14, is not possible due to Proposition 3.12.
- (P2) Does equality of ji-parameterized bisimilarity with respect to environments  $e$  and  $f$  coincide with bisimilarity of  $e$  and  $f$ ? Equivalently, does the implication “ $\Leftarrow$ ” hold in the following statement (of which “ $\Rightarrow$ ” is easy to verify), for all environments  $e$  and  $f$ :

$$e \sim f \stackrel{?}{\iff} \sim_{\&e} = \sim_{\&f} .$$

**Future research.** As two lines of research for which the concept of ji-parameterized bisimilarity may lead to new insights we mention: a continuation of Larsen’s work in his thesis [11] towards flexible formal methods for showing compositionality of program correctness (see (F1)), and consequences for finding interesting contextual behavioural metrics as introduced in [6] (see (F2)):

- (F1) An interesting future work is the study of compositionality properties of  $\sim_{\&e}$ , that is the behavior of  $\sim_{\&e}$  up to context. A context is typically defined as a syntactic process  $C$  (expressed in some process algebra) with a hole  $\square$ . Notation  $C[p]$  is used for the process obtained upon substitution of  $p$  for the hole in  $C$ . In general,  $\sim_{\&e}$  is not preserved by contexts: Consider processes  $p = a + b$ ,  $q = a$ , environment  $e = a.b$  and context  $C = a.\square$ . We have that  $p \sim_{\&e} q$ , but  $C[p] = a.a + b \not\sim_{\&e} a.a = C[q]$ . Notice that the above example also applies to  $\sim_e$ . Indeed, including a process in a context intuitively also affects the environment, as shown in a study of the compositionality of  $\sim_e$  in [12]. The idea in that work is to introduce parametric environment-transformer<sup>4</sup>  $T_C$  which preserves  $\sim_e$  in the following sense:

$$p \sim_{T_C(e)} q \implies \langle C, p \rangle \equiv_e \langle C, q \rangle ,$$

where  $\langle C, p \rangle \equiv_e \langle C, q \rangle$  intuitively means that “ $C[p] \sim_{\&e} C[q]$  with  $C$  interacting identically with  $p$  and  $q$ ” [12] (we omit the formal definition for brevity). We speculate that, for  $\sim_{\&e}$ , the requirement “ $C$  interacting identically with  $p$  and  $q$ ” could be removed. If so, the compositionality of  $\sim_{\&e}$  could be expressed as follows (for an appropriate environment- transformer  $T'_C$ ):

$$p \sim_{\&T'_C(e)} q \implies C[p] \sim_{\&e} C[q] .$$

<sup>4</sup>We use a different notation than [12]. There,  $T_C(e)$  is rendered as  $wie_{\mathbb{E}}(C, e)$

- (F2) The relatively recent work [6] shows that from  $\sim_e$  (and quantitative generalizations of it) one can extract a generalized pseudo-metric between processes, where the codomain of the metric is the set of environments (under some closure assumptions). The idea is that the distance  $d(p, q)$  between processes  $p, q$  is defined as the largest environment  $e$  (according to (2.1)) such that  $p \sim_e q$ . An obvious future work is exploring whether a metric can be extracted for  $\sim_{\&e}$ . The main challenge is finding the right notion of “largest environment” for  $\sim_{\&e}$ , which is related to open problem (P1).

**Acknowledgment.** We thank the reviewers for their questions about the literature on parameterized bisimilarity since its inception, and concerning our work on the logical characterization of  $\sim_{\&e}$ . We are also grateful for lists of comments that pointed us to necessary corrections and adaptations of details. We thank Clément Aubert for a direct discussion after the workshop, and in particular for pointing us to interesting connections with work [1, 2] of his and Daniele Varacca.

## References

- [1] Clément Aubert & Daniele Varacca (2021): *Process, Systems and Tests: Three Layers in Concurrent Computation*. In: *ICE, EPTCS* 347, pp. 1–21, doi:10.4204/EPTCS.347.1.
- [2] Clément Aubert & Daniele Varacca (2022): *Processes against tests: On defining contextual equivalences*. *J. Log. Algebraic Methods Program.* 129, p. 100799, doi:10.1016/J.JLAMP.2022.100799.
- [3] Nikola Beneš, Jan Křetínský, Kim G. Larsen, Mikael H. Møller & Jiří Srba (2011): *Parametric Modal Transition Systems*. In Tevfik Bultan & Pao-Ann Hsiung, editors: *Automated Technology for Verification and Analysis*, Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 275–289, doi:10.1007/978-3-642-24372-1\_20.
- [4] J. A. Bergstra & J. W. Klop (1990): *An Introduction to Process Algebra*. In J. C. M. Baeten, editor: *Applications of Process Algebra*, Cambridge Tracts in Theoretical Computer Science, Cambridge University Press, p. 1–22, doi:10.1017/CB09780511608841. Available at [https://dspace.library.uu.nl/bitstream/handle/1874/13555/bergstra\\_90\\_introduction\\_process.pdf?sequence=3](https://dspace.library.uu.nl/bitstream/handle/1874/13555/bergstra_90_introduction_process.pdf?sequence=3).
- [5] Jan A Bergstra & Jan Willem Klop (1986): *Algebra of Communicating Processes*. Technical Report 89-138, CWI Amsterdam. Available at <https://ir.cwi.nl/pub/1778/1778D.pdf>.
- [6] Ugo Dal Lago & Maurizio Murgia (2023): *Contextual Behavioural Metrics*. In Guillermo A. Pérez & Jean-François Raskin, editors: *34th International Conference on Concurrency Theory (CONCUR 2023), Leibniz International Proceedings in Informatics (LIPIcs)* 279, Schloss Dagstuhl – Leibniz-Zentrum für Informatik, Dagstuhl, Germany, pp. 38:1–38:17, doi:10.4230/LIPIcs.CONCUR.2023.38.
- [7] Uli Fahrenberg, Kim Guldstrand Larsen, Axel Legay & Louis-Marie Traonouez (2014): *Parametric and Quantitative Extensions of Modal Transition Systems*. In Saddek Bensalem, Yassine Lakhneck & Axel Legay, editors: *From Programs to Systems. The Systems perspective in Computing*, Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 84–97, doi:10.1007/978-3-642-54848-2\_6.
- [8] Matthew Hennessy & Robin Milner (1985): *Algebraic Laws for Nondeterminism and Concurrency*. *J. ACM* 32(1), p. 137–161, doi:10.1145/2455.2460.
- [9] Michael Huth, Radha Jagadeesan & David Schmidt (2001): *Modal Transition Systems: A Foundation for Three-Valued Program Analysis*. In David Sands, editor: *Programming Languages and Systems*, Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 155–169, doi:10.1007/3-540-45309-1\_11.
- [10] Ugo Dal Lago & Maurizio Murgia (2023): *Contextual Behavioural Metrics (Extended Version)*, doi:10.48550/arXiv.2307.07400. arXiv:2307.07400.
- [11] Kim G. Larsen (1986): *Context-Dependent Bisimulation between Processes*. Ph.D. thesis, University of Edinburgh. Available at <https://era.ed.ac.uk/bitstream/handle/1842/11030/Larsen1986.pdf>.
- [12] Kim G. Larsen (1987): *A Context Dependent Equivalence between Processes*. *Theoretical Computer Science* 49(2), pp. 185–215, doi:10.1016/0304-3975(87)90007-7.

- [13] Kim G. Larsen, Ulrik Larsen & Andrzej Wasowski (2005): *Color-Blind Specifications for Transformations of Reactive Synchronous Programs*. In Maura Cerioli, editor: *Fundamental Approaches to Software Engineering*, Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 160–174, doi:10.1007/978-3-540-31984-9\_13.
- [14] Kim G. Larsen & Robin Milner (1992): *A Compositional Protocol Verification Using Relativized Bisimulation*. *Information and Computation* 99(1), pp. 80–108, doi:10.1016/0890-5401(92)90025-B.
- [15] Kim G. Larsen, Ulrik Nyman & Andrzej Wasowski (2007): *On Modal Refinement and Consistency*. In Luís Caires & Vasco T. Vasconcelos, editors: *CONCUR 2007 – Concurrency Theory*, Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 105–119, doi:10.1007/978-3-540-74407-8\_8.
- [16] Kim Guldstrand Larsen (1990): *Modal Specifications*. In Joseph Sifakis, editor: *Automatic Verification Methods for Finite State Systems*, Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 232–246, doi:10.1007/3-540-52148-8\_19.
- [17] Robin Milner (1980): *A Calculus of Communicating Systems*, 1980 edition. Springer Berlin Heidelberg, doi:10.1007/3-540-10235-3.
- [18] Robin Milner (1985): *Lectures on a Calculus for Communicating Systems*. In Stephen D. Brookes, Andrew William Roscoe & Glynn Winskel, editors: *Seminar on Concurrency*, Springer Berlin Heidelberg, Berlin, Heidelberg, pp. 197–220, doi:10.1007/3-540-15670-4\_10.
- [19] Benjamin C. Pierce & Davide Sangiorgi (2000): *Behavioral Equivalence in the Polymorphic Pi-Calculus*. *J. ACM* 47(3), p. 531–584, doi:10.1145/337244.337261.
- [20] Davide Sangiorgi (2011): *Introduction to Bisimulation and Coinduction*. Cambridge University Press, doi:10.1017/CB09780511777110.

# A Formalization of the Reversible Concurrent Calculus CCSK<sup>P</sup> in Beluga

Gabriele Cecilia 

School of Computer & Cyber Sciences,  
Augusta University, Augusta, USA

`gcecilia@augusta.edu`

Reversible concurrent calculi are abstract models for concurrent systems in which any action can potentially be undone. Over the last few decades, different formalisms have been developed and their mathematical properties have been explored; however, none have been machine-checked within a proof assistant. This paper presents the first Beluga formalization of the Calculus of Communicating Systems with Keys and Proof labels (CCSK<sup>P</sup>), a reversible extension of CCS. Beyond the syntax and semantics of the calculus, the encoding covers state-of-the-art results regarding three relations over proof labels – namely, dependence, independence and connectivity – which offer new insights into the notions of causality and concurrency of events. As is often the case with formalizations, our encoding introduces adjustments to the informal proof and makes explicit details which were previously only sketched, some of which reveal to be less straightforward than initially assumed. We believe this work lays the foundations for future reversible concurrent calculi formalizations.

## 1 Introduction

Concurrency in computer science refers to the simultaneous execution of multiple operations or computations in a shared environment. It is a fundamental aspect of modern computing, with practical use in several domains such as operating systems, networking and distributed systems. Process calculi like CCS [16] and the  $\pi$ -calculus [17] are well-studied and established mathematical models for formally describing and reasoning about concurrent systems.

In recent years, reversing computations in concurrent systems has gained significant attention, with applications in fields like hardware, software and biochemistry [22]. Enriching concurrent systems with reversibility poses its own set of challenges: for instance, it requires providing some kind of history-preserving mechanism to take track of past actions. Additionally, undoing computation steps in a parallel setting is more complex than in a sequential system: as explained in Fig. 1, reversing a specific action performed by a single thread may require knowing, and eventually undoing, the actions of the other threads it has previously interacted with.

Reversible concurrent calculi address such challenges in various ways. For example, Reversible CCS (RCCS) [11] equips processes with a memory that records information about past computations; conversely, CCS with Keys (CCSK) [20] associates unique keys to each forward action. The latter has been recently upgraded to CCSK with Proof labels (CCSK<sup>P</sup>) [4], which features a proved transition system in the fashion of Degano and Priami [12]; proof labels enable the definition of dependence and independence for both forward and backward transitions. In this framework, the contributions brought by Aubert et al. [3] merit attention. The authors are the first to introduce separate axioms for the relations of dependence, independence and connectivity on proof labels: such relations are proved to be sound, interrelated, and linked to the broader notions of concurrency and causality of events. Additionally, the

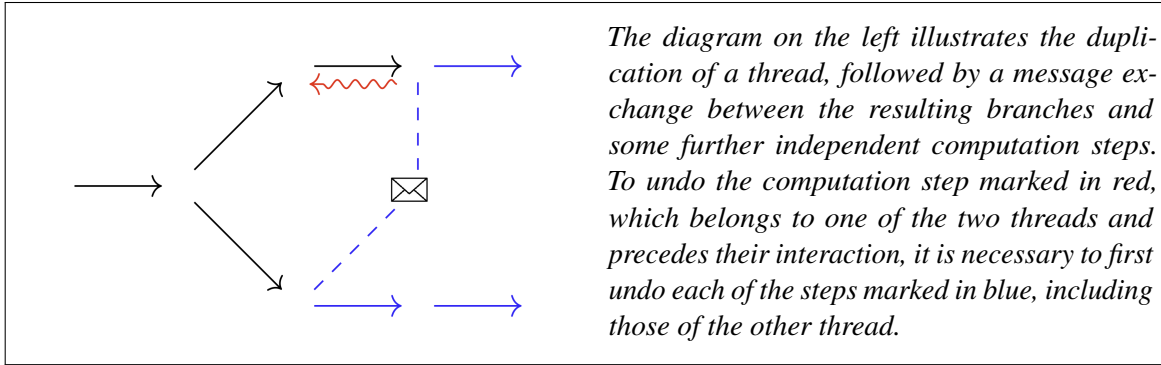


Figure 1: Example of reversal of computation steps in a concurrent setting.

authors outline the difference between various kinds of bisimulations, such as the history preserving bisimulation for CCS.

Formal verification has become a cornerstone in the development of new systems, certifying the correctness of their syntax, semantics and behavioral properties in the most reliable way. In the case of concurrent calculi, there is a long tradition which spans from the early mechanizations by [18] and [15] in HOL to the recent contributions of the Concurrent Calculi Formalisation Benchmark [6]; we also recall the Rocq formalization of the  $\pi$ -calculus by Honsell et al. [14], which has been the baseline for numerous higher-order abstract syntax (HOAS) [19] mechanizations. When it comes to reversible concurrent calculi, however, the landscape looks rather different. Despite the availability of C# and Java implementations of CCSK [10][1], no machine-checked formalization of reversible concurrent calculi currently exists – at least to the best of our knowledge.

This paper presents the first formalization of  $\text{CCSK}^P$  in Beluga [21]. Our encoding covers the core definitions of the system: its syntax, semantics, and the relations of dependence, independence and connectivity on proof labels. Additionally, this work formalizes the central results presented in Sections 3 and 4 of [3], including the complementarity of dependence and independence and the relationship between connectivity of transitions and proof labels. The proofs come along with a library of auxiliary lemmas regarding processes, keys and transitions.

Formalizations typically require small adjustments to fit the proof assistant’s framework, while carefully addressing any of the details which are taken for granted in the informal proof. This encoding is no exception: the formalization process led to minor refinements in the definition of connectivity over proof labels and clarified the proof of one of the aforementioned results, which called for a different approach separating base and inductive cases. Beluga, as a proof assistant, is ideal for reasoning about deductive systems together with their meta-theory, as it naturally supports encodings of object-level binding constructs through higher-order abstract syntax and allows pairing terms with the contexts that give them meaning [21]. Although our HOAS encoding leverages Beluga’s strengths and showcases its versatility, it also deals with limitations such as the lack of syntactic sugar for existentials or conjunctions; considerations on its adoption are further elaborated in the conclusions.

The paper is structured as follows. Section 2 provides an informal description of  $\text{CCSK}^P$  and the results under our study. Section 3 presents the Beluga formalization of such notions and properties. Section 4 contains a technical overview of the formalization, a summary of our contributions and possible future work directions. The artifact is archived in Zenodo [7] and available in the associated GitHub repository <https://github.com/CinRC/A-Beluga-Formalization-of-CCSKP>.

## 2 CCSK<sup>P</sup>

In this section we recall the main definitions and properties of CCSK<sup>P</sup>, as outlined in [3]. We assume familiarity with the basic notions of CCS. Definitions are hyperlinked to their encoding in the repository.

### 2.1 Syntax

As in the standard CCS, we assume the existence of an infinite set  $N$  of *names*, ranged over by  $a, b, c$ , with a bijection  $\bar{\cdot}: N \rightarrow \bar{N}$  denoting the *complement* of a name; names and complementary names respectively denote input and output ports for processes. We define the set of *labels*  $L$  as  $N \cup \bar{N} \cup \{\tau\}$ , where  $\tau$  denotes the interaction of concurrent processes.  $L$  is ranged over by  $\alpha$ , while  $L \setminus \{\tau\}$  is ranged over by  $\lambda$ .

To introduce reversibility, CCSK extends the syntax of CCS with a denumerable set  $K$  of *keys*, ranged over by  $k, m, n$ . Labels are paired with keys to define *keyed labels*, which are elements of the cartesian product  $L \times K$  and are represented as  $a[k], b[m]$ ; the set of keyed labels is also denoted as  $L_K$ .

*Processes* are defined as in the ordinary CCS, with the addition of keyed prefixes and without operators for recursion or replication:

$$\begin{array}{ll}
 X, Y ::= & \mathbf{0} \quad (\text{Inactive}) \quad | \quad \alpha.X \quad (\text{Prefix}) \\
 & | \quad \alpha[k].X \quad (\text{Keyed prefix}) \quad | \quad X + Y \quad (\text{Sum}) \\
 & | \quad X \mid Y \quad (\text{Parallel composition}) \quad | \quad X \setminus a \quad (\text{Restriction})
 \end{array}$$

The set of processes is denoted as  $\mathbb{X}$ . When preceded by a (keyed) prefix, the inactive process  $\mathbf{0}$  is usually omitted; the binding power of the operators, from highest to lowest, is  $\setminus a, \alpha[k], \alpha, |$  and  $+$ . In restrictions  $X \setminus a$ , the occurrences of the name  $a$  in  $X$  are said to be bound; all other occurrences of names and keys in processes are considered free. Processes that are  $\alpha$ -equivalent, i.e., that differ only in the choice of their bound names, will be identified. Unlike [3], restrictions only bind names and not complementary names: this choice does not rule out any significant process (since  $X \setminus a$  or  $X \setminus \bar{a}$  have the same behaviour) and leads to a clearer correspondence between processes and their encoding.

The set of keys occurring in a process is denoted as  $\text{keys}(X)$ . A process for which  $\text{keys}(X)$  is empty is said to be *standard*: in this case, we write that  $\text{std}(X)$  holds.

### 2.2 Semantics

The key feature of CCSK<sup>P</sup> is the notion of *proof keyed labels*:

$$\theta ::= v\alpha[k] \quad | \quad v\langle |_{L}v_1\lambda[k], |_{R}v_2\bar{\lambda}[k] \rangle$$

where  $v, v_1$  and  $v_2$  range over strings of symbols  $\{|_{L}, |_{R}, +_{L}, +_{R}\}$ . We denote the set of proof keyed labels as  $L_K^P$ , and refer to its elements simply as “proof labels” for brevity. The following functions  $\ell$  and  $\mathcal{K}$  map each proof label to its underlying label and key, respectively:

$$\ell(v\alpha[k]) = \alpha \quad \ell(v\langle |_{L}v_1\lambda[k], |_{R}v_2\bar{\lambda}[k] \rangle) = \tau \quad \mathcal{K}(v\alpha[k]) = k \quad \mathcal{K}(v\langle |_{L}v_1\lambda[k], |_{R}v_2\bar{\lambda}[k] \rangle) = k$$

Semantics is given by the *labelled transition system*  $(\mathbb{X}, L_K^P, \overset{\theta}{\rightarrow})$ , where  $\overset{\theta}{\rightarrow}$  denotes the union of the forward and backward transitions displayed in Fig. 2. We will refer to the union of forward and backward transitions as *combined* transitions. Given a transition  $X \overset{\theta}{\rightarrow} Y$ , the process  $X$  is said to be its *source*, while  $Y$  is said to be its *target*.

Prefix and Keyed Prefix	
<p>Forward</p> $\text{std}(X) \frac{}{\alpha.X \xrightarrow{\alpha[k]} \alpha[k].X} \text{pref}$ $\ell(\theta) \neq k \frac{X \xrightarrow{\theta} X'}{\alpha[k].X \xrightarrow{\theta} \alpha[k].X'} \text{kpref}$	<p>Backward</p> $\text{std}(X) \frac{}{\alpha[k].X \xrightarrow{\alpha[k]} \alpha.X} \text{pref}$ $\ell(\theta) \neq k \frac{X' \xrightarrow{\theta} X}{\alpha[k].X' \xrightarrow{\theta} \alpha[k].X} \text{kpref}$
Sum	
<p>Forward</p> $\text{std}(Y) \frac{X \xrightarrow{\theta} X'}{X + Y \xrightarrow{+\theta} X' + Y} +L$	<p>Backward</p> $\text{std}(Y) \frac{X' \xrightarrow{\theta} X}{X' + Y \xrightarrow{+\theta} X + Y} +L$
Parallel Composition	
<p>Forward</p> $\ell(\theta) \notin \text{keys}(Y) \frac{X \xrightarrow{\theta} X'}{X   Y \xrightarrow{ L\theta} X'   Y}  L$ $\frac{X \xrightarrow{v_L \lambda[k]} X' \quad Y \xrightarrow{v_R \bar{\lambda}[k]} Y'}{X   Y \xrightarrow{\langle  L v_L \lambda[k],  R v_R \bar{\lambda}[k] \rangle} X'   Y'} \text{syn}$	<p>Backward</p> $\ell(\theta) \notin \text{keys}(Y) \frac{X' \xrightarrow{\theta} X}{X'   Y \xrightarrow{ L\theta} X   Y}  L$ $\frac{X' \xrightarrow{v_L \lambda[k]} X \quad Y' \xrightarrow{v_R \bar{\lambda}[k]} Y}{X'   Y' \xrightarrow{\langle  L v_L \lambda[k],  R v_R \bar{\lambda}[k] \rangle} X   Y} \text{syn}$
Restriction	
<p>Forward</p> $\ell(\theta) \notin \{a, \bar{a}\} \frac{X \xrightarrow{\theta} X'}{X \setminus a \xrightarrow{\theta} X' \setminus a} \text{nu}$	<p>Backward</p> $\ell(\theta) \notin \{a, \bar{a}\} \frac{X' \xrightarrow{\theta} X}{X' \setminus a \xrightarrow{\theta} X \setminus a} \text{nu}$

Figure 2: Forward and backward transition rules for CCSK<sup>P</sup> (right rules for | and + omitted).

**Example 1** Consider a webpage that allows the user to interact via two independent buttons: one to toggle between light and dark mode, and another to switch between two different languages. The initial state of the system can be modeled as the parallel composition  $m | l$ , where the labels  $m$  and  $l$  represent the actions to switch the visual mode and the language, respectively.

The action of changing the visual mode can be represented by the following forward transition:  $m | l \xrightarrow{|L m[k]} m[k] | l$ . The target process preserves the label  $m$  and pairs it with a fresh key  $k$ . The proof label  $|L m[k]$  not only stores the label  $m$  and key  $k$  used in the transition, but also indicates that the action occurred on the left-hand side of a parallel composition.

Suppose the user now wishes to revert to the previous visual mode: pressing the mode button again can be interpreted as undoing the previously executed action. This can be modeled by the following backward transition, which removes the key associated with the earlier forward step:  $m[k] | l \xrightarrow{|L m[k]} m | l$ .

Two transitions are said to be *composable* if they can be performed consecutively – that is, the target of the first transition is the source of the second transition. A *path* is a (potentially empty) sequence of composable transitions and can be denoted as  $X \mapsto^* Y$ , where  $X$  is the source of the first transition (also called the *source* of the path) and  $Y$  is the target of the last transition (also called the *target* of the path); in

other words,  $\mapsto^*$  is the reflexive and transitive closure of  $\mapsto$ . A process  $X$  is *reachable* if there exists a path whose target is  $X$  and whose source is a standard process. This process, which can be proved to be unique (cf. Lemma B.13 in [2]), is called the *origin* of  $X$  and is denoted as  $O_X$ .

Reachability allows to rule out faulty processes which are syntactically well-formed, but whose particular selection of keys is inconsistent. For instance, this arises when the same key denotes successive actions, as in the process  $a[k].b[k]$ , or when keys internally form a cycle, as in the deadlocked process  $a[k].b[m] \mid \bar{b}[m].\bar{a}[k]$ , where neither action can be undone because of the presence of its associated key in the other thread. From this point on, each process will be assumed to be reachable.

The *loop lemma* (cf. Lemma 3.8 in [3]) is an important result characterizing reversible labelled transition systems. It states that any transition  $X \xrightarrow{\theta} Y$  can be reversed, yielding a transition  $Y \xrightarrow{\theta} X$ ; moreover, the reversing operator is an involution (i.e., reversing a transition twice returns the original transition). The validity of the loop lemma follows directly from the symmetry of the LTS (Labelled Transition System) rules presented in Fig. 2.

Connectivity Relation	
<b>Action</b> $\frac{}{\alpha[k] \Upsilon \theta} A^1 \quad \frac{}{\theta \Upsilon \alpha[k]} A^2$	<b>Choice</b> $\frac{\theta_1 \Upsilon \theta_2}{+_d \theta_1 \Upsilon +_d \theta_2} C_d^1 \quad \frac{}{+_d \theta_1 \Upsilon +_{\bar{d}} \theta_2} C_d^2$
<b>Parallel</b> $\frac{\theta_1 \Upsilon \theta_2}{ _d \theta_1 \Upsilon  _d \theta_2} P_d^1 \quad \frac{}{ _d \theta_1 \Upsilon  _{\bar{d}} \theta_2} P_d^2$	<b>Synchronization</b> $\frac{\theta \Upsilon \theta_d}{ _d \theta \Upsilon \langle  _L \theta_L,  _R \theta_R \rangle} S_d^1 \quad \frac{\theta_d \Upsilon \theta}{\langle  _L \theta_L,  _R \theta_R \rangle \Upsilon  _d \theta} S_d^2$ $\frac{\theta_1 \Upsilon \theta'_1 \quad \theta_2 \Upsilon \theta'_2}{\langle  _L \theta_1,  _R \theta_2 \rangle \Upsilon \langle  _L \theta'_1,  _R \theta'_2 \rangle} S^3$

Dependence Relation	Independence Relation
<b>Action</b> $\frac{}{\alpha[k] \times \theta} A^1 \quad \frac{}{\theta \times \alpha[k]} A^2$	<b>Action</b> <p style="text-align: center;"><i>(empty)</i></p>
<b>Choice</b> $\frac{\theta \times \theta'}{+_d \theta \times +_d \theta'} C_d^1 \quad \frac{}{+_d \theta \times +_{\bar{d}} \theta'} C_d^2$	<b>Choice</b> $\frac{\theta \iota \theta'}{+_d \theta \iota +_d \theta'} C_d^1$
<b>Parallel</b> $\frac{\theta \times \theta'}{ _d \theta \times  _d \theta'} P_d^1 \quad \frac{\mathcal{K}(\theta) = \mathcal{K}(\theta')}{ _d \theta \times  _{\bar{d}} \theta'} P_d^2$	<b>Parallel</b> $\frac{\theta \iota \theta'}{ _d \theta \iota  _d \theta'} P_d^1 \quad \frac{\mathcal{K}(\theta) \neq \mathcal{K}(\theta')}{ _d \theta \iota  _{\bar{d}} \theta'} P_d^2$
<b>Synchronization</b> $\frac{\theta \times \theta_d}{ _d \theta \times \langle  _L \theta_L,  _R \theta_R \rangle} S_d^1 \quad \frac{\theta_d \times \theta}{\langle  _L \theta_L,  _R \theta_R \rangle \times  _d \theta} S_d^2$ $\frac{\theta_i \times \theta'_i \quad \theta_j \Upsilon \theta'_j \quad i, j \in \{1, 2\}, i \neq j}{\langle  _L \theta_1,  _R \theta_2 \rangle \times \langle  _L \theta'_1,  _R \theta'_2 \rangle} S^3$	<b>Synchronization</b> $\frac{\theta \iota \theta_d}{ _d \theta \iota \langle  _L \theta_L,  _R \theta_R \rangle} S_d^1 \quad \frac{\theta_d \iota \theta}{\langle  _L \theta_L,  _R \theta_R \rangle \iota  _d \theta} S_d^2$ $\frac{\theta_1 \iota \theta'_1 \quad \theta_2 \iota \theta'_2}{\langle  _L \theta_1,  _R \theta_2 \rangle \iota \langle  _L \theta'_1,  _R \theta'_2 \rangle} S^3$

Figure 3: Causality relations on proof labels.

Finally, the binary relations of *connectivity*, *dependence* and *independence* on proof labels, respectively denoted as  $\Upsilon$ ,  $\times$  and  $\iota$ , are defined by the rules displayed in Fig. 3, where the label  $d$  ranges over  $\{L, R\}$  and  $\bar{d}$  denotes the opposite of  $d$  (i.e.,  $\bar{L} = R$  and  $\bar{R} = L$ ). Such relations will be referred to as *causality relations* for brevity. Compared to [3], the rule  $A^2$  for connectivity and dependence has been slightly modified, ensuring that each relation is symmetric and simplifying their encoding. This comes at the cost of losing uniqueness in derivations of judgements such as  $\theta_1 \Upsilon \theta_2$ ; however, this property has been shown not to be required for the purposes of our development.

**Example 2** The process  $m \mid l$ , introduced in Example 1 to model a webpage, can perform a forward transition labelled by  $|_L m[k_1]$ , representing the toggling of the visual mode. It can also perform a transition  $m \mid l \xrightarrow{|_R l[k_2]} m \mid l[k_2]$ , denoting the change of the language of the webpage. The two transitions are independent, as the order in which they are executed does not affect the resulting state. This is reflected in the independence of the two proof labels  $|_L m[k_1]$  and  $|_R l[k_2]$ , which follows from the  $P_L^2$  rule in Fig. 3.

Conversely, consider the process  $a.b \mid \bar{b}$ . It can perform a forward transition labelled by  $|_L a[k]$ , followed by another forward transition labelled by  $|_L b[n]$ . However, these transitions cannot be performed in reverse order, since the input action along  $b$  is only enabled after the input action along  $a$  has occurred; the two transitions are thus causally related. This is reflected in the dependence of the two proof labels  $|_L a[k]$  and  $|_L b[n]$ , which follows from the  $P_L^1$  and  $A^1$  rules in Fig. 3.

### 2.3 Properties of causality relations

We now turn to the theorems and lemmas object of our study. Their complete proof can be found in [2], the technical report accompanying [3]. The following theorem specifies the relationship between connectivity of transitions and connectivity of proof labels:

**Theorem 2.1 (cf. Proposition 4.4 in [3])**

- (i) If  $t_1 : X_1 \xrightarrow{\theta_1} X'_1$  and  $t_2 : X_2 \xrightarrow{\theta_2} X'_2$  are connected, then  $\theta_1 \Upsilon \theta_2$ .  $\square$
- (ii) If  $\theta_1 \Upsilon \theta_2$ , then there exist  $t_1 : X_1 \xrightarrow{\theta_1} X'_1$  and  $t_2 : X_2 \xrightarrow{\theta_2} X'_2$  such that  $t_1$  and  $t_2$  are connected.  $\square$

The proof of Theorem 2.1(i) relies on the fact that  $O_{X_1} = O_{X_2}$  and proceeds by induction over such origin process: recall that each process is assumed to be reachable and, therefore, has an origin. The equality of  $O_{X_1}$  and  $O_{X_2}$  follows from the two lemmas:

**Lemma 2.2** For all reachable processes  $X$  and  $Y$ , there exists a path  $X \mapsto^* Y$  iff  $O_X = O_Y$ .

**Lemma 2.3** If  $t_1 : X_1 \xrightarrow{\theta_1} X'_1$  and  $t_2 : X_2 \xrightarrow{\theta_2} X'_2$  are connected, then  $O_{X_1} = O_{X_2}$ .

Conversely, the proof of Theorem 2.1(ii) proceeds by structural induction over the given hypothesis  $\theta_1 \Upsilon \theta_2$  and relies on the following:

**Definition 2.4 (Realisation)** A process  $X$  realises the proof label  $\theta$  if there exist  $X_1$  and  $X_2$  such that  $X \mapsto^* X_1 \xrightarrow{\theta} X_2$ .  $\square$

**Lemma 2.5** For every proof label  $\theta$ , there exists a process that realises it, and we denote it  $r(\theta)$ .  $\square$

Next, the following theorem states the complementarity of the dependence and independence relations:

**Theorem 2.6 (cf. Theorem 4.9 in [3])**

For all  $\theta_1, \theta_2$ ,

- (i) If  $\theta_1 \iota \theta_2$  then  $\theta_1 \Upsilon \theta_2$ .  $\square$
- (ii) If  $\theta_1 \times \theta_2$  then  $\theta_1 \Upsilon \theta_2$ .  $\square$
- (iii) If  $\theta_1 \Upsilon \theta_2$  then either  $\theta_1 \iota \theta_2$  or  $\theta_1 \times \theta_2$ , but not both.  $\square$

This theorem is proved by induction over the structure of the given binary relation.

### 3 Beluga Formalization

In this section we outline the key points of the Beluga formalization of the notions presented in Section 2. Definitions and proofs omitted for brevity are hyperlinked to their encoding in the repository.

#### 3.1 Syntax

Beluga is structured in two layers: the LF (Logical Frameworks [13]) level, which is used to specify the formal system under study, and the computation level, which supports programming with LF data [21]. To encode the syntax of our system, only the former level is deployed. Names, keys, labels and processes are encoded using the LF types displayed in Fig. 4.

<code>LF names: type =;</code>	<code>LF proc: type =</code>	
<code>LF keys: type =</code>	null: proc	% 0
z: keys	pref: labels → proc → proc	% A.X
s: keys → keys;	kpref: labels → keys → proc → proc	% A[k].X
<code>LF labels: type =</code>	sum: proc → proc → proc	% X+Y
inp: names → labels	par: proc → proc → proc	% X Y
out: names → labels	nu: (names → proc) → proc;	% X\ a
tau: labels;		

Figure 4: Encoding of the syntax of CCSK<sup>P</sup>.

Since names in CCSK<sup>P</sup> are an infinite set without any additional assumption, they are represented by a type `names` without constructors; as explained in [8], this type will be dynamically inhabited by variables introduced through contexts. This is enabled by the following line of code:

```
schema ctx = names;
```

This line declares contexts made of a finite collection of distinct variables of type `names`, identified via the keyword `ctx`. Thanks to this setup, we can work with *contextual processes* of the form  $[g \vdash X]$ , i.e., processes  $X$  whose free names are drawn from the context  $g$ . Contextual objects live in the computation level.

Keys are by assumption denumerable, and the LTS rules for keyed prefixes require equality of keys to be decidable. Both conditions are satisfied by encoding keys explicitly as natural numbers.<sup>1</sup> An alternative approach would be to rely on contexts, as is done for names: however, this would require managing mixed contexts of names and keys, and having a more complex encoding of transitions and paths.

Restrictions  $X \setminus a$  are represented by terms of the form  $(\text{nu } \setminus a. (X \ a))$ , where  $\setminus x. (f \ x)$  is Beluga's notation for functions  $f$  mapping  $x$  to  $f(x)$ : following the higher-order abstract syntax (HOAS) paradigm, the bound name  $a$  is represented as the implicit argument of a meta-language function  $\setminus a. (X \ a)$  from `names` to `proc`. In this way, we leverage the meta-language implementation of binders to achieve  $\alpha$ -renaming and capture-avoiding substitutions for free.

An important but often overlooked aspect of formalizations is the *adequacy* of the encoding: the encoding must constitute a faithful representation of the original system into study [9]. Adequacy is generally established by proving the existence of a compositional bijection between the mathematical model and its formalized counterpart. The discussion of the adequacy of our encoding is omitted due to space constraints.

<sup>1</sup>Note that properties such as decidability of equality must be stated and proved manually, as Beluga does not provide a built-in library of properties of natural numbers.

### 3.2 Semantics

Proof labels are encoded by the type `pr_lab` in Fig. 5. Rather than directly modeling the informal definition of proof labels, by defining strings over the symbols  $\{|L, |R, +L, +R\}$  as lists, we are introducing four constructors (`pr_suml`, `pr_sumr`, etc.) that build proof labels incrementally by appending one symbol at a time. This provides a stronger induction principle and streamlines the encoding of LTS rules and subsequent proofs.

```

LF pr_lab: type =
| pr_base: labels → keys → pr_lab
| pr_suml: pr_lab → pr_lab
| pr_sumr: pr_lab → pr_lab
| pr_parl: pr_lab → pr_lab
| pr_parr: pr_lab → pr_lab
| pr_sync: pr_lab → pr_lab → pr_lab;
LF valid: pr_lab → type =
| v_base: valid (pr_base A K)
| v_suml: valid T → valid (pr_suml T)
| v_sumr: valid T → valid (pr_sumr T)
| v_parl: valid T → valid (pr_parl T)
| v_parr: valid T → valid (pr_parr T)
| v_sync: valid T1 → valid T2 → lab T1 (inp A) → key T1 K
  → lab T2 (out A) → key T2 K → valid (pr_sync T1 T2)
| v_syncr: valid T1 → valid T2 → lab T1 (out A) → key T1 K
  → lab T2 (inp A) → key T2 K → valid (pr_sync T1 T2);
LF fstep: proc → pr_lab → proc → type =
| fs_pref: std X → fstep (pref A X) (pr_base A K) (kpref A K X)
| fs_kpref: fstep X T X' → key T M → neq K M
  → fstep (kpref A K X) T (kpref A K X')
| fs_suml: fstep X T X' → std Y → fstep (sum X Y) (pr_suml T) (sum X' Y)
| fs_sumr: fstep Y T Y' → std X → fstep (sum X Y) (pr_sumr T) (sum X Y')
| fs_parl: fstep X T X' → key T K → notin K Y
  → fstep (par X Y) (pr_parl T) (par X' Y)
| fs_parr: fstep Y T Y' → key T K → notin K X
  → fstep (par X Y) (pr_parr T) (par X Y')
| fs_sync: fstep X T1 X' → lab T1 (inp L) → key T1 K
  → fstep Y T2 Y' → lab T2 (out L) → key T2 K
  → fstep (par X Y) (pr_sync T1 T2) (par X' Y')
| fs_syncr: fstep X T1 X' → lab T1 (out L) → key T1 K
  → fstep Y T2 Y' → lab T2 (inp L) → key T2 K
  → fstep (par X Y) (pr_sync T1 T2) (par X' Y')
| fs_nu: ({a:names} fstep (X a) T (X' a)) → fstep (nu X) T (nu X');
LF lab: pr_lab → labels → type =
| lab_base: lab (pr_base A K) A
| lab_suml: lab T A → lab (pr_suml T) A
| lab_sumr: lab T A → lab (pr_sumr T) A
| lab_parl: lab T A → lab (pr_parl T) A
| lab_parr: lab T A → lab (pr_parr T) A
| lab_sync: lab (pr_sync T1 T2) tau;

```

Figure 5: Encoding of the semantics of  $CCSK^P$ .

In Beluga, predicates are encoded as type families, i.e., types parametrized by arguments: a predicate  $P(x_1, \dots, x_n)$  holds iff the corresponding type  $(P \ x_1 \ \dots \ x_n)$  is not empty. Type families are also used to encode functions, identified with their graph, as in the case of the functions  $\ell$  and  $\#$  returning the label and key of a proof label: the former is encoded by the type family `lab` in Fig. 5, while the latter is encoded by the type family `key`, here omitted for brevity. For example, given a proof label  $\theta$  and a label  $\alpha$ , represented as  $T$  and  $A$  in the encoding, the type `lab T A` is inhabited iff  $\ell(\theta) = \alpha$ .

Our encoding of proof labels pays the price of being over-expressive: the constructor `pr_sync` accepts any two proof labels regardless of their key or label, generating terms that fall outside the original definition. For example, the term `(pr_sync (pr_base A K) (pr_base B M))` has type `pr_lab` for any labels A, B and keys K, M, while its counterpart  $\langle |_{L\alpha[k]}, |_{R\beta[m]} \rangle$  is well-defined only if  $\beta = \bar{\alpha}$  and  $k = m$ . While this is harmless in most of our development, since such spurious terms do not label any actual transition, it becomes an issue when proving theorems universally quantified on proof labels, such as Lemma 2.5. To address this problem, we introduce an additional predicate `valid`, displayed in Fig. 5, which filters out the spurious terms. It can be proved the existence of a bijection between proof labels and the set of elements T of type `pr_lab` for which `valid T` holds. From this point forward, we will refer to such terms as *valid* proof labels.

Forward and backward LTS rules are defined through the type families `fstep` and `bstep` in Fig. 5 (with the latter omitted here for brevity). These rules rely on the additional type families `std`, `notin` and `neq`, hyperlinked to their formalization in the repository, which respectively hold when a process is standard, when a key does not occur in a process, and when two keys are not equal. The parameters X and X' in the `fs_nu` rule are functions from `names` to `proc`, whose arguments represent the binders of the restrictions. The universal quantification  $\{a:\text{names}\}$  is used to abstract over the particular choice of the binder; moreover, *a* or  $\bar{a}$  does not occur in the proof label T, since such parameter does not depend on *a* within the body of the universal quantification.

Combined transitions, paths, reachable processes and connected transitions are defined as follows:

```

LF step: proc → pr_lab → proc → type =
  | fw: fstep X T X' → step X T X'
  | bw: bstep X' T X → step X' T X;
LF step*: proc → proc → type =
  | id_s*: step* X X
  | st_s*: step X T Y → step* X Y
  | tr_s*: step* X Y → step* Y Z → step* X Z;
LF reachable: proc → type =
  | rch: std X → step* X Y → reachable Y;
LF conn_tr: step X T1 X' → step Y T2 Y' → type =
  | ct: {S1:step X T1 X'}{S2:step Y T2 Y'} step* X Y' → conn_tr S1 S2;

```

Paths, or multi-step transitions, can be encoded equivalently using only two constructors; in this development, the more verbose version has been adopted as it simplified the proof search.

Finally, the relations of connectivity, dependence and independence are encoded through three type families `conn`, `dep` and `indep`. The former is displayed in Fig. 6. While some of the rules in Fig. 3 are grouped together, by using the label d in the place of L and R, the encoding requires each rule to be stated separately, with its own constructor.

### 3.2.1 Basic properties of keys, proof labels and transitions

Before diving into the theorems related to the causality relations, our encoding requires a small library of properties of keys, proof labels and transitions: these include the *decidability of equality of keys*, the *functionality of  $\ell$  and  $\ell$* , the fact that *standard processes have no keys*, or the *loop lemma*. To provide an overview of how proofs are carried out in Beluga, we will walk through the proof of the following result: “for all proof labels  $\theta$ , there exists a label  $\alpha$  such that  $\ell(\theta) = \alpha$ ”. Its code is displayed in Fig. 7:

```

LF conn: pr_lab → pr_lab → type =
| c_a1: conn (pr_base A K) T
| c_a2: conn T (pr_base A K)
| c_c1l: conn T1 T2 → conn (pr_suml T1) (pr_suml T2)
| c_c1r: conn T1 T2 → conn (pr_sumr T1) (pr_sumr T2)
| c_c2l: conn (pr_suml T1) (pr_sumr T2)
| c_c2r: conn (pr_sumr T1) (pr_suml T2)
| c_p1l: conn T1 T2 → conn (pr_parl T1) (pr_parl T2)
| c_p1r: conn T1 T2 → conn (pr_parr T1) (pr_parr T2)
| c_p2l: conn (pr_parl T1) (pr_parr T2)
| c_p2r: conn (pr_parr T1) (pr_parl T2)
| c_s1l: conn T TL → conn (pr_parl T) (pr_sync TL TR)
| c_s1r: conn T TR → conn (pr_parr T) (pr_sync TL TR)
| c_s2l: conn TL T → conn (pr_sync TL TR) (pr_parl T)
| c_s2r: conn TR T → conn (pr_sync TL TR) (pr_parr T)
| c_s3: conn T1 T1' → conn T2 T2' → conn (pr_sync T1 T2) (pr_sync T1' T2');

```

Figure 6: Encoding of the connectivity relation on proof labels.

```

LF ex_lab: pr_lab → type =
| ex_l: lab T A → ex_lab T;
rec existence_of_lab: (g:ctx) {T:[g ⊢ pr_lab]} [g ⊢ ex_lab T] =
/ total t (existence_of_lab _ t) /
mlam T ⇒ case [_ ⊢ T] of
| [g ⊢ pr_base _ _] ⇒ [g ⊢ ex_l lab_base]
| [g ⊢ pr_suml T'] ⇒ let [g ⊢ ex_l L] = existence_of_lab [g ⊢ T'] in
  [g ⊢ ex_l (lab_suml L)]
| [g ⊢ pr_sumr T'] ⇒ let [g ⊢ ex_l L] = existence_of_lab [g ⊢ T'] in
  [g ⊢ ex_l (lab_sumr L)]
| [g ⊢ pr_parl T'] ⇒ let [g ⊢ ex_l L] = existence_of_lab [g ⊢ T'] in
  [g ⊢ ex_l (lab_parl L)]
| [g ⊢ pr_parr T'] ⇒ let [g ⊢ ex_l L] = existence_of_lab [g ⊢ T'] in
  [g ⊢ ex_l (lab_parr L)]
| [g ⊢ pr_sync _ _] ⇒ [g ⊢ ex_l lab_sync];

```

Figure 7: Proof of the existence of a label in a proof label.

The first two lines of code in Fig. 7 introduce a type family `ex_lab`, which captures the conclusions of the lemma to be proved: the type `ex_lab T` is inhabited whenever there exists a label `A` for which `lab T A` holds. Defining such additional type families is the standard workaround to the lack of syntactic sugar for existentials and conjunctions in Beluga.

Thanks to the Curry-Howard isomorphism, proofs by induction are encoded through recursive functions. In Beluga, these are computation-level entities introduced by the keyword `rec`. The function `existence_of_lab` takes as input a context `g` of schema `ctx` and a contextual object `T` of type `pr_lab` and returns an object of type `ex_lab T`. The second line of the proof asserts that the built function is total and is recursive on the second argument. These conditions are verified by Beluga's totality checker and guarantee that the function constitutes a valid proof.

The proof itself begins by introducing the argument `T` through the keyword `mlam`; the other argument, the context `g`, is implicit due to the use of round brackets in the function declaration. The proof proceeds by

pattern matching on the object  $T$ , which by the Curry-Howard isomorphism corresponds to case analysis on the structure of  $T$  in the informal proof. Underscores are used to omit parameters that Beluga can infer automatically. The `pr_base` and `pr_sync` cases are handled immediately, since we can already provide the required object of type `ex_lab T`; for the remaining four cases, the proof proceeds by recursively applying the function `existence_of_lab` to a subterm  $T'$  of  $T$ , and then using the result to build the desired object. Recursive calls on structurally smaller objects correspond to applications of the inductive hypothesis in the informal proof.

We conclude this subsection by addressing the *symmetry* and *irreflexivity* of the causality relations. We report the signature of the function which proves that connectivity is symmetric:

```
rec symmetric_conn: (g:ctx) [g ⊢ conn T1 T2] → [g ⊢ conn T2 T1] = ...
```

These lemmas are proved by straightforward inductions on the structure of the given predicate.

### 3.3 Properties of causality relations

Although the theorems in Section 2.3 are presented in a different order, here we start with the encoding of Theorem 2.6, as it is straightforward and mirrors the structure of the informal proof.

#### 3.3.1 Encoding of Theorem 2.6

The three statements of the theorem are addressed by four recursive functions: this is because Theorem 2.6(iii) actually consists of two separate assertions, which here we prove separately. Moreover, the disjunction in the conclusions requires defining an additional type family `dep_or_indep`. Fig. 8 displays the proof of the final assertion: “two proof labels cannot be both dependent and independent”. We also present the signatures of the other recursive functions below.

```
rec indep_impl_conn: (g:ctx) [g ⊢ indep T1 T2] → [g ⊢ conn T1 T2] = ...
rec dep_impl_conn: (g:ctx) [g ⊢ dep T1 T2] → [g ⊢ conn T1 T2] = ...
LF dep_or_indep: pr_lab → pr_lab → type =
  | or_dep: dep T1 T2 → dep_or_indep T1 T2
  | or_ind: indep T1 T2 → dep_or_indep T1 T2;
rec conn_impl_dep_or_indep: (g:ctx) [g ⊢ conn T1 T2] → [g ⊢ dep_or_indep T1 T2] = ...
```

The proof in Fig. 8 is an example of proof by contradiction: given two objects of type `dep T1 T2` and `indep T1 T2`, the function `impossible_dep_and_indep` aims to derive an object of the empty type `false`, thereby establishing a contradiction. After introducing the arguments  $d$  and  $i$ , the proof proceeds by pattern matching on  $d$ ; depending on the case, the contradiction is reached in one of three distinct ways.

In case  $d$  is built, e.g., through the constructor `d_a1` (corresponding to the case  $A^1: \alpha[k] \times \theta$  in the informal proof), it is immediately clear that an object  $i$  of type `indep T1 T2` (i.e.,  $\alpha[k] \iota \theta$ ) does not exist: this contradiction is exhibited through the keyword `impossible`. In other subcases, such as when  $d$  is built via `d_c11` (corresponding to  $C_L^1: +_L \theta \times +_L \theta'$ , given  $\theta \times \theta'$ ), the contradiction is obtained by recursively invoking `impossible_dep_and_indep` on smaller arguments. Finally, in the `d_p21` subcase ( $|_L \theta \times |_R \theta'$ , under the assumption  $\#(\theta) = \#(\theta')$ ), we first examine the structure of  $i$  and find that it must have been constructed using `i_p21`. This gives us an object  $N$  witnessing the inequality  $\#(\theta) \neq \#(\theta')$ , which clearly contradicts our assumption; however, to complete the proof in Beluga, it is first necessary to apply the auxiliary function `uniqueness_of_key` for some variable unification, yielding  $\#(\theta) \neq \#(\theta)$ , followed by the function `irreflexive_neq`, which states the irreflexivity of the inequality of keys.

```

rec impossible_dep_and_indep: (g:ctx) [g ⊢ dep T1 T2] → [g ⊢ indep T1 T2]
  → [g ⊢ false] =
/ total d (impossible_dep_and_indep _ _ _ d _) /
fn d,i ⇒ case d of
| [g ⊢ d_a1] ⇒ impossible i
| [g ⊢ d_a2] ⇒ impossible i
| [g ⊢ d_c1l D] ⇒ let [g ⊢ i_c1l I] = i in impossible_dep_and_indep [g ⊢ D] [g ⊢ I]
| [g ⊢ d_c1r D] ⇒ let [g ⊢ i_c1r I] = i in impossible_dep_and_indep [g ⊢ D] [g ⊢ I]
| [g ⊢ d_c2l] ⇒ impossible i
| [g ⊢ d_c2r] ⇒ impossible i
| [g ⊢ d_p1l D] ⇒ let [g ⊢ i_p1l I] = i in impossible_dep_and_indep [g ⊢ D] [g ⊢ I]
| [g ⊢ d_p1r D] ⇒ let [g ⊢ i_p1r I] = i in impossible_dep_and_indep [g ⊢ D] [g ⊢ I]
| [g ⊢ d_p2l H1 H2] ⇒ let [g ⊢ i_p2l H1' H2' N] = i in
  let [g ⊢ refk] = uniqueness_of_key [g ⊢ H1] [g ⊢ H1'] in
  let [g ⊢ refk] = uniqueness_of_key [g ⊢ H2] [g ⊢ H2'] in irreflexive_neq [g ⊢ N]
| [g ⊢ d_p2r H1 H2] ⇒ let [g ⊢ i_p2r H1' H2' N] = i in
  let [g ⊢ refk] = uniqueness_of_key [g ⊢ H1] [g ⊢ H1'] in
  let [g ⊢ refk] = uniqueness_of_key [g ⊢ H2] [g ⊢ H2'] in irreflexive_neq [g ⊢ N]
| [g ⊢ d_s1l D] ⇒ let [g ⊢ i_s1l I] = i in impossible_dep_and_indep [g ⊢ D] [g ⊢ I]
| [g ⊢ d_s1r D] ⇒ let [g ⊢ i_s1r I] = i in impossible_dep_and_indep [g ⊢ D] [g ⊢ I]
| [g ⊢ d_s2l D] ⇒ let [g ⊢ i_s2l I] = i in impossible_dep_and_indep [g ⊢ D] [g ⊢ I]
| [g ⊢ d_s2r D] ⇒ let [g ⊢ i_s2r I] = i in impossible_dep_and_indep [g ⊢ D] [g ⊢ I]
| [g ⊢ d_s3l D1 _] ⇒ let [g ⊢ i_s3 I1 _] = i in
  impossible_dep_and_indep [g ⊢ D1] [g ⊢ I1]
| [g ⊢ d_s3r _ D2] ⇒ let [g ⊢ i_s3 _ I2] = i in
  impossible_dep_and_indep [g ⊢ D2] [g ⊢ I2];

```

Figure 8: Proof of the statement “two proof labels cannot be both dependent and independent”.

### 3.3.2 Encoding of Theorem 2.1

Recall that, throughout our development, each process is assumed to be reachable. Although this hypothesis is not explicitly stated in theorems and lemmas, it is in fact essential for proving Theorem 2.1 and some of its auxiliary lemmas. For this reason, before outlining its encoding, we refine its statement, making the reachability assumption explicit:

#### Theorem 2.1 (Refined)

- (i) If  $t_1 : X_1 \xrightarrow{\theta_1} X'_1$  and  $t_2 : X_2 \xrightarrow{\theta_2} X'_2$  are connected and  $X_1$  is reachable,<sup>2</sup> then  $\theta_1 \curlywedge \theta_2$ .
- (ii) If  $\theta_1 \curlywedge \theta_2$ , then there exist  $t_1 : X_1 \xrightarrow{\theta_1} X'_1$  and  $t_2 : X_2 \xrightarrow{\theta_2} X'_2$ , with  $X_1$  reachable, such that  $t_1$  and  $t_2$  are connected.

Since the two statements are encoded by two distinct functions, we discuss them separately. The proof of Theorem 2.1(i) is given by the following function `conn_rel_one`:

```

rec conn_rel_one: (g:ctx) {S1:[g ⊢ step X1 T1 X1']} {S2:[g ⊢ step X2 T2 X2']}
  [g ⊢ reachable X1] → [g ⊢ conn_tr S1 S2] → [g ⊢ conn T1 T2] = ...

```

<sup>2</sup>The reachability of the only  $X_1$  is enough to deduce the reachability of any other process in the statement, given the existence of a path from  $X_1$  to such processes.

The function takes as inputs two transitions S1 and S2, whose typing judgments introduce the names of each involved parameter, such as the process X1; these are followed by the further assumptions of reachability of X1 and connectivity of S1 and S2. The function returns a derivation of the connectivity of the proof labels T1 and T2.

Although our encoding may appear different – and somewhat longer – than the proof presented in [2], it is, in essence, faithful to the same underlying structure. The original proof leverages Lemma 2.3 to establish the equality of the processes  $O_{X_1}$  and  $O_{X_2}$ , the origins of the sources of the connected transitions  $t_1 : X_1 \xrightarrow{\theta_1} X'_1$  and  $t_2 : X_2 \xrightarrow{\theta_2} X'_2$ . It proceeds by induction on  $O_{X_1}$ , observing that its structure determines that of the processes and transitions in the same environment (e.g., if the outermost operator of  $O_{X_1}$  is a sum, the same applies to  $X_1$ ). The proof then concludes either directly or by applying the induction hypothesis to transitions involving specific subprocesses.

Below, we outline the changes and technical considerations brought by our encoding of this argument:

- The formalized proof proceeds by pattern matching on an object D of type `std OX1`, rather than directly on the process OX1. This is essentially equivalent, since the type family `std proc`, which asserts that a process is standard, is itself defined by pattern matching on the underlying process.
- It is not necessary to encode Lemma 2.3. The reachability of  $X_1$ , together with the existence of a path from  $X_1$  to  $X_2$ , provides us a path between  $O_{X_1}$  and  $X_2$ ; this path is enough to determine the structure of  $X_2$ , known the structure of  $O_{X_1}$ .
- The proof requires analyzing the structure of the given transitions S1 and S2. Since combined transitions are either forward or backward, and each have their own constructors, this results in four levels of nested pattern matching. While most of the subcases can be unified in the informal proof, Beluga requires them to be treated separately: this is the primary reason for the proof’s length. To improve efficiency, certain assertions have been moved earlier in the proof tree compared to their position in the informal version.
- The informal proof takes for granted structural properties such as: “given a path whose source is a sum process, the target is also a sum process”, or “given a path between two sum processes, there exists a path between their left addends”. In the encoding, these results must be explicitly stated and proved, resulting in 16 additional lemmas. Some of these require classical techniques such as mutual recursion or strengthening of contextual judgments, which are described in [8]. We report the signatures of two of these functions:

```
% Type family encoding sum processes
LF is_sum: proc → type =
  | sm: is_sum (sum X Y);
% A path starting from a sum process ends in a sum process
rec step*_from_sum: (g:ctx) [g ⊢ step* (sum X Y) Z] → [g ⊢ is_sum Z] = ...
% Given a path between sum processes, there is a path between their left addends
rec step*_betw_sums_left: (g:ctx) [g ⊢ step* (sum X1 X2) (sum Y1 Y2)]
  → [g ⊢ step* X1 Y1] = ...
```

The formalization of Theorem 2.1(ii) requires encoding Lemma 2.5, which states that every proof label  $\theta$  is realised by some process  $r(\theta)$  – that is, there exist processes  $X_1$ ,  $X_2$  and  $r(\theta)$  such that  $r(\theta) \mapsto^* X_1 \xrightarrow{\theta} X_2$ . The original proof in [2], however, goes further: it builds a process  $r(\theta)$  which is standard and directly performs a single forward transition  $r(\theta) \xrightarrow{\theta} X_2$  (in other words,  $r(\theta)$  and  $X_1$  coincide). Our encoding reflects this stronger formulation by specializing the original Definition 2.4 with the following type family `realised`:

```

LF realised: pr_lab → type =
  | rl: std X → fstep X T X' → realised T;

```

For any proof label  $T$ , `realised T` is non empty iff `std X` and `fstep X T X'` hold for some  $X$  and  $X'$ . The following recursive function `pr_lab_is_realised` encodes the proof of Lemma 2.5:

```

rec pr_lab_is_realised: (g:ctx) [g ⊢ valid T] → [g ⊢ realised T] = ...

```

The proof is a straightforward induction on the structure of the assumption `valid T`.

Other than relying on Lemma 2.5, the proof of Theorem 2.1(ii) in [2] assumes auxiliary results such as the following: “if  $O_X$  realises  $X$  and  $O_Y$  realises  $Y$ , then  $O_X \mid O_Y$  realises  $X \mid Y$ ”. While this result holds in the particular context of Theorem 2.1(ii), where  $X \mid Y$  is known to be reachable and is able to perform a synchronization, it does not hold in general. For instance, consider  $X_1 = a[k]$  and  $X_2 = b[k]$ : the parallel composition  $a[k] \mid b[k]$  is not reachable from  $a \mid b$ . Moreover, even when such conditions are met, building a constructive proof is far from straightforward. These issues led us to revisit the entire argument and develop the following proof strategy for Theorem 2.1(ii):

1. First, we consider the case where neither  $\ell(\theta_1)$  nor  $\ell(\theta_2)$  is  $\tau$  and prove that the diagram in Fig. 9a holds. The hypothesis excludes the cases in which  $\theta_1$  and  $\theta_2$  label synchronizations, thus ruling out the scenarios in which the aforementioned auxiliary lemma occurs. The proved result goes beyond establishing the connectivity of two combined transitions labelled by  $\theta_1$  and  $\theta_2$ : both transitions are forward, and the processes  $X_1$  and  $X'_2$  are either identical or connected by a single combined transition. Additionally, we show that at least one among  $X_1$  and  $X'_2$  is standard.
2. We then move to the general case, proving that for any connected pair of proof labels  $\theta_1$  and  $\theta_2$  the diagram in Fig. 9b holds. Analogously to the previous point, the transitions labelled by  $\theta_1$  and  $\theta_2$  are forward, meaning that our statement is slightly more specific than the original formulation of Theorem 2.1(ii). This refinement helps eliminating non-existent subcases that would arise in the nested pattern matching of combined transitions.




Figure 9: Proof strategy for Theorem 2.1(ii).

In the general case, when  $\theta_1$  and  $\theta_2$  label synchronizations (e.g., when  $\theta_1 = \langle |_{L}\theta_L^1, |_{R}\theta_R^1 \rangle$ ), the labels of their subterms (e.g.,  $\theta_L^1$  and  $\theta_R^1$ ) are not  $\tau$ : this detail allows us to apply the base case of the proof strategy, which provides richer information than the inductive hypothesis of the general Theorem 2.1(ii). That additional information is essential: it enables us to build the desired path between  $X_1$  and  $X'_2$ , actually with at most two transition steps.

The encoding of Theorem 2.1(ii) follows the plan outlined. The base case makes use of the following elements: a type family `lab_not_tau`, characterizing proof keyed labels whose label is not  $\tau$ ;

a type family `max_one_step`, encoding the conclusions of the statement; and the recursive function `conn_rel_two_base`, which proves it.

```

LF lab_not_tau: pr_lab → type =
  | nt_inp: lab T (inp _) → lab_not_tau T
  | nt_out: lab T (out _) → lab_not_tau T;
LF max_one_step: pr_lab → pr_lab → type =
  | c_id: std X1 → fstep X1 T1 X1' → fstep X1 T2 X2 → max_one_step T1 T2
  | c_fw: std X1 → fstep X1 T1 X1' → fstep X1 T3 X2' → fstep X2' T2 X2
    → lab T1 L1 → lab T3 L1 → max_one_step T1 T2
  | c_bw: std X2' → fstep X1 T1 X1' → bstep X1 T3 X2' → fstep X2' T2 X2
    → lab T2 L2 → lab T3 L2 → max_one_step T1 T2;
rec conn_rel_two_base: (g:ctx) [g ⊢ valid T1] → [g ⊢ valid T2] → [g ⊢ conn T1 T2] →
  [g ⊢ lab_not_tau T1] → [g ⊢ lab_not_tau T2] → [g ⊢ max_one_step T1 T2] = ...

```

The proof of this result is given by a long induction on the structure of the given connectivity relation. The predicates  $(\text{lab } T_i \text{ } L_j)$ , for  $i, j$  in  $\{1, 2, 3\}$ , which occur in the type family `max_one_step`, are a technical detail which helps completing few subcases of the proof.

Next, the general case of the proof is addressed by the recursive function `conn_rel_two_fstep` below, which relies on a dedicated type family as well:

```

LF ex_conn_fstep: pr_lab → pr_lab → type =
  | ex_cf: fstep X1 T1 X1' → fstep X2' T2 X2 → step* X1 X2'
    → reachable X1 → ex_conn_fstep T1 T2;
rec conn_rel_two_fstep: (g:ctx) [g ⊢ valid T1] → [g ⊢ valid T2] → [g ⊢ conn T1 T2]
  → [g ⊢ ex_conn_fstep T1 T2] = ...

```

The proof is given by a long induction on the structure of the given connectivity relation. It requires encoding auxiliary lemmas such as the following: “given a path between two processes  $X$  and  $X'$ , there is a path between  $X + \mathbf{0}$  and  $X' + \mathbf{0}$ ”, or: “given a forward transition  $X \xrightarrow{\theta} X'$  where  $X$  is standard and  $\#(\theta) = k$ , then any key  $m \neq k$  does not occur in  $X'$ ”.

Finally, Theorem 2.1(ii) is encoded by the following function `conn_rel_two`. It calls the function `conn_rel_two_fstep`, applies the loop lemma to reverse one of the two forward transitions, and has all the ingredients to conclude:

```

LF ex_conn_tr: pr_lab → pr_lab → type =
  | ex_c: {S1: step X T1 X'} {S2: step Y T2 Y'} reachable X → conn_tr S1 S2
    → ex_conn_tr T1 T2;
rec conn_rel_two: (g:ctx) [g ⊢ valid T1] → [g ⊢ valid T2] → [g ⊢ conn T1 T2]
  → [g ⊢ ex_conn_tr T1 T2] =
/ total c (conn_rel_two _ _ _ _ c) /
fn v1,v2,c ⇒
let [g ⊢ ex_cf F1 F2 S* (rch D S0*)] = conn_rel_two_fstep v1 v2 c in
let [g ⊢ B2] = loop_lemma_one [g ⊢ F2] in
[g ⊢ ex_c (fw F1) (bw B2) (rch D S0*) (ct (fw F1) (bw B2) S*)];

```

## 4 Conclusions and Future Work

We begin with a brief technical overview of the encoding. The complete formalization consists of less than 2000 lines of code and includes a total of 49 theorems and lemmas. Among them, 13 are direct

translations of results stated in Section 2, while the remaining 36 are technical and auxiliary lemmas introduced to support the encoding.

Beluga has proved to be a reliable and expressive proof assistant, well-suited to represent the definitions and properties of  $\text{CCSK}^P$ . Its use of higher-order abstract syntax (HOAS) offers a convenient approach to handling restrictions – even though  $\text{CCSK}^P$  does not feature a particularly complex binding structure, unlike, for instance, the  $\pi$ -calculus. Furthermore, Beluga’s explicit proof style provides a transparency that is often lost in proof assistants that rely heavily on automation.

However, the lack of automation also comes with drawbacks, mainly the increased length of proof terms. This also follows from the lack of syntactic sugar for existentials, conjunctions and disjunctions, which leads to defining additional type families or splitting theorem statements. Additionally, Beluga provides no built-in mechanism to simplify repeated proof patterns, requiring each similar subcase to be handled individually.

Whether the overall outcome is favorable depends largely on the specific system one aims to formalize. For languages with rich binding structures, the benefits of HOAS alone may outweigh the trade-offs. In our case, however, this advantage is less significant, and we believe that other proof assistants (such as Rocq [5]) might be a better fit for formalizing the system at hand.

To the best of our knowledge, this work provides the first formalization of a reversible concurrent calculus in a proof assistant. We have formally verified the correctness of the notions and results presented in Section 2. We gained a deeper understanding of the system itself, leading to refinements in both definitions and proofs; in particular, we provided an alternative way to represent proof labels compared to the informal definition.

This work lays the foundations for future reversible concurrent calculi formalizations. The encoding can be adapted to cover the subsystems of  $\text{CCSK}^P$ , i.e., CCS and CCSK, and can be mapped to existing CCS formalizations. Moreover, it could be extended to include additional portions of [3]. Additionally, it could be translated into other proof assistants, such as Rocq, which are potentially better suited for representing this reversible process calculus. Finally, it can serve as a reference point for future formalizations of other reversible concurrent calculi, such as RCCS.

## Acknowledgments

We warmly thank the anonymous reviewers, as well as Clément Aubert and Deivid Vale, for their insightful comments and suggestions, which greatly contributed to improving the paper. This work is supported by the National Science Foundation under Grant No. 2242786 (SHF:Small:Concurrency In Reversible Computations).

## References

- [1] Clément Aubert & Peter Browning (2023): *Implementation of a Reversible Distributed Calculus*. In Martin Kutrib & Uwe Meyer, editors: *Reversible Computation - 15th International Conference, RC 2023, Giessen, Germany, July 18-19, 2023, Proceedings, Lecture Notes in Computer Science 13960*, Springer, pp. 210–217, doi:10.1007/978-3-031-38100-3\_13.
- [2] Clément Aubert, Iain Phillips & Irek Ulidowski (2024): *Dependence and Independence for Reversible Process Calculi*. CoRR abs/2410.14699, doi:10.48550/ARXIV.2410.14699.
- [3] Clément Aubert, Iain Phillips & Irek Ulidowski (2025): *Independence and Causality in the Reversible Concurrent Setting*. In Robert Glück & Robin Kaarsgaard, editors: *Reversible Computation - 17th International*

- Conference, RC 2025, Odense, Denmark, July 3-4, 2025, Proceedings, Lecture Notes in Computer Science 15716, Springer, pp. 9–26, doi:10.1007/978-3-031-97063-4\_2.
- [4] Clément Aubert (2022): *Concurrencies in Reversible Concurrent Calculi*. In Claudio Antares Mezzina & Krzysztof Podlaski, editors: *Reversible Computation - 14th International Conference, RC 2022, Urbino, Italy, July 5-6, 2022, Proceedings, LNCS 13354*, Springer, pp. 146–163, doi:10.1007/978-3-031-09005-9\_10.
- [5] Yves Bertot & Pierre Castéran (2004): *Interactive Theorem Proving and Program Development - Coq'Art: The Calculus of Inductive Constructions*. Texts in Theoretical Computer Science. An EATCS Series, Springer, doi:10.1007/978-3-662-07964-5.
- [6] Marco Carbone, David Castro-Perez, Francisco Ferreira, Lorenzo Gheri, Frederik Krogsdal Jacobsen, Alberto Momigliano, Luca Padovani, Alceste Scalas, Dawit Tiore, Martin Vassor, Nobuko Yoshida & Daniel Zackon (2024): *The Concurrent Calculi Formalisation Benchmark*. In Ilaria Castellani & Francesco Tiezzi, editors: *Coordination Models and Languages*, Springer Nature Switzerland, Cham, pp. 149–158, doi:10.1007/978-3-031-62697-5\_9.
- [7] Gabriele Cecilia (2025): *A Formalization of the Reversible Concurrent Calculus CCSK<sup>P</sup> in Beluga (artifact)*, doi:10.5281/zenodo.16179366.
- [8] Gabriele Cecilia & Alberto Momigliano (2024): *A Beluga Formalization of the Harmony Lemma in the  $\pi$ -Calculus*. In Florian Rabe & Claudio Sacerdoti Coen, editors: *Proceedings Workshop on Logical Frameworks and Meta-Languages: Theory and Practice*, Tallinn, Estonia, 8th July 2024, *Electronic Proceedings in Theoretical Computer Science* 404, Open Publishing Association, pp. 1–17, doi:10.4204/EPTCS.404.1.
- [9] James Cheney, Michael Norrish & René Vestergaard (2012): *Formalizing Adequacy: A Case Study for Higher-order Abstract Syntax*. *J. Autom. Reason.* 49(2), pp. 209–239, doi:10.1007/S10817-011-9221-6.
- [10] Gavin Cox (2009): *SimCCSK: simulation of the reversible process calculi CCSK*. Master's thesis, University of Leicester. Available at [https://figshare.le.ac.uk/articles/thesis/SimCCSK\\_simulation\\_of\\_the\\_reversible\\_process\\_calculi\\_CCSK/10091681](https://figshare.le.ac.uk/articles/thesis/SimCCSK_simulation_of_the_reversible_process_calculi_CCSK/10091681).
- [11] Vincent Danos & Jean Krivine (2004): *Reversible Communicating Systems*. In Philippa Gardner & Nobuko Yoshida, editors: *CONCUR 2004, LNCS 3170*, Springer, pp. 292–307, doi:10.1007/978-3-540-28644-8\_19.
- [12] Pierpaolo Degano & Corrado Priami (2001): *Enhanced operational semantics*. *ACM Comput. Surv.* 33(2), pp. 135–176, doi:10.1145/384192.384194.
- [13] Robert Harper, Furio Honsell & Gordon Plotkin (1993): *A Framework for Defining Logics*. *J. ACM* 40(1), pp. 143–184, doi:10.1145/138027.138060.
- [14] Furio Honsell, Marino Miculan & Ivan Scagnetto (2001):  *$\pi$ -Calculus in (Co)Inductive Type Theory*. *Theor. Comput. Sci.* 253(2), pp. 239–285, doi:10.1016/S0304-3975(00)00095-5.
- [15] T. F. Melham (1994): *A Mechanized Theory of the  $\pi$ -Calculus in HOL*. *Nordic J. of Computing* 1(1), p. 50–76, doi:10.48456/tr-244.
- [16] Robin Milner (1980): *A Calculus of Communicating Systems*. LNCS, Springer-Verlag, doi:10.1007/3-540-10235-3.
- [17] 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.
- [18] Monica Nesi (1992): *A formalization of the process algebra CCS in high order logic*. Technical Report UCAM-CL-TR-278, University of Cambridge, Computer Laboratory, doi:10.48456/tr-278.
- [19] Frank Pfenning & Conal Elliott (1988): *Higher-Order Abstract Syntax*. In: *Proceedings of the ACM-SIGPLAN Conference on Programming Language Design and Implementation*, ACM Press, pp. 199–208, doi:10.1145/53990.54010.
- [20] Iain Phillips & Irek Ulidowski (2007): *Reversing algebraic process calculi*. *J. Log. Algebr. Program.* 73(1-2), pp. 70–96, doi:10.1016/j.jlap.2006.11.002.
- [21] Brigitte Pientka & Jana Dunfield (2010): *Beluga: A Framework for Programming and Reasoning with Deductive Systems (System Description)*. In Jürgen Giesl & Reiner Hähnle, editors: *Automated Reasoning*,

- 5th International Joint Conference, IJCAR 2010, Edinburgh, UK, July 16-19, 2010. Proceedings, Lecture Notes in Computer Science 6173, Springer, pp. 15–21, doi:10.1007/978-3-642-14203-1\_2.*
- [22] Irek Ulidowski, Iain Phillips & Shoji Yuen (2014): *Concurrency and Reversibility*. In Shigeru Yamashita & Shin-ichi Minato, editors: *Reversible Computation - 6th International Conference, RC 2014, Kyoto, Japan, July 10-11, 2014. Proceedings, LNCS 8507, Springer, pp. 1–14, doi:10.1007/978-3-319-08494-7\_1.*