



HAL
open science

A Behavioral Theory For Crash Failures and Erlang-style Recoveries In Distributed Systems

Giovanni Fabbretti, Ivan Lanese, Jean-Bernard Stefani

► **To cite this version:**

Giovanni Fabbretti, Ivan Lanese, Jean-Bernard Stefani. A Behavioral Theory For Crash Failures and Erlang-style Recoveries In Distributed Systems. RR-9511, Inria. 2023. hal-04123758

HAL Id: hal-04123758

<https://hal.science/hal-04123758v1>

Submitted on 27 Jun 2023

HAL is a multi-disciplinary open access archive for the deposit and dissemination of scientific research documents, whether they are published or not. The documents may come from teaching and research institutions in France or abroad, or from public or private research centers.

L'archive ouverte pluridisciplinaire **HAL**, est destinée au dépôt et à la diffusion de documents scientifiques de niveau recherche, publiés ou non, émanant des établissements d'enseignement et de recherche français ou étrangers, des laboratoires publics ou privés.



Distributed under a Creative Commons Attribution 4.0 International License

Inria

A Behavioral Theory For Crash Failures and Erlang-style Recoveries In Distributed Systems

Giovanni Fabbretti, Ivan Lanese, Jean-Bernard Stefani

**RESEARCH
REPORT**

N° 9511

May 2023

Project-Teams Spades and Focus

ISRN INRIA/RR--9511--FR+ENG

ISSN 0249-6399



A Behavioral Theory For Crash Failures and Erlang-style Recoveries In Distributed Systems

Giovanni Fabbretti *, Ivan Lanese[†], Jean-Bernard Stefani*

Project-Teams Spades and Focus

Research Report n° 9511 — May 2023 — 50 pages

Abstract: Distributed systems can be subject to various kinds of partial failures, and building fault-tolerance or failure mitigation mechanisms for distributed systems remains an important domain of research. In this paper, we present a calculus to formally model distributed systems subject to crash failures, and in which one can encode recovery mechanisms by leveraging a small set of lightweight (in terms of implementation cost) primitives. To the best of our knowledge, our calculus is the first one with support for *all* the following characteristics: i) asynchronous communication; ii) unique location for receivers; iii) dynamic nodes and links; iv) crash failures with recovery; v) nodes with imperfect knowledge of their context. We define a contextual equivalence for our calculus in the classical form of a barbed congruence, and a notion of bisimilarity which we prove fully abstract with respect to our barbed congruence. In addition, we show by means of examples that our calculus can support Erlang-style fault management and recovery, and that our behavioral theory agrees on key instances without recovery with previous work by Francalanza and Hennessy. This paper can be understood as a complete reworking and an extension to tackle recovery of Francalanza and Hennessy's work.

Key-words: formal calculi, distributed systems, concurrency, failures, recoveries

* Univ. Grenoble Alpes, INRIA, CNRS, Grenoble INP, LIG, France

[†] Focus Team, Univ. of Bologna/INRIA, Italy

**RESEARCH CENTRE
GRENOBLE – RHÔNE-ALPES**

Inovallée
655 avenue de l'Europe Montbonnot
38334 Saint Ismier Cedex

Une théorie comportementale des défaillances et des reprise en style Erlang dans les systèmes distribués.

Résumé : Les systèmes distribués peuvent être soumis à différents types de défaillances partielles, et l'élaboration de mécanismes de tolérance aux pannes ou d'atténuation des défaillances pour les systèmes distribués reste un domaine de recherche important. Dans cet article, nous présentons un calcul permettant de modéliser formellement les systèmes distribués soumis à des défaillances accidentelles, et dans lequel il est possible d'encoder des mécanismes de récupération en tirant parti d'un petit ensemble de primitives légères (en termes de coût d'implémentation). À notre connaissance, notre calcul est le premier à prendre en charge toutes les caractéristiques suivantes : i) communication asynchrone ; ii) emplacement unique pour les récepteurs ; iii) nœuds et liens dynamiques ; iv) pannes avec récupération ; v) nœuds ayant une connaissance imparfaite de leur contexte. Nous définissons une équivalence contextuelle pour notre calcul sous la forme classique d'une congruence barrée, ainsi qu'une notion de bisimilarité que nous prouvons totalement abstraite par rapport à notre congruence barrée. En outre, nous montrons à l'aide d'exemples que notre calcul peut prendre en charge la gestion et la récupération des fautes à la manière d'Erlang, et que notre théorie comportementale est en accord avec les travaux antérieurs de Francalanza et Hennessy en ce qui concerne les instances clés sans récupération. Cet article peut être considéré comme un remaniement complet et une extension du travail de Francalanza et Hennessy pour aborder la récupération.

Mots-clés : calcul formel, systèmes distribués, concurrence, défaillances, reprises

Contents

1	Introduction	3
2	Crash and Recovery: Motivating Example	5
3	The Calculus	7
3.1	Names and notations	7
3.2	Syntax	7
3.3	Reduction Semantics	9
3.4	Discussion	12
4	Behavioral Theory	13
4.1	Weak Barbed Congruence	13
4.2	A Labeled Transition Semantics	13
4.3	Full Abstraction	15
5	Examples	15
5.1	Deriving Erlang's Like Constructs	15
5.2	Behavioral Theory In Action	16
6	Related work and conclusion	18
A	Notations and rules	23
A.1	Notations	23
A.2	Calculus syntax and alpha-conversion on systems	23
A.3	Reduction semantics	25
A.4	LTS semantics	26
B	Erlang	30
B.1	Experiments on Erlang's Semantics	30
B.2	Running Example in Erlang	31
C	Modifying networks	33
D	Barb alternative	34
E	Full Abstraction	35
E.1	Soundness	35
E.2	Completeness	44
F	Bisimulation Of Running Example	48

1 Introduction

A key characteristic of distributed computer systems is the occurrence of partial failures, which can affect part of a system, e.g. failures of the system nodes (computers and the processes they support) or failures of the connections between them. The model of crash failures with recovery, where a node can fail by ceasing to operate entirely, and later on recover its operation (typically, after some administrator intervention) is an important model of failure to consider because of its relevance in practice. However the correct design of systems in this model is far from trivial, as the recent study on bugs affecting crash recovery in distributed systems demonstrates [17].

Developing a process calculus analysis of the behavior of distributed systems with crash failures and recovery can help in this respect for it can reveal subtle phenomena in the behavior of such systems, and it can be leveraged for the development of analysis tools, such as verifiers and debuggers. Unfortunately, to the best of our knowledge, the literature contains precious few examples of process calculi accounting for crash failures with recovery, notably [15, 1, 3, 4], and none account simultaneously for the following features which we deem essential to faithfully model actual distributed systems:

1. *Asynchronous communication*: interaction between processes proceeds not by rendez-vous but by an asynchronous exchange of messages, which can possibly be lost or reordered while transiting to their destination.
2. *Unique location for receivers*: each communication targets a single location, either local or remote. This would not be the case, e.g., if communication were via channels, and receivers for a same channel could reside on different locations. This feature ensures that the complexity of message exchange is commensurate with that of simple asynchronous communication used in the Internet, and that no hidden cost, due for instance to leader election or routing protocols, is implied for the implementation of a simple message exchange (see e.g. [14] for a discussion).
3. *Dynamic nodes and links*: the number of nodes in a system, and of communication links between them, is not fixed and may vary. During execution, new nodes and links can be established, existing nodes and links can be removed, either because of failures or by design. This feature is necessary to account for actual distributed systems whose configurations may vary at run-time, notably because of failure management and of performance management (e.g. scaling decisions in cloud systems). The explicit presence of links is important because partial connections often affect large distributed systems and because introducing link failures leads to a different behavioral theory than dealing with node failures only, as noted in [16].
4. *Crash failures*: when a node or a link fails, it does so silently, by ceasing execution. This is a simplifying assumption on the failure model of actual distributed systems, but one which, barring malicious faults, can be relatively well approximated in practice, as with the *let it fail* policy in Erlang.
5. *Recovery*: when a node or link has failed, it can be revived to resume its execution or function. This may imply external intervention (e.g. by a human administrator) and may not be under the control of the running system. Note that we place no strong requirement on the resume operation, just that a recovered node be able to perform some computation: e.g. it need not guarantee that a recovered node operates exactly as prior to its crash. This weak recovery model is consistent e.g. with node recovery in Erlang, or the node failure model adopted in the Verdi framework for distributed systems verification [29].

6. *Imperfect knowledge*: in general, in a distributed system, a node has only a partial and imperfect knowledge of the overall state of the system. We call *local view* the belief a node entertains of the status of its neighbors (nodes it assumes it is connected to). Local views are explicitly maintained by distributed programming language runtimes such as that of Erlang [2] and Elixir [21]. They play an essential part of key distributed algorithms, such as failure detection [19], membership management [7], checkpoint-rollback recovery schemes [9], or gossip-based peer sampling [20].

The two works dealing with crash failures and recovery that come closest to meeting these modeling requirements are [3] and the recent [4]. However, they only consider a fixed number of nodes, and to model recovery they rely on timers and checkpointing constructs (their *save primitives*). Checkpointing primitives are powerful mechanisms that do not match our weak requirements to account for recovery. For instance, recovery in Erlang systems can take place without relying on checkpointing, just by restarting failed nodes and processes in a predefined state, possibly relying on persistence mechanisms for preserving data across node failures.

The goal of this paper, then, is to introduce a process calculus exhibiting all the above features, with no recourse to timers, perfect failure detection or a built-in checkpointing primitive. We have two additional requirements for this calculus. On the one hand, it has to stay close to the behavior of Erlang; on the other hand it has to stay as close as possible to the work by Francalanza and Hennessy [16]. The Erlang programming language and its environment are representative of modern distributed programming facilities. It is a functional, concurrent and distributed language based on the actor model, and it is used in several large distributed projects [12]. It is well known for its “let it fail” policy, whereby a service that is not working as expected is killed as soon as the faulty behavior is detected, if not dead already, and restarted by its supervisor. Staying close to Erlang ensures our modeling remains faithful to actual distributed systems. The work by Francalanza and Hennessy [16] constitutes a good benchmark for comparison. While not tackling recovery and relying on perfect failure detection, it contains several of the features mentioned above, and develops a behavioral theory with node and link failures. Staying close to it allows us to compare our behavioral theory with an established one, when restricting our attention to systems without recovery.

The rest of this paper is organized as follows. Section 2 introduces a motivating example, which serves also as an introduction to our calculus. Section 3 presents the calculus, its reduction semantics, and discusses key design choices. Section 4 equips our calculus with a notion of barbed congruence and characterizes it by a notion of bisimilarity to obtain a proof technique for checking system equivalence. Section 5 shows our calculus can encode constructs similar to Erlang primitives to build systems with hierarchical failure management. Moreover, we show that our behavioral theory agrees with the one in [16] on key examples without recovery. Finally, Section 6 discusses related work and concludes. Details of proofs and complementary material are available in the Appendix.

2 Crash and Recovery: Motivating Example


Before presenting our calculus, we discuss a motivating example which serves as an informal introduction (for the time being, we omit some details). An Erlang implementation of the example is discussed in Appendix B. Consider the following system:

$$\text{servD} = \nu n_r, n_b, r_1, r_2, b. \Delta \triangleright ([I]^{n_i} \parallel [R]^{n_r} \parallel [B]^{n_b})$$

where:

$$\begin{aligned} I &= req(y, z).spawn\ n_r.\bar{r}_1\langle y, z \rangle & R &= (r_1(y, z).spawn\ n_b.\bar{b}\langle y, z \rangle) \mid (r_2(y, z).spawn\ n_i.\bar{z}\langle y \rangle) \\ B &= b(y, z).spawn\ n_r.\bar{r}_2\langle z, w_y \rangle \end{aligned}$$

System **servD** depicts a distributed server running on a network Δ . The network, Δ , whose formal definition we omit for the moment, can be graphically represented as follows,



A location in our calculus represents a locus of computation. It can represent a hardware node in a distributed system, or a virtual machine or a container running on a hardware node. A location constitutes also a unit of crash failure: when a location fails, all processes inside the location cease to function. Located processes take the form $[P]^n$, where P denotes a process and where n is the location name. There can be several processes located at the same location: for instance, $[P]^n \parallel [Q]^n$ denotes two processes P and Q running in parallel inside the same location n .

The (admittedly simplistic) system **servD** behaves as follows. The *interface* process I , running on n_i , awaits a single request from the environment on channel req . Once received, the elements of the request (a parameter y , and a return channel z) are routed to location n_b , which runs the *backend* process B , through location n_r , which hosts the *router* process R , as there is no direct link between n_i and n_b . The router awaits for the elements of the request on a private channel r_1 and forwards them by spawning a message $\bar{b}\langle y, z \rangle$ on n_b , where b is a private channel. Message sending in our calculus can only occur locally. For two remote locations to communicate, like for n_i and n_r , it is necessary for one to asynchronously spawn the message on the other one. The backend handles the information y and returns the answer w_y to the interface, again by routing it through n_r . Finally, the interface emits the answer on z for the client to consume it.

The following is a possible client for **servD**

$$[\text{spawn } n_i.(\bar{req}\langle h, z \rangle \mid z(w).Q)]^c$$

It sends a request to the interface n_i by spawning it on n_i . It also spawns a process on n_i to handle the response, which will take the form of a message on channel z located on n_i .

Now, consider the following system:

$$\text{servDF} = \nu n_r, n_b, n_c, r_1, r_2, b, \text{retry}. \Delta' \triangleright ([J]^{n_i} \parallel [R]^{n_r} \parallel [\text{kill}]^{n_r} \parallel [!B]^{n_b} \parallel [!C]^{n_c})$$

where:

$$\begin{aligned} J &= req(y, z).((spawn\ n_r.\bar{r}_1\langle y, z \rangle) \mid \text{retry}.spawn\ n_r.\bar{r}_1\langle y, z \rangle) \\ C &= \text{create } n_r.(R \mid spawn\ n_i.\text{retry}) \end{aligned}$$

Here, Δ' could be graphically represented like Δ only with an extra link between n_r and n_c . System **servDF** represents a distributed server where location n_r may be subject to one failure, modeled by the primitive `kill`. The system has a recovery mechanism in place to deal with that potential failure: the *controller* $[!C]^{n_c}$ is a location that keeps trying, through the apposite primitive `create`, to recreate n_r (the `!` operator is akin to the π -calculus operator for replication), together with a message to restart the handling of the request. The interface is more sophisticated as it can now send a second request when asked to `retry` if something goes wrong with the first one. If n_r fails, the controller can create another router with a message `retry` for the interface which will start the second attempt.

At first glance, **servDF** seems to correctly handle the failure scenario, but that is not the case. Indeed, consider an execution where the failure happens after the request has already been handled. In such case, the request would be processed again and another response sent back by the interface, a behavior that cannot be exhibited by **servD**. If one considers **servD** as the specification to meet, **servDF** does not satisfy it.

The following system correctly handles recovery, in that it meets the **servD** specification:

$$\mathbf{servDFR} = \nu n_r, n_b, n_c, r_1, r_2, b, c, \mathit{retry}. \Delta' \triangleright ([K]^{n_i} \parallel [R]^{n_r} \parallel [\mathit{kill}]^{n_r} \parallel [!B]^{n_b} \parallel [!C]^{n_c})$$

where:

$$K = \mathit{req}(y, z).((\mathit{spawn } n_r.\bar{r}_1\langle y, c \rangle) \mid c(w).\bar{z}\langle w \rangle \mid \mathit{retry}.\mathit{spawn } n_r.\bar{r}_1\langle y, c \rangle)$$

In system **servDFR** the response from the backend is not directly sent to the client, but goes through a private channel c on which the interface listens only once. This mechanism prevents emitting the answer to the request twice. System **servDFR** can be understood as a masking 1-fault tolerant system, equivalent to the ideal one **servD**. We develop in Section 4 a behavioral theory able to tell apart **servD** from **servDF**, and able to prove equivalent **servD** and **servDFR**.

3 The Calculus

3.1 Names and notations

We assume given mutually disjoint infinite denumerable sets \mathbf{C} , \mathbf{N} and \mathbf{I} . \mathbf{C} is the set of *channel names*, \mathbf{N} is the set of *location names*, and \mathbf{I} is the set of *incarnation variables*. We use the set of integers \mathbb{Z} as the set of *incarnation numbers*. An incarnation number is paired with a location name for recovery purposes, to distinguish the current instance of a location from its past failed instances. We denote by \mathbf{N}^\odot the set $\mathbf{N} \cup \{\odot\}$, where $\odot \notin \mathbf{N}$. As in the π -calculus, channel names can be free or bound in terms. The same holds for location names. Incarnation variables can be bound, but not incarnation numbers. We denote by \tilde{u} a finite (possibly empty) tuple of elements. We write $T\{\tilde{v}/\tilde{u}\}$ for the usual capture-avoiding substitution of elements of \tilde{u} by elements of \tilde{v} in term T , assuming tuples \tilde{u} and \tilde{v} have the same arity. We write u, \tilde{v} or \tilde{v}, u for the tuple \tilde{v} extended with element u as first or last element. Abusing notation, we sometimes identify a tuple \tilde{u} with the set of its elements. We denote by \mathbb{N}^+ the set of strictly positive integers (by definition $0 \notin \mathbb{N}^+$), and by \mathbb{N} the set of positive integers ($0 \in \mathbb{N}$). We denote by $\hat{\mathbf{O}}$ the function $\hat{\mathbf{O}} : \mathbf{N}^\odot \rightarrow \mathbb{Z}$ that maps any $n \in \mathbf{N}$ to 0 and \odot to 1.

3.2 Syntax

Systems in our calculus are defined through three levels of syntax, one for *processes*, one for *configurations*, one for *systems*.

The syntax of *processes* is defined as follows:

$$P, Q ::= \mathbf{0} \mid \bar{x}(\tilde{u}).P \mid x(\tilde{v}).P \mid !x(\tilde{v}).P \mid \nu w.P \mid \text{if } r = s \text{ then } P \text{ else } Q \mid P \mid Q \\ \text{node}(n, \lambda).P \mid \text{remove } n.P \mid \text{spawn } n.P \mid \text{kill} \mid \\ \text{create } n.P \mid \text{link } n.P \mid \text{unlink } n.P$$

$$\text{where: } \tilde{u}, \tilde{r}, \tilde{s} \subset \mathbf{C} \cup \mathbf{N}^\odot \cup \mathbf{I} \cup \mathbb{Z} \quad \tilde{v} \subset \mathbf{C} \cup \mathbf{N} \cup \mathbf{I} \quad w \in \mathbf{C} \cup \mathbf{N} \quad x \in \mathbf{C} \quad n \in \mathbf{N} \quad \lambda \in \mathbf{I}$$

Terms of the form $x(\tilde{u}).P$, $\nu w.P$, and $\text{node}(n, \lambda).P$ are binding constructs for their arguments \tilde{u} , w and n, λ , respectively.

The syntax of processes is that of the π -calculus with matching [26] and replicated receivers (first line of productions), extended with primitives for distributed computing inspired from the Erlang programming language (second line of productions), and three primitives to activate locations, establish and remove links (third line of productions).

$\mathbf{0}$ is the null process which can take no action. Processes can communicate by emitting a message $\bar{x}(\tilde{v}).P$ which can be received by some receiver process of the form $x(\tilde{u}).Q$ residing on the same location. We write $\bar{x}(\tilde{u})$ for $\bar{x}(\tilde{v}).\mathbf{0}$, and just \bar{x} when \tilde{u} is empty. We assume the calculus is well-sorted, so that the arity of receivers always matches that of received messages. The construct $!x(\tilde{u}).P$ is the replicated input construct, which replicates itself when receiving a message on channel x . Note that we use in our examples (as we did in the previous Section) the short-hand $!P$ for $\nu c.(\bar{c} \mid !c.(P \mid \bar{c}))$. The construct $\nu w.P$ is the standard restriction construct, which creates a fresh location or channel name. If \tilde{u} is a (possibly empty) tuple of names, we write $\nu \tilde{u}.P$ for $\nu u_1 \dots \nu u_n.P$ if $\tilde{u} = (u_1, \dots, u_n)$. If \tilde{u} is empty, $\nu \tilde{u}.P$ is just P . The construct $\text{if } r = s \text{ then } P \text{ else } Q$ tests the equality of names r and s and continues as P if the names match and as Q otherwise. The construct $P \mid Q$ is the standard parallel composition for processes. Primitive `node`(n, λ). P substitutes n with the name of the current location, λ with the current incarnation number, and continues as P . Primitive `remove` $n.P$ removes the location with name n from the local view of the current location and continues as P . Primitive `spawn` $n.P$ launches process P at the location named n , if the latter is accessible. Primitive `kill` stops the current location in its current incarnation: no process can execute on a killed location; `kill` also models the crash of a location. Primitive `create` $n.P$ creates a new location n , or reactivates a killed location with a new incarnation number, and launches process P on it. Primitive `link` $n.P$ creates a connection between the current location and n and continues as P , while, `unlink` $n.P$ breaks the link between the current location and n and continues as P ; `unlink` also models the failure of a link.

The syntax of *configurations* is defined as follows:

$$L, M, N ::= \mathbf{0} \mid [P]_{\lambda}^n \mid N \parallel M \quad \text{where:} \quad n \in \mathbb{N}^{\odot} \quad \lambda \in \mathbb{Z}$$

A configuration can be the empty configuration $\mathbf{0}$, a located process $[P]_{\lambda}^n$ or a parallel composition of configurations $N \parallel M$. A located process $[P]_{\lambda}^n$ is a process P running on location (n, λ) , where n is the name of the location and λ is an *incarnation number* (used for recovery). In examples, we may drop incarnation numbers of located processes if they are not relevant. Note that the special name \odot identifies a well-known location which we will assume to be un-killable. Location \odot is used merely for technical purposes to ensure we can simply build appropriate contexts for running systems. Apart from being unkillable, location \odot behaves just as any other locations. We denote by \mathbb{L} the set of configurations.

The syntax of *systems* is defined as follows:

$$S, R ::= \Delta \triangleright N \mid \nu w.S \quad \text{where:} \quad w \in \mathbb{N} \cup \mathbb{C}$$

A system is the composition of a network Δ with a configuration N , or a system under a name (channel or location) restriction. We denote by \mathbb{S} the set of systems. A *network* Δ , is a tuple $\langle \mathcal{A}, \mathcal{L}, \mathcal{V} \rangle$ where

- \mathcal{A} is a function $\mathcal{A} : \mathbb{N}^{\odot} \rightarrow \mathbb{Z}$ such that $\mathcal{A}(\odot) = 1$ and such that the set $\text{supp}(\mathcal{A}) \stackrel{\text{def}}{=} \{n \in \mathbb{N} \mid \mathcal{A}(n) \neq 0\}$ is finite. Function \mathcal{A} may record three types of information on locations. If $\mathcal{A}(n) = \lambda \in \mathbb{N}^+$, then location n is alive and its current incarnation number is λ . If $\mathcal{A}(n) = -\lambda, \lambda \in \mathbb{N}^+$ then location n has been killed and its last incarnation number while alive was λ . If $\mathcal{A}(n) = 0$, then there is no location n in the network, alive or not. Because

of the finiteness condition above, a network can only host a finite number of locations, alive or dead.

- $\mathcal{L} \subseteq \mathbb{N}^\circ \times \mathbb{N}^\circ$ is the set of links between locations. \mathcal{L} is a finite symmetric binary relation over location names such that $\text{dom}(\mathcal{L}) \stackrel{\text{def}}{=} \{n \in \mathbb{N}^\circ \mid \exists m, (n, m) \in \mathcal{L}\}$ is finite.
- $\mathcal{V} : \mathbb{N}^\circ \rightarrow (\mathbb{N}^\circ \rightarrow \mathbb{N})$ is a function that maps location names to their *local view*, which is such that the set $\text{supp}(\mathcal{V}) \stackrel{\text{def}}{=} \{n \in \mathbb{N} \mid \mathcal{V}(n) \neq \hat{\mathbf{0}}\}$ is finite and $\text{supp}(\mathcal{V}) \subseteq \text{supp}(\mathcal{A})$. The local view of a location n is a function $\mathcal{V}(n) : \mathbb{N}^\circ \rightarrow \mathbb{N}$ such that the set $\{m \in \mathbb{N} \mid \mathcal{V}(n)(m) \neq 0\}$ is finite. If $\mathcal{V}(n)(m) = \kappa \in \mathbb{N}^+$, then location m in its incarnation κ is believed by n to be alive. If $\mathcal{V}(n)(m) = 0$, then location n holds no belief on the status of location m .

For convenience we use $\Delta_{\mathcal{A}}$, $\Delta_{\mathcal{L}}$, and $\Delta_{\mathcal{V}}$ to denote the individual components of a network representation Δ , and we use the following notations for extracting information from Δ :

- $\Delta \vdash n_\lambda : \text{alive}$ if $\Delta_{\mathcal{A}}(n) = \lambda$ and $\lambda \in \mathbb{N}^+$.
- $\Delta \vdash n : \text{dead}$ if $\Delta_{\mathcal{A}}(n) \notin \mathbb{N}^+$
- $\Delta \vdash n \leftrightarrow m$ if $(n, m) \in \Delta_{\mathcal{L}}$
- $\Delta \vdash n_\lambda \rightsquigarrow m_\kappa$ if $(n, m) \in \Delta_{\mathcal{L}}$, $\Delta \vdash n_\lambda : \text{alive}$ and $\Delta \vdash m_\kappa : \text{alive}$

We now define update operations over a network Δ .

Definition 1 (Network updates). *Network update operations are defined as follows:*

- $\Delta \oplus n \leftrightarrow m = \langle \Delta_{\mathcal{A}}, \Delta_{\mathcal{L}} \cup \{(n, m), (m, n)\}, \Delta_{\mathcal{V}} \rangle$
- $\Delta \ominus n \leftrightarrow m = \langle \Delta_{\mathcal{A}}, \Delta_{\mathcal{L}} \setminus \{(n, m), (m, n)\}, \Delta_{\mathcal{V}} \rangle$
- $\Delta \oplus (n, \lambda) = \langle \Delta_{\mathcal{A}}[n \mapsto \lambda], \Delta_{\mathcal{L}} \cup \{(n, n)\}, \Delta_{\mathcal{V}}[n \mapsto \hat{\mathbf{0}}] \rangle$
- $\Delta \ominus (n, \lambda) = \langle \Delta_{\mathcal{A}}[n \mapsto -\lambda], \Delta_{\mathcal{L}}, \Delta_{\mathcal{V}} \rangle$
- $\Delta \oplus n \succ (m, \lambda) = \langle \Delta_{\mathcal{A}}, \Delta_{\mathcal{L}}, \Delta_{\mathcal{V}}[n \mapsto \Delta_{\mathcal{V}}(n)[m \mapsto \lambda]] \rangle$, if $n \neq m$
- $\Delta \ominus n \succ m = \langle \Delta_{\mathcal{A}}, \Delta_{\mathcal{L}}, \Delta_{\mathcal{V}}[n \mapsto \Delta_{\mathcal{V}}(n)[m \mapsto 0]] \rangle$, if $n \neq m$
- $\Delta \ominus n \succ n = \Delta \oplus n \succ (n, \lambda) = \Delta$

$\Delta \oplus n \leftrightarrow m$ and $\Delta \ominus n \leftrightarrow m$ add and remove a link, respectively, between n and m . $\Delta \oplus (n, \lambda)$ activates location n with incarnation number λ , and sets its view to the empty one. $\Delta \ominus (n, \lambda)$ kills a location in its incarnation λ . $\Delta \oplus n \succ (m, \lambda)$ adds (m, λ) to the view of n , and $\Delta \ominus n \succ m$ removes any belief on location m from the view of n .

We use below the notion of closed systems. Closed systems are systems which do not have free incarnation variables. The definition of free variables in systems and that of free names in configurations is completely standard. The notion of free names in a system is slightly unconventional because of the presence of a network, but it can be defined straightforwardly (see Appendix A for details).

$$\begin{array}{c}
\text{[S.PAR.C]} \ N \parallel M \equiv M \parallel N \qquad \text{[S.PAR.A]} \ (L \parallel M) \parallel N \equiv L \parallel (M \parallel N) \qquad \text{[S.PAR.N]} \ (N \parallel \mathbf{0}) \equiv N \\
\text{[S.RES.C]} \ \nu u.\nu v.S \equiv \nu v.\nu u.S \qquad \text{[S.RES.NIL]} \ \frac{}{\nu u.S \equiv S} \ u \notin \text{fn}(S) \qquad \text{[S.}\alpha\text{]} \ \frac{S =_\alpha R}{S \equiv R} \\
\text{[S.CTX]} \ \frac{N \equiv M}{C[N] \equiv C[M]}
\end{array}$$

Figure 1: Structural Congruence Rules

Assuming $\Delta \vdash n_\lambda : \text{alive}$

$$\begin{array}{c}
\text{MSG} \quad \frac{}{\Delta \triangleright [\bar{x}(\tilde{v}).Q]_\lambda^n \parallel [x(\tilde{u}).P]_\lambda^n \longrightarrow \Delta \triangleright [Q]_\lambda^n \parallel [P\{\tilde{v}/\tilde{u}\}]_\lambda^n} \qquad \text{BANG} \quad \frac{}{\Delta \triangleright [!x(\tilde{u}).P]_\lambda^n \longrightarrow \Delta \triangleright [x(\tilde{u}).(P \mid !x(\tilde{u}).P)]_\lambda^n} \\
\text{NEW} \quad \frac{}{\Delta \triangleright [\nu u.P]_\lambda^n \longrightarrow \nu u.\Delta \triangleright [P]_\lambda^n} \ u \notin \text{fn}(\Delta) \cup \{n, \lambda\} \qquad \text{FORK} \quad \frac{}{\Delta \triangleright [P \mid Q]_\lambda^n \longrightarrow \Delta \triangleright [P]_\lambda^n \parallel [Q]_\lambda^n} \\
\text{IF-NEQ} \quad \frac{}{\Delta \triangleright [\text{if } r = s \text{ then } P \text{ else } Q]_\lambda^n \longrightarrow \Delta \triangleright [Q]_\lambda^n} \ r \neq s \qquad \text{NODE} \quad \frac{}{\Delta \triangleright [\text{node}(m, \kappa).P]_\lambda^n \longrightarrow \Delta \triangleright [P\{n, \lambda/m, \kappa\}]_\lambda^n} \\
\text{IF-EQ} \quad \frac{}{\Delta \triangleright [\text{if } r = r \text{ then } P \text{ else } Q]_\lambda^n \longrightarrow \Delta \triangleright [P]_\lambda^n} \qquad \text{REMOVE} \quad \frac{}{\Delta \triangleright [\text{remove } m.P]_\lambda^n \longrightarrow \Delta \ominus n \triangleright m \triangleright [P]_\lambda^n}
\end{array}$$

Figure 2: Local Rules

3.3 Reduction Semantics

The operational semantics of our calculus is defined via a reduction semantics given by a binary relation $\longrightarrow \subseteq \mathbb{S} \times \mathbb{S}$ between closed systems, and a structural congruence relation $\equiv \subseteq \mathbb{S}^2 \cup \mathbb{L}^2$, that is a binary equivalence relation between systems and between configurations. Evaluation contexts are “systems with a hole” defined by the following grammar:

$$\mathbb{C} ::= \nu \tilde{w}.\Delta \triangleright \mathbb{E} \qquad \mathbb{E} ::= \cdot \mid (N \parallel \mathbb{E}) \qquad \text{where:} \qquad \tilde{w} \subset \mathbb{N} \cup \mathbb{C}$$

Relation \equiv is the smallest equivalence relation defined by the rules in Fig. 1, where $=_\alpha$ stands for equality up to alpha-conversion, $M, N, L \in \mathbb{L}$, and $S, R \in \mathbb{S}$. Most rules are mundane. Rule S.CTX turns \equiv into a congruence for the parallel and restriction operators. Alpha-conversion on systems, which appears in Rule S. α , is slightly unusual but can be defined straightforwardly (see Appendix A for details).

The reduction relation \longrightarrow is defined by the rules in Fig. 2, 3 and 4. Fig. 2 depicts the *local* reduction rules, i.e., those rules that involve only a single location and that essentially do not modify the network. Rule MSG defines the receipt of a message by an input process. Rule BANG defines the expansion of a replicated input process. Rule NEW performs the scope extrusion of a name from a process to a system. We introduced this rule as a computational step instead of a structural congruence rule for it simplifies our proofs. Rule FORK turns a parallel composition into parallel threads in the same location. Rules IF-EQ, IF-NEQ define the semantics of the branching construct. Rule NODE gets hold of the current location name and its incarnation number for further processing. Finally, rule REMOVE deletes from the local view of the current location the belief it may hold about a given location m .

Fig. 3 depicts the *distributed* rules, i.e., rules that involve several locations or modify the network. Rule SPAWN-S defines a successful spawn, conditional upon the fact that a link exists

Assuming $\Delta \vdash n_\lambda : \text{alive}$

$$\begin{array}{c}
 \text{SPAWN-S} \\
 \frac{}{\Delta \triangleright [\text{spawn } m.P]_\lambda^n \longrightarrow \Delta \oplus m \succ (n, \lambda) \triangleright [P]_\kappa^m} \quad \begin{array}{l} \Delta_n(m) = \kappa \\ \Delta \vdash n_\lambda \leftrightarrow m_\kappa \end{array} \\
 \\
 \text{SPAWN-F} \\
 \frac{}{\Delta \triangleright [\text{spawn } m.P]_\lambda^n \longrightarrow \Delta \ominus n \succ m \triangleright \mathbf{0}} \quad \begin{array}{l} \Delta_n(m) = \kappa \\ \Delta \not\vdash n_\lambda \leftrightarrow m_\kappa \end{array} \\
 \\
 \text{CREATE-S} \\
 \frac{}{\Delta \triangleright [\text{create } m.P]_\lambda^n \longrightarrow \Delta \oplus (m, \kappa + 1) \triangleright [P]_{\kappa+1}^m} \quad \begin{array}{l} \Delta \vdash m : \text{dead} \\ \Delta_{\mathcal{A}}(m) = -\kappa \end{array} \\
 \\
 \text{CREATE-F} \\
 \frac{}{\Delta \triangleright [\text{create } m.P]_\lambda^n \longrightarrow \Delta \triangleright \mathbf{0}} \quad \Delta \not\vdash m : \text{dead} \\
 \\
 \text{KILL} \\
 \frac{}{\Delta \triangleright [\text{kill } m.P]_\lambda^n \longrightarrow \Delta \ominus (n, \lambda) \triangleright \mathbf{0}} \quad n \neq \odot \\
 \\
 \text{LINK} \\
 \frac{}{\Delta \triangleright [\text{link } m.P]_\lambda^n \longrightarrow \Delta \oplus n \leftrightarrow m \triangleright [P]_\lambda^n} \quad \Delta \not\vdash n \leftrightarrow m \\
 \\
 \text{UNLINK} \\
 \frac{}{\Delta \triangleright [\text{unlink } m.P]_\lambda^n \longrightarrow \Delta \ominus n \leftrightarrow m \triangleright [P]_\lambda^n} \quad \Delta \vdash n \leftrightarrow m
 \end{array}$$

Figure 3: Distributed Rules

between the spawning and target locations, and that the spawning location rightly believes the target location, with the incarnation recorded in its view, to be alive, or has no belief on the target location in its local view. This last constraint is captured by the side condition $\Delta_n(m) = \kappa$, which we formally define as follow:

$$\Delta_n(m) = \begin{cases} \kappa & \text{if } n = m \text{ and } \Delta_{\mathcal{A}}(n) = \kappa \\ \kappa & \text{if } n \neq m \text{ and } \Delta_{\mathcal{V}}(n)(m) = \kappa \\ \kappa & \text{if } n \neq m \text{ and } \Delta_{\mathcal{V}}(n)(m) = 0 \text{ and } \Delta_{\mathcal{A}}(m) = \kappa \\ 0 & \text{otherwise} \end{cases}$$

Note that a successful spawn updates the target local view with the belief that the spawning location is alive with the incarnation number it had when it initiated the spawn. Rule `SPAWN-F` defines a failed spawn, which may fail due to a wrong view, to a missing link between the two locations, or because the remote location is not alive. The view of the sender is updated by removing the belief on the target. This behavior is inspired by Erlang whose runtime, in a manner transparent to the user, updates the view of the spawning location if it receives no acknowledgement from the target as part of its distributed protocol.

Rules `LINK` and `UNLINK` define respectively the establishment of a link and the removal of a link. Rule `CREATE-S` defines the successful creation or reactivation of a location, provided it did not already exist in the network or was crashed. The newly activated location has an incarnation number that is the successor of the previous one (0 by convention if the location was not present in the network). Note that we do not require the existence of an alive link between the current location and the one to be activated since we want this operation to also model the possibility of interventions external to the system, such as those performed by human administrators. Note also that a location cannot be created anew each time, otherwise it would be impossible to resume the execution of a service under a well-known name. The use of incarnation numbers provides support for recovery schemes (see the discussion below in Section 3.4). Rule `CREATE-F` defines

$$\begin{array}{c}
\text{PAR} \\
\frac{\Delta \triangleright N \longrightarrow \nu \tilde{u}. \Delta' \triangleright N' \quad \tilde{u} \cap \text{fn}(M) = \emptyset}{\Delta \triangleright N \parallel M \longrightarrow \nu \tilde{u}. \Delta' \triangleright N' \parallel M}
\end{array}
\qquad
\begin{array}{c}
\text{RES} \\
\frac{S \longrightarrow S'}{\nu u. S \longrightarrow \nu u. S'}
\end{array}
\qquad
\begin{array}{c}
\text{STR} \\
\frac{S \equiv S' \quad S' \longrightarrow R' \quad R' \equiv R}{S \longrightarrow R}
\end{array}$$

Figure 4: Contextual rules

the failure of a create operation, which can fail because the location to activate may already be alive. Rule `KILL` defines the killing of a location.

Fig. 4 shows the contextual rules of our calculus. Rules `PAR`, and `RES` are the rules respectively for parallel execution, and execution under restriction. Rule `PAR` is slightly unconventional in its use of restriction. When we consider the case where \tilde{u} is empty in rule `PAR`, we obtain a more standard-looking contextual rule for the parallel operator. However, we also have to consider cases when the active branch in the composition promotes a restriction at system level (via an application of rule `NEW`). In this case we have to avoid name capture by the idle branch in the composition. This way of proceeding, coupled with the systematic presence of the distinguished location \odot in any network, spares us the need to introduce parallel composition between systems. The intuition is that we can always extend a system with a process located on \odot (to ensure it is alive) performing the desired changes on the public part of the network. This intuition is formalized in Appendix C.

3.4 Discussion

Communication. In our calculus, communication is local. Remote communication is obtained using the `spawn` operation (very similar to the `go` operation in [16]) to send to a target node a process performing an output of a message. As a result remote communication is asynchronous (since the `spawn` is so), and there is a single location where the receiver can reside (since such location is specified in the `spawn`). This avoids the need of using a type system to ensure that possible receivers are located on a same node, differently from calculi based on channel-based remote communication such as [1].

Incarnation numbers. Incarnation numbers are called creation numbers in Erlang [11]. We introduce them in our calculus for two main reasons. On the one hand, this ensures we are faithful to Erlang and to the behavior Erlang systems exhibit in presence of failures (more details are provided in Appendix B). On the other hand, this ensures we have basic support in place for encoding different recovery schemes. Incarnation numbers ensure that a message issued by a previous incarnation of a location can safely be dropped, avoiding message duplication across different incarnations of the same location. They are present, under various names such as incarnation or epoch numbers, in several rollback-recovery schemes surveyed in [9], such as optimistic recovery [28, 6] or causal logging schemes [10]. They are also used for scalable distributed failure detection schemes [19] and in the SWIM protocol combining failure detection and membership management [7].

Imperfect knowledge. In distributed systems, the only way for locations to know something about the context that surrounds them is to communicate. If a location n receives a message from a remote one then n learns something on the context, namely that at some point in time the remote location was alive and working since it sent a message to n . Nonetheless, n cannot infer anything on the current status of the remote location or the status of the connection. Indeed, it could have stopped right after sending the message or the link could have broken right after the message was received or both. Erlang systems, like many others, have an optimistic approach: after a first two-way interaction two locations establish a mutual knowledge of their respective

incarnations, typically by means of a shared socket connection. From that point on, they keep using that shared connection until their view changes, rather than setting up a new connection for each message exchange. Reflecting this in our calculus plays a role in the semantics of our spawn primitive whenever the view of the locality is not in sync with the real state of the system. This in turn plays a role in our behavioral theory, as the following example illustrates.

Consider a variant **servDFV** of our running example where n_c is linked to n_i instead of n_r , the network is such that the router n_r is in its incarnation κ and the local view of the interface location n_i contains $n_r \mapsto \kappa$. The controller process running on n_c is defined as follows:

$$C = \text{create } n_r.R \mid \text{spawn } n_i.\overline{\text{retry}}$$

Now, **servDFV** is not equivalent to **servD** due to the fact that, in case of failure of n_r , it is the controller informing the interface to restart the request and not the router n_r , thus failing to update n_i 's local view with the knowledge of n_r 's new incarnation. A message from the interface at this point would fail as its local view contains the previous incarnation of n_r . If not for the imperfect knowledge of the context, **servDFV** would have been equivalent to **servD**. In Appendix B we discuss an implementation of this system and we show that this behavior arises in reality too.

4 Behavioral Theory

4.1 Weak Barbed Congruence

We define a standard notion of contextual equivalence called *weak barbed congruence*, originally proposed in [24]. We denote by \Longrightarrow the reflexive and transitive closure of the reduction relation \longrightarrow . We rely on a notion of observables on systems, called *barbs*, formally defined as follows:

Definition 2 (Barb). *A system S exhibits a barb on channel x at location n in its incarnation λ , in symbols $S \downarrow_{x@n_\lambda}$, iff $S \equiv \nu \tilde{u}.\Delta \triangleright [\bar{x}\langle \tilde{v} \rangle.P]_\lambda^n \parallel N$, for some $x, n, \lambda, \tilde{u}, \tilde{v}, P, N$, where $x, n \notin \tilde{u}$, and $\Delta \vdash n_\lambda : \text{alive}$. Also, $S \downarrow_{x@n_\lambda}$ iff $S \Rightarrow S'$ and $S' \downarrow_{x@n_\lambda}$.*

We now define standard properties expected for a contextual equivalence.

Definition 3 (System congruence). *An equivalence relation \mathcal{R} over closed systems is a system congruence iff, whenever $\nu \tilde{u}.\Delta_1 \triangleright N \mathcal{R} \nu \tilde{v}.\Delta_2 \triangleright M$, for any names \tilde{w} and for any configuration L such that $\text{fn}(L) \cap \tilde{u} = \text{fn}(L) \cap \tilde{v} = \emptyset$, we have:*

$$\nu \tilde{w}.\nu \tilde{u}.\Delta_1 \triangleright N \parallel L \mathcal{R} \nu \tilde{w}.\nu \tilde{v}.\Delta_2 \triangleright M \parallel L$$

Definition 4 (Weak Barb-preserving relation). *A relation \mathcal{R} over closed systems is weak barb-preserving iff whenever $S \mathcal{R} R$ and $S \downarrow_{x@n}$ then $R \downarrow_{x@n}$.*

Definition 5 (Weak reduction-closed relation). *A relation \mathcal{R} over closed systems is weak reduction-closed iff whenever $S \mathcal{R} R$ and $S \longrightarrow S'$ then $R \Longrightarrow R'$ for some R' such that $S' \mathcal{R} R'$.*

Definition 6 (Weak barbed congruence). *Weak barbed congruence, noted \approx , is the largest weak barb-preserving, reduction-closed, system congruence.*

As a first result, we can check that structural congruence is included in weak barbed congruence:

Assuming $\Delta \vdash n_\lambda : \text{alive}$

$$\frac{\text{L-IN}}{\Delta \triangleright [x(\tilde{v}).P]_\lambda^n \xrightarrow{x(\tilde{u})@n_\lambda} \Delta \triangleright [P\{\tilde{u}/\tilde{v}\}]_\lambda^n} \quad \frac{\text{L-OUT}}{\Delta \triangleright [\bar{x}(\tilde{v}).P]_\lambda^n \xrightarrow{\bar{x}(\tilde{v})@n_\lambda} \Delta \triangleright [P]_\lambda^n}$$

Figure 5: Concurrent and distributed Rules

Proposition 1. *Structural congruence \equiv is a weak barb-preserving reduction-closed system congruence.*

Proof. Structural congruence is a barb-preserving system-congruence by definition. The fact that structural congruence is weak reduction-closed follows from propositions 5 and 6 in Appendix E. \square

Much as in [25], a simpler kind of barbs gives rise to the same barbed congruence. The interested reader can find a discussion in Appendix D. We keep the more detailed observables $\downarrow_{x@n_\lambda}$ for convenience, to simplify certain arguments in our proofs.

4.2 A Labeled Transition Semantics

In this section we present a labeled transition semantics for our calculus in order to have a co-inductive characterization of weak barbed congruence. Labels α in our LTS semantics take the following forms:

$$\alpha ::= \tau \mid \nu \tilde{w}.\bar{x}(\tilde{u})@n_\lambda \mid x(\tilde{u})@n_\lambda \mid \text{kill}(n, \lambda) \mid \text{create}(n, \lambda) \mid \oplus n_\lambda \mapsto m \mid \ominus n_\lambda \mapsto m \mid n_\lambda \succ m$$

The first three labels are classical: silent action, output action (possibly with restricted names), and input action. Output and input action mention the name and incarnation of the location performing the action. Labels $\text{kill}(n, \lambda)$ and $\text{create}(n, \lambda)$ indicate respectively the killing and activation of location n in incarnation λ . Labels $\oplus n_\lambda \mapsto m$ and $\ominus n_\lambda \mapsto m$ signal respectively the creation and destruction of a link between n and m , initiated by n at incarnation λ . Finally, $n_\lambda \succ m$ signals that location n at incarnation λ holds the correct belief about location m .

The transition relation labeled by α is denoted $\xrightarrow{\alpha}$. We denote by $\xRightarrow{\tau}$ the reflexive and transitive closure of $\xrightarrow{\tau}$. For $\alpha \neq \tau$, we denote by $\xRightarrow{\alpha}$ the relation $\xRightarrow{\tau} \xrightarrow{\alpha} \xRightarrow{\tau}$.

Transitions relations $\xrightarrow{\alpha}$ in our LTS semantics are defined inductively by several sets of inference rules. The first set of rules, not shown here, contains the equivalent of all the local rules of the reduction relation in Figure 2, except for rule MSG. Fig. 5 depicts part of the the rules for concurrent and distributed primitives. Rules L-IN and L-OUT are as in the standard early instantiation-style LTS for the π -calculus [26]. Rules which are not shown contain the equivalent of local rules in Fig. 3. Rules which are not shown are derived from Figures 2 and 3 by replacing \longrightarrow by $\xrightarrow{\tau}$ (see Appendix A for the full list of rules).

Fig. 6 depicts the rules modeling the possible interactions of a context on the public part of the network. In particular the creation of a new location (L-CREATE-EXT), the killing of a location (L-KILL-EXT), the linking of two locations (L-LINK-EXT) or the unlinking of two locations (L-UNLINK-EXT). Finally, rule L-VIEW imposes equality of views of locations for two equivalent systems.

Fig. 7 depicts composition rules for the labeled transition semantics. Rules L-PAR_L and L-SYNCL have symmetric rules L-PAR_R and L-SYNCR, which are not shown. Most rules are mundane, we only discuss the non standard ones. Rules L-PAR_L is the standard rule for parallel composition allowing independent evolution of one branch of the composition. The side condition on the idle

$$\begin{array}{c}
 \text{L-CREATE-EXT} \\
 \frac{}{\Delta \triangleright N \xrightarrow{\text{create}(n, \kappa+1)} \Delta \oplus (n, \kappa+1) \triangleright N} \quad \Delta \vdash n : \text{dead} \\
 \Delta_{\mathcal{A}}(n) = -\kappa \\
 \\
 \text{L-KILL-EXT} \\
 \frac{}{\Delta \triangleright N \xrightarrow{\text{kill}(n, \lambda)} \Delta \ominus (n, \lambda) \triangleright N} \quad \Delta \vdash n_\lambda : \text{alive} \\
 \\
 \text{L-UNLINK-EXT} \\
 \frac{}{\Delta \triangleright N \xrightarrow{\ominus n_\lambda \mapsto m} \Delta \ominus n \leftrightarrow m \triangleright N} \quad \begin{array}{l} \Delta \vdash n_\lambda : \text{alive} \\ \Delta \vdash n \leftrightarrow m \end{array} \\
 \\
 \text{L-LINK-EXT} \\
 \frac{}{\Delta \triangleright N \xrightarrow{\oplus n_\lambda \mapsto m} \Delta \oplus n \leftrightarrow m \triangleright N} \quad \begin{array}{l} \Delta \vdash n_\lambda : \text{alive} \\ \Delta \not\vdash n \leftrightarrow m \end{array} \\
 \\
 \text{L-VIEW} \\
 \frac{}{\Delta \triangleright N \xrightarrow{n_\lambda \succ m} \Delta \triangleright N} \quad \begin{array}{l} \Delta \vdash n_\lambda : \text{alive} \\ \text{and } (\Delta_{\mathcal{A}}(m) = \Delta_n(m) \neq 0 \\ \text{or } \Delta_n(m) = 0) \end{array}
 \end{array}$$

Figure 6: Net Rules

$$\begin{array}{c}
 \text{L-PAR}_L \\
 \frac{\Delta \triangleright N \xrightarrow{\alpha} \nu \tilde{u}. \Delta' \triangleright N' \quad \tilde{u} \cap \text{fn}(M) = \emptyset}{\Delta \triangleright N \parallel M \xrightarrow{\alpha} \nu \tilde{u}. \Delta' \triangleright N' \parallel M} \\
 \\
 \text{L-SYNCL} \\
 \frac{\Delta \triangleright N \xrightarrow{\bar{x}(\tilde{u})@n_\lambda} \Delta \triangleright N' \quad \Delta \triangleright M \xrightarrow{x(\tilde{u})@n_\lambda} \Delta \triangleright M'}{\Delta \triangleright N \parallel M \xrightarrow{\tau} \Delta \triangleright N' \parallel M'} \\
 \\
 \text{L-RES}_O \\
 \frac{S \xrightarrow{\nu \tilde{v}. \bar{x}(\tilde{u})@n_\lambda} S' \quad w \in \tilde{u} \setminus \tilde{v}, x, n}{\nu w.S \xrightarrow{\nu w. \nu \tilde{v}. \bar{x}(\tilde{u})@n_\lambda} S'} \\
 \\
 \text{L-RES} \\
 \frac{S \xrightarrow{\alpha} S' \quad u \notin \text{fn}(\alpha)}{\nu u.S \xrightarrow{\alpha} \nu u.S'}
 \end{array}$$

Figure 7: Composition Rules

branch is required to avoid name capture when the other branch introduces a restriction. Rule L-RES_O is analogous to the classical OPEN rule in the π -calculus. What is unusual is that there is no corresponding CLOSE rule in our LTS semantics, because rule L-RES_O operates at the system level, and we have no operation for composing systems. Rule L-RES_O is a way to signal that a system is ready to send a message at a given address, possibly bearing private names in its payload.

4.3 Full Abstraction

Before presenting the main result of the paper we introduce the definition of *weak bisimilarity*.

Definition 7 (Weak Bisimilarity). *A binary relation over closed systems $\mathcal{S} \subseteq \mathbb{S}^2$ is a weak simulation iff whenever $(S, R) \in \mathcal{S}$ and $S \xrightarrow{\alpha} S'$, then $R \xRightarrow{\alpha} R'$ for some R' with $(S', R') \in \mathcal{S}$. A binary relation \mathcal{S} over closed systems is a weak bisimulation if both \mathcal{S} and \mathcal{S}^{-1} are weak simulations. Weak bisimilarity, denoted by \approx , is the largest weak bisimulation.*

Our main result states that weak bisimilarity fully characterizes weak barbed congruence.

Theorem 1 (Full Abstraction). $S \approx R$ iff $S \approx R$.

Proof. Details of the proof can be found in Appendix E. □

A few comments are in order. A consequence of our result is that the public part of the network (i.e. those nodes and links whose names are not restricted) of two weak barbed congruent systems must coincide. This is visible directly from the rules in Fig.6. In particular, incarnation numbers of public nodes must coincide. This may seem to be overly discriminative but in fact it is warranted: recovery protocols or failure handling protocols such as SWIM [7], which rely on incarnation numbers, would operate differently in dissimilar systems.

5 Examples

5.1 Deriving Erlang's Like Constructs

This sub-section shows how to implement several Erlang primitives using the ones provided by our calculus. Let us begin with the `ping` definition, which is used to test for accessibility of a remote location (m) by the current one (n). If the `ping` succeeds the view of n is updated with the knowledge of m and vice-versa, moreover the `ping` evaluates to an atom `pong` to signal success. If it fails n removes any belief it may had about the m and evaluates to `pang`.

Since in our calculus we do not have atoms we encode a slight variation of the `ping` that branches according to the result of the test. The successful one mimics the evaluation to `pong` and the negative to `pang`.

$$\text{ping } m.P \text{ else } Q \stackrel{\text{def}}{=} \nu x, t, f. \text{node}(u, \lambda). \left(\begin{array}{l} \bar{x}(f) \mid \text{spawn } m.(\text{spawn } u.\bar{x}(t)) \mid \\ x(y).\text{if } y = t \text{ then } P \text{ else remove } m.Q \end{array} \right)$$

Primitive `monitor` repeatedly tests for a remote location accessibility and continues as P when the accessibility test fails. Since the `ping` can nondeterministically fail at any point the monitor too can nondeterministically fail. We encode it as follows.

$$\text{monitor}(m).P \stackrel{\text{def}}{=} \nu x.(Q \mid !x.Q) \quad \text{where} \quad Q = \text{ping } m.\bar{x} \text{ else } P$$

Primitive `start` creates a new location l , provided that it does not exist already. Location l is said a *slave* node and the creator location is said the *master* node. The slave dies when its master dies, hence we need a monitoring process to implement this behavior. We encode it as follows.

$$\text{start}(l) \stackrel{\text{def}}{=} \text{node}(u, \lambda).\text{create } l.(\text{monitor}(u).\text{kill})$$

Primitive `stop` halts the execution of node m . We encode it as follows.

$$\text{stop}(m) \stackrel{\text{def}}{=} \text{spawn } m.\text{kill}$$

Primitive `disconnect` removes m from the view of n and vice-versa. We encode it as follows.

$$\text{disconnect}(m).P \stackrel{\text{def}}{=} \text{node}(u, \lambda).\text{remove } m.(\text{spawn } m.(\text{remove } u))$$

5.2 Behavioral Theory In Action

We begin this section by applying our behavioral theory to the motivating example in Section 2.

Example 1 (*servD* and *servDFR* are bisimilar.). *To prove $\mathbf{servD} \approx \mathbf{servDFR}$ it suffices to show a candidate bisimulation relation and then play the bisimulation game to its elements. Consider relation $\mathcal{R} = \{(\mathbf{servD}, \mathbf{servDFR})\} \cup \mathcal{S}_0 \cup \mathcal{S}_1 \cup \mathcal{S}_2$ where*

$$\begin{aligned} \mathcal{S}_0 &= \{(\mathbf{servD}, R_0) \mid \mathbf{servDFR} \xrightarrow{\tau} R_0\} \\ \mathcal{S}_1 &= \{(S_1, R_1) \mid (S_0, R_0) \in \mathcal{S}_0, S_0 \xrightarrow{\text{req}(x,y)@n_i} S_1 \wedge R_0 \xrightarrow{\text{req}(x,y)@n_i} R_1\} \\ \mathcal{S}_2 &= \{(S_2, R_2) \mid (S_1, R_1) \in \mathcal{S}_1, S_1 \xrightarrow{\bar{z}(w_\lambda)@n_i} S_1 \wedge R_1 \xrightarrow{\bar{z}(w_\lambda)@n_i} R_2\} \end{aligned}$$

Intuitively, from an external perspective, \mathbf{servD} only inputs the request and exhibits the answer; hence, we need to prove that $\mathbf{servDFR}$ is able to match those two actions. The proof is available in Appendix F. \diamond

We now present two examples inspired from [16] to show that our behavioral theory agrees with the one they present in absence of recovery.

Example 2 (Synchronous moves). *This example declines [16, Example 10] into a version that can nondeterministically fail w.r.t. the original one, due to the fact that we do not have a perfect failure detector. Consider the construct $\text{move } m.P \text{ else } Q$ which attempts to migrate P to m from the current location and if it fails, launches Q locally.*

Assuming $\Delta \vdash (n, \lambda) : \text{alive}$, the behavior could be defined as

$$\frac{\text{MOVE}}{\Delta \triangleright [\text{move } m.P \text{ else } Q]_\lambda^n \rightarrow \Delta \triangleright [P]_\lambda^m} \quad \frac{\Delta_n(m) = \lambda' \quad \Delta \vdash n_\lambda \leftrightarrow m'_\lambda}{\Delta \vdash n_\lambda \leftrightarrow m'_\lambda} \quad \frac{\text{NMOVE}}{\Delta \triangleright [\text{move } m.P \text{ else } Q]_\lambda^n \rightarrow \Delta \triangleright [Q]_\lambda^n}$$

Now one could try to implement the move primitive through the following macro

$$\text{mv } m.P \text{ else } Q \stackrel{\text{def}}{=} \nu a, b. \text{node}(u, \lambda). (\text{spawn } m.(b.P \mid \text{spawn } u.a.(\text{spawn } m.\bar{b})) \mid \bar{a} \mid \text{monitor}_a m.Q)$$

The above code sends $b.P$ to the target location and then goes back to signal the successful arrive of $b.P$, by synchronizing on the private channel a and eventually goes back to release P . The implementation uses mutual exclusion on a together with the macro $\text{monitor}_a m.Q$, which is a slight variation of the one in Section 5.1. This macro acquires the resource a before testing and releases it after the success of the test whereas it releases Q if n' is not accessible. It can be encoded as

$$\text{monitor}_a m.Q \stackrel{\text{def}}{=} \nu \text{test}. (\overline{\text{test}} \mid !\text{test}. a. \text{ping } m. (\overline{\text{test}} \mid \bar{a}) \text{ else } Q)$$

We now show that in our setting, as in [16], the primitive move is not observationally equivalent to mv . To prove this we take advantage of the contextuality property of \approx and show that

$$\nu \tilde{u}. \Delta \triangleright [\text{move } m.P \text{ else } Q]_\lambda^n \parallel [\text{unlink } n]_\kappa^m \not\approx \nu \tilde{u}. \Delta \triangleright [\text{mv } m.P \text{ else } Q]_\lambda^n \parallel [\text{unlink } n]_\kappa^m$$

where $\Delta = \langle \{n \mapsto \lambda, m \mapsto \kappa\}, \{n \leftrightarrow m\}, \{n \mapsto \hat{0}, m \mapsto \hat{0}\} \rangle$

The intuition is that the right-hand side system can reach a point where $b.P$ reached successfully m while $\text{spawn } u.a.(\text{spawn } m.\bar{b})$ manages to go back to n and also synchronizes on a . Now, if at this point the unlink reduces we get to a point where n and m cannot synchronize on b hence the two locations are alive but both P and Q are blocked. Due to the atomicity of the move the same state cannot be achieved on the left-hand side. \diamond

Example 3 (Distributed Server). Here, we rephrase [16, Example 11], where Francalanza and Hennessy show that the behavioral theory they present is able to distinguish a distributed server only able to reach its backend by a direct connection and one that, in addition to the direct connection, has also an indirect connection that goes through a third locality.

The two systems are the following.

$$\begin{aligned} \mathbf{servFHD} &\Leftarrow \nu \text{data}, w_y. \Delta \triangleright \left(\begin{array}{l} [\text{req}(x, y). \text{spawn } n'. \overline{\text{data}}\langle x, y \rangle]^n \parallel \\ [\text{data}(x, y). \text{spawn } n. \overline{x}\langle w_y \rangle]^{n'} \end{array} \right) \\ \mathbf{sFHD2Rt} &\Leftarrow \nu \text{data}, w_y. \Delta \triangleright \left(\begin{array}{l} \left[\text{req}(x, y). \nu \text{sync}. \left(\begin{array}{l} \text{spawn } n'. \overline{\text{data}}\langle \text{sync}, x \rangle \mid \\ \text{spawn } n''. \text{spawn } n'. \overline{\text{data}}\langle \text{sync}, x \rangle \mid \\ \text{sync}(x). \overline{y}\langle x \rangle \end{array} \right) \right]^n \parallel \\ \left[\text{data}(x, y). \left(\begin{array}{l} \text{spawn } n'. \overline{x}\langle f(y) \rangle \mid \\ \text{spawn } n''. \text{spawn } n. \overline{x}\langle w_y \rangle \end{array} \right) \right]^{n'} \end{array} \right) \end{aligned}$$

Now, \mathbf{servD} and $\mathbf{sFHD2Rt}$, like in [16], can be distinguished by the following context

$$C \equiv \cdot \parallel [\text{unlink } n]^{n'} \parallel [\overline{\text{req}}\langle z, h \rangle]^n$$

as \mathbf{servD} would stop working after the break reduces, while $\mathbf{servD2Rt}$ would keep working correctly since it could route the request through m .

◇

6 Related work and conclusion

We have presented in this paper a distributed π -calculus with location and link crash failures and recoveries. The calculus basic constructs, including incarnation numbers to distinguish different versions of the same location across recoveries and local imperfect knowledge, were inspired mainly by the Erlang programming language and environment. To the best of our knowledge, this is the first work that combines these different features, and the first to deal with recovery without relying on some form of checkpointing (as in Erlang-style systems).

As mentioned in the Introduction, there are only a few works in the literature proposing a process calculus analysis of distributed systems with crash failures and recoveries. We discussed the work by Berger and Honda [3] and by Bocchi et al. [4] in the Introduction. The work by Fournet et al. on the join calculus [15] shows how to extend the join calculus with primitives for crash failures and recoveries but does not take into account links and link failures and does not present a behavioral theory for these extensions of the join calculus. The work by Amadio [1] on the π_{1l} -calculus presents an asynchronous π -calculus with unique receivers, located processes, and location failures. It develops a behavioral theory for this calculus by translation of the π_{1l} -calculus into the π_1 calculus (an asynchronous π -calculus with unique receivers), but it relies on perfect failure detectors, does not support links and link failures, and recovery just consists in restarting a stopped process in its exact state at the moment of failure. In its discussion of weaker failure detectors, it does indicate how an extension could support a local view (of failed locations) but it does not elaborate the corresponding calculus and behavioral theory.

Our main inspiration for this paper was the work by Francalanza and Hennessy [16] for their handling of node and link failures, and their behavioral theory. Apart from dealing with recovery, which they do not consider, our development is markedly different from theirs. We have opted for a simpler handling of scope extrusion, with simpler labels in our LTS semantics (just names instead of complex information about the network, including links and liveness of locations),

and no need to make explicit the partial view of a network available to an observer. In effect, we have opted for the alternative design choice they discussed when contrasting their work to that of De Nicola et al. [25]: keep simple labels in the LTS and opt for a classical handling of scope extrusion, possibly at the expense of larger bisimulation relations. We believe the gain in simplicity of exposition and understanding, coupled with the simple handling of recovery our approach allows, is well worth it. Also, we introduce explicit local views, which correspond to the belief that a location has of its neighbours and their current incarnation. Our local views are handled similarly as in Erlang – in particular they may not reflect the current state of the network –, and they have no equivalent in [16]. The notion of partial views for observers in [16] is different: it is a way to filter out information contained in their complex labels so as to obtain full abstraction, and it plays no role in the operational semantics of their model.

Formal models for distributed systems with failures can also be found in recent verification tools for distributed algorithms and distributed systems such as Disel [27], Gobra [30], Perennial [5], Psync [8], TLC [18], Verdi [29]. They can either rely on a specific language (such as Gobra, for Go programs), or domain specific languages for formally specifying algorithms (such as PSync or Disel), or be more general purpose, relying on a mixture of logic and more operational models (such as Perennial, Verdi or TLC). Verdi in particular is interesting for it supports a variety of failure models, including a model of crash failures and recoveries which is quite close to ours, and makes use of simulation relations in its proof techniques. However, to the best of our knowledge, they (Verdi included) do not provide as we do a compositional theory of system equivalence in presence of crash failures and recoveries.

The calculus presented in this paper provides us with a basis for further studies. It would certainly be interesting to further expand it to cater for other failure models, including the kinds of grey failures tackled by Bocchi et al. [4], probably making it parametric in failure models along the lines of Verdi [29]. In its current form, we think it is close enough to the behavior of Erlang systems, and its primitives sufficient to account for a significant subset of Erlang failure handling constructs, to allow us to study reversible debugging for distributed Erlang systems with failures and recoveries, building on the substantial amount of work developed in the past ten years around reversibility and reversible debugging in Erlang [22, 23].

References

- [1] Roberto M. Amadio. An asynchronous model of locality, failure and process mobility. In David Garlan and Daniel Le Métayer, editors, *Coordination Languages and Models, Second International Conference, COORDINATION '97, Berlin, Germany, September 1-3, 1997, Proceedings*, volume 1282 of *Lecture Notes in Computer Science*, pages 374–391. Springer, 1997.
- [2] Joe Armstrong. *Programming Erlang – Software for a Concurrent World*. The Pragmatic Bookshelf, 2007.
- [3] Martin Berger and Kohei Honda. The two-phase commitment protocol in an extended pi-calculus. In Luca Aceto and Björn Victor, editors, *7th International Workshop on Expressiveness in Concurrency, EXPRESS 2000, Satellite Workshop of CONCUR 2000, State College, PA, USA, August 21, 2000*, volume 39 of *Electronic Notes in Theoretical Computer Science*, pages 21–46. Elsevier, 2000.
- [4] Laura Bocchi, Julien Lange, Simon Thompson, and A. Laura Voinea. A model of actors and grey failures. In Maurice H. ter Beek and Marjan Sirjani, editors, *Coordination Models and Languages - 24th IFIP WG 6.1 International Conference, COORDINATION 2022, Held as Part of the 17th International Federated Conference on Distributed Computing Techniques, DisCoTec 2022, Lucca, Italy, June 13-17, 2022, Proceedings*, volume 13271 of *Lecture Notes in Computer Science*, pages 140–158. Springer, 2022.
- [5] Tej Chajed, Joseph Tassarotti, M. Frans Kaashoek, and Nikolai Zeldovich. Verifying concurrent, crash-safe systems with perennial. In Tim Brecht and Carey Williamson, editors, *Proceedings of the 27th ACM Symposium on Operating Systems Principles, SOSP 2019, Huntsville, ON, Canada, October 27-30, 2019*, pages 243–258. ACM, 2019.
- [6] Om P. Damani, Ashis Tarafdar, and Vijay K. Garg. Optimistic recovery in multi-threaded distributed systems. In *The Eighteenth Symposium on Reliable Distributed Systems, SRDS 1999, Lausanne, Switzerland, October 19-22, 1999, Proceedings*, pages 234–243. IEEE Computer Society, 1999.
- [7] Abhinandan Das, Indranil Gupta, and Ashish Motivala. SWIM: scalable weakly-consistent infection-style process group membership protocol. In *2002 International Conference on Dependable Systems and Networks (DSN 2002), 23-26 June 2002, Bethesda, MD, USA, Proceedings*, pages 303–312. IEEE Computer Society, 2002.
- [8] Cezara Dragoi, Thomas A. Henzinger, and Damien Zufferey. Psync: a partially synchronous language for fault-tolerant distributed algorithms. In Rastislav Bodik and Rupak Majumdar, editors, *Proceedings of the 43rd Annual ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, POPL 2016, St. Petersburg, FL, USA, January 20 - 22, 2016*, pages 400–415. ACM, 2016.
- [9] E. N. Elnozahy, Lorenzo Alvisi, Yi-Min Wang, and David B. Johnson. A survey of rollback-recovery protocols in message-passing systems. *ACM Comput. Surv.*, 34(3):375–408, 2002.
- [10] E. N. Elnozahy and Willy Zwaenepoel. Manetho: Transparent rollback-recovery with low overhead, limited rollback, and fast output commit. *IEEE Trans. Computers*, 41(5):526–531, 1992.
- [11] URL: https://www.erlang.org/doc/apps/erts/erl_ext_dist.html#NEW_PID_EXT.

- [12] URL: erlang-solutions.com/blog/which-companies-are-using-erlang-and-why-mytopdogstatus.html, 2019.
- [13] URL: <https://github.com/gfabbretti8/erlang-experiments>, 2023.
- [14] Cédric Fournet and Georges Gonthier. The reflexive CHAM and the join-calculus. In Hans-Juergen Boehm and Guy L. Steele Jr., editors, *Conference Record of POPL'96: The 23rd ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages, Papers Presented at the Symposium, St. Petersburg Beach, Florida, USA, January 21-24, 1996*, pages 372–385. ACM Press, 1996.
- [15] Cédric Fournet, Georges Gonthier, Jean-Jacques Lévy, Luc Maranget, and Didier Rémy. A calculus of mobile agents. In Ugo Montanari and Vladimiro Sassone, editors, *CONCUR '96, Concurrency Theory, 7th International Conference, Pisa, Italy, August 26-29, 1996, Proceedings*, volume 1119 of *Lecture Notes in Computer Science*, pages 406–421. Springer, 1996.
- [16] Adrian Francalanza and Matthew Hennessy. A theory of system behaviour in the presence of node and link failure. *Inf. Comput.*, 206(6):711–759, 2008.
- [17] Yu Gao, Wensheng Dou, Feng Qin, Chushu Gao, Dong Wang, Jun Wei, Ruirui Huang, Li Zhou, and Yongming Wu. An empirical study on crash recovery bugs in large-scale distributed systems. In Gary T. Leavens, Alessandro Garcia, and Corina S. Pasareanu, editors, *Proceedings of the 2018 ACM Joint Meeting on European Software Engineering Conference and Symposium on the Foundations of Software Engineering, ESEC/SIGSOFT FSE 2018, Lake Buena Vista, FL, USA, November 04-09, 2018*, pages 539–550. ACM, 2018.
- [18] Jeremiah Griffin, Mohsen Lesani, Narges Shadab, and Xizhe Yin. TLC: temporal logic of distributed components. *Proc. ACM Program. Lang.*, 4(ICFP):123:1–123:30, 2020.
- [19] Indranil Gupta, Tushar Deepak Chandra, and Germán S. Goldszmidt. On scalable and efficient distributed failure detectors. In Ajay D. Kshemkalyani and Nir Shavit, editors, *Proceedings of the Twentieth Annual ACM Symposium on Principles of Distributed Computing, PODC 2001, Newport, Rhode Island, USA, August 26-29, 2001*, pages 170–179. ACM, 2001.
- [20] Márk Jelasity, Spyros Voulgaris, Rachid Guerraoui, Anne-Marie Kermarrec, and Maarten van Steen. Gossip-based peer sampling. *ACM Trans. Comput. Syst.*, 25(3):8, 2007.
- [21] Sasa Juric. *Elixir in Action*. Manning, 2015.
- [22] Ivan Lanese, Naoki Nishida, Adrián Palacios, and Germán Vidal. A theory of reversibility for Erlang. *Journal of Logical and Algebraic Methods in Programming*, 100:71 – 97, 2018.
- [23] Ivan Lanese, Adrián Palacios, and Germán Vidal. Causal-consistent replay reversible semantics for message passing concurrent programs. *Fundam. Informaticae*, 178(3):229–266, 2021.
- [24] Robin Milner and Davide Sangiorgi. Barbed bisimulation. In Werner Kuich, editor, *Automata, Languages and Programming, 19th International Colloquium, ICALP92, Vienna, Austria, July 13-17, 1992, Proceedings*, volume 623 of *Lecture Notes in Computer Science*, pages 685–695. Springer, 1992.

- [25] Rocco De Nicola, Daniele Gorla, and Rosario Pugliese. Basic observables for a calculus for global computing. *Inf. Comput.*, 205(10):1491–1525, 2007.
- [26] Davide Sangiorgi and David Walker. *The Pi-Calculus - a theory of mobile processes*. Cambridge University Press, 2001.
- [27] Ilya Sergey, James R. Wilcox, and Zachary Tatlock. Programming and proving with distributed protocols. *Proc. ACM Program. Lang.*, 2(POPL):28:1–28:30, 2018.
- [28] Robert E. Strom and Shaula Yemini. Optimistic recovery in distributed systems. *ACM Trans. Comput. Syst.*, 3(3):204–226, 1985.
- [29] James R. Wilcox, Doug Woos, Pavel Panchekha, Zachary Tatlock, Xi Wang, Michael D. Ernst, and Thomas E. Anderson. Verdi: a framework for implementing and formally verifying distributed systems. In David Grove and Stephen M. Blackburn, editors, *Proceedings of the 36th ACM SIGPLAN Conference on Programming Language Design and Implementation, Portland, OR, USA, June 15-17, 2015*, pages 357–368. ACM, 2015.
- [30] Felix A. Wolf, Linard Arqint, Martin Clochard, Wytse Oortwijn, João Carlos Pereira, and Peter Müller. Gobra: Modular specification and verification of go programs. In Alexandra Silva and K. Rustan M. Leino, editors, *Computer Aided Verification - 33rd International Conference, CAV 2021, Virtual Event, July 20-23, 2021, Proceedings, Part I*, volume 12759 of *Lecture Notes in Computer Science*, pages 367–379. Springer, 2021.

A Notations and rules

For ease of reference, we gather in this section all the notations and inference rules used in the paper.

A.1 Notations

We assume given mutually disjoint infinite denumerable sets \mathbf{C} , \mathbf{N} and \mathbf{I} . \mathbf{C} is the set of *channel names*, \mathbf{N} is the set of *location names*, and \mathbf{I} is the set of incarnation number variables. We use the set of integers \mathbb{Z} as the set of *incarnation numbers*. An incarnation number is paired with a location name for recovery purposes, to distinguish the current instance of a location from its past failed instances. We denote by \mathbf{N}^\odot the set $\mathbf{N} \cup \{\odot\}$, where $\odot \notin \mathbf{N}$. As in the π -calculus, channel names can be free or bound in terms. The same holds for location names. Incarnation variables can be bound, but not incarnation numbers. We denote by \tilde{u} a finite (possibly empty) tuple of elements. We write $T\{\tilde{v}/\tilde{u}\}$ for the usual capture-avoiding substitution of elements of \tilde{u} by elements of \tilde{v} in term T , assuming tuples \tilde{u} and \tilde{v} have the same arity. We write u, \tilde{v} or \tilde{v}, u for the tuple \tilde{v} extended with element u as first or last element. Abusing notation, we sometimes identify a tuple \tilde{u} with the set of its elements. We denote by \mathbb{N}^+ the set of strictly positive integers (by definition $0 \notin \mathbb{N}^+$), and by \mathbb{N} the set of positive integers ($0 \in \mathbb{N}$). We denote by $\hat{\mathbf{0}}$ the function $\hat{\mathbf{0}} : \mathbf{N}^\odot \rightarrow \mathbb{Z}$ that maps any $n \in \mathbf{N}$ to 0 and \odot to 1.

A.2 Calculus syntax and alpha-conversion on systems

$$\begin{aligned}
 P, Q ::= & \mathbf{0} \mid \bar{x}(\tilde{u}).P \mid x(\tilde{v}).P \mid !x(\tilde{v}).P \mid \nu w.P \mid \text{if } r = s \text{ then } P \text{ else } Q \mid P \mid Q \\
 & \text{node}(n, \lambda).P \mid \text{remove } n.P \mid \text{spawn } n.P \mid \text{kill} \mid \\
 & \text{create } n.P \mid \text{link } n.P \mid \text{unlink } n.P \\
 L, M, N ::= & \mathbf{0} \mid [P]_\lambda^n \mid N \parallel M \\
 S, R ::= & \Delta \triangleright N \mid \nu w.S
 \end{aligned}$$

where: $\tilde{u}, r, s \subset \mathbf{C} \cup \mathbf{N}^\odot \cup \mathbf{I} \cup \mathbb{Z}$ $\tilde{v} \subset \mathbf{C} \cup \mathbf{N} \cup \mathbf{I}$ $w \in \mathbf{C} \cup \mathbf{N}$ $x \in \mathbf{C}$ $n \in \mathbf{N}$ $\lambda \in \mathbf{I}$

Free names in processes and configurations are defined inductively as follows:

$$\begin{aligned}
 \text{fn}(\mathbf{0}) &= \emptyset \\
 \text{fn}(\bar{x}(\tilde{u}).P) &= \tilde{u} \cup \{x\} \cup \text{fn}(P) \\
 \text{fn}(x(\tilde{v}).P) &= \{x\} \cup \text{fn}(P) \setminus \tilde{v} \\
 \text{fn}(!x(\tilde{v}).P) &= \{x\} \cup \text{fn}(P) \setminus \tilde{v} \\
 \text{fn}(\nu w.P) &= \text{fn}(P) \setminus \{w\} \\
 \text{fn}(\text{if } r = s \text{ then } P \text{ else } Q) &= \text{fn}(P) \cup \text{fn}(Q) \cup \{r, s\} \\
 \text{fn}(P \mid Q) &= \text{fn}(P) \cup \text{fn}(Q) \\
 \text{fn}(\text{node}(n, \lambda).P) &= \text{fn}(P) \setminus \{n\} \\
 \text{fn}(\text{remove } n.P) &= \text{fn}(P) \cup \{n\} \\
 \text{fn}(\text{spawn } n.P) &= \text{fn}(P) \cup \{n\} \\
 \text{fn}(\text{create } n.P) &= \text{fn}(P) \cup \{n\} \\
 \text{fn}([P]_\lambda^n) &= \text{fn}(P) \cup \{n\} \\
 \text{fn}(N \parallel M) &= \text{fn}(N) \cup \text{fn}(M)
 \end{aligned}
 \qquad
 \begin{aligned}
 \text{fn}(\text{kill}) &= \emptyset \\
 \text{fn}(\text{link } n) &= \text{fn}(\text{unlink } n) = \{n\}
 \end{aligned}$$

\mathbb{L} denotes the set of configurations. \mathbb{S} denotes the set of systems. A *network* Δ is a tuple $\langle \mathcal{A}, \mathcal{L}, \mathcal{V} \rangle$ where:

- \mathcal{A} is a function $\mathcal{A} : \mathbb{N}^\odot \rightarrow \mathbb{Z}$ such that $\mathcal{A}(\odot) = 1$ and such that the set $\text{supp}(\mathcal{A}) \stackrel{\text{def}}{=} \{n \in \mathbb{N} \mid \mathcal{A}(n) \neq 0\}$ is finite. Function \mathcal{A} records three types of information on locations. If $\mathcal{A}(n) = \lambda \in \mathbb{N}^+$, then location n is alive and its current incarnation number is λ . If $\mathcal{A}(n) = -\lambda, \lambda \in \mathbb{N}^+$ then location n has been killed and its last incarnation number while alive was λ . If $\mathcal{A}(n) = 0$, then there is no location n in the network, alive or not.
- $\mathcal{L} \subseteq \mathbb{N}^\odot \times \mathbb{N}^\odot$ is the set of links between locations. \mathcal{L} is a finite symmetric binary relation over location names such that $\text{dom}(\mathcal{L}) \stackrel{\text{def}}{=} \{n \in \mathbb{N}^\odot \mid \exists m, (n, m) \in \mathcal{L}\}$ is finite.
- $\mathcal{V} : \mathbb{N}^\odot \rightarrow (\mathbb{N}^\odot \rightarrow \mathbb{N})$ is a function that maps location names to their *local view* which is such that the set $\text{supp}(\mathcal{V}) \stackrel{\text{def}}{=} \{n \in \mathbb{N} \mid \mathcal{V}(n) \neq \hat{\mathbf{0}}\}$ is finite and $\text{supp}(\mathcal{V}) \subseteq \text{supp}(\mathcal{A})$. The local view of a location n is a function $\mathcal{V}(n) : \mathbb{N}^\odot \rightarrow \mathbb{N}$ such that the set $\{m \in \mathbb{N} \mid \mathcal{V}(n)(m) \neq 0\}$ is finite. If $\mathcal{V}(n)(m) = \kappa \in \mathbb{N}^+$, then location m in its incarnation κ is believed by n to be alive. If $\mathcal{V}(n)(m) = 0$, then location n holds no belief on the status of location m .

$\Delta_{\mathcal{A}}, \Delta_{\mathcal{L}},$ and $\Delta_{\mathcal{V}}$ denote the components of a network representation Δ . Notations for extracting information from Δ :

- $\Delta \vdash n_\lambda$: alive if $\Delta_{\mathcal{A}}(n) = \lambda$ and $\lambda \in \mathbb{N}^+$.
- $\Delta \vdash n$: dead if $\Delta_{\mathcal{A}}(n) \notin \mathbb{N}^+$
- $\Delta \vdash n \leftrightarrow m$ if $(n, m) \in \Delta_{\mathcal{L}}$
- $\Delta \vdash n_\lambda \rightsquigarrow m_\kappa$ if $(n, m) \in \Delta_{\mathcal{L}}, \Delta \vdash n_\lambda$: alive and $\Delta \vdash m_\kappa$: alive

Update operations over a network Δ :

- $\Delta \oplus n \leftrightarrow m = \langle \Delta_{\mathcal{A}}, \Delta_{\mathcal{L}} \cup \{(n, m), (m, n)\}, \Delta_{\mathcal{V}} \rangle$
- $\Delta \ominus n \leftrightarrow m = \langle \Delta_{\mathcal{A}}, \Delta_{\mathcal{L}} \setminus \{(n, m), (m, n)\}, \Delta_{\mathcal{V}} \rangle$
- $\Delta \oplus (n, \lambda) = \langle \Delta_{\mathcal{A}}[n \mapsto \lambda], \Delta_{\mathcal{L}} \cup \{(n, n)\}, \Delta_{\mathcal{V}}[n \mapsto \hat{\mathbf{0}}] \rangle$
- $\Delta \ominus (n, \lambda) = \langle \Delta_{\mathcal{A}}[n \mapsto -\lambda], \Delta_{\mathcal{L}}, \Delta_{\mathcal{V}} \rangle$
- $\Delta \oplus n \succ (m, \lambda) = \langle \Delta_{\mathcal{A}}, \Delta_{\mathcal{L}}, \Delta_{\mathcal{V}}[n \mapsto \Delta_{\mathcal{V}}(n)[m \mapsto \lambda]] \rangle$, if $n \neq m$
- $\Delta \ominus n \succ m = \langle \Delta_{\mathcal{A}}, \Delta_{\mathcal{L}}, \Delta_{\mathcal{V}}[n \mapsto \Delta_{\mathcal{V}}(n)[m \mapsto 0]] \rangle$, if $n \neq m$
- $\Delta \ominus n \succ n = \Delta \oplus n \succ (n, \lambda) = \Delta$

Free names of networks and systems are defined inductively as follows:

$$\begin{aligned} \text{fn}(\Delta) &= \text{supp}(\Delta_{\mathcal{A}}) \cup \text{dom}(\Delta_{\mathcal{L}}) \\ \text{fn}(\Delta \triangleright N) &= \text{fn}(\Delta) \cup \text{fn}(N) \\ \text{fn}(\nu \tilde{u}.S) &= \text{fn}(S) \setminus \tilde{u} \end{aligned}$$

To define alpha-conversion on systems, we define capture-avoiding substitution on networks. A capture avoiding substitution $\{v/u\}$ on network Δ is defined as follows: if $\Delta = (\mathcal{A}, \mathcal{L}, \mathcal{V})$, then $\Delta\{v/u\} = (\mathcal{A}\{v/u\}, \mathcal{L}\{v/u\}, \mathcal{V}\{v/u\})$ where:

$$\begin{aligned}\mathcal{A}\{v/u\} &= \begin{cases} \mathcal{A}[v \mapsto \mathcal{A}(u)][u \mapsto 0] & \text{if } u, v \in \mathbf{N} \text{ and } v \notin \text{supp}(\mathcal{A}) \\ \mathcal{A} & \text{otherwise} \end{cases} \\ \mathcal{L}\{v/u\} &= \begin{cases} (\mathcal{L} \setminus \mathcal{L}(u)) \cup \mathcal{L}(u)\{v/u\} & \text{if } u, v \in \mathbf{N} \text{ and } v \notin \text{supp}(\mathcal{A}) \\ \mathcal{L} & \text{otherwise} \end{cases} \\ \mathcal{V}\{v/u\} &= \begin{cases} \mathcal{V}[v \mapsto \mathcal{V}(u)][u \mapsto \hat{0}] & \text{if } u, v \in \mathbf{N} \text{ and } v \notin \text{supp}(\mathcal{A}) \\ \mathcal{V} & \text{otherwise} \end{cases}\end{aligned}$$

with:

$$\begin{aligned}\mathcal{L}(n) &= \{(a, b) \in \mathcal{L} \mid a = n \text{ or } b = n\} \\ \mathcal{L}(n)\{m/n\} &= \{(a\{m/n\}, b\{m/n\}) \mid (a, b) \in \mathcal{L}(n)\} \\ a\{m/n\} &= \begin{cases} m & \text{if } a = n \\ a & \text{otherwise} \end{cases}\end{aligned}$$

The effect of a capture avoiding substitution $\{v/u\}$ on a system $\Delta \triangleright N$ is defined inductively as follows. We define the set occ of name occurrences of a system as follows:

$$\begin{aligned}(\Delta \triangleright N)\{v/u\} &= \Delta\{v/u\} \triangleright N\{v/u\} \\ (\nu w.S)\{v/u\} &= \begin{cases} \nu w.S\{v/u\} & \text{if } u, v \neq w \\ \perp & \text{otherwise} \end{cases}\end{aligned}$$

Equality modulo alpha-conversion $=_\alpha$ on systems is now defined as the smallest equivalence relation on systems defined by the following rules:

$$\frac{S\{v/u\} \neq \perp}{\nu u.S =_\alpha \nu v.S\{v/u\}} \qquad \frac{S =_\alpha T}{\nu u.S =_\alpha \nu u.T}$$

A.3 Reduction semantics

The operational semantics of our calculus is defined via a reduction semantics given by a binary relation $\longrightarrow \subseteq \mathbb{S} \times \mathbb{S}$ between closed systems, and a structural congruence relation $\equiv \subseteq \mathbb{S}^2 \cup \mathbb{L}^2$, that is a binary equivalence relation between systems and between configurations. Evaluation contexts are ‘‘systems with a hole’’ defined by the following grammar:

$$\mathbb{C} ::= \nu \tilde{w}.\Delta \triangleright \mathbb{E} \qquad \mathbb{E} ::= \cdot \mid (N \parallel \mathbb{E}) \qquad \text{where:} \qquad \tilde{w} \subset \mathbf{N} \cup \mathbf{C}$$

Relation \equiv is the smallest equivalence relation defined by the rules in Fig. 8, where $=_\alpha$ stands for equality up to alpha-conversion, $M, N, L \in \mathbb{L}$, and $S, R \in \mathbb{S}$.

Reduction relation \longrightarrow is defined by the rules in Fig. 9, where:

$$\Delta_n(m) = \begin{cases} \kappa & \text{if } n = m \text{ and } \Delta_{\mathcal{A}}(n) = \kappa \\ \kappa & \text{if } n \neq m \text{ and } \Delta_{\mathcal{V}}(n)(m) = \kappa \\ \kappa & \text{if } n \neq m \text{ and } \Delta_{\mathcal{V}}(n)(m) = 0 \text{ and } \Delta_{\mathcal{A}}(m) = \kappa \\ 0 & \text{otherwise} \end{cases}$$

$$\begin{array}{c}
\text{[S.PAR.C]} \ N \parallel M \equiv M \parallel N \qquad \text{[S.PAR.A]} \ (L \parallel M) \parallel N \equiv L \parallel (M \parallel N) \qquad \text{[S.PAR.N]} \ (N \parallel \mathbf{0}) \equiv N \\
\text{[S.RES.C]} \ \nu u. \nu v. S \equiv \nu v. \nu u. S \qquad \text{[S.RES.NIL]} \ \frac{}{\nu u. S \equiv S} \ u \notin \text{fn}(S) \qquad \text{[S.}\alpha\text{]} \ \frac{S =_\alpha R}{S \equiv R} \\
\text{[S.CTX]} \ \frac{N \equiv M}{\mathbf{C}[N] \equiv \mathbf{C}[M]}
\end{array}$$

Figure 8: Structural Congruence Rules

A.4 LTS semantics

Labels α in our LTS semantics take the following forms:

$$\alpha ::= \tau \mid \nu \tilde{w}. \bar{x}(\tilde{u})@n_\lambda \mid x(\tilde{u})@n_\lambda \mid \text{kill}(n, \lambda) \mid \text{create}(n, \lambda) \mid \oplus n_\lambda \mapsto m \mid \ominus n_\lambda \mapsto m \mid n_\lambda \succ m$$

Free names in labels are defined as follows:

$$\begin{array}{lll}
\text{fn}(\tau) = \emptyset & \text{fn}(\nu \tilde{w}. \bar{x}(\tilde{u})@n_\lambda) = (\tilde{u} \cup \{x, n\}) \setminus \tilde{w} & \text{fn}(x(\tilde{u})@n_\lambda) = \tilde{u} \cup \{x, n\} \\
\text{fn}(\text{kill}(n, \lambda)) = \{n\} & \text{fn}(\text{create}(n, \lambda)) = \{n\} & \\
\text{fn}(\oplus n_\lambda \mapsto m) = \{n, m\} & \text{fn}(\ominus n_\lambda \mapsto m) = \{n, m\} & \text{fn}(n_\lambda \succ m) = \{n, m\}
\end{array}$$

Labelled transition relations $\xrightarrow{\alpha}$ of our LTS semantics are defined by the rules in Fig. 10 and 11.

Assuming $\Delta \vdash n_\lambda : \text{alive}$

MSG

$$\frac{}{\Delta \triangleright [\bar{x}\langle \tilde{v} \rangle . Q]_\lambda^n \parallel [x(\tilde{u}) . P]_\lambda^n \longrightarrow \Delta \triangleright [P\{\tilde{v}/\tilde{u}\}]_\lambda^n \parallel [Q]_\lambda^n}$$

NEW

$$\frac{}{\Delta \triangleright [\nu u . P]_\lambda^n \longrightarrow \nu u . \Delta \triangleright [P]_\lambda^n} \quad u \notin \text{fn}(\Delta) \cup \{n, \lambda\}$$

IF-NEQ

$$\frac{}{\Delta \triangleright [\text{if } r = s \text{ then } P \text{ else } Q]_\lambda^n \longrightarrow \Delta \triangleright [Q]_\lambda^n} \quad r \neq s$$

IF-EQ

$$\frac{}{\Delta \triangleright [\text{if } r = r \text{ then } P \text{ else } Q]_\lambda^n \longrightarrow \Delta \triangleright [P]_\lambda^n}$$

SPAWN-S

$$\frac{\Delta_n(m) = \kappa \quad \Delta \vdash n_\lambda \leftrightarrow m_\kappa}{\Delta \triangleright [\text{spawn } m . P]_\lambda^n \longrightarrow \Delta \oplus m \succ (n, \lambda) \triangleright [P]_\kappa^m}$$

SPAWN-F

$$\frac{\Delta_n(m) = \kappa \quad \Delta \not\vdash n_\lambda \leftrightarrow m_\kappa}{\Delta \triangleright [\text{spawn } m . P]_\lambda^n \longrightarrow \Delta \ominus n \succ m \triangleright \mathbf{0}}$$

CREATE-S

$$\frac{\Delta \vdash m : \text{dead} \quad \Delta_{\mathcal{A}}(m) = -\kappa}{\Delta \triangleright [\text{create } m . P]_\lambda^n \longrightarrow \Delta \oplus (m, \kappa + 1) \triangleright [P]_{\kappa+1}^m}$$

CREATE-F

$$\frac{\Delta \not\vdash m : \text{dead}}{\Delta \triangleright [\text{create } m . P]_\lambda^n \longrightarrow \Delta \triangleright \mathbf{0}}$$

RES

$$\frac{S \longrightarrow S'}{\nu u . S \longrightarrow \nu u . S'}$$

BANG

$$\frac{}{\Delta \triangleright [!x(\tilde{u}) . P]_\lambda^n \longrightarrow \Delta \triangleright [x(\tilde{u}) . (P \mid !x(\tilde{u}) . P)]_\lambda^n}$$

FORK

$$\frac{}{\Delta \triangleright [P \mid Q]_\lambda^n \longrightarrow \Delta \triangleright [P]_\lambda^n \parallel [Q]_\lambda^n}$$

NODE

$$\frac{}{\Delta \triangleright [\text{node}(m, \kappa) . P]_\lambda^n \longrightarrow \Delta \triangleright [P\{m, \kappa/n, \lambda\}]_\lambda^n}$$

REMOVE

$$\frac{}{\Delta \triangleright [\text{remove } m . P]_\lambda^n \longrightarrow \Delta \ominus n \succ m \triangleright [P]_\lambda^n}$$

LINK

$$\frac{\Delta \not\vdash n \leftrightarrow m}{\Delta \triangleright [\text{link } m . P]_\lambda^n \longrightarrow \Delta \oplus n \leftrightarrow m \triangleright [P]_\lambda^n}$$

UNLINK

$$\frac{\Delta \vdash n \leftrightarrow m}{\Delta \triangleright [\text{unlink } m . P]_\lambda^n \longrightarrow \Delta \ominus n_\lambda \leftrightarrow m \triangleright [P]_\lambda^n}$$

KILL

$$\frac{n \neq \odot}{\Delta \triangleright [\text{kill}]_\lambda^n \longrightarrow \Delta \ominus (n, \lambda) \triangleright \mathbf{0}}$$

PAR

$$\frac{\Delta \triangleright N \longrightarrow \nu \tilde{u} . \Delta' \triangleright N' \quad \tilde{u} \cap \text{fn}(M) = \emptyset}{\Delta \triangleright N \parallel M \longrightarrow \nu \tilde{u} . \Delta' \triangleright N' \parallel M}$$

STR

$$\frac{S \equiv S' \quad S' \longrightarrow R' \quad R' \equiv R}{S \longrightarrow R}$$

Figure 9: Reduction Rules

Assuming $\Delta \vdash n_\lambda : \text{alive}$

<p>L-BANG</p> $\frac{}{\Delta \triangleright [!x(\tilde{u}).P]_\lambda^n \xrightarrow{\tau} \Delta \triangleright [x(\tilde{u}).(P \mid !x(\tilde{u}).P)]_\lambda^n}$ <p>L-NEW</p> $\frac{u \notin \text{fn}(\Delta) \cup \{n, \lambda\}}{\Delta \triangleright [\nu u.P]_\lambda^n \xrightarrow{\tau} \nu u.\Delta \triangleright [P]_\lambda^n}$ <p>L-IF-NEQ</p> $\frac{r \neq s}{\Delta \triangleright [\text{if } r = s \text{ then } P \text{ else } Q]_\lambda^n \xrightarrow{\tau} \Delta \triangleright [Q]_\lambda^n}$ <p>L-IF-EQ</p> $\frac{}{\Delta \triangleright [\text{if } r = r \text{ then } P \text{ else } Q]_\lambda^n \xrightarrow{\tau} \Delta \triangleright [P]_\lambda^n}$ <p>L-IN</p> $\frac{}{\Delta \triangleright [x(\tilde{v}).P]_\lambda^n \xrightarrow{x(\tilde{u})@n_\lambda} \Delta \triangleright [P\{\tilde{u}/\tilde{v}\}]_\lambda^n}$ <p>L-SPAWN-S</p> $\frac{\Delta_n(m) = \kappa \quad \Delta \vdash n_\lambda \leftrightarrow m_\kappa}{\Delta \triangleright [\text{spawn } m.P]_\lambda^n \xrightarrow{\tau} \Delta \oplus m \succ (n, \lambda) \triangleright [P]_{\kappa}^m}$ <p>L-SPAWN-F</p> $\frac{\Delta_n(m) = \kappa \quad \Delta \not\vdash n_\lambda \leftrightarrow m_\kappa}{\Delta \triangleright [\text{spawn } m.P]_\lambda^n \xrightarrow{\tau} \Delta \ominus n \succ m \triangleright \mathbf{0}}$ <p>L-CREATE-S</p> $\frac{\Delta \vdash m : \text{dead} \quad \Delta_{\mathcal{A}}(m) = -\kappa}{\Delta \triangleright [\text{create } m.P]_\lambda^n \xrightarrow{\tau} \Delta \oplus (m, \kappa + 1) \triangleright [P]_{\kappa+1}^m}$	<p>L-FORK</p> $\frac{}{\Delta \triangleright [P \mid Q]_\lambda^n \xrightarrow{\tau} \Delta \triangleright [P]_\lambda^n \parallel [Q]_\lambda^n}$ <p>L-NODE</p> $\frac{}{\Delta \triangleright [\text{node}(m, \kappa).P]_\lambda^n \xrightarrow{\tau} \Delta \triangleright [P\{m, \kappa/n, \lambda\}]_\lambda^n}$ <p>L-REMOVE</p> $\frac{}{\Delta \triangleright [\text{remove } m.P]_\lambda^n \xrightarrow{\tau} \Delta \ominus n \succ m \triangleright [P]_\lambda^n}$ <p>KILL</p> $\frac{n \neq \ominus}{\Delta \triangleright [\text{kill}]_\lambda^n \xrightarrow{\tau} \Delta \ominus (n, \lambda) \triangleright \mathbf{0}}$ <p>L-OUT</p> $\frac{}{\Delta \triangleright [\bar{x}(\tilde{v}).P]_\lambda^n \xrightarrow{\bar{x}(\tilde{v})@n_\lambda} \Delta \triangleright [P]_\lambda^n}$ <p>L-LINK</p> $\frac{\Delta \not\vdash n \leftrightarrow m}{\Delta \triangleright [\text{link } m.P]_\lambda^n \xrightarrow{\tau} \Delta \oplus n \leftrightarrow m \triangleright [P]_\lambda^n}$ <p>L-UNLINK</p> $\frac{\Delta \vdash n \leftrightarrow m}{\Delta \triangleright [\text{unlink } m.P]_\lambda^n \xrightarrow{\tau} \Delta \ominus n \leftrightarrow m \triangleright [P]_\lambda^n}$ <p>L-CREATE-F</p> $\frac{\Delta \not\vdash m : \text{dead}}{\Delta \triangleright [\text{create } m.P]_\lambda^n \xrightarrow{\tau} \Delta \triangleright \mathbf{0}}$
---	--

Figure 10: LTS Rules (1)

$$\begin{array}{c}
 \text{L-CREATE-EXT} \\
 \frac{\Delta \vdash n : \text{dead} \quad \Delta_{\mathcal{A}}(n) = -\kappa}{\Delta \triangleright N \xrightarrow{\text{create}(n, \kappa+1)} \Delta \oplus (n, \kappa+1) \triangleright N} \\
 \\
 \text{L-UNLINK-EXT} \\
 \frac{\Delta \vdash n_{\lambda} : \text{alive} \quad \Delta \vdash n \leftrightarrow m}{\Delta \triangleright N \xrightarrow{\ominus n_{\lambda} \mapsto m} \Delta \ominus n \leftrightarrow m \triangleright N} \\
 \\
 \text{L-VIEW} \\
 \frac{\Delta \vdash n_{\lambda} : \text{alive} \quad \text{and } (\Delta_{\mathcal{A}}(m) = \Delta_n(m) \neq 0 \quad \text{or } \Delta_n(m) = 0)}{\Delta \triangleright N \xrightarrow{n_{\lambda} \triangleright m} \Delta \triangleright N} \\
 \\
 \text{L-PAR}_L \\
 \frac{\Delta \triangleright N \xrightarrow{\alpha} \nu \tilde{u}. \Delta' \triangleright N' \quad \tilde{u} \cap \text{fn}(M) = \emptyset}{\Delta \triangleright N \parallel M \xrightarrow{\alpha} \nu \tilde{u}. \Delta' \triangleright N' \parallel M} \\
 \\
 \text{L-SYNCL} \\
 \frac{\Delta \triangleright N \xrightarrow{\bar{x}(\tilde{u})@n_{\lambda}} \Delta \triangleright N' \quad \Delta \triangleright M \xrightarrow{x(\tilde{u})@n_{\lambda}} \Delta \triangleright M'}{\Delta \triangleright N \parallel M \xrightarrow{\tau} \Delta \triangleright N' \parallel M'} \\
 \\
 \text{L-SYNCR} \\
 \frac{\Delta \triangleright N \xrightarrow{\bar{x}(\tilde{u})@n_{\lambda}} \Delta \triangleright N' \quad \Delta \triangleright M \xrightarrow{x(\tilde{u})@n_{\lambda}} \Delta \triangleright M'}{\Delta \triangleright M \parallel N \xrightarrow{\tau} \Delta \triangleright M' \parallel N'} \\
 \\
 \text{L-KILL-EXT} \\
 \frac{\Delta \vdash n_{\lambda} : \text{alive}}{\Delta \triangleright N \xrightarrow{\text{kill}(n, \lambda)} \Delta \ominus (n, \lambda) \triangleright N} \\
 \\
 \text{L-LINK-EXT} \\
 \frac{\Delta \vdash n_{\lambda} : \text{alive} \quad \Delta \not\vdash n \leftrightarrow m}{\Delta \triangleright N \xrightarrow{\oplus n_{\lambda} \mapsto m} \Delta \oplus n \leftrightarrow m \triangleright N} \\
 \\
 \text{L-RES} \\
 \frac{S \xrightarrow{\alpha} S' \quad u \notin \text{fn}(\alpha)}{\nu u. S \xrightarrow{\alpha} \nu u. S'} \\
 \\
 \text{L-PARR} \\
 \frac{\Delta \triangleright N \xrightarrow{\alpha} \nu \tilde{u}. \Delta' \triangleright N' \quad \tilde{u} \cap \text{fn}(M) = \emptyset}{\Delta \triangleright M \parallel N \xrightarrow{\alpha} \nu \tilde{u}. \Delta' \triangleright M \parallel N'} \\
 \\
 \text{L-RES}_O \\
 \frac{S \xrightarrow{\nu \tilde{v}. \bar{x}(\tilde{u})@n_{\lambda}} S' \quad w \in \tilde{u} \setminus \tilde{v}, x, n}{\nu w. S \xrightarrow{\nu w. \nu \tilde{v}. \bar{x}(\tilde{u})@n_{\lambda}} \nu S'} \\
 \\
 \text{L-}\alpha \\
 \frac{S =_{\alpha} T \quad T \xrightarrow{\alpha} T' \quad T' =_{\alpha} S'}{S \xrightarrow{\alpha} S'}
 \end{array}$$

Figure 11: LTS Rules (2)


B Erlang

B.1 Experiments on Erlang's Semantics

In this section we will show a typical distributed Erlang behaviors. To simulate a distributed environment Erlang nodes will be running on dockers' container connected through a network. By using docker facilities we will simulate failure and changes in the network structure. The code of this examples can be found at [13].

Pinging a reincarnation of a previous known location

Description. In this scenario l_2 successfully contacts l_1 , thus establishing a connection. Then, l_1 fails and recovers. Finally, l_2 , before detecting that the incarnation of l_1 to which it was connected has failed, tests again its accessibility. Graphically, the network could be represented as follow.



The following commands set up the configuration and attach a remote shell to l_2 .

```

gfabbret@ubuntu:~/ $ docker-compose up -d
gfabbret@ubuntu:~/ $ docker exec -it l2.com erl -name test@l2.com
                        -setcookie cookie -remsh app@l2.com -hidden
  
```

Then, to establish a connection between the two locations in the Erlang console we can ping the remote location.

```

(app@l2.com)2> net_adm:ping('app@l1.com').
pong
  
```

To induce the fault and make it seem like a genuine interruption of services, without third parties being notified, we detach the container from all the networks, restart it, and reconnect it to the networks it was connected to. The disconnection is required as otherwise the other containers would be notified of the restart. Detaching and re-attaching the container to the network takes few milliseconds, hence not enough for the Erlang system to detect it. To do so we get the container and network ids through the docker commands and we feed them to the restart script.

```

gfabbret@ubuntu:~/scenario$ docker container ls
CONTAINER ID   IMAGE          COMMAND          CREATED          STATUS          NAMES
ba5915332d87   erlang:25.0.3 "erl ..."     49 seconds ago   Up 49 seconds   l2.com
76c89a08761c   erlang:25.0.3 "erl ..."     49 seconds ago   Up 49 seconds   l1.com
gfabbret@ubuntu:~/scenario-4$ docker network ls
NETWORK ID     NAME      DRIVER  SCOPE
fec145052743   bridge   bridge  local
1b349490c51c   host     host    local
373b2aa1c5e9   none     null    local
cfc05f57eeee   scenario_net1  bridge  local
gfabbret@ubuntu:~/scenario$ ./restart.sh cfc05f57eeee 76c89a08761c
ba5915332d87
done
  
```

Finally we attempt to ping the remote location after the restarting.

```

...
(app@l2.com)3> spawn('app@l1.com', fun() -> self() end).
<0.107.0>
(app@l2.com)4> =WARNING REPORT==== 2-May-2023::15:16:02.174308 ===
** Can not start erlang:apply, [#Fun<erl_eval.43.3316493>, []] on 'app@l1.com' **
  
```

The test for spawn failed even if there is a live running instance of l_1 . The reason why is that l_2 attempted to spawn on the instance of l_2 that it knew already (in terms of our calculus, location l_2 at incarnation λ), which was stored in its view, and not the one currently alive (location l_2 at

incarnation $\lambda + 1$). If we had given l_2 enough time it would have detected that the incarnation of l_2 it knew was dead since it *would not feel its heartbeat* and in that case the spawn would have succeeded since l_1 would have initiated a new connection.

An example of the corresponding system in our calculus is the following one:

$$\Delta \triangleright [\text{ping } l_1]_{\kappa}^{l_2} \parallel [\text{spawn } l_1.\bar{t}]_{\kappa}^{l_2} \parallel [\text{kill}]_{\lambda}^{l_1} \parallel [\text{create } l_1.\mathbf{0}]^{\odot}$$

where

$$\Delta = \langle \{l_1 \mapsto \lambda, l_2 \mapsto \kappa\}, \{l_1 \leftrightarrow n_2\}, \{l_2 \mapsto \hat{\mathbf{0}}\} \rangle$$

B.2 Running Example in Erlang

In this section we discuss how to reproduce the experiments showing the behaviors of the three servers presented in Section 2. Actually, for simplicity as bugged behavior we show **servDFV** from Section 3.4.

All the code of the examples can be found at [13] (including **servDF**).

We begin by **servD**.

The procedure to set up the three nodes is as above.

Then, we can connect to the interface.

```
gfabbret@ubuntu:~/ $ docker exec -it interface.com erl -name test@interface.com
                        -setcookie cookie -remsh app@interface.com -hidden
```

Then, we can invoke the interface process which will send the request and observe the outcome.

```
(app@interface.com)1> servD:interface().
Initiating request
Response received
```

As expected, everything went well.

Now, let us discuss the more interesting **servDFV**, where the misbehavior is due to a spawn that fails because of a wrong view.

The procedure to set up the system is as before, nonetheless this time two remote consoles are required: one on the interface to start the process and one on the controller to restart the router. The commands to attach the remote consoles are as above, but for replacing the name of the container.

Now, we can initiate the request from the interface.

```
(app@interface.com)1> servDFV:interface().
Initiating request
```

The process gets stuck because the (bugged) router drops the message. Now, we can simulate a failure of the router container by means of

```
gfabbret@ubuntu:~/ $ ./restart servdf_net1 router.com
```

Then, right after, before the interface node detects the absence of the router because of the lack of its *heartbeat* we can restore the service from the controller console.

```
(app@controller.com)1> servDFV:controller().
ok
```

and observe the following behavior on the interface

```
=WARNING REPORT==== 3-May-2023::09:06:33.441402 ===
** Can not start erlang:apply, [#Fun<servDF.1.93666681>, []] on 'app@router.com'
**
```

where the spawn has failed due to wrong view.

Finally, let us discuss **servDFR**.

To set up the scenario and attach consoles to the interface and the controller we proceed as above.

Then we start the request

```
(app@interface.com)1> servDFR:interface().  
Initiating request
```

The process gets stuck like it did in **servDFV**. As in **servDFV** we restart the router and we create a new router through the controller process. In this implementation though is the router to send a *retry* message to the interface, hence its view is correctly updated and the following expected behavior is shown.

```
Retrying request  
Response received  
ok
```

C Modifying networks

We formalize in this section the intuition that we can modify, in essentially arbitrary ways, the observable part of networks in our systems.

Proposition 2. *Consider a system $S = \nu \tilde{u}. \Delta \triangleright N$, where $\Delta = \langle \mathcal{A}, \mathcal{L}, \mathcal{V} \rangle$, a set of live locations \mathcal{A}' , such that $\forall n \in \text{supp}(\mathcal{A}) \setminus \text{supp}(\mathcal{A}')$, $n \notin \tilde{u}$ and a set of live links \mathcal{L}' s.t. $\forall n, m, (n, m) \in \mathcal{L} \setminus \mathcal{L}' \implies n, m \notin \tilde{u}$. Let $\Delta' = \langle \mathcal{A}', \mathcal{L}', \mathcal{V} \rangle$, then, there exists a configuration L with $\text{fn}(L) \cap \tilde{u} = \emptyset$ such that*

$$\nu \tilde{u}. \Delta \triangleright N \parallel L \implies \nu \tilde{u}. \Delta' \triangleright N$$

Proof. Assume to have $\nu \tilde{u}. \Delta \triangleright N$, where $\Delta = \langle \mathcal{A}, \mathcal{L}, \mathcal{V} \rangle$, \mathcal{A}' and \mathcal{L}' .

To transform \mathcal{A} into \mathcal{A}' we need to remove all the elements that are in \mathcal{A} but not in \mathcal{A}' and we need to add all those elements that are in \mathcal{A}' but not in \mathcal{A} . To remove the nodes it suffices to build the following context

$$L_{\mathcal{A}}^- = \prod_{n \in \text{supp}(\mathcal{A}) \setminus \text{supp}(\mathcal{A}')} [\text{kill } n]_{\lambda}^n$$

Then, to add the nodes it suffices to build the following context

$$L_{\mathcal{A}}^+ = \prod_{n \in \text{supp}(\mathcal{A}') \setminus \text{supp}(\mathcal{A})} [\text{create } n. \mathbf{0}]^{\odot}$$

To change \mathcal{L} in to \mathcal{L}' we need to perform the same operations as above. To remove the links it suffices to build the following context

$$L_{\mathcal{L}}^- = \prod_{(n, m) \in \mathcal{L} \setminus \mathcal{L}'} [\text{unlink } m. \mathbf{0}]^n$$

Then, to add the links it suffices to build the following context

$$L_{\mathcal{L}}^+ = \prod_{(n, m) \in \mathcal{L}' \setminus \mathcal{L}} [\text{link } m. \mathbf{0}]^n$$

Finally, we obtain as $L = L_{\mathcal{A}}^- \parallel L_{\mathcal{A}}^+ \parallel L_{\mathcal{L}}^- \parallel L_{\mathcal{L}}^+$. \square

Note that we did not mention incarnation numbers for new nodes in $\text{supp}(\mathcal{A}') \setminus \text{supp}(\mathcal{A})$. If a particular incarnation number $\lambda > 1$ is required for location n , it suffices to repeatedly kill and (re-)create location n in $L_{\mathcal{A}}^+$ until the correct incarnation number is reached.

In the above proposition views updates are not mentioned because there are some limits in extending them. In particular consider the network Δ

$$\langle \{(\odot, 1), (n, \lambda), (n', \lambda')\}, \{n \leftrightarrow n'\}, \{n', \{(n, \lambda_o)\}\} \rangle \quad (1)$$

where $n \neq n'$, $\lambda_o < \lambda$ and the following network Δ'

$$\langle \{(\odot, 1), (n, \lambda), (n', \lambda')\}, \{n \leftrightarrow n'\}, \{n', \{(n, \lambda_o)\}, (n, \{(n', \lambda')\})\} \rangle \quad (2)$$

Then Δ cannot be modified by context into Δ' . The problem lies in the fact that any primitive that adds a location to another location's view requires for the remote location to be in its view. In the above example, n' view is not aligned with the reality, since n' believes to be connected to (n, λ_o) while (n, λ_o) is dead and now (n, λ) , a newer incarnation, is up and running.

D Barb alternative

This section shows that it is possible to simplify our notion of barb, a result akin to that put forward by De Nicola et al. in [25]. The alternate definition of barb we consider in this section is one where a barb just displays the location of messages.

Definition 8 (Location Barb). *We say that a system S exhibits a barb at location n , in symbols $S \downarrow_n$, iff $S \equiv \nu \tilde{u}. \Delta \triangleright [\bar{x} \langle \tilde{v} \rangle . P]_{\lambda}^n \parallel N$, for some $x, n, \lambda, \tilde{u}, \tilde{v}, P, N$, where $x, n \notin \tilde{u}$, and $\Delta \vdash (n, \lambda) : \text{alive}$. Also, $S \downarrow_n$ if $S \Rightarrow S'$ and $S' \downarrow_n$. We denote by \approx_l the weak barbed congruence obtained by considering only barbs at locations.*

Now the two weak barbed congruences \approx and \approx_l coincide

Proposition 3. $\approx_l = \approx$

Proof. In what follows we use the following notation : if $U \equiv \nu \tilde{u}. \Delta \triangleright N$, and L is such that $\text{fn}(L) \cap \tilde{u} = \emptyset$, then $U \parallel L$ denotes the system $\nu \tilde{u}. \Delta \triangleright N \parallel L$.

That $\approx \subseteq \approx_l$ is clear. We show the converse. Let closed systems S and T be such that $S \approx_l T$ and $S \downarrow_{x@n_\lambda}$. By definition of full barbs, we have $S \equiv \nu \tilde{u}. \Delta \triangleright [\bar{x} \langle \tilde{v} \rangle . P]_{\lambda}^n \parallel N$ for some $x, n, \lambda, \tilde{u}, \tilde{v}, P, N$. By definition of location barbs, we have $S \downarrow_n$. Since $S \approx_l T$, we must have $T \Rightarrow T_1 \downarrow_n$ for some T_1 . Assume for the sake of contradiction that $\neg(T \downarrow_{x@n_\lambda})$, and consider the systems $S \parallel L$ and $T \parallel L$, where:

$$L = [\text{create } k. \text{link } n. \text{spawn } n. (x \langle \tilde{y} \rangle . \text{spawn } k. \bar{t})]_1^\circ \quad k, t \notin \text{fn}(S) \cup \text{fn}(T)$$

Since \approx_l is a system congruence, we must have $S \parallel L \approx_l T \parallel L$. Now, by construction, we have

$$S \parallel L \Rightarrow \nu \tilde{u}. \Delta \triangleright N \parallel [\bar{t}]_1^k \quad \text{and} \quad S \downarrow_k$$

But since $\neg(T \downarrow_{x@n_\lambda})$, and t, k are fresh for T , there can be no T' such that $T \parallel L \Rightarrow T'$ and $T' \downarrow_k$, contradicting the fact that $S \approx_l T$. \square

E Full Abstraction

E.1 Soundness

Definition 9 (Strong Simulation). *A binary relation $\mathcal{S} \subseteq \mathbb{S} \times \mathbb{S}$ over closed systems is a strong simulation iff whenever $(P, Q) \in \mathcal{S}$,*

- $P \xrightarrow{\alpha} P'$ implies $Q \xrightarrow{\alpha} Q'$ for some Q' with $(P', Q') \in \mathcal{S}$

Definition 10 (Strong Bisimulation). *A relation \mathcal{S} is a strong bisimulation if both \mathcal{S} and \mathcal{S}^{-1} are strong simulations.*

Definition 11 (Strong Bisimilarity). *Strong bisimilarity, denoted by \sim , is the largest strong bisimulation over systems.*

The fact \sim is an equivalence relation is a standard result for any labelled transition system.

We start with a simple lemma on the preservation of free names by labelled transitions. If α is a label from our LTS semantics, we define $\text{pi}(\alpha)$ and $\text{po}(\alpha)$ as follows:

$$\text{po}(\alpha) = \begin{cases} \tilde{u} \setminus \tilde{w} & \text{if } \alpha = \nu \tilde{w}. \bar{x} \langle \tilde{u} \rangle @ n_\lambda \\ \emptyset & \text{otherwise} \end{cases} \quad \text{pi}(\alpha) = \begin{cases} \tilde{u} & \text{if } \alpha = x(\tilde{u}) @ n_\lambda \\ \emptyset & \text{otherwise} \end{cases}$$

Lemma 1. *Let S, S' be closed systems such that $S \xrightarrow{\alpha} S'$. Then $\text{fn}(S') \subseteq \text{fn}(S) \cup \text{pi}(\alpha)$.*

Proof. By induction on the derivation of $S \xrightarrow{\alpha} S'$. □

We now prove two lemmas relating labelled transitions in the calculus LTS semantics and the structure of systems. In the remainder, we extend \equiv to output actions, and identify equivalent output actions, setting:

$$\nu v. \nu w. \bar{x} \langle \tilde{u} \rangle @ n_\lambda \equiv \nu w. \nu v. \bar{x} \langle \tilde{u} \rangle @ n_\lambda \quad \nu w. \alpha \equiv \nu w. \omega \quad \text{if } \alpha \equiv \omega$$

Lemma 2 (Input/output actions and systems). *Let S, S' be closed systems such that $S \xrightarrow{\alpha} S'$. The following properties hold:*

1. *if $\alpha = \nu \tilde{v}. \bar{x} \langle \tilde{u} \rangle @ n_\lambda$ then*

$$S \equiv \nu \tilde{w}. \Delta \triangleright [\bar{x} \langle \tilde{u} \rangle . P]_\lambda^n \parallel N \quad \text{for some } \Delta, \tilde{w}, P, N \quad \text{with } x, n \notin \tilde{w}, \tilde{v} \subseteq \tilde{w}, \Delta \vdash n_\lambda : \text{alive}$$

2. *if $\alpha = x(\tilde{u}) @ n_\lambda$ then*

$$S \equiv \nu \tilde{w}. \Delta \triangleright [x(\tilde{u}). P]_\lambda^n \parallel N \quad \text{for some } \Delta, \tilde{v}, \tilde{w}, P, N \quad \text{with } \tilde{u}, x, n \cap \tilde{w} = \emptyset, \Delta \vdash n_\lambda : \text{alive}$$

Proof. We show the first assertion, the second one is handled similarly. We reason by induction on the derivation of $S \xrightarrow{\alpha} S'$, where $\alpha = \nu \tilde{v}. \bar{x} \langle \tilde{u} \rangle @ n_\lambda$, considering the last rule used in the proof tree:

- Rule L-OUT: in this case, we have $S = \Delta \triangleright [\bar{x} \langle \tilde{u} \rangle . P]_\lambda^n$, $\Delta \vdash n_\lambda : \text{alive}$, as required.
- Rule L-PARL: in this case, we have $S = \Delta \triangleright N \parallel M$, $\alpha = \bar{x} \langle \tilde{u} \rangle @ n_\lambda$, with $\Delta \triangleright N \xrightarrow{\alpha} \Delta' \triangleright N'$. By induction assumption, we have $N \equiv [\bar{x} \langle \tilde{u} \rangle . P]_\lambda^n \parallel L$ for some L , with $\Delta \vdash n_\lambda : \text{alive}$. Hence $S \equiv \Delta \triangleright [\bar{x} \langle \tilde{u} \rangle . P]_\lambda^n \parallel L \parallel M$, $\Delta \vdash n_\lambda : \text{alive}$, as required.

- Rule $L\text{-PARR}$: same as $L\text{-PARL}$.
- Rule $L\text{-RES}$: in this case we have $S = \nu z.S'$ with $S' \xrightarrow{\alpha} S''$, for some S', S'' , and $v \notin \text{fn}(\alpha)$. By induction assumption, we have $S' \equiv \nu \tilde{w}.\Delta \triangleright [\bar{x}(\tilde{s}).P]_{\lambda}^n \parallel N$, with $\Delta \vdash n_{\lambda} : \text{alive}$, $x, n \notin \tilde{w}$, $\tilde{v} \subseteq \tilde{w}$. Thus $S \equiv \nu z, \tilde{w}.\Delta \triangleright [\bar{x}(\tilde{s}).P]_{\lambda}^n \parallel N$, with $\Delta \vdash n_{\lambda} : \text{alive}$, and $x, n \notin v, \tilde{w}$, $\tilde{v} \subseteq z, \tilde{w}$ as required.
- Rule $L\text{-RES}_O$: in this case, we have $S = \nu z.T$ with $T \xrightarrow{\omega} T'$, $z \in \tilde{u} \setminus x, n$, $\alpha = \nu z.\omega$. By induction assumption, we have

$$T \equiv \nu \tilde{w}.\Delta \triangleright [\bar{x}(\tilde{r}).P]_{\lambda}^n \parallel N \quad \text{for some } \Delta, \tilde{u}, \tilde{w}, P, N \text{ with } x, n \notin \tilde{w}, \tilde{v} \setminus \{z\} \subseteq \tilde{w}, \Delta \vdash n_{\lambda} : \text{alive}$$

Hence, we have

$$S \equiv \nu z, \tilde{w}.\Delta \triangleright [\bar{x}(\tilde{r}).P]_{\lambda}^n \parallel N \quad \text{for some } \Delta, \tilde{u}, \tilde{w}, P, N \text{ with } x, n \notin z, \tilde{w}, \tilde{v} \subseteq z, \tilde{w}, \Delta \vdash n_{\lambda} : \text{alive}$$

as required. □

Lemma 3 (Silent actions and systems). *Let S, S' be closed systems such that $S \xrightarrow{\tau} S'$. Then one of the following properties hold, for some $\tilde{u}, \tilde{w}, u, n, m, \lambda, P, Q, N$ with $\Delta \vdash (n, \lambda) : \text{alive}$:*

$$S \equiv \nu \tilde{w}.\Delta \triangleright [!x(\tilde{u}).P]_{\lambda}^n \parallel N \tag{3}$$

$$S \equiv \nu \tilde{w}.\Delta \triangleright [\nu u.P]_{\lambda}^n \parallel N \tag{4}$$

$$S \equiv \nu \tilde{w}.\Delta \triangleright [\text{if } r = s \text{ then } P \text{ else } Q]_{\lambda}^n \parallel N \tag{5}$$

$$S \equiv \nu \tilde{w}.\Delta \triangleright [P \mid Q]_{\lambda}^n \parallel N \tag{6}$$

$$S \equiv \nu \tilde{w}.\Delta \triangleright [\text{node}(m, \kappa).P]_{\lambda}^n \parallel N \tag{7}$$

$$S \equiv \nu \tilde{w}.\Delta \triangleright [\text{remove } n.P]_{\lambda}^n \parallel N \tag{8}$$

$$S \equiv \nu \tilde{w}.\Delta \triangleright [\text{kill}]_{\lambda}^n \parallel N \quad \text{and} \quad n \neq \odot \tag{9}$$

$$S \equiv \nu \tilde{w}.\Delta \triangleright [\text{spawn } m.P]_{\lambda}^n \parallel N \tag{10}$$

$$S \equiv \nu \tilde{w}.\Delta \triangleright [\text{create } m.P]_{\lambda}^n \parallel N \tag{11}$$

$$S \equiv \nu \tilde{w}.\Delta \triangleright [\text{link } m.P]_{\lambda}^n \parallel N \tag{12}$$

$$S \equiv \nu \tilde{w}.\Delta \triangleright [\text{unlink } m.P]_{\lambda}^n \parallel N \tag{13}$$

$$S \equiv \nu \tilde{w}.\Delta \triangleright [\bar{x}(\tilde{u}).P]_{\lambda}^n \parallel [x(\tilde{u}).Q]_{\lambda}^n \parallel N \tag{14}$$

Proof. We reason by induction on the derivation of $S \xrightarrow{\tau} S'$, considering the last rule used in the derivation.

- Rule $L\text{-RES}$: In this case, we have $S = \nu u.T$ and $S' = \nu u.T'$, with $T \xrightarrow{\tau} T'$. Applying the induction assumption, we have $T \equiv \nu \tilde{w}.\Delta \triangleright L \parallel N$, where L is one of the located processes listed in the lemma assertions ($[!x(\tilde{u}).P]_{\lambda}^n$, $[\nu u.P]_{\lambda}^n$, etc.). Hence $S \equiv \nu u, \tilde{w}.\Delta \triangleright L \parallel N$, as required.
- Rule $L\text{-SYNCL}$: In this case, we have $S = \Delta \triangleright N \parallel M$ with $\Delta \triangleright N \xrightarrow{\bar{x}(\tilde{u})@n_{\lambda}} \Delta \triangleright N'$ and $\Delta \triangleright M \xrightarrow{x(\tilde{u})@n_{\lambda}} \Delta \triangleright M'$. Applying Lemma2 we have $N \equiv [\bar{x}(\tilde{u}).P]_{\lambda}^n \parallel U$ for some P, U with $\Delta \vdash n_{\lambda} : \text{alive}$, and $M \equiv [x(\tilde{u}).Q]_{\lambda}^n \parallel V$ for some Q, V . Hence we have $S \equiv \Delta \triangleright [\bar{x}(\tilde{u}).P]_{\lambda}^n \parallel [x(\tilde{u}).Q]_{\lambda}^n \parallel U \parallel V$, one of the possible forms required.
- Rule $L\text{-SYNCR}$ is handled similarly as rule $L\text{-SYNCL}$.

- Rule $L\text{-PAR}_L$: In this case we have $S = \Delta \triangleright N \parallel M$, with $\Delta \triangleright N \xrightarrow{\tau} \nu \tilde{u}. \Delta' \triangleright N'$, and $\text{fn}(M) \cap \tilde{u} = \emptyset$. By induction assumption, we have $N \equiv L \parallel U$, where L is one of the located processes listed in the lemma assertions. Hence $S \equiv \Delta \triangleright L \parallel U \parallel M$, as required.
- Rule $L\text{-PARR}$ is handled similarly as rule $L\text{-PAR}_L$.
- Rules with $\xrightarrow{\tau}$ conclusion in Fig.10: All these rules conclusion are of the form $\Delta \triangleright L \xrightarrow{\tau} U$, where L is one of the single located processes listed in the lemma assertions. In these cases, we just have $S = \Delta \triangleright L$, as required.

□

Proposition 4 (Structural congruence is a strong bisimulation). *We have $\equiv \subseteq \sim$.*

Proof. Since \equiv is an equivalence relations, it suffices to prove that \equiv is a strong simulation, namely that for any closed systems S, R , if $S \xrightarrow{\alpha} S'$ and $S \equiv R$, then there exists R' such that $R \xrightarrow{\alpha} R'$ and $S' \equiv R'$. We reason by induction on the derivation of $S \equiv R$, considering the last rule used in the proof tree:

Rule $S\text{-RES}_C$: In this case, $S = \nu u. \nu v. T$ and $R = \nu v. \nu u. T$. Since $S \xrightarrow{\alpha} S'$, this can only have been obtained by applying rule $L\text{-RES}$ twice, rule $L\text{-RES}_O$ twice, or a or a combination of rule $L\text{-RES}$ and rule $L\text{-RRES}_O$. We consider the four cases:

Rule $L\text{-RES}$ applied twice: In this case, we have $T \xrightarrow{\omega} T'$, for some T' , $u, v \notin \text{fn}(\alpha)$, $\omega = \alpha$, and $S' = \nu u. \nu v. T'$. Now applying rule $L\text{-RES}$ twice we get $R \xrightarrow{\alpha} \nu v. \nu u. T'$. Applying rule $L\text{-RES}$ twice, we get $R' \equiv S'$, as required.

Rule $L\text{-RES}_O$ applied twice: In this case, we have $T \xrightarrow{\omega} T'$ for some T' , $v \in \pi\omega$, $u \in \pi\omega \setminus \{v\}$, $\alpha = \nu u. \nu v. \omega$, and $S' = T'$. Applying rule $L\text{-RES}_O$ twice we get: $R \xrightarrow{\alpha} T'$, hence we have $R' \equiv S'$, as required.

Rule $L\text{-RES}_O$ followed by rule $L\text{-RES}$: In this case we have $T \xrightarrow{\omega} T'$ for some T' , $v \in \pi\omega$, $u \notin \pi\omega \setminus \{v\}$, $\alpha = \nu v. \omega$, and $S' = \nu u. T'$. Applying rule $L\text{-RES}_O$ we get: $\nu u. T \xrightarrow{\omega} \nu u. T'$. Applying $L\text{-RES}_O$ we get: $R = \nu v. \nu u. T \xrightarrow{\alpha} \nu u. T' = R'$. Hence we have $R' \equiv S'$, as required.

Rule $L\text{-RES}$ followed by rule $L\text{-RES}_O$: Similar to the previous case.

Rule $S\text{-RES}_{NIL}$: In this case, $S = \nu u. R$, $u \notin \text{fn}(R)$. Since $S \xrightarrow{\alpha} S'$, this can only have been obtained by applying rule RESS as the last rule. Hence, we have $R \xrightarrow{\omega} R'$, $S' = \nu u. R'$, $\alpha = \omega \{*/u\}$, $u \notin \text{fn}(\omega) \setminus \text{po}(\omega)$. Now by Lemma 1, we have $\text{fn}(R') \subseteq \text{fn}(R) \cup \text{pi}(\omega)$. Since $\text{po}(\omega) \subseteq \text{fn}(R)$ and $u \notin \text{fn}(R)$, then $u \notin \text{po}(\omega)$, hence $\alpha = \omega$. In addition, since $u \notin \text{fn}(\omega)$, then $u \notin \text{pi}(\omega)$, hence $u \notin \text{fn}(R')$. We can then apply rule $S\text{-RES}_{NIL}$ to get $S' \equiv R'$, as required.

Rule $S\text{-}\alpha$: In this case $S =_{\alpha} R$. By rule $L\text{-}\alpha$ we get directly $R \xrightarrow{\alpha} S'$. Hence we have found $R' = S'$, as required.

Rule $S\text{-CTX}$: In this case, we have $S = \nu \tilde{u}. \Delta \triangleright L \parallel N$ and $R = \nu \tilde{u}. \Delta \triangleright L \parallel M$ with $N \equiv M$. We reason by induction on the structure of \tilde{u} :

$\tilde{u} = \emptyset$: In this case, $S \xrightarrow{\alpha} S'$ can only have been derived by an application of rule $L\text{-PAR}_L$, rule $L\text{-PARR}$, rule $L\text{-SYNCR}$, or rule $L\text{-SYNCL}$. We consider the four cases:

Rule $L\text{-PAR}_L$: In this case, we have $\Delta \triangleright L \xrightarrow{\alpha} \nu \tilde{v}.\Delta' \triangleright L'$ with $\tilde{v} \cap \text{fn}(N)$, and $S' = \nu \tilde{v}.\Delta' \triangleright L' \parallel N$. Applying rule $L\text{-PAR}_L$ we get $R \xrightarrow{\alpha} \nu \tilde{v}.\Delta' \triangleright L' \parallel N = R'$, and by applying rule $S\text{-CTX}$ we have $R' \equiv S'$, as required.

Rule $L\text{-PAR}_R$: In this case we have $\Delta \triangleright N \xrightarrow{\alpha} \nu \tilde{v}.\Delta' \triangleright N'$. We reason according to the last rule used to derive $N \equiv M$:

Rule $S\text{-PAR}_N$: In this case, we have $N = (M \parallel \mathbf{0})$. Now, $\Delta \triangleright N \xrightarrow{\alpha} \nu \tilde{v}.\Delta' \triangleright N'$ can only have been obtained via an application of rule $L\text{-PAR}_L$, with $\Delta \triangleright M \xrightarrow{\alpha} \nu \tilde{v}.\Delta' \triangleright M'$, and $N' = M' \parallel \mathbf{0}$. But then applying $L\text{-PAR}_R$ we get: $R \xrightarrow{\alpha} \nu \tilde{v}.\Delta' \triangleright L \parallel M' = R'$. Since $N' \equiv M'$ by rule $S\text{-PAR}_N$, we have by rule $S\text{-CTX}$ $S' \equiv R'$, as required.

Rule $S\text{-PAR}_C$: In this case, we have $N = U \parallel V$ and $M = V \parallel U$. Now the derivation $\Delta \triangleright N \xrightarrow{\alpha} \nu \tilde{v}.\Delta' \triangleright N'$ can only have been obtained by the application of one of the rules $L\text{-PAR}_L$, $L\text{-PAR}_R$, $L\text{-SYNCL}$ or $L\text{-SYNCR}$. We consider the different cases.

Rule $L\text{-PAR}_L$: In this case we have $\Delta \triangleright U \xrightarrow{\alpha} \nu \tilde{v}.\Delta' \triangleright U'$, $\tilde{v} \cap V = \emptyset$ and $N' = U' \parallel V$. Applying rule $L\text{-PAR}_R$ we obtain $\Delta \triangleright M \xrightarrow{\alpha} \nu \tilde{v}.\Delta' \triangleright M'$ with $M' = V \parallel U'$. Applying rule $L\text{-PAR}_R$ we get $R = \Delta \triangleright L \parallel M \xrightarrow{\alpha} \nu \tilde{v}.\Delta' \triangleright L \parallel (V \parallel U') = R'$. Since we have $S' = \nu \tilde{v}.\Delta' \triangleright L \parallel (U' \parallel V)$, and by rule $S\text{-PAR}_C$ $V \parallel U' \equiv U' \parallel V$, and since \equiv is a congruence, we have by rule $S\text{-CTX}$ $S' \equiv R'$, as required.

The case of rule $L\text{-PAR}_R$ is handled similarly.

Rule $L\text{-SYNCL}$: In this case, we have $\alpha = \tau$, $\vec{v} = \emptyset$, $\Delta \triangleright U \xrightarrow{\bar{\omega}} \Delta \triangleright U'$, $\Delta \triangleright V \xrightarrow{\omega} \Delta \triangleright V'$, where $\bar{\omega}$ is an output action and ω is the matching input action. But then we can apply rule $L\text{-SYNCR}$ to obtain $\Delta \triangleright M \xrightarrow{\tau} \Delta \triangleright M'$ where $M' = V' \parallel U'$, and then apply rule $L\text{-PAR}_R$ to get $R = \Delta \triangleright L \parallel (V \parallel U) \xrightarrow{\tau} \Delta \triangleright L \parallel (V' \parallel U')$. Since we have $S' = \Delta \triangleright L \parallel (U' \parallel V')$, $U' \parallel V' \equiv V' \parallel U'$ by rule $S\text{-PAR}_C$, and since \equiv is a congruence, we obtain by rule $S\text{-CTX}$ $S' \equiv R'$, as required.

The case of rule $L\text{-SYNCR}$ is handled similarly.

Rule $S\text{-PAR}_A$: this case is handled similarly to the case of rule $S\text{-PAR}_C$ above.

Rule $L\text{-SYNCL}$: In this case we have $\alpha = \tau$, $\Delta \triangleright L \xrightarrow{\bar{\omega}} \Delta \triangleright L'$, $\Delta \triangleright N \xrightarrow{\omega} \Delta \triangleright N'$, and $S' = \Delta \triangleright L' \parallel N'$. We reason according to the last rule used to derive $N \equiv M$:

- **Rule** $S\text{-PAR}_N$: In this case, we have $N = (M \parallel \mathbf{0})$. The transition $\Delta \triangleright N \xrightarrow{\bar{\omega}} \Delta \triangleright N'$ can only have been obtained by an application of rule $L\text{-PAR}_L$, with $\Delta \triangleright M \xrightarrow{\omega} \Delta \triangleright M'$ for some M' . Thus we have $S' = \Delta \triangleright L' \parallel (M' \parallel \mathbf{0})$. Now by applying rule $L\text{-SYNCR}$ we get $R = \Delta \triangleright L \parallel M \xrightarrow{\tau} \Delta \triangleright L' \parallel M' = R'$. By rule $S\text{-PAR}_N$, we have $M' \parallel \mathbf{0} \equiv M'$. Since \equiv is a congruence we have $L' \parallel (M' \parallel \mathbf{0}) \equiv L' \parallel M'$, and by rule $S\text{-CTX}$ we get $S' \equiv R'$, as required.
- **Rule** $S\text{-PAR}_C$: we reason exactly as in the subcase $S\text{-PAR}_C$ in the case of rule $L\text{-PAR}_L$ above, except that the only rules to consider are $L\text{-PAR}_L$ and $L\text{-PAR}_R$.
- **Rule** $S\text{-PAR}_A$: this case is handled similarly to the case of rule $S\text{-PAR}_C$ above.

Rule $L\text{-SYNCR}$: this case is handled similarly to the case of rule $L\text{-SYNCL}$ above.

$\vec{u} = v, \vec{w}$: In this case $R = \nu v.U$ with $T \equiv U$ and $S = \nu v.T \xrightarrow{\alpha} S'$. The latter can only have been obtained by rule $L\text{-RES}$ or by rule $L\text{-RES}_O$. We consider the two cases:

Rule $L\text{-RES}$: In this case, we have $T \xrightarrow{\alpha} T'$, $v \notin \pi\alpha$, $S' = \nu v.T'$. By the induction hypothesis, we have $U \xrightarrow{\alpha} U'$ for some $U' \equiv T'$. By rule $L\text{-RES}$ we get $R \xrightarrow{\alpha} \nu v.U' = R'$ and since \equiv is a congruence, we have $R' \equiv S'$, as required.

Rule $L\text{-RES}_O$: In this case, we have $T \xrightarrow{\omega} T'$, $v \in \pi\omega$, $\alpha = \nu v.\omega$, $S' = T'$. By the induction hypothesis, we have $U \xrightarrow{\omega} U'$ for some $U' \equiv T'$. By rule $L\text{-RES}_O$ we get $R \xrightarrow{\alpha} U' = R'$. hence we have $R' \equiv S'$, as required.

□

Proposition 5. *If $S \longrightarrow S'$ then $S \xrightarrow{\tau} \equiv S'$.*

Proof. We proceed by induction on the inference $S \longrightarrow S'$.

Case inferred by IF-EQ Then S is $\Delta \triangleright [\text{if } u = u.P \text{ else } Q]_{\lambda}^n$, and the transition $S \xrightarrow{\tau} S'$ follows by $L\text{-IF-EQ}$.

Case inferred by IF-NEQ Then S is $\Delta \triangleright [\text{if } u = v.P \text{ else } Q]_{\lambda}^n$, where $u \neq v$, and the transition $S \xrightarrow{\tau} S'$ follows by $L\text{-IF-NEQ}$.

Case inferred by BANG Then S is $\Delta \triangleright [!x(\tilde{u}.P)]_{\lambda}^n$, and the transition $S \xrightarrow{\tau} S'$ follows by $L\text{-BANG}$.

Case inferred by NODE Then S is $\Delta \triangleright [\text{node}(m, \kappa).P]_{\lambda}^n$, and the transition $S \xrightarrow{\tau} S'$ follows by $L\text{-NODE}$.

Case inferred by MSG Then S is $\Delta \triangleright [\bar{x}(y)]_{\lambda}^n \parallel [x(z).P]_{\lambda}^n$, and the transition $S \xrightarrow{\tau} S'$ follows by $L\text{-OUT}$, $L\text{-IN}$, and $L\text{-SYNCL}$.

Case inferred by NEW Then S is $\Delta \triangleright [\nu x.P]_{\lambda}^n$, and the transition $S \xrightarrow{\tau} S'$ follows by $L\text{-NEW}$.

Case inferred by SPAWN-S Then S is $\Delta \triangleright [\text{spawn } n.P]_{\lambda}^n$, and the transition $S \xrightarrow{\tau} S'$ follows by $L\text{-SPAWN-S}$.

Case inferred by SPAWN-F Then S is $\Delta \triangleright [\text{spawn } n.P]_{\lambda}^n$, and the transition $S \xrightarrow{\tau} S'$ follows by $L\text{-SPAWN-F}$.

Case inferred by UNLINK Then S is $\Delta \triangleright [\text{unlink } m.P]_{\lambda}^n$, and the transition $S \xrightarrow{\tau} S'$ follows by $L\text{-BREAK}$.

Case inferred by LINK Then S is $\Delta \triangleright [\text{link } m.P]_{\lambda}^n$, and the transition $S \xrightarrow{\tau} S'$ follows by $L\text{-LINK}$.

Case inferred by KILL Then S is $\Delta \triangleright [\text{kill}]_{\lambda}^n$, and the transition $S \xrightarrow{\tau} S'$ follows by $L\text{-KILL}$.

Case inferred by REMOVE Then S is $\Delta \triangleright [\text{remove } m.P]_{\lambda}^n$, and the transition $S \xrightarrow{\tau} S'$ follows by $L\text{-REMOVE}$.

Case inferred by CREATE-S Then S is $\Delta \triangleright [\text{create } m.P]_{\lambda}^n$, and the transition $S \xrightarrow{\tau} S'$ follows by $L\text{-CREATE-S}$.

Case inferred by CREATE-F Then S is $\Delta \triangleright [\text{create } m.P]_{\lambda}^n$, and the transition $S \xrightarrow{\tau} S'$ follows by $L\text{-CREATE-F}$.

Case inferred by PAR Then S is $\nu \tilde{u}.\Delta \triangleright N \parallel M$, and the transition $S \xrightarrow{\tau} S'$ follows by $L\text{-PARL}$ and the inductive hypothesis.

Case inferred by RES Then S is $\nu u.\Delta \triangleright N$, and the transition $S \xrightarrow{\tau} S'$ follows by $L\text{-RESS}$ and the inductive hypothesis.

Case inferred by STR Then we have $S \equiv T$, $T \longrightarrow T'$ and $T' \equiv S'$. By induction hypothesis, we have $T \xrightarrow{\tau} T'$. By Proposition 4 we have $S \xrightarrow{\tau} S''$ with $S'' \equiv T'$. Hence $S \xrightarrow{\tau} S'' \equiv S'$, as required.

□

Proposition 6. *If $S \xrightarrow{\tau} S'$ then $S \longrightarrow S'$.*

Proof. We proceed by induction on the inference $S \xrightarrow{\tau} S'$.

Case L-IF-EQ Then S is $\Delta \triangleright [\text{if } u = u \text{ then } P \text{ else } Q]^n$, and the reduction $S \longrightarrow S'$ follows by IF-EQ.

Case L-IF-NEQ Then S is $\Delta \triangleright [\text{if } u = v \text{ then } P \text{ else } Q]^n$, where $u \neq v$, and the reduction $S \longrightarrow S'$ follows by IF-NEQ.

Case L-BANG Then S is $\Delta \triangleright [!x(\tilde{u}).P]_\lambda^n$, and the reduction $S \longrightarrow S'$ follows by BANG.

Case L-NODE Then S is $\Delta \triangleright [\text{node}(m, \kappa).P]_\lambda^n$, and the reduction $S \longrightarrow S'$ follows by NODE.

Case L-NEW Then S is $\Delta \triangleright [\nu x.P]_\lambda^n$, and the reduction $S \longrightarrow S'$ follows by NEW.

Case L-SPAWN-S Then S is $\Delta \triangleright [\text{spawn } m.P]_\lambda^n$, and the reduction $S \longrightarrow S'$ follows by SPAWN-S.

Case L-SPAWN-F Then S is $\Delta \triangleright [\text{spawn } m.P]_\lambda^n$, and the reduction $S \longrightarrow S'$ follows by SPAWN-F.

Case L-UNLINK Then S is $\Delta \triangleright [\text{unlink } m.P]_\lambda^n$, and the reduction $S \longrightarrow S'$ follows by BREAK.

Case L-LINK Then S is $\Delta \triangleright [\text{link } m.P]_\lambda^n$, and the reduction $S \longrightarrow S'$ follows by LINK.

Case L-KILL Then S is $\Delta \triangleright [\text{kill}]_\lambda^n$, and the reduction $S \longrightarrow S'$ follows by KILL.

Case L-CREATE-S Then S is $\Delta \triangleright [\text{create } n.P]_\lambda^n$, and the reduction $S \longrightarrow S'$ follows by CREATE-S.

Case L-CREATE-F Then S is $\Delta \triangleright [\text{create } n.P]_\lambda^n$, and the reduction $S \longrightarrow S'$ follows by CREATE-F.

Case L-SYNCL Then S is $\Delta \triangleright N_1 \parallel N_2$, and we have

$$\Delta \triangleright N_1 \xrightarrow{\bar{x}(\tilde{v})@n_\lambda} \Delta \triangleright N'_1 \qquad \Delta \triangleright N_2 \xrightarrow{x(\tilde{v})@n_\lambda} \Delta \triangleright N'_2$$

Now, by Lemma 2 we know the following:

$$\Delta \triangleright N_1 \equiv \Delta \triangleright [\bar{x}(\tilde{v})]_\lambda^n \parallel M_1 \qquad \Delta \triangleright N_2 \equiv \Delta \triangleright [x(\tilde{v}).P]_\lambda^n \parallel M_2 \qquad S' = \Delta \triangleright N'_1 \parallel N'_2$$

Hence $S \equiv \Delta \triangleright [\bar{x}(\tilde{v})]_\lambda^n \parallel [x(\tilde{v}).P]_\lambda^n \parallel M_1 \parallel M_2$. Applying MSG and STR, we get $S \equiv \longrightarrow S'$.

Case L-SYNCR Similar to case L-SYNCL.

Case L-PARL Then S is $\Delta \triangleright N_1 \parallel N_2$, S' is $\nu \tilde{v}.\Delta' \triangleright N'_1 \parallel N_2$, where $\Delta \triangleright N_1 \xrightarrow{\tau} \nu \tilde{v}.\Delta' \triangleright N'_1$, and $\tilde{v} \cap \text{fn}(N_2) = \emptyset$. Now, by using the inductive hypothesis and PAR we have $\Delta \triangleright N_1 \parallel N_2 \longrightarrow \nu \tilde{v}.\Delta' \triangleright N'_1 \parallel N_2$.

Case L-PARR Similar to case L-PARL.

Case L-RES Then S is $\nu u.T$, S' is $\nu u.T'$, where $T \xrightarrow{\tau} T'$. Now by using the inductive hypothesis and RES we get $S = \nu u.T \longrightarrow \nu u.T' = S'$.

□

Lemma 4. *Let $S = \nu \tilde{s}.\Delta \triangleright N \parallel L$ be a closed system. If $S \xrightarrow{\alpha} S'$, where α is a silent action, an output action or an input action, then one of the following assertions holds:*

1. $\alpha = \tau$, $\nu \tilde{s}.\Delta \triangleright N \xrightarrow{\tau} \nu \tilde{w}.\nu \tilde{s}.\Delta \triangleright N'$, and $S' \equiv \nu \tilde{w}.\nu \tilde{s}.\Delta \triangleright N' \parallel L$, with $\text{fn}(L) \cap \tilde{w} = \emptyset$, and with rule L-PAR_L the last rule used in the derivation of $S \xrightarrow{\alpha} S'$ before possible applications of rule L-RES .
2. $\alpha = \tau$, $\nu \tilde{s}.\Delta \triangleright L \xrightarrow{\tau} \nu \tilde{w}.\nu \tilde{s}.\Delta \triangleright L'$, and $S' \equiv \nu \tilde{w}.\nu \tilde{s}.\Delta \triangleright N \parallel L'$, with $\text{fn}(N) \cap \tilde{w} = \emptyset$, and with rule L-PAR_R the last rule used in the derivation of $S \xrightarrow{\alpha} S'$ before possible applications of rules L-RES .
3. $\alpha = \tau$, $\nu \tilde{s}.\Delta \triangleright N \xrightarrow{\nu \tilde{w}.\bar{x}(\tilde{u})@n_\lambda} \nu \tilde{r}.\Delta \triangleright N'$, $\tilde{r} = \tilde{s} \setminus \tilde{w}$ and $\tilde{u} \cap \tilde{r} = \emptyset$, $\Delta \triangleright L \xrightarrow{x(\tilde{u})@n_\lambda} \Delta \triangleright L'$, and $S' \equiv \nu \tilde{s}.\Delta \triangleright N' \parallel L'$, with rule L-SYNCL the last rule used in the derivation of $S \xrightarrow{\alpha} S'$ before possible applications of rules L-RES .
4. $\alpha = \tau$, $\Delta \triangleright N \xrightarrow{x(\tilde{u})@n_\lambda} \Delta \triangleright N'$, $\nu \tilde{s}.\Delta \triangleright L \xrightarrow{\nu \tilde{w}.\bar{x}(\tilde{u})@n_\lambda} \nu \tilde{r}.\Delta \triangleright L'$, $\tilde{r} = \tilde{s} \setminus \tilde{w}$ and $\tilde{u} \cap \tilde{r} = \emptyset$, and $S' \equiv \nu \tilde{w}.\nu \tilde{s}.\Delta \triangleright N' \parallel L'$, with rule L-SYNCR the last rule used in the derivation of $S \xrightarrow{\alpha} S'$ before possible applications of rules L-RES .
5. $\alpha = \nu \tilde{w}.\bar{x}(\tilde{u})@n_\lambda$, $\nu \tilde{s}.\Delta \triangleright N \xrightarrow{\nu \tilde{w}.\bar{x}(\tilde{u})@n_\lambda} \nu \tilde{r}.\Delta \triangleright N'$, $\tilde{r} = \tilde{s} \setminus \tilde{w}$, $\tilde{u} \cap \tilde{r} = \emptyset$, and $S' \equiv \nu \tilde{s}.\Delta \triangleright N' \parallel L$ with rule L-PAR_L the last rule used in the derivation of $S \xrightarrow{\alpha} S'$ before possible applications of rules L-RES or L-RES_O .
6. $\alpha = \nu \tilde{w}.\bar{x}(\tilde{u})@n_\lambda$, $\nu \tilde{s}.\Delta \triangleright L \xrightarrow{\nu \tilde{w}.\bar{x}(\tilde{u})@n_\lambda} \nu \tilde{r}.\Delta \triangleright L'$, $\tilde{r} = \tilde{s} \setminus \tilde{w}$, $\tilde{u} \cap \tilde{r} = \emptyset$, and $S' \equiv \nu \tilde{s}.\Delta \triangleright N \parallel L'$ with rule L-PAR_R the last rule used in the derivation of $S \xrightarrow{\alpha} S'$ before possible applications of rules L-RES or L-RES_O .
7. $\alpha = x(\tilde{u})@n_\lambda$, $\nu \tilde{s}.\Delta \triangleright N \xrightarrow{x(\tilde{u})@n_\lambda} \nu \tilde{s}.\Delta \triangleright N'$, $S' \equiv \nu \tilde{s}.\Delta \triangleright N' \parallel L$, and $\tilde{u} \cap \tilde{s} = \emptyset$, with rule L-PAR_L the last rule used in the derivation of $S \xrightarrow{\alpha} S'$ before possible applications of rule L-RES .
8. $\alpha = x(\tilde{u})@n_\lambda$, $\nu \tilde{s}.\Delta \triangleright L \xrightarrow{x(\tilde{u})@n_\lambda} \nu \tilde{s}.\Delta \triangleright L'$, $S' \equiv \nu \tilde{s}.\Delta \triangleright N \parallel L'$, and $\tilde{u} \cap \tilde{s} = \emptyset$, with rule L-PAR_R the last rule used in the derivation of $S \xrightarrow{\alpha} S'$ before possible applications of rule L-RES .

Proof. By case analysis on α .

Case $\alpha = \tau$: in this case we reason by induction on the derivation of $S \xrightarrow{\tau} S'$, considering the last rule used in the proof tree:

Case L-PAR_L : In this case, we have $S = \Delta \triangleright N \parallel L$, $\Delta \triangleright N \xrightarrow{\tau} \nu \tilde{w}.\Delta' \triangleright N'$, and $S' = \Delta \triangleright N \xrightarrow{\tau} \nu \tilde{w}.\Delta' \triangleright N' \parallel L$, with $\text{fn}(L) \cap \tilde{w} = \emptyset$, corresponding to assertion 1.

Case L-PAR_R : In this case, we have $S \equiv \Delta \triangleright N \parallel L$, $\Delta \triangleright L \xrightarrow{\tau} \nu \tilde{w}.\Delta' \triangleright L'$, and $S' = \Delta \triangleright N \xrightarrow{\tau} \nu \tilde{w}.\Delta' \triangleright N \parallel L'$, with $\text{fn}(N) \cap \tilde{w} = \emptyset$, corresponding to assertion 2.

Case L-SYNCL : In this case, we have $S = \Delta \triangleright N \parallel L$, $\Delta \triangleright N \xrightarrow{\bar{x}(\tilde{u})@n_\lambda} \Delta \triangleright N'$, $\Delta \triangleright L \xrightarrow{x(\tilde{u})@n_\lambda} \Delta \triangleright L'$, and $S' = \Delta \triangleright N' \parallel L'$, corresponding to assertion 3.

Case L-SYNCR: In this case, we have $S = \Delta \triangleright N \parallel L, \Delta \triangleright N \xrightarrow{x(\tilde{u})@n_\lambda} \Delta \triangleright N', \Delta \triangleright L \xrightarrow{\bar{x}(\tilde{u})@n_\lambda} \Delta \triangleright L'$, and $S' = \Delta \triangleright N' \parallel L'$, corresponding to assertion 4.

Case L-RES: In this case, we have $S = \nu a.T, S' = \nu a.T', T \xrightarrow{\tau} T'$. By induction hypothesis, we have for T one of the four cases 1 to 4 in the lemma. We consider only the cases L-PAR_L and L-SYNCL, the other ones are handled similarly.

Rule L-PAR_L: In this case we have $T = \nu \tilde{s}.\Delta \triangleright N \parallel L, T' \equiv \nu \tilde{w}.\nu \tilde{s}.\Delta \triangleright N' \parallel L, \nu \tilde{s}.\Delta \triangleright N \xrightarrow{\tau} \nu \tilde{w}.\nu \tilde{s}.\Delta \triangleright N'$, with $\text{fn}(L) \cap \tilde{w} = \emptyset$. Applying rule L-RES, we get $S' \equiv \nu a.\nu \tilde{r}.\nu \tilde{s}.\Delta \triangleright N' \parallel L \equiv \nu a.\nu \tilde{r}.\nu \tilde{s}.\Delta \triangleright N' \parallel L, \nu a.\nu \tilde{s}.\Delta \triangleright N \xrightarrow{\tau} \nu \tilde{w}.\nu a.\nu \tilde{s}.\Delta \triangleright N'$, with $\text{fn}(L) \cap \tilde{w} = \emptyset$, corresponding to assertion 1.

Rule L-SYNCL: In this case we have $T = \nu \tilde{s}.\Delta \triangleright N \parallel L, T' \equiv \nu \tilde{s}.\Delta \triangleright N' \parallel L', \nu \tilde{s}.\Delta \triangleright N \xrightarrow{\nu \tilde{w}.\bar{x}(\tilde{u})@n_\lambda} \nu \tilde{r}.\Delta \triangleright N', \Delta \triangleright L \xrightarrow{x(\tilde{u})@n_\lambda} \Delta \triangleright L'$, with $\tilde{r} = \tilde{s} \setminus \tilde{w}$ and $\tilde{u} \cap \tilde{r} = \emptyset$. Applying rule L-RES, we get $S = \nu w.\nu \tilde{s}.\Delta \triangleright N \parallel L \xrightarrow{\tau} \nu w.\nu \tilde{s}.\Delta \triangleright N' \parallel L' = S'$.

If $a \in \tilde{u}$, then applying L-RES_O, we get $\nu a.\nu \tilde{s}.\Delta \triangleright N \xrightarrow{\nu a.\tilde{w}.\bar{x}(\tilde{u})@n_\lambda} \nu \tilde{r}.\Delta \triangleright N'$, with $\tilde{u} \cap \tilde{r} = \emptyset$ and $\tilde{r} = a, \tilde{s} \setminus a, \tilde{w}$, as required. If $a \notin \tilde{u}$, then applying L-RES we get $\nu a.\nu \tilde{s}.\Delta \triangleright N \xrightarrow{\nu \tilde{w}.\bar{x}(\tilde{u})@n_\lambda} \nu a.\nu \tilde{r}.\Delta \triangleright N'$, with $\tilde{u} \cap a, \tilde{r} = \emptyset$ and $a, \tilde{r} = a, \tilde{s} \setminus \tilde{w}$, as required.

Case $\alpha = \nu \tilde{w}.\bar{x}(\tilde{u})@n_\lambda$: handled by induction on the derivation of $S \xrightarrow{\alpha} S'$ with cases similar to the cases τ .L-PAR_L, τ .L-PAR_R and τ .L-RES above.

Case $\alpha = x(\tilde{u})@n_\lambda$: handled by induction on the derivation of $S \xrightarrow{\alpha} S'$ with cases similar to the cases τ .L-PAR_L, τ .L-PAR_R and τ .L-RES above.

□

Proposition 7 (Weak System Congruence). *Weak bisimilarity is a weak system congruence; that is, if $\nu \tilde{u}.\Delta \triangleright N \approx \nu \tilde{v}.\Delta' \triangleright M$ then for all \tilde{w}, L , with $\text{fn}(L) \cap \tilde{u}, \tilde{v} = \emptyset$, we have $\nu \tilde{w}.\nu \tilde{u}.\Delta \triangleright N \parallel L \approx \nu \tilde{w}.\nu \tilde{v}.\Delta' \triangleright M \parallel L$*

Proof. We prove that the relation

$$\mathcal{S} \stackrel{\text{def}}{=} \{ \langle \nu \tilde{w}.\nu \tilde{s}.\Delta_S \triangleright N_S \parallel L, \nu \tilde{w}.\nu \tilde{r}.\Delta_R \triangleright N_R \parallel L \rangle \mid \nu \tilde{s}.\Delta_S \triangleright N_S \approx \nu \tilde{r}.\Delta_R \triangleright N_R \quad \text{fn}(L) \cap (\tilde{s} \cup \tilde{r}) = \emptyset \}$$

is a weak bisimulation up to \equiv . Since \mathcal{S} is symmetric, it suffices to prove that \mathcal{S} is a weak bisimulation up to \equiv . Define

$$S = \nu \tilde{s}.\Delta_S \triangleright N_S \quad R = \nu \tilde{r}.\Delta_R \triangleright N_R \quad S_L = \nu \tilde{w}.\nu \tilde{s}.\Delta_S \triangleright N_S \parallel L \quad R_L = \nu \tilde{w}.\nu \tilde{r}.\Delta_R \triangleright N_R \parallel L$$

and consider a transition $S_L \xrightarrow{\alpha} U$. We proceed by induction on the structure of \tilde{w} :

Case $\tilde{w} = \emptyset$: We proceed by case analysis on α :

Case τ We consider the different cases listed in Lemma 4 for $S_L = \nu \tilde{s}.\Delta_S \triangleright N_S \parallel L$ and $S_L \xrightarrow{\tau} U$:

Case L-SYNCL In this case, we have:

$$\begin{aligned} S &= \nu \tilde{s}.\Delta_S \triangleright N_S \xrightarrow{\nu \tilde{a}.\bar{x}(\tilde{u})@n_\lambda} \nu \tilde{z}_s.\Delta_S \triangleright N'_S = S' \\ \Delta_S \triangleright L &\xrightarrow{x(\tilde{u})@n_\lambda} \Delta_S \triangleright L' \\ \tilde{z}_s &= \tilde{s} \setminus \tilde{a} \quad \tilde{u} \cap \tilde{z}_s = \emptyset \\ U &= \nu \tilde{s}.\Delta_S \triangleright N'_S \parallel L' \end{aligned}$$

Since $S \approx R$, we have $R \xrightarrow{\nu \tilde{a}.\bar{x}(\tilde{u})@n_\lambda} \nu \tilde{w}_r.\nu \tilde{z}_r.\Delta'_R \triangleright N'_R = R'$ for some $\tilde{w}_r, \Delta'_R, N'_R$, with $\tilde{z}_r = \tilde{r} \setminus \tilde{a}$, and $R' \approx S'$. Now, since $R \xrightarrow{\tau} R_1 \xrightarrow{\bar{x}(\tilde{u})@n_\lambda} R_2 \xrightarrow{\tau} R'$, where $R_1 = \nu \tilde{w}_1.\nu \tilde{z}_r.\Delta^1_R \triangleright N^1_R$ for some Δ^1_R, N^1_R , we are guaranteed that n_λ is alive, and that $\Delta^1_R \triangleright L \xrightarrow{x(\tilde{u})@n_\lambda} \Delta^1_R \triangleright L'$. Also, we have $\tilde{u} \cap \tilde{z}_r = \emptyset$. By repeated applications of rule L-PAR_L and by one application of rule L-SYNC_L, we obtain $R_L \xrightarrow{\tau} V$, with $V = \nu \tilde{w}_r.\nu \tilde{r}.\Delta'_R \triangleright N'_R \parallel L'$ and with $\text{fn}(L') \cap \tilde{w}_r = \emptyset$ because of the conditions of rule L-PAR_L.

Summing up, we have:

$$\begin{aligned} U &= \nu \tilde{a}.\nu \tilde{z}_s.\Delta_S \triangleright N'_S \parallel L' = \nu \tilde{a}.S' \\ V &\equiv \nu \tilde{a}.\nu \tilde{w}_r.\nu \tilde{z}_r.\Delta'_R \triangleright N'_R \parallel L' = \nu \tilde{a}.R' \\ S' &\approx R' \\ \text{fn}(L') &\subseteq \text{fn}(L) \cup \tilde{u} \setminus \tilde{a} \quad \text{by Lemma 1} \\ \text{fn}(L') \cap \tilde{w}_r & \\ \tilde{u} \cap \tilde{z}_s &= \tilde{u} \cap \tilde{z}_r = \emptyset \end{aligned}$$

Hence we have $\text{fn}(L') \cap \tilde{z}_s = \emptyset$ and $\text{fn}(L') \cap (\tilde{w}_r \cup \tilde{z}_r) = \emptyset$, and $S' = \nu \tilde{z}_s.\Delta_S \triangleright N'_S \approx \nu \tilde{w}_r.\nu \tilde{z}_r.\Delta'_R \triangleright N'_R = R'$, which means that $\langle U, V \rangle \in \mathcal{S}$, as required.

Case L-SYNC_R Similar to the case L-SYNC_L, but simpler.

Case L-PAR_L In this case, we have $S_L \xrightarrow{\tau} U$, with $S_L = \Delta_S \triangleright N_S \parallel L$, $U = \nu \tilde{s}.\Delta'_S \triangleright N'_S \parallel L$, $S \xrightarrow{\tau} \nu \tilde{s}.\Delta'_S \triangleright N'_S$, with $\text{fn}(L) \cap \tilde{s} = \emptyset$. Since $S \approx R$, we have $R \xrightarrow{\tau} R'$ with $R' \approx S'$ and $R' = \nu \tilde{r}.\Delta'_R \triangleright N'_R$. Now applying repeatedly rule L-PAR_L, we get $R_L \xrightarrow{\tau} \nu \tilde{r}.\Delta'_R \triangleright N'_R \parallel L = V$, with $\text{fn}(L) \cap \tilde{r} = \emptyset$ because of the conditions of rules L-PAR_L. Now we have $\langle U, V \rangle \in \mathcal{S}$, as required.

Case L-PAR_R Similar to the case L-PAR_L, but simpler.

Case $x(\tilde{u})@n_\lambda$ We consider the different cases listed in Lemma 4 for $S_L = \nu \tilde{s}.\Delta_S \triangleright N_S \parallel L$ and $S_L \xrightarrow{x(\tilde{u})@n_\lambda} U$:

Case L-PAR_L In that case, we have $S = \nu \tilde{s}.\Delta_S \triangleright N_S$, $S \xrightarrow{x(\tilde{u})@n_\lambda} S'$ $S' = \nu \tilde{s}.\Delta_S \triangleright N'_S$. Since $S \approx R$, we have $R \xrightarrow{x(\tilde{u})@n_\lambda} R'$, where $R' = \nu \tilde{z}.\nu \tilde{r}.\Delta'_R \triangleright N'_R$. Now, by repeated application of rule L-PAR_L, we get $R_L \xrightarrow{x(\tilde{u})@n_\lambda} \nu \tilde{w}.\nu \tilde{r}.\Delta'_R \triangleright N'_R \parallel L = V$, with $\text{fn}(L) \cap \nu \tilde{z} = \emptyset$ because of the conditions of rule L-PAR_L. Since we also have $\text{fn}(L) \cap \tilde{r} = \emptyset$ by definition, we have $\langle U, V \rangle \in \mathcal{S}$, as required.

Case L-PAR_R Similar to the case L-PAR_L, but simpler.

Case $\nu \tilde{a}.\bar{x}(\tilde{u})@n_\lambda$ We consider the different cases listed in Lemma 4 for $S_L = \nu \tilde{s}.\Delta_S \triangleright N_S \parallel L$ and $S_L \xrightarrow{\nu \tilde{a}.\bar{x}(\tilde{u})@n_\lambda} U$:

Case L-PAR_L In that case, we have $S = \nu \tilde{s}.\Delta_S \triangleright N_S$, $S \xrightarrow{\nu \tilde{a}.\bar{x}(\tilde{u})@n_\lambda} S'$ $S' = \nu \tilde{s}_a.\Delta_S \triangleright N'_S$, $U = \nu \tilde{s}_a.\Delta_S \triangleright N'_S \parallel L$, $\tilde{s}_a = \tilde{s} \setminus \tilde{a}$. Since $S \approx R$, we have $R \xrightarrow{\nu \tilde{a}.\bar{x}(\tilde{u})@n_\lambda} R'$, where $R' \equiv \nu \tilde{w}.\nu \tilde{r}_a.\Delta'_R \triangleright N'_R$, $\tilde{r}_a = \tilde{r} \setminus \tilde{a}$, and $S' \approx R'$. Now, by repeated application of rule L-PAR_L, we get $R_L \xrightarrow{\nu \tilde{a}.\bar{x}(\tilde{u})@n_\lambda} \nu \tilde{z}.\nu \tilde{r}_a.\Delta'_R \triangleright N'_R \parallel L = V$, with $\text{fn}(L) \cap \nu \tilde{z} = \emptyset$ because of the conditions of rule L-PAR_L. Since we also have $\text{fn}(L) \cap \tilde{r} = \emptyset$ by definition, we have $\langle U, V \rangle \in \mathcal{S}$, as required.

Case L-PAR_R Similar to the case L-PAR_L, but simpler.

Case create(n, λ) Easy since S_L and S have the same network.

Case kill(n, λ) Easy since S_L and S have the same network.

Case $\oplus n \mapsto m$ Easy since S_L and S have the same network.

Case $\ominus n \mapsto m$ Easy since S_L and S have the same network.

Case $n \succ m$ Easy since S_L and S have the same network.

Case $\tilde{w} = w, \tilde{z}$: Let $\langle \nu w., \nu \tilde{z}.S, \nu w., \nu \tilde{z}.R \in \mathcal{S}$. Then we have by construction $\langle \nu \tilde{z}.S, \nu \tilde{z}.R \in \mathcal{S}$. By induction hypothesis, if $\nu \tilde{z}.S \xrightarrow{\alpha} U$, then $\nu \tilde{z}.R \xrightarrow{\alpha} V$ for some V with $\langle U, V \rangle \in \mathcal{S}$. Now using rule L-RES or rule L-RES_O, we obtain either $S_L \xrightarrow{\alpha} \nu w.U$ or $S_L \nu w. \xrightarrow{\nu w.\alpha} U$, which are matched respectively by $R_L \xrightarrow{\alpha} \nu w.V$ or $R_L \xrightarrow{\nu w.\alpha} V$. In both cases, we have $\langle U, V \rangle \in \mathcal{S}$ and $\langle \nu w.U, \nu w.V \rangle \in \mathcal{S}$ (\mathcal{S} is closed under restriction by construction), as required. □

Proposition 8 (Soundness of weak bisimilarity). *Weak bisimilarity is sound with respect to weak barbed congruence, i.e. $\approx \subseteq \approx$.*

Proof. Weak bisimilarity is weak barb-preserving thanks to Lemma 2. It is weak reduction-closed thanks to Proposition 6. It is a system congruence thanks to Proposition 7. Thus $\approx \subseteq \approx$ by definition of weak barbed congruence as the largest of weak barb-preserving, reduction-closed, system congruence. □

E.2 Completeness

In the remainder of this section, we use the following notation: if $S = \nu \tilde{s}.\Delta \triangleright N$ is a closed system, we write $S \parallel M$ for the system $\triangleright \triangleright \nu \tilde{s}.\Delta N \parallel M$ provided $\text{fn}(M) \cap \tilde{s} = \emptyset$.

Lemma 5 (Inducing Network Changes). • Suppose $\Delta \vdash (n, \lambda) : \text{alive}$

- $S \xrightarrow{\text{kill}(n, \lambda)} S'$ implies $S \parallel [\text{kill}]_{\lambda}^n \longrightarrow S'$
- $S \parallel [\text{kill}]_{\lambda}^n \longrightarrow S'$, where $S' \equiv \nu \tilde{u}.\Delta \triangleright N$, $\Delta \not\vdash (n, \lambda) : \text{alive}$ implies $S \xrightarrow{\text{kill}(n, \lambda)} S'$

• Suppose $\Delta \vdash (n, \lambda) : \text{alive}$

- $S \xrightarrow{\text{create}(n, \lambda)} S'$ implies $S \parallel [\text{create } n.P]_{\kappa}^m \longrightarrow S' \parallel [P]_{\lambda}^n$
- $S \parallel [\text{create } n.P]_{\kappa}^m \longrightarrow S' \parallel [P]_{\lambda}^n$, where $S' \equiv \nu \tilde{u}.\Delta \triangleright N$, $\Delta \vdash n_{\lambda} : \text{alive}$ implies $S \xrightarrow{\text{create}(n, \lambda)} S'$

• Suppose $S = \nu \tilde{u}.\Delta.N$ and $\Delta \vdash n \leftrightarrow m$

- $S \xrightarrow{\ominus n_{\lambda} \mapsto m} S'$, where $S' \equiv \nu \tilde{u}.\Delta' \triangleright N$ implies $S \parallel [\text{unlink } m.P]_{\lambda}^n \longrightarrow S' \parallel [P]_{\lambda}^n$
- $S \parallel [\text{unlink } m.P]_{\lambda}^n \longrightarrow S' \parallel [P]_{\lambda}^n$, where $S' \equiv \nu \tilde{u}.\Delta' \triangleright N$, $\Delta' \vdash n \leftrightarrow m$ implies $S \xrightarrow{\ominus n_{\lambda} \mapsto m} S'$

• Suppose $S = \nu \tilde{u}.\Delta.N$ and $\Delta \vdash n \leftrightarrow m$

- $S \xrightarrow{\oplus n_{\lambda} \mapsto m} S'$, where $S' \equiv \nu \tilde{u}.\Delta' \triangleright N$ implies $S \parallel [\text{link } m.P]_{\lambda}^n \longrightarrow S' \parallel [P]_{\lambda}^n$

$$- S \parallel [\text{link } m.P]_{\lambda}^n \longrightarrow S' \parallel [P]_{\lambda}^n, \text{ where } S' \equiv \nu \tilde{u}.\Delta' \triangleright N, \Delta' \vdash n \leftrightarrow m \text{ implies} \\ S \xrightarrow{\oplus n_{\lambda} \mapsto m} S'$$

Proof. The first clause for the action $\text{kill}(n, \lambda)$ is proved by induction on the derivation $S = \nu \tilde{u}.\Delta \triangleright N \xrightarrow{\text{kill}(n, \lambda)} \nu \tilde{u}.\Delta' \triangleright N = S'$. The second clause uses induction on the derivation of $\nu \tilde{u}.\Delta \triangleright N \parallel [\text{kill}]_{\lambda}^n \longrightarrow \nu \tilde{u}.\Delta' \triangleright N'$. The proof for the other clauses is similar. \square

Given a system $S \equiv \nu \tilde{u}.\Delta \triangleright N$ and l s.t. $l \notin \text{fn}(S) \cup \text{bn}(S)$ we define S^l as $\nu \tilde{u}.\Delta \oplus (l, 1) \triangleright N$.

Proposition 9. *If S is a closed system and $l, x \notin \text{fn}(S)$ then $\nu x, l.S^l \parallel [\bar{x}]_{\lambda}^l \sim S$.*

Proof. A direct proof that $\mathcal{R} = \{\langle \nu x, l.S^l \parallel [\bar{x}]_{\lambda}^l, S \rangle \mid S \in \mathbb{S} \text{ closed, } x, l \notin \text{fn}(S)\}$ is a strong bisimulation, noting that any transition $\nu x, l.S^l \parallel [\bar{x}]_{\lambda}^l \xrightarrow{\alpha} T^l$ must have been obtained by a derivation with $S^l \xrightarrow{\alpha} U^l$ as a premise, with $T^l = \nu x, l.U^l$, $x, n \notin \text{fn}(\alpha)$, and that for any such transition we have $S \xrightarrow{\alpha} T$. \square

Proposition 10. *If we have $S^l \parallel [\bar{x}]_{\lambda}^l \approx R^l \parallel [\bar{x}]_{\lambda}^l$ with x, l fresh for S, R then $S \approx R$.*

Proof. If $S^l \parallel [\bar{x}]_{\lambda}^l \approx R^l \parallel [\bar{x}]_{\lambda}^l$ then by system congruence we also have

$$\nu x, l.S^l \parallel [\bar{x}]_{\lambda}^l \approx \nu x, l.R^l \parallel [\bar{x}]_{\lambda}^l$$

and then by Proposition 9, Proposition 8, and transitivity of \approx we have $S \approx R$ as required. \square

Theorem 2 (Completeness of \approx w.r.t. \approx). *If $S \approx R$ then $S \approx R$.*

Proof. To prove the statement it suffices to show that $\approx \cup \approx$ is a bisimulation up-to \equiv . Take $S \approx R$ and suppose that $S \xrightarrow{\alpha} S'$; we reason by case analysis on α .

- Case $\alpha = \tau$. Thanks to Proposition 5 and Proposition 6, the thesis follow by the reduction closure property.
- Case $\alpha = \bar{x}(\tilde{u})@n_{\lambda}$. Consider the context

$$L = [x(\tilde{y}).\text{if } \tilde{y} = \tilde{u} \text{ then create } l.(\overline{\text{fail}} \mid \text{fail}.\overline{\text{succ}})]_{\lambda}^n$$

with l, succ and fail fresh and the reduction $S \parallel L \Longrightarrow T_1$, where $T_1 \equiv S^l \parallel [\overline{\text{succ}}]_{\lambda}^l$. Now, since $S \approx R$, $T_1 \downarrow_{\text{succ}@l}$ and $T_1 \not\downarrow_{\text{fail}@l}$, we must have a transition $R \parallel L \Longrightarrow T_2$ s.t. $T_2 \downarrow_{\text{succ}@l}$ and $T_2 \not\downarrow_{\text{fail}@l}$. We remark here that we have $T_2 \downarrow_{\text{succ}@l}$, with a strong barb, because the only way for $R \parallel L$ to make disappear $\downarrow_{\text{fail}@l}$ is to have consumed it on the fresh location l , thereby showing $\downarrow_{\text{succ}@l}$.

The only way to obtain this is if $R \xrightarrow{\bar{x}(\tilde{u})@n_{\lambda}} R'$, hence $T_2 \equiv R^l \parallel [\overline{\text{succ}}]_{\lambda}^l$. Now, since $(S^l \parallel [\overline{\text{succ}}]_{\lambda}^l, R^l \parallel [\overline{\text{succ}}]_{\lambda}^l) \in \approx$, by Proposition 10 we have $(S', R') \in \approx$ as required.

- Case $\alpha = x(\tilde{u})@n_{\lambda}$. Consider the context

$$L = [\bar{x}(\tilde{u}).\text{create } l.(\overline{\text{fail}} \mid \text{fail}.\overline{\text{succ}})]_{\lambda}^n$$

with $l, \text{fail}, \text{succ}$ fresh and the reduction $S \parallel L \Longrightarrow T_1$, where $T_1 \equiv S^l \parallel [\overline{\text{succ}}]_{\lambda}^l$. Now, since $S \approx R$, $T_1 \downarrow_{\text{succ}@l}$ and $T_1 \not\downarrow_{\text{fail}@l}$, we must have a transition $R \parallel L \Longrightarrow T_2$ s.t. $T_2 \downarrow_{\text{succ}@l}$ and $T_2 \not\downarrow_{\text{fail}@l}$. The only way to obtain this is if $R \xrightarrow{x(\tilde{u})@n_{\lambda}} R'$, hence $T_2 \equiv R^l \parallel [\overline{\text{succ}}]_{\lambda}^l$. Now, since $(S^l \parallel [\overline{\text{succ}}]_{\lambda}^l, R^l \parallel [\overline{\text{succ}}]_{\lambda}^l) \in \approx$, by Proposition 10 we have $(S', R') \in \approx$ as required.

- Case $\alpha = \oplus n_\lambda \mapsto m$. By Lemma 5 we know that $[\text{link } m.P]_\lambda^n$ induces the desired labeled action. Consider the context

$$L = [\text{link } m.\text{create } l.(\overline{fail} \mid fail.\overline{succ})]_\lambda^n$$

with l , $fail$, $succ$ fresh and the reduction $S \parallel L \Longrightarrow T_1$, where $T_1 \equiv S'^l \parallel [\overline{succ}]^l$. Now, since $S \approx R$, $T_1 \downarrow_{succ@l}$ and $T_1 \not\downarrow_{fail@l}$, we must have a transition $R \parallel L \Longrightarrow T_2$ s.t. $T_2 \downarrow_{succ@l}$ and $T_2 \not\downarrow_{fail@l}$. The only way to obtain this is if $R \Longrightarrow R_1 \parallel [\text{link } m.\text{create } l.(\overline{fail} \mid fail.\overline{succ})]_\lambda^n \longrightarrow R_1 \parallel [\text{create } l.(\overline{fail} \mid fail.\overline{succ})]_\lambda^n \Longrightarrow T_2$, where $T_2 \equiv R'^l \parallel [\overline{succ}]^l$.

Then, by Lemma 5 we know that $R \xrightarrow{\tau} R_i \xrightarrow{\ominus n_\lambda \mapsto m} R_j \xrightarrow{\tau} R'$.

Now, since $(S'^l \parallel [\overline{succ}]^l, R'^l \parallel [\overline{succ}]^l) \in \approx$, by Proposition 10 we have $(S', R') \in \approx$ as required.

- Case $\alpha = \ominus n_\lambda \mapsto m$. Consider the context

$$L = [\text{unlink } m.\text{create } l.(\overline{fail} \mid fail.succ)]_\lambda^n$$

with l , $succ$ and $fail$ fresh. The reasoning is similar to above.

- Case $\alpha = \text{kill}(n, \lambda)$. By Lemma 5 we know that $[\text{kill}]_\lambda^n$ is the process that induces the reduction we are looking for. Consider the context

$$L = [\overline{fail}]_\lambda^n \parallel [\text{kill}]_\lambda^n$$

with $fail$ fresh and reduction $S \parallel L \Longrightarrow T_1$.

Now, since $S \approx R$ and $T_1 \not\downarrow_{fail@n}$ there is a matching move $R \parallel L \Longrightarrow T_2$ s.t. $T_2 \not\downarrow_{fail}$. Since $fail$ is fresh the only way to have $T_2 \not\downarrow_{fail@n}$ is to have $R \parallel L \Longrightarrow R_1 \parallel [\overline{fail}]_\lambda^n \parallel [\text{kill}]_\lambda^n \longrightarrow R'_1 \parallel [\overline{fail}]_\lambda^n \Longrightarrow T_2$.

By Lemma 5 we know that $R \xrightarrow{\tau} R_1 \xrightarrow{\text{kill}(n, \lambda)} R'_1 \xrightarrow{\tau} R'$.

Now, we know that $T_1 \approx T_2$, where $T_1 \equiv S' \parallel [\overline{fail}]_\lambda^n$ and $T_2 \equiv R' \parallel [\overline{fail}]_\lambda^n$. We know $S' \parallel [\overline{fail}]_\lambda^n \sim S'$ since n_λ is not alive in S' hence $S \parallel [\overline{fail}]_\lambda^n \approx S$ and by transitivity of \approx we get $S \approx R \parallel [\overline{fail}]_\lambda^n$. By using a similar reasoning we get $(S', R') \in \approx$ as required.

- Case $\alpha = \text{create}(n, \lambda)$. Consider the context

$$L = [\text{create } n.(\text{create } l.(\overline{fail} \mid fail.\overline{succ}))]^\circ$$

with l , $fail$ and $succ$ fresh; the reasoning is similar to above.

- Case $\alpha = n_\lambda \succ m$. Consider the context

$$L = \left[\nu x. \left(x \mid \begin{array}{l} x.\text{link } m.(\text{spawn } m.(\text{unlink } n.\text{create } l.(\overline{fail} \mid fail.\overline{succ}_1)) \mid) \\ x.\text{spawn } m.(\text{create } l.(\overline{fail} \mid fail.\overline{succ}_2)) \end{array} \right) \right]_\lambda^n$$

with l , $fail$, $succ_1$ and $succ_2$ fresh. Now, since $S \approx R$ their public network must agree on aliveness of localities, links, and views. There are two cases to be considered: when the link $n \leftrightarrow m$ is alive in S and when it is not.

- Case $n \leftrightarrow m$.

Consider the transition $S \parallel L \Longrightarrow T_1 \equiv \nu x.S^l \parallel [x.\text{spawn}m.\text{create}l.\overline{fail} \mid fail.\overline{succ}]_\lambda^n \parallel [\overline{succ}_1]^l$. Now, since $S \approx R$ and $T_1 \downarrow_{succ_1@l}$ and $T_1 \not\downarrow_{fail@l}$ we have a matching transition $R \parallel L \Longrightarrow T_2$ s.t. $T_2 \downarrow_{succ_1@l}$ and $T_2 \not\downarrow_{fail@l}$. Hence we have $T_2 \equiv \nu x.R^l \parallel [x.\text{spawn}m.\text{create}l.\overline{fail} \mid fail.\overline{succ}]_\lambda^n \parallel [\overline{succ}_1]^l$ (x is fresh in R). This is possible only if $R \xrightarrow{n_\lambda \succ m} R'$.

Now, $\nu x.S^l \parallel [x.\text{spawn}m.\text{create}l.\overline{fail} \mid fail.\overline{succ}]_\lambda^n \parallel [\overline{succ}_1]^l \sim S^l \parallel [\overline{succ}_1]^l$, because $x \notin fn(S)$.

By using a similar reasoning for T_2 and transitivity of \approx we have $(S^l \parallel [\overline{succ}_1]^l, R^l \parallel [\overline{succ}_1]^l) \in \approx$. Finally by Proposition 10 we have $(S', R') \in \approx$ as required.

- Case $n \not\leftrightarrow m$.

Consider the transition $S \parallel L \Longrightarrow T_1 \equiv S^l \parallel [x.\text{link}m.(\text{spawn}m.(\text{unlink}n.\text{create}l.(\overline{fail} \mid fail.\overline{succ}_1)))]_\lambda^n \parallel [\overline{succ}_2]^l$. Now, since $S \approx R$ and $T_1 \downarrow_{succ_2@l}$ and $T_1 \not\downarrow_{fail@l}$ we have a matching transition $R \parallel L \Longrightarrow T_2$ s.t. $T_2 \downarrow_{succ_2@l}$ and $T_2 \not\downarrow_{fail@l}$. Hence we have $T_2 \equiv R^l \parallel [x.\text{link}m.(\text{spawn}m.(\text{unlink}n.\text{create}l.(\overline{fail} \mid fail.\overline{succ}_1)))]_\lambda^n \parallel [\overline{succ}_2]^l$. This is possible only if $R \xrightarrow{n_\lambda \succ m} R'$.

Now, $\nu x.S^l \parallel [\overline{succ}_2]^l \parallel [x.\text{link}m.(\text{spawn}m.(\text{unlink}n.\text{create}l.(\overline{fail} \mid fail.\overline{succ}_1)))]_\lambda^n \sim S^l \parallel [\overline{succ}_2]^l$, because $x \notin fn(S)$.

By using a similar reasoning for T_2 and transitivity of \approx we have $(S^l \parallel [\overline{succ}_2]^l, R^l \parallel [\overline{succ}_2]^l) \in \approx$. Finally by Proposition 10 we have $(S', R') \in \approx$ as required.

□

F Bisimulation Of Running Example

In this section we show that a bisimulation relating the example without failure and the example with recovery.

We recall briefly the two systems.

$$\begin{aligned} \mathbf{servD} &= \nu n_r, n_b, r_1, r_2, b. \Delta \triangleright [I]^{n_i} \parallel [R]^{n_r} \parallel [B]^{n_b} \\ \mathbf{servDFR} &= \nu n_r, n_b, n_c, r_1, r_2, b, c, \mathit{retry}. \Delta' \triangleright ([K]^{n_i} \parallel [R]^{n_r} \parallel [\mathit{kill}]^{n_r} \parallel [!B]^{n_b} \parallel [C]^{n_c}) \end{aligned}$$

where:

$$\begin{aligned} C &= \nu c. (\bar{c} !c. (\mathit{create} n_r. (R \mid \mathit{spawn} n_i. \overline{\mathit{retry}}) \mid \bar{c}) \\ I &= \mathit{req}(y, z). \mathit{spawn} n_r. \bar{r}_1 \langle y, z \rangle \\ R &= (r_1(y, z). \mathit{spawn} n_b. \bar{b} \langle y, z \rangle) \mid (r_2(y, z). \mathit{spawn} n_i. \bar{z} \langle y \rangle) \\ B &= b(y, z). \mathit{spawn} n_r. \bar{r}_2 \langle z, w_y \rangle \\ K &= \mathit{req}(y, z). ((\mathit{spawn} n_r. \bar{r}_1 \langle y, c \rangle) \mid c(w). \bar{z} \langle w \rangle \mid \mathit{retry}. \mathit{spawn} n_r. \bar{r}_1 \langle y, c \rangle) \end{aligned}$$

Proposition 11. $\mathbf{servD} \approx \mathbf{servDFR}$

Proof. In the following, for simplicity reasons, we will omit in the action label the incarnation number of the location doing the action.

Consider relation $\mathcal{R} = \{(\mathbf{servD}, \mathbf{servDFR})\} \cup \mathcal{S}_0 \cup \mathcal{S}_1 \cup \mathcal{S}_2$ where

$$\begin{aligned} \mathcal{S}_0 &= \{(\mathbf{servD}, R_0) \mid \mathbf{servDFR} \xrightarrow{\tau} R_0\} \\ \mathcal{S}_1 &= \{(S_1, R_1) \mid (S_0, R_0) \in \mathcal{S}_0, S_0 \xrightarrow{\mathit{req}(x,y)@n_i} S_1 \wedge R_0 \xrightarrow{\mathit{req}(x,y)@n_i} R_1\} \\ \mathcal{S}_2 &= \{(S_2, R_2) \mid (S_1, R_1) \in \mathcal{S}_1, S_1 \xrightarrow{\bar{z}\langle w \rangle @n_1} S_2 \wedge R_1 \xrightarrow{\bar{z}\langle w \rangle @n_1} R_2\} \end{aligned}$$

Now we analyze the moves of the various pairs to show that indeed \mathcal{R} is a bisimulation.

- Pair $(\mathbf{servD}, \mathbf{servDFR})$. Now we proceed by case analysis on the possible transitions.
 - Case $x(y, z)@n_i$.

* Case $\mathbf{servD} \xrightarrow{x(y,z)@n_i} T_1$.

Consider transition

$$\mathbf{servD} \xrightarrow{x(y,z)@n_i} T_1 \equiv \nu n_r, n_b, r_1, r_2, b. \Delta \triangleright [\mathit{spawn} n_r. \bar{r}_1 \langle y, z \rangle]^{n_i} \parallel [R]^{n_r} \parallel [B]^{n_b}$$

The move can then be matched by

$$\begin{aligned} \mathbf{servDFR} &\xrightarrow{x(y,z)@n_i} \\ T_2 &\equiv \nu n_r, n_b, n_c, r_1, r_2, b, c, \mathit{retry}. \Delta' \triangleright \\ &\left([(\mathit{spawn} n_r. \bar{r}_1 \langle y, c \rangle) \mid c(w). \bar{z} \langle w \rangle \mid \mathit{retry}. \mathit{spawn} n_r. \bar{r}_1 \langle y, c \rangle]^{n_c} \parallel [R]^{n_r} \parallel [\mathit{kill}]^{n_r} \parallel [!B]^{n_b} \parallel [C]^{n_c} \right) \end{aligned}$$

and $(T_1, T_2) \in \mathcal{R}$.

- * Case **servDFR** $\xrightarrow{x(y,z)@n_i} T_1$. Similar to above.
- Case τ .
- * Consider transition

$$\begin{aligned} \mathbf{servDFR} &\xrightarrow{\tau} \\ T_1 &\equiv \nu n_b, n_r, b, s, \mathit{retry}, n_c, c. \Delta \triangleright \\ &[K]^{n_i} \parallel [R]^{n_r} \parallel [\mathit{kill}]^{n_2} \parallel [!B]^{n_b} \parallel \\ &[\bar{c} | c.(\mathit{create} n_r.(R | \mathit{spawn} n_i.\overline{\mathit{retry}}) | \bar{c})]^{n_c} \end{aligned}$$

Then, the move can be matched by the other system by doing nothing, indeed $(\mathbf{servD}, T_1) \in \mathcal{R}$

- * Consider transition

$$\begin{aligned} \mathbf{servDFR} &\xrightarrow{\tau} T_1 \equiv \nu n_r, n_b, a, b, s, \mathit{retry}, n_c. \Delta - (n_b, \lambda_b) \triangleright \\ &[K]^{n_i} \parallel [R]^{n_r} \parallel [!B]^{n_b} \parallel [C]^{n_c} \end{aligned}$$

Then, the move can be matched by the other system by doing nothing, indeed $(\mathbf{servD}, T_1) \in \mathcal{R}$

- Pairs $(S_1, R_1) \in \mathcal{S}_1$

Case τ Here, any possible τ transition can be matched by the other system by doing nothing.

Case $\bar{z}\langle w \rangle @ n_i$ Here there is only one possible system S that can perform the step $\bar{z}\langle w \rangle @ n_i$, that is

$$\nu \tilde{u}. \Delta \triangleright [\bar{z}\langle w \rangle]^{n_1}$$

Now, there are an unbounded number of different system R_1 that can match the move, according to the state of the recovery. We only show the following example. Consider system

$$\begin{aligned} R_1 &\equiv \nu n_r, n_b, n_c, r_1, r_2, b, c, \mathit{retry}. \Delta'' \triangleright \\ &\left(\begin{array}{l} [c(w).\bar{z}\langle w \rangle]^{n_i} \parallel [\mathit{retry}.\mathit{spawn} n_r.\overline{r_1}\langle y, c \rangle]^{n_i} \parallel \\ [\overline{\mathit{retry}} | R]^{n_r} \parallel [\mathit{kill}]^{n_r} \parallel [!B]^{n_b} \parallel [C]^{n_c} \end{array} \right) \end{aligned}$$

In R_1 the failure has occurred (we omit dead locations) and the recovery process has

started already. By doing the following steps R can match the move.

$$\begin{aligned}
& R_1 \xrightarrow{\tau} \\
\nu \tilde{u}.\Delta'' & \triangleright \frac{[c(w).\bar{z}\langle w \rangle]^{n_i} \parallel [\text{retry.spawn } n_r.\bar{r}_1\langle y, c \rangle]^{n_i} \parallel [\overline{\text{retry}}]^{n_r} \parallel [R]^{n_r} \parallel [\text{kill}]^{n_r} \parallel [!B]^{n_b} \parallel [C]^{n_c}}{[\text{kill}]^{n_r} \parallel [!B]^{n_b} \parallel [C]^{n_c}} \xrightarrow{\tau} \\
\nu \tilde{u}.\Delta'' & \triangleright [c(w).\bar{z}\langle w \rangle]^{n_i} \parallel [\text{spawn } n_r.\bar{r}_1\langle y, c \rangle]^{n_i} \parallel [R]^{n_r} \parallel [\text{kill}]^{n_r} \parallel [!B]^{n_b} \parallel [C]^{n_c} \xrightarrow{\tau} \\
& \nu \tilde{u}.\Delta'' \triangleright [c(w).\bar{z}\langle w \rangle]^{n_i} \parallel [\bar{r}_i\langle y, c \rangle]^{n_r} \parallel [R]^{n_r} \parallel [\text{kill}]^{n_r} \parallel [!B]^{n_b} \parallel [C]^{n_c} \xrightarrow{\tau} \\
& \nu \tilde{u}.\Delta'' \triangleleft \left(\frac{[c(w).\bar{z}\langle w \rangle]^{n_i} \parallel [\bar{r}_1\langle y, c \rangle]^{n_r} \parallel [r_1(y, z).\text{spawn } n_b.\bar{b}\langle y, z \rangle]^{n_r} \parallel [r_2(y, z).\text{spawn } n_i.\bar{z}\langle y \rangle]^{n_r} \parallel [\text{kill}]^{n_r} \parallel [!B]^{n_b} \parallel [C]^{n_c}}{[\text{kill}]^{n_r} \parallel [!B]^{n_b} \parallel [C]^{n_c}} \right) \xrightarrow{\tau} \\
\nu \tilde{u}.\Delta'' & \triangleright \left(\frac{[c(w).\bar{z}\langle w \rangle]^{n_i} \parallel [\text{spawn } n_b.\bar{b}\langle y, c \rangle]^{n_r} \parallel [r_2(y, z).\text{spawn } n_i.\bar{z}\langle y \rangle]^{n_r} \parallel [\text{kill}]^{n_r} \parallel [!B]^{n_b} \parallel [C]^{n_c}}{[\text{kill}]^{n_r} \parallel [!B]^{n_b} \parallel [C]^{n_c}} \right) \xrightarrow{\tau} \\
\nu \tilde{u}.\Delta'' & \triangleright \left(\frac{[c(w).\bar{z}\langle w \rangle]^{n_i} \parallel [\bar{b}\langle y, c \rangle]^{n_b} \parallel [r_2(y, z).\text{spawn } n_i.\bar{z}\langle y \rangle]^{n_r} \parallel [\text{kill}]^{n_r} \parallel [!B]^{n_b} \parallel [C]^{n_c}}{[\text{kill}]^{n_r} \parallel [!B]^{n_b} \parallel [C]^{n_c}} \right) \xrightarrow{\tau} \\
\nu \tilde{u}.\Delta'' & \triangleright \left(\frac{[c(w).\bar{z}\langle w \rangle]^{n_i} \parallel [\bar{b}\langle y, c \rangle]^{n_b} \parallel [r_2(y, z).\text{spawn } n_i.\bar{z}\langle y \rangle]^{n_r} \parallel [\text{kill}]^{n_r} \parallel [!B]^{n_b} \parallel [C]^{n_c}}{[\text{kill}]^{n_r} \parallel [!B]^{n_b} \parallel [C]^{n_c}} \right) \xrightarrow{\tau} \\
\nu \tilde{u}.\Delta'' & \triangleright \left(\frac{[c(w).\bar{z}\langle w \rangle]^{n_i} \parallel [r_2(y, z).\text{spawn } n_i.\bar{z}\langle y \rangle]^{n_r} \parallel [\text{kill}]^{n_r} \parallel [!B]^{n_b} \parallel [C]^{n_c}}{[\text{kill}]^{n_r} \parallel [!B]^{n_b} \parallel [C]^{n_c}} \right) \xrightarrow{\tau} \\
\nu \tilde{u}.\Delta'' & \triangleright \left(\frac{[c(w).\bar{z}\langle w \rangle]^{n_i} \parallel [r_2(y, z).\text{spawn } n_i.\bar{z}\langle y \rangle]^{n_r} \parallel [\text{kill}]^{n_r} \parallel [!B]^{n_b} \parallel [C]^{n_c}}{[\text{kill}]^{n_r} \parallel [!B]^{n_b} \parallel [C]^{n_c}} \right) \xrightarrow{\tau} \\
\nu \tilde{u}.\Delta'' & \triangleright \left(\frac{[c(w).\bar{z}\langle w \rangle]^{n_i} \parallel [r_2(y, z).\text{spawn } n_i.\bar{z}\langle y \rangle]^{n_r} \parallel [\text{kill}]^{n_r} \parallel [!B]^{n_b} \parallel [C]^{n_c}}{[\text{kill}]^{n_r} \parallel [!B]^{n_b} \parallel [C]^{n_c}} \right) \xrightarrow{\tau} \\
\nu \tilde{u}.\Delta'' & \triangleright ([c(w).\bar{z}\langle w \rangle]^{n_i} \parallel [\text{spawn } n_i.\bar{c}\langle w_y \rangle]^{n_r} \parallel [\text{kill}]^{n_r} \parallel [!B]^{n_b} \parallel [C]^{n_c}) \xrightarrow{\tau} \\
\nu \tilde{u}.\Delta'' & \triangleright ([c(w).\bar{z}\langle w \rangle]^{n_i} \parallel [\bar{c}\langle w_y \rangle]^{n_i} \parallel [\text{kill}]^{n_r} \parallel [!B]^{n_b} \parallel [C]^{n_c}) \xrightarrow{\tau} \\
\nu \tilde{u}.\Delta'' & \triangleright ([\bar{z}\langle w_y \rangle]^{n_i} \parallel [\text{kill}]^{n_r} \parallel [!B]^{n_b} \parallel [C]^{n_c}) \xrightarrow{\bar{z}\langle w_y \rangle @ n_i} \\
& T_2 \equiv \nu \tilde{u}.\Delta'' \triangleright ([\text{kill}]^{n_r} \parallel [!B]^{n_b} \parallel [C]^{n_c})
\end{aligned}$$

with $(T_1, T_2) \in \mathcal{R}$

- Pairs $S_2 \times R_2$

Case τ Here there exist infinite R_2 that can perform a τ (the controller still attempting to recreate the location), anyway all the moves can be matched by S_2 by doing nothing.

□

Inria

**RESEARCH CENTRE
GRENOBLE – RHÔNE-ALPES**

Inovallée
655 avenue de l'Europe Montbonnot
38334 Saint Ismier Cedex

Publisher
Inria
Domaine de Voluceau - Rocquencourt
BP 105 - 78153 Le Chesnay Cedex
inria.fr

ISSN 0249-6399