

AperTO - Archivio Istituzionale Open Access dell'Università di Torino

Chaperone Contracts for Higher-Order Sessions

This is a pre print version of the following article:

Original Citation:

Availability:

This version is available <http://hdl.handle.net/2318/1643911> since 2017-07-01T11:21:00Z

Terms of use:

Open Access

Anyone can freely access the full text of works made available as "Open Access". Works made available under a Creative Commons license can be used according to the terms and conditions of said license. Use of all other works requires consent of the right holder (author or publisher) if not exempted from copyright protection by the applicable law.

(Article begins on next page)

Chaperone Contracts for Higher-Order Sessions

Hernán Melgratti, Universidad de Buenos Aires, FCEyN and CONICET-UBA, ICC
Luca Padovani, Università di Torino

July 1, 2017

Abstract

Contracts have proved to be an effective mechanism that helps developers in identifying those modules of a program that violate the contracts of the functions and objects they use. In recent years, sessions have been established as a key mechanism for realizing inter-module communications in concurrent programs. Just like values flow into or out of a function or object, messages are sent on, and received from, a session endpoint. Unlike conventional functions and objects, however, the kind, direction, and properties of messages exchanged in a session may vary over time, as the session progresses. This feature of sessions calls for contracts that evolve along with the session they describe.

In this work, we extend to sessions the notion of chaperone contract (roughly, a contract that applies to a mutable object) and investigate the ramifications of contract monitoring in a higher-order language that features sessions. We give a characterization of a correct module, one that honors the contracts of the sessions it uses, and prove a blame theorem. Guided by the calculus, we describe a lightweight implementation of monitored sessions as an `OCaml` module with which programmers can benefit from static session type checking and dynamic contract monitoring using an off-the-shelf version of `OCaml`.

1 Introduction

The design-by-contract approach to software development Meyer [1992] promotes the usage of executable specifications called *contracts* to describe the mutual obligations that regulate the interaction between different modules. Contracts are embedded in code and checked at runtime to help developers identify faulty modules. Findler and Felleisen [2002] have shown that this approach is applicable also to languages with higher-order functions, despite the fact that checking whether a function satisfies a given contract is undecidable. The key idea is to defer the evaluation of the contract until the function is actually applied to its argument. At this point, the contract is disassembled and used to perform or delay checks on the supplied argument and the returned result. Specifying contracts for mutable objects poses additional problems, though, because checking whether an object satisfies a given contract at a particular point in time does not guarantee that the contract will be also satisfied later on, if the object is modified. To address these situations, Strickland et al. [2012] introduced *chaperone contracts*, an interposition mechanism whereby the object is protected by a proxy that exposes the same interface as the object and that takes actions to perform or delay checks on the values flowing into or out of the object.

In this paper we introduce chaperone contracts for *binary sessions*, which are private communication channels between pairs of processes. Just like mutable objects have operations for inserting and retrieving values, sessions have operations for sending and receiving messages. Unlike conventional objects, though, the *order* in which these operations can be performed is disciplined by a protocol specification called *session type*. Also, the type of exchanged messages may change over time. For these reasons, sessions are more closely related to *non-uniform objects* Ravara and Vasconcelos [2000], Gay et al. [2010], whose interface changes according to their state. Chaperone contracts for sessions can be used to complement session types with precise specifications concerning the content of messages and their relationship with the actual behavior of session participants.

Compared to the functional contracts of Findler and Felleisen [2002] and chaperone contracts of Strickland et al. [2012], the distinguishing feature of chaperone contracts for sessions is that they evolve as the monitored session progresses: not only the contracts of messages exchanged in the session may vary over time, but also the contract of the remainder of a session may depend upon messages that have been exchanged earlier on in the same or a different session.

More in detail, here are the main contributions of this work:

- We define a core functional language called λCoS featuring runtime contract monitoring for higher-order sessions in the style of Findler and Felleisen [2002] and Strickland et al. [2012]. Contracts are dynamically updated along with the session they monitor. We introduce a novel semantics for contract monitoring, called *big-step monitoring*, that is instrumental to the subsequent formal analysis of our monitoring system. The setting allows us to pinpoint unexplored aspects of contract monitoring in a language with linear resources.
- We give an operational characterization of *locally correct modules*, namely modules that honor the contracts of the sessions they use. We argue that this notion can help programmers write correct code. We also prove a *blame soundness* result stating that correct modules cannot be blamed even in the presence of buggy or malicious parties.
- Guided by λCoS , we implement chaperone contracts for sessions on top of an existing `OCaml` library for binary sessions Padovani [2017]. The implementation also supports runtime contract monitoring in the presence of higher-order sessions and its key aspects are portable to other programming languages as well.

Although contracts for sessions and corresponding monitoring techniques have been actively investigated (*cf.* Section 7), our proposal is the first one that follows the approach of Findler and Felleisen [2002], which has well-known assets: contracts are first-class entities written in the same language programmers are already accustomed to, they can be computed, passed as arguments, and returned as results, they can be used gradually within large systems and can describe whole protocols or just parts thereof. Also, monitoring is performed inline and requires no external tools or dedicated middleware.

Structure of the paper. The next section gives an overview of the key ingredients of session contracts and their usage. Although the section is informal, the given examples are “real” in the sense that they run using our implementation. The formal model of λCoS is given in Section 3 and its typing discipline in Section 4. To keep the formalization technically manageable, we make some simplifying assumptions on the operational semantics and the type system which are relaxed in the implementation. Section 5 addresses local correctness and blame soundness. The implementation is detailed in Section 6, whereas Section 7 discusses connections with related work in more detail. Section 8 summarizes our work and hints at some extensions and future developments. Supplementary technical material and proofs of the results presented in Sections 3–6 can be found in the Appendix.

2 A Programmer’s Viewpoint of Contracts for Sessions

In this section we provide a programmer’s viewpoint of contracts for sessions using the process networks depicted in Figure 1 as running examples. In these networks, the aim of module `User` is to obtain an element w resulting from two elements v_1, v_2 produced by `Source` and combined by `Operator`. For the sake of illustration, we assume that all the elements are integer numbers and that $w = v_1 \bmod v_2$.

According to the network on the left-hand side of Figure 1, `User` establishes two sessions x and y with `Source` and `Operator` respectively, it forwards every element v_i received from x on y , and receives the transformed element w from y . Session types allow us to formalize these protocols specifying the type, direction and order of messages exchanged within sessions. In this case, the session




Figure 1: Stream processing networks with first-order (left) and second-order (right) sessions.

endpoints x and y owned by **User** are typed as $x : ?\text{int}.\text{int}.\text{end}$ and $y : !\text{int}.\text{int}.\text{int}.\text{end}$. In words, **User** uses x to receive two integer numbers and y to send two numbers and then to receive another one. What session types do not describe are additional requirements and guarantees concerning the content of exchanged messages. In this scenario, for example, it could be sensible to specify that the second number sent from **User** to **Operator** should be different from zero and that the number sent from **Operator** to **User** is non negative. Such specifications – hereafter called *contracts* – could be used to monitor, at runtime, the interaction between **User** and **Operator** so as to detect contract violations and, hopefully, to guide programmers to the source of the problem.

The approach of Findler and Felleisen [2002] to contract monitoring rests on three key ingredients: (1) a set of combinators for writing contracts; (2) a mechanism for associating contracts with the entities being monitored; (3) a labeling of the modules involved that narrows the source of the problem in case a contract violation is detected. In our setting, the definition

```
let operator_chan = register operator_body operator_c "Operator"
```

registers a new service channel `operator_chan` that can be used to initiate sessions with **Operator**. Registration links together three pieces of information: the body of the process that handles each session initiated with **Operator** (`operator_body`, omitted here), the contract that **Operator** claims to satisfy (`operator_c`), and a symbolic label that identifies **Operator** (the string `"Operator"`). The contract `operator_c` is defined as

```
let operator_c =
  send_c any_c @@ send_c (flat_c (≠ 0)) @@ receive_c (flat_c (≥ 0)) @@ end_c
```

and its structure reflects that of the protocol $!\text{int}.\text{int}.\text{int}.\text{end}$ of the session it describes. Let us focus first on the four sub-contracts separated by `@@` and built with the `send_c`, `receive_c`, and `end_c` combinators. As their name suggest, each of these indicates either an output or an input or the termination of the session. Both `send_c` and `receive_c` have an argument which is itself a contract describing the exchanged message: depending on whether the operation is an input or an output, the contract specifies a requirement or a guarantee for the message. So, the contract `flat_c (≠ 0)` is satisfied by any non-zero integer, `flat_c (≥ 0)` is satisfied by any non-negative integer, and `any_c` is satisfied by any value. Note that the contract is written from the viewpoint of a client of **Operator**, such as **User**. The combinator `@@` plays the same role as the dot `'.'` in session types and composes contracts sequentially: it indicates that the contract associated with the endpoint, and the specified constraints on the exchanged messages, vary after each interaction: **Operator** first accepts an arbitrary integer number, it then expects a non-zero number from the client, and it finally sends back a non-negative number. As it turns out, `@@` is nothing but function application. Therefore, `operator_c` can be alternatively defined as

```
send_c any_c (send_c (flat_c (≠ 0)) (receive_c (flat_c (≥ 0)) end_c))
```

We will keep using `@@` in this section for the sake of readability. In the rest of the paper we will drop `@@` and use parentheses to disambiguate the structure of contracts when necessary.

The definitions

```
let source_c = receive_c any_c @@ receive_c any_c @@ end_c
let source_chan = register source_body source_c "Source"
```

register a service channel for `Source` in an analogous way. In this case, `source_c` does not specify any additional constraints with respect to the session type `?int.?int.end`. With `Source` and `Operator` in place, we are ready to implement `User`:

```

1 let user () =
2   let x = connect source_chan "User" in (* connect with Source *)
3   let y = connect operator_chan "User" in (* connect with Operator *)
4   let v1, x = receive x in (* receive v1 from Source *)
5   let v2, x = receive x in (* receive v2 from Source *)
6   let y = send v1 y in (* send v1 to Operator *)
7   let y = send v2 y in (* send v2 to Operator *)
8   let w, y = receive y in (* receive result *)
9   print_int w; close x; close y (* close sessions *)

```

After initiating the two sessions `x` and `y` (lines 2–3), `User` implements the behavior informally described above making use of the operations `send` and `receive` for session communications. The series of rebindings of the session endpoints `x` and `y` is a common trait of most implementations of sessions for functional languages and allows for the type of `x` and `y` to be appropriately updated after each operation. The `connect` primitive specifies the shared channel on which the sessions are initiated and a label that identifies the requesting party. This label, together with those specified at the time of service registration, is used by the runtime monitoring mechanism to pinpoint the source of a contract violation. Here, `"User"` identifies the module responsible for the messages sent on `x` and `y`, whereas `"Source"` and `"Operator"`, associated with `source_chan` and `operator_chan` at the time of their registration, identify the modules that are responsible for the messages received from `x` and `y`, respectively. So if v_2 turns out to be 0, the runtime monitoring of the `y` session endpoint flags `User` for breaching the contract with `Operator`. If, on the other hand, `Operator` sends a negative remainder back to `User`, then it is `Operator` that violates its own contract.

In this first example, the module responsible for a contract violation is easily identified with the sender of the message that triggers the violation. This is not always the case, though. Consider for instance the network on the right-hand side of Figure 1, which aims to achieve the same computation as before, except that `User` now *delegates* the session endpoint `x` to `Operator` so that `Operator` receives the elements v_1, v_2 directly from `Source`, rather than having them forwarded through `User`. In this case, the session type associated with `y` is `!(?int.?int.end).?int.end`, indicating that `User` first sends `Operator` a session endpoint of type `?int.?int.end` (that is precisely the type of `x`) and then behaves according to `?int.end` to receive the result from `Operator`. Correspondingly, the contract exposed by `Operator` could be the following

```

let operator_deleg_c =
  send_c (receive_c any_c @@ receive_c (flat_c (≠ 0)) @@ end_c) @@
  receive_c (flat_c (≥ 0)) @@ end_c

```

specifying that the first message received by `Operator` is a session endpoint from which the second receive element is supposed to be different from zero. That is, `operator_deleg_c` is a second-order contract for a second-order session.

The implementation of `User` becomes

```

1 let user_deleg () =
2   let x = connect source_chan "User" in
3   let y = connect operator_deleg_chan "User" in
4   let y = send x y in
5   let res, y = receive y in
6   print_int res; close y

```

which establishes the two sessions as before (lines 2–3), delegates `x` on `y` (line 4), and then receives the result from `Operator` (line 5) before closing the session (line 6). As before, the arrival of a zero v_2 element to `Operator` should trigger a contract violation, but establishing which process is to blame is not as easy as in the previous example. On the one hand, the offending element is sent

Table 1: Syntax of λCoS (terms that occur only at runtime and not in user code are shaded).

Process	P, Q	$::=$	$\langle e \rangle_p \mid a \leftarrow_p^c v \mid P \parallel Q \mid (\nu a)P$
Expression	e	$::=$	$v \mid x \mid e_1 e_2 \mid \text{let } x, y = e_1 \text{ in } e_2 \mid \text{case } e \text{ of } e_1 \mid e_2$ $\mid [e_1]^{e_2, p, q} \mid v \triangleleft^p e \mid \text{blame } p$
Value	v, w, C, D	$::=$	$\lambda x. e \mid \varepsilon$ $\mid \mathbf{c} \ v_1 \cdots v_n \quad (0 \leq n \leq \#\mathbf{c})$
Endpoint	ε	$::=$	$a^\iota \mid [\varepsilon]^{c, p, q}$
Constant	\mathbf{c}	$::=$	$() \mid \text{true} \mid \text{false} \mid \text{inl} \mid \text{inr} \mid \text{pair} \mid \text{dual}$ $\mid \text{connect} \mid \text{close} \mid \text{send} \mid \text{receive} \mid \text{branch} \mid \text{left} \mid \text{right}$ $\mid \text{flat_c} \mid \text{end_c} \mid \text{branch_c} \mid \text{choice_c}$ $\mid \text{send_c} \mid \text{send_d} \mid \text{receive_c} \mid \text{receive_d}$

by **Source** so, if we were to follow the simple rule for blame assignment used before, we would conclude that **Source** is to blame. On the other hand, **Source** has never violated the `source_c` contract it was registered with. The problem lies once again within **User**: by delegating `x` on `y`, **User** is claiming that the contract `source_c` associated with `x` entails the contract

```
receive_c any_c @@ receive_c (flat_c ( $\neq 0$ )) @@ end_c
```

but this is not true, because not every message satisfying `any_c` also satisfies `flat_c` ($\neq 0$). In summary, the actual offending operation is the delegation performed by **User**, but establishing this fact would require a decision procedure for contract entailment, which is instead an undecidable relation (contracts may contain arbitrarily complex predicates expressed using a Turing complete language). Therefore, the delegation is provisionally accepted as valid and triggers a suitable rearrangement of contracts and labels for the involved endpoints such that, in the event of a contract violation, **User** and not **Source** is blamed.

In the rest of the paper we extend and formalize the ideas sketched in this section. In particular, we will consider a superset of the contract combinators used in this section that allow us to specify *dependent, possibly branching contracts* whose structure depends on the content of previously exchanged messages. The formalization will culminate with a blame soundness result stating that a process cannot be blamed by the monitor if it always respects the contracts of the endpoints it uses.

3 Syntax and Semantics of λCoS

3.1 Syntax

We use infinite sets of *variables* x, y, z and *service/session channels* a, b, c . We use *polarities* $\iota \in \{+, -\}$ to distinguish the two endpoints of a session channel and write $\bar{\iota}$ for the dual polarity of ι , where $\bar{+} = -$ and $\bar{-} = +$. An *endpoint* is a pair a^ι made of a session channel a and a polarity ι and we say that $a^{\bar{\iota}}$ is the *peer* of a^ι . A *name* u is either a variable or a channel or an endpoint.

The syntax of λCoS comprises *processes* and *expressions* (Table 1). Expressions model the sequential part of programs and processes model parallel threads that invoke services and communicate through sessions. A process is either a *thread* $\langle e \rangle_p$ made of a body e and a label p , a *service* $a \leftarrow_p^c v$ that waits for invocations on the service channel a and spawns new threads with body v , the *parallel composition* $P \parallel Q$ of two processes P and Q , or a *session* $(\nu a)P$ with scope P . Each service advertises a contract \mathbf{c} that describes the intended interaction from the client's viewpoint. For simplicity, we assume that services cannot be created dynamically in λCoS . As we have seen in Section 2, the implementation is more liberal and allows new services to be dynamically created and registered. A *module* is a set of threads and services with the same label. We use *labels* p, q to identify modules and to assign blame.

Expressions include the standard constructs of the λ -calculus, pair splitting `let $x, y = e_1$ in e_2` , and pattern matching `case e of $e_1 \mid e_2$` . Additionally, we have three constructs that support runtime monitoring Findler and Felleisen [2002], Wadler and Findler [2009]. A *monitored expression* $[e_1]^{e_2, p, q}$ wraps e_1 with a contract e_2 and the labels p and q identify the modules responsible for the values respectively flowing out of and into e_1 . Such modules may be blamed if the flowing values do not satisfy the constraints specified by the contract e_2 . In λCoS , the intuition behind these flows of values is best illustrated by thinking of e_1 as a session endpoint. Then, p is the module responsible for the messages *received* from the endpoint, whereas q is the module responsible for the messages *sent* on the endpoint. Sometimes we use σ and ϱ to denote pairs of labels such as p, q , and we write $-\sigma$ for the symmetric pair obtained by swapping the components in σ . We say that p is the *positive* label and q the *negative* label in the pair p, q . Monitors may accumulate on top of each other yielding stacks of the form $[\dots [e]^{c_1, \sigma_1} \dots]^{c_n, \sigma_n}$ that we abbreviate as $[e]^{\overrightarrow{c, \sigma}}$. Values flowing out of e are checked against all the contracts c_1, \dots, c_n starting from c_1 , whereas values flowing into e are checked against all the contracts c_n, \dots, c_1 starting from c_n . It is sometimes necessary to reverse the order of monitors while swapping the elements of the pairs σ_i . In these cases we write $[e]^{\overleftarrow{c, -\sigma}}$ for $[\dots [e]^{c_n, -\sigma_n} \dots]^{c_1, -\sigma_1}$. A *busy monitor* $v \triangleleft^p e$ indicates an ongoing check being performed on value v . If e evaluates to **true** then the check is passed and the busy monitor reduces to v . If e evaluates to **false** then the busy monitor reduces to **blame p** signaling a contract violation of p . So, a busy monitor $v \triangleleft^p e$ is akin to a conditional **if e then v else blame p** .

We use v, w, C and D to range over *values*, which comprise abstractions, (monitored) endpoints, and (applied) constants. We reserve C and D for values that represent contracts, which are described below. Constants comprise some standard data constructors (unit, booleans, pairs, sums) and a standard set of session primitives for connecting to services (**connect**), sending and receiving messages (**send** and **receive**), selecting a choice (**left** and **right**) and offering a choice (**branch**).

Constants of the form $*_c$ and $*_d$ are *contract constructors*. A *flat contract* **flat_c** w is satisfied by every value v such that wv evaluates to **true**. We say that w is the *predicate* of the contract **flat_c** w .

The contract **end_c** describes an endpoint that can only be closed. The *non-dependent contracts* **send_c** $C D$ and **receive_c** $C D$ (sugared as $!C;D$ and $?C;D$) have a prefix C and a continuation D and describe endpoints used for respectively sending and receiving a message that satisfies C and then used according to D . These contracts are qualified as “non dependent” because the continuation contract D does *not* depend on the exchanged message. The *dependent contracts* **send_d** $C w$ and **receive_d** $C w$ (sugared as $!C.w$ and $?C.w$) describe endpoints used for respectively sending and receiving a message v that satisfies C and then used according to wv . That is, w is a function that produces the continuation contract when applied to the exchanged message. In principle, a non-dependent contract such as $!C;D$ could be treated as a degenerate dependent contract $!C.\lambda_.D$, where the function that computes the continuation contract is constant. However, the two contract forms require different typing policies and therefore must be kept distinct, at least in the formal model. In the implementation, which is based on a weaker type system, the two forms can be unified so that non-dependent contracts are indeed degenerate cases of dependent ones. The contracts **choice_c** $C D E$ and **branch_c** $C D E$ (sugared as $!C.D:E$ and $?C.D:E$) describe endpoints used for respectively selecting and offering a choice. The choice is effectively represented and transmitted as a boolean value v that satisfies C . The continuation contracts D and E describe the endpoint after the choice, depending on whether v is **true** or **false**. Compared to dependent contracts, choices also allow the session types of the continuations, and not just the contracts, to depend on the exchanged message. There is a key difference between flat and session contracts. Flat contracts are meant to be checked right away, turning a monitor into a busy monitor. Session contracts chaperone the endpoints they wrap and are checked only if and when the endpoints are used for input/output operations. Also, session contracts are dynamically updated as the endpoint they wrap is used.

The primitive **dual** computes the *dual* of a contract C , which specifies the same constraints as C except for the direction of messages, which is reversed. We will see how it is used later on.

Table 2: Reduction of expressions (contextual rule omitted).

[R1]	$(\lambda x. e)v \rightarrow e\{v/x\}$	[R8]	$\mathbf{dual\ end_c} \rightarrow \mathbf{end_c}$
[R2]	$\mathbf{let\ } x, y = (v, w) \mathbf{ in\ } e \rightarrow e\{v, w/x, y\}$	[R9]	$\mathbf{dual\ } !C; D \rightarrow ?C; \mathbf{dual\ } D$
[R3]	$\mathbf{case\ inl\ } v \mathbf{ of\ } e_1 \mid e_2 \rightarrow e_1 v$	[R10]	$\mathbf{dual\ } ?C; D \rightarrow !C; \mathbf{dual\ } D$
[R4]	$\mathbf{case\ inr\ } v \mathbf{ of\ } e_1 \mid e_2 \rightarrow e_2 v$	[R11]	$\mathbf{dual\ } !C. w \rightarrow ?C. \lambda x. \mathbf{dual\ } (wx)$
[R5]	$[v] \mathbf{flat_c\ } w, p, q \rightarrow v \triangleleft^p wv$	[R12]	$\mathbf{dual\ } ?C. w \rightarrow !C. \lambda x. \mathbf{dual\ } (wx)$
[R6]	$v \triangleleft^p \mathbf{true} \rightarrow v$	[R13]	$\mathbf{dual\ } !C. D : E \rightarrow ?C. \mathbf{dual\ } D : \mathbf{dual\ } E$
[R7]	$v \triangleleft^p \mathbf{false} \rightarrow \mathbf{blame\ } p$	[R14]	$\mathbf{dual\ } ?C. D : E \rightarrow !C. \mathbf{dual\ } D : \mathbf{dual\ } E$

An applied constant $\mathbf{c}\ v_1 \cdots v_n$ is a value only if $0 \leq n \leq \#\mathbf{c}$ where $\#\mathbf{c}$ denotes the arity of \mathbf{c} , that is the maximum number of arguments to which \mathbf{c} can be applied before possibly becoming a redex. In λCoS , $\mathbf{branch_c}$ and $\mathbf{choice_c}$ have arity 3, \mathbf{pair} , $\mathbf{send_c}$, $\mathbf{receive_c}$, $\mathbf{send_d}$ and $\mathbf{receive_d}$ have arity 2, \mathbf{inl} , \mathbf{inr} , \mathbf{send} and $\mathbf{flat_c}$ have arity 1, and all the other constants have arity 0.

We identify processes and expressions modulo α -renaming of bound names, considering that $(\nu a)P$ binds a , a^+ , and a^- in P , and we assume that bound names are all distinct. We say that P is a *user process* if it does not use any runtime syntax as by Table 1.

3.2 Semantics

Expressions reduce according to a call-by-value strategy, for which we define *evaluation contexts* thus:

$$\mathcal{E} ::= [] \mid \mathcal{E}e \mid v\mathcal{E} \mid [\mathcal{E}]^{e,\sigma} \mid [v]^{\mathcal{E},\sigma} \mid v \triangleleft^p \mathcal{E} \mid \mathbf{let\ } x, y = \mathcal{E} \mathbf{ in\ } e \mid \mathbf{case\ } \mathcal{E} \mathbf{ of\ } e_1 \mid e_2$$

The reduction rules for expressions are given in Table 2. Rules [R1–R4] are standard. According to rule [R5], a monitor with a flat contract $\mathbf{flat_c}\ w$ wrapping a value v turns into a busy monitor that checks whether v satisfies the contract by evaluating wv . When the evaluation of wv terminates, the busy monitor reduces to v if the check succeeds (see rule [R6]) or blames p otherwise (see rule [R7]). Note that [R5] *duplicates* v in the reduct. Since endpoints are linear resources, the definition of flat contracts on endpoints (and on linear resources in general) will be forbidden by the type system.

Rules [R8–R14] compute the dual of a contract by reversing the orientation of interactions. For instance, given the contract $C \stackrel{\text{def}}{=} !(\mathbf{flat_c} (\geq 0)); ?(\mathbf{flat_c} (\leq 3)); \mathbf{end_c}$ we have

$$\mathbf{dual\ } C \rightarrow \rightarrow \rightarrow ?(\mathbf{flat_c} (\geq 0)); !(\mathbf{flat_c} (\leq 3)); \mathbf{end_c}$$

Processes reduce according to the rules in Table 3, where we stack parallel threads instead of separating them by \parallel whenever there is not enough space. Rule [R15] models the initiation of a session through the service channel a . In the reduct, the endpoint b^+ of the new session is returned to the client and the endpoint b^- is passed to the body v of the service in a new thread. The endpoints b^+ and b^- are respectively monitored by the contracts C and $\mathbf{dual\ } C$. The labels in the two monitors reflect the direction of the messages exchanged over the endpoints: in the client, p is negative since p is responsible for the messages sent to the service and q is positive since q is responsible for the messages received by the client; in the spawned thread the responsibilities are reversed. Rule [R16] models the closing of a session.

Rule [R17] models a communication from thread p to thread q on session a , when the endpoints of a are monitored by non-dependent contracts. The rule appears more complicated than it really is. If we erase all the monitors, the rule boils down to

$$\langle \mathcal{E}[\mathbf{send\ } v\ a^+] \rangle_p \parallel \langle \mathcal{E}'[\mathbf{receive\ } a^-] \rangle_q \rightarrow \langle \mathcal{E}[a^+] \rangle_p \parallel \langle \mathcal{E}'[(v, a^-)] \rangle_q$$

Table 3: Reduction of processes (contextual rules omitted).

$$\begin{array}{l}
\text{[R15]} \quad \left\langle \mathcal{E}[\text{connect } a] \right\rangle_p \left\{ \begin{array}{l} \langle \mathcal{E}[\text{connect } a] \rangle_p \\ a \leftarrow_q^c v \end{array} \right\} \rightarrow (\nu b) \left(\langle \mathcal{E}[[b^+]^{C,q,p}] \rangle_p \parallel \langle v [b^-]^{\text{dual } C,p,q} \rangle_q \right) \parallel a \leftarrow_q^c v \quad b \text{ fresh} \\
\text{[R16]} \quad (\nu a) \left\langle \mathcal{E}[\text{close } [a^+]^{\text{end}_c, \sigma}] \right\rangle_p \left\langle \mathcal{E}'[\text{close } [a^-]^{\text{end}_c, \hat{\ell}}] \right\rangle_q \rightarrow \langle \mathcal{E}[\langle \rangle] \rangle_p \parallel \langle \mathcal{E}'[\langle \rangle] \rangle_q \\
\text{[R17]} \quad \left\langle \mathcal{E}[\text{send } v [a^t]^{\text{!}C; D, \sigma}] \right\rangle_p \left\langle \mathcal{E}'[\text{receive } [a^{\bar{t}}]^{\text{?}E; F, \hat{\ell}}] \right\rangle_q \rightarrow \left\langle \mathcal{E}[[a^t]^{\overline{D}, \sigma}] \right\rangle_p \left\langle \mathcal{E}'[\langle \langle [v]^{\overleftarrow{C}, \neg\sigma} \rangle^{\overrightarrow{E}, \hat{\ell}}, [a^{\bar{t}}]^{\overrightarrow{F}, \hat{\ell}} \rangle] \right\rangle_q \\
\text{[R18]} \quad \left\langle \mathcal{E}[\text{send } v [a^t]^{\text{!}C.w_1, \sigma}] \right\rangle_p \left\langle \mathcal{E}'[\text{receive } [a^{\bar{t}}]^{\text{?}D.w_2, \hat{\ell}}] \right\rangle_q \rightarrow \left\langle \mathcal{E}[[a^t]^{\overline{w_1 v}, \sigma}] \right\rangle_p \left\langle \mathcal{E}'[\langle \langle [v]^{\overleftarrow{C}, \neg\sigma} \rangle^{\overrightarrow{D}, \hat{\ell}}, [a^{\bar{t}}]^{\overrightarrow{w_2 v}, \hat{\ell}} \rangle] \right\rangle_q \\
\text{[R19]} \quad \left\langle \mathcal{E}[\text{left } [a^t]^{\text{!}C.D; E, \sigma}] \right\rangle_p \left\langle \mathcal{E}'[\text{branch } [a^{\bar{t}}]^{\text{?}F.G; H, \hat{\ell}}] \right\rangle_q \rightarrow \left\langle \mathcal{E}[[a^t]^{\overline{D}, \sigma}] \right\rangle_p \left\langle \mathcal{E}'[\langle \lambda_ . \text{inl } [a^{\bar{t}}]^{\overrightarrow{G}, \hat{\ell}} \rangle \langle [\text{true}]^{\overleftarrow{C}, \neg\sigma} \rangle^{\overrightarrow{F}, \hat{\ell}}] \right\rangle_q \\
\text{[R20]} \quad \left\langle \mathcal{E}[\text{right } [a^t]^{\text{!}C.D; E, \sigma}] \right\rangle_p \left\langle \mathcal{E}'[\text{branch } [a^{\bar{t}}]^{\text{?}F.G; H, \hat{\ell}}] \right\rangle_q \rightarrow \left\langle \mathcal{E}[[a^t]^{\overline{E}, \sigma}] \right\rangle_p \left\langle \mathcal{E}'[\langle \lambda_ . \text{inr } [a^{\bar{t}}]^{\overrightarrow{H}, \hat{\ell}} \rangle \langle [\text{false}]^{\overleftarrow{C}, \neg\sigma} \rangle^{\overrightarrow{F}, \hat{\ell}}] \right\rangle_q
\end{array}$$

where we see that the message v just moves from the sender to the receiver. All the additional machinery in [R17] handles contracts. In particular, a^t in the sender is monitored by a contract of the form $\text{!}C_i; D_i$ and $a^{\bar{t}}$ in the receiver is monitored by a contract of the form $\text{?}E_j; F_j$. In the reduct, these contracts must be updated to D_i and F_j to reflect the fact that the communication has occurred. Also, the message v is wrapped by all the contracts C_i and E_j it is meant to satisfy. Since v is flowing into a^t , the contracts C_i are stacked inside out so that the checks on v are performed in the right order. Also, the blame labels are swapped to reflect the correct responsibilities in case contract violations are detected.

Rule [R18] is similar to [R17], except that the two endpoints before the communication are monitored by dependent contracts $\text{!}C_i.w_{1i}$ and $\text{?}D_j.w_{2j}$ and the continuation contracts are obtained by applying w_{1i} and w_{2j} to the message. Note that v occurs *several times* in the reduct. As we have discussed for [R5], this could be problematic if v contains linear values such as session endpoints. Hence, the typing of `send_d` and `receive_d` will prevent the definition of contracts that depend on linear values. For simplicity, rules [R17] and [R18] are formulated in such a way that communications occur only when the contracts in the monitors stacked around the involved endpoints are all dependent or all non dependent. This restriction has been introduced solely to avoid considering all the possible combinations of dependent and non-dependent contracts and is relaxed in the implementation, where such distinction is blurred (Section 6.1).

Rules [R19] and [R20] model communications whereby *both* the session type *and* the contract of the endpoints in the reduct (may) depend upon the exchanged message, which is implicitly assumed to be `true` in [R19] and `false` in [R20]. As in [R17] and [R18] the message is checked against a suitable stack of contracts. Then, the receiver injects the continuation endpoint using either `inl` or `inr` correspondingly. Observe that not all of the contracts appearing in the redex also appear in the reduct. In prospect of devising a substructural typing discipline for λCoS , this suggests that contracts should not contain linear resources such as session endpoints.

We stress two facts concerning the reduction rules [R17–R20]. First, we observe that the whole stacks of monitors wrapping the endpoints on which the communication takes place are rearranged in a single reduction step. This is what we mean by “big-step” monitoring semantics, as opposed to the “small-step” monitoring semantics where each monitor is rearranged independently of the others Findler and Felleisen [2002], Wadler and Findler [2009]. The big-step semantics allows us to rely on some key properties of monitor configurations (Proposition 1) when proving the blame soundness results (Section 5). Second, we observe that part of the contracts on the sender side migrate to the receiving side after the communication. For delegations – that is the communication of session endpoints – the migration results in the delegated endpoint being wrapped by further

chaperone contracts in addition to those it already has. This is necessary to keep track of the responsibilities of the sender even after the communication has occurred (Example 2). For other messages not containing endpoints the migration essentially means that message contracts are always checked on the receiver side, *after* the communication has occurred. In the implementation (Section 6) we consider the alternative “small-step” monitoring semantics where flat contracts for messages are checked on the sender side, *before* communication takes place.

We omit the obvious rules that lift reduction of expression to processes and that close reductions by (process) contexts and structural congruence. The details can be found in Appendix A. In the following we write \rightarrow^* for the reflexive, transitive closure of the \rightarrow relations for expressions and processes. We also introduce some convenient notation for the evaluation of expressions and predicates:

Definition 1 (evaluation). We write $e \Downarrow v$ if $e \rightarrow^* v$ and we write $v \in w$ if $wv \Downarrow \mathbf{true}$.

Example 1 (function contracts). In this example we show an encoding of function contracts in the style of Findler and Felleisen [2002] as session contracts. Suppose that f is a function and $C \mapsto D$ its contract, where C and D are respectively the contracts for the domain and codomain of f . We may represent the function f as a service and the contract $C \mapsto D$ as the session contract $!C;?D;\mathbf{end_c}$, thus:

$$a_f \leftarrow_p^{!C;?D;\mathbf{end_c}} \lambda x. \mathbf{let } y, x = \mathbf{receive } x \mathbf{ in let } x = \mathbf{send } (f y) x \mathbf{ in close } x$$

The contract $!C;?D;\mathbf{end_c}$ gives a natural description of the intended usage protocol for f : the client must send f an argument that satisfies C and will receive back from f a value that satisfies D . Accordingly, an application $(f v)$ can be modeled as follows:

$$\mathbf{let } x = \mathbf{connect } a_f \mathbf{ in let } x = \mathbf{send } v x \mathbf{ in let } y, x = \mathbf{receive } x \mathbf{ in close } x; \dots y \dots$$

Thanks to delegations, the encoding shown here accounts also for higher-order functions and its generalization to dependent functions is straightforward. However, the encoding is purely conceptual and not meant to suggest a practical implementation. Its purpose is to show that the blame soundness results presented in Section 5 are comprehensive enough to also account for the functional part of the calculus, which we have deliberately neglected to keep the formal development as simple and focused on sessions as possible. ■

4 Type System

We define a session type system for λCoS loosely inspired to that of Gay and Vasconcelos [2010]. Following Tov and Pucella [2011], we use *kinds* to distinguish *unlimited* types, those denoting values that can be discarded and duplicated, from *linear* types, those denoting values (such as endpoints) that must be used exactly once. The syntax of kinds, types, and session types is given below:

$$\begin{array}{ll} \mathbf{Kind} & \kappa ::= 1 \mid \omega \\ \mathbf{Type} & t, s ::= \mathbf{unit} \mid t \times s \mid t + s \mid t \rightarrow^\kappa s \mid [t] \mid \#T \mid T \\ \mathbf{Session Type} & T, S ::= \mathbf{end} \mid ?t.T \mid !t.T \mid T \& S \mid T \oplus S \end{array}$$

Types, ranged over by t, s , include base types and standard type constructors. Arrows \rightarrow^κ have a kind annotation κ that indicates whether a function can be applied an arbitrary number of times ($\kappa = \omega$) or must be applied exactly once ($\kappa = 1$). This latter constraint is necessary if the function contains one or more linear values in its closure. We will abbreviate \rightarrow^ω with \rightarrow . The *contract type* $[t]$ describes contracts for values of type t and the *shared channel type* $\#T$ describes channels for initiating sessions with session type T . Session types have the standard constructors for denoting depleted session endpoints (\mathbf{end}), input/output operations ($?t.T$ and $!t.T$), branches ($T \& S$) and choices ($T \oplus S$).

We now define a relation that assigns kinds to each type. Every type t has kind 1, since using a value of type t exactly once is always safe regardless of t . By contrast, only those types denoting values that can be safely duplicated or discarded have kind ω . Formally:

Table 4: Type schemes of λCoS constants.

$() : \text{unit}$	$\text{pair} : t \rightarrow s \rightarrow^\kappa t \times s$	$t :: \kappa$
$\text{true} : \text{bool}$	$\text{send} : t \rightarrow !t. T \rightarrow^\kappa T$	$t :: \kappa$
$\text{false} : \text{bool}$	$\text{receive} : ?t. T \rightarrow t \times T$	
$\text{inl} : t \rightarrow t + s$	$\text{flat_c} : (t \rightarrow \text{bool}) \rightarrow [t]$	$t :: \omega$
$\text{inr} : s \rightarrow t + s$	$\text{end_c} : [\text{end}]$	
$\text{connect} : \#T \rightarrow T$	$\text{send_c} : [t] \rightarrow [T] \rightarrow [!t. T]$	
$\text{close} : \text{end} \rightarrow \text{unit}$	$\text{receive_c} : [t] \rightarrow [T] \rightarrow [?t. T]$	
$\text{left} : T \circledast S \rightarrow T$	$\text{send_d} : [t] \rightarrow (t \rightarrow [T]) \rightarrow [!t. T]$	$t :: \omega$
$\text{right} : T \circledast S \rightarrow S$	$\text{receive_d} : [t] \rightarrow (t \rightarrow [T]) \rightarrow [?t. T]$	$t :: \omega$
$\text{branch} : T \& S \rightarrow T + S$	$\text{choice_c} : [\text{bool}] \rightarrow [T] \rightarrow [S] \rightarrow [T \circledast S]$	
$\text{dual} : [T] \rightarrow [\bar{T}]$	$\text{branch_c} : [\text{bool}] \rightarrow [T] \rightarrow [S] \rightarrow [T \& S]$	

Definition 2 (kinding). We say that t has kind κ if $t :: \kappa$ is derivable by:

$$\text{unit} :: \omega \quad [t] :: \omega \quad \#T :: \omega \quad t \rightarrow^\kappa s :: \kappa \quad \frac{t :: \kappa \quad s :: \kappa}{t \odot s :: \kappa} \quad \odot \in \{\times, +\} \quad \frac{t :: \omega}{t :: 1}$$

Note that contract types and shared channel types have kind ω , whereas the kind of a sum or product depends on that of its components. In particular, a sum/product has kind ω only if both its components do as well. We say that t is *unlimited* if $t :: \omega$ and that t is *linear* otherwise, namely if $t :: \kappa$ implies $\kappa = 1$. For example, $?int.end$ and $int \rightarrow^1 \text{bool}$ and $int + ?int.end$ are all linear, but $?int.end \rightarrow \text{bool}$ is not.

We introduce *type environments* to keep track of the type of free names in expressions and processes. A type environment Γ is a finite map from names to types written $u_1 : t_1, \dots, u_n : t_n$. We write \emptyset for the empty type environment, $\text{dom}(\Gamma)$ for the domain of Γ , namely the set of names for which there is an association in Γ , and Γ, Γ' for the union of Γ and Γ' when $\text{dom}(\Gamma) \cap \text{dom}(\Gamma') = \emptyset$. We extend the notion of kinding to type environments and write $\Gamma :: \kappa$ if $\Gamma(u) :: \kappa$ for every $u \in \text{dom}(\Gamma)$.

As usual, we need a way of splitting and combining type environments that is more flexible than disjoint union and, at the same time, prevents the duplication of linear resources. We therefore define a (partial) combination operator $+$ analogous to the one given by Kobayashi et al. [1999]:

Definition 3 (environment combination). We write $+$ for the partial operation on type environments such that:

$$\begin{aligned} \Gamma + \Gamma' &\stackrel{\text{def}}{=} \Gamma, \Gamma' && \text{if } \text{dom}(\Gamma) \cap \text{dom}(\Gamma') = \emptyset \\ (\Gamma, u : t) + (\Gamma', u : t) &\stackrel{\text{def}}{=} (\Gamma + \Gamma'), u : t && \text{if } t :: \omega \end{aligned}$$

Note that $\Gamma + \Gamma'$ is undefined if there are associations $u : t \in \Gamma$ and $u : t' \in \Gamma'$ for the same name u such that either t and t' are different or at least one of them is linear. This prevents multiple uses of a linear resource. Note also that $\Gamma + \Gamma$ is always defined and equal to Γ itself when $\Gamma :: \omega$.

To ensure communication safety, threads are required to perform complementary actions on the peer endpoints of the same session. This is enforced by assigning *dual* session types to peer endpoints. Roughly, the dual of a session type is obtained by swapping inputs with outputs and choices with branches. Formally:

Definition 4 (session type duality). The *dual* of T is the session type \bar{T} defined as:

$$\overline{\text{end}} = \text{end} \quad \overline{?t. \bar{T}} = !t. \bar{T} \quad \overline{!t. \bar{T}} = ?t. \bar{T} \quad \overline{T \& \bar{S}} = \bar{T} \circledast \bar{S} \quad \overline{T \circledast \bar{S}} = \bar{T} \& \bar{S}$$

Note that duality is an involution, that is $\overline{\bar{T}} = T$.

Table 5: Typing rules.

Typing rules for expressions

 $\boxed{\Gamma \vdash e : t}$

$\frac{[\text{T-CONST}]}{\Gamma :: \omega \quad \mathbf{c} : t}{\Gamma \vdash \mathbf{c} : t}$	$\frac{[\text{T-NAME}]}{\Gamma :: \omega}{\Gamma, u : t \vdash u : t}$	$\frac{[\text{T-BLAME}]}{\Gamma \vdash \mathbf{blame} \ p : t}$	$\frac{[\text{T-BUSY-MONITOR}]}{\Gamma_1 \vdash e : \mathbf{bool} \quad \Gamma_2 \vdash v : t}{\Gamma_1 + \Gamma_2 \vdash v \triangleleft^p e : t}$
$\frac{[\text{T-FUN}]}{\Gamma, x : t \vdash e : s \quad \Gamma :: \kappa}{\Gamma \vdash \lambda x. e : t \rightarrow^\kappa s}$	$\frac{[\text{T-APP}]}{\Gamma_1 \vdash e_1 : t \rightarrow^\kappa s \quad \Gamma_2 \vdash e_2 : t}{\Gamma_1 + \Gamma_2 \vdash e_1 e_2 : s}$	$\frac{[\text{T-MONITOR}]}{\Gamma_1 \vdash e_1 : t \quad \Gamma_2 \vdash e_2 : [t]}{\Gamma_1 + \Gamma_2 \vdash [e_1]^{e_2, p, q} : t}$	
$\frac{[\text{T-SPLIT}]}{\Gamma_1 \vdash e_1 : t_1 \times t_2 \quad \Gamma_2, x : t_1, y : t_2 \vdash e_2 : t}{\Gamma_1 + \Gamma_2 \vdash \mathbf{let} \ x, y = e_1 \ \mathbf{in} \ e_2 : t}$		$\frac{[\text{T-CASE}]}{\Gamma_1 \vdash e : t_1 + t_2 \quad \Gamma_2 \vdash e_i : t_i \rightarrow^{\kappa_i} t \ (i=1,2)}{\Gamma_1 + \Gamma_2 \vdash \mathbf{case} \ e \ \mathbf{of} \ e_1 \mid e_2 : t}$	

Typing rules for processes

 $\boxed{\Gamma \vdash P}$

$\frac{[\text{T-THREAD}]}{\Gamma \vdash e : \mathbf{unit}}{\Gamma \vdash \langle e \rangle_p}$	$\frac{[\text{T-PAR}]}{\Gamma_i \vdash P_i \ (i=1,2)}{\Gamma_1 + \Gamma_2 \vdash P_1 \parallel P_2}$	$\frac{[\text{T-SESSION}]}{\Gamma, a^+ : T, a^- : \bar{T} \vdash P}{\Gamma \vdash (\nu a)P}$	$\frac{[\text{T-SERVICE}]}{\emptyset \vdash C : [T] \quad \Gamma \vdash v : \bar{T} \rightarrow \mathbf{unit}}{\Gamma + a : \#T \vdash a \leftarrow_p^C v}$
--	--	--	---

Table 4 shows the type schemes of λCoS constants as associations of the form $\mathbf{c} : t$. In general, each constant may have infinitely many types. The types of data constructors are standard, the only quirk being the second arrow in the type of `pair` which has the same kind as the first element of the pair. This accounts for the possibility that such element has a linear type, in which case exactly one pair must be created to avoid duplicating the element. For example, the partial application `pair` a^t is a function which must be applied exactly once in order to preserve the linearity of a^t . The types of the communication primitives are essentially those given by Gay and Vasconcelos [2010] and the types of the contract constructors follow from their semantics. Note that `flat_c` and the dependent contracts `receive_d` and `send_d` require the tested value/message to have an unlimited type, since it may be duplicated.

The typing rules for λCoS (Table 5) are fairly standard, so we only comment on a few notable features. In rules $[\text{T-CONST}]$ and $[\text{T-NAME}]$ the condition $\Gamma :: \omega$ makes sure that the unused part of type environments does not contain linear resources. Rule $[\text{T-FUN}]$ requires the arrow type of a function to have the same kind as the environment in which the function is typed. If there is a linear type in the environment, meaning that the function has a linear value in its closure, then the function can only be applied once. Rule $[\text{T-APP}]$ makes use of the $+$ operator for type environments to keep track of the usage of linear resources occurring in the function and its argument. Rules $[\text{T-SPLIT}]$ and $[\text{T-CASE}]$ are standard. Note once again the use of $+$ for combining type environments. A monitored expression $[e_1]^{e_2, p, q}$ is well typed provided that e_2 is a contract for values of the same type as that of e_1 . A busy monitor is well typed provided that the condition being checked has type `bool`. Rule $[\text{T-BLAME}]$ states that a blame has any type and is always well typed regardless of the environment. We do not treat blame recovery in the formal model.

Concerning processes, $[\text{T-THREAD}]$ and $[\text{T-PAR}]$ are as expected. Rule $[\text{T-SESSION}]$ introduces a session as a pair of peer endpoints with dual session types. Finally, $[\text{T-SERVICE}]$ requires a well-typed service to advertise a contract of the appropriate type and to have a body that accepts one endpoint at each invocation. The body must be unlimited to allow for an arbitrary number of session initiations.

We state subject reduction as a basic sanity check for the type system. Other standard properties – *e.g.* communication safety and session fidelity – also hold, but are irrelevant in this paper.

Theorem 1. *If $\Gamma :: \omega$ and $\Gamma \vdash P$ and $P \rightarrow Q$, then $\Gamma \vdash Q$.*

A suitable generalization of Theorem 1, where Γ is allowed to contain linear names, is given in Appendix B.

5 Blame Soundness

In this section we present the soundness of the monitoring mechanism in λCoS , namely the property that in a well-typed program P in which a module p “behaves well”, p cannot be blamed. We proceed according to the following roadmap. First, we introduce the notion of *contract entailment* to specify when a contract is “more demanding than” another (Section 5.1). Entailment is a natural generalization of subtyping of session types Gay and Hole [2005]. Using entailment, we formalize the notion of *locally correct module p* as a module that always honors the contracts of the endpoints it uses. Locality refers to the fact that the correctness of p solely depends on the actions performed by, and on information known to, the module p itself (Section 5.2). Finally, we characterize the *soundness* of a module p as a set of invariant properties of the (busy) monitors in which the label p occurs. A direct consequence of soundness is that a well-typed, locally correct module p cannot be blamed (Section 5.3).

5.1 Contract Entailment

Contract entailment is a relation \leq such that, when $C \leq D$ holds, each value that satisfies C can be used where a value that satisfies D is expected. When C and D are flat contracts, \leq boils down to the set-theoretic inclusion between the values that satisfy the respective predicates. For example,

$$\text{flat_c } (\geq 3) \leq \text{flat_c } (\geq 0)$$

since every number greater than or equal to 3 is also greater than or equal to 0. To define entailment when C and D are session contracts, it helps to recall the analogy of contracts as specifications for the messages that can be sent on and received from a session endpoint. In this case, $C \leq D$ holds if two conditions are satisfied:

1. Every message that can be received from an endpoint satisfying C can also be received from an endpoint satisfying D .
2. Every message that can be sent on an endpoint satisfying D can also be sent on an endpoint satisfying C .

Note that C and D occur in different orders according to the direction of exchanged messages. With these intuitions, we formalize entailment below:

Definition 5 (contract entailment). We say that e_1 *entails* e_2 , written $e_1 \leq e_2$, if one of following conditions holds:

1. $e_1 \Downarrow \text{flat_c } w_1$ and $e_2 \Downarrow \text{flat_c } w_2$ and $v \in w_1$ implies $v \in w_2$;
2. $e_1 \Downarrow \text{end_c}$ and $e_2 \Downarrow \text{end_c}$;
3. $e_1 \Downarrow !C_1; D_1$ and $e_2 \Downarrow !C_2; D_2$ and $C_2 \leq C_1$ and $D_1 \leq D_2$;
4. $e_1 \Downarrow ?C_1; D_1$ and $e_2 \Downarrow ?C_2; D_2$ and $C_1 \leq C_2$ and $D_1 \leq D_2$;
5. $e_1 \Downarrow !\text{flat_c } v_1.w_1$ and $e_2 \Downarrow !\text{flat_c } v_2.w_2$ and $v \in v_2$ implies $v \in v_1$ and $w_1v \leq w_2v$;
6. $e_1 \Downarrow ?\text{flat_c } v_1.w_1$ and $e_2 \Downarrow ?\text{flat_c } v_2.w_2$ and $v \in v_1$ implies $v \in v_2$ and $w_1v \leq w_2v$;
7. $e_1 \Downarrow !C_1.D_1.E_1$ and $e_2 \Downarrow !C_2.D_2.E_2$ and $C_2 \leq C_1$ and $D_1 \leq D_2$ and $E_1 \leq E_2$;
8. $e_1 \Downarrow ?C_1.D_1.E_1$ and $e_2 \Downarrow ?C_2.D_2.E_2$ and $C_1 \leq C_2$ and $D_1 \leq D_2$ and $E_1 \leq E_2$.

Condition 1 formalizes the set-theoretic inclusion relation between sets of values that satisfy given predicates, whereas condition 2 relates the contract `end_c` with itself. Conditions 3–4 deal with non-dependent contracts. Entailment is covariant on input prefixes, contravariant on output prefixes, and always covariant on continuation contracts. For example, we have

$$!flat_c (\geq 0); end_c \leq !flat_c (\geq 3); end_c$$

because the contract on the left-hand side imposes weaker requirements on the messages that can be sent on the endpoint. On the other hand we have

$$?flat_c (\geq 3); end_c \leq ?flat_c (\geq 0); end_c$$

for the contract on the left-hand side provides stronger guarantees on the messages that can be received from the endpoint.

Conditions 5–6 deal with dependent contracts. The prefix contracts must be flat, since these are the only contracts that can specify constraints on unlimited values. Dependent contracts essentially behave as non-dependent ones, with contravariant outputs and covariant inputs, except that the continuation contracts may depend upon the exchanged message v . For example, we have

$$!flat_c (\geq 0). \lambda x. ?flat_c (\leq x); end_c \leq !flat_c (\geq 3). \lambda_. ?flat_c \lambda_. true; end_c$$

where hereafter we write $_$ for irrelevant names and sub-terms.

Conditions 7–8 deal with choices and branches. Recalling that a choice is akin to an output and a branch is akin to an input, the definition of entailment essentially follows the same contravariant/covariant pattern we have already seen in the other cases.

Overall, contract entailment is analogous to the subtyping relation for session types Gay and Hole [2005]. This analogy provides us with the same substitution principle that drives subtyping: when $C \leq D$, it is safe to use an endpoint with contract C wherever an endpoint with contract D is expected. We will make use of this analogy when defining locally correct modules (Definition 6). There are two traits of entailment that set it apart from subtyping. The first one is that entailment is a relation between terms and therefore is undecidable in general. In our setting this is not an issue because the decision as to whether one contract entails another one is not meant to be computed, but is left to the programmer (Section 5.2). The second difference is that entailment is not reflexive in general: according to Definition 5, it must be possible to evaluate contracts solely using the reduction relation for expressions, whereas it is possible to write (well-typed) contracts that make use of communication primitives and that reduce only when they occur within a suitable process. In this sense, by embracing Definition 5 we take the viewpoint that contracts should be *pure* Meyer [1992] and their evaluation should not involve side effects. This condition cannot be enforced solely by our type system. From now on, we will make the assumption that $C \leq C$ holds whenever C is advertised by a service $a \leftarrow_p^C v$. It is easy to see that entailment is transitive.

5.2 Locally Correct Modules

A *locally correct module* honors the contracts of all the endpoints it uses. Since contracts may depend upon the messages exchanged on such endpoints, we resort to an operational definition of correctness that takes into account all the possible reductions of a process in which label p occurs. The “local” qualification is meant to stress the fact that correctness of a module p only depends on actions performed by, and information known to, the module p itself.

To identify the actions performed by a module p , we introduce *p-contexts* \mathcal{P}_p as terms generated by the grammar

$$\mathcal{P}_p ::= \langle \mathcal{E} \rangle_p \mid (\mathcal{P}_p \parallel P) \mid (P \parallel \mathcal{P}_p) \mid (\nu a) \mathcal{P}_p$$

where \mathcal{E} is a plain evaluation context. As usual, we write $\mathcal{P}_p[e]$ for the process obtained by filling the hole in \mathcal{P}_p with e . Locally correct modules are defined below.

Definition 6 (local correctness). Let \mathcal{C} be the largest relation between blame labels and processes such that $p \mathcal{C} P$ implies:

1. $P = \mathcal{P}_p[\text{send } v \text{ } _]^{\text{!flat_c } w; _ _ _ _ _ _}$ implies $v \in w$, and
2. $P = \mathcal{P}_p[\text{send } v \text{ } _]^{\text{!flat_c } w _ _ _ _ _ _}$ implies $v \in w$, and
3. $P = \mathcal{P}_p[\text{send } _ _]^{\text{C}; _ _ _ _ _ _}$ implies $C \leq D$, and
4. $P = \mathcal{P}_p[\text{left } _ _]^{\text{!flat_c } w _ _ _ _ _ _}$ implies $\text{true} \in w$, and
5. $P = \mathcal{P}_p[\text{right } _ _]^{\text{!flat_c } w _ _ _ _ _ _}$ implies $\text{false} \in w$, and
6. $P \rightarrow Q$ implies $p \mathcal{C} Q$.

We say that module p is (*locally*) *correct* in P if $p \mathcal{C} P$ holds.

Condition 1 requires a message v sent over an endpoint monitored by a non-dependent contract $\text{!flat_c } w; _$ to satisfy the predicate w . Condition 2 is similar, except that the contract is a dependent one. Condition 3 concerns delegations whereby an endpoint with contract C is sent as a message that is supposed to satisfy contract D . According to the substitution principle, this is safe if C entails D . Conditions 4–5 concern choices. In these cases, the selected label (either true or false) is required to satisfy the predicate specified by the choice contract. Finally, condition 6 requires the previous conditions to hold for every possible reduction of P .

Example 2. To illustrate the notion of locally correct module, consider the services

$$\begin{aligned} a &\leftarrow_q^C \lambda x_1. \text{let } y_1 = \text{connect } b \text{ in let } x_2 = \text{send } y_1 \ x_1 \text{ in close } x_2 \\ b &\leftarrow_r^D \lambda y_1. \text{let } z, y_2 = \text{receive } y_1 \text{ in print } z^{-1}; \text{close } y_2 \end{aligned}$$

with labels q and r and the contracts

$$\begin{aligned} C &\stackrel{\text{def}}{=} ?E; \text{end_c} & D &\stackrel{\text{def}}{=} !F; \text{end_c} & E &\stackrel{\text{def}}{=} !\text{any_c}; \text{end_c} & F &\stackrel{\text{def}}{=} \text{flat_c} (\neq 0) \\ \bar{C} &\stackrel{\text{def}}{=} !E; \text{end_c} & \bar{D} &\stackrel{\text{def}}{=} ?F; \text{end_c} & & & & \end{aligned}$$

noting that C and D are the contracts advertised by q and r , respectively. Also note that $\text{dual } C \Downarrow \bar{C}$ and $\text{dual } D \Downarrow \bar{D}$. Service r receives a non-zero number z and prints its inverse. Service q initiates a session with r and then delegates the corresponding endpoint y_1 to its client. Consider also the p -labeled client

$$\langle \text{let } x_1 = \text{connect } a \text{ in} \\ \text{let } y_1, x_2 = \text{receive } x_1 \text{ in} \\ \text{let } y_2 = \text{send } 0 \ y_1 \text{ in close } x_2; \text{close } y_2 \rangle_p$$

which connects to service q , receives a session endpoint y_1 from it, and then sends 0 on y_1 . Observe that p connects to a service whose contract is $?(!\text{any_c}; \text{end_c}); \text{end_c}$. Hence, p believes that the contract of y_1 is $!\text{any_c}; \text{end_c}$ and that it is allowed to send 0 on y_1 .

Working out the reductions of the client in parallel with the two services we obtain the following assignments for the variables occurring in modules p , q , and r :

$$\begin{aligned} p \quad & \begin{cases} x_1 = [c^+]^{C, q, p} \\ y_1 = [[[d^+]^{D, r, q}]^{E, q, p}]^{E, q, p} \\ x_2 = [c^+]^{\text{end_c}, q, p} \\ y_2 = [[[d^+]^{\text{end_c}, r, q}]^{\text{end_c}, q, p}]^{\text{end_c}, q, p} \end{cases} & q \quad & \begin{cases} x_1 = [c^-]^{\bar{C}, p, q} \\ y_1 = [d^+]^{D, r, q} \\ x_2 = [c^-]^{\text{end_c}, p, q} \end{cases} \\ r \quad & \begin{cases} y_1 = [d^-]^{\bar{D}, q, r} \\ z \approx [[[[0]]^{\text{any_c}, p, q}]^{\text{any_c}, p, q}]^{F, q, r}]^{F, q, r} \\ y_2 \approx [d^-]^{\text{end_c}, q, r} \end{cases} \end{aligned}$$

An assignment $x = v$ denotes the fact that the program reduces to a state in which the substitution $\{v/x\}$ is actually performed in the respective thread. The assignment $z \approx e$ in thread r does not materialize into an actual substitution because 0 does not satisfy F and the evaluation of e eventually blames q . A similar observation holds for y_2 .

According to Definition 6, in order to identify the locally correct modules we have to consider all the output operations performed by these threads. We deduce that r is trivially locally correct, since r does not perform any output operation. Concerning p , we see that it eventually performs one output operation `send 0 [[$[d^+]$ D,r,q]E,q,p]` where $E = !\text{any_c}; \text{end_c}$. Since $0 \in \text{any_c}$, we deduce that p is also locally correct. Notice that, in order to decide whether p honors the contract of d^+ , we only look at the topmost monitor wrapping d^+ which, as we will see, reflects the local knowledge of p concerning this endpoint. Concerning q , we observe that it eventually performs the delegation `send [$[d^+]$ D,r,q [$[c^-]$ C,p,q]` where $\bar{C} = !(\text{any_c}; \text{end_c}); \text{end_c}$. In particular, the message being sent is an endpoint with contract $D = !\text{flat_c} (\neq 0); \text{end_c}$ whereas q is supposed to send an endpoint with contract $!\text{any_c}; \text{end_c}$. Observe that $!\text{flat_c} (\neq 0); \text{end_c} \not\leq !\text{any_c}; \text{end_c}$ because $\text{any_c} \not\leq \text{flat_c} (\neq 0)$ and entailment is contravariant with respect to outputs. That is, q violates item 3 of Definition 6 and therefore it is not locally correct. ■

Given that the conditions in Definition 6 are stated by actually running the process P in which label p occurs, one could wonder whether their enforcement demands unreasonable skills (such as divination or omniscience) to the programmer writing the code of p . The “local” qualification in Definition 6 is meant to indicate that this is not the case, namely that the definition provides effective guidance for writing correct code. We make two observations in this respect. First, all the conditions of Definition 6 only concern actions that are actually *performed* by module p , not by other modules. We concede that p may receive code from other modules (functions can be sent as messages) and, if this happens, Definition 6 assumes such code to behave well too. This assumption could be relaxed using a finer mechanism for associating blame labels with code Findler and Felleisen [2002], but we leave this refinement for future work. The second observation is that the contracts `flat_c w`, `C`, and `D` mentioned in conditions 1–5 are always found in the *topmost* monitors wrapping messages and endpoints. Inspection of the rules in Table 3 reveals that these contracts are always (continuations of) those advertised by a service previously invoked by p . We can formalize this claim by showing that the negative label of every topmost monitor occurring in module p is p itself. Given that labels do not migrate between monitors, this means that the contract of the topmost monitor wrapping an endpoint in module p faithfully reflects the knowledge of p about that endpoint.

Proposition 1. *If P is a user process such that $P \rightarrow^* \mathcal{P}_p[\mathbf{c} \ v_1 \cdots v_n \ _\]^{\text{---}q}$ where $\mathbf{c} \in \{\text{send, left, right}\}$, then $q = p$.*

In conclusion, the programmers of module p have all the information to consciously write locally correct code if they know the contracts of the services invoked by p . This is the weakest requirement we could reasonably ask.

Example 3. Consider again the processes of Example 2. The programmer of module p statically knows that the module is establishing a session with the service a whose contract is $C = ?(\text{any_c}; \text{end_c}); \text{end_c}$. According to this contract, the programmer assumes that the endpoint y_1 received from a can be used for sending *any* message, as specified by the contract `any_c`, even though we know, by looking at the assignments determined in Example 2, that the innermost monitor wrapping y_1 specifies a stricter constraint. From the same assignments we also see that `C` will indeed be the contract found in the topmost monitor wrapping y_1 . This means that the assumptions made by the programmer concerning y_1 correspond to the conditions that determine the local correctness of the module p .

The programmer of module q statically knows that the module is establishing a session y_1 with the service b whose contract is $D = !(\text{flat_c} (\neq 0)); \text{end_c}$. The same programmer is also aware that module q is itself a service advertizing the contract `C`, which requires the service to send on x_1 an endpoint with contract $E = !\text{any_c}; \text{end_c}$. The mistake made by the programmer

is detectable by comparing D (the contract of the endpoint y_1 that the service actually sends on x_1) and E (the contract of the endpoint that the service is supposed to send on x_1) and noticing that $D \not\leq E$. \blacksquare

5.3 Local Correctness Implies Blame Soundness

A locally correct module p cannot be blamed. The key insight for proving this result is that *local* correctness of p in a process P grants a number of *global* properties on the (busy) monitors referring to p in P . Crucially, these properties hold *anywhere* in P (not just within module p) and are *invariant* under arbitrary reductions of P .

Henceforth, we write $P \supset e$ (or equivalently $e \subset P$) if the sub-expression e occurs in P .

Definition 7 (module soundness). We say that module p is *sound* in P if:

1. $P \supset [v]^{\text{flat}_c w, p, -}$ implies $v \in w$;
2. $P \supset v \triangleleft^p e$ implies $e \Downarrow \mathbf{true}$;
3. $P \supset [v \triangleleft e]^{\text{flat}_c w, p, -}$ and $e \Downarrow \mathbf{true}$ imply $v \in w$;
4. $P \supset [[_]^{c, -, q}]^{d, r, -}$ and $p \in \{q, r\}$ imply $c \leq d$.

Conditions 1–3 state that all values v wrapped by a (busy) monitor where p is deemed responsible satisfy the contract/condition in the monitor. Condition 4 concerns stacked monitors where p is either the positive or negative label depending on whether p occurs in the outermost or innermost monitor. These configurations roughly correspond to *casts* Wadler and Findler [2009], whereby a value that is supposed to satisfy some contract C is used in a place where a value that satisfies D is expected. Regardless of whether p is the provider ($p = r$) or consumer ($p = q$) of such value, the entailment $C \leq D$ protects p in the sense that, if p risks of being blamed, another module will be blamed beforehand.

The key lemma states that soundness of p is preserved by reductions of P , provided that p is locally correct in P .

Lemma 1 (soundness preservation). *If $\Gamma \vdash P$ where $\Gamma :: \omega$ and p is locally correct and sound in P and $P \rightarrow Q$, then p is sound also in Q .*

Blame soundness is an easy corollary of Lemma 1. Note that p is trivially sound in every user process, since the clauses of Definition 7 solely concern the runtime syntax.

Theorem 2 (blame soundness). *If $\Gamma \vdash P$ where P is a user process and p is locally correct in P , then $P \rightarrow^* Q$ implies $\mathbf{blame} p \not\subset Q$.*

It is worth comparing Theorem 2 with similar results in the literature. The concepts most closely related to blame soundness are *blame correctness* Dimoulas et al. [2011, 2012] and *blame safety* Wadler and Findler [2009], Wadler [2015].

Blame correctness is a general property of a contract system guaranteeing that, whenever a module p may be blamed for a contract violation, it is because a value owned by (*i.e.* generating from) p is being checked against a flat contract which p was supposed to honor. In the literature, this notion has arisen following the observation that alternative monitoring strategies may yield different blame assignments Blume and McAllester [2006]. In contrast, blame soundness is a property of a particular module p and is formulated as the logical transposition of blame correctness: it states that, if p is never responsible for a contract violation, then p will never be blamed. This difference between blame correctness and blame soundness is reflected also in the techniques used for proving the two results. Blame correctness is proved by tracking the ownership of values and the obligations of modules with respect to contracts. In particular, there is no concept akin to that of “correct” module. The proof of blame soundness rests on a characterization of those modules that do not violate contracts and is closer in style to a type safety result, where Lemma 1 plays the role of subject reduction.

The blame calculus Wadler and Findler [2009], Wadler [2015] is a model of programs comprising both more-typed and less-typed modules. The interaction between modules adopting different typing disciplines is governed by type casts that are checked at runtime and may trigger blames. Blame safety is the property guaranteeing that, in case a blame is triggered, it always concerns a less-typed module. A locally correct module can be seen as a software component providing stronger guarantees about its behavior than those granted by the type system alone. In this sense, blame soundness resembles blame safety in that it guarantees that blames always concern modules that provide weaker guarantees about their correctness. Unlike well typing, local correctness is formulated operationally and not through a set of typing rules.

6 OCaml Implementation

In Sections 3–5 we have presented the model of λCoS and studied its metatheory. With this formal background, we now present a practical implementation of λCoS communication primitives and session monitoring in `OCaml`. There are a few differences between the model and the implementation. Some concern the operational semantics and are suggested by practical considerations (Section 6.1). Others concern `OCaml`'s type system, whose support for parametric polymorphism and recursive types allows us to devise a more expressive API to our library of monitored sessions (Section 6.2). We conclude the section with a final example that illustrates all these features at work (Section 6.3).

6.1 Alternative Monitoring Semantics

According to the reduction rule [R15], when a session is initiated with a service $a \Leftarrow_p^C v$, the client endpoint is monitored by the contract C advertised by the service and the service endpoint is monitored by the dual contract computed by `dual` C . We dub this semantics, where each peer of a session has its own monitor, *double-sided monitoring* in contrast to *single-sided monitoring* where only one peer is wrapped. Having monitors on both peers assures the presence of at least one chaperone contract for each endpoint in the system. To see the reason why this is technically convenient, consider item 3 of Definition 6, which concerns delegations: the guarantee that there is a monitor wrapping both the message and the endpoint on which the message is sent is a direct consequence of double-sided monitoring (*cf.* rule [R15]) and allows us to easily refer to the contracts C and D of the involved endpoints. Without this guarantee we would have four versions of item 3, depending on whether each endpoint is monitored or not. Even worse, we would have to *compute* the dual contract of each lone endpoint from that wrapping its peer, which could be located anywhere in the system. With the guarantee that each endpoint is monitored, Definition 6 remains reasonably compact and need not refer to parts of the system other than module p .

In practice, double-sided monitoring induces a useless overhead because each message exchanged through the peer endpoints of a session is checked against the same contract twice, once for each peer. This is a consequence of the fact that `dual` only inverts the direction of the exchanged messages without affecting their contracts (rules [R8–R14]). The overhead caused by double-sided monitoring is clearly illustrated by Example 2, where several values end up being wrapped by adjacent monitors having the same contract and the same labels. We can eliminate this overhead by adopting single-sided monitoring, where session initialization attaches just one monitor to the client endpoint. From an operational standpoint, this amounts to replacing rule [R15] with

$$[\text{R21}] \quad \langle \mathcal{E}[\text{connect } a] \rangle_p \parallel a \Leftarrow_q^C v \mapsto (\nu b) (\langle \mathcal{E}[[b^+]^{C,q,p}] \rangle_p \parallel \langle v b^- \rangle_q) \parallel a \Leftarrow_q^C v$$

that leaves b^- unmonitored. While it is obvious that by removing one monitor we reduce the points in the program that can generate blame, it is not entirely obvious that no new blame is introduced. The next result guarantees that this is indeed the case:

Proposition 2. *Let \mapsto be the relation defined by the rules in Tables 2 and 3 except that [R15] is replaced by [R21]. If P is a user process and $P \mapsto^* \supset \text{blame } p$, then $P \rightarrow^* \supset \text{blame } p$.*

In fact we conjecture that the converse of Proposition 2 also holds, namely that single-sided and double-sided monitoring yield the same blames. Proving this fact turns out not to be trivial as the exact relationship between processes resulting from the two monitoring semantics is hard to formalize. As it stands, Proposition 2 is enough to conclude that single-sided monitoring does not affect the blame soundness result (Theorem 2).

We now turn the attention to the rules [R17–R20], which are defined in such a way that a single reduction step rearranges a whole stack of monitors. This formulation of the operational semantics – which we dub *big-step monitoring* – is convenient to prove various auxiliary results, including Proposition 1 and Lemma 1, which are key ingredients of blame soundness (Theorem 2). The issue with big-step monitoring is that the contracts concerning a communicated message migrate along with the message from the sender to the receiver. Given that contracts are ordinary expressions, this migration is, in effect, a form of code mobility which may pose problems, especially in a distributed system. In addition, the message is not checked against the contracts on the sender side until the communication has completed. This means that contract violations might go undetected if, for instance, a deadlock or livelock prevents the communication from occurring.

For these reasons, it makes sense to consider a *small-step monitoring* semantics that deals with monitors one at a time and that is closer in style to the semantics given by Findler and Felleisen [2002]. We illustrate small-step monitoring for a communication across endpoints monitored by non-dependent contracts. Communication goes through two distinct phases. During the first phase, the monitors on the endpoints are disassembled and distributed on the message and around the whole applications involving **send** and **receive**. Formally:

$$\begin{aligned} \text{[R22]} \quad & \mathbf{send} \ v \ [\varepsilon]^{!C;D,\sigma} \rightsquigarrow [\mathbf{send} \ [v]^{C,-\sigma} \ \varepsilon]^{D,\sigma} \\ \text{[R23]} \quad & \mathbf{receive} \ [\varepsilon]^{?C;D,\sigma} \rightsquigarrow \mathbf{let} \ x, y = \mathbf{receive} \ \varepsilon \ \mathbf{in} \ ([x]^{C,\sigma}, [y]^{D,\sigma}) \end{aligned}$$

Once all the contracts have been stripped off and suitably rearranged/checked, the actual communication may take place:

$$\text{[R24]} \quad \langle \mathcal{E}[\mathbf{send} \ v \ a^+] \rangle_p \parallel \langle \mathcal{E}'[\mathbf{receive} \ a^{\bar{}}] \rangle_q \rightsquigarrow \langle \mathcal{E}[a^+] \rangle_p \parallel \langle \mathcal{E}'[(v, a^{\bar{}})] \rangle_q$$

With small-step monitoring, contract migration only happens during the communication of endpoints (delegations). Arguably, the contexts in which this form of communication is allowed also allow for the migration of contracts. In addition, rule [R22] makes sure that messages different from endpoints are checked against contracts *before* the **send** can reduce further because $[v]^{C,-\sigma}$ is not a value when C is a flat contract. This means that these checks are performed even if the communication does not occur and their cost is charged to the sender. We illustrate these differences between big-step and small-step monitoring in the following example.

Example 4. Let P be the parallel composition

$$\left\langle \begin{array}{l} \mathbf{let} \ x = \mathbf{send} \ 0 \ [a^+]^{!flat_c \ (\neq 0); end_c, q, p} \ \mathbf{in} \\ \mathbf{let} \ y = \mathbf{send} \ 1 \ [b^+]^{!any_c; end_c, q, p} \ \mathbf{in} \\ \mathbf{close} \ x; \mathbf{close} \ y \end{array} \right\rangle_p \parallel \left\langle \begin{array}{l} \mathbf{let} \ _, y = \mathbf{receive} \ [b^-]^{?any_c; end_c, p, q} \ \mathbf{in} \\ \mathbf{let} \ _, x = \mathbf{receive} \ [a^-]^{?flat_c \ (\neq 0); end_c, p, q} \ \mathbf{in} \\ \mathbf{close} \ x; \mathbf{close} \ y \end{array} \right\rangle_q$$

where p is sending two messages (0 and 1) and q is waiting for them. Notice that p is sending a message on a^+ first and then on b^+ , whereas q is waiting for a message from b^- first and then a^- . Because of the different order of communications (and of the synchronous communication model used in λCoS), the process P is stuck according to the big-step monitoring semantics. That is, $P \dashv$. Notice also that p is violating the contract of a^+ , which requires the message to be a number different from 0. Because of the deadlock, this violation is not detected by the big-step monitoring semantics. On the contrary, under small-step monitoring the p thread reduces thus

$$\begin{aligned} & \langle \mathbf{let} \ x = \mathbf{send} \ 0 \ [a^+]^{!flat_c \ (\neq 0); end_c, q, p} \ \mathbf{in} \ \dots \rangle_p \\ & \rightsquigarrow \langle \mathbf{let} \ x = [\mathbf{send} \ [0]^{flat_c \ (\neq 0), p, q} \ a^+]^{end_c, q, p} \ \mathbf{in} \ \dots \rangle_p && \text{by [R22]} \\ & \rightarrow \langle \mathbf{let} \ x = [\mathbf{send} \ (0 \triangleleft^p \ (\neq) \ 0 \ 0) \ a^+]^{end_c, q, p} \ \mathbf{in} \ \dots \rangle_p && \text{by [R5]} \\ & \rightarrow \langle \mathbf{let} \ x = [\mathbf{send} \ (0 \triangleleft^p \ \mathbf{false}) \ a^+]^{end_c, q, p} \ \mathbf{in} \ \dots \rangle_p && \text{semantics of } \neq \\ & \rightarrow \langle \mathbf{let} \ x = [\mathbf{send} \ (\mathbf{blame} \ p) \ a^+]^{end_c, q, p} \ \mathbf{in} \ \dots \rangle_p && \text{by [R7]} \end{aligned}$$

yielding a blame for p even if the communication does not occur. \blacksquare

It is possible to prove a (limited) form of equivalence between big-step and small-step monitoring by considering those configurations where a communication is about to occur:

Proposition 3. *Let \rightsquigarrow be the relation defined by the rules in Tables 2 and 3 except that rule [R17] is replaced by [R22–R24]. If*

$$\begin{aligned} P &\stackrel{\text{def}}{=} \langle \mathcal{E}[\mathbf{send} \ v \ [a^t]^{1C;D,\sigma}] \rangle_p \parallel \langle \mathcal{E}'[\mathbf{receive} \ [a^r]^{?E;F,\varrho}] \rangle_q \\ Q_1 &\stackrel{\text{def}}{=} \langle \mathcal{E}[[a^t]^{D,\sigma}] \rangle_p \parallel \langle \mathcal{E}'[(v, [a^r]^{F,\varrho})] \rangle_q \\ Q_2 &\stackrel{\text{def}}{=} \langle \mathcal{E}[[a^t]^{D,\sigma}] \rangle_p \parallel \langle \mathcal{E}'[(\llbracket v \rrbracket^{\bar{C},-\sigma})_{E,\varrho}, [a^r]^{F,\varrho}] \rangle_q \end{aligned}$$

then either

1. $P \rightsquigarrow^* Q$ and $P \rightarrow^* Q$ for some $Q \in \{Q_1, Q_2\}$, or
2. $P \rightsquigarrow^* \triangleright \mathbf{blame} \ r$ and $Q \rightarrow^* \triangleright \mathbf{blame} \ r$ for some r .

In words, regardless of the monitoring semantics and starting from a configuration involving a sender and a receiver, either the reduction converges to the same state Q_i in which the message has been successfully delivered to the destination, or the same module r is blamed by both semantics. There are two possible resulting states Q_1 and Q_2 depending on whether v is a ground value or a (monitored) endpoint. In the first case, the monitors around v turn into busy monitors and eventually disappear. In the second case, they accumulate around v . Analogous results can be proven for all the other configurations involving communication primitives.

We also note that the small-step semantics provides for greater flexibility in that it allows the reduction of configurations where stacked monitors have a mixture of dependent and non-dependent contracts because each monitor is handled individually.

Despite these differences between big-step and small-step monitoring, the given criterion of local correctness and the blame soundness result are also relevant for our implementation. After all, what small-step monitoring does is to introduce a number of intermediate steps in which the relationships between adjacent monitors are (temporarily) lost and where the actual knowledge of the programmer with respect to the contract of the used endpoints is not as easy to characterize as in Proposition 1. In case communications are guaranteed to occur, Proposition 3 shows that both monitoring semantics either yield the same blame or eventually converge to the same reduct. A difference remains in that the thread from which the blame originates may not be the same: in the big-step monitoring semantics it is always the receiver of a message that possibly detects a violation (*cf.* thread r in Example 2), whereas in the small-step monitoring semantics checks are in general performed on both sides of the session and the sender may yield blames too (*cf.* thread p in Example 4).

6.2 Monitoring Sessions in OCaml

In this section we illustrate the key aspects of an OCaml module that implements λCoS communication primitives. Instead of building the primitives from scratch, we obtain them as wrappers of the corresponding primitives provided by `FuSe`, an OCaml library of binary sessions Padovani [2017]. This way we do not have to delve into low-level details concerning the encoding of session types or the implementation of the session primitives and instead we can focus on the aspects strictly related to contract monitoring. The fact that we can build λCoS primitives on top of `FuSe` is a sign that our monitoring technique is modular and should be portable to other session libraries for possibly different programming languages.

Even though `FuSe` has its own OCaml representation of session types, we keep using the metavariables T and S to improve readability. In particular, we write $T \ \mathbf{st}$ for the OCaml type that denotes a lone (*i.e.*, unmonitored) `FuSe` endpoint with session type T . Similarly, we write $T \ \mathbf{mt}$ for the OCaml type that denotes a possibly monitored endpoint. OCaml's support for parametric polymorphism in conjunction to `FuSe`'s representation of session types makes it possible to also

represent session type variables standing for unknown session types. Hereafter we use A, B to range over session type variables and $A \text{ st}$, $A \text{ mt}$ for their corresponding representations in `FuSe` and our λCoS implementation. The `FuSe` representation of session types makes it easy to switch from a session type to its dual Padovani [2017]. We will write $\overline{A \text{ st}}$ to refer to the dual of $A \text{ st}$.

Concerning the session communication primitives, we use the prefix `FuSe` to disambiguate them from those we are going to implement. As an example,

```
FuSe.send :  $\alpha \rightarrow !\alpha.A \text{ st} \rightarrow A \text{ st}$ 
```

is the signature of the `send` primitive provided by `FuSe`, which we will use for providing a suitable wrapper

```
send :  $\alpha \rightarrow !\alpha.A \text{ mt} \rightarrow A \text{ mt}$ 
```

Following Hinze et al. [2006], we represent the contract forms introduced in Section 3 using the constructors of a generalized algebraic data type (GADT):

```

1 type [_] =
2   | Flat      : ( $\alpha \rightarrow \text{bool}$ )  $\rightarrow [\alpha]$ 
3   | End       : [end]
4   | Receive  : [ $\alpha$ ] * ( $\alpha \rightarrow [A \text{ mt}]$ )  $\rightarrow [?\alpha.A \text{ mt}]$ 
5   | Send     : [ $\alpha$ ] * ( $\alpha \rightarrow [A \text{ mt}]$ )  $\rightarrow [!\alpha.A \text{ mt}]$ 
6   | Branch   : [bool] * [ $A \text{ mt}$ ] * [ $B \text{ mt}$ ]  $\rightarrow [A \& B \text{ mt}]$ 
7   | Choice   : [bool] * [ $A \text{ mt}$ ] * [ $B \text{ mt}$ ]  $\rightarrow [A \circledast B \text{ mt}]$ 

```

According to the `OCaml` syntax for GADTs, the type parameter of the data type is left unspecified (see `_` on line 1) and each constructor is explicitly annotated with its type (lines 2–7). These contract constructors are in close correspondence with those of Table 1 and are kept private in our `OCaml` module for λCoS which instead exports one (curried) function for each of them. For example, `flat_c` is defined thus

```
let flat_c w = Flat w
```

In `OCaml` the function `flat_c` is polymorphic and has type $(\alpha \rightarrow \text{bool}) \rightarrow [\alpha]$, hence it can be used to define flat contracts on values of any type. The typical example is `any_c : [\alpha]` used in several occasions and defined below:

```
let any_c = Flat (fun _  $\rightarrow$  true)
```

Unlike λCoS , `OCaml`'s type system is not substructural and there is no distinction between linear and unlimited types. Consequently, we cannot statically ensure that flat contracts are used only for unlimited values. As we will see shortly, violations on the usage of flat contracts will be detected dynamically using a runtime check.

The remaining constructors are in correspondence with the contract forms for session endpoints. In particular, `End` describes an endpoint that can only be closed. We have:

```
let end_c = End
```

`Send` is the constructor for a dependent output contract, which is built up from a contract $[\alpha]$ for the communicated value and a function $(\alpha \rightarrow [A \text{ mt}])$ that generates the continuation contract. The corresponding combinator `send_d` is defined thus:

```
let send_d k1 k2 = Send (k1, k2)
```

The lack of distinction between linear and unlimited types in `OCaml` allows us to define non-dependent contract combinators as particular instances of the corresponding dependent version. Thus, we have

```
let send_c k1 k2 = send_d k1 (fun _  $\rightarrow$  k2)
```

The primitives `receive_c`, `receive_d`, `choice_c`, and `branch_c` are defined similarly. As pointed out by a reviewer, it is also sensible to provide specializations of `choice_c` and `branch_c` that mandate the selection of a particular branch. For instance

```

let left_choice_c k = choice_c (flat_c id) k any_c
let right_choice_c k = choice_c (flat_c not) any_c k

```

can be used to specify that the user of an endpoint is required to select the left (respectively, right) branch of a choice. Indeed, the `(flat_c id)` contract is satisfied by the `true` label whereas the `(flat_c not)` contract is satisfied by the `false` label.

It is now time to provide the concrete representation of a λCoS endpoints as possibly monitored FuSe endpoints. In general we have to take into account the possibility that an endpoint is wrapped by an arbitrary number of monitors, each carrying a chaperone contract. For this reason, we use a recursive data type:

```

type A mt =
  | Channel of linearity_tag_type * A st
  | Monitor of [A mt] * string * string * A mt

```

Note that this definition and the GADT for contracts are mutually dependent and must be introduced jointly in OCaml. Here we have kept them separate for the sake of a more gradual presentation. The `Channel` constructor represents a lone endpoint (*cf.* Table 1) and consists of a FuSe endpoint and a linearity tag, which will be used for the runtime detection of linearity violations on flat contracts. The definition of `linearity_tag_type` is uninteresting. We only remark that it is a singleton type whose unique inhabitant is kept private to the module and can only occur in a `Channel` constructor. The `Monitor` constructor represents a monitored endpoint, consisting of a contract (of type `[A mt]`), two blame labels (of type `string`), and an underlying, possibly monitored endpoint of type `A mt`.

We now illustrate in detail the implementation of `send`, the other communication primitives are analogous. Recall that `send` has two arguments, a message `v` and an endpoint `ep` on which `v` must be transmitted, and returns the same endpoint (*cf.* Table 3). Keeping this in mind, we have:

```

1 let rec send v =
2   function
3     | Channel (lin, ep) → Channel (lin, FuSe.send v ep)
4     | Monitor (Send (k, w), pos, neg, ep) →
5       wrap (w v) pos neg (send (wrap k neg pos v) ep)
6     | Monitor (Flat _, _, _, _) → assert false (* IMPOSSIBLE *)

```

The case when `ep` is a lone endpoint (line 3) corresponds to the reduction rule [R24] and consists in invoking the corresponding FuSe primitive and turning the resulting FuSe endpoint into a λCoS endpoint through the `Channel` constructor. The case when `ep` is a monitored endpoint (lines 4–5) corresponds to the reduction rule [R22]. In this case, the contract is disassembled into a contract `k` for the message being transmitted and a function `w` that is applied to the message to compute the continuation contract. The auxiliary function `wrap` (described shortly) takes care of checking that contracts are satisfied or attaches chaperone contracts to session endpoints. Note the swapping of blame labels in the inner application of `wrap` and the recursive application of `send`, which takes care of others chaperone contracts in `ep` and eventually performs the communication. The last case in the definition of `send` (line 6) never applies and is only meant to prevent complaints about a non-exhaustive pattern matching. The point is that when `send` is applied to a monitored endpoint `Monitor (k, _, _, ep)` the type of `send` ensures that `k` has type `[!α.A mt]` and the only constructors whose type is unifiable with this one are `Send` and `Flat`, which applies to any value but is never attached as a chaperone contract by `wrap`.

The last key ingredient of the implementation is indeed `wrap`, whose purpose is to check whether a value satisfies a given contract and to blame the guilty party if this is not the case. The structure of `wrap`, shown below, is analogous to that of the homonymous function defined by Findler and Felleisen [2002], the main difference being that, in our case, `wrap` deals with contracts concerning session endpoints.

```

1 let wrap : type a. a ct → string → string → a → a = fun k pos neg v →

```

```

2   match k with
3   | Flat w          → if unlimited v && w v then v else raise (Blame pos)
4   | End as k       → Monitor (k, pos, neg, v)
5   | Receive _ as k → Monitor (k, pos, neg, v)
6   | Send _ as k    → Monitor (k, pos, neg, v)
7   | Branch _ as k → Monitor (k, pos, neg, v)
8   | Choice _ as k  → Monitor (k, pos, neg, v)

```

An application `wrap k pos neg v` inspects the structure of the contract `k` (line 2). When `k` is a flat contract the message is verified to be unlimited and to satisfy the predicate `w` (line 3). If both conditions are satisfied, the value `v` is returned. On the contrary, the sender is blamed for a contract violation. The `unlimited` auxiliary function is implemented by checking that (the runtime representation of) `v` does not contain any occurrence of `linear_tag`. Overall, this case in the definition of `wrap` corresponds to the rules [R5–R7] in Table 2. In the remaining cases (lines 4–8), `wrap` cannot check immediately whether the endpoint satisfies the contract `k`. Therefore, the endpoint is wrapped by a monitor with the contract `k` and the blame labels `pos` and `neg`. Two remarks about the typing of `wrap` are in order. First, it is necessary to provide a type annotation to `wrap` (line 1) to inform OCaml that the type parameter α is *locally abstract* (the explicit quantification `type α` means this). This way, OCaml can refine this type parameter depending on pattern matching Garrigue and Normand [2011]. Second, the lines 5–9 cannot be collapsed into a single case, even if they all bind just one variable `k` and the right hand side of each pattern matching rule is behaviorally the same, for the type of `k` is different in each case.

6.3 Extended Example: List Forwarding

In this section we present a final example that illustrates all the features of the OCaml implementation of λCoS , including dependent contracts. We also take advantage of OCaml’s support for parametric polymorphism and recursive types, which we have omitted in the formal model of λCoS for the sake of simplicity.

The following function

```

1   let forwarder_body x =
2     let rec aux y =
3       function
4         | []      → close (right y)
5         | v :: l → aux (send v (left y)) l in
6     let l, x = receive x in (* receive the elements to be forwarded *)
7     let y, x = receive x in (* receive the destination endpoint *)
8     close x; aux y l

```

models the body of a forwarding service that delivers a list of elements to a given recipient, one element at a time. The service interacts with the client using the endpoint `x`, from which it receives a list `l` of elements (line 6) and another endpoint `y` on which the elements of the list should be forwarded (line 7). The main loop of the service (lines 2–5) iterates over the list: when the list is empty (line 4), the service selects the “right” branch of the protocol and closes the endpoint; when the list contains at least one element `v`, the service selects the “left” branch of the protocol, it sends `v` on `y`, and then iterates over the tail of the list (line 5).

OCaml infers for `forwarder_body` the signature

```
val forwarder_body : ?( $\alpha$  list).?(rec A.(! $\alpha$ .A)  $\circledast$  end).end mt → unit
```

where `rec A.(! α .A) \circledast end` stands for the equi-recursive session type T that satisfies the equation $T = (!\alpha.T) \circledast \text{end}$.¹ As the type suggests, `forwarder_body` performs an arbitrary number of outputs on the delegated endpoint `y`, and by parametricity we also deduce that the elements

¹The type inferred by OCaml is in fact an encoded form of the type shown here. FuSe comes along with a utility that pretty prints encoded OCaml types into the more readable syntax that we have used.

being sent in these outputs must come from the list `l`. However, we do not know whether the number of forwarded elements actually matches the length of the list. The service can advertise this guarantee by means of the following contract:

```

1 let forwarder_c =
2   let rec fwd n =
3     if n > 0 then (* more list elements *)
4       left_choice_c (send_c any_c (fwd (n - 1)))
5     else (* no more list elements *)
6       right_choice_c end_c
7   in send_d any_c (fun l → send_c (fwd (List.length l)) end_c)

```

The contract starts with the specification of the output of a list `l` for which no particular constraints are given (*cf.* the `any_c` on line 7). However, the continuation contract depends on the length of `l`. More precisely, the continuation describes an endpoint that must be used for sending another endpoint with contract `fwd (List.length l)` before being closed. The expression `fwd n` (lines 2–6) yields a contract for an endpoint that must be used for sending exactly `n` messages. This is achieved using the `left_choice_c` and `right_choice_c` contracts. If `n > 0` the forwarder must select the left branch and then send one message followed by `n - 1` more (line 4). If `n` is 0 the forwarder must select the right branch and then terminate the interaction (line 6).

We conclude with two remarks. First, the contract `forwarder_c` allows `forwarder_body` to change the forwarded messages as long as their number is the same as the length of `l`. Parametricity only guarantees that the forwarded elements come from `l`. A stronger version of the contract could be written by making `fwd` a function that operates over the list of remaining messages and adding a condition over sent messages that checks whether the forwarded message is on the list of messages to be delivered. Second, the `forwarder_c` contract ties the behavior of a process on one session (`y`) to data (`l`) exchanged over a different session (`x`). Therefore, besides providing a means for specifying precise constraints over messages exchanged in one session, contracts for higher-order sessions enable the specification of inter-session dependencies.

7 Related Work

Blame Correctness. Blame assignment with higher-order and dependent contracts is subtle because contract verification is deferred until the monitored object is actually used. In fact, alternative monitoring strategies called *lax*, *picky* and *indy* have been proposed Blume and McAllester [2006], Dimoulas et al. [2011], Greenberg et al. [2012]. Such design choices pose the question of whether a particular strategy is reasonable, *i.e.*, if it assigns blames correctly. This problem has been addressed either by showing that principals that can be blamed have responsibility in a contract violation Dimoulas et al. [2011, 2012] or by providing a characterization of components that *satisfy* a given contract Blume and McAllester [2006], Findler and Blume [2006], Dimoulas and Felleisen [2011]. Given a semantics for contracts, contract satisfaction addresses the problem of showing that a contract system is sound and complete with respect to that semantics, *i.e.* that each blame assigned by the system corresponds to a contract violation and each contract violation is detected by the system. Our blame soundness result partially addresses this problem by showing that the contract system does not report blames when a component does not violate its contracts. The exact relationship between blame soundness and contract satisfaction remains to be fully understood.

Contracts have been integrated into type systems as refinement types, which are checked dynamically via cast insertion, like the hybrid types of Knowles and Flanagan [2010]. The relation between contracts and hybrid types has been studied by Greenberg et al. [2012] and Gronski and Flanagan [2007]. Wadler and Findler [2009] provide a unifying view with the blame calculus, but leave dependent contracts for future work. The blame calculus allows for a simple characterization of correct component when typed and untyped code are mixed. The blame theorem shows that blames are always assigned to the lesser-typed components.

Although contracts in λCoS do not interact with the type system, our notions of local correctness and blame soundness have been inspired by Wadler and Findler [2009] and Wadler [2015]. Communication over a monitored endpoint is the only point in the execution of a program that can originate a blame. We interpret communication as an implicit cast in which the message is coerced to the type specified in the contract wrapping the endpoint. However, contract satisfaction is verified dynamically in our approach. We leave the problem of static verification of session contracts, e.g., along the lines of Nguyen et al. [2014], as future work.

Behavioral/Temporal Higher-order Contracts. The interplay between higher-order contracts and behavioral/temporal aspects of modules, such as restricting the order in which functions can be invoked, has been previously addressed by Disney et al. [2011] and Scholliers et al. [2015]. In both approaches, the enforcement of temporal constraints is done dynamically and the contract language allows for the specification of both allowed and disallowed traces. On the contrary, in λCoS , the usage of endpoints is regulated by a combination of static and dynamic constraints: static constraints are enforced by session types, which guarantee that processes use session endpoints according to their protocol; dynamic constraints, which concern the content of exchanged messages and may affect the selection and availability of choices and branches (Section 6.3), are checked at runtime by monitors. None of the previous works on behavioral/temporal contracts Scholliers et al. [2015], Disney et al. [2011] provides a characterisation of module correctness or a formal statement about the correctness of blame assignment akin to our blame soundness result (Section 5).

Swords et al. [2015] proposed λ_{cc} , a language tailored to the implementation of alternative approaches to monitoring, and discussed a possible implementation of the temporal contracts of Disney et al. [2011] on top of λ_{cc} . An interesting question for future work is whether the primitives of λ_{cc} allows for a convenient implementation of λCoS to avoid monitor migration in delegations.

Contracts and Affine/Linear Types. Contract monitoring may duplicate values and discard contracts (Tables 2 and 3). These aspects have not been investigated elsewhere and affect the typing of contract constructors when the type system is substructural. If we allowed contracts to use linear resources, the monitoring of a non-linear object with a contract using linear resources would affect the behavior of the object, which would become linear itself. In this sense, our contracts conform to the definition of “chaperone contracts” given by Strickland et al. [2012], where chaperones are not allowed to affect the behavior of monitored objects.

Tov and Pucella [2010] use stateful contracts for controlling the usage of affine functions (functions that can be applied at most once) when these flow into a region of the program that uses a conventional (*i.e.*, non substructural) type system.

Projections and Duality. Contract duality is tightly related to session type duality Honda [1993], Honda et al. [1998] and should not be confused with the client and server projections of Findler and Blume [2006]. The point is that dual contracts describe the same obligations in a session with respect to complementary directions of the messages being exchanged, whereas projection separates the obligations for values flowing into and out of a monitored expression.

Contracts for Sessions. Bocchi et al. [2010] and Toninho and Yoshida [2016] extend global types with assertions to specify constraints on values communicated in a multiparty session. These approaches are based on a top-down methodology whereby the whole multiparty interaction is designed at once and then projected on the single participants of the session. We work with binary sessions only. Our contracts can be applied gradually to arbitrary subsets of interacting modules and can be used to describe whole protocols or only fragments thereof. The session type is inferred automatically from the structure of the program. Bocchi et al. [2010] use a decidable assertion logic and Toninho and Yoshida [2016] use dependent session types along the lines of Toninho et al. [2011]. These choices make it possible to verify the correctness of participants statically. The soundness result rests on the assumption that all processes participating in a session are

well typed, requiring no runtime monitoring. This assumption is relaxed in some works Chen et al. [2011], Bocchi et al. [2013], where participants may misbehave and a monitor is used to suppress messages that violate the protocol. Neither higher-order sessions nor blame assignment are considered in these works.

Monitoring of untrusted processes has been also used Bartoletti et al. [2012, 2013], Jia et al. [2016] to make sure that processes follow the intended protocol (session fidelity) and to assign blame if this is not the case. In these works monitoring does not concern the content of messages and blame freedom can be guaranteed by typing. In λCoS session fidelity is ensured by typing whereas contracts specify conditions on the content of exchanged messages.

Thiemann [2014] studies a gradual type system for sessions with explicit coercions for branches and choices. Again, the dynamic checks concern session fidelity and particularly the branching structure of protocols.

8 Concluding Remarks

Building on contracts for higher-order functions pioneered by Findler and Felleisen [2002] and later extended to mutable objects by Strickland et al. [2012] we have defined and implemented a monitoring technique for higher-order sessions. Its characterizing aspect is that the contract wrapping a session endpoint must be dynamically updated at runtime, as the session progresses.

The setting provided us with the opportunity to investigate the ramifications of contract monitoring in the presence of linear resources. In particular, we have seen that session endpoints should not occur in contracts, which is in line with the observation that contracts should be “pure” Meyer [1992]. Indeed, the operations involved with session endpoints are inherently impure.

We have proved a blame soundness result stating that modules which do not violate contracts are not blamed. The result rests on the key assumption that messages are either session endpoints or unlimited values (*i.e.* data and functions not containing session endpoints). This assumption allows us to tackle the higher-order case through delegation, which is the idiomatic form of mobility in sessions. Extending the blame soundness result to a setting where functions may contain session endpoints is an intriguing future development.

We intend to incorporate our proof-of-concept implementation of chaperone contracts for sessions into FuSe Padovani [2017]. In this respect, the GADT presented in Section 6 suffers from two important limitations. First, the GADT does not allow the specification of contracts for structured data types (such as pairs) containing session endpoints. This can be overcome either by adding dedicated constructors corresponding to these data types Hinze et al. [2006] or by adding a single, general-purpose constructor that carries a user-provided, type-specific wrapper function. The second current limitation is that recursively defined contracts (such as `forwarder_c` in Section 6.3) must be finite in general. The finiteness of `forwarder_c` was guaranteed by the fact that its structure is isomorphic to that of a (finite) list. However, infinite protocols are commonly found in practice. In these cases, a contract definition relying on OCaml’s support for recursion would diverge. This limitation can be lifted by adding a constructor to the GADT that lazily computes the fixpoint of a recursively defined contract. Once all these ingredients are in place, a proper evaluation of the overhead induced by our technique for contract monitoring is in order.

Acknowledgments. The authors are grateful to the anonymous reviewers for their detailed and thoughtful comments. The first author has been partially supported by World Wide Style project of the University of Torino (second edition) and Fondazione CRT.

References

- Massimo Bartoletti, Emilio Tuosto, and Roberto Zunino. Contract-oriented computing in CO2. *Scientific Annals of Computer Science*, 22(1):5–60, 2012.
- Massimo Bartoletti, Alceste Scalas, Emilio Tuosto, and Roberto Zunino. Honesty by typing. In *Proceedings of FMOODS/FORTE’13*, LNCS 7892, pages 305–320. Springer, 2013.

- Matthias Blume and David A. McAllester. Sound and complete models of contracts. *Journal of Functional Programming*, 16(4-5):375–414, 2006.
- Laura Bocchi, Kohei Honda, Emilio Tuosto, and Nobuko Yoshida. A theory of design-by-contract for distributed multiparty interactions. In *Proceedings of CONCUR'10*, LNCS 6269, pages 162–176. Springer, 2010.
- Laura Bocchi, Tzu-Chun Chen, Romain Demangeon, Kohei Honda, and Nobuko Yoshida. Monitoring networks through multiparty session types. In *Proceedings FMOODS/FORTE'13*, LNCS 7892, pages 50–65. Springer, 2013.
- Tzu-Chun Chen, Laura Bocchi, Pierre-Malo Deniérou, Kohei Honda, and Nobuko Yoshida. Asynchronous distributed monitoring for multiparty session enforcement. In *Proceedings of TGC'11*, LNCS 7173, pages 25–45. Springer, 2011.
- Christos Dimoulas and Matthias Felleisen. On contract satisfaction in a higher-order world. *ACM Transactions on Programming Languages and Systems*, 33(5):16, 2011.
- Christos Dimoulas, Robert Bruce Findler, Cormac Flanagan, and Matthias Felleisen. Correct blame for contracts: no more scapegoating. In *Proceedings of POPL'11*, pages 215–226. ACM, 2011.
- Christos Dimoulas, Sam Tobin-Hochstadt, and Matthias Felleisen. Complete monitors for behavioral contracts. In *Proceedings of ESOP'12*, LNCS 7211, pages 214–233. Springer, 2012.
- Tim Disney, Cormac Flanagan, and Jay McCarthy. Temporal higher-order contracts. In *ACM SIGPLAN Notices*, volume 46, pages 176–188. ACM, 2011.
- Robert Bruce Findler and Matthias Blume. Contracts as pairs of projections. In *Proceedings of FLOPS'06*, LNCS 3945, pages 226–241. Springer, 2006.
- Robert Bruce Findler and Matthias Felleisen. Contracts for higher-order functions. In *Proceedings of ICFP'02*, pages 48–59. ACM, 2002.
- Jacques Garrigue and Jacques Le Normand. Adding GADTs to OCaml: the direct approach. In *Proceedings of ACM SIGPLAN Workshop on ML*, 2011.
- Simon J. Gay and Malcolm Hole. Subtyping for Session Types in the π -calculus. *Acta Informatica*, 42(2-3):191–225, 2005.
- Simon J. Gay and Vasco Thudichum Vasconcelos. Linear type theory for asynchronous session types. *Journal of Functional Programming*, 20(1):19–50, 2010.
- Simon J. Gay, Vasco Thudichum Vasconcelos, António Ravara, Nils Gesbert, and Alexandre Z. Caldeira. Modular session types for distributed object-oriented programming. In *Proceedings of POPL'10*, pages 299–312. ACM, 2010.
- Michael Greenberg, Benjamin C. Pierce, and Stephanie Weirich. Contracts made manifest. *Journal of Functional Programming*, 22(3):225–274, 2012.
- Jessica Gronski and Cormac Flanagan. Unifying hybrid types and contracts. In *Proceedings of TFP'07*, volume 8, pages 54–70. Intellect, 2007.
- Ralf Hinze, Johan Jeuring, and Andres Löf. Typed contracts for functional programming. In *Proceedings of FLOPS'06*, LNCS 3945, pages 208–225. Springer, 2006.
- Kohei Honda. Types for dyadic interaction. In *Proceedings of CONCUR'93*, LNCS 715, pages 509–523. Springer, 1993.

- Kohei Honda, Vasco T. Vasconcelos, and Makoto Kubo. Language primitives and type disciplines for structured communication-based programming. In *Proceedings of ESOP'98*, LNCS 1381, pages 122–138. Springer, 1998.
- Limin Jia, Hannah Gommerstadt, and Frank Pfenning. Monitors and blame assignment for higher-order session types. In *Proceedings of POPL'16*, pages 582–594. ACM, 2016.
- Kenneth Knowles and Cormac Flanagan. Hybrid type checking. *ACM Transactions on Programming Languages and Systems*, 32(2), 2010.
- Naoki Kobayashi, Benjamin C. Pierce, and David N. Turner. Linearity and the pi-calculus. *ACM Transactions on Programming Languages and Systems*, 21(5):914–947, 1999.
- Bertrand Meyer. Design by contract. In *Advances in Object-oriented Software Engineering*, pages 1–50. Prentice-Hall, 1992.
- Phúc C Nguyen, Sam Tobin-Hochstadt, and David Van Horn. Soft contract verification. In *ACM SIGPLAN Notices*, volume 49, pages 139–152. ACM, 2014.
- Luca Padovani. A Simple Library Implementation of Binary Sessions. *Journal of Functional Programming*, 27, 2017. ISSN 0956-7968.
- António Ravara and Vasco Thudichum Vasconcelos. Typing non-uniform concurrent objects. In *Proceedings of CONCUR'00*, LNCS 1877, pages 474–488. Springer, 2000.
- Christophe Scholliers, Éric Tanter, and Wolfgang De Meuter. Computational contracts. *Science of Computer Programming*, 98:360–375, 2015.
- T. Stephen Strickland, Sam Tobin-Hochstadt, Robert Bruce Findler, and Matthew Flatt. Chaperones and impersonators: run-time support for reasonable interposition. In *Proceedings of OOPSLA'12*, pages 943–962. ACM, 2012.
- Cameron Swords, Amr Sabry, and Sam Tobin-Hochstadt. Expressing contract monitors as patterns of communication. In *ACM SIGPLAN Notices*, volume 50, pages 387–399. ACM, 2015.
- Peter Thiemann. Session types with gradual typing. In *Proceedings of TGC'14*, LNCS 8902, pages 144–158. Springer, 2014.
- Bernardo Toninho and Nobuko Yoshida. Certifying data in multiparty session types. In *A List of Successes That Can Change the World - Essays Dedicated to Philip Wadler on the Occasion of His 60th Birthday*, LNCS 9600, pages 433–458. Springer, 2016.
- Bernardo Toninho, Luís Caires, and Frank Pfenning. Dependent session types via intuitionistic linear type theory. In *Proceedings of PPDP'11*, pages 161–172. ACM, 2011.
- Jesse A. Tov and Riccardo Pucella. Stateful Contracts for Affine Types. In *Proceedings of ESOP'10*, LNCS 6012, pages 550–569. Springer, 2010.
- Jesse A. Tov and Riccardo Pucella. Practical affine types. In *Proceedings of POPL'11*, pages 447–458. ACM, 2011.
- Philip Wadler. A complement to blame. In *Proceedings of SNAPL'15*, LIPIcs 32, pages 309–320. Schloss Dagstuhl, 2015.
- Philip Wadler and Robert Bruce Findler. Well-typed programs can't be blamed. In *Proceedings of ESOP'09*, LNCS 5502, pages 1–16. Springer, 2009.

A Supplement to Section 3

We write $\text{fn}(P)$ for the set of free names of P .

Below are the reduction rules of processes that we have omitted from Table 3.

$$\begin{array}{c} \text{[R25]} \\ \frac{e \rightarrow e'}{\langle e \rangle_p \rightarrow \langle e' \rangle_p} \end{array} \quad \begin{array}{c} \text{[R26]} \\ \frac{P \rightarrow Q}{P \parallel R \rightarrow Q \parallel R} \end{array} \quad \begin{array}{c} \text{[R27]} \\ \frac{P \rightarrow Q}{(\nu a)P \rightarrow (\nu a)Q} \end{array} \quad \begin{array}{c} \text{[R28]} \\ \frac{P \equiv P' \quad P' \rightarrow Q' \quad Q' \equiv Q}{P \rightarrow Q} \end{array}$$

Rule [R25] lifts reduction from expressions to processes, and rules [R26–R28] close reductions under parallel compositions, sessions, and structural congruence \equiv , which is the least congruence that satisfies the axioms below:

$$P \parallel Q \equiv Q \parallel P \quad P \parallel (Q \parallel R) \equiv (P \parallel Q) \parallel R \quad \langle () \rangle_p \parallel P \equiv P \quad \frac{a, a^+, a^- \notin \text{fn}(Q)}{(\nu a)P \parallel Q \equiv (\nu a)(P \parallel Q)}$$

In several of the proofs that follow we implicitly use the next results, stating that the reduction of expressions is deterministic.

Proposition 4. *If $e \rightarrow e_1$ and $e \rightarrow e_2$, then $e_1 = e_2$.*

Proof. By inspecting the definition of \rightarrow and that of evaluation contexts. \square

Proposition 5. *If $e \Downarrow v$ and $e \rightarrow e'$, then $e' \Downarrow v$.*

Proof. By definition of $e \Downarrow v$ we have $e = e_0 \rightarrow e_1 \rightarrow \dots \rightarrow e_n = v$ for some $n \geq 0$. From the hypothesis $e \rightarrow e'$ and Proposition 4 we deduce $e' = e_1$. We conclude $e' \Downarrow v$. \square

B Supplement to Section 4

Definition 8 (subkinding). Let \leq be the least total order on kinds such that $\omega \leq 1$. We say that κ_1 is *subkind* of κ_2 if $\kappa_1 \leq \kappa_2$.

The following lemma relates the kind of a value with that of the environment in which it is typed.

Lemma 2. *If $\Gamma \vdash v : t$ and $t :: \kappa$, then $\Gamma :: \kappa$.*

Proof. By induction on the structure of v and by cases on its shape.

Case $v = \mathbf{c}^n v_1 \dots v_n$. Then $\Gamma = \Gamma_0 + \Gamma_1 + \dots + \Gamma_n$ and $\Gamma_0 :: \omega$ and $\Gamma_0 \vdash \mathbf{c}^n : t_1 \rightarrow^{\kappa_1} \dots \rightarrow^{\kappa_{n-1}} t_n \rightarrow^{\kappa_n} t$ and $\Gamma_i \vdash v_i : t_i$ for every $1 \leq i \leq n$. By inspecting Table 4 we observe that $t :: \kappa$ implies $t_i :: \kappa$ for every $1 \leq i \leq n$. By induction hypothesis we deduce $\Gamma_i :: \kappa$ for every $1 \leq i \leq n$, hence we conclude $\Gamma :: \kappa$.

Case $v = \lambda x. e$. From [T-FUN] we deduce $t = t_1 \rightarrow^{\kappa'} t_2$ where $\kappa' \leq \kappa$ and $\Gamma :: \kappa'$, hence we conclude $\Gamma :: \kappa$.

Case $v = a^t$. Then $\Gamma = \Gamma', a^t : t$. From [T-NAME] we deduce $\Gamma' :: \omega$ hence we conclude $\Gamma :: \kappa$.

Case $v = [v_1]^{v_2.p.q}$. From [T-MONITOR] we deduce $\Gamma = \Gamma_1 + \Gamma_2$ and $\Gamma_1 \vdash v_1 : t$ and $\Gamma_2 \vdash v_2 : [t]$. By induction hypothesis we deduce $\Gamma_1 :: \kappa$ and $\Gamma_2 :: \omega$, hence we conclude $\Gamma :: \kappa$. \square

Lemma 3. *If $\Gamma \vdash e : t$, then $\text{fn}(e) \subseteq \text{dom}(\Gamma)$.*

Proof. A simple induction on the derivation of $\Gamma \vdash e : t$. \square

The type system enjoys an almost standard weakening property. Note that weakening is allowed for unlimited environments only.

Lemma 4 (weakening). *If $\Gamma_1 \vdash e : t$ and $\Gamma_1 + \Gamma_2$ is defined and $\Gamma_2 :: \omega$, then $\Gamma_1 + \Gamma_2 \vdash e : t$.*

Proof. A simple induction on the derivation of $\Gamma_1 \vdash e : t$. \square

Lemma 5 (substitution). *If (1) $\Gamma_1, x : s \vdash e : t$ and (2) $\Gamma_2 \vdash v : s$ and (3) $\Gamma_1 + \Gamma_2$ is defined, then $\Gamma_1 + \Gamma_2 \vdash e\{v/x\} : t$.*

Proof. By induction on the derivation of (1) and by cases on the last rule applied. We only discuss a few significant cases.

Case $[\text{T-CONST}]$. Then $e = \mathbf{c}$ for some constant \mathbf{c} and $(\Gamma_1, x : s) :: \omega$. In particular, $s :: \omega$. From Lemma 2 we deduce $\Gamma_2 :: \omega$ and we conclude with one application of $[\text{T-CONST}]$.

Case $[\text{T-NAME}]$. Then $e = u$ and $\Gamma_1, x : s = \Gamma', u : s'$ where $\Gamma' :: \omega$. We distinguish three subcases, depending on whether u is an endpoint or a variable other than x or x itself. If $u = a'$ or $u = y \neq x$, then $s :: \omega$. From Lemma 2 we deduce $\Gamma_2 :: \omega$ and we conclude with one application of $[\text{T-NAME}]$. If $u = x$, then $e\{v/x\} = v$ and we conclude from (2) and Lemma 4.

Case $[\text{T-FUN}]$. Then $e = \lambda y. e'$. Without loss of generality we may assume that $x \neq y$, hence from $[\text{T-FUN}]$ we deduce $t = t' \rightarrow^\kappa t''$ and $\Gamma_1, x : s, y : t' \vdash e' : t''$ and $(\Gamma_1, x : s) :: \kappa$. From Lemma 2 we deduce $\Gamma_2 :: \kappa$. By induction hypothesis we deduce $(\Gamma_1 + \Gamma_2), y : t' \vdash e'\{v/x\} : t''$ where $\Gamma_1 + \Gamma_2 :: \kappa$. We conclude with one application of $[\text{T-FUN}]$.

Case $[\text{T-APP}]$. Then $e = e_1 e_2$ and $\Gamma, x : s = \Gamma_1 + \Gamma_2$ where $\Gamma_1 \vdash e_1 : t' \rightarrow^\kappa t$ and $\Gamma_2 \vdash e_2 : t'$. We distinguish three sub-cases. If $x \in \text{dom}(\Gamma_1) \setminus \text{dom}(\Gamma_2)$, then from Lemma 3 we deduce $x \notin \text{fn}(e_2)$ and $\Gamma_1 = \Gamma'_1, x : s$. By induction hypothesis we deduce $\Gamma'_1 + \Gamma_2 \vdash e_1\{v/x\} : t' \rightarrow^\kappa t$ and we conclude with one application of $[\text{T-APP}]$. If $x \in \text{dom}(\Gamma_2) \setminus \text{dom}(\Gamma_1)$, then from Lemma 3 we deduce $x \notin \text{fn}(e_1)$ and $\Gamma_2 = \Gamma'_2, x : s$. By induction hypothesis we deduce $\Gamma'_2 + \Gamma_1 \vdash e_2\{v/x\} : t'$ and we conclude with one application of $[\text{T-APP}]$. If $x \in \text{dom}(\Gamma_1) \cap \text{dom}(\Gamma_2)$, then $\Gamma_1 = \Gamma'_1, x : s$ and $\Gamma_2 = \Gamma'_2, x : s$. From Definition 3 we deduce $s :: \omega$ and, from Lemma 2, also $\Gamma_2 :: \omega$. Therefore, $\Gamma_1 + \Gamma_2 = \Gamma_2$, again by Definition 3. By induction hypothesis we derive $\Gamma'_1 + \Gamma_2 \vdash e_1\{v/x\} : t' \rightarrow^\kappa t$ and $\Gamma'_2 + \Gamma_2 \vdash e_2\{v/x\} : t'$. We conclude with one application of $[\text{T-APP}]$. \square

Lemma 6 (subject reduction for expressions). *If $\Gamma \vdash e : t$ and $e \rightarrow e'$, then $\Gamma \vdash e' : t$.*

Proof. By cases on the rule used to derive $e \rightarrow e'$. We only discuss a few representative cases.

Case $[\text{R1}]$. Then $e = (\lambda x. e'')v \rightarrow e''\{v/x\} = e'$. From $[\text{T-APP}]$ we deduce that $\Gamma = \Gamma_1 + \Gamma_2$ and $\Gamma_1 \vdash \lambda x. e'' : s \rightarrow^\kappa t$ and $\Gamma_2 \vdash v : s$ and $\Gamma_1 :: \kappa$. Using Lemma 5 we conclude $\Gamma_1 + \Gamma_2 \vdash e''\{v/x\} : t$.

Case $[\text{R5}]$. Then $e = [v]^{\text{flat_c } w, p, q} \rightarrow v \triangleleft^p wv = e'$. From $[\text{T-MONITOR}]$ we deduce $\Gamma = \Gamma_1 + \Gamma_2$ and $\Gamma_1 \vdash v : t$ and $\Gamma_2 \vdash \text{flat_c } w : [t]$. From $[\text{T-APP}]$ and the type scheme of flat_c we deduce $\Gamma_2 = \Gamma_{21} + \Gamma_{22}$ and $\Gamma_{21} \vdash \text{flat_c} : (t \rightarrow \text{bool}) \rightarrow [t]$ and $t :: \omega$ and $\Gamma_{22} \vdash w : t \rightarrow \text{bool}$. From Lemma 2 we deduce $\Gamma_1 :: \omega$, hence $\Gamma_1 + \Gamma_1 = \Gamma_1$ by Definition 3. We conclude with one application of $[\text{T-APP}]$ and one of $[\text{T-BUSY-MONITOR}]$.

Case $[\text{R11}]$. Then $e = \text{dual } !\mathbf{C}.w \rightarrow ?\mathbf{C}.(\lambda x. \text{dual } (wx)) = e'$ and $t = [\overline{T}]$ for some T . From $[\text{T-APP}]$ and the type scheme of dual we deduce $\Gamma = \Gamma_1 + \Gamma_2$ and $\Gamma_1 \vdash \text{dual} : [T] \rightarrow [\overline{T}]$ and $\Gamma_2 \vdash !\mathbf{C}.w : [T]$. From Lemma 2 we deduce $\Gamma_1 :: \omega$ and $\Gamma_2 :: \omega$. From $[\text{T-APP}]$ and the type scheme of receive_d we deduce $\Gamma_2 = \Gamma_{21} + \Gamma_{22} + \Gamma_{23}$ and $T = !s.S$ and $\Gamma_{21} \vdash \text{receive_d} : [s] \rightarrow (s \rightarrow [S]) \rightarrow [T]$ and $\Gamma_{22} \vdash \mathbf{C} : [s]$ and $\Gamma_{23} \vdash w : s \rightarrow [S]$. From Definition 3 we deduce $\Gamma_{2i} :: \omega$ for $1 \leq i \leq 3$. From $[\text{T-FUN}]$ and $[\text{T-APP}]$ we deduce $\Gamma_{23} \vdash \lambda x. \text{dual } (wx) : s \rightarrow [\overline{S}]$. We conclude with two applications of $[\text{T-APP}]$. \square

Next are two results that allow us to reason on the typability of a sub-expression that occurs in the hole of an evaluation context and on the fact that such sub-expression can be replaced by another one having the same type.

Lemma 7 (typability of subterms). *If \mathcal{D} is a derivation of $\Gamma \vdash \mathcal{E}[e] : t$, then there exist Γ_1, Γ_2 , and s such that $\Gamma = \Gamma_1 + \Gamma_2$ and \mathcal{D} has a sub-derivation \mathcal{D}' concluding $\Gamma_1 \vdash e : s$ and the position of \mathcal{D}' in \mathcal{D} corresponds to the position of $[\]$ in \mathcal{E} .*

Proof. By induction on \mathcal{E} . □

Lemma 8 (replacement). *If:*

1. \mathcal{D} is a derivation of $\Gamma_1 + \Gamma_2 \vdash \mathcal{E}[e] : t$,
2. \mathcal{D}' is a sub-derivation of \mathcal{D} concluding that $\Gamma_2 \vdash e : s$,
3. the position of \mathcal{D}' in \mathcal{D} corresponds to the position of $[]$ in \mathcal{E} ,
4. $\Gamma_3 \vdash e' : s$,
5. $\Gamma_1 + \Gamma_3$ is defined,

then $\Gamma_1 + \Gamma_3 \vdash \mathcal{E}[e'] : t$.

Proof. By induction on \mathcal{E} . □

The next result states the obvious property that typing of processes is preserved by structural congruence.

Lemma 9. *If $\Gamma \vdash P$ and $P \equiv Q$, then $\Gamma \vdash Q$.*

Proof. An easy induction on the derivation of $P \equiv Q$ and by cases on the last rule applied. □

In order to state subject reduction for processes, we must take into account the fact that the type associated with session endpoints may change as these endpoints are used for communication. For this reason, we define a reduction relation for type environments that simulates communications as performed by processes at the level of types.

Definition 9 (reduction of type environments). We write \rightarrow for the least relation between type environments such that:

$$\begin{aligned} \text{[R29]} \quad & \Gamma, a^t : !t.T, a^{\bar{t}} : ?t.S \rightarrow \Gamma, a^t : T, a^{\bar{t}} : S \\ \text{[R30]} \quad & \Gamma, a^t : T_1 \oplus T_2, a^{\bar{t}} : S_1 \& S_2 \rightarrow \Gamma, a^t : T_1, a^{\bar{t}} : S_1 \\ \text{[R31]} \quad & \Gamma, a^t : T_1 \oplus T_2, a^{\bar{t}} : S_1 \& S_2 \rightarrow \Gamma, a^t : T_2, a^{\bar{t}} : S_2 \end{aligned}$$

and we write \rightarrow^* for the reflexive, transitive closure of \rightarrow .

Definition 10 (balanced type environment). We say that Γ is *balanced* if $a^t, a^{\bar{t}} \in \text{dom}(\Gamma)$ implies $\Gamma(a^t) = \overline{\Gamma(a^{\bar{t}})}$.

Note that every environment Γ such that $\Gamma :: \omega$ is trivially balanced, since it has no session endpoints in its domain. We need a few properties that concern the reduction of type environments:

Proposition 6. *Let $\Gamma \rightarrow \Gamma'$. The following properties hold:*

1. *If Γ is balanced, then Γ' is balanced as well.*
2. *If $\Gamma + \Gamma''$ is defined, then $\Gamma + \Gamma'' \rightarrow \Gamma' + \Gamma''$.*

Proof. Trivial from Definitions 9 and 10. □

In the proof of the following Lemma we introduce some handy abbreviations for dealing with sets of environments and sets of judgments. In particular, we write $\vec{\Gamma}$ in place of $\Gamma_1 + \dots + \Gamma_n$ where $\Gamma_1 + \dots + \Gamma_n = \Gamma$ and we write $\vec{\Gamma} \vdash e : t$ in place of a set of judgments $\Gamma_i \vdash e_i : t_i$ (usually the t_i are all the same type).

Lemma 10. *If Γ is balanced and $\Gamma \vdash P$ and $P \rightarrow Q$, then there exists Γ' such that $\Gamma \rightarrow^* \Gamma'$ and $\Gamma' \vdash Q$.*

Proof. By induction on the derivation of $P \rightarrow Q$ and by cases on the last rule applied. We only discuss a few representative cases.

Case [R17]. Then we have

$$\begin{aligned} P &= \langle \mathcal{E}[\mathbf{send} \ v \ [a^t]^{\overrightarrow{!C;D;\sigma}}] \rangle_p \parallel \langle \mathcal{E}'[\mathbf{receive} \ [a^{\bar{t}}]^{\overrightarrow{?E;F;\bar{\ell}}}] \rangle_q \\ Q &= \langle \mathcal{E}[[a^t]^{\overrightarrow{D;\sigma}}] \rangle_p \parallel \langle \mathcal{E}'[(\llbracket [v]^{\overrightarrow{C;-\sigma}} \rrbracket^{\overrightarrow{E;\bar{\ell}}}, [a^{\bar{t}}]^{\overrightarrow{F;\bar{\ell}}})] \rangle_q \end{aligned}$$

Let P_1 and P_2 be the two threads in P . From [T-PAR] we deduce $\Gamma = \Gamma_1 + \Gamma_2$ and $\Gamma_i \vdash P_i$ for every $i = 1, 2$. From [T-THREAD] we deduce $\Gamma_1 \vdash \mathcal{E}[\dots] : \mathbf{unit}$ and $\Gamma_2 \vdash \mathcal{E}'[\dots] : \mathbf{unit}$, where we have elided the expressions in the holes of \mathcal{E} and \mathcal{E}' .

From Lemma 7, [T-APP] and the type scheme of **send** we deduce the following facts:

- $\Gamma_1 = \Gamma_{10} + \Gamma_{11} + \Gamma_{12} + \Gamma_{13} + \overrightarrow{\Gamma_{14}} + \overrightarrow{\Gamma_{15}}$;
- $\Gamma_{11} \vdash \mathbf{send} : t \rightarrow !t.T \rightarrow^\kappa T$ where $t :: \kappa$;
- $\Gamma_{12} \vdash v : t$;
- $\Gamma_{13} \vdash a^t : !t.T$, hence $\Gamma_{13} = \Gamma''_{13}, a^t : !t.T$ where $\Gamma''_{13} :: \omega$;
- $\overrightarrow{\Gamma_{14}} \vdash C : [\overrightarrow{t}]$ and $\overrightarrow{\Gamma_{15}} \vdash D : [\overrightarrow{T}]$.

From Lemma 7, [T-APP] and the type scheme of **receive** we deduce the following facts:

- $\Gamma_2 = \Gamma_{20} + \Gamma_{21} + \Gamma_{22} + \overrightarrow{\Gamma_{23}} + \overrightarrow{\Gamma_{24}}$;
- $\Gamma_{21} \vdash \mathbf{receive} : ?s.S \rightarrow s \times S$;
- $\Gamma_{22} \vdash a^{\bar{t}} : ?s.S$, hence $\Gamma_{22} = \Gamma'_{22}, a^{\bar{t}} : ?s.S$ where $\Gamma'_{22} :: \omega$;
- $\overrightarrow{\Gamma_{23}} \vdash E : [\overrightarrow{s}]$ and $\overrightarrow{\Gamma_{24}} \vdash F : [\overrightarrow{S}]$.

From Lemma 2 we deduce $\Gamma_{11} :: \omega$ and $\Gamma_{21} :: \omega$. From the hypothesis that Γ is balanced we deduce that $!t.T = ?s.S$, that is $t = s$ and $T = \overline{S}$ or, equivalently, $\overline{T} = S$.

Let $\Gamma'_{13} \stackrel{\text{def}}{=} \Gamma''_{13}, a^t : T$ and $\Gamma'_{22} \stackrel{\text{def}}{=} \Gamma'_{22}, a^{\bar{t}} : \overline{T}$. From [T-NAME] and [T-MONITOR] we deduce $\Gamma'_{13} + \overrightarrow{\Gamma_{15}} \vdash [a^t]^{\overrightarrow{D;\sigma}} : T$. From [T-MONITOR] we deduce $\Gamma_{12} + \overrightarrow{\Gamma_{14}} + \overrightarrow{\Gamma_{23}} \vdash \llbracket [v]^{\overrightarrow{C;-\sigma}} \rrbracket^{\overrightarrow{E;\bar{\ell}}} : t$. From [T-NAME] and [T-MONITOR] we deduce $\Gamma'_{22} \vdash [a^{\bar{t}}]^{\overrightarrow{F;\bar{\ell}}} : \overline{T}$.

Let $\Gamma'_1 \stackrel{\text{def}}{=} \Gamma_{10} + \Gamma_{11} + \Gamma'_{13} + \overrightarrow{\Gamma_{15}}$ and $\Gamma'_2 \stackrel{\text{def}}{=} \Gamma_{20} + \Gamma_{21} + \Gamma_{12} + \overrightarrow{\Gamma_{14}} + \overrightarrow{\Gamma_{23}} + \Gamma'_{22} + \overrightarrow{\Gamma_{24}}$. From Lemma 8 and Lemma 4 we deduce $\Gamma'_1 \vdash \mathcal{E}[\dots] : \mathbf{unit}$ and $\Gamma'_2 \vdash \mathcal{E}'[\dots] : \mathbf{unit}$ where we have elided the expressions in the holes of Q .

Let $\Gamma' \stackrel{\text{def}}{=} \Gamma'_1 + \Gamma'_2$. We derive $\Gamma' \vdash Q$ with one application of [T-PAR] and we conclude by observing that $\Gamma \rightarrow \Gamma'$.

Case [R25]. Then $P = \langle \mathcal{E}[e] \rangle_p \rightarrow \langle \mathcal{E}[e'] \rangle_p = Q$ where $e \rightarrow e'$. From [T-THREAD] we deduce $\Gamma \vdash \langle \mathcal{E}[e] \rangle_p : \mathbf{unit}$. From Lemma 7 we deduce that there exist Γ_1, Γ_2 , and t such that $\Gamma = \Gamma_1 + \Gamma_2$ and $\Gamma_1 \vdash e : t$. From Lemma 6 we deduce $\Gamma_1 \vdash e' : t$. We conclude by Lemma 8 and one application of [R-THREAD].

Case [R26]. Then $P = P_1 \parallel P_2 \rightarrow P'_1 \parallel P_2 = Q$ where $P_1 \rightarrow P'_1$. From [T-PAR] we deduce $\Gamma = \Gamma_1 + \Gamma_2$ and $\Gamma_i \vdash P_i$ for $i = 1, 2$. By induction hypothesis we derive $\Gamma'_i \vdash P'_i$ for some Γ'_i such that $\Gamma_1 \rightarrow^* \Gamma'_1$. By Proposition 6 we deduce that $\Gamma'_1 + \Gamma_2$ is defined and we conclude by taking $\Gamma' \stackrel{\text{def}}{=} \Gamma'_1 + \Gamma_2$ and one application of [T-PAR].

Case [R27]. Then $P = (\nu a)P' \rightarrow (\nu a)P'' = Q$ where $P' \rightarrow P''$. From [T-SESSION] we deduce $\Gamma, a^+ : T, a^- : \overline{T} \vdash P'$. Note that $\Gamma, a^+ : T, a^- : \overline{T}$ is balanced if so is Γ . By induction hypothesis we deduce that there exist Γ' and S such that $\Gamma, a^+ : T, a^- : \overline{T} \rightarrow^* \Gamma', a^+ : S, a^- : \overline{S}$. We conclude with one application of [T-SESSION].

Case [R28]. Easy consequence of Lemma 9 and the induction hypothesis. \square

Theorem 1. *If $\Gamma :: \omega$ and $\Gamma \vdash P$ and $P \rightarrow Q$, then $\Gamma \vdash Q$.*

Proof. Straightforward corollary of Lemma 10 using the fact that every unlimited type environment is trivially balanced. \square

C Supplement to Section 5

C.1 Properties of \leq

Proposition 7. \leq is transitive.

Proof. This is a standard proof using the coinduction principle to show that $\mathcal{S} \stackrel{\text{def}}{=} \{(e_1, e_2) \mid \exists e : e_1 \leq e \wedge e \leq e_2\} \subseteq \leq$. The result follows easily from the definition of \mathcal{S} and Definition 5. \square

The next result states that entailment is closed under reductions.

Proposition 8. Let $e_1 \leq e_2$. Then the following properties hold:

1. $e_1 \rightarrow e'_1$ implies $e'_1 \leq e_2$;
2. $e_2 \rightarrow e'_2$ implies $e_1 \leq e'_2$.

Proof. An easy consequence of Proposition 5. \square

We now prove a few results concerning the semantics of **dual** and define a correspondence \bowtie akin to duality for session types that relates a contract C with that obtained by evaluating **dual** C .

Proposition 9. Let **dual**(\cdot) be the function inductively defined by

$$\begin{aligned}
 \text{dual}(\text{flat_c } w) &\stackrel{\text{def}}{=} \text{flat_c } w \\
 \text{dual}(\text{end_c}) &\stackrel{\text{def}}{=} \text{end_c} \\
 \text{dual}(!C;D) &\stackrel{\text{def}}{=} ?C; \text{dual}(D) \\
 \text{dual}(?C;D) &\stackrel{\text{def}}{=} !C; \text{dual}(D) \\
 \text{dual}(!C.w) &\stackrel{\text{def}}{=} ?C. \lambda x. \text{dual}(wx) \\
 \text{dual}(?C.w) &\stackrel{\text{def}}{=} !C. \lambda x. \text{dual}(wx) \\
 \text{dual}(!C.E:F) &\stackrel{\text{def}}{=} ?C. \text{dual}(E): \text{dual}(F) \\
 \text{dual}(?C.E:F) &\stackrel{\text{def}}{=} !C. \text{dual}(E): \text{dual}(F)
 \end{aligned}$$

Then **dual** $C \Downarrow \text{dual}(C)$ for every C .

Proof. An easy induction on C . \square

Definition 11 (duality). We write \bowtie for the largest relation between expressions such that $e_1 \bowtie e_2$ implies either:

- $e_1 \Downarrow \text{flat_c } w$ and $e_2 \Downarrow \text{flat_c } w$, or
- $e_1 \Downarrow \text{end_c}$ and $e_2 \Downarrow \text{end_c}$, or
- $e_1 \Downarrow !C;D$ and $e_2 \Downarrow ?C;E$ and $D \bowtie E$, or
- $e_1 \Downarrow ?C;D$ and $e_2 \Downarrow !C;E$ and $D \bowtie E$, or
- $e_1 \Downarrow !\text{flat_c } w.w_1$ and $e_2 \Downarrow ?\text{flat_c } w.w_2$ and $w_1v \bowtie w_2v$ for every $v \in w$, or
- $e_1 \Downarrow ?\text{flat_c } w.w_1$ and $e_2 \Downarrow !\text{flat_c } w.w_2$ and $w_1v \bowtie w_2v$ for every $v \in w$, or
- $e_1 \Downarrow !C.D:E$ and $e_2 \Downarrow ?C.F:G$ and $D \bowtie F$ and $E \bowtie G$, or
- $e_1 \Downarrow ?C.D:E$ and $e_2 \Downarrow !C.F:G$ and $D \bowtie F$ and $E \bowtie G$.

Informally, two contracts C_1 and C_2 are compatible if every output in C_i is matched by a corresponding input in C_{3-i} and the contract in the input prefix of C_{3-i} is entailed by the contract in the corresponding output prefix of C_i . The continuation of dependent contracts is obtained just like in the case of entailment.

The formal relationship between the semantics of **dual** and the relation \bowtie is stated by the following proposition:

Proposition 10. *If $e \leq e'$, then $e \bowtie \mathbf{dual} \ e$.*

Proof. It is enough to show that the relation \mathcal{R} defined by

$$\mathcal{R} \stackrel{\text{def}}{=} \{(e, e') \mid \exists C : C \leq C \wedge e \Downarrow C \wedge e' \Downarrow \mathbf{dual}(C)\}$$

is included in \bowtie . Suppose $e \mathcal{R} \ e'$. From the definition of \mathcal{R} we know that there exists C such that $C \leq C$ and $e \Downarrow C$ and $e' \Downarrow \mathbf{dual}(C)$. We reason by cases on the shape of C , discussing only some representative cases.

Case $C = \mathbf{flat_c} \ w$. We conclude by observing that $\mathbf{dual}(C) = \mathbf{flat_c} \ w$.

Case $C = \mathbf{end_c}$. We conclude by observing that $\mathbf{dual}(C) = \mathbf{end_c}$.

Case $C = ?D;E$. Then $\mathbf{dual}(C) = !D;\mathbf{dual}(E)$. From the hypothesis $C \leq C$ we deduce $E \leq E$. We conclude by observing that $E \mathcal{R} \ \mathbf{dual}(E)$ by definition of \mathcal{R} .

Case $C = ?D.w$. Then $\mathbf{dual}(C) = !D.\lambda x.\mathbf{dual} \ (wx)$. From the hypothesis $C \leq C$ we deduce $D = \mathbf{flat_c} \ w'$ for some w' and $wv \leq wv$ for every $v \in w'$. Take $v \in w'$. Then there exists C' such that $wv \Downarrow C'$. Now we derive $(\lambda x.\mathbf{dual} \ (wx))v \rightarrow \mathbf{dual} \ (wv) \Downarrow \mathbf{dual}(C')$, therefore $wv \mathcal{R} \ (\lambda x.\mathbf{dual} \ (wx))v$ by definition of \mathcal{R} . \square

C.2 Monitor Invariants

We prove an invariant property of the big-step semantics of λCoS , as introduced in Section 3, which is instrumental in several proofs that follow. We formalize the notion of balanced process, one in which the innermost monitors associated with dual endpoints have dual contracts and symmetric pair of blame labels.

Definition 12 (balanced process). We say that P is *balanced* if $P \supset [a^t]^{e_1, \sigma}$ and $P \supset [a^{\bar{t}}]^{e_2, \varrho}$ imply $e_1 \bowtie e_2$ and $\sigma = \neg\varrho$.

The following results show that the property of being balanced is preserved by reductions of the big-step semantics.

Lemma 11. *Let $\Gamma \vdash e : t$. If $\langle e \rangle_p \supset [a^t]^{e_1, p, q}$ and $e \rightarrow e'$ then $\langle e' \rangle_p \supset [a^t]^{e'_1, p, q}$ and $e_1 \rightarrow^* e'_1$.*

Proof. By induction on the derivation of $e \rightarrow e'$. \square

Lemma 12. *Let $\Gamma \vdash P$ and Γ balanced. If P is balanced and $P \rightarrow Q$ and $Q \supset [a^t]^{e_2, p, q}$, then either*

1. $P \supset [a^t]^{e_1, p, q}$ and $P \not\supset a^{\bar{t}}$ and $Q \not\supset a^{\bar{t}}$ and $e_1 \rightarrow^* e_2$;
2. $P \supset [a^t]^{e_1, p, q}$ and $P \supset [a^{\bar{t}}]^{e_1, q, p}$ and $Q \supset [a^{\bar{t}}]^{e_1, q, p}$ and $e_1 \bowtie e_2$;
3. $P \not\supset a^t$ and $P \not\supset a^{\bar{t}}$ and $Q \supset [a^{\bar{t}}]^{e_1, q, p}$ and $e_1 \bowtie e_2$.

Proof. By induction on the derivation of $P \rightarrow Q$ and by cases on the last rule applied. We analyze some representative cases.

Case [R15]. Then,

$$\begin{aligned} P &= \langle \mathcal{E}[\mathbf{connect} \ b] \rangle_p \parallel b \Leftarrow_q^c v \\ Q &= (\nu c)(\langle \mathcal{E}[c^+]^{c, q, p} \rangle_p \parallel \langle (\nu [c^-] \mathbf{dual} \ c, p, q) \rangle_q) \parallel b \Leftarrow_q^c v \end{aligned}$$

where c is fresh. We proceed by cases:

1. Assume $P \supset [a^t]^{e_1, p, q}$ and $P \not\supset a^{\bar{t}}$. Since $\Gamma \vdash P$, both v and C do not contain runtime syntax. Hence, $\langle \mathcal{E}[\mathbf{connect} \ b] \rangle_p \supset [a^t]^{e_1, p, q}$ and $\langle \mathcal{E}[\mathbf{connect} \ b] \rangle_p \not\supset a^{\bar{t}}$. Since c is fresh, $c \neq a$. Therefore, $\langle \mathcal{E}[c^+]^{c, q, p} \rangle_p \supset [a^t]^{e_1, p, q}$ and $\langle \mathcal{E}[c^+]^{c, q, p} \rangle_p \not\supset a^{\bar{t}}$. The proof is completed by taking $e_1 = e_2$.

2. Assume $P \supset [a^t]^{e_2, p, q}$ and $P \supset [a^{\bar{t}}]^{e_1, r, s}$. Since P is balanced, $e_1 \bowtie e_2$ and $p = s$ and $q = r$. As in the previous case, we conclude that $\langle \mathcal{E}[[c^+]^{c, q, p}] \rangle_p \supset [a^t]^{e_2, p, q}$ and $\langle \mathcal{E}[[c^+]^{c, q, p}] \rangle_p \supset [a^{\bar{t}}]^{e_1, q, p}$.
3. Assume $P \not\supset a^t$ and $P \not\supset a^{\bar{t}}$. Hence, $a = c$. Then, the proof is completed by using Proposition 10 to conclude $C \bowtie \text{dual } C$.

Case [R16]. Then

$$\begin{aligned} P &= (\nu c) \langle \langle \mathcal{E}[\text{close } [c^+]^{\overrightarrow{\text{end.c, } \sigma}}] \rangle_p \parallel \langle \mathcal{E}'[\text{close } [c^-]^{\overrightarrow{\text{end.c, } \varrho}}] \rangle_q \rangle \\ Q &= \langle \mathcal{E}[(\text{O})] \rangle_p \parallel \langle \mathcal{E}'[(\text{O})] \rangle_q \end{aligned}$$

The two first cases follow analogously to rule [R15]. The third case is not possible.

Case [R17]. Then,

$$\begin{aligned} P &= \langle \mathcal{E}[\text{send } v [c^t]^{\overrightarrow{\text{!C;D, } \sigma}}] \rangle_p \parallel \langle \mathcal{E}'[\text{receive } [c^{\bar{t}}]^{\overrightarrow{?E;F, } \varrho}] \rangle_q \\ Q &= \langle \mathcal{E}[[c^t]^{\overrightarrow{\text{D, } \sigma}}] \rangle_p \parallel \langle \mathcal{E}'[(\text{C}[v]^{\overrightarrow{\text{C, } \sigma}})^{\overrightarrow{\text{E, } \varrho}}, [c^{\bar{t}}]^{\overrightarrow{\text{F, } \varrho}}] \rangle_q \end{aligned}$$

If $a \neq c$, then the proof follows as for the case 1 in rule [R15]. Otherwise, when $a = c$, we are on case $P \supset a^t$ and $P \supset a^{\bar{t}}$. We know that $\overrightarrow{\text{!C;D}} \bowtie \overrightarrow{?E;F}$ and $\overrightarrow{\sigma} = \overrightarrow{\varrho}$ because P is balanced. Consequently, $\overrightarrow{\text{D}} \bowtie \overrightarrow{\text{F}}$ by duality.

Case [R25]. Then, $P = \langle e \rangle_p$ and $Q = \langle e' \rangle_p$ and $e \rightarrow e'$. Since $\Gamma \vdash P$, $P \supset a^t$ if and only if $Q \supset a^t$. Then, we have the following two cases:

1. If $P \supset [a^t]^{e_1, p, q}$ and $P \not\supset a^{\bar{t}}$, then the proof follows from Lemma 11.
2. If $P \supset [a^t]^{e_3, r, s}$ and $P \supset [a^{\bar{t}}]^{e_4, s, r}$, then $P \supset [a^t]^{e_3, r, s}$ and $P \supset [a^{\bar{t}}]^{e_4, s, r}$ and $e_3 \bowtie e_4$ because P is balanced. By Lemma 11, $Q \supset [a^t]^{e_1, s, r}$ and $e_3 \rightarrow^* e_1$. Similarly, $P \supset [a^{\bar{t}}]^{e_2, s, r}$ and $e_4 \rightarrow^* e_2$. By Lemma 5, $e_3 \Downarrow C$ implies $e_1 \Downarrow C$, and $e_4 \Downarrow D$ implies $e_2 \Downarrow D$. Therefore, $e_1 \bowtie e_2$ implies $e_3 \bowtie e_4$.

Case [R26]. Then, $P = R_1 \parallel R_2$ and $Q = R'_1 \parallel R_2$ and $R_1 \rightarrow R'_1$. Lemma 10, ensures that there exists Γ' such that $\Gamma' \vdash Q$. Moreover, P balanced implies R_1 and R_2 balanced as well. Firstly, we consider the case $R'_1 \supset [a^t]^{e_2, p, q}$. By induction hypothesis we have one of the following cases.

1. $R_1 \supset [a^t]^{e_1, p, q}$ and $R_1 \not\supset a^{\bar{t}}$ and $R'_1 \not\supset a^{\bar{t}}$ and $e_1 \rightarrow^* e_2$. Note that $R_2 \not\supset a^t$ because P is well-typed. If $R_2 \not\supset a^{\bar{t}}$, we are done because $Q \not\supset a^{\bar{t}}$. Otherwise, when $R_2 \supset [a^{\bar{t}}]^{e_3, r, s}$, P balanced implies $e_1 \bowtie e_3$ and $p = s$ and $q = r$. Since $e_1 \rightarrow^* e_2$, $e_1 \Downarrow C$ implies $e_2 \Downarrow C$. Hence, $e_1 \bowtie e_3$ implies $e_2 \bowtie e_3$.
2. $R_1 \supset [a^t]^{e_1, q, p}$ and $R_1 \supset [a^{\bar{t}}]^{e_1, q, p}$ and $R'_1 \supset [a^{\bar{t}}]^{e_1, q, p}$ and $e_1 \bowtie e_2$. Moreover, $R_2 \not\supset a^{\bar{t}}$ because Q is well-typed.
3. $R_1 \not\supset a^t$ and $R_1 \not\supset a^{\bar{t}}$ and $R'_1 \supset [a^{\bar{t}}]^{e_1, q, p}$ and $e_1 \bowtie e_2$. Moreover $R_2 \not\supset a^{\bar{t}}$ because Q is well-typed.

Now we consider the case $R_2 \supset [a^t]^{e_2, p, q}$. Since P is well-typed, $R_1 \not\supset a^t$. We proceed by cases:

1. Assume $R_2 \supset a^{\bar{t}}$. Then, it follows straightforwardly because P is balanced.
2. Assume $R_1 \supset a^{\bar{t}}$. The case is completed by using induction hypothesis on $R_1 \rightarrow R'_1$ and by noting that only case 1 is applicable because $R_1 \not\supset a^t$.
3. Assume $R_1 \not\supset a^{\bar{t}}$ and $R_2 \not\supset a^{\bar{t}}$. Therefore, $P \not\supset a^{\bar{t}}$. Since Q is well-typed, it cannot be the case that $R'_1 \supset a^t$ and $R'_1 \supset a^{\bar{t}}$. Consequently, $Q \not\supset a^{\bar{t}}$. Then, $P \supset [a^t]^{e_2, p, q}$ and $e_2 \rightarrow^* e_2$. \square

Lemma 13. *Let $\Gamma \vdash P$ where Γ is balanced. If P is balanced and $P \rightarrow^* Q$, then Q is balanced as well.*

Proof. By induction on the length of the derivation and using Lemma 12 to conclude that $P \rightarrow Q$ and P balanced imply Q balanced. \square

C.3 Blame Soundness

Definition 13 (reachable expressions and processes). We say that Q is *reachable* if there exists a user process P such that $P \rightarrow^* Q$. We say that e is reachable if there exists a reachable process $P \supset e$.

We introduce the syntactic category of *monitored expressions* as a generalization of monitored endpoints:

Monitored Expression $\bar{e} ::= v \mid \mathbf{blame} \ p \mid v \triangleleft^p e \mid [\bar{e}]^{e,\sigma}$

This category allows us to focus on a small class of expressions that can be wrapped by a monitor. Indeed, it is possible to show that all expressions wrapped by a monitor in a reduct of a user process are monitored expressions:

Lemma 14. *If P is a reachable process and $P \supset [e]$, then e is a monitored expression.*

Proof. An easy inspection of the definition of \rightarrow . \square

We now introduce a generalization of Definition 7 that is necessary in order to prove blame soundness. The main difference between the two definitions is that contracts occurring in monitors are not required to be values, but they can be expressions that need to be evaluated. Also, the process is assumed to be reachable.

Definition 14 (generalized module soundness). Let P be a reachable process. We say that module p is *sound* in P if:

1. $P \supset v \triangleleft^p e$ implies $e \Downarrow \mathbf{true}$;
2. $P \supset [v]^{e_0, p, -}$ and $e_0 \Downarrow \mathbf{flat_c} \ w$ imply $v \in w$;
3. $P \supset [v \triangleleft e]^{e_0, p, -}$ and $e \Downarrow \mathbf{true}$ and $e_0 \Downarrow \mathbf{flat_c} \ w$ imply $v \in w$;
4. $P \supset [[_]^{e_1, -, q}]^{e_2, r, -}$ and $p \in \{q, r\}$ imply $e_1 \leq e_2$;

We say that p is sound in e if p is sound in $\langle e \rangle_p$.

We can now proceed with the results showing that (generalized) soundness is preserved by reductions. To keep the proofs manageable, we split the results for expressions and processes.

Lemma 15. *Let $\Gamma \vdash \mathcal{E}[e] : t$ where $\mathcal{E}[e]$ is reachable. If p is sound in $\mathcal{E}[e]$ and $e \rightarrow e'$, then p is sound also in $\mathcal{E}[e']$.*

Proof. We proceed by induction on the structure of \mathcal{E} and by cases on its shape.

Case $\mathcal{E} = []$. We distinguish several sub-cases, depending on the reduction rule used for deriving $e \rightarrow e'$.

If [R1] is used, then $e = (\lambda x. e'')v \rightarrow e''\{v/x\} = e'$. This rule cannot generate monitored expressions that were not already present in e because x , which is being replaced by v , is not a monitored expression. The sub-cases in which one of [R2-R4] is used are analogous.

If [R5] is used, then $e = [v]^{e_0, p, -} \rightarrow v \triangleleft^p wv = e'$. From the hypothesis that p is sound in e we deduce $v \in w$, hence we conclude $wv \Downarrow \mathbf{true}$ and condition 1 of Definition 14 is satisfied.

The rules [R6-R7] eliminate one busy monitor, hence they do not generate new monitored expressions.

Finally, none of the rules [R8-R14] generates or affects monitored expressions.

Case $\mathcal{E} = [\mathcal{E}']^{e_0, q, r}$. By induction hypothesis we deduce that p is sound in $\mathcal{E}'[e']$. We now consider the possibilities in which the reduct matches one of the sub-expressions in the conditions of Definition 14. In doing so, recall that $\mathcal{E}'[e]$ must be a monitored expression by Lemma 14 and that it cannot be a value or a blame because $e \rightarrow e'$ by hypothesis.

Regarding condition 2, suppose $q = p$ and $\mathcal{E}'[e] = v \leftarrow \mathbf{true} \rightarrow v$. From the hypothesis that p is sound in $\mathcal{E}[e]$ we conclude that $e_0 \Downarrow \mathbf{flat_c} w$ implies $v \in w$.

Regarding condition 3, suppose $q = p$ and that $\mathcal{E}'[e] = [v]^{\mathbf{flat_c} w, \dashv} \rightarrow v \leftarrow wv$. From the hypothesis that p is sound in $\mathcal{E}[e]$ we deduce that $\mathbf{flat_c} w \leq e_0$. From the definition of \leq we deduce $e_0 \Downarrow \mathbf{flat_c} w'$ and $v \in w$ implies $v \in w'$. Hence we conclude that $wv \Downarrow \mathbf{true}$ implies $v \in w'$.

Regarding condition 3, suppose $q = p$ and $\mathcal{E}' = v \leftarrow \mathcal{E}''$ and $\mathcal{E}''[e] \rightarrow \mathcal{E}''[e']$. From the hypothesis that p is sound in $\mathcal{E}[e]$ we deduce that $\mathcal{E}''[e] \Downarrow \mathbf{true}$ and $e_0 \Downarrow \mathbf{flat_c} w$ imply $v \in w$. We conclude that $\mathcal{E}''[e'] \Downarrow \mathbf{true}$ and $e_0 \Downarrow \mathbf{flat_c} w$ imply $v \in w$.

Regarding condition 4, suppose that $\mathcal{E}' = [\mathcal{E}'']^{e_1, \dashv, r}$ and $p \in \{q, r\}$. From the hypothesis that p is sound in $\mathcal{E}[e]$ we conclude that $e_1 \leq e_0$.

Regarding condition 4, suppose that $\mathcal{E}' = [v]^{\mathcal{E}'', \dashv, r}$ and $p \in \{q, r\}$ and $\mathcal{E}''[e] \rightarrow \mathcal{E}''[e']$. From the hypothesis that p is sound in $\mathcal{E}[e]$ we deduce $\mathcal{E}''[e] \leq e_0$. We conclude $\mathcal{E}''[e'] \leq e_0$ by Proposition 8.

Case $\mathcal{E} = [v]^{\mathcal{E}', q, \dashv}$. By induction hypothesis we deduce that p is sound in $\mathcal{E}'[e']$. We now consider the possibilities in which the reduct matches one of the sub-expressions in the conditions of Definition 14.

Regarding condition 2, suppose $q = p$. From the hypothesis that p is sound in $\mathcal{E}[e]$ we deduce that $\mathcal{E}'[e] \Downarrow \mathbf{flat_c} w$ implies $v \in w$. We conclude by observing that $\mathcal{E}'[e] \Downarrow \mathbf{flat_c} w$ implies $\mathcal{E}'[e'] \Downarrow \mathbf{flat_c} w$.

Regarding condition 4, suppose that $v = [_]^{v_0, \dashv, r}$ and $p \in \{q, r\}$. From the hypothesis that p is sound in $\mathcal{E}[e]$ we deduce that $v_0 \leq \mathcal{E}'[e]$. We conclude $v_0 \leq \mathcal{E}'[e']$ by Proposition 8.

Case $\mathcal{E} = v \leftarrow \mathcal{E}'$. By induction hypothesis we deduce that p is sound in $\mathcal{E}'[e']$. Regarding condition 2, suppose $q = p$. From the hypothesis that p is sound in $\mathcal{E}[e]$ we deduce that $\mathcal{E}'[e] \Downarrow \mathbf{true}$, hence we conclude $\mathcal{E}'[e'] \Downarrow \mathbf{true}$.

Cases $\mathcal{E} = \mathcal{E}'e$ or $\mathcal{E} = v\mathcal{E}'$ or $\mathcal{E} = \mathbf{let} x, y = \mathcal{E}' \mathbf{in} e_0$ or $\mathcal{E} = \mathbf{case} \mathcal{E}' \mathbf{of} e_1 \mid e_2$. The result follows immediately by the induction hypothesis. \square

Lemma 16. *Let $\Gamma \vdash P$ where Γ is balanced. If p is locally correct and sound in P and $P \rightarrow Q$, then p is sound also in Q .*

Proof. By induction on the derivation of $P \rightarrow Q$ and by cases on the last rule applied. From the hypothesis $\Gamma \vdash P$ and Theorem 1 we deduce that $\Gamma' \vdash Q$ for some Γ' such that $\Gamma \rightarrow^* \Gamma'$.

Case [R15]. Then

- $P = \langle \mathcal{E}[\mathbf{connect} a] \rangle_q \parallel a \leftarrow_r^C v$ and
- $Q = (\nu b)(\langle \mathcal{E}[[b^+]^{C, r, q}] \rangle_q \parallel \langle (v [b^-]^{\mathbf{dual} C, q, r}) \rangle_r) \parallel a \leftarrow_r^C v$.

Concerning condition 2 of Definition 14, suppose $q = p$. From $\Gamma' \vdash Q$, [T-PAR], [T-SESSION], and Lemma 7 we deduce that there exist Γ_1, Γ_2 , and T such that $\Gamma_1 \vdash [b^+]^{C, q, p} : T$. From [T-MONITOR] and the type scheme of $\mathbf{flat_c}$ we deduce that C cannot be of the form $\mathbf{flat_c} w$, because T is linear. Therefore, condition 2 holds vacuously. A similar reasoning applies if $r = p$.

Case [R16]. This reduction only eliminates monitored expressions, hence there is nothing to discuss.

Case [R17]. Then

- $P = \langle \mathcal{E}[\mathbf{send} v [a^t]^{\overrightarrow{!C; D; \sigma}}] \rangle_q \parallel \langle \mathcal{E}'[\mathbf{receive} [a^r]^{\overrightarrow{?E; F; \varrho}}] \rangle_r$ and
- $Q = \langle \mathcal{E}[[a^t]^{\overrightarrow{D; \sigma}}] \rangle_q \parallel \langle \mathcal{E}'[(\llbracket v \rrbracket^{\overleftarrow{C; \neg \sigma}} \overrightarrow{E; \varrho}, [a^r]^{\overrightarrow{F; \varrho}})] \rangle_r$.

Let $n \geq 1$ be the number of monitors around the endpoint a^l and $m \geq 1$ be the number of monitors around the endpoint a^r . By Proposition 1 we know that q is negative in σ_n and that r is negative in ϱ_m . We now discuss the conditions of Definition 14 that may apply because of the monitored expressions that have been generated or affected by the reduction.

Suppose $\sigma_n = _, p$. We distinguish two sub-cases, according to whether v is a monitored endpoint or not. In the first sub-case we have $v = [_]^{C, -p}$, hence $Q \supset [[_]^{C, -p}]^{C_n, p, -}$. From the hypothesis that p is locally correct in P we deduce $c \leq C_n$, therefore condition 4 is satisfied by this monitored expression. In the second sub-case we have $C_n = \mathbf{flat_c} w$ for some w . From the hypothesis that p is locally correct in P we deduce $v \in w$, therefore condition 2 is satisfied by this monitored expression.

Suppose that $Q \supset [[_]^{C_{i+1}, -, q_0}]^{C_i, r_0, -}$ and $p \in \{q_0, r_0\}$ for some $1 \leq i < n$. Then $P \supset [[_]^{!C_i; D_i, -, r_0}]^{!C_{i+1}; D_{i+1}, q_0, -}$. From the hypothesis that p is sound in P we deduce $!C_i; D_i \leq !C_{i+1}; D_{i+1}$. From the definition of \leq we deduce $C_{i+1} \leq C_i$, therefore condition 4 is satisfied for this monitored expression. Analogous arguments can be used to verify that condition 4 is satisfied also for $Q \supset [[_]^{D_i, -, q_0}]^{D_{i+1}, r_0, -}$ and $Q \supset [[_]^{E_i, -, q_0}]^{E_{i+1}, r_0, -}$ and $Q \supset [[_]^{F_i, -, q_0}]^{F_{i+1}, r_0, -}$ when $p \in \{q_0, r_0\}$.

From Lemma 13 we deduce $!C_1; D_1 \bowtie ?E_1; F_1$. By definition of \bowtie , we deduce $C_1 \leq E_1$ and $D_1 \bowtie F_1$. Also, suppose that $Q \supset [[_]^{C_1, -, q_0}]^{E_1, r_0, -}$ where $p \in \{q_0, r_0\}$. We have already established that $C_1 \leq E_1$, therefore condition 4 is satisfied by this monitored expression.

Cases [R18–R20]. These cases are analogous to [R17].

Case [R25]. Straightforward consequence of Lemma 15.

Cases [R26–R28]. All these cases are proved with a straightforward application of the induction hypothesis. In the case of [R28], observe that structural congruence does not affect expressions, hence monitored expressions in particular. \square

Lemma 1. *If $\Gamma \vdash P$ where $\Gamma :: \omega$ and p is locally correct and sound in P and $P \rightarrow Q$, then p is sound also in Q .*

Proof. This is a particular instance of Lemma 16, since every unlimited Γ is trivially balanced. \square

Theorem 2. *If $\Gamma \vdash P$ where P is a user process and p is locally correct in P , then $P \rightarrow^* Q$ implies $\mathbf{blame} p \notin Q$.*

Proof. Observe that Γ must be unlimited for a user process has no (free) endpoints. Also, p is trivially sound in P for P is a user a process with no (busy) monitors in it. Suppose by contradiction that $P \rightarrow^* Q \supset \mathbf{blame} p$. The only reduction that generates a $\mathbf{blame} p$ term is [R7], hence there must exist Q' such that

$$P \rightarrow^* Q' \rightarrow^* Q \supset \mathbf{blame} p$$

and $Q' \supset v \ll^p \mathbf{false}$. This is absurd, because by Lemma 1 we deduce that p is sound in Q' . \square

C.4 Reduction of Example 2

Let P be the parallel composition of the two services q and r and the client p of Example 2. Tables 6 and 7 shows the relevant reduction steps of P .

D Supplement to Section 6

D.1 Single-Sided Monitoring

The following definition provides a characterization of the redundancy present on monitored expressions when wrapped by several monitors. In what follows we say that an expression is a *user expression* if it does not contain runtime syntax.

$$\begin{aligned}
P &\rightarrow (\nu c) \left(\langle \langle \text{let } x_1 = [c^+]^{c,q,p} \text{ in} \right. \\
&\quad \text{let } y_1, x_2 = \text{receive } x_1 \text{ in} \\
&\quad \text{let } y_2 = \text{send } 0 \ y_1 \text{ in} \\
&\quad \text{close } x_2; \\
&\quad \left. \text{close } y_2 \rangle \rangle_p \parallel \langle \langle \lambda x_1. \text{let } y_1 = \text{connect } b \text{ in} \right. \\
&\quad \quad \text{let } x_2 = \text{send } y_1 \ x_1 \text{ in} \\
&\quad \quad \left. \text{close } x_2 \rangle [c^-]^{\bar{c},p,q} \rangle_q \right) \\
&\rightarrow^* (\nu c) \left(\langle \langle \text{let } y_1, x_2 = \text{receive } [c^+]^{c,q,p} \text{ in} \right. \\
&\quad \text{let } y_2 = \text{send } 0 \ y_1 \text{ in} \\
&\quad \text{close } x_2; \\
&\quad \left. \text{close } y_2 \rangle \rangle_p \parallel \langle \langle \text{let } y_1 = \text{connect } b \text{ in} \right. \\
&\quad \quad \text{let } x_2 = \text{send } y_1 \ [c^-]^{\bar{c},p,q} \text{ in} \\
&\quad \quad \left. \text{close } x_2 \rangle \rangle_q \right) \\
&\rightarrow^* (\nu c)(\nu d) \left(\langle \langle \text{let } y_1, x_2 = \text{receive } [c^+]^{c,q,p} \text{ in} \right. \\
&\quad \text{let } y_2 = \text{send } 0 \ y_1 \text{ in} \\
&\quad \text{close } x_2; \\
&\quad \left. \text{close } y_2 \rangle \rangle_p \parallel \langle \langle \text{let } y_1 = [d^+]^{d,r,q} \text{ in} \right. \\
&\quad \quad \text{let } x_2 = \text{send } y_1 \ [c^-]^{\bar{c},p,q} \text{ in} \\
&\quad \quad \left. \text{close } x_2 \rangle \rangle_q \parallel \langle \langle \lambda y_1. \text{let } z, y_2 = \text{receive } y_1 \text{ in} \right. \\
&\quad \quad \quad \text{print } z^{-1}; \\
&\quad \quad \left. \text{close } y_2 \rangle [d^-]^{\bar{d},q,r} \rangle_r \right) \\
&\rightarrow^* (\nu c)(\nu d) \left(\langle \langle \text{let } y_1, x_2 = \text{receive } [c^+]^{c,q,p} \text{ in} \right. \\
&\quad \text{let } y_2 = \text{send } 0 \ y_1 \text{ in} \\
&\quad \text{close } x_2; \\
&\quad \left. \text{close } y_2 \rangle \rangle_p \parallel \langle \langle \text{let } x_2 = \text{send } [d^+]^{d,r,q} \ [c^-]^{\bar{c},p,q} \text{ in} \right. \\
&\quad \quad \left. \text{close } x_2 \rangle \rangle_q \parallel \langle \langle \text{let } z, y_2 = \text{receive } [d^-]^{\bar{d},q,r} \text{ in} \right. \\
&\quad \quad \quad \text{print } z^{-1}; \\
&\quad \quad \left. \text{close } y_2 \rangle \rangle_r \right)
\end{aligned}$$

Table 6: Reduction of the system described in Example 2.

$$\begin{aligned}
& \rightarrow^* (\nu c)(\nu d) \left(\left\langle \left\langle \left(\text{let } y_1, x_2 = ([[d^+]^{D,r,q}]^{E,q,p}]^{E,q,p}, [c^+]^{\text{end_c},q,p}) \text{ in} \right. \right. \right. \\
& \quad \text{let } y_2 = \text{send } 0 \ y_1 \text{ in} \\
& \quad \text{close } x_2; \\
& \quad \left. \left. \left. \text{close } y_2 \right) \right\rangle_p \right. \\
& \quad \left\| \left\langle \left(\text{let } x_2 = [c^-]^{\text{end_c},p,q} \text{ in} \right. \right. \right. \\
& \quad \quad \left. \left. \left. \text{close } x_2 \right) \right\rangle_q \right. \\
& \quad \left\| \left\langle \left(\text{let } z, y_2 = \text{receive } [d^-]^{\bar{D},q,r} \text{ in} \right. \right. \right. \\
& \quad \quad \left. \left. \left. \text{print } z^{-1}; \right. \right. \right. \\
& \quad \quad \left. \left. \left. \text{close } y_2 \right) \right\rangle_r \right. \\
& \left. \right) \\
& \rightarrow^* (\nu c)(\nu d) \left(\left\langle \left\langle \left(\text{let } y_2 = \text{send } 0 \ [[d^+]^{D,r,q}]^{E,q,p}]^{E,q,p} \text{ in} \right. \right. \right. \\
& \quad \quad \left. \left. \left. \text{close } [c^+]^{\text{end_c},q,p}; \right. \right. \right. \\
& \quad \quad \left. \left. \left. \text{close } y_2 \right) \right\rangle_p \right. \\
& \quad \left\| \left\langle \left(\text{close } [c^-]^{\text{end_c},p,q} \right) \right\rangle_q \right. \\
& \quad \left\| \left\langle \left(\text{let } z, y_2 = \text{receive } [d^-]^{\bar{D},q,r} \text{ in} \right. \right. \right. \\
& \quad \quad \left. \left. \left. \text{print } z^{-1}; \right. \right. \right. \\
& \quad \quad \left. \left. \left. \text{close } y_2 \right) \right\rangle_r \right. \\
& \left. \right) \\
& \rightarrow^* (\nu c)(\nu d) \left(\left\langle \left\langle \left(\text{let } y_2 = ([[d^+]^{\text{end_c},r,q}]^{\text{end_c},q,p}]^{\text{end_c},q,p} \text{ in} \right. \right. \right. \\
& \quad \quad \left. \left. \left. \text{close } [c^+]^{\text{end_c},q,p}; \right. \right. \right. \\
& \quad \quad \left. \left. \left. \text{close } y_2 \right) \right\rangle_p \right. \\
& \quad \left\| \left\langle \left(\text{close } [c^-]^{\text{end_c},p,q} \right) \right\rangle_q \right. \\
& \quad \left\| \left\langle \left(\text{let } z, y_2 = ([[[0]^{\text{any_c},p,q}]^{\text{any_c},p,q}]^{F,q,r}]^{F,q,r}, [d^-]^{\text{end_c},q,r}) \text{ in} \right. \right. \right. \\
& \quad \quad \left. \left. \left. \text{print } z^{-1}; \right. \right. \right. \\
& \quad \quad \left. \left. \left. \text{close } y_2 \right) \right\rangle_r \right. \\
& \left. \right) \\
& \rightarrow^* (\nu c)(\nu d) \left(\left\langle \left\langle \left(\text{close } [c^+]^{\text{end_c},q,p}; \right. \right. \right. \\
& \quad \quad \left. \left. \left. \text{close } [[d^+]^{\text{end_c},r,q}]^{\text{end_c},q,p}]^{\text{end_c},q,p} \right) \right\rangle_p \right. \\
& \quad \left\| \left\langle \left(\text{close } [c^-]^{\text{end_c},p,q} \right) \right\rangle_q \right. \\
& \quad \left\| \left\langle \left(\text{let } z, y_2 = ([[0]^{\text{F},q,r}]^{\text{F},q,r}, [d^-]^{\text{end_c},q,r}) \text{ in} \right. \right. \right. \\
& \quad \quad \left. \left. \left. \text{print } z^{-1}; \right. \right. \right. \\
& \quad \quad \left. \left. \left. \text{close } y_2 \right) \right\rangle_r \right. \\
& \left. \right) \\
& \rightarrow^* (\nu d) \left(\left\langle \left\langle \left(\text{close } [[d^+]^{\text{end_c},r,q}]^{\text{end_c},q,p}]^{\text{end_c},q,p} \right) \right\rangle_p \right. \\
& \quad \left\| \left\langle \left(\right) \right\rangle_q \right. \\
& \quad \left\| \left\langle \left(\text{let } z, y_2 = ([\text{blame } q]^{\text{F},q,r}, [d^-]^{\text{end_c},q,r}) \text{ in} \right. \right. \right. \\
& \quad \quad \left. \left. \left. \text{print } z^{-1}; \right. \right. \right. \\
& \quad \quad \left. \left. \left. \text{close } y_2 \right) \right\rangle_r \right. \\
& \left. \right)
\end{aligned}$$

Table 7: Reduction of the system described in Example 2 (cont.).

Definition 15. A monitored expression is well-stacked if it has one of the following forms:

- v is a user expression.
- $[\mathbf{blame} \ p]^{-\cdot-\cdot-}$,
- $[v \triangleleft^p e]^{\mathbf{flat_c} \ w.p.-}$ and v and w are user expressions and $e \Downarrow v'$ if and only if $wv \Downarrow v'$,
- $[a^t]^{-\cdot-\cdot-}$,
- $[[\bar{e}]^{c,p,q}]^{c,p,q}$ and \bar{e} is well-stacked.

The definition extends homomorphically over (unmonitored) expressions and processes.

The single-sided monitoring semantics avoids the redundancy in well-stacked monitored expressions by removing one duplicated monitor, as formalized by the following erasure function.

Definition 16 (Monitor Erasure). The erasure of monitors from well-stacked monitored expressions is defined as follows.

$$\begin{aligned}
\text{erase}(v) &= v \\
\text{erase}([\mathbf{blame} \ p]^{-\cdot-\cdot-}) &= \mathbf{blame} \ p \\
\text{erase}([v \triangleleft^p e]^{-\cdot-\cdot-}) &= v \triangleleft^p e \\
\text{erase}([a^-]^{-\cdot-\cdot-}) &= a^- \\
\text{erase}([a^+]^{e,p,q}) &= [a^+]^{e,p,q} \\
\text{erase}([[\bar{e}]^{c,p,q}]^{c,p,q}) &= [\text{erase}(\bar{e})]^{c,p,q}
\end{aligned}$$

We also write $\text{erase}(_)$ for the homomorphic extension to expressions and processes.

We now show that substitution and replacement within contexts of well-stacked expressions produce well-stacked expressions.

Lemma 17. *Let e and v be well-stacked expressions. Then,*

1. $e\{v/x\}$ is well-stacked.
2. $\text{erase}(e\{v/x\}) = \text{erase}(e)\{\text{erase}(v)/x\}$.

Proof. By induction on the structure of e . □

Lemma 18. *If e and e' and $\mathcal{E}[e]$ are well-stacked, then $\mathcal{E}[e']$ is well-stacked.*

Proof. By induction on the structure of \mathcal{E} . The only interesting case is $\mathcal{E} = [\mathcal{E}']^{e_1,\sigma}$. Since, $\mathcal{E}[e]$ and e are well-stacked, by case analysis we conclude that the only possibility is $\mathcal{E}' = [\mathcal{E}']^{e_1,\sigma}$ with $e_1 = c$. Then, it is immediate that $\mathcal{E}[e']$ is well-stacked. □

The following auxiliary result shows that a sub-expression obtained after applying the erasure function to a well-stacked expression corresponds to the result of applying the erasure function to a well-stacked sub-expression of the original expression.

Lemma 19. *Let e be a well-stacked expression. If $\text{erase}(e) = \mathcal{E}[e']$ then there exist \mathcal{E}'' and e'' well-stacked such that $e = \mathcal{E}''[e'']$ and $e' = \text{erase}(e'')$.*

Proof. We proceed by induction on e .

Case $e = v$. Then, $\text{erase}(e) = e$. It suffices to take $\mathcal{E} = \mathcal{E}'$ and $e' = e''$.

Case $e = x$. It follows analogously to the previous case.

Case $e = e_1 e_2$. Then, $\text{erase}(e) = \text{erase}(e_1)\text{erase}(e_2)$. Therefore, either $\text{erase}(e) = \mathcal{E}''[e'_1]\text{erase}(e_2)$ or $\text{erase}(e) = \text{erase}(e_1)\mathcal{E}''[e'_2]$. Both cases follow by using induction hypothesis.

Cases $e = \mathbf{let} \ x, y = e_1 \ \mathbf{in} \ e_2$ or $e = \mathbf{case} \ e \ \mathbf{of} \ e_1 \mid e_2$. They follow as in the previous case.

Case $e = [e_1]^{e_2,p,q}$. Since e_2 is a user expression, $\text{erase}(e_2) = e_2$. We proceed by case analysis on e_1 .

- $e_1 = \mathbf{blame} p$. Then, $e = [\mathbf{blame} p]^{e_2, p, q}$ and $\mathbf{erase}(e) = \mathbf{blame} p$. Hence, $\mathcal{E} = []$ and $e' = e_1$. Take $\mathcal{E}' = []$ and $e'' = e$ and note that $\mathbf{erase}(e'') = e_1$.
- $e_1 = v \triangleleft^p e_3$ where $e_2 = \mathbf{flat_c} w$. Then $e = [v \triangleleft^p e_3]^{e_2, p, q}$ and $\mathbf{erase}(e) = v \triangleleft^p e_3$. If $\mathcal{E} = []$, then the proof follows as in the previous case. Otherwise, $\mathcal{E} = v \triangleleft^p \mathcal{E}''$. Take $\mathcal{E}' = [\mathcal{E}]^{e_2, p, q}$ and $e'' = e'$. Note that e'' is well-stacked because e_1 is well-stacked.
- $e_1 = a^+$. Then, $e = [a^+]^{e_2, p, q}$ and $\mathbf{erase}(e) = e$. Take $\mathcal{E}' = []$ and $e'' = e$.
- $e_1 = a^-$. Then, $e = [a^-]^{e_2, p, q}$ and $\mathbf{erase}(e) = e_1$. Take $\mathcal{E}' = []$ and $e'' = e$.
- $e_1 = [\bar{e}]^{c, p, q}$. Then, $e = [[\bar{e}]^{c, p, q}]^{c, p, q}$ and $\mathbf{erase}(e) = [\mathbf{erase}(\bar{e})]^{c, p, q}$. If $\mathcal{E} = []$, then the proof follows as in the previous cases. Otherwise, $\mathcal{E} = [\mathcal{E}''']^{c, p, q}$ and $\mathbf{erase}(e) = [\mathcal{E}'''[e']]^{c, p, q}$ with $\mathbf{erase}(\bar{e}) = \mathcal{E}'''[e']$. By induction hypothesis, $\bar{e} = \mathcal{E}'''[e'']$ and $e' = \mathbf{erase}(e'')$ with e'' well-stacked. The proof is completed by taking $\mathcal{E}'' = [[\mathcal{E}''']]^{c, p, q}$.

Case $e = v \triangleleft^p e_1$. It is not possible because e is not a well-stacked expression.

Case $e = \mathbf{blame} p$. It is not possible because e is not a well-stacked expression. \square

The following two results show that a reduction step of the one-side monitoring semantics is simulated by possibly many reduction steps of the big-step semantics. Several reductions are needed when evaluating redundant monitors.

Lemma 20. *If $\Gamma \vdash e : t$ and e is well-stacked and $\mathbf{erase}(e) \mapsto e'$, then $e \rightarrow^* e''$ and $e' = \mathbf{erase}(e'')$.*

Proof. We proceed by induction on the derivation of $\mathbf{erase}(e) \mapsto e'$.

Case [R1]. Then, $\mathbf{erase}(e) = (\lambda x. e_0)v \mapsto e_0\{v/x\} = e'$. Moreover, $e = (\lambda x. e_1)w$ with $e_0 = \mathbf{erase}(e_1)$ and $v = \mathbf{erase}(w)$. Then, $e \rightarrow e_1\{w/x\} = e''$. Note that e_1 and w are well-stacked because e is well-stacked. Then, $e' = \mathbf{erase}(e'')$ because of Lemma 17.

Cases [R2] – [R4]. Follow analogously to [R1].

Case [R5]. Then, $\mathbf{erase}(e) = [v]^{e_0, p, q}$ and $e_0 = \mathbf{flat_c} w$. Therefore, $\mathbf{erase}(e) \mapsto v \triangleleft^p wv = e'$. Since e is well-stacked,

$$e = [[v]^{e_0, p, q}]^{e_0, p, q}$$

Hence, $e \rightarrow [v \triangleleft^p wv]^{e_0, p, q} = e''$. Since $e_0 = \mathbf{flat_c} w$, it trivially holds that $wv \Downarrow v'$ if and only if $wv \Downarrow v'$. Hence, e'' is well-stacked and $e' = \mathbf{erase}(e'')$.

Case [R6]. $\mathbf{erase}(e) = v \triangleleft^p \mathbf{true} \mapsto v = e'$. Since e is well-stacked, $e = [v \triangleleft^p \mathbf{true}]^{e_0, p, q}$ and $e_0 = \mathbf{flat_c} w$ and $wv \Downarrow \mathbf{true}$. Then,

$$e \rightarrow [v]^{e_0, p, q} \rightarrow v \triangleleft^p wv \rightarrow^* v = e''.$$

Hence, $e' = \mathbf{erase}(e'')$.

Case [R7]. $\mathbf{erase}(e) = v \triangleleft^p \mathbf{false} \mapsto \mathbf{blame} p = e'$. Since e is well-stacked,

$$e = [v \triangleleft^p \mathbf{false}]^{e_0, p, q} \rightarrow [\mathbf{blame} p]^{e_0, p, q} = e''$$

Hence, $e' = \mathbf{erase}(e'')$.

Cases [R8] – [R14]. Follow straightforwardly. \square

Lemma 21. *Let $\Gamma \vdash P$ where Γ is balanced. If P is balanced and well-stacked and $\mathbf{erase}(P) \mapsto Q$, then $P \rightarrow^* Q'$ and $Q = \mathbf{erase}(Q')$.*

Proof. We proceed by induction on the derivation of $\text{erase}(P) \mapsto Q$.

Case [R21]. $\text{erase}(P) = \langle \mathcal{E}[\text{connect } a] \rangle_p \parallel a \Leftarrow_q^c v \mapsto Q$ with

$$Q = (\nu c)(\langle \mathcal{E}[[c^+]^{c,q,p}] \rangle_p \parallel \langle (v \ c^-) \rangle_q) \parallel a \Leftarrow_q^c v$$

Moreover, $P = \langle \mathcal{E}'[\text{connect } a] \rangle_p \parallel a \Leftarrow_q^c v$ and $P \rightarrow Q'$ with

$$Q' = (\nu c)(\langle \mathcal{E}'[[c^+]^{c,q,p}] \rangle_p \parallel \langle (v \ [c^-]^{\text{dual } c,p,q}) \rangle_q) \parallel a \Leftarrow_q^c v$$

Since P is well-stacked, $\mathcal{E}' \neq \mathcal{E}''[[\]^{-\dots}]$. Moreover, service definitions do not have runtime syntax and hence v is also well-stacked. Then, Q' is well-stacked and $Q = \text{erase}(Q')$.

Case [R17]. $\text{erase}(P) \mapsto Q$ with

$$\begin{aligned} \text{erase}(P) &= \langle \mathcal{E}_1[\text{send } v \ [a^t]^{!c;D,\vec{\sigma}}] \rangle_p \parallel \langle \mathcal{E}_2[\text{receive } [a^{\bar{t}}]^{?E;F,\vec{\rho}}] \rangle_q \\ Q &= \langle \mathcal{E}_1[[a^t]^{D,\vec{\sigma}}] \rangle_p \parallel \langle \mathcal{E}_2[(\llbracket v \rrbracket^{c,-\sigma})^{E,\vec{\rho}}, [a^{\bar{t}}]^{F,\vec{\rho}}] \rangle_q \end{aligned}$$

Moreover, $P \rightarrow Q'$ where

$$\begin{aligned} P &= \langle \mathcal{E}'_1[\text{send } v \ [a^t]^{!c';D',\vec{\sigma}'}] \rangle_p \parallel \langle \mathcal{E}'_2[\text{receive } [a^{\bar{t}}]^{?E';F',\vec{\rho}'}] \rangle_q \\ Q' &= \langle \mathcal{E}'_1[[a^t]^{D',\vec{\sigma}'}] \rangle_p \parallel \langle \mathcal{E}'_2[(\llbracket v \rrbracket^{c',-\sigma'})^{E',\vec{\rho}'}, [a^{\bar{t}}]^{F',\vec{\rho}'}] \rangle_q \end{aligned}$$

Since P is well-stacked, $\mathcal{E}'_i \neq \mathcal{E}''_i[[\]^{-\dots}]$ for $i \in \{1, 2\}$. Hence, Q' is well-stacked if $[a^t]^{D',\vec{\sigma}'}$ and $[(\llbracket v \rrbracket^{c',-\sigma'})^{E',\vec{\rho}'}]$ and $[a^{\bar{t}}]^{F',\vec{\rho}'}$ are well-stacked. The subterms $[a^t]^{D',\vec{\sigma}'}$ and $[a^{\bar{t}}]^{F',\vec{\rho}'}$ are well-stacked because $[a^t]^{!c';D',\vec{\sigma}'}$ and $[a^{\bar{t}}]^{?E';F',\vec{\rho}'}$ are so. It remains to show that $[(\llbracket v \rrbracket^{c',-\sigma'})^{E',\vec{\rho}'}]$ is well-stacked. Since P is well-stacked, v is well-stacked. By Lemma 13, $!c'_0;D'_0 \bowtie ?E'_0;F'_0$ and $-\sigma_0 = \rho_0$. By duality, $c_0 = e_0$. Hence, $[(\llbracket v \rrbracket^{c',-\sigma'})^{E',\vec{\rho}'}]$ is well-stacked.

Case [R25]. Then, $\text{erase}(P) = \langle \text{erase}(e) \rangle_p$ and $P \mapsto \langle e' \rangle_p = Q$ with $e \mapsto e'$. By Lemma 20, $e \mapsto e'$ implies $e \rightarrow^* e''$ and $e' = \text{erase}(e'')$. Hence, $P = \langle e \rangle_p \rightarrow^* \langle e'' \rangle_p = Q'$ and $Q = \text{erase}(Q')$.

Case [R26]. Then, $P = R_1 \parallel R_2$ and $\text{erase}(P) = \text{erase}(R_1) \parallel \text{erase}(R_2)$ and $\text{erase}(R_1) \mapsto R'_1$ and $Q = R'_1 \parallel \text{erase}(R_2)$. By induction hypothesis on $\text{erase}(R_1) \mapsto R'_1$, we have $R_1 \rightarrow^* R'_1$ and $R'_1 = \text{erase}(R''_1)$. Therefore, $P \rightarrow^* R'_1 \parallel R_2 = Q'$ and $Q = \text{erase}(Q')$. \square

Lemma 22. *If P is a user process and $P \mapsto^* Q$, then there exists Q' such that $P \rightarrow^* Q'$ and $Q = \text{erase}(Q')$.*

Proof. The proof follows by induction on the length of the derivation $P \mapsto^* Q$. For the base case, *i.e.*, $P = Q$, the proof follows by noting that P is a user process and hence it is well-stacked. Then, take $Q' = P$. For the induction step we consider $P \mapsto^n R \mapsto Q$. By induction hypothesis, there exists R' such that $P \rightarrow^* R'$ and $R = \text{erase}(R')$. Note that R' is well-stacked because $\text{erase}(R')$ is defined. Additionally, R' is balanced, because of Lemma 13. By Lemma 21, $R \mapsto Q$ implies that there exists Q' such that $P \rightarrow^* Q'$ and $Q = \text{erase}(Q')$. \square

Proposition 2. *Let \mapsto be the relation defined by the rules in Tables 2 and 3 except that [R15] is replaced by [R21]. If P is a user process and $P \mapsto^* \supset \text{blame } p$, then $P \rightarrow^* \supset \text{blame } p$.*

Proof. By Lemma 22, $P \mapsto^* Q$ implies $P \rightarrow^* R$ such that $\text{erase}(R) = Q$. By definition of $\text{erase}(_)$, $Q \supset \text{blame } p$ implies $R \supset \text{blame } p$. \square

[R22]	<code>send</code>	v	$[\varepsilon]^{!C;D,\sigma}$	\rightsquigarrow	$[\text{send } [v]^{C,-\sigma} \varepsilon]^{D,\sigma}$
[R23]	<code>receive</code>	$[\varepsilon]^{?C;D,\sigma}$	\rightsquigarrow	<code>let</code> $x, y =$ <code>receive</code> ε <code>in</code> $([x]^{C,\sigma}, [y]^{D,\sigma})$	
[R32]	<code>send</code>	v	$[\varepsilon]^{!C.w,\sigma}$	\rightsquigarrow	$[\text{send } [v]^{C,-\sigma} \varepsilon]^{wv,\sigma}$
[R33]	<code>receive</code>	$[\varepsilon]^{?C.w,\sigma}$	\rightsquigarrow	<code>let</code> $x, y =$ <code>receive</code> ε <code>in</code> $([x]^{C,\sigma}, [y]^{wx,\sigma})$	
[R34]	<code>left</code>	$[\varepsilon]^{!C.D;E,\sigma}$	\rightsquigarrow	$(\lambda_. \text{left } \varepsilon) [\text{true}]^{C,-\sigma}]^{D,\sigma}$	
[R35]	<code>right</code>	$[\varepsilon]^{!C.D;E,\sigma}$	\rightsquigarrow	$(\lambda_. \text{right } \varepsilon) [\text{false}]^{C,-\sigma}]^{E,\sigma}$	
[R36]	<code>branch</code>	$[\varepsilon]^{?C.D;E,\sigma}$	\rightsquigarrow	<code>case branch</code> ε <code>of</code> $\lambda x. (\lambda_. \text{inl } [x]^{D,\sigma}) [\text{true}]^{C,\sigma}$ $ \lambda x. (\lambda_. \text{inr } [x]^{E,\sigma}) [\text{false}]^{C,\sigma}$	
[R37]	<code>close</code>	$[\varepsilon]^{end_c,\sigma}$	\rightsquigarrow	<code>close</code> ε	

Table 8: Small-step reduction of expressions.

[R38]	$(\nu a)(\langle \mathcal{E}[\text{close } a^+] \rangle_p \parallel \langle \mathcal{E}'[\text{close } a^-] \rangle_q) \rightsquigarrow \langle \mathcal{E}[\text{close } a^+] \rangle_p \parallel \langle \mathcal{E}'[\text{close } a^-] \rangle_q$
[R24]	$\langle \mathcal{E}[\text{send } v a^t] \rangle_p \parallel \langle \mathcal{E}'[\text{receive } a^t] \rangle_q \rightsquigarrow \langle \mathcal{E}[a^t] \rangle_p \parallel \langle \mathcal{E}'[v, a^t] \rangle_q$
[R39]	$\langle \mathcal{E}[\text{left } a^t] \rangle_p \parallel \langle \mathcal{E}'[\text{branch } a^t] \rangle_q \rightsquigarrow \langle \mathcal{E}[a^t] \rangle_p \parallel \langle \mathcal{E}'[\text{inl } a^t] \rangle_q$
[R40]	$\langle \mathcal{E}[\text{right } a^t] \rangle_p \parallel \langle \mathcal{E}'[\text{branch } a^t] \rangle_q \rightsquigarrow \langle \mathcal{E}[a^t] \rangle_p \parallel \langle \mathcal{E}'[\text{inr } a^t] \rangle_q$

Table 9: Small-step reduction of processes.

D.2 Small-Step Monitoring

We start by reporting the reduction rules for the small-step semantics of our implementation. We added the rules in Tables 8 and 9 to the original rules for expressions in Table 2. In this section we show that the small-step semantics assigns the same blames as the big-step semantics in configurations in which two threads are about to communicate.

The following auxiliary result states that the two semantics coincide when evaluating monitored expressions. This is an immediate consequence of the fact that the small-step semantics does not change the evaluation rules for monitored expressions.

Lemma 23. *Let \bar{e} be a monitored expression. Then, $\bar{e} \rightsquigarrow \bar{e}'$ implies $\bar{e} \rightarrow \bar{e}'$.*

Proof. By straightforward induction on the structure of \bar{e} . □

The small-step semantics applies monitors one at a time. Additionally, flat contracts are evaluated by the sender before the communication takes place. The next two results characterize the configurations that can be reached when evaluating `send` $v \varepsilon$, depending on whether the involved contracts are flat or not.

Lemma 24. *Let $P \stackrel{\text{def}}{=} \text{send } v [a^t]^{!C;D,\sigma}$. If $C_i \neq \text{flat}_c$ $_$, then $P \rightsquigarrow^* [\text{send } [v]^{C,-\sigma} a^t]^{D,\sigma}$.*

Proof. By induction on the number of stacked monitors. □

Lemma 25. *Let $P \stackrel{\text{def}}{=} \text{send } v [a^t]^{!C;D,\sigma}$. If $C_i = \text{flat}_c$ w_i , then either*

- $P \rightsquigarrow^* \text{send } v a^t$ and $v \in w_i$ for all i .
- $P \rightsquigarrow^* \supset \text{blame } p$ and there exists i such that $v \notin w_i$ and $\sigma_i = p, _$ and $v \in w_j$ for all $j < i$.

Proof. By induction on the number of stacked monitors. □

Analogously, monitors on the receiver side are applied incrementally. We first characterize the configurations that are reachable from the expression `receive` ε . We use the auxiliary notation $\text{let}^{\vec{C};\vec{D},\vec{\sigma}} e$ defined below.

$$\begin{aligned}\text{let}^{\vec{C};\vec{D},\vec{\sigma}} e &= \text{let } x, y = e \text{ in } ([x]^{C,\sigma}, [y]^{D,\sigma}) \\ \text{let}^{\vec{C};\vec{D},\vec{\sigma}} e &= \text{let } x, y = \text{let}^{\vec{E};\vec{F},\vec{\varrho}} e \text{ in } ([x]^{C_n,\sigma_n}, [y]^{D_n,\sigma_n})\end{aligned}$$

with $\vec{\varrho} = \sigma_0, \dots, \sigma_{n-1}$ and $\vec{E} = C_0, \dots, C_{n-1}$, and $\vec{F} = D_0, \dots, D_{n-1}$.

Lemma 26. `receive` $[a^t]^{\vec{C};\vec{D},\vec{\sigma}} \rightsquigarrow^* \text{let}^{\vec{C};\vec{D},\vec{\sigma}} \text{receive } a^t$.

Proof. By induction on the number of stacked monitors. \square

Now we show the possible reductions for the receiver soon after the communication takes place. Also, the reachable configurations depend on whether contracts are flat or not.

Lemma 27. Let $P \stackrel{\text{def}}{=} \text{let}^{\vec{C};\vec{D},\vec{\sigma}} (v, w)$. If $C_i \neq \text{flat_c } _$, then $P \rightsquigarrow^* ([v]^{\vec{C},\vec{\sigma}}, [w]^{\vec{D},\vec{\sigma}})$.

Proof. By induction on the number of stacked monitors. \square

Lemma 28. Let $P \stackrel{\text{def}}{=} \text{let}^{\vec{C};\vec{D},\vec{\sigma}} (v, w)$. If $C_i = \text{flat_c } w_i$, then either

- $P \rightsquigarrow^* (v, [w]^{\vec{D},\vec{\sigma}})$ and $v \in w_i$ for all i .
- $P \rightsquigarrow^* \supset \text{blame } p$ and there exists i such that $v \notin w_i$ and $\sigma_i = p, _$ and $v \in w_j$ for all $j < i$.

Proof. By induction on the number of stacked monitors. \square

Finally, we state the following auxiliary result for the reductions of a monitored value in the big-step semantics.

Lemma 29. Let $P \stackrel{\text{def}}{=} [v]^{\vec{C},\vec{\sigma}}$. If $C_i = \text{flat_c } w_i$, then either

- $P \rightarrow^* v$ and $v \in w_i$ for all i .
- $P \rightarrow^* \supset \text{blame } p$ and there exists i such that $v \notin w_i$ and $\sigma_i = p, _$ and $v \in w_j$ for all $j < i$.

Proof. By induction on the number of stacked monitors. \square

Proposition 3. If

$$\begin{aligned}P &\stackrel{\text{def}}{=} \langle \mathcal{E}[\text{send } v [a^t]^{\vec{C};\vec{D},\vec{\sigma}}] \rangle_p \parallel \langle \mathcal{E}'[\text{receive } [a^t]^{\vec{E};\vec{F},\vec{\varrho}}] \rangle_q \\ Q_1 &\stackrel{\text{def}}{=} \langle \mathcal{E}[[a^t]^{\vec{D},\vec{\sigma}}] \rangle_p \parallel \langle \mathcal{E}'[(v, [a^t]^{\vec{F},\vec{\varrho}})] \rangle_q \\ Q_2 &\stackrel{\text{def}}{=} \langle \mathcal{E}[[a^t]^{\vec{D},\vec{\sigma}}] \rangle_p \parallel \langle \mathcal{E}'[(\llbracket [v]^{\vec{C},\vec{\sigma}} \rrbracket^{\vec{E},\vec{\varrho}}, [a^t]^{\vec{F},\vec{\varrho}}) \rrbracket] \rangle_q\end{aligned}$$

then either

1. $P \rightsquigarrow^* Q$ and $P \rightarrow^* Q$ for some $Q \in \{Q_1, Q_2\}$, or
2. $P \rightsquigarrow^* \supset \text{blame } r$ and $Q \rightarrow^* \supset \text{blame } r$ for some r , or

Proof. There are two cases:

- $C_i \neq \text{flat_c } w_i$. By Lemma 24,

$$P \rightsquigarrow^* \langle \mathcal{E}[\llbracket \text{send } [v]^{\vec{C},\vec{\sigma}} a^t \rrbracket^{\vec{D},\vec{\sigma}}] \rangle_p \parallel \langle \mathcal{E}'[\text{receive } [a^t]^{\vec{E};\vec{F},\vec{\varrho}}] \rangle_q$$

By Lemma 26

$$P \rightsquigarrow^* \langle \mathcal{E}[\llbracket \text{send } [v]^{\vec{C},\vec{\sigma}} a^t \rrbracket^{\vec{D},\vec{\sigma}}] \rangle_p \parallel \langle \mathcal{E}'[\llbracket \text{let}^{\vec{E};\vec{F},\vec{\sigma}} \text{receive } a^t \rrbracket] \rangle_q$$

By rule [R24],

$$P \rightsquigarrow^* \langle \mathcal{E}[[a^t]^{\overrightarrow{D}, \overrightarrow{\sigma}}] \rangle_p \parallel \langle \mathcal{E}'[\mathbf{let}^{\overrightarrow{?E; F}, \overrightarrow{\sigma}} ([v]^{\overleftarrow{C}, \overleftarrow{-\sigma}}, a^t)] \rangle_q$$

By Lemma 27

$$P \rightsquigarrow^* \langle \mathcal{E}[[a^t]^{\overrightarrow{D}, \overrightarrow{\sigma}}] \rangle_p \parallel \langle \mathcal{E}'[(\llbracket [v]^{\overleftarrow{C}, \overleftarrow{-\sigma}} \rrbracket^{\overrightarrow{E}, \overrightarrow{\varrho}}, [a^t]^{\overrightarrow{F}, \overrightarrow{\varrho}})] \rangle_q = Q_2$$

Moreover, $P \rightarrow Q_2$ by rule [R17].

- $C_i = \mathbf{flat_c} \ w_i$. By Lemma 25,

– If $v \in w_i$ for all i .

$$P \rightsquigarrow^* \langle \mathcal{E}[[\mathbf{send} \ v \ a^t]^{\overrightarrow{D}, \overrightarrow{\sigma}}] \rangle_p \parallel \langle \mathcal{E}'[\mathbf{receive} [a^t]^{\overrightarrow{?E; F}, \overrightarrow{\varrho}}] \rangle_q$$

By Lemma 26

$$P \rightsquigarrow^* \langle \mathcal{E}[[\mathbf{send} \ v \ a^t]^{\overrightarrow{D}, \overrightarrow{\sigma}}] \rangle_p \parallel \langle \mathcal{E}'[\mathbf{let}^{\overrightarrow{?E; F}, \overrightarrow{\sigma}} \ \mathbf{receive} \ a^t] \rangle_q$$

By rule [R24],

$$P \rightsquigarrow^* \langle \mathcal{E}[[a^t]^{\overrightarrow{D}, \overrightarrow{\sigma}}] \rangle_p \parallel \langle \mathcal{E}'[\mathbf{let}^{\overrightarrow{?E; F}, \overrightarrow{\sigma}} (v, a^t)] \rangle_q = P_1$$

By rule [R17].

$$P \rightarrow \langle \mathcal{E}[[a^t]^{\overrightarrow{D}, \overrightarrow{\sigma}}] \rangle_p \parallel \langle \mathcal{E}'[(\llbracket [v]^{\overleftarrow{C}, \overleftarrow{-\sigma}} \rrbracket^{\overrightarrow{E}, \overrightarrow{\varrho}}, [a^t]^{\overrightarrow{F}, \overrightarrow{\varrho}})] \rangle_q$$

By Lemma 29,

$$P \rightarrow \langle \mathcal{E}[[a^t]^{\overrightarrow{D}, \overrightarrow{\sigma}}] \rangle_p \parallel \langle \mathcal{E}'[(\llbracket [v]^{\overrightarrow{E}, \overrightarrow{\varrho}} \rrbracket, [a^t]^{\overrightarrow{F}, \overrightarrow{\varrho}})] \rangle_q = P_2$$

By Lemmata 28 and 29, we conclude that one of the following holds

- * $P_1 \rightsquigarrow^* Q_1$ and $P_2 \rightarrow^* Q_1$;
 - * $P_1 \rightsquigarrow^* \supset \mathbf{blame} \ p$ and $P_2 \rightarrow^* \supset \mathbf{blame} \ p$;
- There exists i such that $v \notin w_i$ and $\sigma_i = p, _$ and $v \in w_j$ for all $j < i$. Then, $P \rightsquigarrow^* \supset \mathbf{blame} \ p$. It follows as in the previous case that $P \rightarrow^* \supset \mathbf{blame} \ p$. \square