

AperTO - Archivio Istituzionale Open Access dell'Università di Torino

Exception Handling for Copyless Messaging

This is the author's manuscript

Original Citation:

Availability:

This version is available <http://hdl.handle.net/2318/144527> since 2016-01-08T11:45:32Z

Published version:

DOI:10.1016/j.scico.2013.05.001

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)



UNIVERSITÀ DEGLI STUDI DI TORINO

This Accepted Author Manuscript (AAM) is copyrighted and published by Elsevier. It is posted here by agreement between Elsevier and the University of Turin. Changes resulting from the publishing process - such as editing, corrections, structural formatting, and other quality control mechanisms - may not be reflected in this version of the text. The definitive version of the text was subsequently published in *SCIENCE OF COMPUTER PROGRAMMING*, 84, 2014, 10.1016/j.scico.2013.05.001.

You may download, copy and otherwise use the AAM for non-commercial purposes provided that your license is limited by the following restrictions:

- (1) You may use this AAM for non-commercial purposes only under the terms of the CC-BY-NC-ND license.
- (2) The integrity of the work and identification of the author, copyright owner, and publisher must be preserved in any copy.
- (3) You must attribute this AAM in the following format: Creative Commons BY-NC-ND license (<http://creativecommons.org/licenses/by-nc-nd/4.0/deed.en>), 10.1016/j.scico.2013.05.001

The definitive version is available at:

<http://linkinghub.elsevier.com/retrieve/pii/S0167642313001214>

Exception Handling for Copyless Messaging

Svetlana Jakšić

*Univerzitet u Novom Sadu, Fakultet tehničkih nauka
Trg Dositeja Obradovića 6, 21000 Novi Sad, Serbia*

email: sjaksic@uns.ac.rs

Luca Padovani

*Università di Torino, Dipartimento di Informatica
Corso Svizzera 185, 10149 Torino, Italy*

email: padovani@di.unito.it

Abstract

Copyless messaging is a communication paradigm in which only pointers to messages are exchanged between sender and receiver processes. Because of its intrinsically low overhead, copyless messaging is suitable for the efficient implementation of communication-intensive software systems where processes have access to a shared address space. Unfortunately, the very nature of the paradigm fosters the proliferation of programming errors due to the explicit use of pointers and to the sharing of data. In this paper we study a type discipline for copyless messaging that, together with some minimal support from the runtime system, is able to guarantee the absence of communication errors, memory faults, and memory leaks in presence of exceptions. To formalize the semantics of processes we draw inspiration from software transactional memories: in our case a transaction is a process that is meant to accomplish some exchange of messages and that should either be executed completely, or should have no observable effect if aborted by an exception.

Keywords: Copyless message passing, Session types, Exception handling, Memory leak prevention

1. Introduction

Communication has become a central aspect of all modern software systems, which range from distributed processes connected by wide area networks down

to collections of threads running on different cores within the same processing unit. In all these scenarios, message passing is a flexible paradigm that allows autonomous entities to exchange information and to synchronize with each other. The term “message passing” seems to suggest a paradigm where messages *move* from one entity to another, although more often than not messages are in fact *copied* during communication. While this is inevitable in a distributed setting, the availability of a shared address space makes it possible to implement a *copyless* form of message passing, whereby only pointers to messages are exchanged.

The Singularity Operating System (Singularity OS) [1, 2] is a notable example of system that heavily relies on the copyless paradigm. In Singularity OS, processes have access to a shared region called the *exchange heap*, inter-process communication solely occurs by means of message passing over channels allocated on the exchange heap, and messages are themselves pointers to the exchange heap. As detailed by Hunt et al. [1], Hunt and Larus [2], Finley [3], it is not practical to automatically garbage collect objects on the exchange heap, which therefore must be explicitly managed by processes.

The copyless paradigm has obvious performance advantages over more conventional forms of message passing. At the same time, it fosters the proliferation of subtle programming errors arising from the explicit management of objects and the sharing of data. For this reason the designers of Singularity OS have equipped Sing[#], the programming language used for the development of Singularity OS, with explicit constructs, types, and static analysis techniques to assist programmers in writing code that is free from a number of programming errors, including: *memory faults*, namely the access to unallocated/deallocated objects in the heap; *memory leaks*, that is the accumulation of unreachable allocated objects in the heap; communication errors, which could cause the abnormal termination of processes and trigger the previous kinds of errors.

Some aspects of Sing[#] have already been formalized and studied by Fährdrich et al. [4], Stengel and Bultan [5], Villard et al. [6], Bono and Padovani [7]. In particular, in [7] it was shown that Sing[#] *channel contracts* can be conveniently represented as a variant of session types [8, 9], and that the information given by session types along with a linear type discipline can prevent memory leaks, memory faults, and communication errors. In the present paper we focus on *exceptions* and *exception handling*. The interest in our research stems from the observation that copyless messaging and exceptions are clearly at odds with each other: on the one hand, copyless messaging requires a very disciplined and controlled access to memory; on the other hand, exceptions are in general unpredictable and disrupt the normal control flow of programs. Consequently, and perhaps not sur-

prisingly, these two aspects can be reconciled only with some native support from the runtime system. Below is a summary of the contributions of this research:

- we formalize a calculus of processes that communicate through the copyless paradigm and that can throw exceptions;
- we develop a type system for preventing the aforementioned errors even in the presence of exceptions, if suitable exception handlers are provided;
- we show how to take advantage of the invariants guaranteed by the type system in order to reduce the cost of exception handling.

Origin of the material. An early version of this paper appeared in the proceedings of the PPDP 2012 conference [10]. The present paper corresponds to a revised version of [10] where we provide detailed proofs of all the results and we generalize the process language (with multiple exceptions and exception propagation) as well as the type system (with subtyping). In particular, we show that the exception annotations within types induce an original subtyping relation that sheds light on the differences between exceptions and regular messages.

Structure of the paper. In Section 2 we illustrate the problem we are addressing and informally sketch our solution in terms of types and a revised exception handling construct. In Section 3 we formally define the syntax and the semantics of a language of processes to model Sing[#] programs. The section ends with the definition of *well-behaved processes*, namely of those processes in which memory faults, memory leaks, and communication errors do not occur. Section 4 develops a type system for the process language presented in Section 3 and shows its soundness (well-typed processes are well behaved). Section 5 provides an overview of related work and discusses similarities and differences between our approach and similar ones. Section 6 concludes with a brief summary of the work and hints at further relaxations of the typing rules, in light of the common pattern usage of exception handling mechanisms as found in the source code of Singularity OS. For the sake of readability, the proofs of subject reduction and of type soundness, which are fairly long and require a number of auxiliary lemmas, have been moved to Appendix A and B, respectively.

2. Motivating Example

To introduce the context in which we operate and the kind of problems we have to face we take a look at a real fragment of Singularity OS. In the discussion that

follows it is useful to keep in mind that Singularity channels consist of pairs of related *endpoints*, called the *peers* of the channel. Each endpoint is associated with an unbounded queue containing the messages sent to that endpoint from its peer. Communication is therefore asynchronous and send operations are nonblocking.

Figure 1 shows a `Sing#` function that computes the name for a newly allocated RAM disk.¹ The function has two output parameters, the computed disk name and the endpoint that links the disk to the `DirectoryService` (abbreviated `DS` in the code) which is part of the file system manager. The function begins by retrieving an endpoint `ns` for communicating with `DirectoryService` (line 3). Then the function repeatedly creates a new channel, represented as the peer endpoints `imp` and `exp` which are the output parameters of the `NewChannel` method (lines 6–8), computes a new disk name (line 9), and tries to register the chosen name along with `imp` to `DirectoryService` through `ns` (line 10). The `switch receive` construct (lines 11–24) is used to receive messages and to dispatch control to various cases depending on the type of message that is received. Each `case` block specifies the endpoint from which a message is expected and the tag of the message. In this example, one of two kinds of messages are expected from the `ns` endpoint: either an `AckRegister` message (lines 12–15) or a `NakRegister` message (lines 16–23). In the first case the registration is successful (line 12), so the output parameter `expService` is properly initialized and the function terminates correctly (line 15). In the second case the registration is unsuccessful (line 16), hence a new registration is attempted if the error is recoverable (lines 17–18), otherwise an exception is thrown to abort the execution of the function (line 20). The main loop (lines 5–25) is protected within a `try` block with a `finally` clause that is executed regardless of whether the function terminates correctly or not. In the example, the clause deallocates the `ns` endpoint (line 27).

`Sing#` uses *channel contracts* to detect communication errors. Figure 2 shows (part of) the `DSContract` contract associated with endpoint `ns` in Figure 1. A contract is made of *message specifications* and of *states* connected by *transitions*. Each message specification begins with the keyword `message` and is followed by the *tag* of the message and the type of its arguments. In Figure 2, `DSContract` defines the `Register` message with two arguments (a string and another endpoint)

¹This function has been taken from `./Services/RamDisk/ClientManager/RamDiskClientManager.sg` in the Singularity OS source code available at <http://www.codeplex.com/singularity/>. Here we have shortened some identifiers to fit the available space.

```

1 void GetNextDiskPath(out string! diskName,
2                       out SPContract.Exp! expService) {
3     DSContract.Imp:Ready ns = DS.NewClientEndpoint();
4     try {
5         while (true) {
6             SPContract.Imp! imp;
7             SPContract.Exp! exp;
8             SPContract.NewChannel(out imp, out exp);
9             diskName = pathPrefix + nextDiskNumber.ToString();
10            ns.SendRegister(Bitter.FromString2(diskName), imp);
11            switch receive {
12                case ns.AckRegister():
13                    nextDiskNumber++;
14                    expService = exp;
15                    return;
16                case ns.NakRegister(nakImp, error):
17                    if (error == ErrorCode.AlreadyExists)
18                        nextDiskNumber++;
19                    else
20                        throw new RamDiskErrorException(error);
21                    delete exp;
22                    delete nakImp;
23                    break;
24            }
25        }
26    } finally {
27        delete ns;
28    }
29 }

```

Figure 1: Example of Sing[#] function.

```

contract DSContract {
  out message Success();
  in  message Register(char[]! in ExHeap path,
                      SPContract.Imp:Start! imp);
  out message AckRegister();
  out message NakRegister(SPContract.Imp:Start imp,
                          ErrorCode error);

  // ...more message types

  state Start : one { Success! → Ready; }

  state Ready : one {
    Register? → DoRegister;
    CreateDirectory? → ...
    // ...more transitions
  }

  state DoRegister : one {
    AckRegister! → Ready;
    NakRegister! → Ready;
  }
}

```

Figure 2: Example of Sing[#] contract.

and the AckRegister message with no arguments. The `in` and `out` qualifiers specify the direction of messages from the point of view of the process exporting the contract. The state of the contract gives information about which messages can be sent/received at every given point in time. In DSContract we have a Ready state from which Register, CreateDirectory, and other (here omitted) messages can be received. After receiving a Register message, the contract moves to state DoRegister, from which one of the AckRegister or NakRegister messages can be sent, and then the contract goes back to the Ready state. In fact, each contract has two complementary views – called *exporting* and *importing* views – which are associated with the two peer endpoints of the channel. By convention, a contract declaration like that in Figure 2 specifies the exporting view of the contract: a provider of DSContract must adhere to its exporting view. In contrast,

the function `GetNextDiskPath` in Figure 1 acts as a consumer of `DSCContract`, therefore the function performs complementary actions by sending a `Register` message and then waiting for either an `AckRegister` or a `NakRegister` message. In the code, the importing and exporting views correspond to the types obtained by appending `.Imp` and `.Exp` suffixes to the name of the contract. For example, the declaration on line 3 specifies that `ns` is an endpoint having as type the importing view of `DSCContract` in state `Ready`. After line 10, the type associated with `ns` changes to `DSCContract.Imp:DoRegister` and then it goes back to `DSCContract.Imp:Ready` after any of the receive operations on lines 12 and 16. Note that the changes in the state of the contract associated with `ns` (and therefore of the type of `ns`) are not explicit in the source code. They follow from the initial declaration that brings `ns` into scope (line 3) and from the way `ns` is used in the function. By keeping track of the contract state of `ns`, the compiler can statically check that the actions performed on `ns` (for sending and receiving messages) match corresponding co-actions (for receiving and sending) performed on its peer endpoint, which is in use by some other process in the system.

The code structure in Figure 1, involving channel allocation and deallocation, messaging, delegation (sending endpoints over other endpoints), and exception handling, is in fact typical throughout the whole Singularity OS and shows that these aspects are frequently mixed in non-trivial ways. We can identify two main problems caused by exceptions:

1. Since communication errors are prevented by the complementarity of actions performed by processes accessing peer endpoints, a jump in the control flow of one process, like that caused by an exception, may disrupt the alignment of the peers of a channel and compromise subsequent interactions.
2. When an exception is thrown, messages that have been sent but not yet received and other objects allocated since the beginning of a `try` may become unreachable and therefore turn into memory leaks.

`Sing#` has limited and not fully satisfactory mechanisms for dealing with these problems. Regarding the first, `Sing#` provides an `InState` method through which it is possible to query, at runtime, the actual state of an endpoint. This information can be used to attempt recovery from a possibly inconsistent state of the endpoints. This mechanism implies an overhead for preserving and maintaining typing information at runtime and is unreliable as it depends on the programmer. The second problem seems to have been neglected. For example, the function

in Figure 1 is prone to leak memory on line 20 in the case that the exception is thrown, since neither `exp` nor `nakImp` are properly deallocated (`imp` has been sent away in the call to `SendRegister` so it is not the current thread’s responsibility to deallocate it). In this example it would suffice to move the `delete` instructions on lines 21 and 22 between lines 16 and 17 but, in general, it may be impossible to identify the exact point where an exception can be thrown and therefore when it is appropriate to deallocate resources. Note that it is unreasonable to assume that this clean-up code will be placed in the exception handler, if only because the handler may not be in the scope of the resources to be deallocated: in the example, `exp` and `nakImp` are not visible in the `finally` block so, by the time the exception has been thrown, it is too late to prevent the leak.

In the present paper we put forward a solution that combines static analysis (inspired by existing works on exception handling for sessions by Carbone et al. [11], Capecchi et al. [12]) with a transaction-like, all-or-nothing semantics of `try` blocks. The basic idea is that a `try` block is either executed completely, and then its effects on the heap are committed and become permanent, or it is aborted by an exception. If this happens, *all* the processes involved in the transaction are notified of the exception, so that the types of the endpoints they are using can remain aligned, and the state of the heap is restored to that at the beginning of the `try` block. Our solution relies on the following key ideas:

- (A) Following Carbone et al. [11], Capecchi et al. [12], we add explicit annotations to the types of endpoints used inside a transaction so that all processes involved in the transaction are aware of all the exceptions that can be thrown (possibly by a different process) during the transaction. In addition, these annotations make sure that the queues of endpoints used in a transaction are *empty* at the beginning of the transaction so that heap restoration solely amounts to removing messages from queues.
- (B) Inside `try` blocks, we “seal” the type of any endpoint whose type is not properly annotated and we forbid processes to use endpoints with a sealed type. In this way, the type system can statically ensure that well-typed processes do not modify any portion of the heap outside the restorable one.
- (C) We forbid the deallocation of endpoints inside `try` blocks, unless they have been allocated within the very same block. In this way, state restoration does not involve reallocations, which are difficult to implement correctly.

To prevent memory leaks, it is necessary to dynamically keep track of the

$P ::=$	done $ \text{open}(a, a).P$ $ \text{close}(u).P$ $ \text{u!m}(u).P$ $ \sum_{i \in I} \text{u?m}_i(x_i).P_i$ $ \text{P} \oplus \text{P}$ $ \text{P} \text{P}$ $ \text{try}(U) \{e_i : P_i\}_{i \in I}$ $ \text{throw } e$ $ \text{commit}(U).P$ $ \text{X}(\tilde{u})$	Process (inaction) (open channel) (close endpoint) (send) (receive) (conditional) (parallel) (initiate transaction) (exception) (commit transaction) (invocation)
$D ::=$	$\text{X}(\tilde{u}) \stackrel{\text{def}}{=} P$	Definition (rule)

Table 1: Syntax of processes and definitions.

memory allocated within a `try` block so that this memory can be properly reclaimed in case an exception is thrown. It is unsafe to deallocate an endpoint if its peer is not deallocated simultaneously: mechanism (A) guarantees that these deallocations are safe even if the type of these endpoints would not normally allow it, because transactions define a “closed scope” that includes, for each endpoint used in a transaction, also its peer. Starting with the next section we turn into the technical part of this work, in which we make precise all of the concepts informally introduced so far.

3. Language

Notation. We assume that we are given an infinite set *Pointers* of *heap addresses* ranged over by a, b, \dots , an infinite set *Variables* of *variables* ranged over by x, y, \dots , and a set *Exceptions* of *exceptions* ranged over by e, \dots . We let *names* u, v, \dots range over elements of $\text{Pointers} \cup \text{Variables}$. We use A, B, \dots to denote sets of pointers, \mathcal{E}, \dots to denote sets of exceptions, U to denote sets of names, and \tilde{u}, \tilde{v} to denote sequences of names (we will sometimes use \tilde{u} to denote also the set of names in \tilde{u}). Process variables are ranged over by X, Y, \dots .

Syntax. The process language is essentially a variant of the π -calculus [13], except that names represent heap pointers instead of communicating channels. *Processes* are defined by the grammar in Table 1. The term `done` denotes the idle process that performs no action. The term `open`(a, b). P denotes a process that allocates a new channel, represented as the two peer endpoints a and b , in the heap and continues as P . The term $u!m(v).$ P denotes a process that sends the message $m(v)$ on the endpoint u and then continues as P . A *message* is made of a *tag* m and an argument v . The term $\sum_{i \in I} u?m_i(x_i).$ P_i denotes a process that waits for a message from endpoint u . According to the tag m_i of the received message, the variable x_i is instantiated with the argument of the message in the continuation process P_i . We assume that the set I is always finite and non-empty. The term $P \oplus Q$ denotes a process that nondeterministically decides to behave as either P or Q , while the term $P | Q$ denotes the standard parallel composition of P and Q . The term `try`(U) $\{e_i : Q_i\}_{i \in I} P$ denotes a process willing to initiate a transaction involving the endpoints in U . The process P is the *body* of the transaction and is executed when the transaction is initiated, while the Q_i 's are the *handlers* of the transaction which are activated if the transaction is aborted during the execution of the body by an exception e_i . The term `throw` e denotes the throwing of the exception e , whose effect is to abort the currently running transaction and to execute the appropriate handler. The term `commit`(U). P denotes a process willing to terminate the currently running transaction (involving the endpoints in U). As soon as the transaction has ended, the process continues as P . The term $X\langle \tilde{u} \rangle$ denotes the invocation of the process associated with the process variable X . We assume that we work with a global environment of process definitions of the form

$$X(\tilde{u}) \stackrel{\text{def}}{=} P$$

defining these associations.

The binders of the language are `open`(a, b). P , which binds a and b in P , the input prefix $u?m(x).$ P , which binds x in P , and $X(\tilde{u}) \stackrel{\text{def}}{=} P$ which binds the names \tilde{u} in P . The formal definitions of free and bound names of a process P , respectively denoted by $\text{fn}(P)$ and $\text{bn}(P)$, are standard. We identify processes modulo alpha renaming of bound names.

Syntactic conventions. We adopt some standard conventions regarding the syntax of processes: we sometimes use an infix form for receive operations and write, for example $u?m_1(x_1).P_1 + \dots + u?m_n(x_n).P_n$ instead of $\sum_{i=1..n} u?m_i(x_i).P_i$; we omit message arguments when they are not used; we sometimes use a prefix form for

```

GetNextDiskPath( $DS, ret$ )  $\stackrel{\text{def}}{=} DS?NewClientEndpoint(ns).$ 
   $try(ns) \{RamDiskErrorException : Finally\langle ns, DS, ret \rangle\}$ 
     $Loop\langle ns, DS, ret \rangle$ 

Loop( $ns, DS, ret$ )  $\stackrel{\text{def}}{=} open(imp, exp).ns!Register(imp).$ 
   $ns?AckRegister().commit(ns).$ 
   $ret!SetService(exp).Finally\langle ns, DS, ret \rangle$ 
   $+ ns?NakRegister(nakImp).$ 
     $throw RamDiskErrorException$ 
     $\oplus close(exp).close(nakImp).Loop\langle ns, DS, ret \rangle$ 

Finally( $ns, DS, ret$ )  $\stackrel{\text{def}}{=} close(ns).ret!Result(DS).close(ret)$ 

```

Figure 3: Encoding of the function in Figure 1.

parallel compositions and write, for example, $\prod_{i=1..n} P_i$ instead of $P_1 \mid \dots \mid P_n$; we identify `done` with $\prod_{i \in \emptyset} P_i$ and we omit trailing occurrences of `done`.

To ease the formalization, our process language supports a minimal set of critical features: we focus only on monadic messaging (messages have exactly one endpoint argument) and exception handling, disregarding other constructs and data types of $\text{Sing}^\#$; we assume that receive operations use the same endpoint in every branch, forbidding processes like $u?a(x).P + v?b(y).Q$ which are allowed by the `switch receive` construct in $\text{Sing}^\#$; we work with a purely prefix-based language without sequential composition, encoding `try-catch-finally` blocks in $\text{Sing}^\#$ with transaction bodies and handlers and `commit` processes within bodies; we encode `if-else` commands with the non-deterministic process $P \oplus Q$ omitting the condition that determines the chosen branch. We claim that all the results presented hereafter can be suitably extended to overcome these restrictions.

Example 3.1. Figure 3 shows the encoding of the function in Figure 1 using the syntax of our process language. The structure of the process follows quite closely that of the function, except for some details which we explain here.

The loop on lines 5–25 is encoded as a recursive process `Loop` parameterized on its free names. The `finally` block on lines 26–28 is factored out as a named process `Finally`, since it must be executed regardless of whether the `try` block

is terminated successfully (line 15) or not (line 20). Consequently, `Finally` is invoked twice in the encoding.

The main difference between the function Figure 1 and its encoding concerns parameter passing, which is encoded using explicit communication on the `ret` endpoint. In particular, the initialization of `expService` with `exp` on line 14 corresponds to the output operation `ret!SetService(exp)` in Figure 3.

Note that in Figure 1 the function uses a global name `DS` for accessing a system service. In order to obtain a closed term, in the encoding we explicitly mention a parameter `DS` of the `GetNextDiskPath` process which represents `DS`. Because our type system relies on the linear access to resources, invoking a parametric process such as `GetNextDiskPath` means transferring the ownership of the parameters to the process. To preserve linearity (of `DS` in this case), the `Finally` process sends `DS` back on `ret` before `ret` is closed (more involved examples of function modeling and ownership transfer are described in detail by Bono and Padovani [7]). ■

Operational semantics. In order to describe the operational semantics of processes, we need to represent the *heap* where channels are allocated and through which messages are exchanged. Indeed, channels are accessed through the pointers to their endpoints and message arguments are themselves pointers to heap objects. Intuitively, a heap μ is a finite map from pointers a to endpoint structures $[b, \Omega]$, where b is the *peer endpoint* of a and Ω is the queue of messages waiting to be received from a . In the model, we represent heaps and message queues as terms generated by the grammar in Table 2. The term \emptyset denotes the empty heap, in which no endpoints are allocated. The term $a \mapsto [b, \Omega]$ denotes an endpoint allocated at a pointing to the endpoint structure $[b, \Omega]$. The term μ, μ' denotes the composition of the heaps μ and μ' . We write $\text{dom}(\mu)$ for the *domain* of the heap μ , that is the set of pointers for which there is an allocated endpoint structure. The heap composition μ, μ' is well defined provided that $\text{dom}(\mu) \cap \text{dom}(\mu') = \emptyset$ (there cannot be two endpoint structures allocated at the same address). In the following, we identify queues assuming associativity of composition and the laws $\varepsilon :: \Omega = \Omega :: \varepsilon = \Omega$ and we identify heaps assuming associativity and commutativity of composition and the law $\emptyset, \mu = \mu$. We write $a \mapsto [b, \Omega] \in \mu$ to indicate that the endpoint structure $[b, \Omega]$ is allocated at location a in μ .

Message queues, ranged over by Ω , are also represented as terms: ε denotes the empty queue, $m(c)$ is a queue made of an m message with argument c , and $\Omega :: \Omega'$ is the queue composition of Ω and Ω' . We identify queues modulo associativity of $::$ and we assume that ε is neutral for $::$.

Before defining the operational semantics of processes we formalize two no-

$\mu ::=$	\emptyset $a \mapsto [a, \Omega]$ μ, μ	Heap (empty heap) (endpoint structure) (heap composition)
$\Omega ::=$	ε $m(a)$ $\Omega :: \Omega$	Queue (empty queue) (message) (queue composition)
$P ::=$	\dots $\langle A, B, \{e_i : P_i\}_{i \in I} P \rangle$	Runtime process (as in Table 1) (running transaction)

Table 2: Syntax of heaps, queues, and runtime processes.

tions. The first one is that of peer endpoints:

Definition 3.1 (peer endpoints). We say that a and b are *peer endpoints* in μ , written $a \overset{\mu}{\leftrightarrow} b$, if $a \neq b$ and $a \mapsto [b, \Omega] \in \mu$ and $b \mapsto [a, \Omega'] \in \mu$.

Note that $\overset{\mu}{\leftrightarrow}$ is a symmetric relation. The notion of “closed scope” that we mentioned in Section 2 is formalized as a predicate on sets of pointers:

Definition 3.2 (balanced set of pointers). We say that $A \subseteq \text{dom}(\mu)$ is *balanced* in μ , written μ -balanced(A), if, for every $a \in A$, $a \overset{\mu}{\leftrightarrow} b$ implies $b \in A$.

In words, A is balanced in μ if for every a in A , the peer of a is also in A provided that it is still allocated in μ . Since a message sent over a ends up in the queue of its peer, this means that any communication occurring on one of the endpoints in A remains within the scope identified by A .

In the operational semantics of processes, we need to distinguish between a transaction that has not started yet (and which is represented using the `try` construct of Table 1), and a *running transaction*. This need arises for two reasons: First, a running transaction generally involves more than one process, each with its own set of handlers. Therefore, it is technically convenient to devise an explicit construct that defines the *scope* of the transaction. Second, it is necessary to keep track of the part of the heap that has been allocated since the initiation of the transaction. Table 2 extends the syntax of processes with the term $\langle A, B, \{e_i : P_i\}_{i \in I} P \rangle$

Structural congruence		
[S-PAR IDLE] $P \mid \text{done} \equiv P$	[S-PAR COMM] $P \mid Q \equiv Q \mid P$	[S-PAR ASSOC] $P \mid (Q \mid R) \equiv (P \mid Q) \mid R$
Reduction relation		
[R-OPEN] $\mu \circledast \text{open}(a, b).P \rightarrow \mu, a \mapsto [b, \varepsilon], b \mapsto [a, \varepsilon] \circledast P$	[R-PARALLEL] $\frac{\mu \circledast P \rightarrow \mu' \circledast P'}{\mu \circledast P \mid Q \rightarrow \mu' \circledast P' \mid Q}$	
[R-CLOSE] $\mu, a \mapsto [b, \Omega] \circledast \text{close}(a).P \rightarrow \mu \circledast P$	[R-CHOICE] $\frac{i \in \{1, 2\}}{\mu \circledast P_1 \oplus P_2 \rightarrow \mu \circledast P_i}$	
[R-SEND] $\mu, a \mapsto [b, \Omega], b \mapsto [a, \Omega'] \circledast a!m(c).P \rightarrow \mu, a \mapsto [b, \Omega], b \mapsto [a, \Omega' :: m(c)] \circledast P$		
[R-RECEIVE] $\frac{k \in I}{\mu, a \mapsto [b, m_k(c) :: \Omega] \circledast \sum_{i \in I} a?m_i(x_i).P_i \rightarrow \mu, a \mapsto [b, \Omega] \circledast P_k\{c/x_k\}}$		
[R-INVOKE] $\frac{X(\tilde{u}) \stackrel{\text{def}}{=} P}{\mu \circledast X\langle \tilde{a} \rangle \rightarrow \mu \circledast P\{\tilde{a}/\tilde{u}\}}$	[R-STRUCT] $\frac{P \equiv P' \quad \mu \circledast P' \rightarrow \mu' \circledast Q' \quad Q' \equiv Q}{\mu \circledast P \rightarrow \mu' \circledast Q}$	

Table 3: Operational semantics of processes.

where A is the set of endpoints involved in the transaction, B is the set of endpoints that have been allocated since the transaction has started, P is the (residual) body of the transaction, and the P_i 's represent the handlers of the transaction. In general, P and the P_i 's will be parallel compositions of the bodies and the handlers of the processes that have cooperatively initiated the transaction.

The operational semantics of processes is defined in terms of a structural congruence over processes (identifying structurally equivalent processes) and a reduction relation. Structural congruence is the least relation \equiv including alpha conversion and the laws in Table 3, stating that parallel composition is commutative, associative, and has `done` as neutral element. As process interaction mostly

occurs through the heap, the reduction relation describes the evolution of *configurations* $\mu \wp P$ rather than of processes alone, so that

$$\mu \wp P \rightarrow \mu' \wp P'$$

denotes the fact that process P evolves to P' and, in doing so, it changes the heap from μ to μ' .

Reduction is the smallest relation between configurations defined by the rules in Tables 3 and 4. We explain the rules in the following paragraphs. Rule [R-OPEN] describes the creation of a new channel, which causes the allocation of two new endpoint structures in the heap. The endpoints are initialized with empty queues and are allocated at fresh locations, for otherwise the resulting heap would be ill formed. Since we have assumed that Pointers is infinite, it is always possible to alpha rename a and b to fresh pointers using structural congruence, so that an $\text{open}(a, b).P$ is always able to reduce.

Rule [R-CLOSE] describes the closing of an endpoint, which deallocates its structure from the heap and discards its queue. Note that both endpoints of a channel are created simultaneously by [R-OPEN], but each is closed independently by [R-CLOSE] (this is the same semantics as the one of Sing[#]).

Rule [R-CHOICE] states that a process $P \oplus Q$ nondeterministically reduces to either P or Q .

Rule [R-SEND] describes the sending of a message $m(c)$ on the endpoint a . The message is enqueued at the right end of the queue associated with the peer endpoint b of a . Note that, for this rule to be applicable, it is necessary for both endpoints of a channel to still be allocated.

Rule [R-RECEIVE] describes the receiving of a message from endpoint a . In particular, the message at the left end of the queue associated with a is removed from the queue, its tag m_k is used to select one branch of the process, and its argument c instantiates the corresponding variable x_k .

Rule [R-PARALLEL] describes the independent evolution of parallel processes. Note how the heap is treated globally even if it is only one subprocess to reduce.

Rule [R-INVOKE] describes process invocations simply as the replacement of a process variable with the process it is associated with, modulo the substitution of its parameters. In this rule and in [R-RECEIVE], $P\{\tilde{u}/\tilde{v}\}$ denotes the capture-avoiding substitution of \tilde{u} in place of \tilde{v} in P .

Rule [R-START TRANSACTION] describes the initiation of a transaction by a number of processes. The transaction is identified by a set of endpoints $\bigcup_{i \in I} A_i$ which are distributed among the processes. In order for the transaction to start, this

<div style="border-bottom: 1px solid black; padding-bottom: 5px;"> <p>[R-START TRANSACTION]</p> $\frac{\mu\text{-balanced}(\bigcup_{i \in I} A_i)}{\mu \circ \prod_{i \in I} \text{try}(A_i) \{e_j : Q_{ij}\}_{j \in J} P_i \rightarrow \mu \circ \langle \bigcup_{i \in I} A_i, \emptyset, \{e_j : \prod_{i \in I} Q_{ij}\}_{j \in J} \prod_{i \in I} P_i \rangle}$ </div> <div style="padding: 5px;"> <p>[R-END TRANSACTION]</p> $\mu \circ \langle A, B, \{e_j : Q_j\}_{j \in J} \prod_{i \in I} \text{commit}(A_i).P_i \rangle \rightarrow \mu \circ \prod_{i \in I} P_i$ </div>
<div style="border-bottom: 1px solid black; padding-bottom: 5px;"> <p>[R-RUN TRANSACTION]</p> $\mu \circ P \rightarrow \mu' \circ P'$ </div> <div style="padding: 5px;"> $\mu \circ \langle A, B, \{e_j : Q_j\}_{j \in J} P \rangle \rightarrow \mu' \circ \langle A, \text{track}(B, \text{dom}(\mu), \text{dom}(\mu')), \{e_j : Q_j\}_{j \in J} P' \rangle$ </div>
<div style="border-bottom: 1px solid black; padding-bottom: 5px;"> <p>[R-CATCH EXCEPTION]</p> $k \in J$ </div> <div style="padding: 5px;"> $\mu_1, \{a_i \mapsto [b_i, \Omega_i]\}_{i \in I}, \mu_2 \circ \langle \{a_i\}_{i \in I}, \text{dom}(\mu_2), \{e_j : Q_j\}_{j \in J} \text{throw } e_k \mid P \rangle \rightarrow \mu_1, \{a_i \mapsto [b_i, \varepsilon]\}_{i \in I} \circ Q_k$ </div>
<div style="border-bottom: 1px solid black; padding-bottom: 5px;"> <p>[R-PROPAGATE EXCEPTION]</p> $\forall j \in J : e_j \neq e$ </div> <div style="padding: 5px;"> $\mu_1, \{a_i \mapsto [b_i, \Omega_i]\}_{i \in I}, \mu_2 \circ \langle \{a_i\}_{i \in I}, \text{dom}(\mu_2), \{e_j : Q_j\}_{j \in J} \text{throw } e \mid P \rangle \rightarrow \mu_1, \{a_i \mapsto [b_i, \varepsilon]\}_{i \in I} \circ \text{throw } e$ </div>

Table 4: Operational semantics of transactions.

set of endpoints must be balanced, so that for every endpoint in the set its peer is also in the set. The rule is nondeterministic, in the sense that there can be multiple combinations of processes that can initiate a transaction. We leave the choice of a particular strategy (for example, requiring $\bigcup_{i \in I} A_i$ to be non-empty, minimal, and μ -balanced) to the implementation. The residual process is the tuple

$$\langle \bigcup_{i \in I} A_i, \emptyset, \{e_j : \prod_{i \in I} Q_{ij}\}_{j \in J} \prod_{i \in I} P_i \rangle$$

combining the bodies and the handlers of the processes involved in the transaction. The second component is \emptyset indicating that at this stage no new endpoints have been allocated yet within the transaction. Note that the combined processes must be able to handle the same set $\{e_j\}_{j \in J}$ of exceptions for the reduction to occur. Even if this seems to require a runtime check, the fact that all processes involved in a transaction are able to handle the same set of exceptions will be ensured by the type system (see rule [T-TRY] in Table 8).

Rule [R-END TRANSACTION] reduces a running transaction to its continuation when its body has terminated. The handlers are discarded. The sets A_i play no role in the operational semantics and are used for typing purposes only. In fact, we will see that the type system enforces the invariant that these sets coincide with the ones decorating the `try` blocks corresponding to the `commit` processes and, in particular, $A = \bigcup_{i \in I} A_i$.

Rule [R-RUN TRANSACTION] allows the reduction of a transaction according to the reductions of its body. The rule keeps track of the memory changes occurred during the reduction of the body of the transaction by updating the second component of the transaction to $\text{track}(B, \text{dom}(\mu), \text{dom}(\mu'))$, where

$$\text{track}(B, A_0, A_1) \stackrel{\text{def}}{=} (B \cup (A_1 \setminus A_0)) \setminus (A_0 \setminus A_1)$$

In practice, the pointers to objects allocated during the reduction are added to B , while the pointers to objects deallocated during the reduction are removed from B .

Rule [R-CATCH EXCEPTION] describes the abnormal termination of a running transaction when an exception is thrown and the transaction provides an handler for it. In this case, the queues of all the endpoints involved in the transactions are emptied, the memory allocated within the transaction is reclaimed, and the appropriate handler is run. In a similar way, rule [R-PROPAGATE EXCEPTION] abnormally terminates running transactions when there is no suitable handler for the thrown exception. Also in this case the queues of the endpoints involved in the transactions are emptied and the memory allocated within the transaction is reclaimed, but the exception is propagated (technically, *re-thrown*) at the outer level.

We write $\mu \circledast P \rightarrow$ if $\mu \circledast P \rightarrow \mu' \circledast P'$ for some μ' and P' and $\mu \circledast P \not\rightarrow$ if not $\mu \circledast P \rightarrow$; we write \Rightarrow for the reflexive, transitive closure of \rightarrow .

Well-behaved processes. We conclude this section providing a characterization of *well-behaved processes*, those that are free from memory leaks, memory faults, and communication errors. A *memory leak* occurs when no pointer to an allocated region of the heap is retained by any process. In this case, the allocated region has no owner, it occupies space, but it is no longer accessible. A *memory fault* occurs when a pointer is accessed and the endpoint it points to is not (or no longer) allocated. A *communication error* occurs when some process receives a message of unexpected type. To formalize well-behaved processes, we need to define the reachability of a heap object with respect to a set of *root* pointers. Intuitively, a process P may directly reach any object located at some pointer in the set $\text{fn}(P)$

[ST-INACTIVE] $\mu \circledast \text{done} \downarrow$	[ST-INPUT] $\mu, a \mapsto [b, \varepsilon] \circledast \sum_{i \in I} a?m_i(x_i).P_i \downarrow$	[ST-COMMIT] $\mu \circledast \text{commit}(A).P \downarrow$
[ST-TRY] $\frac{\neg \mu\text{-balanced}(A)}{\mu \circledast \text{try}(A) \{e_j : Q_j\}_{j \in J} P \downarrow}$		[ST-PARALLEL] $\frac{\mu \circledast P \downarrow \quad \mu \circledast Q \downarrow}{\mu \circledast P \mid Q \downarrow}$
[ST-RUNNING TRANSACTION] $\frac{\mu \circledast P \downarrow \quad P \neq \prod_{i \in I} \text{commit}(A_i).P_i}{\mu \circledast \langle A, B, \{e_j : Q_j\}_{j \in J} P \rangle \downarrow}$		

Table 5: Stuck configurations.

(we can think of the pointers in $\text{fn}(P)$ as of the local variables of the process stored on its stack); from these pointers, the process may reach other heap objects by reading messages from the endpoints it can reach, and so forth.

Definition 3.3 (reachable pointers). We say that c is *reachable* from a in μ , notation $c \prec_\mu a$, if $a \mapsto [b, \mathcal{Q} :: m(c) :: \mathcal{Q}'] \in \mu$. We write \preceq_μ for the reflexive, transitive closure of \prec_μ and we define $\mu\text{-reach}(A) = \{c \in \text{Pointers} \mid \exists a \in A : c \preceq_\mu a\}$.

The last auxiliary notion we need provides a syntactic characterization of those configurations that cannot reduce but that do not represent any of the errors described above.

Definition 3.4 (stuck configuration). We say that the configuration $\mu \circledast P$ is *stuck* if the judgment $\mu \circledast P \downarrow$ is inductively derivable by the rules in Table 5.

Rules [ST-INACTIVE] and [ST-PARALLEL] are obvious, while rules [ST-TRY] and [ST-COMMIT] state that transaction initiations and termination are stuck, if taken in isolation. In the former case, the set of involved endpoints must not be balanced, for otherwise the transaction could initiate. Rule [ST-RUNNING TRANSACTION] states that a running transaction is stuck if its body is stuck and different from a combination of processes willing to terminate the transaction, for otherwise the transaction could terminate. Note also that no exception can have been thrown within the body of a stuck running transaction, for the stuckness predicate is undefined for `throw e` processes. Finally, rule [ST-INPUT] states that a process waiting for a message from endpoint a is stuck only if the endpoint a is allocated and its queue is empty. Then, a configuration whose processes are all waiting for a

message corresponds to a genuine deadlock. From these rules we deduce that a process willing to send a message on a is never stuck, and so is a process willing to receive a message from a if the queue associated with a is not empty.

Definition 3.5 (well-behaved process). We say that P is *well behaved* if $\emptyset \varepsilon P \Rightarrow \mu \varepsilon Q$ implies:

1. $\text{dom}(\mu) \subseteq \mu\text{-reach}(\text{fn}(Q))$;
2. $Q \equiv Q_1 \mid Q_2$ and $\mu \varepsilon Q_1 \not\rightarrow$ imply $\mu \varepsilon Q_1 \downarrow$.

In words, a process P is well behaved if every residual Q of P is such that Q can reach every pointer in the heap and every subprocess Q_1 of Q that does not reduce is stuck (recall that the definition of $\mu \varepsilon Q_1 \downarrow$ captures also the possibility that Q_1 is in deadlock). Here are a few examples of ill-behaved processes to illustrate the sort of errors we want to spot with our type system:

- The process `open(a,b).done` violates condition (1), since it leaks endpoints a and b .
- The process `open(a,b).(close(a).close(a) | close(b))` tries to deallocate the same endpoint a twice. This is an example of fault.
- The process `open(a,b).(a!a().close(a) | b?b().close(b))` violates condition (2) since it reduces to a parallel composition of subprocesses where one has sent an a message, but the other one was expecting a b message.
- The process

$$\text{open}(a,b).\text{try}(\emptyset) \{e : \text{done}\} \\ \text{throw } e \oplus \text{commit}(\emptyset).\text{close}(a).\text{close}(b)$$

may leak a and b if the exception is thrown.

Observe that, in item (1) of Definition 3.5, the domain of μ is only required to be *included in* (instead of being *equal to*) the set of pointers reachable from the free names of Q . In particular, it may be the case that Q contains references to unallocated objects, and yet it never attempts to use them. This formulation of leak-freedom, which is slightly more general than the one used by Jakšić and Padovani [10] where equality between the two sets was required, is necessary because the type system that we are about to define allows subtyping, which was

$t ::=$	T	Type (endpoint type)
	$[t]$	(sealed type)
$T ::=$	end	Endpoint type (termination)
	α	(type variable)
	$\{!m_i(T_i).T_i\}_{i \in I}$	(internal choice)
	$\{?m_i(T_i).T_i\}_{i \in I}$	(external choice)
	$\{e_i : T_i\}_{i \in I} \llbracket T$	(initiate transaction)
	$\rrbracket T$	(commit transaction)
	$\text{rec } \alpha : r.T$	(recursive type)
	$\{e_i : T_i\}_{i \in I} T$	(running transaction)

Table 6: Syntax of types and endpoint types.

not considered in [10]. Also note that our notion of leak-freedom does not require a process to eventually deallocate the objects it owns, but only to guarantee the reachability of all the objects it owns. For example, the process $\text{open}(a, b).X\langle a, b \rangle$ where $X(u, v) \stackrel{\text{def}}{=} X\langle u, v \rangle$, maintains the reachability of a and b without ever using them. This process is well behaved according to Definition 3.5 and is also well typed according to the type system that will be developed in Section 4.

4. Type System

We now develop a type system that enforces well-behavedness of processes: in Section 4.1 we introduce the syntax of the type language; in Section 4.2 we define a notion of *type weight* which is used for discriminating between safe and unsafe communications; Section 4.3 is devoted to extending classical subtyping for session types by Gay and Hole [14] so as to take transactions and exceptions into account; Sections 4.4, 4.5, and 4.6 define the actual typing rules; finally, Section 4.7 presents the soundness result.

4.1. Syntax of Types

We assume that we are given an infinite set of *type variables* ranged over by α ; we use t, s, \dots to range over types, and T, S, \dots to range over endpoint types. The syntax of types and endpoint types is defined in Table 6. An endpoint

type describes the behavior of a process with respect to a particular endpoint: the process may send messages over the endpoint, receive messages from the endpoint, deallocate the endpoint, initiate and terminate transactions involving the endpoint. The endpoint type `end` denotes an endpoint that can only be deallocated. An internal choice $\{!m_i(S_i).T_i\}_{i \in I}$ denotes an endpoint on which a process may send any message m_i for $i \in I$. The message has an argument of type S_i and, depending on the tag m_i , the endpoint can be used thereafter according to T_i . In a dual manner, an external choice $\{?m_i(S_i).T_i\}_{i \in I}$ denotes an endpoint from which a process must be ready to receive any message m_i for $i \in I$ and, depending on the tag m_i of the received message, the endpoint is to be used according to T_i . In endpoint types $\{!m_i(S_i).T_i\}_{i \in I}$ and $\{?m_i(S_i).T_i\}_{i \in I}$ we assume that $I \neq \emptyset$ and $m_i = m_j$ implies $i = j$ for every $i, j \in I$. That is, the tag m_i of the message that is sent or received identifies a unique continuation T_i . The endpoint type $\{e_i : S_i\}_{i \in I} \llbracket T$ denotes an endpoint on which it is possible to initiate a transaction. The type T specifies how the endpoint is used within the body of the transaction, whereas each type S_i specifies how the endpoint is used if the transaction is aborted by the exception e_i . The endpoint type $\llbracket T$ denotes the termination of the transaction in which an endpoint with this type is involved. As soon as the transaction is properly terminated, the endpoint can be subsequently used according to T . Terms α and $\text{rec } \alpha : r.T$ can be used to specify recursive behaviors, as usual. The annotation r associated with α represents the rank of α , which will be explained shortly. Finally, the endpoint type $\{e_i : S_i\}_{i \in I} T$ is analogous to $\{e_i : S_i\}_{i \in I} \llbracket T$, except that it specifies the type of an endpoint involved in a transaction which has already been initiated, but has not terminated yet. In fact, this type is needed for technical reasons only, and will be used in conjunction with running transaction processes $\langle A, B, \{e_i : P_i\}_{i \in I} P \rangle$. In no case the programmer is supposed to deal with endpoint types of this form.

Clearly, not every endpoint type written according to the syntax in Table 6 makes sense. For example, it is possible to write terms such as $\{e : \text{end}\} \llbracket \text{end}$ where a transaction is initiated but not terminated or terms where recursions do not respect the intended nesting of transactions, like in $\text{rec } \alpha. \{e : \text{end}\} \llbracket \alpha$ or in $\{e : \text{end}\} \llbracket \text{rec } \alpha. \llbracket \alpha$. As far as our analysis is concerned, the syntax does not even prevent `end` subterms from occurring within transactions, which as we have argued in Section 2 is undesirable since endpoints involved in transactions should not be closed. For all these reasons we define a subset of *well-formed* endpoint types based on a notion of *rank*. Intuitively, the rank of a term T gives the number of transactions within which T may occur, with the proviso that `end` must have rank 0.

$\frac{}{\Theta \vdash \text{end} : 0}$	$\frac{}{\Theta, \alpha : r \vdash \alpha : r}$	$\frac{[\text{WF-REC}] \quad \Theta, \alpha : r \vdash T : r}{\Theta \vdash \text{rec } \alpha : r.T : r}$
$\frac{[\text{WF-PREFIX}] \quad \dagger \in \{?, !\} \quad \Theta \vdash S_i : 0 \quad (i \in I) \quad \Theta \vdash T_i : r \quad (i \in I)}{\Theta \vdash \{\dagger m_i(S_i).T_i\}_{i \in I} : r}$	$\frac{[\text{WF-COMMIT}] \quad \Theta \vdash T : r}{\Theta \vdash \llbracket T \rrbracket : r+1}$	
$\frac{[\text{WF-INITIATE}] \quad \Theta \vdash S_i : r \quad (i \in I) \quad \Theta \vdash T : r+1}{\Theta \vdash \{e_i : S_i\}_{i \in I} \llbracket T \rrbracket : r}$	$\frac{[\text{WF-RUN}] \quad \Theta \vdash S_i : r \quad (i \in I) \quad \Theta \vdash T : r+1}{\Theta \vdash \{e_i : S_i\}_{i \in I} T : r}$	

Table 7: Rank of endpoint types.

In general, we say that the endpoint type T is well formed and has rank r in Θ if $\Theta \vdash T : r$ is inductively derivable by the axioms and rules in Table 7, where Θ ranges over ranking contexts associating ranks to type variables. Then, a derivation of $\emptyset \vdash T : 0$ means that T is a closed endpoint type where transaction initiations and terminations are balanced. Rules [WF-INITIATE], [WF-RUN], and [WF-COMMIT] count the number of nested transactions. Rule [WF-PREFIX] requires all branches of a choice to have the same rank, while rules [WF-REC] and [WF-VAR] deal with recursive types in a standard way, by respectively augmenting and accessing the ranking context. In the following we will omit Θ from judgments $\Theta \vdash T : r$ if Θ is empty.

As welcome side effects of well formedness, note that:

- message types have rank 0 (rule [WF-PREFIX]). Thus, well-typed processes will not be able to send/receive endpoints involved in pending transactions;
- `end` cannot occur inside transactions (rule [WF-END]). Thus, well-typed processes will not be able to close endpoints involved in pending transactions.

The rank annotation r in recursive terms $\text{rec } \alpha : r.T$ guarantees that every well-formed endpoint type has a uniquely determined rank. Without this annotation a term like $\text{rec } \alpha. !m(\text{end}).\alpha$ could be given any rank. The following proposition guarantees that the rank of well-formed endpoint types is unaffected by folding/unfolding of recursions:

Proposition 4.1. *If $\vdash \text{rec } \alpha : r.T : r$, then $\vdash T\{\text{rec } \alpha : r.T/\alpha\} : r$.*

Proof. A simple induction on the derivation of $\Theta, \alpha : r \vdash T : r$. □

In what follows, we will assume that all endpoint types are closed and well formed and we will usually omit the rank annotation from recursive terms with the assumption that they can be properly annotated so that they are well formed; we will also write $\text{rank}(T)$ for the rank of T . We will identify endpoint types modulo alpha renaming of bound type variables (the only binder being rec) and folding/unfolding of recursions knowing that this does not change their rank (Proposition 4.1). In particular, we have $\text{rec } \alpha.T = T\{\text{rec } \alpha.T/\alpha\}$. Finally, we will sometimes use an infix notation for internal and external choices and write $!m_1(S_1).T_1 \oplus \dots \oplus !m_n(S_n).T_n$ instead of $\{!m_i(S_i).T_i\}_{i \in \{1, \dots, n\}}$ and $?m_1(S_1).T_1 + \dots + ?m_n(S_n).T_n$ instead of $\{?m_i(S_i).T_i\}_{i \in \{1, \dots, n\}}$.

Types are possibly sealed endpoint types of the form

$$[\dots[T]\dots]$$

for some arbitrary (possibly zero) number of seals $[\dots]$. Seals protect the endpoints not involved in a transaction: they are applied when the transaction is initiated (the `try` primitive is executed) and are stripped off when the transaction terminates (the `commit` primitive is executed). The type system prevents endpoints with a sealed type from being used, since any change to them would not be undoable in case the currently running transaction is aborted.

Example 4.1. According to the process definitions in Figure 3, the endpoint ns is involved in the transaction around the `Loop` process, it is used for sending a `Register` message and then for receiving either an `AckRegister` or a `NakRegister` message. The same endpoint is then closed regardless of whether the transaction completes successfully or not. We can describe the overall behavior of `GetNextDiskPath`, `Loop`, and `Finally` on ns with the following endpoint type:

$$T_{ns} = \{\text{RamDiskErrorException} : \text{end}\} \llbracket \text{rec } \alpha. !\text{Register}(T_{imp}). \\ (\text{?AckRegister}(). \llbracket \text{end} + \text{?NakRegister}(T_{imp}). \alpha \rrbracket) \rrbracket$$

where T_{imp} is the endpoint type associated with the `imp` and `nakImp` endpoints.

The endpoint ret is not used within the transaction, but its usage differs depending on whether or not the exception is thrown:

$$T_{ret} = !\text{Result}(T_{DS}). \text{end} \oplus !\text{SetService}(T_{exp}). !\text{Result}(T_{DS}). \text{end}$$

If no exception is thrown, ret is used for sending a `SetRegister` message followed by a `Result` one; if an exception is thrown, only the `Result` message is sent. The above type T_{ret} takes into account both possibilities. ■

In order to avoid communication errors, we associate peer endpoints with endpoint types describing complementary actions: if a process sends a message of some kind on one endpoint, another process must be able to receive a message of that kind from the peer endpoint; if one process initiates a transaction involving one endpoint, the other process will do so as well on the peer endpoint; if one process has finished using an endpoint, the process owning the peer endpoint has finished too. We formalize this complementarity of actions by defining a function that, given an endpoint type, computes its dual:

Definition 4.1 (duality). *Duality* is the function $\bar{\cdot}$ on endpoint types defined coinductively by the equations:

$$\begin{aligned} \overline{\mathbf{end}} &= \mathbf{end} \\ \overline{\{?m_i(S_i).T_i\}_{i \in I}} &= \{!m_i(S_i).\bar{T}_i\}_{i \in I} \\ \overline{\{!m_i(S_i).T_i\}_{i \in I}} &= \{?m_i(S_i).\bar{T}_i\}_{i \in I} \\ \overline{\{e_i : S_i\}_{i \in I} \llbracket T \rrbracket} &= \{e_i : \bar{S}_i\}_{i \in I} \llbracket \bar{T} \rrbracket \\ \overline{\llbracket T \rrbracket} &= \llbracket \bar{T} \rrbracket \\ \overline{\{e_i : S_i\}_{i \in I} T} &= \{e_i : \bar{S}_i\}_{i \in I} \bar{T} \end{aligned}$$

Roughly speaking, the dual of an endpoint type T is obtained from T by swapping internal and external choices. For example, the dual of the endpoint type T_{ret} defined in Example 4.1 is

$$\bar{T}_{ret} = ?\mathbf{Result}(T_{DS}).\mathbf{end} + ?\mathbf{SetService}(T_{exp}).?\mathbf{Result}(T_{DS}).\mathbf{end}$$

Note that the dual \bar{T} of T cannot be defined by a simple induction on the structure of T according to this intuition because the type of message arguments is *unaffected* by duality. In particular we have

$$\begin{aligned} \overline{\mathbf{rec} \alpha. ?m(\alpha).\mathbf{end}} &= \overline{?m(\mathbf{rec} \alpha. ?m(\alpha).\mathbf{end}).\mathbf{end}} \\ &= !m(\mathbf{rec} \alpha. ?m(\alpha).\mathbf{end}).\mathbf{end} \\ &\neq \mathbf{rec} \alpha. !m(\alpha).\mathbf{end}. \end{aligned}$$

The interested reader may refer to Bono and Padovani [7] for an equivalent inductive definition of duality.

We list here two important properties of duality, namely that it is an involution and it preserves ranks:

Proposition 4.2. *The following properties hold:*

1. $\overline{\overline{T}} = T$;
2. $\text{rank}(\overline{T}) = \text{rank}(T)$.

Proof. Item (1) is an easy consequence of the definition of duality (Definition 4.1). Item (2) follows from the fact that the rank is only affected by the nesting of transaction types in T and internal/external choices are treated in the same way by rule [WF-PREFIX]. \square

4.2. Type Weight

In previous work by Bono and Padovani [7] it was observed that the delegation of endpoints having some particular types can generate memory leaks even if the delegating process appears to behave correctly with respect to the type of the endpoints it uses. For example, the process

$$P \stackrel{\text{def}}{=} \text{open}(a, b). a!m(b). \text{close}(a) \quad (1)$$

uses a and b according to the endpoint types

$$T = !m(S). \text{end} \quad \text{and} \quad S = \text{rec } \alpha : 0. ?m(\alpha). \text{end} \quad (2)$$

respectively. Note that $\overline{T} = S$, therefore the complementarity of actions performed on the peer endpoints a and b is guaranteed. Now, the process P sends endpoint b over endpoint a . According to T , the process is indeed entitled to send an m message with argument of type S on a and b has precisely that type. After the output operation, the process no longer owns endpoint b and endpoint a is deallocated. Despite its apparent correctness, P generates a leak, as shown by the reduction:

$$\begin{aligned} \emptyset \ ; \ P &\rightarrow a \mapsto [b, \varepsilon], b \mapsto [a, \varepsilon] \ ; \ a!m(b). \text{close}(a) \\ &\rightarrow a \mapsto [b, \varepsilon], b \mapsto [a, m(b)] \ ; \ \text{close}(a) \\ &\rightarrow b \mapsto [a, m(b)] \ ; \ \text{done} \end{aligned}$$

In the final configuration we have $\mu\text{-reach}(\text{fn}(\text{done})) = \emptyset$ while $\text{dom}(\mu) = \{b\}$. In particular, the endpoint b is no longer reachable and therefore this configuration violates condition (1) of Definition 3.5. A closer look at the heap in the reduction above reveals that the problem lies in the cycle involving b resulting from the send operation $a!m(b)$: it is as if the $b \mapsto [a, m(b)]$ region of the heap

needs not be owned by any process because “it owns itself”. Fortunately, it is possible to detect the situations in which these cycles may be generated by looking at the structure of the types of the endpoints that are sent as messages. More specifically, for each endpoint type we compute a value in the set $\mathbb{N} \cup \{\infty\}$, which we call *weight*, that estimates the length of any chain of pointers originating from the queue of the endpoints it denotes. A weight equal to ∞ means that this length can be infinite, in the sense that cycles such as the one shown above may be generated. Then, the type system makes sure that only endpoints having a finite-weight type can be sent as messages, and this has been shown to be enough for preventing these kinds of memory leaks (see Bono and Padovani [7]).

We proceed by recalling here the definition of weight from [7], adapted to our context where we deal also with transaction types:

Definition 4.2 (weight). We say that \mathscr{W} is a *coinductive weight bound* if $(T, n) \in \mathscr{W}$ implies either:

- $T = \text{end}$ or $T = \{e_i : S_i\}_{i \in I} \llbracket T' \text{ or } T = \llbracket T' \text{ or } T = \{!m_i(S_i).T_i\}_{i \in I}$, or
- $T = \{?m_i(S_i).T_i\}_{i \in I}$ and $n > 0$ and $(S_i, n - 1) \in \mathscr{W}$ and $(T_i, n) \in \mathscr{W}$ for every $i \in I$, or
- $T = \{e_i : S_i\}_{i \in I} T'$ and $(T', n) \in \mathscr{W}$.

We write $T :: n$ if $(T, n) \in \mathscr{W}$ for some coinductive weight bound \mathscr{W} . The *weight* of an endpoint type T , denoted by $\|T\|$, is defined by

$$\|T\| = \min\{n \in \mathbb{N} \mid T :: n\}$$

where we let $\min \emptyset = \infty$. When comparing weights we extend the usual total orders $<$ and \leq over natural numbers so that $n < \infty$ for every $n \in \mathbb{N}$ and $\infty \leq \infty$.

The weight of T is defined as the least of its weight bounds, or ∞ if there is no such weight bound. For example we have $\|\text{end}\| = \|\{!m_i(S_i).T_i\}_{i \in I}\| = 0$. Indeed, the queues of endpoints with type `end` and those in a send state are empty and therefore the chains of pointers originating from them have zero length. The same happens for endpoints whose type is $\{e_i : S_i\}_{i \in I} \llbracket T$ and $\llbracket T$, since we will enforce the invariant that when a transaction is initiated or successfully terminated, the endpoints involved in it have empty queues. Endpoint types in a receive state have a strictly positive weight. For instance we have $\|\?m(\text{end}).\text{end}\| = 1$ and $\|\?m(\?m(\text{end}).\text{end}).\text{end}\| = 2$. Indeed, the queue of an endpoint with type

$?m(\text{end}).\text{end}$ may contain another endpoint with an empty queue. Therefore, the chain of pointers originating from the endpoint with type $?m(\text{end}).\text{end}$ has at most length 1. If we go back to the endpoint types in (2) that we used to motivate this discussion, we have $\|T\| = 0$ and $\|S\| = \infty$, from which we deduce that endpoints with type S , like b in (1), are not safe to be used as messages.

4.3. Subtyping

The last notion we need before proceeding with the definition of the type system is a *subtyping relation* for endpoint types. Because of the close relationship between endpoint types and session types, the subtyping relation for endpoint types turns out to be a variant of that for session types [14]. However, the peculiar nature of exceptions has interesting consequences. The original subtyping relation for session types is based on the fundamental duality between input and output actions. In particular, it establishes that subtyping is *covariant* for external choices (inputs) and *contravariant* for internal ones (outputs). For example,

$$T = !a(S_1).T_1 \oplus !b(S_2).T_2 \leq !a(S_1).T_1 = S$$

is a valid subtyping relation between T and S . The underlying intuition is based on the usual principle of safe substitution of an endpoint of type S with another endpoint of type T . If a (well-typed) process is using an endpoint c of type S , then it can only send an a message on c . So, replacing the endpoint c with another one of type T , which allows both a and b messages to be sent, does not compromise communication safety. In a dual manner,

$$T' = ?a(S_1).T_1 \leq ?a(S_1).T_1 + ?b(S_2).T_2 = S'$$

is a valid subtyping relation between T' and S' . In this case, a (well-typed) process using an endpoint of type S' must be capable of handling (at least) a and b messages received from the endpoint. Replacing that endpoint with another one of type T' is safe because from the latter one only a messages can be received.

The covariance and contravariance properties of subtyping with respect to input and output operations follow from the duality of endpoint types associated with peer endpoints: when a process is entitled to send a message on an endpoint, the process using its peer must be ready to receive it, and vice-versa. By contrast, during a transaction, exceptions can be thrown on both peers of a channel. As a consequence, the two transaction types

$$\{e_1 : S_1, e_2 : S_2\} \llbracket T \quad \text{and} \quad \{e_1 : S_1\} \llbracket T$$

cannot be related. Indeed, if we had $\{e_1 : S_1, e_2 : S_2\} \ll T \leq \{e_1 : S_1\} \ll T$, then the process using the endpoint a with type $\{e_1 : S_1\} \ll T$ might not be prepared to handle the exception e_2 thrown by the process using the peer b of a . Similarly, the process using the peer endpoint b might be unable to handle the exception e_2 thrown on a if we had the opposite relation. In the end, because of the bi-directional nature of exceptions thrown during a transaction, subtyping must be *invariant* for transaction types.

We now proceed to define subtyping formally, extending it to possibly sealed endpoint types in the natural way:

Definition 4.3 (subtyping). *Subtyping* is the largest relation \leq such that $t \leq s$ implies either:

- $t = [t']$ and $s = [s']$ and $t' \leq s'$, or
- $t = s = \text{end}$, or
- $t = \{?m_i(T_i).T'_i\}_{i \in I}$ and $s = \{?m_i(S_i).S'_i\}_{i \in I \cup J}$ and $T_i \leq S_i$ and $T'_i \leq S'_i$ for every $i \in I$, or
- $t = \{!m_i(T_i).T'_i\}_{i \in I \cup J}$ and $s = \{!m_i(S_i).S'_i\}_{i \in I}$ and $S_i \leq T_i$ and $T'_i \leq S'_i$ for every $i \in I$, or
- either $(t = \{e_i : T_i\}_{i \in I} \ll T$ and $s = \{e_i : S_i\}_{i \in I} \ll S)$ or $(t = \{e_i : T_i\}_{i \in I} T$ and $s = \{e_i : S_i\}_{i \in I} S)$ and $T_i \leq S_i$ for every $i \in I$ and $T \leq S$, or
- $t = \ll T$ and $s = \ll S$ and $T \leq S$.

According to the definition of \leq , the covariance and contravariance properties for external and internal choices informally introduced earlier are extended to message argument types, in the usual manner. Observe that subtyping is always covariant with respect to continuations. It is easy to show that \leq is a pre-order that is contravariant with respect to duality:

Proposition 4.3. *The following properties hold:*

1. \leq is reflexive and transitive;
2. $T \leq S$ if and only if $\bar{S} \leq \bar{T}$.

Proof. See Gay and Hole [14]. Transaction types do not pose additional issues. \square

In Section 4.2 we have introduced a notion of weight that will be used in the type system for discriminating between safe and unsafe messages. Since the weight is computed on the (static) type of endpoints and subtyping allows for the substitution of endpoints with related but possibly different types, one important question arises whether subtyping and type weight are coherent with each other. This is indeed the case:

Proposition 4.4. $T \leq S$ implies $\|T\| \leq \|S\|$.

Proof. It is easy to see that $\mathcal{W} = \{(T, n) \mid \exists S : T \leq S \ \& \ S :: n\}$ is a coinductive weight bound. In particular, when T is an internal choice we have $T :: 0$ regardless of the number of branches in T . \square

4.4. Typing Processes

We can now proceed to defining a type system for processes. A *type environment* is a finite map $\Gamma = \{u_i : t_i\}_{i \in I}$ from names to types. We write $\text{dom}(\Gamma)$ for the domain of Γ , namely the set $\{u_i\}_{i \in I}$; we write Γ, Γ' for the union of Γ and Γ' when $\text{dom}(\Gamma) \cap \text{dom}(\Gamma') = \emptyset$; finally, we write $\Gamma \vdash u : t$ if $\Gamma(u) = t$. An *exception environment* $\tilde{\mathcal{E}}$ is a finite sequence $\mathcal{E}_1 \cdots \mathcal{E}_n$ of sets of exceptions. We write $e \in \tilde{\mathcal{E}}$ if $e \in \mathcal{E}_k$ for some $k \in \{1, \dots, n\}$. We say that a type t is *local*, written $\text{local}(t)$, if t is not sealed and has a null rank, namely $t = T$ for some T such that $\text{rank}(T) = 0$. Intuitively, a local type denotes an endpoint that can be modified (its type is not sealed) and is not involved in any transaction. We extend the notion of local types to type environments so that $\text{local}(\Gamma)$ holds if every type in the codomain of Γ is local.

The typing rules for processes are inductively defined in Table 8. Judgments have the form

$$\tilde{\mathcal{E}}; \Gamma \vdash P$$

denoting that process P is well typed in the exception environment $\tilde{\mathcal{E}}$ and type environment Γ . In particular, P can only throw exceptions that occur in $\tilde{\mathcal{E}}$. The type system makes use of a global process environment Σ associating process variables X with pairs $(\tilde{t}, \tilde{\mathcal{E}})$ containing the type of the parameters of X as well as the exception environment $\tilde{\mathcal{E}}$ in which X is supposed to be invoked. It is understood that the process environment Σ contains associations for all the global definitions D and that the judgment $\Sigma \vdash D$ defined by

$$\frac{\Sigma(X) = (\tilde{t}, \tilde{\mathcal{E}}) \quad \tilde{\mathcal{E}}; \tilde{u} : \tilde{t} \vdash P}{\Sigma \vdash X(\tilde{u}) \stackrel{\text{def}}{=} P}$$

$\frac{[\text{T-INACTION}]}{\tilde{\mathcal{E}}; \emptyset \vdash \text{done}}$	$\frac{[\text{T-THROW}]}{e \in \tilde{\mathcal{E}} \quad \tilde{\mathcal{E}}; \Gamma \vdash \text{throw } e}$	$\frac{[\text{T-CLOSE}]}{\tilde{\mathcal{E}}; \Gamma \vdash P \quad \tilde{\mathcal{E}}; \Gamma, u : \text{end} \vdash \text{close}(u).P}$
$\frac{[\text{T-INVOKE}]}{\Sigma(X) = (\tilde{s}, \tilde{\mathcal{E}}) \quad \tilde{t} \leq \tilde{s} \quad \tilde{\mathcal{E}}; \tilde{u} : \tilde{t} \vdash X\langle \tilde{u} \rangle}$	$\frac{[\text{T-OPEN}]}{\vdash T : 0 \quad \tilde{\mathcal{E}}; \Gamma, a : T, b : \bar{T} \vdash P \quad \tilde{\mathcal{E}}; \Gamma \vdash \text{open}(a, b).P}$	
$\frac{[\text{T-SEND}]}{k \in I \quad S \leq S_k \quad \ S\ < \infty \quad \tilde{\mathcal{E}}; \Gamma, u : T_k \vdash P \quad \tilde{\mathcal{E}}; \Gamma, u : \{\! m_i(S_i).T_i\}_{i \in I}, v : S \vdash u!m_k(v).P}$		$\frac{[\text{T-CHOICE}]}{\tilde{\mathcal{E}}; \Gamma \vdash P \quad \tilde{\mathcal{E}}; \Gamma \vdash Q \quad \tilde{\mathcal{E}}; \Gamma \vdash P \oplus Q}$
$\frac{[\text{T-RECEIVE}]}{S_i \leq S'_i \quad \tilde{\mathcal{E}}; \Gamma, u : T_i, x_i : S'_i \vdash P_i \quad \tilde{\mathcal{E}}; \Gamma, u : \{\! m_i(S_i).T_i\}_{i \in I} \vdash \sum_{i \in I \cup J} u?m_i(x_i).P_i}$		$\frac{[\text{T-PARALLEL}]}{\tilde{\mathcal{E}}; \Gamma_1 \vdash P \quad \tilde{\mathcal{E}}; \Gamma_2 \vdash Q \quad \tilde{\mathcal{E}}; \Gamma_1, \Gamma_2 \vdash P Q}$
$\frac{[\text{T-TRY}]}{\tilde{\mathcal{E}} \{e_j\}_{j \in J}; [\Gamma], \{u_i : T_i\}_{i \in I} \vdash P \quad \tilde{\mathcal{E}}; \Gamma, \{u_i : S_{ij}\}_{i \in I} \vdash Q_j \quad \tilde{\mathcal{E}}; \Gamma, \{u_i : \{e_j : S_{ij}\}_{j \in J} \llbracket T_i \rrbracket_{i \in I} \text{try}(\{u_i\}_{i \in I}) \{e_j : Q_j\}_{j \in J} P}$		
$\frac{[\text{T-COMMIT}]}{\text{local}(\Gamma_2) \quad \tilde{\mathcal{E}}; \Gamma_1, \{u_i : T_i\}_{i \in I}, \Gamma_2 \vdash P \quad \tilde{\mathcal{E}} \tilde{\mathcal{E}}; [\Gamma_1], \{u_i : \llbracket T_i \rrbracket_{i \in I}\}, \Gamma_2 \vdash \text{commit}(\{u_i\}_{i \in I}).P}$		

Table 8: Typing rules for processes.

holds. In particular, *all* of the free names of P must occur in its binding variable $X(\tilde{u})$.

We describe the typing rules for processes in the following paragraphs. Rule [T-INACTION] states that the idle process is well typed only in the empty type environment. This is a standard rule for linear type systems implying, in our case, that the terminated process has no leaks.

Rule [T-CLOSE] states that a process $\text{close}(u).P$ is well typed provided that u corresponds to an endpoint with type end , on which no further interaction is possible, and P is well typed in the remaining type environment.

Rule [T-OPEN] deals with the creation of a new channel, which is visible in the continuation process as two peer endpoints typed by dual endpoint types. The premise $\vdash T : 0$ means that newly created endpoints have no pending transactions on them.

Rule [T-SEND] states that a process $u!m(v).P$ is well typed if u is associated with an endpoint type T that permits the output of m messages. The type S of the argument v must be unsealed, finite-weight, and has to be a subtype of the expected type in the endpoint type. Finally, the continuation P must be well typed in a type environment where the endpoint u is typed according to the continuation T_k of T and the endpoint v is no longer visible. This models the fact that the ownership of v is transferred to the process that receives the message.

Rule [T-RECEIVE] deals with inputs: a process waiting for a message from an endpoint $u : \{?m_i(S_i).T_i\}_{i \in I}$ is well typed if it can deal with all of the messages m_i . The continuation processes may use the endpoint u according to the endpoint type T_i and can access the message argument x_i of some supertype S'_i of S_i .

Rules [T-CHOICE] and [T-PARALLEL] are standard. In the latter, the type environment is split into two disjoint environments to type the processes being composed.

Rule [T-INVOKE] declares that a process invocation $X\langle\tilde{u}\rangle$ is well typed provided that the number and type of actual parameters \tilde{u} match the number and type of formal parameters in $\Sigma(X)$ and that the process is invoked in the correct exception environment. In this rule we write $\tilde{t} \leq \tilde{s}$ for the pointwise extension of \leq to sequences of types.

We now turn our attention to the constructs dealing with transactions and exceptions.

Rule [T-THROW] states that the process $\text{throw } e$ is well typed in *any* type environment, provided that it occurs within a transaction (the exception being thrown must be among the ones occurring in the exception environment). For this reason, the violation of linearity for the assumptions in the type environment

is only apparent, as control will be transferred at runtime to some appropriate exception handler.

Rule [T-TRY] deals with transaction initiations. All the endpoints in the decoration U must have a type allowing them to be involved in a transaction, while the types of other names are sealed so that P is prevented from using them until the transaction is terminated. Seals are not applied in the type environment for the handlers since they execute only if and when the transaction is aborted and therefore act outside of the transaction. Note that the admitted exceptions are augmented in P but not in Q .

Rule [T-COMMIT] is almost the dual of rule [T-TRY] and deals with transaction termination. Again, the endpoints in the decoration U must have a matching type in the context indicating the end of the transaction. Names with a sealed type must have been inherited from the context surrounding the transaction being terminated, so a seal is stripped off them in the continuation P . Names with a local type must have been created within the transaction being terminated, and can be used in the continuation as well. Note that the rightmost set \mathcal{E} in the exception environment is stripped off when type checking P , since P executes after the transaction has terminated hence outside of the scope where the exceptions in \mathcal{E} can be thrown.

Observe that the type system requires the endpoints specified in a `commit` process to be exactly the same as the ones in the corresponding `try`. This is a consequence of the properties of well-formed endpoint types: endpoints involved in a transaction have a type with a strictly positive rank (see [WF-INITIATE] in Table 7) meaning that they cannot be closed (because `end` has null rank) and they cannot be sent as messages (again because [WF-PREFIX] requires messages to have a type with null rank). For the same reason they cannot be qualified as local, because local endpoints have a type with a null rank. Therefore, the set $\{u_i\}_{i \in I}$ associated with a given `try` process will be exactly the same set associated with the corresponding `commit` process.

Example 4.2. Using the types defined in Example 4.1, the reader can verify that the bodies of the process definitions in Figure 3 for `GetNextDiskPath`, `Loop`, and `Finally` are respectively well typed according to the type environments

$$\begin{aligned}\Gamma_1 &= DS : ?\text{NewClientEndpoint}(T_{ns}).T_{DS}, ret : T_{ret} \\ \Gamma_2 &= ns : T'_{ns}, DS : T_{DS}, ret : T_{ret} \\ \Gamma_3 &= ns : \text{end}, DS : T_{DS}, ret : T_{ret}\end{aligned}$$

where

$$T'_{ns} = !\text{Register}(T_{imp}).(?\text{AckRegister}().)\text{end} + ?\text{NakRegister}(T_{imp}).T_{ns}$$

is an appropriate residual of the unfolding of T_{ns} . Note the role played by subtyping in this example: ret is used according to the type `!Result(T_{DS}).end` in `Finally` and according to the type `!SetService(T_{exp}).!Result(T_{DS}).end` in `Loop`. Since T_{ret} is a subtype of both these types, ret can be passed to `Finally` and `Loop` thanks to the subtyping relation in [T-INVOKe]. ■

4.5. Typing the Heap

The typing rules in Table 8 are not sufficient for proving the soundness of the type system, because they are solely concerned with the static syntax of processes. At runtime, we must take into account running transaction processes (see Tables 2 and 4) as well as the heap. Indeed, since inter-process communication relies on heap-allocated structures, several properties of well-behaved processes depend on properties of the heap saying that its content is consistent with a given type environment. In this section and in the following one we develop a type system for the runtime components of our process language. We remark that the programmer is solely concerned with the typing rules for static processes presented in Section 4.4, while the technical material presented hereafter, which builds on and extends the previous one, is only required for proving that the type system is sound.

Just as we have type checked a process P against a type environment that associates types with the names occurring in P , we also need to check that the heap is consistent with respect to the same environment. This leads to a notion of well-typed heap that we develop in this section. More precisely, well-typedness of a heap μ is checked with respect to a pair $\Gamma_0; \Gamma$ of type environments: the context Γ_0, Γ must provide type information for *all* the allocated structures in μ (that is, $\text{dom}(\Gamma_0, \Gamma) = \text{dom}(\mu)$); the splitting $\Gamma_0; \Gamma$ distinguishes the pointers in $\text{dom}(\Gamma)$ from the pointers in $\text{dom}(\Gamma_0)$ so that Γ contains the *roots* of μ , namely the pointers that are not referenced from any endpoint structure in the heap, while Γ_0 contains pointers that are referenced from some endpoint structure.

Among the properties that must be enforced is the complementarity between the endpoint types associated with peer endpoints. This notion of complementarity does not coincide with duality because the communication model is asynchronous: since messages can accumulate in the queue of an endpoint before they are received, the types of peer endpoints can be misaligned. The two peers are guaranteed to have dual types only when their queues are both empty. In general, we need to compute the actual endpoint type of an endpoint by taking into account the messages in its queue. To this aim we introduce a $\text{tail}(\cdot, \cdot)$ function for

endpoint types such that

$$\text{tail}(T, m_1(S_1) \cdots m_n(S_n)) = T'$$

indicates that messages having tag m_i and an argument of type S_i can be received in the specified order from an endpoint with type T , which can be used according to type T' thereafter. The function is inductively defined by the following rules:

$$\begin{aligned} & \text{tail}(T, \varepsilon) = T \\ & \frac{k \in I \quad S \leq S_k}{\text{tail}(\{?m_i(S_i).T_i\}_{i \in I}, m_k(S)) = T_k} \quad \frac{\text{tail}(T, m(S)) = T'}{\text{tail}(\{e_i : S_i\}_{i \in I} T, m(S)) = T'} \\ & \frac{\text{tail}(T, m_1(S_1)) = T' \quad \text{tail}(T', m_2(S_2) \cdots m_n(S_n)) = T''}{\text{tail}(T, m_1(S_1)m_2(S_2) \cdots m_n(S_n)) = T''} \end{aligned}$$

Note that $\text{tail}(T, m(S))$ is undefined when $T = \text{end}$ or T is an internal choice or T denotes the initiation or the termination of a transaction. This will enforce the property that the queue of endpoints having these types must be empty. In the particular case of transaction initiation, this makes sure that, if an exception is thrown, heap restoration simply amounts to emptying the queues of the endpoints involved in the transaction (mechanism (A) in Section 2). The fact that the queues of the endpoints involved in the transaction are guaranteed to be empty at the end of a transaction is solely motivated by our notion of duality (Definition 4.1), which demands a perfect correspondence between the actions on such endpoints during a transaction. In principle, it would be possible to relax duality in such a way that a message sent within a transaction is received only after the transaction is terminated. However, it would still be necessary for the receive operation to first wait for the actual termination of the transaction, for otherwise the soundness of the transaction would be compromised. This means that this increased flexibility in the syntax of programs would bear no concrete advantage in their semantics.

We now have all the notions to express the well-typedness of a heap μ with respect to a pair $\Gamma_0; \Gamma$ of type environments.

Definition 4.4 (well-typed heap). Let $\text{dom}(\Gamma_0) \cap \text{dom}(\Gamma) = \emptyset$. We write $\Gamma_0; \Gamma \Vdash \mu$ if all of the following conditions hold:

1. If $a \mapsto [b, \Omega] \in \mu$ and $b \mapsto [a, \Omega'] \in \mu$, then either $\Omega = \varepsilon$ or $\Omega' = \varepsilon$.

2. If $a \mapsto [b, \mathfrak{m}_1(c_1) :: \dots :: \mathfrak{m}_n(c_n)] \in \mu$, then

$$\text{tail}(T, \mathfrak{m}_1(S_1) \cdots \mathfrak{m}_n(S_n)) = S$$

where $\Gamma_0, \Gamma \vdash a : [\dots[T]\dots]$ and $\Gamma_0 \vdash c_i : S_i$ and $\|S_i\| < \infty$ and $\vdash S_i : 0$ for $1 \leq i \leq n$ and $b \mapsto [a, \varepsilon] \in \mu$ implies $\Gamma_0, \Gamma \vdash b : [\dots[\bar{S}]\dots]$ and $b \notin \text{dom}(\mu)$ implies $S = \text{end}$.

3. $\text{dom}(\mu) = \text{dom}(\Gamma_0, \Gamma) = \mu\text{-reach}(\text{dom}(\Gamma))$;

4. $A \cap B = \emptyset$ implies $\mu\text{-reach}(A) \cap \mu\text{-reach}(B) = \emptyset$ for every $A, B \subseteq \text{dom}(\Gamma)$.

Condition (1) requires that at least one of the queues of peer endpoints in a well-typed heap is empty. This invariant corresponds to half-duplex communication and is ensured by duality of endpoint types associated with peer endpoints, since a well-typed process cannot send messages on an endpoint until it has read all the pending messages from the corresponding queue (we will see in Example 4.4 how to safely circumvent half-duplex communication thanks to transactions). Condition (2) requires that the content of the queue of an endpoint must be consistent with the type of the endpoint, in the sense that the messages in the queue have the expected tag and an argument with the expected type. In addition, the endpoint types of message arguments must all have finite weight and null rank. Finally, the endpoint types of peer endpoints are dual of each other, modulo the content of the non-empty queue. Condition (3) states that the type environment Γ_0, Γ must specify a type for all of the allocated objects in the heap and, in addition, every object (located at) a in the heap must be reachable from a root $b \in \text{dom}(\Gamma)$. Finally, condition (4) requires the uniqueness of the root for every allocated object. Overall, since the roots are distributed linearly among the processes of the system, conditions (3) and (4) guarantee that every allocated object belongs to one and only one process.

There are a few subtleties regarding conditions (1) and (2) and the fact that, in condition (2), the property $b \mapsto [a, \varepsilon] \in \mu$ is the head of an implication. First of all, condition (2) must hold for both peers of a channel, therefore if a is the peer with the empty queue ($n = 0$) while b has messages in its queue, then the type of a is not necessarily the dual of the type of b . The correct dual correspondence is checked when the symmetric pair of endpoints is considered. Second, it is possible that at some point only one endpoint of a channel is allocated. For example, the well-typed process `open(a, b).close(b).close(a)` reduces to `close(a)` in a configuration where the heap contains only $a \mapsto [b, \varepsilon]$. When this happens,

<div style="border-bottom: 1px solid black; padding-bottom: 5px;"> <div style="display: flex; justify-content: space-between;"> <div style="width: 45%;"> <p>[T-RUNNING PROCESS]</p> $\frac{\Gamma_0; \Gamma_R, \Gamma \Vdash \mu \quad \tilde{\mathcal{E}}; \Gamma \vdash P}{\tilde{\mathcal{E}}; \Gamma_0; \Gamma_R; \Gamma \vdash \mu \ ; P}$ </div> <div style="width: 45%;"> <p>[T-RUNNING PARALLEL]</p> $\frac{\tilde{\mathcal{E}}; \Gamma_0; \Gamma_R, \Gamma_2; \Gamma_1 \vdash \mu \ ; P \quad \tilde{\mathcal{E}}; \Gamma_0; \Gamma_R, \Gamma_1; \Gamma_2 \vdash \mu \ ; Q}{\tilde{\mathcal{E}}; \Gamma_0; \Gamma_R; \Gamma_1, \Gamma_2 \vdash \mu \ ; P \mid Q}$ </div> </div> </div> <div style="padding-top: 10px;"> <p>[T-RUNNING TRANSACTION]</p> $\frac{\begin{array}{c} \mu\text{-balanced}(\{a_i : S_{ij}\}_{i \in I})^{(j \in J)} \quad \mu\text{-balanced}(B) \quad \text{local}(\Gamma_2) \\ \{a_i\}_{i \in I} \cup B = \mu\text{-reach}(\{a_i\}_{i \in I} \cup \text{dom}(\Gamma_2)) \end{array} \quad \tilde{\mathcal{E}}\{e_j\}_{j \in J}; \Gamma_0; \Gamma_R; [\Gamma_1], \{a_i : T_i\}_{i \in I}, \Gamma_2 \vdash \mu \ ; P \quad \tilde{\mathcal{E}}; \Gamma_1, \{a_i : S_{ij}\}_{i \in I} \vdash Q_j^{(j \in J)}}{\tilde{\mathcal{E}}; \Gamma_0; \Gamma_R; \Gamma_1, \{a_i : \{e_j : S_{ij}\}_{j \in J} T_i\}_{i \in I}, \Gamma_2 \vdash \mu \ ; \langle \{a_i\}_{i \in I}, B, \{e_j : Q_j\}_{j \in J} P \rangle}$ </div>
--

Table 9: Typing rules for configurations.

the type of the remaining endpoint forbids any send operation (last property of condition (2)). Note that condition (1) is not implied by condition (2) and both conditions are necessary.

4.6. Typing Configurations

Table 9 defines typing rules for configurations $\mu \ ; P$ as an extension of the typing rules for processes. Judgments have the form

$$\tilde{\mathcal{E}}; \Gamma_0; \Gamma_R; \Gamma \vdash \mu \ ; P$$

and state that the configuration $\mu \ ; P$ is well typed with respect to the exception environment $\tilde{\mathcal{E}}$ and the triple $\Gamma_0; \Gamma_R; \Gamma$ of type environments. Intuitively, Γ is the type environment used to type check P , Γ_R is the type environment describing the type of root pointers owned by processes that are running in parallel with P , and Γ_0 describes the type of pointers that occur in some queue.

Rule [T-RUNNING PROCESS] lifts well-typed processes to well-typed configurations by requiring the heap to be well typed with respect to the pair of environments $\Gamma_0; \Gamma_R, \Gamma$ where Γ_R, Γ represents the whole set of roots obtained from those owned by the process being typed (in Γ) and those owned by processes in parallel with it (in Γ_R).

Rule [T-RUNNING PARALLEL] is similar to [T-PARALLEL], except that it deals with three type environments which are appropriately rearranged for keeping track of the roots of the heap.

Rule [T-RUNNING TRANSACTION] captures the basic properties regarding running transactions $\langle \{a_i\}_{i \in I}, B, \{e_j : Q_j\}_{j \in J} P \rangle$, which we describe here. The rule

makes use of a balancing predicate over type environments that generalizes the notion of balancing for sets of pointers (Definition 3.2):

Definition 4.5. We say that Γ is *balanced* in μ , written μ -balanced(Γ), if $a \in \text{dom}(\Gamma)$ and $a \xleftrightarrow{\mu} b$ imply $b \in \text{dom}(\Gamma)$ and $\Gamma(a) = \overline{\Gamma(b)}$.

First of all, it must be possible to partition the type environment into three parts Γ_1 , $\{a_i : \{e_j : S_{ij}\}_{j \in J} T_i\}_{i \in I}$, and Γ_2 such that the environment Γ_1 corresponds to the endpoints owned by P but which are not involved in the transaction. Consequently, the types of these endpoints are sealed in the judgment corresponding to the typing of P . The environment $\{a_i : \{e_j : S_{ij}\}_{j \in J} T_i\}_{i \in I}$ corresponds to the endpoints involved in the transaction (the first component of the running transaction process), and their type indicates that the transaction is in progress. The environment Γ_2 corresponds to the endpoints that have been allocated inside the transaction. Their type is not sealed in the judgment corresponding to the typing of P . The premises μ -balanced($\{a_i : S_{ij}\}_{i \in I}$) for every $j \in J$ and μ -balanced(B) indicate that the set of all the endpoints to which P has full access is balanced. Therefore, the transaction operates in a closed scope and cannot have “side effects” from the point of view of other processes. The first premise indicates, in addition, that the types S_{ij} associated with peer endpoints are dual of each other (this property is a consequence of well-typedness of the heap before the transaction initiates, but it must be explicitly stated in [T-RUNNING TRANSACTION] where the heap is checked against a type environment where the S_i ’s do not occur any more). The premise $\text{local}(\Gamma_2)$ identifies the Γ_2 partition of the context corresponding to the endpoints that have been created inside the transaction. The premise $\{a_i\}_{i \in I} \cup B = \mu\text{-reach}(\{a_i\}_{i \in I} \cup \text{dom}(\Gamma_2))$ states that all the endpoints allocated within the transaction have not escaped the scope of the transaction. The last two premises correspond to the premises of rule [T-TRY]. In particular, note that the exception environment is properly augmented when typing the body of the transaction.

Since running transaction processes appear only at runtime as the result of [R-START TRANSACTION] reductions, they can never occur behind a prefix and therefore the three rules in Table 9 cover all possible forms of runtime configurations.

4.7. Type Soundness

We can now formulate the two main results about our framework: well-typedness is preserved by reduction, and well-typed processes are well behaved. Subject

reduction takes into account the possibility that types in the environment may change as the process reduces, which is common in behavioral type theories.

Theorem 4.1 (subject reduction). *Let $\tilde{\mathcal{E}}; \Gamma_0; \Gamma_R; \Gamma \vdash \mu \wp P$ and $\mu \wp P \rightarrow \mu' \wp P'$. Then $\tilde{\mathcal{E}}; \Gamma'_0; \Gamma_R; \Gamma' \vdash \mu' \wp P'$ for some Γ'_0 and Γ' .*

Proof. See Appendix A. □

In fact, the proof of this theorem requires to specify a number of additional properties showing the precise relationship between Γ and Γ_0 (before the reduction) and Γ' and Γ'_0 (after the reduction).

Theorem 4.2 (safety). *Let $\emptyset \vdash_0 P$. Then P is well behaved.*

Proof. See Appendix B. □

We conclude this section with two examples showing how transaction types increase the expressiveness of session types by allowing the safe modeling of timeouts and mixed choices.

Example 4.3 (timeouts). In some cases it is possible and desirable to establish a timeout for a receive operation to succeed. Using transaction types and exceptions, it is easy to model timeouts in our process language. As an example, suppose one is interested in modeling a process

$$a?m(x).P + \tau.Q$$

which behaves as P as soon as it receives an m message from endpoint a and reduces to Q through an internal τ move if no message is received after some unspecified amount of time. This can be modeled by means of the process

$$\begin{aligned} & \text{try}(a) \{ \text{timeout} : Q \} \\ & \quad a?m(x).\text{commit}(a).P \\ & \quad | \text{try}(\{\}) \{ \text{ok} : \text{done} \} (\text{throw ok} \oplus \text{throw timeout}) \end{aligned}$$

where the nondeterministic choice between throwing `ok` or `timeout` is the abstract representation of the mechanism that activates the timeout. Note that the modeling uses exception propagation to trigger the timeout. The type associated with a is $\{ \text{timeout} : S \} \llbracket ?m(t) \rrbracket T$ where T describes the behavior of P on a if the m message is received within the timeout, while S describes the behavior of Q on a if the timeout expires. Note that the modeling in this example is not meant to suggest

an actual implementation of the timeout mechanism, but rather to describe its effect in abstract terms. In particular, the `m` message may have already been sent by the time the timeout expires (and the exception `timeout` is thrown). Yet, the sender of the message should be aware of this eventuality, it should be notified in case the timeout expires, and it should provide a suitable recovery action for this eventuality, possibly involving the resending of the `m` message. This is all guaranteed by the fact that the sender, which uses the peer endpoint of `a`, must do so according to the endpoint type $\{\text{timeout} : \bar{S}\} \llbracket !m(t). \bar{T} \rrbracket$. ■

Example 4.4 (mixed choices). The `Sing#` implementation allows the definition of contracts with so-called *mixed choices*, namely states in which there are two (or more) alternative operations involving both inputs and outputs. If mixed choices were allowed in our type language we would have, for example, endpoint types of the form $\{!a(t).T, ?b(s).S\}$ allowing either sending an `a` message or receiving a `b` message. Mixed choices break the half-duplex communication modality and are known to make protocols less robust and prone to deadlock [5, 6]. Yet, they can be safely modeled using transaction types as, for example

$$T_m = \{a : !a(t).T, b : ?b(s).S\} \llbracket \rrbracket \text{end}$$

The intuition is that a process behaving as T_m may throw either an `a` or a `b` exception to notify its party as to which operation (output an `a` message or input a `b` message) it will perform. An example of such process is

$$\text{try}(c) \{a : c!a(u).P, b : c?b(x).Q\} (\text{commit}(c).\text{close}(c) \mid X\langle u \rangle)$$

where

$$X(y) \stackrel{\text{def}}{=} X\langle y \rangle \oplus \text{throw } a$$

represents an internal computation that may eventually throw the `a` exception and cause the sending of the `a` message.

Below is part of the proof derivation showing that the process is well typed in the environment $c : T_m, u : t$. The branch related to the `b` exception has been omitted, but it is analogous to the one for the `a` exception; where necessary, the exception environment $\mathcal{E} = \{a, b\}$ is used:

$$\frac{\frac{\mathcal{E}; \emptyset \vdash \text{done}}{\mathcal{E}; c : \text{end} \vdash \text{close}(c)}}{\mathcal{E}; c : \llbracket \text{end} \vdash \text{commit}(c).\text{close}(c) \quad \mathcal{E}; u : [t] \vdash X\langle u \rangle} \quad \frac{\vdots}{c : T \vdash P}
}{\mathcal{E}; c : \llbracket \text{end}, u : [t] \vdash \text{commit}(c).\text{close}(c) \mid X\langle u \rangle \quad c : !a(t).T, u : t \vdash c!a(u).P \quad \vdots}
\frac{}{c : T_m, u : t \vdash \text{try}(c)\{a : c!a(u).P, b : c?b(x).Q\}(\text{commit}(c).\text{close}(c) \mid X\langle u \rangle)}$$

Finally,

$$\frac{\mathcal{E}; u : [t] \vdash X\langle u \rangle \quad \mathcal{E}; u : [t] \vdash \text{throw } a}{\mathcal{E}; u : [t] \vdash X\langle u \rangle \oplus \text{throw } a}$$

is the proof tree proving that the definition of $X\langle u \rangle$ is well typed. \blacksquare

5. Related Work

Our research fits in the broad spectrum of works developing hybrid static/dynamic techniques for the controlled management of the heap, and relies on invariants regarding the configuration of heap-allocated structures for enabling the efficient implementation of transactional mechanisms. Controlled heap management has fostered the development of a wide spectrum of techniques aimed at the most diverse purposes, of which we provide here a small account.

Pure functional programming languages are excellent candidates for the implementation of implicitly parallel computations. Parallelism is usually achieved by allowing multiple processing units to independently reduce disjoint parts of a program (represented as a graph) stored in a (possibly virtual) shared memory. Both hardware [15] and software [16] architectures have been explored. In these architectures, the crucial aspect is to achieve an optimal distribution of tasks, taking into account the fact that each processing unit often has its own local memory where subgraphs to be reduced must be copied, and that for improved efficiency it is necessary to take into proper account locality properties of the program [17].

Type-based approaches for enforcing and reasoning on properties of the heap are also popular. The seminal work by Tofte and Talpin [18] describes an *effect type system* for region-based memory management that allows for efficient allocation and deallocation of related heap structures. Interestingly, effects can be seen as primitive forms of behavioral types. There exist type systems with resource annotations that allow computing bounds on the heap space usage of functions and methods. Examples of such type systems are given in [19] for (first-order)

functional programs and in [20] for object-oriented programs. These approaches are usually motivated by the need to provide firm guarantees in code meant to be executed in embedded systems with constrained resources. Such forms of analyses can be carried out also for untyped programs. For instance, Albert et al. [21] work directly on Java bytecode, while Hammond [22] studies bounded time/space semantics for a functional and concurrent language.

Suenaga et al. [23] study a type discipline for safe resource deallocation in concurrent programs with shared memory. The idea is to associate each shared resource with a *fractional ownership*, namely a rational number in the interval $[0, 1]$, which denotes the level of sharing of the resource: 0 means that the resource is not owned, 1 means that the resource is owned exclusively, while any intermediate rational number indicates a shared ownership of the resource (often constraining the operations allowed on it). The type system allows a thread to deallocate a resource only if the thread is the only owner of the resource *and* if all of the other resources contained in it have been deallocated or transferred to other threads. In our type system, a similar effect is achieved by the combination of the type rule for the `close` primitive (which allows for the deallocation of endpoints only when they are owned exclusively) and the notion of well-typed heap (which guarantees endpoints with type `end` to have an empty queue).

There are strong analogies between our endpoint types and *usage expressions* defined by Iwama et al. [24], which are used for controlling access to resources (such as files and memory) in a functional language with exceptions. Usage expressions are akin to behavioral types and describe the valid sequence of operations allowed on some resource. In particular, [24] defines a usage constructor $U_1;_E U_2$ where U_1 describes how the resource is accessed under “nominal” conditions while U_2 describes how the resource is accessed if an exception is thrown. The structure of this usage constructor resembles that of a transaction type $\{E : U_2\} \llbracket U_1$. Because usage expressions can be composed in sequence, there is no need to explicitly mark the points where the scope of an exception ends.

The present work continues the type-based formalization of Singularity OS described by Bono and Padovani [7]. To simplify the formal development of the present paper we dropped polymorphism and non-linear types from the type system in [7]. These are orthogonal features that are independent of exception handling and can be added without affecting the results we have presented here. A radically different approach for the static analysis of Singularity processes is explored by Villard et al. [6, 25], where the authors develop a proof system based on a variant of *separation logic*. Exceptions are not taken into account in these works.

The works more closely related to ours, and which we used as starting points, are by Carbone et al. [11] and Capecchi et al. [12]. In [11], which was the first to investigate exceptions in calculi for session-oriented interactions and to propose type constructs to describe explicitly, at the type level, the handling of exceptional events, it is possible to associate an exception handler to a whole (dyadic) session; [12] generalizes this idea to multiparty sessions (those with multiple participants) and allows the same channel to be involved, at different times, in different `try` blocks, each with its own dedicated exception handler. In both [11] and [12] it is possible that messages already present in channel queues at the time an exception occurs are forgotten. In our context, this would immediately yield memory leaks, which we avoid by keeping track of the resources allocated during a transaction and by restoring the system to a consistent configuration in case an exception is thrown. State restoration is made possible in our context because the system is not distributed and the heap is shared by the communicating processes. Neither [11] nor [12] consider session delegation, namely the communication of channels. Also, in [12] the type system forces inner `try` blocks to use a subset of the channels involved in outer blocks. We relax these restrictions and allow locally created channels to be involved in inner transactions. The most notable difference between [12] and the present work regards the semantics of exceptions in nested transactions: in [12], an exception thrown in one transaction is suspended as long as there are active handlers in the nested ones. This semantics is motivated by the observation that, in a distributed setting, it may be desirable to complete the execution of potentially critical handlers before outermost handlers take control. Our semantics allows handlers of outer transactions to take control at any time following the throwing of an exception. As a consequence, more constrained policies, such as the one adopted in [12], can be implemented without invalidating the results presented in our work.

The recent interest on Web services has spawned a number of works investigating (long running) transactions in a distributed setting; a detailed survey with many references is provided by Ferreira et al. [26]. In our context, the components Q_i of a process $\langle A, B, \{e_i : Q_i\}_{i \in I} P \rangle$ are analogous to *compensation handlers*. The main difference between our handlers and those used for compensations is that, in the latter case, it is usually made the assumption that it is not possible to restore the state of the system as it was at the beginning of the transaction. In our case, state restoration is made possible by the fact that the system is local and all the interactions occur through shared memory. In this context, we can rely on some native support from the runtime system to properly cleanup the state of the system and avoid memory leaks.

The operational semantics of exceptions and exception handling in the present paper has been loosely inspired by that of Haskell memory transactions described by Harris et al. [27]. In particular, our semantics describes *what* happens when an exception is thrown but not *how* exception notification and state restoration are implemented. In this sense our semantics is somewhat more abstract than the semantics given in similar works [12]. The semantics of [27] uses a clever combination of small- and big-step reduction rules and is even more abstract than ours. Technically, having an abstract semantics is an advantage because it allows the meaning of programs to be expressed more concisely and in an implementation-independent way. However, because the big-step semantics is affected by diverging programs, it is more appropriate in a functional setting where non-termination is typically considered as a misbehavior. In our context, where non-terminating processes are useful, we had to resort to a more detailed operational semantics that dynamically keeps track of the allocated memory within a transaction.

Donnelly and Fluet [28] put forward a programming abstraction called *transactional events* for the modular composition of communication events into transactions with an all-or-nothing semantics. Their approach focuses on finding synchronization paths between threads communicating synchronously, while in our case transactions are required for preserving type consistency of endpoints and for undoing the effects of asynchronous communication.

Inadequacy of the standard error handling mechanisms provided by mainstream programming languages has already been recognized, even in sequential and communication-free scenarios. Weimer and Necula [29], Weimer [30], Weimer and Necula [31] develop a static analysis technique that spots error handling mistakes concerning proper resource release. Their technique is based on finite-state automata (in other words, a basic form of behavioral type) for keeping track of the state of resources along all possible execution paths. They also propose a more effective mechanism for preventing runtime errors. The basic idea is to accumulate compensation actions regarding resources on a *compensation stack* as resources are allocated. This technique closely resembles dynamic compensations in [26]. Because of their dynamic nature, compensation stacks do not provide any assistance as far as type consistency is concerned.

6. Conclusions and Future Work

We have formalized a core language of processes that communicate and synchronize through the copyless message passing paradigm and can throw exceptions. In this context, where the sharing of data and explicit memory allocation require

controlled policies on the ownership of heap-allocated objects, special care must be taken when exceptions are thrown to prevent communication errors (arising from misaligned states of channel endpoints) and memory leaks (resulting from messages forgotten in endpoint queues). We have studied a type system guaranteeing some safety properties, in particular that well-typed processes are free from communication errors and do not leak memory even in presence of (caught) exceptions. We have taken advantage of invariants guaranteed by the type system for taming the implementation costs of exception handling: the queues of endpoints involved in a transaction are guaranteed to be empty when the transaction starts, so that state restoration in case of exception simply means emptying such queues; also, only endpoints local to a transaction can be freed inside the transaction, so that state restoration in case of exception does not involve re-allocations.

The choice of `Sing#` as our reference language has been motivated by the fact that the Singularity code base provides concrete programming patterns that the formal model is supposed to cover (for example, previous works on exception handling for session-oriented languages by Carbone et al. [11], Capecchi et al. [12] do not consider delegation inside `try`-blocks, which instead is ubiquitous in `Sing#`). In addition, `Sing#` already accommodates channel contracts, which play a crucial role in our formalization. However, we claim that our approach is abstract enough to be applicable to other programming languages and paradigms, provided that suitable type information (possibly in the form of code annotations) is attached to channel endpoints.

To measure the practical impact of our mechanism of exception handling, one can take advantage of the availability of Singularity's source code for verifying whether the constraints imposed by the type discipline are reasonable in practice. Even without thorough investigations, however, we are able to provide some favorable arguments to our type discipline, in particular with respect to the weight and rank restrictions on types. As regards type weights, one has to consider that messages allocated on the exchange heap are explicitly managed by means of reference counting [2], which notoriously falls short in handling cyclic data structures, and that the finite-weight restriction on the type of communicated messages is just aimed at preventing cycles in the exchange heap. Regarding ranks, the subclass of well-formed, null-ranked types are just those in which transactions are properly balanced. In fact, the notion of well formedness arises solely because of our choice of modeling transactions using two matching constructs `try` and `commit` marking their beginning and end. Their balancing arises naturally in a structured language such as `Sing#`. We also plan to work on a prototype implementation of the exception handling so as to explore its practical costs.

Even if the type system has been tailored for a process calculus with a minimal set of critical features, it can be easily extended to incorporate commonly used programming constructs. Some restrictions of the type system can also be relaxed without endangering its soundness. For example, according to rule [T-SEND], only local endpoints (those that have an unsealed and null-ranked type) can be sent as messages inside transactions. This restriction results from the syntax of endpoint types (requiring that message arguments must have an unsealed type) and from rule [WF-PREFIX] regarding well-formed endpoint types (requiring that message argument types must have null rank). Endpoints with a sealed type cannot be accessed from within a transaction, because all the operations that modify the state of the heap require the access to endpoints with unsealed type. Therefore, we claim that endpoints with a sealed type are also safe to be sent as messages. In case an exception is thrown, the only thing that must be restored is their ownership at the beginning of the transaction. Nonetheless, the proof of this relaxed discipline seems to require a non-trivial modification of rule [T-RUNNING TRANSACTION], which is already quite elaborate in the present state, to account for the fact that the state of endpoints with a sealed type can change as the result of concurrent threads that execute *outside* of the transaction.

A. Proof of Theorem 4.1 (Subject Reduction)

The two following lemmas are standard and say that typing is preserved by structural congruence and by substitutions. In the case of substitutions, subtyping may be applied without compromising well-typedness.

Lemma A.1. *Let $\tilde{\mathcal{E}}; \Gamma \vdash P$ and $P \equiv Q$. Then $\tilde{\mathcal{E}}; \Gamma \vdash Q$.*

Proof. By case analysis on the derivation of $P \equiv Q$. □

Lemma A.2 (substitution). *If $\tilde{\mathcal{E}}; \Gamma, u : t \vdash P$ and $v \notin \text{dom}(\Gamma) \cup \text{bn}(P)$ and $s \leq t$, then $\tilde{\mathcal{E}}; \Gamma, v : s \vdash P\{v/u\}$.*

Proof. For notational simplicity we prove the result when $u = x$ and $v = a$. We proceed by induction on the derivation of $\tilde{\mathcal{E}}; \Gamma, x : t \vdash P$ and by cases on the last rule applied. We only prove a few interesting cases.

[T-CLOSE] In this case $P = \text{close}(u).Q$ and $\Gamma, x : t = \Gamma', u : \text{end}$ and $\tilde{\mathcal{E}}; \Gamma' \vdash Q$.

If $x \in \text{dom}(\Gamma')$, then $\Gamma' = \Gamma'', x : t$ and $\Gamma = \Gamma'', x : t, u : \text{end}$. By induction hypothesis we deduce $\tilde{\mathcal{E}}; \Gamma'', a : s \vdash Q\{a/x\}$. Since we know that $a \neq u$, from rule [T-CLOSE] we conclude $\tilde{\mathcal{E}}; \Gamma, a : s \vdash \text{close}(u).Q\{a/x\}$.

If $x = u$, then $\Gamma = \Gamma'$, $t = s = \text{end}$, and $Q = Q\{a/x\}$ because x does not occur free in Q . From rule [T-CLOSE] we conclude $\tilde{\mathcal{E}}; \Gamma, a : \text{end} \vdash \text{close}(a).Q\{a/x\}$.

[T-SEND] In this case $\Gamma, x : t = \Gamma', u : \{!m_i(S_i).T_i\}_{i \in I}, v : S'$ and $P = u!m_k(v).Q$ where $k \in I$ and $S' \leq S_k$ and $\|S'\| < \infty$ and $\tilde{\mathcal{E}}; \Gamma', u : T_k \vdash Q$.

If $x \in \text{dom}(\Gamma')$, then $\Gamma' = \Gamma'', x : t$ and $\Gamma = \Gamma'', u : \{!m_i(S_i).T_i\}_{i \in I}, v : S'$. From $\tilde{\mathcal{E}}; \Gamma', u : T_k \vdash Q$ and by induction hypothesis we obtain $\tilde{\mathcal{E}}; \Gamma'', a : s, u : T_k \vdash Q\{a/x\}$. Since $a \notin \text{dom}(\Gamma)$ we know that $a \notin \{u, v\}$ and from rule [T-SEND] we conclude $\tilde{\mathcal{E}}; \Gamma, a : s \vdash u!m_k(v).Q\{a/x\}$.

If $x = u$, then $\Gamma = \Gamma', v : S'$ and $t = \{!m_i(S_i).T_i\}_{i \in I}$ and $s = \{!m_i(S'_i).T'_i\}_{i \in I \cup J}$ and $S_i \leq S'_i$ and $T'_i \leq T_i$ for every $i \in I$. Then $S' \leq S_k \leq S'_k$. From $\tilde{\mathcal{E}}; \Gamma', u : T_k \vdash Q$ and the induction hypothesis we obtain $\tilde{\mathcal{E}}; \Gamma', a : T'_k \vdash Q\{a/x\}$. We conclude $\tilde{\mathcal{E}}; \Gamma', a : \{!m_i(S'_i).T'_i\}_{i \in I \cup J}, v : S' \vdash a!m_k(v).Q\{a/x\}$ with an application of rule [T-SEND].

If $x = v$, then $t = S'$ and $x \notin \text{dom}(\Gamma') \cup \{u\}$ and $Q = Q\{a/x\}$. From Proposition 4.4 we deduce $\|s\| \leq \|t\| < \infty$. We conclude with an application of rule [T-SEND].

[T-TRY] In this case $\Gamma, x : t = \Gamma', \{u_i : \{e_j : S_{ij}\}_{j \in J} \llbracket T_i \rrbracket\}_{i \in I}$ and $P = \text{try}(\{u_i\}_{i \in I}) \{e_j : Q_j\}_{j \in J} Q$ and $\tilde{\mathcal{E}}; \{e_j\}_{j \in J}; [\Gamma'], \{u_i : T_i\}_{i \in I} \vdash Q$ and $\tilde{\mathcal{E}}; \Gamma', \{u_i : S_{ij}\}_{i \in I} \vdash Q_j$ for every $j \in J$.

If $x \in \text{dom}(\Gamma')$, then $\Gamma' = \Gamma'', x : t$ for some Γ'' . From the hypothesis $s \leq t$ we deduce $[s] \leq [t]$. From $\tilde{\mathcal{E}}; \{e_j\}_{j \in J}; [\Gamma'], \{u_i : T_i\}_{i \in I} \vdash Q$ and by induction hypothesis we deduce $\tilde{\mathcal{E}}; \{e_j\}_{j \in J}; [\Gamma''], a : [s], \{u_i : T_i\}_{i \in I} \vdash Q\{a/x\}$. From $\tilde{\mathcal{E}}; \Gamma', \{u_i : S_{ij}\}_{i \in I} \vdash Q_j$ and by induction hypothesis we deduce $\tilde{\mathcal{E}}; \Gamma'', a : s, \{u_i : S_{ij}\}_{i \in I} \vdash Q_j\{a/x\}$ for every $j \in J$. We conclude with an application of rule [T-TRY].

If $x = u_k$ for some $k \in I$, then $t = \{e_j : S_{kj}\}_{j \in J} \llbracket T_k \rrbracket$ and $s = \{e_j : S'_{kj}\}_{j \in J} \llbracket T'_k \rrbracket$ where $S'_{kj} \leq S_{kj}$ for every $j \in J$ and $T'_k \leq T_k$. By induction hypothesis we deduce $\tilde{\mathcal{E}}; \{e_j\}_{j \in J}; [\Gamma'], \{u_i : T_i\}_{i \in I \setminus \{k\}}, a : T'_k \vdash Q\{a/x\}$ and $\tilde{\mathcal{E}}; \Gamma', \{u_i : S_{ij}\}_{i \in I \setminus \{k\}}, a : S'_{kj} \vdash Q_j\{a/x\}$ for every $j \in J$. We conclude with an application of rule [T-TRY]. \square

The next lemma connects well-typed configurations to well-typed heaps and shows the irrelevance of the first and second components Γ_0 and Γ_R in typing processes.

Lemma A.3. *Let $\tilde{\mathcal{E}}; \Gamma_0; \Gamma_R; \Gamma \vdash \mu \circlearrowright P$. Then:*

- (1) $\Gamma_0; \Gamma_R, \Gamma \Vdash \mu$;
- (2) $\Gamma'_0; \Gamma'_R, \Gamma \Vdash \mu'$ implies $\tilde{\mathcal{E}}; \Gamma'_0; \Gamma'_R; \Gamma \vdash \mu' \circlearrowright P$.

Proof. By induction on the derivation of $\Gamma_0; \Gamma_R; \Gamma \vdash \mu \circledast P$ and by cases on the last rule applied. The only interesting case is [T-RUNNING TRANSACTION], from which we deduce:

- $P = \langle \{a_i\}_{i \in I}, B, \{e_j : R_j\}_{j \in J} Q \rangle$;
- $\Gamma = \Gamma_1, \{a_i : \{e_j : S_{ij}\}_{j \in J} T_i\}_{i \in I}, \Gamma_2$;
- μ -balanced($\{a_i : S_{ij}\}_{i \in I}$) for every $j \in J$;
- $\tilde{\mathcal{E}}\{e_j\}_{j \in J}; \Gamma_0; \Gamma_R; [\Gamma_1], \{a_i : T_i\}_{i \in I}, \Gamma_2 \vdash \mu \circledast Q$.

Regarding (1), from $\tilde{\mathcal{E}}\{e_j\}_{j \in J}; \Gamma_0; \Gamma_R; [\Gamma_1], \{a_i : T_i\}_{i \in I}, \Gamma_2 \vdash \mu \circledast Q$ by induction hypothesis we obtain $\Gamma_0; \Gamma_R, [\Gamma_1], \{a_i : T_i\}_{i \in I}, \Gamma_2 \Vdash \mu$ and then from the definition of function tail and μ -balanced($\{a_i : S_{ij}\}_{i \in I}$) for $j \in J$ we conclude $\Gamma_0; \Gamma_R, \Gamma \Vdash \mu$.

Regarding (2), from $\Gamma'_0; \Gamma'_R, \Gamma \Vdash \mu'$ and the definition of function tail we deduce $\Gamma'_0; \Gamma'_R, [\Gamma_1], \{a_i : T_i\}_{i \in I}, \Gamma_2 \Vdash \mu'$ and then from $\tilde{\mathcal{E}}\{e_j\}_{j \in J}; \Gamma_0; \Gamma_R; [\Gamma_1], \{a_i : T_i\}_{i \in I}, \Gamma_2 \vdash \mu \circledast Q$ by induction hypothesis we obtain $\tilde{\mathcal{E}}\{e_j\}_{j \in J}; \Gamma'_0; \Gamma'_R; [\Gamma_1], \{a_i : T_i\}_{i \in I}, \Gamma_2 \vdash \mu' \circledast Q$. We conclude with an application of rule [T-RUNNING TRANSACTION]. \square

The next lemma shows that the rule [T-RUNNING PARALLEL] for configurations subsumes the rule [T-PARALLEL] for processes. It is used for simplifying some cases in the proof of subject reduction (Theorem 4.1).

Lemma A.4. *If $\tilde{\mathcal{E}}; \Gamma_0; \Gamma_R; \Gamma \vdash \mu \circledast P_1 \mid P_2$ is derivable using [T-PARALLEL] and [T-RUNNING PROCESS], then it is also derivable using [T-RUNNING PROCESS] and [T-RUNNING PARALLEL].*

Proof. From $\tilde{\mathcal{E}}; \Gamma_0; \Gamma_R; \Gamma \vdash \mu \circledast P_1 \mid P_2$ and rule [T-RUNNING PROCESS] we obtain (H.1) $\tilde{\mathcal{E}}; \Gamma \vdash P_1 \mid P_2$ and (H.2) $\Gamma_0; \Gamma_R, \Gamma \Vdash \mu$. From (H.1) and rule [T-PARALLEL] we obtain (T.1) $\Gamma = \Gamma_1, \Gamma_2$ and (P.i) $\tilde{\mathcal{E}}; \Gamma_i \vdash P_i$ for $i \in \{1, 2\}$. From (H.2), (T.1), (P.i) and rule [T-RUNNING PROCESS] we obtain $\tilde{\mathcal{E}}; \Gamma_0; \Gamma_R, \Gamma_{3-i}; \Gamma_i \vdash \mu \circledast P_i$ for $i \in \{1, 2\}$. We conclude with an application of rule [T-RUNNING PARALLEL]. \square

When a new channel is allocated in the heap, it always comes as a pair of peer endpoints. This easy property is formalized thus:

Proposition A.1. *If $\mu \circledast P \rightarrow \mu' \circledast P'$ then μ' -balanced($\text{dom}(\mu') \setminus \text{dom}(\mu)$).*

Proof. Simple induction on the reduction that occurs. \square

The next result is the complete version of subject reduction proving that reductions preserve well typedness and showing the relationship between the contexts used for typing the two configurations involved. In particular, item (1) below says that reductions preserve well typedness; item (2) says that the rank of the type of endpoints are preserved by reductions, and that (de)allocated endpoints have a type with null rank; finally, item (3) formally expresses the concept of process isolation, saying that any portion of the heap that is not reachable by the process being reduced is not affected by the reduction. In this theorem and in its proof, we write $\text{unsealed}(\Gamma)$ if all types in the range of Γ are unsealed. We also write $\bigsqcup_{i \in I} \Gamma_i$ for the disjoint union of the contexts Γ_i .

Theorem 4.1. *Let $\tilde{\mathcal{E}}; \Gamma_0; \Gamma_R; [\Gamma_S], \Gamma \vdash \mu \circledast P$ where $\text{unsealed}(\Gamma)$ and $\mu \circledast P \rightarrow \mu' \circledast P'$. Then there exist Γ'_0 and Γ' such that:*

- (1) $\tilde{\mathcal{E}}; \Gamma'_0; \Gamma_R; [\Gamma_S], \Gamma' \vdash \mu' \circledast P'$, and
- (2) $\text{unsealed}(\Gamma')$ and for every $a \in \text{dom}(\Gamma) \cap \text{dom}(\Gamma')$ we have $\text{rank}(\Gamma(a)) = \text{rank}(\Gamma'(a))$ and for every $a \in \text{dom}(\Gamma) \setminus \text{dom}(\Gamma')$ we have $\text{rank}(\Gamma(a)) = 0$ and for every $a \in \text{dom}(\Gamma') \setminus \text{dom}(\Gamma)$ we have $\text{rank}(\Gamma'(a)) = 0$, and
- (3) for every $\Gamma_I \subseteq \Gamma_R, [\Gamma_S]$ such that μ -balanced(μ -reach($\text{dom}(\Gamma_I, \Gamma)$)) we have

$$\mu\text{-reach}(\text{dom}(\Gamma_R, \Gamma_S) \setminus \text{dom}(\Gamma_I)) = \mu'\text{-reach}(\text{dom}(\Gamma_R, \Gamma_S) \setminus \text{dom}(\Gamma_I)).$$

Proof. By induction on the derivation of $\mu \circledast P \rightarrow \mu' \circledast P'$ and by cases on the last rule applied. We omit trivial and symmetric cases.

[R-OPEN] Then $P = \text{open}(a, b).P'$ and $\mu' = \mu, a \mapsto [b, \varepsilon], b \mapsto [a, \varepsilon]$. From rule [T-RUNNING PROCESS] we obtain:

- (H.1) $\tilde{\mathcal{E}}; [\Gamma_S], \Gamma \vdash \text{open}(a, b).P'$;
- (H.2) $\Gamma_0; \Gamma_R, [\Gamma_S], \Gamma \Vdash \mu$.

From (H.1) and rule [T-OPEN] we obtain:

- $\vdash T : 0$;
- (C.1) $\tilde{\mathcal{E}}; [\Gamma_S], \Gamma, a : T, b : \bar{T} \vdash P'$.

Let $\Gamma'_0 = \Gamma_0$ and $\Gamma' = \Gamma, a : T, b : \bar{T}$. The proof of (C.2) $\Gamma'_0; \Gamma_R, [\Gamma_S], \Gamma' \Vdash \mu'$ is trivial. From (C.1), (C.2) and [T-RUNNING PROCESS] we obtain (1). We conclude by noting that items (2) and (3) hold trivially.

[R-CLOSE] In this case $P = \text{close}(a).P'$ and $\mu = \mu', a \mapsto [b, \Omega]$. From rule [T-RUNNING PROCESS] we obtain:

- (H.1) $\tilde{\mathcal{E}}; [\Gamma_S], \Gamma \vdash \text{close}(a).P'$;
- (H.2) $\Gamma_0; \Gamma_R, [\Gamma_S], \Gamma \Vdash \mu', a \mapsto [b, \Omega]$.

From the hypothesis (H.1) and rule [T-CLOSE] we obtain:

- (L.1) $\Gamma = \Gamma', a : \text{end}$;
- (C.1) $\tilde{\mathcal{E}}; [\Gamma_S], \Gamma' \vdash P'$.

Let $\Gamma'_0 = \Gamma_0$. We only have to show that (C.2) $\Gamma'_0; \Gamma_R, [\Gamma_S], \Gamma' \Vdash \mu'$ and the only interesting case in Definition 4.4 is item 3. The relation $\text{dom}(\mu') = \text{dom}(\Gamma'_0, \Gamma_R, [\Gamma_S], \Gamma')$ is obvious. We need to prove $\text{dom}(\Gamma'_0, \Gamma_R, [\Gamma_S], \Gamma') = \mu'\text{-reach}(\text{dom}(\Gamma'_R, [\Gamma_S], \Gamma'))$. First we show that Ω is empty. Suppose by contradiction that this is not the case. Then the endpoint type associated with a before the reduction occurs must begin with an external choice or running transaction, which contradicts (L.1). So we obtain

$$\begin{aligned} \mu'\text{-reach}(\text{dom}(\Gamma_R, [\Gamma_S], \Gamma')) &= \mu\text{-reach}(\text{dom}(\Gamma_R, [\Gamma_S], \Gamma)) \setminus \{a\} \\ &= \text{dom}(\Gamma_0, \Gamma_R, [\Gamma_S], \Gamma) \setminus \{a\} \\ &= \text{dom}(\Gamma'_0, \Gamma_R, \Gamma') \end{aligned} \quad (\text{H.2})$$

From (C.1), (C.2) and [T-RUNNING PROCESS] we conclude (1). Item (2) is obvious while item (3) holds because $\mu\text{-reach}(\text{dom}(\Gamma_R, \Gamma_S)) = \mu'\text{-reach}(\text{dom}(\Gamma_R, \Gamma_S))$.

[R-PARALLEL] In this case $P = P_1 \mid P_2$ and $\mu \circledast P_1 \rightarrow \mu' \circledast P'_1$ and $P' = P'_1 \mid P_2$. By Lemma A.4 we can assume that $\tilde{\mathcal{E}}; \Gamma_0; \Gamma_R, [\Gamma_S], \Gamma \vdash \mu \circledast P$ was derived by an application of rule [T-RUNNING PARALLEL]. Then:

- $\Gamma = \Gamma_1, \Gamma_2$ and $\Gamma_S = \Gamma_{S1}, \Gamma_{S2}$;
- (P.i) $\tilde{\mathcal{E}}; \Gamma_0; \Gamma_R, [\Gamma_{S3-i}], \Gamma_{3-i}; [\Gamma_{Si}], \Gamma_i \vdash \mu \circledast P_i$ for $i \in \{1, 2\}$.

From (P.1) by induction hypothesis we deduce that there exist Γ'_0 and Γ'_1 such that:

- (1') $\tilde{\mathcal{E}}; \Gamma'_0; \Gamma_R, [\Gamma_{S2}], \Gamma_2; [\Gamma_{S1}], \Gamma'_1 \vdash \mu' \circ P'_1$, and
- (2') $\text{unsealed}(\Gamma'_1)$ and for every $a \in \text{dom}(\Gamma_1) \cap \text{dom}(\Gamma'_1)$ we have $\text{rank}(\Gamma_1(a)) = \text{rank}(\Gamma'_1(a))$ and for every $a \in \text{dom}(\Gamma_1) \setminus \text{dom}(\Gamma'_1)$ we have $\text{rank}(\Gamma_1(a)) = 0$ and for every $a \in \text{dom}(\Gamma'_1) \setminus \text{dom}(\Gamma_1)$ we have $\text{rank}(\Gamma'_1(a)) = 0$, and
- (3') for every $\Gamma_I \subseteq \Gamma_R, [\Gamma_S], \Gamma_2$ such that μ -balanced(μ -reach($\text{dom}(\Gamma_I, \Gamma_1)$)) we have

$$\mu\text{-reach}(\text{dom}(\Gamma_R, \Gamma_S, \Gamma_2) \setminus \text{dom}(\Gamma_I)) = \mu'\text{-reach}(\text{dom}(\Gamma_R, \Gamma_S, \Gamma_2) \setminus \text{dom}(\Gamma_I)).$$

Let $\Gamma' = \Gamma'_1, \Gamma_2$. From (1') and Lemma A.3(1) we obtain (N.1) $\Gamma'_0; \Gamma_R, [\Gamma_S], \Gamma' \Vdash \mu'$. From (P.2), (N.1), and Lemma A.3(2) we deduce (P.2') $\tilde{\mathcal{E}}; \Gamma'_0; \Gamma_R, [\Gamma_{S1}], \Gamma'_1; [\Gamma_{S2}], \Gamma_2 \vdash \mu' \circ P_2$. From (1'), (P.2'), and rule [T-RUNNING PARALLEL] we conclude (1). Regarding (2), just notice that $\text{unsealed}(\Gamma')$. Regarding (3), let $\Gamma_J \subseteq \Gamma_R, [\Gamma_S]$ be such that μ -balanced(μ -reach($\text{dom}(\Gamma_J, \Gamma)$)). Take $\Gamma_I = \Gamma_J, \Gamma_2$ and observe that $\Gamma_I \subseteq \Gamma_R, [\Gamma_S], \Gamma_2$ and μ -balanced(μ -reach($\text{dom}(\Gamma_I, \Gamma_1)$)). From (3') we are able to deduce $\mu\text{-reach}(\text{dom}(\Gamma_R, \Gamma_S, \Gamma_2) \setminus \text{dom}(\Gamma_I)) = \mu'\text{-reach}(\text{dom}(\Gamma_R, \Gamma_S, \Gamma_2) \setminus \text{dom}(\Gamma_I))$ and we conclude (3) by observing that $\text{dom}(\Gamma_R, \Gamma_S) \setminus \text{dom}(\Gamma_J) = \text{dom}(\Gamma_R, \Gamma_S, \Gamma_2) \setminus \text{dom}(\Gamma_I)$.

[R-SEND] In this case $P = a!m(c).P'$ and $\mu = \mu'', a \mapsto [b, \Omega], b \mapsto [a, \Omega']$ and $\mu' = \mu'', a \mapsto [b, \Omega], b \mapsto [a, \Omega' :: m(c)]$. From [T-RUNNING PROCESS] we obtain:

- (H.1) $\tilde{\mathcal{E}}; [\Gamma_S], \Gamma \vdash a!m(c).P'$;
- (H.2) $\Gamma_0; \Gamma_R, [\Gamma_S], \Gamma \Vdash \mu'', a \mapsto [b, \Omega], b \mapsto [a, \Omega']$.

From (H.1) and rule [T-SEND] we deduce:

- (L.1) $\Gamma = \Gamma'', a : \{!m_i(S_i).T_i\}_{i \in I}, c : S$;
- $m = m_k$ for some $k \in I$;
- $S \leq S_k$ and $\|S\| < \infty$;
- (C.1) $\tilde{\mathcal{E}}; [\Gamma_S], \Gamma'', a : T_k \vdash P'$.

Let $\Gamma'_0 = \Gamma_0, c : S_k$ and $\Gamma' = \Gamma'', a : T_k$. We show (C.2) $\Gamma'_0; \Gamma_R, [\Gamma_S], \Gamma' \Vdash \mu'$ by proving the items of Definition 4.4 in order.

1. We only need to show that Ω is empty. Suppose by contradiction that this is not the case. Then the endpoint type associated with a before the reduction must begin with an external choice or running transaction, which contradicts (L.1).
2. Let $\Omega' = m_1(c_1) :: \dots :: m_p(c_p)$. From hypothesis (H.2) we deduce $\Gamma_0, \Gamma_R, [\Gamma_S], \Gamma \vdash b : T_b$ and $\Gamma_0, \Gamma_R, [\Gamma_S], \Gamma \vdash c_i : S'_i$ for $1 \leq i \leq p$ where

$$\text{tail}(T_b, m_1(S'_1) \cdots m_p(S'_p)) = \overline{\{!m_i(S_i).T_i\}_{i \in I}} = \{?m_i(S_i).\overline{T_i}\}_{i \in I}$$

and from $S \leq S_k$ we conclude $\overline{T_k} = \text{tail}(T_b, m_1(S'_1) \cdots m_p(S'_p))m(S)$.

3. From hypothesis (H.2) we have $\text{dom}(\mu) = \text{dom}(\Gamma_0, \Gamma_R, \Gamma)$ and for every $a' \in \text{dom}(\mu)$ there exists $b' \in \text{dom}(\Gamma_R, \Gamma_S, \Gamma)$ such that $a' \preceq_\mu b'$. Clearly $\text{dom}(\mu') = \text{dom}(\Gamma'_0, \Gamma_R, \Gamma_S, \Gamma')$ since $\text{dom}(\mu') = \text{dom}(\mu)$ and $\text{dom}(\Gamma'_0) \cup \text{dom}(\Gamma') = \text{dom}(\Gamma_0) \cup \text{dom}(\Gamma)$. Let $b \preceq_\mu b_0$ and $\Gamma_R, \Gamma_S, \Gamma \vdash b_0 : T_0$. We have $c \prec_{\mu'} b \preceq_{\mu'} b_0$, namely $c \preceq_{\mu'} b_0$. Now

$$\|S\| \leq \|S_k\| < \|\text{tail}(T_b, m_1(S'_1) \cdots m_p(S'_p))\| \leq \|T_b\| \leq \|T_0\|$$

and $\|S\| < \infty$, therefore $c \neq b_0$. We conclude $b_0 \in \text{dom}(\Gamma_R, \Gamma_S, \Gamma')$.

4. Immediate from hypothesis (H.2).

Item (2) holds trivially. Regarding item (3), let $\Gamma_I \subseteq \Gamma_R, [\Gamma_S]$ be such that μ -balanced(μ -reach($\text{dom}(\Gamma_I, \Gamma)$)). From $a \in \text{dom}(\Gamma)$ we are able to deduce that $b \in \mu$ -reach($\text{dom}(\Gamma_I, \Gamma)$) and $c \preceq_{\mu'} b$, therefore μ -reach($\text{dom}(\Gamma_R, \Gamma_S) \setminus \text{dom}(\Gamma_I)$) = μ' -reach($\text{dom}(\Gamma_R, \Gamma_S) \setminus \text{dom}(\Gamma_I)$).

[R-RECEIVE] In this case $P = \sum_{i \in I} a ?m_i(x_i).P_i$ and $\mu = \mu'', a \mapsto [b, m(c) :: \Omega]$ where $\Omega = m_1(c_1) :: \dots :: m_p(c_p)$ and $m = m_k$ for some $k \in I$ and $P' = P_k \{c/x_k\}$ and $\mu' = \mu'', a \mapsto [b, \Omega]$. From rule [T-RUNNING PROCESS] we obtain:

- (H.1) $\tilde{\mathcal{E}}; [\Gamma_S], \Gamma \vdash \sum_{i \in I} a ?m_i(x_i).P_i;$
- (H.2) $\Gamma_0; \Gamma_R, [\Gamma_S], \Gamma \Vdash \mu.$

From (H.1) and rule [T-RECEIVE] we obtain:

- $\Gamma = \Gamma'', a : \{?m_i(S_i).T_i\}_{i \in J}$ with $J \subseteq I;$
- (N.1) $\tilde{\mathcal{E}}; [\Gamma_S], \Gamma'', a : T_k, x_k : S_k \vdash P_k.$

From (H.2) we deduce $\Gamma_0 = \Gamma'_0, c : S$ where $S \leq S_k$ and $k \in J$. Let $\Gamma' = \Gamma'', a : T_k, c : S$. From (N.1) and Lemma A.2 we deduce (C.1) $\tilde{\mathcal{E}}; [\Gamma_S], \Gamma' \vdash P_k\{c/x_k\}$. Now we only have to show (C.2) $\Gamma'_0; \Gamma_R, [\Gamma_S], \Gamma' \Vdash \mu'$ and we do it by proving the items of Definition 4.4 in order.

1. Since the queue associated with a is not empty in μ , the queue associated with its peer endpoint b must be empty. The reduction does not change the queue associated with b , therefore condition (1) of Definition 4.4 is satisfied.
2. From (H.2) we deduce $\Gamma_0, \Gamma_R, [\Gamma_S], \Gamma \vdash b : T_b$ and
$$\bar{T}_b = \text{tail}(\{\text{?m}_i(S_i).T_i\}_{i \in J}, \text{m}(S)\text{m}_1(S'_1) \cdots \text{m}_p(S'_p)) = \text{tail}(T_k, \text{m}_1(S'_1) \cdots \text{m}_p(S'_p))$$
where $\Gamma_0, \Gamma_R, [\Gamma_S], \Gamma \vdash c_i : S'_i$ for $1 \leq i \leq p$.
3. Straightforward by definition of Γ'_0 and Γ' .
4. Immediate from (H.2).

Therefore, from (C.1), (C.2), and rule [T-RUNNING PROCESS] we conclude $\tilde{\mathcal{E}}; \Gamma'_0; \Gamma_R, [\Gamma_S], \Gamma' \vdash \mu' \wp P'$. Regarding (2), observe from (H.2) and condition (2) of Definition 4.4 that $\text{rank}(S_k) = 0$. Regarding (3), it suffices to observe that the only region of the heap that changes is the queue associated with a and that $a \notin \text{dom}(\Gamma_R, \Gamma_S)$.

[R-START TRANSACTION] In this case $P = \prod_{i \in I} \text{try}(A_i) \{e_j : Q_{ij}\}_{j \in J} P_i$ where μ -balanced($\bigcup_{i \in I} A_i$) and $P' = \langle \bigcup_{i \in I} A_i, \emptyset, \{e_j : \prod_{i \in I} Q_{ij}\}_{j \in J} \prod_{i \in I} P_i \rangle$ and $\mu' = \mu$. According to Lemma A.4 we can assume that $\tilde{\mathcal{E}}; \Gamma_0; \Gamma_R; [\Gamma_S], \Gamma \vdash \mu \wp P$ was derived by rule [T-RUNNING PARALLEL]. Then:

- (L.1) $\Gamma_S = \bigsqcup_{i \in I} \Gamma_{S_i}$ and $\Gamma = \bigsqcup_{i \in I} \Gamma_i$;
- (P.i) $\tilde{\mathcal{E}}; \Gamma_0; \Gamma_R, \bigsqcup_{j \in I \setminus \{i\}} [\Gamma_{S_j}], \bigsqcup_{j \in I \setminus \{i\}} \Gamma_j; [\Gamma_{S_i}], \Gamma_i \vdash \mu \wp \text{try}(A_i) \{e_j : Q_{ij}\}_{j \in J} P_i$ for every $i \in I$.

From (L.1), (P.i) and rule [T-RUNNING PROCESS] we obtain:

- (H.1) $\tilde{\mathcal{E}}; [\Gamma_{S_i}], \Gamma_i \vdash \text{try}(A_i) \{e_j : Q_{ij}\}_{j \in J} P_i$ where $\text{unsealed}(\Gamma_i)$ for every $i \in I$;
- (H.2) $\Gamma_0; \Gamma_R, [\Gamma_S], \Gamma \Vdash \mu$.

From (H.1) and rule [T-TRY] we obtain, for every $i \in I$:

- (L.2) $\Gamma_i = \Gamma'_i, \{a : \{e_j : S_{aj}\}_{j \in J} [T_a]_{a \in A_i};$
- $\tilde{\mathcal{E}} \{e_j\}_{j \in J}; [[\Gamma_{S_i}], \Gamma'_i], \{a : T_a\}_{a \in A_i} \vdash P_i;$
- $\tilde{\mathcal{E}}; [\Gamma_{S_i}], \Gamma'_i, \{a : S_{aj}\}_{a \in A_i} \vdash Q_{ij}$ for every $j \in J$.

By rule [T-PARALLEL] we derive:

- (P.1) $\tilde{\mathcal{E}} \{e_j\}_{j \in J}; [[\Gamma_S], \bigsqcup_{i \in I} \Gamma'_i], \{a : T_a\}_{i \in I, a \in A_i} \vdash \prod_{i \in I} P_i;$
- (P.2) $\tilde{\mathcal{E}}; [\Gamma_S], \bigsqcup_{i \in I} \Gamma'_i, \{a : S_{aj}\}_{i \in I, a \in A_i} \vdash \prod_{i \in I} Q_{ij}$ for every $j \in J$.

From (L.2) and the definition of tail it is easy to see that the queue associated with a is empty for every $i \in I$ and $a \in A_i$. Hence, from (H.2), (L.1), and (L.2) and the fact that μ -balanced($\bigcup_{i \in I} A_i$) we have μ -balanced($\{a : S_{aj}\}_{i \in I, a \in A_i}$) for every $j \in J$ and $\bigcup_{i \in I} A_i = \mu$ -reach($\bigcup_{i \in I} A_i$). Also, from (H.2) we have

- (H.2') $\Gamma_0; \Gamma_R, [[\Gamma_S], \bigsqcup_{i \in I} \Gamma'_i], \{a : T_a\}_{i \in I, a \in A_i} \Vdash \mu$.

From (H.2'), (P.1), and rule [T-RUNNING PROCESS] we obtain:

- (T.1) $\tilde{\mathcal{E}} \{e_j\}_{j \in J}; \Gamma_0; \Gamma_R; [[\Gamma_S], \bigsqcup_{i \in I} \Gamma'_i], \{a : T_a\}_{i \in I, a \in A_i} \vdash \mu \circ \prod_{i \in I} P_i$.

We conclude (1) with an application of rule [T-RUNNING TRANSACTION], (T.1), (P.2), and the facts proven above by taking $\Gamma'_0 = \Gamma_0$ and $\Gamma' = \bigsqcup_{i \in I} \Gamma'_i, \{a : \{e_j : S_{aj}\}_{j \in J} T_a\}_{i \in I, a \in A_i}$. Item (2) is trivial and (3) holds since $\mu' = \mu$.

[R-END TRANSACTION] Then $P = \langle A, B, \{e_j : Q_j\}_{j \in J} \prod_{i \in I} \text{commit}(A_i).P_i \rangle$ and $P' = \prod_{i \in I} P_i$ and $\mu' = \mu$. From rule [T-RUNNING TRANSACTION] we obtain:

- $\Gamma = \Gamma_1, \{a : \{e_j : S_{aj}\}_{j \in J} T_a\}_{a \in A}, \Gamma_2;$
- (T.1) $\tilde{\mathcal{E}} \{e_j\}_{j \in J}; \Gamma_0; \Gamma_R; [[\Gamma_S], \Gamma_1], \{a : T_a\}_{a \in A}, \Gamma_2 \vdash \mu \circ \prod_{i \in I} \text{commit}(A_i).P_i;$
- μ -balanced($\{a : S_{aj}\}_{a \in A}$) for every $j \in J$;
- μ -balanced(B);
- local(Γ_2);
- $A \cup B = \mu$ -reach($A \cup \text{dom}(\Gamma_2)$).

From (T.1), [T-RUNNING PARALLEL], and [T-RUNNING PROCESS] we deduce:

- $\Gamma_S = \bigsqcup_{i \in I} \Gamma_{Si}$ and $\Gamma_1 = \bigsqcup_{i \in I} \Gamma_{1i}$ and $\{a : T_a\}_{a \in A} = \bigsqcup_{i \in I} \{a : T_a\}_{a \in B_i}$ and $\Gamma_2 = \bigsqcup_{i \in I} \Gamma_{2i}$ where $\text{rank}(\Gamma_{2i}) = 0$ for every $i \in I$;
- $\tilde{\mathcal{E}}\{e_j\}_{j \in J}; [[\Gamma_{Si}], \Gamma_{1i}], \{a : T_a\}_{a \in B_i}, \Gamma_{2i} \vdash \text{commit}(A_i).P_i$ for every $i \in I$.

From rule [T-COMMIT] we deduce:

- $B_i = A_i$;
- $T_a = \llbracket T'_a \rrbracket$ for every $i \in I$ and $a \in A_i$;
- $\tilde{\mathcal{E}}; [\Gamma_{Si}], \Gamma_{1i}, \{a : T'_a\}_{a \in A_i}, \Gamma_{2i} \vdash P_i$ for every $i \in I$.

From rule [T-PARALLEL] we deduce (H.1) $\tilde{\mathcal{E}}; [\Gamma_S], \Gamma_1, \{a : T'_a\}_{i \in I, a \in A_i}, \Gamma_2 \vdash P'$. Let $\Gamma'_0 = \Gamma_0$ and $\Gamma' = \Gamma_1, \{a : T'_a\}_{a \in A}, \Gamma_2$. From $\tilde{\mathcal{E}}; \Gamma_0; \Gamma_R; [\Gamma_S], \Gamma \vdash \mu \circledast P$ and Lemma A.3(1) we obtain (H.2) $\Gamma_0; \Gamma_R, [\Gamma_S], \Gamma \Vdash \mu$. From (2) we deduce that the queues associated with the pointers $a \in A$ are empty, because Γ includes $\{a : \{e_j : S_{aj}\}_{j \in J} \llbracket T'_a \rrbracket\}_{a \in A}$. Hence we deduce $\Gamma_0; \Gamma_R, [\Gamma_S], \Gamma' \Vdash \mu'$. From (H.1) and rule [T-RUNNING PROCESS] we conclude (1). We observe that (2) holds since $\text{rank}(\{e_j : S_{aj}\}_{j \in J} T_a) = \text{rank}(T'_a)$ for $a \in A$ and (3) holds trivially since the heap has not changed.

[R-RUN TRANSACTION] In this case: $P = \langle A, B, \{e_j : R_j\}_{j \in J} Q \rangle$ and $P' = \langle A, B', \{e_j : R_j\}_{j \in J} Q' \rangle$ where $B' = \text{track}(B, \text{dom}(\mu), \text{dom}(\mu'))$ and $\mu \circledast Q \rightarrow \mu' \circledast Q'$. Let $A = \{a_i\}_{i \in I}$. From rule [T-RUNNING TRANSACTION] we deduce:

- $\Gamma = \Gamma_1, \{a_i : \{e_j : S_{ij}\}_{j \in J} T_i\}_{i \in I}, \Gamma_2$;
- (T.1) $\tilde{\mathcal{E}}\{e_j\}_{j \in J}; \Gamma_0; \Gamma_R; [[\Gamma_S], \Gamma_1], \{a_i : T_i\}_{i \in I}, \Gamma_2 \vdash \mu \circledast Q$;
- (T.2) $\tilde{\mathcal{E}}; \Gamma_1, \{a_i : S_{ij}\}_{i \in I} \vdash R_j$ for every $j \in J$;
- (T.3) μ -balanced($\{a_i : S_{ij}\}_{i \in I}$);
- (T.4) μ -balanced(B);
- (T.5) local(Γ_2);
- (T.6) $\{a_i\}_{i \in I} \cup B = \mu$ -reach($\{a_i\}_{i \in I} \cup \text{dom}(\Gamma_2)$).

Let $\Gamma_3 = \{a_i : T_i\}_{i \in I}, \Gamma_2$. From (T.1) and unsealed(Γ) by induction hypothesis we obtain that there exist Γ'_0 and Γ'_3 such that:

- (1') $\tilde{\mathcal{E}}; \Gamma'_0; \Gamma_R; [[\Gamma_S], \Gamma_1], \Gamma'_3 \vdash \mu' \circledast Q'$, and

- (2') unsealed(Γ'_3) and for every $a \in \text{dom}(\Gamma_3) \cap \text{dom}(\Gamma'_3)$ we have $\text{rank}(\Gamma_3(a)) = \text{rank}(\Gamma'_3(a))$ and for every $a \in \text{dom}(\Gamma_3) \setminus \text{dom}(\Gamma'_3)$ we have $\text{rank}(\Gamma_3(a)) = 0$ and for every $a \in \text{dom}(\Gamma'_3) \setminus \text{dom}(\Gamma_3)$ we have $\text{rank}(\Gamma'_3(a)) = 0$, and
- (3') for every $\Gamma_I \subseteq \Gamma_R, [[\Gamma_S], \Gamma_1]$ such that μ -balanced(μ -reach($\text{dom}(\Gamma_I, \Gamma_3)$)) we have

$$\mu\text{-reach}(\text{dom}(\Gamma_R, \Gamma_S, \Gamma_1) \setminus \text{dom}(\Gamma_I)) = \mu'\text{-reach}(\text{dom}(\Gamma_R, \Gamma_S, \Gamma_1) \setminus \text{dom}(\Gamma_I)).$$

Since $\text{rank}(T_i) > 0$ for all $i \in I$, from (2') we deduce that all the a_i 's are still in the environment for Q' therefore we have $\Gamma'_3 = \{a : T'_i\}_{i \in I}, \Gamma'_2$. Let $\Gamma' = \Gamma_1, \{a_i : \{e_j : S_{ij}\}_{j \in J} T'_i\}_{i \in I}, \Gamma'_2$.

Regarding (1), from (T.4) and Proposition A.1 we obtain (T.4') μ' -balanced(B'). From (2') we deduce (T.5') local(Γ'_2). In order to prove (T.6') $\{a_i\}_{i \in I} \cup B' = \mu'\text{-reach}(\{a_i\}_{i \in I} \cup \text{dom}(\Gamma'_2))$ observe that:

$$\begin{aligned}
(*) \text{ dom}(\mu') &= (\text{dom}(\mu) \cup (B' \setminus B)) \setminus (B \setminus B') && \text{by definition of } B' \\
&= (\mu\text{-reach}(\text{dom}(\Gamma_R, [\Gamma_S], \Gamma_1, \Gamma_2) \cup \{a_i\}_{i \in I}) \cup (B' \setminus B)) \setminus (B \setminus B') \\
&&& \text{by item (3) of Definition 4.4} \\
&= (\mu\text{-reach}(\text{dom}(\Gamma_R, [\Gamma_S], \Gamma_1))) \\
&\quad \uplus \mu\text{-reach}(\text{dom}(\Gamma_2) \cup \{a_i\}_{i \in I}) \cup (B' \setminus B) \setminus (B \setminus B') \\
&&& \text{by item (4) of Definition 4.4} \\
&= (\mu\text{-reach}(\text{dom}(\Gamma_R, [\Gamma_S], \Gamma_1)) \cup \{a_i\}_{i \in I} \cup B \cup (B' \setminus B)) \setminus (B \setminus B') && \text{by (T.6)} \\
&= \mu\text{-reach}(\text{dom}(\Gamma_R, [\Gamma_S], \Gamma_1)) \cup \{a_i\}_{i \in I} \cup B' && \text{by set theory}
\end{aligned}$$

where \uplus denotes disjoint union. In addition, from (1') and Lemma A.3(1) we obtain $\Gamma'_0; \Gamma_R; [[\Gamma_S], \Gamma_1], \Gamma'_3 \Vdash \mu'$, so we have:

$$\begin{aligned}
(**) \text{ dom}(\mu') &= \mu'\text{-reach}(\text{dom}(\Gamma_R, [\Gamma_S], \Gamma_1, \Gamma'_2) \cup \{a_i\}_{i \in I}) \\
&= \mu'\text{-reach}(\text{dom}(\Gamma_R, [\Gamma_S], \Gamma_1)) \uplus \mu'\text{-reach}(\text{dom}(\Gamma'_2) \cup \{a_i\}_{i \in I})
\end{aligned}$$

where the two equalities are respectively justified by items (3) and (4) of Definition 4.4. From (T.3), (T.4) we obtain μ -balanced($\{a_i\}_{i \in I} \cup B$), and then from (T.6) we get μ -balanced(μ -reach($\{a_i\}_{i \in I} \cup \text{dom}(\Gamma_2)$)). Therefore, by taking $\Gamma_I = \emptyset$ in (3') we obtain $\mu\text{-reach}(\text{dom}(\Gamma_R, [\Gamma_S], \Gamma_1)) = \mu'\text{-reach}(\text{dom}(\Gamma_R, [\Gamma_S], \Gamma_1))$ and then from (*) and (**) we obtain (T.6'). We conclude this part of the proof with an application of rule [T-RUNNING TRANSACTION] to (1'), (T.2), (T.3), (T.4'), (T.5') and (T.6').

Regarding (2), we conclude from (2') and rule [WF-RUN] of Table 7. Regarding (3), let $\Gamma_J \subseteq \Gamma_R, [\Gamma_S]$ be such that μ -balanced(μ -reach($\text{dom}(\Gamma_J, \Gamma)$)). Now take $\Gamma_I = \Gamma_J, [\Gamma_1]$ and observe that μ -balanced(μ -reach($\text{dom}(\Gamma_I, \Gamma_3)$)). From (3') we conclude

$$\mu\text{-reach}(\text{dom}(\Gamma_R, \Gamma_S, \Gamma_1) \setminus \text{dom}(\Gamma_I)) = \mu'\text{-reach}(\text{dom}(\Gamma_R, \Gamma_S, \Gamma_1) \setminus \text{dom}(\Gamma_I))$$

which is (3) because $\text{dom}(\Gamma_R, \Gamma_S, \Gamma_1) \setminus \text{dom}(\Gamma_I) = \text{dom}(\Gamma_R, \Gamma_S) \setminus \text{dom}(\Gamma_J)$.

[R-CATCH EXCEPTION] In this case $P = \langle \{a_i\}_{i \in I}, \text{dom}(\mu_2), \{e_j : Q_j\}_{j \in J} \text{throw } e_k \mid P'' \rangle$ and $\mu = \mu_1, \{a_i \mapsto [b_i, \mathfrak{Q}_i]\}_{i \in I}, \mu_2$ where $k \in J$ and $P' = Q_k$ and $\mu' = \mu_1, \{a_i \mapsto [b_i, \varepsilon]\}_{i \in I}$. From rule [T-RUNNING TRANSACTION] we deduce:

- (L.1) $\Gamma = \Gamma_1, \{a_i : \{e_j : S_{ij}\}_{j \in J} T_i\}_{i \in I}, \Gamma_2$;
- (H.1) $\tilde{\mathcal{E}}; [\Gamma_S], \Gamma_1, \{a_i : S_{ik}\}_{i \in I} \vdash P'$;
- (T.1) μ -balanced($\{a_i : S_{ik}\}_{i \in I}$);
- (T.2) local(Γ_2);
- (T.3) $\{a_i\}_{i \in I} \cup \text{dom}(\mu_2) = \mu$ -reach($\{a_i\}_{i \in I} \cup \text{dom}(\Gamma_2)$).

Let $\Gamma'_0 = \Gamma_0 \setminus \text{dom}(\mu_2)$ and $\Gamma' = \Gamma_1, \{a_i : S_{ik}\}_{i \in I}$. We only have to show that (H.2) $\Gamma'_0; \Gamma_R, [\Gamma_S], \Gamma' \Vdash \mu'$ and we prove the items of Definition 4.4. Items (1), (2), and (4) are trivial because μ' has no more pointers than μ , some queues in μ have been emptied in μ' , and duality of endpoint types associated with peer endpoints is preserved by (T.1). Regarding item (3), we have to show that $\text{dom}(\mu') = \text{dom}(\Gamma'_0, \Gamma_R, [\Gamma_S], \Gamma') = \mu'$ -reach($\text{dom}(\Gamma_R, [\Gamma_S], \Gamma')$). The first equality is easy. Regarding the second equality, we derive:

$$\begin{aligned} & \text{dom}(\mu) \\ &= \mu\text{-reach}(\{a_i\}_{i \in I} \uplus \text{dom}(\Gamma_R, \Gamma_1, \Gamma_2)) && \text{by item (3) of Definition 4.4} \\ &= \mu\text{-reach}(\text{dom}(\Gamma_R, \Gamma_1)) \uplus \mu\text{-reach}(\{a_i\}_{i \in I} \cup \text{dom}(\Gamma_2)) \\ & && \text{by item (4) of Definition 4.4} \\ &= \mu\text{-reach}(\text{dom}(\Gamma_R, \Gamma_1)) \uplus \{a_i\}_{i \in I} \uplus \text{dom}(\mu_2) && \text{by (T.3)} \\ &= \text{dom}(\mu_1) \uplus \{a_i\}_{i \in I} \uplus \text{dom}(\mu_2) && \text{by definition of } \mu \end{aligned}$$

where we write \uplus for disjoint union. From the last equality we deduce

$$(*) \quad \text{dom}(\mu_1) = \mu\text{-reach}(\text{dom}(\Gamma_R, [\Gamma_S], \Gamma_1)) = \mu_1\text{-reach}(\text{dom}(\Gamma_R, [\Gamma_S], \Gamma_1))$$

and now we have

$$\begin{aligned}
\text{dom}(\mu') &= \text{dom}(\mu_1) \uplus \{a_i\}_{i \in I} && \text{by definition of } \mu' \\
&= \mu_1\text{-reach}(\text{dom}(\Gamma_R, [\Gamma_S], \Gamma_1)) \cup \{a_i\}_{i \in I} && \text{from (*)} \\
&= \mu'\text{-reach}(\text{dom}(\Gamma_R, [\Gamma_S], \Gamma')) && \text{by definition of } \Gamma' \text{ and } \mu'
\end{aligned}$$

We conclude (1) from (H.1), (H.2) and rule [T-RUNNING PROCESS]. Regarding (2), from (L.1) we know that ranks of all pointers a_i are preserved and from (T.2) that that the rank of all pointers that are no more in the environment is 0. Regarding (3), it holds trivially.

[R-PROPAGATE EXCEPTION] Then $P = \langle \{a_i\}_{i \in I}, \text{dom}(\mu_2), \{e_j : Q_j\}_{j \in J} \text{throw } e \mid P'' \rangle$ and $\mu = \mu_1, \{a_i \mapsto [b_i, \Omega_i]\}_{i \in I}, \mu_2$ and (E.1) $e_j \neq e$ for every $j \in J$ and $P' = \text{throw } e$ and $\mu' = \mu_1, \{a_i \mapsto [b_i, \mathcal{E}]\}_{i \in I}$. From rule [T-RUNNING TRANSACTION] we deduce:

- (L.1) $\Gamma = \Gamma_1, \{a_i : \{e_j : S_{ij}\}_{j \in J} T_i\}_{i \in I}, \Gamma_2$;
- (T.1) $\tilde{\mathcal{E}} \{e_j\}_{j \in J}; \Gamma_0; \Gamma_R; [[\Gamma_S], \Gamma_1], \{a_i : T_i\}_{i \in I}, \Gamma_2 \vdash \mu \ ; \ \text{throw } e \mid P''$
- (T.2) μ -balanced($\{a_i : S_{ij}\}_{i \in I}$);
- (T.3) local(Γ_2);
- (T.4) $\{a_i\}_{i \in I} \cup \text{dom}(\mu_2) = \mu\text{-reach}(\{a_i\}_{i \in I} \cup \text{dom}(\Gamma_2))$.

From (T.1), [T-RUNNING PARALLEL], [T-RUNNING PROCESS], and [T-THROW] we deduce $e \in \tilde{\mathcal{E}}$. Let $\Gamma'_0 = \Gamma_0 \setminus \text{dom}(\mu_2)$ and $\Gamma' = \Gamma_1, \{a_i : S_{ik}\}_{i \in I}$ for some arbitrary $k \in J$. By rule [T-THROW] we derive $\tilde{\mathcal{E}}; [\Gamma_S], \Gamma' \vdash \text{throw } e$. The proof that $\Gamma'_0; \Gamma_R, [\Gamma_S], \Gamma' \Vdash \mu'$ and that items (2) and (3) hold is the same as for the case [R-CATCH EXCEPTION]. \square

B. Proof of Theorem 4.2 (Type Soundness)

The next two lemmas show the relationship between the free names of a process and the names occurring in the context used for typing it.

Lemma B.1. *If $\tilde{\emptyset}; \Gamma \vdash P$, then $\text{dom}(\Gamma) \subseteq \text{fn}(P)$.*

Proof. By induction on the derivation of $\tilde{\emptyset}; \Gamma \vdash P$ and by cases on the last rule applied. Most cases are trivial or follow by a simple induction. Case [T-THROW] is impossible because the exception environment consists of a sequence of empty

sets of exceptions, hence all the exceptions thrown in P are caught. The only interesting case is [T-TRY], when $P = \text{try}(\{u_i\}_{i \in I}) \{e_j : R_j\}_{j \in J} Q$, $\Gamma = \Gamma'$, $\{u_i : \{e_j : S_{ij}\}_{j \in J} \llbracket T_i \rrbracket_{i \in I}, \tilde{\theta} \{e_j\}_{j \in J}; [\Gamma'], \{u_i : T_i\}_{i \in I} \vdash Q$, and $\tilde{\theta}; \Gamma', \{u_i : S_{ij}\}_{i \in I} \vdash R_j$ for all $j \in J$. We distinguish two subcases, according to whether J is empty or not. If $J = \emptyset$, then by induction hypothesis we obtain $\text{dom}([\Gamma'], \{u_i : T_i\}_{i \in I}) \subseteq \text{fn}(Q)$ and then we conclude $\text{dom}(\Gamma) = \text{dom}(\Gamma') \cup \{u_i\}_{i \in I} \subseteq \text{fn}(Q) = \text{fn}(P)$. If $J \neq \emptyset$, then by induction hypothesis we obtain $\text{dom}(\Gamma', \{u_i : S_{ij}\}_{i \in I}) \subseteq \text{fn}(R_j)$ for every $j \in J$. We conclude $\text{dom}(\Gamma) = \text{dom}(\Gamma') \cup \{u_i\}_{i \in I} \subseteq \bigcup_{j \in J} \text{fn}(R_j) \subseteq \text{fn}(P)$. \square

Lemma B.2. *If $\tilde{\theta}; \Gamma_0; \Gamma_R; \Gamma \vdash \mu \ ; P$, then $\text{dom}(\Gamma) \subseteq \text{fn}(P)$.*

Proof. By induction on the derivation of $\tilde{\theta}; \Gamma_0; \Gamma_R; \Gamma \vdash \mu \ ; P$ and by cases on the last rule applied. Cases [T-RUNNING PROCESS] and [T-RUNNING PARALLEL] are easily solved by Lemma B.1 and the induction hypothesis, respectively. Regarding [T-RUNNING TRANSACTION], we have:

- $P = \langle \{a_i\}_{i \in I}, B, \{e_j : R_j\}_{j \in J} Q \rangle$;
- $\Gamma = \Gamma_1, \{a_i : \{e_j : S_{ij}\}_{j \in J} T_i\}_{i \in I}, \Gamma_2$;
- (*) $\{a_i\}_{i \in I} \cup B = \mu\text{-reach}(\{a_i\}_{i \in I} \cup \text{dom}(\Gamma_2))$;
- $\tilde{\theta} \{e_j\}_{j \in J}; \Gamma_0; \Gamma_R; [\Gamma_1], \{a_i : T_i\}_{i \in I}, \Gamma_2 \vdash \mu \ ; Q$;
- $\tilde{\theta}; \Gamma_1, \{a_i : S_{ij}\}_{i \in I} \vdash R_j$ for all $j \in J$.

We distinguish two cases, according to whether J is empty or not. If $J = \emptyset$, then by induction hypothesis we obtain $\text{dom}([\Gamma_1], \{a_i : T_i\}_{i \in I}, \Gamma_2) \subseteq \text{fn}(Q)$ and we conclude $\text{dom}(\Gamma) \subseteq \text{fn}(Q) \subseteq \text{fn}(P)$. If $J \neq \emptyset$, then by induction hypothesis we obtain $\text{dom}(\Gamma_1, \{a_i : S_{ij}\}_{i \in I}) \subseteq \text{fn}(R_j)$ for every $j \in J$. From (*) we deduce $\text{dom}(\Gamma_2) \subseteq \{a_i\}_{i \in I} \cup B$. We conclude $\text{dom}(\Gamma) = \text{dom}(\Gamma_1) \cup \{a_i\}_{i \in I} \cup \text{dom}(\Gamma_2) \subseteq \bigcup_{j \in J} \text{fn}(R_j) \cup \{a_i\}_{i \in I} \cup B \subseteq \text{fn}(P)$. \square

We conclude with the soundness proof of the type system.

Theorem 4.2. *Let $\vdash P$. Then P is well behaved.*

Proof. From the hypothesis $\vdash P$ we deduce $\tilde{\theta}; \emptyset; \emptyset; \emptyset \vdash \emptyset \ ; P$. Consider an arbitrary derivation $\emptyset \ ; P \Rightarrow \mu \ ; Q$. From Theorem 4.1 we deduce that there exist Γ_0 and Γ such that $\tilde{\theta}; \Gamma_0; \emptyset; \Gamma \vdash \mu \ ; Q$ and, from Lemma A.3, we obtain $\Gamma_0; \Gamma \Vdash \mu$.

Regarding condition (1) of Definition 3.5, from $\Gamma_0; \Gamma \Vdash \mu$ and Definition 4.4 we know $\text{dom}(\mu) = \mu\text{-reach}(\text{dom}(\Gamma))$. By Lemma B.2 we have $\text{dom}(\Gamma) \subseteq \text{fn}(Q)$.

We conclude $\text{dom}(\mu) = \mu\text{-reach}(\text{dom}(\Gamma)) \subseteq \mu\text{-reach}(\text{fn}(Q))$ because $\mu\text{-reach}(\cdot)$ is monotone.

Regarding condition (2) of Definition 3.5, suppose that $Q \equiv Q_1 \mid Q_2$ and $\mu \S Q_1 \rightarrow$ and $Q_1 \not\equiv \text{throw } e \mid Q'_1$ (the last hypothesis being granted by the fact that Q is well typed in the $\tilde{\emptyset}$ exception environment). We prove $\mu \S Q_1 \downarrow$ by induction on Q_1 .

- ($Q_1 = \text{done}$) We conclude with an application of rule [ST-INACTIVE].
- ($Q_1 = \text{open}(a, b).R$) Because we have assumed that there are infinitely many pointers and structural congruence includes alpha renaming, we may assume $a, b \notin \text{dom}(\mu)$. Then $\mu \S Q_1 \rightarrow$, which contradicts the hypothesis, therefore this case is impossible.
- ($Q_1 = \text{close}(a).R$) From [T-RUNNING PARALLEL], [T-RUNNING PROCESS], [T-CLOSE] and item (3) of Definition 4.4 we obtain $a \in \text{dom}(\Gamma) \subseteq \text{dom}(\Gamma_0, \Gamma) = \text{dom}(\mu)$, and now $\mu \S Q_1 \rightarrow$ which contradicts the hypothesis, therefore this case is impossible.
- ($Q_1 = R_1 \oplus R_2$) This case is impossible because $\mu \S R_1 \oplus R_2$ always reduces.
- ($Q_1 = a!m(c).R$) From rules [T-RUNNING PROCESS], [T-RUNNING PARALLEL] and [T-SEND] we obtain $\Gamma \vdash a : T$ where T is an internal choice and then from item (3) of Definition 4.4 we have $a \in \text{dom}(\mu)$. From item (2) of Definition 4.4 we deduce that the queue associated with a is empty and also that the peer of a , say b , is still allocated in μ for otherwise T would have to be `end`. Then $\mu \S Q_1 \rightarrow$ which contradicts the hypothesis, therefore this case is impossible.
- ($Q_1 = \sum_{i \in I} a?m_i(x_i).R_i$) Then $a \mapsto [b, \Omega] \in \mu$ and the messages and arguments in Ω are consistent with the type of endpoint a . The only case when $\mu \S \sum_{i \in I} a?m_i(x_i).R_i$ does not reduce is when $\Omega = \varepsilon$, therefore we conclude $\mu \S Q_1 \downarrow$ by an application of rule [ST-INPUT].
- ($Q_1 = R_1 \mid R_2$) From the hypothesis $\mu \S Q_1 \rightarrow$ we deduce $\mu \S R_i \rightarrow$ for $i \in \{1, 2\}$. From the hypothesis $Q_1 \not\equiv \text{throw } e \mid Q'_1$ we deduce $R_i \not\equiv \text{throw } e \mid R'_i$ for $i \in \{1, 2\}$. By induction hypothesis we obtain $\mu \S R_i \downarrow$ for $i \in \{1, 2\}$. We conclude with an application of rule [ST-PARALLEL].
- ($Q_1 = \text{try}(A) \{e_j : R_j\}_{j \in J} R'$) From the hypothesis $\mu \S Q_1 \rightarrow$ we deduce $\neg \mu\text{-balanced}(A)$. We conclude with an application of rule [ST-TRY].

- ($Q_1 = \text{throw } e$) This case is impossible by hypothesis.
- ($Q_1 = \text{commit}(A).R$) We conclude immediately with an application of rule [ST-COMMIT].
- ($Q_1 = X(\tilde{a})$) From rule [T-INVOKE] we deduce $X(\tilde{u}) \stackrel{\text{def}}{=} R$ is a definition and \tilde{a} and \tilde{u} have the same length. Then $\mu \circledast X(\tilde{a}) \rightarrow$ which contradicts the hypothesis, therefore this case is impossible.
- ($Q_1 = \langle A, B, \{e_j : R_j\}_{j \in J} R' \rangle$) From the hypothesis $\mu \circledast Q_1 \rightarrow$ we deduce $\mu \circledast R' \rightarrow$ and $R' \not\equiv \prod_{i \in I} \text{commit}(A_i).R_i$ and $R' \not\equiv \text{throw } e \mid R''$ since these are the cases when $\mu \circledast Q_1$ does reduce. By induction hypothesis we deduce $\mu \circledast R' \downarrow$ and we conclude with an application of rule [ST-RUNNING TRANSACTION]. \square

References

- [1] G. Hunt, J. R. Larus, M. Abadi, M. Aiken, P. Barham, M. Fähndrich, C. Hawblitzel, O. Hodson, S. Levi, N. Murphy, B. Steensgaard, D. Tarditi, T. Wobber, B. Zill, An Overview of the Singularity Project, Technical Report MSR-TR-2005-135, Microsoft Research, 2005.
- [2] G. Hunt, J. R. Larus, Singularity: Rethinking the Software Stack, SIGOPS Operating Systems Review 41 (2007) 37–49.
- [3] S. Finley, Ext2 on Singularity, Technical Report, University of Wisconsin, 2008. Available at <http://pages.cs.wisc.edu/~remzi/Courses/736/Spring2008/Projects/Scott/Ext2%20on%20Singularity.pdf>.
- [4] M. Fähndrich, M. Aiken, C. Hawblitzel, O. Hodson, G. Hunt, J. R. Larus, S. Levi, Language Support for Fast and Reliable Message-based Communication in Singularity OS, in: Proceedings of EuroSys'06, ACM, 2006, pp. 177–190.
- [5] Z. Stengel, T. Bultan, Analyzing Singularity Channel Contracts, in: Proceedings of the 18th International Symposium on Software Testing and Analysis (ISSTA'09), ACM, 2009, pp. 13–24.
- [6] J. Villard, E. Lozes, C. Calcagno, Proving Copyless Message Passing, in: Proceedings of the 7th Asian Symposium on Programming Languages and Systems (APLAS'09), LNCS 5904, Springer, 2009, pp. 194–209.

- [7] V. Bono, L. Padovani, Typing Copyless Message Passing, *Logical Methods in Computer Science* 8 (2012) 1–50.
- [8] K. Honda, Types for Dyadic Interaction, in: *Proceedings of the 4th International Conference on Concurrency Theory (CONCUR'93)*, LNCS 715, Springer, 1993, pp. 509–523.
- [9] K. Honda, V. T. Vasconcelos, M. Kubo, Language Primitives and Type Disciplines for Structured Communication-based Programming, in: *Proceedings of the 7th European Symposium on Programming (ESOP'98)*, LNCS 1381, Springer, 1998, pp. 122–138.
- [10] S. Jakšić, L. Padovani, Exception Handling for Copyless Messaging, in: *Proceedings of the 14th International ACM SIGPLAN Symposium on Principles and Practice of Declarative Programming (PPDP'12)*, ACM, 2012, pp. 151–162.
- [11] M. Carbone, K. Honda, N. Yoshida, Structured Interactional Exceptions in Session Types, in: *Proceedings of the 19th International Conference on Concurrency Theory (CONCUR'08)*, LNCS 5201, Springer, 2008, pp. 402–417.
- [12] S. Capecchi, E. Giachino, N. Yoshida, Global Escape in Multiparty Sessions, in: *Proceedings of the IARCS Annual Conference on Foundations of Software Technology and Theoretical Computer Science (FSTTCS'10)*, 2010, pp. 338–351.
- [13] R. Milner, *Communicating and mobile systems: the π -calculus*, Cambridge University Press, New York, NY, USA, 1999.
- [14] S. Gay, M. Hole, Subtyping for Session Types in the π -calculus, *Acta Informatica* 42 (2005) 191–225.
- [15] S. L. Peyton Jones, C. D. Clack, J. Salkild, M. Hardie, GRIP - A High-Performance Architecture for Parallel Graph Reduction, in: *Proceedings of Functional Programming Languages and Computer Architecture (FPCA'87)*, LNCS 274, Springer, 1987, pp. 98–112.
- [16] P. W. Trinder, K. Hammond, J. S. Mattson Jr., A. Partridge, S. L. Peyton Jones, GUM: A Portable Parallel Implementation of Haskell, in: *Proceedings of the ACM SIGPLAN 1996 Conference on Programming Language Design and Implementation (PLDI'96)*, ACM, 1996, pp. 79–88.

- [17] H.-W. Loidl, The Virtual Shared Memory Performance of a Parallel Graph Reduce, in: Proceedings of the 2nd IEEE International Symposium on Cluster Computing and the Grid (CCGRID'02), 2002, pp. 311–318.
- [18] M. Tofte, J.-P. Talpin, Region-based memory management, *Information and Computation* 132 (1997) 109 – 176.
- [19] M. Hofmann, S. Jost, Static Prediction of Heap Space Usage for First-Order Functional Programs, in: Proceedings of the 30th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL'03), ACM, 2003, pp. 185–197.
- [20] M. Hofmann, S. Jost, Type-Based Amortised Heap-Space Analysis, in: Proceedings of the 15th European Symposium on Programming (ESOP'06), LNCS 3924, Springer, 2006, pp. 22–37.
- [21] E. Albert, S. Genaim, M. Gómez-Zamalloa, Heap Space Analysis for Java Bytecode, in: Proceedings of the 6th International Symposium on Memory Management (ISMM'07), ACM, 2007, pp. 105–116.
- [22] K. Hammond, The Dynamic Properties of Hume: A Functionally-Based Concurrent Language with Bounded Time and Space Behaviour, in: Proceedings of the 12th International Workshop on the Implementation of Functional Languages (IFL'00), LNCS 2011, Springer, 2000, pp. 122–139.
- [23] K. Suenaga, R. Fukuda, A. Igarashi, Type-Based Safe Resource Deallocation for Shared-Memory Concurrency, in: Proceedings of the 27th Annual ACM SIGPLAN Conference on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA'12), ACM, 2012, pp. 1–20.
- [24] F. Iwama, A. Igarashi, N. Kobayashi, Resource Usage Analysis for a Functional Language with Exceptions, in: Proceedings of the 2006 ACM SIGPLAN Workshop on Partial Evaluation and Semantics-based Program Manipulation (PEPM'06), ACM, 2006, pp. 38–47.
- [25] J. Villard, E. Lozes, C. Calcagno, Tracking Heaps That Hop with Heap-Hop, in: Proceedings of the 16th International Conference on Tools and Algorithms for the Construction and Analysis of Systems (TACAS'10), LNCS 6015, Springer, 2010, pp. 275–279.

- [26] C. Ferreira, I. Lanese, A. Ravara, H. T. Vieira, G. Zavattaro, Advanced Mechanisms for Service Combination and Transactions, in: *Rigorous Software Engineering for Service-Oriented Systems*, LNCS 6582, Springer, 2011, pp. 302–325.
- [27] T. Harris, S. Marlow, S. Peyton Jones, M. Herlihy, Composable Memory Transactions, in: *Proceedings of the ACM SIGPLAN Symposium on Principles and Practice of Parallel Programming (PPoPP'05)*, ACM, 2005, pp. 48–60.
- [28] K. Donnelly, M. Fluet, Transactional Events, in: *Proceedings of the 11th ACM SIGPLAN International Conference on Functional Programming (ICFP'06)*, ACM, 2006, pp. 124–135.
- [29] W. Weimer, G. C. Necula, Finding and Preventing Run-Time Error Handling Mistakes, in: *Proceedings of the 19th Annual ACM SIGPLAN Conference on Object-Oriented Programming, Systems, Languages, and Applications (OOPSLA'04)*, ACM, 2004, pp. 419–431.
- [30] W. Weimer, Exception-Handling Bugs in Java and a Language Extension to Avoid Them, in: *Advanced Topics in Exception Handling Techniques*, LNCS 4119, Springer, 2006, pp. 22–41.
- [31] W. Weimer, G. C. Necula, Exceptional situations and program reliability, *ACM Trans. Program. Lang. Syst.* 30 (2008) 1–51.