
COMPOSITIONAL REASONING FOR EXPLICIT RESOURCE MANAGEMENT IN CHANNEL-BASED CONCURRENCY *

ADRIAN FRANCALANZA ^a, EDSKO DEVRIES ^b, AND MATTHEW HENNESSY ^c

^a ICT, University of Malta
e-mail address: adrian.francalanza@um.edu.mt

^b Well-Typed LLP, UK
e-mail address: edsko@well-typed.com

^c Trinity College Dublin, Ireland
e-mail address: matthew.hennessy@cs.tcd.ie

ABSTRACT. We define a π -calculus variant with a costed semantics where channels are treated as resources that must explicitly be allocated before they are used and can be deallocated when no longer required. We use a substructural type system tracking permission transfer to construct coinductive proof techniques for comparing behaviour and resource usage efficiency of concurrent processes. We establish full abstraction results between our coinductive definitions and a contextual behavioural preorder describing a notion of process efficiency *wrt.* its management of resources. We also justify these definitions and respective proof techniques through numerous examples and a case study comparing two concurrent implementations of an extensible buffer.

1. INTRODUCTION

We investigate the *behaviour* and *space efficiency* of concurrent programs with *explicit resource-management*. In particular, our study focuses on *channel-passing concurrent programs*: we define a π -calculus variant, called $R\pi$, where the only resources available are channels; these channels must explicitly be allocated before they can be used, and can be deallocated when no longer required. As part of the operational model of the language, channel allocation and deallocation have costs associated with them, reflecting the respective resource usage.

Explicit resource management is typically desirable in settings where resources are *scarce*. Resource management programming constructs such as explicit deallocation provide fine-grained control over how these resources are used and recycled. By comparison, in automated mechanisms such as garbage collection, unused resources (in this case, memory) tend to remain longer in an unreclaimed state [27, 28]. Explicit resource management constructs such as memory deallocation also carry advantages over automated mechanisms such as garbage collection techniques when it

2012 ACM CCS: [Theory of computation]: Models of Computation—Concurrency—Process Calculi.

Key words and phrases: π -calculus, concurrency, memory management, coinductive reasoning.

* An extended abstract of a preliminary version of the paper has appeared in [11].

^c Supported by SFI project SFI 06 IN.1 1898.

comes to *interactive* and *real-time* programs [10, 27, 28]. In particular, garbage collection techniques require additional computation to determine otherwise explicit information as to which parts of the memory to reclaim and at what stage of the computation; the associated overheads may lead to uneven performance and intolerable pause periods where the system becomes unresponsive [10].

In the case of channel-passing concurrency with explicit memory-management, the analysis of the relative behaviour and efficiency of programs is non-trivial for a number of reasons. Explicit memory-management introduces the risk of either premature or multiple deallocation of resources along separate threads of execution; these are more difficult to detect than in single-threaded programs and potentially result in problems such as wild pointers or corrupted heaps which may, in turn, lead to unpredictable, even catastrophic, behaviour [27, 28]. It also increases the possibility of memory leaks, which are often not noticeable in short-running, terminating programs but subtly eat up resources over the course of long-running programs. In a concurrent settings such as ours, complications relating to the assessment and comparison of resource consumption is further compounded by the fact that the runtime execution of channel-passing concurrent programs can have *multiple interleavings*, is sometimes *non-deterministic* and often *non-terminating*.

1.1. Scenario: Consider a setting with two servers, S_1 and S_2 , which repeatedly listen for service requests on channels srv_1 and srv_2 , respectively. Requests send a *return* channel on srv_1 or srv_2 which is then used by the servers to service the requests and send back answers, v_1 and v_2 . A possible implementation for these servers is given in (1.1) below, where $\text{rec } w.P$ denotes a process P recursing at w , $c?x.P$ denotes a process inputting on channel c some value that is bound to the variable x in the continuation P , and $c!v.P$ outputs a value v on channel c and continues as P :

$$S_i \triangleq \text{rec } w. \text{srv}_i?x. x!v_i. w \quad \text{for } i \in \{1, 2\} \quad (1.1)$$

Clients that need to request service from *both* servers, so as to report back the outcome of both server interactions on some channel, ret , can be programmed in a variety of ways:

$$\begin{aligned} C_0 &\triangleq \text{rec } w. \text{alloc } x_1. \text{alloc } x_2. \text{srv}_1!x_1. x_1?y. \text{srv}_2!x_2. x_2?z. \text{ret}!(y, z). w \\ C_1 &\triangleq \text{rec } w. \text{alloc } x. \text{srv}_1!x. x?y. \text{srv}_2!x. x?z. \text{ret}!(y, z). w \\ C_2 &\triangleq \text{rec } w. \text{alloc } x. \text{srv}_1!x. x?y. \text{srv}_2!x. x?z. \text{free } x. \text{ret}!(y, z). w \end{aligned} \quad (1.2)$$

C_0 corresponds to an idiomatic π -calculus client. In order to ensure that it is the sole recipient of the service requests, it creates *two* new return channels to communicate with S_1 and S_2 on srv_1 and srv_2 , using the command $\text{alloc } x.P$; this command allocates a *new* channel c and binds it to the variable x in the continuation P . Allocating a new channel for each service request ensures that the return channel used between the client and server is *private* for the duration of the service, preventing interferences from other parties executing in parallel.

One important difference between the computational model considered in this paper and that of the standard π -calculus is that channel allocation is an expensive operation *i.e.*, it incurs an additional (*spatial*) cost compared to the other operations. Client C_1 attempts to address the inefficiencies of C_0 by allocating only *one* additional new channel, and *reusing* this channel for both interactions with the servers. Intuitively, this channel reuse is valid, *i.e.*, it preserves the client-server behaviour C_0 had with servers S_1 and S_2 , because the server implementations above use the received return-channels *only once*. This single channel usage guarantees that return channels remain private during the duration of the service, despite the reuse from client C_1 .

Client C_2 attempts to be more efficient still. More precisely, since our computational model does not assume implicit resource reclamation, the previous two clients can be deemed as having *memory leaks*: at every iteration of the client-server interaction sequence, C_0 and C_1 allocate new

channels that are not disposed of, even though these channels are never used again in subsequent iterations. By contrast, C_2 deallocates unused channels at the end of each iteration using the construct $\text{free } c.P$.

In this work we develop a formal framework for comparing the behaviour of concurrent processes that explicitly allocate and deallocate channels. For instance, processes consisting of the servers S_1 and S_2 together with any of the clients C_0 , C_1 or C_2 should be *related*, on the basis that they exhibit the same behaviour. In addition, we would like to *order* these systems, based on their relative efficiencies *wrt.* the (channel) resources used. We note that there are various, at times contrasting, notions of efficiency that one may consider. For instance, one notion may consider acquiring memory for long periods to be less efficient than repeatedly allocating and deallocating memory; another notion of efficiency could instead focus on minimising the allocation and deallocation operations used, as these are considerably more expensive than other operations. In this work, we mainly focus on a notion of efficiency that accounts for the relative memory allocations required to carry out the necessary computations. Thus, we would intuitively like to develop a framework yielding the following preorder, where \sqsubseteq reads "more efficient than":

$$S_1 \parallel S_2 \parallel C_2 \sqsubseteq S_1 \parallel S_2 \parallel C_1 \sqsubseteq S_1 \parallel S_2 \parallel C_0 \quad (1.3)$$

A pleasing property of this preorder would be *compositionality*, which implies that orderings are preserved under larger contexts, *i.e.*, for all (valid) contexts $C[-]$, $P \sqsubseteq Q$ implies $C[P] \sqsubseteq C[Q]$. Dually, compositionality would also improve the scalability of our formal framework since, to show that $C[P] \sqsubseteq C[Q]$ (for some context $C[-]$), it suffices to obtain $P \sqsubseteq Q$. For instance, in the case of (1.3), compositionality would allow us to factor out the common code, *i.e.*, the servers S_1 and S_2 as the context $S_1 \parallel S_2 \parallel [-]$, and focus on showing that

$$C_2 \sqsubseteq C_1 \sqsubseteq C_0 \quad (1.4)$$

1.2. Main Challenges: The details are however far from straightforward. To begin with, we need to assess relative program cost over potentially infinite computations. Thus, rudimentary aggregate measures such as adding up the total computation cost of processes and comparing this total at the end of the computation is insufficient for system comparisons such as (1.3). In such cases, a preliminary attempt at a solution would be to compare the *relative cost* for *every* server interaction (action): in the sense of [4], the preorder would then ensure that every *costed* interaction by the inefficient clients must be matched by a corresponding *cheaper* interaction by the more efficient client (and, dually, costed interactions by the efficient client must be matched by interactions from the inefficient client that are as costly or more).

$$C_3 \triangleq \text{rec } w.\text{alloc } x_1.\text{alloc } x_2. \text{srv}_1!x_1.x_1?y. \text{srv}_2!x_2.x_2?z. \text{free } x_1.\text{free } x_2.\text{ret}!(y,z).w \quad (1.5)$$

There are however problems with this approach. Consider, for instance, C_3 defined in (1.5). Even though this client allocates two channels for every iteration of server interactions, it does not exhibit any memory leaks since it deallocates them both at the end of the iteration. It may therefore be sensible for our preorder to equate C_3 with client C_2 of (1.2) by having $C_2 \sqsubseteq C_3$ as well as $C_3 \sqsubseteq C_2$. However showing $C_3 \sqsubseteq C_2$ would not be possible using the preliminary strategy discussed above, since, C_3 must engage in more expensive computation (allocating two channels as opposed to 1) by the time the interaction with the first server is carried out.

Worse still, an analysis strategy akin to [4] would not be applicable for a comparison involving the clients C_1 and C_3 . In spite of the fact that over the course of its entire computation C_3 requires less resources than C_1 , *i.e.*, it is more efficient, client C_3 appears to be *less efficient* than C_1 after the

interaction with the first server on channel srv_1 since, at that stage, it has allocated two new channels as opposed to one. However, C_1 becomes less efficient for the remainder of the iteration since it never deallocates the channel it allocates whereas C_3 deallocates both channels. To summarise, for comparisons $C_3 \sqsubseteq C_2$ and $C_3 \sqsubseteq C_1$, we need our analysis to allow a process to be *temporarily inefficient* as long as it can recover later on.

In this paper, we use a costed semantics to define an efficiency preorder to reason about the relative cost of processes over potentially infinite computation, based on earlier work by [30, 34]. In particular, we adapt the concept of *cost amortisation* to our setting, used by our preorders to compare processes that are eventually more efficient than others over the course of their entire computation, but are temporarily less efficient at certain stages of the computation.

Issues concerning cost assessment are however not the only obstacles tackled in this work; there are also complications associated with the compositionality aspects of our proposed framework. More precisely, we want to limit our analysis to *safe* contexts, *i.e.*, contexts that use resources in a sensible way, *e.g.*, not deallocating channels while they are still in use. In addition, we also want to consider behaviour *wrt.* a subset of the possible safe contexts. For instance, our clients from (1.2) only exhibit the same behaviour *wrt.* servers that (i) accept (*any number of*) requests on channels srv_1 and srv_2 containing a return channel, which then (ii) use this channel at most *once* to return the requested answer. We can characterise the interface between the servers and the clients using fairly standard channel type descriptions adapted from [31] in (1.6), where $[\mathbf{T}]^\omega$ describes a channel that can be used *any* number of times (*i.e.*, the channel-type attribute ω) to communicate values of type \mathbf{T} , whereas $[\mathbf{T}]^1$ denotes an *affine* channel (*i.e.*, a channel type with attribute $\mathbf{1}$) that can be used *at most once* to communicate values of type \mathbf{T} :

$$\text{srv}_1 : [[\mathbf{T}_1]^1]^\omega, \quad \text{srv}_2 : [[\mathbf{T}_2]^1]^\omega \quad (1.6)$$

In the style of [45, 21], we could then use this interface to abstract away from the actual server implementations described in (1.1) and state that, *wrt.* contexts that observe the channel mappings of (1.6), client C_2 is more efficient than C_1 which is, in turn, more efficient than C_0 . These can be expressed as:

$$\text{srv}_1 : [[\mathbf{T}_1]^1]^\omega, \text{srv}_2 : [[\mathbf{T}_2]^1]^\omega \models C_2 \sqsubseteq C_1 \quad (1.7)$$

$$\text{srv}_1 : [[\mathbf{T}_1]^1]^\omega, \text{srv}_2 : [[\mathbf{T}_2]^1]^\omega \models C_1 \sqsubseteq C_0 \quad (1.8)$$

Unfortunately, the machinery of [45, 21] cannot be easily extended to our costed analysis because of two main reasons. First, in order to limit our analysis to safe computation, we would need to show that clients C_0 , C_1 and C_2 adhere to the channel usage stipulated by the type associations in (1.6). However, the channel reuse in C_1 and C_2 (an essential feature to attain space efficiency) requires our analysis to associate potentially different types (*i.e.*, $[\mathbf{T}_1]^1$ and $[\mathbf{T}_2]^1$) to the same return channel; this channel reuse at different types amounts to a form of *strong update*, a degree of flexibility not supported by [45, 21].

Second, the equivalence reasoning mechanisms used in [45, 21] would be substantially limiting for processes with channel reuse. More specifically, consider the slightly tweaked client implementation of C_2 below:

$$C'_2 \triangleq \text{rec } w.\text{alloc } x.(\text{srv}_1!x \parallel x?y.(\text{srv}_2!x \parallel x?z.\text{free } x.c!(y, z).X)) \quad (1.9)$$

The only difference between the client in (1.9) and the original one in (1.2) is that C_2 *sequences* the service requests before the service inputs, *i.e.*, $\dots \text{srv}_1!x. x?y. \dots$ and $\dots \text{srv}_2!x. x?z. \dots$, whereas C'_2 *parallelises* them, *i.e.*, $\dots \text{srv}_1!x \parallel x?y. \dots$ and $\dots \text{srv}_2!x \parallel x?z. \dots$. Resource-centric type disciplines such as [12, 40] preclude name matching for a particular resource once all the permissions to

use that resource have been used up; this feature is essential to statically reason about a number of basic design patterns for reuse. For such type settings, it turns out that the client implementations C_2 and C'_2 exhibit the same behaviour because the return channel used by both clients for *both* server interactions is private, *i.e.*, unknown to the respective servers; as a result, the servers cannot answer the service on that channel before it is received on either srv_1 or srv_2 .¹ Through *scope extrusion*, theories such as [45, 21] can reason adequately about the first server interaction, and relate $\dots \text{srv}_1!x.x?y.\dots$ of C_2 with $\dots \text{srv}_1!x \parallel x?y.\dots$ of C'_2 . However, they have no mechanism for tracking channel locality post scope extrusion, thereby recovering the information that the return channel *becomes private again* to the client after the first server interaction (since the servers use up the permission to use the return channel once they reply on it). This prohibits [45, 21] from determining that the second server interaction is just an instance of the first server interaction, thus failing to relate these two implementations.

In [12] we developed a substructural type system based around a type attribute describing channel *uniqueness*, and this was used to statically ensure safe computations for $R\pi$. In this work, we weave this type information into our framework, imbuing it with an operational permission-semantics to reason compositionally about the costed behaviour of (safe) processes. More specifically, in (1.2), when C_2 allocates channel x , no other process knows about x : from a typing perspective, but also operationally, x is *unique* to C_2 . Client C_2 then sends x on srv_1 at an *affine* type, which (by definition) limits the server to use x at most once. At this point, from an operational perspective, x is to C_2 , the entity previously “owning” it, *unique-after-1* (communication) use. This means that after one communication step on x , (the derivative of) C_2 recognises that all the other processes apart from it must have used up the single affine permission for x , and hence x becomes once again *unique* to C_2 . This also means that C_2 can safely *reuse* x , possibly at a different object type (strong update), or else safely deallocate it.

The concept of affinity is well-known in the process calculus community. By contrast, uniqueness (and its duality to affinity) is used far less. In a compositional framework, uniqueness can be used to record the guarantee at one end of a channel corresponding to the restriction associated with affine channel usage at the other; an operational semantics can be defined, tracking the *permission transfer* of affine permissions back and forth between processes as a result of communication, addressing the aforementioned complications associated with idioms such as channel reuse. We employ such an operational (costed) semantics to define our efficiency preorders for concurrent processes with explicit resource management, based on the notion of amortised cost discussed above.

1.3. Paper Structure: Section 2 introduces our language with constructs for explicit memory management and defines a costed semantics for it. We illustrate issues relating to resource usage in this language through a case study in Section 3, discussing different implementations for an unbounded buffer. Section 4 develops a labelled-transition system for our language that takes into consideration some representation of the observer and the permissions that are exchanged between the program and the observer; it is a typed transition system similar to [38, 21, 19], nuanced to the resource-focussed type system of [12]. Based on this transition system, the section also defines a coinductive cost-based preorder and proves a number of properties about it. Section 5 justifies the cost-based preorder by relating it with a behavioural contextual preorder defined in terms of the reduction semantics of Section 2. Section 6 applies the theory of Section 4 to reason about the efficiency of the unbounded buffer implementations of Section 3. Finally, Section 7 surveys related work and Section 8 concludes.

¹Analogously, in the π -calculus, $\text{new } d.(c!d \parallel d?x.P)$ is indistinguishable from $\text{new } d.(c!d.d?x.P)$

$P, Q ::= u!\vec{v}.P$	(output)		$u?\vec{x}.P$	(input)
nil	(nil)		$\text{if } u = v \text{ then } P \text{ else } Q$	(match)
$\text{rec } w.P$	(recursion)		x	(process variable)
$P \parallel Q$	(parallel)		$\text{alloc } x.P$	(allocate)
$\text{free } u.P$	(deallocate)			

Figure 1: $R\pi$ Syntax

2. THE LANGUAGE

Figure 1 shows the syntax for our language, the resource π -calculus, or $R\pi$ for short. It has the standard π -calculus constructs with the exception of scoping, which is replaced with primitives for explicit channel allocation, $\text{alloc } x.P$, and deallocation, $\text{free } x.P$. The syntax assumes two separate denumerable sets of channel names $c, d \in \text{CHAN}$, and variables $x, y, z, w \in \text{VAR}$, and lets identifiers u, v range over both sets, $\text{CHAN} \cup \text{VAR}$. The input construct, $c?.x.P$, recursion construct, $\text{rec } w.P$, and channel allocation construct, $\text{alloc } x.P$, are binders whereby free occurrences of the variables x and w in P are bound. As opposed to more standard versions of the π -calculus, we *do not* use name scoping to bind and bookkeep the visibility of names; we shall however use alternative mechanisms to track name knowledge and usage in subsequent development.

$R\pi$ processes run in a resource environment, ranged over by M, N , representing predicates over channel names stating whether a channel is allocated or not. We find it convenient to denote such functions as a list of channels representing the set channels that are allocated, e.g., the list c, d denotes the set $\{c, d\}$, representing the resource environment returning *true* for channels c and d and *false* otherwise - in this representation, the order of the channels in the list is unimportant, but duplicate channels are disallowed; as shorthand, we also write M, c to denote $M \cup \{c\}$ whenever $c \notin M$. In this paper we consider only resource environments with an *infinite* number of deallocated channels, i.e., M is a total function. Models with finite resources can be easily accommodated by making M partial; this also would entail a slight change in the semantics of the allocation construct, which could either block or fail whenever there are no deallocated resources left. Although interesting in its own right, we focus on settings with infinite resources as it lends itself better to the analysis of resource efficiency that follows.

We refer to the pair $M \triangleright P$, consisting of a resource environment M and a *closed* process² P as a *system*; note that *not all* free names in P need to be allocated i.e., present in M : intuitively, any name c used by P and $c \notin M$ represents a *dangling pointer*. Contexts consist of parallel composition of processes; they are however defined over systems, through the grammar and the respective definition at the top of Figure 2. The reduction relation is defined as the least *contextual* relation over systems satisfying the rules in Figure 2. More specifically our reduction relation leaves the following rule implicit:

$$\frac{M \triangleright P \longrightarrow_k M \triangleright Q}{C[M \triangleright P] \longrightarrow_k C[M \triangleright Q]} \text{rCTX}$$

²A closed process has no free variables. Note that the absence of name binders i.e., no name scoping, means that all names are free.

Contexts

$$C ::= [-] \mid C \parallel P \mid P \parallel C$$

$$\begin{aligned} [M \triangleright P] &\stackrel{\text{def}}{=} M \triangleright P \\ C[M \triangleright P] \parallel Q &\stackrel{\text{def}}{=} M' \triangleright (P' \parallel Q) \quad \text{if } C[M \triangleright P] = M' \triangleright P' \\ Q \parallel C[M \triangleright P] &\stackrel{\text{def}}{=} M' \triangleright (Q \parallel P') \quad \text{if } C[M \triangleright P] = M' \triangleright P' \end{aligned}$$

Structural Equivalence

$$\text{sCom} \quad P \parallel Q \equiv Q \parallel P \quad \text{sAss} \quad P \parallel (Q \parallel R) \equiv (P \parallel Q) \parallel R \quad \text{sNil} \quad P \parallel \text{nil} \equiv P$$

Reduction Rules

$$\begin{aligned} &\frac{}{M, c \triangleright c!d.P \parallel c?x.Q \longrightarrow_0 M, c \triangleright P \parallel Q\{d/x\}} \text{rCom} \\ &\frac{}{M, c \triangleright \text{if } c = c \text{ then } P \text{ else } Q \longrightarrow_0 M, c \triangleright P} \text{rThen} \\ &\frac{}{M, c, d \triangleright \text{if } c = d \text{ then } P \text{ else } Q \longrightarrow_0 M, c, d \triangleright Q} \text{rElse} \\ &\frac{}{M \triangleright \text{rec } w.P \longrightarrow_0 M \triangleright P\{\text{rec } w.P/w\}} \text{rRec} \quad \frac{P \equiv P' \quad M \triangleright P' \longrightarrow_k M \triangleright Q' \quad Q' \equiv Q}{M \triangleright P \longrightarrow_k M \triangleright Q} \text{rStr} \\ &\frac{}{M \triangleright \text{alloc } x.P \longrightarrow_{+1} M, c \triangleright P\{c/x\}} \text{rAll} \quad \frac{}{M, c \triangleright \text{free } c.P \longrightarrow_{-1} M \triangleright P} \text{rFree} \end{aligned}$$

Reflexive Transitive Closure

$$\frac{}{M \triangleright P \longrightarrow_0^* M \triangleright P} \quad \frac{M \triangleright P \longrightarrow_k^* M' \triangleright P' \quad M' \triangleright P' \longrightarrow_l M'' \triangleright P''}{M \triangleright P \longrightarrow_{k+l}^* M'' \triangleright P''}$$

Figure 2: $R\pi$ Reduction Semantics

Rule (rSTR) extends reductions to structurally equivalent processes, $P \equiv Q$, i.e., processes that are identified up to superfluous nil processes, and commutativity/associativity of parallel composition (see the structural equivalence rules Figure 2).

Most rules follow those of the standard π -calculus, e.g., (rREC), with the exception of those involving resource handling. For instance, the rule for communication (rCOM) requires the communicating channel to be *allocated*. Allocation (rALL) chooses a deallocated channel, allocates it, and

$a ::= \omega$ (unrestricted)	$\mathbf{1}$ (affine)	(\bullet, i) (unique after i steps)
$\mathbf{T} ::= \mathbf{U}$ (channel type)	\mathbf{proc} (process type)	
$\mathbf{U} ::= [\vec{\mathbf{U}}]^a$ (channel)	$\mu X. \mathbf{U}$ (recursion)	X (variable)

Figure 3: Type Attributes and Types

substitutes it for the bound variable of the allocation construct.³ Deallocation (\mathbf{rFREE}) changes the states of a channel from allocated to deallocated, making it available for future allocations. The rules are annotated with a cost reflecting resource usage; allocation has a cost of $+1$, deallocation has a (negative) cost of -1 while the other reductions carry no cost, *i.e.*, 0 . Figure 2 also shows the natural definition of the reflexive transitive closure of the costed reduction relation. In what follows, we use $k, l \in \mathbb{Z}$ as integer metavariables to range over costs.

Example 2.1. The following reduction sequence illustrates potential unwanted behaviour resulting from resource mismanagement:

$$M, c \triangleright \mathbf{free} c.(c!1 \parallel c?x.P) \parallel \mathbf{alloc} y.(y!42 \parallel y?z.Q) \longrightarrow_{-1} \quad (2.1)$$

$$M \triangleright c!1 \parallel c?x.P \parallel \mathbf{alloc} y.(x!42 \parallel x?z.Q) \longrightarrow_{+1} \quad (2.2)$$

$$M, c \triangleright c!1 \parallel c?x.P \parallel c!42 \parallel c?z.Q \quad (2.3)$$

Intuitively, allocation should yield “fresh” channels *i.e.*, channels that are not in use by any active process. This assumption is used by the right process in system (2.1), $\mathbf{alloc} y.(y!42 \parallel y?z.Q)$, to carry out a *local* communication, sending the value 42 on some local channel y that no other process is using. However, the premature deallocation of the channel c by the left process in (2.1), $\mathbf{free} c.(c!1 \parallel c?x.P)$, allows channel c to be reallocated by the right process in the subsequent reduction, (2.2). This may then lead to unintended behaviour since we may end up with interferences when communicating on c in the residuals of the left and right processes, (2.3).⁴ \square

In [12] we defined a type system that precludes unwanted behaviour such as in Example 2.1. The type syntax is shown in Figure 3. The main type entities are *channel types*, denoted as $[\vec{\mathbf{U}}]^a$, where *type attributes* a range over

- $\mathbf{1}$, for affine, imposing a restriction/obligation on usage;
- (\bullet, i) , for unique-after- i usages ($i \in \mathbb{N}$), providing guarantees on usage;
- ω , for unrestricted channel usage without restrictions or guarantees.

Uniqueness typing can be seen as dual to affine typing [18], and in [12] we make use of this duality to keep track of uniqueness across channel-passing parallel processes: an attribute (\bullet, i) typing an endpoint of a channel c accounts for (at most) i instances of affine attributes typing endpoints of that same channel.

A channel type $[\vec{\mathbf{U}}]^a$ also describes the type of the values that can be communicated on that channel, $\vec{\mathbf{U}}$, which denotes a list of types $\mathbf{U}_1, \dots, \mathbf{U}_n$ for $n \in \mathbb{NAT}$; when $n = 0$, the type list is an

³The expected side-condition $c \notin M$ is implicit in the notation (M, c) used in the system $M, c \triangleright P\{c/x\}$ to which it reduces, since c cannot be present in M for M, c to be valid.

⁴Operationally, we do not describe errors that may result from attempted communications on deallocated channels (we do not have error values). This may occur after reduction (2.1), if the residual of the left process communicate on channel c . Rather, communications on deallocated channels are blocked.

empty list and we simply write $[\]^a$. Note the difference between $[\vec{\mathbf{U}}]^1$, *i.e.*, a channel with an affine usage restriction, and $[\vec{\mathbf{U}}]^{(\bullet,1)}$, *i.e.*, a channel with a unique-after-1 usage guarantee. We denote fully unique channels as $[\vec{\mathbf{U}}]^\bullet$ in lieu of $[\vec{\mathbf{U}}]^{(\bullet,0)}$.

The type syntax also assumes a denumerable set of type variables X, Y , bound by the recursive type construct $\mu X.U$. In what follows, we restrict our attention to *closed, contractive* types, where every type variable is bound and appears within a channel constructor $[\]^a$; this ensures that channel types such as $\mu X.X$ are avoided. We assume an equi-recursive interpretation for our recursive types [36] (see τEQ in Figure 4), characterised as the least type-congruence satisfying rule εREC in Figure 4.

$$\frac{\Gamma \vdash P \quad \text{dom}(\Gamma) \subseteq M \quad \Gamma \text{ is consistent}}{\Gamma \vdash M \triangleright P} \tau\text{SYS}$$

The rules for typing processes are given in Figure 4 and take the usual shape $\Gamma \vdash P$ stating that process P is well-typed with respect to the environment Γ , a list of pairs of identifiers and types. Systems are typed according to (τSYS) above: a system $M \triangleright P$ is well-typed under Γ if P is well-typed wrt. Γ , $\Gamma \vdash P$, and Γ only contains assumptions for channels that have been allocated, $\text{dom}(\Gamma) \subseteq M$. This restricts channel usage in P to allocated channels and is key for ensuring safety.

In [12], typing environments are multisets of pairs of identifiers and types; we do not require them to be partial functions. However, the (top-level) typing rule for systems (τSYS) requires that the typing environment is *consistent*. A typing environment is consistent if whenever it contains multiple assumptions about a channel, then these assumptions can be derived from a *single assumption* using the structural rules of the type system (see the structural rule τCON and the splitting rule εUNQ in Figure 4).

Definition 2.2 (Consistency). A typing environment Γ is *consistent* if there is a partial map Γ' such that $\Gamma' < \Gamma$.

The environment structural rules, $\Gamma_1 < \Gamma_2$, defined in Figure 4, govern the way type environments are syntactically manipulated. For instance, rules τCON and τJOIN state that type assumptions for the same identifier can be split or joined according to the type splitting relation $\mathbf{T} = \mathbf{T}_1 \circ \mathbf{T}_2$, also defined in Figure 4: apart from standard splitting of unrestricted channels, εUNR , and process types, εPROC , we note that a unique-after- i channel may be split into a unique-after- $(i + 1)$ channel and an affine channel; we also note that affine channels are *never* split. The environment structural rules also allow for weakening, τWEAK , equi-recursive manipulation of types, τEQ and εREC , and subtyping, τSUB ; the latter rule is defined in terms of the subtyping relation also stated in Figure 4 (bottom) where, for instance, an unrestricted channel can be used instead of an affine channel (that can be used at most once). The key novel structural rule is however τREV , which allows us to change (revise) the object type of a channel whenever we are guaranteed that the type assumption for that identifier is unique. These rules are recalled from [12] and the reader is encouraged to consult that document for more details.

The consistency condition of Definition 2.2 ensures that there is no mismatch in the duality between the guarantees of unique types and the restrictions of affine types, which allows sound compositional type-checking by our type system. For instance, consistency rules out environments such as

$$c : [\mathbf{U}]^\bullet, c : [\mathbf{U}]^1 \tag{2.4}$$

where a process typed under the guarantee that a channel c is unique now, $c : [\mathbf{U}]^\bullet$, contradicts the fact that some other process may be typed under the affine usage allowed by the assumption $c : [\mathbf{U}]^1$.

Logical rules

$$\begin{array}{c}
\frac{\Gamma, u: [\vec{\mathbf{T}}]^{a-1} \vdash P}{\Gamma, u: [\vec{\mathbf{T}}]^a, \vec{v}: \vec{\mathbf{T}} \vdash u! \vec{v}.P} \text{TOUT} \quad \frac{\Gamma, u: [\vec{\mathbf{T}}]^{a-1}, x: \vec{\mathbf{T}} \vdash P}{\Gamma, u: [\vec{\mathbf{T}}]^a \vdash u? \vec{x}.P} \text{TIN} \quad \frac{\Gamma_1 \vdash P \quad \Gamma_2 \vdash Q}{\Gamma_1, \Gamma_2 \vdash P \parallel Q} \text{TPAR} \\
\\
\frac{u, v \in \Gamma \quad \Gamma \vdash P \quad \Gamma \vdash Q}{\Gamma \vdash \text{if } u = v \text{ then } P \text{ else } Q} \text{TIF} \quad \frac{\Gamma^\omega, x: \mathbf{proc} \vdash P}{\Gamma^\omega \vdash \text{rec } w.P} \text{TREC} \quad \frac{}{x: \mathbf{proc} \vdash x} \text{TVAR} \\
\\
\frac{\Gamma \vdash P}{\Gamma, u: [\vec{\mathbf{T}}]^\bullet \vdash \text{free } u.P} \text{TFFREE} \quad \frac{\Gamma, x: [\vec{\mathbf{T}}]^\bullet \vdash P}{\Gamma \vdash \text{alloc } x.P} \text{TALL} \quad \frac{}{\emptyset \vdash \text{nil}} \text{TNIL} \\
\\
\frac{\Gamma' \vdash P \quad \Gamma < \Gamma'}{\Gamma \vdash P} \text{TSTR}
\end{array}$$

where Γ^ω can only contain unrestricted assumptions and all bound variables are fresh.

Structural rules ($<$) is the least reflexive transitive relation satisfying

$$\begin{array}{c}
\frac{\mathbf{T} = \mathbf{T}_1 \circ \mathbf{T}_2}{\Gamma, u: \mathbf{T} < \Gamma, u: \mathbf{T}_1, u: \mathbf{T}_2} \text{TCON} \quad \frac{\mathbf{T} = \mathbf{T}_1 \circ \mathbf{T}_2}{\Gamma, u: \mathbf{T}_1, u: \mathbf{T}_2 < \Gamma, u: \mathbf{T}} \text{TJOIN} \quad \frac{\mathbf{T}_1 \sim \mathbf{T}_2}{\Gamma, u: \mathbf{T}_1 < \Gamma, u: \mathbf{T}_2} \text{TEQ} \\
\\
\frac{}{\Gamma, u: \mathbf{T} < \Gamma} \text{TWEAK} \quad \frac{\mathbf{T}_1 <_s \mathbf{T}_2}{\Gamma, u: \mathbf{T}_1 < \Gamma, u: \mathbf{T}_2} \text{TSUB} \quad \frac{}{\Gamma, u: [\vec{\mathbf{T}}_1]^\bullet < \Gamma, u: [\vec{\mathbf{T}}_2]^\bullet} \text{TREV}
\end{array}$$

Equi-Recursion

$$\frac{}{\mu X.U \sim U\{\mu X.U/X\}} \text{EREC}$$

Counting channel usage

$$c: [\vec{\mathbf{T}}]^{a-1} \stackrel{\text{def}}{=} \begin{cases} \varepsilon & (\text{empty list}) \text{ if } a = \mathbf{1} \\ c: [\vec{\mathbf{T}}]^\omega & \text{if } a = \omega \\ c: [\vec{\mathbf{T}}]^{(\bullet, i)} & \text{if } a = (\bullet, i + 1) \end{cases}$$

Type splitting

$$\frac{}{[\vec{\mathbf{T}}]^\omega = [\vec{\mathbf{T}}]^\omega \circ [\vec{\mathbf{T}}]^\omega} \text{PUNR} \quad \frac{}{\mathbf{proc} = \mathbf{proc} \circ \mathbf{proc}} \text{PPROC} \quad \frac{}{[\vec{\mathbf{T}}]^{(\bullet, i)} = [\vec{\mathbf{T}}]^\mathbf{1} \circ [\vec{\mathbf{T}}]^{(\bullet, i+1)}} \text{PUNQ}$$

Subtyping

$$\frac{}{(\bullet, i) <_s (\bullet, i + 1)} \text{SINDX} \quad \frac{}{(\bullet, i) <_s \omega} \text{SUNQ} \quad \frac{}{\omega <_s \mathbf{1}} \text{SAFF} \quad \frac{a_1 <_s a_2}{[\vec{\mathbf{T}}]^{a_1} <_s [\vec{\mathbf{T}}]^{a_2}} \text{STYP}$$

Figure 4: Typing processes

For similar reasons, consistency also rules out environments such as

$$c : [\mathbf{U}]^\bullet, c : [\mathbf{U}]^\omega \quad (2.5)$$

However, it does not rule out environments such as (2.6) even though the guarantee provided by $c : [\mathbf{U}]^{(\bullet,2)}$ is too conservative: it states that channel c will become unique after *two* uses but, in actual fact, it becomes unique after one use since the (top-level) environment contains only *one* other affine type assumption, $c : [\mathbf{U}]^1$, that other processes can be typed at.

$$c : [\mathbf{U}]^{(\bullet,2)}, c : [\mathbf{U}]^1 \quad (2.6)$$

A less conservative uniqueness typing guarantee would therefore be $c : [\mathbf{U}]^{(\bullet,1)}$ as shown in (2.7) below; this environment constitutes another case of a consistent environment allowed by Definition 2.2.

$$c : [\mathbf{U}]^{(\bullet,1)}, c : [\mathbf{U}]^1 \quad (2.7)$$

The type system is *substructural*, implying that typing assumptions can be used *only once* during typechecking [37]. This is clearly manifested in the output and input rules, τOUT and τIN in Figure 4. In fact, using the operation $c : [\vec{\mathbf{T}}]^{a-1}$ (see⁵ Figure 4), rule τOUT collapses three different possibilities for typing output processes, which could alternatively have been expressed as the three separate typing rules in (2.8).

$$\begin{array}{c} \frac{\Gamma \vdash P}{\Gamma, u : [\vec{\mathbf{T}}]^1, \vec{v} : \vec{\mathbf{T}} \vdash u! \vec{v}.P} \tau\text{OUTA} \quad \frac{\Gamma, u : [\vec{\mathbf{T}}]^\omega \vdash P}{\Gamma, u : [\vec{\mathbf{T}}]^\omega, \vec{v} : \vec{\mathbf{T}} \vdash u! \vec{v}.P} \tau\text{OUTW} \\ \frac{\Gamma, u : [\vec{\mathbf{T}}]^{(\bullet,i)} \vdash P}{\Gamma, u : [\vec{\mathbf{T}}]^{(\bullet,i+1)}, \vec{v} : \vec{\mathbf{T}} \vdash u! \vec{v}.P} \tau\text{OUTU} \end{array} \quad (2.8)$$

Rule τOUTA states that an output of values \vec{v} on channel u is allowed if the type environment has an *affine* channel-type assumption for that channel, $u : [\vec{\mathbf{T}}]^1$, and the corresponding type assumptions for the values communicated, $\vec{v} : \vec{\mathbf{T}}$, match the object type of the affine channel-type assumption, $\vec{\mathbf{T}}$; in the rule premise, the continuation P must also be typed *wrt.* the *remaining* assumptions in the environment, *without* the assumptions consumed by the conclusion. Rule τOUTW is similar, but permits outputs on u for environments with an *unrestricted* channel-type assumption for that channel, $u : [\vec{\mathbf{T}}]^\omega$. The continuation P is typechecked *wrt.* the remaining assumptions and a *new* assumption, $u : [\vec{\mathbf{T}}]^\omega$; this assumption is identical to the one consumed in the conclusion, so as to model the fact that uses of channel u are unrestricted. Rule τOUTU is again similar, but it allows outputs on channel u for a “*unique after $i+1$* ” channel-type assumption; in the premise of the rule, P is typechecked *wrt.* the remaining assumptions and a *new* assumption $u : [\vec{\mathbf{T}}]^{(\bullet,i)}$, where u is now *unique after i* uses. Analogously, the input rule, τIN , also encodes three input cases (listed below):

$$\frac{\Gamma, x : \vec{\mathbf{T}} \vdash P}{\Gamma, u : [\vec{\mathbf{T}}]^1 \vdash u? \vec{x}.P} \tau\text{INO} \quad \frac{\Gamma, u : [\vec{\mathbf{T}}]^\omega, x : \vec{\mathbf{T}} \vdash P}{\Gamma, u : [\vec{\mathbf{T}}]^\omega \vdash u? \vec{x}.P} \tau\text{INW} \quad \frac{\Gamma, u : [\vec{\mathbf{T}}]^{(\bullet,i)}, x : \vec{\mathbf{T}} \vdash P}{\Gamma, u : [\vec{\mathbf{T}}]^{(\bullet,i+1)} \vdash u? \vec{x}.P} \tau\text{INU} \quad (2.9)$$

⁵This operation on type assumptions, $c : [\vec{\mathbf{T}}]^{a-1}$, defined in Figure 4, describes the cases where, when using an affine type assumption to typecheck a process, the continuation of the process in the rule premise is typed without that assumption (the operation returns no type assumption), whereas when using an unrestricted or unique-after- i assumptions, the premise judgement use *wrt.* (new) unrestricted and unique-after- $(i-1)$ assumptions, respectively. Note that the operation $c : [\vec{\mathbf{T}}]^{a-1}$ is not defined for $a = \bullet$. See [12] for more detail.

Parallel composition (τPAR) enforces the substructural treatment of type assumptions, by ensuring that type assumptions are used by either the left process or the right, but not by both. However, some type assumption can be *split* using contraction, *i.e.*, rules (τSTR) and (τCON). For example, an assumption $c : [\vec{\mathbf{T}}]^{(\bullet,i)}$ can be split as $c : [\vec{\mathbf{T}}]^1$ and $c : [\vec{\mathbf{T}}]^{(\bullet,i+1)}$ —see (pUNQ).

The rest of the rules in Figure 4 are fairly straightforward. Even though these typing rules do not require Γ to be consistent, the consistency requirement at the top level typing judgement (τSYS) ensures that whenever a process is typed *wrt.* a unique assumption for a channel, $[\vec{\mathbf{T}}]^\bullet$, no other process has access to that channel. It can therefore safely deallocate it (τFREE), or change the object type of the channel (τREV). Dually, when a channel is newly allocated it is assumed unique (τALL). Note also that name matching is only permitted when channel permissions are owned, $u, v \in \Gamma$ in (τIF). Uniqueness can therefore also be thought of as “freshness”, a claim we substantiate further in Section 4.2.

In [12] we prove the usual subject reduction and progress lemmas for this type system, given an (obvious) error relation.

Example 2.3. All client implementations discussed in Section 1 typecheck *wrt.* the type environment

$$\Gamma = \text{srv}_1 : [[\mathbf{T}_1]^1]^\omega, \text{srv}_2 : [[\mathbf{T}_2]^1]^\omega, \text{ret} : [\mathbf{T}_1, \mathbf{T}_2]^\omega.$$

For instance, to typecheck C_2 from (1.2), we can apply the typing rules τREC and τALL from Figure 4 to obtain the typing sequent:

$$\Gamma, w : \mathbf{proc}, x : [\mathbf{T}_1]^\bullet \vdash \text{srv}_1!x.x?y. \text{srv}_2!x.x?z. \text{free } x. \text{ret}!(y, z). w \quad (2.10)$$

Using the environment structural rules (*i.e.*, τCON) we can split the type assumption for x :

$$\Gamma, w : \mathbf{proc}, x : [\mathbf{T}_1]^\bullet < \Gamma, w : \mathbf{proc}, x : [\mathbf{T}_1]^1, x : [\mathbf{T}_1]^{(\bullet,1)}$$

Using τSTR and τOUT we can type (2.10) to obtain

$$\Gamma, w : \mathbf{proc}, x : [\mathbf{T}_1]^{(\bullet,1)} \vdash x?y. \text{srv}_2!x.x?z. \text{free } x. \text{ret}!(y, z). w$$

After applying τIN to typecheck the input, we are left with the sequent

$$\Gamma, w : \mathbf{proc}, x : [\mathbf{T}_1]^\bullet, y : \mathbf{T}_1 \vdash \text{srv}_2!x.x?z. \text{free } x. \text{ret}!(y, z). w$$

In particular, we note that the input typing rule stipulates that the input continuation process needs to type *wrt.* the following type assumption for $x : [\mathbf{T}_1]^{(\bullet,1)-1}$ which is equal to $x : [\mathbf{T}_1]^\bullet$. Since x is unique now, we can change the object type from \mathbf{T}_1 to \mathbf{T}_2 using τREV , which allows us to type the interactions with srv_2 in analogous fashion. This leaves us with

$$\Gamma, w : \mathbf{proc}, x : [\mathbf{T}_2]^\bullet, y : \mathbf{T}_1, z : \mathbf{T}_2 \vdash \text{free } x. \text{ret}!(y, z). w$$

which we can discharge using rules τFREE , τOUT and τVAR .

3. A CASE STUDY

Resource management is particularly relevant to programs manipulating (unbounded) regular structures. We consider the concurrent implementation of an unbounded buffer, Buff , receiving values to

queue on channel `in` and dequeuing values by outputting on channel `out`.

$$\begin{aligned} \text{Buff} &\stackrel{\text{def}}{=} \text{in?}y. \text{ alloc } z. (\text{Frn} \parallel b!z \parallel c_1!(y, z)) \parallel c_1?(y, z). \text{ out!}y. (\text{Bck} \parallel d!z) \\ \text{Frn} &\stackrel{\text{def}}{=} \text{rec } w. b?x. \text{ in?}y. \text{ alloc } z. (w \parallel b!z \parallel x!(y, z)) \\ \text{Bck} &\stackrel{\text{def}}{=} \text{rec } w. d?x. x?(y, z). \text{ out!}y. (w \parallel d!z) \end{aligned}$$

In order to decouple input requests from output requests while still preserving the order of inputted values, the process handling inputs in $\text{Buff}, \text{in?}y. \text{ alloc } z. (\text{Frn} \parallel b!z \parallel c_1!(y, z))$, stores inputted values v_1, \dots, v_n as a queue of interconnected outputs

$$c_1!(v_1, c_2) \parallel \dots \parallel c_n!(v_n, c_{n+1}) \quad (3.1)$$

on the internal⁶ channels c_1, \dots, c_{n+1} . The process handling the outputs, $c_1?(y, z). \text{ out!}y. (\text{Bck} \parallel d!z)$, then reads from the head of this queue, *i.e.*, the output on channel c_1 , so as to obtain the first value inputted, v_1 , and the next head of the queue, c_2 . The input and output processes are defined in terms of the recursive processes, Frn and Bck *resp.*, which are parameterised by the channel to output (*resp.* input) on next through the channels b and d .⁷

Since the buffer is *unbounded*, the number of internal channels used for the queue of interconnected outputs, (3.1), is not fixed and these channels cannot therefore be created up front. Instead, they are created on demand by the input process for every value inputted, using the $R\pi$ construct $\text{alloc } z.P$. The newly allocated channel z is then passed on the next iteration of Frn through channel b , $b!z$, and communicated as the next head of the queue when adding the subsequent queue item; this is received by the output process when it inputs the value at the head of the chain and passed on the next iteration of Bck through channel d , $d!z$.

3.1. Typeability and behaviour of the Buffer. Our unbounded buffer implementation, Buff , can be typed *wrt.* the type environment

$$\Gamma_{\text{int}} \stackrel{\text{def}}{=} \text{in}:[\mathbf{T}]^\omega, \text{out}:[\mathbf{T}]^\omega, b:[\mathbf{T}_{\text{rec}}]^\omega, d:[\mathbf{T}_{\text{rec}}]^\omega, c_1:[\mathbf{T}, \mathbf{T}_{\text{rec}}]^\bullet \quad (3.2)$$

where \mathbf{T} is the type of the values stored in the buffer and \mathbf{T}_{rec} is a recursive type defined as

$$\mathbf{T}_{\text{rec}} \stackrel{\text{def}}{=} \mu X. [\mathbf{T}, X]^{(\bullet, 1)}.$$

This recursive type is used to type the internal channels c_1, \dots, c_{n+1} — recall that in (3.1) these channels carry channels of the same kind in order to link to one another as a chain of outputs. In particular, using the typing rules of Section 2 we can prove the following typing judgements:

$$\text{in}:[\mathbf{T}]^\omega, b:[\mathbf{T}_{\text{rec}}]^\omega, c_1:[\mathbf{T}, \mathbf{T}_{\text{rec}}]^\bullet \vdash \text{in?}y. \text{ alloc } z. (\text{Frn} \parallel b!z \parallel c_1!(y, z)) \quad (3.3)$$

$$\text{out}:[\mathbf{T}]^\omega, d:[\mathbf{T}_{\text{rec}}]^\omega, c_1:[\mathbf{T}, \mathbf{T}_{\text{rec}}]^{(\bullet, 1)} \vdash c_1?(y, z). \text{ out!}y. (\text{Bck} \parallel d!z) \quad (3.4)$$

From the perspective of a user of the unbounded buffer, Buff implements the interface defined by the environment

$$\Gamma_{\text{ext}} \stackrel{\text{def}}{=} \text{in}:[\mathbf{T}]^\omega, \text{out}:[\mathbf{T}]^\omega$$

abstracting away from the implementation channels b, d and c_1 .

⁶Subsequent allocated channels are referred to as c_2, c_3 , *etc.*.

⁷This models parametrisable process definitions $\text{Frn}(x)$ and $\text{Bck}(x)$ within our language.

3.2. A resource-conscious Implementation of the Buffer. When the buffer implementation of `Buff` retrieves values from the head of the internal queue, e.g., (3.1), the channel holding the initial value, i.e., c_1 in (3.1), is never reused again even though it is left allocated in memory. This fact will repeat itself for every value that is stored and retrieved from the buffer and amounts to the equivalent of a “*memory leak*”. A more resource-conscious implementation of the unbounded buffer is `eBuff`, defined in terms of the previous input process used for `Buff`, and a modified output process, $c_1?(y, z).\text{free } c_1.\text{out!}y.(eBk \parallel d!z)$, which uses the tweaked recursive process, `eBk`.

$$\begin{aligned} eBuff &\stackrel{\text{def}}{=} \text{in?}y.\text{alloc } z.(\text{Frn} \parallel b!z \parallel c_1!(y, z)) \parallel c_1?(y, z).\text{free } c_1.\text{out!}y.(eBk \parallel d!z) \\ eBk &\stackrel{\text{def}}{=} \text{rec } w. d?x. x?(y, z). \text{free } x. \text{out!}y.(w \parallel d!z) \end{aligned}$$

The main difference between `Buff` and `eBuff` is that the latter deallocates the channel at the head of the internal chain once it is consumed. We can typecheck `eBuff` as safe since no other process uses the internal channels making up the chain after deallocation. More specifically, the typeability of `eBuff` wrt. Γ_{int} of (3.2) follows from (3.3) and the type judgement below:

$$\text{out}:[\mathbf{T}]^\omega, d:[\mathbf{T}_{\text{rec}}]^\omega, c_1:[\mathbf{T}, \mathbf{T}_{\text{rec}}]^{(\bullet, 1)} \vdash c_1?(y, z). \text{free } c_1. \text{out!}y. (\text{Bck} \parallel d!z)$$

Note that by the typing rule τIN of Figure 4, we need to typecheck the continuation of the input process, $\text{free } c_1. \text{out!}y. (\text{Bck} \parallel d!z)$ wrt. the type environment

$$\text{out}:[\mathbf{T}]^\omega, d:[\mathbf{T}_{\text{rec}}]^\omega, c_1:[\mathbf{T}, \mathbf{T}_{\text{rec}}]^\bullet, y:\mathbf{T}, z:\mathbf{T}_{\text{rec}}$$

where, in particular, c_1 is now assigned a *unique* channel type. According to the typing rule τFREE , this suffices to safely type the respective deallocation of c_1 .

4. A COST-BASED PREORDER

We define our cost-based preorder as a *bisimulation relation* that relates two systems $M \triangleright P$ and $N \triangleright Q$ whenever they have equivalent behaviour and when, in addition, $M \triangleright P$ is more efficient than $N \triangleright Q$. We are interested in reasoning about *safe* computations, aided by the type system described in Section 2. For this reason, we limit our analysis to instances of $M \triangleright P$ and $N \triangleright Q$ that are *well-typed*, i.e., that there exist (consistent) environments Δ, Δ' such that $\Delta \vdash M \triangleright P$ and $\Delta' \vdash N \triangleright Q$. In order to preserve safety, we also need to reason under the assumption of safe contexts. Again, we employ the type system described in Section 2 and characterise the (safe) context through a type environment that typechecks it, Γ_{obs} . Thus our bisimulation relations take the form of a typed relation, indexed by type environments [21]:

$$\Gamma_{\text{obs}} \vDash (M \triangleright P) \mathcal{R} (N \triangleright Q) \tag{4.1}$$

Behavioural reasoning for safe systems is achieved by ensuring that the overall type environment $(\Gamma_{\text{sys}}, \Gamma_{\text{obs}})$, consisting of the environment typing $M \triangleright P$ and $N \triangleright Q$, say Γ_{sys} , and the observer environment Γ_{obs} , is *consistent* according to Definition 2.2. This means that there exists a global environment, Γ_{global} , which can be decomposed into Γ_{obs} and Γ_{sys} ; it also means that the observer process, which is universally quantified by our semantic interpretation (4.1), typechecks when composed in parallel with P , resp. Q (see τPAR of Figure 4).

There is one other complication worth highlighting regarding (4.1): although both systems $M \triangleright P$ and $N \triangleright Q$ are related wrt. the same *observer*, Γ_{obs} , they can each be typed under *different*

typing environments. For instance, consider the two clients C_0 and C_1 we would like to relate from the introduction:

$$\begin{aligned} C_0 &\triangleq \text{rec } w. \text{ alloc } x_1. \text{ alloc } x_2. \text{ srv}_1!x_1. x_1?y. \text{ srv}_2!x_2. x_2?z. c!(y, z).w \\ C_1 &\triangleq \text{rec } w. \text{ alloc } x. \text{ srv}_1!x. x?y. \text{ srv}_2!x. x?z. c!(y, z).w \end{aligned} \quad (4.2)$$

Even though, initially, they may be typed by the same type environment, after a few steps, the derivatives of C_0 and C_1 must be typed under different typing environments, because C_0 allocates two channels, while C_1 only allocates a single channel. Our typed relations allows for this by *existentially quantifying* over the type environments typing the respective systems. All this is achieved indirectly through the use of *configurations*.

Definition 4.1 (Configuration). The triple $\Gamma \triangleleft M \triangleright P$ is a configuration if and only if $\text{dom}(\Gamma) \subseteq M$ and there exist some Δ such that (Γ, Δ) is consistent and $\Delta \vdash M \triangleright P$.

Note that, in a configuration $\Gamma \triangleleft M \triangleright P$ (where Γ types some implicit observer):

- $c \in (\text{dom}(\Gamma) \cup \text{names}(P))$ implies $c \in M$ i.e., M is a global resource environment accounting for both P and Γ .
- $c \in M$ and $c \notin (\text{dom}(\Gamma) \cup \text{names}(P))$ denotes a resource leak for channel c .
- $c \notin \text{dom}(\Gamma)$ implies that channel c is not known to the observer; in some sense, this mimics name scoping in more standard π -calculus settings.

Definition 4.2 (Typed Relation). A type-indexed relation \mathcal{R} relates systems under a observer characterized by a context Γ ; we write

$$\Gamma \vDash M \triangleright P \mathcal{R} N \triangleright Q$$

if \mathcal{R} relates $\Gamma \triangleleft M \triangleright P$ and $\Gamma \triangleleft N \triangleright Q$, and both $\Gamma \triangleleft M \triangleright P$ and $\Gamma \triangleleft N \triangleright Q$ are configurations.

4.1. Labelled Transition System. In order to be able to reason coinductively over our typed relations, we define a labelled transition system (LTS) over configurations. Apart from describing the behaviour of the system $M \triangleright P$ in a configuration $\Gamma \triangleleft M \triangleright P$, the LTS also models interactions between the system and an observer typed under Γ . Our LTS is also *costed*, assigning a cost to each form of transition.

The costed LTS, whose actions take the form $\xrightarrow{\mu}_k$, is defined in Figure 5, in terms of a top-level rule, LREN , and a pre-LTS, denoted as $\xrightarrow{\mu}_k$. The rule LREN allows us to rename channels for transitions derived in the pre-LTS, as long as this renaming is invisible to the observer, and is comparable to alpha-renaming of scoped bound names in the standard π -calculus. It relies on the renaming-modulo (observer) type environments given in Definition 4.3.

Definition 4.3 (Renaming Modulo Γ). Let $\sigma_\Gamma : \text{NAME} \mapsto \text{NAME}$ range over bijective name substitutions satisfying the constraint that $c \in \text{dom}(\Gamma)$ implies $c\sigma_\Gamma = c\sigma_\Gamma^{-1} = c$.

The renaming introduced by LREN allows us to relate the clients C_0 and C_1 of (4.2) wrt. an observer environment such as $\text{srv}_1 : [[\mathbf{T}_1]^1]^\omega$, $\text{srv}_2 : [[\mathbf{T}_2]^1]^\omega$ of (1.6) and some appropriate common set of resources M even when, after the initial channel allocations, the two clients communicate potentially different (newly allocated) channels on srv_1 . The rule is particularly useful when, later on, we need to also match the output of a new allocated channel on srv_2 from C_0 with the output on the previously allocated channel from C_1 on srv_2 . The renaming-modulo observer environments function can be used for C_1 at that stage — even though the client reuses a channel previously communicated to the observer — because the respective observer information relating to that channel is lost, i.e., it is not in the domain of the observer environment; see discussion for LIn and LIn

below for an explanation of how observers lose information. This mechanism differs from standard scope-extrusion techniques for π -calculus which assume that, once a name has been extruded, it remains forever known to the observer. As a result, there are more opportunities for renaming in our calculus than there are in the standard π -calculus.

To ensure that only safe interactions are specified, the (pre-)LTS must be able to reason compositionally about resource usage between the process, P , and the observer, Γ . We therefore imbue our type assumptions from Section 2 with a *permission semantics*, in the style of [42, 13]. Under this interpretation, type assumptions constitute *permissions* describing the respective usage of resources. Permissions are woven into the behaviour of configurations giving them an *operational* role: they may either restrict usage or privilege processes to use resources in special ways. In a configuration, the observer and the process each *own* a set of permissions and may *transfer* them to one another during communication. The consistency requirement of a configuration ensures that the guarantees given by permissions owned by the observer are not in conflict with those given by permissions owned by the configuration process, and viceversa.

To understand how the pre-LTS deals with permission transfer and compositional resource usage, consider the rule for output, (LOUT). Since we employ the type system of Section 2 to ensure safety, this rule models the typing rule for output (τ OUT) on the part of the process, and the typing rule for input (τ IN) on the part of the observer. Thus, apart from describing the communication of values \vec{d} from the configuration process to the observer on channel c , it also captures permission transfer between the two parties, mirroring the type assumption usage in τ OUT and τ IN. More specifically, rule (LOUT) employs the operation $c : [\vec{\mathbf{T}}]^{a-1}$ of Figure 4 so as to concisely describe the three variants of the output rule:

$$\begin{array}{c}
 \frac{}{\Gamma, c : [\vec{\mathbf{T}}]^{(\bullet, i+1)} \triangleleft M \triangleright c! \vec{d}. P \quad \frac{c! \vec{d}}{\rightarrow_0} \quad \Gamma, c : [\vec{\mathbf{T}}]^{(\bullet, i)}, \vec{d} : \vec{\mathbf{T}} \triangleleft M \triangleright P} \text{LOUTU} \\
 \frac{}{\Gamma, c : [\vec{\mathbf{T}}]^1 \triangleleft M \triangleright c! \vec{d}. P \quad \frac{c! \vec{d}}{\rightarrow_0} \quad \Gamma, \quad \vec{d} : \vec{\mathbf{T}} \triangleleft M \triangleright P} \text{LOUTA} \\
 \frac{}{\Gamma, c : [\vec{\mathbf{T}}]^\omega \triangleleft M \triangleright c! \vec{d}. P \quad \frac{c! \vec{d}}{\rightarrow_0} \quad \Gamma, c : [\vec{\mathbf{T}}]^\omega, \vec{d} : \vec{\mathbf{T}} \triangleleft M \triangleright P} \text{LOUTW}
 \end{array} \tag{4.3}$$

The first output rule variant, LOUTU, deals with the case where the observer owns a unique-after- $(i+1)$ permission for channel c . Definition 4.1 implies that the process in the configuration is well-typed (*wrt.* some environment) and, since the process is in a position to output on channel c , rule τ OUT must have been used to type it. This typing rule, in turn, states that the type assumptions relating to the values communicated, $\vec{d} : \vec{\mathbf{T}}$, must have been owned by the process and consumed by the output operation. Dually, since the observer is capable of inputting on c , rule τ IN must have been used to type it,⁸ which states that the continuation (after the input) assumes the use the assumptions $\vec{d} : \vec{\mathbf{T}}$. Rule LOUTU models these two usages operationally as the *explicit transfer* of the permissions $\vec{d} : \vec{\mathbf{T}}$ from the process to the observer.

The rule also models the *implicit transfer* of permissions between the observer and the output process. More precisely, Definition 4.1 requires that the process is typed *wrt.* an environment that *does not conflict with* the observer environment, which implies that the process environment must have (necessarily) used an affine permission, $c : [\vec{\mathbf{T}}]^1$, for outputting on channel c .⁹ In fact, any other

⁸More specifically, τ INU of (2.9).

⁹This implies that τ OUTA of (2.8) was used when typing the process

Costed Transitions and pre-Transitions

$$\begin{array}{c}
\frac{\Gamma \triangleleft (M \triangleright P) \sigma_{\Gamma} \xrightarrow{\mu}_k \Gamma' \triangleleft M' \triangleright P'}{\Gamma \triangleleft M \triangleright P \xrightarrow{\mu}_k \Gamma' \triangleleft M' \triangleright P'} \text{LREN} \\
\\
\frac{}{\Gamma, c: [\vec{\mathbf{T}}]^a \triangleleft M \triangleright c! \vec{d}. P \xrightarrow{c! \vec{d}}_0 \Gamma, c: [\vec{\mathbf{T}}]^{a-1}, \vec{d}: \vec{\mathbf{T}} \triangleleft M \triangleright P} \text{LOUT} \\
\\
\frac{}{\Gamma, c: [\vec{\mathbf{T}}]^a, \vec{d}: \vec{\mathbf{T}} \triangleleft M \triangleright c? \vec{x}. P \xrightarrow{c? \vec{d}}_0 \Gamma, c: [\vec{\mathbf{T}}]^{a-1} \triangleleft M \triangleright P\{\vec{d}/\vec{x}\}} \text{LIN} \\
\\
\frac{\Gamma_1 \triangleleft M \triangleright P \xrightarrow{c! \vec{d}}_0 \Gamma'_1 \triangleleft M \triangleright P' \quad \Gamma_2 \triangleleft M \triangleright Q \xrightarrow{c? \vec{d}}_0 \Gamma'_2 \triangleleft M \triangleright Q'}{\Gamma \triangleleft M \triangleright P \parallel Q \xrightarrow{\tau}_0 \Gamma \triangleleft M \triangleright P' \parallel Q'} \text{LCOM-L} \\
\\
\frac{\Gamma \triangleleft M \triangleright P \xrightarrow{\mu}_k \Gamma' \triangleleft M' \triangleright P'}{\Gamma \triangleleft M \triangleright P \parallel Q \xrightarrow{\mu}_k \Gamma' \triangleleft M' \triangleright P' \parallel Q} \text{LPAR-L} \\
\\
\frac{\Gamma < \Gamma'}{\Gamma \triangleleft M \triangleright P \xrightarrow{\text{env}}_0 \Gamma' \triangleleft M \triangleright P} \text{LSTR} \quad \frac{}{\Gamma \triangleleft M \triangleright \text{rec } w. P \xrightarrow{\tau}_0 \Gamma \triangleleft M \triangleright P\{\text{rec } w. P/w\}} \text{LREC} \\
\\
\frac{}{\Gamma \triangleleft M, c \triangleright \text{if } c = c \text{ then } P \text{ else } Q \xrightarrow{\tau}_0 \Gamma \triangleleft M, c \triangleright P} \text{LTHEN} \\
\\
\frac{}{\Gamma \triangleleft M, c, d \triangleright \text{if } c = d \text{ then } P \text{ else } Q \xrightarrow{\tau}_0 \Gamma \triangleleft M, c, d \triangleright Q} \text{LELSE} \\
\\
\frac{}{\Gamma \triangleleft M \triangleright \text{alloc } x. P \xrightarrow{\tau}_{+1} \Gamma \triangleleft M, c \triangleright P\{c/x\}} \text{LALL} \quad \frac{}{\Gamma \triangleleft M \triangleright P \xrightarrow{\text{alloc}}_{+1} \Gamma, c: [\vec{\mathbf{T}}]^{\bullet} \triangleleft M, c \triangleright P} \text{LALLE} \\
\\
\frac{}{\Gamma \triangleleft M, c \triangleright \text{free } c. P \xrightarrow{\tau}_{-1} \Gamma \triangleleft M \triangleright P} \text{LFREE} \quad \frac{}{\Gamma, c: [\mathbf{T}]^{\bullet} \triangleleft M, c \triangleright P \xrightarrow{\text{free } c}_{-1} \Gamma \triangleleft M \triangleright P} \text{LFREEE} \\
\\
\text{Weak (Cost-Accumulating) Transitions} \\
\\
\frac{\Gamma \triangleleft M \triangleright P \xrightarrow{\mu}_k \Delta \triangleleft N \triangleright Q}{\Gamma \triangleleft M \triangleright P \xRightarrow{\mu}_k \Delta \triangleleft N \triangleright Q} \text{wTRA} \quad \frac{\Gamma \triangleleft M \triangleright P \xrightarrow{\tau}_l \Gamma' \triangleleft M' \triangleright P \xRightarrow{\mu}_k \Gamma'' \triangleleft N \triangleright Q}{\Gamma \triangleleft M \triangleright P \xRightarrow{\mu}_{(l+k)} \Gamma'' \triangleleft N \triangleright Q} \text{wLEFT} \\
\\
\frac{\Gamma \triangleleft M \triangleright P \xRightarrow{\mu}_l \Gamma' \triangleleft M' \triangleright P \xrightarrow{\tau}_k \Gamma'' \triangleleft N \triangleright Q}{\Gamma \triangleleft M \triangleright P \xRightarrow{\mu}_{(l+k)} \Gamma'' \triangleleft N \triangleright Q} \text{wRIGHT}
\end{array}$$

Figure 5: LTS Process Moves

type of permission would conflict with the unique-after- $(i+1)$ permission for channel c owned by the observer. Moreover, through the guarantee given by the permission used, $c : [\vec{T}]^{(\bullet, i+1)}$, the observer knows that, after the communication, it is one step closer towards gaining exclusive permission for channel c . Rule LOutU models all this as the (implicit) transfer of the affine permission $c : [\vec{T}]^1$ from the process to the observer, updating the observer's permission for c to $[\vec{T}]^{(\bullet, i)}$ — note that two permissions $c : [\vec{T}]^{(\bullet, i+1)}$, $c : [\vec{T}]^1$ can be consolidated as $c : [\vec{T}]^{(\bullet, i)}$ using the structural rules τJoin and pUnq of Figure 4.

The second output rule variant of (4.3), LOutA , is similar to the first when modelling the explicit transfer of permissions $\vec{d} : \vec{T}$ from the process to the observer. However, it describes a different implicit transfer of permissions, since the observer uses an affine permission to input from the configuration process on channel c . The rule caters for two possible subcases. In the first case, the process could have used a unique-after- $(i+1)$ permission when typed using τOut : this constitutes a dual case to that of rule LOutU , and the rule models the implicit transfer of the affine permission $c : [\vec{T}]^1$ in the *opposite* direction, *i.e.*, from the observer to the process. In the second case, the process could have used an affine or an unrestricted permission instead, which does not result in any implicit permission transfer, but merely the consumption of affine permissions. Since the environment on the process side is existentially quantified in a configuration, this difference is abstracted away and the two subcases are handled by the same rule variant. Note that, in the extreme case where the observer affine permission is the only one relating to channel c , the observer loses all knowledge of channel c .

The explicit permission transfer for LOutW of (4.3), is identical to the other two rule variants. The use of an unrestricted permission for c from the part of the observer, $c : [\vec{T}]^\omega$, implies that the output process could have either used an affine or an unrestricted permission—see (2.5). In either case, there is no implicit permission transfer involved. Moreover, the observer permission is not consumed since it is unrestricted.

The pre-LTS rule LIn can also be expanded into three rule variants, and models analogous permission transfer between the observer and the input process. Importantly, however, the *explicit* permission transfer described is *in the opposite direction* to that of LOut , namely from the observer to the input process. As in the case of LOutA of (4.3), the permission transfer from the observer to the input process may result in the observer losing all knowledge relating to the channels communicated, \vec{d} .

In order to allow an internal communication step through either LCom-L , or its dual LCom-R (elided), the left process should be considered to be part of the “observer” of the right process, and vice versa. However, it is not necessary to be quite so precise; we can follow [19] and consider an arbitrary observer instead. More explicitly, the rule states that if we can find observer environments (Γ_1 and Γ_2) to induce the respective input and output actions from separate constituent processes making up the system, we can then express these separate interactions as a single synchronous interaction; since this interaction is internal, it is independent of the environment representing the observer in the conclusion, Γ . See [19] for more justification.

In our LTS, both the process (LAll , LFree) and the observer (LAlLe , LFreeE) can allocate and deallocate memory. Finally, since the observer is modelled exclusively by the permissions it owns, we must allow the observer to split these permissions when necessary (LStr). The only rules that may alter the observer environment are those corresponding to external actions *i.e.*, LIn , LOut , LAlLe , LFreeE and LStr . The remaining axioms in the pre-LTS model reduction rules from Figure 2 and should be self-explanatory; note that, as in the reduction semantics, the only actions carrying

a cost are those describing allocation and deallocation, where the respective costs associated are inherited directly from the reduction semantics of Section 2.

In Figure 5 we also specify weak costed transitions for configurations, based on the transitions of our LTS (rule $w\text{TRA}$). As is standard, the relation denotes actions padded by τ -transitions to the left and right. However, it also *accumulates* the costs of the respective transitions into one aggregate cost for the entire weak action (rules $w\text{LEFT}$ and $w\text{RIGHT}$).

Technically, the pre-LTS is defined over triples Γ, M, P rather than configurations $\Gamma \triangleleft M \triangleright P$, but we can prove that the pre-LTS rules preserve the requirements for such triples to be configurations; see Lemma 4.5.

Lemma 4.4 (Transition and Structure). $\Gamma \triangleleft M \triangleright P \xrightarrow{\mu}_k \Gamma' \triangleleft M' \triangleright P'$ and for Δ consistent $\Delta \vdash M \triangleright P$ implies the cases:

If $\mu = c!d$: $M = M', k = 0, P \equiv c!d.P_1 \parallel P_2, P' \equiv P_1 \parallel P_2$ and $\Gamma = (\Gamma'', c: [\vec{\mathbf{T}}]^a)$,
 $\Gamma' = (\Gamma'', c: [\vec{\mathbf{T}}]^{a-1}, d: \vec{\mathbf{T}})$ and $\Delta < (\Delta', c: [\vec{\mathbf{T}}]^b, d: \vec{\mathbf{T}})$, $(\Delta', c: [\vec{\mathbf{T}}]^{b-1}) \vdash P'$
 for some $P_1, P_2, \Gamma'', b, \vec{\mathbf{T}}$ and Δ' .

If $\mu = c?d$: $M = M', k = 0, P \equiv c?d.P_1 \parallel P_2, P' \equiv P_1 \{d/\bar{x}\} \parallel P_2$ and
 $\Gamma = (\Gamma'', c: [\vec{\mathbf{T}}]^a, d: \vec{\mathbf{T}}), \Gamma' = (\Gamma'', c: [\vec{\mathbf{T}}]^{a-1})$ and $\Delta < (\Delta', c: [\vec{\mathbf{T}}]^b)$,
 $(\Delta', c: [\vec{\mathbf{T}}]^{b-1}, d: \vec{\mathbf{T}}) \vdash P'$ for some $P_1, P_2, \Gamma'', b, \vec{\mathbf{T}}$ and Δ' .

If $\mu = \tau$: Either of three cases hold :

- $M = M', k = 0$ and $\Gamma = \Gamma'$ and $\Delta \vdash P'$ or;
- $M = (M', c), k = -1$ and $P \equiv \text{free } c.P_1 \parallel P_2, P' \equiv P_1 \parallel P_2, \Gamma = \Gamma'$ and $\Delta < \Delta', c: [\vec{\mathbf{T}}]^\bullet$
 where $\Delta' \vdash P'$ (for some $P_1, P_2, \vec{\mathbf{T}}$ and Δ') or;
- $M' = (M, c), k = +1$ and $P \equiv \text{alloc } x.P_1 \parallel P_2, P' \equiv P_1 \{c/x\} \parallel P_2$ and $\Gamma = \Gamma'$ and $\Delta < \Delta'$
 and $\Delta', c: [\vec{\mathbf{T}}]^\bullet \vdash P'$ (for some $P_1, P_2, \vec{\mathbf{T}}$ and Δ')

If $\mu = \text{free } c$: $M = (M', c), k = -1$ and $\Gamma = \Gamma', c: [\vec{\mathbf{T}}]^\bullet$ and $P = P'$ for some $\vec{\mathbf{T}}$.

If $\mu = \text{alloc}$: $M' = (M, c), k = +1$ and $\Gamma, c: [\vec{\mathbf{T}}]^\bullet = \Gamma'$ and $P = Q$ for some $\vec{\mathbf{T}}$.

If $\mu = \text{env}$: $\Gamma < \Gamma', M = M', k = 0$ and $P = P'$

Proof. By rule induction on $\Gamma \triangleleft M \triangleright P \xrightarrow{\mu}_k \Gamma' \triangleleft M' \triangleright P'$ □

Lemma 4.5 (Subject reduction). If $\Gamma \triangleleft M \triangleright P$ is a configuration and $\Gamma \triangleleft M \triangleright P \xrightarrow{\mu}_k \Delta \triangleleft N \triangleright Q$ then $\Delta \triangleleft N \triangleright Q$ is also a configuration.

Proof. We assume that $\text{dom}(\Gamma) \subseteq M$ and that there exists Δ such that Γ, Δ is consistent and that $\Delta \vdash M \triangleright P$. The rest of the proof follows from Lemma 4.4 (Transition and Structure), by case analysis of μ . □

As a consistency check, we can also show that our LTS semantics is in accordance with the reduction semantics presented in 2. In particular, τ -transitions correspond to reductions modulo renaming and process structural equivalence.

Lemma 4.6 (Reduction and Silent Transitions).

(1) $M \triangleright P \xrightarrow{k} M' \triangleright P'$ implies $\Gamma \triangleleft M \triangleright P \xrightarrow{\tau}_k \Gamma \triangleleft M' \triangleright P'$ for arbitrary Γ where $P'' \equiv P'$.

(2) $\Gamma \triangleleft M \triangleright P \xrightarrow{\tau}_k \Delta \triangleleft M' \triangleright P'$ implies $(M \triangleright P)\sigma_\Gamma \xrightarrow{k} M' \triangleright P'$ for some σ_Γ .

Proof. By rule induction on $M \triangleright P \xrightarrow{k} M' \triangleright P'$ and $\Gamma \triangleleft M \triangleright P \xrightarrow{\tau}_k \Delta \triangleleft M' \triangleright P'$. □

Example 4.7. Recall the buffer implementation `Buff` from Section 3 and the respective external environment Γ_{ext} defined in Section 3.1. The transition rules of Figure 5 allow us to derive the following behaviour for the configuration $\Gamma_{\text{ext}} \triangleleft M, c_1 \triangleright \text{Buff}$ (where $\text{in}, \text{out}, b, d \in M$):

$$\Gamma_{\text{ext}} \triangleleft M, c_1 \triangleright \text{Buff} \xrightarrow{\text{in}^?v_1}_0 \Gamma_{\text{ext}} \triangleleft M, c_1 \triangleright \left(\begin{array}{l} \text{alloc } z. (\text{Frn} \parallel b!z \parallel c_1!(v_1, z)) \\ \parallel c_1?(y, z). \text{out}!y. (\text{Bck} \parallel d!z) \end{array} \right) \quad (4.4)$$

$$\xrightarrow{\tau}_{+1} \Gamma_{\text{ext}} \triangleleft M, c_1, c_2 \triangleright \left(\begin{array}{l} (\text{Frn} \parallel b!c_2 \parallel c_1!(v_1, c_2)) \\ \parallel c_1?(y, z). \text{out}!y. (\text{Bck} \parallel d!z) \end{array} \right) \quad (4.5)$$

$$= \Gamma_{\text{ext}} \triangleleft M, c_1, c_2 \triangleright \left(\begin{array}{l} \text{rec } w. b?x. \text{in}^?y. \text{alloc } z. (w \parallel b!z \parallel x!(y, z)) \\ \parallel b!c_2 \parallel c_1!(v_1, c_2) \\ \parallel c_1?(y, z). \text{out}!y. (\text{Bck} \parallel d!z) \end{array} \right)$$

$$\xrightarrow{\tau}_0 \Gamma_{\text{ext}} \triangleleft M, c_1, c_2 \triangleright \left(\begin{array}{l} b?x. \text{in}^?y. \text{alloc } z. (\text{Frn} \parallel b!z \parallel x!(y, z)) \\ \parallel b!c_2 \parallel c_1!(v_1, c_2) \\ \parallel c_1?(y, z). \text{out}!y. (\text{Bck} \parallel d!z) \end{array} \right) \quad (4.6)$$

$$\xrightarrow{\tau}_0 \Gamma_{\text{ext}} \triangleleft M, c_1, c_2 \triangleright \left(\begin{array}{l} \text{in}^?y. \text{alloc } z. (\text{Frn} \parallel b!z \parallel c_2!(y, z)) \\ \parallel c_1!(v_1, c_2) \\ \parallel c_1?(y, z). \text{out}!y. (\text{Bck} \parallel d!z) \end{array} \right) \quad (4.7)$$

$$\xrightarrow{\text{in}^?v_2}_{+1} \Gamma_{\text{ext}} \triangleleft M, c_1, c_2, c_3 \triangleright \left(\begin{array}{l} \text{in}^?y. \text{alloc } z. (\text{Frn} \parallel b!z \parallel c_3!(y, z)) \\ \parallel c_1!(v_1, c_2) \parallel c_2!(v_2, c_3) \\ \parallel c_1?(y, z). \text{out}!y. (\text{Bck} \parallel d!z) \end{array} \right) \quad (4.8)$$

$$\xrightarrow{\text{out}!v_1}_0 \Gamma_{\text{ext}} \triangleleft M, c_1, c_2, c_3 \triangleright \left(\begin{array}{l} \text{in}^?y. \text{alloc } z. (\text{Frn} \parallel b!z \parallel c_3!(y, z)) \\ \parallel c_2!(v_2, c_3) \\ \parallel c_2?(y, z). \text{out}!y. (\text{Bck} \parallel d!z) \end{array} \right) \quad (4.9)$$

Transition (4.4) describes an input from the user whereas (4.5) allocates a new internal channel, c_2 , followed by a recursive process unfolding, (4.6), and the instantiation of the unfolded process with the newly allocated channel c_2 , (4.7), through a communication on channel b . The weak transition (4.8) is an aggregation of 4 analogous transitions to the ones just presented, this time relating to a second input of value v_2 . This yields an internal output chain of length 2, *i.e.*, $c_1!(v_1, c_2) \parallel c_2!(v_2, c_3)$. Finally, (4.9) is an aggregation of 4 transitions relating to the consumption of the first item in the chain, $c_1!(v_1, c_2)$, the subsequent output of v_1 on channel `out`, and the unfolding and instantiation of the recursive process `Bck` with c_2 — see definition for `Bck`.

4.2. Costed Bisimulation. We define a cost-based preorder over systems as a *typed relation*, *cf.* Definition 4.2, ordering systems that exhibit the same external behaviour at a less-than-or-equal-to cost. We require the preorder to consider client C_1 as more efficient than C_0 *wrt.* an appropriate resource environment M and observers characterised by the type environment stated in (1.6) but also that, *wrt.* the same resource and observer environments, client C_3 of (1.5) is more efficient than C_1 . This latter ordering is harder to establish since client C_1 is at times *temporarily* more efficient than C_3 .

In order to handle this aspect we define our preorder as an *amortized bisimulation* [30]. Amortized bisimulation uses a *credit* n to compare a system $M \triangleright P$ with a less efficient system $N \triangleright Q$ while allowing $M \triangleright P$ to do a more expensive action than $N \triangleright Q$, as long as the credit can make up for the difference. Conversely, whenever $M \triangleright P$ does a cheaper action than $N \triangleright Q$, then the difference

gets *added* to the credit.¹⁰ Crucially, however, the amortisation credit is *never allowed* to become *negative* i.e., $n \in \mathbb{N}_{\text{AT}}$. In general, we refine Definition 4.2 to amortized typed relations with the following structure:

Definition 4.8 (Amortised Typed Relation). An amortized type-indexed relation \mathcal{R} relates systems under an observer characterized by a context Γ , with credit n ($n \in \mathbb{N}_{\text{AT}}$); we write

$$\Gamma \vDash M \triangleright P \mathcal{R}^n N \triangleright Q$$

if \mathcal{R}^n relates $\Gamma \triangleleft M \triangleright P$ and $\Gamma \triangleleft N \triangleright Q$, and both $\Gamma \triangleleft M \triangleright P$ and $\Gamma \triangleleft N \triangleright Q$ are configurations.

Definition 4.9 (Amortised Typed Bisimulation). An amortized type-indexed relation over processes \mathcal{R} is a bisimulation at Γ with credit n if, whenever $\Gamma \vDash (M \triangleright P) \mathcal{R}^n (N \triangleright Q)$,

- If $\Gamma \triangleleft M \triangleright P \xrightarrow{\mu}_k \Gamma' \triangleleft M' \triangleright P'$ then there exist N' and Q' such that $\Gamma \triangleleft N \triangleright Q \xrightarrow{\hat{\mu}}_l \Gamma' \triangleleft N' \triangleright Q'$ where $\Gamma' \vDash (M' \triangleright P') \mathcal{R}^{n+l-k} (N' \triangleright Q')$
- If $\Gamma \triangleleft N \triangleright Q \xrightarrow{\mu}_l \Gamma' \triangleleft N' \triangleright Q'$ then there exist M' and P' such that $\Gamma \triangleleft M \triangleright P \xrightarrow{\hat{\mu}}_k \Gamma' \triangleleft M' \triangleright P'$ where $\Gamma' \vDash (M' \triangleright P') \mathcal{R}^{n+l-k} (N' \triangleright Q')$

where $\hat{\mu}$ is the empty string if $\mu = \tau$ and μ otherwise.

Bisimilarity at Γ with credit n , denoted $\Gamma \vDash M \triangleright P \sqsubseteq_{\text{bis}}^n N \triangleright Q$, is the largest amortized typed bisimulation at Γ with credit n . We sometimes existentially quantify over the credit and write $\Gamma \vDash M \triangleright P \sqsubseteq_{\text{bis}} N \triangleright Q$. We write $\Gamma \vDash M \triangleright P \simeq_{\text{bis}} N \triangleright Q$ to denote the kernel of the preorder (i.e., whenever we have both $\Gamma \vDash M \triangleright P \sqsubseteq_{\text{bis}} N \triangleright Q$ and $\Gamma \vDash N \triangleright Q \sqsubseteq_{\text{bis}} M \triangleright P$), and write $\Gamma \vDash M \triangleright P \sqsubset_{\text{bis}} N \triangleright Q$ whenever $\Gamma \vDash M \triangleright P \sqsubseteq_{\text{bis}} N \triangleright Q$ but $\Gamma \vDash N \triangleright Q \not\sqsubseteq_{\text{bis}} M \triangleright P$.

Example 4.10 (Assessing Client Efficiency). For the (observer) type environment

$$\Gamma_1 \stackrel{\text{def}}{=} \text{srv}_1 : [[\mathbf{T}_1]^1]^\omega, \text{srv}_2 : [[\mathbf{T}_2]^1]^\omega, c : [\mathbf{T}_1, \mathbf{T}_2]^\omega \quad (4.10)$$

and clients C_0 and C_1 defined earlier in (1.2), we can show that $\Gamma_1 \vDash (M \triangleright C_1) \sqsubseteq_{\text{bis}} (M \triangleright C_0)$ by constructing the witness bisimulation (family of) relation(s) \mathcal{R} for $\Gamma_1 \vDash (M \triangleright C_1) \sqsubseteq_{\text{bis}}^0 (M \triangleright C_0)$ stated below:¹¹

$$\mathcal{R} \stackrel{\text{def}}{=} \left\{ \begin{array}{l} \left\langle \Gamma, n, M' \triangleright C_1, N' \triangleright C_0 \right\rangle \\ \left\langle \Gamma, n, M' \triangleright \text{alloc } x. \text{srv}_1 !x. x?y. \text{srv}_2 !x. x?z. c!(y, z). C_1, \right. \\ \quad \left. N' \triangleright \text{alloc } x_1. \text{alloc } x_2. \text{srv}_1 !x_1. x_1?y. \text{srv}_2 !x_2. x_2?z. c!(y, z). C_0 \right\rangle \\ \left\langle \Gamma, n, (M', d) \triangleright \text{srv}_1 !d. d?y. \text{srv}_2 !d. d?z. c!(y, z). C_1, \right. \\ \quad \left. (N', d') \triangleright \text{alloc } x_2. \text{srv}_1 !d'. d'?y. \text{srv}_2 !x_2. x_2?z. c!(y, z). C_0 \right\rangle \\ \left\langle \Gamma, n+1, (M', d) \triangleright \text{srv}_1 !d. d?y. \text{srv}_2 !d. d?z. c!(y, z). C_1, \right. \\ \quad \left. (N', d', d'') \triangleright \text{srv}_1 !d'. d'?y. \text{srv}_2 !d''. d''?z. c!(y, z). C_0 \right\rangle \\ \left\langle (\Gamma, d : [\mathbf{T}_1]^1), n+1, (M', d) \triangleright d?y. \text{srv}_2 !d. d?z. c!(y, z). C_1, \right. \\ \quad \left. (N', d, d'') \triangleright d?y. \text{srv}_2 !d''. d''?z. c!(y, z). C_0 \right\rangle \\ \left\langle \Gamma, n+1, (M', d) \triangleright \text{srv}_2 !d. d?z. c!(v, z). C_1, \right. \\ \quad \left. (N', d, d'') \triangleright \text{srv}_2 !d''. d''?z. c!(v, z). C_0 \right\rangle \\ \left\langle (\Gamma, d : [\mathbf{T}_2]^1), n+1, (M', d) \triangleright d?z. c!(v, z). C_1, (N', d', d) \triangleright d?z. c!(v, z). C_0 \right\rangle \\ \left\langle \Gamma, n+1, (M', d) \triangleright c!(v, v'). C_1, (N', d', d) \triangleright c!(v, v'). C_0 \right\rangle \end{array} \right\} \begin{array}{l} n \geq 0 \\ d \notin \text{dom}(\Gamma) \\ d' \notin \text{dom}(\Gamma) \\ d'' \notin \text{dom}(\Gamma) \\ M' \subseteq N' \\ \text{dom}(\Gamma) \subseteq M' \end{array}$$

¹⁰Stated otherwise, $M \triangleright P$ can do a more expensive action than $N \triangleright Q$ now, as long as it makes up for it later.

¹¹In families of relations ranging over systems indexed by type environments and amortisation credits, such as \mathcal{R} , we represent $\Gamma \vDash (M \triangleright P) \sqsubseteq_{\text{bis}}^n (\Delta \triangleright Q)$ as the quadruple $(\Gamma, n, (M \triangleright P), (\Delta \triangleright Q))$.

It is not hard to see that \mathcal{R} contains the quadruple $\langle \Gamma_1, 0, M \triangleright C_1, M \triangleright C_0 \rangle$. One can also show that it is closed *wrt.* the transfer property of Definition 4.9. The key moves are:

- a single channel allocation by C_1 is matched by two channel allocations by C_0 — from the second up to the fourth quadruple in the definition of \mathcal{R} . Since channel allocations carry a positive cost, the amortisation credit increases from n to $n+2-1$, *i.e.*, $n+1$, but this still yields a quadruple that is in the relation. One thing to note is that the first channel allocated by both systems is allowed to be different, *e.g.*, d and d' , as long as it is not allocated already.
- Even though the internal channels allocated may be different, rule LREN allows us to rename the *resp.* names of the allocated channels (not known to the observer) so as match the channels communicated on srv_1 by the other system (fourth and fifth quadruples). Since these channels are not known to the observer, *i.e.*, they are not in $\text{dom}(\Gamma)$, they all amount to *fresh* names, akin to scope extrusion [35, 19].
- Communicating on the previously communicated channel on srv_1 consumes all of the observer's permissions for that channel (fifth quadruple), which allows rule LREN to be applied again so as to match the channels communicated on srv_2 (sixth quadruple).

We cannot however prove that $\Gamma_1 \vDash (M \triangleright C_0) \stackrel{n}{\approx}_{\text{bis}} (M \triangleright C_1)$ for any n because we would need an *infinite* amortisation credit to account for additional cost incurred by C_0 when it performs the channel extra allocation at every iteration; recall that this credit cannot become negative, and thus no finite credit is large enough to cater for all the additional cost incurred by C_0 over sufficiently large transition sequences.

Similarly, from (1.2), we can show that $\Gamma_1 \vDash (M \triangleright C_2) \sqsubset_{\text{bis}} (M \triangleright C_1)$ but also, from (1.5), that $\Gamma_1 \vDash (M \triangleright C_3) \sqsubset_{\text{bis}} (M \triangleright C_1)$. In particular, we can show $\Gamma_1 \vDash (M \triangleright C_3) \stackrel{1}{\approx}_{\text{bis}} (M \triangleright C_1)$ even though $M \triangleright C_1$ is temporarily more efficient than $M \triangleright C_3$, *i.e.*, during the course of the first iteration. Our framework handles this through the use of the amortisation credit whereby, in this case, it suffices to use a credit of value 1 and show $\Gamma_1 \vDash (M \triangleright C_3) \stackrel{1}{\approx}_{\text{bis}} (M \triangleright C_1)$; we leave the details to the interested reader. Using an amortisation credit of 1 we can also show $\Gamma_1 \vDash (M \triangleright C_3) \stackrel{1}{\approx}_{\text{bis}} (M \triangleright C_2)$ through the bisimulation family-of-relations \mathcal{R}' below — it is easy to check that it observes the transfer property of Definition 4.9; by constructing a similar relation, one can also show that $\Gamma_1 \vDash (M \triangleright C_2) \stackrel{0}{\approx}_{\text{bis}} (M \triangleright C_3)$ which implies that $\Gamma_1 \vDash (M \triangleright C_2) \simeq_{\text{bis}} (M \triangleright C_3)$. We just note that in \mathcal{R}' , the amortisation credit n can be capped $0 \leq n \leq 1$ and revisit this point again in Section 4.4.

$$\mathcal{R}' \stackrel{\text{def}}{=} \left\{ \begin{array}{l} \langle \Gamma, 1, M \triangleright C_3, M \triangleright C_2 \rangle \\ \left\langle \Gamma, 1, M \triangleright \left(\begin{array}{l} \text{alloc } x_1. \text{alloc } x_2. \text{srv}_1!x_1. x_1?y. \\ \text{srv}_2!x_2. x_2?z. \text{free } x_1. \text{free } x_2. \text{c}!(y, z). C_3 \end{array} \right), \right. \\ \quad \left. M \triangleright \text{alloc } x. \text{srv}_1!x. x?y. \text{srv}_2!x. x?z. \text{free } x. \text{c}!(y, z). C_2 \right\rangle \\ \left\langle \Gamma, 1, (M, d) \triangleright \left(\begin{array}{l} \text{alloc } x_2. \text{srv}_1!d. d?y. \\ \text{srv}_2!x_2. x_2?z. \text{free } d. \text{free } x_2. \text{c}!(y, z). C_3 \end{array} \right), \right. \\ \quad \left. (M, d') \triangleright \text{srv}_1!d'. d'?y. \text{srv}_2!d'. d'?z. \text{free } d'. \text{c}!(y, z). C_2 \right\rangle \quad d' \notin \text{dom}(\Gamma) \\ \left\langle \Gamma, 0, (M, d, d'') \triangleright \left(\begin{array}{l} \text{srv}_1!d. d?y. \text{srv}_2!d''. d''?z. \\ \text{free } d. \text{free } d''. \text{c}!(y, z). C_3 \end{array} \right), \right. \\ \quad \left. (M, d') \triangleright \text{srv}_1!d'. d'?y. \text{srv}_2!d''. d''?z. \text{free } d'. \text{c}!(y, z). C_2 \right\rangle \quad d' \notin \text{dom}(\Gamma) \\ \left\langle (\Gamma, d : [\mathbf{T}_1]^1), 0, (M, d, d'') \triangleright \left(\begin{array}{l} d?y. \text{srv}_2!d''. d''?z. \\ \text{free } d. \text{free } d''. \text{c}!(y, z). C_3 \end{array} \right), \right. \\ \quad \left. (M, d) \triangleright d?y. \text{srv}_2!d. d?z. \text{free } d. \text{c}!(y, z). C_2 \right\rangle \quad d'' \notin \text{dom}(\Gamma) \\ \left\langle \Gamma, 0, (M, d, d'') \triangleright \text{srv}_2!d''. d''?z. \text{free } d. \text{free } d''. \text{c}!(v, z). C_3, \right. \\ \quad \left. (M, d) \triangleright \text{srv}_2!d. d?z. \text{free } d. \text{c}!(v, z). C_2 \right\rangle \\ \left\langle (\Gamma, d' : [\mathbf{T}_2]^1), 0, (M, d, d') \triangleright d'?z. \text{free } d. \text{free } d'. \text{c}!(v, z). C_3, \right. \\ \quad \left. (M, d') \triangleright d'?z. \text{free } d'. \text{c}!(v, z). C_2 \right\rangle \quad \text{dom}(\Gamma) \subseteq M \\ \left\langle \Gamma, 0, (M, d, d') \triangleright \text{free } d. \text{free } d'. \text{c}!(v, z). C_3, \right. \\ \quad \left. (M, d') \triangleright \text{free } d'. \text{c}!(v, v'). C_2 \right\rangle \\ \left\langle \Gamma, 0, (M, d') \triangleright \text{free } d'. \text{c}!(v, z). C_3, M \triangleright \text{c}!(v, v'). C_2 \right\rangle \\ \left\langle \Gamma, 1, M \triangleright \text{c}!(v, z). C_3, M \triangleright \text{c}!(v, v'). C_2 \right\rangle \end{array} \right\}$$

□

4.3. Alternatives. The cost model we adhere to in Section 4 is not the only plausible one, but is intended to follow that described by costed reductions of Section 2. There may however be other valid alternatives, some of which can be easily accommodated through minor tweaking to our existing framework.

For instance, an alternative cost model may focus on assessing the runtime execution of programs, whereby operations that access memory such as $\text{alloc } x.P$ and $\text{free } c.P$ have a runtime cost that far exceeds that of other operations. We can model this by considering an LTS that assigns a cost of 1 to both of these operations, which can be attained as a derived LTS from our existing LTS of Section 4.1 through the rule

$$\frac{\Gamma \triangleleft M \triangleright P \xrightarrow{\mu}_k \Gamma' \triangleleft M' \triangleright P'}{\Gamma \triangleleft M \triangleright P \xrightarrow{\mu}_{|k|} \Gamma' \triangleleft M' \triangleright P'} \text{LDER1}$$

where $|k|$ returns the absolute value of an integer. Definition 4.9 extends in straightforward fashion to work with the derived costed LTS $\xrightarrow{\mu}_{|k|}$. This new preorder would allow us to conclude $\Gamma_1 \vDash (M \triangleright C_1) \sqsubseteq_{\text{bis}} (M \triangleright C_2)$ because, according to the new cost model, for every server-interaction iteration, client C_1 uses less expensive memory operations than C_2 .

Another cost model may require us to refine our existing preorder. For instance, consider another client C_4 , defined below, that creates a single channel and keeps on reusing it for all iterations:

$$C_4 \triangleq \text{alloc } x. \text{rec } w. \text{srv}_1!x. x?y. \text{srv}_2!x. x?z. \text{ret}!(y, z). w$$

At present, we are able to equate this client with C_2 and C_3 from (1.2) and (1.5) *resp.*, on the basis that neither client carries any memory leaks.

$$\Gamma_1 \vDash (M \triangleright C_4) \simeq_{\text{bis}} (M \triangleright C_3) \simeq_{\text{bis}} (M \triangleright C_2)$$

However, we may want a finer preorder where C_4 is considered to be (strictly) more efficient than C_2 , which is in turn more efficient than C_3 . The underlying reasoning for this would be that C_4 uses the least amount of expensive operations; by contrast C_2 keeps on allocating (and deallocating) new channels for each iteration, and C_3 allocates (and deallocates) two new channels for every iteration. We can characterise this preorder as follows. First we generate the derived costed LTS using the rule LDER2 below — $[k]$ maps all negative integers to 0, leaving positive integers unaltered.

$$\frac{\Gamma \triangleleft M \triangleright P \xrightarrow{\mu}_k \Gamma' \triangleleft M' \triangleright P'}{\Gamma \triangleleft M \triangleright P \xrightarrow{\mu}_{[k]} \Gamma' \triangleleft M' \triangleright P'} \text{LDER2}$$

Then, after adapting Definition 4.9 to this derived LTS, denoting such a bisimulation relation as $\sqsubseteq_{\text{bis2}}$, we can define the refined preorder, denoted as $\sqsubseteq_{\text{bis3}}$, as follows:

$$\Gamma \vDash M \triangleright P \sqsubseteq_{\text{bis3}} N \triangleright Q \stackrel{\text{def}}{=} \begin{cases} \Gamma \vDash M \triangleright P \sqsubseteq_{\text{bis}} N \triangleright Q \text{ and} \\ \Gamma \vDash N \triangleright Q \sqsubseteq_{\text{bis}} M \triangleright P \text{ implies } \Gamma \vDash M \triangleright P \sqsubseteq_{\text{bis2}} N \triangleright Q \end{cases}$$

The new refined preorder $\sqsubseteq_{\text{bis3}}$ above requires that $M \triangleright P$ is at least as efficient as $N \triangleright Q$ (possibly more) when it comes to memory leaks, *i.e.*, \sqsubseteq_{bis} , and moreover, whenever they are equally efficient *wrt.* these leaks, $M \triangleright P$ must also be as efficient (possibly more) *wrt.* memory allocations, *i.e.*, $\sqsubseteq_{\text{bis2}}$.

4.4. Properties of \sqsubseteq_{bis} . We show that our bisimulation relation of Definition 4.9 observes a number of properties that are useful when reasoning about resource efficiency; see Example 4.24 below. Lemmas 4.11 and 4.12 prove that the relation is in fact a preorder, whereas Lemma 4.14 outlines conditions where symmetry can be recovered. Finally, Theorem 4.23 shows that this preorder is preserved under (valid) context; this is the main result of the section.

First off, we show that \sqsubseteq_{bis} is a preorder following Lemma 4.11 (where σ_Γ would be the identity) and Lemma 4.12.

Lemma 4.11 (Reflexivity upto Renaming). *Whenever the triple $\Gamma \triangleleft M \triangleright P$ is a configuration, then $\Gamma \vDash (M \triangleright P)\sigma_\Gamma \simeq_{\text{bis}} M \triangleright P$*

Proof. By coinduction, by showing that the family of relations

$$\{(\Gamma, 0, (M \triangleright P)\sigma_\Gamma, M \triangleright P) \mid \Gamma \triangleleft M \triangleright P \text{ is a configuration}\}$$

is a bisimulation. □

Lemma 4.12 (Transitivity). *Whenever $\Gamma \vDash M \triangleright P \sqsubseteq_{\text{bis}} M' \triangleright P'$ and $\Gamma \vDash M' \triangleright P' \sqsubseteq_{\text{bis}} M'' \triangleright P''$ then $\Gamma \vDash M \triangleright P \sqsubseteq_{\text{bis}} M'' \triangleright P''$*

Proof. $\Gamma \vDash M \triangleright P \sqsubseteq_{\text{bis}} M' \triangleright P'$ implies that there exists some $n \geq 0$ and corresponding bisimulation relation justifying $\Gamma \vDash M \triangleright P \sqsubseteq_{\text{bis}}^n M' \triangleright P'$. The same applies for $\Gamma \vDash M' \triangleright P' \sqsubseteq_{\text{bis}} M'' \triangleright P''$ and some $m \geq 0$. From these two relations, one can construct a corresponding bisimulation justifying $\Gamma \vDash M \triangleright P \sqsubseteq_{\text{bis}}^{n+m} M'' \triangleright P''$. □

Corollary 4.13 (Preorder). *\sqsubseteq_{bis} is a preorder.*

Proof. Follows from Lemma 4.11 (for the special case where σ_Γ is the identity) and Lemma 4.12. □

We can define a restricted form of amortised typed bisimulation, in analogous fashion to Definition 4.9, whereby the credit is *capped at some upper bound*, i.e., some natural number m . We refer to such relations as *Bounded Amortised Typed-Bisimulations* and write

$$\Gamma \vDash^m M \triangleright P \stackrel{n}{\sim}_{\text{bis}} N \triangleright Q$$

to denote that $\Gamma \triangleleft M \triangleright P$ and $\Gamma \triangleleft N \triangleright Q$ are related by some amortised typed-indexed bisimulation at index Γ and credit n , and where every credit in this relation is less than or equal to m ; whenever the precise credit n is not important we elide it and simply write $\Gamma \vDash^m M \triangleright P \stackrel{}{\sim}_{\text{bis}} N \triangleright Q$. We can show that bounded amortised typed-bisimulations are symmetric.

Lemma 4.14 (Symmetry). $\Gamma \vDash^m M \triangleright P \stackrel{}{\sim}_{\text{bis}} N \triangleright Q$ implies $\Gamma \vDash^m N \triangleright Q \stackrel{}{\sim}_{\text{bis}} M \triangleright P$

Proof. If \mathcal{R} is the bounded amortised typed relation justifying $\Gamma \vDash^m M \triangleright P \stackrel{}{\sim}_{\text{bis}} N \triangleright Q$, we define the amortised typed relation

$$\mathcal{R}_{\text{sym}} = \{ \langle \Gamma, (m-n), N \triangleright Q, M \triangleright P \rangle \mid \langle \Gamma, n, M \triangleright P, N \triangleright Q \rangle \in \mathcal{R} \}$$

and show that it is a bounded amortised typed bisimulation as well. Consider an arbitrary pair of configurations $\Gamma \vDash N \triangleright Q \mathcal{R}_{\text{sym}}^{m-n} M \triangleright P$:

- Assume $\Gamma \triangleleft N \triangleright Q \xrightarrow{l} \Gamma' \triangleleft N' \triangleright Q'$. From the definition of \mathcal{R}_{sym} , it must be the case that $\langle \Gamma, n, M \triangleright P, N \triangleright Q \rangle \in \mathcal{R}$. Since \mathcal{R} is a bounded amortised typed bisimulation, we know that $\Gamma \triangleleft M \triangleright P \xrightarrow{k} \Gamma' \triangleleft M' \triangleright P'$ where $\langle \Gamma', n+k-l, M' \triangleright P', N' \triangleright Q' \rangle \in \mathcal{R}$. We however need to show that $\langle \Gamma', ((m-n)+k-l), N' \triangleright Q', M' \triangleright P' \rangle \in \mathcal{R}_{\text{sym}}$, which follows from the definition of \mathcal{R}_{sym} and the fact that $(m - (n + l - k)) = (m - n) + k - l$.

What is left to show is that \mathcal{R}_{sym} is an amortised typed bisimulation bounded by m , i.e., we need to show that $0 \leq (m-n)+k-l \leq m$. Since \mathcal{R} is an m -bounded amortised typed bisimulation, we know that $0 \leq (n+l-k) \leq m$ from which we can drive $-m \leq -(n+l-k) \leq 0$ and, by adding m throughout we obtain $0 \leq (m - (n + l - k)) = (m - n) + k - l \leq m$ as required.

- The dual case for $\Gamma \triangleleft M \triangleright P \xrightarrow{l} \Gamma' \triangleleft M' \triangleright P'$ is analogous. □

Contextuality is an important property for any behavioural relation. In our case, this means that two systems $M \triangleright P$ and $N \triangleright Q$ related by $\stackrel{}{\sim}_{\text{bis}}$ under Γ , remain related when extended with an additional process, R , whenever this process runs safely over the respective resource environments M and N , and observes the type restrictions and guarantees assumed by Γ (and dually, those of the respective existentially-quantified type environments for $M \triangleright P$ and $N \triangleright Q$). Following Definition 4.1, for these conditions to hold, contextuality requires R to typecheck wrt. a sub-environment of Γ , say Γ_1 where $\Gamma = \Gamma_1, \Gamma_2$, and correspondingly strengthens the relation of $M \triangleright P \parallel R$ and $N \triangleright Q \parallel R$ in $\stackrel{}{\sim}_{\text{bis}}$ under the remaining sub-environment, Γ_2 . Stated otherwise, contextuality requires the transfer of the respective permissions associated with the observer sub-process R from the observer environment Γ ; this is crucial in order to preserve consistency, thus safety, in the respective configurations. The formulation of Theorem 4.23, proving contextuality for $\stackrel{}{\sim}_{\text{bis}}$, follows this reasoning. It relies on a list of lemmas outlined below.

Lemma 4.15 (Weakening). *If $\Gamma \triangleleft M \triangleright P \xrightarrow{k} \Gamma' \triangleleft M' \triangleright P'$ then $(\Gamma, \Delta) \triangleleft M \triangleright P \xrightarrow{k} (\Gamma', \Delta) \triangleleft M' \triangleright P'$. (These may or may not be configurations.)*

Proof. By rule induction on $\Gamma \triangleleft M \triangleright P \xrightarrow{k} \Gamma' \triangleleft M' \triangleright P'$. Note that, in the case of `alloc`, the action can still be performed. □

Lemma 4.16 (Strengthening). *If $(\Gamma, \Delta) \triangleleft M \triangleright P \xrightarrow{\mu}_k (\Gamma', \Delta) \triangleleft M' \triangleright P'$ then $\Gamma \triangleleft M \triangleright P \xrightarrow{\mu}_k \Gamma' \triangleleft M' \triangleright P'$.*

Proof. By rule induction on $\Gamma, \Delta \triangleleft P \xrightarrow{\mu}_k \Gamma', \Delta \triangleleft P'$. Note that strengthening is restricted to the part of the environment that remains unchanged (Δ is the same on the left and right hand side) — otherwise the property does not hold for actions $c!\vec{d}$ and $c?\vec{d}$. \square

Lemma 4.17. *If Γ, Δ is consistent and $\Delta < \Delta'$ then Γ, Δ' is consistent and $\Gamma, \Delta < \Gamma, \Delta'$*

Proof. As in [12]. \square

Lemma 4.18 (Typing Preserved by \equiv). *$\Gamma \vdash P$ and $P \equiv Q$ implies $\Gamma \vdash Q$*

Proof. As in [12]. \square

Lemma 4.19 (Environment Structural Manipulation Preserves Bisimulation).

$$\Gamma \vDash S \stackrel{n}{\sim}_{bis} T \text{ and } \Gamma < \Gamma' \text{ implies } \Gamma' \vDash S \stackrel{n}{\sim}_{bis} T$$

Proof. By coinduction. We define the quaternary relation

$$\{ \langle \Gamma', n, S, T \rangle \mid \Gamma \vDash S \stackrel{n}{\sim}_{bis} T \text{ and } \Gamma < \Gamma' \}$$

and show that it observes the transfer property of Definition 4.9. \square

Lemma 4.20 (Bisimulation and Structural Equivalence).

$$P \equiv Q \text{ and } \Gamma \triangleleft M \triangleright P \xrightarrow{\mu}_k \Delta \triangleleft M' \triangleright P' \text{ implies } \Gamma \triangleleft M \triangleright Q \xrightarrow{\mu}_k \Delta \triangleright M' \triangleright Q' \text{ and } P' \equiv Q'$$

Proof. By rule induction on $P \equiv Q$ and then a case analysis of the rules permitting $\Gamma \triangleleft M \triangleright P \xrightarrow{\mu}_k \Delta \triangleleft M' \triangleright P'$. \square

Corollary 4.21 (Structural Equivalence and Bisimilarity). *$P \equiv Q$ implies $\Gamma \vDash M \triangleright P \stackrel{n}{\sim}_{bis} M \triangleright Q$ for arbitrary n and Γ where $\Gamma \triangleleft M \triangleright P$ and $\Gamma \triangleleft M \triangleright Q$ are configurations.*

Proof. By coinduction and Lemma 4.20. \square

Lemma 4.22 (Renaming). *If $\Gamma, \Delta \vDash (M \triangleright P) \stackrel{n}{\sim}_{bis} (N \triangleright Q)$ then $\Gamma, (\Delta \sigma_\Gamma) \vDash (M \triangleright P) \sigma_\Gamma \stackrel{n}{\sim}_{bis} (N \triangleright Q) \sigma_\Gamma$*

Proof. By coinduction. \square

Theorem 4.23 (Contextuality). *If $\Gamma, \Delta \vDash (M \triangleright P) \stackrel{n}{\sim}_{bis} (N \triangleright Q)$ and $\Delta \vdash R$ then*

$$\Gamma \vDash (M \triangleright P \parallel R) \stackrel{n}{\sim}_{bis} (N \triangleright Q \parallel R) \quad \text{and} \quad \Gamma \vDash (M \triangleright R \parallel P) \stackrel{n}{\sim}_{bis} (N \triangleright R \parallel Q)$$

Proof. We define the family of relations $\mathcal{R}^{\Gamma, n}$ to be the least one satisfying the rules

$$\frac{\Gamma \vDash (M \triangleright P) \stackrel{n}{\sim}_{bis} (N \triangleright Q)}{\Gamma \vDash (M \triangleright P) \mathcal{R}^{\Gamma, n} (N \triangleright Q)} \quad \frac{\Gamma, \Delta \vDash (M \triangleright P) \mathcal{R}^{\Gamma, n} (N \triangleright Q) \quad \Delta \vdash R}{\Gamma \vDash (M \triangleright P \parallel R) \mathcal{R}^{\Gamma, n} (N \triangleright Q \parallel R)} \quad \frac{\Gamma, \Delta \vDash (M \triangleright P) \mathcal{R}^{\Gamma, n} (N \triangleright Q) \quad \Delta \vdash R}{\Gamma \vDash (M \triangleright R \parallel P) \mathcal{R}^{\Gamma, n} (N \triangleright R \parallel Q)}$$

and then show that $\mathcal{R}^{\Gamma, n}$ is a costed typed bisimulation at Γ and n (up to \equiv). Note that the premise of the first rule implies that both $\Gamma, \Delta \triangleleft M \triangleright P$ and $\Gamma, \Delta \triangleleft N \triangleright Q$ are configurations. We consider only the transitions of the left hand configurations for second case of the relation; the first is trivial and the third is analogous to the second. Although the relation is not symmetric, the transition of the right hand configurations are analogous to those of the left hand configurations. There are three cases to consider.

(1) Case the action was instigated by P , i.e., we have:

$$\frac{\frac{\Gamma \triangleleft (M \triangleright P)\sigma_\Gamma \xrightarrow{\mu}_l \Gamma' \triangleleft M' \triangleright P'}{\Gamma \triangleleft (M \triangleright P \parallel R)\sigma_\Gamma \xrightarrow{\mu}_l \Gamma' \triangleleft M' \triangleright P' \parallel R\sigma_\Gamma} \text{LPAR-L}}{\Gamma \triangleleft M \triangleright P \parallel R \xrightarrow{\mu}_l \Gamma' \triangleleft M' \triangleright P' \parallel R\sigma_\Gamma} \text{LREN} \quad (4.11)$$

By Lemma 4.15 (Weakening), LREN and (4.11) we obtain

$$\Gamma, (\Delta\sigma_\Gamma) \triangleleft (M \triangleright P)\sigma_\Gamma \xrightarrow{\mu}_l \Gamma', (\Delta\sigma_\Gamma) \triangleleft M' \triangleright P' \quad (4.12)$$

Lemma 4.22 can be extended to \mathcal{R}^n in straightforward fashion, and from the case assumption $\Gamma, \Delta \vDash M \triangleright P \mathcal{R}^n N \triangleright Q$ (defining $\mathcal{R}^{\Gamma, n}$) and the extension of Lemma 4.22 to \mathcal{R}^n we obtain:

$$\Gamma, (\Delta\sigma_\Gamma) \vDash (M \triangleright P)\sigma_\Gamma \mathcal{R}^n (N \triangleright Q)\sigma_\Gamma \quad (4.13)$$

Hence by (4.13), (4.12) and I.H. there exists a $N' \triangleright Q'$ such that

$$\Gamma, (\Delta\sigma_\Gamma) \triangleleft (N \triangleright Q)\sigma_\Gamma \xrightarrow{\hat{\mu}}_k \Gamma', (\Delta\sigma_\Gamma) \triangleleft N' \triangleright Q' \quad (4.14)$$

$$\text{where } \Gamma', (\Delta\sigma_\Gamma) \vDash (M' \triangleright P') \mathcal{R}^{n+k-l} (N' \triangleright Q') \quad (4.15)$$

By (4.14) and LREN, where $\Gamma_1 = \Gamma, (\Delta\sigma_\Gamma)$, we obtain

$$\Gamma, (\Delta\sigma_\Gamma) \triangleleft ((N \triangleright Q)\sigma_\Gamma)\sigma'_{\Gamma_1} \xrightarrow{(\tau_{k_1}^*)}_{k_1} \xrightarrow{\hat{\mu}}_{k_2} \xrightarrow{(\tau_{k_3}^*)}_{k_3} \Gamma', (\Delta\sigma_\Gamma) \triangleleft N' \triangleright Q' \quad (4.16)$$

where $k = k_1 + k_2 + k_3$. By LPAR Lemma 4.16 (Strengthening) and (4.16) we deduce

$$\Gamma \triangleleft ((N \triangleright Q)\sigma_\Gamma)\sigma'_{\Gamma_1} \parallel R\sigma_\Gamma \xrightarrow{(\tau_{k_1}^*)}_{k_1} \xrightarrow{\hat{\mu}}_{k_2} \xrightarrow{(\tau_{k_3}^*)}_{k_3} \Gamma' \triangleleft N' \triangleright Q' \parallel R\sigma_\Gamma \quad (4.17)$$

From $\Delta \vdash R$ we know

$$\Delta\sigma_\Gamma \vdash R\sigma_\Gamma \quad (4.18)$$

and, from $\Gamma_1 = \Gamma, (\Delta\sigma_\Gamma)$ and Definition 4.3 (Renaming Modulo Environments), we know that $(R\sigma_\Gamma)\sigma'_{\Gamma_1} = R\sigma_\Gamma$ since the renaming does not modify any of the names in the domain of Γ_1 , hence of $\Delta\sigma_\Gamma$. Also, from Definition 4.3, σ'_{Γ_1} is also a substitution modulo Γ and can therefore refer to it as σ'_Γ , thereby rewriting (4.17) as

$$\Gamma \triangleleft (N \triangleright Q \parallel R)\sigma_\Gamma\sigma'_\Gamma \xrightarrow{(\tau_{k_1}^*)}_{k_1} \xrightarrow{\hat{\mu}}_{k_2} \xrightarrow{(\tau_{k_3}^*)}_{k_3} \Gamma' \triangleleft N' \triangleright Q' \parallel R\sigma_\Gamma \quad (4.19)$$

From (4.19) and LREN we thus obtain

$$\Gamma \triangleleft N \triangleright Q \parallel R \xrightarrow{\hat{\mu}}_k \Gamma' \triangleleft N' \triangleright (Q' \parallel R\sigma_\Gamma)$$

This is our matching move since and by (4.15), (4.18) and the definition of \mathcal{R} we obtain $\Gamma' \vDash (M' \triangleright P' \parallel R\sigma_\Gamma) \mathcal{R}^{n+l-k} (N' \triangleright Q' \parallel R\sigma_\Gamma)$.

(2) Case the action was instigated by R , i.e., we have:

$$\frac{\frac{\Gamma \triangleleft (M \triangleright R)\sigma_\Gamma \xrightarrow{\mu}_l \Gamma' \triangleleft M' \triangleright R'}{\Gamma \triangleleft (M \triangleright P \parallel R)\sigma_\Gamma \xrightarrow{\mu}_l \Gamma' \triangleleft M' \triangleright P \parallel R'} \text{LPAR-R}}{\Gamma \triangleleft M \triangleright P \parallel R \xrightarrow{\mu}_l \Gamma' \triangleleft M' \triangleright P \parallel R'} \text{LREN} \quad (4.20)$$

The proof proceeds by case analysis of μ whereby the most interesting cases are when $l = +1$ or $l = -1$. We here show the case for when $l = -1$ (the other case is analogous). By Lemma 4.4 we know that either $\mu = \text{free } c$ and

$$M\sigma_\Gamma = M', c \quad R' \equiv R\sigma_\Gamma \quad \Gamma = \Gamma', c : [\vec{\mathbf{T}}]^\bullet$$

or else that $\mu = \tau$ and

$$M\sigma_\Gamma = M', c \quad (4.21)$$

$$R\sigma_\Gamma \equiv \text{free } c.R_1 \parallel R_2 \text{ and } R' \equiv R_1 \parallel R_2 \quad (4.22)$$

$$\Gamma = \Gamma' \quad (4.23)$$

$$\Delta\sigma_\Gamma < \Delta', c: [\vec{\mathbf{T}}]^\bullet \text{ and } \Delta' \vdash R'. \quad (4.24)$$

We here focus on the latter case, *i.e.*, when $\mu = \tau$. The main complication in finding a matching move for this subcase is that of inferring a pair of resultant systems (one of which is $\Gamma \triangleleft M' \triangleright P \parallel R'$) that are related by \mathcal{R} by using the inductive nature of the relation definition. To be able to do so, we need to mimic the effect of R 's deallocation transition on M in the corresponding system $N \triangleright Q$; we do this with the help of an appropriate external deallocation transition $\text{free } c$.

By the extension of Lemma 4.22 to \mathcal{R}^n we know $\Gamma, \Delta\sigma_\Gamma \vDash (M \triangleright P)\sigma_\Gamma \mathcal{R}^n (N \triangleright Q)\sigma_\Gamma$, and by (4.24) and a straightforward extension of Lemma 4.19 to \mathcal{R} we obtain

$$\Gamma, \Delta', c: [\vec{\mathbf{T}}]^\bullet \vDash (M \triangleright P)\sigma_\Gamma \mathcal{R}^n (N \triangleright Q)\sigma_\Gamma \quad (4.25)$$

and by (4.21) and LFREEE we deduce

$$\Gamma, \Delta', c: [\vec{\mathbf{T}}]^\bullet \triangleleft (M \triangleright P)\sigma_\Gamma \xrightarrow{\text{free } c}_{-1} \Gamma, \Delta' \triangleleft (M' \triangleright P)\sigma_\Gamma$$

and by (4.25) and I.H. there exists a matching move

$$\Gamma, \Delta', c: [\vec{\mathbf{T}}]^\bullet \triangleleft (N \triangleright Q)\sigma_\Gamma \xrightarrow{\text{free } c}_k \Gamma, \Delta' \triangleleft N' \triangleright Q' \quad (4.26)$$

$$\text{and } \Gamma, \Delta' \vDash M' \triangleright P' \mathcal{R}^{n+k-(-1)} N' \triangleright Q' \quad (4.27)$$

By (4.26) and LREN , for $k = k_1 - 1 + k_2$, we know

$$\Gamma, \Delta', c: [\vec{\mathbf{T}}]^\bullet \triangleleft ((N \triangleright Q)\sigma_\Gamma)\sigma'_{\Gamma_2} \xrightarrow{\tau}_{k_1}^* \Gamma, \Delta', c: [\vec{\mathbf{T}}]^\bullet \triangleleft N'' \triangleright Q'' \quad (4.28)$$

$$\text{where } \Gamma_2 = \Gamma, \Delta', c: [\vec{\mathbf{T}}]^\bullet \text{ (used in } \sigma'_{\Gamma_2} \text{ above)} \quad (4.29)$$

$$\Gamma, \Delta', c: [\vec{\mathbf{T}}]^\bullet \triangleleft N'' \triangleright Q'' \xrightarrow{\text{free } c}_{-1} \Gamma, \Delta' \triangleleft N''' \triangleright Q'' \quad (4.30)$$

$$\Gamma, \Delta' \triangleleft N''' \triangleright Q'' \xrightarrow{\tau}_{k_2}^* \Gamma, \Delta' \triangleleft N' \triangleright Q' \quad (4.31)$$

From (4.28), (4.31), LPAR-L and Lemma 4.16 (Strengthening) we obtain:

$$\Gamma \triangleleft ((N \triangleright Q)\sigma_\Gamma)\sigma'_{\Gamma_2} \parallel R\sigma_\Gamma \xrightarrow{\tau}_{k_1}^* \Gamma \triangleleft N'' \triangleright Q'' \parallel (R\sigma_\Gamma) \quad (4.32)$$

$$\Gamma \triangleleft N''' \triangleright Q'' \parallel R' \xrightarrow{\tau}_{k_2}^* \Gamma \triangleleft N' \triangleright Q' \parallel R' \quad (4.33)$$

Also, from (4.30) and Lemma 4.4 (Transition and Structure) we deduce that $N'' = N'''$, c and thus, from (4.22), LFREE , LPAR-R we obtain:

$$\Gamma \triangleleft N'' \triangleright Q'' \parallel R\sigma_\Gamma \xrightarrow{\tau}_{-1} \Gamma \triangleleft N''' \triangleright Q'' \parallel R' \quad (4.34)$$

By (4.24) and (4.29), we know that we can find an alternative renaming function σ''_{Γ_3} , where $\Gamma_3 = \Gamma, (\Delta\sigma_\Gamma)$, in a way that, from (4.32), we can obtain

$$\Gamma \triangleleft ((N \triangleright Q)\sigma_\Gamma)\sigma''_{\Gamma_3} \parallel R\sigma_\Gamma \xrightarrow{\tau}_{k_1}^* \Gamma \triangleleft N'' \triangleright Q'' \parallel (R\sigma_\Gamma) \quad (4.35)$$

Now, by $\Delta \vdash R$ we know $\Delta\sigma_\Gamma \vdash R\sigma_\Gamma$ and subsequently, by Definition 4.3 and (4.29) we know $(R\sigma_\Gamma)\sigma''_{\Gamma_3} = R\sigma_\Gamma$. Thus, we can rewrite $((N \triangleright Q)\sigma_\Gamma)\sigma''_{\Gamma_3} \parallel R\sigma_\Gamma$ in (4.35) as $((N \triangleright Q \parallel R)\sigma_\Gamma)\sigma''_{\Gamma_3}$. Merging (4.35), (4.34) and (4.33) we obtain:

$$\Gamma \triangleleft ((N \triangleright Q \parallel R)\sigma_\Gamma)\sigma''_{\Gamma_3} \xrightarrow{\tau}_{k_1}^* \xrightarrow{\tau}_{-1} \xrightarrow{\tau}_{k_2}^* \Gamma \triangleleft N' \triangleright Q' \parallel R'$$

By Definition 4.3 we know that σ''_{Γ_3} can be rewritten as σ''_{Γ} and thus by LREN we obtain the matching move

$$\Gamma \triangleleft N \triangleright Q \parallel R \xrightarrow{\tau}_k \Gamma \triangleleft N' \triangleright Q' \parallel R'$$

because by (4.27), (4.24) and the definition of \mathcal{R} we know that

$$\Gamma \vDash M' \triangleright P' \parallel R' \mathcal{R}^{n+k-(-1)} N' \triangleright Q' \parallel R'.$$

(3) Case the action resulted from an interaction between P and R , i.e., we have:

$$\frac{\frac{\Gamma_1 \triangleleft (M \triangleright P)\sigma_{\Gamma} \xrightarrow{c!\vec{d}}_0 \Gamma'_1 \triangleleft M' \triangleright P' \quad \Gamma_2 \triangleleft (M \triangleright R)\sigma_{\Gamma} \xrightarrow{c?\vec{d}}_0 \Gamma'_2 \triangleleft M' \triangleright R'}{\Gamma \triangleleft (M \triangleright P \parallel R)\sigma_{\Gamma} \xrightarrow{\tau}_0 \Gamma \triangleleft M' \triangleright P' \parallel R'} \text{LCom-L}}{\Gamma \triangleleft M \triangleright P \parallel R \xrightarrow{\tau}_0 \Gamma \triangleleft M' \triangleright P' \parallel R'} \text{LREN} \quad (4.36)$$

By the two top premises of (4.36) and Lemma 4.4 we know

$$M\sigma_{\Gamma} = M' \quad (4.37)$$

$$P\sigma_{\Gamma} \equiv c!\vec{d}.P_1 \parallel P_2 \quad P' \equiv P_1 \parallel P_2 \quad (4.38)$$

$$R\sigma_{\Gamma} \equiv c?\vec{x}.R_1 \parallel R_2 \quad R' \equiv R_1\{\vec{d}/\vec{x}\} \parallel R_2 \quad (4.39)$$

From $\Delta \vdash R$ we obtain $\Delta\sigma_{\Gamma} \vdash R\sigma_{\Gamma}$, and by (4.39), $\Delta \vdash R$ and Inversion we obtain

$$\Delta\sigma_{\Gamma} < \Delta_1, \Delta_2, c: [\vec{\mathbf{U}}]^a \quad (4.40)$$

$$\Delta_1, c: [\vec{\mathbf{U}}]^{a-1}, \vec{x}: \vec{\mathbf{U}} \vdash R_1 \quad (4.41)$$

$$\Delta_2 \vdash R_2 \quad (4.42)$$

Note that through (4.41) we know that

$$c: [\vec{\mathbf{U}}]^{a-1} \text{ is defined.} \quad (4.43)$$

By (4.41), the Substitution Lemma (Lemma 4.4 from [12]) and (4.42) we obtain

$$\Delta_1, \Delta_2, c: [\vec{\mathbf{U}}]^{a-1}, \vec{d}: \vec{\mathbf{U}} \vdash R_1\{\vec{d}/\vec{x}\} \parallel R_2 \quad (4.44)$$

From the assumption defining \mathcal{R} , and Lemma 4.22 we obtain

$$\Gamma, (\Delta\sigma_{\Gamma}) \vDash (M \triangleright P)\sigma_{\Gamma} \mathcal{R}^n (N \triangleright Q)\sigma_{\Gamma}, \quad (4.45)$$

and by (4.40) and Proposition 4.17 we know that $\Gamma, (\Delta\sigma_{\Gamma}) < \Gamma, \Delta_1, \Delta_2, c: [\vec{\mathbf{U}}]^a$ and also that $\Gamma, \Delta_1, \Delta_2, c: [\vec{\mathbf{U}}]^a$ is consistent. Thus by (4.45) and Lemma 4.19 we deduce

$$\Gamma, \Delta_1, \Delta_2, c: [\vec{\mathbf{U}}]^a \vDash (M \triangleright P)\sigma_{\Gamma} \mathcal{R}^n (N \triangleright Q)\sigma_{\Gamma} \quad (4.46)$$

Now by (4.43), (4.38), (4.37), LOut , LPar-L , LREN and Lemma 4.20 we deduce

$$\Gamma, \Delta_1, \Delta_2, c: [\vec{\mathbf{U}}]^a \triangleleft (M \triangleright P)\sigma_{\Gamma} \xrightarrow{c!\vec{d}}_0 \Gamma, \Delta'_1, \Delta'_2, c: [\vec{\mathbf{U}}]^{a-1}, \vec{d}: \vec{\mathbf{U}} \triangleleft M' \triangleright P' \quad (4.47)$$

and hence by (4.46) and I.H. we obtain

$$\Gamma, \Delta_1, \Delta_2, c: [\vec{\mathbf{U}}]^a \triangleleft (N \triangleright Q)\sigma_{\Gamma} \xrightarrow{c!\vec{d}}_k \Gamma, \Delta'_1, \Delta'_2, c: [\vec{\mathbf{U}}]^{a-1}, \vec{d}: \vec{\mathbf{U}} \triangleleft N' \triangleright Q' \quad (4.48)$$

$$\text{such that } \Gamma, \Delta_1, \Delta_2, c: [\vec{\mathbf{U}}]^{a-1}, \vec{d}: \vec{\mathbf{U}} \vDash (M' \triangleright P') \mathcal{R}^{n+k-0} (N' \triangleright Q') \quad (4.49)$$

From (4.48) and LREN we know

$$\Gamma, \Delta_1, \Delta_2, c: [\vec{\mathbf{U}}]^a \triangleleft ((N \triangleright Q)\sigma_\Gamma)\sigma'_{\Gamma_4} \xrightarrow{\tau}_{k_1}^* \Gamma, \Delta_1, \Delta_2, c: [\vec{\mathbf{U}}]^a \triangleleft N'' \triangleright Q'' \quad (4.50)$$

$$\Gamma, \Delta_1, \Delta_2, c: [\vec{\mathbf{U}}]^a \triangleleft N'' \triangleright Q'' \xrightarrow{c!\vec{d}}_0 \Gamma, \Delta_1, \Delta_2, c: [\vec{\mathbf{U}}]^{a-1}, \vec{d}: \vec{\mathbf{U}} \triangleleft N'' \triangleright Q''' \quad (4.51)$$

$$\Gamma, \Delta_1, \Delta_2, c: [\vec{\mathbf{U}}]^{a-1}, \vec{d}: \vec{\mathbf{U}} \triangleleft N'' \triangleright Q''' \xrightarrow{\tau}_{k_2}^* \Gamma, \Delta_1, \Delta_2, c: [\vec{\mathbf{U}}]^{a-1}, \vec{d}: \vec{\mathbf{U}} \triangleleft N' \triangleright Q' \quad (4.52)$$

$$\text{where } k = k_1 + k_2 \text{ and } \Gamma_4 = \Gamma, \Delta_1, \Delta_2, c: [\vec{\mathbf{U}}]^a \text{ (}\Gamma_4 \text{ is used in (4.50))} \quad (4.53)$$

From (4.50), (4.52), LPAR-L and Lemma 4.16 (Strengthening) we obtain:

$$\Gamma \triangleleft ((N \triangleright Q)\sigma_\Gamma)\sigma'_{\Gamma_4} \parallel R\sigma_\Gamma \xrightarrow{\tau}_{k_1}^* \Gamma \triangleleft N'' \triangleright Q'' \parallel R\sigma_\Gamma \quad (4.54)$$

$$\Gamma \triangleleft N'' \triangleright Q''' \parallel R' \xrightarrow{\tau}_{k_2}^* \Gamma \triangleleft N' \triangleright Q' \parallel R' \quad (4.55)$$

By (4.39), LIN and LPAR-L we can construct (for some Γ_6, Γ_7)

$$\Gamma_6 \triangleleft N'' \triangleright R\sigma_\Gamma \xrightarrow{c!\vec{d}}_0 \Gamma_7 \triangleleft N'' \triangleright R' \quad (4.56)$$

and by (4.51), (4.56) and LCOM-L we obtain

$$\Gamma \triangleleft N'' \triangleright Q'' \parallel R\sigma_\Gamma \xrightarrow{\tau}_0 \Gamma \triangleleft N'' \triangleright Q''' \parallel R' \quad (4.57)$$

By (4.40) and (4.53), we know that we can find an alternative renaming function σ''_{Γ_5} , where $\Gamma_5 = \Gamma, (\Delta\sigma_\Gamma)$, in a way that, from (4.54), we can obtain

$$\Gamma \triangleleft ((N \triangleright Q)\sigma_\Gamma)\sigma''_{\Gamma_5} \parallel R\sigma_\Gamma \xrightarrow{\tau}_{k_1}^* \Gamma \triangleleft N'' \triangleright Q'' \parallel R\sigma_\Gamma \quad (4.58)$$

By Definition 4.3, $\Delta\sigma_\Gamma \vdash R\sigma_\Gamma$, (4.40), (4.53) we know that $(R\sigma_\Gamma)\sigma''_{\Gamma_5} = R\sigma_\Gamma$, and also that σ''_{Γ_5} is also a renaming modulo Γ , so we can denote it as σ''_Γ and rewrite $((N \triangleright Q)\sigma_\Gamma)\sigma''_{\Gamma_5} \parallel R\sigma_\Gamma$ as $((N \triangleright Q \parallel R)\sigma_\Gamma)\sigma''_\Gamma$ in (4.58). Thus, by (4.58), (4.57), (4.55), (4.53) and LREN we obtain the matching move

$$\Gamma \triangleleft N \triangleright Q \parallel R \xRightarrow{\tau}_k \Gamma \triangleleft N' \triangleright Q' \parallel R'$$

since by (4.49), (4.44), (4.39) and the definition of \mathcal{R} we obtain

$$\Gamma \vDash (M' \triangleright P' \parallel R') \mathcal{R}^{n+k-0} (N' \triangleright Q' \parallel R')$$

as required. \square

Example 4.24 (Properties of \sqsubseteq_{bis}). From the proved statements $\Gamma_1 \vDash (M \triangleright C_1) \sqsubseteq_{\text{bis}} (M \triangleright C_0)$ and $\Gamma_1 \vDash (M \triangleright C_2) \sqsubseteq_{\text{bis}} (M \triangleright C_1)$ of Example 4.10, and by Corollary 4.13 (Preorder), we may conclude that

$$\Gamma_1 \vDash (M \triangleright C_2) \sqsubseteq_{\text{bis}} (M \triangleright C_0) \quad (4.59)$$

without the need to provide a bisimulation relation justifying (4.59). We also note that \mathcal{R}' of Example 4.10, justifying $\Gamma_1 \vDash (M \triangleright C_3) \sqsubseteq_{\text{bis}} (M \triangleright C_2)$ is a *bounded* amortised typed-bisimulation, and by Lemma 4.14 we can also conclude

$$\Gamma_1 \vDash (M \triangleright C_2) \sqsubseteq_{\text{bis}} (M \triangleright C_3)$$

and thus $\Gamma_1 \vDash (M \triangleright C_3) \simeq_{\text{bis}} (M \triangleright C_2)$. Finally, by Theorem 4.23, in order to show that

$$c: [\mathbf{T}_1, \mathbf{T}_2]^\omega \vDash (M \triangleright S_1 \parallel S_2 \parallel C_1) \sqsubseteq_{\text{bis}} (M \triangleright S_1 \parallel S_2 \parallel C_0)$$

it suffices to abstract away from the common code, $S_1 \parallel S_2$, and show $\Gamma_1 \vDash (M \triangleright C_1) \sqsubseteq_{\text{bis}} (M \triangleright C_0)$, as proved already in Example 4.10.

5. CHARACTERISATION

In this section we give a sound and complete characterization of bisimilarity in terms of the reduction semantics of Section 2, justifying the bisimulation relation and the respective LTS as a proof technique for reasoning about the behaviour of $R\pi$ processes. Our touchstone behavioural preorder is based on a costed version of families of reduction-closed barbed congruences along similar lines to [20]. In order to limit behaviour to safe computations, these congruences are defined as typed relations (Definition 4.2), where systems are subject to common observers typed by environments.

The observer type-environment delineates the observations that can be made: the observer can only make distinctions for channels that it has a permission for, *i.e.*, at least an affine typing assumption. The observations that can be made in our touchstone behavioural preorder are described as *barbs* [24] that take into account the permissions owned by the observer. We require systems related by our behavioral preorder to exhibit the same barbs *wrt.* a common observer.

Definition 5.1 (Barb). $(\Gamma \triangleleft M \triangleright P) \Downarrow_c^{\text{barb}} \stackrel{\text{def}}{=} (M \triangleright P) \longrightarrow_k^* \equiv (M' \triangleright P' \parallel c! \vec{d}.P'')$ and $c \in \text{dom}(\Gamma)$.

Definition 5.2 (Barb Preservation). A typed relation \mathcal{R} is barb preserving if and only if

$$\Gamma \vDash M \triangleright P \mathcal{R} N \triangleright Q \text{ implies } (\Gamma \triangleleft M \triangleright P \Downarrow_c^{\text{barb}} \text{ iff } \Gamma \triangleleft N \triangleright Q \Downarrow_c^{\text{barb}}).$$

Our behavioural preorder takes cost into consideration; it is defined in terms of families of amortised typed relations that are closed under costed reductions.

Definition 5.3 (Cost Improving). An amortized type-indexed relation \mathcal{R} is *cost improving at credit n* iff whenever $\Gamma \vDash (M \triangleright P) \mathcal{R}^n (N \triangleright Q)$ and

- (1) if $M \triangleright P \longrightarrow_k M' \triangleright P'$ then $N \triangleright Q \longrightarrow_l^* N' \triangleright Q'$ such that $\Gamma \vDash (M' \triangleright P') \mathcal{R}^{n+l-k} (N' \triangleright Q')$;
- (2) if $N \triangleright Q \longrightarrow_l N' \triangleright Q'$ then $M \triangleright P \longrightarrow_k^* M' \triangleright P'$ such that $\Gamma \vDash (M' \triangleright P') \mathcal{R}^{n+l-k} (N' \triangleright Q')$.

Related processes must be related under arbitrary (parallel) contexts; moreover, these contexts must be allowed to allocate new channels. We note that the second clause of our contextuality definition, Definition 5.4, is similar to that discussed earlier in Section 4.4, where we *transfer* the respective permissions held by the observer along with the test R placed in parallel with the processes. This is essential in order to preserve consistency (see Definition 2.2) thus limiting our analysis to safe computations. Definition 5.4 also requires an additional condition, when compared to the contextuality definition discussed in Section 4.4, namely that of *resource extensions* where we consider systems in larger resource contexts (owned exclusively by the observer). This is described by the first clause in the definition; we recall the implicit condition for resource environment representations from Section 2, requiring the channel c not to be present (thus allocated) in M (*resp.* N) for the resource environment to be well-formed — c is therefore fresh. In order to disambiguate between the different contextuality definitions, we refer to Definition 5.4 as *full contextuality*.

Definition 5.4 (Full Contextuality). An amortized type-indexed relation \mathcal{R} is contextual at environment Γ and credit n iff whenever $\Gamma \vDash (M \triangleright P) \mathcal{R}^n (N \triangleright Q)$:

- (1) $\Gamma, c: [\vec{T}]^\bullet \vDash (M, c \triangleright P) \mathcal{R}^n (N, c \triangleright Q)$
- (2) If $\Gamma < \Gamma_1, \Gamma_2$ where $\Gamma_2 \vdash R$ then
 - $\Gamma_1 \vDash (M \triangleright P \parallel R) \mathcal{R}^n (N \triangleright Q \parallel R)$ and
 - $\Gamma_1 \vDash (M \triangleright R \parallel P) \mathcal{R}^n (N \triangleright R \parallel Q)$

We can now define the preorder defining our notion of observational system efficiency:

Definition 5.5 (Behavioral Contextual Preorder). $\lesssim_{\text{beh}}^{\Gamma, n}$ is the largest family of amortized typed relations that is:

- Barb Preserving;
- Cost Improving at credit n ;
- Full contextual at environment Γ .

A system $M \triangleright P$ is said to be behaviourally as efficient as another system $N \triangleright Q$ wrt. an observer Γ , denoted as $\Gamma \vDash M \triangleright P \lesssim_{\text{beh}} N \triangleright Q$, whenever there exists an amortisation credit n such that $\Gamma \vDash M \triangleright P \lesssim_{\text{beh}}^n N \triangleright Q$. Similarly, we can lift our preorder to processes: a process P is said to be as efficient as Q wrt. M and Γ whenever there exists an n such that $\Gamma \vDash M \triangleright P \lesssim_{\text{beh}}^n M \triangleright Q$

5.1. Soundness for \sqsubseteq_{bis} . Through Definition 5.5 we are able to articulate why clients C_2 and C'_2 should be deemed to be behaviourally equally efficient wrt. Γ_1 of (4.10): for an appropriate M , it turns out that we cannot differentiate between the two processes under any context allowed by Γ . Unfortunately, the universal quantification of contexts of Definition 5.4 makes it hard to verify such a statement. Through Theorem 5.10 we can however establish that our bisimulation preorder of Definition 4.9 provides a sound technique for determining behavioural efficiency. This Theorem, in turn, relies on the lemmas we outline below. In particular, Lemma 5.7 and Lemma 5.8 prove that bisimulations are barb-preserving and cost-improving, whereas Lemma 5.9 proves that bisimulations are preserved under resource extensions. The required result then follows from Theorem 4.23 of Section 4.4.

Lemma 5.6 (Reductions and Bijective Renaming). *For any bijective renaming σ , $(M \triangleright P)\sigma \longrightarrow_k (M' \triangleright P')\sigma$ implies $M \triangleright P \longrightarrow_k M' \triangleright P'$*

Proof. By rule induction on $(M \triangleright P)\sigma \longrightarrow_k (M' \triangleright P')\sigma$. □

Lemma 5.7 (Barb Preservation).

$$\Gamma \vDash M \triangleright P \sqsubseteq_{\text{bis}} N \triangleright Q \text{ and } \Gamma \triangleleft M \triangleright P \Downarrow_c^{\text{barb}} \text{ implies } \Gamma \triangleleft N \triangleright Q \Downarrow_c^{\text{barb}}$$

Proof. By Definition 5.1 we know $M \triangleright P \longrightarrow_l^* \equiv (M' \triangleright P' \parallel c!d.P'')$ where $c \in \text{dom}(\Gamma)$. By Lemma 4.6(1) we obtain $\Gamma \triangleleft M \triangleright P \Longrightarrow_l \Gamma \triangleleft M' \triangleright P''$ where $P'' \equiv (P' \parallel c!d.P'')$. Moreover, by LOUT , LPAR-R and Lemma 4.20 we deduce $\Gamma \triangleleft M \triangleright P \xrightarrow{c!d} \equiv \Gamma' \triangleleft M \triangleright P' \parallel P''$. By $\Gamma \vDash M \triangleright P \sqsubseteq_{\text{bis}} N \triangleright Q$ we know that there exists a move $\Gamma \triangleleft N \triangleright Q \xrightarrow{c!d} \equiv \Gamma' \triangleleft N' \triangleright Q'$ and from this matching move, Lemma 4.6(2) (for the initial τ moves of the weak action) and Lemma 4.4 we obtain $(N \triangleright Q)\sigma_\Gamma \longrightarrow_{k_1}^* \equiv (N'' \triangleright Q'' \parallel c!d.Q''')\sigma_\Gamma$, which, together with $c \in \text{dom}(\Gamma)$ and Lemma 5.6, implies $N \triangleright Q \longrightarrow_{k_1}^* \equiv N'' \triangleright Q'' \parallel c!d.Q'''$ i.e., c is unaffected by the renaming σ_Γ , and thus $\Gamma \triangleleft N \triangleright Q \Downarrow_c^{\text{barb}}$. □

Lemma 5.8 (Cost Improving). *$\Gamma \vDash M \triangleright P \sqsubseteq_{\text{bis}}^n N \triangleright Q$ and $M \triangleright P \longrightarrow_l M' \triangleright P'$ then there exist some $N' \triangleright Q'$ such that $N \triangleright Q \longrightarrow_k^* N' \triangleright Q'$ and $\Gamma \vDash M' \triangleright P' \sqsubseteq_{\text{bis}}^{n+k-l} N' \triangleright Q'$*

Proof. By $M \triangleright P \longrightarrow_l M' \triangleright P'$ and Lemma 4.6(1) we know $\Gamma \triangleleft M \triangleright P \xrightarrow{\tau} \equiv \Gamma \triangleleft M \triangleright P''$ where $P'' \equiv P'$. By Definition 4.9 and assumption $\Gamma \vDash M \triangleright P \sqsubseteq_{\text{bis}}^n N \triangleright Q$, this implies that $\Gamma \triangleleft N \triangleright Q \Longrightarrow_k \Gamma \triangleleft M' \triangleright P'$ where

$$\Gamma \vDash M' \triangleright P' \sqsubseteq_{\text{bis}}^{n+k-l} N' \triangleright Q'. \quad (5.1)$$

By Lemma 4.6(2) we deduce $(N \triangleright Q)\sigma_\Gamma \longrightarrow_k^* N' \triangleright Q'$ and by Lemma 5.6 we obtain $N \triangleright Q \longrightarrow^* N'' \triangleright Q''$ where $N'' \triangleright Q'' = (N' \triangleright Q')\sigma_\Gamma$. The required result follows from $\Gamma \vDash M' \triangleright P' \sqsubseteq_{\text{bis}}^0 M' \triangleright P''$, which we obtain from $P' \equiv P''$ and Corollary 4.21 (Structural Equivalence and Bisimilarity), (5.1), $\Gamma \vDash N'' \triangleright Q'' \sqsubseteq_{\text{bis}}^0 N' \triangleright Q'$ which we obtain from Lemma 4.11 (Reflexivity upto Renaming) and $N'' \triangleright Q'' = (N' \triangleright Q')\sigma_\Gamma$, and Lemma 4.12. □

Lemma 5.9 (Resource Extensions).

$$\Gamma \vDash M \triangleright P \stackrel{n}{\approx}_{bis} N \triangleright Q \text{ implies } \Gamma, c : [\mathbf{T}]^{\bullet} \vDash (M, c) \triangleright P \stackrel{n}{\approx}_{bis} (N, c) \triangleright Q$$

Proof. By coinduction. \square

Theorem 5.10 (Soundness). $\Gamma \vDash (M \triangleright P) \stackrel{n}{\approx}_{bis} (N \triangleright Q) \text{ implies } \Gamma \vDash (M \triangleright P) \lesssim_{beh}^n (N \triangleright Q)$.

Proof. Follows from Lemma 5.7 (Barb Preservation), Lemma 5.8 (Cost Improving), Lemma 5.9 (Resource Extensions) and Theorem 4.23 (Contextuality). \square

Corollary 5.11 (Soundness). $\Gamma \vDash (M \triangleright P) \stackrel{n}{\approx}_{bis} (N \triangleright Q) \text{ implies } \Gamma \vDash (M \triangleright P) \lesssim_{beh} (N \triangleright Q)$.

5.2. Full Abstraction of \lesssim_{beh} . To prove completeness, *i.e.*, that for every behavioural contextual preorder there exists a corresponding amortised typed-bisimulation, we rely on the adapted notion of *action definability* [19, 21], which intuitively means that every action (label) used by our LTS can, in some sense, be simulated (observed) by a specific test context. For our specific case, two important aspects need to be taken into consideration:

- the *typeability* of the testing context *wrt.* our substructural type system;
- the *cost* of the action simulation, which has to correspond to the cost of the action being observed.

These aspects are formalised in Definition 5.13, which relies on the functions definitions *doml* and *codl*:

$$\begin{aligned} \text{doml}(\epsilon) &\stackrel{\text{def}}{=} \epsilon & \text{codl}(\epsilon) &\stackrel{\text{def}}{=} \epsilon \\ \text{doml}(\Gamma, c : \mathbf{T}) &\stackrel{\text{def}}{=} \text{doml}(\Gamma), c & \text{codl}(\Gamma, c : \mathbf{T}) &\stackrel{\text{def}}{=} \text{codl}(\Gamma), \mathbf{T} \end{aligned}$$

These two meta-functions take a substructural type environment and returning respectively a *list* of channel names and a *list* of types. For example, for the environment $\Gamma = c : [\mathbf{T}]^{\mathbf{1}}, d : [\mathbf{T}']^{\omega}, c : [\mathbf{T}]^{(\bullet, \mathbf{1})}$, we have $\text{doml}(\Gamma) = c, d, c$ and $\text{codl}(\Gamma) = [\mathbf{T}]^{\mathbf{1}}, [\mathbf{T}']^{\omega}, [\mathbf{T}]^{(\bullet, \mathbf{1})}$.

Before stating cost-definability for actions, Definition 5.13, we prove the technical Lemma 5.12 which allows us to express transitions in a convenient format for the respective definition without loss of generality.

Lemma 5.12 (Transitions and Renaming). $\Gamma \triangleleft M \triangleright P \xrightarrow{\mu}_k \Gamma' \triangleleft M' \triangleright P'$ if and only if $\Gamma \triangleleft M \triangleright P \xrightarrow{\mu}_k (\Gamma'' \triangleleft M'' \triangleright P'') \sigma_{\Gamma}$ for some $\sigma_{\Gamma}, \Gamma'', M'', P''$ where $\Gamma' = \Gamma'' \sigma_{\Gamma}, M' = M'' \sigma_{\Gamma}$ and $P' = P'' \sigma_{\Gamma}$.

Proof. The *if* case is immediate. The proof for the *only-if* is complicated by actions that perform channel allocation (see `LALL` and `LALLE` from Figure 5) because, in such cases, the renaming used in `LREN`'s premise cannot be used directly. More precisely, from the premise we know:

$$\frac{\Gamma \triangleleft (M \triangleright P) \sigma_{\Gamma} \xrightarrow{\mu}_k \Gamma' \triangleleft M' \triangleright P'}{\Gamma \triangleleft M \triangleright P \xrightarrow{\mu}_k \Gamma' \triangleleft M' \triangleright P'} \text{LREN}$$

and the required result follows if we prove the (slightly more cumbersome) sublemma:

Sublemma (Transition and Renaming). $\Gamma \triangleleft (M \triangleright P) \sigma_{\Gamma} \xrightarrow{\mu}_k \Gamma' \triangleleft M' \triangleright P'$ where $\text{fn}(P) \subseteq M$ implies $\Gamma \triangleleft (M \triangleright P) \sigma_{\Gamma} \xrightarrow{\mu}_k (\Gamma'' \triangleleft M'' \triangleright P'') \sigma'_{\Gamma}$ for some $\sigma'_{\Gamma}, \Gamma'', M'', P''$ where

- $\Gamma' = \Gamma'' \sigma'_{\Gamma}, M' = M'' \sigma'_{\Gamma}$ and $P' = P'' \sigma'_{\Gamma}$;
- $c \in \text{dom}(M)$ implies $\sigma_{\Gamma}(c) = \sigma'_{\Gamma}(c)$

The above sublemma is proved by rule induction on $\Gamma \triangleleft (M \triangleright P) \sigma_{\Gamma} \xrightarrow{\mu}_k \Gamma' \triangleleft M' \triangleright P'$. We show one of the main cases:

LALL: We have $\Gamma \triangleleft (M \triangleright \text{alloc } x.P)\sigma_\Gamma \xrightarrow{\tau}_{+1} \Gamma \triangleleft ((M)\sigma_\Gamma, c) \triangleright ((P)\sigma_\Gamma\{c/x\})$. From the fact that $c \notin (M\sigma_\Gamma)$ — it follows because $((M)\sigma_\Gamma, c)$ is defined — we know that $\sigma_\Gamma^{-1}(c) \notin M$. We thus choose some fresh channel d , i.e., $d \notin (M \cup (M\sigma_\Gamma) \cup \text{dom}(\Gamma))$ ¹², and define σ'_Γ as σ_Γ , except that it maps d to c and also maps $\sigma_\Gamma^{-1}(c)$ (i.e., the channel name that mapped to c in σ_Γ) to $\sigma_\Gamma(d)$, since this channel is not mapped to by d anymore (in order to preserve bijectivity):

$$\sigma'_\Gamma(x) \stackrel{\text{def}}{=} \begin{cases} c & \text{if } x = d \\ \sigma_\Gamma(d) & \text{if } x = \sigma_\Gamma^{-1}(c) \\ \sigma_\Gamma(x) & \text{otherwise} \end{cases}$$

We subsequently define

- Γ'' as Γ since $\Gamma\sigma'_\Gamma = \Gamma\sigma_\Gamma = \Gamma$;
- M'' as M, d since $(M, d)\sigma'_\Gamma = ((M)\sigma_\Gamma, c)$; and
- P'' as $P\{d/x\}$ since $P\{d/x\}\sigma'_\Gamma = P\sigma_\Gamma\{c/x\}$ □

Definition 5.13 (Cost Definable Actions). An action μ is cost-definable iff for any pair of type environments¹³ Γ and Γ' , a corresponding substitution σ_Γ , a set of channel names $C \in \text{CHAN}$, and channel names $\text{succ}, \text{fail} \notin C$, there exists a test R such that $\Gamma, \text{succ}:[\text{codl}(\Gamma')]^1, \text{fail}:[\]^1, \text{fail}:[\]^1 \vdash R$ and whenever $M \in C$:

- (1) $\Gamma \triangleleft M \triangleright P \xrightarrow{\mu}_k (\Gamma' \triangleleft M' \triangleright P')\sigma_\Gamma$ implies $M, \text{succ}, \text{fail} \triangleright P \parallel R \xrightarrow{*}_k M', \text{succ}, \text{fail} \triangleright P' \parallel \text{succ}!(\text{doml}(\Gamma'))$.
- (2) $M, \text{succ}, \text{fail} \triangleright P \parallel R \xrightarrow{*}_k M'' \triangleright P''$ where $\text{succ}:[\text{codl}(\Gamma')]^a, \text{fail}:[\]^a \triangleleft M'' \triangleright P'' \Downarrow_{\text{fail}}^{\text{barb}}$ and $\text{succ}:[\text{codl}(\Gamma')]^a, \text{fail}:[\]^a \triangleleft M'' \triangleright P'' \Downarrow_{\text{succ}}^{\text{barb}}$ implies $\Gamma \triangleleft M \triangleright P \xrightarrow{\mu}_k (\Gamma' \triangleleft M' \triangleright P')\sigma_\Gamma$ where $M'' = M', \text{succ}, \text{fail}$ and $P'' \equiv P''' \parallel \text{succ}!(\text{doml}(\Gamma'))$.

Lemma 5.14 (Action Cost-Definability). *External actions $\mu \in \{c!\vec{d}, c?\vec{d}, \text{alloc}, \text{free } c \mid c, \vec{d} \subset \text{CHAN}\}$ are cost-definable.*

Proof. The witness tests for $c!\vec{d}$ and $c?\vec{d}$ are reasonably standard (see [19]), but need to take into account permission transfer. For instance, for the specific case of the action $c!d$ where $d \notin \text{doml}(\Gamma)$, if the transition $\Gamma \triangleleft M \triangleright P \xrightarrow{\mu}_k (\Gamma' \triangleleft M' \triangleright P')\sigma_\Gamma$ holds then we know that, for some Γ_1 and $[\mathbf{T}]^a$:

- $\Gamma = \Gamma_1, c:[\mathbf{T}]^a$;
- $\Gamma'\sigma_\Gamma = \Gamma_1, c:[\mathbf{T}]^{a-1}, d:\mathbf{T}$

In particular, when $a = \mathbf{1}$ (affine), using the permission to input on c implicitly transfers the permission to process P (see Section 4.1), potentially revoking the test's capability to perform name matching on channel name c (see τIf in Figure 4) — this happens if $c \notin \text{dom}(\Gamma_1)$. For this reason, when $a = \mathbf{1}$ the test is defined as

$$\text{fail}! \parallel c?x.\text{if } (x \in \text{doml}(\Gamma_1)) \text{ then nil else fail?}.\text{succ}!(\text{doml}(\Gamma'))$$

where $x \in \text{doml}(\Gamma_1)$ is shorthand for a sequence of name comparisons as in [19]. Otherwise, the respective type assumption is not consumed from the observer environment and the test is defined as

$$\text{fail}! \parallel c?x.\text{if } (x \in \text{doml}(\Gamma)) \text{ then nil else fail?}.\text{succ}!(\text{doml}(\Gamma'))$$

¹²The condition that $d \notin \text{dom}(\Gamma)$ is required since we do not state whether the triple $\Gamma \triangleleft M \triangleright P$ is a configuration; otherwise, it is redundant — see comments succeeding Definition 4.1.

¹³Cost Definability cannot be defined wrt. the first environment only in the case of action alloc , since it non-deterministically allocates a fresh channel name and adds it to the residual environment - see LALLE in Figure 5.

Note that name comparisons on freshly acquired names are typeable since we also obtain the respective permissions upon input, *i.e.*, the explicit permission transfer (see Section 4.1). The reader can verify that these tests typecheck *wrt.* the environment Γ , $\text{succ} : [\text{codl}(\Gamma')]^1$, $\text{fail} : []^1$, $\text{fail} : []^1$ and that they observe clauses (1) and (2) of Definition 5.13. In the case of clause (2), we note that from the typing of the tests above, we know that $c \in \text{doml}(\Gamma)$ must hold (because both tests use channel c for input); this is a key requirement for the transition to fire — see `LOut` of Figure 5.

The witness tests for `alloc` and `free c` involve less intricate permission transfer and are respectively defined as:

$$\text{fail!} \parallel \text{alloc } x.\text{fail}?.\text{succ}!(\text{doml}(\Gamma), x)$$

and

$$\text{fail!} \parallel \text{free } c.\text{fail}?.\text{succ}!(\text{doml}(\Gamma'))$$

We here focus on `alloc` and leave the analogous proof for `free c` for the interested reader:

- (1) If $\Gamma \triangleleft M \triangleright P \xrightarrow{\text{alloc}}_k (\Gamma' \triangleleft M' \triangleright P')\sigma_\Gamma$ we know that, for some $d \notin M$ and $c \notin M\sigma_\Gamma$ where $\sigma_\Gamma(d) = c$, we have $(\Gamma')\sigma_\Gamma = (\Gamma, d : [\mathbf{T}]^\bullet)\sigma_\Gamma = \Gamma, c : [\mathbf{T}]^\bullet$, $M' = (M, d)$ and $P' = P$. We can therefore simulate this action by the following sequence of reductions:

$$\begin{aligned} M \triangleright P \parallel \text{fail!} \parallel \text{alloc } x.\text{fail}?.\text{succ}!(\text{doml}(\Gamma), x) &\longrightarrow \\ M, d \triangleright P \parallel \text{fail!} \parallel \text{fail}?.\text{succ}!(\text{doml}(\Gamma), d) &\longrightarrow M, d \triangleright P \parallel \text{succ}!(\text{doml}(\Gamma), d) \end{aligned}$$

- (2) From the structure of R and the assumption that $\text{fail}, \text{succ} \notin \text{fn}(P)$, we conclude that, if $\text{succ} : [\text{codl}(\Delta)]^a$, $\text{fail} : []^a \triangleleft M' \triangleright P' \Downarrow_{\text{fail}}^{\text{barb}}$ and $\text{succ} : [\text{codl}(\Delta)]^a$, $\text{fail} : []^a \triangleleft M' \triangleright P' \Downarrow_{\text{succ}}^{\text{barb}}$, then it must be the case that, for some $d \notin M$, $P' = P'' \parallel \text{succ}!(\text{doml}(\Gamma), d)$ where $M'' = (M', \text{succ}, \text{fail}, d)$ for some M' .

Since P and R do not share common channels there could not have been any interaction between the two processes in the reduction sequence $M, \text{succ}, \text{fail} \triangleright P \parallel R \xrightarrow{*}_k M' \triangleright P'$. Within this reduction sequence, from every reduction $M_i \triangleright P_i \parallel R' \xrightarrow{k_i} M_{i+1} \triangleright P_{i+1} \parallel R'$ resulting from derivatives of P , *i.e.*, $M_i \triangleright P_i \xrightarrow{k_i} M_{i+1} \triangleright P_{i+1}$ that happened before the allocation of channel d , we obtain a corresponding silent transition

$$\Gamma_i \triangleleft (M_i \setminus \{\text{succ}, \text{fail}\}) \triangleright P_i \xrightarrow{k_i} \tau \Gamma_i \triangleleft (M_{i+1} \setminus \{\text{succ}, \text{fail}\}) \triangleright P_{i+1} \quad (5.2)$$

by Lemma 4.6(1) and an appropriate lemma that uses the fact $\{\text{succ}, \text{fail}\} \cap \text{fn}(P) = \emptyset$ to allow us to shrink the allocated resources from M_i to $(M_i \setminus \{\text{succ}, \text{fail}\})$. A similar procedure can be carried out for reductions that happened after the allocation of d as a result of reductions from P derivatives, and by applying renaming σ_Γ we can obtain

$$(\Gamma_i \triangleleft (M_i \setminus \{\text{succ}, \text{fail}\}) \triangleright P_i)\sigma_\Gamma \xrightarrow{k_i} \tau (\Gamma_i \triangleleft (M_{i+1} \setminus \{\text{succ}, \text{fail}\}) \triangleright P_{i+1})\sigma_\Gamma \quad (5.3)$$

The reduction

$$\begin{aligned} M_i, \text{succ}, \text{fail} \triangleright P_i \parallel \text{alloc } x.\text{fail}?.\text{succ}!(\text{doml}(\Gamma), x) &\longrightarrow_{+1} \\ M_i, \text{succ}, \text{fail}, d \triangleright P_i \parallel \text{fail}?.\text{succ}!(\text{doml}(\Gamma), d) & \end{aligned}$$

can be substituted by the transition

$$\Gamma_i \triangleleft M_i \triangleright P_i \xrightarrow{\text{alloc}}_{+1} \Gamma_i, (d)\sigma_\Gamma : [\mathbf{T}]^\bullet \triangleleft ((M_i)\sigma_\Gamma, (d)\sigma_\Gamma) \triangleright (P_i)\sigma_\Gamma \quad (5.4)$$

This follows from the fact that $d \notin M_i$ and the fact that σ_Γ is a bijection, which implies that $(d)\sigma_\Gamma \notin (M_i)\sigma_\Gamma$ (necessary for $((M_i)\sigma_\Gamma, (d)\sigma_\Gamma)$ to be a valid resource environment). By joining together the transitions from (5.2), (5.4) and (5.3) in the appropriate sequence we obtain the required weak transition. \square

The proof of Theorem 5.18 (Completeness) relies on Lemma 5.14 to simulate a costed action by the appropriate test and is, for the most part, standard. As stated already, one novel aspect is that the cost semantics requires the simulation to incur the same cost as that of the costed action. Through Reduction Closure, Lemma 5.14 again, and then finally the Extrusion Lemma 5.15 we then obtain the matching bisimulation move which preserves the relative credit index. Another novel aspect of the proof for Theorem 5.18 is that the name matching in the presence of our substructural type environment requires a reformulation of the Extrusion Lemma. More precisely, in the case of the output actions, the simulating test requires all of the environment permissions to perform all the necessary name comparisons. We then make sure that these permissions are not lost by communicating them all again on succ; this passing on of permissions then allows us to show contextuality in Lemma 5.15.

Lemma 5.15 (Extrusion). *Whenever $\Gamma \triangleleft M \triangleright P$ and $\Gamma \triangleleft N \triangleright Q$ are configurations and $\vec{d} \notin \text{dom}(\Gamma)$:*

$$\text{succ} : [\text{codl}(\Gamma)]^{(\bullet,1)} \vDash (M, \text{succ}, \vec{d}) \triangleright P \parallel \text{succ}!(\text{doml}(\Gamma)) \lesssim_{\text{beh}}^n (N, \text{succ}, \vec{d}) \triangleright Q \parallel \text{succ}!(\text{doml}(\Gamma))$$

implies $\Gamma \vdash M \triangleright P \lesssim_{\text{beh}}^n N \triangleright Q$

Proof. By coinduction we show that a family of amortized typed relations $\Gamma \vdash M \triangleright P \mathcal{R}^n N \triangleright Q$ observes the required properties of Definition 5.5. Note that the environment $\text{succ} : [\text{codl}(\Gamma)]^{(\bullet,1)}$ ensures that $\text{succ} \notin \text{names}(P, Q)$ since both $P \parallel \text{succ}!(\text{doml}(\Gamma))$ and $Q \parallel \text{succ}!(\text{doml}(\Gamma))$ must typecheck wrt. a type environment that is consistent with $\text{succ} : [\text{codl}(\Gamma)]^{(\bullet,1)}$. Cost improving is straightforward and Barb Preserving and Contextuality follow standard techniques; see [19].

For instance, for barb preservation we are required to show that $\Gamma \triangleleft M \triangleright P \Downarrow_c^{\text{barb}}$ implies $\Gamma \triangleleft N \triangleright Q \Downarrow_c^{\text{barb}}$ (and viceversa). From $\Gamma \triangleleft M \triangleright P \Downarrow_c^{\text{barb}}$ and Definition 5.1 we know that $c : [\vec{\mathbf{T}}]^a \in \Gamma$ at some index i . We can therefore define the process $R \triangleq \text{succ}?\vec{x}.x_i?\vec{y}.ok!$ where $|\vec{\mathbf{T}}| = |\vec{y}|$; this test process typechecks wrt. $\text{succ} : [\text{codl}(\Gamma)]^{(\bullet,1)}, ok : []^1$. Now by Definition 5.4(1) we know

$$\begin{aligned} \text{succ} : [\text{codl}(\Gamma)]^{(\bullet,1)}, ok : []^1 \vDash (M, \text{succ}, \vec{d}, ok) \triangleright P \parallel \text{succ}!(\text{doml}(\Gamma)) \\ \lesssim_{\text{beh}}^n (N, \text{succ}, \vec{d}, ok) \triangleright Q \parallel \text{succ}!(\text{doml}(\Gamma)) \end{aligned}$$

and thus, by Definition 5.4(2) and $\text{succ} : [\text{codl}(\Gamma)]^{(\bullet,1)}, ok : []^1 \vdash R$

$$\begin{aligned} ok : []^{(\bullet,1)} \vDash (M, \text{succ}, \vec{d}, ok) \triangleright P \parallel \text{succ}!(\text{doml}(\Gamma)) \parallel R \\ \lesssim_{\text{beh}}^n (N, \text{succ}, \vec{d}, ok) \triangleright Q \parallel \text{succ}!(\text{doml}(\Gamma)) \parallel R \end{aligned} \quad (5.5)$$

Clearly, if $\Gamma \triangleleft M \triangleright P \Downarrow_c^{\text{barb}}$ then $(ok : []^{(\bullet,1)} \triangleleft (M, \text{succ}, \vec{d}, ok) \triangleright (P \parallel \text{succ}!(\text{doml}(\Gamma)) \parallel R)) \Downarrow_{ok}^{\text{barb}}$. By (5.5) and Definition 5.2 we must have $(ok : []^{(\bullet,1)} \triangleleft (N, \text{succ}, \vec{d}, ok) \triangleright (Q \parallel \text{succ}!(\text{doml}(\Gamma)) \parallel R)) \Downarrow_{ok}^{\text{barb}}$ as well, which can only happen if $N \triangleright Q \longrightarrow^* \equiv Q' \parallel c!\vec{d}.Q''$. This means that $\Gamma \triangleleft N \triangleright Q \Downarrow_c^{\text{barb}}$. \square

Lemma 5.16. $\Gamma \vDash M \triangleright P \lesssim_{\text{beh}}^n N \triangleright Q$ and $\Gamma \triangleleft \Gamma'$ implies $\Gamma' \vDash M \triangleright P \lesssim_{\text{beh}}^n N \triangleright Q$

Proof. By coinduction. \square

Lemma 5.17. $\Gamma \vDash M \triangleright P \lesssim_{\text{beh}}^n N \triangleright Q$ and σ is a bijective renaming implies $\Gamma\sigma \vDash (M \triangleright P)\sigma \lesssim_{\text{beh}}^n (N \triangleright Q)\sigma$

Proof. By coinduction. \square

Theorem 5.18 (Completeness). $\Gamma \vDash (M \triangleright P) \lesssim_{\text{beh}}^n (N \triangleright Q)$ implies $\Gamma \vDash (M \triangleright P) \sqsubseteq_{\text{bis}}^n (N \triangleright Q)$.

Proof. By coinduction, we show that for arbitrary Γ, n , the family of relations included in $\Gamma \vDash M \triangleright P \lesssim_{\text{beh}}^n N \triangleright Q$ observes the transfer properties of Definition 4.9 at Γ, n . Assume

$$\Gamma \triangleleft M \triangleright P \xrightarrow{\mu}_k (\Gamma' \triangleleft M' \triangleright P') \sigma_{\Gamma} \quad (5.6)$$

If $\mu = \tau$, the matching move follows from Lemma 4.6, Definition 5.3 and Definition 5.5.

If $\mu \in \{c!\vec{d}, c?\vec{d}, \text{alloc}, \text{free } c \mid c, \vec{d} \in \text{CHAN}\}$, by Lemma 5.14 we know that there exists a test process that can simulate it; we choose one such test R with channel names $\text{succ}, \text{fail} \notin M, N$. By Definition 5.4(1) we know

$$\Gamma, \text{succ} : [\text{codl}(\Gamma)]^{\bullet}, \text{fail} : []^{\bullet} \vDash M, \text{succ}, \text{fail} \triangleright P \lesssim_{\text{beh}}^n N, \text{succ}, \text{fail} \triangleright Q$$

and by Definition 5.4(2) and $\Gamma, \text{succ} : [\text{codl}(\Delta)]^1, \text{fail} : []^1, \text{fail} : []^1 \vdash R$ (Definition 5.13) we obtain

$$\text{succ} : [\text{codl}(\Gamma)]^{(\bullet,1)}, \text{fail} : []^{(\bullet,2)} \vDash (M, \text{succ}, \text{fail}) \triangleright P \parallel R \lesssim_{\text{beh}}^n (N, \text{succ}, \text{fail}) \triangleright Q \parallel R \quad (5.7)$$

From (5.6) and Definition 5.13(1), we know

$$(M, \text{succ}, \text{fail}) \triangleright P \parallel R \xrightarrow{*}_k (M', \text{succ}, \text{fail}) \triangleright P' \parallel \text{succ}! \text{doml}(\Gamma')$$

By (5.7) and Definition 5.3 (Cost Improving) we know

$$(N, \text{succ}, \text{fail}) \triangleright Q \parallel R \xrightarrow{*}_i N'' \triangleright Q''$$

where

$$\text{succ} : [\text{codl}(\Gamma)]^{(\bullet,1)}, \text{fail} : []^{(\bullet,2)} \vDash (M', \text{succ}, \text{fail}) \triangleright P' \parallel \text{succ}! \text{doml}(\Gamma') \lesssim_{\text{beh}}^{n+l-k} N'' \triangleright Q'' \quad (5.8)$$

By Definition 5.2 (Barb Preservation), this means that $\text{succ} : [\text{codl}(\Gamma)]^{(\bullet,1)}, \text{fail} : []^{(\bullet,2)} \triangleleft N' \triangleright Q' \Downarrow_{\text{fail}}^{\text{barb}}$ and also that $\text{succ} : [\text{codl}(\Gamma)]^{(\bullet,1)}, \text{fail} : []^{(\bullet,2)} \triangleleft N' \triangleright Q' \Downarrow_{\text{succ}}^{\text{barb}}$. By Definition 5.13(2) we obtain

$$Q'' \equiv Q' \parallel \text{succ}! \text{doml}(\Gamma') \text{ and } N'' = (N', \text{succ}, \text{fail}) \quad (5.9)$$

$$\Gamma \triangleleft N \triangleright Q \xRightarrow{\mu}_l (\Gamma' \triangleleft N' \triangleright Q') \sigma_{\Gamma} \quad (5.10)$$

Transition (5.10) is the matching move because by (5.8) and Lemma 5.16 we obtain

$$\text{succ} : [\text{codl}(\Gamma)]^{(\bullet,1)} \vDash (M', \text{succ}, \text{fail}) \triangleright P' \parallel \text{succ}! \text{doml}(\Gamma') \lesssim_{\text{beh}}^{n+l-k} N'' \triangleright Q''$$

By (5.9), and Lemma 5.15 we obtain $\Gamma' \vDash M' \triangleright P' \lesssim_{\text{beh}}^{n+l-k} N' \triangleright Q'$ and subsequently by Lemma 5.17 we obtain

$$\Gamma' \sigma_{\Gamma} \vDash (M' \triangleright P') \sigma_{\Gamma} \lesssim_{\text{beh}}^{n+l-k} (N' \triangleright Q') \sigma_{\Gamma}$$

as required. \square

6. REVISITING THE CASE STUDY

We can formally express that eBuff is (strictly) more efficient than Buff in terms of the reduction semantics outlined in Section 2 through the following statements:

$$\Gamma_{\text{ext}} \vDash M \triangleright \text{eBuff} \lesssim_{\text{beh}} M \triangleright \text{Buff} \quad (6.1)$$

$$\Gamma_{\text{ext}} \vDash M \triangleright \text{Buff} \not\lesssim_{\text{beh}} M \triangleright \text{eBuff} \quad (6.2)$$

In order to show that the second statement (6.2) holds, we need to prove that *there is no amortisation credit* n for which $\Gamma_{\text{ext}} \vDash M \triangleright \text{Buff} \lesssim_{\text{beh}}^n M \triangleright \text{eBuff}$. By choosing the set of inductively defined contexts R_n where:¹⁴

$$R_0 \triangleq \text{nil} \qquad R_{n+1} \triangleq \text{in!}v.\text{out?}x.R_n$$

we can argue by analysing the reduction graph of the respective systems that, for any $n \geq 0$:

$$\Gamma_{\text{ext}} \vDash M \triangleright (\text{Buff} \parallel R_{n+1}) \not\lesssim_{\text{beh}}^n M \triangleright (\text{eBuff} \parallel R_{n+1})$$

since it violates the Cost Improving property of Definition 5.5.

Another way how to prove (6.2) is by exploiting completeness of our bisimulation proof technique *wrt.* our behavioural preorder, Theorem 5.18, and work at the level of the transition system of Section 4 showing that, for all $n \geq 0$, the following holds:

$$\Gamma_{\text{ext}} \vDash M \triangleright \text{Buff} \not\mathcal{F}_{\text{bis}}^n M \triangleright \text{eBuff} \quad (6.3)$$

We prove the above statement as Theorem 6.3 of Section 6.1.

Property (6.1), *prima facie*, seems even harder to prove than (6.2), because we are required to show that Barb Preservation and Cost Improving hold under every possible valid context interacting with the two buffer implementations. Once again, we use the transition system of Section 4 and show instead that:

$$\Gamma_{\text{ext}} \vDash M \triangleright \text{eBuff} \sqsubseteq_{\text{bis}}^0 M \triangleright \text{Buff} \quad (6.4)$$

The required result then follows from Theorem 5.10. The proof for this statement is presented in Section 6.2.

In order to make the presentation of these proofs more manageable, we define the following macro definitions for sub-processes making up the derivatives of $\Gamma_{\text{ext}} \triangleleft M \triangleright \text{Buff}$ and $\Gamma_{\text{ext}} \triangleleft M \triangleright \text{eBuff}$.

$$\begin{aligned} \text{Frn}' &\stackrel{\text{def}}{=} b?x.\text{in?}y.\text{alloc } z.(\text{Frn} \parallel b!z \parallel x!(y, z)) & \text{Bck}' &\stackrel{\text{def}}{=} d?x.x?(y, z).\text{out!}y.(\text{Bck} \parallel d!z) \\ \text{Frn}''(x) &\stackrel{\text{def}}{=} \text{in?}y.\text{alloc } z.(\text{Frn} \parallel b!z \parallel x!(y, z)) & \text{Bck}''(x) &\stackrel{\text{def}}{=} x?(y, z).\text{out!}y.(\text{Bck} \parallel d!z) \\ \text{Frn}'''(x, y) &\stackrel{\text{def}}{=} \text{alloc } z.(\text{Frn} \parallel b!z \parallel x!(y, z)) & \text{Bck}'''(y, z) &\stackrel{\text{def}}{=} \text{out!}y.(\text{Bck} \parallel d!z) \\ \text{eBk}' &\stackrel{\text{def}}{=} d?x.x?(y, z).\text{free } x.\text{out!}y.(\text{eBk} \parallel d!z) & \text{eBk}''(x) &\stackrel{\text{def}}{=} x?(y, z).\text{free } x.\text{out!}y.(\text{eBk} \parallel d!z) \\ \text{eBk}'''(x, y, z) &\stackrel{\text{def}}{=} \text{free } x.\text{out!}y.(\text{eBk} \parallel d!z) & \text{eBk}''''(y, z) &\stackrel{\text{def}}{=} \text{out!}y.(\text{eBk} \parallel d!z) \end{aligned}$$

We can thus express the definitions for Buff and eBuff as:

$$\text{Buff} \stackrel{\text{def}}{=} \text{Frn}''(c_1) \parallel \text{Bck}''(c_1) \qquad \text{eBuff} \stackrel{\text{def}}{=} \text{Frn}''(c_1) \parallel \text{eBk}''(c_1) \quad (6.5)$$

¹⁴Note that $\Gamma_{\text{ext}} \vdash R_n$ for any n .




Figure 6: Transition graph for $\Gamma_{\text{ext}} \triangleleft M \triangleright \text{Buff}$ restricted to $\xrightarrow{\text{in?v} \cdot \text{out!v}}$

6.1. Proving Strict Inefficiency. In order to prove (6.3), we do not need to explore the entire state space for $\Gamma_{\text{ext}} \triangleleft M \triangleright \text{Buff}$ and $\Gamma_{\text{ext}} \triangleleft M \triangleright \text{eBuff}$. Instead, it suffices to limit external interactions with the observer to traces of the form $(\xrightarrow{\text{in?v} \cdot \text{out!v}})^*$, which simulate interactions with the observing processes R_n discussed in Section 6. It is instructive to visualise the transition graphs for both $\Gamma_{\text{ext}} \triangleleft M \triangleright \text{Buff}$ and $\Gamma_{\text{ext}} \triangleleft M \triangleright \text{eBuff}$ for a single iteration $\xrightarrow{\text{in?v} \cdot \text{out!v}}$ as depicted in Figure 6 and Figure 7: due to lack of space, the nodes in these graphs abstract away from the environment Γ_{ext} and appropriate resource environments M, N, \dots containing internal channels c_1, c_2, \dots as required.¹⁵ For instance the first node of the graph in Figure 6, $\text{Frn}''(c_1) \parallel \text{Bck}''(c_1)$, i.e., Buff , stands for $\Gamma_{\text{ext}} \triangleleft M \triangleright (\text{Frn}''(c_1) \parallel \text{Bck}''(c_1))$, where $c_1 \in M$, whereas the third node in the same graph, $\text{Frn} \parallel b!c_2 \parallel c_1!(v, c_2) \parallel \text{Bck}''(c_1)$, stands for $\Gamma_{\text{ext}} \triangleleft N \triangleright (\text{Frn} \parallel b!c_2 \parallel c_1!(v, c_2) \parallel \text{Bck}''(c_1))$, where $c_1, c_2 \in N$.

For instance, the graph in Figure 6 shows that after the input action and the channel allocation for c_2 τ -action (with a cost of +1) the inefficient buffer implementation reaches a state where it can perform a number of internal transitions: either the subcomponent Frn may take a recursion unfold step (the first right τ -action) followed by an input on channel b that instantiates the continuation with channel c_2 (the second right τ -action), or else the subcomponent $\text{Bck}''(c_1)$ reads from the head of the buffer $c_1!(v, c_2)$ (the first downwards τ -action). These τ -actions may be interleaved, but no other silent transitions are possible until an output action is performed, after which the backend subcomponent can perform an unfold τ -action (the first downwards τ -action following action out!v) followed by an instantiation communication on channel d (the first downwards τ -action following action out!v). When all of these actions are completed we reach again the starting process, instantiated with channel c_2 instead. The transitions in Figure 7 are analogous, but include a deallocation transition with a cost of -1 .

¹⁵The transition graph also abstracts away from environment moves.




Figure 7: Transition graph for $\Gamma_{\text{ext}} \triangleleft M \triangleright \text{eBuff}$ restricted to $\xrightarrow{\text{in?v} \cdot \text{out!v}}$

Theorem 6.3, which proves (6.3), relies on two lemmas. The main one is Lemma 6.2, which establishes that a number of derivatives from the configurations $\Gamma_{\text{ext}} \triangleleft M \triangleright \text{Buff}$ and $\Gamma_{\text{ext}} \triangleleft M \triangleright \text{eBuff}$ cannot be related for *any* amortisation credit. This Lemma, in turn, relies on Lemma 6.1, which establishes that, for a particular amortisation credit n , if some pair of derivatives of the configurations $\Gamma_{\text{ext}} \triangleleft M \triangleright \text{Buff}$ and $\Gamma_{\text{ext}} \triangleleft M \triangleright \text{eBuff}$ *resp.* cannot be related, then other pairs of derivatives cannot be related either. Lemma 6.1 is used again by Theorem 6.3 to derive that, from the unrelated pairs identified by Lemma 6.2, the required pair of configurations $\Gamma_{\text{ext}} \triangleleft M \triangleright \text{Buff}$ and $\Gamma_{\text{ext}} \triangleleft M \triangleright \text{eBuff}$ cannot be related for any amortisation credit. Upon first reading, the reader who is only interested in the eventual result may safely skip to the statement of Theorem 6.3 and treat Lemma 6.2 and Lemma 6.1 as black-boxes.

In order to be able to state Lemma 6.1 and Lemma 6.2 more succinctly, we find it convenient to delineate groups of processes relating to derivatives of Buff and eBuff . For instance, we can partition the processes depicted in the transition graph of Figure 7 (derivatives of eBuff) into three

sets:

$$\begin{aligned} \text{Prc}_A &\stackrel{\text{def}}{=} \left\{ \begin{array}{l} (\text{Frn} \| b!c_2 \| c_1!(v, c_2) \| \text{eBk}''(c_1)), \\ (\text{Frn}' \| b!c_2 \| c_1!(v, c_2) \| \text{eBk}''(c_1)), \\ (\text{Frn}''(c_2) \| c_1!(v, c_2) \| \text{eBk}''(c_1)), (\text{Frn} \| b!c_2 \| \text{eBk}'''(c_1, v, c_2)), \\ (\text{Frn}' \| b!c_2 \| \text{eBk}'''(c_1, v, c_2)), (\text{Frn}''(c_2) \| \text{eBk}'''(c_1, v, c_2)) \end{array} \middle| \begin{array}{l} c_1 \neq c_2 \in \\ \text{CHAN} \setminus \{\text{in}, \text{out}, b, d\} \end{array} \right\} \\ \text{Prc}_B &\stackrel{\text{def}}{=} \left\{ \begin{array}{l} (\text{Frn} \| b!c_2 \| \text{eBk}'''(v, c_2)), (\text{Frn}' \| b!c_2 \| \text{eBk}'''(v, c_2)), \\ (\text{Frn}''(c_2) \| \text{eBk}'''(v, c_2)) \end{array} \middle| \begin{array}{l} c_2 \in \text{CHAN} \setminus \{\text{in}, \text{out}, b, d\} \end{array} \right\} \\ \text{Prc}_C &\stackrel{\text{def}}{=} \left\{ \begin{array}{l} (\text{Frn} \| b!c_2 \| \text{eBk} \| d!c_2), (\text{Frn}' \| b!c_2 \| \text{eBk} \| d!c_2), \\ (\text{Frn}''(c_2) \| \text{eBk} \| d!c_2), (\text{Frn} \| b!c_2 \| \text{eBk}' \| d!c_2), \\ (\text{Frn}' \| b!c_2 \| \text{eBk}' \| d!c_2), (\text{Frn}''(c_2) \| \text{eBk}' \| d!c_2), \\ (\text{Frn} \| b!c_2 \| \text{eBk}''(c_2)), (\text{Frn}' \| b!c_2 \| \text{eBk}''(c_2)), \\ (\text{Frn}''(c_2) \| \text{eBk}''(c_2)) \end{array} \middle| \begin{array}{l} c_2 \in \text{CHAN} \setminus \{\text{in}, \text{out}, b, d\} \end{array} \right\} \end{aligned}$$

With respect to the transition graph of Figure 7, Prc_A groups the processes *after the allocation* of an (arbitrary) internal channel c_2 but not before any deallocation, *i.e.*, the second and third rows of the graph. The set Prc_B groups the processes *after the deallocation* of the (arbitrary) internal channel c_1 , *i.e.*, the fourth row of the graph. Finally, the set Prc_C groups processes *after the output action* $\text{out}!v$ is performed (before an input action is performed), *i.e.*, the last three rows of the graph.

Lemma 6.1 (Related Negative Results).

- (1) *For any amortisation credit n and appropriate M, N , whenever:*
 - $\Gamma_{\text{ext}} \vDash M \triangleright \text{Frn}'''(c_1, v) \| \text{Bck}''(c_1) \not\sim_{\text{bis}}^n N \triangleright \text{Frn}'''(c'_1, v) \| \text{eBk}''(c'_1)$
 - *For any $Q \in \text{Prc}_A$ we have $\Gamma_{\text{ext}} \vDash M \triangleright \text{Frn}'''(c_1, v) \| \text{Bck}''(c_1) \not\sim_{\text{bis}}^{n+1} N \triangleright Q$*
 - *For any $Q \in \text{Prc}_B$ we have $\Gamma_{\text{ext}} \vDash M \triangleright \text{Frn}'''(c_1, v) \| \text{Bck}''(c_1) \not\sim_{\text{bis}}^n N \triangleright Q$**then, for any $P \in \text{Prc}_C$, we have $\Gamma_{\text{ext}} \vDash M \triangleright \text{Frn}''(c_1) \| \text{Bck}''(c_1) \not\sim_{\text{bis}}^n N \triangleright P$.*
- (2) *For any amortisation credit n and appropriate M, N , and for any $Q \in \text{Prc}_C$:*
 - (a) $\Gamma_{\text{ext}} \vDash M \triangleright \text{Frn}''(c_1) \| \text{Bck}''(c_1) \not\sim_{\text{bis}}^n N \triangleright Q$ *implies*
for any $P \in \text{Prc}_C$ $\Gamma_{\text{ext}} \vDash M \triangleright \text{Frn}' \| b!c_1 \| \text{Bck}''(c_1) \not\sim_{\text{bis}}^n N \triangleright P$
 - (b) $\Gamma_{\text{ext}} \vDash M \triangleright \text{Frn}' \| b!c_1 \| \text{Bck}''(c_1) \not\sim_{\text{bis}}^n N \triangleright Q$ *implies*
for any $P \in \text{Prc}_C$ $\Gamma_{\text{ext}} \vDash M \triangleright \text{Frn} \| b!c_1 \| \text{Bck}''(c_1) \not\sim_{\text{bis}}^n N \triangleright P$
 - (c) $\Gamma_{\text{ext}} \vDash M \triangleright \text{Frn} \| b!c_1 \| \text{Bck}''(c_1) \not\sim_{\text{bis}}^n N \triangleright Q$ *implies*
for any $P \in \text{Prc}_C$ $\Gamma_{\text{ext}} \vDash M \triangleright \text{Frn} \| b!c_1 \| \text{Bck}' \| d!c_1 \not\sim_{\text{bis}}^n N \triangleright P$
 - (d) $\Gamma_{\text{ext}} \vDash M \triangleright \text{Frn} \| b!c_1 \| \text{Bck}' \| d!c_1 \not\sim_{\text{bis}}^n N \triangleright Q$ *implies*
for any $P \in \text{Prc}_C$ $\Gamma_{\text{ext}} \vDash M \triangleright \text{Frn} \| b!c_1 \| \text{Bck} \| d!c_1 \not\sim_{\text{bis}}^n N \triangleright P$
- (3) *For any amortisation credit n and appropriate M, N , and for any $R \in \text{Prc}_B$, $Q \in \text{Prc}_C$:*
 - (a) $\Gamma_{\text{ext}} \vDash M \triangleright \text{Frn} \| b!c_2 \| \text{Bck} \| d!c_2 \not\sim_{\text{bis}}^n N \triangleright Q$ *implies*
for any $P \in \text{Prc}_B$ $\Gamma_{\text{ext}} \vDash M \triangleright \text{Frn} \| b!c_2 \| \text{Bck}'''(v, c_2) \not\sim_{\text{bis}}^n N \triangleright P$
 - (b) $\Gamma_{\text{ext}} \vDash M \triangleright \text{Frn} \| b!c_2 \| \text{Bck}'''(v, c_2) \not\sim_{\text{bis}}^n N \triangleright R$ *implies*
for any $P \in \text{Prc}_B$ $\Gamma_{\text{ext}} \vDash M \triangleright \text{Frn} \| b!c_2 \| c_1!(v, c_2) \| \text{Bck}''(c_1) \not\sim_{\text{bis}}^n N \triangleright P$
- (4) *For any amortisation credit n and appropriate M, N , and for any $Q \in \text{Prc}_C$:*
 - (a) $\Gamma_{\text{ext}} \vDash M \triangleright \text{Frn} \| b!c_1 \| \text{Bck} \| d!c_1 \not\sim_{\text{bis}}^n N \triangleright Q$ *implies*
for any $P \in \text{Prc}_A$ $\Gamma_{\text{ext}} \vDash M \triangleright \text{Frn} \| b!c_1 \| \text{Bck}'''(v, c_1) \not\sim_{\text{bis}}^{n+1} N \triangleright P$
 - (b) $\Gamma_{\text{ext}} \vDash M \triangleright \text{Frn} \| b!c_1 \| \text{Bck} \| d!c_1 \not\sim_{\text{bis}}^n N \triangleright Q$ *implies*
 $\Gamma_{\text{ext}} \vDash M \triangleright \text{Frn} \| b!c_1 \| \text{Bck}'''(v, c_1) \not\sim_{\text{bis}}^n N \triangleright \text{Frn}'''(c'_1, v) \| \text{eBk}''(c'_1)$
 - (c) $\Gamma_{\text{ext}} \vDash M \triangleright \text{Frn} \| b!c_1 \| \text{Bck} \| d!c_1 \not\sim_{\text{bis}}^n N \triangleright Q$ *implies*
 $\Gamma_{\text{ext}} \vDash M \triangleright \text{Frn} \| b!c_2 \| c_1!(v, c_2) \| \text{Bck}''(c_1) \not\sim_{\text{bis}}^n N \triangleright \text{Frn}'''(c'_1, v) \| \text{eBk}''(c'_1)$

Proof. Each case is proved by contradiction:

- (1) Assume the premises together with the inverse of the conclusion, *i.e.*,

$$\Gamma_{\text{ext}} \vDash M \triangleright \text{Frn}''(c_1) \parallel \text{Bck}''(c_1) \stackrel{n}{\sim}_{\text{bis}} N \triangleright P.$$

Consider the transition from the left-hand configuration:

$$\Gamma_{\text{ext}} \triangleleft M \triangleright \text{Frn}''(c_1) \parallel \text{Bck}''(c_1) \xrightarrow{\text{in}^?v}_0 \Gamma_{\text{ext}} \triangleleft M \triangleright \text{Frn}'''(c_1, v) \parallel \text{Bck}''(c_1).$$

For any $P \in \text{PrC}_C$, this can only be matched by the right-hand configuration, $\Gamma_{\text{ext}} \triangleleft N \triangleright P$, through either of the following cases:

- (a) $\Gamma_{\text{ext}} \triangleleft N \triangleright P \xrightarrow{\text{in}^?v}_0 \Gamma_{\text{ext}} \triangleleft N \triangleright \text{Frn}'''(c'_1, v) \parallel \text{eBk}''(c'_1)$, *i.e.*, a weak input action without trailing τ -moves after the external action $\text{in}^?v$ — see first row of the graph in Figure 7. But we know $\Gamma_{\text{ext}} \vDash M \triangleright \text{Frn}'''(c_1, v) \parallel \text{Bck}''(c_1) \not\stackrel{n}{\sim}_{\text{bis}} N \triangleright \text{Frn}'''(c'_1, v) \parallel \text{eBk}''(c'_1)$ from the first premise.
- (b) $\Gamma_{\text{ext}} \triangleleft N \triangleright P \xrightarrow{\text{in}^?v}_{+1} \Gamma_{\text{ext}} \triangleleft N \triangleright Q$ for some $Q \in \text{PrC}_A$. However, from the second premise we know that $\Gamma_{\text{ext}} \vDash M \triangleright \text{Frn}'''(c_1, v) \parallel \text{Bck}''(c_1) \not\stackrel{n+1}{\sim}_{\text{bis}} N \triangleright Q$
- (c) $\Gamma_{\text{ext}} \triangleleft N \triangleright P \xrightarrow{\text{in}^?v}_0 \Gamma_{\text{ext}} \triangleleft N \triangleright Q$ for some $Q \in \text{PrC}_B$. Again, from the third premise we know that $\Gamma_{\text{ext}} \vDash M \triangleright \text{Frn}'''(c_1, v) \parallel \text{Bck}''(c_1) \not\stackrel{n}{\sim}_{\text{bis}} N \triangleright Q$

Since $\Gamma_{\text{ext}} \triangleleft N \triangleright P$ cannot perform a matching move, we obtain a contradiction.

- (2) We here prove case (a). The other cases are analogous.

Assume $\Gamma_{\text{ext}} \vDash M \triangleright \text{Frn}' \parallel b!c_1 \parallel \text{Bck}''(c_1) \stackrel{n}{\sim}_{\text{bis}} N \triangleright P$ and consider the action

$$\Gamma_{\text{ext}} \triangleleft M \triangleright \text{Frn}' \parallel b!c_1 \parallel \text{Bck}''(c_1) \xrightarrow{\tau}_0 \Gamma_{\text{ext}} \triangleleft M \triangleright \text{Frn}''(c_1) \parallel \text{Bck}''(c_1).$$

For our assumption to hold, $\Gamma_{\text{ext}} \triangleleft N \triangleright P$ would need to match this move by a (weak) silent action leading to a configuration that can match $\Gamma_{\text{ext}} \triangleleft M \triangleright \text{Frn}''(c_1) \parallel \text{Bck}''(c_1)$. The only matching move can be

$$\Gamma_{\text{ext}} \triangleleft N \triangleright P \implies_0 \Gamma_{\text{ext}} \triangleleft N \triangleright Q \quad \text{for some } Q \in \text{PrC}_C.$$

However, from our premise we know $\Gamma_{\text{ext}} \vDash M \triangleright \text{Frn}''(c_1) \parallel \text{Bck}''(c_1) \not\stackrel{n}{\sim}_{\text{bis}} N \triangleright Q'$ for any amortisation credit n and $Q' \in \text{PrC}_C$ and therefore conclude that the move cannot be matched, thereby obtaining a contradiction.

- (3) We here prove case (a). Case (b) is analogous.

Assume $\Gamma_{\text{ext}} \vDash M \triangleright \text{Frn} \parallel b!c_2 \parallel \text{Bck}'''(v, c_2) \stackrel{n}{\sim}_{\text{bis}} N \triangleright P$ and consider the action

$$\Gamma_{\text{ext}} \triangleleft M \triangleright \text{Frn} \parallel b!c_2 \parallel \text{Bck}'''(v, c_2) \xrightarrow{\text{out}!v}_0 \Gamma_{\text{ext}} \triangleleft M \triangleright \text{Frn} \parallel b!c_2 \parallel \text{Bck} \parallel d!c_2$$

This action can only be matched by a transition of the form

$$\Gamma_{\text{ext}} \triangleleft N \triangleright P \xrightarrow{\text{out}!v}_0 \Gamma_{\text{ext}} \triangleleft N \triangleright Q \quad \text{for some } Q \in \text{PrC}_C.$$

However, from our premise we know $\Gamma_{\text{ext}} \vDash M \triangleright \text{Frn} \parallel b!c_2 \parallel \text{Bck} \parallel d!c_2 \not\stackrel{n}{\sim}_{\text{bis}} N \triangleright Q$ for any amortisation credit n and $Q \in \text{PrC}_C$. Thus we conclude that the move cannot be matched, thereby obtaining a contradiction.

- (4) Cases (a) and (b) are analogous to 3(a) and 3(b). We here outline the proof for case (c).

First we note that from the premise $\Gamma_{\text{ext}} \vDash M \triangleright \text{Frn} \parallel b!c_1 \parallel \text{Bck} \parallel d!c_1 \stackrel{n}{\mathcal{F}}_{\text{bis}} N \triangleright Q$ (for any $Q \in \text{Prc}_C$) and Lemma 6.1.3(a), Lemma 6.1.4(a) and Lemma 6.1.4(b) *resp.* we obtain:

$$\Gamma_{\text{ext}} \vDash M \triangleright \text{Frn} \parallel b!c_2 \parallel \text{Bck}'''(v, c_2) \stackrel{n}{\mathcal{F}}_{\text{bis}} N \triangleright P \quad \text{for any } P \in \text{Prc}_B \quad (6.6)$$

$$\Gamma_{\text{ext}} \vDash M \triangleright \text{Frn} \parallel b!c_1 \parallel \text{Bck}'''(v, c_1) \stackrel{n+1}{\mathcal{F}}_{\text{bis}} N \triangleright P \quad \text{for any } P \in \text{Prc}_A \quad (6.7)$$

$$\Gamma_{\text{ext}} \vDash M \triangleright \text{Frn} \parallel b!c_1 \parallel \text{Bck}'''(v, c_1) \stackrel{n}{\mathcal{F}}_{\text{bis}} N \triangleright \text{Frn}'''(c'_1, v) \parallel \text{eBk}''(c'_1) \quad (6.8)$$

We assume $\Gamma_{\text{ext}} \vDash M \triangleright \text{Frn} \parallel b!c_2 \parallel c_1!(v, c_2) \parallel \text{Bck}''(c_1) \stackrel{n}{\mathcal{F}}_{\text{bis}} N \triangleright \text{Frn}'''(c'_1, v) \parallel \text{eBk}''(c'_1)$ and then showing that this leads to a contradiction. Consider the move

$$\Gamma_{\text{ext}} \triangleleft M \triangleright \text{Frn} \parallel b!c_2 \parallel c_1!(v, c_2) \parallel \text{Bck}''(c_1) \xrightarrow{\tau}_0 \Gamma_{\text{ext}} \triangleleft M \triangleright \text{Frn} \parallel b!c_2 \parallel \text{Bck}'''(v, c_2)$$

This can be matched by $\Gamma_{\text{ext}} \triangleleft N \triangleright \text{Frn}'''(c'_1, v) \parallel \text{eBk}''(c'_1)$ using either of the following moves:

- $\Gamma_{\text{ext}} \triangleleft N \triangleright \text{Frn}'''(c'_1, v) \parallel \text{eBk}''(c'_1) \Rightarrow_0 \Gamma_{\text{ext}} \triangleleft N \triangleright \text{Frn}'''(c'_1, v) \parallel \text{eBk}''(c'_1)$. But (6.8) prohibits this from being the matching move.
- $\Gamma_{\text{ext}} \triangleleft N \triangleright \text{Frn}'''(c'_1, v) \parallel \text{eBk}''(c'_1) \Rightarrow_{+1} \Gamma_{\text{ext}} \triangleleft N \triangleright Q$ for some $Q \in \text{Prc}_A$. But (6.7) prohibits this from being the matching move.
- $\Gamma_{\text{ext}} \triangleleft N \triangleright \text{Frn}'''(c'_1, v) \parallel \text{eBk}''(c'_1) \Rightarrow_0 \Gamma_{\text{ext}} \triangleleft N \triangleright Q$ for some $Q \in \text{Prc}_B$. But (6.6) prohibits this from being the matching move.

This contradicts our earlier assumption. \square

Lemma 6.2. *For all $n \in \mathbb{N}_{\text{AT}}$ and appropriate M, N :*

- (1) *For any $Q \in \text{Prc}_A$ we have $\Gamma_{\text{ext}} \vDash M \triangleright \text{Frn} \parallel b!c_2 \parallel \text{Bck}'''(v, c_2) \stackrel{n}{\mathcal{F}}_{\text{bis}} N \triangleright Q$*
- (2) *For any $Q \in \text{Prc}_A$ we have $\Gamma_{\text{ext}} \vDash M \triangleright \text{Frn} \parallel b!c_2 \parallel c_1!(v, c_2) \parallel \text{Bck}''(c_1) \stackrel{n}{\mathcal{F}}_{\text{bis}} N \triangleright Q$*
- (3) *For any $Q \in \text{Prc}_A$ we have $\Gamma_{\text{ext}} \vDash M \triangleright \text{Frn}'''(c_1, v) \parallel \text{Bck}''(c_1) \stackrel{n+1}{\mathcal{F}}_{\text{bis}} N \triangleright Q$*
- (4) *For any $Q \in \text{Prc}_B$ we have $\Gamma_{\text{ext}} \vDash M \triangleright \text{Frn}'''(c_1, v) \parallel \text{Bck}''(c_1) \stackrel{n}{\mathcal{F}}_{\text{bis}} N \triangleright Q$*
- (5) *$\Gamma_{\text{ext}} \vDash M \triangleright \text{Frn}'''(c_1, v) \parallel \text{Bck}''(c_1) \stackrel{n}{\mathcal{F}}_{\text{bis}} N \triangleright \text{Frn}'''(c'_1, v) \parallel \text{eBk}''(c'_1)$*

Proof. We prove statements (1) to (5) simultaneously, by induction on n .

$n = 0$: We prove each clause by contradiction:

- (1) Assume $\Gamma_{\text{ext}} \vDash M \triangleright \text{Frn} \parallel b!c_2 \parallel \text{Bck}'''(v, c_2) \stackrel{0}{\mathcal{F}}_{\text{bis}} N \triangleright Q$ for some $Q \in \text{Prc}_A$ and consider the transition

$$\Gamma_{\text{ext}} \triangleleft M \triangleright \text{Frn} \parallel b!c_2 \parallel \text{Bck}'''(v, c_2) \xrightarrow{\text{out}!v}_0 \Gamma_{\text{ext}} \triangleleft M \triangleright \text{Frn} \parallel b!c_2 \parallel \text{Bck} \parallel d!c_2$$

For any $Q \in \text{Prc}_A$, this cannot be matched by any move from $\Gamma_{\text{ext}} \triangleleft N \triangleright Q$ since output actions must be preceded by a channel deallocation, which incurs a *negative* cost — see second and third rows of the graph in Figure 7. Stated otherwise, every matching move can only be of the form

$$\Gamma_{\text{ext}} \triangleleft N \triangleright Q \xrightarrow{\text{out}!v}_{-1} \Gamma_{\text{ext}} \triangleleft N' \triangleright Q'$$

where $N = (N', c'_1)$ for some c'_1 and $Q' \in \text{Prc}_C$. However, since the amortisation credit cannot be negative, we can never have $\Gamma_{\text{ext}} \vDash M \triangleright \text{Frn} \parallel b!c_2 \parallel \text{Bck} \parallel d!c_2 \stackrel{-1}{\mathcal{F}}_{\text{bis}} N' \triangleright Q'$. We therefore obtain a contradiction.

- (2) Assume $\Gamma_{\text{ext}} \vDash M \triangleright \text{Frn} \parallel b!c_2 \parallel c_1!(v, c_2) \parallel \text{Bck}''(c_1) \stackrel{0}{\mathcal{F}}_{\text{bis}} N \triangleright Q$ for some $Q \in \text{Prc}_A$ and consider the transition

$$\Gamma_{\text{ext}} \triangleleft M \triangleright \text{Frn} \parallel b!c_2 \parallel c_1!(v, c_2) \parallel \text{Bck}''(c_1) \xrightarrow{\tau}_0 \Gamma_{\text{ext}} \triangleleft M \triangleright \text{Frn} \parallel b!c_2 \parallel \text{Bck}'''(v, c_2)$$

Since the amortisation credit can never be negative, the matching move can only be of the form

$$\Gamma_{\text{ext}} \triangleleft N \triangleright Q \Rightarrow_0 \Gamma_{\text{ext}} \triangleleft N \triangleright Q'$$

for some $Q' \in \text{Prc}_A$. But then we get a contradiction since, from the previous clause, we know that $\Gamma_{\text{ext}} \vDash M \triangleright \text{Frn} \parallel b!c_2 \parallel \text{Bck}''(v, c_2) \not\approx_{\text{bis}}^0 N \triangleright Q'$.

- (3) Assume $\Gamma_{\text{ext}} \vDash M \triangleright \text{Frn}''(c_1, v) \parallel \text{Bck}''(c_1) \not\approx_{\text{bis}}^1 N \triangleright Q$ for some $Q \in \text{Prc}_A$ and consider the transition

$$\Gamma_{\text{ext}} \triangleleft M \triangleright \text{Frn}''(c_1, v) \parallel \text{Bck}''(c_1) \xrightarrow{\tau}_{+1} \Gamma_{\text{ext}} \triangleleft M, c_2 \triangleright \text{Frn} \parallel b!c_2 \parallel c_1!(v, c_2) \parallel \text{Bck}''(c_1)$$

for some newly allocated channel c_2 . As in the previous case, since the amortisation credit can never be negative, the matching move can only be of the form

$$\Gamma_{\text{ext}} \triangleleft N \triangleright Q \Rightarrow_0 \Gamma_{\text{ext}} \triangleleft N \triangleright Q'$$

for some $Q' \in \text{Prc}_A$. But then we get a contradiction since, from the previous clause, we know that $\Gamma_{\text{ext}} \vDash (M, c_2) \triangleright \text{Frn} \parallel b!c_2 \parallel c_1!(v, c_2) \parallel \text{Bck}''(c_1) \not\approx_{\text{bis}}^0 N \triangleright Q'$.

- (4) Analogous to the previous case.
 (5) Assume $\Gamma_{\text{ext}} \vDash M \triangleright \text{Frn}''(c_1, v) \parallel \text{Bck}''(c_1) \not\approx_{\text{bis}}^0 N \triangleright \text{Frn}''(c'_1, v) \parallel \text{eBk}''(c'_1)$ and consider the transition

$$\Gamma_{\text{ext}} \triangleleft M \triangleright \text{Frn}''(c_1, v) \parallel \text{Bck}''(c_1) \xrightarrow{\tau}_{+1} \Gamma_{\text{ext}} \triangleleft M, c_2 \triangleright \text{Frn} \parallel b!c_2 \parallel c_1!(v, c_2) \parallel \text{Bck}''(c_1)$$

Since the transition incurred a cost of +1 and the current amortisation credit is 0, the matching weak transition must also incur a cost of +1 and thus $\Gamma_{\text{ext}} \triangleleft N \triangleright \text{Frn}''(c'_1, v) \parallel \text{eBk}''(c'_1)$ can only match this by the move

$$\Gamma_{\text{ext}} \triangleleft N \triangleright \text{Frn}''(c'_1, v) \parallel \text{eBk}''(c'_1) \xRightarrow{\tau}_{+1} \Gamma_{\text{ext}} \triangleleft N, c'_2 \triangleright Q$$

for some $Q \in \text{Prc}_A$. But then we still get a contradiction since, from clause (2), we know $\Gamma_{\text{ext}} \vDash M \triangleright \text{Frn} \parallel b!c_2 \parallel c_1!(v, c_2) \parallel \text{Bck}''(c_1) \not\approx_{\text{bis}}^0 N \triangleright Q$.

$n = k + 1$: We prove each clause by contradiction. However before we tackle each individual clause, we note that from clauses (3), (4) and (5) of the I.H. we know

$$\text{For any } Q \in \text{Prc}_A \text{ we have } \Gamma_{\text{ext}} \vDash M \triangleright \text{Frn}''(c_1, v) \parallel \text{Bck}''(c_1) \not\approx_{\text{bis}}^{k+1} N \triangleright Q$$

$$\text{For any } Q \in \text{Prc}_B \text{ we have } \Gamma_{\text{ext}} \vDash M \triangleright \text{Frn}''(c_1, v) \parallel \text{Bck}''(c_1) \not\approx_{\text{bis}}^k N \triangleright Q$$

$$\Gamma_{\text{ext}} \vDash M \triangleright \text{Frn}''(c_1, v) \parallel \text{Bck}''(c_1) \not\approx_{\text{bis}}^k N \triangleright \text{Frn}''(c_1, v) \parallel \text{eBk}''(c_1)$$

By Lemma 6.1.1 we obtain, for any $Q' \in \text{Prc}_C$ and appropriate N' :

$$\Gamma_{\text{ext}} \vDash M \triangleright \text{Frn}''(c_1) \parallel \text{Bck}''(c_1) \not\approx_{\text{bis}}^k N' \triangleright Q'$$

and by Lemma 6.1.2(a), Lemma 6.1.2(b), Lemma 6.1.2(c) and Lemma 6.1.2(d) we obtain, for any $Q' \in \text{Prc}_C$ and appropriate N' :

$$\Gamma_{\text{ext}} \vDash M \triangleright \text{Frn} \parallel b!c_2 \parallel \text{Bck} \parallel d!c_2 \not\approx_{\text{bis}}^k N \triangleright Q' \quad (6.9)$$

Also, by (6.9), Lemma 6.1.3(a) and Lemma 6.1.3(b) we obtain, for any $Q'' \in \text{Prc}_B$:

$$\Gamma_{\text{ext}} \vDash M \triangleright \text{Frn} \parallel b!c_2 \parallel \text{Bck}''(v, c_2) \not\approx_{\text{bis}}^k N \triangleright Q'' \quad (6.10)$$

$$\Gamma_{\text{ext}} \vDash M \triangleright \text{Frn} \parallel b!c_2 \parallel c_1!(v, c_2) \parallel \text{Bck}''(c_1) \not\approx_{\text{bis}}^k N \triangleright Q'' \quad (6.11)$$

Moreover, by (6.9), Lemma 6.1.4(a), Lemma 6.1.4(b) and Lemma 6.1.4(c) we obtain:

$$\Gamma_{\text{ext}} \vDash M \triangleright \text{Frn} \parallel b!c_2 \parallel c_1!(v, c_2) \parallel \text{Bck}''(c_1) \not\approx_{\text{bis}}^k N \triangleright \text{Frn}''(c'_1, v) \parallel \text{eBk}''(c'_1) \quad (6.12)$$

The proofs for each clause are as follows:

- (1) Assume $\Gamma_{\text{ext}} \vDash M \triangleright \text{Frn} \parallel b!c_2 \parallel \text{Bck}'''(v, c_2) \stackrel{\sqsubseteq_{\text{bis}}^{k+1}}{\sim} N \triangleright Q$ for some $Q \in \text{PrC}_A$ and consider the transition

$$\Gamma_{\text{ext}} \triangleleft M \triangleright \text{Frn} \parallel b!c_2 \parallel \text{Bck}'''(v, c_2) \xrightarrow{\text{out}!v}_0 \Gamma_{\text{ext}} \triangleleft M \triangleright \text{Frn} \parallel b!c_2 \parallel \text{Bck} \parallel d!c_2$$

For any $Q \in \text{PrC}_A$, this can (only) be matched by any move of the form

$$\Gamma_{\text{ext}} \triangleleft N \triangleright Q \xrightarrow{\text{out}!v}_{-1} \Gamma_{\text{ext}} \triangleleft N' \triangleright Q'$$

where $N = (N', c'_1)$ for some c'_1 , $Q' \in \text{PrC}_C$, and the external action $\text{out}!v$ is preceded by a τ -move deallocating c'_1 . For our initial assumption to hold we need to show that *at least one* of these configurations $\Gamma_{\text{ext}} \triangleleft N' \triangleright Q'$ satisfies the property

$$\Gamma_{\text{ext}} \vDash M \triangleright \text{Frn} \parallel b!c_2 \parallel \text{Bck} \parallel d!c_2 \stackrel{\sqsubseteq_{\text{bis}}^k}{\sim} N' \triangleright Q'.$$

But by (6.9) we know that no such configuration exists, thereby contradicting our initial assumption.

- (2) Assume $\Gamma_{\text{ext}} \vDash M \triangleright \text{Frn} \parallel b!c_2 \parallel c_1!(v, c_2) \parallel \text{Bck}''(c_1) \stackrel{\sqsubseteq_{\text{bis}}^{k+1}}{\sim} N \triangleright Q$ for some $Q \in \text{PrC}_A$ and consider the transition

$$\Gamma_{\text{ext}} \triangleleft M \triangleright \text{Frn} \parallel b!c_2 \parallel c_1!(v, c_2) \parallel \text{Bck}''(c_1) \xrightarrow{\tau}_0 \Gamma_{\text{ext}} \triangleleft M \triangleright \text{Frn} \parallel b!c_2 \parallel \text{Bck}'''(v, c_2)$$

This transition can be matched by $\Gamma_{\text{ext}} \triangleleft N \triangleright Q$ through either of the following moves:

- (a) $\Gamma_{\text{ext}} \triangleleft N \triangleright Q \Rightarrow_0 \Gamma_{\text{ext}} \triangleleft N \triangleright Q'$ for some $Q' \in \text{PrC}_A$. However, from the previous clause, *i.e.*, clause (1) when $n = k + 1$, we know that this cannot be the matching move since $\Gamma_{\text{ext}} \vDash M \triangleright \text{Frn} \parallel b!c_2 \parallel \text{Bck}'''(v, c_2) \not\stackrel{\sqsubseteq_{\text{bis}}^{k+1}}{\sim} N \triangleright Q'$.
- (b) $\Gamma_{\text{ext}} \triangleleft N \triangleright Q \Rightarrow_{-1} \Gamma_{\text{ext}} \triangleleft N' \triangleright Q'$ for some $Q' \in \text{PrC}_B$ and $N = (N', c'_1)$. However, from (6.10), we know that this cannot be the matching move since $\Gamma_{\text{ext}} \vDash M \triangleright \text{Frn} \parallel b!c_2 \parallel \text{Bck}'''(v, c_2) \not\stackrel{\sqsubseteq_{\text{bis}}^k}{\sim} N' \triangleright Q'$.

Thus, we obtain a contradiction.

- (3) Assume $\Gamma_{\text{ext}} \vDash M \triangleright \text{Frn}'''(c_1, v) \parallel \text{Bck}''(c_1) \stackrel{\sqsubseteq_{\text{bis}}^{k+2}}{\sim} N \triangleright Q$, where $Q \in \text{PrC}_A$, and consider the transition:

$$\Gamma_{\text{ext}} \triangleleft M \triangleright \text{Frn}'''(c_1, v) \parallel \text{Bck}''(c_1) \xrightarrow{\tau}_{+1} \Gamma_{\text{ext}} \triangleleft M, c_2 \triangleright \text{Frn} \parallel b!c_2 \parallel c_1!(v, c_2) \parallel \text{Bck}''(c_1)$$

for some newly allocated channel c_2 . This can be matched by $\Gamma_{\text{ext}} \triangleleft N \triangleright Q$ through either of the following moves:

- (a) $\Gamma_{\text{ext}} \triangleleft N \triangleright Q \Rightarrow_0 \Gamma_{\text{ext}} \triangleleft N \triangleright Q'$ for some $Q' \in \text{PrC}_A$. However, from the previous clause, *i.e.*, clause (2) when $n = k + 1$, we know that this cannot be the matching move since $\Gamma_{\text{ext}} \vDash M, c_2 \triangleright \text{Frn} \parallel b!c_2 \parallel c_1!(v, c_2) \parallel \text{Bck}''(c_1) \not\stackrel{\sqsubseteq_{\text{bis}}^{k+1}}{\sim} N \triangleright Q'$.
- (b) $\Gamma_{\text{ext}} \triangleleft N \triangleright Q \Rightarrow_{-1} \Gamma_{\text{ext}} \triangleleft N' \triangleright Q'$ for some $Q' \in \text{PrC}_B$ and $N = (N', c'_1)$. However, from (6.11), we know that this cannot be the matching move since $\Gamma_{\text{ext}} \vDash M \triangleright \text{Frn} \parallel b!c_2 \parallel c_1!(v, c_2) \parallel \text{Bck}''(c_1) \not\stackrel{\sqsubseteq_{\text{bis}}^k}{\sim} N' \triangleright Q'$.

Thus, we obtain a contradiction.

- (4) Analogous to the proof for the previous clause and relies on (6.11) again.
- (5) Assume $\Gamma_{\text{ext}} \vDash M \triangleright \text{Frn}'''(c_1, v) \parallel \text{Bck}''(c_1) \stackrel{\sqsubseteq_{\text{bis}}^{k+1}}{\sim} N \triangleright \text{Frn}'''(c'_1, v) \parallel \text{eBk}''(c'_1)$ and consider the transition

$$\Gamma_{\text{ext}} \triangleleft M \triangleright \text{Frn}'''(c_1, v) \parallel \text{Bck}''(c_1) \xrightarrow{\tau}_{+1} \Gamma_{\text{ext}} \triangleleft M, c_2 \triangleright \text{Frn} \parallel b!c_2 \parallel c_1!(v, c_2) \parallel \text{Bck}''(c_1)$$

for some newly allocated channel c_2 . This can be matched by the right-hand configuration $\Gamma_{\text{ext}} \triangleleft N \triangleright \text{Frn}'''(c'_1, v) \parallel \text{eBk}''(c'_1)$ through either of the following moves:

- (a) $\Gamma_{\text{ext}} \triangleleft N \triangleright \text{Frn}'''(c'_1, v) \parallel \text{eBk}''(c'_1) \Rightarrow_0 \Gamma_{\text{ext}} \triangleleft N \triangleright \text{Frn}'''(c'_1, v) \parallel \text{eBk}''(c'_1)$, i.e., no transitions. However, from (6.12), this cannot be the matching move since $\Gamma_{\text{ext}} \vDash M \triangleright \text{Frn} \parallel b!c_2 \parallel c_1!(v, c_2) \parallel \text{Bck}''(c_1) \not\stackrel{k}{\mathcal{F}}_{\text{bis}} N \triangleright \text{Frn}'''(c'_1, v) \parallel \text{eBk}''(c'_1)$.
- (b) $\Gamma_{\text{ext}} \triangleleft N \triangleright \text{Frn}'''(c'_1, v) \parallel \text{eBk}''(c'_1) \xrightarrow{\tau}_{+1} \Gamma_{\text{ext}} \triangleleft N, c'_2 \triangleright Q'$ for some $Q' \in \text{Prc}_A$ and $c'_2 \notin N$. However, from clause (2) when $n = k + 1$, this cannot be the matching move since $\Gamma_{\text{ext}} \vDash M, c_2 \triangleright \text{Frn} \parallel b!c_2 \parallel c_1!(v, c_2) \parallel \text{Bck}''(c_1) \not\stackrel{k+1}{\mathcal{F}}_{\text{bis}} N, c'_2 \triangleright Q'$.
- (c) $\Gamma_{\text{ext}} \triangleleft N \triangleright \text{Frn}'''(c'_1, v) \parallel \text{eBk}''(c'_1) \xrightarrow{\tau}_0 \Gamma_{\text{ext}} \triangleleft (N', c'_2) \triangleright Q'$ for some $Q' \in \text{Prc}_B$, $N = (N', c'_1)$ and $c'_2 \notin N$. However, from (6.11), this cannot be the matching move since $\Gamma_{\text{ext}} \vDash M \triangleright \text{Frn} \parallel b!c_2 \parallel c_1!(v, c_2) \parallel \text{Bck}''(c_1) \not\stackrel{k}{\mathcal{F}}_{\text{bis}} (N', c'_2) \triangleright Q'$. \square

Theorem 6.3 (Strict Inefficiency). *For all $n \geq 0$ and appropriate M we have*

$$\Gamma_{\text{ext}} \vDash M \triangleright \text{Buff} \not\stackrel{n}{\mathcal{F}}_{\text{bis}} M \triangleright \text{eBuff}$$

Proof. Since:

$$\text{Buff} \stackrel{\text{def}}{=} \text{Frn}''(c_1) \parallel \text{Bck}''(c_1) \qquad \text{eBuff} \stackrel{\text{def}}{=} \text{Frn}''(c_1) \parallel \text{eBk}''(c_1)$$

we need to show that

$$\Gamma_{\text{ext}} \vDash M \triangleright \text{Frn}''(c_1) \parallel \text{Bck}''(c_1) \not\stackrel{n}{\mathcal{F}}_{\text{bis}} M \triangleright \text{Frn}''(c_1) \parallel \text{eBk}''(c_1)$$

for any arbitrary n . By Lemma 6.2.3, Lemma 6.2.4 and Lemma 6.2.5 we know that for any n :

$$\text{For any } Q \in \text{Prc}_A \text{ we have } \Gamma_{\text{ext}} \vDash M \triangleright \text{Frn}'''(c_1, v) \parallel \text{Bck}''(c_1) \not\stackrel{n+1}{\mathcal{F}}_{\text{bis}} M \triangleright Q \quad (6.13)$$

$$\text{For any } Q \in \text{Prc}_B \text{ we have } \Gamma_{\text{ext}} \vDash M \triangleright \text{Frn}'''(c_1, v) \parallel \text{Bck}''(c_1) \not\stackrel{n}{\mathcal{F}}_{\text{bis}} M \triangleright Q \quad (6.14)$$

$$\Gamma_{\text{ext}} \vDash M \triangleright \text{Frn}'''(c_1, v) \parallel \text{Bck}''(c_1) \not\stackrel{n}{\mathcal{F}}_{\text{bis}} M \triangleright \text{Frn}'''(c_1, v) \parallel \text{eBk}''(c_1) \quad (6.15)$$

Since $(\text{Frn}''(c_1) \parallel \text{eBk}''(c_1)) \in \text{Prc}_C$, by Lemma 6.1.1, (6.13), (6.14) and (6.15) we conclude

$$\Gamma_{\text{ext}} \vDash M \triangleright \text{Frn}''(c_1) \parallel \text{Bck}''(c_1) \not\stackrel{n}{\mathcal{F}}_{\text{bis}} M \triangleright \text{Frn}''(c_1) \parallel \text{eBk}''(c_1)$$

as required. \square

6.2. Proving Relative Efficiency. As opposed to Theorem 6.3, the proof for (6.4) requires us to consider the entire state-space of $\Gamma_{\text{ext}} \triangleleft M \triangleright \text{Buff}$ and $\Gamma_{\text{ext}} \triangleleft M \triangleright \text{eBuff}$. Fortunately, we can apply the compositionality result of Theorem 4.23 to prove (6.1) and focus on a subset of this state-space. More precisely, we recall from (6.5) that

$$\text{Buff} \stackrel{\text{def}}{=} \text{Frn}''(c_1) \parallel \text{Bck}''(c_1) \qquad \text{eBuff} \stackrel{\text{def}}{=} \text{Frn}''(c_1) \parallel \text{eBk}''(c_1)$$

where both buffer implementations share the common sub-process $\text{Frn}''(c_1)$. We also recall from (3.3) that this common sub-process was typed wrt. the type environment

$$\Gamma_{\text{Frn}} = \text{in} : [\mathbf{T}]^\omega, b : [\mathbf{T}_{\text{rec}}]^\omega, c_1 : [\mathbf{T}, \mathbf{T}_{\text{rec}}]^1.$$

Theorem 4.23 thus states that in order to prove (6.1), it suffices to abstract away from this common code and prove Theorem 6.4

Theorem 6.4 (Relative Efficiency). $(\Gamma_{\text{ext}}, \Gamma_{\text{Frn}}) \vDash M \triangleright \text{eBk}''(c_1) \not\stackrel{0}{\mathcal{F}}_{\text{bis}} M \triangleright \text{Bck}''(c_1)$

Proof. We prove $\Gamma_{\text{ext}}, \Gamma_{\text{Frm}} \vDash M \triangleright \text{eBk}''(c_1) \sqsubseteq_{\text{bis}}^0 M \triangleright \text{Bck}''(c_1)$ through the family of relations \mathcal{R} defined below, which includes the required quadruple $\langle (\Gamma_{\text{ext}}, \Gamma_{\text{Frm}}), 0, (M \triangleright \text{eBk}''(c_1)), (M \triangleright \text{Bck}''(c_1)) \rangle$.

$$\mathcal{R} \stackrel{\text{def}}{=} \left\{ \begin{array}{l} \langle (\Gamma, \Delta), n, (M' \triangleright \text{eBk}''(c)), (N' \triangleright \text{Bck}''(c)) \rangle \\ \langle (\Gamma, \Delta), n, (M' \triangleright \text{eBk}'''(c, v, c')), (N' \triangleright \text{Bck}'''(v, c')) \rangle \\ \langle (\Gamma, \Delta), n, (M'' \triangleright \text{eBk}''''(v, c')), (N' \triangleright \text{Bck}'''(v, c')) \rangle \\ \langle (\Gamma, \Delta), n, (M'' \triangleright \text{eBk} \parallel d!c'), (N' \triangleright \text{Bck} \parallel d!c') \rangle \\ \langle (\Gamma, \Delta), n, (M'' \triangleright \text{eBk}' \parallel d!c'), (N' \triangleright \text{Bck}' \parallel d!c') \rangle \end{array} \middle| \begin{array}{l} (\Gamma_{\text{ext}}, \Gamma_{\text{Frm}}) < \Gamma \\ n \geq 0, M' \subseteq N' \\ c \notin M'', M'' \subset N'', c \in N'' \end{array} \right\}$$

Note that, in the quadruples of \mathcal{R} our observer environment is not limited to derived environments Γ obtained from restructurings of $\Gamma_{\text{ext}}, \Gamma_{\text{Frm}}$, but may include also additional entries, denoted by the environment Δ ; these originate from observer channel allocations and uses through the transition rules `LALL` and `LSTR` from Figure 5. \mathcal{R} observes the transfer property of Definition 4.9. We here go over some key transitions:

- Consider a tuple from the first clause of the relation, for some Γ, Δ, n and c i.e.,

$$(\Gamma, \Delta) \vDash (M' \triangleright \text{eBk}''(c)) \mathcal{R}^n (N' \triangleright \text{Bck}''(c))$$

We recall from the macros introduced in Section 6 that

$$\begin{aligned} \text{eBk}''(c) &= c?(y, z).\text{free } c.\text{out}!y.(\text{eBk} \parallel d!z) \\ \text{Bck}''(c) &= c?(y, z).\text{out}!y.(\text{Bck} \parallel d!z) \end{aligned}$$

Whenever (Γ, Δ) allows it, the left hand configuration can perform an input transitions

$$(\Gamma, \Delta) \triangleleft M' \triangleright \text{eBk}''(c) \xrightarrow{c?(v, c')}_0 (\Gamma', \Delta') \triangleleft M' \triangleright \text{eBk}'''(c, v, c')$$

where $\Gamma = \Gamma', c: [\mathbf{T}, \mathbf{T}_{\text{rec}}]^1$ and $\Delta = \Delta', v: \mathbf{T}, c': \mathbf{T}_{\text{rec}}$. This can be matched by the transition

$$(\Gamma, \Delta) \triangleleft N' \triangleright \text{Bck}''(c) \xrightarrow{c?(v, c')}_0 (\Gamma', \Delta') \triangleleft N' \triangleright \text{Bck}'''(v, c')$$

where we have $(\Gamma', \Delta') \vDash (M' \triangleright \text{eBk}'''(c, v, c')) \mathcal{R}^n (N' \triangleright \text{Bck}'''(v, c'))$ from the second clause of \mathcal{R} . The matching move for an input action from the right-hand configuration is dual to this. Matching moves for `env`, `alloc` and `free c` actions are analogous.

- Consider a tuple from the first clause of the relation, for some Γ, Δ, n, c, v and c' i.e.,

$$(\Gamma, \Delta) \vDash (M' \triangleright \text{eBk}'''(c, v, c')) \mathcal{R}^n (N' \triangleright \text{Bck}'''(v, c'))$$

Since $\text{eBk}'''(c, v, c') = \text{free } c.\text{out}!v.(\text{eBk} \parallel d!c')$, a possible transition by the left-hand configuration is the deallocation of channel c :

$$(\Gamma, \Delta) \triangleleft M' \triangleright \text{eBk}'''(c, v, c') \xrightarrow{\tau}_{-1} (\Gamma, \Delta) \triangleleft M'' \triangleright \text{eBk}''''(v, c')$$

where $M' = M'', c$. In this case, the matching move is the empty (weak) transition, since we have $(\Gamma, \Delta) \vDash (M'' \triangleright \text{eBk}''''(v, c')) \mathcal{R}^{n+1} (N' \triangleright \text{Bck}'''(v, c'))$ by the third clause of \mathcal{R} . Dually, if (Γ, Δ) allows it, the right hand configuration may perform an output action

$$(\Gamma, \Delta) \triangleleft N' \triangleright \text{Bck}'''(v, c') \xrightarrow{\text{out}!v}_0 (\Gamma, \Delta, v: \mathbf{T}) \triangleleft N' \triangleright \text{Bck} \parallel d!c'$$

This can be matched by the weak output action

$$(\Gamma, \Delta) \triangleleft M' \triangleright \text{eBk}''''(v, c') \xrightarrow{\text{out}!v}_{-1} (\Gamma, \Delta, v: \mathbf{T}) \triangleleft M'' \triangleright \text{eBk} \parallel d!c'$$

where $M' = M'', c$; by the fourth clause of \mathcal{R} , we know that this a matching move because $(\Gamma, \Delta, v: \mathbf{T}) \vDash (M'' \triangleright \text{eBk} \parallel d!c') \mathcal{R}^{n+1} (N' \triangleright \text{Bck} \parallel d!c')$. \square

7. RELATED WORK

A note on terminology: From a logical perspective, a *linear* assumption is one that cannot be weakened nor contracted, while an *affine* assumption cannot be contracted but can be weakened. This leads to a reading of linear as “used exactly once” and of affine as “used at most once”. However, in the presence of divergence or deadlock, most linear type systems do not in fact guarantee that a linear resource will be used exactly once. In the discussion below, we will classify such type systems as affine instead.

Linear logic was introduced by Girard [14]; its use as a type system was pioneered by Wadler [43]. Uniqueness typing was introduced by Barendsen and Smetsers [5]; the relation to linear logic has since been discussed in a number of papers (see [17]).

Although there are many substructural (linear or affine) type systems for process calculi [1, 2, 3, 25, 32, 44, and others], some specifically for resources [33], the literature on *behaviour* of processes typed under such type systems is much smaller.

Kobayashi *et al.* [31] introduce an affine type system for the π -calculus. Their channels have a polarity (input, output, or input/output) as well as a multiplicity (unrestricted or affine), and an affine input/output can be split as an affine input and an affine output channel. Communication on an affine input/affine output channel is necessarily deterministic, like communication on an affine/unique-after-1 channel in our calculus; however, both processes lose the right to use the channel after the communication, limiting reuse. Although the paper gives a definition of reduction closed barbed congruence, no compositional proof methods are presented.

Yoshida *et al.* [45, 23] define a linear type system, which uses “action types” to rule out deadlock. The use of action types means that the type system can provide some guarantees that we cannot; this is however an orthogonal aspect of the type system and it would be interesting to see if similar techniques can be applied in our setting. Their type system does not have any type that corresponds to uniqueness; instead, the calculus is based on πI to control dynamic sharing of names syntactically, thereby limiting channel reuse. The authors give compositional proof techniques for their behavioural equivalence, but give no complete characterization.

Teller [41] introduces a π -calculus variant with “finalizers”, processes that run when a resource has been deallocated. The deallocation itself however is performed by a garbage collector. The calculus comes with a type system that provides bounds on the resources that are used, although the scope of channel reuse is limited in the absence of some sort of uniqueness information. Although the paper defines a bisimulation relation, this relation does not take advantage of type information, and no compositionality results or characterization is given.

Hoare and O’Hearn [22] give a trace semantics for a variant of CSP with point-to-point communication and explicit allocation and deallocation of channels, which relies on separation of permissions. However, they do not consider any behavioural theories. Pym and Tofts [39] similarly give a semantics for SCCS with a generic notion of resource, based on separation of permissions; they do however consider behaviour. They define a bisimulation relation, and show that it can be characterized by a modal logic. These approaches do not use a type system but opt for an operational interpretation of permissions, where actions may block due to lack of permissions. Nevertheless, our consistency requirements for configurations (Definition 4.1) can be seen as separation criteria for permission environments. A detailed comparison between this untyped approach and our typed approach would be worthwhile.

Apart from the Clean programming language [6], from where uniqueness types originated, static analysis relating to uniqueness has recently been applied to (more mainstream) Object-Oriented programming languages [15] as well. In such cases, it would be interesting to investigate

whether the techniques developed in this work can be applied to a behavioural setting such as that in [26].

Our unique-after- i type is related to fractional permissions, introduced in [9] and used in settings such as separation logic for shared-state concurrency [8]. A detailed survey of this field is however beyond the scope of this paper.

The use of substitutions in our LTS (Definition 4.3) is reminiscent of the name-bijections carried around in spi-calculus bisimulations [7]. In the spi-calculus however this substitution is carried through the bisimulation, and must remain a bijection throughout. Since processes may lose the permission to use channels in our calculus, this approach is too restrictive for us.

Finally, amortisation for coinductive reasoning was originally developed by Keihn *et al.*, [30] and Lüttgen *et al.* [34]. It is investigated further by Hennessy in [20], whereby a correspondence with (an adaptation of) reduction-barbed congruences is established. However, neither work considers aspects of resource misuse nor the corresponding use of typed analysis in their behavioural and coinductive equivalences.

8. CONCLUSION

We have presented a compositional behavioural theory for $R\pi$, a π -calculus variant with mechanisms for explicit resource management; a preliminary version of the work appeared in [11]. The theory allows us to compare the efficiency of concurrent channel-passing programs *wrt.* their resource usage. We integrate the theory with a substructural type system so as to limit our comparisons to safe programs. In particular, we interpret the type assertions of the type system as permissions, and use this to model (explicit and implicit) permission transfer between the systems being compared and the observer during compositional reasoning. Our contributions are as follows:

- (1) We define a costed semantic theory that orders systems of safe $R\pi$ programs, based on their costed extensional behaviour when deployed in the context of larger systems; Definition 5.5. Apart from cost, formulations relating to contextuality are different from those of typed congruences such as [21], because of the kind of type system used *i.e.*, substructural.
- (2) We define a bisimulation-based proof technique that allows us to order $R\pi$ programs coinductively, without the need to universally quantify over the possible contexts that these programs may be deployed in; Definition 4.9. As far as we are aware, the combination of actions-in-context and costed semantics, used in unison with implicit and explicit transfer of permissions so as to limit the efficiency analysis to safe programs, is new.
- (3) We prove a number of properties for our bisimulation preorder of Definition 4.9, facilitating the proof constructions for related programs. Whereas Corollary 4.13 follows [30, 20], Theorem 4.23 extends the property of compositionality for amortised bisimulations to a typed setting. Lemma 4.14, together with the concept of bounded amortisation, appears to be novel altogether.
- (4) We prove that the bisimulation preorder of Definition 4.9 is a sound and complete proof technique for the costed behavioural preorder of Definition 5.5; Theorem 5.10 and Theorem 5.18. In order to obtain completeness, the LTS definitions employ non-standard mechanisms for explicit renaming of channel names not known to the context. Also, the concept of (typed) action definability [21, 19] is different because it needs to take into consideration cost and typeability *wrt.* a substructural type system; the latter aspect also complicated the respective Extrusion Lemma — see Lemma 5.15.
- (5) We demonstrate the utility of the semantic theory and its respective proof technique by applying them to reason about the client-server systems outlined in the Introduction and a case study, discussed in Section 3.

Future Work. The extension of our framework to a higher-order and distributed setting seems worthwhile. Also, the amalgamation of our uniqueness types with modalities for input and output [38] would give scope for richer notions of subtyping involving covariance and contravariance, affecting the respective behavioural theory; it would be interesting to explore how our notions of permission transfer extend to such a setting. It is also worth pursuing the applicability of the techniques developed in this work to nominal automata such as Variable Automata [16] and Finite-Memory Automata [29].

Acknowledgements. We would like to thank the referees for their incisive comments.

REFERENCES

- [1] Lucia Acciai and Michele Boreale. Type abstractions of name-passing processes. In *FSEN'07*, pages 302–317, 2007.
- [2] Lucia Acciai and Michele Boreale. Responsiveness in process calculi. *Theor. Comput. Sci.*, 409(1):59–93, 2008.
- [3] Roberto M. Amadio, Gérard Boudol, and Cédric Lhoussaine. The receptive distributed π -calculus. *ACM Trans. Program. Lang. Syst.*, 25(5):549–577, 2003.
- [4] S. Arun-Kumar and Matthew Hennessy. An efficiency preorder for processes. *Acta Inf.*, 29(9):737–760, December 1992.
- [5] Erik Barendsen and Sjaak Smetsers. Uniqueness typing for functional languages with graph rewriting semantics. *MSCS*, 6:579–612, 1996.
- [6] Erik Barendsen and Sjaak Smetsers. Uniqueness typing for functional languages with graph rewriting semantics. *Mathematical Structures in Computer Science*, 6(6):579–612, 1996.
- [7] Michele Boreale, Rocco De Nicola, and Rosario Pugliese. Proof techniques for cryptographic processes. *SIAM J. Comput.*, 31(3):947–986, 2001.
- [8] Richard Bornat, Cristiano Calcagno, Peter O’Hearn, and Matthew Parkinson. Permission accounting in separation logic. *SIGPLAN Not.*, 40(1):259–270, 2005.
- [9] John Boyland. Checking interference with fractional permissions. In R. Cousot, editor, *Static Analysis: 10th International Symposium*, volume 2694 of *LNCS*, pages 55–72. Springer, 2003.
- [10] Dov Bulka and David Mayhew. *Efficient C++: performance programming techniques*. Addison-Wesley Longman Publishing Co., Inc., Boston, MA, USA, 2000.
- [11] Edsko DeVries, Adrian Francalanza, and Matthew Hennessy. Reasoning about explicit resource management (extended abstract). In *PLACES*, pages 15–21. ETAPS, April 2011. <http://places11.di.fc.ul.pt/>.
- [12] Edsko DeVries, Adrian Francalanza, and Matthew Hennessy. Uniqueness typing for resource management in message-passing concurrency. *Journal of Logic and Computation*, June 2012.
- [13] Adrian Francalanza, Julian Rathke, and Vladimiro Sassone. Permission-based separation logic for message-passing concurrency. *Logical Methods in Computer Science*, 7(3), 2011.
- [14] Jean-Yves Girard. Linear logic. *Theoretical Computer Science*, 50(1):1–102, 1987.
- [15] Colin S. Gordon, Matthew J. Parkinson, Jared Parsons, Aleks Bromfield, and Joe Duffy. Uniqueness and Reference Immutability for Safe Parallelism. In *Proceedings of the 2012 ACM International Conference on Object Oriented Programming, Systems, Languages, and Applications (OOPSLA’12)*, Tucson, AZ, USA, October 2012.
- [16] Orna Grumberg, Orna Kupferman, and Sarai Sheinvald. Variable automata over infinite alphabets. In Adrian-Horia Dediu, Henning Fernau, and Carlos Martn-Vide, editors, *Language and Automata Theory and Applications*, volume 6031 of *LNCS*, pages 561–572. Springer, 2010.
- [17] Jurriaan Hage, Stefan Holdermans, and Arie Middelkoop. A generic usage analysis with subeffect qualifiers. In *Proceedings of the 12th ACM SIGPLAN International Conference on Functional Programming (ICFP)*, pages 235–246. ACM, 2007.
- [18] Dana Harrington. Uniqueness logic. *Theoretical Computer Science*, 354(1):24–41, 2006.
- [19] Matthew Hennessy. *A Distributed Picalculus*. Cambridge University Press, Cambridge, UK., 2008.
- [20] Matthew Hennessy. A calculus for costed computations. *Logical Methods in Computer Science*, 7(1), 2011.
- [21] Matthew Hennessy and Julian Rathke. Typed behavioural equivalences for processes in the presence of subtyping. *Mathematical Structures in Computer Science*, 14:651–684, 2004.
- [22] Tony Hoare and Peter O’Hearn. Separation logic semantics for communicating processes. *ENTCS*, 212:3–25, 2008.
- [23] Kohei Honda. From process logic to program logic. In *ICFP ’04*, pages 163–174, 2004.

- [24] Kohei Honda and Mario Tokoro. On asynchronous communication semantics. In Mario Tokoro, Oscar Nierstrasz, and Peter Wegner, editors, *Proceedings of the ECOOP'91 Workshop on Object-Based Concurrent Computing*, volume 612 of *LNCS*, pages 21–51. Springer-Verlag, 1992.
- [25] Atsushi Igarashi and Naoki Kobayashi. A generic type system for the pi-calculus. *Theor. Comput. Sci.*, 311(1-3):121–163, 2004.
- [26] Alan Jeffrey and Julian Rathke. Java Jr: Fully abstract trace semantics for a core java language. In Shmuel Sagiv, editor, *ESOP*, volume 3444 of *LNCS*, pages 423–438. Springer, 2005.
- [27] Richard Jones. *Garbage Collection: Algorithms for Automatic Dynamic Memory Management*. John Wiley and Sons, July 1996. With a chapter on Distributed Garbage Collection by Rafael Lins. Reprinted 1997 (twice), 1999, 2000.
- [28] Richard Jones, Anthony Hosking, and Eliot Moss. *The Garbage Collection Handbook: The Art of Automatic Memory Management*. Applied Algorithms and Data Structures. Chapman and Hall/CRC, 1 edition, 2011.
- [29] Michael Kaminski and Nissim Francez. Finite-memory Automata. *Theoretical Computer Science*, 134(2):329 – 363, 1994.
- [30] Astrid Kiehn and S. Arun-Kumar. Amortised bisimulations. In *FORTE 2005*, volume 3731 of *LNCS*, pages 320–334, 2005.
- [31] Naoki Kobayashi, Benjamin C. Pierce, and David N. Turner. Linearity and the pi-calculus. *ACM Trans. Program. Lang. Syst.*, 21(5):914–947, 1999.
- [32] Naoki Kobayashi and Davide Sangiorgi. A hybrid type system for lock-freedom of mobile processes. *ACM Trans. Program. Lang. Syst.*, 32(5):1–49, 2010.
- [33] Naoki Kobayashi, Kohei Suenaga, and Lucian Wischik. Resource usage analysis for the pi-calculus. *Logical Methods in Computer Science*, 2(3), 2006.
- [34] Gerald Lüttgen and Walter Vogler. Bisimulation on speed: A unified approach. *Theor. Comput. Sci.*, 360(1-3):209–227, 2006.
- [35] R. Milner. *Communicating and mobile systems: the π -calculus*. Cambridge Univ., 1999.
- [36] Benjamin C. Pierce. *Types and programming languages*. MIT Press, Cambridge, MA, USA, 2002.
- [37] Benjamin C. Pierce. *Advanced Topics in Types and Programming Languages*. The MIT Press, Cambridge, MA, USA, 2004.
- [38] Benjamin C. Pierce and Davide Sangiorgi. Typing and subtyping for mobile processes. *Mathematical Structures in Computer Science*, 6(5):409–453, 1996.
- [39] David Pym and Chris Tofts. A calculus and logic of resources and processes. *Form. Asp. Comput.*, 18(4):495–517, 2006.
- [40] Frederick Smith, David Walker, and J. Gregory Morrisett. Alias types. In *ESOP*, volume 1782 of *LNCS*, pages 366–381. Springer, 2000.
- [41] David Teller. Recollecting resources in the pi-calculus. In *Proceedings of IFIP TCS 2004*, pages 605–618. Kluwer Academic Publishing, 2004.
- [42] T. Terauchi and A. Aiken. A capability calculus for concurrency and determinism. *TOPLAS*, 30(5):1–30, 2008.
- [43] Philip Wadler. Is there a use for linear logic? In *PEPM*, pages 255–273, 1991.
- [44] Nobuko Yoshida. Channel dependent types for higher-order mobile processes. *SIGPLAN Not.*, 39(1):147–160, 2004.
- [45] Nobuko Yoshida, Kohei Honda, and Martin Berger. Linearity and bisimulation. *Journal of Logic and Algebraic Programming*, 72(2):207 – 238, 2007.