

A CALCULUS FOR COSTED COMPUTATIONS

MATTHEW HENNESSY

Department of Computer Science, Trinity College Dublin, Ireland
e-mail address: matthew.hennessy@cs.tcd.ie

ABSTRACT. We develop a version of the picalculus *Picost* where channels are interpreted as *resources* which have costs associated with them. Code runs under the financial responsibility of *owners*; they must pay to *use* resources, but may profit by *providing* them.

We provide a proof methodology for processes described in *Picost* based on bisimulations. The underlying behavioural theory is justified via a contextual characterisation. We also demonstrate its usefulness via examples.

1. INTRODUCTION

The purpose of this paper is to develop a behavioural theory of processes, in which computations depend on the ability to fund the resources involved. The theory will be based on the well-known concept of bisimulations, [Mil99], which automatically gives a powerful co-inductive proof methodology for establishing properties of processes; here these properties will include the cost of behaviour.

We take as a starting point the well-known picalculus, [SW01, Mil99], a language for describing mobile processes which has a well-developed behavioural theory. In the picalculus a process is described in terms of its ability to input and output on *communication channels*. Here we interpret these channels as *resources*, or services, as for example in [CGP08]. So input along a channel, written as $c?(x).P$ in the picalculus, is now interpreted as *providing* the service c , while output, written $c!\langle v \rangle.P$, is interpreted as a request to *use* the service c . A process is now determined by the manner in which it *provides* services and *uses* them.

Viewed from this perspective, we extend the picalculus in two ways. Firstly we associate a *cost* with resources; specifically for each resource we assume that a certain amount of funds k_u is charged to *use* it, and an amount k_p is also required to *provide* it. Secondly we introduce *principals* or *owners* who provide the funds necessary for the functioning of resources. The novel construct in the language is $[P]_o$, representing the (picalculus) process P running under the financial responsibility of o . For example in $[c!\langle v \rangle.Q]_o$ the use of the resource c is only possible if o can fund the charges. Similarly with $[c?(x).Q]_o$, but

1998 ACM Subject Classification: F.3.1 [Specifying and Verifying and Reasoning about programs]: bisimulations for distributed resources, F.3.2 [Semantics of Programming Languages]: operational semantics, F.3.3 [Studies of Program Constructs]: resource usage.

Key words and phrases: resources, cost, picalculus, bisimulations, amortisation.

The financial support of SFI is gratefully acknowledged.

here there is also the potential for gain for owner \mathfrak{o} ; in our formulation \mathfrak{o} profits from any difference between the cost in providing the resource and the charge made to use it.

Our language *Picost* is presented in Section 2, and is essentially a variation on *Dpi*, a typed distributed version of the picalculus, [Hen07]. The reduction semantics is given in terms of judgements of the form

$$(\Gamma \triangleright M) \longrightarrow (\Delta \triangleright N)$$

where Γ, Δ are *cost environments*. These have a static component, giving the costs associated with resources, and a dynamic part, which gives the funds available to owners and also records expenditure. The usefulness of the language is demonstrated by a series of simple examples.

But the main achievement of the paper is a behavioural theory, expressed as judgements

$$(\Gamma \triangleright M) \sqsubseteq_{\text{awgt}} (\Delta \triangleright N) \tag{1.1}$$

indicating that, informally speaking,

- (i) the process M running relative to the cost environment Γ is bisimilar, in the standard sense [Mil89], with process N running relative to Δ
- (ii) the costs associated with $(\Delta \triangleright N)$ are no more, and possibly less, than those associated with $(\Gamma \triangleright M)$.

Influenced by [KAK05] we first develop a general framework of *weighted labelled transition systems* or *wLTSs*, in which actions, including internal actions, may have multiple weights associated with them. We then define a notion of *amortised weighted bisimulations* between their states, giving rise to a preorder $s \sqsubseteq_{\text{awgt}} t$, meaning that s, t are bisimilar but in some sense the behaviours of t are *lighter* than those of s . From this we obtain, in the standard manner, a co-inductive proof methodology for proving that two systems are related; it is sufficient to find, or construct, a particular amortised weighted bisimulation containing the pair (s, t) .

This proof methodology is applied to *Picost* by first interpreting the language as an LTS, in agreement with the reduction semantics, and then interpreting this LTS as a wLTS, giving rise to (parametrised versions of) the judgements (1.1) above. But as we will see these judgements can be interpreted in two ways. If the recorded expenditure represents *costs* then $(\Delta \triangleright N)$ can be considered an improvement on $(\Gamma \triangleright M)$. On the other hand if it represents *profits* then we have the reverse; $(\Gamma \triangleright M)$ is an improvement on $(\Delta \triangleright N)$ as it has the potential to be heavier.

The details of this theory are given in Section 3, and the resulting proof methodology is illustrated by examples. However in Section 4 we re-examine this proof methodology, in the light of reasonable properties we would expect of it; and these are found wanting. It turns out that the manner in which we generate the wLTS for *Picost* from its operational semantics is too coarse. We show how to generate a somewhat more abstract wLTS, and prove that the resulting proof methodology is satisfactory, in a precise technical sense, by adapting the notion of reduction barbed congruence, [HT92, SW01, HR04, Hen07].

2. THE LANGUAGE *Picost*

2.1. Syntax: We assume a set of *channel* or *resource* names Chan , ranged over by a, b, c, \dots, r, \dots whose use requires some cost, a distinct set of (value) variables Var , ranged over by x, y, \dots , and a further distinct set of recursion variables, X, Y, \dots ; u ranges over

$M, N ::=$	$[T]_{\circ}$	Owned code
	$M \mid N$	Composition
	$(\text{new } r:\mathbf{R})M$	Scoped resource
	0	Identity
$T, U ::=$		
	$u?(x).T$	Provide resource u
	$u!\langle v \rangle.T$	Use resource u
	$\text{if } v = v \text{ then } T \text{ else } U$	Matching
	$(\text{new } r:\mathbf{R})T$	Resource creation
	$T \mid U$	Concurrency
	$\text{rec } X. T$	Recursion
	X	Recursion variable
	stop	Termination

Figure 1: Syntax of Picost

identifiers, which may be either resource names or (value) variables. We also assume a set of *principals* or *owners* Own containing at least two elements, ranged over by \circ, u, p , who are implicitly registered for these resources and who finance their provision and use. The syntax of Picost is then given in Figure 1, and is essentially a very minor variation on Dpi , [Hen07]. The main syntactic category represents code running under responsibility, with $[P]_{\circ}$ being the novel construct. As explained in the Introduction this represents the code P running under the responsibility of the owner \circ ; intuitively \circ is financially responsible for the computation P . Thus in general a system is simply a collection of computation threads each running under the responsibility of an explicit owner, which may share private resources. The syntax for these threads is a version of the well-known picalculus, [SW01].

The type \mathbf{R} of a resource describes the costs associated with that resource. There is a cost associated with *using* a resource, and a cost associated with *providing* it; therefore types take the form $\langle k_u, k_p \rangle$ where k_u, k_p are elements from some *cost domain* K . Here we take K simply to be \mathbb{N} ordered in the standard manner, but most of our results apply equally well to variations.

We employ the standard abbreviations associated with the picalculus, and associated terminology. In particular we assume *Barendregt's convention*, which implies that bound variables used in terms or definitions are distinct, and different from any free variables in use in the current context. In Figure 1 meta-variable v range over *value expressions*, whose specification we omit; but they include at least resource names $a \in \text{Chan}$, variables x from Var , and elements of K . As usual we omit every occurrence of a trailing stop and abbreviate $u?(\cdot).T, u!\langle \cdot \rangle.T$ to $u?.T, u!.T$ respectively. We are only interested in *closed code terms*, those which contain no free occurrences of variables, which are ranged over by P, Q, \dots ; we use $\text{fn}(P)$ to denote the set of names from Chan which occur freely in P . In the sequel we assume all terms are closed.

2.2. Cost environments: Since computations have financial implications, the execution of processes is now relative to a *cost environment* Γ . This records the financial resources available to principals, and the cost of providing and using resources; in order to be able to compare the cost of computations we also assume a component which records the expenditure as a computation proceeds. Thus judgements of the reduction semantics take the form

$$\Gamma \triangleright M \longrightarrow \Delta \triangleright N$$

where Γ, Δ are cost environments.

There are many possibilities for *cost environments*; see [HG08] for an example which directly associates funds with resources. In the present paper we define them in such a way that the owners retain total control over their own funds.

Definition 2.1 (Cost environments). A *cost environment* Γ consists of a 4-tuple $\langle \Gamma^o, \Gamma^u, \Gamma^p, \Gamma^{\text{rec}} \rangle$ where

- $\Gamma^u : \text{Chan} \rightarrow K$
 $\Gamma^u(a)$ records the cost of *using* resource a ; this is a *static* component, and will not vary during computations
- $\Gamma^p : \text{Chan} \rightarrow K$
 $\Gamma^p(a)$ records the cost of *providing* resource a ; again this is a static component
- $\Gamma^o : \text{Own} \rightarrow K$
 $\Gamma^o(o)$ records the funds available to owner o ; this will vary as computations proceed, as owners will need to fund their interactions with resources
- $\Gamma^{\text{rec}} \in K$
 Γ^{rec} keeps an account of the expenditure occurred during a computation; of course this also will vary as a computation proceeds.

We assume that both functions Γ^u, Γ^p have the same finite domain, but not necessarily that $\Gamma^u(a) \geq \Gamma^p(a)$ whenever these are defined. \blacksquare

We now define some operations on cost environments which will enable us to reflect their impact on the semantics of our language. The most important is a partial function, $\Gamma \xrightarrow{(u,a,p)} \Delta$, which informally means that in Γ owner u has sufficient funds to cover the cost of using resource a and owner p has sufficient funds to provide it. Then Δ records the result of the expenditure of both o and p of those funds. There is also considerable scope as to what happens to these funds, and how their expenditure is recorded. Here we take the view that the provider p gains the cost which the user expends, to offset p 's cost in providing the resource.

Definition 2.2 (Resource charging). Let $\xrightarrow{(u,a,p)}$ be the partial function over cost environments defined as follows: $\Gamma \xrightarrow{(u,a,p)} \Delta$ if

- (i) $\Gamma^o(u) \geq \Gamma^u(a)$ and $\Gamma^o(p) \geq \Gamma^p(a)$
- (ii) Δ is the cost environment obtained from Γ by
 - (a) decreasing $\Gamma^o(u)$ by the amount $\Gamma^u(a)$
 - (b) increasing $\Gamma^o(p)$ by the amount $\Gamma^u(a) - \Gamma^p(a)$, which may of course be negative
- (iii) Finally there is considerable flexibility in how this resource expenditure is recorded in Δ^{rec} . We call resource charging for a *standard* when this is set to $\Gamma^{\text{rec}} + \Gamma^u(a) - \Gamma^p(a)$; that is we add to the record the gain obtained in using resource a . But in general

we allow functions $\text{rec}_a(-, -)$, for each resource a , in which case we define Δ^{rec} to be $\Gamma^{\text{rec}} + \text{rec}_a(\Gamma^u(a), \Gamma^p(a))$. ■

In general we allow the owners u and p in this definition to coincide. So, for example if $\Gamma \xrightarrow{(\circ, a, \circ)} \Delta$, then the effect of performing (a) above, followed by (b), is that $\Delta^\circ(\circ)$ is set to $\Gamma^\circ(\circ) - \Gamma^p(a)$.

The use of two independent charges for each resource, Γ^u and Γ^p , may seem overly complex. A simpler model can be obtained by having only one combined charge; effectively we could assume $\Gamma^p(a)$ to be 0 for every a , and so resource charging simply transfers the appropriate amount of funds from the user to the provider; this could be achieved by restricting attention to *simple types*, resource types R of the form $\langle k_u, 0 \rangle$. Indeed this simplification will be quite useful in order to achieve some theoretical properties of our proof methodology; see Definition 4.18 and Section 4.2. Nevertheless the use of the two independent charges $\Gamma^p(-)$ and $\Gamma^u(-)$ allows scope for more interesting examples. In particular it provides considerable scope for variation in the manner in which resource expenditure is recorded in the component Γ^{rec} ; see Example 2.8 for an instance.

We also need to extend cost environments with new resources.

Definition 2.3 (Resource registration). The cost environment $\Gamma, a:R$, is only defined if a is *fresh* to Γ , that is, if a is neither in $\text{dom}(\Gamma^u)$ nor in $\text{dom}(\Gamma^p)$. In this case it gives the new cost environment Δ obtained by adding the new resource, with the capabilities determined by R . Formally the dynamic components of Δ , namely Δ° and Δ^{rec} , are inherited directly from Γ , while the static components have the obvious definition; for example if R is the type $\langle k_u, k_p \rangle$ then Δ^u is given by

$$\Delta^u(x) = \begin{cases} k_u & \text{if } x = a \\ \Gamma^u(x) & \text{otherwise} \end{cases}$$

We also assume that the resource charging for a in $(\Gamma, a:R)$ is always standard. ■

Note that every cost environment may be written in the form

$$\Gamma_{\text{dyn}}, a_1:R_1, \dots, a_n:R_n$$

where Γ_{dyn} is a *basic* environment; that is the static components Γ_{dyn}^u and Γ_{dyn}^p are both empty, and so it only contains non-trivial dynamic components.

2.3. Reduction semantics: The pair $(\Gamma \triangleright M)$ is called a *configuration* if $\text{fn}(M) \subseteq \text{dom}(\Gamma^u) = \text{dom}(\Gamma^p)$, that is every free resource name in M is known to the *cost environment* Γ . The reduction semantics for Picost is then defined as the least relation over configurations which satisfies the rules in Figure 2. The majority of the rules come directly from the reduction semantics of Dpi, [Hen07], and are housekeeping in nature. The only rule of interest is (R-COMM), representing the communication along the channel a , or in Picost the *use* of the resource a by owner u which is *provided* by owner p . However this reduction is only possible whenever the premise $\Gamma \xrightarrow{(u, a, p)} \Delta$ is satisfied. As we have seen, this means that in Γ owner u has sufficient funds to cover the cost of using resource a and owner p has sufficient funds to provide it; and further Δ records the result of the expenditure of both u and p of those funds.

The remainder of the rules are borrowed directly from the standard reduction semantics of Dpi; note that (R-STRUCT) requires a structural equivalence between terms; this again is

$$\begin{array}{c}
\text{(R-COMM)} \\
\frac{\Gamma \xrightarrow{(u,a,p)} \Delta}{\Gamma \triangleright [a!\langle v \rangle . Q]_u \mid [a?(x) . P]_p \longrightarrow \Delta \triangleright [Q \mid P\{v/x\}]_p} \\
\text{(R-SPLIT)} \\
\Gamma \triangleright [M \mid N]_o \longrightarrow \Gamma \triangleright [M]_o \mid [N]_o \\
\text{(R-EXPORT)} \\
\Gamma \triangleright [(\text{new } r : \mathbf{R})P]_o \longrightarrow \Gamma \triangleright (\text{new } r : \mathbf{R})[P]_o \\
\text{(R-UNWIND)} \\
\Gamma \triangleright [\text{rec } x . T]_o \longrightarrow \Gamma \triangleright [T\{\text{rec } x . T/x\}]_o \\
\text{(R-MATCH)} \\
\Gamma \triangleright [\text{if } a = a \text{ then } P \text{ else } Q]_o \longrightarrow \Gamma \triangleright [P]_o \\
\text{(R-MISMATCH)} \\
\Gamma \triangleright [\text{if } a = b \text{ then } P \text{ else } Q]_o \longrightarrow \Gamma \triangleright [Q]_o \quad a \neq b \\
\text{(R-STRUCT)} \\
\frac{M \equiv M', \Gamma \triangleright M \longrightarrow \Delta \triangleright N, N \equiv N'}{\Gamma \triangleright M' \longrightarrow \Delta \triangleright N'} \\
\text{(R-CNTX)} \\
\frac{\Gamma \triangleright M \longrightarrow \Delta \triangleright M'}{\Gamma \triangleright M \mid N \longrightarrow \Delta \triangleright M' \mid N} \\
\text{(R-NEW)} \\
\frac{\Gamma, b : \mathbf{R} \triangleright M \longrightarrow \Delta, b : \mathbf{R} \triangleright N}{\Gamma \triangleright (\text{new } b : \mathbf{R})M \longrightarrow \Delta \triangleright (\text{new } b : \mathbf{R})N}
\end{array}$$

Figure 2: Reduction semantics

$$\begin{array}{ll}
\text{(S-EXTR)} & (\text{new } r : \mathbf{R})(M \mid N) \equiv M \mid (\text{new } r : \mathbf{R})N, \text{ if } r \notin \text{fn}(M) \\
\text{(S-COM)} & M \mid N \equiv N \mid M \\
\text{(S-ASSOC)} & (M \mid N) \mid O \equiv M \mid (N \mid O) \\
\text{(S-ZERO)} & M \mid \mathbf{0} \equiv M \\
& [\text{stop}]_o \equiv \mathbf{0} \\
\text{(S-FLIP)} & (\text{new } r : \mathbf{R})(\text{new } r' : \mathbf{R}')M \equiv (\text{new } r' : \mathbf{R}')(\text{new } r : \mathbf{R})M
\end{array}$$

Figure 3: Structural equivalence of Picost

the standard one from Dpi, the definition of which is given in Figure 3. Also the final rule (R-NEW) uses the registration operation on cost environments, given in Definition 2.3.

Proposition 2.4. *If $(\Gamma_1 \triangleright M_1)$ is a configuration and $(\Gamma_1 \triangleright M_1) \longrightarrow (\Gamma_2 \triangleright M_2)$ then $(\Gamma_2 \triangleright M_2)$ is also a configuration.*

Proof. Straightforward, by induction on the proof that $(\Gamma_1 \triangleright M_1) \longrightarrow (\Gamma_2 \triangleright M_2)$. When handling the rule (R-STRUCT) it uses the obvious fact that $M \equiv N$ implies that M and N have the same set of free names; this in turn means that $M \equiv N$ implies $\Gamma \triangleright M$ is a configuration if and only if $\Gamma \triangleright N$ is. \square

The reductions of a configuration affect its cost environment, and as a sanity check we can describe precisely the kinds of changes which are possible:

Proposition 2.5. *Suppose $(\Gamma_1 \triangleright M_1) \longrightarrow (\Gamma_2 \triangleright M_2)$. Then*

$$\begin{aligned} \text{Sys} &\Leftarrow ([\text{Reader}]_{\text{pub}} \mid [\text{Library} \mid \text{Store}]_{\text{lib}}) \\ \text{where} & \\ \text{Reader} &\Leftarrow \text{rec } R. \text{goLib?}(\text{name}) . (\text{new } r) \text{reqR!}\langle r, \text{name} \rangle . \\ &\quad r?(b) . \text{goHome!}\langle b \rangle . R \\ \text{Library} &\Leftarrow \text{rec } L. \text{reqR?}(y, z) . \quad y!\langle \text{book}(z) \rangle . L \\ &\quad \oplus (\text{new } r) \text{reqS!}\langle r, z \rangle . r?(b) . y!\langle b \rangle . L \\ \text{Store} &\Leftarrow \text{rec } S. \text{reqS?}(y, z) . y!\langle \text{book}(z) \rangle . S \end{aligned}$$

Figure 4: Running a library

- (i) $\Gamma_1 = \Gamma_2$, and $(\Delta \triangleright M_1) \longrightarrow (\Delta \triangleright M_2)$ whenever $(\Delta \triangleright M_1)$ is a configuration
- (ii) or $\Gamma_1 \xrightarrow{(u,a,p)} \Gamma_2$, for some resource a and owners u, p , and whenever $(\Delta \triangleright M_1)$ is a configuration $\Delta \xrightarrow{(u,a,p)} \Delta'$ implies $(\Delta \triangleright M_1) \longrightarrow (\Delta' \triangleright M_2)$
- (iii) or $\Gamma_1, a:\mathbf{R} \xrightarrow{(u,a,p)} \Gamma_2, a:\mathbf{R}$, for some (fresh) resource a , resource type \mathbf{R} and owners u, p , and whenever $(\Delta \triangleright M_1)$ is a configuration $\Delta, a:\mathbf{R} \xrightarrow{(u,a,p)} \Delta', a:\mathbf{R}$ implies $(\Delta \triangleright M_1) \longrightarrow (\Delta' \triangleright M_2)$

Proof. Again this is a simple proof by rule induction on the premise $(\Gamma_1 \triangleright M_1) \longrightarrow (\Gamma_2 \triangleright M_2)$. Intuitively possibility (i) corresponds to a move where no communication occurs, (ii) is when the move is a communication along a channel a known to Γ_1 , and (iii) when the communication is along a private internal channel. \square

2.4. Examples: Formally Picost has only unary communication, but in these examples we will informally allow the communication of tuples along channels. In addition we will use the standard abbreviations associated with the picalculus. We also omit types for channels when they are not relevant; in such cases we assume that they cost nothing to provide, and that there is no charge for using them. It will be convenient to have an *internal choice* operator, with $P \oplus Q$ representing an internal choice between P and Q . This can be taken to be short-hand notation for $(\text{new } c)(c!\langle \rangle \mid c?() . P \mid c?() . Q)$, where c is a fresh channel.

Example 2.6 (Running a library). Consider the system **Sys** from Figure 4, which consists of three recursive components, a library user **Reader**, running under the responsibility of the principal **pub**, standing for **public**, a library interface **Library** and an auxiliary book depository **Store**, both running under some other principal **lib**.

The programming of these components involves the systematic generation of *reply channels*. Thus for example the **Reader** gets the name of a book with which to go to the library, generates a new reply channel r and submits this together with the name of the book via **reqR**; it awaits the book and then returns home. The **Store** is also very simple; it recursively awaits a request on **reqS**, consisting of a reply channel and a **name** and returns the appropriate book on the channel. Finally the **Library** service requests at **reqR** consisting of a reply channel and a name. The book may be immediately available, in which case it is returned, or it may be necessary to send a request to the **Store**.

Let us now consider the behaviour of these systems relative to two cost environments Γ_{local} , Γ_{central} representing two different strategies for providing library services. To focus on the relative cost of providing these services let us assume that their use is free, that is $\Gamma_*^u(a) = 0$ for every resource a , where $*$ ranges over **local**, **central**, and that the amount of funds available is not an issue, that is $\Gamma_*^o(\text{pub}) = \Gamma_*^o(\text{lib}) = \infty$. The cost of providing the services, Γ_*^p is given in the table below, reflecting on the one hand the relative convenience to the Reader of the local services, and on the other the relative convenience to the authorities in providing central services.

	local	central
goLib	1	5
goHome	1	5
reqR	3	1
reqS	5	1

Finally let us take the counters Γ_*^{rec} to be initially set to 0. Note that Γ_{local} can be written as

$$\Gamma_{\text{dyn}}, \text{goLib}:R_l^g, \text{goHome}:R_l^h, \text{reqR}:R_l^r, \text{reqS}:R_l^s$$

where $R_l^g, R_l^h, R_l^r, R_l^s$ are the types $\langle 0, 1 \rangle, \langle 0, 1 \rangle, \langle 0, 3 \rangle, \langle 0, 5 \rangle$ respectively, and Γ_{dyn} is a basic environment; Γ_{central} has a similar representation, with a slightly different sequence of types.

To exercise the system we use

$$\text{Book} \leftarrow [\text{goLib}!(\text{str}).\text{goHome}?(x).\text{stop}]_{\text{pub}}$$

to prod the Reader into action, where str is the name of some book. Consider the configuration

$$\mathcal{C}_1 = \Gamma_{\text{local}} \triangleright (\text{Book} \mid \text{Sys}),$$

and let us ignore the computation steps involved in generating reply channels, and general housekeeping such as the unwinding of recursive definitions, which in any event cost nothing. Because of the internal non-determinism in the library service there are essentially two computations from \mathcal{C}_1 . If the **Store** is not used then after three computation steps which require funds it is in the state $\Delta_{\text{local}} \triangleright \text{Sys}$, where $\Delta_{\text{local}}^{\text{rec}} = 5$. This represents the overall cost of this transaction, 2 of which is paid by **pub** and 3 by **lib**.

On the other hand if the **Store** is used, then there are four computation steps which require funding, after which the state $\Theta_{\text{local}} \triangleright \text{Sys}$ is reached, where $\Theta_{\text{local}}^{\text{rec}} = 10$. However using the central cost environment Γ_{central} the two possibilities are $\Delta_{\text{central}}^{\text{rec}} = 11$ and $\Theta_{\text{central}}^{\text{rec}} = 12$ respectively. In each eventuality the local implementation is more efficient, in the sense that the costs are systematically lower. ■

The charging regime for resources is such that their use effectively means a transfer of funds to the *provider* from the *user*, provided the cost of providing the resource is less than the charge for its use. This enables us to implement a systematic way of transferring funds between owners.

$$\begin{aligned} \text{Sys} &\Leftarrow [P]_{\mathfrak{p}} \mid [N]_{\mathfrak{n}} \mid [A]_{\mathfrak{a}} \mid [R]_{\mathfrak{r}} \\ \text{where} \\ P &\Leftarrow \text{rec } P. (\text{new } r_1) \text{news}! \langle r_1 \rangle. (\text{new } r_2) \text{adv}! \langle r_2 \rangle. \\ &\quad r_1? \langle n \rangle. r_2? \langle d \rangle. \text{publish}? \langle z \rangle. z! \langle n, d \rangle. P \\ N &\Leftarrow \text{rec } N. \text{news}? \langle r \rangle. (\text{new } n) r! \langle n \rangle. N \\ A &\Leftarrow \text{rec } A. \text{adv}? \langle r \rangle. (\text{new } d) r! \langle d \rangle. A \\ R &\Leftarrow \text{rec } R. (\text{new } r) \text{publish}! \langle r \rangle. r? \langle n, d \rangle. R \end{aligned}$$

Figure 5: Publishing

Example 2.7 (Fund transfer). Consider the systems defined as follows:

$$\begin{aligned} \text{Sys} &\Leftarrow [D]_{\text{dad}} \mid [K]_{\text{kate}} \\ \text{where} \\ D &\Leftarrow \text{req}? \langle x \rangle. (\text{new } s : R_s) x! \langle s \rangle. s!..S \\ K &\Leftarrow (\text{new } r) \text{req}! \langle r \rangle. r? \langle y \rangle. y?.H \end{aligned}$$

The size of the transfer from dad to kate depends on the type R_s at which the new channel s is declared. Suppose this type is $\langle 0, k \rangle$, and let Γ be a cost environment in which $\Gamma^o(\text{dad})$ is at least k . Then there is a computation

$$(\Gamma \triangleright \text{Sys}) \longrightarrow^* (\Delta \triangleright [S]_{\text{dad}} \mid [H]_{\text{kate}})$$

in which $\Delta^o(\text{dad}) = \Gamma^o(\text{dad}) - k$ and $\Delta^o(\text{kate}) = \Gamma^o(\text{kate}) + k$. ■

Example 2.8 (Publishing). Consider the system Sys in Figure 5, which has four components:

- (a) publisher: *uses* a news service via the resource **news**, *uses* an advertising agency via the resource **adv** and *provides* the resource **publish**
- (b) news service: *provides* a service via **news**
- (c) ad agency: *provides* a service via **adv**
- (d) reader: *uses* the resource **publish**

The viability of publishing depends of course on the cost associated with these resources. As an example consider an environment Γ_{327} , of the form $\Gamma_{\text{dyn}}, \text{news} : R_n, \text{adv} : R_a, \text{publish} : R_p$, where these types are $\langle 3, 1 \rangle, \langle 2, 0 \rangle, \langle 7, 1 \rangle$ respectively, and let us assume $\Gamma_{327}^{\text{rec}}$ is initialised to 0. Furthermore, since we are concentrating on the publisher, let us assume that the resource charging is defined so that only the effect on the owner \mathfrak{p} is recorded. Referring to Definition 2.2 this means that resource charging is standard for **publish** but we need to set $\text{rec}_a(k_u, k_p)$ to be $-k_u$, if a is either **news** or **adv**.

Now consider a computation from the configuration $\Gamma_{317} \triangleright \text{Sys}$. Provided the owners have sufficient funds, specifically $\Gamma^o(\mathfrak{p}), \Gamma^o(\mathfrak{n})$ and $\Gamma^o(\mathfrak{r})$ must be at least 5, 1, 7 respectively, then we have a computation

$$(\Gamma_{317} \triangleright \text{Sys}) \longrightarrow^* (\Delta_1 \triangleright \text{Sys})$$

where $\Delta_1^{\text{rec}} = 1$; the record part of the initial environment was set to 0, during the computation it was set to -3 after the publisher uses the **news** resource, then to -5 after using

adv; finally, when the reader uses the `publish` resource, this is increased by $(7 - 1)$ to give 1. Because we have defined expenditure recording to reflect the point of view of the publisher, this represents the fact that the publisher has made a profit of 1 as a result of this sequence of transactions. Note also that at this point $\Delta_1^o(\mathbf{p})$ is $\Gamma_{327}^o(\mathbf{p}) + 1$.

We can also see what happens when the costs of using resources is changed. Let Γ_{216} be the environment in which the cost of all three resources are decreased by 1. Then we have the computation

$$(\Gamma_{216} \triangleright \text{Sys}) \longrightarrow^* (\Delta_2 \triangleright \text{Sys})$$

where now $\Delta_2^{\text{rec}} = 2$; this represents an increase in profits for the publisher. \blacksquare

Example 2.9 (Kickbacks). Suppose in Figure 5 we change the situation so that the publisher obtains a kickback from the ad agency when an ad is downloaded. The modified code is given by

$$\begin{aligned} P_K &\Leftarrow \text{rec } P. (\text{new } r_1) \text{news}! \langle r_1 \rangle. (\text{new } r_2) (\text{new } k : K) \text{adv}! \langle k, r_2 \rangle. \\ &\quad r_1? \langle n \rangle. r_2? \langle d \rangle. \text{publish}? \langle z \rangle. k? . z! \langle n, d \rangle. P \\ A_K &\Leftarrow \text{rec } A. \text{adv}? \langle k, r \rangle. (\text{new } d) r! \langle d \rangle. (A \mid k!) \end{aligned}$$

and let Sys_K denote the revised system. The size of the kickback depends on the parameters in the type K . In Sys the ad agency receives the benefit 2 for supplying the ad; if we set K to be $\langle 1, 0 \rangle$ then in Sys_K this benefit is split equally with the publisher. Under the same assumptions as in Example 2.8 we have the computations

$$(\Gamma_{327} \triangleright \text{Sys}_K) \longrightarrow^* (\Phi_1 \triangleright \text{Sys}_K) \quad \text{and} \quad (\Gamma_{216} \triangleright \text{Sys}_K) \longrightarrow^* (\Phi_2 \triangleright \text{Sys}_K)$$

where now $\Phi_1^{\text{rec}}, \Phi_2^{\text{rec}}$ are 2, 3 respectively, indicating more profit in each case for the publisher. \blacksquare

3. COMPOSITIONAL REASONING

The aim of this section is to develop a proof methodology for `Picost`. The idea is to define a *behavioural preorder*

$$(\Gamma \triangleright M) \sqsubseteq (\Delta \triangleright N), \tag{3.1}$$

meaning that in some sense $(\Gamma \triangleright M)$ and $(\Delta \triangleright N)$ offer the same behaviour, but the latter is at least as efficient as the former, and possibly more. We follow the standard approach of defining the preorder (3.1) as the largest relation between `Picost` configurations satisfying a transfer property, associated with the ability of processes to interact with their peers. We thereby automatically get a co-inductive proof methodology for establishing relationships between configurations.

In fact, referring to (3.1), it is better to move away from terminology such as *efficiency* as the interpretation depends very much on the nature of the units being recorded. In Example 2.6 these are *costs* and in such a scenario it is reasonable to interpret (3.1) as saying $(\Delta \triangleright N)$ is an improvement on $(\Gamma \triangleright M)$ as it potentially involves less cost. On the other hand in Example 2.8 the units are *profit* (for the publisher), and here $(\Gamma \triangleright M)$ would be considered to be an improvement on $(\Delta \triangleright N)$, as there is potential for more profit (for the publisher).

We therefore move to the more neutral terminology of *weights*. However we can not simply base the formulation of (3.1) on the relative weight associated with each individual action, as the following example shows.

Example 3.1 (Amortising costs). Consider the simple system

$$\text{UD} \Leftarrow [\text{rec } x. \text{up}!. \text{down}!.x]_{\circ}$$

and let Γ_{25} be an environment in which the unique owner \circ has unlimited funds, the use of **up** costs 2 and the use of **down** costs 5. If we compare $(\Gamma_{25} \triangleright \text{UD})$ with $(\Gamma_{42} \triangleright \text{UD})$, where Γ_{42} is defined analogously, then intuitively the latter is more efficient than the former, despite the fact that in the latter the action **up** is more expensive; this is compensated for by the relative costs of the other action **down**. ■

The remainder of this section is divided into three subsections. In the first we present a theory of amortised weighted bisimulations, based on so-called *weighted labelled transition systems*, wLTSs. This gives rise to a parametrised behavioural preorder, which we call the *amortised weighted bisimulation preorder*. The aim is to apply this theory to **Picost**; with this in mind, in the second subsection we present a (detailed) labelled transition semantics for **Picost**, and show that it is in agreement with the reduction semantics given in Figure 2. In the third section we show how this automatically generates a wLTS, which in turn gives us an amortised weighted bisimulation preorder between **Picost** configurations. We demonstrate the usefulness of the resulting proof methodology by re-examining the examples from Section 2.4.

3.1. Amortised weighted bisimulations: Here we generalise the concepts of [KAK05]; our aim is to apply them to **Picost** but our formulation is at a more abstract level.

Definition 3.2 (Weighted labelled transition systems). An *weighted labelled transition system* or wLTS is a 4-tuple $\langle S, \text{Act}_{\tau}, W, \longrightarrow \rangle$ where S is a set of states, W set of weights, and $\longrightarrow \subseteq S \times \text{Act}_{\tau} \times W \times S$. Here Act_{τ} denotes a set of action names **Act** to which is added an extra distinct name τ which will represent internal action. We normally write $s \xrightarrow{\mu}_w s'$ to mean $(s, \mu, w, s') \in \longrightarrow$. As a default we take the set of weights to be \mathbb{Z} , the set of integers, both negative and positive. ■

A wLTS is called *standard* whenever there is a cost function $\text{weight} : \text{Act} \rightarrow W$ with the property that $s \xrightarrow{a}_w s'$ if and only if $w = \text{weight}(a)$ for every $a \in \text{Act}$. So in a standard wLTS there is a unique weight associated with external actions, although internal actions may have multiple possible associated weights, reflecting the different ways in which these actions may be generated from external moves. The wLTS which we will (eventually) generate for **Picost** will be standard, but the development below will not require that we are working with standard wLTSs.

Relative to a given wLTS *weak moves* are generated in the standard manner, although the associated weights need to be accumulated: $s \xRightarrow{\mu}_w s'$ is the least relation satisfying:

- $s \xrightarrow{\mu}_w s'$ implies $s \xRightarrow{\mu}_w s'$
- $s \xrightarrow{\mu}_w s'', s'' \xrightarrow{\tau}_v s'$ implies $s \xRightarrow{\mu}_{(w+v)} s'$
- $s \xrightarrow{\tau}_w s'', s'' \xrightarrow{\mu}_v s'$ implies $s \xRightarrow{\mu}_{(w+v)} s'$

We also use a variation on the standard notation $s \xrightarrow[\mu]{w} t$ from [Mil89]; when μ is any action other than τ this denotes $s \xrightarrow{\mu} t$, but when it is τ it means either that $s \xrightarrow{\tau} t$ or that s is t and $w = 0$.

Definition 3.3 (Amortised weighted bisimulations). A family of relations $\{\mathcal{R}^n \mid n \in \mathbb{N}\}$ over the states in a wLTS is called an *amortised weighted bisimulation* whenever $s \mathcal{R}^n t$:

- (i) $s \xrightarrow{\mu} s'$ implies $t \xrightarrow[\mu]{w} t'$ for some t', w such that $s' \mathcal{R}^{(n+v-w)} t'$
- (ii) conversely, $t \xrightarrow{\mu} t'$ implies $s \xrightarrow[\mu]{w} s'$ for some s', v such that $s' \mathcal{R}^{(n+v-w)} t'$ ■

Here the parametrisation with respect to \mathbb{N} puts an extra requirement on the standard *transfer properties* associated with bisimulations. In (i) and (ii) above the index $(n+v-w)$ must be in \mathbb{N} , that is must be non-negative. So for example if the amortisation n is 0 then v , the weight of the left hand action, must be greater than or equal to w , the weight of the right hand action. For this reason a standard bisimulation, which ignores the weights, may not be an amortised weighted bisimulation. But the more general effect of the parameter n in the definition is to allow a relaxation in the comparison between the actual weights of the actions in the processes being compared; this point is explained in detail in Example 3.6.

We can mimic the standard development of bisimulations and write $s \sqsubseteq_{\text{wgt}}^m s'$ to say that there is some amortised bisimulation $\{\mathcal{R}^n \mid n \in \mathbb{N}\}$ such that $s \mathcal{R}^m s'$. Weighted bisimulations are (point-wise) closed under unions, and therefore we can mimic the standard development of *bisimulation equivalence*, [Mil89], to obtain the following:

Proposition 3.4.

- (a) The family of relations $\{\sqsubseteq_{\text{wgt}}^n \mid n \in \mathbb{N}\}$ is an amortised weighted bisimulation.
- (b) This family is the largest (point-wise) amortised weighted bisimulation.
- (c) If $s \sqsubseteq_{\text{wgt}}^m t$ and $s \xrightarrow{\mu} s'$ then $t \xrightarrow[\mu]{w} t'$ for some t', v such that $s' \sqsubseteq_{\text{wgt}}^{(m+v-w)} t'$.

Proof. Straightforward, using standard techniques. □

When we are uninterested in the exact amortisation used we write simply $s \sqsubseteq_{\text{wgt}} t$, meaning that there is some $k \geq 0$ such that $s \sqsubseteq_{\text{wgt}}^k t$, and we refer to this preorder as the *amortised weighted bisimulation preorder*.

Proposition 3.5.

- (a) The relations $\sqsubseteq_{\text{wgt}}^n$ are reflexive
- (b) $s_1 \sqsubseteq_{\text{wgt}}^m s_2, s_2 \sqsubseteq_{\text{wgt}}^n s_3$ implies $s_1 \sqsubseteq_{\text{wgt}}^{(m+n)} s_3$
- (c) $\sqsubseteq_{\text{wgt}}^m \subseteq \sqsubseteq_{\text{wgt}}^n$ whenever $m \leq n$.

Proof. In each case it is sufficient to exhibit a suitable amortised weighted bisimulation, that is a suitable family of relations over states. For example to prove (b) we let \mathcal{R}^k , for $k \geq 0$, be the set of pairs $\langle s_1, s_2 \rangle$ such that $s_1 \sqsubseteq_{\text{wgt}}^n s_3$ and $s_3 \sqsubseteq_{\text{wgt}}^m s_2$ for some state s_3 and some numbers n, m such that $k = n + m$.

To show $\{\mathcal{R}^k \mid k \in \mathbb{N}\}$ is an amortised weighted bisimulation let us suppose $s_1 \mathcal{R}^k s_2$ and $s_1 \xrightarrow{\mu} s'_1$; we have to prove

$$s_2 \xrightarrow[\mu]{w} s'_2 \text{ for some } s'_2 \text{ satisfying } s'_1 \mathcal{R}^{(k+v-w)} s'_2 \quad (3.2)$$

(The proof of the symmetric requirement is similar.)

- (i) From $s_1 \sqsubseteq_{\text{wgt}}^n s_3$ we know $s_3 \xrightarrow[\mu]{u} s'_3$ such that $s'_1 \sqsubseteq_{\text{wgt}}^{(n+v-u)} s'_3$

(ii) From $s_3 \sqsubseteq_{\text{wgt}}^m s_2$, and the final part of the previous Proposition, we know $s_2 \xrightarrow{w}^{\hat{\mu}} s'_2$ such that $s'_3 \sqsubseteq_{\text{wgt}}^{(m+u-w)} s'_2$.

But since $(n+v-u) + (m+u-w) = (k+v-w)$ we have $s'_1 \mathcal{R}^{(k+v-w)} s'_2$ and the requirement (3.2) follows.

The proof of part (c) is similar using the family of relations $\{\mathcal{R}^n \mid n \in \mathbb{N}\}$, where $s \mathcal{R}^n t$ whenever $s \sqsubseteq_{\text{wgt}}^m t$ for some $m \leq n$, while the proof of part (a) uses the family where each \mathcal{R}^n is the identity relation. \square

Example 3.6 (Amortising costs continued). Here we continue with Example 3.1. Shortly we will see a systematic way of associating weights with actions in Picost. But informally we can simply say

$$\mathcal{C}_{25} \xrightarrow{2}^{\text{up}!} \mathcal{D}_{25} \xrightarrow{5}^{\text{down}!} \mathcal{C}_{25}$$

where \mathcal{C}_{25} , \mathcal{D}_{25} are abbreviations for the configurations $(\Gamma_{25} \triangleright \text{UD})$ and $(\Gamma_{25} \triangleright [\text{down!.rec } x. \text{up!.down!.x}]_o)$ respectively, and analogously for $(\Gamma_{42} \triangleright \text{UD})$. Then relative to this induced wLTS we can show that the following is a weighted bisimulation:

$$\mathcal{R}^n = \{\langle \mathcal{D}_{25}, \mathcal{D}_{42} \rangle\} \cup \{\langle \mathcal{C}_{25}, \mathcal{C}_{42} \rangle \mid n \geq 2\}$$

It follows that

$$(\Gamma_{25} \triangleright \text{UD}) \sqsubseteq_{\text{wgt}}^2 (\Gamma_{42} \triangleright \text{UD})$$

However $(\Gamma_{42} \triangleright \text{UD}) \not\sqsubseteq_{\text{wgt}}^k (\Gamma_{25} \triangleright \text{UD})$ for any k . To see this suppose $\{\mathcal{R}^n \mid n \geq 0\}$ is a weighted bisimulation; we prove by induction on k that

$$\begin{aligned} \langle \mathcal{D}_{42}, \mathcal{D}_{25} \rangle &\notin \mathcal{R}^{(k+2)} \\ \langle \mathcal{C}_{42}, \mathcal{C}_{25} \rangle &\notin \mathcal{R}^k \end{aligned} \tag{3.3}$$

First notice that the pair $\langle \mathcal{D}_{42}, \mathcal{D}_{25} \rangle$ can not be in \mathcal{R}^2 ; this is because the move $\mathcal{D}_{42} \xrightarrow{2}^{\text{down}!} \mathcal{C}_{42}$ can not be matched by a move $\mathcal{D}_{42} \xrightarrow{w}^{\text{down}!} \mathcal{C}_{42}$ such that $\mathcal{C}_{42} \mathcal{R}^{(2+2-w)} \mathcal{C}_{25}$. The only possible candidate is the move $\mathcal{D}_{42} \xrightarrow{5}^{\text{down}!} \mathcal{C}_{42}$ and \mathcal{R}^{-1} does not exist.

From this fact it follows immediately that the pair $\langle \mathcal{C}_{42}, \mathcal{C}_{25} \rangle$ can not be in \mathcal{R}^0 ; for matching the move $\mathcal{C}_{42} \xrightarrow{4}^{\text{up}!} \mathcal{D}_{42}$ would require the impossible, that $\langle \mathcal{D}_{42}, \mathcal{D}_{25} \rangle$ be \mathcal{R}^2 . In other words we have shown (3.3) in the case when $k = 0$.

Suppose it is true for k ; the proof that it follows for $(k+1)$ is also straightforward. This is because

- for $\langle \mathcal{D}_{42}, \mathcal{D}_{25} \rangle$ to be in $\mathcal{R}^{(k+3)}$ we would require that $\langle \mathcal{C}_{42}, \mathcal{C}_{25} \rangle$ be in $\mathcal{R}^{(k+3+2-5)}$ which contradicts the induction hypothesis
- for $\langle \mathcal{C}_{42}, \mathcal{C}_{25} \rangle$ to be in $\mathcal{R}^{(k+1)}$ we would require $\langle \mathcal{D}_{42}, \mathcal{D}_{25} \rangle$ to be in $\mathcal{R}^{(k+3)}$, which we have just shown not to be possible.

It is important that the set of natural numbers \mathbb{N} is used in Definition 3.3, or at least that the family of relations be parametrised relative to a well-founded order. If instead we allowed families of relations $\{R^z \mid z \in \mathbb{Z}\}$, where \mathbb{Z} is the set of all integers, positive and negative, then $(\Gamma_{42} \triangleright \text{UD}) \sqsubseteq_{\text{wgt}}^0 (\Gamma_{25} \triangleright \text{UD})$ would follow. Simply letting $\mathcal{R}^z = \{\langle \mathcal{C}_{42}, \mathcal{C}_{25} \rangle, \langle \mathcal{D}_{42}, \mathcal{D}_{25} \rangle\}$ for every $z \in \mathbb{Z}$, we would obtain an extended family of relations trivially satisfying the requirements in Definition 3.3. Indeed in general, using \mathbb{Z} in place of \mathbb{N} , there would be no

difference between amortised weighted bisimulations and standard bisimulations (where all weights are ignored). \blacksquare

3.2. An operational semantics for Picost. As a first step in applying the theory of *amortised weighted bisimulations* to Picost we give an operational semantics for the language in terms of a (standard) LTS.

In Figure 6 and Figure 7 we give a set of rules for deriving judgements of the form

$$(\Gamma \triangleright M) \xrightarrow{\lambda} (\Delta \triangleright N),$$

where λ can take one of the forms

- (i) internal action, τ
- (ii) input, $(u, (\tilde{r}:\tilde{R})a?v, p)$: input by resource a of a known or fresh name, or value, where p is the provider of the resource and u the user
- (iii) output: $(u, (\tilde{r}:\tilde{R})a!v, p)$: delivery of a known or fresh name, to resource a , where again p is the provider of the resource and u the user.

We restrict attention to well-formed λ , that is, in the input and output actions each r_i must occur somewhere in v , and applications of the rules must preserve well-formedness. However note that because Picost only uses unary communication the vectors $(\tilde{r}), (\tilde{b})$ will have length either 0 or 1.

The rules are inherited directly from the corresponding ones for Dpi, [Hen07], and for the sake of clarity obvious symmetric rules, such as for (L-COMM) and (L-CNTX), are omitted; *Barendregt's convention* is also liberally applied, for example in omitting side-conditions to (L-CNTX). The only point of interest is the use of the preconditions $\Gamma \xrightarrow{(o_1, a, o_2)} \Delta$ in (L-IN) and (L-OUT); communication is only deemed to be possible if it can be paid for in some manner. Note that u in (L-IN), and p in (L-OUT) are free meta-variables. So for example the simple process $[a!\langle v \rangle.P]_o$ can perform the actions $[a!\langle v \rangle.P]_o \xrightarrow{(o, a!v, o')} \Delta \triangleright [P]_o$ for every owner $o' \in \text{Own}$ such that $\Gamma \xrightarrow{(o, a, o')} \Delta$. Also in the communication rule (L-COMM) any new resources used in the communication, $\tilde{r} : \tilde{R}$ remain private but in general the resulting cost environment Δ will be different from Γ ; the internal communication involves the use of a resource, and the change from Γ to Δ will reflect the associated costs.

We can perform a number of sanity checks on these rules. For example one can show that if $(\Gamma_1 \triangleright P_1) \xrightarrow{(b:\tilde{R})\alpha} (\Gamma_2 \triangleright P_2)$ then $\Gamma_2 = \Delta, b:\tilde{R}$ for some Δ such that $\Gamma_1 \xrightarrow{(u, a, p)} \Delta$, for some u, p , where a is the channel used in α ; a more detailed analysis of the possible judgements is given in the two lemmas below. The actions also preserve configurations:

Proposition 3.7. *If $(\Gamma_1 \triangleright M_1)$ is a configuration and $(\Gamma_1 \triangleright M_1) \xrightarrow{\lambda} (\Gamma_2 \triangleright M_2)$ then $(\Gamma_2 \triangleright M_2)$ is also a configuration.*

Proof. A straightforward induction on the inference of the judgements. \square

We also have a consistency check with respect to the reduction semantics of Section 2, stated in the theorem below; the proof requires two technical lemmas.

Lemma 3.8 (Deriv-output). *Suppose $\Gamma \triangleright M \xrightarrow{(u, (\tilde{r}:\tilde{R})a!v, p)} \Delta \triangleright N$. Then*

- (i) $\Delta = (\Gamma', \tilde{r}:\tilde{R})$ for some Γ'

$$\begin{array}{c}
\text{(L-IN)} \\
\frac{\Gamma \xrightarrow{(u,a,o)} \Delta}{\Gamma \triangleright [a?(x).P]_{\circ} \xrightarrow{(u,a?v,o)} \Delta \triangleright [P\{v/x\}]_{\circ}} \quad v \in \text{dom}(\Gamma^u) \text{ or } v \text{ not a channel} \\
\\
\frac{\Gamma \xrightarrow{(u,a,o)} \Delta}{\Gamma \triangleright [a?(x).P]_{\circ} \xrightarrow{(u,(b:\mathbb{R})a?b,o)} \Delta, b : \mathbb{R} \triangleright [P\{b/x\}]_{\circ}} \quad b \notin \text{dom}(\Gamma^u) \\
\\
\text{(L-OUT)} \\
\frac{\Gamma \xrightarrow{(o,a,p)} \Delta}{\Gamma \triangleright [a!\langle v \rangle.P]_{\circ} \xrightarrow{(o,a!v,p)} \Delta \triangleright [P]_{\circ}} \\
\\
\text{(L-COMM)} \\
\frac{\Gamma \triangleright M \xrightarrow{(u,(\tilde{r}:\tilde{\mathbb{R}})a?v,p)} \Delta, \tilde{r}:\tilde{\mathbb{R}} \triangleright M', \quad \Gamma \triangleright N \xrightarrow{(u,(\tilde{r}:\tilde{\mathbb{R}})a!v,p)} \Delta, \tilde{r}:\tilde{\mathbb{R}} \triangleright N'}{\Gamma \triangleright M \mid N \xrightarrow{\tau} \Delta \triangleright (\text{new } \tilde{r}:\tilde{\mathbb{R}})(M' \mid N')}
\end{array}$$

Figure 6: An action semantics for Picost: main rules

$$\begin{array}{c}
\text{(L-OPEN)} \\
\frac{\Gamma, b:\mathbb{R} \triangleright M \xrightarrow{(u,a!b,p)} \Gamma' \triangleright M'}{\Gamma \triangleright (\text{new } b:\mathbb{R})M \xrightarrow{(u,(b:\mathbb{R})a!b,p)} \Gamma' \triangleright M'} \quad a \neq b \\
\\
\text{(L-EXPORT)} \\
\Gamma \triangleright [(\text{new } r:\mathbb{R})P]_{\circ} \xrightarrow{\tau} \Gamma \triangleright (\text{new } r:\mathbb{R})[P]_{\circ} \\
\\
\text{(L-SPLIT)} \\
\Gamma \triangleright [M \mid N]_{\circ} \xrightarrow{\tau} \Gamma \triangleright [M]_{\circ} \mid [N]_{\circ} \\
\\
\text{(L-UNWIND)} \\
\Gamma \triangleright [\text{rec } x. T]_{\circ} \xrightarrow{\tau} \Gamma \triangleright [T\{\text{rec } x. T/x\}]_{\circ} \\
\\
\text{(L-MATCH)} \\
\frac{}{\Gamma \triangleright [\text{if } a = a \text{ then } P \text{ else } Q]_{\circ} \xrightarrow{\tau} \Gamma \triangleright [P]_{\circ}} \\
\\
\text{(L-MISMATCH)} \\
\frac{}{\Gamma \triangleright [\text{if } a = b \text{ then } P \text{ else } Q]_{\circ} \xrightarrow{\tau} \Gamma \triangleright [Q]_{\circ}} \quad a \neq b \\
\\
\text{(L-CNTX)} \\
\frac{\Gamma \triangleright M \xrightarrow{\lambda} \Gamma' \triangleright M'}{\Gamma \triangleright M \mid N \xrightarrow{\lambda} \Gamma' \triangleright M' \mid N} \\
\\
\text{(L-CNTX)} \\
\frac{\Gamma, b:\mathbb{R} \triangleright M \xrightarrow{\lambda} \Gamma', b:\mathbb{R} \triangleright M'}{\Gamma \triangleright (\text{new } b:\mathbb{R})M \xrightarrow{\lambda} \Gamma' \triangleright (\text{new } b:\mathbb{R})M'} \quad b \notin \text{n}(\lambda)
\end{array}$$

Figure 7: An action semantics for Picost: more rules

- (ii) $\Gamma \xrightarrow{(u,a,p)} \Gamma'$
- (iii) $M \equiv (\text{new } \tilde{r}:\tilde{\mathbb{R}})(M' \mid [a!\langle v \rangle.Q]_{\text{u}})$
- (iv) $N \equiv (M' \mid [Q]_{\text{u}})$
- (v) $\Theta \triangleright M \xrightarrow{(u,(\tilde{r}:\tilde{\mathbb{R}})\alpha,p')} \Theta', \tilde{r}:\tilde{\mathbb{R}} \triangleright N$ whenever $\Theta \xrightarrow{(u,a,p')} \Theta'$, for any owner p' .

Proof. By induction on the derivation of $\Gamma \triangleright M \xrightarrow{(u, (\tilde{r}:\tilde{R})a!v, \rho)} \Delta \triangleright N$. \square

Lemma 3.9 (Deriv-input). *Suppose $\Gamma \triangleright M \xrightarrow{(u, (\tilde{r}:\tilde{R})a?v, \rho)} \Delta \triangleright N$. Then*

- (i) $\Delta = (\Gamma', \tilde{r}:\tilde{R})$ for some Γ'
- (ii) $\Gamma \xrightarrow{(u, a, \rho)} \Gamma'$
- (iii) $M \equiv (\text{new } \tilde{c}:\mathbb{C})([a?(x).T]_{\rho} \mid M')$
- (iv) $N \equiv (\text{new } \tilde{c}:\mathbb{C})([T\{v/x\}]_{\rho} \mid M')$
- (v) $\Theta \triangleright M \xrightarrow{(u', (\tilde{r}:\tilde{R}')\alpha, \rho)} \Theta', \tilde{r}:\tilde{R}' \triangleright N$ whenever $\Theta \xrightarrow{(u', a, \rho)} \Theta'$, for any owner u' , and types (\tilde{R}') .

Proof. Again a straightforward induction on the derivation $\Gamma \triangleright M \xrightarrow{(u, (\tilde{r}:\tilde{R})a?v, \rho)} \Delta \triangleright N$. Note that in part (v) arbitrary types (\tilde{R}') can be used because there is no restriction on the type R in the second part of the rule (L-IN) in Figure 6. \square

Theorem 3.10. $\Gamma \triangleright M \longrightarrow \Delta \triangleright N$ if and only if $\Gamma \triangleright M \xrightarrow{\tau} \Delta \triangleright N'$ for some N' such that $N \equiv N'$.

(Outline). First we need to show the auxiliary result that structural equivalence is preserved by actions. That is $\Gamma \triangleright M \xrightarrow{\lambda} \Delta \triangleright M'$ and $M \equiv N$ implies $\Gamma \triangleright N \xrightarrow{\lambda} \Delta \triangleright N'$ for some N' such that $M' \equiv N'$; this is proved by induction on the proof of the fact that $M \equiv N$ from the rules in Figure 3. Then a straightforward proof by induction on the derivation of $\Gamma \triangleright M \longrightarrow \Delta \triangleright N$ from the rules in Figure 2 will show that this implies $\Gamma \triangleright M \xrightarrow{\tau} \Delta \triangleright N'$ with $N \equiv N'$; the auxiliary result is required when considering the rule (R-STRUCT).

To prove the converse we also employ the two previous lemmas, giving the structure of input and output actions. Suppose $\Gamma \triangleright M \xrightarrow{\tau} \Delta \triangleright N$; we prove by rule induction that $\Gamma \triangleright M \longrightarrow \Delta \triangleright N$. The only non-trivial case is when this judgement is inferred using the rule (L-COMM), or its dual. So without loss of generality we know

- $M = M_1 \mid M_2$
- $N = (\text{new } \tilde{r}:\tilde{R})(N_1 \mid N_2)$
- $\Gamma \triangleright M_1 \xrightarrow{(u, (\tilde{r}:\tilde{R})a?v, \rho)} \Delta, \tilde{r}:\tilde{R} \triangleright N_1$
- $\Gamma \triangleright M_2 \xrightarrow{(u, (\tilde{r}:\tilde{R})a!v, \rho)} \Delta, \tilde{r}:\tilde{R} \triangleright N_2$

The previous two lemmas can now be applied to obtain the structure of M_1 , M_2 , N_1 and N_2 , up to structural equivalence; by rearranging $M_1 \mid M_2$, again using the structural equivalence rules, an application of (R-COMM) followed by one of (R-STRUCT) gives the required $\Gamma \triangleright M \longrightarrow \Delta \triangleright N$. \square

3.3. A proof methodology for Picost. The operational semantics given in the previous subsection can be used in a straightforward way to obtain a wLTS for Picost configurations. It suffices to attach a weight to the actions, which can be done in a systematic manner: we write

$$(\Gamma \triangleright M) \xrightarrow{\mu}_w (\Delta \triangleright N)$$

whenever

- $(\Gamma \triangleright M) \xrightarrow{\mu} (\Delta \triangleright N)$ can be deduced from the rules in Figure 6 and Figure 7

- $w = (\Delta^{\text{rec}} - \Gamma^{\text{rec}})$

Note that the weight associated with an action is ultimately determined by the manner in which expenditure is recorded in the cost environments; this may reflect the cost of providing the resource in question, as in Example 2.6, the profit to be gained by a particular owner in the use of the resource, as in Example 2.8, or combinations of such concerns.

We can now apply Definition 3.3 to this wLTS to obtain a family of preorders

$$(\Gamma \triangleright M) \sqsubseteq_{\text{wgt}}^n (\Delta \triangleright N) \quad (3.4)$$

between Picost configurations. However we must be somewhat careful here, as some of the actions used involve bound names; but by a systematic application of *Barendregt's convention*, mentioned on page 3, confusions between these and free names can be avoided.

As is well-known, the relations (3.4) come equipped with a powerful co-inductive proof methodology. In order to prove $(\Gamma \triangleright M) \sqsubseteq_{\text{wgt}}^k (\Delta \triangleright N)$ for a particular k it is sufficient to exhibit a family of relations $\{\mathcal{R}^n \mid n \in \mathbb{N}\}$ which satisfy the transfer properties of Definition 3.3, such that \mathcal{R}^k contains the pair $(\Gamma \triangleright M, \Delta \triangleright N)$. In the remainder of this section we apply this proof methodology to the examples in Section 2. This allows us to now reason about the behaviour of systems, how they interact with other systems, rather than reason simply about their computation runs.

Example 3.11 (Running a library, revisited). Referring to the definitions in Example 2.6, by exhibiting a witness weighted bisimulation it is possible to show

$$(\Gamma_{\text{central}} \triangleright [\text{Reader}]_{\text{pub}}) \sqsubseteq_{\text{wgt}}^0 (\Gamma_{\text{local}} \triangleright [\text{Reader}]_{\text{pub}})$$

This is despite the fact that the local use of the service `reqR` is more expensive than the central use; this is compensated for by the fact that both `goLib` and `goHome` are less expensive locally. It is also worth noting that although the *use* of resources in both Γ_{central} and Γ_{local} is free, in the generated wLTS the output actions actually have non-zero weights associated with them. For example, a typical run in this wLTS from $(\Gamma_{\text{central}} \triangleright [\text{Reader}]_{\text{pub}})$ takes the form

$$(\Gamma_{\text{central}} \triangleright [\text{Reader}]_{\text{pub}}) \xrightarrow{\text{goLib}?n}_{\rightarrow 5} \dots \xrightarrow{(r)\text{reqR}!(r,n)}_{\rightarrow 1} \dots \xrightarrow{\text{goHome}!b}_{\rightarrow 5} \dots$$

whereas the corresponding local run is

$$(\Gamma_{\text{local}} \triangleright [\text{Reader}]_{\text{pub}}) \xrightarrow{\text{goLib}?n}_{\rightarrow 1} \dots \xrightarrow{(r)\text{reqR}!(r,n)}_{\rightarrow 3} \dots \xrightarrow{\text{goHome}!b}_{\rightarrow 1} \dots$$

To compare the efficiency of the library service itself we consider the following definitions

$$\begin{aligned} \text{Lib}_{\text{local}} &\Leftarrow (\text{new reqS}:\mathbf{R}_s^l)([\text{Library} \mid \text{Store}]_{\text{lib}}) \\ \text{Lib}_{\text{central}} &\Leftarrow (\text{new reqS}:\mathbf{R}_s^c)([\text{Library} \mid \text{Store}]_{\text{lib}}) \end{aligned}$$

where, as explained in Example 2.6, \mathbf{R}_s^l , \mathbf{R}_s^c , are the types $\langle 0, 5 \rangle$, $\langle 0, 1 \rangle$ respectively; here the interaction between the library and the store has been internalised, with types reflecting the relative cost of local and central access. Both these configurations simply *provide* the service `reqR`, and viewed in isolation the local service is not more efficient than the central one; no matter what n we choose, we have

$$(\Gamma_{\text{central}} \triangleright \text{Lib}_{\text{central}}) \not\sqsubseteq_{\text{wgt}}^n (\Gamma_{\text{local}} \triangleright \text{Lib}_{\text{local}}) \quad (3.5)$$

However if we combine the library service with the reader then the overall systems is locally more efficient than the centralised one:

$$(\Gamma_{\text{central}} \triangleright \text{Sys}_{\text{central}}) \sqsubseteq_{\text{wgt}}^2 (\Gamma_{\text{local}} \triangleright \text{Sys}_{\text{local}}) \quad (3.6)$$

where

$$\begin{aligned} \text{Sys}_{\text{local}} &\Leftarrow (\text{new reqR:R}_r^l)([\text{Reader}]_{\text{pub}} \mid \text{Lib}_{\text{local}}) \\ \text{Sys}_{\text{central}} &\Leftarrow (\text{new reqR:R}_r^c)([\text{Reader}]_{\text{pub}} \mid \text{Lib}_{\text{central}}) \end{aligned}$$

We should point out that in (3.5) and (3.6) we have used the full cost environments Γ_{local} , Γ_{central} , despite the fact that some of the resources have been restricted in the systems; this is simply in order to avoid the definition of even more environments.

As an example of how such statements can be proved see the Section A.1 in the appendix for a witness bisimulation which establishes (3.6). \blacksquare

4. CONTEXTUAL CHARACTERISATION

In the previous section we have demonstrated that the preorders $\sqsubseteq_{\text{wgt}}^n$ provide a useful co-inductive methodology for comparing the behaviour of processes, relative to resource costs. In this section we critically review its formulation, revealing some significant inadequacies, and offer a revised version where these are addressed.

Informally we would expect at least the following two properties of a proof methodology:

- (a) It should support *compositional reasoning*, whereby the analysis of process behaviour can be carried out structurally.
- (b) *Soundness*: Any relationship established between the behaviour of processes using the proof methodology should be justifiable in some independent manner.

Further we could hope for:

- (c) *Completeness*: any pair of processes which are intuitively behaviourally related, should be provably related using our methodology.

Relative to our language **Picost** the first criteria, (a), is straightforward to formalise, as a property of the preorders $\sqsubseteq_{\text{wgt}}^n$.

Definition 4.1 (Compositional). A relation \mathcal{R} over **Picost** configurations is said to be *compositional* whenever $(\Gamma \triangleright M) \mathcal{R} (\Delta \triangleright N)$ implies

- (i) $(\Gamma \triangleright M \mid O) \mathcal{R}^m (\Delta \triangleright N \mid O)$, provided $(\Gamma \triangleright M \mid O)$ and $(\Delta \triangleright N \mid O)$ are configurations
- (ii) $(\Gamma, r:\mathbf{R} \triangleright M) \mathcal{R}^m (\Delta, r:\mathbf{R} \triangleright N)$. \blacksquare

We could of course demand that the relation \mathcal{R} should be preserved by all the operators in the language, but for the purposes of the discussion to follow it is sufficient to concentrate on the two most important ones.

Our first remark is that the relations $\sqsubseteq_{\text{wgt}}^n$ are not compositional, and therefore our proposed proof methodology does not support compositional reasoning.

Example 4.2 (Non-compositionality). Let Γ be a cost environment with two owners $\mathfrak{o}, \mathfrak{p}$ and two resources a, b . Suppose further that $\Gamma^{\mathfrak{o}}(\mathfrak{o}) = \Gamma^{\mathfrak{p}}(\mathfrak{p}) = \infty$, while $\Gamma^u(a) = 20$, $\Gamma^u(b) = 10$; the remaining fields in Γ are unimportant, but to be definite let us say that $\Gamma^{\mathfrak{p}}(a) = \Gamma^{\mathfrak{p}}(b) = 0$. Let Δ be another cost environment with the same resources, with both usage costs being 10, and the same owners, but with the difference that $\Delta^{\mathfrak{o}}(\mathfrak{o}) = 10$. Then it is easy to check that

$$\Gamma \triangleright [a!]_{\mathfrak{o}} \sqsubseteq_{\text{wgt}}^0 \Delta \triangleright [a!]_{\mathfrak{o}}$$

However one can also show that

$$\Gamma \triangleright [a!]_{\mathfrak{o}} \mid [b!]_{\mathfrak{o}} \not\leq_{\text{wgt}}^0 \Delta \triangleright [a!]_{\mathfrak{o}} \mid [b!]_{\mathfrak{o}}$$

The problem occurs when we consider the action $(\Gamma \triangleright [a!]_{\mathfrak{o}} \mid [b!]_{\mathfrak{o}}) \xrightarrow{(\mathfrak{o}, b!, \mathfrak{p})}_{10} (\Gamma_1 \triangleright [a!]_{\mathfrak{o}} \mid [\text{stop}]_{\mathfrak{o}})$. This can be matched by the action $(\Delta \triangleright [a!]_{\mathfrak{o}} \mid [b!]_{\mathfrak{o}}) \xrightarrow{(\mathfrak{o}, b!, \mathfrak{p})}_{10} (\Delta_1 \triangleright [a!]_{\mathfrak{o}} \mid [\text{stop}]_{\mathfrak{o}})$ but at the expense of exhausting all of \mathfrak{o} 's funds. $\Delta_1^{\mathfrak{o}}(\mathfrak{o})$ is now set to 0 and therefore the action $(\Gamma_1 \triangleright [a!]_{\mathfrak{o}} \mid [\text{stop}]_{\mathfrak{o}}) \xrightarrow{(\mathfrak{o}, a!, \mathfrak{p})}_{20} (\Gamma_1 \triangleright [\text{stop}]_{\mathfrak{o}} \mid [\text{stop}]_{\mathfrak{o}})$ can not be matched by any action from $(\Delta_1 \triangleright [a!]_{\mathfrak{o}} \mid [\text{stop}]_{\mathfrak{o}})$.

The other criteria, (b) and (c) above, are more difficult to formalise. But even in the absence of a precise formalisation we can also show that our proof methodology runs into difficulties with them, by considering a proposed touchstone family of preorders $\sqsubseteq_{\text{behav}}^n$, $n \geq 0$, which incorporate some intuitive properties which we would expect. First an easy example, essentially taken from [HR04].

Example 4.3 (Problem with output types). Consider the two configurations \mathcal{C} and \mathcal{D} , denoted by

$$\Gamma \triangleright (\text{new } r : \mathbf{R}_1)([a!\langle r \rangle. \text{stop}]_{\mathfrak{o}}), \quad \Gamma \triangleright (\text{new } r : \mathbf{R}_2)([a!\langle r \rangle. \text{stop}]_{\mathfrak{o}})$$

respectively, where the types \mathbf{R}_1 , \mathbf{R}_2 are different, and Γ has sufficient resources for a to be exercised; that is $\Gamma \xrightarrow{(\mathfrak{o}, a, \mathfrak{p})} \Gamma'$ for some owner \mathfrak{p} and some Γ' .

Then it is easy to see that $\mathcal{C} \not\leq_{\text{wgt}}^k \mathcal{D}$ for any k because the only actions which the configurations can perform are different; they are labelled $(\mathfrak{p}, (r : \mathbf{R}_1)a!r, \mathfrak{o})$ and $(\mathfrak{p}, (r : \mathbf{R}_2)a!r, \mathfrak{o})$ respectively.

However it is difficult to envisage any context in which these two configurations can be distinguished; for any reasonable definition of the touchstone relations we would expect $\mathcal{C} \sqsubseteq_{\text{behav}}^k \mathcal{D}$ to be true. Thus our proof methodology will not be *complete*. ■

Our next example focuses on some of the novel features of Picost.

Example 4.4 (Problem with owner identification). Let \mathcal{C} , \mathcal{D} denote the configurations

$$\Gamma \triangleright [a!]_{\mathfrak{o}_1}, \quad \Gamma \triangleright [a!]_{\mathfrak{o}_2}$$

respectively, where \mathfrak{o}_1 , \mathfrak{o}_2 are two different owners, and $\Gamma^{\mathfrak{o}}(\mathfrak{o}_1) = \Gamma^{\mathfrak{o}}(\mathfrak{o}_2)$.

Here again we would expect $\mathcal{C} \sqsubseteq_{\text{behav}}^k \mathcal{D}$ to be true because there is no mechanism in Picost which would enable an observer to discover who was funding the use of the resource a . However assuming some owner \mathfrak{p} has sufficient funds in Γ to provide the resource a , we have $\mathcal{C} \not\leq_{\text{wgt}}^0 \mathcal{D}$ again because the configurations perform different actions, labelled $(\mathfrak{o}_1, a!, \mathfrak{p})$ and $(\mathfrak{o}_2, a!, \mathfrak{p})$ respectively. ■

4.1. Behavioural preorders. In order to address the inadequacies with our proof methodology let us first give one possible formalisation of the touchstone family of behavioural preorders which we have been referring to as $\sqsubseteq_{\text{behav}}^n$, $n \geq 0$; we adapt the theory of *reduction barbed congruences*, [HT92, SW01, HR04] to Picost, often referred to informally as *contextual equivalences*. For simplicity we assume that resource charging is always standard, and that the only values used are channel/resource names.

We first need to introduce into the reduction semantics some record of the costs being expended. Let us write $\Gamma \triangleright M \xrightarrow{c} \Delta \triangleright N$ whenever $\Gamma \triangleright M \longrightarrow \Delta \triangleright N$ can be deduced from

the reduction rules, in Figure 2, and $(\Delta^{\text{rec}} - \Gamma^{\text{rec}}) = c$. This is generalised in the obvious manner to $\Gamma \triangleright M \xrightarrow*_d \Delta \triangleright N$ by the accumulation of costs.

Definition 4.5 (Cost improving). We say that the family of relations $\{\mathcal{R}^n \mid n \in \mathbb{N}\}$ over configurations is *cost improving* whenever $\mathcal{C} \mathcal{R}^m \mathcal{D}$ for any m , then

- (i) $\mathcal{C} \xrightarrow*_c \mathcal{C}'$ implies $\mathcal{D} \xrightarrow*_d \mathcal{D}'$ such that $\mathcal{C}' \mathcal{R}^{(m+c-d)} \mathcal{D}'$
- (ii) conversely, $\mathcal{D} \xrightarrow*_d \mathcal{D}'$ implies $\mathcal{C} \xrightarrow*_c \mathcal{C}'$ such that $\mathcal{C}' \mathcal{R}^{(m+c-d)} \mathcal{D}'$. ■

This is a natural generalisation of the notion of *reduction closure* or *reduction bisimulation* from LTSs to weighted LTSs; for a justification of its use in defining behavioural preorders see Chapter 2 of [SW01].

Definition 4.6 (Observations). Let us write $(\Gamma \triangleright M) \Downarrow a?$ whenever $(\Gamma \triangleright M) \xrightarrow{*} (\Delta \triangleright N)$ where for some owner \circ

- (i) $N \equiv (\text{new } \tilde{c})([a?(x).T]_{\circ} \mid N')$, and a does not occur in (\tilde{c})
- (ii) $\Delta \xrightarrow{(u,a,\circ)} \Delta'$ for some u and Δ' .

The predicate $(\Gamma \triangleright M) \Downarrow a!$ is defined in an analogous manner. Note that here the owner \circ has to be able to pay the appropriate costs for the barb.

Then we say that the family of relations $\{\mathcal{R}^n \mid n \in \mathbb{N}\}$ over configurations *preserves observations* whenever, for any n , $\mathcal{C}_1 \mathcal{R}^n \mathcal{C}_2$ $\mathcal{C}_1 \Downarrow o$ if and only if $\mathcal{C}_2 \Downarrow o$. ■

Note that unlike [HG08] we do not record the cost of making observations; nor do we observe the owner responsible for the observation. This means that our notion of barb is more elementary.

Example 4.2 demonstrates that demanding a behavioural preorder to be compositional, in particular that it be preserved by arbitrary parallel contexts, is very problematic as intuitively it gives observers or external users of a system access to all the funds available to owners of the system. Here we address this issue by defining a relativised version of compositionality, relativised to the set of owners whose funds are available to external users.

Definition 4.7 (O-contextual). Let \mathbf{O} be a subset of the owners Own . A relation \mathcal{R} over Picost configurations is said to be *O-contextual* whenever $(\Gamma \triangleright M) \mathcal{R} (\Delta \triangleright N)$ implies

- (i) $(\Gamma \triangleright M \mid [P]_{\circ}) \mathcal{R} (\Delta \triangleright N \mid [P]_{\circ})$ for every $\circ \in \mathbf{O}$, provided $(\Gamma \triangleright M \mid [P]_{\circ})$ and $(\Delta \triangleright N \mid [P]_{\circ})$ are configurations.
- (ii) $(\Gamma, r:\mathbf{R} \triangleright M) \mathcal{R} (\Delta, r:\mathbf{R} \triangleright N)$. ■

Combining these three properties we obtain:

Definition 4.8 (The contextual improvement preorder). Let $\{\sqsubseteq_{\mathbf{O}:\text{cxt}}^n \mid n \in \mathbb{N}\}$ be the largest family (point-wise) of \mathbf{O} -contextual relations over configurations which preserves observations, and is cost improving. ■

The idea here is that we only consider the behaviour of systems relative to contexts in which observers, or users of the systems, can use code running under the financial authority of the owners in \mathbf{O} . At one extreme we can take \mathbf{O} to be the entire set of owners Own and then observers have access to all owners, and their funds; this gives Compositionality, as expressed in Definition 4.1. The other extreme is when observers have access to none of the owners users in the systems under observation; in this case the observers have to provide their own funds, to support observations.

We now set ourselves the task of modifying the proof methodology of Section 3.3 so that the informal properties (a), (b), and (c) are enforced, relative to the touchstone preorders $\sqsubseteq_{\mathbf{O}:cxt}^n$. First note that Example 4.4 and Example 4.3 still apply when the informal relations $\sqsubseteq_{\text{behav}}^n$ are instantiated by the formal $\sqsubseteq_{\mathbf{O}:cxt}^n$. But the problems presented in Example 4.2 depend on the choice of observers \mathbf{O} :

Example 4.9 (Unsoundness). Let Γ, Δ be as defined in Example 4.2. Then we have already argued that $\Gamma \triangleright [a!]_{\mathfrak{o}} \sqsubseteq_{\text{wgt}}^0 \Delta \triangleright [a!]_{\mathfrak{o}}$. Here we argue that $\Gamma \triangleright [a!]_{\mathfrak{o}} \not\sqsubseteq_{\mathbf{O}:cxt}^0 \Delta \triangleright [a!]_{\mathfrak{o}}$ whenever $\mathfrak{o} \in \mathbf{O}$. For otherwise, this would imply

$$\Gamma \triangleright [a!]_{\mathfrak{o}} \mid [P]_{\mathfrak{o}} \sqsubseteq_{\mathbf{O}:cxt}^0 \Delta \triangleright [a!]_{\mathfrak{o}} \mid [P]_{\mathfrak{o}}$$

for any process P which ensures that the configurations are still well-formed.

However for a contradiction take P to be $a?.(b! \mid b?.\omega!)$ where ω is some cost-free fresh channel. Then we can make the observation $\omega!$ on the left hand configuration but not on the right hand one. \blacksquare

This example shows that in general \mathbf{O} -observers can deplete the resources of any owner in \mathbf{O} , which is important if those owners have only finite funds. A significant consequence is given in the next proposition, which limits the applicability of this behavioural preorder for arbitrary \mathbf{O} .

Proposition 4.10. *If $(\Gamma \triangleright M) \sqsubseteq_{\mathbf{O}:cxt}^n (\Delta \triangleright N)$ for any n , then $\Gamma^{\mathfrak{o}}(\mathfrak{o}) = \Delta^{\mathfrak{o}}(\mathfrak{o})$ for every \mathfrak{o} in \mathbf{O} .*

Proof. Suppose $(\Gamma \triangleright M) \sqsubseteq_{\mathbf{O}:cxt}^n (\Delta \triangleright N)$ for some n , with \mathfrak{o} an owner in \mathbf{O} . We prove that $k \leq \Gamma^{\mathfrak{o}}(\mathfrak{o})$ if and only if $k \leq \Delta^{\mathfrak{o}}(\mathfrak{o})$.

Consider the process $O = [(\text{new } r:\mathbf{R})r! \mid r?.\omega!(\langle \rangle)]_{\mathfrak{o}}$, where ω is a fresh cost-free channel, where \mathbf{R} is the resource type $\langle k, 0 \rangle$; so r costs k to use but is free to provide. Then by compositionality we know

$$\Gamma, \omega : \mathbf{E} \triangleright M \mid O \sqsubseteq_{\mathbf{O}:cxt}^n \Delta, \omega : \mathbf{E} \triangleright N \mid O$$

where \mathbf{E} denotes the trivial type $\langle 0, 0 \rangle$.

If $k \leq \Gamma^{\mathfrak{o}}(\mathfrak{o})$, we have $\Gamma, \omega : \mathbf{E} \triangleright M \mid O \Downarrow \omega!$ and therefore, by the preservation of observations, $\Delta, \omega : \mathbf{E} \triangleright N \mid O \Downarrow \omega!$. But this is only possible if $k \leq \Delta^{\mathfrak{o}}(\mathfrak{o})$.

The converse argument is similar. \square

In effect this means that the behavioural preorders $\sqsubseteq_{\mathbf{O}:cxt}^n$ can not be used to differentiate between configurations in which owners from \mathbf{O} accrue different levels of funds; a typical case in point occurs with the systems in Example 2.9. For this reason we are primarily interested in the extreme case, when the observers have no access to the funds of the owners in the systems under investigation. Let us introduce some special notation for these situations.

Let \mathfrak{e} denote some arbitrary owner, intuitively taken to be external to the systems under observation. For an arbitrary cost environment Γ we use $\Gamma^{\mathfrak{e}}$ to denote the extended cost environment obtained by adding \mathfrak{e} to the domain of $\Gamma^{\mathfrak{o}}$ and setting $\Gamma^{\mathfrak{e}}(\mathfrak{e})$ to be ∞ ; in particular $\Gamma^{\mathfrak{e}}$ is only defined whenever \mathfrak{e} is new to the domain of $\Gamma^{\mathfrak{o}}$. Finally we use the notation

$$\Gamma \triangleright M \sqsubseteq_{\text{ext}}^n \Delta \triangleright N$$

as an abbreviation for

$$\Gamma^{\mathfrak{e}} \triangleright M \sqsubseteq_{\{\mathfrak{e}\}:cxt}^n \Delta^{\mathfrak{e}} \triangleright N$$

Here the observer has no access to the owners' resources used in the configurations \mathcal{C} , \mathcal{D} but has an infinite amount of resources with which to run experiments.

Our revised proof methodology is based on endowing *Picost* with the structure of a different, more abstract, wLTS, which takes into account the set of owners whose funds are available to observers, and employing Definition 3.3 to obtain a more abstract family of co-inductive preorders. In order to obtain our more abstract wLTS we forget some of the details in the labels of the actions of the operational semantics for *Picost*, given in Figure 6 and Figure 7, so that they reflect not what processes can do, but rather what external observers with access to the funds in \mathcal{O} can observe them doing. This leads to *abstract labels* of the following form, ranged over by μ :

- (a) internal label τ as before
- (b) input label $(u, (\tilde{r} : \tilde{R})a?v)$
- (c) output label $((\tilde{r})a!v, p)$

Here only one owner is recorded in the external actions; for input we note the user of the resource u while for output it is the producer p .

Definition 4.11 (*O*-actions). For each abstract label μ let the corresponding *O*-action $\mathcal{C} \xrightarrow{\mu}_w^{\mathcal{O}} \mathcal{D}$ be defined by

- (a) $(\Gamma_1 \triangleright M) \xrightarrow{\tau}_w^{\mathcal{O}} (\Gamma_2 \triangleright N)$ whenever $(\Gamma_1 \triangleright M) \xrightarrow{\tau} (\Gamma_2 \triangleright N)$ can be deduced from the rules, where $(\Gamma_2^{\text{rec}} - \Gamma_1^{\text{rec}}) = w$.
- (b) $(\Gamma_1 \triangleright M) \xrightarrow{((\tilde{r})a!b, p)}_w^{\mathcal{O}} (\Gamma_2 \triangleright N)$ whenever $p \in \mathcal{O}$ and $(\Gamma_1 \triangleright M) \xrightarrow{(\tilde{r} : \tilde{R})a!b, p} (\Gamma_2 \triangleright N)$ can be deduced from the rules for some (\tilde{R}) , and some owner u , where $(\Gamma_2^{\text{rec}} - \Gamma_1^{\text{rec}}) = w$.
- (c) $(\Gamma_1 \triangleright M) \xrightarrow{(u, (\tilde{r} : \tilde{R})a?v)}_w^{\mathcal{O}} (\Gamma_2 \triangleright N)$ whenever $u \in \mathcal{O}$ and $(\Gamma_1 \triangleright M) \xrightarrow{(u, (\tilde{r} : \tilde{R})a?v} (\Gamma_2 \triangleright N)$ can be deduced from the rules for some owner p , where $(\Gamma_2^{\text{rec}} - \Gamma_1^{\text{rec}}) = w$.

Note that in (a) the set of owners \mathcal{O} plays no role, but we leave it there for the sake of uniformity. ■

This endows *Picost* configurations with the structure of a more abstract wLTS, whose actions depend on the set of owners \mathcal{O} . We refer to this as the *O*-wLTS and we write $\mathcal{C} \sqsubseteq_{\mathcal{O}\text{wgt}}^n \mathcal{D}$ whenever there is an amortised weighted bisimulation $\{\mathcal{R}^n \mid n \in \mathbb{N}\}$ in this *O*-wLTS such that $\mathcal{C} \mathcal{R}^n \mathcal{D}$. When \mathcal{O} is the singleton set $\{e\}$ where the owner e is fresh, that is external to the configurations being compared, we abbreviate this to $\mathcal{C} \sqsubseteq_{\text{ewgt}}^n \mathcal{D}$.

Example 4.12 (*Publishing, revisited*). Here we use the notation and definitions from Example 2.8 and Example 2.9.

First we can compare the profits gained by running the publishing system in different cost environments. As before let Γ_{327} represent any cost environment of the form $\Gamma_{dyn}, \text{news} : R_n, \text{adv} : R_a, \text{publish} : R_p$, where these types are $\langle 3, 1 \rangle, \langle 2, 0 \rangle, \langle 7, 1 \rangle$ respectively, and let Γ_{216} be the same environment but with these types changed to $\langle 2, 1 \rangle, \langle 1, 0 \rangle, \langle 6, 1 \rangle$. Then it is straightforward to exhibit a witness bisimulation to establish

$$(\Gamma_{216} \triangleright [P]_p) \sqsubseteq_{\text{ewgt}}^0 (\Gamma_{327} \triangleright [P]_p)$$

Recall from Example 2.8 that in these cost environments we record the costs of the actions relative to their effect on the funds of p the publisher. So this means that that more profit can be gained by the publisher p by using the cost regime underlying the environment Γ_{216} .

To investigate the effect of implementing the *kickback* we consider the two systems

$$\begin{aligned} \text{PA} &\Leftarrow (\text{new adv:R}_a)([P]_{\text{p}} \mid [A]_{\text{a}}) \\ \text{PA}_K &\Leftarrow (\text{new adv:R}_a)([P_K]_{\text{p}} \mid [A_K]_{\text{a}}) \end{aligned}$$

Both these systems *use* the **news** resource and *provide* the **publish** resource. Here we can show, for example, that

$$(\Gamma_{327} \triangleright \text{PA}_K) \sqsubseteq_{\text{ewgt}}^0 (\Gamma_{327} \triangleright \text{PA})$$

provided $\Gamma_{327}^o(\text{p})$ is at least 5. See Section A.2 of the appendix for a description of a witness bisimulation. Again because of the way in which we have set up the accounting in the cost environments this means that the code PA_K is more profitable for the publisher than PA . ■

The abstract \mathcal{O} -wLTS has precisely enough information about actions to characterise the touchstone contextual behavioural preorder, at least in the extreme case of $\mathcal{O} = \{\text{e}\}$.

Theorem 4.13 (Full-abstraction, external case). *For every $n \in \mathbb{N}$, $(\Gamma \triangleright M) \sqsubseteq_{\text{ecxt}}^n (\Delta \triangleright N)$ if and only if $(\Gamma \triangleright M) \sqsubseteq_{\text{ewgt}}^n (\Delta \triangleright N)$.*

Proof. This will follow from the more general full-abstraction result, given in Theorem 4.19. □

Unfortunately this result is not true for an arbitrary set of external owners \mathcal{O} . Example 4.9 can be used to show that the \mathcal{O} -wLTS has not taken into account the fact that observers have access to the funds of arbitrary owners in \mathcal{O} .

Example 4.14. We use the notation from Example 4.9, which in turn is inherited from Example 4.2. Let \mathcal{O} be a set of owners which includes o and the fresh e . Then it is easy to check that $\Gamma \triangleright [a!]_{\text{o}} \sqsubseteq_{\mathcal{O}\text{wgt}}^0 \Delta \triangleright [a!]_{\text{o}}$. But we have already argued in Example 4.9 that $\Gamma \triangleright [a!]_{\text{o}} \not\sqsubseteq_{\mathcal{O}:\text{cxt}}^0 \Delta \triangleright [a!]_{\text{o}}$. ■

So we have to revise the \mathcal{O} -wLTS to take into account the access which observers may have to funds being used by the systems under investigation.

Definition 4.15 (Fund transfer). For every $k \in \mathbb{N}$ let $\xrightarrow{(u,k,p)}$ be the partial function over cost environments defined by letting $\Gamma \xrightarrow{(u,k,p)} \Delta$ whenever Δ can be obtained from Γ by transferring k funds from owner u to owner p . Formally this partial function is only defined when $\Gamma^o(u) \geq k$, in which case $\Delta^o(u) = \Gamma^o(u) - k$, $\Delta^o(p) = \Gamma^o(p) + k$, when $p \neq u$ and all other components of Δ are inherited directly from Γ ; when $p = u$ the operation leaves Δ unchanged. This leads to a new action over configurations, with a new abstract label $\text{ext}(u, k, p)$: we let

$$(\Gamma_1 \triangleright M) \xrightarrow{\text{ext}(u,k,p)_w^{\text{o}}} (\Gamma_2 \triangleright M)$$

whenever $\Gamma_1 \xrightarrow{(u,k,p)} \Gamma_2$, and u, p are owners in \mathcal{O} , where $w = (\Gamma_2^{\text{rec}} - \Gamma_1^{\text{rec}})$. ■

This gives rise to yet another LTS whose states are Picost configurations, which we refer to as \mathcal{O} -awLTS, which induces another bisimulation preorder. But we also need to take Proposition 4.10 into account.

Definition 4.16 (Abstract weighted bisimulation preorder). A family of relations over Picost configurations $\{\mathcal{R}^n \mid n \in \mathbb{N}\}$ is said to be a \mathcal{O} -abstract amortised weighted bisimulation whenever

- (i) $\Gamma \triangleright M \mathcal{R}^n \Delta \triangleright M'$ implies $\Gamma^o(o) = \Delta^o(o)$ for every o in O
- (ii) $\{\mathcal{R}^n \mid n \in \mathbb{N}\}$ is an amortised weighted bisimulation in O -awLTS.

We write $\mathcal{C} \sqsubseteq_{O\text{awgt}}^n \mathcal{D}$ to denote the maximal family of such relations. ■

Note that these relations $\{\sqsubseteq_{O\text{awgt}}^n \mid n \in \mathbb{N}\}$ actually coincide with $\{\sqsubseteq_{\text{ewgt}}^n \mid n \in \mathbb{N}\}$ when O is the singleton external observer $\{e\}$; this follows because the extra fund transfer actions have no effect: $(\Gamma_1^e \triangleright M) \xrightarrow{\text{ext}(u,k,p), \{e\}} (\Gamma_2^e \triangleright M)$ if and only if $\Gamma_1^e = \Gamma_2^e$.

It also coincides with the preorders used in Section 3.3, under certain conditions.

Proposition 4.17. *Let O be the set of owners used in the two configurations Γ and Δ and suppose that all owners in O have indefinite funds; that is $\Gamma(o) = \Delta(o) = \infty$ for every owner $o \in O$. Then $\Gamma \triangleright M \sqsubseteq_{\text{wgt}}^n \Delta \triangleright N$ implies $\Gamma \triangleright M \sqsubseteq_{O\text{awgt}}^n \Delta \triangleright N$.*

Proof. Straightforward. When funds are unlimited the constraint (i) in Definition 4.16 is vacuous, as is the requirement to match the fund actions labelled $\text{ext}(u,k,p)$. The result now follows because every concrete action in the wLTS used in Section 3.3 is automatically also an abstract action in O -awLTS. □

It follows that the work of Section 3.3 has not been in vain; the proofs in the examples can be taken to be about the more abstract preorders $\sqsubseteq_{O\text{awgt}}^n$.

The remainder of this section is devoted to showing that, subject to a minor restriction, the co-inductive proof methodology based on $\{\sqsubseteq_{O\text{awgt}}^n \mid n \in \mathbb{N}\}$ satisfies the informal criteria (a), (b), and (c) set out at the beginning of this section. It has certain advantages over that used in Section 3.3; in matching input and output moves the principles involved do not have to match up exactly. However in the general case it also has a disadvantage with cost environments in which certain owners have finite funds. If the observer has access to such owners then it is necessary to establish that the proposed relations between configurations are invariant under the transfer of funds between them. Of course in the particular case of a purely external observer, where O is taken to be $\{e\}$, which is possibly the most interesting case, then this requirement is vacuous.

Definition 4.18 (Simple types). The type $R = \langle k_u, k_p \rangle$ is *simple* whenever $k_p = 0$, meaning that resources of type R cost nothing to provide. A cost environment is called *simple* whenever it can be written as $\Gamma_{\text{dyn}}, a_1 : R_1, \dots, a_n : R_n$ where Γ_{dyn} is a basic environment and all R_i are simple.

Restricting attention to simple types we know that for every resource name a there is some $k \in \mathbb{N}$ such that $\Gamma \xrightarrow{(u,a,p)} \Delta$ if and only if $\Gamma \xrightarrow{(u,k,p)} \Delta$. ■

Theorem 4.19 (Full-abstraction). *Assuming simple cost environments, for every set of observers O and every $n \in \mathbb{N}$, $(\Gamma \triangleright M) \sqsubseteq_{O:\text{cxt}}^n (\Delta \triangleright N)$ if and only if $(\Gamma \triangleright M) \sqsubseteq_{O\text{awgt}}^n (\Delta \triangleright N)$.*

The proof of this result is the subject of the remainder of this section; we will also see how the restriction to simple types can be lifted, at the expense of a generalisation of the fund action from Definition 4.15.

4.2. Full abstraction. First let us consider criteria (a) above, Compositionality. In fact we now have a parametrised version of this, O -contextuality from Definition 4.7, which we tackle in two steps. First we require a lemma.

Lemma 4.20.

- (1) Suppose $\Gamma \triangleright M \xrightarrow{\lambda} \Delta \triangleright N$. Then $\Gamma, r:\mathbf{R} \triangleright M \xrightarrow{\lambda} \Delta, r:\mathbf{R} \triangleright N$.
- (2) Conversely, suppose $\Gamma, r:\mathbf{R} \triangleright M \xrightarrow{\lambda} \Delta, r:\mathbf{R} \triangleright N$, where the label λ does not describe a communication along the channel r . Then
- (a) $\Gamma \triangleright M \xrightarrow{\lambda} \Delta \triangleright N$
- (b) or the concrete action label λ is of the form $(u, a?r, \mathbf{p})$, in which case $\Gamma \triangleright M \xrightarrow{(u, (r:\mathbf{R})a?r, \mathbf{p})} \Delta, r:\mathbf{R} \triangleright N$.
- (3) $\Gamma \triangleright M \xrightarrow{(u, (r:\mathbf{R})a?r, \mathbf{p})} \Delta, r:\mathbf{R} \triangleright N$ implies $\Gamma, r:\mathbf{R} \triangleright M \xrightarrow{(u, a?r, \mathbf{p})} \Delta, r:\mathbf{R} \triangleright N$

Proof. Each statement is proved by induction on the derivation of the judgement. Note that for any a in the domain of Γ , $\Gamma \xrightarrow{(u, a, \mathbf{p})} \Delta$ if and only if $\Gamma, r:\mathbf{R} \xrightarrow{(u, a, \mathbf{p})} \Delta, r:\mathbf{R}$. \square

Proposition 4.21 (O-contextual). $(\Gamma \triangleright M) \sqsubseteq_{\text{Oawgt}}^n (\Delta \triangleright N)$ implies $(\Gamma, r:\mathbf{R} \triangleright M) \sqsubseteq_{\text{Oawgt}}^n (\Delta, r:\mathbf{R} \triangleright N)$.

Proof. Let $\{\mathcal{R}^n \mid n \in \mathbb{N}\}$ be the family of relations over Picost configurations defined by letting $(\Gamma, r:\mathbf{R} \triangleright M) \mathcal{R}^n (\Delta, r:\mathbf{R} \triangleright N)$ whenever

- (i) either $(\Gamma \triangleright M) \sqsubseteq_{\text{Oawgt}}^n (\Delta \triangleright N)$
- (ii) or $(\Gamma, r:\mathbf{R} \triangleright M) \sqsubseteq_{\text{Oawgt}}^n (\Delta, r:\mathbf{R} \triangleright N)$.

It is sufficient to show that this satisfies the conditions in Definition 4.16. Note that condition (i) of this definition is trivial.

So suppose $(\Gamma, r:\mathbf{R} \triangleright M) \mathcal{R}^n (\Delta, r:\mathbf{R} \triangleright N)$ and $(\Gamma, r:\mathbf{R} \triangleright M) \xrightarrow{\mu \circ \circ} (\Gamma', r:\mathbf{R} \triangleright M')$ is an abstract action. We have to find a matching abstract move $(\Delta, r:\mathbf{R} \triangleright N) \xrightarrow{\hat{\mu} \circ \circ} (\Delta', r:\mathbf{R} \triangleright N')$.

Let us look at the concrete action underlying this abstract action, $(\Gamma, r:\mathbf{R} \triangleright M) \xrightarrow{\lambda} (\Gamma', r:\mathbf{R} \triangleright M')$. Since we know $(\Gamma \triangleright M)$ is a configuration λ can not describe a communication along r , and so we can apply part (2) of the previous lemma, to obtain two cases:

- (a) $\Gamma \triangleright M \xrightarrow{\lambda} \Gamma' \triangleright M'$. In this case the required matching move can be obtained using the fact that $(\Gamma \triangleright M) \sqsubseteq_{\text{Oawgt}}^n (\Delta \triangleright N)$, together with an application of part (1) of Lemma 4.20.
- (b) λ is the input action $(u, a?r, \mathbf{p})$, and $\Gamma \triangleright M \xrightarrow{(u, (r:\mathbf{R})a?r, \mathbf{p})} \Gamma', r:\mathbf{R} \triangleright N$. Here we again use the fact that $(\Gamma \triangleright M) \sqsubseteq_{\text{Oawgt}}^n (\Delta \triangleright N)$ to find a matching weak concrete move from $(\Delta \triangleright N)$ labelled $(u, (r:\mathbf{R})a?r, \mathbf{p}')$ for some owner \mathbf{p}' . Part (3) of Lemma 4.20 can now be used to transform this into a required matching move from $(\Delta, r:\mathbf{R} \triangleright N)$. In this case the matching will be because of clause (ii) in the definition of the family \mathcal{R}^n . \square

Theorem 4.22 (O-contextual). Suppose $(\Gamma \triangleright M \mid [P]_{\circ})$ and $(\Delta \triangleright N \mid [P]_{\circ})$ are both configurations, where $\circ \in \mathbf{O}$. Then $(\Gamma \triangleright M) \sqsubseteq_{\text{Oawgt}}^k (\Delta \triangleright N)$ implies $(\Gamma \triangleright M \mid [P]_{\circ}) \sqsubseteq_{\text{Oawgt}}^k (\Delta \triangleright N \mid [P]_{\circ})$.

Proof. We follow the standard proof structure, see Section 2.3 of [SW01], Proposition 6.4 of [HR04], Proposition 2.21 of [Hen07]; however the precise details are somewhat different. Let $\{\mathcal{R}^n \mid n \in \mathbb{N}\}$ be the smallest family of relations which satisfies:

- (i) $\Gamma \triangleright M \sqsubseteq_{\text{Oawgt}}^n \Delta \triangleright N$ implies $\Gamma \triangleright M \mathcal{R}^n \Delta \triangleright N$
- (ii) $\Gamma \triangleright M \mathcal{R}^n \Delta \triangleright N$ implies $(\Gamma \triangleright M \mid [P]_{\circ}) \mathcal{R}^n (\Delta \triangleright N \mid [P]_{\circ})$, whenever $\circ \in \mathbf{O}$ and both $(\Gamma \triangleright M \mid [P]_{\circ})$ and $(\Delta \triangleright N \mid [P]_{\circ})$ are configurations
- (iii) $\Gamma, r:\mathbf{R}_1 \triangleright M \mathcal{R}^n \Delta, r:\mathbf{R}_2 \triangleright N$ implies $\Gamma \triangleright (\text{new } r:\mathbf{R}_1)M \mathcal{R}^n \Delta \triangleright (\text{new } r:\mathbf{R}_2)N$.

We show that this family satisfies the requirements of Definition 4.16, up to structural equivalence, from which the result will follow.

First note that for any n ,

$$\Gamma \triangleright M \mathcal{R}^n \Delta \triangleright N \text{ implies } \Gamma, r:\mathbb{R} \triangleright M \mathcal{R}^n \Delta, r:\mathbb{R} \triangleright N \quad (4.1)$$

This can be proved by induction on why $\Gamma \triangleright M \mathcal{R}^n \Delta \triangleright N$, with the base case being provided by Proposition 4.21.

So suppose $\Gamma \triangleright M \mathcal{R}^n \Delta \triangleright N$ and $\Gamma \triangleright M \xrightarrow{\mu}_{\circ} \Gamma' \triangleright M_d$; we have to find a matching abstract move $\Delta \triangleright N \xrightarrow{\hat{\mu}}_{\circ} \Delta' \triangleright N_d$ such that $\Gamma' \triangleright M_d \mathcal{R}^{(n+v-w)} \Delta' \triangleright N_d$; the symmetric requirement, of matching a move from $\Delta \triangleright N$ by a corresponding one from $\Gamma \triangleright M$, is treated in an analogous fashion.

We proceed by induction on why $\Gamma \triangleright M \mathcal{R}^n \Delta \triangleright N$, there being three cases, (i), (ii) and (iii) above, to consider. In the first case the requirement comes from Proposition 3.4. We concentrate on case (ii), where we know M, N have the form $(M' \mid [P]_{\circ})$, $(N' \mid [P]_{\circ})$ respectively, where $\circ \in \mathbb{O}$ and we know by induction that $\Gamma \triangleright M' \mathcal{R}^n \Delta \triangleright N'$. We now examine why $\Gamma \triangleright M' \mid [P]_{\circ} \xrightarrow{\mu}_{\circ} \Gamma' \triangleright M_d$, and to start let us assume that μ is the label $\text{ext}(u, k, \mathfrak{p})$, where the reasoning is straightforward. This means, by definition, that M_d is $M \mid [P]_{\circ}$, u, \mathfrak{p} are in \mathbb{O} and $\Gamma \xrightarrow{(u, k, \mathfrak{p})} \Gamma'$, which in turn implies $\Gamma \triangleright M' \xrightarrow{\text{ext}(u, k, \mathfrak{p})}_{\circ} \Gamma' \triangleright M'$; moreover incidently k and v must coincide, although this fact is not required here. By induction this can be matched by an action $\Delta \triangleright N' \xrightarrow{\text{ext}(u, k, \mathfrak{p})}_{\circ} \Delta' \triangleright N''$ such that $(\Gamma \triangleright M') \mathcal{R}^{(n+v-w)} (\Delta \triangleright N'')$. This matching action can now be transformed into an action of the form $\Delta \triangleright N' \mid [P]_{\circ} \xrightarrow{\text{ext}(u, k, \mathfrak{p})}_{\circ} \Delta' \triangleright N'' \mid [P]_{\circ}$ which is easily seen to be the required matching abstract move.

Having disposed of this simple case we now know that there is a derivation using the rules from Figure 6, Figure 7 of the underlying action

$$\Gamma \triangleright M' \mid [P]_{\circ} \xrightarrow{\lambda} \Gamma' \triangleright M_d, \quad (4.2)$$

where $v = (\Gamma'^{\text{rec}} - \Gamma^{\text{rec}})$, and λ is the more concrete version of the label μ . If M' is responsible for the concrete action (4.2), then a straightforward application of the induction hypothesis will provide the required corresponding move. Suppose instead that $[P]_{\circ}$ is responsible, that is (4.2) takes the form

$$\Gamma \triangleright M' \mid [P]_{\circ} \xrightarrow{\lambda} \Gamma' \triangleright M' \mid [P']_{\circ} \quad (4.3)$$

because $\Gamma \triangleright [P]_{\circ} \xrightarrow{\lambda} \Gamma' \triangleright [P']_{\circ}$; here the reasoning needs to be more involved.

(a) First suppose this move is an external move, say an output with label λ being $(\circ, (\tilde{r}:\tilde{\mathbb{R}})a!v, \mathfrak{p})$ for some owner \mathfrak{p} . Because we are actually matching \mathbb{O} -actions we know that this \mathfrak{p} is actually in \mathbb{O} .

Applying Lemma 3.8 we know that Γ' has the form $\Gamma'', \tilde{r}:\tilde{\mathbb{R}}$, where $\Gamma \xrightarrow{(u, a, \mathfrak{p})} \Gamma''$. The use of simple types means that $\Gamma^{\mathfrak{p}}(a) = 0$ and $\Gamma^u(a) = k$ for some k , and standard resource charging implies that this k is actually v . Thus we have the external move $\Gamma \triangleright M' \xrightarrow{\text{ext}(u, k, \mathfrak{p})}_{\circ} \Gamma'' \triangleright M'$ and we know by induction this move can be matched by some $\Delta \triangleright N' \xrightarrow{\text{ext}(u, k, \mathfrak{p})}_{\circ} \Delta'' \triangleright N'$ such that $\Gamma'' \triangleright M' \mathcal{R}^{(n+v-w)} \Delta'' \triangleright N'$. This matching

move actually has the form

$$\Delta \triangleright N' \xrightarrow{\tau_{w_1}^{\circ}} \Delta_1 \triangleright N'_1 \xrightarrow{\text{ext}(u,k,p)_k^{\circ}} \Delta_2 \triangleright N'_1 \xrightarrow{\tau_{w_3}^{\circ}} \Delta'' \triangleright N' \quad (4.4)$$

with $w = w_1 + k + w_3$.

An application of part (iv) of Lemma 3.9 or Lemma 3.8 gives the move $\Delta_1 \triangleright [P]_{\circ} \xrightarrow{(u,(\tilde{r}:\tilde{R})\alpha,p)} \Delta_2, \tilde{r}:\tilde{R} \triangleright [P']_{\circ}$ which can be combined with the pre- and post- τ moves in (4.4) to give $\Delta \triangleright N' \mid [P]_{\circ} \xrightarrow{\lambda} \Delta'', \tilde{r}:\tilde{R} \triangleright N' \mid [P']_{\circ}$. This is the required matching move since we know $\Gamma'' \triangleright M' \mathcal{R}^{(n+v-w)} \Delta'' \triangleright N'$, from which $\Gamma'', \tilde{r}:\tilde{R} \triangleright M' \mid [P']_{\circ} \mathcal{R}^{(n+v-w)} \Delta'', \tilde{r}:\tilde{R} \triangleright N' \mid [P']_{\circ}$ follows by the remark (4.1) above and the definition of the family $\{\mathcal{R}^k \mid k \geq 0\}$.

When the label λ in the move (4.3) above is an input the argument is very much the same but with an application of Lemma 3.9 in place of Lemma 3.8; it is therefore omitted.

- (b) Now suppose the move from $[P]_{\circ}$ we are examining is an internal move, taking the form $\Gamma \triangleright [P]_{\circ} \xrightarrow{\tau} \Gamma' \triangleright [P']_{\circ}$. Here we apply Theorem 3.10 and Proposition 2.5, which tell us that there are in principle three possibilities, (i), (ii) or (iii). But an analysis of the proof will show that for processes of the form $[P]_{\circ}$ case (i) is actually the only possibility. Here Γ' coincides with Γ , implying incidently that $v = 0$. As we know $\Delta \triangleright [P]_{\circ}$ is a configuration we also get $\Delta \triangleright [P]_{\circ} \xrightarrow{\tau} \Delta \triangleright [P']_{\circ}$ and therefore that $\Delta \triangleright N' \mid [P]_{\circ} \xrightarrow{\tau} \Delta \triangleright N' \mid [P']_{\circ}$. It is easy to now check that this is the required matching move, since by definition $\Gamma \triangleright M \mid [P]_{\circ} \mathcal{R}^n \Delta \triangleright N \mid [P']_{\circ}$.

We are left with the possibility that the underlying action to be matched, (4.2) above, involves communication and therefore takes the form

$$\Gamma \triangleright M' \mid [P]_{\circ} \xrightarrow{\tau} \Gamma' \triangleright (\text{new } \tilde{r}:\tilde{R})(M'' \mid [P']_{\circ})$$

There are two cases, depending on whether M' performs an input or an output. Let us consider the latter, the former being similar but slightly easier. So we have

$$\begin{aligned} \Gamma \triangleright M' &\xrightarrow{\lambda} \Gamma', \tilde{r}:\tilde{R} \triangleright M'' \\ \Gamma \triangleright [P]_{\circ} &\xrightarrow{\bar{\lambda}} \Gamma', \tilde{r}:\tilde{R} \triangleright [P']_{\circ} \end{aligned} \quad (4.5)$$

with $\lambda, \bar{\lambda}$ taking the forms $(u, (\tilde{r}:\tilde{R})a!v, \circ)$, $(u, (\tilde{r}:\tilde{R})a?v, \circ)$ respectively, for some owner u . By induction the first move, or rather its abstract version, can be matched because \circ is an owner in \mathcal{O} , giving

$$\Delta \triangleright N' \xrightarrow{\tau} \Delta_1 \triangleright N'_1 \xrightarrow{(u',(\tilde{r}:\tilde{R}')a!v,\circ)} \Delta_2, \tilde{r}:\tilde{R}' \triangleright N'_2 \xrightarrow{\tau} \Delta', \tilde{r}:\tilde{R}' \triangleright N'' \quad (4.6)$$

for some owner u' , such that $(\Gamma', \tilde{r}:\tilde{R}' \triangleright M'') \mathcal{R}^{(n+v-w)} (\Delta', \tilde{r}:\tilde{R}' \triangleright N'')$, where $w = (\Delta'^{\text{rec}} - \Delta^{\text{rec}})$. Note that the type of the extruded names, \tilde{R}' , may in general be different than the types at which they were extruded by M' , and the owner u' may also be different, thereby a priori complicating matters when we try to combine this action with that from $[P]_{\circ}$, in (4.5) above.

However an application of part (ii) of Lemma 3.8, gives $\Delta_1 \xrightarrow{(u',a,\circ)} \Delta_2$, and therefore from (4.5) and part (v) of Lemma 3.9 we get $\Delta_1 \triangleright [P]_{\circ} \xrightarrow{(u',(\tilde{r}:\tilde{R}')a?v,\circ)} \Delta_2, \tilde{r}:\tilde{R}' \triangleright [P']_{\circ}$. This

concrete move can now be combined with the concrete move (4.6) to give the required matching abstract move $\Delta \triangleright N' \mid [P]_{\circ} \xrightarrow{\tau}_w (\text{new } \tilde{r}:\tilde{R}')(N' \mid [P']_{\circ})$. \square

The attentive reader will have noticed that the restriction to *simple* types was necessary in order to be able to model the use of a resource by the observers using actions based on the transfer function $\Gamma \xrightarrow{\text{ext}(u,k,p)} \Delta$, which records the transfer of k funds, the cost of using the resource, from the user to the provider. If we drop the restriction to *simple types*, then the effect of using a resource is more complicated; a certain amount will be debited to the user, while another amount, possibly negative, will be credited to the user. This can be accommodated by a more general transfer function $\Gamma \xrightarrow{\text{ext}(u,(k_1,k_2),p)} \Delta$, leading in turn to a more general abstract arrow in part (d) of Definition 4.11. With this adjustment compositionality can also be established for arbitrary types.

This contextual results leads in a straightforward manner to establishing the second informal criteria, (b):

Theorem 4.23 (Soundness). *For every $n \in \mathbb{N}$ and every set of owners \mathcal{O} , $(\Gamma \triangleright M) \sqsubseteq_{\mathcal{O}\text{awgt}}^n (\Delta \triangleright N)$ implies $(\Gamma \triangleright M) \sqsubseteq_{\mathcal{O}:\text{cxt}}^n (\Delta \triangleright N)$.*

Proof. (Outline) It is sufficient to show that the family of relations $\{\sqsubseteq_{\text{awgt}}^n \mid n \in \mathbb{N}\}$ satisfies the three defining properties of the family of contextual equivalences. *Cost improving* follows by definition, at least up to structural induction, in view of Theorem 3.10, and the two preceding results establish \mathcal{O} -contextuality. The final property, *Preservation of observations*, is also straightforward, since, for example, the ability to observe $a!$ from a configuration coincides with its ability to perform some output action on the resource a . \square

The final criteria (c), *Completeness*, depends as usual on the ability to define contexts which capture the effect of each of the abstract \mathcal{O} -actions described in Definition 4.11. We first make this precise.

We use two fresh cost-free resources, *succ*, *fail* to record the success or failure of tests, and a third *req* for housekeeping purposes. For any Γ we use Γ^t to denote the cost environment obtained by adding on these resources. Now let μ be an abstract action which uses the bound names (\tilde{r}) . Then we say μ is definable relative to \mathcal{O} if for every finite set of names F there exists a system T_{μ}^F using only the owners from \mathcal{O} such that

- (i) if $\text{dom}(\Gamma^u) \subseteq F$ and $\Gamma \triangleright M \xrightarrow{\mu}_w^{\circ} \Delta, \tilde{r}:\tilde{R} \triangleright N$ then $\Gamma^t \triangleright M \mid T_{\mu}^F \xrightarrow{\tau}_w^{\circ} \Delta^t \triangleright (\text{new } r:\mathbb{R})(\text{succ}!\langle \tilde{r} \rangle \mid N)$ where $M' \Downarrow \text{succ}!$ and $R \not\Downarrow \text{fail}!$
- (ii) conversely, whenever $\text{dom}(\Gamma^u) \subseteq F$, $\Gamma^t \triangleright M \mid T_{\mu}^F \xrightarrow{\tau}_w^{\circ} \Delta^t \triangleright M'$ where $M' \Downarrow \text{succ}!$ and $M' \not\Downarrow \text{fail}!$ implies $M' \equiv (\text{new } \tilde{r}:\tilde{R})(\text{succ}!\langle \tilde{r} \rangle \mid N)$, where $\Gamma \triangleright M \xrightarrow{\mu}_w^{\circ} \Delta, \tilde{r}:\tilde{R} \triangleright N$

Theorem 4.24 (Definability). *All input, output and external actions are definable.*

Proof. (Outline) Let us look at two examples. First suppose that μ is the label $\text{ext}(u, k, p)$ where u and p are both in \mathcal{O} ; here (\tilde{r}) is empty and the set of names F plays no role. The definition of T_{μ}^F uses a variation on Example 2.7. We use

$$[\text{fail}! \mid (\text{new } r:\mathbb{R}_k)\text{req}!\langle r \rangle.r!. \text{stop}]_u \mid [\text{req}?(x).y?.\text{fail}?.\text{succ}!]_p$$

where \mathbb{R}_k is the type $(k, 0)$. This ensures that whenever $(\Gamma \triangleright M \mid T_{\mu}^F)$ evolves at cost w to a configuration \mathcal{C} such that $\mathcal{C} \Downarrow \text{succ}$ but $\mathcal{C} \not\Downarrow \text{fail}$ then the newly generated resource r must have been used by u and provided by p . This is only possible if $\Gamma^t \triangleright M$ can evolve to a

configuration in which a transfer of k can be made from u to p ; that is a configuration $\Gamma^{t'} \triangleright M'$ such that $\Gamma^{t'} \xrightarrow{(u,k,p)} \Gamma^{t''}$. This in turns implies that we must have $\Gamma \triangleright M \xrightarrow{\text{ext}(u,k,p)_w^o} \Delta \triangleright N$ for some configuration $\Delta \triangleright N$. Note the cost here is w because all of the resources used by the test T_μ^F are cost-free.

For the second example consider the abstract output action label $((r)a!r, p)$, where we know p is in O . Here we let T_μ^F be

$$[\text{fail!} \mid a?(x) . \text{if } x \in F \text{ then stop else fail?} . \text{succ!}]_p$$

where $x \in F$ is an abbreviation for a series of tests deciding whether or not x is in the finite set of names F . Intuitively whenever this is used in a cost environment Γ satisfying $\text{dom}(\Gamma^u) \subseteq F$ this test will fail only when x is instantiated by a fresh name.

Once more it is easy to say that the ability of $\Gamma^t \triangleright M \mid T_m^F$ to evolve to a configuration \mathcal{C} satisfying $\mathcal{C} \Downarrow \text{succ}$ but $\mathcal{C} \not\Downarrow \text{fail}$ coincides with the ability of $\Gamma \triangleright M$ to do a weak concrete move labelled $(u, (r : R)a!r, p)$ for some owner u and type R . Moreover the cost of this weak concrete action will be exactly the same as the evolution from $\Gamma^t \triangleright M \mid T_m^F$, because the interactions with the test T_μ^F is free. \square

Theorem 4.25 (Completeness). *For every $n \in \mathbb{N}$ and every set of owners O , $(\Gamma \triangleright M) \sqsubseteq_{O:\text{cxt}}^n (\Delta \triangleright N)$ implies $(\Gamma \triangleright M) \sqsubseteq_{O:\text{awgt}}^n (\Delta \triangleright N)$.*

Proof. (Outline) It suffices to show that the family $\{\sqsubseteq_{O:\text{cxt}}^n \mid n \in \mathbb{N}\}$ satisfies the conditions in Definition 4.16. Note that condition (i) is already established by Proposition 4.10. Now suppose $\Gamma \triangleright M \sqsubseteq_{O:\text{cxt}}^n \Delta \triangleright N$ and $\Gamma \triangleright M \xrightarrow{\mu}_v^o \Gamma' \triangleright M'$. We have to find a matching move from $\Delta \triangleright N$, which is relatively straightforward because of Theorem 4.24. As an example suppose μ is the output label $((r)a!r, p)$, and so Γ' has the structure $\Gamma'', r:R$ for some R . Because of Compositionality we know $\Gamma^t \triangleright M \mid T_\mu \sqsubseteq_{\text{cxt}}^n \Delta^t \triangleright N \mid T_\mu$. Using the first part of the Definability Theorem we know that, up to structural equivalence,

$$\Gamma^t \triangleright M \mid T_\mu^F \xrightarrow{v}^* \Gamma^{t''} \triangleright (\text{new } r:R)(\text{succ!}\langle r \rangle \mid M').$$

Using the properties of the family $\{\sqsubseteq_{O:\text{cxt}}^n \mid n \in \mathbb{N}\}$ this move must be matched by move

$$\Delta^t \triangleright N \mid T_\mu^F \xrightarrow{w}^* \Delta^{t''} \triangleright N''$$

where

$$\Gamma^{t''} \triangleright (\text{new } r:R)(\text{succ!}\langle r \rangle \mid M') \sqsubseteq_{\text{cxt}}^{(n+v-w)} \Delta^{t''} \triangleright N'' \quad (4.7)$$

Moreover we know $N'' \Downarrow \text{succ!}$ and $N'' \not\Downarrow \text{fail!}$ and so the Definability theorem tells us that $N'' \equiv (\text{new } r:R')(\text{succ!}\langle r \rangle \mid N')$ where

$$\Delta \triangleright N \xrightarrow{w}^o \Delta'', r:R' \triangleright N'$$

This would be the required matching move, if we had

$$\Gamma'', r:R \triangleright M' \sqsubseteq_{O:\text{cxt}}^{(n+v-w)} \Delta'', r:R \triangleright N' \quad (4.8)$$

whereas (4.7) only gives us, up to structural equivalence,

$$\Gamma^{t''} \triangleright (\text{new } r:R)(\text{succ!}\langle r \rangle \mid M') \sqsubseteq_{\text{cxt}}^{(n+v-w)} \Delta^{t''} \triangleright (\text{new } r:R')(\text{succ!}\langle r \rangle \mid N') \quad (4.9)$$

However the so-called *Extrusion Lemma*, see Proposition 6.7 of [HR04] and Lemma 2.38 of [Hen07], can easily be adapted to Picost, to show that the required (4.8) does indeed follow from (4.9) \square

5. CONCLUSION

In this paper we have developed a behavioural theory based on bisimulations for a version of the picalculus, *Picost*, in which

- resources have costs associated with them
- code runs under the financial responsibility of owners, or principals
- code can only be executed if the owner responsible for it can finance the available transactions.

The behavioural theory gives rise to a co-inductive proof methodology for comparing the costed behaviour of systems. We have demonstrated the usefulness of the methodology by treating some examples, and we have offered at least a preliminary justification for the theory in terms of contextual requirements, parametrised on sets of owners. We have provided some evidence that the most appropriate theory emerges when this set of observers is taken to be some single external observer, external to the owners funding the systems being investigated. In particular with this particular set of observers there is no need to consider the extra actions $\text{ext}(u, k, \mathbf{p})$ when establishing bisimulations.

The language could be extended in many ways without unduely affecting the underlying theory. Perhaps the most obvious extension would be the introduction of *ownership types*, to control which owners can use which resources; this would help in the modularisation of systems. One could also introduce a scoping mechanism for owners, limiting the range within systems of their financial responsibility. One effect of such extensions would be that owners would play a much more significant role in the (abstract) actions on which bisimulations are based. Such investigations we leave for future work.

The language could also be extended with mechanisms whereby processes could be aware of which owners are funding which resources, and more importantly base their behaviour on such knowledge. More ambitiously the semantics of the language could be generalised so that behaviour is now dependent on some *dynamic cost model*. There is considerable scope here for inventing more realistic cost models, whereby for example costs associated with producing/consuming resources could vary according to *market dynamics*. It is likely that a probabilistic setting would be most appropriate for developing such models.

The underlying theory of *weighted bisimulations* also deserves attention. For example it is not clear if the theory is decidable, even for finite-state systems. More generally it would be interesting to have techniques which would calculate the costs necessary to assign to actions in order to ensure the equivalence of two systems. There is already an extensive literature on *weighted automata* [DKV09] and decidability issues concerned with them, which may help in this regard.

Related work: The research reported in the current paper grew out of preliminary work reported in [HG08]. There a language π_{cost} was defined and also given a semantics relative to *cost environments*. But there are significant differences. At the language level the construct central to *Picost*, $[P]_{\mathbf{o}}$, is absent in π_{cost} ; indeed in the latter there is no representation of owners being responsible for specific computations. The cost environments used are also quite different; in π_{cost} funds are associated directly with resources, which complicates considerably the reduction semantics as the resource types need to be dynamic. Here all funds are retained by owners, which simplifies matters considerably, and this facilitates the introduction of *charges* for resource usage and *benefits* for resource provision. Finally the behavioural theories are different. The concept of *weighted bisimulation* is considerably

$$\begin{array}{l}
\text{Reader:} \\
R_1 \Leftarrow \text{goLib?}(\text{name}) . (\text{newr}) R_2(r, \text{name}) \\
R_2(r, \text{name}) \Leftarrow \text{reqR!}\langle r, \text{name} \rangle . R_3(r) \\
R_3(r) \Leftarrow r?(b) . R_4(b) \\
R_4(b) \Leftarrow \text{goHome!}\langle b \rangle . R_1 \\
\\
\text{Library:} \\
L_1 \Leftarrow \text{reqR?}(y, z) . L_2(y, z) \\
L_2(y, z) \Leftarrow L_3(y, z) \oplus (\text{newr}) L_4(r, y, z) \\
L_3(y, z) \Leftarrow y!\langle \text{book}(z) \rangle . L_1 \\
L_4(r, y, z) \Leftarrow \text{reqS!}\langle r, z \rangle . L_5(y) \\
L_5(y) \Leftarrow r?(b) . L_6(y, b) \\
L_6(y, b) \Leftarrow y!\langle b \rangle . L_1 \\
\\
\text{Store:} \\
S_1 \Leftarrow \text{reqS?}(y, z) . S_2(y, z) \\
S_2(y, z) \Leftarrow y!\langle \text{book}(z) \rangle . S_1
\end{array}$$

Figure 8: Notation for library code

more flexible than the *cost bisimulations* of [HG08], as the latter simply compares the relative cost of performing each particular action.

Weighted bisimulations are a direct generalisation of the notion of *amortised bisimulations* from [KAK05]; these were originally defined for a version of CCS, [Mil89], in which only external actions have associated with them a cost. Nevertheless we believe that our generalisation is significant, at least in that it will make the concepts more generally applicable. However similar ideas have a long history in the field of *timed* process calculi; see for example [Tof94]. A good survey of the use of *amortisation* for timed processes can be found in [LV06].

Other resource-aware calculi have already appeared in the literature. A typical example is the variant of *mobile ambients* [CG00] from [BBDCS03] in which the resource in question is *space*, and the processes in the calculi have a *bounded capacity* to host incoming ambients. Another interesting example may be found in [Tel04], and related publications, which develops a version of the picalculus in which unused resources/channels may be garbage collected. Of particular interest to us is the general theory of *resource-based* computation being developed in [CP07], and related publications. In future work we hope to adapt their resource-based modal logic to Picost.

APPENDIX A. SOME WITNESS BISIMULATIONS

A.1. The library. Here we revisit the example on running a library, discussed in Example 2.6 and Example 3.11, and prove

$$(\Gamma_{\text{central}} \triangleright \text{Sys}_{\text{central}}) \sqsubseteq_{\text{wgt}}^2 (\Gamma_{\text{local}} \triangleright \text{Sys}_{\text{local}}) \quad (\text{A.1})$$

$$\begin{aligned}
N_1 &\Leftarrow (\text{new reqR}:\mathcal{R}_r^c)([R_1]_{\text{pub}} \mid (\text{new reqS}:\mathcal{R}_s^c)([L_1]_{\text{lib}} \mid [S_1]_{\text{lib}})) \\
N_2(n) &\Leftarrow (\text{new reqR}, r)([R_2(r, n)]_{\text{pub}} \mid (\text{new reqS})([L_1]_{\text{lib}} \mid [S_1]_{\text{lib}})) \\
N_3(n) &\Leftarrow (\text{new reqR}, r)([R_3(r)]_{\text{pub}} \mid (\text{new reqS})([L_2(r, n)]_{\text{lib}} \mid [S_1]_{\text{lib}})) \\
N_{41}(n) &\Leftarrow (\text{new reqR}, r)([R_3(r)]_{\text{pub}} \mid (\text{new reqS})([L_3(r, n)]_{\text{lib}} \mid [S_1]_{\text{lib}})) \\
N_{51}(b) &\Leftarrow (\text{new reqR}, r)([R_4(b)]_{\text{pub}} \mid (\text{new reqS})([L_1]_{\text{lib}} \mid [S_1]_{\text{lib}})) \\
N_{42}(n) &\Leftarrow (\text{new reqR}, r, r')([R_3(r)]_{\text{pub}} \mid (\text{new reqS})([L_4(r, r', n)]_{\text{lib}} \mid [S_1]_{\text{lib}})) \\
N_{52}(n) &\Leftarrow (\text{new reqR}, r, r')([R_3(r)]_{\text{pub}} \mid (\text{new reqS})([L_5(r, r')]_{\text{lib}} \mid [S_2(r', n)]_{\text{lib}})) \\
N_{53}(b) &\Leftarrow (\text{new reqR}, r, r')([R_3(r)]_{\text{pub}} \mid (\text{new reqS})([L_6(r, b')]_{\text{lib}} \mid [S_1]_{\text{lib}}))
\end{aligned}$$

Figure 9: Library systems

by exhibiting a witness bisimulation. For convenience we work up to *structural equivalence* and modulo β -moves; essentially these are moves which have no effect on the overall behaviour of systems; see [Hen07, GS96] for details. In Picost these include the actions generated by the rules (L-EXPORT), (L-UNWIND), (L-SPLIT), (L-MATCH), (L-MISMATCH). Let us assume a set of book names BN , ranged over by n and a set of books BK , ranged over by b .

Let us write $\Gamma \sim \Delta$ whenever

- (a) Γ has the form Γ_{dyn} , $\text{goLib}:\langle 0, 5 \rangle$, $\text{goHome}:\langle 0, 5 \rangle$, $\text{reqR}:\langle 0, 1 \rangle$, $\text{reqS}:\langle 0, 1 \rangle$ for some basic environment Γ_{dyn}
 - (b) Δ has the form Δ_{dyn} , $\text{goLib}:\langle 0, 1 \rangle$, $\text{goHome}:\langle 0, 1 \rangle$, $\text{reqR}:\langle 0, 3 \rangle$, $\text{reqS}:\langle 0, 5 \rangle$ where again Δ_{dyn} is some basic environment.
 - (c) $\text{dom}(\Gamma^o) = \text{dom}(\Delta^o) = \{\text{pub}, \text{lib}\}$, with $\Gamma^o(\alpha) = \Delta^o(\alpha) = \infty$, for every α in its domain.
- So effectively Γ must be like Γ_{central} with perhaps a different record filed Γ^{rec} , and Δ must be like Γ_{local} . Our witness bisimulation will contain pairs of the form

$$\Gamma \triangleright N \leftrightarrow \Delta \triangleright M \quad \text{where } \Gamma \sim \Delta$$

The allowed forms of N are described in Figure 9, where for convenience we have omitted the explicit occurrence of the local types \mathcal{R}_r^c , \mathcal{R}_s^c after the first line. These in turn use notation given in Figure 8 for the various processes. The allowed forms for M are identical except for the use of the local types \mathcal{R}_r^l , \mathcal{R}_s^l in place of \mathcal{R}_r^c , \mathcal{R}_s^c .

Let the family of relations over configurations $\{\mathcal{R}^k \mid k \in \mathbb{N}\}$ be determined by the following constraints, where we assume in each clause that $\Gamma \sim \Delta$:

$$\begin{aligned}
\Gamma \triangleright N_1 \mathcal{R}^k \Delta \triangleright M_1 & \quad \text{whenever } k \geq 2 \\
\Gamma \triangleright N_2(n) \mathcal{R}^k \Delta \triangleright M_2(n) & \quad \text{whenever } k \geq 6, n \in \text{BN} \\
\Gamma \triangleright N_i(n) \mathcal{R}^k \Delta \triangleright M_i(n) & \quad \text{whenever } k \geq 4, n \in \text{BN}, i = 3, 41, 51, 42 \\
\Gamma \triangleright N_i(n) \mathcal{R}^k \Delta \triangleright M_i(n) & \quad \text{whenever } k \geq 4, n \in \text{BN}, i = 3, 41, 42 \\
\Gamma \triangleright N_i(b) \mathcal{R}^k \Delta \triangleright M_i(b) & \quad \text{whenever } k \geq 0, b \in \text{BK}, i = 51, 52, 53
\end{aligned}$$

$$\begin{array}{l}
\text{Publisher:} \\
P_1(r_1) \Leftarrow \text{news!}\langle r_1 \rangle . (\text{new}r_2)P_2(r_1, r_2) \\
P_2(r_1, r_2) \Leftarrow \text{adv!}\langle r_2 \rangle . P_3(r_1, r_2) \\
P_3(r_1, r_2) \Leftarrow r_1?(n) . P_4(n, r_2) \\
P_4(n, r_2) \Leftarrow r_2?(d) . P_5(n, d) \\
P_5(n, d) \Leftarrow \text{publish?}\langle z \rangle . P_6(n, d, z) \\
P_6(n, d, z) \Leftarrow z!\langle n, d \rangle . (\text{new}r_1)P_1(r_1) \\
\\
\text{Advertiser:} \\
A_1 \Leftarrow \text{adv?}\langle r \rangle . (\text{new}d)A_2(r, d) \\
A_2 \Leftarrow r!\langle d \rangle . A_1 \\
\\
\text{Publisher with kickback:} \\
P_{K1}(r_1) \Leftarrow \text{news!}\langle r_1 \rangle (\text{new}r_2, k)P_{K2}(r_1, r_2, k) \\
P_{K2}(r_1, r_2, k) \Leftarrow \text{adv!}\langle k, r_2 \rangle . P_{K3}(r_1, r_2, k) \\
P_{K3}(r_1, r_2, k) \Leftarrow r_1?(n) . P_{K4}(n, r_2, k) \\
P_{K4}(n, r_2, k) \Leftarrow r_2?(d) . P_{K5}(n, d, k) \\
P_{K5}(n, d, k) \Leftarrow \text{publish?}\langle z \rangle . P_{K6}(n, d, k, z) \\
P_{K6}(n, d, k, z) \Leftarrow k? . P_{K7}(n, d, z) \\
P_{K7}(n, d, z) \Leftarrow z!\langle n, d \rangle . (\text{new}r_1)P_{K1}(r_1) \\
\\
\text{Advertiser with kickback:} \\
A_{K1} \Leftarrow \text{adv?}\langle k, r \rangle . (\text{new}d)A_{K2}(k, r, d) \\
A_{K2}(k, r, d) \Leftarrow r!\langle d \rangle . (A_{K1} \mid k!)
\end{array}$$

Figure 10: Notation for publisher code

It is fairly straightforward, although tedious, to prove that $\{\mathcal{R}^k \mid k \in \mathbb{N}\}$ satisfies the requirements of being a weak bisimulation in the wLTS of Section 3.3, up to structural equivalence and β -moves. This is facilitated by the fact that the code in each component of the pairs is identical.

Note that the configuration $\Gamma_{\text{central}} \triangleright \text{Sys}_{\text{central}}$ β -reduces to a configuration of the form $\Gamma \triangleright N_1$ and $(\Gamma_{\text{local}} \triangleright \text{Sys}_{\text{local}})$ β -reduces to one of the form $\Delta \triangleright M_1$, where $\Gamma \sim \Delta$, and thus (A.1) above follows.

A.2. The publisher. Here we revisit the publishing example developed in Example 2.8, Example 2.9 and Example 4.12; by exhibiting a witness bisimulation, again up to structural equivalence and β -moves, we show that

$$(\Gamma_{327} \triangleright \text{PA}_K) \sqsubseteq_{\text{ewgt}}^0 (\Gamma_{327} \triangleright \text{PA}) \quad (\text{A.2})$$

subject to minor constraints on Γ ; these constraints allow $\Gamma^o(\mathbf{p})$ to be finite. The systems PA and PA_K , in addition to cost-free communications,

Standard publisher:	$PA_1 \Leftarrow (\text{newadv}, r_1)([P_1(r_1)]_p \mid [A_1]_a)$
	$PA_2(r_1) \Leftarrow (\text{newadv}, r_2)([P_2(r_1, r_2)]_p \mid [A_1]_a)$
	$PA_3(r_1) \Leftarrow (\text{newadv}, r_2, d)([P_3(r_1, r_2)]_p \mid [A_2(r_2, d)]_a)$
	$PA_4(n) \Leftarrow (\text{newadv}, r_2, d)([P_4(n, r_2)]_p \mid [A_2(r_2, d)]_a)$
	$PA_5(n) \Leftarrow (\text{newadv}, d)([P_5(n, d)]_p \mid [A_1]_a)$
	$PA_6(n) \Leftarrow (\text{newadv}, d)([P_6(n, d, r)]_p \mid [A_{K1}]_a)$
Publisher with kickback:	$PA_{K1} \Leftarrow (\text{newadv}, r_1)([P_{K1}(r_1)]_p \mid [A_{K1}]_a)$
	$PA_{K2}(r_1) \Leftarrow (\text{newadv}, r_2, k)([P_{K2}(r_1, r_2, k)]_p \mid [A_{K1}]_a)$
	$PA_{K3}(r_1) \Leftarrow (\text{newadv}, k, r_2, d)([P_{K3}(r_1, r_2, k)]_p \mid [A_{K2}(k, r_2, d)]_a)$
	$PA_{K4}(n) \Leftarrow (\text{newadv}, k, r_2, d)([P_{K4}(n, r_2, k)]_p \mid [A_{K2}(k, r_2, d)]_a)$
	$PA_{K5}(n) \Leftarrow (\text{newadv}, k, r_2, d)([P_{K5}(n, d, k)]_p \mid [A_{K1} \mid k!]_a)$
	$PA_{K6}(n) \Leftarrow (\text{newadv}, k, r_2, d)([P_{K6}(n, d, r, k)]_p \mid [A_{K1} \mid k!]_a)$
	$PA_{K7}(n) \Leftarrow (\text{newadv}, d)([P_{K7}(n, d, r)]_p \mid [A_{K1}]_a)$

Figure 11: Publishing systems

- use resource `news`; in the definition of the cost environment from Example 2.8 this is recorded as a loss of 3, the cost of using `news`. In the abstract `wLTS` we are using this loss is paid for by the funds in $\Gamma_{327}^o(\mathfrak{p})$, while it costs nothing to provide
- provide resource `publish`; in the cost environment this is recorded as a gain of 6, namely the difference between providing it 7 and using it 1. Also this gain is added to the funds of $\Gamma_{327}^o(\mathfrak{p})$.

There are also internal communications which have costs associated with them, namely the use and provision of `adv`; again this is recorded as a loss of 2 which must be funded by $\Gamma_{327}^o(\mathfrak{p})$.

In order to describe the witness bisimulation we use the code abbreviations in Figure 10 and the system definitions in Figure 11. All environments we use have the form $\Gamma_{dyn}, \text{news}:\mathbb{R}_n, \text{publish}:\mathbb{R}_p$, and in order to fund the advertising we assume $\Gamma^o(\mathfrak{a}) = \infty$. In the witness bisimulation $\{\mathcal{R}^k \mid k \in \mathbb{N}\}$ all \mathcal{R}^k are identical and this unique relation \mathcal{R} is characterised by the following constraints:

$$\begin{array}{ll}
\Gamma^e \triangleright PA_{K1} \quad \mathcal{R} \quad \Delta^e \triangleright PA_1 & 5 \leq \Gamma^o(\mathfrak{p}), 5 \leq \Delta^o(\mathfrak{p}) \\
\Gamma^e \triangleright PA_{K2}(r) \quad \mathcal{R} \quad \Delta^e \triangleright PA_2(r) & 2 \leq \Gamma^o(\mathfrak{p}), 2 \leq \Delta^o(\mathfrak{p}), r \in \text{Chan} \\
\Gamma^e \triangleright PA_{K3}(r) \quad \mathcal{R} \quad \Delta^e \triangleright PA_3(r) & r \in \text{Chan} \\
\Gamma^e \triangleright PA_{Ki}(n) \quad \mathcal{R} \quad \Delta^e \triangleright PA_i(n) & 4 \leq i \leq 6, n \in \text{News} \\
\Gamma^e \triangleright PA_{K7}(n) \quad \mathcal{R} \quad \Delta^e \triangleright PA_6(n) & n \in \text{News}
\end{array}$$

Here we use `News` to denote some set of news stories.

It is straightforward to show that this is indeed a weak amortised bisimulation in the abstract wLTS relative to the single external observer e . Since $\Gamma_{327} \triangleright \text{PA}_K$ β -reduces to $\Gamma_{327} \triangleright \text{PA}_{K1}$ and $\Gamma_{327} \triangleright \text{PA}$ β -reduces to $\Gamma_{327} \triangleright \text{PA}_1$, and $\Gamma_{327} \triangleright \text{PA}_{K1} \mathcal{R} \Gamma_{327} \triangleright \text{PA}_1$, the required (A.2) above follows.

ACKNOWLEDGMENTS

The author would like to thank the referees for their very useful comments.

REFERENCES

- [BBDCS03] Franco Barbanera, Michele Bugliesi, Mariangiola Dezani-Ciancaglini, and Vladimiro Sassone. A calculus of bounded capacities. In Vijay A. Saraswat, editor, *ASIAN*, volume 2896 of *Lecture Notes in Computer Science*, pages 205–223. Springer, 2003.
- [CG00] Luca Cardelli and Andrew D. Gordon. Mobile ambients. *Theor. Comput. Sci.*, 240(1):177–213, 2000.
- [CGP08] Giuseppe Castagna, Nils Gesbert, and Luca Padovani. A theory of contracts for web services. In *POPL '08, 35th ACM Symposium on Principles of Programming Languages*, Jan 2008.
- [CP07] Matthew Collinson and David Pym. Algebra and logic for resource-based systems modelling. Technical report, Hewlett-Packard Laboratories, 2007. Submitted for Publication.
- [DKV09] Manfred Droste, Werner Kuich, and Heiko Vogler, editors. *Handbook of Weighted Automata*. EATCS Monographs in Theoretical Computer Science. Springer-Verlag, 2009.
- [GS96] Jan Friso Groote and M. P. A. Sellink. Confluence for process verification. *Theor. Comput. Sci.*, 170(1-2):47–81, 1996.
- [Hen07] Matthew Hennessy. *A distributed picalculus*. Cambridge University Press, 2007.
- [HG08] Matthew Hennessy and Manish Gaur. Counting the cost in the picalculus (extended abstract). *Electr. Notes Theor. Comput. Sci.*, 2008. To appear. Preliminary version presented at *First Interaction and Concurrency Experience (ICE'08)*, Reykjavik, July 2008.
- [HR04] Matthew Hennessy and Julian Rathke. Typed behavioural equivalences for processes in the presence of subtyping. *Mathematical Structures in Computer Science*, 14:651–684, 2004.
- [HT92] Kohei Honda and Mario Tokoro. On asynchronous communication semantics. In P. Wegner M. Tokoro, O. Nierstrasz, editor, *Proceedings of the ECOOP '91 Workshop on Object-Based Concurrent Computing*, volume 612 of *LNCS 612*. Springer-Verlag, 1992.
- [KAK05] Astrid Kiehn and Sak Arun-Kumar. Amortised bisimulations. In Farn Wang, editor, *FORTE*, volume 3731 of *Lecture Notes in Computer Science*, pages 320–334. Springer, 2005.
- [LV06] Gerald Lüttgen and Walter Vogler. Bisimulation on speed: a unified approach. *Theor. Comput. Sci.*, 360(1):209–227, 2006.
- [Mil89] Robin Milner. *Communication and Concurrency*. Prentice-Hall, 1989.
- [Mil99] Robin Milner. *Communicating and mobile systems: the π -calculus*. Cambridge University Press, 1999.
- [SW01] Davide Sangiorgi and David Walker. *The π -calculus: A Theory of Mobile Processes*. Cambridge University Press, 2001.
- [Tel04] David Teller. Recollecting resources in the pi-calculus. In *Proceedings of IFIP TCS 2004*, pages 605–618. Kluwer Academic Publishing, 2004.
- [Tof94] Chris M. N. Tofts. Processes with probabilities, priority and time. *Formal Asp. Comput.*, 6(5):536–564, 1994.