

D-Fusion: a Distinctive Fusion Calculus

Michele Boreale¹, Maria Grazia Buscemi², and Ugo Montanari²

¹ Dipartimento di Sistemi e Informatica, Università di Firenze, Italy.

² Dipartimento di Informatica, Università di Pisa, Italy.

boreale@dsi.unifi.it {buscemi, ugo}@di.unipi.it

Abstract. Fusion calculus is commonly regarded as a generalisation of pi-calculus. Actually, we prove that there is no uniform fully abstract embedding of pi-calculus into Fusion. This fact motivates the introduction of a new calculus, D-Fusion, with two binders, λ and ν . We show that D-Fusion is strictly more expressive than both pi-calculus and Fusion. The expressiveness gap is further clarified by the existence of a fully abstract encoding of mixed guarded choice into the choice-free fragment of D-Fusion.

1 Introduction

A recent trend in the design of certain distributed applications like Web services [13] or business-to-business systems [5] based on XML sees the emergence of a message-passing programming style. Languages like Highwire [4] provide, in a concurrency setting, sophisticated data structures; these allow programmers to describe and manipulate complex messages and interaction patterns. If one looks for ‘foundational’ counterparts of these programming languages, both the pi-calculus [6, 7] and the Fusion calculus [11], seem very promising candidates. Both of them, indeed, convey the idea of message-passing in a distilled form, and come equipped with a rich and elegant meta-theory.

The main novelty of Fusion when compared to the pi-calculus is the introduction of *fusions*. A fusion is a name equivalence that, when applied onto a term, has the effect of a (possibly non-injective) name substitution. Fusions are ideal for representing, e.g., forwarders for objects that migrate among locations [2], or forms of pattern matching between pairs of messages [4]. Computationally, a fusion is generated as a result of a synchronisation between two complementary actions, and it is atomically propagated to processes running in parallel with the active one. This happens in much the same way as, in logic programming, term substitutions resulting from a unification step on a subgoal can be forced on the other subgoals.

If compared to pi-calculus name-passing, fusions enable a more general name matching mechanism during synchronisation. However, differently from the pi-calculus, the binding mechanism of Fusion just ignores the issue of unicity of newly generated names. One of the goals of this paper is to show that this fact severely limits the expressiveness of Fusion. Overcoming this limitation calls for co-existence of two name binders, λ and ν : the former analogous to the only binder of Fusion, and the latter

Research supported in part by FET Global project *PROFUNDIS*.

imposing unicity. This implies combining, in a consistent way, the somehow conflicting mechanisms of fusions and *distinctions* à la open pi-calculus [12]. The resulting *distinctive* Fusion calculus, or *D-Fusion*, is at least as expressive as pi-calculus and Fusion separately, and in fact, we strongly argue, *more* expressive than both. A more precise account of our work follows.

The binding mechanism of pi-calculus generalises that of λ -calculus in several ways. Input prefix binds like λ , and name passing takes place in pi-calculus in a way typical of functional programming, i.e., formal names are assigned their actual counterpart. The *restriction* binder ν , however, is very different from λ , as a restricted name can be exported (extruded), with the guarantee that it will never be identified to anything else. Open pi-calculus [12] takes an important step forward over the original proposal, allowing for input actions whose formal parameters can be instantiated, hence ‘fused’, in a ‘lazy’ fashion – at any time when needed for synchronisation. In order to preserve the intended semantics of restriction, however, newly extruded names are kept distinct from ‘old’ names, using special relations called distinctions.

Fusion calculus is presented in [11] as a more uniform and more expressive version of pi-calculus. The main idea is to decompose input prefix $a(x)$ into a binder (x) and a prefix $a\langle x \rangle$. In the polyadic case, matching between the input list and the output list of variables induces name unification, i.e. a fusion. The latter is propagated across processes, but one or more binders can be used to control the scope (i.e. propagation) of the fusion. Thus one achieves both a perfect symmetry between input and output and a more general name passing mechanism.

At first sight, Fusion is more general than pi-calculus. And, indeed, in [11] it is stated that the pi-calculus transition system can be embedded into Fusion’s, provided that one identifies restriction (νx) with the (x) binder of Fusion.

Our first move is to argue that this embedding breaks down if comparing the two calculi on the basis of abstract semantics. We prove that no ‘uniform’ encoding exists of pi-calculus into Fusion that preserves any ‘reasonable’ behavioural equivalence (at least as fine as trace equivalence). Here ‘uniform’ means homomorphic with respect to parallel composition and name substitution, and mapping (νx) to (x) , and preserving (a subset of) weak traces. As hinted before, the ultimate reason is that in Fusion all names are like logical variables, i.e., their unification always succeeds.

The above failure motivates the introduction of a new calculus, D-Fusion, with two binders, λ and ν : the first generalises input prefix, and the second models restriction. Also, any issue of symmetry between input and output is preempted, since we have just one kind of prefix (no polarisation); polarised prefixes can be easily encoded, though. As in open pi-calculus, in D-Fusion distinctions are employed to constrain possible fusions. In logical terms, this corresponds to allow unification not only among variables, but also among variables and dynamically generated constants (that is, ν -extruded names). However, unification fails whenever one tries to identify two distinct constants, or to identify a ‘recent’ constant with an ‘old’ variable. We show that the additional expressive power achieved in this way is relevant. Both pi-calculus and Fusion can be uniformly encoded into D-Fusion. Moreover, the combined mechanism of restriction and unification yields additional expressive power: it allows to express a form of pat-

tern matching which cannot be expressed in the other two calculi. As a consequence, we prove, D-Fusion cannot be uniformly encoded neither into Fusion, nor into pi-calculus.

Next, the gap between D-Fusion and Fusion/pi-calculus is explored from a more concrete perspective. First, we exhibit a simple security protocol and a related *correlation* property that are readily translated into D-Fusion. The property breaks down if uniformly translating the protocol into Fusion. The failure is illuminating: in Fusion, one has no way of declaring unique fresh names to correlate different messages of the protocol.

Palamidessi has shown [9, 10] that nondeterministic *guarded choice* cannot be simulated in the choice (+) -free pi-calculus in a fully abstract way, while preserving any ‘reasonable’ semantics. This is due to the impossibility of atomically performing, in the absence of +, an external synchronisation and an internal exclusive choice among a number of alternatives. We prove that in D-fusion, under mild typing assumptions, guarded choice can actually be simulated in a fully abstract way in the choice-free fragment. The encoding preserves a reasonable semantics in the sense of [9, 10]. Informally, branches of a choice are represented as concurrent processes. Synchronisation is performed in the ordinary way, but it forces a fusion between a global λ -name and a ν -name local to the chosen branch. Excluded branches are atomically inhibited, since any progress would lead them to fusing two distinct ν -names.

The rest of the paper is organised as follows. Section 2 contains the proof that the pi-calculus cannot be uniformly encoded into Fusion. In Section 3 we introduce the D-Fusion calculus and in Section 4 we give a notion of bisimulation for the calculus, and characterise it via a more handy symbolic semantics; we also introduce a notion of *sorting* and sorted bisimulation. In Section 5 we show that the D-Fusion calculus is more expressive than both pi-calculus and Fusion. We further explore this expressiveness gap in Sections 6 by means of an example and in Section 7 by encoding the mixed guarded choice into the choice-free calculus. Finally, Section 8 contains a brief overview of some related work and a few concluding remarks.

2 Fusion and Pi

The aim of this section is illustrating the difference between pi-calculus and Fusion, and to show that the former cannot be uniformly encoded in the latter.

The crucial difference between the pi-calculus and Fusion shows up in synchronisations: in Fusion, the effect of a synchronisation is not necessarily local, and is regulated by the scope of the binder (x). For example, an interaction between $\bar{u}v.P$ and $ux.Q$ will result in a fusion of ν and x . This fusion will also affect any further process R running in parallel, as illustrated by the example below:

$$R|\bar{u}v.P|ux.Q \xrightarrow{\{x=\nu\}} (R|P|Q)[\nu/x].$$

The binding operator (x) can be used to limit the scope of the fusion, e.g.:

$$R|(x)(\bar{u}v.P|ux.Q) \xrightarrow{\tau} R|(P|Q)[\nu/x].$$

where τ denotes the identity fusion. For a full treatment of pi-calculus and Fusion we refer to [7] and to [11], respectively.

Below, we show that there is no ‘reasonably simple’ encoding of the pi-calculus Π into Fusion \mathcal{F} . We focus on encodings $\llbracket \cdot \rrbracket$ that have certain compositional properties and preserve (a subset of) weak traces. As to the latter, we shall only require that moves of P are reflected in $\llbracket P \rrbracket$, not vice-versa. We also implicitly require that the encoding preserves arity of I/O actions (the length of tuples carried on each channel). This is sometimes not the case for process calculi encodings; however, it is easy to relax this condition by allowing names of $\llbracket P \rrbracket$ to carry *longer* tuples. We stick to the simpler correspondence just for notational convenience. Note that the encoding presented in [11] does satisfy our criterion. We shall assume here the standard pi-calculus late operational semantics [7] and, for the purpose of our comparison, we shall identify the late input pi-action $a(\tilde{x})$ with the Fusion input action $(\tilde{x})a\tilde{x}$.

Definition 1. A translation $\llbracket \cdot \rrbracket : \Pi \rightarrow \mathcal{F}$ is uniform if for each $P, Q \in \Pi$ it holds that:

- for each trace of actions s not containing bound outputs, $P \xrightarrow{s}$ implies $\llbracket P \rrbracket \xrightarrow{s}$;
- $\llbracket P|Q \rrbracket = \llbracket P \rrbracket | \llbracket Q \rrbracket$;
- for each y , $\llbracket (\nu y)P \rrbracket = (y)\llbracket P \rrbracket$;
- for each substitution σ , $\llbracket P\sigma \rrbracket = \llbracket P \rrbracket\sigma$.

The next proposition generalises an example from [11]. Below, we fix an arbitrary Π -equivalence included in trace equivalence, \sim_{Π} , and an arbitrary \mathcal{F} -equivalence which is included in trace equivalence *and* is preserved by parallel composition, $\sim_{\mathcal{F}}$ (like, e.g., hyperequivalence of [11]).

Proposition 1. There is no uniform translation $\llbracket \cdot \rrbracket : \Pi \rightarrow \mathcal{F}$ such that for each $P, Q \in \Pi$:

$$P \sim_{\Pi} Q \text{ implies } \llbracket P \rrbracket \sim_{\mathcal{F}} \llbracket Q \rrbracket.$$

PROOF: Suppose that there exists such a translation $\llbracket \cdot \rrbracket$. Let P and Q be the following two pi-agents:

$$\begin{aligned} P &= (\nu u, v) (\bar{a}\langle u, v \rangle | \bar{u}|v.\bar{w}) \\ Q &= (\nu u, v) (\bar{a}\langle u, v \rangle | (\bar{u}.(v.\bar{w}) + v.(\bar{u}|\bar{w}))). \end{aligned}$$

Obviously, $P \sim_{\Pi} Q$ (e.g. they are strongly late bisimilar). Suppose $\llbracket P \rrbracket \sim_{\mathcal{F}} \llbracket Q \rrbracket$. Let $R = a(x, y).(\bar{c}x|cy)$ and A and B be as follows:

$$\begin{aligned} A &= \llbracket P \rrbracket | R = (u, v)(\llbracket \bar{a}\langle u, v \rangle \rrbracket | \llbracket \bar{u} \rrbracket | \llbracket v.\bar{w} \rrbracket) | R \\ B &= \llbracket Q \rrbracket | R = (u, v)(\llbracket \bar{a}\langle u, v \rangle \rrbracket | \llbracket \bar{u}.(v.\bar{w}) + v.(\bar{u}|\bar{w}) \rrbracket) | R. \end{aligned}$$

Since $\sim_{\mathcal{F}}$ is preserved by $|$, A and B are $\sim_{\mathcal{F}}$ -equivalent. By uniformity of the encoding, it is easy to check that $A \xrightarrow{\bar{w}}$. On the other hand, a careful case analysis shows that $B \not\xrightarrow{\bar{w}}$. This is a contradiction. \square

3 The Distinctive Fusion Calculus, D-Fusion

Syntax We consider a countable set of names \mathcal{N} ranged over by $a, b, \dots, u, v, \dots, z$. We write \tilde{x} for a finite tuple x_1, \dots, x_n of names. The set \mathcal{DF} of D-Fusion *processes*, ranged over by P, Q, \dots , is defined by the syntax:

$$P ::= \mathbf{0} \mid \alpha.P \mid P|P \mid P+P \mid [x=y]P \mid !P \mid \lambda x P \mid (\nu x)P$$

where *prefixes* α are defined as $\alpha ::= a\tilde{v}$. The occurrences of x in $\lambda x P$ and $(\nu x)P$ are *bound*, thus notions of *free names* and *bound names* of a process P arise as expected and are denoted by $\text{fn}(P)$ and $\text{bn}(P)$, respectively. The notion of *alpha-equivalence* also arises as expected. In the rest of the paper we will identify alpha-equivalent processes.

Note that we consider one kind of prefix, thus ignoring polarities. However, a sub-calculus with polarities can be easily retrieved, as shown at the end of this section.

The main difference from Fusion is the presence of two distinct binding constructs, λ and ν . The λ -abstraction operator corresponds to the binding construct of Fusion and generalises input binding of the pi-calculus. The restriction operator (ν) corresponds to the analogous operator of the pi-calculus: it allows a process to create a fresh, new name that will be kept distinct from other names.

Operational Semantics For R a binary relation over \mathcal{N} , let R^* denote the reflexive, symmetric and transitive closure of R with respect to \mathcal{N} . We use σ, σ' to range over substitutions, i.e. finite partial functions from \mathcal{N} onto \mathcal{N} . Domain and co-domain of σ , denoted $\text{dom}(\sigma)$, $\text{cod}(\sigma)$ are defined as expected. We denote by $t\sigma$ the result of applying σ onto a term t . Given a set/tuple of names \tilde{x} , we define $\sigma_{\tilde{x}}$ as $\sigma \cap (\tilde{x} \times \mathcal{N})$, and $\sigma_{-\tilde{x}}$ as $\sigma - (\tilde{x} \times \mathcal{N})$.

Definition 2 (fusions). We let ϕ, χ, \dots range over fusions, that is total equivalence relations on \mathcal{N} with only finitely many non-singular equivalence classes. We let:

- $\text{n}(\phi)$ denote $\{x : x\phi y \text{ for some } y \neq x\}$;
- τ denote the identity fusion (i.e., $\text{n}(\tau) = \emptyset$);
- $\phi + \psi$ denote the finest fusion which is coarser than ϕ and ψ , that is $(\phi \cup \psi)^*$;
- ϕ_{-z} denote $(\phi - (\{z\} \times \mathcal{N} \cup \mathcal{N} \times \{z\}))^*$;
- $\{x=y\}$ denote $\{(x, y)\}^*$;
- $\phi[x]$ denote the equivalence class of x in ϕ .

A fusion ϕ arises as the result of equating two lists of names in a synchronisation. In general, names in the original lists might be either λ -abstracted or ν -extruded, or free. Informally, the effect of ϕ should be that of a substitution, mapping equivalent names onto representatives. However, in the presence of ν -, λ -abstracted and free ϕ -equivalent names, some care is needed to ensure that the representative be chosen appropriately. This concept is made precise in the next definition.

Definition 3 (induced substitutions). Let \tilde{x} and \tilde{k} be two disjoint tuples of names and ϕ be a fusion such that for any two distinct $k_1, k_2 \in \tilde{k}$, not $(k_1 \phi k_2)$. We say that a substitution σ is induced by $\lambda\tilde{x}, (\nu\tilde{k})$ and ϕ , if $\text{dom}(\sigma) = \text{n}(\phi)$ and for each equivalence class E of ϕ , σ maps all names in E onto one and the same name $y \in E$ such that:

- if $\tilde{k} \cap E = \{h\}$ then $y = h$;
- if $\tilde{k} \cap E = \emptyset$ and $E - \tilde{x} \neq \emptyset$ then $y \in E - \tilde{x}$.

For example, suppose ϕ has only two non-singular equivalence classes $\{x, w, k\}$ and $\{y, h\}$, then a substitution σ induced by λxy , (vk) and ϕ is the one mapping x, w, k to k and y, h to h . It is important to stress that, by definition, no induced σ exists if for distinct $k_1, k_2 \in \tilde{k}$ it holds $k_1 \phi k_2$. Also note that in general there may be more than one induced substitution. E.g. $\{a = b\}$ (with empty $\lambda\tilde{x}$ and $(v\tilde{k})$) induces both $[a/b]$ and $[b/a]$. We introduce the labelled transition system for D-Fusion.

Definition 4 (labelled transition system). *The transition relation $P \xrightarrow{\mu} Q$, for μ a label of the form $(v\tilde{x})\lambda\tilde{y}a\tilde{v}$ (action) or of the form $(v\tilde{x})\phi$ (effect), is defined in Table 1.*

Some notations for actions and effects. The bound names of μ are written $\text{bn}(\mu)$ and are defined as expected, while $\text{subj}(\mu)$ and $\text{obj}(\mu)$ denote the subject and object part of μ , if μ is an action, otherwise they both denote conventional value ‘–’. Moreover, $\text{n}(\mu)$ denotes all the names of μ . We use abbreviations such as $\text{n}(\phi, \mu)$ to denote $\text{n}(\phi) \cup \text{n}(\mu)$ and $(vz)\mu$ for $(v\tilde{x}z)\phi$, if $\mu = (v\tilde{x})\phi$. Furthermore, we shall identify actions and effects up to reordering of the tuple \tilde{x} in $(v\tilde{x})$ and $\lambda\tilde{x}$.

The rules in Table 1 deserve some explanation. As mentioned, we have two kinds of labels, actions and effects. Apart from the absence of polarities, *actions* are governed by rules similar to those found in pi-calculus. The main difference is that on the same action one can find both v- and λ -extruded names. On the other hand, *effects* are similar to those found in Fusion. A major difference is that our effects can also v-extrude names. An effect $(v\tilde{x})\phi$ can be created as a result of a communication (rule COM), and be propagated across parallel components, until a λ that binds a fused name z is encountered (rule $\lambda\text{-OPEN}_r$). At that point, the corresponding substitution $\sigma_{|z}$ is applied onto the target process and z is discarded from the fusion (the result is ϕ_{-z}). Any v-extruded name $w \in \tilde{x}$ which is equivalent to z must be v-closed back, provided that no free name is in that equivalence class. In rule COM , handling of λ - and v-extruded names can be explained similarly, but note that those λ -extruded names \tilde{z} that are not instantiated (by $\sigma_{|\tilde{y}\tilde{y}'}$) must be λ -closed (this happens when in some equivalence class all names are λ -abstracted). Finally, note that the side condition ‘ $\phi[z] \cap \tilde{x} = \emptyset$ ’ in rule v-OPEN prevent effects from equating two distinct v-extruded names.

Let us illustrate the rules with some examples. We shall use $\{\tilde{x} = \tilde{y}\}.P$ as an abbreviation for $(vc)(c\tilde{x}|c\tilde{y}.P)$ (for a fresh name c) and we write $\tau.P$ when \tilde{x} and \tilde{y} are empty.

Example 1.

1. Let $P = (vx)(vc)(cx.P_1 | cy.P_2)$. The interaction between $cx.P_1$ and $cy.P_2$ will result into a fusion $\{x = y\}$, that causes x to be extruded:

$$(vx)(vc)(cx.P_1 | cy.P_2) \xrightarrow{(vx)\{x=y\}} (vc)(P_1 | P_2).$$

Now consider $Q = \lambda y P$. The effect of λ -abstracting y in P is that of removing ϕ and getting the induced substitution $[x/y]$ applied onto the continuation:

$$\lambda y (vx)(vc)(cx.P_1 | cy.P_2) \xrightarrow{\tau} (vx)((vc)(P_1 | P_2)[x/y]).$$

$$\begin{array}{c}
\text{(ACT)} \quad \alpha.P \xrightarrow{\alpha} P \qquad \text{(MATCH)} \quad \frac{P \xrightarrow{\mu} Q}{[a = a]P \xrightarrow{\mu} Q} \qquad \text{(SUM)} \quad \frac{P_1 \xrightarrow{\mu} Q}{P_1 + P_2 \xrightarrow{\mu} Q} \\
\text{(v-PASS)} \quad \frac{P \xrightarrow{\mu} Q}{(vz)P \xrightarrow{\mu} (vz)Q} \quad z \notin \mathfrak{n}(\mu) \qquad \text{(\lambda-PASS)} \quad \frac{P \xrightarrow{\mu} Q}{\lambda z P \xrightarrow{\mu} \lambda z Q} \quad z \notin \mathfrak{n}(\mu) \\
\text{(v-OPEN)} \quad \frac{P \xrightarrow{\mu} Q}{(vz)P \xrightarrow{(vz)\mu} Q} \quad \left\{ \begin{array}{l} z \in \mathfrak{n}(\mu) \\ \mu \text{ an action implies } z \neq \text{subj}(\mu) \\ \mu = (v\tilde{x})\phi \text{ implies } \phi[z] \cap \tilde{x} = \emptyset \end{array} \right. \\
\text{(\lambda-OPEN}_a\text{)} \quad \frac{P \xrightarrow{(v\tilde{x})\lambda\tilde{y}a\tilde{v}} Q}{\lambda z P \xrightarrow{(v\tilde{x})\lambda\tilde{y}z a\tilde{v}} Q} \quad z \in \tilde{v} - (\{a\} \cup \tilde{x}\tilde{y}) \\
\text{(\lambda-OPEN}_f\text{)} \quad \frac{P \xrightarrow{(v\tilde{x})\phi} Q}{\lambda z P \xrightarrow{(v\tilde{x}')\phi_{-z}} (v\hat{w})(Q\sigma_{|z})} \quad \left\{ \begin{array}{l} z \in \mathfrak{n}(\phi) \\ \sigma \text{ induced by } \lambda z, (v\tilde{x}) \text{ and } \phi \\ \tilde{x}' = \tilde{x} \cap \mathfrak{n}(\phi_{-z}) \\ \hat{w} = \tilde{x} - \tilde{x}' \end{array} \right. \\
\text{(COM)} \quad \frac{P_1 \xrightarrow{(v\tilde{x}')\lambda\tilde{y}'a\tilde{u}} Q_1 \quad P_2 \xrightarrow{(v\tilde{x}'')\lambda\tilde{y}''a\tilde{v}} Q_2}{P_1 | P_2 \xrightarrow{(v\tilde{x})\{\tilde{v}=\tilde{u}\}_{-\tilde{y}'\tilde{y}''}} (v\tilde{w})\lambda\tilde{z}((Q_1 | Q_2)\sigma_{|\tilde{y}'\tilde{y}''})} \quad \left\{ \begin{array}{l} |\tilde{u}| = |\tilde{v}| \\ \sigma \text{ induced by } \lambda\tilde{y}'\tilde{y}'', (v\tilde{x}'\tilde{x}'') \text{ and } \{\tilde{v} = \tilde{u}\} \\ \tilde{z} = \tilde{y}'\tilde{y}'' \cap \text{cod}(\sigma_{|\tilde{y}'\tilde{y}''}) \\ \tilde{x} = \tilde{x}'\tilde{x}'' \cap \mathfrak{n}(\{\tilde{v} = \tilde{u}\}_{-\tilde{y}'\tilde{y}''}) \\ \tilde{w} = \tilde{x}'\tilde{x}'' - \tilde{x} \end{array} \right. \\
\text{(PAR)} \quad \frac{P \xrightarrow{\mu} Q}{P | R \xrightarrow{\mu} Q | R} \qquad \text{(REP)} \quad \frac{P | P \xrightarrow{\mu} Q}{!P \xrightarrow{\mu} Q}
\end{array}$$

Symmetric rules for (SUM) and (PAR) are not shown. Usual conventions about freshness of bound names apply.

Table 1. Actions and effects transitions in D-Fusion.

It is worth noting that our operational rules are preserved by a form of structural equivalence. For instance, let $R = (\nu x)\lambda y(\nu c)(cx.P_1 | cy.P_2)$. R has the same τ -transition as Q above.

2. Another example is as follows (P below is some continuation):

$$\begin{aligned} \lambda z(\nu w)\{z, a = w, b\}.P &\xrightarrow{\{a=b\}} (\nu w)P[w/z], \quad \text{but} \\ \lambda z(\nu w)\{z, z = w, b\}.P &\xrightarrow{(\nu w)\{w=b\}} P[w/z]. \end{aligned}$$

Encoding I/O polarities We can encode polarities as follows:

$$\bar{c}(\tilde{v}).P = (\nu x)\lambda y c\tilde{v}xy.P \quad c(\tilde{v}).P = (\nu x)\lambda y c\tilde{v}yx.P$$

for some chosen fresh x and y . The position of the ν -name x forbids fusions between actions with the same polarity and, hence, communication. For instance, the process $\bar{c}(\tilde{v}).P | \bar{c}(\tilde{u}).Q$ has no τ -transition, since the latter would force the fusion of two distinct ν -names, which is forbidden by the operational rules. We denote by \mathcal{DF}^P , *polarised D-Fusion*, the subset of \mathcal{DF} in which every prefix can be interpreted as an input or output, in the above sense.

4 Bisimulation

Like in the case of Fusion, a ‘natural’ semantics of D-Fusion is required to be closed under substitutions. However, one should be careful in keeping ν -extruded names distinct from others (λ -extruded and free). To this purpose, we introduce below a concept of *distinction*, akin to [12]. Distinctions keep track of ν -extruded names, for which fusions should be forbidden. Differently from [12], however, our distinctions keep track of λ -extruded names as well, and of the order of their extrusion relative to ν -extruded names. The idea is to treat ν -extruded names as constants, while allowing a fusion between two λ -extruded names, or a fusion of a ‘recent’ λ -extruded name to an ‘old’ ν -extruded name. Technically, we find it convenient to model distinctions as sequences.

Definition 5 (distinctions). A distinction D is a sequence of labelled names of the form $x_1^{e_1} \cdots x_n^{e_n}$, with labels $e_i \in \{\nu, \lambda\}$ and x_i pairwise distinct, for $1 \leq i \leq n$. We let $\lambda(D)$ (resp. $\nu(D)$) denote the set of λ -labelled (resp. ν -labelled) names in D and let $\mathfrak{n}(D) = \lambda(D) \cup \nu(D)$. Given a substitution σ and a distinction $D = x_1^{e_1} \cdots x_n^{e_n}$, we say that σ respects D , written $\sigma \vdash D$, if for each $x_i^{e_i}$ in D it holds that: $x_i\sigma = x_i$ and moreover for each $x_j^{e_j}$ in D , $x_j\sigma = x_i$ implies $j > i$. We let $D\sigma$ denote the result of applying σ to D and then erasing all duplicated (i.e. following the leftmost) occurrences of names in the sequence.

Given a fusion ϕ and a distinction D , we say that ϕ and D induce σ , written $\phi, D \rightsquigarrow \sigma$, if σ is a substitution induced by $\lambda\tilde{x}, (\nu\tilde{k})$ and ϕ , where $\tilde{x} = \lambda(D)$ and $\tilde{k} = \nu(D)$, and σ respects D .

Example 2. Consider $D = x_0^\lambda \cdot x_1^\nu \cdot x_2^\lambda \cdot x_3^\lambda \cdot x_4^\nu$ and $\phi = \{x_1, x_2, x_4 = x_2, x_3, x_5\}$. We have that $\phi, D \rightsquigarrow \sigma$, where σ maps x_1, x_2, x_3 to x_1 , and x_5 to x_4 : in fact, we have that $\sigma \vdash D$.

Note that $D\sigma = x_0^\lambda \cdot x_1^\nu \cdot x_4^\nu$. On the other hand, if one considers D and $\{x_4 = x_0\}$ then there is no σ' induced by this effect and D : indeed, σ' should have to map x_0 to x_4 , but this implies $\sigma' \not\vdash D$.

We use the following abbreviation: $P \xrightarrow{(\nu\bar{x})\lambda\bar{y}(\alpha,\phi)} P'$ means that either $P \xrightarrow{(\nu\bar{x})\lambda\bar{y}\alpha} P'$ and $\phi = -$, or $P \xrightarrow{(\nu\bar{x})\phi} P'$ and $\alpha = -$ (here $-$ is just a conventional 'null' value). Moreover, we stipulate that $- , D \rightsquigarrow \text{id}$, where id is the identity substitution.

Definition 6 (open bisimulation). A set $\mathcal{R} = \{R_D\}_D$ of process relations indexed by distinctions is an indexed simulation if for each D , whenever $P R_D Q$ and $P \xrightarrow{(\nu\bar{x})\lambda\bar{y}(\alpha,\phi)} P'$ and $\phi, D \cdot \bar{x}^\nu \rightsquigarrow \sigma$, then a transition $Q \xrightarrow{(\nu\bar{x})\lambda\bar{y}(\alpha,\phi)} Q'$ exists such that $P'\sigma R_{D'} Q'\sigma$, where $D' = D\sigma \cdot \bar{x}^\nu \cdot \bar{y}^\lambda$.

\mathcal{R} is an indexed bisimulation if both \mathcal{R} and $\mathcal{R}^{-1} = \{R_D^{-1}\}_D$ are indexed simulations. Open bisimulation, \sim , is the largest indexed bisimulation preserved by respectful substitutions, i.e.: for each σ and D , if $P \sim_D Q$ and σ respects D then $P\sigma \sim_{D\sigma} Q\sigma$.

When no confusion arises, we write $P \sim Q$ for $P \sim_\varepsilon Q$, where ε is the empty distinction.

Proposition 2 (basic properties of \sim_D). Let P and Q be two processes. Then:

- $P \sim_D Q$ and $x \notin n(D)$ imply $\lambda x P \sim_D \lambda x Q$.
- $P \sim_D Q$ implies $(\nu x) P \sim_{D-x} (\nu x) Q$.

Prefix, parallel composition, matching, sum and replication operators preserve \sim_D .

Example 3. 1. An example of 'expansion' for parallel composition is as follows:

$$((\nu k)ak.ak)|av \sim (\nu k)ak.(ak.av + av.ak + \{k = v\}) + av.((\nu k)ak.ak) + (\nu k)\{k = v\}.ak.$$

Note that, after the extrusion of k , free name v can *still* be fused to k . On the contrary, if we consider a distinction D and let $v \in n(D)$, no fusion at all is possible:

$$((\nu k)ak.ak)|av \sim_D (\nu k)ak.(ak.av + av.ak) + av.((\nu k)ak.ak).$$

2. The following two examples show the effect of fusing a λ -abstracted name with a free name and with another λ -abstracted name, respectively:

$$\lambda v \{k = v\}.P \sim \tau.P[k/v] \quad (\lambda k, v) \{k = v\}.P \sim \lambda k \tau.P[k/v].$$

Symbolic bisimulation The definition of open bisimulation contains an implicit universal quantification over substitutions, which makes it not handy. The next definition and the subsequent theorem provide us with an alternative formulation of this equivalence akin to the 'efficient characterisation' of [12]. The new formulation is based on symbolic transitions of the form $P \xrightarrow{\mu}_\chi Q$, where χ is a logical condition necessary for the transition to take place. For technical convenience, we shall represent conditions as fusions. The defining clauses for $\xrightarrow{\mu}_\chi$ are obtained from those in Table 1 by simply

attaching the appropriate χ . In particular, action prefix is defined as $\alpha.P \xrightarrow{\alpha}_\tau P$, while communication and matching rules are defined as follows:

$$(\text{MATCH}_S) \frac{P \xrightarrow{\mu}_\chi Q}{[a=b]P \xrightarrow{\mu}_{\chi+\{a=b\}} Q} \quad (\text{COM}_S) \frac{P_1 \xrightarrow{(v\tilde{x}')\lambda\tilde{y}'a\tilde{u}}_{\chi_1} Q_1 \quad P_2 \xrightarrow{(v\tilde{x}'')\lambda\tilde{y}''b\tilde{v}}_{\chi_2} Q_2}{P_1|P_2 \xrightarrow{(v\tilde{x})\{\tilde{v}=\tilde{u}\}_{-\tilde{y}'\tilde{y}''}}_{\chi_1+\chi_2+\{a=b\}} (v\tilde{w})\lambda\tilde{z}((Q_1|Q_2)\sigma_{|\tilde{y}'\tilde{y}''})},$$

where the side condition of COM_S is the same as in rule COM . The other rules just propagate χ , with the proviso that λ - and ν - extruded names do not occur in χ . Given fusions ϕ and ψ , we write $\phi \models \psi$ if $\psi \subseteq \phi$.

Definition 7 (symbolic bisimulation). A set $\mathcal{R} = \{R_D\}_D$ of process relations indexed by distinctions is an indexed symbolic simulation if whenever $P R_D Q$ and $P \xrightarrow{(v\tilde{x})\lambda\tilde{y}(\alpha,\phi)}_\chi P'$ and $(\phi+\chi), D \cdot \tilde{x}^\nu \rightsquigarrow \sigma$, then a transition $Q \xrightarrow{(v\tilde{x})\lambda\tilde{y}(\beta,\psi)}_{\chi'} Q'$ exists such that:

1. $\chi \models \chi'$,
2. $\chi \models \{\text{subj}(\alpha) = \text{subj}(\beta)\}$ and $\chi \models \{\text{obj}(\alpha) = \text{obj}(\beta)\}$,
3. $\phi + \chi = \psi + \chi$,
4. $P'\sigma R_{D'} Q'\sigma$, where $D' = D\sigma \cdot \tilde{x}^\nu \cdot \tilde{y}^\lambda$.

\mathcal{R} is a indexed symbolic bisimulation if both \mathcal{R} and $\mathcal{R}^{-1} = \{R_D^{-1}\}_D$ are indexed open simulations. The largest indexed symbolic bisimulation is written \sim^{symb} .

Theorem 1. For each P, Q and D , $P \sim_D Q$ if and only if $P \sim_D^{\text{symb}} Q$.

PROOF: For any ϕ and σ , let $\phi\sigma$ denote the fusion $\{(x\sigma, y\sigma) \mid x\phi y\}^*$, and let the fusion induced by σ (seen as a binary relation) be defined as σ^* . The proof of the theorem is based on the following operational correspondences:

1. $P \xrightarrow{(v\tilde{x})\lambda\tilde{y}(\alpha,\phi)}_\chi P'$ implies there is a transition $P\sigma \xrightarrow{(v\tilde{x})\lambda\tilde{y}(\alpha\sigma,\phi\sigma)} P'\sigma$, for each substitution σ induced by χ .
2. $P\sigma \xrightarrow{(v\tilde{x})\lambda\tilde{y}(\beta,\psi)}_{P_1}$ implies there is a fusion χ induced by σ and a transition $P \xrightarrow{(v\tilde{x})\lambda\tilde{y}(\alpha,\phi)}_{\chi'} P'$ s.t. $\chi \models \chi'$, $\beta = \alpha\sigma$ and $\psi = \phi\sigma$ and $P_1 = P'\sigma$.

□

Example 4. Fusions and conditions play pretty different roles in bisimulation semantics. As an example, let $P = [a=b](ca \mid \bar{c}b)$ and $Q = ca \mid \bar{c}b$. Both in P and Q the interaction between ca and $\bar{c}b$ results in a fusion $\{a=b\}$, but $P \not\sim Q$ since P can perform an action only under the condition $\chi = \{a=b\}$, i.e. $P \xrightarrow{\{a=b\}}_{\{a=b\}} \mathbf{0}$ while $Q \xrightarrow{\{a=b\}}_\tau \mathbf{0}$.

Sorting Sorting systems are used in polyadic pi-calculus to discipline the use of channels [6]. We adapt the concept of sorting from pi-calculus, and slightly extend it in order to also control ‘patterns’ that can arise in the object part of actions. In a sorting system \mathbf{S} , the set of names is partitioned into a family of countable sorts, denoted as S, S', \dots . Moreover, there is a sorting function that maps each sort S to a pattern

$(x_1 : S_1, \dots, x_n : S_n)$, written also $\mathbf{S} : S \mapsto \tilde{x} : \tilde{S}$, with the x_i 's not necessarily distinct: the intended meaning is that names of sort S should only carry tuples \tilde{v} of type \tilde{S} s.t. $x_i = x_j$ implies $v_i = v_j$.

Like in polyadic pi-calculus, one can devise typing systems which guarantee that well-typed processes will never perform ‘illegal’ (in the above sense) actions. However, here we are primarily interested in defining the open bisimilarity \sim^S naturally induced by a sorting system \mathbf{S} . Let us say that a substitution σ is \mathbf{S} -*respectful* if, for each x , $x\sigma$ and x belong to the same sort. Similarly, we say a fusion ϕ is \mathbf{S} -*respectful* if each equivalence class is entirely included in one sort (hence any substitution induced by ϕ is respectful). In a \mathbf{S} -sorted D-fusion, for a given sorting system \mathbf{S} , we assume that in every prefix $a\tilde{v}$ each v_i belongs to the appropriate sort, and we *only consider respectful substitutions and fusions*. For $\alpha = a\tilde{v}$ an action, where $a \in S \mapsto \tilde{x} : \tilde{S}'$, the fusion induced by α , denoted ϕ_α , is defined as $\sum_{x_i, x_j \in \tilde{x}, x_i = x_j} \{v_i = v_j\}$ (here ‘ Σ ’ denotes sum of fusions defined in Def. 2). Furthermore, we shall denote by \equiv equivalence over actions induced by the law: $\lambda x \mu \equiv \mu$ if $x \notin n(\mu)$.

With the above conventions and notations, the definition of \sim^S is obtained from Definition 6 by making the clauses of actions and effects explicitly distinct, and changing the clause for *actions* as follows:

$$\text{if } P \xrightarrow{(\tilde{v}\tilde{x})\lambda\tilde{y}a\tilde{v}} P' \text{ and } \phi_{a\tilde{v}}, D \cdot \tilde{x}^v \cdot \tilde{y}^\lambda \rightsquigarrow \sigma \text{ then a transition } Q \xrightarrow{(\tilde{v}\tilde{x})\lambda\tilde{y}'a\tilde{v}'} Q' \text{ exists s.t. } \phi_{a\tilde{v}'}, D \cdot \tilde{x}^v \cdot \tilde{y}'^\lambda \rightsquigarrow \sigma' \text{ and } \lambda\tilde{y}'(a\tilde{v}'\sigma) \equiv \lambda\tilde{y}'(a\tilde{v}'\sigma'), \sigma_{-\tilde{y}} = \sigma'_{-\tilde{y}'} \text{ and } P'\sigma R_{D'} Q'\sigma', \text{ where } D' = (D \cdot \tilde{x}^v \cdot \tilde{y}^\lambda)\sigma.$$

Note that \sim^S is in general more generous than ordinary open bisimilarity \sim . For example, assuming $a \in S \mapsto (x : S', x : S')$, then $(\nu n)\lambda x a x n.P \sim^S (\nu n)\lambda x a n x.P \sim^S (\nu n)ann.P[n/x]$ and $(\nu n)(\nu m)amn.P \sim^S \mathbf{0}$. We will see applications of sorted bisimilarity in Section 7.

5 Expressiveness of D-Fusion

Below, we provide the two obvious uniform and fully abstract translations from Π and \mathcal{F} to \mathcal{DF} . The definition of uniformity can be extended to the case of encodings from \mathcal{F}/Π into \mathcal{DF} in the obvious way: in particular, by requiring that (x) and (νx) be mapped to λx and (νx) , respectively.

Definition 8 ($\llbracket \cdot \rrbracket_\pi$). *The translation $\llbracket \cdot \rrbracket_\pi : \Pi \rightarrow \mathcal{DF}$ is defined by extending in the expected homomorphic way the following clauses:*

$$\llbracket \bar{a}(x).P \rrbracket_\pi = \bar{a}(x).\llbracket P \rrbracket_\pi \quad \llbracket a(x).P \rrbracket_\pi = \lambda x a(x).\llbracket P \rrbracket_\pi \quad \llbracket (\nu x)P \rrbracket_\pi = (\nu x)\llbracket P \rrbracket_\pi.$$

For any distinction $D = x_1^{e_1} \dots x_n^{e_n}$, we can generate an equivalent Π -distinction D_π (that is, an irreflexive, symmetric relation on \mathcal{N} , [12]) by setting $x_i D_\pi x_j$ iff: either $(e_i = e_j = \nu \text{ and } i \neq j)$ or $(e_i = \nu \text{ and } e_j = \lambda \text{ and } j < i)$ or $(e_j = \nu \text{ and } e_i = \lambda \text{ and } i < j)$. It is reasonable to consider here only Π -distinctions that can be generated in this way. For any such Π -distinction D_π , let us denote by $\llbracket D_\pi \rrbracket$ a chosen \mathcal{DF} -distinction that generates D_π . Let $\sim_{D_\pi}^0$ denote open D_π -bisimilarity over Π ([12]).

Proposition 3. $P \sim_{D\pi}^o Q$ iff $\llbracket P \rrbracket_\pi \sim_{\llbracket D\pi \rrbracket} \llbracket Q \rrbracket_\pi$.

We turn now our attention to Fusion.

Definition 9 ($\llbracket \cdot \rrbracket_f$). *The translation $\llbracket \cdot \rrbracket_f : \mathcal{F} \rightarrow \mathcal{DF}$ is defined by extending in the expected homomorphic way the following clauses:*

$$\llbracket \bar{a}\langle x \rangle.P \rrbracket_f = \bar{a}\langle x \rangle.\llbracket P \rrbracket_f \quad \llbracket a\langle x \rangle.P \rrbracket_f = a\langle x \rangle.\llbracket P \rrbracket_f \quad \llbracket (x)P \rrbracket_f = \lambda x \llbracket P \rrbracket_f$$

Let \sim^{he} denote hyper-equivalence over \mathcal{F} (see [11]).

Proposition 4. $P \sim^{\text{he}} Q$ iff $\llbracket P \rrbracket_f \sim \llbracket Q \rrbracket_f$.

Most interesting, we now show that D-Fusion cannot be uniformly encoded into Π . The intuitive reason is that, in D-Fusion, the combined use of fusions and restrictions allows one to express a form of input with pattern matching. This is not possible in Π without breaking atomicity of certain actions (e.g. choice). To show this, we restrict our attention to polarised D-Fusion, \mathcal{DF}^P .

Given $P \in \Pi$ and a trace of actions s , let us write $P \xrightarrow{s}$ if $P \xrightarrow{s'}$ for some trace s' that has the same sequence of subject names, with the same polarity, as s (e.g., $s = a\langle \tilde{x} \rangle \cdot \lambda \tilde{y} \bar{b}\langle \tilde{v} \rangle$ and $s' = a\langle \tilde{z} \rangle \cdot \bar{b}\langle \tilde{w} \rangle$). The reference semantics for Π is again the late operational semantics.

Definition 10. *A translation $\llbracket \cdot \rrbracket : \mathcal{DF}^P \rightarrow \Pi$ is uniform if for each $P, Q \in \mathcal{DF}^P$:*

- for each trace s , $P \xrightarrow{s}$ implies $\llbracket P \rrbracket \xrightarrow{s}$;
- $\llbracket P|Q \rrbracket = \llbracket P \rrbracket | \llbracket Q \rrbracket$;
- for each y , $\llbracket (\nu y)P \rrbracket = (\nu y) \llbracket P \rrbracket$;
- for each substitution σ , $\llbracket P\sigma \rrbracket = \llbracket P \rrbracket \sigma$.

Below, we denote by $\sim_{\mathcal{DF}^P}$ any fixed equivalence over \mathcal{DF}^P which is contained in trace semantics (defined in the obvious way), and by \sim_Π any fixed equivalence over Π which is contained in trace equivalence.

Proposition 5. *There is no uniform translation $\llbracket \cdot \rrbracket : \mathcal{DF}^P \rightarrow \Pi$ such that $\forall P, Q \in \mathcal{DF}^P$:*

$$P \sim_{\mathcal{DF}^P} Q \Rightarrow \llbracket P \rrbracket \sim_\Pi \llbracket Q \rrbracket.$$

PROOF: Suppose that there exists such a translation $\llbracket \cdot \rrbracket$. Let us consider the following two \mathcal{DF}^P -processes P and Q :

$$P = (\nu c, k, h) (c\langle k \rangle.\bar{a}|c\langle h \rangle.\bar{b}|\bar{c}\langle k \rangle) \quad Q = \tau.\bar{a}.$$

It holds that $P \sim Q$ in \mathcal{DF}^P : the reason is that, in P , synchronisation between prefixes $c\langle h \rangle$ and $\bar{c}\langle k \rangle$, which carry different *restricted* names h and k , is forbidden (see rule ν -OPEN). Thus P can only make $c\langle k \rangle$ and $\bar{c}\langle k \rangle$ synchronise, and then perform \bar{a} . Thus, $P \sim_{\mathcal{DF}^P} Q$ holds too.

On the other hand, due to uniformity, it is easily checked that $\llbracket P \rrbracket \xrightarrow{\bar{b}}$, while $\llbracket Q \rrbracket \not\xrightarrow{\bar{b}}$ (the proof of the latter relies on $b \notin \text{fn}(Q)$ and on the uniformity with respect to substitutions). Thus $\llbracket P \rrbracket \not\sim_{\Pi} \llbracket Q \rrbracket$. \square

Of course, it is also true that D-Fusion cannot be uniformly encoded into \mathcal{F} , as this would imply the existence of a uniform fully abstract encoding from Π to \mathcal{F} , which does not exist (Proposition 1).

The conclusion is that there is some expressiveness gap between D-Fusion on one side and the other two calculi on the other side, at least, as far as our simple notion of uniform encoding is concerned. This gap is further explored by means of more elaborate examples in the next two sections.

6 Example: Correlation

This example aims at illustrating the gap between D-Fusion and Fusion from a more concrete perspective. Consider the following simple protocol. An agent A asks a trusted server S for two keys, to be used to access two distinct services (e.g. A might be a proxy requiring remote connections on behalf of two different users). Communication between A and S takes place over an insecure public channel, controlled by an adversary, but it is protected by encryption and challenge-response nonces. Informally, the dialogue between A and S is as follows:

1. $A \rightarrow S : n$
2. $S \rightarrow A : \{n, k\}_{k_S}$
- 1'. $A \rightarrow S : n'$
- 2'. $S \rightarrow A : \{n', k'\}_{k_S}$

Here $\{\cdot\}_{(\cdot)}$ is symmetric encryption and k_S is a secret master key shared by A and S . A simple property of this protocol is that A should never receive k and k' in the wrong order (k' and then k), even in case S accepts new requests before completing old ones. Indeed, nonces n and n' are intended to help A to contrast attackers replying old, compromised keys or trying to get distinct sessions confused. In other words, nonces do *correlate* each request to S with the appropriate reply of S .

Below, we show that the above small protocol and the related ordering property can be readily translated and verified in D-Fusion. Next, we show that the property breaks down when (uniformly) translating the protocol into Fusion.

D-Fusion Encryption is not a primitive operation in D-Fusion. However, in the present case, it is sensible to model an encrypted message $\{n, k\}_{k_S}$ as an output action $\bar{k}_S\langle n, k \rangle$: only knowing the master key k_S , and further specifying a session-specific nonce, it is possible to acquire the key k (similarly for $\{n', k'\}_{k_S}$, of course). Thus, assuming A concludes the protocol with a conventional ‘commit’ action and that p is the public channel, A , S and the whole protocol P might be specified as follows (below, we abbreviate

$\lambda\tilde{x}p\langle\tilde{x}\rangle.X$ as $p\langle\tilde{x}\rangle.X$:

$$\begin{aligned} A &= (\nu n) \left(\overline{p}\langle n \rangle | \lambda y k_S \langle n, y \rangle . (\nu n') \left(\overline{p}\langle n' \rangle | \lambda y' k_S \langle n', y' \rangle . \overline{commit}\langle y, y' \rangle \right) \right) \\ S &= p(x) . \left(\overline{k_S}\langle x, k \rangle | p(x') . \overline{k_S}\langle x', k' \rangle \right) \\ P &= (\nu k_S) (A | S). \end{aligned}$$

Let A_{spec} be the process defined like A , except that the $\overline{commit}\langle y, y' \rangle$ action is replaced by $\overline{commit}\langle k, k' \rangle$, and let $P_{\text{spec}} = (\nu k_S) (A_{\text{spec}} | S)$. The property that A should never receive k and k' in the wrong order is stated as: $P \sim P_{\text{spec}}$.

Informally, equivalence holds true because the second input action in A/A_{spec} , that is $\lambda y' k_S \langle n', y' \rangle$, can only get synchronised with the second output action in P , that is $\overline{k_S}\langle x', k' \rangle$. In fact, n' can be extruded only *after* x has been received, hence the appropriate distinction will forbid fusion of x and n' . The equivalence is formally shown by exhibiting the appropriate symbolic indexed bisimulation, composed basically by all pairs of derivatives P' and Q' s.t. $P \xrightarrow{s} P'$ and Q' is obtained from P' by replacing $\overline{commit}\langle y, y' \rangle$ with $\overline{commit}\langle k, k' \rangle$, with distinctions as given by the bound names in s .

Fusion Suppose P^f and P_{spec}^f are obtained by some uniform encoding of P and P_{spec} above into Fusion. It is not difficult to show that P^f can be ‘attacked’ by an adversary R that gets n and n' and fuse them together, $R = p(x) . (\overline{p}\langle x \rangle | p(y) . \overline{p}\langle y \rangle)$. Formally, for $\alpha = \overline{commit}\langle k', k \rangle$,

$$P^f | R \xrightarrow{\alpha} \text{ and, thus, } P^f | R \not\sim^{\text{he}} P_{\text{spec}}^f | R,$$

which proves that P^f and P_{spec}^f are not hyper-equivalent.

This example illustrates the difficulty of modelling fresh, indistinguishable quantities (nonces) in Fusion. This makes apparent that Fusion is not apt to express security properties based on correlation.

7 Encoding guarded choice

In this section we show how the combined mechanisms of fusions and restrictions can be used to encode different forms of guarded choice *via* parallel composition, in a clean and uniform way.

Informally, different branches of a guarded choice will be represented as concurrent processes. The encodings adds pairs of extra names to the object part of each action: these extra names are used as ‘side-channels’ for atomic coordination among the different branches. We start by looking at a simple example.

Example 5. Consider the guarded choice $A = \lambda x (\nu n) a \langle xn \rangle . P + \lambda x (\nu m) a \langle xm \rangle . Q$. Its intended ‘parallel’ implementation is the process:

$$B = \lambda x \left((\nu n) a \langle xn \rangle . P | (\nu m) a \langle xm \rangle . Q \right)$$

(here, $x, n, m \notin \text{fn}(a, P, Q)$). Assume now a sorting discipline by which messages traveling on channel a must carry two identical names (i.e. $a \in S \mapsto (x : S', x : S')$, according

to Section 4). In B , the parallel component that first consumes any such message, forces fusion of x either to n or to m , and consequently inhibits the other component. E.g.:

$$\lambda u \bar{a}(uu)|B \xrightarrow{\tau} \sim (vn)(P|(vm)a\langle mn\rangle.Q) \sim P|(vn,m)a\langle mn\rangle.Q.$$

Under the mentioned sorting assumption, $(vm,n)a\langle mn\rangle.Q$ is equivalent to $\mathbf{0}$, because there is no way of fusing m and n . Thus the process on the right of \sim is equivalent to P . In other words, choice between P and Q has been resolved atomically.

The above line of reasoning can be formalised in two ways. One way is considering bisimilarity induced by the above sorting, say \sim^S , which only takes into account transitions that obey the sorting discipline (the rightmost two names in the object part of a/\bar{a} -actions are the same). The other way is keeping standard bisimilarity \sim , while inserting processes inside a ‘firewall’ that filters out a -messages not respecting the given sorting. The latter can be easily defined in D-Fusion relying on ‘non-linear’ inputs:

$$F_{a,a'}[\cdot] = (\nu a') \left(\lambda z a z z. \bar{a}' \langle z z \rangle | [\cdot] \right).$$

We state the result in both forms below.

Proposition 6. *Let A and B be as in Example 5.*

1. $A \sim^S B$;
2. Let A' and B' be the processes obtained from A and B , respectively, by replacing the outermost occurrences of a with a fresh a' . Then $F_{a,a'}[A'] \sim F_{a,a'}[B']$.

Note that the result above exploits in a crucial way features of both Fusion (non-linearity of input actions, in the firewall, and sharing of input variable x , in B) and of D-Fusion (restricted input).

Proposition 6 can be generalised to fully abstract encodings of different forms of guarded choice. For the sake of simplicity, we will state the results in terms of sorted bisimilarity. We believe the results can also be stated in terms of untyped bisimilarity, at the cost of breaking uniformity of the encoding and of introducing more sophisticated forms of ‘firewalls’. We examine two cases, input-guarded choice and mixed choice.

Input-guarded (ig) choice Let us fix, as a source language the fragment of polarised D-Fusion with guarded choice, $\mathcal{DF}^{\text{p,ig}}$. In this language, input prefix and summation $+$ are replaced by input-guarded choice $\sum_{i \in I} a_i \langle \tilde{x}_i \rangle . P_i$. We also assume a sorting system on this source language. The target language is the fragment of polarised D-Fusion with no form of summation. The relevant clauses of the encoding are:

$$\llbracket \sum_{i \in I} a_i \langle \tilde{x}_i \rangle . P_i \rrbracket_{\text{ig}} = \lambda z \Pi_{i \in I} (\nu n) \lambda \tilde{x}_i a_i \langle \tilde{x}_i z n \rangle . \llbracket P_i \rrbracket_{\text{ig}} \quad \llbracket \bar{a} \langle \tilde{v} \rangle . P \rrbracket_{\text{ig}} = \lambda z \bar{a} \langle \tilde{v} z z \rangle . \llbracket P \rrbracket_{\text{ig}},$$

where $\Pi_{i \in I} X_i$ denotes the parallel composition of all X_i 's. The encoding acts as a homomorphism over the remaining operators. Below, $\mathbf{S1}$ indicates the sorting that extends the sorting system \mathbf{S} of the source language by imposing that the two extra object names introduced by the encoding be the same (i.e., if $\mathbf{S}: S \mapsto \tilde{x} : \tilde{S}'$, then $\mathbf{S1}$:

$S \mapsto (\tilde{x} : \tilde{S}', y : S_0, y : S_0)$, for a fresh y and a new sort S_0) and \sim^{S1} is the corresponding open bisimilarity. The proof of the following theorem is straightforward, given that there is a 1-to-1 correspondence between moves of P and moves of $\llbracket P \rrbracket_{\text{ig}}$, under the sortings **S** and **S1**.

Theorem 2 (full abstraction for ig choice). *Let $P, Q \in \mathcal{DF}^{\text{p,ig}}$. It holds that $P \sim^S Q$ if and only if $\llbracket P \rrbracket_{\text{ig}} \sim^{S1} \llbracket Q \rrbracket_{\text{ig}}$.*

Of course, the above theorem also yields a fully abstract encoding of input-guarded choice for pi-calculus, which may be viewed as a sub-calculus of $\mathcal{DF}^{\text{p,ig}}$.

Mixed choice in a sorted pi-calculus As a source language we fix here a sorted version of polyadic pi-calculus [6] with ‘mixed’ choice, Π^{mix} . In this language, prefixes and $+$ are replaced by mixed summation, $\sum_{i \in I} a_i(\tilde{x}_i).P_i + \sum_{j \in J} \bar{b}_j(\tilde{v}_j).Q_j$. The target language is again the fragment of polarised D-Fusion with no summation at all. The encoding is a bit more complex than in the previous case, as it implies adding *two* pairs of extra names to coordinate different branches. The relevant clause is:

$$\begin{aligned} \llbracket \sum_{i \in I} a_i(\tilde{x}_i).P_i + \sum_{j \in J} \bar{b}_j(\tilde{v}_j).Q_j \rrbracket_{\text{mix}} &= \\ (\lambda z, u) (\Pi_{i \in I}(\nu n) \lambda \tilde{x}_i a_i(\tilde{x}_i z n u u). \llbracket P_i \rrbracket_{\text{mix}} \mid \Pi_{j \in J}(\nu n) \bar{b}_j(\tilde{v}_j u u z n). \llbracket Q_j \rrbracket_{\text{mix}}). \end{aligned}$$

Note that the relative positions of ν -names correctly forbid communication between branches of opposite polarities within the same choice (no ‘incestuous’ communication, according to the terminology of [8]). The encoding acts as a homomorphism over the remaining operators of Π^{mix} .

Below, \sim^S indicates sorted open bisimilarity in Π^{mix} . We denote by **S2** the D-Fusion sorting induced by the above translation of the source sorting (i.e., if **S**: $S \mapsto \tilde{x} : \tilde{S}'$, then **S2**: $S \mapsto (\tilde{x} : \tilde{S}', y : S_0, y : S_0, z : S_0, z : S_0)$, where S_0 is a new sort and y and z are fresh). The corresponding bisimilarity is written \sim^{S2} .

Theorem 3 (full abstraction for mixed choice). *Let $P, Q \in \Pi^{\text{mix}}$. It holds that $P \sim^S Q$ if and only if $\llbracket P \rrbracket_{\text{mix}} \sim^{S2} \llbracket Q \rrbracket_{\text{mix}}$.*

In a pi-calculus setting, it is known that mixed choice cannot be encoded into the choice-free fragment, if one requires the encoding be uniform and preserve a ‘reasonable semantics’ [9, 10, 8]. The theorem above shows that, under mild typing conditions, pi-calculus mixed choice *can* be implemented into the choice-free fragment of D-Fusion. The encoding is uniform, deadlock- and divergence-free, and preserves a ‘reasonable semantics’. This is yet another evidence of the expressiveness gap between D-Fusion and pi-calculus.

8 Conclusions and Future Work

We have proposed the D-Fusion calculus, an extension of the fusion calculus where two distinct binders coexist, one analogous to the (x) binder in fusion, the other imposing

name freshness. We have shown that D-Fusion is more expressive than both Fusion and pi-calculus.

Our expressiveness results seem to suggest that an efficient distributed implementation of D-Fusion might be nontrivial to design. This design would certainly involve the introduction of a distributed model of the calculus, including, e.g., explicit fusions [3] for broadcasting fusions asynchronously, and primitives for handling explicit localities. We leave this task for future work. For the time being, we just note that distributed implementations of pi/fusion-like calculi do exist (e.g., the fusion machine of [2]). One may expect that, starting from these, building a distributed implementation of D-Fusion should not be much harder.

In [1] the synchronisation mechanism of the pi-calculus is extended to allow for polyadic synchronisation, where channels are vectors of names. The expressiveness of polyadic synchronisation, matching and mixed choice is compared and it is shown how the degree of synchronisation of a calculus increases its expressive power.

We plan to extend the D-Fusion calculus by generalising name fusions to arbitrary substitutions over a signature of terms. We believe that the extended D-Fusion would be strictly more expressive than Logic Programming, the intuition being that restriction (creation of new fresh names) cannot be modelled in LP.

Finally, it would also be interesting to study a coalgebraic model for D-Fusion. A coalgebraic framework would present several advantages. For instance, morphisms between coalgebras enjoy the property of “reflecting behaviours” and thus they allow to characterise bisimulation equivalences in more abstract terms as kernels of morphisms.

References

1. M. Carbone and S. Maffei. On the Expressive Power of Polyadic Synchronisation in Pi-Calculus. To appear in *Nordic Journal of Computing*.
2. P. Gardner, C. Laneve, and L. Wischik. The fusion machine (extended abstract). In *Proc. of CONCUR '02*, LNCS 2421. Springer-Verlag, 2002.
3. P. Gardner and L. Wischik. Explicit Fusions. *Theoretical Computer Science*. To appear.
4. L. G. Meredith, S. Bjorg, and D. Richter. Highwire Language Specification Version 1.0. Unpublished manuscript.
5. Microsoft Corp. Biztalk Server - <http://www.microsoft.com/biztalk>.
6. R. Milner. The Polyadic pi-Calculus: a Tutorial. Technical Report, Computer Science Dept., University of Edinburgh, 1991.
7. R. Milner, J. Parrow, and D. Walker. A calculus of mobile processes (parts I and II). *Information and Computation*, 100(1):1–77, 1992.
8. U. Nestmann and B. C. Pierce. Decoding choice encodings. *Information and Computation*, 163(1):1–59, 2000.
9. C. Palamidessi. Comparing the Expressive Power of the Synchronous and the Asynchronous pi-calculus. In *Conf. Rec. of POPL'97*, 1997.
10. C. Palamidessi. Comparing the Expressive Power of the Synchronous and the Asynchronous pi-calculus. *Mathematical Structures in Computer Science*, 13(5):685–719, 2003.
11. J. Parrow and B. Victor. The Fusion Calculus: Expressiveness and Symmetry in Mobile Processes. In *Proc. of LICS'98*. IEEE Computer Society Press, 1998.
12. D. Sangiorgi. A Theory of Bisimulation for the pi-Calculus. *Acta Informatica*, 33(1): 69–97, 1996.
13. World Wide Web Consortium (W3C) - <http://www.w3.org/TR/wsd112>.