

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. We study the relative expressive power of Fusion and pi-calculus. Fusion 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

The design of distributed applications based on XML, like Web services [22] or business-to-business systems [6], sees the emergence of a message-passing programming style. Languages like Highwire [5] provide, in a concurrent setting, sophisticated data structures, which allow programmers to describe and manipulate complex messages and interaction patterns. If one looks for ‘foundational’ counterparts of these programming languages the pi-calculus [9, 10] and the Fusion calculus [16] 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 [3], or forms of pattern matching between pairs of messages [5]. 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, fusion name-passing enables a more general binding mechanism. However, differently from the pi-calculus, the binding mechanism of Fusion ignores the issue of unicity of newly generated names. One of the goals of this paper is to show that this fact limits the expressiveness of Fusion. Overcoming

^{*} Research partially supported by FET-Global Computing projects *PROFUNDIS* and *MIKADO*.

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 imposing unicity, like in pi-calculus. 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. However, the main motivation of the present study is not to push for adoption of yet another calculus, but rather to clarify the relationship between two fundamental concepts like binding and fusions, and propose a consistent and simple semantical framework to reason on them. A more precise account of our work follows.

The binding mechanism of pi-calculus generalises that of λ -calculus in several ways. Input prefix $a(x)$. binds like λx , and name passing takes place in pi-calculus in a way typical of functional programming, i.e., formal names are assigned their actual counterparts. 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.

Fusion calculus is presented in [16] as a more uniform and more expressive evolution of the pi-calculus. The main idea is to decompose input prefix $a(x)$. into a binder (x) and a prefix $a(x)$. In the polyadic case, matching between the input list and the output list of arguments 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, the pi-calculus transition system can be embedded into Fusion's, provided that one identifies restriction (νx) with the (x) binder of Fusion [16].

Our first move is to argue that this embedding breaks down if comparing the two calculi on the basis of behavioural 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, mapping (νx) to (x) and preserving (a subset of) weak traces. As hinted before, the ultimate reason for this failure is that in Fusion all names are like logical variables, i.e., unification always succeeds, which is not true in the pi-calculus.

The above considerations motivate the introduction of a new calculus, D-Fusion, with two binders, λ and ν : the first generalises input prefix, and the second corresponds to restriction. Also, any issue of symmetry between input and output is preempted, since the calculus has just one kind of prefix (no polarisation); polarised prefixes can be easily encoded, though. In D-Fusion, while lambdas are used to control the propagation of fusions, restrictions are used to possibly inhibit fusions. In logical terms, this corresponds to consider unification not only among variables, but also among variables and dynamically generated constants (that is, ν -extruded names). As expected, unification fails whenever one tries to identify two distinct constants. We show that the additional expressive power achieved in this way is relevant. Both pi-calculus and Fusion are subcalculi of D-Fusion. Moreover, the combined mechanism of restriction and unification yields additional expressive power: it allows to express a form of pattern 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 [13, 14] that nondeterministic *guarded choice* cannot be simulated in the choice (+) -free pi-calculus in a fully abstract way, while preserving any ‘reasonable’ semantics. The reason is that it is not possible to atomically perform, 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, defined in terms of barbed equivalence ([11]). Informally, branches of a choice are represented as concurrent processes. Synchronisation is performed in the ordinary way, but it forces a fusion between a λ -name global to all branches 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.

In the present paper, we are mainly interested in assessing the expressive power of D-Fusion compared to other calculi. The principal tool for this study will be barbed bisimilarity and the induced equivalences, as they enjoy a uniform definition based only on a reduction relation, on an observation predicate and on context-closure. We defer the study of alternative, more tractable semantics for D-Fusion, like a form of ‘labelled’ bisimulation, to a forthcoming work.

The rest of the paper is organised as follows. Section 2 contains a proof that the pi-calculus cannot be uniformly encoded into Fusion. In Section 3 we introduce the D-Fusion calculus, its operational semantics and barbed congruence. In Section 4 we show that D-Fusion calculus is strictly more expressive than both pi-calculus and Fusion. We further explore this expressiveness gap in Section 5, by means of an example concerning a security protocol, and in Section 6, by encoding mixed guarded choice into the choice-free calculus. Section 7 contains a brief overview of some related work and a few concluding remarks.

2 Fusion and Pi

The aim of this section is to illustrate 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:

$$R | \bar{u}v.P | ux.Q \xrightarrow{\{x=\nu\}} R | P | Q.$$

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

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

where τ denotes the identity fusion. For a full treatment of pi-calculus and Fusion we refer to [10] and to [16], 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 [16] does satisfy our criterion. We shall assume here the standard pi-calculus late operational semantics [10] 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 , $\llbracket P\sigma \rrbracket = \llbracket P \rrbracket\sigma$.

Note that the above notion of uniform encoding is stronger than the one introduced in [14].

The next proposition generalises an example from [16]. 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 [16]).

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:

$$P = (\nu u, v)(\bar{a}(u, v)|\bar{u}|v.\bar{w}) \quad Q = (\nu u, v)(\bar{a}(u, v)|(\bar{u}.(v.\bar{w}) + v.(\bar{u}\bar{w}))).$$

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:

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

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:

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

where *prefixes* α are defined as $\alpha ::= a\tilde{v}$. The occurrences of x in λxP 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. A *context* $C[\cdot]$ is a process with a hole that can be filled with any process P , thus yielding a process $C[P]$.

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.

Definition 2 (Structural Congruence). *The structural congruence \equiv is the least congruence on processes satisfying the abelian monoid laws for Summation and Composition (associativity, commutativity and $\mathbf{0}$ as identity), the scope laws*

$$\begin{array}{lll} (\nu x)\mathbf{0} \equiv \mathbf{0} & (\nu x)(\nu y)P \equiv (\nu y)(\nu x)P & (\nu x)(P+Q) \equiv (\nu x)P + (\nu x)Q \\ \lambda x\mathbf{0} \equiv \mathbf{0} & \lambda x\lambda yP \equiv \lambda y\lambda xP & \lambda x(P+Q) \equiv \lambda xP + \lambda xQ \end{array}$$

plus the scope extrusion laws

$$(\nu x)(P|Q) \equiv P|(\nu x)Q \text{ where } x \notin \text{fn}(P) \quad \lambda x(P|Q) \equiv P|\lambda xQ \text{ where } x \notin \text{fn}(P)$$

and the swapping law

$$\lambda x(\nu y)P \equiv (\nu y)\lambda xP \text{ where } x \neq y.$$

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 $\sigma - (\tilde{x} \times \mathcal{N})$.

Below, we define fusions, that is, name equivalences that arise as the result of equating two lists of names in a synchronisation.

Definition 3 (Fusions). *We let ϕ, χ, \dots range over fusions, that is total equivalence relations on \mathcal{N} with only finitely many non-singleton equivalence classes. We let:*

- $n(\phi)$ denote $\{x : x\phi \text{ for some } y \neq x\}$;
- τ denote the identity fusion (i.e., $n(\tau) = \emptyset$);
- 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 .

We now introduce a labelled transition system for D-Fusion. This will be useful in order to compare D-Fusion with other calculi and, in particular, to prove that D-Fusion cannot be encoded into pi-calculus (see Section 4). The reduction relation coincides with the identity fusion transition $\xrightarrow{\tau}$ of the labelled transition system.

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; when $\mu = (v\tilde{x})\lambda\tilde{y}a\tilde{v}$ is an action we let $\text{subj}(\mu) = a$ and $\text{obj}(\mu) = \tilde{v}$ denote the subject and object part of μ , otherwise they both denote a conventional value ‘-’. Moreover, $n(\mu)$ denotes all names in μ . We use abbreviations such as $n(\phi, \mu)$ to denote $n(\phi) \cup n(\mu)$ 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: an effect of the form ϕ 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 $\lambda\text{-OPEN}_f$). At that point, a substitutive effect $[w/z]$ is applied onto the target process and z is discarded from the fusion (the result is ϕ_{-z}). A major difference with respect to Fusion is that our effects can also v -extrude names. The side condition ‘ $\phi[z] \cap \tilde{x} = \emptyset$ ’ in rule $v\text{-OPEN}$ prevents effects from equating two distinct v -extruded names. Note that the premises of rules COM and $\lambda\text{-OPEN}$ only deal with binder-free actions and effects, while $\lambda\text{-OPEN}_a$ only deals with v -binder-free actions. However, structural congruence permits freely moving λ ’s and v ’s, so the general case is covered via STRUCT .

Let us illustrate the rules with some examples. We shall write $\{\tilde{x} = \tilde{y}\}.P$ for $(vc)(c\tilde{x}|c\tilde{y}.P)$ (for a fresh name c).

Example 1.

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

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

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

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

Table 1. Actions and effects transitions in D-Fusion

$\text{(PREF)} \quad \alpha.P \xrightarrow{\alpha} P$	$\text{(SUM)} \quad \frac{P_1 \xrightarrow{\mu} Q}{P_1 + P_2 \xrightarrow{\mu} Q}$
$\text{(COM)} \quad \frac{P_1 \xrightarrow{a\tilde{u}} Q_1 \quad P_2 \xrightarrow{a\tilde{v}} Q_2}{P_1 P_2 \xrightarrow{\{\tilde{v}=\tilde{u}\}} Q_1 Q_2} \quad \tilde{u} = \tilde{v} $	$\text{(PAR)} \quad \frac{P \xrightarrow{\mu} Q}{P R \xrightarrow{\mu} Q R}$
$\text{(v-PASS)} \quad \frac{P \xrightarrow{\mu} Q}{(vz)P \xrightarrow{\mu} (vz)Q} \quad z \notin n(\mu)$	$\text{(\lambda-PASS)} \quad \frac{P \xrightarrow{\mu} Q}{\lambda z P \xrightarrow{\mu} \lambda z Q} \quad z \notin n(\mu)$
$\text{(v-OPEN)} \quad \frac{P \xrightarrow{\mu} Q}{(vz)P \xrightarrow{(vz)\mu} Q} \quad \left\{ \begin{array}{l} z \in n(\mu) \\ \mu \text{ an action implies } z \neq \text{subj}(\mu) \\ \mu = (v\tilde{x})\phi \text{ implies } \phi[z] \cap \tilde{x} = \emptyset \end{array} \right.$	
$\text{(\lambda-OPEN}_a\text{)} \quad \frac{P \xrightarrow{\lambda\tilde{y}a\tilde{v}} Q}{\lambda z P \xrightarrow{\lambda\tilde{y}z a\tilde{v}} Q} \quad z \in \tilde{v} - (\{a\} \cup \tilde{y})$	$\text{(\lambda-OPEN}_f\text{)} \quad \frac{P \xrightarrow{\phi} Q}{\lambda z P \xrightarrow{\phi \cdot z} Q[w/z]} \quad z \phi w, w \neq z$
$\text{(MATCH)} \quad \frac{P \xrightarrow{\mu} Q}{[a = a]P \xrightarrow{\mu} Q}$	$\text{(STRUCT)} \quad \frac{P_1 \equiv P \quad P \xrightarrow{\mu} Q \quad Q \equiv Q_1}{P_1 \xrightarrow{\mu} Q_1}$

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

2. Another example illustrating interplay between v-bound, λ-bound and free names:

$$\lambda z (vw) \{z, a = w, b\}.P \xrightarrow{\{a=b\}} (vw)P[w/z], \quad \text{but} \quad \lambda z (vw) \{z, z = w, b\}.P \xrightarrow{(vw)\{w=b\}} P[w/z].$$

Encoding I/O Polarities. We can encode input and output polarities as follows:

$$\bar{c}(\tilde{v}).P \triangleq (vx)\lambda y c\tilde{v}xy.P \quad c(\tilde{v}).P \triangleq (vx)\lambda y c\tilde{v}yx.P$$

for some chosen fresh x and y . The position of the v-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 v-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.

Barbed Congruence. We now define our main tools for assessing the expressive power of D-Fusion compared to other calculi, that is barbed bisimulation and barbed congruence.

Definition 5 (Barbs). We write $P \downarrow a$ if and only if there exist an action μ and a process Q such that $P \xrightarrow{\mu} Q$ and $\text{subj}(\mu) = a$.

Definition 6 (Barbed Bisimulation). A barbed bisimulation is a symmetric binary relation \mathcal{R} between processes such that $P \mathcal{R} Q$ implies:

1. whenever $P \xrightarrow{\tau} P'$ then $Q \xrightarrow{\tau} Q'$ and $P' \mathcal{R} Q'$;
2. for each name a , if $P \downarrow a$ then $Q \downarrow a$.

P is a barbed bisimilar to Q , written $P \sim Q$, if $P \mathcal{R} Q$ for some barbed bisimulation \mathcal{R} .

Definition 7 (Barbed Congruence). Two processes P and Q are barbed congruent, written $P \sim Q$, if for all contexts $C[\cdot]$, it holds that $C[P] \sim C[Q]$.

Example 2. 1. An example of ‘expansion’ for parallel composition is as follows:

$$(\nu k) ak.\mathbf{0} | av.\mathbf{0} \sim (\nu k) ak.av.\mathbf{0} + av.(\nu k) ak.\mathbf{0} + (\nu k) \{k = v\}.\mathbf{0}$$

On the other side,

$$\begin{aligned} ((\nu k) ak.ak.\mathbf{0}) | av.\mathbf{0} \not\sim & (\nu k) ak.(ak.av.\mathbf{0} + av.ak.\mathbf{0}) + av.((\nu k) ak.ak.\mathbf{0}) \\ & + (\nu k) \{k = v\}.ak.\mathbf{0}, \end{aligned}$$

since the two processes are not barbed bisimilar within a context $C[\cdot] = (\lambda x, v)([\cdot] | ax.\mathbf{0})$.

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:

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

On the contrary,

$$(\nu v) \{k = v\}.P \not\sim \tau.P[k/v],$$

since the two processes are not barbed bisimilar within a context $C[\cdot] = (\nu k)[\cdot]$.

4 Expressiveness of D-Fusion

Pi-calculus and Fusion are subcalculi of D-Fusion, because their respective labelled transition systems are embedded into polarised D-Fusion’s, under the two obvious uniform translations from Π and \mathcal{F} to \mathcal{DF}^P given below. The definition of uniformity can be extended to the case of encodings from \mathcal{F}/Π into \mathcal{DF}^P in the obvious way: in particular, by requiring that (x) and (νx) be mapped to λx and (νx) , respectively.

Definition 8. The translations $[[\cdot]]_\pi : \Pi \rightarrow \mathcal{DF}^P$ and $[[\cdot]]_f : \mathcal{F} \rightarrow \mathcal{DF}^P$ are defined by extending in the expected homomorphic way the following clauses, respectively:

$$\begin{aligned} [[\bar{a}\langle x \rangle.P]]_\pi &= \bar{a}\langle x \rangle.[[P]]_\pi & [[a\langle x \rangle.P]]_\pi &= \lambda x a\langle x \rangle.[[P]]_\pi & [[(vx)P]]_\pi &= (vx) [[P]]_\pi \\ [[\bar{a}\langle x \rangle.P]]_f &= \bar{a}\langle x \rangle.[[P]]_f & [[a\langle x \rangle.P]]_f &= a\langle x \rangle.[[P]]_f & [[(x)P]]_f &= \lambda x [[P]]_f \end{aligned}$$

As expected, inclusion in term of labelled transition systems naturally lifts to bisimulation equivalences. Let \sim^π and \sim^f denote barbed congruence, respectively, over Π ([18]) and over \mathcal{F} (see [21]) (Note that, over image-finite processes, \sim^π is pi-calculus early congruence [18], and \sim^f is Fusion hyper-equivalence [21]). Also, let $\sim^{[\pi]}$ and $\sim^{[f]}$ be the equivalences on \mathcal{DF}^P obtained by closing barbed bisimulation \sim only under translated pi- and Fusion-contexts, respectively (e.g., $P \sim^{[\pi]} Q$ iff for each Π -context $C[\cdot]$, $[[C]]_\pi[P] \sim [[C]]_\pi[Q]$).

Proposition 2.

1. Let P and Q be two pi-calculus processes. $P \sim^\pi Q$ iff $[[P]]_\pi \sim^{[\pi]} [[Q]]_\pi$.
2. Let P and Q be two Fusion processes. $P \sim^f Q$ iff $[[P]]_f \sim^{[f]} [[Q]]_f$.

More 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 action prefix, fusions and restrictions allows one to express a form of pattern matching. This is not possible in Π , at least not atomically. To show this fact, we restrict our attention to polarised D-Fusion, \mathcal{DF}^P . The reference semantics for Π is again the late operational semantics.

Given $P \in \Pi$ and a trace of actions s , let us write $P \xrightarrow{s}^{\hat{s}}$ if $P \xrightarrow{s'}$ for some trace s' that exhibits the same sequence of subject names, with the same polarity, as s (e.g., $s = a\langle \bar{x} \rangle \cdot \lambda \bar{y} \bar{b}\langle \bar{v} \rangle$ and $s' = a\langle \bar{z} \rangle \cdot \bar{b}\langle \bar{w} \rangle$).

Definition 9. A translation $[[\cdot]] : \mathcal{DF}^P \rightarrow \Pi$ is uniform if for each $P, Q \in \mathcal{DF}^P$:

- for each trace s , $P \xrightarrow{s}$ implies $[[P]] \xrightarrow{\hat{s}}$;
- $[[P|Q]] = [[P]] | [[Q]]$;
- for each y , $[[(vy)P]] = (vy) [[P]]$;
- for each substitution σ , $[[P\sigma]] = [[P]]\sigma$.

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. Note that barbed congruence over \mathcal{DF}^P , \sim , is contained in trace equivalence.

Proposition 3. There is no uniform translation $[[\cdot]] : \mathcal{DF}^P \rightarrow \Pi$ such that $\forall P, Q \in \mathcal{DF}^P$:

$$P \sim_{\mathcal{DF}^P} Q \Rightarrow [[P]] \sim_\Pi [[Q]].$$

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

$$P = (vc, k, h) (c\langle k \rangle \cdot \bar{a} \cdot \mathbf{0} | c\langle h \rangle \cdot \bar{b} \cdot \mathbf{0} | \bar{c}\langle k \rangle \cdot \mathbf{0}) \quad Q = \tau \cdot \bar{a} \cdot \mathbf{0}.$$

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 v-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, by Definition 9, for any uniform encoding $[[\cdot]]$, c and \bar{c} in $[[P]]$ can synchronise and, thus, $[[P]] \xrightarrow{\bar{b}}$, while $[[Q]] \not\xrightarrow{\bar{b}}$ (because of $b \notin \text{fn}(Q)$ and of the uniformity with respect to substitutions). Thus $[[P]] \not\sim_{\Pi} [[Q]]$. \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.

5 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 avoid confusion of distinct sessions. 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 $\lambda \tilde{x} p(\tilde{x}).X$ as $p(\tilde{x}).X$):

$$\begin{aligned}
A &= (\nu n) \left(\overline{p}\langle n \rangle . \mathbf{0} \mid \lambda y k_S \langle n, y \rangle . (\nu n') \left(\overline{p}\langle n' \rangle . \mathbf{0} \mid \lambda y' k_S \langle n', y' \rangle . \overline{\text{commit}}\langle y, y' \rangle . \mathbf{0} \right) \right) \\
S &= p(x) . \left(\overline{k_S}\langle x, k \rangle . \mathbf{0} \mid p(x') . \overline{k_S}\langle x', k' \rangle . \mathbf{0} \right) \\
P &= (\nu k_S) (A \mid S).
\end{aligned}$$

Let A_{spec} be the process defined like A , except that the $\overline{\text{commit}}\langle y, y' \rangle$ action is replaced by $\overline{\text{commit}}\langle k, k' \rangle$, and let $P_{\text{spec}} = (\nu k_S) (A_{\text{spec}} \mid S)$. 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 S , that is $\overline{k_S}\langle x', k' \rangle$. In fact, n' can be extruded only *after* x has been received, hence fusion of x and n' is forbidden. Note that the above protocol specification would not be easily translated in pi-calculus, because in A the input prefix $k_S \langle n, y \rangle$ has a ν -bound name n .

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, $R = p(x) . (\overline{p}\langle x \rangle . \mathbf{0} \mid p(y) . \overline{p}\langle y \rangle . \mathbf{0})$. Formally, for $\alpha = \overline{\text{commit}}\langle k', k \rangle$,

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

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.

6 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 add 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 3. 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:

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

(here, $x, n, m \notin \text{fn}(a, P, Q)$). Assume a channel discipline by which output actions on channel a must carry two identical names. 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.:

$$\lambda u \overline{a}\langle uu \rangle \mid B \xrightarrow{\tau} \sim (\nu n) (P \mid (\nu m) a \langle mn \rangle . Q) \quad \sim \quad P \mid (\nu n, m) a \langle mn \rangle . Q.$$

Under the mentioned assumption, $(\nu m, n) a \langle mn \rangle . Q$ 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 a new ‘disciplined’ equivalence \sim^d , obtained by closing barbed bisimilarity only with respect to contexts $C[\cdot]$ obeying the mentioned channel discipline (i.e. for each $\bar{c}\langle\tilde{v}\rangle$ in $C[\cdot]$ with $|\tilde{v}| \geq 2$, $\tilde{v} = \tilde{w}uu$, for some \tilde{w} and u). The other way is keeping standard barbed congruence \sim , but inserting processes inside a ‘firewall’ that filters out \bar{a} -messages not respecting the given channel discipline. The latter can be easily defined in D-Fusion relying on ‘non-linear’ inputs:

$$F_{a,a'}[\cdot] = (\nu a') \left(\lambda z a z z . \bar{a}' \langle z z \rangle . \mathbf{0} | [\cdot] \right).$$

We state the result in both forms below.

Proposition 4. *Let A and B be as in Example 3.*

1. $A \sim^d B$;
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 4 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 ‘disciplined’ equivalences. We believe the results can also be stated in terms of ordinary barbed congruence, 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, $\mathcal{DF}^{p,ig}$. In this language, input prefix and summation $+$ are replaced by input-guarded choice $\sum_{i \in I} a_i \langle \tilde{x}_i \rangle . P_i$. The target language is the fragment of polarised D-Fusion with no form of summation. The relevant clauses of the encoding are:

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

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, we denote by $\sim^{p,ig}$ barbed congruence over $\mathcal{DF}^{p,ig}$, and denote by $\sim^{\llbracket p,ig \rrbracket}$ the equivalence over D-Fusion obtained by closing barbed bisimulation under translated contexts (i.e. $P \sim^{\llbracket p,ig \rrbracket} Q$ iff for each $\mathcal{DF}^{p,ig}$ -context $C[\cdot]$, it holds $\llbracket C \rrbracket_{ig}[P] \sim \llbracket C \rrbracket_{ig}[Q]$); note that both equivalences are *reasonable* semantics in the sense of [13]. The proof of the following theorem is straightforward, given that there is a 1-to-1 correspondence between reductions and barbs of R and of $\llbracket R \rrbracket_{ig}$, for any R , and given that the encoding is compositional, in particular, for any context $C[\cdot]$, it holds $\llbracket C \rrbracket_{ig}[\llbracket P \rrbracket_{ig}] = \llbracket C[P] \rrbracket_{ig}$.

Theorem 1 (Full Abstraction for ig Choice). *Let $P, Q \in \mathcal{DF}^{p,ig}$. It holds that $P \sim^{p,ig} Q$ if and only if $\llbracket P \rrbracket_{ig} \sim^{\llbracket p,ig \rrbracket} \llbracket Q \rrbracket_{ig}$.*

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 $\mathcal{DF}^{p,ig}$.

Mixed Choice in a Sorted pi-Calculus. As a source language we fix here a sorted version of polyadic pi-calculus [9] 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:

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

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 [12]). The encoding acts as a homomorphism over the remaining operators of Π^{mix} .

Below, \sim^{mix} denotes barbed congruence over Π^{mix} , and $\sim^{\llbracket \text{mix} \rrbracket}$ the equivalence over D-Fusion obtained by closing barbed bisimulation under translated Π^{mix} -contexts. Both equivalences are reasonable semantics in the sense of [13]. The proof of the following theorem is again straightforward by correspondence on reductions and barbs, and by compositionality of the encoding.

Theorem 2 (Full Abstraction for Mixed Choice). *Let $P, Q \in \Pi^{\text{mix}}$. It holds that $P \sim^{\text{mix}} Q$ if and only if $\llbracket P \rrbracket_{\text{mix}} \sim^{\llbracket \text{mix} \rrbracket} \llbracket Q \rrbracket_{\text{mix}}$.*

In a pi-calculus setting, it is well-known that mixed choice cannot be encoded into the choice-free fragment, if one requires the encoding be uniform and preserve a reasonable semantics [13, 14, 12]. The theorem above shows that 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.

7 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 strictly more expressive than both Fusion and pi-calculus.

Our expressiveness results seem to suggest that the design of an efficient distributed implementation of D-Fusion might be nontrivial. This design would probably involve the introduction of a distributed model of the calculus, including, e.g., explicit fusions [4] for broadcasting fusions asynchronously and rollback mechanisms for handling fusion failures. 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 [3]) and may represent a good starting point for distributed implementations of D-Fusion.

Another point that deserves further study is characterization of D-Fusion barbed congruence in terms of a more tractable, labelled bisimulation, which would avoid universal quantification on all contexts. Preliminary results indicate that definition of this

equivalence would require a (nontrivial) integration of *substitutive effects* à la fusion calculus [16], i.e. name substitutions resulting from fusions, with *distinctions* à la open pi-calculus [19].

We also plan to extend the D-Fusion calculus by generalising name fusions to substitutions over an arbitrary signature of terms. It would be interesting to compare the expressive power of this extended D-Fusion to systems of Concurrent Constraint or Logic Programming that allow creation of fresh names, such as lambda-Prolog [8], and CCP [17, 20].

In [7] Merro gives an encoding from asynchronous Chi calculus (a variant of Fusion, independently introduced by Fu, [2]) to (asynchronous) pi-calculus. However, no result on the other direction is proven. Here, we have proved that pi-calculus cannot be encoded into Fusion.

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.

Acknowledgments. The authors wish to thank Bjorn Victor for stimulating discussions. The anonymous referees provided valuable suggestions.

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. Y. Fu. A Proof Theoretical Approach to Communication. In *Proc. of ICALP '97*, LNCS 1256. Springer-Verlag, 1997.
3. P. Gardner, C. Laneve, and L. Wischik. The fusion machine (extended abstract). In *Proc. of CONCUR '02*, LNCS 2421. Springer-Verlag, 2002.
4. P. Gardner and L. Wischik. Explicit Fusions. *Theoretical Computer Science*. To appear.
5. L. G. Meredith, S. Bjorg, and D. Richter. Highwire Language Specification Version 1.0. Unpublished manuscript.
6. Microsoft Corp. Biztalk Server - <http://www.microsoft.com/biztalk>.
7. M. Merro. On the Expressiveness of Chi, Update, and Fusion calculi. In *Proc. of EXPRESS '98*, ENTCS 16(2), Elsevier Science, 1998.
8. D. Miller. Unification under a mixed prefix. *Journal of Symbolic Computation*, 14(4):321–358, 1992.
9. R. Milner. The Polyadic pi-Calculus: a Tutorial. Technical Report, Computer Science Dept., University of Edinburgh, 1991.
10. R. Milner, J. Parrow, and D. Walker. A calculus of mobile processes (parts I and II). *Information and Computation*, 100(1):1–77, 1992.
11. R. Milner and D. Sangiorgi. Barbed Bisimulation. In *Proc. of ICALP '92*, LNCS 623, Springer-Verlag, 1992.
12. U. Nestmann and B. C. Pierce. Decoding choice encodings. *Information and Computation*, 163(1):1–59, 2000.
13. C. Palamidessi. Comparing the Expressive Power of the Synchronous and the Asynchronous pi-calculus. In *Conf. Rec. of POPL'97*, 1997.
14. C. Palamidessi. Comparing the Expressive Power of the Synchronous and the Asynchronous pi-calculus. *Mathematical Structures in Computer Science*, 13(5):685–719, 2003.

15. J. Parrow and B. Victor. The Update Calculus. In *Proc. of AMAST'97*, LNCS 1349, Springer-Verlag, 1997.
16. J. Parrow and B. Victor. The Fusion Calculus: Expressiveness and Symmetry in Mobile Processes. In *Proc. of LICS'98*. IEEE Computer Society Press, 1998.
17. E. Shapiro. The Family of Concurrent Logic Programming Languages. *ACM Computing Surveys*, 21(3):413-510, 1989.
18. D. Sangiorgi. Expressing Mobility in Process Algebras: First-Order and Higher-Order Paradigms. PhD thesis, Department of Computer Science, University of Edinburgh, 1992.
19. D. Sangiorgi. A Theory of Bisimulation for the pi-Calculus. *Acta Informatica*, 33(1): 69-97, 1996.
20. V. Saraswat. Concurrent Constraint Programming. The MIT Press, 1993.
21. B. Victor. The Fusion Calculus: Expressiveness and Symmetry in Mobile Processes. PhD thesis, Department of Computer Systems, Uppsala University, 1998.
22. World Wide Web Consortium (W3C) - <http://www.w3.org/TR/wsd112>.