



Laboratoire IBISC

Biologie Intégrative et Systèmes Complexes

Pigraphs with Replicators: Finiteness and Boundedness Results

Frédéric Peschanski, Hanna Klaudel, and Raymond Devillers

IBISC University of Evry-Val d'Essonne, Genopole, France



RAPPORT DE RECHERCHE

IBISC-RR-2011-04

mai 2011

Pigraphs with Replicators

Finiteness and Boundedness Results

(Technical Report)

Frédéric Peschanski¹, Hanna Klaudel², and Raymond Devillers³

¹ UPMC – LIP6

² Université d'Évry–IBISC

³ Université Libre de Bruxelles

Abstract. This paper describes the pi-graphs, a diagrammatic variant of the pi-calculus. The semantics is based on graph relabelling techniques, giving a much more concrete characterization compared to standard characterizations. This allows us to derive very precise results, for example the prediction of upper bounds for the dynamic resources produced and consumed by (thread-bounded) pi-graphs.

1 Introduction

This paper describes the *pi-graph calculus*, a variant of the pi-calculus [1, 2]. The most important feature of the formalism is that each construct possesses a simple, if not natural, diagrammatic interpretation. This emphasizes our first goal: the conception of (visual *and* textual) *modelling* artefacts for mobile systems.

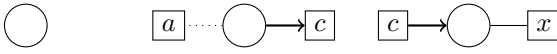
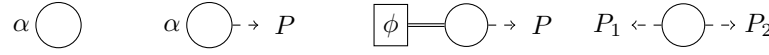
Unlike previous pi-graphs variants [3–5], the calculus presented in this paper is quite a standard pi-calculus. Thus, it can be semantically characterized effortlessly by reusing one of the various known semantics (early, late, symbolic, etc.) and behavioral equivalences (barbed, open, etc.) for the pi-calculus. However, beyond modelling, our second goal is to develop efficient *implementation* and *verification* techniques for mobile systems. This lead us to develop more *concrete* semantics for the calculus as a *token-game* inspired by the Petri nets.

Technically we build on *graph relabelling* techniques which (unlike most graphical encodings of the pi-calculus, e.g. [6]) disallow node or edge creation or destruction in graph rewrites: the graph structure is invariant. This has practical advantages : the semantics is almost directly implementable, and structural invariance reflects a major trait of most popular modelling languages (e.g. finite state machines, statecharts, Petri nets, etc.). Most importantly, the essential notion of *thread* naturally emerges from the characterization. Although reasoning is less easy than in more abstract settings, we are able to derive quite precise results. For example, we exhibit upper bounds for the dynamic resources produced and consumed by systems in which the maximum number of simultaneously running threads is bounded, so-called *thread-bounded* pi-graphs. As a consequence, the “natural” behavioral equivalence on pi-graphs - namely *stateful bisimilarity* - is decidable for thread-bounded systems.

The outline of the paper is as follows. In Section 2 we introduce the diagram calculus and its informal semantics. The operational semantics, in terms of graph relabelling, is detailed in Section 4. The finiteness and boundedness results are presented in Section 5. The related and future works are discussed in Section 6.

2 The diagram calculus

The pi-graph calculus is a relatively standard variant of the π -calculus, and *also* a language of *diagrams*. The syntax of the calculus, as well as the diagrammatic interpretation of each construct, is given below.

Action $\alpha ::=$	Guard $\phi ::= [a = b] \mid [a \neq b]$
	
τ (silent)	$\bar{c}\langle a \rangle$ (output)
	$c(x)$ (input)
Process $P ::=$	
	
α (suffix)	$\alpha.P$ (prefixing)
	ϕP (guard)
	$P_1 + P_2$ (choice)
Replicator $R ::= r[k] : \widetilde{\nu a} !\alpha.P$	
	Diagram $\pi ::= \widetilde{\nu A} R_1 \dots R_n$

In the graphical interpretation, each action corresponds to a node similar to a place in the Petri nets terminology, and indeed will contain *tokens*. The rectangular-shaped nodes are holders for *names* that can be used as channels. Of course, names can be transmitted dynamically as in the pi-calculus. The structure of the control-flow is encoded by dashed arrows. For instance, the choice operator involves two possible continuations. Unlike previous models of pi-graphs [3–5], we reuse the familiar pi-calculus concept of *replication* to model repetitive (and potentially infinite) behaviors. A difference with traditional replication is that in pi-graphs, the so-called *replicators* are not processes *per se* but can only be defined at the top-level. At the cost of some extra occurrences of tau transitions, nested replications can be encoded as top-level ones in the pi-calculus⁴. Relaxing this constraint would lead to semantic quirks (namely ϵ -divergence), and intuitively the reason is that replicated processes are spawned in a demand-driven way. Each replicator - identified as r - is associated to a bound k that limits the number of running *threads*. Control-infinite systems can be modelled with bound $k = \omega$. Note also that the replicated process must be prefixed, which also avoids semantic difficulties. If in the pi-calculus one can “play” with the scope of restrictions using *scope-extrusion* laws, the syntax is here constrained so that restrictions can only occur at the global level, or alternatively within a given replicator. The former are global restrictions, whereas the latter are thread-specific (although they might be shared). Note also the

⁴ For a pi-calculus context C and a process P , $C[!P]$ can be encoded as e.g. $\nu c(c(x).!P \mid C[\bar{c}\langle c \rangle.0])$ provided c is fresh in $C[P]$ and x is not free in P .

absence of a 0 terminator (thus any action can be a suffix), which is simply not required (and indeed undesirable because tokens may accumulate indefinitely in the terminal nodes).

To illustrate the language, its diagrammatic interpretation and its informal semantics, we consider the following example.

Example 1. A (k -bounded) multi-session server (with initial redexes)

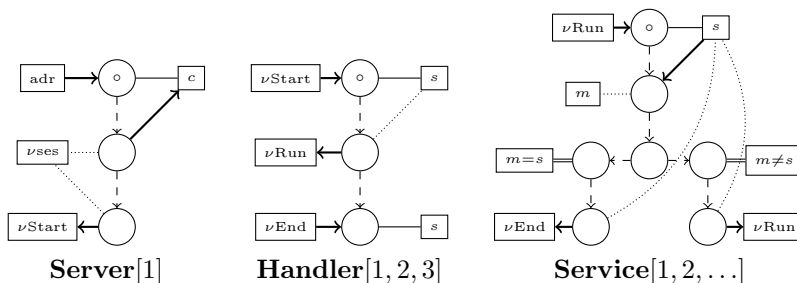
$$\nu(\text{Start}, \text{Run}, \text{End})$$

$$\text{Service}[\omega] :! \boxed{\text{Run}(s)} .s(m).([m = s]\overline{\text{End}}(s) + [m \neq s]\overline{\text{Run}}(s))$$


$$\text{Handler}[k] :! \boxed{\text{Start}(s)} .\overline{\text{Run}}(s).\text{End}(s)$$

$$\text{Server}[1] : (\nu \text{ses}) ! \boxed{\text{adr}(c)} .\overline{c}(\text{ses}).\overline{\text{Start}}(\text{ses})$$

The corresponding diagram (for $k = 3$) is depicted below.



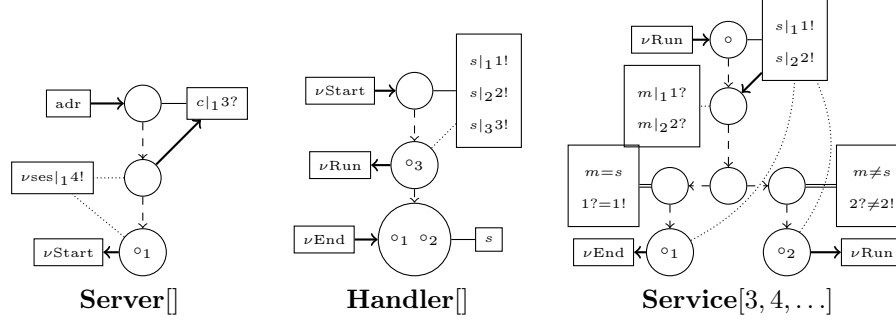
The tokens \circ in the top-level place of each replicator represent the control state. In the textual representation, each token is represented as a frame surrounding an active part of process, a so-called *redex*. Each replicator label (Server, etc.) is associated to the set of available threads. For example, there is at most one thread (indexed 1) running in the server, three in the handler and an infinity in the service. In the first step, the server receives through the public channel adr a request for connection from a client.



The client request is a name bound to variable c , and it is given a *fresh identity* denoted $1?$. The use of logical clock values as fresh names is a particularity of pi-graphs, which as explained in [3] gives a natural interpretation of freshness. Note also that the token within the server is labelled 1, which identifies the actual thread running the process, so the value of c is only $1?$ from the point of view of thread 1, hence the notation $c|_1 1?$. Distinct threads may “see” different values.

In the second step, the locally private channel νses is sent back to the client. This channel will be used for further communications between the client and its specific service. Since the client is unknown (i.e. in the external environment), a *fresh output* name $1!$ is generated. This provides a concrete interpretation of a *bound output* in the pi-calculus terminology.

We depict below a reachable configuration of the system with two service threads running in parallel.



In the thread with identity 1 the session channel is $1!$ (bound to s) and a message $1?$ (bound to m) has been received. The match branch of the choice has been taken and thus $1?$ is asserted equal to $1!$ (i.e. $m = s$ for thread 1). The received message is thus the session channel itself, which triggers the termination of the service. In the second thread (identity 2), the message $2?$ is indeed distinct from the session channel $2!$ (i.e. $m \neq s$ for thread 2), and the service can be reentered at once through νRun (an example of self-recursion). In the Handler threads 1 and 2 are waiting for the end of their corresponding service. A further thread 3 is ready to start a service, which means a third client connection has been effected. We can see that any other potential client is now waiting for the handler to release at least one session (since $k = 3$).

3 Graphical representation

Process calculi are most often presented as term-rewrite systems whereas the pi-graphs follow the principle of *graph relabelling*. In this section we discuss the graphical model underlying the process algebra. Firstly, a precise definition of *names* manipulated by pi-graphs is required.

Definition 1. *The set of names \mathcal{N} is the disjoint union of the following sets:*

- \mathcal{N}_f the set of free names a, b, \dots
- \mathcal{N}_r the set of global restrictions $\nu A, \nu B, \dots$
- \mathcal{N}_v the set of input variables x, y, \dots
- \mathcal{N}_p the set of local private names $\nu a, \nu b, \dots$
- \mathcal{N}_s the set of shared names $\nu a.\ell, \nu b.\ell' \dots$ ($\ell, \ell' \in \mathbb{N}$)
- $\mathcal{N}_o \stackrel{\text{def}}{=} \{n! \mid n \in \mathbb{N}\}$ the set of output names
- $\mathcal{N}_i \stackrel{\text{def}}{=} \{n? \mid n \in \mathbb{N}\}$ the set of input names

We define $\text{Priv} \stackrel{\text{def}}{=} \mathcal{N}_r \cup \mathcal{N}_p \cup \mathcal{N}_s$ (*private names*), $\text{Pub} \stackrel{\text{def}}{=} \mathcal{N} \setminus \text{Priv}$ (*public names*), $\text{Local} \stackrel{\text{def}}{=} \mathcal{N}_p \cup \mathcal{N}_v$ and $\text{Global} \stackrel{\text{def}}{=} \mathcal{N}_r \cup \mathcal{N}_f \cup \mathcal{N}_s$.

This categorization is required because names are *globalized* in the semantics (i.e. all names have global scope). The sets \mathcal{N}_i and \mathcal{N}_o are names whose identity is generated fresh by construction. As illustrated before, they play a prominent role in the model. Together with the shared names in \mathcal{N}_s - private names that have been communicated internally - they represent the *dynamic resources* produced and consumed by pi-graphs.

Definition 2. a *pi-graph* π is a tuple $\langle R, J, V, L, E, M \rangle$ with

- R a set of replicator labels
- $J \in R \rightarrow \mathbb{P}(\mathbb{N})$ the available thread identifiers,
- $V \stackrel{\text{def}}{=} \bigcup_{r \in R} V_r$ a set of vertices (with all subsets V_r assumed disjoint),
- $L \in V \times \{o, i, g=, g\neq\} \rightarrow \mathcal{N}^2$ a vertex labelling,
- $E \subset V \times V$ a set of edges,
- $M \in V \rightarrow \mathbb{P}(\mathbb{N})$ a marking function,

A pi-graph π is composed of a set of replicators, each one identified by a label r in set R^π . The vertices of the pi-graph define set $V^\pi \stackrel{\text{def}}{=} \bigcup_{r \in R^\pi} V_r^\pi$ where each set V_r^π identifies the vertices within replicator r (the vertex sets are assumed disjoint). The tokens in the diagrams - and equivalently the textual redexes - correspond in graphical terms to markings of vertices. We denote $M^\pi(v) \subseteq \mathbb{N}$ the set of tokens for vertex $v \in V^\pi$. In a diagram, if a node v (of replicator r) contains a token \circ_i , then $i \in M_r^\pi(v)$ represents the identity of the corresponding thread. We denote $J_r^\pi \subseteq \mathbb{N} \cup \{\omega\}$ the set of remaining thread identifiers in iterator r . If $J_r^\pi \neq \emptyset$ then a token \circ is put in the initial place of the replicator. The maximum number of running threads of replicator r is denoted k_r^π .

Particularly important is the *context* in which the graphs are rewritten.

Definition 3. A *pi-graph context* Δ is a triple $\beta; \gamma; \kappa$ with:

- $\beta \in \mathcal{N} \rightarrow \mathcal{N}$ the **name environment**
- $\kappa \in \mathcal{N}_o \rightarrow \mathbb{P}(\mathcal{N}_i)$ a **causal clock**, and
- $\gamma \stackrel{\text{def}}{=} (\gamma^=, \gamma^\neq) \in (\mathcal{N} \times \mathcal{N})^2$ a **dynamic match/mismatch partition**.

The name environment β relates the syntactic - and static - occurrences of names in the diagrams to their dynamic counterpart. To isolate a name x whose scope is bound to a given replicator r and thread j (e.g. a private name or an input variable), the manipulation functions of Table 1 reference such name as x_j^r . The **reset** function is used when a thread j finishes, and restores the identity for all its local names⁵. A special case is when a private name has been communicated to another thread and in such a situation it is replaced by a shared private name attached to a fresh identity $\ell \in \mathbb{N}$. The causal clock κ is used to

⁵ For the sake of compactness, the identity is not explicitly encoded in the name environment β , but it is enforced by its manipulation functions.

Name environment β	
reference	$\beta_j^r(x) \stackrel{\text{def}}{=} \begin{cases} \beta(x) & \text{if } x \in \text{dom}(\beta), \\ x_j^r & \text{if } x \in \text{Local}, \\ x & \text{otherwise} \end{cases}$
update ⁶	$\beta_j^r[y/x] \stackrel{\text{def}}{=} \begin{cases} \beta \triangleright \{x \mapsto y\} \cup \{z \mapsto y \mid \beta(z) = x\} & \text{if } x \in \text{Global} \\ \beta \triangleright \{x_j^r \mapsto y\} \cup \{z \mapsto y \mid \beta(z) = x_j^r\} & \text{otherwise} \end{cases}$
reset	$\text{reset}_j^r(\beta) \stackrel{\text{def}}{=} \beta' \setminus \{x_j^r \mapsto y \mid x_j^r \in \text{dom}(\beta)\}$
	$\text{with } \beta' \stackrel{\text{def}}{=} \beta \triangleright \{z \mapsto x:\ell \mid z \mapsto x_j^r \in \beta, \ell = \min(\mathbb{N}^+ \setminus \{\ell' \mid x_{:\ell'} \in \beta\})\}$
Causal clock κ	
fresh output	$\text{out}(\kappa) \stackrel{\text{def}}{=} \kappa \cup \{\text{next}_o(\kappa)! \mapsto \emptyset\}$
fresh input	$\text{in}(\kappa) \stackrel{\text{def}}{=} \{o \mapsto (\kappa(o) \cup \{\text{next}_i(\kappa)?\}) \mid o \in \text{dom}(\kappa)\}$
freshness (in)	$\text{next}_i(\kappa) \stackrel{\text{def}}{=} \min(\mathbb{N}^+ \setminus \{n \mid n? \in \bigcup \text{cod}(\kappa)\})$
freshness (out)	$\text{next}_o(\kappa) \stackrel{\text{def}}{=} \min(\mathbb{N}^+ \setminus \{n \mid n! \in \text{dom}(\kappa)\})$
r-w causality	$x \prec_\kappa y \stackrel{\text{def}}{=} x \in \text{dom}(\kappa) \wedge y \in \kappa(x)$
unify	$\kappa_{\triangleleft x=y} \stackrel{\text{def}}{=} \left\{ n! \mapsto \begin{cases} \kappa(n!) \setminus \{x\} & \text{if } y \notin \kappa(n!) \\ \kappa(n!) \setminus \{y\} & \text{if } x \notin \kappa(n!) \\ \kappa(n!) & \text{otherwise} \end{cases} \mid n! \in \text{dom}(\kappa) \right\}$
Partition $\gamma \stackrel{\text{def}}{=} (\gamma^=, \gamma^\neq)$	
match	$x =_\gamma y \text{ iff } \begin{cases} x = y \vee (x, y) \in \gamma^= \vee (y, x) \in \gamma^= \\ \vee \exists z, x =_\gamma z \wedge z =_\gamma y \end{cases}$
mismatch	$x \neq_\gamma^\kappa y \text{ iff } \begin{cases} (x, y) \in \gamma^\neq \vee (y, x) \in \gamma^\neq \vee y \neq_\gamma^\kappa x \\ \vee (\exists z_1 z_2, x =_\gamma z_1 \wedge z_1 \neq_\gamma^\kappa z_2 \wedge z_2 =_\gamma y) \\ \vee (x, y \in \text{Priv} \cup \mathcal{N}_o \wedge x \neq y) & \text{- private } \neq \text{ private} \\ \vee (x \in \mathcal{N}_f \wedge y \in \text{Priv} \cup \mathcal{N}_o) & \text{- public } \neq \text{ private} \\ \vee (x \in \mathcal{N}_o \wedge y \in \mathcal{N}_i \wedge x \not\prec_\kappa y) & \text{- read-write order} \end{cases}$
compatibility	$x \sim_\kappa y \text{ iff } \begin{cases} \neg(x \neq_\gamma^\kappa y) \wedge (x =_\gamma y \vee y \sim_\kappa x) \\ \vee (x, y \in \mathcal{N}_f \cup \mathcal{N}_i) & \text{- free and input compatibles} \\ \vee (x \in \mathcal{N}_o \wedge y \in \mathcal{N}_i \wedge x \prec_\kappa y) & \text{- write-read order} \end{cases}$
refine	$\gamma_{\triangleleft x=y} \stackrel{\text{def}}{=} (\gamma^= \cup \{(x, y)\}, \gamma^\neq) \text{ if } \neg(x =_\gamma y), \gamma \text{ otherwise}$
refine set	$\gamma \triangleleft (\{(x, y)\} \cup E) \stackrel{\text{def}}{=} \gamma_{\triangleleft x=y} \triangleleft E \text{ and } \gamma \triangleleft \emptyset \stackrel{\text{def}}{=} \gamma$
discriminate	$\gamma_{\triangleleft x \neq^\kappa y} \stackrel{\text{def}}{=} (\gamma^=, \gamma^\neq \cup \{(x, y)\}) \text{ if } \neg(x \neq_\gamma^\kappa y), \gamma \text{ otherwise}$

Table 1. The pi-graph context and associated operations.

generate fresh identities for names received from or sent to the environment, and it also implements *read-write causality* [7]. Finally, the γ component allows to record which names are *effectively* equal (in $\gamma^=$) (resp. unequal in γ^\neq) after a successful match (resp. mismatch). The functions manipulating the context, as defined in Table 1, will be discussed in the next section.

⁶ For partial functions f_1 and f_2 , $f_1 \triangleright f_2 \stackrel{\text{def}}{=} f_1 \setminus \{x \mapsto f_1(x) \mid x \in \text{dom}(f_2)\} \cup f_2$

The interpretation of a diagram $D \stackrel{\text{def}}{=} \nu(A_1, \dots, A_p) R_1 \dots R_n$ is a graph π such that the statement $D \models \pi$ is provable by the following rule:

$$\frac{\begin{array}{l} P_1 \models_{A_1 \mapsto \nu A_1, \dots, A_p \mapsto \nu A_p, a_{11} \mapsto \nu a_{11}, \dots, a_{1m_1} \mapsto \nu a_{1m_1}} \langle V_1, i_{P_1}, L_1, E_1 \rangle \\ \dots \\ P_n \models_{A_1 \mapsto \nu A_1, \dots, A_p \mapsto \nu A_p, a_{n1} \mapsto \nu a_{n1}, \dots, a_{nm_n} \mapsto \nu a_{nm_n}} \langle V_n, i_{P_n}, L_n, E_n \rangle \\ V_1 \cap \dots \cap V_n = \emptyset \end{array}}{\nu(A_1, \dots, A_p) r_1[k_1] : \nu(a_{11}, \dots, a_{1m_1}) !P_1 \dots r_n[k_n] : \nu(a_{n1}, \dots, a_{nm_n}) !P_n \models \langle \{r_1, \dots, r_n\}, \{r_1 \mapsto [1, k_1], \dots, r_n \mapsto [1, k_n]\}, V_1 \cup \dots \cup V_n, L_1 \cup \dots \cup L_n, E_1 \cup \dots \cup E_n, \{v \mapsto \emptyset \mid v \in V_1 \cup \dots \cup V_n\} \rangle}$$

In the resulting pi-graph π the (disjoint) sets of vertices V_1, \dots, V_n , their labelling L_1, \dots, L_n and their edge sets E_1, \dots, E_n are obtained by proving, for each sub-process P_1, \dots, P_n an assertion of the form $P_j \models_\sigma \langle V_j, i_{P_j}, L_j, E_j \rangle$ where i_{P_j} is the node corresponding to the initial action of P_j . The substitution σ maps the bound names (the global restrictions and the private names) to their value (in \mathcal{N}_s or \mathcal{N}_p). The rules for sub-processes are as follows :

$$\frac{\alpha \models_\sigma \langle v_\alpha, L_\alpha \rangle \quad P \models_\sigma \langle V_P, i_P, L_P, E_P \rangle \quad v_\alpha \notin V_P}{\alpha.P \models_\sigma \langle \{v_\alpha\} \cup V_P, v_\alpha, L_\alpha \cup L_P, \{v_\alpha \mapsto i_P\} \cup E_P \rangle}$$

$$\frac{P \models_\sigma \langle V_P, i_P, L_P, E_P \rangle \quad L_\phi \stackrel{\text{def}}{=} \begin{cases} (v_\phi, g_=) \mapsto (\sigma(x), \sigma(y)) \text{ if } \phi = [x = y] \\ (v_\phi, g_\neq) \mapsto (\sigma(x), \sigma(y)) \text{ if } \phi = [x \neq y] \end{cases} \quad v_\phi \notin V_P}{\phi P \models_\sigma \langle \{v_\phi\} \cup V_P, v_\phi, L_\phi \cup L_P, \{v_\phi \mapsto i_P\} \cup E_P \rangle}$$

$$\frac{\begin{array}{l} P_1 \models_\sigma \langle V_{P_1}, i_{P_1}, L_{P_1}, E_{P_1} \rangle \\ P_2 \models_\sigma \langle V_{P_2}, i_{P_2}, L_{P_2}, E_{P_2} \rangle \end{array} \quad v_+ \notin V_{P_1} \cup V_{P_2} \quad V_{P_1} \cap V_{P_2} = \emptyset}{P_1 + P_2 \models_\sigma \langle \{v_+\} \cup V_{P_1} \cup V_{P_2}, v_+, L_{P_1} \cup L_{P_2}, \{v_+ \mapsto i_{P_1}, v_+ \mapsto i_{P_2}\} \cup E_{P_1} \cup E_{P_2} \rangle}$$

Each action α generates a single vertex v_α yielding an assertion $\alpha \models_\sigma \langle v_\alpha, L_\alpha \rangle$ provable by one of the following rules:

$$\frac{}{\tau \models_\sigma \langle v_\tau, \emptyset \rangle}$$

$$\frac{}{\bar{c}(a) \models_\sigma \langle v_o, \{(v_o, c) \mapsto (\sigma(c), \sigma(c)), (v_o, o) \mapsto (\sigma(a), \sigma(a))\} \rangle}$$

$$\frac{}{c(x) \models_\sigma \langle v_i, \{(v_i, c) \mapsto (\sigma(c), \sigma(c)), (v_i, i) \mapsto (x, x)\} \rangle}$$

The most important property of the graphical translation is that it generates graphs that have a size comparable to that of the initial (syntactic) diagrams.

Definition 4. Let D be a diagram, we denote $|D|$ its size defined as follows:

$$\begin{cases} |\nu \tilde{A} r_1[k_1] : \nu \tilde{a}_1 !P_1 \dots r_n[k_n] : \nu \tilde{a}_n !P_n| \stackrel{\text{def}}{=} \sum_{i=1}^n |P_i| \\ |\alpha.P| \stackrel{\text{def}}{=} 1 + |P| \\ |\phi P| \stackrel{\text{def}}{=} 1 + |P| \\ |P_1 + P_2| \stackrel{\text{def}}{=} 1 + |P_1| + |P_2| \end{cases}$$

The size of a pi-graph π , denoted $|\pi|$, with vertices V is the cardinal of V .

In the diagrams we omit the restrictions and private names because they do not generate nodes in the translation. Moreover, in the translated graphs the name holders (square nodes) are not taken into account because at most two of these can be associated to each vertex.

Theorem 1. *Let D be a well-formed diagram. Then there exists a graph π such that $D \models \pi$. Moreover, $|\pi| = |D|$.*

Proof. By a simple structural induction on the translation rules □

4 Operational semantics

The operational semantics manipulates terms of the form $\Delta \vdash \pi$ where π is pi-graph (with explicit redexes) and $\Delta \stackrel{\text{def}}{=} \beta; \gamma; \kappa$ a context. The *initial context* has all empty components, it is denoted $\Delta_0 \stackrel{\text{def}}{=} \emptyset; (\emptyset, \emptyset); \emptyset$.

The rewrite rules are decomposed into layers. At the lower-level, the *commitment rules* (cf. Table 2) explain the atomic actions of pi-graphs. Above these, the *process rules* (cf. Table 3) describe the *local* behavior of processes. The *replicator rules* (cf. Table 4) describe what can happen for a prefix or suffix action of a replicator. Finally, there is a single *diagram rule* acting at the global-level.

[silent]	$\beta; \gamma; \kappa \vdash r[j] : \tau \dashrightarrow \beta; \gamma; \kappa$
[out]	$\beta; \gamma; \kappa \vdash r[j] : \bar{c}(a) \dashrightarrow \beta; \gamma; \kappa \quad \text{if } c' \stackrel{\text{def}}{=} \beta_j^r(c), a' \stackrel{\text{def}}{=} \beta_j^r(a) \in \text{Pub}$
[o-fresh]	$\beta; \gamma; \kappa \vdash r[j] : \bar{c}(a) \dashrightarrow \beta_j^r[\text{next}_o(\kappa)!/a]; \gamma; \text{out}(\kappa)$ if $c' \stackrel{\text{def}}{=} \beta_j^r(c) \in \text{Pub}$ and $\beta_j^r(a) \in \text{Priv}$
[i-fresh]	$\beta; \gamma; \kappa \vdash r[j] : c(x) \dashrightarrow \beta_j^r[\text{next}_i(\kappa)?/x]; \gamma; \text{in}(\kappa)$ if $c' \stackrel{\text{def}}{=} \beta_j^r(c) \in \text{Pub}$
[com]	$\beta; \gamma; \kappa \vdash r_1[j_1] : \bar{c}(a) \parallel r_2[j_2] : d(x) \dashrightarrow \beta_{j_2}^{r_2}[a'/x]; \gamma_{\triangleleft c'=d'}; \kappa_{\triangleleft c'=d'}$ if $c' \stackrel{\text{def}}{=} \beta_{j_1}^{r_1}(c)$, $a' \stackrel{\text{def}}{=} \beta_{j_1}^{r_1}(a)$ and $d' \stackrel{\text{def}}{=} \beta_{j_2}^{r_2}(d)$

Table 2. Operational semantics: commitment rules.

Each atomic action performed by a pi-graph is handled by a specific commitment rule. The [silent] rule explains that a τ action performed in the context of a replicator r and thread j yields a commitment denoted \dashrightarrow . The context remains unmodified. The [out] rule generates a commitment $\bar{c}(a) \dashrightarrow$, where c' and a' are the values bound to “variables” c and a in the environment β . Both names must be public to trigger the rule. The rule [o-fresh] explains the emission

[prefix]	$\beta; \gamma; \kappa \vdash r : \boxed{\alpha}_j.P \xrightarrow{\mu} \beta'; \gamma'; \kappa' \vdash r : \alpha.\boxed{P}_j$ if $\beta; \gamma; \kappa \vdash r[j] : \alpha \xrightarrow{\mu} \beta'; \gamma'; \kappa'$	
[match]	$\beta; \gamma; \kappa \vdash r : \boxed{[a = b]}_j.P \xrightarrow{\varepsilon} \beta; \gamma_{\triangleleft a' = b'}; \kappa_{\triangleleft a' = b'} \vdash r : [a = b]\boxed{P}_j$ if $a' \mathcal{L}_\kappa b'$ with $a' \stackrel{\text{def}}{=} \beta_j^r(a)$ and $b' \stackrel{\text{def}}{=} \beta_j^r(b)$	
[miss]	$\beta; \gamma; \kappa \vdash r : \boxed{[a \neq b]}_j.P \xrightarrow{\varepsilon} \beta; \gamma_{\triangleleft a' \neq \kappa b'}; \kappa \vdash r : [a \neq b]\boxed{P}_j$ if $\neg(a' =_\gamma b')$ with $a' \stackrel{\text{def}}{=} \beta_j^r(a)$ and $b' \stackrel{\text{def}}{=} \beta_j^r(b)$	
[sync]	$\beta; \gamma; \kappa \vdash r_1 : \boxed{\alpha_1}_{j_1}.P \parallel r_2 : \boxed{\alpha_2}_{j_2}.Q \xrightarrow{\tau} \beta'; \gamma'; \kappa' \vdash r_1 : \alpha_1.\boxed{P}_{j_1} \parallel r_2 : \alpha_2.\boxed{Q}_{j_2}$ if $\beta; \gamma; \kappa \vdash r_1[j_1] : \alpha_1 \parallel r_2[j_2] : \alpha_2 \xrightarrow{\tau} \beta'; \gamma'; \kappa'$	
[sum _L]	$\beta; \gamma; \kappa \vdash r : P_1 \boxed{+}_j.P_2 \xrightarrow{\varepsilon} \beta; \gamma; \kappa \vdash r : \boxed{P_1}_j + P_2$	– also [sum _R]

Table 3. Operational semantics: process rules.

of a private name (either a restriction νA , a private name νa or a shared name $\nu a.\ell$) over a public channel. As illustrated previously, a fresh value - denoted $\text{next}_o(\kappa)!$ - is generated for the output name. The clock itself is modified so that the freshness guarantee is preserved. Part of the [o-fresh] commitment is the update of the environment to reflect the renaming by a fresh value. Note that if the name a is a global restriction, then the update is *permanent*. If it is a private name local to a thread, then the binding is tied to the life-cycle of that thread, except if it is shared with another thread. The [i-fresh] rule is almost the symmetric, it deals with names received from the environment, picked up freshly as $\text{next}_i(\kappa)?$. Finally, the [com] rule explains the commitment of value-passing between two threads (here separated by a \parallel ; they may arise from two different replicators and/or threads). For the sake of *compositionality*, a communication can take place between an output and an input on two *distinct* channels c' and d' , as long as they can be equated explicitly, which we denote $c' \mathcal{L}_\kappa d'$ (cf. the [match] case below). The environment is updated to reflect this new equality.

Each process rule of Table 3 provides the description of a local rewrite of the form: $\beta; \gamma; \kappa \vdash \theta \xrightarrow{\mu} \beta'; \gamma'; \kappa' \vdash \theta'$ where $\beta; \gamma; \kappa$ is the context before the rewrite, and θ is a “subterm” of the pi-graph π with some redex(es) inside. In graphical terms, θ identifies the replicator r and a set of vertices $V_\theta \subseteq V_r^\pi$ of the graph π where the rewrite takes place. The [prefix] rule, for instance, describes the execution of an action in prefix position within a process, if the corresponding commitment can be proved. Here, the subterm θ is $\alpha.P$ and V_θ is defined as $\{v_\alpha\} \cup V_P$ where v_α is the single vertex corresponding to α , and V_P the vertices of (the graphical representation of) P . The notation $\boxed{\alpha}_j$ means that α is a *redex* for a thread identified as j . In the underlying graph, j is in $M^\pi(v_\alpha)$, i.e. the vertex v_α is marked with a token indexed by j . More generally, a redex can be any action $\boxed{\alpha}_j$, a guard $\boxed{\phi}_j.P$ or a sum $P_1 \boxed{+}_j.P_2$. The notation extends to

arbitrary processes beyond action suffixes with $\boxed{\alpha.P}_j \stackrel{\text{def}}{=} \boxed{\alpha}_j.P$, $\boxed{\phi P}_j \stackrel{\text{def}}{=} \boxed{\phi}_j.P$ and $\boxed{P_1 + P_2}_j \stackrel{\text{def}}{=} P_1 \boxed{+}_j.P_2$. A prefixed (resp. guarded) process is thus in its initial state for thread j if its initial prefix (resp. guard) is a redex. So in the right-hand side of the [prefix] rule the initial action of the continuation P (i.e. a place $v \in V_p$) received the token. The label μ of the transition as well as the update of the context are given by the commitment triggered for the action.

The [match] and [miss] rules are similar to a mix of commitment and prefixing, but cannot be decomposed because they are not “visible” commitment. Importantly, such *invisible steps*, labelled ϵ , cannot commit the choice of a given branch in a sum. Upon matching, the two names to equate - denoted a' and b' - must satisfy the *compatibility* relation $a' \mathcal{Z}_\kappa b'$ (cf. Table 1). It is often the case that names are compatible, e.g. two free names are compatible as long as they are not explicitly inequated (by a mismatch). For dynamic names in \mathcal{N}_i and \mathcal{N}_o read-write causality is enforced by the causal clock κ . An output name $o!$ can only be equated to an input name $i?$ if the latter is *causally dependent* on the former, i.e. $o!$ “escaped” to the environment *before* $i?$ was received. This is the *write-read order* when $i? \in \kappa(o!)$. In the case of *read-write order* $?i \notin \kappa(o!)$, i.e. if $i?$ is received before $o!$ escapes, then the equality is forbidden. There are other subtleties, e.g. to reflect the equality in both γ and κ , but the formalization is routine. The mismatch is often omitted from pi-calculi (note that our introductory example motivates its use), mostly due to its non-trivial characterization. A natural treatment - similar to match - results from the concrete pi-graph semantics. The [miss] rule is very close to [match] except for the condition of name *incompatibility*. In Table 1 the relation \neq_γ lists all the “natural” incompatibilities among names: e.g. restricted names are in essence distinguished, etc. Once again, read-write causality is enforced.

The synchronization rule [sync] involves two local contexts : an emission and a reception. These can be within the same replicator (albeit in distinct threads) or between two replicators. The communication must be committed (through [com]) and the tokens are “pushed” as in the [prefix] case.

Finally, the choice construct is interpreted by the [sum_L] rule (and its symmetric [sum_R]). The choice is made *arbitrarily* in the local context, but the chosen branch is *not* committed since an invisible ϵ -transition is taken. The notion of abstracted transition (cf. below) allows to recover the usual interpretation of the choice, which is to find a valid commitment within the branch.

Above the process rules are defined replicator rules involving at least a replicator-level context, i.e. its initial or one of its final actions. The [spawn] rule, for instance, is a variant of [prefix] that initiates a new thread with identity j , if available. The [term] rule explains the (normal) termination of a thread by its suffix action. The thread identity j is given back to the replicator, and the local names are reset. As explained previously, the private names of j that have been communicated to other threads become explicitly shared upon termination. The other rules are variants of [spawn] and [term] involving synchronization, which we interpret as “procedure” calls. The [call_L] (resp. [call_R]) rule synchronizes a

[spawn]	$\beta; \gamma; \kappa \vdash r[J] : !\boxed{\alpha}.P \xrightarrow{\mu} \beta'; \gamma'; \kappa' \vdash r[J \setminus \{j\}] : !\alpha.\boxed{P}_j$	with $j \stackrel{\text{def}}{=} \min(J)$
	$\text{if } \beta; \gamma; \kappa \vdash r[j] : \alpha.P \xrightarrow{\mu} \beta'; \gamma'; \kappa'$	
[term]	$\beta; \gamma; \kappa \vdash r[J] : \boxed{\alpha}_j \xrightarrow{\mu} \text{reset}_j^r(\beta'); \gamma'; \kappa' \vdash r[J \cup \{j\}] : \alpha$	
	$\text{if } \beta; \gamma; \kappa \vdash r[j] : \alpha \xrightarrow{\mu} \beta'; \gamma'; \kappa'$	
[call _L]	$\beta; \gamma; \kappa \vdash r_1[J_1] : !\boxed{\alpha_1}.P \parallel r_2 : \boxed{\alpha_2}_{j_2}.Q$	with $j_1 \stackrel{\text{def}}{=} \min(J_1)$
	$\xrightarrow{\tau} \beta'; \gamma'; \kappa' \vdash r_1[J_1 \setminus \{j_1\}] : !\alpha_1.\boxed{P}_{j_1} \parallel r_2 : \alpha_2.\boxed{Q}_{j_2}$	
	$\text{if } \beta; \gamma; \kappa \vdash r_1[j_1] : \alpha_1 \parallel r_2[j_2] : \alpha_2 \xrightarrow{\tau} \beta'; \gamma'; \kappa' - \text{also } [\text{call}_{R}] \text{ and } [\text{call}_{LR}]$	
[call _{Lt}]	$\beta; \gamma; \kappa \vdash r_1[J_1] : !\boxed{\alpha_1}.P \parallel r_2[J_2] : \boxed{\alpha_2}_{j_2}$	with $j_1 \stackrel{\text{def}}{=} \min(J_1)$
	$\xrightarrow{\tau} \text{reset}_{j_2}^{r_2}(\beta'); \gamma'; \kappa' \vdash r_1[J_1 \setminus \{j_1\}] : !\alpha_1.\boxed{P}_{j_1} \parallel r_2[J_2 \cup \{j_2\}] : \alpha_2$	
	$\text{if } \beta; \gamma; \kappa \vdash r_1[j_1] : \alpha_1 \parallel r_2[j_2] : \alpha_2 \xrightarrow{\tau} \beta'; \gamma'; \kappa' - \text{also } [\text{call}_{Rt}]$	
[sync _{Lt}]	$\beta; \gamma; \kappa \vdash r_1 : \boxed{\alpha_1}_{j_1}.P \parallel r_2[J_2] : \boxed{\alpha_2}_{j_2}$	
	$\xrightarrow{\tau} \text{reset}_{j_2}^{r_2}(\beta'); \gamma'; \kappa' \vdash r_1 : \alpha_1.\boxed{P}_{j_1} \parallel r_2[J_2 \cup \{j_2\}] : \alpha_2$	
	$\text{if } \beta; \gamma; \kappa \vdash r_1[j_1] : \alpha_1 \parallel r_2[j_2] : \alpha_2 \xrightarrow{\tau} \beta'; \gamma'; \kappa' - \text{also } [\text{sync}_{Rt}] \text{ and } [\text{sync}_{tt}]$	

Table 4. Operational semantics: replicator rules.

replicator entry on the left (resp. on the right) with a local process on the right (resp. on the left), and intertwines the communication with the spawn of a new thread (i.e. it is a mix of [sync] and [spawn]). The rule [call_{LR}] is when two replicator entries synchronize. The special case [call_{Lt}] (resp. [call_{Rt}]) is when the right-hand (resp. left-hand) is the suffix of another thread, which means there is a simultaneous [sync], [spawn] and [term]. Finally, the [sync_X] rules ($X \in \{Lt, Rt, tt\}$) involve a local process and a termination, or two simultaneous terminations. The redundancy here is the price to pay for *concreteness*, and each rule corresponds to a situation worth studying.

As explained previously, the left-hand side subterm θ in each non-global rule corresponds to a set V_θ of vertices where the local rewrite takes place in the pigraph π . We write $\pi[\theta]$ to identify explicitly this set of vertices in the context of the whole graph π . The unique global-level rule reflects the application of a process or replicator rule at the global level:

Definition 5. A *global transition* $\beta; \gamma; \kappa \vdash \pi[\theta] \xrightarrow{\mu} \text{gc}(\beta'; \gamma'; \kappa') \vdash \pi[\theta']$ is produced if $\beta; \gamma; \kappa \vdash \theta \xrightarrow{\mu} \beta'; \gamma'; \kappa' \vdash \theta'$ can be inferred by either a process or a replicator rule.

The gc function “garbage collects” all the *inactive* names from the context. This clean-up allows to strictly control the usage of dynamic resources, and it is a key player of our main results (cf. Section 5).

Definition 6. In a context $\beta; \gamma; \kappa$, the set of **inactive names** is:

$$\text{inact}(\beta; \gamma; \kappa) \stackrel{\text{def}}{=} \{n \mid (n \in \mathcal{N}_i \cup \mathcal{N}_s \wedge n \notin \text{cod}(\beta)) \\ \vee (n \in \mathcal{N}_o \wedge n \notin \text{cod}(\beta) \wedge \neg(\exists m \in \text{cod}(\beta), n =_\gamma m))\}$$

The input and shared names are considered inactive if they are not used in β . On the contrary, it is possible that an output name - not explicitly referenced in β - is indeed equated to another name in γ , which means it is still active. For example, if two names $o!$ and $i?$ are considered equal (i.e. $o! =_\gamma i?$), even if $o!$ is not referenced (i.e. $o! \notin \text{cod}(\beta)$), any other $o!$ must be considered distinct from $i?$ as long as the latter remains active.

Now, the garbage collection function simply consists in removing all the occurrences of the inactive names in the context.

Definition 7. In a pi-graph context $\beta; \gamma; \kappa$, let $E \stackrel{\text{def}}{=} \text{inact}(\beta; \gamma; \kappa)$, then:

$$\text{gc}(\beta; \gamma; \kappa) \stackrel{\text{def}}{=} \beta; \gamma'; \kappa' \text{ such that } \begin{cases} \gamma' \stackrel{\text{def}}{=} (\emptyset \triangleleft \{(x, y) \mid x \neq y \wedge x =_\gamma y \wedge x, y \notin E\}, \\ \{(x, y) \in \gamma \neq \mid x, y \notin E\}) \\ \kappa' \stackrel{\text{def}}{=} \{(x, \kappa(x) \setminus E) \mid x \in \text{dom}(\kappa) \setminus E\} \end{cases}$$

Abstracted transitions Despite their key role in describing the semantics of pi-graphs, the non-committing ϵ -transitions must be abstracted away when comparing behaviors at the global level. However, a “blind” abstraction of ϵ -transitions allows incorrect states to be reached, e.g. a state in which a branch in a choice has been taken (through rule $[\text{sum}_L]$ yielding an ϵ) but with a first action that cannot be committed (i.e. no commitment rule applies). A solution is to enforce a causal dependence on abstracted ϵ -transitions, using the notion of *causal sequence*.

Definition 8. Let $\Delta_1 \vdash \pi_1 \xrightarrow{\alpha_1} \dots \xrightarrow{\alpha_{n-1}} \Delta_n \vdash \pi_n \xrightarrow{\alpha_n} \Delta_{n+1} \vdash \pi_{n+1}$ be a sequence of transitions. The sequence is **causal** iff for any i , $1 \leq i < n$, there exists j , $i < j \leq n$, and $\theta_i, \theta'_i, \theta_j, \theta'_j$, with $\Delta_i \vdash \pi_i[\theta_i] \xrightarrow{\alpha_i} \Delta_{i+1} \vdash \pi_i[\theta'_i]$ and $\Delta_j \vdash \pi_j[\theta_j] \xrightarrow{\alpha_j} \Delta_{j+1} \vdash \pi_j[\theta'_j]$ such that $M^{\pi_i}(V_{\theta'_i}) \cap M^{\pi_j}(V_{\theta'_j}) \neq \emptyset$.

In such a causal sequence, each transition but the last one produces tokens that are necessarily consumed by a further transition in the sequence. In particular the last transition is causally dependent on all the previous ones. The abstraction principle is then expressed as follows:

Definition 9. An **abstracted transition** $\Delta_1 \vdash \pi_1 \xrightarrow{\alpha} \Delta_n \vdash \pi_n$ is inferred iff there is a causal sequence $\Delta_1 \vdash \pi_1 \xrightarrow{\epsilon} \dots \xrightarrow{\epsilon} \Delta_n \vdash \pi_n \xrightarrow{\alpha} \Delta_{n+1} \vdash \pi_{n+1}$ ($\alpha \neq \epsilon$).

The definition says that an abstracted transition with label α corresponds to a path of length n beginning with $n - 1$ ϵ -transitions, and ending with a single non- ϵ -transition labelled α . Moreover, all the intermediate invisible transitions are causally needed to produce the visible one. The following strengthens the definition.

Theorem 2. Every causal ϵ -path starting from a state $\Delta \vdash \pi$ is finite.

Proof. The syntax requires that the each replicator process is prefixed with an action, which cannot yield an ϵ -transition; hence an ϵ -path cannot spawn threads. Moreover, only finite sequences of ϵ can be triggered by a given thread (in fact, a bound can be easily computed from the syntax). Finally, causal paths involve at most the interleaving of two threads (in the case of [com] commitments), which allows to conclude \square

Stateful bisimilarity The standard notion of bisimulation is not sufficient to compare pi-graph behaviours precisely enough. For example a process $[a = b] \bar{c}\langle a \rangle$ (put in an adequate replicator) will be distinguished with e.g. $[a = b] \bar{c}\langle b \rangle$ because the generated labels, respectively $\bar{c}\langle a \rangle$ and $\bar{c}\langle b \rangle$, are distinct! Bisimulation is indeed a *stateless* principle: only transition labels are compared and states are abstracted away. In a pi-graph, the only part of the state that must be inspected to decide whether labels are equal or not is the partition γ . The following definition makes this precise.

Definition 10. In a context $\Delta = \beta; \gamma; \kappa$, the labels μ_1 and μ_2 are said equated, which we note $\mu_1 =_\Delta \mu_2$ iff
$$\begin{cases} \mu_1 = \tau \wedge \mu_2 = \tau \\ \mu_1 = \bar{c}_1\langle a_1 \rangle \wedge \mu_2 = \bar{c}_2\langle a_2 \rangle \wedge c_1 =_\gamma c_2 \wedge a_1 =_\gamma a_2 \\ \mu_1 = c_1(x_1) \wedge \mu_2 = c_2(x_2) \wedge c_1 =_\gamma c_2 \wedge x_1 =_\gamma x_2 \end{cases}$$

The bisimilarity relation we are looking for is thus a *stateful* variant of the former, defined as follows.

Definition 11. A (*stateful*) *bisimulation* S is a symmetric relation on pi-graph states $\Delta \vdash \pi_1$ and $\Delta \vdash \pi_2$ such that $(\Delta \vdash \pi_1, \Delta \vdash \pi_2) \in S$ and $\Delta_1 \vdash \pi_1 \xrightarrow{\mu_1} \Delta' \vdash \pi'_1$ imply there is some μ_2 such that $\Delta \vdash \pi_2 \xrightarrow{\mu_2} \Delta' \vdash \pi'_2$, $[\mu_1]_\Delta = [\mu_2]_\Delta$ and $(\Delta' \vdash \pi'_1, \Delta' \vdash \pi'_2) \in S$

The (*stateful*) *bisimilarity* \sim is the largest stateful bisimulation.

5 Finiteness and boundedness results

Definition 12. Let π be a pi-graph. We denote $\text{lts}(\pi) \stackrel{\text{def}}{=} \langle Q, T \rangle$ its labelled transition system with Q the states reachable from $\Delta_0 \vdash \pi$, and T the set of triplets of the form $(\Delta' \vdash \pi', \alpha, \Delta'' \vdash \pi'')$ such that we can infer $\Delta' \vdash \pi' \xrightarrow{\alpha} \Delta'' \vdash \pi''$.

Replicators allow to characterize infinite systems. For example, the diagram $r[\omega] : a(x).\bar{b}\langle x \rangle$ (with a single replicator) yields an infinite transition system. In the intermediate state after the input and before the output, an unbounded number of threads can accumulate, so the number of reachable states is infinite.

An interesting subclass of the pi-graphs is when the number of threads is bounded, i.e. when in each iterator there is a finite number of tokens, which we denote $T_r^\pi \stackrel{\text{def}}{=} |\text{cod}(M_r^\pi)|$.

Definition 13. A pi-graph π is *thread-bounded* iff $\forall r \in R^\pi, \exists b_r < \omega$ such that $\forall (\Delta' \vdash \pi') \in \text{lts}(\pi), T_r^{\pi'} \leq b_r$. The *thread-bound* is then $\sum_{r \in R^\pi} b_r$.

The server of Example 1 can be shown thread-bounded with bounds $(1, k, k + 1)$, and the thread-bound is thus $2k + 2$. These notions are semantic, and are obviously not decidable in the general case. A syntactic condition is *control-finiteness* when each replicator has an explicit finite bound k_r^π .

Definition 14. A pi-graph π such that $\forall r \in R^\pi, k_r^\pi < \omega$ is said **control-finite**.

Proposition 1. Let π be a pi-graph. If π is control-finite then π is thread-bounded with thread-bound $\sum_{r \in R^\pi} k_r^\pi$.

A decisive advantage of the concrete semantics developed for the pi-graphs is that very precise results can be derived. For instance, we can exhibit upper bounds for the amount of the dynamic resources - the number of output, input and shared names - manipulated by thread-bounded pi-graphs. For a state q , we denote $\text{dyn}(q)$ the dynamic names (i.e. in set $\mathcal{N}_i \cup \mathcal{N}_o \cup \mathcal{N}_s$) that occur in q . We also identify an upper bound for the number of variables (input actions) simultaneously active in a given process:

Definition 15. A process P has **variable-bound** $\mathcal{V}(P)$ such that:

$$\begin{cases} \mathcal{V}(\alpha) \stackrel{\text{def}}{=} 1 \text{ if } \alpha = c(x), 0 \text{ otherwise} \\ \mathcal{V}(\alpha.P) \stackrel{\text{def}}{=} \mathcal{V}(\alpha) + \mathcal{V}(P), \mathcal{V}(\phi P) = \mathcal{V}(P), \mathcal{V}(P_1 + P_2) = \max(\mathcal{V}(P_1), \mathcal{V}(P_2)) \end{cases}$$

We shall denote $\mathcal{V}_r \stackrel{\text{def}}{=} \mathcal{V}(\alpha.P)$ the maximum number of active variables of a replicator $r[k] : \bar{v}a !\alpha.P$. Now we may state an important boundedness result concerning the input and shared private names.

Lemma 1. Let π be a pi-graph with thread-bounds b_r ($r \in R^\pi$) and $\text{lts}(\pi) = \langle Q, T \rangle$. Then, for all $q \in Q$ we have $|\text{dyn}(q) \cap (\mathcal{N}_i \cup \mathcal{N}_s)| \leq \sum_{r \in R^\pi} b_r \mathcal{V}_r$.

Proof. The proof proceeds by contraposition and we first only consider the case for input names. Let $B \stackrel{\text{def}}{=} \sum_{r \in R^\pi} b_r \mathcal{V}_r$ and assume there are two states in Q , $q_{n-1} \stackrel{\text{def}}{=} \Delta_{n-1} \vdash \pi_{n-1}$ and $q_n \stackrel{\text{def}}{=} \Delta_n \vdash \pi_n$, such that $|\text{dyn}(q_{n-1}) \cap \mathcal{N}_i| = B$, $|\text{dyn}(q_n) \cap \mathcal{N}_i| = 1 + B$ and we can infer $q_{n-1} \xrightarrow{c(n?)}$ q_n induced by the [i-fresh] commitment: $\beta_{n-1}; \gamma_{n-1}; \kappa_{n-1} \vdash \rho[j] : \alpha_n \xrightarrow{c(n?)}$ $\beta_n; \gamma_n; \kappa_n$. By Definition 5, the gc function (cf. Def. 7) ensures that $\forall i? \in \text{dyn}(q_{n-1}) \cap \mathcal{N}_i, i? \in \text{cod}(\beta_{n-1})$. This implies that the number of active variables in q_{n-1} (i.e. with an image in β_{n-1}) is exactly B . Moreover, the [i-fresh] commitment requires the existence of a variable x of replicator ρ such that x is inactive in q_{n-1} and active in q_n for thread j . More precisely, $\beta_{n-1}(x_j^\rho) = x_j^\rho$ and $\beta_n(x_j^\rho) = n?$. Thus, the number of active variables in q_n would be $B + 1$, which leads to a contradiction. In conclusion the state q_n is not reachable and thus $|\text{dyn}(q_n) \cap \mathcal{N}_i| \leq B$. For the shared name the reasoning steps are similar, because the gc function requires the names to exist in $\text{cod}(\beta) \cap \mathcal{N}_s$, and as seen above this codomain is bound by B ; moreover, a variable name may not be associated to both an input name and a shared private name, hence the formula \square

For the output names the situation is more subtle. First, the restrictions and private names can be associated to output names in the environment β . However, we need to take into account only those names who are objects of output prefixes, hence the following definition.

Definition 16. *The **escape-bound** of:*

- a pi-graph $\pi \stackrel{\text{def}}{=} \nu \tilde{A} r_1[k_1] : \nu \tilde{a}_1! \alpha_1.P_1 \dots r_n[k_n] : \nu \tilde{a}_n! \alpha_n.P_n$ is:

$$\mathcal{E}_\pi \stackrel{\text{def}}{=} \sum_{A \in \tilde{A}} \begin{cases} 1 & \text{if } \exists i \text{ s.t. } \mathcal{E}^A(\alpha_i.P_i) = 1 \\ 0 & \text{otherwise} \end{cases}$$
- a replicator $r[k] : \nu \tilde{a}! \alpha.P$ is $\mathcal{E}_r \stackrel{\text{def}}{=} \sum_{a \in \tilde{a}} \mathcal{E}^a(\alpha.P)$

$$\text{where } \begin{cases} \mathcal{E}^\sigma(\alpha) \stackrel{\text{def}}{=} 1 & \text{if } \alpha = \bar{c}(\sigma), 0 & \text{otherwise} \\ \mathcal{E}^\sigma(\alpha.P) \stackrel{\text{def}}{=} 1 & \text{if } \mathcal{E}^\sigma(\alpha) = 1, \mathcal{E}^\sigma(P) & \text{otherwise} \\ \mathcal{E}^\sigma(\phi P) \stackrel{\text{def}}{=} \mathcal{E}^\sigma(P), \mathcal{E}^\sigma(P_1 + P_2) \stackrel{\text{def}}{=} \max(\mathcal{E}^\sigma(P_1), \mathcal{E}^\sigma(P_2)) \end{cases}$$

Note that each restriction νA can escape *at most once* and thus the escape bound of the pi-graph cannot exceed the number of restrictions. Similarly, a (local) private name can escape at most once per spawned thread.

Lemma 2. *Let π be a pi-graph with thread-bounds b_r ($r \in R^\pi$) and $\text{Its}(\pi) = \langle Q, T \rangle$. Then, for all $q \in Q$ we have $|\text{dyn}(q) \cap \mathcal{N}_o| \leq \mathcal{E}_\pi + \sum_{r \in R^\pi} b_r(\mathcal{E}_r + \mathcal{V}_r)$.*

Proof. Beyond the (global and local) escape bounds, the situation appears as more subtle at first sight than for Lemma 1. By Definition 6 an output name $o!$ can exist in the partition γ even if it is not in the codomain of β . This is because the name can serve as a discriminant for an input name $i?$. However, there is at most one output name $o!$ equated to each input name $i?$ in γ , so that those extra cases are limited by the number of variables (not already used for output names), hence the formula derives from a similar proof scheme \square

As an illustration our bounds predict that Example 1 generates at most $4k+3$ input names (exactly $k+1$) or shared names (exactly 0), and $4k+4$ output names (exactly $k+1$). The following is a direct consequence of the existence of bounds for the dynamic resources and the fact that we always choose the minimum available identity when generating a dynamic name (cf.[4]).

Lemma 3. *A thread-bounded system is finite-state.*

Finally, this obviously impacts bisimilarity as follows.

Theorem 3. *Bisimilarity is decidable for thread-bounded pi-graphs.*

Beyond decidability, we may remark that bisimilarity checking can exploit the following properties: (1) only (a subset of) the states and transition labels must be inspected (cf. Def. 11), and (2) for control-finite systems, the dynamic resources can be precomputed according to Lemmas 1 and 2.

6 Related and future works

In [3] we introduced a dynamic variant of the pi-graphs based on general graph rewriting techniques, similarly to [6]. The variants discussed in [4, 5] introduced *structural invariance* as in the present paper. Instead of replicators, we used previously an explicit parallel operator and a model of iterated processes. The introduction of replicators has several advantages. First they increase the expressivity of the language, previously restricted to control-finite systems. Moreover, they yield much more compact encodings since multiple threads can share common subgraphs. In [5] we show that pi-graphs represent a good intermediate language when translating pi-calculi into Petri-nets. The adaptation to the replicator-based variant should derive easily, although this is left as a future work. Finally the calculus presented in this paper adopts a standard (although constrained) syntax, which makes it a “real” pi-calculus unlike previous variants.

In [8] a verification framework is proposed for a pi-calculus with general recursion but no match nor mismatch. The translation produces Place/Transition Petri nets with sub-processes (so-called *fragments*) as places and either *reactions* between reachable fragments or communications through public channels as transitions. The translated nets produce the *reduction semantics* of the processes. In comparison we do not perform a semantic but a syntactic (graphical) translation of the pi-calculus, clearly much more compact. Most importantly, we characterize the *transition semantics* of the pi-graphs. Indeed, without the match (and mismatch) and considering the reduction semantics only, then most of the non-trivial components of the semantics disappear: the dynamics resources, hence the causal clock and the partition as well as all the commitment rules except [silent] and [com]. Interesting classes of finitely characterizable infinite systems are proposed in [8]. A further study of such classes in the setting of pi-graphs seems quite appealing. Another semantic translation of pi-calculi is proposed in the *history dependent automata* (HDA) framework [9]. The transitions of HDA provide injective correspondences between names, which ensures locally the freshness of the generated names. In the pi-graphs, the freshness property is enforced at the global level by the use of the causal clock. One advantage of the global approach is the possibility to implement non-trivial phenomena such as read-write causality [7]. We think the pi-graphs also enjoy a simpler notion of bisimilarity than HD-automata, in particular there is no renaming concern and the detection of inactive names is purely local (cf. Definition 6).

Our current work is the development of a verification framework for pi-graphs. Stateful bisimilarity as introduced in this paper seems amenable to the usual bisimulation checking techniques, using a mix of state-based and transition-based algorithms (cf.[10]). For control-finite systems Lemmas 1 and 2 can be exploited for precomputations (e.g. “perfect” bitstate hashing), which may thus yield particularly efficient algorithms. Another line of work is the reinterpretation of the graph relabelling semantics such that the LTS states would be plain processes and not graphs. This would offer a simpler setting to demonstrate our conjecture that the pi-graph semantics presented in this paper is fully *compositional*. As a witness, consider the emblematic example of [11]. The pi-calculus

processes $[x = b]\bar{b}\langle b \rangle$ and 0 are both early and late bisimilar, although in a context with x equal to b they are discriminated. In the pi-graphs, the problem does not occur since $[x = b]\bar{b}\langle b \rangle$ produces a transition and is thus *not* (stateful) bisimilar to the deadlocked process.

References

1. Milner, R.: Communicating and Mobile Systems: The π -Calculus. Cambridge University Press (1999)
2. Sangiorgi, D., Walker, D.: The π -calculus: a Theory of Mobile Processes. Cambridge University Press (2001)
3. Peschanski, F., Bialkiewicz, J.A.: Modelling and verifying mobile systems using pi-graphs. In: SOFSEM. Volume 5404 of LNCS., Springer (2009) 437–448
4. Peschanski, F., Klaudel, H., Devillers, R.: A decidable characterization of a graphical pi-calculus with iterators. In: Infinity. Volume 39 of EPTCS. (2010) 47–61
5. Peschanski, F., Klaudel, H., Devillers, R.: A Petri net interpretation of open reconfigurable systems. In: Petri Nets. Volume to appear of LNCS. (2011)
6. Gadducci, F.: Graph rewriting for the π -calculus. Mathematical Structures in Computer Science **17** (2007) 407–437
7. Degano, P., Priami, C.: Causality for mobile processes. In: ICALP. Volume 944 of LNCS., Springer (1995) 660–671
8. Meyer, R., Gorrieri, R.: On the relationship between π -calculus and finite place/transition Petri nets. In: CONCUR. Volume 5710 of LNCS., Springer (2009) 463–480
9. Montanari, U., Pistore, M.: History-dependent automata: An introduction. In: SFM. Volume 3465 of LNCS., Springer (2005) 1–28
10. Baier, C., Katoen, J.P.: Chapter 7. Equivalence and Abstraction. In: Principles of Model Checking. MIT Press (2008)
11. Sangiorgi, D.: A theory of bisimulation for the pi-calculus. Acta Inf. **33** (1996) 69–97