Controlling resource access in Directed Bigraphs

Davide Grohmann[1], **Marino Miculan**[2]

[1] grohmann@dimi.uniud.it, [2] miculan@dimi.uniud.it
Department of Mathematics and Computer Science, University of Udine, Italy

Abstract: We study *directed bigraph with negative ports*, a bigraphical framework for representing models for distributed, concurrent and ubiquitous computing. With respect to previous versions, we add the possibility that components may *govern* the access to resources, like (web) servers control requests from clients. This framework encompasses many common computational aspects, such as name or channel creation, references, client/server connections, localities, etc, still allowing to derive systematically labelled transition systems whose bisimilarities are congruences.

In order to illustrate the expressivity of this framework, we give the encodings of client/server communications through firewalls, of (compositional) Petri nets and of chemical reactions.

Keywords: Bigraphs, reactive systems, Petri nets, graph-based approaches to service-oriented applications.

1 Introduction

Bigraphical reactive systems (BRSs) are an emerging graphical framework proposed by Milner and others [Mil01, Mil06] as a unifying theory of process models for distributed, concurrent and ubiquitous computing. A bigraphical reactive system consists of a category of *bigraphs* (usually generated over a given *signature of controls*) and a set of *reaction rules*. Bigraphs can be seen as representations of the possible configurations of the system, and the reaction rules specify how these configuration can evolve, i.e., the reaction relation between bigraphs. Often, bigraphs represent terms up-to structural congruence and reaction rules represent term rewrite rules.

Many process calculi have successfully represented as bigraphical reactive systems: λ-calculus [Mil07], CCS [Mil06], π-calculus [BS06, JM04], Mobile Ambients [Jen08], Homer [BH06], Fusion [GM07c], Petri nets [LM06], and context-aware systems [BDE$^+$06]. The advantage of using bigraphical reactive systems is that they provide powerful general results for deriving a labelled transition system *automatically* from the reaction rules, via the so-called *IPO construction*. Notably, the bisimulation on this transition system is always a congruence; thus, bigraphical reactive systems provide general tools for compositional reasoning about concurrent, distributed systems.

Bigraphs are the key structures supporting these results. A bigraph is a set of nodes (the *controls*), endowed with two independent graph structures, the *place graph* and the *link graph* (Figure 1). The place graph is a tree over the nodes, representing the spatial arrangement (i.e., nesting) of the various components of the system. The link graph represents the communication connections between the components, possibly traversing the place structure. A bigraph may be "not ground", in the sense that it may have one or more "holes", or *sites* (the gray boxes) to

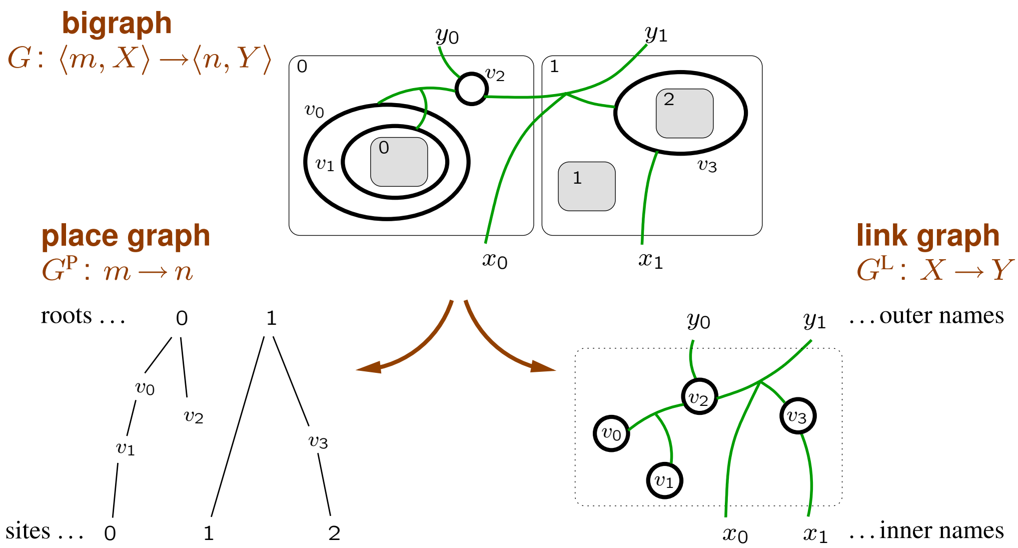

Figure 1: Example of pure bigraph (from [Mil06]).

be instantiated; these holes are specific leaves of the place graph, where other bigraphs can be grafted, respecting the connection links. This operation gives rise to a notion of composition between bigraphs, and hence to a categorical structure.

In Milner's "pure bigraphs" [Mil06], connections are represented by hyper-arcs between nodes (Figure 1). This model has been successfully used to represent many calculi, such as CCS, and (with a small variant) λ-calculus, π-calculus. Nevertheless, other calculi, such as Fusion [PV98], seem to escape this framework. Aiming to a more expressive framework, in previous work [GM07b, GM07c], we have introduced *directed bigraphs*. Pure and directed bigraphs differ only on the link structure: in the directed variant, we distinguish "edges" from "connections". Intuitively, edges represent *(delocalized) resources*, or *knowledge tokens*, which can be *accessed* by controls. Arcs are arrows from ports of controls to edges (possibly through names on the interfaces of bigraphs); moreover, in the version considered in the present paper, we allow arcs to point to other control's ports (Figure 2). Outward ports on a control represent the capability of the control to access to (external) resources; insted, inward ports represent the capability of the control to "stop" or "govern" other node's requests. The presence of both kinds of capabilities is common in distributed scenarios, such as client/server communications, firewalls, web services etc; for instance a system may ask to access to some data, but this attempt may be blocked, checked and possibly redirected by a guarding mechanism. Moreover, controls with inward ports can represent *localized* resources, that is, resources with a position within the place hierarchy; this cannot be represented easily by edges, which do not appear in the place graph.

Notably, these extended have RPO and IPO constructions, there is a notion of normal form, and a sound a complete axiomatization can be given. Therefore, these bigraphs can be conveniently used for building wide reaction systems from which we can synthesize labelled transition systems via the IPO construction, and whose bisimilarity is still a congruence.

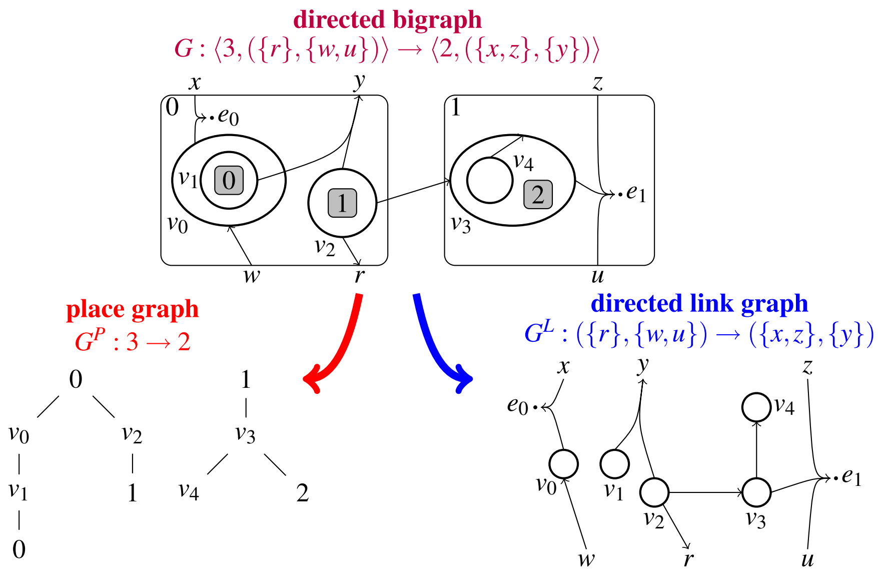

Figure 2: An example of directed bigraph, with negative ports.

Due to lack of space, in this paper we can only skim over these theoretical results; we prefer to focus on some important applications of this framework. In Section 2 we give the basic definitions about directed bigraphs. In Sections 3 we present the elementary bigraphs, which are enough to generate all possible bigraphs. Section 4 is devoted to example applications, highlighting the expressive power of this framework: we show how distributed services and protocols can be represented, by describing a three-tier architecture with a firewall; we will present an encoding of Petri nets, and finally we apply this framework to the representation of chemical reactions. Conclusions and direction for future work are in Section 6. Constructions of RPOs and IPOs, a notion of normal form and a complete axiomatization appear in [GM08].

2 Directed bigraphs over polarized signatures

In this section we introduce directed bigraphs, with inward (“negative”) ports on controls, extending [GM07b]. Following previous developments about pure and directed bigraphs, we work in *supported monoidal precategories*; we refer to [JM04, §3] for an introduction.

A *polarized signature* is a signature of controls, which may have two kind of ports: *negative* and *positive*. Let $\mathcal{K}$ be a polarized signature; we denote with $ar^n, ar^p : \mathcal{K} \to \mathbb{N}$ the arity functions of the negative and positive ports, respectively. Thus, for $k \in \mathcal{K}$, the arity function is $ar(k) \triangleq (ar^n(k), ar^p(k))$. A control k is *positive* if it has only positive ports (i.e., $ar^n(k) = 0$); it is *negative* if it has only negative ports (i.e., $ar^p(k) = 0$).

Definition 1 A *polarized interface* X is a pair of sets of names $X = (X^-, X^+)$; the two components are called *downward* and *upward* interfaces, respectively.

A *directed link graph* $A : X \to Y$ is $A = (V, E, ctrl, link)$ where X, Y are the *inner* and *outer* interfaces, V is the set of *nodes*, E is the set of *edges*, $ctrl : V \to \mathcal{K}$ is the *control map*, and $link : \mathsf{Pnt}(A) \to \mathsf{Lnk}(A)$ is the *link map*, where *ports*, *points* and *links* of A are defined as follows:

$$\mathsf{Prt}^n(A) \triangleq \sum_{v \in V} ar^n(ctrl(v)) \quad \mathsf{Prt}^p(A) \triangleq \sum_{v \in V} ar^p(ctrl(v)) \quad \mathsf{Prt}(A) \triangleq \mathsf{Prt}^n(A) \cup \mathsf{Prt}^p(A)$$

$$\mathsf{Pnt}(A) \triangleq X^+ \uplus Y^- \uplus \mathsf{Prt}^p(A) \qquad \mathsf{Lnk}(A) \triangleq X^- \uplus Y^+ \uplus \mathsf{Prt}^n(A) \uplus E$$

The link map cannot connect downward and upward names of the same interface, i.e., the following condition must hold: $(link(X^+) \cap X^-) \cup (link(Y^-) \cap Y^+) = \emptyset$; moreover the link map cannot connect positive and negative ports of the same node.

Directed link graphs are graphically depicted much like ordinary link graphs, with the difference that edges are explicit objects, and not hyper-arcs connecting points and names; points and names are associated to links (that is edges or negative ports) or other names by (simple, non hyper) directed arcs. An example are given in Figure 2. This notation aims to make explicit the “resource request flow”: positive ports and names in the interfaces can be associated either to internal or to external resources. In the first case, positive ports and names are connected to an edge or a negative port; these names are “inward” because they offer to the context the access to an internal resource. In the second case, the positive ports and names are connected to an “outward” name, which is waiting to be plugged by the context into a resource.

In the following, by “signature”, “interface” and “link graphs” we will intend “polarized signature”, “polarized interface” and “directed link graphs” respectively, unless otherwise noted.

Definition 2 The precategory of *directed link graphs* has polarized interfaces as objects, and directed link graphs as morphisms.

Given two directed link graphs $A_i = (V_i, E_i, ctrl_i, link_i) : X_i \to X_{i+1}$ $(i = 0, 1)$, the composition $A_1 \circ A_0 : X_0 \to X_2$ is defined when the two link graphs have disjoint nodes and edges. In this case, $A_1 \circ A_0 \triangleq (V, E, ctrl, link)$, where $V \triangleq V_0 \uplus V_1$, $ctrl \triangleq ctrl_0 \uplus ctrl_1$, $E \triangleq E_0 \uplus E_1$ and

$$link : X_0^+ \uplus X_2^- \uplus \mathsf{Prt}^p(A_0) \uplus \mathsf{Prt}^p(A_1) \to X_0^- \uplus X_2^+ \uplus E \uplus \mathsf{Prt}^n(A_0) \uplus \mathsf{Prt}^n(A_1)$$

is defined as follows:

$$link(p) \triangleq \begin{cases} link_0(p) & \text{if } p \in X_0^+ \uplus \mathsf{Prt}^p(A_0) \text{ and } link_0(p) \in X_0^- \uplus E_0 \uplus \mathsf{Prt}^n(A_0) \\ link_1(x) & \text{if } p \in X_0^+ \uplus \mathsf{Prt}^p(A_0) \text{ and } link_0(p) = x \in X_1^+ \\ link_1(p) & \text{if } p \in X_2^- \uplus \mathsf{Prt}^p(A_1) \text{ and } link_1(p) \in X_2^+ \uplus E_1 \uplus \mathsf{Prt}^n(A_1) \\ link_0(x) & \text{if } p \in X_2^- \uplus \mathsf{Prt}^p(A_1) \text{ and } link_1(p) = x \in X_1^-. \end{cases}$$

The identity link graph of X is $id_X \triangleq (\emptyset, \emptyset, \emptyset_{\mathcal{K}}, id_{X^-} \cup id_{X^+}) : X \to X$.

It is easy to check that composition is associative, and that given a link graph $A : X \to Y$, the compositions $A \circ id_X$ and $id_Y \circ A$ are defined and equal to A.

Definition 1 forbids connections between names of the same interface in order to avoid undefined link maps after compositions. Similarly, links between ports on the same node are forbidden, because these graphs cannot be obtained by composing an “unlinked” node and a context.

It is easy to see that the precategory ′DLG is self-dual, that is ′DLG $\cong$ ′DLG^{op}.

The notions of openness, closeness, leanness, etc. defined in [GM07b] can be easily extended to the new framework, considering negative ports as a new kind of resources. Moreover, the definition of tensor product can be derived extending to negative ports the one given in [GM07b],

Finally, we can define the *(extended) directed bigraphs* as the composition of standard place graphs (see [JM04, §7] for definitions) and directed link graphs.

Definition 3 A *directed bigraph* with signature $\mathcal{K}$ is $G = (V, E, ctrl, prnt, link) : I \to J$, where $I = \langle m, X \rangle$ and $J = \langle n, Y \rangle$ are its inner and outer interfaces respectively. An interface is composed by a *width* (a finite ordinal) and by a pair of finite sets of names. V and E are the sets of nodes and edges respectively, and $prnt$, $ctrl$ and $link$ are the parent, control and link maps, such that $G^P \triangleq (V, ctrl, prnt) : m \to n$ is a place graph and $G^L \triangleq (V, E, ctrl, link) : X \to Y$ is a directed link graph.

We denote G as combination of G^P and G^L by $G = \langle G^P, G^L \rangle$. In this notation, a place graph and a (directed) link graph can be put together iff they have the same sets of nodes.

Definition 4 The precategory ′DBIG of directed bigraph with signature $\mathcal{K}$ has interfaces $I = \langle m, X \rangle$ as objects and directed bigraphs $G = \langle G^P, G^L \rangle : I \to J$ as morphisms. If $H : J \to K$ is another directed bigraph with sets of nodes and edges disjoint from the respectively ones of G, then their composition is defined by composing their components, i.e.:

$$H \circ G \triangleq \langle H^P \circ G^P, H^L \circ G^L \rangle : I \to K.$$

The identity directed bigraph of $I = \langle m, X \rangle$ is $\langle id_m, id_X \rangle : I \to I$.

Analogously, the tensor product of two bigraphs can be defined tensoring their components.

It is easy to check that for every signature $\mathcal{K}$, the precategory ′DBIG is wide monoidal; the origin is $\varepsilon = \langle 0, (\emptyset, \emptyset) \rangle$ and the interface $\langle n, X \rangle$ has width n. Hence, ′DBIG can be used for applying the theory of *wide reaction systems* and *wide transition systems* as developed by Jensen and Milner; [JM04, §4, §5]. To this end, we need to show that ′DBIG has RPOs and IPOs. Since place graphs are as usual, it suffices to show that directed link graphs have RPOs and IPOs.

Theorem 1 *If a pair $\vec{A}$ of link graphs has a bound $\vec{D}$, there exists an RPO $(\vec{B}, B)$ for $\vec{A}$ to $\vec{D}$.*

As a consequence, ′DLG has IPOs too. See [GM08] for the constructions for RPOs and IPOs in directed bigraphs with negative ports, extending the construction given in [GM07b].

Actually, often we do not want to distinguish bigraphs differing only on the identity of nodes and edges. To this end, we introduce the category DBIG of *abstract directed bigraphs*, which is constructed from ′DBIG forgetting the identity of nodes and edges and any idle edge. More precisely, abstract bigraphs are bigraphs taken up-to an equivalence $\risingdotseq$ (see [JM04] for details).

Definition 5 Two concrete directed bigraphs G and H are *lean-support equivalent*, written $G \risingdotseq H$, if there exists an iso between their nodes and edges sets after removing any idle edges.

The category DBIG of abstract directed bigraphs has the same objects as ′DBIG, and its arrows are lean-support equivalence classes of directed bigraphs.

3 Algebra and Axiomatization

As for directed bigraphs, also in the case of polarized signature it is possible to give a sound and complete axiomatization. In this section, due to lack of space, we describe only the main classes of bigraphs and the elementary bigraphs which can generate all bigraphs according to a well-defined normal form. Due to lack of space, the definition of normal form and the normalization theorem is given in [GM08]. We refer the reader to [GM07a] for a complete presentation of the notation used here.

First, we introduce two distinct and complementary subclasses of bigraphs: *wirings* and *discrete bigraphs*. that are strongly used in defining the normal form and the axiomatization.

Definition 6 A *wiring* is a bigraph whose interfaces have zero width (and hence has no nodes). The wirings ω are generated by the composition or tensor product of three elements: substitutions $\sigma : (\emptyset, X^+) \to (\emptyset, Y^+)$, fusions $\delta : (Y^-, \emptyset) \to (X^-, \emptyset)$, and closures $\rotatebox[origin=c]{180}{$\curlyvee$}^x_y : (\emptyset, y) \to (x, \emptyset)$.

Definition 7 An interface is *prime* if its width is 1. Often we abbreviate a prime interface $\langle 1, (X^-, X^+)\rangle$ with $\langle (X^-, X^+)\rangle$, in particular $1 = \langle(\emptyset, \emptyset)\rangle$. A prime bigraph $P : \langle m, (Y^-, Y^+)\rangle \to \langle (X^-, X^+)\rangle$ has a prime outer interface and the names in Y^+, X^- are linked to negative ports of P.

An important prime bigraph is $merge_m : m \to 1$, it has no nodes and maps m sites to one root.

Definition 8 A bigraph is *discrete* if it has no edges and every open link has exactly one point.

The discreteness is well-behaved, and preserved by composition and tensor. It is easy to see that discrete bigraphs form a monoidal sub-precategory of $'$DBIG.

Definition 9 Let K be any non atomic control with arity (k^-, k^+), let $\vec{x}^-, \vec{x}^+$ be two sequences of distinct names, and let $\vec{Y}^+, \vec{Y}^-$ be two sequences of (possibly empty) sets of distinct names, such that: $|\vec{x}^-| + |\vec{x}^+| = k^+$ and $|\vec{Y}^-| = |\vec{Y}^+| = k^-$. For a K-node v, we define the discrete *ion*

$$K(v,l) : \langle(\vec{x}^-, \vec{Y}^+)\rangle \to \langle(\vec{Y}^-, \vec{x}^+)\rangle$$

as the bigraph with exactly a node v and l is a pair of maps: an iso map $l^p : \vec{x}^- \cup \vec{x}^+ \to \mathsf{Prt}^p(v)$ describing the linking among positive ports and names in $\vec{x}^-$ or $\vec{x}^+$, and another iso map $l^n : \vec{Y}^- \cup \vec{Y}^+ \to \mathsf{Prt}^n(v)$ describing the linking among negative ports and sets of upward inner names (in $\vec{Y}^+$) and sets of downward outer names (in $\vec{Y}^-$). We omit v when it can be understood.

For a prime discrete bigraph P with outer names in (Z^-, Z^+), we define a discrete *molecule* as:

$$(K(l) \otimes id_{(Z^- \setminus \vec{x}^-, Z^+ \setminus \vec{Y}^+)}) \circ P.$$

If K is atomic, we define the discrete *atom*, as an ion without sites:

$$K(l) : (\vec{x}^-, \vec{Y}^+) \to \langle(\vec{Y}^-, \vec{x}^+)\rangle.$$

An arbitrary (non-discrete) ion, molecule or atom is formed by the composition of $\omega \otimes id_1$ with a discrete one. Often we omit $\ldots \otimes id_I$ in the compositions, when there is no ambiguity.

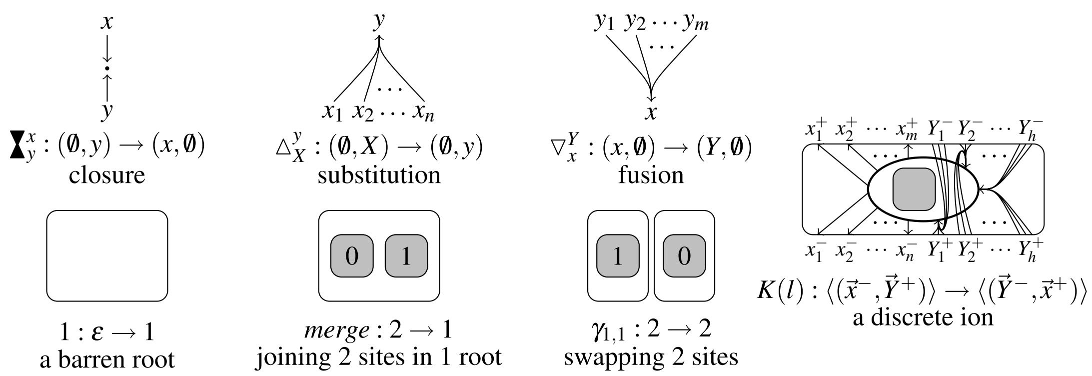

Figure 3: Elementary Bigraphs over polarized signatures.

Figure 3 shows the *algebraic signature*, that is a set of elementary bigraphs able to define any other bigraph using composition and tensor product. The various *sharing products* are the intuitive generalization of the ones defined in [GM07a]; see [GM08] for a detailed description.

4 Applications

4.1 Three-tier interaction with access control

As mentioned before, directed bigraphs over polarized signatures allow to represent resource access control, by means of negative ports. This is particularly useful for representing access policies between systems, possibly in different locations; the edges can represent access tokens (or keys), which are global (although known possibly to only some controls). An example and quite common scenario is a client-server connection, where the access to the server is subject to authentication; after the request has been accepted, the server can route it to a back-end service (e.g., a DBMS); see Figure 4. The security policy is implemented by the firewall control, which allows a query to reach the server only if the client knows the correct key (rule AUTH). The server routes the query to the correct back-end service using rules like ROUTE; finally the back-end service provides the data (rule GET). An example computation is shown in Figure 5.

4.2 Compositional Petri Nets

In this section we recall briefly what a Petri net is and we give an encoding of these nets as directed bigraphs; to this end it is preferable to work with sorted links, as in [LM06]. Notice that this encoding yields naturally a notion of composition between Petri nets.

Definition 10 A *place transition net* (*P/T net*) is a 5-tuple (P,T,F,M_i) $(P \cap T = \emptyset)$, where:

- P is the set of *places;* T is the set of *transitions;*

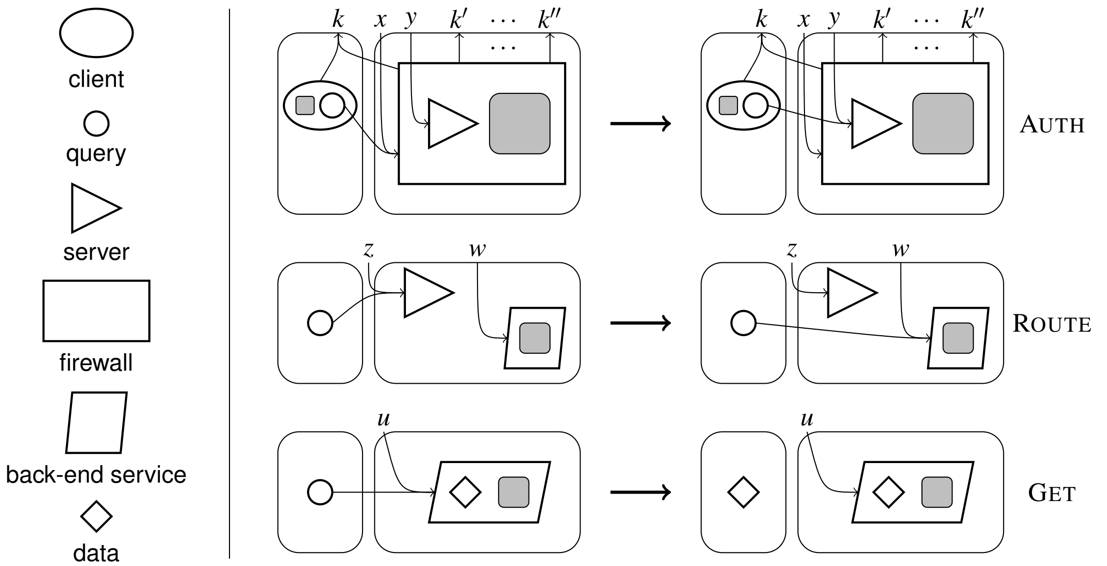

Figure 4: Signatures and rules for three-tier architecture services through a firewall.

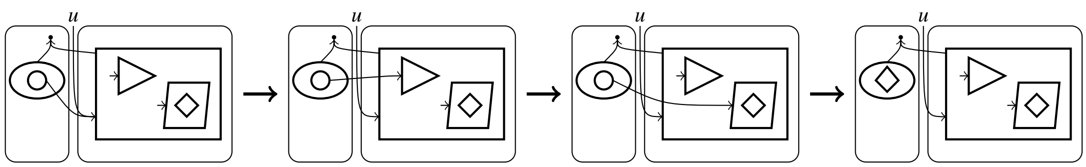

Figure 5: An example of client-server interaction through a firewall.

- F is the multiset of *arcs,* linking places to transitions and vice versa: $F \triangleq \langle (P \times T) \cup (T \times P), f : (P \times T) \cup (T \times P) \to \mathbb{N} \rangle$, with the constrain $\forall t \in T.\ \exists p,q \in P.\ (p,t) \in F \wedge (t,q) \in F$;
- $M : P \to \mathbb{N}$ is a *marking,* giving to each place a number of tokens, a place p is *marked* by M if $M(p) > 0$ and *unmarked* if $M(p) = 0$; M_i is the *initial marking.*

We define ${}^\bullet t \triangleq \{p \mid (p,t) \in F\}$ to be the *pre-multiset* of the transition t, and $t^\bullet \triangleq \{p \mid (t,p) \in F\}$ the *post-multiset* of the transition t.

A transition t is *enabled* by a marking M if M marks every place in ${}^\bullet t$; a transition *fires* from a marking M to a marking M', written $M \xrightarrow{t} M'$, iff for all $p \in P : M'(p) = M(p) - \sharp({}^\bullet p) + \sharp(p^\bullet)$, where $\sharp({}^\bullet p)$ and $\sharp(p^\bullet)$ are the number of occurrences of p in ${}^\bullet t$, $t^\bullet$, respectively.

Notice that we allow multiple connections between a place and a transition, that is analogous to assign a weight to an arc representing the token that have to be consumed to fire the reaction.

Definition 11 Let $N = (P,T,F,M)$ and $N' = (P',T',F',M')$ be two P/T nets, we say that N and N' are *isomorphic*, if there exist two bijections $\alpha : P \to P'$ and $\beta : T \to T'$, such that:

- $(p,t) \in F$ iff $(\alpha(p), \beta(t)) \in F'$;

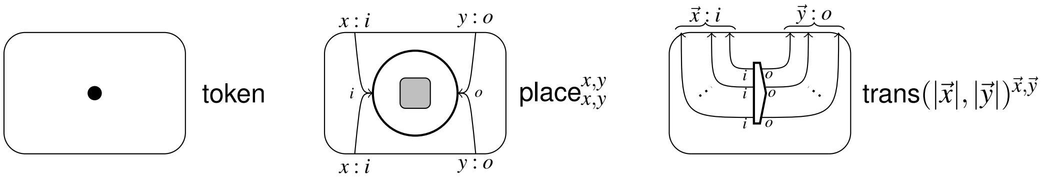

Figure 6: Signature for the encoding of compositional Petri nets.

- $(t,p) \in F$ iff $(\beta(t),\alpha(p)) \in F'$;
- $M = M' \circ \alpha$.

We recall, as defined in [LM06], the definition of link sorting.

Definition 12 A *link sorting* is a triple $\Sigma = (\Theta,\mathcal{K},\Phi)$, where Φ is a set of sorts, and $\mathcal{K}$ is a sorted signature (that is, a signature enriched with a sort to ports of each control). Furthermore, each name in the interface (X^-,X^+) is given a sort, so the interfaces take the form $(\{x_1^- : \theta_1^-,\dots,x_n^- : \theta_n^-\},\{x_1^+ : \theta_1^+,\dots,x_m^+ : \theta_m^+\})$. Finally, Φ is a rule on such enriched bigraphs, that is preserved by identities, composition and tensor product.

We denote the precategory and category of, respectively, concrete and abstract Σ-sorted directed bigraphs with $'\mathrm{DB}\textsc{ig}(\Sigma)$ and $\mathrm{DB}\textsc{ig}(\Sigma)$.

Definition 13 A *positive-negative sorting* $\Sigma = (\Theta,\mathcal{K},\Phi)$ has sorts: $\Theta = \{\theta_1,\dots,\theta_n\}$. The signature $\mathcal{K}$ assigns sorts to ports arbitrarily. The unique Φ rule is: a point and a link (except of edges) can be connected if they are equally sorted.

In order to define an encoding for compositional Petri nets, we introduce a positive-negative sorting Σ_{petri}, having sort $\Theta_{petri} \triangleq \{i,o\}$ and sorted signature:

$$\mathcal{K}_{petri} \triangleq \{\mathsf{token} : (0,0),\ \mathsf{place} : (\{1:i,1:o\},0),\ \mathsf{trans}(h,k) : (0,\{h:i,k:o\})\} \qquad \text{where } h,k>0$$

where the controls token and trans are both atomic, while the control place is passive. Finally, the Φ rule ensures that the linking is allowed only among ports having the same sort. An example of use of this sorted signature is shown in Figure 6. The encoding function $[\![\cdot]\!]$ is defined as follows:

$$[\![(P,T,F,M)]\!] = merge_{(|P|+|T|)} \circ \Big(id_{|P|} \curlyvee id_{(P\times\{i,o\},\emptyset)} \curlyvee \Big(\curlyvee_{t\in T} trans(|{}^\bullet t|,\ |t^\bullet|)_{({}^\bullet t\times\{i\},\ t^\bullet\times\{o\})} \Big)\Big) \circ \Big(\sum_{p\in P} place^{(p,i),\ (p,o)}_{(p,i),\ (p,o)} \circ (merge_{(|M(p)|+1)} \circ (\triangle^{\{(p,i),\ (p,o)\}} \otimes (\sum_{i=0}^{M(p)} token) \otimes 1))\Big).$$

where, with an abuse of notation, $trans(|{}^\bullet t|,\ |t^\bullet|)_{({}^\bullet t\times\{i\},\ t^\bullet\times\{o\})}$ means that if in the multisets there are some repetitions of places then the ports of $trans$ are linked to the same downward inner name (i.e., (p,i) or (p,o)), an alternative definition is to link every port of $trans$ to a different downward inner name and then (eventually) "equate" these names using fusions.

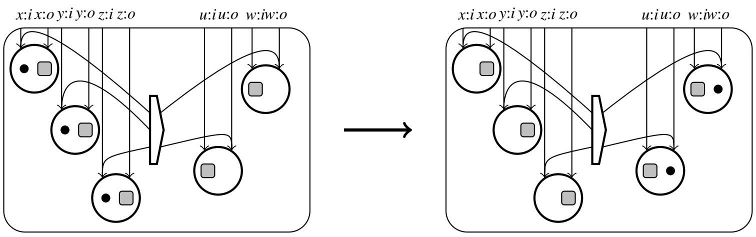

Figure 7: Example of reaction rule in the case of 3 input and 2 output places.

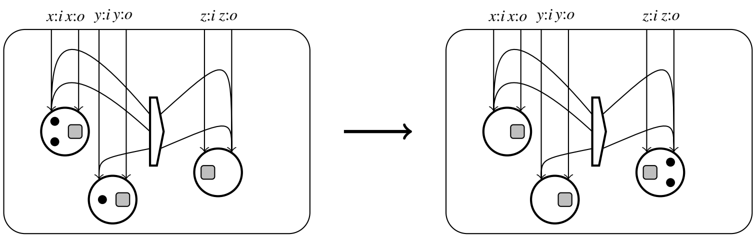

Figure 8: Example of reaction rule in the case of 2 input and 1 output places (with multiple arcs).

Proposition 1 *Let N, N' be two P/T nets, N is isomorphic to N' iff $[\![N]\!] = [\![N']\!]$ up to iso.*

We have a different reaction rule for any pair (h,k) associated to the control trans, in Figure 7 we show the reaction rule for the pair $(3,2)$, that is a transition having 3 inputs and 2 outputs. Moreover, we allow multiple connections between places and transitions, as in Figure 8, and we can have transitions using some places as inputs and outputs, see Figure 9.

Now we can show that the given translation is adequate.

Theorem 2 *Let (P,T,F,M_i) be a P/T net, $M \xrightarrow{t} M'$ iff $[\![(P,T,F,M)]\!] \longrightarrow [\![(P,T,F,M')]\!]$.*

Proof. ($\Rightarrow$) Suppose $M \xrightarrow{t} M'$, so M enable the transition t, then there exists a trans-node in $[\![(P,T,F,M)]\!]$ encoding the transition t, and the corresponding place-node of ${}^\bullet t$ contain the necessary tokens to fire the transition (by translation of M), then we can apply the appropriate rule to perform the reaction reaching the configuration $[\![(P,T,F,M')]\!]$.

($\Leftarrow$) If $[\![(P,T,F,M)]\!] \longrightarrow [\![(P,T,F,M')]\!]$, there exists a matching of a rule with a sub-bigraph of $[\![(P,T,F,M)]\!]$, in particular the matched nodes have a counter part into the P/T net (P,T,F,M), so the marking M enables a transition t (corresponding to the trans-node), and then $M \xrightarrow{t} M'$. □

An interesting future work is to study the bisimulation induced by the IPO LTS over these compositional Petri nets. We remark however, that this notion of composition is different from that in Open Petri nets, since in the latter the interfaces express also behavioural properties, while in the bigraphical encoding the interfaces express resource requests and offerings.

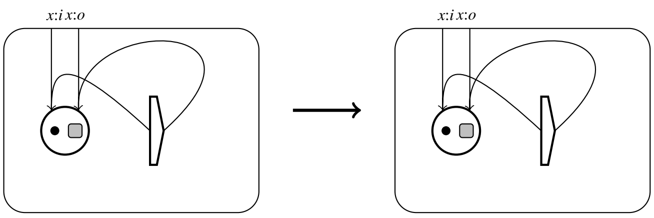

Figure 9: Example of reaction rule in the case of 1 place used as input and output.

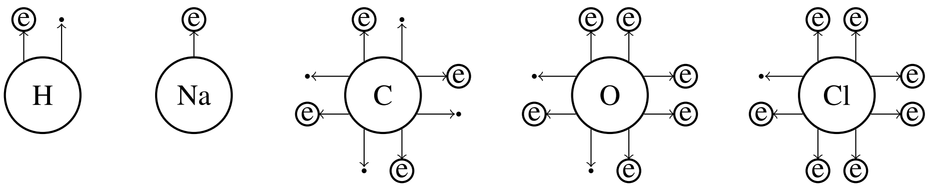

Figure 10: Example of atom encodings in directed link graphs.

5 Chemical Reactions

A chemical reactions is a process describing the conversions of a chemical compositions. Always, the chemical changes caused by a reaction involve the motion of electrons in the forming or breaking of *chemical bonds*. For example, the *octet rule* says that atoms tend to gain, lose or share electrons so as to have eight electrons in their outer electron shell.

In this section, we give an encoding of atoms into directed link graphs, as shown in Figure 10, inspired by the well-known *Lewis structures*. We describe the atoms as nodes, and those nodes have a number of positive ports equal to the number of valence electrons. Each of these ports are linked to an electron, represented as a node having a negative port (accepting incoming connections; for sake of simplicity we identify the node representing the electron with its port, that is, we do not force all incoming connections to be linked to a precise point of the node). Moreover, some nodes can have extra ports, that are initially linked to edges, hydrogen and oxygen can be two examples, the idea is that such a configuration describes the aim of the atom to "capture" electrons to complete its external shell; e.g. an oxygen atom has two missing electrons, so it tries to share these two electrons with a pair of hydrogen atoms forming the water molecule.

We apply this model describing the forming and breaking of bonds among atoms, here we deal with *strong bonds*, that is *covalent* and *ionic bonds*.

Some examples of covalent bonds are shown in Figure 11, the first shows how to hydrogen atoms can share their electron. The second one is well-know and describes the generation of a water molecule from two hydrogen atoms and an oxygen one: the oxygen shares two electrons: one with each hydrogen, in this way it gets the two missing electrons in its external orbit, conversely each hydrogen atom completes its orbit sharing an electron with the oxygen. The latter describes a more complicate situation, where the two carbon atoms (each needing four electron

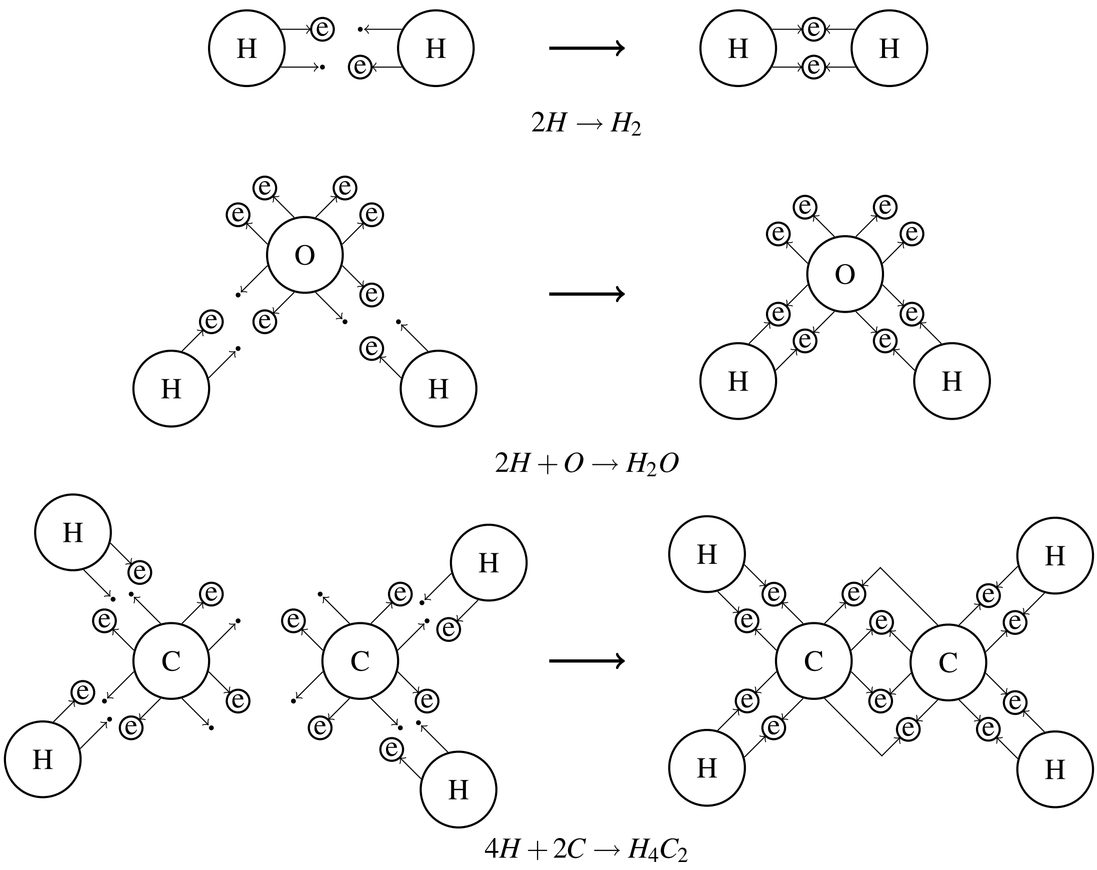

Figure 11: Example of covalent bonds among atoms.

to complete its orbit) share two electron with the other carbon atom, and the remaining two missing electrons are provided by a pair of hydrogen.

In Figure 12, we show an example of ionic bond: given an atom of sodium and a chlorine one, it may happen (by octet rule) that the external electron of sodium is lost by the atom and "captured" by the chlorine, forming a sodium (positive) ion and a chlorine (negative) ion. These two ions attract each other by the electrostatic force caused by the electron exchange. Finally the ions can be composed to form sodium-chloride molecule, that is the common salt.

An interesting future work concern to represent the *weak bonds*, i.e. *hydrogen bonds* and *van der Waals bonds*, using the same representation as much as possible.

6 Conclusions

In this paper, we have considered directed bigraphs over *polarized signatures*, a bigraphical model for concurrent, distributed system with resources and controls. The main difference with previous versions of bigraphs is the capability of nodes (i.e., systems) to ask for resource access

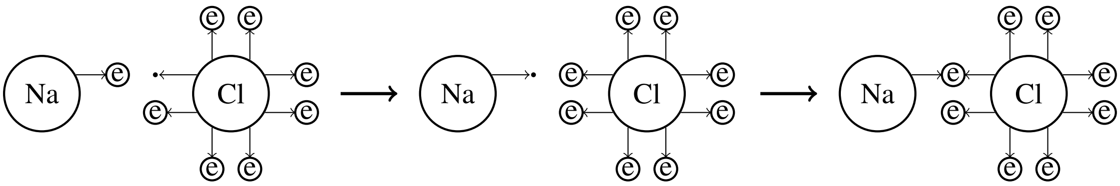

$$Na + Cl \rightarrow Na^{+} + Cl^{-} \rightarrow NaCl$$

Figure 12: Examples of ion bonds among atoms.

(via the "positive ports") and to control other's requests, providing access to own resources (via the negative ports). These bigraphs have RPO and IPO constructions, thus allowing to derive systematically labelled transition systems from reactive systems, as in [JM03, GM07c]; notably the bisimilarities induced by these labelled transition systems are always congruences. These directed bigraphs admit also a notion of normal form, and a complete axiomatization.

We have exhibited the expressive power of this framework, by applying it some interesting cases: a three-tier interaction between client, server and back-end service through a firewall, the Petri nets, and chemical reactions. All these cases are faithfully encoded as directed bigraphs with polarized signatures (possibly with sorting).

An interesting future work is to develop properly the treatment of web service interactions, extending the ideas shown in Section 4.1. In particular, we would like to give a bigraphical semantics of some formal calculus for web services, such as SCC or CC-Pi [BBC+06, BM07].

Another future development is to use this kind of bigraphs as a general framework for systems biology. Some preliminary experiment about the representation of biochemical reactions, not shown in this paper due to lack of space, are promising: ions, electrons, chemical links can be represented as controls and arcs, and the place structure can be fruitfully used to represent nesting of chemical compounds. It would be interesting to encode in directed bigraphs some important formalism for systems biology, such as the κ-calculus [DL04]. Along this line, also the possibility of adding quantitative aspects (i.e., reaction rates) sounds very promising.

Bibliography

[BBC+06] M. Boreale, R. Bruni, L. Caires, R. D. Nicola, I. Lanese, M. Loreti, F. Martins, U. Montanari, A. Ravara, D. Sangiorgi, V. T. Vasconcelos, G. Zavattaro. SCC: A Service Centered Calculus. In Bravetti et al. (eds.), *Proc. WS-FM*. Lecture Notes in Computer Science 4184, pp. 38–57. Springer, 2006.

[BDE+06] L. Birkedal, S. Debois, E. Elsborg, T. Hildebrandt, H. Niss. Bigraphical Models of Context-Aware Systems. In Aceto and Ingólfsdóttir (eds.), *Proc. FoSSaCS*. Lecture Notes in Computer Science 3921, pp. 187–201. Springer, 2006.

[BH06] M. Bundgaard, T. T. Hildebrandt. Bigraphical Semantics of Higher-Order Mobile Embedded Resources with Local Names. *Electr. Notes Theor. Comput. Sci.* 154(2):7–29, 2006.

[BM07] M. G. Buscemi, U. Montanari. CC-Pi: A Constraint-Based Language for Specifying Service Level Agreements. In Nicola (ed.), *Proc. ESOP*. Lecture Notes in Computer Science 4421, pp. 18–32. Springer, 2007.

[BS06] M. Bundgaard, V. Sassone. Typed polyadic pi-calculus in bigraphs. In Bossi and Maher (eds.), *Proc. PPDP*. Pp. 1–12. ACM, 2006.

[DL04] V. Danos, C. Laneve. Formal molecular biology. *Theoretical Computer Science* 325, 2004.

[GM07a] D. Grohmann, M. Miculan. An Algebra for Directed Bigraphs. In Mackie and Plump (eds.), *Pre-proceedings of TERMGRAPH 2007*. Electronic Notes in Theoretical Computer Science. Elsevier, 2007.

[GM07b] D. Grohmann, M. Miculan. Directed bigraphs. In *Proc. XXIII MFPS*. Electronic Notes in Theoretical Computer Science 173, pp. 121–137. Elsevier, 2007.

[GM07c] D. Grohmann, M. Miculan. Reactive Systems over Directed Bigraphs. In Caires and Vasconcelos (eds.), *Proc. CONCUR 2007*. Lecture Notes in Computer Science 4703, pp. 380–394. Springer-Verlag, 2007.

[GM08] D. Grohmann, M. Miculan. Controlling resource access in Directed Bigraphs. Technical report, Department of Mathematics and Computer Science, University of Udine, 2008. Available at http://www.dimi.uniud.it/miculan/Papers/.

[Jen08] O. H. Jensen. *Mobile Processes in Bigraphs*. PhD thesis, University of Aalborg, 2008. To appear.

[JM03] O. H. Jensen, R. Milner. Bigraphs and transitions. In *Proc. POPL*. Pp. 38–49. 2003.

[JM04] O. H. Jensen, R. Milner. Bigraphs and mobile processes (revised). Technical report UCAM-CL-TR-580, Computer Laboratory, University of Cambridge, 2004.

[LM06] J. J. Leifer, R. Milner. Transition systems, link graphs and Petri nets. *Mathematical Structures in Computer Science* 16(6):989–1047, 2006.

[Mil01] R. Milner. Bigraphical Reactive Systems. In Larsen and Nielsen (eds.), *Proc. 12th CONCUR*. Lecture Notes in Computer Science 2154, pp. 16–35. Springer, 2001.

[Mil06] R. Milner. Pure bigraphs: Structure and dynamics. *Information and Computation* 204(1):60–122, 2006.

[Mil07] R. Milner. Local Bigraphs and Confluence: Two Conjectures. In *Proc. EXPRESS 2006*. Electronic Notes in Theoretical Computer Science 175(3), pp. 65–73. Elsevier, 2007.

[PV98] J. Parrow, B. Victor. The Fusion Calculus: Expressiveness and Symmetry in Mobile Processes. In *Proceedings of LICS '98*. Pp. 176–185. Computer Society Press, July 1998.
http://www.docs.uu.se/~victor/tr/fusion.shtml