

Bigraphical Models of Context-Aware Systems

L. Birkedal, S. Debois, E. Elsborg, T. Hildebrandt, and H. Niss

IT University of Copenhagen (ITU)
{birkedal, debois, elsborg, hilde, hniss}@itu.dk

Abstract. As part of ongoing work on evaluating Milner’s bigraphical reactive systems, we investigate bigraphical models of *context-aware systems*, a facet of ubiquitous computing. We find that naively encoding such systems in bigraphs is somewhat awkward; and we propose a more sophisticated modeling technique, introducing *Plato-graphical models*, alleviating this awkwardness. We argue that such models are useful for simulation and point out that for reasoning about such bigraphical models, the bisimilarity inherent to bigraphical reactive systems is not enough in itself; an equivalence between the bigraphical reactive systems themselves is also needed.

1 Introduction

The theory of *bigraphical reactive systems*, due to Milner and co-workers, is based on a graphical model of mobile computation that emphasizes both locality and connectivity [15, 19, 21]. A bigraph comprises a place graph, representing locations of computational nodes, and a link graph, representing interconnection of these nodes. We give dynamics to bigraphs by defining reaction rules that rewrite bigraphs to bigraphs; roughly, a bigraphical reactive system (BRS) is a set of such rules. Based on methods of the seminal [16], any BRS has a labelled transition system, the behavioural equivalence (bisimilarity) of which is a congruence.

There are two principal aims for the theory of bigraphical reactive systems: (1) to model ubiquitous systems [28], capturing mobile locality in the place graph and mobile connectivity in the link graph; and (2) to be a meta-theory encompassing existing calculi for concurrency and mobility. To date, the theory has been evaluated only wrt. the second aim: We have bigraphical understanding of Petri nets [18], π -calculus [13, 15, 14], CCS [21], mobile ambients [13], HOMER [5], and λ -calculus [19, 20].

The present paper initiates the evaluation of the first aim. We investigate modeling of *context-aware systems*, a vital aspect of ubiquitous systems. A context-aware application is an application that adapts its behaviour depending on the context at hand [26], interpreting “context” to mean the situation in which the computation takes place [10]. The canonical example of such a situation is the location of the device performing the computation; systems sensitive to location are called *location-aware*. As an example, a location-aware printing system could

send a user’s print job to a printer close by. (For notions of context different from location, refer to [27]; for large-scale practical examples, see [1].)

To observe changes in the context, context-aware systems typically include a separate context sensing component that maintains a model of the current context. Such models are known as context models [12] or, more specifically, location models [2]. The above-mentioned location-aware printing system would need to maintain a model of the context that supports finding the printer closest to a given device. Such models are informal. There are only very few formal models of context-aware computing (refer to [11] for an overview). We point out Context Unity [25]; in spirit, our proposal is somewhat closer to process calculi than Context Unity is. However, bigraphs differ from traditional process calculi in that we get to write our own reaction rules.

In overall terms, our contribution is two-fold.

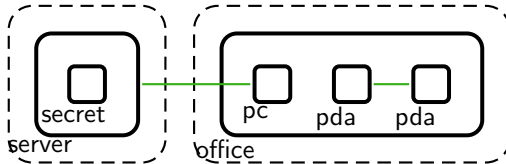
- We find, perhaps surprisingly, that naively modeling context-aware systems as BRSs is somewhat awkward; and
- we propose a more sophisticated modeling technique, in which the perceived and actual context are both explicitly represented as distinct but overlapping BRSs. We call such models *Plato-graphical*.

The remainder of this paper is organized as follows. In Section 2, we introduce bigraphs and bigraphical reactive systems. In Section 3, we discuss naive bigraphical models of location-aware systems. In Section 4, we introduce our Plato-graphical models of context-aware systems. In Section 5, we present two example models. In Section 6, we discuss. Finally, in Section 7, we conclude and note future work.

2 Bigraphs and Bigraphical Reactive Systems

We introduce bigraphs by example (the reader can find the relevant formal definitions of [15, 21] in Appendix A of [3]). Readers acquainted with bigraphs may skip this section.

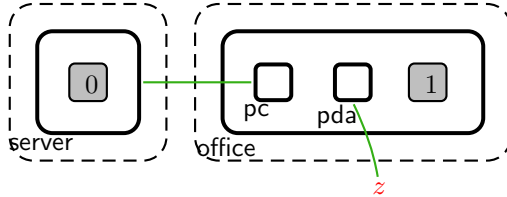
Here is a bigraph, *A*:



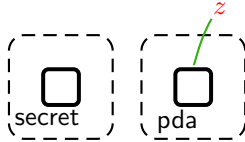
It has *nodes* (vertices), indicated by solid boxes. Each node has a *control*, written in sans serif. Each control has a number of *ports*; ports can be linked by *edges*, indicated by lines. Here, the controls *secret* and *office* have no ports, all other controls have one port. Nodes can be nested, indicated by containment. The two outermost dashed boxes indicate *roots*. Roots have no controls; they serve solely to separate different nesting hierarchies.

The bigraph A ostensibly models two physically separate locations (because of the two roots). The first contains a server, which in turn contains secret data; the second contains an office, which in turn contains a PC and two PDAs. The server and the PC are connected, as are the PDAs.

Here is another bigraph, B :



B resembles A , except that the content of `server` has been replaced with a *site* $-_0$, one of the `pdas` has been replaced by a site $-_1$, and there is an *inner name* z connected to the remaining `pda`. Using sites and names, we can define composition of bigraphs. For that, here is yet another bigraph C :



C has an *outer name* z . The bigraphs B and C compose to form A , i.e., $A = B \circ C$. Composition proceeds by plugging the roots of C into the sites of B (in order), and fusing together the connections `pda` \rightarrow z (in C) and $z \rightarrow$ `pda` (in B) removing the name z in the process.

One cannot compose arbitrary bigraphs. For $U \circ V$ to be defined, U must have exactly as many sites as V has roots, and the inner names of U must be precisely the outer names of V . The sites and inner names are collectively called the *inner face*; similarly, the roots and outer names are called the *outer face*. A has inner face $\langle 0, \emptyset \rangle$ (no holes, no inner names) and outer face $\langle 2, \emptyset \rangle$ (two roots, no outer names). We write $A : \langle 0, \emptyset \rangle \rightarrow \langle 2, \emptyset \rangle$. Similarly, $B : \langle 2, \{z\} \rangle \rightarrow \langle 2, \emptyset \rangle$ and $C : \langle 0, \emptyset \rangle \rightarrow \langle 2, \{z\} \rangle$.

The graphical representation used above is handy for modeling, but unwieldy for reasoning. Fortunately, bigraphs have an associated term language [7, 17], which we use (albeit in a sugared form) in what follows. The language is summarized in Table 1. Here are, in order of increasing complexity, term representations of the bigraphs A , B and C .

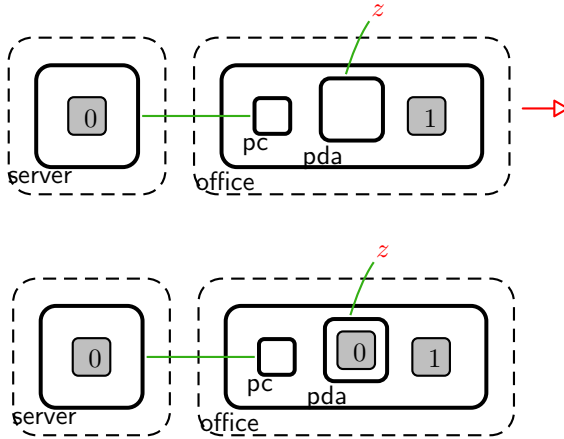
$$\begin{aligned}
 C &= \text{secret} \parallel \text{pda}_z \\
 A &= /x./y.\text{server}_x(\text{secret}) \parallel \text{office}(\text{pc}_x \mid \text{pda}_y \mid \text{pda}_y) \\
 B &= /x./y.\text{server}_x(-_0) \parallel \text{office}(\text{pc}_x \mid \text{pda}_y \mid -_1) \mid y/z
 \end{aligned}$$

Notice how, in B , edges are specified by first linking nodes to the same name, then converting that name to an edge using the closure $'/'$.

Table 1. Sugared term language for bigraphs

Term	Meaning
$U \parallel V$	Concatenation (juxtaposition) of roots.
$U V$	Concatenation (juxtaposition) of children. (collect the children of U and V under one root.)
$U \circ V$	Composition.
$U(V)$	Nesting. U contains V .
$K_{\vec{x}}(U)$	Ion. Node with control K of arity $ \vec{x} $, ports connected to the outer names of vector \vec{x} . The node contains U .
1	The <i>barren</i> (empty) root.
$-_i$	Site numbered i .
$/x.U$	U with outer name x replaced by an edge.
x/y	Connection from inner name y to outer name x .

We give dynamics to bigraphs by defining reaction rules. Example:



$$\begin{aligned}
 & /x.server_x(-_0) \parallel office(pc_x | pda_z | -_1) \\
 & \longrightarrow /x.server_x(-_0) \parallel office(pc_x | pda_z(-_0) | -_1)
 \end{aligned}$$

This rule might model that if a PC in some office is linked to a server, a PDA in the same office may use the PC as a gateway to copy data from the server. The rule matches the bigraph A above, taking $secret$ to the site $-_0$ and pda_y to the site $-_1$, rewriting A to

$$A' = /x./y.server_x(secret) \parallel office(pc_x | pda_y(secret) | pda_y)$$

(We omit details on what it means to match connections; refer to one of [15, 21].)

It is occasionally convenient to limit the contexts in which a reaction rule applies [4], i.e., we might want to limit the above example reaction rule to apply only in the left wing of the building. To this end, bigraphs can be equipped

with a *sorting* [13, 21, 18]. A sorting consists of a set of *sorts* (or types); all inner and outer faces are then enriched with such a sort. Further, a sorting must stipulate some condition on bigraphs, we then restrict our attention to the bigraphs that satisfy that condition, thus outlawing some contexts. Obviously, removing contexts may ruin the congruence property of the induced bisimilarity; [13] and [21] give different sufficient conditions for a sorting to preserve that congruence property.

This concludes our informal overview of bigraphs. Now on to the models.

3 Naive Models of Location-Aware Systems

In this section, we attempt to model location-aware systems naively in bigraphs. We will find the naive approach to be somewhat awkward. Due to space constraints we do not discuss other forms of context.

We use the place and link graphs for describing locations and interconnections directly, and we use reaction rules to implement both *reconfiguration* of the context and *queries* on the context. The former is simply a non-deterministic change in the context; the latter is a computation on the context that does not change the context, except for producing an answer to some question. In a location-aware system, a device moving would be a reconfiguration, whereas computing the answer to the question “what devices are currently at the location l ” is a query.

We discuss the implementation of this query. (An implementation of the query can be found in Appendix B in [3].) Incidentally, a query such as “find nearest neighbor”, which conceptually is only slightly harder, is significantly harder to implement. (Other examples plagued by essentially the same difficulties can be found in [9].)

Consider the following bigraph representing devices (e.g., PDAs) located at locations (e.g., offices, meeting rooms) within a building.

$$l = /w./x./y./z.\text{loc}(\text{loc}(\text{loc}(\text{dev}_w) \mid \text{loc}(\text{dev}_x \mid \text{dev}_y))) \mid \text{loc}() \mid \text{loc}(\text{dev}_z)$$

Off-hand, finding all devices, say, beneath the root, looks straightforward: We should simply recursively traverse the nesting tree. Unfortunately, such traversal is quite complicated for the following reasons.

- The bigraphical reaction rules do not support recursion directly, so we must encode a runtime stack by means of additional controls.
- Bigraphical reaction rules can be applied in *any* context, but when implementing an operation such as the query we consider now, we need more refined control over when rules can be applied; one may achieve this more refined control by again using additional nodes and controls, essentially implementing what corresponds to a program counter. This still leaves great difficulty in handling concurrent operations, though.
- As a special case of the previous item, it is particularly difficult to express that a reaction rule is intended to apply only in case something is *not* present in the context.

Summing up, the bigraphical rules that model physical action do not in general provide the power to compute directly with a model of that action (because of a lack of control structures). The slogan is “reconfiguring is easy, querying is hard”.

In earlier work on evaluating bigraphs as a meta-theory (aim (2) mentioned in the Introduction), reaction rules were used to encode the operational semantics of a calculus or programming language. However, above we attempt to implement a query *directly* as reaction rules. This seemingly innocuous difference will turn out to have major implications for reasoning methods; more on this in Section 6.

We imagine that adding more flexibility to the reaction rules might make it easier to program directly with bigraphs. One possible attempt is to use spatial logics for bigraphs [6] in combination with sorting, to get control of the contexts in which a particular reaction rule applies.

In the following sections, we propose another way to model context-aware systems in bigraphs, where the reaction rules are not used to program directly with but instead they are used (1) to represent transitions happening in the real world and (2) to encode operational semantics of programming languages, within which one can then implement queries on representations of the real world.

4 Plato-Graphical Models of Context-Aware Systems

The naive model of the previous section shares an important characteristic with recent proposals of formal models for context-aware computation [4, 8, 25] that comprise agents and contexts only: These models take the agent’s ability to determine *what is* the present context as given. We contend that for some systems, it is natural to model not only the actual context but also the agent’s representation of the actual context. We shall see that pursuing this idea will partially alleviate the awkwardness seen in the previous Section.

We shall need some notation and definitions.

Notation 1. We write $\mathbf{B} = (\mathcal{K}, \mathcal{R})$ to indicate that \mathbf{B} is a bigraphical reactive system with controls \mathcal{K} and rules \mathcal{R} , and write $f \in \mathbf{B}$ to mean that f is a bigraph of \mathbf{B} .

Definition 1 (Independence). Let $\mathbf{B} = (\mathcal{K}, \mathcal{R})$ and $\mathbf{B}' = (\mathcal{K}', \mathcal{R}')$ be bigraphical reactive systems. Say that \mathbf{B} and \mathbf{B}' are *independent* and write $\mathbf{B} \perp \mathbf{B}'$ iff \mathcal{K} and \mathcal{K}' are disjoint.

Definition 2 (Composite bigraphical reactive systems). Let $\mathbf{B} = (\mathcal{K}, \mathcal{R})$ and $\mathbf{B}' = (\mathcal{K}', \mathcal{R}')$ be bigraphical reactive systems. Define the union $\mathbf{B} \cup \mathbf{B}'$ point-wise, i.e., $\mathbf{B} \cup \mathbf{B}' = (\mathcal{K} \cup \mathcal{K}', \mathcal{R} \cup \mathcal{R}')$, when \mathcal{K} and \mathcal{K}' agree on the arities of the controls in $\mathcal{K} \cap \mathcal{K}'$.

Be aware that there are two ways of taking the union of two sets of *parametrized* reaction rules: (1) merge the rules and then ground them, or (2) first ground the rules and then merge them. In general, the resulting rule set of (1) is a superset of the rule set of (2). We use approach (1).

We propose a model of context-aware computing that comprises three bigraphical reactive systems: the context \mathbf{C} ; its representation or proxy \mathbf{P} ; and the computational agents \mathbf{A} . Drawing on classical work [23], specifically The Allegory of the Cave, we call such a model *Plato-graphical*.

Definition 3 (Plato-graphical model). A *Plato-graphical model* is a triple $(\mathbf{C}, \mathbf{P}, \mathbf{A})$ of bigraphical reactive systems, such that $\mathcal{M} = \mathbf{C} \cup \mathbf{P} \cup \mathbf{A}$ is itself a bigraphical reactive system and $\mathbf{C} \perp \mathbf{A}$. A *state* of the model is a bigraph of \mathcal{M} on the form $/\vec{x}.(C \parallel P \parallel A)$, where $C \in \mathbf{C}$, $P \in \mathbf{P}$, $A \in \mathbf{A}$, and \vec{x} is some vector of names.

We emphasize the intended difference between \mathbf{C} and \mathbf{P} : Whereas an element of \mathbf{C} models a possible context, an element of \mathbf{P} models *a model* of a possible context. The independence condition ensures that agents can only directly observe or manipulate the proxy; not the context itself. (In the parlance of [25], the independence condition ensures separability.) To query or alter the context, agents must use the proxy as a sensor and actuator.

Using bigraphs as our basic formalism gives us two things. First, we can write our own reaction rules. We claim that because of this ability, models become remarkably straightforward and intuitive; hopefully, the reader will agree after seeing our example models in the next section. Second, we automatically get a bisimilarity that is a congruence. Thus, bisimilarity of agents is a very fine equivalence: No state of the context and proxy can distinguish bisimilar agents.

Proposition 1. *Let \sim denote the bisimilarity in \mathcal{M} , and let $A, A' \in \mathbf{A}$ with $A \sim A'$. For any $C \in \mathbf{C}$, $P \in \mathbf{P}$, and \vec{x} , we have $/\vec{x}.(C \parallel P \parallel A) \sim / \vec{x}.(C \parallel P \parallel A')$.*

To get a less discriminating equivalence we can consider agents under a particular state of the context, or a particular state of the system.

Definition 4. Let \sim denote the bisimilarity in \mathcal{M} , and let $A, A' \in \mathbf{A}$, $C \in \mathbf{C}$ and $P \in \mathbf{P}$. We say A and A' are *equivalent w.r.t. P* iff $P \parallel A \sim P \parallel A'$, and we say A and A' are *equivalent w.r.t. C, P* iff $C \parallel P \parallel A \sim C \parallel P \parallel A'$.

We conjecture that the above forms of derived equivalences will prove useful for reasoning about a given Plato-graphical system.

Working within the Plato-graphical model, we are free to emphasize any of its three components, perhaps modeling \mathbf{P} in great detail, but keeping \mathbf{C} and \mathbf{A} abstract.

Definition 3 above does not impose any restriction on composition of states. For example, assume that we have a Plato-graphical model $\mathcal{M} = (\mathbf{C}, \mathbf{P}, \mathbf{A})$, that c , p and a are controls of \mathbf{C} , \mathbf{P} and \mathbf{A} , respectively, and that p is *not* a control of \mathbf{C} . Then the bigraphs

$$F = c(-_0 \mid -_1) \parallel p \parallel a(-_2) \quad \text{and} \quad G = c \parallel p \parallel a$$

are both states of \mathcal{M} , but their composite $F \circ G = c(c \mid p) \parallel p \parallel a(a)$ is not a state of \mathcal{M} . This example implies that bisimilarity of states of a Plato-graphical system may be too fine a relation: Conceivably, when comparing two states s

and s' , we may wish to take into account only contexts C such that $C \circ s$ and $C \circ s'$ are themselves states, i.e., we might want to outlaw F as a possible context for G . We can achieve this finer control using place-sorting. So, we define a place-sorted Plato-graphical model. The intuition behind our sorting is that we want to keep controls of \mathbf{C} , \mathbf{P} and \mathbf{A} separate when composing contexts of form $\mathbf{C} \parallel \mathbf{P} \parallel \mathbf{A}$.

Notation 2. Denote by $S_{i \leq m}$ a vector m_0, \dots, m_{n-1} of sorts. We will write $S_{i \leq m}$ for a sorted interface $\langle m, X, S_{i \leq m} \rangle$ when we do not care about names.

Definition 5 (Sorted Plato-graphical model). Let $\mathcal{M} = \mathbf{C} \cup \mathbf{P} \cup \mathbf{A}$ be a Plato-graphical model with $\mathbf{C} = (\mathcal{K}_{\mathbf{C}}, \mathcal{R}_{\mathbf{C}})$, $\mathbf{P} = (\mathcal{K}_{\mathbf{P}}, \mathcal{R}_{\mathbf{P}})$ and $\mathbf{A} = (\mathcal{K}_{\mathbf{A}}, \mathcal{R}_{\mathbf{A}})$. Define a sorting discipline on \mathcal{M} by taking sorts $\Theta = \{\mathcal{K}_{\mathbf{C}}, \mathcal{K}_{\mathbf{P}}, \mathcal{K}_{\mathbf{A}}\}$ and, for primes, sorting condition $\Phi(f : S_{i \leq n} \rightarrow S) = \text{ctrl}(f) \subseteq S \wedge \forall i \leq n. S_i = S$, lifting to an arbitrary bigraph f' by decomposing f into primes $f' = f_0 \dots f_{n-1}$ and declaring f' well-sorted iff all the f_i are. Let ϕ be an assignment of Θ -sorts to the rules of $\mathcal{R}_{\mathbf{C}}$, $\mathcal{R}_{\mathbf{P}}$, and $\mathcal{R}_{\mathbf{A}}$, such that every rule is well-sorted under Φ . Define \mathcal{M}' to be \mathcal{M} sorted by (Θ, Φ) (using ϕ to lift the reaction rules). In this case, we call \mathcal{M}' a *sorted Plato-graphical model*, and define the states of \mathcal{M}' to be the well-sorted bigraphs with outer face $\mathcal{K}_{\mathbf{C}}, \mathcal{K}_{\mathbf{P}}, \mathcal{K}_{\mathbf{A}}$.

The condition Φ essentially requires that (1) the controls of a prime (bigraph) are elements of the sort of its outer face, and (2) the sort of the outer face is exactly the sort of each of the sites. Under this sorting discipline and new definition of state, if G is assigned a sort such that it is a state, then F cannot be assigned a sort that makes it composable with G .

Is the bisimilarity in the sorted system \mathcal{M}' a congruence? The sorting discipline of \mathcal{M}' is in general not homomorphic in the sense of Milner [21, Definition 10.4]: we cannot give a sort to controls in $\mathcal{K}_{\mathbf{C}} \cap \mathcal{K}_{\mathbf{P}}$. (If \mathbf{C} , \mathbf{P} and \mathbf{A} are pairwise independent, the sorting is homomorphic; however, such a model is pathologic.) Neither is the sorting safe in the sense of Jensen [13, Definition 4.30]; condition (4) cannot be met. Counterexample: Suppose $f : \mathcal{K}_{\mathbf{C}} \rightarrow \mathcal{K}_{\mathbf{C}}$ is well-sorted; take $g = f \otimes 1 : \mathcal{K}_{\mathbf{C}} \rightarrow \mathcal{K}_{\mathbf{C}}, \mathcal{K}_{\mathbf{A}}$ (recall that $1 : \epsilon \rightarrow \langle 1, \emptyset \rangle$ denotes the barren root). Clearly, $\mathcal{U}(f) = (-_0 \mid -_1) \circ \mathcal{U}(f \otimes 1)$. However, if $\mathcal{K}_{\mathbf{C}} \neq \mathcal{K}_{\mathbf{A}}$ then $(-_0 \mid -_1) : \mathcal{K}_{\mathbf{C}}, \mathcal{K}_{\mathbf{A}} \rightarrow \mathcal{K}_{\mathbf{C}}$ is not well-sorted.

Nevertheless, the sorting of Definition 5 does give rise to a bisimilarity that is a congruence; we prove so in Appendix C in [3].

5 Examples

5.1 A Simple Context-Aware Printing System

We model the simple context-aware printing system of [4]. An office-building contains both modern PCL-5e compatible printers and old-fashioned raw-printers. Occasionally, the IT-staff at the building removes or replaces either type of printers. Each printer can process only one job; queueing is done by a central print server. The print server dispatches jobs to raw-printers only if it knows no PCL-printers; if there are PCL-printers, but they are all busy, the job will

simply have to wait. This system is context-aware: The type and number of printers physically available determine the meaning of the action “to print”. We give a model \mathbf{B} of this system in Figure 1. Looking at the controls of \mathbf{B} , it is straightforward to verify that \mathbf{B} is Plato-graphical.

Proposition 2. *The model \mathbf{B} of Figure 1 is Plato-graphical.*

We take a detailed look at the model. A state of the context \mathbf{C} consists of nested physical locations loc , within which printers prt are placed. We distinguish between PCL- and raw-printers by putting a token pcl and raw within them, respectively. Each printer has a single port, intended to link the printer to the proxy. Here is a state of the context with a PCL-printer and a raw-printer at adjacent locations; the PCL-printer is idle whereas the raw-printer is busy.

$$C = \text{loc}(\text{loc}(\text{prt}_x(\text{raw} \mid \text{dat}_z)) \mid \text{loc}(/y.\text{prt}_y(\text{pcl})))$$

Setting C in parallel with some proxy P will allow P access to the raw printer through the shared link x , but not to the PCL-printer, because it is in a closed link. The dynamics of \mathbf{C} allow printers to appear (1, 2), disappear (3), and finish printing (4).

A state of the proxy \mathbf{P} consists of a pool of pending jobs jobs and two tables of printers prts ; one contains a token raw , the other a token pcl , indicating what type of printer the table lists. The prts is a table in the sense that its only port is linked to all the printers in the context that the table knows about. Here is an example state of the proxy which knows one raw-printer, knows no PCL-printers and has two pending jobs.

$$P = \text{prts}_x(\text{raw}) \mid /y.\text{prts}_y(\text{pcl}) \mid \text{jobs}(/z.\text{doc}_z \mid /z'.\text{doc}_{z'})$$

Setting C and P above in parallel by \parallel , and closing the link x , we get a system $/x.C \parallel P$, where the table $\text{prts}_x(\text{raw})$ and the physical printer $\text{prt}_x(\text{raw} \mid \text{dat})$ are linked. The dynamics of \mathbf{P} state that if there is a job and a known, idle PCL-printer, the proxy may activate this printer (5); that if there is a job, no known PCL-printer, and an idle raw-printer, the context may activate that printer (6); and finally, that the proxy may discover a previously unknown printer (7, 8).

The dynamics of \mathbf{A} allow the agents to spontaneously spool documents (9).

Notice how the two printing rules (5) and (6) do not observe the context directly. Instead, the proxy observes the context (rules (7) and (8)) and records its observations in the tables $\text{prts}_x(\text{raw})$ and $\text{prts}_y(\text{pcl})$; the printing rules (5) and (6) then consults the tables. It is straightforward to determine whether there are no known PCL-printers: simply check if the table of PCL-printers has the form $/y.\text{prts}_y(\text{pcl})$.

As observed in Section 3 and [4], it is generally very difficult, if not impossible, to observe the *absence* of something in the context directly. An interesting but rather natural consequence of the indirect observation is that it becomes asynchronous, i.e., it is possible that a PCL-printer exists but has not yet been observed.

		Control	Activity	Arity	Comment	
Context C .	loc	active		0	Nested location	
	prt	passive		1	Physical printer	
	pcl	atomic		0	Printer-type token	
	raw	atomic		0	Printer-type token	
	dat	atomic		1	Binary data for printer	
					$\text{loc}(-_0) \longrightarrow \text{loc}(-_0 \mid /x.\text{prt}_x(\text{raw}))$	(1)
					$\text{loc}(-_0) \longrightarrow \text{loc}(-_0 \mid /x.\text{prt}_x(\text{pcl}))$	(2)
					$\text{loc}(-_0 \mid \text{prt}_x(-_1)) \longrightarrow \text{loc}(-_0) \mid x/$	(3)
					$\text{prt}_x(\text{dat}_z \mid -_0) \longrightarrow \text{prt}_x(-_0) \mid z/$	(4)
		Control	Activity	Arity	Comment	
Proxy P .	prt	passive		1	Physical printer	
	pcl	atomic		0	Printer-type token	
	raw	atomic		0	Printer-type token	
	dat	atomic		1	Binary data for printer	
	prts	passive		1	Known devices	
	jobs	passive		0	Pending documents	
	doc	atomic		1	Document	
					$\text{jobs}(\text{doc}_z \mid -_0) \parallel \text{prts}_y(\text{pcl}) \parallel \text{prt}_y(\text{pcl}) \longrightarrow$	(5)
					$\text{jobs}(-_0) \parallel \text{prts}_y(\text{pcl}) \parallel \text{prt}_y(\text{pcl} \mid \text{dat}_z)$	
					$\text{jobs}(\text{doc}_z \mid -_0) \parallel /x.\text{prts}_x(\text{pcl}) \mid \text{prts}_y(\text{raw}) \parallel \text{prt}_y(\text{raw}) \longrightarrow$	(6)
					$\text{jobs}(-_0) \parallel /x.\text{prts}_x(\text{pcl}) \mid \text{prts}_y(\text{raw}) \parallel \text{prt}_y(\text{raw} \mid \text{dat}_z)$	
					$/x.\text{prt}_x(\text{pcl}) \parallel \text{prts}_y(\text{pcl}) \longrightarrow \text{prt}_y(\text{pcl}) \parallel \text{prts}_y(\text{pcl})$	(7)
					$/x.\text{prt}_x(\text{raw}) \parallel \text{prts}_y(\text{raw}) \longrightarrow \text{prt}_y(\text{raw}) \parallel \text{prts}_y(\text{raw})$	(8)
		Control	Activity	Arity	Comment	
Agents A .	jobs	passive		0	Pending documents	
	doc	atomic		1	Document	
					$\text{jobs}(-_0) \longrightarrow \text{jobs}(-_0 \mid /z.\text{doc}_z)$	(9)

Fig. 1. Example Plato-graphical model **B**

Context C	Proxy P	Agent A
(1) : \mathcal{K}_C	(5) : $\mathcal{K}_A, \mathcal{K}_P, \mathcal{K}_C$	(9) : \mathcal{K}_A
(2) : \mathcal{K}_C	(6) : $\mathcal{K}_A, \mathcal{K}_P, \mathcal{K}_C$	
(3) : \mathcal{K}_C	(7) : $\mathcal{K}_P, \mathcal{K}_C$	
(4) : \mathcal{K}_C	(8) : $\mathcal{K}_P, \mathcal{K}_C$	

Fig. 2. Sorts for the rules of **C**, **P**, and **A**

This model **B** can be lifted to a sorted one by adding the sorts given in Figure 2; the figure assigns sorts to the outer face of both the redexes and reactums of the indicated rules. It is straightforward to verify that all of the rules are well-sorted.

Proposition 3. *The model **B** with the sorting assignment of Figure 2 is a sorted Plato-graphical model.*

5.2 A Location-Aware Printing System

Suppose we extend the printing system with location-awareness, by stipulating that a print job is not printed until the printer and the device submitting the job are co-located. To model this extended system, we introduce a new control dev for devices (PCs or PDAs) with one port and change doc to include an extra port so we can link submitted jobs to the devices submitting them. The linking is reflected in the following modified rule (9) for spooling print jobs:

$$\text{loc}(\text{dev}_x \mid -_0) \parallel \text{jobs}(-_1) \longrightarrow \text{loc}(\text{dev}_x \mid -_0) \parallel \text{jobs}(-_1 \mid /z.\text{doc}_{z,x}) \quad (9')$$

We must also modify rules (5) and (6) to insist that the device and printer are co-located. Rule (5) becomes

$$\begin{aligned} \text{jobs}(\text{doc}_{z,x} \mid -_0) \parallel \text{prts}_y(\text{pcl}) \parallel \text{loc}(\text{dev}_x \mid \text{prt}_y(\text{pcl})) \longrightarrow \\ \text{jobs}(-_0) \parallel \text{prts}_y(\text{pcl}) \parallel \text{loc}(\text{dev}_x \mid \text{prt}_y(\text{pcl} \mid \text{dat}_z)). \end{aligned} \quad (5')$$

(We suppress the new Rule (6').)

Modifying the system once again, instead of insisting that device and printer have to be actually co-located, we just require the print job to end at a printer close to the device. The print server will need to query the proxy for the printer nearest a given device. We saw in Section 3 that implementing such queries is awkward, so we will need to use the proxy. In the preceding Section, we did so directly in bigraphs; this time around, we transfer the expressive convenience of a general-purpose programming language to bigraphs for ease of implementation. We use bigraphs directly for modeling the actual context **C**, whereas we will exploit bigraphs as a meta-calculus for modeling the proxy **P**.

In detail, the whole model is $\mathbf{B} = \mathbf{C} \cup \mathbf{P} \cup \mathbf{A}$, with $\mathbf{P} = \mathbf{S} \cup \mathbf{L}$. Here **C** is intended to be a bigraphical model of the “real world”, the proxy **P** is comprised of a location sensor **S** and a location model **L** and **A** is the location-based application (the “computational agent”).

A state *C* of **C** could look like this:

$$C = \text{loc}(\text{loc}(\text{loc}(\text{loc}(\text{dev}_w) \mid \text{loc}(\text{dev}_x \mid \text{dev}_y))) \mid \text{loc} \mid \text{loc}(\text{dev}_z))$$

Changes in the real world are modeled by reaction rules that reconfigure such states. If we want to model, say, that a devices may move from one location to another, we include the reaction rule

$$\text{loc}(\text{dev}_x \mid -_0) \parallel \text{loc}(-_1) \longrightarrow \text{loc}(-_0) \parallel \text{loc}(\text{dev}_x \mid -_1). \quad (10)$$

To implement the proxy, encode as a BRS a programming language \mathcal{L} with data structures, communication primitives, and concurrency, e.g., Pict [22] or CML [24]. (We return to this assumption below.) That is, define a translation from terms of \mathcal{L} to bigraphs, and add reaction rules encoding the operational semantics of \mathcal{L} . *Then implement the location model, the sensor, and the agents in \mathcal{L} and use the encoding to transfer that model to bigraphs.* In particular, a state of the location model \mathbf{L} will have a data structure *representing* the current state of \mathbf{C} . If \mathcal{L} is an even half-way decent programming language, it should be straightforward to implement queries such as one of Section 3 or the “find closest printer” we need above.

The sensor informs the location model about changes in \mathbf{C} . We extend the above rule (10) moving a device to

$$(\text{loc}(\text{dev}_x \mid -_0) \parallel \text{loc}(-_1)) \mid S \mid L \longrightarrow (\text{loc}(-_0) \parallel \text{loc}(\text{dev}_x \mid -_1)) \mid S' \mid L, \quad (10')$$

where S' is an \mathcal{L} -encoding of “send a notification to \mathbf{L} that device x has moved”. Upon receiving the notification, \mathbf{L} updates its representation of the world. Agents of \mathbf{A} can in turn query \mathbf{L} when they need location information.

6 Discussion

We consider the following questions.

1. What languages \mathcal{L} can we encode?
2. How close are Plato-graphical models to real systems?
3. What challenges have we found for bigraphical models?
4. What uses do we envision for Plato-graphical models?
5. How do we reason about Plato-graphical models?

Ad. 1. As mentioned, there exist bigraphical encodings of various π -calculi [13, 15, 14] and of the λ -calculus [19, 20]. Using ideas of the latter encodings, we have encoded Mini-ML (call-by-value λ -calculus with pairs and lists) in local bigraphs [19]. Based on our experiences with this encoding, we find it palatable to encode CML or Pict¹.

Ad. 2. The model closely reflects how some actual location-aware systems work, for instance the one running at the ITU. Here, a sensor system (made by Ekahau) computes every two seconds the physical location of every device on the WLAN. The sensor system informs a location model about updates to locations; location-aware services then interact with the location model. In our sketched Plato-graphical model, the location model \mathbf{L} may lag behind the actual \mathbf{C} , if \mathbf{L} 's representation of \mathbf{C} does not reflect some recent reconfiguration of \mathbf{C} . But that also happens in the real system at the ITU – when a location-aware service asks the location model for the whereabouts of a device, it obtains not the position

¹ We are presently working on implementing an interpreter for bigraphical reactive systems; such an interpreter will make it easier to experiment with these and other encodings.

of the device, but the position of the device the last time the sensor checked. In the mean time, the device may have moved.

Ad. 3. When modeling the physical world, we have made use of both the place and link graphs, the place graph modeling the location hierarchy of a building. As argued in [2], DAGs or graphs are more natural models of location. Thus, systems such as the ones we have considered here suggest generalizing the place graphs part of bigraphs, or consider ways to encode DAGs or general graphs naturally as place graphs.

Ad. 4. Given an implementation of bigraphical reactive systems, one could *simulate* the behaviour of a location-aware system, and thus allow for experimentation with different designs of location-aware and context-aware systems. Likewise, one could experiment with different choices for the \mathcal{L} language of Section 5.2. Such simulation suggests further extensions of the bigraphical model: In actual context-aware systems, one is generally interested in timing aspects (e.g., the sensor samples only every two seconds), continuous space (e.g., the sensor really produces continuous data), and probabilistic models (e.g., to accurately simulate sensors and sensor failure).

Ad. 5. What about using Plato-graphical models for *formal reasoning* about context-aware systems? One use of formal models is to prove an abstract specification model equivalent to a concrete implementation model. In π -calculus, we come with π -terms i, s , one for the implementation and one for the specification. The terms i and s are themselves the models; we take (π -) bisimilarity as equivalence, so to prove i and s equivalent, we merely prove them bisimilar. We can play the same game within any BRS: Simply come up with a bigraph I (the implementation model) and a bigraph S (the specification model), and prove them bisimilar within the labelled transition system of the BRS. Because that bisimulation is a congruence, such reasoning should be tractable, e.g. with the bisimulation in Definition 4.

Unfortunately, bisimulation within a single BRS is not always enough wrt. Plato-graphical models. Suppose we want a specification model \mathcal{M} with an abstract view of the context, and an implementation model \mathcal{M}' with a detailed view of the context. We express this by having \mathcal{M} and \mathcal{M}' differ only in their context sub-BRSs, that is,

$$\mathcal{M} = \mathbf{C} \cup \mathbf{P} \cup \mathbf{A} \quad \mathcal{M}' = \mathbf{C}' \cup \mathbf{P} \cup \mathbf{A}.$$

The trouble is that because \mathbf{C} and \mathbf{C}' may have different controls and reaction rules, bisimulation between their respective labelled transition systems is meaningless! What we need is a notion of equivalence of BRSs, not just equivalence of bigraphs of a single BRS. At the time of writing, we know of no such equivalence². Thus, our investigation of bigraphical models for context-aware systems

² The reader may suggest that we just define a common language for modeling both the abstract and detailed view, and define a translation from this language into a single BRS. However, in this case we are no longer modeling a ubiquitous system directly in bigraphs (aim 1 of the Introduction), but using bigraphs as a meta-calculus (aim 2 of the Introduction).

suggests that equivalence of BRSs is a key notion currently missing. One possible direction would be to try recover from the notion of WRS-functor [16] — functors that preserve reaction rules — a notion of a BRS implementing another BRS.

7 Conclusion and Future Work

We have initiated an evaluation of the use of bigraphical reactive systems for models of context-aware computing in ubiquitous systems. We found that BRSs, in their current form, are not suitable for directly modeling context queries, but on the other hand suitable for modeling reconfigurations of the actual context.

In response, we proposed Plato-graphical models, where both state and dynamics are logically divided in three parts: the actual context, the observed context (or proxy), and the computational agents, respectively. The computational agents and the actual context are separated, and interact only through the proxy. This separation into different BRSs makes it possible to encode different parts of the system using domain-specific languages. Moreover, we have shown how the context-aware printing system of [4] can be modeled straightforwardly in the Plato-graphical model.

Further, we have argued that Plato-graphical models are useful for simulating context-aware systems, and we are currently working on an implementation of BRSs at ITU to allow such experimentation. Only through such experimentation will it be clear how useful Plato-graphical models really are. For simulation purposes it will be important to extend bigraphs with timing aspects, continuous space, and probabilities.

Finally, we have pointed out that establishing a notion of equivalence between BRSs, as opposed to bisimilarity within a BRS, is important future work.

Acknowledgments

We gratefully acknowledge discussions with the other members of the BPL group at ITU, in particular Arne Glenstrup, Troels Damgaard and Mikkel Bundgaard; and with Robin Milner. This work was funded in part by the Danish Research Agency (grant no.: 2059-03-0031) and the IT University of Copenhagen (the LaCoMoCo project).

References

1. M. Addlesee, R. Curwen, S. Hodges, J. Newman, P. Steggles, A. Ward, and A. Hopper. Implementing a sentient computing system. *IEEE Computer*, 2001.
2. C. Becker and F. Dürr. On location models for ubiquitous computing. *Personal and Ubiquitous Computing*, 9:20–31, 2005. Springer.
3. L. Birkedal, S. Debois, E. Elsborg, T. Hildebrandt, and H. Niss. Bigraphical Models of Context-aware Systems. Technical Report 74, IT Univ. of Copenhagen, 2005.

4. P. Braione and G. P. Picco. On calculi for context-aware coordination. In *Proc. of COORDINATION'04*, vol. 2949 of *LNCS*, pages 38–54. 2004.
5. M. Bundgaard and T. Hildebrandt. Bigraphical semantics of higher-order mobile embedded resources with local names. In *Proc. of GT-VC'05*, 2005.
6. G. Conforti, D. Macedonio, and V. Sassone. Spatial Logics for Bigraphs. In *Proc. of ICALP'05*, vol. 3580 of *LNCS*, pages 766–778. 2005.
7. T. C. Damgaard and L. Birkedal. Axiomatizing binding bigraphs (revised). Technical Report TR-2005-71, IT University of Copenhagen, 2005.
8. R. De Nicola, D. Gorla, and R. Pugliese. Basic observables for a calculus for global computing. In *Proc. of ICALP'05*, volume 3580 of *LNCS*, pages 1226–1238. Springer, 2005.
9. S. Debois and T. C. Damgaard. Bigraphs by Example. Technical Report TR-2005-61, IT University of Copenhagen, March 2005.
10. A. K. Dey and G. D. Abowd. Towards a better understanding of context and context-awareness. In *Workshop on The What, Who, Where, When, and How of Context-Awareness*, 2000.
11. M. Hennessy. Context-awareness: Models and analysis. Talk at 2nd UK-UbiNet Workshop, 2004.
12. K. Henriksen, J. Indulska, and A. Rakotonirainy. Modeling context information in pervasive computing systems. In *Proc. of Pervasive'02*, vol. 2414 of *LNCS*, 2002.
13. O. H. Jensen. *Mobile Processes in Bigraphs*. PhD thesis, 2005. Forthcoming.
14. O. H. Jensen and R. Milner. Bigraphs and Transitions. In *Proc. of POPL'03*, 2003.
15. O. H. Jensen and R. Milner. Bigraphs and mobile processes (revised). Technical Report UCAM-CL-TR-580, University of Cambridge, 2004.
16. J. J. Leifer and R. Milner. Deriving bisimulation congruences for reactive systems. In *Proc. of CONCUR'00*, 2000.
17. R. Milner. Axioms for bigraphical structure. Technical Report UCAM-CL-TR-581, University of Cambridge, 2004.
18. R. Milner. Bigraphs for Petri Nets. In *Lectures on Concurrency and Petri Nets: Advances in Petri Nets*, vol. 3098 of *LNCS*, 2004.
19. R. Milner. Bigraphs whose names have multiple locality. Technical Report UCAM-CL-TR-603, 2004.
20. R. Milner. Bigraphs: A tutorial. Slides, April 2005. Available at <http://www.cl.cam.ac.uk/users/rm135/bigraphs-tutorial.pdf>.
21. Robin Milner. Pure bigraphs: Structure and dynamics. To appear in *Information and Computation*, 2005.
22. B. C. Pierce and D. N. Turner. Pict: A prog. lang. based on the pi-calculus. In *Proof, Language and Interaction: Essays in Honour of R. Milner*, MIT, 2000.
23. Plato. The republic, book vii, 360 B.C. Translation by Benjamin Jowett.
24. J H. Reppy. *Concurrent Programming in ML*. Cambridge University Press, 1999.
25. G.-C. Roman, C. Julien, and J. Payton. A formal treatment of context-awareness. In *Proc. of FASE'04*, vol. 2984 of *LNCS*, 2004.
26. B. Schilit, N. Adams, and R. Want. Context-aware computing applications. In *Proc. of IEEE Workshop on Mobile Computing Systems and Applications*, 1994.
27. A. Schmidt, M. Beigl, and H.-W. Gellersen. There is more to context than location. *Computers & Graphics Journal*, 1999.
28. M. Weiser. Hot topics – ubiquitous computing. *IEEE Computer*, 1993.