

BigActors - A Model for Structure-aware Computation ^{*}

Eloi Pereira
Systems Engineering
UC Berkeley, CA, USA
eloi@berkeley.edu

Christoph M. Kirsch
Dept. of Computer Sciences
Univ. of Salzburg, Austria
ck@cs.uni-salzburg.at

Raja Sengupta
Systems Engineering
UC Berkeley, CA, USA
sengupta@ce.berkeley.edu

João Borges de Sousa
School of Engineering
Univ. of Porto, Portugal
jtasso@fe.up.pt

ABSTRACT

This paper describes a model of computation for structure-aware computing called the BigActor model. The model is a hybrid. It combines the Actor model [1] and the Bigraph model [10]. The contributions of this paper are an operational semantics, an example illustrating how the model supports the concise programming of a mobile agent working in a ubiquitous computing world, a query language enabling a bigActor to observe the world around it, and a definition giving semantics to the feedback loop in control theory in the context of this model. This is followed by three theorems showing how the operational semantics supports the programming of concurrent mobile agents in the semantics of feedback control.

Categories and Subject Descriptors

D.3.1 [Programming Languages]: Formal Definitions and Theory—*Semantics*; F.1.1 [Computation by Abstract Devices]: Models of Computation—*Relations between models, Self-modifying machines*

1. INTRODUCTION

This paper describes a model of computation for structure-aware computing. The model is a hybrid. It combines the Actor model [1] and the Bigraph model [10]. Hence the name *BigActor model*. In this model, computation is modelled using the Actor model, enriched by three new semantics rules (Figure 8) to make it structure-aware. Structure is a bigraph. The dynamics of the structure is modelled as a Bigraph Reactive System [9]. The three extra semantic rules

^{*}Research supported in part by the National Science Foundation (CNS1136141), by the Fundação para a Ciência e Tecnologia (SFRH/BD/43596/2008), by the Portuguese MoD - project PITVANT, and by the National Research Network RiSE on Rigorous Systems Engineering (Austrian Science Fund S11404-N23)

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. To copy otherwise, to republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee.

ICCPs '13, April 08 - 11 2013, Philadelphia, PA, USA
Copyright 2013 ACM 978-1-4503-1996-6/13/04 \$15.00.

provide means for an actor to observe and control a bigraph, and migrate from one host to another. Hence it is called a bigActor. When the structure models a physical world as in Example 3 of Section 4, bigActors become a cyber entities in a physical world.

The contributions of this paper are an operational semantics for the BigActor model stated in Section 4.1. BigActors are linked to a model of the structure (bigraph) by a hosting relation. Without a structure (or context as in context-aware computing), a bigActor is just an actor. But when a bigActor is hosted in a bigraph it can leverage the three new semantic rules enabling it to observe the bigraph, control it, and migrate in it. The set of observations a bigActor can make in the bigraph hosting it, is formalized as a query language. Section 5 formalizes one possible query language. The operational semantics is well defined for a class of query languages as expressed by Equation 1. Section 4.1 includes an example program (Example 3) to illustrate how this model contributes to the concise programming of a mobile agent (`app@sp` in Figure 7) computing in a ubiquitous computing environment (Figure 10).

The BigActor model is defined over a bigraph transition system. This is formally a sequence of bigraphs generated by applying Bigraph Reaction Rules [10] to an initial bigraph. Thus, structure can be dynamic and the three new semantic rules are designed to program actors that adapt to changing structure. Context-Aware programming needs a semantics for the interaction of program and context. In this paper we borrow the one underpinning control theory (see Figure 1), because it is widely used. In control theory the system is separated into plant and controller with the controller being composed in feedback. It observes, then controls, and the control generates new observations, repeating the cycle.




Figure 1: BigActors: actors composed as a feedback loop with a bigraph.

Section 6 shows the BigActor model is able to support this semantics of adaptation. The controller is restricted to being a bigActor and the plant a bigraph. Definition 5 is the semantics of the feedback loop. We define a feedback bigActor analogous to a feedback controller in Definition 11. Theorem 1 shows that in a system with one bigActor, being a feedback bigActor is sufficient for the control-theoretic semantics. This is called *correctness* for want of a better term. Theorems 2 and 3 show another sufficient condition is required for correctness when there are multiple concurrent bigActors. The additional conditions required to realize the feedback control semantics are the standard ones used to exclude race conditions in concurrent programming. Theorems 1, 2, and 3 formalize the sense in which the bigActor model supports the programming of concurrent agents working in a world with changing structure by adapting to it.

We situate the BigActor model in the Actor and Bigraph literatures as follows. The contributions of this model to the Actor literature are to provide the actor with a reflection of the structure of the system. This approach has been adopted to reflect back to actor systems low-level properties of the system. Nielsen, Ren, and Agha [11, 12] provide a real-time semantics to actor systems using timed graphs to constrain the execution of actors. Agha introduces locality to actors semantics [13] for modelling multiple mobile agents. The location model used by Agha is simply a host relation between actors and hosts. The location model is flat in the sense that there is no information of where a host stands in relation to another host. Moreover, there is no sense of connectivity between hosts. Using bigraphs in the BigActor model we provide a richer model of the structure, capable of modelling nested locality of components, their connectivity and also the way the structure evolves with time.

From a bigraph research stand-point, the BigActor model provides means for embedding computation into bigraphs that model the structure of the world. This is addressed in the literature by [5, 4] which model the structure of the world and computation using different Bigraph Reaction Systems (BRS) which are composed together. It is known that querying bigraphs exclusively using bigraphs reactive systems is difficult [4]. This is addressed by [4] by introducing three different BRSs: context, proxy, and agents. Our approach differs from [5] and [4] in the sense that we model computation using the actors model, which is coupled from an observation stand point using a query language. This removes the burden of modelling queries using BRSs. Moreover, by combining the Actor model with bigraphs we intent to leverage the use of bigraphs together with the large spectrum of actor languages and frameworks (e.g. Erlang [3], Scala [7], Cloud Haskell [6], Dart, Akka, etc.).

2. BIGRAPHICAL FRAMEWORK

In this section we introduce Bigraphs. A reader familiar with Bigraphs should be able to skip this section. The examples, however are used throughout the paper. This explanation follows [10].

As the name suggests a bigraph is a mathematical structure with two graphs, the place graph - a forest that represents nested locality of components and a link graph - a hypergraph that models connectivity between components. Figure 2 presents an example of a bigraph.

Place graphs are contained inside *regions* (dashed rectangles) and may also contain *holes* (dark grey empty rect-




Figure 2: Example of a bigraph and the corresponding place and link graphs.

angles) Regions and holes enable composition of placing graphs, i.e. a hole of a given bigraph can be replaced by a region of another bigraph using the composition operator. We explain composition later.

A link graph may contain hyperedges and *inner names* and *outer names*. Names are graphically represented by a line connected at one end to a port or an edge and the other end is left loose. Just as one can fit regions inside holes, one can also merge inner names and outer names using the composition operator.

A node can have *ports* (black dots) which are points for connections to edges or names. The kinds of nodes and their number of ports (arity) are the signature of the bigraph.

The *signature* takes the form (\mathcal{K}, ar) where \mathcal{K} is a set of kinds of nodes called *controls* and $ar : \mathcal{K} \rightarrow \mathbb{N}$ assigns an arity (i.e. a natural number) to each control. Each node in the bigraph is assigned a *control*. For example, the bigraph of Figure 2(a) has the following signature: $\mathcal{K} = \{M : 1, Q : 2\}$ where m_i has kind M with arity 1 and q_i has kind Q with arity 2. By convention we start kind names with upper-case characters and node names with lower-case characters.

A bigraph B is called *concrete* when each node and each edge is assigned a unique identifier (known as *support*). We denote the set of node identifiers of B as V_B and the set of edge identifiers as E_B (e.g. $V_B = \{m_0, m_1, m_2, q_0, q_1\}$ and $E_B = \{e\}$ in the example of Figure 2).

A bigraph without support is called *abstract*. In abstract bigraphs, nodes are denoted by their control while edges are kept anonymous. Abstract bigraphs are defined using an algebra. In this paper we almost exclusively use concrete bigraphs and thus we devote this section to present them formally. For a formal introduction to abstract bigraphs see [10], Chapter 3.

In order to define bigraphs formally we need to introduce the concept of a bigraph interface. An **interface** is a pair $\langle n, X \rangle$ where n in the interface denotes the set $\{0, 1, \dots, n-1\}$ of holes (or regions) and X denotes the set of inner names (or outer names). Holes and inner names are collectively called an *inner face* while regions and outer names are called an *outer face*.

Definition 1. A **bigraph** is a 5-tuple of the form:

$$(V, E, ctrl, prnt, link) : \langle m, X \rangle \rightarrow \langle n, Y \rangle$$

where V is a set of nodes, E is the set of hyperedges, $ctrl : V \rightarrow \mathcal{K}$ is a control map that assigns controls to nodes,

$prnt : m \uplus V \rightarrow V \uplus n^1$ is the parent map and defines the nested place structure, $link : X \uplus P \rightarrow E \uplus Y$ is the link map and defines the link structure. P denotes the set of ports of the bigraph and is formalized as $P = \{(v, i) \mid i \in \{0, 1, \dots, ar(ctrl(v)) - 1\}\}$. For convenience we introduce a map $Pts : V_B \rightarrow \mathcal{P}(\mathbb{N})$ that takes a node and returns the set of ports of that node. $\langle m, X \rangle \rightarrow \langle n, Y \rangle$ provides the inner face and outer face of the bigraph, i.e. $\langle m, X \rangle$ is the inner face and $\langle n, Y \rangle$ is the outer face.

The composition of bigraphs is defined by matching interfaces. Like function composition, a bigraph $A : I \rightarrow J$ composed with a bigraph $B : K \rightarrow I$ is a bigraph $C : K \rightarrow J$.

Definition 2. Let $A : \langle m_a, X_a \rangle \rightarrow \langle n_a, Y_a \rangle$ and $B : \langle m_b, X_b \rangle \rightarrow \langle n_b, Y_b \rangle$. The **composition** of bigraphs A and B , denoted by $A \circ B$, is a bigraph $C : \langle m_b, X_b \rangle \rightarrow \langle n_a, Y_a \rangle$ where $V_C = V_A \uplus V_B$, $E_C = E_A \uplus E_B$, $ctrl_C = ctrl_A \uplus ctrl_B$. $prnt_C$ is obtained by “filling” the holes of A with regions of B while $link_C$ is obtained by “merging” the inner names of A with the outer names of B . By convention, the regions are matched to holes with the same indices and inner names are matched with outer names with the same name. For a formal definition of $prnt_C$ and $link_C$ see [10], page 17. The composition $A \circ B$ is well defined iff the outer face of B is equal to the inner face of A , i.e. $\langle n_b, Y_b \rangle = \langle m_a, X_a \rangle$.

Example 1. Consider the bigraphs of Figure 3. The bi-

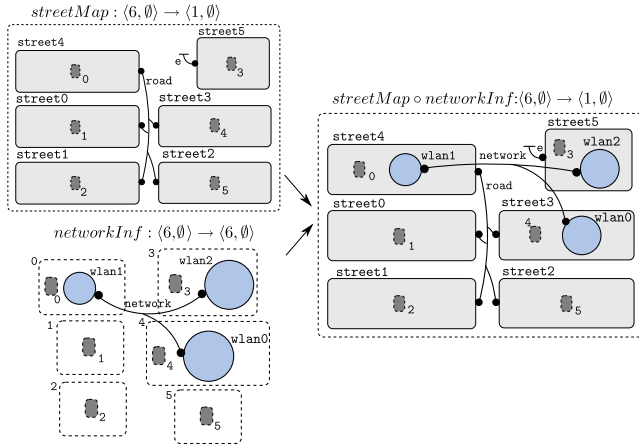


Figure 3: Composition of the bigraph $streetMap$ with the bigraph $networkInf$.

graph $streetMap : \langle 6, \emptyset \rangle \rightarrow \langle 1, \emptyset \rangle$ models a street map where the grey nodes represent streets and links represent physical adjacency between street nodes. Bigraph $networkInf : \langle 6, \emptyset \rangle \rightarrow \langle 6, \emptyset \rangle$ models a network infrastructure where the blue nodes model wireless hotspots and links represent connectivity which is linked to the edge **network** modelling a network. The bigraph resulting from the composition of both bigraphs is $streetMap \circ networkInf : \langle 6, \emptyset \rangle \rightarrow \langle 1, \emptyset \rangle$. Note that since the $networkInf$ keeps the holes inside street nodes one could compose $streetMap \circ networkInf$ with other features (e.g. cars, pedestrians, utilities network, etc.). The signature for this example and subsequent ones in this

¹The symbol \uplus denotes the exclusive union operator for sets.




Figure 4: Abstract BBRs MOVE that moves a Comp node from one street to another, and CONNECT that connects the Comp node to the network infrastructure.

paper is $\mathcal{K} = \{\text{Street} : 1, \text{Wlan} : 1, \text{Comp} : 1, \text{Feat} : 0\}$. In Figure 3, **street_i** are of kind **Street**, and **wlan_i** are of kind **Wlan**. The kinds **Street** and **Wlan** denote respectively streets on a city and locations with wireless connectivity. The kinds **Comp** and **Feat** are going to be used in examples later and represent, respectively computational devices and features of the world to be observed.

2.1 Dynamics of bigraphs

Milner defines Bigraph Reaction Rules (BRR) to create dynamics on bigraphs [10].

A *bigraph reaction rule* is a tuple (R, R', η) where R and R' are bigraphs called respectively *redex*² and *reactum*. The redex is the portion of the bigraph to be matched and the reactum is the bigraph that replaces the matched portion. η is called the instantiation map and indicates how holes in R correspond to holes in R' . If η is the identity map, then we represent the rule as $R \rightarrow R'$. For the remainder of the paper we always assume η is the identity map, i.e. hole i in R matches with hole i in R' for all $i \in \mathbb{N}$. If R and R' are abstract bigraphs, then $R \rightarrow R'$ is an abstract BRR. If R and R' are concrete then the BRR is concrete. It is often convenient to define BBRs to be abstract even when working with concrete bigraphs. This defines rules that can be applied to several contexts.

Let $r = R \rightarrow R'$ be a BRR and B a bigraph. In order to perform the reaction r in B we need first to decompose B into $C \circ R \circ d$ where C represents the context and d represents the parameters inside the holes of R . Assuming that η is the identity map, we compose C with the reactum R' and with d to get the resulting B' , i.e. $B = C \circ R \circ d \Rightarrow B' = C \circ R' \circ d$. This application of BRRs works both for abstract and concrete BRRs.

Example 2. Consider the two abstract BRR of Figure 4. The first rule, **MOVE**, models a computational device (e.g. a user with a smartphone) denoted by a star with kind **Comp**) moving from one street to another. The second rule, **CONNECT**, models computational device to move inside a hotspot with network connectivity (denoted by a node with kind **Wlan**) and connecting to the network. Note that the rules are parametric, i.e. they can be applied regardless the nodes inside the street nodes. For example, this rule allows modelling a computational device to move with or without other users in a street. Figure 5 shows an example of the application of BRR **MOVE**. The bigraphs in this example are concrete. Thus, we need to first concretize **MOVE**. Since we want to move **sp0** of kind **Comp** from **street2** to **street3** we

²Redex stands for “reducible expression”.

denote the concrete BRR as $\text{MOVE}(\text{street2}, \text{street3})$. Figure 5 depicts the context C , and parameters d where the rule is applied. Note that the parameters allow the rule to be applied with sp1 inside street3 . The upper part of Fig-

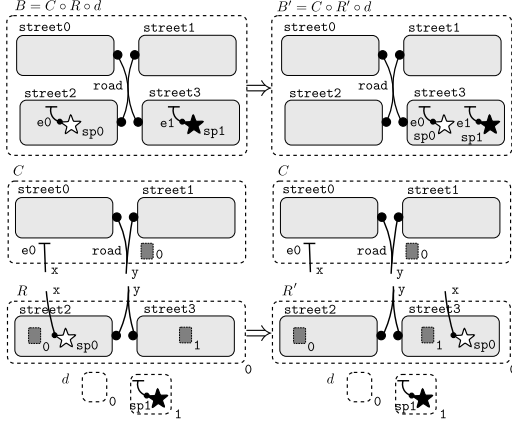


Figure 5: Example of the application of the MOVE reaction rule.

ure 5 shows the bigraph B transitioning to B' . Under B and B' we show the decomposition of each bigraph in the context C , redex R , reactum R' , and parameters d .

3. ACTOR MODEL

The Actors model of computation is a model for distributed concurrent computing entities [1, 2]. An actor system is composed of autonomous objects called actors. Actors communicate using asynchronous message passing. Messages that have been sent but not yet received are queued up in the receiver actor's mailbox. The receiver eventually removes the message and processes it. An actor encapsulates a state and a thread. Each actor has a mail-address used by other actors to send it messages.

As a response to a message an actor may: **compute** and change state; **send** a message; or **create** new actors.

The Actor model has been adopted as the concurrency model in several programming languages such as Erlang [3], Scala [7], Dart, Could Haskell [6], and in embedded systems design [8].

Next we present an operational semantics (using the contextual style) for the actors model. We follow the operational semantics approach taken in [11, 2, 13].

The Actor model operational semantics is formalized as a transition relation over the set of actors configurations.

Definition 3. The Actor model configuration is a tuple $\langle \alpha \mid \mu \rangle$ where α is a set of actor names, and μ is the set of pending messages. An actor $a \in \alpha$ can be either *busy* or *inactive*. An actor is busy if it is currently executing. $[E \vdash R \sqsubset e \sqsupset]_a$ denotes a **busy actor** a , with environment E (i.e. local state), and reduction context R filled with expression e which is currently being executed. An actor is *inactive* if it is waiting for a message. $(E \vdash b)_a$ denotes an **inactive actor** a with environment E and next behaviour b . When a message arrives, the actor's behaviour b is applied to the corresponding message. $\langle a \Leftarrow v \rangle$ denotes a message to a with content v .

The operational semantics is defined using the five rules³ of Figure 6. The rule $\langle \text{fun} : \mathbf{a} \rangle$ models the execution of an ex-

$$\begin{aligned} \langle \text{fun} : \mathbf{a} \rangle & \frac{[E \vdash e]_a \rightarrow_\lambda [E' \vdash e']_a}{\langle \alpha, [E \vdash p]_a \mid \mu \rangle \rightarrow \langle \alpha, [E' \vdash p']_a \mid \mu \rangle} \\ \langle \text{new} : \mathbf{a}, \mathbf{a}' \rangle & \langle \alpha, [E \vdash R \sqsubset \text{new}(\square)]_a \mid \mu \rangle \rightarrow \langle \alpha, [E \vdash R \sqsubset \text{nil} \square]_a, (E \vdash b)_{a'} \mid \mu \rangle \\ \langle \text{term} : \mathbf{a} \rangle & \langle \alpha, [E \vdash R \sqsubset \square]_a \mid \mu \rangle \rightarrow \langle \alpha, (E \vdash b)_a \mid \mu \rangle \\ \langle \text{rcv} : \mathbf{a}, \langle \mathbf{a} \Leftarrow \mathbf{v} \rangle \rangle & \langle \alpha, (E \vdash b)_a \mid \mu, \langle \mathbf{a} \Leftarrow \mathbf{v} \rangle \rangle \rightarrow \langle \alpha, [E[FV(b) \mapsto v] \vdash b \sqsubset \text{nil} \square]_a \mid \mu \rangle \\ \langle \text{snd} : \mathbf{a}, \langle \mathbf{a}' \Leftarrow \mathbf{m} \rangle \rangle & \langle \alpha, [E \vdash R \sqsubset \text{send}(\mathbf{a}', \mathbf{m}) \square]_a \mid \mu \rangle \rightarrow \langle \alpha, [E \vdash R \sqsubset \text{nil} \square]_a \mid \mu, \langle \mathbf{a}' \Leftarrow \mathbf{v} \rangle \rangle \end{aligned}$$

Figure 6: Actor operational semantics.

pression e of an actor a . \rightarrow_λ denotes the semantics of a given language that specifies the local computation of the actor. Note that the execution of e to e' can produce side-effects in the actor's local environment E . This rule is also responsible for removing *nil* and making the actor ready to be terminated by rule **term**. The rule $\langle \text{new} : \mathbf{a}, \mathbf{a}' \rangle$ models the spawn of a new actor a' from a . The new actor a' becomes inactive after creation. The rule $\langle \text{term} : \mathbf{a} \rangle$ models the termination of the execution of a command, making the actor changing from busy to inactive. The rule $\langle \text{rcv} : \mathbf{a}, \langle \mathbf{a} \Leftarrow \mathbf{v} \rangle \rangle$ models an inactive actor a removing a message from its mailbox and updating its local environment with the contents \mathbf{v} . The rule $\langle \text{snd} : \mathbf{a}, \langle \mathbf{a}' \Leftarrow \mathbf{v} \rangle \rangle$ models an actor a sending a message \mathbf{v} to a' . The new message is included in the set of pending messages.

4. BIGACTOR MODEL

This section formalizes the BigActor Model. The operational semantics is presented in Section 4.1. We explain the semantics by deriving a trace of the program in Example 3.

Example 3. Consider the pseudo-code in Figure 7. The pseudo-code is in the style of actor languages. The program

```
app@sp
method start()
  control(MOVE(street2,street1));
  observe(children.parent.host);
method rcvObs(B_obs: Bigraph)
  control(MOVE(street1,street3));
  control(CONNECT(wlan0));
  send(server@srv,B_obs);
```

Figure 7: Pseudo-code for the appBA BigActor.

models the following scenario. An app is hosted on a smartphone (app@sp) in street2 . It wants to photograph the green boxes (Figure 10) in street1 . Hence the smartphone (bigraph node sp) moves. It specifies the picture to be taken by a query and waits till the picture is bound to its free variable. In order to send the picture to server@srv the app

³The original semantics includes two more rules for handling communication with external actors, i.e. actors not defined in the configuration. External actors are not an essential feature for bigActors and thus we decided to remove them for the sake of simplifying the exposition.

moves to one of the blue circles modeling wireless hotspots, and connects to the hyperlink **network**. It then sends the message to the server.

4.1 Operational semantics

The operational semantics is entirely in Figure 8. We use the same contextual style as the actor semantics in Section 3. A BigActor model executes by transitioning through a sequence of configurations. Definition 4 defines a configuration.

Definition 4. A BigActor **configuration** is a tuple $\langle \alpha \mid \mu \mid \eta \mid B \rangle$ where α is a set of bigActors, μ is a set of pending messages, η is a set of pending requests, and B is a bigraph. A BigActor in α is of the form $a@h$, with a an actor, and $h \in V_B$ a node in the bigraph B . h is called the host of a . The set of pending messages μ is as in the actor semantics in Section 3. A request in the set of requests η is a tuple $(a@h, r)$ meaning that the operation r is requested by $a@h$. The operation r can be any one of the expressions $\mathbf{send}(a'@h', m)$, $\mathbf{observe}(q)$, $\mathbf{control}(u)$, $\mathbf{migrate}(h')$. These are the premises of semantic rule $\langle \mathbf{req} : a@h, r \rangle$ in Figure 8. $a'@h'$ can be any BigActor in α , m any actor message, q a query resulting in an observation by a BigActor of the bigraph B as expressed by Equation 1, u any concrete bigraph reaction rule as described in Section 2, and h' another node in the bigraph B .

The interpretation of a query into an observation of a given bigraph B relative to a host h is specified by the map:

$$\llbracket \cdot \rrbracket_h^B : \mathbf{query} \rightarrow \mathbb{B} \quad (1)$$

Queries and observations are addressed in Section 5.

Of the four elements in the BigActor configuration, α is as in the actor model except we write the $a@h$ instead of just a for the elements of α , μ is exactly as in the actor model, and η and B are new. Definition 5 defines an execution of the BigActor model as a sequence of configurations.

Definition 5. A BigActor execution is a sequence c_0, c_1, c_2, \dots where each $c_i \xrightarrow{\lambda_i} c_{i+1}$ and $\xrightarrow{\lambda_i}$ is derived by a semantic rule labelled λ_i . The permissible λ_i are those in Figure 8.

The semantics can be understood as follows. A BigActor can do actor computations. This is asserted by the inclusion of the rules $\langle \mathbf{fun} : a@h \rangle$, $\langle \mathbf{new} : a@h, a'@h \rangle$, $\langle \mathbf{term} : a@h \rangle$, and $\langle \mathbf{rcv} : a@h, m \rangle$. This rules are the same as their actor counterparts in Figure 6. The augmentation of an actor a to $a@h$ and the expansion of the configuration to include η and B has no significance in the rules \mathbf{fun} , \mathbf{term} and \mathbf{rcv} . In the rule \mathbf{new} , the $@h$ is used to assert that a new actor will have the same host as its creator. Thus to an actor programmer \mathbf{new} has the usual actor semantics, though in the BigActor model \mathbf{new} where the $@h$ has the semantics of a host, the rule prevents the creation of remote bigActors. The environment E in Figure 8 is the same as the E in the actor model in all rules.

In addition to executing expressions by \mathbf{fun} , \mathbf{term} , \mathbf{new} and \mathbf{rcv} as an actor would, a BigActor can also execute the expressions $\mathbf{send}(a'@h', m)$, $\mathbf{observe}(q)$, $\mathbf{control}(u)$ and $\mathbf{migrate}(h')$ as syntactically formalized in Definition 4. The semantics of these expressions are unique to this model and given first by the rule $\langle \mathbf{req} : a@h, r \rangle$ and then by the rules labelled as \mathbf{snd} , \mathbf{obs} , \mathbf{ctr} and \mathbf{mgrt} , respectively. The \mathbf{req}

rule asserts that when any of the expressions $\mathbf{send}(a'@h', m)$, $\mathbf{observe}(q)$, $\mathbf{control}(u)$ or $\mathbf{migrate}(h')$ execute in a BigActor program, the program advances exactly as an actor program by the rule \mathbf{fun} in Figure 6 but with a side-effect. The side-effect is the addition of the request $(a@h, r)$ to the set η . This creates a new configuration. When η contains a request with an expression $\mathbf{send}(a'@h', m)$, $\mathbf{observe}(q)$, $\mathbf{control}(u)$ or $\mathbf{migrate}(h')$ a BigActor model can further advance by application of the rules \mathbf{snd} , \mathbf{obs} , \mathbf{ctr} or \mathbf{mgrt} as described next.

To make the rest comprehensible we analyze an execution trace of the program in Figure 7 from Example 3.

The $\mathbf{control}$, $\mathbf{observe}$ and \mathbf{send} instructions have the syntax and semantics in Figure 8 and Definition 4. These instructions are encapsulated in two behaviors named the method \mathbf{start} and method \mathbf{rcvObs} . The syntax of the method construct and sequentiality are not formalized. We informally assume instructions execute in the order in which they appear in the text and the BigActor $\mathbf{app@sp}$ starts with the method \mathbf{start} . The rest follows formally from the semantics in figure 8. Figure 9 is a formal execution of this program in the sense of definition 5.

All configurations in Figure 9 follow the notation in Definition 4, Definition 5, and two short-hands for brevity. The Figure has 9 lines. Each line has two configurations. Each line has two transitions except the last one which has only one. In each configuration the reduction context, denoted by R in the semantic rules, is abstracted with a \cdot . The $\mathbf{app@sp}$ program instruction executing at each step is explicitly typed in the square brackets in each configuration. The environment E in the semantic rules is also abstracted by \cdot in the configurations in the first 3 lines and the first configuration in line 4. The environment is a valuation of the variables of $\mathbf{app@sp}$ just as in the actor model. $\mathbf{app@sp}$ has only one variable $\mathbf{B.obs}$. This has no valuation in the first three lines and acquires a value in the second configuration in line 4. It is denoted with an E thereafter. B_0 through B_3 are bigraphs. These are visualized in Figure 10. This figure visualizes a projection of the trace in Figure 9. It projects away all transitions representing applications of \mathbf{fun} , \mathbf{term} and \mathbf{req} . The \mathbf{fun} and \mathbf{term} transitions advance the program $\mathbf{app@sp}$ just like an actor program and do nothing particular to this model. The \mathbf{req} transitions advance like \mathbf{fun} but also have the side-effect of adding a request to the set η . Then other semantic rules can be applied to consume the elements in η as mentioned earlier to produce the configurations visualized in Figure 10.

The execution trace in Figure 9 can now be understood as follows. We assume the BigActor $\mathbf{app@sp}$ starts with the behavior \mathbf{start} and advances to the first instruction by the rule \mathbf{fun} . The next transition is by the \mathbf{req} rule and puts $\mathbf{control}(\mathbf{move}(\mathbf{street2}, \mathbf{street1}))$ in η . $\mathbf{MOVE}(\mathbf{street2}, \mathbf{street1})$ is a concretization of the abstract BRR \mathbf{MOVE} in Figure 4. This enables use of the rule \mathbf{ctr} to change the bigraph B_0 to B_1 . This is visualized in the top row of Figure 10. One can see the white star move. Everything else in the two bigraphs is the same. The second transition in line 2 simply advances the actor program by rule \mathbf{fun} to the instruction $\mathbf{observe}(\mathbf{children.parent.host})$. The first transition in line 3 puts this imperative into η by the rule \mathbf{req} . The first behavior is now exhausted and the actor become inactive by application of the rule \mathbf{fun} and then \mathbf{term} resulting in the two configurations in line 4. In the second configuration

$$\begin{array}{c}
\langle \text{fun} : \mathbf{a@h}, \langle \text{new} : \mathbf{a@h}, \mathbf{a'@h} \rangle \mid \langle \text{req} : \mathbf{a@h}, \mathbf{r} \rangle \\
\langle \text{term} : \mathbf{a@h}, \langle \text{rcv} : \mathbf{a@h}, \mathbf{m} \rangle \mid \frac{\mathbf{r} \in \{\text{send}(\mathbf{a'@h'}, \mathbf{m}), \text{observe}(\mathbf{q}), \text{control}(\mathbf{u}), \text{migrate}(\mathbf{h}')\}}{\langle \alpha, [E \vdash R \sqsubset \mathbf{r} \sqsupset]_{\mathbf{a@h}} \mid \mu \mid \eta \mid B \rangle \rightarrow \langle \alpha, [E \vdash R \sqsubset \text{nil} \sqsupset]_{\mathbf{a@h}} \mid \mu \mid \eta, (\mathbf{a@h}, \mathbf{r}) \mid B \rangle} \\
\langle \text{obs} : \mathbf{a@h}, \mathbf{q}, \mathbf{o} \rangle \\
\frac{B_{\text{obs}} = \llbracket \mathbf{q} \rrbracket_h^B \quad \mathbf{q} \in \mathcal{Q}_a \quad E' = E[FV(b) \mapsto B_{\text{obs}}]}{\langle \alpha, (E \vdash b)_{\mathbf{a@h}} \mid \mu \mid \eta, (\mathbf{a@h}, \text{observe}(\mathbf{q})) \mid B \rangle \rightarrow \langle \alpha, E' \vdash b \sqsubset \text{nil} \sqsupset]_{\mathbf{a@h}} \mid \mu \mid \eta \mid B \rangle} \\
\langle \text{ctr} : \mathbf{a@h}, \mathbf{u} \rangle \quad \mathbf{u} = (R \rightarrow R') h \in V_R \cap V_{R'} \quad B = C \circ R \circ d B' = C \circ R' \circ d \quad \left| \frac{\langle \text{mgrt} : \mathbf{a@h}, \mathbf{h}' \rangle}{\mathbf{h}' \in V_B \quad \exists pt \in Pts(h). \exists pt' \in Pts(h'). \text{link}(pt) = \text{link}(pt')} \right. \\
\frac{\langle \alpha \mid \mu \mid \eta, (\mathbf{a@h}, \text{control}(\mathbf{u})) \mid B \rangle \rightarrow \langle \alpha \mid \mu \mid \eta \mid B' \rangle}{\langle \alpha, \mathbf{a@h} \mid \mu \mid \eta, (\mathbf{a@h}, \text{migrate}(\mathbf{h}')) \mid B \rangle \rightarrow \langle \alpha, \mathbf{a@h}' \mid \mu \mid \eta \mid B \rangle} \\
\langle \text{snd} : \mathbf{a@h}, \langle \mathbf{a'@h}' \leftarrow \mathbf{m} \rangle \rangle \\
\frac{\exists pt \in Pts(h). \exists pt' \in Pts(h'). \text{link}(pt) = \text{link}(pt') \vee (h = h')}{\langle \alpha \mid \mu \mid \eta, (\mathbf{a@h}, \text{send}(\mathbf{a'@h'}, \mathbf{m})) \mid B \rangle \rightarrow \langle \alpha \mid \mu, \mathbf{m} \mid \eta \mid B \rangle}
\end{array}$$

Figure 8: BigActors operational semantics.

$$\begin{array}{c}
\langle [\cdot \vdash \cdot \sqsubset \text{start}() \sqsupset]_{\text{app@sp}} \mid \emptyset \mid \emptyset \mid B_0 \rangle \xrightarrow{\langle \text{fun} : \text{app@sp} \rangle} \langle [\cdot \vdash \cdot \sqsubset \text{control}(\text{MOVE}(\text{street2}, \text{street1})) \sqsupset]_{\text{app@sp}} \mid \emptyset \mid \emptyset \mid B_0 \rangle \xrightarrow{\langle \text{req} : \text{app@sp}, \text{MOVE}(\cdot, \cdot) \rangle} \\
\langle [\cdot \vdash \cdot \sqsubset \text{nil} \sqsupset]_{\text{app@sp}} \mid \emptyset \mid (\text{app@sp}, \text{MOVE}(\text{street2}, \text{street1})) \mid B_0 \rangle \xrightarrow{\langle \text{ctr} : \text{app@sp}, \text{MOVE}(\cdot, \cdot) \rangle} \langle [\cdot \vdash \cdot \sqsubset \text{nil} \sqsupset]_{\text{app@sp}} \mid \emptyset \mid \emptyset \mid B_1 \rangle \xrightarrow{\langle \text{fun} : \text{app@sp} \rangle} \\
\langle [\cdot \vdash \cdot \sqsubset \text{observe}(\mathbf{q}) \sqsupset]_{\text{app@sp}} \mid \emptyset \mid \emptyset \mid B_1 \rangle \xrightarrow{\langle \text{req} : \text{app@sp}, \text{observe}(\mathbf{q}) \rangle} \langle [\cdot \vdash \cdot \sqsubset \text{nil} \sqsupset]_{\text{app@sp}} \mid \emptyset \mid (\text{app@sp}, \text{observe}(\mathbf{q})) \mid B_1 \rangle \xrightarrow{\langle \text{fun} : \text{app@sp} \rangle} \\
\langle [\cdot \vdash \cdot \sqsubset]_{\text{app@sp}} \mid \emptyset \mid (\text{app@sp}, \text{observe}(\mathbf{q})) \mid B_1 \rangle \xrightarrow{\langle \text{term} : \text{app@sp} \rangle} \langle (E \vdash \text{rcvObs})_{\text{app@sp}} \mid \emptyset \mid (\text{app@sp}, \text{observe}(\mathbf{q})) \mid B_1 \rangle \xrightarrow{\langle \text{obs} : \text{app@sp}, \mathbf{q}, B_{\text{obs}} \rangle} \\
\langle [E[FV(\text{rcvObs}) \mapsto B_{\text{obs}}] \vdash \text{rcvObs}]_{\text{app@sp}} \mid \emptyset \mid \emptyset \mid B_1 \rangle \xrightarrow{\langle \text{fun} : \text{app@sp} \rangle} \langle [E \vdash \cdot \sqsubset \text{control}(\text{MOVE}(\text{street1}, \text{street3})) \sqsupset]_{\text{app@sp}} \mid \emptyset \mid \emptyset \mid B_1 \rangle \xrightarrow{\langle \text{req} : \text{app@sp}, \text{MOVE}(\cdot, \cdot) \rangle} \\
\langle [E \vdash \cdot \sqsubset \text{nil} \sqsupset]_{\text{app@sp}} \mid \emptyset \mid (\text{app@sp}, \text{MOVE}(\text{street1}, \text{street3})) \mid B_1 \rangle \xrightarrow{\langle \text{ctr} : \text{app@sp}, \text{MOVE}(\cdot, \cdot) \rangle} \langle [E \vdash \cdot \sqsubset \text{nil} \sqsupset]_{\text{app@sp}} \mid \emptyset \mid \emptyset \mid B_2 \rangle \xrightarrow{\langle \text{fun} : \text{app@sp} \rangle} \\
\langle [E \vdash \cdot \sqsubset \text{control}(\text{CONN.}(\text{wlan0})) \sqsupset]_{\text{app@sp}} \mid \emptyset \mid \emptyset \mid B_2 \rangle \xrightarrow{\langle \text{req} : \text{app@sp}, \text{CONN.}(\cdot) \rangle} \langle [E \vdash \cdot \sqsubset \text{nil} \sqsupset]_{\text{app@sp}} \mid \emptyset \mid (\text{app@sp}, \text{CONN.}(\text{wlan0})) \mid B_2 \rangle \xrightarrow{\langle \text{ctr} : \text{app@sp}, \text{CONN.}(\cdot) \rangle} \\
\langle [E \vdash \cdot \sqsubset \text{nil} \sqsupset]_{\text{app@sp}} \mid \emptyset \mid \emptyset \mid B_3 \rangle \xrightarrow{\langle \text{fun} : \text{app@sp} \rangle} \langle [E \vdash \cdot \sqsubset \text{send}(\text{server@srv}, \text{B_obs}) \sqsupset]_{\text{app@sp}} \mid \emptyset \mid \emptyset \mid B_3 \rangle \xrightarrow{\langle \text{req} : \text{app@sp}, \text{send}(\cdot) \rangle} \\
\langle [E \vdash \cdot \sqsubset \text{nil} \sqsupset]_{\text{app@sp}} \mid \emptyset \mid (\text{app@sp}, \text{send}(\text{server@srv}, \text{B_obs})) \mid B_3 \rangle \xrightarrow{\langle \text{snd} : \text{app@sp}, \mathbf{h}, \mathbf{m} \rangle} \langle [E \vdash \cdot \sqsubset \text{nil} \sqsupset]_{\text{app@sp}} \mid \langle \text{server@srv} \Rightarrow \text{B_obs} \rangle \mid \emptyset \mid B_3 \rangle
\end{array}$$

Figure 9: Execution of BigActor program from Example 3.

in line 4 the BigActor `app@sp` has advanced to the second behavior called `rcvObs` in Figure 7.

The `observe(q)` imperative is still present in η which allows us to apply the rule `obs` taking us to the first configuration in line 5. The query q in the program is `children.parent.host` and is interpreted on bigraph B_1 with respect to sp . This transition is visualized in Figure 10 as the transition from the top row to the middle one. Since the host node of `app` is `sp` denoted by the white star, its parent is `street1` and the children are the green nodes and `sp` itself. The interpretation of queries is formalized in Section 5. The application of rule `obs` produces the bigraph enclosed by the dashed region in the left-hand figure of the middle row of Figure 10. The transmission of the observed bigraph to the BigActor in the rule `obs` is handled much like an actor receives a message. The inactive actor become busy and the observation is bound to one of its free variables, `FV(rcvObs)` in line 5 of Figure 9. The behavior has only one free variable named `B_obs` in Figure 7. The environment E is now the observed bigraph being the valuation of `B_obs`. η once again becomes empty. The rest of line 5, lines 6 and 7 in Figure 9 are derived like the prior lines. The rule `CONNECT(wlan0)` is a concretization of the abstract BRR in Figure 4.

Lines 8 and 9 illustrate the semantics of the BigActor `send`. In the second configuration in line 8 the BigActor is busy at the expression `send(server@srv, B_obs)`. Like the actor model `server@srv` is intended to denote another Bi-

gActor supposed to receive the message. The message is the value of `B_obs`. The actor semantics would put this message in μ . We put the expression in η just like the other requests and then apply the `send` rule to complete message delivery. We do this because the premises of the `send` rule requires a link in the bigraph between the hosts of the sending and receiving bigActors. If the existence of this link can be proven, the message is transferred from η to μ . It can in this case due to prior execution of the expression `connect(wlan0)` as illustrated by the transition connecting the middle and last rows of Figure 10, It would then be received by the other actor by application of the rule `rcv` exactly as in the actor model.

The BigActor rule `send` is the counterpart of the actor rule with the same name, and the BigActor expression `send(a'@h', m)` the counterpart of the actor `send(a', m)`. However, the semantics is different since it requires the existence of a link between two hosts in the bigraph. This means a BigActor sending a message to a correct mailbox address could fail to communicate. This is a violation of actor semantics but not of BigActor semantics. A good analogy might be accessing a webpage in the internet, it is not sufficient to know the URL, one must also be connected to the internet by some means (e.g. wireless connection, 3G, etc.).

The example does not illustrate the rule `mgrt` for the expression `migrate`. This rule behaves much like the others. A BigActor can request a migration by the expression

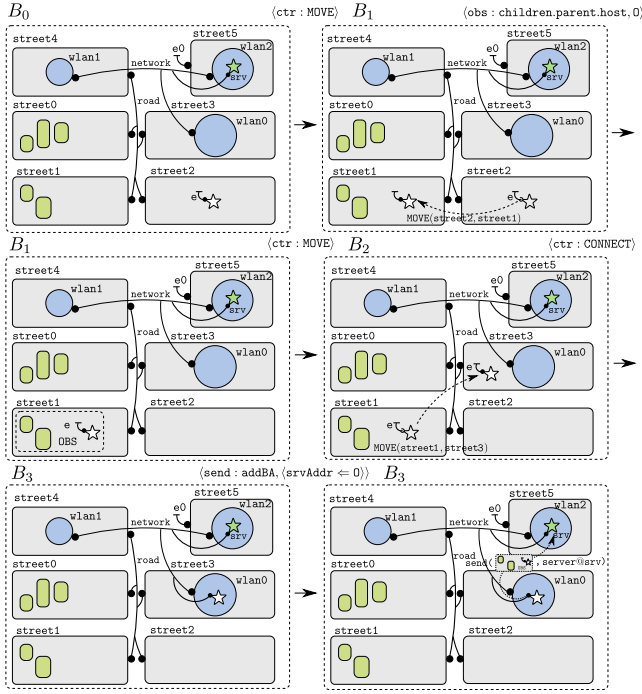


Figure 10: A projection of the execution trace of the BigActor `app@sp`.

`migrate(h')` where h' is a node in the bigraph. If the host of the BigActor and h' are connected by a link in the bigraph, the request can be consumed by the rule `mgmt` and the host of the BigActor will be changed. Note the bigraph remains unchanged. There is only a change in the hosting relationship between the actors and the bigraph. Once again, we are thinking of the link as the physical network required for the flow of a nomadic program.

We conclude this section with a few remarks.

Remark 1. The program in Figure 7 is precise in the BigActor model. We see the conciseness of the program as a contribution of the model.

Remark 2. We see the relationship between BigActor and bigraph like the relationship between a program and its machine. Actor programs are concurrent programs and so the machine can be a distributed one. The actors model the programs and the bigraph a shared distributed machine. The programs can observe the machine (rule `obs`) and change the machine (rule `ctr`). Control can move, add, or remove nodes and do the same with connections (links). The hosting relationship is the distribution of the program over the machine. The programs can change their distribution (rule `mgmt`). The programs can exchange messages (rules `snd`, `rcv`) but only up to the machine (bigraph links). The programs can compute (rules `fun`, `term`) and create new programs (rule `new`) and locate them in the machine (locally). They could migrate thereafter.

Remark 3. In this semantics the coupling of BigActor and bigraph is asynchronous. Program execution puts requests in η . The requests affect the bigraph, hosting relationship, or observations at a later instant of model time. This means a BigActor can make requests executable in the bigraph at

request time but unexecutable at the later time. Section 6 addresses this issue.

Remark 4. Control is local. If a BigActor seeks to replace component R of a bigraph with R' then the rule `ctr` requires the BigActor to be located (hosted) in R . We show in Section 5 that observation is also local.

Remark 5. In the actor model state is local. The memory of an actor influences another only up to the messages flowing between them. In the BigActor model the Bigraph is a shared structure, and some might argue this globalizes state. For example, one may communicate between bigActors by placing and removing nodes in the bigraph and never send any inter-actor messages at all. We would consider this misuse of the model, but have nevertheless elected to provide the bigraph as a shared structure to keep the model expressive. The BigActor formalism is a model for actors operating in the physical world, which may in many applications be a shared space to the actors.

Thus the bigraph can be used as a back-channel for communication between actors. There are two ways to close this channel. The first is rather trivial. One may simply make the bigraph into an actor. Then all `req`, `obs`, `ctr`, and `mgmt` actions would appear as the flow of messages between actors and the entire physical world would become the local state of one actor. This solves the problem but does not really address any of the fundamental difficulties of coordinating observing and changing a shared physical machine. Therefore we have presented a semantics in which the bigraph is not an actor, because by making the coordination difficulties visible in the model, the programmer gets the opportunity to program for them.

However, the model does support a sufficient condition, derived from its semantics closing the bigraph as a backchannel for communication. Basically, one requires two bigActors never control and observe the same region of the bigraph. The semantics of a BigActor permits us to derive the area of influence of a BigActor as expressed in Definition 6.

Definition 6. Let $a@h$ be a BigActor and $\mathcal{R}_a = \{R_0 \rightarrow R'_0, R_1 \rightarrow R'_1, \dots, R_n \rightarrow R'_n\}$ the set of all concrete BRRs used by it. The **area of influence** of $a@h$ over a bigraph B is the set of nodes A such that $v \in A \Rightarrow \exists i. B = C \circ R_i \circ d \wedge v, h \in V_{R_i}$.

Example 4. Consider Example 3 where the host of the app BigActor is inside `street2` (bigraph B_0 of Figure 10). The area of influence under the current bigraph is:

$$A^{appBA} = \{\text{street0}, \text{street1}, \text{street2}, \text{street3}, \text{street4}\}$$

Note that due to the current location of the host the only rule that can be applied is `MOVE`.

Let $c_0 \rightarrow c_1 \rightarrow \dots$ be an execution trace where $c_i = \langle \alpha_i \mid \mu_i \mid \eta_i \mid B_i \rangle$ is a BigActor configuration and \rightarrow is given by the BigActor semantics. Let $A_i^{a@h}$ be the area of influence of $a@h$ over the bigraph B_i and $B_{obs_j}^{a@h}$ be an observation of $a@h$ over the bigraph B_j . If for any two bigActors $a@h$ and $a'@h'$, $A_i^{a@h}$ and $V_{B_{obs_j}^{a'@h'}}$ never overlap for any time

$i < j$ then the bigActors cannot use the bigraph as a back channel for communication, e.g. $a@h$ might place a node in the bigraph but $a'@h'$ would not observe it because the sets $A_i^{a@h}$ and $V_{B_{obs_j}^{a'@h'}}$ would be disjoint.

5. QUERY LANGUAGE AND OBSERVATIONS

In this section we introduce a query language for querying the bigraph for local observations.

The query language syntax is specified by the following grammar:

```
query ::= node|children.node|linkedTo.node
node  ::= host|parent.node
```

Observations are local with respect to the host node of the BigActor that is querying. We think of “local” in terms of the placing graph (parents and children of a given node) and in terms of the linking graph (nodes linked to a given node). Our intention is to introduce a query language that is expressive enough for the examples in the paper and that entails the property of providing observations that are local with respect to a given host. Our aim is not to introduce a general purpose query language for Bigraphs.

The interpretation of a query q over a bigraph B with respect to a host h is defined such that a given observation $B_{obs} = \llbracket q \rrbracket_h^B$ is composable with the remaining context of B , i.e. there exists a context C and parameters d such that $B = C \circ B_{obs} \circ d$.

Example 5. Figure 11 presents an example of querying the bigraph B_0 of Figure 10 with respect to sp . On the right-hand side of Figure 11 one can see the interpretation of the queries $host$, $parent.host$, $children.parent.parent.host$.

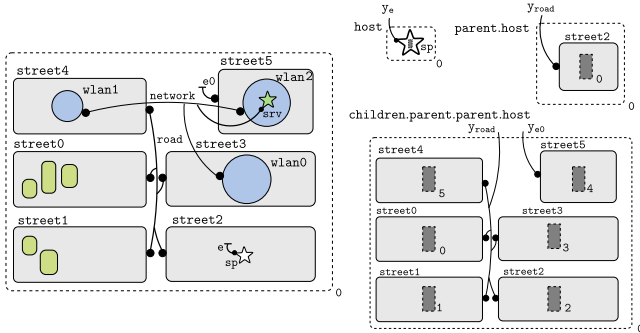


Figure 11: BigActor observing the street map using query $children.parent.parent.host$.

Let B be given by $(V_B, E_B, ctrl_B, prnt_B, link_B) : \langle m_B, X_B \rangle \rightarrow \langle n_B, Y_B \rangle$. The interpretation of the query language goes as follows.

$$\llbracket host \rrbracket_h^B \mapsto (\{h\}, \emptyset, ctrl_h, prnt_h, link_h) : \langle 1, \emptyset \rangle \rightarrow \langle 1, Y_h \rangle$$

where $ctrl_h(h) = ctrl_B(h)$,

$$prnt_h(v) = \begin{cases} 0 & \text{if } v = h \\ h & \text{if } v = 0 \end{cases}$$

$$link_h(pt) = \begin{cases} link_B(pt) & \text{if } link_B(pt) \in Y_B \\ y_e & \text{if } e = link_B(pt) \in E_B \end{cases}$$

$$Y_h = \{y_e \mid \forall pt. e = link_h(pt) \wedge e \in E_B\} \uplus \{y \mid \forall pt. y = link_h(pt) \wedge y \in Y_B\}$$

Edges are abstracted by unique names to make the composition $B = C \circ B_{obs} \circ d$ valid. The edges are included in the

context and are also connected to the same names. Composition would bind the names in the context and in the observation restoring the edges. Also note that the contents of node h are abstracted with a hole 0 .

Let the interpretation of $\llbracket node \rrbracket_h^B$ be denoted as the bigraph B_n . Note that B_n only has one node which we denote as n , i.e. $V_{B_n} = \{n\}$. The interpretation if $parent.node$ is given by:

$$\llbracket parent.node \rrbracket_h^B \mapsto (\{p\}, \emptyset, ctrl_p, prnt_p, link_p) : \langle 1, \emptyset \rangle \rightarrow \langle 1, Y \rangle$$

where $p = prnt_B(n)$. The remainder of the bigraph definition is provided by replacing p for h in the definition of $\llbracket host \rrbracket$.

The interpretation of $children.node$ is given by:

$$\llbracket children.node \rrbracket_h^B \mapsto (V_C, \emptyset, ctrl_C, prnt_C, link_C) : \langle \mid V_C \mid, \emptyset \rangle \rightarrow \langle \mid V_C \mid, Y \rangle$$

where $V_C = \{v \mid v \in V_B, n = prnt_B(v)\}$ and

$$prnt_C(v) = \begin{cases} 0 & \text{if } v \in V_C \\ v' & \text{if } v = holeOf(v') \end{cases}$$

$holeOf : V_C \rightarrow \mathbb{N}$ is a function that assigns to each children node a unique hole to abstract its contents. $ctrl_C$, $link_C$, and the set of external names Y_C are given as before although now ranging over the set of children nodes and respective ports.

Finally, the interpretation of $linkedTo$ is given by:

$$\llbracket linkedTo.node \rrbracket_h^B \mapsto (V_L, \emptyset, ctrl_L, prnt_L, link_L) : \langle \mid V_L \mid, \emptyset \rangle \rightarrow \langle \mid V_L \mid, Y \rangle$$

where $V_L = \{v \mid \exists pt \in ports(v). \exists pt' \in ports(n). link_B(pt) = link_B(pt')\}$,

$$prnt_h(v) = \begin{cases} regionOf(v) & \text{if } v \in V_L \\ v' & \text{if } v = holeOf(v') \end{cases}$$

Note that for each node in V_L we create a unique region using the function $regionOf : V_L \rightarrow \mathbb{N}$. By putting each node in a unique region we can restore the parenthood by composing again with the context. The remainder of the bigraph definition is as per the prior cases.

In this paper we use the query language **query** to define bigraphs which are local under a given bigraph and with respect to a host node.

Definition 7. Given two bigraphs B and B' and a host node h . B' is *local* with respect to h if there exists a q such that $\llbracket q.h \rrbracket_h^B = B'$ or $\llbracket q.h \rrbracket_h^B = C \circ B' \circ d$ for an arbitrary context C and parameters d . In other words, B' is local with respect to h if one can find a query with an interpretation containing B' .

An observation B_{obs} obtained by a BigActor $a@h$ querying a bigraph B using the local query language is, by Definition 7, a local bigraph over B with respect to h ($\llbracket q.h \rrbracket_h^B = C \circ B_{obs} \circ d$ where C and d are empty).

6. CORRECT EXECUTION

In this section we address the correctness of bigActors executions. Correctness is stated as a safety property over

configurations. In this section we assume that requests are served as per a First-Come First-Served (FCFS) scheduling semantics.

Definition 8. Let $c = \langle \alpha \mid \mu \mid \eta \mid B \rangle$ be a BigActor configuration. c is **terminal** if $\mu = \emptyset$, $\eta = \emptyset$, and all bigActors in α are inactive.

Definition 9. Let $c = \langle \alpha \mid \mu \mid \eta \mid B \rangle$ be a BigActor configuration. c is **correct** if $\exists \lambda.c \xrightarrow{\lambda} c'$ or c is terminal.

Let Λ denote the set of labels as per the semantics stated in Figure 8. Let Λ_η denote the set of request labels and Λ_π denote the set of **obs**, **ctr**, **mgrt**, and **snd** labels. Consider a function $schd_t : \Lambda_\pi \rightarrow \Lambda_\eta$ where, for a given execution $t = c_0 \xrightarrow{\lambda_0} c_1 \xrightarrow{\lambda_1} \dots$, $schd_t(\lambda_\pi)$ returns the request served by λ_π . Let $\prec_t \subseteq \Lambda \times \Lambda$ be a total order induced by the index i of λ_i in t .

Definition 10. An execution $t = c_0 \xrightarrow{\lambda_0} c_1 \xrightarrow{\lambda_1} \dots$ follows a First-Come First-Served (FCFS) discipline if:

$$\lambda_i \prec_t \lambda_j \Rightarrow schd(\lambda_i) \prec_t schd(\lambda_j)$$

where $\lambda_i, \lambda_j \in \Lambda_\pi$.

In other words, FCFS discipline requests that requests are served in the order that are requested.

6.1 Correctness of a single bigActor

The first step towards correctness is to require that a single bigActor itself never requests operations that will lead to an incorrect configuration.

Definition 11. Let $t = c_0 \xrightarrow{\lambda_0} c_1 \xrightarrow{\lambda_1} \dots$ be a bigActor execution. For each λ_i the operational semantics identifies a unique actor $a@h$. For clarity we augment λ_i to $\lambda_i^{a@h}$. $a@h$ is called **feedback** for all \mathbf{r} and \mathbf{r}' that are either **control**, **migrate**, or **send** the following are true:

1. $\forall j. \lambda_j = \langle \mathbf{req} : a@h, \mathbf{r} \rangle \Rightarrow \exists i < j. \lambda_i = \langle \mathbf{obs} : a@h, \mathbf{q}, B_{\mathbf{obs}} \rangle$
2. $\forall i. \forall k > i. \lambda_i = \langle \mathbf{req} : a@h, \mathbf{r} \rangle \wedge \lambda_k = \langle \mathbf{req} : a@h, \mathbf{r}' \rangle \Rightarrow \exists j. i < j < k \wedge \lambda_j = \langle \mathbf{obs} : a@h, \mathbf{q}, B_{\mathbf{obs}} \rangle$
3. For any λ_i let $B_{\mathbf{obs}}^{a@h, \lambda_i}$ denote the last observation of $a@h$ preceding λ_i . $\forall i. \lambda_i = \langle \mathbf{req} : a@h, \mathbf{r} \rangle \Rightarrow \exists \mu. \exists B. (\alpha \mid \emptyset \mid \{\mathbf{r}\} \mid B_{\mathbf{obs}}^{a@h, \lambda_i}) \xrightarrow{(a@h, \mathbf{r})} \langle \alpha \mid \mu \mid \emptyset \mid B \rangle$

The first requirement states that before any request \mathbf{r} by $a@h$ that is not an observation, there must be an observation by $a@h$. The second requirement states that between two requests \mathbf{r} and \mathbf{r}' by $a@h$ that are not observations there must be an observation by $a@h$. The last requirement states that, given a configuration with $B_{\mathbf{obs}}^{a@h, \lambda_i}$ as the bigraph and one request \mathbf{r} then there is a new configuration obtained by consuming \mathbf{r} . This means that \mathbf{r} does not make a configuration with the observed bigraph incorrect.

THEOREM 1. Consider a BigActor execution $t = c_0 \xrightarrow{\lambda_0} c_1 \xrightarrow{\lambda_1} \dots \xrightarrow{\lambda_n} c_n$ that follows a FCFS discipline (Assumption 1.1). Assume that α_i contains a single BigActor $a@h$ for all i and $a@h$ is a feedback BigActor (Assumption 1.2). Assume that η and μ are empty at c_0 (Assumption 1.3). Then c_n is correct.

PROOF. By mathematical induction in the the number of transitions.

Base case Given $c_0 \xrightarrow{\lambda_0} c_1$, c_1 is correct.

By Assumption 1.3 $a@h$ must be busy otherwise c_0 is terminal (Definition 8) and $\xrightarrow{\lambda}$ can not be triggered. If $a@h$ is busy then it can trigger any of the following rules leading to c_1 : **fun**, **new**, **term**, and **req**. Since the overall execution is assumed to have only one bigActor **new** is excluded. By Assumption 1.3, the rules **ctr**, **mgrt**, **obs**, and **snd** can not be executed since η is still empty. **fun**, and **term** are internal to the actor and can not generate a transition to an incorrect configuration. By Definition 11 item 1, the rule **req** can only request an observation because a control, migrate, or a send request require previous observations.

Inductive step Give $c_{i-1} \xrightarrow{\lambda_{i-1}} c_i$ where c_i is correct then $c_i \xrightarrow{\lambda_i} c_{i+1}$ and c_{i+1} is correct. Since c_i is correct then λ_i can be given by one of the following rules: **fun**, **rcv**, **term**, **req**, **ctr**, **mgrt**, **obs**, and **snd** (**new** is excluded by Assumption 2).

By the semantics definition of Figure 8, the rules **fun**, **rcv**, and **term** are internal to the bigActor and do not produce nor consume requests that can put the configuration in an incorrect state. The rule **req** can produce a new request, i.e. $\eta_{i+1} = \eta_i \cup \{\mathbf{r}\}$. If \mathbf{r} corresponds to an observation then c_{i+1} is correct since the observations do not change the bigraph. If \mathbf{r} is either a control, a migration or a send request then it must produce a correct configuration under $B_{\mathbf{obs}}^{a@h, \mathbf{r}}$. By Definition 11, $a@h$ alternates between observation and control/migrate/send then the bigraph B_i at configuration c_i can be decomposed as $B_i = C \circ B_{\mathbf{obs}}^{a@h, \mathbf{r}} \circ d$ where $B_{\mathbf{obs}}^{a@h, \mathbf{r}}$ is the last observation prior to \mathbf{r} . By Definition 11 item 3, \mathbf{r} can be applied to $B_{\mathbf{obs}}^{a@h, \mathbf{r}}$. Since $B_i = C \circ B_{\mathbf{obs}}^{a@h, \mathbf{r}} \circ d$ also produces a correct configuration under B_i . The rules **ctr**, **mgrt**, and **obs** consume one request \mathbf{r} out of η_i . By Assumption 1.1, all requests of $a@h$ are served in order. Since all requests in η_i were generated by a (single) feedback BigActor, then $\eta_{i+1} = \eta_i \setminus \{\mathbf{r}\}$ do not make c_{i+1} incorrect. \square

6.2 Correctness of concurrent bigActors

In the multiple bigActor case, the feedback property is necessary but not sufficient for correctness. A bigActor can request an operation executable at request time by the operational semantics, but not executable when the request is served. This is a consequence of the fact that in the bigActor operational semantics (Figure 8) the time which a request is generated is necessarily different from the time which the request is served. We take two strategies to avoid these concurrency issues and ensure a correct execution of multiple bigActors: prevent undesired interleaving (time strategy); and/or prevent shared resources (spatial strategy). These lead to theorems 2 and 3.

6.2.1 Atomic read-write semantics

Definition 12. Let Λ_θ denote the set of all **ctr**, **mgrt**, and **snd** labels and let $\Lambda_\theta^{a@h}$ be the subset of labels from Λ_θ produced by serving requests from $a@h$. An execution $t = c_0 \xrightarrow{\lambda_0} c_1 \xrightarrow{\lambda_1} \dots$ follows an atomic read-write semantics if for any $\lambda_i = \langle \mathbf{obs} : a@h, \cdot \rangle$ for which there exists a $\lambda_{i+k} \in \Lambda_\theta^{a@h}$ then $\lambda_{i+j} \notin \Lambda_\theta^{a'@h'}$, $j \in \{1, \dots, k-1\}$ where $k = \min\{k \mid \lambda_{i+k} \in \Lambda_\theta^{a@h}\}$ and $a@h \neq a'@h'$.

In other words, an atomic read-write semantics requires that between any **obs** from $a@h$ and subsequent **ctr**, **mgrt**, or

snd from the same bigActor, no other ctr, mgrt, or snd can occur from another bigActor $a'@h'$.

THEOREM 2. Consider a BigActor execution $t = c_0 \xrightarrow{\lambda_0} c_1 \xrightarrow{\lambda_1} \dots \xrightarrow{\lambda_n} c_n$ that follows a FCFS discipline (Assumption 2.1) and the execution follows an atomic read-write semantics (Assumption 2.2). Assume that all bigActors are feedback (Assumption 2.3). Assume that η and μ are empty at c_0 (Assumption 2.4). Then c_n is correct.

PROOF. By Assumption 2.3 all bigActors are feedback. Since the feedback property is only sufficient for correct executions of single bigActor configurations we can, without loss of generality, assume that the configurations at execution t contains only two bigActors $a@h$ and $a'@h'$. By Assumption 2.1 and Assumption 2.2 the execution has no interleaving between observations and `ctr/mgrt/send` of two different bigActors. Thus, one can partition the sequence t into subsequences $t_0^{a@h}, t_0^{a'@h'}, t_1^{a@h}, t_1^{a'@h'} \dots$ where each $t_k^{a@h}$ denotes the execution of a single BigActor $a@h$ (and possibly labels regarding rules `fun`, `term`, `new` and `rcv` which are internal to bigActors and thus do not make the configuration incorrect). Since each bigActor $a@h$ is feedback, by Theorem 1, each $t_k^{a@h} = c \rightarrow \dots \rightarrow c', c'$ is correct. Thus, c_n is also correct. \square

6.2.2 Partitioning semantics

Definition 13. Consider an execution $t = c_0 \xrightarrow{\lambda_0} c_1 \xrightarrow{\lambda_1} \dots$. Let $A_i^{a@h}$ denote the area of influence of $a@h$ at c_i . An execution follows a *partitioning semantics* if for every two bigActors $a@h$ and $a'@h'$

1. $\forall i. \forall j. A_i^{a@h} \cap A_j^{a'@h'} = \emptyset$
2. $\forall i. \forall j. \forall v \in A_i^{a@h}. \forall v' \in A_j^{a'@h'}. Lks(v) \cap Lks(v') = \emptyset$

where $Lks : V_B \rightarrow \mathcal{P}(E_B \uplus Y_B)$ is a function that returns the set of links connected to ports of a given node.

In other words, partitioning semantics requires the areas of influence of any two bigActors and their links be disjoint at every time.

THEOREM 3. Consider a BigActor execution $t = c_0 \xrightarrow{\lambda_0} c_1 \xrightarrow{\lambda_1} \dots \xrightarrow{\lambda_{n-1}} c_n$ that follows a FCFS discipline (Assumption 3.1) and the execution follows a partitioning semantics (Assumption 3.2). Assume that all bigActors are feedback (Assumption 3.3). Assume that η and μ are empty at c_0 (Assumption 3.4). Then c_n is correct.

PROOF. Without loss of generality assume the execution of two bigActors $a@h$ and $a'@h'$. By Definition 13, the areas of influence and corresponding links are disjoint throughout the overall trace $t = c_0 \xrightarrow{\lambda_0} c_1 \xrightarrow{\lambda_1} \dots \xrightarrow{\lambda_{n-1}} c_n$ then for each B_i at configuration c_i can be decomposed into $B_i = B_i^{a@h} \circ B_i^{a'@h'}$ where $B_i^{a@h}$ and $B_i^{a'@h'}$ are in two different bigraph regions without sharing any edges nor names. Thus, each configuration $c_i = \langle \{a@h, a'@h'\} \mid \mu_i \mid \eta_i \mid B_i \rangle$ can be decomposed into two configurations $c_i^{a@h} = \langle \{a@h\} \mid \mu_i^{a@h} \mid \eta_i^{a@h} \mid B_i^{a@h} \rangle$ and $c_i^{a'@h'} = \langle \{a'@h'\} \mid \mu_i^{a'@h'} \mid \eta_i^{a'@h'} \mid B_i^{a'@h'} \rangle$ where $\mu_i = \mu_i^{a@h} \cup \mu_i^{a'@h'}$ and $\eta_i = \eta_i^{a@h} \cup \eta_i^{a'@h'}$. Since $a@h$ and $a'@h'$ are feedback, then by Theorem 1, there exists two traces where $c_0^{a@h} \rightarrow \dots \rightarrow c_n^{a@h}$ and $c_0^{a'@h'} \rightarrow \dots \rightarrow c_n^{a'@h'}$ where $c_n^{a@h}$ and $c_n^{a'@h'}$ are correct. Thus, c_n is correct. \square

7. CONCLUSIONS

In this paper we introduce the BigActor model for modelling structure-aware computation. The model combines the Actor model [1] and the Bigraph model [10]. We introduce an operational semantics for the BigActor model as an augmentation of the Actor model semantics. The semantics is enriched with three rules that provide means for bigActors to migrate and to observe and control the bigraph. We introduce a query language and introduce means for preventing concurrent issues that arise due to the fact that bigActors operate in a shared bigraph.

BigActor model provide means for defining agents that observe, control, and migrate over a structure of the world modelled as a bigraph. The BigActor model also introduces a model of reflection of the structure in for actor system. BigActors can also be used to embedded actors into bigraphs.

As per the future work, we are aiming at developing a theory of controllability and observability for bigActors in a control theory stand-point. We are also interested in implementing a prototype for a bigActor domain specific language for autonomous vehicles.

8. REFERENCES

- [1] G. Agha. *Actors: a model of concurrent computation in distributed systems*. MIT Press, Cambridge, MA, USA, 1986.
- [2] G. Agha, I. Mason, S. Smith, and C. Talcott. A foundation for actor computation. *Journal of Functional Programming*, 7(1):1–72, 1997.
- [3] J. Armstrong, R. Virding, C. Wikstr, M. Williams, et al. *Concurrent programming in erlang*. 1996.
- [4] L. Birkedal, S. Debois, E. Elsborg, T. Hildebrandt, and H. Niss. Bigraphical models of context-aware systems. In *Foundations of software science and computation structures*, pages 187–201. Springer, 2006.
- [5] S. Debois. Computation in the informatic jungle. *To appear*. Draft available at <http://www.itu.dk/people/debois/pubs/computation.pdf>, 2010.
- [6] J. Epstein, A. Black, and S. Peyton-Jones. Towards haskell in the cloud. In *Proceedings of the 4th ACM symposium on Haskell*, pages 118–129. ACM, 2011.
- [7] P. Haller and M. Odersky. Scala actors: Unifying thread-based and event-based programming. *Theoretical Computer Science*, 410(2-3):202–220, 2009.
- [8] E. A. Lee, S. Neuendorffer, and M. J. Wirthlin. Actor-oriented design of embedded hardware and software systems. *Journal of Circuits, Systems, and Computers*, 2002.
- [9] R. Milner. Bigraphical reactive systems. *CONCUR 2001 - Concurrency Theory*, pages 16–35, 2001.
- [10] R. Milner. *The Space and Motion of Communicating Agents*. Cambridge University Press, 2009.
- [11] B. Nielsen and G. Agha. Semantics for an actor-based real-time language. *Proceedings of the 4th International Workshop on Parallel and Distributed Real-Time Systems*, pages 223–228, 1996.
- [12] B. Nielsen, S. Ren, and G. Agha. Specification of real-time interaction constraints. *Proc. of First Int. Symposium on Object-Oriented Real-Time Computing, IEEE Computer Society*, 1998.
- [13] G. Weiss. *Multiagent systems: a modern approach to distributed artificial intelligence*. The MIT press, 1999.