

# Mobility and Security in Worldwide Computing

Robin D. Toll  
Rensselaer Polytechnic Institute  
110 8th Street  
Troy, NY 12180-3590, U.S.A.  
tollr@cs.rpi.edu

Carlos Varela  
Rensselaer Polytechnic Institute  
110 8th Street  
Troy, NY 12180-3590, U.S.A.  
cvarela@cs.rpi.edu

## ABSTRACT

Modern distributed computing requires a secure framework capable of free code mobility. In this paper, we present a simple lambda-based actor language with extensions for mobility and security, as well as the operational semantics to reason about these topics in distributed systems. Finally, we describe our preliminary implementation results.

## 1. INTRODUCTION

Internet based distributed computing systems benefit from open dynamically reconfigurable designs as these designs allow the systems to be used in constantly changing heterogeneous environments. Consider a data mining application running on an open distributed system using hundreds of thousands of computing devices to discover useful patterns in scientific data sets (e.g., protein folding, SETI, weather forecasting, etc.). If a system can make use of new computing resources as they become available and is both secure and resilient to failures in a subset of its computing nodes, it has the potential to leverage the power of idle computing resources around the world.

In this paper, we study program component mobility and security as fundamental stepping stones towards robust distributed computing systems. Mobility and security introduce new requirements on software, e.g., it is critical to devise strategies for secure and controlled distributed resource management. We specify an actor-based model and formalism to reason about program component mobility and secure resource access in worldwide computing systems.

### *Paper Outline*

In Section 2, we further motivate and informally introduce abstractions for programming worldwide computing applications. Section 3 specifies our model by providing an operational semantics for a simple actor language and extensions for mobility and security. Lastly, section 4 talks about other systems and future work.

## 2. PROGRAMMING ABSTRACTIONS FOR WORLDWIDE COMPUTING

### 2.1 Actors

The Actor model of computation is based around the concept of encapsulating state and process into a single entity. Actors are therefore inherently independent, concurrent and autonomous which enables efficiency in parallel execution [KA95] and facilitates mobility [AJ99]. Each actor is a unit of computation encapsulating data and behavior. The behavior defines how the actor reacts on receipt of a message. Each actor has a unique name, which can be used as a reference by other actors.

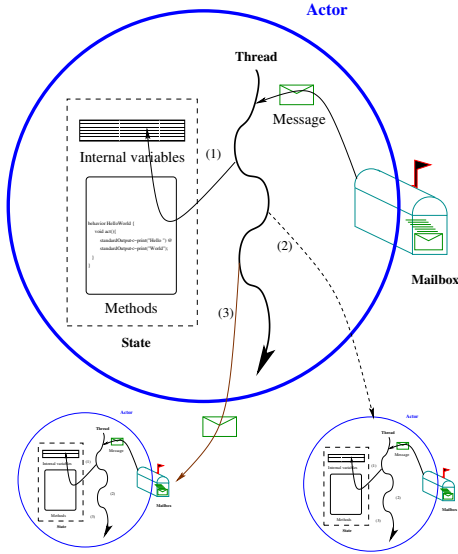
Actors only process information in reaction to messages. While processing a message, an actor can carry out any of three basic operations: alter its state, create new actors, or send messages to peer actors (see Figure 1).

Communication between actors is purely asynchronous and guaranteed. That is, when a message is sent, the model guarantees that the destination actor will receive the message; however, it does *not* guarantee the order of message arrival or, therefore, the order of processing. A side effect of this is that since actors can change their own behavior based on incoming messages, unless the actor's name is known only to the sender, its behavior could change significantly before a message arrives and is processed.

The actor model and languages provide a very useful framework for understanding and developing open distributed systems. For example, among other applications, actor systems have been used for enterprise integration [TCMW93], real-time programming [RAS96], fault-tolerance [AFPS93], and distributed artificial intelligence [FB88].

### 2.2 Universal Actors

In considering mobile computation, it becomes useful to not only model the interactions of actors with each other, but also to model the interactions of actors with their environments. In the actor model, locations are not represented, therefore, it does not matter if two actors are in the same memory space, or on two computers on opposite ends of the earth. However, when considering the problems associated with worldwide computing, it becomes important to represent the actor's environment. Otherwise it is not possible to model the behavior of an actor, e.g., when its computation environment is unreliable, or when different resources are available in different locations.



**Figure 1: Actors are reactive entities. In response to a message, an actor can (1) change its internal state, (2) create new actors, and/or (3) send messages to peer actors**

Universal actors extend actors with locations, mobility, and the concept of universal names and universal locators. Names represent actor references that do not change with actor migration. Locators represent references that enable communication with universal actors at a specific location.

An actor's location abstracts over its position relative to other actors. Each location represents an actor's run-time environment and serves as an encapsulation unit for local resources. Ubiquitous resources have a generic representation –actors keep references which get updated upon migration to resources at new locations. Actors can also keep references to non-ubiquitous resources –scarce or not generally available– by using resource attachment and detachment operations. For example, a standard output stream is ubiquitous and can always refer to the current actor's execution environment. Conversely, an actor needs to attach to a robot resource in Mars, so that the reference remains the same upon migration.

### 2.3 Secure Actors

Mobile code can pose a serious danger to any environment it executes in and, conversely, any environment can prove dangerous to actors executing inside. Therefore, it is important to consider the security of host resources and actors. Secure actors restrict communication and migration behaviors to actors within specific access control lists.

The access control list for an actor or resource contains every actor allowed to send messages to it. The access control list for a location contains every actor allowed to migrate into the location. These lists can only be altered by the resource or actor in consideration, or by a resident actor in case of passive locations. Using this method, no unprivileged actor can gain access to a resource, since any unauthorized

communication or migration request is rejected.

## 3. PROGRAMMING LANGUAGES AND SEMANTICS

### 3.1 A Simple Actor Language and Its Operational Semantics

Agha, Mason, Smith, and Talcott introduce a simple actor language as an extension to the call-by-value lambda calculus, with primitives for actor communication [AMST97].

The actor language, named here *AL*, formally defines three primitives:

- $new(b)$ , which creates an actor which has behavior  $b$  and returns the new actor's name.
- $send(v_0, v_1)$ , which sends a message with contents  $v_1$  to actor  $v_0$ .
- $ready(b')$ , which signals the end of the current execution and makes the actor ready to receive a new message using behavior  $b'$ .

### 3.2 Actor Configurations

They assume as given two sets  $At(Atoms)$  and  $X(Variables)$ , and then define the set of *values*,  $V$ , *expressions*,  $E$ , and *messages*,  $M$ , as:

$$V = At \cup X \cup \lambda X.E \cup pr(V, V)$$

$$E = V \cup app(E, E) \cup F_n(E^n)$$

$$M = \langle V \Leftarrow V \rangle$$

where  $F_n(E^n)$  is all arity- $n$  primitives.

Variables are used for actor names. At any given point, an actor can either be ready to receive a message (denoted  $ready(e)$ , where  $e$  is a lambda abstraction); or currently executing some expression  $e$ . A message sent to actor  $v_0$  with contents  $v_1$  is written as  $\langle v_0 \Leftarrow v_1 \rangle$ .

An actor configuration is a global snapshot of a group of actors. It includes the concept of an *actor mapping*, where each actor name is mapped to a behavior; a *message set* of messages in transit; a set of *receptionists* (internal actors known to the outside world); and a set of *external actors* (known actors not in the configuration).

An *actor configuration* with actor map,  $\alpha$ , multi-set of messages,  $\mu$ , receptionists,  $\rho$ , and external actors,  $\chi$ , is written<sup>1</sup>

$$\langle \alpha \mid \mu \rangle_\chi^\rho$$

where  $\rho, \chi \in \mathbf{P}_\omega[X]$ ,  $\alpha \in X \xrightarrow{f} E$ ,  $\mu \in \mathbf{M}_\omega[M]$ , and let  $A = \text{Dom}(\alpha)$ , then:

<sup>1</sup>Let  $\mathbf{P}_\omega[X]$  be the set of finite subsets (Power Set) of  $X$ ,  $\mathbf{M}_\omega[M]$  be the set of (finite) multi-sets with elements in  $M$ ,  $X_0 \xrightarrow{f} X_1$  be the set of finite maps from  $X_0 \xrightarrow{f} X_1$ ,  $\text{Dom}(f)$  be the domain of  $f$  and  $\text{FV}(e)$  be the set of free variables in  $e$ .

(0)  $\rho \subseteq A$  and  $A \cap \chi = \emptyset$ ,

(1) if  $a \in A$ , then  $\text{FV}(\alpha(a)) \subseteq A \cup \chi$ , and if  $\langle v_0 \Leftarrow v_1 \rangle \in \mu$  then  $\text{FV}(v_i) \subseteq A \cup \chi$  for  $i < 2$ .

### 3.3 Operational Semantics

We define a transition relation between actor configurations as the least relation satisfying the rules in Figure 2<sup>2</sup>. To describe the internal transitions between configurations other than message receipt, an expression is decomposed into a reduction context filled with a redex. The notation  $R[e]$  represents a redex  $e$  in a reduction context  $R$ , as described by Honsell et al. [HMST95] and used by Agha et al. [AMST97]. For a formal definition of reduction contexts, expressions with a unique hole; and for the definition of functional progress within an actor ( $\xrightarrow{\lambda}_A$ ), we refer the reader to [AMST97]. The actor redexes are: **newactor**( $e$ ), **send**( $v_0, v_1$ ), and **ready**( $v$ ).

**<fun :  $a$ >**  
 $e \xrightarrow{\lambda}_{\text{Dom}(\alpha) \cup \{a\}} e' \Rightarrow \langle \alpha\{[e]_a\} \mid \mu \rangle_{\chi}^{\rho} \mapsto \langle \alpha\{[e']_a\} \mid \mu \rangle_{\chi}^{\rho}$

**<new :  $a, a'$ >**  
 $\langle \alpha\{[R[\text{new}(e)]]_a\} \mid \mu \rangle_{\chi}^{\rho} \mapsto \langle \alpha\{[R[a']_a, [e]_{a'}]\} \mid \mu \rangle_{\chi}^{\rho} \quad a' \text{ fresh}$

**<send :  $a, v_0, v_1$ >**  
 $\langle \alpha\{[R[\text{send}(v_0, v_1)]]_a\} \mid \mu \rangle_{\chi}^{\rho} \mapsto \langle \alpha\{[R[\text{nil}]]_a\} \mid \mu \uplus \langle v_0 \Leftarrow v_1 \rangle \rangle_{\chi}^{\rho}$

**<receive :  $v_0, v_1$ >**  
 $\langle \alpha\{[R[\text{ready}(v)]]_a\} \mid \langle a \Leftarrow v_0 \rangle \uplus \mu \rangle_{\chi}^{\rho} \mapsto \langle \alpha\{[R[\text{app}(v, v_0)]]_a\} \mid \mu \rangle_{\chi}^{\rho}$

**<out :  $v_0, v_1$ >**  
 $\langle \alpha \mid \mu \uplus \langle a \Leftarrow v_0 \rangle \rangle_{\chi}^{\rho} \mapsto \langle \alpha \mid \mu \rangle_{\chi}^{\rho'}$   
 if  $a \in \chi$  and  $\rho' = \rho \cup (\text{FV}(v_0) \cap \text{Dom}(\alpha))$

**<in :  $v_0, v_1$ >**  
 $\langle \alpha \mid \mu \rangle_{\chi}^{\rho} \mapsto \langle \alpha \mid \mu \uplus \langle a \Leftarrow v_0 \rangle \rangle_{\chi \cup (\text{FV}(v_0) - \text{Dom}(\alpha))}^{\rho}$   
 if  $a \in \rho$  then  $\text{FV}(v_0) \cap \text{Dom}(\alpha) \subseteq \rho$

Figure 2: Actor Language (AL) Semantics

## 3.4 Mobile Actor Language and Its Operational Semantics

### 3.4.1 Resources

When we introduce the concept of locations, we also introduce the need to model resources in those locations. One example is the standard output stream in a run-time environment. While the concept remains the same, the actual

<sup>2</sup>We define  $\alpha\{[e]_a\}$  as an extended mapping which maps  $a$  into  $e$ , and all other actor names  $a'$  into  $\alpha(a')$ , and  $\alpha\{[e]_a, [e']_{a'}\}$  as  $\alpha\{\{[e]_a\}\{[e']_{a'}\}\}$  for  $a \neq a'$ .

implementation may change when the actor migrates across different locations.

Each resource is referred to by a universally understood resource name. This resource name is used to contact the implementing resource actor without necessarily ever knowing the specific name of the service providing that resource, or the implementation of that resource. These resources may be ubiquitous, such as a standard output stream. The name 'standard output' may apply to output on a console, or a text field on a graphical user interface, a log file, or even a printer interface, but almost any executing program has access to a primary output stream. If an actor migrates, the primary output stream changes but the transition is transparent to the actor. Referring to some standard resource name allows the environment to handle requests properly.

### 3.4.2 Resource Maps

Resource-to-actor translations are stored in resource maps. These functions are maps between global names and the names of local actors who fill a resource's roll. When an actor at a location sends a message, resource maps are searched for the target's name; first the actor's resource map followed by the location's resource map. If the target is found, the message is diverted to the resolved actor- otherwise conventional message sending proceeds.

Remote access to resources is permitted because both actors and locations have independent resource maps. Consider the case when a actor migrates but needs to retain a reference to a resource at the original location, such as a output device; the actor's map will direct messages to the original location, where the resource's name is resolved to the implementing actor.

## 3.5 Mobile Actor Language

To reflect mobility in the universal actor model, we add four primitives to *AL*, forming the Mobile Actor Language (*MAL*). *MAL* thus formally defines seven primitives:

- *new*( $b$ ), which creates an actor which has behavior  $b$  and returns the new actor's name.
- *send*( $v_0, v_1$ ), which sends a message with contents  $v_1$  to actor  $v_0$ .
- *ready*( $b'$ ), which signals the end of the current execution and makes the actor ready to receive a new message using behavior  $b'$ .
- *newloc*( $Y$ ), which indicates the appearance or creation of a new location, with an initial resource map denoted by  $Y$ .
- *migrate*( $l'$ ), which moves an actor from its current location to one denoted by  $l'$ .
- *attach*( $v$ ), which saves a resource denoted by  $v$  into the actor's resource map.
- *detach*( $v$ ), which removes a resource denoted by  $v$  from the actor's resource map.
- *register*( $v_0, v_1, l$ ), which adds the mapping of a resource name to an actor in a location.

- $deregister(v_0, v_1, l)$ , which removes the mapping of a resource name to an actor in a location.

We assume as given two sets  $At(Atoms)$  and  $X(Variables)$ , and we then define the set of *values*,  $V$ , *expressions*,  $E$ , and *messages*,  $M$ , as:

$$V = At \cup X \cup \lambda X. E \cup pr(V, V)$$

$$E = V \cup app(E, E) \cup F_n(E^n)$$

$$M = \langle V \Leftarrow V \rangle$$

where  $F_n(E^n)$  is all arity- $n$  primitives.

The set  $X$  of variables represents both actors –which includes all resources– and locations. We introduce a set  $L$ , of locations, with  $L \subseteq X$ . We also introduce a set  $R$ , of resource identifiers, with  $R \subseteq X$ . We modify the definition of  $\alpha$  so that  $\alpha \in X \xrightarrow{t} (E \times L \times (R \xrightarrow{t} X))$ .<sup>3</sup> We extend the definition of an actor configuration to include a mapping,  $\pi$ , from locations to resource maps, with  $\pi \in L \xrightarrow{t} (R \xrightarrow{t} X)$ . A universal actor configuration is written

$$\langle \alpha \mid \mu \mid \pi \rangle_{\chi}^{\rho}$$

We add another rule to the definition of actor configurations to denote that actor names and locations are disjoint<sup>4</sup>:

$$(0) \quad \rho \subseteq A \text{ and } A \cap \chi = \emptyset,$$

$$(1) \quad \text{if } a \in A, \text{ then } FV(\alpha(a)) \subseteq A \cup \chi, \text{ and if } \langle v_0 \Leftarrow v_1 \rangle \in \mu \text{ then } FV(v_i) \subseteq A \cup \chi \text{ for } i < 2.$$

$$(2) \quad Range(\alpha) \downarrow_2 \cap A = \emptyset$$

### 3.6 Operational Semantics

We define a transition relation between universal actor configurations as the least relation satisfying the rules in Figures 3 and 4.

### 3.7 Secure Mobile Actor Language

We extend the mobile actor language to form the Secure Mobile Actor Language (*SMAL*), which defines eleven primitives:

- $new(b)$ , which creates an actor which has behavior  $b$  and returns the new actor's name.
- $send(v_0, v_1)$ , which sends a message with contents  $v_1$  to actor  $v_0$ .
- $ready(b')$ , which signals the end of the current execution and makes the actor ready to receive a new message using behavior  $b'$ .

<sup>3</sup> $\alpha(a) = e \times l \times r$  implies that actor  $a$  has behavior  $e$  and is currently executing at location  $l$  with resource map  $r$ .

<sup>4</sup> $\downarrow_i$  indicates the projection of a cross product onto its  $i^{th}$  coordinate.

**<fun : a >**

$$e \xrightarrow{\lambda}_{\text{Dom}(\alpha) \cup \{a\}} e' \Rightarrow \langle \alpha\{[e, l, r]_a\} \mid \mu \mid \pi \rangle_{\chi}^{\rho} \mapsto \langle \alpha\{[e', l, r]_a\} \mid \mu \mid \pi \rangle_{\chi}^{\rho}$$

**<new : a, a' >**

$$\langle \alpha\{[R[new(e)], l, r]_a\} \mid \mu \mid \pi \rangle_{\chi}^{\rho} \mapsto \langle \alpha\{[R[a'], l, r]_a, [e, l, \emptyset]_{a'}\} \mid \mu \mid \pi \rangle_{\chi}^{\rho}$$

$a'$  fresh

**<send : a, v<sub>0</sub>, v<sub>1</sub> >**

$$\langle \alpha\{[R[send(v_0, v_1)], l, r]_a\} \mid \mu \mid \pi \rangle_{\chi}^{\rho} \mapsto \langle \alpha\{[R[nil], l, r]_a\} \mid \mu \uplus \langle v_2 \Leftarrow v_1 \rangle \mid \pi \rangle_{\chi}^{\rho}$$

if  $r(v_0) = v_2$

$$\langle \alpha\{[R[nil], l, r]_a\} \mid \mu \uplus \langle v_3 \Leftarrow v_1 \rangle \mid \pi \rangle_{\chi}^{\rho}$$

if  $\pi(l)(v_0) = v_3$

$$\langle \alpha\{[R[nil], l, r]_a\} \mid \mu \uplus \langle v_0 \Leftarrow v_1 \rangle \mid \pi \rangle_{\chi}^{\rho}$$

if  $v_0 \notin \text{Dom}(r) \ \& \ v_0 \notin \text{Dom}(\pi)$

**<receive : v<sub>0</sub>, v<sub>1</sub> >**

$$\langle \alpha\{[R[ready(v)], l, r]_{v_0}\} \mid \langle v_0 \Leftarrow v_1 \rangle \uplus \mu \mid \pi \rangle_{\chi}^{\rho} \mapsto \langle \alpha\{[R[app(v, v_1)], l, r]_{v_0}\} \mid \mu \mid \pi \rangle_{\chi}^{\rho}$$

**Figure 3: MAL Semantics, Part I**

- $newloc(Y)$ , which indicates the appearance or creation of a new location, with an initial resource map denoted by  $Y$ .
- $migrate(l)$ , which moves an actor from its current location to one denoted by  $l$ .
- $attach(v)$ , which saves a resource denoted by  $v$  into the actor's resource map.
- $detach(v)$ , which removes a resource denoted by  $v$  from the actor's resource map.
- $register(v_0, v_1, l)$ , which adds the mapping of a resource name to an actor in a location.
- $deregister(v_0, v_1, l)$ , which removes the mapping of a resource name to an actor in a location.
- $allow(v)$  changes the actor's access list to include actor  $v$ .
- $allowloc(v)$  changes the actor location's access list to include actor  $v$ .
- $disallow(v)$  changes the actor's access list to exclude actor  $v$ .
- $disallowloc(v)$  changes the actor location's access list to exclude actor  $v$ .

$\langle \text{out} : v_0, v_1 \rangle$   
 $\langle \alpha \mid \mu \uplus \langle a \Leftarrow v_0 \rangle \mid \pi \rangle_{\chi}^{\rho} \mapsto \langle \alpha \mid \mu \mid \pi \rangle_{\chi}^{\rho'}$   
 if  $a \in \chi$ , and  $\rho' = \rho \cup (\text{FV}(v_0) \cap \text{Dom}(\alpha))$

$\langle \text{in} : v_0, v_1 \rangle$   
 $\langle \alpha \mid \mu \mid \pi \rangle_{\chi}^{\rho} \mapsto \langle \alpha \mid \mu \uplus \langle a \Leftarrow v_0 \rangle \mid \pi \rangle_{\chi \cup (\text{FV}(v_0) - \text{Dom}(\alpha))}^{\rho}$   
 if  $a \in \rho$  then  $\text{FV}(v_0) \cap \text{Dom}(\alpha) \subseteq \rho$

$\langle \text{migrate} : l' \rangle$   
 $\langle \alpha \{ [R[\text{migrate}(l')], l, r]_a \} \mid \mu \mid \pi \rangle_{\chi}^{\rho} \mapsto \langle \alpha \{ [R[\text{nil}], l', r]_a \} \mid \mu \mid \pi \rangle_{\chi}^{\rho}$   
 $l' \in \text{Dom}(\pi)$

$\langle \text{newloc} : Y \rangle$   
 $\langle \alpha \{ [R[\text{newloc}(Y)], l, r]_a \} \mid \mu \mid \pi \rangle_{\chi}^{\rho} \mapsto \langle \alpha \{ [R[l'], l, r]_a \} \mid \mu \mid \pi \cup (l' \rightarrow Y) \rangle_{\chi}^{\rho}$   
 $l$  fresh,  $Y \in (X \rightarrow X)$

$\langle \text{attach} : a' \rangle$   
 $\langle \alpha \{ [R[\text{attach}(a')], l, r]_a \} \mid \mu \mid \pi \rangle_{\chi}^{\rho} \mapsto \langle \alpha \{ [R[\text{nil}], l, r \cup (a' \rightarrow a'')]_a \} \mid \mu \mid \pi \rangle_{\chi}^{\rho}$   
 if  $(a' \rightarrow a'') \in \pi(l)$

$\langle \text{detach} : a' \rangle$   
 $\langle \alpha \{ [R[\text{detach}(a')], l, r \cup (a' \rightarrow a'')]_a \} \mid \mu \mid \pi \rangle_{\chi}^{\rho} \mapsto \langle \alpha \{ [R[\text{nil}], l, r]_a \} \mid \mu \mid \pi \rangle_{\chi}^{\rho}$

$\langle \text{register} : a', a'', l' \rangle$   
 $\langle \alpha \{ [R[\text{register}(a', a'', l')], l, r]_a \} \mid \mu \mid \pi \rangle_{\chi}^{\rho} \mapsto$   
 $\langle \alpha \{ [R[\text{nil}], l, r \cup (a' \rightarrow a'')]_a \} \mid \mu \mid \pi \cup (l' \rightarrow (a' \rightarrow a'')) \rangle_{\chi}^{\rho}$

$\langle \text{unregister} : a', a'', l' \rangle$   
 $\langle \alpha \{ [R[\text{unregister}(a', a'', l')], l, r]_a \} \mid \mu \mid \pi \cup (l' \rightarrow (a' \rightarrow a'')) \rangle_{\chi}^{\rho} \mapsto$   
 $\langle \alpha \{ [R[\text{nil}], l, r]_a \} \mid \mu \mid \pi \rangle_{\chi}^{\rho}$

Figure 4: MAL Semantics, Part II

The *allow* and *allowloc* primitives can receive as an argument a null access list, represented by  $\perp$ . This indicates the absence of restrictions on messaging and migration.

We assume as given two sets  $At(\text{Atoms})$  and  $X(\text{Variables})$ , and we then define the set of *values*,  $V$ , *expressions*,  $E$ , and *messages*,  $M$ , as:

$$V = At \cup X \cup \lambda X. E \cup pr(V, V)$$

$$E = V \cup app(E, E) \cup F_n(E^n)$$

$$M = \langle V \Leftarrow V \rangle_X$$

where  $F_n(E^n)$  is all arity- $n$  primitives.

The set of variables,  $X$ , including resource names,  $R$ , and locations,  $L$ , remains the same. The structure of  $M$  allows selection of valid senders via access control lists. We define a set of access control lists as  $ACL \in \mathbf{P}_{\omega}[\text{Dom}(\alpha) \cup \{\perp\}]$ . We alter the actor configuration definition to change the

actor and location maps,  $\alpha$  and  $\pi$  respectively. The actor configuration is written

$$\langle \alpha \mid \mu \mid \pi \rangle_{\chi}^{\rho}$$

where  $\rho$ ,  $\mu$  and  $\chi$  are as in *MAL*,  $\alpha \in X^{\dagger} \rightarrow (E \times L \times (R^{\dagger} \rightarrow X) \times ACL)$ , and  $\pi \in L^{\dagger} \rightarrow ((R^{\dagger} \rightarrow X) \times ACL)$ . Changes in  $\alpha$  and  $\pi$  reflect the inclusion of access control lists for actors and locations.<sup>5</sup>

### 3.8 Operational Semantics

We define a transition relation between secure universal actor configurations as the least relation satisfying the rules in Figures 5, 6 and 7.

## 4. DISCUSSION AND FUTURE WORK

<sup>5</sup> $\alpha(a) = e \times l \times r \times c$  implies that actor  $a$  has behavior  $e$  and is currently executing at location  $l$  with resource map  $r$  and access control list  $c$ .

$\langle \text{migrate} : l' \rangle$   
 $\langle \alpha\{[R[\text{migrate}(l')], l, r, c]_a\} \mid \mu \mid \pi \uplus [c', r']_{l'} \rangle_{\chi}^{\rho} \mapsto$   
 $\langle \alpha\{[R[\text{nil}], l', r, c]_a\} \mid \mu \mid \pi \uplus [c', r']_{l'} \rangle_{\chi}^{\rho}$   
 if  $a \in c'$  or  $c' = \perp$

$\langle \text{newloc} : Y \rangle$   
 $\langle \alpha\{[R[\text{newloc}(Y)], l, r, c]_a\} \mid \mu \mid \pi \rangle_{\chi}^{\rho} \mapsto$   
 $\langle \alpha\{[R[l'], l, r, c]_a\} \mid \mu \mid \pi \uplus [\{a\}, Y]_{l'} \rangle_{\chi}^{\rho}$   
 $l'$  fresh

$\langle \text{attach} : a' \rangle$   
 $\langle \alpha\{[R[\text{attach}(a')], l, r, c]_a\} \mid \mu \mid \pi \rangle_{\chi}^{\rho} \mapsto$   
 $\langle \alpha\{[R[\text{nil}], l, r \cup (a' \rightarrow a''), c]_a\} \mid \mu \mid \pi \uplus [c', r' \cup (a' \rightarrow a'')]_{l'} \rangle_{\chi}^{\rho}$

$\langle \text{detach} : a' \rangle$   
 $\langle \alpha\{[R[\text{detach}(a')], l, r \cup (a' \rightarrow a''), c]_a\} \mid \mu \mid \pi \rangle_{\chi}^{\rho} \mapsto$   
 $\langle \alpha\{[R[\text{nil}], l, r, c]_a\} \mid \mu \mid \pi \rangle_{\chi}^{\rho}$

$\langle \text{register} : a', a'', l' \rangle$   
 $\langle \alpha\{[R[\text{register}(a', a'', l')], l, r]_a\} \mid \mu \mid \pi \uplus [c', r']_{l'} \rangle_{\chi}^{\rho} \mapsto$   
 $\langle \alpha\{[R[\text{nil}], l, r \cup (a' \rightarrow a'')]_a\} \mid \mu \mid \pi \uplus [c', r' \cup (a' \rightarrow a'')]_{l'} \rangle_{\chi}^{\rho}$   
 if  $(a' \rightarrow a'') \in r$

$\langle \text{unregister} : a', a'', l' \rangle$   
 $\langle \alpha\{[R[\text{unregister}(a', a'', l')], l, r]_a\} \mid \mu \mid \pi [c', r' \cup (a' \rightarrow a'')]_{l'} \rangle_{\chi}^{\rho} \mapsto$   
 $\langle \alpha\{[R[\text{nil}], l, r]_a\} \mid \mu \mid \pi \rangle_{\chi}^{\rho}$   
 if  $(a' \rightarrow a'') \in r$

**Figure 6: SMAL Semantics, Part II: Mobility**

The actor model was first created by Hewitt and his group at MIT [Hew77] in the late 1970s. The model has been further developed by Agha [Agh86] and his group at UIUC. Agha, Mason, Smith and Talcott [AMST97] have developed a simple actor language as an extension to the lambda calculus, its operational semantics and they have studied a family of equivalence relations on actor expressions. Talcott has developed an interaction semantics for actor systems [Tal96, Tal98]. These forms of actor semantics have been the basis of many studies on extensions to the actor model (e.g., for coordination [FA93, VA99, AJV01, FV02], real-time [RAS96], software architectures [AA98], fault-tolerance [SA94], adaptive and meta-level architectures [VTA01], and artificial intelligence [AJ99]).

Several research groups have been trying to achieve distributed computing on a large scale. Berkeley's NOW project has been effectively distributing computation in a "building-wide" scale [ACP95], and Berkeley's Millennium project is exploiting a hierarchical cluster structure to provide distributed computing on a "campus-wide" scale [BGC98]. The Globus project seeks to enable the construction of larger *computational grids* [FK98]. Caltech's Infospheres project

has a vision of a worldwide pool of millions of objects (or agents) much like the pool of documents on the World-Wide Web today [CRS<sup>+</sup>96]. WebOS seeks to provide operating system services, such as client authentication, naming, and persistent storage, to wide area applications [VAD<sup>+</sup>98]. UIUC's 2K is an integrated operating system architecture addressing the problems of resource management in heterogeneous networks, dynamic adaptability, and configuration of component-based distributed applications [KCMN99].

Security for distributed systems has been looked at for a number of other agent systems. Safe Mobile Ambients [LS00] restrict mobile ambients [CG98] so that sensitive operations such as entering, exiting, and opening an ambient are performed with common agreement. While safe ambients preserve the expressibility of mobile ambients, they prevent programming mistakes by controlling undesirable *grave* interferences. The Seal calculus [VC98], resembles ambients, with two important exceptions. First, *seals* can only move with the environment's control, and the  $\pi$ -calculus is used as a basis for computation, rather than mobility itself. Other methods have been used for securing code on different levels. Ideas such as proof-carrying code [Nec97] and stack inspec-

$\langle \text{allow} : a' \rangle$   
 $\langle \alpha\{[R[\text{allow}(a')], l, r, c]_a\} \mid \mu \mid \pi \rangle_X^{\rho} \mapsto \langle \alpha\{[R[\text{nil}], l, r, c \cup \{a'\}]_a\} \mid \mu \mid \pi \rangle_X^{\rho}$

$\langle \text{allow} : \perp \rangle$   
 $\langle \alpha\{[R[\text{allow}(\text{nil})], l, r, c]_a\} \mid \mu \mid \pi \rangle_X^{\rho} \mapsto \langle \alpha\{[R[\text{nil}], l, r, \perp]_a\} \mid \mu \mid \pi \rangle_X^{\rho}$

$\langle \text{allowloc} : a' \rangle$   
 $\langle \alpha\{[R[\text{allowloc}(a')], l, r, c]_a\} \mid \mu \mid \pi \uplus [c', r']_l \rangle_X^{\rho} \mapsto$   
 $\langle \alpha\{[R[\text{nil}], l, r, c]_a\} \mid \mu \mid \pi \uplus [c' \cup \{a'\}, r']_l \rangle_X^{\rho}$

$\langle \text{allowloc} : \perp \rangle$   
 $\langle \alpha\{[R[\text{allowloc}(\text{nil})], l, r, c]_a\} \mid \mu \mid \pi \uplus [c', r']_l \rangle_X^{\rho} \mapsto$   
 $\langle \alpha\{[R[\text{nil}], l, r, c]_a\} \mid \mu \mid \pi \uplus [\perp, r']_l \rangle_X^{\rho}$

$\langle \text{disallow} : a' \rangle$   
 $\langle \alpha\{[R[\text{disallow}(a')], l, r, c \cup \{a'\}]_a\} \mid \mu \mid \pi \rangle_X^{\rho} \mapsto$   
 $\langle \alpha\{[R[\text{nil}], l, r, c]_a\} \mid \mu \mid \pi \rangle_X^{\rho}$

$\langle \text{disallow} : \perp \rangle$   
 $\langle \alpha\{[R[\text{disallow}(\text{nil})], l, r, c]_a\} \mid \mu \mid \pi \rangle_X^{\rho} \mapsto$   
 $\langle \alpha\{[R[\text{nil}], l, r, \emptyset]_a\} \mid \mu \mid \pi \rangle_X^{\rho}$

$\langle \text{disallowloc} : a' \rangle$   
 $\langle \alpha\{[R[\text{disallowloc}(a')], l, r, c]_a\} \mid \mu \mid \pi \uplus [c' \cup \{a'\}, r']_l \rangle_X^{\rho} \mapsto$   
 $\langle \alpha\{[R[\text{nil}], l, r, c]_a\} \mid \mu \mid \pi \uplus [c', r']_l \rangle_X^{\rho}$

$\langle \text{disallowloc} : \perp \rangle$   
 $\langle \alpha\{[R[\text{disallowloc}(\text{nil})], l, r, c]_a\} \mid \mu \mid \pi \uplus [c', r']_l \rangle_X^{\rho} \mapsto$   
 $\langle \alpha\{[R[\text{nil}], l, r, c]_a\} \mid \mu \mid \pi \uplus [\emptyset, r']_l \rangle_X^{\rho}$

Figure 7: SMAL Semantics, Part III: Access Control

tion [WF98] are methods of protecting hosts, also discussed in [ST98].

While there are excellent algorithms for load balancing in clusters and other more static environments, e.g., random stealing and cluster-aware random stealing [vNKB01], the dynamic and heterogeneous nature of the nodes on the WWC make such algorithms much less efficient, especially when IO's peer-to-peer nature is taken into account. Currently, IO's peer-to-peer network is a variant of Gnutella [Cli00]; however, in the future, implementing IO on top of an already existing peer-to-peer network such as JXTA[TAD<sup>+</sup>02] may prove to be a more interesting option.

## 5. ACKNOWLEDGEMENTS

We would like to thank the Worldwide Computing group at RPI for input on these semantics. We would especially like to thank Carolyn Talcott for her help in finalizing the semantics.

## 6. REFERENCES

- M. Astley and G. Agha. Modular construction and composition of distributed software architectures. In *Proceedings of the International Symposium on Software Engineering for Parallel and Distributed Systems (PDSE '98)*, 1998.
- Thomas E. Anderson, David E. Culler, and David A. Patterson. A Case for Networks of Workstations: NOW. *IEEE Micro*, February 1995.
- G. Agha, S. Frølund, R. Panwar, and D. Sturman. A linguistic framework for dynamic composition of dependability protocols. In *Dependable Computing for Critical Applications III*, pages 345–363. International Federation of Information Processing Societies (IFIP), Elsevier Science Publisher, 1993.
- G. Agha. *Actors: A Model of Concurrent Computation in Distributed Systems*. MIT

- Press, 1986.
- G. Agha and N. Jamali. Concurrent programming for distributed artificial intelligence. In G. Weiss, editor, *Multiagent Systems: A Modern Approach to DAI.*, chapter 12. MIT Press, 1999.
- G. Agha, N. Jamali, and C. Varela. Agent Naming and Coordination: Actor Based Models and Infrastructures. In A. Ominici, F. Zambonelli, M. Klusch, and R. Tolksdorf, editors, *Coordination of Internet Agents*, chapter 9, pages 225–248. Springer-Verlag, 2001.
- G. Agha, I. A. Mason, S. F. Smith, and C. L. Talcott. A foundation for actor computation. *Journal of Functional Programming*, 7:1–72, 1997.
- Philip Buonadonna, Andrew Geweke, and David E. Culler. An implementation and analysis of the virtual interface architecture. In *Proceedings of Supercomputing '98*, Orlando, FL, November 1998.
- L. Cardelli and A.D. Gordon. Mobile ambients. In *Foundations of System Specification and Computational Structures*, LNCS 1378, pages 140–155. Springer Verlag, 1998.
- Clip2.com. The gnutella protocol specification v0.4, 2000.
- K. M. Chandy, A. Rifkin, P. A. G. Sivilotti, J. Mandelson, M. Richardson, W. Tanaka, and L. Weisman. A World-Wide Distributed System Using Java and the Internet. In *Proceedings of the Fifth IEEE International Symposium on High Performance Distributed Computing*, New York, U.S.A., Aug 1996.
- S. Frølund and G. Agha. A language framework for multi-object coordination. In *Proceedings of ECOOP 1993*. Springer Verlag, 1993. LNCS 707.
- J. Ferber and J. Briot. Design of a concurrent language for distributed artificial intelligence. In *Proceedings of the International Conference on Fifth Generation Computer Systems*, volume 2, pages 755–762. Institute for New Generation Computer Technology, 1988.
- I. Foster and C. Kesselman. The Globus Project: A Status Report. In J. Antonio, editor, *Proceedings of the Seventh Heterogeneous Computing Workshop (HCW '98)*, pages 4–18. IEEE Computer Society, March 1998.
- John Field and Carlos Varela. Toward a programming model for building reliable systems with distributed state. In *Proceedings of the First International Workshop on Foundations of Coordination Languages and Software Architectures (FOCLASA)*, Brno, Czech Republic, August 2002.
- C. Hewitt. Viewing control structures as patterns of passing messages. *Journal of Artificial Intelligence*, 8-3:323–364, June 1977.
- Furio Honsell, Ian A. Mason, Scott F. Smith, and Carolyn L. Talcott. A variable typed logic of effects. *Information and Computation*, 119(1):55–90, 1995.
- W. Kim and G. Agha. Efficient Support of Location Transparency in Concurrent Object-Oriented Programming Languages. In *Proceedings of Supercomputing'95*, 1995.
- F. Kon, R. Campbell, M. Dennis Mickunas, and K. Nahrstedt. 2K: A Distributed Operating System for Dynamic Heterogeneous Environments. Technical report, Department of Computer Science, University of Illinois at Urbana-Champaign, December 1999.
- Francesca Levi and Davide Sangiorgi. Controlling interference in ambients. In *Symposium on Principles of Programming Languages*, pages 352–364, 2000.
- George C. Necula. Proof-carrying code. In *Proceedings of the 24th ACM SIGPLAN-SIGACT Symposium on Principles of Programming Languages (POPL '97)*, pages 106–119, Paris, January 1997.
- S. Ren, G. A. Agha, and M. Saito. A modular approach for programming distributed real-time systems. *Journal of Parallel and Distributed Computing*, 36:4–12, 1996.
- D. C. Sturman and G. Agha. A protocol description language for customizing failure semantics. In *Proceedings of the 13th Symposium on Reliable Distributed Systems*. IEEE Computer Society Press, October 1994.
- Tomas Sander and Christian F. Tschudin. Protecting mobile agents against malicious hosts. *Lecture Notes in Computer Science*, 1419:44–??, 1998.
- Bernard Traversat, Mohamed Abdelaziz, Mike Duigou, Jean-Christophe Hugly, Eric Pouyoul, and Bill Yeager. Project jxta virtual network, February 2002.
- C. Talcott. Interaction semantics for components of distributed systems. In *First IFIP workshop on Formal Methods for Open Object-based Distributed Systems (FMOODS '96)*, Paris, France, March 1996.
- C. L. Talcott. Composable semantic models for actor theories. *Higher-Order and Symbolic Computation*, 11(3), 1998.



C. Tomlinson, P. Cannata, G. Meredith, and D. Woelk. The extensible services switch in Carnot. *IEEE Parallel and Distributed Technology*, 1(2):16–20, May 1993.

C. Varela and G. Agha. A Hierarchical Model for Coordination of Concurrent Activities. In P. Ciancarini and A. Wolf, editors, *Third International Conference on Coordination Languages and Models (COORDINATION '99)*, LNCS 1594, pages 166–182, Berlin, April 1999. Springer-Verlag.  
<http://osl.cs.uiuc.edu/Papers/Coordination99.ps>.

Amin Vahdat, Thomas Anderson, Michael Dahlin, David Culler, Eshwar Belani, Paul Eastham, and Chad Yoshikawa. WebOS: Operating System Services For Wide Area Applications. In *Proceedings of the Seventh IEEE Symposium on High Performance Distributed Computing*, July 1998.

Jan Vitek and Giuseppe Castagna. Seal: A framework for secure mobile computations. In *ICCL Workshop: Internet Programming Languages*, pages 47–77, 1998.

Rob V. van Nieuwpoort, Thilo Kielmann, and Henri E. Bal. Efficient load balancing for wide-area divide-and-conquer applications. *ACM*, 36:34–43, 2001.

Nalini Venkatasubramanian, Carolyn Talcott, and Gul A. Agha. A formal model for reasoning about adaptive qos-enabled middleware. In *Formal Methods Europe (FME 2001)*. Humboldt-Universitt, Berlin, Germany, March 2001.

Dan S. Wallach and Edward W. Felten. Understanding Java stack inspection. In *19th IEEE Symposium on Security and Privacy*, pages 52–63, 1998.

```

<fun : a>
  e  $\xrightarrow{\lambda}_{\text{Dom}(\alpha) \cup \{a\}}$  e'  $\Rightarrow \langle \alpha\{[e, l, r, c]_a\} \mid \mu \mid \pi \rangle_{\mathcal{X}}^{\rho} \mapsto$ 
     $\langle \alpha\{[e', l, r, c]_a\} \mid \mu \mid \pi \rangle_{\mathcal{X}}^{\rho}$ 
<new : a, a'>
   $\langle \alpha\{[R[\text{new}(e)], l, r, c]_a\} \mid \mu \mid \pi \rangle_{\mathcal{X}}^{\rho} \mapsto$ 
     $\langle \alpha\{[R[a'], l, r, c]_a, [e, l, r, \{a, a'\}]_{a'}\} \mid \mu \mid \pi \rangle_{\mathcal{X}}^{\rho}$ 
    a' fresh
<send : a, v0, v1>
   $\langle \alpha\{[R[\text{send}(v_0, v_1)], l, r, c]_a\} \mid \mu \mid \pi \rangle_{\mathcal{X}}^{\rho} \mapsto$ 
     $\langle \alpha\{[R[\text{nil}], l, r, c]_a\} \mid \mu \uplus \langle v_2 \Leftarrow v_1 \rangle_a \mid \pi \rangle_{\mathcal{X}}^{\rho}$ 
    if r(v0)=v2
     $\langle \alpha\{[R[\text{nil}], l, r, c]_a\} \mid \mu \uplus \langle v_3 \Leftarrow v_1 \rangle_a \mid \pi \rangle_{\mathcal{X}}^{\rho}$ 
    if  $\pi(l)(v_0)=v_3$ 
     $\langle \alpha\{[R[\text{nil}], l, r, c]_a\} \mid \mu \uplus \langle v_0 \Leftarrow v_1 \rangle_a \mid \pi \rangle_{\mathcal{X}}^{\rho}$ 
    if v0  $\notin$  Dom( $\pi$ ) & v0  $\notin$  Dom( $\pi$ )
<receive : v0, v1>
   $\langle \alpha\{[R[\text{ready}(v)], l, r, c]_{v_0}\} \mid \langle v_0 \Leftarrow v_1 \rangle_a \uplus \mu \mid \pi \rangle_{\mathcal{X}}^{\rho} \mapsto$ 
     $\langle \alpha\{[R[\text{app}(v, v_1)], l, r, c]_{v_0}\} \mid \mu \mid \ell \rangle_{\mathcal{X}}^{\rho}$ 
    if a  $\in$  c or c =  $\perp$ 
<out : v0, v1>
   $\langle \alpha \mid \mu \uplus \langle a \Leftarrow v_0 \rangle_a \mid \pi \rangle_{\mathcal{X}}^{\rho} \mapsto \langle \alpha \mid \mu \mid \pi \rangle_{\mathcal{X}}^{\rho'}$ 
    if a  $\in$   $\mathcal{X}$ , and  $\rho' = \rho \cup (\text{FV}(v_0) \cap \text{Dom}(\alpha))$ 
<in : v0, v1>
   $\langle \alpha \mid \mu \mid \ell \rangle_{\mathcal{X}}^{\rho} \mapsto$ 
     $\langle \alpha \mid \mu \uplus \langle a \Leftarrow v_0 \rangle_{a'} \mid \pi \rangle_{\mathcal{X} \cup (\text{FV}(v_0) - \text{Dom}(\alpha))}^{\rho}$ 
    if a  $\in$   $\rho$  then  $\text{FV}(v_0) \cap \text{Dom}(\alpha) \subseteq \rho$ 

```

Figure 5: SMAL Semantics, Part I: Actor Semantics