

A Foundation for Actor Computation

Gul Agha

University of Illinois agha@cs.uiuc.edu

Ian A. Mason

Stanford University iam@sail.stanford.edu

Scott F. Smith

The Johns Hopkins University scott@cs.jhu.edu

Carolyn L. Talcott

Stanford University clt@sail.stanford.edu

Abstract

We present an actor language which is an extension of a simple functional language, and provide a precise operational semantics for this extension. Actor configurations represent open distributed systems, by which we mean that the specification of an actor system explicitly takes into account the interface with external components. We study the composability of such systems. We define and study various notions of testing equivalence on actor expressions and configurations. The model we develop provides fairness. An important result is that the three forms of equivalence, namely, convex, must, and may equivalences, collapse to two in the presence of fairness. We further develop methods for proving laws of equivalence and provide example proofs to illustrate our methodology.

Contents

1	Introduction	2
1.1	Overview	3
1.2	Related Work	4
2	Actor Language Constructs	7
2.1	Trivial Examples	8
2.2	Actor Cells	9
2.3	Join Continuations	9
2.4	Notation.	12
3	A Simple Lambda-Based Actor Language	12
3.1	Syntax	13
3.2	Reduction Semantics for Actor Configurations	13
3.3	Composition of Actor Configurations	19
4	Equivalence of Expressions	20
4.1	Events	20
4.2	Three Equivalences	21

4.3	Partial Collapse	22
4.4	Equivalence of Configurations	24
5	Laws of Expression Equivalence	25
5.1	Functional Laws	25
5.2	Basic Laws for Actor Primitives	27
6	Proving Expression Equivalence	30
6.1	The General Method	32
6.2	Common Reduct Case	34
6.3	Equivalence by Two Stage Reduction	47
6.4	Equivalence of Reduction Contexts	50
7	Discussion	53
7.1	Future Directions	53
	References	54
8	Index of Notations (Not including §5)	57

1 Introduction

The modern computing environment is becoming increasingly open and distributed. Research on semantics to support reasoning about components and configuration in *open distributed systems* is in its early stages. The main characteristics of an open distributed system are that such systems allow the addition of new components, the replacement of existing components, and changes in interconnections between components, largely without disturbing the functioning of the system. Open distributed systems require a discipline in which a component may not have any direct control over other components with which it is connected. Instead, the behavior of a component is locally determined by its initial state and the history of its interactions with the environment. Moreover, interactions between components may occur only through their interfaces. Thus, the internal state of a component must only be accessible through operations provided by the interface.

The actor model of computation has a built-in notion of local component and interface, and thus it is a natural model to use as a basis for a theory of open distributed computation. Specifically, we view actors as a model of coordination between autonomous interacting components. The local computation carried out by the components may be specified in any sequential language. However, we carry out the development of an actor semantics in a framework where local computation is specified using a functional language. We extend a functional programming language with actor coordination primitives to model open distributed systems. Our semantic theory uses techniques first developed in a functional setting. The transition system operational semantics extends the reduction system semantics of the functional language. The notion of observational equivalence studied generalizes the now standard notion of operational equivalence for functional languages. The resulting equational theory embodies that of the computational lambda calculus and preserves many of the advantages provided by functional programming for reasoning about programs and program transformations. The actor language we

study provides an alternative approach to concurrent extensions of functional languages that is object based rather than process based. Our components are reactive instead of active processes more common in other models of concurrency. Active processes correspond to a thread of control that disappears when execution of the thread is complete; unlike reactive processes, activity in them is not initiated by the receipt of a message. Note that receiving a message is analogous to a function being invoked.

1.1 Overview

In this paper we present a study of a particular actor language. Our actor language is an extension of a simple call-by-value functional language (lambda calculus plus arithmetic primitives and structure constructors, recognizers, and destructors) by primitives for creating and manipulating actors. Our approach is motivated by a desire to bridge the gap between theory and practice. The semantic theory we develop is intended to be useful for justifying program transformations for real languages, and to formalize intuitive arguments and properties used by programmers.

In our model we make explicit the notion of open system component through the notion of an *actor configuration*. An actor configuration is a collection of individually named concurrently executing actors, a collection of messages in route, a set of receptionist names and a set of external actor names. Receptionists are the externally visible actors of the configuration, and names of external actors are references to actors outside the configuration. The receptionist and external actor names explicitly define the interface to the environment.

A common criticism of the actor model of computation is that actors do not compose. However actor configurations can be combined in a variety of ways to form new configurations. As a first step towards an algebra of operations on actor configurations, we define a composition operator on configurations. Composition is associative, commutative, and has a unit. This allows large complex configurations to be studied in parts and composed to form the full system. Unlike most notions of modularity and composability, which are static, the notions we define are fundamentally dynamic ones that allow for the interface between components to evolve over time.

Following the tradition of (Morris, 1968; Plotkin, 1975; Mason and Talcott, 1991; Felleisen and Hieb, 1992; Felleisen and Wright, 1991) we develop the semantics in two stages. The first stage consists in giving an operational semantics for actor configurations. In the second stage various notions of equivalence are investigated, both of expressions and of configurations. The operational semantics extends that of the embedded functional language in such a way that the equational theory of the functional language is preserved. In particular the equational laws of the computational lambda calculus (Moggi, 1988) as well as the usual laws for pairing and arithmetic hold. This provides a basis for a rich set of equational reasoning principles. There are also numerous equational laws that relate to actor computations, for instance allowing two adjacent message send operations to be permuted.

The operational semantics of actor configurations is defined by a transition re-

lation on configurations. An important aspect of the actor model is the fairness requirement: message delivery is guaranteed, and individual actor computations are guaranteed to progress. We make the fairness requirement explicit in our semantics by requiring infinite sequences of transitions on actor configurations to be fair. Although fairness makes some aspects of reasoning more complicated, it simplifies others. Certain classes of intuitively obvious equations fail to hold without the fairness requirement. The assumption of fairness allows equational specification of some liveness properties, and it is particularly important for modular reasoning about liveness properties.

Our equational theory is based on the notion of observational equivalence. Two expressions/configurations are said to be observationally equivalent if they give rise to the same observations, suitably defined, inside all observing contexts. This notion is closely related to testing equivalence (de Nicola and Hennessy, 1984). We prove that in the presence of fairness, the three standard notions of observational equivalence collapse to two. Observational equivalence provides a semantic basis for developing sound transformation rules for expressions.

Outline

The remainder of this paper is organized as follows. The rest of this section describes previous work on Actors and related models of concurrent computation. §2 gives the syntax and operational semantics of our actor language, and describes a composability result. In §3 we study the notion of observational equivalence for actor expressions. In §4 we state a variety of basic equational laws along with an intuitive explanation of how these laws are established. In §5 we develop methods for establishing expression equivalence, and use these methods to prove the laws of §4. §6 summarizes the highlights of this paper, and discusses future work.

1.2 Related Work

We discuss related work in actors, process algebras and concurrent functional languages.

1.2.1 Research on Actors

We may briefly summarize the principles underlying the actor model of computation that we use as follows. Actors are self-contained, concurrently interacting entities of a computing system. They communicate via message passing which is asynchronous and fair. Actors can be dynamically created and the topology of actor systems can change dynamically. The actor model is a primitive model of computation which nonetheless easily expresses a wide range of computation paradigms. It directly supports encapsulation and sharing, and provides a natural extension of both functional programming and object style data abstraction to concurrent open systems. (See (Agha, 1986a; Agha, 1990; Agha et al., 1993a) for more discussion of the actor model, and for many examples of programming with actors.

The actor model was originally proposed by Hewitt and the meaning of the term has evolved over time in the work of Hewitt and associates. We briefly describe the history of actor research, necessarily omitting some of the significant work.

In his early work on planner (Hewitt, 1971), Hewitt used the term actor to describe active entities which, unlike functions, went around looking for patterns to match in order to trigger activity. This concept was later developed into the scientific community metaphor where *sprites* examined facts and added to them in a monotonically growing knowledge base (Kornfeld and Hewitt, 1981). In (Hewitt et al., 1973), the notion of actors was close to that of an agent in Distributed Artificial Intelligence: actors have intentions, resources, contain message monitors and a scheduler. Irene Greif (Greif, 1975) developed an abstract model of actors in terms of event diagrams which recorded local events at each actor and the causal relations between events.

Baker and Hewitt (Baker and Hewitt, 1977) then formalized a set of axioms for concurrent computation which stated properties that events in actor systems must obey in order to prevent causality violations. The work in (Hewitt, 1977) contains the insight that the usual control structures could be represented as patterns of message passing between simple actors which had a conditional construct but no local state. It provided a basis for the continuation passing style of programming as found in Scheme (Abelson and Sussman, 1985).

In (Hewitt and Atkinson, 1979), the concept of serializer is described: a serializer localizes conditions for resumption of waiting processes and thus improves on monitors which require explicit signaling of dormant processes. A related notion, namely, that of *guardians*, was defined in (Attardi and Hewitt, 1978). A guardian regulates the use of shared resources, scheduling their access and providing protection and “recovery” boundaries. Guardians thus explicitly incorporated the notion of state. Lieberman implemented an actor language, Act1, incorporating guardians, serializers, and ‘rock bottom’ actors which is best described in (Lieberman, 1987).

Will Clinger (Clinger, 1981) developed a semantics of actor systems, showing the consistency of axioms proposed in (Baker and Hewitt, 1977). A key accomplishment of Clinger’s work was to show that a powerdomain semantics could be developed despite the fact that the underlying domain was incomplete due to fairness. The work did not develop a theory of actor systems – specifically, no notion of equivalence of actor systems was defined.

The model of actors we use is most closely resembles (Agha, 1986a) which defined a simple transition system for actors, and developed a notion of configurations, receptionists and external actors. The model was implemented by Carl Manning at MIT in the Acore programming language (Manning, 1987) and by Tomlinson and others at the Microelectronics and Computer Technology Consortium in the Rosette programming language (Tomlinson et al., 1989; Tomlinson et al., 1993). It has also provided a basis for dozens of other projects (Agha et al., 1989; Agha et al., 1991; Agha et al., 1993b). Some of the more recent research on actors has focused on coordination structures and meta-architectures (Agha et al., 1993a).

1.2.2 Other Related Research

We consider two areas of related work: process algebras and concurrent functional languages. In all of the work considered here the underlying model uses synchronous communication as primitive instead of asynchronous communication. Synchronous communication can be simulated with asynchronous primitives (Amadio, 1994) and vice-versa, and forms of both are required in a realistic programming language. Also in models considered, although processes may be dynamically created, they are not first class entities that can be directly manipulated. Channels are first class entities that can be dynamically created, communicated as values, and used for communication. One process can communicate with another only if they happen to share a communication channel. In general any number of different processes can send or receive on a given channel, thus processes can interfere with one another. In actor and other object based models, actors/objects are first class, history sensitive, entities whose identity can be communicated and used for communication with the identified object. The communication medium is not explicitly represented. Again, to some degree each choice of primitives can simulate the other. Further work is needed to clarify precisely the relative expressive powers of the two approaches. A more important general distinction is that of these models only (Reppy, 1992) incorporates fairness. Without fairness, specifications fail to compose — a process may behave correctly in isolation, but may fail to do so in the presence of other processes (even if it is not interacting with them).

Much of the work on process algebras is based on Milner's Calculus of Concurrent Systems (CCS) (Milner, 1983), Hoare's CSP (Hoare, 1985), and Milner's π -calculus (Milner et al., 1989). The work on process algebras and the π -calculus has focused on understanding elementary communications between processes, abstracting away other programming language issues. The CSP model assumes a fixed interconnection topology of processes, supports only static storage allocation, and disallows recursive procedures. The π -calculus is to concurrent programming as the pure λ -calculus is to functional programming in the sense that it provides a bare minimum of primitive notions: functions and function application in the case of the λ -calculus; and channels and communication in the case of the π -calculus. A variety of equivalence relations have been defined and studied. These include bisimulation (Milner, 1983; Milner, 1989) which defines a back-and-forth simulation relation between two processes, and trace-based equivalence (Brookes et al., 1984). Various logics have been developed for process algebras, including (Hoare, 1985; Milner, 1989; Bergstra and Klop, 1986).

Facile (Giacalone et al., 1989; Prasad et al., 1990; Thomsen et al., 1992) is perhaps the first language design based on the principle of combining a functional language with concurrency primitives. It differs from our approach in using the typed lambda calculus as the functional component and using concurrency primitives based on CSP and CCS. The underlying model of concurrency is one of processes communicating by synchronized message passing over typed channels. Any value of appropriate type, including higher-order values can be communicated over a channel. In (Giacalone et al., 1989) an operational semantics for Facile is

given using a labelled transition system, and an abstract model of implementation is defined. In (Prasad et al., 1990) an algebraic semantics for Facile is given based on the operational semantics. The central notion here is bisimulation relations. Bisimulations are indexed by sets of channels called windows. This explicitly accounts for the interface of a system to its environment in much the same way as receptionists do for actor systems. Basic process algebra equations are established using bisimulations as well as the beta-value rule. The authors point out that establishing expression equivalence using bisimulations is much more complicated than establishing process algebra equations.

An early version of actor semantics presented in this paper appeared in (Agha et al., 1992). There we defined a notion of operational bisimulation that incorporated fairness. Operational bisimulations were intended to serve as tools for establishing observational equivalence, not as equivalence relations per se. We also found that operational bisimulations were not an effective tool for establishing expression equivalence, since it was difficult to find suitable bisimulations. This led us to develop the alternative methods presented in this paper.

Recently there have been two efforts (Reppy, 1991; Berry et al., 1992) to combine the work on ML with the work on process algebras to obtain a concurrent version of ML. The underlying concurrency model here is the same as that of Facile. In CML (Reppy, 1991) the channel interference problem can be avoided with the use of ML data abstraction mechanisms. A structured operational semantics and type inference system for a small kernel language contained in CML is described in (Reppy, 1991) and a type safety theorem is proved. In addition an implementation of CML is described. Basic properties of the operational semantics of similar kernel language are presented in (Berry et al., 1992). At this point, little has been done towards developing an equational theory concurrent extensions of ML.

2 Actor Language Constructs

Our actor language is an extension of a simple functional language which provides primitives for coordinating components which carry out local computation. An individual actor represents the smallest unit of coordination in the model. Our language provides a mechanism for specifying the creation and manipulation of actors. An actor's behavior is described by a lambda abstraction which embodies the code to be executed when a message is received. The actor primitives are: `send`; `become`; `newadr`; and `initbeh`.

`send` is for sending messages; `send(a, v)` creates a new message with receiver *a* and contents *v* and puts the message into the message delivery system.

`become` is for changing behavior; `become(b)` creates an anonymous actor to carry out the rest of the current computation, alters the behavior of the actor executing the `become` to be *b*, and frees that actor to accept another message. This provides additional parallelism. The anonymous actor may send messages or create new

actors in the process of completing its computation, but will never receive any messages as its address can never be known.

Note that in open distributed systems, the order of arrival of messages from different external sources is nondeterministic. The `become` primitive is necessary to provide the history-sensitive behavior necessary to model asynchronous access to shared resources in such systems. A canonical example of the use of `become` is in modelling a shared bank account accessible by a two or more automatic teller machines. On the other hand, observe that the current behavior of an actor always remains a deterministic function of the sequence of messages that the actor has thus far received.

`newadr` and `initbeh` are for actor creation. `newadr()` creates a new (uninitialized) actor and returns its address. `initbeh(a, b)` initializes the behavior of a newly created actor with address *a* to be *b*. An uninitialized actor can only be initialized by the actor which created it. Without this restriction composability of actor configurations is problematic, as it would permit an external actor to initialize an internally created actor. The allocation of a new address and initialization of the actor's behavior have been separated in order to allow an actor to learn its own address upon initialization. An alternative would be to have a `letactor` construct similar to the Scheme `letrec` construct.

2.1 Trivial Examples

A simple actor behavior *b* that expects its message to be an actor address, sends the message 5 to that address, and becomes the same behavior, may be expressed as follows.

$$b = \text{rec}(\lambda y. \lambda x. \text{seq}(\text{send}(x, 5), \text{become}(y)))$$

where `seq` is syntactic sugar for expressing sequential composition, and `rec` is a definable (in the pure λ fragment) call-by-value fixed-point combinator (cf. (Mason and Talcott, 1991)). An equivalent expression of this behavior is:

$$b = \text{rec}(\lambda y. \lambda x. \text{seq}(\text{become}(y), \text{send}(x, 5)))$$

since the order of executing the `become` and the `send` cannot be observed. An expression that would create an actor with behavior *b* and send it another actor address *a* is

$$e = \text{let}\{z := \text{newadr}()\}\text{seq}(\text{initbeh}(z, b), \text{send}(z, a)).$$

The behavior of a sink, an actor that ignores its messages and becomes this same behavior, is defined by

$$\text{sink} = \text{rec}(\lambda b. \lambda m. \text{become}(b)).$$

2.2 Actor Cells

It is easy to represent objects with local state in our language. As an example of this we describe an actor akin to an ML reference cell. The actor responds to two sorts of messages. The first sort of message is a `get` message, recognized by the `get?` operation. A `get` message contains the address (customer) to which the response (i.e the current contents of the cell) should be sent. This address is accessed via the `cust` operation. A `mkget(a)` constructs a `get` message with customer `a`. The second message is a `set` message, recognized by the `set?` operation. A `set` message should contain the desired new contents of the cell-actor. This value is accessed via the `contents` operation. A `mkset(c)` constructs a `set` message with new contents `c`. Using these operations, the behavior of a cell is given by:

```
Bcell = rec(λb.λc.λm.
    if(get?(m),
      seq(become(b(c)), send(cust(m), c))
    if(set?(m),
      become(b(contents(m))),
      become(b(c))))))
```

Evaluating

```
let{a := newadr()}seq(initbeh(a, Bcell(0)), e)  where
e = seq(send(a, mkset(3)), send(a, mkset(4)), send(a, mkget(a)))
```

will result in the actor `a` receiving a message containing either 0, 3 or 4, depending on the arrival order. A cell is one of the simplest kinds of history sensitive object. The `become` primitive is the key to expressing history sensitivity by allowing new behaviors to be installed in response to messages. Accumulators, counters and gensym actors are easily constructed in a similar manner.

2.3 Join Continuations

Simple forms of recursion are often amenable to concurrent execution. A typical example is tree recursion. Consider the problem of determining the product of the leaves of a tree (whose leaves are numbers). The problem can be recursively subdivided into the problem of computing the result for the two subtrees, and multiplying the results. The product is then returned.

```
treeprod = rec(λf.λtree.
    if(isnat(tree),
      tree,
      f(left(tree)) * f(right(tree))))
```

In the above code, a `tree` is passed to `treeprod` which tests to see if the `tree` is a number (i.e. a leaf). If so it returns the tree, otherwise it subdivides the problem into two recursive calls. `left` and `right` are functions which pick off the left and

right branches of the tree. It is clear that the arguments to $*$ may be evaluated concurrently. It is also clear that if a zero is encountered then the computation can terminate. In this example we only deal with the former optimization. The latter optimization, made using continuations, is treated in detail in (Talcott, 1993b; Talcott, 1985; Agha, 1990)

Such concurrency can be implemented by using a *join continuation* which synchronizes the evaluation of the different arguments. For example, the `treeprod` program given above can be expressed in terms of actor primitives as:

```

Btreeprod =
  rec( $\lambda b.\lambda self.\lambda m.$ 
    seq(become(b(self)),
      if(isnat(tree(m)),
        send(cust(m), tree(m)),
        let{newcust := newadr()}
          seq(initbeh(newcust, Bjoincont(cust(m), 0, nil)),
            send(self, <left(tree(m), newcust>),
            send(self, <right(tree(m), newcust>))))))
Bjoincont = rec( $\lambda b.\lambda cust.\lambda nargs.\lambda firstnum.\lambda num$ 
  if(eq(nargs, 0),
    become(b(cust, 1, num)),
    seq(become(sink),
      send(cust, firstnum * num))))

```

When an *tree product* actor (with behavior B_{treeprod}) receives a tree of numbers that is not a leaf it creates a customer, called a *join continuation*, and sends two messages to itself to evaluate the two halves of the tree. These messages have the join continuation as customer. The join continuation (with behavior B_{joincont}) expects to receive two numbers representing the computation of the products of each of the two subtrees. When both numbers have arrived, the join continuation multiplies them and sends the result to its customer. Figure 1 shows some stages in the computation of a tree production. The join continuation's customer can be the original requester (root of the tree), or the join continuation of a higher branch point. Because multiplication is commutative, we need not be concerned about matching the responses to the order of the parameters. If we were dealing with an operator which was not commutative, we would need to tag the message corresponding to each argument and this tag would be returned with the response from the corresponding subcomputation. The replacement behavior of the join continuation would then depend on the order in which the evaluation of arguments was completed.

An advantage of explicit join continuations is that they provide considerable flexibility – they can be used to control the evaluation order, to do partial computations, and to do handle special cases or errors. For example, if the number 0

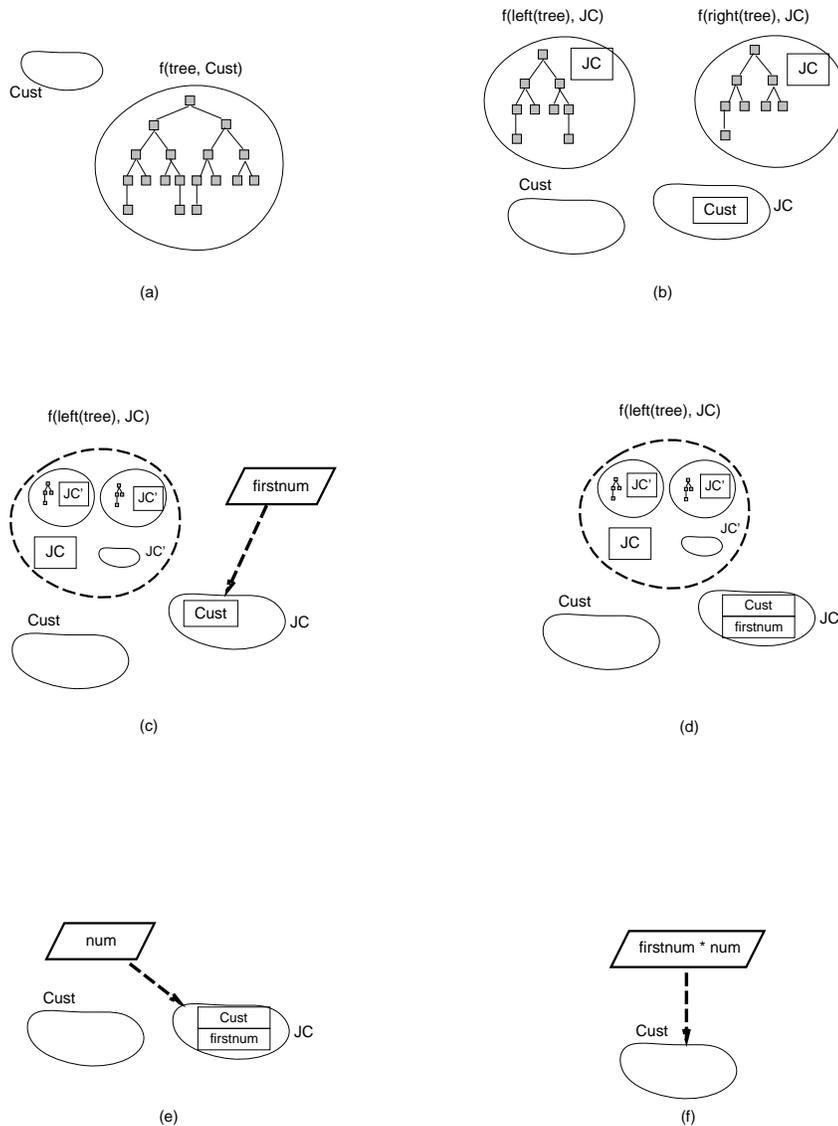


Fig. 1. The leaves of tree, contain numbers to be multiplied and Cust is the actor to which the product is to be sent. (b) the tree is subdivided and two asynchronous messages are sent to compute the product on the two halves (concurrently). JC represents a newly created actor to which the two results should be sent. Each subtree will be recursively subdivided and new join-continuations will be successively created. (c) When the product of one of the subtrees has been computed, the value is sent to JC. (d) JC does a become to store the first number it receives. (e) The second product sends a value to JC. (f) JC multiplies the second number it receives with the number it had stored earlier and sends the result to Cust.

is encountered, the join continuation can immediately return a 0 – in some cases without waiting for the results of evaluating the other subtree.

The above program may not be optimal in other respects. For example `sends` may be quite expensive. Consequently it may be prudent to check that the sub-computation is worth dispatching. This observation together with the fact, proved in §4, that `sends` of values commute leads to the possibly faster version:

```

Btreeproduct =
  rec( $\lambda b. \lambda self. \lambda m.$ 
    seq(become(b),
      if(isnat(tree(m)),
        send(cust(m), tree(m)),
        let{newcust := newadr()}
          seq(initbeh(newcust, Bjoincont(cust(m), 0, nil))
            let{l := left(tree(m)), r := right(tree(m))}
              if(or(isnat(l), isnat(r)),
                send(cust(m), treeprod(l) * treeprod(r)),
                seq(send(self, <l, newcust>
                  send(self, <r, newcust>)))))))))

```

2.4 Notation.

We use the usual notation for set membership and function application. Let Y, Y_0, Y_1 be sets. We specify meta-variable conventions in the form: let y range over Y , which should be read as: the meta-variable y and decorated variants such as y', y_0, \dots , range over the set Y . Y^n is the set of sequences of elements of Y of length n . Y^* is the set of finite sequences of elements of Y . $\bar{y} = [y_1, \dots, y_n]$ is the sequence of length $\text{Len}(\bar{y}) = n$ with i th element y_i . (Thus $[]$ is the empty sequence.) $u * v$ denotes the concatenation of the sequences u and v . If u is a non-empty sequence, then $\text{Last}(u)$ is the last element of u . $\mathbf{P}_\omega[Y]$ is the set of finite subsets of Y . $\mathbf{M}_\omega[Y]$ is the set of (finite) multi-sets with elements in Y . $Y_0 \xrightarrow{f} Y_1$ is the set of finite maps from Y_0 to Y_1 . $[Y_0 \rightarrow Y_1]$ is the set of total functions, f , with domain Y_0 and range contained in Y_1 . We write $\text{Dom}(f)$ for the domain of a function and $\text{Rng}(f)$ for its range. For any function f : $f\{y := y'\}$ is the function f' such that $\text{Dom}(f') = \text{Dom}(f) \cup \{y\}$, $f'(y) = y'$, and $f'(z) = f(z)$ for $z \neq y, z \in \text{Dom}(f)$; and $f \upharpoonright Y$ is the restriction of f to the set Y .

3 A Simple Lambda-Based Actor Language

In this section we give the syntax and operational semantics of our actor language.

3.1 Syntax

We take as given countable sets \mathbb{X} (variables), At (atoms), and for each $n \in \mathbb{N}$ a (possibly empty) set of n -ary operations, \mathbb{G}_n on At . \mathbb{F}_n is the set of primitive operations of arity n , which includes \mathbb{G}_n , and $\mathbb{F} = \bigcup_{n \in \mathbb{N}} \mathbb{F}_n$. We assume At contains t, nil for booleans, as well as constants for natural numbers, \mathbb{N} . \mathbb{F} contains arithmetic operations, recognizers isatom for atoms, and isnat for numbers (arities 1, 1), branching br (arity 3), pairing $\text{ispr}, \text{pr}, 1^{\text{st}}, 2^{\text{nd}}$ (arities 1, 2, 1, 1), and actor primitives $\text{newadr}, \text{initbeh}, \text{send},$ and become (arities 0, 2, 2, 1). The sets of expressions, \mathbb{E} , value expressions (or just values), \mathbb{V} , and contexts (expressions with holes), \mathbb{C} , are defined inductively as follows.

Definition ($\mathbb{E} \ \mathbb{V} \ \mathbb{C}$):

$$\mathbb{V} = \text{At} \cup \mathbb{X} \cup \lambda \mathbb{X}. \mathbb{E} \cup \text{pr}(\mathbb{V}, \mathbb{V})$$

$$\mathbb{E} = \text{At} \cup \mathbb{X} \cup \lambda \mathbb{X}. \mathbb{E} \cup \text{app}(\mathbb{E}, \mathbb{E}) \cup \mathbb{F}_n(\mathbb{E}^n)$$

$$\mathbb{C} = \text{At} \cup \mathbb{X} \cup \lambda \mathbb{X}. \mathbb{C} \cup \text{app}(\mathbb{C}, \mathbb{C}) \cup \mathbb{F}_n(\mathbb{C}^n) \cup \{\bullet\}$$

We let x, y, z range over \mathbb{X} , v range over \mathbb{V} , e range over \mathbb{E} , and C range over \mathbb{C} . Since we are working with a syntactic reduction semantics, there is no distinction between a value expression and the value it denotes. Hence we use the terms *value* and *value expression* interchangeably. $\lambda x.e$ binds the variable x in the expression e . Two expressions are considered equal if they are the same up to renaming of bound variables. We say that a variable is *fresh* with respect to a context of use if it does not occur free or bound in any syntactic entity. We write $\text{FV}(e)$ for the set of free variables of e . We write $e[x := e']$ to denote the expression obtained from e by replacing all free occurrences of x by e' , avoiding the capture of free variables in e' . Contexts are expressions with holes. We use \bullet to denote a hole. $C[e]$ denotes the result of replacing any holes in C by e . Free variables of e may become bound in this process. $\text{let}, \text{if},$ and seq are the usual syntactic sugar, seq being a sequencing primitive.

$$\text{let}\{x := e_0\}e_1 \quad \text{abbreviates} \quad \text{app}(\lambda x.e_1, e_0)$$

$$\text{seq}(e_0, e_1) \quad \text{abbreviates} \quad \text{app}(\text{app}(\lambda z.\lambda x.x, e_0), e_1)$$

$$\text{if}(e_0, e_1, e_2) \quad \text{abbreviates} \quad \text{app}(\text{br}(e_0, \lambda z.e_1, \lambda z.e_2), \text{nil}) \quad \text{for } z \text{ fresh}$$

3.2 Reduction Semantics for Actor Configurations

The operational semantics for actor systems is given by a transition relation on configurations. A configuration can be thought of as representing a global snapshot of an actor system with respect to some idealized observer (Agha, 1986b). It contains a collection of actors, messages, external actor names, and receptionist names. The sets of receptionists and external actors are the interface of an actor configuration to its environment. They specify what actors are visible and what actor connections must be provided for the configuration to function. Both the set of receptionists and the set of external actors may grow as the configuration evolves.

The state of the actors in a configuration is given by an actor map. An actor map is a finite map from actor addresses to actor states. Each actor state is one of

- (? a) uninitialized, having been newly created by an actor named a ;
- (b) ready to accept a message, where b is its behavior, a lambda abstraction; or
- [e] busy executing e , here e represents the actor's current (local) processing state.

A message m contains the address of the actor to whom it is sent and the message contents, $\langle a \Leftarrow v \rangle$. The contents v can be any value constructed from atoms and actor addresses using the pairing constructor. Lambda abstractions and structures containing lambda abstractions are not allowed to be communicated in messages. There are two reasons for this restriction. Firstly, allowing lambda abstractions to be communicated in values violates the actor principle that only an actor can change its own behavior, because a become in a lambda message may change the receiving actor behavior. Secondly, if lambda abstractions are communicated to external actors, there is no reasonable way to control what actor addresses are actually exported. This has unpleasant consequences in reasoning about equivalence, amongst other things. This restriction is not a serious limitation since the address of an actor whose behavior is the desired lambda abstraction can be passed in a message.

The transition relation determines the set of possible future configurations. We classify actor configuration transitions as internal or external to the configuration. The internal transitions are:

rcv: receipt of a message by an actor not currently busy computing; and
exec: an actor executing a step of its current computation.

The internal transitions involve a single active actor, which we call the *focus actor* for the transition. **exec** transitions may be purely local (a λ -transition), or messages may be sent, or a new actor may be created, or a newly created actor may be initialized. **rcv** transitions consume a message, putting the focus actor in a busy state.

In addition to the internal transitions of a configuration, there are i/o transitions that correspond to interactions with external agents:

in: arrival of a message to a receptionist from the outside; and
out: passing a message out to an external actor.

To capture actor semantics, we require that computations be fair.

3.2.1 Actor Configurations

We assume that we are given a countable set \mathbb{A} of actor addresses. To simplify notation, we identify \mathbb{A} with \mathbb{X} and call variables used in this way actor names. This pun is useful for two reasons: it allows us to use expressions to describe actor states and message contents; and it allows us to avoid problems of choice of names for newly created actors by appealing to an extended form of alpha conversion. (See (Mason and Talcott, 1991; Felleisen and Hieb, 1992) for use of this pun to represent reference cells.)

Definition ($c\mathbb{V}$ \mathbb{A} \mathbb{M}): The set of communicable values, $c\mathbb{V}$, the set of actor

states, $\mathbb{A}s$, and the set of messages, \mathbb{M} , are defined as follows.

$$c\mathbb{V} = \text{At} \cup \mathbb{X} \cup \text{pr}(c\mathbb{V}, c\mathbb{V}) \quad \mathbb{A}s = (?_{\mathbb{X}}) \cup (\mathbb{V}) \cup [\mathbb{E}] \quad \mathbb{M} = \langle \mathbb{V} \Leftarrow \mathbb{V} \rangle$$

We let cv range over $c\mathbb{V}$ and m range over \mathbb{M} . Note that actor behaviors are not syntactically restricted to be lambda abstractions, nor are messages syntactically restricted to be of the form $\langle \text{Ad} \Leftarrow c\mathbb{V} \rangle$. The reduction system will prevent use of any ill-formed behavior or message. This is in keeping with the untyped nature of our language.

Definition (Actor Configurations (\mathbb{K})): An actor configuration with actor map, α , multi-set of messages, μ , receptionists, ρ , and external actors, χ , is written

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

where $\rho, \chi \in \mathbf{P}_{\omega}[\mathbb{X}]$, $\alpha \in \mathbb{X} \xrightarrow{f} \mathbb{A}s$, and $\mu \in \mathbf{M}_{\omega}[\mathbb{M}]$. Further, it is required that, letting $A = \text{Dom}(\alpha)$, the following constraints are satisfied:

- (0) $\rho \subseteq A$ and $A \cap \chi = \emptyset$,
- (1) if $\alpha(a) = (?_{a'})$, then $a' \in A$,
- (2) 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 $i < 2$.

We let \mathbb{K} denote the set of actor configurations and let κ range over \mathbb{K} . The *receptionists* ρ are names of actors within the configuration that are externally visible; all other actors in the (actor) configuration are local and thus inaccessible from the outside. The *external actors* χ are names of actors that are outside this configuration but to which messages may be directed. A configuration in which both the receptionist and external actor sets are empty is said to be *closed*. For closed configurations we may omit explicit mention of the empty ρ and χ sets. The actor map portion of a configuration is presented as a list of actor states each subscripted by the actor address which is mapped to this state. If $\alpha'(a) = (b)$, and α is α' with a omitted from its domain, we write α' as $(\alpha, (b)_a)$ to focus attention on a . We follow a similar convention for other states subscripted with addresses.

The set of possible computations of an actor configuration is defined in terms of the labelled transition relation \mapsto on configurations. Although this is, on the surface an interleaving semantics, it is easy to modify our transition system to obtain a truly concurrent semantics either by forming a labelled transition system with independence (Sassone et al., 1993), or by using concurrent rewriting (Meseguer, 1992).

3.2.2 Decomposition and Reduction

To describe the internal transitions other than message receipt, a non-value expression is decomposed uniquely into a reduction context filled with a redex. Reduction contexts identify the subexpression of an expression that is to be evaluated next, they correspond to the standard reduction strategy (left-first, call-by-value) of (Plotkin, 1975) and were first introduced in (Felleisen and Friedman, 1986). For further discussion of this method of defining reduction relations see (Honsell et al.,

1994). In order to distinguish holes used for different purposes, we use the sign \square for the hole occurring in a reduction context, and call such holes redex holes (although they may in fact be filled with non-redex expressions).

Definition (\mathbb{E}_{rdx} \mathbb{R}): The set of redexes, \mathbb{E}_{rdx} , and the set of reduction contexts, \mathbb{R} , are defined by

$$\begin{aligned}\mathbb{E}_{\text{rdx}} &= \mathbf{app}(\mathbb{V}, \mathbb{V}) \cup (\mathbb{F}_n(\mathbb{V}^n) - \mathbf{pr}(\mathbb{V}, \mathbb{V})) \\ \mathbb{R} &= \{\square\} \cup \mathbf{app}(\mathbb{R}, \mathbb{E}) \cup \mathbf{app}(\mathbb{V}, \mathbb{R}) \cup \mathbb{F}_{n+m+1}(\mathbb{V}^n, \mathbb{R}, \mathbb{E}^m)\end{aligned}$$

We let R range over \mathbb{R} .

An expression e is either a value or it can be decomposed uniquely into a reduction context filled with a redex. Thus, local actor computation is deterministic.

Lemma (Unique decomposition):

- (0) $e \in \mathbb{V}$, or
- (1) $(\exists! R, r)(e = R[r])$

Proof: An easy induction on the structure of e . \square

Redexes can be split into two classes, purely functional and actor redexes. The actor redexes are: $\mathbf{newadr}()$, $\mathbf{initbeh}(v_0, v_1)$, $\mathbf{become}(v)$, and $\mathbf{send}(v_0, v_1)$. Reduction rules for the purely functional case are given by a relation $\xrightarrow{\lambda}$ on *expressions*.

Definition ($\xrightarrow{\lambda}$):

- (beta-v) $R[\mathbf{app}(\lambda x.e, v)] \xrightarrow{\lambda} R[e[x := v]]$
- (delta) $R[\delta(v_1, \dots, v_n)] \xrightarrow{\lambda} R[v']$
where $\delta \in \mathbb{G}_n$, $v_1, \dots, v_n \in \mathbb{At}^n$, and $\delta(v_1, \dots, v_n) = v'$.
- (br) $R[\mathbf{br}(v, v_1, v_2)] \xrightarrow{\lambda} \begin{cases} R[v_1] & \text{if } v \in \mathbb{V} - \{\mathbf{nil}\} \\ R[v_2] & \text{if } v = \mathbf{nil} \end{cases}$
- (ispr) $R[\mathbf{ispr}(v)] \xrightarrow{\lambda} \begin{cases} R[\mathbf{t}] & \text{if } v \in \mathbf{pr}(\mathbb{V}, \mathbb{V}) \\ R[\mathbf{nil}] & \text{if } v \in \mathbb{V} - \mathbf{pr}(\mathbb{V}, \mathbb{V}) \end{cases}$
- (fst) $R[\mathbf{1}^{\text{st}}(\mathbf{pr}(v_0, v_1))] \xrightarrow{\lambda} R[v_0]$
- (snd) $R[\mathbf{2}^{\text{nd}}(\mathbf{pr}(v_0, v_1))] \xrightarrow{\lambda} R[v_1]$
- (eq) $R[\mathbf{eq}(v_0, v_1)] \xrightarrow{\lambda} \begin{cases} R[\mathbf{t}] & \text{if } v_0 = v_1 \in \mathbb{At} \\ R[\mathbf{nil}] & \text{if } v_0, v_1 \in \mathbb{At} \text{ and } v_0 \neq v_1 \end{cases}$

The rules for \mathbf{isatom} and \mathbf{isnat} are analogous to that for \mathbf{ispr} . The single-step transition relation \mapsto on actor configurations is generated by the following rules. Each rule is given a label l consisting of a tag indicating the primitive instruction, and additional parameters. In all cases other than i/o transitions (with tags in, out) the first parameter names the *focus* actor of the transition.

Definition (\mapsto):

$\langle \mathbf{fun} : a \rangle$

$$\begin{aligned}
e &\xrightarrow{\lambda} e' \Rightarrow \langle\langle \alpha, [e]_a \mid \mu \rangle\rangle_{\chi}^{\rho} \mapsto \langle\langle \alpha, [e']_a \mid \mu \rangle\rangle_{\chi}^{\rho} \\
\langle \text{new} : a, a' \rangle & \\
&\langle\langle \alpha, [R[\text{newadr}()]]_a \mid \mu \rangle\rangle_{\chi}^{\rho} \mapsto \langle\langle \alpha, [R[a']]_a, (?_a)_{a'} \mid \mu \rangle\rangle_{\chi}^{\rho} \quad a' \text{ fresh} \\
\langle \text{init} : a, a' \rangle & \\
&\langle\langle \alpha, [R[\text{initbeh}(a', v)]]_a, (?_a)_{a'} \mid \mu \rangle\rangle_{\chi}^{\rho} \mapsto \langle\langle \alpha, [R[\text{nil}]]_a, (v)_{a'} \mid \mu \rangle\rangle_{\chi}^{\rho} \\
\langle \text{bec} : a, a' \rangle & \\
&\langle\langle \alpha, [R[\text{become}(v)]]_a \mid \mu \rangle\rangle_{\chi}^{\rho} \mapsto \langle\langle \alpha, [R[\text{nil}]]_{a'}, (v)_a \mid \mu \rangle\rangle_{\chi}^{\rho} \quad a' \text{ fresh} \\
\langle \text{send} : a, m \rangle & \\
&\langle\langle \alpha, [R[\text{send}(v_0, v_1)]]_a \mid \mu \rangle\rangle_{\chi}^{\rho} \mapsto \langle\langle \alpha, [R[\text{nil}]]_a \mid \mu, m \rangle\rangle_{\chi}^{\rho} \\
&\quad m = \langle v_0 \leftarrow v_1 \rangle \\
\langle \text{rcv} : a, cv \rangle & \\
&\langle\langle \alpha, (v)_a \mid \langle a \leftarrow cv \rangle, \mu \rangle\rangle_{\chi}^{\rho} \mapsto \langle\langle \alpha, [\text{app}(v, cv)]_a \mid \mu \rangle\rangle_{\chi}^{\rho} \\
\langle \text{out} : m \rangle & \\
&\langle\langle \alpha \mid \mu, m \rangle\rangle_{\chi}^{\rho} \mapsto \langle\langle \alpha \mid \mu \rangle\rangle_{\chi}^{\rho'} \\
&\quad \text{if } m = \langle a \leftarrow cv \rangle, a \in \chi, \text{ and } \rho' = \rho \cup (\text{FV}(cv) \cap \text{Dom}(\alpha)) \\
\langle \text{in} : m \rangle & \\
&\langle\langle \alpha \mid \mu \rangle\rangle_{\chi}^{\rho} \mapsto \langle\langle \alpha \mid \mu, m \rangle\rangle_{\chi \cup (\text{FV}(cv) - \text{Dom}(\alpha))}^{\rho} \\
&\quad \text{if } m = \langle a \leftarrow cv \rangle, a \in \rho \text{ and } \text{FV}(cv) \cap \text{Dom}(\alpha) \subseteq \rho
\end{aligned}$$

In the rules for `newadr` and `become`, a' fresh abbreviates $a' \notin \text{Dom}(\alpha) \cup \{a\} \cup \chi$. We write $\kappa_0 \xrightarrow{l} \kappa_1$ if $\kappa_0 \mapsto \kappa_1$ according to the rule labelled by l . We call this triple a labelled transition. We say l is *enabled in configuration* κ if there is some κ' such that $\kappa \xrightarrow{l} \kappa'$. The transitions are labelled to allow us to reason about sequences of transitions in terms of the rules applied, and to allow for alternative representation of computations, including: sequences of configurations; sequences of labelled transitions; and sequences of labels. Note that we have chosen the labels to include sufficient information that κ' is uniquely determined by κ and l .

i/o transitions are transitions with tags `in` or `out`; `rcv` transitions are transitions with tag `rcv`. The remaining transitions are called `exec` transitions. The `exec` transitions correspond to the execution of functional or actor redexes.

As mentioned above, we allow ill-formed messages to be created, but such messages can never be delivered. The last three rules assure this by restricting the form of the message: the target must be an actor and the contents must be a communicable value. In the case of input, the actor is further restricted to be a receptionist.

We could easily prevent the formation of ill-formed messages and actor states if so desired.

A clone produced to carry on after a become is not allowed to initialize an actor created by its cloner. This is a technical simplification. With some additional book-keeping we could keep track of cloners and allow clones to initialize. Alternatively, this technical detail would disappear if we used the `letactor` construct for actor creation.

These choices affect the details of expression equivalence, but not the basic properties. Such choices will become more important if we want to model an implemented language and consider matters such as signaling of exceptions.

3.2.3 Computation Sequences and Paths

Definition (Computation trees): If κ is a configuration, then we define the *computation tree* for κ , $\mathcal{T}(\kappa)$, to be the set of all finite sequences of labelled transitions of the form $[\kappa_i \xrightarrow{l_i} \kappa_{i+1} \mid i < n]$ for some $n \in \mathbb{N}$. We call such sequences *computation sequences* and let ν range over them.

Lemma (Anonymity): If $\kappa \xrightarrow{\langle \text{bec} : a, a' \rangle} \kappa'$ and ν is any computation sequence in the computation tree, $\mathcal{T}(\kappa)$, then ν contains no transitions with label of the form $\langle \text{send} : a', v \rangle$.

Definition (Computation paths): The sequences of a computation tree are partially ordered by the initial segment relation. A *computation path* from κ is a maximal linearly ordered set of computation sequences in the computation tree, $\mathcal{T}(\kappa)$. Note that a path can also be regarded as a (possibly infinite) sequence of labelled transitions. We use $\mathcal{T}^\infty(\kappa)$ to denote the set of all paths from κ , and let π range over computation paths. When thinking a path as a possibly infinite sequence we write $[\kappa_i \xrightarrow{l_i} \kappa_{i+1} \mid i < \infty]$ where $\infty \in \mathbb{N} \cup \{\omega\}$ is the length of the sequence.

Since the result of a transition is uniquely determined by the starting configuration and the transition label, computation sequences and paths can be equally represented by their initial configuration and the sequence of transitions labels. The sequence of configurations can be computed by induction on the index of occurrence.

Definition (Cfig): Let κ be a configuration and let $L = [l_i \mid i < \infty]$ be a sequence of labels corresponding to a computation from κ . The i th configuration of the computation from κ determined by L , $\text{Cfig}(\kappa, L, i)$, is defined by induction on i as follows.

- (0) $\text{Cfig}(\kappa, L, 0) = \kappa$
- (1) $\text{Cfig}(\kappa, L * [l_i], i + 1) = \kappa'$ where $\text{Cfig}(\kappa, L, i) \xrightarrow{l_i} \kappa'$.

Thus, the path π determined by κ, L is the sequence

$$[\text{Cfig}(\kappa, L, i) \xrightarrow{l_i} \text{Cfig}(\kappa, L, i + 1) \mid i < \infty]$$

This notation has the advantage that when an initial starting configuration is fixed, either implicitly or explicitly, computation sequences in the computation tree can be

identified with sequences of labels. When the sequence L is finite we let $\text{Cfig}(\kappa, L)$ denote the final configuration: $\text{Cfig}(\kappa, L) = \text{Cfig}(\kappa, L, \text{Len}(L))$.

Definition (multi-step transition): Let $L = [l_j \mid j < n]$ be a finite sequence of transition labels (possibly empty). L is a *multi-step transition* $\kappa \xrightarrow{L} \kappa'$ just if $\text{Cfig}(\kappa, L) = \text{Cfig}(\kappa, L, \text{Len}(L)) = \kappa'$. Or in other words, we can find a $[\kappa_j \mid j \leq n]$ such that $\kappa = \kappa_0$, $\kappa' = \kappa_n$, and $[\kappa_j \xrightarrow{l_j} \kappa_{j+1} \mid j < n]$.

3.2.4 Fairness

We do not consider all paths admissible. We rule out those computations that are unfair, i.e. those in which there is some transition that should eventually happen and does not.

Definition (Fair paths): A path $\pi = [\nu_i \xrightarrow{l_i} \nu_{i+1} \mid i < \infty]$ in the computation tree $\mathcal{T}^\infty(\kappa)$ is fair if each enabled transition eventually happens or becomes permanently disabled. That is, if l is enabled in κ_i and is not of the form $\langle \text{in} : m \rangle$, then $\kappa_j \xrightarrow{l} \kappa_{j+1}$ for some $j \geq i$, or l has the form $\langle \text{rcv} : a, cv \rangle$ and for some $j \geq i$ a is busy and never again becomes ready to accept a message. For a configuration κ we define $\mathcal{F}(\kappa)$ to be the subset of $\mathcal{T}^\infty(\kappa)$ that are fair.

Note that finite computation paths are fair, since by maximality all of the enabled transitions must have happened.

3.3 Composition of Actor Configurations

Actor configurations can be composed to form new actor configurations. This composition operation is commutative, associative, and has the empty configuration as unit. This is made precise by the following definitions and lemmas.

Definition (Composable): Two configurations $\kappa_i = \langle\langle \alpha_i \mid \mu_i \rangle\rangle_{\chi_i}^{\rho_i}$, $i < 2$ are composable if $\text{Dom}(\alpha_0) \cap \text{Dom}(\alpha_1) = \emptyset$, $\chi_0 \cap \text{Dom}(\alpha_1) \subseteq \rho_1$, and $\chi_1 \cap \text{Dom}(\alpha_0) \subseteq \rho_0$.

Definition (Composition, decomposition): The composition $\kappa_0 \parallel \kappa_1$ of composable configurations $\kappa_i = \langle\langle \alpha_i \mid \mu_i \rangle\rangle_{\chi_i}^{\rho_i}$, $i < 2$ is defined by

$$\kappa_0 \parallel \kappa_1 = \langle\langle \alpha_0 \cup \alpha_1 \mid \mu_0 \cup \mu_1 \rangle\rangle_{(\chi_0 \cup \chi_1) - (\rho_0 \cup \rho_1)}^{\rho_0 \cup \rho_1}$$

(κ_0, κ_1) is a decomposition of κ if κ_i , $i < 2$ are composable configurations, and $\kappa = \kappa_0 \parallel \kappa_1$.

Lemma (AC): Let $\kappa_i = \langle\langle \alpha_i \mid \mu_i \rangle\rangle_{\chi_i}^{\rho_i}$, $i < 3$ be pairwise composable configurations. And let $\kappa_\emptyset = \langle\langle \emptyset \mid \emptyset \rangle\rangle$ be the empty configuration. Then

$$\kappa_0 \parallel \kappa_1 = \kappa_1 \parallel \kappa_0$$

$$\kappa_0 \parallel \kappa_\emptyset = \kappa_0$$

$$(\kappa_0 \parallel \kappa_1) \parallel \kappa_2 = \kappa_0 \parallel (\kappa_1 \parallel \kappa_2)$$

Proof : Using the AC properties of set union, the only thing to check is the

equality of external actors for the two associations. It is easy to see that

$$\begin{aligned} & (((\chi_0 \cup \chi_1) - (\rho_0 \cup \rho_1)) \cup \chi_2) - (\rho_0 \cup \rho_1 \cup \rho_2) \\ &= (\chi_0 \cup \chi_1 \cup \chi_2) - (\rho_0 \cup \rho_1 \cup \rho_2) \\ &= (\chi_0 \cup ((\chi_1 \cup \chi_2) - (\rho_1 \cup \rho_2))) - (\rho_0 \cup \rho_1 \cup \rho_2) \end{aligned}$$

□

Furthermore, it is possible to define composition operations on trees

$$\mathcal{T}(\kappa_0) \parallel \mathcal{T}(\kappa_1) \quad \text{and} \quad \mathcal{F}(\kappa_0) \parallel \mathcal{F}(\kappa_1)$$

such that

- (1) The computation tree of the composition of actor configurations is the composition of the computation trees of the components.

$$\mathcal{T}(\kappa_0 \parallel \kappa_1) = \mathcal{T}(\kappa_0) \parallel \mathcal{T}(\kappa_1)$$

- (2) The set of fair computation paths of the composition of actor configurations is the composition of the fair computation paths of the components.

$$\mathcal{F}(\kappa_0 \parallel \kappa_1) = \mathcal{F}(\kappa_0) \parallel \mathcal{F}(\kappa_1)$$

4 Equivalence of Expressions

In this section we study equivalence of expressions of our actor language. Our notion of equivalence is a combination of the now classic *operational equivalence* of (Plotkin, 1975) and *testing equivalence* of (de Nicola and Hennessy, 1984). For the deterministic functional languages of the sort Plotkin studied, this equivalence is defined as follows. Two program expressions are said to be equivalent if they behave the same when placed in any observing context. An observing context is some complete program with a hole, such that all of the free variables in the expressions being observed are captured when the expressions are placed in the hole. The notion of “behave the same” is (for deterministic functional languages) typically equi-termination, i.e. either both converge or both diverge.

4.1 Events

The first step is to find proper notions of “observing context” and “behave the same” in an actor setting. The analogue of an observing context is an observing actor configuration: a configuration that contains an actor state with a hole. Since termination is not relevant for actor configurations, we instead introduce an observer primitive, `event` and observe whether or not in a given computation, `event` is executed. Our approach is identical in spirit to that used in defining testing equivalence for CCS (de Nicola and Hennessy, 1984).

Definition (event): Formally, the language of observing contexts is obtained

by introducing a new 0-ary primitive operator, `event`. We extend the reduction relation \mapsto by adding the following rule.

$$\langle e : a \rangle \quad \left\langle \alpha, [R[\text{event}()]]_a \mid \mu \right\rangle_x^\rho \mapsto \left\langle \alpha, [R[\text{nil}]]_a \mid \mu \right\rangle_x^\rho$$

Definition (\mathbb{O}): For an expression e , the observing configurations are configurations over the extended language of the form $\left\langle \alpha, [C]_a \mid \mu \right\rangle$, such that filling the hole in C with e results in a closed configuration. We use \mathbb{O} to denote the set of observing configurations, and let O range over \mathbb{O} .

In our definition of observing configuration, the holes appear in the current state of an single executing actor. It is not hard to see that allowing holes in any actor state does not change the resulting notion of equivalence. A generalization of this fact, (`ocx`) is proved in §5.

Since the language is nondeterministic, three different outcomes are possible in place of the two in the deterministic case: either `event` occurs for all possible executions, it occurs in some executions but not others, or it never occurs. We observe `event` transitions in the fair paths. We say that a path succeeds, **s**, if an `event` transition occurs in it. This is the basic unit of observation; on top of this derived notions can be defined. We say a computation is observed to fail, **f**, if it is not observed to succeed. $obs(\pi)$ is **s** if π succeeds, and **f** otherwise, and $Obs(\kappa)$ is the set of observations possible for all paths of a closed actor configuration.

Definition (observations): Let κ be a configuration of the extended language, and let $\pi = [\kappa_i \xrightarrow{l_i} \kappa_{i+1} \mid i < \heartsuit]$ be a fair path, i.e. $\pi \in \mathcal{F}(\kappa)$. Define

$$obs(\pi) = \begin{cases} \mathbf{s} & \text{if } (\exists i < \heartsuit, a)(l_i = \langle e : a \rangle) \\ \mathbf{f} & \text{otherwise} \end{cases}$$

$$Obs(\kappa) = \begin{cases} \mathbf{s} & \text{if } (\forall \pi \in \mathcal{F}(\kappa))(obs(\pi) = \mathbf{s}) \\ \mathbf{sf} & \text{if } (\exists \pi \in \mathcal{F}(\kappa))(obs(\pi) = \mathbf{s}) \text{ and } (\exists \pi \in \mathcal{F}(\kappa))(obs(\pi) = \mathbf{f}) \\ \mathbf{f} & \text{if } (\forall \pi \in \mathcal{F}(\kappa))(obs(\pi) = \mathbf{f}) \end{cases}$$

4.2 Three Equivalences

The natural notion of observational equivalence is that equal observations are made in all closing configuration contexts. However, it is possible in some cases to use a weaker equivalence. An **sf** observation may be considered as good as an **s** observation, and a new equivalence arises if these observations are equated. Similarly, an **sf** observation may be as bad as an **f** observation. We define the following three equivalences.

Definition ($\cong_{1,2,3}$):

- (1) (testing or convex or Plotkin or Egli-Milner)

$$e_0 \cong_1 e_1 \quad \text{iff} \quad Obs(O[e_0]) = Obs(O[e_1]) \text{ for all observing contexts } O \in \mathbb{O}$$
- (2) (must or upper or Smyth)

$$e_0 \cong_2 e_1 \quad \text{iff} \quad Obs(O[e_0]) = \mathbf{s} \Leftrightarrow Obs(O[e_1]) = \mathbf{s} \text{ for all observing contexts } O \in \mathbb{O}$$
- (3) (may or lower or Hoare)

$e_0 \cong_3 e_1$ iff $Obs(O[e_0]) = \mathbf{f} \Leftrightarrow Obs(O[e_1]) = \mathbf{f}$ for all observing contexts $O \in \mathbb{O}$

By construction each of these equivalence relations is a congruence.

Theorem (congruence):

$$e_0 \cong_j e_1 \Rightarrow C[e_0] \cong_j C[e_1] \quad \text{for } j \in \{1, 2, 3\}$$

4.3 Partial Collapse

Note that may-equivalence (\cong_3) is determined by computation trees (that is by quantification over finite sequences rather than paths), since all events are observed after some finite amount of time. Consequently this relation is independent of whether or not fairness is required. Since fairness sometimes makes proving equivalences more difficult, it is useful that may-equivalence can always be proved ignoring the fairness assumption. The other two equivalences are sensitive to choice of paths admitted as computations. In particular when fairness is required, as in our model, \cong_2 is in fact the same as \cong_1 . In models without the fairness requirement, they are distinct. In either case, \cong_3 is distinct from \cong_1 and \cong_2 .

Theorem (partial collapse):

- (1 = 2) $e_0 \cong_2 e_1$ iff $e_0 \cong_1 e_1$
- (1 \Rightarrow 3) $e_0 \cong_1 e_1$ implies $e_0 \cong_3 e_1$
- (3 $\not\Rightarrow$ 1) $e_0 \cong_3 e_1$ does not imply $e_0 \cong_1 e_1$

To demonstrate (1 = 2) we consider a fixed pair of expressions e_0, e_1 and categorize their closing configuration contexts O according to what observations are made by $O[e_0]$ and $O[e_1]$. We say O is labelled $o : o'$ for $o, o' \in \{\mathbf{s}, \mathbf{sf}, \mathbf{f}\}$ just if $Obs(O[e_0]) = o$ and $Obs(O[e_1]) = o'$. This partitions the observing configuration contexts of e_0 and e_1 into nine sets labeled $o : o'$ for $o, o' \in \{\mathbf{s}, \mathbf{sf}, \mathbf{f}\}$.

Lemma (sets) characterizes the various possibilities for equivalence in terms of which sets must be empty.

Lemma (sets):

- $e_0 \cong_1 e_1$ iff at most the sets labeled $\mathbf{s} : \mathbf{s}$, $\mathbf{sf} : \mathbf{sf}$, and $\mathbf{f} : \mathbf{f}$ are non-empty.
- $e_0 \cong_2 e_1$ iff the sets labeled $\mathbf{s} : \mathbf{sf}$, $\mathbf{sf} : \mathbf{s}$, $\mathbf{s} : \mathbf{f}$, $\mathbf{f} : \mathbf{s}$ are all empty.
- $e_0 \cong_3 e_1$ iff the sets labeled $\mathbf{s} : \mathbf{f}$, $\mathbf{f} : \mathbf{s}$, $\mathbf{sf} : \mathbf{f}$, $\mathbf{f} : \mathbf{sf}$ are all empty.

This characterization is summarized in the picture below. Here \times indicates that the set must be empty, and \checkmark indicates that the set might be non-empty. The two *'d cases in \cong_2 are cases in which sets are allowed in theory to be non-empty, but lemma (f.sf) below shows they in fact never arise.

		\cong_1			\cong_2			\cong_3		
		e_1			e_1			e_1		
		\mathbf{s}	\mathbf{sf}	\mathbf{f}	\mathbf{s}	\mathbf{sf}	\mathbf{f}	\mathbf{s}	\mathbf{sf}	\mathbf{f}
e_0	\mathbf{s}	\checkmark	\times	\times	\checkmark	\times	\times	\checkmark	\checkmark	\times
	\mathbf{sf}	\times	\checkmark	\times	\times	\checkmark	*	\checkmark	\checkmark	\times
	\mathbf{f}	\times	\times	\checkmark	\times	*	\checkmark	\times	\times	\checkmark

The key to collapsing \cong_2 into \cong_1 is the observation that if $Obs(O[e_0]) = \mathbf{f}$ and

$Obs(O[e_1]) = \mathbf{sf}$ it is always possible to construct a O^* such that $Obs(O^*[e_0]) = \mathbf{s}$, and $Obs(O^*[e_1]) = \mathbf{sf}$.

Lemma (f.sf): For some e_0, e_1 , if the set labeled $\mathbf{f} : \mathbf{sf}$ is non-empty then the set labeled $\mathbf{s} : \mathbf{sf}$ is non-empty. Symmetrically, if the set labeled $\mathbf{sf} : \mathbf{f}$ is non-empty then the set labeled $\mathbf{sf} : \mathbf{s}$ is non-empty.

Proof (f.sf): Let $O \in \mathbf{f} : \mathbf{sf}$. Form O' by replacing all occurrences of `event()` in O by `send(a, nil)` for some fresh variable a . Let O^* be obtained by adding to O' a message $\langle a \leftarrow t \rangle$ and an actor a where a has the following behavior: If a receives the message t , it executes `event()` and becomes a sink, and if a receives the message `nil`, it just becomes a sink. Recall that a sink is an actor that ignores its message and becomes a sink. We claim $O^* \in \mathbf{s} : \mathbf{sf}$. If $O[e_0]$ never executes `event()`, then in any fair complete computation, the t message will be received by a , so $O^*[e_0]$ will always execute `event()`. If $O[e_1]$ executes `event()` in some computation, then in the corresponding computations for $O^*[e_1]$, sometimes `nil` will be received by a before t is received and sometimes it won't, hence $O^*[e_1]$ will execute `event()` in some computations, but not in all. \square

Proof (partial collapse):

$1 = 2$ Assume $e_0 \cong_2 e_1$. Then by (sets) the sets labeled $\mathbf{s} : \mathbf{sf}$, $\mathbf{sf} : \mathbf{s}$, $\mathbf{s} : \mathbf{f}$, $\mathbf{f} : \mathbf{s}$ are all empty. By (f.sf) $\mathbf{f} : \mathbf{sf}$ and $\mathbf{sf} : \mathbf{f}$ must also be empty. Hence by (sets), $e_0 \cong_1 e_1$. $\square_1 = 2$

$1 \Rightarrow 3$ By (sets) $\square_1 \Rightarrow 3$

$3 \not\Rightarrow 1$ We construct expressions e_0, e_1 such that $e_0 \cong_3 e_1$, but $\neg(e_0 \cong_2 e_1)$. Let e_0 create an actor that sends a message (say `nil`) to an external actor a and becomes a sink, and let e_1 create an actor that may or may not send a message `nil` to a depending on a coin flip (there are numerous methods of constructing coin flipping actors), and also then becomes a sink. Let O be an observing configuration context that with an actor a that executes `event` just if `nil` is received. Then $Obs(O[e_0]) = \mathbf{s}$ but $Obs(O[e_1]) = \mathbf{sf}$, so $\neg(e_0 \cong_2 e_1)$. To show that $e_0 \cong_3 e_1$, show for arbitrary O that some path in the computation of $O[e_0]$ contains an event iff some path in the computation of $O[e_1]$ contains an event. This is easy, because when e_1 's coin flip indicates `nil` is sent, the computation proceeds identically to e_0 's computation. $\square_3 \not\Rightarrow 1$

\square

Hereafter, \cong (observational equivalence) will be used as shorthand for either \cong_1 or \cong_2 .

The fairness requirement is critical in the proof of $(1 = 2)$. For example in CCS, where fairness is not assumed, no such collapse of \cong_2 to \cong_1 occurs. If we omitted the fairness requirement we could make more \cong -distinctions between actors. For example, let a_0 be a sink. Let a_1 be an actor that also ignores its messages and becomes the same behavior, but it continues executing an infinite loop. The infinite looping actor could prevent the rest of the configuration's computation from progressing. In the presence of fairness this could not happen, so the two are equivalent. Thus fairness allows modular reasoning about liveness properties: one can reason about the behavior of individual actors without worrying about whether composition with another would cause such failures.

4.4 Equivalence of Configurations

Now we extend the notion of observational equivalence to configurations.

Definition (Observing Configurations): For an actor configuration $\kappa = \langle\langle \alpha \mid \mu \rangle\rangle_x^\rho$ the observing configurations are configurations over the extended language of the form $\kappa' = \langle\langle \alpha' \mid \mu' \rangle\rangle_x^\rho$ such that κ' is composable (in the sense of §2.3) with κ .

We are interested in observing internal event transitions rather than interactions with the environment. Thus we define an operation $\text{Hide}(\kappa)$ hiding all the receptionists of a configuration.

Definition ($\text{Hide}(\kappa)$): $\text{Hide}(\langle\langle \alpha \mid \mu \rangle\rangle_x^\rho) = \langle\langle \alpha \mid \mu \rangle\rangle_x^\emptyset$

Definition ($\kappa_0 \cong \kappa_1$): For $\kappa_0 = \langle\langle \alpha_0 \mid \mu_0 \rangle\rangle_x^\rho$ and $\kappa_1 = \langle\langle \alpha_1 \mid \mu_1 \rangle\rangle_x^\rho$, $\kappa_0 \cong \kappa_1$ iff $\text{Obs}(\text{Hide}(\kappa_0 \parallel \kappa')) = \text{Obs}(\text{Hide}(\kappa_1 \parallel \kappa'))$ for all observing configurations κ' for κ_j , $j < 2$.

Note the following. Firstly, $\text{Hide}(\kappa \parallel \kappa')$ is a closed configuration for observer κ' for κ . Secondly, no two closed configurations can be distinguished by an external observer.

Using observational equivalence configurations we can extend the property of congruence with respect to expression construction to congruence with respect to configuration construction. Namely, replacing an expression occurring in a configuration by an observationally equivalent one yields an equivalent configuration.

Theorem (exp-cfig): If $e_0 \cong e_1$ then

- (i) $\kappa_0 = \langle\langle \alpha, [C[e_0]]_a \mid \mu \rangle\rangle_x^\rho \cong \langle\langle \alpha, [C[e_1]]_a \mid \mu \rangle\rangle_x^\rho = \kappa_1$
- (ii) $\kappa'_0 = \langle\langle \alpha, (\lambda x.C[e_0])_a \mid \mu \rangle\rangle_x^\rho \cong \langle\langle \alpha, (\lambda x.C[e_1])_a \mid \mu \rangle\rangle_x^\rho = \kappa'_1$

Proof: (i) We need to show that $\text{Obs}(\text{Hide}(\kappa_0 \parallel \kappa')) = \text{Obs}(\text{Hide}(\kappa_1 \parallel \kappa'))$ for any observing κ' . Note however that $\text{Hide}(\kappa_0 \parallel \kappa') = O[e_0]$ and $\text{Hide}(\kappa_1 \parallel \kappa') = O[e_1]$ for some $O \in \mathbb{O}$, so the result follows directly from the definition of \cong .

(ii) We need to show that $\text{Obs}(\text{Hide}(\kappa'_0 \parallel \kappa')) = \text{Obs}(\text{Hide}(\kappa'_1 \parallel \kappa'))$ for any observing κ' . Pick any $\pi_0 \in \mathcal{F}(\text{Hide}(\kappa'_0 \parallel \kappa'))$, then we must find $\pi_1 \in \mathcal{F}(\text{Hide}(\kappa'_1 \parallel \kappa'))$ such that $\text{obs}(\pi_0) = \text{obs}(\pi_1)$. There are two cases to consider. Either actor a never becomes active, or it becomes active first after k steps of computation. In the first case, the e_i are never touched, so both computations proceed uniformly, thus their observation and fairness behavior both correspond. In the second case, consider the step where a receives its first message:

$$\begin{aligned} & \langle\langle \alpha', (\lambda x.C[e_0])_a \mid \mu, \langle a \leftarrow cv \rangle \rangle\rangle \\ & \xrightarrow{\langle \text{rcv} : a, cv \rangle} \langle\langle \alpha', [\text{app}(\lambda x.C[e_0], cv)]_a \mid \mu \rangle\rangle = O[e_0] \end{aligned}$$

Factor $\pi_0 = \nu[e_0] * \pi'_0$, where $\pi'_0 \in \mathcal{F}(O[e_0])$ and $\nu[]$ denotes a sequence where each configuration in the sequence contains a hole that computes uniformly in the hole. Thus, $\text{Obs}(O[e_0]) = \text{Obs}(O[e_1])$ because $e_0 \cong e_1$ and O is a configuration

context. This means by the definition of *Obs* there is a path $\pi'_1 \in \mathcal{F}(O[e_1])$, such that $obs(\pi'_0) = obs(\pi'_1)$. Let $\pi_1 = \nu[e_1] * \pi'_1$. Then, $\pi_1 \in \mathcal{F}(\text{Hide}(\kappa_1 \parallel \kappa'))$, since by construction it is a computation for $(\text{Hide}(\kappa_1 \parallel \kappa'))$, and because π fair implies $\nu * \pi$ is fair for any ν such that $\nu * \pi$ is a computation path. Moreover, $obs(\pi_0) = obs(\pi_1)$ because any event transitions in $\nu[e_0]$ also occur in $\nu[e_1]$, and $obs(\pi'_0) = obs(\pi'_1)$ by hypothesis.

□

5 Laws of Expression Equivalence

With a notion of equivalence on actor expressions defined, a library of useful equivalences can be established. The first part of this section contains a collection of purely functional laws that continue to hold in the actor setting. The second part contains a collection of laws for manipulating expressions that involve actor primitives. These laws are established in §5.

5.1 Functional Laws

Since our reduction rules preserve the evaluation semantics of the embedded functional language, many of the equational laws for this language (cf. (Talcott, 1993b)) continue to hold in the full actor language. A first simple observation is two communicable values are observationally equivalent iff they are the same value expression.

Lemma (cv):

$$cv_0 \cong cv_1 \Leftrightarrow cv_0 = cv_1$$

Proof : The if direction is trivial. The only-if direction is proved by exhibiting an observing context that distinguishes expressions that are not equal. Clearly both must be atoms, or variables, or pairs, otherwise they can be distinguished using *eq* and *ispr*. For example,

$$O = \left\langle \left\langle [\text{if}(\text{eq}(cv_0, \bullet), \text{event}, \text{nil})]_a \mid \right\rangle \right\rangle$$

distinguishes the atom cv_0 from any non-atom (and any other atom). Similarly,

$$O = \left\langle \left\langle [\text{let}\{x := 0\}\text{let}\{y := 1\}\text{if}(\text{eq}(x, \bullet), \text{event}, \text{nil})]_a \mid \right\rangle \right\rangle$$

distinguishes the variables x, y . Similarly, if both are pairs, we can construct contexts to distinguish differences in the components. □

The laws of the computational lambda calculus (Moggi, 1988), and the laws for conditional and pairing continue to hold in the actor setting.

Theorem (functional laws):

- (beta-v) $\text{app}(\lambda x.e, v) \cong e[x := v]$
- (ift) $\text{if}(v, e_1, e_2) \cong e_1$ if $v \in (\text{At} - \{\text{nil}\}) \cup \mathbb{L} \cup \text{pr}(\mathbb{V}, \mathbb{V})$
- (ifn) $\text{if}(\text{nil}, e_1, e_2) \cong e_2$
- (ifelim) $\text{if}(v, e, e) \cong e$

- (iflam) $\lambda x.\mathbf{if}(v, e_1, e_2) \cong \mathbf{if}(v, \lambda x.e_1, \lambda x.e_2) \quad x \notin \text{FV}(v)$
- (isprt) $\mathbf{ispr}(\mathbf{pr}(v_0, v_1)) \cong \mathbf{t}$,
- (isprn) $\mathbf{ispr}(v) \cong \mathbf{nil} \quad v \in \mathbb{A}\mathbf{t} \cup \mathbb{L}$
- (fst) $1^{\text{st}}(\mathbf{pr}(v_0, v_1)) \cong v_0$
- (snd) $2^{\text{nd}}(\mathbf{pr}(v_0, v_1)) \cong v_1$

Each of these laws (except for **(iflam)**) is a consequence of the following operational law.

Theorem (red-exp):

$$e_0 \xrightarrow{\lambda} e_1 \Rightarrow e_0 \cong e_1$$

The theorem **(rcx)** is a special case of a theorem proved in (Talcott, 1989).

Theorem (rcx): If R is a reduction context, then

- (letx) $\mathbf{let}\{x := e\}R[x] \cong R[e]$
- (let.dist) $R[\mathbf{let}\{x := e\}e_0] \cong \mathbf{let}\{x := e\}R[e_0]$
- (if.dist) $R[\mathbf{if}(e, e_1, e_2)] \cong \mathbf{if}(e, R[e_1], R[e_2])$

In fact these laws can be derived from **(beta-v)**, the **if** laws and the following special instances.

- (app) $e_0(e_1) \cong (\lambda f.f(e_1))(e_0)$
- (cmps) $f(g(e)) \cong (\lambda x.f(g(x)))(e)$
- (id) $\mathbf{app}(\lambda x.x, e) \cong e$

Some useful corollaries of **(rcx)** are the following.

Corollary (uni-rcx):

- (let.arg) $v(\mathbf{let}\{x := e_0\}e_1) \cong \mathbf{let}\{x := e_0\}v(e_1)$
- (if.if) $\mathbf{if}(\mathbf{if}(e_0, e_1, e_2), e_a, e_b) \cong \mathbf{if}(e_0, \mathbf{if}(e_1, e_a, e_b), \mathbf{if}(e_2, e_a, e_b))$

The above laws are really about equivalence of reduction contexts. They are instances of the operational law **(red-rcx)**. Two reduction contexts are considered equivalent if placing an arbitrary expression in the redex hole results in equivalent expressions. The law **(red-rcx)** says that if two reduction contexts have a common λ -reduct when the redex hole is filled with a fresh variable (standing for an arbitrary value expression), then they are equivalent.

Theorem (red-rcx): If there is some e' such that $R_0[x] \xrightarrow{\lambda} e'$ and $R_1[x] \xrightarrow{\lambda} e'$ where x is a fresh variable, then $R_0[e] \cong R_1[e]$ for any e .

We also note that any expressions that hang (reduce in a finite number of lambda steps to a stuck state) or have infinite lambda computations are observationally equivalent. If the reductions could involve non-lambda steps the result clearly fails since they could have different effects such as message sends that other actors in the configuration could observe. We let $\mathbf{stuck} \in \mathbf{Hang}$ be a prototypical stuck

expression, for example $\text{app}(0, 0)$, and let $\text{bot} \in \text{Infin}$ be a prototypical expression with infinite computation, for example $\text{app}(\lambda x. \text{app}(x, x), \lambda x. \text{app}(x, x))$.

To make these ideas more precise we define *Hang* and *Infin* as follows.

Definition (Hang): Let *Hang* be the set of non-value expressions such that (every closed instance) lambda reduces (i.e. \mapsto^λ in possibly 0 steps) to a stuck state – an expression e' that decomposes as $R[r]$ where r is a functional redex (i.e any non-actor redex) that does not reduce.

Definition (Infin): Let *Infin* the the set of (non-value) expressions e such that (every closed instance) has an infinite lambda reduction sequence. Thus $e \in \text{Infin}$ just if we can find e_j for $j \in \mathbb{N}$ such that $e_0 = e$ and $e_j \mapsto^\lambda e_{j+1}$.

Theorem (hang-infin): If $e_0, e_1 \in \text{Hang} \cup \text{Infin}$, then $e_0 \cong e_1$.

5.2 Basic Laws for Actor Primitives

Now we consider the equational properties of the actor primitives, `send`, `become`, `newadr`, and `initbeh`. As is the case for a language with state operations, $\text{seq}(e, e) \cong e$ fails to hold because the computation e could have effects such as message sends. A stronger analogy exists between the actor primitives and the reference primitives $\{\text{mk}, \text{get}, \text{set}\}$: `newadr` is an allocation primitive analogous to `mk`; `initbeh` and `become` update or alter state; while the effect of `send` depends on the state in a way analogous to `get`. There are limits to this analogy, for example `send` does not return anything of interest. In fact `send`, `become`, and `initbeh` all return `nil` as values. Thus

$$\text{(triv)} \quad \vartheta(\bar{x}) \cong \text{seq}(\vartheta(\bar{x}), \text{nil}) \quad \text{for } \vartheta \in \{\text{send}, \text{become}, \text{initbeh}\}$$

That `newadr` is an allocation primitive analogous to `mk` manifests itself in its properties. However, because we have not allowed clones to initialize newly spawned actors, one characteristic property of allocation fails to holds. Namely,

(non-delay)

$$\text{let}\{y := e_0\}\text{let}\{x := \text{newadr}()\}e_1 \not\cong \text{let}\{x := \text{newadr}()\}\text{let}\{y := e_0\}e_1$$

where x is not free in e_0 . Since, if evaluation of e_0 executes a `become` and e_1 is of the form `initbeh`(x, v), then the left-hand side evaluation of e_1 will succeed, while the right-hand side evaluation of e_1 will suspend.

Once allocated, an actor behavior is updated by either `become` or `initbeh`. In analogy with `set` both `become` and `initbeh` satisfy certain, slightly different, cancellation laws:

$$\text{(can-b)} \quad \text{seq}(\text{become}(v_0), \text{become}(v_1)) \cong \text{seq}(\text{become}(v_0), \text{nil}) \cong \text{become}(v_0)$$

$$\begin{aligned} \text{(can-i)} \quad \text{seq}(\text{initbeh}(v, v_0), \text{initbeh}(v, v_1)) &\cong \text{seq}(\text{initbeh}(v, v_0), \text{stuck}) \\ &\cong \text{seq}(\text{initbeh}(v, v_0), \text{bot}) \end{aligned}$$

Note the difference between these two principles. In the case of `become` the second call is equivalent to `nil`, while in the case of `initbeh` it is stuck (which is equivalent to diverging).

5.2.1 Commuting Operations

How the effects of the actor primitives interact with one another is of paramount importance. We have seen some aspects of this interaction above. We now study the interactions more systematically.

Definition (commutes): We say two operations ϑ_0 and ϑ_1 commute if

$$\mathbf{let}\{x_0 := \vartheta_0(\bar{y})\}\mathbf{let}\{x_1 := \vartheta_1(\bar{z})\}e \cong \mathbf{let}\{x_1 := \vartheta_1(\bar{z})\}\mathbf{let}\{x_0 := \vartheta_0(\bar{y})\}e$$

for all $e \in \mathbb{E}$, $x_0 \notin \bar{z}$, $x_1 \notin \bar{y}$ and x_0 distinct from x_1 . Similarly we say two expressions e_0 and e_1 commute iff

$$\mathbf{let}\{x_0 := e_0\}\mathbf{let}\{x_1 := e_1\}e \cong \mathbf{let}\{x_1 := e_1\}\mathbf{let}\{x_0 := e_0\}e$$

if the obvious hygiene conditions hold.

`newadr` commutes with every operation except `become`. For example the expressions

$$e_0 = \mathbf{let}\{y := \mathbf{newadr}()\}\mathbf{let}\{z := \mathbf{become}(b)\}\mathbf{initbeh}(y, b')$$

$$e_1 = \mathbf{let}\{z := \mathbf{become}(b)\}\mathbf{let}\{y := \mathbf{newadr}()\}\mathbf{initbeh}(y, b')$$

are not equivalent, since the first will always fail to execute the initialization and the second will always succeed. A distinguishing context is

$$\langle\langle (\lambda x.\mathbf{event}())_{a_0}, [\mathbf{seq}(\bullet, \mathbf{send}(a_0, 0))]_a \mid \emptyset \rangle\rangle$$

If we allowed clones to initialize, then `newadr` would also commute with `become`. On the other hand, by (**can-b**) and (**can-1**) neither `become` nor `initbeh` commute with themselves, since this amounts to equivalence of two becomes (or initializations) with different behaviors. The remaining operation `send`, like `newadr`, does commute with itself:

$$\text{(com-ss)} \quad \mathbf{seq}(\mathbf{send}(v_0, v_1), \mathbf{send}(v_2, v_3)) \cong \mathbf{seq}(\mathbf{send}(v_2, v_3), \mathbf{send}(v_0, v_1))$$

`send` also commutes with `become`

$$\text{(com-sb)} \quad \mathbf{seq}(\mathbf{send}(a_0, v_0), \mathbf{become}(v_1)) \cong \mathbf{seq}(\mathbf{become}(v_1), \mathbf{send}(a_0, v_0))$$

The question of whether or not two distinct operations commute is simplified by the observation, captured in (**partial**), that a computation may have observable effects even if a subcomputation diverges. This is in contrast to the sequential case, where an effect of a subcomputation is only observable if the computation completes.

Lemma (partial): If ϑ is not a total operation, then ϑ does not commute with `send`, `become` or `initbeh`.

Proof (partial): If $\vartheta(\bar{y})$ diverges, then

$$\mathbf{let}\{x := \vartheta(\bar{y})\}\mathbf{let}\{x_1 := \mathbf{send}(a, v)\}e$$

will not execute the `send`, whereas

$$\mathbf{let}\{x_1 := \mathbf{send}(a, v)\}\mathbf{let}\{x := \vartheta(\bar{y})\}e$$

will execute the `send`. Consequently the two expressions are easily distinguished. Similarly with the two operations `become` and `initbeh`. \square

Since `initbeh` is partial, it does not commute with either `send` or `become`. For example

$$e_0 = \text{seq}(\text{initbeh}(a_0, b_0), \text{send}(a_1, 0))$$

$$e_1 = \text{seq}(\text{send}(a_1, 0), \text{initbeh}(a_0, b_0))$$

are distinguished by

$$O_0 = \langle\langle (\lambda x.\text{event}())_{a_1}, (?_{a_1})_{a_0}, [\bullet]_a \mid \mu \rangle\rangle$$

for $a \neq a_1$, or by

$$O_1 = \langle\langle (\lambda x.\text{event}())_{a_1}, (b)_{a_0}, [\bullet]_a \mid \mu \rangle\rangle$$

Also,

$$e_0 = \text{seq}(\text{initbeh}(a_0, b_0), \text{become}(\lambda x.\text{send}(a_1, 0)))$$

$$e_1 = \text{seq}(\text{become}(\lambda x.\text{send}(a_1, 0)), \text{initbeh}(a_0, b_0))$$

are distinguished by O_0, O_1 if μ contains $\langle a \Leftarrow 0 \rangle$.

(partial) emphasizes that the valid equations for actor primitives are sensitive to the details of when we check for ill-formed redexes. For example if we restricted the `send` redex to avoid ill-formed messages (**com-ss,com-sb**) would no longer hold.

We summarize these results in the following:

Lemma (commutes):

	n	s	i	b
n	+	+	+	-
s	+	+	-	+
i	+	-	-	-
b	-	+	-	-

(n) `newadr` commutes with `send`, `newadr`, and `initbeh`, but not with `become`.

(s) `send` commutes with `send`, `become`, and `newadr`, but not with `initbeh`.

(i) `initbeh` commutes with `newadr`, but not with `send`, `become`, and `initbeh`.

(b) `become` commutes with `send`, but not with `initbeh`, `newadr`, or `become`.

Note that the remaining operations in \mathbb{F} (i.e. the arithmetic operations, branching `br`, and the pairing `ispr, pr, 1st, 2nd`) are all context insensitive, and thus those that are total commute with all other operations. In the case if this is perhaps worth pointing out:

Lemma (commutes-if): If ϑ commutes with e_0 and e_1 , then it also commutes with `if`(z, e_0, e_1)

Proof : This follows from **(if-lam,if-dist)**. \square

Using these basic principles we can prove more complex properties, the following theorem being the most obvious.

Theorem (commutes): Suppose that e_0 and e_1 are built up from \mathbb{V} using only the operations, **if** and **let**. Furthermore suppose every operation occurring in e_0 commutes with every operation occurring in e_1 . Then

$$\mathbf{let}\{x_0 := e_0\}\mathbf{let}\{x_1 := e_1\}e \cong \mathbf{let}\{x_1 := e_1\}\mathbf{let}\{x_0 := e_0\}e$$

provided x_j is not free in e_{1-j} for $j < 2$,

Proof (commutes): The proof is by induction on the complexity of e_0 . We sketch the induction step. We may assume, without loss of generality, that e_0 decomposes into $R[\vartheta(\bar{y})]$ with R being non-trivial. Then

$$\begin{aligned} & \mathbf{let}\{x_0 := e_0\}\mathbf{let}\{x_1 := e_1\}e \cong \\ & \cong \mathbf{let}\{x_0 := R[\vartheta(\bar{y})]\}\mathbf{let}\{x_1 := e_1\}e \\ & \quad \text{by hypothesis} \\ & \cong \mathbf{let}\{x_0 := \mathbf{let}\{z := \vartheta(\bar{y})\}R[z]\}\mathbf{let}\{x_1 := e_1\}e \\ & \quad \text{by (cong) and (letx)} \\ & \cong \mathbf{let}\{z := \vartheta(\bar{y})\}\mathbf{let}\{x_0 := R[z]\}\mathbf{let}\{x_1 := e_1\}e \\ & \quad \text{by (let.dis)} \\ & \cong \mathbf{let}\{z := \vartheta(\bar{y})\}\mathbf{let}\{x_1 := e_1\}\mathbf{let}\{x_0 := R[z]\}e \\ & \quad \text{by the induction hypothesis and (cong)} \\ & \cong \mathbf{let}\{x_1 := e_1\}\mathbf{let}\{z := \vartheta(\bar{y})\}\mathbf{let}\{x_0 := R[z]\}e \\ & \quad \text{by the induction hypothesis and (cong)} \\ & \cong \mathbf{let}\{x_1 := e_1\}\mathbf{let}\{x_0 := \mathbf{let}\{z := \vartheta(\bar{y})\}R[z]\}e \\ & \quad \text{by (let.dis)} \\ & \cong \mathbf{let}\{x_1 := e_1\}\mathbf{let}\{x_0 := e_0\}e \\ & \quad \text{by (cong) and (letx)} \end{aligned}$$

□

Note that the theorem fails in the case when the expressions contain **app** and λ due to the possibility of divergence.

6 Proving Expression Equivalence

In this section we develop methods for proving expressions observationally equivalent. The methods are a variation on the techniques used to prove the **(ciu)** theorem (a form of context lemma) in (Mason and Talcott, 1991).

Before outlining the methods we motivate the difficulty of establishing equivalences. A relatively simple case is $\mathbf{succ}(0) \cong 1$. It is simple for two reasons: there are no free variables occurring, and only one step of computation separates $\mathbf{succ}(0)$

and 1. By the definition of \cong , we need to establish

$$Obs(O[\text{succ}(0)]) = Obs(O[1])$$

for all observing contexts O . To establish this, we construct, for each computation path $\pi_0 \in \mathcal{F}(O[\text{succ}(0)])$, a $\pi_1 \in \mathcal{F}(O[1])$ such that $obs(\pi_0) = obs(\pi_1)$ and conversely. We call such a construction a *path correspondence*. Informally, the path correspondence is constructed as follows. First consider how from a path π_0 in $\mathcal{F}(O[\text{succ}(0)])$ we obtain a path in $\mathcal{F}(O[1])$. At each point in π_0 where the $\text{succ}(0)$ is reduced to 1, we remove this step, giving path π_1 . Describing this operation in detail requires care, for there could be other independent occurrences of the computation of subterm $\text{succ}(0)$ in O which are not to be removed. We then can show that π_1 is a computation for $O[1]$, with the same observable outcome and same fairness property as π_0 . The converse construction is similar, except steps computing $\text{succ}(0)$ are inserted into π_0 each time the 1 first appears in a reduction context.

These two expressions differ by only one step of computation; in general they could differ by more than one step, and could both reduce to a common reduct rather than one reducing to the other, e.g. $\text{pred}(\text{succ}(1)) \cong \text{succ}(0)$. The complication arising from this case is the two-step execution of $\text{pred}(\text{succ}(1))$ may be interleaved with computations of other actors and a local replacement thus not possible. To solve this problem the computation path is put in an equivalent canonical form with both steps adjacent. In general we place as many steps of an individual actor adjacent as possible.

A complication also arises in proving equations that may contain free variables, for instance $1^{\text{st}}(\text{pr}(x, 0)) \cong x$. Such expressions may be self-substituted: if $1^{\text{st}}(\text{pr}(x, 0))$ occurs in the local context $\text{app}(\lambda y.\text{app}(y, y), \lambda x.\bullet)$, upon computing the free x in $1^{\text{st}}(\text{pr}(x, 0))$ will be replaced with $\lambda x.1^{\text{st}}(\text{pr}(x, 0))$. This means the necessary replacements are not flat but may be nested. A notion of generalized hole is introduced to account for this nesting.

In general we give methods for establishing three different varieties of expression equivalence; the above informal description describes only the first variant. The three variants are as follows.

- (1) The first variant treats equivalence of expressions that have a common reduct – i.e. expressions that step in 0 or more steps to the same expression having the same effects (sends, becomes, creation of new actors, initializing new actors). This is called the *common reduct case*.
- (2) The second variant is an elaboration of the first, treating expressions that reduce to lambda abstractions that are application equivalent – i.e. have a common reduct when applied to any value. This is called the *two-stage reduction case*.
- (3) The third variant treats equivalence of reduction contexts. This is called the *equivalence of reduction contexts case*.

We provide examples of the use of these techniques by using them to establish the equational laws stated in §4.

6.1 The General Method

Each of these three methods is based on the idea of using *configuration templates* to establish a correspondence between the fair computations of configurations containing the entities to be proved equivalent. A configuration template is simply a configuration with holes, i.e. schematic variables, that may be instantiated by various sorts of syntactic entities. Observing contexts correspond to a special case of configuration templates.

The first step then is to choose a class of configuration templates CT such that $e_0 \cong e_1$ if $Obs(ct[e_0]) = Obs(ct[e_1])$ for all templates $ct \in CT$. To establish the equality of observations, it is sufficient to construct, for each $\pi_0 \in \mathcal{F}(ct[e_0])$, a $\pi_1 \in \mathcal{F}(ct[e_1])$ such that $obs(\pi_0) = obs(\pi_1)$ and conversely. We call such a construction a *path correspondence*. The crucial fact concerning configuration templates is that one can compute symbolically with them in the sense that computation is parametric in the holes. We call this form of computation *uniform computation* or *uniform reduction*.

A suitable class of configuration templates is obtained by extending each syntactic class to allow holes and defining appropriate notions of hole filling. Decomposition theorems and schematic reduction rules are then developed. In each of the three methods the only essential difference is the type and number of holes needed:

- (1) For the common reduct case we define templates by adding a single hole, \circ , for expressions. We call this hole an *expression hole*.
- (2) For the two stage reduction case we need not only a hole for expressions, but also a countable family of holes for lambda abstractions. We call these holes *abstraction holes* and they are denoted by \triangleright_j for $j \in \mathbb{N}$. Note that these holes are filled by values, specifically by lambda abstractions, not simply by expressions.
- (3) For the equivalence of reduction contexts we need an entirely new kind of hole, \diamond , for reduction contexts. We call it a *reduction context hole*. Note that occurrences of holes will be filled by reduction contexts and are *not* to be confused with redex holes. As far as we are aware the introduction of holes that are filled by contexts is completely novel.

For each variant, syntactic classes X are annotated with the signs of the sorts of holes they contain: ${}^\circ X$ for expression holes; ${}^{\circ\triangleright} X$ for expression and lambda abstraction holes; and ${}^\diamond X$ for reduction context holes. We prefix the names of these classes by E-, LE-, or R- respectively. Thus E-expressions are expression templates with holes for expressions, ${}^\circ\mathbb{E}$ is the set of E-expressions, and we let ${}^\circ e$ range over ${}^\circ\mathbb{E}$. A similar convention holds for the other syntactic classes and hole types.

The idea underlying the construction of a path correspondence to establish equivalence is the same for each of the three cases. It relies on the ability to localize differences in computations as multi-steps (§2.2.3), and to use holes to formalize the aspects of computation that are independent of the local differences. Consider the case of proving expressions equivalent using templates with expression holes. We consider fair computation paths starting from an E-configuration with holes filled by one of the expressions, say e_0 . For each such path, π_0 , we show how to obtain a sequence of E-configurations satisfying two conditions. The first is that filling the

holes in the sequence of E-configurations with e_0 (and filling in transition labels) yields π_0 . The second is that filling the holes in the sequence of E-configurations with e_1 (and filling in transition labels) yields a fair computation path with the same observation. The other two cases are simple variations on this idea.

One of the keys ideas in uniform computation is to insure that transitions commute with hole filling; except of course when the *hole is touched*, i.e. information about the contents of the hole is required to carry out the step. Consider the schematic redex $\text{app}(\lambda x.\bullet, v)$. We need a notation that allows us to carry out this reduction in such a way that filling the hole and then reducing gives the same result as reducing and then filling the hole. For this purpose we associate with each hole a substitution to be applied when the hole is filled. The domain of the substitution also determines the variables of an expression that are trapped at the hole. This localizes trapping and allows renaming of lambda-variables even in the presence of holes (which is not the case for traditional notions of expression context). A detailed development of this notation can be found in (Talcott, 1991; Talcott, 1993a). We use $\circ[\sigma]$ to denote an expression hole with associated substitution σ (which may in turn have expressions holes in its range), a similar notation holds for the other classes of holes: $\triangleright_j[\sigma]$ for abstraction holes, and $\diamond[\sigma]$ for reduction context holes.

To simplify definitions of syntactic classes we treat app on a par with elements of \mathbb{F}_2 . We use Θ_n for syntactic operations of arity n , and Θ_n^e to indicate the operations of the extended language (i.e. Θ_0 extended to include event). Thus:

Definition ($\Theta_n \Theta_n^e$):

$$\begin{aligned} \Theta_2 &= \mathbb{F}_2 \cup \{\text{app}\} & \Theta_n &= \mathbb{F}_n & \text{for } n \neq 2 \\ \Theta_0^e &= \Theta_0 \cup \{\text{event}\} & \Theta_n^e &= \Theta_n & \text{for } n \neq 0 \end{aligned}$$

As the last technical detail, we make precise the sense in which we are able to localize differences in computations as multi-steps. We first define the notion of *thread fragment*, and then show that any family of disjoint thread fragments in a computation path can be regroup as multi-step without effecting the observation made of the path.

A thread fragment, I , at a in π is a finite subsequence of exec transitions of π with focus actor a or a clone of a created by a become such that any gaps in the sequence are transitions with some other focus, or at a after a new message receipt.

Definition (thread fragment): Let

$$\begin{aligned} \mathcal{F}(\kappa) &= [\kappa_i \xrightarrow{i} \kappa_{i+1} \mid i < \infty], \text{ and } \pi \in \mathcal{F}(\kappa), \\ I &= [i_j \mid j < n] \text{ such that } j < j' < n \Rightarrow i_j < i_{j'}, \\ L(I, \pi) &= [l_j \mid j < n], \text{ the transition sequence corresponding to } I \text{ in } \pi. \end{aligned}$$

Then I is a *thread fragment* at a in π if

- (1) $l_{i_0} = \text{rcv}(a, cv)$ for some cv ,
- (2) $L(I, \pi)$ contains no rcv , in or out , and
- (3) $L(I, \pi)$ is a computation for $\langle\langle \alpha_{i_0} \{a\} \mid \emptyset \rangle\rangle$.

Condition (3) makes explicit that a thread fragment is essentially running the focus actor in a configuration with only itself. With no receives only that actor or its become clones can execute.

Theorem (infinite-macro-steps): Let $\pi = [\kappa_i \xrightarrow{l_i} \kappa_{i+1} \mid i < \infty] \in \mathcal{F}(\kappa)$. Let $I_j = [i_{j,0}, \dots, i_{j,n_j}]$ for $j < J \leq \infty$ be a family (possibly infinite) of thread fragments in π such that

- (a) if $j < j'$ then $i_{j,0} < i_{j',0}$, and
- (b) if $j \neq j'$ the I_j and $I_{j'}$ have empty intersection.

Then there is a bijection, ξ , on ∞ such that letting $\pi' = \text{Cfig}(\kappa_0, [\xi(l_i) \mid i < \infty])$ (recall the definition of Cfig from §2.2.3)

- (1) $\pi' \in \mathcal{F}(\kappa_0)$
- (2) $\xi(i_{j,k+1}) = \xi(i_{j,k}) + 1$ for $j \in J$, and $0 \leq k < n_j$.

Part (2) says that in π' the thread fragments of π marked by I_j for $j \in J$ occur as multi-steps, that is, with no interleaved computation steps. Note that $\text{obs}(\pi') = \text{obs}(\pi)$.

Proof: ξ is computed by induction on the index set, permuting the `exec` steps of each successive fragment across interleaved steps the obvious way. By the definition of thread fragment and the disjointness requirement, we see that permutations only involve moving `exec` steps before steps with distinct focus. Hence the resulting sequence of labels defines a computation. Also the enabledness is not effected by such permutations (except possibly enabling a transition earlier). All transitions that occur in π also occur in π' this means that fairness is also preserved. \square

6.2 Common Reduct Case

We now treat the common reduct case in depth. The other two cases follow in the same manner and we allow ourselves to be a little more terse.

6.2.1 E-Syntax

As mentioned above, syntactic classes, X , with expression holes are indicated by the mark ${}^\circ X$. Metavariables ranging over these classes are indicated by the same mark, and we prefix the names of these classes by E-. Thus we have E-expressions where ${}^\circ e$ ranges over ${}^\circ \mathbb{E}$, E-configurations where ${}^\circ \kappa$ ranges over ${}^\circ \mathbb{K}$, etc. We first define the E- analogs of expression, value expression, and value substitution.

Definition (${}^\circ \mathbb{E}$ ${}^\circ \mathbb{V}$ ${}^\circ \mathbb{S}$):

$$\begin{aligned} {}^\circ \mathbb{V} &= \text{At} \cup \mathbb{X} \cup \lambda \mathbb{X}. {}^\circ \mathbb{E} \cup \text{pr}({}^\circ \mathbb{V}, {}^\circ \mathbb{V}) \\ {}^\circ \mathbb{E} &= {}^\circ \mathbb{V} \cup \Theta_n^e({}^\circ \mathbb{E}^n) \cup \circ[{}^\circ \mathbb{S}] \\ {}^\circ \mathbb{S} &= \mathbb{X} \xrightarrow{f} {}^\circ \mathbb{V} \end{aligned}$$

As before, λ is the only binding operator, and free variables of E-expressions are defined as follows:

Definition ($\text{FV}({}^\circ e)$ $\text{FV}({}^\circ \sigma)$):

$$\text{FV}({}^\circ e) = \begin{cases} \text{FV}({}^\circ \sigma) & \text{if } {}^\circ e = \circ[{}^\circ \sigma] \\ \{{}^\circ e\} & \text{if } {}^\circ e \in \mathbb{X} \\ \text{FV}({}^\circ e_0) - \{z\} & \text{if } {}^\circ e = \lambda z. {}^\circ e_0 \\ \text{FV}({}^\circ e_1) \cup \dots \cup \text{FV}({}^\circ e_n) & \text{if } {}^\circ e = \theta({}^\circ e_1, \dots, {}^\circ e_n) \text{ and } \theta \in \Theta_n^e \end{cases}$$

$$\text{FV}({}^\circ\sigma) = \bigcup_{x \in \text{Dom}({}^\circ\sigma)} \text{FV}({}^\circ\sigma(x))$$

The variables in the domain of occurrences of ${}^\circ\sigma$ are neither free or bound. In particular, renaming of bound variables only applies to the range of a substitution associated with a hole, not to its domain.

Definition (${}^\circ e[{}^\circ\sigma]$ ${}^\circ\sigma_1 \odot {}^\circ\sigma_2$): Substitution is extended to E-expressions as follows:

$${}^\circ e[{}^\circ\sigma] = \begin{cases} \circ[{}^\circ\sigma \odot {}^\circ\sigma'] & \text{if } {}^\circ e = \circ[{}^\circ\sigma'] \\ {}^\circ e & \text{if } {}^\circ e \in \mathbb{X} - \text{Dom}({}^\circ\sigma) \\ {}^\circ\sigma({}^\circ e) & \text{if } {}^\circ e \in \text{Dom}({}^\circ\sigma) \\ \lambda z. {}^\circ e_0[{}^\circ\sigma](\text{Dom}({}^\circ\sigma) - \{z\}) & \text{if } {}^\circ e = \lambda z. {}^\circ e_0 \text{ and } z \notin \text{FV}({}^\circ\sigma) \\ \theta({}^\circ e_0[{}^\circ\sigma], \dots, {}^\circ e_n[{}^\circ\sigma]) & \text{if } {}^\circ e = \theta({}^\circ e_1, \dots, {}^\circ e_n) \text{ and } \theta \in \Theta_n^e \end{cases}$$

$${}^\circ\sigma_1 \odot {}^\circ\sigma_2 = \lambda x \in \text{Dom}({}^\circ\sigma_2). {}^\circ\sigma_2(x)[{}^\circ\sigma_1]$$

As defined here substitution is a partial operation. Using renaming substitutions we can define α renaming in the usual way. We consider E-expressions (and entities containing them) to be equivalent if they differ only by α renaming. Thus, for any substitution we can always choose an α variant so that substitution is defined. Note that such renaming is not possible in the case of traditional contexts where holes have no associated substitution (c.f. (Talcott, 1993a)).

Expression hole filling is defined by induction on the structure of ${}^\circ e$. We let ${}^\circ e[\circ := e]$ be the result of filling expression holes in ${}^\circ e$ with e . Like substitution, we avoid capture of free variables in e by lambda binding. All capture is done at hole occurrences by the associated substitution.

Definition (${}^\circ e[\circ := e]$):

$${}^\circ e[\circ := e] = \begin{cases} e[{}^\circ\sigma[\circ := e]] & \text{if } {}^\circ e = \circ[{}^\circ\sigma] \\ \theta({}^\circ e_1[\circ := e], \dots, {}^\circ e_n[\circ := e]) & \text{if } {}^\circ e = \theta({}^\circ e_1, \dots, {}^\circ e_n) \text{ and } \theta \in \Theta_n^e \\ {}^\circ v & \text{if } {}^\circ e = {}^\circ v \in \text{At} \cup \mathbb{X} \\ \lambda x. {}^\circ e'[\circ := e] & \text{if } {}^\circ e = \lambda x. {}^\circ e' \text{ and } x \text{ not free in } e \end{cases}$$

$${}^\circ\sigma[\circ := e] = \lambda x \in \text{Dom}({}^\circ\sigma). {}^\circ\sigma(x)[\circ := e]$$

The following example illustrates hole filling and variable scoping. Let

$${}^\circ v = \lambda x. \text{if}(x, \circ[\{y := 0\}], z)$$

$${}^\circ\sigma = \{y := {}^\circ v\}$$

$${}^\circ e = \lambda z. \circ[{}^\circ\sigma]$$

$$e = \theta(x, y)$$

then

$$\begin{aligned} {}^\circ v[\circ := e] &= \lambda w. \text{if}(w, \circ[\{y := 0\}][\circ := e], z) && \text{note the change in bound variable} \\ &= \lambda w. \text{if}(w, e[\{y := 0\}], z) \\ &= \lambda w. \text{if}(w, \theta(x, 0), z) = v \end{aligned}$$

$$\begin{aligned} {}^\circ\sigma[\circ := e] &= \{y := {}^\circ v[\circ := e]\} \\ &= \{y := \lambda w. \text{if}(w, \theta(x, 0), z)\} \end{aligned}$$

$$\begin{aligned}
{}^\circ e[\circ := e] &= \lambda z.(\circ[{}^\circ\sigma][\circ := e]) \\
&= \lambda z.(e[{}^\circ\sigma[\circ := e]]) \\
&= \lambda z.\theta(x, \lambda w.\mathbf{if}(w, \theta(x, 0), z))
\end{aligned}$$

The following lemma is the key to developing a notion of uniform computation.

Lemma (fil-subst): Hole filling and substitution commute.

$${}^\circ e[{}^\circ\sigma][\circ := e'] = {}^\circ e[\circ := e'][{}^\circ\sigma[\circ := e']]$$

if $\text{Dom}({}^\circ\sigma) \cap \text{FV}(e') = \emptyset$.

Proof : By induction on the structure of ${}^\circ e$. We assume the names of bound variables in ${}^\circ e$ have been chosen not to conflict with any free variables in e' , or the range of ${}^\circ\sigma$, or the domain of ${}^\circ\sigma$. As examples, we consider the cases where ${}^\circ e$ is a lambda abstraction or a hole. If ${}^\circ e = \lambda z.{}^\circ e_0$ then

$$\begin{aligned}
{}^\circ e[{}^\circ\sigma][\circ := e'] &= (\lambda z.{}^\circ e_0[{}^\circ\sigma])[\circ := e'] \\
&= \lambda z.({}^\circ e_0[{}^\circ\sigma][\circ := e']) \quad \text{by hygiene assumptions} \\
&= \lambda z.({}^\circ e_0[\circ := e'][{}^\circ\sigma[\circ := e']]) \quad \text{by the Induction Hypothesis} \\
&= (\lambda z.{}^\circ e_0[\circ := e'])[{}^\circ\sigma[\circ := e']] \quad \text{by hygiene assumptions} \\
&= {}^\circ e[\circ := e'][{}^\circ\sigma[\circ := e']]
\end{aligned}$$

If ${}^\circ e = \circ[{}^\circ\sigma']$ then

$$\begin{aligned}
{}^\circ e[{}^\circ\sigma][\circ := e'] &= \circ[{}^\circ\sigma'][{}^\circ\sigma][\circ := e'] \\
&= (\circ[\lambda z \in \text{Dom}({}^\circ\sigma').{}^\circ\sigma'(z)[{}^\circ\sigma]])[\circ := e'] \\
&= e'[\lambda z \in \text{Dom}({}^\circ\sigma').{}^\circ\sigma'(z)[{}^\circ\sigma][\circ := e']] \\
&= e'[\lambda z \in \text{Dom}({}^\circ\sigma').{}^\circ\sigma'(z)[\circ := e'][{}^\circ\sigma[\circ := e']]] \quad \text{by the Induction Hypothesis} \\
&= e'[{}^\circ\sigma'[\circ := e'] \odot {}^\circ\sigma[\circ := e']] \\
&= \circ[{}^\circ\sigma'][\circ := e'][{}^\circ\sigma[\circ := e']] \quad \text{by hygiene assumptions} \\
&= {}^\circ e[\circ := e'][{}^\circ\sigma[\circ := e']]
\end{aligned}$$

□

Next we define analogs of redex and reduction context.

Definition (${}^\circ\mathbb{R}$ ${}^\circ\mathbb{E}_{\text{rdx}}$):

$$\begin{aligned}
{}^\circ\mathbb{R} &= \{\square\} \cup \Theta_{m+n+1}({}^\circ\mathbb{V}^m, {}^\circ\mathbb{R}, {}^\circ\mathbb{E}^n) \\
{}^\circ\mathbb{E}_{\text{rdx}} &= \Theta_n^e({}^\circ\mathbb{V}^n)
\end{aligned}$$

Note that E-reduction contexts possess two types of holes, consequently we must disambiguate the process of hole filling. Note that the unique occurrence of a redex hole is not adorned with a substitution, consequently the process of filling the

redex hole, \square , with the E-expression, ${}^\circ e$, remains unchanged, and we denote it by ${}^\circ R[\square := {}^\circ e]$.

In the case of multiple hole filling we write ${}^\circ R[\circ := e_0][\square := e]$ for the result of filling the expression holes with e_0 , and the redex hole with e .

Lemma (E-properties):

- (1) ${}^\circ R[\circ := e_0][\square := e] = {}^\circ R[\square := e][\circ := e_0]$
- (2) Filling an E-expression, E-reduction context, or E-redex with an expression yields an expression, reduction context, or redex, respectively.

6.2.2 E-Expression Decomposition

We now give a decomposition lemma for E-expressions: An E-expression ${}^\circ e$ is either an E-value (element of ${}^\circ \mathbb{V}$) or it can be decomposed uniquely into an E-reduction context filled with either an E-redex or with an expression hole.

Lemma (E-expression decomposition):

- (0) ${}^\circ e \in {}^\circ \mathbb{V}$, or
- (1) $(\exists ! {}^\circ R, {}^\circ r)({}^\circ e = {}^\circ R[\square := {}^\circ r])$, or
- (2) $(\exists ! {}^\circ R, {}^\circ \sigma)({}^\circ e = {}^\circ R[\square := \circ[{}^\circ \sigma]])$

Proof: An easy induction on the structure of ${}^\circ e$. We consider two example cases. First, suppose ${}^\circ e = \circ[{}^\circ \sigma]$. Then we have case (2) with ${}^\circ R = \square$. Second, suppose ${}^\circ e = \theta({}^\circ e_1, \dots, {}^\circ e_n)$. If ${}^\circ e_i \in {}^\circ \mathbb{V}$ for $1 \leq i \leq n$, then we have case (1) with ${}^\circ R = \square$ (and ${}^\circ r = {}^\circ e$). If ${}^\circ e_i \notin {}^\circ \mathbb{V}$ for some $1 \leq i \leq n$, assume k to be the least such i . Then by the induction hypothesis, ${}^\circ e_k$ decomposes either as (i) ${}^\circ R'[\square := {}^\circ r]$, or as (ii) ${}^\circ R'[\square := \circ[{}^\circ \sigma]]$. Taking ${}^\circ R = \theta({}^\circ e_1, \dots, {}^\circ e_{k-1}, {}^\circ R', {}^\circ e_{k+1}, \dots, {}^\circ e_n)$ we obtain the desired decomposition of ${}^\circ e$. \square

6.2.3 E-Configurations

An E-configuration, ${}^\circ \kappa$, is formed in the same manner as a simple configuration, using E-expressions and E-values instead of simple expressions and values.

Definition (${}^\circ \mathbb{K}$):

$${}^\circ \mathbb{K} = \left\langle \left\langle {}^\circ \mathbb{A}c \mid {}^\circ \mathbb{M} \right\rangle_x \right\rangle_\rho$$

where

$${}^\circ \mathbb{A}c = \mathbb{A}d \xrightarrow{f} {}^\circ \mathbb{A}s$$

$${}^\circ \mathbb{A}s = ({}^\circ \mathbb{V}) \cup [{}^\circ \mathbb{E}] \cup \{({}^\circ \mathbb{A}d)\}$$

$${}^\circ \mathbb{M} = \langle {}^\circ \mathbb{V} \Leftarrow {}^\circ \mathbb{V} \rangle$$

and the constraints specified in the definition of actor configurations in §2 are satisfied.† We let ${}^\circ \kappa$ range over ${}^\circ \mathbb{K}$, and ${}^\circ \alpha$ range over ${}^\circ \mathbb{A}c$. Filling expression holes

† The only condition whose meaning is altered in this general setting is (2), where the

of an E-configuration, E-actor map, E-actor state, E-multiset of messages, and E-messages is defined in the obvious manner. Let ${}^\circ X$ stand generically for an element of one of these E-syntactic classes, then we define ${}^\circ X[\circ := e]$ as follows:

Definition (${}^\circ X[\circ := e]$):

$${}^\circ X[\circ := e] = \begin{cases} \left\langle \left\langle {}^\circ \alpha[\circ := e] \mid {}^\circ \mu[\circ := e] \right\rangle \right\rangle_{\chi}^{\rho} & \text{if } {}^\circ X = \left\langle \left\langle {}^\circ \alpha \mid {}^\circ \mu \right\rangle \right\rangle_{\chi}^{\rho} \\ \lambda x \in \text{Dom}({}^\circ \alpha). {}^\circ \alpha(x)[\circ := e] & \text{if } {}^\circ X = {}^\circ \alpha \\ ((\lambda x. {}^\circ e)[\circ := e]) & \text{if } {}^\circ X = (\lambda x. {}^\circ e) \\ [{}^\circ e[\circ := e]] & \text{if } {}^\circ X = [{}^\circ e] \\ (?_a) & \text{if } {}^\circ X = (?_a) \\ \{ {}^\circ m[\circ := e] \mid {}^\circ m \in {}^\circ \mu \} & \text{if } {}^\circ X = {}^\circ \mu \\ \langle {}^\circ v_0[\circ := e] \Leftarrow {}^\circ v_1[\circ := e] \rangle & \text{if } {}^\circ X = \langle {}^\circ v_0 \Leftarrow {}^\circ v_1 \rangle \end{cases}$$

An E-configuration, ${}^\circ \kappa$, is *closing* for e if ${}^\circ \kappa[\circ := e]$ is a closed configuration. Dually an expression e is a *valid filling* for an E-configuration, ${}^\circ \kappa$, if ${}^\circ \kappa[\circ := e]$ is a configuration. As for atoms and variables, the notion of communicable value remains unchanged and we do not introduce new notation for these. In particular, although messages may have holes, a message with a hole can effectively be ignored. This is because holes in E-values must occur inside λ 's and hence filling these holes cannot yield communicable values or actor addresses. Thus a message with a hole can never be processed. The next lemma expresses the fact that closing E-configurations make just the same observations as simple observing contexts.

Lemma (ocx): $e_0 \cong e_1$ iff $\text{Obs}({}^\circ \kappa[\circ := e_0]) = \text{Obs}({}^\circ \kappa[\circ := e_1])$ for all ${}^\circ \kappa$ that close e_0, e_1 .

Proof: The backward implication is easy to see, since \circledast is (with suitable translation to account for trapping at holes rather than at lambdas) a subset of ${}^\circ \mathbb{K}$. The idea for the proof of the forward implication is to define for each configuration context ${}^\circ \kappa$, an observing context O whose computations give rise to the same set of observations. In fact O evolves to ${}^\circ \kappa$ in a finite number of steps. For an E-expression ${}^\circ e$ we define ${}^\circ e^*$ to be the result of recursively replacing decorated holes $\circ[\sigma]$ by applications $\text{app}(\dots \text{app}(\lambda x_1 \dots \lambda x_n. \bullet, {}^\circ \sigma(x_1)^*), \dots, {}^\circ \sigma(x_n)^*)$ where $\{x_1, \dots, x_n\} = \text{Dom}({}^\circ \sigma)$. Let ${}^\circ \kappa = \left\langle \left\langle {}^\circ \alpha \mid {}^\circ \mu \right\rangle \right\rangle$, let $A = [a_i \mid i < n] = \text{Dom}({}^\circ \alpha)$, and define $O = \left\langle \left\langle [e_{\circ \kappa}]_{\hat{a}} \mid \emptyset \right\rangle \right\rangle$ where $\hat{a} \notin A$, and $e_{\circ \kappa}$ is constructed as follows. Let

$$E = \{i < n \mid (\exists e_i)({}^\circ \alpha(a_i) = [{}^\circ e_i])\}, \text{ and let } n_E \text{ be the cardinality of } E.$$

$$I = \{i < n \mid (\exists v_i)({}^\circ \alpha(a_i) = ({}^\circ v_i))\}.$$

$$B_i(a_0, \dots, a_{n-1}) = {}^\circ v_i^*, \text{ if } {}^\circ \alpha(a_i) = ({}^\circ v_i).$$

$$B_i(a_0, \dots, a_{n-1}) = \lambda a. \text{seq}(\text{send}(a, 0), {}^\circ e_i^*), \text{ if } {}^\circ \alpha(a_i) = [{}^\circ e_i].$$

$$\mu = \{ \langle z_j \Leftarrow {}^\circ v_j' \rangle \mid j < m \}$$

Define

$$W_{\circ \kappa} = \text{rec}(\lambda b. \lambda k. \lambda m. \text{if}(\text{eq}(k, 0), \text{seq}(\text{send}(z_j, ({}^\circ v_j')^*)_{j < m}), \text{become}(b(k-1))))$$

free variables of any hole occurrences (namely the free variables in the range of the associated substitution) must be taken into consideration.

$$\begin{aligned}
e_{\circ\kappa} = & \text{let}\{a_i := \text{newadr}()\}_{i < n} \\
& \text{seq}(\text{initbeh}(a_i, B_i(a_0, \dots, a_{n-1}))_{i \in I \cup E}, \\
& \quad \text{send}(a_i, \hat{a})_{i \in E}, \\
& \quad \text{become}(W_{\circ\kappa}(n_E)))
\end{aligned}$$

Now, we claim that any computation of ${}^{\circ}\kappa[o := e]$ has a corresponding computation (with same observations) of $O[e]$ obtained by accepting all the startup messages, sending and accepting the acks, and completing the computation of the initializing actor (which can then be ignored). Conversely any computation of $O[e]$ has a corresponding computation of ${}^{\circ}\kappa[o := e]$ obtained by ignoring the finite amount of initializing activity. A more detailed proof can be given along the lines of the proof of the theorem (**fun-red-eq**) below. \square

6.2.4 E-Reduction

The reduction relations \mapsto^{λ} and \mapsto are extended to the generalized domains in the obvious fashion, simply by liberally annotating metavariables with ${}^{\circ}$'s, modulo the extension of substitution to E-expressions. As examples, we give the (**beta-v**), (**br**), and (**eq**) clauses of \mapsto^{λ} and the internal transitions for \mapsto on closed E-configurations.

Definition (\mapsto^{λ}):

$$\begin{aligned}
(\text{beta-v}) \quad & {}^{\circ}R[\square := \text{app}(\lambda x. {}^{\circ}e, {}^{\circ}v)] \mapsto^{\lambda} {}^{\circ}R[\square := {}^{\circ}e[x := {}^{\circ}v]] \\
(\text{br}) \quad & {}^{\circ}R[\square := \text{br}({}^{\circ}v_1, {}^{\circ}v_2, {}^{\circ}v)] \mapsto^{\lambda} \begin{cases} {}^{\circ}R[\square := {}^{\circ}v_1] & \text{if } {}^{\circ}v \in {}^{\circ}\mathbb{V} - \{\text{nil}\} \\ {}^{\circ}R[\square := {}^{\circ}v_2] & \text{if } {}^{\circ}v = \text{nil} \end{cases} \\
(\text{eq}) \quad & {}^{\circ}R[\square := \text{eq}({}^{\circ}v_0, {}^{\circ}v_1)] \mapsto^{\lambda} \begin{cases} {}^{\circ}R[\square := \text{t}] & \text{if } {}^{\circ}v_0 = {}^{\circ}v_1 \in \text{At} \\ {}^{\circ}R[\square := \text{nil}] & \text{if } {}^{\circ}v_0, {}^{\circ}v_1 \in \text{At} \text{ and } {}^{\circ}v_0 \neq {}^{\circ}v_1 \end{cases}
\end{aligned}$$

Definition (\mapsto):

$$\begin{aligned}
\langle \text{fun} : a \rangle \quad & {}^{\circ}e \mapsto^{\lambda} {}^{\circ}e' \Rightarrow \langle\langle {}^{\circ}\alpha, [{}^{\circ}e]_a \mid {}^{\circ}\mu \rangle\rangle \mapsto \langle\langle {}^{\circ}\alpha, [{}^{\circ}e']_a \mid {}^{\circ}\mu \rangle\rangle \\
\langle \text{new} : a, a' \rangle \quad & \langle\langle {}^{\circ}\alpha, [{}^{\circ}R[\square := \text{newadr}()]]_a \mid {}^{\circ}\mu \rangle\rangle \mapsto \\
& \langle\langle {}^{\circ}\alpha, [{}^{\circ}R[\square := a']]_a, (?_a)_{a'} \mid {}^{\circ}\mu \rangle\rangle \quad a' \text{ fresh} \\
\langle \text{init} : a, a' \rangle \quad & \langle\langle {}^{\circ}\alpha, [{}^{\circ}R[\square := \text{initbeh}(a', {}^{\circ}v)]]_a, (?_a)_{a'} \mid {}^{\circ}\mu \rangle\rangle \mapsto \\
& \langle\langle {}^{\circ}\alpha, [{}^{\circ}R[\square := \text{nil}]]_a, ({}^{\circ}v)_{a'} \mid {}^{\circ}\mu \rangle\rangle \\
\langle \text{bec} : a, a' \rangle \quad & \langle\langle {}^{\circ}\alpha, [{}^{\circ}R[\square := \text{become}({}^{\circ}v)]]_a \mid {}^{\circ}\mu \rangle\rangle \mapsto \\
& \langle\langle {}^{\circ}\alpha, [{}^{\circ}R[\square := \text{nil}]]_{a'}, ({}^{\circ}v)_a \mid {}^{\circ}\mu \rangle\rangle \quad a' \text{ fresh} \\
\langle \text{send} : a, m \rangle \quad & \langle\langle {}^{\circ}\alpha, [{}^{\circ}R[\square := \text{send}({}^{\circ}v_0, {}^{\circ}v_1)]]_a \mid {}^{\circ}\mu \rangle\rangle \mapsto
\end{aligned}$$

$$\langle \text{rcv} : a, cv \rangle \quad \left\langle \circ\alpha, (\circ v)_a \mid \langle a \Leftarrow cv \rangle, \circ\mu \right\rangle \mapsto \left\langle \circ\alpha, [\text{app}(\circ v, cv)]_a \mid \circ\mu \right\rangle$$

6.2.5 E-Uniform Computation

The notion of E-uniform computation is made precise in the following definitions and lemmas. The basic idea is that given a decomposition of a configuration as an E-configuration with holes filled by a given expression, then any transition step leading from that configuration is either independent of what appears in the holes, or it explicitly uses information about the contents of some hole occurrence.

Definition (E-hole touching): Let $\circ\kappa = \left\langle \circ\alpha \mid \circ\mu \right\rangle$. We say that $\circ\kappa$ touches a hole at a if $\circ\alpha(a) = [\circ R[\square := \circ[\circ\sigma]]]$ for some $\circ R, \circ\sigma$.

We say that a transition $\kappa \xrightarrow{l} \kappa'$ touches a hole relative to a decomposition $\kappa = \circ\kappa[\circ := e]$ if l has focus a and $\circ\kappa$ touches a hole at a .

Lemma (E-Uniform Computation):

- (1) If $\circ\kappa \xrightarrow{l} \circ\kappa'$, then $\circ\kappa[\circ := e] \xrightarrow{l} \circ\kappa'[\circ := e]$ for any valid filling expression e .
- (2) If $\circ\kappa$ has no transition with focus a (and a is an actor of $\circ\kappa$), then either $\circ\kappa$ touches a hole at a or $\circ\kappa[\circ := e]$ has no transition with focus a for any valid filling expression e .
- (3) If $\kappa \xrightarrow{l} \kappa'$ and $\kappa = \circ\kappa[\circ := e]$, then either the transition touches a hole or we can find $\circ\kappa'$ such that $\kappa' = \circ\kappa'[\circ := e]$ and $\circ\kappa \xrightarrow{l} \circ\kappa'$.

Proof (1): This is proved by considering cases on the transition rule applied. The only interesting case is (**beta-v**). This follows from (**fil-subst**) \square_1

Proof (2): Assume $\circ\kappa = \left\langle \circ\alpha \mid \circ\mu \right\rangle$ has no transition with focus a , and $\circ\kappa$ does not touch a hole at a . Then one of the following holds:

- (i) $\circ\alpha(a) = (?_{a'})$
- (ii) $\circ\alpha(a) = (\circ v)$ and $\circ\mu$ contains no messages deliverable to a
- (iii) $\circ\alpha(a) = [\circ v]$
- (iv) $\circ\alpha(a) = [\circ R[\square := \text{initbeh}(\circ v_0, \circ v_1)]]$ where $\circ v_0$ is not the address of an uninitialized actor created by a
- (v) $\circ\alpha(a) = [\circ R[\square := \circ r]]$ where $\circ r$ is a non-actor redex that is stuck.

In each of these cases, it easy to see that there will be no transition with focus a enabled when the expressions holes are filled. \square_2

Proof (3): Assume $\kappa = \left\langle \alpha \mid \mu \right\rangle \xrightarrow{l} \kappa' = \left\langle \alpha' \mid \mu' \right\rangle$, $\kappa = \circ\kappa[\circ := e]$, and the transition does not touch the hole. Thus $\circ\kappa = \left\langle \circ\alpha \mid \circ\mu \right\rangle$ where $\alpha = \circ\alpha[\circ := e]$ and $\mu = \circ\mu[\circ := e]$. We want to find $\circ\alpha', \circ\mu'$ such that $\circ\kappa \xrightarrow{l} \circ\kappa' = \left\langle \circ\alpha' \mid \circ\mu' \right\rangle$, $\alpha' = \circ\alpha'[\circ := e]$, and $\mu' = \circ\mu'[\circ := e]$. Since we are considering closed configurations there are no i/o transitions. Thus, we need to consider only two cases rcv transitions and exec transitions. We split the exec transitions into functional and actor primitives.

Receive: $l = \langle \text{rcv} : a, cv \rangle$, $\langle a \Leftarrow cv \rangle \in \mu$, and $\alpha(a) = (v)$. Thus $\circ\alpha(a) = (\circ v)$

with $v = {}^\circ v[\circ := e]$. Thus we let ${}^\circ\alpha' = {}^\circ\alpha\{a := [\mathbf{app}({}^\circ v, cv)]\}$, and ${}^\circ\mu = {}^\circ\mu' \cup \{<a \leftarrow cv>\}$.

Execution-lambda: $l = \langle \mathbf{fun} : a \rangle$, $\alpha(a) = [R[\square := r]]$ and $r \xrightarrow{\lambda} e'$. Thus ${}^\circ\alpha(a) = [{}^\circ R[\square := {}^\circ r]]$ with $R = {}^\circ R[\circ := e]$, and $r = {}^\circ r[\circ := e]$. Thus we want to find ${}^\circ e'$ such that ${}^\circ r \xrightarrow{\lambda} {}^\circ e'$. Then ${}^\circ\alpha' = {}^\circ\alpha\{a := [{}^\circ R[\square := {}^\circ e']]\}$ and ${}^\circ\mu' = {}^\circ\mu$. If $r = \mathbf{app}(\lambda z.e_0, v)$ (z chosen fresh), then $e' = e_0[z := v]$, and ${}^\circ r = \mathbf{app}(\lambda z.{}^\circ e_0, {}^\circ v)$ where $e_0 = {}^\circ e_0[\circ := e]$ and $v = {}^\circ v[\circ := e]$. Take ${}^\circ e' = {}^\circ e_0[z := {}^\circ v]$ and use (**fil-subst**). If $r = \mathbf{eq}(v_0, v_1)$, then ${}^\circ r = \mathbf{eq}({}^\circ v_0, {}^\circ v_1)$ where $v_j = {}^\circ v_j[\circ := e]$ for $j < 2$. e' is **t** or **nil** and we may take ${}^\circ e' = e'$. The remaining cases are similar.

Execution-actor: If $l = \langle \mathbf{send} : a \rangle$, then $\alpha(a) = [R[\square := \mathbf{send}(v_0, v_1)]]$ $\alpha'(a) = [R[\square := \mathbf{nil}]]$, $\mu' = \mu \cup \{<v_0 \leftarrow v_1>\}$. Also ${}^\circ\alpha(a) = [{}^\circ R[\square := \mathbf{send}({}^\circ v_0, {}^\circ v_1)]]$ where $R = {}^\circ R[\circ := e]$, and $v_j = {}^\circ v_j[\circ := e]$ for $j < 2$. Take ${}^\circ\alpha' = {}^\circ\alpha\{a := [{}^\circ R[\square := \mathbf{nil}]]\}$ and ${}^\circ\mu' = {}^\circ\mu \cup \{<{}^\circ v_0 \leftarrow {}^\circ v_1>\}$.

If $l = \langle \mathbf{become} : a, a' \rangle$, then a' is fresh, $\alpha(a) = [R[\square := \mathbf{become}(v)]]$ $\alpha'(a) = (v)$, $\alpha'(a') = [R[\square := \mathbf{nil}]]$, and $\mu' = \mu$. Also ${}^\circ\alpha(a) = [{}^\circ R[\square := \mathbf{become}({}^\circ v)]]$ where $R = {}^\circ R[\circ := e]$, and $v = {}^\circ v[\circ := e]$. Take ${}^\circ\alpha' = {}^\circ\alpha\{a := ({}^\circ v), a' := [{}^\circ R[\square := \mathbf{nil}]]\}$ and ${}^\circ\mu' = {}^\circ\mu$.

If $l = \langle \mathbf{new} : a, a' \rangle$, then a' is fresh, $\alpha(a) = [R[\square := \mathbf{newadr}()]]$ $\alpha'(a) = [R[\square := \mathbf{nil}]]$, $\alpha'(a') = (?_a)$, and $\mu' = \mu$. Also ${}^\circ\alpha(a) = [{}^\circ R[\square := \mathbf{newadr}()]]$ where $R = {}^\circ R[\circ := e]$. Take ${}^\circ\alpha' = {}^\circ\alpha\{a := [{}^\circ R[\square := \mathbf{nil}]]\}$, $a' := (?_a)\}$ and ${}^\circ\mu' = {}^\circ\mu$.

If $l = \langle \mathbf{init} : a, a' \rangle$, then $\alpha(a) = [R[\square := \mathbf{initbeh}(a', v)]]$, $\alpha(a') = (?_a)$, $\alpha'(a) = [R[\square := \mathbf{nil}]]$, $\alpha'(a') = (v)$, and $\mu' = \mu$. Also ${}^\circ\alpha(a) = [{}^\circ R[\square := \mathbf{initbeh}(a', {}^\circ v)]]$, where $R = {}^\circ R[\circ := e]$, and $v = {}^\circ v[\circ := e]$, and ${}^\circ\alpha(a') = (?_a)_{..}$. Take ${}^\circ\alpha' = {}^\circ\alpha\{a := [{}^\circ R[\square := \mathbf{nil}]]\}$, $a' := ({}^\circ v)\}$ and ${}^\circ\mu' = {}^\circ\mu$.

□₃ □**E**-uniform

6.2.6 The Common Expression Reduct Theorem

Now we have enough notation and tools to describe the construction of path correspondences for expressions with uniform common reducts. We first consider the case of expressions that reduce via purely functional reductions. Then we show how this construction can be modified to allow for reduction of actor primitives.

Theorem (fun-red-eq): If for each ${}^\circ\sigma$ with domain containing the free variables of e_0, e_1 , either $e_j[{}^\circ\sigma]$ hangs for $j < 2$, or there is some ${}^\circ e_c$ such that $e_j[{}^\circ\sigma]$ reduces in 0 or more $\xrightarrow{\lambda}$ steps to ${}^\circ e_c$ uniformly, then $e_0 \cong e_1$.

Corollary (fun-red-eq): The following laws are instances of (**fun-red-eq**): (**red-exp**), (**beta-v**), (**ift**), (**ifn**), (**ifelim**), (**isprt**), (**isprn**), (**fst**), and (**snd**).

Proof: Assume that for each closing ${}^\circ\sigma$ there is ${}^\circ e_{c,j}$ such that $e_j[{}^\circ\sigma] \xrightarrow{\lambda} \dots \xrightarrow{\lambda} {}^\circ e_{c,j}$ $j < 2$, uniformly, and either ${}^\circ e_{c,j}$ is (uniformly) stuck for $j < 2$, or ${}^\circ e_{c,0} = {}^\circ e_{c,1}$. In either case we call ${}^\circ e_{c,j}$ the common reduct. We want to show that $e_0 \cong e_1$. By (**ocx**) it is sufficient to show that $Obs({}^\circ\kappa[\circ := e_0]) = Obs({}^\circ\kappa[\circ := e_1])$ for any ${}^\circ\kappa$ that is a closing E-configuration for e_0 and e_1 . To do this, we show that for any $\pi_0 \in \mathcal{F}({}^\circ\kappa[\circ := e_0]) = [\kappa_i \xrightarrow{l_i} \kappa_{i+1} \mid i \in \mathbb{N}]$ we can find $\pi_1 \in \mathcal{F}({}^\circ\kappa[\circ := e_1])$ such that $obs(\pi_0) = obs(\pi_1)$. (The case with 0 and 1 interchanged is symmetric.)

Informally, by the uniformity property of computations, we see that replacing occurrences of e_0 by e_1 has no effect on a computation except where a hole is touched. Using (**infinite macro-steps**) we can localize non-uniform steps so that when a hole is touched, reduction to a common reduct occurs in a single multi-step (which involves no event transitions). Thus we may obtain a computation for ${}^{\circ}\kappa[\circ := e_1]$ by replacing occurrences of e_0 by e_1 and replacing multi-step transitions reducing $e_0[\circ\sigma]$ to its common reduct by multi-step transitions reducing $e_1[\circ\sigma]$ to its common reduct. To insure completeness/fairness of the result, we need to take account of the case where a hole $\circ[\circ\sigma]$ is exposed, but the multi-step for $e_0[\circ\sigma]$ is trivial and hence does not appear as a transition. We do this by inserting the corresponding multi-step for $e_1[\circ\sigma]$ at the point where the hole is first exposed. Such holes then effectively disappear, since they are either filled with a stuck expression or with the same expression. Now we make this informal argument more rigorous, by the following steps:

- (1) We analyze the configurations occurring in π_0 and record occurrences of e_0 in holes descending from ${}^{\circ}\kappa$. This gives us decompositions ${}^{\circ}\kappa_i[\circ := e_0]$ of κ_i . In the cases where a hole is touched such that e_0 is its common reduct, we fill that hole with e_0 giving a new E-configuration ${}^{\circ}\kappa'_i$ with one less hole, such that ${}^{\circ}\kappa'_i[\circ := e_0]$ is κ_i . This process of filling holes with common reducts continues until the transition l_i is either uniform or touches a hole in which e_0 is not its common reduct. We also record subsequences of transitions corresponding to uniform reduction of such occurrences of e_0 to its common reduct.
- (2) Using (**infinite-macro-steps**) we may assume that the path is expressed in terms of multi-step transitions such that the recorded subsequences of transitions corresponding to non-trivial reduction to a common reduct are single multi-steps. We also insert copies of κ_i for each hole that is filled with a common reduct, remembering the corresponding decomposition, and insert empty multi-steps between these copies. We also insert a copy of κ_i and a connecting empty multi-step for each hole that occurs in a reduction context that is not touched – because the occurrence of e_0 is stuck, or because it is a value and placing it in the redex hole produces either a value or a stuck state.
- (3) Form π_1 by filling the holes of ${}^{\circ}\kappa_i$ with e_1 and replacing multi-steps for e_0 by corresponding multi-steps for e_1 . Note that empty multi-steps may expand to non-trivial reductions of occurrences of e_1 to its common reduct.

It is easy to see that π_1 is a computation path. The argument that it is complete and fair relies on the insertion of multi-steps, and uses the same case analysis that was used in the uniform computation lemma.

Step (1) We analyze and decompose π_0 to obtain

- (i) for each $i < \infty$, an integer n_i and a sequence of decompositions ${}^{\circ}\kappa_{i,j}$ for $j \leq n_i$ such that $\kappa_i = {}^{\circ}\kappa_{i,j}[\circ := e_0]$ for $j \leq n_i$ and such that ${}^{\circ}\kappa_{i,n_i} \xrightarrow{l_i} {}^{\circ}\kappa_{i+1,0}$ uniformly, or the transition touches a hole in which e_0 has non-trivial reduction to its common reduct. (We call this entering the hole.) n_i will be 0 except in the case of a hole touched in which e_0 is its own common reduct. Then we fill that hole with the common reduct and re-decompose.

- (ii) The set I of indices of transitions that enter holes
- (iii) The map J from I to the sequence of indices of transitions corresponding to the thread of computation that carries out the reduction to the common reduct.

This is done incrementally by defining sequences I_n, J_n by induction on n and taking $I = \bigcup_{n < \infty} I_n$, and $J = \bigcup_{n < \infty} J_n$. At stage (i, j) we have defined I_i, J_i , and ${}^\circ\kappa_{i,j}$. If $j = n_i$, then the next stage is $(i + 1, 0)$ otherwise it is $(i, j + 1)$.

Stage $(0, 0)$ $I_0 = \emptyset$, and J_0 is the empty map. ${}^\circ\kappa_{0,0} = {}^\circ\kappa$.

At stage (i, j) There are four cases to consider:

- (1) l_i is execution in a hole;
- (2) l_i is uniform wrt ${}^\circ\kappa_{i,j}$
- (3) l_i touches a hole and
 - (3.1) enters the hole
 - (3.2) does not enter the hole

Case 1: This case occurs if i is an element of $M = J_i(m)$ for some $m \in I_i$. Thus $n_i = j$ and $l_i = \langle \text{fun} : a \rangle$ for some a . We move to stage $(i + 1, 0)$ with $I_{i+1} = I_i$, $J_{i+1} = J_i$, ${}^\circ\mu_{i+1,0} = {}^\circ\mu_{i,j}$, and ${}^\circ\alpha_{i+1,0} = {}^\circ\alpha_{i,j}\{a := [{}^\circ R[\square := {}^\circ e_{m,k+1}]]\}$ where ${}^\circ R$, and ${}^\circ e_{m,k+1}$ are obtained as follows. Let k be the index of i in M . The hole is entered at stage (m, n_m) with ${}^\circ\alpha_{m,n_m}(a) = [{}^\circ R[\square := \circ[{}^\circ\sigma]]]$. Let ${}^\circ e_{m,0} = e_0[{}^\circ\sigma]$, let n be the length of M , and let $[{}^\circ e_{m,k} \xrightarrow{\lambda} {}^\circ e_{m,k+1} \mid k < n]$ be the thread of computation reducing $e_0[{}^\circ\sigma]$ to its common reduct ${}^\circ e_{m,n}$. Note that ${}^\circ\alpha_{i,j}(a) = [{}^\circ R[\square := {}^\circ e_{m,k}]]$.

Case 2: $n_i = j$ and we move to stage $(i + 1, 0)$ with $I_{i+1} = I_i$, $J_{i+1} = J_i$, ${}^\circ\kappa_{i+1,0}$ such that ${}^\circ\kappa_{i,j} \xrightarrow{l_i} {}^\circ\kappa_{i+1,0}$ uniformly according to the uniformity lemma.

Case 3.1: In this case $l_i = \langle \text{fun} : a \rangle$ for some a . $n_i = j$ and we move to stage $(i + 1, 0)$ with $I_{i+1} = I_i \cup \{i\}$, $J_{i+1} = J_i\{i := M\}$, ${}^\circ\mu_{i+1} = {}^\circ\mu_i$, and ${}^\circ\alpha_{i+1} = {}^\circ\alpha_i\{a := [{}^\circ R[\square := {}^\circ e_{i,1}]]\}$ where M , ${}^\circ R$, and ${}^\circ e_{i,1}$ are obtained as follows. Let ${}^\circ\alpha_{i,j}(a) = [{}^\circ R[\square := \circ[{}^\circ\sigma]]]$, let ${}^\circ e_{i,0} = e_0[{}^\circ\sigma]$, and let $[{}^\circ e_{i,k}[{}^\circ\sigma] \xrightarrow{\lambda} {}^\circ e_{i,k+1} \mid k < n + 1]$ be the thread of computation reducing $e_0[{}^\circ\sigma]$ to its common reduct ${}^\circ e_{i,n+1}$. By fairness, there is a sequence of indices $M = [i_k \mid k < n + 1]$ with $i_0 = i$, $i_k < i_{k+1}$ for $k < n + 1$ such that M is the macro step corresponding to the above lambda reduction.

Case 3.2: We move to stage $(i, j + 1)$ with $I_{i+1} = I_i$, $J_{i+1} = J_i$, ${}^\circ\mu_{i,j+1} = {}^\circ\mu_{i,j}$, and ${}^\circ\alpha_{i,j+1} = {}^\circ\alpha_{i,j}\{a := [{}^\circ R[\square := e_0[{}^\circ\sigma]]]\}$ where a , ${}^\circ R$, ${}^\circ\sigma$ are obtained as follows. a is the focus of l_i , and ${}^\circ\alpha_{i,j}(a) = [{}^\circ R[\square := \circ[{}^\circ\sigma]]]$, with $e_0[{}^\circ\sigma]$ equal to its common reduct.

Step (2) The family $J(i)$ for $i \in I$ satisfies the conditions of **(infinite-macro-step)**. Hence we may assume that π_0 has the form

$$[{}^\circ\kappa_{i,0}[\circ := e_0] \xrightarrow{[]} \dots \xrightarrow{[]} {}^\circ\kappa_{i,n_i} \xrightarrow{L_i} {}^\circ\kappa_{i+1,0}[\circ := e_0]] \mid i < \infty$$

where each L_i is either a single (uniform) transition, or a multi-step reduction of an occurrence of e_0 to its common reduct, and the E-configurations obtained by the above decomposition method. For each i , if ${}^\circ\alpha_{i,0}(a) = [{}^\circ R[\square := \circ[{}^\circ\sigma]]]$, and there are no transitions L_j for $i \leq j$ with focus a , and i is the least such index, we insert before each transition leading from ${}^\circ\alpha_{i,0}$ an empty transition

$${}^\circ\kappa_{i,0}[\circ := e_0] \xrightarrow{[]} {}^\circ\kappa_{i,0}[\circ := e_0].$$

Step (3) We let π_1 be the path

$$[[\circ\kappa_{i,0}[\circ := e_1] \xrightarrow{L_{i,0}} \dots \xrightarrow{L_{i,n_i-1}} \circ\kappa_{i,n_i}[\circ := e_1] \xrightarrow{L'_i} \circ\kappa_{i+1,0}[\circ := e_1]] \mid i < \infty]$$

where L'_i is L_i if L_i is a single (uniform) transition; L'_i is the corresponding macro-step reduction of the occurrence of e_1 to its common reduct, if L_i is a macro-step or an empty transition.

Clearly π_1 is a complete computation path. Also the transitions are the same except for points where holes are touched, but these differences are not observable. Thus $obs(\pi_0) = obs(\pi_1)$.

It remains to check that fairness has been preserved. Suppose some transition l is enabled at stage i in π_1 . We have three cases:

Receive: Suppose l is receipt of $\langle a \Leftarrow cv \rangle$. Then $\circ\alpha_i(a) = (\lambda x. \circ e)$ and hence l is enabled in π_0 at stage i . If l fires in π_0 at stage $i' \geq i$, then it also fires at this stage in π_1 . Suppose l never fires in π_0 . Then by fairness, there is some $i' > i$ such that $\circ\alpha_{j'}(a)$ is an executing state for $j' \geq i'$. By construction l is permanently disabled at i' in π_1 as well.

Uniform Execution: Suppose l is an execution step with focus a where $\circ\alpha_i(a) = [\circ R[\square := \circ r]]$. Then l is enabled in π_0 at i , it can not be disabled, and must occur in π_0 at some stage $i' \geq i$ and hence will occur in π_1 at that stage.

Hole Touching: Suppose l is an execution step by a with $\circ\alpha_i(a) = [\circ e]$ where $\circ e = \circ R[\square := \circ[\circ\sigma]]$. First assume $e_1[\circ\sigma]$ reduces. If $\circ e[\circ := e_0]$ does not reduce, then by construction, the transition is taken in π_1 as soon as it is enabled. If $\circ e[\circ := e_0]$ reduces, then a transition will eventually be taken at a in π_0 , and the l will be taken at the corresponding point in π_1 . Suppose $e_1[\circ\sigma]$ does not reduce. Then it must be a value, hence the common reduct. Hence the reduction of e_0 is enabled in π_0 and will eventually be taken. $\circ R$ has the form $\circ R_0[\square := \theta(\circ v^m, \circ[\circ\sigma], \circ e^n)]$. If all the E-expressions $\circ e^n$ are E-values, then l must be reduction of the redex in $\circ R_0$ and this is also enabled now, in π_0 . Otherwise consider decomposition of the first non-value element of $\circ e^n$ and repeat this argument. Since we are now looking at a smaller E-expression, we eventually reach the point where the step enabled in π_1 corresponds to one in π_0 and hence will occur eventually. \square

6.2.7 The Proof of the Equivalence of Hanging and Lambda-Divergence

We now prove (**hang-infin**), which says that any two expressions that hang or have infinite computations are observationally equivalent.

Proof (hang-infin): Assume $e_0, e_1 \in \text{Hang} \cup \text{Infin}$. We want to show that $e_0 \cong e_1$. Let $\circ\kappa = \langle \langle \circ\alpha \mid \circ\mu \rangle \rangle$ be a closing E-configuration for e_0 and e_1 . Assume $\pi_0 \in \mathcal{F}(\circ\kappa[\circ := e_0]) = [\kappa_i \xrightarrow{L_i} \kappa_{i+1} \mid i \in \infty]$. We want to find $\circ\alpha_i$, and L_i , such that $\kappa_i = \circ\kappa_i[\circ := e_0]$ where $\circ\kappa_i = \langle \langle \circ\alpha_i \mid \circ\mu_i \rangle \rangle$ and, letting $\pi_1 = [\circ\kappa_i[\circ := e_1] \xrightarrow{L_i} \circ\kappa_{i+1}[\circ := e_1] \mid i \in \infty]$, we have $\pi_1 \in \mathcal{F}(\circ\kappa[\circ := e_1])$ and $obs(\pi_0) = obs(\pi_1)$. (Actually, we let holes in π_j be filled by any expression of the same class as e_j .) For the base case we have $\circ\alpha_0 = \circ\alpha$. Assume we have $\circ\alpha_i$. Suppose $e_0 \in \text{Hang}$. Let a be

the focus of l_i . We first consider each a' other than a such that ${}^\circ\alpha_i(a') = {}^\circ R[\square := \circ[\circ\sigma]]$. If $e_1 \in \text{Hang}$ then we just insert any steps needed to reach the stuck state (we assume that they are already macroized for π_0). If $e_1 \in \text{Infin}$, then insert the step to reach the next element of its infinite sequence. Now we consider the transition label l_i . If it does not touch a hole, then ${}^\circ\alpha_{i+1}$ is given by the uniform transition lemma. Suppose ${}^\circ\alpha_i(a) = {}^\circ R[\square := \circ[\circ\sigma]]$. Then ${}^\circ\alpha_{i+1}$ has the same decomposition, just possibly different expressions (of the same class) filling the holes. $\square_{\mathbf{h-i}}$

6.2.8 The Proofs of the Actor Primitive Laws

We show how to modify the construction for the purely functional case to establish equivalence where reductions may involve actor primitives. The notion of common reduct is a fragment of an E-configuration which must be merged with the parent configuration. In the equivalences considered below, with one exception, no extra steps need to be inserted since in every instance either both expressions step, or both hang. Thus we only need to construct ${}^\circ\alpha_{i+1}$ for the hole touching case. Fairness follows easily using the same argument as for the functional case. The exception is the commuting of **newadr** with **initbeh**, which will be treated in more detail below.

Proof (triv): Suppose e_0, e_1 are taken from the pair $\vartheta(\bar{x}), \text{seq}(\vartheta(\bar{x}), \text{nil})$ where $\vartheta \in \{\text{send}, \text{become}, \text{initbeh}\}$. Suppose l_i is an execution by a with ${}^\circ\alpha_i(a) = [{}^\circ R[\square := \circ[\circ\sigma]]]$. We define ${}^\circ\alpha_{i+1}$ for each form $\vartheta(\bar{x})$.

(**send**(x_0, x_1)) ${}^\circ\alpha_{i+1} = {}^\circ\alpha_i\{a := [{}^\circ R[\square := \text{nil}]]\}$. Also, ${}^\circ\mu_{i+1} = {}^\circ\mu_i \cup \{\langle v_0 \Leftarrow v_1 \rangle\}$, where ${}^\circ v_j = x_j[\circ\sigma]$. (Recall that any holes appearing in messages must be inside lambdas and hence the message is ill-formed and can be ignored.)

(**become**(x_0)) ${}^\circ\alpha_{i+1} = {}^\circ\alpha_i\{a := ({}^\circ\sigma(x_0)), a' := [{}^\circ R[\square := \text{nil}]]\}$ where a' is a new (anonymous) actor identifier.

(**initbeh**(x_0, x_1)) ${}^\circ\alpha_{i+1} = {}^\circ\alpha_i\{a := [{}^\circ R[\square := \text{nil}]], {}^\circ\sigma(x_0) := ({}^\circ\sigma(x_1))\}$. Since the step occurs in π_0 , we may assume $a' = {}^\circ\sigma(x_0) \in \text{Dom}({}^\circ\alpha_i)$, and ${}^\circ\alpha_i(a') = (?_a)$.

$\square_{\mathbf{triv}}$

Proof (can-b): Suppose e_0, e_1 are taken from the pair

$$\text{seq}(\text{become}(v_0), \text{become}(v_1)), \quad \text{seq}(\text{become}(v_0), \text{nil}).$$

Assume l_i is an execution by a with ${}^\circ\alpha_i(a) = [{}^\circ R[\square := \circ[\circ\sigma]]]$. Define

$${}^\circ\alpha_{i+1} = {}^\circ\kappa_i\{a := (v_0[\circ\sigma]), a' := [{}^\circ R[\square := \text{nil}]]\}$$

where a' is a new anonymous actor identifier. We ignore the additional anonymous actor with state $(v_1[\circ\sigma])$ in the case of two becomes, since it is carried along unchanged in the remainder of the computation. \square

Proof (can-i): Suppose e_0, e_1 are taken from the pair

$$\text{seq}(\text{initbeh}(v, v_0), \text{initbeh}(v, v_1)), \quad \text{seq}(\text{initbeh}(v, v_0), \text{stuck}).$$

(The equivalence with **stuck** replaced by **bot** follows from (**hang-infin**).) Assume l_i is an execution by a with ${}^\circ\alpha_i(a) = [{}^\circ R[\square := \circ[\circ\sigma]]]$. Define

$${}^\circ\alpha_{i+1} = {}^\circ\alpha_i\{a := [{}^\circ R[\square := \text{bot}]], {}^\circ\sigma(v) := ({}^\circ\sigma(v_0))\}$$

Since the step occurs in π_0 , we may assume $a' = {}^\circ\sigma(v) \in \text{Dom}({}^\circ\alpha_i)$, and ${}^\circ\alpha_i(a') = (?_a)$. Officially, in π_1 ${}^\circ\alpha_{i+1}(a)$ should be $[{}^\circ R[\square := \text{initbeh}(v, v_1)[{}^\circ\sigma]]]$. But, as in the proof of (**hang-infin**) we treat stuck expressions as indistinguishable. \square

Proof (commutes): Suppose e_0, e_1 are taken from one of the pairs in the commutes lemma, except (**n-i**). Assume l_i is an execution by a with ${}^\circ\alpha_i(a) = [{}^\circ R[\square := {}^\circ\sigma]]]$.

(s-s)

$$\text{let}\{x_0 := \text{send}(v_0, v_1)\}\text{let}\{x_1 := \text{send}(v_2, v_3)\}e$$

$$\text{let}\{x_1 := \text{send}(v_2, v_3)\}\text{let}\{x_0 := \text{send}(v_0, v_1)\}e$$

Define

$${}^\circ\alpha_{i+1} = {}^\circ\alpha_i\{a := [{}^\circ R[\square := e\{x_0 := \text{nil}, x_1 := \text{nil}\}[{}^\circ\sigma]]]\}$$

$${}^\circ v_j = v_j[{}^\circ\sigma] \quad \text{for } j < 4$$

$${}^\circ\mu_{i+1} = {}^\circ\mu_i \cup \{\langle {}^\circ v_0 \Leftarrow {}^\circ v_1 \rangle, \langle {}^\circ v_2 \Leftarrow {}^\circ v_3 \rangle\}$$

$\square_{\text{s-s}}$

(s-n)

$$\text{let}\{x_0 := \text{send}(v_0, v_1)\}\text{let}\{x_1 := \text{newadr()}\}e$$

$$\text{let}\{x_1 := \text{newadr()}\}\text{let}\{x_0 := \text{send}(v_0, v_1)\}e$$

Let a_1 be fresh and define

$${}^\circ\alpha_{i+1} = {}^\circ\alpha_i\{a := [{}^\circ R[\square := e\{x_0 := \text{nil}, x_1 := a_1\}[{}^\circ\sigma]]], a_1 := (?_a)\}$$

$${}^\circ v_j = v_j[{}^\circ\sigma] \quad \text{for } j < 2$$

$${}^\circ\mu_{i+1} = {}^\circ\mu_i \cup \{\langle {}^\circ v_0 \Leftarrow {}^\circ v_1 \rangle\}$$

$\square_{\text{s-n}}$

(n-n)

$$\text{let}\{x_0 := \text{newadr()}\}\text{let}\{x_1 := \text{newadr()}\}e$$

$$\text{let}\{x_1 := \text{newadr()}\}\text{let}\{x_0 := \text{newadr()}\}e$$

Let a_0, a_1 be fresh and define

$${}^\circ\alpha_{i+1} = {}^\circ\alpha_i\{a := [{}^\circ R[\square := e\{x_0 := a_0, x_1 := a_1\}[{}^\circ\sigma]]] a_0 := (?_a) a_1 := (?_a)\}$$

$\square_{\text{n-n}}$

(n-i) To complete the proof, we treat the case (**n-i**). Here e_0, e_1 are taken from the pair

$$\text{let}\{x_0 := \text{newadr()}\}\text{let}\{x_1 := \text{initbeh}(v, v_1)\}e$$

$$\text{let}\{x_1 := \text{initbeh}(v, v_1)\}\text{let}\{x_0 := \text{newadr()}\}e.$$

Again, we follow the construction of §5.2.6. In step 1. the only problematic situation is case 3 where a hole is touched but not entered. This can only happen if e_0

attempts the initialization first. Suppose ${}^{\circ}\alpha_{i,j}(a) = [{}^{\circ}R[\square := \circ[{}^{\circ}\sigma]]]$, such that when $\text{initbeh}(v, v_1)$ is placed in the expression hole no transition is enabled. We insert a conditional multi-step which is the empty sequence of steps when filling holes with e_0 and the newadr transition when filling holes with e_1 . ${}^{\circ}\alpha_{i,j+1}(a)$ is obtained by filling the hole with the stuck initbeh . As before we do not distinguish between different stuck expressions placed in the hole. Also, in π_1 there will be an extra uninitialized actor carried along untouched for each such insertion. These extra uninitialized actors can also be ignored. The remainder of the construction is as before. $\square_{\mathbf{n-i}}$

$\square_{\mathbf{c}}$ **commutes**

6.3 Equivalence by Two Stage Reduction

There is one remaining equivalence to establish using common reducts:

(if.lam) $\lambda x.\text{if}(v, e_1, e_2) \cong \text{if}(v, \lambda x.e_1, \lambda x.e_2) \quad x \notin \text{FV}(v)$

The intuitive reasoning behind this equivalence is that for any closing substitution (allowing holes, and actor addresses in the range) the two expressions reduce to equivalent lambda expressions. In fact these lambda expressions have the property that when applied to any argument they reduce to a common expression.

The method developed so far requires reduction to a common local configuration in one stage. Thus we must elaborate the notion of a template to provide for two stages. Specifically, we add a family of holes for lambda-abstractions, which we denote by \triangleright_j for $j \in J$ for some $J \in \mathbb{N} \cup \{\omega\}$.

6.3.1 LE-Syntax

Syntactic classes X with both expression and lambda holes are indicated by the mark ${}^{\circ\triangleright}X$, and we prefix the names of these classes by **LE-**, thus we have **LE-**expressions, **LE-**configurations, etc. The defining clauses are as before with two exceptions: lambda holes are added to the clause generating values; and $\text{app}(\triangleright_j[{}^{\circ\triangleright}\sigma], {}^{\circ\triangleright}v)$ is omitted from the class of **LE-**redexes. The latter exception is made in order to preserve the property that redexes reduce uniformly.

Definition (${}^{\circ\triangleright}\mathbb{V}$, ${}^{\circ\triangleright}\mathbb{E}$, ${}^{\circ\triangleright}\mathbb{S}$, ${}^{\circ\triangleright}\mathbb{R}$, ${}^{\circ\triangleright}\mathbb{E}_{\text{rdx}}$):

$${}^{\circ\triangleright}\mathbb{V} = \text{At} \cup \mathbb{X} \cup \lambda \mathbb{X}.{}^{\circ\triangleright}\mathbb{E} \cup \text{pr}({}^{\circ\triangleright}\mathbb{V}, {}^{\circ\triangleright}\mathbb{V}) \cup \triangleright_{\mathbb{N}}[{}^{\circ\triangleright}\mathbb{S}]$$

$${}^{\circ\triangleright}\mathbb{E} = {}^{\circ\triangleright}\mathbb{V} \cup \Theta_n^e({}^{\circ\triangleright}\mathbb{E}^n) \cup \{\circ[{}^{\circ\triangleright}\mathbb{S}]\}$$

$${}^{\circ\triangleright}\mathbb{S} = \mathbb{X} \xrightarrow{f} {}^{\circ\triangleright}\mathbb{V}$$

$${}^{\circ\triangleright}\mathbb{R} = \{\square\} \cup \Theta_{m+n+1}({}^{\circ\triangleright}\mathbb{V}^m, {}^{\circ\triangleright}\mathbb{R}, {}^{\circ\triangleright}\mathbb{E}^n)$$

$${}^{\circ\triangleright}\mathbb{E}_{\text{rdx}} = \Theta_n^e({}^{\circ\triangleright}\mathbb{V}^n) - \text{app}(\triangleright_j[{}^{\circ\triangleright}\mathbb{S}], {}^{\circ\triangleright}\mathbb{V})$$

Note that lambda holes can occur in the range of a value substitution, and as arguments in redices, except in the function position of an application. Using the double index convention, we write ${}^{\circ\triangleright}e[\triangleright_j := \varphi_j]$ to indicate the simultaneous filling of the holes \triangleright_j with the corresponding lambdas φ_j from some previously specified family

$\{\varphi_j\}_{j \in J}$ of lambda abstractions. The definitions of substitution, free variables, and hole filling are entirely analogous to the expression hole case and we omit them.

The decomposition lemma is modified as follows. An LE-expression ${}^{\circ}e$ is either an LE-value expression (element of ${}^{\circ}\mathbb{V}$), or it can be decomposed uniquely into an LE-reduction context with redex hole filled with either an LE-redex, an LE-expression hole, or an application of a lambda hole (to an LE-value).

Lemma (LE-expression decomposition):

- (0) ${}^{\circ}e \in {}^{\circ}\mathbb{V}$, or
- (1) $(\exists! {}^{\circ}R, {}^{\circ}r)({}^{\circ}e = {}^{\circ}R[\square := {}^{\circ}r])$, or
- (2) $(\exists! {}^{\circ}R, {}^{\circ}\sigma)({}^{\circ}e = {}^{\circ}R[\square := \circ[{}^{\circ}\sigma]])$, or
- (3) $(\exists! {}^{\circ}R, {}^{\circ}\sigma, {}^{\circ}v)({}^{\circ}e = {}^{\circ}R[\square := \mathbf{app}(\triangleright_j[{}^{\circ}\sigma], {}^{\circ}v)])$

6.3.2 LE-computation

The definition of LE-configurations and LE-reduction are the natural extensions of E-configurations and E-reduction to the situation with lambda abstraction holes added. The definition of hole touching and the uniform computation lemmas generalize easily to this situation.

Definition (${}^{\circ}\mathbb{K}$):

$${}^{\circ}\mathbb{K} = \left\langle \left\langle {}^{\circ}\mathbb{A}c \mid {}^{\circ}\mathbb{M} \right\rangle \right\rangle_{\chi}^{\rho}$$

where

$${}^{\circ}\mathbb{A}c = \mathbb{A}d \xrightarrow{f} {}^{\circ}\mathbb{A}s$$

$${}^{\circ}\mathbb{A}s = ({}^{\circ}\mathbb{V}) \cup [{}^{\circ}\mathbb{E}] \cup \{(\mathbb{?}_{\mathbb{A}d})\}$$

$${}^{\circ}\mathbb{M} = \langle {}^{\circ}\mathbb{V} \leftarrow {}^{\circ}\mathbb{V} \rangle$$

and the constraints specified in the definition of actor configurations in §2. are satisfied. We let ${}^{\circ}\kappa$ range over ${}^{\circ}\mathbb{K}$, and ${}^{\circ}\alpha$ range over ${}^{\circ}\mathbb{A}c$. Filling expression and abstraction holes of an LE-configuration, LE-actor map, LE-actor state, LE-multiset of messages, and LE-messages is defined in the obvious manner.

An LE-configuration, ${}^{\circ}\kappa$, is closing for e and a family $\{\varphi_j\}_{j \in J}$ of lambda abstractions if $({}^{\circ}\kappa[\circ := e])[\triangleright_j := \varphi_j]$ is a closed configuration.

Definition (LE-Reduction): The reduction relations $\xrightarrow{\lambda}$ and \mapsto are extended to the generalized domains in the obvious fashion, simply by liberally annotating metavariables with ${}^{\circ}$'s. We omit the details.

Definition (LE-hole touching): If ${}^{\circ}\kappa = \left\langle \left\langle {}^{\circ}\alpha \mid {}^{\circ}\mu \right\rangle \right\rangle$, then ${}^{\circ}\kappa$ touches a hole at a if ${}^{\circ}\alpha(a) = [{}^{\circ}e]$ and either ${}^{\circ}e = {}^{\circ}R[\square := \circ[{}^{\circ}\sigma]]$ or ${}^{\circ}e = {}^{\circ}R[\square := \mathbf{app}(\triangleright_j[{}^{\circ}\sigma], {}^{\circ}v)]$. A transition from ${}^{\circ}\kappa[\circ := e][\triangleright_j := \varphi_j]$ touches a hole at a if the focus actor of the transition is a and ${}^{\circ}\kappa$ touches a hole at a .

Note that since an abstraction hole must be filled with a value, they are not touched in the same ways as arbitrary expression holes, in particular if the transition

is a $\langle \text{fun} : a \rangle$ execution step where ${}^\circ\alpha(a) = [{}^\circ e]$ and ${}^\circ e = {}^\circ\text{R}[\square := \text{app}(\lambda x. {}^\circ e', \triangleright_j [{}^\circ\sigma])]$, then this is not considered touching the hole, \triangleright_j .

The **(E-Uniform Computation)** lemma generalizes to the situation with added abstraction holes.

Lemma (LE-Uniform Computation):

- (1) If ${}^\circ\kappa \xrightarrow{l} {}^\circ\kappa'$, then ${}^\circ\kappa[\circ := e][\triangleright_j := \varphi_j] \xrightarrow{l} {}^\circ\kappa'[\circ := e][\triangleright_j := \varphi_j]$ for any valid filling expression e and family of lambda abstractions φ_j .
- (2) If ${}^\circ\kappa$ has no transition with focus a (and a is an actor of ${}^\circ\kappa$), then either ${}^\circ\kappa$ touches a hole at a or ${}^\circ\kappa[\circ := e][\triangleright_j := \varphi_j]$ has no transition with focus a for any valid filling expression e and family of lambda abstractions φ_j .
- (3) If $\kappa \xrightarrow{l} \kappa'$ and $\kappa = {}^\circ\kappa[\circ := e][\triangleright_j := \varphi_j]$, then either the transition touches a hole or we can find ${}^\circ\kappa'$ such that $\kappa' = {}^\circ\kappa'[\circ := e][\triangleright_j := \varphi_j]$ and ${}^\circ\kappa \xrightarrow{l} {}^\circ\kappa'$.

Proof: Similar to the proof of **(E-uniform computation)** Now there are two cases in which a hole is touched in the decomposition of ${}^\circ e$, namely cases (2) and (3) of the decomposition lemma. \square

6.3.3 LE-Main Theorem

Now we have developed sufficient notation and machinery to state and prove a general result giving equivalence via two-stage reduction.

Theorem (eq-reduct): Let $e_0, e_1, \varphi_{0,j}, \varphi_{1,j}$ for $j < J$ be such that for each ${}^\circ\sigma$ we can find $j \in J$ such that $e_i[{}^\circ\sigma]$ reduces uniformly via $\xrightarrow{\lambda}$ steps to $\varphi_{i,j}[{}^\circ\sigma]$ for $i < 2$, and that for each ${}^\circ\sigma, {}^\circ v$, and $j \in J$ we can find ${}^\circ e_c$ such that $\text{app}(\varphi_{i,j}[{}^\circ\sigma], {}^\circ v)$ reduces uniformly via $\xrightarrow{\lambda}$ steps to ${}^\circ e_c$ for $i < 2$. Then $e_0 \cong e_1$.

Corollary (eq-reduct): **(if.lam)** is an example. Here we take

$$e_0 = \lambda x. \text{if}(v, e_a, e_b)$$

$$e_1 = \text{if}(v, \lambda x. e_a, \lambda x. e_b) \quad \text{where } x \notin \text{FV}(v)$$

$$J = \{a, b\}$$

$$\varphi_{0,j} = \lambda x. \text{if}(v, e_a, e_b)$$

$$\varphi_{1,j} = \lambda x. e_j \text{ for } j \in J$$

Proof: Let ${}^\circ\kappa = \langle \langle {}^\circ\alpha \mid {}^\circ\mu \rangle \rangle$ be a closing LE-configuration for $e_0, e_1, \varphi_{0,j}, \varphi_{1,j} \mid j \in J$. Assume $\pi_0 \in \mathcal{F}({}^\circ\kappa[\circ := e_0][\triangleright_j := \varphi_{0,j}]) = [\kappa_i \xrightarrow{l_i} \kappa_{i+1} \mid i \in \mathbb{N}]$. We want to find ${}^\circ\kappa_i$, and L_i , such that $\kappa_i = {}^\circ\kappa_i[\circ := e_0]$ and, letting $\pi_1 = [{}^\circ\kappa_i[\circ := e_1][\triangleright_j := \varphi_{1,j}] \xrightarrow{L_i} {}^\circ\kappa_{i+1}[\circ := e_1][\triangleright_j := \varphi_{1,j}] \mid i \in \mathbb{N}]$, we have $\pi_1 \in \mathcal{F}({}^\circ\kappa[\circ := e_1])$ and $\text{obs}(\pi_0) = \text{obs}(\pi_1)$. At each stage i we first consider dangling steps. Suppose ${}^\circ\alpha_i(a) = [{}^\circ\text{R}[\square := \circ[{}^\circ\sigma]]]$ and the reduction of e_0 is trivial, i.e. $e_0[{}^\circ\sigma] = \varphi_{0,j}[{}^\circ\sigma]$ for some j . Then we prefix L_i with the transitions for $e_1[{}^\circ\sigma] \xrightarrow{\lambda} \varphi_{1,j}[{}^\circ\sigma]$ and convert the E-hole to \triangleright_j . This is done for each $a \in \text{Dom}({}^\circ\alpha_i)$ meeting the condition.

Now we consider the decomposition of the configuration at stage $i+1$ in the case l_i touches a hole. Suppose l_i is an execution by a with ${}^\circ\alpha_i(a) = [{}^\circ\text{R}[\square := \circ[{}^\circ\sigma]]]$. By

the elimination of ‘dangling steps’ we may assume that e_0 is not a value expression and hence the execution occurs at the hole. Suppose also that $e_i[\circ\sigma] \xrightarrow{\lambda} \varphi_{i,j}[\circ\sigma]$. Define $\circ\alpha_{i+1} = \circ\alpha_i\{a := [\circ R[\square := \triangleright_j[\circ\sigma]]]\}$.

Suppose l_i is an execution by a with $\circ\alpha_i(a) = [\circ R[\square := \mathbf{app}(\triangleright_j[\circ\sigma], \circ v)]]$. Suppose also that $\varphi_{i,j}[\circ\sigma]$ reduces to $e_c[\circ\sigma]$ by $\xrightarrow{\lambda}$ steps. Define $\circ\alpha_{i+1} = \circ\alpha_i\{a := [\circ R[\square := e_c[\circ\sigma]]]\}$.

It is easy to check (as in the proof of **(fun-red-eq)**) that fairness is preserved. \square

6.4 Equivalence of Reduction Contexts

To establish the equivalence of reduction contexts, we define templates for syntactic entities — expressions, reduction context, redexes, configurations — with holes to be filled by a reduction context. We then proceed as before, to show how configurations can be suitably decomposed in order to define the desired path correspondences.

6.4.1 R-Syntax

We use \diamond for reduction context holes and signify the corresponding syntactic entities with a mark \diamond . We prefix names of templates for syntactic classes by R-, thus expression templates are called R-expressions, etc.

Definition ($\diamond\mathbb{E}$ $\diamond\mathbb{V}$):

$$\begin{aligned} \diamond\mathbb{V} &= \mathbf{At} \cup \mathbf{X} \cup \lambda\mathbf{X}. \diamond\mathbb{E} \cup \mathbf{pr}(\diamond\mathbb{V}, \diamond\mathbb{V}) \\ \diamond\mathbb{E} &= \diamond\mathbb{V} \cup \Theta_n^e(\diamond\mathbb{E}^n) \cup \diamond\mathbb{S}[\square := \diamond\mathbb{E}] \\ \diamond\mathbb{S} &= \mathbf{X} \xrightarrow{f} \diamond\mathbb{V} \end{aligned}$$

$\diamond e[\diamond := R]$ is the result of filling R-holes in $\diamond e$ with R . We give only the clause for the hole case in the recursive definition of filling.

Definition ($\diamond e[\diamond := R]$):

$$\begin{aligned} (\diamond^\sigma[\square := \diamond e])[\diamond := R] &= R[\sigma[\square := \diamond e[\diamond := R]]] \\ \text{where } \sigma &= \diamond\sigma[\diamond := R] = \lambda\mathbf{x} \in \text{Dom}(\diamond\sigma). \diamond\sigma(\mathbf{x})[\diamond := R] \end{aligned}$$

Definition ($\diamond\mathbb{R}$ $\diamond\mathbb{E}_{\text{rdx}}$):

$$\begin{aligned} \diamond\mathbb{R} &= \{\square\} \cup \Theta_{m+n+1}(\diamond\mathbb{V}^m, \diamond\mathbb{R}, \diamond\mathbb{E}^n) \cup \diamond\mathbb{S}[\square := \diamond\mathbb{R}] \\ \diamond\mathbb{E}_{\text{rdx}} &= \Theta_n^e(\diamond\mathbb{V}^n) \end{aligned}$$

The clauses directly involving holes in the definitions of hole filling for R-reduction contexts are:

$$\begin{aligned} \square[\diamond := R] &= \square \\ \square[\square := \diamond e] &= \diamond e \\ (\diamond^\sigma[\square := \diamond R])[\diamond := R] &= R^\sigma[\square := \diamond R[\diamond := R]] \\ \text{where } \sigma &= \diamond\sigma[\diamond := R] = \lambda\mathbf{x} \in \text{Dom}(\diamond\sigma). \diamond\sigma(\mathbf{x})[\diamond := R] \end{aligned}$$

$$(\diamond^\sigma[\square := \diamond R])[\square := \diamond e] = \diamond^\sigma[\square := \diamond R[\square := \diamond e]]$$

Note that filling R-holes in R-expressions, R-reduction contexts, or R-redexes with a reduction context yields an expression, a reduction context, or redex, respectively.

An R-expression $\diamond e$ is either an R-value (element of $\diamond\mathbb{V}$) or it can be decomposed uniquely into an R-reduction context filled with either an R-redex or an R-hole.

Lemma (R-expression decomposition):

- (0) $\diamond e \in \diamond\mathbb{V}$, or
- (1) $(\exists! \diamond R, \diamond r)(\diamond e = \diamond R[\square := \diamond r])$, or
- (2) $(\exists! \diamond R, \diamond \sigma, \diamond v)(\diamond e = \diamond R[\square := \diamond^\sigma[\square := \diamond v]])$

Proof : An easy induction on the structure of $\diamond e$. \square

6.4.2 R-Configurations

Definition ($\diamond\mathbb{K}$): An R-configuration for reduction contexts $\diamond\kappa$ is formed like a configuration, but using R-expressions instead of simple expressions.

$$\diamond\mathbb{K} = \left\langle \left\langle \diamond\mathbb{A}c \mid \diamond\mathbb{M} \right\rangle_x \right\rangle_\rho$$

$$\diamond\mathbb{A}c = \mathbb{A}d \xrightarrow{f} \diamond\mathbb{A}s$$

$$\diamond\mathbb{A}s = (\diamond\mathbb{V}) \cup [\diamond\mathbb{E}] \cup \{(\diamond\mathbb{A}d)\}$$

$$\diamond\mathbb{M} = \langle \diamond\mathbb{V} \leftarrow \diamond\mathbb{V} \rangle$$

We let $\diamond\kappa$ range over $\diamond\mathbb{K}$, and $\diamond\alpha$ range over $\diamond\mathbb{A}c$. Filling holes of an R-configuration is analogous to filling holes of an E-configuration. An R-configuration $\diamond\kappa$ is closing for R if $\diamond\kappa[\diamond := R]$ is a closed configuration.

6.4.3 R-Reduction

Definition (touching R-holes): A transition l from $\diamond\kappa[\diamond := R]$ touches an R-hole at a if l is an execution transition with focus a and execution state of $\diamond\kappa$ at a decomposes according to case (2) of the decomposition lemma.

6.4.4 R-Uniform Computation

Lemma (Uniform Computation): An R-redex reduces or hangs uniformly (for a given enabling occurrence in a configuration). Hence transitions not touching an R-hole are uniform. More precisely, if $\diamond\kappa_i[\diamond := R] \xrightarrow{l_i} \kappa_{i+1}$, with focus a , that does not touch an R-hole at a , then l_i is either a receive or an execution transition in which the execution state of $\diamond\kappa_i$ at a decomposes according to case (1) of the decomposition lemma. Let $\diamond\kappa_i = \left\langle \left\langle \diamond\alpha_i \mid \diamond\mu_i \right\rangle \right\rangle$. Then the decomposition of $\kappa_{i+1} = \left\langle \left\langle \diamond\alpha_{i+1} \mid \diamond\mu_{i+1} \right\rangle \right\rangle$ is defined as follows.

Receive: In the receive case we must have that $\circ\alpha_i(a) = (\lambda x.\circ e)$. Thus $\circ\alpha_{i+1} = \circ\alpha_i\{a := [\mathbf{app}(\lambda x.\circ e, cv)]\}$, and $\circ\mu_i = \circ\mu_{i+1} + \langle a \Leftarrow cv \rangle$.

Uniform execution: In the uniform execution case $\circ\alpha_i(a) = [\circ e]$, where $\circ e$ has the form $\circ R[\square := \circ r]$. In this case the step is independent of what fills the holes.

Thus we can find $\circ\kappa_{i+1}$ such that $\circ\kappa_i \xrightarrow{l_i} \circ\kappa_{i+1}$ uniformly.

Now we show how to establish equivalence of expressions of the form $R_j[\square := e]$ for $j < 2$ where the $R_j[\square := v]$ have a common reduct for any value expression.

Theorem (eq-r): If for z fresh, there is some e such that $R_j[\square := z]$ reduces uniformly via 0 or more $\xrightarrow{\lambda}$ steps to e for $j < 2$, then $R_0[\square := e] \cong R_1[\square := e]$ for any e .

Corollary (eq-r): (**app**), (**cmps**), (**id**), (**letx**), (**let.dist**), (**if.dist**) are instances of (**eq-r**).

In fact we prove a slightly more general result, since we use a weaker assumption on the reduction contexts: for each $\circ\sigma$, $\circ v$, we can find $\circ e_c$ such that $R_j^{\circ\sigma}[\square := \circ v]$ reduces uniformly via 0 or more $\xrightarrow{\lambda}$ steps to $\circ e_c$ for $j < 2$.

Proof (eq-r): Suppose R_0, R_1 are reduction contexts that we wish to establish the observational equivalence of. That is, we want to show $R_0[\square := e] \cong R_1[\square := e]$ for all expressions e . Let $\circ\kappa = \langle\langle \circ\alpha \mid \circ\mu \rangle\rangle$ be a closing R-configuration for R_0, R_1 .

Assume $\pi_0 \in \mathcal{F}(\circ\kappa[\diamond := R_0]) = [\kappa_i \xrightarrow{l_i} \kappa_{i+1} \mid i \in \mathbb{N}]$. We want to find $\circ\kappa_i$, and L_i , such that $\kappa_i = \circ\kappa_i[\diamond := R_0]$ and, letting $\pi_1 = [\circ\kappa_i[\diamond := R_1] \xrightarrow{L_i} \circ\kappa_{i+1}[\diamond := R_1] \mid i \in \mathbb{N}]$, we have $\pi_1 \in \mathcal{F}(\circ\kappa[\diamond := R_1])$ and $obs(\pi_0) = obs(\pi_1)$. As in the expression context case, we can focus our attention on the construction of the actor configuration part, since here also deliverable messages cannot have holes. Assume we have $\circ\alpha_i$ and consider cases on the transition label l_i . We have three cases to consider. For the base case we have $\circ\alpha_0 = \circ\alpha$. At stage i suppose $\circ\alpha_i(a') = \circ R[\square := \diamond^{\circ\sigma}[\square := \circ v]]$. If $R_0^{\circ\sigma}[\square := \circ v]$ is the common form (i.e. the reduction is trivial), then we prefix L_i with steps for reduction of $R_1^{\circ\sigma}[\square := \circ v]$ to common form and remove this hole. This is carried out for each a' in the domain of $\circ\alpha_i$.

Now, suppose l_i is an execution with focus a that touches a hole. Thus $\circ\alpha_i(a) = [\circ e]$ where $\circ e$ has the form $\circ R[\square := \diamond^{\circ\sigma}[\square := \circ v]]$. Suppose also that $R_j^{\circ\sigma}[\square := \circ v] \xrightarrow{\lambda} \circ e_c$ for $j < 2$. Then we define $\circ\alpha_{i+1} = \circ\alpha_i\{a := [\circ R[\square := \circ e_c]]\}$.

Let $\pi_1 = [\circ\kappa_i[\diamond := R_1] \xrightarrow{L_i} \circ\kappa_{i+1}[\diamond := R_1] \mid i < \mathbb{N}]$ be the constructed computation path. Note that the transitions are the same except for the points where holes are touched, but these differences are not observable. Clearly, under the assumptions on R_j , π_1 is a computation path. It remains to show that the construction preserves fairness.

Suppose some transition l is enabled at stage i in π_1 . If l is a receive or uniform execution, then l is also enabled in π_0 at stage i and will eventually occur, uniformly. Suppose l is an execution step by a with $\circ\alpha_i(a) = [\circ R[\square := \diamond^{\circ\sigma}[\square := \circ v]]]$. Either $R_0^{\circ\sigma}[\square := \circ v]$ is in common form and the transition occurs as soon as it is enabled in π_1 , or a transition is enabled at this hole in π_0 . This transition will eventually occur. If $R_1^{\circ\sigma}[\square := \circ v]$ is not in common form, l will happen in π_1 at the same stage. If $R_1^{\circ\sigma}[\square := \circ v]$ is in common form, then we must consider whether l occurs at the hole

– i.e. whether or not the common form is an R-value expression. If l occurs at the hole, then the same transition is now enabled in π_0 and will eventually occur in both paths. Otherwise consider the decomposition after the transition in π_0 . As shown before, this reduces to considering a proper subexpression of $\diamond R[\square := \diamond^\sigma[\square := \diamond v]]$ and the process will eventually terminate with a uniform execution.

□eq-r

7 Discussion

In this paper we presented an operational semantics of actor computation. The actor language is an extension of a call-by-value functional language by primitives for creating and manipulating actors. Central to the theory is the concept of an *actor configuration* that makes explicit the notion of open system component. A composition operation on actor configurations is defined. It is associative, commutative, has a unit, and is thus a first step towards an algebra of configurations. The operational semantics is defined by a labelled transition relation on configurations. This notion of computation incorporates fairness in an essential way. This operational semantics is used to define a notion of observational equivalence of expressions based on traditional operational and testing equivalence. An interesting consequence of fairness is that the classic three testing equivalences collapse to two. Methods for establishing equivalence of expressions are developed and a plethora of laws of expression equivalence that incorporates the equational theory of the embedded functional language are presented. We expect that these methods will be useful in developing equational theories for other concurrent extensions of functional languages such as CML or FACILE.

7.1 Future Directions

The theory presented is perhaps best viewed as a starting point for further research rather than a final product. In particular there are at least four directions for further research. First, work is needed to develop an algebra of operations on configurations. Treating configurations as objects would allow us to abstract over specific linguistic constructs. Second, a systematic study of notions of equivalence within the actor theory, including notions of configuration equivalence. Third, developing a logic for specifying components as actor configurations which would provide methods for verifying that programs implementing components meet their specifications as well as methods for refining specifications into implementations. Finally, it would be useful to determine whether or not there is a context lemma (similar to (ciu)) for observational equivalence of expressions? .

In hindsight, the theory as presented here might benefit from replacing the `newadr`, `initbeh` choice of language primitives by a `letactor` primitive, and by making the operational semantics truly concurrent. Work is in progress to formalize the semantics presented in this paper in the PVS proof development system (Owre et al., 1992) and to develop mechanically checked proofs of the main theorems. This is a first step towards building of a set of tools for computer aided specifica-

from
gul: what does this
mean? Sounds like
more of the same.
Should say what
this would give us.

is
everyone supposed to
know ciu and why
they should care?

tion and design of open distributed systems. The formalization work incorporates these modifications.

Acknowledgements

The authors would like to thank Carl Hewitt and Richard Weyhrauch for many discussions about actor computation that served as a foundation for this work.

This research was partially supported by ARPA contract NAG2-703, NSF grants CCR-8917606, CCR-8915663, CCR-9109070 and INT-89-20626, and ARPA and NSF joint contract CCR 90-07195, ONR contract N00014-90-J-1899, and by the Digital Equipment Corporation.

References

- Abelson, H. and Sussman, G. J. (1985). *Structure and Interpretation of Computer Programs*. The MIT Press, McGraw-Hill Book Company.
- Agha, G. (1986a). *Actors: A Model of Concurrent Computation in Distributed Systems*. MIT Press, Cambridge, Mass.
- Agha, G. (1986b). *Actors: A Model of Concurrent Computation in Distributed Systems*. MIT Press, Cambridge, Mass.
- Agha, G. (1990). Concurrent object-oriented programming. *Communications of the ACM*, 33(9):125–141.
- Agha, G., Frølund, S., Kim, W., Panwar, R., Patterson, A., and Sturman, D. (1993a). Abstraction and Modularity Mechanisms for Concurrent Computing. *IEEE Parallel and Distributed Technology: Systems and Applications*, 1(2):3–14.
- Agha, G., Hewitt, C., Wegner, P., and Yonezawa, A., editors (1991). *Proceedings of the ACM SIGPLAN Workshop on Object-Based Concurrent Programming*. OOPS Messenger, vol. 2, no. 2.
- Agha, G., Mason, I. A., Smith, S. F., and Talcott, C. L. (1992). Towards a theory of actor computation. In *The Third International Conference on Concurrency Theory (CONCUR '92)*, volume 630 of *Lecture Notes in Computer Science*, pages 565–579. Springer Verlag.
- Agha, G., Wegner, P., and Yonezawa, A., editors (1989). *Proceedings of the ACM SIGPLAN Workshop on Object-Based Concurrent Programming*. Special Issue of SIGPLAN Notices.
- Agha, G., Wegner, P., and Yonezawa, A., editors (1993b). *Research Directions in Concurrent Object Oriented Programming*. MIT Press.
- Amadio, R. M. (1994). Translating core facile. Technical Report ECRC-1994-3, European Computer-Industry Research Centre.
- Attardi, G. and Hewitt, C. (1978). Specifying and proving properites of guardians for distributed systems.
- Baker, H. G. and Hewitt, C. (1977). Laws for communicating parallel processes. In *IFIP Congress*, pages 987–992. IFIP.
- Bergstra, J. A. and Klop, J. W. (1986). *Process Algebra: Specification and Verification in Bisimulation Semantics*. Number CWI Monograph 4 in Mathematics and Computer Science, II. North-Holland.
- Berry, D., Milner, R., and Turner, D. (1992). A semantics for ML concurrency primitives. In *Conference record of the 19th annual ACM symposium on principles of programming languages*, pages 105–129.
- Brookes, S., Hoare, C., and Roscoe, A. (1984). A Theory of Communicating Sequential Processes. *Journal of the ACM*, 31(3):560–599.

- Clinger, W. D. (1981). Foundations of actor semantics. AI-TR- 633, MIT Artificial Intelligence Laboratory.
- de Nicola, R. and Hennessy, M. C. B. (1984). Testing equivalences for processes. *Theoretical Computer Science*, 34:83–133.
- Felleisen, M. and Friedman, D. (1986). Control operators, the SECD-machine, and the λ -calculus. In Wirsing, M., editor, *Formal Description of Programming Concepts III*, pages 193–217. North-Holland.
- Felleisen, M. and Hieb, R. (1992). The revised report on the syntactic theories of sequential control and state. *Theoretical Computer Science*, 103:235–271.
- Felleisen, M. and Wright, A. K. (1991). A syntactic approach to type soundness. Technical Report Rice COMP TR91-160, Rice University Computer Science Department. To appear, *Information and Computation*.
- Giacalone, A., Mishra, P., and Prasad, S. (1989). Facile: A symmetric integration of concurrent and functional programming. *International Journal of Parallel Programming*, 18(2):121–160.
- Greif, I. (1975). Semantics of communicating parallel processes. Technical Report 154, MIT, Project MAC.
- Hewitt, C. (1971). *Description and Theoretical Analysis (Using Schemata) of PLANNER: A Language for Proving Theorems and Manipulating Models in a Robot*. PhD thesis, MIT.
- Hewitt, C. (1977). Viewing control structures as patterns of passing messages. *Journal of Artificial Intelligence*, 8(3):323–364.
- Hewitt, C. and Atkinson, R. (1979). Specifications and proof techniques for serializers. *IEEE Transactions on Software Engineering*, SE-5(1):10–23.
- Hewitt, C., Bishop, P., and Steiger, R. (1973). A universal modular actor formalism for artificial intelligence. In *Proceedings of 1973 International Joint Conference on Artificial Intelligence*, pages 235–245.
- Hoare, C. A. R. (1985). *Communicating Sequential Processes*. Prentice-Hall.
- Honsell, F., Mason, I. A., Smith, S. F., and Talcott, C. L. (1994). A Variable Typed Logic of Effects. *Information and Computation*. to appear.
- Kornfeld, W. A. and Hewitt, C. (1981). The scientific community metaphor. *IEEE Transactions on Systems, Man, and Cybernetics*, SMC-11(1).
- Lieberman, H. (1987). Concurrent object-oriented programming in act-1. In Yonezawa, A. and Tokoro, M., editors, *Object-Oriented Concurrent Programming*, pages 9–36. MIT Press.
- Manning, C. (1987). ACORE: The design of a core actor language and its compiler. Master's thesis, MIT, Artificial Intelligence Laboratory.
- Mason, I. A. and Talcott, C. L. (1991). Equivalence in functional languages with effects. *Journal of Functional Programming*, 1:287–327.
- Meseguer, J. (1992). Conditional rewriting logic as a unified model of concurrency. *Theoretical Computer Science*, 96(1):73–155.
- Milner, R. (1983). Calculi for synchrony and asynchrony. *Theoretical Computer Science*, 25:267–310.
- Milner, R. (1989). *Communication and Concurrency*. Prentice Hall.
- Milner, R., Parrow, J. G., and Walker, D. J. (1989). A calculus of mobile processes, Parts I and II. Technical Report ECS-LFCS-89-85, -86, Edinburgh University.
- Moggi, E. (1988). Computational lambda-calculus and monads. Technical Report ECS-LFCS-88-86, University of Edinburgh.
- Morris, J. H. (1968). *Lambda calculus models of programming languages*. PhD thesis, Massachusetts Institute of Technology.
- Owre, S., Rushby, J., and Shankar, N. (1992). PVS: A prototype verification system.
- Plotkin, G. (1975). Call-by-name, call-by-value and the lambda calculus. *Theoretical Computer Science*, 1:125–159.

- Prasad, S., Giacalone, A., and Mishra, P. (1990). Operational and algebraic semantics for facile: A symmetric integration of concurrent and functional programming. In Paterson, M. S., editor, *17th International Colloquium on Automata, Languages and Programming*, volume 443 of *Lecture Notes in Computer Science*, pages 765–780.
- Reppy, J. H. (1991). An operational semantics of first-class synchronous operations. Technical Report TR 91-1232, Cornell University.
- Reppy, J. H. (1992). *Higher-Order Concurrency*. PhD thesis, Cornell University. available as Technical Report 92-1285.
- Sassone, V., Nielsen, M., and Winskel, G. (1993). A hierarchy of models for concurrency. In *The Fourth International Conference on Concurrency Theory (CONCUR '93)*, Lecture Notes in Computer Science. Springer Verlag.
- Talcott, C. (1993a). A theory of binding structures and its applications to rewriting. *Theoretical Computer Science*, 112:99–143.
- Talcott, C. L. (1985). *The essence of Rum: A theory of the intensional and extensional aspects of Lisp-type computation*. PhD thesis, Stanford University.
- Talcott, C. L. (1989). Programming and proving with function and control abstractions. Technical Report STAN-CS-89-1288, Stanford University Computer Science Department.
- Talcott, C. L. (1991). Binding structures. In Lifschitz, V., editor, *Artificial Intelligence and Mathematical Theory of Computation*. Academic Press.
- Talcott, C. L. (1993b). A theory for program and data specification. *Theoretical Computer Science*, 104:129–159.
- Thomsen, B., Leth, L., and Giacalone, A. (1992). Some issues in the semantics of facile distributed programming. Technical Report ECRC-1992-32, European Computer-Industry Research Centre.
- Tomlinson, C., Cannata, P., Meredith, G., and Woelk, D. (1993). The extensible services switch in carnot. *IEEE Parallel and Distributed Technology: Systems and Applications*, 1(2):16–20.
- Tomlinson, C., Kim, W., Schevel, M., Singh, V., Will, B., and Agha, G. (1989). Rosette: An object oriented concurrent system architecture. *Sigplan Notices*, 24(4):91–93.

8 Index of Notations (Not including §5)

Symbol	Description	§
\mathbb{N}	The natural numbers, $i, j, \dots, n \in \mathbb{N}$	1.3
Y^n	Sequences from Y of length n , $\bar{y} = [y_1, \dots, y_n] \in Y^n$	1.3
Y^*	Finite sequences from Y	1.3
$[]$	The empty sequence	1.3
$\text{Len}(\bar{y})$	The length of the sequence \bar{y}	1.3
$u * v$	The concatenation of the sequences u and v	1.3
$\text{Last}(u)$	The last element of the sequence u	1.3
$\mathbf{P}_\omega[Y]$	Finite subsets of Y	1.3
$\mathbf{M}_\omega[Y]$	Finite multi-sets with elements in Y	1.3
$Y_0 \xrightarrow{f} Y_1$	Finite maps from Y_0 to Y_1	1.3
$Y_0 \rightarrow Y_1$	Total functions from Y_0 to Y_1	1.3
$\text{Dom}(f)$	The domain of the function f	1.3
$\text{Rng}(f)$	The range of the function f	1.3
$f\{y := y'\}$	An extension to, or alteration of, the function f	1.3
$f Y$	The restriction of f to the set Y	1.3
\mathbb{X}	A countably infinite set of variables, $x, y, z \in \mathbb{X}$	2.1
At	Atoms	2.1
\mathbf{t}, \mathbf{nil}	Atoms playing the role of booleans	2.1
\mathbb{G}_n	n -ary algebraic operations	2.1
\mathbb{F}	Operations, $\delta \in \mathbb{F}$	2.1
\mathbb{F}_n	n -ary operation symbols	2.1
\mathbb{F}_0	Zero-ary operation symbols $\supseteq \{\mathbf{newadr}\}$	2.1
\mathbb{F}_1	Unary operation symbols $\supseteq \{\mathbf{isatom}, \mathbf{isnat}, \mathbf{ispr}, \mathbf{1^{st}}, \mathbf{2^{nd}}, \mathbf{become}\}$	2.1
\mathbb{F}_2	Binary operation symbols $\supseteq \{\mathbf{pr}, \mathbf{initbeh}, \mathbf{send}\}$	2.1
\mathbb{F}_3	Ternary operation symbols $\supseteq \{\mathbf{br}\}$	2.1
\mathbb{L}	λ -abstractions, $\lambda x.e \in \mathbb{L}$	2.1
\mathbb{V}	Value expressions, $v \in \mathbb{V}$	2.1
\mathbb{E}	Expressions, $e \in \mathbb{E}$	2.1
$\lambda x.e$	Abstractions	2.1
$\mathbf{app}(e_0, e_1)$	Application	2.1
$\delta(\bar{e})$	Application of operations	2.1
$\mathbf{if}(e_0, e_1, e_2)$	Conditional branching	2.1
$\mathbf{let}\{x := e_0\}e_1$	Lexical variable binding	2.1
$\mathbf{seq}(e_1, \dots, e_n)$	Sequencing construct	2.1

Symbol	Description	§
$FV(e)$	The free variables of the expression e	2.1
$e[x := e']$	The result of substituting e' for x in e	2.1
\mathbb{C}	Contexts, $C \in \mathbb{C}$	2.1
\bullet	The hole in contexts	2.1
$C[e]$	The result of filling the context with e	2.1
$\mathbb{A}d$	Actor addresses, identified with \mathbb{X}	2.2
$\mathbb{A}s$	Actor states, $(?_a), (b), [e] \in \mathbb{A}s$	2.2
$(?_a)$	An uninitialized actor created by a	2.2
(b)	An actor with behavior b ready to accept a message	2.2.1
$[e]$	A processing actor with current computation e	2.2.1
\mathbb{M}	Messages, $\langle V \Leftarrow V \rangle \in \mathbb{M}$	2.2.1
cV	Communicable values, $cv \in cV$	2.2.1
$\langle\langle \alpha \mid \mu \rangle\rangle_x^\rho$	An actor configuration with: α – an actor map μ – a multi-set of messages ρ – the receptionists χ – the external actors	2.2.1
\mathbb{K}	Actor configurations, $\kappa \in \mathbb{K}$	2.2.2
\mathbb{E}_{rdx}	The set of redexes, $r \in \mathbb{E}_{rdx}$	2.2.2
\mathbb{R}	The set of reduction contexts, $R \in \mathbb{R}$	2.2.2
\square	The reduction context hole	2.2.2
\mapsto^λ	The reduction relation for functional redexes, $e_0 \mapsto^\lambda e_1$	2.2.2
\mapsto	The reduction relation for configurations, $\kappa_0 \mapsto \kappa_1$	2.2.2
$\kappa_0 \xrightarrow{l} \kappa_1$	$\kappa_0 \mapsto \kappa_1$ via the rule labelled by l	2.2.2
Labels :	Transition labels, $l \in \text{Labels}$	2.2.2
$\langle \text{fun} : a \rangle$	A functional transition	2.2.2
$\langle \text{new} : a, a' \rangle$	newadr redex transition with focus a	2.2.2
$\langle \text{init} : a, a' \rangle$	initbeh redex transition with focus a	2.2.2
$\langle \text{bec} : a, a' \rangle$	become redex transition with focus a	2.2.2
$\langle \text{send} : a, m \rangle$	send redex transition with focus a	2.2.2
$\langle \text{rcv} : a, cv \rangle$	The receipt of a message with focus a	2.2.2
$\langle \text{out} : m \rangle$	A message exiting the configuration	2.2.2
$\langle \text{in} : m \rangle$	A message entering the configuration	2.2.2

Symbol	Description	§
$\mathcal{T}(\kappa)$	All finite sequences of labeled transitions from κ , $\nu \in \mathcal{T}(\kappa)$	2.2.3
$\mathcal{T}^\infty(\kappa)$	the set of all computation paths in $\mathcal{T}(\kappa)$, $\pi \in \mathcal{T}^\infty(\mathbb{K})$	2.2.3
$\bowtie \in \mathbb{N} \cup \{\omega\}$	The length of a finite or infinite sequence	2.2.3
$\text{Cfig}(\kappa, L, i)$	The i th configuration of the computation from κ via L	2.2.3
$\kappa \xrightarrow{L} \kappa'$	A multi-step transition	2.2.3
$\mathcal{F}(\kappa)$	the fair subset of $\mathcal{T}^\infty(\kappa)$	2.2.4
event	A zero-ary primitive/observation	3.1
$\langle e : a \rangle$	An observation transition	3.1
$\langle\langle \alpha, [C]_a \mid \mu \rangle\rangle$	An observing configuration	3.1
\mathbb{O}	The set of observing configurations, $O \in \mathbb{O}$	3.1
s	Signifies that an event transition occurs	3.1
f	Signifies that an event transition does not occur	3.1
$obs(\pi)$	The s/f classification of the path π	3.1
$Obs(\kappa)$	The s/f classification of the configuration, $\kappa, \in \{\mathbf{s}, \mathbf{f}, \mathbf{sf}\}$	3.1
$e_0 \cong_1 e_1$	Testing or Convex or Plotkin or Egli-Milner equivalence	3.2
$e_0 \cong_2 e_1$	Must or Upper or Smyth equivalence	3.2
$e_0 \cong_3 e_1$	May or Lower or Hoare equivalence	3.2
$e_0 \cong e_1$	Operational equivalence (either \cong_1 or equivalently \cong_2)	3.2
<i>Hang</i>	The set of all stuck expressions, $\mathbf{stuck} \in \text{Hang}$	4.1
<i>Infn</i>	The set of all diverging expressions, $\mathbf{bot} \in \text{Infn}$	4.1