

## Actor Model of Computation

Carl Hewitt  
<http://carlhewitt.info>

*This paper is dedicated to Alonzo Church and Dana Scott.*

The Actor model is a mathematical theory that treats “Actors” as the universal primitives of concurrent digital computation. The model has been used both as a framework for a theoretical understanding of concurrency, and as the theoretical basis for several practical implementations of concurrent systems.

Unlike previous models of computation, the Actor model was inspired by physical laws. It was also influenced by the programming languages Lisp, Simula 67 and Smalltalk-72, as well as ideas for Petri Nets, capability-based systems and packet switching. The advent of massive concurrency through client-cloud computing and many-core computer architectures has galvanized interest in the Actor model.

An Actor is a computational entity that, in response to a message it receives, can concurrently:

- send a finite number of messages to other Actors;
- create a finite number of new Actors;
- designate the behavior to be used for the next message it receives.

There is no assumed order to the above actions and they could be carried out concurrently. In addition two messages sent concurrently can arrive in either order.

Decoupling the sender from communications sent was a fundamental advance of the Actor model enabling asynchronous communication and control structures as patterns of passing messages.

## Contents

Introduction.....	3
Fundamental concepts.....	3
Illustrations .....	3
Modularity thru Direct communication and asynchrony .....	3
Indeterminacy and Quasi-commutativity.....	4
Locality and Security .....	4
Computational Representation Theorem .....	4
iAdaptive™ concurrency using ActorScript™ .....	5
Other models of message-passing concurrency .....	5
Historical Background .....	5
Concurrency versus Turing’s Model.....	5
$\lambda$ -Calculus .....	6
Petri nets.....	6
Simula .....	6
Planner.....	7
Smalltalk-72 .....	7
Concurrency .....	8
Actors .....	8
Actor Semantics .....	9
Logic Programming .....	10
Reasoning about Actors .....	10
Early Actor Programming languages .....	10
Causal Consistency .....	10
Extension versus Specialization .....	12
Early Actor Programming languages .....	12
Language constructs versus Library APIs .....	13
Hairy Control Structure Redux .....	13
Promises .....	14
Garbage Collection.....	14
Cosmic Cube .....	14
J-Machine .....	15
$\pi$ -Calculus .....	15
Was the Actor model premature?.....	15
Acknowledgment .....	16
Bibliography .....	16
End Notes.....	21

## Introduction<sup>i</sup>

The Actor model is a mathematical theory of computation that treats “Actors” as the universal primitives of concurrent digital computation [Hewitt, Bishop, and Steiger 1973; Hewitt 1977]. The model been used both as a framework for a theoretical understanding of concurrency, and as the theoretical basis for several practical implementations of concurrent systems.

Unlike previous models of computation, the Actor model was inspired by physical laws. It was also influenced by the programming languages Lisp [McCarthy *et. al.* 1962], Simula 67 [Dahl and Nygaard 1967] and Smalltalk-72 [Kay 1975], as well as ideas for Petri Nets [Petri 1962], capability-based systems [Dennis and van Horn 1966] and packet switching [Baran 1964]. The advent of massive concurrency through client-cloud computing and many-core computer architectures has galvanized interest in the Actor model [Hewitt 2009b].

## Fundamental concepts

An Actor is a computational entity that, in response to a message it receives, can concurrently:

- send a finite number of messages to other Actors;
- create a finite number of new Actors;
- designate the behavior to be used for the next message it receives.

There is no assumed order to the above actions and they could be carried out concurrently. In addition two messages sent concurrently can arrive in either order.

Decoupling the sender from communications sent was a fundamental advance of the Actor model enabling asynchronous communication and control structures as patterns of passing messages [Hewitt 1977].

An Actor can only communicate with Actors to which it is connected. It can directly obtain information only from other Actors to which it is directly connected. Connections can be implemented in a variety of ways:

- direct physical attachment
- memory or disk addresses
- network addresses
- email addresses

The Actor model is characterized by inherent concurrency of computation within and among Actors, dynamic creation of Actors, inclusion of Actor addresses in messages, and interaction only through direct asynchronous message passing with no restriction on message arrival order.

## Illustrations

The Actor model can be used as a framework for modeling, understanding, and reasoning about, a wide range of concurrent systems. For example:

- Electronic mail (e-mail) can be modeled as an Actor system. Accounts are modeled as Actors and email addresses as Actor addresses.
- Web Services can be modeled with SOAP endpoints modeled as Actor addresses.
- Objects with locks (*e.g.* as in Java and C#) can be modeled as Actors.

## Modularity thru Direct communication and asynchrony

Modularity in the Actor model comes from one-way asynchronous communication. Once a message has been sent, it is the responsibility of the receiver.<sup>ii</sup>

Messages in the Actor model are decoupled from the sender and are delivered by the system on a best efforts basis. This was a sharp break with previous approaches to models of concurrent computation in which message sending is tightly coupled with the sender and sending a message synchronously transfers it someplace, *e.g.*, to a buffer, queue, mailbox, channel, broker, server, *etc.* or to the “*ether*” or “*environment*” where it temporarily resides. The lack of synchronicity caused a great deal of misunderstanding at the time of the development of the Actor model and is still a controversial issue.

Because message passing is taken as fundamental in the Actor model, there cannot be any inherent overhead, *e.g.*, any requirement for buffers, pipes, queues, classes, channels, *etc.* Prior to the Actor model, concurrency was defined in low level machine terms.

It certainly is the case that implementations of the Actor model typically make use of these hardware capabilities. However, there is no reason that the model could not be implemented directly in hardware without exposing any hardware threads, locks, queues, channels, tasks, *etc.* Also, there is no necessary relationship between the number of Actors and the number threads, locks, tasks, queues, *etc.* that might be in use. Implementations of the Actor model are free to make use of threads, locks, tasks, queues, global assignment, coherent memory, transactional memory, cores, *etc.* in any way that is compatible with the laws for Actors [Baker and Hewitt 1977].

As opposed to the previous approach based on composing sequential processes, the Actor model was developed as an inherently concurrent model. In the Actor model sequential

ordering is a special case that derived from concurrent computation.

A natural development of the Actor model was to allow Actor addresses in messages. A computation might need to send a message to a recipient from which it would later receive a response. The way to do this is to send a communication which has the message along with the address of another Actor called the *customer* along with the message. The recipient could then cause a response message to be sent to the customer.

Of course, any Actor could be used as a customer to receive a response message. By using customers, common control structures such a recursion, co-routines, hierarchical parallelism, futures [Baker and Hewitt 1977], *etc.* can be implemented.

Actors is based on “behavior” as opposed to the “class” concept of object-oriented programming [Dahl and Nygaard 1967] in which the operations on objects are “inherited” in a hierarchy of classes that is expected to be shared by the caller and implementation of an operation.

### Indeterminacy and Quasi-commutativity

The Actor model supports indeterminacy because the arrival order of messages can affect future behavior. Operations are said to be quasi-commutative to the extent that it doesn’t matter in which order they occur. To the extent possible, quasi-commutativity is used to tame indeterminacy. <sup>iii</sup>

### Locality and Security

Another important characteristic of the Actor model is locality. Locality means that in processing a message: an Actor can send messages only to addresses for which it has information by the following means:

1. that it receives in the message
2. that it already had before it received the message
3. that it creates while processing the message.

In the Actor model, there is no simultaneous change in multiple locations. In this way it differs from some other models of concurrency, *e.g.*, the Petri net model in which tokens are simultaneously removed from multiple locations and placed in other locations.

The security of Actors can be protected in the following ways:

- hardwiring in which Actors are physically connected
- tagged memory as in Lisp machines, *etc.*
- virtual machines as in Java virtual machine, Common Language Runtime, *etc.*
- signing and/or encryption of Actors and their addresses

A delicate point in the Actor model is the ability to synthesize the address of an Actor. In some cases security can be used to prevent the synthesis of addresses. However, if an Actor address is simply a bit string then clearly it can be synthesized given eons of time although it can be made unguessable in practice if the bit strings are long enough.

### Computational Representation Theorem

There is a *Computational Representation Theorem* in the Actor model for systems which are closed in the sense that they do not receive communications from outside. A closed system is defined to be one which does not communicate with the outside. Actor model theory provides the means to characterize all the possible computations of a closed system in terms of the Computational Representation Theorem [Clinger 1981; Hewitt 2006]:

The denotation  $\text{Denote}_S$  of a closed system  $S$  represents all the possible behaviors of  $S$  as

$$\text{Denote}_S = \bigsqcup_{i \in \mathbb{N}} \text{Progressions}_S^i(\perp_S)$$

where  $\text{Progressions}$  is an approximation function that takes a set of partial behaviors to their next stage and  $\perp_S$  is the initial behavior of  $S$ .

In this way, the behavior of  $S$  can be mathematically characterized in terms of all its possible behaviors.<sup>iv</sup>

Although  $\text{Denote}_S$  is not an implementation of  $S$ , it can be used to prove a generalization of the Church-Turing-Rosser-Kleene thesis [Kleene 1943]:

**Enumeration Theorem:** If the primitive Actors of a closed Actor System  $S$  are effective, then the possible outputs of  $S$  are recursively enumerable.

*Proof:* Follows immediately from the Representation Theorem.

The upshot is that *concurrent systems can be represented and characterized by logical deduction but cannot be implemented.* Thus, the following practical problem arose:

How can practical programming languages be rigorously defined since the proposal by Scott and Strachey [1971] to define them in terms  $\lambda$ -calculus failed because the  $\lambda$ -calculus cannot implement concurrency?<sup>v</sup>

## iAdaptive™ concurrency using ActorScript™

iAdaptive concurrency is the ability to expression computations with discretionary activity that can be adapted s to fit available resources in terms of demand and available capacity (*e.g.* cores).

The ActorScript programming languages has been developed to express iAdaptive concurrency as follows:

- Universality
  - Everything in the language is accomplished using message passing including the very definition of the language itself.
  - Directly express discretionary push and pull concurrency
  - Functional and Logic Programming are integrated into general concurrent programming.
  - Programs do not expose low-level implementation mechanisms such as threads, tasks, locks, cores, *etc.* Messages can be directly communicated without requiring indirection through brokers, class hierarchies, channels, mailboxes, pipes, ports, queues *etc.*
  - Program resource control of energy, processing, storage, and communications using sponsors.
  - Binary XML and JSON are data types
  - Applications cannot directly harm each other.
- Efficiency
  - Allow execution to be dynamically adjusted for system load and capacity (cores)
  - Minimize latency along critical paths

ActorScript attempts to achieve the highest level of performance, scalability, and expressibility with a minimum of primitives.

## Other models of message-passing concurrency

The Actor Model does not have the following restrictions of other models of message-passing concurrency:<sup>vi</sup>

- *Single threadedness*: There are no restrictions on the use of threads in implementations.
- *Message delivery order*: There no restrictions on message delivery order.
- *Independence of sender*: The semantics of a message in the Actor model is independent of the sender.
- *Lack of garbage collection (automatic storage reclamation)*: The Actor Model can be used in the following systems:

- CLR and extensions (Microsoft)
- JVM (Sun, IBM, Oracle, SAP)
- LLVM (Apple)
- Dalvik (Google)

In due course, we will need to extend the above systems with a tagged extension of the X86 architecture. Many-core architecture has made a tagged extension necessary in order to provide the following:

- concurrent, nonstop, no-pause automatic storage reclamation (garbage collection) and relocation to improve efficiency,
- prevention of memory corruption that otherwise results from programming languages like C and C++ using thousands of threads in a process,
- nonstop migration of Actors (while they are in operation) within a computer and between distributed computers

Imposing restrictions on threadedness, message order arrival, and garbage collection (automatic storage management) can greatly diminish the performance and scalability of Actor systems.

## Historical Background<sup>vii</sup>

The Actor model builds on previous models of nondeterministic computation. Several models of nondeterministic computation were developed including the following:

### *Concurrency versus Turing's Model*

Turing's model of computation was intensely psychological.<sup>viii</sup> Sieg [2008] formalized it as follows:

- *Boundedness*: A computer can immediately recognize only a bounded number of configurations.
- *Locality*: A computer can change only immediately recognizable configurations.

In the above, computation is conceived as being carried out in a single place by a device that proceeds from one well-defined state to the next.

By contrast, in the Actor model [Hewitt, Bishop and Steiger 1973; Hewitt 2010b], computation is conceived as distributed in space where computational devices communicate asynchronously and the entire computation is not in any well-defined state.<sup>ix</sup> Properties of *Locality* and *Local Finiteness* hold as follows:

- *Locality*: In response to a message received, an Actor can change only its local storage that can include addresses only of Actors provided when it was created and those that it has received in messages.

- *Local Finiteness*: At any point, an Actor has finite storage. In response to a message received, an Actor can directly
  - send only a finite number of messages to addresses in the message and its local storage
  - create only a finite number of Actors

Actor systems can perform computations that are impossible by Turing Machines as illustrated by the following example:

1. There is a bound on the size of integer that can be computed by an *always halting* nondeterministic Turing Machine starting on a blank tape.
2. By contrast, there are always-halting Actor systems with no inputs that can compute an integer of unbounded size. (An Actor can be created with local storage that is initialized with a count of 0 that concurrently sends itself both a *stop* and a *go* message. When it receives a *go* message, it increments its count by 1 and sends itself a *go* message. When it receives a *stop* message, it stops with an unbounded number in its local storage.)

Another point of departure from Turing's model is that concurrency violates a narrowly conceived "public processes" [Hofstadter 1980] criterion for computation. Actor systems make use of hardware devices called arbiters to decide arrival orderings of messages. The internal processes of arbiters are not public processes.<sup>x</sup> Instead of observing the internals of arbitration processes, we necessarily await outcomes.

## *λ-Calculus*

The  $\lambda$ -calculus was originally developed as part of a system for the foundations of logic [Church 1932-33]. However, the system was soon shown to be inconsistent. Church removed logical propositions from the system leaving a purely procedural  $\lambda$ -calculus [Church 1941].

The  $\lambda$ -calculus can be viewed as the earliest message passing programming language [Hewitt, Bishop, and Steiger 1973] building on previous work. For example the  $\lambda$ -expression below implements a tree data structure when supplied with parameters for a leftSubTree and rightSubTree. When such a tree is given a parameter message "*getLeft*", it returns leftSubTree and likewise when given the message "*getRight*" it returns rightSubTree:

```

λ(leftSubTree, rightSubTree)
  λ(message)
    if(message = "getLeft")
      then leftSubTree
    else if(message == "getRight")
      then rightSubTree

```

However, the semantics of the  $\lambda$ -calculus were expressed using variable substitution in which the values of parameters were substituted into the body of an invoked  $\lambda$  expression. The substitution model is unsuitable for concurrency because it does not allow the capability of sharing of changing resources. Inspired by the  $\lambda$ -calculus, the interpreter for the programming language Lisp [McCarthy *et. al.* 1962] made use of a data structure called an environment so that the values of parameters did not have to be substituted into the body of an invoked  $\lambda$ -expression. This allowed for sharing of the effects of updating shared data structures but did not provide for concurrency.

Scott and Strachey [1971] proposed to develop a mathematical semantics for programming languages based on the procedural  $\lambda$ -calculus. Milne and Strachey [1976]. continued this work attempting to reduce all computation to the  $\lambda$ -calculus. In this work, they invented a style of sequential programming that was later formalized using monads [Moggi 1989].<sup>xi</sup>

## *Petri nets*

Prior to the development of the Actor model, Petri nets<sup>xiii</sup> were widely used to model nondeterministic computation. However, they were widely acknowledged to have an important limitation: they modeled control flow but not data flow. Consequently they were not readily composable thereby limiting their modularity. Hewitt pointed out another difficulty with Petri nets: simultaneous action, *i.e.*, the atomic step of computation in Petri nets is a transition in which tokens simultaneously disappear from the input places of a transition and appear in the output places. The physical basis of using a primitive with this kind of simultaneity seemed questionable to him. Despite these apparent difficulties, Petri nets continue to be a popular approach to modeling nondeterminism, and are still the subject of active research.

## *Simula*

Simula 1 [Nygaard 1962] pioneered nondeterministic discrete event simulation using a global clock:

*In this early version of Simula a system was modeled by a (fixed) number of "stations", each with a queue of "customers". The stations were the active parts, and each was controlled by a program that could "input" a customer from the station's queue, update variables (global, local in station, and local in customer), and transfer the customer to the queue of another station. Stations could discard customers by not transferring them to another queue, and could generate new customers. They could also wait a given period (in simulated time) before starting the next action. Custom types were declared as data records, without any actions (or procedures) of their own. [Krogdahl 2003]*

Thus at each time step, the program of the next station to be simulated would update the variables.

Kristen Nygaard and Ole-Johan Dahl developed the idea (first described in an IFIP workshop in 1967) of organizing objects into “classes” with “subclasses” that could inherit methods for performing operations from their superclasses. In this way, Simula 67 considerably improved the modularity of nondeterministic discrete event simulations.

According to [Krogdahl 2003]:

Objects could act as processes that can execute in “quasi-parallel” that is in fact a form of nondeterministic sequential execution in which a simulation is organized as “independent” processes. Classes in Simula 67 have their own procedures that start when an object is generated. However, unlike Algol procedures, objects may choose to temporarily stop their execution and transfer the control to another process. If the control is later given back to the object, it will resume execution where the control last left off. A process will always retain the execution control until it explicitly gives it away. When the execution of an object reaches the end of its statements, it will become “terminated”, and can no longer be resumed (but local data and local procedures can still be accessed from outside the object).

The quasi-parallel sequencing is essential for the simulation mechanism. Roughly speaking, it works as follows: When a process has finished the actions to be performed at a certain point in simulated time, it decides when (again in simulated time) it wants the control back, and stores this in a local “next-event-time” variable. It then gives the control to a central “time-manager”, which finds the process that is to execute next (the one with the smallest next-event-time), updates the global time variable accordingly, and gives the control to that process.

The idea of this mechanism was to invite the programmer of a simulation program to model the underlying system by a set of processes, each describing some natural sequence of events in that system (e.g. the sequence of events experienced by one car in a traffic simulation).

Note that a process may transfer control to another process even if it is currently inside one or more procedure calls. Thus, each quasi-parallel process will have its own stack of procedure calls, and if it is not executing, its “reactivation point” will reside in the innermost of these calls. Quasi-parallel sequencing is analogous to the notion of co-routines [Conway 1963].

Note that a class operation operated on the global state of the simulation and not just on the local variables of the class in which the operation is defined.<sup>xiii</sup> Also Simula-67 lacked formal interfaces and instead relied on inheritance from an abstract class thereby placing limitations to the ability to invoke behavior independent of the class hierarchy.

Also, although Simula had nondeterminism, it did not have concurrency.<sup>xiv</sup>

### **Planner**

The two major paradigms for constructing semantic software systems were procedural and logical. The procedural paradigm was epitomized by using Lisp [McCarthy *et al.* 1962; Minsky, *et al.* 1968] recursive procedures operating on list structures. The logical paradigm was epitomized by uniform resolution theorem provers [Robinson 1965].

Planner [Hewitt 1969] was a kind of hybrid between the procedural and logical paradigms.<sup>xv</sup> An implication of the form (*P implies Q*) was procedurally interpreted as follows:<sup>xvi</sup>

- *when assert P, assert Q*
- *when goal Q, goal P*
- *when assert(not Q), assert(not P)*
- *when goal(not P), goal(not Q)*

Planner was the first programming language based on the pattern-directed invocation of procedural plans from assertions and goals. ***It represented a rejection of the resolution uniform proof procedure paradigm.***

### **Smalltalk-72**

Planner, Simula 67, Smalltalk-72 [Kay 1975; Ingalls 1983] and computer networks had previously used message passing. However, they were too complicated to use as the foundation for a mathematical theory of concurrency. Also they did not address fundamental issues of concurrency.

Alan Kay was influenced by message passing in the pattern-directed invocation of Planner in developing Smalltalk-71. Hewitt was intrigued by Smalltalk-71 but was put off by the complexity of communication that included invocations with many fields including **global, sender, receiver, reply-style, status, reply, operator, etc.**

In November 1972 Kay visited MIT and discussed some of his ideas for Smalltalk-72 building on the Logo work of Seymour Papert and the “little person” metaphor of computation used for teaching children to program. Smalltalk-72 made important advances in graphical user interfaces including code browsing.

However, the message passing of Smalltalk-72 was quite complex [Kay 1975]. Code in the language was viewed by the interpreter as simply a stream of tokens. As Dan Ingalls [1983] later described it:<sup>xvii</sup>

*The first (token) encountered (in a program) was looked up in the dynamic context, to determine the receiver of the subsequent message. The name lookup began with the class dictionary of the current activation. Failing there, it moved to the sender of that activation and so on up the sender chain. When a binding was finally found for the token, its value became the receiver of a new message, and the interpreter activated the code for that object's class.*<sup>xviii</sup>

Thus the message passing model in Smalltalk-72 was closely tied to a particular machine model and programming language syntax that did not lend itself to concurrency. Also, although the system was bootstrapped on itself, the language constructs were not formally defined as objects that respond to *eval* messages as in the definition of ActorScript [Hewitt 2010a].

### Concurrency

Arguably, the first concurrent programs were interrupt handlers. During the course of its normal operation, a computer needed to be able to receive information from outside (characters from a keyboard, packets from a network, *etc.*). So when the information arrived, execution of the computer was “interrupted” and special code called an interrupt handler was called to *put* the information in a buffer where it could be subsequently retrieved.

In the early 1960s, interrupts began to be used to simulate the concurrent execution of several programs on a single processor. Having concurrency with shared memory gave rise to the problem of concurrency control. Originally, this problem was conceived as being one of mutual exclusion on a single computer. Edsger Dijkstra developed semaphores and later, [Brinch Hansen 1996; Hoare 1974] developed monitors to solve the mutual exclusion problem. However, neither of these solutions provided a programming language construct that encapsulated access to shared resources [Hewitt and Atkinson 1977, 1979; Atkinson 1980].

The first models of computation (*e.g.* Turing machines, Post productions, the  $\lambda$ -calculus, *etc.*) were based on mathematics and made use of a global state to represent a computational *step* (later generalized in [McCarthy and Hayes 1969] and [Dijkstra 1976]). Each computational step was from one global state of the computation to the next global state. The global state approach was continued in automata theory for finite state machines and push down stack machines, including their nondeterministic versions.

Such nondeterministic automata have the property of bounded nondeterminism; that is, if a machine always halts when started in its initial state, then there is a bound on the number of states in which it halts.

Edsger Dijkstra further developed the nondeterministic global state approach. Dijkstra's model gave rise to a controversy concerning *unbounded nondeterminism*. Unbounded nondeterminism (also called *unbounded indeterminacy*), is a property of concurrency by which the amount of delay in servicing a request can become unbounded as a result of arbitration of contention for shared resources *while still guaranteeing that the request will eventually be serviced*. The Actor model provides the guarantee of service. In Dijkstra's model, although there could be an unbounded amount of time between the execution of sequential instructions on a computer, a (parallel) program that started out in a well-defined state could terminate in only a bounded number of states [Dijkstra 1976]. He believed that it was impossible to implement unbounded nondeterminism.

However, there is no bound that can be placed on how long it takes a computational circuit called an *arbiter* to settle. Arbiters are used in computers to deal with the circumstance that computer clocks operate asynchronously with input from outside, *e.g.* keyboard input, disk access, network input, *etc.* So it could take an unbounded time for a message sent to a computer to be received and in the meantime the computer could traverse an unbounded number of states. Thus the computer may not be in any defined stable state for an unbounded period of time [Hewitt 2006]. Therefore computers have the property of unbounded nondeterminism. So there is an inconsistency between the nondeterministic state model of computation and the circuit model of arbiters.<sup>xix</sup>

### Actors

The invention of digital computers caused a decisive paradigm shift when the notion of an interrupt was invented so that input that arrived asynchronously from outside could be incorporated in an ongoing computation. At first concurrency was conceived using low level machine implementation concepts like threads, locks, channels, cores, queues, *etc.*

The Actor Model [Hewitt, Bishop, and Steiger 1973] was based on message passing that was different from previous models of computation because the sender of a message is not intrinsic to the semantics of a communication: only the message and recipient are relevant.<sup>xx</sup> In this way, the Actor Model integrated the following:



1. the lambda calculus
2. interrupts
3. blocking method invocation
4. imperative programming using locks
5. capability systems
6. co-routines
7. packet networks
8. email systems
9. Petri nets
10. Smalltalk-72
11. Simula-67
12. pattern-directed invocation (from Planner)

The Actor model has laws that govern privacy and security [Baker and Hewitt 1977].<sup>xxi</sup>

The Actor model was a decisive break with previous models of computation because asynchronous communication cannot be implemented by Turing machines etc. because the order of arrival of messages cannot be logically inferred. Message passing is the foundation of many-core and client-cloud computing.

### Actor Semantics

Consider the following program written in CSP<sup>xxii</sup> [Hoare 1978]:

```
[X :: Z!stop( )
  ||
  Y :: guard: boolean;
      guard := true;
      *[guard →
          Z!go( );
          Z?guard]
  ||
  Z :: n: integer;
      n:= 0;
      continue: boolean;
      continue := true;
      *[X?stop( )→
          continue := false;
          [ ]
          Y?go( ) →
              n := n+1;
              Y!continue]]
```

According to Clinger [1981]:

*this program illustrates global nondeterminism, since the nondeterminism arises from incomplete specification of the timing of signals between the three processes X, Y, and Z. The repetitive guarded command in the definition of Z has two alternatives: either the stop message is accepted from X, in which case continue is set to false, or a go message is accepted from Y, in which case n is*

*incremented and Y is sent the value of continue. If Z ever accepts the stop message from X, then X terminates. Accepting the stop causes continue to be set to false, so after Y sends its next go message, Y will receive false as the value of its guard and will terminate. When both X and Y have terminated, Z terminates because it no longer has live processes providing input.*

*As the author of CSP points out, therefore, if the repetitive guarded command in the definition of Z were required to be fair, this program would have unbounded nondeterminism: it would be guaranteed to halt but there would be no bound on the final value of n.<sup>xxiii</sup> In actual fact, the repetitive guarded commands of CSP are not required to be fair, and so the program may not halt [Hoare 1978]. This fact may be confirmed by a tedious calculation using the semantics of CSP [Francez, Hoare, Lehmann, and de Roever 1979] or simply by noting that the semantics of CSP is based upon a conventional power domain and thus does not give rise to unbounded nondeterminism.*

However, the following concurrent program in ActorScript [Hewitt 2010a] will return an integer of unbounded size:

Unbounded ===

```
behavior{
  ① methods
  start() → ① a start message is implemented by
              let(Counter::c = create SimpleCounter(n=0)),
              ① let c be a Counter that is a created by
              ① SimpleCounter with count equal 0
              {c.go(), ① send c a go message and concurrently
               c.stop(),} ① return the value of sending c a stop
```

SimpleCounter ===

```
behavior implements Counter {
  Integer::n , ① n is the current count
  | ① methods for messages are below
  stop() → n,} ① stop returns count n
  this.go() → ① go message to self does
              exit(this.go(),) alsobecome (n=n+1)
              ① exit sending self a go message concurrently
              ① incrementing the count n
```

By the semantics of the Actor model of computation [Clinger 1981; Hewitt 2006], the result of sending Unbounded a start message is an integer of unbounded size.

The model of CSP deliberately had bounded nondeterminism [Francez, Hoare, Lehmann, and de Roever 1979] whereas the Actor model has unbounded nondeterminism. Influence by Dijkstra, Hoare had been

convinced that a programming language with unbounded nondeterminism could not be implemented. Consequently it was not possible to guarantee that servers implemented using CSP would provide service to multiple clients.<sup>xxiv</sup>

In 1975, Irene Greif published the first operational model of Actors in her dissertation. Two years after Greif published her operational model, Carl Hewitt and Henry Baker published the Laws for Actors [Baker and Hewitt 1977]. Other major milestones include [Clinger 1981] introducing a denotational semantics based on power domains, and [Agha 1985] which further developed a complementary transition-based semantic model.

### **Logic Programming<sup>xxv</sup>**

Robert Kowalski developed the thesis that “*computation could be subsumed by deduction*” [Kowalski 1988a] that he states was first proposed by Hayes [1973] in the form “*Computation = controlled deduction.*”<sup>xxvi</sup> [Kowalski 1979] Kowalski forcefully stated:

*There is only one language suitable for representing information -- whether declarative or procedural -- and that is first-order predicate logic. There is only one intelligent way to process information -- and that is by applying deductive inference methods.* [Kowalski 1980]

The gauntlet was officially thrown in *The Challenge of Open Systems* [Hewitt 1985] to which [Kowalski 1988b] replied in *Logic-Based Open Systems*. This was followed up with [Hewitt and Agha 1988] in the context of the Japanese Fifth Generation Project.

According to Hewitt, *et. al.* and contrary to Kowalski and Hayes, computation in general cannot be subsumed by deduction and contrary to the quotation (above) attributed to Hayes computation in general is not controlled deduction. Hewitt and Agha [1991] and other published work argued that mathematical models of concurrency did not determine particular concurrent computations because they make use of arbitration for determining which message is next in the arrival order when multiple messages concurrently. For example Arbiters can be used in the implementation of the arrival order. Since arrival orders are in general indeterminate, they cannot be deduced from prior information by mathematical logic alone. Therefore mathematical logic cannot implement concurrent computation in open systems.

In concrete terms, typically we cannot observe the details by which the arrival order of messages determined. Attempting to do so affects the results and can even push the indeterminacy elsewhere. Instead of observing the internals of arbitration processes, we await outcomes. The

reason that we await outcomes is that we have no alternative because of indeterminacy.<sup>xxvii</sup>

### **Reasoning about Actors**

The principle of Actor induction is:

1. Suppose that an Actor  $x$  has property  $P$  when it is created
2. Further suppose that if  $x$  has property  $P$  when it processes a message, then it has property  $P$  when it processes the next message.
3. Then  $x$  always has the property  $P$ .

In his doctoral dissertation, Aki Yonezawa developed further techniques for proving properties of Actor systems including those that make use of migration. Russ Atkinson developed techniques for proving properties of Actors that are guardians of shared resources. Gerry Barber's 1981 doctoral dissertation concerned reasoning about change in knowledgeable office systems.

### **Early Actor Programming languages**

Henry Lieberman, Dan Theriault, *et al.* developed Act1, an Actor programming language. Subsequently for his master's thesis, Dan Theriault developed Act2. These early proof of concept languages were rather inefficient and not suitable for applications. In his doctoral dissertation, Ken Kahn developed Ani, which he used to develop several animations. Bill Kornfeld developed the Ether programming language for the Scientific Community Metaphor in his doctoral dissertation. William Athas and Nanette Boden [1988] developed Cantor which is an Actor programming language for scientific computing. Jean-Pierre Briot [1988, 1999] developed means to extend Smalltalk 80 for Actor computations. Darrell Woelk [1995] at MCC developed an Actor programming language for InfoSleuth agents in Rosette.

Hewitt, Attardi, and Lieberman [1979] developed proposals for delegation in message passing. This gave rise to the so-called inheritance anomaly controversy in object-oriented concurrent programming languages [Satoshi Matsuoka and Aki Yonezawa 1993, Giuseppe Milicia and Vladimiro Sassone 2004].

### **Causal Consistency**

Some kinds of sequential consistency<sup>xxviii</sup> are special cases of *causal consistency* which is an expression of the laws for Actors [Hewitt and Baker 1977].

Consider the following implementation of a cell: in ActorScript [Hewitt 2010a]:<sup>xxix</sup>

### Implementation of a Cell:

Below SimpleCell is defined that implements Cell

```
SimpleCell===  
behavior{  
  ?::contents  
  ⓘ with instance variable contents  
| implements Cell  
  read→ ⓘ read message does  
          contents, ⓘ return contents  
  write( ?::nextContents) →  
    void alsoBecome (contents=nextContents)}  
CreateCell( ?::x ) ===  
  create SimpleCell (contents=x)
```

### Illustration of use of a Cell:

The following expression creates a cell x with initial contents 5 and then concurrently writes to it with the values 7 and 9 while concurrently reading the contents.

```
let(Cell::x = CreateCell(5))
```

```
{x.write(7), x.write(9), x.read}
```

ⓘ “,” above means discretionary concurrency

The value of the above expression is 5, 7 or 9.

On the other hand sequential evaluation proceeds as follows:

```
let(Cell::x = CreateCell(5))
```

```
{x.write(7); x.write(9); x.read}
```

The value of the above expression is 9.

Of course the following expression:

```
let(Cell::x = CreateCell(5))
```

```
{x.write(7), x.write(9); x.read}
```

ⓘ “,” above means discretionary concurrency evaluates to 7 or 9.

### First Illustration of causal consistency

```
let(Cell::x = CreateCell(0),  
    Cell::y = CreateCell(0))
```

```
list[{x.write(1); y.read},
```

ⓘ “,” above means discretionary concurrency

```
{y.write(1); x.read}]
```

The above expression returns either `list[1 0]` or `list[0, 1]`. (These results can be calculated using the Computational Representation Theorem.)

It is impossible for the expression to return `list[0 0]` for the following reason: When y receives the `read` message there are two possibilities:

1. It has sent an acknowledgment for the `write(1)` message, in which case it must return 1 for the `read` message.
2. It has not sent an acknowledgment for a `write` message and so can return 0 for the `read` message. However, this means that x must have previously sent an acknowledgment for the `write(1)` message and so when x subsequently receives a `read` message it must return 1.

### Second Illustration of causal consistency

```
let(Cell::x = CreateCell(0),  
    Cell::y = CreateCell(0))
```

```
[{x.read ?? {1 : y.write(1); 0 : void};  
 x.read},
```

ⓘ “,” above means discretionary concurrency

```
{y.read ?? {1 : x.write(1); 0 : void};  
 y.read}]
```

The above expression returns `[0 0]` because it is impossible to read a value from a cell that was not the initial value and never written to the cell.

### Third Illustration of causal consistency

```
let (Cell:x = CreateCell(0),
    Cell:y = CreateCell(0))
{x.write(1),
  ① “,” above means discretionary concurrency
  [let(?:x1=x.read); {y.write(1), x1},
  ① “,” above means discretionary concurrency
  [y.read; x.read]]}
  ① “;” means x is sent read after y responds
```

The above expression can never return [1 [1 0]] because if x returned 0 from the second listed *read* message then y must have previously acknowledged the *write* (1) message which happened after the first listed x *read* message returned 1 which must have happened after x acknowledged the *write* (1) message.

### Locking illustration:

Below is an (overly) simple implementation of a spin lock. More sophisticated implementations must be used in practical applications.

```
SimpleSpinLock(Atomic:x, ?:value)===
x.conditionalWrite(value) ?? {
  ① conditionalWrite message can be
  ① implemented as an atomic instruction to
  ① perform store in x if it differs from value
  false:
    ① false that x was updated
    ① in a more practical implementation,
    ① there would be backoff here
  SimpleSpinLock(x, value); ① retry
  true: void} ① true that x was atomically
    ① updated with value
```

### Atomic operation illustration:

Below is an illustration of an atomic operation to swap the contents of two cells guarded by a lock where the contents of the lock are 0 if unlocked and 1 if locked:

```
Swap(Cell:x, Cell:y, Atomic:lock) ===
{SimpleSpinLock(lock,1);
  ① acquire the lock
  let(?:temp=x.read)
  {x.write(y.read);
  y.write(temp)};
  ① swap the contents
  lock.write(0);} ① release the lock
```

### Extension versus Specialization

Programming languages like ActorScript [Hewitt 2010] take the approach of extending behavior in contrast to the approach of specializing behavior:

- Type specialization: If type t1 is a subtype of type t2, then instances of t1 have all of the properties that are provable from the definition of type t2 [Liskov 1987, Liskov and Wing 2001].
- Type extension: A type can be extended to have additional (perhaps incompatible) properties from the type that it extends. An extension type can make use of the implementation of the type that it extends. Type extension is commonly used to extend the behavior of operating system software as well as applications.

The term “inheritance” in object-oriented programming has been used (sometimes ambiguously) to mean both specialization and extension.

### Early Actor Programming languages

Henry Lieberman, Dan Theriault, *et al.* developed Act1, an Actor programming language. Subsequently for his master’s thesis, Dan Theriault developed Act2. These early proof of concept languages were rather inefficient and not suitable for applications. In his doctoral dissertation, Ken Kahn developed Ani, which he used to develop several animations. Bill Kornfeld developed the Ether programming language for the Scientific Community Metaphor in his doctoral dissertation. William Athas and Nanette Boden [1988] developed Cantor which is an Actor programming language for scientific computing. Jean-Pierre Briot [1988, 1999] developed means to extend Smalltalk 80 for Actor computations. Darrell Woelk [1995] at MCC developed an Actor programming language for InfoSleuth agents in Rosette.

Hewitt, Attardi, and Lieberman [1979] developed proposals for delegation in message passing. This gave rise to the so-called inheritance anomaly controversy in object-oriented concurrent programming languages [Satoshi

Matsuoka and Aki Yonezawa 1993, Giuseppe Milicia and Vladimiro Sassone 2004].

### Language constructs versus Library APIs

Library Application Programming Interfaces (APIs) are an alternative way to introduce concurrency. For example,

- Futures in C++0x [Stroustrup *et. al.* 2008] provided additional capability for C++. Unfortunately, the C++0x addition is verbose, inefficient and overly rigid.
- Message Passing Interface (MPI) [Gropp *et. al.* 1998] provides message passing capabilities. Unfortunately, MPI is also verbose, inefficient and overly rigid.
- GCD queues were an important advance in commercial software. (See an implementation below.) However, in general, queues are plumbing that are better placed inside abstractions rather than being exposed externally.

In general, appropriately defined language constructs provide greater power, flexibility, and efficiency than library APIs.<sup>xxx</sup>

#### Implementation GCD queues:

Below is an implementation of GCD queues in ActorScript [Hewitt 2010a] using the Swiss Cheese model [Hewitt and Atkinson 1977, 1979; Atkinson 1980] (holes in the cheese are highlighted in grey):

```
SimpleGCDqueue ===
behavior queues{q} {
| implements GCDqueue
dispatch_sync(theBlock) →
  passThru(q) {hole (theBlock();)}
  ① Semicolon above means hole is not
  ① exited until theBlock responds.
  responded(r) {
    void alsoDequeue(q)};
this.dispatch_async(theBlock) →
  {this ← passThru(q) {
  ① ← means that this is pushed the value
  ① computed by the following expression
    hole (theBlock();)}
  ① Semicolon above means that the hole
  ① is not exited until theBlock responds
    responded r {
      void alsoDequeue(q)}},
  ① Comma above means that above is
  ① is performed concurrently with below
void}} ① dispatch_async method
  ① immediately returns void
```

### Hairy Control Structure Redux

McDermott and Sussman in [1972] developed the Lisp-based language Conniver based on “hairy control structure” (first introduced in [Landin 1965]) that could implement non-chronological backtracking that was more general than the chronological backtracking in Planner. Hewitt and others were skeptical of hairy control structure. Pat Hayes [1974] remarked:

*Their [Sussman and McDermott] solution, to give the user access to the implementation primitives of Planner, is however, something of a retrograde step (what are Conniver's semantics?).*

In the 1960's at the MIT AI Lab a remarkable culture grew up around “hacking” that concentrated on remarkable feats of programming.<sup>xxx1</sup> Growing out of this tradition, Gerry Sussman and Guy Steele decided to try to understand Actors by reducing them to machine code that they could understand and so developed a “Lisp-like language, Scheme, based on the lambda calculus, but extended for side effects, multiprocessing, and process synchronization.” [Sussman and Steele 1975] Unfortunately, their reductionist approach included primitives like START!PROCESS, STOP!PROCESS and EVALUATE!UNINTERRUPTIBLEY that had the following explanation.<sup>xxxii</sup>

*This is the synchronization primitive. It evaluates an expression uninterruptedly; i.e. no other process may run until the expression has returned a value.*

Of course, the above reductionist approach is very unsatisfactory because it missed a crucial aspect of the Actor model: *the arrival ordering of messages.*

McDermott, and Sussman [1972] developed the Lisp-based language Conniver based on “hairy control structure” that could implement non-chronological backtracking that was more general than the chronological backtracking in Planner. However, Hewitt and others were skeptical of hairy control structure. Pat Hayes remarked:

*Their [Sussman and McDermott] solution, to give the user access to the implementation primitives of Planner, is however, something of a retrograde step (what are Conniver's semantics?). [Hayes 1974]*

Hewitt had concluded:

*One of the most important results that has emerged from the development of Actor semantics has been the further development of techniques to semantically analyze or synthesize control structures as patterns of passing messages. As a result of this work, we have found that we can do without the paraphernalia of “hairy control structure.”<sup>xxxiii</sup> (emphasis in original)*

Sussman and Steele [1975] noticed some similarities between Actor programs and the  $\lambda$ -calculus. They

mistakenly concluded that they had reduced Actor programs to a “continuation-passing programming style”:

*It is always possible, if we are willing to specify explicitly what to do with the answer, to perform any calculation in this way: rather than reducing to its value, it reduces to an application of a continuation to its value. That is, in this continuation-passing programming style, a function always “returns” its result by “sending” it to another function.* (emphasis in original)

However, some Actor programming language constructs are not reducible to a continuation-passing style. For example, push concurrency (i.e. ... $\leftarrow$ ...) is not reducible to continuation-passing style.<sup>xxxiv</sup>

[Reynolds 1972] introduced hairy control structure continuations using a primitive called *escape*. Sussman and Steele called their variant of *escape* by the name “*call with current continuation*.”<sup>xxxv</sup> Unfortunately, general use of *escape* is not compatible with usual hardware stack discipline introducing considerable operational inefficiency. Also, using *escape* can leave customers stranded.<sup>xxxvi</sup> Consequently, use of *escape* is generally avoided these days and exceptions<sup>xxxvii</sup> are used instead so that clean up can be performed.

On the basis of their experience, Sussman and Steele developed the general thesis that Actors were merely the  $\lambda$ -calculus in disguise. Steele [1976] in the section “Actors  $\equiv$  Closures (mod Syntax)” disagreed with Hewitt who had “*expressed doubt as to whether these underlying continuations can themselves be expressed as lambda expressions.*” However, customers cannot be expressed as  $\lambda$ -expressions because doing so would preclude being able to enforce the requirement that a customer will process at most one response (i.e. exception or value return). Also implementing customers as  $\lambda$  expressions can leave customers stranded.

In summary, Sussman and Steele [1975] mistakenly concluded “*we discovered that the ‘Actors’ and the lambda expressions were identical in implementation.*”<sup>xxxviii</sup> The actual situation is that the  $\lambda$ -calculus is capable of expressing some kinds of sequential and parallel control structures but, in general, *not* the concurrency expressed in the Actor model.<sup>xxxix</sup> On the other hand, the Actor model is capable of expressing everything in the  $\lambda$ -calculus and more [Hewitt 2008f].

### **Promises**

Promises [Liskov and Shriram 1988] are a lower level version of futures [Baker and Hewitt 1977] in the following respects:

1. **Direct use:** Futures can be used directly<sup>xl</sup> whereas promises must first be resolved<sup>xli</sup> to extract the outcome of a promise before it can be used.
2. **Derivatives:** According to [Miller 2006]: “*Messages sent to [a] promise cannot be delivered until the promise is resolved, so they are buffered in FIFO order within the promise.*” Derivative futures are different from promises in that a deferrable message sent to a derivative future results in the creation of a *new* future for the outcome of the message.
3. **Resource control:** Unlike futures, promises do not make use of sponsors for resource control.

### **Garbage Collection**

Garbage collection (the automatic reclamation of unused storage) was an important theme in the development of the Actor model.

In his doctoral dissertation, Peter Bishop developed an algorithm for garbage collection in distributed systems. Each system kept lists of links of pointers to and from other systems. Cyclic structures were collected by incrementally migrating Actors (objects) onto other systems which had their addresses until a cyclic structure was entirely contained in a single system where the garbage collector could recover the storage.

Henry Baker developed an algorithm for real-time garbage collection in his doctoral dissertation. The fundamental idea was to interleave collection activity with construction activity so that there would not have to be long pauses while collection takes place.

Lieberman and Hewitt [1983] developed a real time garbage collection based on the lifetimes of Actors (Objects). The fundamental idea was to allocate Actors (objects) in generations so that only the latest generations would have to be examined during a garbage collection.

### **Cosmic Cube**

The Cosmic Cube was developed by Chuck Seitz *et al.* at Caltech providing architectural support for Actor systems. A significant difference between the Cosmic Cube and most other parallel processors is that this multiple instruction multiple-data machine used message passing instead of shared variables for communication between concurrent processes. This computational model was reflected in the hardware structure and operating system, and also the explicit message passing communication seen by the programmer.

## J-Machine

The J-Machine was developed by Bill Dally *et al.* at MIT providing architectural support suitable for Actors. This included the following:

- Asynchronous messaging
- A uniform space of Actor addresses to which messages could be sent concurrently regardless of whether the recipient Actor was local or nonlocal
- A form of Actor pipelining

Concurrent Smalltalk (which can be modeled using Actors) was developed to program the J Machine.

## $\pi$ -Calculus

Robin Milner's initial published work on concurrency [Milner 1973] was notable in that it positions mathematical semantics of communicating processes as a framework to understand a variety of interaction agents including the computer's interaction with memory. The framework of modeling was based on Scott's model of domains and as such was not based on sequential processes. His work differed from the independently developed Actor model in the following ways:

- There are a fixed number of processes as opposed to the Actor model which allows the number of Actors to vary dynamically
- The only quantities that can be passed in messages are integers and strings as opposed to the Actor model which allows the addresses of Actors to be passed in messages
- The processes have a fixed topology as opposed to the Actor model which allows varying topology
- Communication is synchronous as opposed to the Actor model in which an unbounded time can elapse between sending and receiving a message.
- Unlike the Actor model, there is no arrival ordering and consequently there is only bounded nondeterminism. However, with bounded nondeterminism is impossible for a server to guarantee service to its clients, *i.e.*, a client might starve.

Building on the Actor model, Milner [1993] removed some of these restrictions in his work on the  $\pi$ -calculus:

*Now, the pure lambda-calculus is built with just two kinds of thing: terms and variables. Can we achieve the same economy for a process calculus? Carl Hewitt, with his Actors model, responded to this challenge long ago; he declared that a value, an operator on values, and a process should all be the same kind of thing: an Actor.*

*This goal impressed me, because it implies the homogeneity and completeness of expression ... But it was long before I could see how to attain the goal in terms of an algebraic calculus...*

*So, in the spirit of Hewitt, our first step is to demand that all things denoted by terms or accessed by names--values, registers, operators, processes, objects--are all of the same kind of thing; they should all be processes. Thereafter we regard access-by-name as the raw material of computation...*

However, some fundamental differences remain between the Actor model and the  $\pi$ -calculus.

- The Actor model is founded on physics whereas the  $\pi$ -calculus is founded on algebra.
- Semantics of the Actor model is based on message orderings in the Computational Representation Theorem. Semantics of the  $\pi$ -calculus is based on structural congruence in various kinds of bisimulations and equivalences.

Communication in the  $\pi$ -calculus takes the following form:

- *input:*  $u(x).P$  is a process that gets a message from a communication channel  $u$  before proceeding as  $P$ , binding the message received to the identifier  $x$ . In ActorScript [Hewitt 2010a], this can be modeled as follows:  $\{let\ x=u.get(); P\}$ <sup>xlii</sup>
- *output:*  $\bar{u}(m).P$  is a process that puts a message  $m$  on communication channel  $u$  before proceeding as  $P$ . In ActorScript, this can be modeled as follows:  $\{u.put(m); P\}$ <sup>xliii</sup>

The above operations of the  $\pi$ -calculus can be implemented in Actor systems using a two-phase commit protocol [Knabe 1992; Reppy, Russo, and Xiao 2009]. The overhead of communication in the  $\pi$ -calculus presents tremendous difficulties to its use in practical applications.

## Was the Actor model premature?

The history of the Actor model raises the question of whether it was premature.

### Original definition of prematurity

As originally defined by Gunther Stent [1972], "A discovery is premature if its implications cannot be connected by a series of simple logical steps to contemporary canonical or generally accepted knowledge." Ilana Lövy 2002 glossed the phrase "series of simple logical steps" in Stent's definition as referring to the "target community's ways of asking relevant questions, of producing experimental results, and of examining new evidence." Michael Ghiselin [2002] argued that if a "minority of scientists accept a discovery, or even pay serious attention to it, then the discovery is not altogether premature in the Stentian sense." In accord with Ghiselin's argument, the Actor model was not premature. Indeed it

enjoyed initial popularity and for a couple of decades underwent steady development.

However, Stent in his original article also referred to a development as premature such that when it occurred contemporaries did not seem to be able to do much with or build on. This is what happened after a while with the Actor model. The reasons were twofold:

1. For over 30 years after the first publication of the Actor model, widely deployed computer architectures developed in the direction of making a single sequential thread of execution run faster.
2. For over 25 years after the first publication, there was no agreed standard by which software could communicate high level data structures across organizational boundaries.

### *Before its time?*

According to Elihu M. Gerson [2002], phenomena that lead people to talk about discoveries being before their time can be analyzed as follows:

*We can see the phenomenon of 'before its time' as composed of two separate steps. The first takes place when a new discovery does not get tied to the conventional knowledge of its day and remains unconnected in the literature. The second step occurs when new events lead to the 'rediscovery' of the unconnected results in a changed context that enables or even facilitates its connection to the conventional knowledge of the rediscovering context.*

Now both of the above circumstances that held back the Actor model have changed with the development of

- (1) many-core computer architectures a
- (2) client-cloud computing

By the criteria of Gerson, the Actor model might be described by some as *before its time*.

According to Hadasa Zuckerman and Joshua Lederberg [1986], premature discoveries are those that were made but neglected. By their criteria it remains to be seen whether or not the Actor model was premature.

Gerson 2002 argued,

*But histories and sociological studies repeatedly show that we do not have a discovery until the scientific community accepts it as such and stops debating about it. Until then the proposed solution is in an intermediate state."*

By his argument, the Actor model is a discovery but since its practical importance is not yet accepted by the community, its practical importance is not yet a discovery.

## **Acknowledgment**

Important contributions to the semantics of Actors have been made by: Gul Agha, Beppe Attardi, Henry Baker, Will Clinger, Irene Greif, Carl Manning, Ian Mason, Ugo Montanari, Maria Simi, Scott Smith, Carolyn Talcott, Prasanna Thati, and Aki Yonezawa.

Important contributions to the implementation of Actors have been made by: Bill Athas, Russ Atkinson, Beppe Attardi, Henry Baker, Gerry Barber, Peter Bishop, Nanette Boden, Jean-Pierre Briot, Bill Dally, Peter de Jong, Jessie Dedecker, Ken Kahn, Henry Lieberman, Carl Manning, Mark S. Miller, Tom Reinhardt, Chuck Seitz, Richard Steiger, Dan Theriault, Mario Tokoro, Darrell Woelk, and Carlos Varela.

Research on the Actor model has been carried out at Caltech Computer Science, Kyoto University Tokoro Laboratory, MCC, MIT Artificial Intelligence Laboratory, SRI, Stanford University, University of Illinois at Urbana-Champaign Open Systems Laboratory, Pierre and Marie Curie University (University of Paris 6), University of Pisa, University of Tokyo Yonezawa Laboratory and elsewhere.

The members of the Silicon Valley Friday AM group made valuable suggestions for improving this paper. Blaine Garst found numerous bugs and made valuable comments including information on the historical development of interfaces. Patrick Beard found bugs and suggested improvements in presentation.

The Actor model is intended to provide a foundation for information integration in privacy-friendly client-cloud computing [Hewitt 2009b].

## **Bibliography**

- Hal Abelson and Gerry Sussman *Structure and Interpretation of Computer Programs* 1984.
- Sarita Adve and Hans-J. Boehm. *Memory Models: A Case for Rethinking Parallel Languages and Hardware* CACM. July 2010.
- Gul Agha. *Actors: A Model of Concurrent Computation in Distributed Systems* Doctoral Dissertation. 1986.
- Gul Agha, Ian Mason, Scott Smith, and Carolyn Talcott. "A foundation for Actor computation." *Journal of Functional Programming*. 1997.
- Joe Armstrong *History of Erlang* HOPL III. 2007.
- Joe Armstrong. *Erlang*. CACM. September 2010.
- William Athas and Charles Seitz *Multicomputers: message-passing concurrent computers* IEEE Computer August 1988.
- William Athas and Nanette Boden *Cantor: An Actor Programming System for Scientific Computing* in



- Proceedings of the NSF Workshop on Object-Based Concurrent Programming. 1988. Special Issue of SIGPLAN Notices.
- Russ Atkinson. *Automatic Verification of Serializers* MIT Doctoral Dissertation. June, 1980.
- Henry Baker. *Actor Systems for Real-Time Computation* MIT EECS Doctoral Dissertation. January 1978.
- Henry Baker and Carl Hewitt *The Incremental Garbage Collection of Processes* Proceeding of the Symposium on Artificial Intelligence Programming Languages. SIGPLAN Notices 12, August 1977.
- Paul Baran. *On Distributed Communications Networks* IEEE Transactions on Communications Systems. March 1964.
- Gerry Barber. *Reasoning about Change in Knowledgeable Office Systems* MIT EECS Doctoral Dissertation. August 1981.
- Peter Bishop *Very Large Address Space Modularly Extensible Computer Systems* MIT EECS Doctoral Dissertation. June 1977.
- Andreas Blass, Yuri Gurevich, Dean Rosenzweig, and Benjamin Rossman (2007a) *Interactive small-step algorithms I: Axiomatization* Logical Methods in Computer Science. 2007.
- Andreas Blass, Yuri Gurevich, Dean Rosenzweig, and Benjamin Rossman (2007b) *Interactive small-step algorithms II: Abstract state machines and the characterization theorem*. Logical Methods in Computer Science. 2007.
- Per Brinch Hansen *Monitors and Concurrent Pascal: A Personal History* CACM 1996.
- Don Box, David Ehnebuske, Gopal Kakivaya, Andrew Layman, Noah Mendelsohn, Henrik Nielsen, Satish Thatte, Dave Winer. *Simple Object Access Protocol (SOAP) 1.1* W3C Note. May 2000.
- Jean-Pierre Briot. *Acttalk: A framework for object-oriented concurrent programming-design and experience* 2nd France-Japan workshop. 1999.
- Jean-Pierre Briot. *From objects to Actors: Study of a limited symbiosis in Smalltalk-80* Rapport de Recherche 88-58, RXF-LITP. Paris, France. September 1988.
- Luca Cardelli, James Donahue, Lucille Glassman, Mick Jordan, Bill Kalsow, Greg Nelson. *Modula-3 report (revised)* DEC Systems Research Center Research Report 52. November 1989.
- Luca Cardelli and Andrew Gordon *Mobile Ambients* FoSSaCS'98.
- Arnaud Carayol, Daniel Hirschhoff, and Davide Sangiorgi. *On the representation of McCarthy's amb in the  $\pi$ -calculus* "Theoretical Computer Science" February 2005.
- Alonzo Church "A Set of postulates for the foundation of logic (1&2)" *Annals of Mathematics*. Vol. 33, 1932. Vol. 34, 1933.
- Alonzo Church *The Calculi of Lambda-Conversion* Princeton University Press. 1941.
- Will Clinger. *Foundations of Actor Semantics* MIT Mathematics Doctoral Dissertation. June 1981.
- Tyler Close *Web-key: Mashing with Permission* WWW'08.
- Melvin Conway. *Design of a separable transition-diagram compiler* CACM. 1963.
- Ole-Johan Dahl and Kristen Nygaard. "Class and subclass declarations" *IFIP TC2 Conference on Simulation Programming Languages*. May 1967.
- Ole-Johan Dahl and Tony Hoare. *Hierarchical Program Structures* in "Structured Programming" Prentice Hall. 1972.
- William Dally and Wills, D. *Universal mechanisms for concurrency* PARLE '89.
- William Dally, et al. *The Message-Driven Processor: A Multicomputer Processing Node with Efficient Mechanisms* IEEE Micro. April 1992.
- Jack Dennis and Earl Van Horn. *Programming Semantics for Multiprogrammed Computations* CACM. March 1966.
- A. Fine *The Shaky Game: Einstein Realism and the Quantum Theory* University of Chicago Press, Chicago, 1986.
- Frederic Fitch. *Symbolic Logic: an Introduction*. Ronald Press. 1952
- Nissim Francez, Tony Hoare, Daniel Lehmann, and Willem-Paul de Roever. *Semantics of nondeterminism, concurrency, and communication* Journal of Computer and System Sciences. December 1979.
- Christopher Fuchs *Quantum mechanics as quantum information (and only a little more)* in A. Khrenikov (ed.) *Quantum Theory: Reconstruction of Foundations* (Vaxjo: Vaxjo University Press, 2002).
- Blaine Garst. *Origin of Interfaces* Email to Carl Hewitt on October 2, 2009.
- Elihu M. Gerson. *Prematurity and Social Worlds in Prematurity in Scientific Discovery*. University of California Press. 2002.
- Andreas Glausch and Wolfgang Reisig. *Distributed Abstract State Machines and Their Expressive Power* Informatik Berichete 196. Humboldt University of Berlin. January 2006.
- Adele Goldberg and Alan Kay (ed.) *Smalltalk-72 Instruction Manual* SSL 76-6. Xerox PARC. March 1976.
- Dina Goldin and Peter Wegner. *The Interactive Nature of Computing: Refuting the Strong Church-Turing Thesis* Minds and Machines March 2008.
- Cordell Green. *Application of Theorem Proving to Problem Solving* IJCAI'69.
- Irene Greif and Carl Hewitt. *Actor Semantics of PLANNER-73* Conference Record of ACM Symposium on Principles of Programming Languages. January 1975.
- Irene Greif. *Semantics of Communicating Parallel Professes* MIT EECS Doctoral Dissertation. August 1975.

- Pat Hayes *Some Problems and Non-Problems in Representation Theory* AISB. Sussex. July, 1974
- Werner Heisenberg. *Physics and Beyond: Encounters and Conversations* translated by A. J. Pomerans (Harper & Row, New York, 1971), pp. 63 – 64.
- Carl Hewitt, Peter Bishop and Richard Steiger. *A Universal Modular Actor Formalism for Artificial Intelligence* IJCAI'73.
- Carl Hewitt, et al. *Actor Induction and Meta-evaluation* Conference Record of ACM Symposium on Principles of Programming Languages, January 1974.
- Carl Hewitt and Henry Lieberman. *Design Issues in Parallel Architecture for Artificial Intelligence* MIT AI memo 750. Nov. 1983.
- Carl Hewitt, Tom Reinhardt, Gul Agha, and Giuseppe Attardi *Linguistic Support of Receptionists for Shared Resources* MIT AI Memo 781. Sept. 1984.
- Carl Hewitt, et al. *Behavioral Semantics of Nonrecursive Control Structure* Proceedings of Colloque sur la Programmation, April 1974.
- Carl Hewitt. *How to Use What You Know* IJCAI. September, 1975.
- Carl Hewitt. *Viewing Control Structures as Patterns of Passing Messages* AI Memo 410. December 1976. Journal of Artificial Intelligence. June 1977.
- Carl Hewitt and Henry Baker *Laws for Communicating Parallel Processes* IFIP-77, August 1977.
- Carl Hewitt and Russ Atkinson. *Specification and Proof Techniques for Serializers* IEEE Journal on Software Engineering. January 1979.
- Carl Hewitt, Beppe Attardi, and Henry Lieberman. *Delegation in Message Passing* Proceedings of First International Conference on Distributed Systems Huntsville, AL. October 1979.
- Carl Hewitt and Gul Agha. *Guarded Horn clause languages: are they deductive and Logical?* in Artificial Intelligence at MIT, Vol. 2. MIT Press 1991.
- Carl Hewitt and Jeff Inman. *DAI Betwixt and Between: From "Intelligent Agents" to Open Systems Science* IEEE Transactions on Systems, Man, and Cybernetics. Nov./Dec. 1991.
- Carl Hewitt and Peter de Jong. *Analyzing the Roles of Descriptions and Actions in Open Systems* Proceedings of the National Conference on Artificial Intelligence. August 1983.
- Carl Hewitt. (2006). "What is Commitment? Physical, Organizational, and Social" *COIN@AAMAS'06*. (Revised version to be published in Springer Verlag Lecture Notes in Artificial Intelligence. Edited by Javier Vázquez-Salceda and Pablo Noriega. 2007) April 2006.
- Carl Hewitt (2007a). "Organizational Computing Requires Unstratified Paraconsistency and Reflection" *COIN@AAMAS*. 2007.
- Carl Hewitt (2008a) [Norms and Commitment for iOrgs™ Information Systems: Direct Logic™ and Participatory Argument Checking](#) ArXiv 0906.2756.
- Carl Hewitt (2008b) "Large-scale Organizational Computing requires Unstratified Reflection and Strong Paraconsistency" *Coordination, Organizations, Institutions, and Norms in Agent Systems III* Jaime Sichman, Pablo Noriega, Julian Padget and Sascha Ossowski (ed.). Springer-Verlag. <http://organizational.carlhewitt.info/>
- Carl Hewitt (2008c) [Middle History of Logic Programming: Resolution, Planner, Edinburgh Logic for Computable Functions, Prolog and the Japanese Fifth Generation Project](#) ArXiv 0904.3036.
- Carl Hewitt (2008e). *ORGs for Scalable, Robust, Privacy-Friendly Client Cloud Computing* IEEE Internet Computing September/October 2008.
- Carl Hewitt (2008f) [Common sense for concurrency and inconsistency robustness using Direct Logic™ and the Actor Model](#) ArXiv 0812.4852.
- Carl Hewitt (2009a) *Perfect Disruption: The Paradigm Shift from Mental Agents to ORGs* IEEE Internet Computing. Jan/Feb 2009.
- Carl Hewitt (2009b) [A historical perspective on developing foundations for client-cloud computing: iConsult™ & iEntertain™ Apps using iInfo™ Information Integration for iOrgs™ Information Systems](#) (Revised version of "Development of Logic Programming: What went wrong, What was done about it, and What it might mean for the future" AAAI Workshop on What Went Wrong. AAAI-08.) ArXiv 0901.4934.
- Carl Hewitt (2009c) [Middle History of Logic Programming: Resolution, Planner, Prolog and the Japanese Fifth Generation Project](#) ArXiv 0904.3036
- Carl Hewitt (2010a) [ActorScript™ extension of C#®, Java® and Objective C®; iAdaptive™ concurrency for antiCloud™-privacy and security](#) ArXiv 1008.2748
- Tony Hoare *Quick sort* Computer Journal 5 (1) 1962.
- Tony Hoare *Monitors: An Operating System Structuring Concept* CACM. October 1974.
- Tony Hoare. *Communicating sequential processes* CACM. August 1978.
- Tony Hoare. *Communicating Sequential Processes* Prentice Hall. 1985.
- W. Horwat, Andrew Chien, and William Dally. *Experience with CST: Programming and Implementation* PLDI. 1989.
- Daniel Ingalls. *Design Principles Behind Smalltalk*. Byte. August 1981.
- M. Jammer *The EPR Problem in Its Historical Development* in Symposium on the Foundations of Modern Physics: 50 years of the Einstein-Podolsky-Rosen Gedankenexperiment, edited by P. Lahti and P. Mittelstaedt. World Scientific. Singapore. 1985.
- Simon Peyton Jones, Andrew Gordon, Sigbjorn Finne. *Concurrent Haskell*, POPL'96.
- Ken Kahn. *A Computational Theory of Animation* MIT EECS Doctoral Dissertation. August 1979.

- Matthias Kaiser and Jens Lemcke *Towards a Framework for Policy-Oriented Enterprise Management* AAAI 2008.
- Alan Kay. "Personal Computing" in *Meeting on 20 Years of Computing Science* Instituto di Elaborazione della Informazione, Pisa, Italy. 1975.  
<http://www.mprove.de/diplom/gui/Kay75.pdf>
- Frederick Knabe *A Distributed Protocol for Channel-Based Communication with Choice* PARLE'92.
- Bill Kornfeld and Carl Hewitt. *The Scientific Community Metaphor* IEEE Transactions on Systems, Man, and Cybernetics. January 1981.
- Bill Kornfeld. *Parallelism in Problem Solving* MIT EECS Doctoral Dissertation. August 1981.
- Robert Kowalski *Algorithm = Logic + Control* CACM. July 1979.
- Robert Kowalski. *Response to questionnaire* Special Issue on Knowledge Representation. SIGART Newsletter. February 1980.
- Robert Kowalski (1988a) *The Early Years of Logic Programming* CACM. January 1988.
- Robert Kowalski (1988b) *Logic-based Open Systems Representation and Reasoning*. Stuttgart Conference Workshop on Discourse Representation, Dialogue tableaux and Logic Programming. 1988.
- Stein Krogdahl. *The birth of Simula* HiNC 1 Conference. Trondheim. June 2003.
- Leslie Lamport *How to make a multiprocessor computer that correctly executes multiprocess programs* IEEE Transactions on Computers. 1979.
- Peter Landin. *A Generalization of Jumps and Labels* UNIVAC Systems Programming Research Report. August 1965. (Reprinted in *Higher Order and Symbolic Computation*. 1998)
- Peter Landin *A correspondence between ALGOL 60 and Church's lambda notation* CACM. August 1965.
- Edward Lee and Stephen Neuendorffer (June 2004). *Classes and Subclasses in Actor-Oriented Design*. Conference on Formal Methods and Models for Codesign (MEMOCODE).
- Steven Levy *Hackers: Heroes of the Computer Revolution* Doubleday. 1984.
- Henry Lieberman. *An Object-Oriented Simulator for the Apiary* Conference of the American Association for Artificial Intelligence, Washington, D. C., August 1983
- Henry Lieberman. *Thinking About Lots of Things at Once without Getting Confused: Parallelism in Act 1* MIT AI memo 626. May 1981.
- Henry Lieberman. *A Preview of Act 1* MIT AI memo 625. June 1981.
- Henry Lieberman and Carl Hewitt. *A real Time Garbage Collector Based on the Lifetimes of Objects* CACM June 1983.
- Barbara Liskov *Data abstraction and hierarchy* Keynote address. OOPSLA'87.
- Barbara Liskov and Liuba Shrira *Promises: Linguistic Support for Efficient Asynchronous Procedure Calls* SIGPLAN'88.
- Barbara Liskov and Jeannette Wing. *Behavioral subtyping using invariants and constraints* in "Formal methods for distributed processing: a survey of object-oriented approaches" Cambridge University Press. 2001.
- Carl Manning. *Traveler: the Actor observatory* ECOOP 1987. Also appears in *Lecture Notes in Computer Science*, vol. 276.
- Carl Manning. *Acore: The Design of a Core Actor Language and its Compile* Master's Thesis. MIT EECS. May 1987.
- Satoshi Matsuoka and Aki Yonezawa. *Analysis of Inheritance Anomaly in Object-Oriented Concurrent Programming Languages Research Directions in Concurrent Object-Oriented Programming* MIT Press. 1993.
- John McCarthy *Programs with common sense* Symposium on Mechanization of Thought Processes. National Physical Laboratory, UK. Teddington, England. 1958.
- John McCarthy. *A Basis for a Mathematical Theory of Computation* Western Joint Computer Conference. 1961.
- John McCarthy, Paul Abrahams, Daniel Edwards, Timothy Hart, and Michael Levin. *Lisp 1.5 Programmer's Manual* MIT Computation Center and Research Laboratory of Electronics. 1962.
- John McCarthy. *Situations, actions and causal laws* Technical Report Memo 2, Stanford University Artificial Intelligence Laboratory. 1963.
- John McCarthy and Patrick Hayes. *Some Philosophical Problems from the Standpoint of Artificial Intelligence* Machine Intelligence 4. Edinburgh University Press. 1969.
- Giuseppe Milicia and Vladimiro Sassone. *The Inheritance Anomaly: Ten Years After SAC*. Nicosia, Cyprus. March 2004.
- Mark S. Miller *Robust Composition: Towards a Unified Approach to Access Control and Concurrency Control* Doctoral Dissertation. John Hopkins. 2006.
- George Milne and Robin Milner. "Concurrent processes and their syntax" *JACM*. April, 1979.
- Robert Milne and Christopher Strachey. *A Theory of Programming Language Semantics* Chapman and Hall. 1976.
- Robin Milner *Processes: A Mathematical Model of Computing Agents* Proceedings of Bristol Logic Colloquium. 1973.
- Robin Milner *Elements of interaction: Turing award lecture* CACM. January 1993.
- Marvin Minsky (ed.) *Semantic Information Processing* MIT Press. 1968.
- Eugenio Moggi *Computational lambda-calculus and monads* IEEE Symposium on Logic in Computer Science. Asilomar, California, June 1989.

- Allen Newell and Herbert Simon. *The Logic Theory Machine: A Complex Information Processing System*. Rand Technical Report P-868. June 15, 1956
- Kristen Nygaard. *SIMULA: An Extension of ALGOL to the Description of Discrete-Event Networks* IFIP'62.
- Carl Petri. *Kommunikation mit Automate* Ph. D. Thesis. University of Bonn. 1962.
- Simon Peyton Jones, Alastair Reid, Fergus Henderson, Tony Hoare, and Simon Marlow. *A semantics for imprecise exceptions* Conference on Programming Language Design and Implementation. 1999.
- Gordon Plotkin. *A powerdomain construction* SIAM Journal of Computing. September 1976.
- George Polya (1957) *Mathematical Discovery: On Understanding, Learning and Teaching Problem Solving Combined Edition* Wiley. 1981.
- Karl Popper (1935, 1963) *Conjectures and Refutations: The Growth of Scientific Knowledge* Routledge. 2002.
- John Reppy, Claudio Russo, and Yingqi Xiao *Parallel Concurrent ML* ICFP'09.
- John Reynolds. *Definitional interpreters for higher order programming languages* ACM Conference Proceedings. 1972.
- Bill Roscoe. *The Theory and Practice of Concurrency* Prentice-Hall. Revised 2005.
- Dana Scott and Christopher Strachey. *Toward a mathematical semantics for computer languages* Oxford Programming Research Group Technical Monograph. PRG-6. 1971
- Charles Seitz. *The Cosmic Cube* CACM. Jan. 1985.
- Peter Sewell, et. al. *x86-TSO: A Rigorous and Usable Programmer's Model for x86 Microprocessors* CACM. July 2010.
- Michael Smyth. *Power domains* Journal of Computer and System Sciences. 1978.
- Guy Steele Jr. *Lambda: The Ultimate Declarative* MIT AI Memo 379. November 1976.
- Gunther Stent. *Prematurity and Uniqueness in Scientific Discovery* Scientific American. December, 1972.
- Gerry Sussman and Guy Steele *Scheme: An Interpreter for Extended Lambda Calculus* AI Memo 349. December, 1975.
- Daniel Theriault. *A Primer for the Act-1 Language* MIT AI memo 672. April 1982.
- Daniel Theriault. *Issues in the Design and Implementation of Act 2* MIT AI technical report 728. June 1983.
- Hayo Thielecke *An Introduction to Landin's "A Generalization of Jumps and Labels"* Higher-Order and Symbolic Computation. 1998.
- Dave Thomas and Brian Barry. *Using Active Objects for Structuring Service Oriented Architectures: Anthropomorphic Programming with Actors* Journal of Object Technology. July-August 2004.
- Darrell Woelk. *Developing InfoSleuth Agents Using Rosette: An Actor Based Language* Proceedings of the CIKM '95 Workshop on Intelligent Information Agents. 1995.
- Akinori Yonezawa, Ed. *ABCL: An Object-Oriented Concurrent System* MIT Press. 1990.
- Aki Yonezawa *Specification and Verification Techniques for Parallel Programs Based on Message Passing Semantics* MIT EECS Doctoral Dissertation. December 1977.
- Hadasa Zuckerman and Joshua Lederberg. *Postmature Scientific Discovery?* Nature. December, 1986.

## End Notes

<sup>i</sup> In part, this section extends some material that was submitted to Wikipedia.

<sup>ii</sup> The receiver might be on another computer and in any case can make use of threads, locks, queues, cores, channels, etc. as it sees fit.

Messages in the Actor model are generalizations of packets in Internet computing in that they need not be received in the order sent. Not providing the order of delivery guarantee allows packet switching to buffer packets, use multiple paths to send packets, resend damaged packets, and to provide other optimizations.

For example, Actors are allowed to pipeline the processing of messages. What this means is that in the course of processing a message *m1*, an Actor can designate the behavior to be used to process the next message, and then in fact begin processing another message *m2* before it has finished processing *m1*. Just because an Actor is allowed to pipeline the processing of messages does not mean that it *must* pipeline the processing. Whether a message is pipelined is an engineering tradeoff.

<sup>iii</sup> Below is a simple implementation of an account that receives balance and withdraw messages:

```
SimpleAccount ===
```

```
  ① SimpleAccount is defined
```

```
behavior {
```

```
  Currency::theBalance, ① theBalance is an
```

```
  ① instance variable of type Currency
```

```
| implements Account
```

```
balance → theBalance, ① return balance
```

```
withdraw(amount) →
```

```
  ① for a withdraw message with amount
```

```
amount ?? {
```

```
  (> theBalance) :
```

```
  ① amount is greater than the value of theBalance  
    throw OverdrawnException;
```

```
  else :
```

```
    void, also Become (theBalance=  
                       theBalance - amount)}}
```

```
  ① return void also the next message is
```

```
  ① processed with amount deducted
```

```
CreateAccount(Currency::initialBalance) ===  
  create SimpleAccount(balance=initialBalance)
```

---

The operations of reading the balance and making withdrawals are commutative to the extent that their order doesn't matter. For example the following expression returns 2 even though the withdrawals can occur in either order:

```
let (Account::a = CreateAccount(5))  
  {a.withdraw(1), a.withdraw(2); a.balance}
```

<sup>iv</sup> including those involving unbounded nondeterminism

<sup>v</sup> One solution is to develop a concurrent variant of the Lisp meta-circular definition [McCarthy, Abrahams, Edwards, Hart, and Levin 1962] that was inspired by Turing's Universal Machine [Turing 1936]. If *exp* is a Lisp expression and *env* is an environment that assigns values to identifiers, then the procedure Eval with arguments *exp* and *env* evaluates *exp* using *env*. In the concurrent variant, *eval*(*env*) is a message that can be sent to *exp* to cause *exp* to be evaluated. Using such messages, modular meta-circular definitions can be concisely expressed in the Actor model for universal concurrent programming languages (e.g. ActorScript [Hewitt 2010a]).

<sup>vi</sup> E.g. processes in Erlang [Armstrong 2007] and vats in the object-capability model [Miller 2006].

<sup>vii</sup> In part, this section extends some material that was submitted to Wikipedia and [Hewitt 2008f].

<sup>viii</sup> Turing [1936] stated:

*the behavior of the computer at any moment is determined by the symbols which he [the computer] is observing, and his 'state of mind' at that moment" and "there is a bound B to the number of symbols or squares which the computer can observe at one moment. If he wishes to observe more, he must use successive observations."*

Gödel's conception of computation was formally the same as Turing but more reductionist in motivation:

*There is a major difference between the historical contexts in which Turing and Gödel worked. Turing tackled the Entscheidungsproblem [recursive decidability of provability] as an interesting mathematical problem worth solving; he was hardly aware of the fierce foundational debates. Gödel on the other hand, was passionately interested in the foundations of mathematics. Though not a student of Hilbert, his work was nonetheless deeply entrenched in the framework of Hilbert's finitistic program, whose main goal was to provide a meta-theoretic finitary proof of the consistency of a formal system "containing a certain amount of finitary number theory." Shagrir [2006]*

---

<sup>ix</sup> An Actor can have information about other Actors that it has received in a message about what it was like when the message was sent. See section of this paper on unbounded nondeterminism in ActorScript.

<sup>x</sup> Attempting to observe the internal processes of a physical arbiter affects its outcome because of indeterminacy.

<sup>xi</sup> However, the nondeterministic  $\lambda$ -calculus has bounded nondeterminism [Plotkin 1976] and thus is incapable of implementing concurrency.

<sup>xii</sup> Petri [1962]

<sup>xiii</sup> Consequently in Simula-76 there was no required locality of operations unlike the laws for locality in the Actor mode [Baker and Hewitt 1977].

<sup>xiv</sup> The ideas in Simula became widely known by the publication of [Dahl and Hoare 1972] at the same time that the Actor model was being invented to formalize concurrent computation using message passing [Hewitt, Bishop, and Steiger 1973].

<sup>xv</sup> The development of Planner was inspired by the work of Karl Popper [1935, 1963], Frederic Fitch [1952], George Polya [1954], Allen Newell and Herbert Simon [1956], John McCarthy [1958, *et. al.* 1962], and Marvin Minsky [1968].

<sup>xvi</sup> This turned out later to have a surprising connection with Direct Logic. See the Two-Way Deduction Theorem below.

<sup>xvii</sup> Subsequent versions of the Smalltalk language largely followed the path of using the virtual methods of Simula in the message passing structure of programs. However Smalltalk-72 made primitives such as integers, floating point numbers, etc. into objects. The authors of Simula had considered making such primitives into objects but refrained largely for efficiency reasons. Java at first used the expedient of having both primitive and object versions of integers, floating point numbers, etc. The C# programming language (and later versions of Java, starting with Java 1.5) adopted the more elegant solution of using boxing and unboxing, a variant of which had been used earlier in some Lisp implementations.

<sup>xviii</sup> According to the Smalltalk-72 Instruction Manual [Goldberg and Kay 1976]:

There is not one global message to which all message “fetches” (use of the Smalltalk symbols eyeball,  $\blacktriangleleft$ ; colon,  $\blacktriangleright$ , and open colon,  $\circ$ ) refer; rather, messages form a hierarchy which we explain in the following way--

---

suppose I just received a message; I read part of it and decide I should send my friend a message; I wait until my friend reads his message (the one I sent him, not the one I received); when he finishes reading his message, I return to reading my message. I can choose to let my friend read the rest of my message, but then I cannot get the message back to read it myself (note, however, that this can be done using the Smalltalk object *apply* which will be discussed later). I can also choose to include permission in my message to my friend to ask me to fetch some information from my message and to give that information to him (accomplished by including  $\blacktriangleright$  or  $\circ$  in the message to the friend). However, anything my friend fetches, I can no longer have. In other words,

- 1) An object (let's call it the CALLER) can send a message to another object (the RECEIVER) by simply mentioning the RECEIVER's name followed by the message.
- 2) The action of message sending forms a stack of messages; the last message sent is put on the top.
- 3) Each attempt to receive information typically means looking at the message on the top of the stack.
- 4) The RECEIVER uses the eyeball,  $\blacktriangleleft$ , the colon,  $\blacktriangleright$ , and the open colon,  $\circ$ , to receive information from the message at the top of the stack.
- 5) When the RECEIVER completes his actions, the message at the top of the stack is removed and the ability to send and receive messages returns to the CALLER. The RECEIVER may return a value to be used by the CALLER.
- 6) This sequence of sending and receiving messages, viewed here as a process of stacking messages, means that each message on the stack has a CALLER (message sender) and RECEIVER (message receiver). Each time the RECEIVER is finished, his message is removed from the stack and the CALLER becomes the current RECEIVER. The now current RECEIVER can continue reading any information remaining in his message.
- 7) Initially, the RECEIVER is the first object in the message typed by the programmer, who is the CALLER.
- 8) If the RECEIVER's message contains an eyeball,  $\blacktriangleleft$ ; colon,  $\blacktriangleright$ , or open colon,  $\circ$ , he can obtain further information from the CALLER's message. Any information successfully obtained by the RECEIVER is no longer available to the CALLER.
- 9) By calling on the object *apply*, the CALLER can give the RECEIVER the right to see all of the CALLER's remaining message. The CALLER can no longer get information that is read by the RECEIVER; he can, however, read anything that remains after the RECEIVER completes its actions.

- 
- 10) There are two further special Smalltalk symbols useful in sending and receiving messages. One is the keyhole, `⌈`, that lets the RECEIVER “peek” at the message. It is the same as the `⌈` except it does not remove the information from the message. The second symbol is the hash mark, `#`, placed in the message in order to send a reference to the next token rather than the token itself.

<sup>xix</sup> Of course the same limitation applies to the Abstract State Machine (ASM) model [Blass, Gurevich, Rosenzweig, and Rossman 2007a, 2007b; Glausch and Reisig 2006]. In the presence of arbiters, the global states in ASM are mythical.

<sup>xx</sup> The sender is an intrinsic component of communication in the following previous models of computation:

- *Petri Nets*: the input places of a transition are an intrinsic component of a computational step (transition).
- *$\lambda$ -Calculus*: the expression being reduced is an intrinsic component of a computational step (reduction).
- *Simula*: the stack of the caller is an intrinsic component of a computation step (method invocation).
- *Smalltalk 72*: the invoking token stream is an intrinsic component of a computation step (message send).

Process calculi (*e.g.* [Milner 1993; Cardelli and Gordon 1998]) are closely related to the Actor model. There are many similarities between the two approaches, but also several differences (some philosophical, some technical):

- There is only one Actor model (although it has numerous formal systems for design, analysis, verification, modeling, etc.); there are numerous process calculi, developed for reasoning about a variety of different kinds of concurrent systems at various levels of detail (including calculi that incorporate time, stochastic transitions, or constructs specific to application areas such as security analysis).
- The Actor model was inspired by the laws of physics and depends on them for its fundamental axioms, *i.e.* physical laws (see Actor model theory); the process calculi were originally inspired by algebra [Milner 1993].
- Unlike the Actor model, the sender is an intrinsic component of process calculi because they are defined in terms of reductions (as in the  $\lambda$ -calculus).
- Processes in the process calculi are anonymous, and communicate by sending messages either through named channels (synchronous or asynchronous), or via ambients (which can also be used to model channel-

like communications [Cardelli and Gordon 1998]). In contrast, Actors in the Actor model possess an identity, and communicate by sending messages to the mailing addresses of other Actors (this style of communication can also be used to model channel-like communications).

Process calculi can be modeled in the Actor model using a two-phase commit protocol [Knabe 1992].

<sup>xxi</sup> These laws can be enforced by a proposed extension of the X86 architecture that will support the following operating environments:

- CLR and extensions (Microsoft)
- JVM (Sun, IBM, Oracle, SAP)
- Dalvik (Google)

Many-core architecture has made the above extension necessary in order to provide the following:

- concurrent nonstop automatic storage reclamation (garbage collection) and relocation to improve efficiency,
- prevention of memory corruption that otherwise results from programming languages like C and C++ using thousands of threads in a process,
- nonstop migration of iOrgs (while they are in operation) within a computer and between distributed computers

<sup>xxii</sup> In summary, CSP differed from the Actor model in the following respects:

- *The concurrency primitives of CSP were input, output, guarded commands, and parallel composition* whereas the Actor model is based on asynchronous one-way messaging.
- *The fundamental unit of execution was a sequential process* in contrast to the Actor model in which execution was fundamentally concurrent. Sequential execution is problematical because multi-processor computers are inherently concurrent.
- *The processes had a fixed topology of communication* whereas Actors had a dynamically changing topology of communications. Having a fixed topology is problematical because it precludes the ability to dynamically adjust to changing conditions.
- *The processes were hierarchically structured using parallel composition* whereas Actors allowed the creation of non-hierarchical execution using futures [Baker and Hewitt 1977]. Hierarchical parallel composition is problematical because it precludes the ability to create a process that outlives its creator. Also message passing is the fundamental mechanism for generating parallelism in the Actor model; sending more messages generates the possibility of more parallelism.

- *Communication was synchronous* whereas Actor communication was asynchronous. Synchronous communication is problematical because the interacting processes might be far apart.
- *Communication was between processes* whereas in the Actor model communications are one-way to Actors. Synchronous communication between processes is problematical by requiring a process to wait on multiple processes.
- *Data structures consisted of numbers, strings, and arrays* whereas in the Actor model data structures were Actors. Restricting data structures to numbers, strings, and arrays is problematical because it prohibits programmable data structures.
- *Messages contain only numbers and strings* whereas in the Actor model messages could include the addresses of Actors. Not allowing addresses in messages is problematical because it precludes flexibility in communication because there is no way to supply another process with the ability to communicate with an already known process.

<sup>xxiii</sup> Of course, **n** would not survive the termination of **Z** and so the value cannot actually be exhibited after termination!

In the ActorScript [Hewitt 2010a] program below, the unbounded count is sent to the customer of the **start** message so that it appears externally.

<sup>xxiv</sup> Hoare [1985] developed a revised version of CSP with unbounded nondeterminism [Roscoe 2005].

<sup>xxv</sup> See [Hewitt 2008c] for a more complete history.

<sup>xxvi</sup> This thesis was also implicit in one interpretation of Cordell Green's earlier work [Green 1969].

<sup>xxvii</sup> Although [Goldin and Wegner 2008] may seem superficially similar, it unfortunately failed to comprehend previous publications on the Actor model (e.g. [Hewitt, Bishop and Steiger 1973], [Hewitt 1977], and [Hewitt and Agha 1988]).

<sup>xxviii</sup> For example see [Lampert 1979; Sewell, *et. al.* 2010].

<sup>xxix</sup> The cell interface can be defined as follows:

```
Cell=== ① Cell is defined
interface {
  ① to be an interface with the following methods
  ? ←read( ),           ① read message
  void ← write( ?::nextContents ) }
  ① write message with nextContents returns void
```

<sup>xxx</sup> However, they come with additional commitment. Inappropriate language constructs are difficult to leave behind.

<sup>xxxi</sup> Notable members of this community included Bill Gosper, Richard Greenblatt, Jack Holloway, Tom Knight, Stuart Nelson, Peter Samson, Richard Stallman, *etc.*. See [Levy 1984].

<sup>xxxii</sup> For semantics, Sussman and Steele [1975] proposed an interpreter for a time-sliced processor.

<sup>xxxiii</sup> Published in [Hewitt 1976, 1977].

<sup>xxxiv</sup> In addition to being less general than the Actor model, analysis in terms of continuation-passing style is also less rigorous and more complex than Actor semantics.

<sup>xxxv</sup> A restricted form of the escape expression [Reynolds 1972] can be defined as follows in ActorScript [Hewitt 2010a]:

```
escape identifier in body ===
behavior { implements Expression
  eval( Environment::e ) with Customer::c →
  body.
  eval( e.bind( identifier,
    to: (( ?::v ) with Customer::stranded →
      { c.return( v ), noResponse } ) ) )
  ① the return value v has been
  ① received as an argument to the
  ① continuation in a λ method
  ① The customer stranded will
  ① most likely be left hanging
```

Note that the continuation customer *stranded* is most likely stranded. And if the customer *stranded* is actually used then the caller of the escape continuation will most likely be rather surprised at the return! This is an important reason that customers are not  $\lambda$  expressions contra [Steele 1976].

Sussman and Steele [1975] developed a version of escape for Scheme named *call with current continuation*, a restricted version of which can be implemented in ActorScript as follows:

```
callWithCurrentContinuation expression ===
behavior { implements Expression
  eval( Environment::e ) with Customer::c →
  expression.eval( e )
  ( ?::v ) with Customer::stranded →
  { c.return( v ), noResponse } }
```



---

***However, customers differ from continuations because sending more than one response to a customer causes an exception to be thrown.***

Consequently, ActorScript [Hewitt 2010a] does not allow violation of stack discipline meaning that customers do not implement “*hairy control structure*” (the name given in [McDermott and Sussman 1972] for the control structure invented in [Landin 1965, Thielecke 1998]). Unfortunately, hairy control structure causes inefficiency and can make programs difficult to comprehend. ActorScript has much better ways than hairy control structure to implement generators, co-routines, *etc.* (See discussion below.)

<sup>xxxvi</sup> See implementations of continuations in ActorScript above.

<sup>xxxvii</sup> missing from Scheme

<sup>xxxviii</sup> This error was partially acknowledged in some of their subsequent work.

<sup>xxxix</sup> The  $\lambda$ -calculus includes the following limitations:

- Message arrival ordering cannot be implemented.
- Actors cannot be implemented whose behavior evolves with time.
- The  $\lambda$ -calculus does not have exceptions and consequently neither does Scheme [Sussman and Steele 1975].
- Attempting to reduce Actor customers to continuation functions was unsuccessful, e.g., it led to the hanging customer issue (see above).

<sup>xl</sup> as in `3+future(f(4))`

<sup>xli</sup> as in the following:

`resolve promise(f(4)) {value v : 3+v,  
                          exception e : throw e}`

<sup>xlii</sup> Note that there is a limitation on concurrency because `u.get()` must complete before P starts

<sup>xliii</sup> As above, there is a limitation on concurrency because `u.put(m)` must complete before P starts