

MASSACHUSETTES INSTITUTE OF TECHNOLOGY
ARTIFICIAL INTELLIGENCE LABORATORY

AI.MEMO-428

June 1977

MODELLING DISTRIBUTED SYSTEMS¹

by

AKINORI YONEZAWA and CARL HEWITT

Abstract

Distributed systems are multi-processor information processing systems which do not rely on the central shared memory for communication. This paper presents ideas and techniques in modelling distributed systems and its application to Artificial Intelligence. In section 2 and 3, we discuss a model of distributed systems and its specification and verification techniques. We introduce a simple example of air line reservation systems in Section 4 and illustrate our specification and verification techniques for this example in the subsequent sections. Then we discuss our further work.

This report describes research done at the Artificial Intelligence Laboratory of the Massachusetts Institute of Technology. Support for the laboratory's artificial intelligence research is provided in part by the Advance Research Projects Agency of the Department of Defence under Office of Naval Research contract N00014-75-C0522

1. This paper will appear in the proceeding of the 5-th International Joint Conference on Artificial Intelligence, Cambridge, U.S.A., August, 1977.

1. Introduction

Distributed systems are multi-processor information processing systems which do not rely on the central shared memory for communication. The importance of distributed systems has been growing with the advent of "computer networks" of a wide spectrum: Networks of geographically distributed computers at one end, and tightly coupled systems built with a large number of inexpensive physical processors at the other end. Both kinds of distributed systems are made available by the rapid progress in the technology of large scale integrated circuits. Yet little has been done in the research on semantics and programming methodologies for distributed information processing systems.

Our main research goal is to understand and describe the behavior of such distributed systems in seeking the maximum benefit of employing multi-processor computation schemata.

The contribution of such research to Artificial Intelligence is manifold. We advocate an approach to modelling intelligence in terms of cooperation and communication between knowledge-based problem solving experts. In this approach, we present a coherent methodology for the distribution of active knowledge as a knowledge representation theory. Also this methodology provides flexible control structures which we believe are well suited for organizing distributed active knowledge. Furthermore we hope to make technical contributions to the central issues of problem solving such as parallel versus serial processing, centralization versus decentralization of control and information storage, and the "declarative-procedural" controversy.

This paper presents ideas and techniques in modelling distributed systems and its application to Artificial Intelligence. In section 2 and 3, we discuss a model of distributed systems and its specification and verification techniques. We introduce a simple example of air line reservation systems in Section 4 and illustrate our specification and verification techniques for this example in the subsequent sections. Then we discuss our further work.

2. A Model of Distributed Systems

The actor model of computation [Greif&Hewitt75, Greif75, Hewitt&Baker77] has been developed as a model of communicating parallel processes. The fundamental objects in the model of computation are actors. An actor is a potentially active piece of knowledge (procedure) which is activated when it is sent a message which is also an actor. Actors interact by sending messages to other actors. More than one transmission of messages may take place concurrently. Each actor decides how to respond to messages sent to it. An actor is defined by its two parts, a script and a set of acquaintances. Its script is a description of how it should behave when it is sent a message. Its acquaintances are a finite set of actors that it directly knows about. If an actor A knows about another actor B, A can send a message to B directly. The concept of an event is fundamental in the actor model of computation. An event is an arrival of a message actor M at a target actor T and is denoted by the expression $(T \leftarrow M)$. A computation is expressed as a partially ordered set of events. We call this partial order the "precedes" ordering. Events which are unordered in the computation can be concurrent. Thus the partial order of events naturally generalizes the notion of serial computation (which is a sequence of events) to that of parallel computation.

A collection of actors which communicate and cooperate with each other in a goal oriented fashion can be implemented as a single actor. In essence actors are procedural objects which may or may not have local storage. Some may behave like procedures and some may behave like data structures. Modules in distributed systems are modelled by actors and systems of actors. In this regard, IC chips can be viewed as actors.

Knowledge and intelligence can be embedded as actors in a modular and distributed fashion. For example, frames [Minsky75, Kuipers75], units [Bobrow&Winograd76], beings [Lenat75], stereotypes [Hewitt75] e.t.c. which represent modular knowledge with procedural attachments are modelled and implemented as actors. In the context of electronic mail systems and business information systems, objects such as forms, documents, customers, mail collecting stations, and mail

distributing stations are easily modelled and implemented as actors.

Messages which are sent to target actors usually contain *continuation* actors to indicate where the result of the receipt of the message should be sent. By virtue of continuations in messages, the message-passing in the actor model of computation realizes a universal and yet flexible control structure without using implicit mechanisms such as push down stacks. Various forms of control structures such as go-to's, procedure calls, and coroutines can be viewed as particular patterns of message passing [Hewitt76].

This model of computation has been implemented as a programming language PLASMA[Hewitt76]. The script of an actor can be written as a PLASMA program. We believe that PLASMA will provide a basis for programming languages for distributed systems. In section 5, an example of PLASMA programs is given as a script of a flight-data actor in the model of a simple air line reservation system.

3. Techniques for Specification and Verification

In designing and implementing a distributed (message-passing) system, it is desirable to have a precise specification of the intended behavior of the distributed system. Also we need sound techniques for demonstrating that implementations of the system meet its specification. Below we give some of the central ideas of our specification and verification techniques based on the model introduced in the previous section. The more detailed work will be found in [Yonezawa77].

In specifying the behavior of a distributed system, it is not only practically infeasible, but also irrelevant to use global states of the entire system or the global time axis which governs the uniform time reference throughout the system. We are concerned with states of modular components of a distributed system which interact with each other by sending messages. Thus we are interested in the states of actors participating in an event at the instance at which the message is received.

In our specification language, conceptual representations are used to express

local states of actors (modules). Conceptual representations were originally developed to specify the behavior of actors which behave like data structures[Yonezawa&Hewitt76]. We have found them very useful to express states of modules in distributed systems at varying levels of abstraction and also from various view points. The basic motivation of conceptual representations is to aid in providing a specification language which serves as a good interface between programmers and the computer and also between users and implementors. Conceptual representations are intuitive clear and easy to understand, yet their rigorous interpretations are provided. Instead of going into details of syntactic constructs of conceptual representations, we give examples. Below `!<exp>` is the unpack operation on `<exp>` which means writing out all elements denoted by `<exp>` individually.

```
(CELL A) ;a cell containing A as its contents.
(QUEUE A B C) ;a queue with elements A B C.
(NODE (car: A)(cdr: B)) ;a LISP node containing A and B.
(CUSTOMER (letters: {!m})(#-of-stamps-needed: n)) ;a customer visiting a post office
;who carries letters !m and wants n stamps.
(POST-OFFICE (customer: {!c}) (collector: {!cl})) ;a post office which contains customers !c and mail collectors !cl.
```

It should be noted that a conceptual representation does not represent the identity of an actor. It only provides a description of the state of an actor. Thus to state that an actor Q is in the state expressed by a conceptual representation `(QUEUE A B C)`, an assertion of the following form:

`(Q is-a (QUEUE A B C))`

is used. Some examples of specification using conceptual representation are given in the later sections.

Symbolic evaluation is a process which interprets a module on abstract data to demonstrate that the module satisfies its specification. Symbolic evaluation differs from ordinary evaluation in that 1) the only properties of input that can be used are the ones

specified in the pre-requisites, and 2) if the symbolic evaluation of a module M encounters an invocation of some module N, the specification of N is used to continue the symbolic evaluation. The implementation of N is not used. The technique of symbolic evaluation has been studied by a number of reseachers, for example [Boyer&Moore75, Burstall&Darlington75, Hewitt&Smith75, Yonezawa75, King76].

Our method for symbolic evaluation of distributed systems is an extention of the one developed for symbolic evaluation of programs written in SIMULA-like languages[Yonezawa&Hewitt76]. One of the main techinques we employ in symbolic evaluation is the introduction of a notion of *situations*[McCarthy&Hayes69]. A situation is the *local* state of an actor system at a give moment. The precise definition of locality in the actor model of computation is found in [Hewitt&Baker77]. By relativizing states of modules with *situational tags* which denote situations, relations and assertions about states of modules in different situations can be expressed. Explicit uses of situational tags seem to be very powerful in symbolic evaluation of distributed systems. A simple example is given in Section 7.

Another technique we employ in symbolic evaluation is the use of *actor induction* to prove properties holding in a computation. Actor induction is a computational induction based on the *precedes* ordering (cf. Section 2) among events. It can be stated intuitively as follows:

"For each event E in a computation C, if *preconditions* for E imply *preconditions* for each event E' which is immediately caused by E, then the computation C is carried out according to the overall specifications."

The precedes ordering has two kinds of suborderings, 1) the activation ordering, "*activates*", which is the causal relation among events, and 2) the arrival ordering, "*arrives-before*", which expresses ordering among events which have the same target actor. Thus there are two kinds of actor induction according to these suborderings. An example of the induction based on the arrival ordering is used in Section 7.

4. Modelling an Air Line Reservation System

-- A specification of an Air Line Reservation System --

As an illustrative example of distributed systems, let us consider a very simple air line reservation system. Suppose we have just one flight which has a non-negative number of seats. A number of travel agencies (parallel processes) independently try to reserve or cancel seats for this flight, possibly concurrently. We model an air line reservation system as a flight actor *F* which behaves as follows. The flight actor *F* accepts two kinds of message, (*reserve-a-seat:*) and (*cancel-a-seat:*). When *F* receives (*reserve-a-seat:*), if the number of free seats is zero, a message (*no-more-seats:*) is returned. Otherwise a message (*ok-its-reserved:*) is returned and the number of free seats is decreased by one. When *F* receives (*cancel-a-seat:*), if the number of free seats is less than the maximum number of seats of the flight, a message (*ok-its-cancelled:*) is returned and the number of free seats is increased by one, otherwise (*too-many-cancels:*) is returned. Furthermore requests by (*reserve-a-seat:*) and (*cancel-a-seat:*) are served on a first-come-first-served base.

To write a formal specification of the air line reservation system, we need to describe the states of the flight actor. For this purpose, we use the following conceptual representation

(FLIGHT (seats-free: <m>) (size: <s>))

which describes the state of a flight actor. The number of free seats is *<m>* and *<s>* is the size of the flight in terms of the total number of seats. The formal specification of the air line reservation system using this conceptual representation is depicted in Figure 1 below.

```
<event: (create-flight (= S)
  <pre-cond: (S > 0) >
  <return: F* >
  <post-cond: (F is-a (FLIGHT (seats-free: S) (size: S))))>

<event: (F (= (reserve-a-seat:))
  (case-1:
    <pre-cond: (F is-a (FLIGHT (seats-free: 0) (size: S))))>
    <next-cond: (F is-a (FLIGHT (seats-free: 0) (size: S))))>
    <return: (no-more-seats:) >)
  (case-2:
    <pre-cond:
      (F is-a (FLIGHT (seats-free: N) (size: S)))
      (N > 0) >
    <next-cond: (F is-a (FLIGHT (seats-free: N - 1) (size: S))))>
    <return: (ok-its-reserved:) >))

<event: (F (= (cancel-a-seat:))
  (case-1:
    <pre-cond: (F is-a (FLIGHT (seats-free: S) (size: S))))>
    <next-cond: (F is-a (FLIGHT (seats-free: S) (size: S))))>
    <return: (too-many-cancels:) >)
  (case-2:
    <pre-cond:
      (F is-a (FLIGHT (seats-free: N) (size: S)))
      (N < S) >
    <next-cond: (F is-a (FLIGHT (seats-free: N + 1) (size: S))))>
    <return: (ok-its-cancelled:) > >)

<for-events: E, E'
  where E = (F (= M), E' = (F (= M'))
  <pre-cond:
    (F is-a (FLIGHT (seats-free: ...) (size:...)))
    (E arrives-before E')>
  <caused-events: reply-for[E], reply-for[E']>
  <post-cond: (reply-for[E] precedes reply-for[E']) >>
```

Figure 1 A Specification of the Air Line Reservation System
(A Specification for the Flight Actor)

The first <event:...>-clause states that a new flight actor F is created by an event where the create-flight actor receives a positive number S. <actor>* means that <actor> is newly created. The second <event:...>-clause has two cases according to the number of free seats at the moment when the flight actor F receives (reserve-a-seat:). When the number of free seats is zero (Case-1), the state of F does not change. When it is positive (Case-2), the

number of free seats decreases by one as stated by the assertion in the *<next-cond:...>*-clause. The notation in Figure 1:

```
<event: (T <= M)  
<pre-cond: ... >  
<next-cond: ... <assertion> ... >  
<return: <actor> >>
```

means that when an event ($T \leq M$) takes place, if the preconditions are satisfied, *<assertion>*s in the *<next-cond:...>*-clause hold immediately after the event until the next message arrives at T . *<actor>* in the *<return:...>*-clause is returned as the result of the event. *<next-cond:...>* differs from *<post-cond:...>* in that assertions in *<post-cond:...>*-clause hold at the time *<actor>* is returned, whereas assertions in *<next-cond:...>*-clause hold at the time the next message arrives. The next message may arrive at T before or after a reply for the previous message is returned. The third *<event:...>*-clause is for the cancelling event, which is interpreted in a similar way. The *<for-events: ...>*-clause states that requests (messages) received by the flight actor are served on the first-come-first-served base. Namely, the replying events for events E and E' take place in the same order as E and E' .

5. Implementing the Air Line Reservation System

Our strategy to implement the air line reservation system (specified in the previous section) is as follows. First, we implement a flight-data actor which satisfies the specification in Figure 1 on the condition that it is always activated *serially*. Then we put some protecting (or scheduling) mechanism on the flight-data actor so that the protected flight-data actor may satisfy the specification of the air line reservation system.

In Figure 2 below we give an implementation of the flight-data actor in PLASMA.

```

(create-flight-data =s) ≡
  (let (size initially s)
    (seats-free initially s)
    then
      (cases
        (⇒ (reserve-a-seat:)
          (rules seats-free
            (⇒ 0
              (no-more-seats:)))
          (⇒ else
            (seats-free - (seats-free - 1))
            (ok-its-reserved:))))
        (⇒ (cancel-a-seat:)
          (rules seats-free
            (⇒ size
              (too-many-cancels:))
            (⇒ else
              (seats-free - (seats-free + 1))
              (ok-its-cancelled:))))))
      ;create-flight-data receives a size s of flight.
      ;a variable size is set to s.
      ;a variable seats-free is set to s.
      ;the following cases-clause is
      ;returned as an actor which behaves as a flight-data.
      ;when a (reserve-...) message is received,
      ;if seats-free is zero,
      ;(no-...) message is returned.
      ;otherwise
      ;seats-free is decreased by one.
      ;(ok-...) message is returned.
      ;when a (cancel-...) message is received,
      ;if seats-free is equal to size,
      ;(too-...) is returned,
      ;otherwise
      ;seats-free is increased by one.
      ;(ok-...) is returned.

```

Figure

It is fairly straightforward to write a specification for this flight-data FD by using a conceptual representation:

(FLIGHT-DATA (seats-free: <m>) (size: <s>))

which describes the state of a flight-data actor. The number of free seats is <m> and <s> is the size of the flight in terms of the number of seats. Note that if FD were sent more than one message concurrently, anomalous results would be caused. For example, in the implementation in Figure 2, if *(reserve-a-seat:)* and *(cancel-a-seat:)* messages are sent concurrently, *(no-more-seats:)* message might be returned even if there are vacant seats. Therefore in order to model the air line reservation system by using the above implementation of a flight-data actor, the way it is used must be restricted so that interference between different activations does not take place. As suggested in the beginning of this section, the restriction we impose is that FD must be used serially in the sense that FD is not allowed to receive a message until the activation by the previous message is completed. Now the flight-data actor can be used to implement the air line

reservation system under this restriction. We give a formal specification for the flight-data actor in Figure 3 below.

```
<event: (create-flight-data <= S)
  <pre-cond: (S > 0)>
  <return: FD* >
  <post-cond: (FD is-a (FLIGHT-DATA (seats-free: S) (size: S)))>>

<event: (FD <= (reserve-a-seat:))
  (case-1:
    <pre-cond:
      (FD is-used-serially)
      (FD is-a (FLIGHT-DATA (seats-free: 0) (size: S)))>
    <return: (no-more-seats:) >
    <post-cond: (FD is-a (FLIGHT-DATA (seats-free: 0) (size: S))) >
  (case-2:
    <pre-cond:
      (FD is-used-serially)
      (FD is-a (FLIGHT-DATA (seats-free: N) (size: S)))
      (N > 0) >
    <return: (ok-its-reserved:) >
    <post-cond: (FD is-a (FLIGHT-DATA (seats-free: N - 1) (size: S))) >>

<event: (FD <= (cancel-a-seat:))
  (case-1:
    <pre-cond:
      (FD is-used-serially)
      (FD is-a (FLIGHT-DATA (seats-free: S) (size: S))) >
    <return: (too-many-cancels:) >
    <post-cond: (FD is-a (FLIGHT-DATA (seats-free: S) (size: S))) >
  (case-2:
    <pre-cond:
      (FD is-used-serially)
      (FD is-a (FLIGHT-DATA (seats-free: N) (size: S)))
      (N < S) >
    <return: (ok-its-cancelled:) >
    <post-cond: (FD is-a (FLIGHT-DATA (seats-free: N + 1) (size: S))) >>
```

Figure 3 A Specification for the Flight-data Actor

In this specification, the restriction of the serial use is expressed in the following notation,

(FD is-used-serially)

stated as a precondition for events. In contrast to the specification above, there are no such

preconditions in the specification of the air line reservation system (the flight actor) in Figure 1. Thus the reservation system is specified to work properly even if it is accessed concurrently. Also notice that the specification above has no statements about scheduling such as the first-come-first-served scheduling which is stated as *<for-events:...>*-clause in the specification of the air line reservation system.

6. One-at-a-time

In this section, we consider how the serial use of a flight-data actor is realized in environments where communicating parallel processes try to use the flight-data actor. Our approach is to surround a flight-data actor FD with some mechanism which arbitrates parallel requests to the flight-data actor FD and passes these requests to FD in the serial fashion. We call this protection mechanism a one-at-a-time guardian. A one-at-a-time guardian can be easily implemented by a *serializer*[Atkinson&Hewitt77] which is a general synchronization mechanism in the actor model of computation.

Now we give a specification for one-at-a-time guardians. A one-at-a-time guardian is created in an event where an actor one-at-a-time receives a resource (a flight-data actor in this case). The one-at-a-time guardian thereby created will then contain the received resource. The following *<event:...>*-clause expresses this.

```
<event: (one-at-a-time (= RESOURCE)
  <return: G* >
  <post-cond: (G is-a (ONE-AT-A-TIME RESOURCE)) >>
```

where *(ONE-AT-A-TIME <resource>)* is the conceptual representation for a one-at-a-time guardian which contains *<resource>*. Next, we specify how a one-at-a-time guardian G behaves. In general a *request* to the guardian G, which is an arrival of a message M at G, eventually causes an invocation (or use) of RESOURCE. The invocation of RESOURCE begins with an *access* to RESOURCE which is an arrival of the same message M at RESOURCE and ends with a *reply* for the *access* which is a return of some result of the invocation. (See

the Figure 4 below.)


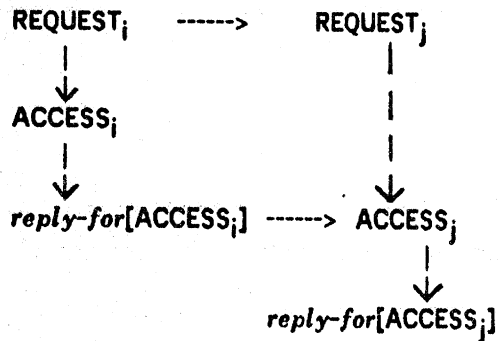


Figure 4

Our aim of using a *one-at-a-time* guardian G is to control invocations of $RESOURCE$ by parallel requests so that only one invocation of $RESOURCE$ takes place at a time. In order to do so, if we have two concurrent requests, the end of the invocation by one request should always precede the beginning of the invocation of the other request. This intuitive description of the desired behavior of a *one-at-a-time* guardian can be described in terms of the order of the events *request*, *access* and *reply* introduced above. Suppose we have two requests, $REQUEST_i$ which is an arrival of a message M_i at G , and $REQUEST_j$ which is an arrival of a message M_j at G . Then $REQUEST_k$ causes $ACCESS_k$ which is an arrival of M_k at $RESOURCE$ resulting in *reply-for* $[ACCESS_k]$, in this order (where k stands for either i or j). To ensure the *one-at-a-time* property of invocations of a resource, the following ordering relation must be satisfied:

"if $REQUEST_i$ *precedes* $REQUEST_j$,
then *reply-for* $[ACCESS_i]$ *must precede* $ACCESS_j$ ".

Since $REQUEST_k$ always precedes $ACCESS_k$ and $ACCESS_k$ always precedes *reply-for* $[ACCESS_k]$, the desired ordering relation can be expressed by the following diagram.



This behavior of the one-at-a-time guardian is formally described as a specification in Figure 5 below. Note that RESOURCE must be guaranteed to reply.

```

<event: (one-at-a-time <= RESOURCE)
  <return: G* >
  <post-cond: (G is-a (ONE-AT-A-TIME RESOURCE)) >>

<for-events: REQUESTi, REQUESTj
  where REQUESTi = (G <= Mi), REQUESTj = (G <= Mj)
  <pre-cond:
    (G is-a (ONE-AT-A-TIME RESOURCE))
    (RESOURCE is-guaranteed-to-reply)
    (REQUESTi precedes REQUESTj) >

<caused-events: ACCESSi, ACCESSj, reply-for[ACCESSi], reply-for[ACCESSj]
  where ACCESSi = (RESOURCE <= Mi), ACCESSj = (RESOURCE <= Mj)>
  <post-cond:
    (REQUESTi precedes ACCESSi)
    (REQUESTj precedes ACCESSj)
    (ACCESSi precedes reply-for[ACCESSi])
    (ACCESSj precedes reply-for[ACCESSj])
    (reply-for[ACCESSi] precedes ACCESSj)>>
  
```

Figure 5 A Specification for the One-at-a-Time Actor

7. Symbolic Evaluation of the Air Line Reservation System

Our implementation of the air line reservation system is expressed by the following simple code:

```
(create-flight =s) ≡ (one-at-a-time (create-flight-data s))
```

(Equivalently, $(\text{create-flight } =s) \equiv (\text{one-at-a-time } \leq (\text{create-flight-data } \leq s)).$)

In this section we demonstrate that the above code meets the specification of the air line reservation system given in Figure 1. Our method for the demonstration is symbolic evaluation.

The symbolic evaluation of the code


$(\text{one-at-a-time } (\text{create-flight-data } s))$

reveals the following facts:

- 1) an actor FD is created by $(\text{create-flight-data } \leq s)$,
- 2) G is created by $(\text{one-at-a-time } \leq \text{FD})$ and returned, and
- 3) the two actors satisfy the following assertions immediately after the creation of G

$(\text{FD is-a } (\text{FLIGHT-DATA } (\text{seats-free: } s) (\text{size: } s)))$

$(\text{G is-a } (\text{ONE-AT-A-TIME } \text{FD}))$.



This means that the flight actor is created as a **one-at-a-time** guardian G which contains a flight-data actor FD with s free seats. In what follows, we will establish that the **one-at-a-time** guardian G satisfies the specification for the flight actor in Figure 1.

The $\langle \text{event:} \dots \rangle$ -clause in the specification for the flight actor in Figure 1 specifies the behavior of G in terms of the conceptual representation

$(\text{G is-a } (\text{FLIGHT } (\text{seats-free:} \dots)(\text{size:} \dots)))$

(Notice that F in the specification for the flight actor is instantiated as G.) On the other hand, G is implemented as a **one-at-a-time** guardian which contains the flight-data actor FD.

This means that we have two views of G and correspondingly two different conceptual representations are used to describe the state of G. In order to show that the implementation satisfies the specification, we need to establish some relation between the state of G expressed by

(FLIGHT (seats-free:...) (size:...))

and the state of FD expressed by

(FLIGHT-DATA (seats-free:...)(size:...)).

The relation we need is:

"If G satisfies the assertion

(G is-a (FLIGHT (seats-free: n) (size: s)))

in a situation where G receives a message M, then FD always satisfies the assertion

(FD is-a (FLIGHT-DATA (seats-free: n) (size: s)))

in the situation where FD receives the same message M (through the one-at-a-time guardian), and vice versa."

This relation is expressed formally as follows:

<implementation-commentary:

(G is-a (FLIGHT (seats-free: n) (size: s))) in S

where S = Sit[(G <= M)]

if-and-only-if

(FD is-a (FLIGHT-DATA (seats-free: n) (size: s))) in S'

where S' = Sit[(FD <= M)] >.

Sit[E] expresses the situation where an event E takes place. The above implementation commentary formally describes the basic idea of the implementation. It can be viewed as the counterpart of an "invariant" in parallel process environments, which was first introduced by [Hoare 1972] to show correctness of implementation of data structures which are supposed to be used serially.

It should be noted that the first-come-first-served based scheduling by the

guardian **G** guarantees the above relation. If the guardian does more complicated scheduling, the relation needed for the demonstration may not be so simple. For more general scheduling cases, see [Yonezawa77].

I. Establishing the $\langle \text{event: } (G \Leftarrow (\text{reserve-a-seat:})) \dots \rangle$ -clause

There are two cases to be considered. We only consider the (Case-2...)-clause.

Case-2: $(G \text{ is-a } (FLIGHT \text{ (seats-free: } n) \text{ (size: } s))), (n > 0)$

The guardian **G** receives a (reserve-a-seat:) message **M**. To know the result of this event, the specification for **one-at-a-time** in Figure 5 is used. Since the flight-data actor **FD** is guaranteed to reply, the specification for **one-at-a-time** guarantees that the (reserve-a-seat:) message **M** is received by **FD**. To know the state of the flight-data actor **FD** at the time of the arrival of **M**, the above implementation commentary is used. Since the state of **G** at the time of the arrival of **M** at **G** is described as:

$(G \text{ is-a } (FLIGHT \text{ (seats-free: } n) \text{ (size: } s))),$

the state of **FD** at the arrival of **M** at **FD** is described as

$(FD \text{ is-a } (FLIGHT-DATA \text{ (seats-free: } n) \text{ (size: } s))).$

Then the (Case-2...)-clause in the $\langle \text{event:} \dots \rangle$ -clause of the specification for flight-data actors in Figure 3 is referred to. Since the precondition that **FD** must be used serially is satisfied (because **FD** is contained inside the **one-at-a-time G**), the (Case-2...)-clause of the specification for flight-data actors tells us that

(1) $(\text{ok-its-reserved:})$ is returned, and

(2) the state of **FD** is now expressed as:

$(FD \text{ is-a } (FLIGHT-DATA \text{ (seat-free: } n - 1) \text{ (size: } s))).$

(I) is what the *<return:...>*-clause in the specification for the flight in Figure 1 requires. Since the state of FD expressed as

(FD is-a (FLIGHT-DATA (seat-free: n - 1) (size: s)))

remains unchanged until the next message M' arrives at FD, by using the implementation commentary in the other direction this time, we know that the state of G remains unchanged as

(G is-a (FLIGHT (seats-free: n - 1) (size: s)))

until the message M' arrives at G. This is what *<next-cond:...>* clause in the specification for the flight actor in Figure 1 requires. Thus Case-2 is shown. Case-1 may be shown analogously. It should be noted that induction on the order of arrival of messages is used.

II. Establishing the *<event: (G <= (cancel-a-seat:))...>*-clause

The demonstration for this event is analogous to that of I.

III. Establishing the *<for-events:...>*-clause

The event where the flight actor G receives a message means that the one-at-a-time guardian receives the same message. Suppose that M and M' arrive at G in this order. The specification for the one-at-a-time guardian specifies that M' is not received by FD until the reply from FD for M is completed. Therefore the reply to M' always takes place after the reply to M. This is what the specification requires.

IV. Establishing the Confinement of the flight-data actor FD

The discussion in I, II and III above assumes that no one can access the flight-data actor FD except through the guardian G. This assumption always holds because the flight-data actor FD created by *(create-flight-data <= s)* is never released outside the

one-at-a-time actor.

8. Further Work

We are currently working to establish a coherent methodology for demonstrating that a distributed message-passing system will meet its specifications. By using the technique of symbolic evaluation, we would like to analyze the relationships and dependencies between modules in a distributed system. This approach will be instrumental in assisting us with the evolutionary development of distributed systems.

We are also working on the application of procedural objects (such as actors) to the area of business automation. In order to replace paper forms and paper documents, we use "active" forms and "active" documents which are displayed as images on the TV terminal accompanied by procedures. Active forms and documents are sent from one site to another whereby clerks are requested to provide necessary information with the guidance of the accompanying procedures. Such procedures may also check the consistency of filled items and point out errors and inconsistencies to persons who are processing forms. Thus active forms and documents accompanied by procedures enormously increase the flexibility and security of message and document systems. Furthermore we propose to use the "language" of forms and documents as the basis for the user to communicate with the information processing system. One of the ultimate objectives in our research is to develop a methodology for the construction of real-time distributed systems which can be efficiently and effectively used by non-programmers.

9. Acknowledgements

Conversations with Jeff Rulifson have been helpful in further developing the notion of an "active" document. This research was conducted at the Artificial Intelligence Laboratory and Laboratory for Computer Science (formerly Project MAC), Massachusetts

Institute of Technology under the sponsorship of the Office of Naval Research, contract number N00014-75C0522.

10. Bibliography

- Atkinson, R. and Hewitt, C. "Synchronization in Actor Systems" 4-th SIGPLAN-SIGACT Symposium on Principles of Programming Languages, Los Angeles., January, 1977.
- Birtwistle, G., Dahl, O.-J., Myhrhang, B. and Nygaard, K. SIMULA Begin Auerbach, Philadelphia., 1973.
- Bobrow, R. S. and Winograd, T. "An Overview of KRL, a Knowledge Representation Language" CSL-76-4. Xerox Palo Alto Research Center., July, 1976.
- Boyer, R. S. and Moore, J. S. "Proving Theorems about LISP Functions" JACM. Vol.22. No.1. January, 1975.
- Burstall, R. M and Darlington, J "Some Transformations for Developing Recursive Functions" Proc. of International Conference on Reliable Software, Los Angeles., April, 1975.
- Greif, I. "Semantics of Communicating Parallel Processes" Ph.D Thesis MIT, also Technical Report TR-154., Laboratory for Computer Science (formerly Project MAC)., September, 1975.
- Greif, I. and Hewitt, C. "Actor Semantics of PLANNER-73" Proc. of ACM SIGPLAN-SIGACT Conference. Palo Alto, California. January, 1975.
- Hewitt, C. "How to Use What You Know" IJCAI-75, USSR., September, 1977.
- Hewitt, C. "Viewing Control Structures as Patterns of Passing Messages" AI-MEMO No.410., Artificial Intelligence Laboratory, MIT., December. 1976., also to appear in the Journal of Artificial Intelligence.
- Hewitt, C. and Baker, H. "Laws for Communicating Parallel Processes" Proceeding of IFIP-77, Toronto, August. 1977, also Working Paper 134. Artificial Intelligence Laboratory, MIT., December, 1976.
- Hewitt, C. and Smith, B.C. "Towards a Programming Apprentice" IEEE Transaction on Software Engineering, Vol. SE-1 No. 1., March, 1975.
- Hoare, C.A.R. "Proof of Correctness of Data Representation" Acta Informatica Vol. 1., pp271-281. 1972

- King, J. "Symbolic Execution and Program Testing" CACM. Vol.19 No. 7., July, 1976.
- Kulpers, B. J. "A Frame for Frames: Representating Knowledge for Recognition" in D. G. Bobrow & A. M. Collins (Ed.) Representation and Understanding Academic Press, New York., 1975.
- Learning Research Group "Personal Dynamic Media" SSL-76-1. Xerox Palo Alto Research Center., April, 1976
- Lenat, D. B. "Beings: Knowledge as Interacting Experts" IJCAI-75, USSR., September, 1975.
- McCarthy, J. and Hayes, P. "Some Philosophical Problems from the Standpoint of Artificial Intelligence" Machine Intelligence Vol.4. American Elsevier, New York., 1969.
- Minsky, M. "A Framework for Representing Knowledge" in P. H. Winston (Ed.) The Psychology of Computer Vision, McGraw-Hill, New York., 1975.
- Rich, C. and Shrobe, H.E. "Understanding Lisp Programs: Towards a Programmer's Apprentice" Masters' Thesis, Electrical Engineering and Computer Science, MIT August, 1975., also AI-TR No.354, Artificial Intelligence Laboratory, MIT.
- Steiger, R. "Actor Machine Architecture" Master's Thesis., Department of Electrical Engineering and Computer Science, MIT., June 1974.
- Yonezawa, A. "Meta-evaluation of Actors with Side-effects" Working paper 101., Artificial Intelligence Laboratory, MIT., June, 1975.
- Yonezawa, A. Forthcoming Ph.D. Thesis., Department of Electrical Engineering and Computer Science, MIT., 1977.
- Yonezawa, A. and Hewitt, C. "Symbolic Evaluation using Conceptual Representations for Programs with Side-effects." AI-Memo. No.399., Artificial Intelligence Laboratory, MIT., December, 1976.

