

**A SIMULATOR FOR MESSAGE-BASED  
DISTRIBUTED SYSTEMS**

William F. Quinlivan III

Department of Computer Sciences  
The University of Texas at Austin  
Austin, Texas 78712-1188

TR-87-33

August 1987

**Abstract**

A broad class of physical systems including queueing, communication, and computer networks can be modelled as a collection of computing nodes, or processes, connected over directed arcs representing communication paths. As a means of executing the algorithm from such a distributed system, a sequential program, named SIM, has been constructed which simulates the execution of the distributed program. This simulation can be used to study the operation of the distributed algorithm, and to obtain its performance data. This report explains the structure and operation of this program, and suggests applications.

A SIMULATOR FOR MESSAGE-BASED DISTRIBUTED SYSTEMS

BY

William Francis Quinlivan III, B.S., B.A.

REPORT

Presented to the Faculty of the Graduate School of  
The University of Texas at Austin  
in Partial Fulfillment  
of the Requirements  
for the Degree of

MASTER OF SCIENCE IN COMPUTER SCIENCE

THE UNIVERSITY OF TEXAS AT AUSTIN

August, 1981

## ACKNOWLEDGEMENTS

I would like to thank Dr. J. Misra and Dr. M. Chandy for their help and advice on this project, and suggestions that improved this report.

## TABLE OF CONTENTS

1.0	INTRODUCTION . . . . .	5
2.0	THE PROCESS . . . . .	10
2.1	The Process in CSP . . . . .	10
2.2	The Process in SIM . . . . .	13
2.3	Mapping CSP Programs into SIM Programs . . . . .	19
	2.3.1 Mapping a CSP Parallel Program Segment . . . . .	19
	2.3.2 Mapping CSP Processes to SIM Process- Procedures . . . . .	20
3.0	THE IMPLEMENTATION OF SIM . . . . .	28
3.1	SIM Data Structures . . . . .	28
3.2	Flow of Control in SIM . . . . .	32
3.3	Statistics Collected by SIM . . . . .	38
3.4	Debugging Considerations and Simulation Termination . . . . .	39
4.0	AN APPLICATION -- DISTRIBUTED SIMULATION . . . . .	43
4.1	Some Process Types for Distributed Simulation . . . . .	44
	4.1.1 The SOURCE Process Type . . . . .	45
	4.1.2 The SINK Process Type . . . . .	47
	4.1.3 The DELAY Process Type . . . . .	47
	4.1.4 The QUEUE20 Process Type . . . . .	50
	4.1.5 The FORK2 Process Type . . . . .	54
	4.1.6 The MERGE2 Process Type . . . . .	55
4.2	Deadlocks in Distributed Simulation . . . . .	56
4.3	Deadlock Resolution in Distributed Simulation . . . . .	57
5.0	SUGGESTIONS FOR FURTHER WORK . . . . .	61
6.0	SUMMARY . . . . .	63
	APPENDIX 1 TEXT OF SIM PROGRAM . . . . .	64
	APPENDIX 2 SOME TEST RESULTS . . . . .	97
	REFERENCES . . . . .	102
	VITA . . . . .	103

## 1.0 INTRODUCTION

A broad class of physical systems including queueing, communication, and computer networks can be modeled as a collection of computing nodes, or processes, connected over directed arcs representing communication paths. As a means of executing the algorithm from such a distributed system, I have constructed a sequential program, named SIM, which simulates the execution of the distributed program. The text of this program appears in appendix 1. This simulation can be used to study the operation of the distributed algorithm, and to obtain its performance data. This report explains the structure and operation of this program, and suggests applications.

The execution model of the computing node is based on Hoare's paper, "Communicating Sequential Processes" [HOA78]. In the language model (referred to as CSP) suggested in this paper, processes share no data. Instead, all interaction between processes is via messages passed between a sender and a receiving process, which name each other explicitly.

This model of parallel execution and communication is attractive because of the availability of low-cost processors

which can each be required to carry out the computation of a single process. These could be interconnected over communication lines so that the topology of a particular distributed algorithm could be embedded in the processor communication graph. From a program verification point of view, this model has merit because a process literally participates in every transaction that has any opportunity for an outside process to change its data - namely, communications. This simplifies proofs of programs because of the reduced chance of interaction between processes.

In order to study these algorithms, it is generally necessary to either have such an embedding architecture for execution, or to map the problem onto a simpler architecture. Such mapping leads to execution of a parallel algorithm on a single sequential machine. This is the approach taken in this program. SIM was developed and run on the University of Texas' DEC-10 computer, using the programming language, Pascal.

At the current stage of development, no source language processing is supported by SIM; its function is rather the run-time support for programs written in CSP-like languages. The notion of a process in SIM is a Pascal procedure which is a single thread running through local variable initialization, parallel execution with other processes (simulated), algorithm result reporting, and process termination.

In CSP, a process is always in one of three distinct states: executing, waiting for communication, or terminated. Ordinarily, a process alternates between executing, and communicating with other processes. At some point, the process may enter the terminated state after an execution phase. The communication consists of an indefinite wait for communication, followed by an instantaneous passing of a message. Such a message passing is loosely referred to as a message firing, or a port firing. This three state behavior is summarized in the state graph shown in fig 1.1.

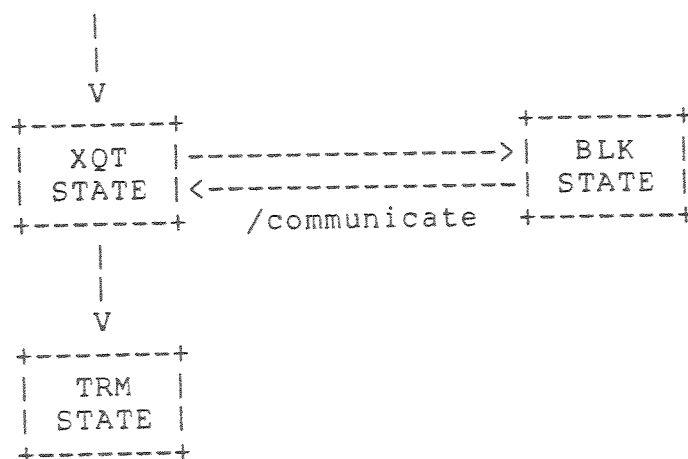


fig 1.1

The SIM implementation uses a central CONTROLLER which resumes the processes in the network, after they communicate, to execute their own part of the distributed algorithm. In the SIM implementation, the role of the process is played by a Pascal procedure which contains the procedural description of the process it emulates. Hence, the action of resuming a process after communication is accomplished by

calling the procedure which implements the process.

A central network-time is maintained which orders the execution and communication of the network, and permits collection of global performance statistics. As long as the individual processes never refer to this global network time, it remains a purely auxiliary variable, serving only to meter the time lost to processes while communicating, or waiting to communicate. Since one definition of a distributed system is one in which communication introduces non-negligible time delays, [LAM78] it seems beneficial to model such delays in a network simulation.

Hoare's paper made no mention of time dependent behavior, except to suggest that, as a fairness issue, process pairs awaiting communication with each other should not be delayed indefinitely often; this is appropriate for a paper making a preliminary language definition. However, to obtain any kind of performance statistics for an algorithm written in such a language, it is necessary to place bounds on wait times incurred by process pairs which are waiting to communicate with each other. In SIM, this bound is synonymous with the "time unit" in which time periods are measured. This is equivalent to saying that communication begins in the first time unit following the event that both partners of a communication pair become ready to communicate. The performance statistics derived from such an approach therefore



represent best-case times, which could only be realized in a physical system that similarly bounded mutual-wait times.

## 2.0 THE PROCESS

The process plays a central role in both CSP and SIM. The nature of the process in both CSP and SIM, and the mapping of CSP processes into the "process-procedures" of SIM is the subject of this section.

### 2.1 The Process in CSP

In the source text, a process in CSP consists of three items:

- . a name, which can either be a simple identifier, or an identifier followed by index limits, specifying an array of similar processes.

- . a local variable declaration part which defines the local data structures and specifies initial values.

- . a procedural description of the actions of the process. This program is specified in terms of six control structures: guarded commands, assignment commands, parallel commands, repetitive commands, alternative commands, and the I/O commands that Hoare maintains are primitive.

The reader is referred to the original CSP paper for a complete description of the process textual structure, and the semantics and a recommended syntax for the command types,

together with illustrative examples. An important characteristic of any I/O command that appears in the text of a process,  $P$ , is that it specifies the identities of the processes with which  $P$  will be eligible to communicate, when the command is encountered during  $P$ 's execution.

The CSP model of a process is a three state finite state machine, (FSM) which makes some number of transitions between executing and waiting for communication, followed optionally by a transition from the executing to the terminated state. Communication is treated as if it occurred instantaneously after an indeterminate period of waiting.

The entire address space of the process is local to the process. The process communicates with other processes by naming them explicitly in an input or output (I/O) command, and subsequently having that communication selected for firing by scheduler whose selection policy is arbitrary. The process may wait simultaneously for communication with many other processes, but the scheduler never chooses more than one message for firing per process. As part of the communication, the parties involved are informed of the identity of the processes with whom they communicated.

When a message fires, the effect is the same as an assignment statement where the expression yielding the value assigned is evaluated in the sending process, and selection of the receiving variable is performed in the receiving

process. As in an assignment statement in a strongly typed language, the type of the value sent must match the type of the receiving variable.

Although the process is never involved in communication with more than one other process at a time, parallelism is introduced in the I/O handling because a process may wait for communication with more than one process at a time. This arises in two ways. The first is that a single I/O statement may name an entire array of processes as eligible for communication, and the second is that the process tries to execute some I/O commands in parallel. Actually, only one of the commands will be selected for execution. Three separate statement types give rise to this latter type of parallelism: the parallel command, the repetitive command, and the alternative command.

The repetitive and alternative commands are borrowed from similar commands suggested by Dijkstra in [DIJ77]. The repetitive command:

```
* [ P1 ? y --> x1 := x1 + 1 ;
  □ P2 ? y --> x2 := x2 + 1 ]
```

causes the process to wait for input from either P1 or P2, and depending on which one does send a message, either x1 or x2 is incremented, and the statement repeats. Here, only one message fires at a time, but the process always waits in parallel. The alternative command:

```
[ P1 ? y --> x1 := x1 + 1 ;
  [] P2 ? y --> x2 := x2 + 1 ]
```

is similar to the repetitive command except that once either process sends a value for  $y$ , the respective increment of  $x_1$  or  $x_2$  is performed, and the statement terminates. The parallel command:

```
[ P1 ! x || P2 ! x ]
```

forces the containing process to wait to send the value of  $x$  to both  $P_1$  and  $P_2$  in any order. Hence, before either message fires, the process is waiting in parallel.

This then is the kernel of process functionality that is required in an implementation: the basic structure and execution behavior of the process as described above, including parallel waiting for communication. Hoare has further requirements such as termination of a repetitive command when all of the named-for-communication processes have terminated, but this is not supported in the implementation. Hoare also disallowed the presence of output statements (i.e. message sending) in the guards of the repetitive and alternative commands, but the SIM program does not make this restriction. SIM can do this because the states of all processes are available in a single memory to which the scheduler has instantaneous access.

## 2.2 The Process in SIM

In SIM, the process appears as a Pascal procedure

describing not a named process or an array of such processes, but rather a generic process type. These procedures are all compiled and mapped with the SIM code. In a particular execution of SIM, a given process type may be instantiated zero, one, or more times. Most SIM runs will exercise instances of more than one process kind. All process instantiations are assigned a unique identifier, known locally as ID, that is an input parameter in the calls to the procedure containing the process definition. Such calls correspond to a resumption of the process being modelled, and are typically handled by means of a call to the RESUME procedure.

All of the local data of all the processes is contained in the array OWN. This array is indexed by the values of the process identifiers. In an actual distributed system, this would be all of the state information necessary. However, since SIM uses a central scheduler, information on the process' current state is kept in a separate array LPS, also indexed by the process identifiers. The structure of the LPS entries is the same for all of the processes, but the OWN entry for a process is determined by the kind of process that owns it.

In SIM, communication lines are known as ports, and each port has a fixed sender, receiver, and type of message. Since it would be very restrictive to require the processes to know the identity of all the other processes at compile

time, unique port names, instead of process names are given in the I/O commands which give rise to communication. The unique port names exist as the values of local variables in a process. Hence, a process may have an output command which specifies a variable OUTPORT as the port over which a message should be sent. The particular port over which this command specified message passing would depend on the value of this variable at a given moment.

During the initial phase of a SIM execution, all such local port names are bound to particular ports; also, the sending and receiving process identifiers, and the message type allowed are bound to the particular port. This latter information is kept in the array PORTS, which is indexed by the port identifiers. If all such binding is made consistent with some communication graph,  $G$ , then the topology of  $G$  is reflected in the connectivity of the processes and ports. This scheme effectively isolates the procedural specification of the process from the particular topology of the network to be simulated. In other words, the process semantics, but not the network topology is bound at compile time; the topology is specified at run time, and may vary between executions of SIM. Use of port names instead of process names was also suggested by Hoare in the original paper.

The actions of the process may be viewed as responding to a set of significant events by deciding what the next

event should be, and updating its local variables. The events here are the initial invocation of the process, the sending of a particular output, and the receiving of a particular input. In SIM, each event is associated with a label preceeding the code which responds to the event. During any particular execution, the process specifies the set of next events, and then returns to the central CONTROLLER. To the process, this is equivalent to invoking an "oracle" which selects one of the events and resumes the process at the label associated with that event.

The labels then are the points at which the process may potentially be resumed. As such, they may be viewed as the addresses of a program of the form:

ADDRESS	ACTION
label1	respond to event corresponding to label1
label2	respond to event corresponding to label2
.	.
.	.

A variable that ranges over these labels is known as a meta-program counter, abbreviated MPC. Each resumption of a process gives as one of the parameters the particular MPC value at which execution should continue.

A special procedure, named PARWAIT ( for parallel wait ) is provided for the purpose of telling the CONTROLLER



about the events for which the process will next wait. To wait for communication over a set of ports, one PARWAIT call is made for each port. If a process returns to the controller without invoking PARWAIT, it is marked "terminated", and dropped from further simulation.

PARWAIT has three parameters: the port name over which communication may take place, the identity of the requesting process, and the MPC value at which control should resume if the communication described in the current PARWAIT invocation turns out to be the next communication that the calling process is involved in.

Several of the MPC values are reserved for events which are common to all of the process kinds. MPC = 0 selects code activated at process creation. Here local variables are initialized which do not change throughout the entire life of the process. These include things like the local variables that name ports, and parameterizing variables, such as the buffer size of a bounded queue, or the delay parameters for a process which simulates a time delay process. MPC = 1 corresponds to initial process activation. Here local variables are given values which may change during execution. Activation and creation have been kept separate to allow several activations of a process that only needs to be created once. During the activation, the process decides and informs the CONTROLLER which events will be the first

ones allowed by the process. MPC = 1000 is reserved for code which allows reporting of results after the simulation has terminated, and do things like close output files if this is necessary. This reporting is strictly of things known locally; the overall performance reporting for the entire network is handled by the central CONTROLLER. Two other MPC values, 1010, and 1020 have dedicated functions relating to a deadlock recovery technique discussed in section 4.3.

All other MPC values are provided to specify the code to respond to events. The only other restriction on these MPC values is that they be unique within a particular process.

Each process instance has a message buffer that is used for communication. This buffer is part of the process' activation record in the LPS array. When process I wants to send a message over port X, where MPC = Z corresponds to the code segment that responds to this communication, it first writes the message into its message buffer, and then invokes PARWAIT with the appropriate parameters: port = X, requester = I, and MPC = Z.

If process I were to be the receiver, instead of the sender of this message, the message is not written into the buffer (since I does not yet know its contents), but the PARWAIT call is the same. When I is resumed after the message is passed, it will find the message in its buffer.

Thus, processes are supported with all of the necessary functionality. As long as a user of this program were content to code the process descriptions in the necessary form for execution, SIM would be usable as described. Typically, however, the user will wish to write the process code in a language that allows concentration on the semantics of the process, and not worry about MPC values and PARWAIT parameters. The next section gives the outline of a translation procedure to convert the CSP commands into code compatible with SIM.

### 2.3 Mapping CSP Programs into SIM Programs

The process of converting a program written in CSP into the equivalent program suitable for execution on SIM consists of translation of the main program, and all of the individual processes.

#### 2.3.1 Mapping a CSP Parallel Program Segment

The typical form of a parallel program segment in CSP is:

$$[ P_1 \parallel p_2 \parallel \dots \parallel P_n ]$$

where the  $P_1 \dots P_n$  are the processes defined in the program. This just says that they all should run in parallel. This structure is the implicit execution model for a SIM main program. Main program structures that depart from this form require the definition of a special process that performs the

function desired in the main program.

Such a "main program process" would run just like the other processes, but would be the only one that performed any meaningful computing before becoming blocked. The normal call-return structure familiar in sequential programs is easily implemented by having the processes pass special "call", and "return" messages, where the content of the message is made up of calling and return parameters, respectively.

### 2.3.2 Mapping CSP Processes to SIM Process-Procedures

In order to convert a CSP process into the equivalent SIM process-procedure, it is necessary to map the three parts of the CSP process. These are the name, the local variable declaration, and the programming language statements that state the execution behavior of the process.

The name conversion consists of placing the process name into the scalar list defining the type PROCKIND, and into the PRPROCKIND procedure that prints out the name of a process in SIM. The name must also appear in the RESUME procedure that invokes a process after it communicates.

Mapping the local variable declaration section means placing the structural definition into the definition of the data type, OWNDATA. The structure of a variable of this type is determined by the kind of process that owns it; this

information is kept in the field OPROCKIND. In general, many data templates will appear in the definition of OWNDATA, but the inclusion of a particular local data type is straightforward, as can be verified from the example in the appendix. This inclusion makes the definition needed by the process. To implement the initialization of this data, the code must appear in the process initialization part of the definition of the process. By convention, this function occurs at MPC = 1.

An implicit data type exists in the form of the messages passed by the process. In CSP, this is handled by insisting that the type of all messages match the type of the variables which receive their content. This is not supported in SIM. The messages passed by a process must have a defined structure. The structure for all message types is made in the definition of the SIM data type, MESSAGE. The actual structure of a message is user defined, and a particular network simulation may contain messages of more than one type, e.g. DATA and ACKNOWLEDGE. In this case, this type has a variable structure determined by the value of its MSGKIND field. Hence all message kinds in a simulation must have their name listed in the definition of the scalar type, MSGKIND, and the corresponding message structure must appear in the MESSAGE definition.

The information about the graph topology that is

explicit in a CSP process is handled in two stages in SIM. At compile time, the ports are assigned local variable names in the processes. Then, at network specification time during the actual running of the program, the values of these variables are bound to specific ports in the network.

The translation of the parts of the process described above is fairly mechanical. This leaves the procedural description of the actions of the CSP process as the major challenge in converting a CSP process into its SIM equivalent.

The only parallelism or nondeterminism supported by SIM is the parallel waiting for communication. This is "parallel" because the process waits for more than one process at a time, even though the waiting will be terminated whenever any one of the named ports does actually fire. The nondeterministic element is introduced because the process does not know in advance which of its potential communications will be the one to fire.

CSP allows other forms of local parallelism in commands like:

```
[ x := x + 1 || y := y + 1 ]    ...and...
[ guard1 --> x := x + 1 ;
  □ guard2 --> y := y + 1      ]
```

where both guards are true. An implementation for this

requires a local scheduler to evaluate guards and select eligible actions. This low-level nondeterminism is not supported by SIM, so a translator would have to generate code to arbitrarily select one of the assignments in the first example to perform, and to select one of the guards in the second to evaluate first, and if found to be true, perform the indicated action.

Hence, the main difficulty in translating CSP processes into equivalent SIM Pascal procedures is to translate the parallel, repetitive and alternative commands in such a way that the parallel waiting for communication is preserved. These will be discussed in turn.

The general form of the repetitive command is:

```
* [ <boolean expression 1> ; <I/O command 1> -->
    <action 1> ;
  □ <boolean expression 2> ; <I/O command 2> -->
    <action 2> ;
    .
    .
  □ <boolean expression n> ; <I/O command n> -->
    <action n> ]
```

Here, the process will wait for any of the I/O commands which have the corresponding boolean expression true. Once one of them fires, perform the corresponding action. Since the set of awaited lines must be specified

when the process returns to the controller, it is necessary for the process to evaluate all of the boolean expressions, and request parallel waiting for the corresponding I/O command for each one that is true. This also requires that a condition, once evaluated to be true, must not be subsequently falsified, e.g. through a side effect of evaluating a subsequent expression. Consider DECIDENEXT to be a shorthand notation for the Pascal statements:

```

if not (<boolean expression 1> or
        <boolean expression 2> or
        .
        .
        <boolean expression n>)
    then go to L1
else for i := 1 to n do
    if <boolean condition i>
        then PARWAIT(porti,id,mpci);

```

Here L1 is the label of the next statement after the statements associated with this repetitive command, and porti is the port named in the i-th I/O command, and mpci is the MPC label associated with the code which should follow I/O command i. The equivalent SIM coding for this command is:



```

                DECIDENEXT;
mpc1:   begin <action 1> ; DECIDENEXT ; end;
mpc2:   begin <action 2> ; DECIDENEXT ; end;
        .           .           .
        .           .           .
mpcn:   begin <action n> ; DECIDENEXT ; end;
L1  :   <rest of program> ;

```

With this interpretation, the command would terminate when all of the boolean conditions were false.

The general form of the alternative command is:

```

[ <boolean expression 1> ; <I/O command 1> -->
    <action 1> ;
□ <boolean expression 2> ; <I/O command 2> -->
    <action 2> ;
    .           .
    .           .
□ <boolean expression n> ; <I/O command n> -->
    <action n> ]

```

This CSP command is supposed to begin waiting in parallel for all I/O statements with a true guard. If none are true, the statement should fail. Unlike the repetitive command, the action is only performed once. SIM has no feature corresponding to having a statement fail. Hence, if it is desirable to have the process be terminated in response

to execution of this statement with all guards false, then the DECIDENEXT procedure defined above should simply return to the central CONTROLLER instead of the statement fragment:

```
go to L
```

Since the process will have returned without specifying any I/O, the CONTROLLER will mark the process as terminated. Otherwise, the DECIDENEXT can be used as is, and the semantics of the repetitive command become: if there is no true boolean condition, then just go on to the next statement, otherwise, begin waiting in parallel for an I/O statement with a corresponding true guard; after one of them has fired, execute the associated action.

With this possible modification to the DECIDENEXT procedure, the code for a SIM interpretation of the repetitive statement becomes:

```

DECIDENEXT;

mpc1:  begin <action 1> ; go to L1 ; end;
mpc2:  begin <action 2> ; go to L1 ; end;
      .           .           .
      .           .           .
mpcn:  begin <action n> ; go to L1 ; end;
L1 :   < rest of program >
```

The general form of a parallel command that may result in parallel waiting for I/O is:

```
[ <I/O command 1> || <I/O command 2> || ...
      ... || <I/O command n> ]
```

A possible means of handling this is to establish an n-element boolean array, DONE, and convert this parallel command to the equivalent CSP statements:

```
{ set all DONE[i] to false }
* [ DONE[i] --> DONE[i] := false ]
* [ not DONE[1] ; <I/O command 1> --> DONE[1] := TRUE;
      not DONE[2] ; <I/O command 2> --> DONE[2] := TRUE;
      . . . .
      . . . .
      not DONE[n] ; <I/O command n> --> DONE[n] := TRUE; ]
```

The resulting repetitive command would then be translated according to the rules described above into the equivalent SIM statements.

### 3.0 THE IMPLEMENTATION OF SIM

This chapter will describe in considerable detail the data structures, flow-of-control, termination conditions, statistics collected, and debugging aids of the SIM program.

#### 3.1 SIM Data Structures

SIM allows execution of networks consisting of several varieties, or types, of processes. To qualify the kind of particular processes, a scalar type, PROCKIND, is defined which ranges over all possible process kinds. The current execution state of a process is contained in entries of two arrays which have already been mentioned. These are the OWN and LPS arrays, which contain the local variables and activation records, respectively, of the processes. Both of these arrays are indexed by the unique process identifiers.

The element of the LPS array is of a structured record type ACTREC which contains the following information:

. the type and instance of the process. Together, these two fields uniquely identify the process in a way that has mnemonic value to the user. As an example, they could specify a [server,6] or [queue,2] in a network where process types server, and queue are supported.

. the current state and the next state of the process. The current state is one of executing (XQT), blocked (BLK), communicating (CMN), or terminated (TRM). The next state can only be BLK or TRM, and is used for a process to inform the central CONTROLLER of the state it will enter following its current phase of execution.

. the time-left field is used as a count down timer. When this reaches zero, it signals the end of the current state, either CMN or XQT. Initial values for this timer come from the process itself for execution time, and from the port records for communication delays ( see description of PORTS, below).

. three accumulators record the total process time spent in the states XQT, CMN, and BLK. In case a process has terminated, a separate field contains the network time at which this occurred.

. one buffer each is provided to hold an MPC, a port identifier, and a message. These are used in communicating the actual message passed, the MPC value at which a process should resume, and the identity of the port over which a message has just been passed, when the central CONTROLLER resumes the process after a message firing.

In short, the LPS entries contain everything the CONTROLLER needs to know to change the state of the process, handle the details of message passing, and collect process-specific statistics on its performance.

The OWN array also has elements of a structured type, called OWNDATA. This record type contains a field of the PROCKIND type which tells the type of the process with process identifier equal to this element's index in the array. The rest of the OWNDATA structure is variable, and depends only on the kind of process that owns it. Typically, this contains the local variables that name ports incident on the process which are referenced in I/O statements. This may also contain a local variable for local time, if this is desired, and any data for holding locally-maintained statistics, such as queue length information that might be kept by a process implementing a queue. Basically, this array must hold all of the process' data which must survive between calls to the procedure that implements the process. Strictly temporary variables can be handled by the regular Pascal local variables.

Since all of the field names of Pascal record types with variable fields must be unique, some care is required to prevent collisions in the name space of the variants for different process kinds. A practical solution which has been exploited in the example SIM program in the appendix, is to prefix all field names with a two- or three-letter combination which is suggestive of, and unique to the owning process kind.

Port identifiers in the simulated network are also

unique, and serve to index the PORTS array. PORTS[j] contains all of the information held about the port with unique identifier, j. The element type of this array, called PORTREC, contains the following information:

- . the unique process identifier of the sending and receiving process, which remains constant throughout an execution of SIM.

- . four boolean<sup>8</sup> fields tell whether the sender and receiver are ready and/or eligible for communication. Eligible means the process named this port for communication during its last execution phase by means of PARWAIT calls, and ready means that the process is currently blocked. Actual message passing on this port will never take place unless the sender and receiver are both ready and eligible.

- . the MPC values of both the sender and receiver which should be returned to when and if this port fires. Since the processes may wait on many messages, each with a separate MPC return point, a separate MPC must be kept for every possible communication.

- . the amount of time that communication over this port delays both the sender and receiver. The interesting cases seem to be when the send-time is less than or equal to the receive time. The send time is related to the size of the message, and the communication rate. The difference between send and receive time corresponds to communication delay. Hence message propagation delays can be explicitly simulated,

e.g. in a simulation of communication over satellite links.

- . a count of the total number of messages that have been sent over this port since the beginning of the run.

- . a buffer that holds the message, if any, that was specified most recently by the sender process, just prior to the PARWAIT call which marked the sender as eligible for communication over this port. The MESSAGE record type, contains in addition to the actual message, a message type. The type in the message associated with this port binds the port to a single message type. Attempts by a process to send a message with a type different from that recorded in the port's message buffer results in an error message.

The global variables maintained by the central controller include the global time, accumulators for time spent by processes in each of the states XQT, BLK, TRM, and CMN, and a count of the total number of messages sent. Several trace variables are used to turn the run-time trace off and on. This trace is useful for debugging both the process procedures and the SIM program itself, and will be discussed below. Other variables keep termination thresholds for time and message counts, and deadlock occurrences.

### 3.2 Flow of Control in SIM

The sequence of events in a run of SIM consists of the following steps:

- . an interactive session which solicits inputs from the



operator about the connectivity of the network to be simulated. This information could optionally be read from a file. This session obtains information in two categories:

- . bind each unique process identifier to a particular process kind and instance by instantiating the defined process kinds zero, one, or more times.

- . for each of the ports in the network, bind the sender and receiver processes; this implicitly specifies the topology of the network.

- . simulate the execution of the network, collecting statistics as progress is made.

- . print the statistics gathered in the last step.

The execution model of SIM during the simulation of the target network is the familiar operating systems concept of multiprogramming, where a number of tasks (processes) that are not blocked (whose state is XQT) are each allocated one time quantum of compute time in round-robin fashion. At the limit where the quantum size becomes zero, this becomes an instance of processor sharing, and an external observer sees all executing processes making steady progress. During each time unit in a true distributed system, the volume of processing performed is the product of one time unit and the number of processes that are executing at that time. This of course assumes that all processes that are executing accomplish the same amount of processing in a time unit. Hence, it seems reasonable to assert that one time unit of simulated time has

passed every time that the central CONTROLLER has made one pass through the list of executable processes; equivalently, this grants one time quantum to each ready process each time the network time is incremented.

An important point here is that the process is not actually resumed during every time unit in which it is in the XQT state. Processes are only resumed when they make a transition into the XQT state. The actual amount of time required for the process to complete its actual computation is immaterial. All that matters is that the process inform the CONTROLLER of the simulated time required to complete the execution being performed. This is used for performance purposes only. In cases where only the execution of an algorithm, as opposed to its simulation, is important, this time value can be ignored altogether.

Hence, a process simulating a time delay of 25 time units need only tell the CONTROLLER that it will require 25 time units. The time required to return this result does not matter. The effect of resumption, from the CONTROLLER'S point of view, is that an oracle is invoked which provides this time estimate, and the set of ports over which this process will be waiting to communicate during the block state which will follow the current execution state. In the case that no ports will be awaited, the action simulated is that the process is getting ready to terminate. Furthermore, the

oracle can provide the contents of any message that the process will wait to send. Lastly, the oracle can assert that the process has updated its local variables as if the execution had already taken place, so for the rest of the simulated time period, nothing at all needs to be computed. These are exactly the effects of resuming a process, and as a result, the CONTROLLER can merely simulate the passage of time rather than actually giving a time quantum to the process every time the network time is incremented.

The estimate of the total compute time is returned by the processes in the XQTTIME field of their LPS entry. The central CONTROLLER copies this value into the TIMELEFT field; at every subsequent time unit, this value is decremented until it reaches zero. At this point, the process enters the next state, always one of terminated, or waiting for communication.

Similarly, any process in the CMN state has its time left field decremented, and the process reenters the XQT state as soon as the count reaches zero. Here the total time is provided not by the process, but rather by the user during system specification early in a run of the program, who presumably knows the intrinsic characteristics of the communication port, and the size of the messages that pass over the port.

Processes in the BLK state have no such bound on the

amount of time they will remain in this state. They simply wait until one or more of the processes that they are waiting to communicate with become ready for message passing. If a process is found to be waiting only for processes that have terminated, it is marked terminated.

The above mentioned processing takes place in the procedure TICK, which charges each process according to its state, and initiates those state changes which occur at predictable times. This routine is called once each time unit.

The procedure PASSMESSAGES is also called once every time unit, and its action is to fire some of the ports which have both sender and receiver both ready and eligible for communication over the same port. The current implementation uses a fair scheduler which guarantees that no such ready-to-fire port will be passed over more times than there are ports. This is accomplished by keeping a variable, called netfair which ranges over all the port identifiers, favoring them for a message pass if their sender and receiver are prepared to communicate, and incrementing to the next port at each time unit.

The scheduler is also deterministic in the sense that if two processes are both ready and eligible for a message pass, over say port J, then at least one of them will be involved in some message passing, either over port J, or

over one of the other ports that is incident on the process. The only reason port J would not be selected for firing is that the scheduler already selected a different port incident on the port J sender or receiver. In this case, port J could not be fired because that would mean some process was involved in more than one message firing.

The scheduler has been coded as a separate procedure so that a user may provide a different algorithm for selecting among the ready ports to decide which, if any, of them should actually fire.

When a message does fire, the contents which have been buffered in the port's records are copied to the receiver process' message buffer. Both the sender and receiver are placed into the CMN state to wait out the simulated communication time. This takes place in the procedure, FIREPORT.

The simulation view of the processes is now complete. Compare the state graph for the SIM process-procedures in fig 3.2 with the state graph that describes only the logical view of a process' states, fig 1.1. The nodes correspond to the states, and the arcs are the transitions. The procedures named for the arcs are those invoked to make the transition for the processes. Not shown in the state graph is a self loop from each of the states to itself, taken implicitly in the TICK procedure. The lengths of time spent in each state are determined as follows:

. time in the CMN state is determined by the properties of the port handling the message, and the length of the message. This information is held in the PORTS array.

. time spent in the BLK state is a function of the readiness-to-communicate of the processes named for potential communication by a process.

. time spent in the XQT state is determined by the process itself.

. time spent in the TRM state is from the time the state is entered until the end of the simulation.

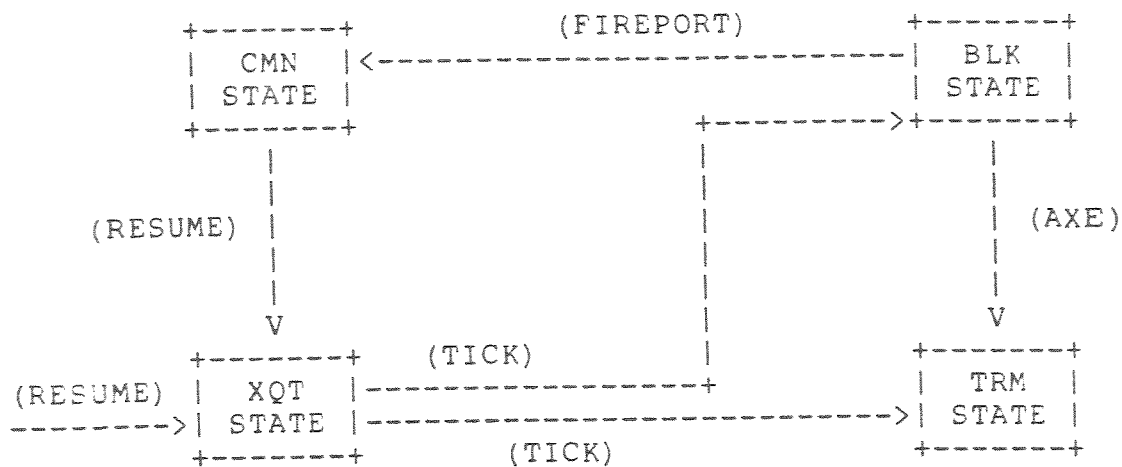


fig 3.2

### 3.3 Statistics collected by SIM

As part of the activation record of a process, SIM keeps a running sum of the total amount of time that the process is in each of its states, and the network time at which the process terminated, if this has occurred. For each

port, SIM records the total number of messages passed along the port. SIM also keeps the total elapsed network time.

The procedure PRINTSTATISTICS is used at the end of a run to print out the statistics collected during the run. This routine can be extended to also print out any other statistics which may be deemed necessary for implementation of a target network.

It is appropriate for the processes themselves to keep track of some types of performance statistics, and to print out this information at the end of the simulation. For this latter purpose, the MPC label 1000 is provided. Processes are resumed at this label one last time when the simulation is complete for the purpose of printing this data. As an example, queue length distribution histograms could be printed out by processes implementing a waiting queue.

#### 3.4 Debugging Considerations and Simulation Termination

In order to facilitate implementation of SIM and the various target networks that may be developed using this program, provisions have been made to force SIM to produce various amounts of auxiliary, or trace, information while it is simulating the network. This output is controlled by the values of seven "trace" variables which the user is able to change at critical points during the run. Each of these variables controls the output within a particular area. The

output produced is none if the variables are set to zero, and generally more output is produced the larger the variables are assigned. The trace variable name and area of execution trace controlled by the variable is summarized in the following table:

<u>Variable</u>	<u>Trace Area Controlled</u>
SIMTRACE	General execution of SIM
NETXQTTRACE	Processes whose state is XQT
NETBLKTRACE	Processes whose state is BLK
NETCMNTRACE	Processes whose state is CMN
NETTRMTRACE	Processes whose state is TRM
NETDEADTRACE	Network deadlocks
TARGETTRACE	Operation of target network

Discussion of network deadlocks is delayed until section four.

During SIM execution, the user provides execution parameters in response to specific questions presented by the program. Two of these are the time and message limit. When either of these user-supplied limits is exceeded, rather than automatically terminating, SIM gives the operator a chance to set new limits in order to continue the simulation. The program will terminate immediately if the operator does not increase the limit which was exceeded.

The time and message limits, and the trace variables



mentioned above are all solicited together by a procedure named SETTRACE. Accordingly, by judicious choice of the limits and the trace variables, the operator can select particular trace options between particular times or message counts. For example, if a target network has been observed to blow up near network time = 100, then the operator may leave the trace values low between times zero through 95, and then, when the time limit of 95 is exceeded, increase the time limit to say 105, and increase the trace variables so that a large volume of trace information about the target network will be printed out between the current time and the new time limit.

This is the exact use intended for the trace variable TARGETTRACE. Of course, since much of the implementation of the output from the target networks is necessarily handled by the processes themselves, it is necessary for the code to be included in the processes so that a higher value of TARGETTRACE does in fact produce more output from the network. This is evident in the example processes in the SIM listing in the appendix.

As an aid in including such code in the individual processes, P, of a new type, the following routines are provided:

<u>Routine</u>	<u>Prints Out</u>
PRSIGNATURE(P)	P's identity, kind and instance
SHOWOWN(P)	P's local variables
SHOWPROCESS(P)	P's activation records from <u>LPS</u>
SHOWMSG(M)	Message M's contents
SHOWNETWORK	Port - Process connectivity graph
TRACELP(P)	Useful process information
DUMPTTY	All global data, <u>PORTS</u> , <u>LPS</u> . etc.

#### 4.0 AN APPLICATION -- DISTRIBUTED SIMULATION

Distributed Simulation [CHA81] has been proposed as a fruitful applications area for computer architectures exhibiting parallel processing capabilities. Here, the parallelism in a computation mirrors the parallelism inherent in a set of physical entities to be modelled. Interactions between the entities are simulated by messages passed between the processes that simulate the entities.

This is an instance of a message-based distributed program that is amenable to simulation on the SIM program. Here, the communication lines are modelled as ports, and the roles of the processing nodes are filled by SIM procedures.

In the paper presenting this concept, Chandy and Misra define the notions of process-times and line times. These are the times in the entity-level through which a process or port has been simulated. A key correctness requirement is that the events in a process or a port must be totally ordered in time, even where the time is in a "logical" sense, i.e. measured by a counter, as opposed to time measured by a clock. A port exhibiting this property is said to be "chronological". Lamport [LAM78] also treats the sub-

ject of local and global clocks, and describes the need for an extension of the partial ordering imposed by the collection of individual local clocks, to a total ordering of the events within the network if they are to be executed in a consistent manner. This is exactly the function of SIM.

In the physical system being simulated, the analog of line and process times are all trivially equal because time is always constant across the entire network. This is not so at the logical level of the simulation, and in general processes and ports may have radically different local times. This corresponds to the situation where some parts of the simulation are running far ahead of others, and the possibility of this asynchrony is the motivation for Distributed Simulation in the first place. The asynchrony represents parallelism in the simulation that is not possible in the entity level network.

In order to be able to simulate an interesting set of networks, it is necessary to define the individual process types which will comprise these networks. The next section will define six different types and explain the semantics and waiting rules for each of the types.

#### 4.1 Some Process Types for Distributed Simulation

A very rich class of networks can be simulated by using a small set of process types. The processes construct-

ed for this application are as follows:

- . the SOURCE type which generates "jobs"
- . the SINK type which consumes jobs
- . a DELAY type which makes a job wait for some time
- . a FORK2 type which accepts jobs on an input line, makes a decision, and sends the job out over whichever of its two output lines is selected by the decision.
- . a MERGE2 accepts jobs over two input lines and sends them out over a single output port while guaranteeing the chronological condition on the output line times
- . a QUEUE20 models a waiting queue with a maximum size of 20 jobs.

In the following sections, each of these process types will be explained. The QUEUE20 process type is described in greater detail so that the relationship between its name, message types, local variable structure, procedural definition, and the facilities of SIM can be better understood.

#### 4.1.1 The SOURCE Process Type

A process of type SOURCE is connected to the rest of the network via a single port, known to the process as the value of the local variable, SOOUTPORT. When the process is created, i.e. resumed at MPC = 0, this variable is bound to one of the ports in the network being simulated.

A SOURCE type process sends messages out over its output port with an interdeparture time determined by the code implementing the process. One application could read the departure times from a file. This might be used to perform trace-driven simulations, where the input to a simulation was a set of actual values measured in a real physical system. The SOURCE process listed in the appendix uses as the interdeparture time the sum of a constant, contained in the local variable SOCON, and a "discretized" exponential random variable, with parameter SOMU generated using the RANDOM function available on the DEC-10.

When the process is created during a run, the operator supplies values for the variables SOMU and SOCON, and binds the output port to a specific port in the target network by specifying its unique port identifier as the value of the variable SOOUTPORT. Such a SOURCE "node" also keeps a local time, in the local variable, SOTIME, which is always the time that the last message was sent out. The mechanism for sending a message after the passage of some delay, say  $l_0$  time units, is to create a new message with a timestamp of SOTIME +  $l_0$ , and place the message into the process' message buffer. This value also replaces the current value of SOTIME. Then the interdeparture time,  $l_0$ , is also used as the XQTTIME returned to the central CONTROLLER. This tells the controller that the process will be executing for  $l_0$  time units before it tries to send out any message over its port.

Finally, PARWAIT is invoked to indicate that the process will eventually wait to send a message over the port whose identifier is the value of SOOUTPORT.

After this port does fire, this simple program is merely repeated; each time a positive delay is computed, and used as the time lapse until the next message will depart. Since the successive message times are increasing, the chronological condition is met.

#### 4.1.2 The SINK Process Type

The SINK process type has a single input port, known locally as SIINPORT. The SINK process always waits for input, and when a message is received, it is disposed of. In the current implementation, nothing is done with this message. Other uses of the SINK process might have the SINK process record some data from the message, or possibly print out a message as it is received. A SINK has no output ports, and so trivially guarantees monotonicity on its output ports.

#### 4.1.3 The DELAY Process Type

A DELAY process waits initially for input over an input port, known as DINPORT, and after a message is received, computes a time delay in the same way as the process type SOURCE. After this time delay is waited out, the process waits to send out the same message over a port named DOOUTPORT. Waiting out the delay is accomplished by leaving

the process in the XQT state for the computed delay period. After the message is sent out, the process again waits for input, and the cycle repeats. There is no internal buffering in a DELAY node for processes awaiting service. Once a message enters the node, it begins its wait period. However, it is buffered in the process while the process waits to send it out.

The particular semantics of this DELAY process models a single server in the sense that the delay time is only "charged" to a single message at a time. This is illustrated in the following example:

Suppose that some source of messages will try to send three messages to a DELAY node, which is initially ready to receive, with successive timestamps of 10, 20 and 80. Since the processes do not know how long they have been waiting to pass the message, all that can be asserted is that the arrivals to the DELAY node were at times at least as large as the three timestamps. Suppose that the delay values used for these three messages are 50, 5, and 30 respectively. The first message will leave the DELAY node at some time after time 60, because it arrived at time 10 or later, and had to wait out a delay of 50 time units. The second message will leave at time 65, or later. The reason for this is that in this type of DELAY node, a message does not enter and begin its waiting period until previous messages have left



the node. The last message leaves the node at some time at least as great as 110, since the node was "free" to begin its delay (i.e. begin its processing) when it arrived at time 80, or later, and the delay used was 30 time units. This behavior is summarized in the following table:

Message	Arrival Time	Delay Time	Departure Time
msg 1	$\geq 10$	50	$\geq 60$
msg 2	$\geq 20$	5	$\geq 65$
msg 3	$\geq 80$	30	$\geq 110$

This behavior is guaranteed by beginning the time delay for a message,  $M$ , at local time, DTIME, which is maintained as the maximum of the timestamp  $M$  had when it arrived, and that of the last message sent out. DTIME corresponds to local wall-clock time inside of a DELAY node, with the unusual property that the time displayed is always a lower bound of the real time. According to this clock, the arrivals take place at times 10, 60, and 80.

A separate time accumulator, DSUMPTIME, keeps the sum of all the computed delays. Hence, DSUMPTIME is the total "busy" time of the node. The value DTIME - DSUMPTIME is a measure of the time lost because of starvation to the node. The ratio of DSUMPTIME to DTIME is roughly the utilization of the node.

If messages had fields representing their sizes,

then the DELAY node could be constructed to return some function of the message size as its estimate of compute time. This provides a very easy way of simulating an algorithm whose time complexity is a function of problem size, e.g.  $O(\text{size}^2)$ .

#### 4.1.4 The QUEUE20 Process Type

The process that simulates a FIFO waiting queue in this implementation is named QUEUE20 because the procedure that implements the process is compiled with a hard limit of 20 as the maximum queue size that can be accommodated. The actual run time queue size maximum, called Q20MAX, is constrained to fall between one and 20, inclusive, and is solicited from the operator at process-creation.

The following quasi-CSP program is presented to demonstrate the translation of a program from the original CSP into the procedural description required by SIM.

```

QUEUE20 ::

    { LOCAL VARIABLE DECLARATION AND INITIALIZATION }

    Q20INPORT,Q20OUTPORT:PORTNAME;
    Q20INPTR,Q20OUTPTR:INTEGER;
    Q20BUFFER:ARRAY[0..19]OF INTEGER;

    { initialization of local variables --@ MPC = 1 }
    Q20INPTR:=1; Q20OUTPTR:=1;

    { PROCEDURAL DESCRIPTION }

    * [
      { IF WE HAVE ROOM IN THE BUFFER }
      Q20INPTR<Q20OUTPTR+20;
      { THEN TRY TO RECEIVE A MESSAGE }
      Q20INPORT?Q20BUFFER[Q20INPTR MOD 20] -->
      { WHEN RECEIVED, INCR INPUT PTR --@ MPC = 2 }
      Q20INPTR:=Q20INPTR+1;
      { AT THE SAME TIME, }

      [
        { IF WE HAVE A MESSAGE TO SEND, }
        Q20INPTR>Q20OUTPTR;
        { TRY TO SEND IT OUT }
        Q20OUTPORT!Q20BUFFER[Q20OUTPTR MOD 20] -->
        { WHEN SENT, INCR OUTPUT PTR --@ MPC = 3 }
        Q20OUTPTR:=Q20OUTPTR+1
      ]
    ]

```

Here it is assumed that the run-time queue size is always 20, and that naming a port for communication is equivalent to naming the process tied to the other end of the port. Q20INPORT is the name of the input port, and Q20OUTPORT is the name of the output port for this process. These are assumed to be bound to actual port numbers elsewhere.

The behavior of a QUEUE20 process is to initialize the local variables, and then any time it has a message to send out, it waits to send the message. At the same time,

whenever there is room in the queue, the process waits to receive another message.

In the SIM implementation, the name of the process kind, `QUEUE20`, is bound at SIM compile time and is distributed throughout the program. The name appears in the list of allowable values for the scalar type `PROCKIND`, along with all other process names in the simulation. This same name, preceded by "LP" (for logical process) is by convention the name of the Pascal procedure that emulates this process. "LPQUEUE20" also appears in the procedure `RESUME`, by means of which all processes are activated when in the simulation, they make a transition into the `XQT` state. See fig 3.2.

The process name also appears in the definition of the `OWNDATA` record type. This is where the definition of the local, or "own" variables, is bound to the type of the owning process kind, in this case, the `QUEUE20`.

As each instance of the process kind `QUEUE20` is created, a fresh local address space, and activation record for the new `QUEUE20` process are allocated in the form of previously unused elements in the `OWN`, and `LPS` arrays. The `OWNDATA` records, which comprise the base type of the `OWN` array, have their structure determined by the type of their owner.

An argument is provided in the calls to the `CREATE`

pointers. Also at this label, if the queue is not full, the process waits for another input message.

Whenever a message is successfully sent out from this process, it is resumed at MPC = 3. Here it is certain that there is room for at least one more message, so the process waits for this input. Depending on the relative values of the pointers, there may be another message to send out, so the process may actually wait in parallel for both input and output.

#### 4.1.5 The FORK2 Process Type

A process of kind FORK2 has a single input port, known as F2INPORT, and two output ports, known as F2OUT1PORT, and F2OUT2PORT. Initially, a FORK2 process waits for input. When this is received, the process decides which of its two output ports to send the message over, and begins waiting to do so. Once the message is sent out, the process again waits for input and the cycle repeats. Currently, the decision about which output port to try to send over is determined by performing a single Bernoulli trial, where the probability of a success, defined as deciding to send over F2OUT1PORT, is a parameter that is entered by the operator at process-create time.

A different decision policy would result in a process with the same connectivity topology, but different com-

munication semantics. For example, the decision could be based on message traffic intensity, or perhaps an attribute of the particular message itself.

The chronological condition is satisfied because the inputs to the process are assumed to be chronological, and the process never holds more than one message, hence it has no opportunity to exchange two messages before they are sent out.

#### 4.1.6 The MERGE2 Process Type

The MERGE2 process has two input ports, named M2IN1PORT, and M2IN2PORT, and a single output port named M2OUTPORT. The waiting rules of this type of process are pretty much fixed because of the requirement that the output messages be chronological, and the fact that there is minimal queueing inside of the process.

Initially, the process waits for input from either ports. As soon as any message is received over either port, it is buffered, and the process continues to wait for a message to come in over the other port. When this second message is received, the timestamps of the two messages are compared, and the process waits to send the message with the smaller timestamp over its output port. When the message is sent out, the process waits for input over the port which provided the message that was last sent. Once this message

is received, the timestamps are again compared, and the cycle repeats. If two timestamps are found to be equal upon comparison, one message is arbitrarily selected to be sent first, and when it is sent, the process begins waiting for the port that provided this message.

This section has described six process varieties which have been modelled using SIM. Many others are possible, and some of these have been suggested. Each of these fills a role as a building block for constructing fairly general networks of processes. The processes as presented are useful for simulation of networks that create, delay, enqueue, and kill messages or jobs. The exact processing performed within a network depends on characteristic parts of the process' code.

#### 4.2 Deadlocks in Distributed Simulation

A problem with distributed algorithms following the waiting rules inherent in the processes described above is the occurrence of deadlocks. Of course, deadlocks that arise in the modelled system will show up in the simulation; this is desirable. In fact, the purpose of a simulation might be to gain insights into the deadlock characteristics of a network. However, other deadlocks, not arising in the simulated system can and do show up strictly as a result of the waiting rules proposed in [CHA79].

#### 4.3 Deadlock Resolution in Distributed Simulation

A solution proposed in this same paper was to introduce special "null" messages, whose sole purpose was to advance the line times of the ports, so that some process in a deadlocked system might be able to change the ports upon which it was waiting, while still guaranteeing the chronological condition for all of the ports, and hence break the deadlock. This algorithm was implemented in [SEE79] and found to work well in systems without feedback paths. In systems with feedback paths, the number of null messages grew so large as to flood the communication capacity of the network.

A separate solution was proposed in [CHA81] that computes a least upper bound on the line times. The algorithm consists of  $N$  phases, where a phase is comprised of an execution part followed by a communication part.  $N$  represents the total number of processes in the network. The reader is referred to the original article for a complete discussion of the problem of deadlocks in distributed simulation, with examples, and a description of the deadlock recovery algorithm together with a proof of its correctness.

As a testbed for general purpose distributed simulation, this algorithm has been implemented in the processes described above. During a simulation run, the processes compute and communicate until deadlock occurs, at which



point, the entire system typically backs up and ceases to make progress. In the current implementation, the deadlock is detected by the central CONTROLLER, which initiates the recovery algorithm, and synchronizes the N phases. Initially, the line time used for each port is infinitely large. Then at each phase, every process either revises downward, or leaves constant its estimate of the earliest time ( = smallest line time ) at which it may try to send out a message over each of its output ports. At the k-th computation step, the process is able to do this based on the presence of the (k-1)st such estimate from all processes that send messages to it.

In a pure form, this would have been passed as a message to it along all of its input ports during the (k-1)st communication phase. This would be complicated in SIM, however, because the ports can only pass messages of a single type. Hence a pure implementation in SIM would require a network of communication lines for these time estimates that was parallel to the regular message ports. In the SIM implementation, this pure structure has been compromised to the extent of providing a dedicated word of the port's records to carry this information. This corresponds in the physical system to giving a little of the bandwidth of the port over to the deadlock recovery messages, as was needed. Of course, the ports would not be trying to carry regular messages at the same time, because if they were, the system would not be

deadlocked.

A process executing its k-th step of this algorithm is resumed at MPC = 1010. There it reads the (k-1)st estimate on all of its input ports, and writes its own k-th estimate to its output ports. When the network consists of N processes, the algorithm computes its terminal values within N steps. Typically, the actual number of steps required is much lower, and seems tied to the length of the longest closed chain of processes.

After the N steps, the central CONTROLLER resumes the processes one last time in order to let them try to change the ports they are waiting for. According to [CHA81], there is always at least one process that can change its waiting pattern. At this point, if a process can change the lines that it is waiting for, it does so by means of PARWAIT calls.

After this k-step algorithm runs, the CONTROLLER attempts to make some of the ports fire, as if there had been no deadlock. If any firing is possible, at least one message is fired, and the simulation can continue. Otherwise, the deadlock cannot be eliminated, due to the fact that it is one which arose in the physical system, not one which appeared as a result of the waiting rules.

The results of SIM runs with this algorithm are

encouraging. The algorithm works as expected to recover from deadlocks in both feedback, and feedforward networks. The interpretation of performance statistics for the algorithm requires only that the amounts of time for communication and processing be stated for a particular network. The relative cost would be greater in communication-limited, as opposed to compute-bound systems. As long as the time required for a process to communicate with all other  $N-1$  processes is no worse than  $O(N)$ , the overall performance is  $O(N^2)$ , because it repeats this  $O(N)$  communication for each of the  $N$  steps of the algorithm. This result would be somewhat harder to establish for a totally distributed implementation, e.g. because of the time required to detect the deadlock initially, but should prove to be true.

As a last point on this subject, it is very easy to examine the effect of queue size on the frequency of deadlocks in such a distributed simulation. In an environment where execution of the recovery algorithm were very undesirable, one could see the advantage to be gained in avoided execution that could be obtained at the expense of extra memory buffering for the processes. The results from such a simulation for two particular networks appear in appendix 2.

## 5.0 SUGGESTIONS FOR FURTHER WORK

An obvious extension of the SIM program would be to provide the front-end language processing that would be required to implement a very useful subset of CSP. Such a translator would output entire SIM programs with the procedures embedded in them. This report has suggested some of the strategy that could be used in such a program.

In the current implementation, the topology of a network is constant throughout the execution of the simulation after it is entered during the network specification part of a run. This is not a general constraint imposed by SIM however. The connectivity of the network is contained in the global data of the system, and by changing this data during a run, the effective topology could be changed. In addition to changing the records about who talks to whom, new instances of the processes can be created, initialized, and executed, all on the fly. In this way, a network could change its topology to adapt to a particular processing problem. This would allow very general forking and joining in the processes of the network, as is required by many of the languages supporting concurrency. By spawning a number of subprocesses, a process could implement nondeterministic

algorithms.

A very simple mapping would allow an instance of the SIM program and its processes to become the program that operated a true distributed system with a central controlling process, communicating e.g. over wires. The "MPC label" concept would map directly onto any of the typical machines with a vectored interrupt structure. With some modifications, the "oracle" functions provided by the CONTROLLER could be distributed and placed in the processes they served, possibly in the form of hardware bus arbitration.

## 6.0 SUMMARY

This report has described the SIM program, which has been written to assist one in the development of programs with distributed control in the form of a number of processes executing in parallel, and strong interactions in the form of messages passed between the processes. The program was motivated by the requirements of a language proposed by Hoare, that has proven useful in the expression of programs based on message passing.

The data, control structures, and special debug facilities of the program were detailed as an aid for anyone wishing to use it for future work.

A particular application, distributed simulation, was used as an example of the process structure and general nature of networks supported by SIM. Handling of the problem of deadlocks in this type of system was accommodated at modest expenditure of effort by introducing a dependence of the algorithm upon the central CONTROLLER.

# APPENDIX 1

## TEXT OF SIM PROGRAM

```

00050 PROGRAM DSIM(INPUT*,OUTPUT)
00100
00150
00200 (* IN THIS LISTING, 'TARDEP' MEANS TARGET DEPENDENT, AND IMPLIES
00250 THAT THE REFERENCED CODE OR DATA STRUCTURES ARE IMPACTED
00300 BY CHANGES IN THE TARGET NETWORK TO BE SIMULATED, *)
00350
00400 CONST
00450
00500 (* TARDEP * TOTAL NETWORK PORT COUNT *)
00550 PORTMAX = 100
00600
00650 (* TARDEP * TOTAL NETWORK PROCESS COUNT *)
00700 PROCMAX = 100
00750
00800 TYPE
00850
00900 (* TARDEP * SCALAR PROCESS TYPE LIST*)
00950 PROCKIND = (SOURCE,SINK,FORK2,MERGE2,DELAY,QUEUE20)
01000
01050 (* EACH LOGICAL PROCESS IS ALWAYS IN ONE OF FOUR STATES,
01100 ACTUALLY EXECUTING (XDT), WAITING FOR COMMUNICATION (BLK),
01150 ACTUALLY COMMUNICATING (CHN), OR TERMINATED (TRM), *)
01200 STATETYPE = (XDT,BLK,TRM,CMN)
01250
01300 (* SMALL INTEGERS PORTID AND PROCID TAKE ON VALUES THAT UNIQUELY
01350 IDENTIFY A PARTICULAR PORT OR PROCESS, RESPECTIVELY, *)
01400 PORTID = 1,PORTMAX
01450 PROCID = 1,PROCMAX
01500
01550 (* PORTDIRECTIONTYPES ARE USED AS PARAMETERS TO DEFINEPORT *)
01600 PORTDIRECTIONTYPE = (INN,OUT)
01650
01700 (* TARDEP * SCALAR TARGET MESSAGE TYPE LIST*)
01750 MSGKIND = (UNDEFINED,JOB)
01800
01850 (* TARDEP * TARGET MESSAGE-RECORD TYPE, A 'MESSAGE' CONSISTS OF
01900 AN MTIME FIELD FOR THE TIME COMPONENT OF THE MESSAGE, AND AN MKIND
01950 FIELD TO DENOTE THE TARGET DEPENDENT MSGKIND OF THE MESSAGE,
02000 E.G. 'SIGNAL' & 'ACK', THE MKIND VARIABLE ALSO DETERMINES THE
02050 REMAINING STRUCTURE OF THE MESSAGE, EACH NEW ELEMENT IN THE
02100 SCALAR MSGKIND LIST MUST HAVE AN ASSOCIATED CASE MKIND
02150 DEFINITION IN THE FOLLOWING LINES, *)
02200 MESSAGE = RECORD
02250 MTIME:INTEGER
02300 CASE MKIND : MSGKIND OF
02350 UNDEFINED : (UDATA:INTEGER)
02400 JOB : (JOBNUMBER:INTEGER)
02450 END
02500
02550 (* EACH LOGICAL PROCESS HAS AN ASSOCIATED ACTIVITY RECORD OF TYPE
02600 ACTREC, THAT RETAINS THAT PART OF THE PROCESS' STATE VECTOR THAT
02650 IS MAINTAINED BY THE INTERPRETER. READ 'A' FOR ACTIVITY IN
02700 INTERPRETING THE FIELDS. THE TARGET PROCESSES INFORM THE
02750 INTERPRETER OF THEIR INTENDED NEXT STATE VIA THE VARIABLE
02800 ANEXTSTATE, WHICH ONLY HAS MEANINGFUL VALUES OF BLK OR TRM,
02850 ASTATE AND ANEXTSTATE THEN ARE THE CURRENT AND NEXT STATES OF
02900 THE BOUND PROCESS RESPECTIVELY, THE ATIMELEFT FIELD HAS
02950 MEANING IN BOTH THE XDT AND CHN STATES, AND DENOTES THE AMOUNT

```

```

03000 OF TIME REQUIRED TO CONCLUDE THE CURRENT EXECUTION AND COM-
03050 MUNICATION RESPECTIVELY, THE SUMS VARIABLES RETAIN THE AMOUNT
03100 OF TIME SPENT SO FAR IN EACH OF THE STATES XOT, BLK, AND CMN.
03150 IF THE CURRENT STATE OF A PROCESS IS TRM, THEN ATRMTIME HAS
03200 THE NETWORK CLOCK TIME THAT THE PROCESS BECAME TERMINATED,
03250 AMESSAGE IS A ONE-MESSAGE BUFFER USED FOR MESSAGES TRANSFERED
03300 TO AND FROM THE PORTS FOR COMMUNICATION,
03350 ATYPE AND AINSTANCE CONTAIN THE TARGET NETWORK PROCKIND AND
03400 INSTANCE OF THE BOUND PROCESS, E.G. 'SERVER' '7', OR 'MERGE' '2'.
03450 AMPC (FOR META PROGRAM COUNTER) IS THE MEANS OF TELLING THE
03500 INTERPRETER WHAT THE PROCESS POINT-OF-RETURN SHOULD BE FOLLOWING
03550 COMMUNICATION OVER A PARTICULAR PORT. THIS IS PASSED AS A
03600 PARAMETER IN THE PARWAIT CALL THAT DECLARES THAT PORT TO BE
03650 ONE OF THE AWAITED PORTS FOR THE CALLING PROCESS DURING THE NEXT
03700 BLK STATE FOR THE PROCESS.
03750 ACCORDINGLY, WHEN THE INTERPRETER HAS CHOSEN A PARTICULAR PORT
03800 FOR FIRING, THE COMMUNICATING PROCESSES ARE INFORMED OF THEIR
03850 RETURN POINT THAT CORRESPONDS TO ACTION TO FOLLOW THIS PORT
03900 FIRING.
03950
03950 ACTREC = RECORD *)
04000   ATYPE:PROCKIND;
04050   AINSTANCE:1,.,PROCMAX;
04100   ASTATE,ANEXTSTATE:STATETYPE;
04150   ATIMELEFT,ASUMXUTTIME,ASUMCMNTIME,ASUMBLKTIME:INTEGER;
04200   ATRMTIME,AMPC,AXOTTIME:INTEGER;
04250   APURT:PORTID;
04300   AMESSAGE:MESSAGE;
04350   END;
04400
04450 (* EACH PORT HAS AN ASSOCIATED PORT RECORD OF TYPE PDRTRC,
04500 READ 'P' FOR PORT IN INTERPRETING THE FIELD NAMES.
04550 A PORT HAS A FIXED SENDER AND RECEIVER WHICH ARE NAMED AS
04600 PSENDER AND PRECEIVER RESPECTIVELY, FOR EACH OF SENDER AND
04650 RECEIVER, THE PORT MAINTAINS BOOLEAN VARIABLES INDICATING
04700 READY AND ELIGIBLE, WHERE X IS ONE OF S FOR SENDER OR
04750 R FOR RECEIVER, PXWAITING <====> ( X IS READY TO COMMUNICATE ),
04800 AND PXELIGIBLE <====> ( X MARKED THIS PORT AS ELIGIBLE DURING
04850 ITS LAST EXECUTION PHASE BY MEANS OF A PARWAIT CALL ) .
04900 PORTS ARE NEVER FIRED UNLESS BOTH NAMED PROCESSES ARE
04950 BOTH ELIGIBLE AND WAITING,
05000 WHEN THE MESSAGE DOES FIRE, THE SENDER AND RECEIVER
05050 ARE INFORMED OF THEIR RESUME POINTS AS PSMPC AND PRMPC RESP,
05100 PSTIME AND PRTIME ARE THE NUMBER OF TIME UNITS REQUIRED TO
05150 SEND AND RECEIVE A MESSAGE ON THIS PORT, COMMUNICATION TIME
05200 IS EXPLICITLY MODELLED WHEN THESE VALUES ARE SET NON-ZERO,
05250 A VIRTUAL FIELD, PCAPACITY, IS A MEASURE OF THE COMMUNICATION
05300 CAPACITY REQUIRED FOR THE 'LINE OF COMMUNICATION' MODELLED
05350 BY THIS PORT:
05400   PCAPACITY = (SIZE OF MESSAGE)/(TIME REQUIRED FOR TRANS)
05450 SIMILARLY, PDELAY, FOR PORT TRANSMISSION DELAY COULD BE:
05500   PDELAY = PRTIME - PSTIME
05550 THAT IS, THE ADDITIONAL AMOUNT OF TIME THAT COMMUNICATION
05600 DETAINS THE RECEIVER COMPARED WITH THE SENDER.
05650 HENCE FAIRLY GENERAL COMMUNICATION CAN BE MODELLED EXPLICITLY,
05700 PWIJ IS USED IN THE DEADLOCK RECOVERY COMPUTATION OF W SUB I,J
05750 THIS IS BEING HANDLED FOR THESE EXPERIMENTS
05800 AS PART OF THE PORT RECORD INSTEAD OF THE LOCAL DATA OF THE LPS,
05850 AN IMPLEMENTATION MIGHT MULTIPLEX THE CHANNEL INTO A DATA PART,
05900 AND A PART USED TO COMMUNICATE THE INFORMATION NECESSARY TO

```



```

05950          BREAK DEADLOCK,
06000
06050          PORTREC = RECORD
06100              PSENDER,PRECEIVER:PROCID;
06150              PSELIGIBLE,PSWAITING,PRWAITING,PRELIGIBLE:BOOLEAN;
06200              PSMPC,PRMPC,PSTIME,PRTIME,PMSGCOUNT:INTEGER;
06250              PMESSAGE:MESSAGE;
06300          PWIJ:INTEGER;
06350          END;
06400
06450          (* TANDEP * PROCESS' OWN DATA RECORD TYPES
06500          EACH PROCESS HAS A DATA STRUCTURE FOR ITS OWN DATA, WHOSE
06550          STRUCTURE IS DETERMINED BY THE TARGET NETWORK PROCESS KIND
06600          E.G. A PROCESS OF KIND 'QUEUE' WOULD REQUIRE OWN DATA THAT
06650          RETAINED INFORMATION ABOUT JOBS IN THE QUEUE, AND HEAD AND
06700          TAIL POINTERS, WHERE A BASIC PROCESS KIND HAS BEEN IN-
06750          STANTIATED SEVERAL TIMES, SEPARATE OWN DATA IS MAINTAINED
06800          FOR EACH ONE.
06850          EACH SUCH RECORD HAS AN DPROCKIND FIELD TO RECORD THE OWNER
06900          PROCESS KIND, AND A VARIABLE PART WHOSE STRUCTURE IS
06950          DETERMINED BY DPROCKIND.
07000          SEE THE DISCUSSION ABOUT THE VARIABLE, OWN, FOR EXAMPLES OF
07050          REFERENCING OWN DATA.
07100
07150          (* QUEUE TYPE IS A TYPE DEFINED FOR USE BY THE QUEUE20 LP KIND *)
07200          QUEUE TYPE = ARRAY[0..19] OF INTEGER;
07250
07300          OWN DATA = RECORD
07350              CASE DPROCKIND : PROCKIND OF
07400                  SOURCE:(SOUTPORT:PORTID;SOMU:REAL;SOTIME,SOMSGCOUNT:INTEGER;
07450                      SOCON:REAL;SUSEQ:INTEGER);
07500                  SINK:(SINPORT:PORTID;SITIME,SIJOBCOUNT:INTEGER);
07550                  MERGE2:(M2IN1PORT,M2IN2PORT,M2OUTPORT:PORTID;
07600                      M2IN1TIME,M2IN2TIME,M2JOB1COUNT,
07650                      M2JOB2COUNT:INTEGER;M2HAVE1,M2HAVE2:BOOLEAN;
07700                      M2IN1MSG,M2IN2MSG:MESSAGE);
07750                  FORK2:(F2INPORT,F2OUT1PORT,F2OUT2PORT:PORTID;
07800                      F2TIME,F2OUT1COUNT,F2OUT2COUNT:INTEGER;F2RHO:REAL);
07850                  QUEUE20:(Q20INPORT,Q20OUTPORT:PORTID;
07900                      Q20BUFFER,Q20BUFFER2:QUEUE TYPE;Q20TIME:INTEGER;
07950                      Q20INPTR,Q20OUTPTR,Q20JOB1COUNT,Q20QMAX:INTEGER);
08000                  DELAY:(DINPORT,DOUTPORT:PORTID;DEMU,UCON:REAL;
08050                      DTIME,DSUMPTIME,DJOB1COUNT:INTEGER);
08100          END;
08150
08200          (* A TALLEYBYSTATE VARIABLE HAS 4 INTEGER FIELDS TO RETAIN A COUNT
08250          OF THE NUMBER OF PROCESSES IN EACH STATE, SUCH A VARIABLE
08300          IS USED AS A RESULT PARAMETER OF THE COUNTBYSTATE ROUTINE WHICH
08350          COUNTS THE NUMBER OF PROCESSES IN EACH STATE AND PUTS THE TOTALS
08400          INTO THE APPROPRIATE FIELDS. FOLLOWING THIS PROCEDURE CALL, THE
08450          COUNTS REMAIN VALID AS LONG AS NO PROCESS CHANGES STATE, AND
08500          NO ROUTINES CHANGE THE VALUES. COUNTBYSTATE IS THE ONLY PLACE
08550          THESE VARIABLES ARE WRITTEN INTO.
08600          TALLEYBYSTATE = RECORD
08650              XQTING,CMNING,BLKED,TRMED:INTEGER;
08700          END;
08750
08800          (* A PORT POINTER IS AN ARRAY WITH ONE ENTRY PER PORT, WHOSE
08850          ELEMENTS ARE POINTERS TO A PORT. ( SMALL INTEGER INDICIES INTO

```

```

08900      THE PORTS' DATA RECORD). THIS IS A CONVENIENT MECHANISM TO
08950      DEFINE AN ORDERING ON THE PORTS SUCH AS THE USE BY THE SCHEDULE
09000      PROCEDURE TO DEFINE THE ORDER IN WHICH THE PORTS SHOULD FIRE, *)
09050      PORTPOINTER = ARRAY[1,,PORTMAX] OF 0,,PORTMAX)
09100
09150
09200  VAR      (* GLOBAL VARIABLES *)
09250
09300
09350      (* PORTS IS A TABLE OF PORTRECORDS, ONE FOR EACH PORT IN THE
09400      TARGET NETWORK, A PORT'S PORTID INDEXES THE TABLE, *)
09450      PORTS I ARRAY[1,,PORTMAX] OF PORTREC)
09500
09550      (* LPS IS A TABLE OF ACTIVATION RECORDS, ONE FOR EACH LOGICAL
09600      PROCESS IN THE TARGET NETWORK, A PROCESS' PROCID INDEXES
09650      INTO THE TABLE, *)
09700      LPS I ARRAY[1,,PROCMAX] OF ACTREC)
09750      (* OWN IS A TABLE CONTAINING THE PROCESS' OWNED DATA RECORDS,
09800      ONE FOR EACH LOGICAL PROCESS, PROCESS I ACCESSES ITS DATA
09850      AS OWN[I],HEAD , OWN[I].TAIL AND OWN[I],Q[I] FOR THE
09900      DEFINITIONS OF OWN DATA FOR THE PROCESS KIND 'QUEUE'
09950      DISCUSSED WITH THE OWNDATA TYPE DEFINITION, *)
10000      OWN I ARRAY[1,,PROCMAX] OF OWNDATA)
10050      (* NETTIME IS THE GLOBAL NETWORK CLOCK, A DISCRETE-VALUED
10100      CLOCK, IT ADVANCES ONE TIME UNIT EVERY TIME THE GLOBAL
10150      CLOCK 'TICKS', SEE THE PROCEDURE DEFINITION FOR TICK,
10200
10250      NETSUMXQTIME,NETSUMCMNTIME, AND NETSUMBLKTIME ACCUMULATE
10300      THE AMOUNT OF TIME THAT PROCESSES SPEND IN STATES
10350      EXECUTING, COMMUNICATING, AND BLOCKED RESPECTIVELY,
10400
10450      NETMSGLIMIT AND NETTIMELIMIT ARE THE TIME AND MESSAGE BOUNDS
10500      THAT THE PROGRAM OPERATOR HAS SPECIFIED, IF EITHER LIMIT
10550      IS EXCEEDED, THE SIMULATION IS STOPPED AND A MESSAGE PRINTED,
10600      BEFORE THIS HAPPENS HOWEVER, THE OPERATOR GETS THE OPTION OF
10650      ENTERING NEW, HIGHER VALUES SO THAT THE SIMULATION MAY
10700      PROCEED, NETMSGCOUNT CONTAINS THE TOTAL NUMBER OF MESSAGES
10750      THAT HAVE BEEN SENT SINCE THE BEGINNING OF THE SIMULATION,
10800
10850      PORTJ IS AN INDEX VARIABLE THAT RANGES OVER THE DEFINED PORTS
10900
10950      PROC I AND PROCK ARE INDEX VARIABLES THAT RANGE OVER THE PROCS
11000
11050      NETDEADLOCK AND NETTERM ARE BOOLEAN VARIABLES WITH DEFINITIONS:
11100      ALL NONTERMINATED PROCESSES ARE DEADLOCKED, AND ALL PROCESSES
11150      ARE TERMINATED RESPECTIVELY,
11200
11250      SEVEN DIFFERENT TRACE VARIABLES TURN ON SIMULATION RUN-TIME
11300      OUTPUT FOR THE PURPOSE OF DEBUGGING THE INTERPRETER OR THE
11350      TARGET LOGICAL PROCESSES, OR OBSERVING THE PROGRESS OF A
11400      SIMULATION, IN GENERAL THESE VARIABLES PRODUCE MORE LISTING
11450      WHEN THEY HAVE LARGER VALUES, ALL ARE INTEGER, AND A ZERO
11500      MEANS NO-TRACE, THESE ARE:
11550      NAME GOVERNS LISTING WITHIN THE SUBJECT AREA
11600      NETXQTRACE THE PROCESS STATE = EXECUTING
11650      NETCMNTRACE THE PROCESS STATE = COMMUNICATING
11700      NETBLKTRACE THE PROCESS STATE = BLOCKED
11750      NETTRMTRACE THE PROCESS STATE = TERMINATED
11800      INTERTRACE ACTION OF THE INTERPRETER, PROCEDURE CALLS, ETC.

```

```

11850 TARGETTRACE ACTIVITIES OF THE TARGET NETWORK
11900 NETDEADTRACE DEADLOCK DETECTION AND RECOVERY
11950
12000 COUNT IS INITIALIZED BY COUNDBYSTATE AND PROVIDES A CONVENIENT
12050 WAY OF DETERMINING THE NUMBER OF PROCESSES THAT ARE IN EACH STATE
12100 AT ANY GIVEN TIME. SEE THE TESTS FOR TERMINATION AND DEADLOCK
12150 DETECTION IN THE MAIN PROGRAM.
12200
12250 TTY IS THE PROGRAM NAME OF THE OPERATOR'S I/O DEVICE, ASSUMED
12300 TO BE A TELETYPE-LIKE DEVICE.
12350
12400 NETFAIR IS AN OWNED VARIABLE OF THE SCHEDULE PROCEDURE THAT
12450 IS USED TO ENSURE FAIRNESS AMONG THE WAITING PORTS, SO THAT
12500 NO MESSAGE WITH ELIGIBLE AND WAITING SENDER AND RECEIVER IS PASSED
12550 OVER INFINITELY OFTEN FOR FIRING. SEE PROCEDURE SCHEDULE.
12600
12650 DEADLOCKCOUNT IS A COUNT OF THE NUMBER OF DEADLOCKS DETECTED
12700 IN THE SEQUENTIAL SIMULATION. NOTE THAT THIS INCLUDES ANY
12750 THAT ARISE IN THE TARGET SIMULATION, AND TYPICALLY MANY MORE
12800 THAT ARTIFICIALLY ARISE BECAUSE OF THE SEQUENTIAL SIMULATION
12850 OF THE WAITING RULES FOR MERGE PROCESSES, AND POSSIBLY OTHERS.
12900
12950 HIGHPROC AND HIGHPORT ARE THE HIGHEST NUMBERED PROC AND PORT
13000 RESPECTIVELY, THAT ARE USED IN THE CURRENT SIMULATION. THIS
13050 DEPENDS ON WHAT NETWORK THE OPERATOR HAS SPECIFIED.
13100
13150
13200 NETTIME,NETSUMXQTIME,NETSUMCMNTIME,NETSUMBLKTIME(INTEGER)
13250 NETMSGLIMIT,NETTIMELIMIT,NETMSGCOUNT(INTEGER)
13300 PORTJ: 1.,PORTMAX)
13350 PROCJ,PROCK: 1.,PROCMAX)
13400 NETDEADLOCK,NETTERM(BOOLEAN)
13450 NETCMNTRACE,NETXQTRACE,NETBLKTRACE,INTERTRACE,NETTRMTRACE(INTEGER)
13500 TARGETTRACE(INTEGER) COUNT:ALLEYBYSTATE)
13550 NETDEADTRACE(INTEGER)
13600 TTY: ATEXT)
13650 NETFAIR:PORTID) (* PORT FAIRNESS ==> PROCESS FAIRNESS *)
13700 DEADLOCKCOUNT(INTEGER)
13750 DEADLOCKLIMIT(INTEGER)
13800 HIGHPROC:PROCID)HIGHPORT:PORTID)
13850 BUFFERSIZE(INTEGER) (* FOR BUFFERSIZE VS DEADLOCK EXPERIMENT *)
13900
13950 PROCEDURE CONTINUE;
14000 (* PRINTS A MESSAGE AND SOLICITS DUMMY INPUT AS A DEBUG TOOL *)
14050 VAR DUMMY:CHAR)
14100 BEGIN
14150 WRITELN(TTY,' CONTINUE CALLED, ENTER ANY CHAR ');
14200 BREAK;
14250 RESET(TTY);
14300 READ(TTY,DUMMY);
14350 END;
14400 FUNCTION MIN(ARG1,ARG2:INTEGER):INTEGER)
14450 BEGIN
14500 IF ARG1<=ARG2 THEN MINI=ARG1 ELSE MINI=ARG2;
14550 END;
14600
14650 FUNCTION MAX(ARG1,ARG2:INTEGER):INTEGER)
14700 BEGIN
14750 IF ARG1>=ARG2 THEN MAXI=ARG1 ELSE MAXI=ARG2;

```

```

14800 END)
14850
14900 PROCEDURE INCR(VAR ARG:INTEGER)
14950 BEGIN
15000     ARG:=ARG+1;
15050 END)
15100
15150 PROCEDURE DECR(VAR ARG:INTEGER)
15200 BEGIN
15250     ARG:=ARG-1;
15300 END)
15350
15400 PROCEDURE SHOWMSG(MSG:MESSAGE)
15450 (* TARDEP * PROCEDURE TO WRITE THE PERTINANT INFORMATION FROM MSG
15500 TO THE TTY. THE ALLOWABLE VALUES OF MSG,MSGKIND MUST BE ACCOUNTED
15550 FOR IN THE CASE STATEMENT SO THAT ALL ALLOWABLE MESSAGES CAN BE
15600 DUMPED TO TTY. *)
15650 BEGIN
15700     WRITELN(TTY);
15750     WITH MSG DO BEGIN
15800         WRITE(TTY,' MTIME = ',MTIME:6,' HAS KIND = ');
15850         CASE MKIND OF
15900             UNDEFINED:BEGIN
15950                 WRITELN(TTY,' UNDEFINED, MESSAGE UDATA = ',UDATA);
16000             END;
16050             JOB :BEGIN
16100                 WRITELN(TTY,' JOB, JOBNUMBER = ',JOBNUMBER);
16150             END;
16200         END;
16250     END;
16300 END)
16350
16400 PROCEDURE COPYMSGTOFROM(VAR DEST:MESSAGE;FROM:MESSAGE)
16450 (* TARDEP * COPIES THE FROM MESSAGE TO THE DEST MESSAGE. *)
16500 (* MUST BE ABLE TO COPY ALL THE NECESSARY PARTS OF ALL POSSIBLE
16550 MESSAGE VARIETIES, AS DETERMINED IN ANY PARTICULAR CALL BY THE
16600 VALUE OF FROM,MKIND *)
16650 BEGIN
16700     DEST,MTIME := FROM,MTIME;
16750     DEST,MKIND := FROM,MKIND;
16800     CASE FROM,MKIND OF
16850         UNDEFINED:DEST,UDATA:=FROM,UDATA;
16900         JOB :DEST,JOBNUMBER:=FROM,JOBNUMBER;
16950     END;
17000     IF (INTERTRACE>0) THEN
17050         WRITELN(TTY,' MESSAGE COPIED');
17100     IF (INTERTRACE>50) OR (NETCMNTRACE>50) THEN CONTINUE;
17150 END)
17200
17250 PROCEDURE SHOWPORT(ID:PORTID)
17300 (* DUMTTY STATE OF PORT *)
17350 BEGIN
17400     WITH PORTS[ID] DO BEGIN
17450         WRITELN(TTY);
17500         WRITE(TTY,' PORT ',ID:3,' PSENDER=',PSENDER:3,' PRECEIVER=');
17550         WRITELN(TTY,PRECEIVER:3,' PSELIGIBLE=',PSELIGIBLE);
17600         WRITE(TTY,' PRELIGIBLE=',PRELIGIBLE,' PSWAITING=',PSWAITING);
17650         WRITE(TTY,' PRWAITING=',PRWAITING,' PSMPC=',PSMPC:3,' PRMPC=');
17700

```

```

17750      WRITELN(TTY,PRMPC13))
17800      WRITELN(TTY,' PWIJ = ',PWIJ16))
17850      WRITE(TTY,' PSTIME=',PSTIME18,' PRTIME=',PRTIME18))
17900      WRITELN(TTY,' PMSGCOUNT=',PMSGCOUNT16,' PMESSAGE IS '))
17950      SHOWMSG(PMESSAGE))
18000      END) (* WITH *)
18050      END)

18100      PROCEDURE PRSTATETYPE(ST:STATETYPE);
18150      (* PRINT PROCESS STATE GIVEN BY ST          FIELD WIDTH IS 5 *)
18200      BEGIN
18250          CASE ST OF
18300              XQT:WRITE(TTY,' XQT ');
18350              BLK:WRITE(TTY,' BLK ');
18400              CHN:WRITE(TTY,' CHN ');
18450              TRM:WRITE(TTY,' TRM ');
18500          END)
18550      END)

18600      PROCEDURE PRMSGKIND(MK:MSGKIND);
18650      (* TARDEP * WRITE THE MESSAGE KIND PRINTNAME AT THE CURRENT CURSOR
18700      ON TTY
18750      BEGIN
18800          CASE MK OF
18850              UNDEFINED:WRITE(TTY,' UNDEFINED ');
18900              JOB:WRITE(TTY,' JOB ');
18950          END)
19000      END)

19050      PROCEDURE PRPROCKIND(TY:PROCKIND);
19100      (* TARDEP * PRINT LOGICAL PROCESS KIND GIVEN BY TY AT CURRENT CURSOR
19150      FIELD WIDTH IS 9.
19200      BEGIN
19250          CASE TY OF
19300              SOURCE :WRITE(TTY,' SOURCE ');
19350              SINK :WRITE(TTY,' SINK ');
19400              FORK2 :WRITE(TTY,' FORK2 ');
19450              MERGE2 :WRITE(TTY,' MERGE2 ');
19500              DELAY :WRITE(TTY,' DELAY ');
19550              QUEUE20:WRITE(TTY,' QUEUE20 ');
19600          END)
19650      END)

19700      PROCEDURE PRSIGNATURE(ID:PROCID);
19750      (* PRINT PROCESS KIND, INSTANCE, AND PROCESS ID FOR THE PROCESS
19800      NUMBER PASSED AS ID, AT THE CURRENT CURSOR POSITION, FIELD WIDTH
19850      IS 31
19900      BEGIN
19950          PRPROCKIND(LPS[ID],ATYPE);
20000          WRITE(TTY,' INSTANCE = ',LPS[ID],AINSTANCE14,' UNIQUE ID = ',ID14))
20050      END)

20100      PROCEDURE SHOWPROCESS(ID:PROCID);
20150      (* DUMPTTY STATE OF THE PROCESS NAMED ID.
20200      BEGIN
20250          WRITELN(TTY);
20300          WRITE(TTY,' SHOWPROCESS');
20350          PRSIGNATURE(ID);

```

```

20700 WITH LPS(ID) DO BEGIN
20750   WRITELN(TTY);
20800   WRITE(TTY, ' ASTATE = ');
20850   PRSTATETYPE(ASTATE);
20900   WRITE(TTY, ' ANEXTSTATE = ');
20950   PRSTATETYPE(ANEXTSTATE);
21000   WRITELN(TTY);
21050   WRITE(TTY, ' ATIMELEFT = ', ATIMELEFT:14, ' ASUMXQTTIME = ');
21100   WRITELN(TTY, ASUMXQTTIME:16, ' ASUMCHNTIME = ', ASUMCHNTIME:16);
21150   WRITE(TTY, ' ASUMBLKTIME = ', ASUMBLKTIME:16);
21200   WRITE(TTY, ' ATRMTIME = ', ATRMTIME:18, ' AMPC = ', AMPC:14);
21250   WRITELN(TTY, ' AXQTTIME = ', AXQTTIME:16);
21300   WRITELN(TTY, ' ATIMELEFT = ', ATIMELEFT:15);
21350   WRITE(TTY, ' APORT = ', APORT:14, ' MESSAGE BUFFER ');
21400   WRITELN(TTY, 'CONTAINS THE MESSAGE');
21450   SHOWMSG(AMESSAGE);
21500   END;
21550 END;
21600
21650 PROCEDURE SHOWOWN(PROCS:PROCID);
21700 (* TARDEF * PRINT OWNED DATA FOR PROCESS GIVEN BY PROCS *)
21750 VAR I:INTEGER;
21800 BEGIN
21850   WRITELN(TTY);
21900   WRITE(TTY, ' PROCESS ', PROCS:14, ' HAS OWNED DATA FOR OPROCKIND=');
21950   PRPROCKIND(OWN[PROCS], OPROCKIND);
22000   WRITELN(TTY);
22050   WITH DWN[PROCS] DO
22100     CASE OPROCKIND OF
22150       SOURCE:BEGIN
22200         WRITE(TTY, ' SOOUTPORT = ', SOOUTPORT:13);
22250         WRITE(TTY, ' SOCON = ', SOCON);
22300         WRITELN(TTY, ' SOMU = ', SOMU);
22350         WRITE(TTY, ' SOTIME = ');
22400         WRITE(TTY, ' ', SOTIME:16, ' SOMSGCOUNT = ', SOMSGCOUNT:16);
22450         WRITELN(TTY, ' SSEQ = ', SSEQ:16);
22500         END;
22550       SINK:BEGIN
22600         WRITE(TTY, ' SIINPORT = ', SIINPORT:13, ' LOCAL TIME = ');
22650         WRITELN(TTY, SITIME:16, ' SIJOBCCOUNT = ', SIJOBCCOUNT:16);
22700         END;
22750       QUEUE:BEGIN
22800         WRITE(TTY, ' Q20INPORT = ', Q20INPORT:13, ' Q20OUTPORT = ');
22850         WRITELN(TTY, Q20OUTPORT:13, ' Q20TIME = ', Q20TIME:16);
22900         WRITE(TTY, ' Q20INPTR = ', Q20INPTR:17, ' Q20OUTPTR = ');
22950         WRITELN(TTY, Q20OUTPTR:17, ' Q20JOBCCOUNT = ', Q20JOBCCOUNT:15);
23000         Q20QMAX:=MAX(0, Q20QMAX);
23050         Q20QMAX:=MIN(19, Q20QMAX);
23100         WRITELN(TTY, ' Q20QMAX = ', Q20QMAX:2);
23150         IF (TARGETTRACE>10) THEN BEGIN
23200           FOR I:=0 TO Q20QMAX DO BEGIN
23250             WRITELN(TTY, ' Q20BUFFER[', I:2, '] CONTAINS MESSAGE ');
23300             WRITE(TTY, ' MESSAGE TIME = ', Q20BUFFER[I]:7);
23350             WRITELN(TTY, ' JOB SEQUENCE NUMBER = ',
23400               Q20BUFFER2[I]:14);
23450             END;
23500           WRITE(TTY, ' Q20BUFFER CONTAINS ', Q20INPTR-Q20OUTPTR:12);
23550           WRITELN(TTY, ' JOBS ');
23600           END;

```

```

23650      END;
23700      MERGE2IBEGIN
23750          WRITE(TTY, ' M2IN1PORT = ', M2IN1PORT:3, ' M2IN2PORT = ');
23800          WRITELN(TTY, M2IN2PORT:3, ' M2OUTPORT = ', M2OUTPORT:3);
23850          WRITE(TTY, ' M2IN1TIME = ', M2IN1TIME:6, ' M2IN2TIME = ');
23900          WRITELN(TTY, M2IN2TIME:6);
23950          WRITE(TTY, ' M2HAVE1 = ', M2HAVE1);
24000          WRITELN(TTY, ' M2HAVE2 = ', M2HAVE2);
24050          WRITE(TTY, ' M2JOB1COUNT = ', M2JOB1COUNT:6);
24100          WRITELN(TTY, ' M2JOB2COUNT = ', M2JOB2COUNT:6);
24150          IF (TARGETTRACE>20) THEN BEGIN
24200              WRITELN(TTY, ' M2IN1MSG CONTAINS THE MESSAGE : ');
24250              SHOWMSG(M2IN1MSG);
24300              WRITELN(TTY, ' M2IN2MSG CONTAINS THE MESSAGE : ');
24350              SHOWMSG(M2IN2MSG);
24400          END;
24450      END;
24500      FORK2IBEGIN
24550          WRITE(TTY, ' F2INPORT = ', F2INPORT:3, ' F2OUT1PORT = ');
24600          WRITELN(TTY, F2OUT1PORT:3, ' F2OUT2PORT = ', F2OUT2PORT:3);
24650          WRITE(TTY, ' F2TIME = ', F2TIME:6, ' F2OUT1COUNT = ');
24700          WRITELN(TTY, F2OUT1COUNT:6, ' F2OUT2COUNT = ', F2OUT2COUNT:6);
24750          WRITE(TTY, ' F2RHD = ', F2RHO);
24800      END;
24850      DELAY1IBEGIN
24900          WRITE(TTY, ' DINPORT = ', DINPORT:3, ' DOUTPORT = ');
24950          WRITELN(TTY, DOUTPORT:3, ' DTIME = ', DTIME:6);
25000          WRITE(TTY, ' DJOBCOUNT = ', DJOBCOUNT:6, ' DEMU = ', DEMU);
25050          WRITELN(TTY, ' DCON = ', DCON);
25100          WRITE(TTY, ' DSUMPTIME = ', DSUMPTIME:6);
25150      END;
25200      END; (* CASE *)
25250      WRITELN(TTY);
25300  END;
25350
25400  PROCEDURE SHOWNETWORK;
25450  (* SHOWS THE PROCESS IDENTITIES, TYPES AND INSTANCES, AND CONNECTIVITY
25500  OF PROCESSES AND PORTS *)
25550  VAR PROCI:PROCID;PORTJ:PORTID;
25600  BEGIN
25650      WRITELN(TTY);
25700      WRITELN(TTY, ' NETWORK DEFINITION FOLLOWS. HERE ARE THE PROCESSES ');
25750      FOR PROCI:=1 TO HIGHPROC DO BEGIN
25800          WRITE(TTY, ' ');
25850          PRSIGNATURE(PROCI);
25900          WRITELN(TTY);
25950      END;
26000      WRITELN(TTY);
26050      WRITELN(TTY, ' HERE ARE THE PORTS ');
26100      WRITE(TTY, ' SENDER RECIIVER PORTID MSGKIND ');
26150      WRITELN(TTY);
26200      FOR PORTJ:=1 TO HIGHPORT DO WITH PORTS[PORTJ] DO BEGIN
26250          WRITE(TTY, ' ', PSENDER:3, ' ', PRECIIVER:3);
26300          WRITE(TTY, ' ', PORTJ:3, ' ');
26350          PRMSGKIND(PMESSAGE, MKIND);
26400          WRITELN(TTY);
26450      END;
26500      WRITELN(TTY);
26550  END;

```

```

26600
26650 PROCEDURE PARWAIT(WHICHPORT:PORTID;REQUESTER:PROCID;MPC:INTEGER);
26700 (* REQUESTING PROCESS INVOKES PROCEDURE PARWAIT TO INITIATE
26750 PARALLEL WAITING FOR A MESSAGE OVER PORT # WHICHPORT.
26800 SHOULD THIS BE THE PORT OVER WHICH THIS PROCESS NEXT
26850 COMMUNICATES, EXECUTION SHOULD RESUME AT THE CONTROL
26900 POINT MPC ( META PROGRAM COUNTER ), THIS PROC WILL
26950 CHECK THAT THE REQUESTING PROCESS IS NAMED AS A PARTY
27000 TO COMMUNICATION OVER THE NAMED PORT, AND THAT THE TYPE OF
27050 MESSAGE SENT TO THIS PORT AGREES WITH THE TYPE NAMED IN THIS
27100 PORT'S PORTREC. AN ERROR MESSAGE IS PRINTED UPON CALL
27150 IF THESE REQUIREMENTS ARE NOT SATISFIED. *)
27200
27250 BEGIN
27300 IF (INTERTRACE+NETBLKTRACE+NETCMNTRACE+TARGETTRACE) > 0 THEN
27350 WRITE(TTY,' PARWAIT CALLED');
27400 IF (INTERTRACE>10) OR (NETBLKTRACE>10) OR (TARGETTRACE>10) THEN BEGIN
27450 WRITE(TTY,' PROCESS ',REQUESTER13,' WILL AWAIT PORT ',WHICHPORT13);
27500 IF (PORTS[WHICHPORT].PSENDER = REQUESTER) THEN BEGIN
27550 WRITE(TTY,' TO SEND A MESSAGE WITH MTIME = ');
27600 WRITELN(TTY,LPS[REQUESTER],AMESSAGE,MTIME);
27650 END;
27700 WRITE(TTY,' AT NETTIME = ',NETTIME+LPS[REQUESTER],AXQTIME16);
27750 WRITELN(TTY,' WITH RESUME POINT ',MPC14);
27800 END;
27850 IF (PORTS[WHICHPORT].PSENDER=REQUESTER) AND
27900 (PORTS[WHICHPORT].PMESSAGE,MKIND<>LPS[REQUESTER],AMESSAGE,MKIND)
27950 THEN BEGIN
28000 WRITE(TTY,' SERIOUS ERROR *** PROCESS ',REQUESTER13);
28050 WRITELN(TTY,' HAS ALTERED THE MESSAGE TYPE OF PORT ');
28100 WRITE(TTY,' NUMBER ',WHICHPORT13,' EFFECTIVELY CHANGING ');
28150 WRITELN(TTY,' ITS TYPE ** ERROR DETECTED IN PARWAIT');
28200 SHOWPROCESS(REQUESTER);
28250 SHOWPORT(WHICHPORT);
28300 WRITE(TTY,' THIS VIOLATES AN INVARIANT OF THE PORT');
28350 WRITELN(TTY,' , FIXED MESSAGE TYPE');
28400 END;
28450 IF PORTS[WHICHPORT].PSENDER = REQUESTER THEN BEGIN (* SEND *)
28500 COPYMSGTOFROM(PORTS[WHICHPORT],PMESSAGE,LPS[REQUESTER],AMESSAGE);
28550 PORTS[WHICHPORT].PSELIGIBLE := TRUE;
28600 PORTS[WHICHPORT].PSWAITING := FALSE;
28650 PORTS[WHICHPORT].PSMPC := MPC;
28700 LPS[REQUESTER].ANEXTSTATE := BLK;
28750 END ELSE
28800 IF PORTS[WHICHPORT].PRECEIVER = REQUESTER THEN BEGIN
28850 (* REQUEST TO RECEIVE OVER THIS LINE *)
28900 PORTS[WHICHPORT].PRELIGIBLE := TRUE;
28950 PORTS[WHICHPORT].PRMPC := MPC;
29000 LPS[REQUESTER].ANEXTSTATE := BLK;
29050 END ELSE BEGIN
29100 WRITE(' THIS ILLEGAL REQUEST IGNORED');
29150 WRITE(TTY,' , REQUESTER ',REQUESTER13,' IS NOT ');
29200 WRITELN(TTY,' CONNECTED');
29250 WRITE(TTY,' TO THE PORT=',WHICHPORT13,' NAMED IN');
29300 WRITELN(TTY,' PARWAIT CALL');
29350 END;
29400 IF (PORTS[WHICHPORT].PSENDER=REQUESTER) AND (NETCMNTRACE>10) THEN BEGIN
29450 WRITELN(TTY,' THE MESSAGE TO BE SENT IS');
29500 SHOWMSG(LPS[REQUESTER],AMESSAGE);

```



procedure that creates the processes, which should be unique to each call. The value of this "uniqueid" becomes the process-identifier of the resulting process created. This is the value that is used as the index into the OWN and LPS arrays for the process. If all the processes are well behaved, and never access the OWN and LPS arrays with arguments other than their own unique identifier, then some level of protection is maintained for the process' local data. Since processes are resumed with this same argument as a calling parameter, "ID", some enforcement of protection would be afforded by disallowing any reference to OWN or LPS with arguments other than this ID, and also flagging as an error any attempt to alter the value of this variable within a process. It would also be necessary to keep the OPROCKIND field of the OWNDATA entry for the process from changing during execution since this would alter the structure of the local data, possibly providing access to the data of another process, so any code that attempted to change this value should be disallowed.

In the SIM QUEUE20 type, initialization occurs at MPC = 1. Here it can be asserted that there is room in the queue; in fact, it is empty. Accordingly, the process initially waits to receive its first message. Whenever input is received, the process resumes at MPC = 2. At this label, it is known that there is at least one message to send, so the process can wait to send it out even without checking the

```

29550      END;
29600      IF (INTERTRACE+NETBLKTRACE+NETCMNTRACE+TARGETTRACE >0) THEN BEGIN
29650          WRITELN(TTY,' PARWAIT RETURNED ');
29700          END;
29750      END;
29800
29850
29900      PROCEDURE DUMPTTY;
29950      (* DISPLAY THE WORLD *)
30000      VAR PROCI:PROCID;PORTJ:PORTID;
30050      BEGIN
30100          WRITELN(TTY);
30150          WRITELN(TTY,' ** PROCEDURE DUMPTTY CALLED ** HERE ARE THE PORTS');
30200          FOR PORTJ:=1 TO HIGHPORT DO SHOWPORT(PORTJ);
30250          WRITELN(TTY,' HERE ARE THE PROCESS RECORDS');
30300          FOR PROCI := 1 TO HIGHPROC DO BEGIN
30350              SHOWPROCESS(PROCI);
30400              SHOWOWN(PROCI);
30450          END;
30500          WRITELN(TTY);
30550          WRITE(TTY,' NETTIME=',NETTIME:8,' NETSUMXQTTIME=');
30600          WRITE(TTY,NETSUMXQTTIME:8,' NETSUMCMNTIME=',NETSUMCMNTIME:8);
30650          WRITELN(TTY,' NETSUMBLKTIME=',NETSUMBLKTIME:8);
30700          WRITE(TTY,' NETMSGLIMIT=',NETMSGLIMIT:6,' NETTIMELIMIT=');
30750          WRITE(TTY,NETTIMELIMIT:8,' NETMSGCOUNT=',NETMSGCOUNT:6);
30800          WRITELN(TTY,' PORTJ=',PORTJ:4,' PROCI=',PROCI:4);
30850          WRITE(TTY,' NETDEADLOCK=',NETDEADLOCK,' NETTERM=',NETTERM);
30900          WRITELN(TTY,' NETCMNTRACE=',NETCMNTRACE:3);
30950          WRITE(TTY,' NETXQTTTRACE=',NETXQTTTRACE:3,' NETBLKTRACE=');
31000          WRITELN(TTY,' INTERTRACE=',INTERTRACE:3);
31050          WRITELN(TTY,' DEADLOCKCOUNT = ',DEADLOCKCOUNT:5);
31100          WRITELN(TTY,' END DUMPTTY ');
31150      END;
31200
31250      PROCEDURE TRACELP(ID:PROCID);
31300      (* PRINT INFORMATION ABOUT THE STATE OF THE LOGICAL PROCESS NAMED
31350      ID. MORE INFORMATION IS PRINTED THE LARGER THE CALLING VALUES
31400      OF THE TRACE VARIABLES, NOTHING IS PRINTED AT ALL IF ALL TRACE
31450      VARS ARE ZERO. CALLERS SHOULD PRINT PERTINANT DETAILS ABOUT
31500      THE LOCATION FROM WHICH THIS ROUTINE IS CALLED, E.G. WHEN
31550      CALLED AT THE BEGINNING OF THE PROCEDURAL DESCRIPTION OF AN LP,
31600      THE CODE MIGHT BE
31650          WRITE(TTY,' ENTER TARGET PROCESS');
31700          TRACELP(ID);
31750          WRITELN(TTY);
31800      END;
31850
31900      BEGIN
31950          IF (INTERTRACE>0) OR (TARGETTRACE>0) OR (NETXQTTTRACE>0) THEN BEGIN
32000              PRSIGNATURE(ID);
32050              WRITELN(TTY);
32100              END;
32150          IF (INTERTRACE>9 ) THEN SHOWPROCESS(ID);
32200          IF (TARGETTRACE>9) THEN SHOWOWN(ID);
32250      END;
32300
32350      PROCEDURE DEFINEPORT(DIRECTION:PORTDIRECTIONTYPE/OWNER:PROCID;
32400          VAR PORTNUM:PORTID);
32450      (* THIS ROUTINE WILL SOLICIT INFORMATION FROM THE USER ABOUT

```

```

32500      A PORT WHOSE DIRECTION, ONE OF INN OR OUT, IS PASSED IN AS
32550      DIRECTION, AND WHOSE OWNER PROCESS IS NAMED AS OWNER, THE
32600      PORTNUMBER WILL BE RETURNED AS PORTNUM, AND THE PORTS ARRAY
32650      WILL BE UPDATED WITH THE USER-SUPPLIED DATA, THIS INCLUDES
32700      THE PORTID, MESSAGE KIND, AND THE SEND OR RECEIVE TIME AS
32750      DETERMINED BY WHETHER THIS IS AN INN, OR AN OUT PORT.  *)
32800
32850      VAR INDATA,I,INTEGER;KIND,MSGKIND;
32900      BEGIN
32950          WRITELN(TTY,' ENTER PORT ID NUMBER');
33000          BREAK;
33050          RESET(TTY);
33100          READ(TTY,PORTNUM);
33150          (* DON'T GET THIS FROM OPERATOR ...
33200          WRITE(TTY,' ENTER THE INTEGER CODE FOR THE TYPE OF MESSAGE ');
33250          WRITELN(TTY,'SENT OVER THIS PORT');
33300          I:=0;
33350          (* TARDEP * NEXT LINE CONTAINS 1ST AND LAST MSGKIND SCALARS  ENDCOMMENT
33400          FOR KIND:=UNDEFINED TO JOB DO BEGIN
33450              WRITE(TTY,'      ',I,2,' * ');
33500              PRMSGKIND(KIND);
33550              INCR(I);
33600          END;
33650          WRITELN(TTY);
33700          BREAK;
33750          RESET(TTY);
33800          READ(TTY,INDATA);
33850
33900
33950          INDATA:=1;
34000          CASE INDATA OF
34050              (* ALL MESSAGE KINDS MUST BE ACCOUNTED FOR IN CASE RANGE *)
34100              0:KIND:=UNDEFINED;
34150              1:KIND:=JOB;
34200              END;
34250          (* WRITE(TTY,' ENTER TIME UNITS FOR ');
34300          CASE DIRECTION OF
34350              INNI:WRITELN(TTY,'INPUT FROM THIS PORT ');
34400              OUTI:WRITELN(TTY,'OUTPUT TO THIS PORT ');
34450              END;
34500          BREAK;
34550          RESET(TTY);
34600          READ(TTY,INDATA); *)
34650          INDATA:=1;
34700          WITH PORTS[PORTNUM] DO BEGIN
34750              CASE DIRECTION OF
34800                  INNI:BEGIN
34850                      RECEIVER:=OWNER;
34900                      PRTIME:=INDATA;
34950                      END;
35000                  OUTI:BEGIN
35050                      SENDE:=OWNER;
35100                      PSTIME:=INDATA;
35150                      END;
35200                  END;
35250              PMESSAGE,MKIND:=KIND;
35300              END;
35350          IF (INTERTRACE>30) OR (TARGETTRACE>30) THEN BEGIN
35400              WRITE(TTY,' PORT ',PORTNUM,3,' IS INCIDENT ON');
              PRSIGNATURE(OWNER);

```

```

WRITELN(TTY, ' HERE IS THE PORT');
SHOWPORT(PORTNUM);
END;

END;

FUNCTION AWAIT(S(PROC:PROCID;PORT:PORTID):BOOLEAN;
(* RETURNS TRUE IFF PROC IS WAITING FOR COMMUNICATION OVER PORT *)
BEGIN
  WITH PORTS(PORT) DO
    AWAIT(S:=(((PRECEIVER=PROC) AND PWAITING AND PRELIGIBLE) OR
              ((SENDER=PROC) AND PSWAITING AND PSELIGIBLE)));
END;

(* *****
*
*   DEFINITION OF TARGET PROCESSES VIA PASCAL PROCEDURES
*
* *****
*)

PROCEDURE LPSOURCE(ID:PROCID);
(* THE TYPE OF THIS LOGICAL PROCESS IS OBTAINED BY REMOVING THE
   'LP' FROM ITS NAME. THIS DETERMINES THE OWNDATA ENTRY FOR
   PROCESSES OF THE TYPE DESCRIBED BY THE FOLLOWING CODE. *)
VAR DELAY:INTEGER;

PROCEDURE CREATEMESSAGE(VAR MSG:MESSAGE;TIME:INTEGER;KIND:MSGKIND;
                        DATA:INTEGER);
(* INITIALIZE THE MESSAGE WITH THE SUPPLIED TIME, KIND, AND SINGLE
   WORD OF DATA. *)
BEGIN
  WITH MSG DO BEGIN
    MTIME:=TIME;
    MKIND:=KIND;
    CASE MKIND OF
      (* TARDEP * DEPENDING ON KIND OF MSG *)
      (* MUST ACCOUNT FOR EVERY POSSIBLE VALUE OF KIND EXP LATER *)
      JOB|JOBNUMBER:=DATA;
      UNDEFINED:=UDATA:=DATA;
    END;
  END;
END;

BEGIN
  IF (TARGETTRACE>0) OR (INTERTRACE>0) THEN BEGIN
    WRITE(TTY, ' (RE)ENTERING ');
    PRSIGNATURE(ID);WRITELN(TTY);END;
    WITH LPS[ID] DO
      WITH OWN[ID] DO BEGIN
        AXQTIME:=1;
        CASE AMPC OF
          0|BEGIN
            (* CREATE THIS LP SOURCE INSTANCE *)
            WRITE(TTY, ' CREATE');
            PRSIGNATURE(ID);
            WRITELN(TTY);
            WRITE(TTY, ' FOR THE OUTPUT PORT,');
            DEFINEPORT(OUT, ID, SOUTPORT);
            WRITE(TTY, ' ENTER REAL EXPONENTIAL DELAY ');

```

```

38650          WRITELN(TTY, ' PARAMETER SOMU ');
38700          BREAK;
38750          RESET(TTY);
38800          READ(TTY, SOMU);
38850          WRITE(TTY, ' ENTER REAL CONSTANT DELAY ');
38900          WRITELN(TTY, 'PARAMETER SOCON');
38950          BREAK;
39000          RESET(TTY);
39050          READ(TTY, SOCON);
39100          END;
39150          11BEGIN      (* FIRST CALL TO THE PROCESS NUMBER ID *)
39200              SOMSGCOUNT:=0;
39250              DELAY:=TRUNC(-SOMU*LN(RANDOM(0))) + TRUNC(SOCON);
39300              IF (SOCON#0,0) THEN INCR(DELAY);
39350              AXGTIME:=DELAY;
39400              SOTIME:=DELAY;
39450              SUSEQ:=1;
39500              CREATEMESSAGE (AMESSAGE, SOTIME, JOB, SOSEQ);
39550              PARWAIT(SOOUTPORT, ID, 2);
39600              END;
39650          21BEGIN      (* JUST SENT OVER SOOUTPORT *)
39700              DELAY:=TRUNC(-SOMU*LN(RANDOM(0))) + TRUNC(SOCON);
39750              IF (SOCON#0,0) THEN INCR(DELAY);
39800              AXGTIME:=DELAY;
39850              SOTIME:=SOTIME+DELAY;
39900              INCR(SOMSGCOUNT);
39950              INCR(SUSEQ);
40000              CREATEMESSAGE (AMESSAGE, SOTIME, JOB, SOSEQ);
40050              PARWAIT(SOOUTPORT, ID, 2);
40100              END;
40150          10011BEGIN  (* LAST CALL FOR ANY STATISTICS *)
40200              WRITE(TTY, ' REPORT FROM');
40250              PRSIGNATURE(ID);
40300              WRITELN(TTY);
40350              SHOWDOWN(ID);
40400              END;
40450          10101BEGIN  (* W=SUB-IJ COMPUTATION *)
40500              WITH PORTS(SOOUTPORT) DO BEGIN
40550                  (* A SOURCE ALWAYS WAITS TO OUTPUT *)
40600                  PHIJ:=SOTIME;
40650                  END;
40700              END;
40750          10201      (* NO ACTION REQUIRED AFTER WIJ FOR SOURCE *)
40800          OTHERS:BEGIN
40850              PRSIGNATURE(ID);
40900              WRITELN(TTY, ' CALLED WITH BAD AMPC = ', AMPC14);
40950              SHOWPROCESS(ID);
41000              SHOWDOWN(ID);
41050              IF (INTERTRACE>20) OR (TARGETTRACE>20) THEN DUMPTY;
41100              END;
41150          END;      (* CASE *)
41200          END;
41250          IF (INTERTRACE>0) OR (TARGETTRACE>0) THEN BEGIN
41300              WRITE(TTY, ' LEAVING ');
41350              TRACELP(ID);
41400              END;
41450          END;      (* LPSOURCE PROCEDURAL DESCRIPTION *)
41500          PROCEDURE LPSINK(ID:PHOCID);
41550

```

```

41600 (* THE TYPE OF THIS LOGICAL PROCESS IS OBTAINED BY REMOVING THE
41650 'LP' FROM ITS NAME, THIS DETERMINES THE OWNDATA ENTRY FOR
41700 PROCESSES OF THE TYPE DESCRIBED BY THE FOLLOWING CODE. *)
41750
41800 VAR DELAY(INTEGER);
41850
41900 PROCEDURE CONSUMEMESSAGE(VAR MSG(MESSAGE));
41950 BEGIN
42000 (* TAKE DEPEND * DO AS YOU WILL WITH THIS MESSAGE BEFORE IT DIES *)
42050 MSG, MKINDI=UNDEFINED;
42100 MSG, MTIME:=0;
42150 END;
42200
42250
42300 BEGIN
42350 IF (TARGETTRACE>0) OR (INTERTRACE>0) THEN BEGIN
42400 WRITE(TTY, ' (RE)-ENTERING ');
42450 PRSIGNATURE(ID); WRITELN(TTY); END;
42500 WITH LPS(ID) DO
42550 WITH OWN(ID) DO BEGIN
42600 AXQTIME:=1;
42650 CASE AMPC OF
42700 0: BEGIN (* CREATE THIS INSTANCE OF A SINK LP *)
42750 WRITE(TTY, ' CREATE ');
42800 PRSIGNATURE(ID);
42850 WRITELN(TTY);
42900 WRITE(TTY, ' FOR THE INPUT PORT, ');
42950 DEFINEPORT(INN, ID, SIINPORT);
43000 END;
43050 1: BEGIN (* FIRST CALL TO THIS SINK PROCESS *)
43100 SITIME:=0;
43150 SIJOBCCOUNT:=0;
43200 PARWAIT(SIINPORT, ID, 2);
43250 END;
43300 2: BEGIN (* JUST RECEIVED A JOB, DO SOMETHING & DISCARD *)
43350 SITIME:=AMESSAGE, MTIME;
43400 INCR(SIJOBCCOUNT);
43450 CONSUMEMESSAGE(AMESSAGE);
43500 PARWAIT(SIINPORT, ID, 2);
43550 END;
43600 1000: BEGIN (* LAST CALL, PRINT LOCAL REPORT *)
43650 WRITE(TTY, ' REPORT FROM ');
43700 PRSIGNATURE(ID);
43750 WRITELN(TTY);
43800 SHOWDOWN(ID);
43850 END;
43900 1010: (* SINK COMPUTES NO W SUB IJ FOR ANY PORTS *)
43950 1020: (* AND CANNOT CHANGE THE AWAITED PORTS *)
44000 OTHERS: BEGIN
44050 PRSIGNATURE(ID);
44100 WRITELN(TTY, ' CALLED WITH BAD AMPC = ', AMPC(4));
44150 SHOWPROCESS(ID);
44200 SHOWDOWN(ID);
44250 IF (INTERTRACE>20) OR (TARGETTRACE>20) THEN DUMPTTY;
44300 END;
44350 END; (* CASE *)
44400 END;
44450 IF (INTERTRACE>0) OR (TARGETTRACE>0) THEN BEGIN
44500 WRITE(TTY, ' LEAVING ');

```

```

44550         TRACEP(ID);
44600         END;
44650     END) (* LPSINK PROCEDURAL DESCRIPTION *)
44700
44750
44800
44850     PROCEDURE LPDELAY(ID;PROCID);
44900     (* THE TYPE OF THIS LOGICAL PROCESS IS OBTAINED BY REMOVING THE
44950     *LP* FROM ITS NAME, THIS DETERMINES THE OWNDATA ENTRY FOR
45000     PROCESSES OF THE TYPE DESCRIBED BY THE FOLLOWING CODE.  *)
45050
45100     VAR DELAY;INTEGER;
45150
45200     BEGIN
45250         IF (TARGETTRACE>0) OR (INTERTRACE>0) THEN BEGIN
45300             WRITE(TTY,' (RE)=ENTERING ');
45350             PRSIGNATURE(ID);WRITELN(TTY);END;
45400         WITH LPS(ID) DO
45450             WITH OWN(ID) DO BEGIN
45500                 AXQTIME:=1;
45550                 CASE AMPC OF
45600                     0:BEGIN (* CREATE THIS LP SINK INSTANCE *)
45650                         WRITE(TTY,' CREATE ');
45700                         PRSIGNATURE(ID);
45750                         WRITELN(TTY);
45800                         WRITE(TTY,' FOR INPUT PORT, ');
45850                         DEFINEPORT(INN,ID,DINPORT);
45900                         WRITE(TTY,' FOR OUTPUT PORT, ');
45950                         DEFINEPORT(OUT,ID,DOUTPORT);
46000                         WRITELN(TTY,' ENTER CONSTANT DELAY REAL PARAMETER DCON');
46050                         BREAK;
46100                         RESET(TTY);
46150                         READ(TTY,DCON);
46200                         WRITELN(TTY,' ENTER EXPONENTIAL DELAY REAL PARAMETER DEMU');
46250                         BREAK;
46300                         RESET(TTY);
46350                         READ(TTY,DEMU);
46400                         END;
46450                     1:BEGIN (* FIRST CALL TO THIS PROCESS *)
46500                         DSUMPTIME:=0;
46550                         DTIME:=0;
46600                         DJOBCOUNT:=0;
46650                         PARWAIT(DINPORT,ID,2);
46700                         END;
46750                     2:BEGIN (* JUST RECEIVED INPUT *)
46800                         DTIME:=MAX(DTIME,AMESSAGE,MTIME);
46850                         INCR(DJOCOUNT);
46900                         DELAY:=TRUNC(-DEMU*LN(RANDOM(0))) + TRUNC(DCON);
46950                         IF (DCON=0,0) THEN INCR(DELAY);
47000                         AXQTIME:=DELAY;
47050                         DSUMPTIME:=DSUMPTIME+DELAY;
47100                         DTIME:=DTIME+DELAY;
47150                         AMESSAGE,MTIME:=DTIME;
47200                         PARWAIT(DOUTPORT,ID,3);
47250                         END;
47300                     3:BEGIN (* JUST SENT OUTPUT *)
47350                         PARWAIT(DINPORT,ID,2);
47400                         END;
47450                     100:BEGIN (* LAST CALL, PRINT REPORT *)

```

```

47500      WRITE(TTY,' REPORT FROM');
47550      PRSIGNATURE(ID);
47600      WRITELN(TTY);
47650      SHOWOWN(ID);
47700      END;
47750      1010:BEGIN      (* W-SUB IJ COMPUTATION *)
47800          IF AWAITS(ID,DDOUTPORT) THEN BEGIN
47850              PORTS(DDOUTPORT),PWIJI=DTIME;
47900          END ELSE BEGIN      (* AWAITING INPUT *)
47950              PORTS(DDOUTPORT),PWIJI=PORTS(DINPORT),PWIJ;
48000          END;
48050      END;
48100      1020: (* NO CHANGE OF AWAITED PORTS CAN OCCUR *)
48150      OTHERS:BEGIN
48200          PRSIGNATURE(ID);
48250          WRITELN(TTY,' CALLED WITH BAD AMPC = ',AMPC);
48300          SHOWPROCESS(ID);
48350          SHOWOWN(ID);
48400          IF (INTERTRACE>20) OR (TARGETTRACE>20) THEN DUMPTTY;
48450          END;
48500      END;      (* CASE *)
48550      END;
48600      IF (INTERTRACE>0) OR (TARGETTRACE>0) THEN BEGIN
48650          WRITE(TTY,' LEAVING ');
48700          TRACELP(ID);
48750          END;
48800      END;      (* LPDELAY PROCEDURAL DESCRIPTION *)
48850
48900
48950
49000
49050      PROCEDURE LPMERGE2(ID:PROCID);
49100      (* THE TYPE OF THIS LOGICAL PROCESS IS OBTAINED BY REMOVING THE
49150      'LP' FROM ITS NAME. THIS DETERMINES THE OWNDATA ENTRY FOR
49200      PROCESSES OF THE TYPE DESCRIBED BY THE FOLLOWING CODE. *)
49250
49300      PROCEDURE DECIDENEXT(ID:PROCID);
49350      (* DECIDES THE NEXT AWAITED LINES FOR MERGE2
49400      ACCORDING TO THE FOLLOWING ACTION TABLE
49450
49500      * HAVE1 * HAVE2 * INITIME<IN2TIME * INITIME=IN2TIME * IN2TIME<INITIME
49550      *****
49600      * FALSE * FALSE * AWAIT BOTH INPT * AWAIT BOTH INPT * AWAIT BOTH
49650      * FALSE * TRUE * WAIT 1 INPUT * WAIT JOB2 OUT * WAIT JOB2 OUT
49700      * TRUE * FALSE * WAIT JOB1 OUT * WAIT JOB1 OUT * WAIT 2 INPUT
49750      * TRUE * TRUE * WAIT JOB1 OUT * WAIT JOB2** OUT * WAIT JOB2 OUT
49800
49850      THE ** ENTRY COULD ALSO BE WAIT JOB1 OUT,
49900      IT IS ASSUMED THAT ALL PREVIOUSLY SCHEDULED OUTPUT HAS BEEN SENT
49950      OUT. *)
50000      BEGIN
50050          WITH OWN[ID] DO
50100          WITH LPS[ID] DO
50150          IF NOT (M2HAVE1 OR M2HAVE2) THEN BEGIN
50200              (* AWAIT INPUT FROM BOTH INPUT PORTS, 1 AND 2 *)
50250              PARWAIT(M2IN1PORT,ID,2);
50300              PARWAIT(M2IN2PORT,ID,3);
50350          END ELSE
50400              IF (NOT M2HAVE1 AND (M2IN1TIME<M2IN2TIME)) THEN BEGIN

```



```

50450      (* AWAIT INPUT FROM PORT 1 *)
50500      PARWAIT(M2IN1PORT, ID, 2);
50550      END ELSE IF (M2IN1TIME < M2IN2TIME) THEN BEGIN
50600          (* SEND MESSAGE FROM PORT 1 *)
50650          M2HAVE1 := FALSE;
50700          COPYMSGTOFROM(AMESSAGE, M2IN1MSG);
50750          AMESSAGE, MTIME := MIN(M2IN1TIME, M2IN2TIME);
50800          PARWAIT(M2OUTPORT, ID, 4);
50850      END ELSE IF M2HAVE2 THEN BEGIN
50900          (* SEND MESSAGE FROM PORT 2 *)
50950          M2HAVE2 := FALSE;
51000          COPYMSGTOFROM(AMESSAGE, M2IN2MSG);
51050          AMESSAGE, MTIME := MIN(M2IN1TIME, M2IN2TIME);
51100          PARWAIT(M2OUTPORT, ID, 4);
51150      END ELSE IF (M2IN1TIME = M2IN2TIME) THEN BEGIN
51200          (* SEND MESSAGE FROM PORT 1 *)
51250          M2HAVE1 := FALSE;
51300          COPYMSGTOFROM(AMESSAGE, M2IN1MSG);
51350          AMESSAGE, MTIME := MIN(M2IN1TIME, M2IN2TIME);
51400          PARWAIT(M2OUTPORT, ID, 4);
51450      END ELSE BEGIN
51500          (* AWAIT INPUT FROM 2 *)
51550          PARWAIT(M2IN2PORT, ID, 3);
51600      END;
51650  END; (* DECIDENEXT PROCEDURE...USED ONLY BY MERGE2 *)
51700
51750  BEGIN
51800      IF (TARGETTRACE > 0) OR (INTERTRACE > 0) THEN BEGIN
51850          WRITE(TTY, " (RE)-ENTERING ");
51900          PRSIGNATURE(ID); WRITELN(TTY); END;
51950      WITH LPS(ID) DO
52000          WITH OWN(ID) DO BEGIN
52050              AXQTIME := 1;
52100              CASE AMPC OF
52150                  0: BEGIN          (* CREATE THIS INSTANCE *)
52200                      WRITE(TTY, " CREATE ");
52250                      PRSIGNATURE(ID);
52300                      WRITELN(TTY);
52350                      WRITE(TTY, " FOR INPUT PORT 1, ");
52400                      DEFINEPORT(INN, ID, M2IN1PORT);
52450                      WRITE(TTY, " FOR INPUT PORT 2, ");
52500                      DEFINEPORT(INN, ID, M2IN2PORT);
52550                      WRITE(TTY, " FOR OUTPUT PORT, ");
52600                      DEFINEPORT(OUT, ID, M2OUTPORT);
52650                      END;
52700                  1: BEGIN          (* FIRST CALL *)
52750                      M2IN1TIME := 0;
52800                      M2JOB1COUNT := 0;
52850                      M2JOB2COUNT := 0;
52900                      M2IN2TIME := 0;
52950                      M2HAVE1 := FALSE;
53000                      M2HAVE2 := FALSE;
53050                      M2IN1MSG, MKIND1 := JOB;
53100                      M2IN2MSG, MKIND2 := JOB;
53150                      DECIDENEXT(ID);
53200                      END;
53250                  2: BEGIN          (* JUST RECEIVED OVER IN1 *)
53300                      M2HAVE1 := TRUE;
53350                      M2IN1TIME := AMESSAGE, MTIME;

```

```

53400      INCR(M2JOB1COUNT)
53450      COPYMSGTOFROM(M2IN1MSG, AMESSAGE)
53500      DECIDENEXT(ID)
53550      END
53600      3IBEGIN      (* JUST RECEIVED OVER IN2 *)
53650      M2HAVE2:=TRUE
53700      M2IN2TIME:=AMESSAGE, MTIME
53750      INCR(M2JOB2COUNT)
53800      COPYMSGTOFROM(M2IN2MSG, AMESSAGE)
53850      DECIDENEXT(ID)
53900      END
53950      4IBEGIN      (* JUST SENT OVER OUTPUT *)
54000      DECIDENEXT(ID)
54050      END
54100      1000IBEGIN   (* PRINT FINAL REPORT *)
54150      WRITE(TTY, ' REPORT FROM')
54200      PRSIGNATURE(ID)
54250      WRITELN(TTY)
54300      SHOWUWN(ID)
54350      END
54400      1010IBEGIN   (* W-SUB IJ COMPUTATION *)
54450      IF AWAITS(ID, M2OUTPORT) THEN (* WAITING TO OUTPUT *)
54500      WITH PORTS(M2OUTPORT) DO
54550      PWIJ:=MIN(M2IN1TIME, M2IN2TIME)
54600      ELSE PORTS(M2OUTPORT), PWIJ:=
54650      MIN( PORTS(M2IN1PORT), PWIJ,
54700      PORTS(M2IN2PORT), PWIJ)
54750      END (* W SUB IJ COMPUTATION *)
54800      1020IBEGIN   (* ATTEMPT TO CHANGE LINES AWAITED *)
54850      IF NOT AWAITS(ID, M2OUTPORT) THEN BEGIN
54900      (* MOVE LINE TIMES FORWARD IF POSSIBLE *)
54950      M2IN1TIME:=MAX(M2IN1TIME, PORTS(M2IN1PORT), PWIJ)
55000      M2IN2TIME:=MAX(M2IN2TIME, PORTS(M2IN2PORT), PWIJ)
55050      (* TRY FOR A DIFFERENT SET OF LINES *)
55100      PORTS(M2IN1PORT), PRELIGIBLE:=FALSE
55150      PORTS(M2IN2PORT), PRELIGIBLE:=FALSE
55200      DECIDENEXT(ID)
55250      PORTS(M2IN1PORT), PRWAITING:=TRUE
55300      PORTS(M2IN2PORT), PRWAITING:=TRUE
55350      PORTS(M2OUTPORT), PSWAITING:=TRUE
55400      END
55450      END
55500      OTHERSIBEGIN
55550      PRSIGNATURE(ID)
55600      WRITELN(TTY, ' CALLED WITH BAD AMPC = ', AMPC(4))
55650      SHOWPROCESS(ID)
55700      SHOWDOWN(ID)
55750      IF (INTERTRACE>20) OR (TARGETTRACE>20) THEN DUMPTTY
55800      END
55850      END (* CASE *)
55900      END
55950      IF (INTERTRACE>0) OR (TARGETTRACE>0) THEN BEGIN
56000      WRITE(TTY, ' LEAVING ')
56050      TRACE(LP, ID)
56100      END
56150      END
56200      END (* LPMERGE2 PROCEDURAL DESCRIPTION *)
56250
56300

```

```

56350 PROCEDURE LPFORK2(IDIPROCID)
56400 (* THE TYPE OF THIS LOGICAL PROCESS IS OBTAINED BY REMOVING THE
56450 "LP" FROM ITS NAME. THIS DETERMINES THE OWNDATA ENTRY FOR
56500 PROCESSES OF THE TYPE DESCRIBED BY THE FOLLOWING CODE. *)
56550
56600
56650 BEGIN
56700 IF (TARGETTRACE>0) OR (INTERTRACE>0) THEN BEGIN
56750 WRITE(TTY,' (RE)-ENTERING ');
56800 PRSIGNATURE(ID);WRITELN(TTY);END;
56850 WITH LPS[ID] DO
56900 WITH OWN[ID] DO BEGIN
56950 AXOTIME:=1;
57000 CASE AMPC OF
57050 0:BEGIN (* CALL TO CREATE THE PROCESS *)
57100 WRITE(TTY,' CREATE');
57150 PRSIGNATURE(ID);
57200 WRITELN(TTY);
57250 WRITE(TTY,' FOR INPUT PORT,');
57300 DEFINEPORT(INN,ID,F2INPORT);
57350 WRITE(TTY,' FOR OUTPUT PORT 1,');
57400 DEFINEPORT(OUT,ID,F2OUT1PORT);
57450 WRITE(TTY,' FOR OUTPUT PORT 2,');
57500 DEFINEPORT(OUT,ID,F2OUT2PORT);
57550 WRITE(TTY,' WHAT IS PROBABILITY OF A BRANCH ');
57600 WRITELN(TTY,' TO OUTPUT PORT 1 ?');
57650 BREAK;
57700 RESET(TTY);
57750 READ(TTY,F2RHO);
57800 END;
57850 1:BEGIN (* FIRST CALL TO THIS INSTANCE OF FORK2 *)
57900 F2TIME:=0;
57950 F2OUT1COUNT:=0;
58000 F2OUT2COUNT:=0;
58050 PARWAIT(F2INPORT,ID,2);
58100 END;
58150 2:BEGIN (* JUST RECEIVED INPUT OVER F2INPORT *)
58200 F2TIME:=AMESSAGE,MTIME;
58250 IF (RANDOM(0)<=F2RHO) THEN BEGIN (* SEND ON 1 *)
58300 INCR(F2OUT1COUNT);
58350 PARWAIT(F2OUT1PORT,ID,3);
58400 END ELSE BEGIN (* SEND ON 2 *)
58450 INCR(F2OUT2COUNT);
58500 PARWAIT(F2OUT2PORT,ID,3);
58550 END;
58600 3:BEGIN (* JUST SENT A JOB OUT *)
58650 PARWAIT(F2INPORT,ID,2);
58700 END;
58750 1000:BEGIN (* FINAL CALL FOR REPORTS *)
58800 WRITE(TTY,' REPORT FROM');
58850 PRSIGNATURE(ID);
58900 WRITELN(TTY);
58950 SHOWOWN(ID);
59000 END;
59100 1010:BEGIN (* W SUB IJ COMPUTATION *)
59150 IF AWAITS(ID,F2INPORT) THEN BEGIN (* WAITING INPUT *)
59200 PURTS(F2OUT1PORT),PWIJ:=PURTS(F2INPORT),PWIJ;
59250

```

```

59300      PORTS(F2OUT2PORT),PHIJI=PORTS(F2INPORT),PHIJI
59350      END ELSE IF AWAITS(ID,F2OUT1PORT) THEN BEGIN
59400      (* AWAITING OUTPUT ON OUTPUT PORT 1 *)
59450      PORTS(F2OUT1PORT),PHIJI=F2TIME)
59500      PORTS(F2OUT2PORT),PHIJI=PORTS(F2INPORT),PHIJI
59550      END ELSE BEGIN (* AWAITING OUTPUT OVER PORT 2 *)
59600
59650      PORTS(F2OUT2PORT),PHIJI=F2TIME)
59700      PORTS(F2OUT1PORT),PHIJI=PORTS(F2INPORT),PHIJI
59750      END)
59800      END) (* W SUB IJ COMPUTATION *)
59850      10201) (* AFTER W SUB IJ COMPUTATION, FORK AWAITS THE
59900      SAME PORTS IT DID BEFORE *)
59950
60000
60050      OTHERS:BEGIN
60100      PRSIGNATURE(ID);
60150      WRITELN(TTY,' CALLED WITH BAD AMPC = ',AMPC);
60200      SHOWPROCESS(ID);
60250      SHOWDOWN(ID);
60300      IF (INTERTRACE>20) OR (TARGETTRACE>20) THEN DUMPTTY)
60350      END)
60400      END) (* CASE *)
60450      END)
60500      IF (INTERTRACE>0) OR (TARGETTRACE>0) THEN BEGIN
60550      WRITE(TTY,' LEAVING ');
60600      TRACE(LP(ID));
60650      END)
60700      END) (* LPFORK2 PROCEDURAL DESCRIPTION *)
60750
60800
60850      PROCEDURE LPQUEUE20(ID:PROCID);
60900      (* THE TYPE OF THIS LOGICAL PROCESS IS OBTAINED BY REMOVING THE
60950      'LP' FROM ITS NAME, THIS DETERMINES THE OWNDATA ENTRY FOR
61000      PROCESSES OF THE TYPE DESCRIBED BY THE FOLLOWING CODE, *)
61050
61100      VAR I:INTEGER;
61150
61200      BEGIN
61250      IF (TARGETTRACE>0) OR (INTERTRACE>0) THEN BEGIN
61300      WRITE(TTY,' (RE)-ENTERING ');
61350      PRSIGNATURE(ID);WRITELN(TTY);END)
61400      WITH LP8[ID] DO
61450      WITH OWN[ID] DO BEGIN
61500      AXQTIME:=1;
61550      CASE AMPC OF
61600      0:BEGIN (* CREATE *)
61650      WRITE(TTY,' CREATE');
61700      PRSIGNATURE(ID);
61750      WRITELN(TTY);
61800      WRITE(TTY,' FOR INPUT PORT,');
61850      DEFINEPORT(INN,ID,Q20INPORT);
61900      WRITE(TTY,' FOR OUTPUT PORT,');
61950      DEFINEPORT(OUT,ID,Q20OUTPORT);
62000      WRITELN(TTY,' ENTER QUEUE CAPACITY, 1,,20 ');
62050      BREAK;
62100      RESET(TTY);
62150      READ(TTY,Q20QMAX);
62200      Q20QMAX:=Q20QMAX-1;

```

```

62250          Q20QMAX1=MAX(0,Q20QMAX)
62300          Q20QMAX1=MIN(19,Q20QMAX)
62350          END)
62400          1)BEGIN (* FIRST CALL TO THIS INSTANCE OF A QUEUE20 *)
62450              Q20INPTR1=1)
62500              Q20OUTPTR1=1)
62550              Q20TIME1=0)
62600              Q20JOB COUNT1=0)
62650              PARWAIT(Q20INPORT, ID, 2)
62700              END)
62750          2)BEGIN (* JUST RECEIVED OVER INPORT *)
62800              Q20TIME1=AMESSAGE, MTIME)
62850              INCR(Q20JOB COUNT)
62900              Q20BUFFER(Q20INPTR MOD 20)1=AMESSAGE, MTIME)
62950              Q20BUFFER2(Q20INPTR MOD 20)1=AMESSAGE, JOB NUMBER)
63000              INCR(Q20INPTR)
63050              IF (Q20INPTR<=Q20OUTPTR+Q20QMAX) THEN
63060                  PARWAIT(Q20INPORT, ID, 2)
63100              AMESSAGE, MTIME1=Q20TIME)
63150              PARWAIT(Q20OUTPORT, ID, 3)
63200              END)
63250          3)BEGIN (* JUST SENT OVER OUTPORT *)
63300              INCR(Q20OUTPTR)
63350              IF (Q20OUTPTR<Q20INPTR) THEN BEGIN
63400                  AMESSAGE, MTIME1=Q20TIME)
63450                  AMESSAGE, MKIND1=JOB)
63500                  AMESSAGE, JOB NUMBER1=Q20BUFFER2(Q20OUTPTR MOD 20)
63550                  PARWAIT(Q20OUTPORT, ID, 3)
63600                  END)
63650              PARWAIT(Q20INPORT, ID, 2)
63700              END)
63750          1000)BEGIN (* FINAL CALL FOR REPORT *)
63800              WRITE(TTY, ' REPORT FROM')
63850              PRSIGNATURE(ID)
63900              WRITELN(TTY)
63950              SHOWOWN(ID)
64000              END)
64050          1010)BEGIN (* W SUB IJ COMPUTATION *)
64100              IF AWAITS(ID, Q20OUTPORT) THEN BEGIN (* WAIT OUTPUT *)
64150                  PORTS(Q20OUTPORT), PWIJ1=Q20TIME)
64200              END ELSE BEGIN
64250                  PORTS(Q20OUTPORT), PWIJ1=PORTS(Q20INPORT), PWIJ)
64300                  END)
64350              END)
64400          1020) (* A QUEUE CANNOT CHANGE AWAITED LINES *)
64450
64500
64550          OTHERS)BEGIN
64600              PRSIGNATURE(ID)
64650              WRITELN(TTY, ' CALLED WITH BAW AMPC = ', AMPC(14))
64700              SHOWPROCESS(ID)
64750              SHOWOWN(ID)
64800              IF (INTERTRACE>20) OR (TARGETTRACE>20) THEN DUMPTTY)
64850              END)
64900          END) (* CASE *)
64950          END)
65000          IF (INTERTRACE>0) OR (TARGETTRACE>0) THEN BEGIN
65050              WRITE(TTY, ' LEAVING ')
65100              TRACELP(ID)

```

```

65150         END)
65200     END) (* LPQUEUE20 PROCEDURAL DESCRIPTION *)
65250
65300     (* HERE IS A BABY LP TO PLAY WITH ...
65350
65400
65450     PROCEDURE LPNAME(ID)
65500     (* THE TYPE OF THIS LOGICAL PROCESS IS OBTAINED BY REMOVING THE
65550       *LP* FROM ITS NAME. THIS DETERMINES THE OWNDATA ENTRY FOR
65600       PROCESSES OF THE TYPE DESCRIBED BY THE FOLLOWING CODE.      ENDCOMMENT
65650
65700     VAR
65750
65800     BEGIN
65850         IF (TARGETTRACE>0) OR (INTERTRACE>0) THEN BEGIN
65900             WRITE(TTY, ' (RE)-ENTERING ')
65950             PRSIGNATURE(ID)/WRITELN(TTY)/END)
66000         WITH LPB(ID) DO
66050             WITH OWN(ID) DO BEGIN
66100                 AXQTIME:=1
66150                 CASE AMPC OF
66200                     0:BEGIN
66250                         PRSIGNATURE(ID)
66300                     END)
66350
66400
66450
66500
66550             OTHERS:BEGIN
66600                 PRSIGNATURE(ID)
66650                 WRITELN(TTY, ' CALLED WITH BAD AMPC = ',AMPC(4))
66700                 SHOWPROCESS(ID)
66750                 SHOWOWN(ID)
66800                 IF (INTERTRACE>20) OR (TARGETTRACE>20) THEN DUMPTTY)
66850                 END)
66900             END) (* CASE ENDCOMMENT
66950             END)
67000         IF (INTERTRACE>0) OR (TARGETTRACE>0) THEN BEGIN
67050             WRITE(TTY, ' LEAVING ')
67100             TRACELP(ID)
67150             END)
67200     END) (* LPNAME PROCEDURAL DESCRIPTION ENDCOMMENT
67250
67300
67350     END OF THE COMMENT CONTAINING THE PROCEDURAL TEMPLATES      *)
67400
67450
67500
67550
67600
67650
67700     PROCEDURE INITIALIZE)
67750     BEGIN
67800         NETOEADLOCK:=FALSE)
67850         NETTERM:=FALSE)
67900         NETSUMXQTIME:=0)
67950         NETSUMCMNTIME:=0)
68000         NETSUMBLKTIME:=0)
68050         DEADLOCKCOUNT:=0)

```

```

68100     NETTIME:=0!
68150     NETFAIR:=1!
68175     NETMSGCOUNT:=0!
68200     END!
68250
68300     PROCEDURE SETTRACE!
68350     (* SOLICITS VALUES FOR TRACE VARS AND TIME AND MESSAGE LIMITS
68400     FROM THE OPERATOR *)
68450     BEGIN
68500         WRITELN(TTY,' SET TRACE VALUES, 0==> NO TRACE, BIG==> MORE TRACE'!)
68550         WRITELN(TTY,' REENTER THE INTERPRETER TRACE VALUE'!)
68600         BREAK!
68650         RESET(TTY)!
68700         READ(TTY,INTERTRACE)!
68750         WRITELN(TTY,' ECHO ',INTERTRACE!4,' ENTER TARGET TRACE '!)
68800         BREAK!
68850         RESET(TTY)!
68900         READ(TTY,TARGETTRACE)!
68950         WRITELN(TTY,' ECHO ',TARGETTRACE!4,' ENTER DEADLOCK TRACE '!)
69000         BREAK!
69050         RESET(TTY)!
69100         READ(TTY,NETDEADTRACE)!
69150         WRITELN(TTY,' ECHO ',NETDEADTRACE!4,' ENTER COMMUNICATION TRACE'!)
69200         BREAK!
69250         RESET(TTY)!
69300         READ(TTY,NETCMNTRACE)!
69350         WRITELN(TTY,' ECHO ',NETCMNTRACE!3,' ENTER EXECUTION TRACE'!)
69400         BREAK!
69450         RESET(TTY)!
69500         READ(TTY,NETXQTRACE)!
69550         WRITELN(TTY,' ECHO ',NETXQTRACE!3,' ENTER BLOCKING TRACE'!)
69600         BREAK!
69650         RESET(TTY)!
69700         READ(TTY,NETBLKTRACE)!
69750         WRITE(TTY,' ECHO ',NETBLKTRACE!3,' MESSAGE COUNT= ',NETMSGCOUNT)!
69800         WRITELN(TTY,' ENTER MESSAGE LIMIT'!)
69850         BREAK!
69900         RESET(TTY)!
69950         READ(TTY,NETMSGLIMIT)!
70000         WRITELN(TTY,' ECHO ',NETMSGLIMIT,' TIME NOW = ',NETTIME)!
70050         WRITELN(TTY,' ENTER TIME LIMIT'!)
70100         BREAK!
70150         RESET(TTY)!
70200         READ(TTY,NETTIMELIMIT)!
70250         WRITE(TTY,' ECHO ',NETTIMELIMIT!6,' THERE HAVE BEEN '!)
70300         WRITELN(TTY,DEADLOCKCOUNT!4,' DEADLOCKS, ENTER NEW LIMIT'!)
70350         BREAK!
70400         RESET(TTY)!
70450         READ(TTY,DEADLOCKLIMIT)!
70500         WRITELN(TTY,' ECHO ',DEADLOCKLIMIT!4,' END SETTRACE '!)
70550     END!
70600
70650     PROCEDURE RESUME(PROCNUM!PROCID)!
70700
70750     (* RESUME IS INVOKED WHEN THE PROCESS PROCNUM ENTERS THE EXECUTING
70800     STATE. PROCNUM ALSO INDEXES THE ARRAYS LPS AND OWN CONTAINING
70850     THE PROCESS' STATE AND OWNED DATA RESPECTIVELY. *)
70900
70950

```

```

71000 BEGIN
71050 IF (INTERTRACE+NETCMNTRACE+NETXQTRACE+TARGETTRACE) > 0 THEN
71100 WRITELN(TTY,' PROCESS ',PROCNUM:3,' RESUMED');
71150 LPS[PROCNUM],ASTATE := XQT;
71200 LPS[PROCNUM],ANEXTSTATE := TRM; (* PROCESS MAY SET TO BLK WITH PARWAIT *)
71250 (* IN W-SUB IJ COMPUTATION, PROCESSES REMAIN BLOCKED UNLESS
71300 IN THE PROCESS OF COMPUTATION, A PROCESS DETERMINES ITS NEXT
71350 STATE WILL BE TERMINATED, AND SETS ITS ANEXTSTATE ACCORDINGLY *)
71400 IF (LPS[PROCNUM],AMPC=1010) OR (LPS[PROCNUM],AMPC=1020) THEN BEGIN
71450 LPS[PROCNUM],ASTATE:=BLK;
71500 LPS[PROCNUM],ANEXTSTATE:=BLK;
71550 END;
71600 CASE LPS[PROCNUM],ATYPE OF
71650
71700 (* TARDEP * RESUME PROCESS PROCEDURE'S PROGRAM *)
71750
71800 SOURCE:=LPSOURCE[PROCNUM];
71850 SINK:=LPSINK[PROCNUM];
71900 FORK2:=LPFORK2[PROCNUM];
71950 MERGE2:=LPMERGE2[PROCNUM];
72000 DELAY:=LPDELAY[PROCNUM];
72050 QUEUE2:=LPQUEUE2[PROCNUM];
72100 END;
72150 WITH LPS[PROCNUM] DO BEGIN
72200 ATIMELEFT := AXQTIME;
72250 END;
72300 END;
72350
72400 PROCEDURE FIREPORT(PORTNUM:=PORTID);
72450
72500 (* THE PORT PORTNUM IS COMMITTED TO FIRE, THE MESSAGE IS COPIED
72550 FROM THE PORT MESSAGE BUFFER INTO THE RECEIVER'S BUFFER, THE
72600 SENDING AND RECEIVING PROCESSES ARE MARKED AS STATE = COM-
72650 MUNICATING, AND THEIR TIME REMAINING FIELDS GET THE SENDTIME
72700 AND RECEIVETIME SPECIFIC TO THIS PORT, THE MESSAGE COUNTS ARE
72750 INCREMENTED FOR THE PORT, AND NETWORK TOTAL, THE SENDER AND
72800 RECEIVER ARE DISQUALIFIED FROM FURTHER COMMUNICATION BY SETTING
72850 APPROPRIATE ELIGIBLE FIELDS FOR ALL PORTS NAMING THEM, THE
72900 PROCESSES WILL NOT BE ACTUALLY RESUMED UNTIL THEY HAVE WAITED OUT
72950 THEIR COMMUNICATION TIME, *)
73000
73050 VAR I:=PORTID;
73100 BEGIN
73150 IF (INTERTRACE+NETBLKTRACE+NETCMNTRACE+TARGETTRACE) > 0 THEN BEGIN
73200 WRITELN(TTY);
73250 WRITELN(TTY,' PORT ',PORTNUM:3,' FIRED');
73300 END;
73350 IF (INTERTRACE>10) OR (NETBLKTRACE>10)
73400 OR (NETCMNTRACE>10) OR (TARGETTRACE>10) THEN
73450 SHOWMSG(PORTS[PORTNUM],PMESSAGE);
73500 INCR(NETHSGCOUNT);
73550 INCR(PORTS[PORTNUM],PMSGCOUNT);
73600 WITH PORTS[PORTNUM] DO BEGIN
73650 WITH LPS[SENDER] DO BEGIN
73700 IF (INTERTRACE>10) OR (NETBLKTRACE>10)
73750 OR (NETCMNTRACE>10) OR (TARGETTRACE>10) THEN BEGIN
73800 WRITE(TTY,' PROCESS ',PSENDER:3,' GOES FROM ');
73850 WRITELN(TTY,' BLK TO CMN AS SENDER');
73900 END;

```



```

73950         AMPC := PSMPC;
74000         ATIMELEFT := PSTIME;
74050         ASTATE := CMN;
74100         APORT := PORTNUM;
74150         END;
74200     WITH LPS[PRECEIVER] DO BEGIN
74250         IF (INTERTRACE>10) OR (NETBLKTRACE>10)
74300         OR (NETCMNTRACE>10) OR (TARGETTRACE>10) THEN BEGIN
74350             WRITE(TTY, ' PROCESS ', PRECEIVER, ' GOES FROM');
74400             WRITELN(TTY, ' BLK TO CMN AS RECEIVER');
74450             END;
74500         ATIMELEFT := PRTIME;
74550         ASTATE := CMN;
74600         APORT := PORTNUM;
74650         AMPC := PRMPC;
74700         COPYMSGTOFROM(AMESSAGE, PMESSAGE);
74750         END;
74800     END;
74850     (* TURN OFF ELIGIBLE BITS FOR PORTS NAMING PROCESS I *)
74900     FOR I:=1 TO HIGHPORT DO BEGIN
74950         PORTS[I], PSELIGIBLE := (PORTS[I], PSELIGIBLE) AND
75000         (PORTS[I], PSENDER<>PORTS[PORTNUM], PSENDER);
75050         PORTS[I], PRELIGIBLE := (PORTS[I], PRELIGIBLE) AND
75100         (PORTS[I], PRECEIVER<>PORTS[PORTNUM], PRECEIVER);
75150     END;
75200 END;
75250
75300
75350 PROCEDURE SCHEDULE(VAR PORTFIRINGORDER:PORTPOINTERS);
75400
75450     (* USER MAY PROVIDE THE SCHEDULER PROCEDURE TO CHANGE THE WAY THE
75500     PORTS ARE SELECTED TO FIRE. AT END, PORTFIRINGORDER[I]=0 IFF
75550     THERE ARE NOT AS MANY AS I READY PORTS, AND PORTFIRINGORDER[I]=J
75600     IFF J IS TO BE THE ITH PORT TO FIRE.
75650     THIS IMPLEMENTATION IS FAIR BECAUSE IT FAVORS THE PORT = NETFAIR
75700     IF THIS PORT IS READY TO FIRE, AND NETFAIR IS ALWAYS INCREMENTED
75750     MODULO HIGHPORT WHEN PROCEDURE SCHEDULE IS INVOKED. HENCE HIGHPORT
75800     IS AN UPPER LIMIT TO THE TIME THAT A READY PORT CAN WAIT FOR
75850     COMMUNICATION. NO PROCESS IS SCHEDULED FOR
75900     MORE THAN ONE MESSAGE FIRING. *)
75950
76000     VAR OKFIRE:ARRAY[1..PROCMAX] OF BOOLEAN;NEXT,PROCI:PROCID;PORTJ:PORTID;
76050
76100
76150
76200     BEGIN
76250     IF (INTERTRACE>0)OR(NETCMNTRACE>0) THEN WRITELN(TTY, ' SCHEDULE CALLED');
76300     IF (INTERTRACE>30) OR (NETCMNTRACE>30) THEN BEGIN
76350         WRITELN(TTY, ' HERE ARE THE PORTS');
76400         FOR PORTJ:=1 TO HIGHPORT DO SHOWPORT(PORTJ);
76450         END;
76500     FOR PORTJ:=1 TO PORTMAX DO PORTFIRINGORDER[PORTJ]:=0;
76550     FOR PROCI:=1 TO PROCMAX DO OKFIRE[PROCI]:=TRUE;
76600     NEXT := 1;
76650     WITH PORTS[NETFAIR] DO
76700         IF PRWAITING AND PSWAITING AND PRELIGIBLE AND PSELIGIBLE THEN BEGIN
76750             PORTFIRINGORDER[NEXT]:=NETFAIR;
76800             NEXT:=2;
76850             IF NETFAIR=HIGHPORT THEN NETFAIR:=1 ELSE INCR(NETFAIR);

```

```

76900         OKFIRE[PSENDER] := FALSE;
76950         OKFIRE[PRECEIVER] := FALSE;
77000         END;
77050     FOR PORTJ := 1 TO HIGHPORT DO WITH PORTS[PORTJ] DO
77100         IF PRWAITING AND PSWAITING AND PRELIGIBLE AND PSELIGIBLE
77150         AND OKFIRE[PSENDER] AND OKFIRE[PRECEIVER] THEN BEGIN
77200             PORTFIRINGORDER[NEXT] := PORTJ;
77250             INCR(NEXT);
77300             OKFIRE[PSENDER] := FALSE;
77350             OKFIRE[PRECEIVER] := FALSE;
77400         END;
77450     END;
77500
77550     PROCEDURE PASSMESSAGES;
77600
77650     (* SCAN THE COMMUNICATIONS TABLES, IN THE GLOBAL PORTS, AND
77700     SELECT THE PORTS THAT WILL FIRE DURING THE CURRENT MOMENT
77750     OF TIME, AND FIRE THEM. PASSMESSAGES HANDLES THINGS PORTS DO *)
77800
77850     VAR PORTSTOBEFIRED: ARRAY[1..PORTMAX] OF INTEGER; NEXT: INTEGER;
77900
77950     BEGIN
78000         IF (INTERTRACE > 0) OR (NETCMNTRACE > 0) THEN
78050             WRITELN(TTY, ' PASSMESSAGES CALLED');
78100             SCHEDULE(PORTSTOBEFIRED);
78150             NEXT := 1;
78200             WHILE (PORTSTOBEFIRED[NEXT] <> 0) AND (NEXT <= HIGHPORT) DO BEGIN
78250                 FIREPORT(PORTSTOBEFIRED[NEXT]);
78300                 INCR(NEXT);
78350             END;
78400     END;
78450
78500     PROCEDURE AXE(PROCI: PROCID);
78550     (* AXE CAN BE CALLED TO MARK A PROCESS AS TERMINATED FOR ANY
78600     ABNORMAL TERMINATION NOT DICTATED BY THE LOGICAL PROCESS CODE *)
78650
78700     VAR PORTJ: PORTID;
78750
78800     BEGIN
78850         IF (INTERTRACE + NETXOTTRACE + NETTRMTRACE + TARGETTRACE) > 0 THEN BEGIN
78900             WRITE(TTY, ' AXE CALLED TO TERMINATE PROCESS ', PROCI);
78950             WRITELN(TTY);
79000             END;
79050         WITH LPS[PROCI] DO BEGIN
79100             ASTATE := TRM;
79150             FOR PORTJ := 1 TO HIGHPORT DO WITH PORTS[PORTJ] DO BEGIN
79200                 (* IF PORTJ NAMES PROCI, SET FALSE THE READY AND ELIGIBLES *)
79250                 PSWAITING := PSWAITING AND (PSENDER <> PROCI);
79300                 PRWAITING := PRWAITING AND (PRECEIVER <> PROCI);
79350                 PSELIGIBLE := PSELIGIBLE AND (PSENDER <> PROCI);
79400                 PRELIGIBLE := PRELIGIBLE AND (PRECEIVER <> PROCI);
79450             END;
79500         END;
79550     END;
79600
79650     PROCEDURE TICK;
79700     (* TICK UPDATES THE GLOBAL CLOCK NETTIME, INCREMENTING IT ONCE
79750     PER CALL. ALSO, CHARGE ALL PROCESSES WITH ONE TIME UNIT
79800

```

```

79850      ACCORDING TO THEIR STATE, PROCESSES WITH ASTATE= CMN OR XQT
79900      ARE COUNTING DOWN THEIR ATIMELEFT FIELDS, IF IT NOW BECOMES
79950      ZERO, ENTER STATES BLK OR TRM, OR XQT RESPECTIVELY, A PROCESS
80000      ENTERING THE BLK STATE HAS ALREADY SPECIFIED THE PORTS OVER
80050      WHICH IT IS ELIGIBLE TO COMMUNICATE VIA PARWAIT CALLS, MARK
80100      THEM AS XWAITING NOW, X=S SENDER OR R RECEIVER, TICK BASICALLY
80150      SEES ALL OF THE PROCESSES THROUGH THE CURRENT TIME UNIT, WHATEVER
80200      THEIR CURRENT STATE, TICK HANDLES THINGS PROCESSES DO      *)
80250
80300      VAR PROCI:PROCID;PORTJ:PORTID;
80350          ALIVE:BOOLEAN;
80400
80450      BEGIN
80500          IF INTERTRACE#0 THEN WRITELN(TTY,' TICK CALLED, NETTIME = ',NETTIME1);
80550          INCR(NETTIME);
80600          FOR PROCI:=1 TO HIGHPROC DO
80650              WITH LPS[PROCI] DO BEGIN
80700                  CASE ASTATE OF
80750                      XQT : BEGIN
80800                          IF (INTERTRACE>20) OR (NETXQTTTRACE>20)
80850                              OR (TARGETTRACE>20) THEN BEGIN
80900                                  WRITE(TTY,' PROCESS ',PROCI13,' IS EXECUTING ');
80950                                  WRITELN(TTY,' TIMELEFT IS ',ATIMELEFT13);
81000                                  END;
81050                                  DECR(ATIMELEFT);
81100                                  INCR(ASUMXQTTIME);
81150                                  INCR(NETSUMXQTTIME);
81200                                  IF ATIMELEFT<=0 THEN BEGIN
81250                                      (* POSSIBLE NEXT STATES ARE TRM AND BLK      *)
81300                                      ASTATE:=ANEXTSTATE;
81350                                      IF ASTATE=TRM THEN BEGIN
81400                                          ATRMTIME:=NETTIME;
81450                                          IF (INTERTRACE>10) OR (NETXQTTTRACE>10)
81500                                              OR (NETTRMTRACE>10) OR (TARGETTRACE>10)
81550                                              THEN BEGIN
81600                                                  WRITE(TTY,' PROCESS ',PROCI13,' GOES ');
81650                                                  WRITELN(TTY,' FROM XQTING TO TRMED');
81700                                                  END;
81750                                          END ELSE      (* ASTATE = BLK      *)
81800                                              IF (INTERTRACE>10) OR (NETXQTTTRACE>10)
81850                                                  OR (NETBLKTRACE>10)
81900                                                  OR (TARGETTRACE>10) THEN BEGIN
81950                                                      WRITE(TTY,' PROCESS ',PROCI13,' GOES');
82000                                                      WRITELN(TTY,' FROM XQTING TO BLKED');
82050                                                      END;
82100                                                  END; (* IF ATIMELEFT<=0 *)
82150                                  END; (* PROCI XQTING      *)
82200                      CMN : BEGIN
82250                          IF (INTERTRACE>20) OR (NETCMNTRACE>20)
82300                              OR (TARGETTRACE>20) THEN BEGIN
82350                                  WRITE(TTY,' PROCESS ',PROCI13,' IS COMMUNICA');
82400                                  WRITELN(TTY,' TING, TIMELEFT IS ',ATIMELEFT13);
82450                                  END;
82500                                  INCR(NETSUMCMNTIME);
82550                                  DECR(ATIMELEFT);
82600                                  INCR(ASUMCMNTIME);
82650                                  IF (ATIMELEFT=0) THEN RESUME(PROCI);
82700                                  END;
82750                      BLK : BEGIN

```

```

82000          IF (INTERTRACE>20) OR (NETBLKTRACE>20)
82050          OR (TARGETTRACE>20) THEN BEGIN
82400          WRITE(TTY,' PROCESS ',PROCI:3,' BLOCKED');
82950          WRITELN(TTY);
83000          END;
83050          INCR(NETSUMBLKTIME);
83100          INCR(ASUMBLKTIME);
83150          END;
83200          TRM ;;
83250          END;
83300          END; (* WITH *)
83350          FOR PROCI:=1 TO HIGHPROC DO
83400          IF (LPS[PROCI],ASTATE = BLK) THEN BEGIN
83450          (* ALL BLOCKED PROCESSES MUST BE WAITING FOR AT LEAST ONE
83500          PROCESS THAT HAS NOT TRMED *)
83550          ALIVE:=FALSE; (* IF THERE IS NONE, THIS PROCESS IS AXED *)
83600          FOR PORTJ:=1 TO HIGHPORT DO WITH PORTS[PORTJ] DO BEGIN
83650          (* IF PORTJ NAMES PROCI THEN MARK THE APPR, XWAITING TRUE *)
83700          PSWAITING:=PSWAITING
83750          OR (PSENDER=PROCI);
83800          PRWAITING:=PRWAITING
83850          OR (PRECEIVER=PROCI);
83900          (* IF PROCI IS WAITING A LIVE PARTNER ON THIS PORT, THEN *)
83950          ALIVE:=ALIVE OR (* MARK IT AS LIVING *)
84000          ((PRECEIVER=PROCI) AND (LPS[PSENDER],ASTATE<>TRM));
84050          ALIVE:=ALIVE OR
84100          ((PSENDER=PROCI) AND (LPS[PRECEIVER],ASTATE<>TRM));
84150          END;
84200          IF NOT ALIVE THEN BEGIN
84250          AXE(PROCI);
84300          IF (INTERTRACE+NETBLKTRACE+NETTRMTRACE+TARGETTRACE) > 0 THEN
84350          WRITELN(TTY,' BECAUSE ALL AWAITED PROCESSES ARE TRMED');
84400          END;
84450          END; (* IF BLKED *)
84500          END; (* TICK *)
84550
84600          PROCEDURE COUNTBYSTATE(VAR RESULT:TABLEBYSTATE);
84650          (* COUNT THE NUMBER OF PROCESSES IN EACH STATE *)
84700          VAR PROCI:PROCID;
84750          BEGIN
84800          IF INTERTRACE<>0 THEN WRITELN(TTY,' COUNTBYSTATE CALLED');
84850          WITH RESULT DO BEGIN
84900          XQTING:=0;
84950          CMNING:=0;
85000          BLKED :=0;
85050          TRMED :=0;
85100          FOR PROCI:=1 TO HIGHPROC DO
85150          CASE LPS[PROCI],ASTATE OF
85200          XQT:INCR(XQTING);
85250          BLK:INCR(BLKED);
85300          CMN:INCR(CMNING);
85350          TRM:INCR(TRMED);
85400          END;
85450          IF INTERTRACE>10 THEN BEGIN
85500          WRITE(TTY,' COUNTS ARE XQTING = ',XQTING:3,' CMNING = ');
85550          WRITELN(TTY,CMNING:3,' BLKED = ',BLKED:3,' TRMED = ',TRMED:3);
85600          END;
85650          END; (* WITH *)
85700          END;

```

```

05750
05800 PROCEDURE PRINTSTATISTICS)
05850 (* CAN BE EXTENDED AS THE NEED ARISES *)
05900 VAR PROCID:PROCID;PORTJ:PORTID;SUM:INTEGER;FACIOR:REAL;
05950 BEGIN
06000   WRITELN(TTY,' NETWORK PERFORMANCE STATISTICS FOLLOW!');
06050   WRITE(TTY,' ELAPSED TIME = ',NETTIME:17,' TOTAL EXECUTION TIME = ');
06100   WRITELN(TTY,NETSUMXQTTIME:18,' FOR A MULTIPROCESSING FACTOR ');
06150   FACTOR:=NETSUMXQTTIME/NETTIME;
06200   WRITE(TTY,' OF ',FACTOR,' TOTAL MESSAGE COUNT = ');
06250   WRITELN(TTY,NETMSGCOUNT:17,' TOTAL COMMUNICATION TIME =1');
06300   WRITE(TTY,NETSUMCMNTIME:16);
06350   WRITELN(TTY,' NETSUMBLKTIME=',NETSUMBLKTIME:16,' PORT SUMMARIES: ');
06400   WRITELN(TTY,' PORTID SENDER RECEIVER STIME RTIME MSG-COUNT');
06450   FOR PORTJ:=1 TO HIGHPORT DO WITH PORTS(PORTJ) DO BEGIN
06500     WRITE(TTY,' ',PORTJ:14,' ',PSENDER:14,' ',PRECEIVER:14);
06550     WRITELN(TTY,' ',PSTIME:14,' ',PRTIME:14,' ',PMSGCOUNT:15);
06600   END;
06650   IF (INTERTRACE >10) OR (TARGETTRACE >10) THEN BEGIN
06700     WRITELN(TTY,' PROCESS SUMMARY');
06750     FOR PROCID:=1 TO HIGHPROC DO SHOWPROCESS(PROCID);
06800     WRITELN(TTY);
06850   END;
06900   WRITE(TTY,' THERE WERE ',DEADLOCKCOUNT:5,' DEADLOCKS IN');
06950   WRITELN(TTY,' THE SEQUENTIAL SIMULATION ');
07000 END;
07050
07100 PROCEDURE INITPORTS)
07150 (* TARDEP * INITIALIZE THE PORT TABLES *)
07200 VAR I:PORTID;
07250 BEGIN
07300   IF INTERTRACE <> 0 THEN WRITELN(TTY,' INITPORTS CALLED');
07350   FOR I:=1 TO PORTMAX DO
07400     WITH PORTS[I] DO BEGIN
07450       PSELIGIBLE:=FALSE;
07500       PSWAITING:=FALSE;
07550       PRWAITING:=FALSE;
07600       PRELIGIBLE:=FALSE;
07650       PSMPC:=1;
07700       PRMPC:=1;
07750       PMSGCOUNT:=0;
07800     END;
07850   WRITELN(TTY,' ENTER THE TOTAL NUMBER OF PORTS USED IN THIS NET');
07900   BREAK;
07950   RESET(TTY);
08000   READ(TTY,HIGHPORT);
08050   WRITELN(TTY,' ECHO ',HIGHPORT:14);
08100 END;
08150
08200 PROCEDURE INITPROCS)
08250 (* TARDEP * INITIALIZE PROCEDURE TABLES *)
08300 VAR I,INDEX,IDXNEXT,KINDMAX:0..PROCMAX;KIND:PROCKIND;
08350 BEGIN
08400   IF INTERTRACE <> 0 THEN WRITELN(TTY,' INITPROCS CALLED');
08450   FOR I:=1 TO PROCMAX DO WITH LPS[I] DO BEGIN
08500     ASTATE:=XQT;
08550     ASUMXQTTIME:=0;
08600     ASUMBLKTIME:=0;
08650     ASUMCMNTIME:=0;

```

```

88700      ATNMTIME:=0)
88750      END)
88800      IDNEXT:=1)
88850      FOR KIND:= SOURCE TO QUEUE2E DO BEGIN
88900          WRITE(TTY,' HOW MANY PROCESSES OF TYPE ');
88950          PRPROCKIND(KIND);
89000          WRITELN(TTY,' ');
89050          BREAK)
89100      RESET(TTY);
89150      READ(TTY,KINDMAX);
89200      IF (KINDMAX>0) THEN
89250          FOR I:=1 TO KINDMAX DO BEGIN
89300              DWN[IDNEXT],OPROCKIND:=KIND)
89350              LPS[IDNEXT],ATYPE:=KIND;
89400              LPS[IDNEXT],AINSTANCE:=I)
89450              LPS[IDNEXT],AMPCI:=0)
89500              (* THAT ASSIGNMENT DISTINGUISHED THIS CALL AS A CREATE *)
89550              RESUME(IDNEXT);
89600              INCR(IDNEXT);
89650              END)
89700      END)
89750      HIGHPROC:=IDNEXT-1)
89800      END)
89850
89900      PROCEDURE BREAKDEADLOCK)
89950      (* ATTEMPT TO BREAK DEADLOCK BY LETTING ALL PROCESSES COMPUTE W SUB IJ
90000      BY RESUMING THEM AT AMPC = 1010. EACH CALL TO PROCI LETS PROCI
90050      REVISE DOWNWARD OR LEAVE CONSTANT ITS ESTIMATE OF THE EARLIEST
90100      PROCI COULD TRY TO SEND ON THE ARC TO PROCJ. WHEN PROCI IS
90150      SO RESUMED THE K*TH TIME, IT COMPUTES THE K*TH ESTIMATE OF
90200      W SUB IJ FOR ALL OUTPUT PORTS TO PROCJ BASED ON THE AVAILABILITY
90250      OF THE (K-1)TH W SUB HI ON ALL INPUT PORTS FROM PROCH, THIS
90300      HAVING BEEN COMPUTED ON THE PREVIOUS WIJPASS,
90350      PORTS HAVE A .PHIJ FIELD WHICH IS WRITTEN BY THE SENDER, AND
90400      READ BY THE RECEIVER, OF MESSAGES ON THE BOUND PORT. *)
90450
90500      VAR WIJPASS,PROCI:PKUCID;PORTP:PORTID)
90550      BEGIN
90600      IF (INTERTRACE+NETDEADTRACE+NETCMNTRACE+NETXQTTTRACE+NETBLKTRACE>0)
90650      THEN BEGIN
90700          WRITELN(TTY,' BREAKDEADLOCK ATTEMPTED AT NETTIME = ',NETTIME);
90750          END)
90800      FOR PROCI:= 1 TO HIGHPROC DO LPS(PROCI),AMPCI:=1010)
90850      FOR PORTP:= 1 TO HIGHPORT DO WITH PORTS(PORTP) DO BEGIN
90900          PHIJ:=MAXINT)
90950          END)
91000      FOR WIJPASS:= 1 TO HIGHPROC DO BEGIN
91050          IF (NETDEADTRACE>10) THEN BEGIN
91100              WRITELN(TTY,' COMPUTE W (',WIJPASS:',') FOR ALL PORTS');
91150              END)
91200          FOR PROCI:= 1 TO HIGHPROC DO BEGIN
91250              RESUME(PROCI);
91300              END)
91350          IF (NETDEADTRACE>20) THEN BEGIN
91400              WRITE(TTY,' HERE ARE W (K=',WIJPASS:',') SUB IJ FOR ');
91450              WRITELN(TTY,' ALL PORTS IJ ');
91500              WRITE(TTY,'          I          J          W(K) SUB IJ ');
91550              WRITELN(TTY);
91600

```

```

91650         FOR PORTP:=1 TO HIGHPORT DO WITH PORTS[PORTP] DO BEGIN
91700             WRITE(TTY,'          ',PSENDER13,'          ',PRECEIVER13);
91750             WRITELN(TTY,'          ',PW1J10,'          ');
91800             END;
91850         END;
91900     END;
91950     FOR PROC1:=1 TO HIGHPROC DO BEGIN
92000         LPS[PROC1],AMPCI=1020;
92050         RESUME(PROC1);
92100     END;
92150 END;

92200
92250 (*          MAIN          PROGRAM          *)
92300
92350 BEGIN
92400     WRITELN(TTY,' BEGIN PROGRAM DSIM');
92450     INITPROCS I          (* INIT TARGET NETWORK PROCESS RECORDS *)
92475     INITPORTS          (* INIT TARGET NETWORK PORT RECORDS *)
92500     WRITELN(TTY,' END OF NETWORK SPECIFICATION, INSPECT AND VERIFY ');
92550     SHOWNETWORK;          (* DISPLAY NETWORK FOR USER *)
92600     BUFFERSIZE1=1)      (* BUFFERSIZES WILL BE 1,2,4,8, & 16 *)
92650     REPEAT              (* RUN THE PROGRAM WITH A NEW BUFFERSIZE *)
92700     INITIALIZE;          (* INITIALIZE INTERPRETER VARS FOR FRESH RUN *)
92750     INITPORTS;          (* INITIALIZE THE PORT RECORDS *)
92800     SETTRACE;          (* INITIALIZE THE TRACE AND LIMIT VARIABLES *)
92850     WRITELN(TTY);
92900     WRITE(TTY,' BEGIN THE SIMULATION, NETTIME = ',NETTIME10);
92925     WRITELN(TTY,' BUFFERSIZE = ',BUFFERSIZE12);
92937     BREAK;
92950     FOR PROC1:=1 TO HIGHPROC DO BEGIN (* ACTIVATE PROCESS I *)
93000         LPS[PROC1],AMPCI= 1;
93050             (* THAT ASSIGNMENT DISTINGUISH THIS CALL TO
93100             THE PROCESSES AS THEIR INITIAL ACTIVATION *)
93150         RESUME(PROC1);
93200             (* FORCING THE QUEUE20 BUFFERSIZES FOR EXPT. *)
93250         IF LPS[PROC1],ATYPE=QUEUE20 THEN OWN[PROC1],Q20QMAX1=BUFFERSIZE-1;
93300     END;          (* ALL PROCESSES ARE EXECUTING *)
93350     IF (INTERTRACE>50) THEN DUMPTY;
93400     REPEAT          (* SIMULATE PASSAGE OF ONE TIME UNIT FOR NETWORK *)
93450     TICK;          (* ALL LPS DO THEIR THING FOR ONE TIME UNIT *)
93500     PASSMESSAGES; (* FIRE SOME READY PORTS IF POSSIBLE *)
93550     COUNTBYSTATE(COUNT); (* COUNT THE SURVIVORS, AND OTHERS *)
93600     IF (COUNT,XOTING=0) AND (COUNT,CMNING=0) THEN (* NO ONE ALIVE *)
93650         IF (COUNT,TRMED>=HIGHPROC) THEN NETTERMI=TRUE (* ALL DEAD *)
93700         ELSE BEGIN (* NEITHER DEAD NOR ALIVE-- DEADLOCKED! *)
93750             IF (NETDEADTRACE>0) THEN BEGIN
93800                 WRITELN(TTY,' NET DEADLOCKED, NETTIME = ',NETTIME);
93850             END;
93900             NETDEADLOCK:=TRUE;
93950             INCR(DEADLOCKCOUNT);
94000             BREAKDEADLOCK;
94050             PASSMESSAGES;
94100             COUNTBYSTATE(COUNT);
94150             NETDEADLOCK:= (COUNT,XOTING = 0) AND (COUNT,CMNING = 0);
94200             IF (INTERTRACE>0) OR (NETDEADTRACE>0)
94300             OR (TARGETTRACE>0) THEN BEGIN
94350                 WRITE(TTY,' PROCEDURE BREAKDEADLOCK WAS ');
94400                 CASE NETDEADLOCK OF
94450                     TRUE:WRITELN(TTY,'UNABLE TO BREAK DEADLOCK');

```

```

94500             FALSE;WRITELN(TTY,'ABLE TO BREAK DEADLOCK');
94550             END;
94600             END;
94650             END;
94700             END;
94750             (* IF TIME OR MESSAGE LIMITS EXCEEDED, GIVE OPERATOR A CHANCE
94800             TO ENTER NEW LIMITS, CHANGE THE TRACE VARIABLES, ETC. *)
94850             IF (NETTIME>NETTIMELIMIT) OR (NETMSGCOUNT>NETMSGLIMIT)
94900             OR (INTERTRACE>30) OR (DEADLOCKCOUNT>DEADLOCKLIMIT) THEN BEGIN
94950                 SETTRACE;
95000                 WRITELN(TTY,' RESUME SIMULATION, WATCH YOUR TIME ');
95050                 BREAK;
95100             END;
95150             UNTIL NETDEADLOCK OR NETTERM OR (NETTIME>NETTIMELIMIT)
95200             OR (NETMSGCOUNT>NETMSGLIMIT) OR (DEADLOCKCOUNT>DEADLOCKLIMIT);
95250             WRITE(TTY,' TARGET SIMULATION TERMINATED BECAUSE');
95300             IF (DEADLOCKCOUNT>DEADLOCKLIMIT) THEN
95350                 WRITELN(TTY,' DEADLOCK LIMIT EXCEEDED ');
95400             ELSE IF NETDEADLOCK THEN
95450                 WRITELN(TTY,' OF UNRESOLVABLE DEADLOCK ');
95500             IF NETTERM THEN WRITELN(TTY,' ALL PROCESSES TERMINATED');
95550             IF (NETTIME>NETTIMELIMIT) THEN WRITELN(TTY,' TIMELIMIT EXCEEDED');
95600             IF NETMSGCOUNT>NETMSGLIMIT THEN WRITELN(TTY,' MSG LIMIT EXCEEDED');
95650             WRITELN(TTY);
95700             (* RESUME PROCESSES ONE LAST TIME TO PRINT THEIR SUMMARIES ETC. *)
95750             FOR PROCI=1 TO HIGHPROC DO BEGIN
95800                 LPS(PROCI),AMPC=1000;
95850                 RESUME(PROCI);      (* BY CONVENTION, THE LAST CALL TO PROCI *)
95900             END;
95950             PRINTSTATISTICS;
95960             BUFFERSIZE:=BUFFERSIZE+BUFFERSIZE;
95970             UNTIL (BUFFERSIZE>20);
96050             WRITELN(TTY,' END PROGRAM DSIM');
96100             END.

```




## APPENDIX 2

### SOME TEST RESULTS

This appendix reports the results of simulation runs on two different networks where the sizes of all the buffers ( = queue sizes in processes of type QUEUE20) were successively set to 1, 2, 4, 8, and 16.

The connectivity graph of network 1 is shown in fig A-1. All SOURCE type nodes emitted jobs every 10 time units. The server, or DELAY, nodes had a service time with an exponential distribution, and a mean service time of 7.0 time units. All FORK nodes had a probability of 0.4 of sending a received job out to its respective SINK node.

The relative frequency of job arrivals from SOURCE nodes and departures to SINK nodes was such that the network tended to fill up with jobs over a period of time, and finally encounter a deadlock where all nodes except the SINK nodes were waiting to send out a job. The times at which this occurred depended on the buffer sizes. The following table summarizes the behavior of this network. Elapsed time is the total time from the beginning of the simulation to the final deadlock. MPF represents the "multiprocessing factor", taken as the ratio of the total amount of execution time for all



NETWORK 1

fig A-1

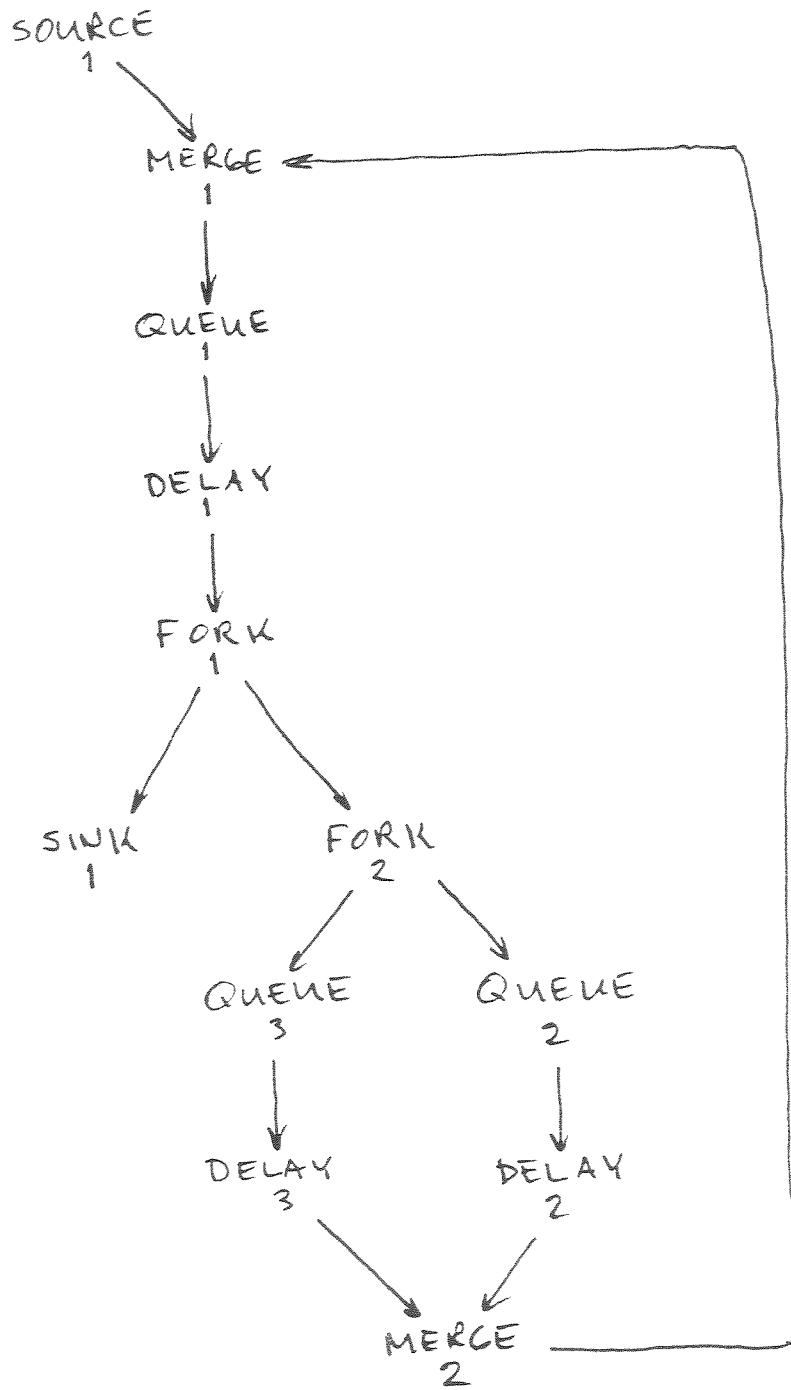
nodes in the network, to the elapsed time for the network.

Buffer Size	Elapsed Time	Execution Time	Deadlock Count	MPF
1	630	2876	2	4.57
2	439	2146	2	4.88
4	891	4794	4	5.38
8	1864	12300	2	6.60
16	1580	10092	2	6.38

In all cases, at least one deadlock was detected and recovered from. With identical job arrivals, service times, and departure times, one would expect the MPF and total execution times to be larger for networks with larger buffer sizes. As the results show, this was not always the case in this run, although the trend is there. This may have been caused in part by short-term fluctuations in the RANDOM function on the DEC-10. This is possible because these results were obtained in a single run of the program, without ever resetting the "seed" of the pseudo-random number generator.

The graph of the second network tested is shown in fig A-2. Here, the SOURCE emitted a job every 20 time units. The node labelled DELAY-1 had a constant service time of 2.0 time units. The other DELAY nodes had exponential service times with mean value = 2.0. Jobs entering FORK 1 were sent to the SINK with probability 0.2. Jobs entering FORK 2 were equally likely to go to QUEUE 2 or QUEUE 3.

Again, the buffer sizes were set to 1, 2, 4, 8, and



NETWORK 2

fig A-2

16, and statistics collected for each size. This network did not fill up with jobs like network 1. Each run continued until the elapsed time was 1001. All of the deadlocks encountered were those arising as a result of the waiting rules for the processes. The following table summarizes the behavior of this network.

Buffer Size	Execution Time	Number of Deadlocks
1	1758	29
2	1724	27
4	1677	24
8	1773	23
16	1702	21

In this network, the larger buffers always resulted in fewer deadlocks, although the net-wide sum of execution time seems uncorrelated.

## REFERENCES

- [CHA79] Chandy, K.M., J. Misra, "Distributed Simulation: A Case Study in Design and Verification of Distributed Programs," IEEE Transactions on Software Engineering (September, 1979), pp. 440-452.
- [CHA81] Chandy, K.M., J. Misra, "Asynchronous Distributed Simulation via a Sequence of Parallel Computations," CACM (April, 1981), pp. 198-206.
- [HOA78] Hoare, C.A.R., "Communicating Sequential Processes," CACM (August, 1978), pp. 666-677.
- [LAM78] Lamport, Leslie, "Time, Clocks, and the Ordering of Events in a Distributed System," CACM (July, 1978), pp. 558-565
- [SEE79] Seethalakshmi, M., "A Study and Analysis of Performance of Distributed Simulation," M.S. Report, 1979, Computer Science Department, University of Texas, Austin, Texas.
- [DIJ77] Dijkstra, Edsger W., "Guarded Commands, Nondeterminacy and Formal Derivation of Programs," Current Trends in Programming Methodology, Vol I. Software Specification and Design, (Yeh, Raymond T., editor) Prentice Hall, Englewood Cliffs, N.J., 1977, pp. 233-242.

## VITA

William Francis Quinlivan III was born in Portsmouth, Virginia, on February 4, 1950, the son of Faye Elizabeth Quinlivan and William Francis Quinlivan, Jr. After attending Central Catholic High School, he graduated from Melbourne High School in 1967. He attended The University of Florida and subsequently served in the U.S. Army. He attended Austin Community College for a year in 1974, and in 1975 he entered The University of Texas. In June, 1978, he received a Bachelor of Science degree in Mathematics, and a Bachelor of Arts in Computer Science. During the following years, he was employed as a Systems Programmer at Tracor, Inc. In June, 1979, he entered the Graduate School of The University of Texas. He is currently employed as a Systems Analyst by Information Research Associates in Austin, Texas.

Permanent address: 609 Orland Blvd.  
Austin, Texas. 78745

This report was typed by William F. Quinlivan.