



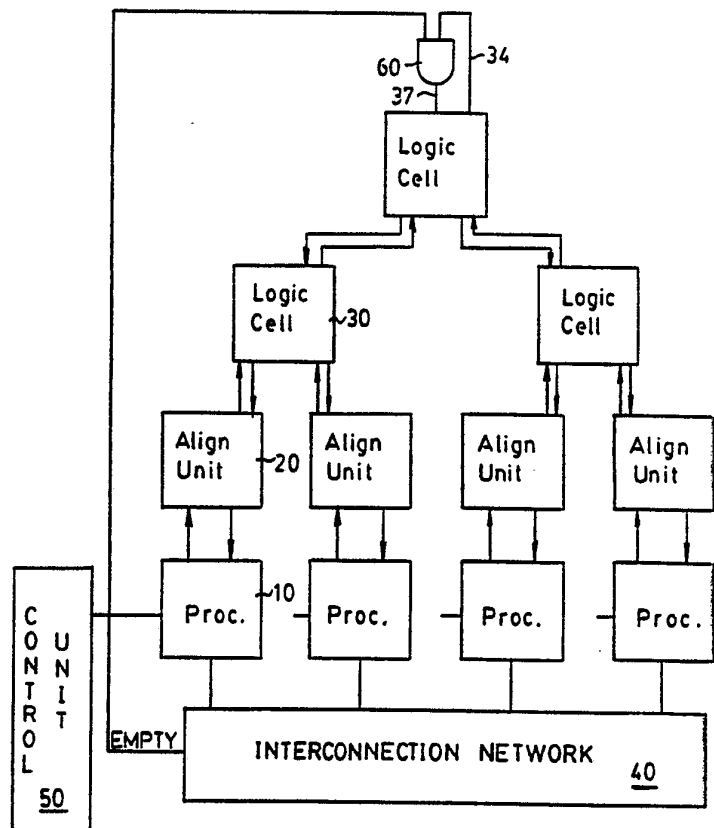
INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

<p>(51) International Patent Classification <sup>4</sup> : <b>G06F 15/16</b></p>	<p>A1</p>	<p>(11) International Publication Number: <b>WO 89/ 07299</b> (43) International Publication Date: 10 August 1989 (10.08.89)</p>
<p>(21) International Application Number: PCT/US89/00395 (22) International Filing Date: 1 February 1989 (01.02.89) (31) Priority Application Number: 151,386 (32) Priority Date: 2 February 1988 (02.02.88) (33) Priority Country: US  (71) Applicant: THINKING MACHINES CORPORATION [US/US]; 245 First Street, Cambridge, MA 02142 (US). (72) Inventor: HILLIS, W., Daniel ; 135 St. Paul Street, Brookline, MA 02146 (US). (74) Agent: MORRIS, Francis, E.; Pennie &amp; Edmonds, 1155 Avenue of the Americas, New York, NY 10036 (US).</p>		<p>(81) Designated States: AT (European patent), BE (European patent), CH (European patent), DE (European patent), FR (European patent), GB (European patent), IT (European patent), JP, LU (European patent), NL (European patent), SE (European patent).  <b>Published</b> <i>With international search report.</i></p>

(54) Title: METHOD AND APPARATUS FOR ALIGNING THE OPERATION OF A PLURALITY OF PROCESSORS

(57) Abstract

A method and apparatus are disclosed for aligning a plurality of multi-processors (10). The apparatus preferably comprises an alignment unit (20) associated with each processor and a logic network (30) for combining the output of the alignment unit and for broadcasting information to these units. Alignment is achieved by inserting in the instruction stream of each processor that is to be aligned a request for alignment, by testing for prior completion of any instructions that must be completed and by causing all processors to wait until they have all made the request for alignment and completed necessary prior instructions. The alignment unit associated with each processor monitors the instruction stream to detect a request for alignment. When all the units have made such a request, the logic network informs the alignment units; and the alignment units inform the processors.



**FOR THE PURPOSES OF INFORMATION ONLY**

Codes used to identify States party to the PCT on the front pages of pamphlets publishing international applications under the PCT.

<b>AT</b> Austria	<b>FR</b> France	<b>ML</b> Mali
<b>AU</b> Australia	<b>GA</b> Gabon	<b>MR</b> Mauritania
<b>BB</b> Barbados	<b>GB</b> United Kingdom	<b>MW</b> Malawi
<b>BE</b> Belgium	<b>HU</b> Hungary	<b>NL</b> Netherlands
<b>BG</b> Bulgaria	<b>IT</b> Italy	<b>NO</b> Norway
<b>BJ</b> Benin	<b>JP</b> Japan	<b>RO</b> Romania
<b>BR</b> Brazil	<b>KP</b> Democratic People's Republic of Korea	<b>SD</b> Sudan
<b>CF</b> Central African Republic	<b>KR</b> Republic of Korea	<b>SE</b> Sweden
<b>CG</b> Congo	<b>LI</b> Liechtenstein	<b>SN</b> Senegal
<b>CH</b> Switzerland	<b>LK</b> Sri Lanka	<b>SU</b> Soviet Union
<b>CM</b> Cameroon	<b>LU</b> Luxembourg	<b>TD</b> Chad
<b>DE</b> Germany, Federal Republic of	<b>MC</b> Monaco	<b>TG</b> Togo
<b>DK</b> Denmark	<b>MG</b> Madagascar	<b>US</b> United States of America
<b>FI</b> Finland		

METHOD AND APPARATUS FOR  
ALIGNING THE OPERATION OF A PLURALITY OF PROCESSORS

5 Background of the Invention

This relates to multi-processors and in particular to a system for aligning the operation of a plurality of processors in a multi-processor.

10 Numerous multi-processors are known in the art. See Wang, A. and Brigg, F.A., Computer Architecture in Parallel Processing, (McGraw Hill, 1984). Illustrative such systems include the Illiac-IV, the Burroughs Scientific Processor (BSP) manufactured by Burroughs Corporation (now  
15 Unisys), the Massively Parallel Processor (MPP) developed at the NASA Goddard Space Flight Center and manufactured by Goodyear Aerospace, the DAP manufactured by ICL of England, the STARAN designed by Goodyear Aerospace, the IBM 370/168 MP and IBM 3081/3084 both manufactured by International Business  
20 Machines, the Univac 1100/80 manufactured by Sperry Univac (now Unisys), Tandem-16 Nonstop, D-825 manufactured by Burroughs Corporation (now Unisys), the HEP manufactured by Denelcor Inc., the Cray X-MP and Cray-2 each manufactured by Cray Research Inc., and the Cm\* and C.mmp, each developed at  
25 Carnegie Mellon University.

Recently, multi-processors have been introduced in which thousands of processors are operated in parallel. One such processor is the Connection Machine Computer being  
manufactured and sold by the assignee of the present application and described more fully in U.S.  
30 Patent 4,598,400, which is incorporated herein by reference. The Connection Machine Computer comprises a central computer, a microcontroller, and an array of as many as 65,536 parallel processors in presently available embodiments. The central  
35 computer may be a suitably programmed commercially available

computer such as a Symbolics 3600-series LISP Machine. The microcontroller is an instruction sequencer of conventional design for generating a sequence of instructions that are applied to the array of parallel processors by means of a  
5 thirty-two bit parallel bus. The microcontroller receives from the parallel processing array a signal which is a general purpose or GLOBAL signal that can be used for data output and status information.

Common to all multi-processors are problems of  
10 synchronizing the operations of the multi-processor so as to make it possible for individual multi-processors to exchange information. These problems can be avoided by operating all the processors of the multi-processor in a single instruction stream, multiple data stream (SIMD) environment in which all  
15 processing units receive the same instruction stream in parallel although they operate on different data. It is advantageous, however, to be able to operate the processors of a multi-processor independently to the extent that each processor operates on its own instruction stream in a  
20 multiple instruction stream, multiple data stream (MIMD) environment.

Synchronization may be effected by having each processor set a bit in a register when the processor reaches a point in its program where it is to be aligned with other  
25 processors. Upon reaching this synchronization point, each processor tests the register to determine if all bits have been properly set and no processor goes forward in its program until all such bits have been set. As a result of  
30 this technique, processors may be brought into strict time synchronization such that at each synchronization point each processor performs a predetermined operation at the same time.

### Summary of the Invention

In many operations, it is not necessary to achieve strict time synchronization in the operations of a plurality of processors. Rather it is sufficient to ensure simply that each processor has performed a specified task before any processor can initiate another specified task. This less strict form of synchronization is referred to herein as alignment. In accordance with the invention, a method and apparatus have been devised for aligning a plurality of multi-processors. The apparatus preferably comprises an alignment unit associated with each processor and a logic network for combining the output of the alignment units and for broadcasting information to these units. Alignment is achieved by inserting in the instruction stream from each processor that is to be aligned a request for alignment and by causing all such processors to wait until they have all made the request for alignment and completed any necessary operations.

The alignment unit associated with each processor monitors the instruction stream to detect a request for alignment. The logic network illustratively is an array of AND gates that tests each alignment unit to determine if it has detected a request for alignment and tests the processors and/or their interconnection network to determine if all necessary operations have been completed. When all the units have made such a request and all necessary operations have been completed, the logic network informs the alignment units; and the alignment units inform the processors.

### Brief Description of the Drawings

These and other objects, features and advantages of the invention will become more readily apparent from the following detailed description of the invention in which:

Fig. 1 is a block diagram illustrating an illustrative embodiment of a multi-processor which incorporates the invention;

Fig. 2 is a block diagram illustrating an illustrative embodiment of the alignment unit of the present invention;

Fig. 3 is an illustrative embodiment of a unit of combinatorial logic of the present invention;

Fig. 4 is a state diagram illustrating the operation of the alignment unit of Fig. 2; and

Fig. 5 is a block diagram of a second illustrative embodiment of the invention.

#### Detailed Description of the Invention

As shown in Fig. 1, illustrative apparatus of the present invention comprises an array of processors 10, an array of alignment units 20, an array of combinatorial logic cells 30, an interconnection network 40 and a control unit 50. Illustratively, each processor includes an arithmetic/logic unit, a read/write memory, and input and output means. Numerous such processors are well known in the art. Illustratively, each processor is an 80386-type microprocessor such as that made by Intel Corporation.

Interconnection network 40 provides for communication between the processors. In particular, it permits the exchange of data and control information between one processor and another. Numerous interconnection networks are known in the art. See for example Wu, C. and Feng, T, Tutorial: Interconnection Network for Parallel and Distributed Processing (IEEE, 1984). Illustrative communication networks are arrangements which connect each processor to its nearest neighbor in an orthogonal array of

two or more dimensions. Two such networks are described in the above referenced '400 patent and numerous others are known in the art.

5 The processors and interconnection network are typically operated under the control of some type of control unit 50. In the case of the processors and interconnection network of the Connection Machine Computer, an illustrative control unit is a Symbolics 3600 Series LISP machine and a microcontroller. Numerous other control units are used with  
10 other multi-processors.

If there is communication among the processors of a multi-processor, it may be necessary to assure that one processor has reached a certain step in its operation before another processor can perform a specified operation. For  
15 example, if data is to be fetched from a first processor to a second, it is necessary to ensure that the data be in the first processor before the fetch occurs. To ensure that the data is there when needed, it is necessary to align the individual processors. This can be achieved by operating all  
20 the processors strictly in parallel in a SIMD environment. However, in many cases it may be advantageous to operate the processors in a MIMD environment and align only where needed.

In accordance with the invention, alignment is achieved among a group of processors by inserting in the  
25 instruction stream of each of these processors two additional instructions: a request for alignment and a test if alignment has been achieved. Typically, the request for alignment instruction is issued at or after the time the instruction stream issues an instruction (such as a  
30 communication instruction) which will require the processors to be in alignment at some later time; and the test instruction is issued immediately before the point in the instruction stream where alignment must have been achieved.

For example, let us assume that data is to be transferred from each processor in a group of processors to its neighboring processor and combined with data at that processor. If we identify a pair of neighboring processors in the group of processors as processor  $i-1$  and processor  $i$ , the following instruction stream may be provided to each processor in the group to combine such data at each processor:

- 01 Send Data from Processor  $i-1$  to Register A of Processor  $i$
- 02 Move Data from Register B to Register C of Processor  $i$
- 03 Add Contents of Register A to Contents of Register C of Processor  $i$

It is clearly necessary that the proper data be in Register A before it is added to the contents of Register C. To ensure this, it is necessary to align the operations of the processors.

In accordance with the present invention, such alignment is accomplished by inserting request and test instructions in the instruction stream so that for the above example the instruction stream is as follows:

- 01 Send Data from Processor  $i-1$  to Register A of Processor  $i$
- 02 Request alignment
- 03 Move Data from Register B to Register C of Processor  $i$
- 04 Test if aligned
- 05 If false, reexecute Test
- 06 If true, Add Contents of Register A to Contents of Register C of Processor  $i$



In accordance with the invention, the request and test instructions are inserted in the instruction stream from each processor to be aligned and all the processors wait at the test instruction until all the processors have asserted the request instruction.

5 Preferred apparatus for processing such request and test instructions comprises alignment units 20 and combinatorial logic cells 30 of the apparatus of Fig. 1. As shown in Fig. 1, each processor 10 is connected to an alignment unit so as to provide signals to such unit and to receive signals from it. As depicted in Fig. 2, each alignment unit 20 comprises four D-type flip flops 21, 22, 10 23, 24 and combinatorial logic 26 which implements the input/output table set forth in Table I. As shown in Fig. 3, each logic cell 30 comprises an AND gate 33 and optionally 15 two D-type flip flops 35, 36. The logic cells are connected in the form of a binary tree with each logic cell in the lowest level of the tree being connected to a pair of alignment units 20.

20

25

30

35

TABLE I

INPUTS		OUTPUTS				
State	INST	GRANT	State+1	ALIGN	ASSERT	
	IDLE	TEST	LOW	ERROR	X	X
5	IDLE	TEST	HIGH	ERROR	X	X
	IDLE	REQ	LOW	PEND	LOW	LOW
	IDLE	REQ	HIGH	ERROR	X	X
	IDLE	OTHER	LOW	IDLE	LOW	LOW
	IDLE	OTHER	HIGH	ERROR	X	X
	PEND	TEST	LOW	PEND	LOW	HIGH
	PEND	TEST	HIGH	BLOCK	HIGH	HIGH
	PEND	REQ	LOW	ERROR	X	X
10	PEND	REQ	HIGH	ERROR	X	X
	PEND	OTHER	LOW	PEND	LOW	HIGH
	PEND	OTHER	HIGH	GOT	HIGH	HIGH
	GOT	TEST	LOW	IDLE	HIGH	LOW
	GOT	TEST	HIGH	BLOCK	HIGH	LOW
	GOT	REQ	LOW	ERROR	X	X
	GOT	REQ	HIGH	ERROR	X	X
15	GOT	OTHER	LOW	SAVE	HIGH	LOW
	GOT	OTHER	HIGH	GOT	HIGH	LOW
	SAVE	TEST	LOW	IDLE	HIGH	LOW
	SAVE	TEST	HIGH	ERROR	X	X
	SAVE	REQ	LOW	ERROR	X	X
	SAVE	REQ	HIGH	ERROR	X	X
	SAVE	OTHER	LOW	SAVE	HIGH	LOW
	SAVE	OTHER	HIGH	ERROR	X	X
20	BLOCK	TEST	LOW	ERROR	X	X
	BLOCK	TEST	HIGH	ERROR	X	X
	BLOCK	REQ	LOW	PEND	LOW	LOW
	BLOCK	REQ	HIGH	QUEUE	LOW	LOW
	BLOCK	OTHER	LOW	IDLE	LOW	LOW
	BLOCK	OTHER	HIGH	BLOCK	LOW	LOW
	QUEUE	TEST	LOW	PEND	LOW	LOW
	QUEUE	TEST	HIGH	QUEUE	LOW	LOW
25	QUEUE	REQ	LOW	ERROR	X	X
	QUEUE	REQ	HIGH	ERROR	X	X
	QUEUE	OTHER	LOW	PEND	LOW	LOW
	QUEUE	OTHER	HIGH	QUEUE	LOW	LOW

Each alignment unit receives from the processor to which it is connected the instruction stream INST being executed by the processor. In accordance with the input/output table of Table I as described more fully below, each alignment unit provides to its processor a signal ALIGN indicating whether or not alignment has been achieved. In addition, each alignment unit also provides an output signal

ASSERT to an input of the logic cell to which it is connected and it receives from that logic cell a READY signal which is applied to flip-flop 21 and then to combinatorial logic 26.

AND gate 33 of logic cell 30 receives ASSERT  
5 signals on input lines 31, 32 and produces an output ASSERT signal on output line 34 (optionally, via flip-flop 36). In the specific embodiment of the invention shown in Fig. 3, a logic high signal is not asserted on an output line 34 until logic high signals are asserted on both input signal lines  
10 31, 32 to AND gate 33. The output ASSERT signals from two cells are applied as the two inputs to a logic cell 30 in the next higher level of the binary tree.

The uppermost cell in the AND tree has an output line 34 which is input to an AND gate 60 along with a signal  
15 from interconnection network 10 which indicates that the interconnection network is empty. The network is empty whenever there are no messages being routed through the network. A test for this condition is conventional in some interconnection networks and the provision of a signal  
20 indicating that the network is empty is within the ordinary skill of the art. In the preferred embodiment shown in Fig. 1, a test for an empty network is made to ensure that message traffic on the interconnection network is completed.

Upon coincidence of a logic high empty signal and a  
25 logic high signal on the assert output 34 of the uppermost cell, a ready signal is applied by AND gate 60 to an input ready line 37 to the uppermost cell of the binary tree. This signal is provided to two output ready lines 38, 39  
(optionally, via flip-flop 35) which are connected  
30 respectively to the input ready lines of two cells in the next lower level of the binary tree. At the lowest level of the binary tree, the READY signal is provided by each cell to two alignment units 20 so that the READY signal is broadcast to the entire array of alignment units 20.

In general, the operation of alignment units 20 and logic cells 30 is such that a logic high signal is output on the ASSERT line from an alignment unit when that unit detects a REQUEST for alignment signal in the instruction stream received at the alignment unit. If and only if each alignment unit detects a REQUEST for alignment signal, output line 34 from the uppermost logic cell 30 in the binary tree will be logic high. If message traffic on processor interconnection network 10 is completed at some point after all the processors issue a REQUEST for alignment, then the empty signal will become logic high as well and the output of AND gate 60 will become logic high. As a result, a logic high READY signal will be broadcast down through logic cells 30 to all the alignment units where the logic high READY signal is recognized as a GRANT signal.

The GRANT signal is processed by each alignment unit so as to output to the processor a single logic high ALIGN instruction for each TEST instruction in the instruction stream from the processor. If the GRANT signal is received before the TEST signal in the instruction stream, each alignment unit simply stores the GRANT signal and waits. If the TEST signal is received first, the alignment unit likewise just waits. Upon receipt of both TEST and GRANT signals, the alignment unit issues a logic high ALIGN signal to the processor, and the processor is then enabled to perform the operation for which alignment was required.

Further details of the operation of the alignment unit are best explained in conjunction with Table I and the state diagram of Fig. 4, which depict the operation of alignment unit 20 and its combinatorial logic 26 and flip-flops 22-24. As shown in Fig. 4, the alignment unit can be in any one of six states: IDLE, PENDING, BLOCKING, QUEUE, GOT and SAVE. The state of the alignment unit is stored in binary encoded form in flip-flops 22, 23, 24 and is provided as an input to combinatorial logic 26 along with the

instruction stream INST from the processor and the GRANT signal from flip-flop 21. The instruction stream is the sequence of instructions that is provided to the processor and, in accordance with the invention includes three types of instructions: REQUEST for alignment, TEST, and all OTHER instructions which are represented in Fig. 4 by the symbols R, T and O respectively. The logic high GRANT signal is represented in Fig. 4 by G and the logic low GRANT signal by  $\bar{G}$ .

The different combinations of inputs to combinatorial logic 26 are set forth in the three left-hand columns of Table I. The outputs produced by combinatorial logic 26 in response to these inputs are set forth in the three right-hand columns of Table I. As indicated, one output specifies the next state of the alignment unit and the other two specify the logic level of the ALIGN and ASSERT lines. In cases where some combination of inputs is not possible, the next state has been indicated to be an ERROR and the ALIGN and ASSERT levels have been left undefined.

The apparatus of Fig. 1 is initialized so that the alignment units are in the IDLE state. Ordinarily, each alignment unit is in the IDLE state and remains so upon receipt of any instruction in the instruction stream from the processor except for a REQUEST instruction. This is indicated by the loop labelled  $O\bar{G}$ . While in the IDLE state both the ALIGN and ASSERT outputs are logic low. The alignment unit changes to the PENDING state upon receipt of a REQUEST signal as detailed below; and in the PENDING state causes the ASSERT output to be logic high. If and when the ASSERT outputs of all the alignment units become logic high and the interconnection network is empty, the GRANT signal will become logic high and causing the GRANT signal to be logic high each alignment unit will producing output signals

as specified in Table I. The alignment unit will ultimately return to the IDLE state if a TEST instruction and a logic low GRANT signal are received.

As indicated, if a REQUEST signal is detected in  
5 the instruction stream and if a logic high GRANT signal has not been received from the logic cells, the alignment unit changes from the IDLE state to the PENDING state as indicated by the line  $\overline{RG}$  in Fig. 4 and the third line of Table I. Once in the PENDING state, the alignment unit will stay there  
10 unless a logic high GRANT signal is received. While in the PENDING state, combinatorial logic 26 will issue a logic high ASSERT signal for any of the possible combination of input signals. The ASSERT signals are applied to the inputs to logic cells 30; and if all the ASSERT signals are logic high  
15 and if the connection network is empty a logic high READY will be broadcast to all the alignment units indicating that all the processors have completed the instruction (such as a communication instruction) they were required to complete and accordingly are aligned. This logic high READY signal is the  
20 logic high GRANT signal of Table I. If a logic high GRANT signal is received along with an OTHER instruction, the unit changes to the GOT state as indicated by the line OG in Fig. 4 and the twelfth line of Table I and issues a logic high ALIGN signal while continuing to issue a logic high  
25 ASSERT signal.

If any OTHER instruction continues to be received and the GRANT signal continues to be logic high, the combinatorial logic remains in the GOT state, but drops the ASSERT signal to logic low which causes the GRANT signal to  
30 become logic low. If the GRANT signal is logic low while an OTHER instruction is received, the alignment unit switches to the SAVE state where it remains until a TEST instruction is received and then returns to the IDLE state.

If, however, a TEST instruction is received while the unit is in the GOT state, the unit changes to the BLOCKING state if the GRANT signal is logic high or to the IDLE state if the GRANT signal is logic low and causes the  
5 ASSERT signal to become logic low.

The BLOCKING state prevents other alignment operations from being performed. Upon receipt of any OTHER instruction at the alignment unit while it is in the BLOCKING state and the GRANT signal is low, the unit returns to the  
10 IDLE state as indicated by the line  $O\bar{G}$  between the BLOCKING and IDLE states. If, however, another REQUEST instruction is received, the alignment unit returns to the PENDING state if the GRANT signal is logic low or to the QUEUE state if the GRANT signal is logic high. The unit will remain in QUEUE as  
15 long as the GRANT signal is high and will return to PENDING when the GRANT signal goes low.

Alternatively, while in the PENDING state, if the alignment unit receives the TEST instruction while the GRANT signal is logic high, it will change to the BLOCKING state.  
20 Further operations in the BLOCKING state will be the same as described above.

An alternative embodiment of the invention is disclosed in Fig. 5. As shown therein, the apparatus comprises an array of processors 110, an array of alignment units 130, an array of combinatorial logic cells 130 and an  
25 interconnection network 140. The processors 110, alignment units 120 and combinatorial logic cells 130 may be the same as processors 10, alignment units 20 and logic cells 30 of Fig. 1 and processors 110 may be controlled by a control unit that is not shown. Interconnection network 140, however, is  
30 different from that of network 40 in that it is implemented in the form of separate routers 142 and connection paths 144 between each pair of processors.

In such a case, it is possible to test for the absence of message traffic in the interconnection network by testing on a processor by processor basis. Thus as shown in Fig. 5, the router for each processor produces an output signal which is a logic high when there is no message traffic on the connection paths to or from the processor with which it is associated. This signal is applied as one input to an AND gate 146 which has as a second input the output of the alignment unit 130 associated with that processor. As a result, a high logic signal is provided by AND gate 146 to an input to logic cell 130 when there is no message traffic in the portion of the interconnection network associated with that processor and a logic high ASSERT signal has been generated by the alignment unit. As will be apparent, this changes the location of the test for no message traffic from the uppermost cell in the binary tree of logic cells 30 of Fig. 1 to the lowermost cells in the binary tree of logic cells 130 of Fig. 5. In other respects the operation of the alignment units and logic cells is the same as described for Figs. 1-4.

The invention may be practiced in numerous environments. For example, if the processors are Intel 80386 microprocessors, the invention may be practiced using the IN and OUT instructions of the 80386 instruction set to provide instructions to an alignment unit and receive an ALIGN signal from such unit.

30

35



## Illustrative code comprises:

	1.	MOV	AL, 5	load an address, 5, of the processor to which data is to be sent
5	2.	OUT	sending device, AL	send data via sending device
	3.	ADD	AL, CL	compute data of message
	4.	OUT	sending device, AL	send data via sending device
10	5.	MOV	AL, 1	load into AL a 1-bit representing a request for alignment
	6.	OUT	align unit, AL	send the 1-bit in register AL to the alignment unit
15	.			20
	.			
	.			
20	20.	IN	AL, align unit	load the ALIGN signal from the alignment unit into register AL
	21.	TEST	AL, 1	test if the signal in register AL is a 1-bit
	22.	INE	20	if not, return to instruction at line 20
25	23.	IN	AL, receiving device	if it is a 1-bit, proceed with program
				illustratively by loading signals from receiving device into register AL
30				

In the practice of the invention, the request for alignment is made by each processor wherever it is necessary for inter-processor communications to have been completed before the processor can continue with its program. In the

illustrative code set forth above, the request for alignment is made at line 6. The requests for alignment are processed by the alignment units 20 associated with each processor and are forwarded to logic cells 30 as logic high ASSERT signals. These requests are ANDed together by logic cells 30 and when all processors have made such a request a logic high signal is applied to AND gate 60. If the interconnection network is empty, a logic high READY signal is broadcast by AND gate 60 to all the alignment units 20 via logic cells 30.

In the illustrative code set forth above, the logic level of the READY or GRANT signal is monitored by the instructions at lines 20, 21, and 22, which respectively read this signal level, test if it is a 1-bit or logic high, and recycle if the 1-bit has not been received. These instructions provide the function of the TEST instruction and the ALIGN signal of Table I.

The processors are all aligned at the time the READY signal goes logic high which occurs at some point between the request for alignment at line 6 and the time each processor initiates a successful test for a 1-bit in the AL register as at lines 20-23. It is not necessary, however, that each processor make such a test at the same time; and it is not necessary that the first such test be made after the READY signal goes logic high. If the logic high READY signal is received before the program reaches lines 20-21, the alignment unit will save that information until those lines are reached; and if the READY signal is still logic low when the program reaches line 20, the program will simply loop through the instructions at lines 20-22 until the READY signal goes logic high.

As will be apparent to those skilled in the art, numerous modifications may be made in the practice of the invention. For example, while the invention has been described in terms of apparatus (AND gate 60) which tests for completion of communication on an interconnection network, it

may be desirable to make other tests in place of or in addition to such a test of the communication network. Thus, if the computer systems includes mathematics co-processors or the software includes various subroutines whose outputs might  
5 alter data to be communicated from one processor to another, it will be desirable to test such equipment or operations to make sure they have completed their tasks before issuing the READY signal to the processors. Such a test can readily be made by using a global OR network to test the state of an  
10 activity flag associated with each hardware or software element whose operation must be completed before the program can proceed.

15

20

25

30

35

What is claimed:

1. In a computer comprising a plurality of processors each of which operates on its own stream of instructions, at least one instruction requiring prior completion of a prior instruction for proper execution, a method of aligning said processors comprising the steps of:
  - inserting in the instruction stream of each processor an instruction requesting alignment;
  - testing for each processor that requests alignment whether all other processors to be aligned have requested alignment;
  - testing for prior completion of any instructions that must be completed; and
  - before proceeding, waiting at each processor until all other processors to be aligned have requested alignment and any instructions that must be completed have been completed.
2. The method of claim 1 wherein the instruction requesting alignment is inserted in the instruction stream at or after an instruction in the instruction stream which creates a situation that requires alignment.
3. The method of claim 1 wherein the step of testing for prior completion of any instructions comprises the step of testing for the absence of message traffic in an interprocessor communication means before proceeding.
4. In a computer comprising a plurality of processors each of which operates on its own stream of instructions, at least one instruction requiring prior completion of a prior instruction for proper execution, means for aligning said processors comprising:

35

means at each processor for testing the instruction stream of each processor for the presence of a request for alignment, a test, or some other instruction;

5 means for determining that all the testing means have detected a request for alignment in the instruction streams of all the processors to be aligned;

means for testing for prior completion of any instructions that must be completed; and

10 means for indicating to each processor to be aligned that all the testing means have detected a request for alignment in the instruction streams of all the processors to be aligned and any instructions that must be completed have been completed.

15

20

25

30

35

FIG. 1

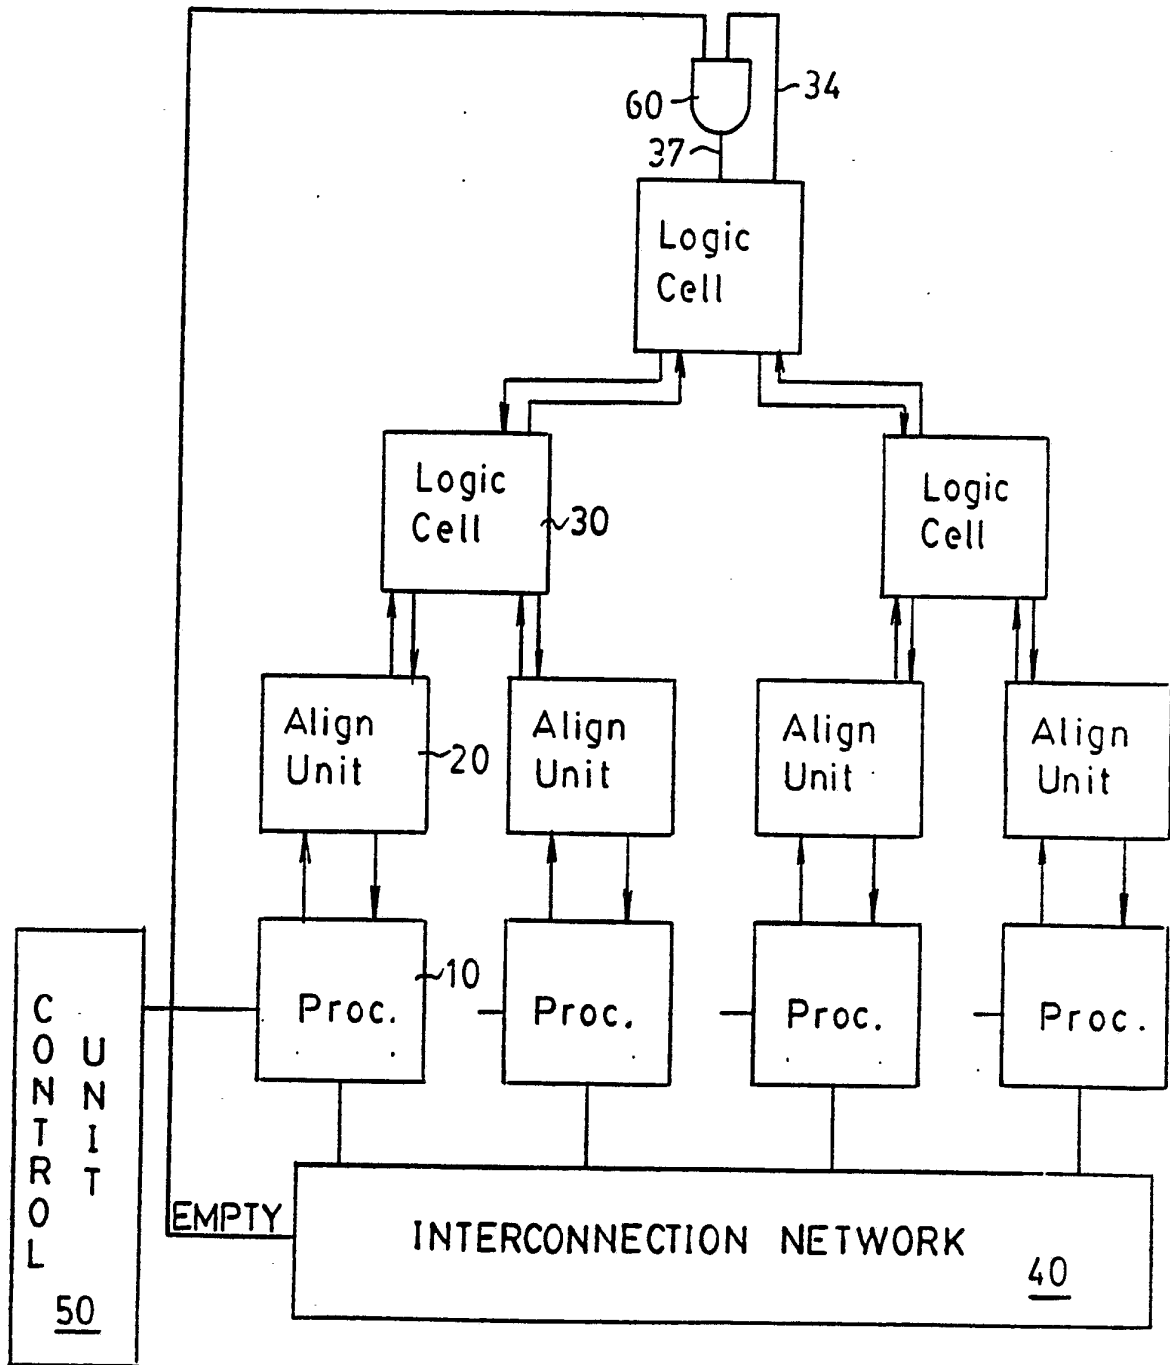


FIG. 2

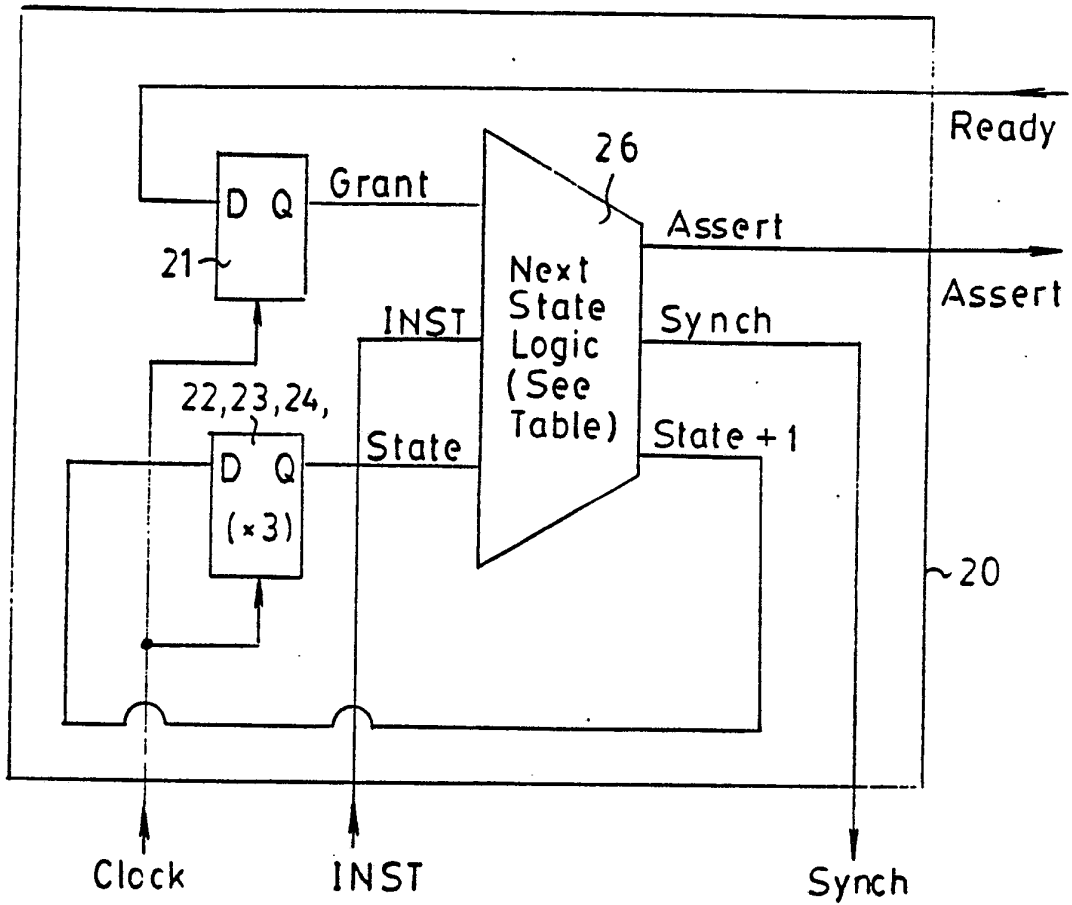
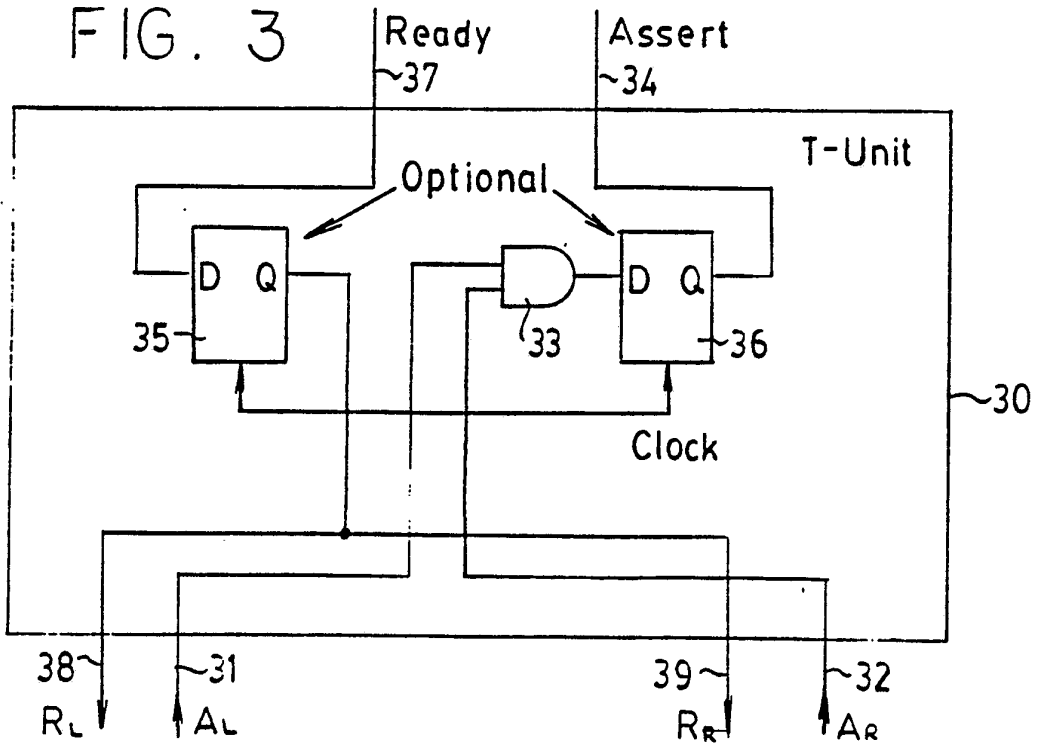


FIG. 3



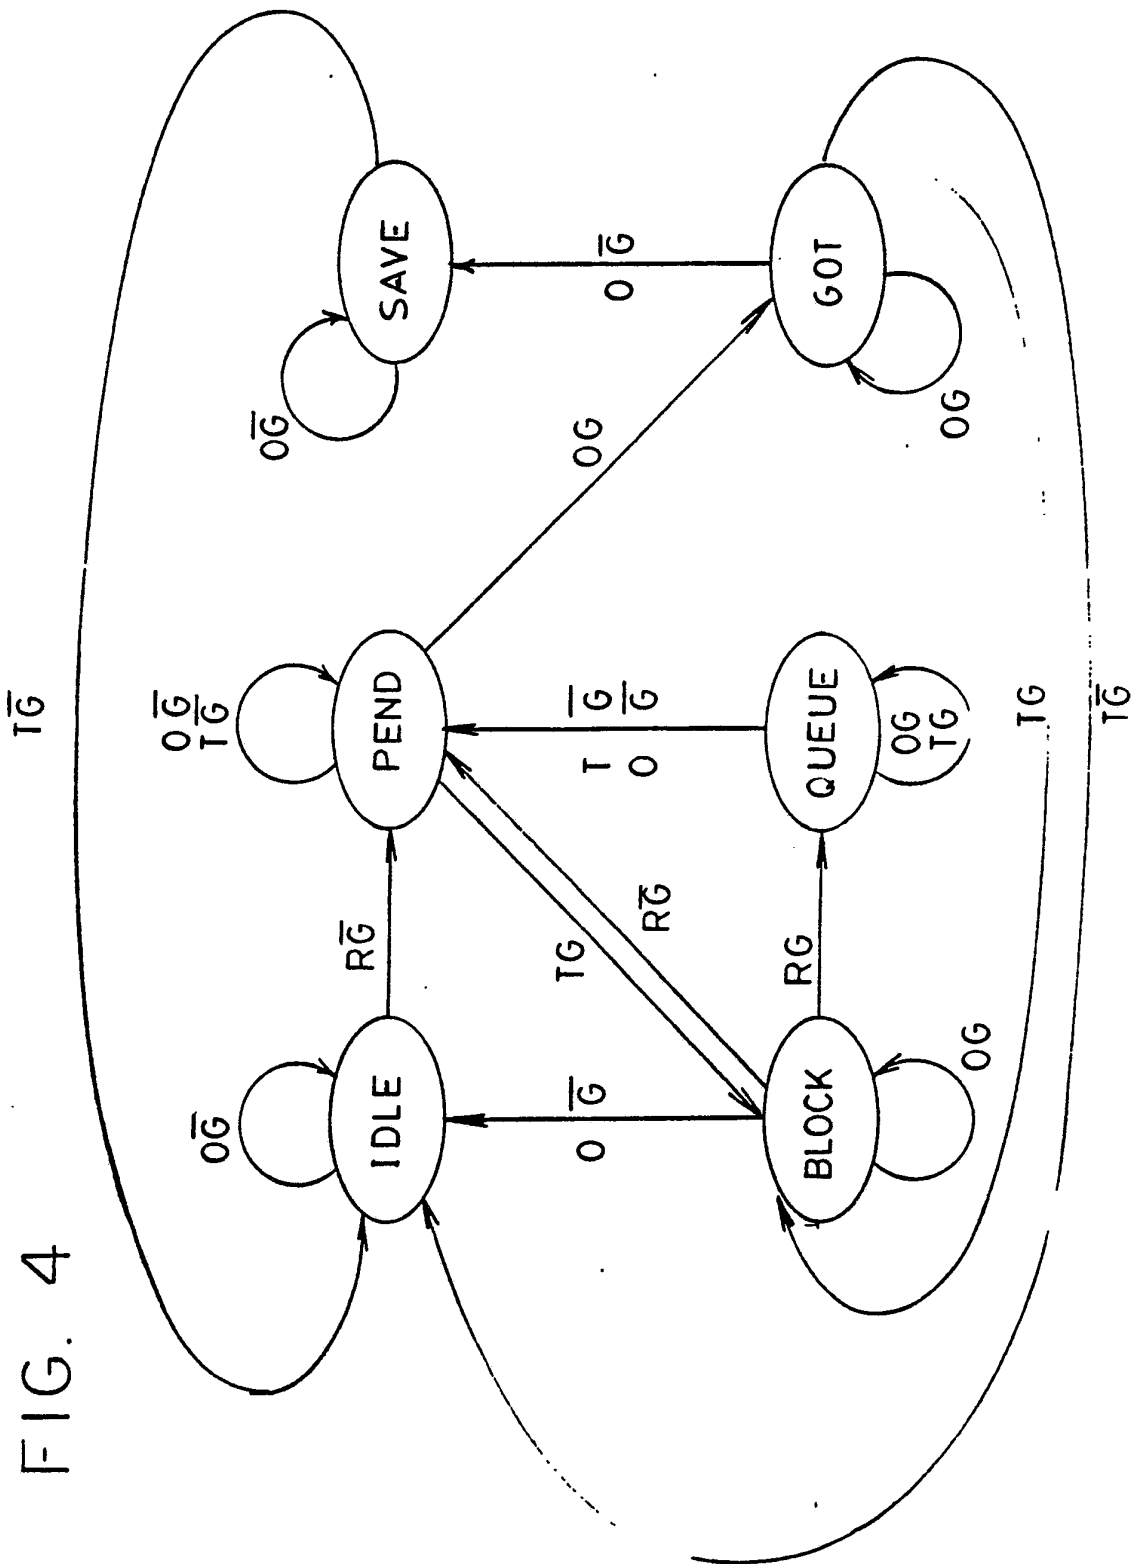
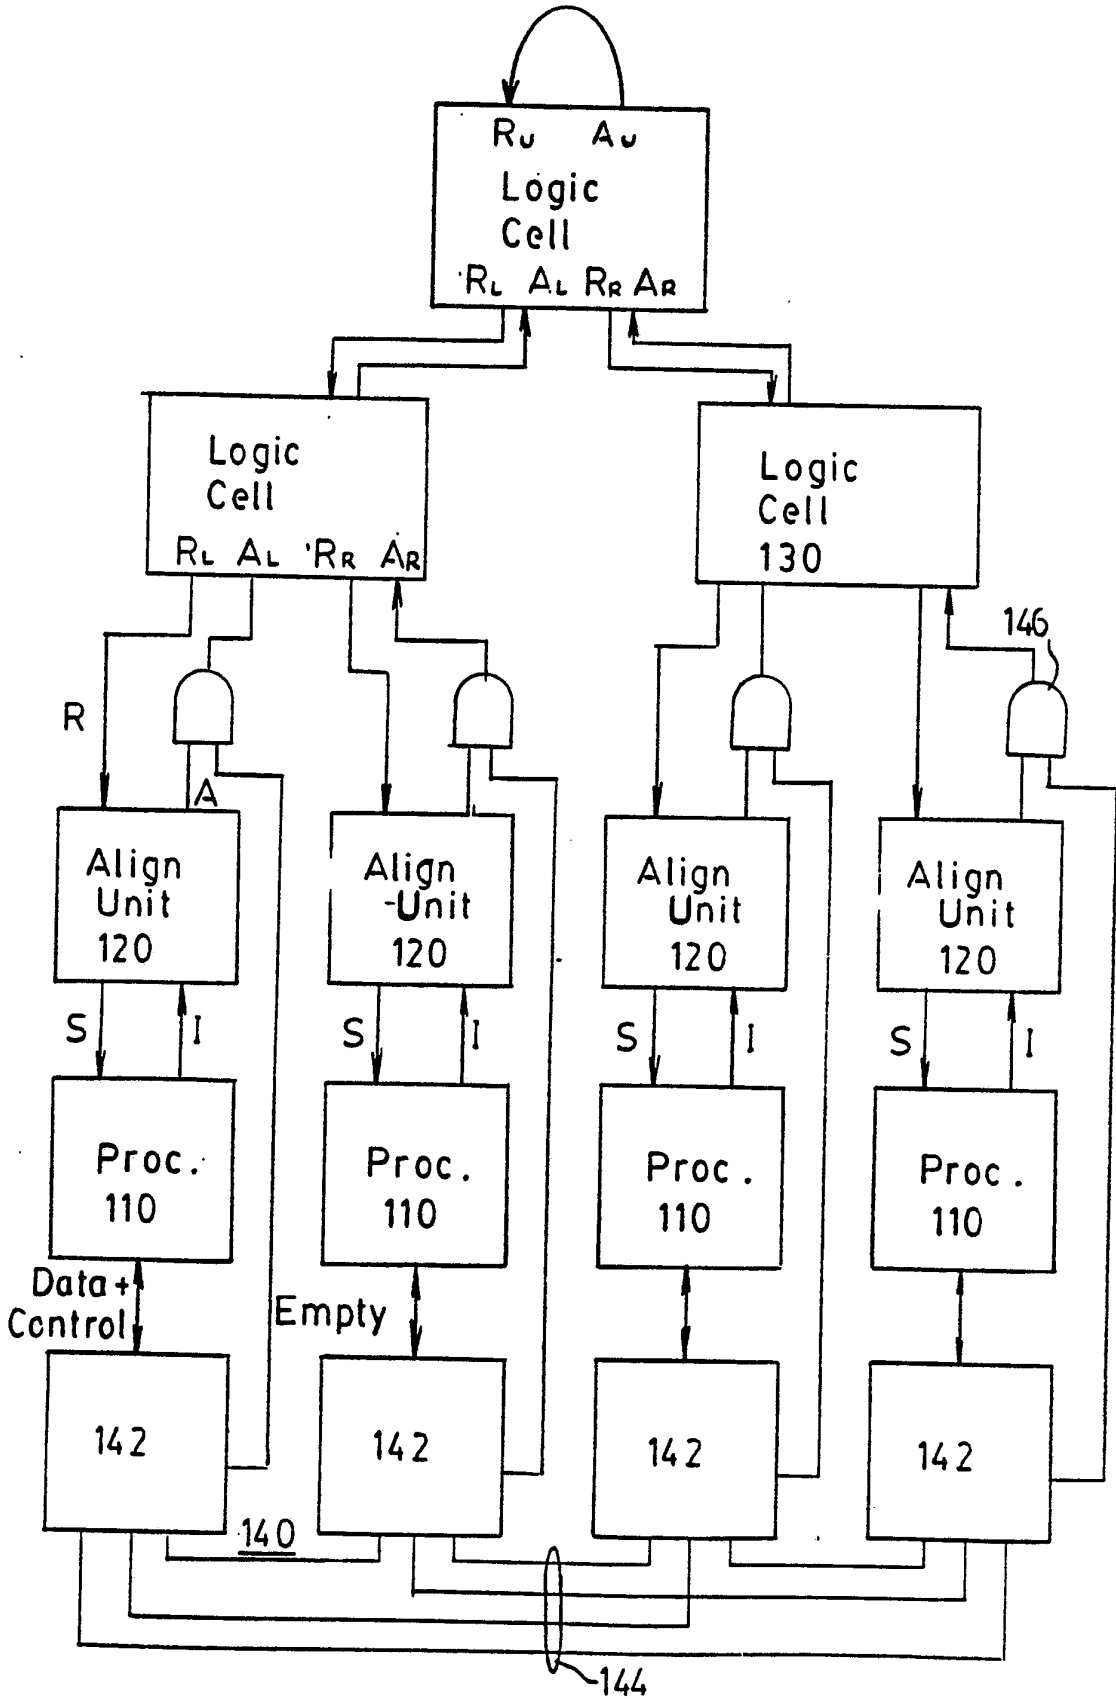


FIG. 4



FIG. 5



# INTERNATIONAL SEARCH REPORT

International Application No. PCT/US89/00395

<b>I. CLASSIFICATION OF SUBJECT MATTER</b> (if several classification symbols apply, indicate all) <sup>6</sup>		
According to International Patent Classification (IPC) or to both National Classification and IPC		
IPC(4): G06F 15/16		
U.S. Cl. 364/200		
<b>II. FIELDS SEARCHED</b>		
Minimum Documentation Searched <sup>7</sup>		
Classification System	Classification Symbols	
U.S.	364/200,900 371/47,19	
Documentation Searched other than Minimum Documentation to the extent that such Documents are included in the Fields Searched <sup>8</sup>		
<b>III. DOCUMENTS CONSIDERED TO BE RELEVANT <sup>9</sup></b>		
Category <sup>*</sup>	Citation of Document, <sup>11</sup> with indication, where appropriate, of the relevant passages <sup>12</sup>	Relevant to Claim No. <sup>13</sup>
Y P	US,A, 4,733,353 (JASWA) 22 March 1988. See entire document.	1-4
Y P	US,A, 4,775,934 (HOURI et al.) 4 October 1988 See entire document.	1-4
Y	US,A, 4,229,790 (GILLILAND et al.) 21 October 1980. See entire document.	1-4
A	US,A, 4,554,626 (KATZ et al.) 19 November 1985.	1-4
A	US,A, 4,399,504 (OBERMARCK et al.) 16 August 1983.	1-4
A	US,A, 4,394,725 (BIENVENU et al.) 19 June 1983.	1-4
A	US,A, 4,330,826 (WHITESIDE et al.) 18 May 1982.	1-4
<p><sup>*</sup> Special categories of cited documents: <sup>10</sup></p> <p>"A" document defining the general state of the art which is not considered to be of particular relevance</p> <p>"E" earlier document but published on or after the international filing date</p> <p>"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)</p> <p>"O" document referring to an oral disclosure, use, exhibition or other means</p> <p>"P" document published prior to the international filing date but later than the priority date claimed</p> <p>"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention</p> <p>"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step</p> <p>"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.</p> <p>"&amp;" document member of the same patent family</p>		
<b>IV. CERTIFICATION</b>		
Date of the Actual Completion of the International Search		Date of Mailing of this International Search Report
16 FEBRUARY 1989		12 APR 1989
International Searching Authority		Signature of Authorized Officer
ISA/US		PAUL KULIK 