

(12) INTERNATIONAL APPLICATION PUBLISHED UNDER THE PATENT COOPERATION TREATY (PCT)

(19) World Intellectual Property Organization
International Bureau




(43) International Publication Date
22 July 2010 (22.07.2010)

(10) International Publication Number
WO 2010/082067 A2

- (51) **International Patent Classification:**
G06F 15/173 (2006.01) H04L 12/43 (2006.01)
- (21) **International Application Number:**
PCT/GB20 10/050064
- (22) **International Filing Date:**
15 January 2010 (15.01.2010)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (30) **Priority Data:**
0900667.7 16 January 2009 (16.01.2009) GB
- (71) **Applicant (for all designated States except US):** **THE UNIVERSITY OF READING** [GB/GB]; Whiteknights, P.O. Box 217, Reading Berkshire RG6 6AH (GB).
- (72) **Inventors; and**
- (75) **Inventors/Applicants (for US only):** **LEACH, Stephen Frederick Knight** [GB/GB]; 32 Rodney Road, Backwell, North Somerset BS48 3HB (GB). **ANDERSON, James Arthur Dean Wallace** [GB/GB]; 88 Lower Henley Road, Caversham, Reading RG4 5LE (GB).
- (74) **Agent:** **BARKER BRETTELL LLP**; 100 Hagley Road, Edgbaston, Birmingham B16 8QQ (GB).
- (81) **Designated States (unless otherwise indicated, for every kind of national protection available):** AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PE, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.
- (84) **Designated States (unless otherwise indicated, for every kind of regional protection available):** ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).
- Published:**
— without international search report and to be republished upon receipt of that report (Rule 48.2(g))

(54) **Title:** PROCESSORS

Fig. 4



(57) **Abstract:** A processing system comprises a plurality of processors (12) and communication means (20) arranged to carry messages between the processors, wherein each of the processors (12) has an operating instruction memory field (32, 34, 36) arranged to hold stored operating instructions including a re-routing target address. Each processor is arranged to receive a message (38) including operating instructions including a target address. On receipt of the message, each processor is arranged to: check the target address in the message to determine whether it corresponds to an address associated with the processor; if the target address in the message does correspond to an address associated with the processor, to check the operating instructions in the message to determine whether the message is to be re-routed; and, if the message is to be re-routed, to replace operating instructions within the message with the stored operating instructions, and place the message on the communication means for delivery to the re-routing target address.



WO 2010/082067 A2

PROCESSORS

The present invention relates to communication between processors.

5 Typically communication between processors is arranged as the transmission of a packet composed of a header and data. The packet is sent from a single source processor and may be sent to a single destination processor; to a fixed number of destination processors, commonly two; to all processors in an acyclic graph (tree) of processors;
10 to all processors in a cyclic graph (loop) of processors; or it may be broadcast to all processors in an arbitrary graph in the expectation that some of the processors are prepared to accept the packet. Communication may also be via memory when a processor stores data in memory and a processor fetches data from memory.

15

Some prior art systems will now be described. These use a number of common principles, seven of which are identified, and features of the systems described which use the same principles are identified by the same reference numerals from (i) to (vii). When a packet is sent from its
20 source to a single destination, the header may encode the destination and/or the source. Source processors may (i) forward the packet without modifying it. Alternatively, they may use information in the header to look up an intermediate destination. The packet may then be (ii) sent directly to the processor at the intermediate destination. This implies that
25 a source processor has access to a dedicated route to an intermediate processor. Alternatively, (iii) the address of the intermediate processor may be added to the header, before the packet is sent on its way via (ii) a dedicated or (iv) a widely accessible route. Adding information to the header makes the header longer. It can be seen that sending a packet from
30 a single source to a single destination, using methods known in the art, is either (i) a fixed operation or else (ii) it implies having a dedicated route

or else (iii) it implies that the header is of variable size. The routing of a packet from a single source to a fixed number of destinations is essentially similar.

5 When a packet is sent from its source to all processors in a tree it may be (ii) sent on a dedicated route or (iii) all of the destinations in the tree may be encoded in a variable length header. These alternatives, (ii) and (iii), are the same as above. The routing of a packet from a single source to all processors in a loop is essentially similar.

10

When a packet is broadcast from its source to all processors in a graph (v) the packet is delivered to many processors but, generally, only some of the processors accept the packet. This generally implies that routes to some or all of the non-accepting processors are wasted.

15

When a processor sends data to memory it commonly sends the data to the destination address via (iv) a widely accessible route. It may send the data to a single address or it may send data indirectly via a second and a third distinct addresses. In principle (vi) any number of indirections that is
20 fixed and known at the source could be used. But, in the art, this does not exceed three addresses: two indirect addresses and one final address, all of which addresses are distinct. The alternative, (iv), is the same as above. The fetching of data from memory is similar.

25 A processor may also send data to memory via a sequence of caches. In this case (i) a packet composed of the address and data is unmodified, and is (ii) sent via a direct route to a fixed sequence of caches until the memory transaction is complete. These alternatives, (i) and (ii), are the same as above. The fetching of data from memory via a sequence of
30 caches is essentially similar.

In all of the above communication methods, (vii) a packet may be delivered to an intermediate processor and/or it may be passed on to a successor processor. However, in some embodiments of a computer architecture of our own devising, which we have simulated in software, we provide an apparatus for communicating between processors which we find reduces, by up to 30%, the number of instructions needed by our parallel computer to execute a given program, as compared to any combination of the above communication methods known in the art. More generally, some embodiments of our invention improve the performance of a class of parallel computers. Some embodiments of our apparatus allow a packet to be modified or to be (i) fixed; to be sent via (ii) a dedicated or (iv) a widely accessible route to one or more processors; but do not require the header to (iii) vary in size; they do allow (v) a packet to be broadcast. They allow any number of, not necessarily distinct, indirections to be used, that number not being known or fixed at the source, in contrast to (vi). They do allow (vii) a packet to be delivered to an intermediate processor and/or to a successor processor. In short, in some embodiments of our invention, many of the communication methods known in the art are allowed, but are made more efficient for applications in parallel computing, by using a fixed size of header, whose contents may be modified, and by allowing a number of, not necessarily distinct, indirections which are generally not known or fixed at the source. In this way an arbitrarily complex route is defined which can be followed by a message having a fixed length header.

25

Accordingly the present invention provides a processing system comprising a plurality of processors and communication means arranged to carry messages between the processors. Each of the processors may have an operating instruction memory field arranged to hold stored operating instructions, which may including a re-routing target address. Each processor may also be arranged to receive a message including

30

operating instructions, which may include a target address. On receipt of the message each processor may be arranged to perform any one or more of the following steps: check the target address in the message to determine whether it corresponds to an address associated with the processor; if the target address in the message does correspond to a target address associated with the processor, to check the operating instructions in the message to determine whether the message is to be re-routed; and if the message is to be re-routed to replace operating instructions within the message with the stored operating instructions, and place the message on the communication means for delivery to the re-routing target address.

The target address may comprise a processor address, or it may comprise the address of a register or other memory field associated with a particular processor. For example it may comprise a high level address field indicating a processor, and a low level address field which may indicate a register or other memory field within the processor, or may not be associated with any specific register or memory field within the processor. In this case it includes both a processor address and a register or memory field address.

The present invention further provides a processing system comprising a plurality of processors, which may form an array, and a communication means, which may be a pipe, arranged to carry messages between the processors, wherein each of the processors has a re-routing address memory field arranged to hold a processor address. Each processor may also have a delivery instruction memory field arranged to hold a delivery instruction. Each processor may be arranged to receive a message including a target address and a re-routing instruction. Each processor may be arranged, on receipt of the message to check the target address to determine whether it corresponds to the processor's own address. If the target address does correspond to the processor's own address, the

processor may be arranged to check the re-routing instruction to determine whether the message is to be re-routed. If the message is to be re-routed the processor may be arranged to put the re-routing address into the message. This may replace a re-routing address in the message. It
5 may be arranged to put the delivery instruction into the message to indicate to the processor at the re-routing address whether it is to re-route the message. It may then place the message on the communication means for delivery to the re-routing address.

10 Each processor may be arranged, if the message is to be re-routed, to replace an address, which may be a processor address, and which may be a used address, within the message with the re-routing address. The used address may be the processor's own address or associated with the processor. The message may include a routing address field holding at
15 least one address, which may comprise a processor address, for use in routing the message, including the target address, and if the message is to be re-routed the address which the processor is arranged to replace with the re-routing address may be an address held in the routing address field. This enables the length of the message to remain the same, for example
20 including a fixed number of address fields, which may be one field or a plurality of fields, containing a fixed number of addresses, which again may be one or a plurality, as it is transmitted between processors. The length of the message can therefore be independent of the length of the route which it is to follow. If there is only one address field, then the
25 address in that field will generally be the target address, and the processor may be arranged to replace the target address in the received message, which will generally be its own address or associated with it, with the re-routing address, so that the re-routing address forms the target address in the re-routed message. If there is more than one address field
30 within the message, then the re-routing address stored at the processor may replace the processor's own address as target. Alternatively it may

replace another address, for example an address which has previously been used as a target address. In other words, the re-routing address may be used to form the target address of the re-routed message, or it may be inserted into the message for use as a target address in a subsequent re-
5 routing.

The operating or delivery instruction in the message may include a re-routing instruction. The stored operating instruction may include a further re-routing instruction. Each processor may be arranged, if the message is
10 to be re-routed, to replace a re-routing instruction within the message, which may be a used re-routing instruction, with the further re-routing instruction from the operating or delivery instruction memory field.

The operating or delivery instruction may comprise a store instruction
15 indicating whether the message is to be stored at the processor at the re-routing address.

Each processor may be arranged to determine from the message whether the message is to be stored, and if it is, to store at least a part of the
20 message. Each message may be arranged to determine from a delivery instruction or operating instruction in the received message whether the message is to be stored. The operating instruction or delivery instruction may include a re-routing instruction.

25 Each of the processors may be further arranged to perform an instruction using a part of the message as an input.

The processor may have a further address memory field arranged to hold a further re-routing address and, if the message is to be re-routed, the
30 processor may be arranged to place the message on the communication means for delivery, or re-routing, to the further re-routing address.

The delivery instruction or re-routing instruction may have a fixed number of possible states indicating, preferably independently of the processor at which it is stored and the re-routing address to which it is to be sent, whether the message is to be re-routed at the processor at the re-
5 routing address. The delivery instruction, or the store instruction, may also indicate, preferably in a similarly consistent way, whether the message is to be stored at the processor at the re-routing address.

10 The communication means may comprise a pipe and may comprise a plurality of frames and be arranged to move the data and/or the control tokens between frames to move them along the communication means. Each processor may have one or more of the frames associated with it from which data can be written to the processor.

15

The data may be carried on the communication means in the form of data tokens.

Each processor may be arranged to perform the same instruction as every
20 other processor. Each processor may be arranged to perform only one instruction. Each time that each processor performs its instruction it may be able to identify no, one, or more than one, data target processor, and no, one, or more than one control target processor. This can enable it to transmit data to multiple processors in parallel. The communication
25 means may be arranged to transmit the result to the identified data target processor so that it can be written into the data target processor.

Each processor may be arranged to transmit a control token by writing it to the communication means together with the address of the control
30 target processor to which it is to be transmitted. Each processor may be

able, on performing the instruction, to identify a plurality of control target processors to which it can transmit control tokens in parallel.

Each processor is preferably arranged, when it has transmitted the result
5 and a control token to any identified target processors, to give up possession of its control token so that it will not perform the instruction again until it receives another control token.

The instruction may be a multiply-add of the form $a \times b + c \rightarrow r'$.

10

Each processor may be arranged to select any control processors on the basis of the result r' . For example each processor may be arranged to determine whether the result r' is less than zero, zero, greater than zero, or nullity (as defined in J. A. D. W. Anderson, Norbert Völker, Andrew
15 A. Adams "Perspex Machine VIII: Axioms of Trarisreal Arithmetic" in *Vision Geometry XV*, Longin Jan Lateki, David M. Mount, Angela Y. Wu, Editors, *Proceedings of SPIE* Vol. 6499 (2007), and select any control target processors or data target processors accordingly.

20 Each processor may comprise a plurality of memory cells in which the inputs to the instruction, which may be received over the communication means, can be stored. Each processor may comprise a plurality of memory cells in which addresses of control target processors can be stored. All memory in every processor may be set to a fixed value on
25 power up so that the processors do not execute arbitrary programs as would otherwise be the case with arbitrary values being set on power up.

The processors may be each arranged to perform an instruction once in each of a series of operating cycles, and the communication means may
30 comprise a plurality of pipes each arranged to transmit data in a respective direction. Each processor may comprise a plurality of output

means each arranged to output data onto a respective one of the pipes in the form of tokens. At least one of the output means may be arranged to output two tokens onto the same pipe in one operating cycle of the processor.

5

Indeed the present invention further provides a processing system comprising a plurality of processors each arranged to perform an instruction once in each of a series of operating cycles, and communication means arranged to carry messages between the processors, wherein the communication means comprises a plurality of pipes each arranged to transmit data in a respective direction, and each processor comprises a plurality of output means each arranged to output data onto a respective one of the pipes in the form of tokens, at least one of the output means being arranged to output two tokens onto the same pipe in one operating cycle of the processor.

10
15

Each pipe may comprise a plurality of memory frames. The communication means may be arranged to transfer tokens between adjacent pipe frames, to carry the messages, at a pipe frequency which may be higher than that of the operating cycle of the processors.

20

The processor may be arranged to output two tokens to the same pipe frame at different times within the operating cycle of the processor and the communication means may be arranged to transmit the first of the tokens to another pipe frame before the second token is output onto the pipe.

25

The present invention further provides a method of setting up a processing system according to the invention. The method may comprise loading a respective re-routing address into the re-routing address memory field of a plurality of the processors. The method may comprise

30

loading a respective delivery instruction into the delivery instruction memory field of a plurality of the processors. The method may thereby define a route along which data can be communicated between a plurality of the processors,

5

Preferred embodiments of the present invention will now be described by way of example only with reference to the accompanying drawings in which:

10 **Figure 1** is a schematic diagram of a processor chip according to an embodiment of the invention;

Figure 2 is a schematic view of one processor and sections of pipes of the chip of Figure 1;

15

Figure 3 is a schematic diagram of a part of one of the pipes of the chip of Figure 1;

20 **Figure 4** is a diagram showing the input and re-routing output of the processor of Figure 2;

Figure 5 is a diagram showing the input and output of the processor 12 of Figure 2 operating in a different mode;

25 **Figure 6** is a schematic diagram of a processor forming part of a further embodiment of the invention; and

Figure 7 is a schematic diagram of a processing system according to a further embodiment of the invention.

30

Referring to Figure 1, a processor chip 10 comprises a two dimensional rectangular array of processors 12. Each processor or tuple 12 has an address defined by perpendicular coordinates x , y . The main array is of finite size with addresses having y coordinates centred on the origin and extending from $-y_{\max}$ to $+y_{\max}$ and x coordinates also centred on the origin and extending from $-x_{\max}$ to $+x_{\max}$. At each end of each row and each column of the main array of processors 12 is an input/output processor or I/O device 14 which has either an x coordinate of $+\infty$ or $-\infty$ or a y coordinate of $+\infty$ or $-\infty$. Only one quarter of the chip 10 is shown in Figure 1, which is in the positive quadrant of the coordinate system. A number of chips can be combined to form a single machine with data moving from one chip to the next via the I/O devices 14.

A series of pipes 20 are laid out in a rectangular grid between the processors 12. Between each column of processors 12, there are a pair of pipes extending in the y direction, one $+Y$ being arranged to transfer data in the positive Y direction and the other $-Y$ being arranged to transfer data in the negative Y direction. Between each row of processors 12, there are a pair of pipes extending in the X direction, one $+X$ being arranged to transfer data in the positive X direction and the other $-X$ being arranged to transfer data in the negative X direction. Each pair of pipes 20 is shown as a single line in Figure 1, but the parts of the individual pipes 20 surrounding one of the processors 12 are shown in Figure 2. Figure 2 therefore shows a single processor tile 22 which forms a basic unit which is repeated across the chip 10 to build up the entire chip. Each processor 12 is connected to each of the four pipes 20 that are adjacent to it on its four sides, so that each processor can put data onto the appropriate pipe for it to be transmitted in any of the four directions.

Referring to Figure 3, each of the uni-directional pipes 20 comprises a series of pipe frames 24 each of which is made up of a number of memory

cells, and a parallel series of temporary frames 26 each of which is made up of the same number of memory cells. Each of the temporary frames 26 is connected to two adjacent pipe frames 24 so that it can receive data from one of them and output data into the other of them. Therefore data
5 can be transmitted along the pipe by being transferred from one pipe frame 24 to the next one in the direction of the pipe, via the appropriate temporary frame 26. Each processor 12 is connected to one of the pipe frames 24 on each of the pipes that pass adjacent to it, so that it can accept data from the pipe, and write data to the pipe at that point. Each
10 processor 12 has two output channels 29a 29b which allow the processor to put two outputs onto the pipe so that they will be transmitted in different pipe frames 24. These are shown schematically as parallel channels formed of parallel wires, communicating with adjacent pipe frames 24, but can, in practice, be in the form of a single wire which can
15 transmit two outputs using time multiplexing to provide the two channels. In that arrangement the processor is arranged to output two tokens to the same pipe frame at different times within the operating cycle of the processor and the pipe is arranged to transmit the first of the tokens to another pipe frame before the second token is output onto the pipe.

20 All of the processors 12 and all of the pipe frames 24 and temporary frames 26 are connected to a common clock line 28, on which a clock signal is transmitted. This is used to coordinate the timing of the transfer of data along the pipes 20 and between the pipes 20 and processors 12, as
25 well as the performance of instructions by the processors. The data in a pipe frame is copied to its adjacent pipe frame, via a temporary pipe frame, every clock tick. Generally the moving of data between pipe frames 24 occurs at a frequency that is higher than the frequency at which each of the processor 12 performs its instruction. This allows a processor
30 12 to transmit more than one datum along a single pipe in one processor operating cycle, with each datum being put onto the pipe 20 via a

respective one of the output channels 29a, 29b. The processors can each perform their instruction once every operating cycle. All the processors have the same operating cycle length, or 'processor clock tick', which may be equal in length to two or more pipe clock ticks. Barring
5 fabrication errors, processors 12 have a neighbouring tuple on either side, and I/O tuples have a neighbouring tuple on exactly one side.

Each processor 12 is arranged to perform its instruction, which may be the same single instruction for all of the processors, having the form:

10

$$\begin{aligned} &u \times v + w \rightarrow r' \\ &\text{write}(r', r) \\ &\text{jump}(r', i, z, g, n) \end{aligned}$$

15 as described in international patent application No. PCT/GB2008/001821. Alternatively the processors 12 may perform different instructions and each processor may be arranged to perform more than one instruction.

Each processor 12 has a buffer that accepts tokens from the pipe 20.
20 When the processor 12 computes it copies the buffer into internal registers and operates on the internal registers.

Messages on the pipe 20 are transmitted in the form of tokens. Each token in this embodiment is either a data token or a control token, or both, as
25 indicated by an appropriate flag or flags within the token. Each processor 12 is arranged to receive data in the form of data tokens, and to perform its instruction when it receives a control token. When it has performed its instruction, each processor 12 can send on data as part of an output message in the form of a token which may include the result of the
30 instruction, to one or more other processors 12 in the form of data tokens, and pass control to one or more other processors 12 using a

control token. Therefore it will be appreciated that each processor 12 needs to be able to output two or more tokens onto the pipes 20 each time it performs its instruction, i.e. in each processor operating cycle. These tokens may be placed onto different pipes 20 to be transmitted in different directions, but they may be placed on the same pipe 20 for transmission
5 in the same direction, and this is enabled by the multiple output channels 29a, 29b. In some cases more than two output channels may be required on each edge of the processor.

10 Referring to Figure 4, each processor 12 includes a memory store 50 in which received data can be stored and processing components 52 which are arranged to perform an instruction producing a result which can be included in a processing output message 54 for transmission to another processor. The format of the processing output message can be the same
15 as the received message, but will generally include a header h_0 and data d_0 . In addition to its ability to perform one or more instructions to generate the processing output, each processor is arranged to perform a re-routing function on a received token. To achieve this, each of the processors 12 includes an 'own address' memory field 30 in which its
20 own address a_1 is stored, a hopping address field 32 in which a hopping address a_2 can be stored, a delivery instruction field 34, 36 which includes a store instruction field 34 and hop instruction field 36, each of which contains a one bit flag having two states: a 'set' state and a 'cleared' state. More generally the delivery instruction can be considered
25 as a single instruction that can have any of a fixed number of states, in this case four, indicating whether the message is to be hopped (re-routed) at the next processor and whether the message is to be stored at the next processor. The delivery instruction is independent of the processor at which it is stored and the re-routing address to which it is to be sent.

The processor 12 is arranged to receive a message in the form of a token 38 which includes a data field 46 including data d , and a header which holds operating instructions indicating what operations the processor is to perform on the data. The header includes an address field 40, a delivery instruction field including a store instruction field 42, and a hop instruction field 44. The address field 40, which may be considered a target address field, includes the address a_j of the target processor to which the token is being sent, the store instruction field 42 contains a store instruction which, in this case, is in the form of a one bit flag, which can be set or cleared, indicating whether the data is to be stored at the target processor, and the hop instruction field 44 contains a hop instruction which, in this case, is in the form of a one bit flag, which can be set or cleared, indicating whether the data is to be 'hopped' i.e. sent on from the target processor to another processor.

15

Each processor is therefore arranged so that it can act as a 'signpost' which receives a message addressed to it together with a flag indicating whether it is to be sent on. If the message is to be sent on, the processor inserts a new address into the message, inserts a further delivery instruction into the message indicating whether it is to be stored and/or sent on again from the next processor, and sends the message on to the next processor. In order to achieve this, when the message 38 arrives at the processor 12, the processor 12 is arranged to compare the address a_j in address field 40 of the message with its own address. If these addresses do not match, the message is not addressed to the processor, so it simply allows the message to move on along the pipe. If the addresses do match, the processor 12 checks the store instruction field 42 and the hop instruction field 44 in the received message. If the store instruction flag S_j is set indicating that the data in the message is to be stored, the processor 12 stores the data within its own memory 50, for example so that it can perform an instruction on the received data using the processing

components 52 to produce a result which can also be placed on the pipe as a 'processing output' 54 for transmission to another processor. If the store instruction flag S_1 is clear indicating that the data is not to be stored, then it is not stored. If the hop instruction flag h_1 is set to indicate that the data is to be hopped to another processor, then the processor 12 is arranged to insert the address a_2 from its hop address field 32 into the message as a new target address in the address field 40, replacing its own address a_1 which is currently there, and insert the delivery instruction from its delivery instruction field 34, 36 into the message. Specifically this includes updating the header fields 40, 42, 44 of the message by inserting the store instruction flag S_2 from its store instruction field 34 and the hop instruction flag h_2 from its hop instruction field 36 into the appropriate fields of the message to replace the delivery instruction which is currently there. Then, the processor 12 is arranged to place the message 38 onto the appropriate one of the pipes 20, as a re-routing output, for delivery to the processor identified in the hop address field. The target processor address a_2 and the store and hop instructions S_2 , h_2 are retained in the appropriate memory fields of the processor 12, so that any further tokens targeted to the processor can also be re-routed along the same route.

If the hop instruction flag h_1 in the received message is clear indicating that the message is not to be transmitted on to another processor, then the processor 12 can either not re-transmit the message at all, or it can place it on the pipe as an empty token, for transmission to, and re-use by, another processor. Marking both the hopping instruction field and the store instruction field of the re-transmitted token as clear can be used to mark the token as empty.

The format of the message may vary, but in this case it only includes one address field 40, one store instruction field 42 and one hop instruction

field 44. Therefore, when the message is received and sent on by a processor, the address of that processor in the message is simply replaced by the new target processor address a_2 taken from the hopping address field 32, and the store instruction flag and hopping instruction flag in the incoming message are replaced by the new store and hopping instructions taken from the memory fields 34, 36. This means that the message can be sent along an arbitrarily long route between processors without the size of the message increasing as the length of the route increases, as each time it is put on the pipe it only needs to include the address and delivery instructions for the next processor along the route. Also, because each processor can produce two output messages, one processing output including the result of performing its instruction on received data, and one re-routing or hopping output which re-routes the received data to another processor, this allows a branching of the line of communication through the processor. More generally, the processing output and the hopping output can be sent to a plurality of processors, for example any number of copies of each type of output can be sent to different processors, either being sent out in the same direction or in different directions. The number of output channels on each edge of the processor may be greater than two if required. This supports a wide range of parallel programming idioms which may be supported by a compiler of parallel processing languages.

It will be appreciated that, with the address a_2 stored in the hopping address field 32, and the delivery instruction stored in the delivery instruction field 34, 36, a route can be defined from one of the processors to another via any number of intermediate processors. The route can be arbitrarily complex and arbitrarily long because the route is defined by data stored at the processors on the route, with the token being communicated along the route at any one time carrying data that defines a fixed small part only, typically of one or two steps, of the route. In this

embodiment the chip 10 is programmable, and the address a_2 stored in the hopping address field 32, and the delivery instruction $S_2 h_2$ stored in the delivery instruction field 34, 36, can be set as required. In some programs the route is fixed by the compiler so that any data sent by the first processor in the route to the second, with a hopping instruction to send the data on, will proceed along the route to the final target processor at the end of the route. In other programs the compiler generates a program which can over-write addresses and delivery instructions so as to switch a data flow. This may be done, for example, to divert a data stream when some target data in the stream have been detected. In this way, the string processing functions of an ordinary computer can be implemented in a parallel computer.

In other embodiments, such as in a field programmable gate array (FPGA) system or an application specific integrated circuit (ASIC), while the program is running the route remains fixed so that any data sent by the first processor in the route to the second, with a hopping instruction to send the data on, will proceed along the route to the final target processor at the end of the route. However the chip can be re-programmed by loading a different set of addresses and delivery instructions into the appropriate fields of the processors, so that different programs can be run. In other embodiments the chip may be non-programmable with the stored addresses and delivery instructions being fixed.

In other embodiments, the message can be made more complex. For example it can include two address fields within a routing address field, each address field holding one processor address. The two addresses can then be updated alternately as the message is transmitted from processor to processor along the route. Similarly there may be two delivery instruction fields each associated with one of the address fields which can

be updated alternately. This still keeps the routing part of the message to a fixed length, by replacing a used address and delivery instructions with new addresses and delivery instructions. In this case the re-routing address at the processor may not replace the processor's own address but
5 may replace a previous address which has already been used as a target address. Alternatively the re-routing address may be inserted so as not to form the target address in the immediately re-transmitted message, but so as to be used as a target address in a subsequent hop, the target address for the immediate re-transmission being already present in the message.
10 Therefore the processor may place the message on the pipe for delivery either directly or indirectly to the re-routing address.

Referring to Figure 5, in a further embodiment any one of the processors
12 can be arranged to re-transmit the received message as two separate
15 messages or tokens 66, 68 to two target processors, which may be on the same or different pipes. This provides a further mechanism for the communication line within the system to branch. In order to do this, each processor includes two hopping address/instruction fields 62, 64 which can contain the addresses a_2 , a_3 of two separate target processors, and two
20 associated two-part delivery instructions S_2 , h_2 , S_3 , h_3 for the respective target processor. In this case, if the received message has the hop instruction field flagged, then the processor 12 is arranged to place two output tokens on the pipe, one addressed to each of the target processors identified by the addresses stored in the two address/instruction fields 62,
25 64. It will be appreciated that in this arrangement the size of the received message does not need to change, depending on how many processors it is to be 'hopped' or retransmitted to. Each processor that receives the message only needs the message to flag whether it is to be hopped on or not, and if it is, all of the information required to perform the re-
30 transmission is stored at the processor.

Referring to Figure 6, in another embodiment the processor 12 has a number of internal registers, in this case four 71, 72, 73, 74 are shown. Three of these 71, 72, 73 are used to store the inputs to the instruction which the processor is arranged to perform, and one 74 is arranged to store the result of the instruction. Each register has its own address, a_{1x} where a_1 is the address of the processor 12 and x indicates which register within the processor is being addressed. The processor 12 further comprises processing components 76 arranged to perform the instruction, and a hopping address field 78 arranged to store a hopping address, in the form a_{2x} where a_2 identifies the target processor to which the token is to be forwarded, and x indicates the register within the target processor to which the token is addressed. Each received token 84 includes a target address a_{1x} and data d . On receipt of a token addressed to the processor 12, the processor copies the token into a buffer 80. If the token is addressed to one of the registers 71, 72, 73 then it is copied into that register. If it is addressed to the result register 74, then it is not copied into that register, but is forwarded on to another processor by replacing the target address in the token with the address stored in the hopping address field 78. The result register address of each processor is therefore used as a sink register address. This means that the token can be made to follow a route which can be arbitrarily long and complex, by being forwarded, or hopped, repeatedly with the sink address, in this case the result register address, of each processor on the route being used as the hopping address. When the target address of a non-sink register, such as registers 1 to 3 in the processor of Figure 6, is inserted into the token, the processor at which it is inserted will be the penultimate processor in the route, and the token will be delivered to the final target register at the final processor in the route, and will not be forwarded further. Therefore in this embodiment the full register address, as stored and as inserted in the token, includes one high level part which identifies the processor to which the token is to be sent, and another low level part which determines

both whether the token is to be hopped or stored and, if stored, in which register. This therefore enables the data in the token to be forwarded along a route of any length and complexity, but without being delivered to any of the processors along the route, except the final one.

5

In a modification to the embodiment of Figure 6, each processor has a delivery instruction memory field containing a hopping instruction, which can be a single bit as in the embodiment of Figure 4. In this case each incoming token includes a target register address, a hopping instruction,
10 and data. The hopping instruction determines whether the data is hopped on to another processor or not. The target register address will either be for one of the non-sink registers, in which case the processor is arranged to store the data in the appropriate register, or it will be the sink register address, in which case the data will not be stored. Therefore the default is
15 to store the data in the register the token is addressed to, but if that is the sink register then the data will not be stored. Therefore, in this arrangement also, the full register address, as carried in the token and as stored in each processor for insertion into the token, includes a high level part which identifies the processor to which the token is to be forwarded
20 and a low level part which determines whether the token is to be stored and, if so, in which register. This arrangement therefore allows the data to be forwarded along an arbitrarily complex route with the option of storing the data at any of the processors along the route.

25 It will be appreciated that both the system of Figure 6, and the modification to that system described above, use the presence of a sink register in each processor, to which incoming data will never need to be input. If no such sink register is present, then each processor can still have sink address associated with it, which is equivalent to a register
30 address but is not associated with any specific register. The processor therefore has more register addresses associated with it than it has

registers. Alternatively a system such as that of Figure 4 can be used in which delivery instructions are used with one bit to control storing and one to control hopping.

- 5 Each of the embodiments described above can be modified by the inclusion in the token of a fetch bit (f), or a longer fetch instruction, which instructs the processor whether or not to replace the datum in an outgoing token with the contents of one of the processor's memory locations. In this case each processor is arranged to include a fetch
10 instruction memory field so that the processor can insert the fetch instruction into the token with the target address, and hop and store instructions if appropriate, when forwarding the token along the route.

Similarly any of the embodiments described above can modified so that
15 the token includes a control bit (c) which instructs the processor the token is addressed to whether or not to execute its instruction or program. In this case each processor also has a control instruction memory field containing a control instruction which can be inserted into the token when forwarding the token along the route.

20

It will be appreciated that the embodiments described above are examples of the more general case in which any operation code may be used which forms part of the token and which is arranged to instruct the processor to which it is addressed to perform an associated operation, the operation
25 code to be forwarded to each processor on the route being stored at a previous processor. The operation may be performed on the token, or some data it contains, as it is delivered to a processor along the route, as it is forwarded, or as it is delivered to the final destination of the route.

- 30 The re-routing or 'hopping' methods described above can be used to perform a number of functions within the processing system. The simplest

process is to re-route a token at one of the processors 12 to change the direction in which it is travelling. For example in a simple square processor array a message can in theory be sent from any processor to any other by travelling the required distance in the X direction and then
5 the required distance in the Y direction, with the processor at the point of change of direction performing the re-routing function. Clearly more complex routes can be laid out as required.

Other applications make use of the dual function of storing the received
10 data and generating a re-routing output including the received data. One such application of the re-routing function is to initialize a vector, or linear row of processors to the same datum. This can be achieved by sending a message from one end of the row (or column) with delivery instructions for each processor to store the datum from the message and
15 send the message on to the next processor in the row (or column). A two-dimensional array of processors can be initialized in a similar way, by further generating a re-routing output from each processor in the initial row, with the message being sent a long one row, and each processor in the row being arranged to store the datum, send the message on along the
20 row, and also re-route it in the perpendicular direction down the respective column.

Referring to Figure 7, in another application, one of the processors 12a is designated as a working processor and another 12b is designated as a
25 sentinel processor. The working processor 12a is arranged to perform an instruction to generate a result, and to send the result to another processor 12c, for example for further processing, and also to send the result on to the sentinel processor 12b (whose address is stored at the working processor 12a). The three processors 12a, 12b, 12c are shown as
30 being adjacent to each other but they can be further apart. The sentinel processor 12b is arranged to check the result that it receives and respond

in one of a number of ways depending on the result. For example the sentinel processor 12b may be arranged, if the result meets one or more conditions, to transmit a message back to the working processor 12a to stop it outputting any further results. Alternatively the sentinel processor

5 12b may be arranged to send a message to a further processor if the result meets one or more conditions. The advantage of the use of the embodiments described above in this application is that the result can be stored and sent to the sentinel processor 12b without the need for any message headers to include more than one address and delivery

10 instructions made up of at most two bits.

CLAIMS

1. A processing system comprising a plurality of processors and communication means arranged to carry messages between the processors, wherein each of the processors has an operating instruction memory field
5 arranged to hold stored operating instructions including a re-routing target address, and is arranged to receive a message including operating instructions including a target address, and on receipt of the message to:
 - check the target address in the message to determine whether it corresponds to an address associated with the processor;
 - 10 if the target address in the message does correspond to an address associated with the processor, to check the operating instructions in the message to determine whether the message is to be re-routed; and
 - if the message is to be re-routed to replace operating instructions within the message with the stored operating instructions, and place the
15 message on the communication means for delivery to the re-routing target address.
2. A system according to claim 1 wherein the message includes an operating instruction field holding at least one address including the target
20 address, and if the message is to be re-routed the processor is arranged to replace an address in the operating instruction field with the re-routing target address.
3. A system according to claim 1 or claim 2 wherein the address
25 within the message that is replaced is the address associated with the processor.
4. A system according to any foregoing claim wherein the operating instructions in the message include a re-routing instruction indicating
30 whether the message is to be re-routed at the processor.

5. A processing system according to claim 4 wherein each processor is arranged, if the received message is to be re-routed, to replace a used re-routing instruction in the message with the re-routing instruction from the operating instruction memory field.
- 5
6. A system according to any foregoing claim wherein the operating instructions in the message include a store instruction indicating whether the message is to be stored at the processor associated with the re-routing target address.
- 10
7. A system according to any foregoing claim wherein each processor is arranged to determine from the message whether the message is to be stored, and if it is, to store at least a part of the message.
- 15
8. A system according to any foregoing claim wherein each processor has a plurality of addresses associated with it, one of the addresses being a sink address.
9. A system according to claim 8 wherein each processor is arranged, if the target address in the message is its sink address, to re-route the message.
- 20
10. A system according to claim 8 or claim 9 wherein each processor is arranged, if the target address in the message is one of its addresses other than the sink address, to store at least a part of the message at the target address.
- 25
11. A system according to any foregoing claim wherein each of the processors is further arranged to perform an instruction using a part of the message as an input.
- 30

12. A system according to claim 11 wherein each processor is arranged to produce an output as a result of performing the instruction and to place the output on the communication means for delivery to another processor.
- 5 13. A system according to any foregoing claim wherein the operating instruction field of each processor includes a further address memory field arranged to hold a further re-routing target address and, if the message is to be re-routed, to also place the message on the communication means for delivery to any further re-routing address
10 stored in the further address memory field.
14. A system according to claim 13 wherein each processor has a further re-routing instruction memory field and, if the message is to be re-routed, to insert any further delivery instruction from the further re-
15 routing instruction field into the message to indicate to the processor at the further re-routing address how the message is to be delivered.
15. A system according to claim 4 or claim 5 wherein the re-routing instruction has a fixed number of possible states indicating, independently
20 of the processor at which it is stored and the re-routing address to which it is to be sent, whether the message is to be re-routed at the processor at the re-routing address.
16. A method of communicating messages between processors of a
25 processing system, the method comprising receiving a message at a processor, the message including operating instructions which comprise a target address, and on receipt of the message:
- checking the target address to determine whether it corresponds to an address associated with the processor;

if the target address does correspond to an address associated with the processor, checking the operating instructions to determine whether the message is to be re-routed;

5 if the message is to be re-routed, replacing an address within the message with a re-routing target address to enable re-routing of the message to the re-routing target address.

10 17. A method according to claim 16 wherein the address, within the message, that is replaced, is the address associated with the processor.

18. A method according to claim 17 wherein the operating instructions within the received message include a re-routing instruction.

15 19. A method according to claim 18 including determining whether the message is to be re-routed and, if the message is to be re-routed, replacing the re-routing instruction within the message with a further re-routing instruction.

20 20. A method according to any of claims 16 to 19 wherein the operating instructions comprise a store instruction indicating whether the message is to be stored at the processor.

25 21. A method according to claim 20 including determining whether the message is to be stored and, if it is, storing at least a part of the message.

22. A method according to any of claims 16 to 21 further comprising performing an instruction using a part of the message as an input.

23. A method according to claim 22 including producing an output as a result of performing the instruction and delivering the output to another processor.
- 5 24. A method according to any of claims 16 to 23 further comprising, if the message is to be re-routed inserting a further re-routing target address into a copy of the message to enable delivery of the message to the further re-routing target address.
- 10 25. A method according to claim 24 further comprising, if the message is to be re-routed, inserting a further re-routing instruction into the copy of the message to indicate to the processor at the further re-routing target address whether the message is to be re-routed at that processor.
- 15 26. A method according to claim 18 or claim 19 wherein the re-routing instruction has a fixed number of possible states indicating, independently of the processor at which it is stored and the re-routing address to which it is to be sent, whether the message is to be re-routed at the processor at the re-routing address.
- 20 27. A method of setting up a processing system according to any of claims 1 to 15, the method comprising loading a respective re-routing target address into the re-routing address memory field of a plurality of the processors thereby to define a route along which data can be
25 communicated between a plurality of the processors,
28. A processing system according to any of claims 1 to 15 wherein the processors are each arranged to perform an instruction once in each of a series of operating cycles, and the communication means comprises a
30 plurality of pipes each arranged to transmit data in a respective direction, and each processor comprises a plurality of output means each arranged

to output data onto a respective one of the pipes in the form of tokens, at least one of the output means being arranged to output two tokens onto the same pipe in one operating cycle of the processor.

5 29. A processing system comprising a plurality of processors each arranged to perform an instruction once in each of a series of operating cycles, and communication means arranged to carry messages between the processors, wherein the communication means comprises a plurality of
10 pipes each arranged to transmit data in a respective direction, and each processor comprises a plurality of output means each arranged to output data onto a respective one of the pipes in the form of tokens, at least one of the output means being arranged to output two tokens onto the same pipe in one operating cycle of the processor.

15 30. A processing system according to claim 28 or claim 29 wherein each pipe comprises a plurality of pipe frames and the communication means is arranged to transfer tokens between adjacent pipe frames, to carry the messages, at a pipe frequency which is higher than that of the operating cycle of the processors.

20 31. A processing system according to claim 30 wherein the processor is arranged to output two tokens to the same pipe frame at different times within the operating cycle of the processor and the communication means is arranged to transmit the first of the tokens to another pipe frame before
25 the second token is output onto the pipe.

32. A processing system substantially as hereinbefore described with reference to any one or more of the accompanying drawings.

33. A method of communicating messages between processors substantially as hereinbefore described with reference to any one or more of the accompanying drawings.





Fig. 3




Fig. 4




Fig. 5




Fig. 6




Fig. 7

