

(43) Date of A Publication 18.10.2000

(21) Application No 9908199.4

(22) Date of Filing 09.04.1999

(71) Applicant(s)
Pixelfusion Limited
(Incorporated in the United Kingdom)
2440 The Quadrant, Aztec West, Almondsbury,
BRISTOL, BS32 4AQ, United Kingdom

(72) Inventor(s)
Trey Greer
Phil Atkin

(74) Agent and/or Address for Service
Haseltine Lake & Co
Imperial House, 15-19 Kingsway, LONDON,
WC2B 6UD, United Kingdom

(51) INT CL⁷
G06F 13/16

(52) UK CL (Edition R)
G4A AMB

(56) Documents Cited
WO 96/41250 A2 US 4538226 A

(58) Field of Search
UK CL (Edition Q) **G4A AMB**
INT CL⁶ **G06F 12/00 12/08 12/14 13/16**
ONLINE: COMPUTER EPODOC INSPEC JAPIO WPI

(54) Abstract Title
Accessing a memory unit in a parallel data processing system

(57) A method of accessing a memory unit in a data processing apparatus including a plurality of processing units, particularly a SIMD (single instruction multiple data) array, comprising selecting a processing element requiring access to an address, determining which other elements require access to that address, and transmitting the required data to or from each element with the appropriate transaction identification stored. Different processing elements may, in a write action, store data items in respective regions of the target memory address.




FIG. 1


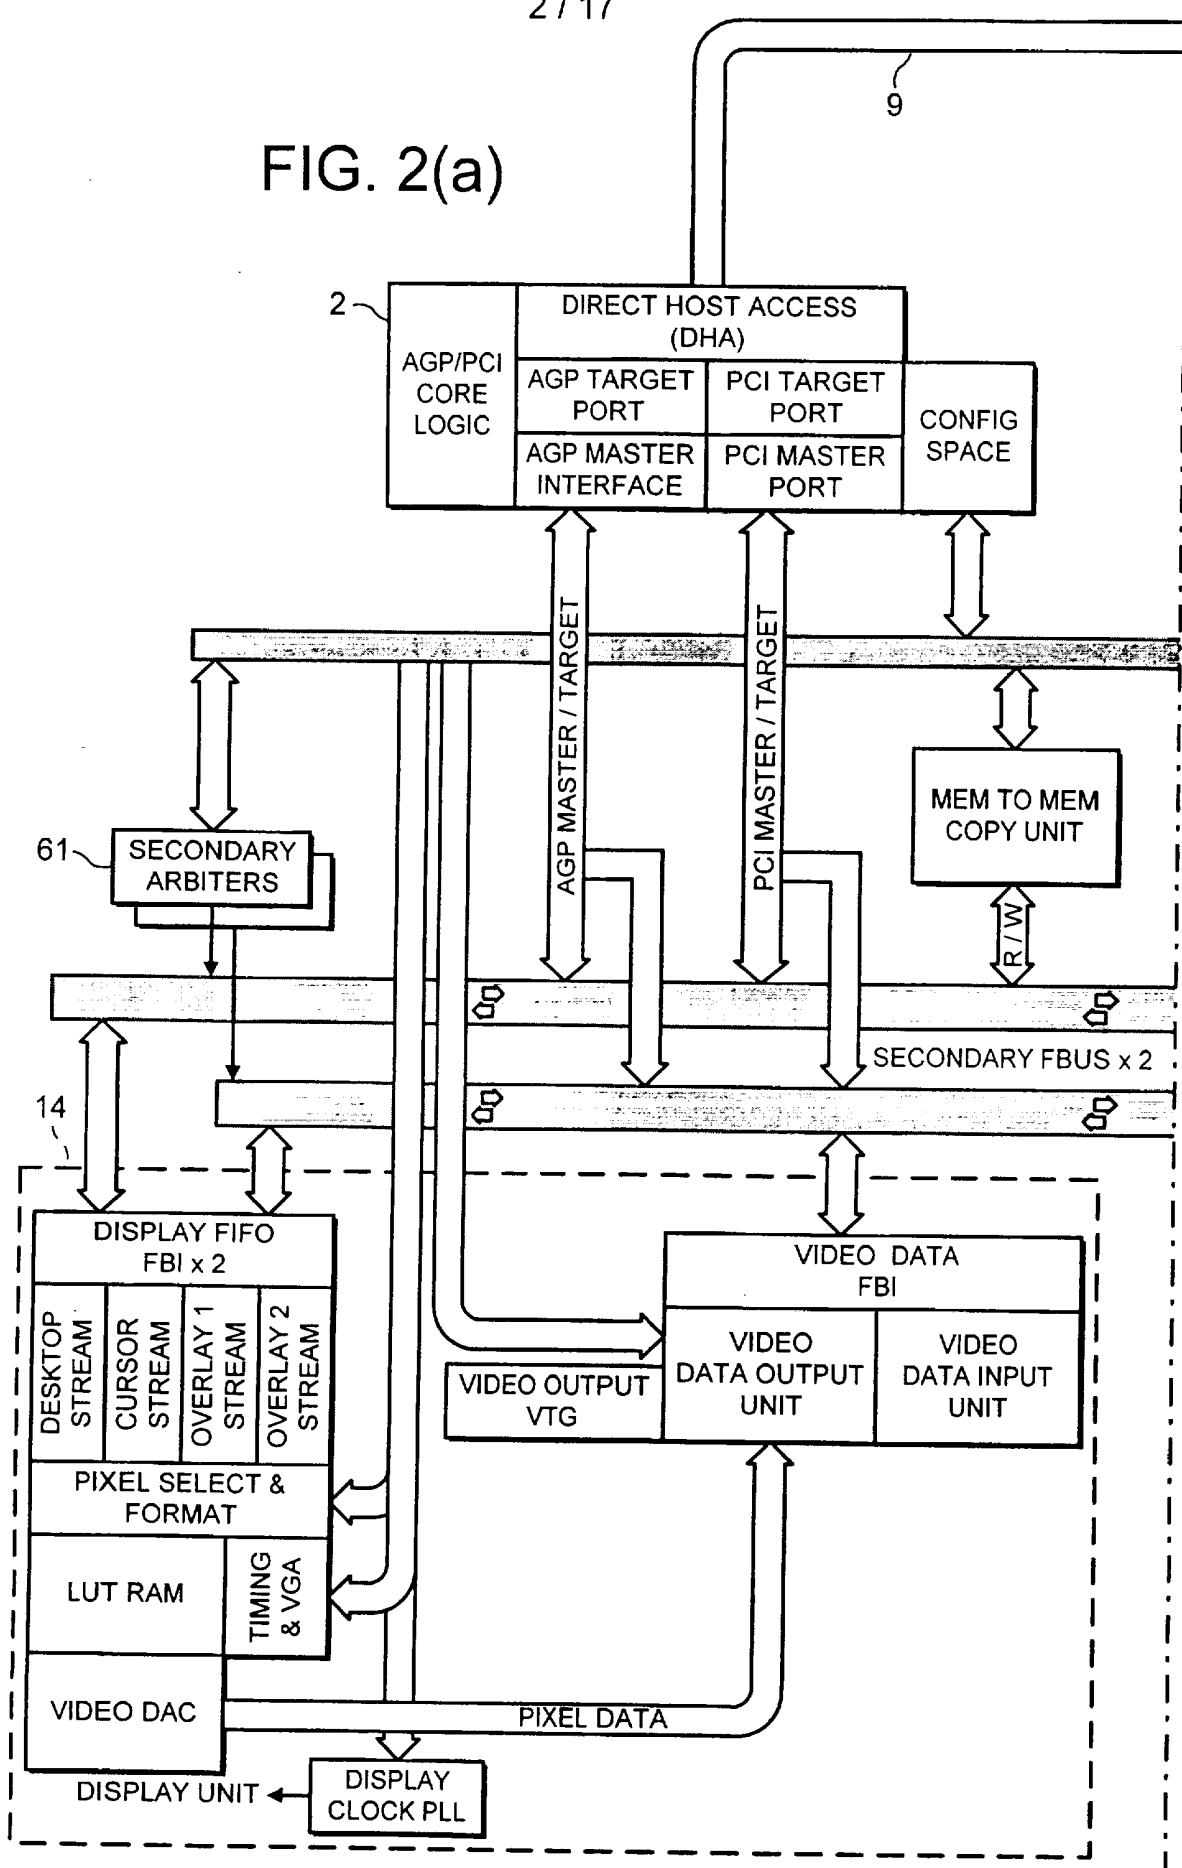


FIG. 1

FIG. 2(a)



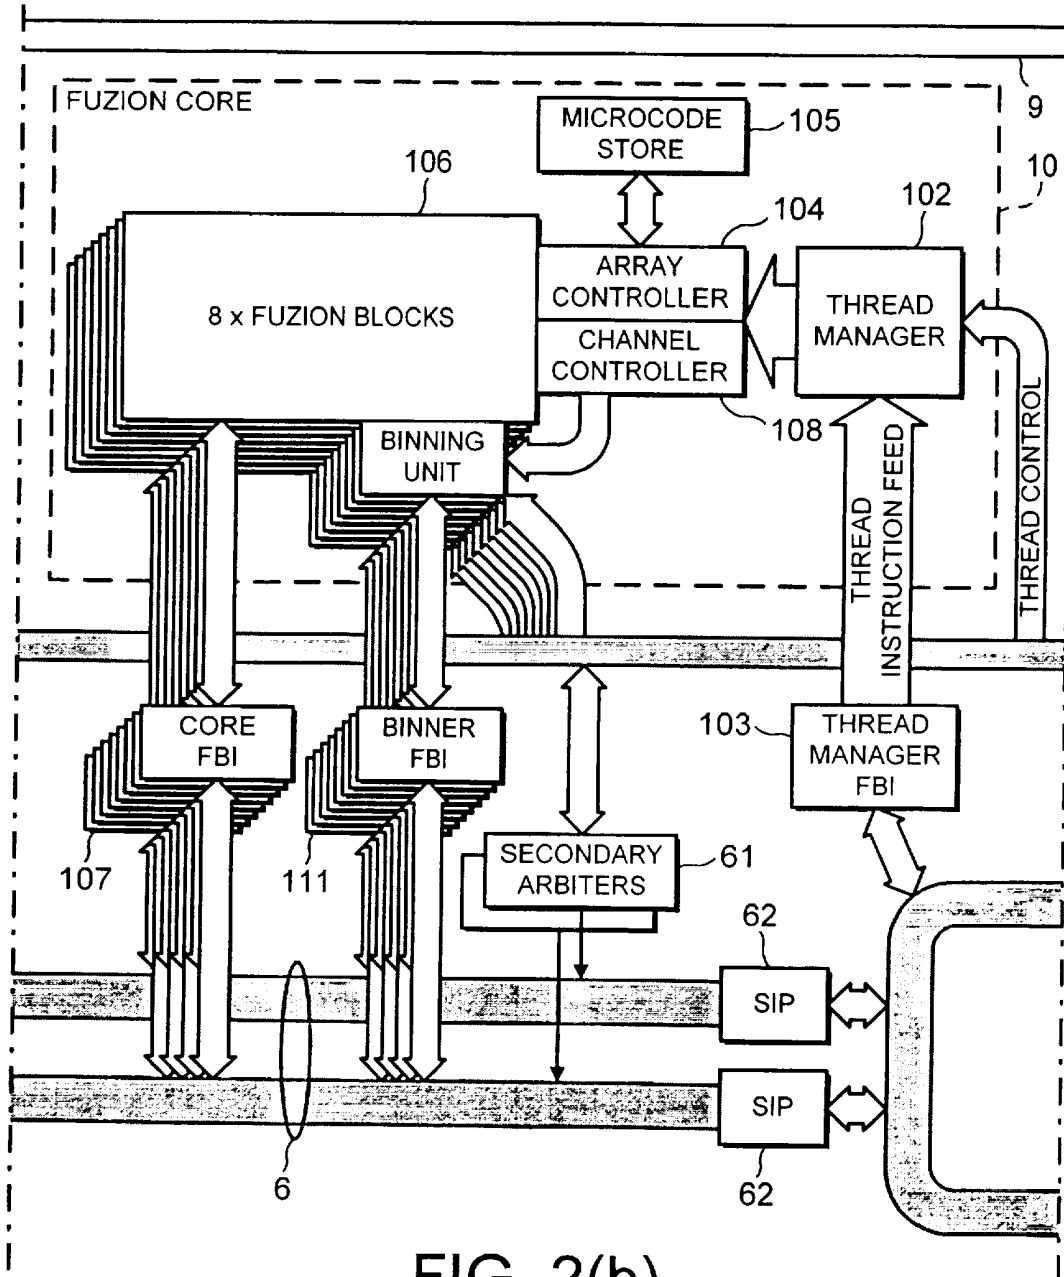


FIG. 2(b)




FIG. 2(c)

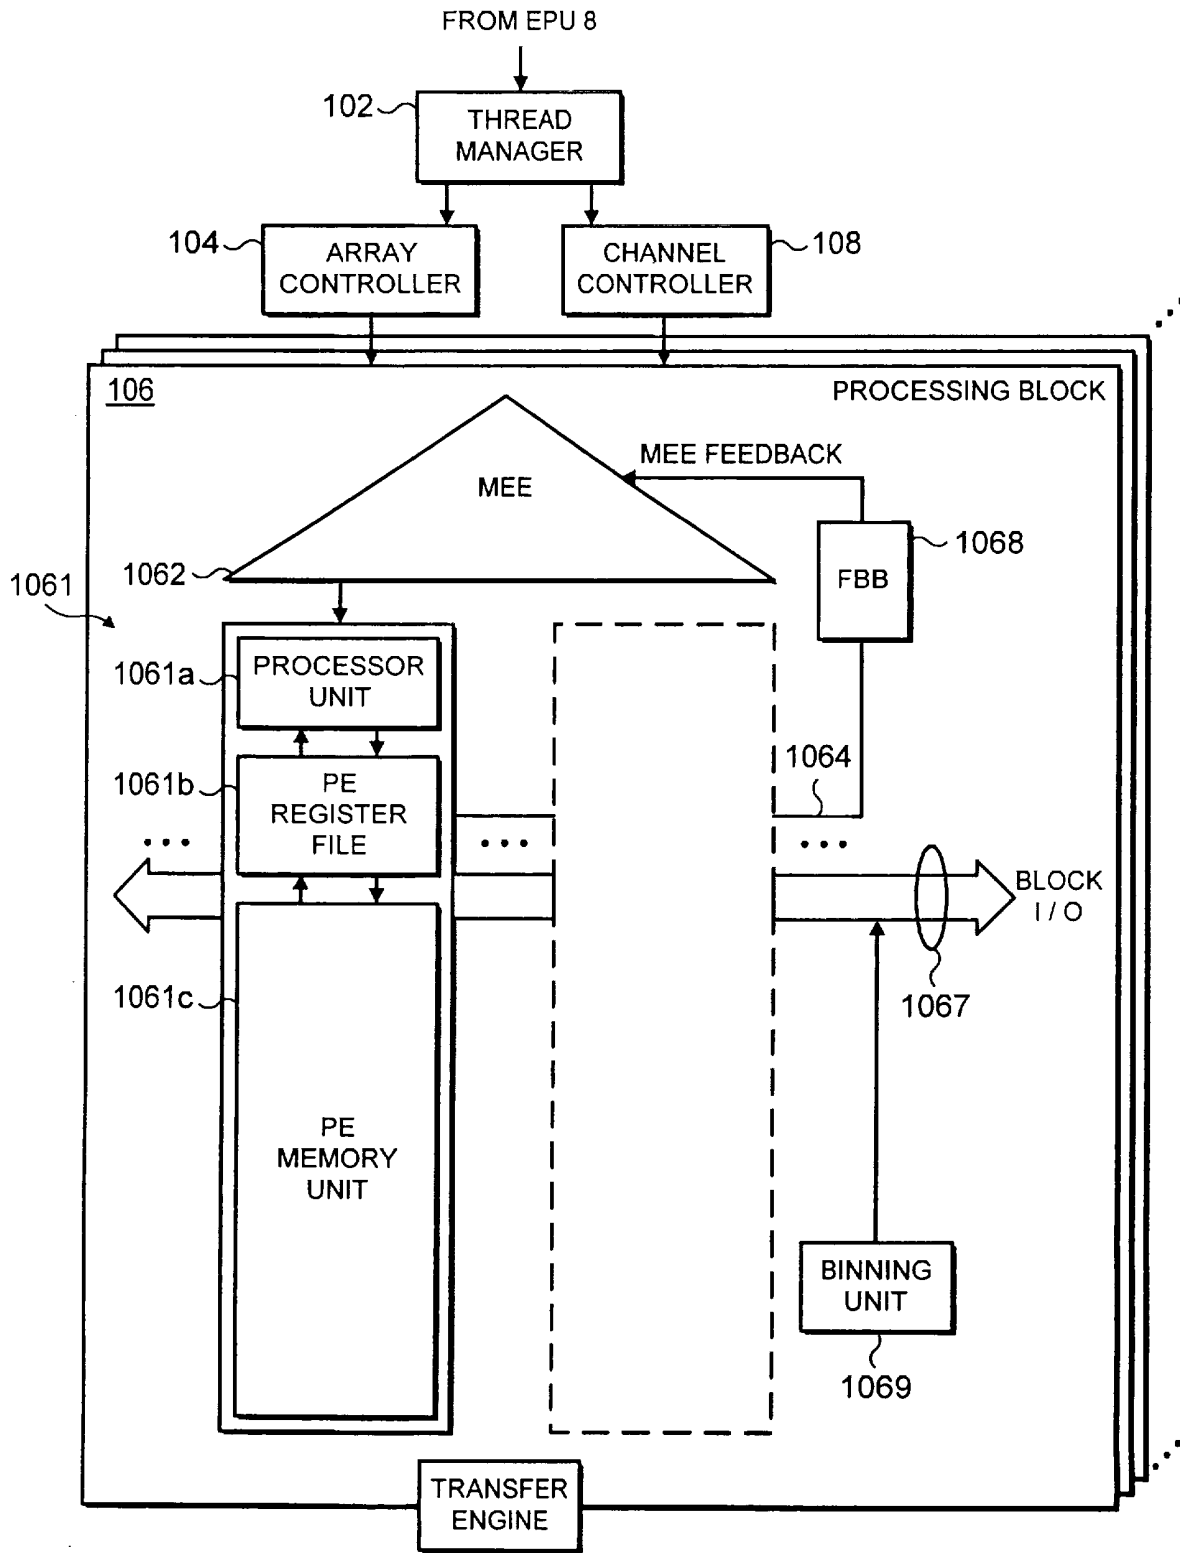



FIG. 3

FIG. 4



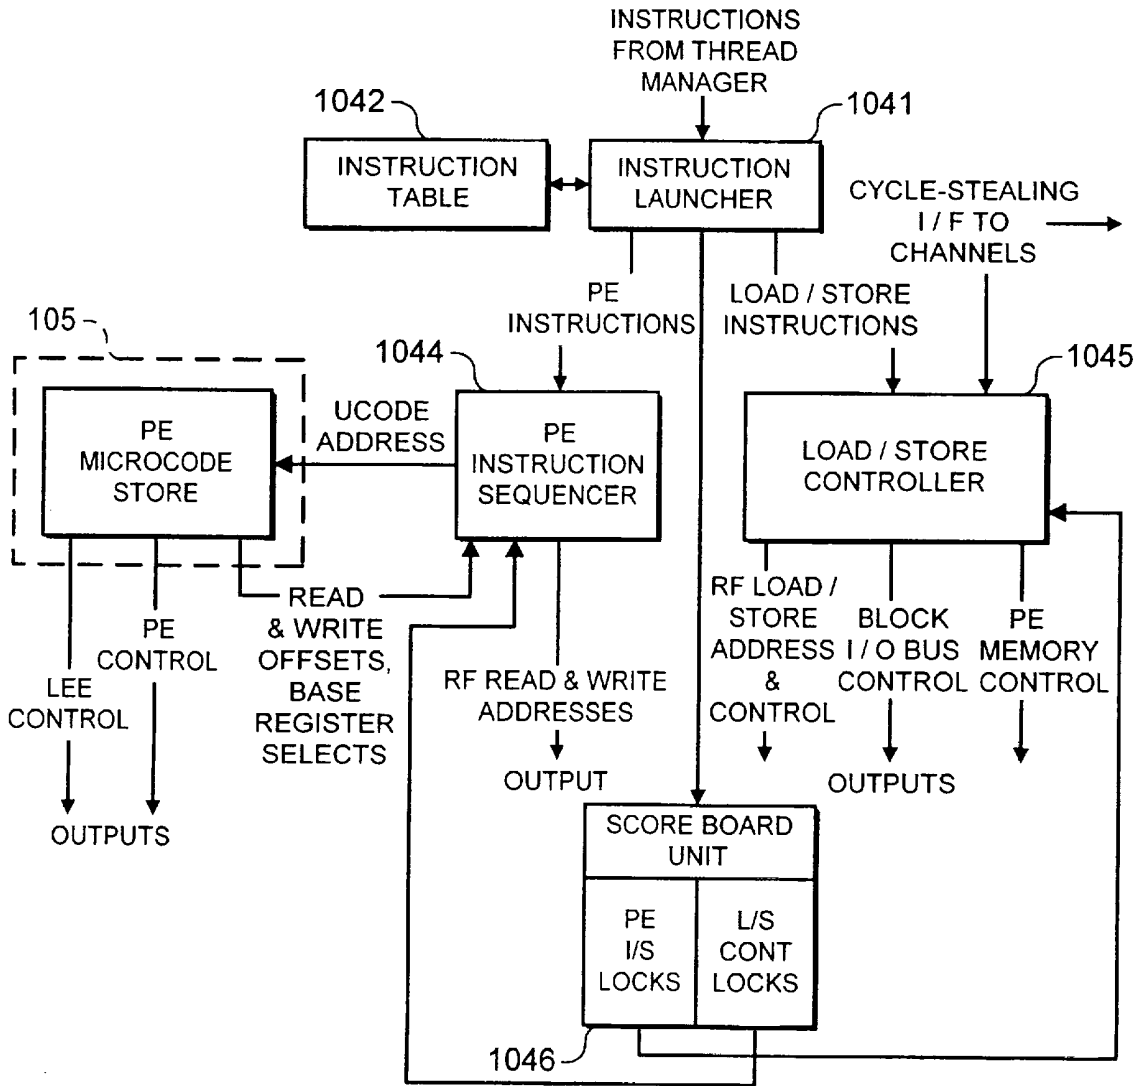


FIG. 5

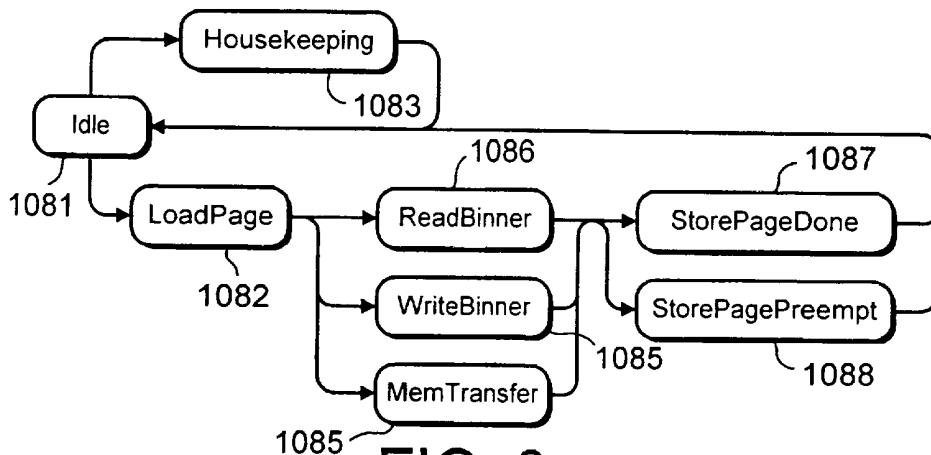


FIG. 6

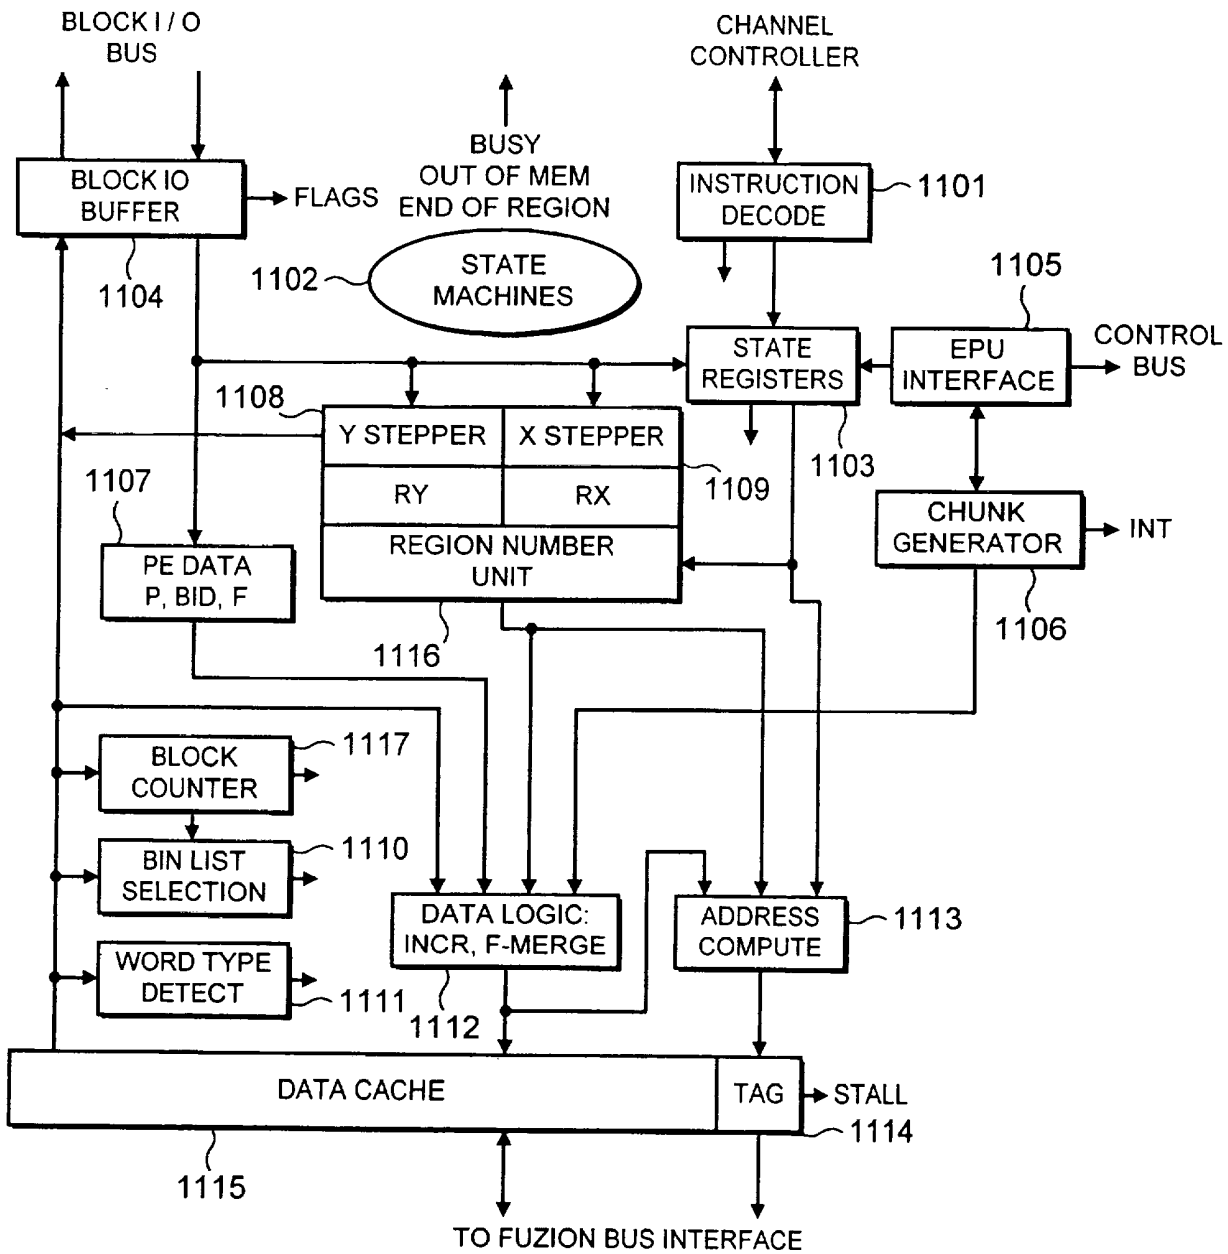


FIG. 7

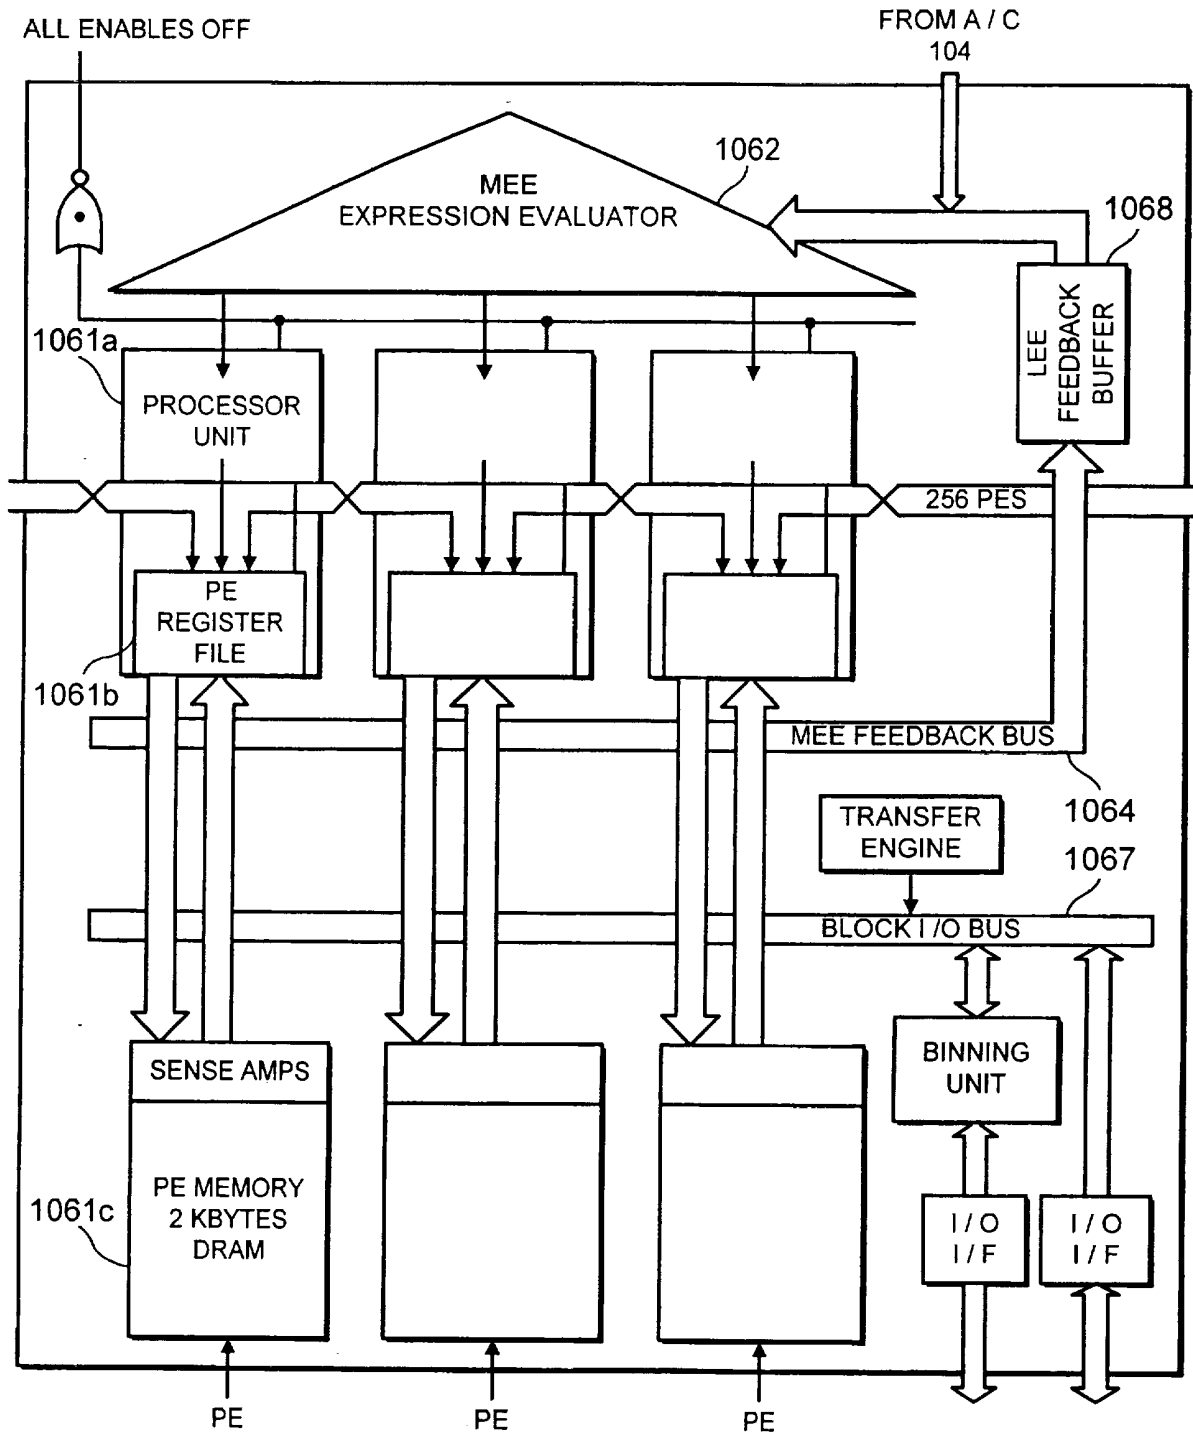


FIG. 8

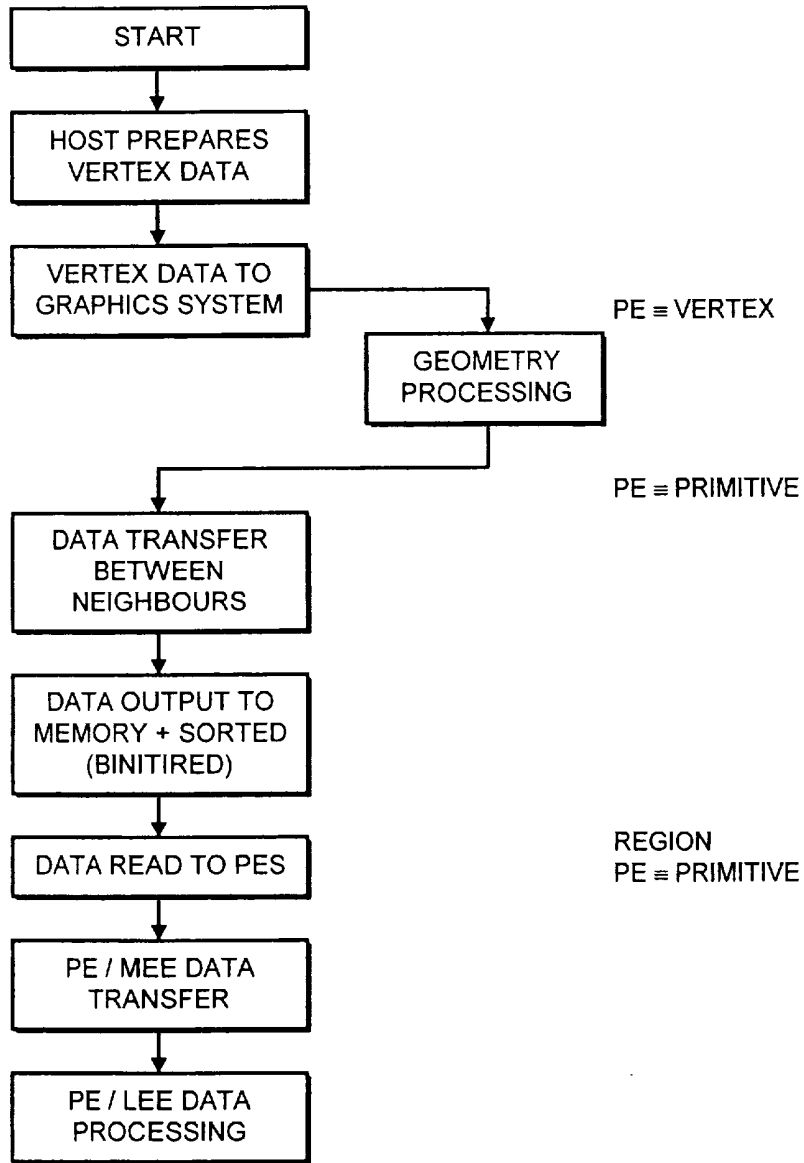


FIG. 9




FIG. 10(b)




FIG. 11




FIG. 12

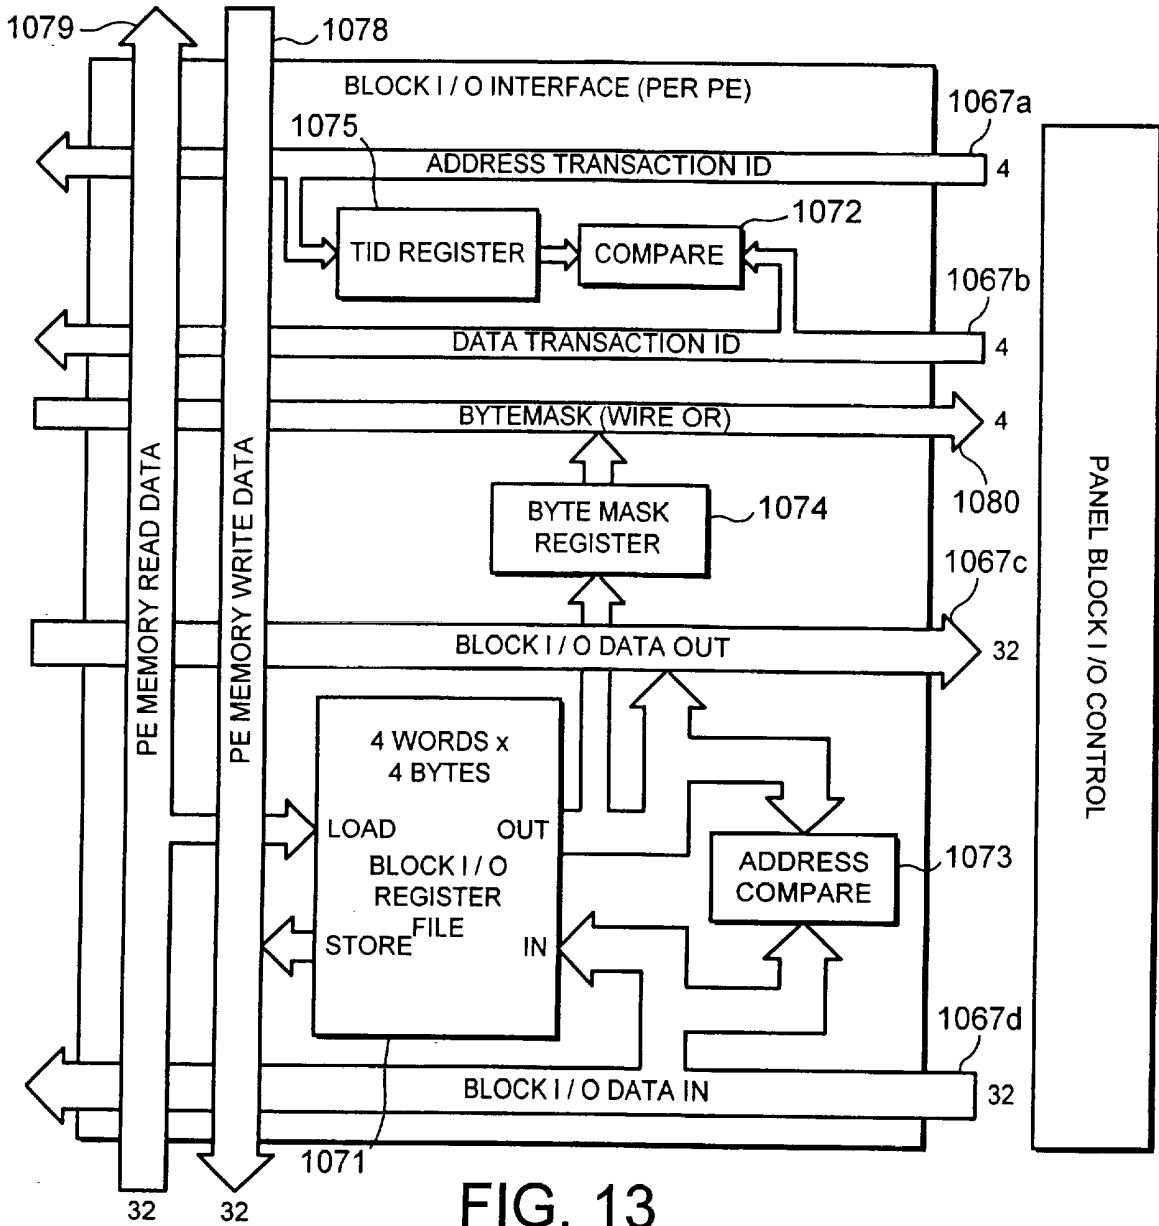


FIG. 13

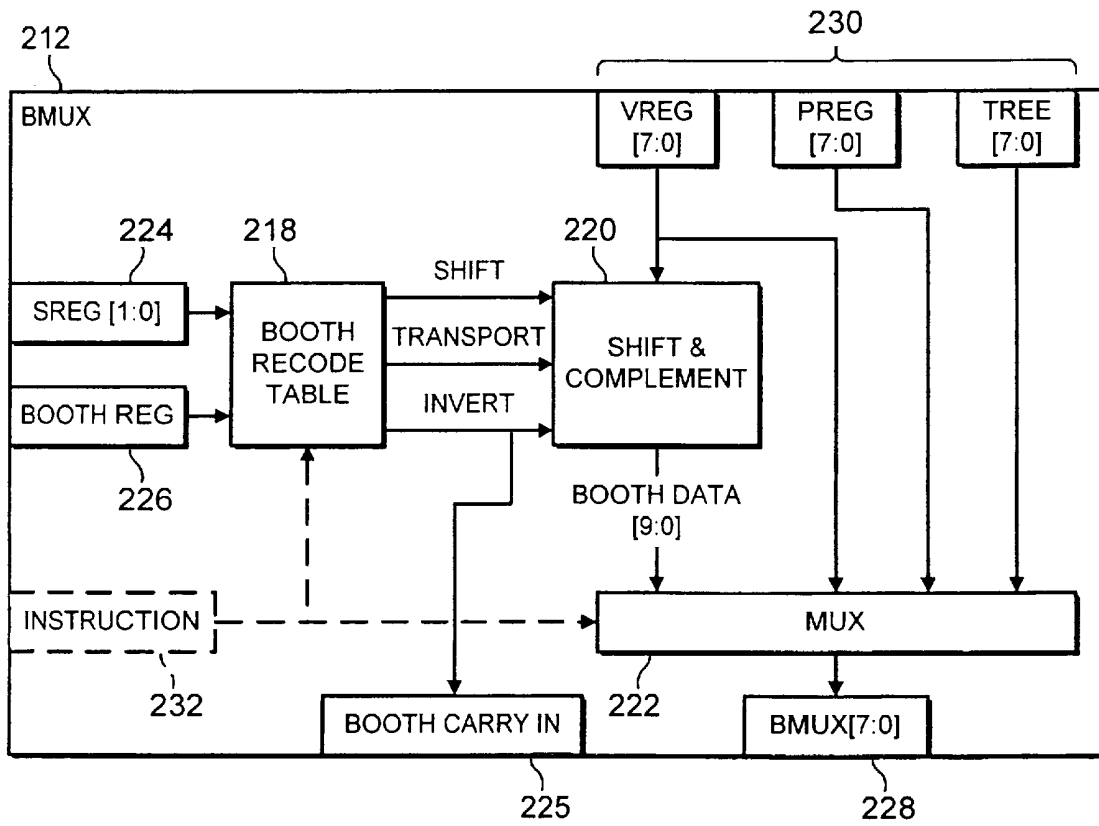


FIG. 14




FIG. 15

214

PARALLEL DATA PROCESSING SYSTEMS

The present invention relates to parallel data
processing systems, and in particular to SIMD (single
5 instruction multiple data) data processing apparatus.

BACKGROUND OF THE INVENTION

Increasingly, data processing systems are required to
10 process large amounts of data. In addition, users of
such systems are demanding that the speed of data
processing is increased. One particular example of the
need for high speed processing of massive amounts of
data is in the computer graphics field. In computer
15 graphics, large amounts of data are produced that
relate to, for example, geometry, texture, and colour
of objects and shapes to be displayed on a screen.
Users of computer graphics are increasingly demanding
more lifelike and faster graphical displays which
20 increases the amount of data to be processed and
increases the speed at which the data must be
processed.

A previously proposed processing architecture for
25 processing large amounts of data in a computer system
uses a Single Instruction Multiple Data (SIMD) array of
processing elements. In such an array all of the
processing elements receive the same instruction
stream, but operate on different respective data items.
30 Such an architecture can thereby process data in
parallel, but without the need to produce parallel
instruction streams. This can be an efficient and
relatively simple way of obtaining good performance
from a parallel processing machine.

35

However, the provision of a large number of processing elements results in a large number of memory accesses being required for processing the data stored in memory. In known systems, such as that described in US Patent No.5,539,898, one of the processing elements can be selected and the address to which it requires access is supplied on an address bus to the memory. The other processing elements check whether the address on the address bus is the address to which they require access, and if so access the memory at the same time as the selected processing element.

However, the performance of the system is severely limited by the need for a memory access to be completed, albeit for a number of processing elements in the same access, before the next memory access is able to be made.

Accordingly, it is desirable to provide an architecture that can overcome this disadvantage.

SUMMARY OF THE INVENTION

According to one aspect of the present invention, there is provided a method of controlling data read access to memory in a data processing apparatus including a plurality of processing elements, the method comprising:

selecting a processing element that requires access to the memory,
retrieving a target address from the selected processing element,
transmitting the target address to the plurality of processing elements,

transmitting transaction identification
information to the processing elements, which
information identifies the target address access
operation concerned,
5 storing the transaction identification information
in the or each processing element that requires
access to the target address,
transmitting data obtained from the target address
together with the transaction identification
10 information to the plurality of processing
elements, and
storing the obtained data in the or each
processing element in which the transaction
identification information is stored.

15

According to another aspect of the present invention,
there is provided a method of retrieving a data item
from a memory unit in a data processing apparatus which
includes a memory unit in which data items are stored
20 at addresses therein, and a plurality of processing
elements which have access to the memory unit, the
method comprising:

25

for each processing element in the array which
requires access to the memory unit, setting an
access indicator to indicate that the processing
element concerned requires access to the memory
unit, and storing a target address of the memory
unit to which such access is required;

30

selecting one of the processing elements having
the access indicator set, and retrieving the
stored target address from that selected
processing element;

35

transmitting the retrieved target address and transaction identification information to the processing elements in the array;

5 for each processing element having the access indicator set, comparing the transmitted target address with the stored target address, and if the stored and transmitted target addresses are equivalent, clearing the access indicator and
10 storing the transaction identification information;

retrieving at least one data item stored at the transmitted target address in the memory unit;

15 transmitting the or each retrieved data item and associated transaction identification information to the processing elements in the array; and

20 for each processing element having stored transaction identification information, comparing the stored transaction identification information with the transmitted transaction identification information, and if the stored transaction
25 information is equivalent to the transmitted transaction information, receiving the or each retrieved data item.

30 According to another aspect of the present invention, there is provided method of writing data items to a memory unit in a data processing apparatus including the memory unit in which data items are stored at addresses therein, and a plurality of processing
35 elements which have access to the memory unit, the method comprising:

5 for each processing element in the array which
requires access to the memory unit, setting an
access indicator to indicate that the processing
element concerned requires access to the memory
unit, and storing a target address of the memory
unit to which such access is required;

10 selecting one of the processing elements having
the access indicator set, and retrieving the
stored target address from that selected
processing element;

15 transmitting the retrieved target address and
transaction identification information to all the
processing elements in the array;

20 for each processing element having the access
indicator set, comparing the transmitted target
address with the stored target address, and if the
stored and transmitted target addresses are
equivalent, clearing the access indicator and
storing the transaction identification
information;

25 transmitting transaction identification
information to the processing elements in the
array;

30 for each processing element having stored
transaction identification information, comparing
the stored transaction identification information
with the transmitted transaction identification
information, and if the stored transaction
information is equivalent to the transmitted
35 transaction information, transmitting at least one

data item to be stored in the memory unit at the target address; and

5 storing the or each transmitted data item at the target address in the memory unit.

Various further aspects of the present invention are exemplified by the attached claims.

10 BRIEF DESCRIPTION OF THE DRAWINGS

Figure 1 is a block diagram illustrating a graphics data processing system;
Figure 2 is a more detailed block diagram illustrating the graphics data processing system of Figure 1;
15 Figure 3 is a block diagram of a processing core of the system of Figure 2;
Figure 4 is a block diagram of a thread manager of the system of Figure 3;
20 Figure 5 is a block diagram of a array controller of the system of Figure 3;
Figure 6 is a block diagram of an instruction issue state machine of the channel controller of Figure 3;
Figure 7 is a block diagram of a binning unit of the system of Figure 3;
25 Figure 8 is a block diagram of a processing block of the system of Figure 3;
Figure 9 is a flowchart illustrating data processing using the system of Figures 1 to 8;
30 Figure 10 is a more detailed block diagram of a thread processor of the thread manager of Figure 4;
Figure 11 is a block diagram of a processor unit of the processing block of Figure 8;
35 Figure 12 is a block diagram illustrating a processing element interface;

Figure 13 is a block diagram illustrating a block I/O interface;

Figure 14 is a block diagram of part of the processor unit of Figure 11; and

5 Figure 15 is a block diagram of another part of the processor unit of Figure 11.

DESCRIPTION OF THE PREFERRED EMBODIMENT

10 The data processing system described below is a graphics data processing system for producing graphics images for display on a screen. However, this embodiment is purely exemplary, and it will be readily apparent that the techniques and architecture described
15 here for processing graphical data are equally applicable to other data types, such as video data. The system is of course applicable to other signal and/or data processing techniques and systems.

20 An overview of the system will be given, followed by brief descriptions of the various functional units of the system. A graphics processing method will then be described by way of example, followed by detailed description of the functional units.

25

OVERVIEW

Figure 1 is a system level block diagram illustrating a graphics data processing system 3. The system 3
30 interfaces with a host system (not shown), such as a personal computer or workstation, via an interface 2. Such a system can be provided with an embedded processor unit (EPU) for control purposes. For example, the specific graphics system 3 includes an
35 embedded processing unit (EPU) 8 for controlling the

overall function of the graphics processor and for
interfacing with the host system. The system includes
a processing core 10 which processes the graphical data
for output to the display screen via a video output
5 interface 14. Local memory 12 is provided for the
graphics system 3.

Such a data processing can be connected for operation
to a host system or could provide a stand alone
10 processing system, without the need for a specific host
system. Examples of such application include a "set
top box" for receiving and decoding digital television
and internet signals.

15 Figure 2 illustrates the graphics processing system in
more detail. In one particular example, the graphics
system connects to the host system via an advanced
graphics port (AGP) or PCI interface 2. The PCI
interface and AGP 2 are well known.

20 The host system can be any type of computer system,
for example, a PC 99 specification personal computer or
a workstation.

25 The AGP 2 provides a high bandwidth path from the
graphics system to host system memory. This allows
large texture databases to be held in the host system
memory, which is generally larger than local memory
associated with the graphics system. The AGP also
30 provides a mechanism for mapping memory between a
linear address space on the graphics system and a
number of potentially scattered memory blocks in the
host system memory. This mechanism is performed by a
graphics address re-mapping table (GART) as is well
35 known.

The graphics system described below is preferably implemented as a single integrated circuit which provides all of the functions shown in Figure 1. However, it will be readily apparent that the system
5 may be provided as separate circuit card carrying several different components, or as a separate chipset provided on the motherboard of the host, or integrated with the host central processing unit (CPU), or in any suitable combination of these and other
10 implementations.

The graphics system includes several functional units which are connected to one another for the transfer of data by way of a dedicated bus system. The bus system
15 preferably includes a primary bus 4 and a secondary bus 6. The primary bus is used for connection of latency intolerant devices, and the secondary bus is used for connection of latency tolerant devices. The bus architecture is preferably as described in detail in
20 the Applicant's co-pending UK patent applications, particularly GB 9820430.8. It will be readily appreciated that any number of primary and secondary buses can be provided in the bus architecture in the system. The specific system shown in Figure 2 includes
25 two secondary buses.

Referring mainly to Figure 2, access to the primary bus 4 is controlled by a primary arbiter 41, and access to the secondary buses 6 by a pair of secondary arbiters
30 61. Preferably, all data transfers are in packets of 32 bytes each. The secondary buses 6 are connected with the primary bus 4 by way of respective interface units (SIP) 62.

35 An auxiliary control bus 7 is provided in order to

enable control signals to be communicated to the various units in the system.

5 The AGP/PCI interface is connected to the graphics system by way of the secondary buses 6. This interface can be connected to any selection of the secondary buses, in the example shown, to both secondary buses 6. The graphics systems also includes an embedded processing unit (EPU) 8 which is used to control operation of the graphics system and to communicate with the host system. The host system has direct access to the EPU 8 by way of a direct host access interface 9 in the AGP/PCI 2. The EPU is connected to the primary bus 4 by way of a bus interface unit (EPU FBI) 90.

10

15

Also connected to the primary bus is a local memory system 12. The local memory system 12 includes a number, in this example four, of memory interface units 121 which are used to communicate with the local memory itself. The local memory is used to store various information for use by the graphics system.

20

The system also includes a video interface unit 14 which comprises the hardware needed to interface the graphics system to the display screen (not shown), and other devices for exchange of data which may include video data. The video interface unit is connected to the secondary buses 6, via bus interface units (FBI).

25

30 The graphics processing capability of the system is provided by a processing core 10. The core 10 is connected to the secondary buses 6 for the transfer of data, and to the primary bus 4 for the transfer of instructions. As will be explained in more detail

35

below, the secondary bus connections are made by a core
bus interface (Core FBI) 107, and a binner bus
interface (Binner FBI) 111, and the primary bus
connection is made by a thread manager bus interface
5 (Thread Manager FBI) 103.

As will be explained in greater detail below, the
processing core 10 includes a number of control units:
thread manager 102, array controller 104, channel
10 controller 108, a binning unit 1069 per block and a
microcode store 105. These control units control the
operation of a number of processing blocks 106 which
perform the graphics processing itself.

15 In the example shown in Figure 2, the processing core
10 is provided with eight processing blocks 106. It
will be readily appreciated that any number of
processing blocks can be provided in a graphics system
using this architecture.

20 PROCESSING CORE

Figure 3 shows the processing core in more detail. The
thread manager 102 is connected to receive control
25 signals from the EPU 8. The control signals inform the
thread manager as to when instructions are to be
fetched and where the instructions are to be found.
The thread manager 102 is connected to provide these
instructions to the array controller 104 and to the
30 channel controller 108. The array and channel
controllers 104 and 108 are connected to transfer
control signals to the processing blocks 106 dependent
upon the received instructions.

35 Each processing block 106 comprises an array 1061 of

processor elements (PEs) and a mathematical expression evaluator (MEE) 1062. As will be described in more detail below, a path 1064 for MEE coefficient feedback is provided from the PE memory, as is an input/output channel 1067. Each processing block includes a binning unit 1069 unit 1068 and a transfer engine 1069 for controlling data transfers to and from the input/output channel under instruction from the channel controller 108.

10 The array 1061 of processor elements provides a single instruction multiple data (SIMD) processing structure. Each PE in the array 1061 is supplied with the same instruction, which is used to process data specific to
15 the PE concerned.

Each processing element (PE) 1061 includes a processor unit 1061a for carrying out the instructions received from the array controller, a PE memory unit 1061c for storing data for use by the processor unit 1061a, and a
20 PE register file 1061b through which data is transferred between the processor unit 1061a and the PE memory unit 1061c. The PE register file 1061b is also used by the processor unit 1061a for temporarily
25 storing data that is being processed by the processor unit 1061a.

The provision of a large number of processor elements can result in a large die size for the manufacture of
30 the device in a silicon device. Accordingly, it is desirable to reduce the effect of a defective area on the device. Therefore, the system is preferably provided with redundant PEs, so that if one die area is faulty, another can be used in its place.

35

In particular, for a group of processing elements used for processing data, additional redundant processing elements can be manufactured. In one particular example, the processing elements are provided in
5 "panels" of 32 PEs. For each panel a redundant PE is provided, so that a defect in one of the PEs of the panel can be overcome by using the redundant PE for processing of data. This will be described in more detail below.

10

THREAD MANAGER

The array of processing elements is controlled to carry out a series of instructions in an instruction stream.
15 Such instruction streams for the processing blocks 106 are known as "threads". Each thread works cooperatively with other threads to perform a task or tasks. The term "multithreading" refers to the use of several threads to perform a single task, whereas the
20 term "multitasking" refers to the use of several threads to perform multiple tasks simultaneously. It is the thread manager 102 which manages these instruction streams or threads.

25

There are several reasons for providing multiple threads in such a data processing architecture. The processing element array can be kept active, by processing another thread when the current active thread is halted. The threads can be assigned to any
30 task as required. For example, by assigning a plurality of threads for handling data I/O operations for transferring data to and from memory, these operations can be performed more efficiently, by overlapping I/O operations with processing operations.
35 The latency of the memory I/O operations can

effectively be masked from the system by the use of different threads.

5 In addition, the system can have a faster response time to external events. Assigning particular threads to wait on different external events, so that when an event happens, it can be handled immediately.

10 The thread manager 102 is shown in more detail in Figure 4, and comprises a cache memory unit 1024 for storing instructions fetched for each thread. The cache unit 1024 could be replaced by a series of first-in-first-out (FIFO) buffers, one per thread. The thread manager also includes an instruction fetch unit 15 1023, a thread scheduler 1025, thread processors 1026, a semaphore controller 1028 and a status block 1030.

20 Instructions for a thread are fetched from local memory or the EPU 8 by the fetch unit 1023, and supplied to the cache memory 1024 via connecting logic.

25 The threads are assigned priorities relative to one another. Of course, although the example described here has eight threads, any number of threads can be controlled in this manner. At any particular moment in time, each thread may be assigned to any one of a number of tasks. For example, thread zero may be assigned for general system control, thread 1 assigned to execute 2D (two dimensional) activities, and threads 30 2 to 7 assigned to executing 3D activities (such as calculating vertices, primitives or rastering).

35 In the example shown in Figure 4, the thread manager includes one thread processor 1026 for each thread. The thread processors 1026 control the issuance of core

instructions from the thread manager so as to maintain
processing of simultaneously active program threads, so
that each the processing blocks 106 can be active for
as much time as possible. In this particular example
5 the same instruction stream is supplied to all of the
processing blocks in the system.

It will be appreciated that the number of threads could
exceed the number of thread processors, so that each
10 thread processor handles control of more than one
thread. However, providing a thread processor for each
thread reduces the need for context switching when
changing the active thread, thereby reducing memory
accesses and hence increasing the speed of operation.

15 The semaphore controller 1028 operates to synchronise
the threads with one other.

Within the thread manager 102, the status block 1030
20 receives status information 1036 from each of the
threads. The status information is transferred to the
thread scheduler 1025 by the status block 1030. The
status information is used by the thread scheduler 1025
to determine which thread should be active at any one
25 time.

Core instructions 1032 issued by the thread manager 102
are sent to the array controller 104 and the channel
controller 108 (figure 3).

30
ARRAY CONTROLLER

The array controller 104 directs the operation of the
processing block 106, and is shown in greater detail in
35 Figure 5.

The array controller 104 comprises an instruction launcher 1041, connected to receive instructions from the thread manager. The instruction launcher 1041 indexes an instruction table 1042, which provides further specific instruction information to the instruction launcher.

On the basis of the further instruction information, the instruction launcher directs instruction information to either a PE instruction sequencer 1044 or a load/store controller 1045. The PE instruction sequencer receives instruction information relating to data processing, and the load/store controller receives information relating to data transfer operations.

The PE instruction sequencer 1044 uses received instruction information to index a PE microcode store 105, for transferring PE microcode instructions to the PEs in the processing array.

The array controller also includes a scoreboard unit 1046 which is used to store information regarding the use of PE registers by particular active instructions. The score board unit 1046 is functionally divided so as to provide information regarding the use of registers by instructions transmitted by the PE instruction sequencer 1044 and the load/store controller 1045 respectively.

In general terms, the PE instruction sequencer 1044 handles instructions that involve data processing in the processor unit 1061a. The load/store controller 1045, on the other hand, handles instructions that involve data transfer between the registers of the processor unit 1061a and the PE memory unit 1061c. The

load/store controller 1045 will be described in greater detail later.

5 The instruction launcher 1041 and the score board unit 1046 maintain the appearance of serial instruction execution whilst achieving parallel operation between the PE instruction sequencer 1044 and the load/store controller 1045.

10 The remaining core instructions 1032 issued from the thread manager 102 are fed to the channel controller 108. This controls transfer of data between the PE memory units and external memory (either local memory or system memory in AGP or PCI space).

15

CHANNEL CONTROLLER

The channel controller 108 operates asynchronously with respect to the execution of instructions by the array controller 104. This allows computation and external I/O to be performed simultaneously and overlapped as much as possible. Computation (PE) operations are synchronised with I/O operations by means of semaphores in the thread manager, as will be explained in more detail below.

25

The channel controller 108 also controls the binning units 1068 which are associated with respective processing blocks 106. This is accomplished by way of channel controller instructions.

30

Figure 6 shows the channel controller's instruction issue state machine, which lies at the heart of the channel controller's operation, and which will be described in greater detail later.

35

Each binning unit 1069 (Figure 3) is connected to the I/O channels of its associated processing block 106. The purpose of the binning unit 1069 is to sort primitive data by region, since the data is generally not provided by the host system in the correct order for region based processing.

The binning units 1068 provide a hardware implemented region sorting system, (shown in Figure 7), which removes the sorting process from the processing elements, thereby releasing the PEs for data processing.

MEMORY ACCESS CONSOLIDATION

In a computer system having a large number of elements which require access to a single memory, or other addressed device, there can be a significant reduction in processing speed if accesses to the storage device are performed serially for each element.

The graphics system described above is one example of such a system. There are a large number of processor elements, each of which requires access to data in the local memory of the system. Since the number of elements requiring memory access exceeds the number of memory accesses that can be made at any one time, accesses to the local and system memory involves serial operation. Thus, performing memory access for each element individually would cause a degradation in the speed of operation of the processing block.

In order to reduce the effect of this problem on the speed of processing of the system, the system of

Figures 1 and 2 includes a memory access consolidating function.

5 The memory access consolidation is also described below with reference to figures 12 and 13. In general, however, the processing elements that require access to memory indicate that this is the case by setting an indication flag or mark bit. The first such marked PE is then selected, and the memory address to which it
10 requires access is transmitted to all of the processing elements of the processing block. The address is transmitted with a corresponding transaction ID. Those processing elements which require access (ie. have the indication flag set) compare the transmitted address
15 with the address to which they require access, and if the comparison indicates that the same address is to be accessed, those processing elements register the transaction ID for that memory access and clear the indication flag.

20 When the transaction ID is returned to the processing block, the processing elements compare the stored transaction ID with the incoming transaction ID, in order to recover the data.

25 Using transaction IDs in place of simply storing the accessed address information enables multiple memory accesses to be carried, and then returned in any order. Such a "fire and forget" method of recovering data can
30 free up processor time, since the processors do not have to await return of data before continuing processing steps. In addition, the use of transaction ID reduces the amount of information that must be stored by the processing elements to identify the data
35 recovery transaction. Address information is generally

of larger size than transaction ID information.

5 Preferably, each memory address can store more data
than the PEs require access to. Thus, a plurality of
PEs can require access to the same memory address, even
though they do not require access to the same data.
This arrangement can further reduce the number of
memory accesses required by the system, by providing a
hierarchical consolidation technique. For example,
10 each memory address may store four quad bytes of data,
with each PE requiring one quad byte at any one access.

15 This technique can also allow memory write access
consolidation for those PEs that require write access
to different portions of the same memory address.

20 In this way the system can reduce the number of memory
accesses required for a processing block, and hence
increase the speed of operation of the processing
block.

25 The indication flag can also be used in another
technique for writing data to memory. In such a
technique, the PEs having data to be written to memory
signal this fact by setting the indication flag. Data
is written to memory addresses for each of those PEs in
order, starting at a base address, and stepped at a
predetermined spacing in memory. For example, if the
step size is set to one, then consecutive addresses are
30 written with data from the flagged PEs.

PROCESSING BLOCKS

35 One of the processing blocks 106 is shown in more

detail in Figure 8. The processing block 106 includes an array of processor elements 1061 which are arranged to operate in parallel on respective data, items but carrying out the same instruction (SIMD). Each
5 processor element 1061 includes a processor unit 1061a, a PE register file 1061b and a PE memory unit 1061c. The PE memory unit 1061c is used to store data items for processing by the processor unit 1061a. Each
10 processor unit 1061a can transfer data to and from its PE memory unit 1061c via the PE register file 1061b. The processor unit 1061a also uses the PE register file 1061b to store data which is being processed. Transfer of data items between the processor unit 1061a and the memory unit 1061c is controlled by the array controller
15 104.

Each of the processing elements is provided with a data input from the mathematical expression evaluator (MEE) 1062. The MEE operates to evaluate a mathematical
20 expression for each of the PEs. The mathematical expression can be a linear, bi-linear, cubic, quadratic or more complex expression depending upon the particular data processing application concerned.

25 One particular example of a mathematical expression evaluator is the linear expression evaluator (LEE). The LEE is a known device for evaluating the bi-linear expression:

$$ax_i + by_j + c$$

30 for a range of values of x_i and y_j .

The LEE is described in detail in US Patent No.

4,590,465. The LEE is supplied with the coefficient values a, b and c for evaluating the bi-linear expression, and produces a range of outputs corresponding to different values of xi and yj. Each processing element 1061 represents a particular (xi, yj) pair and the LEE produces a specific value of the bi-linear expression for each processor element.

The bi-linear expression could, for example, define a line bounding one side of a triangle that is to be displayed. The linear expression evaluator then produces a value to indicate to the processor element whether the pixel for which the processor element is processing data lies on the line, to one side or the other of the line concerned. Further processing of the graphical data can then be pursued.

The mathematical expression evaluator 1062 is provided with coefficients from a feedback buffer (FBB) 1068 or from a source external to the processing block (known as immediates). The feedback buffer 1068 can be supplied with coefficients from a PE register file 1061b, or from a PE memory unit 1061c.

The bus structure 1064 is used to transfer data from the processor elements (register file or memory unit) to the FBB 1068. Each PE is controlled in order to determine if it should supply coefficient data to the MEE.

In one example, only one PE (at a time is enabled) to transfer data to the feedback buffer FBB 1068. The FBB queues the data to be fed to the MEE 1062. In another example, multiple PEs can transfer data to the FBB at the same time, and so the handling of the transfer of

data would then depend upon the nature of the MEE
feedback bus structure 1064. For example, the bus
could be a wired-OR so that if multiple data is
written, the logical OR of the data is supplied to the
5 MEE 1062.

The MEE operand feedback path can also effectively be
used to communicate data from one processor element to
all the others in the block concerned, by setting the a
10 and b coefficients to zero, and supplying the data to
be communicated as the c coefficient. All of the MEE
results would then be equal to the coefficient c, thus
transferring the data to the other processor elements.

15 In the present system the processing blocks 106 are
provided with opcodes (instructions) and operands (data
items) for the expression evaluator separately from one
another. Previously, instructions and data are
provided in a single instruction stream. This stream
20 must be produced during processing which can result in
a slowing of processing speed, particularly when the
operands are produced in the array itself.

In the present system, however, since the opcode is
25 separated from the operand, opcodes and operands can be
produced by different sources and are only combined
when an operation is to be performed by the MEE 1062.

30 GRAPHICS DATA PROCESSING

Figure 9 illustrates simplified steps in a graphics
data processing method using the system of Figures 1 to
8. The host system prepares data concerning the
vertices of the primitive graphical images to be
35 processed and displayed by the graphics system. The

data is then transferred, either as a block of vertex data, or vertex by vertex as it is prepared by the host system to the graphics system.

5 The data is loaded into the PEs of the graphics system so that each PE contains data for one vertex. Each PE then represents a vertex of a primitive that can be at an end of a line or part of a two dimensional shape such as a triangle.

10 The received data is then processed to transform it from the host system reference space to the required screen space. For example, three dimensional geometry, view, lighting and shading etc. is performed to produce data depending upon the chosen viewpoint.

15 Each PE then copies its vertex data to its neighbouring PEs so that each PE then has at least one set of vertex data that corresponds to a graphical primitive, be that a line, a triangle or a more complex polygon. The data is then organised on a primitive per PE basis.

20 The primitive data is then output from the PEs to the local memory in order that it can be sorted by region. This is performed by the binning unit 1069 of Figure 3, as will be described in more detail below. The binning unit 1069 sorts primitive data by region, since the data is generally not provided by the host system in the correct order for region based processing.

30 The binning units 1068 provide a hardware implemented region sorting system which removes the sorting process from the processing elements, thereby releasing the PEs for data processing.

35

All of the primitive data is written into local memory, each primitive having one entry. When data for a particular primitive is written, its extent is compared with the region definitions. Information regarding the primitives that occur in each region is stored in local memory. For each region in which at least part of a primitive occurs, a reference is stored to the part of local memory in which the primitive data is stored. In this way, each set of primitive data need only be stored once.

Once the primitive information has been stored in local memory, it is read back into the individual PEs. However, at this stage, all of the PEs in one processing block contain data concerning respective primitives occurring in a single region. From this point, a given processing block operates on data associated with a single region of the display.

Each PE then transfers, in turn, its data concerning its primitive to the MEE for processing into pixel data. For example, a PE will supply coefficient data to the MEE which define a line that makes up one side of a triangular primitive. The MEE will then evaluate all of the pixel values on the basis of the coefficients, and produce results for each pixel which indicate whether a pixel appears above, below or on the line. For a triangle, this is carried out three times, so that it can be determined whether or not a pixel occurs within the triangle, or outside of it. Each PE then also includes data about a respective pixel (i.e., data is stored on a pixel per PE basis).

Once each pixel is determined to be outside or inside the triangle (primitive) concerned, the processing for

the primitive can be carried out only on those pixels occurring inside the primitive. The remainder of the PEs in the processing block do not take any further part in the processing until that primitive is processed.

DETAILED DESCRIPTION OF THE FUNCTIONAL UNITS DESCRIBED ABOVE

10 THREAD MANAGER

A detailed description will now be given of the thread manager 102, which as mentioned above with reference to Figure 4, comprises a cache memory unit 1024 for storing instructions fetched for each thread. The cache unit 1024 could be replaced by a series of first-in-first-out (FIFO) buffers, one per thread. The thread manager also includes an instruction fetch unit 1023, a thread scheduler 1025, thread processors 1026, a semaphore controller 1028 and a status block 1030.

Instructions for a thread are fetched from local external memory 103 or from the EPU 8 by the fetch unit 1023, and supplied to the cache memory 1024 via connecting logic.

At a given time, only one thread is executing, and the scheduling of the time multiplexing between threads is determined by the dynamic conditions of the program execution. This scheduling is performed by a thread scheduler in the thread manager 102, which ensures that each processor block 106 is kept busy as much as possible. The switching from one thread to another involves a state saving and restoring overhead. Therefore, the priority of threads is used to reduce

the number of thread switches, thereby reducing the associated overheads.

5 Core instructions issued by the thread manager 102 are sent to one of two controller units, the array controller 104 or channel controller 108.

Determining which thread should be active
10 The thread scheduler, when running, recalculates which thread should be active whenever one of the following scheduling triggers occur:

15 A thread with higher priority than the current active thread is READY, or
The thread is (not Ready) and YIELDING.

20 The thread scheduler is able to determine this because each thread reports the status of whether it is READY or YIELDING back to the thread scheduler, and are examined in a register known as the Scheduler-Status register.

25 In determining the above, a thread is always deemed to be READY, unless it is:

- waiting on an instruction cache miss,
- waiting on a zero semaphore;
- waiting on a busy execution unit, or
- waiting on a HALT instruction.

30 When a thread stops operation, for example because it requires memory access, it can be "yielding" or "not yielding". If the thread is yielding, then if another thread is ready, then that other thread can become
35 active. If the thread is not yielding, then other

threads are prevented from becoming active, even though ready. A thread may not yield, for example, if that thread merely requires a short pause in operation. This technique avoids the need to swap between active threads unnecessarily, particularly when a high priority thread simply pauses momentarily.

In the event that a scheduling trigger occurs as described above, the scheduler comes into effect, and carries out the following. First, it stops the active thread from running, and waits a cycle for any semaphore decrements to propagate.

If the previously active thread is yielding, the scheduler activates the highest priority READY thread, or the lowest priority thread if no thread is ready (since this will cause another immediate scheduling trigger).

If the previously active thread is not yielding, the scheduler activates the highest priority thread which is READY which has higher priority than the previously active thread. If there is no such thread, the scheduler reactivates the previously active thread (which will cause another scheduling trigger if that thread has not become READY).

The thread scheduler can be disabled through the EPU interface. When the scheduler is disabled the EPU is able to control activation of the threads. For example, the EPU could start and stop the active thread, set the active thread pointer to a particular thread, and single step through the active thread.

The thread manager 102 only decodes thread manager

instructions or semaphore instructions. In addition,
each thread has its own thread processor 1026, as shown
in Figure 10. The thread processor 1026 can be divided
into several parts in order to aid understanding of its
5 operation.

Each thread processor comprises a byte alu 540, a
predicate alu 550, a branch unit 520, an instruction
cache 530, an instruction assembler 510 and an enable
10 unit 500.

The purpose of the thread processor 1026 is to allow
high level flow control to be performed for a thread,
(such as looping and conditional branches), and to
15 assemble instructions to be issued to the array
controller 104 and channel controller 108.

An enable unit 500 is used to determine whether a
thread is READY, as outlined in the text above.

20 The instruction cache 530 receives addresses for
instructions from the branch unit 520 and fetches them
from the cache 5301. During start up, the EPU can
program the program counters in the branch unit. If
25 the cache 5301 does not contain the instruction, a
cache miss is signalled, and an instruction fetch from
local memory is initiated. If there is no miss, the
instruction is latched into the instruction register
5302.

30 The branch adder 520 controls the address of the next
instruction. In the normal course of events, it simply
increments the last address, thus stepping sequentially
through the instructions in memory. However, if a
35 branch is requested, it calculates the new address by

adding an offset (positive or negative) to the current
address, or by replacing the current address with an
absolute address in memory. If the thread processor is
halted, a PC0 register 5201 provides the last address
5 requested, as a PC1 register 5202 will already have
been changed.

The byte alu section 540 provides a mechanism for
performing mathematical operations on the 16-bit
10 registers contained in the thread processor 102. The
programmer can use thread manager instructions to add,
subtract and perform logical operations on the thread
processor general registers 5402, thereby enabling
loops to be written. Information can also be passed to
15 the array controller 104 from the general registers by
using the byte alu 540 and the instruction assembler
510.

The predicate alu 550 contains sixteen 1 bit predicate
20 registers 5501. These represent true or false
conditions. Some of these predicates indicate carry,
overflow, negative, most significant bit status for the
last byte alu operation. The remaining predicates can
be used by the programmer to contain conditions. These
25 are used to condition branches (for loop termination),
and can receive status information from the array
controller 104 indicating "all enable registers off"
(AEO) in the array.

The instruction assembler 510 assembles instructions
30 for the various controllers such as channel controller
108 and array controller 104. Most instructions are
not modified and are simply passed on to the respective
controllers. However, sometimes fields in the various
35 instructions can be replaced with the contents of the

general registers. The instruction assembler 510 does this before passing the instruction to the relevant controller. The instruction assembler 510 also calculates the yield status, the wait status and the controller signal status sent to the enable unit 500 and the scheduler in the thread manager 102.

SEMAPHORE CONTROLLER

10 Synchronisation of threads and control of access to other resources is provided by the semaphore controller 1028.

15 Semaphores are used to achieve synchronisation between threads, by controlling access to common resources. If a resource is in use by a thread, then the corresponding semaphore indicates this to the other threads, so that the resource is unavailable to the other threads. The semaphore can be used for queueing access to the resource concerned.

20 In a particular example, the semaphore controller 1028 uses a total of eighty semaphores, split into four groups in dependence upon which resources the semaphores relate to.

Semaphore Count and Overflow

The semaphores have an eight bit unsigned count. However, the msb (bit7) is used as an overflow bit, and thus should never be set. Whenever any semaphore's bit 7 is set, the semaphore overflow flag in the thread manager status register is set. If the corresponding interrupt enable is set the EPU is interrupted. The semaphore overflow flag remains set until cleared by the EPU.

Semaphore Operations

The following operations are provided for each semaphore:

5 Preset: A thread can preset the semaphore value. The thread should issue a preset instruction only when it is known that there are no pending signals for the semaphore.

10 Wait: A thread can perform a wait operation on the semaphore by issuing a wait instruction. If the semaphore is nonzero the semaphore is decremented. If it is zero the thread is paused waiting to issue the wait instruction.

15 Signal: The semaphore is incremented. This operation can be performed by the threads, the PE Sequencer, the Load/Store Unit, or the Channel Controller. But in general a semaphore can only be
20 signalled by one of these, as discussed below.

The EPU 8 can read and write the thread semaphore counts anytime. In general, the core should not be
25 executing instructions when the EPU accesses the other semaphore values.

SEMAPHORE GROUPS

The semaphores are broken into four groups according to which execution units they can be signalled by.

30

group id	number of sems in group	semaphore group name	semaphores in group can be signalled by
0	32	Thread	threads and EPU
1	16	Channel	channel controller

2	16	Load/Store	load/store unit
3	16	PE	PE sequencer

5 The EPU can read and write all semaphore values when
the core is frozen. In addition, the EPU can preset,
increment, and decrement a thread semaphore at any time
as follows:

10 Increment: the EPU can atomically increment
the semaphore by writing its
increment register (an atomic
operation is an operation that
cannot be interrupted by other
operations, as is well known).

15 Decrement: the EPU can atomically decrement
the semaphore by reading its
decrement register. If the
semaphore is nonzero before
decrementing the read returns TRUE.
20 Otherwise the read returns FALSE
and the semaphore is left at zero.

Each thread semaphore has a separately enabled nonzero
interrupt. When this interrupt is enabled the
25 semaphore interrupts the EPU when nonzero. The EPU
would typically enable this interrupt after receiving a
FALSE from a semaphore decrement. Upon receiving the
interrupt, it is preferable to attempt the decrement
again.

30

ARRAY CONTROLLER

A detailed description will now be given of the array
controller 104, as shown in Figure 5. The array

controller 104 directs the operation of the processing block 106. The array controller 104 comprises an instruction launcher 1041, connected to receive instructions from the thread manager. The instruction launcher 1041 indexes an instruction table 1042, which provides further specific instruction information to the instruction launcher.

On the basis of the further instruction information, the instruction launcher directs instruction information to either a PE instruction sequencer 1044 or a load/store controller 1045. The PE instruction sequencer receives instruction information relating to data processing, and the load/store controller receives information relating to data transfer operations.

The PE instruction sequencer 1044 uses received instruction information to index a PE microcode store 105, for transferring PE microcode instructions to the PEs in the processing array.

The array controller also includes a scoreboard unit 1046 which is used to store information regarding the use of PE registers by particular active instructions. The scoreboard unit 1046 is functionally divided so as to provide information regarding the use of registers by instructions transmitted by the PE instruction sequencer 1044 and the load/store controller 1045 respectively.

The instruction launcher 1041 and the scoreboard unit 1046 maintain the appearance of serial instruction execution whilst achieving parallel operation between the PE instruction sequencer 1044 and the load/store

controller 1045.

5 The remaining core instructions 1032 issued from the thread manager 102 are fed to the channel controller 108. This controls transfer of data between the PE memory units and external memory (either local memory or system memory in AGP or PCI space).

10 In order to maintain the appearance of serial instruction execution, the PE instruction sequencer or Load/store controller stalls the execution of an instruction when that instruction accesses a PE register which is locked by a previously launched, still executing instruction from the load/store
15 controller and PE instruction sequencer respectively. This mechanism does not delay the launching of instructions. Instruction execution is stalled only when a lock is encountered in the instruction execution.

20 The PE register accesses which cause a stall are:

Any access to a locked register

Write to the enable stack (used as enable for load/store)

25 Write to a P register (Figure 4) (used as indexed address for load/store)

Write to a V register (Figure 4) (used as enable for MEE feedback)

30 The Instruction Launcher 1041 determines which registers an instruction accesses and locks these registers as the instruction is launched. The registers are unlocked when the instruction completes. For load/store instructions, determining the accessed

registers is straight forward. This is because the
accessed registers are encoded directly in the
instruction. For PE instructions the task is more
complex because the set of accessed registers depends
5 on the microcode. This problem is solved by using nine
bits of the PE instruction to address the instruction
table 1042 (which is preferably a small memory), which
gives the byte lengths of the four operands accessed by
the instruction.

10 The instruction table 1042 also determines whether the
instruction modifies the enable stack, P register, or V
register. Furthermore, it also contains the microcode
start address for the instruction.

When a PE instruction is launched, the instruction
15 table 1042 is accessed to determine the set of
registers accessed. These registers are marked in the
scoreboard 1046 as locked by that instruction. The
registers are unlocked when the instruction completes.
Load/Store instructions are stalled when they access or
20 use a register locked by the PE instruction sequencer
1044.

When a load/store instruction is launched, all register
file registers (R31-R0) which are loaded or stored by
that instruction are locked. The registers are
25 unlocked when the instruction completes. PE
instructions are stalled when they access a register
locked by the load/store controller.

Writes to the P registers stall execution of the
Load/Store unit as follows (V register and enable stack
30 are similar). When a PE instruction is launched, it
locks the P register if the instruction table lookup
indicates that the instruction modifies the P register.
The P register remains locked until the instruction

completes. A load/store instruction stalls while the P register is locked if the load/store instruction's Indirect bit is set. A load/store instruction stalls while the V register is locked if the load/store instruction writes the feedback buffer. A load/store instruction stalls while the enable stack is locked if the load/store instruction's Condition bit is set.

As mentioned earlier, the instruction table 1042 may be a small memory (RAM), 512 words deep by 64 bits wide. The table is addressed by the instruction index field of PE instructions to determine the instruction start address and type. The table is written with the Load Address and Load Data housekeeping instructions and is read via I address and I data registers on the EPU bus.

LOAD/STORE CONTROLLER

A detailed description will now be given of the load/store controller 1045.

In a particular example, PE memory cycles are nominally at one quarter of the PE clock rate, but can be geared to any desired rate, such as one sixth of the PE clock rate. The memory is 128 bits wide (a page), and has a quadbyte (32-bit) wide interface to the PE register file. This register file interface runs at four times the memory cycle rate, so the register file interface runs at full memory speed.

Load/store controller instructions execute in one memory cycle (nominally four PE cycles) unless they are

stalled by the instruction launcher 1041 or by cycles stolen for refresh or I/O.

Each load/store instruction transfers part or all of a single memory page. No single load/store instruction accesses more than one page.

Memory operations performed by the Load/Store Controller

The load/store controller 1045 performs the following operations on PE memory 1063:

10 loads and stores from PE memory 1063 to PE register files

 reads from PE memory 1063 to the MEE feedback buffers

 copies from PE memory to PE memory

15 PE memory refresh

 I/O channel transfers

Loading and storing from PE memory to PE register files

20 The Load and Store instructions transfer the number of bytes indicated between a single memory page and four quadbytes of the register file as follows:

25 The memory access begins at the indicated memory byte address (after applying address manipulations, see below) and proceeds for the indicated number of bytes, wrapping from the end of the page (byte 15) to the start of the page (byte 0).

 The register file access is constrained to four quadbytes of the register file. The access begins at the indicated register and proceeds through four

quadbytes, then wraps to byte 0 of the first quadbyte accessed.

Once the transfer is initiated it executes in one memory cycle.

5

Reading from PE memory to the LEE feedback buffers

10 All or part of a memory page may be copied to the MEE feedback buffer. The page address can be modified with the Memory Base Register mechanism (see below). Each quadbyte of the page can be copied into any subset of the A, B, or C parts of the MEE feedback buffer, with a feedback buffer push available after each quadbyte.

Cycle Priorities

15 Memory refresh has priority over all other memory operations. The Load/Store versus I/O Channels priority is selected by a status register bit.

Refresh

20 The PE Memory is dynamic and must be refreshed. This may be achieved in software by ensuring all pages are read every refresh period. However, the preferred method is to include a hardware refresh in the architecture.

25

Address Manipulations

The memory addresses used by the load/store controller 1045 can be manipulated with either or both of the following two mechanisms:

Memory Base Register (MBR)

5 The Memory Base Register is optionally added to the page address specified by appropriate instructions, conditioned by a bit in the instruction.

10 Each thread has its own MBR in the array controller. Threads load their MBR with a housekeeping instruction. The MBR can be read over the EPU bus.

Address Indexing

15 When an instruction's Index bit is set, the low five bits of the instruction's memory quadbyte address are ORed per PE with the low five bits of the PE's P register.

CHANNEL CONTROLLER

20 A detailed description now follows of the channel controller 108. As mentioned above, the channel controller controls the transfer of data between external memory and PE memory. At each processing block 106, a transfer engine carries out Direct Memory Access DMA transfers between the block I/O registers and the bus architecture. Depending upon the channel instruction, the data transfers go through a binning unit 1069, or directly to/from external memory.

30 The channel controller 108 operates on an instruction set which is spilt into three fundamental parts:

Read instructions which transfer data from external memory to PE memory,

Write instructions which transfer data from PE memory to external memory,

5 Housekeeping instructions which manipulate register values within the channels and binning units.

Instructions from the thread manager 102 are pushed into three separate instruction FIFOs for low priority, high priority, and binner instructions. Each FIFO has
10 its own "full" indication which is sent to the thread manager 102, so that a thread blocked on a full instruction FIFO will not prevent another thread from pushing an instruction into a non-full instruction FIFO.

15

Figure 6 shows an instruction state machine which controls the operation of the channel controller 108.

All instructions are launched from the idle state 1081.
20 The highest priority ready instruction is launched, where the instruction readiness is determined according to preset rules.

There are three priorities for channel instructions:
25 Addressed and Strided instructions can be specified as low or high priority. Binning instructions are always treated as very high priority. Lower priority instructions may be interrupted or pre-empted by higher priority ones. When a transfer instruction is pre-empted,
30 the contents of the PE page registers are returned to the PE memory pages from which they came.

They can then be restarted at a later time when the higher priority instruction has completed.

5 Addressed instruction are data transfers between PE memory and external memory where every PE specifies the external memory address of the data it wishes to read or write.

10 The data transfer is subject to the consolidation process, so that, for example, four PEs that each write to different bytes of a 32 byte packet address result in a single memory access of 32 bytes, any subset of which may contain valid data to be written to external memory. Also, any number of PEs which wish to read
15 data from the same packet address have their accesses consolidated into a single access to external memory.

20 In a Write Addressed instruction, each PE supplies 8 bytes of data together with the external memory address it is to be written to, and 8 bits which serve as byte enables. Any number of PEs which wish to write data to the same packet address have their accesses
consolidated into a single access to external memory.

25 In a Read Addressed instruction, each PE supplies an address for the data it wishes to read, and sixteen bytes of data (one half of a memory packet) are delivered back to the PE.

30 "Strided" memory accesses are data transfers between PE

memory and external memory where the external memory address of each PEs data is generated by the transfer engine. Addresses are stepped from a base register by a predetermined step size, such that the selected PEs
5 send to or receive from spaced external memory addresses. For example, if the step size is set to one, then the selected PEs access consecutive memory addresses. This has the advantage over "Addressed" transfers in that PEs can use all their I/O page
10 register data, instead of using some of it for address information. The base address for the transfer can be specified with a channel controller instruction or written by the EPU.

15 For a Write Strided instruction, each PE outputs 16 bytes of data. Data from two PEs is combined into a 32 byte data packet and written to an external memory address generated by the transfer engine. Consequently packets are written to incrementing addresses.
20 Optionally in the instruction, the external address that each PE's data was written to can be returned to the PE I/O page registers.

For potential Read Strided instructions, each PE in
25 turn receives 16 bytes of data from stepped addresses under control of the transfer engine.

Binning instructions relate to data transfers between PE memory and external memory where the data flows
30 through the binning unit of each core block between the block I/O bus and a system bus to external memory. The binning unit contains a number of control registers

that are set with special instructions. It generates external memory addresses for all the data being written to or read from external memory. It contains logic for the support of binning primitives into the regions that they fall in, and for merging multiple bin lists that are held in external memory. It also performs management of bin lists in external memory. Data flow between PEs and the binning unit are buffered in a FIFO.

5
10

BINNING FUNCTION

As mentioned above, each processing block 106 has an associated binning unit 1069, which is attached between the block I/O bus and the system bus 6. The binning unit provides specific support for the writing and reading of primitive pointers in bin lists in external memory.

15

The binning process must maintain primitive order between the geometry and rasterisation phases due to requirements of most host systems. Since both phases are block parallel, there needs to be a mechanism for transferring data between any block to any of the bins and between any bin and any block. This is implemented by creating multiple bin lists per region, one for every processing block 106 that is processing geometry data. This allows the geometry output phase to proceed in block parallel mode. Then, during the rastering phase, each region is processed by just one processing block 106, and a merge sort of the multiple bin lists in memory for that region is performed.

20

25

30

The binning unit 1069 only handles pointers. Primitive data itself can be written to memory using normal channel write operations. It can also be read using normal channel read operations once the binner hardware has provided each PE with a primitive pointer.

A record is kept of how many primitives are written to each bin, so that regions can be sorted into similar size groups for block parallel rasterisation. In addition, primitive "attribute" flags are recorded per region. This allows optimisation of rasterisation and shade code per region by examining the bitwise "OR" of a number of defined flags of every primitive in a region. In this way regions requiring similar processing can be grouped for parallel processing, which results in reduced processing time.

After the PE array 1061 has computed bounding boxes for primitives, the binner hardware offloads the binitization process from the PE array 1061, and turns it into a pure I/O operation. This enables it to be overlapped with some further data processing , for example the next batch of processing geometry data.

Writing - On writing the primitive pointers at the end of a geometry pass, the PEs output the pointers, flags and bounding box information for primitives on the channel. The binning unit 1069 appends the pointer to the bin list of every region included in the bounding box for that primitive. It also updates the primitive count and attribute flags for that region. The binner is responsible for maintaining the bin lists only for

its processing block 106, and the bin list state is preserved across multiple geometry passes.

5 Reading - The binning unit 1069 supplies ordered primitive pointers to the processing block 106, one per PE that requests, for a specific region. It traverses the multiple bin lists for that region, with a merge sort to restore original primitive order. Bin list state is preserved across multiple rasterisation
10 passes.

Binning Memory organisation

The bin lists are created in external memory, by outputting list data to memory. The bin lists indicate
15 the locations of the contents of the bin within memory. Maintenance of such linked list structures requires additional storage in the form of pointer arrays. The binner hardware accesses these structures in memory directly.

20

BINNING HARDWARE

The binning hardware is shown in detail in Figure 7, and is responsible for handling the computation
involved in the binning process needed to enable
25 the PE array 1061 to read and write primitive pointers to external memory.

Instruction decoder 1101 receives instructions from the
channel controller 108, and triggers the state machine
30 1102 into operation. The state machine 1102 is the

logic that sequences the other parts of the binning unit to perform a particular function such as reading or writing primitive pointers to or from external memory. The state machine 1102 may be implemented as several communicating state machines. Control signals to all other parts of the binning unit are not shown.

The binning function is executed by the binning unit according to a set of internal registers 1103 that define the current binning context, that is the location of bin lists in external memory, the region to be rasterised next, the operation mode and so on. This set of "state" registers 1103 is multiplexed to the channel controller 108, the block I/O bus and the EPU 8 (ie. the registers have a number of ports that can be used simultaneously).

Between the block I/O bus and the binning unit 1069 itself there is a data buffer FIFO 1104, which is regarded as being part of the binning unit 1069. The purpose of the data buffer 1104 is to buffer data flowing between the PE I/O page registers and the binning unit 1069, to smooth out the indeterminate timing of the binning unit 1069. Data is transferred to/from the binning unit 1069 in bursts of size that depends on the buffer depth. The binning unit 1069 presents the status of this buffer to the rest of the block control logic, and by looking at the status of all the binning unit buffers 1104, the channel controller 108 can schedule data transfer bursts to the binning units 1068 in an efficient way.

The binning unit 1069 of each block has its own register set interface 1105 to the EPU 8. The EPU 8 performs the following set of binning unit 1069 tasks via the interface 1105:

- 5 Initialisation
- Allocation of bin list memory
- Save and restore of binning state on context switch

10 When the binning unit 1069 is executing a Write Binner instruction, it needs an unknown amount of memory to be allocated for the creation of bin lists. It requests this memory a portion at a time from the EPU 8, and assigns it to whichever bin lists require it. The binner unit 1068 assigns small chunks (portions) of 32
15 bytes to bin lists, but this would load the EPU intolerably if it were to be allocated at this level. Instead, the EPU provides large portions of data of whatever size it decides is appropriate (for example, 64kBytes, but any convenient multiple of 32 bytes) and
20 the binner unit 1068 divides this up into individual chunks, using the chunk generator 1106. The transfer of large amounts of data from the EPU is more efficient for the EPU, and the processing of small amounts of
25 data for the binning unit 1069 is more efficient for the binning unit 1069.

 During pointer writing, primitive data from PEs is lodged in a register set 1107, and passed to the data logic 1112 as required.

30

 A Y stepper 1108 is used to step the y axis region co-

ordinate across the primitive bounding box during pointer writing as part of the binitization process. It comprises a counter and register pair with an equality comparator.

5

A X stepper 1109 is used to step the X axis region coordinate across the primitive bounding box during pointer writing as part of the binitization process. It also comprises a counter and register pair with an equality comparator. However, since the X stepper must also run the same sequence of values for every value of the Y stepper 1108, the counter is loaded and reloaded from an extra register that contains the initial value.

10

15

To merge block bin lists for a region during the pointer read process, there is provided a dedicated hardware section 1110. So that primitives can be ordered through the binning process, a batch id code is added to the bin lists. The batch id code relates to the geometry ordering, since host requires geometry to be returned in the correct order. Under control of the state machine 1102, and aided by a block counter 1117, the binning unit 1069 evaluates which bin list has the lowest batch ID and directs pointer reading from that list.

20

25

When a further batch ID is encountered in that list, or a NULL terminator encountered, the block selection is re-evaluated. The block counter 1117 provides a loop counter for the state machine 1102 when it is evaluating the next bin list to process (in conjunction with the bin list selection unit 1110).

30

The Data logic unit 1112 is the data processing block of the binning unit 1069. It is able to increment pointers, merge attribute flags and format different data types for writing to external memory via the data cache 1115.

A region number unit 1116 computes a linear region number from the X and Y region co-ordinates outputted from the X/Y steppers 1108/1109. This number, together with the output of the data logic unit 1112 and state registers 1103, are used by an address compute unit 1113, to compute a memory address for bin list array entries.

The data cache 1115 is provided for decoupling all memory references from the external memory bus. It exploits the address coherence of the binning unit memory accesses to reduce the external memory bandwidth, and to reduce the stall time that would be caused by waiting for data to arrive.

The data cache 1115 has an address tag section 1114. This indicates to the binning unit 1069 whether any particular external memory access is a hit or a miss in the data cache. On miss, the binning unit 1069 is stalled until the required data packet is fetched from memory.

PROCESSING ELEMENTS

Figure 11 shows a processor unit 1061a and PE register file 1061b which form part of the processing element

shown in Figures 3 and 8. The PE 1061 includes an arithmetic logic unit (alu) 214 which is connected to receive data values from a block of 8 bit registers 202, 204, 206, 208 (designated R, S, V and P) via
5 multiplexers 210 and 212 (A and B).

The PE register file 1061b which operates to buffer data between the PE and its associated PE memory, and to store temporarily data on which the processor unit
10 1061a is processing.

The RSVP registers 202, 204, 206, 208 operate to supply operands to the alu 214. The A multiplexer 210 receives data values from the R and S registers and so
15 controls which of those register values is supplied to the alu 214. The B multiplexer 212 is connected to receive data values from the V and P registers and also from the MEE 1062, and so controls which of those values is to be supplied to the alu.

20

The processor unit 1061a further includes a shifter 200 which can perform a left or right shift on the data output from the S, V and P registers.

25

The R register can hold its previous value, or can be loaded with a byte from the register file, or the result from the alu. The alu result is 10 bits wide, and so the R register can receive the first 8 bits (bits 7 to 0) or bits 9 to 2, for a Booth multiply
30 step. Booth multiplication is a well known way of providing multiplication results in one clock cycle.

The S register can hold its previous value, or can be loaded with a shifted version of its previous value. The S register can also be loaded with the alu result, a bit from the register file or the low 2 bits from the alu concatenated with the high 6 bits of the S registers previous value (for the Booth multiply step).

The V and P registers can both be loaded with the alu result, or a byte from the register file. The lsb of the V register is used to determine the set of processor elements which are participating in MEE feedback transfer. The five low bits of the P register are used to modify the memory address in memory accesses.

Using four registers R, S, V and P provides the system with improved performance over previously known systems because any of the registers are able to provide data to the alu 214. In addition, any of the registers can be loaded with data from the PE register file 1061b, which improves the generality of the system, and provides better support for floating point operations. Since the R register input is never shifted, the R register can be used to store and modify the exponent of floating point numbers.

The alu 214 receives instructions from the array controller (not shown) and supplies its output to the PE register file 1061b. The PE register file 1061b is used to store data for immediate use by the PE, for example, the register file 1061b can store 16 words of 16 bits in length.

Data to be written to the register file is transferred via a write port, and data to be read from the register file is transferred via a read port. Data is transferred to and from the register file from the PE memory via a load/store port under the control of the load/store controller.

The PE register file 1061b can receive data to be stored through its write port in a number of ways: a 16 bit value can be received from the processor element which form the element's left or right neighbour, a 16 bit value can be received from a status/enable register, or an 8 bit value can be received from the alu result. In the case that the alu result is supplied to the register file, the 8 bit value is copied into both the high and low bytes of the register file entry concerned.

The write port is controlled on the basis of the source of data, and is usually controlled by way of the contents of the enable stack. It is possible to force a register file write regardless of the enable stack contents.

The processor unit 1061a also includes an enable stack which is used to determine when the alu 214 can process data. The enable stack provides 8 enable bits which indicate if the alu can operate on the data supplied to it. In a preferred example, the alu 214 will only operate if all 8 bits are set to logical 1. A stack of enable bits is particularly useful when the alu is to perform nested conditional instructions. Such nested

instructions tend to occur most often in IF, ELSE, ENDIF instruction sequences.

5 By providing an enable stack of multiple bits in hardware, it is possible to remove the need for software to save and load the contents of a single enable bit when the alu is processing a nested instruction sequences.

10 The read and write ports of the PE register file 1061b enable a 16 bit data word to be copied to the PE register file of at least one of the neighbouring PEs.

The load and store operations can be issued in parallel with microcoded alu instructions from the array controller. The PE register file 1061b provides
15 several performance advantages over previous systems in which the alu has directly accessed a memory device. The PE register file 1061b provides faster access to frequently used data values than a processor element to
20 memory or memory to memory architecture can provide. In addition, there are no restrictions on the order in which data values are ordered in the register file, which further aids speed of processing and programming flexibility.

25 Figure 12 is a block diagram illustrating a processing element, and data input and output lines to that element. As previously described, the processing element includes a processor unit 1061a, a PE register
30 file 1061b, and a PE memory unit 1061c. The memory unit 1061c is preferably DRAM which is able to store 128 pages of 16 bytes. Alternatively, other memory

configurations could be used for the PE memory unit.
Data items can be transferred between the PE register
file 1061b and the PE memory unit 1061c by way of
memory read data and memory write data lines 1078 and
1079.

5

In addition, data can be transferred out of the
processor element, and indeed out of the processor
block in which the element is situated, by way of a
block I/O data out bus 1067d, and can be transferred
into the processor block by way of a block I/O data in
bus 1067c. Address transaction ID and data transaction
ID information can be transferred to the processor
block by way of busses 1067a and 1067b. The MEE
feedback data is transferred from the PE memory unit
1061c or the PE register file 1061b to the MEE feedback
buffer (not shown) by way of a MEE feedback data out
bus 1064.

10

15

20

Figure 13 shows the block I/O interface in more detail.
PE memory read and write data buses 1078 and 1079
interface with a block I/O register file 1071 for
transferring data between the register and the
processing unit and the memory unit. Data to be read
out from the processing element is output from the
block I/O register file 1071 onto the block I/O data
out bus 1067c, and data to be read into the processing
element concerned is input to the block I/O register
file 1071 from the block I/O in bus 1067d.

25

30

The processing elements that require access to memory
indicate that this is the case by setting an indication

flag or mark bit. The first such marked PE is then selected, and the memory address to which it requires access is transmitted to all of the processing elements of the processing block. The address is transmitted with a corresponding transaction ID. Those processing elements which require access (ie. have the indication flag set) compare the transmitted address with the address to which they require access, and if the comparison indicates that the same address is to be accessed, those processing elements register the transaction ID for that memory access and clear the indication flag.

All those PEs requiring access to memory (including the selected PE) then compare the required address with the address transmitted on the block I/O inbus 1067d, by way of an address compare unit 1073. If the result of the address compare demonstrates that the selected address is required for use, then the byte mask is unset and the transaction ID for the memory access concerned is stored in a transaction ID register 1075. The address transaction ID is supplied on the address transaction ID bus 1067a. Later, the required data carrying the same transaction ID returned along the block I/O data inbus 1067d. Simultaneously, or just before the data is returned, the transaction ID is returned along the data transaction ID bus 1067b all of the processor elements compare the returned data transaction ID with transaction ID stored in the transaction ID register 1075 by means of comparator 1076. If the comparison indicates that the returned transaction ID is equivalent to the stored transaction ID, the data arriving on the block I/O data inbus 1067d is input into the PE register file 1061b. When the

transaction ID is returned to the processing block, the processing elements compare the stored transaction ID with the incoming transaction ID, in order to recover the data.

5

Using transaction IDs in place of simply storing the accessed address information enables multiple memory accesses to be carried, and then returned in any order.

10

Booth multiplication is achieved using the B multiplexer 212, which is shown in more detail in Figure 14. The B multiplexer 212 receives inputs 230 from the V and P registers and from the MEE 1602. The B multiplexer 212 includes a Booth recode table 218 and a shift and complement unit 220. The Booth recode table 218 receives inputs 224, 226 from the two least significant bits of the S register and from a Booth register (S reg and Boothreg). Booth recoding is based on these inputs and the Booth recode table transforms these bits into shift, transport and invert control bits which are fed to the shift and complement unit 220. The shift and complement unit 220 applies shift, transport and invert operations to the contents of the V register. The shift operation shifts the V register one bit to the left, shifting in a 0, and the transport and invert bits cause the possibly shifted result to be transported, inverted or zeroed or a combination of those.

20

25

30 Figure 15 shows a block diagram of the alu 214 of the processor element shown in figure 13. The alu 214 receives 10 bit inputs 234 from the A and B

..

multiplexers 210 and 212, and also receives inputs 244 and 246 from the BoothCarryIn and CarryReg registers. The alu 214 also receives instructions from the controller. The alu 214 includes a carry propagate unit 236, a carry generate unit 238 and a carry select unit 242. The alu also includes an exclusive OR (XOR) gate 250 for determining the alu result output. A CarryChain unit 240 receives inputs from Carry propagate unit 236 and the carry generate unit 238, and outputs a result to the XOR gate 250.

The various units in the alu 214 operate to carry out instructions issued by the controller.

CLAIMS:

1. A method of controlling data read access to memory
in a data processing apparatus including a
5 plurality of processing elements, the method
comprising:

selecting a processing element that requires
access to the memory,

10 retrieving a target address from the selected
processing element,

transmitting the target address to the plurality
of processing elements,

15 transmitting transaction identification
information to the processing elements, which
information identifies the target address access
operation concerned,

20 storing the transaction identification information
in the or each processing element that requires
access to the target address,

transmitting data obtained from the target address
together with the transaction identification
information to the plurality of processing
elements, and

25 storing the obtained data in the or each
processing element in which the transaction
identification information is stored.

2. A method of retrieving a data item from a memory

unit in a data processing apparatus which includes
a memory unit in which data items are stored at
addresses therein, and a plurality of processing
elements which have access to the memory unit, the
method comprising:

5

for each processing element in the array which
requires access to the memory unit, setting an
access indicator to indicate that the processing
element concerned requires access to the memory
unit, and storing a target address of the memory
unit to which such access is required;

10

selecting one of the processing elements having
the access indicator set, and retrieving the
stored target address from that selected
processing element;

15

transmitting the retrieved target address and
transaction identification information to the
processing elements in the array;

20

for each processing element having the access
indicator set, comparing the transmitted target
address with the stored target address, and if the
stored and transmitted target addresses are
equivalent, clearing the access indicator and
storing the transaction identification
information;

25

30

retrieving at least one data item stored at the transmitted target address in the memory unit;

5 transmitting the or each retrieved data item and associated transaction identification information to the processing elements in the array; and

10 for each processing element having stored transaction identification information, comparing the stored transaction identification information with the transmitted transaction identification information, and if the stored transaction information is equivalent to the transmitted transaction information, receiving the or each
15 retrieved data item.

3. A method as claimed in claim 1 or 2, wherein the required data is returned in the order in which the transaction identification information is
20 produced.

4. A method as claimed in claim 1 or 2, wherein the required data is returned in the order in which it is retrieved from the memory.
25

5. A method of writing data items to a memory unit in a data processing apparatus including the memory unit in which data items are stored at addresses therein, and a plurality of processing elements
30 which have access to the memory unit, the method

comprising:

5 for each processing element in the array which
requires access to the memory unit, setting an
access indicator to indicate that the processing
element concerned requires access to the memory
unit, and storing a target address of the memory
unit to which such access is required;

10 selecting one of the processing elements having
the access indicator set, and retrieving the
stored target address from that selected
processing element;

15 transmitting the retrieved target address and
transaction identification information to all the
processing elements in the array;

20 for each processing element having the access
indicator set, comparing the transmitted target
address with the stored target address, and if the
stored and transmitted target addresses are
equivalent, clearing the access indicator and
storing the transaction identification
25 information;

30 transmitting transaction identification
information to the processing elements in the
array;

for each processing element having stored transaction identification information, comparing the stored transaction identification information with the transmitted transaction identification information, and if the stored transaction information is equivalent to the transmitted transaction information, transmitting at least one data item to be stored in the memory unit at the target address; and

5

10

storing the or each transmitted data item at the target address in the memory unit.

15

6. A method as claimed in claim 3, wherein processing elements store data items at respective regions of the target memory address.

20

7. A data processing apparatus comprising a SIMD (single instruction multiple data) array of processing elements, data storage means for storing data items for access by the processing elements, and control means for controlling access to the storage means in accordance with a method as claimed in any one of the preceding claims.

25

8. An apparatus as claimed in claim 7, provided on a single integrated circuit.

30

9. A graphical data processing system comprising a host general data processing apparatus and a data

processing apparatus as claimed in claim 7 for processing graphical data.

- 5 10. A system as claimed in claim 9, provided on a single integrated circuit.



Application No: GB 9908199.4
Claims searched: 1 to 10

Examiner: Julyan Elbro
Date of search: 5 November 1999

**Patents Act 1977
Search Report under Section 17**

Databases searched:

UK Patent Office collections, including GB, EP, WO & US patent specifications, in:
UK CI (Ed.Q): G4A (AMB)
Int Cl (Ed.6): GO6F 12/00, 12/08, 12/14, 13/16
Other: ONLINE: COMPUTER EPODOC INSPEC JAPIO WPI

Documents considered to be relevant:

Category	Identity of document and relevant passage	Relevant to claims
A	WO 96/41250 A2 S MOS	
A	US 4538226 NEC	

X	Document indicating lack of novelty or inventive step	A	Document indicating technological background and/or state of the art.
Y	Document indicating lack of inventive step if combined with one or more other documents of same category.	P	Document published on or after the declared priority date but before the filing date of this invention.
&	Member of the same patent family	E	Patent document published on or after, but with priority date earlier than, the filing date of this application.