

(19)



(11)

EP 1 936 501 B1

(12)

EUROPEAN PATENT SPECIFICATION

(45) Date of publication and mention of the grant of the patent:
08.08.2018 Bulletin 2018/32

(51) Int Cl.:
G06F 5/10 (2006.01) G06F 9/54 (2006.01)

(21) Application number: **06756828.7**

(86) International application number:
PCT/JP2006/310908

(22) Date of filing: **31.05.2006**

(87) International publication number:
WO 2007/020740 (22.02.2007 Gazette 2007/08)

(54) BUFFER MANAGEMENT METHOD AND BUFFER MANAGEMENT DEVICE

PUFFERVERWALTUNGSVERFAHREN UND PUFFERVERWALTUNGSEINRICHTUNG

MÉTHODE DE GESTION DE TAMPON ET DISPOSITIF DE GESTION DE TAMPON

(84) Designated Contracting States:
AT BE BG CH CY CZ DE DK EE ES FI FR GB GR HU IE IS IT LI LT LU LV MC NL PL PT RO SE SI SK TR

• **OHKAWA, Yasukichi**
Tokyo 107-0062 (JP)

(30) Priority: **15.08.2005 JP 2005235580**

(74) Representative: **Müller Hoffmann & Partner**
Patentanwälte mbB
St.-Martin-Strasse 58
81541 München (DE)

(43) Date of publication of application:
25.06.2008 Bulletin 2008/26

(56) References cited:
WO-A-03/019350 JP-A- 08 087 478
JP-A- 08 241 186 JP-A- 2006 040 285
US-A- 5 867 734 US-B1- 6 173 307
US-B1- 6 304 924

(73) Proprietor: **Sony Interactive Entertainment Inc.**
Tokyo 108-0075 (JP)

(72) Inventors:
• **INOUE, Keisuke**
Tokyo 107-0062 (JP)

EP 1 936 501 B1

Note: Within nine months of the publication of the mention of the grant of the European patent in the European Patent Bulletin, any person may give notice to the European Patent Office of opposition to that patent, in accordance with the Implementing Regulations. Notice of opposition shall not be deemed to have been filed until the opposition fee has been paid. (Art. 99(1) European Patent Convention).

Description

TECHNICAL FIELD

[0001] The present invention relates to a method for managing a buffer which is intended to exchange data between processing entities in a system that may have a plurality of processing entities, and a buffer management apparatus.

BACKGROUND ART

[0002] A multiprocessor system, or a system that includes a plurality of processors, can perform processing in parallel or in a cooperative fashion to achieve speedup of the entire processing. Parallel cooperative processing entails data exchange between processors. Processors called producers generate data, which are passed to processors called consumers and are processed by the consumers. The efficiency of the entire system varies depending on how the data exchange between the producers and the consumers is devised.

[0003] Aside from multiprocessor systems, data exchange also occurs between tasks in a multitask environment (including multiprocesses and multithreads). In the following description of this specification, processors, tasks, and the like that exchange data with each other will be referred to as processing entities. In the case of multitasking, tasks serve as either of producers and consumers. As with the data exchange between processors, the efficiency of the entire system varies depending on how the data exchange between tasks is devised. US 6 173 307 B1 (DREWS PAUL [US]) discloses a mechanism for producers (writers) of fixed-size data items to deliver those items to consumers (readers) using a shared multiple-reader, multiple-writer circular queue. The mechanism provides two markers for manipulating the slots of the circular queue: a next-available marker that indicates the next slot available to the producer for writing (but not yet acquired) and a next-ready marker to indicate the next slot to be made available to the consumer (but not yet made ready). In the disclosed mechanism, a producer obtains an available slot for writing by incrementing the next-available marker, writes the data item into the available slot (producer-busy slot), and, when the writing is completed, makes the available slot ready for reading by incrementing the next-ready marker as far as possible (e.g. until a producer-busy slot is found).

DISCLOSURE OF THE INVENTION

PROBLEMS TO BE SOLVED BY THE INVENTION

[0004] The present invention has been achieved in view of the foregoing circumstances. It is thus a general purpose of the present invention to provide a buffer management technology capable of managing a buffer intended to exchange data between processing entities,

thereby improving the processing efficiency.

MEANS TO SOLVE THE PROBLEMS

[0005] A first embodiment according to the present invention is a method for managing a buffer which is divided into a plurality of blocks, the blocks being used cyclically in a predetermined order by producers and consumers as a temporary storage location for data to be transferred between processing entities. The producers are processing entities for writing data. The consumers are processing entities for reading data written by the producers. This method includes: providing the first block to be written with a leading pointer and a following pointer which designate the block; and retaining block state information indicating which states respective associated blocks are in, busy, write-completed, or read-completed. Then, after the first block to be written starts to be written, the positions of the pointers and the block state information are updated in accordance with the progress of writing and reading.

[0006] The pointers are moved under the constraint that the two pointers move in the same direction and do not pass each other. After the block designated by the leading pointer starts to be written, the leading pointer is moved to a next block only if the next block is in the read-completed state. After the block designated by the following pointer starts to be read, the following pointer is moved to a next block only if the next block is in the write-completed state.

[0007] In the present invention, a "processing entity" refers to an entity that processes data, either being capable of reading data from a temporary storage location and processing the same or capable of writing processed data to a temporary storage location. This processing entity is not limited to an individual processor in a multiprocessor system, but may cover a task, a process, a thread, and the like.

[0008] A second embodiment according to the present invention is also a method for managing a buffer which is divided into a plurality of blocks, the blocks being used cyclically in predetermined order by producers and consumers as a temporary storage location for data to be transferred between processing entities. The producers are processing entities for writing data. The consumers are processing entities for reading data written by the producers.

[0009] This method includes updating and retaining block state information including write completed-or-not information and read completed-or-not information. The write completed-or-not information indicates whether associated blocks are write-completed or not. The read completed-or-not information indicates whether associated blocks are read-completed or not.

[0010] Then, the first block to be written is provided with a first leading pointer, a second leading pointer, a first following pointer, and a second following pointer which designate the block. Under the constraint that all

the pointers move in the same direction, and the first leading pointer, the second leading pointer, the first following pointer, and the second following pointer do not pass each other in this order: after a block starts to be written, the first leading pointer designating the block is moved to the next block. After the block designated by the second leading pointer finishes being written, the second leading pointer is moved to a block next to a block that is farthest from the second leading pointer among consecutive write-completed blocks subsequent to the block designated by the second leading pointer. Moreover, after a block starts to be read, the first following pointer designating the block is moved to a next block. After the block designated by the second following pointer finishes being read, the second following pointer is then moved to a block next to a block that is farthest from the second following pointer among consecutive write-completed blocks subsequent to the block designated by the second following pointer.

[0011] Arbitrary combinations of the aforementioned components, and implementations of the present invention in the form of systems, programs, and program-containing recording media may also be practiced as applicable embodiments of the present invention.

ADVANTAGES OF THE INVENTION

[0012] The present invention is advantageous when exchanging data between processing entities in a multiprocessor, multitask, or other system that may have a plurality of processing entities.

BRIEF DESCRIPTION OF THE DRAWINGS

[0013]

Fig. 1 is a diagram showing a multiprocessor system which is a first embodiment according to the present invention;

Fig. 2 is a diagram showing a shared memory of the multiprocessor system shown in Fig. 1;

Fig. 3 is a diagram showing the initial state of a buffer which is formed in the shared memory shown in Fig. 2;

Fig. 4 is a diagram showing the initial state of a bitmap which is retained in the shared memory shown in Fig. 2;

Fig. 5 is a diagram for explaining the movement of pointers;

Fig. 6 is a diagram for explaining the movement of the pointers and the updating of bitmap information; Fig. 7 is a diagram showing the relationship between the bit values of each bit of the bitmap and block states;

Fig. 8 is a diagram showing a multiprocessor system which is a second embodiment according to the present invention;

Fig. 9 is a diagram showing the shared memory of

the multiprocessor system shown in Fig. 8;

Fig. 10 is a diagram showing the initial state of the buffer which is formed in the shared memory shown in Fig. 9;

Fig. 11 is a diagram showing the initial states of two bit strings which are retained in the shared memory shown in Fig. 9;

Fig. 12 is a diagram (1) for explaining the movement of the pointers and the updating of the bit strings;

Fig. 13 is a diagram (2) for explaining the movement of the pointers and the updating of the bit strings;

Fig. 14 is a diagram for explaining how the destination of the second following pointer is determined and the bit string is updated after the block designated by the second following pointer finishes being read;

Fig. 15 is a diagram showing the bit string updated after the block designated by the second following pointer finishes being read;

Fig. 16 is a diagram showing the state of the buffer and the bit string after the second following pointer is moved;

Fig. 17 is a diagram showing a multiprocessor system which is a third embodiment of the present invention; and

Fig. 18 is a diagram showing the shared memory of the multiprocessor system shown in Fig. 17.

DESCRIPTION OF REFERENCE NUMERALS

[0014] 10A ... processing unit, 10B ... processing unit, 12A ... producer, 12B ... consumer, 14A ... local memory, 14B ... local memory, 20 ... buffer, 30 ... bitmap, 44 ... leading pointer, 48 ... following pointer, 50 ... shared memory, 100 ... multiprocessor system, 110 ... processing unit, 112 ... processor, 114 ... local memory, 120 ... buffer, 130a ... bit string, 130b ... bit string, 144 ... first leading pointer, 145 ... second leading pointer, 148 ... first following pointer, 149 ... second following pointer, 150 ... shared memory, 200 ... multiprocessor system, 210 ... processing unit, 212 ... processor, 214 ... local memory, 220 ... buffer, 230 ... bit string, 240 ... pointer queue, and 250 ... shared memory.

BEST MODE FOR CARRYING OUT THE INVENTION

[0015] The embodiments will be overviewed initially. A detailed description will then be given with reference to the drawings.

[0016] A first embodiment according to the present invention is a method for managing a buffer which is divided into a plurality of blocks, the blocks being used cyclically in a predetermined order by producers and consumers as a temporary storage location for data to be transferred between processing entities. The producers are processing entities for writing data. The consumers are processing entities for reading data written by the producers. This method includes: providing the first block to be written

with a leading pointer and a following pointer which designate the block; and retaining block state information indicating which states respective associated blocks are in, busy, write-completed, or read-completed. Then, after the first block to be written starts to be written, the positions of the pointers and the block state information are updated in accordance with the progress of writing and reading.

[0017] The pointers are moved under the constraint that the two pointers move in the same direction and do not pass each other. After the block designated by the leading pointer starts to be written, the leading pointer is moved to a next block only if the next block is in the read-completed state. After the block designated by the following pointer starts to be read, the following pointer is moved to a next block only if the next block is in the write-completed state.

[0018] In the present invention, a "processing entity" refers to an entity that processes data and is capable of reading data from a temporary storage location and processing the same or capable of writing processed data to a temporary storage location. This processing entity is not limited to an individual processor in a multiprocessor system, but may cover a task, a process, a thread, and the like.

[0019] In the following description, this method will be referred to as a first method.

[0020] In a system that may have a plurality of processing entities, a buffer can be used when exchanging data between the processing entities. The use of the buffer allows producers to write processing-completed data into the buffer, and consumers to read data from the buffer for processing. This makes it possible to improve the processing efficiency of the entire system.

[0021] With buffer management of this type, the buffer may be locked for data protection while one of the processing entities uses it, so that data in the buffer will not be modified by the other processing entities. Locking the buffer, however, precludes consumers from reading data that has finished being written by producers, if any, until the lock is released. The same applies to producers. When the buffer is locked, producers must wait for unlocking even if the buffer contains data-writable blocks. A further improvement in the processing efficiency of the system can be expected if it is made possible for a plurality of processing entities to use the buffer simultaneously, i.e., if it is possible to facilitate multiaccessing.

[0022] The first method of the present invention realizes secure multiaccess to a buffer that has a plurality of blocks to be used cyclically in a predetermined order. This method includes: updating and retaining the block state information which indicates a busy, write-completed, or read-completed state; and moving the leading pointer and the following pointer under the constraint that the two pointers move in the same direction and do not pass each other.

[0023] Here, retaining the block state information provides a great advantage when achieving multiaccess.

[0024] After the block designated by the leading pointer starts to be written, the leading pointer is moved to a next block only if the next block is in the read-completed state. Here, the timing for movement of the leading pointer depends on the design of the system.

[0025] For example, in one possible system, if the block designated by the leading pointer starts to be written when the next block is already in the read-completed state, then the leading pointer is moved to the next block simultaneously with the start of the writing. On the other hand, if the block designated by the leading pointer starts to be written when the next block is not yet in the read-completed state, then the leading pointer may be moved to the next block at a point in time when the next block enters the read-completed state. In this case, the block designated by the leading pointer is in any of: the read-completed state; the busy state which indicates that the last block having started to be written is being written; and the write-completed state which indicates that this block has finished being written. The next producer searches for the leading pointer, checks whether or not this block is in the read-completed state, and can write if in the read-completed state.

[0026] Alternatively, in a simpler system, the leading pointer is configured to designate the last block that has started to be written. The next producer searches for the leading pointer, checks whether or not a block next to a block that is designated by the leading pointer is in the read-completed state, and moves the leading pointer to the next block and starts writing it if it is in the read-completed state.

[0027] In either system, even if one producer is writing, other producers can securely write other blocks based on the leading pointer and the block state information.

[0028] The same also applies to consumers. The combined use of the following pointer and the block state information allows a plurality of consumers to securely make a simultaneous read without locking the buffer.

[0029] Consequently, it is possible to realize secure multi-access using the processing entities.

[0030] A second embodiment according to the present invention is also a method for managing a buffer which is divided into a plurality of blocks, the blocks being used cyclically in a predetermined order by producers and consumers. The producers are processing entities for writing data. The consumers are processing entities for reading data written by the producers. This method also realizes secure multiaccess. Hereinafter, this method will be referred to as a second method.

[0031] Like the first method, this second method uses leading pointers and following pointers when specifying blocks for producers to write and specifying blocks for consumers to read. Here, for the sake of distinction from a second leading pointer and a second following pointer to be described later, a leading pointer and a following pointer that perform the same role as the leading pointer and the following pointer in the first method will be referred to as a first leading pointer and a first following

pointer, respectively.

[0032] The second method includes: providing the first block to be written with a first leading pointer, a second leading pointer, a first following pointer, and a second following pointer all of which designate the block; and retaining block state information indicating whether associated blocks are write-completed or not, and whether they are read-completed or not. Then, after the first block to be written starts to be written, the positions of the pointers and the block state information are updated in accordance with the progress of writing and reading.

[0033] The second leading pointer follows the first leading pointer and is moved forward when the block designated thereby finishes being written.

[0034] The second following pointer follows the first following pointer and is moved forward when the block designated thereby finishes being read.

[0035] Moreover, in the second method, these four pointers are moved under the constraint that all the pointers move in the same direction, and the first leading pointer, the second leading pointer, the first following pointer, and the second following pointer do not pass each other in this order.

[0036] That is, according to this second method, blocks lying between the first following pointer and the second leading pointer in the moving direction of the pointers are in the write-completed state. Blocks lying between the first leading pointer and the second following pointer are ones in the read-completed state.

[0037] Consequently, when moving the first leading pointer forward, it is unnecessary to check whether the destination block is write-completed or not because the blocks up to the second following pointer are in the write-completed state. Similarly, when moving the first following pointer forward, it is unnecessary to check whether the destination block is read-completed or not because the blocks up to the second leading pointer are in the read-completed state.

[0038] In other words, the first leading pointer and the second following pointer can be used to designate blocks to be written and to protect data that is being read without locking the buffer.

[0039] The same applies consumers. The first following pointer and the second leading pointer can be used to designate blocks for consumers to read and to protect data that is being written without locking the buffer.

[0040] Now, the destinations of the second leading pointer and the second following pointer will be examined.

[0041] The second leading pointer is to be moved after the block designated thereby finishes being written. Blocks lying between the second leading pointer and the first leading pointer include blocks being written and blocks that have finished being written. The farther they are from the first leading pointer, the earlier the blocks have started to be written. Take, for example, the case where block A and block F are designated by the second leading pointer and the first leading pointer, respectively,

and four blocks B, C, D, and E are arranged therebetween in the moving direction of the pointers. These six blocks have started to be written in order of A, B, C, D, E, and F.

[0042] Blocks that started to be written first will not necessarily finish being written first. For example, blocks B and C may finish being written while block A which is designated by the second leading pointer is still being written. When block A finishes being written, block C which has already finished being written cannot be used by consumers if the second leading pointer is simply moved to the next block (block B).

[0043] Accordingly, the second method of the present invention includes retaining information that indicates whether write-completed or not, included in the block state information. When block A finishes being written, the information indicating whether write-completed or not is referred to, and the second leading pointer is moved to a block (block D) next to a block (block C) that lies farthest from the second leading pointer among consecutive write completed blocks (in this example, blocks B and C) subsequent to the block designated by the second leading pointer (in this example, block A). Consequently, even if the blocks do not finish being written in the starting order, it is possible to move the second leading pointer to the foremost block within the movable range, thereby providing write-completed blocks for consumers earlier.

[0044] The same applies to the second following pointer. Since information that indicates whether read-completed or not is retained, it is possible to move the second following pointer to the foremost block even if the blocks do not finish being read in the starting order. This makes it possible to provide read-completed blocks for producers earlier.

[0045] Furthermore, producers may sometimes fail to terminate processing even after the completion of writing if there is no information that indicates whether a block is write-completed or not. Take, for example, the producer that writes block B in the foregoing case. Even after the completion of writing, the producer must move the second leading pointer and thus cannot end processing until block A finishes being written and the second leading pointer can be moved to block B. The same applies to the producer of block C. In such situations, the producers of blocks B and C cannot be released from the processed data even after the completion of writing, which lowers the processing efficiency of the system.

[0046] Now, as in the second method of the present invention, information indicating whether write-completed or not shall be provided and producers that have finished writing, such as the producer of block B, shall update this information when finishing writing. Then, the producer of block A can refer to the information and move the second leading pointer accordingly. As a result, the producer of block B need not wait until block A finishes being written, and can end processing immediately after finishing writing.

[0047] The same also applies to reading.

[0048] As has been described, the second method of

the present invention can also realize secure multiaccess. In addition, it is possible to avoid a drop in the processing efficiency of the system even if writing or reading is not finished in the starting order.

[0049] Fig. 1 shows the configuration of a multiprocessor system 100 which is a first embodiment of the present invention. The multiprocessor system 100 has a plurality of processing units 10A, a plurality of processing units 10B, and a shared memory 50. Each of the processing units is connected to the shared memory 50.

[0050] In the multiprocessor system, individual processors are included in the respective processing units. These processing units are classified into main processing units and sub processing units. The sub processing units may all be formed using an identical architecture, or may have different configurations. The main processing units may be positioned locally to the sub processing units, such as on the same chip, the same package, the same circuit board, or the same product as the sub processing units. The main processors may alternatively be positioned remotely from the sub processing units, such as on a product that is connectable over a bus, the Internet, or other communication networks. Similarly, the sub processing units may be positioned locally to or remotely from each other.

[0051] The plurality of processing units 10A and 10B in the multiprocessor system 100 shown in Fig. 1 may include main processing units.

[0052] The processing units 10A each have a processor 12A and a local memory 14A. The processing units 10B each have a processor 12B and a local memory 14B. The processor 12A is capable of reading and writing data from/to the local memory 14A. The processor 12B is capable of reading and writing data from/to the local memory 14B.

[0053] The multiprocessor system 100 pertains to the transfer of stream data. Each of the processing units 10A writes a predetermined transfer unit of stream data (hereinafter, referred to simply as data) to a buffer 20 to be described later, which is formed in the shared memory 50. Each of the processing units 10B reads data written by the processing units 10A and transfers the same. Here, since the processing units 10A write data to the buffer 20, the processors 12A serve as producers. Since the processing units 10B read data, the processors 12B serve as consumers.

[0054] One of the processing units 10A or the processing units 10B performs the role of a service unit with respect to the other processing units when using the buffer 20. The role of the service unit includes, for example, the initial setup of the buffer 20 and the initialization of pointers to be described later. Any of the processing units may be in charge of this service unit. If the processing units 10A or the processing units 10B include any main processing unit, it is preferable, though not restrictive, that the main processing unit be in charge of this service unit.

[0055] Fig. 2 shows the shared memory 50. The shared

memory 50 has a buffer 20 and a bitmap 30. The buffer 20 is provided with a leading pointer 44 and a following pointer 48.

[0056] Fig. 3 shows the buffer 20. The buffer 20 is divided into a plurality of consecutive blocks. The blocks are used cyclically in succession in a predetermined order as shown by the arrow L in the diagram, for example. Here, the first block to be written is the zeroth block shown in Fig. 3. The leading pointer 44 and the following pointer 48 are initially positioned at the block used immediately before the zeroth block when the blocks are used cyclically (in the diagram, the nth block).

[0057] The service unit determines which block to start writing first, sets the pointers, determines the initial positions of the pointers, etc.

[0058] Fig. 4 shows the bitmap 30 in an initial state. As shown in the diagram, the bitmap 30 has a bit string 30a and a bit string 30b, in both of which each single bit is assigned to one block of the buffer 20. The bitmap 30 is the block state information for indicating whether the blocks are busy (being written or read), write-completed, or read-completed.

[0059] As shown in Fig. 4, each bit of the bit string 30a has an initial value of 1. Each bit of the bit string 30b has an initial value of 0.

[0060] Fig. 5 is a diagram showing processing whereby any one of the producers 12A moves the leading pointer 44 when writing stream data to the buffer 20. This producer 12A initially searches for the leading pointer 44 of the buffer 20. Since the leading pointer 44 designates the nth block, the producer 12 refers to the bit string 30a to check the value of the bit corresponding to a block next to the nth block, i.e., the zeroth block. Here, the bit corresponding to the zeroth block has the initial value of 1. The producer 12A therefore moves the leading pointer 44 to the zeroth block, starts writing to the zeroth, and changes the value of the bit corresponding to the zeroth block in the bit string 30a to 0.

[0061] At this point, if another producer 12A is ready for a write, this producer 12A also initially searches for the leading pointer 44 of the buffer 20. Since the leading pointer 44 is positioned at the zeroth block, this producer refers to the bit string 30a to find that the bit corresponding to the first block is 1 in value. The producer moves the leading pointer 44 to the first block, starts writing to the first block, and changes the value of the bit corresponding to the first block in the bit string 30a to 0.

[0062] Similarly, another producer 12A that is ready for a write can also perform the processing of moving the leading pointer 44 to the second block, starting writing to the second block, and changing the value of the bit corresponding to the second block in the bit string 30a to 0.

[0063] Moreover, if any of the blocks has finished being written, the producer 12A having written this changes the value of the bit corresponding to this block in the bit string 30b to 1. Suppose here that the zeroth and second blocks have finished being written, and the bits corresponding to the zeroth and second blocks in the bit string 30b are

changed to 1 in value.

[0064] When any one of the consumers 12B starts reading, it initially searches for the following pointer 48 of the buffer 20. In an initial state, the following leading pointer 48 designates the nth block. The producer 12 thus refers to the bit string 30b to check the value of the bit corresponding to a block next to the nth block, i.e., the zeroth block. At this point, if the bit corresponding to the zeroth block is 0 in value, it indicates that the zeroth block is in an unwritten initial state or is still being written. If the bit value is 1, it indicates that the zeroth block has finished being written. Only if the bit corresponding to the zeroth block in the bit string 30b is 1 in value, the consumer 12B moves the following pointer 48 to the zeroth block, starts reading from the zeroth block, and changes the value of the bit corresponding to the zeroth block in the bit string 30b to 0.

[0065] Fig. 6 shows the state of the buffer 20 and the bit strings 30a and 30b here. At this point in time, as shown in the diagram, the leading pointer 44 designates the second block, which is the latest to start being written. The following pointer 48 designates the zeroth block, which is the latest to start being read. The zeroth and first blocks are busy, and the second block is in the write-completed state. This state can be read from the bit string 30a and the bit string 30b, which will be described later.

[0066] Subsequently, the producers 12A and the consumers 12B repeat writing and reading. With the progress of writing and reading, the positions of the leading pointer 44 and the following pointer 48, the bit string 30a, and the bit string 30b are updated.

[0067] Updating of the bit string 30a, specifically, is such that when a block finishes being read, the bit value of this block is set to 1 by the consumer 12B that has read it. Then, when this block starts to be written, the bit value is reset to 0 by the producer 12A that starts writing it.

[0068] Updating of the bit string 30b, specifically, is such that when a block finishes being written, the bit value of this block is set to 1 by the producer 12A that has written it. Then, when this block starts to be read, the bit value is reset to 0 by the consumer that starts reading it.

[0069] Consequently, as shown in Fig. 7, the bit values of the bit string 30a and the bit values of the bit string 30b can indicate which states the blocks corresponding to the bits are in, busy, write-completed, or read-completed.

[0070] Moreover, the leading pointer 44 designates the last block that has started to be written. Producers 12A, when writing, search for the leading pointer 44 of the buffer 20 and refer to the bit string 30a to check whether or not a block next to the block that is designated by the leading pointer 44 is in the read-completed state. If the value of the bit corresponding to this block in the bit string 30a is 1, i.e., if this block is in the read-completed state, they move the leading pointer 44 to this block, start writing, and reset the value of the bit corresponding to this block in the bit string 30a to 0.

[0071] The following pointer 48 designates the last block that has started to be read. Consumers 12B, when

reading, search for the following pointer 48 of the buffer 20 and refer to the bit string 30b to check whether or not a block next to the block that is designated by the following pointer 48 is in the write-completed state. If the value of the bit corresponding to this block in the bit string 30b is 1, i.e., if this block is in the write-completed state, they move the following pointer to this block, start reading, and reset the value of the bit corresponding to this block in the bit string 30b to 0.

[0072] The producers 12A and the consumers 12B move the pointers under the constraint that the two pointers move in the same direction and do not pass each other.

[0073] As above, according to the multiprocessor system 100 shown in Fig. 1, the combined use of the bitmap 30 and the two pointers makes it possible to realize secure multiaccess.

[0074] Fig. 8 shows the configuration of a multiprocessor system 200 which is a second embodiment of the present invention. The multiprocessor system 200 has a plurality of processing units 110 and a shared memory 150. Each of the processing units 110 is connected to the shared memory 150. The plurality of processing units 110 in the multiprocessor system 200 may include main processing units.

[0075] The processing units 110 each have a processor 112 and a local memory 114. The processor 112 is capable of reading and writing data from/to the local memory 114.

[0076] In the multiprocessor system 200, the processing units 110 perform processing in parallel. The processors 112 individually write processed data to a buffer 120 to be described later, and then copy data written by other processing units 110 from the buffer 120 into their respective local memories 114 for processing. That is, the processors 112 included in the processing units 110 can serve as both producers and consumers. In the following description, an identical processor may be referred to as a producer or a consumer depending on whether it writes or reads data.

[0077] Any one of the processing units 110 performs the role of a service unit with respect to the other processing units when using the buffer 120. The role of the service unit includes, for example, the initial setup of the buffer 120 and the initialization of pointers to be described later. Any of the processing units may be in charge of this service unit. If the processing units 110 include any main processing unit, it is preferable, though not restrictive, that the main processing unit be in charge of this service unit.

[0078] Fig. 9 shows the shared memory 150. The shared memory 150 has the buffer 120, a bit string 130a, and a bit string 130b. The buffer 120 is provided with a first leading pointer 144, a second leading pointer 145, a first following pointer 148, and a second following pointer 149. The bit string 130a has a bit width equal to the number of producers which can access at a time. The bit string 130b has a bit width equal to the number of con-

sumers which can access at a time. For example, suppose here that there are eight processing units 110. Since each processor 112 can serve as both a producer and a consumer, the bit string 130a and the bit string 130b here have a bit width of 8 bits.

[0079] Fig. 10 shows the buffer 120. The buffer 120 is divided into a plurality of consecutive blocks. The blocks are used cyclically in succession in a predetermined order as shown by the arrow L in the diagram, for example. Here, the first block to be written is the zeroth block shown in Fig. 10. All the pointers are initially positioned at the block used immediately before the zeroth block when the blocks are used cyclically (in the diagram, the nth block).

[0080] The service unit determines which block to start writing first, sets the pointers, and determines the initial positions of the pointers.

[0081] Each bit of the bit string 130a, when it has a value of 1, indicates that the block corresponding to the bit is in a write-completed state. When it has a value of 0, it indicates that the block corresponding to the bit is in a state other than the write-completed state.

[0082] Each bit of the bit string 130b, when it has a value of 1, indicates that the block corresponding to the bit is in a read-completed state. When it has a value of 0, it indicates that the block corresponding to the bit is in a state other than the read-completed state.

[0083] Fig. 11 shows the bit string 130a and the bit string 130b in an initial state. As shown in the diagram, the bits of both the two bit strings have a value of 0 in the initial state.

[0084] Fig. 12 is a diagram showing processing to be performed by the service unit before any one of the processing units 110 starts writing. As shown in the diagram, before any one of the processing units writes data to the buffer 120, the service unit moves the first leading pointer 144 and the second leading pointer 145 from their initial positions, or the nth block, to the zeroth block at the top. It also assigns the bits of the bit string 130a to a total of eight blocks from the zeroth block to the seventh block.

[0085] At this point, in order to write data, one of the producers 112 initially searches for the first leading pointer 144 of the buffer 120. In this case, the first leading pointer 144 designates the zeroth block. The producer 112 thus copies data from its local memory 114 to the zeroth block for a write. When it starts writing to the zeroth block, this producer 112 moves the first leading pointer 144 to the next block, i.e., the first block.

[0086] At this point, if another producer 112 is ready for a write, this producer 112 also initially searches for the leading pointer 144 of the buffer 120.

Since the first leading pointer 144 is positioned at the first block, this producer 112 starts writing to the first block and moves the first leading pointer 144 to the next block, i.e., the second block. Fig. 13 shows the positions of the respective pointers and the values of the respective bits of the bit string 130a at this point. The first leading pointer 144 is on the second block, and the second leading point-

er 145 is on the zeroth block. The bits of the bit string 130a remain 0 in value, indicating that none of blocks from the zeroth block to the seventh block is write-completed.

[0087] When the zeroth block finishes being written, the producer 112 that has written it moves the second leading pointer 145. The destination will be described later. When moving the second leading pointer 145, it also assigns the bit string 130a to eight consecutive blocks starting from the destination block.

[0088] When the bit of the zeroth block in the bit string 130a is set to 1, the service unit moves the first following pointer 148 and the second following pointer 149 to the zeroth block and assigns the bits of the bit string 130b to a total of eight blocks from the zeroth block to the seventh block. Subsequently, the management of the buffer 120 is passed over to the individual processing units.

[0089] Producers 112 and consumers 112 then repeat writing and reading. With the progress of writing and reading, the positions of the four pointers, the bit string 130a, and the bit string 130b are updated. In the following description, the producers 112 and the consumers 112 shall move the pointers under the constraints that the four pointers move in the same direction, and the first leading pointer 144, the second leading pointer 145, the first following pointer 148, and the second following pointer 149 do not pass each other in this order.

[0090] When a producer 112 writes data, it searches for the first leading pointer 144 of the buffer 120, and writes to the block that is designated by the first leading pointer 144. When it starts writing, it also moves the first leading pointer 144 to the next block.

[0091] When a block other than that designated by the second leading pointer 145 finishes being written, the producer 112 that has written this sets the value of the corresponding bit in the bit string 130a to 1, and ends processing.

[0092] When a consumer 112 reads data, it searches for the first following pointer 148 of the buffer 120 and reads from the block that is designated by the first following pointer 148. When it starts reading, it also moves the first following pointer 148 to the next block.

[0093] When a block other than that designated by the second following pointer 149 finishes being read, the producer 112 that has read this sets the value of the corresponding bit in the bit string 130b to 1, and ends processing.

[0094] The second leading pointer 145 is to be moved after the block designated by this pointer finishes being written. The second following pointer 149 is to be moved after the block designated by this pointer finishes being read. Now, taking the state of Fig. 14 as an example, description will be given of the processing to be performed by a consumer 112 when the block designated by the second following pointer 149 finishes being read.

[0095] In the example of Fig. 14, the first leading pointer 144 is moved to the ninth block, and the second leading pointer 145 is moved to the seventeenth block. The first

following pointer 148 is moved to the fifteenth block, indicating that the blocks up to the fourteenth have started to be read. The second following pointer 149 is moved to the eleventh block, indicating that the blocks up to the tenth have finished being read.

[0096] Suppose here that the four blocks eleventh to fourteenth have already started to be read but none have finished. The bit string 130b is assigned to eight blocks starting from the eleventh, and all the bits are 0 in value. The bit string 130b is intended to determine the destination of the second following pointer 149. Since the second following pointer will not pass the first following pointer 148, all the bits of the blocks ahead of the first following pointer are given a value of 0.

[0097] The four blocks eleventh to fourteenth have started to be read in order of 11, 12, 13, and 14. However, the finishing order is not necessarily the same as the starting order. Suppose here that the order of finishing reading is 13, 12, 11, and 14.

[0098] When the thirteenth block finishes being read while the eleventh, twelfth, and fourteenth blocks are still being read, the consumer 112 that has read it sets the value of the bit corresponding to the thirteenth block in the bit string 130b (in this case, the third bit) to 1, and ends processing.

[0099] When the consumer 112 of the twelfth block finishes reading, it also sets the value of the bit corresponding to the twelfth block in the bit string 130b (in this case, the second bit) to 1, and ends processing.

[0100] Subsequently, the eleventh block finishes being read. Since the eleventh block is designated by the second following pointer 149, the consumer 112 that has read it sets the value of the bit corresponding to the eleventh block in the bit string 130b (the first bit) to 1, determines the destination of the second following pointer, and moves it.

[0101] Here, as shown in Fig. 15, the bits of the bit string 130b are such that the bits corresponding to the read-completed blocks, or the eleventh, twelfth, and thirteenth blocks, are 1 in value and the other bits are 0 in value.

[0102] When determining the destination, the consumer 112 uses an atomic command such as `clz` to determine the number of consecutive bits having a value of 1, starting from the top bit of the bit string 130b. In the example shown in Fig. 15, the result obtained is 3. The consumer 112 moves the second following pointer 149 ahead by a number of blocks equal to the result obtained. In the example shown in Fig. 14, the second following pointer 149 is moved from the eleventh block to the fourteenth block.

[0103] Moreover, after moving the second following pointer 149, this consumer 112 assigns the bit string 130b to eight blocks starting from the current position of the second following pointer 149, and ends processing.

[0104] In other words, when the block designated by the second following pointer 149 finishes being read, the consumer 112 that has read it sets the value of the bit corresponding to this block in the bit string 132b to 1. It

also moves the second following pointer 149 to a block next to a block that is farthest from the second following pointer among consecutive read-completed blocks subsequent to the block designated by the second following pointer 149. It then assigns the bit string 130b to eight blocks starting from the current position of the second following pointer 149.

[0105] Fig. 16 shows the pointer positions and the bit string 130b at this time. The second following pointer 149 is moved to the fourteenth block, and the bit string 130b is assigned to eight blocks starting from the fourteenth.

[0106] The same applies to the second leading pointer 145. When a block designated by the second leading pointer 145 finishes being written, the producer 112 that has written it sets the value of the bit corresponding to this block in the bit string 132a to 1. It also moves the second leading pointer 145 to a block next to a block that is farthest from the second leading pointer among consecutive write-completed blocks subsequent to the block designated by the second leading pointer 145. It then assigns the bit string 130a to eight blocks starting from the current position of the second leading pointer 145, and ends processing.

[0107] As above, according to the multiprocessor system 200 shown in Fig. 8, the four pointers and the two bit strings are used to update the pointer positions and the bit strings according to the progress of writing and reading. This realizes secure multiaccess without locking the buffer. Furthermore, it is possible to avoid a drop in system efficiency even if writing or reading is not completed in the same order as the starting order.

[0108] Moreover, the information that indicates whether blocks are write-completed or not (the bit string 130a) is retained for only eight consecutive blocks subsequent to the second leading pointer 145. This makes it possible to provide information necessary for moving the second leading pointer 145 while reducing the bit width of the bit string 130a. The same applies to the bit string 130b. This also contributes to improved processing efficiency of the system.

[0109] Furthermore, moving the second leading pointer 145 requires the bit string 130a alone, and moving the second following pointer 149 requires the bit string 130b alone. Here, since the bit string 130a and the bit string 130b are retained separately, producers refer only to the area that contains the bit string 130a when moving the second leading pointer 145. Consumers refer only to the area that contains the bit string 130b when moving the second following pointer 149. This facilitates simple processing.

[0110] The buffer managing method used in this second embodiment may be applied not only to a multiprocessor system as shown in Fig. 8, but also to any multiprocessor systems in which data is exchanged between processors, such as the multiprocessor system for transferring stream data shown in Fig. 1.

[0111] Fig. 17 shows a multiprocessor system 300 which is a third embodiment of the present invention.

[0112] The multiprocessor system 300 has a plurality of processing units 210 and a shared memory 250. Each of the processing units 210 is connected to the shared memory 250. The plurality of processing units 210 in the multiprocessor system 300 may include main processing units.

[0113] The processing units 210 each have a processor 212 and a local memory 214. The processor 212 is capable of reading and writing data from/to the local memory 214.

[0114] In the multiprocessor system 300, the processing units 210 perform processing in parallel. The processors 212 individually write processed data to a buffer 220 to be described later, and then copy data written by other processing units 210 from the buffer 220 into their respective local memories 214 for processing. That is, the processors 212 included in the processing units 210 can serve as both producers and consumers. In the following description, an identical processor may be referred to as a producer or a consumer depending on whether it writes or reads data.

[0115] Any one of the processing units 210 performs the role of a service unit with respect to the other processing units when using the buffer 220. The role of the service unit includes, for example, the initial setup of the buffer 220 and the initialization of a bit string to be described later. Any one of the processing units may be in charge of this service unit. If the processing units 210 include any main processing unit, it is preferable, though not imperative, that the main processing unit be in charge of this service unit.

[0116] Fig. 18 shows the shared memory 250. The shared memory 250 has the buffer 220, a bit string 230, and a pointer queue 240.

[0117] The buffer 220 is divided into a plurality of consecutive blocks. Identifiers (here, block numbers) are given to the respective blocks.

[0118] The bit string 230 has a bit width equal to the number of blocks included in the buffer 220, and the individual bits correspond to the respective blocks. The initial values of all the bits are 0.

[0119] The pointer queue 240 retains the numbers of write-completed blocks in order of completion of writing. In an initial state, it is empty.

[0120] When a producer 212 is starting to write, it selects a block that corresponds to a bit having a value of 0 in the bit string 230 as a writable block, and starts writing to this block. When it starts writing to the block, the producer 212 sets the value of the bit corresponding to this block in the bit string 230 to 1.

[0121] When it finishes writing, the producer 212 puts the number of the write-completed block into the pointer queue 240, and ends processing.

[0122] When a consumer 212 is starting to read, it refers to the pointer queue 240, deletes the number of the earliest write-completed block from the pointer queue 240, and reads data from the block corresponding to this number. When it finishes reading, it resets the value of

the bit corresponding to this block in the bit string 230 to 0.

[0123] In other words, the bit string 230 indicates whether each block is in the read-completed state or in the other states. Whether put in the pointer queue 240 or not indicates whether a block is in the write-completed state or in the other states. If there is a plurality of blocks in the write-completed state, the order of completion is also indicated.

[0124] In this way, the multiprocessor system 300 also realizes secure multiaccess.

[0125] It will be understood that the buffer managing method used in this third embodiment may be applied not only to a multiprocessor system as shown in Fig. 17, but also to any multiprocessor systems in which data is exchanged between processors, such as the multiprocessor system for transferring stream data shown in Fig. 1.

[0126] Up to this point, the present invention has been described in conjunction with the embodiments thereof. The foregoing embodiments have been given solely by way of illustration. It will be understood by those skilled in the art that a variety of alternate and/or equivalent implementations may be substituted for the specific embodiments shown and described without departing from the scope of the described embodiments. This application is intended to cover any adaptations or variations of the specific embodiments discussed herein.

[0127] For example, in the multiprocessor systems according to the embodiments shown in Figs. 1 and 8, the bit strings are updated and retained in atomic areas on the shared memory so that the buffer on the shared memory is synchronized as well. However, the functions of these bit strings may be implemented as a library, so that the buffer is implemented on the local memories of the producers or the local memories of the consumers.

[0128] Considering that producers and consumers require different information, the shared memory may be omitted in an alternative mode of implementation. For example, in the system according to the embodiment shown in Fig. 1, the producer-required information is the information indicating whether blocks are read-completed or not (in the embodiment shown in Fig. 1, the bit string 30a) and the leading pointer. The consumer-required information is the information indicating whether blocks are write-completed or not (in the embodiment shown in Fig. 1, the bit string 30b) and the following pointer. The leading pointer and the following pointer thus need not be shared between producers and consumers. If there are one producer and one consumer, the producer and the consumer can use message passing techniques to transmit each other's necessary information to each other, thereby sharing information that must be shared. More specifically, when the producer finishes writing, it updates write completion information and transmits the same to the consumer. The consumer ORs the received information and its own write completion information into, for example, a bit string for cumulative update. The same applies to read completion information. When the consumer completes reading, it updates the read completion infor-

mation and transmits the same to the producer. The producer updates the received information and its own read completion information by, for example, using an OR command. Such a mode of implantation can achieve buffer synchronization without using a shared memory for storing bit strings.

[0129] Furthermore, the foregoing embodiments have dealt with examples where the processing entities are processors, and the system itself is configured as a multiprocessor system. However, the buffer management technology according to the present invention is also applicable to multitask systems in which the processing entities are tasks (including processes, threads, and so on).

INDUSTRIAL APPLICABILITY

[0130] As above, the present invention may be applied to electronic equipment which processes a plurality of tasks in parallel, such as a computer, a cellular phone, or a game console.

Claims

1. A method for managing a buffer (120) which is divided into a plurality of blocks, the blocks being used cyclically in a predetermined order by producers and consumers as a temporary storage location for data to be transferred between processing entities, the producers being processing entities for writing data, the consumers being processing entities for reading data written by the producers, the method comprising:

updating and retaining block state information (130a, 130b) including write completed-or-not information (130a) and read completed-or-not information (130b), the write completed-or-not information (130a) indicating whether associated blocks are write-completed or not, the read completed-or-not information (130b) indicating whether associated blocks are read-completed or not;

providing the first block to be written in the predetermined order with a first leading pointer (144), a second leading pointer (145), a first following pointer (148), and a second following pointer (149) which designate the block; and after a block starts to be written, moving the first leading pointer (144) designating the block to a next block,

after the block designated by the second leading pointer (145) finishes being written, moving the second leading pointer (145) to a block next to a block that is farthest from the second leading pointer (145) among consecutive write-completed blocks subsequent to the block designated by the second leading pointer (145),

after a block starts to be read, moving the first following pointer (148) designating the block to a next block, and

after the block designated by the second following pointer (149) finishes being read, moving the second following pointer (149) to a block next to a block that is farthest from the second following pointer (149) among consecutive read-completed blocks subsequent to the block designated by the second following pointer (149), under a constraint that all the pointers (144, 145, 148, 149) move in the same direction, and the first leading pointer (144), the second leading pointer (145), the first following pointer (148), and the second following pointer (149) do not pass each other in this order.

2. The method for managing a buffer (120) according to claim 1, wherein:

the write completed-or-not (130a) information is retained only for as many consecutive blocks equal to the number of producers capable of accessing the buffer (120) at a time, the consecutive blocks starting from the block designated by the second leading pointer (145); and the read completed-or-not information (130b) is retained only for as many consecutive blocks equal to the number of consumers capable of accessing the buffer (120) at a time, the consecutive blocks starting from the block designated by the second following pointer (149).

3. The method for managing a buffer (120) according to claims 1 or 2, wherein:

the processing entities are processors (112); and the block state information (130a, 130b) is updated and/or the pointers are moved by processors (112) that write or read, or processors (112) that have written or read.

4. The method for managing a buffer (120) according to any one of claims 1 to 3, wherein:

the block state information (130a, 130b) is retained in the form of a bitmap in which two bits are assigned to each block; and the bitmap is updated by an atomic operation.

5. An apparatus (200) for managing a buffer (120) which is divided into a plurality of blocks, the blocks being used cyclically in a predetermined order by producers and consumers as a temporary storage location for data to be transferred between processing entities, the producers being processing entities for writing data, the consumers being processing en-

ties for reading data written by the producers, the apparatus comprising:

a block state information retaining part (150) which retains block state information including write completed-or-not information (130a) and read completed-or-not information (130b), the write completed-or-not information (130a) indicating whether associated blocks are write-completed or not, the read completed-or-not information (130b) indicating whether associated blocks are read-completed or not;

a block state information updating part (112) which updates the block state information;

a pointer setting part (112) which provides the first block to be written in the predetermined order with a first leading pointer (144), a second leading pointer (145), a first following pointer (148), and a second following pointer (149) which designate the block; and

a pointer moving part (112) which, after a block starts to be written, moves the first leading pointer (144) designating the block to a next block, after the block designated by the second leading pointer (145) finishes being written, moves the second leading pointer (145) to a block next to a block that is farthest from the second leading pointer (145) among consecutive write-completed blocks subsequent to the block designated by the second leading pointer (145), after a block starts to be read, moves the first following pointer (148) designating the block to a next block, and after the block designated by the second following pointer (149) finishes being read, moves the second following pointer (149) to a block next to a block that is farthest from the second following pointer (149) among consecutive read-completed blocks subsequent to the block designated by the second following pointer (149), under a constraint that all the pointers (144, 145, 148, 149) move in the same direction, and the first leading pointer (144), the second leading pointer (145), the first following pointer (148), and the second following pointer (149) do not pass each other in this order.

- 6. The apparatus (200) for managing a buffer (120) according to claim 5, wherein the block state information retaining part (150) retains the write completed-or-not information (130a) only for as many consecutive blocks equal to the number of producers capable of accessing the buffer (120) at a time, the consecutive blocks starting from the block designated by the second leading pointer (145), and retains the read completed-or-not information (130b) only for as many consecutive blocks equal to the

number of consumers capable of accessing the buffer (120) at a time, the consecutive blocks starting from the block designated by the second following pointer (149).

- 7. The apparatus (200) for managing a buffer (120) according to any one of claims 5 to 6, wherein:

the processing entities are processors (112); and

the block state information updating part (112) and/or the pointer moving part (112) is/are composed of processors (112) that write or read, or processors (112) that have written or read.

- 8. The apparatus (200) for managing a buffer (120) according to any one of claims 5 to 7, wherein:

the block state information retaining part (150) retains the block state information (130a, 130b) in the form of a bitmap in which two bits are assigned to each block; and

the block state updating part (112) updates the bitmap by an atomic operation.

- 9. A computer program product for managing a buffer (120) which is divided into a plurality of blocks, the blocks being used cyclically in a predetermined order by producers and consumers as a temporary storage location for data to be transferred between processing entities, the producers being processing entities for writing data, the consumers being processing entities for reading data written by the producers, the computer program product comprising making a computer perform the functions of:

a program code module for retaining block state information (130a, 130b) to be updated in accordance with progress of writing and reading, the block state information (130a, 130b) associating blocks with whether they are write-completed or not, and whether they are read-completed or not;

a program code module for providing the first block to be written in the predetermined order with a first leading pointer (144), a second leading pointer (145), a first following pointer (148), and a second following pointer (149) which designate the block; and

a program code module for moving the first leading pointer (144) designating a block to a next block, after the block designated by the first leading pointer (144) starts to be written,

a program code module for moving the second leading pointer (145) to a block next to a block that is farthest from the second leading pointer (145) among consecutive write-completed blocks subsequent to the block designated by

the second leading pointer (145), after the block designated by the second leading pointer (145) finishes being written,

a program code module for moving the first following pointer (148) designating a block to a next block, after the block designated by the first following pointer (148) starts to be read, and

a program code module for moving the second following pointer (149) to a block next to a block that is farthest from the second following pointer (149) among consecutive read-completed blocks subsequent to the block designated by the second following pointer (149), after the block designated by the second following pointer (149) finishes being read,

under a constraint that all the pointers (144, 145, 148, 149) move in the same direction, and the first leading pointer (144), the second leading pointer (145), the first following pointer (148), and the second following pointer (149) do not pass each other in this order.

10. A recording medium containing a program for managing a buffer (120) which is divided into a plurality of blocks, the blocks being used cyclically in a predetermined order by producers and consumers as a temporary storage location for data to be transferred between processing entities, the producers being processing entities for writing data, the consumers being processing entities for reading data written by the producers, the program making a computer perform the functions of:

retaining block state information (130a, 130b) to be updated in accordance with progress of writing and reading, the block state information (130a, 130b) associating blocks with whether they are write-completed or not, and whether they are read-completed or not;

providing the first block to be written in the predetermined order with a first leading pointer (144), a second leading pointer (145), a first following pointer (148), and a second following pointer (149) which designate the block; and

after a block starts to be written, moving the first leading pointer (144) designating the block to a next block,

after the block designated by the second leading pointer (145) finishes being written, moving the second leading pointer (145) to a block next to a block that is farthest from the second leading pointer (145) among consecutive write-completed blocks subsequent to the block designated by the second leading pointer (145),

after a block starts to be read, moving the first following pointer (148) designating the block to a next block, and

after the block designated by the second following pointer (149) finishes being read, moving the second following pointer (149) to a block next to a block that is farthest from the second following pointer (149) among consecutive read-completed blocks subsequent to the block designated by the second following pointer (149),

under a constraint that all the pointers (144, 145, 148, 149) move in the same direction, and the first leading pointer (144), the second leading pointer (145), the first following pointer (148), and the second following pointer (149) do not pass each other in this order.

Patentansprüche

1. Verfahren zur Verwaltung eines Puffers (120), der in mehrere Blöcke unterteilt ist, wobei die Blöcke zyklisch in einer vorbestimmten Reihenfolge von Erzeugern und Verbrauchern als ein temporärer Speicherort für Daten, die zwischen Verarbeitungsentitäten übertragen werden sollen, verwendet werden, wobei die Erzeuger Verarbeitungsentitäten zum Schreiben von Daten sind, die Verbraucher Verarbeitungsentitäten zum Lesen von Daten sind, die von den Erzeugern geschrieben werden, wobei das Verfahren Folgendes umfasst:

Aktualisieren und Halten von Blockzustandsinformationen (130a, 130b) einschließlich Schreiben-beendet-oder-nicht-Informationen (130a) und Lesen-beendet-oder-nicht-Informationen (130b), wobei die Schreiben-beendet-oder-nicht-Informationen (130a) angeben, ob assoziierte Blöcke Schreiben-beendet sind oder nicht, und die Lesen-beendet-oder-nicht-Informationen (130b) angeben, ob assoziierte Blöcke Lesen-beendet sind oder nicht;

Versehen des ersten Blocks, der in der vorbestimmten Reihenfolge geschrieben werden soll, mit einem ersten führenden Zeiger (144), einem zweiten führenden Zeiger (145), einem ersten folgenden Zeiger (148) und einem zweiten folgenden Zeiger (149), die den Block designieren; und,

nachdem das Schreiben eines Blocks gestartet wird, Bewegen des ersten führenden Zeigers (144), der den Block bezeichnet, zu einem nächsten Block,

nachdem das Schreiben des Blocks, der von dem zweiten führenden Zeiger (145) designiert wird, abgeschlossen wird, Bewegen des zweiten führenden Zeigers (145) zu einem Block neben einem Block, der sich am weitesten entfernt vom zweiten führenden Zeiger (145) befindet, unter aufeinanderfolgenden Schreiben-beendet-Blöcken anschließend an den Block, der von

- dem zweiten führenden Zeiger (145) designiert wird,
 nachdem das Lesen eines Blocks gestartet wird,
 Bewegen des ersten folgenden Zeigers (148),
 der den Block designiert, zu einem nächsten Block, und,
 nachdem das Lesen des Blocks, der von dem
 zweiten folgenden Zeiger (149) designiert wird,
 abgeschlossen wird, Bewegen des zweiten folgen-
 5 genden Zeigers (149) zu einem Block neben einem Block,
 der sich am weitesten entfernt vom
 zweiten folgenden Zeiger (149) befindet, unter
 10 aufeinanderfolgenden Lesen-beendet-Blöcken
 anschließend an den Block, der von dem zweiten
 folgenden Zeiger (149) designiert wird,
 15 unter einer Bedingung, dass sich alle Zeiger
 (144, 145, 148, 149) in die gleiche Richtung be-
 wegen und der erste führende Zeiger (144), der
 zweite führende Zeiger (145), der erste folgende
 20 Zeiger (148) und der zweite folgende Zeiger
 (149) einander nicht in dieser Reihenfolge pas-
 sieren.
2. Verfahren zur Verwaltung eines Puffers (120) nach
 Anspruch 1, wobei:
 25 die Schreiben-beendet-oder-nicht-Informationen
 (130a) nur für so viele aufeinanderfolgende
 Blöcke gleich der Anzahl von Erzeugern, die in
 der Lage sind, zu einer Zeit auf den Puffer (120)
 30 zuzugreifen, gehalten werden, wobei die aufei-
 nanderfolgenden Blöcke von dem Block starten,
 der von dem zweiten führenden Zeiger (145) de-
 signiert wird; und
 35 die Lesen-beendet-oder-nicht-Informationen
 (130b) nur so viele aufeinanderfolgende Blöcke
 gleich der Anzahl von Verbrauchern, die in der
 Lage sind, zu einer Zeit auf den Puffer (120)
 40 zuzugreifen, gehalten werden, wobei die aufei-
 nanderfolgenden Blöcke von dem Block starten,
 der von dem zweiten folgenden Zeiger (149) de-
 signiert wird.
3. Verfahren zur Verwaltung eines Puffers (120) nach
 den Ansprüchen 1 oder 2, wobei:
 45 die Verarbeitungsentitäten Prozessoren (112)
 sind und
 die Blockzustandsinformationen (130a, 130b)
 durch die Prozessoren (112), die schreiben oder
 50 lesen, oder Prozessoren (112), die geschrieben
 oder gelesen haben, aktualisiert werden
 und/oder die Zeiger durch diese bewegt werden.
4. Verfahren zur Verwaltung eines Puffers (120) nach
 einem der Ansprüche 1 bis 3, wobei:
 55 die Blockzustandsinformationen (130a, 130b) in
- Form eines Bitmaps aufbewahrt werden, bei
 dem jedem Block zwei Bits zugewiesen werden;
 und
 das Bitmap durch eine atomare Operation aktu-
 5 alisiert wird.
5. Vorrichtung (200) zur Verwaltung eines Puffers
 (120), der in mehrere Blöcke unterteilt ist, wobei die
 Blöcke zyklisch in einer vorbestimmten Reihenfolge
 von Erzeugern und Verbrauchern als ein temporärer
 10 Speicherort für Daten, die zwischen Verarbeitungs-
 entitäten übertragen werden sollen, verwendet wer-
 den, wobei die Erzeuger Verarbeitungsentitäten
 zum Schreiben von Daten sind, die Verbraucher Ver-
 15 arbeitungsentitäten zum Lesen von Daten sind, die
 von den Erzeugern geschrieben werden, wobei die
 Vorrichtung Folgendes umfasst:
- einen Blockzustandsinformationen-Halteteil
 (150), der Blockzustandsinformationen einschließ-
 20 lich Schreiben-beendet-oder-nicht-
 Informationen (130a) und Lesen-beendet-oder-
 nicht-Informationen (130b) hält, wobei die
 Schreiben-beendet-oder-nicht-Informationen
 (130a) angeben, ob assoziierte Blöcke Schrei-
 25 ben-beendet sind oder nicht, und die Lesen-be-
 endet-oder-nicht-Informationen (130b) ange-
 ben, ob assoziierte Blöcke Lesen-beendet sind
 oder nicht;
 30 einen Blockzustandsinformationen-Aktualisie-
 rungsteil (112), der die Blockzustandsinformati-
 onen aktualisiert;
 35 einen Zeigereinstellungsteil (112), der den ers-
 ten Block, der in der vorbestimmten Reihenfolge
 geschrieben werden soll, mit einem ersten füh-
 40 renden Zeiger (144), einem zweiten führenden
 Zeiger (145), einem ersten folgenden Zeiger
 (148) und einem zweiten folgenden Zeiger
 (149), die den Block designieren, versieht; und
 einen Zeigerbewegungsteil (112), der, nach-
 45 dem das Schreiben eines Blocks gestartet wird,
 den ersten führenden Zeiger (144), der den
 Block designiert, zu einem nächsten Block be-
 50 wegt,
 nachdem das Schreiben des Blocks, der von
 dem zweiten führenden Zeiger (145) designiert
 wird, abgeschlossen wird, den zweiten führen-
 55 den Zeiger (145) zu einem Block neben einem
 Block, der sich am weitesten entfernt vom zwei-
 ten führenden Zeiger (145) befindet, unter auf-
 einanderfolgenden Schreiben-beendet-Blö-
 cken anschließend an den Block, der von dem
 zweiten führenden Zeiger (145) designiert wird,
 bewegt,
 nachdem das Lesen eines Blocks gestartet wird,
 den ersten folgenden Zeiger (148), der den
 Block designiert, zu einem nächsten Block be-
 60 wegt, und,

nachdem das Lesen des Blocks, der von dem zweiten folgenden Zeiger (149) designiert wird, abgeschlossen wird, den zweiten folgenden Zeiger (149) zu einem Block neben einem Block, der sich am weitesten entfernt vom zweiten folgenden Zeiger (149) befindet, unter aufeinanderfolgenden Lesen-beendet-Blöcken anschließend an den Block, der von dem zweiten folgenden Zeiger (149) designiert wird, bewegt, unter einer Bedingung, dass sich alle Zeiger (144, 145, 148, 149) in die gleiche Richtung bewegen und der erste führende Zeiger (144), der zweite führende Zeiger (145), der erste folgende Zeiger (148) und der zweite folgende Zeiger (149) einander nicht in dieser Reihenfolge passieren.

6. Vorrichtung (200) zur Verwaltung eines Puffers (120) nach Anspruch 5, wobei:

der Blockzustandsinformationen-Halteteil (150) die Schreiben-beendet-oder-nicht-Informationen (130a) nur für so viele aufeinanderfolgende Blöcke gleich der Anzahl von Erzeugern, die in der Lage sind, zu einer Zeit auf den Puffer (120) zuzugreifen, hält, wobei die aufeinanderfolgenden Blöcke von dem Block starten, der von dem zweiten führenden Zeiger (145) designiert wird, und die Lesen-beendet-oder-nicht-Informationen (130b) nur für so viele aufeinanderfolgende Blöcke gleich der Anzahl von Verbrauchern, die in der Lage sind, zu einer Zeit auf den Puffer (120) zuzugreifen, hält, wobei die aufeinanderfolgenden Blöcke von dem Block starten, der von dem zweiten folgenden Zeiger (149) designiert wird.

7. Vorrichtung (200) zur Verwaltung eines Puffers (120) nach einem der Ansprüche 5 bis 6, wobei:

die Verarbeitungsentitäten Prozessoren (112) sind und der Blockzustandsinformationen-Aktualisierungsteil (112) und/oder der Zeigerbewegungsteil (112) aus Prozessoren (112), die schreiben oder lesen, oder Prozessoren (112), die geschrieben oder gelesen haben, bestehen.

8. Vorrichtung (200) zur Verwaltung eines Puffers (120) nach einem der Ansprüche 5 bis 7, wobei:

der Blockzustandsinformationen-Halteteil (150) die Blockzustandsinformationen (130a, 130b) in Form eines Bitmaps hält, bei dem jedem Block zwei Bits zugewiesen werden; und der Blockzustand-Aktualisierungsteil (112) das Bitmap durch eine atomare Operation aktualisiert.

9. Computerprogrammprodukt zur Verwaltung eines Puffers (120), der in mehrere Blöcke unterteilt ist, wobei die Blöcke zyklisch in einer vorbestimmten Reihenfolge von Erzeugern und Verbrauchern als ein temporärer Speicherort für Daten, die zwischen Verarbeitungsentitäten übertragen werden sollen, verwendet werden, wobei die Erzeuger Verarbeitungsentitäten zum Schreiben von Daten sind, die Verbraucher Verarbeitungsentitäten zum Lesen von Daten sind, die von den Erzeugern geschrieben werden, wobei das Computerprogrammprodukt umfasst, einen Computer zum Durchführen der folgenden Funktionen zu veranlassen:

ein Programmcodemodul zum Halten von Blockzustandsinformationen (130a, 130b), die gemäß einem Schreib- und Lesefortschritt aktualisiert werden sollen, wobei die Blockzustandsinformationen (130a, 130b) Blöcke damit assoziieren, ob sie Schreiben-beendet sind oder nicht, und, ob sie Lesen-beendet sind oder nicht;

ein Programmcodemodul zum Versehen des ersten Blocks, der in der vorbestimmten Reihenfolge geschrieben werden soll, mit einem ersten führenden Zeiger (144), einem zweiten führenden Zeiger (165), einem ersten folgenden Zeiger (148) und einem zweiten folgenden Zeiger (149), die den Block designieren; und

ein Programmcodemodul zum Bewegen des ersten führenden Zeigers (144), der einen Block designiert, zu einem nächsten Block, nachdem das Schreiben des Blocks, der von dem ersten führenden Zeiger (144) designiert wird, gestartet wird,

ein Programmcodemodul zum Bewegen des zweiten führenden Zeigers (145) zu einem Block neben einem Block, der sich am weitesten entfernt vom zweiten führenden Zeiger (145) befindet, unter aufeinanderfolgenden Schreiben-beendet-Blöcken anschließend an den Block, der von dem zweiten führenden Zeiger (145) designiert wird, nachdem das Schreiben des Blocks, der von dem zweiten führenden Zeiger (145) designiert wird, abgeschlossen wird,

ein Programmcodemodul zum Bewegen des ersten folgenden Zeigers (148), der einen Block designiert, zu einem nächsten Block, nachdem das Lesen des Blocks, der von dem ersten folgenden Zeiger (148) designiert wird, gestartet wird, und

ein Programmcodemodul zum Bewegen des zweiten folgenden Zeigers (149) zu einem Block neben einem Block, der sich am weitesten entfernt vom zweiten folgenden Zeiger (149) befindet, unter aufeinanderfolgenden Lesen-beendet-Blöcken anschließend an den Block, der von dem zweiten folgenden Zeiger (149) designiert

wird, nachdem das Lesen des Blocks, der von dem zweiten folgenden Zeiger (149) designiert wird, abgeschlossen wird,
 unter einer Bedingung, dass sich alle Zeiger (144, 145, 148, 149) in die gleiche Richtung bewegen und der erste führende Zeiger (144), der zweite führende Zeiger (145), der erste folgende Zeiger (148) und der zweite folgende Zeiger (149) einander nicht in dieser Reihenfolge passieren.

10. Aufzeichnungsmedium, das ein Programm zur Verwaltung eines Puffers (120) enthält, der in mehrere Blöcke unterteilt ist, wobei die Blöcke zyklisch in einer vorbestimmten Reihenfolge von Erzeugern und Verbrauchern als ein temporärer Speicherort für Daten, die zwischen Verarbeitungsentitäten übertragen werden sollen, verwendet werden, wobei die Erzeuger Verarbeitungsentitäten zum Schreiben von Daten sind, die Verbraucher Verarbeitungsentitäten zum Lesen von Daten sind, die von den Erzeugern geschrieben werden,
 wobei das Programm einen Computer zum Durchführen der folgenden Funktionen veranlasst:

Halten von Blockzustandsinformationen (130a, 130b), die gemäß einem Schreib- und Lesefortschritt aktualisiert werden sollen, wobei die Blockzustandsinformationen (130a, 130b) Blöcke damit assoziieren, ob sie Schreiben-beendet sind oder nicht, und, ob sie Lesen-beendet sind oder nicht;

Versehen des ersten Blocks, der in der vorbestimmten Reihenfolge geschrieben werden soll, mit einem ersten führenden Zeiger (144), einem zweiten führenden Zeiger (145), einem ersten folgenden Zeiger (148) und einem zweiten folgenden Zeiger (149), die den Block designieren; und,

nachdem das Schreiben eines Blocks gestartet wird, Bewegen des ersten führenden Zeigers (144), der den Block designiert, zu einem nächsten Block,

nachdem das Schreiben des Blocks, der von dem zweiten führenden Zeiger (145) designiert wird, abgeschlossen wird, Bewegen des zweiten führenden Zeigers (145) zu einem Block neben einem Block, der sich am weitesten entfernt vom zweiten führenden Zeiger (145) befindet, unter aufeinanderfolgenden Schreiben-beendet-Blöcken anschließend an den Block, der von dem zweiten führenden Zeiger (145) designiert wird,

nachdem das Lesen eines Blocks gestartet wird, Bewegen des ersten folgenden Zeigers (148), der den Block designiert, zu einem nächsten Block, und,
 nachdem das Lesen des Blocks, der von dem

zweiten folgenden Zeiger (149) designiert wird, abgeschlossen wird, Bewegen des zweiten folgenden Zeigers (149) zu einem Block neben einem Block, der sich am weitesten entfernt vom zweiten folgenden Zeiger (149) befindet, unter aufeinanderfolgenden Lesen-beendet-Blöcken anschließend an den Block, der von dem zweiten folgenden Zeiger (149) designiert wird,
 unter einer Bedingung, dass sich alle Zeiger (144, 145, 148, 149) in die gleiche Richtung bewegen und der erste führende Zeiger (144), der zweite führende Zeiger (145), der erste folgende Zeiger (148) und der zweite folgende Zeiger (149) einander nicht in dieser Reihenfolge passieren.

Revendications

1. Procédé pour gérer une mémoire tampon (120) qui est divisée en une pluralité de blocs, les blocs étant utilisés de manière cyclique dans un ordre prédéterminé par des producteurs et des consommateurs comme emplacement de stockage temporaire pour des données qui doivent être transférées entre des entités de traitement, les producteurs étant des entités de traitement pour écrire des données, les consommateurs étant des entités de traitement pour lire des données écrites par les producteurs, le procédé consistant :

à mettre à jour et à conserver des informations d'état de bloc (130a, 130b) comportant des informations d'écriture terminée ou non (130a) et des informations de lecture terminée ou non (130b), les informations d'écriture terminée ou non (130a) indiquant si des blocs associés sont des blocs d'écriture terminée ou non, les informations de lecture terminée ou non (130b) indiquant si des blocs associés sont des blocs de lecture terminée ou non ;

à fournir, au premier bloc qui doit être écrit dans l'ordre prédéterminé, un premier pointeur de tête (144), un second pointeur de tête (145), un premier pointeur suivant (148) et un second pointeur suivant (149) qui désignent le bloc ; et après qu'un bloc commence à être écrit, à déplacer le premier pointeur de tête (144) désignant le bloc jusqu'à un prochain bloc, après que le bloc désigné par le second pointeur de tête (145) finit d'être écrit, à déplacer le second pointeur de tête (145) jusqu'à un bloc à côté d'un bloc qui est le plus éloigné du second pointeur de tête (145) parmi des blocs d'écriture terminée consécutifs après le bloc désigné par le second pointeur de tête (145), après qu'un bloc commence à être lu, à déplacer le premier pointeur suivant (148) désignant le

- bloc jusqu'à un prochain bloc et après que le bloc désigné par le second pointeur suivant (149) finit d'être lu, à déplacer le second pointeur suivant (149) jusqu'à un bloc à côté d'un bloc qui est le plus éloigné du second pointeur suivant (149) parmi des blocs de lecture terminée consécutifs après le bloc désigné par le second pointeur suivant (149), sous une contrainte que tous les pointeurs (144, 145, 148, 149) se déplacent dans la même direction et que le premier pointeur de tête (144), le second pointeur de tête (145), le premier pointeur suivant (148) et le second pointeur suivant (149) ne se dépassent pas les uns les autres dans cet ordre.
2. Procédé pour gérer une mémoire tampon (120) selon la revendication 1, dans lequel :
- les informations d'écriture terminée ou non (130a) sont conservées seulement pour autant de blocs consécutifs que le nombre de producteurs capables d'avoir accès à la mémoire tampon (120) en une fois, les blocs consécutifs commençant à partir du bloc désigné par le second pointeur de tête (145) ; et les informations de lecture terminée ou non (130b) sont conservées seulement pour autant de blocs consécutifs que le nombre de consommateurs capables d'avoir accès à la mémoire tampon (120) en une fois, les blocs consécutifs commençant à partir du bloc désigné par le second pointeur suivant (149).
3. Procédé pour gérer une mémoire tampon (120) selon les revendications 1 ou 2, dans lequel :
- les entités de traitement sont des processeurs (112) ; et les informations d'état de bloc (130a, 130b) sont mises à jour et/ou les pointeurs sont déplacés par des processeurs (112) qui écrivent ou lisent, ou par des processeurs (112) qui ont écrit ou lu.
4. Procédé pour gérer une mémoire tampon (120) selon l'une quelconque des revendications 1 à 3, dans lequel :
- les informations d'état de bloc (130a, 130b) sont conservées sous la forme d'une table de bits dans laquelle deux bits sont attribués à chaque bloc ; et la table de bits est mise à jour au moyen d'une opération atomique.
5. Appareil (200) pour gérer une mémoire tampon (120) qui est divisée en une pluralité de blocs, les blocs étant utilisés de manière cyclique dans un ordre pré-
- déterminé par des producteurs et des consommateurs comme emplacement de stockage temporaire pour des données qui doivent être transférées entre des entités de traitement, les producteurs étant des entités de traitement pour écrire des données, les consommateurs étant des entités de traitement pour lire des données écrites par les producteurs, le procédé consistant :
- une partie de conservation d'informations d'état de bloc (150) qui conserve des informations d'état de bloc comportant des informations d'écriture terminée ou non (130a) et des informations de lecture terminée ou non (130b), les informations d'écriture terminée ou non (130a) indiquant si des blocs associés sont des blocs d'écriture terminée ou non, les informations de lecture terminée ou non (130b) indiquant si des blocs associés sont des blocs de lecture terminée ou non ;
- une partie de mise à jour d'informations d'état de bloc (112) qui met à jour les informations d'état de bloc ;
- une partie de réglage de pointeur (112) qui fournit, au premier bloc qui doit être écrit dans l'ordre prédéterminé, un premier pointeur de tête (144), un second pointeur de tête (145), un premier pointeur suivant (148) et un second pointeur suivant (149) qui désignent le bloc ; et
- une partie de déplacement de pointeur (112) qui, après qu'un bloc commence à être écrit, déplace le premier pointeur de tête (144) désignant le bloc jusqu'à un prochain bloc, après que le bloc désigné par le second pointeur de tête (145) finit d'être écrit, déplace le second pointeur de tête (145) jusqu'à un bloc à côté d'un bloc qui est le plus éloigné du second pointeur de tête (145) parmi des blocs d'écriture terminée consécutifs après le bloc désigné par le second pointeur de tête (145), après qu'un bloc commence à être lu, déplace le premier pointeur suivant (148) désignant le bloc jusqu'à un prochain bloc et après que le bloc désigné par le second pointeur suivant (149) finit d'être lu, déplace le second pointeur suivant (149) jusqu'à un bloc à côté d'un bloc qui est le plus éloigné du second pointeur suivant (149) parmi des blocs de lecture terminée consécutifs après le bloc désigné par le second pointeur suivant (149), sous une contrainte que tous les pointeurs (144, 145, 148, 149) se déplacent dans la même direction et que le premier pointeur de tête (144), le second pointeur de tête (145), le premier pointeur suivant (148) et le second pointeur suivant (149) ne se dépassent pas les uns les autres dans cet ordre.

6. Appareil (200) pour gérer une mémoire tampon (120) selon la revendication 5, dans lequel :

la partie de conservation d'informations d'état de bloc (150) conserve les informations d'écriture terminée ou non (130a) seulement pour autant de blocs consécutifs que le nombre de producteurs capables d'avoir accès à la mémoire tampon (120) en une fois, les blocs consécutifs commençant à partir du bloc désigné par le second pointeur de tête (145) ; et
 conserve les informations de lecture terminée ou non (130b) seulement pour autant de blocs consécutifs que le nombre de consommateurs capables d'avoir accès à la mémoire tampon (120) en une fois, les blocs consécutifs commençant à partir du bloc désigné par le second pointeur suivant (149).

7. Appareil (200) pour gérer une mémoire tampon (120) selon l'une quelconque des revendications 5 à 6, dans lequel :

les entités de traitement sont des processeurs (112) ; et
 la partie de mise à jour d'informations d'état de bloc (112) et/ou la partie de déplacement de pointeur (112) sont composées de processeurs (112) qui écrivent ou lisent, ou de processeurs (112) qui ont écrit ou lu.

8. Appareil (200) pour gérer une mémoire tampon (120) selon l'une quelconque des revendications 5 à 7, dans lequel :

la partie de conservation d'informations d'état de bloc (150) conserve les informations d'état de bloc (130a, 130b) sous la forme d'une table de bits dans laquelle deux bits sont attribués à chaque bloc ; et
 la partie de mise à jour d'état de bloc (112) met à jour table de bits au moyen d'une opération atomique.

9. Produit-programme d'ordinateur pour gérer une mémoire tampon (120) qui est divisée en une pluralité de blocs, les blocs étant utilisés de manière cyclique dans un ordre prédéterminé par des producteurs et des consommateurs comme emplacement de stockage temporaire pour des données qui doivent être transférées entre des entités de traitement, les producteurs étant des entités de traitement pour écrire des données, les consommateurs étant des entités de traitement pour lire des données écrites par les producteurs, le produit-programme d'ordinateur consistant à faire qu'un ordinateur réalise les fonctions :

d'un module de code de programme pour conserver des informations d'état de bloc (130a, 130b) qui doivent être mises à jour en fonction de la progression de l'écriture et de la lecture, les informations d'état de bloc (130a, 130b) associant des blocs au fait de savoir s'ils sont des blocs d'écriture terminée ou non, et s'ils sont des blocs de lecture terminée ou non ;

d'un module de code programme pour fournir, au premier bloc qui doit être écrit dans l'ordre prédéterminé, un premier pointeur de tête (144), un second pointeur de tête (145), un premier pointeur suivant (148) et un second pointeur suivant (149) qui désignent le bloc ; et

d'un module de code programme pour déplacer le premier pointeur de tête (144) désignant le bloc jusqu'à un prochain bloc, après que le bloc désigné par le premier pointeur de tête (144) commence à être écrit,

d'un module de code de programme pour déplacer le second pointeur de tête (145) jusqu'à un bloc à côté d'un bloc qui est le plus éloigné du second pointeur de tête (145) parmi des blocs d'écriture terminée consécutifs après le bloc désigné par le second pointeur de tête (145), après que le bloc désigné par le second pointeur de tête (145) finit d'être écrit,

d'un module de code de programme pour déplacer le premier pointeur suivant (148) désignant un bloc jusqu'à un prochain bloc, après que le bloc désigné par le premier pointeur suivant (148) commence à être lu, et

d'un module de code de programme pour déplacer le second pointeur suivant (149) jusqu'à un bloc à côté d'un bloc qui est le plus éloigné du second pointeur suivant (149) parmi des blocs de lecture terminée consécutifs après le bloc désigné par le second pointeur suivant (149), après que le bloc désigné par le second pointeur suivant (149) finit d'être lu,

sous une contrainte que tous les pointeurs (144, 145, 148, 149) se déplacent dans la même direction et que le premier pointeur de tête (144), le second pointeur de tête (145), le premier pointeur suivant (148) et le second pointeur suivant (149) ne se dépassent pas les uns les autres dans cet ordre.

10. Support d'enregistrement contenant un programme pour gérer une mémoire tampon (120) qui est divisée en une pluralité de blocs, les blocs étant utilisés de manière cyclique dans un ordre prédéterminé par des producteurs et des consommateurs comme emplacement de stockage temporaire pour des données qui doivent être transférées entre des entités de traitement, les producteurs étant des entités de traitement pour écrire des données, les consommateurs étant des entités de traitement pour lire des

données écrites par les producteurs,
le programme faisant qu'un ordinateur réalise les
fonctions consistant :

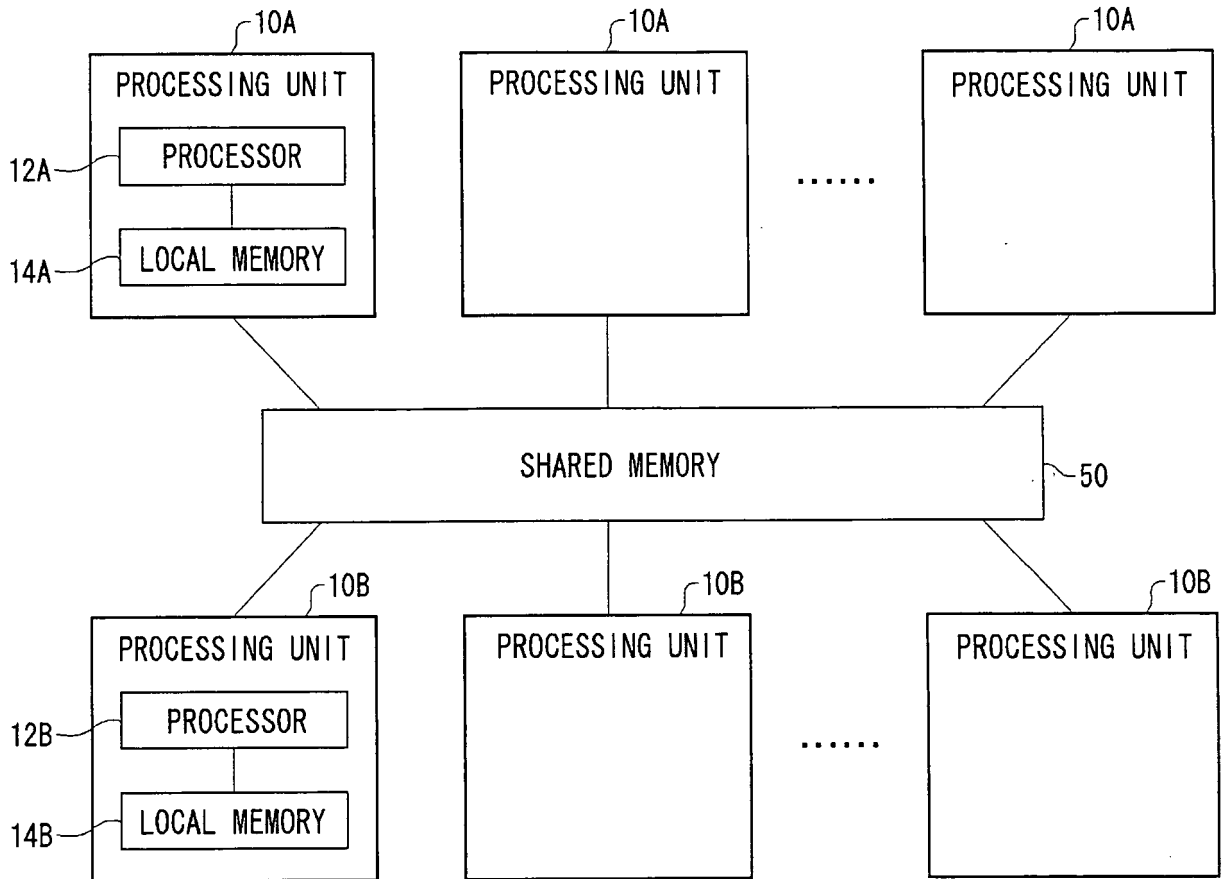
à conserver des informations d'état de bloc 5
(130a, 130b) qui doivent être mises à jour en
fonction de la progression de l'écriture et de la
lecture, les informations d'état de bloc (130a,
130b) associant des blocs au fait de savoir s'ils 10
sont des blocs d'écriture terminée ou non, et s'ils
sont des blocs de lecture terminée ou non ;
à fournir, au premier bloc qui doit être écrit dans
l'ordre prédéterminé, un premier pointeur de tête (144), un second pointeur de tête (145), un 15
premier pointeur suivant (148) et un second
pointeur suivant (140) qui désignent le bloc; et
après qu'un bloc commence à être écrit, à dé-
placer le premier pointeur de tête (144) dési-
gnant le bloc jusqu'à un prochain bloc,
après que le bloc désigné par le second pointeur 20
de tête (145) finit d'être écrit, à déplacer le se-
cond pointeur de tête (145) jusqu'à un bloc à
côté d'un bloc qui est le plus éloigné du second
pointeur de tête (145) parmi des blocs d'écriture 25
terminée consécutifs après le bloc désigné par
le second pointeur de tête (145),
après qu'un bloc commence à être lu, à déplacer
le premier pointeur suivant (148) désignant le
bloc jusqu'à un prochain bloc et
après que le bloc désigné par le second pointeur 30
suivant (149) finit d'être lu, à déplacer le second
pointeur suivant (149) jusqu'à un bloc à côté
d'un bloc qui est le plus éloigné du second poin-
teur suivant (149) parmi des blocs de lecture 35
terminée consécutifs après le bloc désigné par
le second pointeur suivant (149),
sous une contrainte que tous les pointeurs (144,
145, 148, 149) se déplacent dans la même di-
rection et que le premier pointeur de tête (144), 40
le second pointeur de tête (145), le premier poin-
teur suivant (148) et le second pointeur suivant
(149) ne se dépassent pas les uns les autres
dans cet ordre.

45

50

55

FIG. 1



100

FIG. 2




FIG. 3




FIG. 4




FIG. 5




FIG. 6




FIG. 7

STATE \ BITMAP	INITIAL	WRITING	WRITE-COMPLETED	READING	READ-COMPLETED
30a	1	0	0	0	1
30b	0	0	1	0	0

FIG. 8




FIG. 9




FIG. 10




FIG. 11




FIG. 12




FIG. 13




FIG. 14




FIG. 15




FIG. 16




FIG. 17





FIG. 18



REFERENCES CITED IN THE DESCRIPTION

This list of references cited by the applicant is for the reader's convenience only. It does not form part of the European patent document. Even though great care has been taken in compiling the references, errors or omissions cannot be excluded and the EPO disclaims all liability in this regard.

Patent documents cited in the description

- US 6173307 B1 [0003]