

PATENT ASSIGNMENT COVER SHEET

Electronic Version v1.1
Stylesheet Version v1.2

EPAS ID: PAT7211005

SUBMISSION TYPE:	NEW ASSIGNMENT
NATURE OF CONVEYANCE:	ASSIGNMENT
CONVEYING PARTY DATA	
Name	Execution Date
STUART JOHN INGLIS	08/19/2019
TIMOTHY KELLY DAWSON	08/19/2019
SHERIDAN JOHN LAMBERT	08/19/2019
RECEIVING PARTY DATA	
Name:	NYRIAD LIMITED
Street Address:	LEWIS' CNR DICK & ALPHA STREETS
City:	CAMBRIDGE
State/Country:	NEW ZEALAND
Postal Code:	3434
PROPERTY NUMBERS Total: 1	
Property Type	Number
Application Number:	16544605
CORRESPONDENCE DATA	
Fax Number:	
<i>Correspondence will be sent to the e-mail address first; if that is unsuccessful, it will be sent using a fax number, if provided; if that is unsuccessful, it will be sent via US Mail.</i>	
Phone:	9727312288
Email:	dallaspatents@dfw.conleyrose.com
Correspondent Name:	CONLEY ROSE, P.C.
Address Line 1:	4965 PRESTON PARK BOULEVARD
Address Line 2:	SUITE 195E
Address Line 4:	PLANO, TEXAS 75093
ATTORNEY DOCKET NUMBER:	4592-01100
NAME OF SUBMITTER:	MICHAEL W. PIPER
SIGNATURE:	/Michael W. Piper/
DATE SIGNED:	03/07/2022
Total Attachments: 47	
source=4592-01100 NYR040P13468US01 Assignment#page1.tif	
source=4592-01100 NYR040P13468US01 Assignment#page2.tif	
source=4592-01100 NYR040P13468US01 Assignment#page3.tif	

[illegible]

Stuart John Inglis
Sheridan John Lambert
Timothy Kelly Dawson

Nyriad Limited

**ASSIGNMENT OF
INTELLECTUAL PROPERTY RIGHTS**

PARTIES

Stuart John Inglis, a New Zealand citizen care of Nyriad Limited, LEWIS', Cnr Dick & Alpha Streets, Cambridge, 3434, New Zealand

Sheridan John Lambert, a New Zealand citizen care of Nyriad Limited, LEWIS', Cnr Dick & Alpha Streets, Cambridge, 3434, New Zealand

Timothy Kelly Dawson, a New Zealand citizen care of Nyriad Limited, LEWIS', Cnr Dick & Alpha Streets, Cambridge, 3434, New Zealand

(each individually the Assignor)

Nyriad Limited, a New Zealand company having its registered office at LEWIS', Cnr Dick & Alpha Streets, Cambridge, 3434, New Zealand

(the Assignee)

INTRODUCTION

- A. The Assignor has devised or contributed to the Invention during employment by or under a commission from the Assignee.
- B. The Assignor acknowledges that the Assignee is or should be the legal and beneficial owner of the Invention and the Intellectual Property Rights.
- C. The Assignor has agreed to assign the Invention and the Intellectual Property Rights to the Assignee subject to the terms and conditions of this assignment.

COVENANTS

1. DEFINITIONS

- 1.1 In this assignment, including the Introduction, the following terms will have the following meanings:

Copyright means:

- (a) all copyrights and all rights in the nature of copyright, in any original artistic, literary and other works; and
- (b) any database rights,

comprising or relating to the Invention and/or the Design, that exist or may in the future exist anywhere in the world;

Design means the designs described in the Schedule (as improved, modified, developed or amended at any time before execution of this assignment by the Assignor);

Design Rights means all rights in and to the designs to be applied to articles of, or relating to, the Invention and/or the Design, that exist or may in the future exist anywhere in the world including, but not limited to:

- (a) the right to apply for and obtain protection for such designs in relation to such articles and the rights conferred by such protection when granted;

- (b) the right to claim priority under any international convention or agreement, including the Paris Convention (as amended), from any such application referred to in paragraph (a) above;
- (c) all unregistered design rights arising from the Invention and/or the Design; and
- (d) all semi-conductor topography or integrated circuit layout rights arising from the Invention and/or the Design;

Improvement means an improvement in, modification of, or addition to:

- (a) the Invention and/or the Design;
- (b) any original artistic, literary or other works relating to the Invention and/or the Design; or
- (c) any designs to be applied to articles of or relating to the Invention and/or the Design,

devised, created, designed, contributed to or acquired by the Assignor, while the Assignor is employed by or under any commission from the Assignee;

Intellectual Property Rights means all industrial and intellectual property rights (whether protectable by statute, at common law or in equity) in and to the Invention and/or the Design that exist or may in the future exist anywhere in the world and whether or not registered or registrable, including the Patent Rights, Copyright and Design Rights;

Invention means any and all inventions described in the Schedule (as improved, modified, developed or amended at any time before execution of this assignment by the Assignor);

Patent Applications means the patent applications identified in the Schedule; and

Patent Rights means all patent rights in and to the Invention and/or the Design that exist or may in the future exist anywhere in the world including, but not limited to:

- (a) the right to apply for and obtain patents or other similar forms of protection for the Invention and/or the Design in any country;
- (b) the right to claim priority under any international convention or agreement, including the Paris Convention (as amended) from any application referred to in paragraph (a) above;
- (c) all rights conferred by any patents or similar forms of protection obtained from any applications referred to in paragraphs (a) and (b) above and
- (d) the Patent Applications together with:
 - (i) any patents that may be granted pursuant to the Patent Applications; and
 - (ii) any future patent(s) and patent application(s) that are based on or derive priority from or have equivalent claims to the Patent Applications in any country in the world (including divisionals, continuations, continuations in part, patents of addition, supplementary protection certificates, reissues, extensions, innovation patents, utility models and petty patents).

2. ASSIGNMENT

2.1 **Assignment:** In consideration of the sum of NZ\$1.00 paid by the Assignee to the Assignor, the receipt and sufficiency of which is hereby acknowledged by the Assignor, the Assignor assigns to the Assignee absolutely all of the Assignor's rights, title and interest in and to:

- (a) the Invention;
- (b) the Design; and
- (c) the Intellectual Property Rights.

2.2 **Rights of action:** The assignments effected by clause 2.1 include:

- (a) the assignment and transfer of:
 - (i) the right to sue for damages for infringement or misuse of the Intellectual Property Rights; and
 - (ii) all other rights of action, powers, legal and equitable remedies, and benefits arising from ownership of the Intellectual Property Rights,
- in relation to all causes of action arising before, on or after the date of this assignment.

2.3 **Moral rights:** The Assignor:

- (a) waives all of the Assignor's moral rights arising from the Invention and the Design throughout the world; and
- (b) consents to all uses and treatments of the Invention, the Design, and the Intellectual Property Rights, to the extent that the Assignor may lawfully do so.

2.4 **Authority to add application particulars:**

- (a) The Assignor acknowledges that it is the intention of the Assignee to file one or more patent and/or design applications in various countries around the world in respect of the Invention and/or the Design. The Assignor hereby irrevocably authorises the Assignee or its agent to add in the Schedule particulars of any patent or design applications in respect of the Invention or Design that may be filed following the date of execution by the Assignor of this assignment.
- (b) The Assignor agrees and acknowledges that any addition under clause 2.4(a) will be valid and binding upon the Assignor as if the addition had been made prior to execution by the Assignor of this assignment.

3. CONFIDENTIALITY

3.1 **Confidentiality obligations:** The Assignor agrees to treat all information relating to the Invention and/or the Design and the Intellectual Property Rights as secret and confidential. The Assignor will not use, disclose or publish such information without the Assignee's prior written consent.

3.2 **Exceptions to confidentiality:** The obligations of confidentiality in clause 3.1 do not extend to any information that is or becomes generally available to the public through no act or default of the Assignor. However, the Assignor will promptly advise the Assignee if the Assignor becomes aware of any actual or potential loss of secrecy or confidentiality of the information referred to in clause 3.1.

4. IMPROVEMENTS

4.1 **Disclosure:** The Assignor will immediately disclose to the Assignee all Improvements as and when such Improvements arise.

4.2 Ownership of Improvements:

- (a) The Assignee will exclusively own all Improvements and all intellectual property rights in those Improvements.
- (b) If any Improvements or intellectual property rights in those Improvements are not owned by the Assignee on their creation but are owned by the Assignor, the Assignor will hold those Improvements and intellectual property rights on trust for the Assignee.
- (c) The Assignor will at any time reasonably requested by the Assignee, and at the Assignee's expense, appropriately execute all documents necessary to:
 - (i) confirm the Assignee's ownership of the Improvements and all intellectual property rights in those Improvements; or
 - (ii) file, prosecute and defend any protective application for the Improvements and all intellectual property rights in those Improvements.

5. EXECUTION OF DOCUMENTS AND FURTHER ACTIONS

5.1 **Further actions:** If requested by the Assignee the Assignor will, at the Assignee's expense, execute all documents, give such assistance and do all other acts and things that the Assignee considers necessary or desirable to:

- (a) apply for and obtain, or (if the Assignee thinks fit) join with the Assignee in applying for and obtaining, protection of the Intellectual Property Rights. The Assignor acknowledges that all rights, title and interest in and to such applications and any granted protection will be owned by the Assignee;
- (b) vest in the Assignee ownership of any protection referred to in paragraph (a) above;
- (c) amend, maintain, or renew any protection referred to in paragraph (a);
- (d) enforce any of the Intellectual Property Rights (including by obtaining any remedies that may be available for infringement of the Intellectual Property Rights);
- (e) defeat any challenge to the validity of any of the Intellectual Property Rights;
- (f) defend any opposition proceedings brought by a third party in relation to the Intellectual Property Rights;
- (g) conduct opposition proceedings regarding any application for intellectual property protection by a third party, where such application may adversely affect the Assignee's ability to exploit the Intellectual Property Rights; and
- (h) otherwise implement and carry out the Assignor's obligations under this assignment.

5.2 **Recordal:** The Assignor authorises and requests the Commissioner of Patents and Trademarks of the United States, and any official of any other country or state, whose duty it is to issue patents, to issue the same to the Assignee in accordance with this assignment.

5.3 **Power of attorney:** The Assignor irrevocably appoints the Assignee as the Assignor's attorney with full power to act in the Assignor's name and on the Assignor's behalf to do any of the things required to be done by the Assignor under clause 5.1 and 5.2:

- (a) to the extent that the Assignor fails to do any of those things within a reasonable time after being called upon to do so by the Assignee; or
- (b) if the Assignee is unable, after making reasonable and proper inquiries, to locate the Assignor to ask the Assignor to do those things.

6. GENERAL

6.1 **Assignment binding and delivered:** This assignment:

- (a) is intended to be immediately and unconditionally binding upon the Assignor as soon as the Assignor executes and delivers this assignment; and
- (b) without limiting any other mode of delivery, will be delivered by the Assignor immediately upon the earlier of:
 - (i) physical delivery of an original form of this assignment executed by the Assignor; or
 - (ii) transmission (whether by facsimile or email) of a copy of this assignment executed by the Assignor,

to the Assignee or its solicitor, patent attorney, or agent.

6.2 **Waiver:** No failure or delay by any party in exercising any right, power or privilege under this assignment will operate as a waiver of such right, power or privilege, nor will any single or partial exercise preclude any other or further exercise of any right, power or privilege under this assignment.


6.3 **Counterparts:**

- (a) This assignment may be executed in any number of counterparts (including facsimile or electronically scanned copies) all of which, when taken together, will constitute one and the same instrument.
- (b) A party may enter into this assignment by executing any counterpart. The parties acknowledge that this assignment may be executed on the basis of an exchange of facsimile or electronically scanned copies and confirm that their respective execution of this assignment by such means will be a valid and sufficient execution.

6.4 **Governing law:** The formation, validity, construction and performance of this assignment will be governed by and construed in accordance with the laws of New Zealand. The parties irrevocably agree that the Courts of New Zealand will have non-exclusive jurisdiction to hear and determine all disputes under or in connection with this assignment. The parties irrevocably waive any objections to New Zealand as the forum for proceedings on the grounds of forum non-conveniens or any similar grounds.


SIGNED AS A DEED

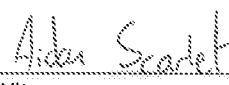
Stuart John Inglis:

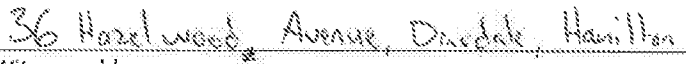

Signature

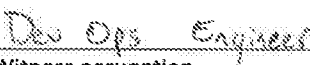

Date

in the presence of

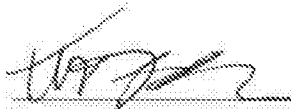

Witness signature


Witness name


Witness address

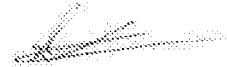

Witness occupation

Timothy Kelly Dawson:


Signature

19/08/2019
Date

in the presence of

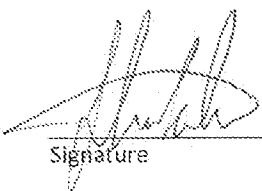

Witness signature

Aidan Scarlet
Witness name

36 Hazelwood Avenue, Dinsdale, Hamilton
Witness address

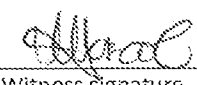
Dev Ops Engineer
Witness occupation

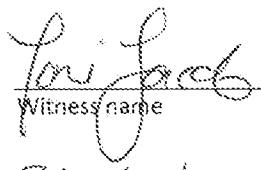
Sheridan John Lambert:


Signature

19/8/19
Date

In the presence of


Witness signature


Witness name

36 Lake St, Cambridge.
Witness address

Admin
Witness occupation

SCHEDULE

INVENTION

The invention described in the attached specification.

This is filed as one or more patent applications according to the following particulars (to be completed under clause 2.4):

Country	Number	Filing date
----------------	---------------	--------------------

PERSISTENT STORAGE DEVICE MANAGEMENT

FIELD

[0001] This relates to persistent storage device management.

BACKGROUND

5 [0002] Data may be stored on a persistent storage device, such as a hard drive, which has a number of blocks. A request to write data to the drive may specify a physical location (such as a block) of the persistent storage device at which to write the data. A request to read data from the drive may specify a physical location (such as a block) of the persistent storage device from which to read the data.

10 SUMMARY

[0003] In a first example embodiment, there is provided a method comprising: receiving a request to write data at a virtual location; writing the data to a physical location on a persistent storage device; and recording a mapping from the virtual location to the physical location; wherein the physical location corresponds to a
15 next free block in a sequence of blocks on the persistent storage device.

BRIEF DESCRIPTION

[0004] The description is framed by way of example with reference to the drawings which show certain embodiments. However, these are provided for illustration only.

20 [0005] Figure 1 shows an example approach for handling a write request.

[0006] Figure 2 shows a worked example of the approach of Figure 1.

[0007] Figure 3 shows an example approach for handling a write request where data is to be stored in stripes.

[0008] Figures 4A and 4B show a worked example of the approach of Figure 3.

25 [0009] Figure 5 shows an example approach for handling a read request.

[0010] Figure 6 shows an example approach for handling block reconstruction of a missing block.

[0011] Figure 7 shows an example approach for handling an update request.

[0012] Figure 8 shows an example approach for garbage collection.

5 [0013] Figure 9 shows a worked example of the approach of Figure 8.

[0014] Figure 10 shows a block diagram of an example computer system 900.

DETAILED DESCRIPTION

[0015] In some embodiments, there is provided a method for a driver to receive a request to write data at a virtual location, writing data at the next free physical
10 location in sequence, and recording a mapping between the virtual and physical locations. By always storing data on a next free block on the disk, this can be used to avoid a disk going through a read-modify-write cycle. This can improve throughput and write speed.

[0016] In this case, “next” may refer to the next in a sequence of blocks. The
15 sequence may be determined based on the configuration of the particular persistent storage device. For example, “next” may be based, at least in part, on the block having the lowest expected seek time from the current block.

[0017] This method may be implemented via a driver. The driver may provide a functionality to write data, read data, and update data. In this way, the driver
20 provides all the necessary functionality for a block device, such as a persistent storage device. In particular, the driver may be a userspace application which receives system calls from a kernel module and in turn writes data to the underlying persistent devices.

Writing

[0018] Figure 1 shows an example approach for the driver handling a write request to a persistent storage device, such as a hard drive. This may be implemented in a driver running on a system.

5 [0019] At step 101, a request is received to write data at a first location of the persistent storage device. The request may be received by a function call of the driver. The request may comprise the data or may comprise a pointer to where the data may be retrieved from (for example, a memory address).

[0020] The first location is a virtual location, since it may not correspond to a physical location on a persistent device. However, the first location is structured as
10 if it were a physical location, since the redirection by the driver is hidden from a user of the driver. The request may therefore comprise an identifier of the first location, such as a block index and disk identifier. A subsequent request to read the data from the virtual location should therefore return the same data
15 irrespective of the physical location of the data.

[0021] In some cases, the data may be part of a stream of data. A request to write data is therefore received (or treated as being received) for each block in the stream.

[0022] At step 102, the data is written to a second location. The second location
20 is a physical location which corresponds to a block on the persistent storage device. However, this may be different from any block indicated by the virtual location. The virtual location may not be correlated with the physical location.

[0023] In particular, the physical location may be a head of the persistent storage device, that is, the next free block of the persistent storage device or a sequence
25 of next free blocks of the persistent storage device. This can be recorded at the persistent storage device by a head counter which records the index of the next free block of the persistent storage device. A free block is a block in which no data is stored, or if data is stored, there is no need for that data to be retrievable. In this

way, the block is appended to the sequence of blocks which have already been written.

[0024] By writing to a next free block (which, by definition, is free), there is no need to determine whether the data at the indicated block has changed. This
5 avoids the need to initially read from the block, and therefore can result in improved performance.

[0025] Moreover, because the physical location need not be correlated with the virtual location, there is no need for the requestor (such as a program) to determine where the next free block is to achieve this. The function of a location
10 as a label to retrieve data is decoupled from the function of a location as a mechanism for managing a persistent storage device. These functions are performed by the virtual location and the physical location respectively.

[0026] At step 103, a checksum for the data at the block indicated by the physical location is calculated. The checksum may be a cryptographic hash. This can be used
15 to verify that data subsequently retrieved from the block is the same as the data stored at the block. In some cases, the checksum may comprise parity information which can correct some errors in the block.

[0027] At step 104, metadata for the persistent storage device is updated.

[0028] The metadata comprises a mapping from the virtual location to the
20 physical location. This links the virtual location (which the requestor or another process can use to indicate the intended data to be retrieved) to the physical location (where the actual data is stored). The metadata further comprises the checksum for each block.

[0029] The metadata may be stored on the persistent storage device. For
25 example, this may be a predetermined part of the persistent storage device. Additionally or alternatively, the metadata is stored on a separate persistent storage device which may be optimized for use as a cache.

[0030] In some cases, updating metadata for the block comprises determining if metadata for the block indicated by the first location already exists. If it exists, the original metadata may be updated, and replacement metadata inserted (or the new metadata overwritten on the original metadata). This can prevent the metadata having multiple entries for the same data. The physical location corresponding to the previous location may be marked as dirty, indicating that it is free to be written over.

[0031] At step 105, the head counter is updated to reflect the next free block.

[0032] This may comprise incrementing the head counter to the next block in the sequence of blocks of the persistent storage device (or wrapping around to the first block, once the final block has been reached). However, in some cases, certain blocks may be skipped (for example, if they correspond to faulty portions of the persistent storage device or if the persistent storage device is missing).

[0033] This may involve calculating a tail. The tail may correspond to the earliest block in the sequence of blocks before the head where there has been data written. The earliest block may not have the lowest index. If the incremented head counter is equal to the tail, this indicates that the nominal next free block indicated by the head counter is actually not free: it has data written to it. In this case, a notification, warning, or error may be raised and/or a garbage collection process may be performed. An example garbage collection process is shown with reference to Figure 8. Until the garbage collection process is completed, the driver may refuse to handle any more requests to write data.

[0034] Alternatively, the tail may be the location at which data could be written sequentially. In some cases, this may correspond to a garbage collected region.

[0035] In some cases, the tail may be recorded as a tail counter which is updated based on a garbage collection procedure.

[0036] This approach provides a high-performance approach to writing data to persistent storage device, since the read-modify-write cycle of some conventional

write approaches may be avoided. In addition, because data is written sequentially, this typically results in a faster write speed.

[0037] A further benefit is that the lifespan of persistent storage devices may be prolonged using the current approach. In conventional approaches, when a single
 5 block fails or reaches its maximum number of writes, the whole device is often replaced. For example, in the case of a hard disk, a single hot-spot (corresponding to one physical address) can cause a persistent storage device to reach the end of its life even if the bulk of the persistent storage device has never been written to. However, in the current approach, the head moves across the whole persistent
 10 storage device. Because data is always written to the head, this will naturally lead to equal levels of writing across the persistent storage device. That is, each block will generally be written to once before any block is written to a second time, irrespective of the virtual address of the data being written. This approach therefore may avoid premature wearing of persistent storage devices.

15 **Example of Writing**

[0038] Figure 2 shows a worked example of how the approach shown in Figure 1 can be used to write a stream of data to a persistent storage device. In Figure 2, the driver writes data to a single persistent storage device (such as a hard drive).

[0039] At 200, the system is shown in its initial state. The head counter indicates
 20 location 0 as the next free block to be written to. The mapping is empty.

[0040] At 201, the driver receives a request to write data **a** to virtual location 6. Because the head counter indicates physical location 0 as the next free block, data **a** is written at physical location 0, along with a checksum **h(a)**. The mapping is updated to show a mapping from virtual location 6 to physical location 0. The head
 25 counter is incremented to physical location 1. The tail is computed to be at physical location 0, which is different from the head counter.

[0041] At 202, the driver receives a request to write data **b** to virtual location 3. Because the head counter indicates physical location 1 as the next free block, data

b is written at physical location 1, along with a checksum **h(b)**. The mapping is updated to show a mapping from virtual location 3 to physical location 1. The head counter is incremented to physical location 2. The tail is computed to be at physical location 0, which is different from the head counter.

5 [0042] At 203, the driver receives a request to write data **c** to virtual location 2. Because the head counter indicates physical location 2 as the next free block, data **c** is written at physical location 2, along with a checksum **h(c)**. The mapping is updated to show a mapping from virtual location 2 to physical location 2. The head counter is incremented to physical location 3. The tail is computed to be at physical
10 location 0, which is different from the head counter.

[0043] At 204, the driver receives a request to write data **b'** to virtual location 3. Because the head counter indicates physical location 3 as the next free block, data **d** is written at physical location 3, along with a checksum **h(b')**. The mapping is updated to show a mapping from virtual location 3 to physical location 3. This
15 replaces the previous mapping from virtual location 3 to physical location 1. Physical location 1 is therefore marked as dirty and can be freed in a garbage collection process. However, the data at physical location 1 remains at physical location 1 until this garbage collection process occurs.

[0044] The head counter is incremented to physical location 4. The tail is
20 computed to be at physical location 0, which is different from the head counter.

[0045] This example demonstrates the sequential nature of writing under the current approach. While the driver receives requests to write data at varied virtual locations, the driver actually writes in sequence. This avoids the need for the driver to read from the persistent storage device during write operations, and therefore
25 may provide a higher performance for writing random data.

Writing to Arrays

[0046] A similar approach may be used to write data as a stripe across an array of persistent storage devices. In this way, a subset of the devices can be used for data and the remaining devices can be used for parity.

5 [0047] During configuration, the driver may be configured to have a particular resilience. That is, among an array of n persistent storage devices, data is stored in stripes across the data. Each stripe comprises n blocks, one of which is on each persistent storage device. Within each stripe, k of the blocks include data and m of the block are parity, such that $n = k + m$. Consequently, the data can be recovered
10 if any k of the blocks within a stripe is available (whether the block are data or parity). This may be described as $k + m$ parity. One common arrangement is $8 + 2$ parity, where within a stripe there are 8 blocks of data and 2 blocks of parity, however any value for k and m may be used.

[0048] One approach for generating the parity data in such a scheme is to use
15 erasure encoding. Erasure encoding is described in more detail below.

[0049] The blocks of each stripe are stored in the same indexed physical location on all of the persistent storage devices. For example, stripe 0 consists of the blocks at location 0 on every device. The driver can therefore maintain a single head counter that is common to all of the devices.

20 [0050] While a single head counter is described for simplicity, in practice each persistent storage device may maintain a separate head counter. This can allow for garbage collection to occur on one device while a write is occurring on another device.

[0051] It can be useful for subsequent stripes to use different drives for parity.
25 One way to administer this is to nominate a starting drive for each stripe: stripe 0 starts at drive 0, stripe 1 starts at drive 1, and so on. This will naturally cause the parity blocks (which are written after the data blocks) to be on different drives.

[0052] Figure 3 shows an example approach for handling a write request where data is to be stored in such stripes. This may be implemented in a driver running on a system.

[0053] At step 301, a request to write k blocks of data is received. The request
5 corresponds to a request received at step 101.

[0054] At step 302, after k data blocks have been received, the driver calculates parity data based on written data. This results in m blocks of parity data.

[0055] At step 303, each of the k blocks of data and m blocks of parity data is written to corresponding persistent storage devices at the physical location
10 indicated by the head counter. This results in the stripe being completed. In addition, $k+m$ checksums are written to the metadata, each corresponding to a block.

[0056] At step 304, the head counter is updated to reflect the physical location of the next free block at each device, in a similar manner to step 105. Once a block
15 has been written to each persistent storage device at the physical location indicated by the head counter, the head counter may be incremented to the next free location in the sequence.

[0057] This approach provides a high-performance approach to writing data to an array of persistent storage devices, since the read–modify–write cycle of some
20 conventional write approaches may be avoided.

[0058] This also avoids the penalty that comes with higher parity in conventional approaches. Where data can be modified, the parity must be updated after every modification. Thus, if there are two parity blocks in a stripe, modifying a piece of data requires a read–modify–write cycle for the data itself, plus two further read–
25 modify–write cycles for each parity block. This amplification places a performance limit on the level of parity that is practical within convention systems. In practice, parity is limited to 2 in conventional systems because of this.

[0059] In contrast, the approach shown in Figure 3 does not involve any read–modify–write cycles. Increasing parity therefore increases the overhead at a much lower rate than conventional approaches. Accordingly, parity may be almost unlimited, and in practice 2, 4, or 12 parity may be used.

- 5 [0060] Similarly, the approach shown in Figure 3 is not limited to any particular size of array, and there may be hundreds of persistent storage devices in the array. In practice, an array having around 10 persistent storage devices may be used.

[0061] Moreover, because the number of operations at each device is relatively low, even slower persistent storage devices may provide high levels of
10 performance.

[0062] In some cases, one device in the array may fail or otherwise be unavailable. When this occurs, a block which may have intended to be written to that device may simply be written to the next device in the array. In this manner, a failed device does not substantially impede writing.

- 15 [0063] In such cases, the system may still be said to use $k + m$ parity. However, for each stripe, there may end up being less than k data blocks (due to the failed device) written in a stripe to maintain m blocks of parity.

Example of Writing to Arrays

[0064] Figures 4A and 4B shows an example process by which the approach
20 shown in Figure 3 can be used to write a stream of data to an array of persistent storage device. In Figures 4A and 4B, the driver writes data to an array of five persistent devices. The driver is configured to use 3+2 parity (that is, 2 blocks of parity for every 3 blocks of data).

[0065] At 400, the system is shown in its initial state. The head counter indicates
25 location 0 as the next free block to be written to. The mapping is empty. Each of drives 0, 1, 2, 3, and 4 are empty.

[0066] At 401, the driver receives a request to write data **a** to virtual location 6. Because the head counter indicates physical location 0 as the next free block, data **a** is written at physical location 0 on drive 0, along with a checksum **h(a)**. The mapping is updated to show a mapping from virtual location 6 to physical location 0 on drive 0. Because the stripe has not yet been completed, the head counter is not incremented.

[0067] At 402, the driver receives a request to write data **b** to virtual location 3. Because the head counter indicates physical location 0 as the next free block, data **b** is written at physical location 0 on drive 1, along with a checksum **h(b)**. The mapping is updated to show a mapping from virtual location 3 to physical location 0 on drive 1. Because the stripe has not yet been completed, the head counter is not incremented.

[0068] At 403, the driver receives a request to write data **c** to virtual location 2. Because the head counter indicates physical location 0 as the next free block, data **c** is written at physical location 0 on drive 2, along with a checksum **h(c)**. The mapping is updated to show a mapping from virtual location 2 to physical location 0 on drive 2. Because the stripe has not yet been completed, the head counter is not incremented.

[0069] Now that 3 data blocks have been written, the stripe is completed by adding parity blocks to the 2 remaining devices. The parity data is computed based on the data that has been written to the devices. This may be read from the devices, or it may be stored in memory. This results in parity data **d** and parity data **e**.

[0070] At 404, because the head counter indicates physical location 0 as the next free block, parity data **d** is written at physical location 0 on drive 3, along with a checksum **h(d)**, and parity data **e** is written at physical location 0 on drive 4, along with a checksum **h(e)**. In addition, now that the stripe has been completed, the head counter is incremented to the next free location in the sequence, and is now 1.

[0071] In this way, a stripe may be written without the need for the driver to read from the persistent storage device during write operations, and therefore may provide a higher performance for writing random data with any level of parity.

Reading

5 [0072] Figure 5 shows an example approach for handling a read request. This may be implemented in the same manner as the approach of Figure 1.

[0073] At step 501, a request is received to read data from a virtual location. The request may be received in the same manner as the request at step 101. For example, this may be received by a function call of the driver. The request may
10 comprise a pointer to where the data should be stored once retrieved (for example, a memory address).

[0074] At step 502, a physical location for the data corresponding to the first location is obtained from the metadata. This may be obtained by finding the virtual location in the mapping of the metadata and finding the corresponding physical
15 location which was previously stored at step 104.

[0075] If the metadata does not contain references to the virtual location, the read request may fail, and an error may be raised.

[0076] At step 503, the data is retrieved by reading from the physical location at the appropriate permanent storage device. The data may then be stored in a
20 memory location indicated in the request or another location.

[0077] At step 504, the retrieved data is validated to ensure that the retrieved data is the data which was originally stored. This may involve calculating a checksum for the retrieved data and comparing it to the checksum stored with the data. If the checksums match, it is highly likely the data was successfully retrieved.

25 [0078] If the checksums do not match, but the data can be recovered, this may be performed. This may make use of parity information stored in the same stripe

as the data. If the errors cannot be corrected, a notification, warning, or error may be raised.

[0079] The retrieved data can then be returned to the process which issued the read request. This may comprise providing a success value (such as 0) as a return value to a function call.

[0080] In some cases, step 504 may be omitted. For example, on reads where it is desirable to maximize throughput, step 504 may be selectively skipped.

[0081] From the point of view of the other requestor, the data was stored at, and retrieved from, the virtual location. The management and configuration of the persistent storage device is therefore obscured from the process which attempts to read the data.

Block Reconstruction

[0082] Occasionally, it may not be possible to retrieve data at a particular physical location, or that physical location may be overly slow. This may occur if step 503 is unsuccessful due to the hardware failure of a block or device, or step 504 is unsuccessful due to retrieved, but unverified, data. When this occurs, a block reconstruction process may be undertaken.

[0083] Figure 6 shows an example approach for handling block reconstruction of a missing block. In this example, the data is stored at an array of persistent storage devices which are configured to use $k + m$ parity.

[0084] At step 601, the driver reads data from at least k blocks in the same stripe as the missing block. Since each block in the stripe is from a different persistent storage device, it is highly unlikely in practice that multiple blocks within a stripe will become unavailable at the same time.

[0085] At step 602, the driver uses the available data to reconstruct the data from the missing block. This is possible since, for $k + m$ parity, any k blocks of a stripe can be used to reconstruct all the blocks within that stripe. If the blocks in a stripe

have been encoded using erasure encoding to produce the parity, any missing blocks can be derived by erasure decoding, as described below.

[0086] At step 603, the reconstructed data is written at the head (that is, the next free block, as indicated by the head counter). This may be implemented in the same manner at steps 102 and 103.

[0087] At step 604, the metadata is updated such that the virtual location corresponding to the missing block is now mapped to the block at which the reconstructed data was written. The missing block can then be marked as dirty (if there is no hardware issue) or be recorded as unusable (if there is a hardware issue).

[0088] In this manner, the data can be recovered immediately after a single failed read. The block reconstruction process occurs quickly since only the failed block needs to be rewritten. This allows minimal speed loss: while there may be a slight delay in reconstructing the data after the first read, any subsequent reads will revert to the original speed even if the hardware issue has not been resolved.

[0089] When block reconstruction is required due to a failed persistent storage device, all blocks may be reconstructed and rewritten to non-failed devices. Once this process occurs, the speed of the array recovers from a degraded condition without the need to re-write any data of the other non-failed persistent storage devices.

Duplication

[0090] In some cases, it may be desirable to have a single virtual address corresponding to multiple physical addresses. The same data is duplicated at each physical address.

[0091] In this manner, a read request may read from any of the physical addresses corresponding to the virtual address. If the block at one of the physical addresses fails, the read request may resolve from any of the other physical addresses. The failed block may be reconstructed using a block reconstruction

technique noted above. This can avoid the delay of reconstruction while still providing that reconstruction can occur.

[0092] Alternatively, in some cases duplicates may be stored at once of more immediately following physical addresses. If a first physical address is unavailable, the data can be retrieved from the next physical address where the data is duplicated.

[0093] Thus, from the point of view of a user, the array may appear to never be slow, even if there is a failure of a persistent storage device within the array.

Update

10 [0094] Figure 7 shows an example approach for handling an update request. This may be implemented in the same manner as the approach of Figure 1.

[0095] At step 701, a request is received to update the data from a virtual location. The request may be received in the same manner as the request at step 101. For example, this may be received by a by a function call of the driver.

15 [0096] At step 702, references to the virtual location are removed from the metadata and the corresponding physical location is marked as dirty. This indicates that the physical location is free for future use.

[0097] Once the metadata is updated, there is no need for further persistent storage device operations. If the virtual location is not in the metadata, then the corresponding data at the physical location could not be retrieved. The physical location will eventually be reused after a garbage collection process has occurred.

Dirty Blocks

[0098] Updating data has been described in terms of selected blocks being marked as dirty. This may be implemented by each stripe maintaining a counter of the number of unused blocks in the stripe. When a block is marked dirty or is inaccessible (for example, due to a corresponding device failing), the number of gaps for the stripe may be incremented. When each data block is rewritten to the

head (for example, during garbage collection), the number of gaps for the stripe may also be incremented. Once the counter reaches a predetermined limit (which may be the number of non-parity blocks stored in a stripe), the stripe may

Garbage Collection

5 [0099] When data indicated by a virtual location is updated, the corresponding physical location may be marked dirty indicating that the physical location can be reused. Thus, although data may initially be written sequentially on a persistent storage device, over time gaps in the sequence of data may appear.

[0100] One way to mitigate this is to periodically perform garbage collection. This
10 aims to move data from the tail (that is, the earliest non-dirty block in the sequence of written blocks) to the head, until there are no gaps in the sequence.

[0101] Garbage collection may be performed in response to one or more garbage collection conditions being met.

[0102] A first garbage collection condition may be that a number of gaps on the
15 persistent storage device exceeds a predetermined threshold. For example, where the number of gaps exceeds 1% of the capacity of the persistent storage device, the garbage collection condition may be met.

[0103] A second garbage collection condition may be that an activity level of persistent storage device falls below a threshold. For example, if the persistent
20 storage device has not been used for several minutes, a garbage collection process may commence.

[0104] A third garbage collection condition may be that a free capacity of the persistent storage device falls below a threshold. For example, when the head counter gets near to the tail, it may be useful to try to recover some usable space.

25 [0105] A fourth garbage collection condition may be that a certain time period has passed. That is, it may be useful for a garbage collection process to commence periodically.

[0106] Additionally or alternatively, garbage collection may occur any time the driver is not processing a request and there exists one or more dirty blocks.

[0107] Figure 8 shows an example garbage collection process.

[0108] At step 801, the tail is calculated. The tail is the non-dirty block (that is, the block which is storing data) that is earliest in the sequence of written blocks. In other words, this is the oldest written non-dirty block.

[0109] At step 802, the data stored at the tail is written to the block indicated by the head counter.

[0110] At step 803, the metadata for the persistent storage device is updated. A new entry mapping the virtual address for the newly written block to the head counter replaces the previous entry for that block.

[0111] At step 804, the head counter is updated to reflect the next free block in the same manner as at step 105.

[0112] This process is repeated until one of the garbage collection conditions is no longer met or until there are no dirty blocks left. By iteratively moving blocks from the tail to the head, eventually there will be no remaining gaps. This avoids any

[0113] Moreover, this process can be stopped at any time. Once step 804 terminates, the driver is ready to process further requests (including writing new blocks). In this way, garbage collection can be started and stopped without causing significant downtime.

[0114] In this manner, overall storage capacity can be recovered when data is updated without limiting the performance of other requests.

Example of Garbage Collection

[0115] Figure 9 shows a worked example approach for performing garbage collection. This may be performed by the same driver as that which performs the method of Figure 1, for example.

- 5 [0116] In Figure 9, the driver is configured to use 2+1 parity (that is, 1 block of parity for each 2 blocks of data). In practice, the same approach may be applied to any level of parity. The head counter and mapping are omitted from Figure 9.

[0117] At 900, the system is shown with data that has previously been written. Drive 0 has blocks marked dirty at physical locations 0 and 2. Drive 1 has a block marked dirty at physical location 1. Drive 2 has a block marked dirty at physical location 1. Blocks at locations 6 and 7 in all drives are empty. All other blocks in all drives have non-dirty data. The head counter is at physical location 6, since that represents the next free block. The tail is at 0, since the block at location 0 on drive 1 is the oldest non-parity non-dirty block.

- 15 [0118] At 901, the data in all non-dirty, non-parity blocks at the tail (location 0) has been moved to the head. The data at location 0 on drive 0 (**a**) is not moved, as it has been marked dirty. The data at location 0 on drive 2 (**ab**) is not moved, as it is parity data. Only the data at location 0 on drive 1 (**b**) is moved to the head. This has been moved to location 6 on drive 0. Location 6 is derived from the head counter (which indicates location 6 as the next free block). Drive 0 is derived by cycling through the starting drive of each stripe: since the stripe at location 5 started with drive 2, the stripe at location 6 starts with drive 0.

- [0119] After the move, the mapping is updated so that the virtual location which previously mapped to location 0 on drive 1 now maps to location 6 on drive 0. Location 0 on drive 1 is marked as dirty, since the data is no longer needed, and the block can be reused. The head counter remains at 6. The tail moves to 2, since the next oldest non-dirty, non-parity data is at location 2 on drive 2. All of location 1 is dirty or parity on all drives (or in other words, the number of gaps in stripe 1 is 2).

[0120] At 902, the data in all non-dirty, non-parity blocks at the tail (location 2) has been moved to the head. The data at location 1 on drive 1 (**c**), location 1 on drive 2 (**d**), and location 2 on drive 0 (**f**) is not moved, as it has been marked dirty. The data at location 1 on drive 0 (**cd**) and location 2 on drive 1 (**ef**) is not moved,
 5 as it is parity data. Only the data at location 2 on drive 2 (**e**) is moved to the head. This has been moved to location 6 on drive 1. Location 6 is derived from the head counter (which indicates location 6 as the next free block). Drive 1 is derived from the previous written block in the stripe: since drive 0 was previously written to, drive 1 is the next drive in the sequence.

10 [0121] After the move, the mapping is updated so that the virtual location which previously mapped to location 2 on drive 2 now maps to location 6 on drive 1. Location 2 on drive 2 is marked as dirty, since the data is no longer needed, and the block can be reused.

[0122] Since the data in stripe 6 is now completed, parity data (**be**) is calculated
 15 and written to the remaining block in stripe 6 (that is, at location 6 on drive 1). The head counter moves to 7, since the next free block is in location 7. The tail moves to 3, since the next oldest non-dirty, non-parity data is at location 3 on drives 0 and 1.

[0123] Following this garbage collection process, all of locations 0, 1, and 2 on all
 20 drives is now free for use. Once stripe 7 is completed, the head counter will loop to stripe 0. Garbage collection therefore allows a persistent storage device or array of persistent storage devices to be more fully utilized.

Erasure Coding

[0124] As noted above, erasure coding is an approach in which a piece of data is
 25 encoded to form n data fragments. The original data is recoverable from any k data fragments ($k \leq n$). Each data fragment is a w -bit word, where w can be any positive integer. For example, 1-bit and 8-bit are common. Generally, the higher the values for n and w , the more computationally intensive the erasure encoding and decoding process is.

[0125] First, the original data is segmented into $w \times k$ -bit data segments.

[0126] Second, k data fragments $d_1 \dots d_k$ are formed by separating each data segment into k w -bit sized data words. These are data words.

[0127] Third, $m = n - k$ data fragments $c_1 \dots c_m$ are formed based on linear combinations of the data words with coefficients, such that:

$$c_i = a_{i,0}d_0 + \dots + a_{i,k}d_k$$

where the arithmetic is defined over the Galois field $GF(2^w)$. These are coding words.

[0128] The data fragments can be stored in a stripe over an array of persistent storage devices, such that the data fragments are distributed over the persistent storage devices. In some cases, different persistent storage devices can hold different numbers of data fragments. Optionally, a checksum based on the data segment is also stored. Once all data segments are stored, a further checksum based on the original data can also be stored.

[0129] There are multiple ways to select the coefficients. One approach is to use a Cauchy Reed–Solomon code. This occurs by selecting n numbers in $GF(2^w)$ and partitioning them into two distinct sets $X = \{x_1, \dots, x_m\}$ and $Y = \{y_1, \dots, y_k\}$, with $X \cap Y = \emptyset$ such that:

$$a_{i,j} = \frac{1}{x_i + y_j}$$

where the arithmetic is defined over the Galois field $GF(2^w)$.

[0130] Once the data is stored, any k of the data fragments can be used to retrieve the data. For example, if the data words $d_1 \dots d_k$ were obtained, these can be combined using the inverse of how they were segmented. If a coding word c_i is used, a system of equations based on the equation noted above can be solved to obtain the missing data fragment.

[0131] In some situations, erasure encoding can be performed efficiently using one or more GPUs while the request itself is handled by one or more CPUs. This can result in improved overall performance.

[0132] Erasure encoding provides an approach which allows for high resilience and flexibility. The data fragments can be stored across multiple independent systems. If m or fewer fragments are lost, the corresponding data segment can be recovered, and thus the original data can be calculated. In addition, any k data fragments can be used to recover the original data, which allows any of the storing systems to respond.

10 System

[0133] Figure 10 shows a block diagram of an example computer system 1000.

[0134] The system 1000 comprises a CPU (central processing unit) 1010, a main memory 1020, one or more peripherals 1030, and a GPU (graphics processing unit) 1050. The CPU 1010, main memory 1020, peripherals 1030, and GPU 1050 are connected by a bus 1040. In the case of a PCIe (PCI Express) topology, the bus 1040 includes a root complex, and each of the peripherals 1030 and the GPU 1050 are PCIe endpoints. The GPU 1050 can communicate with the peripherals 1030 directly via the bus 1040. In some cases, more than one CPU 1010 and/or more than one GPU 1050 is provided.

[0135] The peripherals 1030 may include persistent storage devices (such as hard disk drives, solid state drives, or the like), storage controllers (such as RAID controllers), network controllers (such as network interface cards), switches (such as PCIe switches configured to connect further peripherals), or any of a variety of devices.

[0136] The peripherals 1030 and the GPU 1050 may have access to the main memory 1020 via the bus 1040. Through DMA (direct memory access), this can allow the peripheral 1030 or the GPU 1050 to read data to and write data from the

main memory. This may involve copying the data from main memory 1020 to a memory local to the peripheral 1030 or the GPU 1050.

[0137] In some cases, the memory space for the CPU 1010 and the GPU 1050 can be shared via unified virtual memory. This allows for a pointer to a single memory address to be passed between the CPU and the GPU without the need to copy data between a main memory 1020 and a local memory.

[0138] Computer system 500 may be used to implement the methods noted above. For example, CPU 1010 or GPU 1050 may run a program which operates as a driver to one or more of the storage device peripherals 1030.

[0139] The storage device peripherals 1030 may be of different types. A relatively fast storage device peripheral 1030 may be used to store metadata (such as a mapping between virtual and physical locations) and an array of relatively slow storage device peripherals 1030 may be used to store the underlying data. In some cases, some of the storage device peripherals 1030 used to store the underlying data can be of mixed types, including one or more hard disk drives, solid state drives, RAM devices, or non-volatile RAM devices.

Interpretation

[0140] A number of methods have been described above. It will be appreciated that any of these methods may be embodied by a series of instructions, which may form a computer program. These instructions, or this computer program, may be stored on a computer readable medium, which may be non-transitory. When executed, these instructions or this program may cause a processor, such as a CPU or GPU, to perform the described methods.

[0141] Where an approach has been described as being implemented by a processor, this may comprise a plurality of processors. That is, at least in the case of processors, the singular should be interpreted as including the plural. Where methods comprise multiple steps, different steps or different parts of a step may be performed by different processors. In some cases, one step may be performed

by a first type of processor (for example, processing of large sets of data may be performed by a GPU) and another step may be performed by a second type of processor (such as metadata lookup being performed by a CPU).

[0142] The order of steps within methods may be altered, such that steps are performed out of order or in parallel, except where one step is dependent on another having been performed, or the context otherwise requires.

[0143] The term “comprises” and other grammatical forms is intended to have an inclusive meaning unless otherwise noted. That is, they should be taken to mean an inclusion of the listed components, and possibly of other non-specified components or elements.

[0144] While the present invention has been explained by the description of certain embodiments and with reference to the drawings, the invention is not intended to be restricted to such details. Additional advantages and modifications will readily appear to those skilled in the art. Therefore, the invention in its broader aspects is not limited to the specific details, representative apparatuses and methods, and illustrative examples shown and described. Accordingly, departures may be made from such details without departure from the spirit or scope of the general inventive concept.

CLAIMS

1. A method comprising:
 - receiving a request to write data at a virtual location;
 - writing the data to a physical location on a persistent storage device; and
 - 5 recording a mapping from the virtual location to the physical location;wherein the physical location corresponds to a next free block in a sequence of blocks on the persistent storage device.
2. The method of claim 1, wherein the virtual location and the physical location are not correlated.
- 10 3. The method of claim 1 or 2, wherein the physical location is determined by a head counter.
4. The method of claim 3, further comprising:
 - after writing the data to the physical location, updating the head counter.
5. The method of any of claims 1 to 4, further comprising:
 - 15 receiving one or more requests to write further data to corresponding virtual locations;
 - wherein the data and the further data form a data stream.
6. The method of any of claims 1 to 5, wherein writing the data to a second location occurs without performing a read from the persistent storage device.
- 20 7. The method of any of claims 1 to 6, further comprising:
 - calculating a checksum for the data; and

recording the checksum in metadata associated with the persistent storage device.

8. The method of any of claims 1 to 7, further comprising:
determining that a garbage collection condition is met.
- 5 9. The method of claim 8, wherein the garbage collection condition comprises one or more of:

a number of gaps on the persistent storage device exceeds a predetermined threshold;

an activity level of the persistent storage device falls below a threshold;
10 a free capacity of the persistent storage device falls below a threshold; or

a time period has passed.
10. The method of claim 8 or 9, further comprising, in response to determining that a garbage collection condition is met:

determining a block at the tail of the sequence of blocks;
15 writing the data at the block to the head of the sequence of blocks; and

updating the mapping based on the writing.
11. The method of any of claims 1 to 10, wherein writing the data to a physical location on a persistent storage device comprises:

writing the data to a physical location on a persistent storage device, the
20 physical location corresponding to a block within a stripe, the stripe comprising corresponding a plurality of blocks, each block being a physical location on each of a plurality of persistent storage device.

12. The method of claim 11, further comprising:
- determining that a predetermined number of blocks within the stripe have been written;
- calculating parity corresponding to the data written to the predetermined number of blocks within the stripe; and
- 5 writing the parity data in one or more blocks within the stripe.
13. The method of claim 12, further comprising:
- retrieving a checksum for the retrieved data;
- calculating a checksum for the retrieved data; and
- 10 comparing the calculated checksum with the retrieved checksum.
14. The method of any of claims 11 to 13, further comprising:
- determining that a block of a stripe is missing;
- reconstructing the data at the missing block from one or more other blocks in the stripe;
- 15 writing the data to the head of the sequence of blocks; and
- updating the mapping based on the writing.
15. The method of any of claims 11 to 14, further comprising:
- determining that a persistent storage device among the plurality of persistent storage devices is unavailable; and
- 20 wherein writing the data to a physical location on a persistent storage device comprises:
- writing the data to a physical location on a different one of the persistent storage devices.

16. The method of any of claims 1 to 14, wherein recording a mapping from the virtual location to the physical location comprises:
- identifying a mapping from the virtual location to a previous physical location; and
 - 5 updating the mapping to record a mapping from the virtual location to the physical location.
17. The method of claim 16, further comprising:
- recording the block corresponding to the previous physical location as dirty.
18. The method of any of claims 1 to 17, further comprising:
- 10 writing a duplicate of the data to a subsequent physical location.
19. A system comprising:
- one or more processors; and
 - a memory;
- 15 wherein the memory comprises instructions which, when executed by the one or more processors, configure the one or more processors to perform the method of any of claims 1 to 18.
20. One or more non-transitory computer readable media comprising instructions which, when executed by one or more processors, cause the one or more processors to perform the method of any of claims 1 to 18.
- 20 21. A computer program which, when executed by one or more processors, cause the one or more processors to perform the method of any of claims 1 to 18.

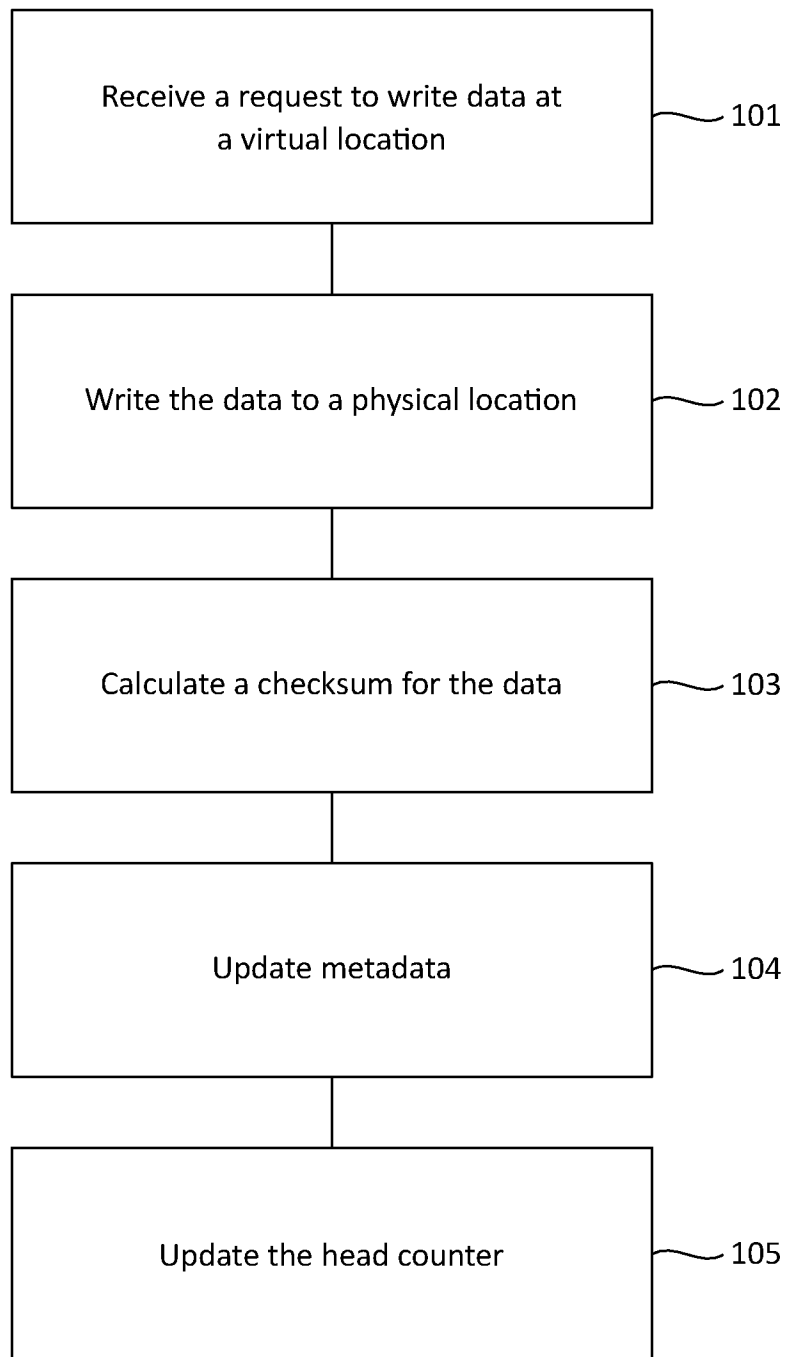


FIG. 1

	0	1	2	3	4	5	6	7	HEAD	MAP
200									0	
201	a h(a)								1	6→0
202	a h(a)	b h(b)							2	6→0 3→1
203	a h(a)	b h(b)	c h(c)						3	6→0 3→1 2→2
204	a h(a)	b h(b)	c h(c)	b' h(b')					4	6→0 3→3 2→2

FIG. 2

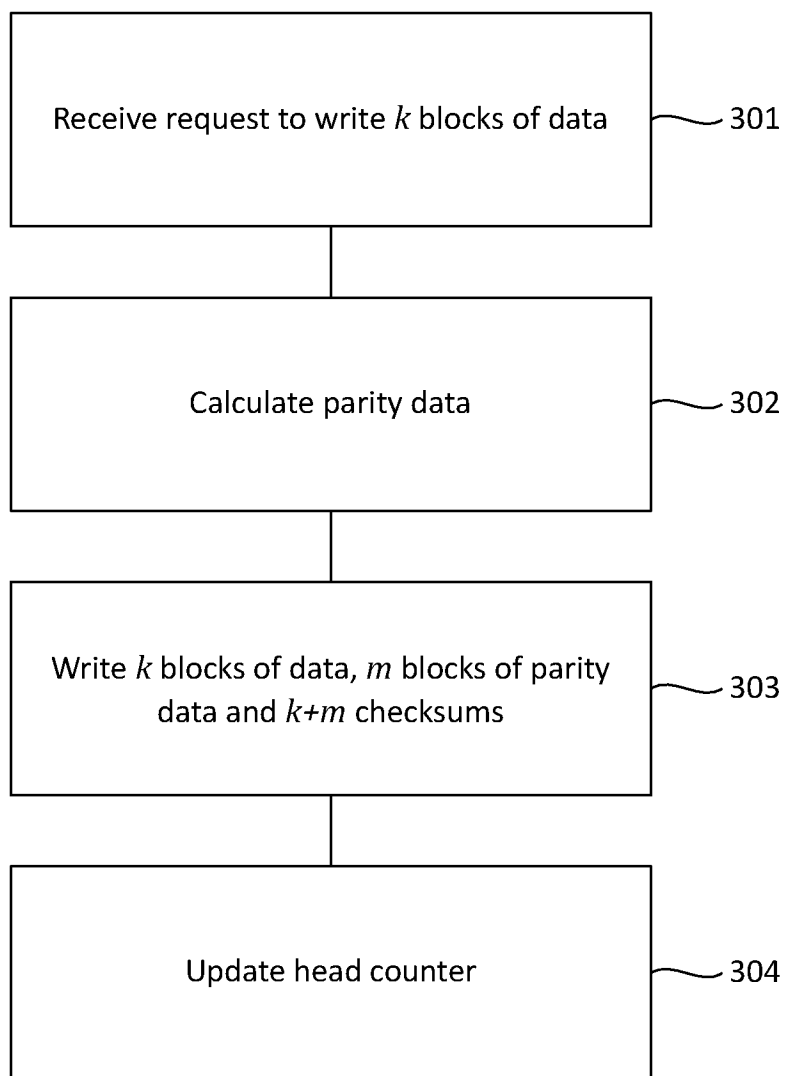


FIG. 3

		0	1	2	3	4	5	6	7	HEAD	MAP
400	Drive 0									0	
	Drive 1										
	Drive 2										
	Drive 3										
	Drive 4										

		0	1	2	3	4	5	6	7	HEAD	MAP
401	Drive 0	a h(a)								0	6→0:0
	Drive 1										
	Drive 2										
	Drive 3										
	Drive 4										

		0	1	2	3	4	5	6	7	HEAD	MAP
402	Drive 0	a h(a)								0	6→0:0 6→1:0
	Drive 1	b h(b)									
	Drive 2										
	Drive 3										
	Drive 4										

FIG. 4A

		0	1	2	3	4	5	6	7	HEAD	MAP
403	Drive 0	a h(a)								0	6→0:0 3→1:0 2→2:0
	Drive 1	b h(b)									
	Drive 2	c h(c)									
	Drive 3										
	Drive 4										
404	Drive 0	a h(a)								1	6→0:0 3→1:0 2→2:0
	Drive 1	b h(b)									
	Drive 2	c h(c)									
	Drive 3	d h(d)									
	Drive 4	e h(e)									

FIG. 4B

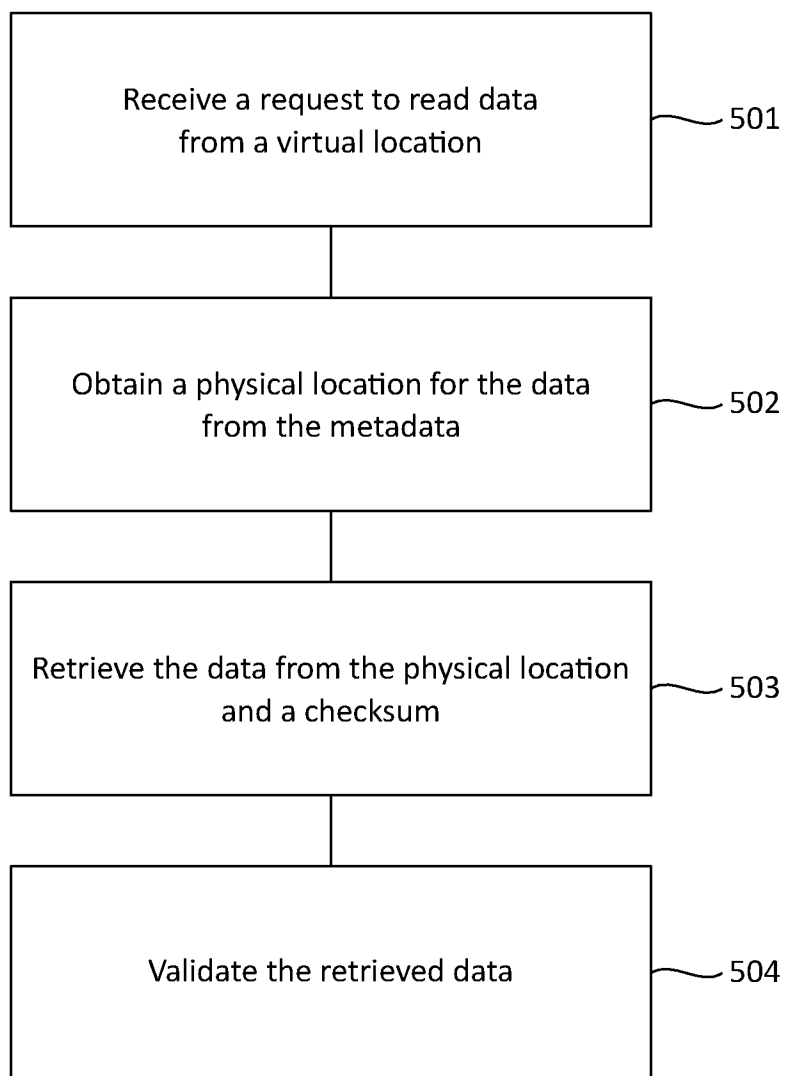


FIG. 5

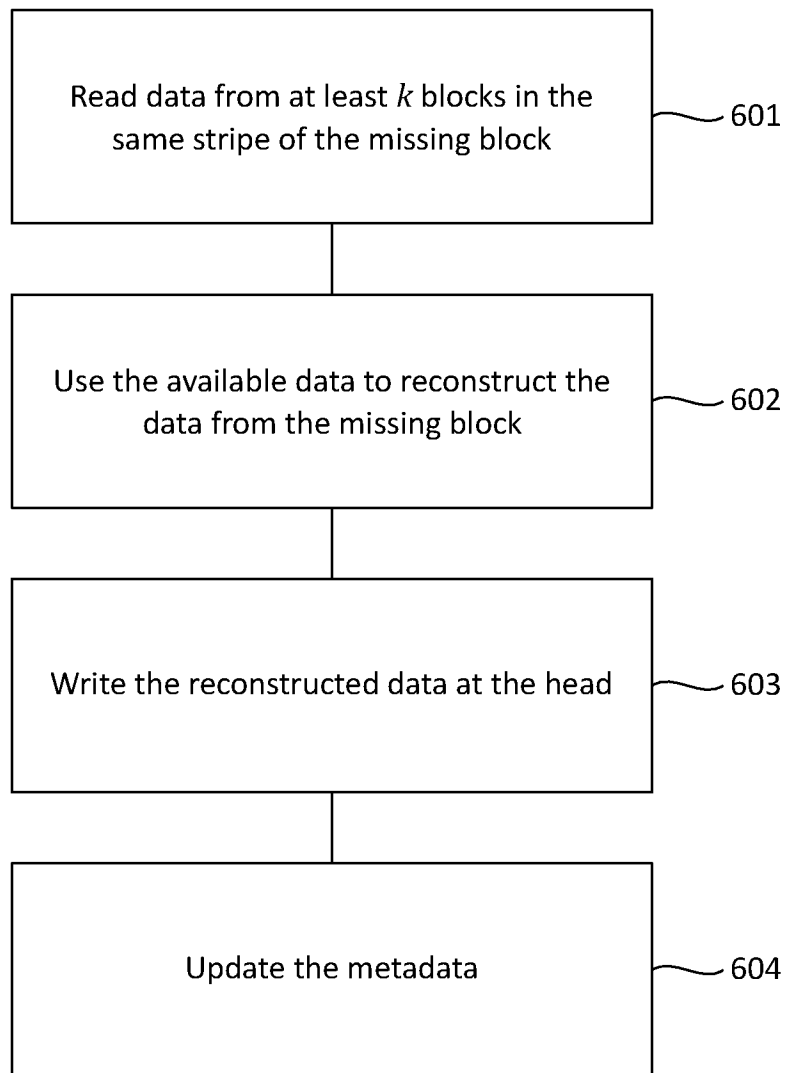
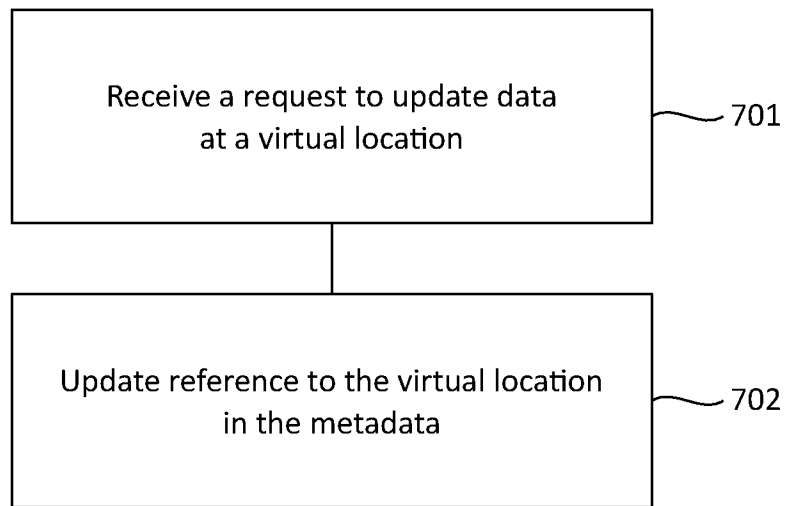
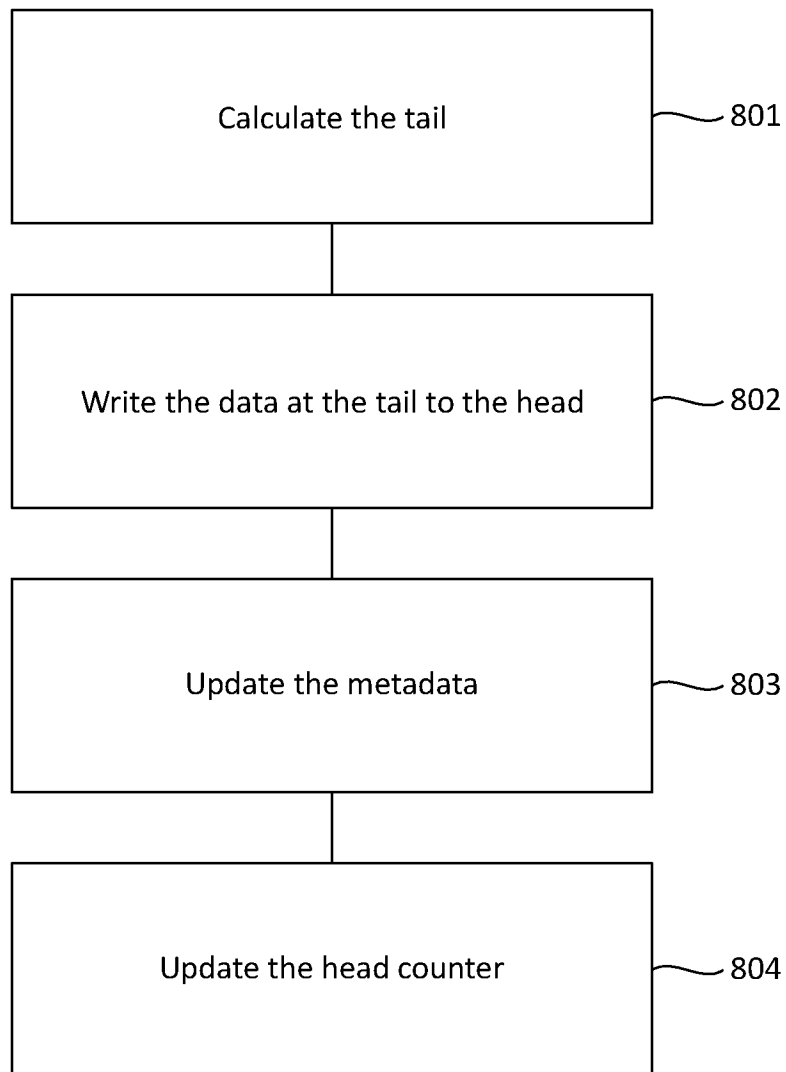


FIG. 6

**FIG. 7****FIG. 8**

		0	1	2	3	4	5	6	7
900	Drive 0	a h(a)	cd h(cd)	f h(f)	g h(g)	ij h(ij)	l h(l)		
	Drive 1	b h(b)	c h(c)	ef h(e)	h h(h)	i h(i)	kl h(kl)		
	Drive 2	ab h(ab)	d h(d)	e h(e)	gh h(gh)	j h(j)	k h(k)		

		0	1	2	3	4	5	6	7
901	Drive 0	a h(a)	cd h(cd)	f h(f)	g h(g)	ij h(ij)	l h(l)	b h(b)	
	Drive 1	b h(b)	c h(c)	ef h(e)	h h(h)	i h(i)	kl h(kl)		
	Drive 2	ab h(ab)	d h(d)	e h(e)	gh h(gh)	j h(j)	k h(k)		

		0	1	2	3	4	5	6	7
902	Drive 0	a h(a)	cd h(cd)	f h(f)	g h(g)	ij h(ij)	l h(l)	b h(b)	
	Drive 1	b h(b)	c h(c)	ef h(e)	h h(h)	i h(i)	kl h(kl)	e h(e)	
	Drive 2	ab h(ab)	d h(d)	e h(e)	gh h(gh)	j h(j)	k h(k)	be h(be)	

FIG. 9

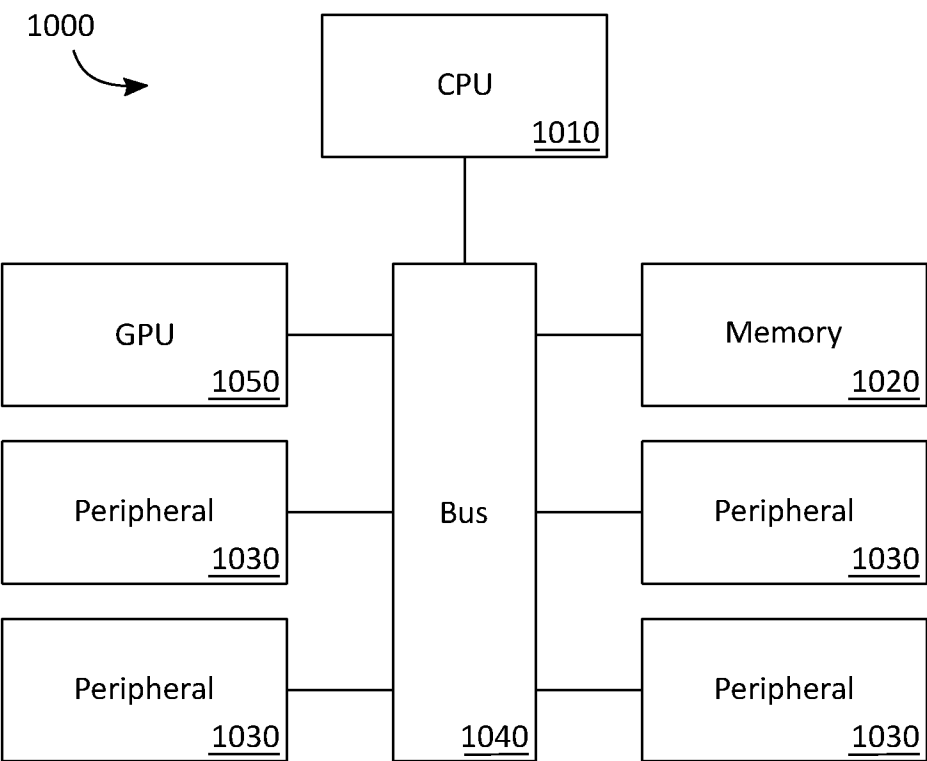


FIG. 10