

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
12 June 2008 (12.06.2008)

PCT

(10) International Publication Number
WO 2008/070800 A1

(51) International Patent Classification:
G06F 3/06 (2006.01) G11C 16/06 (2006.01)
G06F 12/02 (2006.01)

(71) Applicants and
(72) Inventors: FLYNN, David [US/US]; 8856 Shady Meadow Drive, Sandy, Utah 84093 (US). LAGERST-EDT, Bert [US/US]; 1234 E. 50 N., Pleasant Grove, UT 84062 (US). STRASSER, John [US/US]; 2323 South, Syracuse, UT 84075 (US). THATCHER, Jonathan [US/US]; 2259 North 2080 West, Lehi, UT 84043 (US). ZAPPE, Michael [US/US]; 4615 Simms Street, Wheatridge, CO 80033 (US). WALKER, John [US/US]; 11037 S. 1280 E., Sandy, UT 84094 (US).

(21) International Application Number:
PCT/US2007/086688

(22) International Filing Date:
6 December 2007 (06.12.2007)

(74) Agents: NEEDHAM, Bruce et al.; 8 East Broadway, Suite 600, Salt Lake City, UT 84111 (US).

(25) Filing Language: English

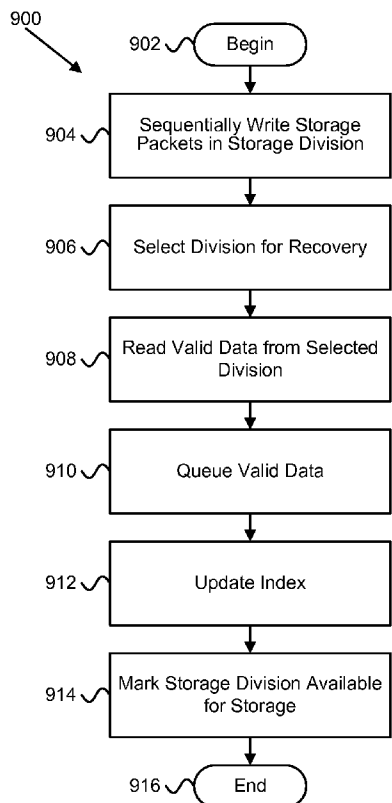
(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PG, PH, PL,

(26) Publication Language: English

(30) Priority Data:
60/873,111 6 December 2006 (06.12.2006) US
60/974,470 22 September 2007 (22.09.2007) US

[Continued on next page]

(54) Title: APPARATUS, SYSTEM, AND METHOD FOR STORAGE SPACE RECOVERY IN SOLID-STATE STORAGE



(57) Abstract: An apparatus, system and method are disclosed for storage space recovery in solid-state storage 110. A sequential storage module 802 sequentially writes data packets in a storage division. The storage division includes a portion of a solid-state storage 110. The data packets are derived from an object. The data packets are sequentially stored by order of processing. A storage division.selection module 804 selects a storage division for recovery. A data recovery module 806 reads valid data packets from the storage division selected for recovery, queues the valid data packets with other data packets to be written sequentially, and updates an index with a new physical address of the valid data. The index includes a mapping of physical addresses of data packets to object identifiers. A storage division recovery module 808 marks the storage division selected for recovery as available for sequentially written data packets in response to completing copying valid data from the storage division.

WO 2008/070800 A1



PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, SV, SY,
TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA,
ZM, ZW.

FR, GB, GR, HU, IE, IS, IT, LT, LU, LV, MC, MT, NL, PL,
PT, RO, SE, SI, SK, TR), OAPI (BF, BJ, CF, CG, CI, CM,
GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

(84) Designated States (*unless otherwise indicated, for every kind of regional protection available*): ARIPO (BW, GH, GM, KE, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI,

Published:

- *with international search report*
- *with amended claims*

Date of publication of the amended claims: 31 July 2008

AMENDED CLAIMS**received by the International Bureau on 27 May 2008 (27.05.2008)**5 **CLAIMS**

1. An apparatus for storage recovery, the apparatus comprising:
 - a sequential storage module that sequentially writes data packets in a storage division, the storage division comprising a portion of a solid-state storage in a solid-state storage device, each data packet derived from an object;
 - 10 a storage division selection module that selects a storage division for recovery;
 - a data recovery module that reads valid data packets from the storage division selected for recovery, queues the valid data packets with other data packets to be written sequentially by the sequential storage module, and updates an index with a new physical address of the valid data written by the sequential storage module, the index comprising a mapping of physical addresses of data packets stored in the
 - 15 solid-state storage to object identifiers corresponding to the data packets; and
 - a storage division recovery module that marks the storage division selected for recovery as available to the sequential storage module for sequentially writing data packets, the storage division marked for recovery in response to the data recovery module completing copying valid data from the storage division,
 - 20 wherein the data packets are stored sequentially on the solid-state storage device in an order that the data packets are received from one or more clients and from a storage division marked for recovery such that the sequentially stored data packets are stored contiguously and comprise the data packets received from the one or more clients interspersed with the data packets received from a storage
 - 25 division marked for recovery.
2. The apparatus of claim 1, wherein the sequential storage module, the storage division selection module, the data recovery module, and the storage division recovery module are in a solid-state storage device controller within the solid-state storage device and the
- 30 object index is located in the solid-state storage device and created and maintained by the solid-state storage device controller.
3. The apparatus of claim 1, wherein the storage division comprises an erase block and further comprising an erase module that erases an erase block selected for recovery prior to the storage division recovery module marking the erase block available for recovery,

erasing an erase block selected for recovery in response to the data recovery module copying valid data packets from the selected erase block.

4. The apparatus of claim 3, wherein the solid-state storage is arranged in two or more banks and the erase module erases an erase block in a first bank while the sequential storage module writes data packets to a storage division in a second bank.

17. The apparatus of claim 17, wherein identifying a storage division with a high amount of wear comprises identifying a storage division with a high number of erase cycles.
18. The apparatus of claim 17, wherein the storage division recovery module marks the storage division selected for recovery as unavailable for storing data.
19. The apparatus of claim 17, wherein identifying a storage division with a high amount of wear comprises identifying a storage division with a high bit error rate.
20. The apparatus of claim 1, wherein identifying a storage division with a high bit error rate comprises identifying a storage division with a non-recoverable error-correcting code ("ECC") block.
21. The apparatus of claim 1, wherein the storage division selection module selects a storage division for recovery by identifying a storage division with a high amount of invalid data.
22. The apparatus of claim 1, wherein the storage division selection module selects a storage division for recovery by identifying a storage division using a combination of two or more of program count, erase cycle count, bit error rate, amount of invalid data, erase cycles, and a non-recoverable ECC block.
23. The apparatus of claim 1, wherein the solid-state storage is arranged in an array of storage elements and a storage division comprises a virtual storage division that includes a storage division within each storage element in a row of storage elements accessed together.
24. A system for storage recovery, the system comprising:
 - a solid-state storage device comprising a solid-state storage, the solid-state storage comprising one or more storage divisions, each storage division comprising a portion of the solid-state storage in a solid-state storage device, the solid-state storage device comprising
 - a sequential storage module that sequentially writes data packets in a storage division, each data packet derived from an object;
 - a storage division selection module that selects a storage division for recovery;
 - a data recovery module that reads valid data packets from the storage division selected for recovery, queues the valid data packets with other data packets to be written sequentially by the sequential storage module, and updates an index with a new physical address of the valid data written by the

sequential storage module, the index comprising a mapping of physical addresses of data packets stored in the solid-state storage to object identifiers corresponding to the data packets; and

a storage division recovery module that marks the storage division selected for recovery as available to the sequential storage module for sequentially writing data packets, the storage division marked for recovery in response to the data recovery module completing copying valid data from the storage division;

wherein the data packets are stored sequentially on the solid-state storage device in an order that the data packets are received from one or more clients and from a storage division marked for recovery such that the sequentially stored data packets are stored contiguously and comprise the data packets received from the one or more clients interspersed with the data packets received from a storage division marked for recovery.

25. A computer program product comprising a computer readable medium having computer usable program code executable to perform operations for storage recovery, the operations of the computer program product comprising:
- sequentially writing data packets in a storage division, the storage division comprising a portion of a solid-state storage in a solid-state storage device, each data packet derived from an object;
 - selecting a storage division for recovery;
 - reading valid data packets from the storage division selected for recovery;
 - queuing the valid data packets with other data packets queued to be written sequentially;
 - updating an index with a new physical address of the valid data written by the sequential storage module, the index comprising a mapping of physical addresses of data packets stored in the solid-state storage to object identifiers corresponding to the data packets; and
 - marking the storage division selected for recovery as available for sequentially writing data packets, the storage division marked for recovery in response to completing copying valid data from the storage division;

wherein the data packets are stored sequentially on the solid-state storage device in an order that the data packets are received from one or more clients and from a storage division marked for recovery such that the sequentially stored data packets are stored contiguously and comprise the data packets received from the one or more clients interspersed with the data packets received from a storage division marked for recovery.

26. The computer program product of claim 25, wherein the storage division comprises an erase block and further comprising
- erasing an erase block selected for recovery prior to the storage division recovery module marking the erase block available for recovery, erasing an erase block selected for recovery in response to copying valid data packets from the selected erase block;
 - and
 - identifying a data packet in a storage division as invalid in response to an operation indicating that the data packet is no longer valid.