



(43) International Publication Date
20 November 2014 (20.11.2014)

- (51) International Patent Classification:
G06F 11/08 (2006.01)
- (21) International Application Number:
PCT/US2014/038510
- (22) International Filing Date:
16 May 2014 (16.05.2014)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
13/895,928 16 May 2013 (16.05.2013) US

(71) Applicants: WESTERN DIGITAL TECHNOLOGIES, INC. [US/US]; 3355 Michelson Drive, Suite 100, Irvine, California 92612 (US). SKYERA, INC. [US/US]; 1704 Automation Parkway, San Jose, California 95131 (US).

(72) Inventors: FLINSBAUGH, Jack W.; 350 River Oaks Parkway, #1456, San Jose, California 95134 (US). MULLENDORE, Rodney M.; 1704 Automation Parkway, San Jose, California 95131 (US).

(74) Agents: CHAN, David H. et al.; Western Digital Technologies, Inc., 3355 Michelson Drive, Suite 100, Irvine, California 92612 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT,

HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, 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, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Declarations under Rule 4.17:

- as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))
- as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))

Published:

- with international search report (Art. 21(3))
- before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments (Rule 48.2(h))

(54) Title: HIGH PERFORMANCE READ-MODIFY-WRITE SYSTEM PROVIDING LINE-RATE MERGING OF DATA-FRAME SEGMENTS IN HARDWARE




FIG. 1A

(57) Abstract: A method of merging data frames includes: receiving a first data frame having a plurality of sectors; receiving a second data frame having a plurality of sectors; generating a merged output data frame by merging, using a plurality of data paths including a plurality of multiplexers, sectors of the second data frame with sectors of the first data frame; and performing an error check on at least one check-data frame having sectors corresponding to those in the first data frame or the second data frame, where at least some of the sectors in the check-data frame are transmitted on a subset of the plurality of data paths that transmits sectors of the merged output data frame, and where the error check verifies the merged output data frame.

WO 2014/186779 A1

HIGH PERFORMANCE READ-MODIFY-WRITE SYSTEM PROVIDING LINE-RATE MERGING OF DATAFRAME SEGMENTS IN HARDWARE

BACKGROUND

1. Technical Field

[0001] This disclosure relates to data storage systems for computer systems. More particularly, the disclosure relates to ensuring that no undetectable data errors are introduced during a data modification process.

2. Related Art

[0002] Data storage systems provide storage for data of a host system. As solid-state storage capacity increases, semiconductor device features shrink. The smaller devices are more susceptible to “soft” errors caused by, for example, alpha-particles. While these soft errors do not cause permanent device damage, they corrupt data stored in devices such as registers and memories. Error checking codes such as cyclic redundancy codes (CRC) are used to determine whether a data error has occurred over data for which the CRC was calculated.

BRIEF DESCRIPTION OF THE DRAWINGS

[0003] Aspects and features of the present inventive concept will be more apparent by describing example embodiments with reference to the accompanying drawings, in which:

[0004] FIGS. 1A and 1B are block diagrams illustrating a read-modify-write system according to example embodiments of the present inventive concept;

[0005] FIG. 2 is a schematic illustration showing details of the read-modify-write system according to an example embodiment of the present inventive concept;

[0006] FIG. 3 is a diagram illustrating a method of merging of a first data frame and a second data frame using a merge mask according to an example embodiment of the present inventive concept;

[0007] FIG. 4 is a flowchart illustrating a method of generating a merged output data frame according to an example embodiment of the present inventive concept;

[0008] FIG. 5 is a flowchart illustrating a method of generating a merged output data frame based on a merge mask according to an example embodiment of the present inventive concept;

[0009] FIG. 6 is a flow chart illustrating read-modify-write operations according to an example embodiment of the present inventive concept; and

[0010] FIG. 7 is a flowchart illustrating merge operations according to an example embodiment of the present inventive concept.

DETAILED DESCRIPTION

[0011] While certain embodiments are described, these embodiments are presented by way of example only, and are not intended to limit the scope of protection. The methods and systems described herein may be embodied in a variety of other forms. Furthermore, various omissions, substitutions, and changes in the form of the example methods and systems described herein may be made without departing from the scope of protection.

Overview

[0012] Solid state drives (SSDs) are employed as data storage systems for computer systems. To modify data stored in a memory location of a data storage system, a read-modify-write operation may be employed to read a memory location and write a new value into it. As storage capacity of SSDs increases, features sizes of the semiconductor devices forming the SSDs decreases. As a result of the smaller device feature sizes, SSDs are increasingly susceptible to soft errors which may corrupt data during read-modify-write operations.

System Overview

[0013] Some embodiments of the present inventive concept are directed to a Read Modify Write system and method which ensure that no undetectable data errors are introduced during a data modification process.

[0014] In one example embodiment, the Read-Modify-Write system (RMW) receives L-Page frames (i.e., logical page data frames having a plurality of sectors) via two channels. The RMW system operates on the received frames based on their types as determined by a merge mask, applying hardware optimization to what were formerly software processes. The RMW system can merge data frame sectors on-

the-fly across frames and performs the appropriate bit-modifications based on the merge mask. The RMW system may minimize the firmware cycles required to perform read-modify-write operations.

[0015] FIG. 1A is a block diagram illustrating a RMW system according to an example embodiment of the present inventive concept. As illustrated in FIG. 1A, the read-modify-write system 100 includes circuitry including a storage 110, input logic 120, a plurality of multiplexers 130, a first logic device 140, and a second logic device 150. The RMW system 100 is synchronized to a system clock 160 and may be scalable with system clock speed. In one embodiment, FIG. 1A offers a hardware-based approach that provides significant speed improvement in handling read-modify-write operations, which helps meet the high throughput requirements in many data storage devices today. As illustrated in FIG. 1B, the read-modify-write system according to an example embodiment may further include a processor 170 which controls at least some operations of the RMW system. In one example embodiment, the firmware executed on the processor 170 would control at least a portion of the logic checking described herein.

Read-Modify-Write

[0016] In operation, in one example embodiment the storage 110 receives a first data frame having a plurality of sectors and stores the first data frame in the storage 110 until a second data frame having a plurality of sectors is received at a data port 122. The second data frame is passed through the RMW system 100 sector by sector at about the system clock speed and is merged with the first data frame which is read from the storage 110. The first and second data frames are processed by the first logic device 140 to determine whether the merged output data frame is valid.

[0017] FIG. 2 is a schematic illustration showing details of the read-modify-write system according to an example embodiment of the present inventive concept. Referring to FIG. 2, the input logic 120 receives the first and second data frames and verifies CRCs contained in the headers of each data frame. It must be noted that although CRCs are mentioned in the example embodiments in this disclosure, the use of other error detection mechanisms for example, but not limited to, checksums and hash functions, etc., are possible in some embodiments. The first and second

data frames are transferred sector by sector from the input logic 120 in synchronization with the system clock 160 to the multiplexers 130 where the sectors of the first and second data frames are merged at about the same speed as the system clock speed based on a merge mask.

[0018] FIG. 3 is a diagram illustrating merging of a first data frame and a second data frame using a merge mask according to an example embodiment of the present inventive concept. In the example illustration of FIG. 3, the merge mask indicates a sector to merge into the output data frame with a "1" and a sector to discard with a "0" with respect to the second data frame, and indicates a sector to merge into the output data frame with a "0" and a sector to discard with a "1" with respect to the first data frame. One of ordinary skill in the art will understand that the designations "1" and "0" are merely exemplary and that other designations are possible without departing from the scope of the inventive concept.

[0019] Referring to FIG. 3, a merge mask 370 specifies sectors of the second data frame 340 which will be merged with sectors of the first data frame 310 in place of sectors of the first data frame 310 which will be discarded (shown as sectors to discard 320). Sectors of the second data frame 340 which are not specified to merge into the output data frame by the merge mask 370 are discarded (shown as sectors to discard 350) while sectors of the second data frame 340 which are specified to merge into the output data frame by the merge mask 370 (shown as sectors to merge 360) are merged with the first data frame 310 to generate a merged output data frame 380. Conversely, sectors of the first data frame 310 which are not specified to merge into the output data frame by the merge mask 370 are discarded (shown as sectors to discard 320) while sectors of the first data frame 310 which are specified to merge into the output data frame (shown as sectors to merge 330) by the merge mask 370 are merged with the sectors to merge 360 of the second data frame 340 to generate the merged output data frame 380. In other words, sectors of the first data frame specified by the merge mask are replaced by corresponding sectors of the second data frame in order to generate the merged output data frame and the unused sectors of the first and second data frames are discarded.

[0020] While the above merge operations are described in terms of specified sectors of the second data frame which may be merged with sectors of the first data

frame, one of ordinary skill in the art will appreciate that the scope of the present inventive concept also includes merging specified sectors of the first data frame with sectors of the second data frame. Also, while illustrated as such in FIG. 3, the number of sectors in the plurality of sectors of the first data frame and the second data frame may not be identical.

[0021] Referring again to FIG. 2, in one example embodiment each sector of the second data frame is transmitted to a first multiplexer 5 on a sector by sector basis in synchronization with the system clock. In one example embodiment, the first multiplexer 5 transmits each sector of the second data frame to a second multiplexer 6 and a merge multiplexer 8. Simultaneously, the first data frame is transmitted to a third multiplexer 4. In one example embodiment, the third multiplexer 4 transmits each sector of the first data frame to a fourth multiplexer 15 and the merge multiplexer 8. The merge multiplexer 8 is controlled by the merge mask to select a sector from the first or second data frame to be merged into the output data frame.

[0022] In one example embodiment, the sector from the first or second data frame selected by the merge mask for merging is output by the merge multiplexer 8 as a sector of the merged output data frame and is also transmitted from the output of the merge multiplexer 8 to the second logic device 150. The second logic device 150 calculates and accumulates a CRC for the merged output data frame. Each sector of the merged output data frame is output from the RMW system 100 by an output multiplexer 17 on a sector by sector basis. Generation of the merged output data frame is performed at about the same speed as the system clock speed.

[0023] In one example embodiment, the second multiplexer 6 and the fourth multiplexer 15 select either the sector that is merged into the output data frame or the sector that is discarded and transmit the selected or discarded sector to the first logic device 140. For example, if the merge mask causes the merge multiplexer 8 to select a sector from the first data frame, the corresponding sector from the second data frame is discarded. The merge multiplexer 8 transmits the selected sector from the first data frame to the output multiplexer 17 and the output multiplexer 17 outputs the selected sector from the first data frame as a sector of the merged output data frame and also transmits the selected sector from the first data frame to the fourth

multiplexer 15. The fourth multiplexer 15 selects the selected sector from the first data frame and transmits it to the first logic device 140. Simultaneously, the discarded sector of the second data frame is transmitted from the second multiplexer 6 to the first logic device 140.

[0024] In one example embodiment, the first logic device 140 reconstructs the first and second data frames as first and second check-data frames, respectively, and the CRCs of the check-data frames are calculated and accumulated by the first logic device 140 on a sector by sector basis. The first check-data frame may include sectors that are transmitted on a first subset of data paths different from the subset of data paths that sectors of the first data frame merged into the output data frame are transmitted on, and the second check-data frame may include sectors that are transmitted on a second subset of data paths different from the subset of data paths that sectors of the second data frame merged into the output data frame are transmitted on.

[0025] After processing the last sector of the first and second data frames, the CRC of the first data frame is transmitted through the third multiplexer 4 and fourth multiplexer 15 to the first logic device 140. Simultaneously, the CRC of the second data frame is transmitted through the first multiplexer 5 and second multiplexer 6 to the first logic device 140.

[0026] In one example embodiment, the first logic device 140 performs an error check by comparing the transmitted CRC of the first data frame to the CRC calculated for the first check-data frame and compares the transmitted CRC of the second data frame to the CRC calculated for the second check-data frame. If the received data frame CRCs matches the CRCs of the check-data frames, the merged output data frame will be valid since sectors of the merged output data frame passed through the same logic as the sectors of the check-data frames and no errors were introduced. If the CRCs do not match, a message is generated indicating that the merged output data frame is invalid. The validity can be verified this way because at least some of the sectors in the check-data frame are transmitted on a subset of the plurality of data paths that also transmits sectors of the merged output data frame, so that the error check of the at least one check-data frame verifies the merged output data frame. The use of output data from the actual merged data frame as part of the

calculation of the check-data frame CRC ensures that any error introduced along the path of the frame merging mechanism is detected.

[0027] Concurrently with the generation of the merged output data frame, the second logic device 150 generates a CRC for the merged output data frame on a sector by sector basis. After the final sector of the merged output data frame is outputted, the CRC for the merged output data frame generated by the second logic device 150 is selected and output by the output multiplexer 17.

[0028] FIG. 4 is a flowchart illustrating a method of generating a merged output data frame according to an example embodiment of the present inventive concept. Referring to FIG. 4, a first data frame comprising a plurality of sectors is received by the RMW system (410). The first data frame may be stored in a data storage such as in volatile or non-volatile memory. A second data frame comprising a plurality of sectors is received by the RMW system (420). The second data frame may include, for example, new data from a host system to be written to the data storage. Sectors of the first and second data frames are merged via a plurality of data paths and multiplexers to generate a merged output data frame (430). In addition, a CRC for the merged output data frame is generated by calculating and accumulating on a sector by sector basis CRC data of the first and second data frame sectors merged to generate the merged output data frame (440).

[0029] The sectors of least one of the first and second data frames are reconstructed into a check-data frame as the merged output frame is being generated (450). Alternatively, both of the first and second data frames may be reconstructed into check-data frames as the merged output frame as being generated. A CRC is calculated and accumulated on a sector by sector basis for the check-data frame, and an error check is performed on the at least one check-data frame comprising sectors corresponding to those of the first data frame or the second data frame to verify the corresponding sectors of the merged output data frame (460). Since at least some of the sectors in the check-data frame are transmitted on a subset of the plurality of data paths that transmits sectors of the merged output data frame, the error check of the at least one check-data frame verifies the sectors from the at least one check-data frame in the merged output data frame. In one example embodiment, the error check of both check-data frames

verifies the sectors of the first and second check-data frames in the merged output data frame.

[0030] FIG. 5 is a flowchart illustrating a method of generating a merged output data frame based on a merge mask according to an example embodiment of the present inventive concept. Referring to FIG. 5, a first data frame comprising a plurality of sectors is received by the RMW system (510). The first data frame may be stored in a data storage such as in volatile or non-volatile memory. A second data frame comprising a plurality of sectors is received by the RMW system (520).

[0031] A merge mask which specifies sectors of the first and second data frames to be merged is received (530). The merge mask specifies which sectors of the first and second data frames occupy each sector of the output data frame. The merge mask may be received as part of the first received data frame or as part of the second received data frame. Alternatively, the merge mask may be received independently of the received data frames. A merged output data frame is generated by merging selected sectors from the first and second data frames based on the merge mask (540).

[0032] FIG. 6 is a flow chart illustrating read-modify-write operations according to an example embodiment of the present inventive concept. Referring to FIG. 6, a read-data frame to be modified is input from memory external to the RMW system and stored in an internal storage (605) until a write-data frame containing sectors for modifying the read-data frame is received (610). Each of the read-data frame and the write-data frame has a plurality of sectors. The number of sectors in the plurality of sectors of the read-data frame and the write-data frame may not be identical.

[0033] A merge mask is received (615) and the sectors of the read-data frame and the write-data frame to be merged in an output data frame are determined based on the merge mask (620). If a sector of the write-data frame is selected (625-Y), the selected write-data frame sector is merged into the output data frame (630). Alternatively, if the write-data frame sector is not selected (625-N), a corresponding read-data frame sector is merged into the output data frame (635). CRC data is calculated and accumulated for the selected sector (640) and concurrently the selected sector merged into the output data frame is outputted from the RMW system (645). The unselected data frame sector is discarded.

[0034] Concurrently with the merge process, the read-data frame and the write-data frame are reconstructed with their respective sectors (650) and CRC data is calculated and accumulated for each sector of the reconstructed read-data frame and the write-data frame (655). If the reconstructed sector is not the last sector to be processed (660-N), the process at operations 625 through 655 is repeated. If the reconstructed sector is the last sector to be processed (660-Y), the CRCs of the reconstructed read-data frame and the write-data frame are compared to the CRCs of the originally received read-data frame and the write-data frame to verify that no errors were introduced into the data during the read-modify-write process (665).

[0035] If either of the CRCs for the reconstructed data frames does not match the CRCs of the originally received data frames (670-N), a message is generated indicating that the merged output data frame is invalid (675). Since at least some of the sectors in the reconstructed data frames are transmitted on a subset of the plurality of data paths that transmits sectors of the merged output data frame, the error check of the reconstructed frames verify the merged output data frame. Otherwise, if CRCs for the reconstructed data frames match the CRCs of the originally received data frames (670-Y), the merged data frame is verified to be correct.

[0036] Finally, FIG. 7 is a flowchart illustrating merge operations according to an example embodiment of the present inventive concept. Referring to FIGS. 3 and 7, after initializing a sector counter (710), a sector of the read-data frame and a corresponding sector of a write-data frame are read (720). The write-data sector address is compared to a merge mask (730), and if the write-data sector corresponds to the merge mask (740-Y) the write-data sector is merged into the address slot of the output data frame (750). Alternatively, if the write-data sector does not correspond to the merge mask (740-N) the read-data sector is merged into the address slot of the output data frame (760). The sectors of the read-data frame and the write-data frame are transmitted to respective CRC test modules to reconstruct the original read-data frame and the original write-data frame (770). If the reconstructed sector is not the last sector to be processed (780-N), the process at operations 720 through 770 is repeated.

[0037] While the examples explained with reference to FIG. 6 and FIG. 7 are described in terms of a read-data frame and a write-data frame, one of ordinary skill in the art will appreciate that the present inventive concept may be generalized to a first data frame and a second data frame, either of which may correspond to a read-data frame and a write-data frame.

[0038] Further, operations of the present inventive concept may be performed in the order described, in a different order, or operations may be combined. One of ordinary skill in the art will appreciate that the foregoing process is exemplary and that other variations are possible without departing from the inventive concept.

[0039] While certain embodiments have been described, these embodiments have been presented by way of example only, and are not intended to limit the scope of the protection. The methods and systems described herein may be embodied in a variety of other forms. Various omissions, substitutions, and/or changes in the form of the example methods and systems described herein may be made without departing from the spirit of the protection.

[0040] The accompanying claims and their equivalents are intended to cover such forms or modifications as would fall within the scope and spirit of the protection. For example, the example systems and methods disclosed herein can be applied to solid-state drives, hard disk drives, hybrid hard drives, and the like. In addition, other forms of storage, for example, but not limited to, DRAM or SRAM, battery backed-up volatile DRAM or SRAM devices, EPROM, EEPROM memory, etc., may additionally or alternatively be used. As another example, the various components illustrated in the figures may be implemented as software and/or firmware on a processor, ASIC/FPGA, or dedicated hardware. Also, the features and attributes of the specific example embodiments disclosed above may be combined in different ways to form additional embodiments, all of which fall within the scope of the present disclosure.

[0041] Although the present disclosure provides certain example embodiments and applications, other embodiments that are apparent to those of ordinary skill in the art, including embodiments which do not provide all of the features and advantages set forth herein, are also within the scope of this disclosure. Accordingly, the scope of the present disclosure is intended to be defined only by reference to the appended claims.

WHAT IS CLAIMED IS:

1. A method of merging data frames, the method comprising:
 - receiving a first data frame comprising a plurality of sectors;
 - receiving a second data frame comprising a plurality of sectors;
 - generating a merged output data frame by merging, using a plurality of data paths comprising a plurality of multiplexers, sectors of the second data frame with sectors of the first data frame; and
 - performing an error check on at least one check-data frame that comprises sectors corresponding to those in the first data frame or the second data frame, wherein at least some of the sectors in the check-data frame are transmitted on a subset of the plurality of data paths that transmits sectors of the merged output data frame, and
 - wherein the error check verifies the merged output data frame.

2. The method of claim 1, wherein performing the error check further comprises:
 - performing an error check on two check-data frames including:
 - a first check-data frame including sectors corresponding to those in the first data frame; and
 - a second check-data frame including sectors corresponding to those in the second data frame.

3. The method of claim 2, wherein the first check-data frame and the second check-data frame are reconstructed with their respective sectors that correspond to the first and second data frames after the sectors are transmitted through selected ones of the plurality of multiplexers based on a merge mask that indicates which sectors of the first and second data frames are to be merged into the output data frame.

4. The method of claim 2, wherein performing the error check further comprises:
 - comparing cyclic redundancy codes (CRCs) of the first data frame and second data frame with CRCs generated for the first and second check-data frames, respectively, and
 - if either CRC comparison does not match, generating a message indicating the merged output data frame is invalid.
5. The method of claim 4, wherein the CRC of the first check-data frame is generated by accumulating CRC data of the first check-data frame on a sector-by-sector basis and the CRC of the second check-data frame is generated by accumulating CRC data of the second check-data frame on a sector-by-sector basis.
6. The method of claim 2, wherein each sector in the merged output data frame is in either the first check-data frame or the second check-data frame.
7. The method of claim 2, wherein:
 - the first check-data frame further comprises sectors that are transmitted on a first subset of data paths different from the subset of data paths that sectors of the first data frame merged into the output data frame are transmitted on, and
 - the second check-data frame further comprises sectors that are transmitted on a second subset of data paths different from the subset of data paths that sectors of the second data frame merged into the output data frame are transmitted on.
8. The method of claim 1, wherein the at least one check-data frame further comprises sectors that are transmitted on a subset of data paths different from the subset of data paths that sectors of the merged output data frame are transmitted on.
9. The method of claim 1, wherein the sectors of at least one of the first and second data frames to be merged into the output data frame are specified by a merge mask.

10. The method of claim 9, wherein the merge mask specifies which sectors of the first and second data frames occupy each sector of the output data frame.

11. The method of claim 1, further comprising reconstructing the at least one check-data frame as the merged output data frame is generated.

12. The method of claim 1, further comprising generating a cyclic redundancy code (CRC) for the merged output data frame by accumulating CRC data of the sectors merged into the output data frame on a sector-by-sector basis.

13. An apparatus for merging data frames, the apparatus comprising:
one or more devices configured to receive a first data frame comprising a plurality of sectors and a second data frame comprising a plurality of sectors; and
circuitry comprising a plurality of data paths comprising a plurality of multiplexers, the circuitry configured to:

generate, using the plurality of data paths, a merged output data frame by merging sectors of the second data frame with sectors of the first data frame; and

perform an error check on at least one check-data frame that comprises sectors corresponding to those in the first data frame or the second data frame,

wherein at least some of the sectors in the check-data frame are transmitted on a subset of the plurality of data paths that transmits sectors of the merged output data frame,

wherein the error check verifies the merged output data frame.

14. The apparatus of claim 13, wherein the one or more devices further comprises:

a storage configured to store the first data frame; and

a data port configured to receive the second data frame.

15. The apparatus of claim, 13 wherein the circuitry is further configured to perform the error check by at least:

performing an error check on two check-data frames including:

a first check-data frame including sectors corresponding to those in the first data frame; and

a second check-data frame including sectors corresponding to those in the second data frame.

16. The apparatus of claim 15, wherein the first check-data frame and the second check-data frame are reconstructed with their respective sectors that correspond to the first and second data frames by at least one first logic device after the sectors are transmitted through selected ones of the plurality of multiplexers based on a merge mask that indicates which sectors of the first and second data frames are to be merged into the output data frame.

17. The apparatus of claim 16, wherein the at least one first logic device is configured to:

compare a cyclic redundancy code (CRC) of the first data frame with a CRC generated for the first check-data frame;

compare a CRC of the second data frame with a CRC generated for the second check-data frame; and

if either CRC comparison does not match, generate a message indicating the merged output data frame is invalid.

18. The apparatus of claim 17, wherein the at least one first logic device is further configured to:

generate the CRC of the first check-data frame by accumulating CRC data of the first check-data frame on a sector-by-sector basis; and

generate the CRC of the second check-data frame by accumulating CRC data of the second check-data frame on a sector-by-sector basis.

19. The apparatus of claim 15, wherein each sector in the merged output data frame is in either the first check-data frame or the second check-data frame.

20. The apparatus of claim 15, wherein:

the first check-data frame further comprises sectors that are transmitted on a first subset of data paths different from the subset of data paths that sectors of the first data frame merged into the output data frame are transmitted on, and

the second check-data frame further comprises sectors that are transmitted on a second subset of data paths different from the subset of data paths that sectors of the second data frame merged into the output data frame are transmitted on.

21. The apparatus of claim 13, wherein the at least one check-data frame further comprises sectors that are transmitted on from a subset of data paths different from the subset of data paths that sectors of the merged output data frame are transmitted on.

22. The apparatus of claim 13, wherein the sectors of the at least one of the first and second data frames to be merged into the output data frame are specified by a merge mask.

23. The apparatus of claim 22, wherein the merge mask specifies which sectors of the first and second data frames occupy each sector of the output data frame.

24. The apparatus of claim 22, wherein the merge mask is incorporated into one of the first and the second data frames.

25. The apparatus of claim 22, wherein the merge mask is received separately from the first and second data frames.

26. The apparatus of claim 13, further comprising at least one first logic device configured to reconstruct the at least one check-data frame as the merged output data frame is generated.

27. The apparatus of claim 13, further comprising a second logic device configured to generate a cyclic redundancy code (CRC) for the merged output data frame by accumulating CRC data of the sectors merged into the output data frame on a sector-by-sector basis.

28. The apparatus of claim 13, wherein the circuitry further comprises a processor, and wherein generating the merged output data frame and performing the error check are performed under the control of the processor.

29. A method of merging data frames, the method comprising:
receiving a first data frame comprising a plurality of sectors;
receiving a second data frame comprising a plurality of sectors; and
generating a merged output data frame based on the first data frame and the second data frame, the generating comprising, for each slot in the merged output data frame, selecting a sector from either the first data frame or the second data frame based on a merge mask.

30. The method of claim 29, wherein the merge mask is incorporated into the one of the first and second data frames.

31. The method of claim 29, wherein the merge mask is received separately from the first and second data frames.

32. An apparatus for merging data frames, the apparatus comprising:
a first circuit configured to receive a first data frame comprising a plurality of sectors;
a second circuit configured to receive a second data frame comprising a plurality of sectors; and
a plurality of data paths configured to generate a merged output data frame based on the first data frame and the second data frame,

wherein the generating comprises, for each slot in the merged output data frame, selecting a sector from either the first data frame or the second data frame based on a merge mask.

33. The apparatus of claim 32, wherein the first circuit comprises a storage configured to store the first data frame; and
the second circuit comprises a data port configured to receive the second data frame.

34. The apparatus of claim 32, wherein the merge mask is incorporated into the one of the first and second data frames.

35. The apparatus of claim 32, wherein the merge mask is received separately from the first and second data frames.




FIG. 1A




FIG. 1B

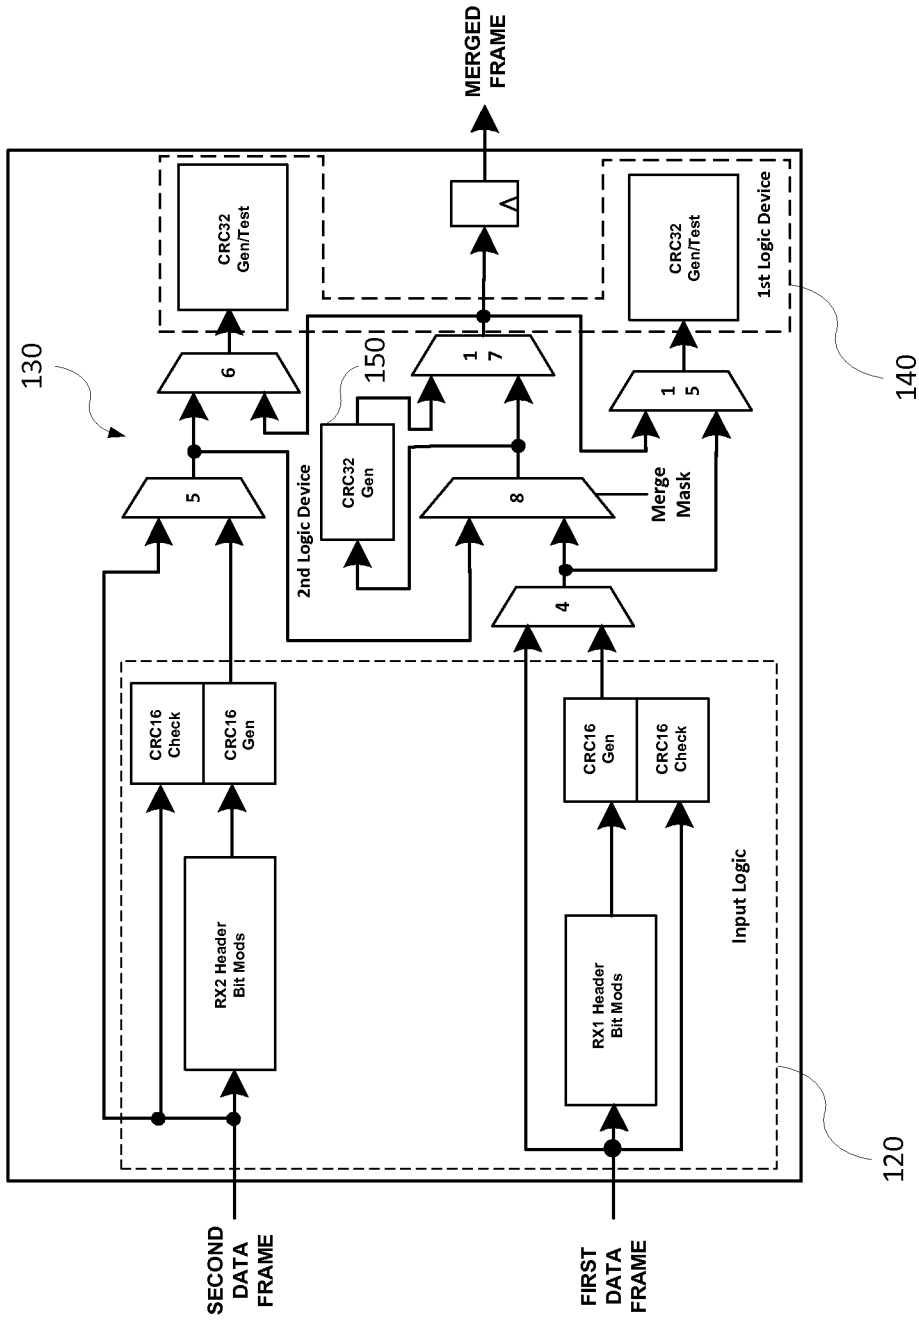


FIG. 2




FIG. 3

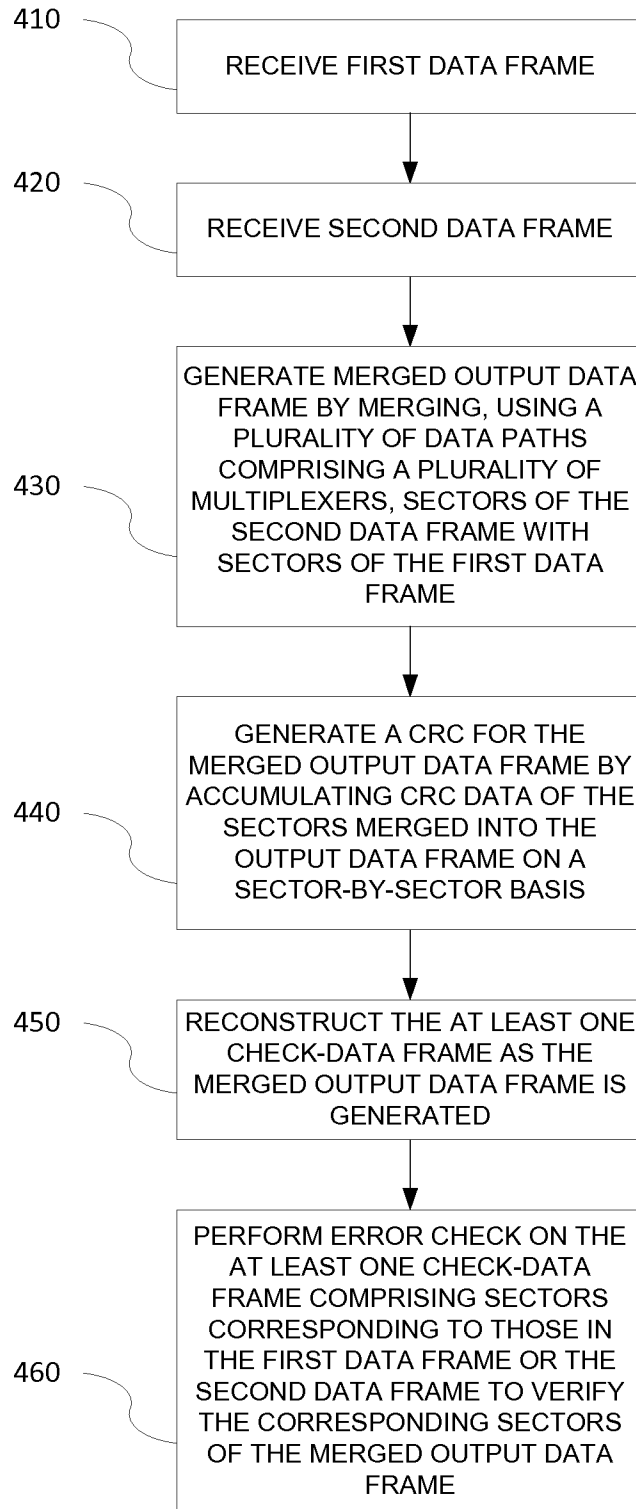
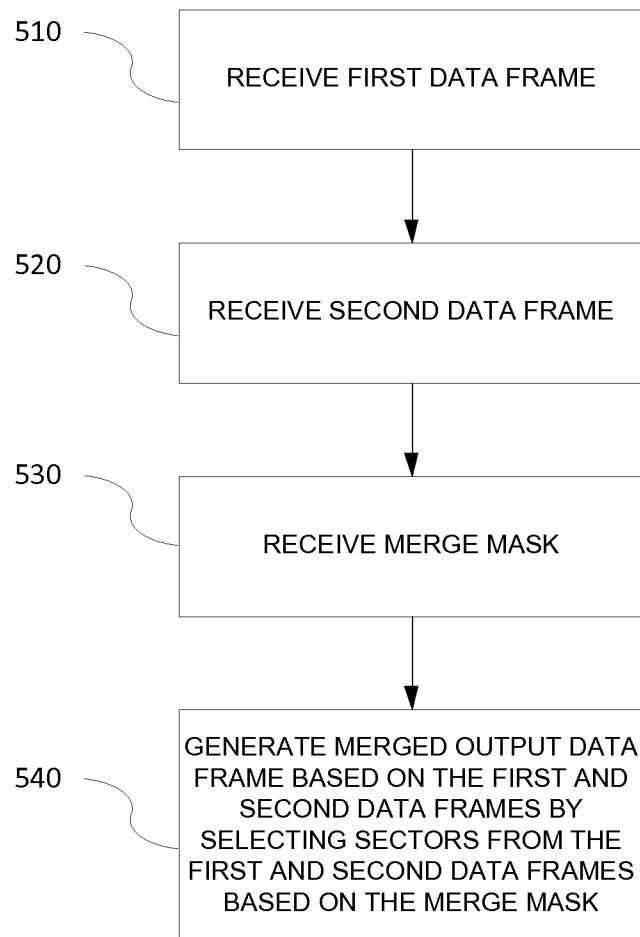


FIG. 4

**FIG. 5**

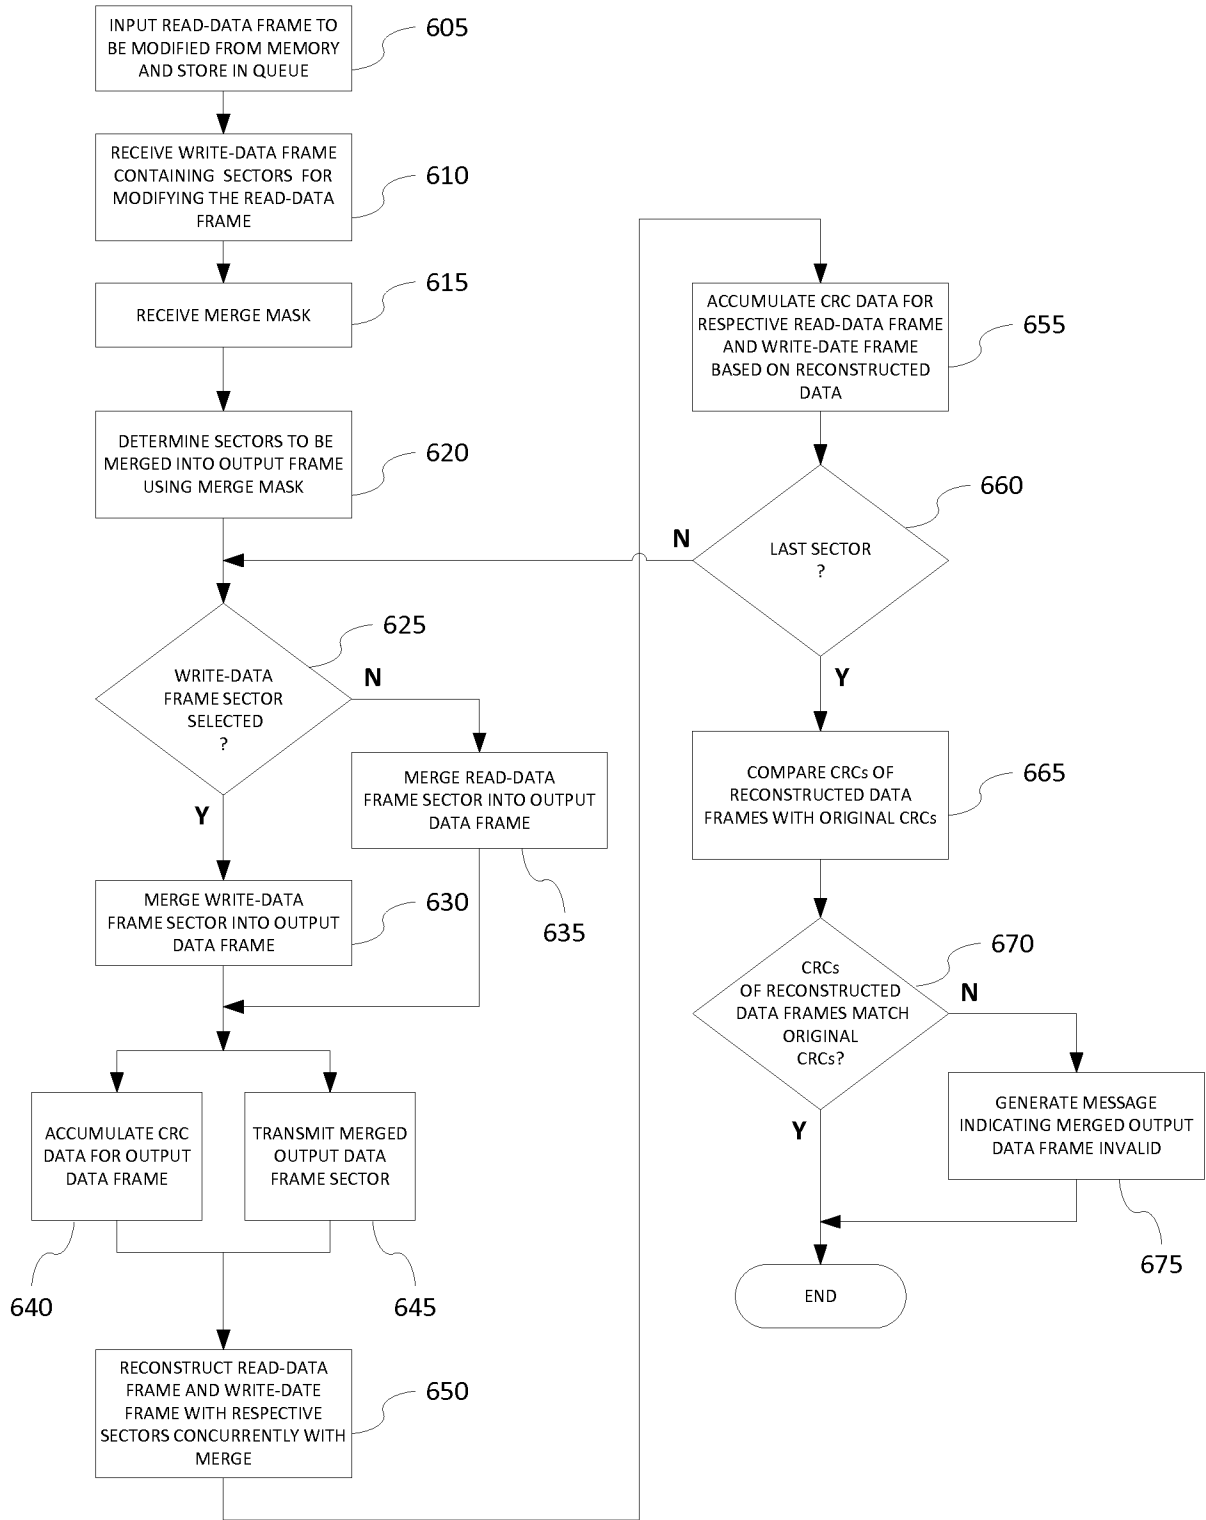


FIG. 6

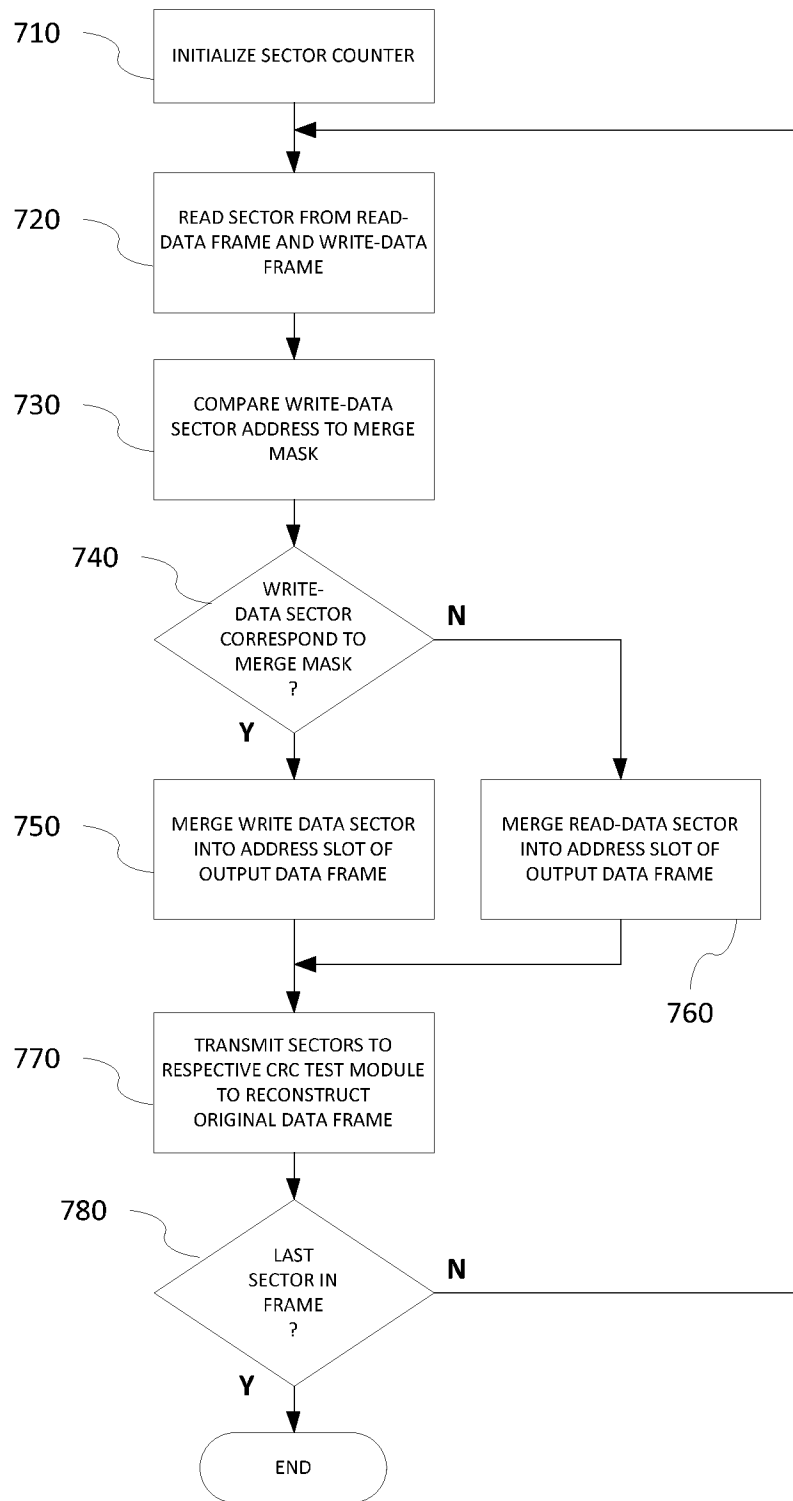


FIG. 7

A. CLASSIFICATION OF SUBJECT MATTER**G06F 11/08(2006.01)i**

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHEDMinimum documentation searched (classification system followed by classification symbols)
G06F 11/08; H03M 13/09; H03M 13/00; H03M 13/05; G06F 11/10; G06F 12/00; G06F 7/38Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched
Korean utility models and applications for utility models
Japanese utility models and applications for utility modelsElectronic data base consulted during the international search (name of data base and, where practicable, search terms used)
eKOMPASS(KIPO internal) & Keywords: data frame, merge, sector, path, CRC, error check, multiplexer**C. DOCUMENTS CONSIDERED TO BE RELEVANT**

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 2011-0239085 A1 (KWOK W. YEUNG et al.) 29 September 2011 See paragraphs [0019]-[0025], [0030]-[0035]; claim 1; and figure 1.	1-35
A	US 2012-0151172 A1 (PAUL W. COTEUS et al.) 14 June 2012 See paragraphs [0006]-[0010], [0030]-[0034]; claim 1; and figure 3.	1-35
A	US 2009-0222708 A1 (AKIRA YAMAGA) 03 September 2009 See paragraphs [0009]-[0010], [0042]-[0058]; and figure 8.	1-35
A	US 2010-0241936 A1 (RYAN E. HIRTH) 23 September 2010 See paragraphs [0005], [0029]-[0040]; claim 1; and figure 10.	1-35
A	US 2007-0136411 A1 (TOSHIHIRO KATASHITA et al.) 14 June 2007 See paragraphs [0018]-[0020]; and figure 8.	1-35

 Further documents are listed in the continuation of Box C. See patent family annex.

* Special categories of cited documents:

"A" document defining the general state of the art which is not considered to be of particular relevance

"E" earlier application or patent but published on or after the international filing date

"L" document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

"O" document referring to an oral disclosure, use, exhibition or other means

"P" document published prior to the international filing date but later than the priority date claimed

"T" later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

"X" document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

"Y" document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art

"&" document member of the same patent family

Date of the actual completion of the international search

23 October 2014 (23.10.2014)

Date of mailing of the international search report

24 October 2014 (24.10.2014)

Name and mailing address of the ISA/KR

International Application Division
Korean Intellectual Property Office
189 Cheongsu-ro, Seo-gu, Daejeon Metropolitan City, 302-701,
Republic of Korea

Facsimile No. +82-42-472-7140

Authorized officer

YU, JAE CHON

Telephone No. +82-42-481-8647



INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2014/038510

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2011-0239085 A1	29/09/2011	US 7958427 B1 US 8413009 B2	07/06/2011 02/04/2013
US 2012-0151172 A1	14/06/2012	CN 101300556 A CN 101300556 B EP 1958073 A1 JP 2009-517725 A JP 5186382 B2 TW I399649 B US 2007-160053 A1 US 2007-183331 A1 US 2007-286078 A1 US 2007-286199 A1 US 2012-151171 A1 US 7685392 B2 US 8145868 B2 US 8151042 B2 US 8327105 B2 US 8495328 B2 WO 2007-060250 A1	05/11/2008 19/05/2010 20/08/2008 30/04/2009 17/04/2013 21/06/2013 12/07/2007 09/08/2007 13/12/2007 13/12/2007 14/06/2012 23/03/2010 27/03/2012 03/04/2012 04/12/2012 23/07/2013 31/05/2007
US 2009-0222708 A1	03/09/2009	JP 2009-211742 A JP 4672743 B2 US 8312348 B2	17/09/2009 20/04/2011 13/11/2012
US 2010-0241936 A1	23/09/2010	US 8661314 B2	25/02/2014
US 2007-0136411 A1	14/06/2007	JP 2007-166031 A JP 4764973 B2 US 7590916 B2	28/06/2007 07/09/2011 15/09/2009