

Office Action in Ex Parte Reexamination	Control No. 90/019,777	Patent Under Reexamination 7739544	
	Examiner WILLIAM H WOOD	Art Unit 3992	AIA (First Inventor to File) Status No

-- The MAILING DATE of this communication appears on the cover sheet with the correspondence address --

- a. Responsive to the communication(s) filed on 11 July 2025.
 A declaration(s)/affidavit(s) under **37 CFR 1.130(b)** was/were filed on _____.
- b. This action is made FINAL.
- c. A statement under 37 CFR 1.530 has not been received from the patent owner.

A shortened statutory period for response to this action is set to expire 2 month(s) from the mailing date of this letter. Failure to respond within the period for response will result in termination of the proceeding and issuance of an *ex parte* reexamination certificate in accordance with this action. 37 CFR 1.550(d). **EXTENSIONS OF TIME ARE GOVERNED BY 37 CFR 1.550(c)**. If the period for response specified above is less than thirty (30) days, a response within the statutory minimum of thirty (30) days will be considered timely.

Part I THE FOLLOWING ATTACHMENT(S) ARE PART OF THIS ACTION:

1. Notice of References Cited by Examiner, PTO-892. 3. Interview Summary, PTO-474.
2. Information Disclosure Statement, PTO/SB/08. 4. _____.

Part II SUMMARY OF ACTION

- 1a. Claims 1-22 are subject to reexamination.
- 1b. Claims _____ are not subject to reexamination.
2. Claims _____ have been canceled in the present reexamination proceeding.
3. Claims _____ are patentable and/or confirmed.
4. Claims 1-22 are rejected.
5. Claims _____ are objected to.
6. The drawings, filed on _____ are acceptable.
7. The proposed drawing correction, filed on _____ has been (7a) approved (7b) disapproved.
8. Acknowledgment is made of the priority claim under 35 U.S.C. 119(a)-(d) or (f).
a) All b) Some* c) None of the certified copies have
1 been received.
2 not been received.
3 been filed in Application No. 11/369,736.
4 been filed in reexamination Control No. _____.
5 been received by the International Bureau in PCT application No. _____.
- * See the attached detailed Office action for a list of the certified copies not received.
9. Since the proceeding appears to be in condition for issuance of an *ex parte* reexamination certificate except for formal matters, prosecution as to the merits is closed in accordance with the practice under *Ex parte* Quayle, 1935 C.D. 11, 453 O.G. 213.
10. Other: _____

cc: Requester (if third party requester)



UNITED STATES PATENT AND TRADEMARK OFFICE

UNITED STATES DEPARTMENT OF COMMERCE
United States Patent and Trademark Office
Address: COMMISSIONER FOR PATENTS
P.O. Box 1450
Alexandria, Virginia 22313-1450
www.uspto.gov

Table with 5 columns: APPLICATION NO., FILING DATE, FIRST NAMED INVENTOR, ATTORNEY DOCKET NO., CONFIRMATION NO.
90/019,777 12/13/2024 7739544 63549-414701 1702

7590 08/15/2025
Quinn Emanuel Urquhart & Sullian, LLP
50 California Street
22nd Floor
San Francisco, CA 94111

EXAMINER

WOOD, WILLIAM H

ART UNIT PAPER NUMBER

3992

MAIL DATE DELIVERY MODE

08/15/2025

PAPER

Please find below and/or attached an Office communication concerning this application or proceeding.

The time period for reply, if any, is set in the attached communication.



UNITED STATES PATENT AND TRADEMARK OFFICE

Commissioner for Patents
United States Patent and Trademark Office
P.O. Box 1450
Alexandria, VA 22313-1450
www.uspto.gov

DO NOT USE IN PALM PRINTER

(THIRD PARTY REQUESTER'S CORRESPONDENCE ADDRESS)

BARNES & THORNBURG, LLP
555 12TH STREET, NW
SUITE 1200
WASHINGTON, DC 20004

***EX PARTE* REEXAMINATION COMMUNICATION TRANSMITTAL FORM**

REEXAMINATION CONTROL NO. 90/019,777 .

PATENT UNDER REEXAMINATION 7739544 .

ART UNIT 3992 .

Enclosed is a copy of the latest communication from the United States Patent and Trademark Office in the above identified *ex parte* reexamination proceeding (37 CFR 1.550(f)).

Where this copy is supplied after the reply by requester, 37 CFR 1.535, or the time for filing a reply has passed, no submission on behalf of the *ex parte* reexamination requester will be acknowledged or considered (37 CFR 1.550(g)).

DETAILED ACTION

Reexamination (*ex parte*) has been requested by a third party (Unified Patents, LLC) for claims 1-22 of US 7,739,544 to Yamato et al., which issued on 06/15/2010 (hereinafter, the '544 patent). The Request for *ex parte* reexamination (herein *Request*) was received on 12/13/2024.

A substantial new question of patentability affecting claims 1-22 of United States Patent Number 7,739,544 to Yamato et al. is raised by the request for *ex parte* reexamination. As such the filed request for reexamination was granted and reexamination order on 02/11/2025.

Notice of Pre-AIA or AIA Status

The present application is being examined under the pre-AIA first to invent provisions.

Reexamination

Extensions of time under 37 CFR 1.136(a) will not be permitted in these proceedings because the provisions of 37 CFR 1.136 apply only to "an applicant" and not to parties in a reexamination proceeding. Additionally, 35 U.S.C. 305 requires that *ex parte* reexamination proceedings "will be conducted with special dispatch" (37 CFR 1.550(a)). Extensions of time in *ex parte* reexamination proceedings are provided for in 37 CFR 1.550(c).

The patent owner is reminded of the continuing responsibility under 37 CFR 1.565(a) to apprise the Office of any litigation activity, or other prior or concurrent proceeding, involving Patent No. 7,739,544 throughout the course of this reexamination proceeding. The third party requester is also reminded of the ability to similarly apprise the Office of any such activity or proceeding throughout the course of this reexamination proceeding. See MPEP §§ 2207, 2282 and 2286.

Patent Owner is notified that any proposed amendment to the specification and/or claims in this reexamination proceeding must comply with 37 CFR 1.530(d)-(j), must be formally presented pursuant to 37 CFR 1.52(a) and (b), and must contain any fees required by 37 CFR 1.20(c). See MPEP 2250.

After filing of a request for ex parte reexamination by a third party requester, any document filed by either the patent owner or the third party requester must be served on the other party (or parties where two or more third party requester proceedings are merged) in the reexamination proceeding in the manner provided in 37 CFR 1.248. The document must reflect service or the document may be refused consideration by the Office. See 37 CFR 1.550(f).

Claim Rejections - 35 USC § 103

In the event the determination of the status of the application as subject to AIA 35 U.S.C. 102 and 103 (or as subject to pre-AIA 35 U.S.C. 102 and 103) is incorrect, any correction of the statutory basis (i.e., changing from AIA to pre-AIA) for the rejection will not be considered a new ground of rejection if the prior art relied upon, and the rationale supporting the rejection, would be the same under either status.

The following is a quotation of pre-AIA 35 U.S.C. 103(a) which forms the basis for all obviousness rejections set forth in this Office action:

(a) A patent may not be obtained though the invention is not identically disclosed or described as set forth in section 102, if the differences between the subject matter sought to be patented and the prior art are such that the subject matter as a whole would have been obvious at the time the invention was made to a person having ordinary skill in the art to which said subject matter pertains. Patentability shall not be negated by the manner in which the invention was made.

Claims 1, and 11-12 is/are rejected under pre-AIA 35 U.S.C. 103(a) as being unpatentable over US 2004/0037120 A1, filed 08/23/2002, published 02/26/2004 (herein **Uysal**) in view of US 2003/0018657 A1, filed 07/18/2001, published 01/23/2003 (herein **Monday**) in further view of US 7,024,595 B2, filed 08/29/2003, patented 04/04/2006 (herein **Fujimoto**).

Claim 1	Uysal, Monday, and Fujimoto
A disk array system comprising:	<i>Uysal discloses a disk array system (Uysal: paragraph 0002, "Disk arrays are used to provide storage for computer applications that need increased reliability in the face of component failures, as well as high performance for normal use ..."; and figure 1, element 140).</i>
a plurality of disk drives constituting a disk array;	<i>Uysal discloses a plurality of disk drives constituting a disk array (Uysal: figure 1, element 140; paragraph 0002, "... [t]he disks in the disk arrays are often arranged as a redundant array of independent disks (RAID) ..."; and paragraph 0017, "FIG. 1 illustrates a computer system 100, according to an embodiment of the invention. The computer system 100 includes a storage device 110 connected to at least one client 120 (e.g., a server) via a network 130. The storage device 10 includes multiple magnetic disks 140 (which may include an array) and multiple fast storage devices 150 connected to a controller 160, which manages data operations for the disks 140 and the fast storage devices 150").</i>

	<p>FIG. 1 (Uysal: figure 1)</p>
<p>a RAID module for performing a normal read/write access to the disk array;</p>	<p><i>Uysal discloses a RAID module (Uysal: figure 1, element 160; paragraph 0017, “a controller 160, which manages data operations for the disks 140 and the fast storage devices 150”; paragraph 0022, “[i]n one embodiment, the disks 140 and the fast storage devices 150 are configured as one or more RAIDs ... [t]he disks 140 and the fast storage devices 150 may be configured according to one of a variety of RAID layouts”; controller 160 managing a RAID is thus a RAID module) for performing a normal read/write access to the disk array (Uysal: paragraph 0031, “[t]he controller 160 may handle requests from the clients 120 and perform data operations for writing and reading data from the disks 140 and the fast storage devices 150”; figures 3-4; showing read/write data operations).</i></p>
<p>a backup storage device;</p>	<p><i>Uysal discloses a backup storage device (Uysal: figure 1, element 150; paragraph 0022, “[r]edundant data may be stored in the fast storage devices 150, which enables failure recovery and may improve performance”, showing the fast</i></p>

	<p><i>storage devices being used as a backup for at least failure recovery).</i></p>
<p>a backup module for reading data from said disk array and writing the read data onto said backup storage device;</p>	<p><i>Monday</i> demonstrates that it was known at the time of invention for a backup module (<i>Monday: figure 2, element 32 in backup device 10; see also paragraphs 0027-0028</i>) to read data (<i>Monday: paragraph 0034, "archive backup data"</i>) from a disk array (<i>Monday: figure 2, elements 36</i>) and write that data onto backup storage (<i>Monday: paragraph 0034, "[a]s data disks 36 become populated with data, backup module 32 may archive backup data from data disks 36 to a tape-based storage device [42]", context suggests tape storage device 42 instead of logical volume 40</i>).</p> <p><i>Uysal</i> does not explicitly state "a backup module for reading data from said disk array and writing the read data onto said backup storage device". <i>Monday</i> demonstrates that it was known at the time of invention for a backup module to read data from a disk array and write that data onto backup storage (<i>Monday: as discussed above, paragraphs 0027-0028, 0034 and figure 2</i>). It would have been obvious to one of ordinary skill in the art at the time of invention to implement the controller of <i>Uysal</i> with a backup module for reading data from said disk array and writing the read data onto the backup storage device of <i>Uysal</i> as suggested by the teachings <i>Monday</i>. This implementation would have been obvious because one of ordinary skill in the art would have found: the implementation of is an application/addition of one known element and technique with another yielding a predictable result using an acceptable piece of prior art; both references are directed to data and disk operations using RAID (<i>Uysal: paragraph 0022; Monday: paragraphs</i></p>

	<p>0027 and 0034); and <i>Monday</i> indicates the desirability of as needed, periodic backup operations (<i>Monday: paragraph 0028, "Archiving is useful, for example, for performing full backup operations on a weekly or monthly basis, or as needed. Long-term backup device 20 offers the advantage of using removable media, which can be stored off-site for use in disaster scenarios. Local disk-based storage for incremental backup operations substantially shrinks the full backup window and reduces operator intervention for backup to long-term backup device 20 or other secondary media"</i>), which would make use of reading the disk array and writing to the backup storage.</p>
<p>a replacement disk drive for replacing one of the plurality of disk drives after one of the plurality of disk drives fails;</p>	<p><i>Uysal</i> discloses a replacement disk for replacing a failed disk of the plurality of disks (<i>Uysal: figure 5; paragraph 0046, "[i]n step 505, the controller 160 determines whether a disk of the disks 140 failed"; paragraph 0049, "[i]n step 535, the failed device is replaced ... [i]n step 540, the current version of the data is copied back to the replacement device"</i>).</p>
<p>a rebuild module for performing rebuilding of data on the replacement disk drive based on the data backed up in said backup storage device; and</p>	<p><i>Uysal</i> shows a rebuild module (<i>Uysal: figure 5 and paragraphs 0046-0049, controller 160 managing disk failure, replacement and reconstruction</i>) for performing rebuilding of data on the replacement disk drive based on the data backed up in said backup storage device (<i>Uysal: paragraph 0037, "When a failed disk is replaced, the current version of the data is copied back to the replaced disk. <u>The data can be read from either the copy in the fast storage devices 150</u> or, if it has been reconstructed to the spares on the surviving disks, from there", emphasis added</i>). As discussed above, it would have been obvious to use the backup module of</p>

	<p><i>Monday</i> (and corresponding backup data, which is stored in the fast storage devices 150 of <i>Uysal</i>).</p>
<p>means for rebuilding data on the replacement disk drive from the backup storage device through the rebuild module while simultaneously providing normal read/write access to the disk drives that have not failed through the RAID module.</p>	<p>On its face, <i>Uysal</i> shows the broadest reasonable interpretation of rebuilding data from backup while simultaneously providing normal read/write access to disks that haven't failed (<i>Uysal</i>: paragraphs 0036, 0038, and 0050, discussing reconstruction operations occurring in the background while giving priority to external I/O requests, allowing for reading and writing to be asynchronous, or in other words reconstruction does not block normal reading/writing). <i>Uysal</i> shows giving priority to normal read/writing with reconstruction operations happening in the background, which is the broadest reasonable interpretation of the claimed "simultaneous" as they are both in a state of processing for a given period.</p> <p>However, in accordance with 35 USC 112(f), the broadest reasonable interpretation of the means for "rebuilding data on the replacement disk drive from the backup storage device through the rebuild module while simultaneously providing normal read/write access to the disk drives that have not failed through the RAID module" is provided by "disk access bus/switch" (7,739,544: column 3, lines 33-39; column 10, lines 45-47).</p> <p><i>Uysal</i> does not explicitly state "disk access bus/switch". <i>Fujimoto</i> demonstrates that it was known at the time of invention to connect multiple storage devices/disks (<i>Fujimoto</i>: figure 1, elements 300; column 4, lines 38-40) for access via a bus/switch (<i>Fujimoto</i>: figure 1, element 150; column 5, lines 39-41, "[t]he connection portion 150 is a</p>

	<p><i>high-speed bus such as an ultra high-speed crossbar switch for performing data transmission by high-speed switching”), which allows for independent operations on those storage devices/disks (Fujimoto: column 4, lines 1-12). It would have been obvious to one of ordinary skill in the art at the time of invention to implement the system of simultaneous operations of Uysal with connecting the disks of Uysal with a bus/switch for allowing independent access operations as suggested by the teachings Fujimoto. This implementation would have been obvious because one of ordinary skill in the art would have found: the implementation of is an application/substitution of one known element and technique with another yielding a predictable result using an acceptable piece of prior art; both references are directed to RAID storage device/disk operations (Uysal: paragraph 0022; Fujimoto: column 4, lines 45-49).</i></p>
<p>Claim 11</p>	<p>Uysal, Monday, and Fujimoto</p>
<p>The disk array system according to claim 1, further comprising: a disk access bus/switch for providing the RAID module and the rebuild module with parallel access to the disk array and the replacement disk, respectively;</p>	<p>As discussed above for claim 1, <i>Uysal</i> shows the broadest reasonable interpretation of rebuilding data from backup while simultaneously providing normal read/write access to disks that haven’t failed, that is to say parallel access to the various disks, replacements, and backup storages (<i>Uysal: paragraphs 0036, 0038, and 0050, discussing reconstruction operations occurring in the background while giving priority to external I/O requests, allowing for reading and writing to be asynchronous, or in other words reconstruction does not block normal reading/writing</i>). <i>Uysal</i> shows giving priority to normal read/writing with reconstruction operations happening in the background, which is the broadest reasonable interpretation of the claimed “simultaneous” as they are both in a state of processing for a given period.</p>

	<p><i>Uysal</i> does not explicitly state “a disk access bus/switch for providing the RAID module and the rebuild module with parallel access to the disk array and the replacement disk”. <i>Fujimoto</i> demonstrates that it was known at the time of invention to connect multiple storage devices/disks (<i>Fujimoto: figure 1, elements 300; column 4, lines 38-40</i>) for access via a bus/switch (<i>Fujimoto: figure 1, element 150; column 5, lines 39-41, “[t]he connection portion 150 is a high-speed bus such as an ultra high-speed crossbar switch for performing data transmission by high-speed switching”),</i> which allows for independent/parallel operations on those storage devices/disks (<i>Fujimoto: column 4, lines 1-12</i>). It would have been obvious to one of ordinary skill in the art at the time of invention to implement the system of simultaneous operations of <i>Uysal</i> with a disk access bus/switch for providing the RAID module and the rebuild module with parallel access to the disk array and the replacement disk as suggested by the teachings <i>Fujimoto</i>. This implementation would have been obvious because one of ordinary skill in the art would have found: the implementation of is an application/substitution of one known element and technique with another yielding a predictable result using an acceptable piece of prior art; both references are directed to RAID storage device/disk operations (<i>Uysal: paragraph 0022; Fujimoto: column 4, lines 45-49</i>).</p>
<p>wherein the rebuilding of said replacement disk drive from said backup storage is performed through said rebuild module and said disk</p>	<p><i>Uysal</i> and <i>Fujimoto</i> show the broadest reasonable interpretation of the rebuilding of said replacement disk drive from said backup storage is performed through said rebuild module and said disk access bus/switch, thereby the</p>

<p>access bus/switch, thereby the rebuilding of said replacement disk drive being prohibited from influencing a normal access by said RAID module to said disk array.</p>	<p>rebuilding of said replacement disk drive being prohibited from influencing a normal access by said RAID module to said disk array (<i>Uysal</i>: paragraphs 0036, 0038, and 0050, discussing reconstruction operations occurring in the background while giving priority to external I/O requests, allowing for reading and writing to be asynchronous, or in other words reconstruction does not block normal reading/writing). <i>Uysal</i> shows giving priority to normal read/writing with reconstruction operations happening in the background, which is the broadest reasonable interpretation of the claimed the rebuilding of said replacement disk drive being prohibited from influencing a normal access by said RAID module to said disk array. <i>Uysal</i> in combination with <i>Fujimoto</i> show the mechanism allowing the disk drives/storages to be accessed independently (disk access bus/switch of <i>Fujimoto</i>) thus providing the rebuilding of said replacement disk drive being prohibited from influencing a normal access by said RAID module to said disk array.</p>
<p>Claim 12</p>	<p><i>Uysal</i>, <i>Monday</i>, and <i>Fujimoto</i></p>
<p>The disk array system according to claim 1, wherein the rebuild module and the RAID module are separate from each other.</p>	<p><i>Uysal</i> does not state “wherein the rebuild module and the RAID module are separate from each other”. However, <i>Uysal</i> demonstrates that it was known at the time of invention to implement the exemplary embodiments using alternatives, modifications, and variations (<i>Uysal</i>: paragraphs 0016 and 0051) including some method steps simultaneously (<i>Uysal</i>: paragraph 0051). <i>Uysal</i> further shows separate and independent methods, two for RAID module operations (<i>Uysal</i>: figures 3-4), and one for rebuild module operations (<i>Uysal</i>: figure 5). It would have been obvious to one of ordinary skill in the art at the time of invention to implement</p>

	<p>a rebuild module of <i>Uysal</i> as separate from a RAID module of <i>Uysal</i> as suggested by the variation teachings of <i>Uysal</i> and the fact that they are separate and independent operations. This implementation would have been obvious because one of ordinary skill in the art would have found: the implementation of is a variation the disclosed separate methods and the technique is a predictable result; having already disclosed the functionality is separate, the implementation in separate modules is trivial.</p>
--	---

Claims 2-7 is/are rejected under pre-AIA 35 U.S.C. 103(a) as being unpatentable over US 2004/0037120 A1, filed 08/23/2002, published 02/26/2004 (herein ***Uysal***) in view of US 2003/0018657 A1, filed 07/18/2001, published 01/23/2003 (herein ***Monday***) in view of US 7,024,595 B2, filed 08/29/2003, patented 04/04/2006 (herein ***Fujimoto***) in further view of US 6,460,054 B2, filed 12/16/1999, patented 10/01/2002 (herein ***Grummon***).

<i>Claim 2</i>	<i>Uysal, Monday, Fujimoto, and Grummon</i>
<p>The disk array system according to claim 1, further comprising a modify map for recording information corresponding to an update location, when the data in said disk array is updated.</p>	<p><i>Uysal</i> does not explicitly state “further comprising a modify map for recording information corresponding to an update location, when the data in said disk array is updated”. <i>Grummon</i> demonstrates that it was known at the time of invention to use a modified-bit-map table recording information corresponding to an update when updating storage (<i>Grummon: figure 2, element 214; column 7, lines 26-30, 42-45</i>). It would have been obvious to one of ordinary skill in the art at the time of invention to implement the storage system of <i>Uysal</i> with a modify map for recording information corresponding to an update location, when the</p>

	<p>data in said disk array is updated as suggested by the teachings <i>Grummon</i>. This implementation would have been obvious because one of ordinary skill in the art would have found: the implementation of is an application of one known element and technique with another yielding a predictable result using an acceptable piece of prior art; both references are directed to data and disk operations using RAID (<i>Uysal: paragraph 0022; Grummon: column 5, lines 54-56</i>); and desirably a modified-bit-map table can allow for an incremental backup, providing a reduced backup time (<i>Grummon: column 4, lines 21-24</i>).</p>
<p>Claim 3</p>	<p><i>Uysal, Monday, Fujimoto, and Grummon</i></p>
<p>The disk array system according to claim 2, wherein, when the data in said disk array is updated by an access from a host after the backup onto said backup storage device, said backup module sends to said backup storage device the data updated after the backup onto said backup storage device, said modify map records the information corresponding to the updated data; and said data updated after backup is stored as the backup data in said backup storage device.</p>	<p><i>Uysal</i> in view of <i>Monday</i> and <i>Grummon</i>, shows a backup module sending updated data, that arrives to the disk array after a backup, to the backup storage device (<i>Uysal: paragraph 0031, ensuring incoming data is written to both the disk array 140 and the backup storage 150</i>).</p> <p><i>Uysal</i> does not explicitly state “modify map records the information corresponding to the updated data”. <i>Grummon</i> demonstrates that it was known at the time of invention to use a modified-bit-map table recording information corresponding to an update when updating storage (<i>Grummon: figure 2, element 214; column 7, lines 26-30, 42-45</i>). It would have been obvious to one of ordinary skill in the art at the time of invention to implement the storage system of <i>Uysal</i> with a modify map recording the information corresponding to the updated data as suggested by the teachings <i>Grummon</i>. This implementation would have been obvious because one of ordinary skill in the art would have found: the implementation of is an application</p>

	<p>of one known element and technique with another yielding a predictable result using an acceptable piece of prior art; both references are directed to data and disk operations using RAID (<i>Uysal: paragraph 0022; Grummon: column 5, lines 54-56</i>); and desirably a modified-bit-map table can allow for an incremental backup, providing a reduced backup time (<i>Grummon: column 4, lines 21-24</i>).</p>
<p>Claim 4</p>	<p><i>Uysal, Monday, Fujimoto, and Grummon</i></p>
<p>The disk array system according to claim 3, wherein said rebuild module reads from said backup storage device data stored in said failed disk drive and backed up in said backup storage device, writes the data to said replacement disk drive; and</p>	<p><i>Uysal</i> shows a rebuild module reading from the backup to write to the replacement drive of a failed drive (<i>Uysal: figure 5; paragraphs 0046-0049; and paragraph 0037, “[w]hen a failed disk is replaced, the current version of the data is copied back to the replaced disk. The data can be read from either the copy in the fast storage devices 150 or, if it has been reconstructed to the spares on the surviving disks, from there”</i>).</p>
<p>after the completion of writing of the data in said failed disk drive to said replacement disk drive from said backup storage device, said rebuild module refers to said modify map, and generates data for said failed disk drive which have been updated after the backup, based on the data on other disk drives constituting said disk array, to write the generated data to said replacement disk drive.</p>	<p><i>Uysal</i> as discussed above, shows a backup module sending updated data, that arrives to the disk array after a backup, to the backup storage device (<i>Uysal: paragraph 0031, ensuring incoming data is written to both the disk array 140 and the backup storage 150</i>).</p> <p><i>Uysal</i> shows generating data from other disks for a failed disk to write to a replacement disk (<i>Uysal: paragraph 0037, “[w]hen a failed disk is replaced, the current version of the data is copied back to the replaced disk. The data can be read from either the copy in the fast storage devices 150 or, if it has been reconstructed to the spares on the surviving disks, from there”, emphasis added</i>).</p>

	<p>As discussed above, <i>Uysal</i> in view of <i>Grummon</i> shows it would have been obvious to implement <i>Uysal</i> with a modified-bit-map table of <i>Grummon</i>. For the same reasons, it would have been obvious to one of ordinary skill in the art at the time of invention to implement the storage system of <i>Uysal</i> with a modify map used to generate data not yet updated for a replacement drive (that is to determine if a replacement drive's data is current).</p>
<p><i>Claim 5</i></p>	<p><i>Uysal, Monday, Fujimoto, and Grummon</i></p>
<p>The disk array system according to claim 2, wherein said modify map has an update flag corresponding to a unit region for access in each of said disk drives in said disk array, the update flag indicating whether the unit region has been updated or not after the backup.</p>	<p><i>Grummon</i> shows a modify map having an update flag (<i>Grummon: block references 232</i>) corresponding to regions in each disk (<i>Grummon: column 8, lines 8-17, blocks</i>), the flag indicating whether the unit region has been updated or not after the backup (<i>Grummon: column 6, lines 56-63, "Including Read-Write Snapshot Half In performing the copy-on-write procedure, the 1/0 subsystem 112, which reads from and writes to an original read-write container 210, now creates a snapshot container 206, a read-write snapshot container 208 and a backing store or 'backup' container, that is actually termed a snapshot information container for reasons to be described further below 212"; column 7, lines 26-30, "If the request is a storage request, the system checks the modified-bit-map table 214 to determine if the read-write container's block of data was modified after the read-write snapshot container 208 was created. If the block has been modified, the modified bit is set. Therefore, the snapshotted"; column 7, lines 42-45, "During execution, backup processes 204 forward 1/0 requests for files to the snapshot container 208. The read write snapshot container 208 determines whether the file has been modified by</i></p>

	<i>checking the modified-bit-map table 214 for the block where the file is stored. If the block has been”).</i>
Claim 6	Uysal, Monday, Fujimoto, and Grummon
The disk array system according to claim 5, wherein said modify map has a rebuild flag corresponding to the unit region for access in said each of said disk drives in said disk array, the rebuild flag indicating whether rebuilding of the unit region has been performed or not.	<i>Uysal in combination with Grummon shows the rebuild flag indicating whether rebuilding of the unit region has been performed or not (Uysal: paragraph 0037, “When a failed disk is replaced, the <u>current version of the data</u> is copied back to the replaced disk”, emphasis added; paragraph 0038, “A (potentially fault-tolerant) bitmap or watermark pointers can be used to keep track of the progress of the reconstruction”).</i>
Claim 7	Uysal, Monday, Fujimoto, and Grummon
The disk array system according to claim 6, wherein said rebuild module refers to said modify map after the completion of writing of the data in said failed disk drive to said replacement disk drive from said backup storage device, and generates data on the unit region with the update flag indicating updating and with the rebuild flag indicating that the rebuilding has not been performed yet, based on the data in other disk drives constituting said disk array, to write the generated data to said replacement disk drive.	<i>Uysal and Grummon show wherein said rebuild module refers to said modify map after the completion of writing of the data in said failed disk drive to said replacement disk drive from said backup storage device (as above for claim 4), and generates data on the unit region with the update flag indicating updating (as above for claim 5) and with the rebuild flag indicating that the rebuilding has not been performed yet (as above for claim 6), based on the data in other disk drives constituting said disk array (Uysal: paragraph 0037, “[w]hen a failed disk is replaced, the <u>current version of the data</u> is copied back to the replaced disk. The data can be read from either the copy in the fast storage devices 150 or, <u>if it has been reconstructed to the spares on the surviving disks, from there</u>”, emphasis added), to write the generated data to said replacement disk drive (as above for claim 4).</i>

Claim 8 is/are rejected under pre-AIA 35 U.S.C. 103(a) as being unpatentable over US 2004/0037120 A1, filed 08/23/2002, published 02/26/2004 (herein **Uysal**) in view of US 2003/0018657 A1, filed 07/18/2001, published 01/23/2003 (herein **Monday**) in view of US 7,024,595 B2, filed 08/29/2003, patented 04/04/2006 (herein **Fujimoto**) in view of US 6,460,054 B2, filed 12/16/1999, patented 10/01/2002 (herein **Grummon**) in further view of US 2006/0041782 A1, filed 10/20/2004, published 02/23/2006 (herein **Ali**).

Claim 8	Uysal, Monday, Fujimoto, Grummon, and Ali
<p>The disk array system according to claim 1, wherein said rebuild module refers to a modify map, and performs control so that writing back from said backup storage device to said replacement disk drive is not performed on a unit region of the failed disk drive updated after the backup.</p>	<p><i>Uysal</i> does not state, “wherein said rebuild module refers to a modify map, and performs control so that writing back from said backup storage device to said replacement disk drive is not performed on a unit region of the failed disk drive updated after the backup”.</p> <p><i>Grummon</i> shows using a modified-bit-map table to determine if data has been modified (<i>Grummon: column 7, lines 26-30, “If the request is a storage request, the system checks the modified-bit-map table 214 to determine if the read-write container's block of data was modified after the read-write snapshot container 208 was created. If the block has been modified, the modified bit is set. Therefore, the snapshotted”, emphasis added</i>).</p> <p><i>Ali</i> demonstrates that it was known at the time of invention for a rebuild module to refer to a modify map, and perform control so that writing back from a backup storage device to a replacement disk drive is not performed on a unit region of the failed disk drive updated after the backup (<i>Ali: paragraph 0018, “During step 30, a write command from a</i></p>

server may be directed at the replacement disk. The write command is executed on the replacement disk even though step 30 may be incomplete at the time that the command is executed. If a write command is executed with respect to the replacement disk before the completion of step 30, the replacement disk will include current data that should not be overwritten by data included on the spare storage array. In the event that such a write occurs to the replacement disk before the completion of step 30, the corresponding memory locations in the spare storage array are invalidated and are not copied back to the replacement disk during step 30. Identifying the invalidated memory locations in the spare storage array can be accomplished through a look-up table or a bitmap that identifies the range of invalidated memory locations”, emphasis added).

It would have been obvious to one of ordinary skill in the art at the time of invention to implement the storage system of *Uysal* with a modify map (as discussed above for claim 2) such that the rebuild module does not perform write back from backup storage device to the replacement disk drive on a unit region of the failed disk drive updated after the backup as suggested by the teachings *Ali*. This implementation would have been obvious because one of ordinary skill in the art would have found: the implementation of is an application of one known element and technique with another yielding a predictable result using an acceptable piece of prior art; the references are directed to data and disk operations using RAID (*Uysal*: paragraph 0022; *Grummon*: column 5, lines 54-56; *Ali*: paragraph 0006); the desirability of ensuring current data is not overwritten by invalid data from the backup (*Ali*:

	<p><i>paragraph 0018</i>); and the desirability of improving rebuild times that reduce the amount of time that the storage array exists in a non-redundant state (<i>Ali: paragraph 0019</i>).</p>
--	--

Claims 9-10 is/are rejected under pre-AIA 35 U.S.C. 103(a) as being unpatentable over US 2004/0037120 A1, filed 08/23/2002, published 02/26/2004 (herein **Uysal**) in view of US 2003/0018657 A1, filed 07/18/2001, published 01/23/2003 (herein **Monday**) in view of US 7,024,595 B2, filed 08/29/2003, patented 04/04/2006 (herein **Fujimoto**) in further view of US 5,390,187, filed 02/16/1993, patented 02/14/1995 (herein **Stallmo**).

Claim 9	Uysal, Monday, Fujimoto, and Stallmo
<p>The disk array system according to claim 1, wherein an update to the failed disk during rebuilding of the replacement disk is written to the replacement disk drive without alteration.</p>	<p><i>Uysal</i> does not explicitly state, “an update to the failed disk during rebuilding of the replacement disk is written to the replacement disk drive without alteration”.</p> <p><i>Stallmo</i> demonstrates that it was known at the time of invention for an update (<i>Stallmo: figure 5, element 50</i>) to the failed disk during rebuilding of the replacement disk (<i>Stallmo: figure 5, element 52; column 7, lines 9-16</i>) is written to the replacement disk drive without alteration (<i>Stallmo: figure 5, element 55; column 9, lines 28-30, 57-68, “If the Write operation requested by the CPU 1 is for a data block where either the data block or its corresponding parity bock are on the replacement storage unit S1' (step 52), the procedure is different ... [t]he new data block overwrites the old data block and the new parity block overwrites the old parity block (Step 55)”, emphasis added</i>).</p>


	<p>It would have been obvious to one of ordinary skill in the art at the time of invention to implement the storage system of <i>Uysal</i> such that an update to the failed disk during rebuilding of the replacement disk is written to the replacement disk drive without alteration as suggested by the teachings <i>Stallmo</i>. This implementation would have been obvious because one of ordinary skill in the art would have found: the implementation of is an application of one known element and technique with another yielding a predictable result using an acceptable piece of prior art; the references are directed to data and disk operations using RAID (<i>Uysal: paragraph 0022; Stallmo: column 6, lines 33-34</i>); the desirability of a RAID system in which reconstruction of a failed storage unit can be conducted with “on-line” access to the entire redundancy group, while general system operation continues in a normal fashion (<i>Stallmo: column 4, line 63 to column 5, line 2</i>); and the desirability of “the data storage subsystem to always maintain and present correct data without having to keep track of whether or not a requested data block is ‘above’ or ‘below’ the reconstruction stripe” (<i>Stallmo: column 10, lines 31-35</i>).</p>
<p>Claim 10</p>	<p><i>Uysal, Monday, Fujimoto, and Stallmo</i></p>
<p>The disk array system according to claim 1, wherein, even before the rebuilding is finished, when an access request is made to latest data among the data in said replacement disk drive under the rebuilding, access to said replacement disk drive is permitted.</p>	<p><i>Stallmo</i> shows wherein, even before the rebuilding is finished (<i>Stallmo: column 18, lines 18-23, “If a Read operation is requested by the CPU 1 during general reconstruction, a Read process begins as an independent but concurrent task (FIG. 4, Step 40). This means that the reconstruction task continues to execute, but that the Read task and reconstruction task contend for system resources, in known fashion”</i>), when an access request is made to latest data among the data in said replacement disk drive under</p>

	<p>the rebuilding (<i>Stallmo: column 9, lines 18-22, "if the requested block is 'above' the reconstruction line, the block has already been reconstructed. Therefore, the controller 3 performs a normal Read of the data block and transfers the data block to the CPU 1 (Step 43)"</i>), access to said replacement disk drive is permitted (<i>Stallmo: column 9, line 57 to column 10, line 3; column 10, lines 22-26</i>).</p>
--	--

Claims 13 and 22 is/are rejected under pre-AIA 35 U.S.C. 103(a) as being unpatentable over US 2004/0037120 A1, filed 08/23/2002, published 02/26/2004 (herein **Uysal**) in view of US 7,024,595 B2, filed 08/29/2003, patented 04/04/2006 (herein **Fujimoto**).

Claim 13	Uysal, and Fujimoto
<p>A method of rebuilding a disk array system, said method comprising:</p>	<p><i>Uysal</i> discloses a disk array system (<i>Uysal: paragraph 0002, "Disk arrays are used to provide storage for computer applications that need increased reliability in the face of component failures, as well as high performance for normal use ..."; and figure 1, element 140</i>).</p> <p><i>Uysal</i> discloses a plurality of disk drives constituting a disk array (<i>Uysal: figure 1, element 140; paragraph 0002, "... [t]he disks in the disk arrays are often arranged as a redundant array of independent disks (RAID) ..."; and paragraph 0017, "FIG. 1 illustrates a computer system 100, according to an embodiment of the invention. The computer system 100 includes a storage device 110 connected to at least one client 120 (e.g., a server) via a network 130. The storage device 10 includes multiple magnetic disks 140 (which may include an array) and multiple fast storage</i></p>

devices 150 connected to a controller 160, which manages data operations for the disks 140 and the fast storage devices 150”).



Uysal discloses a RAID module (Uysal: figure 1, element 160; paragraph 0017, “a controller 160, which manages data operations for the disks 140 and the fast storage devices 150”; paragraph 0022, “[i]n one embodiment, the disks 140 and the fast storage devices 150 are configured as one or more RAIDs ... [t]he disks 140 and the fast storage devices 150 may be configured according to one of a variety of RAID layouts”; controller 160 managing a RAID is thus a RAID module) for performing a normal read/write access to the disk array (Uysal: paragraph 0031, “[t]he controller 160 may handle requests from the clients 120 and perform data operations for writing and reading data from the disks 140 and the fast storage devices 150”; figures 3-4; showing read/write data operations).

	<p><i>Uysal shows a rebuild module (Uysal: figure 5 and paragraphs 0046-0049, controller 160 managing disk failure, replacement and reconstruction) for performing rebuilding of data on the replacement disk drive based on the data backed up in said backup storage device (Uysal: paragraph 0037, “When a failed disk is replaced, the current version of the data is copied back to the replaced disk. <u>The data can be read from either the copy in the fast storage devices 150 or, if it has been reconstructed to the spares on the surviving disks, from there</u>”, emphasis added).</i></p>
<p>backing up data in a disk array including a plurality of disk drives onto a backup storage device; and</p>	<p><i>Uysal discloses backing up data from the disks to a backup storage device (Uysal: figure 1, element 150; paragraph 0022, “[r]edundant data may be stored in the fast storage devices 150, which enables failure recovery and may improve performance”, showing the fast storage devices being used as a backup for at least failure recovery).</i></p>
<p>when a failed disk drive among said disk drives constituting said disk array is replaced with a replacement disk drive, rebuilding data in said replacement disk drive from the backed-up data in the backup storage device while simultaneously providing</p>	<p><i>Uysal discloses a replacement disk for replacing a failed disk of the plurality of disks (Uysal: figure 5; paragraph 0046, “[i]n step 505, the controller 160 determines whether a disk of the disks 140 failed”; paragraph 0049, “[i]n step 535, the failed device is replaced ... [i]n step 540, the current version of the data is copied back to the replacement device”).</i></p>

<p>other devices with access to the disk drives that have not failed.</p>	<p>On its face, <i>Uysal</i> shows the broadest reasonable interpretation of rebuilding data from backup while simultaneously providing normal read/write access to disks that haven't failed (<i>Uysal</i>: paragraphs 0036, 0038, and 0050, discussing reconstruction operations occurring in the background while giving priority to external I/O requests, allowing for reading and writing to be asynchronous, or in other words reconstruction does not block normal reading/writing). <i>Uysal</i> shows giving priority to normal read/writing with reconstruction operations happening in the background, which is the broadest reasonable interpretation of the claimed "simultaneous" as they are both in a state of processing for a given period.</p> <p>However, in accordance with 35 USC 112(f), the broadest reasonable interpretation of the means for "rebuilding data on the replacement disk drive from the backup storage device through the rebuild module while simultaneously providing normal read/write access to the disk drives that have not failed through the RAID module" is provided by "disk access bus/switch" (7,739,544: column 3, lines 33-39; column 10, lines 45-47).</p> <p><i>Uysal</i> does not explicitly state "disk access bus/switch". <i>Fujimoto</i> demonstrates that it was known at the time of invention to connect multiple storage devices/disks (<i>Fujimoto</i>: figure 1, elements 300; column 4, lines 38-40) for access via a bus/switch (<i>Fujimoto</i>: figure 1, element 150; column 5, lines 39-41, "[t]he connection portion 150 is a high-speed bus such as an ultra high-speed crossbar switch for performing data transmission by high-speed switching"), which allows for independent operations on those storage</p>
---	--

	<p>devices/disks (<i>Fujimoto: column 4, lines 1-12</i>). It would have been obvious to one of ordinary skill in the art at the time of invention to implement the system of simultaneous operations of <i>Uysal</i> with connecting the disks of <i>Uysal</i> with a bus/switch for allowing independent access operations as suggested by the teachings <i>Fujimoto</i>. This implementation would have been obvious because one of ordinary skill in the art would have found: the implementation of is an application/substitution of one known element and technique with another yielding a predictable result using an acceptable piece of prior art; both references are directed to RAID storage device/disk operations (<i>Uysal: paragraph 0022; Fujimoto: column 4, lines 45-49</i>).</p>
<p>Claim 22</p>	<p><i>Uysal</i>, and <i>Fujimoto</i></p>
<p>The method according to claim 13, further comprising: providing a host with normal read/write access to the disk array in parallel with the access to the replacement disk drive for the rebuilding of data in the replacement disk.</p>	<p>As discussed above for claim 1, <i>Uysal</i> shows the broadest reasonable interpretation of rebuilding data from backup while simultaneously providing normal read/write access to disks that haven't failed, that is to say parallel access to the various disks, replacements, and backup storages (<i>Uysal: paragraphs 0036, 0038, and 0050, discussing reconstruction operations occurring in the background while giving priority to external I/O requests, allowing for reading and writing to be asynchronous, or in other words reconstruction does not block normal reading/writing</i>). <i>Uysal</i> shows giving priority to normal read/writing with reconstruction operations happening in the background, which is the broadest reasonable interpretation of the claimed "simultaneous" as they are both in a state of processing for a given period.</p> <p><i>Uysal</i> does not explicitly state "a disk access bus/switch for providing the RAID module and the rebuild module with</p>

	<p>parallel access to the disk array and the replacement disk".</p> <p><i>Fujimoto</i> demonstrates that it was known at the time of invention to connect multiple storage devices/disks (<i>Fujimoto: figure 1, elements 300; column 4, lines 38-40</i>) for access via a bus/switch (<i>Fujimoto: figure 1, element 150; column 5, lines 39-41, "[t]he connection portion 150 is a high-speed bus such as an ultra high-speed crossbar switch for performing data transmission by high-speed switching"</i>), which allows for independent/parallel operations on those storage devices/disks (<i>Fujimoto: column 4, lines 1-12</i>). It would have been obvious to one of ordinary skill in the art at the time of invention to implement the system of simultaneous operations of <i>Uysal</i> with a disk access bus/switch for providing the RAID module and the rebuild module with parallel access to the disk array and the replacement disk as suggested by the teachings <i>Fujimoto</i>. This implementation would have been obvious because one of ordinary skill in the art would have found: the implementation of is an application/substitution of one known element and technique with another yielding a predictable result using an acceptable piece of prior art; both references are directed to RAID storage device/disk operations (<i>Uysal: paragraph 0022; Fujimoto: column 4, lines 45-49</i>).</p>
--	--

Claims 14-18 is/are rejected under pre-AIA 35 U.S.C. 103(a) as being unpatentable over US 2004/0037120 A1, filed 08/23/2002, published 02/26/2004 (herein **Uysal**) in view of US 7,024,595 B2, filed 08/29/2003, patented 04/04/2006 (herein **Fujimoto**) in further view of US 6,460,054 B2, filed 12/16/1999, patented 10/01/2002 (herein **Grummon**).

<i>Claims 14</i>	<i>Uysal, Fujimoto, and Grummon</i>
<p>The method according to claim 13, further comprising: recording information corresponding to an update location, when the data in said disk array is updated by an access from a host after the backup onto said backup storage device; sending only the updated data to said backup storage device after the backup onto said backup storage device; and storing the updated data in said backup storage device as the backup data.</p>	<p>The limitations of claim 14 correspond to the limitations of claims 2-3. As such, the limitations of claim 14 are rejected in a corresponding manner as claims 2-3 in view of <i>Grummon</i>.</p>
<i>Claim 15</i>	<i>Uysal, Fujimoto, and Grummon</i>
<p>The method according to claim 14, further comprising: reading from said backup storage device data stored in said failed disk drive and backed up in said backup storage device, and writing the data to said replacement disk drive; and referring to a modify map, after writing of the data in said failed disk drive to said replacement disk drive from said backup storage device is finished, generating data on said failed disk drive updated after the backup, based on the data on other disk drives constituting said disk array</p>	<p>The limitations of claim 15 correspond to the limitations of claim 4. As such, the limitations of claim 15 are rejected in a corresponding manner as claim 4 in view of <i>Grummon</i>.</p>

and writing the generated data to said replacement disk drive.	
<i>Claim 16</i>	<i>Uysal, Fujimoto, and Grummon</i>
The method according to claim 14, wherein a modify map has an update flag corresponding to a unit region for access in each of said disk drives in said disk array, the update flag indicating whether the unit region has been updated or not after the backup.	The limitations of claim 16 correspond to the limitations of claim 5. As such, the limitations of claim 16 are rejected in a corresponding manner as claim 5 in view of <i>Grummon</i> .
<i>Claim 17</i>	<i>Uysal, Fujimoto, and Grummon</i>
The method according to claim 16, wherein said modify map has a rebuild flag corresponding to the unit region for access in said each of said disk drives in said disk array, the rebuild flag indicating whether rebuilding of the unit region has been performed or not.	The limitations of claim 17 correspond to the limitations of claim 6. As such, the limitations of claim 17 are rejected in a corresponding manner as claim 6 in view of <i>Grummon</i> .
<i>Claim 18</i>	<i>Uysal, Fujimoto, and Grummon</i>
The method according to claim 17, further comprising: referring to said modify map, after writing of the data in said failed disk drive to said replacement disk drive from said backup storage device is finished; generating data on the unit region with the update flag indicating	The limitations of claim 18 correspond to the limitations of claim 7. As such, the limitations of claim 18 are rejected in a corresponding manner as claim 7 in view of <i>Grummon</i> .

<p>updating and with the rebuild flag indicating that the rebuilding has not been performed yet, based on the data in other disk drives constituting said disk array; and writing the generated data to said replacement disk drive.</p>	
--	--

Claims 19 is/are rejected under pre-AIA 35 U.S.C. 103(a) as being unpatentable over US 2004/0037120 A1, filed 08/23/2002, published 02/26/2004 (herein **Uysal**) in view of US 7,024,595 B2, filed 08/29/2003, patented 04/04/2006 (herein **Fujimoto**) in view of US 6,460,054 B2, filed 12/16/1999, patented 10/01/2002 (herein **Grummon**) in further view of US 2006/0041782 A1, filed 10/20/2004, published 02/23/2006 (herein **Ali**).

Claim 19	Uysal, Fujimoto, Grummon, and Ali
<p>The method according to claim 13, the rebuilding data in said replacement disk further comprising referring to a modify map and performing control so that writing back from said backup storage device is not performed on a unit region of the failed disk drive updated after the backup.</p>	<p>The limitations of claim 19 correspond to the limitations of claim 8. As such, the limitations of claim 19 are rejected in a corresponding manner as claim 8 in view of <i>Ali</i>.</p>

Claims 20-21 is/are rejected under pre-AIA 35 U.S.C. 103(a) as being unpatentable over US 2004/0037120 A1, filed 08/23/2002, published 02/26/2004 (herein **Uysal**) in view of US 7,024,595 B2,

Art Unit: 3992

filed 08/29/2003, patented 04/04/2006 (herein **Fujimoto**) in further view of US 5,390,187, filed

02/16/1993, patented 02/14/1995 (herein **Stallmo**).

Claim 20	Uysal, Fujimoto, and Stallmo
<p>The method according to claim 13, further comprising performing control so that a content updated during the rebuilding of said replacement disk drive is written to said replacement disk drive without using said backup storage device.</p>	<p><i>Uysal</i> does not explicitly state, “further comprising performing control so that a content updated during the rebuilding of said replacement disk drive is written to said replacement disk drive without using said backup storage device”.</p> <p><i>Stallmo</i> demonstrates that it was known at the time of invention for an update (<i>Stallmo: figure 5, element 50</i>) to the failed disk during rebuilding of the replacement disk (<i>Stallmo: figure 5, element 52; column 7, lines 9-16</i>) is written to said replacement disk drive without using said backup storage device (<i>Stallmo: figure 5, element 55; column 9, lines 28-30, 57-68, “If the Write operation requested by the CPU 1 is for a data block where either the data block or its corresponding parity bock are on the replacement storage unit S1' (step 52), the procedure is different ... [t]he new data block overwrites the old data block and the new parity block overwrites the old parity block (Step 55)”</i>, emphasis added, as opposed to column 9, lines 40-49).</p> <p>It would have been obvious to one of ordinary skill in the art at the time of invention to implement the storage system of <i>Uysal</i> such that an update to the failed disk during rebuilding of the replacement disk is written to said replacement disk drive without using said backup storage device as suggested by the teachings <i>Stallmo</i>. This implementation would have</p>

	<p>been obvious because one of ordinary skill in the art would have found: the implementation of is an application of one known element and technique with another yielding a predictable result using an acceptable piece of prior art; the references are directed to data and disk operations using RAID (<i>Uysal: paragraph 0022; Stallmo: column 6, lines 33-34</i>); the desirability of a RAID system in which reconstruction of a failed storage unit can be conducted with “on-line” access to the entire redundancy group, while general system operation continues in a normal fashion (<i>Stallmo: column 4, line 63 to column 5, line 2</i>); and the desirability of “the data storage subsystem to always maintain and present correct data without having to keep track of whether or not a requested data block is ‘above’ or ‘below’ the reconstruction stripe” (<i>Stallmo: column 10, lines 31-35</i>).</p>
<p><i>Claim 21</i></p>	<p><i>Uysal, Fujimoto, and Stallmo</i></p>
<p>The method according to claim 13, further comprising performing control so that even before the rebuilding is finished, when an access request is made to latest data among the data in said replacement disk drive under rebuilding, access to said replacement disk drive is permitted.</p>	<p>The limitations of claim 21 correspond to the limitations of claim 10. As such, the limitations of claim 21 are rejected in a corresponding manner as claim 10 in view of <i>Stallmo</i>.</p>

Declaration Exhibit

The Declaration of Paul S. Min, PH.D. submitted on 07/11/2025 is considered by the examiner.

See the following Response to Arguments for any comments regarding the Declaration.

Response to Arguments

Patent Owner's arguments filed 07/11/2025 (herein *Remarks*) have been fully considered but they are not persuasive. The *Remarks* argues: (1) *Uysal* does not show simultaneously providing read/write access to disk drives that have not been failed while rebuilding data on the replacement disk occurs (*Remarks: pages 5-16, section A.*); (2) *Fujimoto's* disclosure of "independent NAS" does not show simultaneously providing access to disk drives that have not failed while rebuilding of data occurs (*Remarks: pages 16-20, section B.*); and (3) the implementation of *Fujimoto's* crossbar switch into *Uysal* does not show simultaneously providing read/write access to disk drives that have not failed while rebuilding of data occurs (*Remarks: pages 20-23, section C.*).

Issue 1 (*Remarks* section A.)

These arguments are not persuasive due to the breadth of the claim language in question. Not only is the term "simultaneous" broad, but so is the word "access".

The relevant language of independent claim 1 is "... rebuilding data on the replacement disk drive from the backup storage device through the rebuild module while simultaneously providing normal read/write access to the disk drives that have not failed ...". The relevant language of independent claim 13 is "... rebuilding data in said replacement disk drive from the backed-up data in the backup storage device while simultaneously providing other devices with access to the disk drives that have not failed". Both versions require providing simultaneous access to the disk drives. Neither require simultaneous individual memory operations. Considering the respective phrases of the independent claims as a whole (in reference to the Min Declaration, paragraph 11), reveals the broad claim language reads upon the cited prior art.

Uysal shows the broadest reasonable interpretation of rebuilding data from backup while simultaneously providing normal read/write/device access to disks that haven't failed (*Uysal*:

paragraphs 0036, 0038, and 0050, discussing reconstruction operations occurring in the background while giving priority to external I/O requests, allowing for reading and writing to be asynchronous, or in other words reconstruction does not block normal reading/writing). *Uysal* shows giving priority to normal read/writing with reconstruction operations happening in the background, which is the broadest reasonable interpretation of the claimed “simultaneous” as both the read/write/device access process and the rebuilding data process are in a state of processing for the same given period of time (note, claim 13 simply recites device access instead of read/write access). This is true even if the individual operations of the processes alternate or themselves do not occur at the same time.

The *Remarks* first argue the architecture of *Uysal* does not permit conducting both normal read/write operations and rebuilding operations simultaneously because the way the stripe units are allocated. However, the claim language does not require individual read/write operations and individual rebuilding operations to be performed simultaneously. Instead, the claims call for read/write disk access to be simultaneous to rebuilding data.

As an initial matter, the argument presented here is to whether *Uysal* shows simultaneous individual operations. The *Remarks* assert that in *Uysal*, at a disk in the array, either a normal operation or a rebuilding operation is performed, not both operations simultaneously. The arguments do not address the broader claim language that allows for read/write operations of the larger read/write access process to alternate with rebuilding operations of the larger rebuilding data process.

Further, the argument presented here posits a limited scenario in *Uysal* without addressing the larger situation of *Uysal*. For example, the *Remarks* suggest each individual disclosed disk has only one read/write head and therefore each disk can only perform either a normal/read write operation or a rebuilding read/write operation. The *Remarks* then conclude if a regular disk fails, a fast/backup disk can only be supporting normal read/write or rebuilding the failed disk. This argument does not account for

the larger disclosure of *Uysal* that the regular disks are part of an array of disks and that the fast/backup disk is part of an array of disks. As the arguments point out, the stripes used for backup are spread across multiple fast/backup disks. This very feature allows for normal read/write access and rebuilding to be simultaneous, at least with regard to each disk having one read/write head, since a stripe can be normally read/write from one disk in the fast/backup array while another stripe is being rebuilt by using another disk in the fast/backup array. The arguments do not seem to account for stripes across several disks means multiple disk operations (read/write access or rebuilding data) can occur simultaneously because there are multiple disks (in reference to the Min Declaration, paragraph 22).

For both of these reasons, this argument is not persuasive.

The *Remarks* second argue the act of giving priority to one operation over another operation does not disclose processing two operations simultaneously. However, the claim language does not require specific and individual read/write operations and rebuilding operations to be performed simultaneously. Instead, the claims call for read/write disk access to be simultaneous to rebuilding data. In fact, *Uysal*'s description of "reconstruction operations can occur in the background, giving priority to external I/O requests" is a standard description of simultaneous execution of two processes (processes that may be made of many smaller individual operations or actions). The claim language's generalized recitation of read/write access and rebuilding data merely equates to two simultaneous processes that each may be made of several individual operations/actions. The processes are simultaneous, even if the operations/actions of those processes are interleaved and not performed at the exact same time. This analysis of simultaneous is consistent with the '544 specification's use of "in parallel with". Further, *Uysal*'s giving priority to external I/O requests is consistent with the '544 specification's description of access to the disks that have not failed is non-blocked (in reference to the Min Declaration, paragraph 28). Therefore, this argument is not persuasive.

The *Remarks* third argue the cache of *Uysal* does not provide the ability to rebuild data from a backup storage device while simultaneously providing normal read/write access to disks that have not failed. However, as previously stated, the claim language does not require individual read/write operations and individual rebuilding operations to be performed simultaneously. Instead, the claims call for read/write disk access to be simultaneous to rebuilding data. Since the claim language is directed toward the simultaneous performance of a read/write/device access process and rebuilding data process, whether the individual cache/buffer transfer operations are interrupted is not relevant. Further, in *Uysal*, the presence of caches/buffers holding intermediate data states between disks and data access requests further shows simultaneously processing the read/write access and rebuilding data processes (the individual operations, reads and writes, may not be processed at the same time, but as a whole the read/write access process and the rebuilding data process are simultaneous). Further, “simultaneous” and “parallel” are shown as indicated above. Therefore, this argument is not persuasive.

Issue 2 (Remarks section B.)

These arguments are not persuasive since *Uysal* shows read/write/device access to disks while simultaneously rebuilding data as discussed above under Issue 1 and the rejections. *Fujimoto* shows a “disk access bus/switch” as discussed above in the rejections of claims 1 and 13. In other words, the rejections rely on *Uysal* to show the broadest reasonable interpretation of “simultaneously”. Therefore, this argument is not persuasive.

Issue 3 (Remarks section C.)

These arguments are not persuasive since *Uysal* shows read/write/device access to disks while simultaneously rebuilding data as discussed above under Issue 1 and the rejections. *Fujimoto* shows a

“disk access bus/switch” as discussed above in the rejections of claims 1 and 13. In other words, the rejections rely on *Uysal* to show the broadest reasonable interpretation of “simultaneously”.

The relevant language of independent claim 1 is “... rebuilding data on the replacement disk drive from the backup storage device through the rebuild module while simultaneously providing normal read/write access to the disk drives that have not failed ...”. The relevant language of independent claim 13 is “... rebuilding data in said replacement disk drive from the backed-up data in the backup storage device while simultaneously providing other devices with access to the disk drives that have not failed”.

Further, the *Remarks* argue the bus/switch of *Fujimoto* operate on a fixed, uniform time interval, and as such “the combination would not show or suggest the final limitations of claims 1 and 13” (*Remarks: page 21*). But, the claims merely require access to multiple disk drives. The bus/switch of *Fujimoto* provides a mechanism to connect multiple disks drives.

Additionally, it is not apparent from the arguments why or how the alleged uniform time interval or the alleged disks systems not being of uniform length is applicable to the broad recitation of the claim language. *Fujimoto* is not cited for reassembly/rebuilding of a failed disk. There is no clear reason why switch collision must be the result of the combination. *Uysal* deals with stripe sizes (*Uysal: paragraphs 0008-0009*) and how they relate to the regular disk array and the fast/backup disk array. *Fujimoto* simply provides a bus/switch mechanism to connect devices. Any issues with stripe size and/or collision are addressed by *Uysal* alone, regardless of the introduction of *Fujimoto*.

Finally, the *Remarks* assert the combination would result in one stripe unit used at a time or one operation at a time. This would still read upon the broadest reasonable interpretation of “simultaneous” as discussed above.

Therefore, these arguments are not persuasive.

The dependent claims are not argued separately and as such the above comments are applicable to them as well. Likewise, the arguments of the Min Declaration are addressed in the same manner as the above comments. Having addressed all the outstanding issues, the arguments are not persuasive for all the reasons discussed here.

Conclusion

THIS ACTION IS MADE FINAL.

A shortened statutory period for response to this action is set to expire 2 from the mailing date of this action.

Extensions of time under 37 CFR 1.136(a) do not apply in reexamination proceedings. The provisions of 37 CFR 1.136 apply only to “an applicant” and not to parties in a reexamination proceeding. Further, in 35 U.S.C. 305 and in 37 CFR 1.550(a), it is required that reexamination proceedings “will be conducted with special dispatch within the Office.”

Extensions of time in reexamination proceedings are provided for in 37 CFR 1.550(c). A request for extension of time must specify the requested period of extension and it must be accompanied by the petition fee set forth in 37 CFR 1.17(g)(1). Any request for an extension in a third party requested ex parte reexamination must be filed on or before the day on which action by the patent owner is due, and the mere filing of a request will not effect any extension of time. A request for an extension of time in a third party requested ex parte reexamination will be granted only for sufficient cause, and for a reasonable time specified. Any request for extension in a patent owner requested ex parte reexamination (including reexamination ordered under 35 U.S.C. 257) for up to two months from the time period set in the Office action must be filed no later than two months from the expiration of the time period set in the Office action. A request for an extension in a patent owner requested ex parte reexamination for more than two months from the time period set in the Office action must be filed on

or before the day on which action by the patent owner is due, and the mere filing of a request for an extension for more than two months will not effect the extension. The time for taking action in a patent owner requested *ex parte* reexamination will not be extended for more than two months from the time period set in the Office action in the absence of sufficient cause or for more than a reasonable time.

The filing of a timely first response to this final rejection will be construed as including a request to extend the shortened statutory period for an additional two months. In no event, however, will the statutory period for response expire later than SIX MONTHS from the mailing date of the final action.

See MPEP § 2265.

Correspondence Information

All correspondence relating to this *ex parte* reexamination proceeding may be submitted via:

Electronically: Registered users may submit via Patent Center
<https://patentcenter.uspto.gov/>.

By Mail to: Mail Stop *Ex Parte* Reexam
Central Reexamination Unit
Commissioner for Patents
United States Patent & Trademark Office
P.O. Box 1450
Alexandria, VA 22313-1450

By FAX to: (571) 273-9900
Central Reexamination Unit

By hand: Customer Service Window
Knox Building
501 Dulany Street
Alexandria, VA 22314

For Patent Center transmissions, 37 CFR 1.8(a)(1)(i)(C) and (ii) states that correspondence (except for a request for reexamination and a corrected or replacement request for reexamination) will be considered timely filed if (a) it is transmitted via the Office's electronic filing system in accordance with 37 CFR 1.6(a)(4), and (b) includes a certificate of transmission for each piece of correspondence stating the date of transmission, which is prior to the expiration of the set period of time in the Office action.

Any inquiry concerning this communication or earlier communications from the examiner should be directed to WILLIAM H WOOD whose telephone number is (571)272-3736. The examiner can normally

Art Unit: 3992

be reached Monday-Friday 7am-3pm. If attempts to reach the examiner by telephone are unsuccessful, the examiner's supervisor, Alexander Kosowski can be reached on (571)272-3744. For all general inquiries the CRU main phone number is (571)272-7705.

Examiner interviews are available via telephone, in-person, and video conferencing using a USPTO supplied web-based collaboration tool. To schedule an interview, applicant is encouraged to use the USPTO Automated Interview Request (AIR) at <http://www.uspto.gov/interviewpractice>.

Information regarding the status of published or unpublished applications may be obtained from Patent Center. Unpublished application information in Patent Center is available to registered users. To file and manage patent submissions in Patent Center, visit: <https://patentcenter.uspto.gov>. Visit <https://www.uspto.gov/patents/apply/patent-center> for more information about Patent Center and <https://www.uspto.gov/patents/docx> for information about filing in DOCX format. For additional questions, contact the Electronic Business Center (EBC) at 866-217-9197 (toll-free). If you would like assistance from a USPTO Customer Service Representative, call 800-786-9199 (IN USA OR CANADA) or 571-272-1000.

/William H. Wood/

Reexamination Specialist, Art Unit 3992

Conferees:

/RACHNA S DESAI/

Reexamination Specialist, Art Unit 3992

/ALEXANDER J KOSOWSKI/

Supervisory Patent Examiner, Art Unit 3992