



- (51) **International Patent Classification:**
G06F 17/00 (2006.01) *G06F 17/40* (2006.01)
- (21) **International Application Number:**
PCT/US2013/053239
- (22) **International Filing Date:**
1 August 2013 (01.08.2013)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (30) **Priority Data:**
13/564,622 1 August 2012 (01.08.2012) US
- (71) **Applicant:** NETAPP, INC. [US/US]; 495 East Java Drive, Sunnyvale, California 94089 (US).
- (72) **Inventor:** HORN, Gustav; 495 East Java Drive, Sunnyvale, California 94089 (US).
- (74) **Agents:** SARATHY, Rajiv P. et al.; PERKINS COIE LLP, P.O. Box 1208, Seattle, Washington 98111-1208 (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, 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, 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).

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))

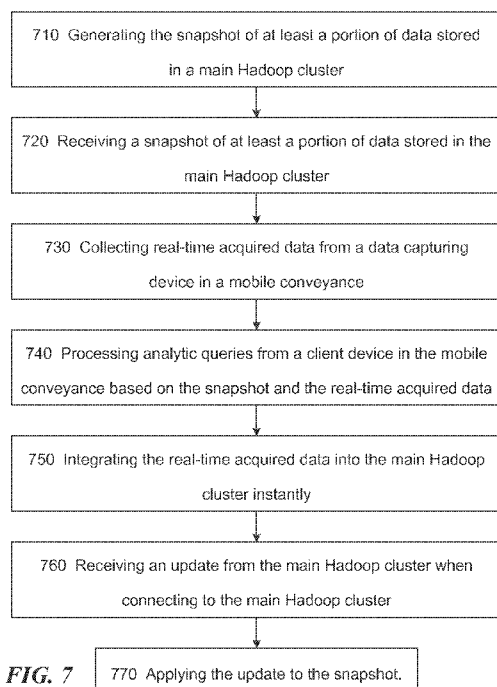
(54) **Title:** MOBILE HADOOP CLUSTERS

FIG. 7

(57) **Abstract:** Techniques for mobile clusters for collecting telemetry data and processing analytic tasks, are disclosed herein. The mobile cluster includes a processor, a plurality of data nodes and an analysis module. The data nodes receive and store a snapshot of at least a portion of data stored in a main Hadoop storage cluster and real-time acquired data received from a data capturing device. The analysis module is operatively coupled to the processor to process analytic tasks based on the snapshot and the real-time acquired data when the storage cluster is not connected to the main storage cluster.

MOBILE HADOOP CLUSTERS

CROSS-REFERENCE TO RELATED APPLICATION

5 **[0001]** This application claims priority to U.S. Patent Application No. 13/564,622 filed 1 August 2012, which is hereby incorporated by reference in its entirety.

FIELD OF THE INVENTION

[0002] At least one embodiment of the present invention pertains to data-intensive clusters, and more particularly, to a mobile Hadoop cluster capable of
10 collecting data and providing autonomous data analysis in real time.

BACKGROUND

[0003] Developments in computer and networking technology have given rise to applications that require massive amounts of data storage. For example, tens of
15 millions of users can create web pages and upload images and text to a social media website. Consequently, a social media website can accumulate massive amounts of data each day and therefore need a highly scalable system for storing and processing data. Various tools exist to facilitate such mass data storage.

[0004] Hadoop is a popular open source framework that supports large-scale
20 data-intensive distributed applications, by enabling applications to interact with a cluster of thousands of computers (also referred to as nodes) and petabytes of data. The Hadoop framework utilizes a distributed, scalable, portable file system, called Hadoop Distributed File System (HDFS), to distribute a massive amount of data among data nodes (also referred to as slave nodes) in a Hadoop cluster. In order to
25 reduce the adverse impact of a data node power outage or network failure (including

switch failure), data in an HDFS is typically replicated on different data nodes. The HDFS can include at least one metadata node to host a file system index.

[0005] However, in order to store the huge amount of data in the Hadoop cluster, the cluster needs a large number of nodes. The number of nodes requires a
5 significant space to maintain the cluster. In addition, all the hardware running within the cluster consumes a significant amount of power. These factors limit the usability of a conventional Hadoop cluster.

[0006] For example, certain applications employing real-time data collection and analysis, such as signal intelligence gathering, may require both mobility and very
10 large data storage capacity. Yet a conventional Hadoop cluster, which has very large storage capacity, is not mobile due to its size and power consumption limitations. Therefore, unless the data collection equipment for such an application has a communication link (wired or wireless) to a conventional Hadoop cluster while collecting data “in the field,” which may not be practical or desirable, the application
15 cannot take advantage of the benefits of a conventional Hadoop cluster.

SUMMARY

[0007] The technology introduced here provides the benefits of mobility along with the very high data storage capacity associated with a large scale storage cluster such as a Hadoop cluster. In particular, the technology introduced here includes a mobile data storage cluster (hereinafter "mobile cluster") and method that can be used for, among other things, collecting telemetry data and processing analytic tasks in an autonomous fashion. The mobile cluster is capable of storing a snapshot of some or all of the data of a main cluster, which can be a Hadoop cluster. The mobile cluster has a sufficiently small footprint that it can be contained within a vehicle, such as a truck. While inside of the vehicle and disconnected from the main cluster, the mobile cluster collects telemetry data from capturing devices of the vehicle and processes analytic tasks based on the snapshot and newly acquired telemetry data in real time. The mobile cluster acquires real time data and provides analytic computations in an autonomous way, without having to be connected to the main cluster while acquiring the data or performing the computations. Therefore, the mobile cluster can be used for data storage and real time analytic inquiries while collecting telemetry data in the field.

[0008] When the vehicle returns back to a depot, the mobile cluster is capable of integrating the newly acquired telemetry data into the main cluster immediately, or even integrating its own data nodes into the main cluster as a part of the main cluster.

[0009] In accordance with the techniques introduced here, therefore, a mobile cluster for collecting telemetry data and processing analytic tasks is provided. The mobile cluster in one embodiment includes a processor, a memory, a plurality of data

- nodes and an analysis module stored in the memory. The data nodes store a snapshot of at least a portion of the data stored in a main cluster and data acquired in real-time from a data capturing device after the snapshot has been stored. The analysis module is executable by the processor to process analytic tasks based on
- 5 the snapshot and the real-time acquired data without a need to connect to the main cluster. The mobile cluster may further include an integration module stored in the memory that is executable by the processor to instantly integrate the real-time acquired data into the main Hadoop cluster when being connected to the main cluster.
- 10 **[0010]** Upon reintegration of the storage of the mobile cluster to the main cluster, the telemetry data and transaction information processed by the mobile cluster are seamlessly reintegrated into the main cluster without any loss of data or transaction information. In some embodiments, the storage of the mobile cluster can be removed and replaced into the main cluster.
- 15 **[0011]** Other aspects of the technology introduced here will be apparent from the accompanying figures and from the detailed description which follows.

BRIEF DESCRIPTION OF THE DRAWINGS

[0012] These and other objects, features and characteristics of the present invention will become more apparent to those skilled in the art from a study of the following detailed description in conjunction with the appended claims and drawings, all of which form a part of this specification. In the drawings:

[0013] Figure 1 illustrates an example of a Hadoop cluster including a metadata node;

[0014] Figure 2 illustrates an example of a process according to a MapReduce framework of a Hadoop cluster;

10 [0015] Figure 3 illustrates an example of a Hadoop cluster having a JobTracker managing a MapReduce task;

[0016] Figure 4 illustrates an example of a mobile Hadoop cluster connecting to a main Hadoop cluster;

[0017] Figure 5 illustrates an example of a mobile Hadoop cluster disposed in a vehicle disconnected from the main cluster;

[0018] Figure 6 illustrates an example of a mobile Hadoop cluster returning with a vehicle to a depot;

[0019] Figure 7 illustrates an example of a process of a Hadoop cluster interacting with a main Hadoop cluster and a vehicle before, during, and after field use; and

20 [0020] Figure 8 is a high-level block diagram showing an example of the architecture of a node of a Hadoop cluster.

DETAILED DESCRIPTION

[0021] References in this specification to “an embodiment,” “one embodiment,” or the like, mean that the particular feature, structure, or characteristic being described

is included in at least one embodiment of the present invention. All occurrences of such phrases in this specification do not necessarily refer to the same embodiment.

[0022] The mobile Hadoop cluster disclosed herein can be utilized for real time data acquisition and analytic field studies, including applications for military defense and law enforcement. The mobile Hadoop cluster functions as a high performance Hadoop analytic cluster while still acquiring and ingesting field specific content. Upon returning to a home depot, the mobile Hadoop cluster can rapidly update its new information to the master system and refresh the main Hadoop cluster. This allows for mobile acquisition of analytic data as well as mobile autonomous Hadoop analytic processing. There are a variety of potential applications for the technique introduced here, including public sector applications such as law enforcement and Department of Defense, as well as internet companies in need of updating map data and other mobile data, such as a WiFi hot spot detection vehicle and the telemetry collection aircraft, as explained further below.

[0023] Figure 1 illustrates an example of a Hadoop cluster. In Figure 1, the Hadoop cluster 100 includes a metadata node 110A and a plurality of data nodes 110B, 110C and 110D. The nodes can communicate with each other through an interconnect 120. The interconnect 120 may be, for example, a local area network (LAN), wide area network (WAN), metropolitan area network (MAN), global area network such as the Internet, a Fibre Channel fabric, or any combination of such interconnects. In some embodiments, the interconnect 120 can include a network switch for processing and routing data between the nodes under network protocols, including TCP/IP. Clients 130A and 130B may communicate with the Hadoop cluster 100 via a network 140, which can be, for example, the Internet, a LAN, or any other type of network or combination of networks. Each of the clients may be, for

example, a conventional personal computer (PC), server-class computer, workstation, handheld computing/communication device, or the like. In some embodiments, the Hadoop cluster is implemented using one or more racks of commodity-class servers.

5 **[0024]** Files and data are stored in a distributed manner across the nodes in the Hadoop cluster 100 in a Hadoop Distributed File System (HDFS). To the clients of the cluster 100, including clients 130A and 130B, the HDFS serves functions of a conventional hierarchical file system. Files and data blocks of files can be created, deleted, or moved in the HDFS. There is a NameNode service 150 running on the
10 metadata node 110A to provide metadata services within HDFS, including maintaining the file system namespace and controlling access by external clients. The NameNode service can store the file system index, including the mapping of blocks to files and file system properties, in a file 160 called FsImage in the metadata node. In some embodiments, there may be a secondary metadata node running a
15 secondary NameNode service. The secondary metadata node serves as a backup in case the metadata node fails.

[0025] Each data node 110 is responsible for storing files for the HDFS. Files stored in the HDFS are divided into subsets, referred to herein as “blocks”. In one embodiment, the size of a block is 64 MBs. The blocks are typically replicated to
20 multiple data nodes. Therefore, the HDFS in the Hadoop cluster 100 can, but does not necessarily, use a conventional RAID architecture to achieve data reliability. File operations are controlled by the NameNode service 150 running on the metadata node 110A. In some embodiments, data nodes 110B, 110C, 110D are organized into racks where all nodes are connected via a network switch. The network speed
25 between nodes within a rack can be faster than the network speed between nodes in

different racks. The Hadoop cluster may take that fact into account when assigning tasks. A DataNode service 170 runs on each data node for responding read and write requests for blocks. The DataNode services 170 also respond to requests from the metadata node for creating, deleting and replicating blocks.

5 **[0026]** In some embodiments, the data nodes 110B, 110C, 110D send periodic heartbeat messages including block reports to the metadata node 110A. The metadata node 110A validates its block mapping and other file system metadata using the periodic block reports.

10 **[0027]** When a client 130A or 130B seeks to write a file to the Hadoop cluster 100, the client sends a file creation request to the metadata node 110A. The metadata node 110A responds to the client with the identity of one or more assigned data node(s) and the destination locations of the blocks of the files. The client sends data blocks of the file to the assigned data node(s); and the cluster can replicate the data blocks on one or more data nodes. Once all blocks are sent, the metadata
15 node records the file creation in its metadata, including the FsImage file.

20 **[0028]** A Hadoop cluster works as a parallel data processing engine based on a framework called MapReduce. A MapReduce framework divides a problem into multiple sub-problems, receives intermediate outputs by solving these sub-problems, and combines the intermediate outputs into an output for the problem. A
25 MapReduce application written based on the MapReduce framework can contain at least three elements: a map function, a reduce function and a main function that combines job control and file input/output. Hadoop architecture provides interfaces for a MapReduce application to run on a Hadoop cluster. Figure 2 illustrates an example of a process according to the MapReduce framework of a Hadoop cluster.
As illustrated in Figure 2, the map function takes a problem for a set of data (typically

a large set of data) and divides the problem into smaller sub-problems for smaller sub-set of data. An example application of MapReduce is a process of counting appearances of each different word in a set of documents. Each sub-problem can be responsible for counting appearances of one word in one document.

5 **[0029]** The sub-problems are distributed to nodes within the Hadoop cluster for processing. After the processing, the answers to the sub-problems (intermediate outputs) are collected and combined by the reduce function to form the output, i.e. the answer to the problem that was originally trying to solve. In the previous example of counting word appearances, the reduce function collects the intermediate
10 outputs of word appearances in each document, and sum the intermediate outputs to find the total appearances of each word in the set of documents.

[0030] The Hadoop cluster includes a JobTracker to implement the MapReduce function. As illustrated in Figure 3, a JobTracker can be implemented as a dedicated server (JobTracker node 310A) within a Hadoop cluster 300. The Hadoop cluster
15 300 includes two racks 342 and 344 interconnected via a switch 326. Rack 342 includes JobTracker node 310A, metadata node 310B, data nodes 310C-310D and a switch 322 to interconnect the nodes within rack 342. Rack 344 includes data nodes 310E-310H and a switch 324 to interconnect the nodes within rack 342. In some other embodiments, a JobTracker can be implemented as a service that shares the
20 same metadata node with a NameNode service. The JobTracker node 310A, which is dedicated to control MapReduce jobs, receives a request from a client 330 to launch a MapReduce job. Once the MapReduce job (also referred to as a MapReduce application or a MapReduce task) is submitted to the JobTracker 310A, JobTracker 310A identifies input and output files and/or directories in the HDFS for
25 the job. The input file for the MapReduce task can include a plurality of input file

blocks that contain input data for the MapReduce task. The JobTracker 310A uses knowledge of the input file blocks (including physical quantity of blocks and where the blocks are located) to decide how many subordinate tasks will be created. The MapReduce application is copied to every processing node where the input file

5 blocks are present. For each assigned processing node, JobTracker 310A creates at least one subordinate task. On each assigned processing node, a TaskTracker service monitors the status of the subordinate task on that node and reports the status and intermediate output back to the JobTracker. The Hadoop cluster 300 distributes the subordinate tasks based on the knowledge of the file blocks.

10 Therefore, instead of moving storage to the processing locations, a Hadoop cluster moves the processing tasks to the storage locations. This enables efficient processing of the data by scaling the processing with the number of nodes within the cluster.

[0031] Although the nodes 310A-310H are illustrated as single units in Figure 3,

15 each node can have a distributed architecture. For example, a node can be designed as a combination of a network module (e.g., "N-blade") and disk module (e.g., "D-blade") (not shown), which may be physically separate from each other and which may communicate with each other over a physical interconnect. Such an architecture allows convenient scaling, such as by deploying two or more N-modules

20 and D-modules, all capable of communicating with each other through the interconnect.

[0032] In accordance with the technique introduced here, to achieve a small footprint for mobility while still retaining data storage and analytic processing abilities, a mobile Hadoop cluster is provided. Figure 4 illustrates an example of a mobile

25 Hadoop cluster receiving a snapshot 490 from a main Hadoop cluster. In Figure 4,

the main Hadoop cluster 410 is a large cluster for storing a large amount of data, for instance, more than one petabyte (PB) of data. The data stored in the main Hadoop cluster 410 in one illustrative embodiment is used for analytic inquiries (also referred to as analytic tasks). For example, an analytic task can be a request to identify a particular person based on the existing data, or a request to calculate a next route.

Mobile Hadoop cluster 420 is a Hadoop cluster that has a much smaller footprint than the main Hadoop cluster 410 in terms of the number of physical nodes and the overall physical space that it occupies. For example, a main cluster may have 100 nodes occupying 20 racks, while a mobile cluster may only have 10 nodes occupying 2 racks. The mobile Hadoop cluster 420 can be installed in a vehicle 430. In one embodiment, the mobile Hadoop cluster 420 is connected with the main Hadoop cluster 410 via an interconnect 470. For example, the interconnect 470 can be a Gigabit Ethernet or other high speed local area networks (LANs). The vehicle 430 can be a land vehicle, an aircraft, a train, a watercraft, a submarine, a spacecraft or any other kind of vehicle. The vehicle 430 can be manned or unmanned.

[0033] When the mobile Hadoop cluster 420 is connected to the main Hadoop cluster 410, a snapshot 490 of the data store in the main Hadoop cluster 410 is generated and migrated into the mobile Hadoop cluster 420 via the interconnect 470. A “snapshot” is a point-in-time image of a set of data. In some other embodiments, the mobile cluster 420 has an architecture other than a Hadoop architecture.

[0034] In some embodiments, when the mobile Hadoop cluster 420 is connected to the main Hadoop cluster 410, the mobile Hadoop cluster 420 couples to the main Hadoop cluster 410 as a backup cluster of the main Hadoop cluster 410. A snapshot 490 of the file system metadata of the main cluster 410, including the FsImage file, is copied to the metadata node of the mobile Hadoop cluster 420. Data blocks are

copied to the data nodes of the mobile Hadoop cluster 420. After the data blocks and metadata are copied to the mobile Hadoop cluster 420, the mobile Hadoop cluster 420 is ready to be disconnected from the main Hadoop cluster 410, the mobile Hadoop cluster 420 contains a snapshot 490 of the data stored in the main Hadoop cluster 410. The snapshot process illustrated in Figure 4 can happen at a central base, a depot, or a place near the main cluster for a fast connection between the mobile and main clusters.

[0035] In one embodiment, a mobile Hadoop cluster is much smaller than a main Hadoop cluster. The mobile Hadoop cluster can still hold a snapshot of the data stored in the main Hadoop cluster partially because the mobile Hadoop cluster uses a smaller data replication factor, or even no data replication.

[0036] Figure 5 illustrates an example of a vehicle carrying a mobile Hadoop cluster “in the field”, i.e., disconnected from the main cluster 410. The vehicle 430 moves in the field and collects telemetry data from its telemetry data capturing device(s) 540. The telemetry data can include wireless radio signals, cellular phone signals, wireless network signals, video signals, audio signals, photos, location information, telecommand information or other telemetry data can be captured in the field. The telemetry data capturing device(s) 540 may include a camera (still and/or video) or other type of image sensor, microphone, antenna, signal sensor, signal processing devices and/or other data capturing elements. In some embodiments, the vehicle 430 can include more than one telemetry data capturing device. The telemetry data are acquired and stored in the mobile Hadoop cluster 420. Further, while in the field, the vehicle 430 and the mobile Hadoop cluster 420 do not have any physical connection to the main Hadoop cluster and need not have any type of connection (including wireless) to the main wireless cluster. The mobile Hadoop

cluster 420 storing the snapshot data from the main Hadoop cluster serves as an autonomous cluster for real time analytics. One or more client devices 550 in the vehicle 430 are connected to the mobile Hadoop cluster 420 to input analytic tasks into the cluster in real time when the vehicle 430 is in the field. The mobile Hadoop cluster 420 processes the analytic tasks based on the snapshot data and/or newly acquired telemetry data, and provides the analytic results to the client device 550. The client devices can be operated by human operators or operated automatically. In one embodiment, the client devices can be computers connected to the mobile Hadoop cluster in a SAN or NAS network in the vehicle.

10 **[0037]** For instance, the vehicle 430 can be a car or truck (or a military vehicle such as a Mine Resistant Ambush Protected vehicle) driving around a city district. The car can have a camera that captures street photos, a wireless receiver that detects WiFi hot spots, and a GPS receiver that determines the location of the vehicle. The telemetry data such as street photos, WiFi hot spots information and

15 GPS coordinates are stored in real time into the HDFS file system of the mobile Hadoop cluster. Besides the newly acquired telemetry information, the mobile Hadoop cluster in the vehicle 430 has a snapshot data of a main Hadoop cluster. The snapshot data contains telemetry data for some of the streets in the city district that has been collected previously. In one example, by using the previously-loaded

20 snapshot data and the newly acquired telemetry data, a client computer 550 in the vehicle 430 can request the mobile Hadoop cluster 420 to process an analytic task of an optimal route for the vehicle 430 for detecting more WiFi hot spots. The optimal route can be determined in real time by the history data of WiFi hot spots and street photos. For instance, a fast changing street (based on the photo analysis) with new

25 buildings may have a better possibility of detecting new WiFi hot spots; or photo

recognition may identify certain places (e.g. malls, cafeteria, transportation stations) that may have WiFi hot spots. Therefore, the mobile Hadoop cluster serves as an autonomous analytic server for analyzing the stored data and processing analytic tasks for the client device(s) in the vehicle in the field, while there is no need and/or
5 feasibility to connect to the main Hadoop cluster.

[0038] In another example, the vehicle containing the mobile Hadoop cluster can be an aircraft, such as an unmanned aerial vehicle (UAV), collecting telemetry data such as cell phone signals and video signals. The data in the mobile Hadoop cluster can include a huge amount unstructured data, i.e. data that does not fit well into
10 relational tables or pre-defined data models. The mobile Hadoop cluster in the aircraft stores the newly acquired telemetry data and compares the new data with the data base in the snapshot. If, for example, a particular voice pattern and a facial pattern of a target are identified based on the comparison, the aircraft can send an alert to a human operator or perform some other programmed action. The mobile
15 Hadoop cluster in the aircraft serves as an autonomous analytic server for processing analytic tasks including voice and facial recognition in the field.

[0039] In some embodiments, the main Hadoop cluster can contain telemetry data applicable to a large geographic region, such as an entire country; whereas the snapshot on the mobile Hadoop cluster only needs to contain the (portion of the data
20 for the) particular region to which the vehicle containing it is assigned (e.g., a particular city).

[0040] Figure 6 illustrates an example of a vehicle carrying a mobile Hadoop cluster back from the field into a depot. The mobile Hadoop cluster 420 has been gathering and storing the newly acquired telemetry data while the vehicle 430 is in
25 the field. Eventually the vehicle 430 carrying the mobile Hadoop cluster 420 returns

to a depot 680. The depot 680 has a connection 690 to the main Hadoop cluster 410. For example, the connection 690 can be a high speed Ethernet connection routed by a Gigabyte switch. Thus, when the vehicle 430 is back in the depot 680, the mobile Hadoop cluster is connected to the main Hadoop cluster 410 via the connection 690. Once the mobile Hadoop cluster establishes a connection, the mobile Hadoop cluster transfers the newly acquired telemetry data 695 to the main Hadoop cluster 410 via the connection 690 and interconnect 470. In addition, the data on the main Hadoop cluster 410 may have been updated while the vehicle is on the field. The main Hadoop cluster 410 can transfer the changes to the data on the main Hadoop cluster 410 since the last time cluster 410 was connected to the mobile Hadoop cluster 420. Therefore, the mobile Hadoop cluster 420 rapidly offloads and refreshes data while the vehicle 430 is in the depot 680. In another embodiment, the main Hadoop cluster 410 can transfer the actual data as modified on the main cluster since the last time cluster 410 was connected to the mobile Hadoop cluster 420.

[0041] In one embodiment, The Fslmage of the mobile Hadoop cluster will be integrated into the main Hadoop cluster in a substantially instant fashion, once the mobile system is brought on-line in the main cluster, as described further below. All new data and processed information are synchronized.

[0042] In some embodiments, in addition to the newly acquired telemetry data, the information generated from the real-time analysis of the mobile Hadoop cluster is transferred to the main mobile Hadoop cluster. In the above-described example of WiFi spot detection, the mobile Hadoop cluster in the vehicle can analyze the data and generate a list of possible places for finding WiFi hotspots in real time. This list of possible places can be transferred back to the main Hadoop to be used for future

exploration. In the previous above-described example of an aircraft, the mobile Hadoop cluster in the aircraft can store recognized identities of certain individuals and the location and time when the individuals are identified. That information can be transferred back to the main Hadoop cluster as, for example, entries of a
5 database for identified individuals.

[0043] In some embodiments, the nodes in the mobile Hadoop cluster 420 are modular components (e.g. server “blades”) that are built from commodity hardware. The data nodes in the mobile Hadoop cluster 420 can be easily removed from a rack of the cluster 420 and then placed in a rack as a part of the main Hadoop cluster
10 410. Then the main Hadoop cluster 410 recognizes the newly placed data nodes and starts to replicate the blocks from newly placed nodes if necessary. The differences of the file system metadata between the main and mobile Hadoop cluster will be applied to update the file system of the main Hadoop cluster. The ingestion of the new data is instant, because the new data has been already stored in the data
15 nodes that are now a part of the main Hadoop cluster. Upon receiving an integrating command, the main Hadoop cluster is instantly aware of the new data. The integrating command updates the metadata maintained by the NameNode service to reflect the integration of the new data and/or modified data.

[0044] Typically data blocks on the main Hadoop cluster are replicated by a factor
20 of at least three, i.e., two copies are created for each stored data block. To maintain the mobile Hadoop cluster's small footprint, the blocks on the mobile Hadoop cluster can be replicated by a factor smaller than the factor of the main Hadoop cluster replication. In some embodiments, the replication factor on mobile Hadoop cluster is one (i.e. no block replication). In some embodiments, the replication factor on a
25 Hadoop cluster is two, i.e. one copy is created for each stored data block. In some

embodiments, the data on the mobile Hadoop cluster can be protected by other redundancy architectures including RAID.

[0045] In some embodiments, in order to achieve a small footprint, the mobile Hadoop cluster has less processor (CPU) power than the main Hadoop cluster.

5 While the mobile cluster may be responsible for performing some real-time analytics in the field, other analytics that are more computationally intensive can be done in the main Hadoop cluster.

[0046] In some embodiments, the storage devices (e.g. storage disks or non-volatile solid state memory such as flash) of the mobile Hadoop cluster are

10 connected to the CPUs via Serial Attached SCSI (SAS) to achieve fast communication between the storage units and CPUs, while the storage units can be detached from the mobile Hadoop cluster and be inserted in a rack of the main Hadoop cluster. When the vehicle returns to the depot, the storage units can be detached from the mobile Hadoop cluster (and the vehicle), and then be inserted into
15 and be a part of the main Hadoop cluster.

[0047] Figure 7 illustrates an example process of a mobile Hadoop cluster interacting with a main Hadoop cluster and a vehicle before, during, and after field use. At 710, the main Hadoop cluster or the mobile Hadoop cluster generates a snapshot of at least a portion of data stored in the main Hadoop cluster. Then the

20 snapshot is transferred to the mobile Hadoop cluster (720). Once the vehicle carrying the mobile Hadoop cluster drives off, at 730, the mobile Hadoop cluster starts collecting real-time acquired data from a data capturing device in a vehicle. While collecting the real-time data, the mobile Hadoop cluster may process analytic queries from one or more client devices in the vehicle; the analysis performed by the
25 mobile cluster in response to such queries is based on the snapshot and the real-

time acquired data (740). For example, the analysis performed can be voice or facial recognition in the example of a telemetry collecting aircraft.

[0048] When the vehicle returns to a depot (750), the mobile Hadoop cluster is connected to the main Hadoop cluster via a connection, and the real-time acquired data is integrated into the main Hadoop cluster with minimal delay or essentially instantly. In one embodiment, the mobile Hadoop cluster only moves the data blocks and associated metadata for the real-time acquired data to a corresponding data node within the main Hadoop cluster. Once the mobile Hadoop cluster has moved the data blocks for the real-time acquired data, the main Hadoop cluster may initiate a process for replicating the newly removed data blocks based on the predetermined replication factor of the main Hadoop cluster. In some other embodiments, instead of connecting the mobile and main Hadoop clusters with a connection, the data nodes or the storage units storing the real-time acquired data are physically removed from their mounting locations in the mobile Hadoop cluster and integrated into mounting locations in the main Hadoop cluster. Further, at 760 of the example process, an update from the main Hadoop cluster is transferred to the mobile Hadoop cluster, wherein the update includes changes of the data stored in the main Hadoop cluster since last time of connecting to the main Hadoop cluster. The update is applied to the snapshot stored in the mobile Hadoop cluster (770).

[0049] Figure 8 is a high-level block diagram showing an example of the architecture of a node 800, which may represent any of nodes 110A-110D or 310A-310H. The node 800 includes one or more processors 810 and memory 820 coupled to an interconnect 830. The interconnect 830 shown in Figure 8 is an abstraction that represents any one or more separate physical buses, point to point connections, or both connected by appropriate bridges, adapters, or controllers. The

interconnect 830, therefore, may include, for example, a system bus, a Peripheral Component Interconnect (PCI) bus or PCI-Express bus, a HyperTransport or industry standard architecture (ISA) bus, a small computer system interface (SCSI) bus, a universal serial bus (USB), IIC (I2C) bus, or an Institute of Electrical and Electronics Engineers (IEEE) standard 1394 bus, also called "Firewire".

[0050] The processor(s) 810 is/are the central processing unit (CPU) of the storage controller 800 and, thus, control the overall operation of the node 800. In certain embodiments, the processor(s) 810 accomplish this by executing software or firmware stored in memory 820. The processor(s) 810 may be, or may include, one or more programmable general-purpose or special-purpose microprocessors, digital signal processors (DSPs), programmable controllers, application specific integrated circuits (ASICs), programmable logic devices (PLDs), trusted platform modules (TPMs), or the like, or a combination of such devices.

[0051] The memory 820 is or includes the main memory of the node 800. The memory 820 represents any form of random access memory (RAM), read-only memory (ROM), flash memory, or the like, or a combination of such devices. In use, the memory 820 may contain, among other things, code 870 embodying at least a portion of a storage operating system of the node 800. Code 870 may also include a snapshot module stored in the memory that is executable by the processor to generate a snapshot of at least a portion of the data stored in the data nodes, an analysis module stored in the memory that is executable by the processor to process analytic tasks based on the snapshot and real-time acquired data without a need to connect to a main Hadoop cluster, and/or an integration module stored in the memory that is executable by the processor to virtually instantly integrate real-time acquired data into a main Hadoop cluster.

[0052] Also connected to the processor(s) 810 through the interconnect 830 are a network adapter 840 and a storage adapter 850. The network adapter 840 provides the node 800 with the ability to communicate with mobile devices, such as clients 130A or 130B, over a network and may be, for example, an Ethernet adapter or
5 Fibre Channel adapter. The network adapter 840 may also provide the node 800 with the ability to communicate with other nodes within the data storage cluster. In some embodiments, a node may use more than one network adapter to deal with the communications within and outside of the data storage cluster separately. The storage adapter 850 allows the node 800 to access a persistent storage, and may
10 be, for example, a Fibre Channel adapter or SCSI adapter. Node 800 can further include a data input adapter 860 for receiving acquired data from a data capturing device on or in a vehicle.

[0053] The code 870 stored in memory 820 may be implemented as software and/or firmware to program the processor(s) 810 to carry out actions described
15 below. In certain embodiments, such software or firmware may be initially provided to the node 800 by downloading it from a mobile system through the node 800 (e.g., via network adapter 840).

[0054] The techniques introduced herein can be implemented by, for example, programmable circuitry (e.g., one or more microprocessors) programmed with
20 software and/or firmware, or entirely in special-purpose hardwired circuitry, or in a combination of such forms. Special-purpose hardwired circuitry may be in the form of, for example, one or more application-specific integrated circuits (ASICs), programmable logic devices (PLDs), field-programmable gate arrays (FPGAs), etc.

[0055] Software or firmware for use in implementing the techniques introduced
25 here may be stored on a machine-readable storage medium and may be executed

by one or more general-purpose or special-purpose programmable microprocessors.

A "machine-readable storage medium", as the term is used herein, includes any mechanism that can store information in a form accessible by a machine (a machine may be, for example, a computer, network device, cellular phone, personal digital assistant (PDA), manufacturing tool, any device with one or more processors, etc.).

For example, a machine-accessible storage medium includes recordable/non-recordable media (e.g., read-only memory (ROM); random access memory (RAM); magnetic disk storage media; optical storage media; flash memory devices; etc.), etc.

10 **[0056]** The term "logic", as used herein, can include, for example, programmable circuitry programmed with specific software and/or firmware, special-purpose hardwired circuitry, or a combination thereof.

[0057] In addition to the above mentioned examples, various other modifications and alterations of the invention may be made without departing from the invention.

15 Accordingly, the above disclosure is not to be considered as limiting and the appended claims are to be interpreted as encompassing the true spirit and the entire scope of the invention.

CLAIMS

What is claimed is:

1. A storage cluster comprising:
 - a processor;
 - 5 a plurality of data nodes for receiving and storing a snapshot of at least a portion of data stored in a main storage cluster and acquired data received from a data capturing device; and
 - an analysis module operatively coupled to the processor to process analytic tasks based on the snapshot and the acquired data when the storage cluster is not
 - 10 connected to the main storage cluster.
2. The storage cluster of claim 1, further comprising:
 - a connection; and
 - an integration module to integrate the acquired data into the main storage cluster via the connection.
- 15 3. The storage cluster of claim 1, wherein the storage cluster is disposed in a vehicle, and the data capturing device is disposed in or on the vehicle.
4. The storage cluster of claim 3, wherein the analysis module is configured to process analytic tasks from a client device within the vehicle based on the snapshot and the acquired data without being connected to the main storage cluster.
- 20 5. The storage cluster of claim 1, further comprising:
 - a connection to retrieve incoming data from the main storage cluster; and
 - a snapshot module to generate the snapshot of at least the portion of the data store in the main storage cluster based on the incoming data.

6. The storage cluster of claim 1, wherein the acquired data includes telemetry data, and the data capturing device comprises a telemetry data collecting device.
7. The storage cluster of claim 1, wherein the data nodes are removable from the storage cluster and configured to be integrated into the main storage cluster.
- 5 8. The storage cluster of claim 1, wherein the storage cluster is a Hadoop storage cluster, which is separated from the main storage cluster when the storage cluster receives the acquired data.
9. The storage cluster of claim 5, wherein, when the storage cluster is connected to the main storage cluster via the connection, the snapshot module is configured to
- 10 further receive an update via the connection from the main storage cluster and to apply the update to the snapshot; and wherein the update includes changes of the data stored in the main storage cluster since a last time the storage cluster was connected to the main storage cluster.
10. The storage cluster of claim 1, wherein the main storage cluster is a Hadoop
- 15 storage cluster.
11. A Hadoop cluster comprising:
- a processor;
 - a plurality of data nodes for storing data;
 - a snapshot module operatively coupled to the processor to generate a
- 20 snapshot of at least a portion of the data stored in the data nodes; and
- a connection configured to transfer the snapshot to a mobile cluster.
12. The Hadoop cluster of claim 11, further comprising:
- an integration module stored in a memory of the Hadoop cluster that is executable by the processor to instantly integrate real-time acquired data from the

mobile cluster into the Hadoop cluster via the connection, when the mobile cluster is connected to the Hadoop cluster via the connection.

13. The Hadoop cluster of claim 12, wherein the real-time acquired data is telemetry data collected into the mobile cluster by a telemetry data collecting device when the mobile cluster is not connected to the Hadoop cluster.

14. The Hadoop cluster of claim 11, wherein the mobile cluster is disposed in a vehicle when the mobile cluster is not connected to the Hadoop cluster, and the real-time acquired data is received from a data capturing device within the vehicle.

15. The Hadoop cluster of claim 11, wherein the snapshot enables the mobile cluster to process analytic tasks based on the snapshot without a need to connect to the Hadoop cluster.

16. The Hadoop cluster of claim 11, wherein the Hadoop cluster is configured to integrate a plurality of mobile data nodes removable from the mobile cluster.

17. The Hadoop cluster of claim 11, wherein, when the mobile cluster is connected to the main Hadoop cluster via the connection, the snapshot module is executable by the processor to further generate an update of the data and transmit the update via the connection to the mobile cluster; and wherein the update includes changes of the data stored in the Hadoop cluster since last time the mobile cluster was connected to the Hadoop cluster.

18. A computer implemented method comprising:

receiving, at a mobile storage cluster, a snapshot of at least a portion of data stored in a main storage cluster;

collecting, at the mobile storage cluster, real-time acquired data from a data capturing device; and

processing analytic tasks in the mobile storage cluster based on the snapshot and the real-time acquired data when the mobile storage cluster is not connected to the main storage cluster.

19. The method of claim 18, further comprising:

5 integrating the real-time acquired data into the main storage cluster instantly via a connection when the mobile storage cluster is connected to the main storage cluster.

20. The method of claim 18, wherein said collecting real-time acquired data comprises:

10 collecting real-time acquired data from a data capturing device disposed in or on a vehicle.

21. The method of claim 18, wherein said processing analytic tasks comprises:

 processing analytic tasks from a client device in the mobile storage cluster based on the snapshot and the real-time acquired data when the mobile storage
15 cluster is not connected to the main storage cluster, when the device and the mobile storage cluster are disposed in a vehicle.

22. The method of claim 18, further comprising:

 generating the snapshot of at least the portion of data stored in the main storage cluster.

20 23. The method of claim 18, wherein the real-time acquired data is telemetry data, and the data capturing device is a telemetry data collecting device.

24. The method of claim 18, further comprising:

 removing data nodes storing the real-time acquired data from the mobile storage cluster; and

25 integrating the data nodes into the main storage cluster.

25. The method of claim 18, further comprising:

receiving an update from the main storage cluster when the mobile storage cluster is connected to the main storage cluster, wherein the update includes

5 changes of the data stored in the main storage cluster since a last time when the mobile storage cluster was connected to the main storage cluster; and

applying the update to the snapshot stored in the mobile storage cluster.

26. The method of claim 18, wherein the main storage cluster is a Hadoop storage cluster.

10 27. A data storage node comprising:

a processor;

a connection interface through which to contact one or more other data storage nodes within a mobile storage cluster for accessing a snapshot of at least a portion of data stored in a main storage cluster; and

15 an analysis module operatively coupled to the processor to process analytic tasks based on the snapshot and acquired data received from a data capturing device, when the mobile storage cluster is not connected to the main storage cluster.

28. The data storage node of claim 27, further comprising:

20 an integration module to integrate the acquired data into the main storage cluster when the mobile storage cluster is connected to the main storage cluster.

29. The data storage node of claim 27, wherein the analysis module is configured to process analytic tasks from a client device within a vehicle based on the snapshot and the acquired data when the mobile storage cluster is not connected to the main storage cluster.

25

30. The data storage node of claim 27, wherein the mobile storage cluster is a mobile Hadoop storage cluster, which is separated from the main storage cluster when the mobile Hadoop storage cluster receives the acquired data.

31. A computer implemented method comprising:

5 receiving, at a data storage node of a mobile storage cluster, an analytic task from a client device;

 contacting one or more other data storage nodes within a main storage cluster for accessing a snapshot of at least a portion of data stored in the main storage cluster; and

10 processing the analytic task, at the mobile storage cluster, based on the snapshot and acquired data from a data capturing device, when the mobile storage cluster is not connected to the main storage cluster.







FIG. 1

**FIG. 2**

**FIG. 3**






FIG. 5


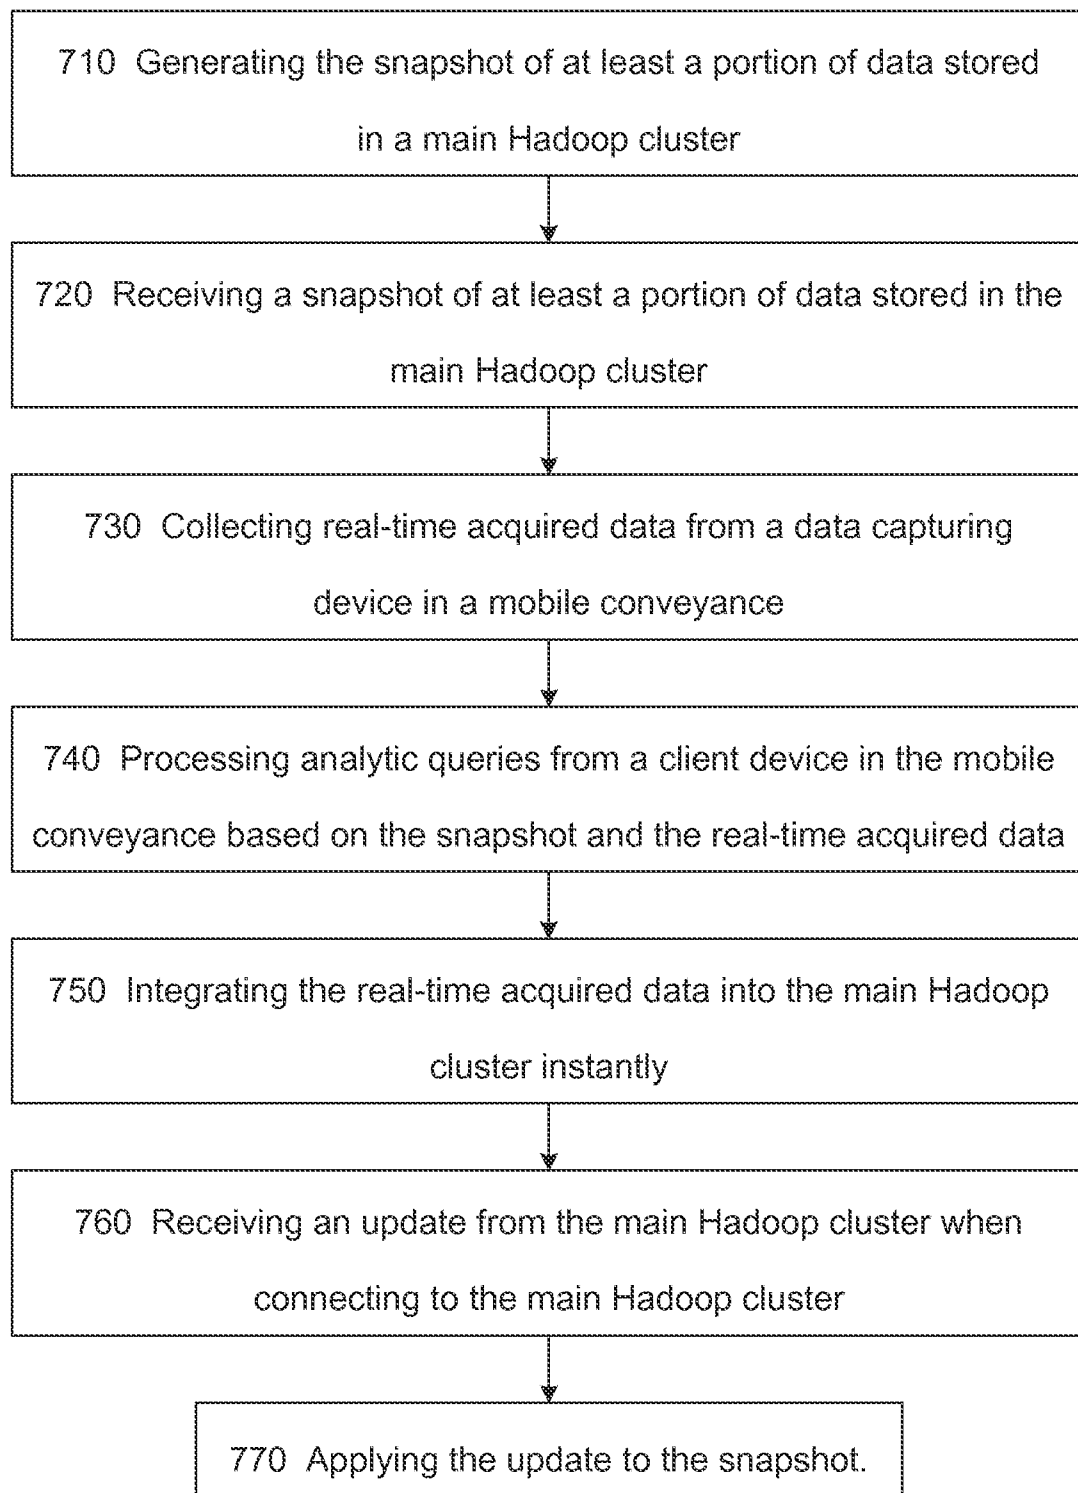



FIG. 6

**FIG. 7**

**FIG. 8**

A. CLASSIFICATION OF SUBJECT MATTER**G06F 17/00(2006.01)i, G06F 17/40(2006.01)i**

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

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

G06F 17/00; G06F 12/16; G06F 17/30; G06F 9/455; G06F 21/00; G06F 17/40

Documentation 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 models

Electronic data base consulted during the international search (name of data base and, where practicable, search terms used)

eKOMPASS(KIPO internal) & Keywords: data node, snapshot, cluster, analysis, Hadoop, mobile, vehicle, storage, connection, disconnection, and similar terms

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 2011-0252208 A1 (ALI, ABID et al.) 13 October 2011 See paragraphs [0011]-[0036] and [0064]-[0065]; claims 1-2 and 11; and figures 1-6.	11,16
A		1-10, 12-15, 17-31
Y	US 2011-0313973 A1 (SRIVAS, MANDAYAM C. et al.) 22 December 2011 See paragraphs [0045]-[0083] and [0105]; claims 1-5; and figures 1-7.	11,16
A	WO 2011-159842 A2 (NIMBULA, INC.) 22 December 2011 See paragraphs [0165]-[0174] and [0201]-[0211]; claims 1-10; and figures 1-7.	1-31
A	US 2008-0243866 A1 (PANDEY, MANU) 2 October 2008 See paragraphs [0028]-[0051] and figures 2-5.	1-31
A	US 2012-0005164 A1 (PATANKAR, SWAPNIL DILIP et al.) 5 January 2012 See paragraphs [0024]-[0030] and [0061]-[0072]; claims 1-3; and figures 1-8.	1-31



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 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

26 November 2013 (26.11.2013)

Date of mailing of the international search report

27 November 2013 (27.11.2013)

Name and mailing address of the ISA/KR



 Korean Intellectual Property Office
 189 Cheongsa-ro, Seo-gu, Daejeon Metropolitan City,
 302-701, Republic of Korea

Facsimile No. +82-42-472-7140

Authorized officer

NHO, Ji Myong

Telephone No. +82-42-481-8528



INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2013/053239

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2011-0252208 A1	13/10/2011	CN 102834822 A EP 2558949 A2 WO 2011-129987 A2 WO 2011-129987 A3	19/12/2012 20/02/2013 20/10/2011 12/01/2012
US 2011-0313973 A1	22/12/2011	EP 2583180 A1 WO 2011-160054 A1	24/04/2013 22/12/2011
WO 2011-159842 A2	22/12/2011	EP 2583211 A2 US 2012-110055 A1 US 2012-110056 A1 US 2012-110180 A1 US 2012-110188 A1 US 2012-110636 A1 US 2012-110650 A1 US 2012-110651 A1 US 2012-116937 A1 US 2012-117229 A1 US 2013-060839 A1 WO 2011-159842 A3	24/04/2013 03/05/2012 03/05/2012 03/05/2012 03/05/2012 03/05/2012 03/05/2012 03/05/2012 10/05/2012 10/05/2012 07/03/2013 01/03/2012
US 2008-0243866 A1	02/10/2008	US 8112593 B2 WO 2008-121284 A1	07/02/2012 09/10/2008
US 2012-0005164 A1	05/01/2012	None	