



- (51) **International Patent Classification:**
H04L 12/26 (2006.01) *H04L 12/24* (2006.01)
- (21) **International Application Number:**
PCT/US2015/038532
- (22) **International Filing Date:**
30 June 2015 (30.06.2015)
- (25) **Filing Language:** English
- (26) **Publication Language:** English
- (30) **Priority Data:**
6287/CHE/2014 12 December 2014 (12.12.2014) IN
- (71) **Applicant:** HEWLETT PACKARD ENTERPRISE DEVELOPMENT LP [US/US]; 11445 Compaq Center Drive West, Houston, TX 77070 (US).
- (72) **Inventors:** JAGANNATH, Kishore; Sy.No.192, Whitefield Road, Mahadevapura Post, Bangalore 560048 (IN). SRINIVAS, Santhosh; Sy.No.192, Whitefield Road, Mahadevapura Road, Bangalore 560048 (IN).
- (74) **Agent:** WOODS, Ariana G.; Hewlett Packard Enterprise, 3404 E. Harmony Road, Mail Stop 79, Fort Collins, CO 80528 (US).
- (81) **Designated States** (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY,

BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SA, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

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

Declarations under Rule 4.17:

- as to the identity of the inventor (Rule 4.17(i))
- as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))

Published:

- with international search report (Art. 21(3))

(54) **Title:** CLOUD SERVICE TUNING




Figure 1

(57) **Abstract:** According to an example, to tune a cloud service, an event and an event type on a cloud server application are detected by monitoring at least one of a hardware value and a software value. An application parameter to tune is determined based on the event type, and a tuning priority of the application parameter based on historical tuning data is fetched. The application parameter is tuned by modifying a software variable. In the event that tuning the application parameter resolves the event type, the tuning priority of the parameter is increased. In the event that tuning the parameter fails to resolve the event type, the parameter priority is decreased, the tuning is rolled back, and a default tune setting is invoked.

CLOUD SERVICE TUNING

BACKGROUND

[0001] Cloud computing services have become an increasingly popular alternative to purchasing and maintaining computer equipment locally, particularly at the enterprise level. The cloud computing market now includes a myriad of service types and service levels of varying scope, scalability, and performance.

BRIEF DESCRIPTION OF THE DRAWINGS

[0002] FIG. 1 illustrates a cloud system including a developer, cloud servers, and an operations site, according to an example of the present disclosure;

[0003] FIG. 2 illustrates a flowchart of tuning a cloud service application, according to an example of the present disclosure;

[0004] FIG. 3 illustrates a flowchart of deploying a tunable cloud service application, according to an example of the present disclosure; and

[0005] FIG. 4 illustrates a schematic representation of a computing device that may be used as a platform for implementing or executing at least one of the processes depicted in FIGS. 2-3, according to an example of the present disclosure.

DETAILED DESCRIPTION

[0006] Various embodiments described below provide for automatic or “self” tuning of a cloud server or cloud application on a cloud server, or auto-remediation of events detected on a cloud server by tuning a parameter or parameters based on a priority level and other data. In some embodiments, tuning parameters are exposed by application developers, along with other data, during the application development

and/or deployment process to enable self-tuning or auto-remediation based on certain criteria.

[0007] Cloud computing services (hereinafter simply “cloud services”) may comprise servers, databases, networking equipment, and other computing equipment maintained by a cloud computing provider, commonly referred to as a cloud service provider or CSP. Cloud services may include virtual machines or other distributed processing systems that may be used to host servers, websites, software, databases, and other cloud applications in the consumer, commercial, and industrial sectors.

[0008] Cloud services and cloud applications may be accessed by computing systems, devices, and electronic components such as desktop computers, laptop computers, servers, thin clients, tablets, smartphones, digital video recorders, retail point of sale devices, and other computing equipment (hereinafter “device” or “devices”). Access may be provided via public or private networks, such as the internet or intranets.

[0009] In response to the variable demand for cloud services at any given time, techniques such as auto scaling (also known as “autoscaling” or “auto-scaling”) may be used to dynamically adjust available computing resources available to a cloud service application at any given time, so that an application does not need to provision a high-level of resources even during times of low demand. For example, auto scaling may allow a cloud service to dynamically increase or decrease the number of virtual machines running a particular cloud service application. Auto scaling may be based on a trigger such as server performance, application performance, response times, or site traffic, as examples.

[0010] While auto scaling may be beneficial to a subscriber of cloud services by reducing overall cost to run a cloud application, auto scaling may not be the most efficient technique for adapting to events on a server such as performance degradation or increased traffic, either from the perspective of computational efficiency or cost-effectiveness. For example, auto scaling may be a generic approach to resolving performance issues on a cloud application that may come at a significant financial cost. In contrast, a more tailored approach, such as tuning application parameters, may resolve the same issues at no cost.

[0011] As a specific example, a monitoring tool may detect that a CPU utilization rate on a server running a cloud application is high and reaching or exceeding a threshold. A generic approach to resolving the high CPU utilization rate may be to auto scale the cloud services to allow for more computing resources, e.g., to allocate more CPUs, to the particular cloud application. However, the primary cause of the high CPU utilization rate may be an issue or event such as a high number of threads configured in a thread pool causing a context switching overhead issue. In this example, lowering the number of threads in the thread pool by tuning an application parameter, at no financial cost, may be an effective or more effective solution to resolving the issue than simply increasing overall computing resources available to the application, e.g., by auto scaling.

[0012] However, tuning a particular application parameter may require knowledge of the development and intricacies of the application. In cases where an application is developed by a development team and then passed to an operations team, the development details are often lost, restricting the ability to optimally self-

tune the application once deployed, or to do so in a timely fashion to prevent, e.g., application outages resulting from an unexpected load or traffic increase.

[0013] According to an example, to tune a cloud service, an event and an event type on a cloud server application are detected by monitoring at least one of a hardware value and a software value. An application parameter to tune is determined based on the event type, and a tuning priority of the application parameter based on historical tuning data is fetched. The application parameter is tuned by modifying a software variable. In the event that tuning the application parameter resolves the event type, the tuning priority of the parameter is increased. In the event that tuning the parameter fails to resolve the event type, the parameter priority is decreased, the tuning is rolled back, and a default tune setting is invoked.

[0014] FIG. 1 illustrates a cloud system including a developer, cloud servers, and an operations site, according to an example of the present disclosure.

[0015] In an example, the operations site includes a monitoring server 102 which may be coupled to a monitoring database 104. Server 102 and database 104 may monitor, receive, calculate, and/or store data related to various cloud services, such as names, descriptions, service types, performance metrics, historical data, and other information. In some examples, users of the cloud servers discussed below may grant permission or opt-in to be monitored by monitoring server 102.

[0016] Tuning server 106 may be coupled to a tuning parameters database 108. Server 106 and database 108 may receive, calculate, and/or store data related to tunable parameters that may be tuned on a particular cloud server or in a particular cloud application, including historical tuning data and priority levels,

discussed below in more detail. Tuning server 106 may also store tuning parameters set by a developer pre-deployment or post-deployment.

[0017] In some embodiments, the functions of monitoring server 102 and tuning server 106 may be incorporated directly into a cloud server running a cloud application, as discussed below in more detail with respect to servers 110A-C.

[0018] Network or cloud 114 may connect the operations site with cloud service providers and customers, discussed in more detail below. Network or cloud 114 may be any public or private network such as the internet or intranets.

[0019] Cloud servers 110A-C may represent any cloud service hardware and/or software available to run a cloud application. Various applications and service types may run on cloud servers 110A-C. For example, cloud servers 110A-C may offer web host, database, and/or software applications or service types. Cloud servers 110A-C may also store tuning parameters and auto scaling parameters 112A-C, including those set by a developer pre-deployment or post-deployment.

[0020] Developer device 118 may be any device such as a computer used by a device to access cloud servers 110A-C or to deploy a cloud application to cloud servers 110A-C. Developer device 118 may also be used to transmit data such as developer tuning preferences or priorities, discussed in more detail below, to cloud servers 110A-C or to the operations site.

[0021] FIG. 2 illustrates a flowchart of tuning a cloud service application, according to an example of the present disclosure. In block 202, an event is detected on a cloud server application. In some examples, the event may be a performance degradation event type, while in other examples the event may be a threshold level alert. The event may be detected based on monitoring of a hardware variable or a

software variable, or other metrics, values, or inputs, such as the CPU utilization or response time examples discussed herein.

[0022] In block 204, application parameters available to tune may be determined based on the event type. During development and/or deployment of a cloud application, as discussed below in more detail with respect to FIG. 3, various parameters may be identified as tunable, and more specifically as tunable in a production environment. For example, as discussed above, the number of threads in a thread pool configuration file for a cloud application may be identified as a parameter (or value or variable) that is tunable once a cloud application is deployed to a test or production cloud server. As other examples, a database connection size or a caching value may be tunable parameters. Parameters may be related to the application itself, or other parameters that interface with the application such as system parameters.

[0023] As discussed below in more detail with respect to FIG. 3, the tunable parameters may be associated with an event type. For example, a CPU performance degradation event type detected on a cloud server may be associated with the thread pool parameter and other CPU-related parameters, while a database threshold level event type may be associated with a database connection size parameter or caching parameters. In each case, the tunable or “exposed” parameters may be determined in block 204 by, for example, retrieving a list of developer-exposed parameters or scanning the cloud application.

[0024] In block 206, a tuning priority for each parameter may be determined. A tuning priority may be a value to indicate the order of parameters to be tuned, from a first priority to a last priority, and may be grouped by each event type. For

example, a database-related event type may have a database connection size parameter and a caching parameter exposed for tuning in production, with the database connection size having a priority of 1, and the caching parameter having a priority of 2. In examples where two parameters have the same priority level, a secondary rule such as time of creation may be used.

[0025] In some examples, the tuning priority may have been adjusted previously, either by the flow of FIG. 2 or by a developer setting, and thus the current priority may be based on historical data.

[0026] In block 208, self-tuning or “auto-remediation” of the parameters may commence, with the highest-priority parameter the first to be tuned. Tuning a parameter may comprise adjusting a software or firmware variable or value on the cloud server, such as increasing a number of worker threads or decreasing a cache size.

[0027] In block 210, a determination may be made as to whether the tuned parameter resolved the event type. For example, a determination may be made as to whether increasing the number of worker threads resolved a CPU-related event type by determining whether the CPU utilization rate has been lowered or whether CPU performance has improved. The determination of block 210 may be based on an updated monitoring or detection, as discussed above with respect to block 202. In some examples, the updated monitoring or detection may occur at a preset time interval after tuning a parameter, to allow the parameter tuning to take effect.

[0028] If the tuned parameter resolved the event type, and if the parameter priority is not already at the highest level, the priority of the parameter may be increased in block 212 for future tuning. For example, if a first loop of the flow of FIG.

2 did not resolve the event type by tuning a database connection size with a priority 1, but did resolve the event type by tuning a caching parameter having a priority of 2, the caching parameter priority may be upgraded or promoted to a priority level of 1. Similarly, if the tuned parameter did not resolve the event type, and if the priority is not already at the lowest level, the priority of the parameter may be decreased in block 214 for future tuning. The priority levels may be stored as historical data, e.g., on tuning server 108 or cloud servers 110A-C, or transmitted back to a developer.

[0029] In block 216, if the tuned parameter did not resolve the event type, the parameter tuning may be rolled back. In some examples, the parameter may be rolled back by adjusting a single software variable or value. In other examples, the roll back may require adjusting several settings, or rolling back to an earlier version or state of a server image, virtual server image, operating system image, firmware image, or application image. Rolling back the tuning may allow for a more accurate reading of whether or not a particular parameter tuning was effective in resolving an event type. In some examples, the roll back step may be omitted such that parameter tuning may be “on top” of a previous tuning of another parameter.

[0030] In some examples, block 216 may flow back to block 206 if additional parameters of lower priority are to be tuned. In other examples, the flow from block 216 may loop back to block 208, i.e., the same parameter may be tuned incrementally or iteratively. As one example, a database cache size parameter may be tuned in steps before progressing to a lower-priority parameter, if necessary. A maximum allowable count of tuning attempts may be set for a single parameter, or across all parameters.

[0031] In block 218, if tuning all available parameters as shown in the loop of blocks 206 through 216 has not resolved the event type, a default tuning action may be invoked on the cloud server. For example, an auto scaling function may be called to increase overall computing resources on the cloud server or service available to the cloud application.

[0032] In some examples, a report of the tuning parameters, priorities, and changes to the priorities may be output to a developer. The report may be useful to developers in identifying and setting parameters and tuning preferences both pre-deployment and post-deployment, which may later be transmitted back to the cloud services or operations site through the flow of FIG. 2.

[0033] FIG. 3 illustrates a flowchart of deploying a tunable cloud service application, according to an example of the present disclosure. In some examples, the flow of FIG. 3 may occur pre-deployment or post-deployment of a cloud application by a software developer.

[0034] In an example, in block 302, tunable application parameters may be associated with event types, as discussed above. The association may be based on a rule or set of rules based on a known association between an event type and a parameter or parameters.

[0035] In block 304, the tunable parameters may be set for exposure or visibility to a cloud server, or to an operations team. Exposing the parameter may involve setting the cloud application to accept incoming requests to adjust or tune a particular parameter, which may not otherwise be accessible or visible to a cloud server or operations team, for example. In some examples, exposing tunable

parameters may also comprise a security or signature verification to restrict the entities able to read, modify, or tune the parameters.

[0036] In block 306, the priority levels for each tunable parameter may be set, and in some examples, the priorities for each parameter may be adjustable based on a particular event type or a developer preference. For example, a database connection size parameter may be set with a priority 1, while a caching parameter may be set with a priority of 2, for a database performance degradation event type. In the example of a CPU performance degradation event type, the priority levels may be set lower for such database-related parameters.

[0037] The priority levels may be based on a known association between an event type and the relative success or failure of a particular parameter in resolving the event type. In some examples, the priority levels may be set with input from the report of FIG. 2 discussed above.

[0038] In block 308, a rollback point for each tunable parameter may be set. In some examples, the rollback point may be the point in time immediately before the last parameter was tuned. In other examples, the rollback point may be the last time a service patch or other update was installed.

[0039] In block 310, monitor thresholds for each event type may be set. For example, for a particular cloud application, a CPU monitor threshold of 80% utilization may be set, or a connection response time threshold of 100 milliseconds may be set. The thresholds, ranges, or other variables of block 310 may be used on the cloud server to determine when to trigger an event and invoke self-tuning or auto-remediation as discussed with respect to FIG. 2.

[0040] In block 312, a default tune setting or action may be set. In some examples, the default tune action may be an auto scaling function. In other examples, the default tune action may comprise alerting an operations team.

[0041] In block 314, the application may be deployed to the operations team or to a cloud service or server, e.g., for release to production.

[0042] FIG. 4 illustrates a schematic representation of a computing device that may be used as a platform for implementing or executing at least one of the processes depicted in FIGS. 2-3, according to an example of the present disclosure.

[0043] In an example, device 402 comprises a processing resource such as processor or CPU 404; a memory resource 406 such as memory, RAM, ROM, or Flash memory; a disk drive 410 such as a hard disk drive or a solid state disk drive; a firmware 412; an operating system 414; and a network interface 416 such as a Local Area Network LAN card, a wireless 802.11x LAN card, a 3G or 4G mobile WAN, or a WiMax WAN card. Each of these components may be operatively coupled to a bus.

[0044] Some or all of the operations set forth in the figures may be contained as a utility, program, or subprogram in any desired computer readable storage medium, or embedded on hardware. The computer readable medium may be any suitable medium that participates in providing instructions to the processing resource 404 for execution. For example, the computer readable medium may be non-volatile media, such as an optical or a magnetic disk, or volatile media, such as memory. The computer readable medium may also store other machine-readable instructions, including instructions downloaded from a network or the internet.


[0045] In addition, the operations may be embodied by machine-readable instructions. For example, they may exist as machine-readable instructions in source code, object code, executable code, or other formats.

[0046] Device 402 may comprise, for example, a computer readable medium that may comprise instructions 408A-C to associate at least one tunable application parameter with at least one event type; expose the at least one tunable application parameter to a production cloud server for auto-remediation on the production cloud server; set an adjustable priority level for the at least one application parameter; set a rollback point for the at least one application parameter; set a monitor threshold for the at least one event type; set a default tune action; and deploy the tunable cloud services application to the production cloud server.

[0047] The computer-readable medium may also store an operating system such as Microsoft Windows, Mac OS, Unix, or Linux; network applications such as network interfaces and/or cloud interfaces; and a cloud broker service, monitoring tool, or metrics tool, for example. The operating system may be multi-user, multiprocessing, multitasking, and/or multithreading. The operating system may also perform basic tasks such as recognizing input from input devices, such as a keyboard or a keypad; sending output to a display; keeping track of files and directories on a medium; controlling peripheral devices, such as drives, printers, or image capture devices; and/or managing traffic on a bus. The network applications may include various components for establishing and maintaining network connections, such as machine readable instructions for implementing communication protocols including, but not limited to, TCP/IP, HTTP, Ethernet, USB, and FireWire.

[0048] In certain examples, some or all of the processes performed herein may be integrated into the operating system. In certain examples, the processes may be at least partially implemented in digital electronic circuitry, in computer hardware, in machine readable instructions (such as firmware and/or software), or in any combination thereof.

[0049] The above discussion is meant to be illustrative of the principles and various embodiments of the present disclosure. Numerous variations and modifications will become apparent to those skilled in the art once the above disclosure is fully appreciated. It is intended that the following claims be interpreted to embrace all such variations and modifications.



We claim:

1. A method of tuning a cloud service, comprising:
 - detecting an event and an event type on a cloud server application by monitoring at least one of a hardware value and a software value on a cloud server;
 - determining at least one exposed application parameter to tune based on the event type;
 - determining a tuning priority of the at least one application parameter based on historical tuning data;
 - automatically tuning the at least one application parameter by modifying a software variable on the cloud server; and
 - in the event that tuning the at least one application parameter resolves the event type, increasing the tuning priority of the at least one application parameter, and
 - in the event that tuning the at least one application parameter fails to resolve the event type, decreasing the priority of the at least one application parameter, rolling back the tuning of the at least one application parameter, and transmitting an instruction to invoke a default tune setting.
2. The method according to claim 1, wherein automatically tuning the at least one application parameter further comprises looping through an order of priority.
3. The method according to claim 1, wherein the event type comprises a performance degradation alert.
4. The method according to claim 1, wherein the event type comprises a threshold level alert.
5. The method according to claim 1, wherein the at least one application parameter comprises at least one of a database connection size, a caching value, a connection pool thread value, and a worker threads value.

6. The method according to claim 1, wherein automatically tuning the application parameter comprises tuning the application parameter incrementally.

7. The method according to claim 1, where modifying a software variable on the cloud server comprises at least one of increasing a variable, decreasing a variable, and toggling a variable.

8. The method according to claim 1, wherein the default tune setting comprises an autoscaling function.

9. A system for tuning a cloud service, the system comprising a processing resource in communication with a network interface and a computer readable medium, wherein the computer readable medium contains a set of instructions and wherein a processing unit is designed to carry out a set of instructions to:

detect, on the processing resource, a first monitoring event and an event type on a cloud server application;

determine at least one application parameter exposed in the cloud server application;

determine an order of priority for tuning the at least one application parameter based on the detected event type;

tune the application parameter;

detect, on the processing resource, a second monitoring event on the cloud server application; and

determine whether the second monitoring event indicates that the first monitoring event has been resolved, wherein

in the event that the first monitoring event is not resolved, the processing resource is to decrease the priority of the at least one application parameter, roll back the tuned application parameter to a previous state, and invoke an auto-scaling function.

10. The system according to claim 9, wherein the first monitoring event comprises at least one of a CPU alert, a memory alert, a disk drive alert, a bandwidth alert, and an application alert.

11. The system according to claim 9, wherein the previous state comprises a virtual machine image.

12. The system according to claim 9, wherein the instructions are to detect, on the processing resource, the second monitoring event on the cloud server application at a pre-set time after the application parameter is tuned.

13. A non-transitory computer readable storage medium on which is embedded a computer program, said computer program to deploy a tunable cloud services application, said computer program comprising a set of instructions to:

associate at least one tunable application parameter with at least one event type;

expose the at least one tunable application parameter to a production cloud server for auto-remediation on the production cloud server;

set an adjustable priority level for the at least one application parameter;

set a rollback point for the at least one application parameter;

set a monitor threshold for the at least one event type;

set a default tune action; and

deploy the tunable cloud services application to the production cloud server.

14. The computer readable storage medium according to claim 13, wherein the event type is based on an application response time.

15. The computer readable storage medium according to claim 13, wherein the default tune action comprises an instruction to increase the resources available to the production cloud server.




Figure 1




Figure 2




Figure 3

J.N.




Figure 4

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2015/038532**A. CLASSIFICATION OF SUBJECT MATTER****H04L 12/26(2006.01)i, H04L 12/24(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)

H04L 12/26; H04L 12/24; G06F 9/46; G06F 12/00; G06F 9/50

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: tuning parameter, application, priority, performance, cloud

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 2014-0143402 A1 (HUAWEI TECHNOLOGIES CO., LTD.) 22 May 2014 See paragraphs [0068]–[0105] and figure 1.	1–15
A	US 2009-0183162 A1 (CHARLES E. KINDEL et al.) 16 July 2009 See paragraphs [0035]–[0059], claim 1 and figure 1.	1–15
A	WO 2014-025584 A1 (ADVANCED MICRO DEVICES, INC.) 13 February 2014 See paragraphs [0095]–[0096], [00165]–[00166] and figures 1, 6.	1–15
A	EP 2669799 A1 (ALCATEL LUCENT) 04 December 2013 See paragraphs [0053]–[0059] and figures 5–6.	1–15
A	US 2006-0101224 A1 (PUNIT B. SHAH et al.) 11 May 2006 See paragraphs [0044]–[0046] and figures 4A–4B.	1–15



Further documents are listed in the continuation of Box C.



See patent family annex.

* Special categories of cited documents:

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

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

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

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

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

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

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

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

"&" document member of the same patent family

Date of the actual completion of the international search

16 September 2015 (16.09.2015)

Date of mailing of the international search report

13 October 2015 (13.10.2015)

Name and mailing address of the ISA/KR


International Application Division
Korean Intellectual Property Office
189 Cheongsu-ro, Seo-gu, Daejeon Metropolitan City, 35208,
Republic of Korea

Facsimile No. +82-42-472-7140

Authorized officer

KIM, Seong Woo

Telephone No. +82-42-481-3348



INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2015/038532

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2014-0143402 A1	22/05/2014	CN 102916906 A EP 2709321 A1 EP 2709321 A4 WO 2013-017001 A1	06/02/2013 19/03/2014 25/06/2014 07/02/2013
US 2009-0183162 A1	16/07/2009	US 8473956 B2	25/06/2013
WO 2014-025584 A1	13/02/2014	CN 104541247 A EP 2883140 A1 KR 10-2015-0043377 A US 2014-0047095 A1 US 2014-0047227 A1	22/04/2015 17/06/2015 22/04/2015 13/02/2014 13/02/2014
EP 2669799 A1	04/12/2013	None	
US 2006-0101224 A1	11/05/2006	US 2012-0110592 A1 US 2013-0042087 A1 US 8145872 B2 US 8285966 B2 US 8700876 B2	03/05/2012 14/02/2013 27/03/2012 09/10/2012 15/04/2014