




- (51) International Patent Classification:
G06F 17/30 (2006.01) *G06F 13/40* (2006.01)
- (21) International Application Number:
PCT/US2017/043182
- (22) International Filing Date:
20 July 2017 (20.07.2017)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
62/365,969 22 July 2016 (22.07.2016) US
62/376,859 18 August 2016 (18.08.2016) US
62/427,268 29 November 2016 (29.11.2016) US
15/654,615 19 July 2017 (19.07.2017) US
- (71) Applicant: INTEL CORPORATION [US/US]; 2200 Mission College Boulevard, Santa Clara, California 95054 (US).
- (72) Inventors: KUMAR, Mohan J.; 18680 SW Marko Lane, Aloha, Oregon 97007 (US). NACHIMUTHU, Murugasamy K.; 8441 SW 186th Avenue, Beaverton, Oregon 97007 (US). GORIUS, Aaron; 8 Riverbend Road, Upton, Massachusetts 01568 (US). ADILETTA, Matthew J.; 244 Sawyer Road, Bolton, Massachusetts 01740 (US). WILDE, Myles; 48 Russell Street, Charlestown, Massachusetts 02129 (US). CROCKER, Michael; c/o 10915 SW Fairhaven Way, Portland, Oregon 97223 (US). ZIAKAS, Dimitrios; Jones Farm 5, Hillsboro, Oregon 97124 (US).
- (74) Agent: RAYMOND, Andrew D.; Kacvinsky Daisak Bluni PLLC, 430 Davis Drive, Suite 150, Morrisville, North Carolina 27560 (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, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JO, JP, KE, KG, KH, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, LU, LY, MA, MD, ME,

(54) Title: AUTOMATED DATA CENTER MAINTENANCE

FIG. 1

Data Center 100



(57) Abstract: Techniques for automated data center maintenance are described. In an example embodiment, an automated maintenance device may comprise processing circuitry and non-transitory computer-readable storage media comprising instructions for execution by the processing circuitry to cause the automated maintenance device to receive an automation command from an automation coordinator for a data center, identify an automated maintenance procedure based on the received automation command, and perform the identified automated maintenance procedure. Other embodiments are described and claimed.



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).

Published:

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

AUTOMATED DATA CENTER MAINTENANCE**RELATED CASE**

This application claims priority to United States Patent Application Serial Number
5 15/654,615 filed July 19, 2017, United States Provisional Patent Application Number
62/365,969, filed July 22, 2016, United States Provisional Patent Application Number
62/376,859, filed August 18, 2016, and United States Provisional Patent Application Number
62/427,268, filed November 29, 2016, each of which is hereby incorporated by reference in its
entirety.

BACKGROUND

10 In the course of ordinary operation of a data center, various types of maintenance are
typically necessary in order to maintain desired levels of performance, stability, and reliability.
Examples of such maintenance include testing, repair, replacement, and/or reconfiguration of
components, installing new components, upgrading existing components, repositioning
15 components and equipment, and other tasks of such a nature. A large modern data center may
contain great numbers of components and equipment of various types, and as a result, may have
the potential to impose a fairly substantial maintenance burden.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 illustrates an embodiment of a first data center.

20 **FIG. 2** illustrates an embodiment of a logical configuration of a rack.

FIG. 3 illustrates an embodiment of a second data center.

FIG. 4 illustrates an embodiment of a third data center.

FIG. 5 illustrates an embodiment of a connectivity scheme.

FIG. 6 illustrates an embodiment of first rack architecture.

25 **FIG. 7** illustrates an embodiment of a first sled.

FIG. 8 illustrates an embodiment of a second rack architecture.

FIG. 9 illustrates an embodiment of a rack.

FIG. 10 illustrates an embodiment of a second sled.

FIG. 11 illustrates an embodiment of a fourth data center.

30 **FIG. 12** illustrates an embodiment of a first logic flow.

FIG. 13 illustrates an embodiment of a fifth data center.

FIG. 14 illustrates an embodiment of an automated maintenance device.

FIG. 15 illustrates an embodiment of a first operating environment.

FIG. 16 illustrates an embodiment of a second operating environment.

FIG. 17 illustrates an embodiment of a third operating environment.

FIG. 18 illustrates an embodiment of a fourth operating environment.

FIG. 19 illustrates an embodiment of a fifth operating environment.

FIG. 20 illustrates an embodiment of a sixth operating environment.

5 **FIG. 21** illustrates an embodiment of a first logic flow.

FIG. 22 illustrates an embodiment of a second logic flow.

FIG. 23 illustrates an embodiment of a third logic flow.

FIG. 24A illustrates an embodiment of a first storage medium.

FIG. 24B illustrates an embodiment of a second storage medium.

10 **FIG. 25** illustrates an embodiment of a computing architecture.

FIG. 26 illustrates an embodiment of a communications architecture.

FIG. 27 illustrates an embodiment of a communication device.

FIG. 28 illustrates an embodiment of a first wireless network.

FIG. 29 illustrates an embodiment of a second wireless network.

15 **DETAILED DESCRIPTION**

Various embodiments may be generally directed to techniques for automated data center maintenance. In one embodiment, for example, an automated maintenance device may comprise processing circuitry and non-transitory computer-readable storage media comprising instructions for execution by the processing circuitry to cause the automated maintenance device to receive
20 an automation command from an automation coordinator for a data center, identify an automated maintenance procedure based on the received automation command, and perform the identified automated maintenance procedure. Other embodiments are described and claimed.

Various embodiments may comprise one or more elements. An element may comprise any structure arranged to perform certain operations. Each element may be implemented as
25 hardware, software, or any combination thereof, as desired for a given set of design parameters or performance constraints. Although an embodiment may be described with a limited number of elements in a certain topology by way of example, the embodiment may include more or less elements in alternate topologies as desired for a given implementation. It is worthy to note that any reference to “one embodiment” or “an embodiment” means that a particular feature,
30 structure, or characteristic described in connection with the embodiment is included in at least one embodiment. The appearances of the phrases “in one embodiment,” “in some embodiments,” and “in various embodiments” in various places in the specification are not necessarily all referring to the same embodiment.

FIG. 1 illustrates a conceptual overview of a data center 100 that may generally be representative of a data center or other type of computing network in/for which one or more techniques described herein may be implemented according to various embodiments. As shown in FIG. 1, data center 100 may generally contain a plurality of racks, each of which may house computing equipment comprising a respective set of physical resources. In the particular non-limiting example depicted in FIG. 1, data center 100 contains four racks 102A to 102D, which house computing equipment comprising respective sets of physical resources (PCRs) 105A to 105D. According to this example, a collective set of physical resources 106 of data center 100 includes the various sets of physical resources 105A to 105D that are distributed among racks 102A to 102D. Physical resources 106 may include resources of multiple types, such as – for example – processors, co-processors, accelerators, field-programmable gate arrays (FPGAs), memory, and storage. The embodiments are not limited to these examples.

The illustrative data center 100 differs from typical data centers in many ways. For example, in the illustrative embodiment, the circuit boards (“sleds”) on which components such as CPUs, memory, and other components are placed are designed for increased thermal performance. In particular, in the illustrative embodiment, the sleds are shallower than typical boards. In other words, the sleds are shorter from the front to the back, where cooling fans are located. This decreases the length of the path that air must travel across the components on the board. Further, the components on the sled are spaced further apart than in typical circuit boards, and the components are arranged to reduce or eliminate shadowing (i.e., one component in the air flow path of another component). In the illustrative embodiment, processing components such as the processors are located on a top side of a sled while near memory, such as DIMMs, are located on a bottom side of the sled. As a result of the enhanced airflow provided by this design, the components may operate at higher frequencies and power levels than in typical systems, thereby increasing performance. Furthermore, the sleds are configured to blindly mate with power and data communication cables in each rack 102A, 102B, 102C, 102D, enhancing their ability to be quickly removed, upgraded, reinstalled, and/or replaced. Similarly, individual components located on the sleds, such as processors, accelerators, memory, and data storage drives, are configured to be easily upgraded due to their increased spacing from each other. In the illustrative embodiment, the components additionally include hardware attestation features to prove their authenticity.

Furthermore, in the illustrative embodiment, the data center 100 utilizes a single network architecture (“fabric”) that supports multiple other network architectures including Ethernet and Omni-Path. The sleds, in the illustrative embodiment, are coupled to switches via optical fibers,

which provide higher bandwidth and lower latency than typical twister pair cabling (e.g., Category 5, Category 5e, Category 6, etc.). Due to the high bandwidth, low latency interconnections and network architecture, the data center 100 may, in use, pool resources, such as memory, accelerators (e.g., graphics accelerators, FPGAs, ASICs, etc.), and data storage drives that are physically disaggregated, and provide them to compute resources (e.g., processors) on an as needed basis, enabling the compute resources to access the pooled resources as if they were local. The illustrative data center 100 additionally receives usage information for the various resources, predicts resource usage for different types of workloads based on past resource usage, and dynamically reallocates the resources based on this information.

The racks 102A, 102B, 102C, 102D of the data center 100 may include physical design features that facilitate the automation of a variety of types of maintenance tasks. For example, data center 100 may be implemented using racks that are designed to be robotically-accessed, and to accept and house robotically-manipulable resource sleds. Furthermore, in the illustrative embodiment, the racks 102A, 102B, 102C, 102D include integrated power sources that receive a greater voltage than is typical for power sources. The increased voltage enables the power sources to provide additional power to the components on each sled, enabling the components to operate at higher than typical frequencies. **FIG. 2** illustrates an exemplary logical configuration of a rack 202 of the data center 100. As shown in FIG. 2, rack 202 may generally house a plurality of sleds, each of which may comprise a respective set of physical resources. In the particular non-limiting example depicted in FIG. 2, rack 202 houses sleds 204-1 to 204-4 comprising respective sets of physical resources 205-1 to 205-4, each of which constitutes a portion of the collective set of physical resources 206 comprised in rack 202. With respect to FIG. 1, if rack 202 is representative of – for example – rack 102A, then physical resources 206 may correspond to the physical resources 105A comprised in rack 102A. In the context of this example, physical resources 105A may thus be made up of the respective sets of physical resources, including physical storage resources 205-1, physical accelerator resources 205-2, physical memory resources 205-3, and physical compute resources 205-5 comprised in the sleds 204-1 to 204-4 of rack 202. The embodiments are not limited to this example. Each sled may contain a pool of each of the various types of physical resources (e.g., compute, memory, accelerator, storage). By having robotically accessible and robotically manipulable sleds comprising disaggregated resources, each type of resource can be upgraded independently of each other and at their own optimized refresh rate.

FIG. 3 illustrates an example of a data center 300 that may generally be representative of one in/for which one or more techniques described herein may be implemented according to

various embodiments. In the particular non-limiting example depicted in FIG. 3, data center 300 comprises racks 302-1 to 302-32. In various embodiments, the racks of data center 300 may be arranged in such fashion as to define and/or accommodate various access pathways. For example, as shown in FIG. 3, the racks of data center 300 may be arranged in such fashion as to define and/or accommodate access pathways 311A, 311B, 311C, and 311D. In some embodiments, the presence of such access pathways may generally enable automated maintenance equipment, such as robotic maintenance equipment, to physically access the computing equipment housed in the various racks of data center 300 and perform automated maintenance tasks (e.g., replace a failed sled, upgrade a sled). In various embodiments, the dimensions of access pathways 311A, 311B, 311C, and 311D, the dimensions of racks 302-1 to 302-32, and/or one or more other aspects of the physical layout of data center 300 may be selected to facilitate such automated operations. The embodiments are not limited in this context.

FIG. 4 illustrates an example of a data center 400 that may generally be representative of one in/for which one or more techniques described herein may be implemented according to various embodiments. As shown in FIG. 4, data center 400 may feature an optical fabric 412. Optical fabric 412 may generally comprise a combination of optical signaling media (such as optical cabling) and optical switching infrastructure via which any particular sled in data center 400 can send signals to (and receive signals from) each of the other sleds in data center 400. The signaling connectivity that optical fabric 412 provides to any given sled may include connectivity both to other sleds in a same rack and sleds in other racks. In the particular non-limiting example depicted in FIG. 4, data center 400 includes four racks 402A to 402D. Racks 402A to 402D house respective pairs of sleds 404A-1 and 404A-2, 404B-1 and 404B-2, 404C-1 and 404C-2, and 404D-1 and 404D-2. Thus, in this example, data center 400 comprises a total of eight sleds. Via optical fabric 412, each such sled may possess signaling connectivity with each of the seven other sleds in data center 400. For example, via optical fabric 412, sled 404A-1 in rack 402A may possess signaling connectivity with sled 404A-2 in rack 402A, as well as the six other sleds 404B-1, 404B-2, 404C-1, 404C-2, 404D-1, and 404D-2 that are distributed among the other racks 402B, 402C, and 402D of data center 400. The embodiments are not limited to this example.

FIG. 5 illustrates an overview of a connectivity scheme 500 that may generally be representative of link-layer connectivity that may be established in some embodiments among the various sleds of a data center, such as any of example data centers 100, 300, and 400 of FIGs. 1, 3, and 4. Connectivity scheme 500 may be implemented using an optical fabric that features a

dual-mode optical switching infrastructure 514. Dual-mode optical switching infrastructure 514 may generally comprise a switching infrastructure that is capable of receiving communications according to multiple link-layer protocols via a same unified set of optical signaling media, and properly switching such communications. In various embodiments, dual-mode optical switching infrastructure 514 may be implemented using one or more dual-mode optical switches 515. In various embodiments, dual-mode optical switches 515 may generally comprise high-radix switches. In some embodiments, dual-mode optical switches 515 may comprise multi-ply switches, such as four-ply switches. In various embodiments, dual-mode optical switches 515 may feature integrated silicon photonics that enable them to switch communications with significantly reduced latency in comparison to conventional switching devices. In some embodiments, dual-mode optical switches 515 may constitute leaf switches 530 in a leaf-spine architecture additionally including one or more dual-mode optical spine switches 520.

In various embodiments, dual-mode optical switches may be capable of receiving both Ethernet protocol communications carrying Internet Protocol (IP packets) and communications according to a second, high-performance computing (HPC) link-layer protocol (e.g., Intel's Omni-Path Architecture's, Infiniband) via optical signaling media of an optical fabric. As reflected in FIG. 5, with respect to any particular pair of sleds 504A and 504B possessing optical signaling connectivity to the optical fabric, connectivity scheme 500 may thus provide support for link-layer connectivity via both Ethernet links and HPC links. Thus, both Ethernet and HPC communications can be supported by a single high-bandwidth, low-latency switch fabric. The embodiments are not limited to this example.

FIG. 6 illustrates a general overview of a rack architecture 600 that may be representative of an architecture of any particular one of the racks depicted in FIGs. 1 to 4 according to some embodiments. As reflected in FIG. 6, rack architecture 600 may generally feature a plurality of sled spaces into which sleds may be inserted, each of which may be robotically-accessible via a rack access region 601. In the particular non-limiting example depicted in FIG. 6, rack architecture 600 features five sled spaces 603-1 to 603-5. Sled spaces 603-1 to 603-5 feature respective multi-purpose connector modules (MPCMs) 616-1 to 616-5.

Included among the types of sleds to be accommodated by rack architecture 600 may be one or more types of sleds that feature expansion capabilities. **FIG. 7** illustrates an example of a sled 704 that may be representative of a sled of such a type. As shown in FIG. 7, sled 704 may comprise a set of physical resources 705, as well as an MPCM 716 designed to couple with a counterpart MPCM when sled 704 is inserted into a sled space such as any of sled spaces 603-1 to 603-5 of FIG. 6. Sled 704 may also feature an expansion connector 717. Expansion

connector 717 may generally comprise a socket, slot, or other type of connection element that is capable of accepting one or more types of expansion modules, such as an expansion sled 718. By coupling with a counterpart connector on expansion sled 718, expansion connector 717 may provide physical resources 705 with access to supplemental computing resources 705B residing on expansion sled 718. The embodiments are not limited in this context.

FIG. 8 illustrates an example of a rack architecture 800 that may be representative of a rack architecture that may be implemented in order to provide support for sleds featuring expansion capabilities, such as sled 704 of FIG. 7. In the particular non-limiting example depicted in FIG. 8, rack architecture 800 includes seven sled spaces 803-1 to 803-7, which feature respective MPCMs 816-1 to 816-7. Sled spaces 803-1 to 803-7 include respective primary regions 803-1A to 803-7A and respective expansion regions 803-1B to 803-7B. With respect to each such sled space, when the corresponding MPCM is coupled with a counterpart MPCM of an inserted sled, the primary region may generally constitute a region of the sled space that physically accommodates the inserted sled. The expansion region may generally constitute a region of the sled space that can physically accommodate an expansion module, such as expansion sled 718 of FIG. 7, in the event that the inserted sled is configured with such a module.

FIG. 9 illustrates an example of a rack 902 that may be representative of a rack implemented according to rack architecture 800 of FIG. 8 according to some embodiments. In the particular non-limiting example depicted in FIG. 9, rack 902 features seven sled spaces 903-1 to 903-7, which include respective primary regions 903-1A to 903-7A and respective expansion regions 903-1B to 903-7B. In various embodiments, temperature control in rack 902 may be implemented using an air cooling system. For example, as reflected in FIG. 9, rack 902 may feature a plurality of fans 919 that are generally arranged to provide air cooling within the various sled spaces 903-1 to 903-7. In some embodiments, the height of the sled space is greater than the conventional “1U” server height. In such embodiments, fans 919 may generally comprise relatively slow, large diameter cooling fans as compared to fans used in conventional rack configurations. Running larger diameter cooling fans at lower speeds may increase fan lifetime relative to smaller diameter cooling fans running at higher speeds while still providing the same amount of cooling. The sleds are physically shallower than conventional rack dimensions. Further, components are arranged on each sled to reduce thermal shadowing (i.e., not arranged serially in the direction of air flow). As a result, the wider, shallower sleds allow for an increase in device performance because the devices can be operated at a higher thermal

envelope (e.g., 250W) due to improved cooling (i.e., no thermal shadowing, more space between devices, more room for larger heat sinks, etc.).

MPCMs 916-1 to 916-7 may be configured to provide inserted sleds with access to power sourced by respective power modules 920-1 to 920-7, each of which may draw power from an external power source 921. In various embodiments, external power source 921 may deliver alternating current (AC) power to rack 902, and power modules 920-1 to 920-7 may be configured to convert such AC power to direct current (DC) power to be sourced to inserted sleds. In some embodiments, for example, power modules 920-1 to 920-7 may be configured to convert 277-volt AC power into 12-volt DC power for provision to inserted sleds via respective MPCMs 916-1 to 916-7. The embodiments are not limited to this example.

MPCMs 916-1 to 916-7 may also be arranged to provide inserted sleds with optical signaling connectivity to a dual-mode optical switching infrastructure 914, which may be the same as – or similar to – dual-mode optical switching infrastructure 514 of FIG. 5. In various embodiments, optical connectors contained in MPCMs 916-1 to 916-7 may be designed to couple with counterpart optical connectors contained in MPCMs of inserted sleds to provide such sleds with optical signaling connectivity to dual-mode optical switching infrastructure 914 via respective lengths of optical cabling 922-1 to 922-7. In some embodiments, each such length of optical cabling may extend from its corresponding MPCM to an optical interconnect loom 923 that is external to the sled spaces of rack 902. In various embodiments, optical interconnect loom 923 may be arranged to pass through a support post or other type of load-bearing element of rack 902. The embodiments are not limited in this context. Because inserted sleds connect to an optical switching infrastructure via MPCMs, the resources typically spent in manually configuring the rack cabling to accommodate a newly inserted sled can be saved.

FIG. 10 illustrates an example of a sled 1004 that may be representative of a sled designed for use in conjunction with rack 902 of FIG. 9 according to some embodiments. Sled 1004 may feature an MPCM 1016 that comprises an optical connector 1016A and a power connector 1016B, and that is designed to couple with a counterpart MPCM of a sled space in conjunction with insertion of MPCM 1016 into that sled space. Coupling MPCM 1016 with such a counterpart MPCM may cause power connector 1016 to couple with a power connector comprised in the counterpart MPCM. This may generally enable physical resources 1005 of sled 1004 to source power from an external source, via power connector 1016 and power transmission media 1024 that conductively couples power connector 1016 to physical resources 1005.

Sled 1004 may also include dual-mode optical network interface circuitry 1026. Dual-mode optical network interface circuitry 1026 may generally comprise circuitry that is capable of communicating over optical signaling media according to each of multiple link-layer protocols supported by dual-mode optical switching infrastructure 914 of FIG. 9. In some embodiments, dual-mode optical network interface circuitry 1026 may be capable both of Ethernet protocol communications and of communications according to a second, high-performance protocol. In various embodiments, dual-mode optical network interface circuitry 1026 may include one or more optical transceiver modules 1027, each of which may be capable of transmitting and receiving optical signals over each of one or more optical channels. The embodiments are not limited in this context.

Coupling MPCM 1016 with a counterpart MPCM of a sled space in a given rack may cause optical connector 1016A to couple with an optical connector comprised in the counterpart MPCM. This may generally establish optical connectivity between optical cabling of the sled and dual-mode optical network interface circuitry 1026, via each of a set of optical channels 1025. Dual-mode optical network interface circuitry 1026 may communicate with the physical resources 1005 of sled 1004 via electrical signaling media 1028. In addition to the dimensions of the sleds and arrangement of components on the sleds to provide improved cooling and enable operation at a relatively higher thermal envelope (e.g., 250W), as described above with reference to FIG. 9, in some embodiments, a sled may include one or more additional features to facilitate air cooling, such as a heatpipe and/or heat sinks arranged to dissipate heat generated by physical resources 1005. It is worthy of note that although the example sled 1004 depicted in FIG. 10 does not feature an expansion connector, any given sled that features the design elements of sled 1004 may also feature an expansion connector according to some embodiments. The embodiments are not limited in this context.

FIG. 11 illustrates an example of a data center 1100 that may generally be representative of one in/for which one or more techniques described herein may be implemented according to various embodiments. As reflected in FIG. 11, a physical infrastructure management framework 1150A may be implemented to facilitate management of a physical infrastructure 1100A of data center 1100. In various embodiments, one function of physical infrastructure management framework 1150A may be to manage automated maintenance functions within data center 1100, such as the use of robotic maintenance equipment to service computing equipment within physical infrastructure 1100A. In some embodiments, physical infrastructure 1100A may feature an advanced telemetry system that performs telemetry reporting that is sufficiently robust to support remote automated management of physical infrastructure 1100A. In various

embodiments, telemetry information provided by such an advanced telemetry system may support features such as failure prediction/prevention capabilities and capacity planning capabilities. In some embodiments, physical infrastructure management framework 1150A may also be configured to manage authentication of physical infrastructure components using
5 hardware attestation techniques. For example, robots may verify the authenticity of components before installation by analyzing information collected from a radio frequency identification (RFID) tag associated with each component to be installed. The embodiments are not limited in this context.

As shown in FIG. 11, the physical infrastructure 1100A of data center 1100 may comprise
10 an optical fabric 1112, which may include a dual-mode optical switching infrastructure 1114. Optical fabric 1112 and dual-mode optical switching infrastructure 1114 may be the same as – or similar to – optical fabric 412 of FIG. 4 and dual-mode optical switching infrastructure 514 of FIG. 5, respectively, and may provide high-bandwidth, low-latency, multi-protocol connectivity among sleds of data center 1100. As discussed above, with reference to FIG. 1, in various
15 embodiments, the availability of such connectivity may make it feasible to disaggregate and dynamically pool resources such as accelerators, memory, and storage. In some embodiments, for example, one or more pooled accelerator sleds 1130 may be included among the physical infrastructure 1100A of data center 1100, each of which may comprise a pool of accelerator resources – such as co-processors and/or FPGAs, for example – that is available globally
20 accessible to other sleds via optical fabric 1112 and dual-mode optical switching infrastructure 1114.

In another example, in various embodiments, one or more pooled storage sleds 1132 may be included among the physical infrastructure 1100A of data center 1100, each of which may comprise a pool of storage resources that is available globally accessible to other sleds via
25 optical fabric 1112 and dual-mode optical switching infrastructure 1114. In some embodiments, such pooled storage sleds 1132 may comprise pools of solid-state storage devices such as solid-state drives (SSDs). In various embodiments, one or more high-performance processing sleds 1134 may be included among the physical infrastructure 1100A of data center 1100. In some embodiments, high-performance processing sleds 1134 may comprise pools of high-performance
30 processors, as well as cooling features that enhance air cooling to yield a higher thermal envelope of up to 250W or more. In various embodiments, any given high-performance processing sled 1134 may feature an expansion connector 1117 that can accept a far memory expansion sled, such that the far memory that is locally available to that high-performance processing sled 1134 is disaggregated from the processors and near memory comprised on that

sled. In some embodiments, such a high-performance processing sled 1134 may be configured with far memory using an expansion sled that comprises low-latency SSD storage. The optical infrastructure allows for compute resources on one sled to utilize remote accelerator/FPGA, memory, and/or SSD resources that are disaggregated on a sled located on the same rack or any other rack in the data center. The remote resources can be located one switch jump away or two-switch jumps away in the spine-leaf network architecture described above with reference to FIG. 5. The embodiments are not limited in this context.

In various embodiments, one or more layers of abstraction may be applied to the physical resources of physical infrastructure 1100A in order to define a virtual infrastructure, such as a software-defined infrastructure 1100B. In some embodiments, virtual computing resources 1136 of software-defined infrastructure 1100B may be allocated to support the provision of cloud services 1140. In various embodiments, particular sets of virtual computing resources 1136 may be grouped for provision to cloud services 1140 in the form of SDI services 1138. Examples of cloud services 1140 may include – without limitation – software as a service (SaaS) services 1142, platform as a service (PaaS) services 1144, and infrastructure as a service (IaaS) services 1146.

In some embodiments, management of software-defined infrastructure 1100B may be conducted using a virtual infrastructure management framework 1150B. In various embodiments, virtual infrastructure management framework 1150B may be designed to implement workload fingerprinting techniques and/or machine-learning techniques in conjunction with managing allocation of virtual computing resources 1136 and/or SDI services 1138 to cloud services 1140. In some embodiments, virtual infrastructure management framework 1150B may use/consult telemetry data in conjunction with performing such resource allocation. In various embodiments, an application/service management framework 1150C may be implemented in order to provide QoS management capabilities for cloud services 1140. The embodiments are not limited in this context.

FIG. 12 illustrates an example of a logic flow 1200 that may be representative of a maintenance algorithm for a data center, such as one or more of data center 100 of FIG. 1, data center 300 of FIG. 3, data center 400 of FIG. 4, and data center 1100 of FIG. 11. As shown in FIG. 12, data center operation information may be collected at 1202. In various embodiments, the collected data center operation information may include information describing various characteristics of ongoing operation of the data center, such as resource utilization levels, workload sizes, throughput rates, temperature measurements, and so forth. In some embodiments, the collected data center operation information may additionally or alternatively

include information describing other characteristics of the data center, such as the types of resources comprised in the data center, the locations/distributions of such resources within the data center, the capabilities and/or features of those resources, and so forth. The embodiments are not limited to these examples.

5 Based on data center operation information such as may be collected at 1202, a maintenance task to be completed may be identified at 1204. In one example, based on data center operation information indicating that processing resources on a given sled are non-responsive to communications from resources on other sleds, it may be determined at 1204 that the sled is to be pulled for testing. In another example, based on data center operation
10 information indicating that a particular DIMM has reached the end of its estimated service life, it may be determined that the DIMM is to be replaced. At 1206, a set of physical actions associated with the maintenance task may be determined, and those physical actions may be performed at 1208 in order to complete the maintenance task. For instance, in the
15 aforementioned example in which it is determined at 1204 that a DIMM is to be replaced, the physical actions identified at 1206 and performed at 1208 may include traveling to a particular rack in order to access a sled comprising the DIMM, removing the DIMM from a socket on the sled, and inserting a replacement DIMM into the socket. The embodiments are not limited to this example.

FIG. 13 illustrates an overhead view of an example data center 1300. According to
20 various embodiments, data center 1300 may be representative of a data center in which various operations associated with data center maintenance – such as operations associated with one or more of blocks 1202, 1204, 1206, and 1208 in logic flow 1200 of FIG. 12 – are automated using the capabilities of robotic maintenance equipment. According to some embodiments, data center 1300 may be representative of one or more of data center 100 of FIG. 1, data center 300 of FIG.
25 3, data center 400 of FIG. 4, and data center 1100 of FIG. 11. The embodiments are not limited in this context.

 In various embodiments, according to an automated maintenance scheme implemented in data center 1300, robots 1360 may be used to service, repair, replace, clean, test, configure, upgrade, move, position, and/or otherwise manipulate equipment housed in racks 1302. Racks
30 1302 may be arranged in such fashion as to define and/or accommodate access pathways via which robots 1360 can physically access such equipment. Robots 1360 may traverse such access pathways in conjunction with moving around in data center 1300 to perform various tasks. Physical features of equipment housed in racks 1302 may be designed to facilitate robotic manipulation/handling. It is to be appreciated that in various embodiments, the equipment

housed in racks 1302 may include some equipment that is not robotically accessible/serviceable. Further, in some embodiments, there may be some equipment within data center 1300 that is robotically accessible/serviceable but is not housed in racks 1302. The embodiments are not limited in this context.

5 **FIG. 14** illustrates a block diagram of an automated maintenance device 1400 that may be representative of any given robot 1360 in data center 1300 of FIG. 13 according to various embodiments. As shown in FIG. 14, automated maintenance device 1400 may comprise a variety of elements. In the non-limiting example depicted in FIG. 14, automated maintenance
10 device 1400 comprises locomotion elements 1462, manipulation elements 1463, sensory elements 1464, communication elements 1465, interfaces 1466, memory/storage elements 1467, and operations management and control (OMC) elements 1468.

Locomotion elements 1462 may generally comprise physical elements enabling automated maintenance device 1400 to move around within a data center. In various embodiments, locomotion elements 1462 may comprise wheels. In some embodiments, locomotion elements
15 1462 may comprise caterpillar tracks. In various embodiments, automated maintenance device 1400 may provide the motive power/force required for motion. For example, in some embodiments, automated maintenance device 1400 may feature a battery that provides power to drive wheels or tracks used by automated maintenance device 1400 for moving around in a data center. In various other embodiments, the motive power/force may be provided by an external
20 source. The embodiments are not limited in this context.

Manipulation elements 1463 may generally comprise physical elements that are usable to manipulate various types of equipment in a data center. In some embodiments, manipulation elements 1463 may include one or more robotic arms. In various embodiments, manipulation elements 1463 may include one or more multi-link manipulators. In some embodiments,
25 manipulation elements 1463 may include one or more end effectors usable for gripping various types of equipment, components, and/or other objects within the data center. In various embodiments, manipulation elements 1463 may include one or more end effectors comprising impactful grippers, such as jaw or claw grippers. In some embodiments, manipulation elements 1463 may include one or more end effectors comprising ingressive grippers, which may feature
30 pins, needles, hackles, or other elements that are to physically penetrate the surface of an object being gripped. In various embodiments, manipulation elements 1463 may include one or more end effectors comprising astrictive grippers, which may grip objects using air suction, magnetic adhesion, or electroadhesion. The embodiments are not limited to these examples.

Sensory elements 1464 may generally comprise physical elements that are usable to sense various aspects of ambient conditions within a data center. Examples of sensory elements 1464 may include cameras, alignment guides/sensors, distance sensors, proximity sensors, barcode readers, RFID/NFC readers, temperature sensors, airflow sensors, air quality sensors, humidity
5 sensors, and pressure sensors. The embodiments are not limited to these examples.

Communication elements 1465 may generally comprise a set of electronic components and/or circuitry operable to perform functions associated with communications between automated maintenance device 1400 and one or more external devices. In a given embodiment, such communications may include wireless communications, wired communications, or both. In
10 various embodiments, communication elements 1465 may include elements operative to generate/construct packets, frames, messages, and/or other information to be wirelessly communicated to external device(s), and/or to process/deconstruct packets, frames, messages, and/or other information wirelessly received from external device(s). In various embodiments, for example, communication elements 1465 may include baseband circuitry supporting wireless
15 communications according to one or more wireless communication protocols/standards. In some embodiments, communication elements 1465 may include elements operative to generate, process, construct, and/or deconstruct packets, frames, messages, and/or other information communicated over wired media. In various embodiments, for example, communication elements 1465 may include network interface circuitry supporting wired communications
20 according to one or more wired communication protocols/standards. The embodiments are not limited in this context.

In various embodiments, interfaces 1466 may include one or more communication interfaces 1466A. As reflected in FIG. 14, examples of interfaces 1466 that automated maintenance device 1400 may feature in various embodiments may include – without limitation
25 – communication interfaces 1466A, testing interfaces 1466B, power interfaces 1466C, and user interfaces 1466D.

Communication interfaces 1466A may generally comprise interfaces usable to transmit and/or receive signals via one or more communication media, which may include wired media, wireless media, or both. In various embodiments, communication interfaces 1466A may include
30 one or more wireless communication interfaces, such as radio frequency (RF) interfaces and/or optical wireless communication (OWC) interfaces. In some embodiments, communication interfaces may additionally or alternatively include one or more wired communication interfaces, such as interface(s) for communicating over media such as coaxial cable, twisted pair, and optical fiber. The embodiments are not limited to these examples.

In various embodiments, interfaces 1466 may include one or more testing interfaces 1466B. Testing interfaces 1466B may generally comprise interfaces via which automated maintenance device 1400 is able to test physical components/resources of one or more types, which may include – without limitation – one or more of physical storage resources 205-1, physical accelerator resources 205-2, physical memory resources 205-3, and physical compute resources 205-4 of FIG. 2. In an example embodiment, interfaces 1466 may include a testing interface 1466B that enables automated maintenance device 1400 to test the functionality of a DIMM inserted into a testing slot. The embodiments are not limited to these examples.

In various embodiments, interfaces 1466 may include one or more power interfaces 1466C. Power interfaces 1466C may generally comprise interfaces via which automated maintenance device 1400 can draw and/or source power. In various embodiments, power interfaces 1466C may include one or more interfaces via which automated maintenance device 1400 can draw power from external source(s). In some embodiments, automated maintenance device 1400 may feature one or more power interfaces 1466C configured to provide charge to one or more batteries (not shown), and automated maintenance device may draw its operating power from those one or more batteries. In various embodiments, automated maintenance device 1400 may feature one or more power interfaces 1466C via which it can directly draw operating power. In various embodiments, automated maintenance device 1400 may feature one or more power interfaces 1466C via which it can source power to external devices. For example, in various embodiments, automated maintenance device 1400 may feature a power interface 1466C via which it can source power to charge a battery of a second automated maintenance device. The embodiments are not limited to this example.

In some embodiments, interfaces 1466 may include one or more user interfaces. User interfaces 1466D may generally comprise interfaces via which information can be provided to human technicians and/or user input can be accepted from human technicians. Examples of user interfaces 1466D may include displays, touchscreens, speakers, microphones, keypads, mice, trackballs, trackpads, joysticks, fingerprint readers, retinal scanners, buttons, switches, and the like. The embodiments are not limited to these examples.

Memory/storage elements 1467 may generally comprise a set of electronic components and/or circuitry capable of retaining data, such as any of various types of data that may be generated, transmitted, received, and/or used by automated maintenance device 1400 during normal operation. In some embodiments, memory/storage elements 1467 may include one or both of volatile memory and non-volatile memory. For example, in various embodiments, memory/storage elements 1467 may include one or more of read-only memory (ROM), random-

access memory (RAM), dynamic RAM (DRAM), Double-Data-Rate DRAM (DDRAM), synchronous DRAM (SDRAM), static RAM (SRAM), programmable ROM (PROM), erasable programmable ROM (EPROM), electrically erasable programmable ROM (EEPROM), flash memory, polymer memory such as ferroelectric polymer memory, ovonic memory, phase change or ferroelectric memory, silicon-oxide-nitride-oxide-silicon (SONOS) memory, magnetic or optical cards, hard disks, an array of devices such as Redundant Array of Independent Disks (RAID) drives, solid state memory devices, solid state drives (SSDs), or any other type of media suitable for storing information. The embodiments are not limited to these examples.

OMC elements 1468 may generally comprise a set of components and/or circuitry capable of performing computing operations required to implement logic for managing and controlling the operations of automated maintenance device 1400. In various embodiments, OMC elements 1468 may include processing circuitry, such as one or more processors/processing units. In some embodiments, an automation engine 1469 may execute on such processing circuitry. Automation engine 1469 may generally be operative to conduct overall management, control, coordination, and/or oversight of the operations of automated maintenance device 1400. In various embodiments, this may include management, coordination, control, and/or oversight of the operations/usage of various other elements within automated maintenance device 1400, such as any or all of locomotion elements 1462, manipulation elements 1463, sensory elements 1464, communication elements 1465, interfaces 1466, and memory/storage elements 1467. The embodiments are not limited in this context.

FIG. 15 illustrates an example of an operating environment 1500 that may be representative of the implementation of an automated maintenance scheme in data center 1300 according to various embodiments. According to such an automated maintenance scheme, an automation coordinator 1555 may centrally manage/coordinate various aspects of automated maintenance operations in data center 1300. In some embodiments, automation coordinator 1555 may centrally manage/coordinate various aspects of automated maintenance operations in data center 1300 based in part on telemetry data 1571 provided by a telemetry framework 1570. According to various embodiments, telemetry framework 1570 may be representative of an advanced telemetry system that performs telemetry reporting for physical infrastructure 1100A in data center 1100 of FIG. 11, and automation coordinator 1555 may be representative of automated maintenance coordination functionality of physical infrastructure management framework 1150A. The embodiments are not limited in this context.

In some embodiments, management/coordination functionality of automation coordinator 1555 may be provided by a coordination engine 1572. In various embodiments, coordination

engine 1572 may execute on processing circuitry of automation coordinator 1555. In various embodiments, coordination engine 1572 may generate automation commands 1573 for transmission to robots 1360 in order to instruct robots 1360 to perform automated maintenance tasks and/or actions associated with such tasks. In some embodiments, robots 1360 may provide automation coordinator 1555 with various types of feedback 1574 in order to – for example – acknowledge automation commands 1573, report the results of attempted maintenance tasks, provide information regarding the statuses of components, resources, and/or equipment, provide information regarding information regarding the statuses of robots 1360 themselves, and/or report measurements of one or more aspects of ambient conditions in the data center. The
5
10
embodiments are not limited to these examples.

In some embodiments, coordination engine 1572 may consider various types of information in conjunction with automated maintenance coordination/management. As reflected in FIG. 15, examples of such types of information may include physical infrastructure information 1575, data center operations information 1576, maintenance task information 1577,
15
and maintenance equipment information 1579.

Physical infrastructure information 1575 may generally comprise information identifying equipment, devices, components, interconnects, physical resources, and/or other infrastructure elements that comprise portions of the physical infrastructure of data center 1300, and describing characteristics of such elements. Data center operations information 1576 may generally
20
comprise information describing various aspects of ongoing operations within data center 1300. In some embodiments, for example, data center operations information 1576 may include information describing one or more workloads currently being processed in data center 1300. In various embodiments, data center operations information 1576 may include metrics characterizing one or more aspects of current operations in data center 1300. For example, in
25
some embodiments, data center operations information 1576 may include performance metrics characterizing the relative level of performance currently being achieved in data center 1300, efficiency metrics characterizing the relative level of efficiency with which the physical resources of data center 1300 are being used to handle the current workloads, and utilization metrics generally indicative of current usage levels of various types of resources in data center
30
1300. In various embodiments, data center operations information 1576 may include telemetry data 1571, such as automation coordinator 1555 may receive via telemetry framework 1570 or from robots 1360. The embodiments are not limited in this context.

Maintenance task information 1577 may generally comprise information identifying and describing ongoing and pending maintenance tasks of data center 1300. Maintenance task

information 1577 may also include information identifying and describing previously completed maintenance tasks. In various embodiments, maintenance task information 1577 may include a pending task queue 1578. Pending task queue 1578 may generally comprise information identifying a set of maintenance tasks that need to be performed in data center 1300.

- 5 Maintenance equipment information 1579 may generally comprise identifying and describing automated maintenance equipment – such as robots 1360 – of data center 1300. In some embodiments, maintenance equipment information 1579 may include a candidate device pool 1580. Candidate device pool 1580 may generally comprise information identifying a set of robots 1360 that are currently available for use in data center 1300. The embodiments are not
10 limited in this context.

- In various embodiments, based on telemetry data 1571, automation coordinator 1555 may identify automated maintenance tasks to be performed in data center 1300 by robots 1360. For example, based on telemetry data 1571 indicating a high bit error rate at a DIMM, automation coordinator 1555 may determine that a robot 1360 should be assigned to replace that DIMM. In
15 some embodiments, automation coordinator 1555 may use telemetry data 1571 to prioritize among automated maintenance tasks, such as tasks comprised in pending task queue 1578. For example, automation coordinator 1555 may use telemetry data 1571 to assess the respective expected performance impacts of multiple automated maintenance tasks in pending task queue 1578, and may assign out an automated maintenance task with the highest expected performance
20 impact first. In some embodiments, in identifying and/or prioritizing among automated maintenance tasks, automation coordinator 1555 may consider any or all of physical infrastructure information 1575, data center operations information 1576, maintenance task information 1577, and maintenance equipment information 1579 in addition to – or in lieu of – telemetry data 1571.

- 25 In a first example, automation coordinator 1555 may assign a low priority to an automated maintenance task involving replacement of a malfunctioning compute sled based on physical infrastructure information 1575 indicating that another sled in a different rack can be used as a substitute without need for replacing the malfunctioning compute sled. In a second example, automation coordinator 1555 may assign a high priority to an automated maintenance task
30 involving replacing a malfunctioning memory sled based on data center operation information 1576 indicating that a scarcity of memory constitutes a performance bottleneck with respect to workloads being processed in data center 1300. In a third example, automation coordinator 1555 may determine not to add a new maintenance task to pending task queue 1578 based on a determination that a maintenance task already present in pending task queue 1578 may render

the new maintenance task unnecessary and/or moot. In a fourth example, in determining an extent to which to prioritize an automated maintenance task that requires the use of particular robots 1360 featuring specialized capabilities, automation coordinator 1555 may consider maintenance equipment information 1579 indicating whether any robots 1360 featuring such specialized capabilities are currently available. The embodiments are not limited to these examples.

In various embodiments, based on telemetry data 1571, automation coordinator 1555 may control the positioning and/or movement of robots 1360 within data center 1300. For example, having used telemetry data 1571 to identify a region of data center 1300 within which a greater number of hardware failures have been and/or are expected to be observed, automation coordinator 1555 may position robots 1360 more densely within that identified region than within other regions of data center 1300. The embodiments are not limited in this context.

In some embodiments, in response to automated maintenance decisions – such as may be reached based on any or all of telemetry data 1571, physical infrastructure information 1575, data center operations information 1576, maintenance task information 1577, and maintenance equipment information 1579 – automation coordinator 1555 may send automation commands 1573 to robots 1360 in order to instruct robots 1360 to perform operations associated with automated maintenance tasks. For example, upon determining that a particular compute sled should be replaced, automation coordinator 1555 may send an automation command 1573 in order to instruct a robot 1360 to perform a sled replacement procedure to replace the sled. In various embodiments, automation coordinator 1555 may inform robots 1360 of various parameters characterizing assigned automated maintenance tasks by including such parameters in automation commands 1573. For instance, in the context of the preceding example, the automation command 1573 may contain fields specifying a sled ID uniquely identifying the sled to be replaced and a rack ID and/or sled space ID identifying the location of that sled within the data center, as well as analogous parameters associated with the replacement sled. The embodiments are not limited to this example.

It is worthy of note that in various embodiments, with respect to some aspects of automated maintenance operations, decision-making may be handled in a distributed – rather than centralized – fashion. In such embodiments, robots 1360 may make some automated maintenance decisions autonomously. In some such embodiments, as illustrated in FIG. 15, robots 1360 may perform such autonomous decision-making based on telemetry data 1571 received from telemetry framework 1570. In an example embodiment, a robot 1360 may determine based on analysis of telemetry data 1571 that a particular CPU is malfunctioning, and

autonomously decide to replace that malfunctioning CPU. In various embodiments, some or all of the robots 1360 in data center 1300 may have access to any or all of physical infrastructure information 1575, data center operations information 1576, maintenance task information 1577, and maintenance equipment information 1579, and may consider such information as well in conjunction with autonomous decision-making. In various embodiments, distributed coordination functions may be implemented to enable some types of maintenance tasks to be completed via collaborative maintenance procedures involving cooperation between multiple robots. The embodiments are not limited in this context.

FIG. 16 illustrates an example of an operating environment 1600 that may be representative of various embodiments. In operating environment 1600, in conjunction with automated maintenance operations in data center 1300, robots 1360 may provide automation coordinator 1555 with feedback 1574 that includes one or more of position data 1681, assistance data 1682, and environmental data 1683. The embodiments are not limited to these examples. It is worthy of note that in some embodiments, although not depicted in FIG. 16, robots 1360 may gather various types of telemetry data 1571 in conjunction with automated maintenance operations and include such gathered telemetry data 1571 in the feedback 1574 provided to automation coordinator 1555. The embodiments are not limited in this context.

Position data 1681 may generally comprise data for use by automation coordinator 1555 to determine/track the positions and/or movements of robots 1360 within data center 1300. In some embodiments, position data 1681 may comprise data associated with an indoor positioning system. In some such embodiments, the indoor positioning system may be a radio-based system, such as a Wi-Fi-based or Bluetooth-based indoor positioning system. In some other embodiments, a non-radio based positioning system, such as a magnetic, optical, or inertial indoor positioning system may be used. In various embodiments, the indoor positioning system may be a hybrid system, such as one that combines two or more of radio-based, magnetic, optical, and inertial indoor positioning techniques. The embodiments are not limited in this context.

Assistance data 1682 may generally comprise data for use by automation coordinator 1555 to provide human maintenance personnel with information aiding them in the identification and/or performance of manual maintenance tasks. In various embodiments, a given robot 1360 may generate assistance data 1682 in response to identifying a maintenance issue that it cannot correct/resolve in an automated fashion. For instance, after identifying a component that needs to be replaced and determining that it cannot perform the replacement itself, a robot 1360 take a picture of the component and provide assistance data 1682 comprising that picture to automation

coordinator 1555. Automation coordinator 1555 may then cause the picture to be presented on a display for reference by human maintenance personnel in order to aid visual identification of the component to be replaced. The embodiments are not limited to this example.

5 In some embodiments, the performance and/or reliability of various types of hardware in data center 1300 may potentially be affected by one or more aspects of the ambient conditions within data center 1300, such as ambient temperature, pressure, humidity, and air quality. For example, a rate at which corrosion occurs on metallic contacts of components such as DIMMs may depend on the ambient temperature and humidity. In various embodiments, it may thus be desirable to monitor various types of environmental parameters at various locations during
10 ongoing operations of data center 1300.

In some embodiments, robots 1360 may be configured to support environmental condition monitoring by measuring one or more aspects of ambient conditions within the data center during ongoing operations and providing those collected measurements to automation coordinator 1555 in the form of environmental data 1683. In various embodiments, robots 1360
15 may collect environmental data 1683 using sensors or sensor arrays comprising sensory elements such as sensory elements 1464 of FIG. 14. Examples of conditions/parameters that robots 1360 may measure and report to automation coordinator 1555 in the form of environmental data 1683 may include – without limitation – temperature, pressure, humidity, and air quality. In some embodiments, in conjunction with providing environmental condition measurements in the form
20 of environmental data 1683, robots 1360 may also provide corresponding position data 1681 that indicates the locations at which the associated measurements were performed. The embodiments are not limited in this context.

In various embodiments, access to dynamic, continuous, and location-specific measurements of such parameters may enable a data center operator to predict failures,
25 dynamically configure systems for best performance, and dynamically move resources for data center optimization. In some embodiments, based on environmental data 1683 provided by robots 1360, a data center operator may be able to predict accelerated failure of parts versus standard factory specification and replace parts earlier (or move to lower priority tasks). In various embodiments, environmental data 1683 provided by robots 1360 may enable a data
30 center operator to initiate service tickets ahead of predicted failure timelines. For example, a cleaning of DIMM contacts may be initiated in order to avoid corrosion build-up to the level where failures start occurring. In some embodiments, environmental data 1683 provided by robots 1360 may enable a data center operator to continuously and dynamically configure servers based on, for example, altitude, pressure and other parameters that may be important to such

things as fan speeds and cooling configurations which in turn may affect performance of a server in a given environment and temperature. In various embodiments, environmental data 1683 provided by robots 1360 may enable a data center operator to detect and move data center resources automatically from zones/locations of the data center that may be affected by equipment failures or environment variations detected by the robot's sensors. For example, based on environmental data 1683 indicating an excessive temperature or air quality deterioration in a particular data center region, servers and/or other resources may be relocated from the affected region to a different region. The embodiments are not limited to these examples.

10 **FIG. 17** illustrates an example of an operating environment 1700 that may be representative of the implementation of an automated data center maintenance scheme according to some embodiments. In operating environment 1700, a robot 1760 may perform one or more automated maintenance tasks at a rack 1702. According to some embodiments, robot 1760 may be representative of a robot 1360 that performs operations associated with automated data center maintenance in data center 1300 of FIGs. 13, 15, and 16. In various embodiments, robot 1760 may be implemented using automated maintenance device 1400 of FIG. 14. In various embodiments, as reflected by the dashed line in FIG. 17, robot 1760 may move to a location of rack 1702 from another location in order to perform one or more automated maintenance tasks at rack 1702. In some embodiments, robot 1760 may perform one or more such tasks based on automation commands 1773 received from automation coordinator 1555. In various embodiments, robot 1760 may additionally or alternatively perform one or more such tasks autonomously, without intervention on the part of automation coordinator 1555. The embodiments are not limited in this context.

In some embodiments, robot 1760 may perform one or more automated maintenance tasks involving the installation and/or removal of sleds at racks of a data center such as data center 1300. In various embodiments, for example, robot 1760 may be operative to install a sled 1704 at rack 1702. In some embodiments, robot 1760 may install sled 1704 by inserting it into an available sled space of rack 1702. In various embodiments, in conjunction with inserting sled 1704, robot 1760 may grip particular physical elements designed to accommodate robotic manipulation/ handling. In some embodiments, robot 1760 may use image recognition and/or other location techniques to locate the elements to be gripped, and may insert sled 1704 while gripping those elements. In various embodiments, rather than installing sled 1704, robot 1760 may instead remove sled 1704 from rack 1702 and install a replacement sled 1704B. In some embodiments, robot 1760 may install replacement sled 1704B in a same sled space as was

occupied by sled 1704, once it has removed sled 1704. In various other embodiments, robot 1760 may install replacement sled 1704B in a different sled space, such that it does not need to remove sled 1704 before installing replacement sled 1704B. The embodiments are not limited in this context.

5 In some embodiments, robot 1760 may perform one or more automated maintenance tasks involving upkeep, repair, and/or replacement of particular components on sleds of a data center such as data center 1300. In various embodiments, robot 1760 may be used to power up a component 1706 in accordance with a scheme for periodically powering up components in the data center on a periodic basis in order to improve the reliability of such components. In some
10 embodiments, for example, storage and/or memory components may tend to malfunction when left idle for excessive periods of time, and thus robots may be used to power up such components according to a defined cycle. In such an embodiment, robot 1760 may be operative to power up an appropriate component 1706 by plugging that component 1706 into a powered interface/slot. The embodiments are not limited to this example.

15 In various embodiments, robot 1760 may be operative to manipulate a given component 1706 in accordance with a scheme for automated upkeep of pooled memory resources of a data center. According to such a scheme, robots may be used to assess/troubleshoot apparently malfunctioning memory resources such as DIMMs. In some embodiments, according to such a scheme, robot 1760 may identify a component 1706 comprising a memory resource such as a
20 DIMM, remove that component 1706 from a slot on sled 1704, and clean the component 1706. Robot 1760 may then test the component 1706 to determine whether the issue has been resolved, and may determine to pull sled 1704 for “back-room” servicing if it finds that the problem persists. In various embodiments, robot 1760 may test the component 1706 after reinserting it into its slot on sled 1704. In some other embodiments, robot 1760 may be configured with a
25 testing slot into which it can insert the component 1706 for the purpose of testing. The embodiments are not limited in this context.

FIG. 18 illustrates an example of an operating environment 1800 that may be representative of the implementation of an automated data center maintenance scheme according to some embodiments. In operating environment 1800, a robot 1860 may perform automated
30 CPU cache servicing for a sled 1804 at a rack 1802. According to some embodiments, robot 1860 may be representative of a robot 1360 that performs operations associated with automated data center maintenance in data center 1300 of FIGs. 13, 15, and 16. In various embodiments, robot 1860 may be implemented using automated maintenance device 1400 of FIG. 14. In some embodiments, as reflected by the dashed line in FIG. 18, robot 1860 may move to a location of

rack 1802 from another location in order to perform the automated CPU cache servicing for sled 1804. In various embodiments, robot 1860 may perform such automated CPU cache servicing based on automation commands 1873 received from automation coordinator 1555. In some other embodiments, robot 1860 may perform the automated CPU cache servicing autonomously, without intervention on the part of automation coordinator 1555. The embodiments are not limited in this context.

As shown in FIG. 18, sled 1804 may comprise components 1806 that include a CPU 1806A, cache memory 1806B for the CPU 1806A, and a heatsink 1806C for the CPU 1806A. In various embodiments, cache memory 1806B may underlie CPU 1806A, and CPU 1806A may underlie heatsink 1806C. In some embodiments, cache memory 1806B may comprise one or more cache memory modules. In various embodiments, the automated CPU cache servicing that robot 1860 performs in operating environment 1800 may involve replacing cache memory 1806B. For example, in some embodiments, cache memory 1806B may comprise one or more cache memory modules that robot 1860 removes from sled 1804 and replaces with one or more replacement cache modules. In various embodiments, the determination to perform automated CPU cache servicing and thus replace cache memory 1806B may be based on a determination that cache memory 1806B is not functioning properly or is outdated. For example, in some embodiments, automation coordinator 1555 may determine – based on telemetry data 1571 of FIG. 15 – that cache memory 1806B is not functioning, and may use robot 1860 to replace cache memory 1806B in response to that determination. The embodiments are not limited to this example.

In various embodiments, according to a procedure for automated CPU cache servicing, robot 1860 may remove CPU 1806A and heat sink 1806C from sled 1804 in order to gain physical access to cache memory 1806B. In some embodiments, robot 1860 may remove sled 1804 from rack 1802 prior to removing CPU 1806A and heat sink 1806C from sled 1804. In various other embodiments, robot 1860 may remove CPU 1806A and heat sink 1806C from sled 1804 while sled 1804 remains seated within a sled space of rack 1802. In some embodiments, robot 1860 may first remove heat sink 1806C, and then remove CPU 1806A. In various other embodiments, robot 1860 may remove both heat sink 1806C and CPU 1806A simultaneously and/or as a collective unit (i.e., without removing heat sink 1806C from CPU 1806A). In some embodiments, after replacing cache memory 1806B, robot 1860 may reinstall CPU 1806A and heat sink 1806C upon sled 1804, which it may then reinsert into a sled space of rack 1802 in embodiments in which it was previously removed. The embodiments are not limited in this context.

FIG. 19 illustrates an example of an operating environment 1900 that may be representative of the implementation of an automated data center maintenance scheme according to some embodiments. In operating environment 1900, a robot 1960 may perform automated storage and/or transfer of a compute state of a compute sled 1904 at a rack 1902. According to some embodiments, robot 1760 may be representative of a robot 1360 that performs operations associated with automated data center maintenance in data center 1300 of FIGs. 13, 15, and 16. In various embodiments, robot 1960 may be implemented using automated maintenance device 1400 of FIG. 14. In some embodiments, as reflected by the dashed line in FIG. 19, robot 1960 may move to a location of rack 1902 from another location in order to perform the automated storage and/or transfer of the compute state of compute sled 1904. In various embodiments, robot 1960 may perform such automated compute state storage and/or transfer based on automation commands 1973 received from automation coordinator 1555. In some other embodiments, robot 1960 may perform the automated compute state storage and/or transfer autonomously, without intervention on the part of automation coordinator 1555. The embodiments are not limited in this context.

As shown in FIG. 19, compute sled 1904 may comprise components 1906 that include one or more CPUs 1906A and a connector 1906B. In various embodiments, compute sled 1904 may comprise two CPUs 1906A. In some other embodiments, compute sled 1904 may comprise more than two CPUs 1906A, or only a single CPU 1906A. Connector 1906B may generally comprise a slot, socket, or other connective component designed to accept a memory daughter card for use to store a compute state of compute sled 1904. In various embodiments, compute sled 1904 may comprise two CPUs 1906A and connector 1906B may be located between those two CPUs 1906A. The embodiments are not limited in this context.

In some embodiments, according to a procedure for automated compute state storage and/or transfer, robot 1960 may insert a memory card 1918 into connector 1906B. In various embodiments, robot 1960 may remove compute sled 1904 from rack 1902 prior to inserting memory card 1918 into connector 1906B. In some other embodiments, robot 1960 may insert memory card 1918 into connector 1906B while compute sled 1904 remains seated within a sled space of rack 1902. In still other embodiments, memory card 1918 may be present and coupled with connector 1906B prior to initiation of the automated compute state storage and/or transfer procedure. In various embodiments, memory card 1918 may comprise a set of physical memory resources 1906C. In some embodiments, once memory card is inserted into/coupled with connector 1906B, a compute state 1984 of compute sled 1904 may be stored on memory card 1918 using one or more of the physical memory resources 1906C comprised thereon. In various

embodiments, compute state 1984 may include respective states of each CPU 1906A comprised on compute sled 1904. In some embodiments, compute state 1984 may also include states of one or more memory resources comprised on compute sled 1904. The embodiments are not limited in this context.

5 In various embodiments, robot 1960 may perform an automated compute state storage/transfer procedure in order to preserve the compute state of compute sled 1904 during upkeep/repair of compute sled 1904. In some such embodiments, once compute state 1984 is stored on memory card 1918, robot 1960 may remove memory card 1918 from connector 1906B, perform upkeep/repair of compute sled 1904, reinsert memory card 1918 into connector 1906B,
10 and then restore compute sled 1904 to the compute state 1984 stored on memory card 1918. For instance, in an example embodiment, robot 1960 may remove a CPU 1906A from a socket on compute sled 1904 and insert a replacement CPU into that socket, and then cause compute sled 1904 to be restored to the compute state 1984 stored on memory card 1918. In various other embodiments, robot 1960 may perform an automated compute state storage/transfer procedure in
15 order to replace compute sled 1904 with another compute sled. In some such embodiments, once compute state 1984 is stored on memory card 1918, robot 1960 may remove memory card 1918 from connector 1906B, insert memory card 1918 into a connector on a replacement compute sled, insert the replacement compute sled into a sled space of rack 1902 or another rack, and cause the replacement compute sled to realize the compute state 1984 stored on memory card
20 1918. The embodiments are not limited in this context.

FIG. 20 illustrates an example of an operating environment 2000. According to various embodiments, operating environment 2000 may be representative of the implementation of an automated data center maintenance scheme according to which some aspects of automated maintenance operations involve collaboration/cooperation between robots. In operating
25 environment 2000, in conjunction with performing a collaborative maintenance task, robots 2060A and 2060B may coordinate with each other by exchanging interdevice coordination information 2086A and 2086B via one or more communication links 2085. Communication links 2085 may comprise wireless communication links, wired communication links, or a combination of both. According to some embodiments, robots 2060A and 2060B may be
30 representative of robots 1360 that perform operations associated with automated data center maintenance in data center 1300 of FIGs. 13, 15, and 16. In various embodiments, one or both of robots 2060A and 2060B may be implemented using automated maintenance device 1400 of FIG. 14.

It is worthy of note that the absence of automation coordinator 1555 in FIG. 20 is not intended to indicate that no aspects of automated maintenance would/could be centrally coordinated in operating environment 2000. It is both possible and contemplated that in various embodiments, distributed coordination may be implemented for some aspects of automated maintenance in a data center in which other aspects of automated maintenance are centrally coordinated by an entity such as automation coordinator 1555. For example, in operating environment 2000, a central automation coordinator may determine the need for performance of the collaborative maintenance task, select robots 2060A and 2060B as the robots that are to perform the collaborative maintenance task, and send automation commands to cause robots 2060A and 2060B to initiate the collaborative maintenance task. Robots 2060A and 2060B may then coordinate directly with each other in conjunction with performing the physical actions necessary to complete the collaborative maintenance task. The embodiments are not limited to this example.

FIG. 21 illustrates an example of a logic flow 2100 that may be representative of the implementation of one or more of the disclosed techniques according to some embodiments. For example, logic flow 2100 may be representative of operations that automation coordinator 1555 may perform in any of operating environments 1500, 1600, 1700, 1800, 1900, and 2000 of FIGs. 15-20 according to various embodiments. As shown in FIG. 21, at 2102, a maintenance task that is to be performed in a data center may be identified. For example, in operating environment 1500 of FIG. 15, automation coordinator 1555 may identify a maintenance task that is to be performed in data center 1300.

At 2104, a determination may be made to initiate automated performance of the maintenance task. For example, having added an identified maintenance task to pending task queue 1578 in operating environment 1500 of FIG. 15, automation coordinator 1555 may determine at a subsequent point in time that that maintenance task constitutes the highest priority task in the pending task queue 1578 and thus that its performance should be initiated. In another example, rather than adding the identified maintenance task to pending task queue 1578, automation coordinator 1555 may determine to initiate performance of the maintenance task immediately after it is identified.

At 2106, an automated maintenance device to which to assign the maintenance task may be selected. For example, among one or more robots 1360 comprised in candidate device pool 1580 in operating environment 1500 of FIG. 15, automation coordinator 1555 may select a robot 1360 to which to assign an identified maintenance task. It is worthy of note that in some embodiments, the identified maintenance task may be handled by multiple robots according to a

collaborate maintenance procedure. In such cases, more than one automated maintenance device may be selected at 2106 as an assignee of the maintenance task. For example, in operating environment 1500 of FIG. 15, automation coordinator 1555 may select multiple robots 1360 among those comprised in candidate device pool 1580 that are to work together according to a collaborative maintenance procedure to complete a maintenance task.

At 2108, one or more automation commands may be sent to cause an automated maintenance device selected at 2106 to perform an automated maintenance procedure associated with the maintenance task. For example, in operating environment 1500 of FIG. 15, automation coordinator 1555 may send one or more automation commands 1573 to cause a robot 1360 to perform an automated maintenance procedure associated with a maintenance task to which that robot 1360 has been allocated. In some embodiments in which multiple automated maintenance devices are selected at 2106 as assignees of the same maintenance task, automation commands may be sent to multiple automated maintenance devices at 2108. For example, in operating environment 1500 of FIG. 15, automation coordinator 1555 may send respective automation command(s) 1573 to multiple robots 1360 to cause those robots to perform a collaborative maintenance procedure associated with the maintenance task to be completed. The embodiments are not limited to these examples.

FIG. 22 illustrates an example of a logic flow 2200 that may be representative of the implementation of one or more of the disclosed techniques according to some embodiments. For example, logic flow 2200 may be representative of operations that may be performed in various embodiments by a robot such as a robot 1360 in one or both of operating environments 1500 and 1600 of FIGs. 15 and 16 and/or any of robots 1760, 1860, 1960, 2060A, and 2060B in operating environments 1700, 1800, 1900, and 2000 of FIGs. 17-20. As shown in FIG. 22, one or more automation commands may be received from an automation coordinator of a data center at 2202. For example, in operating environment 1500 of FIG. 15, a robot 1360 may receive one or more automation commands 1573 from automation coordinator 1555.

At 2204, an automated maintenance procedure may be identified based on the one or more automation commands received at 2202. For example, based on one or more automation commands 1573 received from automation coordinator 1555 in operating environment 1500 of FIG. 15, a robot 1360 may identify an automated maintenance procedure that it is to perform. The automated maintenance procedure identified at 2204 may then be performed at 2206. In various embodiments, the identification of the automated maintenance procedure at 2204 may be based on a maintenance task code that is comprised in at least one of the received automation commands, and is defined to correspond to a particular automated maintenance procedure. For

example, based on a maintenance task code comprised in an automation command 1573 received from automation coordinator 1555, a robot 1360 in operating environment 1500 of FIG. 15 may identify an automated DIMM testing procedure as an automated maintenance procedure to be performed. In various embodiments, the one or more automation commands received at 2202
5 may collectively contain one or more maintenance task parameters specifying particular details of the automated maintenance task, and such details may also be identified at 2204. For instance, in the context of the preceding example, the robot 1360 may identify – based on maintenance task parameters comprised in one or more automation commands 1573 received from automation coordinator 1555 – details such as a physical resource ID of a DIMM to be
10 tested, an identity and location of a sled on which that DIMM resides, and an identity of a particular DIMM slot on that sled that currently houses the DIMM. The embodiments are not limited to these examples.

FIG. 23 illustrates an example of a logic flow 2300 that may be representative of the implementation of one or more of the disclosed techniques according to some embodiments. For
15 example, logic flow 2300 may be representative of operations that may be performed by robot 2060A or robot 2060B in operating environment 2000 of FIG. 20. As shown in FIG. 23, a collaborative maintenance procedure that is to be performed in a data center may be identified at an automated maintenance device at 2302. For example, in operating environment 2000 of FIG. 20, robot 2060A may determine that a collaborative CPU replacement procedure is to be
20 performed. In some embodiments, the identification of the collaborative maintenance procedure at 2302 may be based on one or more automation commands received by the automated maintenance device from a centralized automation coordinator such as automation coordinator 1555. In various other embodiments, the identification of the collaborative maintenance procedure at 2302 may be performed autonomously. For example, in operating environment
25 1500 of FIG. 15, a robot 1360 may determine based on analysis of telemetry data 1571 that a particular CPU is malfunctioning, and may then identify a collaborative maintenance procedure to be performed in order to replace that malfunctioning CPU. The embodiments are not limited to this example.

A second automated maintenance device with which to collaborate during performance of
30 the collaborative maintenance procedure may be identified at 2304, and interdevice coordination information may be sent to the second automated maintenance device at 2306 in order to initiate the collaborative maintenance procedure. For example, in operating environment 2000 of FIG. 20, robot 2060A may determine that it is to collaborate with robot 2060B in conjunction with a collaborative CPU replacement procedure, and may send interdevice coordination information

2086A to robot 2086B in order to initiate that collaborative CPU replacement procedure. In some embodiments, the identification of the second automated maintenance device may be based on information received from a centralized automation coordinator such as automation coordinator 1555. For example, in some embodiments, a centralized automation coordinator
5 may be responsible for selecting the particular robots that are to work together to perform the collaborative maintenance procedure, and the identity of the second automated maintenance device may be indicated by a parameter comprised in an automation command received from the centralized automation coordinator. In other embodiments, the identification performed at 2304 may correspond to an autonomous selection of the second automated maintenance device. For
10 example, in operating environment 1500 of FIG. 15, a first robot 1360 may select a second robot 1360 that is comprised among those in candidate device pool 1580 as the second automated maintenance device that is to participate in the collaborative maintenance procedure. The embodiments are not limited to these examples.

FIG. 24A illustrates an embodiment of a storage medium 2400. Storage medium 2400
15 may comprise any computer-readable storage medium or machine-readable storage medium, such as an optical, magnetic or semiconductor storage medium. In some embodiments, storage medium 2400 may comprise a non-transitory storage medium. In various embodiments, storage medium 2400 may comprise an article of manufacture. In some embodiments, storage medium 2400 may store computer-executable instructions, such as computer-executable instructions to
20 implement logic flow 2100 of FIG. 21. Examples of a computer-readable storage medium or machine-readable storage medium may include any tangible media capable of storing electronic data, including volatile memory or non-volatile memory, removable or non-removable memory, erasable or non-erasable memory, writeable or re-writeable memory, and so forth. Examples of computer-executable instructions may include any suitable type of code, such as source code,
25 compiled code, interpreted code, executable code, static code, dynamic code, object-oriented code, visual code, and the like. The embodiments are not limited to these examples.

FIG. 24B illustrates an embodiment of a storage medium 2450. Storage medium 2450 may comprise any computer-readable storage medium or machine-readable storage medium, such as an optical, magnetic or semiconductor storage medium. In some embodiments, storage
30 medium 2450 may comprise a non-transitory storage medium. In various embodiments, storage medium 2450 may comprise an article of manufacture. According to some embodiments, storage medium 2450 may be representative of a memory/storage element 1467 comprised in automated maintenance device 1400 of FIG. 14. In some embodiments, storage medium 2450 may store computer-executable instructions, such as computer-executable instructions to

implement one or both of logic flow 2200 of FIG. 22 and logic flow 2300 of FIG. 23. Examples of a computer-readable storage medium or machine-readable storage medium and of computer-executable instructions may include any of the respective examples identified above in reference to storage medium 2400 of FIG. 24A. The embodiments are not limited to these examples.

5 **FIG. 25** illustrates an embodiment of an exemplary computing architecture 2500 that may be suitable for implementing various embodiments as previously described. In various embodiments, the computing architecture 2500 may comprise or be implemented as part of an electronic device. In some embodiments, the computing architecture 2500 may be representative, for example, of a computing device suitable for use in conjunction with
10 implementation of one or more of robots 1360, 1760, 1860, 1960, 2060A, and 2060B, automated maintenance device 1400, automation coordinator 1555, and logic flows 2100, 2200, and 2300. The embodiments are not limited in this context.

As used in this application, the terms “system” and “component” and “module” are intended to refer to a computer-related entity, either hardware, a combination of hardware and
15 software, software, or software in execution, examples of which are provided by the exemplary computing architecture 2500. For example, a component can be, but is not limited to being, a process running on a processor, a processor, a hard disk drive, multiple storage drives (of optical and/or magnetic storage medium), an object, an executable, a thread of execution, a program, and/or a computer. By way of illustration, both an application running on a server and the server
20 can be a component. One or more components can reside within a process and/or thread of execution, and a component can be localized on one computer and/or distributed between two or more computers. Further, components may be communicatively coupled to each other by various types of communications media to coordinate operations. The coordination may involve the uni-directional or bi-directional exchange of information. For instance, the components may
25 communicate information in the form of signals communicated over the communications media. The information can be implemented as signals allocated to various signal lines. In such allocations, each message may be a signal. Further embodiments, however, may alternatively employ data messages. Such data messages may be sent across various connections. Exemplary connections include parallel interfaces, serial interfaces, and bus interfaces.

30 The computing architecture 2500 includes various common computing elements, such as one or more processors, multi-core processors, co-processors, memory units, chipsets, controllers, peripherals, interfaces, oscillators, timing devices, video cards, audio cards, multimedia input/output (I/O) components, power supplies, and so forth. The embodiments, however, are not limited to implementation by the computing architecture 2500.

As shown in FIG. 25, according to computing architecture 2500, a computer 2502 comprises a processing unit 2504, a system memory 2506 and a system bus 2508. In some embodiments, computer 2502 may comprise a server. In some embodiments, computer 2502 may comprise a client. The processing unit 2504 can be any of various commercially available processors, including without limitation an AMD® Athlon®, Duron® and Opteron® processors; ARM® application, embedded and secure processors; IBM® and Motorola® DragonBall® and PowerPC® processors; IBM and Sony® Cell processors; Intel® Celeron®, Core (2) Duo®, Itanium®, Pentium®, Xeon®, and XScale® processors; and similar processors. Dual microprocessors, multi-core processors, and other multi processor architectures may also be employed as the processing unit 2504.

The system bus 2508 provides an interface for system components including, but not limited to, the system memory 2506 to the processing unit 2504. The system bus 2508 can be any of several types of bus structure that may further interconnect to a memory bus (with or without a memory controller), a peripheral bus, and a local bus using any of a variety of commercially available bus architectures. Interface adapters may connect to the system bus 2508 via a slot architecture. Example slot architectures may include without limitation Accelerated Graphics Port (AGP), Card Bus, (Extended) Industry Standard Architecture ((E)ISA), Micro Channel Architecture (MCA), NuBus, Peripheral Component Interconnect (Extended) (PCI(X)), PCI Express, Personal Computer Memory Card International Association (PCMCIA), and the like.

The system memory 2506 may include various types of computer-readable storage media in the form of one or more higher speed memory units, such as read-only memory (ROM), random-access memory (RAM), dynamic RAM (DRAM), Double-Data-Rate DRAM (DDRAM), synchronous DRAM (SDRAM), static RAM (SRAM), programmable ROM (PROM), erasable programmable ROM (EPROM), electrically erasable programmable ROM (EEPROM), flash memory, polymer memory such as ferroelectric polymer memory, ovonic memory, phase change or ferroelectric memory, silicon-oxide-nitride-oxide-silicon (SONOS) memory, magnetic or optical cards, an array of devices such as Redundant Array of Independent Disks (RAID) drives, solid state memory devices (e.g., USB memory, solid state drives (SSD) and any other type of storage media suitable for storing information. In the illustrated embodiment shown in FIG. 25, the system memory 2506 can include non-volatile memory 2510 and/or volatile memory 2512. A basic input/output system (BIOS) can be stored in the non-volatile memory 2510.

The computer 2502 may include various types of computer-readable storage media in the form of one or more lower speed memory units, including an internal (or external) hard disk drive (HDD) 2514, a magnetic floppy disk drive (FDD) 2516 to read from or write to a removable magnetic disk 2518, and an optical disk drive 2520 to read from or write to a removable optical disk 2522 (e.g., a CD-ROM or DVD). The HDD 2514, FDD 2516 and optical disk drive 2520 can be connected to the system bus 2508 by a HDD interface 2524, an FDD interface 2526 and an optical drive interface 2528, respectively. The HDD interface 2524 for external drive implementations can include at least one or both of Universal Serial Bus (USB) and IEEE 1394 interface technologies.

The drives and associated computer-readable media provide volatile and/or nonvolatile storage of data, data structures, computer-executable instructions, and so forth. For example, a number of program modules can be stored in the drives and memory units 2510, 2512, including an operating system 2530, one or more application programs 2532, other program modules 2534, and program data 2536.

A user can enter commands and information into the computer 2502 through one or more wire/wireless input devices, for example, a keyboard 2538 and a pointing device, such as a mouse 2540. Other input devices may include microphones, infra-red (IR) remote controls, radio-frequency (RF) remote controls, game pads, stylus pens, card readers, dongles, finger print readers, gloves, graphics tablets, joysticks, keyboards, retina readers, touch screens (e.g., capacitive, resistive, etc.), trackballs, trackpads, sensors, styluses, and the like. These and other input devices are often connected to the processing unit 2504 through an input device interface 2542 that is coupled to the system bus 2508, but can be connected by other interfaces such as a parallel port, IEEE 1394 serial port, a game port, a USB port, an IR interface, and so forth.

A monitor 2544 or other type of display device may also be connected to the system bus 2508 via an interface, such as a video adaptor 2546. The monitor 2544 may be internal or external to the computer 2502. In addition to the monitor 2544, a computer typically includes other peripheral output devices, such as speakers, printers, and so forth.

The computer 2502 may operate in a networked environment using logical connections via wire and/or wireless communications to one or more remote computers, such as a remote computer 2548. The remote computer 2548 can be a workstation, a server computer, a router, a personal computer, portable computer, microprocessor-based entertainment appliance, a peer device or other common network node, and typically includes many or all of the elements described relative to the computer 2502, although, for purposes of brevity, only a memory/storage device 2550 is illustrated. The logical connections depicted include

wire/wireless connectivity to a local area network (LAN) 2552 and/or larger networks, for example, a wide area network (WAN) 2554. Such LAN and WAN networking environments are commonplace in offices and companies, and facilitate enterprise-wide computer networks, such as intranets, all of which may connect to a global communications network, for example, the Internet.

When used in a LAN networking environment, the computer 2502 may be connected to the LAN 2552 through a wire and/or wireless communication network interface or adaptor 2556. The adaptor 2556 can facilitate wire and/or wireless communications to the LAN 2552, which may also include a wireless access point disposed thereon for communicating with the wireless functionality of the adaptor 2556.

When used in a WAN networking environment, the computer 2502 can include a modem 2558, or may be connected to a communications server on the WAN 2554, or has other means for establishing communications over the WAN 2554, such as by way of the Internet. The modem 2558, which can be internal or external and a wire and/or wireless device, connects to the system bus 2508 via the input device interface 2542. In a networked environment, program modules depicted relative to the computer 2502, or portions thereof, can be stored in the remote memory/storage device 2550. It will be appreciated that the network connections shown are exemplary and other means of establishing a communications link between the computers can be used.

The computer 2502 may be operable to communicate with wire and wireless devices or entities using the IEEE 802 family of standards, such as wireless devices operatively disposed in wireless communication (e.g., IEEE 802.16 over-the-air modulation techniques). This includes at least Wi-Fi (or Wireless Fidelity), WiMax, and Bluetooth™ wireless technologies, among others. Thus, the communication can be a predefined structure as with a conventional network or simply an ad hoc communication between at least two devices. Wi-Fi networks use radio technologies called IEEE 802.11x (a, b, g, n, etc.) to provide secure, reliable, fast wireless connectivity. A Wi-Fi network can be used to connect computers to each other, to the Internet, and to wire networks (which use IEEE 802.3-related media and functions).

FIG. 26 illustrates a block diagram of an exemplary communications architecture 2600 suitable for implementing various embodiments as previously described. The communications architecture 2600 includes various common communications elements, such as a transmitter, receiver, transceiver, radio, network interface, baseband processor, antenna, amplifiers, filters, power supplies, and so forth. The embodiments, however, are not limited to implementation by the communications architecture 2600.

As shown in FIG. 26, the communications architecture 2600 comprises includes one or more clients 2602 and servers 2604. The clients 2602 and the servers 2604 are operatively connected to one or more respective client data stores 2608 and server data stores 2610 that can be employed to store information local to the respective clients 2602 and servers 2604, such as cookies and/or associated contextual information. Any one of clients 2602 and/or servers 2604 may implement one or more of robots 1360, 1760, 1860, 1960, 2060A, and 2060B, automated maintenance device 1400, automation coordinator 1555, logic flows 2100, 2200, and 2300, and computing architecture 2500.

The clients 2602 and the servers 2604 may communicate information between each other using a communication framework 2606. The communications framework 2606 may implement any well-known communications techniques and protocols. The communications framework 2606 may be implemented as a packet-switched network (e.g., public networks such as the Internet, private networks such as an enterprise intranet, and so forth), a circuit-switched network (e.g., the public switched telephone network), or a combination of a packet-switched network and a circuit-switched network (with suitable gateways and translators).

The communications framework 2606 may implement various network interfaces arranged to accept, communicate, and connect to a communications network. A network interface may be regarded as a specialized form of an input output interface. Network interfaces may employ connection protocols including without limitation direct connect, Ethernet (e.g., thick, thin, twisted pair 10/100/1000 Base T, and the like), token ring, wireless network interfaces, cellular network interfaces, IEEE 802.11a-x network interfaces, IEEE 802.16 network interfaces, IEEE 802.20 network interfaces, and the like. Further, multiple network interfaces may be used to engage with various communications network types. For example, multiple network interfaces may be employed to allow for the communication over broadcast, multicast, and unicast networks. Should processing requirements dictate a greater amount speed and capacity, distributed network controller architectures may similarly be employed to pool, load balance, and otherwise increase the communicative bandwidth required by clients 2602 and the servers 2604. A communications network may be any one and the combination of wired and/or wireless networks including without limitation a direct interconnection, a secured custom connection, a private network (e.g., an enterprise intranet), a public network (e.g., the Internet), a Personal Area Network (PAN), a Local Area Network (LAN), a Metropolitan Area Network (MAN), an Operating Missions as Nodes on the Internet (OMNI), a Wide Area Network (WAN), a wireless network, a cellular network, and other communications networks.

As used herein, the term "circuitry" may refer to, be part of, or include an Application Specific Integrated Circuit (ASIC), an electronic circuit, a processor (shared, dedicated, or group), and/or memory (shared, dedicated, or group) that execute one or more software or firmware programs, a combinational logic circuit, and/or other suitable hardware components that provide the described functionality. In some embodiments, the circuitry may be implemented in, or functions associated with the circuitry may be implemented by, one or more software or firmware modules. In some embodiments, circuitry may include logic, at least partially operable in hardware. Embodiments described herein may be implemented into a system using any suitably configured hardware and/or software.

FIG. 27 illustrates an embodiment of a communication device 2700 that may implement one or more of robots 1360, 1760, 1860, 1960, 2060A, and 2060B, automated maintenance device 1400, automation coordinator 1555, logic flows 2100, 2200, and 2300, storage media 2400 and 2450, computing architecture 2500, clients 2602, and servers 2604. In various embodiments, device 2700 may comprise a logic circuit 2728. The logic circuit 2728 may include physical circuits to perform operations described for one or more of robots 1360, 1760, 1860, 1960, 2060A, and 2060B, automated maintenance device 1400, automation coordinator 1555, logic flows 2100, 2200, and 2300, computing architecture 2500, clients 2602, and servers 2604 for example. As shown in FIG. 27, device 2700 may include a radio interface 2710, baseband circuitry 2720, and computing platform 2730, although the embodiments are not limited to this configuration.

The device 2700 may implement some or all of the structure and/or operations for one or more of robots 1360, 1760, 1860, 1960, 2060A, and 2060B, automated maintenance device 1400, automation coordinator 1555, logic flows 2100, 2200, and 2300, storage media 2400 and 2450, computing architecture 2500, clients 2602, servers 2604, and logic circuit 2728 in a single computing entity, such as entirely within a single device. Alternatively, the device 2700 may distribute portions of the structure and/or operations for one or more of robots 1360, 1760, 1860, 1960, 2060A, and 2060B, automated maintenance device 1400, automation coordinator 1555, logic flows 2100, 2200, and 2300, storage media 2400 and 2450, computing architecture 2500, clients 2602, servers 2604, and logic circuit 2728 across multiple computing entities using a distributed system architecture, such as a client-server architecture, a 3-tier architecture, an N-tier architecture, a tightly-coupled or clustered architecture, a peer-to-peer architecture, a master-slave architecture, a shared database architecture, and other types of distributed systems. The embodiments are not limited in this context.

In one embodiment, radio interface 2710 may include a component or combination of components adapted for transmitting and/or receiving single-carrier or multi-carrier modulated signals (e.g., including complementary code keying (CCK), orthogonal frequency division multiplexing (OFDM), and/or single-carrier frequency division multiple access (SC-FDMA) symbols) although the embodiments are not limited to any specific over-the-air interface or modulation scheme. Radio interface 2710 may include, for example, a receiver 2712, a frequency synthesizer 2714, and/or a transmitter 2716. Radio interface 2710 may include bias controls, a crystal oscillator and/or one or more antennas 2718-f. In another embodiment, radio interface 2710 may use external voltage-controlled oscillators (VCOs), surface acoustic wave filters, intermediate frequency (IF) filters and/or RF filters, as desired. Due to the variety of potential RF interface designs an expansive description thereof is omitted.

Baseband circuitry 2720 may communicate with radio interface 2710 to process receive and/or transmit signals and may include, for example, a mixer for down-converting received RF signals, an analog-to-digital converter 2722 for converting analog signals to digital form, a digital-to-analog converter 2724 for converting digital signals to analog form, and a mixer for up-converting signals for transmission. Further, baseband circuitry 2720 may include a baseband or physical layer (PHY) processing circuit 2726 for PHY link layer processing of respective receive/transmit signals. Baseband circuitry 2720 may include, for example, a medium access control (MAC) processing circuit 2727 for MAC/data link layer processing. Baseband circuitry 2720 may include a memory controller 2732 for communicating with MAC processing circuit 2727 and/or a computing platform 2730, for example, via one or more interfaces 2734.

In some embodiments, PHY processing circuit 2726 may include a frame construction and/or detection module, in combination with additional circuitry such as a buffer memory, to construct and/or deconstruct communication frames. Alternatively or in addition, MAC processing circuit 2727 may share processing for certain of these functions or perform these processes independent of PHY processing circuit 2726. In some embodiments, MAC and PHY processing may be integrated into a single circuit.

The computing platform 2730 may provide computing functionality for the device 2700. As shown, the computing platform 2730 may include a processing component 2740. In addition to, or alternatively of, the baseband circuitry 2720, the device 2700 may execute processing operations or logic for one or more of robots 1360, 1760, 1860, 1960, 2060A, and 2060B, automated maintenance device 1400, automation coordinator 1555, logic flows 2100, 2200, and 2300, storage media 2400 and 2450, computing architecture 2500, clients 2602, servers 2604,

and logic circuit 2728 using the processing component 2740. The processing component 2740 (and/or PHY 2726 and/or MAC 2727) may comprise various hardware elements, software elements, or a combination of both. Examples of hardware elements may include devices, logic devices, components, processors, microprocessors, circuits, processor circuits, circuit elements (e.g., transistors, resistors, capacitors, inductors, and so forth), integrated circuits, application specific integrated circuits (ASIC), programmable logic devices (PLD), digital signal processors (DSP), field programmable gate array (FPGA), memory units, logic gates, registers, semiconductor device, chips, microchips, chip sets, and so forth. Examples of software elements may include software components, programs, applications, computer programs, application programs, system programs, software development programs, machine programs, operating system software, middleware, firmware, software modules, routines, subroutines, functions, methods, procedures, software interfaces, application program interfaces (API), instruction sets, computing code, computer code, code segments, computer code segments, words, values, symbols, or any combination thereof. Determining whether an embodiment is implemented using hardware elements and/or software elements may vary in accordance with any number of factors, such as desired computational rate, power levels, heat tolerances, processing cycle budget, input data rates, output data rates, memory resources, data bus speeds and other design or performance constraints, as desired for a given implementation.

The computing platform 2730 may further include other platform components 2750. Other platform components 2750 include common computing elements, such as one or more processors, multi-core processors, co-processors, memory units, chipsets, controllers, peripherals, interfaces, oscillators, timing devices, video cards, audio cards, multimedia input/output (I/O) components (e.g., digital displays), power supplies, and so forth. Examples of memory units may include without limitation various types of computer readable and machine readable storage media in the form of one or more higher speed memory units, such as read-only memory (ROM), random-access memory (RAM), dynamic RAM (DRAM), Double-Data-Rate DRAM (DDRAM), synchronous DRAM (SDRAM), static RAM (SRAM), programmable ROM (PROM), erasable programmable ROM (EPROM), electrically erasable programmable ROM (EEPROM), flash memory, polymer memory such as ferroelectric polymer memory, ovonic memory, phase change or ferroelectric memory, silicon-oxide-nitride-oxide-silicon (SONOS) memory, magnetic or optical cards, an array of devices such as Redundant Array of Independent Disks (RAID) drives, solid state memory devices (e.g., USB memory, solid state drives (SSD) and any other type of storage media suitable for storing information.

Device 2700 may be, for example, an ultra-mobile device, a mobile device, a fixed device, a machine-to-machine (M2M) device, a personal digital assistant (PDA), a mobile computing device, a smart phone, a telephone, a digital telephone, a cellular telephone, user equipment, eBook readers, a handset, a one-way pager, a two-way pager, a messaging device, a computer, a
5 personal computer (PC), a desktop computer, a laptop computer, a notebook computer, a netbook computer, a handheld computer, a tablet computer, a server, a server array or server farm, a web server, a network server, an Internet server, a work station, a mini-computer, a main frame computer, a supercomputer, a network appliance, a web appliance, a distributed computing system, multiprocessor systems, processor-based systems, consumer electronics, programmable
10 consumer electronics, game devices, display, television, digital television, set top box, wireless access point, base station, node B, subscriber station, mobile subscriber center, radio network controller, router, hub, gateway, bridge, switch, machine, or combination thereof. Accordingly, functions and/or specific configurations of device 2700 described herein, may be included or omitted in various embodiments of device 2700, as suitably desired.

15 Embodiments of device 2700 may be implemented using single input single output (SISO) architectures. However, certain implementations may include multiple antennas (e.g., antennas 2718-f) for transmission and/or reception using adaptive antenna techniques for beamforming or spatial division multiple access (SDMA) and/or using MIMO communication techniques.

 The components and features of device 2700 may be implemented using any combination
20 of discrete circuitry, application specific integrated circuits (ASICs), logic gates and/or single chip architectures. Further, the features of device 2700 may be implemented using microcontrollers, programmable logic arrays and/or microprocessors or any combination of the foregoing where suitably appropriate. It is noted that hardware, firmware and/or software elements may be collectively or individually referred to herein as “logic” or “circuit.”

25 It should be appreciated that the exemplary device 2700 shown in the block diagram of FIG. 27 may represent one functionally descriptive example of many potential implementations. Accordingly, division, omission or inclusion of block functions depicted in the accompanying figures does not infer that the hardware components, circuits, software and/or elements for implementing these functions would be necessarily be divided, omitted, or included in
30 embodiments.

FIG. 28 illustrates an embodiment of a broadband wireless access system 2800. As shown in FIG. 28, broadband wireless access system 2800 may be an internet protocol (IP) type network comprising an internet 2810 type network or the like that is capable of supporting mobile wireless access and/or fixed wireless access to internet 2810. In one or more

embodiments, broadband wireless access system 2800 may comprise any type of orthogonal frequency division multiple access (OFDMA)-based or single-carrier frequency division multiple access (SC-FDMA)-based wireless network, such as a system compliant with one or more of the 3GPP LTE Specifications and/or IEEE 802.16 Standards, and the scope of the claimed subject matter is not limited in these respects.

In the exemplary broadband wireless access system 2800, radio access networks (RANs) 2812 and 2818 are capable of coupling with evolved node Bs (eNBs) 2814 and 2820, respectively, to provide wireless communication between one or more fixed devices 2816 and internet 2810 and/or between one or more mobile devices 2822 and Internet 2810. One example of a fixed device 2816 and a mobile device 2822 is device 2700 of FIG. 27, with the fixed device 2816 comprising a stationary version of device 2700 and the mobile device 2822 comprising a mobile version of device 2700. RANs 2812 and 2818 may implement profiles that are capable of defining the mapping of network functions to one or more physical entities on broadband wireless access system 2800. eNBs 2814 and 2820 may comprise radio equipment to provide RF communication with fixed device 2816 and/or mobile device 2822, such as described with reference to device 2700, and may comprise, for example, the PHY and MAC layer equipment in compliance with a 3GPP LTE Specification or an IEEE 802.16 Standard. eNBs 2814 and 2820 may further comprise an IP backplane to couple to Internet 2810 via RANs 2812 and 2818, respectively, although the scope of the claimed subject matter is not limited in these respects.

Broadband wireless access system 2800 may further comprise a visited core network (CN) 2824 and/or a home CN 2826, each of which may be capable of providing one or more network functions including but not limited to proxy and/or relay type functions, for example authentication, authorization and accounting (AAA) functions, dynamic host configuration protocol (DHCP) functions, or domain name service controls or the like, domain gateways such as public switched telephone network (PSTN) gateways or voice over internet protocol (VoIP) gateways, and/or internet protocol (IP) type server functions, or the like. However, these are merely example of the types of functions that are capable of being provided by visited CN 2824 and/or home CN 2826, and the scope of the claimed subject matter is not limited in these respects. Visited CN 2824 may be referred to as a visited CN in the case where visited CN 2824 is not part of the regular service provider of fixed device 2816 or mobile device 2822, for example where fixed device 2816 or mobile device 2822 is roaming away from its respective home CN 2826, or where broadband wireless access system 2800 is part of the regular service provider of fixed device 2816 or mobile device 2822 but where broadband wireless access

system 2800 may be in another location or state that is not the main or home location of fixed device 2816 or mobile device 2822. The embodiments are not limited in this context.

Fixed device 2816 may be located anywhere within range of one or both of eNBs 2814 and 2820, such as in or near a home or business to provide home or business customer broadband access to Internet 2810 via eNBs 2814 and 2820 and RANs 2812 and 2818, respectively, and home CN 2826. It is worthy of note that although fixed device 2816 is generally disposed in a stationary location, it may be moved to different locations as needed. Mobile device 2822 may be utilized at one or more locations if mobile device 2822 is within range of one or both of eNBs 2814 and 2820, for example. In accordance with one or more embodiments, operation support system (OSS) 2828 may be part of broadband wireless access system 2800 to provide management functions for broadband wireless access system 2800 and to provide interfaces between functional entities of broadband wireless access system 2800. Broadband wireless access system 2800 of FIG. 28 is merely one type of wireless network showing a certain number of the components of broadband wireless access system 2800, and the scope of the claimed subject matter is not limited in these respects.

FIG. 29 illustrates an embodiment of a wireless network 2900. As shown in FIG. 29, wireless network comprises an access point 2902 and wireless stations 2904, 2906, and 2908. Any one of access point 2902 and wireless stations 2904, 2906, and 2908 may potentially implement one or more of robots 1360, 1760, 1860, 1960, 2060A, and 2060B, automated maintenance device 1400, automation coordinator 1555, logic flows 2100, 2200, and 2300, storage media 2400 and 2450, computing architecture 2500, clients 2602, servers 2604, and communication device 2700.

In various embodiments, wireless network 2900 may comprise a wireless local area network (WLAN), such as a WLAN implementing one or more Institute of Electrical and Electronics Engineers (IEEE) 802.11 standards (sometimes collectively referred to as “Wi-Fi”). In some other embodiments, wireless network 2900 may comprise another type of wireless network, and/or may implement other wireless communications standards. In various embodiments, for example, wireless network 2900 may comprise a WWAN or WPAN rather than a WLAN. The embodiments are not limited to this example.

In some embodiments, wireless network 2900 may implement one or more broadband wireless communications standards, such as 3G or 4G standards, including their revisions, progeny, and variants. Examples of 3G or 4G wireless standards may include without limitation any of the IEEE 802.16m and 802.16p standards, 3rd Generation Partnership Project (3GPP) Long Term Evolution (LTE) and LTE-Advanced (LTE-A) standards, and International Mobile

Telecommunications Advanced (IMT-ADV) standards, including their revisions, progeny and variants. Other suitable examples may include, without limitation, Global System for Mobile Communications (GSM)/Enhanced Data Rates for GSM Evolution (EDGE) technologies, Universal Mobile Telecommunications System (UMTS)/High Speed Packet Access (HSPA) technologies, Worldwide Interoperability for Microwave Access (WiMAX) or the WiMAX II technologies, Code Division Multiple Access (CDMA) 2000 system technologies (e.g., CDMA2000 1xRTT, CDMA2000 EV-DO, CDMA EV-DV, and so forth), High Performance Radio Metropolitan Area Network (HIPERMAN) technologies as defined by the European Telecommunications Standards Institute (ETSI) Broadband Radio Access Networks (BRAN), Wireless Broadband (WiBro) technologies, GSM with General Packet Radio Service (GPRS) system (GSM/GPRS) technologies, High Speed Downlink Packet Access (HSDPA) technologies, High Speed Orthogonal Frequency-Division Multiplexing (OFDM) Packet Access (HSOPA) technologies, High-Speed Uplink Packet Access (HSUPA) system technologies, 3GPP Rel. 8-12 of LTE/System Architecture Evolution (SAE), and so forth. The embodiments are not limited in this context.

In various embodiments, wireless stations 2904, 2906, and 2908 may communicate with access point 2902 in order to obtain connectivity to one or more external data networks. In some embodiments, for example, wireless stations 2904, 2906, and 2908 may connect to the Internet 2912 via access point 2902 and access network 2910. In various embodiments, access network 2910 may comprise a private network that provides subscription-based Internet-connectivity, such as an Internet Service Provider (ISP) network. The embodiments are not limited to this example.

In various embodiments, two or more of wireless stations 2904, 2906, and 2908 may communicate with each other directly by exchanging peer-to-peer communications. For example, in the example of FIG. 29, wireless stations 2904 and 2906 communicate with each other directly by exchanging peer-to-peer communications 2914. In some embodiments, such peer-to-peer communications may be performed according to one or more Wi-Fi Alliance (WFA) standards. For example, in various embodiments, such peer-to-peer communications may be performed according to the WFA Wi-Fi Direct standard, 2010 Release. In various embodiments, such peer-to-peer communications may additionally or alternatively be performed using one or more interfaces, protocols, and/or standards developed by the WFA Wi-Fi Direct Services (WFDS) Task Group. The embodiments are not limited to these examples.

Various embodiments may be implemented using hardware elements, software elements, or a combination of both. Examples of hardware elements may include processors,

microprocessors, circuits, circuit elements (e.g., transistors, resistors, capacitors, inductors, and so forth), integrated circuits, application specific integrated circuits (ASIC), programmable logic devices (PLD), digital signal processors (DSP), field programmable gate array (FPGA), logic gates, registers, semiconductor device, chips, microchips, chip sets, and so forth. Examples of software may include software components, programs, applications, computer programs, application programs, system programs, machine programs, operating system software, middleware, firmware, software modules, routines, subroutines, functions, methods, procedures, software interfaces, application program interfaces (API), instruction sets, computing code, computer code, code segments, computer code segments, words, values, symbols, or any combination thereof. Determining whether an embodiment is implemented using hardware elements and/or software elements may vary in accordance with any number of factors, such as desired computational rate, power levels, heat tolerances, processing cycle budget, input data rates, output data rates, memory resources, data bus speeds and other design or performance constraints.

One or more aspects of at least one embodiment may be implemented by representative instructions stored on a machine-readable medium which represents various logic within the processor, which when read by a machine causes the machine to fabricate logic to perform the techniques described herein. Such representations, known as “IP cores” may be stored on a tangible, machine readable medium and supplied to various customers or manufacturing facilities to load into the fabrication machines that actually make the logic or processor. Some embodiments may be implemented, for example, using a machine-readable medium or article which may store an instruction or a set of instructions that, if executed by a machine, may cause the machine to perform a method and/or operations in accordance with the embodiments. Such a machine may include, for example, any suitable processing platform, computing platform, computing device, processing device, computing system, processing system, computer, processor, or the like, and may be implemented using any suitable combination of hardware and/or software. The machine-readable medium or article may include, for example, any suitable type of memory unit, memory device, memory article, memory medium, storage device, storage article, storage medium and/or storage unit, for example, memory, removable or non-removable media, erasable or non-erasable media, writeable or re-writeable media, digital or analog media, hard disk, floppy disk, Compact Disk Read Only Memory (CD-ROM), Compact Disk Recordable (CD-R), Compact Disk Rewriteable (CD-RW), optical disk, magnetic media, magneto-optical media, removable memory cards or disks, various types of Digital Versatile Disk (DVD), a tape, a cassette, or the like. The instructions may include any suitable type of

code, such as source code, compiled code, interpreted code, executable code, static code, dynamic code, encrypted code, and the like, implemented using any suitable high-level, low-level, object-oriented, visual, compiled and/or interpreted programming language.

The following examples pertain to further embodiments:

5 Example 1 is a method for automated data center maintenance, comprising processing, by processing circuitry of an automated maintenance device, an automation command received from an automation coordinator for a data center, identifying an automated maintenance procedure based on the received automation command, and performing the identified automated maintenance procedure.

10 Example 2 is the method of Example 1, the identified automated maintenance procedure to comprise a sled replacement procedure.

Example 3 is the method of Example 2, the sled replacement procedure to comprise replacing a compute sled.

15 Example 4 is the method of Example 3, the sled replacement procedure to comprise removing the compute sled from a sled space, removing a memory card from a connector slot of the compute sled, inserting the memory card into a connector slot of a replacement compute sled, and inserting the replacement compute sled into the sled space.

Example 5 is the method of Example 4, the memory card to store a compute state of the compute sled.

20 Example 6 is the method of Example 5, the sled replacement procedure to comprise initiating a restoration of the stored compute state on the replacement compute sled.

Example 7 is the method of Example 2, the sled replacement procedure to comprise replacing an accelerator sled.

25 Example 8 is the method of Example 2, the sled replacement procedure to comprise replacing a memory sled.

Example 9 is the method of Example 2, the sled replacement procedure to comprise replacing a storage sled.

Example 10 is the method of Example 1, the identified automated maintenance procedure to comprise a component replacement procedure.

30 Example 11 is the method of Example 10, the component replacement procedure to comprise removing a component from a socket of a sled, and inserting a replacement component into the socket.

Example 12 is the method of Example 11, the component to comprise a processor.

Example 13 is the method of Example 11, the component to comprise a field-programmable gate array (FPGA).

Example 14 is the method of Example 11, the component to comprise a memory module.

5 Example 15 is the method of Example 11, the component to comprise a non-volatile storage device.

Example 16 is the method of Example 15, the non-volatile storage device to comprise a solid-state drive (SSD).

Example 17 is the method of Example 16, the SSD to comprise a three-dimensional (3D) NAND SSD.

10 Example 18 is the method of Example 10, the component replacement procedure to comprise a cache memory replacement procedure.

Example 19 is the method of Example 18, the cache memory replacement procedure to comprise replacing one or more cache memory modules of a processor on a sled.

15 Example 20 is the method of Example 19, the cache memory replacement procedure to comprise removing a heat sink from atop the processor, removing the processor from a socket to facilitate access to one or more cache memory modules underlying the processor, removing the one or cache memory modules, inserting one or more replacement cache memory modules, reinserting the processor into the socket, and reinstalling the heat sink.

20 Example 21 is the method of Example 1, the identified automated maintenance procedure to comprise a component servicing procedure.

Example 22 is the method of Example 21, the component servicing procedure to comprise servicing a component on a sled.

Example 23 is the method of Example 22, the component servicing procedure to comprise removing the sled from a sled space of a rack.

25 Example 24 is the method of any of Examples 22 to 23, the component servicing procedure to comprise removing the component from the sled.

Example 25 is the method of any of Examples 22 to 24, the component servicing procedure to comprise testing the component.

30 Example 26 is the method of any of Examples 22 to 25, the component servicing procedure to comprise cleaning the component.

Example 27 is the method of any of Examples 22 to 26, the component servicing procedure to comprise power-cycling the component.

Example 28 is the method of any of Examples 22 to 27, the component servicing procedure to comprise capturing one or more images of the component.

Example 29 is the method of Example 28, comprising sending the one or more captured images to the automation coordinator.

Example 30 is the method of any of Examples 22 to 29, the component to comprise a processor.

5 Example 31 is the method of any of Examples 22 to 29, the component to comprise a field-programmable gate array (FPGA).

Example 32 is the method of any of Examples 22 to 29, the component to comprise a memory module.

10 Example 33 is the method of any of Examples 22 to 29, the component to comprise a non-volatile storage device.

Example 34 is the method of Example 33, the non-volatile storage device to comprise a solid-state drive (SSD).

Example 35 is the method of Example 34, the SSD to comprise a three-dimensional (3D) NAND SSD.

15 Example 36 is the method of any of Examples 1 to 35, comprising identifying the automated maintenance procedure based on a maintenance task code comprised in the received automation command.

Example 37 is the method of any of Examples 1 to 36, comprising performing the identified automated maintenance procedure based on one or more maintenance task parameters.

20 Example 38 is the method of Example 37, the one or more maintenance task parameters to be comprised in the received automation command.

Example 39 is the method of Example 37, at least one of the one or more maintenance task parameters to be comprised in a second automation command received from the automation coordinator.

25 Example 40 is the method of any of Examples 37 to 39, the one or more maintenance task parameters to include one or more location parameters.

Example 41 is the method of Example 40, the one or more location parameters to include a rack identifier (ID) associated with a rack within the data center.

30 Example 42 is the method of any of Examples 40 to 41, the one or more location parameters to include a sled space identifier (ID) associated with a sled space within the data center.

Example 43 is the method of any of Examples 40 to 42, the one or more location parameters to include a slot identifier (ID) associated with a connector socket on a sled within the data center.

Example 44 is the method of any of Examples 37 to 43, the one or more maintenance task parameters to include a sled identifier (ID) associated with a sled within the data center.

Example 45 is the method of any of Examples 37 to 44, the one or more maintenance task parameters to include a component identifier (ID) associated with a component on a sled within
5 the data center.

Example 46 is the method of any of Examples 1 to 45, the automation command to be comprised in signals received via a communication interface of the automated maintenance device.

Example 47 is the method of Example 46, the communication interface to comprise a radio
10 frequency (RF) interface, the signals to comprise RF signals.

Example 48 is the method of any of Examples 1 to 47, comprising sending a message to the automation coordinator to acknowledge the received automation command.

Example 49 is the method of any of Examples 1 to 48, comprising sending a message to the automation coordinator to report a result of the automated maintenance procedure.

Example 50 is the method of any of Examples 1 to 49, comprising sending position data to
15 the automation coordinator, the position data to indicate a position of the automated maintenance device within the data center.

Example 51 is the method of any of Examples 1 to 50, comprising sending assistance data to the automation coordinator, the assistance data to comprise an image of a component that is to
20 be manually replaced or serviced.

Example 52 is the method of any of Example 1 to 51, comprising sending environmental data to the automation coordinator, the environmental data to comprise measurements of one or more aspects of ambient conditions within the data center.

Example 53 is the method of Example 52, comprising one or more sensors to generate the
25 measurements comprised in the environmental data.

Example 54 is the method of any of Examples 52 to 53, the environmental data to comprise one or more temperature measurements.

Example 55 is the method of any of Examples 52 to 54, the environmental data to comprise one or more humidity measurements.

Example 56 is the method of any of Examples 52 to 55, the environmental data to
30 comprise one or more air quality measurements.

Example 57 is the method of any of Examples 52 to 56, the environmental data to comprise one or more pressure measurements.

Example 58 is a computer-readable storage medium storing instructions that, when executed, cause an automated maintenance device to perform a method according to any of Examples 1 to 57.

5 Example 59 is an automated maintenance device, comprising processing circuitry and computer-readable storage media storing instructions for execution by the processing circuitry to cause the automated maintenance device to perform a method according to any of Examples 1 to 57.

10 Example 60 is a method for coordination of automated data center maintenance, comprising identifying, by processing circuitry, a maintenance task to be performed in a data center, determining to initiate automated performance of the maintenance task, selecting an automated maintenance device to which to assign the maintenance task, and sending an automation command to cause the automated maintenance device to perform an automated maintenance procedure associated with the maintenance task.

15 Example 61 is the method of Example 60, comprising identifying the maintenance task based on telemetry data associated with one or more physical resources of the data center.

Example 62 is the method of Example 61, comprising receiving the telemetry data via a telemetry framework of the data center.

Example 63 is the method of any of Examples 61 to 62, the telemetry data to include one or more telemetry metrics associated with a physical compute resource.

20 Example 64 is the method of any of Examples 61 to 63, the telemetry data to include one or more telemetry metrics associated with a physical accelerator resource.

Example 65 is the method of any of Examples 61 to 64, the telemetry data to include one or more telemetry metrics associated with a physical memory resource.

25 Example 66 is the method of any of Examples 61 to 65, the telemetry data to include one or more telemetry metrics associated with a physical storage resource.

Example 67 is the method of any of Examples 60 to 66, comprising identifying the maintenance task based on environmental data received from one or more automated maintenance devices of the data center.

30 Example 68 is the method of Example 67, the environmental data to include one or more temperature measurements.

Example 69 is the method of any of Examples 67 to 68, the environmental data to include one or more humidity measurements.

Example 70 is the method of any of Examples 67 to 69, the environmental data to include one or more air quality measurements.

Example 71 is the method of any of Examples 67 to 70, the environmental data to include one or more pressure measurements.

Example 72 is the method of any of Examples 60 to 71, comprising adding the maintenance task to a pending task queue following identification of the maintenance task.

5 Example 73 is the method of Example 72, comprising determining to initiate automated performance of the maintenance task based on a determination that the maintenance task constitutes a highest priority task among one or more maintenance tasks comprised in the pending task queue.

10 Example 74 is the method of any of Examples 60 to 73, comprising selecting the automated maintenance device from among one or more automated maintenance devices in a candidate device pool.

Example 75 is the method of any of Examples 60 to 74, comprising selecting the automated maintenance device based on one or more capabilities of the automated maintenance device.

15 Example 76 is the method of any of Examples 60 to 75, comprising selecting the automated maintenance device based on position data received from the automated maintenance device.

Example 77 is the method of any of Examples 60 to 76, the automation command to comprise a maintenance task code indicating a task type associated with the maintenance task.

20 Example 78 is the method of any of Examples 60 to 77, the automation command to comprise location information associated with the maintenance task.

Example 79 is the method of Example 78, the location information to include a rack identifier (ID) associated with a rack within the data center.

25 Example 80 is the method of any of Examples 78 to 79, the location information to include a sled space identifier (ID) associated with a sled space within the data center.

Example 81 is the method of any of Examples 78 to 80, the location information to include a slot identifier (ID) associated with a connector socket on a sled within the data center.

Example 82 is the method of any of Examples 60 to 81, the automation command to comprise a sled identifier (ID) associated with a sled within the data center.

30 Example 83 is the method of any of Examples 60 to 82, the automation command to comprise a physical resource identifier (ID) associated with a physical resource within the data center.

Example 84 is the method of any of Examples 60 to 81, the maintenance task to comprise replacement of a sled.

Example 85 is the method of Example 83, the sled to comprise a compute sled, an accelerator sled, a memory sled, or a storage sled.

Example 86 is the method of any of Examples 60 to 81, the maintenance task to comprise replacement of one or more components of a sled.

5 Example 87 is the method of any of Examples 60 to 81, the maintenance task to comprise repair of one or more components of a sled.

Example 88 is the method of any of Examples 60 to 81, the maintenance task to comprise testing of one or more components of a sled.

10 Example 89 is the method of any of Examples 60 to 81, the maintenance task to comprise cleaning of one or more components of a sled.

Example 90 is the method of any of Examples 60 to 81, the maintenance task to comprise power cycling one or more memory modules.

Example 91 is the method of any of Examples 60 to 81, the maintenance task to comprise power cycling one or more non-volatile storage devices.

15 Example 92 is the method of any of Examples 60 to 81, the maintenance task to comprise storing a compute state of a compute sled, replacing the compute sled with a second compute sled, and transferring the stored compute state to the second compute sled.

Example 93 is the method of any of Examples 60 to 81, the maintenance task to comprise replacing one or more cache memory modules of a processor.

20 Example 94 is a computer-readable storage medium storing instructions that, when executed by an automation coordinator for a data center, cause the automation coordinator to perform a method according to any of Examples 60 to 93.

Example 95 is an apparatus, comprising processing circuitry and computer-readable storage media storing instructions for execution by the processing circuitry to perform a method
25 according to any of Examples 60 to 93.

Example 96 is a method for automated data center maintenance, comprising identifying, by processing circuitry of an automated maintenance device, a collaborative maintenance procedure to be performed in a data center, identifying a second automated maintenance device with which to collaborate during performance of the collaborative maintenance procedure, and sending
30 interdevice coordination information to the second automated maintenance device to initiate the collaborative maintenance procedure.

Example 97 is the method of Example 96, comprising identifying the collaborative maintenance procedure based on telemetry data associated with one or more physical resources of the data center.

Example 98 is the method of Example 97, the telemetry data to include one or more telemetry metrics associated with a physical compute resource.

Example 99 is the method of any of Examples 97 to 98, the telemetry data to include one or more telemetry metrics associated with a physical accelerator resource.

5 Example 100 is the method of any of Examples 97 to 99, the telemetry data to include one or more telemetry metrics associated with a physical memory resource.

Example 101 is the method of any of Examples 97 to 100, the telemetry data to include one or more telemetry metrics associated with a physical storage resource.

10 Example 102 is the method of any of Examples 96 to 101, comprising identifying the collaborative maintenance procedure based on environmental data comprising measurements of one or more aspects of ambient conditions within the data center.

Example 103 is the method of Example 102, comprising one or more sensors to generate the measurements comprised in the environmental data.

15 Example 104 is the method of any of Examples 102 to 103, the environmental data to comprise one or more temperature measurements.

Example 105 is the method of any of Examples 102 to 104, the environmental data to comprise one or more humidity measurements.

Example 106 is the method of any of Examples 102 to 105, the environmental data to comprise one or more air quality measurements.

20 Example 107 is the method of any of Examples 102 to 106, the environmental data to comprise one or more pressure measurements.

Example 108 is the method of Example 96, comprising identifying the collaborative maintenance procedure based on an automation command received from an automation coordinator for the data center.

25 Example 109 is the method of Example 108, comprising identifying the collaborative maintenance procedure based on a maintenance task code comprised in the received automation command.

Example 110 is the method of any of Examples 96 to 109, comprising selecting the second automated maintenance device from among a plurality of automated maintenance devices in a candidate device pool for the data center.

30 Example 111 is the method of any of Examples 96 to 110, comprising identifying the second automated maintenance device based on a parameter comprised in a command received from an automation coordinator for the data center.

Example 112 is the method of any of Examples 96 to 111, the collaborative maintenance procedure to comprise replacing a sled.

Example 113 is the method of Example 112, the sled to comprise a compute sled.

Example 114 is the method of Example 113, the collaborative maintenance procedure to
5 comprise removing the compute sled from a sled space, removing a memory card from a connector slot of the compute sled, inserting the memory card into a connector slot of a replacement compute sled, and inserting the replacement compute sled into the sled space.

Example 115 is the method of Example 114, the memory card to store a compute state of the compute sled.

10 Example 116 is the method of Example 115, the collaborative maintenance procedure to comprise initiating a restoration of the stored compute state on the replacement compute sled.

Example 117 is the method of Example 112, the sled to comprise an accelerator sled, a memory sled, or a storage sled.

Example 118 is the method of any of Examples 96 to 111, the collaborative maintenance
15 procedure to comprise replacing a component on a sled.

Example 119 is the method of Example 118, the component to comprise a processor.

Example 120 is the method of Example 118, the component to comprise a field-programmable gate array (FPGA).

Example 121 is the method of Example 118, the component to comprise a memory
20 module.

Example 122 is the method of Example 118, the component to comprise a non-volatile storage device.

Example 123 is the method of Example 122, the non-volatile storage device to comprise a solid-state drive (SSD).

25 Example 124 is the method of Example 123, the SSD to comprise a three-dimensional (3D) NAND SSD.

Example 125 is the method of any of Examples 96 to 111, the collaborative maintenance procedure to comprise replacing one or more cache memory modules of a processor on a sled.

Example 126 is the method of Example 125, the collaborative maintenance procedure to
30 comprise removing a heat sink from atop the processor, removing the processor from a socket to facilitate access to one or more cache memory modules underlying the processor, removing the one or cache memory modules, inserting one or more replacement cache memory modules, reinserting the processor into the socket, and reinstalling the heat sink.

Example 127 is the method of any of Examples 96 to 111, the collaborative maintenance procedure to comprise servicing a component on a sled.

Example 128 is the method of Example 127, the collaborative maintenance procedure to comprise removing the sled from a sled space of a rack.

5 Example 129 is the method of any of Examples 127 to 128, the collaborative maintenance procedure to comprise removing the component from the sled.

Example 130 is the method of any of Examples 127 to 129, the collaborative maintenance procedure to comprise testing the component.

10 Example 131 is the method of any of Examples 127 to 130, the collaborative maintenance procedure to comprise cleaning the component.

Example 132 is the method of any of Examples 127 to 131, the collaborative maintenance procedure to comprise power-cycling the component.

Example 133 is the method of any of Examples 127 to 132, the collaborative maintenance procedure to comprise capturing one or more images of the component.

15 Example 134 is the method of any of Examples 127 to 133, the component to comprise a processor.

Example 135 is the method of any of Examples 127 to 133, the component to comprise a field-programmable gate array (FPGA).

20 Example 136 is the method of any of Examples 127 to 133, the component to comprise a memory module.

Example 137 is the method of any of Examples 127 to 133, the component to comprise a non-volatile storage device.

Example 138 is the method of Example 137, the non-volatile storage device to comprise a solid-state drive (SSD).

25 Example 139 is the method of Example 138, the SSD to comprise a three-dimensional (3D) NAND SSD.

Example 140 is the method of any of Examples 96 to 139, the interdevice coordination information to comprise a rack identifier (ID) associated with a rack within the data center.

30 Example 141 is the method of any of Examples 96 to 140, the interdevice coordination information to comprise a sled space identifier (ID) associated with a sled space within the data center.

Example 142 is the method of any of Examples 96 to 141, the interdevice coordination information to comprise a slot identifier (ID) associated with a connector socket on a sled within the data center.

Example 143 is the method of any of Examples 96 to 142, the interdevice coordination information to comprise a sled identifier (ID) associated with a sled within the data center.

Example 144 is the method of any of Examples 96 to 143, the interdevice coordination information to comprise a component identifier (ID) associated with a component on a sled
5 within the data center.

Example 145 is a computer-readable storage medium storing instructions that, when executed, cause an automated maintenance device to perform a method according to any of Examples 96 to 144.

Example 146 is an automated maintenance device, comprising processing circuitry and
10 computer-readable storage media storing instructions for execution by the processing circuitry to cause the automated maintenance device to perform a method according to any of Examples 96 to 144.

Example 147 is an automated maintenance device, comprising means for receiving an automation command from an automation coordinator for a data center, means for identifying an
15 automated maintenance procedure based on the received automation command, and means for performing the identified automated maintenance procedure.

Example 148 is the automated maintenance device of Example 147, the identified automated maintenance procedure to comprise a sled replacement procedure.

Example 149 is the automated maintenance device of Example 148, the sled replacement
20 procedure to comprise removing a compute sled from a sled space, removing a memory card from a connector slot of the compute sled, inserting the memory card into a connector slot of a replacement compute sled, and inserting the replacement compute sled into the sled space.

Example 150 is the automated maintenance device of Example 149, the memory card to store a compute state of the compute sled.

Example 151 is the automated maintenance device of Example 150, the sled replacement
25 procedure to comprise initiating a restoration of the stored compute state on the replacement compute sled.

Example 152 is the automated maintenance device of Example 148, the sled replacement procedure to comprise replacing an accelerator sled, a memory sled, or a storage sled.

Example 153 is the automated maintenance device of Example 147, the identified
30 automated maintenance procedure to comprise a component replacement procedure.

Example 154 is the automated maintenance device of Example 153, the component replacement procedure to comprise removing a component from a socket of a sled, and inserting a replacement component into the socket.

Example 155 is the automated maintenance device of Example 154, the component to comprise a processor, a field-programmable gate array (FPGA), a memory module, or a solid-state drive (SSD).

5 Example 156 is the automated maintenance device of Example 153, the component replacement procedure to comprise a cache memory replacement procedure.

Example 157 is the automated maintenance device of Example 156, the cache memory replacement procedure to comprise replacing one or more cache memory modules of a processor on a sled.

10 Example 158 is the automated maintenance device of Example 157, the cache memory replacement procedure to comprise removing a heat sink from atop the processor, removing the processor from a socket to facilitate access to one or more cache memory modules underlying the processor, removing the one or cache memory modules, inserting one or more replacement cache memory modules, reinserting the processor into the socket, and reinstalling the heat sink.

15 Example 159 is the automated maintenance device of Example 147, the identified automated maintenance procedure to comprise a component servicing procedure.

Example 160 is the automated maintenance device of Example 159, the component servicing procedure to comprise servicing a component on a sled.

Example 161 is the automated maintenance device of Example 160, the component servicing procedure to comprise removing the sled from a sled space of a rack.

20 Example 162 is the automated maintenance device of any of Examples 160 to 161, the component servicing procedure to comprise removing the component from the sled.

Example 163 is the automated maintenance device of any of Examples 160 to 162, the component servicing procedure to comprise testing the component.

25 Example 164 is the automated maintenance device of any of Examples 160 to 163, the component servicing procedure to comprise cleaning the component.

Example 165 is the automated maintenance device of any of Examples 160 to 164, the component servicing procedure to comprise power-cycling the component.

Example 166 is the automated maintenance device of any of Examples 160 to 165, the component servicing procedure to comprise capturing one or more images of the component.

30 Example 167 is the automated maintenance device of any of Examples 160 to 166, the component to comprise a processor, a field-programmable gate array (FPGA), a memory module, or a solid-state drive (SSD).

Example 168 is the automated maintenance device of any of Examples 147to 167, comprising means for identifying the automated maintenance procedure based on a maintenance task code comprised in the received automation command.

5 Example 169 is the automated maintenance device of any of Examples 147to 168, comprising means for performing the identified automated maintenance procedure based on one or more maintenance task parameters.

Example 170 is the automated maintenance device of Example 169, the one or more maintenance task parameters to be comprised in the received automation command.

10 Example 171 is the automated maintenance device of Example 169, at least one of the one or more maintenance task parameters to be comprised in a second automation command received from the automation coordinator.

Example 172 is the automated maintenance device of any of Examples 169 to 171, the one or more maintenance task parameters to include one or more location parameters.

15 Example 173 is the automated maintenance device of Example 172, the one or more location parameters to include a rack identifier (ID) associated with a rack within the data center.

Example 174 is the automated maintenance device of any of Examples 172 to 173, the one or more location parameters to include a sled space identifier (ID) associated with a sled space within the data center.

20 Example 175 is the automated maintenance device of any of Examples 172 to 174, the one or more location parameters to include a slot identifier (ID) associated with a connector socket on a sled within the data center.

Example 176 is the automated maintenance device of any of Examples 169 to 175, the one or more maintenance task parameters to include a sled identifier (ID) associated with a sled within the data center.

25 Example 177 is the automated maintenance device of any of Examples 169 to 176, the one or more maintenance task parameters to include a component identifier (ID) associated with a component on a sled within the data center.

30 Example 178 is the automated maintenance device of any of Examples 147to 177, the automation command to be comprised in signals received via a communication interface of the automated maintenance device.

Example 179 is the automated maintenance device of Example 178, the communication interface to comprise a radio frequency (RF) interface, the signals to comprise RF signals.

Example 180 is the automated maintenance device of any of Examples 147to 179, comprising means for sending a message to the automation coordinator to acknowledge the received automation command.

5 Example 181 is the automated maintenance device of any of Examples 147to 180, comprising means for sending a message to the automation coordinator to report a result of the automated maintenance procedure.

Example 182 is the automated maintenance device of any of Examples 147to 181, comprising means for sending position data to the automation coordinator, the position data to indicate a position of the automated maintenance device within the data center.

10 Example 183 is the automated maintenance device of any of Examples 147to 182, comprising means for sending assistance data to the automation coordinator, the assistance data to comprise an image of a component that is to be manually replaced or serviced.

Example 184 is the automated maintenance device of any of Example 147to 183, comprising means for sending environmental data to the automation coordinator, the
15 environmental data to comprise measurements of one or more aspects of ambient conditions within the data center.

Example 185 is the automated maintenance device of Example 184, comprising means for generating the measurements comprised in the environmental data.

20 Example 186 is the automated maintenance device of any of Examples 184 to 185, the environmental data to comprise one or more temperature measurements.

Example 187 is the automated maintenance device of any of Examples 184 to 186, the environmental data to comprise one or more humidity measurements.

Example 188 is the automated maintenance device of any of Examples 184 to 187, the environmental data to comprise one or more air quality measurements.

25 Example 189 is the automated maintenance device of any of Examples 184 to 188, the environmental data to comprise one or more pressure measurements.

Example 189 is an apparatus for coordination of automated data center maintenance, comprising means for identifying a maintenance task to be performed in a data center, means for determining to initiate automated performance of the maintenance task, means for selecting an
30 automated maintenance device to which to assign the maintenance task, and means for sending an automation command to cause the automated maintenance device to perform an automated maintenance procedure associated with the maintenance task.

Example 190 is the apparatus of Example 189, comprising means for identifying the maintenance task based on telemetry data associated with one or more physical resources of the data center.

5 Example 191 is the apparatus of Example 190, comprising means for receiving the telemetry data via a telemetry framework of the data center.

Example 192 is the apparatus of any of Examples 190 to 191, the telemetry data to include one or more telemetry metrics associated with a physical compute resource.

Example 193 is the apparatus of any of Examples 190 to 192, the telemetry data to include one or more telemetry metrics associated with a physical accelerator resource.

10 Example 194 is the apparatus of any of Examples 190 to 193, the telemetry data to include one or more telemetry metrics associated with a physical memory resource.

Example 195 is the apparatus of any of Examples 190 to 194, the telemetry data to include one or more telemetry metrics associated with a physical storage resource.

15 Example 196 is the apparatus of any of Examples 189 to 195, comprising means for identifying the maintenance task based on environmental data received from one or more automated maintenance devices of the data center.

Example 197 is the apparatus of Example 196, the environmental data to include one or more temperature measurements.

20 Example 198 is the apparatus of any of Examples 196 to 197, the environmental data to include one or more humidity measurements.

Example 199 is the apparatus of any of Examples 196 to 198, the environmental data to include one or more air quality measurements.

Example 200 is the apparatus of any of Examples 196 to 199, the environmental data to include one or more pressure measurements.

25 Example 201 is the apparatus of any of Examples 189 to 200, comprising means for adding the maintenance task to a pending task queue following identification of the maintenance task.

Example 202 is the apparatus of Example 201, comprising means for determining to initiate automated performance of the maintenance task based on a determination that the maintenance task constitutes a highest priority task among one or more maintenance tasks
30 comprised in the pending task queue.

Example 203 is the apparatus of any of Examples 189 to 202, comprising means for selecting the automated maintenance device from among one or more automated maintenance devices in a candidate device pool.

Example 204 is the apparatus of any of Examples 189 to 203, comprising means for selecting the automated maintenance device based on one or more capabilities of the automated maintenance device.

5 Example 205 is the apparatus of any of Examples 189 to 204, comprising means for selecting the automated maintenance device based on position data received from the automated maintenance device.

Example 206 is the apparatus of any of Examples 189 to 205, the automation command to comprise a maintenance task code indicating a task type associated with the maintenance task.

10 Example 207 is the apparatus of any of Examples 189 to 206, the automation command to comprise location information associated with the maintenance task.

Example 208 is the apparatus of Example 207, the location information to include a rack identifier (ID) associated with a rack within the data center.

Example 209 is the apparatus of any of Examples 207 to 208, the location information to include a sled space identifier (ID) associated with a sled space within the data center.

15 Example 210 is the apparatus of any of Examples 207 to 209, the location information to include a slot identifier (ID) associated with a connector socket on a sled within the data center.

Example 211 is the apparatus of any of Examples 189 to 210, the automation command to comprise a sled identifier (ID) associated with a sled within the data center.

20 Example 212 is the apparatus of any of Examples 189 to 211, the automation command to comprise a physical resource identifier (ID) associated with a physical resource within the data center.

Example 213 is the apparatus of any of Examples 189 to 212, the maintenance task to comprise replacement of a sled.

25 Example 214 is the apparatus of Example 213, the sled to comprise a compute sled, an accelerator sled, a memory sled, or a storage sled.

Example 215 is the apparatus of any of Examples 189 to 212, the maintenance task to comprise replacement of one or more components of a sled.

Example 216 is the apparatus of any of Examples 189 to 212, the maintenance task to comprise repair of one or more components of a sled.

30 Example 217 is the apparatus of any of Examples 189 to 212, the maintenance task to comprise testing of one or more components of a sled.

Example 218 is the apparatus of any of Examples 189 to 212, the maintenance task to comprise cleaning of one or more components of a sled.

Example 219 is the apparatus of any of Examples 189 to 212, the maintenance task to comprise power cycling one or more memory modules.

Example 220 is the apparatus of any of Examples 189 to 212, the maintenance task to comprise power cycling one or more non-volatile storage devices.

5 Example 221 is the apparatus of any of Examples 189 to 212, the maintenance task to comprise storing a compute state of a compute sled, replacing the compute sled with a second compute sled, and transferring the stored compute state to the second compute sled.

Example 222 is the apparatus of any of Examples 189 to 212, the maintenance task to comprise replacing one or more cache memory modules of a processor.

10 Example 223 is an automated maintenance device, comprising means for identifying a collaborative maintenance procedure to be performed in a data center, means for identifying a second automated maintenance device with which to collaborate during performance of the collaborative maintenance procedure, and means for sending interdevice coordination information to the second automated maintenance device to initiate the collaborative
15 maintenance procedure.

Example 224 is the automated maintenance device of Example 223, comprising means for identifying the collaborative maintenance procedure based on telemetry data associated with one or more physical resources of the data center.

Example 225 is the automated maintenance device of Example 224, the telemetry data to
20 include one or more telemetry metrics associated with a physical compute resource.

Example 226 is the automated maintenance device of any of Examples 224 to 225, the telemetry data to include one or more telemetry metrics associated with a physical accelerator resource.

Example 227 is the automated maintenance device of any of Examples 224 to 226, the
25 telemetry data to include one or more telemetry metrics associated with a physical memory resource.

Example 228 is the automated maintenance device of any of Examples 224 to 227, the telemetry data to include one or more telemetry metrics associated with a physical storage resource.

30 Example 229 is the automated maintenance device of any of Examples 223 to 228, comprising means for identifying the collaborative maintenance procedure based on environmental data comprising measurements of one or more aspects of ambient conditions within the data center.

Example 230 is the automated maintenance device of Example 229, comprising one or more sensors to generate the measurements comprised in the environmental data.

Example 231 is the automated maintenance device of any of Examples 229 to 230, the environmental data to comprise one or more temperature measurements.

5 Example 232 is the automated maintenance device of any of Examples 229 to 231, the environmental data to comprise one or more humidity measurements.

Example 233 is the automated maintenance device of any of Examples 229 to 232, the environmental data to comprise one or more air quality measurements.

10 Example 234 is the automated maintenance device of any of Examples 229 to 233, the environmental data to comprise one or more pressure measurements.

Example 235 is the automated maintenance device of Example 223, comprising means for identifying the collaborative maintenance procedure based on an automation command received from an automation coordinator for the data center.

15 Example 236 is the automated maintenance device of Example 235, comprising means for identifying the collaborative maintenance procedure based on a maintenance task code comprised in the received automation command.

Example 237 is the automated maintenance device of any of Examples 223 to 236, comprising means for selecting the second automated maintenance device from among a plurality of automated maintenance devices in a candidate device pool for the data center.

20 Example 238 is the automated maintenance device of any of Examples 223 to 237, comprising means for identifying the second automated maintenance device based on a parameter comprised in a command received from an automation coordinator for the data center.

Example 239 is the automated maintenance device of any of Examples 223 to 238, the collaborative maintenance procedure to comprise replacing a sled.

25 Example 240 is the automated maintenance device of Example 239, the sled to comprise a compute sled.

Example 241 is the automated maintenance device of Example 240, the collaborative maintenance procedure to comprise removing the compute sled from a sled space, removing a memory card from a connector slot of the compute sled, inserting the memory card into a connector slot of a replacement compute sled, and inserting the replacement compute sled into the sled space.

Example 242 is the automated maintenance device of Example 241, the memory card to store a compute state of the compute sled.

Example 243 is the automated maintenance device of Example 242, the collaborative maintenance procedure to comprise initiating a restoration of the stored compute state on the replacement compute sled.

5 Example 244 is the automated maintenance device of Example 239, the sled to comprise an accelerator sled, a memory sled, or a storage sled.

Example 245 is the automated maintenance device of any of Examples 223 to 238, the collaborative maintenance procedure to comprise replacing a component on a sled.

10 Example 246 is the automated maintenance device of Example 245, the component to comprise a processor, a field-programmable gate array (FPGA), a memory module, or a solid-state drive (SSD).

Example 247 is the automated maintenance device of any of Examples 223 to 238, the collaborative maintenance procedure to comprise replacing one or more cache memory modules of a processor on a sled.

15 Example 248 is the automated maintenance device of Example 247, the collaborative maintenance procedure to comprise removing a heat sink from atop the processor, removing the processor from a socket to facilitate access to one or more cache memory modules underlying the processor, removing the one or cache memory modules, inserting one or more replacement cache memory modules, reinserting the processor into the socket, and reinstalling the heat sink.

20 Example 249 is the automated maintenance device of any of Examples 223 to 238, the collaborative maintenance procedure to comprise servicing a component on a sled.

Example 250 is the automated maintenance device of Example 249, the collaborative maintenance procedure to comprise removing the sled from a sled space of a rack.

Example 251 is the automated maintenance device of any of Examples 249 to 250, the collaborative maintenance procedure to comprise removing the component from the sled.

25 Example 252 is the automated maintenance device of any of Examples 249 to 251, the collaborative maintenance procedure to comprise testing the component.

Example 253 is the automated maintenance device of any of Examples 249 to 252, the collaborative maintenance procedure to comprise cleaning the component.

30 Example 254 is the automated maintenance device of any of Examples 249 to 253, the collaborative maintenance procedure to comprise power-cycling the component.

Example 255 is the automated maintenance device of any of Examples 249 to 254, the collaborative maintenance procedure to comprise capturing one or more images of the component.

Example 256 is the automated maintenance device of any of Examples 249 to 255, the component to comprise a processor, a field-programmable gate array (FPGA), a memory module, or a solid-state drive (SSD).

5 Example 257 is the automated maintenance device of any of Examples 223 to 256, the interdevice coordination information to comprise a rack identifier (ID) associated with a rack within the data center.

Example 258 is the automated maintenance device of any of Examples 223 to 257, the interdevice coordination information to comprise a sled space identifier (ID) associated with a sled space within the data center.

10 Example 259 is the automated maintenance device of any of Examples 223 to 258, the interdevice coordination information to comprise a slot identifier (ID) associated with a connector socket on a sled within the data center.

Example 260 is the automated maintenance device of any of Examples 223 to 259, the interdevice coordination information to comprise a sled identifier (ID) associated with a sled
15 within the data center.

Example 261 is the automated maintenance device of any of Examples 223 to 260, the interdevice coordination information to comprise a component identifier (ID) associated with a component on a sled within the data center.

20 Numerous specific details have been set forth herein to provide a thorough understanding of the embodiments. It will be understood by those skilled in the art, however, that the embodiments may be practiced without these specific details. In other instances, well-known operations, components, and circuits have not been described in detail so as not to obscure the embodiments. It can be appreciated that the specific structural and functional details disclosed herein may be representative and do not necessarily limit the scope of the embodiments.

25 Some embodiments may be described using the expression "coupled" and "connected" along with their derivatives. These terms are not intended as synonyms for each other. For example, some embodiments may be described using the terms "connected" and/or "coupled" to indicate that two or more elements are in direct physical or electrical contact with each other. The term "coupled," however, may also mean that two or more elements are not in direct contact
30 with each other, but yet still co-operate or interact with each other.

Unless specifically stated otherwise, it may be appreciated that terms such as "processing," "computing," "calculating," "determining," or the like, refer to the action and/or processes of a computer or computing system, or similar electronic computing device, that manipulates and/or transforms data represented as physical quantities (e.g., electronic) within the computing

system's registers and/or memories into other data similarly represented as physical quantities within the computing system's memories, registers or other such information storage, transmission or display devices. The embodiments are not limited in this context.

It should be noted that the methods described herein do not have to be executed in the order described, or in any particular order. Moreover, various activities described with respect to the methods identified herein can be executed in serial or parallel fashion.

Although specific embodiments have been illustrated and described herein, it should be appreciated that any arrangement calculated to achieve the same purpose may be substituted for the specific embodiments shown. This disclosure is intended to cover any and all adaptations or variations of various embodiments. It is to be understood that the above description has been made in an illustrative fashion, and not a restrictive one. Combinations of the above embodiments, and other embodiments not specifically described herein will be apparent to those of skill in the art upon reviewing the above description. Thus, the scope of various embodiments includes any other applications in which the above compositions, structures, and methods are used.

It is emphasized that the Abstract of the Disclosure is provided to comply with 37 C.F.R. § 1.72(b), requiring an abstract that will allow the reader to quickly ascertain the nature of the technical disclosure. It is submitted with the understanding that it will not be used to interpret or limit the scope or meaning of the claims. In addition, in the foregoing Detailed Description, it can be seen that various features are grouped together in a single embodiment for the purpose of streamlining the disclosure. This method of disclosure is not to be interpreted as reflecting an intention that the claimed embodiments require more features than are expressly recited in each claim. Rather, as the following claims reflect, inventive subject matter lies in less than all features of a single disclosed embodiment. Thus the following claims are hereby incorporated into the Detailed Description, with each claim standing on its own as a separate preferred embodiment. In the appended claims, the terms "including" and "in which" are used as the plain-English equivalents of the respective terms "comprising" and "wherein," respectively. Moreover, the terms "first," "second," and "third," etc. are used merely as labels, and are not intended to impose numerical requirements on their objects.

Although the subject matter has been described in language specific to structural features and/or methodological acts, it is to be understood that the subject matter defined in the appended claims is not necessarily limited to the specific features or acts described above. Rather, the specific features and acts described above are disclosed as example forms of implementing the claims.

CLAIMS

What is claimed is:

1. An automated maintenance device, comprising:
processing circuitry; and
5 non-transitory computer-readable storage media comprising instructions for execution by the processing circuitry to cause the automated maintenance device to:
receive an automation command from an automation coordinator for a data center;
identify an automated maintenance procedure based on the received automation
command; and
10 perform the identified automated maintenance procedure in the data center.
2. The automated maintenance device of claim 1, the automated maintenance procedure to comprise replacing a compute sled in the data center.
3. The automated maintenance device of claim 2, the automated maintenance procedure to comprise:
15 removing the compute sled from a sled space within a rack;
removing a memory card from a connector slot of the compute sled, the memory card to store a compute state of the compute sled;
inserting the memory card into a connector slot of a replacement compute sled;
inserting the replacement compute sled into the sled space; and
20 initiating a restoration of the stored compute state on the replacement compute sled.
4. The automated maintenance device of claim 1, the automated maintenance procedure to comprise replacing one or more cache memory modules of a processor on a sled.
5. The automated maintenance device of claim 4, the automated maintenance procedure to comprise:
25 removing the processor from a socket to facilitate access to one or more cache memory modules underlying the processor;
removing the one or cache memory modules;
inserting one or more replacement cache memory modules; and
reinserting the processor into the socket.
- 30 6. The automated maintenance device of claim 5, the automated maintenance procedure to comprise:
removing a heat sink from atop the processor prior to removing the processor from the socket; and
reinstalling the heat sink after reinserting the processor into the socket.

7. The automated maintenance device of claim 1, comprising a radio frequency (RF) interface to receive a wireless signal comprising the automation command.
8. An apparatus for coordination of automated data center maintenance, comprising:
- 5 processing circuitry; and
- non-transitory computer-readable storage media comprising instructions for execution by the processing circuitry to:
- identify a maintenance task to be performed in a data center;
- determine to initiate automated performance of the maintenance task;
- 10 select an automated maintenance device to which to assign the maintenance task; and
- send an automation command to cause the automated maintenance device to perform an automated maintenance procedure associated with the maintenance task.
9. The apparatus of claim 8, the non-transitory computer-readable storage media comprising instructions for execution by the processing circuitry to identify the maintenance task based on
- 15 telemetry data associated with one or more physical resources of the data center.
10. The apparatus of claim 8, the non-transitory computer-readable storage media comprising instructions for execution by the processing circuitry to identify the maintenance task based on environmental data received from one or more automated maintenance devices of the data center.
- 20 11. The apparatus of claim 8, the non-transitory computer-readable storage media comprising instructions for execution by the processing circuitry to add the maintenance task to a pending task queue following identification of the maintenance task.
12. The apparatus of claim 11, the non-transitory computer-readable storage media comprising instructions for execution by the processing circuitry to determine to initiate automated
- 25 performance of the maintenance task based on a determination that the maintenance task constitutes a highest priority task among one or more maintenance tasks comprised in the pending task queue.
13. The apparatus of claim 12, the non-transitory computer-readable storage media comprising instructions for execution by the processing circuitry to select the automated maintenance device
- 30 from among one or more automated maintenance devices in a candidate device pool.
14. A method for automated data center maintenance, comprising:
- receiving, at an automated maintenance device, an automation command from an automation coordinator for a data center;

- identifying, by processing circuitry of the automated maintenance device, an automated maintenance procedure based on the received automation command; and
- performing the identified automated maintenance procedure in the data center.
15. The method of claim 14, the automated maintenance procedure to comprise replacing a
5 compute sled in the data center.
16. The method of claim 15, the automated maintenance procedure to comprise:
- removing the compute sled from a sled space within a rack;
- removing a memory card from a connector slot of the compute sled, the memory card to store
a compute state of the compute sled;
- 10 inserting the memory card into a connector slot of a replacement compute sled;
- inserting the replacement compute sled into the sled space; and
- initiating a restoration of the stored compute state on the replacement compute sled.
17. The method of claim 14, the automated maintenance procedure to comprise replacing one or
more cache memory modules of a processor on a sled.
- 15 18. The method of claim 17, the automated maintenance procedure to comprise:
- removing the processor from a socket to facilitate access to one or more cache memory
modules underlying the processor;
- removing the one or cache memory modules;
- inserting one or more replacement cache memory modules; and
- 20 reinserting the processor into the socket.
19. The method of claim 18, the automated maintenance procedure to comprise:
- removing a heat sink from atop the processor prior to removing the processor from the
socket; and
- reinstalling the heat sink after reinserting the processor into the socket.
- 25 20. At least one non-transitory computer-readable storage medium comprising a set of
instructions that, when executed by an automation coordinator for a data center, cause the
automation coordinator to:
- identify a maintenance task to be performed in a data center;
- determine to initiate automated performance of the maintenance task;
- 30 select an automated maintenance device to which to assign the maintenance task; and
- send an automation command to cause the automated maintenance device to perform an
automated maintenance procedure associated with the maintenance task.
21. The at least one non-transitory computer-readable storage medium of claim 20, comprising
instructions that, when executed by the automation coordinator, cause the automation

coordinator to identify the maintenance task based on telemetry data associated with one or more physical resources of the data center.

22. The at least one non-transitory computer-readable storage medium of claim 20, comprising instructions that, when executed by the automation coordinator, cause the automation
- 5 coordinator to identify the maintenance task based on environmental data received from one or more automated maintenance devices of the data center.

23. The at least one non-transitory computer-readable storage medium of claim 20, comprising instructions that, when executed by the automation coordinator, cause the automation
- coordinator to add the maintenance task to a pending task queue following identification of the
- 10 maintenance task.

24. The at least one non-transitory computer-readable storage medium of claim 23, comprising instructions that, when executed by the automation coordinator, cause the automation
- coordinator to determine to initiate automated performance of the maintenance task based on a determination that the maintenance task constitutes a highest priority task among one or more
- 15 maintenance tasks comprised in the pending task queue.

25. The at least one non-transitory computer-readable storage medium of claim 24, comprising instructions that, when executed by the automation coordinator, cause the automation
- coordinator to select the automated maintenance device from among one or more automated maintenance devices in a candidate device pool.

20


FIG. 1***Data Center 100***


FIG. 2

FIG. 3

Data Center 300





FIG. 4

FIG. 5

500




FIG. 6

600

Rack Access Region 601




FIG. 7

Sled 704





FIG. 8**800**

FIG. 9




FIG. 10




FIG. 11

Data Center 1100





FIG. 12**1200**

FIG. 13

1300





FIG. 14


FIG. 15**1500**

FIG. 16


1600

FIG. 17

1700




FIG. 18

1800




FIG. 19

1900




FIG. 20

2000





FIG. 21**2100**


FIG. 22**2200**


FIG. 23**2300**

FIG. 24A




FIG. 24B




FIG. 25


2500


FIG. 26**2600**

FIG. 27

Device 2700




FIG. 28

Broadband Wireless Access System

2800





FIG. 29



A. CLASSIFICATION OF SUBJECT MATTER**G06F 17/30(2006.01)i, G06F 13/40(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 17/30; G06F 9/44; B25J 5/02; G06F 9/06; G06F 11/00; B25J 9/16; G06F 11/20; G06F 13/40

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 center, automation, maintenance, server, rack, replace, assign, task, queue, 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.
X	US 2007-0094535 A1 (DAVID PRASEE) 26 April 2007 See paragraphs [0016], [0030]-[0034], and [0046]-[0054]; claims 1, 4, and 17; and figures 1, 3, and 8A-8B.	1-10, 14-22
Y		11-13, 23-25
Y	WO 2012-033773 A2 (MICROSOFT CORPORATION) 15 March 2012 See paragraphs [0002]-[0003] and [0021]-[0046]; and figures 2-4.	11-13, 23-25
A	US 2016-0107312 A1 (CENTURYLINK INTELLECTUAL PROPERTY LLC) 21 April 2016 See paragraphs [0010]-[0018]; and figures 1 and 7A-7B.	1-25
A	US 2009-0222686 A1 (JOHN P. NIBARGER et al.) 03 September 2009 See paragraphs [0009]-[0021]; and figures 1 and 4.	1-25
A	US 2014-0173336 A1 (INTERNATIONAL BUSINESS MACHINES CORPORATION) 19 June 2014 See paragraphs [0006] and [0057]-[0070]; and figures 1-4.	1-25



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

31 October 2017 (31.10.2017)

Date of mailing of the international search report

31 October 2017 (31.10.2017)

Name and mailing address of the ISA/KR

International Application Division

Korean Intellectual Property Office

189 Cheongsa-ro, Seo-gu, Daejeon, 35208, Republic of Korea



Facsimile No. +82-42-481-8578

Authorized officer

NHO, Ji Myong

Telephone No. +82-42-481-8528



INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2017/043182

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2007-0094535 A1	26/04/2007	US 7725212 B2	25/05/2010
WO 2012-033773 A2	15/03/2012	AU 2011-299337 A1	21/03/2013
		AU 2011-299337 B2	15/05/2014
		CN 102385541 A	21/03/2012
		CN 102385541 B	03/08/2016
		EP 2614436 A2	17/07/2013
		EP 2614436 A4	02/12/2015
		EP 2614436 B1	09/08/2017
		JP 2014-506342 A	13/03/2014
		JP 6132766 B2	24/05/2017
		KR 10-2013-0136449 A	12/12/2013
		US 2012-0066541 A1	15/03/2012
		US 8365009 B2	29/01/2013
		WO 2012-033773 A3	24/05/2012
US 2016-0107312 A1	21/04/2016	US 2017-144305 A1	25/05/2017
		US 9597801 B2	21/03/2017
US 2009-0222686 A1	03/09/2009	None	
US 2014-0173336 A1	19/06/2014	US 9122652 B2	01/09/2015