



(19) **United States**

(12) **Patent Application Publication**
Kalsi et al.

(10) **Pub. No.: US 2020/0076192 A1**

(43) **Pub. Date: Mar. 5, 2020**

(54) **DISTRIBUTED HIERARCHICAL CONTROL ARCHITECTURE FOR INTEGRATING SMART GRID ASSETS DURING NORMAL AND DISRUPTED OPERATIONS**

Publication Classification

(51) **Int. Cl.**
H02J 3/16 (2006.01)
G05F 5/00 (2006.01)
H02J 3/00 (2006.01)
H02J 13/00 (2006.01)
H02J 3/06 (2006.01)

(71) Applicant: **Battelle Memorial Institute**, Richland, WA (US)

(72) Inventors: **Karanjit Kalsi**, Richland, WA (US);
Jason C. Fuller, Richland, WA (US);
Abhishek Somani, Richland, WA (US);
Robert G. Pratt, Kennewick, WA (US);
David P. Chassin, Pasco, WA (US);
Donald J. Hammerstrom, West Richland, WA (US)

(52) **U.S. Cl.**
CPC **H02J 3/16** (2013.01); **G05F 5/00** (2013.01); **H02J 3/00** (2013.01); **H02J 3/008** (2013.01); **H02J 13/0006** (2013.01); **H02J 3/382** (2013.01); **Y04S 10/123** (2013.01); **Y04S 10/12** (2013.01); **Y04S 50/10** (2013.01); **Y02E 40/72** (2013.01); **H02J 3/06** (2013.01)

(73) Assignee: **Battelle Memorial Institute**, Richland, WA (US)

(57) **ABSTRACT**

(21) Appl. No.: **16/671,353**

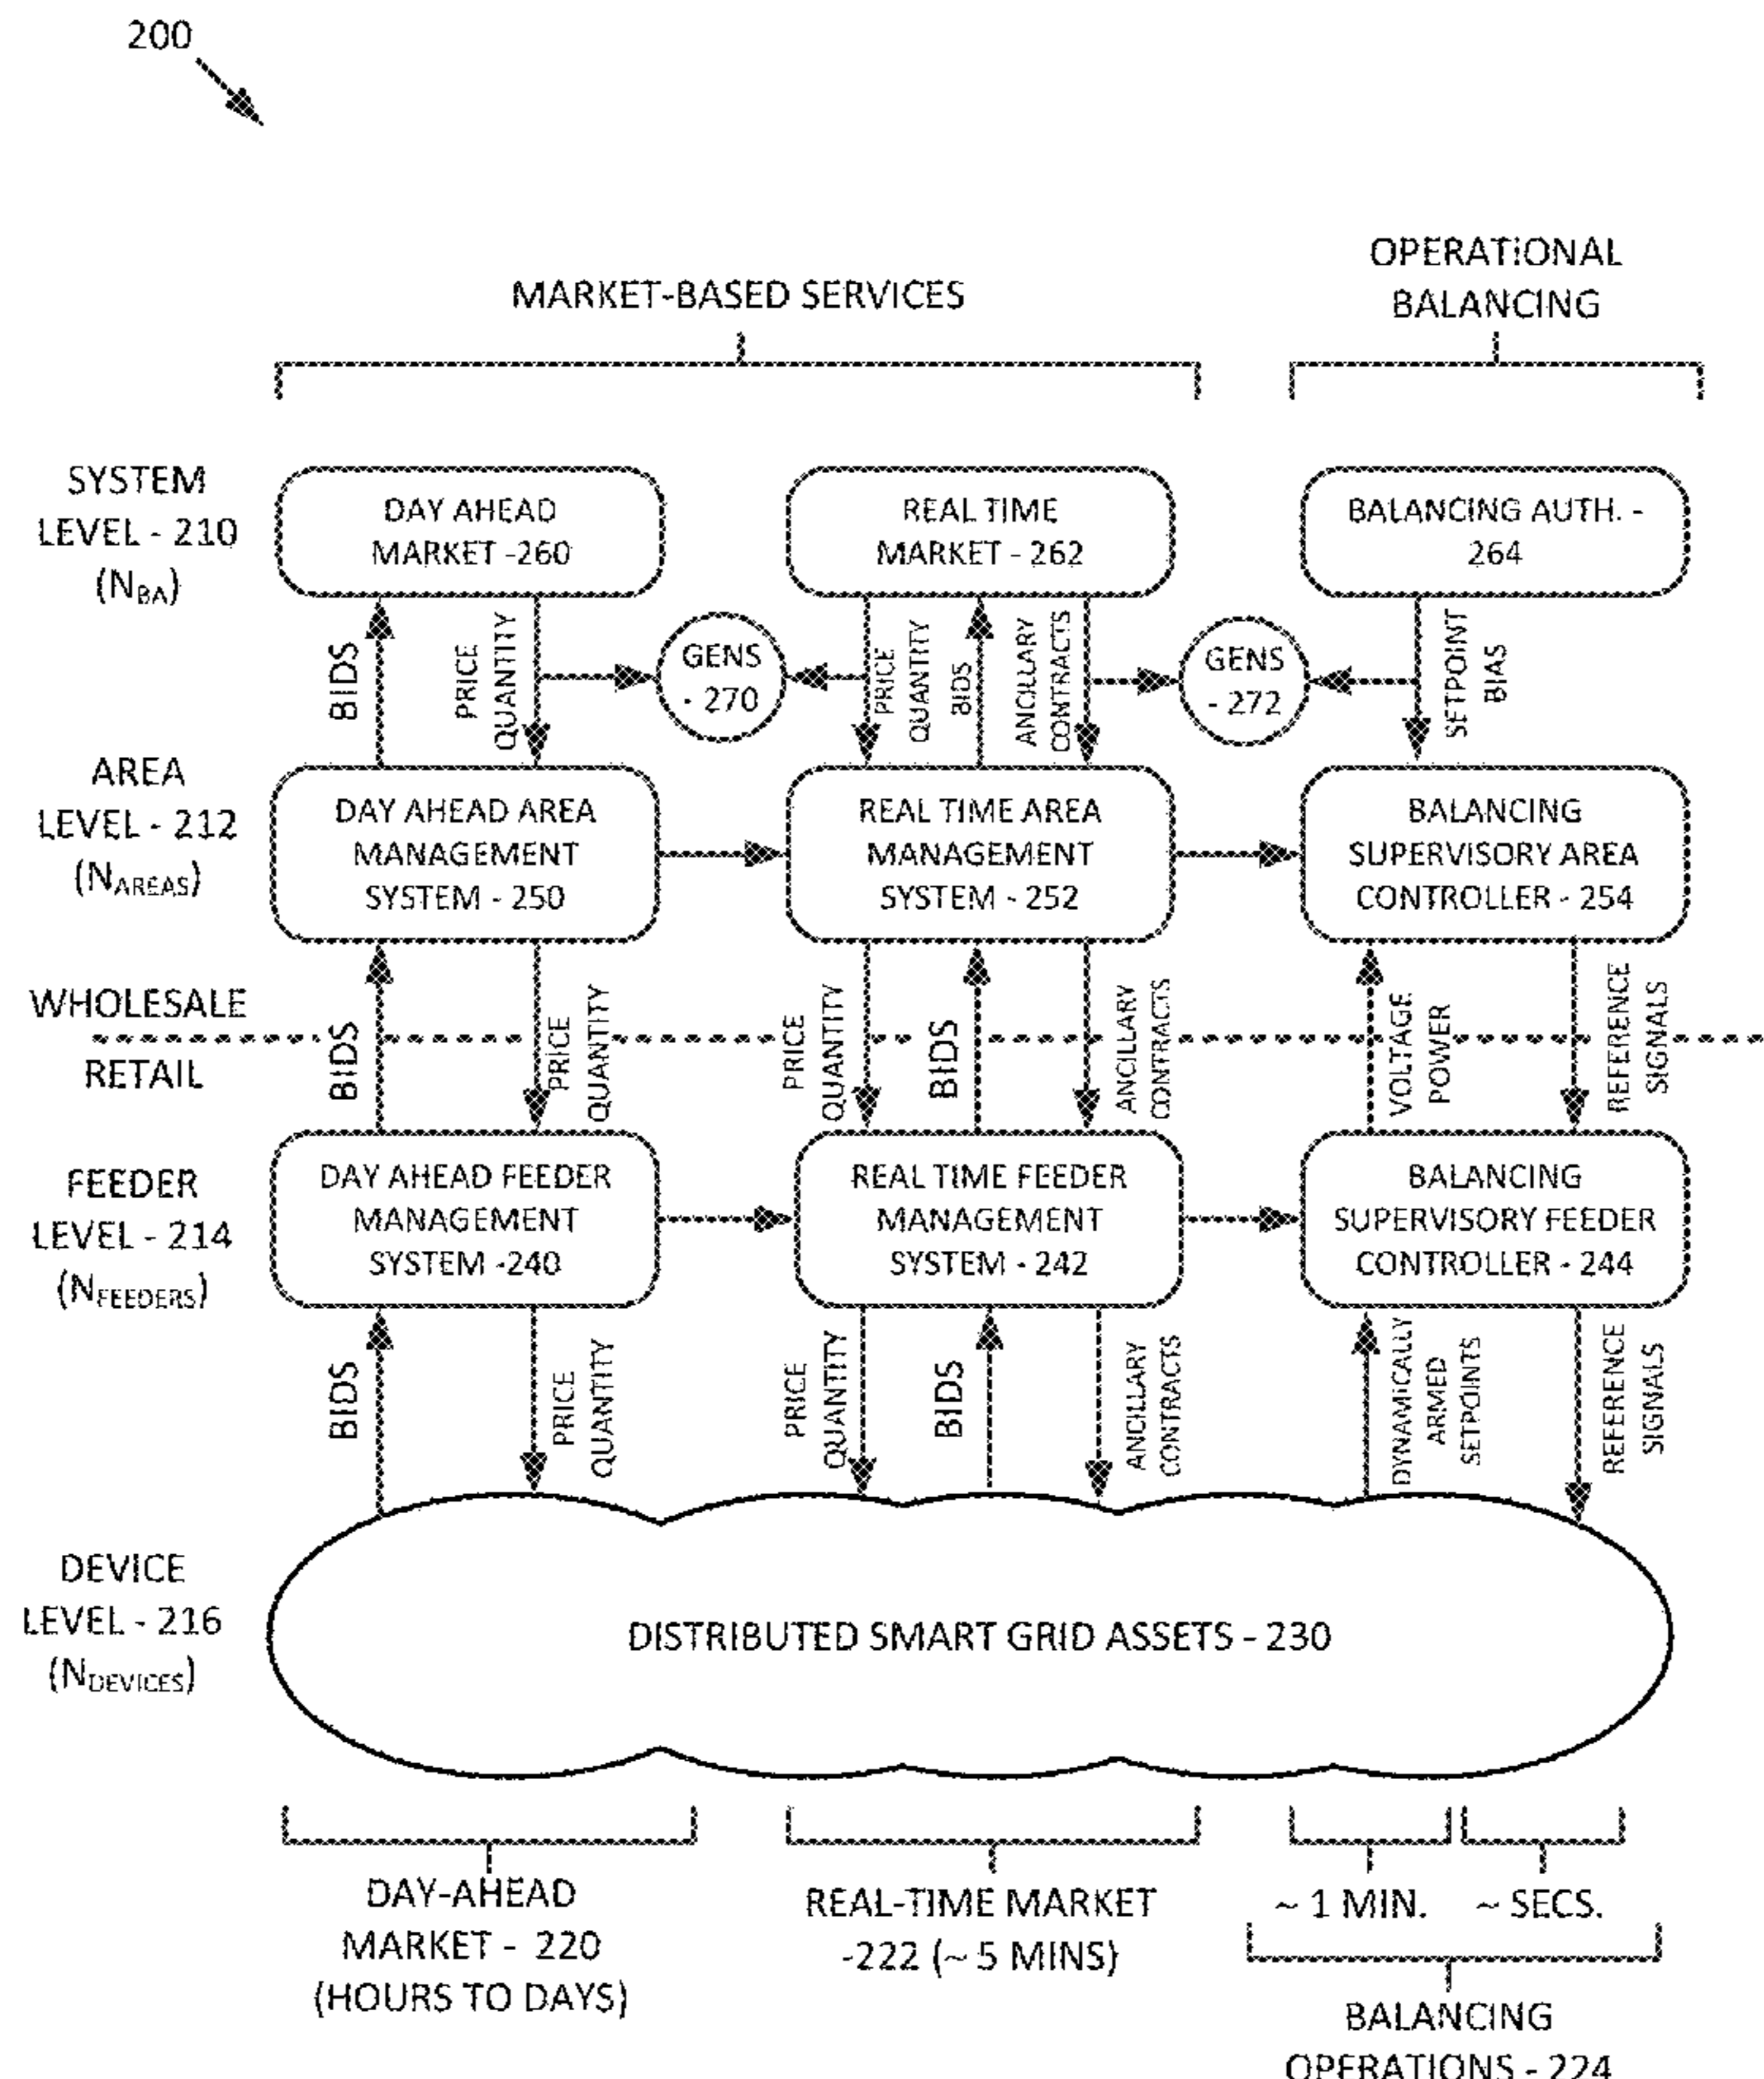
(22) Filed: **Nov. 1, 2019**

Related U.S. Application Data

(60) Continuation of application No. 15/693,413, filed on Aug. 31, 2017, now Pat. No. 10,498,141, which is a division of application No. 14/145,742, filed on Dec. 31, 2013, now Pat. No. 9,762,060.

(60) Provisional application No. 61/748,059, filed on Dec. 31, 2012.

Disclosed herein are representative embodiments of methods, apparatus, and systems for facilitating operation and control of a resource distribution system (such as a power grid). Among the disclosed embodiments is a distributed hierarchical control architecture (DHCA) that enables smart grid assets to effectively contribute to grid operations in a controllable manner, while helping to ensure system stability and equitably rewarding their contribution. Embodiments of the disclosed architecture can help unify the dispatch of these resources to provide both market-based and balancing services.



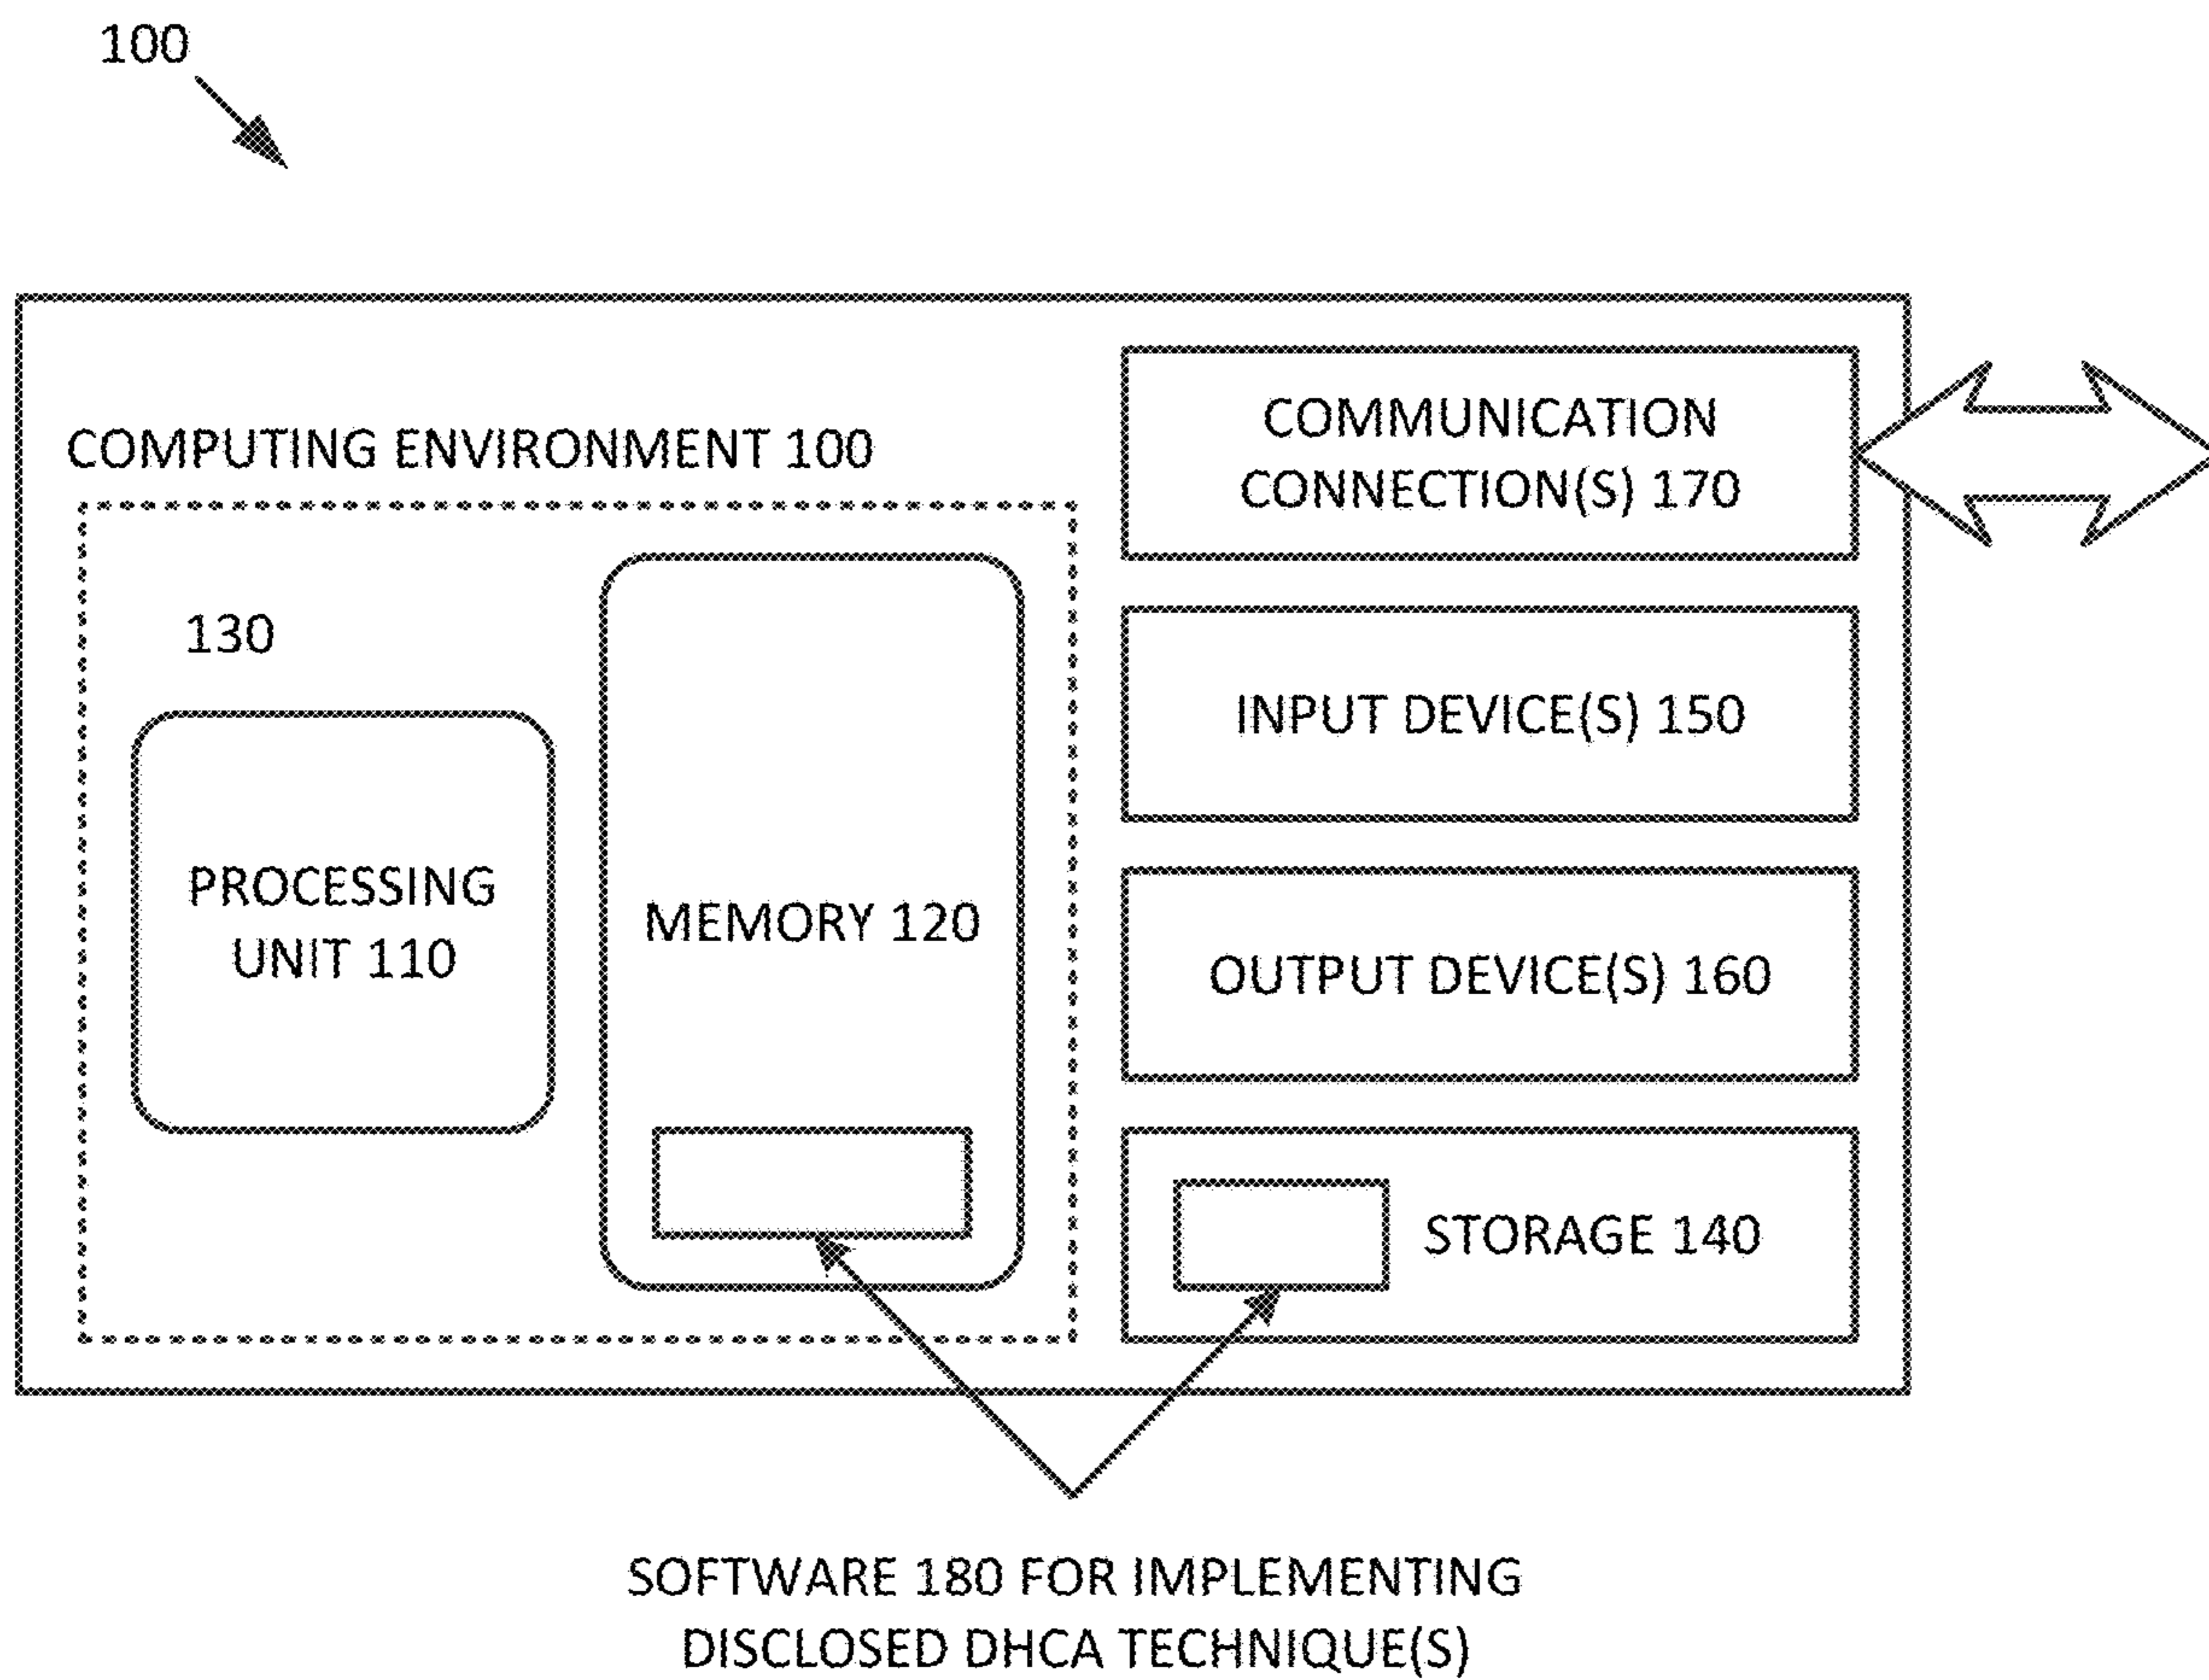


FIG. 1

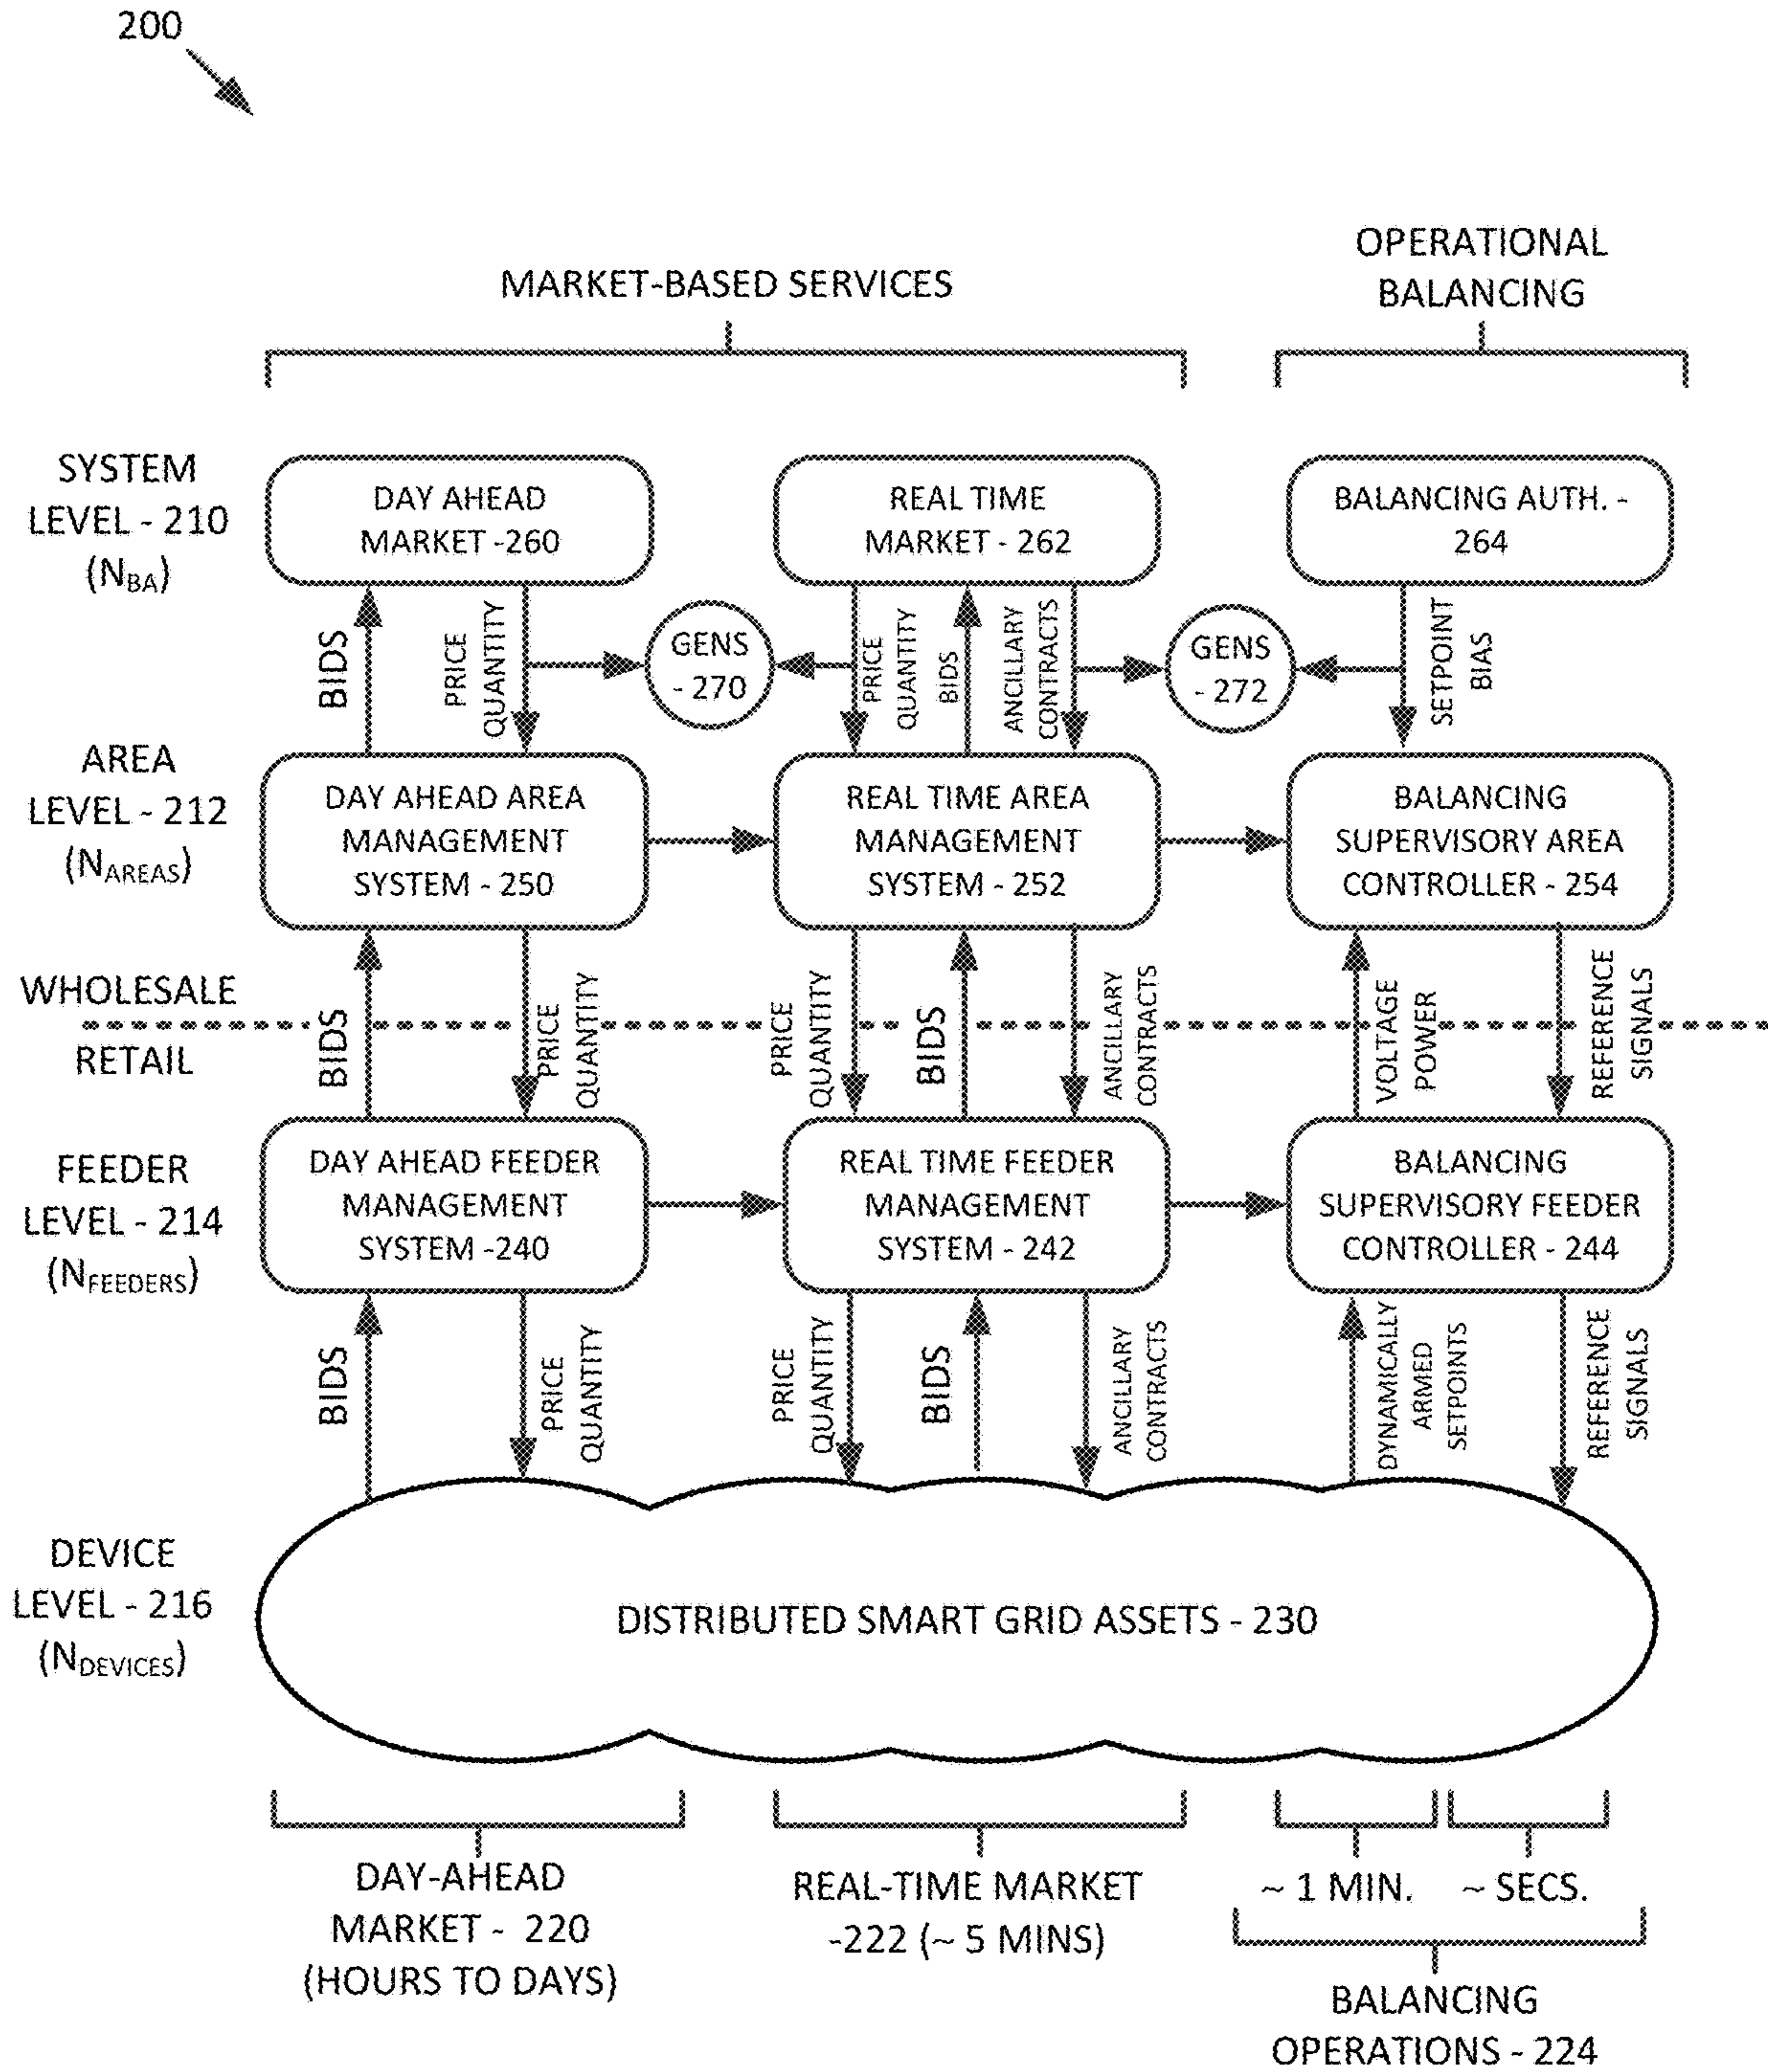


FIG. 2




FIG. 3

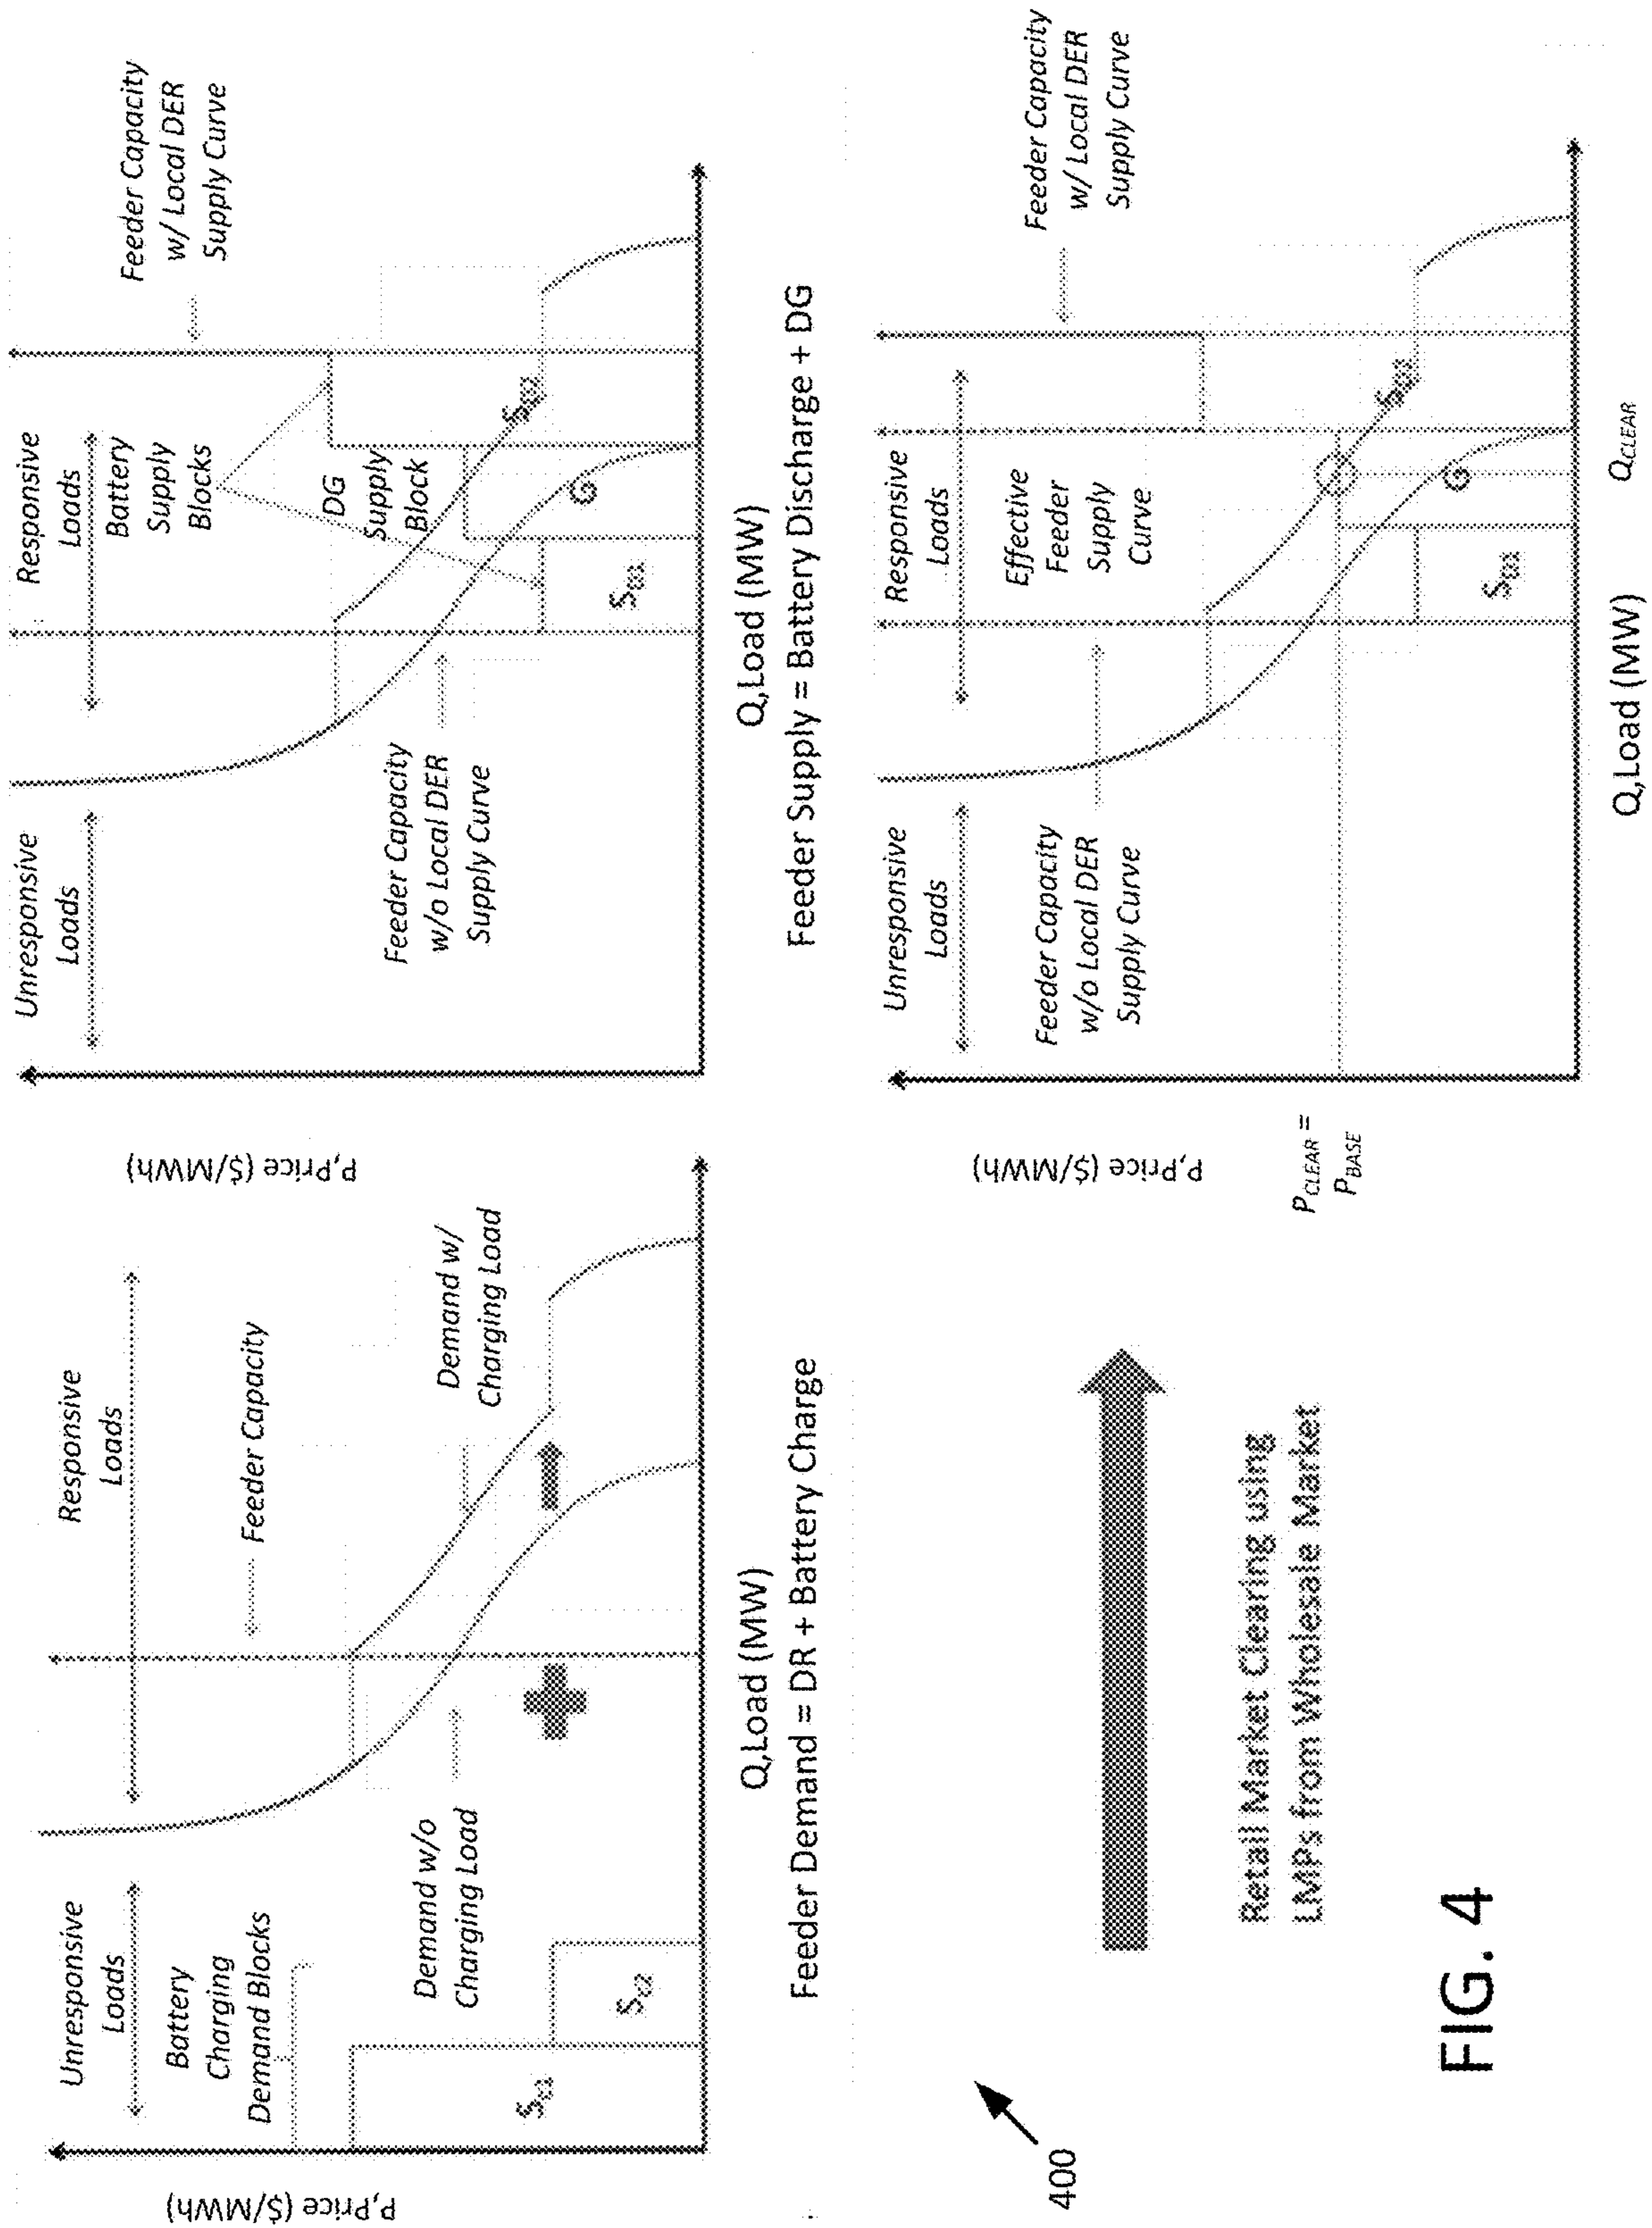


FIG. 4

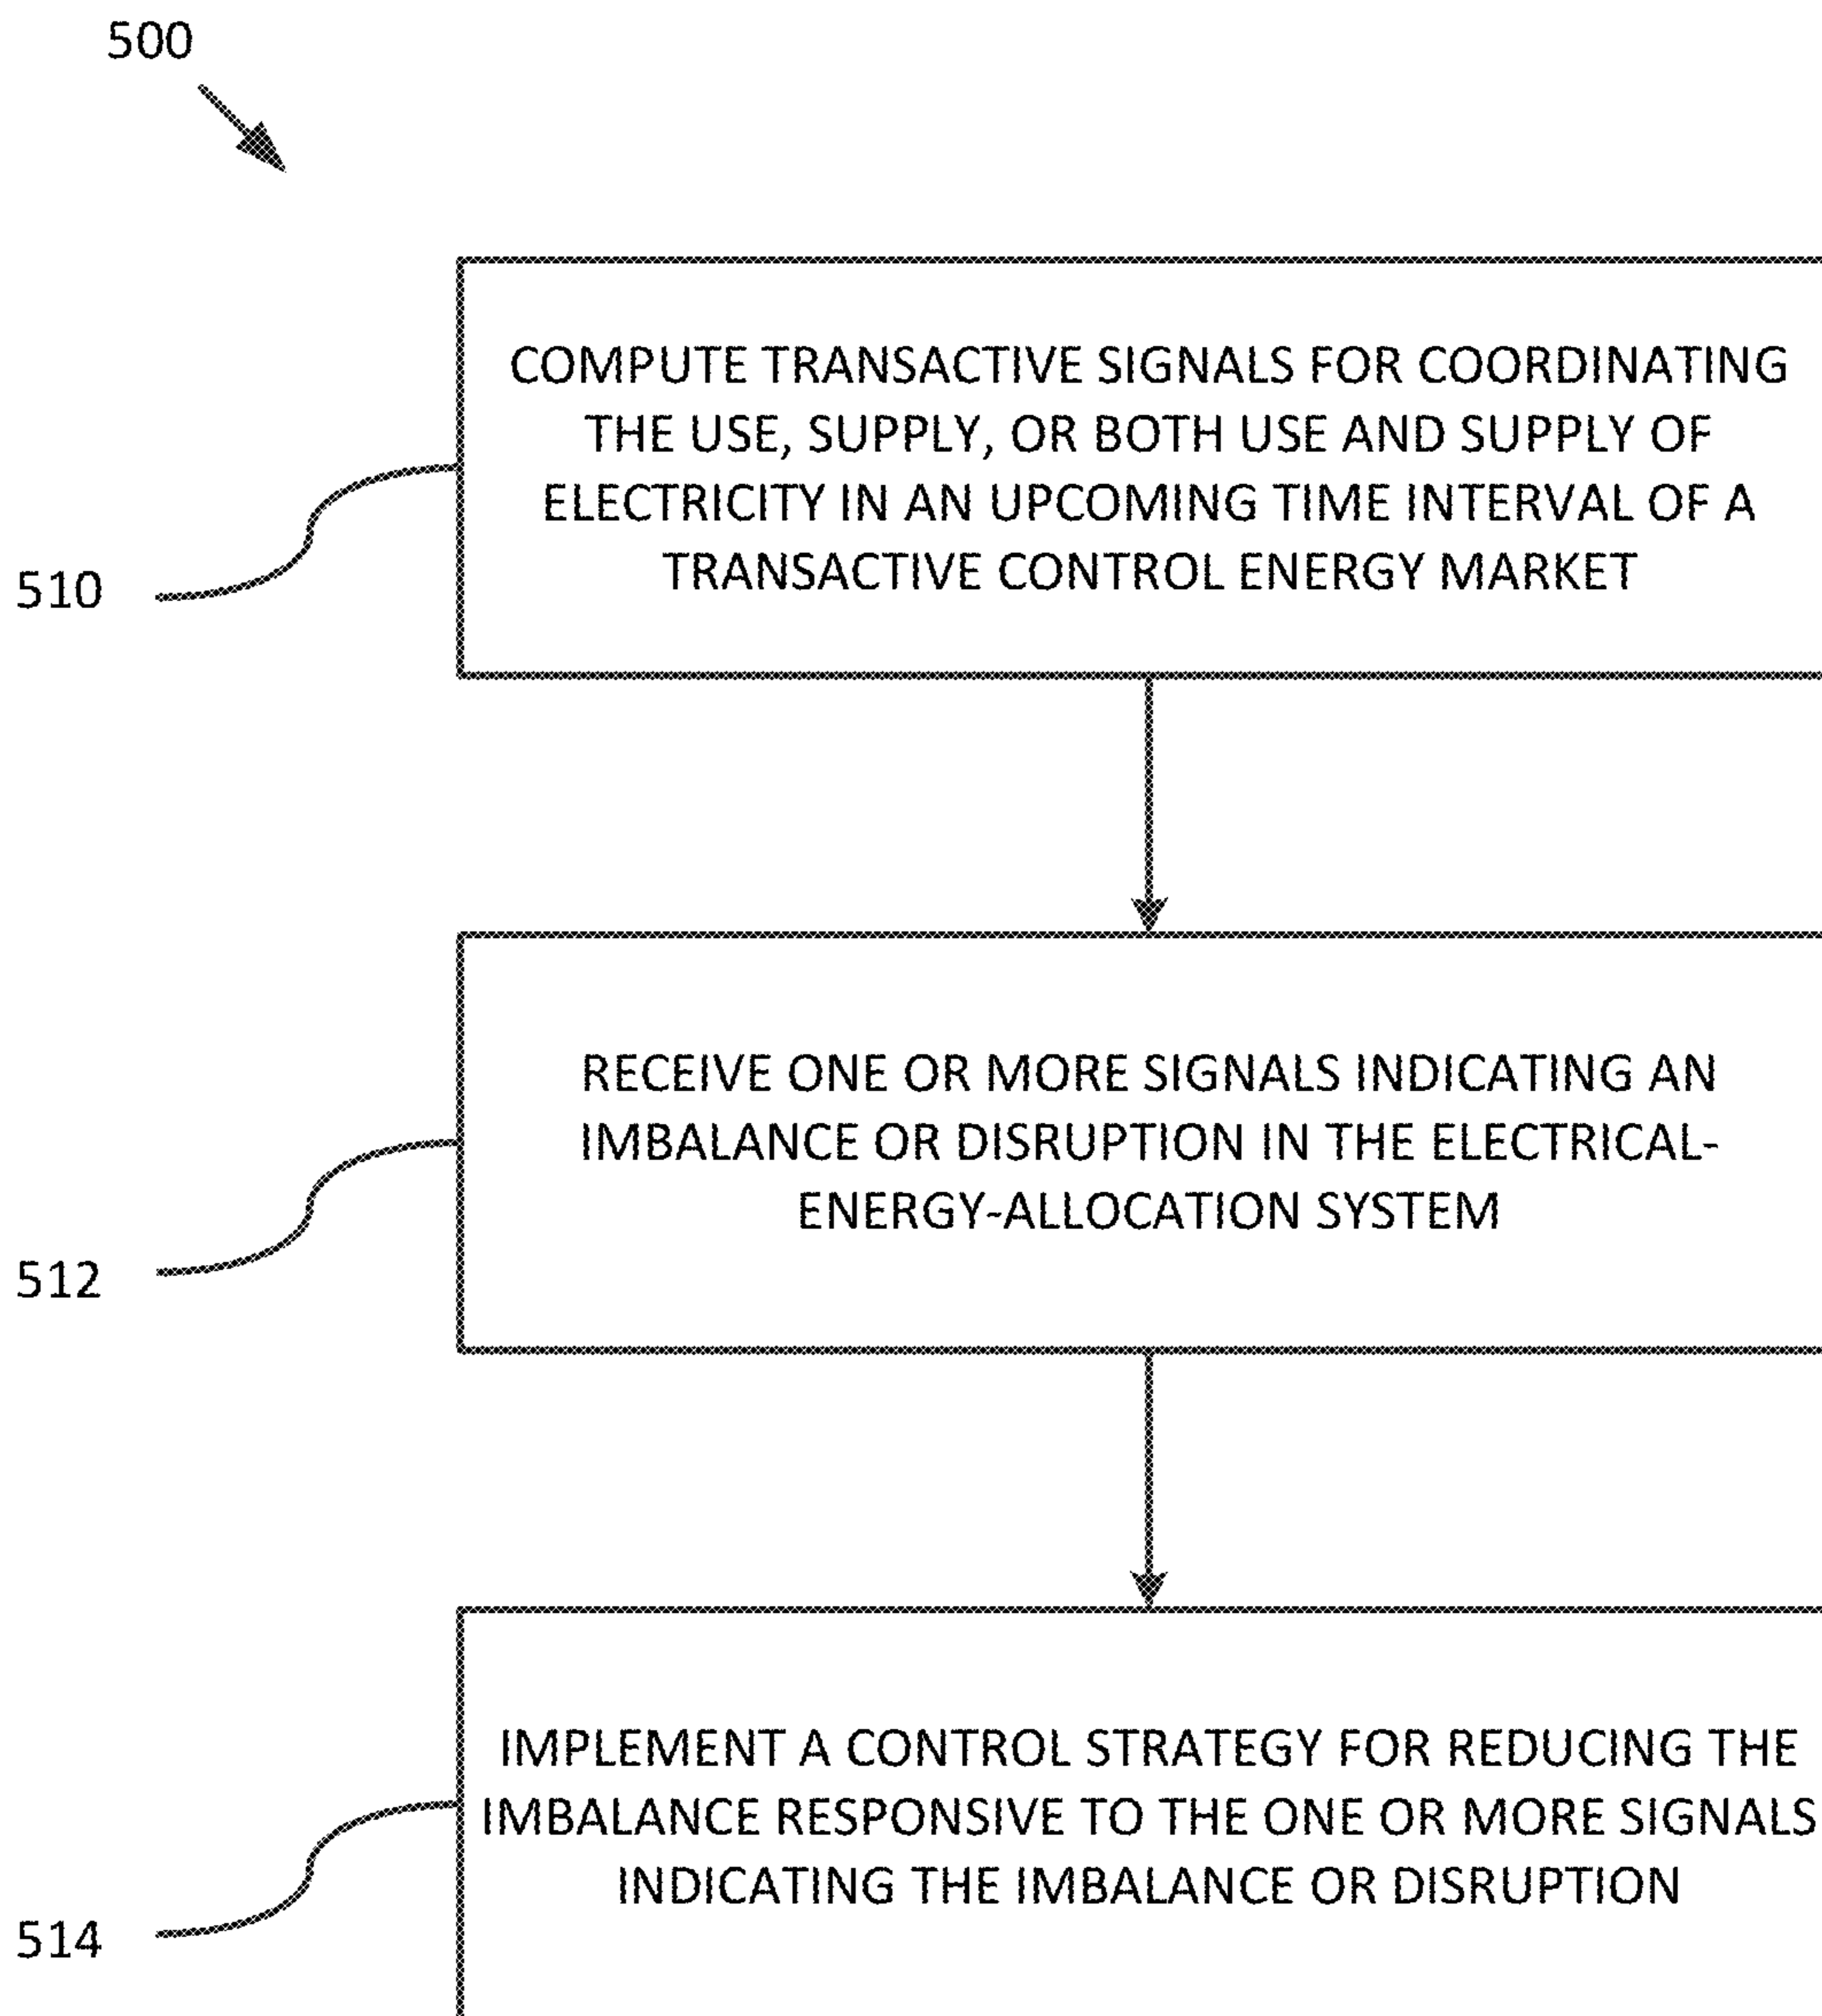


FIG. 5

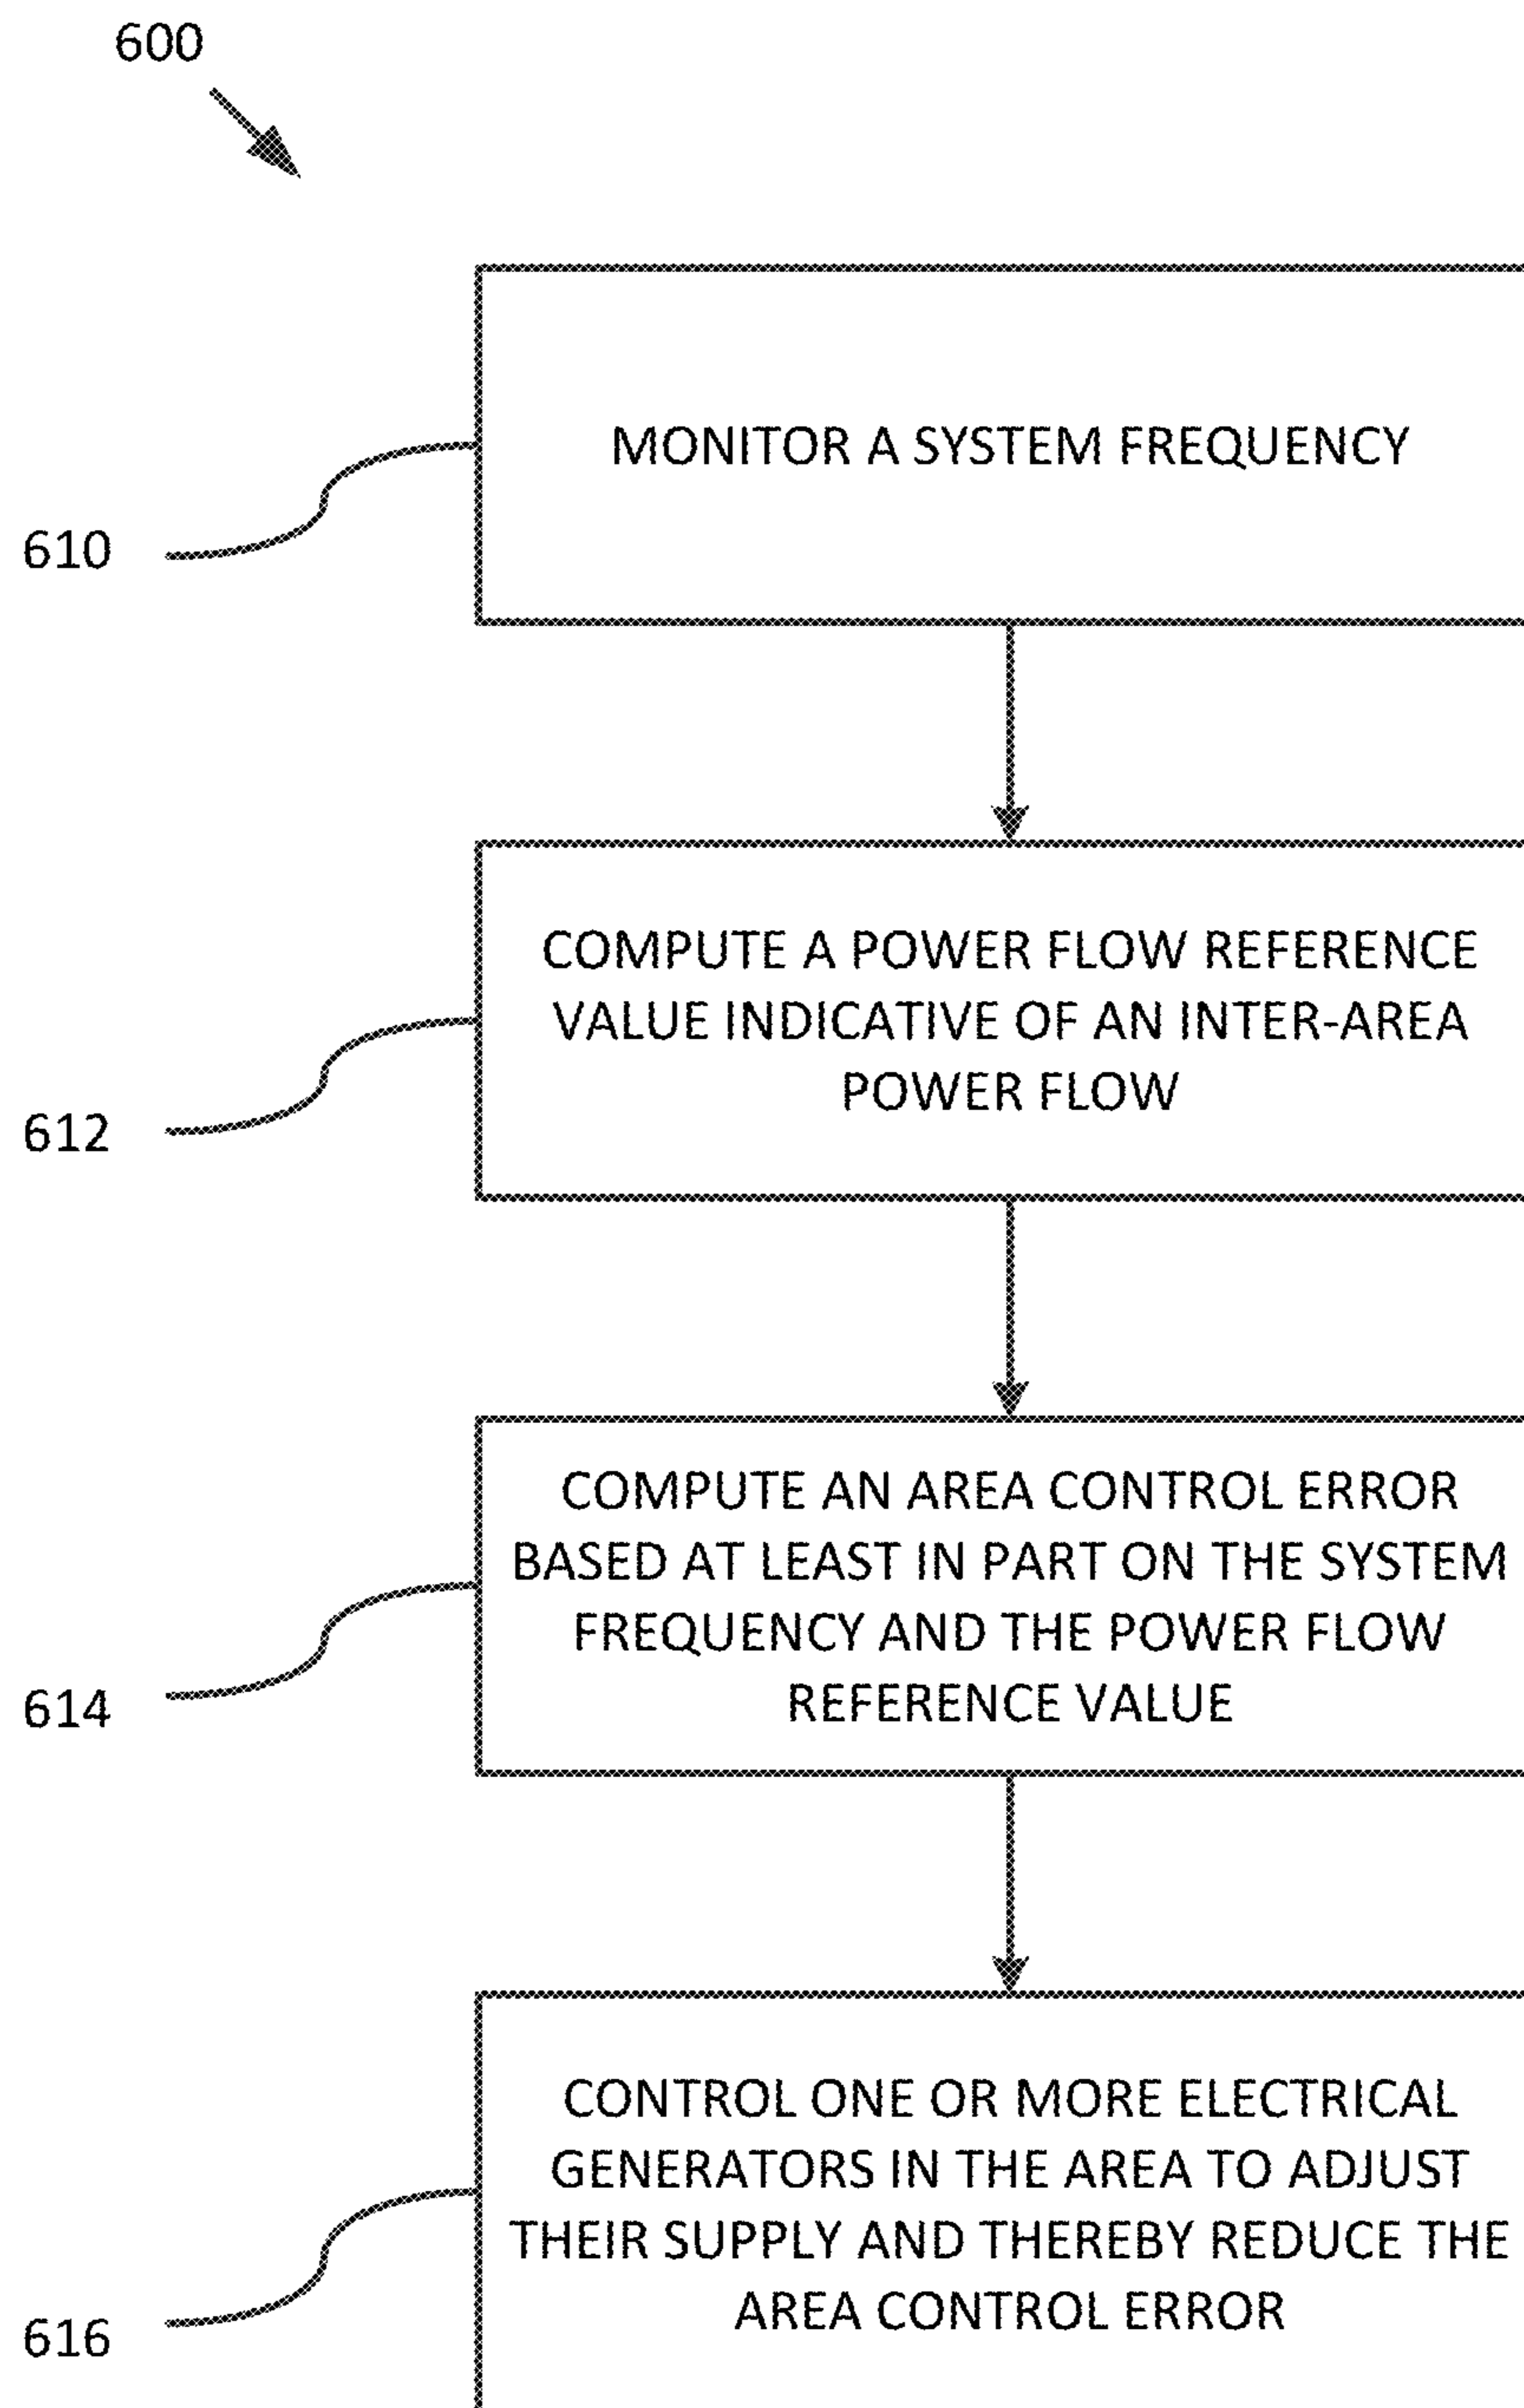


FIG. 6

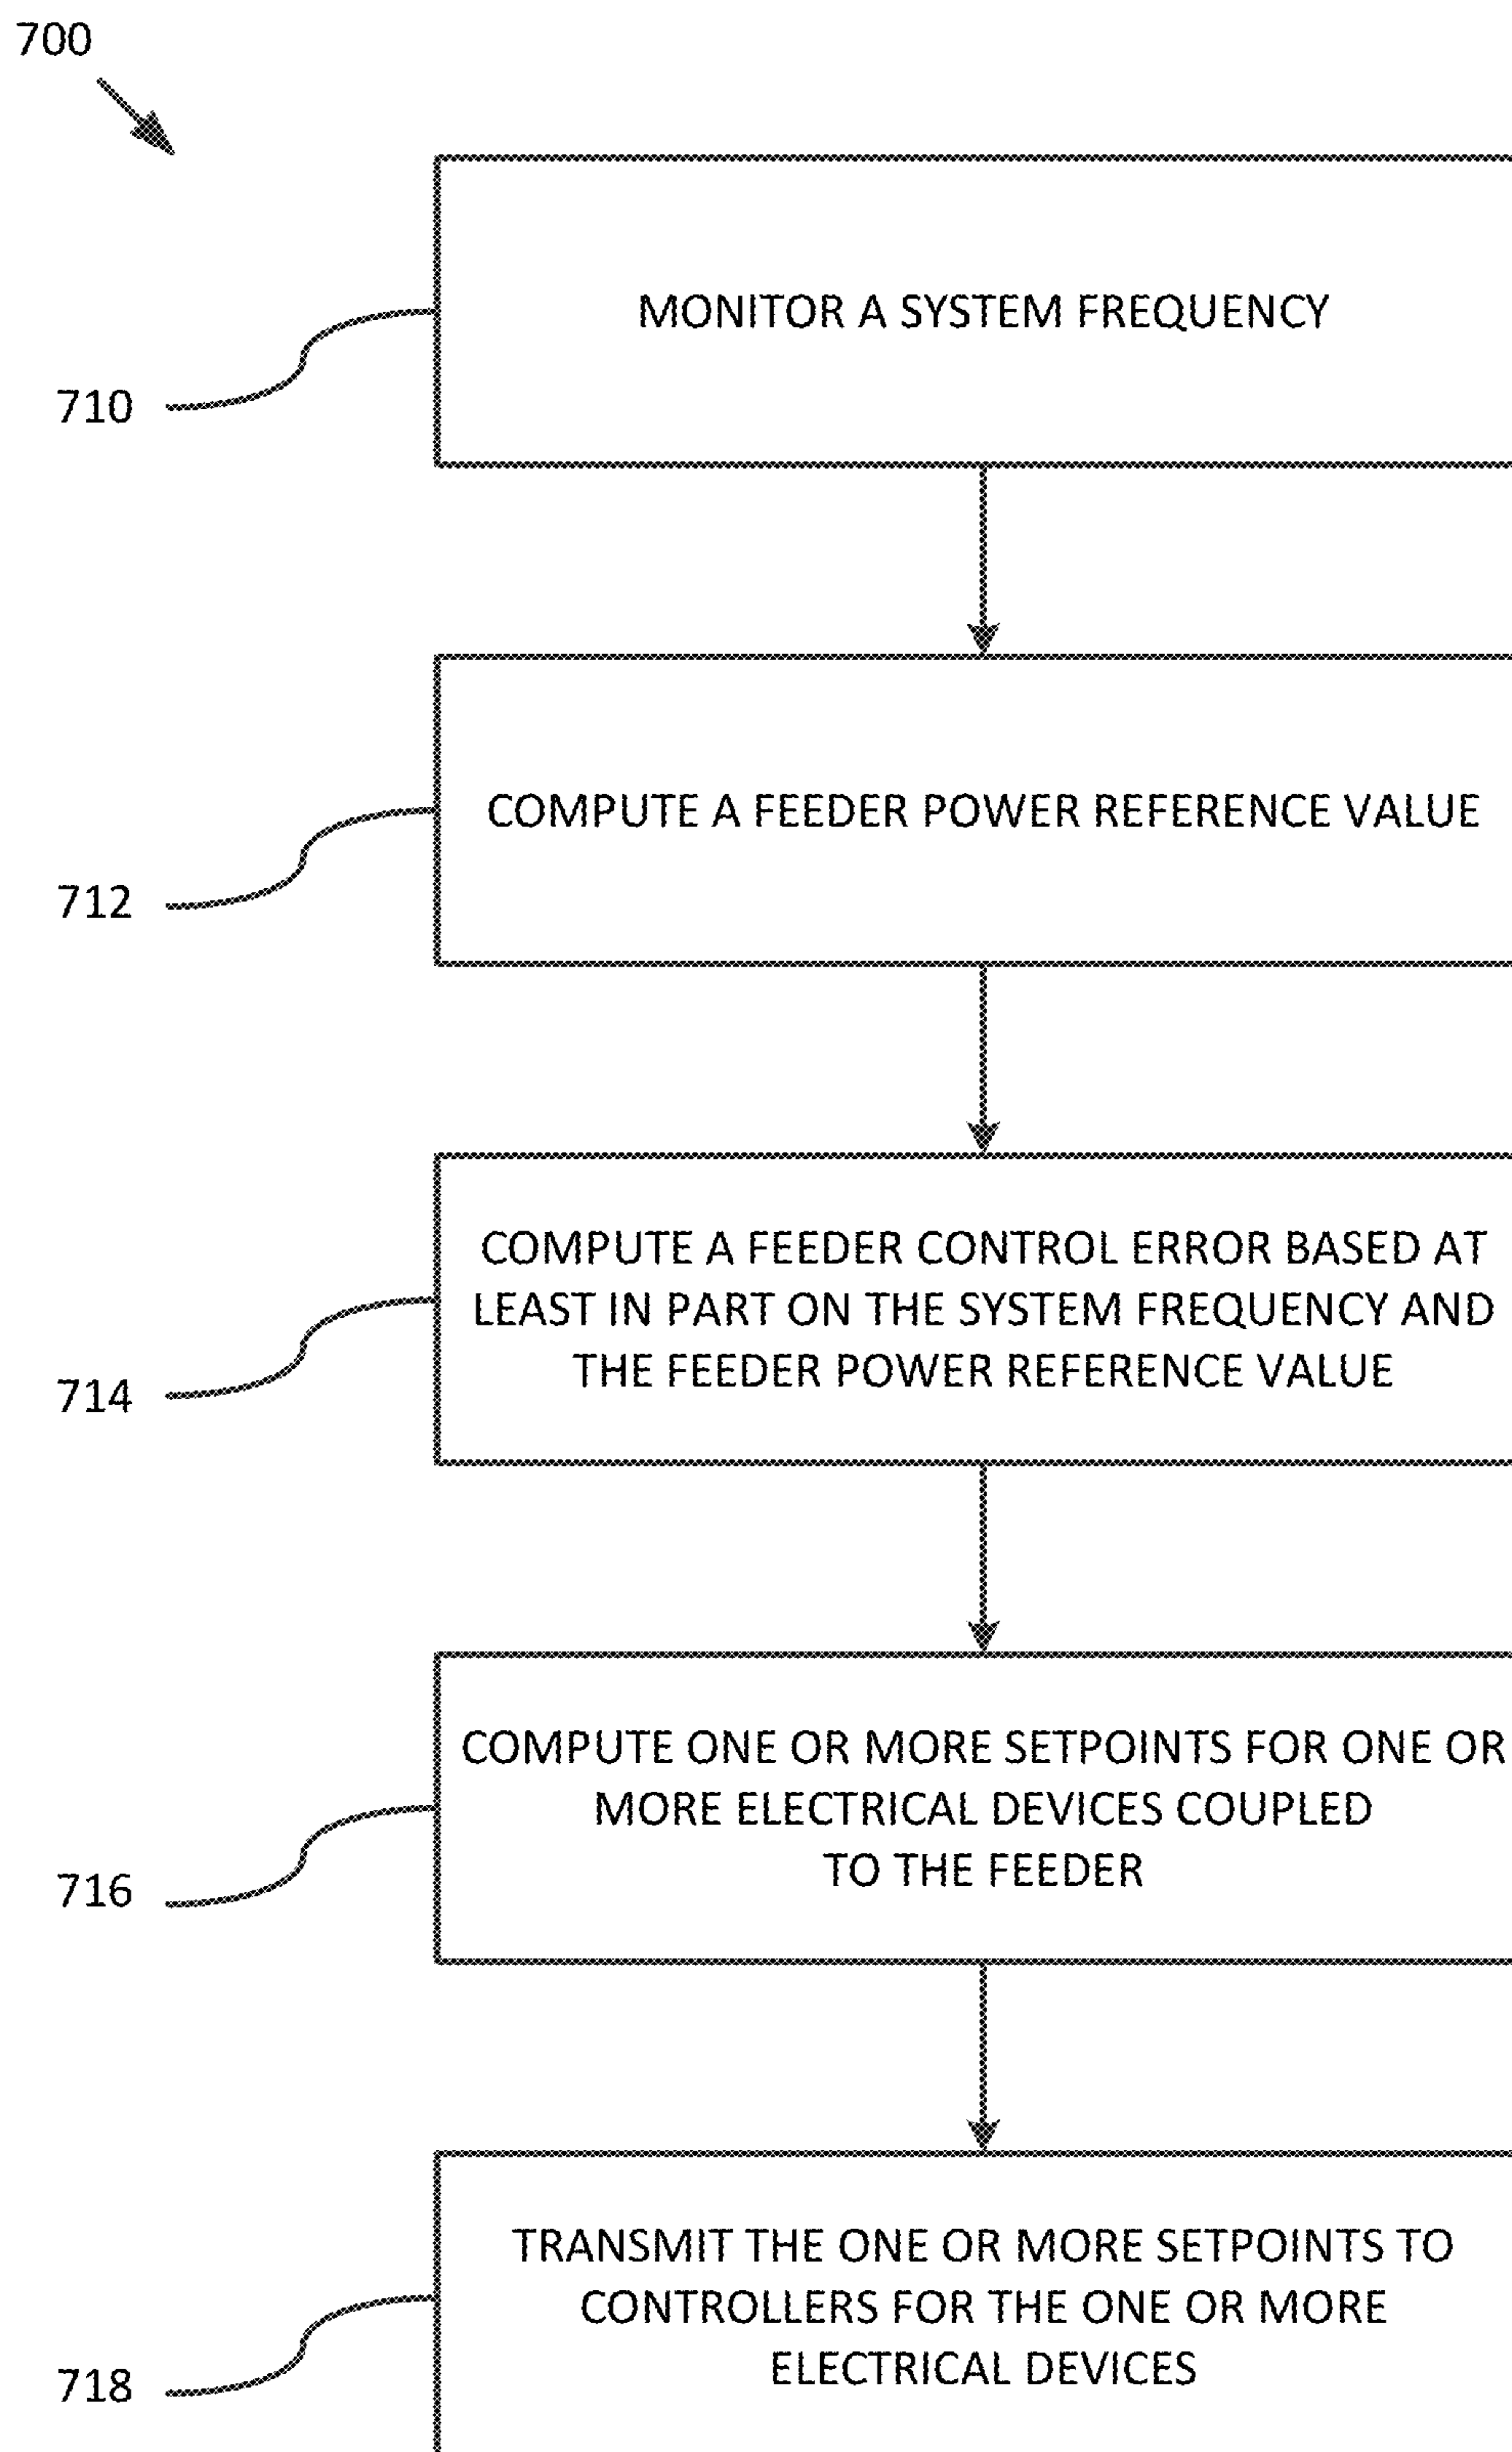


FIG. 7

**DISTRIBUTED HIERARCHICAL CONTROL
ARCHITECTURE FOR INTEGRATING
SMART GRID ASSETS DURING NORMAL
AND DISRUPTED OPERATIONS**

CROSS-REFERENCE TO RELATED
APPLICATIONS

[0001] This application claims the benefit of U.S. Provisional Application 61/748,059 filed on Dec. 31, 2013, and entitled “DISTRIBUTED HIERARCHICAL CONTROL ARCHITECTURE FOR INTEGRATING SMART GRID ASSETS DURING NORMAL AND DISRUPTED OPERATIONS”, which is hereby incorporated herein by reference in its entirety.

ACKNOWLEDGMENT OF GOVERNMENT
SUPPORT

[0002] This invention was made with government support under DE-AC05-76RL01830 awarded by the U.S. Department of Energy. The government has certain rights in the invention.

FIELD

[0003] This application relates generally to the field of power grid management and control.

SUMMARY

[0004] Disclosed below are representative embodiments of methods, apparatus, and systems for facilitating operation and control of a resource distribution system (such as a power grid). Among the disclosed embodiments herein is a distributed hierarchical control architecture (DHCA) that enables smart grid assets to effectively contribute to grid operations in a controllable manner, while helping to ensure system stability and equitably rewarding their contribution. The architecture can help unify the dispatch of these resources to provide both market-based and balancing services. Means to dynamically select and arm the autonomous responses from these assets are also disclosed, enabling the assets to offer significant reliability benefits under the full range of grid operating conditions. Further, transmission-level controls can be integrated with new and existing distribution-level control strategies, within a market structure, under both normal and disrupted operations (e.g., disrupted communications and other unforeseen events).

[0005] In certain embodiments, for example, a transactive scheme is used to clear markets at multiple levels (e.g., at one or more of an electronic device level (such as at the level of smart grid assets), feeder level, area level, or system level (such as at the level of an independent system operator (ISO) or regional transmission organization (RTO))). Further, in some embodiments, the framework is designed for clearing markets at multiple temporal scales (e.g., the framework is designed for coordinating day-ahead and real-time energy market). Still further, in some embodiments, the framework is designed for integrating retail and wholesale energy markets.

[0006] In some embodiments, the architecture is further configured to provide ancillary or balancing services. For example, the architecture can comprise a nested hierarchical set of co-optimized energy and balancing markets. In operation, embodiments of the disclosed technology can be used to help ensure continued operations of the asset network

under disrupted conditions (e.g., when communications are incomplete or unavailable) and during grid contingencies when distributed smart grid assets (DSGAs) may be useful elements of local- and wide-area recovery schemes.

[0007] In some example implementations, the overall hierarchical control architecture is divided into four levels: device, feeder, area, and system levels. The architecture can also integrate operations across multiple time scales: day-ahead (24-hr), real-time (–5-min), and near instantaneous (–sec) balancing requirements (e.g., regulation, spinning reserve, etc.). At the device level, decentralized control schemes allow distributed smart grid assets to provide both economic and reliability responses. The control schemes can be based, for example, on the self-sensing of frequency, voltage, broadcasts of the imbalance signal, current and future prices, and/or device conditions like state-of-charge (batteries) or temperature and setpoint (thermostats, smart appliances, etc.) during normal operation. The distributed smart grid assets can also be dynamically armed via centralized control signals (setpoints) while responding autonomously and instantaneously.

[0008] In some embodiments, a balancing supervisory area controller solves an area wide optimization problem, including bids and then allocates a portion of its requirements to the feeder controller. It also works to maintain adequate system-wide support in terms of frequency and tie-line flows. A balancing feeder controller then dispatches setpoints (e.g., –sec to 1-min) to the distributed smart grid assets and primary devices to meet the requirements of the balancing area controller and the market-based feeder management system.

[0009] In certain embodiments, market-based feeder and area management systems aim to minimize or reduce total costs in the real-time and day-ahead markets by incorporating the smart grid resources into standard ISO/RTO market structures. This can be done, for example, by solving a large-scale optimization problem subject to local and area operational constraints and uncertainties of renewables and distributed smart grid assets. Demonstrations have been performed that showed how devices can bid their capabilities into local feeder-level markets in a transactive scheme that reflects both wholesale prices and capacity constraints. Embodiments of the disclosed technology can close the loop by affecting the price at the system or area levels. At the system level, for example, the ISOs and BAs receive aggregated net load demand, supply bids for smart grid assets from the real-time and day-ahead area management systems, and generator power supply offers from generator aggregators. In certain implementations, the cleared or scheduled power setpoints are then dispatched to the area controllers, which in turn dispatch requirements or price signals to the feeder controllers and distributed smart grid assets.

[0010] Embodiments of the disclosed technology can have a direct impact in the efficiency of generation and delivery of electric power, in the reduction of greenhouse gas emissions from the electric power system, and/or in ensuring that the U.S. maintains its leadership in developing and deploying advanced energy technologies. The disclosed technology can be used, for example, to effectively harness the potential of smart grid approaches to affect actual, moment-by-moment grid operations in very substantial ways. Embodiments of the disclosed technology overcome the balkanized approaches being pursued by various elements of the industry by providing a unified approach to obtaining increased

benefits at multiple levels of the grid infrastructure from many types of distributed smart grid assets once they are purchased, whether they are owned by utilities, customers, or third parties. As such, embodiments of the disclosed technology can be used to form the basic, overarching architecture that can give focus to priority, minimal (as opposed to endlessly expansive) communications, and interoperability standards. Such a unified approach can significantly help realize the transformation envisioned by the smart grid.

[0011] Embodiments of the disclosed technology can help alleviate concerns of system planners and operators regarding controllability of distributed smart grid assets, allowing them to be fully incorporated into system operations to achieve multiple objectives. Among the possible benefits that can be realized by embodiments of the disclosed technology are one or more of the following:

[0012] Higher utilization of generation, transmission, and distribution assets, by reducing peak loads;

[0013] Lower wholesale market costs or power production costs, especially during high price periods, from DSGAs response to wholesale prices;

[0014] Lower ancillary service costs by engaging distributed assets to supply them; and/or

[0015] Lower cost for integrating new solar and wind generation them into system operations by mitigating their variability and uncertainty.

[0016] Another impact that can be realized in embodiments of the disclosed technology is the increased reliability at both the bulk grid and distribution levels, from coordinating the engagement of distributed assets by multiple operating entities by:

[0017] Increasing available reserve margins;

[0018] Incorporating them into bulk grid wide-area control schemes; and

[0019] Integrating them with distribution level voltage control and reconfiguration schemes.

[0020] Embodiments of the disclosed technology can also increase the penetration of demand response and other distributed assets by allowing them to provide a complete set of services on par with traditional large-scale generation plants. The increased penetration of demand response and other distributed assets can also help to keep utility revenue requirements, and hence customer rates, as low as possible by utilizing lower cost distributed assets to displace the need for additional traditional infrastructure, thereby increasing the overall cost effectiveness of the grid infrastructure. By minimizing (or otherwise reducing) the information content of data transferred, embodiments of the distributed hierarchical control architecture can also enhance cyber-security and customer privacy.

[0021] Embodiments of the disclosed methods can be performed using computing hardware, such as a computer processor or an integrated circuit. For example, embodiments of the disclosed methods can be performed by software stored on one or more non-transitory computer-readable media (e.g., one or more optical media discs, volatile memory components (such as DRAM or SRAM), or non-volatile memory or storage components (such as hard drives or solid state drives (e.g., solid state drives based on flash memory))). Such software can be executed on a single computer or on a networked computer (e.g., via the Internet, a wide-area network, a local-area network, a client-server network, a cloud-based network, or other such network).

Embodiments of the disclosed methods can also be performed by specialized computing hardware (e.g., one or more application specific integrated circuits (ASICs) or programmable logic devices (such as field programmable gate arrays (FPGAs)) configured to perform any of the disclosed methods). Additionally, any intermediate or final result created or modified using any of the disclosed methods can be stored on a non-transitory storage medium (e.g., one or more optical media discs, volatile memory or storage components (such as DRAM or SRAM), or nonvolatile memory or storage components (such as hard drives)) and are considered to be within the scope of this disclosure. Furthermore, any of the software embodiments (comprising, for example, computer-executable instructions which when executed by a computer cause the computer to perform any of the disclosed methods), intermediate results, or final results created or modified by the disclosed methods can be transmitted, received, or accessed through a suitable communication means.

[0022] The foregoing and other objects, features, and advantages of the invention will become more apparent from the following detailed description, which proceeds with reference to the accompanying figures.

BRIEF DESCRIPTION OF THE DRAWINGS

[0023] FIG. 1 is a block diagram illustrating an example computing hardware environment which can be used to implement embodiments of the disclosed technology.

[0024] FIG. 2 is a block diagram of a system for coordinating distribution of electricity according to a distributed hierarchical control architecture (DHCA)

[0025] FIG. 3 is a block diagram of a graph that represents an HVAC bidding/response strategy that translates real-time price into a thermostat setpoint, and current air temperature into a demand bid.

[0026] FIG. 4 is a block diagram of an example retail market clearing process.

[0027] FIG. 5 is a flowchart showing an exemplary embodiment for operating a transactive controller in a market-based electrical-energy-allocation system.

[0028] FIG. 6 is a flowchart showing an exemplary embodiment for operating an area controller configured to coordinate the distribution of electricity to and from multiple feeder networks of an area.

[0029] FIG. 7 is a flowchart showing an exemplary embodiment for operating a feeder controller configured to coordinate the distribution of electricity to and from multiple electrical devices (e.g., distributed smart grid assets) coupled to a feeder controlled by the feeder controller.

DETAILED DESCRIPTION

I. General Considerations

[0030] Disclosed below are representative embodiments of methods, apparatus, and systems for facilitating operation and control of a resource distribution system (such as a power grid). The disclosed methods, apparatus, and systems should not be construed as limiting in any way. Instead, the present disclosure is directed toward all novel and non-obvious features and aspects of the various disclosed embodiments, alone and in various combinations and subcombinations with one another. Furthermore, any one or more features or aspects of the disclosed embodiments can be

used alone or in various combinations and subcombinations with one another. The disclosed methods, apparatus, and systems are not limited to any specific aspect or feature or combination thereof, nor do the disclosed embodiments require that any one or more specific advantages be present or problems be solved.

[0031] Although the operations of some of the disclosed methods are described in a particular, sequential order for convenient presentation, it should be understood that this manner of description encompasses rearrangement, unless a particular ordering is required by specific language set forth below. For example, operations described sequentially may in some cases be rearranged or performed concurrently. Moreover, for the sake of simplicity, the attached figures may not show the various ways in which the disclosed methods can be used in conjunction with other methods. Additionally, the description sometimes uses terms like “determine” and “generate” to describe the disclosed methods. These terms are high-level abstractions of the actual operations that are performed. The actual operations that correspond to these terms may vary depending on the particular implementation and are readily discernible by one of ordinary skill in the art. Furthermore, in general, and as used herein, the term “optimal” describes a solution that satisfies some set of criteria better than other solutions according to some parameterization or modeling, which may or may not be optimal in absolute terms depending on circumstances, and the term “optimize” or “optimization” is used to indicate the process of finding such a solution. Furthermore, as used herein, the term “and/or” means any one item or combination of items in the phrase.

[0032] Any of the embodiments disclosed herein can be used to with any of the transactive control schemes and architectures described in U.S. Nonprovisional application Ser. No. 12/587,008 filed on Sep. 29, 2009, and entitled “ELECTRIC POWER GRID CONTROL USING A MARKET-BASED RESOURCE ALLOCATION SYSTEM,” (published as U.S. Patent Application Publication No. 2010/0114387); U.S. Nonprovisional application Ser. No. 12/686,243 filed on Jan. 12, 2010, and entitled “NESTED, HIERARCHICAL RESOURCE ALLOCATION SCHEMA FOR MANAGEMENT AND CONTROL OF AN ELECTRIC POWER GRID” (published as U.S. Patent Application Publication No. 2010/0179862); and U.S. Nonprovisional application Ser. No. 14/108,078 filed on Dec. 16, 2013, and entitled “TRANSACTIVE CONTROL AND COORDINATION FRAMEWORK AND ASSOCIATED TOOLKIT FUNCTIONS”, all of which are hereby incorporated herein by reference in their entirety.

[0033] Any of the disclosed methods can be implemented using computer-executable instructions stored on one or more computer-readable media (e.g., non-transitory computer-readable media, such as one or more optical media discs, volatile memory components (such as DRAM or SRAM), or nonvolatile memory components (such as hard drives)) and executed by a processor in a computing device (e.g., a computer, such as any commercially available computer). Any of the computer-executable instructions for implementing the disclosed techniques as well as any intermediate or final data created and used during implementation of the disclosed systems can be stored on one or more computer-readable media (e.g., non-transitory computer-readable media). The computer-executable instructions can be part of, for example, a dedicated software application or

as part of a software agent’s transport payload that is accessed or downloaded via a network (e.g., a local-area network, a wide-area network, a client-server network, or other such network).

[0034] Such software can be executed on a single computer (e.g., a computer embedded in or electrically coupled to a sensor, controller, or other device in the power grid) or in a network environment. For example, the software can be executed by a computer embedded in or communicatively coupled to a sensor for measuring electrical parameters of a power line, a synchrophasor sensor, a smart meter, a control unit for a home or household appliance or system (e.g., an air-conditioning unit; heating unit; heating, ventilation, and air conditioning (HVAC) system; hot water heater; refrigerator; dish washer; washing machine; dryer; oven; microwave oven; pump; home lighting system; electrical charger; electric vehicle charger; home electrical system; or any other electrical system having variable performance states), a control unit for a distributed generator (e.g., photovoltaic arrays, wind turbines, or electric battery charging systems), a control unit for controlling the distribution or generation of power along the power grid (e.g., a transformer, switch, circuit breaker, generator, resource provider, or any other device on the power grid configured to perform a control action), and the like. These household appliances and distributed generators (along with any associated control unit) are examples of the “electrical devices”, “distributed smart grid assets”, “DSGAs”, “distributed assets”, or “assets” discussed below. As more fully explained below, these devices can be controlled and coordinated in embodiments of the disclosed DHCA.

[0035] For clarity, only certain selected aspects of the software-based embodiments are described. Other details that are well known in the art are omitted. For example, it should be understood that the software-based embodiments are not limited to any specific computer language or program. For instance, embodiments of the disclosed technology can be implemented by software written in C++, Java, Perl, JavaScript, Adobe Flash, Python, JINI, .NET, Lua or any other suitable programming language. Likewise, embodiments of the disclosed technology are not limited to any particular computer or type of hardware. Details of suitable computers and hardware are well known and need not be set forth in detail in this disclosure.

[0036] Furthermore, any of the software-based embodiments (comprising, for example, computer-executable instructions which when executed by a computer cause the computer to perform any of the disclosed methods) can be uploaded, downloaded, or remotely accessed through a suitable communication means. Such suitable communication means include, for example, the Internet, the World Wide Web, an intranet, software applications, cable (including fiber optic cable), magnetic communications, electromagnetic communications (including RF, microwave, and infrared communications), electronic communications, or other such communication means.

[0037] The disclosed methods can also be implemented by specialized computing hardware that is configured to perform any of the disclosed methods. For example, the disclosed methods can be implemented by a computing device comprising an integrated circuit (e.g., an application specific integrated circuit (ASIC) or programmable logic device (PLD), such as a field programmable gate array (FPGA)). The integrated circuit or specialized computing hardware

can be embedded in or directly coupled to a sensor, control unit, or other device in the power grid. For example, the integrated circuit can be embedded in or otherwise coupled to a synchrophasor sensor, smart meter, control unit for a home or household appliance or system, a control unit for a distributed generator, a control unit for controlling power distribution on the grid, or other such device.

[0038] FIG. 1 illustrates a generalized example of a suitable computing hardware environment **100** for a computing device with which several of the described embodiments can be implemented. The computing environment **100** is not intended to suggest any limitation as to the scope of use or functionality of the disclosed technology, as the techniques and tools described herein can be implemented in diverse general-purpose or special-purpose environments that have computing hardware.

[0039] With reference to FIG. 1, the computing environment **100** includes at least one processing unit **110** and memory **120**. In FIG. 1, this most basic configuration **130** is included within a dashed line. The processing unit **110** executes computer-executable instructions. In a multi-processing system, multiple processing units execute computer-executable instructions to increase processing power. The memory **120** may be volatile memory (e.g., registers, cache, RAM), non-volatile memory (e.g., ROM, EEPROM, flash memory), or some combination of the two. The memory **120** stores software **180** for implementing one or more of the described techniques for operating or using the disclosed systems. For example, the memory **120** can store software **180** for implementing any of the disclosed techniques.

[0040] The computing environment can have additional features. For example, the computing environment **100** includes storage **140**, one or more input devices **150**, one or more output devices **160**, and one or more communication connections **170**. An interconnection mechanism (not shown) such as a bus, controller, or network interconnects the components of the computing environment **100**. Typically, operating system software (not shown) provides an operating environment for other software executing in the computing environment **100**, and coordinates activities of the components of the computing environment **100**.

[0041] The storage **140** can be removable or non-removable, and includes magnetic disks, magnetic tapes or cassettes, CD-ROMs, DVDs, or any other tangible storage medium which can be used to store information in a non-transitory manner and which can be accessed within the computing environment **100**. The storage **140** can also store instructions for the software **180** implementing any of the described techniques, systems, or environments.

[0042] The input device(s) **150** can be a touch input device such as a keyboard, mouse, touch screen, pen, or trackball, a voice input device, a scanning device, or another device that provides input to the computing environment **100**. The output device(s) **160** can be a display, touch screen, printer, speaker, or another device that provides output from the computing environment **100**.

[0043] The communication connection(s) **170** enable communication over a communication medium to another computing entity. The communication medium conveys information such as computer-executable instructions, an agent transport payload, or other data in a modulated data signal. A modulated data signal is a signal that has one or more of its characteristics set or changed in such a manner as to encode information in the signal. By way of example,

and not limitation, communication media include wired or wireless techniques implemented with an electrical, optical, RF, infrared, acoustic, or other carrier.

[0044] The various methods, systems, and interfaces disclosed herein can be described in the general context of computer-executable instructions stored on one or more computer-readable media. Computer-readable media are any available media that can be accessed within or by a computing environment and do not encompass transitory carrier waves. By way of example, and not limitation, with the computing environment **100**, computer-readable media include tangible non-transitory computer-readable media, such as memory **120** and storage **140**.

[0045] The various methods, systems, and interfaces disclosed herein can also be described in the general context of computer-executable instructions, such as those included in program modules, being executed in a computing environment on a target processor. Generally, program modules include routines, programs, libraries, objects, classes, components, data structures, and the like that perform particular tasks or implement particular abstract data types. The functionality of the program modules may be combined or split between program modules as desired in various embodiments. Computer-executable instructions for program modules may be executed within a local or distributed computing environment.

[0046] As noted, the disclosed technology is implemented at least part using a network of computing devices (e.g., any of the computing device examples described above). The network can be implemented at least in part as a Local Area Network (“LAN”) using wired networking (e.g., the Ethernet IEEE standard 802.3 or other appropriate standard) or wireless networking (e.g. one of the IEEE standards 802.11a, 802.11b, 802.11g, or 802.11n or other appropriate standard). Furthermore, at least part of the network can be the Internet or a similar public network.

II. Introduction to the Disclosed Technology

[0047] Distributed generation, demand response, distributed storage, smart appliances, electric vehicles, and other emerging distributed smart grid assets are expected to play a key part in the transformation of the American power system. The nation has made a major investment in deploying smart grid technology via the smart grid investment grants (SGIG) and other utility projects, resulting in rapidly increasing penetration of distributed smart assets. With large-scale plans to integrate renewable generation driven mainly by state-level renewable portfolio requirements, more resources will be needed to compensate for the uncertainty and variability associated with intermittent generation resources. Distributed assets can be used to mitigate the concerns associated with renewable energy resources and to keep costs down. To date, however, the focus has been on obtaining these assets and demonstrating use in small pilot studies, rather than how they will be controlled, coordinated, and compensated when they are present in large numbers.

[0048] For example, some solutions for dispatching demand response are not integrated with those that will be used for distributed storage and generation, nor are they coordinated with the use of more traditional utility assets. The primary use of these assets has been for limiting peak demand, or responding to wholesale price spikes. Using these assets for this purpose is akin to buying a peaking power plant that is utilized only a few hours per year.

Distributed assets have not been widely used to provide additional value streams such as managing capacity at the distribution level, responding to wholesale market fluctuations on a continuous basis, or fast-action regulation services. Not tapping in to their ability to provide these other valuable services limits their cost effectiveness, and therefore slows their penetration and long term viability.

[0049] Centralized command and control approaches that purport to “optimize” the use of these assets are sometimes proposed as the solution to this problem. However, such approaches are impractical because: (1) continuous optimization of the system with large numbers of distributed assets is computationally intractable; (2) they require large transfers of possibly sensitive grid and personal information; and (3) they require the optimizer to make tradeoffs between customer service and value of response that are best made locally at the device level. As a consequence of this desire to support centralized optimization, the path to develop interoperability standards for distributed assets is greatly complicated by the necessity to support the seamless transfer of such a large variety of information from many different types of systems and devices.

[0050] Many distributed assets can provide the same benefits to the grid, and therefore they will compete with each other for the opportunity to provide them. However, retail markets have not gone through a restructuring process similar to what has occurred at the wholesale level and, hence, a level playing field for distributed assets does not currently exist, presenting a need going forward. Currently, there is no unified theory or architecture for how to approach the distributed control problem for distributed assets, nor the requisite market mechanisms to coordinate and incentivize them. New technologies provide the potential for secure, lower-cost power, but lack the control architecture to make them effective.

[0051] Additionally, there is not an integrated simulation environment to develop, test, and validate the distributed assets and the control architecture, with the necessary fidelity to accurately model the physics of the grid in combination with models of wholesale markets.

[0052] Embodiments of the disclosed technology comprise distributed hierarchical control architectures that provides an equitable means for coordinating distributed assets to serve both economic and operational/balancing (reliability) objectives. In certain embodiments, the architecture uses a distributed control and decision making approach, called transactive control, to coordinate distributed assets, as opposed to centralized command and control or a price-responsive only approach. The premise of transactive-based control is that interactions between various components in a complex energy system can be controlled by negotiating immediate and contingent contracts (e.g., using transactive controllers) on a regular basis in lieu of, or in addition to, the conventional command and control. For example, an electrical device (e.g., a distributed smart grid asset) can be given the ability to negotiate deals with its peers, suppliers, and customers to maximize (or otherwise increase) revenues while minimizing (or otherwise reducing) costs. The transactive control paradigm can be implemented through design and integration of retail and wholesale energy markets, creating a feedback loop between the two. The framework can enable the flow of information between different (hierarchical) levels of the electrical system, and can be similar at every time scale (e.g., day-ahead or real-time markets). In

addition to vertical information flow between the layers, the framework can allow information to flow across markets at different time scales—for instance, between day-ahead and real-time markets. Hence, bulk energy will be contracted in day-ahead markets, requiring transacting of residual quantities and remedial actions to be taken into account in real-time markets and balancing operations, respectively. Finally, embodiments of the disclosed technology can allow closer integration of energy markets and balancing operations at every time scale and across different levels of the electrical system. In general, embodiments of the disclosed technology provide a feedback mechanism between the distributed assets via the retail markets up to ISO/RTO-operated wholesale market, and across energy and balancing markets. Hence, embodiments of the DHCA enable a predictable, reliable, smooth, and/or stable response from the entire collection of assets.

[0053] Embodiments of the disclosed technology allow the electrical infrastructure to tap all the values that can be provided by distributed assets for reducing wholesale costs, to manage congestion and capacity at various levels, and/or to manage the intermittency due to renewables. This will make distributed assets more valuable to the grid and more valuable to the customers who provide it. Embodiments of the disclosed technology can be extended to include the provision of ancillary and balancing services with these assets by dynamically selecting them and arming their autonomous response. This enables them to offer significant reliability benefits under disrupted operations.

III. Detailed Description of Example Embodiments of the Disclosed Hierarchical Control Architecture

A. Overview of Distributed Hierarchical Control Architectures

[0054] Embodiments of the disclosed technology employ a distributed control and decision making approach based on a transactive control paradigm as the central tenet. For example, certain embodiments use contract networks and agent-based systems to coordinate agent interactions within and across different hierarchical layers. In particular implementations, the following two elements of the infrastructure help facilitate smooth and reliable operations. The first is the inter-level infrastructure that allows agents at various levels to cooperate to determine the efficient allocation of the available resources. The second is the inter-temporal infrastructure that allows agents to shape the allocation they have received within the time horizon in which it is allocated. The agents in these embodiments can be implemented by transactive controllers comprising computing hardware as described above. The transactive controllers can be used to coordinate the distribution of electricity (e.g., the supply and/or use of electricity) at a variety of hierarchical levels in the power system.

[0055] The inter-level infrastructure addresses resource allocation and is used to reconcile supply resource limits with demand requirements (e.g., feeder constraints versus consumer comfort settings at the retail level). This can be accomplished by using real-time prices. In particular embodiments, for example, this infrastructure establishes a retail market that discovers the price at which supply equals demand at each feeder in the system given the current day-ahead price and prevailing conditions on the feeder and in the homes equipped with price-responsive devices.

[0056] FIG. 2 is a block diagram 200 of an exemplary distributed hierarchical control architecture (DHCA) having a multi-level infrastructure. In FIG. 2, for example, the inter-level infrastructure comprises a device level 210, feeder level 212, area level 214, and system (ISO/RTO) level 216. The interaction of device level agents (collectively shown in FIG. 2 as distributed smart grid assets 230) with the feeder level controllers is facilitated through the newly formulated retail markets. In FIG. 2, the feeder level controllers are illustrated as comprising day ahead feeder management system 220 and real-time feeder management system 222. The area level controllers and the ISO/RTO interact in the wholesale market. In FIG. 2, the area level controllers are illustrated as comprising day ahead area management system 250 and real time area management system 252.

[0057] Block diagram 200 of FIG. 2 also shows that the exemplary DHCA comprises multiple temporal markets. In certain embodiments, for example, the inter-temporal infrastructure comprises a day-ahead market 220 (e.g., 24-hr), a real-time market 222 (e.g., ~5-15 min), and near instantaneous (e.g., ~sec) balancing operations 224.

[0058] In embodiments of the illustrated architecture, the real-time and day-ahead market-based feeder and area management systems 240, 242 aim to maximize (or otherwise increase) distributed asset participation in the real-time and day-ahead markets by incorporating the smart grid assets 230 into standard ISO/RTO market structures. This can be done by solving an optimization problem subject to feeder and area level operational constraints and uncertainties of intermittent renewables and distributed smart grid assets. At the wholesale system level, the ISOs and BAs (which operate respective system level controllers 260, 262 for the day-ahead and real-time markets) receive aggregated net load demand, supply bids for distributed assets from the area management systems 250, 252, and generator power supply offers from generator aggregators 270, 272.

[0059] In the illustrated embodiment, the ISO/RTO also runs balancing reserve markets (shown as balancing operations 224) in parallel with the energy markets to procure reserve capacity to maintain system stability. The cleared, or scheduled, power setpoints and reserve capacity requirements can then be dispatched to area controllers, which in turn dispatch requirements or price signals to the feeder controllers and distributed assets in their respective retail markets.

[0060] More specifically, at the area level, to maintain operational balance at any given instant, an area-level controller (illustrated in FIG. 2 as comprising balancing supervisory area controller 254) solves an area wide optimization problem and then allocates a portion of its reserve capacity requirements to the feeder controller (illustrated in FIG. 2 as comprising balancing supervisory feeder controller 244). It also works to maintain adequate area-wide support in terms of frequency and tie-line flows.

[0061] At the feeder level, the feeder-level controller (e.g., the balancing supervisory feeder controller 244 in FIG. 2) then dispatches setpoints to the distributed assets and primary devices to meet the requirements of the area and feeder controllers. The setpoints can be dispatched at a relatively high frequency on the order of, for example, seconds to 1 minute.

[0062] At the device level, decentralized control schemes are used for the distributed assets 230 to provide both economic and reliability responses. These can be based on self-sensing of frequency, voltage, broadcasts of the imbal-

ance signal, current and future prices, and/or device conditions. The distributed assets can be dynamically influenced via centralized control signals (such as setpoints dispatched from a feeder controller or an RTO/ISO controller) while responding autonomously and instantaneously (e.g., on the order of seconds, such as less than 60 seconds) during disrupted operations.

[0063] The following section introduced exemplary control strategies and market mechanism designs for operating components of the DHCA illustrated in FIG. 2.

B. Control Strategies and Market Mechanism Designs for DHCA

[0064] Control strategies, market mechanisms, and incentive structures can be implemented for a full range of distributed assets in which they express their ability, willingness, and/or intent to modify their consumption or generation. These strategies, mechanisms, and structures set prices or incentives that coordinate the device's response to meet the needs of the grid, as a function of time and location, from the lowest-cost resources available. More specifically, a range of incentives and price schemes can be used to align with operational and capital costs, applicable in both vertically-integrated and restructured market environments, to ensure appropriate levels of customer incentives and the utility revenues needed to invest in and operate the network. The responsiveness of the smart grid assets to both balancing and market-based services is described next.

[0065] 1. Market-Based Services

[0066] Current market structures do not support a level playing field for distributed assets when compared to conventional generation. Embodiments of the disclosed technology provide an equitable market mechanism for coordinating and controlling system assets through a distributed, self-organizing control paradigm which maintains customer free will, but incentivizes and coordinates participation. This is performed in some embodiments using a transactive control paradigm. Distributed smart grid asset participation in the wholesale market can be coordinated through a hierarchical architecture of nested market mechanisms, which involves designing retail markets, while leaving actual functional control at the device level. This allows load serving entities (LSEs) to play their natural role as a resource aggregator in the retail markets and allows for connections to independent third party aggregators for development of optimal portfolios.

[0067] The implementation of embodiments of the disclosed technology typically does not necessitate complete structural changes to current ISO/RTO day-ahead and real-time structures (system level). Instead, embodiments of the disclosed technology complement such structures by providing a mechanism to integrate retail and wholesale markets using a continuous feedback mechanism. In some embodiments, for example, at each of the hierarchical levels (e.g., feeder level 214 and area level 212), available resources—whether demand, supply, or ancillary services—are aggregated from the level below while considering local constraints, such as capacity limits. In particular, device level bids are aggregated by a feeder level controller (e.g., a feeder level management system) while applying local constraints (which enables clearing of a retail market). Further, the feeder level bids are aggregated by an area controller (e.g., an area level management system), which submits an aggregated bid curve into the ISO/RTO whole-

sale market. Conversely, the area and feeder controllers receive the cleared price and dispatch quantities from the ISO/RTO, which are eventually passed down to the end-use customers. This structure forms a feedback mechanism for a closed-loop, multi-level optimization problem, which allows distributed assets to be engaged in the wholesale market.

[0068] The same structural formulation can be applied in both day-ahead and real-time markets. In certain implementations, the only change between day-ahead and real-time markets is the formulation of agents' optimization problems. Exemplary methods for coupling control strategies with economic objectives to derive device level demand/supply bids will be described in the next section.

[0069] 2. Device Level Controls and Demand/Supply Bids

[0070] In certain embodiments of the disclosed technology, at the primary level, the electrical devices 230 use price (and other information) to autonomously determine appropriate actions and apply their own constraints in a transactive manner. FIG. 3, for example, is a block diagram 300 of a graph that represents an HVAC bidding/response strategy that translates real-time price into a thermostat setpoint, and current air temperature into a demand bid. In certain implementations, customers are actively engaged with a simple user interface that allows users to choose along a range between "more comfort" and "more economic" with a simple slider bar. The consumer can also be able to choose the level of market interaction and can override the response, as long as they are willing to pay higher prices. This helps to ensure customer free-will, while incentivizing participation. Exemplary implementations of such strategies are described in more detail in U.S. Nonprovisional application Ser. No. 12/587,008 filed on Sep. 29, 2009, and entitled "ELECTRIC POWER GRID CONTROL USING A MARKET-BASED RESOURCE ALLOCATION SYSTEM," (published as U.S. Patent Application Publication No. 2010/0114387), which is hereby incorporated herein by reference.

[0071] Similar device bid and response mechanisms can be created for other distributed assets, including distributed storage, distributed generation, and smart appliances. Transactive strategies can be generated for a variety of devices. In certain implementations, for example, an environment and a set of rules for participation is created and used where vendors can create additional bidding and control strategies, depending upon the goals of the customer, ranging from relatively simple to highly complex optimization routines or predictive algorithms. Design of device level controls and bidding strategies forms the basis for their participation in retail markets. Equitable treatment of distributed assets in the wholesale markets can be accomplished through retail-wholesale integration as described in the following section.

[0072] 3. Retail-Wholesale Integration

[0073] Embodiments of the disclosed technology provide a framework for fully integrated retail and wholesale power markets. The framework can provide a way for end-users (comprising distributed electrical devices, such as distributed smart grid assets) to participate in the wholesale markets. Embodiments of the retail market design are described first.

[0074] The retail market design facilitates interactions between end-users (distributed electrical devices or assets 230) and a feeder level controller (e.g., comprising the illustrated feeder level management systems 240, 242). The feeder level controller is designed to coordinate the behaviors of the distributed assets, within their respective retail

markets. This provides an avenue to inject local constraints, which are often overlooked when solving system-wide problems. Feeder level optimization and control for a real-time retail market have been used in certain demonstration systems, whose technical aspects can be used to implement embodiments of the disclosed technology. For example, any of the control schemes described in U.S. Nonprovisional application Ser. No. 12/587,008 filed on Sep. 29, 2009, and entitled "ELECTRIC POWER GRID CONTROL USING A MARKET-BASED RESOURCE ALLOCATION SYSTEM," (published as U.S. Patent Application Publication No. 2010/0114387); U.S. Nonprovisional application Ser. No. 12/686,243 filed on Jan. 12, 2010, and entitled "NESTED, HIERARCHICAL RESOURCE ALLOCATION SCHEMA FOR MANAGEMENT AND CONTROL OF AN ELECTRIC POWER GRID" (published as U.S. Patent Application Publication No. 2010/0179862); or U.S. Nonprovisional application Ser. No. 14/108,078 filed on Dec. 16, 2013, and entitled "TRANSACTIVE CONTROL AND COORDINATION FRAMEWORK AND ASSOCIATED TOOLKIT FUNCTIONS", all of which are hereby incorporated herein by reference in their entirety, can be used in embodiments of the disclosed DHCA. Sometimes, in those systems, system-wide constraints (in the form of, for example, wholesale market prices and/or locational marginal prices (LMPs)) are coupled with local constraints (local feeder capacity) to clear retail markets and provide both local and system-wide benefits. Effectively, the systems enable customers to reduce their energy consumption during high price events to reduce energy costs, while coordinating HVAC response during localized congestion events to decrease demand, providing a system to equitably reward customers for participation.

[0075] Distributed generation and storage may similarly bid into the retail market, subject to runtime constraints (e.g., the number of allowable run hours). While successfully showing that distributed assets can participate in retail level markets, the distributed assets in the demonstration projects did not affect the wholesale price; they only reacted to wholesale prices and local constraints.

[0076] In embodiments of the disclosed technology, the feedback loop is closed by integration of retail and wholesale energy markets. This integration allows distributed assets to interact with and affect the wholesale market through the feeder (retail market) and area level controllers (e.g., the feeder and area level management systems 240, 242, 250, 252 illustrated in FIG. 2). In particular implementations, price and availability information flow from the device level to the feeder level (e.g., to feeder level controllers, such as systems 204, 242, which act as retail aggregators). For example, in certain embodiments, the feeder level controller combines individual demand bids, including battery demand bids for charging. Similarly, the feeder level controller combines the supply bids from distributed generators and battery discharging to form feeder level supply curves. The aggregate demand and supply bids are reported to an area level controller (e.g., area level management systems 250, 252 as in FIG. 2), which combines various feeder level bids to report to the ISO/RTO (wholesale market). Once the wholesale market clears, the cleared prices and quantities are reported back to the area and feeder level controllers, which apply their local constraints to clear the respective "markets."

[0077] One such retail market clearing process is demonstrated in block diagram 400 of FIG. 4. In the example illustrated in FIG. 4, based on wholesale market clearing (LMP determination), only battery block SD1 and DG are cleared to supply energy while both supply block SD2 and demand block SC2 are not cleared. The bids submitted by distributed assets are used in the clearing of the wholesale market, and by extension, the retail markets as well. Hence, the optimal bids of distributed assets, as a function of their desired control strategies, can be determined using the feedback mechanism created through integration of retail and wholesale markets.

[0078] In certain implementations, the end-users' bidding process are modeled using learning capabilities. The end users can simultaneously (or substantially simultaneously) report supply bids to participate in ancillary services markets. The demand curves for ancillary services can be derived by respective entities at the different levels and markets can be cleared similar to the energy market.

[0079] In certain embodiments, the wholesale markets are operated by an ISO/RTO and facilitate interactions between the ISO/RTO and the area level controllers (e.g., the area level management systems 250, 252). Conventional and grid-level renewable generation can also operate directly in the wholesale markets. The ISO/RTO can clear the wholesale market using the processes of security constrained unit-commitment and economic dispatch. Given the high percentage of renewables and distributed assets to be modeled in certain environments, it may be desirable to use stochastic versions of market clearing processes.

[0080] In some embodiments, at the area level controllers (which act as wholesale aggregators), feeder level bids are aggregated and area constraints are applied to derive wholesale market bids. This has the advantage that, at the area level, distributed asset constraints are no longer considered but rather handled at the lower level of control. In return, the ISO can establish wholesale market prices, which are delivered to the area level management system. The optimal bidding strategies problem of the area level management system can be modeled as a mathematical program with equilibrium constraints. The outer problem of the bi-level problem is the area level management system's optimization problem while the inner problem is the ISO's optimization problem. Hence, the wholesale market clearing process will depend on bids provided by agents and entities at different levels. The bids, in turn, can be formulated based on market clearing processes at both wholesale and retail layers. The integration of retail and wholesale markets in this manner facilitates greater participation by distributed assets. The design of the retail and wholesale markets can be kept the same in both day-ahead and real-time markets. The interaction of agents in day-ahead and real-time markets is described next.

[0081] 4. Interaction of Day-Ahead and Real-Time Markets

[0082] Day-ahead markets are operated as pure financial markets, allowing participants to enter financially binding contracts that hedge against price volatility in real-time markets. The real-time markets serve as imbalance markets.

[0083] The principle of the two-settlement system can be preserved in embodiments of the disclosed technology. For example, in certain embodiments, distributed assets are able to enter into contracts to procure or sell most of their power "needs" in the day-ahead markets. The residual amount of

power is transacted in the real-time markets, when more information of the prevailing weather and system conditions become available. The real-time markets can also serve to correct the imbalance between contracted day-ahead positions and actual real-time positions. The horizontal information flow between retail and wholesale market, and their respective entities (feeder-level and area-level management systems) can be modeled similar to distributed assets.

[0084] In the presence of high renewable and distributed assets penetration, economics cannot be the only objective for using distributed resources in an effective manner (unless reliability can be translated into costs). As illustrated in FIG. 2, markets and control strategies for balancing services can be run in parallel with retail and wholesale markets in both day-ahead and real-time markets. The market clearing mechanism used in the wholesale markets can clear balancing reserves by co-optimizing energy and balancing needs. For example, embodiments of the framework illustrated in FIG. 2 can allow clearing balancing markets at a nodal level, rather than system or zonal levels as is the norm today. Hence, the prices for balancing reserves more realistically reflect the demand for reserve capacity; for instance, areas with greater wind penetration would require more reserve capacity, reflected in the nodal prices. The framework also allows for balancing markets to be cleared at the retail level with greater penetration of distributed assets.

[0085] Further, the balancing reserves can be determined endogenously, based on energy demand and supply, rather than set as hard limits as is the norm in most markets today. The balancing service needs will be described in more detail in the next section.

[0086] 5. Balancing Services

[0087] In certain embodiments of the disclosed technology, a reliability safety net is created, comprising fast-acting smart grid assets (e.g., distributed smart grid assets 230). These distributed assets can be aggregated into a grid-friendly network of actively-configured, autonomous responses to self-sense frequency and voltage fluctuations, and broadcast signals for control area imbalance. This enables the assets to provide the full range of today's ancillary services and more: virtual inertia, regulation, ramping, spinning reserve, and/or emergency curtailment capabilities. Example control strategies at each level are discussed in more detail in the subsections below.

[0088] a. Balancing Supervisory Area Controller

[0089] The balancing supervisor (e.g., balancing supervisory area controller 254) enables DSGAs to provide balancing services, reducing the burden on conventional generation, particularly when increasing balancing requirements are expected due to the high integration of renewable variable generation. The balancing supervisory area controller can be implemented as part of an area controller that also computes the transactive signals for participating in the real-time or day-ahead market or can be implemented separately. Thus, the area controller may comprise a system of multiple computing devices (each implementing one or more of the day-ahead area management system 250, real time area management system 252, or balancing supervisory area controller 254) or comprise a single computing device that performs all of those functions. For ease of discussion, however, all such implementations are encompassed by the term "area controller".

[0090] At the area level, and according to one exemplary embodiment of the disclosed technology, the goal of the

balancing supervisor is to minimize (or otherwise reduce) a DSGA Area Control Error (DSGA-ACE) signal, which is a weighted sum of the deviations of the system frequency and the inter-area power flow. The balancing supervisor can be coordinated with Automatic Generator Control (AGC) at the transmission level to provide frequency and tie line interchange support (minimizing or otherwise reducing balancing authority ACE). In particular implementations, the balancing supervisor, in coordination with AGC, aims to maintain the system frequency at 60 Hz during the normal load demand fluctuation and to restore the system frequency gradually when a contingency occurs in the system. In particular, the balancing supervisor desirably maintains the inter-area power flow at the desired level. The inter-area power flow reference value is calculated based on differences of measured total area real and reactive power and reference values the power reference set by a Balancing Authority (BA) and the cleared power reference from a Real Time Area Management System (e.g., a real-time controller configured to participate in the real-time market). The balancing supervisor can weight these two power references according to the current system operating conditions.

[0091] FIG. 6 is a flowchart 600 illustrating one exemplary method for performing balancing operations at an area controller (e.g., using the balancing supervisory area controller 254). The method of FIG. 6 can be performed using computing hardware (e.g., a computer processor or an integrated circuit). For instance, the method can be performed by a computer or specialized hardware of a utility or power system operator and be configured to exchange signals with feeders in the area and with one or more ISO/RTO computing entities. Furthermore, it should be understood that the method acts in FIG. 6 do not necessarily occur in the illustrated sequence.

[0092] At 610, a system frequency is monitored. For example, the system frequency can be monitored at one or more points in the area (e.g., across multiple feeders in the area).

[0093] At 612, a power flow reference value indicative of an inter-area power flow is computed. In certain embodiments, and as discussed above, the power flow reference value is based on one or more of: (a) differences between measured area real and reactive power and power reference values set by a system-wide Balancing Authority (BA); or (b) differences between measured area real and reactive power and power reference values for a current cleared time interval of a transactive control energy market in which the area controller operates. For example, the power flow reference value can be based on a weighted sum of (a) and (b), and (a) can be weighted more heavily during contingencies and (b) can be weighted more heavily during normal operation.

[0094] At 614, an area control error is computed based at least in part on the system frequency and the power flow reference value. For example, the area control error can be computed as described above.

[0095] At 616, one or more electrical generators in the area are controlled to adjust their supply and to thereby reduce the area control error.

[0096] In certain embodiments, the power flow reference value is transmitted to controllers for the multiple feeder networks (e.g., to a feeder level controller implementing a balancing supervisory feeder controller 244). As explained above, the power flow reference value can then be used by

the feeder networks to compute feeder-level reference values that are used autonomously by the feeders to control electrical devices (e.g., distributed smart grid assets) to help address any imbalance.

[0097] In some embodiments, the method is performed in cooperation with a transactive control energy market in which the area controller also participates. Thus, as the method of FIG. 6 is being performed, the area controller can additionally operate as a transactive node in the energy market (e.g., the area controller can implement one or more of the day-ahead area management system 250 or real-time area management system 252). This participation can involve computing one or more transactive signals representing a bid for using or supplying electricity in an upcoming time interval of a transactive control energy market, the time interval being a periodic time interval (e.g., a time interval associated with a real-time market, a day-ahead market, or both).

[0098] b. Balancing Supervisory Feeder Controller

[0099] The feeder controller (e.g., comprising balancing supervisory feeder controller 244) connects the area controller with individual devices. The balancing supervisory feeder controller can be implemented as part of a feeder controller that also computes the transactive signals for participating in the real-time or day-ahead market or can be implemented separately. Thus, a feeder-level controller may comprise a system of multiple computing devices (each implementing one or more of the day-ahead feeder management system 240, real time feeder management system 242, or balancing supervisory feeder controller 244) or comprise a single computing device that performs all of those functions. For ease of discussion, however, all such implementations are encompassed by the term “feeder controller”.

[0100] In particular embodiments, the feeder controller’s role is bidirectional in the sense that: (i) it “translates” the balancing need specified by the area control to individual control actions for the devices; and (ii) it estimates elasticity of the assets and sends this information to the area controller.

[0101] In certain embodiments, the feeder controller has two objectives. The first one is to minimize (or otherwise reduce) a Feeder Control Error (FCE) signal which represents the weighted sum of the differences between the measured total feeder real and reactive power and the reference values. The second objective is to coordinate various devices to provide VAR support for voltage regulation.

[0102] In certain embodiments, the feeder real power reference is calculated based on the power reference received from an area controller (e.g., from a balancing supervisory area controller 254) and the power reference for an imminent time interval for the real-time market (e.g., from a market-based real time feeder management system 242). As noted, the balancing supervisory feeder controller 244 may be implemented as part of the feeder controller that also implements the real time feeder management system 242, either through separate computing hardware or as integrated hardware. The feeder controller can weight these two power references according to the current system operating conditions. During normal conditions, the cleared power reference from the market-based entity will be given more weight, while during contingencies, the signal received from the balancing entity will be weighted more. Additionally, the feeder controller can account for variability and uncertainty of local distributed renewable generation (such

as rooftop photovoltaic panels), and local constraints such as feeder congestion. After the feeder real power reference is determined, the feeder controller dispatches setpoints (e.g., optimal setpoints) to the devices to maintain adequate power support.

[0103] Additionally, in some embodiments, the feeder controller collects real time feeder voltage information from devices involved in VAR support. Then the feeder controller can coordinate with local device-level controllers by dispatching voltage setpoints and, if necessary, can also override the local decentralized control signal to avoid excessive voltage regulation.

[0104] FIG. 7 is a flowchart 700 showing one exemplary method for performing balancing operations at a feeder controller (e.g., using the balancing supervisory feeder controller 244). The method of FIG. 7 can be performed using computing hardware (e.g., a computer processor or an integrated circuit). For instance, the method can be performed by a computer or specialized hardware of a utility or power system operator and be configured to exchange signals with electrical devices (e.g., assets 230) served by the feeder and with one or more area controllers (e.g., a balancing supervisory area controller 254). Furthermore, it should be understood that the method acts in FIG. 7 do not necessarily occur in the illustrated sequence.

[0105] At 710, a system frequency at the feeder is monitored.

[0106] At 712, a feeder power reference value is computed. In some embodiments, the feeder power reference value is based on one or more of: (a) differences between measured feeder real and reactive power and power reference values set by an area controller; or (b) differences between measured feeder real and reactive power and power reference values for the feeder for a current cleared time interval of a transactive control energy market in which the feeder operates. For example, the feeder power reference value can be based on a weighted sum of (a) and (b), and (a) can be weighted more heavily during contingencies and (b) can be weighted more heavily during normal operation.

[0107] At 714, a feeder control error is computed based at least in part on the system frequency and the feeder power reference value.

[0108] At 716, one or more setpoints are computed for one or more electrical devices coupled to and served by the feeder. In certain embodiments, the setpoints are configured to adjust performance of the one or more electrical devices so that the feeder control error is reduced.

[0109] At 718, the one or more setpoints are transmitted to controllers for the one or more electrical devices (e.g., to transactive controllers for the devices).

[0110] In some embodiments, the method is performed in cooperation with a transactive control energy market in which the feeder controller also participates. Thus, as the method of FIG. 7 is being performed, the feeder controller can additionally operate as a transactive node in the energy market. This participation can involve computing one or more transactive signals representing a bid for using or supplying electricity in an upcoming time interval of a transactive control energy market, the time interval being a periodic time interval (e.g., a time interval associated with a real-time market, a day-ahead market, or both).

[0111] c. Device Level Control

[0112] At the device level, distributed assets (e.g., distributed smart grid assets 230) desirably provide multiple ser-

vices at different time scales, including, for instance, one or more of the following services: (1) responding to market prices, (2) responding to an imbalance signal (e.g., from the feeder controller), and/or (3) responding autonomously to reliability needs.

[0113] Autonomous responses are desirable for many reliability purposes where there may not be time to communicate needed actions through a wide-area network. Appliance and equipment manufacturers are rapidly moving toward mass production of devices with smart grid capabilities that can be leveraged for this purpose. However, utilities and balancing authorities have been hesitant to support such deployments because the response of fleets of such devices has not been fully integrated with their control schemes for grid stability.

[0114] Distributed assets in embodiments of the disclosed technology are equipped with autonomous controllers with settings armed according to instructions from feeder, area or system levels. In this way, the autonomous immediate response of devices can be tailored to system needs, such as low system inertia due to high on-line renewable generation.

[0115] To provide the multiple services, at the device level, distributed assets can be equipped with multi-objective control strategies developed to enable single resources to provide multiple benefits to the system. These control strategies can be accounted for in the coordination problem, for example, by using a receding horizon optimization technique (e.g., model predictive control).

[0116] d. Integrating Balancing and Market-Based Services

[0117] In certain embodiments, the real-time market management systems (both area and feeder level) play an additional role. As part of the co-optimization problem developed at each level, ancillary service contracts can also be formed on the same time scale as the real-time energy market, and weighted against real-time energy market requirements. During normal operations, the management systems typically take a purely economic perspective to maximize (or otherwise increase) reward (profit) by dispatching smart grid assets either towards real-time energy needs, ancillary service needs (such as frequency or voltage regulation, spinning reserve, etc.), or a combination of both.

[0118] In certain embodiments, during disrupted or stressed system conditions, weighting functions are adjusted to focus on system stability requirements rather than economic concerns. Effectively, during each real-time market cycle (e.g., at a ~5 min, or other real-time market cycle), the management systems at each level (implemented by transactive controllers at each level) form a “contract” for real-time energy and balancing/regulation services, dispatching resources subject to local constraints and availability provided by device bidding. This allows the smart grid assets to participate in multiple market revenue streams under a multi-objective control problem (e.g., a storage device participating in both energy markets and frequency regulation), capturing multiple revenue streams to increase profitability and long-term sustainability. Also of note, is that at this time-scale, “contracts” are formed for ancillary services, while the control is performed on a much faster scale (as described in the section on operational and balancing services).

[0119] e. Control Strategies During Disrupted Operations

[0120] In embodiments of the disclosed technology, the distributed smart grid assets 230 (e.g., the end-use electrical

devices) are dynamically armed for instantaneous autonomous response during normal and disrupted operations. Disrupted operations may occur, for example, due to a loss of communications or an emergency event. This allows each asset (electrical device) to respond to the correct extent to avoid amplifying frequency oscillation. For non-controllable loads that cannot continuously adjust their power use (such as water heaters, HVACs, and the like), the devices can be switched on/off probabilistically so that the aggregation of a large number of these loads will provide the desired amount of power droop. Distributed control strategies can be designed to coordinate the different devices to respond autonomously while maintaining the overall stability of the system. One exemplary approach is to use the mean-field game (MFG) theoretic framework. The basic premise behind the MFG framework is that it provides insights on the behavior of large population of “agents”, where each agent optimizes his action taking into account his own objectives as well as global macroscopic information (e.g., aggregate impact of the actions of all other players). The resulting problem has a dynamic game structure, since each player optimizes independently based on his available information. Under limiting conditions, such as population size approaching infinity, the common information available to the players reduces to a deterministic quantity and the game can be analyzed using dynamic programming techniques. In the context of distributed assets, each resource—be it a distributed generator or load (demand response) or storage device—observes the frequency signal that is common to the entire system and which depends on the power injections/withdrawals of all the resources in the system. The resource then adjusts its injection/withdrawal to correct the deviation of the frequency from its nominal value.

[0121] f. Other Cases: Feeder Level System Owns and Manages Distributed Assets

[0122] The previous discussion has assumed that the distributed assets are owned by the end-users. In the case where the feeder level entity owns and operates the distributed assets, optimization goals may be re-ordered. Embodiments of the disclosed technology can also account for cases where the feeder level entity owns the distributed assets. To be consistent with the completely decentralized (transactive) control and decision making paradigm, the total capacity of distributed assets in the system can be kept the same between the two cases. The feeder level management system can use local constraints and distributed asset control strategies to form the feeder level demand and supply curves to facilitate retail market clearing. The same formulation can be used in both day-ahead and real-time markets.

[0123] As seen in block diagram 400 of FIG. 4, the market-clearing mechanism determines which assets get cleared to participate in the market and the corresponding reward they earn. The distributed asset control strategy used by a feeder level management system can directly influence retail and wholesale market clearing and, hence, the rewards of distributed assets. It is therefore desirable that the entity appropriately weight risks and rewards in selecting an appropriate (e.g., optimal) control strategy. A control strategy that weighs risk more than reward would adversely impact revenues by clearing fewer resources, while an objective weighted higher with customer rewards would expose the aggregator to undue risks. The risk for a feeder level management system manifests itself in the form of under-supply of energy or reserves relative to the contracted

amounts. Additionally, the entity would be burdened with procuring additional capacity at presumably higher costs. The reward, on the other hand, is derived from supply of energy and reserve capacity to the system. This model can also be extended in a manner where the aggregator contracts with the end users to manage their DSGs in return for energy cost savings and/or revenues from the supply of energy and reserves. The aggregator desirably selects a control strategy to properly align end user economic incentives with its own, while minimizing (or otherwise reducing) risk exposure.

C. General Embodiments of the Disclosed Technology

[0124] It should be understood that the architecture shown in block diagram 200 of FIG. 2 is a specific example of a DHCA and that variations exist without departing from the principles underlying the example architecture. This section describes general embodiments of the disclosed technology exemplified by FIG. 2 and the discussions above. For instance, in certain embodiments, the system can be generally described as a system comprising two energy markets—wholesale and retail—that includes a means for integrating the two markets, such as through transactive signals that are exchanged between the two markets.

[0125] More generally, then, embodiments of the disclosed technology include a system comprising one or more transactive controllers configured to operate a wholesale market for coordinating the distribution of electricity between bulk power generation resources and area controllers that distribute electricity to and from multiple feeder networks, and one or more transactive controllers configured to operate a retail market for coordinating the distribution of electricity between one or more of the feeder networks and one or more electrical devices. Further, at least one of the transactive controllers in the wholesale market is configured to exchange transactive control signals with at least one of the transactive controllers in the retail market. The retail market is thereby integrated with the wholesale market. The integration of the retail market with the wholesale market allows one or more electrical devices to participate in the wholesale market via the exchanged transactive control signals between the at least one of the transactive controllers in the wholesale market and the at least one of the transactive controllers in the retail market.

[0126] Further, in certain implementations (and as illustrated in FIG. 2 and discussed above), at least one of the transactive controllers in the retail market is a transactive controller associated with a respective one of the feeder networks and is configured to: (a) exchange transactive control signals with transactive controllers for multiple electrical devices served by the respective one of the feeder networks; (b) aggregate bids for supply, demand, or both supply and demand from the electrical devices; and (c) generate bids for supply, demand, or both supply and demand for the wholesale market based at least in part on the aggregated bids.

[0127] In further implementations (and as illustrated in FIG. 2 and discussed above), at least one of the transactive controllers in the wholesale market is a transactive controller associated with one of the area controllers and is configured to: (a) exchange transactive control signals with transactive controllers for multiple feeder networks in a respective area; (b) aggregate bids for supply, demand, or both supply and demand from the transactive controllers for the multiple

feeder networks; and (c) generate bids for supply, demand, or both supply and demand for the wholesale market based at least in part on the aggregated bids.

[0128] In certain implementations, one or more of the transactive controllers configured to operate in the wholesale market comprises a system-level transactive controller associated with an ISO or RTO. The ISO or RTO can clear the wholesale market using one or more of security constrained unit-commitment or economic dispatch. Further, a system-level transactive controller can be configured to transmit signals indicative of a cleared price and a quantity of electricity to a respective area-level transactive controller (an area controller).

[0129] In some implementations, the system-level transactive controller is further configured to compute and transmit a setpoint for use by one or more electrical devices in the event of a disruption to the electrical system.

[0130] In certain implementations, the one or more transactive controllers in the wholesale market and in the retail market are further configured to operate in multiple temporal markets. For example, the transactive controllers can operate in a real-time market and a day-ahead market. In particular implementations, the real-time market comprises a market for settling a price for electricity in an imminent time interval (e.g., an upcoming (or next-to-occur) time interval, which may be of any suitable duration, such as 5 minutes, 10, minutes, 15 minutes or other such duration). Further, the day-ahead market can comprise a market for settling power supply commitment for the following day.

[0131] In some implementations, at least one of the transactive controllers configured to operate in the retail market is further configured to receive a signal indicative of an imbalance or disruption in the system, and, in response thereto, adjust electrical supply or demand of one or more electrical devices associated with the at least one of the transactive controllers.

[0132] In a more specific embodiment, the example architecture of FIG. 2 illustrates a system for coordinating distribution of electricity according to a distributed hierarchical control architecture (DHCA), comprising one or more device-level transactive controllers at a device level of the DHCA; one or more feeder-level transactive controllers at a feeder level of the DHCA; one or more area-level transactive controllers at an area level of the DHCA; and one or more system-level transactive controllers of the DHCA.

[0133] In the illustrated embodiment, at least one of the device-level transactive controllers and at least one of the feeder-level transactive controllers is configured to communicate with one another, at least one of the feeder-level transactive controllers and at least one of the area-level transactive controllers is configured to communicate with one another, at least one of the area-level transactive controllers and at least one of the system-level transactive controllers is configured to communicate with one another.

[0134] Further, the transactive controllers can be configured to participate in a transactive control energy market as well as to perform balancing operations designed to improve the stability of the power system. For instance, the device-level transactive controllers, feeder-level transactive controllers, area-level transactive controllers, and system-level transactive controllers can be configured to participate in a real-time market for settling a price for electricity in one or more upcoming time intervals for the real-time market. Further, at least some of the device-level transactive con-

trollers, feeder-level transactive controllers, area-level transactive controllers, and system-level transactive controllers can also be configured to participate in a balancing market, the balancing market being configured to maintain system stability during periods of imbalance in the system. It is to be understood that the phrase “transactive controller” as used in this context may comprise a system of computing devices, where one computing device is responsible for participating in the real-time market and another computing device is responsible for participating in balancing operations, or may comprise a single computing device configured to perform all the operations.

[0135] In particular implementations, the device-level transactive controllers operate at a faster frequency in the balancing market than in the real-time market. Further, in certain implementations, one or more of the device-level transactive controllers participating in the balancing market are configured to control a controllable load or a controllable energy supply. For example, in certain implementations, one or more of the device-level transactive controllers participating in the balancing market are configured to control discharging and charging of an electric battery. In some implementations, one or more of the device-level transactive controllers participating in the balancing market are configured to provide balancing services by reacting autonomously to one or more of a frequency of the power system, voltage of the power system, a broadcast imbalance signal, current real-time market price, day-ahead price, or device conditions. Further, in certain implementations, the device-level transactive controllers, feeder-level transactive controllers, area-level transactive controllers, and system-level transactive controllers are further configured to participate in a day-ahead market for settling a price for electricity for a following day.

[0136] FIG. 5 is a flowchart 500 showing another exemplary embodiment for operating a transactive controller in a market-based electrical-energy-allocation system. The method shown in FIG. 5 can be implemented, for example, by a device-level transactive controller used to control one or more electrical devices (e.g., distributed smart grid assets) in a system, but this usage should not be construed as limiting. The electrical device can be a variety of devices or assets, such as an air-conditioning unit; heating unit; heating, ventilation, and air conditioning (HVAC) system; hot water heater; refrigerator; dish washer; washing machine; dryer; oven; microwave oven; pump; home lighting system; electrical charger; electric vehicle charger; home electrical system; or any other electrical system having variable performance states. The method of FIG. 5 can be performed using computing hardware (e.g., a computer processor or an integrated circuit). For instance, the method can be performed by a computer at an end-user’s locale or home, a computer coupled to an electrical device, or by specialized hardware (e.g., an ASIC or programmable logic device) coupled to the electrical device. Furthermore, it should be understood that the method acts in FIG. 5 do not necessarily occur in the illustrated sequence.

[0137] At 510, one or more transactive signals are computed for coordinating the use, supply, or both use and supply of electricity in an upcoming time interval of a transactive control energy market. The time interval can be a periodic time interval (e.g., every 5 minutes, 10 minutes, 15 minutes, or other such interval).

[0138] At **512**, one or more signals are received indicating an imbalance in the electrical-energy-allocation system.

[0139] At **514**, a control strategy for reducing the imbalance is implemented responsive to the one or more signals indicating the imbalance. For example, electrical use or supply of the one or more electrical devices controlled by the transactive controller can be adjusted. The adjusting can comprise a number of difference responses (e.g., depending on whether the electrical device(s) controlled by the transactive controller have widely variable performance states (controllable electrical devices) or simply two operational states, on and off (non-controllable electrical devices). Thus, the adjusting can comprise shutting off the electrical device, turning on an electrical device, reducing an electricity consumption rate of the electrical device, or increasing an electrical supply from the electrical device.

[0140] In certain implementations, the receiving and the implementing occur independent and autonomously of the transactive signals exchanged during the periodic time interval. For example, in certain implementations, the receiving and the implementing occur asynchronously to the fixed interval frequency. Thus, the transactive controller is able to quickly respond to imbalances in the system at a rate that is faster than the frequency at which prices and electricity distribution is settled in the real-time market. In some implementations, the control strategy that is implemented is based at least in part on a setpoint dispatched from a controller associated with a feeder with which the electrical device is connected. For instance, in the event of a disruption, the electrical devices controlled by the transactive controller can be set to respective setpoints dispatched from the feeder controller.

[0141] In some implementations, the one or more signals indicating the imbalance comprise one or more of a frequency of the power system, voltage of the power system, or a broadcast imbalance signal (e.g., broadcast from a feeder controller, an area controller, or an ISO/RTO controller).

[0142] In certain implementations, a disruption to the system is detected (e.g., as a result of a major contingency or communication outage) and a control strategy for responding to the disruption is automatically implemented until the disruption ends. For instance, the control strategy can comprise shutting off the electrical device, reducing an electricity consumption rate of the electrical device, or increasing an electrical supply from the electrical device. This helps create a predictable response among distributed smart grid assets for the ISO/RTO.

D. Simulation Platform for DHCA

[0143] Modeling of smart grid applications is desirably performed using an integrated modeling approach that portrays the performance of the assets. This is in context with the surrounding bulk generation, transmission infrastructure, market systems, reliability coordination, and other aspects of utility planning and operations. Power system modeling and simulation tools are available for a wide range of these applications, but do not effectively capture the rich interactions between the different domains. To address this problem, and to provide a platform for rapid control algorithm prototyping, embodiments of the disclosed technology comprise a simulation platform that simultaneously models the integrated transmission and distribution systems, incorporates wholesale and retail markets (real-time and day-ahead), and/or considers the communication requirements

needed to realize the integrated system. In one particular implementation, the simulation platform comprises a transmission solver (e.g., PowerWorld), a distribution and smart grid asset simulator (e.g., GridLAB-D), a market-based simulator (e.g., SuperOPF) and/or a communications simulator (e.g., ns-3). The platform can also include “hardware-in-the-loop” functionality for validating the final DHCA by tying directly into existing devices where applicable.

[0144] In particular implementations, the four simulation tools provide a testbed for evaluating and designing embodiments of the disclosed DHCA. The tools can be used independently or in an integrated fashion to facilitate the development of the DHCA.

E. Validation and Testing of DHCA Embodiments

[0145] In certain implementations, validation of one or more components and/or control strategies can be part of the process of implementing the disclosed technology for a particular area. Where available, hardware-in-the-loop and existing market and visualization tools can be incorporated into the simulation environment to test the validity of the particular implementation being developed.

F. Technology Commercialization

[0146] In certain implementations, technology commercialization can be part of the process of implementing the disclosed technology for a particular area. For example, a commercialization path can be developed by evaluating the results from validation. The commercialization path can include evaluating the cost effectiveness of system-wide deployment versus retrofit to existing systems, market software applications, device control algorithms, and support infrastructure requirements. Paths for participation in technology deployment by additional entities can be identified (e.g., under the requirement they follow interoperability standards defined for the various components of the system).

IV. Concluding Remarks

[0147] Having illustrated and described the principles of the disclosed technology, it will be apparent to those skilled in the art that the disclosed embodiments can be modified in arrangement and detail without departing from such principles. For example, any one or more aspects of the disclosed technology can be applied in other embodiments.

[0148] In view of the many possible embodiments to which the principles of the disclosed technologies can be applied, it should be recognized that the illustrated embodiments are only preferred examples of the technologies and should not be taken as limiting the scope of the invention. Rather, the scope of the invention is defined by the following claims and their equivalents. We therefore claim all that comes within the scope and spirit of these claims and their equivalents.

1. A system for coordinating distribution of electricity according to a distributed hierarchical control architecture (DHCA), comprising:

one or more transactive controllers configured to operate a wholesale market for coordinating the distribution of electricity between power generation resources and area controllers that distribute electricity to and from multiple feeder networks; and

one or more transactive controllers configured to operate a retail market for coordinating the distribution of

electricity between one or more of the feeder networks and one or more electrical devices, wherein at least one of the transactive controllers in the wholesale market is further configured to exchange transactive control signals with at least one of the transactive controllers in the retail market, and to thereby integrate the retail market with the wholesale market.

2. The system of claim **1**, wherein the integration of the retail market with the wholesale market allows one or more of the electrical devices to participate in the wholesale market via the exchanged transactive control signals between the at least one of the transactive controllers in the wholesale market and the at least one of the transactive controllers in the retail market.

3. The system of claim **1**, wherein the at least one of the transactive controllers in the retail market is a transactive controller associated with a respective one of the feeder networks and is configured to:

- (a) exchange transactive control signals with transactive controllers for multiple electrical devices served by the respective one of the feeder networks;
- (b) aggregate bids for supply, demand, or both supply and demand from the electrical devices; and
- (c) generate bids for supply, demand, or both supply and demand for the wholesale market based at least in part on the aggregated bids.

4. The system of claim **3**, wherein the generated bids are further based at least in part on constraints at the respective one of the feeder networks.

5. The system of claim **1**, wherein the at least one of the transactive controllers in the wholesale market is a transactive controller associated with one of the area controllers and is configured to:

- (a) exchange transactive control signals with transactive controllers for multiple feeder networks in a respective area;
- (b) aggregate bids for supply, demand, or both supply and demand from the transactive controllers for the multiple feeder networks; and
- (c) generate bids for supply, demand, or both supply and demand for the wholesale market based at least in part on the aggregated bids.

6. The system of claim **5**, wherein the generated bids are further based at least in part on constraints in the respective area.

7. The system of claim **1**, wherein the one or more transactive controllers configured to operate in the wholesale market comprises a system-level controller associated with an independent system operator or regional transmission organization.

8. The system of claim **7**, wherein the system-level transactive controller is configured to transmit signals indicative of a cleared price and a quantity of electricity to a respective transactive controller associated with an area controller.

9. The system of claim **8**, wherein the system-level transactive controller is further configured to compute and transmit a setpoint for use by one or more electrical devices in the event of a disruption to the system.

10. The system of claim **1**, wherein the one or more transactive controllers configured to operate in the wholesale market and the one or more transactive controllers config-

ured to operate in the retail market are further configured to operate in both a real-time market and a day-ahead market, wherein the real-time market comprises a market for settling a price for electricity in an imminent time interval, and

wherein the day-ahead market comprises a market for settling power supply commitment for the following day.

11. The system of claim **1**, wherein at least one of the transactive controllers configured to operate in the retail market is further configured to receive a signal indicative of an imbalance or disruption in the system, and, in response thereto, adjust electrical supply or demand of one or more electrical devices associated with the at least one of the transactive controllers.

12.-17. (canceled)

18. A method for operating a controller in a market-based electrical-energy-allocation system, comprising:

by computing hardware:

computing one or more transactive signals for coordinating the use, supply, or both use and supply of electricity in an upcoming time interval of a transactive control energy market, the time interval being a periodic time interval;

receiving one or more signals indicating an imbalance in the electrical-energy-allocation system; and

responsive to the one or more signals indicating the imbalance, implementing a control strategy for reducing the imbalance, wherein the receiving and the implementing occur independent and autonomously of the transactive signals exchanged during the periodic time interval.

19. The method of claim **18**, wherein the implementing the control strategy comprises adjusting electrical use or supply of the one or more electrical devices controlled by the controller.

20. The method of claim **18**, wherein the receiving and the implementing occur independent and autonomously of the real-time market.

21. The method of claim **18**, wherein the controller is a device-level controller for controlling one or more distributed smart grid assets.

22. The method of claim **18**, wherein the control strategy is based at least in part on a setpoint dispatched from a controller associated with a feeder with which the electrical device is connected.

23. The method of claim **18**, wherein the one or more signals indicating the imbalance comprise one or more of a frequency of the power system, voltage of the power system, or a broadcast imbalance signal.

24. The method of claim **18**, further comprising detecting a disruption in the system, and automatically implementing a control strategy for a system disruption until the disruption ends.

25. The method of claim **24**, wherein the control strategy for the system disruption comprises one of shutting off the electrical device, reducing an electricity consumption rate of the electrical device, or increasing an electrical supply from the electrical device.

26. One or more non-transitory computer-readable media storing computer-readable instructions for causing computer to perform the method of claim **18**.

27.-40. (canceled)