



(12) **DEMANDE DE BREVET CANADIEN**
CANADIAN PATENT APPLICATION


(13) **A1**

(86) Date de dépôt PCT/PCT Filing Date: 2020/07/31
(87) Date publication PCT/PCT Publication Date: 2021/02/04
(85) Entrée phase nationale/National Entry: 2022/01/24
(86) N° demande PCT/PCT Application No.: US 2020/044539
(87) N° publication PCT/PCT Publication No.: 2021/022175
(30) Priorité/Priority: 2019/08/01 (US16/529,402)

(51) Cl.Int./Int.Cl. *G06Q 50/06* (2012.01),
G06Q 30/02 (2012.01)
(71) Demandeur/Applicant:
LANCIUM LLC, US
(72) Inventeurs/Inventors:
MCNAMARA, MICHAEL T., US;
CLINE, RAYMOND E. JR., US
(74) Agent: ROBIC

(54) Titre : MODIFICATION D'OPERATIONS DE SYSTEME INFORMATIQUE SUR LA BASE DE CONDITIONS DE COUT
ET DE PUISSANCE

(54) Title: MODIFYING COMPUTING SYSTEM OPERATIONS BASED ON COST AND POWER CONDITIONS



PRIOR ART
FIGURE 1

(57) **Abrégé/Abstract:**

Examples relate to flexible datacenters or other power loads tolerant of intermittent operation and configured to use power received behind-the-meter. A system may include a set of computing systems powered by behind-the-meter power and a datacenter control system. The datacenter control system may be configured to monitor a set of conditions that includes behind-the-meter power availability at the set of computing systems and perform a weighted analysis using the set of monitored conditions. Based on the weighted analysis, the datacenter control system may be configured to modulate operating attributes of one or more computing systems of the set of computing systems. In some examples, the datacenter control system is a remote master control system in communication with a datacenter control system of a flexible datacenter housing the set of computing systems.

Date Submitted: 2022/01/24

CA App. No.: 3145486

Abstract:

Examples relate to flexible datacenters or other power loads tolerant of intermittent operation and configured to use power received behind-the-meter. A system may include a set of computing systems powered by behind-the-meter power and a datacenter control system. The datacenter control system may be configured to monitor a set of conditions that includes behind-the-meter power availability at the set of computing systems and perform a weighted analysis using the set of monitored conditions. Based on the weighted analysis, the datacenter control system may be configured to modulate operating attributes of one or more computing systems of the set of computing systems. In some examples, the datacenter control system is a remote master control system in communication with a datacenter control system of a flexible datacenter housing the set of computing systems.

MODIFYING COMPUTING SYSTEM OPERATIONS BASED ON COST AND POWER CONDITIONS

CROSS REFERENCE TO RELATED APPLICATIONS

[0001] This application claims priority to and is a Continuation of U.S. Patent Application No. 16/529,402, filed August 1, 2019, the contents of which are incorporated herein by reference in their entirety.

FIELD

[0002] This specification relates to a system using intermittent behind-the-meter power.

BACKGROUND

[0003] “Electrical grid” or “grid,” as used herein, refers to a Wide Area Synchronous Grid (also known as an Interconnection), and is a regional scale or greater electric power grid that that operates at a synchronized frequency and is electrically tied together during normal system conditions. An electrical grid delivers electricity from generation stations to consumers. An electrical grid includes: (i) generation stations that produce electrical power at large scales for delivery through the grid, (ii) high voltage transmission lines that carry that power from the generation stations to demand centers, and (iii) distribution networks carry that power to individual customers.

[0004] Figure 1 illustrates a typical electrical grid, such as a North American Interconnection or the synchronous grid of Continental Europe (formerly known as the UCTE grid). The electrical grid of Figure 1 can be described with respect to the various segments that make up the grid.

[0005] A generation segment 102 includes one or more generation stations that produce utility-scale electricity (typically >50MW), such as a nuclear plant 102a, a coal plant 102b, a wind power station (i.e., wind farm) 102c, and/or a photovoltaic power station (i.e., a solar farm) 102d. Generation stations are differentiated from building-mounted and other decentralized or local wind or solar power applications because they supply power at the utility level and scale (>50MW), rather than to a local user or users. The primary purpose of generation stations is to produce power for distribution through the grid, and in exchange for payment for the supplied electricity. Each of the generation stations 102a-d includes power generation equipment 102e-h, respectively, typically capable of supply utility-scale power (>50MW). For example, the power generation equipment 102g at wind power station 102c includes wind

turbines, and the power generation equipment 102h at photovoltaic power station 102d includes photovoltaic panels.

[0006] Each of the generation stations 102a-d may further include station electrical equipment 102i-l respectively. Station electrical equipment 102i-l are each illustrated in Figure 1 as distinct elements for simplified illustrative purposes only and may, alternatively or additionally, be distributed throughout the power generation equipment, 102e-h, respectively. For example, at wind power station 102c, each wind turbine may include transformers, frequency converters, power converters, and/or electrical filters. Energy generated at each wind turbine may be collected by distribution lines along strings of wind turbines and move through collectors, switches, transformers, frequency converters, power converters, electrical filters, and/or other station electrical equipment before leaving the wind power station 102c. Similarly, at photovoltaic power station 102d, individual photovoltaic panels and/or arrays of photovoltaic panels may include inverters, transformers, frequency converters, power converters, and/or electrical filters. Energy generated at each photovoltaic panel and/or array may be collected by distribution lines along the photovoltaic panels and move through collectors, switches, transformers, frequency converters, power converters, electrical filters, and/or other station electrical equipment before leaving the photovoltaic power station 102d.

[0007] Each generation station 102a-d may produce AC or DC electrical current which is then typically stepped up to a higher AC voltage before leaving the respective generation station. For example, wind turbines may typically produce AC electrical energy at 600V to 700V, which may then be stepped up to 34.5kV before leaving the generation station 102d. In some cases, the voltage may be stepped up multiple times and to a different voltage before exiting the generation station 102c. As another example, photovoltaic arrays may produce DC voltage at 600V to 900V, which is then inverted to AC voltage and may be stepped up to 34.5kV before leaving the generation station 102d. In some cases, the voltage may be stepped up multiple times and to a different voltage before exiting the generation station 102d.

[0008] Upon exiting the generation segment 102, electrical power generated at generation stations 102a-d passes through a respective Point of Interconnection (“POI”) 103 between a generation station (e.g., 102a-d) and the rest of the grid. A respective POI 103 represents the point of connection between a generation station’s (e.g. 102a-d) equipment and a transmission system (e.g., transmission segment 104) associated with electrical grid. In some cases, at the POI 103, generated power from generation stations 102a-d may be stepped up at transformer systems 103e-h to high voltage scales suitable for long-distance transmission along transmission lines 104a. Typically, the generate electrical energy leaving the POI 103 will be

at 115 kV AC or above, but in some cases it may be as low as, for example, 69 kV for shorter distance transmissions along transmission lines 104a. Each of transformer systems 103e-h may be a single transformer or may be multiple transformers operating in parallel or series and may be co-located or located in geographically distinct locations. Each of the transformer systems 103e-h may include substations and other links between the generation stations 102a-d and the transmission lines 104a.

[0009] A key aspect of the POI 103 is that this is where generation-side metering occurs. One or more utility-scale generation-side meters 103a-d (e.g., settlement meters) are located at settlement metering points at the respective POI 103 for each generation station 102a-d. The utility-scale generation-side meters 103a-d measure power supplied from generation stations 102a-d into the transmission segment 104 for eventual distribution throughout the grid.

[0010] For electricity consumption, the price consumers pay for power distributed through electric power grids is typically composed of, among other costs, Generation, Administration, and Transmission & Distribution (“T&D”) costs. T&D costs represent a significant portion of the overall price paid by consumers for electricity. These costs include capital costs (land, equipment, substations, wire, etc.), costs associated with electrical transmission losses, and operation and maintenance costs.

[0011] For utility-scale electricity supply, operators of generation stations (e.g., 102a-d) are paid a variable market price for the amount of power the operator generates and provides to the grid, which is typically determined via a power purchase agreement (PPA) between the contracting parties to the PPA or locational marginal pricing (LMP). The amount of power the generation station operator generates and provides to the grid is measured by utility-scale generation-side meters (e.g., 103a-d) at settlement metering points. As illustrated in Figure 1, the utility-scale generation-side meters 103a-d are shown on a low side of the transformer systems 103e-h), but they may alternatively be located within the transformer systems 103e-h or on the high side of the transformer systems 103e-h. A key aspect of a utility-scale generation-side meter is that it is able to meter the power supplied from a specific generation station into the grid. As a result, the grid operator can use that information to calculate and process payments for power supplied from the generation station to the grid. That price paid for the power supplied from the generation station is then subject to T&D costs, as well as other costs, in order to determine the price paid by consumers.

[0012] After passing through the utility-scale generation-side meters in the POI 103, the power originally generated at the generation stations 102a-d is transmitted onto and along the transmission lines 104a in the transmission segment 104. Typically, the electrical energy is

transmitted as AC at 115 kV+ or above, though it may be as low as 69kV for short transmission distances. In some cases, the transmission segment 104 may include further power conversions to aid in efficiency or stability. For example, transmission segment 104 may include high-voltage DC (“HVDC”) portions (along with conversion equipment) to aid in frequency synchronization across portions of the transmission segment 104. As another example, transmission segment 104 may include transformers to step AC voltage up and then back down to aid in long distance transmission (e.g., 230kV, 500 kV, 765 kV, etc.).

[0013] Power generated at the generation stations 104a-d is ultimately destined for use by consumers connected to the grid. Once the energy has been transmitted along the transmission segment 104, the voltage will be stepped down by transformer systems 105a-c in the step down segment 105 so that it can move into the distribution segment 106.

[0014] In the distribution segment 106, distribution networks 106a-c take power that has been stepped down from the transmission lines 104a and distribute it to local customers, such as local sub-grids (illustrated at 106a), industrial customers, including large EV charging networks (illustrated at 106b), and/or residential and retail customers, including individual EV charging stations (illustrated at 106c). Customer meters 106d, 106f measure the power used by each of the grid-connected customers in distribution networks 106a-c. Customer meters 106d are typically load meters that are unidirectional and measure power use. Some of the local customers in the distribution networks 106a-d may have local wind or solar power systems 106e owned by the customer. As discussed above, these local customer power systems 106e are decentralized and supply power directly to the customer(s). Customers with decentralized wind or solar power systems 106e may have customer meters 106f that are bidirectional or net-metering meters that can track when the local customer power systems 106e produce power in excess of the customer’s use, thereby allowing the utility to provide a credit to the customer’s monthly electricity bill. Customer meters 106d, 106f differ from utility-scale generation-side meters (e.g., settlement meters) in at least the following characteristics: design (electro-mechanical or electronic vs current transformer), scale (typically less than 1600 amps vs. typically greater than 50MW; typically less than 600V vs. typically greater than 14kV), primary function (use vs. supply metering), economic purpose (credit against use vs payment for power), and location (in a distribution network at point of use vs. at a settlement metering point at a Point of Interconnection between a generation station and a transmission line).

[0015] To maintain stability of the grid, the grid operator strives to maintain a balance between the amount of power entering the grid from generation stations (e.g., 102a-d) and the amount

of grid power used by loads (e.g., customers in the distribution segment 106). In order to maintain grid stability and manage congestion, grid operators may take steps to reduce the supply of power arriving from generation stations (e.g., 102a-d) when necessary (e.g., curtailment). Particularly, grid operators may decrease the market price paid for generated power to dis-incentivize generation stations (e.g., 102a-d) from generating and supplying power to the grid. In some cases, the market price may even go negative such that generation station operators must pay for power they allow into the grid. In addition, some situations may arise where grid operators explicitly direct a generation station (e.g., 102a-d) to reduce or stop the amount of power the station is supplying to the grid.

[0016] Power market fluctuations, power system conditions (e.g., power factor fluctuation or generation station startup and testing), and operational directives resulting in reduced or discontinued generation all can have disparate effects on renewal energy generators and can occur multiple times in a day and last for indeterminate periods of time. Curtailment, in particular, is particularly problematic.

[0017] According to the National Renewable Energy Laboratory's Technical Report TP-6A20-60983 (March 2014):

[0018] [C]urtailment [is] a reduction in the output of a generator from what it could otherwise produce given available resources (e.g., wind or sunlight), typically on an involuntary basis. Curtailments can result when operators or utilities command wind and solar generators to reduce output to minimize transmission congestion or otherwise manage the system or achieve the optimal mix of resources. Curtailment of wind and solar resources typically occurs because of transmission congestion or lack of transmission access, but it can also occur for reasons such as excess generation during low load periods that could cause baseload generators to reach minimum generation thresholds, because of voltage or interconnection issues, or to maintain frequency requirements, particularly for small, isolated grids. Curtailment is one among many tools to maintain system energy balance, which can also include grid capacity, hydropower and thermal generation, demand response, storage, and institutional changes. Deciding which method to use is primarily a matter of economics and operational practice.

[0019] "Curtailment" today does not necessarily mean what it did in the early 2000s. Two separate changes in the electric sector have shaped curtailment practices since that time: the utility-scale deployment of wind power, which has no fuel

cost, and the evolution of wholesale power markets. These simultaneous changes have led to new operational challenges but have also expanded the array of market-based tools for addressing them.

[0020] Practices vary significantly by region and market design. In places with centrally-organized wholesale power markets and experience with wind power, manual wind energy curtailment processes are increasingly being replaced by transparent offer-based market mechanisms that base dispatch on economics. Market protocols that dispatch generation based on economics can also result in renewable energy plants generating less than what they could potentially produce with available wind or sunlight. This is often referred to by grid operators by other terms, such as “downward dispatch.” In places served primarily by vertically integrated utilities, power purchase agreements (PPAs) between the utility and the wind developer increasingly contain financial provisions for curtailment contingencies.

[0021] *****

[0022] Some reductions in output are determined by how a wind operator values dispatch versus non-dispatch. Other curtailments of wind are determined by the grid operator in response to potential reliability events. Still other curtailments result from overdevelopment of wind power in transmission-constrained areas.

[0023] Dispatch below maximum output (curtailment) can be more of an issue for wind and solar generators than it is for fossil generation units because of differences in their cost structures. The economics of wind and solar generation depend on the ability to generate electricity whenever there is sufficient sunlight or wind to power their facilities.

[0024] Because wind and solar generators have substantial capital costs but no fuel costs (i.e., minimal variable costs), maximizing output improves their ability to recover capital costs. In contrast, fossil generators have higher variable costs, such as fuel costs. Avoiding these costs can, depending on the economics of a specific generator, to some degree reduce the financial impact of curtailment, especially if the generator's capital costs are included in a utility's rate base.

[0025] Curtailment may result in available energy being wasted because solar and wind operators have zero variable cost (which may not be true to the same extent for fossil generation units which can simply reduce the amount of fuel that is being used). With wind generation, in particular, it may also take some time for a wind farm to become fully operational following

curtailment. As such, until the time that the wind farm is fully operational, the wind farm may not be operating with optimum efficiency and/or may not be able to provide power to the grid.

SUMMARY

[0026] In an example, a system includes a plurality of computing systems and a datacenter control system. The plurality of computing systems receive behind-the-meter (“BTM”) power from a BTM power source. The datacenter control system is configured to monitor a set of conditions. The set of monitored conditions may include BTM power availability at the plurality of computing systems. The datacenter control system is further configured to perform a weighted analysis using the set of monitored conditions, and based on the weighted analysis, modulate operating attributes of one or more computing systems of the plurality of computing systems.

[0027] In another example, a method includes monitoring, at a control system, a set of conditions. The set of monitored conditions includes BTM power availability at a plurality of computing systems, and the plurality of computing systems receive BTM power from a BTM power source. The method further includes performing a weighted analysis using the set of monitored conditions, and based on the weighted analysis, modulating operating attributes of one or more computing systems of the plurality of computing systems.

[0028] In a further example, non-transitory computer readable medium is described that is configured to store instructions, that when executed by a computing system, causes the computing system to perform functions consistent with the method steps described above.

[0029] Other aspects of the present invention will be apparent from the following description and claims.

BRIEF DESCRIPTION OF THE FIGURES

[0030] Figure 1 shows a typical electrical grid.

[0031] Figure 2 shows a behind-the-meter arrangement, including one or more flexible datacenters, according to one or more example embodiments.

[0032] Figure 3 shows a block diagram of a remote master control system, according to one or more example embodiments.

[0033] Figure 4 a block diagram of a generation station, according to one or more example embodiments.

[0034] Figure 5 shows a block diagram of a flexible datacenter, according to one or more example embodiments.

[0035] Figure 6A shows a structural arrangement of a flexible datacenter, according to one or more example embodiments.

[0036] Figure 6B shows a set of computing systems arranged in a straight configuration, according to one or more example embodiments.

[0037] Figure 7 shows a control distribution system for a flexible datacenter, according to one or more example embodiments.

[0038] Figure 8 shows a control distribution system for a fleet of flexible datacenters, according to one or more example embodiments.

[0039] Figure 9 shows a queue distribution system for a traditional datacenter and a flexible datacenter, according to one or more example embodiments.

[0040] Figure 10A shows a method of dynamic power consumption at a flexible datacenter using behind-the-meter power, according to one or more example embodiments.

[0041] Figure 10B shows a method of dynamic power delivery at a flexible datacenter using behind-the-meter power, according to one or more example embodiments.

[0042] Figure 11 shows a block diagram of a system for modulating operating attributes of computing systems at a flexible datacenter according to one or more embodiments.

[0043] Figure 12 shows a method for modulating operating attributes of one or more computing systems according to one or more embodiments.

DETAILED DESCRIPTION

[0044] Disclosed examples will now be described more fully hereinafter with reference to the accompanying drawings, in which some, but not all of the disclosed examples are shown. Different examples may be described and should not be construed as limited to the examples set forth herein.

[0045] As discussed above, the market price paid to generation stations for supplying power to the grid often fluctuates due to various factors, including the need to maintain grid stability and based on current demand and usage by connected loads in distribution networks. Due to these factors, situations can arise where generation stations are offered substantially lower prices to deter an over-supply of power to the grid. Although these situations typically exist temporarily, generation stations are sometimes forced to either sell power to the grid at much lower prices or adjust operations to decrease the amount of power generated. Furthermore,

some situations may even require generation stations to incur costs in order to offload power to the grid or to shut down generation temporarily.

[0046] The volatility in the market price offered for power supplied to the grid can be especially problematic for some types of generation stations. In particular, wind farms and some other types of renewable resource power producers may lack the ability to quickly adjust operations in response to changes in the market price offered for supplying power to the grid. As a result, power generation and management at some generation stations can be inefficient, which can frequently result in power being sold to the grid at low or negative prices. In some situations, a generation station may even opt to halt power generation temporarily to avoid such unfavorable pricing. As such, the time required to halt and to restart the power generation at a generation station can reduce the generation station's ability to take advantage of rising market prices for power supplied to the grid.

[0047] Example embodiments provided herein aim to assist generation stations in managing power generation operations and avoid unfavorable power pricing situations like those described above. In particular, example embodiments may involve providing a load that is positioned behind-the-meter ("BTM") and enabling the load to utilize power received behind-the-meter at a generation station in a timely manner.

[0048] For purposes herein, a generation station is considered to be configured for the primary purpose of generating utility-scale power for supply to the electrical grid (e.g., a Wide Area Synchronous Grid or a North American Interconnect).

[0049] In one embodiment, equipment located behind-the-meter ("BTM equipment") is equipment that is electrically connected to a generation station's power generation equipment behind (i.e., prior to) the generation station's POI with an electrical grid.

[0050] In one embodiment, behind-the-meter power ("BTM power") is electrical power produced by a generation station's power generation equipment and utilized behind (i.e., prior to) the generation station's POI with an electrical grid.

[0051] In another embodiment, equipment may be considered behind-the-meter if it is electrically connected to a generation station that is subject to metering by a utility-scale generation-side meter (e.g., settlement meter), and the BTM equipment receives power from the generation station, but the power received by the BTM equipment from the generation station has not passed through the utility-scale generation-side meter. In one embodiment, the utility-scale generation-side meter for the generation station is located at the generation station's POI. In another embodiment, the utility-scale generation-side meter for the generation

station is at a location other than the POI for the generation station – for example, a substation between the generation station and the generation station's POI.

[0052] In another embodiment, power may be considered behind-the-meter if it is electrical power produced at a generation station that is subject to metering by a utility-scale generation-side meter (e.g., settlement meter), and the BTM power is utilized before being metered at the utility-scale generation-side meter. In one embodiment, the utility-scale generation-side meter for the generation station is located at the generation station's POI. In another embodiment, the utility-scale generation-side meter for the generation station is at a location other than the POI for the generation station – for example, a substation between the generation station and the generation station's POI.

[0053] In another embodiment, equipment may be considered behind-the-meter if it is electrically connected to a generation station that supplies power to a grid, and the BTM equipment receives power from the generation station that is not subject to T&D charges, but power received from the grid that is supplied by the generation station is subject to T&D charges.

[0054] In another embodiment, power may be considered behind-the-meter if it is electrical power produced at a generation station that supplies power to a grid, and the BTM power is not subject to T&D charges before being used by electrical equipment, but power received from the grid that is supplied by the generation station is subject to T&D charges.

[0055] In another embodiment, equipment may be considered behind-the-meter if the BTM equipment receives power generated from the generation station and that received power is not routed through the electrical grid before being delivered to the BTM equipment.

[0056] In another embodiment, power may be considered behind-the-meter if it is electrical power produced at a generation station, and BTM equipment receives that generated power, and that generated power received by the BTM equipment is not routed through the electrical grid before being delivered to the BTM equipment.

[0057] For purposes herein, BTM equipment may also be referred to as a behind-the-meter load ("BTM load") when the BTM equipment is actively consuming BTM power.

[0058] Beneficially, where BTM power is not subject to traditional T&D costs, a wind farm or other type of generation station can be connected to BTM loads which can allow the generation station to selectively avoid the adverse or less-than optimal cost structure occasionally associated with supplying power to the grid by shunting generated power to the BTM load.

[0059] An arrangement that positions and connects a BTM load to a generation station can offer several advantages. In such arrangements, the generation station may selectively choose

whether to supply power to the grid or to the BTM load, or both. The operator of a BTM load may pay to utilize BTM power at a cost less than that charged through a consumer meter (e.g., 106d, 106f) located at a distribution network (e.g., 106a-c) receiving power from the grid. The operator of a BTM load may additionally or alternatively pay less than the market rate to consume excess power generated at the generation station during curtailment. As a result, the generation station may direct generated power based on the “best” price that the generation station can receive during a given time frame, and/or the lowest cost the generation station may incur from negative market pricing during curtailment. The “best” price may be the highest price that the generation station may receive for its generated power during a given duration, but can also differ within embodiments and may depend on various factors, such as a prior PPA. In one example, by having a behind-the-meter option available, a generation station may transition from supplying all generated power to the grid to supplying some or all generated power to one or more BTM loads when the market price paid for power by grid operators drops below a predefined threshold (e.g., the price that the operator of the BTM load is willing to pay the generation station for power). Thus, by having an alternative option for power consumption (i.e., one or more BTM loads), the generation station can selectively utilize the different options to maximize the price received for generated power. In addition, the generation station may also utilize a BTM load to avoid or reduce the economic impact in situations when supplying power to the grid would result in the generation station incurring a net cost.

[0060] Providing BTM power to a load can also benefit the BTM load operator. A BTM load may be able to receive and utilize BTM power received from the generation station at a cost that is lower than the cost for power from the grid (e.g., at a customer meter 106d, 106f). This is primarily due to avoidance in T&D costs and the market effects of curtailment. As indicated above, the generation station may be willing to divert generated power to the BTM load rather than supplying the grid due to changing market conditions, or during maintenance periods, or for other non-market conditions. Furthermore, in some situations, the BTM load may even be able to obtain and utilize BTM power from a generation station at no cost or even at negative pricing since the generation station may be receiving tax credits (e.g., Production Tax Credits) for produced wind or is slow to self-curtail.

[0061] Another example of cost-effective use of BTM power is when the generation station 202 is selling power to the grid at a negative price that is offset by a production tax credit. In certain circumstances, the value of the production tax credit may exceed the price the generation station 202 would have to pay to the grid power to offload generation’s station 202 generated power. Advantageously, one or more flexible datacenters 220 may take the generated

power behind-the-meter, thereby allowing the generation station 202 to produce and obtain the production tax credit, while selling less power to the grid at the negative price.

[0062] Another example of cost-effective behind-the-meter power is when the generation station 202 is selling power to the grid at a negative price because the grid is oversupplied and/or the generation station 202 is instructed to stand down and stop producing altogether. A grid operator may select and direct certain generation stations to go offline and stop supplying power to the grid. Advantageously, one or more flexible datacenters may be used to take power behind-the-meter, thereby allowing the generation station 202 to stop supplying power to the grid, but still stay online and make productive use of the power generated.

[0063] Another example of beneficial behind-the-meter power use is when the generation station 202 is producing power that is, with reference to the grid, unstable, out of phase, or at the wrong frequency, or the grid is already unstable, out of phase, or at the wrong frequency. A grid operator may select certain generation stations to go either offline and stop producing power, or to take corrective action with respect to the grid power stability, phase, or frequency. Advantageously, one or more flexible datacenters 220 may be used to selectively consume power behind-the-meter, thereby allowing the generation station 202 to stop providing power to the grid and/or provide corrective feedback to the grid.

[0064] Another example of beneficial behind-the-meter power use is that cost-effective behind-the-meter power availability may occur when the generation station 202 is starting up or testing. Individual equipment in the power generation equipment 210 may be routinely offline for installation, maintenance, and/or service and the individual units must be tested prior to coming online as part of overall power generation equipment 210. During such testing or maintenance time, one or more flexible datacenters may be intermittently powered by the one or more units of the power generation equipment 210 that are offline from the overall power generation equipment 210.

[0065] Another example of beneficial behind-the-meter power use is that datacenter control systems at the flexible datacenters 220 may quickly ramp up and ramp down power consumption by computing systems in the flexible datacenters 220 based on power availability from the generation station 202. For instance, if the grid requires additional power and signals the demand via a higher local price for power, the generation station 202 can supply the grid with power nearly instantly by having active flexible datacenters 220 quickly ramp down and turn off computing systems (or switch to a stored energy source), thereby reducing an active BTM load.

[0066] Another example of beneficial behind-the-meter power use is in new photovoltaic generation stations 202. For example, it is common to design and build new photovoltaic generation stations with a surplus of power capacity to account for degradation in efficiency of the photovoltaic panels over the life of the generation stations. Excess power availability at the generation station can occur when there is excess local power generation and/or low grid demand. In high incident sunlight situations, a photovoltaic generation station 202 may generate more power than the intended capacity of generation station 202. In such situations, a photovoltaic generation station 202 may have to take steps to protect its equipment from damage, which may include taking one or more photovoltaic panels offline or shunting their voltage to dummy loads or the ground. Advantageously, one or more flexible datacenters (e.g., the flexible datacenters 220) may take power behind-the-meter at the Generations Station 202, thereby allowing the generation station 202 to operate the power generation equipment 210 within operating ranges while the flexible datacenters 220 receive BTM power without transmission or distribution costs.

[0067] Thus, for at least the reasons described herein, arrangements that involves providing a BTM load as an alternative option for a generation station to direct its generated power to can serve as a mutually beneficial relationship in which both the generation station and the BTM load can economically benefit. The above-noted examples of beneficial use of BTM power are merely exemplary and are not intended to limit the scope of what one of ordinary skill in the art would recognize as benefits to unutilized BTM power capacity, BTM power pricing, or BTM power consumption.

[0068] Within example embodiments described herein, various types of utility-scale power producers may operate as generation stations 202 that are capable of supplying power to one or more loads behind-the-meter. For instance, renewable energy sources (e.g., wind, solar, hydroelectric, wave, water current, tidal), fossil fuel power generation sources (coal, natural gas), and other types of power producers (e.g., nuclear power) may be positioned in an arrangement that enables the intermittent supply of generated power behind-the-meter to one or more BTM loads. One of ordinary skill in the art will recognize that the generation station 202 may vary based on an application or design in accordance with one or more example embodiments.

[0069] In addition, the particular arrangement (e.g., connections) between the generation station and one or more BTM loads can vary within examples. In one embodiment, a generation station may be positioned in an arrangement wherein the generation station selectively supplies power to the grid and/or to one or more BTM loads. As such, power cost-analysis and other

factors (e.g., predicted weather conditions, contractual obligations, etc.) may be used by the generation station, a BTM load control system, a remote master control system, or some other system or enterprise, to selectively output power to either the grid or to one or more BTM loads in a manner that maximizes revenue to the generation station. In such an arrangement, the generation station may also be able to supply both the grid and one or more BTM loads simultaneously. In some instances, the arrangement may be configured to allow dynamic manipulation of the percentage of the overall generated power that is supplied to each option at a given time. For example, in some time periods, the generation station may supply no power to the BTM load.

[0070] In addition, the type of loads that are positioned behind-the-meter can vary within example embodiments. In general, a load that is behind-the-meter may correspond to any type of load capable of receiving and utilizing power behind-the-meter from a generation station. Some examples of loads include, but are not limited to, datacenters and electric vehicle (EV) charging stations.

[0071] Preferred BTM loads are loads that can be subject to intermittent power supply because BTM power may be available intermittently. In some instances, the generation station may generate power intermittently. For example, wind power station 102c and/or photovoltaic power station 102d may only generate power when resource are available or favorable. Additionally or alternatively, BTM power availability at a generation station may only be available intermittently due to power market fluctuations, power system conditions (e.g., power factor fluctuation or generation station startup and testing), and/or operational directives from grid operators or generation station operators.

[0072] Some example embodiments of BTM loads described herein involve using one or more computing systems to serve as a BTM load at a generation station. In particular, the computing system or computing systems may receive power behind-the-meter from the generation station to perform various computational operations, such as processing or storing information, performing calculations, mining for cryptocurrencies, supporting blockchain ledgers, and/or executing applications, etc. Multiple computing systems positioned behind-the-meter may operate as part of a “flexible” datacenter that is configured to operate only intermittently and to receive and utilize BTM power to carry out various computational operations similar to a traditional datacenter. In particular, the flexible datacenter may include computing systems and other components (e.g., support infrastructure, a control system) configured to utilize BTM power from one or more generation stations. The flexible datacenter may be configured to use particular load ramping abilities (e.g., quickly increase or decrease power usage) to effectively

operate during intermittent periods of time when power is available from a generation station and supplied to the flexible datacenter behind-the-meter, such as during situations when supplying generated power to the grid is not favorable for the generation station. In some instances, the amount of power consumed by the computing systems at a flexible datacenter can be ramped up and down quickly, and potentially with high granularity (i.e., the load can be changed in small increments if desired). This may be done based on monitored power system conditions or other information analyses as discussed herein. As recited above, this can enable a generation station to avoid negative power market pricing and to respond quickly to grid directives. And by extension, the flexible datacenter may obtain BTM power at a price lower than the cost for power from the grid.

[0073] Various types of computing systems can provide granular behind-the-meter ramping. Preferably, the computing systems utilizing BTM power is utilized to perform computational tasks that are immune to, or not substantially hindered by, frequent interruptions or slow-downs in processing as the computing systems ramp down or up. In some embodiments, a control system may be used to activate or de-activate one or more computing systems in an array of computing systems sited behind the meter. For example, the control system may provide control instructions to one or more blockchain miners (e.g., a group of blockchain miners), including instructions for powering on or off, adjusting frequency of computing systems performing operations (e.g., adjusting the processing frequency), adjusting the quantity of operations being performed, and when to operate within a low power mode (if available).

[0074] Within examples, a control system may correspond to a specialized computing system or may be a computing system within a flexible datacenter serving in the role of the control system. The location of the control system can vary within examples as well. For instance, the control system may be located at a flexible datacenter or physically separate from the flexible datacenter. In some examples, the control system may be part of a network of control systems that manage computational operations, power consumption, and other aspects of a fleet of flexible datacenters.

[0075] Some embodiments may involve using one or more control systems to direct time-insensitive (e.g., interruptible) computational tasks to computational hardware, such as central processing units (CPUs) and graphics processing units (GPUs), sited behind the meter, while other hardware is sited in front of the meter (i.e., consuming metered grid power via a customer meter (e.g., 106d, 106f)) and possibly remote from the behind-the-meter hardware. As such, parallel computing processes, such as Monte Carlo simulations, batch processing of financial transactions, graphics rendering, machine learning, neural network processing, queued

operations, and oil and gas field simulation models, are good candidates for such interruptible computational operations.

[0076] Figure 2 shows a behind-the-meter arrangement, including one or more flexible datacenters, according to one or more example embodiments. Dark arrows illustrate a typical power delivery direction. Consistent with Figure 1, the arrangement illustrates a generation station 202 in the generation segment 102 of a Wide-Area Synchronous Grid. The generation station 202 supplies utility-scale power (typically >50MW) via a generation power connection 250 to the Point of Interconnection 103 between the generation station 202 and the rest of the grid. Typically, the power supplied on connection 250 may be at 34.5kV AC, but it may be higher or lower. Depending on the voltage at connection 250 and the voltage at transmission lines 104a, a transformer system 203 may step up the power supplied from the generation station 202 to high voltage (e.g., 115kV+ AC) for transmission over connection 252 and onto transmission lines 104a of transmission segment 104. Grid power carried on the transmission segment 104 may be from generation station 202 as well as other generation stations (not shown). Also consistent with Figure 1, grid power is consumed at one or more distribution networks, including example distribution network 206. Grid power may be taken from the transmission lines 104a via connector 254 and stepped down to distribution network voltages (e.g., typically 4kV to 26kV AC) and sent into the distribution networks, such as distribution network 206 via distribution line 256. The power on distribution line 256 may be further stepped down (not shown) before entering individual consumer facilities such as a remote master control system 262 and/or traditional datacenters 260 via customer meters 206A, which may correspond to customer meters 106d in Figure 1, or customer meters 106f in Figure 1 if the respective consumer facility includes a local customer power system, such as 106e (not shown in Figure 2).

[0077] Consistent with Figure 1, power entering the grid from generation station 202 is metered by a utility-scale generation-side meter. A utility-scale generation-side meter 253 is shown on the low side of transformer system 203 and an alternative location is shown as 253A on the high side of transformer system 203. Both locations may be considered settlement metering points for the generation station 202 at the POI 103. Alternatively, a utility-scale generation-side meter for the generation station 202 may be located at another location consistent with the descriptions of such meters provided herein.

[0078] Generation station 202 includes power generation equipment 210, which may include, as examples, wind turbines and/or photovoltaic panels. Power generation equipment 210 may further include other electrical equipment, including but not limited to switches, busses,

collectors, inverters, power quality and conditioning equipment, and power unit transformers (e.g., transformers in wind turbines).

[0079] As illustrated in Figure 2, generation station 202 is configured to connect with BTM equipment which may function as BTM loads. In the illustrated embodiment of Figure 2, the BTM equipment includes flexible datacenters 220. Various configurations to supply BTM power to flexible datacenters 220 within the arrangement of Figure 2 are described herein.

[0080] In one configuration, generated power may travel from the power generation equipment 210 over one or more connectors 230A, 230B to one or more electrical busses 240A, 240B, respectively. Each of the connectors 230A, 230B may be a switched connector such that power may be routed independently to 240A and/or 240B. For illustrative purposes only, connector 230B is shown with an open switch, and connector 230A is shown with a closed switch, but either or both may be reversed in some embodiments. Aspects of this configuration can be used in various embodiments when BTM power is supplied without significant power conversion to BTM loads.

[0081] In various configurations, the busses 240A and 240B may be separated by an open switch 240C or combined into a common bus by a closed switch 240C.

[0082] In another configuration, generated power may travel from the power generation equipment 210 to the high side of a local step-down transformer 214. The generated power may then travel from the low side of the local step-down transformer 214 over one or more connectors 232A, 232B to the one or more electrical busses 240A, 240B, respectively. Each of the connectors 232A, 232B may be a switched connector such that power may be routed independently to 240A and/or 240B. For illustrative purposes only, connector 232A is shown with an open switch, and connector 232B is shown with a closed switch, but either or both may be reversed in some embodiments. Aspects of this configuration can be used when it is preferable to connect BTM power to the power generation equipment 210, but the generated power must be stepped down prior to use at the BTM loads.

[0083] In another configuration, generated power may travel from the power generation equipment 210 to the low side of a local step-up transformer 212. The generated power may then travel from the high side of the local step-up transformer 212 over one or more connectors 234A, 234B to the one or more electrical busses 240A, 240B, respectively. Each of the connectors 234A, 234B may be a switched connector such that power may be routed independently to 240A and/or 240B. For illustrative purposes only, both connectors 234A, 234B are shown with open switches, but either or both may be closed in some embodiments.

Aspects of this configuration can be used when it is preferable to connect BTM power to the outbound connector 250 or the high side of the local step-up transformer 212.

[0084] In another configuration, generated power may travel from the power generation equipment 210 to the low side of the local step-up transformer 212. The generated power may then travel from the high side of the local step-up transformer 212 to the high side of local step-down transformer 213. The generated power may then travel from the low side of the local step-down transformer 213 over one or more connectors 236A, 236B to the one or more electrical buses 240A, 240B, respectively. Each of the connectors 236A, 236B may be a switched connector such that power may be routed independently to 240A and/or 240B. For illustrative purposes only, both connectors 236A, 236B are shown with open switches, but either or both may be closed in some embodiments. Aspects of this configuration can be used when it is preferable to connect BTM power to the outbound connector 250 or the high side of the local step-up transformer 212, but the power must be stepped down prior to use at the BTM loads.

[0085] In one embodiment, power generated at the generation station 202 may be used to power a generation station control system 216 located at the generation station 202, when power is available. The generation station control system 216 may typically control the operation of the generation station 202. Generated power used at the generation station control system 216 may be supplied from bus 240A via connector 216A and/or from bus 240B via connector 216B. Each of the connectors 216A, 216B may be a switched connector such that power may be routed independently to 240A and/or 240B. While the generation station control system 216 can consume BTM power when powered via bus 240A or bus 240B, the BTM power taken by the generation station control system 216 is insignificant in terms of rendering an economic benefit. Further, the generation station control system 216 is not configured to operate intermittently, as it generally must remain always on. Further still, the generation station control system 216 does not have the ability to quickly ramp a BTM load up or down. In some instances, the generation station control system 216 may receive and use power from the electrical grid.

[0086] In another embodiment, grid power may alternatively or additionally be used to power the generation station control system 216. As illustrated here, metered grid power from a distribution network, such as distribution network 206 for simplicity of illustration purposes only, may be used to power generation station control system 216 over connector 216C. Connector 216C may be a switched connector so that metered grid power to the generation station control system 216 can be switched on or off as needed. More commonly, metered grid

power would be delivered to the generation station control system 216 via a separate distribution network (not shown), and also over a switched connector. Any such grid power delivered to the generation station control system 216 is metered by a customer meter 206A and subject to T&D costs.

[0087] In another embodiment, when power generation equipment 210 is in an idle or off state and not generating power, grid power may backfeed into generation station 202 through POI 103 and such grid power may power the generation station control system 216.

[0088] In some configurations, an energy storage system 218 may be connected to the generation station 202 via connector 218A, which may be a switched connector. For illustrative purposes only, connector 218A is shown with an open switch but in some embodiments it may be closed. The energy storage system 218 may be connected to bus 240A and/or bus 240B and store energy produced by the power generation equipment 210. The energy storage system may also be isolated from generation station 202 by switch 242A. In times of need, such as when the power generation equipment is in an idle or off state and not generating power, the energy storage system may feed power to, for example, the flexible datacenters 220. The energy storage system may also be isolated from the flexible datacenters 220 by switch 242B.

[0089] In a preferred embodiment, as illustrated, power generation equipment 210 supplies BTM power via connector 242 to flexible datacenters 220. The BTM power used by the flexible datacenters 220 was generated by the generation station 202 and did not pass through the POI 103 or utility-scale generation-side meter 253, and is not subject to T&D charges. Power received at the flexible datacenters 220 may be received through respective power input connectors 220A. Each of the respective connectors 220A may be a switched connector that can electrically isolate the respective flexible datacenter 220 from the connector 242. Power equipment 220B may be arranged between the flexible datacenters 220 and the connector 242. The power equipment 220B may include, but is not limited to, power conditioners, unit transformers, inverters, and isolation equipment. As illustrated, each flexible datacenter 220 may be served by a respective power equipment 220B. However, in another embodiment, one power equipment 220B may serve multiple flexible datacenter 220.

[0090] In one embodiment, flexible datacenters 220 may be considered BTM equipment located behind-the-meter and electrically connected to the power generation equipment 210 behind (i.e., prior to) the generation station's POI 103 with the rest of the electrical grid.

[0091] In one embodiment, BTM power produced by the power generation equipment 210 is utilized by the flexible datacenters 220 behind (i.e., prior to) the generation station's POI with an electrical grid.

[0092] In another embodiment, flexible datacenters 220 may be considered BTM equipment located behind-the-meter as the flexible datacenters 220 are electrically connected to the generation station 202, and generation station 202 is subject to metering by utility-scale generation-side meter 253 (or 253A, or another utility-scale generation-side meter), and the flexible datacenters 220 receive power from the generation station 202, but the power received by the flexible datacenters 220 from the generation station 202 has not passed through a utility-scale generation-side meter. In this embodiment, the utility-scale generation-side meter 253 (or 253A) for the generation station 202 is located at the generation station's 202 POI 103. In another embodiment, the utility-scale generation-side meter for the generation station 202 is at a location other than the POI for the generation station 202 – for example, a substation (not shown) between the generation station 202 and the generation station's POI 103.

[0093] In another embodiment, power from the generation station 202 is supplied to the flexible datacenters 220 as BTM power, where power produced at the generation station 202 is subject to metering by utility-scale generation-side meter 253 (or 253A, or another utility-scale generation-side meter), but the BTM power supplied to the flexible datacenters 220 is utilized before being metered at the utility-scale generation-side meter 253 (or 253A, or another utility-scale generation-side meter). In this embodiment, the utility-scale generation-side meter 253 (or 253A) for the generation station 202 is located at the generation station's 202 POI 103. In another embodiment, the utility-scale generation-side meter for the generation station 202 is at a location other than the POI for the generation station 202 – for example, a substation (not shown) between the generation station 202 and the generation station's POI 103.

[0094] In another embodiment, flexible datacenters 220 may be considered BTM equipment located behind-the-meter as they are electrically connected to the generation station 202 that supplies power to the grid, and the flexible datacenters 220 receive power from the generation station 202 that is not subject to T&D charges, but power otherwise received from the grid that is supplied by the generation station 202 is subject to T&D charges.

[0095] In another embodiment, power from the generation station 202 is supplied to the flexible datacenters 220 as BTM power, where electrical power is generated at the generation station 202 that supplies power to a grid, and the generated power is not subject to T&D charges before being used by flexible datacenters 220, but power otherwise received from the connected grid is subject to T&D charges.

[0096] In another embodiment, flexible datacenters 220 may be considered BTM equipment located behind-the-meter because they receive power generated from the generation station 202 intended for the grid, and that received power is not routed through the electrical grid before being delivered to the flexible datacenters 220.

[0097] In another embodiment, power from the generation station 202 is supplied to the flexible datacenters 220 as BTM power, where electrical power is generated at the generation station 202 for distribution to the grid, and the flexible datacenters 220 receive that power, and that received power is not routed through the electrical grid before being delivered to the flexible datacenters 220.

[0098] In another embodiment, metered grid power may alternatively or additionally be used to power one or more of the flexible datacenters 220, or a portion within one or more of the flexible datacenters 220. As illustrated here for simplicity, metered grid power from a distribution network, such as distribution network 206, may be used to power one or more flexible datacenters 220 over connector 256A and/or 256B. Each of connector 256A and/or 256B may be a switched connector so that metered grid power to the flexible datacenters 220 can be switched on or off as needed. More commonly, metered grid power would be delivered to the flexible datacenters 220 via a separate distribution network (not shown), and also over switched connectors. In some instances, grid power may be distributed to the flexible datacenters 220 via a backfeed process that involves measuring the amount of grid power via a subtraction meter. Any such grid power delivered to the flexible datacenters 220 is metered by customer meters 206A and subject to T&D costs. In one embodiment, connector 256B may supply metered grid power to a portion of one or more flexible datacenters 220. For example, connector 256B may supply metered grid power to control and/or communication systems for the flexible datacenters 220 that need constant power and cannot be subject to intermittent BTM power. Connector 242 may supply solely BTM power from the generation station 202 to high power demand computing systems within the flexible datacenters 220, in which case at least a portion of each flexible datacenters 220 so connected is operating as a BTM load. In another embodiment, connector 256A and/or 256B may supply all power used at one or more of the flexible datacenters 220, in which case each of the flexible datacenters 220 so connected would not be operating as a BTM load.

[0099] In another embodiment, when power generation equipment 210 is in an idle or off state and not generating power, grid power may backfeed into generation station 202 through POI 103 and such grid power may power the flexible datacenters 220. Backfeed may enable power

generation equipment 210 to maintain a safe state using minimal backfed power until operations resume at the power generation equipment 210.

[0100] The flexible datacenters 220 are shown in an example arrangement relative to the generation station 202. Particularly, generated power from the generation station 202 may be supplied to the flexible datacenters 220 through a series of connectors and/or busses (e.g., 232B, 240B, 242, 220A). As illustrated, in other embodiments, connectors between the power generation equipment 210 and other components may be switched open or closed, allowing other pathways for power transfer between the power generation equipment 210 and components, including the flexible datacenters 220. Additionally, the connector arrangement shown is illustrative only and other circuit arrangements are contemplated within the scope of supplying BTM power to a BTM load at generation station 202. For example, there may be more or fewer transformers, or one or more of transformers 212, 213, 214 may be transformer systems with multiple steppings and/or may include additional power equipment including but not limited to power conditioners, filters, switches, inverters, and/or AC/DC-DC/AC isolators. As another example, metered grid power connections to flexible datacenters 220 are shown via both 256A and 256B; however, a single connection may connect one or more flexible datacenters 220 (or power equipment 220B) to metered grid power and the one or more flexible datacenters 220 (or power equipment 220B) may include switching apparatus to direct BTM power and/or metered grid power to control systems, communication systems, and/or computing systems as desired.

[0101] In some examples, BTM power may arrive at the flexible datacenters 220 in a three-phase AC format. As such, power equipment (e.g., power equipment 220B) at one or more of the flexible datacenters 220 may enable each flexible datacenter 220 to use one or more phases of the power. For instance, the flexible datacenters 220 may utilize power equipment (e.g., power equipment 220B, or alternatively or additionally power equipment that is part of the flexible datacenter 220) to convert BTM power received from the generation station 202 for use at computing systems at each flexible datacenter 220. In other examples, the BTM power may arrive at one or more of the flexible datacenters 220 as DC power. As such, the flexible datacenters 220 may use the DC power to power computing systems. In some such examples, the DC power may be routed through a DC-to-DC converter that is part of power equipment 220B and/or flexible datacenter 220.

[0102] In some configurations, a flexible datacenter 220 may be arranged to only have access to power received behind-the-meter from a generation station 202. In the arrangement of Figure 2, the flexible datacenters 220 may be arranged only with a connection to the generation

station 202 and depend solely on power received behind-the-meter from the generation station 202. Alternatively or additionally, the flexible datacenters 220 may receive power from energy storage system 218.

[0103] In some configurations, one or more of the flexible datacenters 220 can be arranged to have connections to multiple sources that are capable of supplying power to a flexible datacenter 220. To illustrate a first example, the flexible datacenters 220 are shown connected to connector 242, which can be connected or disconnected via switches to the energy storage system 218 via connector 218A, the generation station 202 via bus 240B, and grid power via metered connector 256A. In one embodiment, the flexible datacenters 220 may selectively use power received behind-the-meter from the generation station 202, stored power supplied by the energy storage system 218, and/or grid power. For instance, flexible datacenters 220 may use power stored in the energy storage system 218 when costs for using power supplied behind-the-meter from the generation station 202 are disadvantageous. By having access to the energy storage system 218 available, the flexible datacenters 220 may use the stored power and allow the generation station 202 to subsequently refill the energy storage system 218 when cost for power behind-the-meter is low. Alternatively, the flexible datacenters 220 may use power from multiple sources simultaneously to power different components (e.g., a first set and a second set of computing systems). Thus, the flexible datacenters 220 may leverage the multiple connections in a manner that can reduce the cost for power used by the computing systems at the flexible datacenters 220. The flexible datacenters 220 control system or the remote master control system 262 may monitor power conditions and other factors to determine whether the flexible datacenters 220 should use power from either the generation station 202, grid power, the energy storage system 218, none of the sources, or a subset of sources during a given time range. Other arrangements are possible as well. For example, the arrangement of Figure 2 illustrates each flexible datacenter 220 as connected via a single connector 242 to energy storage system 218, generation station 202, and metered grid power via 256A. However, one or more flexible datacenters 220 may have independent switched connections to each energy source, allowing the one or more flexible datacenters 220 to operate from different energy sources than other flexible datacenters 220 at the same time.

[0104] The selection of which power source to use at a flexible datacenter (e.g., the flexible datacenters 220) or another type of BTM load can change based on various factors, such as the cost and availability of power from both sources, the type of computing systems using the power at the flexible datacenters 220 (e.g., some systems may require a reliable source of power for a long period), the nature of the computational operations being performed at the flexible

datacenters 220 (e.g., a high priority task may require immediate completion regardless of cost), and temperature and weather conditions, among other possible factors. As such, a datacenter control system at the flexible datacenters 220, the remote master control system 262, or another entity (e.g., an operator at the generation station 202) may also influence and/or determine the source of power that the flexible datacenters 220 use at a given time to complete computational operations.

[0105] In some example embodiments, the flexible datacenters 220 may use power from the different sources to serve different purposes. For example, the flexible datacenters 220 may use metered power from grid power to power one or more systems at the flexible datacenters 220 that are configured to be always-on (or almost always on), such as a control and/or communication system and/or one or more computing systems (e.g., a set of computing systems performing highly important computational operations). The flexible datacenters 220 may use BTM power to power other components within the flexible datacenters 220, such as one or more computing systems that perform less critical computational operations.

[0106] In some examples, one or more flexible datacenters 220 may be deployed at the generation station 202. In other examples, flexible datacenters 220 may be deployed at a location geographically remote from the generation station 202, while still maintaining a BTM power connection to the generation station 202.

[0107] In another example arrangement, the generation station 202 may be connected to a first BTM load (e.g., a flexible datacenter 220) and may supply power to additional BTM loads via connections between the first BTM load and the additional BTM loads (e.g., a connection between a flexible datacenter 220 and another flexible datacenter 220).

[0108] The arrangement in Figure 2, and components included therein, are for non-limiting illustration purposes and other arrangements are contemplated in examples. For instance, in another example embodiment, the arrangement of Figure 2 may include more or fewer components, such as more BTM loads, different connections between power sources and loads, and/or a different number of datacenters. In addition, some examples may involve one or more components within the arrangement of Figure 2 being combined or further divided.

[0109] Within the arrangement of Figure 2, a control system, such as the remote master control system 262 or another component (e.g., a control system associated with the grid operator, the generation station control system 216, or a datacenter control system associated with a traditional datacenter or one or more flexible datacenters) may use information to efficiently manage various operations of some of the components within the arrangement of Figure 2. For example, the remote master control system 262 or another component may manage distribution

and execution of computational operations at one or more traditional datacenters 260 and/or flexible datacenters 220 via one or more information-processing algorithms. These algorithms may utilize past and current information in real-time to manage operations of the different components. These algorithms may also make some predictions based on past trends and information analysis. In some examples, multiple computing systems may operate as a network to process information.

[0110] Information used to make decisions may include economic and/or power-related information, such as monitored power system conditions. Monitored power system conditions may include one or more of excess power generation at a generation station 202, excess power at a generation station 202 that a connected grid cannot receive, power generation at a generation station 202 subject to economic curtailment, power generation at a generation station 202 subject to reliability curtailment, power generation at a generation station 202 subject to power factor correction, low power generation at a generation station 202, start up conditions at a generation station 202, transient power generation conditions at a generation station 202, or testing conditions where there is an economic advantage to using behind-the-meter power generation at a generation station 202. These different monitored power system conditions can be weighted differently during processing and analysis.

[0111] In some examples, the information can include the cost for power from available sources (e.g., BTM power at the generation station 202 versus metered grid power) to enable comparisons to be made as to which power source costs less. In some instances, the information may include historic prices for power to enable the remote master control system 262 or another system to predict potential future prices in similar situations (e.g., the cost of power tends to trend upwards for grid power during warmer weather and peak-use hours). The information may also indicate the availability of power from the various sources (e.g., BTM power at the generation station 262, the energy storage system 218 at the generation station 262, and/or metered grid power).

[0112] In addition, the information may also include other data, including information associated with operations at components within the arrangement. For instance, the information may include data associated with performance of operations at the flexible datacenters 220 and the traditional datacenters 260, such as the number of computational tasks currently being performed, the types of tasks being performed (e.g., type of computational operation, time-sensitivity, etc.), the number, types, and capabilities of available computing systems, the amount of computational tasks awaiting performance, and the types of computing systems at one or more datacenters, among others. The information may also include data

specifying the conditions at one or more datacenters (e.g., whether or not the temperatures are in a desired range, the amount of power available within an energy storage system such as 218), the amount of computational tasks awaiting performance in the queue of one or more of the datacenters, and the identities of the entities associated with the computational operations at one or more of the datacenters. Entities associated with computational operations may be, for example, owners of the datacenters, customers who purchase computational time at the datacenters, or other entities.

[0113] The information used by the remote master control system 262 or another component may include data associated with the computational operations to be performed, such as deadlines, priorities (e.g., high vs. low priority tasks), cost to perform based on required computing systems, the optimal computing systems (e.g., CPU vs GPU vs ASIC; processing unit capabilities, speeds, or frequencies, or instructional sets executable by the processing units) for performing each requested computational task, and prices each entity (e.g., company) is willing to pay for computational operations to be performed or otherwise supported via computing systems at a traditional datacenter 260 or a flexible datacenter 220, among others. In addition, the information may also include other data (e.g., weather conditions at locations of datacenters or power sources, any emergencies associated with a datacenter or power source, or the current value of bids associated with an auction for computational tasks).

[0114] The information may be updated in-real time and used to make the different operational decisions within the arrangement of Figure 2. For instance, the information may help a component (e.g., the remote master control system 262 or a control system at a flexible datacenter 220) determine when to ramp up or ramp down power use at a flexible datacenter 220 or when to switch one or more computing systems at a flexible datacenter 220 into a low power mode or to operate at a different frequency, among other operational adjustments. The information can additionally or alternatively help a component within the arrangement of Figure 2 to determine when to transfer computational operations between computing systems or between datacenters based on various factors. In some instances, the information may also be used to determine when to temporarily stop performing a computational operation or when to perform a computational operation at multiple sites for redundancy or other reasons. The information may further be used to determine when to accept new computational operations from entities or when to temporarily suspend accepting new tasks to be performed due to lack of computing system availability.

[0115] The remote master control system 262 represents a computing system that is capable of obtaining, managing, and using the information described above to manage and oversee one or

more operations within the arrangement of Figure 2. As such, the remote master control system 262 may be one or more computing systems configured to process all, or a subset of, the information described above, such as power, environment, computational characterization, and economic factors to assist with the distribution and execution of computing operations among one or more datacenters. For instance, the remote master control system 262 may be configured to obtain and delegate computational operations among one or more datacenters based on a weighted analysis of a variety of factors, including one or more of the cost and availability of power, the types and availability of the computing systems at each datacenter, current and predicted weather conditions at the different locations of flexible datacenters (e.g., flexible datacenters 220) and generation stations (e.g., generation stations 202), levels of power storage available at one or more energy storage systems (e.g., energy storage system 218), and deadlines and other attributes associated with particular computational operations, among other possible factors. As such, the analysis of information performed by the remote master control system 262 may vary within examples. For instance, the remote master control system 262 may use real-time information to determine whether or not to route a computational operation to a particular flexible datacenter (e.g., a flexible datacenter 220) or to transition a computational operation between datacenters (e.g., from traditional datacenter 260 to a flexible datacenter 220).

[0116] As shown in Figure 2, the generation station 202 may be able to supply power to the grid and/or BTM loads such as flexible datacenters 220. With such a configuration, the generation station 202 may selectively provide power to the BTM loads and/or the grid based on economic and power availability considerations. For example, the generation station 202 may supply power to the grid when the price paid for the power exceeds a particular threshold (e.g., the power price offered by operators of the flexible datacenters 220). In some instances, the operator of a flexible datacenter and the operator of a generation station capable of supplying BTM power to the flexible datacenter may utilize a predefined arrangement (e.g., a contract) that specifies a duration and/or price range when the generation station may supply power to the flexible datacenter.

[0117] The remote master control system 262 may be capable of directing one or more flexible datacenters 220 to ramp-up or ramp-down to desired power consumption levels, and/or to control cooperative action of multiple flexible datacenters by determining how to power each individual flexible datacenter 220 in accordance with operational directives. Desired power consumption levels can depend on monitored conditions, such as power availability, prices for power from one or more sources, and timelines for completing computational operations.

[0118] The configuration of the remote master control system 262 can vary within examples as further discussed with respect to Figures 2, 3, and 7-9. The remote master control system 262 may operate as a single computing system or may involve a network of computing systems. Preferably, the remote master control system 262 is implemented across one or more servers in a fault-tolerant operating environment that ensures continuous uptime and connectivity by virtue of its distributed nature. Alternatively, although the remote master control system 262 is shown as a physically separate component arrangement for Figure 2, the remote master control system 262 may be combined with another component in other embodiments. To illustrate an example, the remote master control system 262 may operate as part of a flexible datacenter (e.g., a computing system or a datacenter control system of the flexible datacenter 220), including sharing components with a flexible datacenter, sharing power with a flexible datacenter, and/or being co-located with a flexible datacenter.

[0119] In addition, the remote master control system 262 may communicate with components within the arrangement of Figure 2 using various communication technologies, including wired and wireless communication technologies. For instance, the remote master control system 262 may use wired (not illustrated) or wireless communication to communicate with datacenter control systems or other computing systems at the flexible datacenters 220 and the traditional datacenters 260. The remote master control system 262 may also communicate with entities inside or outside the arrangement of Figure 2 and other components within the arrangement of Figure 2 via wired or wireless communication. For instance, the remote master control system 262 may use wireless communication to obtain computational operations from entities seeking support for the computational operations at one or more datacenters in exchange for payment. The remote master control system 262 may communicate directly with the entities or may obtain the computational operations from the traditional datacenters 260. For instance, an entity may submit jobs (e.g., computational operations) to one or more traditional datacenters 260. The remote master control system 262 may determine that transferring one or more of the computational operations to a flexible datacenter 220 may better support the transferred computational operations. For example, the remote master control system 262 may determine that the transfer may enable the computational operations to be completed quicker and/or at a lower cost. In some examples, the remote master control system 262 may communicate with the entity to obtain approval prior to transferring the one or more computational operations. In some instances, transferring may involve migrating computational operations being performed and/or moving computational operations prior to initiation of the computational operations.

[0120] The remote master control system 262 may also communicate with grid operators and/or an operator of generation station 202 to help determine power management strategies when distributing computational operations across the various datacenters. In addition, the remote master control system 262 may communicate with other sources, such as weather prediction systems, historical and current power price databases, and auction systems, etc.

[0121] In further examples, the remote master control system 262 or another computing system within the arrangement of Figure 2 may use wired or wireless communication to submit bids within an auction that involves a bidder (e.g., the highest bid) obtaining computational operations or other tasks to be performed. Particularly, the remote master control system 262 may use the information discussed above to develop bids to obtain computing operations for performance at available computing systems at flexible datacenters (e.g., flexible datacenters 220).

[0122] In the example arrangement shown in Figure 2, the flexible datacenters 220 represent example loads that can receive power behind-the-meter from the generation station 202. In such a configuration, the flexible datacenters 220 may obtain and utilize power behind-the-meter from the generation station 202 to perform various computational operations. Performance of a computational operation may involve one or more computing systems providing resources useful in the computational operation. For instance, the flexible datacenters 220 may include one or more computing systems configured to store information, perform calculations and/or parallel processes, perform simulations, mine cryptocurrencies, and execute applications, among other potential tasks. The computing systems can be specialized or generic and can be arranged at each flexible datacenter 220 in a variety of ways (e.g., straight configuration, zig-zag configuration) as further discussed with respect to Figures 6A, 6B. Furthermore, although the example arrangement illustrated in Figure 2 shows configurations where flexible datacenters 220 serve as BTM loads, other types of loads can be used as BTM loads within examples.

[0123] The arrangement of Figure 2 includes the traditional datacenters 260 coupled to metered grid power. The traditional datacenters 260 using metered grid power to provide computational resources to support computational operations. One or more enterprises may assign computational operations to the traditional datacenters 260 with expectations that the datacenters reliably provide resources without interruption (i.e., non-intermittently) to support the computational operations, such as processing abilities, networking, and/or volatile storage. Similarly, one or more enterprises may also request computational operations to be performed by the flexible datacenters 220. The flexible datacenters 220 differ from the traditional

datacenters 260 in that the flexible datacenters 220 are arranged and/or configured to be connected to BTM power, are expected to operate intermittently, and are expected to ramp load (and thus computational capability) up or down regularly in response to control directives. In some examples, the flexible datacenters 220 and the traditional datacenters 260 may have similar configurations and may only differ based on the source(s) of power relied upon to power internal computing systems. Preferably, however, the flexible datacenters 220 include particular fast load ramping abilities (e.g., quickly increase or decrease power usage) and are intended and designed to effectively operate during intermittent periods of time.

[0124] Figure 3 shows a block diagram of the remote master control system 300 according to one or more example embodiments. Remote master control system 262 may take the form of remote master control system 300, or may include less than all components in remote master control system 300, different components than in remote master control system 300, and/or more components than in remote master control system 300.

[0125] The remote master control system 300 may perform one or more operations described herein and may include a processor 302, a data storage unit 304, a communication interface 306, a user interface 308, an operations and environment analysis module 310, and a queue system 312. In other examples, the remote master control system 300 may include more or fewer components in other possible arrangements.

[0126] As shown in Figure 3, the various components of the remote master control system 300 can be connected via one or more connection mechanisms (e.g., a connection mechanism 314). In this disclosure, the term “connection mechanism” means a mechanism that facilitates communication between two or more devices, systems, components, or other entities. For instance, a connection mechanism can be a simple mechanism, such as a cable, PCB trace, or system bus, or a relatively complex mechanism, such as a packet-based communication network (e.g., LAN, WAN, and/or the Internet). In some instances, a connection mechanism can include a non-tangible medium (e.g., where the connection is wireless).

[0127] As part of the arrangement of Figure 2, the remote master control system 300 (corresponding to remote master control system 262) may perform a variety of operations, such as management and distribution of computational operations among datacenters, monitoring operational, economic, and environment conditions, and power management. For instance, the remote master control system 300 may obtain computational operations from one or more enterprises for performance at one or more datacenters. The remote master control system 300 may subsequently use information to distribute and assign the computational operations to one or more datacenters (e.g., the flexible datacenters 220) that have the resources (e.g., particular

types of computing systems and available power) available to complete the computational operations. In some examples, the remote master control system 300 may assign all incoming computational operation requests to the queue system 312 and subsequently assign the queued requests to computing systems based on an analysis of current market and power conditions.

[0128] Although the remote master control system 300 is shown as a single entity, a network of computing systems may perform the operations of the remote master control system 300 in some examples. For example, the remote master control system 300 may exist in the form of computing systems (e.g., datacenter control systems) distributed across multiple datacenters.

[0129] The remote master control system 300 may include one or more processors 302. As such, the processor 302 may represent one or more general-purpose processors (e.g., a microprocessor) and/or one or more special-purpose processors (e.g., a digital signal processor (DSP)). In some examples, the processor 302 may include a combination of processors within examples. The processor 302 may perform operations, including processing data received from the other components within the arrangement of Figure 2 and data obtained from external sources, including information such as weather forecasting systems, power market price systems, and other types of sources or databases.

[0130] The data storage unit 304 may include one or more volatile, non-volatile, removable, and/or non-removable storage components, such as magnetic, optical, or flash storage, and/or can be integrated in whole or in part with the processor 302. As such, the data storage unit 304 may take the form of a non-transitory computer-readable storage medium, having stored thereon program instructions (e.g., compiled or non-compiled program logic and/or machine code) that, when executed by the processor 302, cause the remote master control system 300 to perform one or more acts and/or functions, such as those described in this disclosure. Such program instructions can define and/or be part of a discrete software application. In some instances, the remote master control system 300 can execute program instructions in response to receiving an input, such as from the communication interface 306, the user interface 308, or the operations and environment analysis module 310. The data storage unit 304 may also store other information, such as those types described in this disclosure.

[0131] In some examples, the data storage unit 304 may serve as storage for information obtained from one or more external sources. For example, data storage unit 304 may store information obtained from one or more of the traditional datacenters 260, a generation station 202, a system associated with the grid, and flexible datacenters 220. As examples only, data storage 304 may include, in whole or in part, local storage, dedicated server-managed storage, network attached storage, and/or cloud-based storage, and/or combinations thereof.

[0132] The communication interface 306 can allow the remote master control system 300 to connect to and/or communicate with another component according to one or more protocols. For instance, the communication interface 306 may be used to obtain information related to current, future, and past prices for power, power availability, current and predicted weather conditions, and information regarding the different datacenters (e.g., current workloads at datacenters, types of computing systems available within datacenters, price to obtain power at each datacenter, levels of power storage available and accessible at each datacenter, etc.). In an example, the communication interface 306 can include a wired interface, such as an Ethernet interface or a high-definition serial-digital-interface (HD-SDI). In another example, the communication interface 406 can include a wireless interface, such as a cellular, satellite, WiMAX, or WI-FI interface. A connection can be a direct connection or an indirect connection, the latter being a connection that passes through and/or traverses one or more components, such as such as a router, switcher, or other network device. Likewise, a wireless transmission can be a direct transmission or an indirect transmission. The communication interface 306 may also utilize other types of wireless communication to enable communication with datacenters positioned at various locations.

[0133] The communication interface 306 may enable the remote master control system 300 to communicate with the components of the arrangement of Figure 2. In addition, the communication interface 306 may also be used to communicate with the various datacenters, power sources, and different enterprises submitting computational operations for the datacenters to support.

[0134] The user interface 308 can facilitate interaction between the remote master control system 300 and an administrator or user, if applicable. As such, the user interface 308 can include input components such as a keyboard, a keypad, a mouse, a touch-sensitive panel, a microphone, and/or a camera, and/or output components such as a display device (which, for example, can be combined with a touch-sensitive panel), a sound speaker, and/or a haptic feedback system. More generally, the user interface 308 can include hardware and/or software components that facilitate interaction between remote master control system 300 and the user of the system.

[0135] In some examples, the user interface 308 may enable the manual examination and/or manipulation of components within the arrangement of Figure 2. For instance, an administrator or user may use the user interface 308 to check the status of, or change, one or more computational operations, the performance or power consumption at one or more datacenters, the number of tasks remaining within the queue system 312, and other operations. As such,

the user interface 308 may provide remote connectivity to one or more systems within the arrangement of Figure 2.

[0136] The operations and environment analysis module 310 represents a component of the remote master control system 300 associated with obtaining and analyzing information to develop instructions/directives for components within the arrangement of Figure 2. The information analyzed by the operations and environment analysis module 310 can vary within examples and may include the information described above with respect predicting and/or directing the use of BTM power. For instance, the operations and environment analysis module 310 may obtain and access information related to the current power state of computing systems operating as part of the flexible datacenters 220 and other datacenters that the remote master control system 300 has access to. This information may be used to determine when to adjust power usage or mode of one or more computing systems. In addition, the remote master control system 300 may provide instructions a flexible datacenter 220 to cause a subset of the computing systems to transition into a low power mode to consume less power while still performing operations at a slower rate. The remote master control system 300 may also use power state information to cause a set of computing systems at a flexible datacenter 220 to operate at a higher power consumption mode. In addition, the remote master control system 300 may transition computing systems into sleep states or power on/off based on information analyzed by the operations and environment analysis module 310.

[0137] In some examples, the operations and environment analysis module 310 may use location, weather, activity levels at the flexible datacenters or the generation station, and power cost information to determine control strategies for one or more components in the arrangement of Figure 2. For instance, the remote master control system 300 may use location information for one or more datacenters to anticipate potential weather conditions that could impact access to power. In addition, the operations and environment analysis module 310 may assist the remote master control system 300 determine whether to transfer computational operations between datacenters based on various economic and power factors.

[0138] The queue system 312 represents a queue capable of organizing computational operations to be performed by one or more datacenters. Upon receiving a request to perform a computational operation, the remote master control system 300 may assign the computational operation to the queue until one or more computing systems are available to support the computational operation. The queue system 312 may be used for organizing and transferring computational tasks in real time.

[0139] The organizational design of the queue system 312 may vary within examples. In some examples, the queue system 312 may organize indications (e.g., tags, pointers) to sets of computational operations requested by various enterprises. The queue system 312 may operate as a First-In-First-Out (FIFO) data structure. In a FIFO data structure, the first element added to the queue will be the first one to be removed. As such, the queue system 312 may include one or more queues that operate using the FIFO data structure.

[0140] In some examples, one or more queues within the queue system 312 may use other designs of queues, including rules to rank or organize queues in a particular manner that can prioritize some sets of computational operations over others. The rules may include one or more of an estimated cost and/or revenue to perform each set of computational operations, an importance assigned to each set of computational operations, and deadlines for initiating or completing each set of computational operations, among others. Examples using a queue system are further described below with respect to Figure 9.

[0141] In some examples, the remote master control system 300 may be configured to monitor one or more auctions to obtain computational operations for datacenters to support. Particularly, the remote master control system 300 may use resource availability and power prices to develop and submit bids to an external or internal auction system for the right to support particular computational operations. As a result, the remote master control system 300 may identify computational operations that could be supported at one or more flexible datacenters 220 at low costs.

[0142] Figure 4 is a block diagram of a generation station 400, according to one or more example embodiments. Generation station 202 may take the form of generation station 400, or may include less than all components in generation station 400, different components than in generation station 400, and/or more components than in generation station 400. The generation station 400 includes a power generation equipment 401, a communication interface 408, a behind-the-meter interface 406, a grid interface 404, a user interface 410, a generation station control system 414, and power transformation equipment 402. power generation equipment 210 may take the form of power generation equipment 401, or may include less than all components in power generation equipment 401, different components than in power generation equipment 401, and/or more components than in power generation equipment 401. Generation station control system 216 may take the form of generation station control system 414, or may include less than all components in generation station control system 414, different components than in generation station control system 414, and/or more components than in generation station control system 414. Some or all of the components generation station 400

may be connected via a communication interface 516. These components are illustrated in Figure 4 to convey an example configuration for the generation station 400 (corresponding to generation station 202 shown in Figure 2). In other examples, the generation station 400 may include more or fewer components in other arrangements.

[0143] The generation station 400 can correspond to any type of grid-connected utility-scale power producer capable of supplying power to one or more loads. The size, amount of power generated, and other characteristics of the generation station 400 may differ within examples. For instance, the generation station 400 may be a power producer that provides power intermittently. The power generation may depend on monitored power conditions, such as weather at the location of the generation station 400 and other possible conditions. As such, the generation station 400 may be a temporary arrangement, or a permanent facility, configured to supply power. The generation station 400 may supply BTM power to one or more loads and supply metered power to the electrical grid. Particularly, the generation station 400 may supply power to the grid as shown in the arrangement of Figure 2.

[0144] The power generation equipment 401 represents the component or components configured to generate utility-scale power. As such, the power generation equipment 401 may depend on the type of facility that the generation station 400 corresponds to. For instance, the power generation equipment 401 may correspond to electric generators that transform kinetic energy into electricity. The power generation equipment 401 may use electromagnetic induction to generate power. In other examples, the power generation equipment 401 may utilize electrochemistry to transform chemical energy into power. The power generation equipment 401 may use the photovoltaic effect to transform light into electrical energy. In some examples, the power generation equipment 401 may use turbines to generate power. The turbines may be driven by, for example, wind, water, steam or burning gas. Other examples of power production are possible.

[0145] The communication interface 408 enables the generation station 400 to communicate with other components within the arrangement of Figure 2. As such, the communication interface 408 may operate similarly to the communication interface 306 of the remote master control system 300 and the communication interface 503 of the flexible datacenter 500.

[0146] The generation station control system 414 may be one or more computing systems configured to control various aspects of the generation station 400.

[0147] The BTM interface 406 is a module configured to enable the power generation equipment 401 to supply BTM power to one or more loads and may include multiple components. The arrangement of the BTM interface 406 may differ within examples based on

various factors, such as the number of flexible datacenters 220 (or 500) coupled to the generation station 400, the proximity of the flexible datacenters 220 (or 500), and the type of generation station 400, among others. In some examples, the BTM interface 406 may be configured to enable power delivery to one or more flexible datacenters positioned near the generation station 400. Alternatively, the BTM interface 406 may also be configured to enable power delivery to one or more flexible datacenters 220 (or 500) positioned remotely from the generation station 400.

[0148] The grid interface 404 is a module configured to enable the power generation equipment 401 to supply power to the grid and may include multiple components. As such, the grid interface 404 may couple to one or more transmission lines (e.g., transmission lines 404a shown in Figure 2) to enable delivery of power to the grid.

[0149] The user interface 410 represents an interface that enables administrators and/or other entities to communicate with the generation station 400. As such, the user interface 410 may have a configuration that resembles the configuration of the user interface 308 shown in Figure 3. An operator may utilize the user interface 410 to control or monitor operations at the generation station 400.

[0150] The power transformation equipment 402 represents equipment that can be utilized to enable power delivery from the power generation equipment 401 to the loads and to transmission lines linked to the grid. Example power transformation equipment 402 includes, but is not limited to, transformers, inverters, phase converters, and power conditioners.

[0151] Figure 5 shows a block diagram of a flexible datacenter 500, according to one or more example embodiments. Flexible datacenters 220 may take the form of flexible datacenter 500, or may include less than all components in flexible datacenter 500, different components than in flexible datacenter 500, and/or more components than in flexible datacenter 500. In the example embodiment shown in Figure 5, the flexible datacenter 500 includes a power input system 502, a communication interface 503, a datacenter control system 504, a power distribution system 506, a climate control system 508, one or more sets of computing systems 512, and a queue system 514. These components are shown connected by a communication bus 528. In other embodiments, the configuration of flexible datacenter 500 can differ, including more or fewer components. In addition, the components within flexible datacenter 500 may be combined or further divided into additional components within other embodiments.

[0152] The example configuration shown in Figure 5 represents one possible configuration for a flexible datacenter. As such, each flexible datacenter may have a different configuration when implemented based on a variety of factors that may influence its design, such as location

and temperature that the location, particular uses for the flexible datacenter, source of power supplying computing systems within the flexible datacenter, design influence from an entity (or entities) that implements the flexible datacenter, and space available for the flexible datacenter. Thus, the embodiment of flexible datacenter 220 shown in Figure 2 represents one possible configuration for a flexible datacenter out of many other possible configurations.

[0153] The flexible datacenter 500 may include a design that allows for temporary and/or rapid deployment, setup, and start time for supporting computational operations. For instance, the flexible datacenter 500 may be rapidly deployed at a location near a source of generation station power (e.g., near a wind farm or solar farm). Rapid deployment may involve positioning the flexible datacenter 500 at a target location and installing and/or configuring one or more racks of computing systems within. The racks may include wheels to enable swift movement of the computing systems. Although the flexible datacenter 500 could theoretically be placed anywhere, transmission losses may be minimized by locating it proximate to BTM power generation.

[0154] The physical construction and layout of the flexible datacenter 500 can vary. In some instances, the flexible datacenter 500 may utilize a metal container (e.g., a metal container 602 shown in Figure 6A). In general, the flexible datacenter 500 may utilize some form of secure weatherproof housing designed to protect interior components from wind, weather, and intrusion. The physical construction and layout of example flexible datacenters are further described with respect to Figures 6A-6B.

[0155] Within the flexible datacenter 500, various internal components enable the flexible datacenter 500 to utilize power to perform some form of operations. The power input system 502 is a module of the flexible datacenter 500 configured to receive external power and input the power to the different components via assistance from the power distribution system 506. As discussed with respect to Figure 2, the sources of external power feeding a flexible datacenter can vary in both quantity and type (e.g., the generation stations 202, 400, grid-power, energy storage systems). Power input system 502 includes a BTM power input sub-system 522, and may additionally include other power input sub-systems (e.g., a grid-power input sub-system 524 and/or an energy storage input sub-system 526). In some instances, the quantity of power input sub-systems may depend on the size of the flexible datacenter and the number and/or type of computing systems being powered.

[0156] In some embodiments, the power input system 502 may include some or all of flexible datacenter Power Equipment 220B. The power input system 502 may be designed to obtain power in different forms (e.g., single phase or three-phase behind-the-meter alternating current

(“AC”) voltage, and/or direct current (“DC”) voltage). As shown, the power input system 502 includes a BTM power input sub-system 522, a grid power input sub-system 524, and an energy input sub-system 526. These sub-systems are included to illustrate example power input sub-systems that the flexible datacenter 500 may utilize, but other examples are possible. In addition, in some instances, these sub-systems may be used simultaneously to supply power to components of the flexible datacenter 500. The sub-systems may also be used based on available power sources.

[0157] In some implementations, the BTM power input sub-system 522 may include one or more AC-to-AC step-down transformers used to step down supplied medium-voltage AC to low voltage AC (e.g., 120V to 600V nominal) used to power computing systems 512 and/or other components of flexible datacenter 500. The power input system 502 may also directly receive single-phase low voltage AC from a generation station as BTM power, from grid power, or from a stored energy system such as energy storage system 218. In some implementations, the power input system 502 may provide single-phase AC voltage to the datacenter control system 504 (and/or other components of flexible datacenter 500) independent of power supplied to computing systems 512 to enable the datacenter control system 504 to perform management operations for the flexible datacenter 500. For instance, the grid power input sub-system 524 may use grid power to supply power to the datacenter control system 504 to ensure that the datacenter control system 504 can perform control operations and communicate with the remote master control system 300 (or 262) during situations when BTM power is not available. As such, the datacenter control system 504 may utilize power received from the power input system 502 to remain powered to control the operation of flexible datacenter 500, even if the computational operations performed by the computing system 512 are powered intermittently. In some instances, the datacenter control system 504 may switch into a lower power mode to utilize less power while still maintaining the ability to perform some functions.

[0158] The power distribution system 506 may distribute incoming power to the various components of the flexible datacenter 500. For instance, the power distribution system 506 may direct power (e.g., single-phase or three-phase AC) to one or more components within flexible datacenter 500. In some embodiments, the power distribution system 506 may include some or all of flexible datacenter Power Equipment 220B.

[0159] In some examples, the power input system 502 may provide three phases of three-phase AC voltage to the power distribution system 506. The power distribution system 506 may controllably provide a single phase of AC voltage to each computing system or groups of

computing systems 512 disposed within the flexible datacenter 500. The datacenter control system 504 may controllably select which phase of three-phase nominal AC voltage that power distribution system 506 provides to each computing system 512 or groups of computing systems 512. This is one example manner in which the datacenter control system 504 may modulate power delivery (and load at the flexible datacenter 500) by ramping-up flexible datacenter 500 to fully operational status, ramping-down flexible datacenter 500 to offline status (where only datacenter control system 504 remains powered), reducing load by withdrawing power delivery from, or reducing power to, one or more of the computing systems 512 or groups of the computing systems 512, or modulating power factor correction for the generation station 300 (or 202) by controllably adjusting which phases of three-phase nominal AC voltage are used by one or more of the computing systems 512 or groups of the computing systems 512. The datacenter control system 504 may direct power to certain sets of computing systems based on computational operations waiting for computational resources within the queue system 514. In some embodiments, the flexible datacenter 500 may receive BTM DC power to power the computing systems 512.

[0160] One of ordinary skill in the art will recognize that a voltage level of three-phase AC voltage may vary based on an application or design and the type or kind of local power generation. As such, a type, kind, or configuration of the operational AC-to-AC step down transformer (not shown) may vary based on the application or design. In addition, the frequency and voltage level of three-phase AC voltage, single-phase AC voltage, and DC voltage may vary based on the application or design in accordance with one or more embodiments.

[0161] As discussed above, the datacenter control system 504 may perform operations described herein, such as dynamically modulating power delivery to one or more of the computing systems 512 disposed within flexible datacenter 500. For instance, the datacenter control system 504 may modulate power delivery to one or more of the computing systems 512 based on various factors, such as BTM power availability or an operational directive from a generation station 262 or 300 control system, a remote master control system 262 or 300, or a grid operator. In some examples, the datacenter control system 504 may provide computational operations to sets of computing systems 512 and modulate power delivery based on priorities assigned to the computational operations. For instance, an important computational operation (e.g., based on a deadline for execution and/or price paid by an entity) may be assigned to a particular computing system or set of computing systems 512 that has the capacity, computational abilities to support the computational operation. In addition, the datacenter

control system 504 may also prioritize power delivery to the computing system or set of computing systems 512.

[0162] In some example, the datacenter control system 504 may further provide directives to one or more computing systems to change operations in some manner. For instance, the datacenter control system 504 may cause one or more computing systems 512 to operate at a lower or higher frequency, change clock cycles, or operate in a different power consumption mode (e.g., a low power mode). These abilities may vary depending on types of computing systems 512 available at the flexible datacenter 500. As a result, the datacenter control system 504 may be configured to analyze the computing systems 512 available either on a periodic basis (e.g., during initial set up of the flexible datacenter 500) or in another manner (e.g., when a new computational operation is assigned to the flexible datacenter 500).

[0163] The datacenter control system 504 may also implement directives received from the remote master control system 262 or 300. For instance, the remote master control system 262 or 300 may direct the flexible datacenter 500 to switch into a low power mode. As a result, one or more of the computing systems 512 and other components may switch to the low power mode in response.

[0164] The datacenter control system 504 may utilize the communication interface 503 to communicate with the remote master control system 262 or 300, other datacenter control systems of other datacenters, and other entities. As such, the communication interface 503 may include components and operate similar to the communication interface 306 of the remote master control system 300 described with respect to Figure 4.

[0165] The flexible datacenter 500 may also include a climate control system 508 to maintain computing systems 512 within a desired operational temperature range. The climate control system 508 may include various components, such as one or more air intake components, an evaporative cooling system, one or more fans, an immersive cooling system, an air conditioning or refrigerant cooling system, and one or more air outtake components. One of ordinary skill in the art will recognize that any suitable heat extraction system configured to maintain the operation of computing systems 512 within the desired operational temperature range may be used.

[0166] The flexible datacenter 500 may further include an energy storage system 510. The energy storage system 510 may store energy for subsequent use by computing systems 512 and other components of flexible datacenter 500. For instance, the energy storage system 510 may include a battery system. The battery system may be configured to convert AC voltage to DC voltage and store power in one or more storage cells. In some instances, the battery system may

include a DC-to-AC inverter configured to convert DC voltage to AC voltage, and may further include an AC phase-converter, to provide AC voltage for use by flexible datacenter 500.

[0167] The energy storage system 510 may be configured to serve as a backup source of power for the flexible datacenter 500. For instance, the energy storage system 510 may receive and retain power from a BTM power source at a low cost (or no cost at all). This low-cost power can then be used by the flexible datacenter 500 at a subsequent point, such as when BTM power costs more. Similarly, the energy storage system 510 may also store energy from other sources (e.g., grid power). As such, the energy storage system 510 may be configured to use one or more of the sub-systems of the power input system 502.

[0168] In some examples, the energy storage system 510 may be external to the flexible datacenter 500. For instance, the energy storage system 510 may be an external source that multiple flexible datacenters utilize for back-up power.

[0169] The computing systems 512 represent various types of computing systems configured to perform computational operations. Performance of computational operations include a variety of tasks that one or more computing systems may perform, such as data storage, calculations, application processing, parallel processing, data manipulation, cryptocurrency mining, and maintenance of a distributed ledger, among others. As shown in Figure 5, the computing systems 512 may include one or more CPUs 516, one or more GPUs 518, and/or one or more Application-Specific Integrated Circuits (ASIC's) 520. Each type of computing system 512 may be configured to perform particular operations or types of operations.

[0170] Due to different performance features and abilities associated with the different types of computing systems, the datacenter control system 504 may determine, maintain, and/or relay this information about the types and/or abilities of the computing systems, quantity of each type, and availability to the remote master control system 262 or 300 on a routine basis (e.g., periodically or on-demand). This way, the remote master control system 262 or 300 may have current information about the abilities of the computing systems 512 when distributing computational operations for performance at one or more flexible datacenters. Particularly, the remote master control system 262 or 300 may assign computational operations based on various factors, such as the types of computing systems available and the type of computing systems required by each computing operation, the availability of the computing systems, whether computing systems can operate in a low power mode, and/or power consumption and/or costs associated with operating the computing systems, among others.

[0171] The quantity and arrangement of these computing systems 512 may vary within examples. In some examples, the configuration and quantity of computing systems 512 may

depend on various factors, such as the computational tasks that are performed by the flexible datacenter 500. In other examples, the computing systems 512 may include other types of computing systems as well, such as DSPs, SIMDs, neural processors, and/or quantum processors.

[0172] As indicated above, the computing systems 512 can perform various computational operations, including in different configurations. For instance, each computing system may perform a particular computational operation unrelated to the operations performed at other computing systems. Groups of the computing systems 512 may also be used to work together to perform computational operations.

[0173] In some examples, multiple computing systems may perform the same computational operation in a redundant configuration. This redundant configuration creates a back-up that prevents losing progress on the computational operation in situations of a computing failure or intermittent operation of one or more computing systems. In addition, the computing systems 512 may also perform computational operations using a check point system. The check point system may enable a first computing system to perform operations up to a certain point (e.g., a checkpoint) and switch to a second computing system to continue performing the operations from that certain point. The check point system may also enable the datacenter control system 504 to communicate statuses of computational operations to the remote master control system 262 or 300. This can further enable the remote master control system 262 300 to transfer computational operations between different flexible datacenters allowing computing systems at the different flexible datacenters to resume support of computational operations based on the check points.

[0174] The queue system 514 may operate similar to the queue system 312 of the remote master control system 300 shown in Figure 3. Particularly, the queue system 514 may help store and organize computational tasks assigned for performance at the flexible datacenter 500. In some examples, the queue system 514 may be part of a distributed queue system such that each flexible datacenter in a fleet of flexible datacenter includes a queue, and each queue system 514 may be able to communicate with other queue systems. In addition, the remote master control system 262 or 300 may be configured to assign computational tasks to the queues located at each flexible datacenter (e.g., the queue system 514 of the flexible datacenter 500). As such, communication between the remote master control system 262 or 300 and the datacenter control system 504 and/or the queue system 514 may allow organization of computational operations for the flexible datacenter 500 to support.

[0175] Figure 6A shows another structural arrangement for a flexible datacenter, according to one or more example embodiments. The particular structural arrangement shown in Figure 6A may be implemented at flexible datacenter 500. The illustration depicts the flexible datacenter 500 as a mobile container 702 equipped with the power input system 502, the power distribution system 506, the climate control system 508, the datacenter control system 504, and the computing systems 512 arranged on one or more racks 604. These components of flexible datacenter 500 may be arranged and organized according to an example structural region arrangement. As such, the example illustration represents one possible configuration for the flexible datacenter 500, but others are possible within examples.

[0176] As discussed above, the structural arrangement of the flexible datacenter 500 may depend on various factors, such as the ability to maintain temperature within the mobile container 602 within a desired temperature range. The desired temperature range may depend on the geographical location of the mobile container 602 and the type and quantity of the computing systems 512 operating within the flexible datacenter 500 as well as other possible factors. As such, the different design elements of the mobile container 602 including the inner contents and positioning of components may depend on factors that aim to maximize the use of space within mobile container 602, lower the amount of power required to cool the computing systems 512, and make setup of the flexible datacenter 500 efficient. For instance, a first flexible datacenter positioned in a cooler geographic region may include less cooling equipment than a second flexible datacenter positioned in a warmer geographic region.

[0177] As shown in Figure 6A, the mobile container 602 may be a storage trailer disposed on permanent or removable wheels and configured for rapid deployment. In other embodiments, the mobile container 602 may be a storage container (not shown) configured for placement on the ground and potentially stacked in a vertical or horizontal manner (not shown). In still other embodiments, the mobile container 602 may be an inflatable container, a floating container, or any other type or kind of container suitable for housing a mobile flexible datacenter. As such, the flexible datacenter 500 may be rapidly deployed on site near a source of unutilized behind-the-meter power generation. And in still other embodiments, the flexible datacenter 500 might not include a mobile container. For example, the flexible datacenter 500 may be situated within a building or another type of stationary environment.

[0178] Figure 6B shows the computing systems 512 in a straight-line configuration for installation within the flexible datacenter 500, according to one or more example embodiments. As indicated above, the flexible datacenter 500 may include a plurality of racks 604, each of which may include one or more computing systems 512 disposed therein. As discussed above,

the power input system 502 may provide three phases of AC voltage to the power distribution system 506. In some examples, the power distribution system 506 may controllably provide a single phase of AC voltage to each computing system 512 or group of computing systems 512 disposed within the flexible datacenter 500. As shown in Figure 6B, for purposes of illustration only, eighteen total racks 604 are divided into a first group of six racks 606, a second group of six racks 608, and a third group of six racks 610, where each rack contains eighteen computing systems 512. The power distribution system (506 of Figure 5) may, for example, provide a first phase of three-phase AC voltage to the first group of six racks 606, a second phase of three-phase AC voltage to the second group of six racks 608, and a third phase of three-phase AC voltage to the third group of six racks 610. In other embodiments, the quantity of racks and computing systems can vary.

[0179] Figure 7 shows a control distribution system 700 of the flexible datacenter 500 according to one or more example embodiments. The system 700 includes a grid operator 702, a generation station control system 216, a remote master control system 300, and a flexible datacenter 500. As such, the system 700 represents one example configuration for controlling operations of the flexible datacenter 500, but other configurations may include more or fewer components in other arrangements.

[0180] The datacenter control system 504 may independently, or cooperatively with one or more of the generation station control system 414, the remote master control system 300, and the grid operator 702, modulate power at the flexible datacenter 500. During operations, the power delivery to the flexible datacenter 500 may be dynamically adjusted based on conditions or operational directives. The conditions may correspond to economic conditions (e.g., cost for power, aspects of computational operations to be performed), power-related conditions (e.g., availability of the power, the sources offering power), demand response, and/or weather-related conditions, among others.

[0181] The generation station control system 414 may be one or more computing systems configured to control various aspects of a generation station (not independently illustrated, e.g., 216 or 400). As such, the generation station control system 414 may communicate with the remote master control system 300 over a networked connection 706 and with the datacenter control system 704 over a networked or other data connection 708.

[0182] As discussed with respect to Figures 2 and 3, the remote master control system 300 can be one or more computing systems located offsite, but connected via a network connection 710 to the datacenter control system 504. The remote master control system 300 may provide

supervisory controls or override control of the flexible datacenter 500 or a fleet of flexible datacenters (not shown).

[0183] The grid operator 702 may be one or more computing systems that are configured to control various aspects of the power grid (not independently illustrated) that receives power from the generation station. The grid operator 702 may communicate with the generation station control system 300 over a networked or other data connection 712.

[0184] The datacenter control system 504 may monitor BTM power conditions at the generation station and determine when a datacenter ramp-up condition is met. The BTM power availability may include one or more of excess local power generation, excess local power generation that the grid cannot accept, local power generation that is subject to economic curtailment, local power generation that is subject to reliability curtailment, local power generation that is subject to power factor correction, conditions where the cost for power is economically viable (e.g., low cost to obtain power), low priced power, situations where local power generation is prohibitively low, start up situations, transient situations, or testing situations where there is an economic advantage to using locally generated behind-the-meter power generation, specifically power available at little to no cost and with no associated transmission or distribution losses or costs. For example, a datacenter control system may analyze future workload and near term weather conditions at the flexible datacenter.

[0185] In some instances, the datacenter ramp-up condition may be met if there is sufficient behind-the-meter power availability and there is no operational directive from the generation station control system 414, the remote master control system 300, or the grid operator 702 to go offline or reduce power. As such, the datacenter control system 504 may enable 714 the power input system 502 to provide power to the power distribution system 506 to power the computing systems 512 or a subset thereof.

[0186] The datacenter control system 504 may optionally direct one or more computing systems 512 to perform predetermined computational operations (e.g., distributed computing processes). For example, if the one or more computing systems 512 are configured to perform distributed computing operations (e.g., hashing operations), the datacenter control system 504 may direct them to perform the distributed computing operations for a specific blockchain application, such as, for example, Bitcoin, Litecoin, or Ethereum. Alternatively, one or more computing systems 512 may be configured to perform high-throughput computing operations and/or high performance computing operations.

[0187] The remote master control system 300 may specify to the datacenter control system 504 what sufficient behind-the-meter power availability constitutes, or the datacenter control

system 504 may be programmed with a predetermined preference or criteria on which to make the determination independently. For example, in certain circumstances, sufficient behind-the-meter power availability may be less than that required to fully power the entire flexible datacenter 500. In such circumstances, the datacenter control system 504 may provide power to only a subset of computing systems, or operate the plurality of computing systems in a lower power mode, that is within the sufficient, but less than full, range of power that is available or to maximize profitability. In addition, the computing systems 512 may adjust operational frequency, such as performing more or less processes during a given duration.

[0188] While the flexible datacenter 500 is online and operational, a datacenter ramp-down condition may be met when there is insufficient or anticipated to be insufficient, behind-the-meter power availability or there is an operational directive from the generation station control system 414, the remote master control system 300, or the grid operator 702. The datacenter control system 504 may monitor and determine when there is insufficient, or anticipated to be insufficient, behind-the-meter power availability. As noted above, sufficiency may be specified by the remote master control system 300 or the datacenter control system 504 may be programmed with a predetermined preference or criteria on which to make the determination independently.

[0189] An operational directive may be based on current dispatchability, forward looking forecasts for when behind-the-meter power is, or is expected to be, available, economic considerations, reliability considerations, operational considerations, or the discretion of the generation station control system 414, the remote master control system 300, or the grid operator 702. For example, the generation station control system 414, the remote master control system 300, or the grid operator 702 may issue an operational directive to flexible datacenter 500 to go offline and power down. When the datacenter ramp-down condition is met, the datacenter control system 504 may disable power delivery to the plurality of computing systems (e.g., 512). The datacenter control system 504 may disable 714 the power input system 502 from providing power (e.g., three-phase nominal AC voltage) to the power distribution system 506 to power down the computing systems 512 while the datacenter control system 504 remains powered and is capable of returning service to operating mode at the flexible datacenter 500 when behind-the-meter power becomes available again.

[0190] While the flexible datacenter 500 is online and operational, changed conditions or an operational directive may cause the datacenter control system 504 to modulate power consumption by the flexible datacenter 500. The datacenter control system 504 may determine, or the generation station control system 414, the remote master control system 300, or the grid

operator 702 may communicate, that a change in local conditions may result in less power generation, availability, or economic feasibility, than would be necessary to fully power the flexible datacenter 500. In such situations, the datacenter control system 504 may take steps to reduce or stop power consumption by the flexible datacenter 500 (other than that required to maintain operation of datacenter control system 504).

[0191] Alternatively, the generation station control system 414, the remote master control system 300, or the grid operator 702, may issue an operational directive to reduce power consumption for any reason, the cause of which may be unknown. In response, the datacenter control system 504 may dynamically reduce or withdraw power delivery to one or more computing systems 512 to meet the dictate. The datacenter control system 504 may controllably provide three-phase nominal AC voltage to a smaller subset of computing systems (e.g., 512) to reduce power consumption. The datacenter control system 504 may dynamically reduce the power consumption of one or more computing systems by reducing their operating frequency or forcing them into a lower power mode through a network directive.

[0192] One of ordinary skill in the art will recognize that datacenter control system 504 may be configured to have a number of different configurations, such as a number or type or kind of the computing systems 512 that may be powered, and in what operating mode, that correspond to a number of different ranges of sufficient and available behind-the-meter power. As such, the datacenter control system 504 may modulate power delivery over a variety of ranges of sufficient and available unutilized behind-the-meter power availability.

[0193] Figure 8 shows a control distribution system 800 of a fleet of flexible datacenters according to one or more example embodiments. The control distribution system 800 of the flexible datacenter 500 shown and described with respect to Figure 7 may be extended to a fleet of flexible datacenters as illustrated in Figure 8. For example, a first generation station (not independently illustrated), such as a wind farm, may include a first plurality of flexible datacenters 802, which may be collocated or distributed across the generation station. A second generation station (not independently illustrated), such as another wind farm or a solar farm, may include a second plurality of flexible datacenters 804, which may be collocated or distributed across the generation station. One of ordinary skill in the art will recognize that the number of flexible datacenters deployed at a given station and the number of stations within the fleet may vary based on an application or design in accordance with one or more example embodiments.

[0194] The remote master control system 300 may provide directive to datacenter control systems of the fleet of flexible datacenters in a similar manner to that shown and described

with respect to Figure 7, with the added flexibility to make high level decisions with respect to fleet that may be counterintuitive to a given station. The remote master control system 300 may make decisions regarding the issuance of operational directives to a given generation station based on, for example, the status of each generation station where flexible datacenters are deployed, the workload distributed across fleet, and the expected computational demand required for one or both of the expected workload and predicted power availability. In addition, the remote master control system 300 may shift workloads from the first plurality of flexible datacenters 802 to the second plurality of flexible datacenters 804 for any reason, including, for example, a loss of BTM power availability at one generation station and the availability of BTM power at another generation station. As such, the remote master control system 300 may communicate with the generation station control systems 806A, 806B to obtain information that can be used to organize and distribute computational operations to the fleets of flexible datacenters 802, 804.

[0195] Figure 9 shows a queue distribution arrangement for a traditional datacenter 902 and a flexible datacenter 500, according to one or more example embodiments. The arrangement of Figure 9 includes a flexible datacenter 500, a traditional datacenter 902, a queue system 312, a set of communication links 916, 918, 920A, 920B, and the remote master control system 300. The arrangement of Figure 9 represents an example configuration scheme that can be used to distribute computing operations using a queue system 312 between the traditional datacenter 902 and one or more flexible datacenters. In other examples, the arrangement of Figure 9 may include more or fewer components in other potential configurations. For instance, the arrangement of Figure 9 may not include the queue system 312 or may include routes that bypass the queue system 312.

[0196] The arrangement of Figure 9 may enable computational operations requested to be performed by entities (e.g., companies). As such, the arrangement of Figure 9 may use the queue system 312 to organize incoming computational operations requests to enable efficient distribution to the flexible datacenter 500 and the critical traditional datacenter 902. Particularly, the arrangement of Figure 9 may use the queue system 312 to organize sets of computational operations thereby increasing the speed of distribution and performance of the different computational operations among datacenters. As a result, the use of the queue system 312 may reduce time to complete operations and reduce costs.

[0197] In some examples, one or more components, such as the datacenter control system 504, the remote master control system 300, the queue system 312, or the control system 936, may be configured to identify situations that may arise where using the flexible datacenter 500 can

reduce costs or increase productivity of the system, as compared to using the traditional datacenter 902 for computational operations. For example, a component within the arrangement of Figure 9 may identify when using behind-the-meter power to power the computing systems 512 within the flexible datacenter 500 is at a lower cost compared to using the computing systems 934 within the traditional datacenter 902 that are powered by grid power. Additionally, a component in the arrangement of Figure 9 may be configured to determine situations when offloading computational operations from the traditional datacenter 902 indirectly (i.e., via the queue system 312) or directly (i.e., bypassing the queue system 312) to the flexible datacenter 500 can increase the performance allotted to the computational operations requested by an entity (e.g., reduce the time required to complete time-sensitive computational operations).

[0198] In some examples, the datacenter control system 504 may monitor activity of the computing systems 512 within the flexible datacenter 500 and use the respective activity levels to determine when to obtain computational operations from the queue system 312. For instance, the datacenter control system 504 may analyze various factors prior to requesting or accessing a set of computational operations or an indication of the computational operations for the computing systems 512 to perform. The various factors may include power availability at the flexible datacenter 500 (e.g., either stored or from a BTM source), availability of the computing systems 512 (e.g., percentage of computing systems available), type of computational operations available, estimated cost to perform the computational operations at the flexible datacenter 500, cost for power, cost for power relative to cost for grid power, and instructions from other components within the system, among others. The datacenter control system 504 may analyze one or more of the factors when determining whether to obtain a new set of computational operations for the computing systems 512 to perform. In such a configuration, the datacenter control system 504 manages the activity of the flexible datacenter 500, including determining when to acquire new sets of computational operations when capacity among the computing systems 512 permit.

[0199] In other examples, a component (e.g., the remote master control system 300) within the system may assign or distribute one or more sets of computational operations organized by the queue system 312 to the flexible datacenter 500. For example, the remote master control system 300 may manage the queue system 312, including the distribution of computational operations organized by the queue system 312 to the flexible datacenter 500 and the traditional datacenter 902. The remote master control system 300 may utilize the information described

with respect to the Figures above to determine when to assign computational operations to the flexible datacenter 500.

[0200] The traditional datacenter 902 may include a power input system 930, a power distribution system 932, a datacenter control system 936, and a set of computing systems 934. The power input system 930 may be configured to receive power from a power grid and distribute the power to the computing systems 934 via the power distribution system 932. The datacenter control system 936 may monitor activity of the computing systems 934 and obtain computational operations to perform from the queue system 312. The datacenter control system 936 may analyze various factors prior to requesting or accessing a set of computational operations or an indication of the computational operations for the computing systems 934 to perform. A component (e.g., the remote master control system 300) within the arrangement of Figure 9 may assign or distribute one or more sets of computational operations organized by the queue system 312 to the traditional datacenter 902.

[0201] The communication link 916 represents one or more links that may serve to connect the flexible datacenter 500, the traditional datacenter 902, and other components within the system (e.g., the remote master control system 300, the queue system 312 – connections not shown). In particular, the communication link 916 may enable direct or indirect communication between the flexible datacenter 500 and the traditional datacenter 902. The type of communication link 916 may depend on the locations of the flexible datacenter 500 and the traditional datacenter 902. Within embodiments, different types of communication links can be used, including but not limited to WAN connectivity, cloud-based connectivity, and wired and wireless communication links.

[0202] The queue system 312 represents an abstract data type capable of organizing computational operation requests received from entities. As each request for computational operations are received, the queue system 312 may organize the request in some manner for subsequent distribution to a datacenter. Different types of queues can make up the queue system 312 within embodiments. The queue system 312 may be a centralized queue that organizes all requests for computational operations. As a centralized queue, all incoming requests for computational operations may be organized by the centralized queue.

[0203] In other examples, the queue system 312 may be distributed consisting of multiple queue sub-systems. In the distributed configuration, the queue system 312 may use multiple queue sub-systems to organize different sets of computational operations. Each queue sub-system may be used to organize computational operations based on various factors, such as according to deadlines for completing each set of computational operations, locations of

enterprises submitting the computational operations, economic value associated with the completion of computational operations, and quantity of computing resources required for performing each set of computational operations. For instance, a first queue sub-system may organize sets of non-intensive computational operations and a second queue sub-system may organize sets of intensive computational operations. In some examples, the queue system 312 may include queue sub-systems located at each datacenter. This way, each datacenter (e.g., via a datacenter control system) may organize computational operations obtained at the datacenter until computing systems are able to start executing the computational operations. In some examples, the queue system 312 may move computational operations between different computing systems or different datacenters in real-time.

[0204] Within the arrangement of Figure 9, the queue system 312 is shown connected to the remote master control system 300 via the communication link 918. In addition, the queue system 312 is also shown connected to the flexible datacenter via the communication 920A and to the traditional datacenter 902 via the communication link 920B. The communication links 918, 920A, 920B may be similar to the communication link 916 and can be various types of communication links within examples.

[0205] The queue system 312 may include a computing system configured to organize and maintain queues within the queue system 312. In another example, one or more other components of the system may maintain and support queues within the queue system 312. For instance, the remote master control system 300 may maintain and support the queue system 312. In other examples, multiple components may maintain and support the queue system 312 in a distributed manner, such as a blockchain configuration.

[0206] In some embodiments, the remote master control system 300 may serve as an intermediary that facilitates all communication between flexible datacenter 500 and the traditional datacenter 902. Particularly, the traditional datacenter 902 or the flexible datacenter 500 might need to transmit communications to the remote master control system 300 in order to communicate with the other datacenter. As also shown, the remote master control system 300 may connect to the queue system 312 via the communication link 918. Computational operations may be distributed between the queue system 312 and the remote master control system 300 via the communication link 918. The computational operations may be transferred in real-time and mid-performance from one datacenter to another (e.g., from the traditional datacenter 902 to the flexible datacenter 500). In addition, the remote master control system 300 may manage the queue system 312, including providing resources to support queues within the queue system 312.

[0207] As a result, the remote master control system 300 may offload some or all of the computational operations assigned to the traditional datacenter 902 to the flexible datacenter 500. This way, the flexible datacenter 500 can reduce overall computational costs by using the behind-the-meter power to provide computational resources to assist traditional datacenter 902. The remote master control system 300 may use the queue system 312 to temporarily store and organize the offloaded computational operations until a flexible datacenter (e.g., the flexible datacenter 500) is available to perform them. The flexible datacenter 500 consumes behind-the-meter power without transmission or distribution costs, which lowers the costs associated with performing computational operations originally assigned to the traditional datacenter 902. The remote master control system 300 may further communicate with the flexible datacenter 500 via communication link 922 and the traditional datacenter 902 via the communication link 924.

[0208] Figure 10A shows method 1000 of dynamic power consumption at a flexible datacenter using behind-the-meter power according to one or more example embodiments. Other example methods may be used to manipulate the power delivery to one or more flexible datacenters.

[0209] In step 1010, the datacenter control system, the remote master control system, or another computing system may monitor behind-the-meter power availability. In some embodiments, monitoring may include receiving information or an operational directive from the generation station control system or the grid operator corresponding to behind-the-meter power availability.

[0210] In step 1020, the datacenter control system or the remote master control system 300 may determine when a datacenter ramp-up condition is met. In some embodiments, the datacenter ramp-up condition may be met when there is sufficient behind-the-meter power availability and there is no operational directive from the generation station to go offline or reduce power.

[0211] In step 1030, the datacenter control system may enable behind-the-meter power delivery to one or more computing systems. In some instances, the remote master control system may directly enable BTM power delivery to computing systems within the flexible system without instructing the datacenter control system.

[0212] In step 1040, once ramped-up, the datacenter control system or the remote master control system may direct one or more computing systems to perform predetermined computational operations. In some embodiments, the predetermined computational operations may include the execution of one or more distributed computing processes, parallel processes, and/or hashing functions, among other types of processes.

[0213] While operational, the datacenter control system, the remote master control system, or another computing system may receive an operational directive to modulate power consumption. In some embodiments, the operational directive may be a directive to reduce power consumption. In such embodiments, the datacenter control system or the remote master control system may dynamically reduce power delivery to one or more computing systems or dynamically reduce power consumption of one or more computing systems. In other embodiments, the operational directive may be a directive to provide a power factor correction factor. In such embodiments, the datacenter control system or the remote master control system may dynamically adjust power delivery to one or more computing systems to achieve a desired power factor correction factor. In still other embodiments, the operational directive may be a directive to go offline or power down. In such embodiments, the datacenter control system may disable power delivery to one or more computing systems.

[0214] Figure 10B shows method 1050 of dynamic power delivery to a flexible datacenter using behind-the-meter power according to one or more embodiments. In step 1060, the datacenter control system or the remote master control system may monitor behind-the-meter power availability. In certain embodiments, monitoring may include receiving information or an operational directive from the generation station control system or the grid operator corresponding to behind-the-meter power availability.

[0215] In step 1070, the datacenter control system or the remote master control system may determine when a datacenter ramp-down condition is met. In certain embodiments, the datacenter ramp-down condition may be met when there is insufficient behind-the-meter power availability or anticipated to be insufficient behind-the-meter power availability or there is an operational directive from the generation station to go offline or reduce power.

[0216] In step 1080, the datacenter control system may disable behind-the-meter power delivery to one or more computing systems. In step 1090, once ramped-down, the datacenter control system remains powered and in communication with the remote master control system so that it may dynamically power the flexible datacenter when conditions change.

[0217] One of ordinary skill in the art will recognize that a datacenter control system may dynamically modulate power delivery to one or more computing systems of a flexible datacenter based on behind-the-meter power availability or an operational directive. The flexible datacenter may transition between a fully powered down state (while the datacenter control system remains powered), a fully powered up state, and various intermediate states in between. In addition, flexible datacenter may have a blackout state, where all power consumption, including that of the datacenter control system is halted. However, once the

flexible datacenter enters the blackout state, it will have to be manually rebooted to restore power to datacenter control system. Generation station conditions or operational directives may cause flexible datacenter to ramp-up, reduce power consumption, change power factor, or ramp-down.

[0218] Figure 11 illustrates a system for modulating operating attributes of computing systems at a flexible datacenter according to one or more embodiments. The system 1100 is shown with the remote master control system 262 in wireless communication 1130 with multiple datacenters (i.e., a set of flexible datacenters 1102, 1104 and a traditional datacenter 1106). This arrangement is shown for illustration purposes, but can differ within other examples. In other embodiments, remote master control system 262 could be implemented as remote master control system 300 and/or flexible datacenters 220 or 500 could be implemented as flexible datacenters 1102 and/or 1104. For example, the remote master control system 262 may be configured to manage operations at a fleet of datacenters that are positioned at various locations and coupled to different power generation sources.

[0219] Within the system 1100, the remote master control system 262 may be configured to help manage operations at the flexible datacenters 1102, 1104, and the traditional datacenter 1106. Each datacenter may be electrically coupled to one or more power generation sources. For example, the flexible datacenter 1102 may be electrically coupled to a first power generation source from which flexible datacenter 1102 receives BTM power and the flexible datacenter 1104 may be electrically coupled to a second power generation source from which flexible datacenter 1104 receives BTM power. In addition, the traditional datacenter 1106 may be electrically coupled to the power grid from which it receives metered grid power. Other arrangements are possible.

[0220] Each datacenter may be configured with one or more computing systems to perform computational operations. As shown in Figure 11, the flexible datacenter 1102 includes computing systems 1108 arranged into a first set 1114A, a second set 1114B, and a third set 1114C. Similarly, the flexible datacenter 1104 includes computing systems 1110 arranged into a first set 1116A, a second set 1116B, and a third set 1118B and the traditional datacenter 1106 includes computing systems 1112 arranged into a first set 1118A, a second set 1118B, and a third set 1118C.

[0221] The remote master control system 262 or another computing system may be configured to modulate operating attributes of computing systems at one or more of the datacenters based on one or more monitored conditions. Particularly, the remote master control system 262 may monitor and analyze a set of conditions to determine when to adjust operating attributes at one

or more computing systems. An operating attribute can represent computational or performance aspects of one or more computing systems. Example operating attributes may include, but are not limited to, the processing frequency, performance speed, or hash rate performed at a computing system (or set of computing systems), the rate of power consumption at one or more computing systems (e.g., a high-power mode, a standard-power mode, and a low-power mode), parameters related to performing or supporting computational operations, and the current state of one or more computing systems (e.g.: powered on or off; actively computing or not computing).

[0222] Different types of computing systems may have operating attributes that vary. For instance, CPUs, GPUs, and ASICs may each have different operating attributes. In addition, the setup and other external factors may influence the operating attributes of a computing system. For instance, a computing system's source of power, networking with other computing systems, and software can impact the operating attributes of the computing system.

[0223] The operating attributes of a computing system may be modulated in various ways. The modulation of a computing system's operating attributes can depend on the type of computing system and other factors. As such, modulation of a computing system's operating attributes may involve adjusting a processing speed or frequency of the computing system. For instance, a computing system may have its processing speed turned down through an underclocking technique or power consumption adjustment. Similarly, the computing system may also have its processing speed turned up via an overclocking technique or power consumption adjustment.

[0224] In some examples, modulation of a computing system's operating attributes may adjust the power consumption at the computing system. For instance, modulation may involve powering a computing system off or on depending on the current state of the computing system. In addition, some computing systems may have other power modes, such as a low-power mode, a mid-power mode, and a high-power mode. These modes may vary in at least the amount of power consumed when performing computational operations.

[0225] As indicated above, the remote master control system 262 or another computing system may be configured to modulate operating attributes at one or more computing systems. In some instances, modulating one or more operating attributes may involve a direct modulation via instructions transmitted from the remote master control system 262 (or another computing system) to the computing systems at the datacenter. To illustrate an example, the remote master control system 262 may modulate one or more computing systems 1108 at the flexible datacenter 1102, one or more computing systems 1110 at the flexible datacenter 1104, and one

or more computing systems 1112 at the traditional datacenter 1106 based on various monitored conditions. In other instances, modulating operating attributes may involve an indirect modulation via instructions transmitted from the remote master control system 262 (or another computing system) to a datacenter control system of a datacenter. For example, the remote master control system 262 may transmit instructions to a datacenter control system (not shown) at the flexible datacenter 1104 to modulate operating attributes of one or more of the first set 1116A, the second set 1116B, and the third set 1116C of the computing systems 1110.

[0226] As shown above, modulation of the operating abilities can involve modulating a single computing system or sets of computing systems. A set of computing systems may include one, a few, dozens, or even more computing systems. For instance, the remote master control system 262 may modulate the operating attributes of only the first set 1114A of the computing systems 1108 at the flexible datacenter 1102. In another example, the remote master control system 262 may be configured to modulate the operating attributes of the first set 1114A of the computing systems 1108 in a first way (e.g., switch to a low power mode) and the operating attributes of the second set 1114B in a second, different way (e.g., power off).

[0227] In addition, modulation of the operating abilities of a computing system may be performed in real-time or scheduled in advanced. When the modulation occurs in real-time, the computing system may adjust operating abilities immediately in response to receiving modulation instructions. When the modulation of the operating abilities is scheduled, the modulation change may occur when the scheduled time occurs. For example, the remote master control system 262 may provide instructions to a datacenter control system at the flexible datacenter 1104 to modulate operating abilities of the first set 1116A and the second set 1116B at a scheduled time. In turn, the datacenter control system may cause the modulation of the operating abilities of the first set 1116A and the second set 1116B at the scheduled time.

[0228] In addition, the arrangement, type of computing systems, and operating attributes associated with the computing systems within each set can vary. For instance, the first set 1114A of the flexible datacenter 102 may include only CPUs, the second set 1114B may only include ASICs, and the third set 1114C may include a combination of GPUs and CPUs. In another example, the sets may be arranged based on other operating attributes. For instance, the first set 1116A of the flexible datacenter 1104 may only include computing systems capable of switching and operating in a low power mode while the other sets do not include computing systems capable of operating in the low power mode. As such, computing systems may have different operating attributes. Some computing systems may only operate in a single mode and lack the ability to adjust power consumption or processing frequency. As such, as further

discussed below, the remote master control system 262 may obtain computational systems parameters from datacenters to gain an understanding of the different operating attributes of the computing systems positioned at the flexible datacenters 1102, 1104 and the traditional datacenter 1106.

[0229] Some situations may arise that can benefit from the remote master control system 262 modulating operating attributes at the various datacenters. The situations may vary within examples and can involve modulating operating attributes at one or more sets of computing systems based on detected changes in various conditions. These various conditions are also referred to herein as monitored conditions, may include but are not limited to, power availability at each datacenter (i.e., power availability 1120) and the price of the different available power options (i.e., the power prices 1122), parameters related to the computing systems (e.g., abilities, quantity available, estimated time until availability of occupied computing systems) at the various datacenters (i.e., the computing systems parameters 1124), and parameters related to the computational operations requiring performance and currently being performed (e.g., quantity of operations, estimated time to complete, cost to perform each computational operation, deadlines and priorities associated with each computational operation) (i.e., the computational operations parameters 1128). Other monitored conditions may include prices of cryptocurrencies 1126, any power disruption due to emergencies or other events, and changes in temperatures or weather conditions at power generators or datacenters (i.e., the weather conditions 1129). As such, the operations and environment analysis module (or another component) of the remote master control system 262 may be configured to monitor one or more conditions described above.

[0230] The various monitored conditions described above as well as other potential conditions may change dynamically and with great frequency. Thus, to enable efficient distribution and performance of the computational operations at the datacenters, the remote master control system 262 may be configured to monitor changes in the various conditions to assist with the efficient management and operations of the computing systems at each datacenter. For instance, the remote master control system 262 may engage in wired or wireless communication 1130 with datacenter control systems (e.g., datacenter control system 504) at each datacenter as well as other sources to monitor for changes in the conditions. In turn, the remote master control system 262 may analyze the different conditions in real-time to modulate operating attributes of one or more computing systems at one or more of the datacenters. By using the monitored conditions, the remote master control system 262 may increase revenue, decrease costs, and/or increase performance of computational operations via various

modifications, such as transferring computational operations between datacenters or sets of computing systems within a datacenter and adjusting performance at one or more sets of computing systems (e.g., switching to a low power mode).

[0231] The remote master control system 262 may be configured to modulate operating attributes at one or more computing systems in various ways. The changes initiated at a datacenter may depend on the analysis of the monitored conditions. In addition, the remote master control system 262 may initiate multiple changes at a given time, such as transferring computational operations between datacenters and powering off some computing systems while switching the processing frequency on others.

[0232] In one respect, the remote master control system 262 or another datacenter control system may cause one or more computing systems to switch into a low-power mode. In another respect, the remote master control system 262 or another datacenter control system may cause one or more computing systems to operate in another mode (e.g., high-power consumption mode).

[0233] In some situations, modulating operating attributes may involve adjusting the processing speed, frequency, hash rate or another aspect of performance at one or more computing systems. The remote master control system may also modulate operating attributes at one or more sets of computing systems by switching one or more computing systems (or sets of computing systems) into a low-power mode that reduces power consumption at the computing systems or powering off one or more computing systems (or sets of computing systems).

[0234] When analyzing the monitored conditions, the remote master control system may perform a weighted analysis to determine instructions for modulating operating attributes of one or more computing systems at one or more datacenters. The weighted analysis may involve assigning weights to each monitored condition such that some of the conditions have more influence on the outcome. In some examples, the weights may be predefined by an administrator or another entity. When a weight is predefined, the remote master control system may use that weight during analysis. Alternatively, some weights may be developed using analysis of the monitored conditions and performance at the datacenters over time.

[0235] Figure 12 illustrates a method for modulating operating attributes of one or more computing systems according to one or more embodiments. The method serves as an example and may include other steps within other embodiments. A datacenter control system (e.g., the remote master control system 262) may be configured to perform one or more steps of the method 1200.

[0236] At step 1202, the method 1200 involves monitoring a set of conditions. For instance, a datacenter control system at a flexible datacenter, the remote master control system, or another computing system may be configured to monitor the set of conditions alone or in combination within one or more other computing systems. For instance, the remote master control system may monitor the set of conditions with assistance from datacenter control systems (e.g., datacenter control system 504) positioned at each datacenter in communication with the remote master control system. Thus, the set of monitored conditions may include a variety of information obtained from one or more external sources, such as one or more datacenters, databases, power generation stations, or types of sources.

[0237] In some embodiments, the set of monitored conditions may include an indication of BTM power availability at one or more datacenters. Particularly, computing systems at a flexible datacenter (e.g., the flexible datacenter 500 of Figure 5) may be configured to receive and use BTM power from a power generation source. As such, the remote master control system (or another computing system) may monitor the availability of BTM power at the flexible datacenter. By extension, the remote master control system may monitor BTM power availability at a fleet of flexible datacenters. In addition, the monitored conditions may also include monitoring other types of power availability at datacenters, such as grid power and power from energy storage.

[0238] Monitoring the set of conditions may also include monitoring the prices associated with using different types of power available at each datacenter. More specifically, the remote master control system (or another computing system) may monitor the prices of grid power and BTM power at datacenters available to perform computational operations. By monitoring prices of the different power sources, the remote master control system may factor cost for supporting computational resources at different locations. The cost analysis may also influence the way the computational operations are performed at one or more datacenters.

[0239] In addition, monitored conditions may also include factoring other prices and costs. For example, the current prices of one or more cryptocurrencies may be monitored as a factor when distributing computational operations. Particularly, when revenue associated with a price of a cryptocurrency exceeds the cost associated with powering computing systems to mine the cryptocurrency, the remote master control system (or another computing system) may determine to use the computing systems to mine the cryptocurrency unless the computing systems can be better utilized supporting other computational operations with higher revenue streams.

[0240] Monitored conditions may also include monitoring parameters related to computing systems at one or more datacenters. Particularly, the remote master control system may monitor parameters associated with computing systems, such as a quantity of computing systems available at each datacenter, an estimated time of availability for currently occupied computing systems at each datacenter, and the types and particular abilities of computing systems at the datacenters. As discussed above, different types of computing systems may have various operating abilities that can be adjusted, such as changes in the processing speed and power consumption. As such, the remote master control system may use these parameters when distributing and managing performance of computational operations at the datacenters.

[0241] In some embodiments, the monitored conditions may also include an indication of parameters associated with computational operations. These computational operations may be waiting for performance in a queue or may be currently undergoing performance at one or more computing systems. The parameters may assist the remote master control system (or another computing system) to manage distribution and performance of the computational operations. Example parameters include, but are not limited to, type of computational operation, estimated time required to perform the computational operation, deadline associated with completing the computational operation (if a projected deadline is provided), quantity of processing required to perform the computational operation, entity and payment associated with the computational operation, and an indication of a hash rate or a difficulty level associated with performance of the computational operation.

[0242] At step 1204, the method 1200 involves performing a weighted analysis using the set of monitored conditions. The remote master control system or another computing system may be configured to analyze the various monitored conditions. Particularly, analyzing may involve using a weighted analysis that assigns greater weights to some conditions relative to the weights assigned to other conditions. For example, the weighted analysis may involve assigning a greater weight to BTM power availability such that assigning computational operations to computing systems at a flexible datacenter is prioritized over traditional datacenters when BTM power is available at the flexible datacenter. As such, the weight assigned to conditions can differ within various situations and may be refined to increase productivity across available datacenters while minimizing the costs associated with powering the computing systems required to drive the productivity.

[0243] In some embodiments, performing the weighted analysis may involve a comparison between the price of BTM power available a datacenter (e.g., a flexible datacenter) and the price of grid power at the same datacenter. The weighted analysis may also involve comparing

the price of power at multiple datacenters to determine how to route computational operations and adjust operating abilities at computing systems to increase productivity while decreasing costs. To illustrate an example, the remote master control system may compare the price of power at a traditional datacenter relative to the price of power at a flexible datacenter to determine when to transfer computational operations between the traditional datacenter and the flexible datacenter to lower power costs. In additional embodiments, the weighted analysis may involve factoring performance capabilities, prices, power availability, and other monitored conditions to manage operations among one or more datacenters.

[0244] At step 1206, the method 1200 involves modulating operating attributes of one or more computing systems of the plurality of computing systems based on the weighted analysis. As discussed above, the operating attributes of a computing system (or a set of computing systems) may be modulated in response to changes of monitored conditions. Some example changes to the operating attributes of a computing system (or a set of computing systems) includes adjusting processing speed, power consumption, or performance of the computing system. In addition, other changes may include transferring computational operations between computing systems, delaying performance of a computational operation, and forming new networks among computing systems, etc. As such, the remote master control system (or another computing system) may modulate operating abilities of one or more computing systems (or sets of computing systems) based on the analysis of the monitored conditions described above. In turn, the remote master control system may manage distribution and performance of various computational operations across one or more datacenters in a manner that increases efficiency while reducing associated costs.

[0245] In some embodiments, the remote master control system (or another computing system) may transfer computational operations between computing systems. For instance, the weighted analysis may indicate that performance of a computational operation is more cost efficient at a flexible datacenter rather than at a traditional datacenter. Particularly, the weighted analysis may show that BTM power at the flexible datacenter costs less than grid power at the datacenter and that the flexible datacenter includes computing systems capable and available to perform the computational operation. As a result, the remote master control system may transfer the computational operation from the traditional datacenter to the flexible datacenter. While the flexible datacenter performs the computational operation, the remote master control system may continue to monitor the set of conditions and may subsequently transfer the computational operation again if the analysis of changes in the set of conditions warrants another transfer.

[0246] In some embodiments, the analysis of the set of monitored conditions may indicate that modulation of operating abilities at a set of computing systems may effectively take advantage of one or more changes in the set of conditions. For example, the analysis may indicate that causing the set of computing systems to switch into a high power consumption mode that enables the computing systems to increase performance frequency (e.g., by overclocking the processors) may enable the completion of the computational operation more quickly while taking advantage of a decrease in the cost of power at the computing systems.

[0247] In further examples, the analysis may indicate that a combination of transferring a computational operation from a first datacenter to a second datacenter and further modifying operating attributes at the datacenters may increase productivity while increasing revenue and/or decreasing costs. For instance, the remote master control system may cause the transfer of the computational operation to enable high performance computing systems at the second datacenter to use low cost power to complete the computational operation. In such an example, the first datacenter may not have access to low-cost power or the computing systems at the first datacenter may lack the ability to execute at a high-performance level. As such, the remote master control system may initiate the transfer to speed up completion of the computational operation while using power at a decreased cost. In addition, the remote master control system may also cause the computing systems that transferred the computational operation to switch into a low power mode (e.g., a standby mode) or power off to stop power consumption until another computational operation may be performed using these computing systems.

[0248] In some embodiments, weather conditions or emergencies may also influence the outcome of the analysis of the monitored conditions. Particularly, an emergency at a generation station may cause the remote master control system to transfer computational operations away from a particular datacenter configured to use power from that generation station. Similarly, some weather conditions, such as low temperatures or storms, may increase or decrease the ability of a datacenter to perform computational operations at decreased or increased costs. As such, the analysis may factor these conditions such that management of datacenters and computing systems at the datacenters take advantage in changes in the monitored conditions.

[0249] Advantages of one or more embodiments of the present invention may include one or more of the following:

[0250] One or more embodiments of the present invention provides a green solution to two prominent problems: the exponential increase in power required for growing blockchain operations and the unutilized and typically wasted energy generated from renewable energy sources.

[0251] One or more embodiments of the present invention allows for the rapid deployment of mobile datacenters to local stations. The mobile datacenters may be deployed on site, near the source of power generation, and receive low cost or unutilized power behind-the-meter when it is available.

[0252] One or more embodiments of the present invention provide the use of a queue system to organize computational operations and enable efficient distribution of the computational operations across multiple datacenters.

[0253] One or more embodiments of the present invention enable datacenters to access and obtain computational operations organized by a queue system.

[0254] One or more embodiments of the present invention allows for the power delivery to the datacenter to be modulated based on conditions or an operational directive received from the local station or the grid operator.

[0255] One or more embodiments of the present invention may dynamically adjust power consumption by ramping-up, ramping-down, or adjusting the power consumption of one or more computing systems within the flexible datacenter.

[0256] One or more embodiments of the present invention may be powered by behind-the-meter power that is free from transmission and distribution costs. As such, the flexible datacenter may perform computational operations, such as distributed computing processes, with little to no energy cost.

[0257] One or more embodiments of the present invention provides a number of benefits to the hosting local station. The local station may use the flexible datacenter to adjust a load, provide a power factor correction, to offload power, or operate in a manner that invokes a production tax credit and/or generates incremental revenue.

[0258] One or more embodiments of the present invention allows for continued shunting of behind-the-meter power into a storage solution when a flexible datacenter cannot fully utilize excess generated behind-the-meter power.

[0259] One or more embodiments of the present invention allows for continued use of stored behind-the-meter power when a flexible datacenter can be operational but there is not an excess of generated behind-the-meter power.

[0260] One or more embodiments of the present invention allows for management and distribution of computational operations at computing systems across a fleet of datacenters such that the performance of the computational operations take advantages of increased efficiency and decreased costs.

[0261] It will also be recognized by the skilled worker that, in addition to improved efficiencies in controlling power delivery from intermittent generation sources, such as wind farms and solar panel arrays, to regulated power grids, the invention provides more economically efficient control and stability of such power grids in the implementation of the technical features as set forth herein.

[0262] While the present invention has been described with respect to the above-noted embodiments, those skilled in the art, having the benefit of this disclosure, will recognize that other embodiments may be devised that are within the scope of the invention as disclosed herein. Accordingly, the scope of the invention should be limited only by the appended claims.

CLAIMS

What is claimed is:

1. A system comprising:
a plurality of computing systems, wherein the plurality of computing systems receive behind-the-meter (“BTM”) power from a BTM power source; and
a datacenter control system configured to:
monitor a set of conditions, wherein the set of monitored conditions includes BTM power availability at the plurality of computing systems;
perform a weighted analysis using the set of monitored conditions; and
based on the weighted analysis, modulate operating attributes of one or more computing systems of the plurality of computing systems.
2. The system of claim 1, wherein the BTM power source comprises a generation station configured to generate utility-scale electrical power for supply to an electrical grid.
3. The system of claim 1, wherein the plurality of computing system is electrically connected to a generation station at a location behind a Point of Interconnection between the generation station and an electrical grid.
4. The system of claim 1, wherein the BTM power received by the plurality of computing systems is electrical power produced by a generation station and transmitted to the plurality of computing systems behind the generation station’s Point of Interconnection with an electrical grid.
5. The system of claim 1, wherein the plurality of computing systems is electrically connected to a generation station that is subject to metering by one or more utility-scale generation-side meter, wherein the BTM power received by the plurality of computing systems has not passed through the one or more utility-scale generation-side meters.
6. The system of claim 1, wherein the power received by the plurality of computing systems is received from a generation station that is subject to metering by one or

more utility-scale generation-side meters, and wherein the BTM power received by the plurality of computing systems is not metered by the one or more utility-scale generation-side meter.

7. The system of claim 1, wherein the plurality of computing systems is electrically connected to a generation station that supplies utility-scale power to a grid, wherein power supplied by the grid is subject to Transmission and/or Distribution charges, and wherein the BTM power received by the plurality of computing systems from the generation station is not subject to Transmission and/or Distribution charges.

8. The system of claim 1, wherein the set of monitored conditions further comprises:

a price of BTM power available at the plurality of computing systems and a price of grid power; and

wherein performing the weighted analysis using the set of monitored conditions involves a comparison between the price of BTM power available at the plurality of computing systems and the price of grid power.

9. The system of claim 8, wherein the set of monitored conditions further comprises:

a current price of one or more cryptocurrencies.

10. The system of claim 9, wherein the set of monitored conditions further comprises:

an indication of a hash rate, a difficulty level, and a deadline associated with performance of a computational operation.

11. The system of claim 10, wherein the set of monitored conditions further comprises:

an indication of a quantity and respective types of computing systems of the plurality available to perform the computational operation.

12. The system of claim 1, wherein modulating operating attributes of one or more computing systems of the plurality of computing systems comprises:

adjusting an operating frequency or processing speed of one or more computing systems of the plurality of computing systems.

13. The system of claim 12, wherein each of the one or more computing systems is an application-specific integrated circuits (ASIC).

14. The system of claim 1, wherein modulating operating attributes of one or more computing systems of the plurality of computing systems comprises:
causing one or more computing systems to switch to a low power mode.

15. The system of claim 1, wherein modulating operating attributes of one or more computing systems of the plurality of computing systems comprises:
causing a set of computing systems to power off.

16. The system of claim 1, wherein the datacenter control system is a remote master control system; and
wherein the plurality of computing systems, the power input system, and the power distribution system are located at a flexible datacenter positioned remotely from the remote master control system.

17. A method comprising:
monitoring, at a control system, a set of conditions, wherein the set of monitored conditions includes behind-the-meter (“BTM”) power availability at a plurality of computing systems, wherein the plurality of computing systems receive BTM power from a BTM power source;
performing a weighted analysis using the set of monitored conditions; and
based on the weighted analysis, modulating operating attributes of one or more computing systems of the plurality of computing systems.

18. The system of claim 1, wherein the BTM power source comprises a generation station configured to generate utility-scale electrical power for supply to an electrical grid.

19. The system of claim 1, wherein the plurality of computing system is electrically connected to a generation station at a location behind a Point of Interconnection between the generation station and an electrical grid.

20. The system of claim 1, wherein the BTM power received by the plurality of computing systems is electrical power produced by a generation station and transmitted to the plurality of computing systems behind the generation station's Point of Interconnection with an electrical grid.

21. The system of claim 1, wherein the plurality of computing systems is electrically connected to a generation station that is subject to metering by one or more utility-scale generation-side meter, wherein the BTM power received by the plurality of computing systems has not passed through the one or more utility-scale generation-side meters.

22. The system of claim 1, wherein the power received by the plurality of computing systems is received from a generation station that is subject to metering by one or more utility-scale generation-side meters, and wherein the BTM power received by the plurality of computing systems is not metered by the one or more utility-scale generation-side meter.

23. The system of claim 1, wherein the plurality of computing systems is electrically connected to a generation station that supplies utility-scale power to a grid, wherein power supplied by the grid is subject to Transmission and/or Distribution charges, and wherein the BTM power received by the plurality of computing systems from the generation station is not subject to Transmission and/or Distribution charges.

24. The method of claim 17, wherein monitoring the set of conditions comprises monitoring at least one condition and detecting a breach of a threshold value by the at least one condition, and wherein performing the weighted analysis using the set of monitored conditions is in response to detecting the breach of the threshold by the at least one condition.

25. The method of claim 17, wherein performing the weighted analysis using the set of monitored conditions comprises determining that BTM power is available at the

plurality of computing systems at a particular price and comparing the particular price to a price for grid power, and

wherein modulating operating attributes of one or more computing systems of the plurality of computing systems comprises transferring a computational operation from a datacenter to the plurality of computing systems.

26. The method of claim 17, wherein modulating operating attributes of one or more computing systems of the plurality of computing systems comprises:

causing a set of computing systems of the plurality of computing systems to operate at a higher frequency.

27. The method of claim 17, wherein modulating operating attributes of one or more computing systems of the plurality of computing systems comprises:

transitioning a computational operation from a first set of computing systems of the plurality of computing systems to a second set of computing systems of the plurality of computing systems; and

causing the second set of computing systems to perform the computational operation at an increased frequency rate.

28. The method of claim 17, wherein the datacenter control system is a remote master control system; and

wherein the plurality of computing systems are located at a flexible datacenter positioned remotely from the remote master control system.


29. A non-transitory computer-readable medium configured to store instructions, that when executed by a computing system, causes the computing system to perform functions comprising:

monitoring, at a datacenter control system, a set of conditions, wherein the set of monitored conditions includes behind-the-meter (BTM) power availability at a plurality of computing systems, and wherein the plurality of computing systems receive BTM power from a BTM power source;

performing a weighted analysis using the set of monitored conditions; and

based on the weighted analysis, modulating operating attributes of one or more computing systems of the plurality of computing systems.

30. The system of claim 1, wherein the BTM power source comprises a generation station configured to generate utility-scale electrical power for supply to an electrical grid, and wherein the plurality of computing system is electrically connected to the generation station at a location behind a Point of Interconnection between the generation station and an electrical grid.



PRIOR ART
FIGURE 1




FIGURE 2




FIGURE 3




FIGURE 4





FIGURE 5






FIGURE 6B




FIGURE 7




FIGURE 8




FIGURE 9




FIGURE 10A




FIGURE 10B




FIGURE 11





FIGURE 12



PRIOR ART
FIGURE 1