

© DIGITAL STOCK, PHOTO F/X, & DIGITAL VISION

By Alec Brooks, Ed Lu, Dan Reicher, Charles Spirakis, and Bill Weihl



DEMAND RESPONSE (DR) TRADITIONALLY REFERS TO THE

ability to curtail some electrical loads at peak times to alleviate the need for peaking generation sources. Basically, it means being able to turn loads off on command. Progress in communication protocols and technology has been extraordinary in the past decade, making cheap, fast communication widespread. Over the next decade, we expect inexpensive broadband to become ubiquitous. In addition, more and more electrical loads are equipped for communication as well as control. Together, these trends enable a new way of thinking about DR, which we call *demand dispatch*.

Demand dispatch is the capability to aggregate and precisely control (or dispatch) individual loads on command. Unlike traditional DR, demand dispatch is active and deployed all the time, not just at peak times. Demand dispatch represents a qualitatively different approach to balancing generation and load for a power grid. We believe that demand dispatch will be an important enabling technology for incorporating ever higher levels of intermittent renewable generation on the grid.

In this article we touch on some background requirements for demand dispatch and how the Internet can be used for communication and control. In addition, we review some of the basics of the operation of the electric power grid. We show how loads that meet the communication and control requirements can be aggregated and dispatched—turned on or off—to help manage the grid. Aggregated loads will be able to perform many of the same ancillary services for the grid that are provided by power plants today. We describe some benefits of load-based ancillary services, such as the potential for very fast response, and explain how some characteristics of load-based services differ from power plants. Finally, we give a concrete example of demand dispatch as it can be applied to plug-in electric vehicles: smart charging.

Background

One of the requirements for demand dispatch is a low-latency, moderatebandwidth communication path to an electrical device. *Low latency* in this context refers to the time delay from when a request is made by a control entity to when the electrical device receives the request and can act on it. Ideally, that latency should be less than about 500 ms. *Bandwidth* refers to the data-transfer rate required by each device. On average, we expect this to be quite small for demand dispatch.

A variety of communication mechanisms meet these latency and bandwidth requirements. In this paper, we illustrate how the Internet (and Internet network protocols) can be used as the underlying communication network for demand dispatch. Note that many networks being built by utilities for smart grid applications do not support the low latency and moderate bandwidth needed by demand dispatch.

The Internet is a network of networks that use simple, common protocols (one of which is the Internet Protocol, or IP) to enable virtually any device

Digital Object Identifier 10.1109/MPE.2010.936349

We touch on some background requirements for demand dispatch and how the Internet can be used for communication and control.

(e.g., a computer, a cell phone, or a smart appliance in a home) to communicate with virtually any other device connected to the Internet. The way the actual bits get from here to there can be very different based on the underlying mechanisms (fiber optics, wireless, copper wire, and so on), but regardless of the underlying layer the protocols remain the same. Many of these protocols are employed in applications where low latency is required (e.g., instant messaging, voice over IP, video chat, and online gaming). These same low-latency protocols could be used for certain grid-related functions in homes that have a high-speed connection to the Internet (via DSL, cable modem, or another type of broadband connection).

In the near future, many electrical devices will include two-way communication capabilities that could allow grid operators to influence when those devices consume energy. It is not unreasonable to postulate that those devices will be able to communicate via low-latency Internet protocols. Indeed, given the success of those protocols over nearly 30 years of enormous technological change, it seems far better to use such a proven approach than to invent entirely new protocols.

Power Grid Fundamentals

The power grid must be operated so that there is a real-time balance of generation and load. Without this balance, the grid frequency will drift up or down from the nominal value of 60 Hz. The grid has an effective physical rotating inertia from the combined total of all the rotating generators connected to the grid. If more power is coming in from generation than going out to loads, there will be a net torque on the system that causes all the generators to speed up, increasing the grid frequency. Figure 1 shows a sample time history of the frequency on the grid in the western United States, sampled six times a second. The slope of the frequency trace is a measure of the overall imbalance of generation and load at any given moment.

The actual grid frequency tends to oscillate slightly around 60 Hz. The frequency error from 60 Hz is used to fine-tune the generation level through regulation, an ancillary service employed by grid operators. Regulation provides grid operators with the ability to quickly ramp power up or down from a defined baseline.

Power grids today are operated for the most part by controlling generation to match load at any particular time. The overall daily profile of load in a given area can be predicted reasonably well, and from that prediction a day-ahead generating schedule can be developed. Today, loads are not generally controlled directly, except for the case when there is insufficient generation available on peak days, at which point load may be reduced through DR programs. The overall approach is that generation is controlled—or "dispatched"—to follow load. In other words, generation is "load-following" in nature.

The load-following strategy becomes more difficult as more renewable generation is added to the grid. Intermittent renewable energy sources like wind and solar generation can't be scheduled and can't be predicted with certainty. Solar energy generation can change quickly because of passing clouds. Wind energy has a different profile every day and can change output quickly. As a result, increases in wind and solar generation capacity on the grid will also require increases in conventional generation sources like peaking gas turbines. These conventional sources are needed to provide ancillary services like spinning reserves and regulation to compensate for the variability of the renewable sources.

With more and more intermittent renewable generation, it will become ever more difficult for the remaining dispatchable generation capacity (such as natural gas and coal plants) to provide the needed services and fast ramp-

> ing that ensure that generation follows load. Fortunately, there is a new option enabled by information technology: direct control of loads. As *generation* becomes less dispatchable overall, we can compensate in part by making *loads* more dispatchable.

Beyond DR: Demand Dispatch

Demand dispatch is a generalization and extension of DR. DR



figure 1. Grid power imbalance causes changes in frequency.

The value of demand dispatch to the electricity grid results from aggregating a large number of dispatchable loads.

usually refers to the ability of some loads to be shed on command in order to reduce the overall load at peak times of peak days when load is approaching the available generating capacity. In most areas, relatively few customers participate in DR programs, and DR is called upon relatively infrequently.

Demand dispatch is similar to DR in that it involves turning loads on and off. But unlike DR, which is used rarely and typically only to shed load during periods of peak demand, demand dispatch is intended to be used actively at all times to contribute services that support the operation of the grid. We call it "demand dispatch" because we are dispatching loads in real time, much as today's grid dispatches generation.

Turning on a load or increasing the demand on the grid has the same effect on the balance of power on the grid as reducing generation. Similarly, turning off a load has the same effect as increasing generation. Under demand dispatch, many loads become "generation-following."

Demand dispatch takes advantage of the scheduling flexibility of some kinds of loads. For example, a dishwasher could have a flexible start time overnight as long as the cycle was finished by morning. A washer or dryer might similarly be flexible in terms of when it runs (though not always—a household with several small children might not want to wait a day for clean laundry!).

Demand dispatch is enabled by real-time targeted communication with individual loads to provide remote load control. We believe that the coordinated control of millions of individual loads could provide extraordinarily useful grid management services.

Loads that can potentially be turned on or off remotely fall into two broad categories:

- loads that, when remotely controlled, could result in some inconvenience or discomfort for the electricity customer (e.g., lights, air-conditioning)
- loads that, when remotely controlled, would go largely unnoticed by the electricity customer (e.g., charging a plug-in vehicle at night).

The first category is largely where traditional DR has been focused, with particular emphasis on large loads such as air-conditioning systems. The second category—where the exact timing of load dispatch is likely to go unnoticed—is where demand dispatch can best be implemented.

Loads that are good candidates for demand dispatch are those in which *all* of the following are true:

It is reasonably well known how much energy will be needed over the course of the next day or so.

- The load needs to be draw power for only a fraction of that period (i.e., there is "slack time").
- The time when the load is turned on is not critical or may not even be apparent to the consumer, i.e., the main need is to draw a specific amount of energy by a specific time.
- Good candidates for demand dispatch include:
- ✓ dishwashers
- ✓ washers and dryers
- ✓ electric hot water heaters
- ✓ HVAC systems with thermal storage
- some aspects of refrigerator operation (e.g., the defrost cycle)
- ✓ battery chargers for consumer electronics
- ✓ plug-in vehicles, both battery electric and plug-in hybrid.

We have estimated that up to 33% of all loads could have at least some level of demand dispatch control without a significant impact on end users.

Applications of Demand Dispatch

The value of demand dispatch to the electricity grid results from aggregating a large number of dispatchable loads. The combined aggregated loads under dispatch control will have a minimum and maximum power-draw capability and a maximum energy-draw capability projected out into the future. Unlike a traditional power plant that can keep on generating as long as there is fuel available, dispatchable loads have a limited amount of total energy draw. For example, a battery being charged is eventually fully charged, and after that the battery charger is incapable of drawing any more energy from the grid.

Each individual load will have certain parameters that define minimum and maximum power draw, how much energy is needed and by what time, and (perhaps) limitations on the number of times it can be cycled on and off in a given amount of time. An aggregator would assemble this information from dispatchable loads and serve as the intermediary between a utility or grid operator and the individual loads. The aggregator would simultaneously meet the needs of each individual load while providing a useful aggregated demand dispatch service to the grid operator.

Loads could be aggregated within a city, in a utility service area, or over an entire control area, depending on the service being provided and the customer for that service.

The aggregated loads under demand dispatch control could provide existing ancillary services including spinning

reserves and regulation. New types of grid services could also be developed for dispatchable loads that could take into account their energy limitations as well as their capability for fast and accurate response.

There are several general classes of services that could be provided. The key characteristic of each service class would be the time scale of the load control, from hours to subseconds.

Hours

Dispatchable load could be controlled in a "generationfollowing" mode to follow some or all wind generation in a particular region. Wind generation profiles are different every day, with the time of day for peak generation varying by many hours from day to day (see Figure 2 for one month's worth of daily wind generation profiles in the PG&E territory). Dispatchable loads can respond very quickly and would have no difficulty in following the fast ramp rates sometimes seen in wind generation. This would reduce the capacity of fast-response generation assets that would need to be online.

Seconds to Minutes

The existing regulation ancillary service is typically dispatched at four-second intervals. Regulation plays the key role in fine-tuning the balance of generation and load on the grid. Regulation is mostly performed today by power plants. There are also pilot tests of regulation using energy storage (by means of batteries or flywheels). Regulation could also be accomplished with aggregated controllable loads in a demand dispatch system. Providing regulation by controlling loads could be more effective and efficient than regulation with generation or storage. Power plants that have to ramp up and down to follow a regulation dispatch profile (for an example, see Figure 3) incur additional wear and tear and exhibit higher emissions and lower thermal efficiency. Power plants also respond relatively slowly and have limited ramp rates. Energy storage, such as that provided by flywheels or batteries, can respond quickly but incurs energy throughput losses in the process of providing regulation. Dispatchable loads can respond just as quickly as energy storage systems but would incur no additional losses in providing the regulation service. The load would have been drawn anyway-only the timing of the load draw is changed to provide the service.

Under a Second

Many loads have the potential to respond very quickly upon command—faster in fact than is practical to dispatch remotely from a grid operator through an aggregator. In this case, loads could respond to locally sensed changes in grid frequency. For example, loads could be programmed to shut off upon a sudden drop in grid frequency. This could happen independently, without the need for a dispatch command from the



figure 2. Tehachapi California wind production for each day of April 2009. (Image courtesy of Cal ISO.)

aggregator. In the frequency trace example shown in Figure 4 (taken from the grid in the western United States), a sudden sharp drop in frequency is apparent. Loads that monitor grid frequency could shut off in less than 1 s when such an event is sensed.

Another load function that could be enabled by local frequency sensing is for loads to have a built-in frequency response characteristic. In this case, the load would automatically vary linearly with grid frequency error, much like a traditional thermal power plant droop characteristic in which power generation increases or decreases depending on grid frequency. As traditional power plants are supplanted with renewable generation this power plant droop characteristic will in many cases be lost; many renew-



figure 3. A regulation dispatch example for PJM over a 24-hour period. The regulation command is typically updated every 4 s.





able generation technologies do not have a droop characteristic. For example, the wind doesn't blow any harder when the grid frequency droops. This loss of power plant droop characteristic will be destabilizing to the grid. A load-based droop characteristic can help to compensate.

Note that a frequency-based droop characteristic is closely related to regulation ancillary service. A key element of the regulation dispatch command is based on the frequency error. In fact, PJM has a regulation ancillary service that is based solely on frequency error. There is the potential for a load-based regulation service that is very accurate and fast and also does not require much, if any, communication for dispatch, as the frequency can be sensed right at the load.

Example: Regulation Ancillary Service with Demand Dispatch

The regulation ancillary service is usually the most valuable as it requires the fastest control and response of all the ancillary services. Regulation is a contract to provide the capability to remotely control power up or down from a nominal value. The amount of regulation contracted is the total amount by which power can deviate from a baseline level. The baseline is often called the preferred operating point, or POP. Contracts are typically over a one-hour period. Figure 5 shows regulation as it could be performed by an aggregated controllable load. Note that the shaded area in the figure represents the energy draw from the grid for that period of time. The energy will typically be approximately the POP multiplied by the duration of the time period, as the regulation dispatch signal usually averages to the POP value over time.



figure 5. Regulation signals.

Demand dispatch is well suited for providing regulation. An aggregation of loads will be capable of some range of power draw between a minimum and maximum value. The midpoint of this range would typically be bid in as the POP, with regulation-up dispatch representing decreasing the load and regulation-down dispatch representing increasing the load. (Note that it is often incorrectly assumed that loads are only capable of regulation down service. This is not the case; equal amounts of up and down regulation are possible with the POP set as described above.)

Demand Dispatch Example: Smart Charging

Plug-in vehicles are being developed by almost all major automakers and will be introduced in volume within the Plug-in vehicles could be an excellent demand dispatch resource, given their potential for rapid response and expected rapid growth in the number deployed.

next year or two. Plug-in vehicles derive some or all of their energy from grid electricity. The daily energy use for driving 30 mi electrically for a typical passenger car will be on the order of 10 kWh, making plug-in vehicles one of the biggest energy-consuming devices in a household.

Plug-in vehicles will usually be charged at night or at work. Most vehicles are driven only one or two hours a day on average and are parked the rest of the time. The time needed to recharge 10 kWh of energy to a vehicle battery will typically be only 2–5 hours. But most vehicles can reasonably be expected to be plugged in for 10–15 hours a day (more for vehicles that are regularly plugged in both at home and at work). The difference between the elapsed time needed for actual charging and the time that the vehicle is plugged in results in timing flexibility that can be harnessed to provide grid services while at the same time meeting the needs of the driver. We call this "smart charging."

Smart charging is more than just charging at off-peak times. It involves fine-grained control of the charging of each vehicle to meet both the needs of the vehicle owner (charging the vehicle by a certain time) and the needs of the grid (matching generation and load, providing frequency regulation, and perhaps also avoiding overload in distribution networks from many vehicles being charged at the same time).

Some vehicles will have tighter constraints on charging than others. A plug-in hybrid electric vehicle (PHEV), which can be viewed as an electric vehicle with an onboard "range extender" that uses a conventional fuel (e.g., gasoline) to provide energy when the battery has been discharged, has enormous flexibility in charging: if the battery is not completely charged when the vehicle is needed, the vehicle will simply use a little more onboard fuel. A battery-electric vehicle will likely have tighter constraints on charging, depending on its range and its owner's driving patterns.

Most new plug-in vehicles will be equipped with Internetenabled communications through one or more means: Wi-Fi, power line carrier, cell data network, or WiMAX. With this capability, plug-in vehicles could be aggregated to provide a dispatchable load resource. When a driver parks her vehicle and plugs it in, actual charging would generally not start immediately. Rather, plugging in the vehicle would initiate a sign-on with the aggregator. The vehicle would communicate status information such as location, required energy, the maximum and minimum recharge power draw, and when the driver wanted the recharging of the vehicle to be complete. The driver would be in control of this process and would set up a normal vehicle usage profile via a secure Web page. After the initial setup, the normal sign-on each time the vehicle is plugged in would occur without driver input. Exceptions such as "charge now" could easily be entered through the driver's mobile phone or through a dashboard control panel.

Once a vehicle was plugged in and logged on to the aggregator, the aggregator would determine how and when the vehicle would be charged. By individually controlling the charging of each vehicle, the aggregator could provide a demand-dispatch service to a utility or grid operator. The aggregator would employ algorithms to optimize the provision of services while at the same time meeting the charging requirements for each vehicle. Vehicles logging on with a greater amount of "slack time" would offer more charging flexibility to the aggregator and thus would be expected to reap a greater value in return. A vehicle logging on with a "charge now" condition would offer no flexibility for control of the recharge timing and would be expected to pay regular time-of-use prices for the electricity consumed.

Plug-in vehicles could be an excellent demand dispatch resource, given their potential for rapid response and expected rapid growth in the number deployed. It could even be possible that all regulation ancillary service could be provided by plug-in vehicles and other dispatched loads just 20 years from now.

Figure 6 shows a scenario in which plug-in vehicles alone could provide all of the regulation in the PJM control area. The upper graph in the figure shows a representative day of regulation dispatch in PJM. The range is about +/-1 GW. The lower graph shows how this regulation could be served by demand-dispatched (smart-charged) plug-in vehicles. For this example, the POP has been set at 1,300 MW of load. Regulation up is accomplished by reducing the load from the 1,300-MW level and regulation down by increasing the load beyond 1,300 MW. The blue-shaded area in the graph represents the total amount of energy drawn by the loads providing regulation. The energy in the shaded area is about 32 GWh, enough to charge 3.2 million vehicles with 10 kWh a day each, good for about 30 miles of travel. (PJM manages the delivery of power to more than 50 million people.)

Figure 7 shows how smart charging might work. Plug-in vehicles are connected to the grid for power and to the Internet for data communications. The network connection could be accomplished through a number of different physical networks (such as Wi-Fi, power line carrier, Zig-Bee, cellular data, etc). In order to provide the fast response that is valued in regulation, the communication channel should support round-trip timing of charging commands and



figure 6. Example of providing regulation with load. (a) Regulation dispatch in PJM over a 24-hour period. (b) How all regulation could be provided with dispatchable loads, in this case 3.2 million plug-in cars.

acknowledgment on the order of one or two seconds.

The grid operator continuously evaluates grid operations data and determines an aggregate load dispatch command, typically every four seconds. The command is sent to the aggregator, and the aggregator then determines which connected vehicles to contact in order to comply with the load change represented by that particular command. The process repeats every four seconds. It is not necessary to communicate with every connected vehicle every four seconds-only with the vehicles that need to change their charging rates. With a large number of vehicles participating, it would be practical to simply turn charging on or off rather than trying to modulate the charging rate of each vehicle. With on-off control, only a small subset of the connected



figure 7. System diagram for plug-in vehicles providing regulation. (Grid data and operator courtesy of CAL ISO; vehicles courtesy of Google; transmission tower courtesy of Wikimedia Commons.)

vehicles would need to be contacted at each step.

Figure 8 shows the operation of a prototype smart charging system developed by Google. The smart charging algorithm is structured to provide as much regulation ancillary service as possible while meeting the energy needs of each individual vehicle while it is connected to the grid. The prototype system is set up to control both real and simulated vehicles. We have equipped our fleet of eight plug-in hybrid vehicles with a smart charging communication system and have run simulations and tests with both the real vehicles and a larger suite of simulated vehicles.

The system is set up such that the aggregate total load under control can be dispatched in a blended fashion, both to follow wind generation and to provide regulation ancillary service. Figure 8 shows the operation of the system using data from ERCOT (the Texas grid). The upper graph shows the wind Over the next few years a large number of consumer devices from televisions to household appliances to cars—will have communications built in, most likely using Internet protocols.



figure 8. Plug-in vehicle smart charging example. The individual vehicle charging is dispatched on and off. The overall vehicle charging load profile is controlled to be a combination of wind following and regulation.

generation profile over a 24-hour period. The second graph shows the regulation dispatch profile in ERCOT over the same period. The third graph shows a calculated blended load dispatch command with a combination of wind following and regulation. Note that regulation-up dispatch means decreasing load: regulation up has been traditionally considered as increasing generation output—hence "up"—but decreasing load has the same net effect on the grid. As a result, the short-period deviations evident in the overall load profile are inverted as compared with the regulation dispatch profile.

The graphic on the bottom shows a simulation of several hundred plug-in vehicles while they are smart-charged under aggregator control. Each row in the table represents one vehicle. Places in a row with a colored line indicate that the vehicle's battery charger is turned on and drawing power. The sections of a row with just the background color represent periods when the vehicle charger is turned off by the aggregator. The color in a row represents the vehicle's state of charge: dark blue represents a low state of charge, and white represents a full state of charge.

The aggregator sends out dispatch commands every four seconds, or 900 times an hour. Any given vehicle gets a command only four to eight times every hour. At each new command received from the grid operator, the aggregator determines which individual vehicles to turn on or off based on their status and how soon particular vehicles need to arrive at a fully charged state. At each new dispatch point (i.e., at four-second intervals), the aggregator sends out "on" or "off" commands to a small subset of the connected vehicles.

The historical average price of regulation in PJM is about US\$35 per MW per hour, according to a presentation made at the Plug-In 2009 conference. In PJM, regulation up and regulation down are combined into a single service. A contract for 1 MW of regulation means that power can be commanded to increase or decrease by 1 MW from the POP. Providing 1 MW of regulation with demand dispatch could be accomplished

with aggregated load that would be capable of being dispatched between 0 MW and 2 MW of load. Since regulation is generally symmetrical about the POP over time, the expected average load would be halfway between 0 MW and 2 MW, or 1 MW. Over one hour, the aggregated loads would thus be expected to draw 1 MWh of energy. With US\$35 payment for this service for the hour, the value of the payment per unit of energy draw would be US\$0.035/kWh. This could significantly reduce the cost of energy for those loads that have the timing flexibility to participate in demand dispatch programs.

Figure 9(a) shows some of Google's fleet of plug-in Priuses smart-charging under the solar carport on the company campus in Mountain View, California. In Figure 9(b), the charging graphs compare conventional charging with smart charging. With smart charging, the charging power is turned on or off remotely by the aggregator. This results in slower overall charging, but the aggregator takes into account the time the vehicle driver has specified that the battery must be fully charged and controls the charging profile accordingly. The needs of the driver take first priority, and the available "slack time" or charge timing flexibility is optimally utilized by the aggregator to provide the greatest possible amount of grid services.



figure 9. (a) Google PHEV fleet and (b) a comparison of conventional and smart charging. (Google fleet courtesy of Google.)

Security

Along with the benefits of demand dispatch, there are some associated risks. Demand dispatch by its nature involves the coordinated control of a large number of loads. When operating as planned, this coordinated control offers the prospect of substantial grid benefits. However, this ability also has the potential to be harmful to the grid. As expected in any complex system, there is the potential for human error, software bugs, or malicious intent to disrupt the intended control of demand dispatch–enabled loads in a way that could be destabilizing to the grid.

Although the initial demand dispatch trials will be small, security needs to be designed in from the beginning. Mitigating risk can be done using a variety of approaches, from legal to policy to technical, and the best solution is likely to involve a combination. Some examples include:

- ✓ limiting the size of any given aggregation entity
- certifying aggregators and aggregation software, much as ancillary service providers do today
- designing systems for fault detection and flagging and inhibiting highly unusual commands (e.g., prohibiting turning all on or all off at once)
- designing end-user loads to do a sanity check on the commands received (e.g., a command to turn on to full load when the load detects that grid frequency is falling fast might be flagged as an invalid command and ignored).

A full security discussion is beyond the scope of this article, but Google is engaged with NIST and various smart grid task forces promoting and working on security as an integral part of the smart grid standards discussions.

Conclusions

The proliferation of high-speed, low-latency communication networks enables large-scale, low-latency individual communication with loads. Over the next few years a large number of consumer devices—from televisions to household appliances to cars—will have communications built in, most likely using Internet protocols. This new ecosystem of interconnected loads will enable the valuable new capability of demand dispatch, which can play a major role in the management of the grid.

In the future, the grid will evolve away from the current model, in which generation is controlled to follow load, to a new approach that includes demand dispatch, in which some of the generation and some of the load are both controlled to achieve an overall balance of generation and load. Demand dispatch will be especially useful to assist in the integration of an ever greater share of intermittent renewable generation and could also be critical to the successful deployment of large numbers of plug in vehicles.

For Further Reading

K. Huber. (2009, Dec.). Testimony to Senate Energy and Natural Resources Committee [Online]. Available: http:// www.pjm.com/documents/~/media/documents/congressional-material/20091210-senate-testimony-huber.ashx

MAGICC – Mid Atlantic Grid Interactive Cars Consortium [Online]. Available: http://www.magicconsortium.org/

Pacific Northwest National labs (PNNL) smart charging controller [Online]. Available: http://availabletechnologies. pnl.gov/technology.asp?id=284

B. Kirby. Frequency regulation basics and trends. [Online]. Available: http://www.consultkirby.com/files/TM2004-291_Frequency_Regulation_Basics_and_Trends.pdf

Lawrence Berkeley National Lab. (2007). Loads providing ancillary services: review of international experience, [Online]. Available: http://www.oe.energy.gov/DocumentsandMedia/Loads_providing_Ancillary_Services_main_ report_62701.pdf

Biographies

Alec Brooks is a renewable energy engineer at Google.

Ed Lu runs a number of projects at Google, including its energy information efforts.

Dan Reicher is Google's director of climate change and energy initiatives.

Charles Spirakis is a software engineer at Google. *Bill Weihl* is Google's green energy czar.