Lawrence Berkeley National Laboratory

Recent Work

Title

DC Local Power Distribution: Technology, Deployment, and Pathways to Success

Permalink

<https://escholarship.org/uc/item/1xp921jv>

Journal

IEEE Electrification Magazine, 4(2)

ISSN

2325-5897

Authors

Nordman, B
Christensen, K

Publication Date

2016-06-01

DOI

10.1109/MELE.2016.2544218

Peer reviewed

DC Local Power Distribution

Technology, deployment,
and pathways to success.

DIRECT-CURRENT (DC) POWER distribution has been used ever since electric grids were invented, but, for the last century, low-voltage dc has been largely limited to a variety of niche applications such as rail transport, vehicles, telecommunications, and off-grid buildings. Recent years have seen a variety of innovations in dc distribution technology, notably, standards for 380-V dc cabling and connectors and increases in power that can be carried over Ethernet and universal serial bus (USB). There are increasing calls for much more use of dc distribution and dc microgrids in buildings, and there are potential advantages of both. However, open questions remain about the directions this might take, what policy makers could and should do in this area, and technology developments that would be most useful. This article considers potential pathways for increased use of dc and identifies those pathways that seem most beneficial and likely to succeed. We limit the scope of consideration to distribution within (or between) buildings.

DC Power in Buildings

When an alternating-current (ac) utility grid power is available to a building, directly distributing and using

power as ac is the obvious initial choice. If costs and benefits are similar, ac has a great advantage as the incumbent technology in product and parts availability and its familiarity with designers, electricians, contractors, building owners, and local planning departments. There has always been clear advantages of ac power within utility grids, particularly for the ease of changing voltages, and those advantages remain. Thus, for those who would like to see more use of dc distribution in buildings, the focus needs to be on cases for which the net benefits of dc are both substantial and compelling.

An increasing fraction of electrical load in buildings is natively dc, most recently with the rise of light-emitting

IMAGE LICENSED BY GRAPHIC STOCK

diode (LED) lighting and increasing loads from electronic devices such as computers and displays. Some devices, including fluorescent lighting and variable speed drive motors, convert ac to dc (and then to high-frequency ac), so are also dc internally. A large portion of the remainder of loads in buildings could readily be converted to dc at equal or greater efficiency than with ac. Vehicle charging with dc is not the norm, but it is available.

DC distribution today is highly successful in specific niche applications but is rarely used outside of them. It is used in 12- and 24-V distribution in vehicles, USB for mobile device charging, Ethernet for desktop phones in office buildings, and off-grid buildings; these are all very specific applications in which dc distribution and use have advantages over ac. In data centers, 380-V dc is seeing increasing uptake. The use of dc may not be competitive today for mainstream or generic use, but that is beginning to change with the rise of lighting powered by Ethernet.

The use of dc distribution in buildings is growing, primarily at the very edge of the building power-distribution tree. Power over Ethernet port shipments has consistently risen during the past decade, from under 10 million/year in 2005 to more than 100 million/year in 2015. With the expected increase soon in Ethernet power to nearly 100 W (more than six times the original power limit), many more devices will be brought into the scope of being dc powered, and shipments should thus grow even more rapidly. USB annual port shipments are approaching 5 billion. With USB recently enhanced to provide 100 W, it also can now support many more product types than before. With this capability, USB can now provide 40 times as much power as it did when originally introduced more than 20 years ago.

There are three basic approaches to power-native dc end-use devices, with two involving dc distribution. Many buildings contain two or even all three of these architectures in different places. In the present paradigm of building operation (ac distribution), the source of all power is

originally ac from the grid, with conversion to dc occurring inside of each dc device (with an internal or external power supply). There is no alteration or addition to the building wiring. This is shown in Figure 1(a). The second option, central dc distribution, moves the location of the ac–dc conversion to a central device and distributes dc to end-use devices, as shown in Figure 1(b). This may be more efficient [i.e., save energy compared to the scheme of Figure 1(a)] in some uses because it reduces the number of front-end components that must engage in power conditioning to deal with variations on the input ac bus and provide stable dc output. Central dc distribution is often used to obtain cost savings, greater reliability, and convenience. If any local power generation and/or storage are present, they are connected to the ac infrastructure.

More interesting and useful are power architectures in buildings with local generation [most commonly, solar photovoltaic (PV)] and local storage, as shown in the right of Figure 1(c). Direct dc enables power to flow from generation to end use—through storage as needed—without ever having to be converted to or from ac. This saves energy from the avoided conversion losses and saves capital by requiring less conversion hardware, thus saving money from both. Direct dc increases reliability by avoiding potentially unreliable conversion hardware and by the presence of battery storage, because storage is less expensive to include in an all-dc system. Savings estimates from direct dc vary dramatically, from 2 to 14%, in part due to differences in the application contexts and baseline assumptions.

The ac–dc conversion in direct dc can be bidirectional, to enable exporting excess power out of the dc domain; this provides operational flexibility. A simple alternative is to install a unidirectional interface, as shown in Figure 1(c), in which electricity is never exported from the dc domain; we call this “semidetached direct dc.” Power is imported across the ac–dc link as needed, but for long periods of time no power may be flowing across this link. In this model, a unidirectional link means that dc generation and storage are effectively invisible to the utility grid; the dc system is just another load. This approach should reduce permitting and regulatory burdens of installing direct dc systems. In cases in which more generation exists than can be stored or used for ordinary purposes, inexpensive resistance heating of space or water could productively use excess power.

Savings estimates from direct dc vary dramatically, from 2 to 14%, in part from differences in the application contexts and baseline assumptions.

Figure 1. The (a) ac distribution, (b) central dc conversion, and (c) direct dc.

When generation and loads are compatibly sized, the vast majority of electricity generated in the dc domain can remain there, and the vast majority of energy used in the dc domain is also generated there. A result is that only small amounts of energy cross the ac–dc link, and the electrical capacity of the link can thus be significantly downsized.

Many end uses such as lighting and electronics require relatively consistent daily use of electricity, so generation and storage for such local dc grids can be sized to the needed consumption level. In Figure 1(c), if either the generation or storage is eliminated, some of the benefits of direct dc are retained, but with decreasing costs, it is likely that systems with both generation and storage will be the norm.

DC Technology Architectures

Although direct dc is possible to implement today—and is effectively used in some installations—it is uncommon and lacks standard designs on which to draw. Some voltages (e.g., 12 V) lack standard connectors, although in recent years progress in the area of connectors has been made by the Emerge Alliance for 380- and 24-V dc (Figure 2 shows standard dc connectors for three technologies). Even if common deployment schemes were available, it is likely that the use of direct dc would remain uncommon if the only benefit delivered was energy savings. Key to making direct dc more successful is to focus on other benefits and develop new technology to create more benefits and/or decrease costs.

DC power distribution exists today in two primary forms—and a third form (networked dc) could be developed—in terms of how it is organized and structured and the way in which it operates within buildings.

- 1) *Traditional dc* moves power through circuits as determined by Kirchhoff's laws, similarly to the way in which ac systems operate. For example, in vehicles, power distribution is organized into circuits, similar to power distribution of circuits in ac building infrastructure.
- 2) *Managed dc* moves power across a single cable from one device to another, with characteristics (including voltage and current) determined principally by digital management as communicated between the two devices and implemented with modern power electronics. USB and Ethernet both perform this process and regularly increase capability. Managed dc derives from modern communications technologies more so than from traditional power distribution (ac or dc).
- 3) *Networked dc* extends managed dc from single power links to a network or mesh of power entities

Figure 2. The standard dc connectors for (a) a USB, (b) an Ethernet, and (c) a 380-V dc.

(nanogrids) of arbitrary topology and scale. Products are unavailable today that implement networked dc, but the local power distribution (LPD) technology described below would make this possible. Similar to managed dc, on which it is built, networked dc has more in common with network communications technologies than with traditional power distribution technologies. Figure 3 shows an example network of dc nanogrids, including local generation and vehicles, with a connection to ac infrastructure.

Networked dc makes use of a relatively new feature of some managed dc technologies—bidirectional power—for which the direction of power flow may be different at different times. This exists in the latest version of the USB Power Delivery specification as well as with HDBaseT, a variant of Ethernet. Networked dc does not require any innovation in how electricity is transferred, but it does require new capabilities in communications and control. Figure 4 shows a USB hub and an Ethernet switch; with the addition of modest additional communications technology, each of these could become a fully functional nanogrid controller, although adding some electricity storage as well will likely be the norm.

A key principle of LPD is to decouple how power distribution is managed from functional control protocols in which end-use devices participate, similar to how physical layer and application layer protocols are separated from each other with the Internet Protocol (IP) in IP networks.

Figure 3. An example network of dc nanogrids.

Figure 4. (a) A USB hub and (b) an Ethernet switch, devices that could be readily modified to be nanogrid controllers.

These separations make each part simpler and more effective. In addition, the topologies of power distribution, data communications, and building structure can all be disjoint. Communication is needed to determine where power should flow in these networks. In networked dc, power flows are determined by computation.

What is LPD?

LPD is a network model of power. Electricity is managed in arbitrary and dynamically changing topologies of sources, storage, and end-use devices. Power connections are peer to peer between two devices, not via buses with many devices attached. All power exchanges are digitally mediated. End-use devices are organized into nanogrids—single domains of power for voltage, capacity, reliability, and management. A nanogrid controller manages the power distribution to end-use devices connected to the controller and exchanges power with other nanogrids; local generation (or a utility grid connection) is a special form of a nanogrid controller. Controllers almost always have local storage to aid in balancing supply and demand over time. In LPD, there are only two device types and two types of links, paralleling the architecture of the Internet. Figure 5 shows a diagram of a nanogrid controller with attached end-use devices, internal storage, and connections to other controllers.

In data networks, all packets are different, addresses provide their destination, and routing mechanisms

Figure 5. A nanogrid controller.

determine the path required to get to the destination. In contrast, all electrons are the same, so how do they know where to flow in a power network? The answer is that they should flow from places at which they are more available (plentiful supply) to those where they are less available (more in demand). Balancing supply and demand is a basic function of any grid—and of our economy in general. Elsewhere in our economy, price is the basic mechanism used to accomplish resource allocation. LPD, based on a local price set by each nanogrid controller, directs electricity to flow toward nanogrids with higher prices, much as gravity forces water to flow downhill.

With generation and storage in many places in the network, power might flow in different directions at different times, as supply and demand change. The prices are local to each nanogrid. Local pricing is the core mechanism used for coordinating among devices on when to generate power, when to charge or discharge storage, and how much and when to use electricity in end-use devices. The price includes a nonbinding forecast of future prices to help inform decisions about modulating device behaviors over time.

LPD capabilities can be added to managed dc technologies, such as USB and Ethernet, and to new physical layers of power that implement managed dc. Some will have higher power capacity. A key feature of LPD is that it will enable both storage and generation of electricity to become plug and play, i.e., able to be connected and disconnected at will by anyone without the safety risk otherwise present when using higher voltage levels. This makes changing power infrastructure something that can be done simply, inexpensively, and frequently, as local conditions require.

How LPD Can Enable DC Success

Direct dc with LPD together provide a compelling value proposition and a plausible deployment path that can tip dc power from being a small player on the edges of electricity distribution to a sizeable component—to become a success in the market by delivering value to users and saving energy. The following are benefits that LPD can bring to power systems in any building.

Optimal Operation with a Local Price

A core element of the definition of a nanogrid is a local price that correctly indicates the relative scarcity of power and so drives efficient operation of end-use devices, local generation, and local storage. It can also be used to optimize exchange of power with a utility grid. Analog voltage levels can be used to indicate scarcity, but this is not accurate and precludes including a forecast of future prices and negotiating between devices. Forecasts are essential to moving generation, storage, and end use across time for better system operation.

Communications are also needed for devices to know when they should switch the direction of power flow in a

link. Some technologies allow for bidirectional power today; others could accomplish reversing power flow with two parallel links, one for each direction. Power links to vehicles will become a common example of bidirectional power flows. We can also expect vehicles to be able to connect directly to one another, so that one can charge another anywhere. Management of power distribution with and within vehicles should use the same technology as is used everywhere else.

Plug-and-Play End-Use Devices

In existing managed dc technologies, initially only a small amount of power is provided to enable communications. Each device communicates to the other about its capabilities and preferences to allow for adjusting and optimizing the link. The voltage can be negotiated to the highest that can be provided and then used and is also appropriate to the cable (based on its capacity and length). This minimizes resistance losses and maximizes power capacity. Communications can be used to determine cable length. In some USB links, the cable itself reports its characteristics to the devices, and a variety of combinations of voltage and current are available that the electricity-providing and electricity-consuming devices might support. This allows for the most efficient combination that both devices can use to be selected. When a device is plugged in that is not compatible electrically for any reason, this can be detected and reported.

Improved Safety

By communicating before delivering power, conditions that would otherwise be unsafe can be avoided. If a cable is cut or improperly connected, this will be detected by an interruption in the communication. This signal can be used to terminate ongoing power delivery or to stop delivery from the beginning. In such a case, one proprietary mechanism cuts power within 3 ms to enable 400-V delivery over Ethernet cabling. Capacity and voltage limitations of the cable or either device can also be respected automatically. Cable overheating could be detected through the two devices comparing the supplied and received power levels and recognizing when the difference between these is inappropriately large.

Managed Limited Supply

Each electrical circuit has a maximum current or power level that it can manage, due to wire size, length, temperature, and other factors. In ac systems, a circuit breaker

A key feature of LPD is that it will enable both storage and generation of electricity to become plug and play, i.e., able to be connected and disconnected at will by anyone without the safety risk otherwise present when using higher voltage levels.

will cut power to the entire circuit (and hence all devices on it) when the limit is reached. This is inconvenient if it occurs frequently. To avoid it, ordinary practice is to substantially oversize wires and circuits, and building codes generally require this, even on circuits that will never carry such high current. This is a waste of the copper and capital and increases installation costs.

With managed dc, actual power levels being used can be tracked. In addition, what devices can potentially use at maximum and what they are actually consuming at any particular moment can be considered to then install equipment of appropriate capacity. That is, capacity limits can be respected through pricing and communication rather than oversizing and blowing breakers. In emergency situations, devices can be summarily disconnected on

an individual basis as needed to reduce total demand rather than interrupting power to the entire circuit.

Plug-and-Play Generation

Traditional electricity systems (ac and dc) use generators that can vary their output to follow electricity demand based on voltage changes, thus maintaining balance in the system. Communications enable systems in which demand can respond to supply conditions, and in systems with more than one generator, communications can determine generator on-times and output levels. When connected to a utility grid, the best balance of utility power and local generation can be determined. Generation technologies are often nondispatchable (as with many renewables), lose power on conversion between ac and dc and between different voltages, and have variable part-load efficiencies, minimum times to be on or off, and losses on each cycle up and down. With communications, these factors can be coordinated to maximize efficiency and equipment use and improve reliability and safety. Communication can ensure that a generator can safely deliver power of the quantity and type desired, before it begins doing so. Without this ability, safe and efficient operation requires careful system design and management, as well as extra hardware. Optimal system operation is impossible without effective communications.

Plug-and-Play Storage

Electricity storage as a common feature in electricity systems is only now emerging. Existing use of storage has been for reliable or disconnected operation of individual

devices or in uninterruptible power supply (UPS) systems; usually the battery is only ever used when the primary supply is lost, and then it supplies all demand. LPD enables the best use of batteries all the time, taking into consideration lifetime impacts of battery cycling. Without communications, a storage device will generally not know if it should be charging or discharging (or neither) and at what rate. Some systems use changes in voltage levels to communicate about system status, but this is not always reliable and greatly limits the scope of possible information passed. In complex systems with multiple local generation and multiple local storage devices, only with effective communications can proper, efficient, and economic use of electricity storage be accomplished.

Privacy and Cybersecurity

Communication in LPD is only ever point to point—there is no multihop communication as in data networks. This dramatically reduces the scope for concern for security and privacy, because only devices with a direct wired connection can attempt to disrupt or spy on one another. In addition, because the communications are so simple, there is little opportunity for mischief. This is greatly different from the risks and security vulnerabilities that occur in ac grid-tied systems, particularly as large numbers of devices owned by different people are connected.

New Powering Models

In the past, end-use devices were only ever connected to a single external power source. There are rare exceptions, such as electronic devices in data centers or telecom

facilities. However, managed dc creates the possibility of devices easily able to use power from more than one source, at different times, or at the same time. This is particularly useful when resiliency is of concern, because a device can be powered through one means most of the time, but by another means when the first is not working or is more expensive. For example, it would be convenient if refrigerators could take in dc power from a vehicle or local generation or storage during times when the utility grid is down, or expensive, and use grid ac power otherwise.

AC circuits are usually multidrop—many devices can be attached to a single wire. DC topologies more commonly provide only a single device per power port. This ensures that capacity limits on the cable, and cable length, are not exceeded. However, with communications, it could be possible to create multidrop capability for technologies such as Ethernet, with cost advantages for many devices, including lighting, to enable easy daisy-chaining.

The recent 380-V dc standard lacks any mechanism for communicating about power. A good option to remedy this is to use standard Ethernet links, in parallel to each power link. The Ethernet link could provide small amounts of power (relative to what the 380-V path could provide) to energize the end-use device for communications to negotiate aspects of the 380-V line before it is energized.

DC Deployment Architectures

Direct dc, particularly when combined with LPD, raises the question of how LPD should be deployed in buildings with most economy and benefit. The dc portion of building infrastructure can be treated as a single cloud of technology for this purpose, including generation, storage, and connections to vehicles. DC infrastructure can be integrated with ac buildings in several ways that differ primarily in how reliable device operation is accomplished. Power reliability and quality are often key drivers for dc adoption.

The first way to integrate dc is to simply attach dc systems to the existing ac infrastructure with a rectifier or inverter/rectifier, as shown in Figure 6(a). The ac system is unchanged in topology or capability. The dc system can be reliable in the face of grid failure, as long as power is prevented from being exported to the ac side during these times. The data line between the meter and dc grids passes price information; it must come from the meter, not the grid, in cases when the price is different depending on whether the building is buying power from or selling power to the grid, as some regions are soon planning to do.

A second method is to create some reliable ac infrastructure as strictly subsidiary to the dc grids, as shown in Figure 6(b). The reliable ac power is always produced from dc and thus can also be insulated from power quality issues that may be present on the utility grid; needs for quality and reliability are highly correlated. The reliable ac is decoupled from the ac grid at all times. Typically, the amount of power in the reliable ac domain is small in

Figure 6. Alternative methods to integrate dc grids and reliable ac. (a) Attaching dc systems to the existing ac infrastructure with a rectifier or inverter/rectifier, (b) creating reliable ac infrastructure that is strictly subsidiary to the dc grids, and (c) carving out a portion of the ac infrastructure that can be powered from the dc side during times of grid failure.

comparison to that in the dc grids. This is convenient in that the ac infrastructure is changed minimally.

A variant of this is to carve out a portion of the ac infrastructure that can be powered from the dc side during times of grid failure, through the inverter, as shown in Figure 6(c). This adds an islanding switch to assure that power does not flow to the utility grid when it is not operating. There does need to be a mechanism to extend the price-based management of the dc grids to the reliable ac domain, to ensure that supply and demand can be best balanced. In all of these cases, dc systems are reliable, and easy reliability is a key benefit. In LPD, the degree of reliability can be varied from nanogrid to nanogrid and device to device. Changes to the ac infrastructure are minimized as much as possible to minimize costs. A long-term principle that can be adopted is that all devices for which power reliability and quality are of particular concern should be dc powered. This would mean that no capital or other expense would need to be expended to make ac infrastructure reliable, and over time, the utility grid itself could optimally be tuned to lower reliability and quality levels, saving considerable capital and energy.

DC Deployment Paths

One approach to dc deployment is to install a large amount of direct dc infrastructure at one time. This is possible for new construction or major renovations, but for most buildings it is impractical in the near term. An alternative is to slowly evolve toward major use of dc and direct dc in a building, through a series of many steps. DCs can thus be deployed incrementally and organically and can be introduced as opportunity or need dictate. The evolutionary approach can be used in circumstances such as those in the following examples:

- ▶ spot reliability; when there is a need or desire to make some devices reliable in the face of a grid outage, installing direct dc infrastructure can be done for those particular devices, rather than for a large part of the building
- ▶ modest remodeling projects; small projects can be used to introduce dc hardware
- ▶ large device replacement, such as an appliance or climate control system; although it may not make economic sense to replace ac devices that are functioning properly, it is a much smaller hurdle to shift to a new dc device rather than buy a new ac device
- ▶ occupancy changes; occupant needs change over time, as do the occupants themselves, introducing more opportunities for changed infrastructure.

A building owner may be interested in dc power but reluctant to move a lot of infrastructure over to it at all or

all at once. The evolutionary approach allows modest investment and risk to allow dc to prove itself in cost effectiveness and other benefits, so that financial and other risks are minimized. DC technology is very much in flux, and so for many end uses in many buildings it is premature to convert to dc. However, this need not preclude installation of dc devices for which the merit of doing so is clear today. DC devices can be introduced with external central or local ac–dc conversion, and dc sources added later to produce direct dc.

This model of technology deployment is familiar for information technology (IT) systems, where conversion of functions from analog to digital, or upgrading or adding new functions, is generally done on an ad hoc basis rather than through single large-scale upgrades. The nature of IT networking is such that swapping in new hardware and connectivity is not as burdensome in the way that it

is for modifying traditional electrical systems. This is another way in which networked dc is inspired by architectures and capabilities of IT.

Reliable Communications

Another current opportunity for dc distribution is in reliable communications. The U.S. Federal Communications Commission (FCC) has been concerned in recent years with ensuring residential communications continuity during utility grid outages. For many decades, telephones were reliable in such circumstances because they were very low-power devices

powered through the communications wiring from the telephone central office. Today, communications is more commonly a combination of text messaging from mobile phones, e-mail from PCs, and voice-over IP phone calls, requiring a combination of modems, handsets, network equipment, mobile phones, and computers, all of which must be continuously powered or periodically recharged.

It is possible to supply ac power to these devices via a generator (or battery), but this would be cumbersome and relatively expensive, particularly given the relatively low power levels involved and the expense of hardware to automatically make the switch and ensure continuous power delivery. Another alternative is to place batteries into each end-use device, but this would make them costlier, bulkier, and more failure prone. A better solution is to power all devices needed for communication via USB from a central hub, eliminating ac–dc conversions within each device—because all of the involved devices are dc internally. The hub could have a battery recharged from the ac power and, even better, have a link to a PV panel to provide power when the utility grid is down and displace grid power when it is up. USB could add a mechanism to indicate when power is scarce so that the end-use

The economics of daisy-chaining devices on a power cable are significant, and developing ways to accomplish this in the context of LPD will be valuable.

devices could scale back or power down less critical features; related to this, the current draft of the update to the Ethernet standard has such a feature, a local price index, showing the practicality of the LPD approach. This example of powering devices needed for reliable communication is another example of how dc distribution can be introduced for specific needs to gain multiple benefits. In cases such as these, dc distribution for communications devices could be initially a dc island but later connected to other dc domains, including local generation, for greater efficiency and reliability.

Challenges and Opportunities

The key near-term challenge is to do the research necessary to determine the communications needed between grid controllers and end-use devices and, secondly, between grid controllers. Once this is done, the features required can be added to all physical layers that implement managed dc. Managed dc link technologies above 100 W are needed; one implementation of that is from Voltserver, Inc., which offers technology for managed dc over 1 kW, over relatively thin cable by using much higher voltages while maintaining safety. The economics of daisy-chaining devices on a power cable are significant, and developing ways to accomplish this in the context of LPD will be valuable. LPD is based on a model of peer-to-peer power exchange. It may be possible in some contexts (perhaps for higher power levels) to create managed dc power buses, with many devices able to put power on or take power off at the same time. The electrical engineering and communications aspects of this may be substantial but could offer compelling efficiency and other benefits. A final challenge will be to ensure that sufficient end-use devices are available that implement LPD. The electronics industry shows that quick innovation and product introduction are possible when manufacturers see a viable market.

Summary

DC power, particularly with direct dc, has the potential to save energy, and offers many other benefits. Making this a reality requires developing new technology. Existing managed dc creates the possibility for new mechanisms to choose how to manage generation, storage, and end use, but because it is limited to a single power links, it has great limits in how it can be leveraged. Only networked dc greatly expands the degree of capability and utility of dc over ac, and LPD is a simple and powerful way to create this. Although direct dc will sometimes be introduced into buildings in single large installations, more commonly it may be added in small pieces over time, building an evolving network of power entities inside buildings.

For Further Reading

B. Nordman, "Local Grid Definitions," prepared for the Home-to-Grid Domain Expert Working Group (H2G DEWG), Smart Grid Interoperability Panel (SGIP), Jan. 5, 2016.

B. Nordman, K. Christensen, and A. Meier, "Think globally, distribute power locally: The promise of nanogrids," *IEEE Computer*, vol. 45, no. 9, pp. 89–91, Sept. 2012.

B. Nordman and K. Christensen, "DC local power distribution with microgrids and nanogrids," in *Proc. 1st Int. Conf. DC Microgrids*, Atlanta, GA, 2015, pp. 199–204.

B. Nordman and K. Christensen, "The need for communications to enable DC power to be successful," in *Proc. 1st Int. Conf. DC Microgrids*, Atlanta, GA, 2015, pp. 108–112.

USB Implementers Forum. (2012, July 5). Universal Serial Bus: Power Delivery Specification, Revision 1.0. [Online]. Available: <http://www.usb.org/developers/docs/>

K. Garbesi, V. Vossos, A. Sanstad, and G. Burch, "Optimizing energy savings from direct-DC in US residential buildings," Lawrence Berkeley Nat. Laboratory, Berkeley, CA, Rep. LBNL-5193E, 2011.

LAN/MAN Standards Committee. (2009, Oct. 2). IEEE Std. 802.3at-2009, DTE Power Enhancements. [Online]. Available: <http://www.ieee802.org/3/at/>

Biographies

Bruce Nordman (BNordman@LBL.gov) received his B.A. degree in architecture and his M.A. degree in energy and resources from the University of California, Berkeley, in 1984 and 1990, respectively. He is a research scientist in the Buildings Technology Department at Lawrence Berkeley National Laboratory. His research interests include energy use and efficiency of electronics and networks, low-power-mode energy use, user interfaces, energy policy, and power distribution technology. He has contributed to many technology standards organizations including the IEEE, the International Electrotechnical Commission, Ecma International, the Consumer Technology Association, the Internet Engineering Task Force, and Energy Star specifications and test procedures.

Ken Christensen (christen@cse.usf.edu) received his electrical and computer engineering B.S. degree from the University of Florida, Gainesville, in 1981 and his M.S. and Ph.D. degrees from North Carolina State University, Raleigh, in 1983 and 1991, respectively. He is a professor and associate chair of the Department of Computer Science and Engineering at the University of South Florida, Tampa. His research interest is in performance evaluation of computer networks. In the past 10 years, he has made significant contributions toward proxying for network connectivity and energy efficient Ethernet. This research has contributed to Ecma International, IEEE standards, and Energy Star specifications. From 1983 to 1995, he was employed at IBM Research Triangle Park in Durham, North Carolina, as an advisory engineer. He has written more than 100 journal and conference publications and holds 13 U.S. patents. He is a licensed professional engineer in the state of Florida, a member of the ACM and the ASEE, and a Senior Member of the IEEE.