

Annex 33F

(informative)

Power Price Index Supplemental Information

This Annex presents information about the intended use of the PSE Power Price Index element in the Power via MDI Measurements TLV.

33F.1 Purpose

The purpose of the Power Price Index (PPI) is to provide a simple, effective, and universal mechanism for a PSE to communicate to a PD the relative *availability* of power in the PSE domain so that the PD can change its behavior to best meet the needs of the entire system.

33F.2 Example PD operation

Figure 1 shows how a PD that is a light source might commonly respond to changes in the PPI; two operational modes are shown (full-on and dimmed). The light participates in other (application layer) protocols to help it to determine when it should be on at all, and what dimming or other changes in behavior it should engage at any particular time.

When power is available at a normal level (nominal PPI or lower) then operation is unchanged. As power is less available the light (PD) dims proportionately to reduce power use and help the PSE better balance supply and demand. When power is sufficiently scarce in the PSE domain so that the PPI reaches a sufficiently high level, the light turns itself off entirely. While many lights might exhibit behavior similar to this, lights designated as “emergency lights” might ignore the PPI entirely.

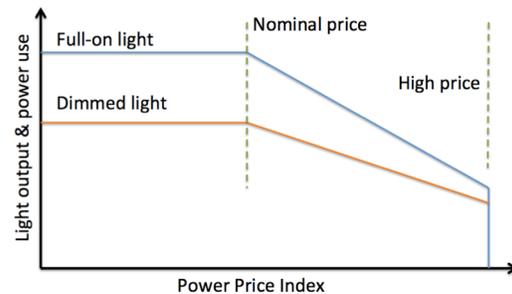


Figure 1. Lighting PD response to PPI

33F.3 Example PSE Operation

Three cases of PSE deployment illustrate how it can be effectively used. The first is a PSE that receives AC power directly from a utility grid. An increasing number of locations allow for, or require, the price of grid power to change over the course of a day, in predictable or unpredictable ways, to enable the utility to better balance supply and demand. The PSE has a prior understanding of a usual or normal price by other means, or by simply observing the grid price it receives over a period of time and averaging. It computes the PPI that is local to the PSE by dividing the current price by the nominal price, and multiplying by 1000.

The second simple case is a PSE that has local battery storage and is presently unconnected to a utility grid. When the battery is full, the PPI will be the nominal PPI multiplied by a factor to account for the round-trip efficiency loss through charging and discharging the battery. As the battery state of charge drops, the PSE will increase the PPI over time, taking into consideration the power draw of all PDs, any expectations on when a grid connection may be reinstated, and the number of hours of battery power availability the PSE seeks to maintain.

The third case could apply whether the PSE is grid-connected or not. As total power demand of PDs rises toward the electrical capacity limit of the PSE to deliver power, the PSE can increase the PPI to moderate demand to insure that the limit is respected. While an emergency priority ranking for when to make

unannounced removal of power from selected PDs is needed when other mechanisms fail and there is a sudden change in demand or in supply of power, accomplishing this balance is better carried out through voluntary changes in device operation as with changing the PPI. Later, as total PD demand drops, the PPI can be consequently reduced until it again reaches its base level.

33F.4 Theory of Operation

Figure 2 shows the domain of the PPI. The PPI expresses the availability of power within the PSE device, including any internal electricity storage. The PSE dictates the PPI to attached PDs. The PSE is a “nanogrid controller” (see references). The PSE also has power links to other nanogrid controllers, some of which may be bi-directional. These could be Ethernet or some other physical layer, though operating as an “uplink”, most PSEs may have electrical links of much higher power capacity than Ethernet presently provides for. Two parallel Ethernet links (one for each direction of power) can enable bi-directional power links between two devices; other physical layers enable switching the direction of power flow. Such links may be highly managed, or have little or no communication, as from the utility grid, or from local renewable (for example, photovoltaic) generation. Regardless, the nature of such other links and how they should be managed is outside the scope of the PPI.

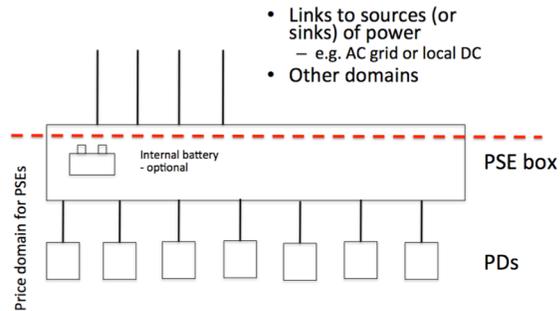


Figure 2. Theory of operation

Figure 3 shows schematically how nanogrid controllers can be deployed in a wider network, similar in many ways to how switches and routers are connected in IP data networks (note that power is not “routed”; all nanogrid controllers are comparable to switches). Each grey box is a nanogrid controller; the orange boxes are attached loads. Links can be made or broken at any time. The microgrid controller is only distinct in that it forms a boundary of the local grid network. Power links can be made to adjacent buildings.

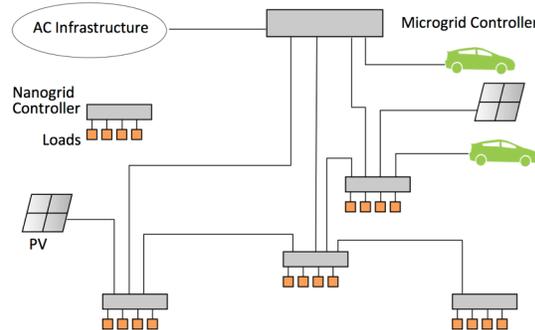


Figure 3. A network of Nanogrids

A notable feature of this architecture is that all communications about power are only across single links; there is no multi-hop power communication. This is highly beneficial for security, privacy, and simplicity.

33F.5 Intended Applications

The Power Price Index (PPI) is intended for use in *all* building types (residential, commercial, industrial, vehicles, and outside infrastructure), and *all* countries, in places with long-established utility grids, and in places with no utility grid. Networks of nanogrids may commonly also employ power links of physical layers other than Ethernet, with both lower and higher power levels.

33F.6 Other PD Types

Some devices can modulate operation over time to shift when consumption occurs; examples include those with thermal storage, such as refrigerators or HVAC systems, and those with internal batteries. A

refrigerator could raise its internal temperature setpoint when the PPI rises above the nominal value, and reduce it when it drops below it. An example of doing this (with a price forecast rather than only a single instantaneous value) can be found in (Nordman and Bugossi, 2014). Similarly, PD internal batteries can be discharged when the PPI is high and charged when the PPI is low.

33F.7 Implementation Notes

The PPI is disjoint from nearly all other aspects of PoE operation, and consequently has no interaction with any state machines defined by this standard. Determination of the PPI will generally use information from the voltage and current measurements available from the PSE or the PD via LLDP to determine power levels flowing across each link.

Algorithms used by PSEs and PDs are strictly internal to the individual device and need no standardization. PDs that use the PPI will likely be shipped with default behaviors that can be adjusted by an installer or user as needed. Behaviors can be static or dynamic based on other information the PD has from internal sources or from other protocols.

The PPI contains one bit of the 16 bits allocated that is reserved for an unspecified potential future use.

33F.9 References

Nordman, Bruce, Ken Christensen, and Alan Meier, “Think Globally, Distribute Power Locally”, *IEEE Computer (Green IT Column)*, September 2012.

Nordman, Bruce, and Mattia Bugossi, “Optimizing Device Operation with a Local Electricity Price”, *Micro-Energy International: Innovating Energy Access for Remote Areas: Discovering Untapped Resources*, Berkeley, CA, April 2014.

Nordman, Bruce, and Ken Christensen, “DC Local Power Distribution with Microgrids and Nanogrids,” *IEEE First International Conference on DC Microgrids*, Atlanta, GA, June 2015.

Nordman, Bruce, and Ken Christensen, “The Need for Communications to Enable DC Power to be Successful,” *IEEE First International Conference on DC Microgrids*, Atlanta, GA, June 2015.

Nordman, Bruce, “Local Grid Definitions”, prepared for Smart Grid Interoperability Panel, February 2016.