Network connectivity and low-power mode energy consumption

Bruce Nordman[1], Hans-Paul Siderius[2], Lloyd Harrington[3], Mark Ellis[4], Alan Meier[1]

[1]Lawrence Berkeley National Laboratory
[2]SenterNovem
[3]Energy Efficient Strategies
[4]Mark Ellis & Associates

ABSTRACT

Since 'standby' was recognized as an energy efficiency issue, a growing number of countries have put policies in place that reduce low power mode energy consumption. However, most of these policies, e.g. the EU Ecodesign regulation on standby and off modes, only target the most simple low power modes. Meanwhile, a rapidly increasing number of products have greater numbers and increasingly complex low power modes. The complexity arises from the fact that many products are already or will be in future connected to a network and will maintain a connection to the network when the product is not performing one of its main functions. Most standby definitions and levels do not take this complexity into account. As a result, only a few policies cover network-connected modes, and not in a consistent or comprehensive way.

This paper will review the state of the art regarding low power modes generally, network connectivity, and its relation to standby and other low power modes:

- recent developments in standby / low-power policy
- the "functional" approach to low-power modes
- how to test and regulate low-power mode consumption
- how networks affect low-power mode consumption
- existing efficiency policies and standards
- A horizontal approach to (network) low power modes
- future directions and guiding principles

The aim of the paper is to provide policy makers with guidelines for the development of effective policies to reduce power consumption in networked products and modes.

Introduction

Low power mode energy use has developed from an obscure issue 20 years ago to being widely recognized, and from no recognition within the policy environment to being common within many energy efficiency policies. An infrastructure has been built of ideas, test procedures, measurements, policies, and institutions. Initially, it was seen as a "mode" of a product, and used as a mode name on some products ("standby"). When officially defined, in IEC 62301 [9], standby became a power level ("minimum power level while connected to mains"[1]) that could occur in any operating mode, depending on the product. As it became clear that much low-power mode energy use was due to modes other than the minimum one, the term "standby" became used as a proxy for all low-power modes, and most recently, only covering those modes that are not designated as "off" or as having network connectivity. Modes with connectivity are now called "network modes" in a draft revision of IEC 62301 [8][2].

---

[1] Edition 1 of IEC62301 defined standby as the lowest power while connected to the mains. This was intended as an internal definition within the standard itself. The limitations of this definition are widely acknowledged and Edition 2, due in 2009 will have this definition removed.
[2] This to differentiate products with network functions from those within the traditional scope of "standby", in the interests of better clarity of understanding in future.

Past ways of addressing low power mode consumption are generally not adequate to deal with network issues, so low power mode policy needs additional evolution, along with new content that is network-specific. As an example, the EU Ecodesign process on standby did not incorporate network connectivity into its requirements (but has started a new process to address it) [3, 4].

In the first part of the paper, we review the "functional" approach to low power modes, and describe in detail how it is used in practice in testing and policy. The second part then builds on this by including network functionality and its many interesting dimensions. Finally, we have recommendations for research and policy.

The functional approach to low-power modes

Definition: list of functions

Recent years have seen increasing interest in grounding low-power modes into listings of "functions" present in any particular mode as a way of defining and evaluating their energy performance. Collections of functions have been previously listed in the EuP Lot 6 study, in the current draft of IEC 62301, in the EU Codes of Conduct on Broadband Equipment and Set-top Boxes [1, 2], and in several Energy Star specifications [14, 15, 16].

Ultimately, a mode is a collection of functions.

For many reasons, it is desirable to have a list and definitions of functions that spans all product categories, and can be applied globally. Harmonization of power levels associated with functions is desirable, but this may well vary across products and will definitely vary with time as functions evolve and technology allows them to become more efficient.

There are various ways to organize a list of functions, and Table 1 lists one such approach. The functions listed in Table 1 are certainly the majority of those of interest for low power modes in user-oriented products.

Table 1. A proposed taxonomy of low-power mode functions

Communication – Devices

- remote control – power only, 1-way, 2-way
- data connectivity – wired, wireless
- network connectivity – wired, wireless

Communication – People and Environment

- display – indicator, alphanumeric, graphic, audio, tactile;
  functional information / status, power state, time
- electronic controls and switches (energizing a keypad); power control
- sensors – occupancy, ambient light, internal

Time

- clock (keeping time), timer (tracking relative time, switching), and schedules

Power

- EMC filters; ability to power-up on timer; battery-related
- power source (ambient, battery, AC, DC);
  ability to power other devices (e.g. via USB)

Functions required and functions present

The functional approach is a critical development, but needs one additional feature to reach its full potential, which clarifies a core duality of functions. In casual language this can be described as "**what you want** vs. **what you get**". It is easily explained in the context of a test (and a test procedure). A specific test of a particular project will identify a number of specific modes to be evaluated. For each mode, there will be a list of functions that **must** be present in the product during the test. Since a product may have many more functions than distinct modes (and even more combinations of functions than modes), many collections of functions will not have an exact incarnation in the product being tested with *only* those functions present.

For example, a microwave oven being tested has a display and an internal light (that comes on when the door is open). Any time the light is on, the display is on. The test of the mode with the light being on only requested that function, but the display is included in the mode as it occurs. So, depending on the product design, functions of direct interest may be always bundled with other functions of indirect or no interest. This example may seem trivial, but the principle is important: each measurement of a mode has two sets of functions: those required, and those actually present when tested. Any measurement or regulation needs to be clear which is being addressed.

An example closer to the core of this paper is the presence of network connectivity as a function. Some products have network connectivity in all low-power modes; that is, so long as it is connected to power, you always get that connectivity, with the functionality and energy use that implies.

Testing low power modes

A test procedure based on functions must list all required functions for a particular mode, annotated in some cases with standard conditions such as expected network speeds, traffic levels, etc. Some modes could list "prohibited" functions[3]. For many products, there may be a need for some configuration of the product prior to testing to accomplish particular modes (or reflect the intention of the test). Details of this configuration (and the as-shipped conditions as well) need to be carefully reported. An example would be to enable Wi-Fi connectivity during sleep mode (which might be disabled as the default).

As products increasingly alter their energy-using behavior in response to environmental conditions, these will need to be specified, and established for tests for these products. Examples include products with light sensors that adjust display brightness, that change behaviour depending on assessment of room occupancy, or have fans that are thermostatically controlled (and so affected by the ambient room temperature).

Thus, the procedure that a test laboratory would follow is:

- Begin with a product as shipped
- Review a list of required functions from the test procedure (or specification) for each mode of interest
- Modify the configuration and setup as per required functions
- Measure power for each relevant mode and combination of functions
- Report actual functions present in each mode (and level of activity if relevant)

Beginning with a product as-shipped is important to have an unambiguous test procedure, and as many people may never change many or any of the settings. The order of testing can be important; for example, some products enter a different mode when entering their off mode via an automatic power-down feature than they do when done with a manual power control. For testing with this scheme, the forthcoming second edition of IEC 62301 is sound and does not need particular alteration, other than being supplemented with additional reporting, and possibly standard terminology to refer to particular types of functions.

---

[3] Whether any prohibitions are needed is not yet known, though one can imagine a product "cheating", as by making strategic use of stored battery power to distort a test result.

Regulating low power modes

This testing method is designed to facilitate voluntary and mandatory limits ("regulation" here used to cover both) on the energy use of products in low-power modes.

Program requirements can address "any mode with *X* characteristics and not *Y* characteristics" and can require that a product "shall have such a mode with a power level of ≤ *x* watts" or "shall have *all* low power modes require ≤ *x* watts"[4]. This is best understood in the context of policy collections of horizontal and vertical efficiency standards [5]. A vertical standard is one specific to a single product type (or collection of like ones). A horizontal standard is one that covers a particular function, set of functions, or characteristic across many or all product types (e.g. power supplies, low power mode consumption, network characteristics, battery characteristics, user interfaces, etc.) [11].

There are two basic methods of setting thresholds on consumption: modal (power) and annual (energy). The **modal** approach specifies one or more low-power modes that are measured and establishes individual limits for each mode. The **annual** approach (or any convenient period of time) sums up the consumption of all modes deemed significant, with an operating pattern that is seen as representative of typical consumption and limits *only* the sum, *not* the components. The annual total often also includes active consumption and is sometimes referred to as a 'duty cycle' or 'typical energy consumption' approach. Both of these have their advantages and for the foreseeable future, both will have good uses for specific product types. The functional approach with adders applies equally well to both.

For a particular mode, a regulation can specify a base power level with additional allowances ("adders") provided for functions beyond the assumed base for the mode. Adders have been widely used in the Energy Star imaging and set-top box specifications, and the Code of Conduct DIgital TV Services. More recently it has been used in the Energy Star computer specification, and the Code of Conduct on Broadband Equipment.

Some view such adders as a way for manufacturers of products that shouldn't qualify for a limit or label to do so by loading up with many spurious functions. This could work if the adder values are too large, but if they are set to the minimum level required for the functionality (and no more) this problem does not arise. A regulator can place a limit on which allowances are valid for a particular mode, and how many can be used[5].

Most efficiency test procedures and policies are limited to AC-powered products. However, there are many reasons to include low-power DC products in these as they can have efficiency (and other) advantages and should be allowed to compete on an even basis with those powered by traditional AC sources.

Network low-power modes

Introduction; general definition

The first step in addressing network modes is to define what actually constitutes a network and a network connection.

A network connection is a digital connection that allows exchange of digital[6] data among a set of (more than 2) devices. Devices connected by the Internet Protocol are the most common and obvious example of those with network connections, but there are other digital networks (e.g. HDMI). In addition to "true" network connections, there are also simple digital data communications

---

[4] Special consideration is needed to consider the possibility that a product being tested has no mode of the type being specified so that how these products treatment is clear.

[5] This doesn't preclude a manufacturer from including additional functions, but there will not be additional power for doing so; if there are many such functions, their power levels have to be very low to compensate. Also, in setting adders economies of scale need to be considered, i.e. an adder being additional (on top of other adders) might be lower than a single adder.

[6] In contrast to analogue electronic communications mechanisms that have existed for well over 100 years.

mechanisms that connect only two devices. Examples are the VESA connections between a computer and its display, USB links, and the (one way) infrared communication from a remote control to a TV.

Many network technologies are based on simple data links at lower layers, but higher layers enable communications to move transparently across many links to the destination device. In this paper, the design is for network connections, but the conclusions generally also apply to all types of data links so they are implicitly included in “networks”, including in Table 1. We ignore analogue connections, as they are diminishing as a determinant of energy use.

The IEC 62301 draft refers to a “network function” as “reactivation via network command or network integrity communication”. The latter is communication that is essential to having the network connection continue (lack of routine traffic is a way for networks to flush out devices that are no longer present or responding). This in line with the notion of “standby” as “readiness to act”. This definition includes data connections under the general rubric of true network connections, as does EuP Lot 6.

States of networked devices

Some network connections are also qualified by values such as speed, in addition to the core physical layer technology in use. Ultimately, any test procedure or regulation has to be informed by an assessment of what combinations of functions are particularly useful or occur in typical use. Typical or normal use is particularly important – specifying configurations or combinations of functions that are never used in practice is a dubious basis for a regulation or other requirement. The other danger is the stifling of innovation; if broad mode requirements are stringent, suppliers may be forced to eliminate higher levels of functionality from their default settings. Of course the ultimate objective of any efficiency policy is to minimise energy consumption of products while maintaining acceptable functionality.

Network connectivity is not a single function, with a single impact on power and energy use. Rather, we need to deal with a variety of potential conditions that can occur that have inherent implications for energy use, dictate potentials for functionality, and are measureable. These all can depend on the nature of the physical layer interface in use, the status of the connection, and the quantity and nature of any data flowing across the link. Mere data links are almost always much simpler; for example, the infrared sensor (receiver) on a television for a conventional remote control has a power requirement that is essentially constant, and low.

While most functions are either present or not in a mode, for network connections, there are three levels of “presence” that can be distinguished:

- the function **exists** (but is unavailable as configured)
- the function is **available** (but not active)
- the function is **active**

The following discussion refers to a single network connection, though a device may have multiple network connections with the same or different physical layers. A network interface can be in one of several states as follows:

Disabled

The first possibility is that a device has the **capability** to be network connected but the feature is disabled through configuration. In principle, all electronics associated with the connection might be removed from power, though in practice, some power may be expended by the interface or related electronics (including power supply losses). Thus, even a disabled capability may require some power. In the language of functions, a disabled interface only *exists*.

Absent - Wired

For a wired network (e.g. Ethernet), there are several possible paths to no data connection to another device:

- the interface might be enabled but no cable connected;
- a cable might be connected with no device at the other end; or
- a device might be at the other end but not enabled or turned on.

These cases of absent connection are not likely a central focus of standards developers and product designers (who quite naturally focus on designing interfaces and products to actually be in use). Nevertheless, these non-functional conditions are common in practice, particularly for products that have network connections only intermittently, are portable devices (e.g. notebooks PCs), or have multiple network connections (and so may have several commonly inactive even if most of the time at least one is active). An absent connection implies that it is *available*.

Absent - Wireless

For wireless networks, no cable is involved, but that does not mean the end of complications. With an enabled wireless interface, it may find or not find one or more other devices (particularly access points) to connect to, and may be able to make a connection to it or not depending on physical conditions and security limitations. Some types of wireless interfaces expend more power trying to establish a connection than they do to maintain a connection. Since wireless devices commonly can move, there is an inherent dynamism in wireless connections not present in wired ones. As with wired, an absent connection is *available*.

Linked

The purpose of any wired or wireless interface, is to actually establish a link to other devices. Links may be capable of several different operational modes that affect the amount of power required, including different throughput capabilities (e.g. Ethernet which supports speeds that vary by three orders of magnitude and the various “flavours” of Wi-Fi), the length of the data link, or the ability of a link to go to sleep for short periods of time when utilization is low. The ability of the device being tested to enter these speeds or modes can depend on the capabilities of the device to which it is connected. Signal strength and interference can also be important and can dictate the connection speed. A linked connection is *active*, regardless of how much data is flowing across the link (even none).

Full Connectivity

A link can be accomplished with fairly limited communication, but maintaining the device as connected to the network in a more general sense involves much more, including some network infrastructure activity, as well as maintaining presence in an application sense. This requires some additional power (energy) over just maintaining a link — sometimes much more. This is associated more with active consumption.

Special Modes

There are special modes in certain networks that provide for sending a wake-up signal but not ordinary connectivity. Wake-On-LAN is the most widely known of these. The EU Code of Conduct for Broadband references a DSL state in which the interfaces are only looking for wake-up signal. Some sensor network technologies use a “wakeup radio” signal (that can be listened to at very low power) that is separate from the one used to actually transmit data. Other modes useful when a device is relatively inactive are those that introduce latencies in communication, based on knowledge that these are acceptable in the usage context. These modes can be useful and may be a way to leverage significant savings. Thus, they need to be accommodated and recognized by test procedures.

Defining network modes

So, what **is** a network mode? One possible definition is any mode in which a network characteristic changes how much energy or power the device should reasonably use. Whether this include the disabled or absent cases described above is a question of technology and policy. The proposed revision to IEC 62301 defines “Network Mode(s)” as including “any product modes where the energy using product is connected to a mains power source and at least one network function is activated (such as reactivation via network command or network integrity communication)”.

A final aspect of network technology that is unique among energy-using equipment is that the behaviour of one device affect energy consumption of others, either directly, in the nature of the physical layer link between the devices, or more indirectly, in the content of the data sent among them. This interdependency means that efficiency standards need to take into account these behaviours independently of how energy consumption of the product itself is assessed. This also has implications for technology research and development.

Power (energy) impacts of networks

Network low power modes must be considered in the larger context of how networks affect energy use. Networks increase energy consumption in two ways. The first is *direct* consumption of network interface circuitry, and network devices. The second is the *induced* consumption of devices that are in a higher power mode by virtue of being network-connected. PCs and set-top boxes are the most notable examples of this currently. Often devices will wait in an active state for long periods of time on the chance that some useful network communication may occur. This is the most important consequence of network-related energy consumption and the one that requires the closest attention.

Today ‘network mode’ is relatively small, but it is poised to grow very rapidly over the coming decade as networked products become much more prevalent. Network mode occurs in several types of products.

Products that are already widely or universally networked. These include PCs and set-top boxes, and today they spend little time in a sleep mode, and for PCs the sleep mode lacks real network connectivity. Thus, as a percent of total energy for these products, the network modes are small. In many cases, these devices are left fully on to provide for full network connectivity so that going to sleep is not considered viable. This active mode consumption is a shortcoming of existing technology and effective energy management could result in a large energy savings. We expect that technology advances will allow for these types of devices in future to spend most of their time asleep, greatly increasing time in this mode.

Products that today are rarely digitally networked but likely will be in future. These include many displays, other audio/video devices, appliances, climate controls, security devices, and lighting. Some devices are relatively new, such as digital picture frames, but may usually be networked in the future. In some cases, such as displays or A/V devices, the network connectivity will be core to the product’s function, with it replacing more limited data link or analogue connections. For others, the network capability exports the user interface to devices that are more convenient to use or in a different room or part of the world. For most of these, it is not convenient to manually power the device on and off to engage or disengage the network connectivity, so that the default behaviour will be to toggle between on and sleep, resulting in large amounts of network mode time.

The rise of network mode consumption is not all bad. Large savings can result as time in low power modes replaces large amounts of full-on time. The investment in power required to maintain connectivity may enable devices (e.g. lights) to better match their service delivery to people’s needs. On the other hand, the *possibility* of savings does not mean their inevitability; it is quite possible that pervasive networking of buildings will lead to substantial increase in consumption rather than the reverse unless strong efforts are made to fundamentally build in energy management into networked product behaviour.

Network technologies have understandably focused primarily on functionality – mostly speed and to some extent versatility and security. The associated issue of energy consumption has received little attention in terms of enabling power management capabilities in network protocols. Power consumption due to networks is strongly affected by the throughput capacity of data links and network devices. The requirements for data flow vary by many orders of magnitude among devices, with video streams requiring the most, and appliances the least. When more capacity is provisioned than is needed, energy is wasted. Some network technologies allow for speed changes, but often only when the connection is initially established. An example of what is needed in general is the IEEE

802.3az process on Energy Efficient Ethernet which changes the paradigm to allow power consumption to scale with data throughput rather than capacity (see: http://ieee802.org/3/az).

Existing policies for network modes

Policy is taken here in a broad sense, including mandatory and voluntary energy standards, policy-driven technology development, and test procedures.

The Energy Star program has dealt with network issues for over ten years, with early versions of the computer specification providing for different sleep power levels depending on the nature of the connectivity provided during sleep. In 2007, the version 1.0 imaging specification (printers, copiers, etc.) went into effect which provided for a range of different allowances for network connectivity in sleep depending on the speed of the physical layer technology, and whether the interface was active or not during the test. As printers often have a variety of connection types present (data and network), sometimes multiple, this scheme allowed the manufacturer flexibility in which were connected so long as at least one was and at most three.

The most important aspect of the V5.0 computer specification is the inclusion of “proxying” as a capability that makes it easier to qualify as Energy Star. This technology increases the functionality of the device when it is asleep so that many people can use sleep instead of leaving it fully on [12]. It also provides new functionality to people who had previously powered down their computer when not in active use. This may require small increases in power consumption in sleep, but enabled avoiding large amounts of on time. This type of more functional low-power mode poses new challenges to policy in this area.

The EuP process on standby specifically excluded “network standby” from its scope. However, a new study on network standby is being undertaken in 2009.

As noted previously the EU Codes of Conduct on Broadband Equipment and Digital TV Services include provision for extra power for additional network functionality of products, and specify some network-specific test conditions.

The IEA Implementing Agreement on Efficient Electrical End-Use Equipment (4E) will address network connectivity and low-power modes through work within its Standby Power Annex [7]. This international project will engage stakeholders from the electronics industry, technology standards organizations, and governments to identify problems and policy solutions to many of the questions raised in this paper.

A horizontal approach to (network) low power modes

The infrastructure for policy in this area relies on network technology, test procedures, and product energy requirements. A key need is a single *horizontal* approach to measuring and regulating low-power modes.

In support of efficiency policies, the first need is a catalogue of the power requirements for different types of network interfaces, both as shown from measurements of actual devices, and as indicated by the requirements of the technologies and standards. This will provide the best basis for efficiency programs.

A workable functional adder approach applicable to all or most products needs to be developed for network functions. This would be derived from detailed assessments of the energy requirements for each of the main network related functions

For testing, a common international reference of standard conditions to apply to network connections will likely be necessary and desirable, to minimize the amount specified in individual procedures and to facilitate timely updating with new technologies. There will also likely be some reporting necessary of additional details of network conditions in testing a product to ensure that it can be reproduced and understood.

In choosing test details when there are many options, there is a tension among the goals of conditions which:

- are typical (which may be changing over time);
- fully exercise the product (e.g. use all ports); and
- provide simple comparability (and so use a minimum of ports).

The number and speed of ports connected for a test are the most immediately apparent detail to specify, but some others can be relevant, such as the distance from a product to the other end of the network link for wired or wireless interfaces (as some can modulate their transmit power to just what is needed). In future, versatility and ability to change connection speed may need to be assessed as well. As an example, the EU Code of Conduct for Broadband Equipment states that for testing, the cable on an Ethernet port should be 5m in length, and makes even more specific requirements for testing standard telephone line equipment.

For *vertical* standards (specific to one product type), many can adopt language and requirements from the horizontal standard to address network issues in the various modes and testing they cover.

The “minimum power mode while connected to mains” concept is still useful for testing and policy, and so should be retained. A good name for this is “minimum power mode”.

Guiding Principles and Conclusions

Guiding Principles

Managing energy consequences of network connectivity requires an overall approach to technology and policy. These were summarized for a workshop in 2007 as “Guiding Principles” of behaviour that underpin networks [10] and are reproduced below (a smaller set of principles are in the EU Code of Conduct on Broadband Equipment.

To ensure that digital networks and network-connected devices support the minimisation of direct and induced energy consumption, the following the principles should be adopted:

Digital Networks:

- All network technologies should actively support power management.
- Connection to a network should not impede a device from power management activities.
- The network should be designed such that a legacy or incompatible device does not prevent the rest of the network from effective power management.
- Connections should have the ability to modulate their own energy use in response to
- the amount of the service required by the system.
- Terminology and concepts relating to energy management used in the design of all networks should be internationally harmonised.

Network connected devices:

- Devices should not impede power management activities in other connected devices.
- Devices should expose their own power state to the network and be able to report estimated or actual energy use.
- User interfaces should follow (international) energy management standard principles and designs.

- Devices and connections should have the ability to modulate their own energy use in response to the amount of the service required by the system.
- Terminology and concepts relating to energy management used in the design of all devices should be internationally harmonised.
- The behaviour and communication of devices relevant to energy consumption should adhere to (international) standards.

Energy Efficiency Policy:

- Governments should ensure that electronic devices enter low-power modes automatically after a reasonable period when not being used.
- Governments should ensure that network-connected electronic devices minimise energy consumption, with a priority placed on the establishment of industry-wide protocols for power management.
- Energy efficiency efforts should not favour any particular hardware or software technology.
- Energy efficiency policy should identify digital networks as a promising method for attaining energy efficiency.

Conclusions

Network connectivity is an increasing factor driving energy use in low-power modes (and other modes as well). It poses unique challenges to measurement and policy, and while this will take time to work through, the tasks are not insurmountable, and some principles and the immediate path forward are clear. A key need for networks is one that is extremely helpful for low power modes generally, namely the functional approach to testing and regulation. For networks, it will be key to having common test conditions and methods that can be incorporated into a global horizontal approach.

References

[1] European Commission 2008a, Code of Conduct on Energy Consumption of Broadband Communication Equipment, Version 3, 18 November 2008. http://re.jrc.ec.europa.eu/energyefficiency/

[2] European Commission 2008b, Code of Conduct on Energy Consumption of Digital TV Services, Version 7, 18 November 2008. http://re.jrc.ec.europa.eu/energyefficiency/

[3] Fraunhofer, 2006. Draft Definition Document for Standby and Off-mode Losses (Lot 6, Task 1). Initial report dated 30 August 2006 under the European Commission Eco-design Directive contract. Available from www.ecostandby.org

[4] Fraunhofer, 2007. Preparatory Study Lot 6 “standby and off-mode losses” – report on task 1 (definition), task 2 (market data) and task 3 (consumer behaviour and local infrastructure). Interim reports dated 10 January 2007 under the European Commission Eco-design Directive contract. Available from www.ecostandby.org

[5] Harrington, Lloyd, Jack Brown, Shane Holt, Alan Meier, Bruce Nordman, Mark Ellis, Standby Energy: building a coherent international policy framework – moving to the next level, European Council for an Energy Efficient Economy Summer Study, June 2007.

[6] Harrington, Lloyd, Hans-Paul Siderius, Mark Ellis, “Standby Power: Building a Coherent International Policy Framework”, ACEEE Summer Study on Energy Efficiency in Buildings, 2008.

[7] IEA, Implementing Agreement on Efficient Electrical End-Use Equipment (4E): Programme of Work, January 2008. http://www.iea-4e.org/

[8] IEC, Measurement of Standby Power (IEC 62301), (IEC 59/423/CD – 2008), 2008.

[9] IEC, Measurement of Standby Power (IEC 62301), Edition 1.0, 2005.

[10] Meier, Alan, Bruce Nordman, and Mark Ellis, Buildings as Networks: Danger, Opportunity and Guiding Principles for Energy Efficiency, September 2007 http://www.iea.org/Textbase/work/workshopdetail.asp?WS_ID=285

[11] Nordman, 2006. The Way Forward: A New Standby Framework. Paper presented to the international standby power conference, Canberra, Australia, 6-7 November 2006. http://www.energyrating.gov.au/forums-2006-standby.html

[12] Nordman, Bruce, and Ken Christensen, "Improving the Energy Efficiency of Ethernet-Connected Devices: A Proposal for Proxying," White Paper, Version 1.0, Ethernet Alliance, October 2007. http://ethernetalliance.org/library/white-papers.html

[13] Nordman, Bruce, Alan Meier, Mark Ellis, Draft Principles for Energy Efficient Digital Networks and Network-connected Devices, International Workshop on Energy Efficient Set-top Boxes and Digital Networks, International Energy Agency, July 6, 2007. http://www.iea.org/Textbase/work/workshopdetail.asp?WS_ID=285

[14] U.S. EPA, ENERGY STAR® Program Requirements for Imaging Equipment, Version 1.0, 2006. http://www.energystar.gov/ia/partners/product_specs/program_reqs/Prog_Req.pdf

[15] U.S. EPA, ENERGY STAR® Program Requirements for Computers, Version 4.0, 2006. http://www.energystar.gov/ia/partners/product_specs/program_reqs/Computer_Spec_Final.pdf

[16] U.S. EPA, ENERGY STAR® Program Requirements for Set-top Boxes, Version 2.0, 2008, http://www.energystar.gov/ia/partners/prod_development/revisions/downloads/settop_boxes/Set-top_Boxes_Spec.pdf