xergyconsulting

Building energy solutions

FINAL REPORT:
POWER REQUIREMENTS FOR FUNCTIONS

Prepared for: International Energy Agency 4E Standby Power Annex

Submitted by: Peter May-Ostendorp, PhD, LEED AP (Xergy Consulting)
Katherine Dayem, PhD
Jeffrey Swofford (Ecova)
David Thomsen (Ecova)

September 8, 2013

EXECUTIVE SUMMARY

Consumer electronics and appliances are increasingly designed and produced to provide the user with amenities that go beyond the primary purpose of the device. A smart refrigerator, for example, not only keeps food cold, but also may be networked to send information to a user's smart phone or equipped with an informational display for the user to monitor refrigerator contents. With this increased functionality comes increased standby energy use. Because a number of so-called secondary functions — networking, information display, and power conversion, to name a few — are common to such a large number of products, cross-cutting or horizontal energy efficiency policies have the potential to encourage lower standby power requirements by addressing these common "building blocks" of consumer products. Although much attention has been focused on total standby power for a device, standby power related to secondary functions is largely undocumented. This report, following on previous IEA 4E Annex work entitled *Mapping Secondary Product Functions to Products and Operational Modes* (May-Ostendorp et al., 2012), is a first look at the power requirements and horizontal policy opportunities for a selection of high-priority secondary functions.

Based on the previous report, the 4E Standby Power Annex identified four topic areas for study: networking (including high-speed wired, high-speed wireless, and low-speed wireless communication), informational displays, infrared (IR) sensing, and power supplies. For each function, we present a technology summary and trends, power estimates, and energy savings opportunities. Savings opportunities generally fall into two categories: power scaling (e.g., scaling power and performance to the required activity, such as data transfer speed in network functions) and power management (e.g., reducing or eliminating power when the function is not in use).

High-speed wired network communication, dominated by Ethernet, is one of the most prevalent secondary functions identified in the previous report. Active power draw of Ethernet chipsets increases with data transmission rates and can vary by three times or more depending on the components used. As the market moves toward devices with higher link rates, mitigating increased power draw becomes more important. This is well managed with Energy Efficient Ethernet (EEE) that effectively scales power to the amount of data transferred. To accomplish this power scaling, however, both sides of a connection (e.g., an edge device and a router) must be EEE-capable and have the feature enabled. Horizontal policy should require EEE installed and enabled on new end use products.

Wi-Fi, the dominant high-speed wireless network protocol, is similar to Ethernet in that power draw increases with link rate, and link rates continue to increase in the latest wireless networking hardware. Energy savings opportunities exist in both minimizing power while data is transferred and in reducing power when no data is transferred, although as with Ethernet, the greatest savings opportunity is to simply put the Wi-Fi chipsets into a low-power state when there is no network traffic. Existing Wi-Fi networking standards provide support for this type of power management at the component level, but policy must ensure that it implemented in end use products by setting standby levels that are sufficiently aggressive.

Many low-speed wireless communication protocols exist, and unlike for high-speed communication, no one protocol dominates the market. Here we focus on Zigbee, developed specifically for low-speed wired communication, as the most likely candidate for this secondary function. Current Zigbee chipsets draw on the order of tens of milliwatts during data transfer and fractions of a milliwatt when in standby. For applications that require only low-speed communication, such as in smart appliances or energy monitoring applications, Zigbee is a far more appropriate and efficient choice than Wi-Fi. However, many manufacturers are basing smart appliance communications on Wi-Fi because so many homes and businesses already operate Wi-Fi networks. Policies should encourage use of low power, low-speed networking when it is a sufficient tool for the job at hand.

Active mode power requirements for displays vary by display size, resolution, and technology. Yet, when a display is not being used, its power can be reduced to essentially zero. In secondary displays, therefore, we expect the largest savings opportunities to be related to cutting power to the display when it is not in use. When the display is in use, efficiency improvements can be achieved by power scaling, using more efficient technologies, and using the right technology for the content displayed. Thus, horizontal policies should first encourage powering down secondary displays when the user does not need them, and secondarily encourage reducing power when the display is in use.

We touch briefly on IR communication, a mature technology that is likely to remain prevalent in remote controls for home entertainment systems. Energy savings in IR communication devices exist primarily in the receiver, and can be achieved with power management techniques that allow most of the microprocessor to power down when the device is not receiving an input signal.

Finally, even the best power management and power scaling techniques are moot without a properly designed power supply. In most cases, a separate low voltage power supply is required to deliver standby power efficiently. In an illustrative example, we show that the use of a dedicated standby power supply can reduce losses in the ac-dc conversion stage by about 80% and cut overall ac standby power by more than half compared to a design employing only a high voltage main power supply.

In summary, two broad techniques reduce power draw of secondary functions: power management to disable activity when not needed, and power scaling to actively adapt power draw to the device's needs at any one time. Given the increases in the prevalence of network connectivity and connected standby, network speeds, and the use of secondary displays, power required by secondary functions will continue to grow unless mitigated by, among other measures, horizontal policies that ensure the functions employ best-available technologies and are aggressively power-managed.

However, horizontal policies alone are not a sufficient means of regulating standby power. Much of the orchestration of secondary functions depends on processing capabilities that tend to be application-specific. Thus, we still see the need for vertical policies to address the device-specific implementation issues that inevitably result from orchestrating secondary functions— even for standby and other low-power modes. These policies can be vastly simplified when built on a foundation of horizontal secondary function power requirements like those provided in this report.

ACKNOWLEDGMENTS

The authors would like to thank several individuals for their willingness to share information, particularly on the latest technology developments impacting standby power requirements. Our research would not have been complete without their valuable participation.

- Vojin Zivojnovic (AGGIOS)
- Davorin Mista (AGGIOS)
- Prasan Pai (Broadcom Corporation)
- Philip Walters (East Canyon Services)
- Ben Eckermann (Freescale Semiconductor)
- Eric Mann (Intel Corporation)
- Bruce Nordman (Lawrence Berkeley National Laboratory)
- Steven Lanzisera (Lawrence Berkeley National Laboratory)
- Miji Yang (Realtek Semiconductor Corporation)
- Mike Yeh (Realtek Semiconductor Corporation)
- Thomas Bolioli (Terra Novum)
- Evan Davis (TRENDnet)
- Mike Hack (Universal Display)

TABLE OF CONTENTS

LIST OF FIGURES

LIST OF TABLES

1 INTRODUCTION

Standby power is a general term used to describe the energy consumed during low-power mode(s) that many electrical and electronic products are in when not performing their primary function (Harrington and Nordman, 2010). Standby power has become more complex in recent years, a trend likely to continue into the future. Not only are the number of possible functions and operating modes increasing, but the number of products with significant functionality during low-power modes is also increasing. The presence of electronic controls in most products — the buttons on a laundry machine or the touch screen on a multi-function printer — now provides a platform for the proliferation of these functions, which can provide users with new and useful features when the product is not in active mode. There also may be a range of functions present of which the user is not aware. The appearance of network-related modes in a wider range of products and the expansion of home networks is an important component of this trend. Many sophisticated network functions that are typically associated with computers and networking equipment are now expanding to a wider variety of loads — ranging from traditional white goods like laundry equipment to screw-based lamps and lamp sockets to home entertainment equipment — in attempts to make products smarter, more interconnected, and compatible with smart grid and home automation technologies. Each of these so-called “secondary” product functions contributes to power consumption in standby mode (or other low-power modes).

In 2012, the International Energy Agency’s (IEA) 4E Standby Power Annex commissioned Ecova to conduct an initial investigation of secondary functions that may have an impact on low-power modes. In a report entitled *Mapping Secondary Product Functions to Products and Operational Modes* (May-Ostendorp et al., 2012), researchers:

- Examined the variety of secondary functions available in common consumer products (appliances, home entertainment products, and office products), and assigned those functions to several broad categories;
- Conducted market research to “map” product functions based on their prevalence in major product categories like televisions, white goods, computers, and printers;
- Through further market research and review of product specifications, mapped secondary functions to major operational modes in key products, such as “standby” or “idle” (operational modes vary by product category); and
- Recommended high-priority secondary functions with large “horizontal” applicability — that is, those functions that are prevalent in a large number of consumer goods, that are likely to be operational during low-power modes, and that could potentially have measurable impacts on standby power.

This report builds on the previous work examining in greater detail several high-priority secondary functions to further quantify their power requirements and provide the technical background that might be required to advance horizontal policy action associated with these functions. We do not provide a comprehensive summary of the previous work or detailed justifications for the chosen functions. For detailed information about the framework for secondary functions and a ranking of the most horizontally applicable functions, we direct the reader to the prior report, *Mapping Product Functions to Products and Operational Modes*.

1.1 KEY HORIZONTAL FUNCTIONS

In conjunction with the IEA's 4E Annex, we prioritized four broad topic areas for further investigation: networking, displays, IR sensing, and power supplies. Below, we introduce the functions residing under these categories and provide examples of the technologies that deliver these services.

1.1.1 Networking

Prior secondary function market research identified three key networking functions for further exploration: high-speed wired networking, high-speed wireless networking, and low-speed wireless networking. In this research, the three functions effectively translate into three corresponding, dominant technologies.

In the high-speed wired networking space, Ethernet (IEEE 802.3 standard) dominates. In networking equipment and certain consumer electronics like set-top boxes, other protocols such as DOCSIS, DSL, or Fiber Channel do exist. However, these networking technologies provide primary functionality because they are part of dedicated communications equipment. Since this work focuses on addressing horizontally applicable secondary functions, this report does not cover these other protocols.

For high-speed wireless networking, Wi-Fi (IEEE 802.11 standard) is the dominant technology used in consumer and office products. Many mobile consumer products, such as tablets and smart phones, integrate other wireless networking technologies to connect to wide area data networks, but these alternative technologies are not expected to have a significant foothold in stationary products where standby power is of greater concern.

Although not extremely prevalent in today's end uses, low-speed wireless networking, such as Zigbee (IEEE 802.15.4) and Z-Wave may play a significant role in monitoring and control applications, where high bandwidth is not a strict requirement. Zigbee components are currently offered by several prominent electronic component manufacturers, including Texas Instruments, Analog Devices, Freescale, and Ember. Z-Wave products, on the other hand, are only offered by Sigma Designs and Mitsumi.

No one technology dominates the low-speed networking space. Arguably, any number of wireless networking protocols — Zigbee, Z-Wave and Bluetooth — could provide monitoring and control services. Some might even argue that Wi-Fi networking is a good candidate itself, given its incumbency in homes and businesses as a wireless networking solution. Section 4 of this report addresses the challenges posed by the emerging smart/connected devices market and helps to chart several policy priorities given the uncertain evolution of this space.

1.1.2 Displays

In previous work, we also identified information displays as a secondary function for further investigation (May-Ostendorp et al., 2012). Displays are already common — and in some cases must be present — in a wide range of products, including appliances, home audio/video equipment, smart phones, cars, and even industrial and medical applications. We expect this proliferation to continue as more functionality is built into appliances and the mobile device market continues its growth.

Opportunities for saving energy consumed by secondary displays exist for both active and standby operating modes. Efficiency improvements in active mode can be achieved by power scaling, using more efficient technologies, and using the right display technology for the content displayed. Secondary display standby power is on the order of tens of mW or less, so energy savings are less a function of reducing standby power itself and more a matter of improving power management so that a display shows information only when the user needs it.

1.1.3 IR Sensing

Infrared (IR) remote controls ushered in the concept of standby power in consumer electronics. Prior to remotes, most electronics could be fully powered down with hard off switches. Using a remote control to power on electronics required some amount of standby power to keep the device in a ready state, receive incoming remote signals, and eventually power on. Remote controls and related short range wireless controls have since proliferated. Although radio frequency (RF) communication protocols such as Bluetooth and Zigbee RF4CE are likely to work their way into devices that have traditionally used IR communication, the mature IR technology seems likely to remain prevalent due to its low cost and familiarity. Energy savings in IR communication devices exist primarily in the receiver, and can be achieved with power management techniques that allow most of the microprocessor to power down when the device is not receiving an input signal. In this report, we provide a snapshot of IR sensing rather than a detailed quantification, as presented in the networking and displays categories.

1.1.4 Power Supplies

Power supplies convert high-voltage mains electricity down to lower dc voltages that sensitive electronic components require to operate. As a result, power supplies continue to have far-reaching energy impacts because they are present in all electronic devices. Due to the broader use of displays, electronic controls, and networking in traditional appliances like refrigerators and laundry equipment, power supplies and standby power have expanded their foothold in these end uses as well. This report provides an overview of appropriate standby power supply design and illustrates the impact that poor design can have on standby power values.

2 POWER REQUIREMENTS FOR HIGH-SPEED WIRED NETWORKING

The energy efficiency community has studied Ethernet technology more extensively than other networking technologies in this report, in large part through the efforts of the IEA's 4E Standby Annex (Harrington and Nordman, 2010). Due to its widespread use in enterprise networking and data centers, the technology has also received significant attention from standards organizations seeking to reduce Ethernet's overall energy impact, primarily through the IEEE 802.3az task group on Energy-Efficient Ethernet (EEE) (IEEE, 2012a). Recent research into small networking equipment by the California investor-owned utilities and the Natural Resources Defense Council in support of California Energy Commission (CEC) Title 20 standards has helped characterize typical energy use and savings opportunities at the end use level (Hardy et al., 2013a). However, very little research has been conducted to characterize the power requirements of current Ethernet functionality as a subsystem that exists horizontally across a large number of products.

This research utilized a combination of literature review, product inspections, and manufacturer outreach to determine the power requirements for Ethernet functionality in end use products. We have mainly relied on manufacturer data sheets for recent single-port Ethernet controller products to characterize their power consumption and to identify best-practice approaches. Outreach to Intel Corporation and Realtek Semiconductor Corporation helped to confirm assumptions on the availability of certain energy-saving features in new products.

2.1 TECHNOLOGY SNAPSHOT

High-speed wired networking is one of the most prevalent, horizontally applicable secondary functions identified in earlier functional mapping research. Ethernet (IEEE 802.3 standard) technology dominates this functional category, particularly in client/edge device applications where networking functions are secondary to main product functions. A variety of electronic products in the home entertainment and office products areas use this technology, including televisions, home audio/video products, computers, printers, and multi-function devices (Table 1). Traditional white goods like refrigerators almost exclusively use wireless technologies for their networking functions.

Ethernet functionality in end devices is primarily provided by integrated Ethernet controllers that serve as both the physical (PHY) and link layer (e.g. Media Access Control or MAC) for the edge device, allowing it to physically connect to the broader network and providing a unique hardware address (MAC address) for the end use device containing the controller. Several vendors dominate the market for such controllers, including Intel, Marvell, Broadcom, and Realtek. Figure 1 provides an example of a typical integrated Ethernet controller coupled to a LAN port to form a device's wired network connection — in this case, for a desktop computer.

Table 1: Common devices using Ethernet technology

Product Category	Example Devices
Home audio and video	Televisions Blu-Ray players Set-top boxes Stereo components
Office equipment	Computers Printers Multi-function devices VoIP phones

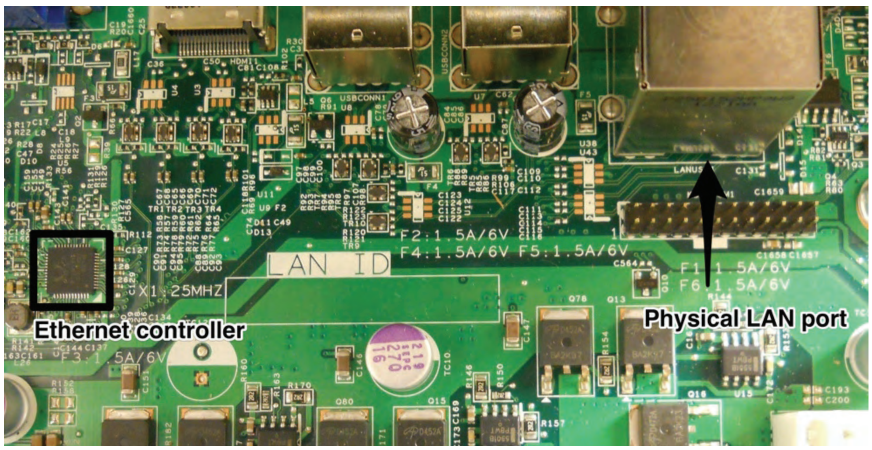

Figure 1: Ethernet controller and LAN port on desktop computer motherboard

Whereas Ethernet connectivity in network gear is an "always available" functionality, this is not necessarily the case in edge devices like computers, televisions, or audio/video equipment. These devices frequently enter low-power modes, and therefore have opportunity to completely power down Ethernet ports.

2.2 FINDINGS

A survey of currently available Ethernet controller products from leading manufacturers (Intel, Marvell, Realtek, Broadcom) indicates that wired networking functionality can consume widely varying amounts of power, even for products with similar overall functionality. For example, Figure 2 provides the range and median dc power consumption of Ethernet controllers while communicating with another Ethernet device. Power consumption can vary by over a factor of three depending on vendor and, more importantly, product vintage. Generally more power is required to maintain connections at higher link rates. Typically the power consumption at gigabit link rates is approximately double the consumption at 10 and 100 Mbps. There are negligible differences between 10 and 100 Mbps.

 xergyconsulting

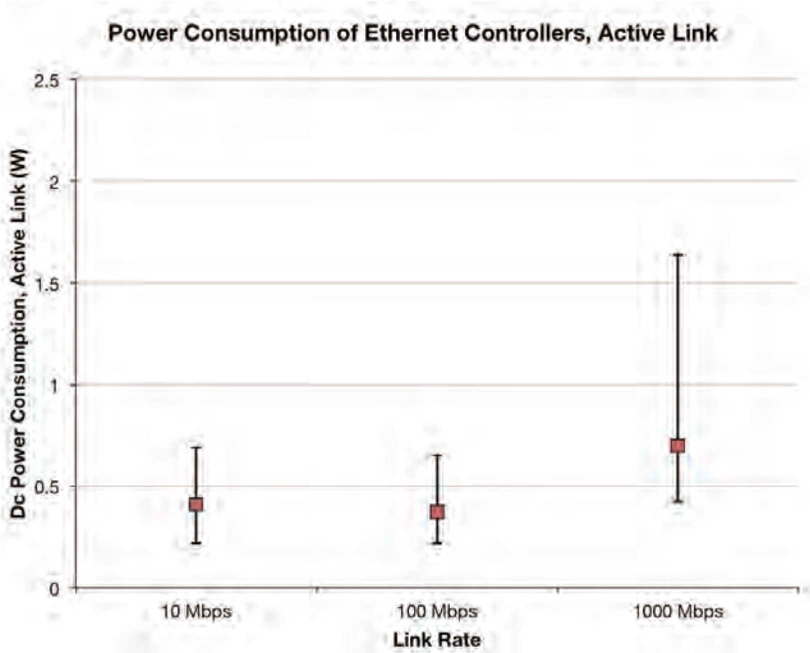

Figure 2: Dc power consumption of Ethernet controllers for products released 2008 and later

The values reported in the chart above are illustrative of Ethernet controllers in the existing stock of networked products today. However, there is evidence to suggest that manufacturers have made significant progress in recent years to bring overall power consumption values down through die shrinkage in the silicon fabrication process. A review of single-port Ethernet controllers from Intel shows that thermal design power (TDP) values — the maximum power consumption limits that hardware designers will use when "budgeting" the power consumption of Ethernet controllers into their products — have been steadily decreasing in products released since the early 2000s, as illustrated in Figure 3. TDP values are for worst-case power consumption, so are not indicative of real-world power consumption per se, but they are useful for illustrating an overall trend toward more efficient silicon. Ethernet controllers released within the past year consume approximately two-thirds less power than parts being offered in the early 2000s using an older fabrication process. The process of feature shrinkage is common among all integrated circuit manufacturers, so the trend we see here for Intel is likely representative of the industry.

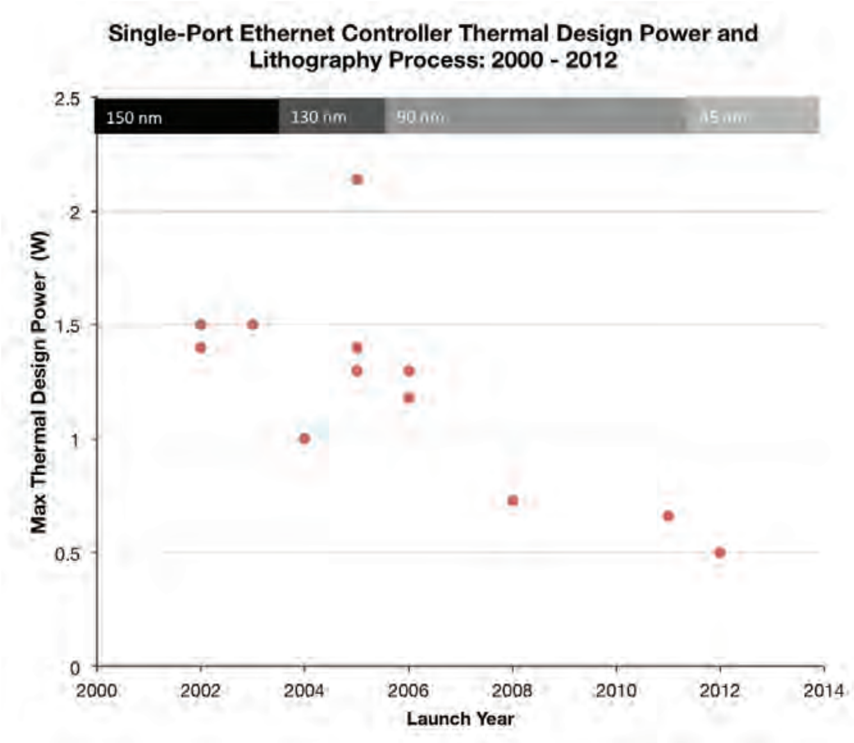

Figure 3: Evolution in thermal design power over time for Intel single-port gigabit Ethernet controller products

Until 2010 the power consumption of Ethernet equipment was relatively constant regardless of the amount of traffic traveling on the network. In September of 2010, the IEEE ratified the 802.3az or Energy-Efficient Ethernet (EEE) provision for the 802.3 standard. The 802.3az task force originally discussed a number of power-saving mechanisms to codify in standards, including adaptive link rate technology that would scale the Ethernet controller's link rate according to user bandwidth needs (Gunaratne et al., 2008). The ratified standard ultimately adopted a related but simpler energy savings mechanism that allows Ethernet controllers to sleep while links are idle. The concept is illustrated in Figure 4.

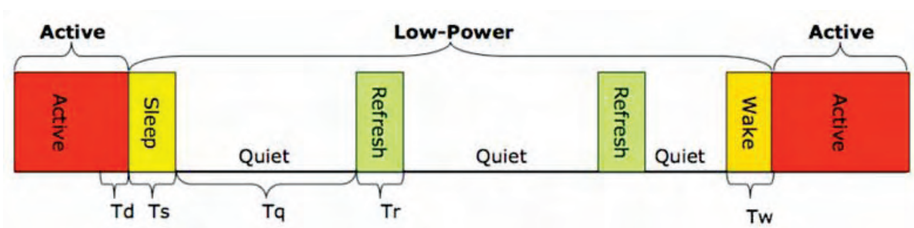

Figure 4: EEE sleep scheduling. (Source: Lawrence Berkeley National Laboratory)

Through a combination of product specification review and manufacturer interviews, our research has confirmed that most Ethernet controllers released into the market during the past year support EEE. Data sheets for mainstream products demonstrate that the energy savings potential touted by EEE advocates does exist. Figure 5 illustrates the dramatic power reductions achievable for an Intel EEE-enabled controller at a variety of link rates. For devices operating on legacy networks with 100 Mbps link rates, the savings between legacy products (no EEE) and EEE-enabled products are modest at about 25%. However, in contemporary gigabit networks with link rates of 1,000 Mbps, EEE can reduce power consumption during idle periods by over 55%. Idle power consumption with EEE enabled is within several milliwatts of "disconnected" power consumption (power consumption of the controller with no physical network connection).

EEE-enabled savings claims can vary depending on manufacturer. Realtek claims that, when operating at gigabit link rates, its implementation of EEE can reduce controller idle power from 332 mW to 48 mW, a reduction of approximately 85%. However even with EEE disabled, current controllers do appear to moderately scale power consumption to network activity, at least according to manufacturer-reported data. Note the reduction in power between active links and idle links in Figure 5, even when EEE is disabled.

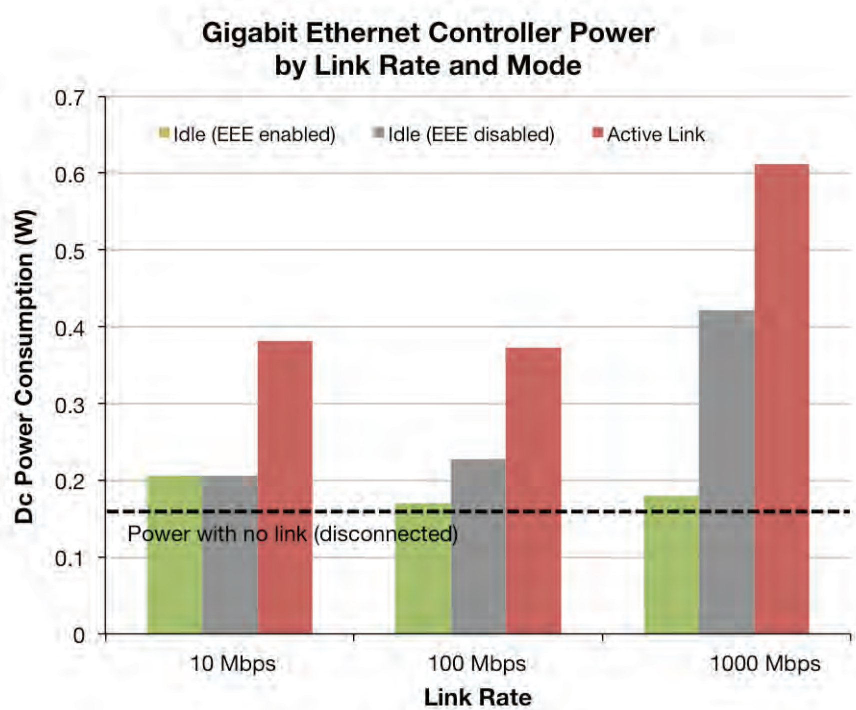

Figure 5: Impacts of link activity with and without EEE enabled for an Intel single-port gigabit Ethernet controller

A great caveat for EEE remains: devices on *both* ends of the link must have EEE enabled in order for controllers to properly scale power during idle periods. Natural Resources Defense Council research conducted by Ecova in 2012 and 2013 revealed that the mere presence of EEE-compatible Ethernet controllers in both a host and client device does not guarantee the intended EEE power scaling behavior. Ecova measured the power consumption of EEE-enabled home routers when connected to a variety of computers with EEE-compliant Ethernet controllers. EEE was disabled in software settings on about half of the computers by default, blocking both host and client devices from scaling power on the associated ports. Enabling EEE functionality on the affected computers involved changes to hardware driver settings that would likely challenge most computer users (Hardy et al., 2013b).

Table 2 provides a summary of expected ac standby power consumption associated with the various modes of operation present in today's Ethernet controller products. We have used a standard set of assumptions to estimate the power conversion impacts and derive ac power estimates from the motherboard-level dc power values reported by industry.[1] We assume that designers have been diligent in prescribing an efficient standby power supply design, as described in the "Power Supplies for Low Standby" section of this report. In general, today's Ethernet components can contribute anywhere from about 1 to 3 W ac to standby power when they are maintaining an active link and passing traffic; however, when ports are disconnected or devices are otherwise idle, the best designs should be able to throttle power back to below 0.2 W ac.

Table 2: Ac standby power impacts of Ethernet controllers

	Dc Power Consumption (W)		**Ac Power Consumption (W)**	
Operating Condition	Min	Max	Min	Max
Active Link, 1,000 Mbps	0.425	1.64	0.738	2.763
Active Link, 100 Mbps	0.22	0.65	0.397	1.113
Active Link, 10 Mbps	0.22	0.69	0.397	1.180
Idle Link (no EEE)	0.206	0.332	0.373	0.583
Idle Link (EEE)	0.048	0.18	0.110	0.330
Disconnected Port	0.013	0.16	0.052	0.297

2.3 BEST AVAILABLE TECHNOLOGIES

Some network hardware vendors also offer various "green Ethernet" products. TRENDnet, a networking hardware manufacturer, touts products with GREENnet technology. In addition to allowing Ethernet controllers to sleep during idle periods as per EEE, GREENnet can also detect when a host port is connected to a client over a short length of cable, such as in a home or small office. Since less power is required to deliver data over this shorter distance, the Ethernet controller scales its power back accordingly. Unlike EEE, this technology does not require that both ends of the Ethernet connection support GREENnet. Vendors like TRENDnet and D-Link can also detect disconnected Ethernet ports and reduce power to associated controllers to minimal levels, a technique that could be applied to edge devices. It should be noted that this would only generate meaningful savings for devices that do not otherwise support EEE; once EEE achieves broader adoption, such technologies would no longer be needed.

[1] We assume here that manufacturers employ a main ac-dc standby power supply with conversion efficiency of 75% along with a secondary dc-dc conversion stage with an efficiency of 80%, for a combined end-to-end efficiency of about 60%.

In edge devices that spend a significant portion of their duty cycle in low-power modes (e.g. televisions, home computers), one of the largest opportunities for energy savings is simply ensuring that Ethernet controllers and ports are powered down to their lowest possible state when end use products enter low-power modes. As an example, current Realtek gigabit Ethernet controllers will reduce power by almost two orders of magnitude depending on their state, ranging from active data transfer to sleep mode (i.e. when the end use product that contains the Ethernet controller is sleeping), as shown in Figure 6. Of course, it is up to end product designers to effectively harness these low-power features and ensure that controllers enter their lowest power states along with their host device.

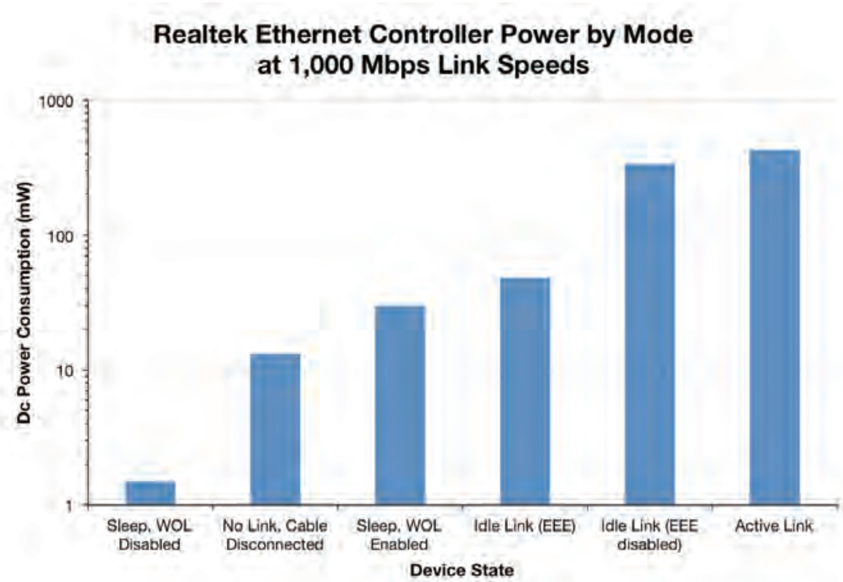

Figure 6: Range of power consumption for Realtek gigabit Ethernet controller by operational mode

Finally, standards continue to evolve, and researchers have continued to identify opportunities to expand energy savings opportunities in Ethernet. Several of the broad strategies proposed currently include:

- Classifying incoming packets and requests so that the Ethernet controller understands which ones it can safely ignore;
- Buffering messages intended for the controller that are low-priority so that they can be processed in bulk, thus eliminating frequent waking; and
- Enabling a small portion of controller hardware to serve as a "proxy" for the main controller on the network, automatically answering routine requests, such as Address Resolution Protocol (ARP), without having to wake the entire device.

Eckermann (2013) provides a useful summary of these new techniques, many of which bear similarity to those introduced into IEEE's 802.11 standard for Wi-Fi (see below).

3 POWER REQUIREMENTS FOR HIGH-SPEED WIRELESS NETWORKING

Wi-Fi technology is effectively synonymous with the so-called high-speed wireless networking functionality described in earlier research. Despite the relative ubiquity of Wi-Fi radios in products as diverse as computers, smart phones, and televisions, there is a surprising dearth of information about their power consumption. Several studies in the literature have examined power-saving techniques applicable to mobile products (Gupta and Mohapatra, 2007), but examinations of power consumption and energy savings mechanisms in stationary devices are relatively rare. Furthermore, manufacturer data sheets are less than forthcoming on the power requirements of Wi-Fi transceiver products. Even Intel, who provides detailed power information on the power consumption of its Ethernet controller products, does not publicly disclose similar information on its Wi-Fi transceivers. Conversations with industry experts illuminate some of the reasons for the current lack of information:

- Given that a large number of Wi-Fi transceivers are deployed in battery-powered, mobile products, there is significant innovation around power efficiency for these market segments. Manufacturers typically only disclose these values under non-disclosure agreements to maintain competitive advantage.
- Typical power consumption of wireless networking products like Wi-Fi transceivers can be difficult to estimate given the large number of variables that enter into the equation, such as distances between communicating transceivers, total number of clients on a network, and even the amount of external interference that transceivers might need to overcome.

This section sheds light on the issue of Wi-Fi power consumption, even in the absence of extensive data.

3.1 TECHNOLOGY SNAPSHOT

Wi-Fi products are based on the IEEE 802.11 body of standards that govern the physical and communication protocol requirements to enable interoperability between Wi-Fi products. The standard has evolved since its inception in the late 1990s. Current variants of the protocol include 802.11n (4th generation) and 802.11ac (5th generation). As illustrated in Figure 7 below, these products now enable data transfer in the 600 to 3,600 Mbps range, using up to three antennae.

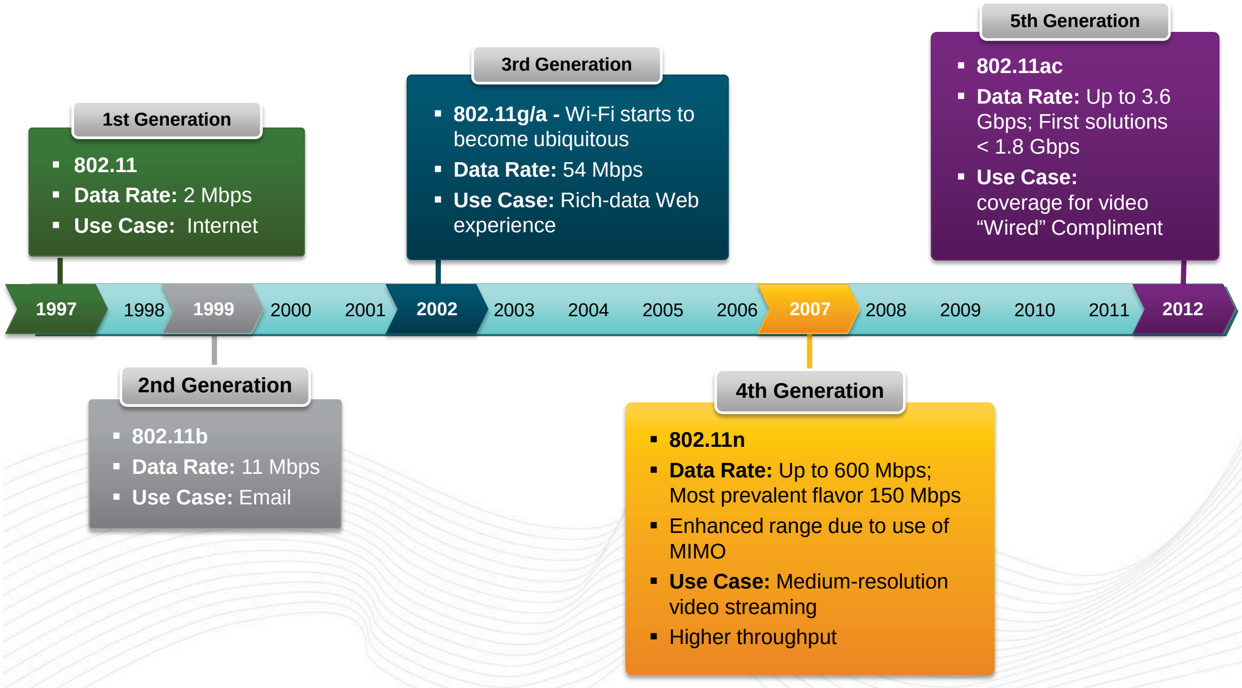

Figure 7: Timeline of Wi-Fi protocol evolution. (Source: Broadcom, 2012)

Wi-Fi functionality is typically provided through transceiver chipsets that enable digital devices to transmit and receive data according to the 802.11 protocol. In current generation products, this can include the use of three separate radios operating on different frequency bands. These packages may include integrated antennae or can also offer connections to external antennae for improved range and reception. In computers, such as notebooks and all-in-one desktops, Wi-Fi functionality may be provided through an adapter card, but in mobile devices and an increasing number of small form factor stationary devices, Wi-Fi transceivers are integrated directly onto the motherboard (Figure 8).

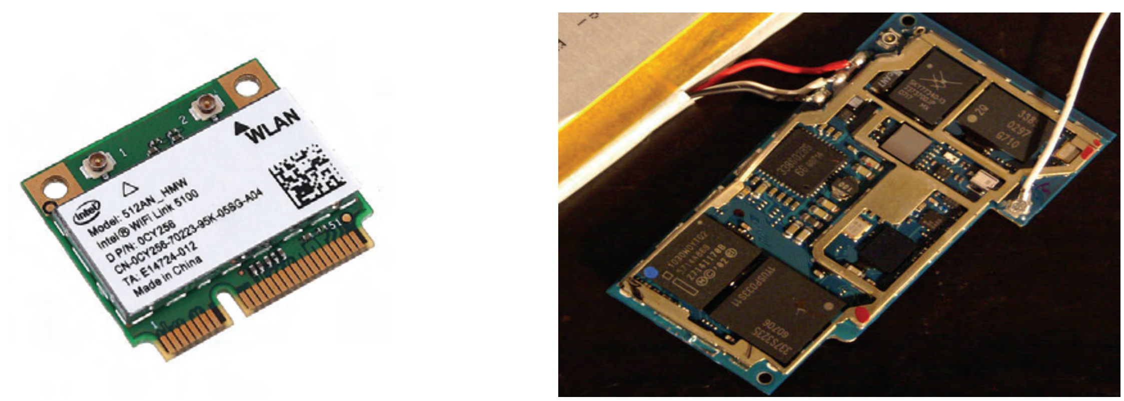

Figure 8: A computer Wi-Fi "half mini" adapter card (left) and integrated iPhone Wi-Fi transceiver (right)

Wi-Fi chipsets used to be found mainly in notebook computers and some desktops, but are now present in a broad array of plug load devices and appliances (Table 3). Many new connected “smart devices” are envisioned in the smart grid and home automation space, but these applications are limited at present.

Table 3: Common devices using Wi-Fi technology

Product Category	**Example Devices**
Major appliances	Clothes washers and dryers Refrigerators Ovens
Home audio and video	Televisions Blu-Ray players Media boxes (e.g. Apple TV and Roku) Stereo components Wireless speaker systems
Office equipment	Computers Printers Multi-function devices
Consumer mobile devices	Smart phones Tablets Digital cameras Music players
Smart grid and home automation	In-home energy displays Smart switches/outlets Home automation “hubs” Wireless sensors Smart thermostats

3.2 FINDINGS

Many of the same factors that influence power consumption in Ethernet-based networks have a similar impact in the wireless realm. For example, power consumption increases with link speed and distance between connected devices. However, extrapolations can be dangerous because a host of new variables can influence operational power consumption in wireless equipment many that are beyond our control to predict, such as “noisy” radio environments with significant interference or construction materials of indoor environments that degrade the wireless signal.

Testing by Ecova conducted on behalf of the Asia-Pacific Partnership and the IEA 4E Annex shed some initial light on which variables might influence Wi-Fi power consumption the most. In 2011, Ecova’s (formerly Ecos) Research and Policy group measured the power consumption of four Wi-Fi routers to examine the impacts of maximum link speed, network activity, and number of connected clients (Calwell et al., 2011). The research showed that the presence of Wi-Fi contributed anywhere from 0.25 to 2 W ac to total device power, depending on the operating frequency and number of radios operating simultaneously (Figure 9). Dual-band, 802.11n routers saw the greatest increase in power because of the presence of multiple, simultaneously operating radios.

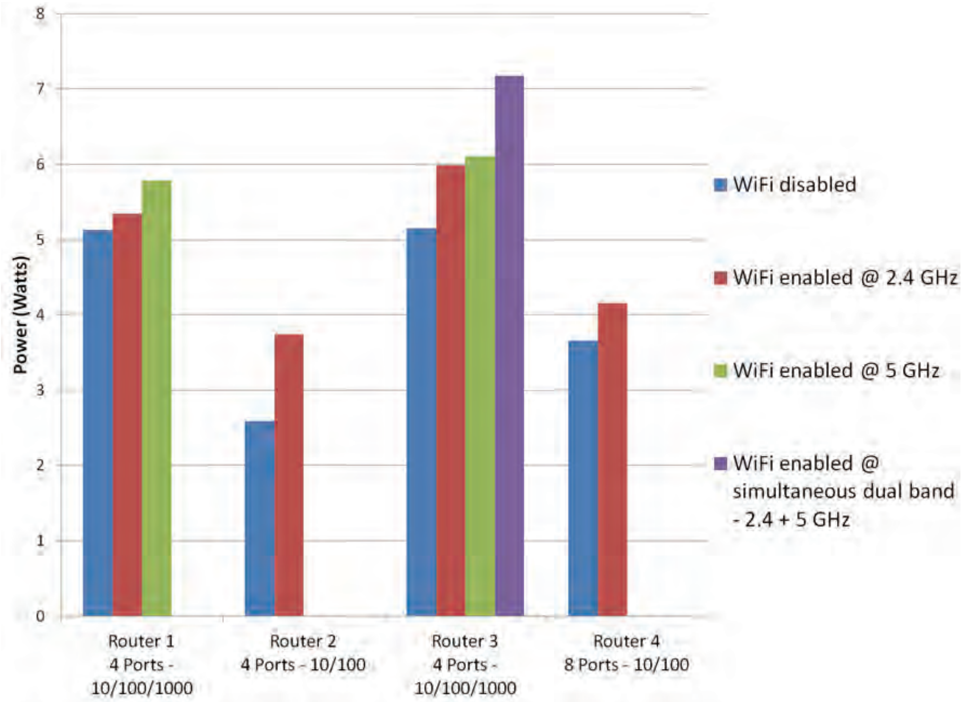

Figure 9: Ac power consumption of consumer Wi-Fi routers (Source: Calwell et al., 2011)

Early in our literature review, we uncovered engineering design materials that provided succinct comparisons between various wireless networking technologies, including Wi-Fi, Bluetooth, Zigbee, and several other more obscure technologies. These comparisons reduce each technology down to its characteristic "power efficiency" or power use per data transfer rate in units of W/Mbps. In the case of Wi-Fi, the metric was derived from power consumption characteristics in an 802.11b era product at a link rate of 40 Mbps and a power consumption of 0.210 W, yielding overall power efficiencies of 5.25 mW/Mbps (Smith, 2011). While this sort of power efficiency analysis might be useful to compare disparate technologies at a point in time, it does not provide a good means for extrapolation to next-generation Wi-Fi products with higher link rates. Current 802.11n and 802.11ac transceivers can achieve theoretic link rates from the hundreds of Mbps up to gigabit range, which would scale power consumption for individual transceivers to over 5 W dc!

Typical transmit/receive power consumption on today's chipsets — including some emerging 802.11ac chipsets — ranges from 0.5 to 1.9 W dc, which could yield ac power consumption values of 0.9 to 3.2 W, depending on the number of power conversion stages upstream of the Wi-Fi transceiver.[2] As shown

[2] We assume that up to two power conversion steps may be required to power a Wi-Fi transceiver: the first to convert ac to dc and potentially a second to further step down to lower voltage dc. Assuming that the typical efficiency of each stage is about 70%, this results in a multiplier effect on dc power consumption from 1.4x to 2x.

in Figure 10, power use for the latest generation of 802.11ac chips is fairly comparable to recent 802.11b/g/n products, according to data furnished by Realtek Semiconductor.

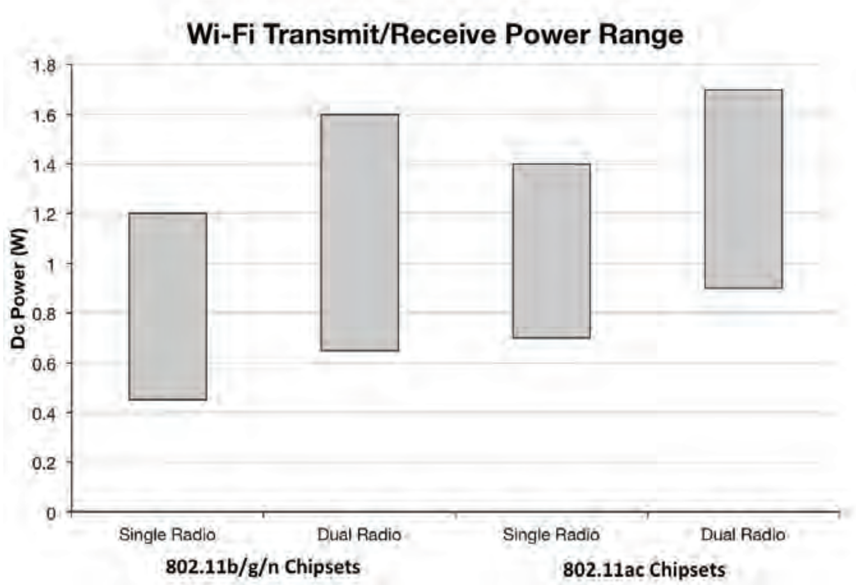

Figure 10: Transmit/receive power for 2013 Wi-Fi chipsets

The deep sleep levels achievable in today's products are fairly independent of the maximum link speed the product supports. Rather, they are more a function of product vintage. The latest generation silicon is capable of idle power (e.g. when a transceiver has no network traffic or might be scanning for an access point) below 1 mW; however, standby or idle power in Wi-Fi chipsets ranges by quite a large amount, from under 1 mW to about 0.25 W ac once power supply losses are considered. Table 4 summarizes the contribution of Wi-Fi to standby power in today's products.

Table 4: Ac standby power impacts of Wi-Fi transceivers

	Dc Power Consumption (W)		Ac Power Consumption (W)	
Operating Condition	Min	Max	Min	Max
Active Transmit/Receive	0.51	1.9	0.88	3.2
Standby/Idle	0.004	0.13	0.036	0.25

3.3 BEST AVAILABLE TECHNOLOGIES

There are two broad categories for saving energy in Wi-Fi, much as with Ethernet. The first scales power back during times of network inactivity or when the transceiver's end use product is in a low-power mode; the second minimizes power consumption during periods of network traffic by scaling power consumption to the current needs of the network. We focus here on the first general strategy, as it is most applicable to client devices where wireless networking is a secondary function. As for scaling power consumption according to “live” network traffic conditions, certain network hardware manufacturers such as TRENDnet have described GREENwifi products that allow routers to reduce broadcast power when client devices have a strong signal (i.e. when all devices talking to the router are located nearby) (TRENDnet, 2011).

Significant potential exists to take advantage of power scaling during idle traffic times or when the end use product is in a low-power mode. Existing 802.11 standards provide support so that wireless routers and connected edge devices can coordinate power saving strategies with one another. Except where specially noted, these mechanisms apply to the state of the Wi-Fi transceiver and not the state of its end use device. For example, a wirelessly networked printer could appear to be “on” to the user, but its Wi-Fi chipset could be operating in a power-saving mode. In most cases, the strategies embedded in the standard benefit client devices the most, leaving more of the “active” network management tasks with the access point or router. Some of the strategies that exist in the current 802.11 standard are, in layman’s terms:

- Allowing access points to answer certain routine requests on behalf of connected devices so that they can remain in a low-power state;
- Allowing access points to determine which messages need to be delivered to a client device and which can be buffered until the device wakes up;
- Enabling client devices to roam between two access points in an efficient manner (mainly for mobile devices);
- Providing mechanisms for access points and routers to negotiate longer “rest” periods between routine communication so that client and potentially access point devices can sleep for longer periods of time; and
- Enabling even longer rest periods for battery-powered client devices with greater power constraints.

The intent of most of these strategies is that client devices will be able to spend a significant portion of their time in a standby mode with minimal noticeable latency impacts on the user. There are great energy savings benefits to these sorts of strategies for client devices. Current Wi-Fi chipsets can achieve idle power of 1 mW and below, several orders of magnitude lower than power consumption in active transmit/receive modes. There have been few if any component-level measurements of these strategies to date, but a study related to power-saving strategies for Wi-Fi voice over IP (VoIP) phones proves illustrative. Researchers increased the interval for communication between client wireless devices and their access point. Wi-Fi power use[3] could be cut by one third by increasing this communication interval from the default of 200 ms up to 1,000 ms or one second, while adding only an additional 800 ms of latency. There are diminishing returns associated with increasing beacon intervals much beyond this point, because the Wi-Fi chipset is already predominantly idle.

[3] In this case, power use is directly proportional to current because the device is operated at a fixed voltage.

The single broadest caveat regarding all of these energy savings mechanisms is that, much like EEE, they require several important prerequisites in order to deliver the intended savings:

- Both the client device and wireless access point must support the desired power management feature. Many 802.11 power management features, such as "power poll" have been a part of the standard since 2007. However, several newer power management features, such as the ability for the access point to buffer traffic for an edge device, were only ratified as of 2012, so many legacy products will not support them.
- Even if the physical Wi-Fi chipset implemented in a product supports the latest 802.11 power management features, it is still up to individual hardware designers to appropriately integrate these products and leverage their power saving features. Clearly, aggressive power management of Wi-Fi chipsets is not on the forefront of a designer's mind when implementing Wi-Fi in stationary plug load products compared to, say, a smart phone or tablet where battery life is a primary design concern.

3.4 TRENDS

We expect 802.11ac products to race into the market at a dramatic pace. Furthermore, market forecasts suggest that global Wi-Fi chipsets will continue their inexorable rise, effectively tripling from the period between 2010 and 2015 (Figure 11). This continued rapid growth of Wi-Fi-connected devices means that it will be increasingly important for policymakers to ensure that stationary wireless devices utilize 802.11's power-saving mechanisms as readily as their battery-powered cousins, especially as power consumption in transmit/receive modes continues to gradually rise (Figure 10).

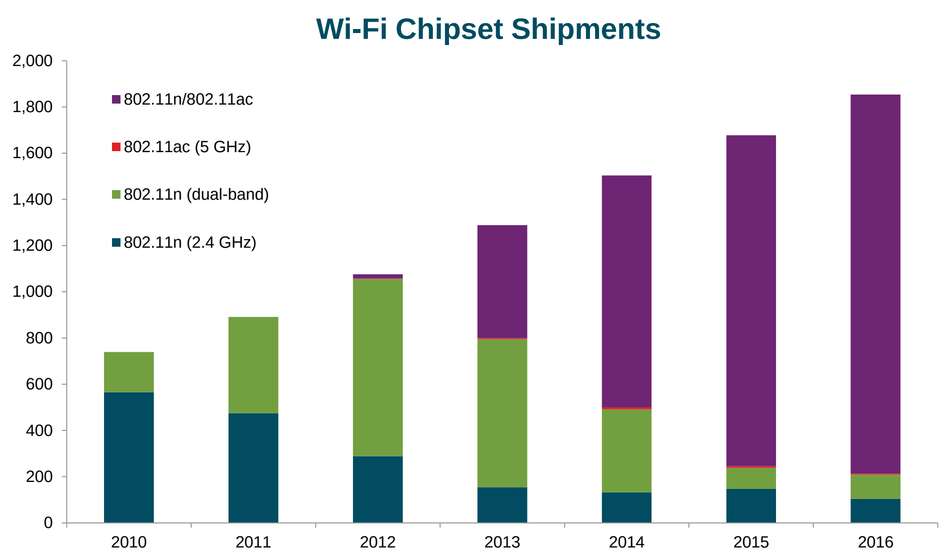

Figure 11: Global Wi-Fi chipset market forecast (Source: Broadcom, 2012)

Even as 802.11ac technology enters the marketplace, Wi-Fi products with still higher link rates are already planned. The 802.11ad standard or "WiGig" will provide wireless networking capabilities at theoretical speeds of up to 7 Gbps — 3,500 times faster than the first Wi-Fi protocol from the late. WiGig is really intended as a “cable-cutting” technology. Anticipated use cases include beaming high definition video from computers to televisions or wirelessly docking laptops in office environments. This means that WiGig products will likely exist alongside the more mainstream Wi-Fi networking products we know today, expanding the foothold of network standby. The first WiGig-enabled end use products are anticipated to hit the market in 2014. Manufacturers currently claim that WiGig chipsets should be able to achieve power consumption below 1 W dc, even with the significantly higher data rates. However, these claims still need to be verified. Even if power consumption levels for the technology are kept in check, it is still always possible that implementation in stationary products like wireless docking station will be more geared toward constant connectivity rather than aggressive power management.

4 POWER REQUIREMENTS FOR LOW-SPEED WIRELESS NETWORKING

As mentioned in the introduction to this report, low-speed wireless networking functions are nascent in the market and could conceivably be served by a variety of technologies and protocols, including Zigbee, Z-Wave, Bluetooth, and even Wi-Fi. We have focused our research efforts in this section on Zigbee for several reasons:

- Zigbee is based on an open IEEE standard (802.15.4) and, as such, has achieved larger adoption among semiconductor manufacturers than Z-Wave transceivers, which are currently only manufactured by Sigma Designs and Mitsumi.
- Zigbee was specifically designed for wireless networking applications with low bandwidth requirements. This aligns particularly well with the low-speed wireless networking secondary function that is meant to capture connectivity between smart appliances.
- Bluetooth has predominantly been used as a “cable-cutting” approach to enable peripheral device connectivity (e.g. wireless mice and keyboards or wireless data transfer from smart phones).

Admittedly, Zigbee faces the steepest competition from Wi-Fi given the relative ubiquity of Wi-Fi routers in homes and businesses. Later in this section, we examine the wisdom of using Wi-Fi to provide smart device connectivity and examine current trends in the smart appliance market to help establish policy priorities.

4.1 TECHNOLOGY SNAPSHOT

Zigbee or 802.15.4 technology is a low-power, radio frequency wireless networking technology ideally suited to automation, control and monitoring applications where bandwidth requirements are relatively low and where devices may be battery-powered. The basic functionality for Zigbee connectivity is encapsulated in integrated transceivers that simply transmit and receive data according to the Zigbee protocol. However, more frequently we see “system-on-chip” or SoC solutions that contain a Zigbee transceiver as well as an onboard, low-power microprocessor for driving custom-programmed applications. The thermal, space and power constraints on many Zigbee applications, such as sensors, result in highly compact products with very low power consumption and robust power management support. As mentioned in previous sections of this report, Z-Wave is a competing technology that is currently exclusively produced by Sigma Designs and Mitsumi.

In *Mapping Secondary Product Functions to Products and Operational Modes*, market research revealed that low-speed wireless networking functions like Zigbee are currently in their extreme infancy. However, growth in the smart appliance market could rapidly increase the use of this secondary function and associated technologies. Zigbee networking solutions can be found today in smart appliances including laundry equipment, room air conditioners, water heaters, dishwashers, refrigerators, and thermostats. Although not discussed in detail here, the basic motivation for most smart appliances is to provide communication and control to facilitate applications such as home automation and coordination with the electric grid (e.g. to allow electric utilities to directly control demand through smart loads during critical peak periods). In essence, this is a technology partly designed with energy efficiency and demand-side management in mind, so it is not surprising that its standby power impacts are minimal.

Table 5: Common devices using Zigbee/Z-Wave technology

Product Category	Example Devices
Major appliances	Clothes washers and dryers Dishwashers Refrigerators Ovens Microwaves Water heaters Window air conditioners
Home audio and video	Potential use in remote control applications (See Section 6)
Smart grid and home automation	In-home energy displays Smart switches/outlets Home automation “hubs” Wireless sensors Smart thermostats

4.2 FINDINGS

Data sheets and product briefs from a variety of manufacturers revealed, not surprisingly, that current Zigbee products consume minimal amounts of dc power, even when transmitting and receiving data. Figure 12 provides a summary of typical reported transmit/receive power consumption for a collection of 23 products from five different manufacturers. Even today’s SoC solutions, which include small, onboard microprocessors, consume less than 0.2 W of dc power during communication modes.[4]

[4] For comparison, we also examined the equivalent power requirements for a Sigma Designs Z-Wave product and found comparable transmit/receive power (approximately 0.07 W).

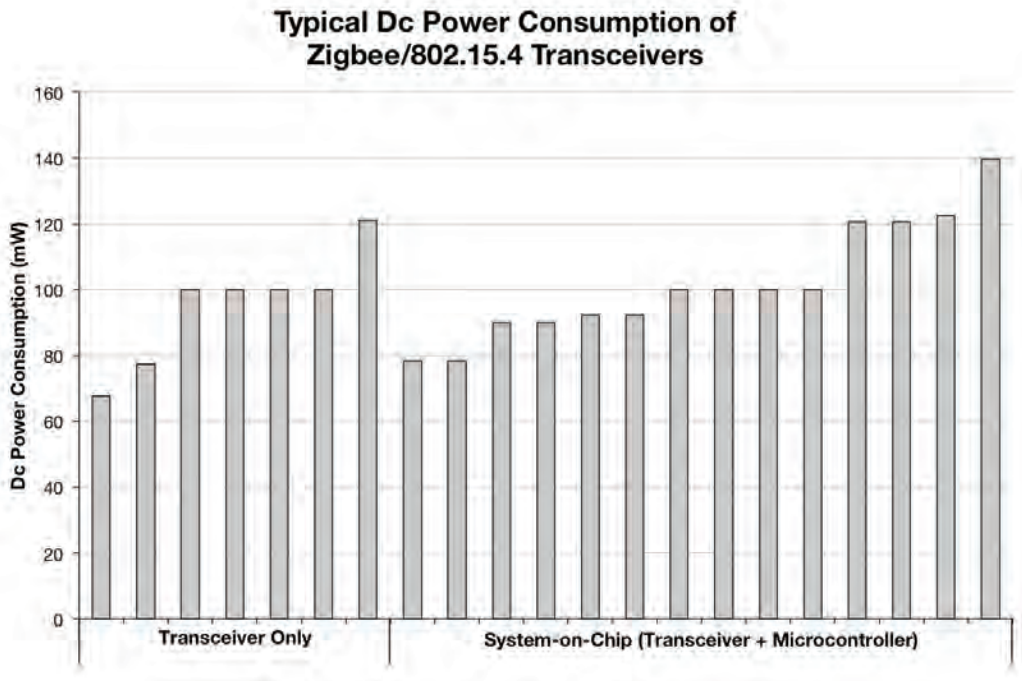

Figure 12: Transmit/receive power consumption for current Zigbee transceiver and SoC solutions

Zigbee products can power down to extremely low levels during periods of network inactivity and can wake with low latency (less than 30 ms). Typical current readings for Zigbee devices while sleeping are in the micro-amp range, resulting in sub-milliwatt power consumption. In most envisioned Zigbee applications, namely automation and monitoring, data will travel in short bursts. The anticipated schedule for network traffic is significantly more sparse than in traditional high-speed data networks. As a result, Zigbee and related products should be capable of sleeping the vast majority of the time.

A summary of Zigbee power consumption values from current products is provided in Table 6.

Table 6: Ac standby impacts of Zigbee technology

	Dc Power Consumption (mW)		Ac Power Consumption (mW)	
Operating Condition	Min	Max	Min	Max
Active Transmit/Receive	68	589	113	982
Sleep	0.004	2.3	0.036	3.9

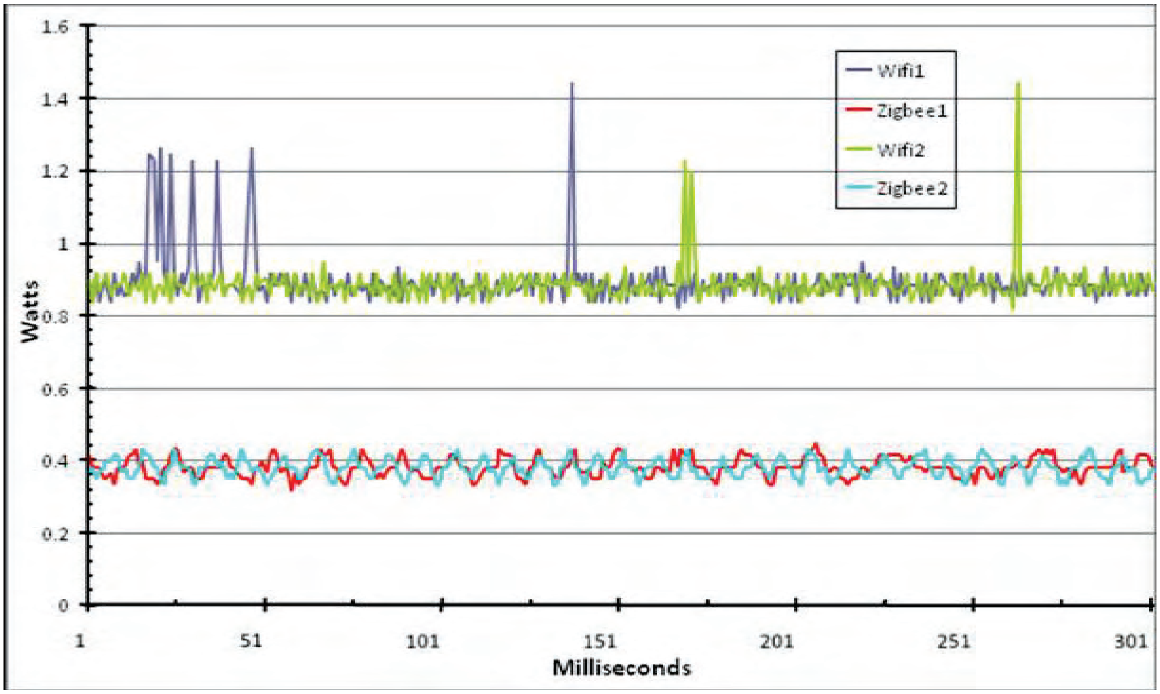

Figure 13: Power consumption comparison of Wi-Fi and Zigbee radios (Source: Drake et al., 2010)

Despite the extremely low standby power consumption of Zigbee radios, their efficiency in operation can again be greatly impacted by hardware and firmware implementation. In a GE-sponsored investigation of wireless radio power consumption, researchers conducted a side-by-side comparison of Wi-Fi and Zigbee radios connected to a smart appliance communications module that received 50-byte packets at 5-minute intervals from a home energy management system (Drake et al., 2010). The base power consumption of the Wi-Fi chipset was over twice as high as the Zigbee technology; however, the Zigbee transceiver was implemented in such a way that it was not allowed to go to sleep. Given Zigbee transceivers' extremely low current requirements in standby mode (on the micro-amp scale), enabling sleep at the hardware and firmware level could have virtually eliminated the power consumption in Figure 13.

4.3 TRENDS AND THE QUESTION OF BEST AVAILABLE TECHNOLOGY

Low-speed data networking functionality is primarily being driven by the growth in smart appliances. Until now, few smart appliances have been available on the market, and their market penetration has been slow to expand. Many smart appliances are still involved in small electric utility pilots and have failed to hit retail outlets in any large number. There are signs that this is about to change. According to a 2010 report from Pike Research, the smart appliance market is transitioning from the development stage to the higher growth commercialization stage (Pike Research, 2010). As illustrated in Figure 14, the global market for smart appliances is expected to grow slowly, with only about USD 2.2 billion in sales by 2013 and USD 6.3 billion by 2015. The global market is expected to experience greater growth in 2013 and is projected to reach USD 26.1 billion by 2019. This will result in approximately 118 million smart devices deployed, which will comprise about 8% of the world's installed appliances later in the decade (St. John 2012).

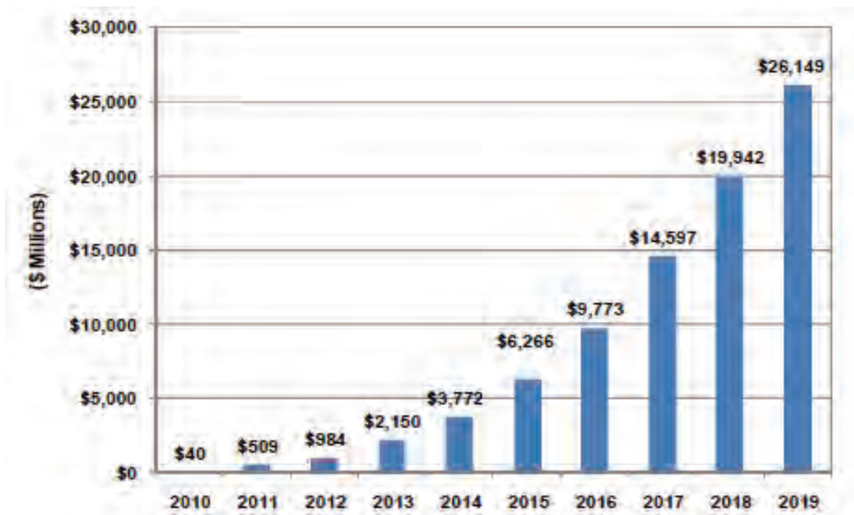

Figure 14: Smart Device Market Value (USD), World Markets: 2010-2019 (Source: Pike Research, 2010)

Of course connected devices need not be appliances, and indeed industry visionaries have been speaking of an "Internet of Things" or IoT for a number of years, in which everyday objects like alarm clocks, watches, light sockets, and even doorknobs will have some inherent connectivity, often by means of a wireless network connection. As an example, Cisco estimates that the IoT — including not just computers, tablets, and smart phones, but any type of connected object — will grow to approximately 50 billion devices by the year 2020. Cisco estimates that this represents only a small fraction — less than 3%, a relatively conservative number — of the total number of connectable objects on earth, but the values are still staggeringly large (Cisco, 2013). Many of these smart and connected devices will, of course, be mobile in nature: watches, smart phones, wearable computers, e-readers, tablets, etc. However, many of the devices in our lives continue to be stationary, and we expect continued proliferation of wireless technology in these devices. Imagine a world in which billions of alarm clocks, coffee makers, microwaves, refrigerators, laundry machines, light sockets, televisions, audio systems, and even garage door openers contain wireless radios of some form or another — and add to that other smart devices that have not been conceived yet — and one begins to appreciate the potentially enormous impact that this trend could have on network standby power.

Of course the story of the IoT is largely an open book, but the direction of the narrative depends largely on the choice of networking technologies. As stated above, Zigbee is a technology that was truly designed for the control and monitoring functions required by many automation and controls applications. However, Zigbee is still far from the level of interoperability that has been achieved in Wi-Fi and Bluetooth, Z-Wave even less so. Home automation and home energy management solutions have largely developed on proprietary “islands” or ecosystems that utilize the basic 802.15.4 physical and link layers, but build proprietary “stacks” of applications on top of this. To achieve true plug-and-play functionality, one must own smart appliances as well as a home energy management system from a single vendor. Even if interoperability were achievable today among Zigbee devices, it is questionable whether homeowners would purchase the necessary hub required to communicate with them.

As a result, many appliance manufacturers have either dropped Zigbee technology altogether or supplemented their smart appliances with Wi-Fi chipsets as well. GE discontinued its Zigbee-based Nucleus products for Brillion, whose current communicating wall ovens connect via Wi-Fi. LG’s THINQ line of smart appliances (pictured below) also connect over a standard Wi-Fi network. Samsung smart laundry products contain Wi-Fi in addition to Zigbee radios. Wi-Fi is naturally the incumbent wireless networking technology and has significant market penetration in homes. However, with data rates in current Wi-Fi technology now approaching gigabit range, it is safe to say that Wi-Fi is complete overkill for many smart device applications, where monitoring and automation are the primary objectives. Zigbee is the more appropriate choice and, as demonstrated by GE testing, the more efficient one as well.

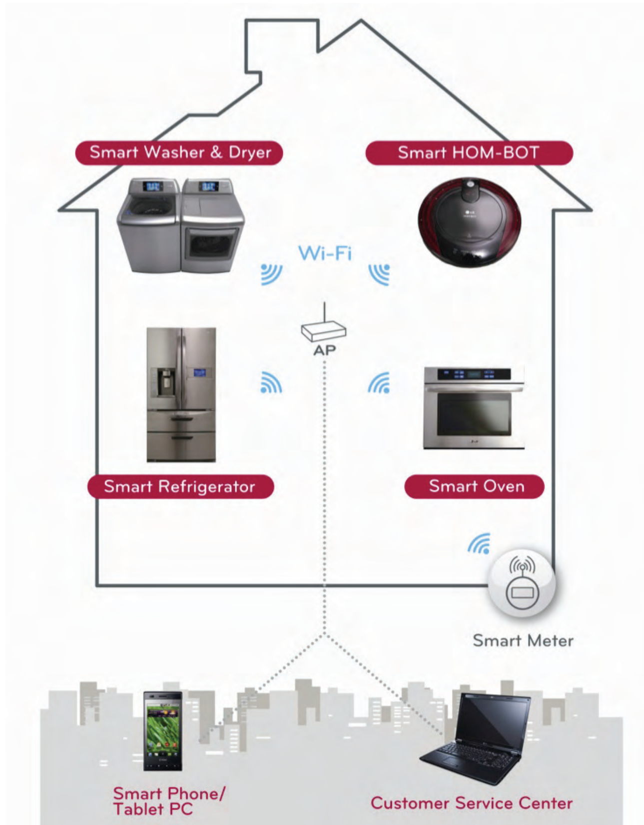

Figure 15: LG Thinq smart appliances with Wi-Fi connectivity (Source: Davies, 2011)

Faced with this smart device connectivity "battle," some might hark back to the VHS vs. Betamax or HD-DVD vs. BluRay format wars. However, the battle for home automation and smart device connectivity is many times more complex due to the array of potential protocols that devices could use. Wi-Fi, Bluetooth, and Zigbee are certainly the top three, but some home automation vendors have chosen to design so-called "god boxes" — wireless networking hubs with over half a dozen radios supporting *up to 10* wireless networking protocols — to avoid having to side with one technology.

With such open competition, the energy efficiency community should maneuver itself safely ahead of this trend and begin outreach now to encourage manufacturers to adopt the most appropriate networking technology for their application.

5 POWER REQUIREMENTS FOR DISPLAY FUNCTIONS

5.1 APPLICATIONS AND MANUFACTURERS

Secondary displays are ubiquitous in many types of mobile and plug-in devices. For the purposes of this study, we group them into the categories shown in Table 4.

Table 7: Categories of devices that use secondary displays and example devices

Product Category	Example Devices
Major appliances	Clothes washers and dryers Refrigerators Dishwashers Water heaters Ovens Microwaves
Home audio and video	Blu-Ray players DVD players CD players Amplifiers
Office equipment	Printers Multi-function devices VoIP phones
Automotive	In-dash information displays
Consumer mobile devices	Smart phones Digital cameras Music players
Other	Industrial applications Medical equipment Military

Liquid crystal display (LCD) panels dominate the small-medium (9 inches in diameter or smaller) display market. In 2013 an estimated 79% of displays shipped will be LCDs, and 9% will be organic light emitting diode (OLED) displays (DisplaySearch, 2012). The small-medium display market is dominated by five major players: Samsung, Japan Display, Sharp, LG Display, and Chimei Innolux (Figure 17). Of the five major manufacturers, only Samsung produces OLED displays as well as LCD panels.

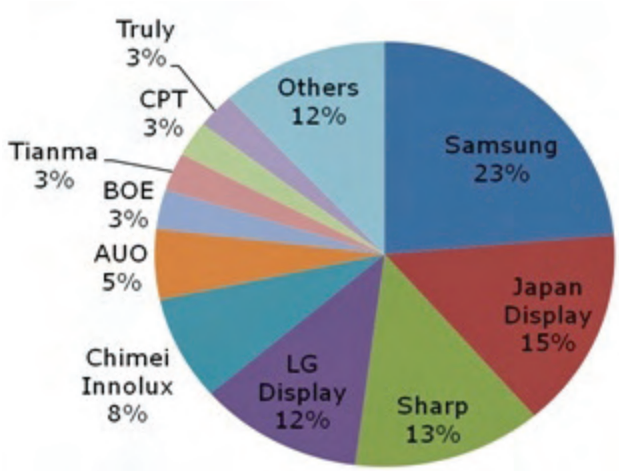

Figure 16: 2012 revenue share for manufacturers of displays 9 inches or less (DisplaySearch 2013).

OLED technology is still in development and is available in only a small number of products. Samsung dominates small-medium OLED display production, manufacturing displays for Samsung phones and tablets (OLED-Info, 2013). Relatively few OLED displays are used outside the mobile market.

5.2 DISPLAY TECHNOLOGIES

In the previous study, we identified two distinct types of displays. Alphanumeric displays show characters and numbers using a small number (7 to 16) of pixels for each character (Figure 17, left). They are used when basic information needs to be relayed to the user, such as time to cook on a microwave, or track number on a CD player. Informational displays use a dot matrix, and thus can display any visual information, such as pictures, video and text. Such displays are typically used on mobile devices, smart appliances, and office equipment (Figure 17, right). The panel and backlight technologies described below may be used for either display type.

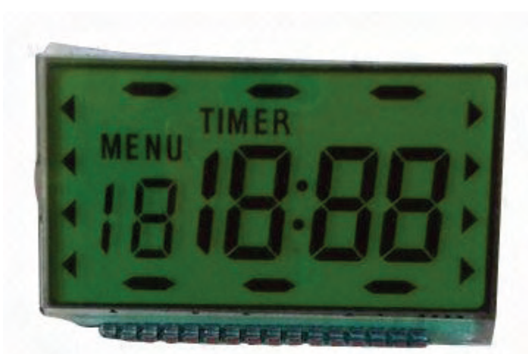

Figure 17: Alphanumeric display (left) and information display (right)

Secondary display technology is dominated by LCD panels backlit with light emitting diodes (LEDs). Cold cathode fluorescent lamp (CCFL) backlights are still used, but appear to be in the process of being phased out for LEDs. Reflective LCD panels are illuminated using reflected ambient light, eliminating the need for a backlight. A transflective LCD panel reflects ambient light when in bright conditions, and uses a backlight to illuminate the display in low ambient light.

Two major types of LCD panels exist, and are appropriate for different types of content. Pixels on a twisted nematic (TN) display are open when no electric field is applied to the liquid crystal, and closed in the presence of an electric field. Consequently TN displays draw less power for white content than for dark content. The converse is true for vertically aligned (VA) and in-plane switching (IPS) panels. Displays showing dark backgrounds or video content draw less power with these “normally black” panels than with a TN, or “normally white” panel.

OLED displays have potential to save energy over LCDs because rather than controlling how much light passes through each pixel with liquid crystals, the pixels emit light, eliminating heat loss in the panel. Because black content requires no light emission by the OLEDs, and dark content requires little light, OLED displays are best suited for video or graphics rather than bright content like internet browsing or email.

These technologies are summarized in Table 8.

Table 8: Summary of display technologies

Technology	Relative power draw	Advantages	Disadvantages
CCFL backlit LCD	High	Low cost, small number of lamps required	Requires more power than LED backlight
LED backlit LCD	Middle	Lower power than CCFL, inexpensive and mature compared to OLED	Requires backlight
Transflective LCD	Low	Takes advantage of ambient light when possible	Low contrast ratio, high cost
Reflective LCD	Lowest	No backlight, low power	Unusable in low light
OLED	Potentially lower than LED backlit LCD	Low power, potential to improve, emissive technology allows light produced only at pixel when needed, better display quality	Limited availability, new technology, may not save energy for bright content

5.3 IMPLEMENTATION

The power budget of a LCD generally consists of two parts: power draw of the backlight (if present) and power draw of the panel as it addresses, opens, and closes pixels. The power budget of an OLED display has parts analogous to those of an LCD: power to address and control each pixel and power to light each pixel.

Display manufacturers generally package the display panel, panel control, backlight, and backlight control into a display module. Figure 18 shows a block diagram for a typical LCD computer monitor. Of the components shown, secondary display modules usually include panel and backlight dc-dc power conversion (blue), panel control and the panel itself (orange), and backlight control and the backlight itself (red). In our literature review we did note, however, that some display modules do not include dc-dc power conversion for the backlight. External to the display module (and not considered in the power estimates below), the device must have a main processor (green) to send data to the LCD and backlight units and main and possibly standby ac-dc power supplies. Standby power for secondary displays can be zero, since neither the panel nor backlight functions are needed when the screen is off. The device's main controller can simply turn off panel and backlight power when the display is not in use.

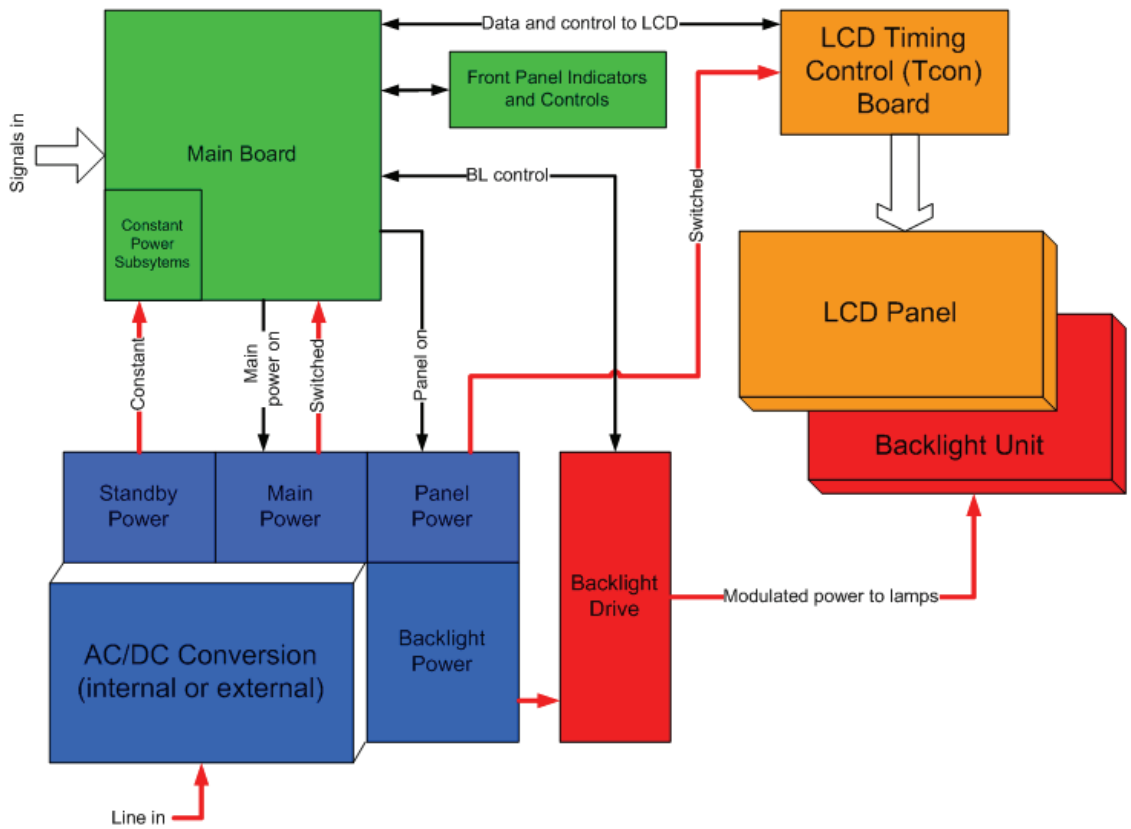

Figure 18: Block diagram of LCD display (Kundu et al., 2013).

As defined above, the display module is a horizontal technology, working in a similar manner no matter the device into which it is installed. The external processor controlling the display module, however, is device-specific and may carry out other primary or secondary functions. Consequently, active and standby power for the main processor is device-specific, and not quantified in the power estimates below.

5.4 POWER BUDGETS

5.4.1 Methods

We estimated power budgets for displays using datasheet information, discussions with experts, and information from articles, blog posts, and papers. We took the majority of quantitative data from datasheets, and backed up those data with information from other sources.

LCD datasheets usually listed typical power for the display and/or for the backlight and panel. If typical power was not listed, we calculated typical power from typical voltage and current of backlight and panel. On most datasheets, manufacturers reported typical voltage, current, and/or power for a multi-bar grayscale pattern. Other product datasheets listed typical power and/or current for black and white screens. In those cases, we estimated power by averaging the power draw for black and white patterns.

We obtained a small number of datasheets for OLED displays. We relied on information from blogs and manufacturer websites and a manufacturer interview to estimate OLED display power.

5.4.2 Findings

Active mode power required by a display is related to several of the display's characteristics such as size, resolution, and brightness. Size and resolution are loosely related; larger displays tend to have higher resolution. Higher resolution displays require more power not only to control more pixels in the panel, but also to produce more light to counteract the smaller apertures created by a denser grid of pixels. Increased screen brightness only involves an increase in backlight power. In the discussion below we normalize power by screen area for the most intuitive metric, but note that we arrive at similar results when we normalize power by resolution.

We found a broad range of display power, from hundredths of a mW for a transflective monochrome LCD for wrist-worn devices, to over 15 W for a CCFL backlit LCD for automobile use (Figure 19). Power and screen size are loosely related, but technology type is a determining factor in power draw.

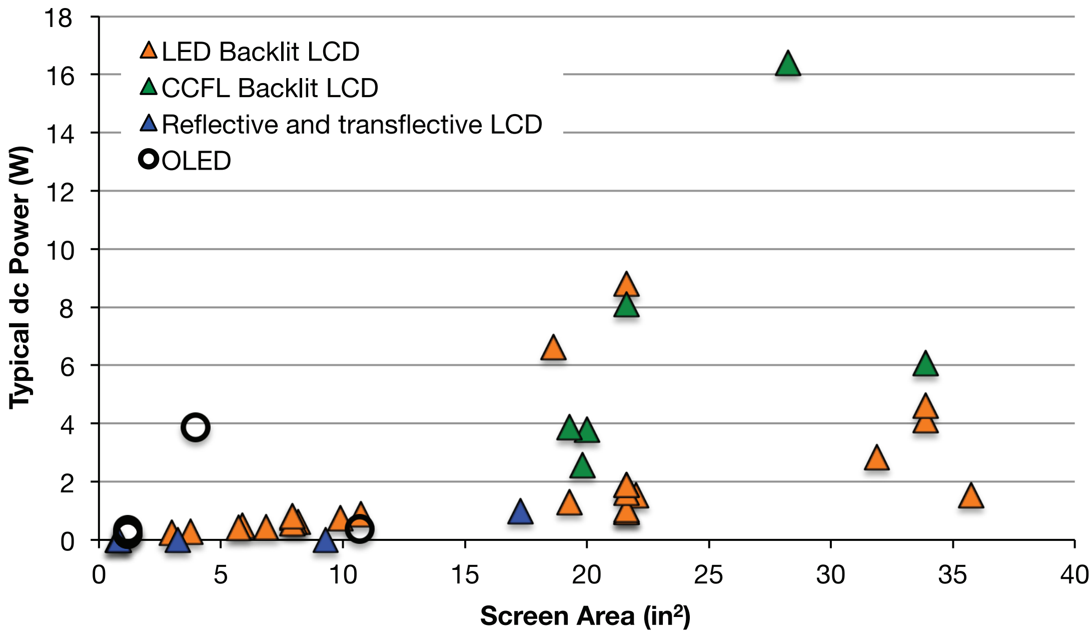

Figure 19: Active mode dc power draw versus screen area for various technologies.

Backlight power is the greatest portion of the power budget of a display. We found that the backlight can draw over 90% of total display module power. LED backlit LCD panels are generally more efficient than CCFL backlit panels, yet a wide spread within these technologies does exist (Figure 19). Displays that use reflected light to illuminate the panel draw considerably lower power to function (Figure 19). To further reduce power, Sharp Memory LCD reflective panels can reduce power to the panel control when the display is static. Without a backlight and with very little power required to control an often static display, these displays draw hundredths of a mW, and are used in battery-powered, wrist-worn devices.

We made only a small number of power draw estimates for OLED displays, which are comparable to the most efficient LED backlit LCD panels (Figure 20). OLED technology is expected to improve, however, and one manufacturer estimates that OLED power requirements have potential to decrease about 45% (Universal Display, 2013).

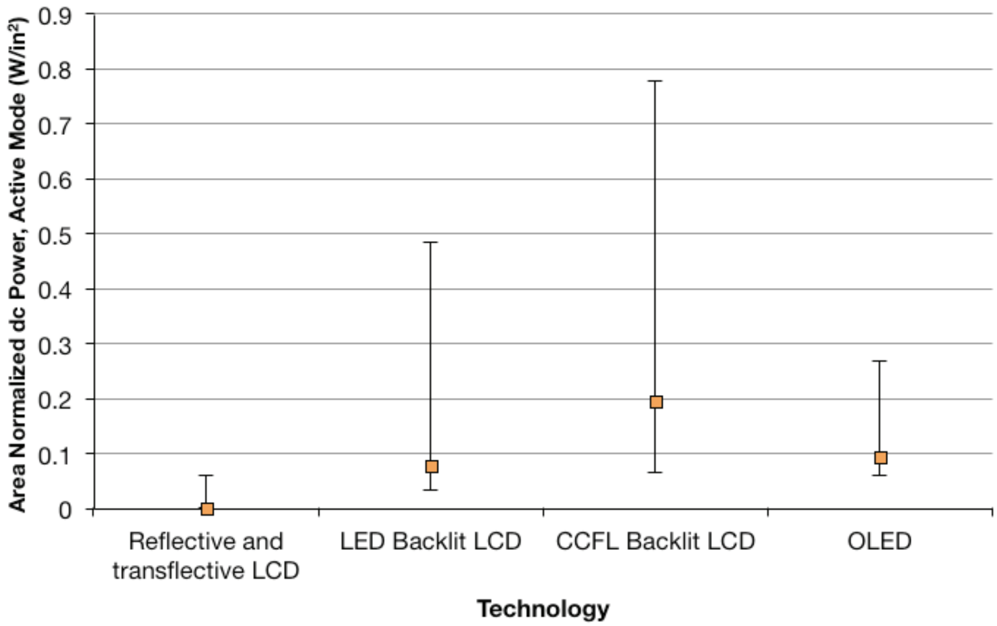

Figure 20: Median (orange boxes) and range (vertical bars) of area normalized, active mode dc power by technology.

5.5 ENERGY SAVINGS OPPORTUNITIES

In addition to using the best-in-class technology appropriate for the application (Table 4), additional techniques can be used to save energy. Even when a display produces light efficiently with LEDs or OLEDs, the greatest savings opportunity is likely related to reducing the amount of light produced. This power scaling can be accomplished with ambient light sensors, which adjust the screen brightness to the ambient light conditions, dimming the screen when ambient light levels are lower. Other power scaling techniques such as global or local dimming, used primarily in TVs, scale backlight power to content brightness. To the best of our knowledge, these techniques are not currently employed in small, secondary displays.

Although LCD panel power is generally much lower than backlight power, significant reductions in LCD power may exist for secondary displays because they often show static images, such as clocks or menus. Cutting power to the panel when the display is static may yield significant savings.

Standby power, in contrast to active mode power, is largely independent of technology, screen size, resolution, and brightness. Standby power for secondary displays is on the order of tenths of mW or less and can be eliminated completely by cutting power to the panel. Energy savings potential is not so much in reducing standby power as it is in increasing the amount of time in standby.

Power management using software and hardware techniques help increase the amount of time a device spends in standby mode. Perhaps the simplest technique is to put the display in standby/off mode after a set amount of time without user input. User input on a touchscreen or button, internal communication (e.g., the end of a countdown cycle), or external communication (e.g., a phone call) can activate the display. Devices can also use presence sensors to determine when a user is nearby to view the display, or in the case of cell phones, too close to the display to be viewing it.

Cell phones and other mobile devices are likely to optimize time in standby to prolong battery life. Devices for which a voluntary specification or mandatory standard based on total energy use has been developed such as printers and multi-function devices are also likely to use power management to cut power to displays when not in use. Of greatest concern are stationary applications that traditionally have not contained displays, but which have recently seen increased use of small color LCD panels, namely major appliances like refrigerators, dishwashers, and laundry equipment.

We have not considered power supply efficiency in the displays data above. For comparison purposes, Table 9 shows estimated ac input power for a 3.5 inch secondary display module by technology.[5]

Table 9: Estimated ac input power for a 3.5 inch display by technology type

	Dc Power Consumption (W)		**Ac Power Consumption (W)**	
Technology	Min	Max	Min	Max
CCFL backlit LCD	0.76	3.4	1.3	5.7
LED backlit LCD	0.25	2.4	0.42	4.0
Reflective and transflective LCD	0.000092	0.34	0.00015	0.57
OLED	0.19	1.0	0.32	1.7

In recent years, a trend away from CCFL to LED backlit LCD panels has increased energy efficiency. Based on our examination of datasheets described above, we believe the majority of secondary displays now produced use LED backlights. OLED displays have the potential to be even more efficient, especially for dark content.

We expect a shift in type and prevalence of secondary displays on appliances. Whereas alphanumeric displays may still be most common on appliances today, network connectivity and additional functionality of new generations of appliances including smart appliances will likely drive an increased prevalence of informational displays.

[5] We assume an ac/dc power conversion efficiency of 75% followed by an 80% efficient dc/dc stage for a combined efficiency of 60%.

6 AN UPDATE ON IR SENSING AND STANDBY

Infrared (IR) communication systems transfer data via infrared frequency radiation and are a mature technology widely used for short-range communications. IR technology is found in remote control systems for televisions, DVD and Blu-ray players and set-top boxes, laptop computers, tablets, cameras, printers and smart phones. IR transmitters are inexpensive, lightweight, reliable and easy to manufacture. In the last thirty years, IR has become commonplace within consumer electronics as the legacy technology of short-range wireless communication (Wells, 2013). IR remote controls have arguably led to consumer demand for wake-on-demand functionality in their electronics and resulting standby power modes.

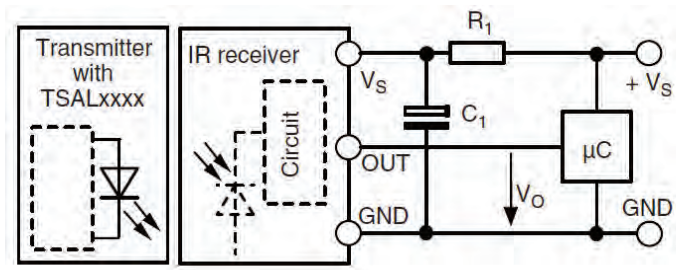

Figure 21: Block diagram of an infrared remote control system. (Source: Mouser Electronics, 2013)

In an IR remote control system (Figure 21), batteries power the transmitter. Energy efficiency, therefore, is likely a primary consideration in transmitter design. IR receivers are on the plugged-in side of the communication link (e.g., in a television or set top box) and may not be specifically designed to be energy efficient, thus we focus our evaluation on them.

IR receivers scan for an incoming signal based on a polling protocol. If unmanaged, this searching function dominates the power budget of the receiver and standby power of the device.[6] To reduce receiver power, most of the microprocessor may be powered down when no signal is being received. Upon detecting an incoming signal, the listening portion of the processor wakes the remainder of the processor. Using this technique can yield a 95% drop in power when the receiver is in standby (Lees and Schelle, 2006). If a typical receiver draws 3.3 mW when active (Gotschlich, 2010), it could draw 0.17 mW in standby by using this technique. Thus, power management can reduce standby power from IR receivers by an order of magnitude or more.

IR has one significant drawback to other wireless communication technologies: its inability to pass through solid objects. In the past decade or more, alternative wireless communication technologies that use radio frequency (RF) technology and thus eliminate the line-of-sight issue, such as Wi-Fi, Bluetooth, and Zigbee, have been developed. RF transmission has been used in some satellite TV receivers and

---

[6] Note that power supply losses associated with standby power also exist. We discuss power supplies, including architectures with an additional small power supply to reduce no- or low-load losses, in Section 7.

high-end stereo systems, remote controls for laptops and smart phones, garage door openers, and car key fobs (Layton 2005). These RF devices have their own disadvantages, namely issues with coexistence (i.e., interference) and network security. Zigbee RF4CE, a new communication protocol that reduces these concerns and networks devices in the home (Zigbee, 2013), has the potential to displace IR in remote controls, but is still a nascent technology. At this time, we expect IR to continue as the dominant remote control technology over RF, as it is a low-cost, mature technology with succinct communication protocols.

More and more, consumer electronic devices are able to be controlled not only with IR or RF remotes, but with smart phone applications via Wi-Fi. These devices must include and power the hardware to receive either signal. We expect this trend of multiple communication strategies on a single device to continue, and if not properly managed, power budgets will increase.

In summary, power draw of wireless communication receivers continues to contribute to standby power in many households and will continue to increase as the market develops additional devices with wireless communication capabilities. Increased standby power can be mitigated, however, if receivers use power management techniques.

7 POWER SUPPLIES FOR LOW STANDBY

This report has shown how many secondary functions can be made incredibly efficient, consuming mere milliwatts (or less) of power within their host devices. If properly integrated and orchestrated, devices like EEE-compliant Ethernet controllers and power-managed IR receivers enable extremely low standby power. However, all secondary functions must still receive some small amount of power in standby, and that power all ultimately passes through various power conversation stages.

Power supplies are the “gates” through which a device’s power must pass, converting high-voltage mains electricity down to lower dc voltages that can be utilized by sensitive electronic components. The highest efficiency power supplies today can achieve peak conversion efficiencies of over 90% when operated in their “sweet spot” between 50% and 100% of their rated output power, but efficiency drops off precipitously at lower load fractions where more than half of the incoming high-voltage power may be converted to waste heat. In standby applications, a device’s main power supply may be grossly oversized to meet the power needs of small standby loads such as IR receivers, resulting in low load fractions and very poor efficiency. This can add as a “multiplier” on standby power, doubling or even tripling standby power on the ac side of the plug. The effect is even greater when multiple power conversion stages (dc-dc power supplies) are required to ramp power down from intermediate dc voltages to even lower dc voltages required for certain integrated circuits.

The preferred power supply architecture for low standby power provides a highly efficient, appropriately sized secondary or standby power supply that operates during standby mode to supply small amounts of power to key secondary functions. This device has the capability to wake the main power supply when the end user wishes to turn the device on. For example, in a television a power on signal from the remote control would be picked up by its IR receiver, that would pass an on signal to the standby power supply, which would in turn energize the main power supply to wake higher levels of functionality in the product (e.g. screen, tuner, speakers, etc.). This staged wake-up process is illustrated in Figure 22. Power flows are indicted in red, and control information is shown in green.

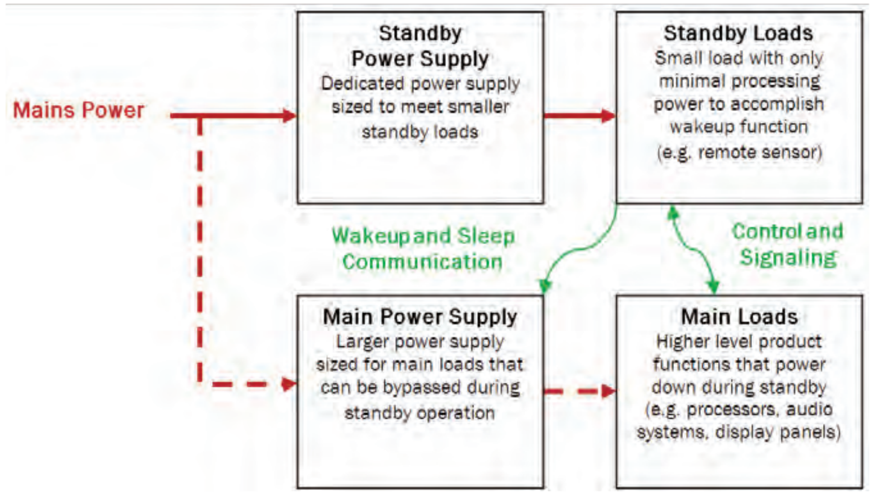

Figure 22: Block diagram for efficient standby power design (Source: APP/Ecova 2010)

Improper power supply design can have significant impacts on standby power and can negate the energy savings of properly designed, downstream secondary functions. Consider a television with an IR receiver, Wi-Fi transceiver, and Ethernet port. It might contain a 100 W ac-dc power supply that converts incoming mains electricity from 120 – 240 V ac down to 12 V dc. Most electronic components would require even lower voltages, so a second dc-dc stage might be required to ramp the 12 V dc down to 1 – 3.3 V dc. Figure 23 illustrates possible standby power flows to secondary functions for this TV. The power requirements of the secondary functions as well as the approximate efficiencies of the various power supply choices are representative of real devices, although not confirmed by actual measurements. The use of a dedicated standby power supply reduces losses through the ac-dc conversion stage by over 80% and cut overall ac standby power by more than half (0.7 W) compared to the design employing the TV's main power supply.

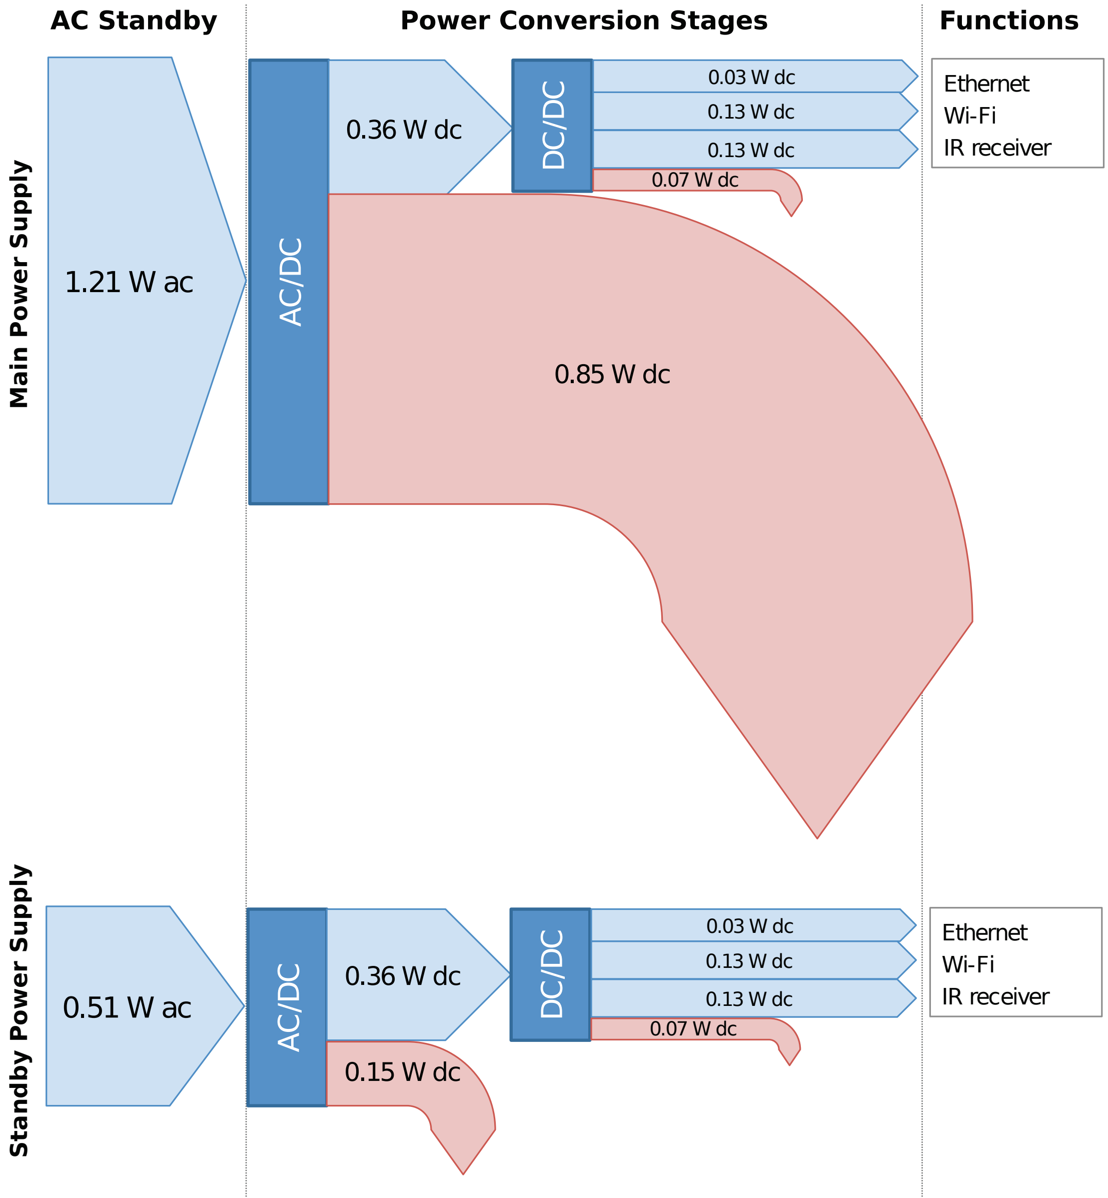

Figure 23: Power flow during standby for two TV power supply designs

8 CONCLUSION

8.1 KEY FINDINGS AND STRATEGIES FOR STANDBY MITIGATION

8.1.1 Advancing Power Management and Power Scaling

All of the secondary functions examined in this report are capable of extreme low standby power, in several cases at levels below 0.1 W ac (Figure 24). Display-related standby power can effectively be eliminated when information does not need to be displayed. However, those same functions can consume orders of magnitude more power if left active during low-power modes, as shown in the chart below. It follows that significant opportunity for mitigating standby power involves maximizing idle time for secondary functions in ways that do not impact the user experience.

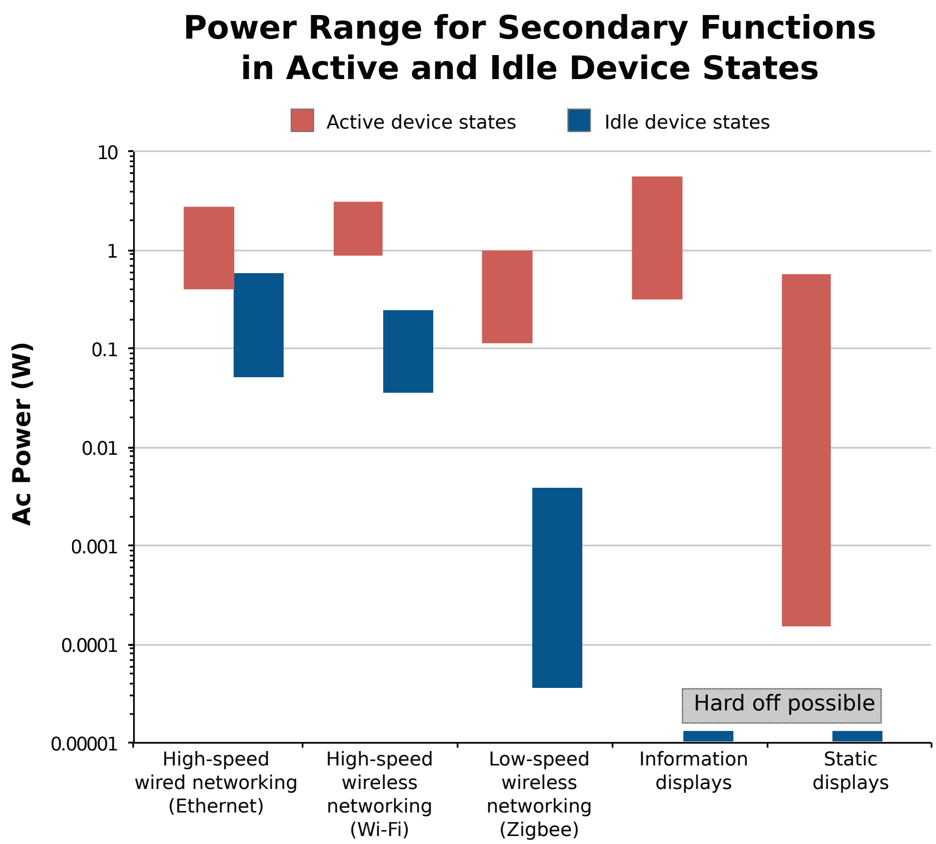

Figure 24: Power ranges across secondary functions

The ultimate strategies to mitigate standby power fall into the following two broad principles:

- **Power management — disable functionality when it is not needed**: Power management techniques transition end use products and their components between high-level operational states — such as "sleep" and "standby" in computers — on timescales of minutes. This coarse approach to modulating power in end products can have noticeable latency impacts to the end user, who must sometimes wait several seconds for a product to awake from a low-power state.

- **Power scaling — adapt a function's performance to current needs**: When functions must remain available to the user throughout low-power modes with low latency, that functionality should scale with actual user workloads. Component-level power scaling techniques like those used to extend battery life in mobile devices can occur at millisecond timescales with almost unnoticeable latency to users. Components like Wi-Fi radios and individual processor cores can be put into extremely low power states in a seamless manner that does not impact the user. In a sense, power scaling is power management, but operating on much shorter — even down to the millisecond — timescales.

Unfortunately, the former, more cumbersome power management paradigm is what currently dominates the design of stationary devices like plug loads and appliances. There are several reasons for this. For one, there is little incentive for OEMs to invest in greater power optimization beyond what is required by regulations because these products can be mains-powered all the time. Secondly, designers for many stationary products are saddled with legacy power management frameworks, like Advanced Configuration and Power Interface (ACPI) in computers, that are not nimble enough to be applied to power scaling. The organizations that have been most successful in achieving power scaling are mobile device OEMs like Apple, Samsung, and now Google. Their engineers developed a power-optimized integration between hardware and software that, prior to devices like smart phones and tablets, had not been achieved. This integrated design approach allows power optimizations at various "layers" of the design, from individual board-level components like Wi-Fi chipsets up to the application/software level itself.

A major challenge now is to bring the power-optimized design framework applied to smart phones and tablets to stationary products like consumer electronics, office products, and even traditional appliances. Organizations like California-based startup AGGIOS are making strides in this direction by developing design tools for stationary products that enable end use hardware designers to better coordinate the power scaling features of disparate components on the motherboard and optimize component-level power for different combinations of user tasks.

8.1.2 Interoperability Barriers

The research in this report was conducted by carefully examining manufacturer claims on components that drive standby power for certain secondary functions. At the component level, we see many examples of thoughtful power management features — like EEE or the power management provisions of 802.11 — that can enable best-in-class standby power values in end use products. However, implementation is absolutely critical to realize these savings mechanisms in reality. As an example, prior research demonstrated that many computers with EEE-compliant Ethernet controllers could not take advantage of their power management features simply because Ethernet driver options were set, by default, to *disable* EEE. When power management and power scaling functions are built into components like Ethernet controllers and Wi-Fi transceivers, it is incumbent upon end product designers to enable them by default, ensuring that the user's power management experience is automatic, seamless, and mostly invisible. Policies for end use devices with Ethernet functionality should simply require that EEE is enabled by default, with no need for intervention by the user.

8.1.3 Challenges and Opportunities for "Clusters" of Secondary Functions

Throughout our research we observed products in which several secondary functions are converging to provide what one might call a "smart" or "connected device" experience. For example, to make smart appliances truly smart, designers need to include much more than just networking functionality. The network function is simply the communication gateway for the appliance. To act on and display that information, smart appliances also require displays (typically large enough to enable a touch interface) and processors to determine how to handle incoming messages. Earlier generations of appliances may have only required a small, low-power logic circuit to register button pushes, whereas today's smart appliances effectively contain small computers.

This increase in functionality clearly presents challenges to reducing standby power. However, the presence of these "clusters" of functions also represents an opportunity to centrally manage power through use of inexpensive, low-power, system-on-chip (SoC) packages that integrate several functions into one part (e.g. Wi-Fi, display output, and general processing capability on a single piece of silicon). Savings can be achieved not only through better power management orchestration on a cluster of secondary functions, but also by reducing the number of microcontrollers and associated power overhead that could be needed to separately power-manage different functions. Although not common in today's market, representatives from Broadcom Corporation anticipate that SoC solutions for connected stationary devices like smart appliances could enter the market in the next two to three years as technology matures.

8.1.4 Standby Power Supply Design Still Matters

Finally, gains in energy efficiency at the component level can be lost through poor standby power supply design. Devices wishing to meet today's more aggressive standby power targets must necessarily incorporate a dedicated standby power supply to optimize power conversion efficiency during low-power modes. As demonstrated in Section 7 of this report, an inefficient power conversion strategy has the potential to more than double standby power use compared to the use of a dedicated, highly efficient standby power supply.

8.2 TRENDS TO WATCH

Our research suggests that the significance of network standby will only continue to grow. Several product trends bear mentioning:

- **Connected standby**: With the proliferation of connected devices and the increased availability of data networks at home, at work, and in public, many consumer products are being designed around a notion of near-constant network connectivity, even in low-power modes. Apple's "Power Nap" and Windows 8 "connected standby" features instruct newer computers to download emails, sync calendars, perform system updates, and even back up files all while the system is otherwise asleep. Power scaling and management techniques can minimize the impact to standby power by scaling power to data transfer speed in this connected standby state and returning the device to deep standby once the download is complete. In the end, however, the reality is that “active” network traffic during standby modes for these products will only increase.

- **Connected everything**: Network connectivity is proliferating rapidly into a large number of consumer devices where it never existed before. Consumer electronics are at the forefront of this shift, to be sure, but traditional appliances like refrigerators and laundry equipment continue to be revamped with “smart appliance” treatments. It will be important to ensure that these newer connected end uses employ network power management just as diligently as many more “mature” edge devices, such as computers and printers. More importantly, policymakers must encourage device manufacturers to employ networking technologies that are appropriate to the expected traffic. In smart appliances in particular, we do not need to build freeways (i.e. Wi-Fi) when country lanes (Zigbee) will suffice.
- **Speed**: Link rates for a variety of networking technologies continue to increase rapidly. In the wireless networking space, we are nearing a new age of gigabit wireless networking through the latest 802.11ac and 802.11ad protocols. Little is known about the standby power implications of these new technologies, particularly on stationary client devices. However, experience with Ethernet technology suggests that transmit/receive power will continue to increase as link rates go up and as devices require multiple radios to achieve the highest link rates.
- **More information, more displays**: With increased device connectivity and information transfer comes increased use of displays. In many “smart” devices, displays and networking functionality form a key “cluster” of secondary functions. We anticipate continued proliferation of displays in white goods that traditionally had few electronics.

8.3 KNOWLEDGE GAPS AND RECOMMENDATIONS

8.3.1 Knowledge Gaps and Recommended Research Activities

Our research uncovered two critical knowledge gaps that will require further research to help illuminate the path forward on network standby. With the emergence of next-generation gigabit Wi-Fi protocols, it is critical that the energy efficiency community better understand the power consumption impacts of these devices, both at the component and system levels. Unfortunately, manufacturer data on the subject is scant, and many of the manufacturers contacted only provided information on their Ethernet products. Even if more manufacturer data were available, little is known about the implementation of Wi-Fi power management features in more stationary edge devices, such as printers, IP set-top boxes (e.g. Apple TV, Roku), and televisions. We recommend the following follow-on actions to help close this gap:

- Continued manufacturer outreach, particularly to OEMs of end use products to understand how energy-saving features in Wi-Fi products are leveraged;
- Procurement and measurement of next-generation 802.11ac wireless networking products, including routers as well as edge devices to determine the impacts of increased link rates on power;

- Detailed, component-level measurement of Wi-Fi transceiver power consumption in stationary end uses to assess the impacts of different power management practices allowed by the 802.11 standard; and
- Investigation of the power implications of emerging high-bandwidth technologies like WiGig (802.11ad) that could significantly increase network standby associated with docking stations and home entertainment products.

A second and very broad gap exists in the realm of smart appliances. As many traditional appliances simultaneously gain networking and display functions, it will be important for the efficiency community to investigate the effectiveness with which power management settings are being implemented in the field and to determine whether power management strategies are being employed as aggressively as in other connected devices (e.g. computers, printers). Even though these devices may only represent a small fraction of sales today, it is advisable that the efficiency community continue to monitor the pulse of this market and ensure that appropriate, low-power wireless networking practices are encouraged through both horizontal and vertical efficiency policies.

8.3.2 Recommendations Toward a Standardized Policy Approach

One of the IEA 4E Annex's original objectives in investigating power requirements for key functions was to use these requirements to inform a horizontal policy framework.[7] This framework could be used to establish standby power targets for disparate classes of devices containing common secondary functions, such as networking and displays. The information provided in this report can inform horizontal power targets for low-power modes to any device containing the features. This is similar to the "functional adder" approach applied in many vertical energy efficiency specifications today, in which certain incremental functionality is allowed for limited additional power or annual energy consumption.

The single largest functional area not covered in this research relates to services that are typically provided by processors: timer/clock capabilities, access to memory, power management, registering user input, and delivering information to a display. Because processors are the "brains" that connect many other functions together and orchestrate power management and power scaling, it is crucial to capture their power requirements as well. Unfortunately, processors are highly application-specific, making them difficult to characterize in a generalized, horizontal manner as we have done with other functions. Processing capabilities in white goods could be provided by devices as simple as logic controllers, whereas consumer electronics would use more powerful processors for running software applications. Most products will have several specialized processors distributed around the motherboard, making it even more difficult to generalize the power requirements for this function.

Thus, a horizontal policy approach has its benefits but also its practical limitations, leading us to recommend an energy efficiency policy approach that ultimately combines horizontal and vertical

[7] Readers interested in a more detailed treatment of this framework are encouraged to review the work of Lloyd Harrington and Bruce Nordman (2010), who have written extensively on the subject. A recent proposal by AGGIOS to the California Energy Commission for horizontal Title 20 standards applicable to all consumer electronics is also instructive and is available at http://www.energy.ca.gov/appliances/2013rulemaking/documents/proposals/12-AAER-2A_Consumer_Electronics/.

considerations. Secondary functions that can be implemented in a similar way from product to product, such as those studied in this project, may be addressed in a horizontal manner (e.g. setting standby power budgets associated with each function). The information processing and control functions of a given end use device, however, are likely product-specific and still must be addressed vertically. In effect, horizontal considerations can form the foundation for low-power mode requirements, but may still need to be addressed in a vertical manner for large classes of products to capture implementation-specific considerations, as shown illustrated in Figure 25. Horizontal policies form a valuable foundation for policymaking, but our research suggests that vertical policy considerations cannot be avoided, even in regulating low-power modes.

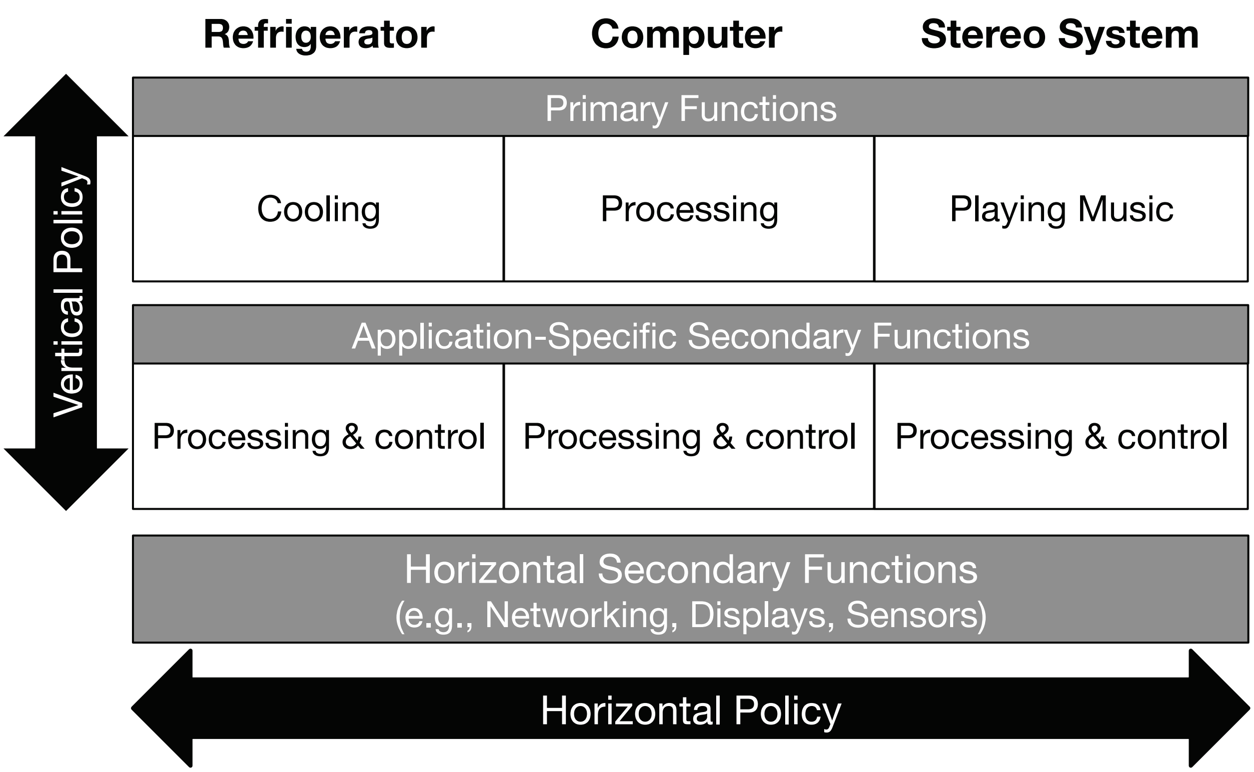

Figure 25: Examples of horizontal and vertical policy scopes

9 REFERENCES

Broadcom, 2012. "World's First 5G WiFi 802.11ac SoC," http://www.broadcom.com/docs/press/80211ac_for_Enterprise.pdf (accessed August 23, 2013).

Calwell, Chris, Jeffrey Swofford, Matthew Krick, Eric Wanless, David Denkenberger, 2011. *Power Scaling in Proportion to Data Processing*, Ecos.

Cisco, 2013. "Connections Counter: The Internet of Everything in Motion," Cisco Corporation, http://newsroom.cisco.com/feature-content?type=webcontent&articleId=1208342 (Accessed August 21, 2013).

Christensen, K, P Reveriego, B Nordman, M Bennett, M Mostowfi, J.A. Maestro, 2010. "IEEE 802.3az: the road to energy efficient Ethernet," *Communications Magazine*, Vol. 48, No. 11.

Davies, Charles, 2011. "LG THINQ Smart Appliances lower power bills & add remote control," *SlashGear*, http://www.slashgear.com/lg-thinq-smart-appliances-lower-power-bills-add-remote-control-03122496/ (Accessed August 21, 2013).

DisplaySearch, 2012. "Small and Medium Display Active Matrix OLED Penetration to More Than Double by 2015, According to NPD DisplaySearch," http://www.displaysearch.com/cps/rde/xchg/displaysearch/hs.xsl/121217_small_and_medium_display_active_matrix_oled_penetration_to_more_than_double.asp (accessed August 5, 2013).

DisplaySearch, 2013. "AMOLED Helps Samsung Display Become the Largest Small-Medium FPD Maker," http://www.displaysearchblog.com/2013/03/amoled-helps-samsung-display-become-the-largest-small-medium-fpd-maker/ (accessed August 5, 2013).

Drake, Jeff, David Najewicz, William Watts, 2010. *Energy Efficiency Comparisons of Wireless Communication Technology Options for Smart Grid Enabled Devices*, report, General Electric Company.

Eckermann, Ben, 2013. *Low-Power Network Standby in the Home and Office*, white paper, Freescale Semiconductor.

Foresman, Chris, 2012. "WiGig promises low-power 2Gbps wireless device communication by 2013," *ArsTechnica*, January 13, 2012, http://arstechnica.com/gadgets/2012/01/wigig-promises-low-power-2gbps-wireless-device-communication-by-2013/ (Accessed August 22, 2013).

Gotschlich, Martin, 2010. "Remote Controls –Radio Frequency or Infrared." Whitepaper, Infineon Technologies AG.

Gunaratne, C, K Christensen, B Nordman, S Suen, 2008. "Reducing the Energy Consumption of Ethernet and Adaptive Link Rate (ALR)," *IEEE Transactions on Computers*, Vol. 57, No. 4.

Gupta, Ashima and Prasant Mohapatra, 2007. "Energy Consumption and Conservation in WiFi Based Phones: A Measurement-Based Study," *4th Annual IEEE Communications Society Conference on Sensor, Mesh and Ad Hoc Communications and Networks*, June 18-21, 2007, San Diego, CA.

Hardy, Gregg, Dave Denkenberger, Jeffrey Swofford, Suzanne Foster Porter, Apurva Pawashe, Katherine Dayem, and Debbie Driscoll, 2013. *Small Network Equipment Codes and Standards Enhancement (CASE) Initiative For PY 2013: Title 20 Standards Development*, Ecova.

Hardy, Gregg, Jeffrey Swofford, Philip Walters, Debbie Driscoll, and Katherine Dayem, 2013. *Small Network Equipment in U.S. Homes: Using Less Energy to Connect Electronic Devices*, Natural Resources Defense Council and Ecova.

Harrington, Lloyd, and Bruce Nordman, 2010. *Standby Power and Low Energy Networks – issues and directions*, Energy Efficient Strategies.

IEEE, 2012. *802.3-2012 - IEEE Standard for Ethernet*, Institute of Electrical and Electronics Engineers.

IEEE, 2012. *802.11-2012 - IEEE Standard for Information technology — Telecommunications and information exchange between systems Local and metropolitan area networks — Specific requirements Part 11: Wireless LAN Medium Access Control (MAC) and Physical Layer (PHY) Specifications*, Institute of Electrical and Electronics Engineers.

Kundu, Bijit, Clancy Donnelly, Katherine Dayem, and Philip Walters, 2013. *Electronic Displays: Codes and Standards Enhancement (CASE) Initiative for PY 2013: Title 20 Standards Development*, Energy Solutions and Ecova.

Layton, Julia, 2005. "HowStuffWorks: How Remote Controls Work," November 10, 2005, http://electronics.howstuffworks.com/remote-control3.htm (accessed August 20, 2013).

Lees, David and Donald Schelle, 2006. "Infrared (IR) receiver circuit minimizes power consumption." *EE|Times*.

May-Ostendorp, Peter, Jeffrey Swofford, Clancy Donnelly, and Philip Walters, 2012. *Mapping Secondary Product Functions to Products and Operational Modes*, Ecova.

Mouser Electronics, 2013. "Newest Products: Vishay TSOP75D25 / TSOP35D25 IR Receivers," http://www.mouser.com/new/vishay/vishay3dirreceivers/ (accessed August 20, 2013).

OLED-Info.com, 2013. "DisplaySearch sees a $11 billion OLED market in 2013, shows OLED vs LCD price comparison," http://www.oled-info.com/displaysearch-sees-11-billion-oled-market-2013-shows-oled-lcd-price-comparison (Accessed August 5, 2013).

Pike Research, 2010. "Smart Appliance Sales to Start off Slow, but 118 Million Units will be Sold Worldwide by 2019," press release, Pike Research, http://www.pikeresearch.com/newsroom/smart-appliance-sales-to-start-off-slow-but-118-million-units-will-be-sold-worldwide-by-2019 (Accessed August 22, 2013).

Smith, Phill, 2011. *Comparisons Between Low Power Wireless Technologies: Bluetooth low energy, ANT, ANT+, RF4CE, ZigBee, Wi-Fi, Nike+, IrDA and NFC*, white paper, CSR plc, http://www.csr.com/sites/default/files/white-papers/comparisons_between_low_power_wireless_technologies.pdf (Accessed August 23, 2013).

St. John, Jeff, 2012, "Energate Launches Broadband Smart-Grid-to-Home Connection," *Greentech Media*, Jan 2012, http://www.greentechmedia.com/articles/read/energate-forges-the-broadband-smart-grid-to-home-connection-in-ontario/ (Accessed August 20, 2013).

TRENDnet, 2011. *TRENDnet Green Initiatives*, white paper, TRENDnet, http://trendnet.com/downloads/GREENnet_Initiatives.pdf (Accessed August 23, 2013).

Universal Display, 2013. http://www.udcoled.com/default.asp?contentID=605 (Accessed August 14, 2013).

Wells, Christopher J., 2013. "Communications Technologies - Infrared Communication Systems," http://www.technologyuk.net/telecommunications/communication_technologies/infrared_communication.shtml (accessed August 20, 2013).

Zigbee, 2013. Understanding Zigbee RF4CE. https://docs.zigbee.org/zigbee-docs/dcn/09-5231.PDF (accessed August 22, 2013