Electronic Devices & Networks Annex
EDNA

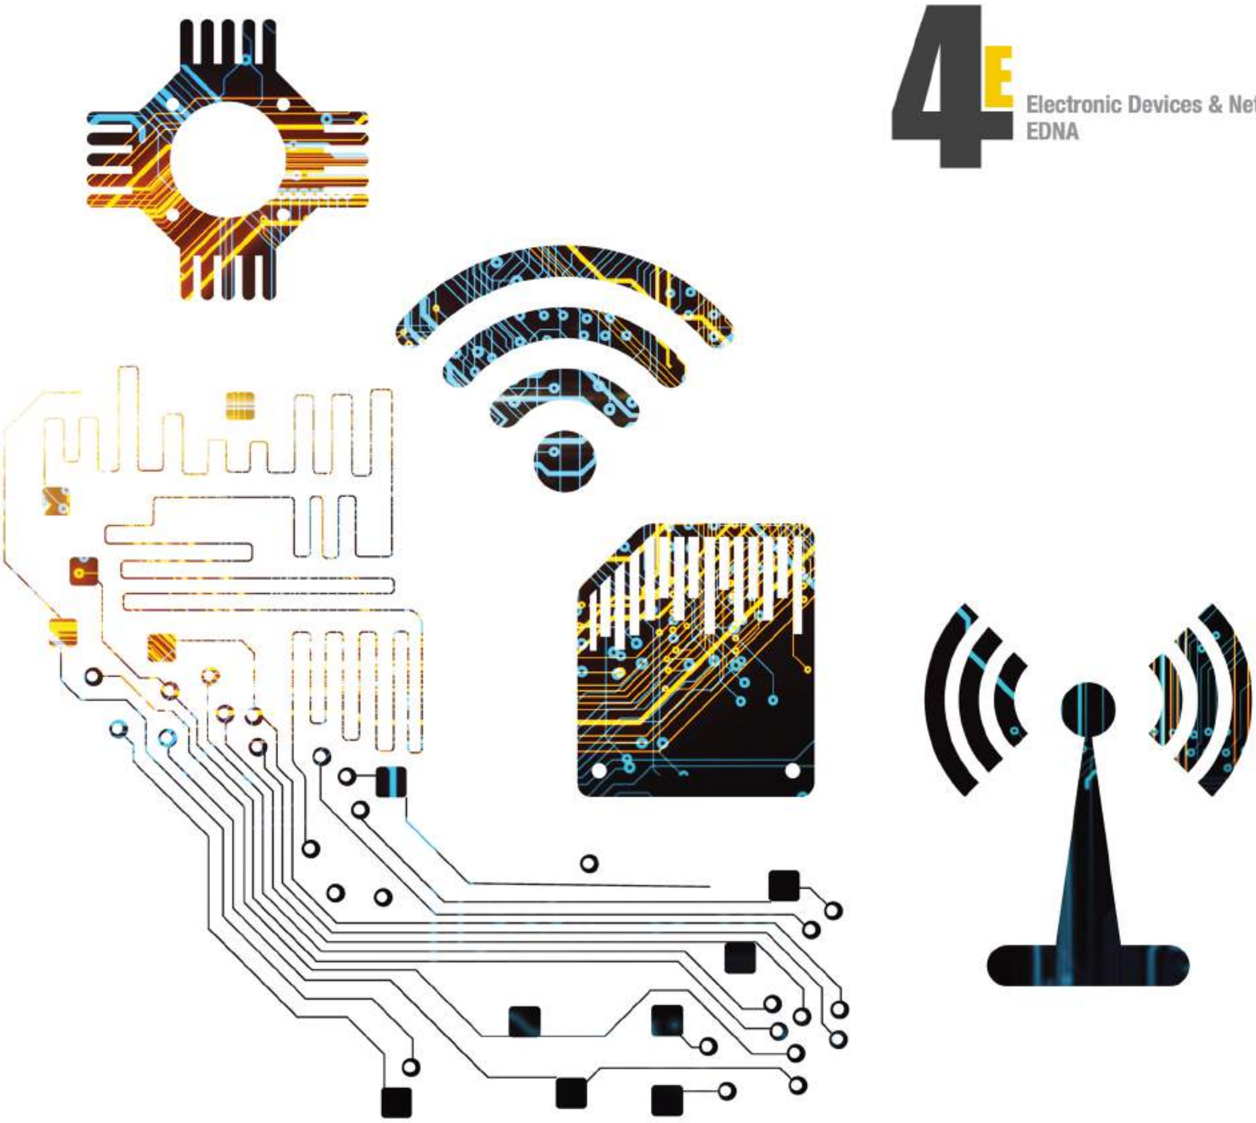

Network Standby Power Basics

Factors Impacting Network Standby Power in Edge Devices

APRIL 2018

4E is the Energy Efficient End-Use Equipment Technology Collaboration Programme, established by the International Energy Agency (IEA) in 2008 to support governments in co-ordinating effective energy efficiency policies. Twelve countries have joined together under the 4E platform to exchange technical and policy information focused on increasing the production and trade in efficient end-use equipment. However 4E is more than a forum for sharing information – it pools resources and expertise on a wide a range of projects designed to meet the policy needs of participating governments. Participants find that is not only an efficient use of available funds, but results in outcomes that are far more comprehensive and authoritative than can be achieved by individual jurisdictions.

Current members of 4E are: Australia, Austria, Canada, Denmark, France, Japan, Korea, Netherlands, Switzerland, Sweden, UK and USA.

Further information on the 4E Implementing Agreement is available from: **www.iea-4e.org**

Network connected devices, including the Internet of Things, are growing rapidly and offer enormous opportunities for improved energy management. At the same time, there is a responsibility to ensure that these devices use a minimal amount of energy to stay connected. 4E's Electronic Devices and Networks Annex (EDNA) works to align government policies in this area and keep participating countries informed as markets for network connected devices develop.

Further information on EDNA is available at: **http://edna.iea-4e.org**

This report is authored by Catherine Mercier, Katherine Dayem, Peter May-Ostendorp and Jason Wagner of Xergy Consulting.

NETWORK STANDBY POWER BASICS: FACTORS IMPACTING NETWORK STANDBY POWER IN EDGE DEVICES

Prepared by:
Catherine Mercier, LEED AP
Katherine Dayem, PhD
Peter May-Ostendorp, PhD, LEED AP O+M
Jason Wagner

Prepared for:
IEA 4E EDNA

27 April 2018

ACKNOWLEDGEMENTS

This study relied on input from a variety of industry and other stakeholders. The authors thank Hans-Paul Siderius (Netherlands Enterprise Agency), Bruce Nordman (Lawrence Berkeley National Laboratory), Miji Yang (Realtek), John Clinger (ICF), Klaus Verschuere (Cisco), Nathan Moin (HP), Ben Eckermann (NXP), Harold Jepsen (Legrand), David Chen (Power Integrations), and Vojin Zivojnovic (Aggios) for productive and enlightening conversations. However, the views expressed in this study remain solely those of the authors.

EXECUTIVE SUMMARY

The IEA 4E Electronic Devices and Networks Annex (EDNA) was established to improve the energy efficiency of network connected devices. As a growing number of products incorporate network connectivity, addressing energy use in network standby operation is an important and growing energy efficiency policy consideration. Network standby, in this report, represents a function that allows a product to maintain a network connection and await a network “trigger” or message to be woken as needed. This report incorporates and organizes our collective understanding of the network standby function, identifies the primary factors that influence its power draw, develops a framework for understanding these factors, and estimates power ranges associated with different network standby function implementations.

After a comprehensive literature review and dialogue with key stakeholders, we have identified three key factors that describe the power needs of the network standby function, comprising a network standby framework. They include: 1) the number and type of network interfaces, 2) the level of network interface functionality required during low-power modes, and 3) the type of network traffic management strategies that may be employed. The first consideration establishes, very simply, the amount of network componentry and its physical capabilities. The second consideration, network interface functionality, factors in the degree to which a device will need to maintain its presence on the network. The third consideration examines the strategies that network interfaces can use to ensure that they only fully wake for important network traffic.

Other product-specific factors and secondary functions outside of this network standby framework, such as power conversion or environmental sensing, will also impact the power draw of products in low-power, network-connected states. Product functions and features, such as on-board cameras and video streaming will also have design implications that could impact the selection or power management of network components.

The components required for network standby functionality are ultimately capable of consuming very low levels of power, although real-world power draw can be heavily influenced by the strategies employed in power-managing these components. Through a survey of common network interface components and various supporting electronics, such as memory, we estimate that most common network interfaces should consume less than 1.2 Wac (including power conversion losses), even when minimally power-managed. Values below 0.5 Wac are achievable when power management and duty cycling are taken into account. AC power consumption ranges by interface type are summarized in the figure below.

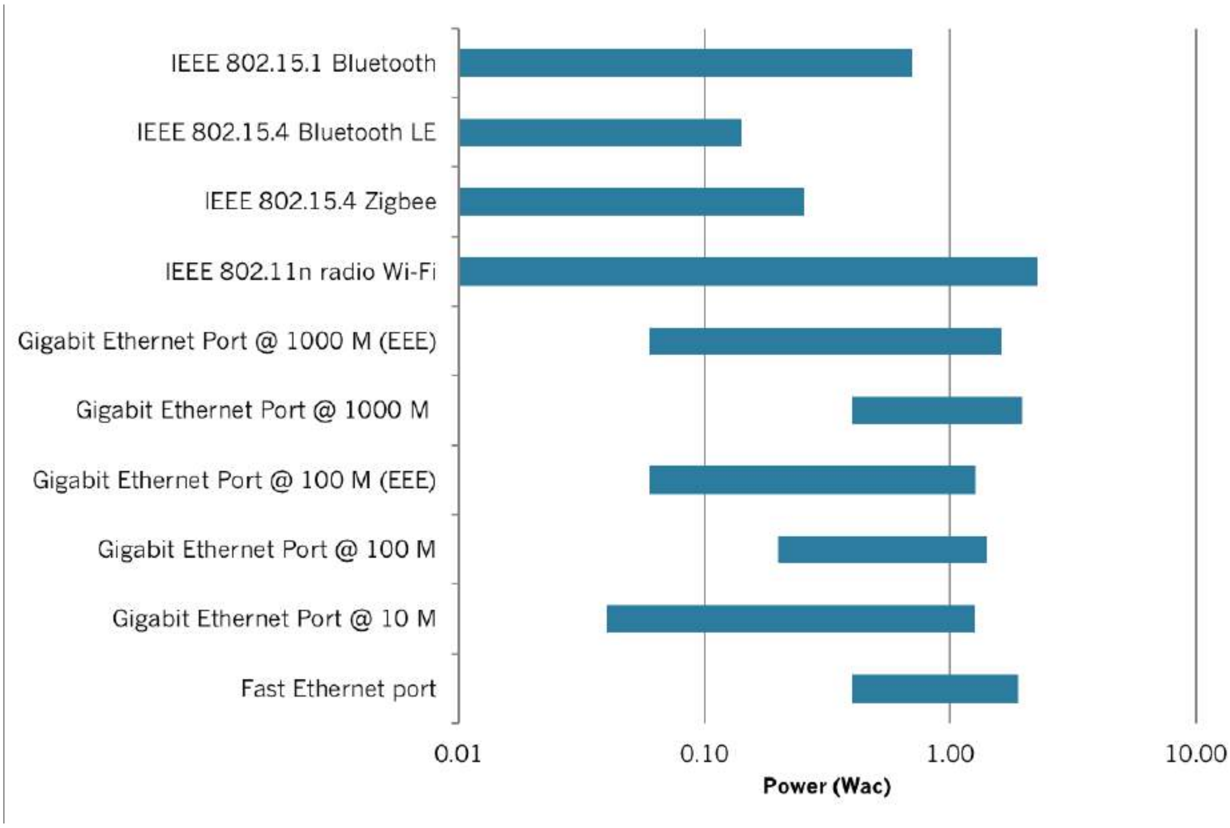

Estimated network standby function AC power draw ranges for various interface types. Note log scale on horizontal axis.

A selection of example end-use products illustrates that these components could constitute the vast majority of power draw in network-connected low-power states in many edge devices, such as over-the-top set-top boxes and office equipment. What is uncertain, however, is the extent to which component-level power management is deployed and optimized in mains-connected products. The duty cycling of these components effectively determines where in the power range presented above a design will operate. Future research efforts should examine the orchestration of network standby components and their power management strategies through product teardowns and in situ measurements to more precisely define achievable power draw ranges and savings potential.

TABLE OF CONTENTS

I. INTRODUCTION

PROJECT ORIGINS AND SCOPE

The IEA 4E Electronic Devices and Networks Annex (EDNA) was established to improve the energy efficiency of network connected devices. This report incorporates and organizes our collective understanding of the network standby function, identifies the primary factors that influence the power draw of the network standby function, develops a framework for understanding them, and estimates power ranges associated with different network standby function implementations.

The research team used a combination of secondary research and outreach to manufacturers and other network product experts to develop and refine this framework. This included a literature review and metastudy of prior research in the network standby field as well as a review of current hardware components associated with providing network standby functionality. This approach was useful to yield an initial framework, but should be followed with additional laboratory investigations and data gathering to more fully quantify and validate the framework.

This report focuses on digital networks and more specifically on devices that operate within the common TCP/IP protocol suite.[1] For the moment, we ignore other digital networks and serial data interfaces such as HDMI (High-Definition Multimedia interface, used to transfer digital audio and video signals between audio-visual equipment) and USB (Universal Serial Bus, used to transfer power and data between computer peripherals and certain consumer electronics). We also have explicitly excluded several product-category-specific network technologies, such as DOCSIS and MoCA (network technologies prevalent in set-top boxes). Finally, this report emphasizes network technologies commonly used in mains-powered equipment, and does not cover mobile communication technologies like 4G cellular networks or less common technologies, such as power line carrier. Other network types may be useful for understanding network principles in a more general way, but were beyond the scope of this project.

The proposed framework applies to edge devices (products that form the network endpoints) and not network equipment (products that serve as the intermediate nodes of the network and primarily route or forward network traffic to edge devices), because the network standby function, as we define it in this report, only applies to edge devices, in which network services are a secondary

[1] Transmission Control Protocol and Internet Protocol (collectively TCP/IP) are the backbone protocols used for communication on the Internet and are broadly supported in data networks for consumer and office products.

function. In network equipment, these services are a primary function of the product and may need to be continuously available.

Similarly, this report also excludes cable and satellite set-top boxes from its scope. These products have a special set of provisioning, security, and quality of service requirements that are less applicable to other common edge devices.

The following sections introduce some key terminology and define the scope of the network standby function. Section II introduces and discusses the key elements of the conceptual framework for network standby function power draw. Section III discusses additional hardware considerations outside the network standby function, including other secondary functions and power supply design, that impact product-level network standby power. Section IV outlines the range of network standby power for a variety of different implementations. Section V discusses project outcomes and provides suggestions on future research to further inform this framework.

TERMINOLOGY

Terms in the literature on network standby and network technologies are often used interchangeably. To avoid confusion, we use the following terms throughout this report:

Bandwidth: The maximum data transfer rate of a network link or connection.

Data transfer rate: The amount of data passing over a network link or connection per unit time.

Latency: The speed at which the equipment can provide its full functionality after having been triggered over the network (Nissen et al 2011, Task 1).

Network interface: A component (hardware and software) whose primary function is to make the connected product capable of communicating over one or more network technologies.

Protocol: A set of rules that allow two or more network-connected devices to transmit information.

Throughput: The observed rate at which data is sent through a link or connection.

NETWORK STANDBY FUNCTION

The network standby function allows an end-use product to resume other functions upon a remotely initiated trigger via a network connection (Nissen et al. 2011, Task 4). Because the network standby function can be available in any product mode (Nissen et al. 2011, Task 4), we are careful here to emphasize the characteristics that impact the network standby *function and its power*, rather than trying to comprehensively define standby and low-power *modes* in which a network connection is present.

We acknowledge that the network standby function is an idealized concept, and may not always represent how network functionality is implemented in products. Many of the stakeholders we interviewed emphasized that in practice, product functionalities are intertwined; the network standby function does not operate in isolation. A few stakeholders argued that the terminology itself was unclear and inconsistent with terminology previously used leading to confusion in the industry.

In this report, we investigate the characteristics that influence the network standby function as a first step to build understanding of the more complex reality: the network standby power draw of the whole product. Whole-product network standby power is influenced not only by the network standby function but also other functions that may need to be minimally available in low-power modes (Figure 1). How efficiently the product delivers and manages power, as well as the requirements of other functions (both the primary and any secondary), also impact whole-product network standby power.

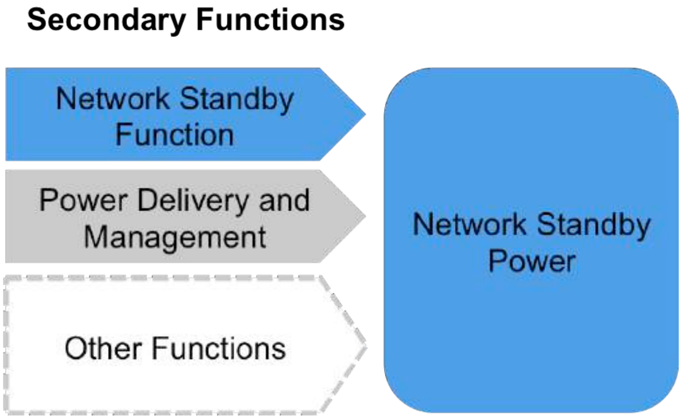

Figure 1: Network standby function and other characteristics including power delivery and management and other functions (primary and secondary) that impact network standby power.

The framework, discussed in Section II, thus focuses on the factors that affect the network standby function. Although not the main focus of this work, we also consider other product-related characteristics such as power supplies that do not impact the network standby function, but do impact network standby power (see Section III).

KEY HARDWARE COMPONENTS ASSOCIATED WITH THE NETWORK STANDBY FUNCTION

Implementation of the network standby function will vary from product to product, but includes common components to maintain the network connection and interpret signals from the network (Figure 2). All network connections require, at a minimum, an interface that can send or receive traffic through a physical medium (over wires or via electromagnetic signals), convert that traffic into a digital form (through a physical layer device or PHY), and then further ensure error-free transmission (through a media access controller or MAC). In general, a network interface contains a PHY and a MAC. In some cases (discussed below), the network controller screens packets for a wake-up signal. In other cases, some amount of intelligence is required to help interpret incoming packets and determine how to respond to them. In some designs, the product's main processor handles this work; in others, a dedicated microcontroller (MCU) handles network processing. Discussed in further detail below, MCUs are critical to using sophisticated, component-level power management strategies, such as proxying.

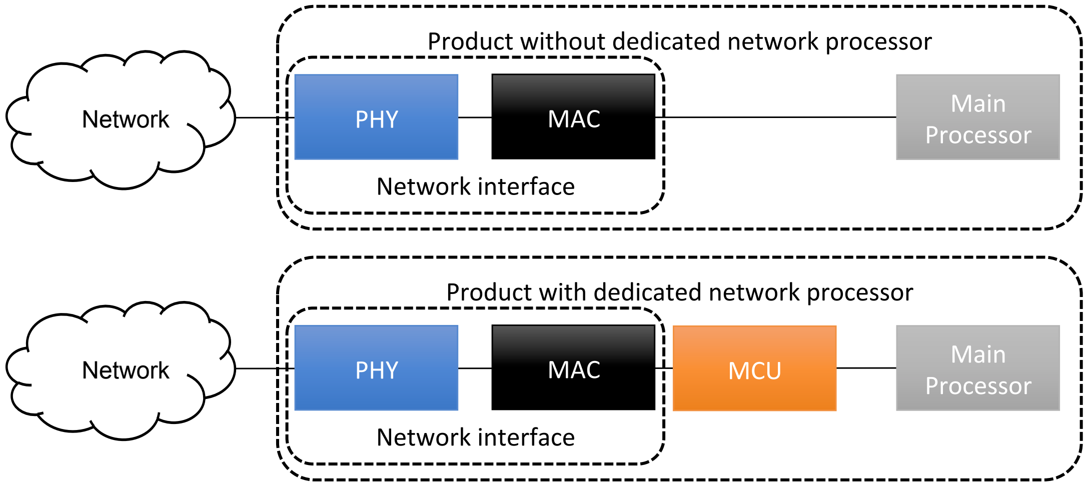

Figure 2: Block diagram of key components associated with network standby functionality.

II. CONCEPTUAL FRAMEWORK

Many researchers have studied the effect of network connectivity on power consumption (e.g., Nordman et al. 2009, Harrington and Nordman 2010, Viegand et al. 2017). To build on these efforts, we drafted a framework based on the present literature and discussed the concepts with industry and policy stakeholders. Based on these conversations, we refined the framework presented below.

Three factors determine the functionality delivered and power required by the network standby function: 1) the number and type of network interfaces, 2) the status (or level of functionality) of the network interfaces, and 3) network traffic management strategies. Though these three factors may be the proximate cause for power consumption by the network standby function, they are influenced by other design considerations, such as anticipated usage, user preferences, and requirements associated with a product's primary function (Figure 3).

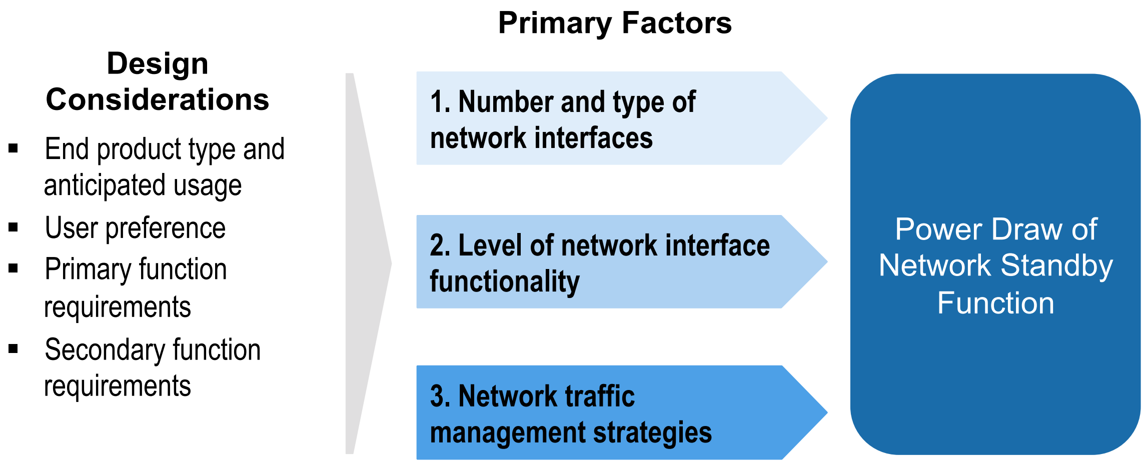

Figure 3: Conceptual framework of factors that impact the network standby function and its power draw. Primary factors are characteristics of the network communication hardware and functionality. Product design considerations often influence how network communication is implemented in a product.

1. Number and type of network interfaces

The network interface, a necessary component to provide connectivity, maintains a product's connection to the network and responds to incoming network traffic. The power consumed by network interfaces depends on the number and type of interfaces present in a product. A wide variety of network interfaces exist today, both wired and wireless. Wired network interfaces transfer data via traditional electrical connections over copper wires or fibre-optic cables, whereas wireless network interfaces use radio wave and even light signals (i.e. Li-Fi). A network interface is also defined by the network protocols it uses to guide the transmission of data across links for different physical configurations and link technologies, and the meaning of data transmitted. Although older network-connected products might have employed a single network interface, today's edge devices often include multiple network interfaces using different network technologies (Nissen et al. 2011, Task 4; Peterson and Davie 2012), for example Ethernet, Wi-Fi, and Bluetooth, potentially with significant impact to network standby power draw.

The primary factors that impact the power draw of network interfaces are related to a measure of the data transfer rate (either bandwidth or throughput) for wired and wireless interfaces, and additionally frequency and range for wireless interfaces. In general, higher bandwidth, throughput, and, for wireless interfaces, wider range, requires more power. Network technologies are available to mitigate the increased power draw by scaling power to throughput. For a detailed discussion of these considerations and how they impact individual network interface types, see Appendix A.

Power draw of network interfaces is not only dictated by a combination of the network technology standards implemented, but also the state of semiconductor technology today. These are not always within control of end product manufacturers and continue to evolve over time. The biggest uncertainty in future power draw for network standby functionality is the development of network technologies that could drive power use either up or down. If trends toward higher bandwidths for data-intense applications continue, for example, power draw will likely increase without improvements in semiconductor technology.

Although we aim to arrive at a framework that is technology-neutral and that derives from the underlying physical first principles driving power consumption, the reality is device manufacturers only produce (and OEMs only integrate) network interfaces based on a limited number of established network technologies, which we discuss in the following section.

Relevant Interfaces and Power Impacts

High-speed wired and wireless communication is predominantly achieved with Ethernet (IEEE 802.3) and various permutations of Wi-Fi (802.11), respectively (May-Ostendorp et al. 2013). Ethernet, one of the oldest interfaces and still the dominant wired one, covers a large family of physical layer interfaces, some historic, and many for specialized applications. A key feature of Ethernet is the many types of interfaces available that can provide different bandwidths (Nordman 2011). Products requiring high bandwidth can use interfaces that offer 1000 Mbps and beyond, while products with lower requirements may use interfaces that offer 100 Mbps.

High-speed wireless communication is dominated by Wi-Fi. Wi-Fi interfaces offer numerous bandwidths, frequency, and range for different applications. For the most part, increased bandwidth has been included in each successive 802.11 protocol; 802.11ac, released in 2013, is capable of 1,000 Mbps. The wireless market continues to evolve with 802.11ad for high-bandwidth, line-of-sight applications and 802.11ah for low-bandwidth, Internet of Things (IoT) applications, as well as emerging proprietary protocols. Similarly to Ethernet, this range of interfaces allows product designers to match product data requirements to an interface of appropriate capability.

Where high bandwidth is not a requirement of the application, lower speed wireless technologies, such as Zigbee, Bluetooth Low Energy (branded Bluetooth Smart), and Z-Wave may suffice (May-Ostendorp et al. 2013). IoT adoption is also driving a new wave of lower bandwidth, higher range, lower power protocols, including Bluetooth 5, Wi-Fi HaLow (802.11ah), and LoRaWAN. In addition, the newly-released IEEE 802.11ba standard allows the use of a dedicated, low power wake-up radio that can be duty cycled to greatly reduce average power (IEEE 2017). These technologies each have their own implications for power required by the network standby function.

Impact of network interfaces on network standby function power draw

Dominant network interfaces used currently can draw less than 0.5 Wdc. A recent study of four Ethernet interface products found an average power draw of 0.59 Wdc, and a minimum of 0.43 Wdc, achieved by two of the products examined (Viegand et al. 2017). May-Ostendorp et al. (2013) estimated power draw for Ethernet interfaces on the order of 0.2 - 0.3 Wdc per port. Wi-Fi network interface power ranges broadly, from less than 0.05 Wdc to several watts (May-Ostendorp et al. 2013, Viegand et al. 2017)[2]. Low-speed wireless technologies such as Zigbee can draw several orders of magnitude less power in idle device states, in the mW range (May-Ostendorp et al. 2013, Viegand et al. 2017). Emerging interface technologies may reduce network standby function power draw. The IEEE 802.11ba wake-up radio, for example, is expected to actively draw 0.1 mW or be duty cycled to draw an average power of thousandths of mW (IEEE 2017).

2. Level of network interface functionality

Network connectivity is not a single function, with a simple impact on power and energy use, but rather involves multiple conditions with different levels of functionality, ranging from a maintained link to full network presence. Some products, like smart clothes washers, may only require of the network connection the ability to receive a simple command to wake up and start a load of laundry.

[2] Note that these power ranges are for Wi-Fi stations only, and do not include wireless access points, which may be a topic of further research.

Other products, like VoIP phones and computers maintain higher functionality to be able to understand more complex wake-up signals that contain, for example, caller ID information. The network interface provides this functionality, which is related to the amount of participation in the network that the product conducts. Different network interfaces provide different levels of processing to provide the functionality required by the product, and in general, more processing requires more power.

Network interface functionality is frequently conceptualized using the Open Systems Interconnection (OSI) reference model, which partitions functionality into seven layers. Low levels of network interface functionality reside on the lower levels of the OSI Model (the physical and data link layers), whereas full network presence resides on higher layers. It follows that power draw increases when more layers and sophisticated protocols that require more processing are supported (Harrington and Nordman 2010). We note that high levels of network interface functionality may exceed the network standby function as defined, but many stakeholders stressed the need to move beyond examination of the network standby function in isolation and consider network functions in the context of the entire product rather than the network standby function concept.

Below we describe three basic states of network interface functionality ranging from low to high: (1) maintained link, (2) maintained link with wake-up signal capability, and (3) full network presence (Figure 4). Network interface power draw differs for each of these conditions; with increasing network capabilities, power and energy use typically increase as well (Harrington and Nordman 2010, Nissen et al. 2011).

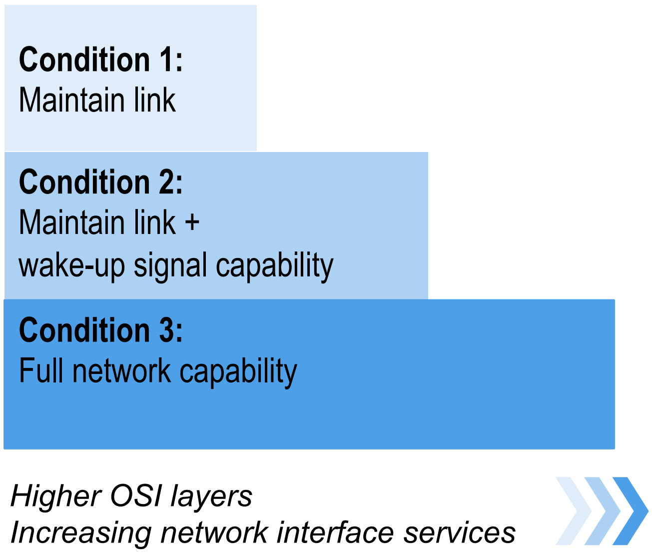

Figure 4: Three network functionality conditions, employing progressively higher OSI layers to provide more network interface services.

Condition 1: Maintain link or connection

A wired or wireless link is established, but no additional network services, including the ability to wake on a network signal, are provided. In this condition, only the physical layer is involved, and wake signals, such as via user input or a schedule, must come from the device itself rather than the network. This condition is the lowest level of network connectivity, but does not provide enough functionality for the network standby function because the interface cannot respond to wake-up signals from the network. The device can, however, wake the network standby function, and thereby enter Condition 2.

Condition 2: Maintain link + Wake up signal

Essential to the network standby function is the ability to receive and process the trigger (wake-up signal) via the network. Multiple wake-up processes exist; the ECMA-393 Standard, for example, identifies three: magic packet, wake on user datagram protocol (UDP), and wake on transmission control protocol (TCP) (ECMA International 2012). These processes fall into two distinct categories: (1) Address wake-up, in which a specialized wake packet targets a specific hardware address and (2) protocol wake-up, in which a sequence of packets corresponding to a protocol signal wake-up (Nissen 2011 Task 1).

Address wake-up is a simple process that uses the data link layer, the lowest of the network stack. Wake-on-LAN (WoL), for example, uses a signal broadcast to the entire LAN to wake a specific hardware address. To work, the network must know that the host is asleep and cannot respond to requests over the network. Instead, the sleeping host supports a wake-up mechanism that monitors messages received over the network for its specific wake-up message. On receipt of the wake-up message, the product wakes in order to respond to an incoming request or command.

WoL and other address wake-up implementations are simple, require minimal processing by the network interface, and draw minimal power: 0.5 Wac or less (Tanizawa et al. 2013, Nissen 2011). Address wake-up does, however, have several drawbacks. It requires a special interaction with the network because the product does not hide the fact that it is sleeping, and, because the only information that can be received is a wake-up signal, products may wake up frequently to processes minor information, or not often enough (Harrington and Nordman 2010). Consequently, address wake-up techniques are no longer commonly used, one component stakeholder noted.

Other wake-up processes are triggered by a sequence of events, defined by a protocol. These wake-up signals may contain more complex information in a sequence of packets and reside on higher network layers. The important distinguishing factor is that, in order to process a protocol wake-up signal, the receiving device must understand the overall protocol context in which the message was sent. A wake-up pattern for a VoIP phone, for example, might include caller ID information appended to the wake-up command. Depending on the complexity of the message and the sophistication of the network interface, the wake signal may need be handled by a processor (Figure 2): either a dedicated network microcontroller or the product's main processor (Siderius 2012). The more sophisticated the protocol (i.e., the more information it relays or the higher the OSI layers it invokes), the more data processing power is necessary to maintain or interpret the network connection.

Wake-up functionality is implemented using network interface hardware components in the host (Figure 2). Network interface circuitry that detects messages addressed to the sleeping host must stay powered on. That circuitry, upon detecting a message directed to the device, triggers the processor or other components of the host to wake up and process the packet. In other cases, the network interface may contain circuitry implementing essentially all the functionality of a network stack. In this scenario, the network interface itself may process received messages without waking other components of the device.

Condition 3: Full network presence

To deliver the expected service to the user (responsiveness/readiness), some products must deliver higher level functionality beyond network standby by maintaining full network presence. Imaging equipment, computers, and telephony, for example, often need to provide certain minimal network services, even when not providing their primary functions. A printer may have to respond to routine data queries on the printer's status that do not involve any actual imaging function. This allows a sleeping printer to still present itself as available on the network when a user wants to print to it. This can be accomplished with an active processor that interacts with the network when the rest of the product is asleep. Similarly, a computer may need to support a number of protocols and applications including responding to periodic Address Resolution Protocol (ARP) requests, maintaining its IP address by releasing periodic dynamic host configuration protocol (DHCP) lease requests, and responding to TCP-SYN (synchronize) packets sent to open ports. The ECMA-393 standard defines this concept as "full network connectivity," the ability of the computer to maintain network presence while in sleep and intelligently wake when further processing is required (ECMA International 2012).

Maintaining this level of network presence involves additional layers of the OSI model (Figure 4), additional active processing, and more power (Harrington and Nordman 2010, see Section IV). Thus a product’s network standby power draw depends on the level of network presence and functionality that it needs to respond to relevant information from the network, such as requests and packets, which can depend on its primary and secondary functions.

3. Network traffic management strategies

How a product manages incoming network traffic also has an impact on network standby power. Some products require some degree of active participation in the network; otherwise, applications fail and the product can not provide its expected functions. In these cases, the network interface may contain additional circuitry implementing all the functionality of the network reference model and allow the connected product or host to respond to some important requests, such as ARP requests or an IP version 4 (IPv4) Internet Control Message Protocol (ICMP) echo requests, while sleeping. Without any network traffic management strategies, a product continually wakes and sleeps, rather than spending significant time in a true network standby state, in which the product is simply waiting for an appropriate wake signal. Traffic management strategies employ a small amount of processing in the network interface itself, reducing the frequency with which other components need to be woken to respond to network traffic.

Three general classes of network traffic management techniques exist and allow the system to reduce the frequency of network wake events (Eckermann 2013):

- **Packet classification** allows systems to parse and drop packets that do not require a response.
- **Packet accumulation** buffers packets for response at a later point in time to extend the time the system can remain in a low power state.
- An **autorespond proxy** utilizes additional processing embedded in the network interface to respond to common, higher-level network protocols and requests without waking the entire product. A detailed discussion of proxy techniques can be found in (Nedevschi 2009). The ECMA 393 standard also provides an overall architecture for proxies and details on responding to a variety of key protocols.

In terms of hardware, packet classification and accumulation techniques can be accomplished by a modern network interface controller with some additional memory for buffering packets. Proxying requires an additional processing element capable of interpreting higher level network protocols. This additional processing could come in the form of a “smart” network interface (network interface with integrated processor) or through the processor on a system-on-chip (SoC) design. In either case, close integration of proxying functionality with the network interface is required.

Most Wi-Fi products can use a packet accumulation technique called Automatic Power Save Delivery (APSD), defined in IEEE 802.11e and required by IEEE 802.11n. APSD allows data to be buffered by the access point (rather than the client), and released when the client sends a message to do so.

Figure 5 conceptually illustrates how traffic management strategies might impact overall network standby power by altering the frequency and cycling of wake events. This cycling occurs on a small timescale (on the order of milliseconds); product power measurements are usually averaged over periods of five minutes or more, and thus is the average of these millisecond-scale wake cycles, which determine overall product power draw. Since traffic management techniques can lengthen the amount of time that the network interface spends in low power states, they directly impact power consumed by the network standby function.

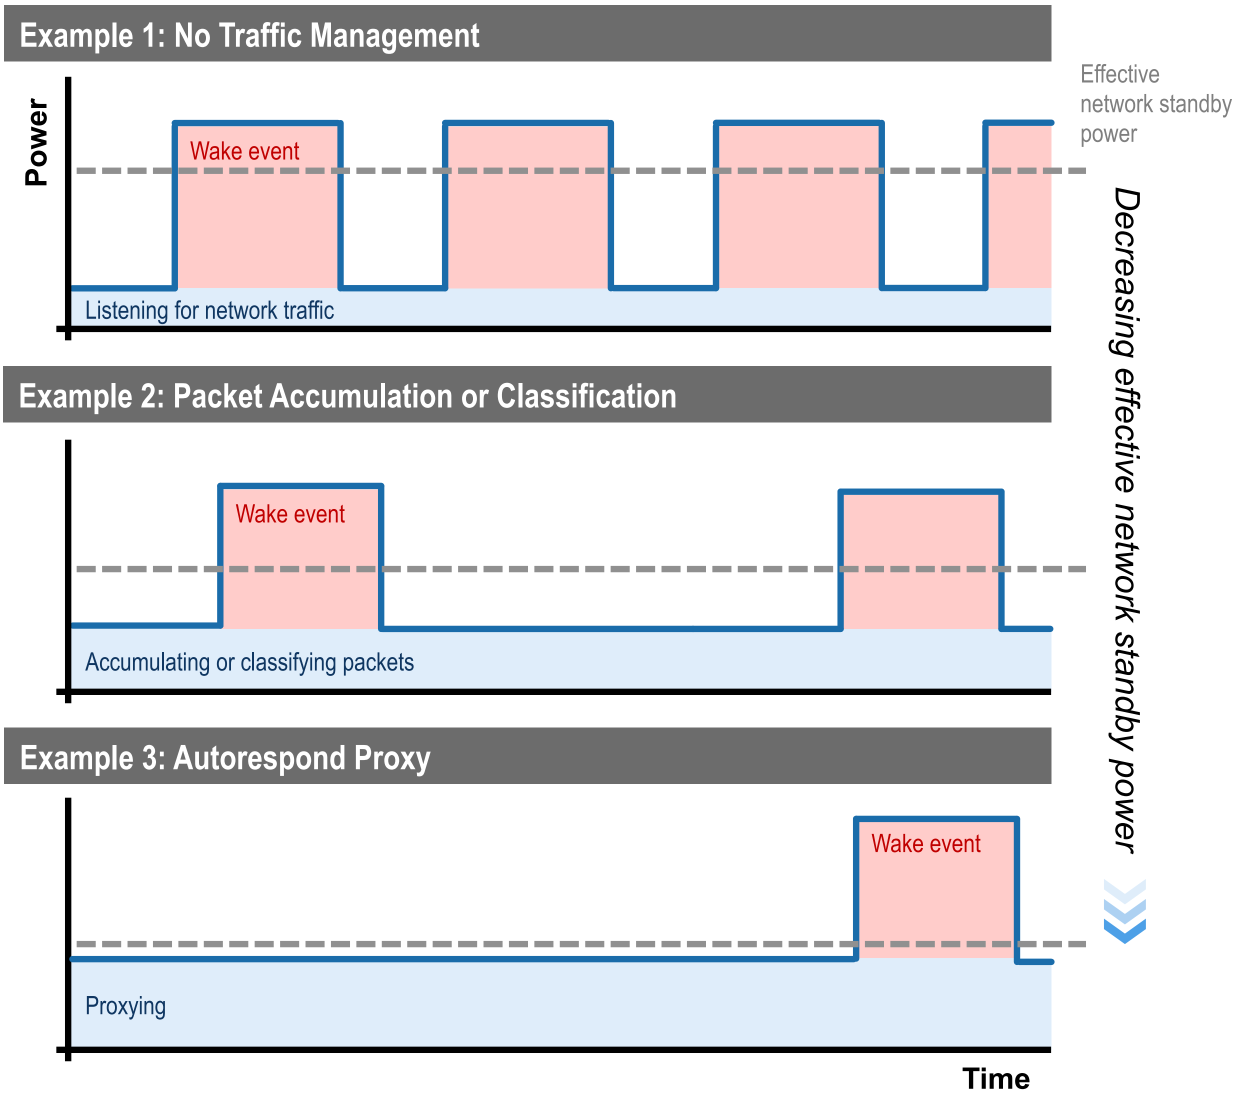

Figure 5: Illustration of relative power impacts of traffic management strategies. Overall time-averaged power draw can decrease as wake events are minimized through increasingly sophisticated traffic management strategies.

Products with no traffic management strategy (Figure 5, Example 1) may be capable of achieving very low power levels for brief periods while they are simply waiting for traffic, but may frequently wake other circuitry or the entire product to deal with spurious network requests (the extent of this problem depends on the amount of traffic on the device's local network). Network interfaces capable of packet accumulation or classification strategies (Figure 5, Example 2) may consume slightly more power while they are waiting for important network traffic, due their additional network standby functionality; they are not simply waiting for traffic but must buffer or filter that traffic. These strategies either delay handling spurious network traffic (accumulation) or ignore it (classification), thereby lengthening periods between wake events and lowering average network standby power levels. Care must be taken in implementing such strategies that high-priority traffic can be addressed and that the buffering process does not introduce unwanted latency. Finally, autorespond proxies (Figure 5, Example 3) employ additional processing capability to fully handle network requests. This can minimize or even eliminate wake events under the right conditions. On a quiet network, the power savings between classification and accumulation methods and autorespond proxies are similar; however, for real networks with increased traffic, the cumulative power savings can be an order of magnitude greater due to the ability of proxies to minimize wake events (Eckermann 2011).

Stakeholder outreach and industry literature confirm the significance of traffic management strategies as a factor in overall network standby power consumption. Eckermann (2013) analysed the power consumption implications of several traffic management strategies based on data for NXP's (then Freescale) network interfaces and processors. On a time-averaged basis, power values of 3.1 Wdc were achievable with classification, 0.44 Wdc with a combination of classification and buffering, and 0.15 Wdc with proxying (Eckermann 2011). Effective implementation of traffic management strategies may be one of the key technical pathways to achieving low overall network standby power while maintaining high levels of network service.

III. OTHER PRODUCT-RELATED ASPECTS THAT CAN INFLUENCE NETWORK STANDBY POWER

To the extent possible, we have discussed the network standby function in isolation to understand the basic components and service that impacts network standby power. In practice, however, additional, product-specific features impact network standby power, and are interrelated with the network standby function. Many of the stakeholders interviewed during this study noted that network functions are often linked to other functions (primary and secondary), both in their requirements for available processing, and in how power is supplied to the functions. The key task is to manage this link, so that power draw of both hardware and software is minimized when the product is not providing these other functions.

POWER MANAGEMENT AND OTHER FUNCTIONS

Product functions other than network standby impact network standby power both directly by drawing power and indirectly through the design decisions they influence.

Direct impacts

Secondary functions directly impact network standby power if they are not powered down, because of either poor design or latency requirements. A copier, for example, might need to maintain the temperature of its drum in order to run a copy job as soon as the user requests it, without having to wait for the drum to heat up. The information display on the copier, however, may be powered down while waiting for user input, because it can return to an active state with little latency. Well-implemented power management techniques power down functions when appropriate, while still providing useful network-enabled services (Harrington and Nordman 2011, Rozite and Siderius 2014), establishing a balance between power savings and efficiency on the one hand and resume times and functionality on the other (Siderius 2012).

Increased embedded processing capabilities, novel sensor solutions, and a drive toward IoT applications is increasing the prevalence of secondary functions in connected devices. Whereas older generation of products were not connected to a network, new products actively seek connection with other products, including “the cloud”, on the network. This trend can make the implementation of proper power management techniques more challenging. Consumer electronics and appliances are increasingly designed with multiple secondary functions like information displays, instant-on memory, and sensors (a smart phone or tablet contains more sensing capability than the rest of the building in which it resides). When secondary functions are poorly implemented from a power management standpoint, they can remain on unnecessarily in low-power modes and drive up network standby power, even if not directly related to network services (May-Ostendorp et al. 2013). Secondary functions in mobile products are managed aggressively to preserve battery power, but similar solutions in mains-connected devices has lagged behind.

Indirect impacts

Functions other than network standby also have indirect impacts on network standby power because they influence design decisions such as the network technology, network interface, bandwidth, and range. If the main function of a product requires a high speed network connection, for example, the network standby function will use this connection. Of course, strategies exist to mitigate power draw of the high speed connection, such as power management to scale power to speed.

POWER SUPPLY DESIGN

The standby power function itself can operate on a very small and efficient power budget; however, that power must still be provisioned and converted through power electronics, usually at some additional loss.

Most products today have one primary power supply, which may be housed within the product enclosure (internal power supply) or outside of the product in a separate enclosure (external power supply). In network standby, a device's main power supply may be grossly oversized to meet the power needs of small network standby loads, resulting in low load fractions and very poor efficiency (on the order of 50%). This effect is even greater when multiple power conversion stages (dc-dc power supplies) are required (Siderius 2012, May-Ostendorp et al. 2013).

One suggested power supply architecture for low network standby power incorporates a highly efficient, appropriately sized, dedicated network standby power supply that supplies small amounts of power to support network standby and other functions when products are in low power modes[3]. This device has the capability to wake the main power supply when additional power is required for the primary or other secondary functions. Power electronics manufacturers such as Power Integrations and Texas Instruments offer such devices to enable very low no-load power in external power supplies, but the same concepts have been recently used in a computer internal power

[3] One product manufacturer has claimed that a single power supply design is also capable of high efficiency at very low loads.

supply prototype as well (Delforge 2016). When appropriately sized, losses through the dedicated standby power supply stage can be limited to tens of milliWatts, according to one stakeholder we interviewed.

Research also continues on power supply architectures that would eliminate losses and enable devices to power down to virtually zero, eliminating the need for a mains connection in low-power modes. This may be possible by using a capacitor to provide the small amounts of power needed to monitor for network or other wake-up signals. This concept is currently being explored in a California-funded research project at Lawrence Berkeley National Laboratory headed by Alan Meier. It has also been explored and implemented by office equipment manufacturer Brother in inkjet printer models (Brother 2017). At present, it is unclear whether either of these efforts will fully apply to network standby or merely traditional standby functionality.

Finally, Power-over-Ethernet (PoE) and USB Power Delivery are becoming more popular options to deliver power to some consumer electronics that require less than 100 Wac of power in their active modes. PoE is commonly used in office equipment such as VoIP, security devices, and wireless access points. USB Power Delivery may see increased use in certain computer peripherals. These solutions eliminate the need for an ac-dc power supply, but they still may require further dc-dc power conversions within products themselves.

Impact on network standby function power draw

Power supply design may have a significant impact on network standby power, but calculating energy consumption in absolute values is difficult because individual device power draw varies so dramatically. May-Ostendorp et al. (2013) considered, for example, a television with a Wi-Fi transceiver and Ethernet port. They found that 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. Two interviewed stakeholders, however, argued that with today's technologies the power supply design should only have a small impact on network standby power.

Initial demonstrations of zero-standby solutions indicate that such solutions can be implemented and that it is possible to maintain network functionalities while letting devices power down to zero watts (Tanizawa et al., 2013).

Preliminary energy consumption estimates of dc systems vary depending on converter efficiencies and study type (modelled or experimental) (LBNL 2017). More research needs to be done to determine whether the end-to-end conversion losses for dc-powered products (from ac sources down to the product) would be smaller than conversion losses in traditional ac-powered products.

IV. A RANGE OF NETWORK STANDBY POWER

Through discussions with stakeholders and review of product datasheets, the study team identified the physical components within a product that provide network standby functionality. The main components investigated included:

- Network interfaces (encompassing both physical and link layer functionality in the OSI network model)
- Controllers/processors needed to interpret and respond to remote triggers, namely a microcontroller unit (MCU) and supporting memory and oscillators
- Power supplies, including ac-dc and dc-dc power conversions

We developed power ranges associated with different network standby function implementations, and applied these estimates to several actual products (client devices, not network equipment itself) to establish potential ranges for network standby function power draw. These ranges are a first step in understanding how much power is needed to provide network functionality, but we strongly encourage future laboratory tear-down research of products to better understand the power implications and management of various components in situ.

KEY COMPONENTS ASSOCIATED WITH NETWORK STANDBY POWER

As noted in Section I, the network standby function can be implemented in a variety of ways, but includes several common components to maintain the network connection and interpret signals from the network, including a network interface and processing on either a dedicated or system processor (Figure 2). To estimate network standby power draw, we consider additional components required by network interface, including an additional oscillator or clock chip for timing as well as some separate memory, depending on the extent to which these components are integrated (Figure 6).

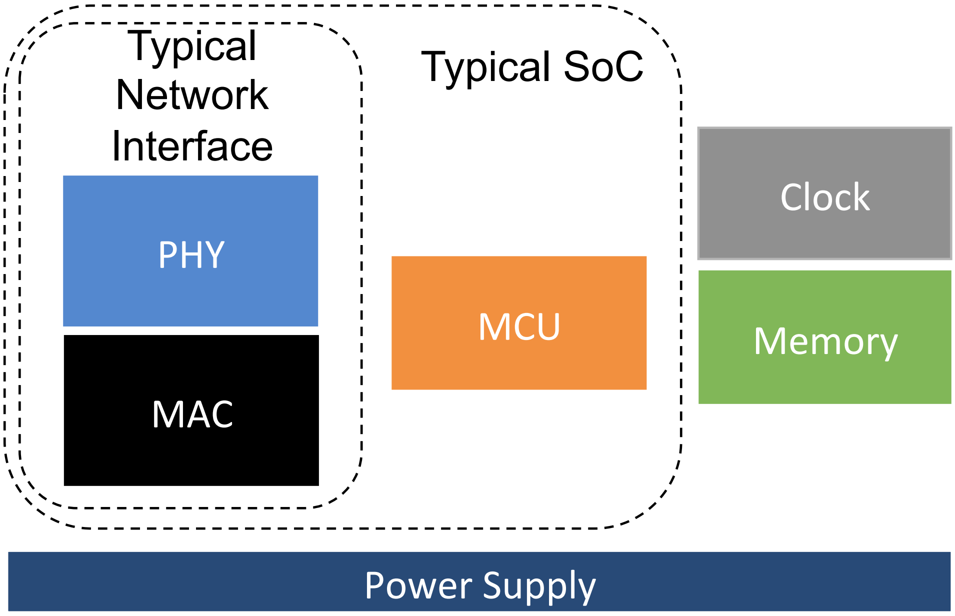

Figure 6: Block diagram of key components associated with network standby functionality. PHY: physical layer, MAC: media access controller, MCU: microcontroller unit.

Many connected products require additional processing capabilities in the form of an MCU to help interpret incoming packets and determine how to respond to them. MCUs are also critical in establishing more sophisticated, component-level power management strategies, such as proxying, because they can provide enough intelligence to determine which packets require wake-up and which can be ignored. MCUs are increasingly integrated alongside network functionalities on system-on-chip (SoC) designs, which may simplify orchestration of power states compared to systems built up from multiple discrete components. In cases where a local MCU is not employed, processing may be handled by a product's main processor. As with the network interface components, additional oscillators and memory may be required, depending on the level of component integration.

Finally, a power supply is required to provide discrete components with the appropriate dc voltage (Figure 6). Although many products of interest will contain an ac-dc power supply at a minimum, they frequently require dc-dc converters to regulate the low voltages necessary for certain MCUs, memory, and other devices. These dc-dc converters are, in some instances, integrated into larger, integrated SoC designs.

The ultimate network standby power depends not only on the underlying functional and hardware requirements, but also the *orchestration* of those hardware components between different power states. Designers may program their applications to wake the MCU and network interface on a frequent basis to achieve low network latency in some products (e.g. a consumer device receiving push notifications from the cloud). The same hardware might be programmed to spend more time idle for less latency-sensitive applications (e.g. a smart appliances), achieving lower average power with similar hardware.

POWER RANGES

To begin to characterize the range of power required by the network standby function in and across products, we consider the power (in Wdc) related to the components in Figure 6. A range of device power states is available to the network interface and the MCU, thus a range of power draw. The low end of the range is defined by the lowest state in which a network link can still be maintained. On the high end of the range, the network interface can be actively communicating and the associated MCU can be actively processing incoming information. The actual network standby function power draw lies somewhere between these extremes.[4] Figure 7 illustrates the concept for the network interface and the microcontroller. For the network interface, the power range will vary with maximum link speed or bandwidth; for the MCU, the power range will vary with increasing processor clock speed.

Currently achievable power draw ranges for key hardware components associated with the network standby function are provided in Tables 1 and 2.[5] Table 3 provides an overall summary of total network standby function power consumption based on interface type, including a breakdown of power conversion losses in dc-dc and ac-dc stages. The interfaces shown are the most common for consumer and office products and have the best supporting data. The ranges represent power draw for the fundamental network standby components referenced above, based on datasheets from and conversations with component manufacturers. In Tables 1 through 3, “low” values are based on the lowest power sleep states that can provide network standby functionality. “High” values are those reported for active mode, and represents the worst-case scenario of providing network standby with active components that are not power managed. We exclude power draw from other components not associated with network standby functionality (e.g. sensing/metrology, indicator lights, or remote control/local signal processing).

[4] Both the high- and low-power device states may persist for brief periods of time—measured in milliseconds—during behavior that otherwise looks like network standby. Averaged over longer periods of time and over a representative duty cycle—measured in seconds or minutes—we can realize the actual network standby power of the product.

[5] Emerging technology and standards, such as the newly-released IEEE 802.11ba wake-up radio (see Appendix A), are not considered in this analysis.

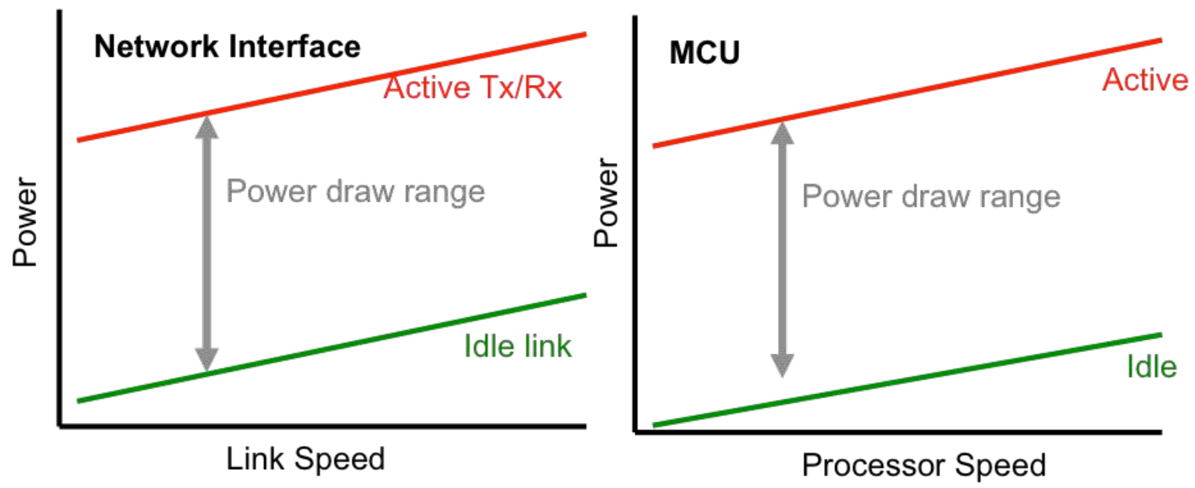

Figure 7: Conceptual illustration of reported power ranges for network interfaces and MCUs.

Note that the low and high values in the tables do not represent the power range of any one specific model, but rather the range found across the components examined, whose applications range from mobile and embedded to desktops and servers. For example, for a Gigabit Ethernet port using EEE at 100 Mbps, a MAC + PHY for mobile and embedded applications could draw as little as 31 mWdc if it utilized a low-power state, while on the high side, an active MAC + PHY for server applications draws 373 mWdc.

Several sources and assumptions factor into our range estimates. First, we assume that products include a dedicated standby power supply, limiting power supply losses to tens of milliWatts (based on personal communications with Power Integrations). Power estimates for Ethernet components are based on conversations with Realtek and datasheets from NXP, Intel, and Microchip.[6] Power estimates for wireless interface components are based on Texas Instruments, NXP, and Microchip datasheets. The table provides additional notes on data sources for other components. For a full listing of manufacturers and products referenced and links to any referenced datasheets, see Appendix B.

[6] Some manufacturers may choose to deliver power using a single-stage power supply, which could decrease conversion efficiency in low-power product states.

Table 1: Range of dc power draw for network interface components by device state

<table>
<tr><th rowspan="4">Network Interface[1,2]</th><th colspan="6">Power by Device State (mW dc)</th></tr>
<tr><th colspan="2" rowspan="2">Integrated/SoC Interface Components[3]</th><th colspan="4">Discrete Interface Components</th></tr>
<tr><th colspan="2">PHY + MAC[4]</th><th colspan="2">MCU[5]</th></tr>
<tr><th>Low</th><th>High</th><th>Low</th><th>High</th><th>Low</th><th>High</th></tr>
<tr><td>Fast Ethernet port</td><td colspan="2" rowspan="6"></td><td>243</td><td>560</td><td rowspan="6">9</td><td rowspan="6">367</td></tr>
<tr><td>Gigabit Ethernet Port @ 10 M</td><td>21</td><td>381</td></tr>
<tr><td>Gigabit Ethernet Port @ 100 M</td><td>98</td><td>373</td></tr>
<tr><td>Gigabit Ethernet Port @ 100 M (EEE)</td><td>31</td><td>373</td></tr>
<tr><td>Gigabit Ethernet Port @ 1000 M</td><td>247</td><td>612</td></tr>
<tr><td>Gigabit Ethernet Port @ 1000 M (EEE)</td><td>32</td><td>612</td></tr>
<tr><td>IEEE 802.11n radio Wi-Fi</td><td>0.8</td><td>1030</td><td>2</td><td>315</td><td>0.3</td><td>83</td></tr>
<tr><td>IEEE 802.15.4 Zigbee</td><td>0.0001</td><td>111</td><td colspan="4" rowspan="3"></td></tr>
<tr><td>IEEE 802.15.4 Bluetooth LE</td><td>0.006</td><td>36</td></tr>
<tr><td>IEEE 802.15.1 Bluetooth</td><td>0.7</td><td>405</td></tr>
</table>

Notes and Sources:

1. Power estimates for Ethernet network components gathered from conversations with Realtek and datasheets from NXP, Intel, and Microchip. Power estimates for wireless interface network components gathered from Texas Instruments, NXP, and Microchip datasheets.

2. For most interface types, total network component power is the sum of power requirements for the PHY, MAC, MCU clock, power supply, and memory. Zigbee, Bluetooth LE and Bluetooth are generally implemented with system-on-chip solutions, and we were unable to find information on the individual network components. "Low" column represents lowest power found for an idle link and MCU (or SoC). "High" column represents highest power found for an active link and processor.

3. Source: Microchip datasheets.

4. Sources: Realtek personal communications, Intel and Microchip datasheets.

5. Data gathered for 400 MHz MCUs. Source: Texas Instruments datasheets.

Table 2: Dc power draw ranges of network interface supporting hardware components

<table>
<tr><th rowspan="2">Component</th><th colspan="2">Power by Device State (mW dc)</th></tr>
<tr><th>Low</th><th>High</th></tr>
<tr><td>Clock or Oscillator[1]</td><td colspan="2">0.15</td></tr>
<tr><td>Serial Flash Memory[2]</td><td>0.01</td><td>50</td></tr>
</table>

Notes and Sources:

1. Data reported for 32 kHz crystal. Source: Texas Instruments.

2. Although other memory may be used, serial flash appears to be typical configuration. Source: Microchip datasheets.

Table 3: Dc and ac power draw ranges for network standby function based on network interface used.

Network Interface	DC Component Power (mWdc)		DC-DC Conversion Losses (mWdc)[1]		AC-DC Conversion Losses (mWdc)[2]		Total Network Standby Component Impact (Wac)	
	Low	*High*	*Low*	*High*	*Low*	*High*	*Low*	*High*
Fast Ethernet port	252	977	13	49	93	462	0.4	1.5
Gigabit Ethernet Port @ 10 M	30	798	2	40	11	377	0.04	1.2
Gigabit Ethernet Port @ 100 M	107	790	5	40	39	373	0.2	1.2
Gigabit Ethernet Port @100 M (EEE)	40	790	2	40	15	373	0.06	1.2
Gigabit Ethernet Port @ 1000 M	256	1029	13	51	94	486	0.4	1.6
Gigabit Ethernet Port @ 1000 M (EEE)	41	1029	2	51	15	486	0.06	1.6
IEEE 802.11n radio Wi-Fi	3	1478	0.2	74	7	698	0.01	2.3
IEEE 802.15.4 Zigbee	0.2	161	0.01	8	7	76	0.01	0.2
IEEE 802.15.4 Bluetooth LE	0.2	86	0.01	4	7	41	0.01	0.1
IEEE 802.15.1 Bluetooth	0.9	455	0.04	23	7	215	0.01	0.7

Notes and Sources:
1. Assumes 90-95% conversion efficiency. Source: Rhom Semiconductor and Digital Europe.
2. Assumes dedicated standby power supply with efficiency of 55-65% and no load losses of 7 mW. Source: Power Integrations Design Evaluation Report 623.

The network interface and MCU components represent the vast majority of power draw associated with the network standby function, assuming that designers have taken care to implement a dedicated standby power supply with low losses. Most network interfaces can achieve power levels in the tens of mW if the product allows the network interface and MCU to power down to their low-power idle states. Products that do not take advantage of these lower power states can draw upwards of 1 Wac to power network standby components. The specific orchestration of the interface and MCU between active and idle states will drive ultimate power consumption of the network standby function. This orchestration directly relates to the level of functionality required of the interface and the traffic management strategies the designers employ (factors 2 and 3 in the proposed framework for network standby functionality).

APPLICATION TO PRODUCT EXAMPLES

To illustrate how this data might be used, we apply it to a handful of mains-connected products by comparing whole-product network standby power measurements to our bottom-up estimates of the network standby function in isolation (Table 4). Lacking information on exactly how network standby power management strategies are implemented in products, it is impossible to accurately estimate the power draw of the network standby function. The ranges provide our best estimate of the power needed for the network standby function based on the information available today. The examples presented in Table 4 are for illustration purposes only; they each apply to one specific product model and do not represent the product category as a whole.

Table 4: Measured, product-level standby power compared to estimated power required by network standby components.

Example product	Network Interface(s)	Product-level network standby power (W)	Power required by network standby function components (W)	Notes and Sources
Inkjet printer	Gigabit Ethernet Wi-Fi	2.2	Ethernet: 0.04-1.2 Wi-Fi: 0.01-2.3	1,2
Desktop computer	Gigabit Ethernet Wi-Fi	1.6	Ethernet: 0.04-1.2 Wi-Fi: 0.01-2.3	3,4
Speakers	Wi-Fi Bluetooth	2.1	Wi-Fi: 0.01-2.3 Bluetooth: 0.01-0.5	5
Smart dishwasher	Wi-Fi	2.0	Wi-Fi: 0.01-2.3 Bluetooth: 0.01-0.5	6
Display	Fast Ethernet	1.8	Ethernet: 0.3-1.2	7
Over-the-top STB	Fast Ethernet Wi-Fi	0.3	Ethernet: 0.3-1.2 Wi-Fi: 0.01-2.3	4

Notes and Sources:

1. For example, see the Epson WF-8090: https://files.support.epson.com/docid/cpd4/cpd41778.pdf
2. Data from US EPA ENERGY STAR Qualifying Products List for imaging equipment
3. For example, see Dell Optiplex 960: https://www.dell.com/downloads/global/products/optix/en/desktop-optiplex-960-technical-guidebook-en.pdf
4. EuP Preparatory Studies Lot 26, Networked Standby Losses, 2011.
5. L. Kaufmann, R. Kyburz, Network Connected Audio Products, Measurements and Analysis of Network Standby Consumption, 2016
6. J. Viegand, B. Huang, L. Maya-Drysdale, Review study on Standby Regulation, 2017.
7. Data from US EPA ENERGY STAR Qualifying Products List for displays

For example, the inkjet printer is a model on the ENERGY STAR qualified products list, with a reported standby power of 2.2 Wac. This particular model contains both Ethernet and Wi-Fi network technologies, and the range of ac power required to provide the network standby function, from Table 2, is listed in the fourth column. The ENERGY STAR test procedure for imaging equipment requires that only one network technology is active during testing, and prioritizes Ethernet first. It follows that standby power for this model was measured with an Ethernet connection, and those network components can use anywhere from 0.04 to 1.2 Wac, according to our estimates in Table 3. Therefore, the network components could account for a small fraction to more than half of the product's total standby power, depending on their implementation.

As described in previous sections, the range of network standby power of a product depends on: (1) the number and type of network interfaces employed, (2) the level of functionality of the interface, and (3) any traffic management strategies employed. Without more precise knowledge of network standby implementation, we can only conjecture on its contribution to overall product standby power in these products. Note that network standby could constitute a majority of standby power for the given products if not well managed. Other significant areas for attention could include power supply design and power management of other secondary functions.

V. CONCLUSIONS

Within the scope of the network standby function, we found three primary factors that influence network standby power: (1) the number and type of network interfaces, (2) the level of network interface functionality, and (3) the use of traffic management strategies. Factors 1 and 2 determine the quantity and type of network interface hardware used, the physical media over which data will be transferred, and which communications protocols will be supported. These choices are dictated, to some degree, by product design considerations, such as the product's primary function, consumer preferences, the network environment in which the product is intended for use, and so on. They broadly establish the range of power that might be required for network standby functionality. The third factor, traffic management strategies, represents hardware and software design choices that determine how to orchestrate the various hardware components involved with network standby functionality. These choices can also have a significant impact on power consumption.

In examining overall product power draw in network-connected low power modes, network standby is just one of several functions that may be present. At a minimum, devices will need to provision power to components, but they may also need to power sensors, indicator lights, memory, and other components. Although this project could not comprehensively address all of these possible functions, we have identified both direct and indirect means by which these functions might influence overall network standby power values. Our research also found that the contributions of power supplies to network standby power overall should be very small (on the order of tens of mW) if efficient, dedicated power conversion stages for low-power modes are used.

Our analysis shows that the network standby function and associated hardware components can be a significant contributor to overall network standby power in a variety of edge products today, possibly accounting for over half of power draw. The network standby function's influence on product-level power draw, however, can still vary over a wide range depending on power management and control strategies. Most components we examined as part of this effort are capable of extremely low-power states, but based on current product-level network standby power values, we can infer that these components do not currently spend significant time in such states. In short, the extent to which mains-powered products utilize low-power device states or whether they implement traffic management techniques to minimize wake events is worthy of further investigation.

To illuminate these issues, the energy efficiency community must invest in detailed laboratory investigations of networked products. Future work should ideally investigate the operation and power draw of network interface components in situ to: (1) more tightly quantify contributions to overall product power draw and (2) assess opportunities for reducing the power implications of the network standby function through strategies described in this research. It will also be useful to continue documenting product-level network standby power for a variety of fast-growing, vertical product categories, such as smart lighting, smart televisions, and network-connected speakers.

VI. REFERENCES

Brother International, 2017. Low Energy Standby. http://www.brotherearth.com/en/story/standby.html. Accessed April 14, 2017.

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

Delforge, P., 2016. Slashing Energy Use in Computers and Monitors While Protecting Our Wallets, Health, and Planet. White paper published by Natural Resources Defense Council. https://www.nrdc.org/sites/default/files/slashing-energy-use-computers-monitors-ib.pdf.

Eckermann, B., 2013. Low-Power Network Standby in the Home and Office. White paper published by Freescale Semiconductor.

ECMA International, 2012. ECMA 393: proxZZZ for sleeping hosts. 2nd edition.

Harrington, L., and B. Nordman, 2010. Standby Power and Low Energy Networks – issues and directions. Prepared for APP and IEA 4E Standby Annex.

EPA, 2017. ENERGY STAR Product Specification for Set-Top Boxes, Eligibility Criteria Version 5.0.

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.

IEEE, 2017. IEEE Technology Report on Wake-Up Radio: An Application, Market, and Technology Impact Analysis of Low-Power/Low-Latency 802.11 Wireless LAN Interfaces. Preview copy.

LBNL, 2017. Direct current as an integrating and enabling platform. Presentation prepared by LBNL, EPRI, CIEE and Arup for TAC Meeting #2.

May-Ostendorp, P., K. Dayem., J. Swofford and D. Thomsen, 2013. Power requirement for functions. Prepared for International Energy Agency 4E Standby Power Annex.

Nedevschi, S., Chandrashekar, J., Liu, J., Nordman, B., Ratnasamy, S., and N. Taft, 2009. "Skilled in the Art of Being Idle: Reducing Energy Waste in Networked Systems." USENIX Symposium on Networked Systems Design and Implementation (NSDI).

Nissen, N.F., 2011. EuP Preparatory Study Lot 26. Networked Standby Losses. http://www.ecostandby.org/. Berlin: Fraunhofer IZM.

Nordman, B. and K. Christensen, 2010. Proxying: the next step in reducing IT energy use. Published by the IEEE Computer Society. https://www.computer.org/cms/ComputingNow/homepage/2010/0210/W_CO_ProxyingTheNextStep.pdf

Nordman, B, 2011. Testing products with network connectivity. Prepared for the Australian Government, Department of Climate Change and Energy Efficiency. http://standby.iea-4e.org/files/otherfiles/0000/0034/nordman_project_C_jun21-m-b.pdf

Nordman, B., H-P Siderius, L. Harrington, M. Ellis and A. Meier, 2009. Network connectivity and low-power mode energy consumption. http://energy.lbl.gov/ea/nordman/docs/eedal-2009-nordman.pdf.

Peterson, L and B. Davie, 2012. Computer networks: a systems approach. 5th Edition. Elsevier, 884 p.

Rozite, V. and Siderius, H-P, 2014. Making smart products and systems efficient. 2014 ACEEE Summer Study on Energy Efficiency in Buildings.

Siderius, H-P, 2012. Networked Standby: A wake-up Call for 1 W. 2012 ACEEE Summer Study on Energy Efficiency in Buildings.

Tanizawa, Y., T. Ishihara, Y. Oyama, T. Murakami, T. Kogawa, E. Kamagata, and K. Ise, 2013. Zero-Watt Networked Standby: Development and Evaluation of a Home A/V Network System. Applications and the Internet, vol. 21, no. 23, pp. 468-480, 2013, doi:10.2197/ipsjjip.21.468.

Viegand, J., B. Huang, and L. Maya-Drysdale, 2017. Review study on standby regulation. Draft final report. Prepared by Viegand Maagøe for the European Commission. http://www.ecostandbyreview.eu/index.html

VII. APPENDIX A – FIRST PRINCIPLES INFLUENCING NETWORK INTERFACE POWER

BANDWIDTH AND THROUGHPUT CAPABILITIES

Network interfaces can operate at different data transfer rates depending on configuration, implementation or physical characteristics of the network connection (Nordman 2011). Generally, the faster the interface transfer rate, the more power required to prepare and transmit, or receive and interpret information (Nordman 2011). Two aspects of data transfer rate can impact network standby power: bandwidth, the capacity or maximum speed of the link, and throughput, the actual amount of traffic on the link.

In well-designed products, bandwidth decisions should be driven by the thoughput required by the application that is using the network connection (Harrington and Nordman 2010). Bandwidth and throughput requirements of today's products are trending in two directions. Increasing throughput demands on networks for video streaming, gaming, and high-performance applications has led toward higher bandwidth network technologies in certain media-intensive edge devices. At the same time, penetration of "smart" or IoT devices continues to increase, and many of these devices require minimal bandwidth to periodically transfer data to and from the cloud. In response, low-bandwidth, long-range wireless technologies are being introduced in some IoT products, such as 802.11ah (Wi-Fi HaLow) and Bluetooth Low Energy.

Whether bandwidth or throughput impacts network standby power draw depends on whether or not the network interface supports a network technology that includes power management strategies that scale power to the amount of traffic on the network link. If the network interface does not include power scaling capacity, network standby power is related to bandwidth, and varies little with throughput (Harrington and Nordman, 2010). Unmanaged Ethernet links, for example, are fully powered even when no data is transferred. In this case, power draw for the network standby function is the same as for an active connection, and increase by about an order of magnitude between bandwidths of 1000 Mbps to 1GBps (Figure 8).

When devices are in a low power state (and for much of the time when they are in an active state), however, the amount of data being communicated is usually small compared to the capacity of the link, so techniques to save energy during periods of low utilization (data throughput) can have significant impact on network standby power. A range of techniques, dictated by link technology standards, exist that allow a link to go to sleep for short periods of time when utilization is low, adjusting the device capability to more closely match the end use data related demands at any point in time, and reducing energy use. IEEE 802.3az, commonly referred to as Energy Efficient Ethernet (EEE), for example, allows the link to rapidly sleep and wake up (Christensen et al. 2010), so power is proportional to data throughput for links where both network interfaces support and implement EEE (Figure 9). As of 2013, most Ethernet controllers on the market supported EEE (May-Ostendorp et al. 2013). Most current wireless technology standards include power scaling techniques, therefore data throughput requirements rather than bandwidth influences active power consumption (Paterson and Davie 2012).

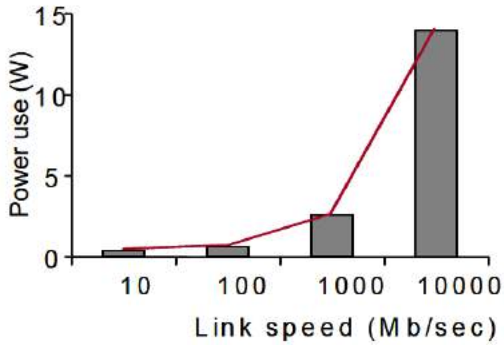

Figure 8: Power as a function of bandwidth for unmanaged Ethernet interfaces. Source: Nordman et al. (2010).

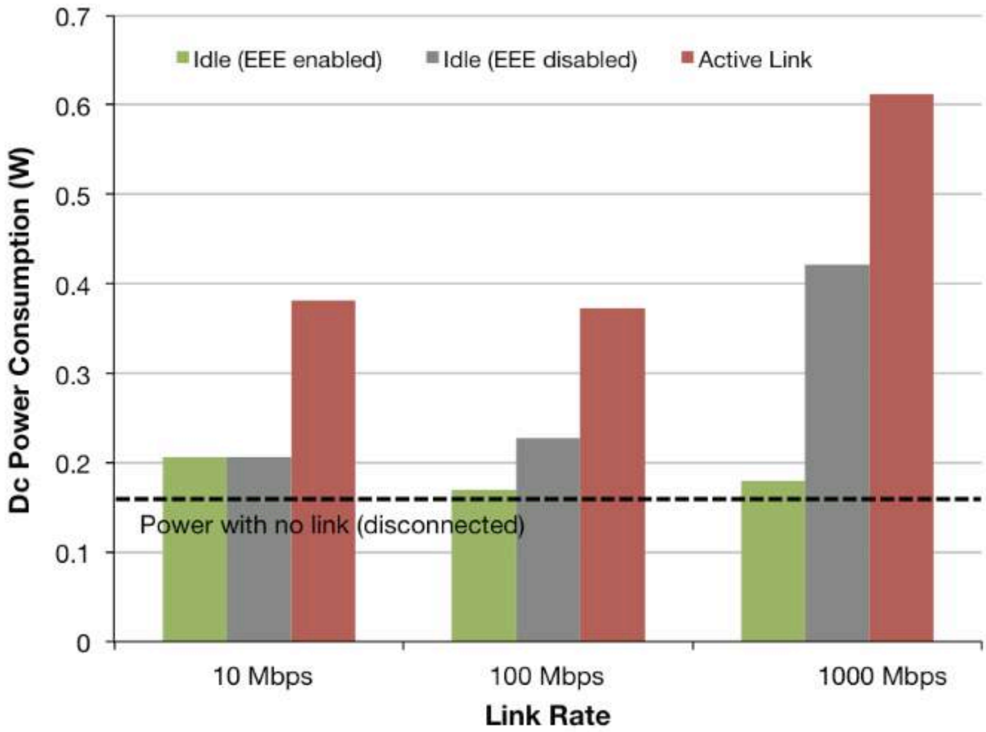

Figure 9: Power as a function of link rate or bandwidth for a Gigabit Ethernet controller. Source: May-Ostendorp et al. (2013).

FREQUENCY AND RANGE

Frequency and range are unique to wireless communication. Early Wi-Fi network standards operated at 2.4 GHz, but starting with the IEEE 802.11ac standard that added coverage in the 5 GHz band, more frequency choices have become available both at higher and lower frequencies. The WiGig or 802.11ad standard operates on the 60 GHz band and will provide peak data transfer speeds of 7 Gb/s. On the other end of the spectrum, the recently developed 802.11ah allows low frequency (less than 1 GHz) communication. Bluetooth, Bluetooth Low Energy, and Zigbee operate on the 2.4 GHz band.

As discussed above more generally, manufacturers and their engineering teams select network technologies based on several product design requirements and need to make tradeoffs between bandwidth, range, and power efficiency. Generally speaking, high frequency technologies like 802.11ad translate into high bandwidth, high power consumption, but low range (802.11ad is essentially a line-of-sight technology and cannot penetrate walls). Lower frequency technologies like 802.11ah provide lower bandwidth, lower overall power consumption, but dramatically broader range. Figure 10 illustrates this concept for a variety of points in the 802.11 standard.

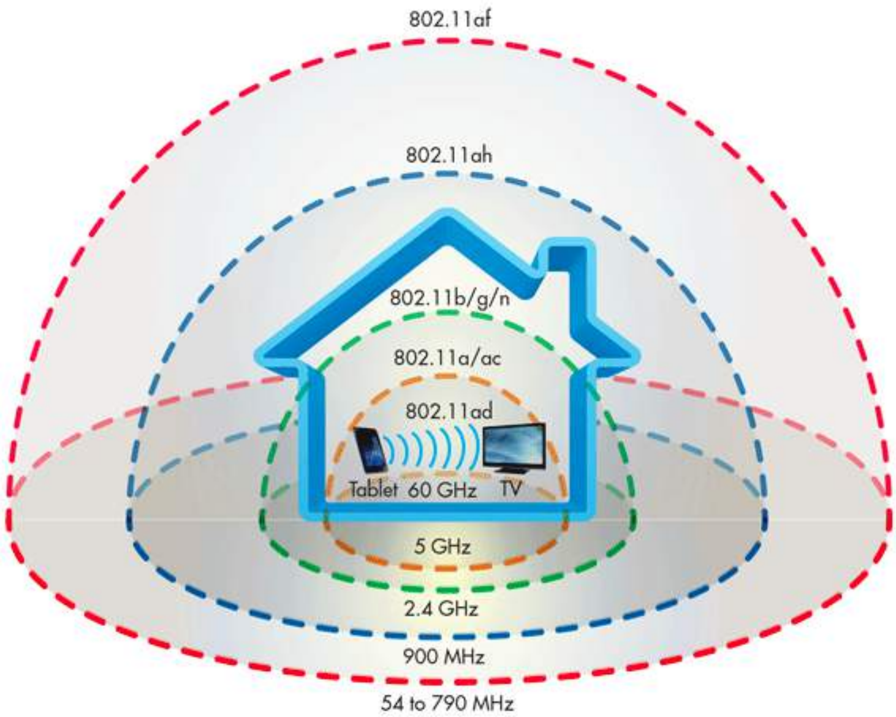

Figure 10: IEEE 802.11 Standards and frequencies and relative ranges. Source: Ars Technica.

Data on the impact of frequency and range on the network standby function power draw is limited. Calwell et al. (2011), in an examination of 4 SNE products, suggests that frequency impacts Wi-Fi power draw to some degree, however that increase is small compared to the overall power draw of the product. Other investigators have studied mobile products, but we lack information that is generally applicable to mains-connected devices.

Wi-Fi cards for several frequency bands (currently 2.4 GHz (most typical) and 5 GHz bands) typically have higher power draw levels, but mobile phones are examples of products with several radios (Wi-Fi 2.4 and 5 GHz, LTE, 2G, 3G, Bluetooth and NFC) and with low consumption levels (Viegand et al. 2017).

VIII. APPENDIX B – COMPONENT DATA SOURCES

The following tables present information on a variety of manufacturer datasheets used to inform power range estimates.

Table 5: Datasheets Referenced for System-on-Chip Products

Manufacturer	Part Number	Link
Wi-Fi		
Texas Instruments	CC3200	http://www.ti.com/product/cc3200
Texas Instruments	CC3200MOD	http://www.ti.com/product/cc3200mod
Texas Instruments	CC3220 R,S	http://www.ti.com/product/cc3220
Texas Instruments	CC3220 SF	http://www.ti.com/product/cc3220
Texas Instruments	CC3220MODS	http://www.ti.com/product/cc3220mod
Texas Instruments	CC3220MODF	http://www.ti.com/product/cc3220mod
Texas Instruments	MRF24WB0MA	http://ww1.microchip.com/downloads/en/DeviceDoc/70632C.pdf
Texas Instruments	MRF24WB0MB	http://ww1.microchip.com/downloads/en/DeviceDoc/70632C.pdf
Texas Instruments	RN131	http://ww1.microchip.com/downloads/en/DeviceDoc/rn-131-ds-v3.2r.pdf
Texas Instruments	RN171	http://ww1.microchip.com/downloads/en/DeviceDoc/70005171B.pdf
Texas Instruments	RN1723	http://ww1.microchip.com/downloads/en/DeviceDoc/70005224A.pdf
Texas Instruments	MRF24WN0MA	http://ww1.microchip.com/downloads/en/DeviceDoc/50002410A.pdf
Texas Instruments	MRF24WN0MB	http://ww1.microchip.com/downloads/en/DeviceDoc/50002410A.pdf
Texas Instruments	RN1810	http://www.ti.com/lit/gpn/LMX9830
Texas Instruments	MRF24WG0MA	http://ww1.microchip.com/downloads/en/DeviceDoc/70686B.pdf
Texas Instruments	MRF24WG0MB	http://ww1.microchip.com/downloads/en/DeviceDoc/70686B.pdf
Texas Instruments	ATSAMW25	http://ww1.microchip.com/downloads/en/DeviceDoc/Atmel-42618-SmartConnect-ATSAMW25-MR210PB_Datasheet.pdf
Texas Instruments	ATWINC1500	http://ww1.microchip.com/downloads/en/DeviceDoc/70005304A.pdf
Zigbee		
Texas Instruments	CC2630	http://www.ti.com/product/CC2630
Texas Instruments	CC2538	http://www.ti.com/product/CC2538
Texas Instruments	CC2531	http://www.ti.com/product/CC2531
Texas Instruments	CC2530	http://www.ti.com/product/CC2530
Texas Instruments	CC2520	http://www.ti.com/product/CC2520
Bluetooth LE		
Microchip	RN4020	http://ww1.microchip.com/downloads/en/DeviceDoc/50002279B.pdf
Nordic Semiconductor	nRF51822	https://www.nordicsemi.com/eng/Products/Bluetooth-low-energy/nRF51822
Nordic Semiconductor	nRF8001	https://www.nordicsemi.com/eng/Products/Bluetooth-low-energy/nRF8001
Bluetooth		
Texas Instruments	CC2560	http://www.ti.com/product/CC2560
Texas Instruments	CC2564	http://www.ti.com/product/CC2564
Texas Instruments	CC2564MODA	http://www.ti.com/product/CC2564MODA
Texas Instruments	CC2564C	http://www.ti.com/product/CC2564C
Texas Instruments	CC2564MODN	http://www.ti.com/product/CC2564MODN

Manufacturer	Part Number	Link
Microchip	BM20	http://ww1.microchip.com/downloads/en/DeviceDoc/BM20_EvaluationBoard_UserGuide.pdf
Microchip	BM23	http://ww1.microchip.com/downloads/en/DeviceDoc/BM20_EvaluationBoard_UserGuide.pdf
Microchip	BM62	http://ww1.microchip.com/downloads/en/DeviceDoc/60001403A.pdf
Microchip	BM63	http://ww1.microchip.com/downloads/en/DeviceDoc/60001431B.pdf
Microchip	BM64	http://ww1.microchip.com/downloads/en/DeviceDoc/60001403A.pdf
Microchip	BM70	http://ww1.microchip.com/downloads/en/DeviceDoc/60001372F.pdf
Microchip	BM71	http://ww1.microchip.com/downloads/en/DeviceDoc/60001372F.pdf
Microchip	BM77	http://ww1.microchip.com/downloads/en/DeviceDoc/BM77%20Data%20Sheet%20v2.1.1%202015-Jun16.pdf
Microchip	BM78	http://ww1.microchip.com/downloads/en/DeviceDoc/60001380A.pdf
Microchip	RN41	http://ww1.microchip.com/downloads/en/DeviceDoc/rn-41-ds-v3.42r.pdf
Microchip	RN41XV	http://ww1.microchip.com/downloads/en/DeviceDoc/rn-41-ds-v3.42r.pdf
Microchip	RN42	http://ww1.microchip.com/downloads/en/DeviceDoc/rn-42-ds-v2.32r.pdf
Microchip	RN42XV	http://ww1.microchip.com/downloads/en/DeviceDoc/rn-42-ds-v2.32r.pdf
Microchip	RN4677	http://ww1.microchip.com/downloads/en/DeviceDoc/50002370A.pdf
Microchip	RN4678	http://ww1.microchip.com/downloads/en/DeviceDoc/50002519A.pdf
Microchip	RN4870	http://ww1.microchip.com/downloads/en/DeviceDoc/50002489A.pdf
Microchip	RN4871	http://ww1.microchip.com/downloads/en/DeviceDoc/50002489A.pdf
Microchip	RN52	http://ww1.microchip.com/downloads/en/DeviceDoc/rn-52-ds-1.0r.pdf

Table 6: Datasheets Referenced for Ethernet PHY + MAC Products

Manufacturer	Model	Link
Realtek	RTL8111HS	http://www.realtek.com.tw/products/productsView.aspx?Langid=1&PFid=5&Level=5&Conn=4&ProdID=11
Intel	I211	https://www.intel.com/content/dam/www/public/us/en/documents/datasheets/i211-ethernet-controller-datasheet.pdf
Intel	I350	https://www.intel.com/content/dam/www/public/us/en/documents/datasheets/ethernet-controller-i350-datasheet.pdf
Intel	82580DB/EB - 2 ports	https://www.intel.com/content/dam/www/public/us/en/documents/datasheets/82580-eb-db-gbe-controller-datasheet.pdf
Intel	82579LM	https://www.intel.com/content/www/us/en/embedded/products/networking/82579-gbe-phy-datasheet-vol-2-1.html
Intel	82583V	https://www.intel.com/content/dam/www/public/us/en/documents/product-briefs/82583v-gbe-controller-brief.pdf
Intel	82473E/V	https://www.intel.com/content/dam/doc/datasheet/82579-gbe-phy-datasheet-vol-2-1.pdf
Intel	82574L/IT	https://www.intel.com/content/dam/doc/datasheet/82574l-gbe-controller-datasheet.pdf
Intel	I210-IS	https://www.intel.com/content/dam/www/public/us/en/documents/datasheets/i210-ethernet-controller-datasheet.pdf
Intel	82599ES	https://www.intel.com/content/dam/www/public/us/en/documents/datasheets/82599-10-gbe-controller-datasheet.pdf

Manufacturer	Model	Link
Microchip	ENC424J600/624J600	http://ww1.microchip.com/downloads/en/DeviceDoc/39935c.pdf
Microchip	LAN9420/LAN9420i	http://ww1.microchip.com/downloads/en/DeviceDoc/9420.pdf
Microchip	LAN9211	http://ww1.microchip.com/downloads/en/DeviceDoc/00002414A.pdf
Microchip	LAN9215	http://ww1.microchip.com/downloads/en/DeviceDoc/00002412A.pdf
Microchip	KSZ8441HL/FHL	http://ww1.microchip.com/downloads/en/DeviceDoc/KSZ8441HL_FHL_ds_1.0.pdf
Microchip	LAN9210	http://ww1.microchip.com/downloads/en/DeviceDoc/9210.pdf
Microchip	LAN9217	http://ww1.microchip.com/downloads/en/DeviceDoc/00002411A.pdf
Microchip	LAN9218	http://ww1.microchip.com/downloads/en/DeviceDoc/00002409A.pdf
Microchip	LAN9220	http://ww1.microchip.com/downloads/en/DeviceDoc/00002417A.pdf
Microchip	LAN9221/LAN9221i	http://ww1.microchip.com/downloads/en/DeviceDoc/00002416A.pdf
Microchip	LAN9420/LAN9420i	http://ww1.microchip.com/downloads/en/DeviceDoc/9420.pdf

Table 7: Datasheets Referenced for Wi-Fi PHY + MAC Products

Manufacturer	Model	Link
Microchip	MRF24WB0MA	http://ww1.microchip.com/downloads/en/DeviceDoc/70632C.pdf
Microchip	MRF24WB0MB	http://ww1.microchip.com/downloads/en/DeviceDoc/70632C.pdf
Microchip	RN131	http://ww1.microchip.com/downloads/en/DeviceDoc/rn-131-ds-v3.2r.pdf
Microchip	RN171	http://ww1.microchip.com/downloads/en/DeviceDoc/70005171B.pdf
Microchip	RN1723	http://ww1.microchip.com/downloads/en/DeviceDoc/70005224A.pdf
Microchip	MRF24WN0MA	http://ww1.microchip.com/downloads/en/DeviceDoc/50002410A.pdf
Microchip	MRF24WN0MB	http://ww1.microchip.com/downloads/en/DeviceDoc/50002410A.pdf
Microchip	RN1810	http://ww1.microchip.com/downloads/en/DeviceDoc/50002460B.pdf
Microchip	MRF24WG0MA	http://ww1.microchip.com/downloads/en/DeviceDoc/70686B.pdf
Microchip	MRF24WG0MB	http://ww1.microchip.com/downloads/en/DeviceDoc/70686B.pdf
Microchip	ATSAMW25	http://ww1.microchip.com/downloads/en/DeviceDoc/Atmel-42618-SmartConnect-ATSAMW25-MR210PB_Datasheet.pdf
Microchip	ATWINC1500	http://ww1.microchip.com/downloads/en/DeviceDoc/70005304A.pdf
Texas Instruments	CC3100	http://www.ti.com/product/cc3100
Texas Instruments	CC3100MOD	http://www.ti.com/product/cc3100mod
Texas Instruments	CC3120	http://www.ti.com/product/cc3120
Texas Instruments	CC3120MOD	http://www.ti.com/product/cc3120mod
Texas Instruments	WL1807MOD	http://www.ti.com/product/WL1807MOD
Texas Instruments	WL1837MOD	http://www.ti.com/product/WL1837MOD
Texas Instruments	WL1831MOD	http://www.ti.com/product/WL1831MOD
Texas Instruments	WL1835MOD	http://www.ti.com/product/WL1835MOD
Texas Instruments	WL1805MOD	http://www.ti.com/product/WL1805MOD
Texas Instruments	WL1801MOD	http://www.ti.com/product/WL1801MOD