ecos™

Making a World of Difference

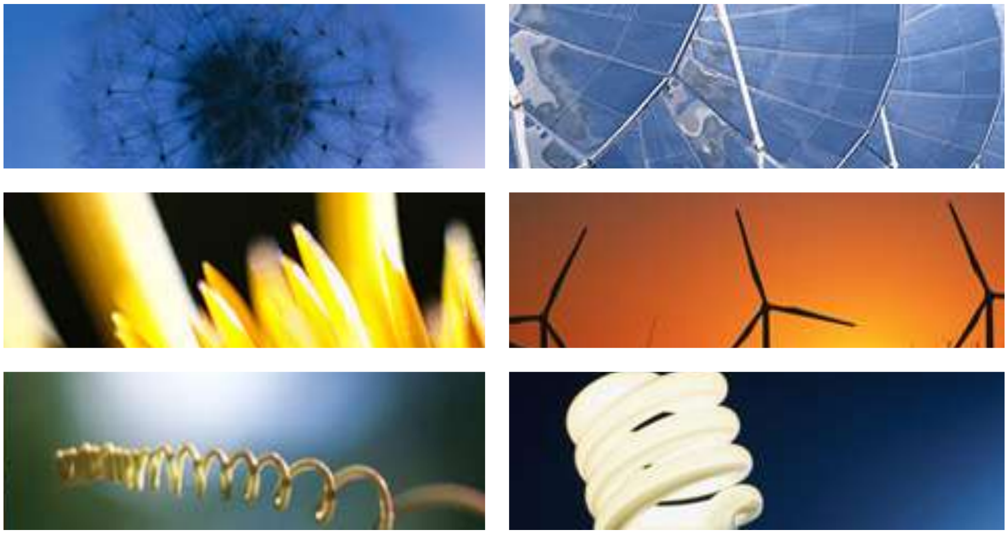

Portland, OR | San Francisco | Seattle | Durango, CO
Detroit | Philadelphia

Prepared for the
Asia-Pacific Partnership

APP Project G2: Power Scaling in Proportion to Data Processing

Submitted by:
Chris Calwell, Jeffrey Swofford, Matthew Krick, Eric Wanless and David Denkenberger, PhD.

Ecos
1199 Main Avenue #242
Durango CO 81301
970 259-6801

September 16, 2011~~June 29, 2011~~

Table of Contents

Executive Summary

The notion of power scaling is not complex or proprietary. Simply stated, power scaling means that a product can dynamically and proportionally vary its power consumption as its workload changes. For this project, Ecos investigated the power scaling abilities of a wide variety of consumer electronics products. To be able to quantify these power use differences, Ecos has identified a set of “universal tasks” that users would routinely expect to perform on their electronic products. We conducted power consumption measurements over time with each universal task performed on each device, in addition to documenting the relative performance capability. For edge devices, we focus on five universal tasks: 1) standard modes available (e.g. active idle and low power), 2) video play, 3) music play, 4) daily computing and 5) game play. Each of these “primary” universal tasks has a number of sub-tasks (e.g. different methods for listening to music). For network equipment, we focus on a separate set of tasks centered on network speeds and network load.

On the following page, we provide an updated version of the “master” graphic that jump-started our initial power scaling investigations over a year ago. The universal tasks listed in the horizontal axis are some of the most popular tasks today, and we have plotted them for nearly all of the edge devices that we tested. Key findings, both from this graph and the project as a whole, include:

- Most of the devices that use a small amount of power in an absolute sense are doing so with fundamental efficiency improvements to their hardware, rather than with power scaling per se.
- The devices that use the most power tend to exhibit some degree of power scaling, but even when idling, they cannot approach the consumption levels of their more efficient competitors.
- The Lenovo laptop was one of the few devices we tested that could perform most of these universal tasks and scale its power consumption by a factor of more than four among the tasks.
- The Mac Mini also demonstrated significant power scaling (a factor of more than three, while drawing significantly less power than most other stationary devices to perform similar tasks.
- For the rest of the products, power scaling across a range of 2:1 or less was the norm, even though some of the tasks were far more computationally intensive than others.
- Game consoles and set-top boxes continue to exhibit very poor power scaling capabilities, even in the newest iterations.

Power scaling remains an idea that can add significantly to the way energy efficiency specifications are developed and implemented. Not only can wattage limits be specified for standby, sleep, and idle modes for various electronic products, for example, but policymakers could also consider introducing latency and performance considerations. For example:

- How many minutes can elapse between when no activity occurs with a device and when it drops into a sleep mode?
- How rapidly must a product return from sleep mode to idle mode?
- By what percentage must the power consumption drop between active or maximum performance mode and idle mode?
- To what extent can products be expected to maintain minimal network connectivity during sleep or hibernate modes?

Finally, we recommend that policymakers pursue additional research on the topic of network power scaling specifically. Standard protocols should be developed that can assess power consumption of network routers, hubs, and switches as the amount of network traffic flowing through them varies. Thereafter, labeling programs and eventually MEPS could be developed to promote those design approaches that meaningfully scale power consumption to network activity.

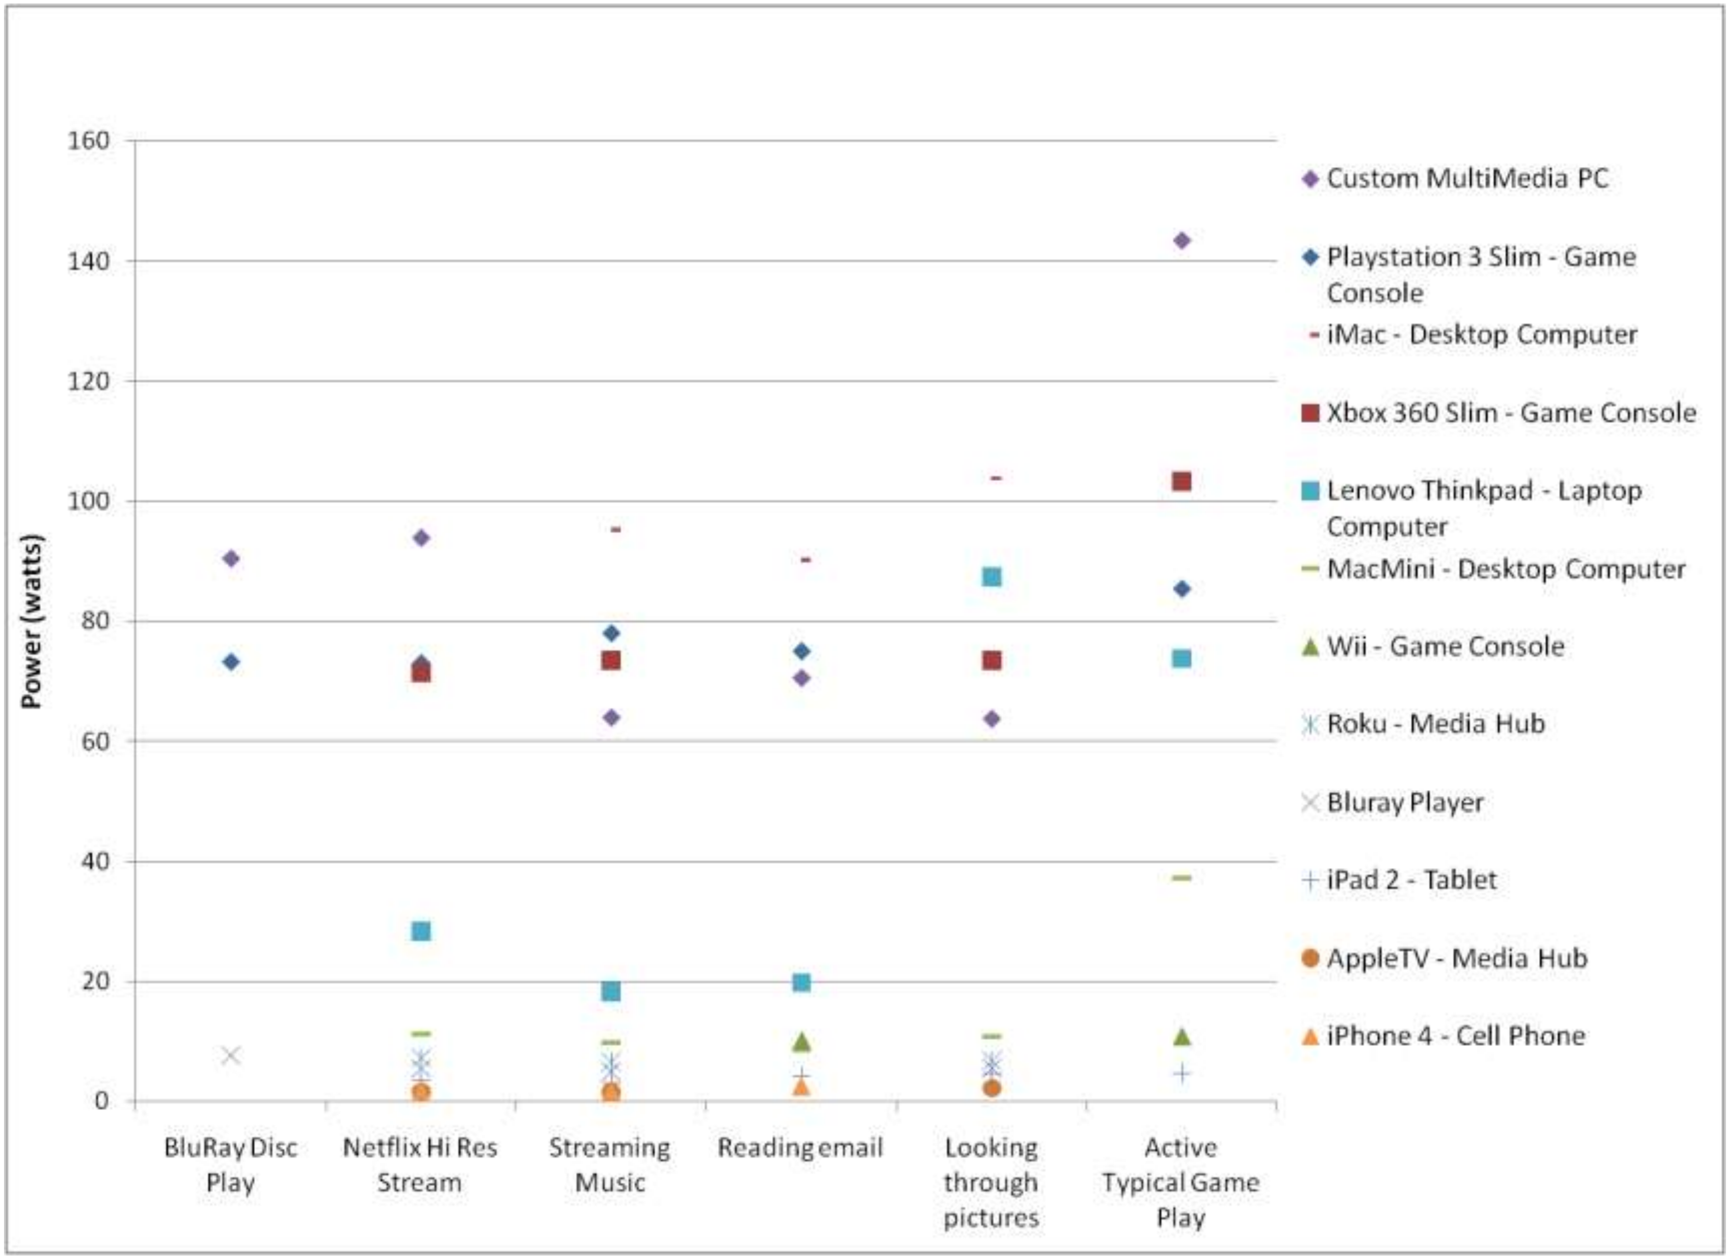
160
140
120
100
80
60
40
20
0
Power (watts)
BluRay Disc Play
Netflix Hi Res Stream
Streaming Music
Reading email
Looking through pictures
Active Typical Game Play
Custom MultiMedia PC
Playstation 3 Slim - Game Console
iMac - Desktop Computer
Xbox 360 Slim - Game Console
Lenovo Thinkpad - Laptop Computer
MacMini - Desktop Computer
Wii - Game Console
Roku - Media Hub
Bluray Player
iPad 2 - Tablet
AppleTV - Media Hub
iPhone 4 - Cell Phone

1. Introduction

As part of the Asia-Pacific Partnership on Clean Development and Climate (APP), the Building and Appliance Task Force has initiated research tasks on a variety of topics related to standby power, resulting in significant progress over recent years. The overall goal of the task force is to "…develop a common approach that all partner countries can endorse, that delivers the lowest feasible standby power for agreed appliances types by 2015". As a subset of the project, the Australian Department of Climate Change and Energy Efficiency tasked Ecos to research power scaling in proportion to data processing for information based appliances and equipment. This work, in addition to other work tasked to Ecos, will provide relevant information to assist policymakers in addressing today's standby power problem without limiting functionality and product innovation.

The notion of power scaling is not complex or proprietary. Simply stated, power scaling means that a product can dynamically and proportionally vary its power consumption as its workload changes. This is achieved fairly well in non-electronics appliances such as ovens. However, the power consumed by electronic equipment that process information is a function of circuit design, power supply, and the capacity or speed of the processor or system. For a given processor speed, the power consumed is more or less constant, irrespective of whether a lot of data is processed or none is processed. So being able to take the foot "off the accelerator" during periods of low demand is an important power saving principle. A related approach is to have separate, "purpose-built" processor cores that are designed for specific functions within a product that can be "turned off" when those functions are not required. While these concepts are simple, to effectively implement them in products requires careful design and coordination of hardware and software.

The example of power scaling in a gasoline-powered automobile is instructive. The power consumption when operating at maximum acceleration is about 10 times higher than when idling (engine on, but vehicle not in motion). However, for a desktop computer, which power scales better than many electronic devices, this ratio is only about a factor of two or three. Yet, for all of its mechanical complexity and number of onboard computers, a vehicle can still be shut completely off and restarted again in a matter of a few seconds, while desktop computers still require roughly 20 to 45 seconds. Vehicles therefore have a lower *latency* than desktop computers. Computers can scale their power consumption down by about two orders of magnitude when in sleep mode and up to three orders of magnitude when in standby mode. But here again, vehicles do better. When switched off, their only power consumption occurs from self-discharge in their battery, and that occurs at roughly five orders of magnitude below full power operation. There is considerable room for improvement in computers to achieve comparable scaling and latency.

Personal computers were one of the first consumer electronics products to demonstrate power scaling capability, automatically adjusting the voltage and frequency of their central processing units (CPUs) in response to changes in workload (processing requirement). Cellular phones and other mobile devices power scale extraordinarily well in order to conserve battery life and limit heat build-up. Mobile devices employ numerous separate cores, each dedicated to a different function such as watch for incoming calls, process video, and operate built-in cameras as required. These cores can be switched on or off as needed (so called heterogeneous cores). Power scaling becomes increasingly important as electronic products acquire more capabilities. Early cellular phones were principally used for voice communication, for example, but are not routinely able to take digital photographs, record high definition video, send and receive text messages, play music, stream movies, and surf the Internet. These tasks place different workloads on the hardware within a phone, and so the phone needs to be able to scale power consumption in real time to maximize battery life during periods of less intensive processing. This imperative continues to drive innovation in mobile devices.

On the other hand, many manufacturers of mains-powered products have not yet chosen to include power scaling in most types of popular consumer electronics products as there is yet no strong driver to encourage the adoption of these advanced power management approaches. These non-mobile devices tend to consume fairly constant amounts of power regardless of the task they are performing, foregoing a significant opportunity to conserve energy.

For this project, Ecos investigated the power scaling abilities of a wide variety of consumer electronics products. To be able to quantify these power use differences, Ecos has identified a set of "universal tasks" that users would routinely expect to perform on their electronic products. Our earlier testing of a limited subset of consumer electronics products revealed power levels that vary by orders of magnitude when the different devices are all performing similar tasks. This suggests that one of the most promising new avenues for saving energy in consumer electronics would be to routinely test their power use across a broad set of universal tasks, and then develop specifications or labeling guidelines encouraging users to purchase the devices that can accomplish those tasks while using the least energy. In some cases, these results will favor the use of a very powerful general purpose device like a computer or smart phone that employs sophisticated power scaling. In other cases, these results will favor the use of multiple, special-purpose devices, each optimized to perform particular tasks as efficiently as possible.

The research summarized in this report aims to expand on previous power scaling research to more systematically characterize power use, performance differences and latency differences across a wide range of products performing a wide range of universal tasks. We focus primarily on electronic products with multimedia, audio/video capabilities (i.e. products that have data flow or information processing as their primary function). We also explore the power scaling capability of network equipment (i.e. equipment that makes up standard IP networks such as modems, routers and switches), and how their power consumption changes in proportion to network speeds and workload.

Even when power use and latency differences can be quantified among otherwise comparable products, it remains vital to characterize performance differences as well, which can be more challenging. This will be critical in moving the discussion forward, beyond the measurements made previously, by identifying the *energy* consequences of power scaling. If the power use of one device is half that of another, but it takes three times longer to perform the task, its energy consumption can actually be higher, for example. Similarly, if one device is far better at power scaling than another but imposes very high latency constraints, few users will bother to place it in its lowest power consuming mode, or, perhaps, even purchase it in the first place.

2. Methodology

For this project, we focus on three primary research tasks:

1. Establish a set of universal tasks that are suited to a wide variety of devices;
2. Conduct detailed measurements of power consumption over time with each function on each device, for purposes of documenting typical power consumption, energy use, latency, and relative performance levels associated with each; and
3. Identify and discuss the *energy* consequences of power scaling.

2.1. Selection of Devices and Test Methodology

To explore each of these tasks, we purchased a variety of consumer electronics products to investigate their power scaling capabilities. The goal in every case is to understand the qualitative and quantitative extent to which each product can vary its energy consumption in response to workload, and do so in a way that meets consumer expectations for latency and functionality.

Table 1: List of edge and network devices used in study

Device Family	Device Type	Device Name	Device Family	Device Type	Device Name
Edge Devices	Game Console	PlayStation 3 Slim	**Network Devices**	Router	D-Link DIR-655
	Game Console	Xbox 360 Slim		Router	D-Link DIR-615
	Game Console	Wii		Router	D-Link DIR-825
	Blu-ray Player	Panasonic BDT210P		Router	D-Link DIR-632
	Blu-ray Player	Samsung BD-D6700		Switch	D-Link DGS-1016D
	DVD Player	Yamaha S659		Switch	D-Link DGS-2205
	Media Hub	Roku		Switch	D-Link DSS-8+
	Media Hub	AppleTV		Switch	D-Link 1008G
	Set-top box	Satellite		Switch	Linksys SD216
	Set-top box	Satellite			
	iPad	iPad2			
	Desktop Computer	iMac			
	Desktop Computer	MacMini			
	Desktop Computer	HP Firebird			
	Desktop Computer	Custom MultiMedia PC			
	Laptop	Lenovo T420 Thinkpad			
	Cell Phone	iPhone 4			

Note that the range of manufacturers and types of devices is intentionally broad among the edge devices but narrow within the network devices. Because D-Link has made specific, detailed claims in product literature and product packaging that its competitors have not made about the ability of its products to scale power in response to changes in network traffic speed, number of connected ports, and even the length of connected Ethernet cables, we wanted to investigate and compare those products in some detail. These D-Link products span a number of years of manufacturing history that allows us to note any trends or changes over time.

Most power measurements were made using a Watts Up Pro meter, averaged over 5 minutes as needed for the power values to stabilize. We measured low-power tasks and devices using a Wattman meter for better precision. We focused our latency measurements around relevant low power modes (boot time from off, wake up time from sleep, etc). We measured the time elapsed for the device to enter into a specific mode, and correlated the mode change with significant differences in power consumption observed at the meter. For most universal tasks, we also documented relative performance in the format of screen size, resolution, and any other qualitative measures worth noting.

2.2. Defining a Set of Universal Tasks

We used a limited set of universal tasks in our previous power scaling work, but for this work we set out to more clearly define and categorize universal tasks that could be tested across a wide range of edge devices. Early on in our thinking, we determined that it would be best to separate the universal tasks related to 1) edge devices and 2) network equipment. While edge devices are capable of a wide range of tasks, network equipment offers much less variety in terms of functionality and end-use. While we exposed edge devices themselves to a variety of different “network” related tasks in their own testing, we limit what we call network testing in this document to measurements made to network equipment directly.

We developed our preliminary list of universal tasks after consultation with internal experts on various consumer electronic topics and review of previous field research findings related to typical consumer behavior and viewing habits. We then expanded the list of sub-tasks to be measured in consultation with other outside international experts, including the DCCEE/APP review committee.

Edge Devices

The rapid pace of innovation occurring in consumer electronics has resulted in an increase in the number of features and functionalities present in some of today’s devices. In addition, several devices are now beginning to offer secondary functionality. For example, game consoles are also now capable of streaming video content in addition to merely playing video games. Knowing this, we spoke with a variety of experts on today’s most common universal tasks related to edge devices. The universal tasks we chose can be divided into five primary tasks, each of which consists of a set of diverse subtasks. This makes it possible, for example, to determine if playing music from one media format is more energy consumptive than another, both within a platform and across platforms. While one of our goals was to maximize cross-functionality (i.e. number of tasks that a device is capable of and vice versa), we recognize that not all devices are capable of performing all tasks, but some, like computers, are. Many of the findings presented here will thus not include all of the devices previously presented in Table 1.

Table 2: Primary and secondary universal tasks for edge devices.

Primary Task	Sub-Task(s)					
Operation Mode	Off/Standby	Sleep	Idle for 5 minutes	Idle for 15 minutes	-	-
Video Play	Blu-Ray Disc - Play	Blu-Ray Disc - Paused	DVD Disc	Netflix Hi Res Stream	Low Res Stream	Pay-TV
Music Play	Streaming Music	Play CD Disc	Play from storage	-	-	-
Daily Computing	Reading email	Internet browsing	Texting/chatting	Archiving data	Looking through pictures	-
Gameplay	Idle - console menu	Idle - game menu	Idle - game paused	Idle - during typical game play	Active - during typical game play	-

Network Equipment

When setting out to define a set of universal tasks for network equipment, we wanted to subject the network devices we purchased to a variety of configurations that are most commonly used in homes today. Unlike our focus for edge devices, which provided us with a variety of end uses to center on, our focus for network equipment was primarily centered on selecting a variety of configurations that varied workload.

The following figures illustrate the range of different network configurations tested. Varying workload was accomplished in two ways: changing the number of ports and changing the throughput speed of the network as a whole.[1] Some additional testing was undertaken to determine the power increment associated with network connectivity of edge devices, and to more carefully compare the power consumption of wireless to wired networks. Each of the following diagrams represents completed network equipment device testing across a variety of switches and routers.

---

[1] We looped a 30 MB file continuously in transfers among networked devices, changing the maximum throughput speed at the edge device. This approach did not control the number of data packets transmitted over time, but did vary their transmit speed.

Baseline Test

One of the first measurements we set out to accomplish was simply the power draw of each network device (i.e. both switches and routers) with no load or traffic, with WiFi capability disabled. We plugged the device into the power meter, with no Ethernet cables plugged into the network gear, and recorded the average power draw over a 5 minute period (Figure 1).

Figure 1: Baseline test configuration for network equipment.

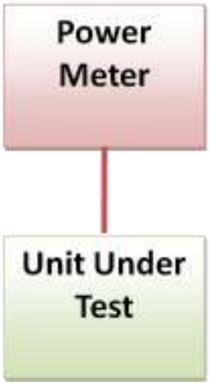

Baseline Test 2

We also wanted to investigate the incremental power for enabling WiFi in routers. For this test, we plugged the device into the power meter and enabled WiFi capability in the router firmware, again with no Ethernet cables plugged into the network gear, and recorded the average power draw over a 5 minute period (Figure 2). We took measurements for enabling the WiFi at 2.4 GHz, at 5 GHz, and for one router that was capable of enabling both frequencies simultaneously (i.e. simultaneous dual band).

Figure 2: Baseline test configuration with WiFi enabled.

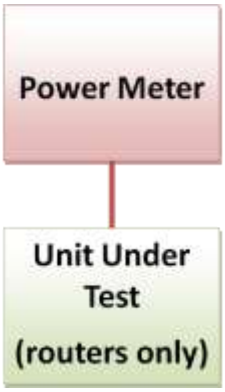

Peer-to-Peer File Transfer Testing

Now that we had established the basic power draw of our network equipment, we shifted our testing to focus on varying workload. We purchased network traffic generator software that allowed us to control and send a custom binary file (30 MB size) continuously from one computer to another. The configuration consisted of the network device connected to the power meter, and two edge devices (one laptop and one desktop) connected to the network equipment with Cat-6 Ethernet cables (Figure 3). We took measurements when 1) no traffic was being generated, 2) connection speed was set at 10 Mbit/s, 3) connection speed was set at 100 Mbit/s and 4) connection speed was set at 1000 Mbit/s. It is important to note that we configured the maximum connection speed at the *edge device*.[2] Unfortunately, the software did not control data rate transfer or connection speeds. We were not monitoring the true rate of data traffic. Due to errors with the network traffic generator software, we only tested this configuration with switches.

Figure 3: Peer-to-peer test configuration for network equipment using Ethernet and network traffic generator.

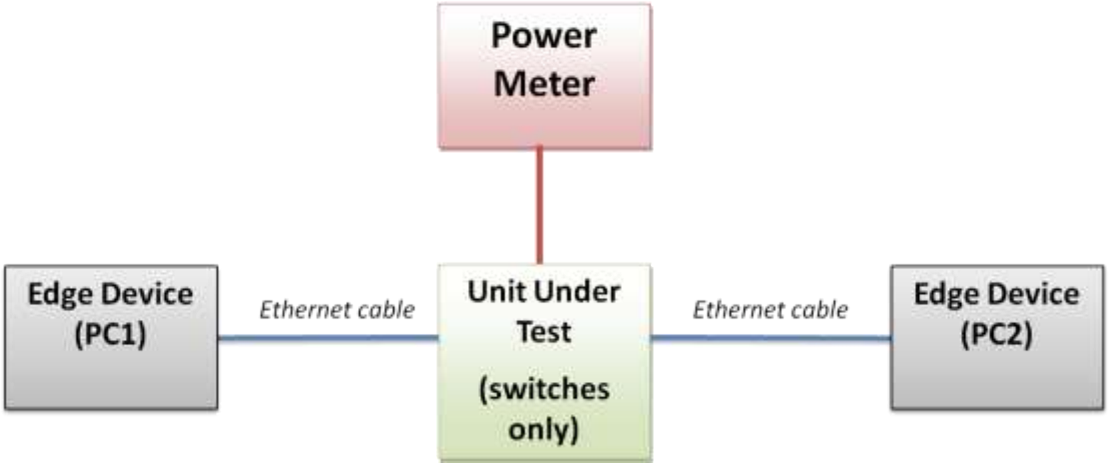

Multiple Edge Devices Served by Network Device Testing (Ethernet)

We also wanted to test network devices subjected to a full load in terms of maximum number of edge devices (i.e. number of ports used). For these tests, we disabled WiFi capability. The network device was connected to the internet via our office LAN network. We connected edge devices via Cat 6 Ethernet cables to the network device, and recorded power as each new device was added to the network. For these tests, the edge devices were all idling with established network connectivity. We did not vary connection speeds.

Figure 4: Multiple edge devices connected to internet via Ethernet.

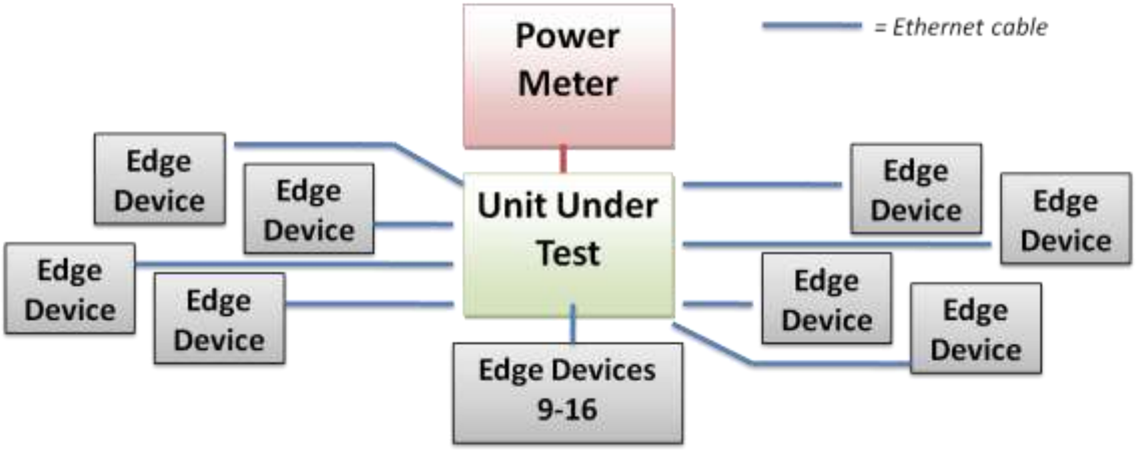

---

[2] As opposed to the network equipment firmware or within the network traffic generator software. For each edge device, we altered the maximum connection speed in Windows 7. In Window 7 device manager, we set the maximum link duplex speed to the associated value (i.e. 10/100/1000 Mbit/s).

Multiple Edge Devices Served by Network Device Testing (WiFi)

Similar to the Ethernet tests shown above, we wanted to test network devices subjected to a full load in terms of edge devices connected to one network device via WiFi. We connected edge devices via WiFi to the network device, and recorded power as each new device was added to the network. For these tests, the edge devices were all idling with established network connectivity. We did not vary connection speeds.

Figure 5: Multiple edge devices connected to internet via Wifi.

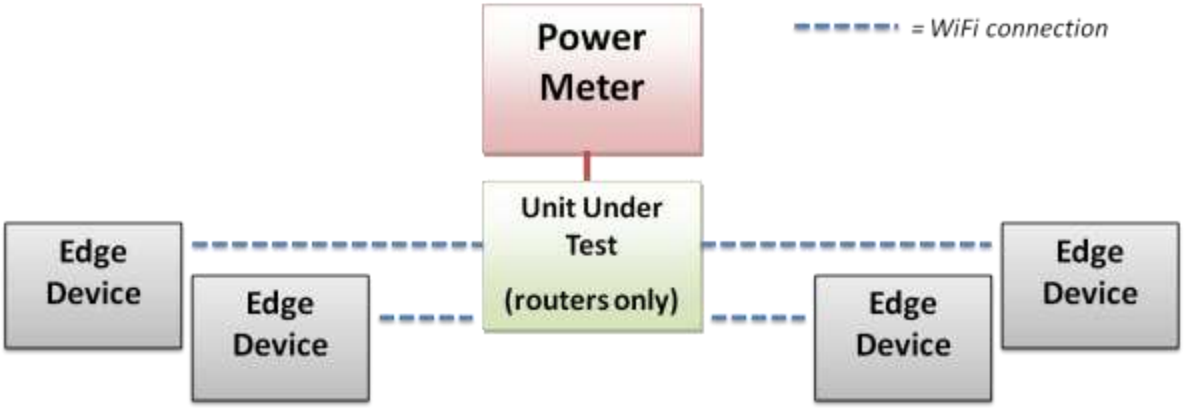

3. Findings

In our prior work on power scaling across universal tasks, we presented results in a single chart illustrating the power consumed by a wide range of devices when performing similar functions. We include a similar broadly rendered chart in our initial findings here, before elaborating on additional functions and capabilities.

The data gathered for this project offered a wealth of information about power scaling and energy consumption of consumer electronics in general. Below we present findings on five universal tasks: 1) standard modes available (e.g. active idle and low power), 2) video play, 3) music play, 4) daily computing and 5) game play. While our main goal was to highlight key findings focused on our previously defined universal tasks, of course other findings presented themselves as noteworthy during the analysis phase. We present those findings here as well.

3.1. Power Consumption of Edge Devices

Figure 6 is an updated version of the “master” graphic that jump-started our initial power scaling investigations over a year ago. The universal tasks listed in the horizontal axis are some of the most popular tasks today, and we have plotted them for nearly all of the edge devices that we tested. Note that the Lenovo laptop was one of the few devices we tested that could perform most of these universal tasks and scale its power consumption by a factor of more than four among the tasks. The Mac Mini also demonstrated significant power scaling (a factor of more than three, while drawing significantly less power than most other stationary devices to perform similar tasks. For the rest of the products, power scaling across a range of 2:1 or less was the norm, even though some of the tasks were far more computationally intensive than others. Game consoles continue to exhibit very poor power scaling capabilities, even in the newest iterations.

Obviously, for some of the most efficient mobile devices tested, a graph at this vertical scale can mask significant power consumption differences from one task to the next. What is evident, however, is the degree to which optimized devices like the iPad2 can draw less power than virtually all other products when performing similar tasks.

Figure 6: Popular universal tasks performed on a variety of devices and associated power consumption.

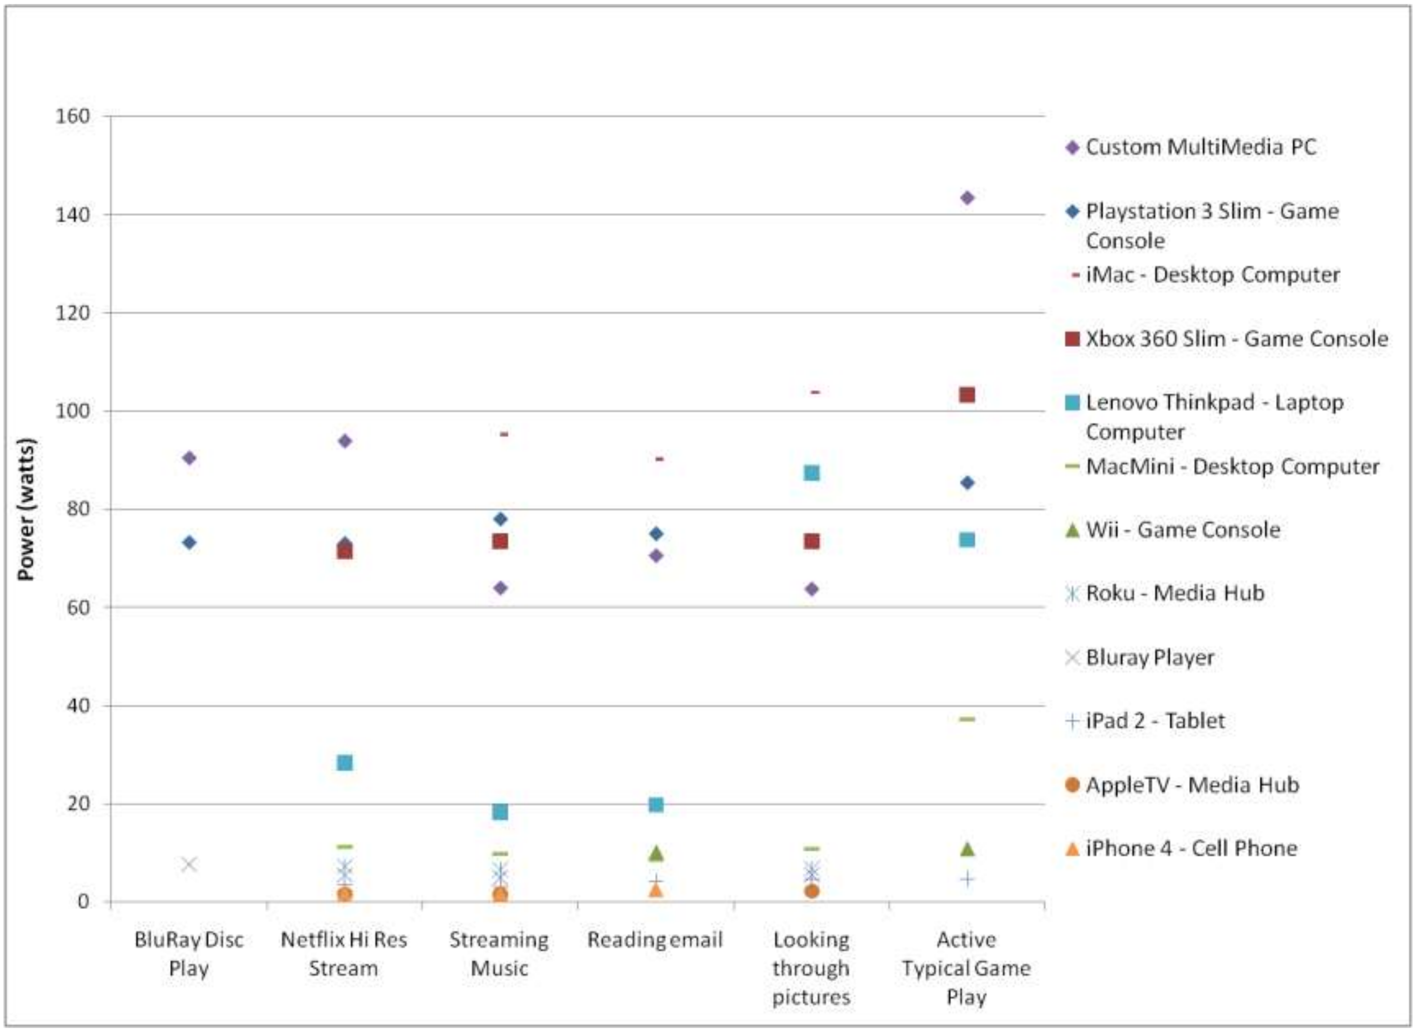

Standard Modes

One of our first set of detailed tests focused on the power consumed during off/standby, sleep, and idle modes for each of the devices examined.

As is evident in Figure 7, idle mode power values varied much more widely than off/standby and sleep mode power values. Likewise, some devices were capable of and configured to go into a sleep or auto power down mode after 15 minutes, while others were not. Computers were, as a rule, capable of the most nuanced power scaling, because they could shut off individual components like their displays or hard drives after particular periods of inactivity.[3] However, whether these computers exhibited that ability was entirely a function of whether, and how, it had been enabled in software. Default settings within Windows XP and Windows 7, for example, often forego major opportunities to scale power whenever a laptop is plugged in to an ac power source, reserving that capability only for situations in which the device is operating on battery power.

The cost of delivering a certain amount of energy from primary batteries like AAs has been well understood to be hundreds of times higher than obtaining that same amount of energy from the mains.

---

[3] Many mobile devices possess extensive power scaling capabilities as well, but do not always allow the user to customize the way they are employed to the extent that computers do.

However, even rechargeable batteries impose significant cost premiums on energy delivery relative to the mains because of power conversion losses, coulombic losses, the cost of the battery itself, and the number of charge/discharge cycles such batteries can withstand before failing. Therefore, it makes sense that power scaling is much better in battery-powered devices, but there are still cost effective opportunities to power scale with mains-powered devices. While these *capabilities* are certainly present at the time of manufacture, many IT products lose them when computer buyers, resellers, or IT managers adjust power management settings. Game consoles and set-top boxes, as noted earlier, generally exhibited very poor power scaling ability.

Figure 7: Power consumption of select edge devices during non-active modes.

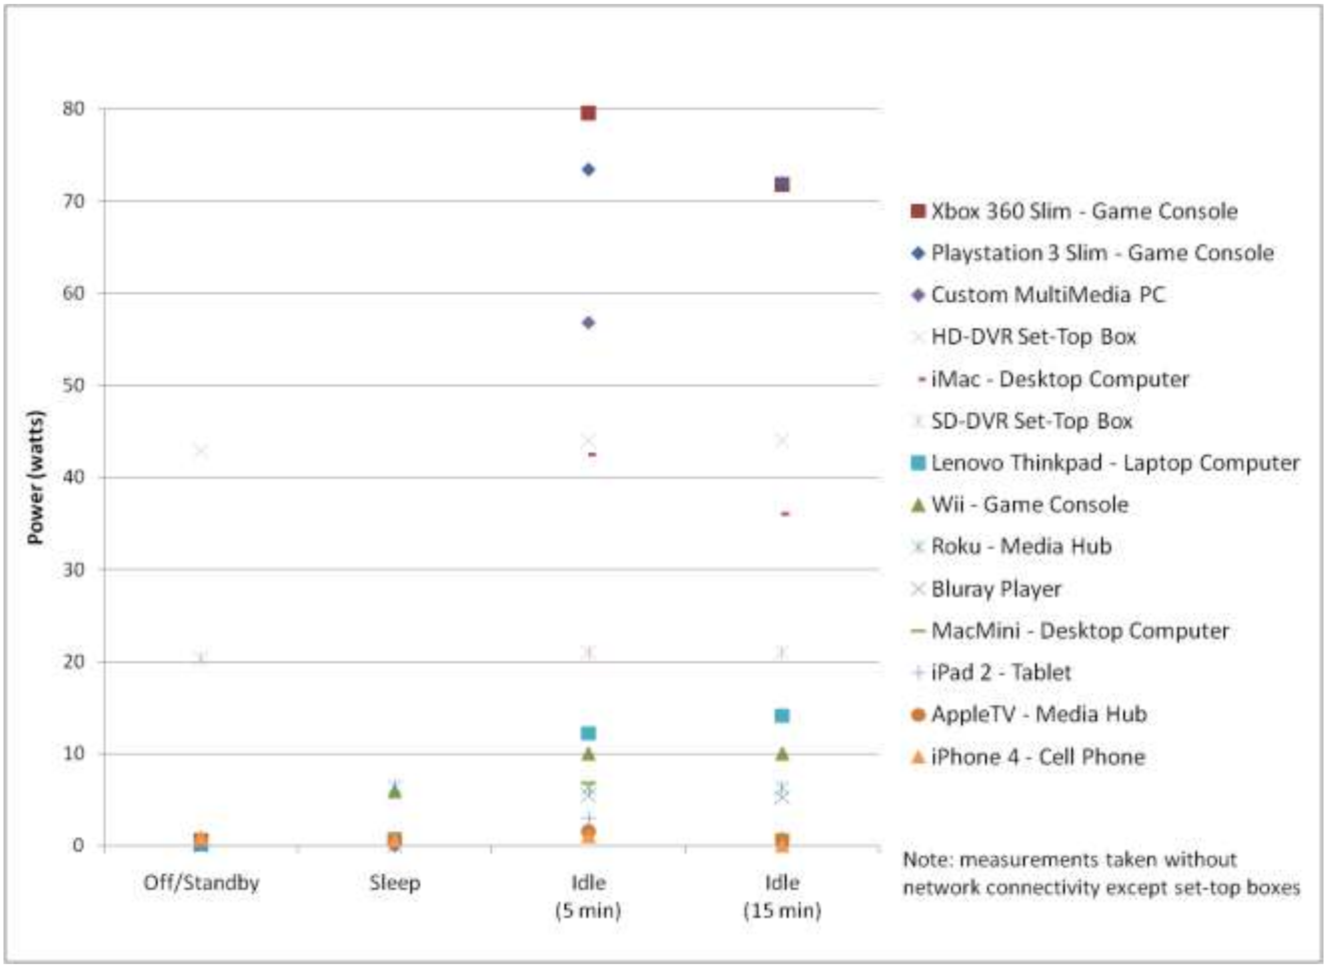

Note: measurements taken without network connectivity except set-top boxes

Off and Sleep mode power values are difficult to resolve at this scale for most of the test devices, so are shown in more detail in the following graph (Figure 8).

Off/standby mode values ranged from 0.0 to 0.9 watts for the devices we tested. Sleep mode power values ranged from 0.03 to 1.2 watts. The Wii is the only game console with sleep mode. Both the Xbox 360 and PlayStation 3 have an auto-power down feature rather than a true sleep mode. Note that the two set-top boxes are no longer present in this graphic due to their lack of a true low power sleep or off mode.[4] For the products where these values are identical to the ones observed after 15 minutes of idle, that indicates that those devices entered sleep mode during that time period. Even though the power use

---

[4] The two set-top boxes measured do not have a hard off switch, similar to most U.S. set-top boxes.

differences observed here are small on an absolute basis, they are large on a percentage basis. On the other hand, these differences could be explained by differences in power supply configuration and design.

Our sense, from the limited measurements we do have or had made for other projects previously, is that a network connection generally adds little to no power consumption during standby and sleep modes, because most products drop off the network in that situation. In idle, it may add some amount of incremental power, but we certainly could not find a big effect in the select devices we measured. For example, network connectivity in idle mode for the iMac added less than 1 watt of additional power. For the PlayStation 3, we observed a 2 watt increase in power during idle mode, one of the largest incremental power values we observed when adding network connectivity.

Figure 8: Power consumption of select edge devices during non-active modes (high resolution).

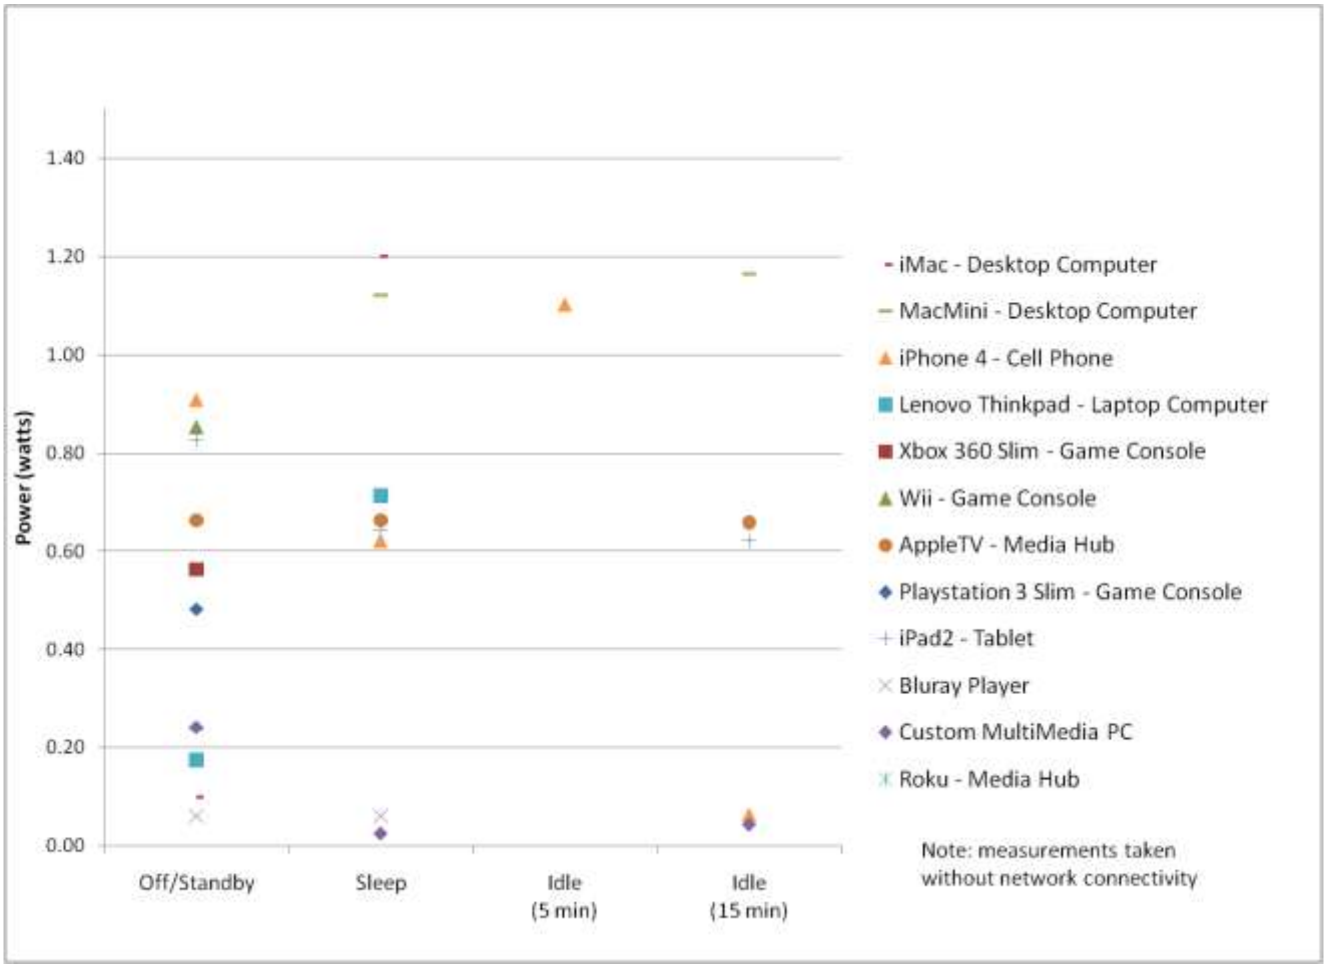

Video Play

It is now possible for video signals to reach playback devices in a number of disparate ways, each of which have their own power consumption consequences. Moreover, some playback devices contain their own built-in display screens, while others simply translate an incoming signal into something that can be readily viewed on an external device. As a result, all of the power values shown below do not always represent fully comparable data. The iMac, ThinkPad, iPad2, and iPhone include the power use of their associated screens which, though relatively small, represent an opportunity to avoid the use of an additional powered display. All of the other devices require an additional external display.

As is evident from the data in Figure 9, multi-purpose devices like computers and game consoles are still not competitive with special purpose devices like standalone Blu-ray players when playing Blu-ray discs. However, many different devices can stream Netflix movies for only about 10% of the power or less consumed by a game console or multimedia desktop PC to do the same thing. The Apple TV consumed 98% less power than the multimedia PC to stream Netflix.

Figure 9: Power consumption of select edge devices during video play tasks.

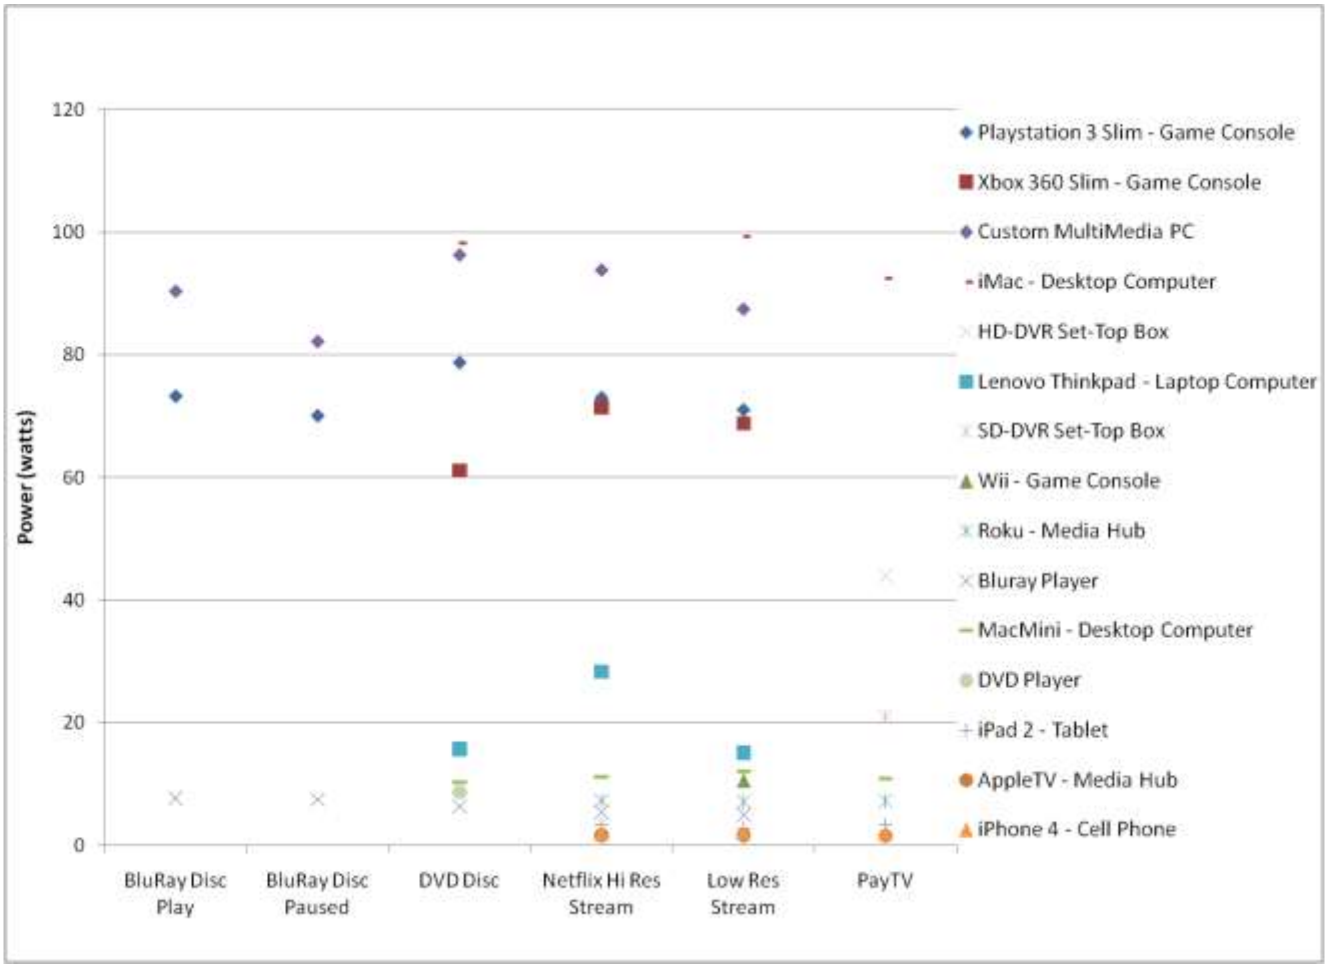

Testing focused on high resolution Netflix streaming revealed the most dramatic power consumption differences among the range of products we tested (Figure 10). Netflix appears to stream at 720p for most devices connected to broadband.[5] Each device we tested is capable of a screen resolution of 720p or greater except for the iPhone 4. As televisions are now shipping with more features embedded, we also chose to test one TV with Netflix capability. The power consumption of this particular TV was nearly 10% higher when streaming Netflix as opposed to playing the same film connected to a Blu-ray player.

Figure 10: Power consumption of select edge devices when streaming HD Netflix video.

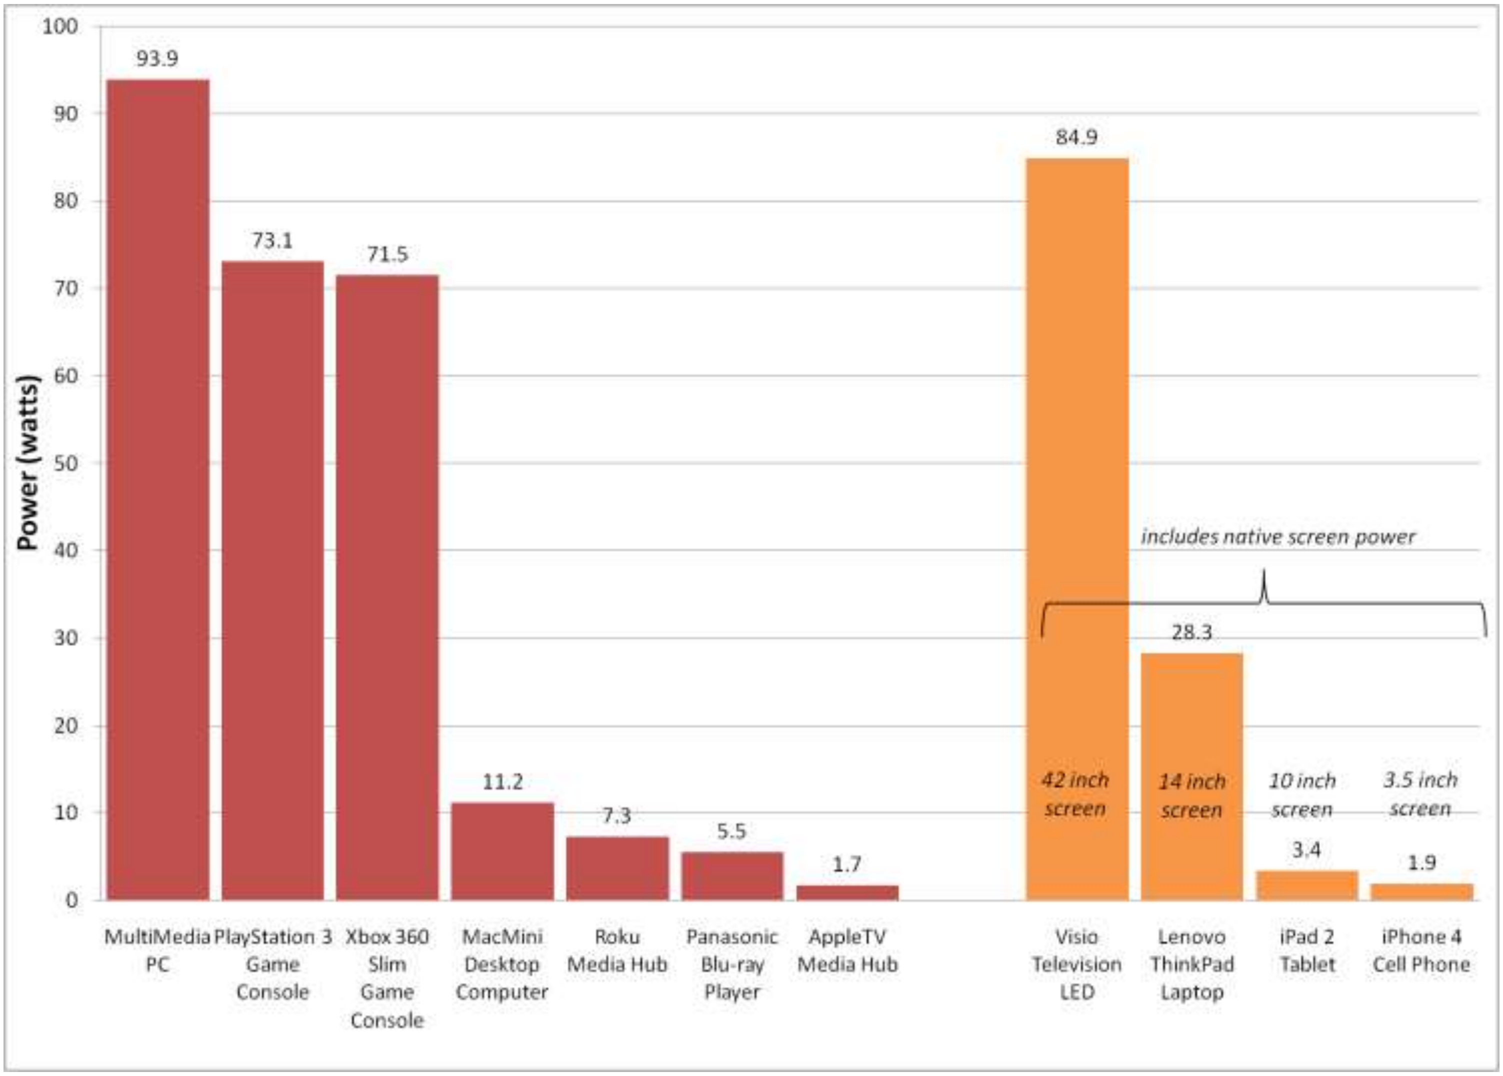

Note: For the devices in red that require an external display, we outputted to a 32” LCD television that drew an average of 77 watts during typical Netflix use.

---

[5] Wikipedia offers the following explanation of Netflix resolution (see http://en.wikipedia.org/wiki/Netflix):

“According to Netflix Tech Support, Netflix's content library is encoded into three bandwidth tiers, in a compression format based on the VC-1 video and Windows Media audio codecs. The lowest tier requires a continuous downstream bandwidth (to the client) of 1.5Mbps, and offers stereo audio and video quality comparable to DVD. The middle tier requires 3Mbps, and offers "better than DVD quality". The highest tier requires 5 Mbps, and offers 720p HD with surround sound audio. As of December 2010[update], the PLAYSTATION 3 is the only device able to stream Netflix at 1080p resolution.”

Music Play

Although music can now reach multimedia devices from a variety of sources and file types, it was evident from our testing that a given device tended to consume a similar amount of power regardless of whether it was streaming music from a network connection, playing a CD, or retrieving music from onboard memory storage. If any one trend was evident, the mechanical process of spinning a CD appeared to use more power than purely electronic processes for retrieving music. Much bigger differences were evident among the devices themselves, with the Mac Mini and dedicated playback devices drawing far less power than game consoles and computers.

Figure 11: Power consumption of select edge devices during music play tasks.

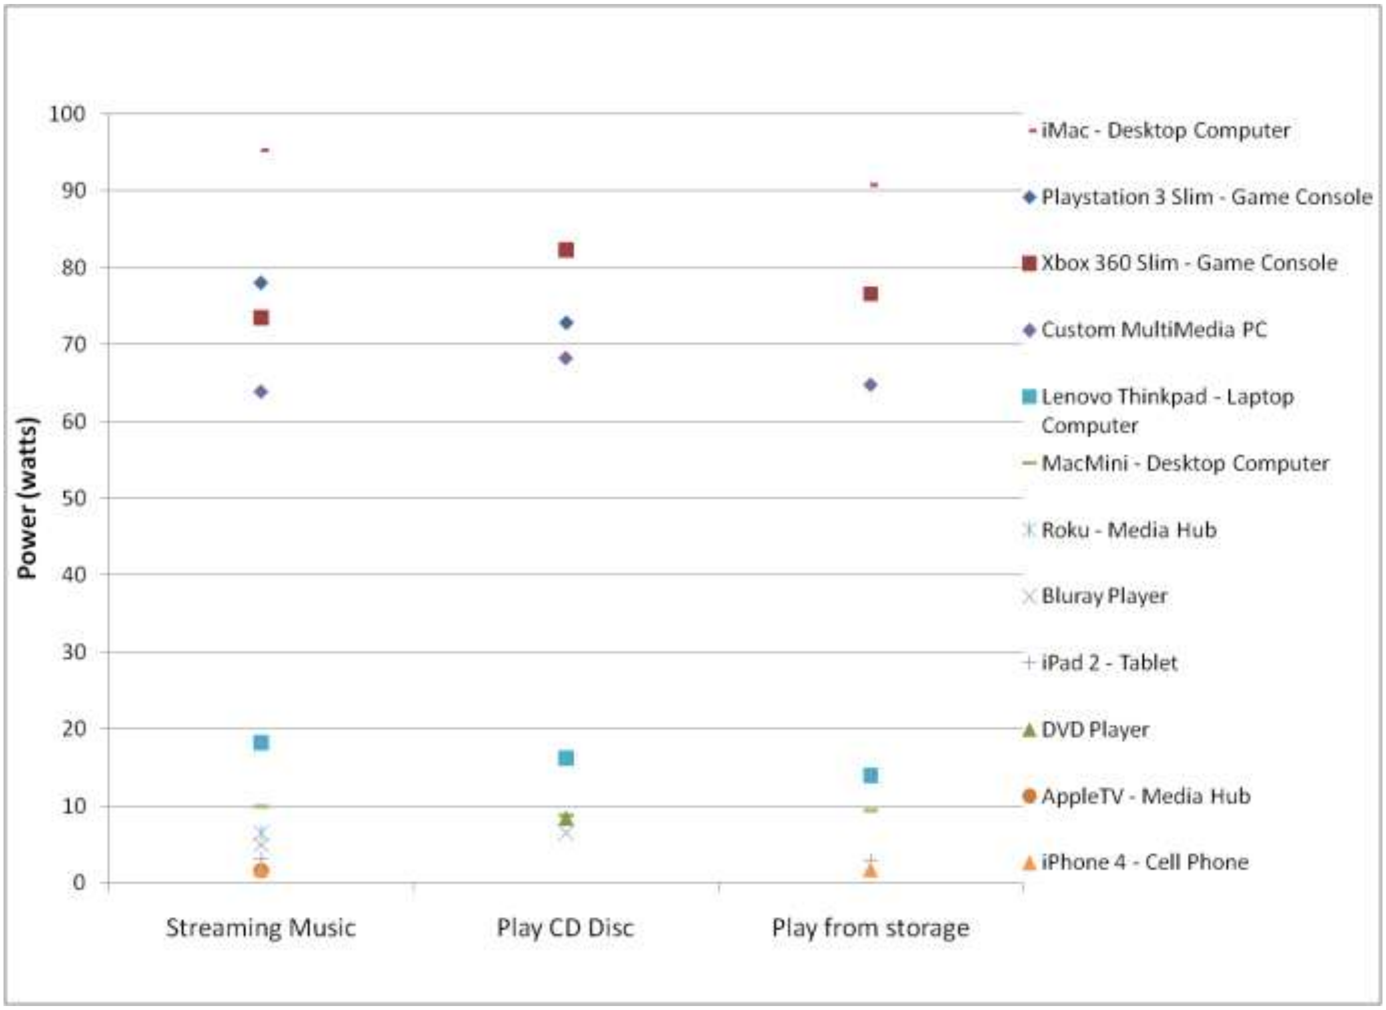

Daily Computing

Figure 12 shows several of the daily computing tasks we defined tested across a number of devices. Although power scaling capability is clearly evident in the multimedia PC, it is not as evident in the iMac, possibly because of screen power being included. The iPad 2 and MacMini (small form factor desktop computer) were both highly efficient (or rather very low power), surpassing the performance of an efficiency-optimized Lenovo laptop. With its screen power included, the iPad2 consumed only 3.8 watts while internet browsing – approximately 96% less power than the custom multimedia PC consumed, *without including the power of an associated monitor or TV*, to perform the same task.[6]

The Lenovo laptop exhibited a wider range of power scaling capability, and, of course, higher performance when executing complex tasks like active game play, but still consumed significantly more power when performing relatively simple tasks like looking through pictures or Internet browsing. More tests are presented below where we modify power management settings and re-measure devices. The Windows 7 standard power management scheme is primarily optimized for battery operation and may need to be manually changed to yield similar savings during ac power operation.

Figure 12: Daily computing tasks across various device types.

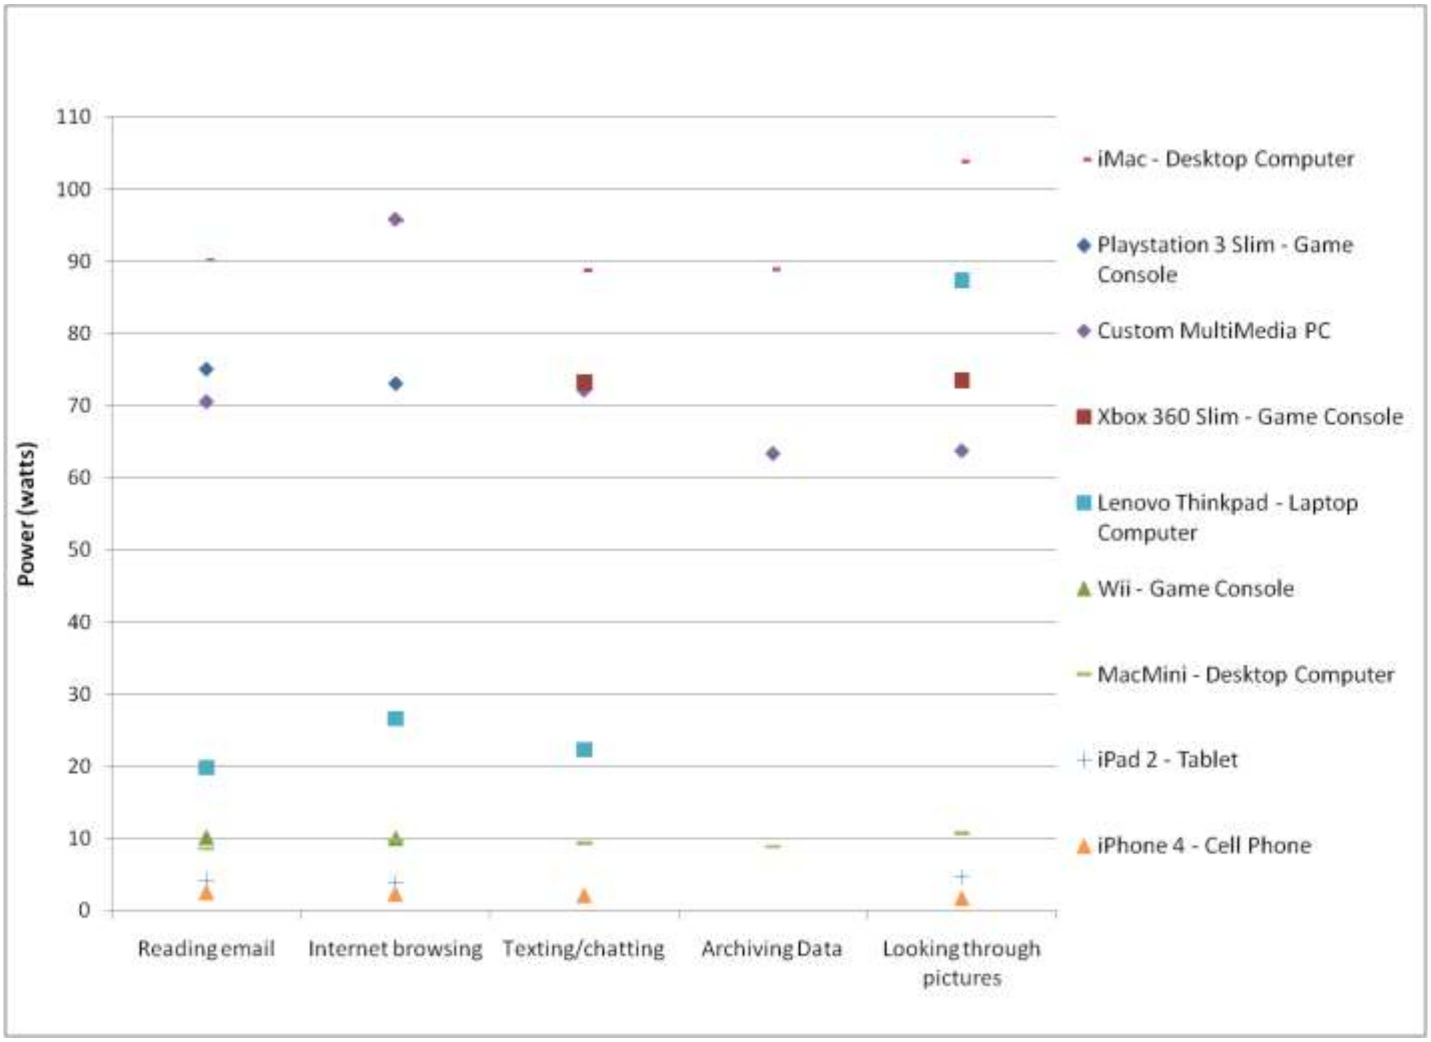

Game Play

Our next set of tests examined game consoles in more detail, with and without motion control capability (Figure 13). As noted in our prior power scaling research, absolute power consumption has already dropped significantly with each subsequent revision of the XBox 360 and PlayStation 3 hardware. As a result, we chose to focus this testing on the most current generations of all three major game console

---

[6] We tested the iPad 2, along with the other battery-powered devices, with the battery at 100% and the battery charger plugged into the power meter.

platforms, rather than comparing new to old consoles. We would note, however, that the PlayStation 3 appears to have improved more than the Xbox 360 in its on mode power. Prior generations of the PlayStation 3 used more power than an Xbox 360 manufactured at the same time, but now use less, even with built-in Blu-ray capability.

Figure 13: Power use of game consoles with and without motion control functionality active.

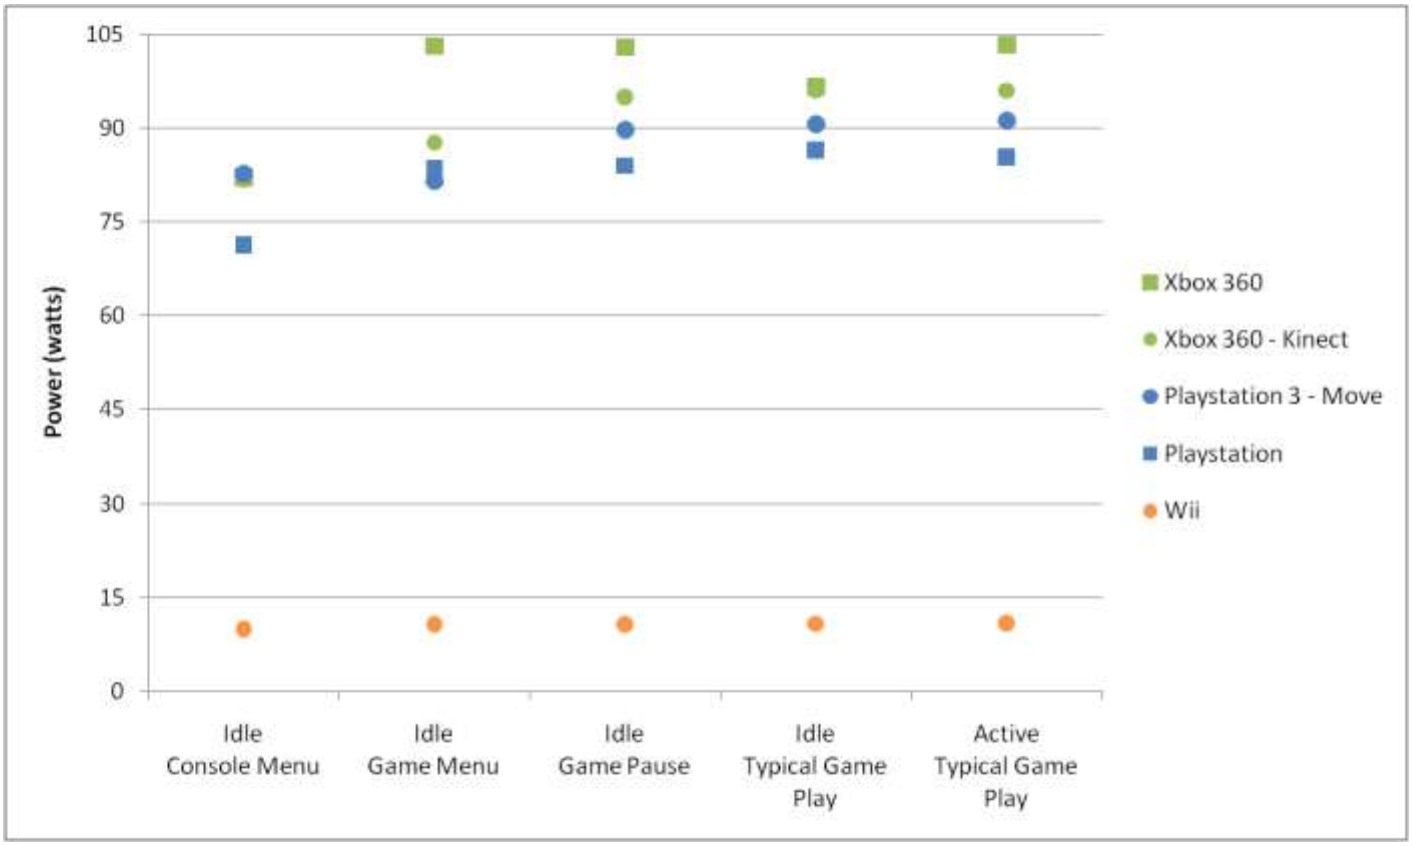

The major finding here is that all three game consoles still lack the ability to do any meaningful power scaling. The consoles' power consumption tended to be fairly constant regardless of the function being performed, though we did observe slight drops in power consumption associated with idling at the console menu. Whether a game was paused, idle, or actively being played made almost no difference to power consumption.

With the PlayStation 3, motion capability added about 4-11 watts except for Game Menu mode, when it inexplicably dropped slightly. With the Xbox 360, adding the Kinect hardware and software capability mysteriously seemed to reduce power consumption slightly. While console menu power consumption (i.e. no game disc loaded) was almost identical with and without Kinect hardware, we observed considerable differences in power draw when playing a motion game. We hypothesize this is dependent on the actual game title which, in this case, could have been far less demanding than the game used in non-motion testing (Call of Duty Black Ops). Further testing could be warranted assessing incremental power use associated with motion controllers in games that can operate with our without them.

Not surprisingly, the Wii's power use is much lower than that of the other game consoles, but performance is significantly lower as well. Screen resolution, refresh rate, user interface, and ease of use were all significantly less satisfying with Wii, to the point where the email capability was virtually unusable.

Acknowledging that more devices than just game consoles are capable of gaming functions, we extended our game play testing to a variety of edge devices (Figure 14). The three dedicated game consoles were tested with the same game title, Call of Duty Block Ops, across all tasks. The multimedia PC and Lenovo laptop were also tested using this game title. For each of the other devices, we targeted what we believed to be a high-performance game title for associated testing. There were, of course, significant differences in frame rate, screen size, resolution, and overall immersiveness across these devices, so it is impossible to attribute all of the power consumption differences to differences in product efficiency. Nevertheless, it was evident from these tests that the amount of power needed to play the same game can vary quite widely from one platform to another, and that only modest power scaling is being achieved by any of the platforms during game pause or game menu conditions.

To the extent users are shifting from playing games on multimedia desktop PCs or dedicated game consoles to mobile devices, however, it is clear that this will lead to significant energy savings. We note as well that the PCs are able to sleep or hibernate during periods of inactivity, yielding significantly greater energy savings (where they are configured to do so). This capability is still less fully achieved on the game consoles, and achieving this in the future will depend on cooperation by not only the hardware developers, but game software developers as well.

Figure 14: Power use of gaming related functions across various platforms.

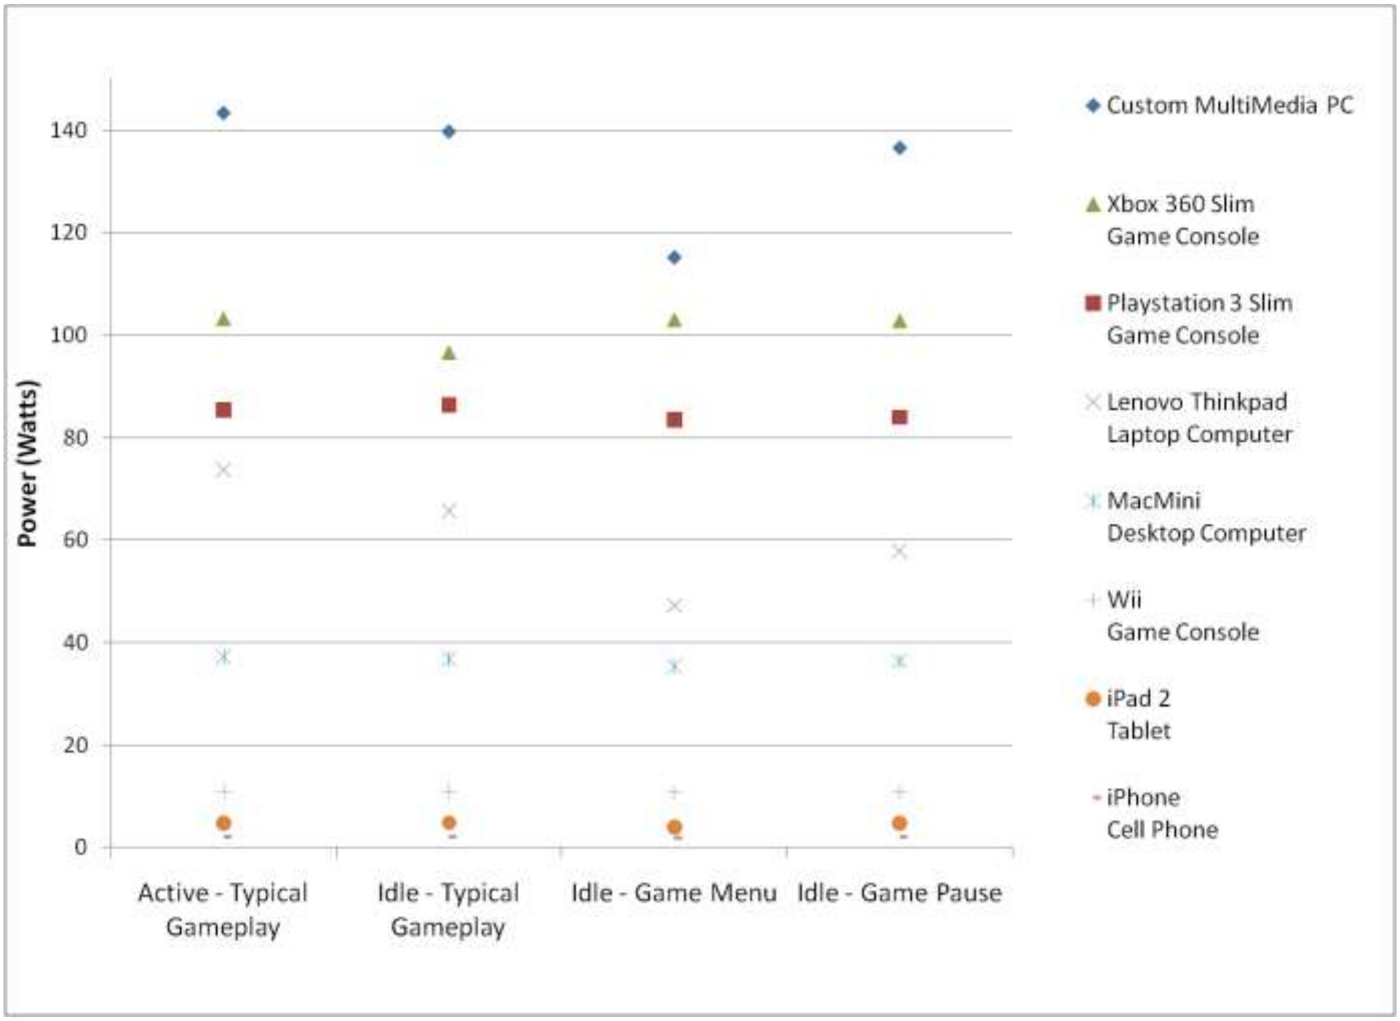

Power Management Configurations

We also wanted to explore the power consumption implications of altering the power management settings set by the operating system in certain devices.

Microsoft’s Windows 7 operating system and Apple’s Snow Leopard operating system both offer customizable power management settings. Named “Power Options”, Windows 7 offers balanced, power saver, and high performance power management schemes (Figure 15). For laptops, each of these three plans can be customized for on battery use as opposed to when the laptop’s external power supply is plugged into a wall outlet. Most importantly, Windows 7 offers advanced settings that are hardware specific, allowing for a higher resolution of modifiable settings.

Figure 15: Windows 7 power management settings.

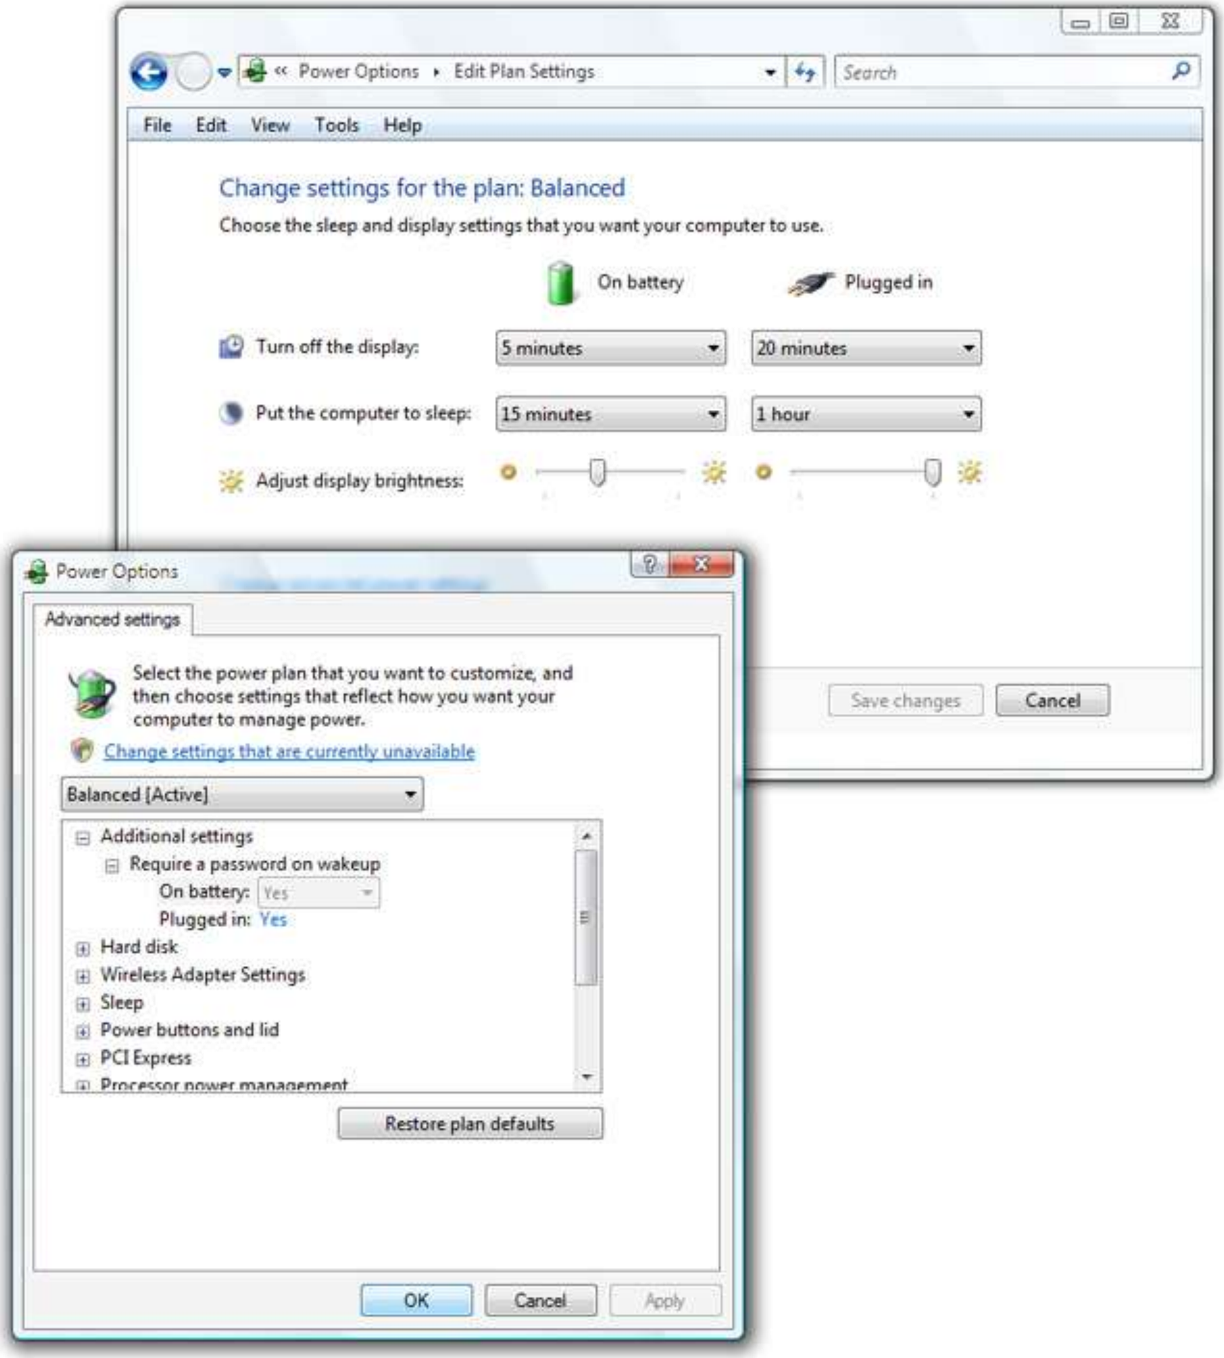

Change settings for the plan: Balanced

Choose the sleep and display settings that you want your computer to use.

	On battery	Plugged in
Turn off the display:	5 minutes	20 minutes
Put the computer to sleep:	15 minutes	1 hour
Adjust display brightness:		

Select the power plan that you want to customize, and then choose settings that reflect how you want your computer to manage power.

Change settings that are currently unavailable

Balanced [Active]

- Additional settings
  - Require a password on wakeup
    - On battery: Yes
    - Plugged in: Yes
- Hard disk
- Wireless Adapter Settings
- Sleep
- Power buttons and lid
- PCI Express
- Processor power management

Apple’s Snow Leopard operating system presents power management settings in a much simpler layout. The two primary options are adjusting sleep time for both the computer and also the display seperately (for the iMac in this case). Apple also offers options for spinning down the harddrive when not in use, wake for network use, setting the power button as a sleep button, automatically reducing the screen brightness upon entering sleep mode automatically and booting up the computer after a power outage.

Figure 16: Apple power management settings.

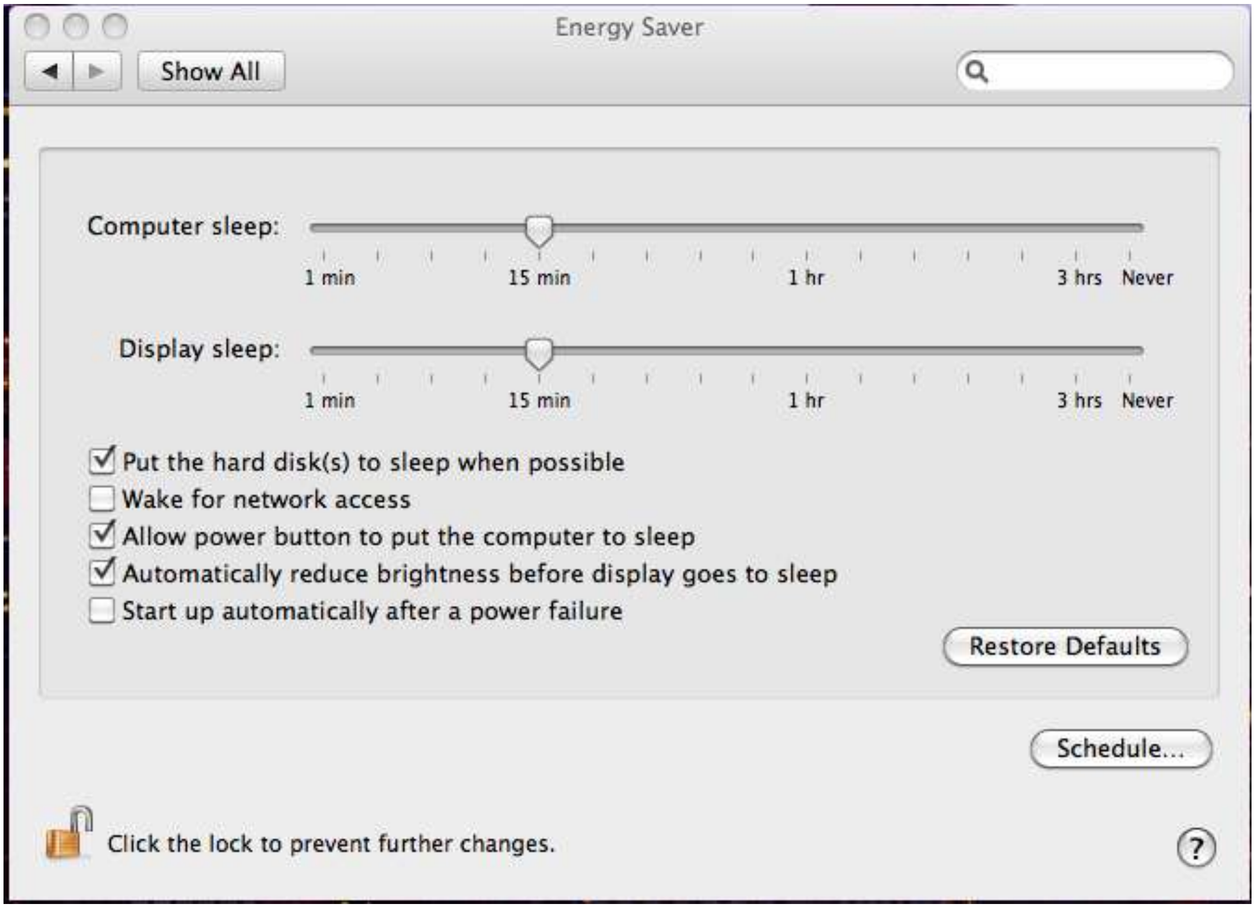

Computer sleep: 15 min

Display sleep: 15 min

☑ Put the hard disk(s) to sleep when possible
☐ Wake for network access
☑ Allow power button to put the computer to sleep
☑ Automatically reduce brightness before display goes to sleep
☐ Start up automatically after a power failure

Click the lock to prevent further changes.

We tested two Windows PCs, one laptop and one desktop, with the same power management schemes applied to observe any differences (Figure 17 and Figure 18). We found that the desktop tested had a clear and dramatic drop in power consumption when entering sleep mode, while the laptop did a much better job of incrementally reducing its power consumption before fully entering sleep mode. As shown in Figures 17 and 18, users have a number of opportunities for reducing power consumption incrementally before their desktop or laptop computers move completely into sleep or hibernate modes. In most cases, the latency associated with these individual steps is quite small, allowing users to capture meaningful energy savings during the work day during the short periods of time between more intensive utilization of their computers. However, manual modification to default power management settings is needed to fully capture the resulting benefits.

Figure 17: Power management example; maximum power saver settings in desktop PC.

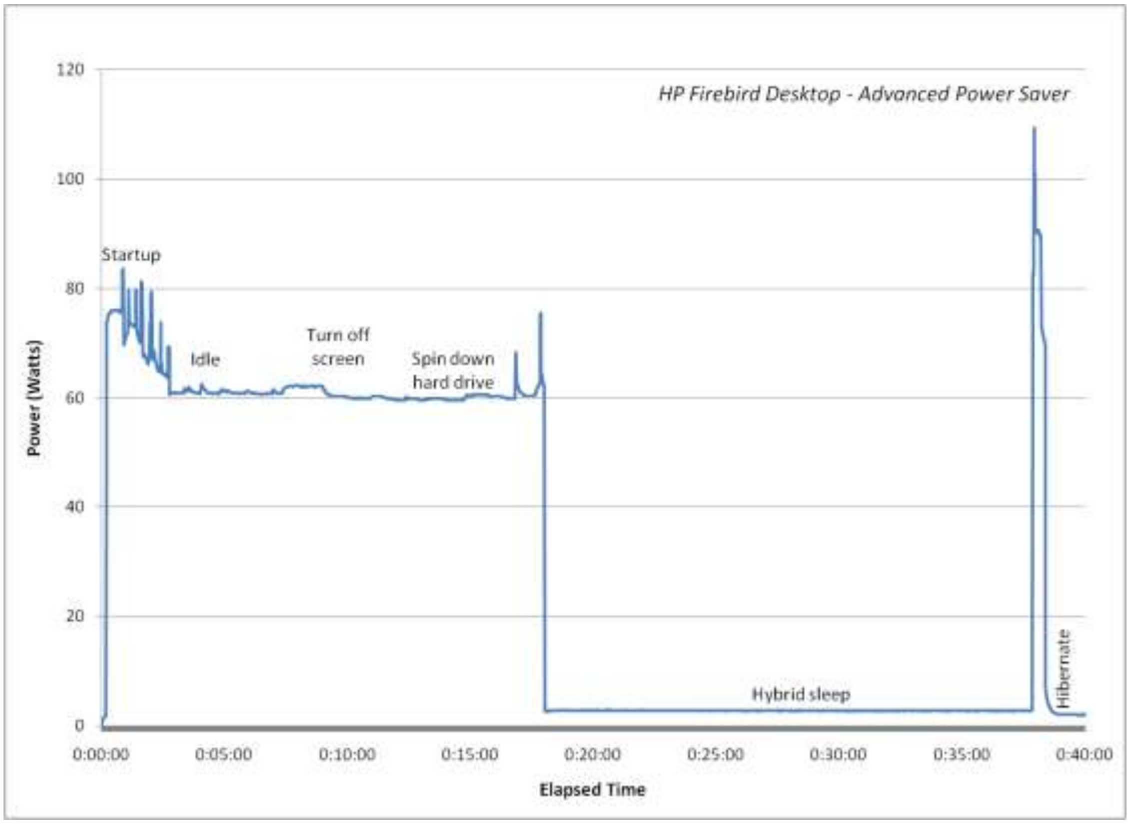

Figure 18: Power management example; maximum power saver settings in laptop PC.

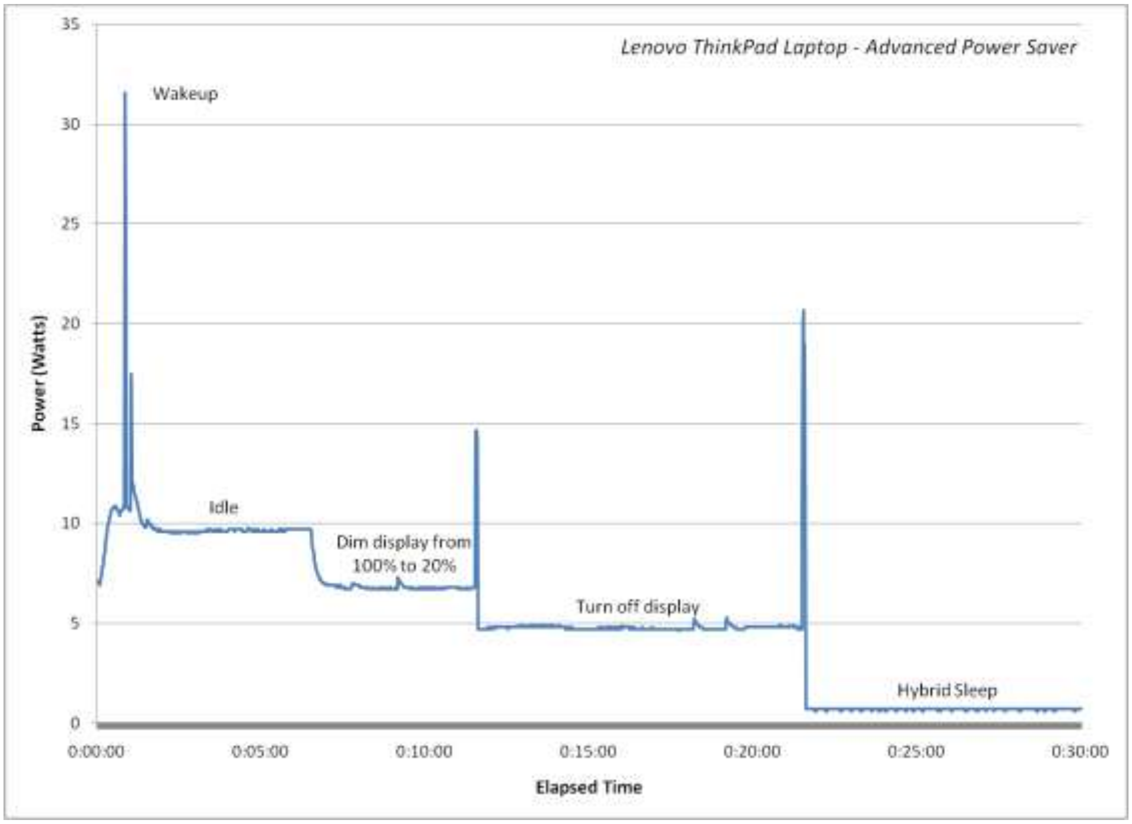

Latency

Separate from the power consumption metrics, we also investigated latencies, to determine the wait time imposed by different devices during boot up, shut down, or sleep modes.

First we focused on the traditional boot times of computers (Figure 19). Even among these four conventional computers, we observed large percentage differences in the latencies associated with going to sleep (via user command), waking from sleep, and booting from “off.” We also considered comparable metrics for less traditional computing devices like the iPad2 and smart phones, for which observed boot times range from a few seconds to virtually instantaneous.

The multimedia PC and Lenovo laptop both utilize solid state hard drives, which probably reduced boot times significantly, but boot and wake times are still fairly long. Operating systems across these platforms can differ significantly in size and complexity, explaining at least some of the observed time differences.

Figure 19: Latency of computers and associated power draw.

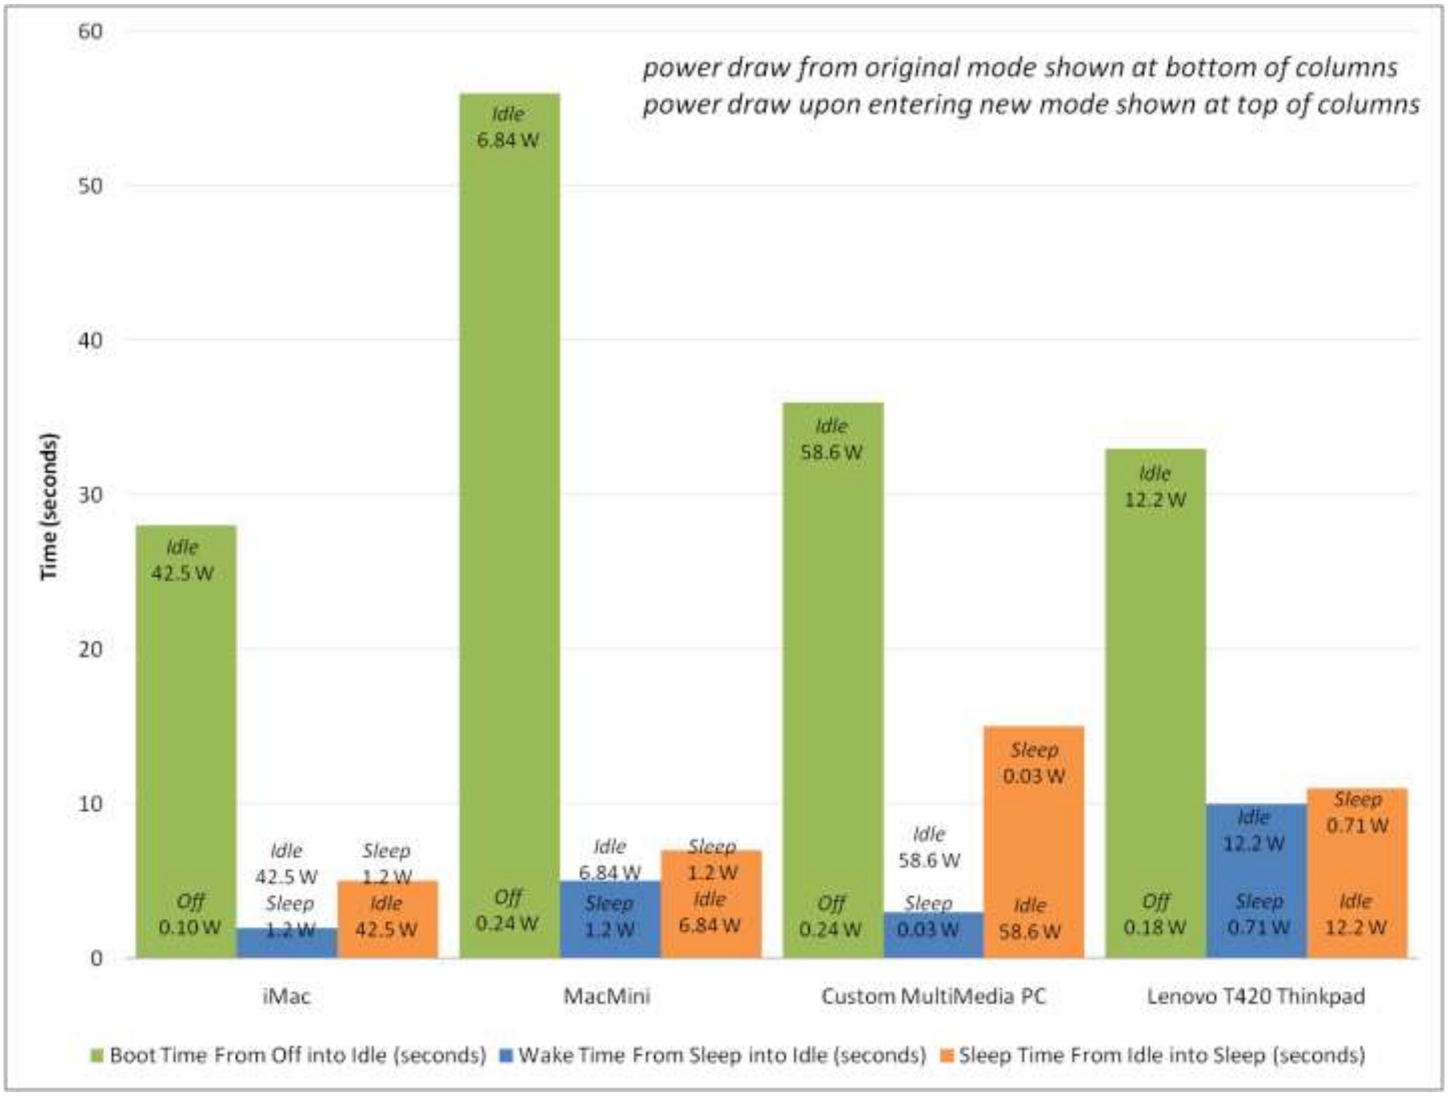

Second, we expanded our latency measurements to other edge devices. Below are latency times for booting a computer from off mode into idle mode, and associated power draw (of idle mode) (Figure 20). The noteworthy finding is that pay-TV set-top boxes have much lengthier boot times than any other edge device we tested. The quickest boot time we examined was the Wii at 10 seconds. Game consoles, in general, had very fast boot times when compared to PCs and set-top boxes.

We also subjected the same devices to a latency test from waking from sleep to idle mode (Figure 21). Unlike their associated boot times from true off mode, pay-TV set-top boxes have very quick wake times from sleep. This is because set-top boxes retained network connectivity during this state (which isn't a true sleep mode) but operate at nearly the same power as in active mode. The three to five minute boot times for set-top boxes are unacceptable for most U.S. consumers, so set-top boxes are almost always left powered on.

Solid state hard drives in both the multimedia PC and the Lenovo laptop significantly reduced boot up and wake times in those computers. Other anecdotal latency measurements, such as opening applications and loading games, showed very fast response times as well. It appears that SSD technology, while costly, offers both energy efficiency and latency benefits. In terms of game consoles, only the Wii has a true sleep mode. The Xbox 360 and PlayStation 3 only have auto-power down capability.

Figure 20: Latency and power consumption of select devices; from off to idle.

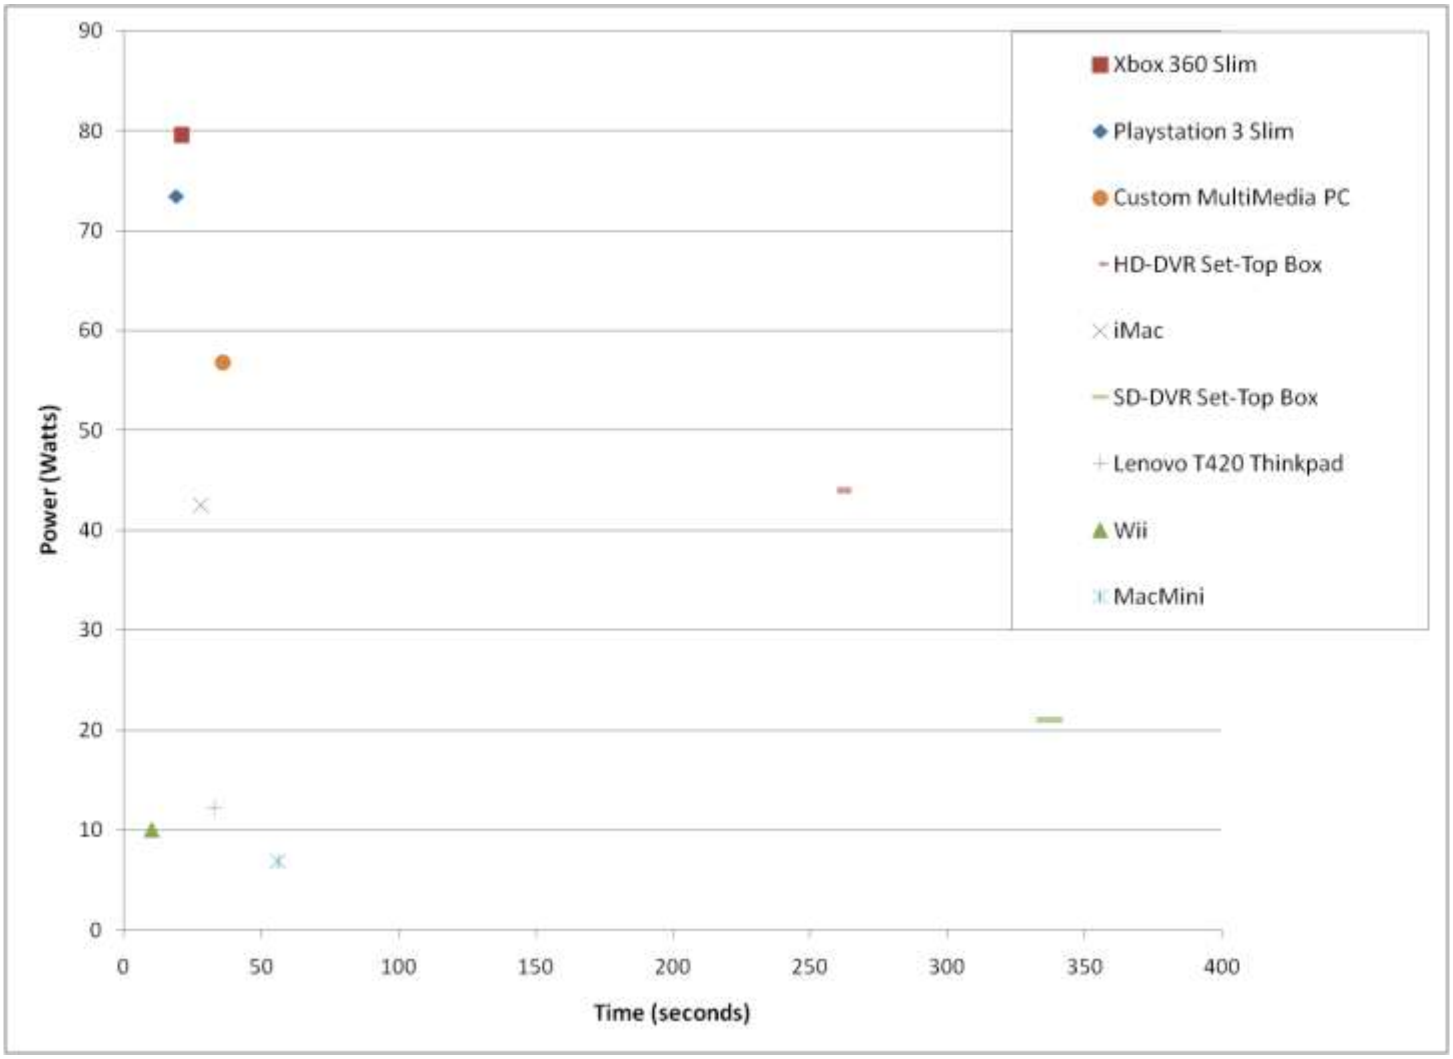

In general, we did not find profound energy impacts of latency, per se. The bigger effect is on usability. If users experience latency times of more than a minute to activate various devices associated with Internet browsing or TV watching, for example, it is increasingly likely that they will never switch those devices fully off. Thus, the real energy impact of latency is associated with duty cycles – users will forgo power management opportunities in favor of greater convenience.

Figure 21: Latency and power consumption of select devices; from sleep to idle.

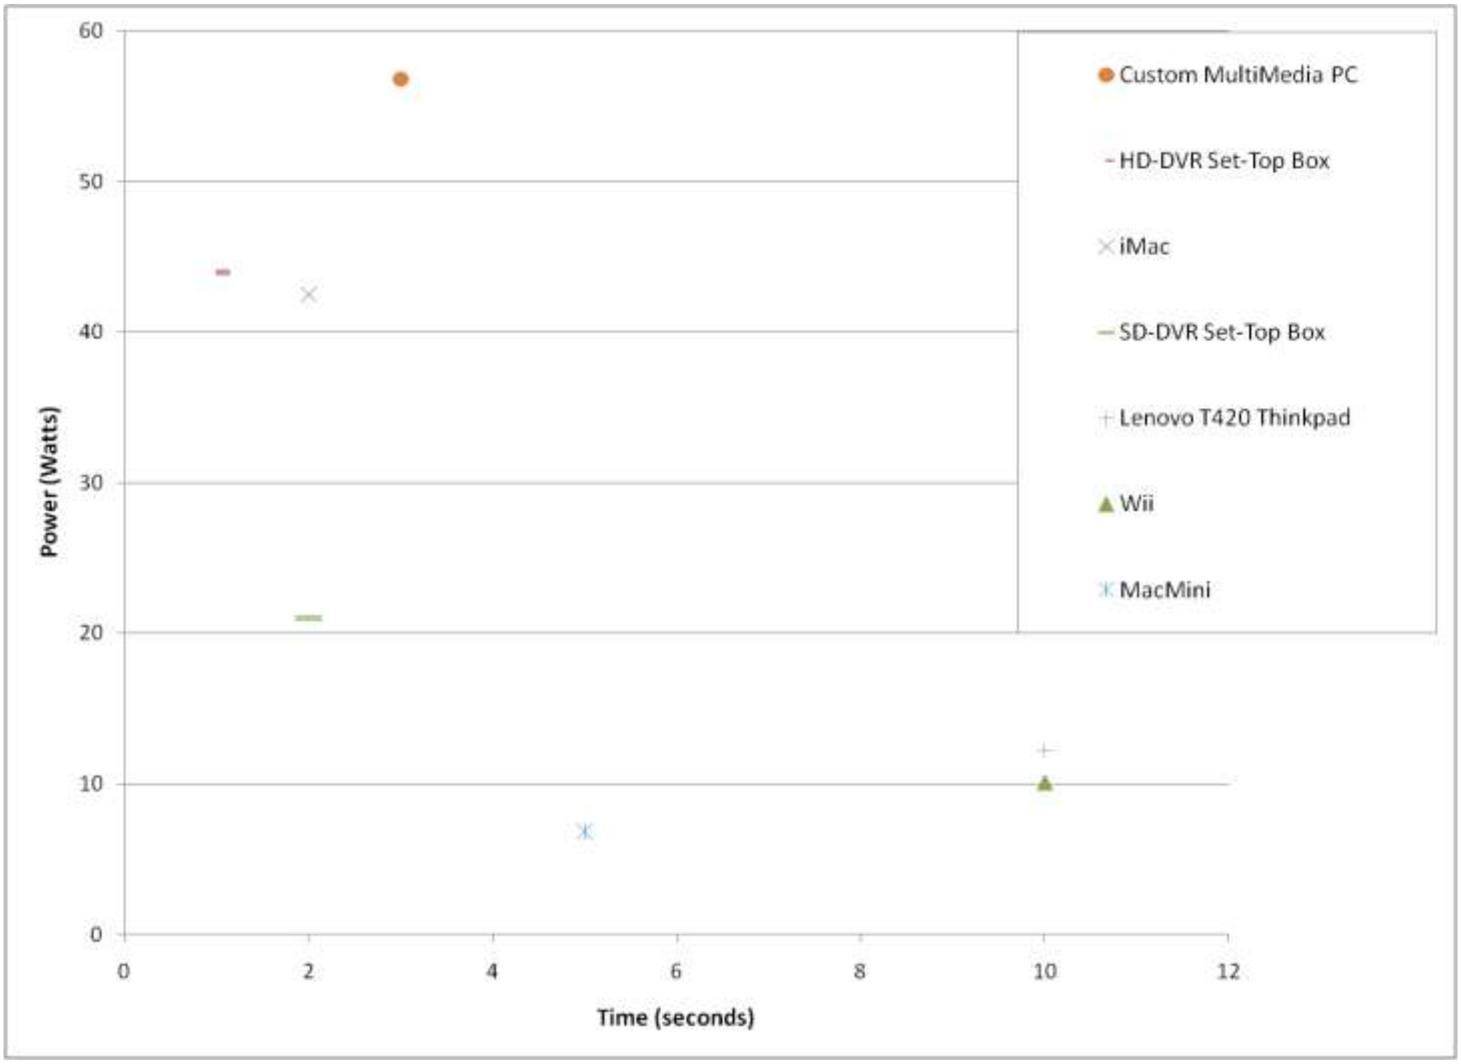

Converting latency times to monetary values and thus real consumer behavior is difficult. However, an order of magnitude analysis is informative. If we assume an opportunity cost at the average wage rate of $20 per hour (personal time is generally valued less, but company time would generally be valued more because of overhead), and a cost of 12 cents/kWh, it is possible to estimate whether energy savings pay for additional time spent.

Putting a computer to sleep versus letting it idle saves approximately 50 W, so if this is for more than 16 hours, this would have a value of 10 cents. The latency of going from idle to sleep and back from sleep to idle is approximately 10 seconds, which has a time value of 7 cents, which is barely cost-effective. This can be made more cost effective by having the computer automatically go to sleep, saving the user time of going from idle to sleep, and only losing a small amount of energy savings. Technically, the time lost is the time of executing the commands to go to sleep, rather than the time it takes to scale the power down, but this introduces uncertainty in measurement.

Turning the computer off versus putting it to sleep saves approximately 1 W, so if this is over 16 hours, this would have a value of 0.2 cents. However, the increased latency of coming from off to idle versus sleep to idle is approximately 30 seconds, which has a time value of $.20, or two orders of magnitude greater. This analysis can be extended to the question of whether consumers will switch to lower power devices that have greater latency.

3.2. Power Consumption of Network Equipment

Findings on network equipment suggest that routers and switches do in fact scale power meaningfully with the number of connected ports, but do so less consistently with changing network traffic speeds. Curiously, many of the units we tested consumed the same amount of power with no traffic on their network as they did when operating at gigabit speeds.

Baseline Tests

As described in the methodology section, we first set out to establish a few baseline tests for the network equipment we purchased. For switches, we simply powered them on with no traffic from edge device connections and recorded their power draw (Figure 22). We recorded the number of ports and Mbit/s capability of each device.

The 16-port Gigabit switch that we purchased, which offered the most functionality, consumed the most power of the five switches. On the other hand, a 4-port switch consumed almost twice the amount of idle power as a 8-port switch of the same 10/100/1000 Mbit/s capability.

Figure 22: Baseline test findings of switches; no traffic.

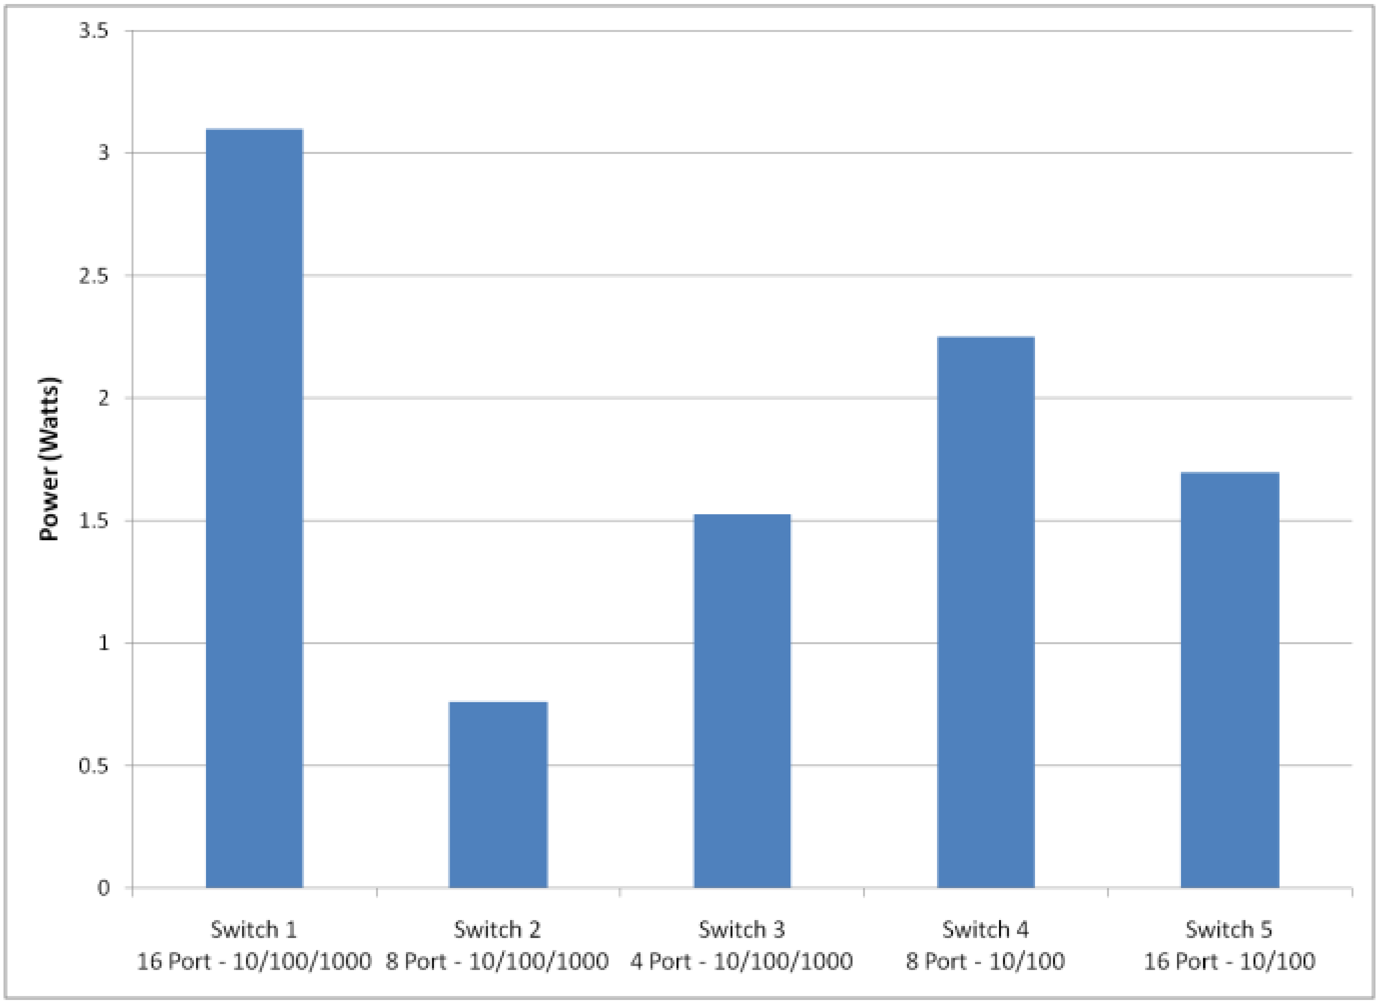

We did the same for routers, but also measured the incremental power consumption when Wi-Fi was enabled – again with no traffic or edge device connections (Figure 23). We recorded power consumption for Wi-Fi disabled, for the 2.4 GHz band, 5 GHz, and both activated (for one router). As we expected, enabled Wi-Fi increased the power consumption of all the routers we tested. Activating the 5 GHz radio band instead of the 2.4 GHz also added power. The primary energy savings opportunity appears to be designing routers to drop into some sort of much lower power state during periods of time when no traffic is occurring.

Figure 23: Baseline test findings of routers with and without WiFi enabled; no traffic.

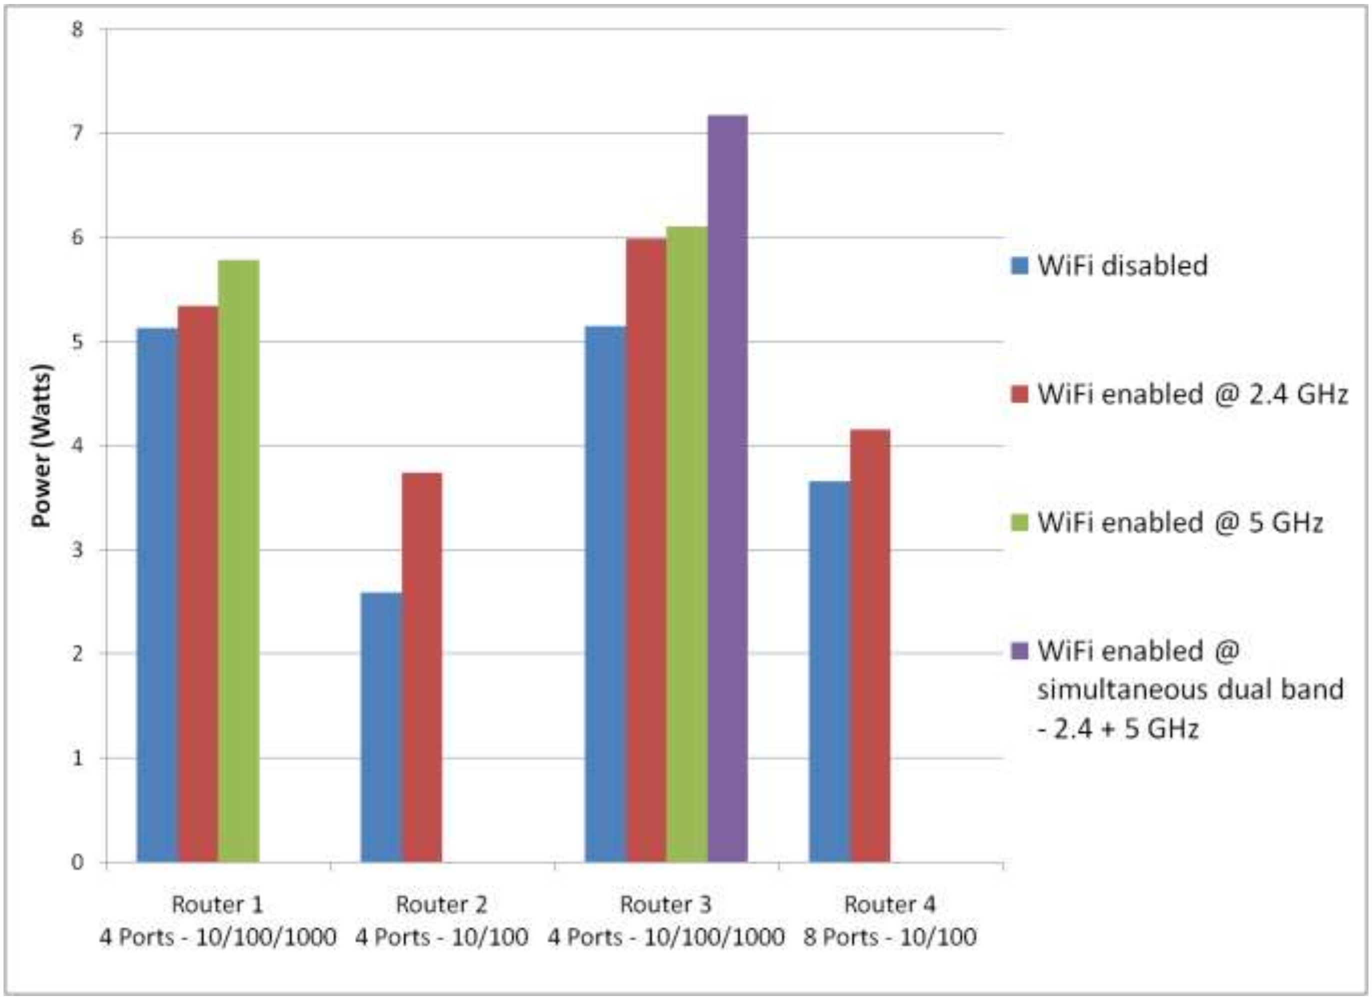

Peer-to-Peer File Transfer Testing

Having established a basic understanding of metering network equipment, we set out to subject each piece of network equipment to different connection speeds and file transfers (Figure 24). We connected two edge devices to one another to create a closed network using a network device. We took power measurements of the network device when 1) no traffic was being generated, 2) connection speed was set at 10 Mbit/s, 3) connection speed was set at 100 Mbit/s and 4) connection speed was set at 1000

Mbit/s. It is important to note that we configured the maximum connection speed at the *edge device*.[7] Unfortunately, the software did not control data rate transfer or connection speeds. We were not monitoring the true rate of data traffic. Due to errors with the network traffic generator software, we only tested this configuration with switches. In general, the power scaling effect was quite modest as connection speeds increased. It was rather perplexing that a number of devices tested with no traffic recorded higher power consumption than when traffic was enabled.

Figure 24: Power consumption of network switches during looped file transfer; varying connection speeds.

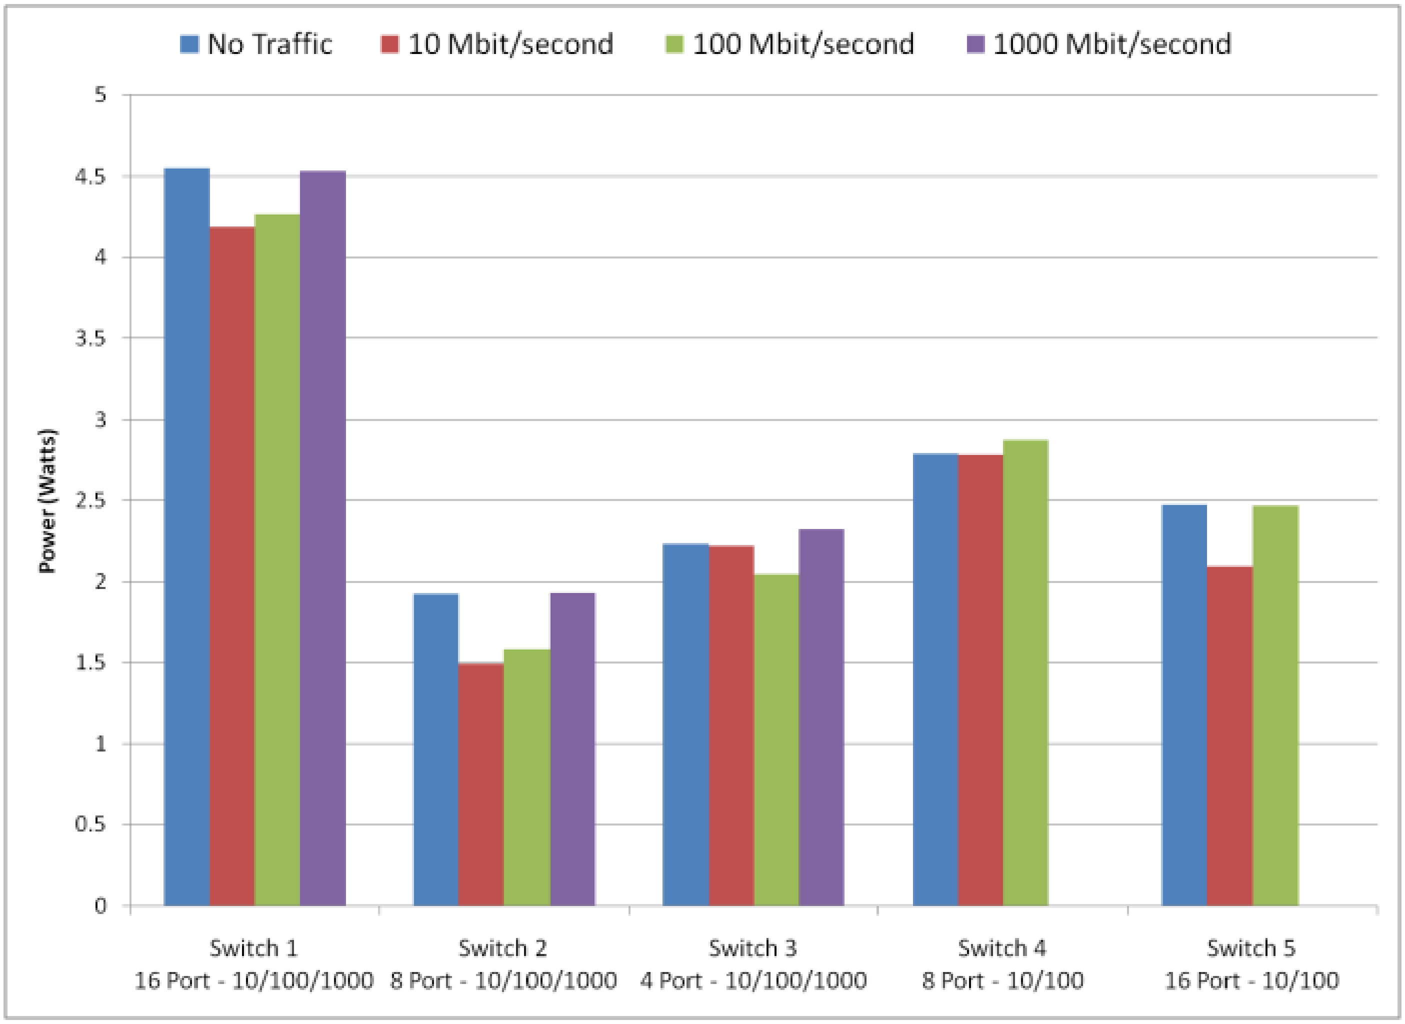

Multiple Edge Devices Served by Network Device Testing (Ethernet)

We also wanted to test network devices subjected to a full load in terms of maximum number of edge devices (i.e. number of ports used). For these tests, we disabled Wi-Fi capability. The network device was connected to the internet via our office LAN network. We connected edge devices via Cat 6 Ethernet cables to the network device, and recorded power as each new device was added to the network. For these tests, the edge devices were all idling with established network connectivity. We did not vary connection speeds.

---

[7] As opposed to the network equipment firmware or within the network traffic generator software. For each edge device, we altered the maximum connection speed in Windows 7. In Window 7 device manager, we set the maximum link duplex speed to the associated value (i.e. 10/100/1000 Mbit/s).

Figure 25: Multiple edge devices connected to network device.

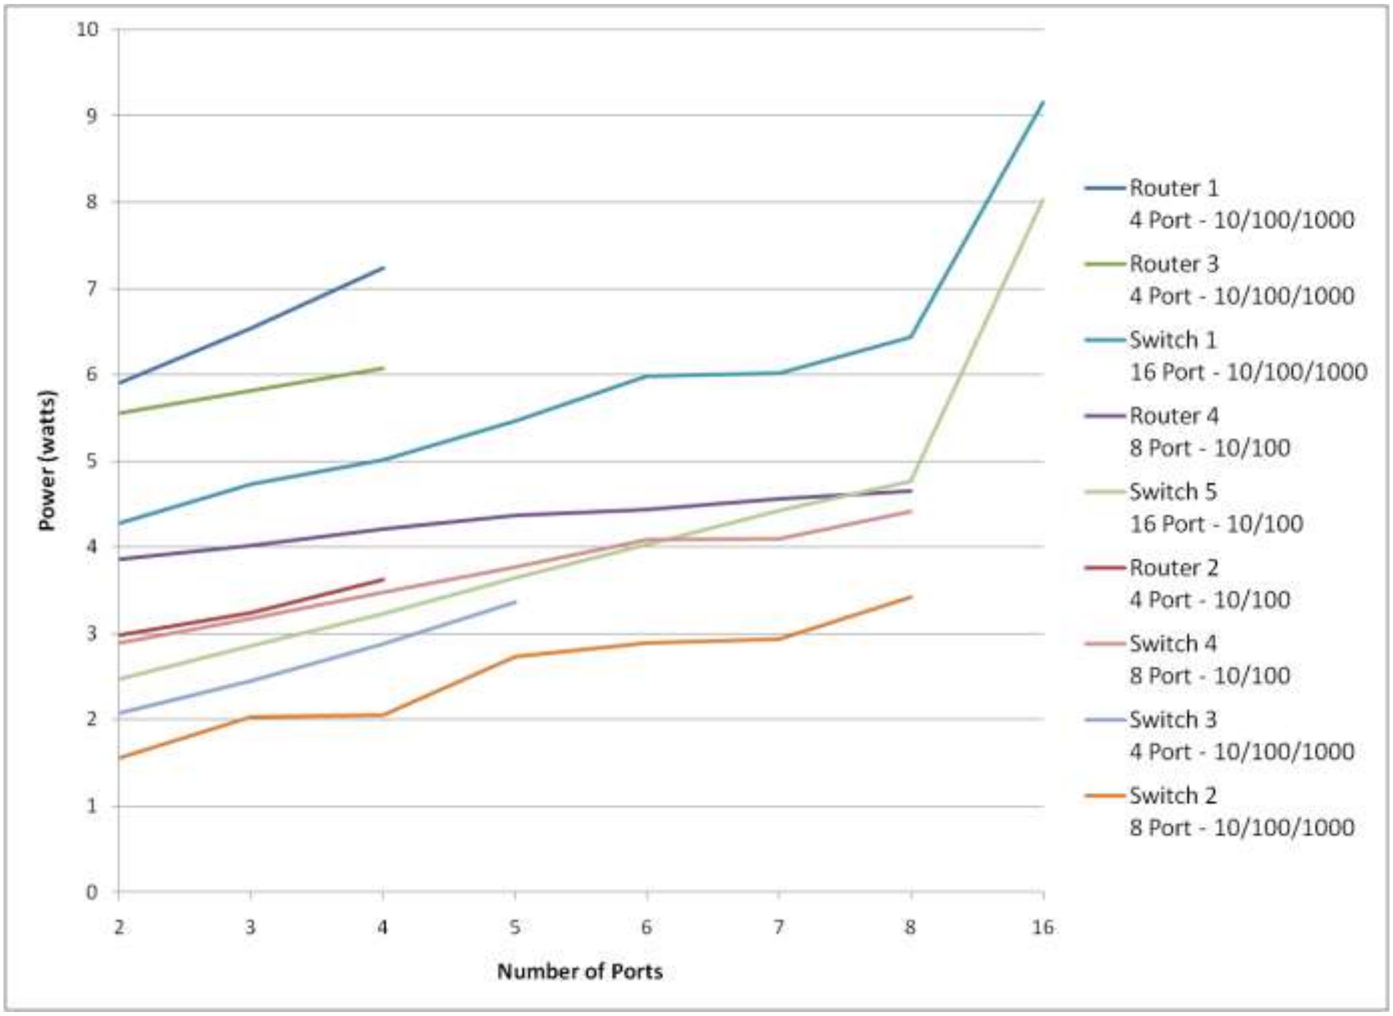

Multiple Edge Devices Served by Network Device Testing (Wi-Fi)

We hypothesized that network device power consumption would not be affected by the number of devices it hosted via Wi-Fi. Instead, we imagined bandwidth would be constrained. We connected edge devices via Wi-Fi to the network device, and recorded power as each new device was added to the network. For these tests, the edge devices were idling with established network connectivity. We did not vary connection speeds. Findings are shown below in Figure 26:

Figure 26: Multiple edge devices served by network device testing (WiFi)

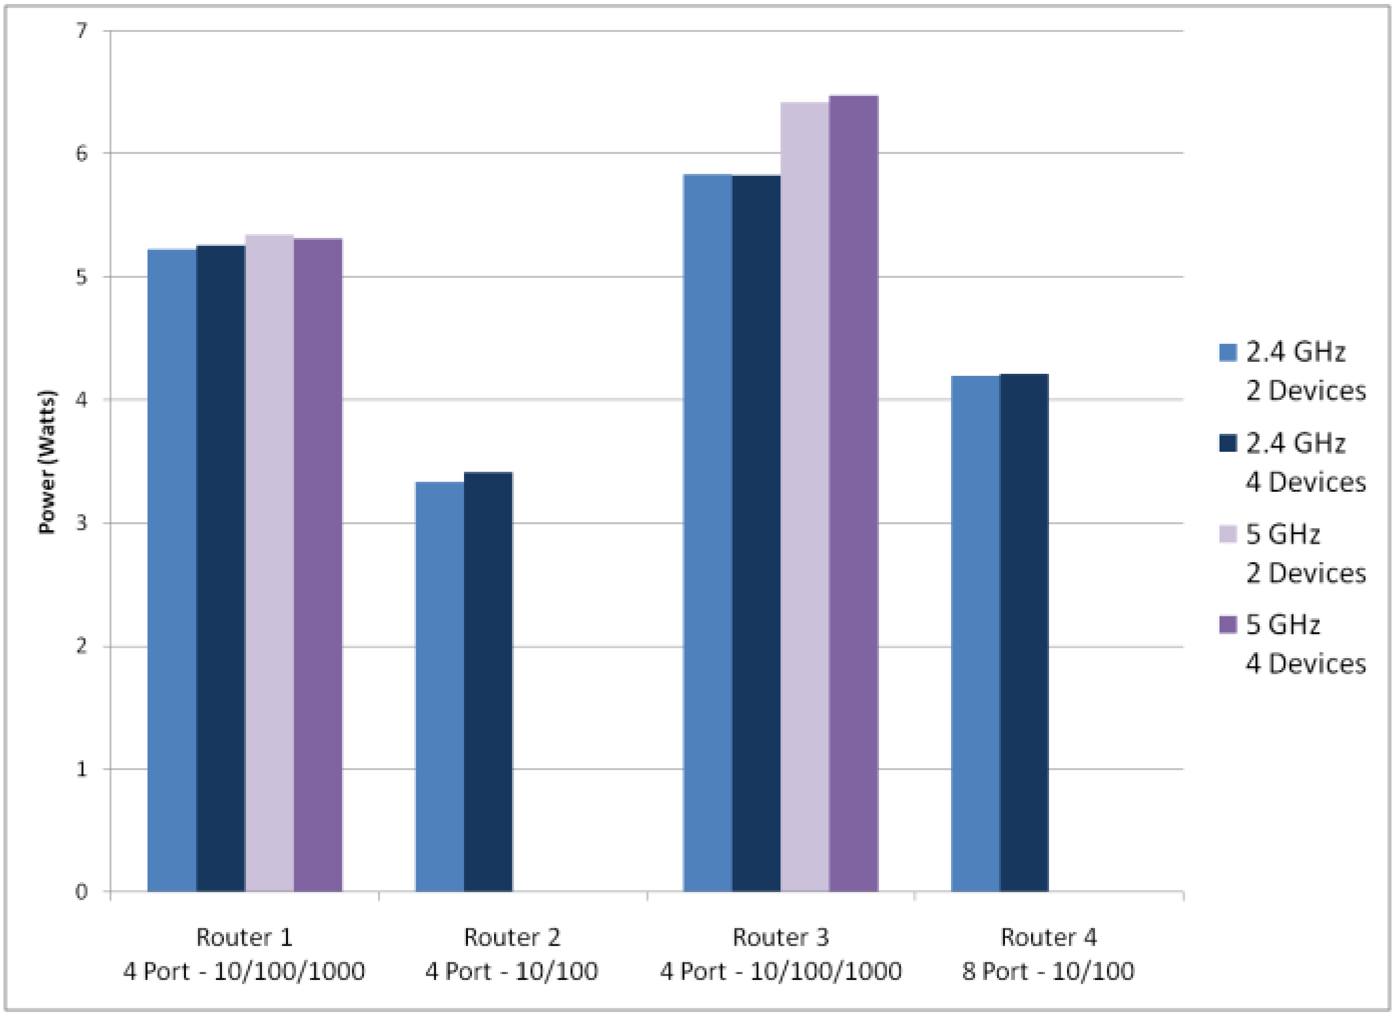

Note: Router 2 and Router 4 were not capable of dual band.

To test Wi-Fi power scaling, we sent a 500 MB file from one edge device, through the router, to another edge device. In this test, the independent variable was the protocol and radio frequency used by the Wi-Fi router. We compared a legacy protocol, 802.11g, with the current 802.11n protocol (Figure 27). Equivalent to “g,” the “n” protocol can transmit via 2.4Ghz, but can additionally use the 5Ghz band.

Similar to our findings for wired networking equipment, only minimal power scaling was evident. Some protocols were able to transmit the same amount of data more rapidly than others, and could theoretically save energy if they dropped into a lower power-consuming state thereafter. However, because power consumption tended to be fairly steady in these devices regardless of the amount of traffic flowing, speed was a performance issue but not really an energy issue.

Figure 27: Power consumption of different WiFi frequencies of one router.

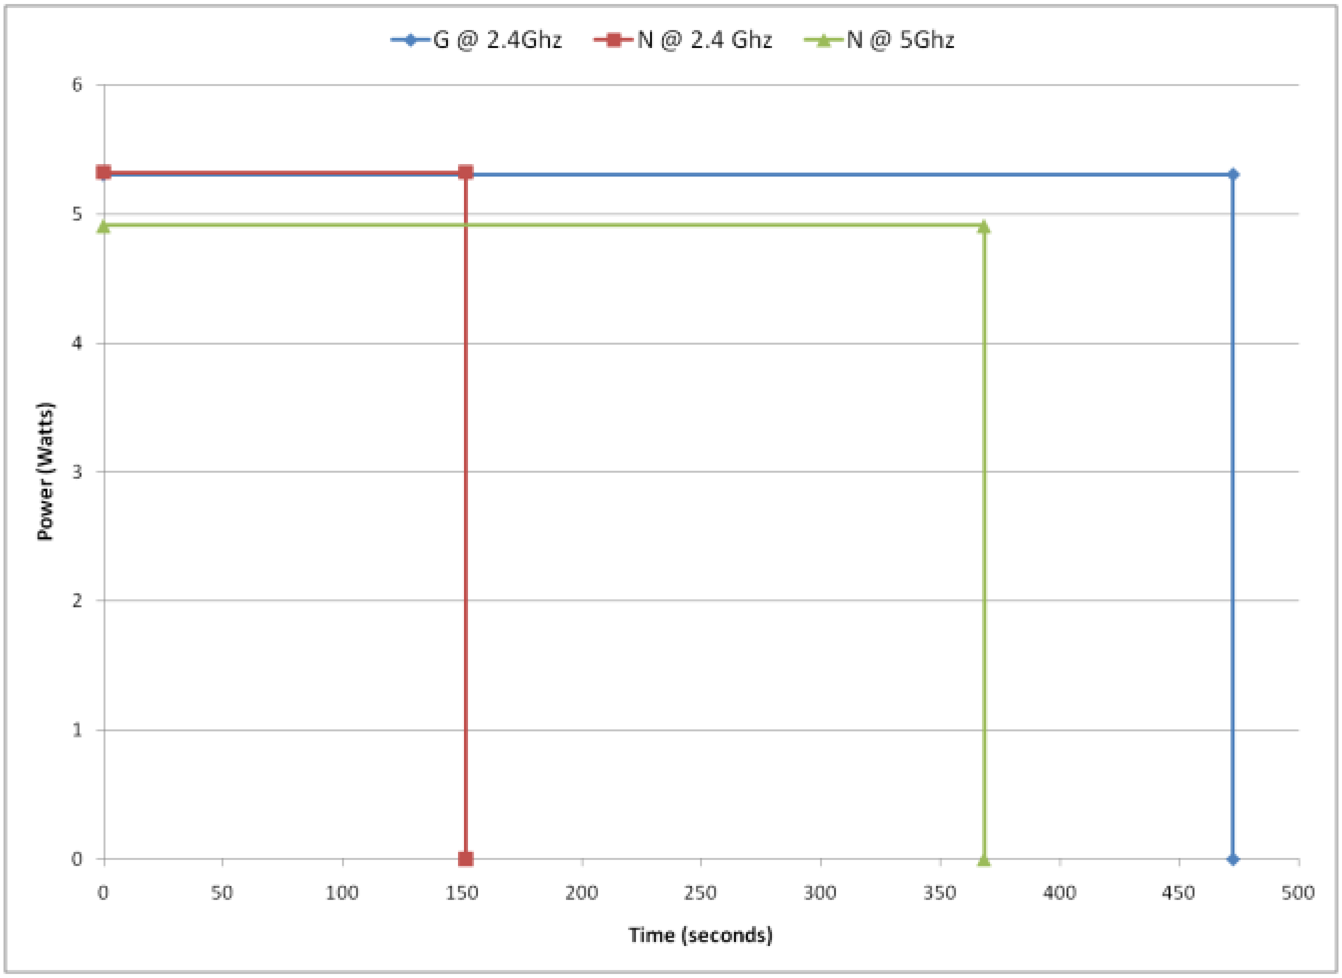

4. Conclusions

We conclude that most of the devices that use a small amount of power in an absolute sense are doing so with fundamental efficiency improvements to their hardware, rather than with power scaling per se. The devices that have the highest power tend to exhibit some degree of power scaling, but even when idling, they cannot approach the power consumption levels of their more efficient competitors. So, while the promise of power scaling is real, little of that promise has been achieved in many of the consumer electronics products we tested.

Power scaling remains an idea that can add much to the way energy efficiency specifications are developed and implemented. Not only can wattage limits be specified for standby, sleep, and idle modes for various electronic products, for example, but policymakers could also consider introducing latency and performance considerations. For example:

- How many minutes can elapse between when no activity occurs with a device and when it drops into a sleep mode?
- How rapidly must a product return from sleep mode to idle mode?
- By what percentage must the power consumption drop between active or maximum performance mode and idle mode?
- To what extent can products be expected to maintain minimal network connectivity during sleep or hibernate modes?

Finally, we recommend that policymakers pursue additional research on the topic of network power scaling specifically. Standard protocols should be developed that can assess power consumption of network routers, hubs, and switches as the amount of network traffic flowing through them varies. Thereafter, labeling programs and eventually MEPS could be developed to promote those design approaches that meaningfully scale power consumption to network activity.