DOCKETED

Docket Number:	14-AAER-02
Project Title:	Computer, Computer Monitors, and Electronic Displays
TN #:	210857
Document Title:	California IOUs Standards Proposal Addendum Computers
Description:	This document is to be filed in 14-AAER-02, not in 12-AAER-2A as stated on the document.
Filer:	Patty Paul
Organization:	Energy Solutions
Submitter Role:	Public
Submission Date:	3/24/2016 3:18:21 PM
Docketed Date:	3/24/2016

Computers

Codes and Standards Enhancement (CASE) Initiative
For PY 2013: Title 20 Standards Development

Analysis *of* Standards Proposal for
Computers

Addendum to submittal on August 6, 2013

Docket #12-AAER-2A

October 27, 2014

Prepared for:

PACIFIC GAS & ELECTRIC COMPANY

SOUTHERN CALIFORNIA EDISON

SAN DIEGO GAS AND ELECTRIC

SOUTHERN CALIFORNIA GAS COMPANY

NATURAL RESOURCES DEFENSE COUNCIL

Prepared by:
Nate Dewart, ENERGY SOLUTIONS
Pierre Delforge, NATURAL RESOURCES DEFENSE COUNCIL
Peter May-Ostendorp, XERGY CONSULTING
Joe Colett, ECOVA
Catherine Mercier, ECOVA
Vojin Zivojnovic, AGGIOS
Davorin Mista, AGGIOS

This report was prepared by the California Statewide Utility Codes and Standards Program and funded by the California utility customers under the auspices of the California Public Utilities Commission.

Table *of* Contents

APPENDIX A: DUTY CYCLE STUDIES

APPENDIX B: REAL-WORLD ADJUSTMENT FACTOR TESTING RESULTS

APPENDIX C: NOTEBOOK DEMONSTRATION PROJECT

1 Executive Summary

The Pacific Gas and Electric Company (PG&E), Southern California Edison (SCE), Southern California Gas (SCG), San Diego Gas & Electric (SDG&E) Codes and Standards Enhancement (CASE) Initiative Project seeks to address energy efficiency opportunities through development of new and updated Title 20 standards. Individual reports document information and data helpful to the California Energy Commission (CEC) and other stakeholders in the development of these new and updated standards. The objective of this project is to develop CASE Reports that provide comprehensive technical, economic, market, and infrastructure information on each of the potential appliance standards. The CASE Team has conducted additional testing and analyses of market information to update key aspects of the standards proposal presented in the initial Computers CASE report submitted in August 2013 (CA IOUs 2013a).The primary adjustment is a shift of the standards levels and effective dates to Q4' 2016 for Tier 1 and Q4' 2018 for Tier 2, as well as the cost-effectiveness. The updates include:

- Duty cycle (increase time in active/idle mode, based on review of duty cycle studies)
- Stock and shipments (decrease, based on review of IDC sales data)
- Real world adjustment Factor (roughly stays the same, supported by additional laboratory testing)
- Market Changes (ENERGY STAR-calculated TEC has decreased since 2012, based on ENERGY STAR QPL analysis)
- Technologies and Best Practices for Energy Use (summary of findings and recommendations from notebook demonstration project)
- Standards proposal (more stringent for Tier 1 and Tier 2, based on market improvements)

3 Energy Usage

To estimate annual energy consumption, ENERGY STAR uses typical energy consumption (TEC) in kWh/yr which is calculated by multiplying the wattage in each of these four modes by the estimated percentage time the computer is in each mode, or mode weighting. The CASE Team recommends aligning with the ENERGY STAR 6.0 metric, test procedure and mode weightings for calculating energy consumption for purposes of the standard. For life-cycle benefit cost analysis for cost-effectiveness and statewide energy savings, the CASE Team recommends using a different mode weighting and account for a real-world adjustment factor discussed below.

3.1 Duty Cycle

The duty cycle of computers is determined both by the extent of the computer's power management settings and by the extent the user manually switches the modes, and therefore varies considerably by user, though general usage trends have been documented. As discussed in the CA IOUs' response to the Invitation to Participate (CA IOUs 2013b), there are several studies which sample PC user behavior in both residential and commercial settings that capture an estimation of daily duty cycles with and without power management (PG&E 2010; Pigg & Bensch 2010; TIAX 2007, Fraunhofer 2011, ECMA-383, Chetty 2009 et al, Microsoft 2008). During the Invitation to Participate, the CASE Team recommended aligning with ENERGY STAR 6.0 for a sector-weighted duty cycle for notebooks, and making a modification to the duty cycle for desktops based on the available research, which would increases the percentage of time in idle mode (short and long).

Given additional studies that have been released since the previous submittal—Greenblatt 2013, Fraunhofer 2014 and the forthcoming 2014 California Plug Load Research Center study—the CASE Team continues to recommend that CEC review all the available studies and develop a mode weighting that accounts for the limitations of these studies including sample size, survey methods, year conducted and a differentiation in short idle and long idle, to most accurately estimate current computers usage. We have performed such a review (see Appendix A) and based on a weighted average by sample size (each study assigned a value of 1-10, rather than absolute sample size) recommend a revised mode weighting for conventional desktops, integrated desktops and thin clients (the other form factors remain the same), however, this revised mode weighting does not account for the forthcoming California Plug Load Research Center study. Upon further review of this study, the CASE team may recommend a further revision to the duty cycle.

See below in Table 3.1 for the current revised mode weighting. Note that these values do not reflect the subdivision between short idle and active modes. This differentiation is accounted for in the real-world adjustment factor, described further in Appendix A.

Table 3.1 Estimated Duty Cycle For Each Form Factor

Mode	Conventional Desktops, Integrated Desktops and Thin Clients	Notebooks	Workstations	Small-scale Servers
Off	30%	25%	35%	0%
Sleep	10%	35%	10%	0%
Long Idle	20%	10%	15%	100%
Short Idle	40%	30%	40%	0%

3.2 Real-World Adjustment Factor To Estimate Computer Energy Use, Savings And Cost-Effectiveness

As highlighted in the CA IOUs computers CASE Report (2013b) and in NRDC's response to the Invitation to Participate (2013), the annual energy consumption of computers is conventionally estimated using the ENERGY STAR Typical Energy Consumption (TEC) metric, which uses idle mode as a proxy for active mode power, however, this is no longer appropriate for two reasons:

1) The new ENERGY STAR 6.0 specification now includes both long and short idle modes (instead of just short idle) increasing the difference between typical active power and the average reported idle power using the ENERGY STAR test procedure;

2) Modern computers are able to better scale power down when inactive than recently, leading to a higher difference between idle and active power.

The IOUs have performed additional testing to supplement the values the CASE Team recommended in the CASE Report summarized in the testing memo, found in Appendix B. Based on this testing, the CASE Team recommends that ENERGY STAR Version 6.0 measurements be increased by:**35% for notebooks**, **26% for integrated desktops and 13% for conventional desktops and workstations** when calculating energy use, savings and cost-effectiveness. The lower adjustment factors for desktops relative to notebooks reflect the testing-based evidence that desktops, while benefitting from power scalability technology, do not implement it to the same extent as notebooks due to a lack of battery life incentive.

Using ENERGY STAR Version 6.0 test method without adjustments would significantly underestimate the actual energy consumption of computers in California, the savings potential from standards and their cost-effectiveness. Given that a revised test method that would appropriately account for this energy consumption does not currently exist, we recommend that CEC uses the ENERGY STAR Version 6.0 test method and adjusts the results using the "real-world adjustment factor" values indicated above.

It is important to note that the standards levels do not account for the real-world adjustment factor; the CASE Team proposes keeping the ENERGY STAR 6.0 TEC calculations as is.

3.3 Market Changes

Technology advances in computers since the CA IOUs originally submitted the 2013 computers CASE report have had a marked impact on energy efficiency. In particular, the adoption of new processor and chipset architectures across all computer form factors and performance categories — namely, Intel's Haswell and AMD's Trinity and Richland architectures — has brought with it continued efficiency gains. As noted in our 2013 CASE report and subsequent supplemental report, Haswell was expected to reduce idle power levels by approximately 30%, and this technology now dominates most mainstream computer categories. This technology was only beginning to become available when the CA IOUs submitted the supplemental technical report in early 2014 and was not reflected in the bulk of measurements that formed the basis of our original standards proposal, which dated from the 2012 *Cost-Effective Computer Efficiency* research (PG&E 2012).

Due to this trend, the average TEC of computers has since decreased. While not representative of the whole market, analysis of the ENERGY STAR computer qualifying product lists (December 2012 and October 2014) by comparing the average measured TEC across integrated graphics categories (0, I1, I2,and I3) indicates that the greatest improvements have been made in notebook products, approximately a 48% reduction over this time. Traditional desktop and integrated desktop products have experienced 38% and 36% reductions in TEC, respectively. Note that, since display power was not measured for qualifying ENERGY STAR v5.0 products in the 2012 timeframe, we have adjusted for display power in our analysis of the current 2014 ENERGY STAR QPL to ensure a fair comparison. The average TEC values by product category for 2012 and 2014 as well as the percent change in TEC is presented in Table 3.2 below.

Table 3.2: ENERGY STAR QPL Comparison 2012 to 2014

	Computer Performance Category				
	0	I1	I2	I3	All
Desktops (Conventional)					
2012 # Datapoints	16	33	25	85	159
2014 # Datapoints	18	83	66	101	268
2012 Average TEC	118.9	139.0	164.9	171.8	158.6
2014 Average TEC	67.6	85.1	110.2	105.7	97.9
% Difference	-43%	-39%	-33%	-38%	-38%
Desktops (Integrated)					
2012 # Datapoints	13	27	9	34	83
2014 # Datapoints	30	66	37	52	185
2012 Average TEC	90.6	126.6	142.5	125.8	122.3
2014 Average TEC	59.0	80.8	83.4	83.5	78.5
% Difference	-35%	-36%	-41%	-34%	-36%
Notebooks					
2012 # Datapoints	11	65	29	9	114
2014 # Datapoints	109	957	319	417	1802
2012 Average TEC	29.5	28.7	31.6	31.2	29.7
2014 Average TEC	14.3	14.3	13.8	19.8	15.5
% Difference	-52%	-50%	-57%	-36%	-48%

Note that average values do not reflect the whole market nor do they reflect real-world energy use. Utilizing the whole market model with the same methodology as the 2013 analysis and including the revised duty cycle and the real-world adjustment factors, we estimate the current shipment average kwh/yr to be 157 for conventional deskops, and 121 for integrated destkops and 19 for notebooks (excludes display energy). Relative to the revised standards levels described below, the average non-qualifying and qualifying products are estimated in Table 3.3.

Table 3.3 Estimate of Energy Use for Non-Qualifying and Qualifying Products in 2014

Product Class	Average Unit Energy Consumption of Non-Qualifying Products (kWh/yr)	Average Unit Energy Consumption of Qualifying Products (kWh/yr)
Desktops – Conventional	217	56
Desktops – Integrated	138	46
Notebooks	20	10

Despite the progress in the underlying hardware platforms used by contemporary computers, we still anticipate significant opportunity for energy savings in this most current computer generation and advocate for additional reductions in TEC requirements relative to those CA IOUs proposed in 2013. The computers CASE report and subsequent 2014 supplemental report (CA IOUs 2014) demonstrated through desktop computer testing that substantial and cost-effective energy savings could be achieved using more efficient hard drives, power supplies, and discrete GPUs, technologies which are available on the market and used by some but not all computers. These efficiency techniques are still relevant in the 2014 market.

4 Market Saturation & Sales

4.1 Current Market Situation

4.1.1 Stock

As pointed out in the CA IOUs 2013 computers CASE Report, a significant dispcrepancy exists between reported stock estimates from Fraunhofer 2011 and calculating stock when multiplying historical shipments by average design life, for both desktops and notebooks. In Fraunhofer 2014, the survey methodology was modified slightly with an added question of whether in the last month had the computer been plugged-in, to more narrowly target the computers being used. This in turn has reduced the estimates, however the stock estimate discrepancy still remains. Given the greater certainty behind this survey data, we recommend utilizing the Fraunhofer 2014 estimates for stock. Utilizing a California / U.S. GDP percentage of 13% (BEA 2012) and an approximate division of 59% to 41% between the commercial and residential PC market (Hamm and Greene 2008) results in an estimate of 23 million conventional desktops, 6 million integrated desktops and 30 million notebooks statewide, respectively in 2013.

Utilizing the calculated approach, the CASE Team estimates there are approximately there at least 28.5 million computers in use in California in 2014: 8.5 million conventional desktops, 5 integrated desktops and 15 million notebooks. This estimate was derived using shipment data from IDC (2012, 2013a and 2013b, 2014), an average design life of 4 years for desktops and 3 years for notebooks, an ENERGY STAR 6.0 QPL (EPA 2014a) split between conventional and integrated desktops (52% and 48%, respectively), and a California / U.S. GDP percentage of 13% (BEA 2012).

Finally, this discrepancy between estimates using the two different methods at roughly a factor of two, at the very least necessitates a re-examination of the estimated design life, especially in the residential sector. These numbers would suggest a higher estimate of up to 7-8 years. One additional consideration is that the Fraunhofer 2014 survey does delineate a distinction between primary computers (53 million desktops and 53 notebooks in the U.S) and secondary (plus) computers (35 million desktops and 40 notebooks in the U.S) in homes. The existence of secondary computers could but not necessarily suggest a reduced duty cycle for some units, as some computers move from primary to secondary over time as they age. Some secondary computers however may get the same amount of use, especially in households with several members. Without additional information, we recommend CEC consider a design life range of 4-8 years for both residential and commercial desktops and notebooks, with perhaps some modification to the duty cycle to account for some secondary computers.

4.1.2 Shipments

Based on 2011 and 2012 growth projections for 2013 and 2014 reported by IDC (2013a, 2014), we estimated 59 million total shipments for desktops, notebooks ("portables" and min-notebooks) and workstations, respectively in the U.S in 2014. There was a 37% / 63% split between desktops and notebooks in 2012 (IDC 2012) in mature markets (U.S. Western Europe, Canada, Japan). Using the California / U.S. GDP percentage of 13% (BEA 2012), not accounting for workstations, there were approximately 3.1 million desktops and 4.9 million notebooks sold in 2013.

5 Technologies and Best Practices for Energy Use

In addition to the energy savings opportunities highlighted in the August 2013 CA IOUs Computers CASE Report, the CASE Team has identified another opportunity described below.

5.1 Real-time Power Management Techniques ("keystroke sleep")

One of the largest opportunities for energy savings in computers may be real-time power management. This consists of optimizing the power state of all components in the device, not just the CPU, at a millisecond level rather than 15-minute level with conventional ACPI. Apple has implemented this technology in its computers and sometimes referred to it as "putting the computer to sleep between keystrokes." AMD is including this approach as one of its three key strategies to meet its "25x20" mobile efficiency goal and AGGIOS has demonstrated similar technology on an IPTV gateway, which shares many components with computers. The CASE Team encourages CEC to consider the ongoing developments of this opportunity in the development of standards. A further description of the IOU-funded real-time power management demonstration project and conclusions are highlighted in Appendix C.

6 Recommended Standards Proposal

Of the energy savings opportunities available, a system-wide, performance-based energy use approach allows manufacturers to select a suite of options, while accommodating for functionality.

Factoring for a revised duty cycle, real-world adjustment factors, experience curves, we estimate the previous standards proposal submitted in 2013 to be even more cost-effective. Moreover, the market improvements discussed in Section 3 suggest that the cost-effective base allowances from the previous proposed standards could be made more stringent, cost-effectively. Revised proposed requirements with a base allowance reduction of approximately 20% more stringent for desktop and notebooks both Tier 1 and Tier 2 and 10% for integrated desktops are shown below in Table 6.1. The IOUs are conducting additional testing on the latest releases of graphics cards to supplement the previous testing results (PG&E 2012b and CLASP 2011), and will submit these when they become available, before the end of the year.

Table 6.1 CASE Team Proposed TEC Base Allowances, and Graphics Adders – Tier 1 and Tier 2

Product Class	Performance Category	Tier 1 – 2016	Tier 2 - 2018	Adder Category(1)	Tier 1 – 2016	Tier 2 – 2018
		Maximum Base TEC (kWh/yr)			Adder	
Conventional Desktops	DT0	52	38	G1	21	17
	DT I1	84	62	G2	25	20
	DT I2	90	66	G3	32	26
	DT I3	101	74	G4	40	32
	DT D1	86	63	G5	48	38
	DT D2	101	74	G6	51	41
	DT D3	180	142	G7	57	46
Integrated Desktops	DT0	38	26	Same as Above		
	DT I1	62	41			
	DT I2	66	44			
	DT I3	74	50			
	DT D1	63	43			
	DT D2	74	50			
Notebooks	NB0	11	9	G1	9	6
	NB I1	17	14	G2	10	7
	NB I2	18	16	G3	13	8
	NB I3	21	18	G4	16	10
	NB D1	12	10	G5	20	13
	NB D2	14	12	G6	21	13
				G7	23	15

(1) The graphics adders are for DT 0, DT D1, DT D2, DT D3 and the categories do not correlate with the rows of the performance category in this table.

With observed market improvements, these revised levels would roughly maintain or increase today’s share of products that meet the standard relative to our 2013 proposal (see Table 6.2), and also do not account for future improvements.As discussed further in Appendix C, current real-time power management in mobile computing device could be possible for all non-mobile devices if the applicable market barriers were removed. The implementation of these efficiency standards could therefore encourage the removal of these barriers, and make the already feasible and cost-effective standards easier to achieve and more cost-effective.

Table 6.2 Estimated Market Saturation of Products Meeting the Proposed Standard

	2012 - Market Saturation of Products Meeting the Proposed Standard (submitted in 2013)		2014 - Market Saturation of Products Meeting the Proposed Standard (submitted in 2014)	
	Tier 1	Tier 2	Tier 1	Tier 2
Desktops - Conventional	25%	13%	37%	16%
Desktops - Integrated	30%	10%	44%	23%
Notebooks	30%	11%	41%	33%

7 References

[BEA] Bureau of Economic Analysis. 2012. Widespread Economic Growth Across States in 2011. http://www.bea.gov/newsreleases/regional/gdp_state/2012/pdf/gsp0612.pdf.

[CA IOUs] California Investor-Owned Utilities. 2013a. Response to CEC Invitation to Submit Proposals. Codes and Standards Enhancement (CASE) Initiative: Title 20 Standards Development - Computers. CEC Docket #12-AAER-2A. Prepared by Energy Solutions and Ecova.

__.California Investor-Owned Utilities. 2013b. Response to CEC Invitation to Participate. Codes and Standards Enhancement (CASE) Initiative For PY 2013: Title 20 Standards Development - Computers. CEC Docket #12-AAER-2A. Prepared by Energy Solutions and Ecova.

__.2014a. Response to CEC Invitation to Participate. Codes and Standards Enhancement (CASE) Initiative: Title 20 Standards Development – Computers Technical Report. CEC Docket #12-AAER-2A. Prepared by Ecova.

Chetty, Marshini, A.J. Bernheim Brush, Brian Meyers, and Paul Johns. 2009. "It's Not Easy Being Green: Understanding Home Computer Power Management." Paper presented at CHI 2009 Conference in Boston, MA, 4-9 April. http://research.microsoft.com/pubs/79362/chihomepowermanagementchetty.pdf

[DOE] Department of Energy. 2012. Market Assessment and Markup Spreadsheet. Microsoft Excel. http://www.regulations.gov/#!documentDetail;D=EERE-2008-BT-STD-0005-0082.

ECMA International. 2010. Standard ECMA-383. 3rd Edition. Measuring the Energy Consumption of Personal Computing Products. http://www.ecma-international.org/publications/files/ECMA-ST/ECMA-383.pdf

[EPA] U.S. Environmental Protection Agency. 2012. ENERGYSTAR 5.2 Qualified Products List. Accessed on December, 2012. http://www.energystar.gov/productfinder/product/certified-computers/results

__. U.S. Environmental Protection Agency. 2014. ENERGYSTAR 6.0 Qualified Products List. Accessed on October 10, 2014. http://www.energystar.gov/productfinder/product/certified-computers/results

__. U.S. Environmental Protection Agency. 2014a. ENERGYSTAR 6.0 Qualified Products List. Accessed on August 26, 2014. http://www.energystar.gov/productfinder/product/certified-computers/results

Greenblatt, J.; S. Pratt, H. Willem, E. Claybaugh, L. Desroches, B. Beraki, M. Nagaraju, S. Price and S. Young. 2013. Final Field data collection of miscellaneous electrical loads in Northern California: Initial results. Lawrence Berkeley National Laboratory. February.

Fraunhofer. 2011. Energy Consumption of Consumer Electronics in U.S. Homes in 2010. Final report to the Consumer Electronics Association (CEA). December 2011. By Bryan Urban, Verena Tiefenbeck and Kurt Roth.

Fraunhofer. 2014. Energy Consumption of Consumer Electronics in U.S. Homes in 2013. Final report to the Consumer Electronics Association (CEA).June 2014. By Bryan Urban, Victoria Shmakova, Brian Lim, and Kurt Roth.

Hamm, Steve and Jay Greene, 2008. That Computer Is So You. January 2, 2008 http://www.businessweek.com/magazine/content/08_02/b4066000313325.htm. Accessed May 2nd, 2013.

IDC. 2012. Press Release. IDC Worldwide Quarter PC Tracker, February 2012. http://www.idc.com/getdoc.jsp?containerId=prUS23371512. Accessed May 5th, 2013.

IDC. 2013a. Press Release. IDC Worldwide Quarter PC Tracker: United States PC Shipments, January 2013. http://www.idc.com/getdoc.jsp?containerId=prUS23903013. Accessed May 5th, 2013

IDC. 2013b.Press Release. IDC Worldwide Quarter PC Tracker, February 2013. http://www.idc.com/getdoc.jsp?containerId=prUS23987313. Accessed May 5th, 2013.

IDC. 2013c. IDC Worldwide Quarter PC Tracker, June 2013. http://www.idc.com/getdoc.jsp?containerId=prUS24180013. Accessed on July 9, 2013.

IDC. 2014.Press Release. IDC Worldwide Quarter PC Tracker, February 2014. http://www.idc.com/getdoc.jsp?containerId=prUS24700314. Accessed on September 20, 2014.

Microsoft. 2008. Microsoft Power Transition Report; Microsoft Windows.

[NRDC] Natural Resources Defense Council. 2013a. Invitation to Participate Response.

Pigg & Bensch. 2010. "Electricity Savings Opportunities for Home Electronic and Other Plug-In Devices in Minnesota Homes: a technical and behavioral field assessment."
Prepared by the Energy Center of Wisconsin.

[PG&E] Pacific Gas & Electric Company. 2010. "Thin Client Investigation including PC and Imaging State Data." DRAFT. Emerging Technologies Program. Prepared by QDI Strategies. June.

[PG&E] Pacific Gas and Electric Company. 2012a. Cost-Effective Computer Efficiency. Prepared by Ecova.http://www.etcc-ca.com/reports/cost-effective-computer-efficiency.

[PG&E] Pacific Gas and Electric Company. 2012b. Updated Discrete Graphics Processing Unit Testing Results. Prepared by Ecova.

http://energystar.gov/products/specs/sites/products/files/California%20IOU%20%20NEEA%20Comments_Public.pdf.

TIAX. 2007. Energy Consumption by Consumer Electronics in U.S. Residences - Prepared by K. Roth, K. McKenney. Final Report to the Consumer Electronics Association. December. https://www.ce.org/CorporateSite/media/Government-Media/Green/Energy-Consumption-by-Consumer-Electronics-in-U-S-Residences.pdf

Appendix A Duty Cycle Studies

Segment	Study Author	Year	Methodology	Sample Size	Active/ Idle	Sleep	Off
Desktop - Commercial	Barr et al., QDI	2010	Automated tracking and collection	91,000	85%	5%	10%
	Ecma-383, 3rd Edition, Annex B	2010	Automated tracking and collection	250	50%	5%	45%
Desktop - Residential	Microsoft customer experience	2008	Automated tracking and collection	37,400	41%	5%	54%
	Greenblatt et al	2013	Automated tracking and collection (1)	39	32%	48%	20%
	Pigg & Bensch	2010	Automated tracking and collection (2)	42	49%	26%	26%
	Fraunhofer	2014	Phone survey	1,060	32%	24%	44%
	Fraunhofer	2011	Phone survey	1,000	39%	25%	36%
	Chetty et al.	2009	Logging, surveys, interviews	24	75%	12%	12%
Notebook-Commercial	Barr et al., QDI	2010	Automated tracking and collection	19,000	63%	15%	22%
	Ecma-383, 3rd Edition, Annex B	2010	Automated tracking and collection	250	40%	35%	25%
Notebook - Residential	Microsoft customer experience	2008	Automated tracking and collection	35,200	27%	15%	58%
	Greenblatt et al.	2013	Automated tracking and collection (3)	11	16%	25%	42%
	Pigg & Bensch	2010	Automated tracking and collection (4)	12	29%	42%	42%
	Fraunhofer	2014	Phone survey (5)	1,060	17%	42%	42%
	Fraunhofer	2011	Phone survey	1,000	33%	13%	48%
	Chetty et al.	2009	Logging, surveys, interviews	35	36%	32%	32%

Notes:(1) and (3) Low-power mode was cut-off of 10 W, so "sleep" is not technically sleep. This study may therefore understate time in active-idle mode.
(3) and (4) Low-power mode cut-off of 15 W, so "sleep" is not technically sleep. This study may therefore understate time in active-idle mode.
(5) estimates of 14% active mode, 5% short idle, 12% long idle

Appendix B Real-World Adjustment Factor Testing Results

Determining a Real World Adjustment Factor for Computer Energy Use: Laboratory Testing the Impact of Real-World Idle, Active Mode and Peripherals

Prepared by:
Peter May-Ostendorp (Xergy Consulting)
Joe Colett (Ecova)
Catherine Mercier (Ecova)

1 Introduction

The 2013 IOU CASE report and standard proposal includes the use of real-world adjustment factors (15% for desktops and 30% for notebooks) to account for the disparities between observed field usage of computers and the total energy consumption (TEC) approach provided in the ENERGY STAR v6.0 specification, which only accounts for idle (short and long), sleep and off mode power without peripherals. The idle conditions used in the ENERGY STAR v6 test procedure have no windows opened, no applications installed or loaded, and therefore are not representative of typical computer usage. The proposed adjustment factors account for energy consumed during typical idle situations and typical low-intensity active tasks that commonly occur in real-world usage, and include energy used by system drivers and components for common external accessories such as docking stations and secondary displays.

From August through October of 2014, the IOU technical team conducted a project to document the difference, between measured long/short idle mode energy use as measured by ENERGY STAR v6.0 and real-world computer energy use in idle and low-intensity active mode and with peripherals in order to validate existing real-world assumptions from the CASE report.

2 Establishing Laboratory Test Scenarios

In the first phase of this work, the IOU technical teamcontacted computer industry experts — such as participants in ECMA 383 and IEC 62623 — and examined existing literature on computer duty cycles[1] to determine representative idle conditions (referred to as "real-world idle"), active mode tasks and peripherals, estimate their prevalence in the marketplace, and establish approximate duty cycles. We prioritized tasks and peripherals based on a combination of their potential energy impacts, relative market prevalence, and duty cycle. Our estimates of market prevalence and duty cycle for individual tasks and peripherals are a combination of measurements/estimates from previous studies and common sense assumptions. We conducted preliminary testing on these most important tasks and peripherals to measure energy impacts (specifically, increases to short and long idle). Next, we combined the power draw with estimated prevalence and duty cycles to estimate the incremental TEC. Based on our rankings, we prioritized a total of 3 real-world idle scenarios, 5 active task scenarios, and two different peripheral configurations for testing. These scenarios and other basic objectives of the project were shared with the CEC and computer industry stakeholders, and we incorporated feedback leading up to the testing phase.

Tables 2.1 and 2.2 below indicate the peripherals, applications, and tasks associated with each usage scenario. The numbers shown in the tables indicate the number of computers tested; letters indicate the ID of the specific system tested (see Table 3.1 for a complete listing of computers, their hardware attributes, and associated IDs). Usage profiles may vary between home and business users so we segmented our task and peripheral choices accordingly. Tables are also segmented into different peripheral-task combinations, with peripheral scenarios shown in the column headings and tasks indicated as rows. For a home user, the low peripheral case would indicate a casual computer user, whereas high peripheral use may indicate a "power user" with an extensive home office or entertainment setup. Similarly, the peripheral scenarios for the business cases indicate different intensity of peripheral use. The impact of peripherals was assumed to be independent of the task, so the tables reflect testing of the low peripheral cases with only one task for each computer. Though not all peripheral, task, and computer combinations could be explored under the scope of this project, we feel that these scenarios are representative and still capture a diverse range of possible usage.

---

[1] Beauvisage, Thomas. "Computer usage in daily life." *Proceedings of the 27th international conference on Human factors in computing systems Boston*, MA, USA, 2009, pp. 575-584.

Table 2.1: Home Scenarios Identified for Laboratory Phase

	Peripheral Combinations			
	Desktops		Notebooks	
Tasks and Applications	**Low usage** Wired keyboard and mouse, non-USB wired speakers, Ethernet connected	**High usage** USB speakers, wired multi-function printer, wireless mouse, second monitor, wired headset, gaming keyboard, Ethernet connected	**Low usage** non-USB wired speakers, Wifi connected	**High usage** External wireless mouse, wireless keyboard, USB speakers, wired multi-function printer, external monitor (built-in screen still active), wireless headset, Wifi connected
1a) Real-world idle: 3 Chrome tabs (Google, GMail, NY times op-ed), 3 Word windows (3 generic Word templates), 1 PDF document open (Copy of Real World memo), 1 Windows Explorer window C: drive), 1 Windows Media Player window, security/anti-virus suite running (but not scanning) in background (AVG)	2 (A,D)	2 (A,D)	2 (F,H)	2 (F,H)
1b) Real-world idle 1a scenario plus web content: 3 additional Chrome tabs with LA Times homepage, Amazon.com homepage, and Facebook newsfeed (user logged in)	2 (A,D)	2 (A,D)	2 (F,H)	2 (F,H)
1c) Real-world idle1b scenario plus Flash-based web content: Pandora music streaming service loaded, station paused.	2 (A,D)	2 (A,D)	2 (F,H)	2 (F,H)
2) Streaming audio: Spotify streaming on a 1,000 Mbps network connection, using popular music track	2 (A,D)	1 (A)	2 (F,H)	1 (F)
3) Streaming video: Netflix running in full-screen mode on a 1,000 Mbps network connection, using episode of popular TV show	2 (A,D)	1 (A)	2 (F,H)	1 (F)
4) Video chat: Skype running in full-screen mode on a 1,000 Mbps network connection. External webcam will be used on desktops when built-in camera not available. Webcams on both clients will be trained on static images.	2 (A,D)	1 (A)	2 (F,H)	1 (F)
5) Virus/security scan: AVG anti-virus software.	2 (A,D)	1 (A)	2 (F,H)	1 (F)
6) Gaming: Active gameplay using Angry Birds for integrated GPU systems (A and F), Team Fortress 2 for discrete GPU systems (D and H)	2 (A,D)	1 (A)	2 (F,H)	1 (F)
7) Virtualization: Windows 7 VM on Mac using VMWare Fusion.	1 (D)	0	1 (F)	0

*Tasks 2 - 7 are performed running the applications specified in task 1a, real-world idle.

Note: All network interfaces, including wireless and Bluetooth, are powered on. In addition, the ENERGY STAR v6.0 test procedure requires that tests are conducted with a minimum of a mouse and keyboard connected (desktops only) and with an active Ethernet network connection.

Table 2.2: Business Scenarios Identified for Laboratory Phase

	Peripheral Combinations			
	Desktops		Notebooks	
	Low usage	**High usage**	**Low usage**	**High usage**
Tasks and Applications	Wired keyboard and mouse, non-USB wired speakers, Ethernet connected	USB speakers, wireless mouse, second monitor, wired headset, wireless keyboard, Ethernet connected	non-USB wired speakers, wired keyboard, wired mouse, Wifi connected	Docking station with external wireless mouse, wireless keyboard, USB speakers, external monitor (built-in screen still active), wireless headset, Wifi connected
1a) Real-world idle: 3 IE tabs (Google, GMail, NY times op-ed), 3 Word windows (3 generic Word templates), 2 Excel templates, 2 Outlook, 2 PDF document open (Copy of Real World memo), 1 PPT, 1 Windows Explorer window C: drive), security/anti-virus suite running (but not scanning) in background (AVG, ClamWin)	3 (B,C,E)	3 (B,C,E)	1 (G)	1 (G)
1b) Real-world idle scenario 1a plus three IE tabs: LA Times homepage, Amazon.com homepage, Facebook newsfeed (user logged in)	3 (B,C,E)	3 (B,C,E)	1 (G)	1 (G)
1c) Real-world idle scenario 1b plus Flash content: Pandora music streaming service open to station, paused.	3 (B,C,E)	3 (B,C,E)	1 (G)	1 (G)
2) Streaming audio: Spotify streaming on a 1,000 Mbps network connection, using popular music track	3 (B,C,E)	0	1 (G)	0
3) Streaming video: YouTube running in full-screen mode on a 1,000 Mbps network connection, playing standard clip	3 (B,C,E)	0	1 (G)	0
4) Video chat: GoToMeeting with web cameras trained on static images	3 (B,C,E)	0	1 (G)	0
5) Virus/security scan: AVG anti-virus software	3 (B,C,E)	0	1 (G)	0
6) Screen sharing: GoToMeeting. Unit under test shares screen with another client.	3 (B,C,E)	0	1 (G)	0

*Tasks 2 - 6 are performed running the applications specified in task 1a, real-world idle.

Note: All network interfaces, including wireless and Bluetooth, are powered on. In addition, the ENERGY STAR v6.0 test procedure requires that tests are conducted with a minimum of a mouse and keyboard connected (desktops only) and with an active Ethernet network connection.

3 Computer Selection

We selected eight computers on which to test the above scenarios. The table below shows the systems tested: three conventional desktops, two integrated desktops, and three notebooks. We have indicated the ENERGY STAR performance category, market sector, and basic hardware configuration for each system.

Table 3.1: Computer Systems

System ID	ENERGY STAR v6 Category	Sector	System Description
Desktops			
A	DT 0	Home	Dell Inspiron 660s (2013): Intel Celeron G465 1.9 GHz, 2GB memory, Windows 7
B	DT I2	Business	Dell OptiPlex (2014): Intel Core i3-4150 3.5 GHz, 4 GB memory, Windows 7
C	DT I3	Business	Lenovo ThinkCentre (2013): AMD A6-5400B 3.8 GHz, 4 GB memory, Windows 7
Integrated Desktops			
D	DT D2	Home	Apple iMac (2013): Intel Core i5-4670 2.9 GHz, NVIDIA GeForce GT 750M graphics, 8 GB memory, 1 TB hybrid hard drive, OS X 10.9 (Mavericks)
E	DT I3	Business	Dell XPS 27 Touch (2014): Intel Core i5-4440S 3.3 GHz, 8 GB memory, Windows 8
Notebooks			
F	NB I1	Home	Apple MacBook Pro (2014): Intel Core i5-4288U 2.4 GHz, 4 GB memory, OS X 10.9 (Mavericks)
G	NB I1	Business	Lenovo ThinkPad X1 (2014): Intel Core i5-4200U 1.6 GHz, 4 GB memory, Windows 7
H	NB D2	Home	Razer Blade (2013): Intel Core i7-4702HQ 2.2 GHz, NVIDIA GeForce GTX 765M, Windows 8

4 Laboratory Measurements

The laboratory tests commenced first by conducting a standard ENERGY STAR v6.0 test to each system, measuring power consumption in off, sleep, long, and short idle modes, providing a baseline power consumption values in different modes against which future tests could be compared. Systems were also measured under three real-world idle conditions, listed as 1a, b, and c in Table 2.1 and Table 2.2. Measurements of real-world idle effectively followed the same ENERGY STAR v6.0 procedure, except that additional peripherals and idle applications were also present.

Given the exploratory nature of this research, the various active mode measurements were conducted by modifying existing procedures. The IEC 62623 test standard (on which ENERGY STAR's test method is largely based) does contain a section related to the measurement of active workloads (5.6.3). We adapted this methodology by recording average active power over an interval of 15 minutes, ensuring that we captured at least 5 minutes of usable data.[2]

5 Formulation of Real-World Adjustment Factors

The real-world adjustment factors or TEC errors calculated for this project are calculated as follows:

$$\mathrm{TEC_{err}} = \frac{\mathrm{TEC_{new}} - \mathrm{TEC_{old}}}{\mathrm{TEC_{old}}}$$

Eq. 1

where $\mathrm{TEC_{new}}$ is the real-world TEC of the test system and $\mathrm{TEC_{old}}$ is the TEC as computed by the IOU team's recommended duty cycles (presented in Table 3.1). The ENERGY STAR TEC is the sum of energy consumption in each measured mode of operation, each of which is a product of power, p, and time, t:

$$\boldsymbol{\mathrm{TEC_{old}} = p_{\mathrm{off}} t_{\mathrm{off}} + p_{\mathrm{sleep}} t_{\mathrm{sleep}} + p_{\mathrm{idle,short}} t_{\mathrm{idle,short}} + p_{\mathrm{idle,long}} t_{\mathrm{idle,long}}}$$

Eq. 2

To calculate the revised TEC, including the impacts of real-world short and long idle as well as the impact of peripherals and active tasks, we have included several additional terms:

$$\mathrm{TEC_{new}} = p_{\mathrm{off}} t_{\mathrm{off}} + p_{\mathrm{sleep}} t_{\mathrm{sleep}} + p_{\mathrm{RWI,short}} t_{\mathrm{RWI,short}} + p_{\mathrm{RWI,long}} t_{\mathrm{RWI,long}} + p_{\mathrm{active}} t_{\mathrm{active}} + p_{\mathrm{periph}} t_{\mathrm{periph}}$$

Eq. 3

where the RWI subscript represents real-world idle, followed by active task energy use, and peripheral energy use.

The power for real-world long/short idle, active mode, and peripherals were measured as described, but a modified duty cycle was created to account for the time allocated to these modes. We used a 2013 study of residential digital media consumption[3] to infer the absolute time spent in homes conducting active tasks (video streaming, web browsing, video chatting, etc.). We conservatively estimate that homes use their computers for about 850 hours per year for these active tasks, which we have called "Active-High" tasks. We estimate that an additional 526 hours per year are spent in "Active-Low" tasks, such as e-mail and productivity applications, but we have treated these hours as equivalent to short idle mode in this study.

---

[2] Background tasks can and do occur during active (and sometimes even idle) power measurements, so testers and analysts must take care to exclude spurious power spikes from measurements. We took great care to identify contiguous, 5-minute periods of "clean" data for averaging.

[3] Short, James. 2013. How Much Media? Report on American Consumers. USC Marshall School of Business; Institute of Communication Technology Management. http://www.marshall.usc.edu/faculty/centers/ctm/research/how-much-media

In the commercial sector, we utilized assumptions from a Dell client computer energy savings calculator and associated white paper[4] to estimate active mode operation in office settings. Dell's calculator assumes that client computers are utilized for so-called "high performance" tasks – what we have called Active-High – about one hour per day or 250 hours per year. This is also a fairly conservative estimate, as it does not include productivity tasks – again what we would call Active-Low tasks – which Dell assumes constitute another 7 hours per day of activity.

Finally, we assumed that active mode operation displaces time that would otherwise be spent in short idle mode. When combining residential and commercial duty cycles using a sector weighting — approximately 60% commercial and 40% residential based on sales and stock estimates — we arrived at the following real-world duty cycles for desktops (standard and integrated) and notebooks. The "baseline" duty cycles on the left-hand sides of the charts represent the duty cycle recommendations presented by the IOU team in Table 3.1 in the main body, although we also have the ability to generate results using the ENERGY STAR Version 6.0 duty cycle as a baseline.

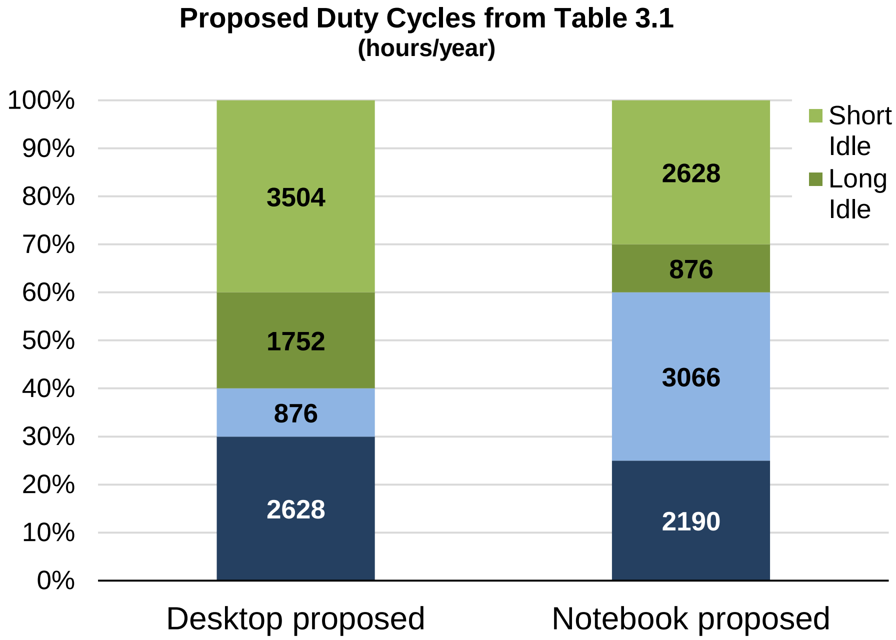

Figure 1: Proposed desktop and notebook duty cycles

---

[4] Dell. 2014. Client Energy Savings Calculator. http://www.dell.com/content/topics/topic.aspx/global/products/landing/en/client-energy-calculator?c=us&l=en

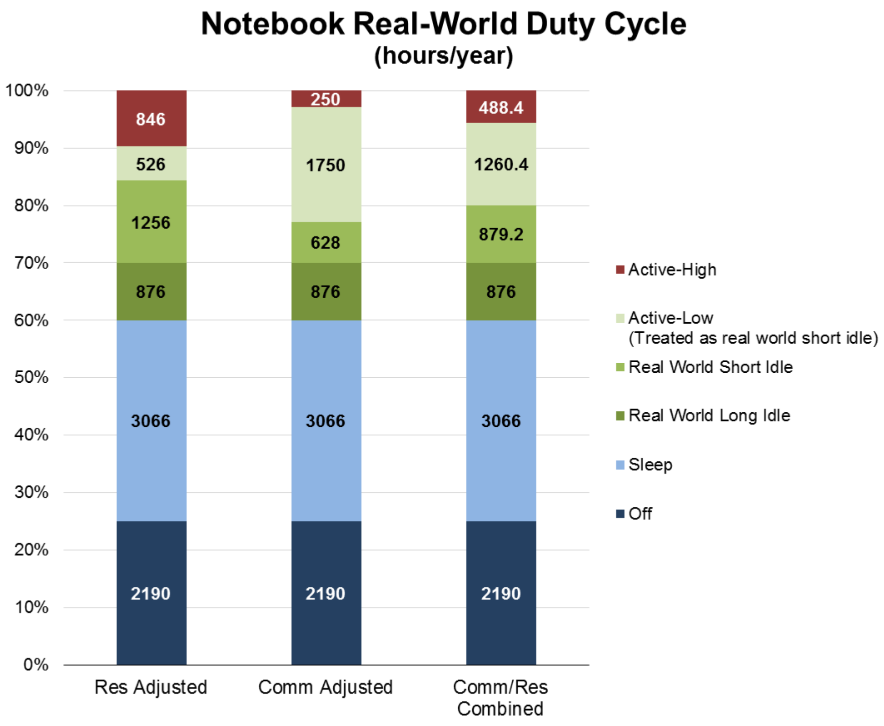

(a)

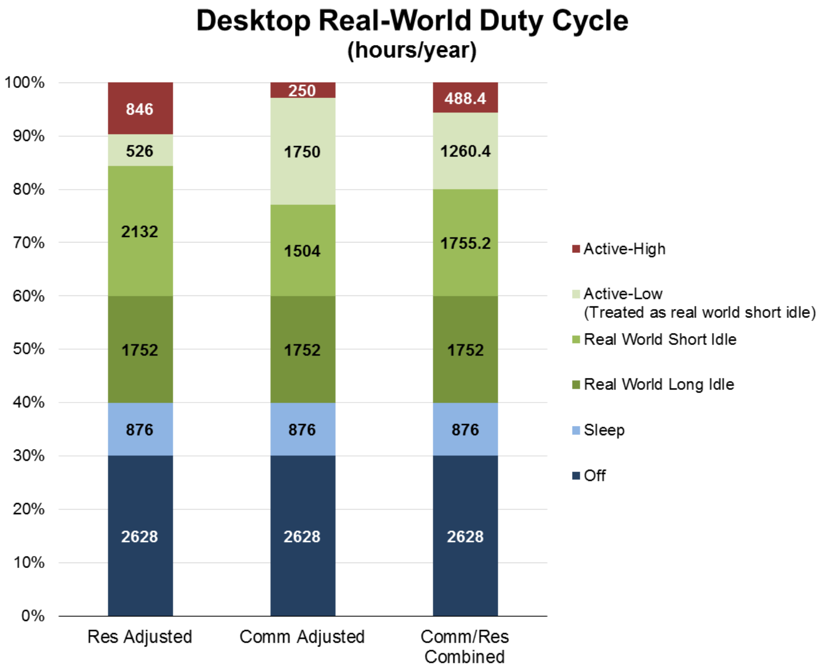

(b)

Figure 2: Desktop (a) and notebook (b) duty cycles adjusted for real-world usage

For peripherals, we assume that any incremental power incurred by peripherals only applies to idle and active modes of operation, so the t_{periph} term above includes long and short real-world idle time as well as active time (4,380 hours per year).

To compute TEC adjustment factors for each system, established typical real-world power consumption values for long idle, short idle, and active-high modes, then provided a peripheral adder. Real-world long and short idle values represent the average power measured from the three real-world idle scenarios (1a-c). Active power values are the average of all active scenarios for a given system (1-6 for most systems, with the addition of Task 7 for Apple systems with virtualization).

The peripheral adder for each system was estimated based on differences in power consumption between ENERGY STAR baseline measurements and our low and high peripheral configurations. We found that the low peripheral configuration generally did not differ significantly from ENERGY STAR test conditions, so the high peripheral configuration is effectively the only result of interest. In developing our peripheral power adder, we assume that only 25% of computer users have such a configuration, so the high peripheral configuration was only given a 25% weighting.

6 Results

Figure 3 provides a comparison of the TEC as computed using standard ENERGY STAR v6.0 measurements and the IOUs' recommended duty cycles versus the estimated real-world TEC based on laboratory measurements. TEC associated with long idle increases across all products due to real-world idle impacts. Although short idle power increases under real-world conditions, sometimes quite dramatically, the overall share of TEC in short idle decreases because we now assume that some of that time is spent in Active. Peripherals generally contribute very little to the TEC of the system, about 3.7 kWh per year on average. Active mode, however, can be another substantial portion of the TEC, particularly with today's highly power-scalable processors.

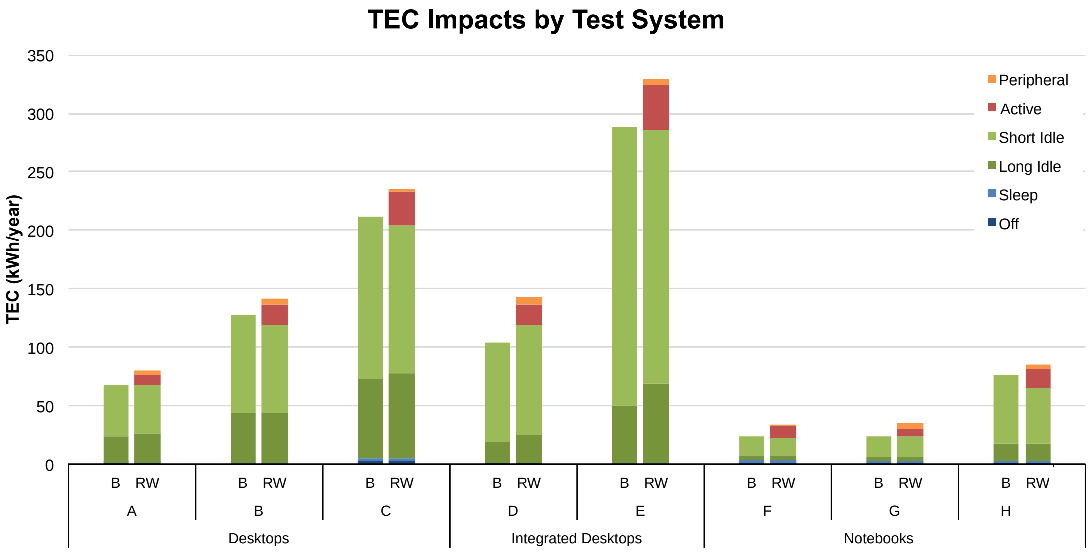

Figure 3: Baseline (B) and real-world (RW) TEC impacts by test system

After averaging results by product form factor, we have illustrated the overall estimated TEC adjustment factors in Figure 4. We have found that the factors used in the 2013 CASE report generally agree well with measurements conducted across the 8 computer systems, 9 tasks, and 2 peripheral scenarios examined in this work. The adjustment factors for integrated desktops and notebooks may, in fact, have been somewhat more conservative than the data indicates.

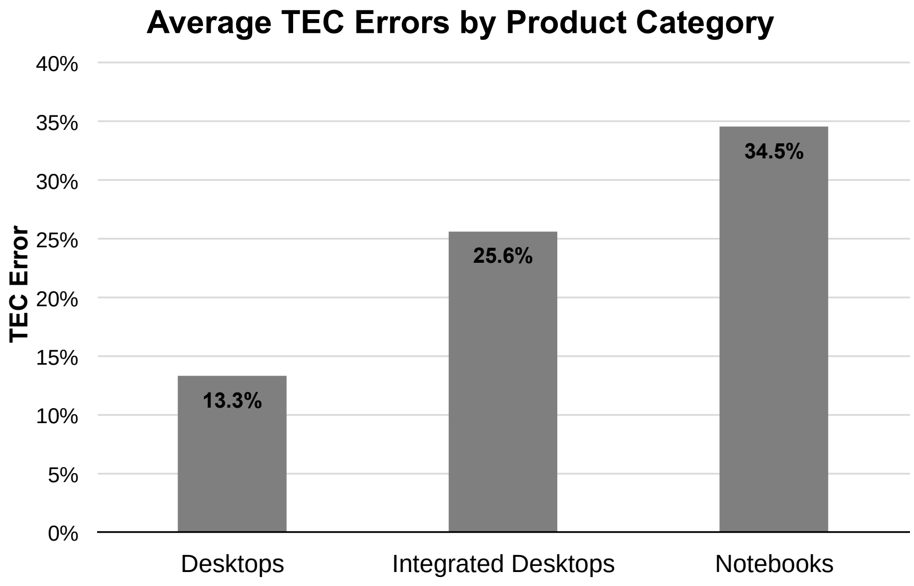

Figure 4: Real-world TEC adjustment factors

7 Applicability of Real-World Adjustment Factors (RWAF)

The real-world adjustment factors developed in the IOU team's research compare measured TEC using the standard test procedure to the measured TEC under real-world usage conditions. Our results show that most systems will experience TECs anywhere from 13% to 35% higher than those estimated using standard test data. This section addresses the question of whether this adjustment factor can be applied not just to energy consumption but to energy savings as well. The following two arguments make the case that most of the RWAF can be applied to energy savings.

The underlying hardware components of the system drive this behavior by providing services directly to the user (i.e. tasks that the user has initiated during active operation) or by facilitating background tasks even when no service is provided (i.e. real-world idle tasks like animated web content that may persist even when the user is unable to view or use them). The components whose power is most increased by real-world conditions include:

- Processors that power scale — by increasing clock speeds or by other means — according to the computational load placed on the system
- GPUs that may be required to animate content even if users are unable to view it; and
- Power supplies that must provide DC power for all of these underlying components.

In the 2013 CASE report, the IOUs presented various cost-effective technology pathways that allow desktop computers to comply with the proposed standard. Approximately 83% of the energy savings opportunity for Tier 1 and 75% of the opportunity in Tier 2 were comprised of modifications to the above components. In other words, 75% to 83% of the technological opportunities identified in test systems were associated with components that are impacted by real-world operating conditions.[1] Consequently, we have applied the real-world adjustment factors so that they only apply to 75% of the estimated savings. For example, in a typical desktop we would expect that real-world energy savings is about 13% higher than those we predicted based on bench measurements. However, the adjustment only applies to the 75% of energy savings that we estimate can be gotten with power-scalable components that are sensitive to real-world effects.

Another way to estimate the applicability of real-world adjustment factor to energy savings is to consider the effects of real-time power management. Techniques such as "keystroke sleep" and not rendering graphics that are hidden by other windows offer a similar savings potential in low-intensity tasks such as reading web content, writing email or using office productivity software, where little processing is required, as they do in idle mode. Real-time power management allows the system to ramp power down during most of the duration of these tasks. Even video streaming, an ongoing low-to-medium intensity task, can be performed for as little as 2 to 3 watts by media players such as Apple TV, Roku, Google Chromecast etc. The only scenarios where real-time power management would provide significantly reduced benefits are high-intensity tasks such as gaming

---

[1] The remainder of savings involved upgrades to hard drives, which have a more constant power profile and are not as dramatically impacted by real-world computational loads.

Appendix C Notebook Demonstration Project

Notebook Demonstration Project

Phase 3 - Notebook Power Management Analysis and Demonstration
Part I - Opening up the Power Management

Version 1.0
10/27/2014

Authors:
Davorin Mista
Vojin Zivojnovic
AGGIOS, Inc.

.

Executive Summary

This report provides a qualitative analysis of the energy consumption of desktop and notebook computers in the long and short idle modes based on a similar analysis conducted for IPTV set top box devices[5]. The goal is to suggest hardware, software and design methodology improvements for the reduction of computer energy consumption by comparing computers and mobile devices in the equivalent and similar idle modes. The proposed solutions are guided by the principle of energy proportionality which requires that the work done by the software and the energy consumed by the hardware are measured and kept in proportions. We are advocating for improvements in the way software is reporting the work it is conducting, the introduction of energy-driven hardware hierarchies, more openness for device integrators to control the power management of the hardware and an industry standard-driven methodology to align software and hardware. Our estimate is that these efforts have the potential to reduce the energy consumption of computers by around 25% in short idle and 50% in long idle. The resulting total savings for the US electricity ratepayers are estimated at $510M per year at the total incremental costs of $54M per year. In order to provide a detailed quantitative analysis of the proposed improvements and define the next steps to move toward the solution, we recommend the initiation of a new project developing a fully functional computer reference design implementing the suggested improvements and conducting detailed measurements, testing and an in-depth costs/benefits analysis. We estimate that such a project would take between 12 and 18 months to complete.

---

[5] Final report of CEC EISG project #57751A/13-14: Unified Power Control for Set Top Box Devices (AGGIOS)

Introduction

The report describes the forward-looking analysis and ideas for the improvement of the energy consumption of personal computers (desktops and notebooks). It is similar in objectives and methods to the IPTV set top box project conducted previously under the CEC EISG grant[6], but it is based on qualitative assessments and not on implemented and tested improvements on an actual device. The first two phases of the notebook demonstration project have led the research team to discover that the analyzed notebooks are built without much openness and adaptability in the areas of power and energy management. Despite a variety of approaches, the attempts to control the power states from the application, OS or firmware software side stayed unsuccessful. This report documents the steps taken and discusses the need and justification for a more open and easier to implement power management in personal computers, like desktops or notebooks.

---

[6] Final report of CEC EISG project #57751A/13-14: Unified Power Control for Set Top Box Devices (AGGIOS)

ACPI and TDP Overview

In January 1992, Intel and Microsoft developed APM (Advanced Power Management) to manage power when a computer system was in idle. At the time, this new technology was driven by the increasing popularity of battery powered notebooks. In December 1996, the successor of APM - the Advanced Configuration and Power Interface (ACPI) specification[7] was developed by Compaq, Microsoft, Intel, Phoenix and Toshiba as the industry open-standard power management interface. The 940 pages long ACPI specification was the consequence of not only the tight cooperation, but also dominance of Microsoft and Intel who controlled the vast majority of personal computers produced. The goal was to remove the third-party BIOS software from the power management activity (and other run-time activities) and establish the Windows OS and its infrastructure as the single system software controlling the run-time of ACPI based computers. The only remaining role of the BIOS software in ACPI based computers became to carry out the boot procedure and then hand over full control to the Windows OS. The side by side comparison of APM and ACPI is shown on Figure 1.

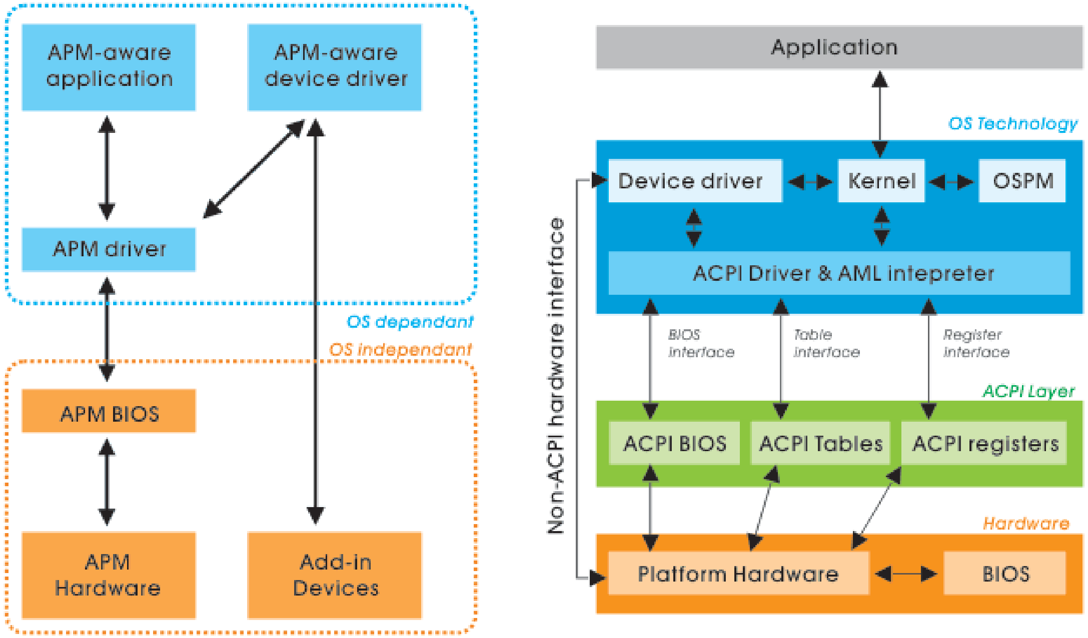

Figure 1 - Comparison of APM (left) and ACPI (right)[8]

The classical BIOS from the APM system was replaced by the OS dependent ACPI layer which was taking over the power management tasks in coordination with the ACPI driver of the OS. As the APM BIOS for each product design has been developed through close collaboration between the BIOS companies and the OEMs, ACPI had to provide a replacement mechanism for the OEMs to manage the power of the non-standard devices they regularly add to the motherboard and the personal computer, as hard disk drives (HDD).

---

[7] http://www.uefi.org/sites/default/files/resources/ACPI_5_1release.pdf

[8] http://www.techarp.com/showarticle.aspx?artno=420&pgno=0

The ACPI specification[9] states: "*Although the classical fixed hardware programming model requires hardware registers to be defined at specific address locations, the generic hardware programming model used in ACPI allows hardware registers to reside in most address spaces and provides system original equipment manufacturers (OEMs) with a wide degree of flexibility in the implementation of specific functions in hardware. The OS based Power Management (OSPM) directly accesses the fixed hardware registers, but relies on OEM-provided ACPI Machine Language (AML) code to access generic hardware registers. AML code allows the OEM to provide the means for OSPM to control a generic hardware feature's control and event logic."* Figure 2 shows the hardware events and control based on ACPI and AML.

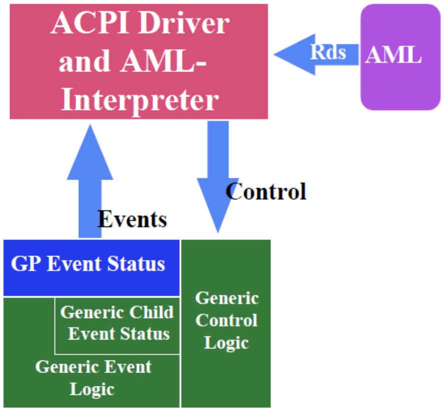

Figure 2 - ACPI and AML for Hardware Control

The specification continues*: "One goal of ACPI is to allow the OEM "value added" hardware to remain basically unchanged in an ACPI configuration. One attribute of value-added hardware is that it is all implemented differently. To enable OSPM to execute properly on different types of value added hardware, ACPI defines higher level "control methods" that it calls to perform an action. The OEM provides AML code, which is associated with control methods, to be executed by OSPM. By providing AML code, generic hardware can take on almost any form."*

The following is an example provided in the ACPI document: *"As an example of a generic hardware control feature, a platform might be designed such that the IDE HDD's D3 state has value-added hardware to remove power from the drive. The IDE drive would then have a reference to the AML PowerResource object (which controls the value added power plane) in its namespace, and associated with that object would be control methods that OSPM invokes to control the D3 state of the drive:*

- *_PS0: A control method to sequence the IDE drive to the D0 state.*
- *_PS3: A control method to sequence the IDE drive to the D3 state.*
- *_PSC: A control method that returns the status of the IDE drive (on or off).*

---

[9] http://www.uefi.org/sites/default/files/resources/ACPI_5_1release.pdf

The control methods under this object provide an abstraction layer between OSPM and the hardware. OSPM understands how to control power planes (turn them on or off or to get their status) through its defined PowerResource object, while the hardware has platform-specific AML code (contained in the appropriate control methods) to perform the desired function. In this example, the platform would describe its hardware to the ACPI OS by writing and placing the AML code to turn the hardware off within the _PS3 control method. This enables the following sequence:

- *When OSPM decides to place the IDE drive in the D3 state, it calls the IDE driver and tells it to place the drive into the D3 state (at which point the driver saves the device's context).*
- *When the IDE driver returns control, OSPM places the drive in the D3 state.*
- *OSPM finds the object associated with the HDD and then finds within that object any AML code associated with the D3 state.*
- *OSPM executes the appropriate _PS3 control method to control the value-added "generic" hardware to place the HDD into an even lower power state."*

The ACPI has faced strong criticism among the OS experts. The creator of Linux, Linus Torvalds, has been a particularly harsh critic of ACPI[10]: His view was shared by other OS experts as well. As a consequence, in 2009 Intel started the new Simple Firmware Interface (SFI)[11], but this project never really took off.

On the other hand, there are some CPU centric technologies for tuning the active mode power of the chipsets through Intel's Dynamic Platform and Thermal Framework[12,13] (DPTF) offering features such as configurable TDP[14] (thermal design power) and LPM (low power mode). These features allow OEMs to customize their designs in regards to thermal and power/performance properties in the active mode, setting thermal limits or limiting the maximum power consumed by the main SoC when the device is in active use. In theory, DPTF can leverage ACPI-based information about devices, thermal zones, thermal sensors, thermal relationships etc., but also can work with information coming directly from drivers, giving the SoC manufacturers the option to hide all of these details in their proprietary device drivers.

---

[10] http://www.linuxjournal.com/article/7279

[11] https://simplefirmware.org/

[12] https://01.org/intel%C2%AE-dynamic-platform-and-thermal-framework-dptf-chromium-os/

[13] V.Srinivasan et al, "An Innovative Approach to Dynamic Platform and Thermal Management for Mobile Platforms

[14] http://en.wikipedia.org/wiki/Thermal_design_power#Configurable_TDP

DELL Notebook Power Management Analysis

In this project the research team explored the level of openness of the ACPI mechanism and level of ease for an OEM or smaller PC manufacturer to tune the power management of the system on chip, motherboard, device and the software in order to reduce the energy consumption of the product. The research team used the Dell Inspiron 5748 notebook to conduct the analysis and draw conclusions, selected based on its high ENERGY STAR reported idle wattage relative to other products.

The hardware features of the notebook's main processing component are as follows:

- Intel i5-4210U chip includes two dies in one package:
    - Haswell CPU (22nm);
    - LynxPoint Chipset (32nm);
- Haswell CPU includes:
    - Dual-core CPU;
    - Graphics and Media accelerators;
    - Fully Integrated Voltage Regulator (FIVR);
        - Two voltage rails are supplied to i5 chip (6+ in IvyBridge):
            - one variable ~1.8V voltage rail supplied to the i5 chip;
            - one variable ~1.3V voltage rail for the memory controller;
    - Power Control Unit (PCU):
        - enables lower power C7 state even while display is on;
- Intel 8 Series PCH (LynxPoint-LP HM86):
    - Included in package of i5 chip (separate die);
    - PCH = Platform Control Hub;
    - PCH replaces Northbridge/Southbridge, features include:
        - PCI-Express Controller;
        - SATA HDD Controller;
        - USB Controller;
        - Gigabit Ethernet Controller;
        - DMA Engine;
        - SMBus/I2C;
        - HD Audio;
        - Serial Ports, Timers, Real Time Clock, GPIO;
    - 4 Voltage Rails: 1.05V, 1.5V, 1.8V, 3.3V;

In the Haswell datasheets, no information about the PCU and FIVR as key power management elements could be found. Specifically, the ACPI machine description tables (SSDT, DSDT) do not expose the critical building blocks needed to analyze or influence power management behavior. There were no power resource objects defined in the description, and the power state change methods (_PS0, _PS2, _PS3) were only defined for two major peripherals, the USB 3.0 and the WiFi, as well as for a handful of the simple low-power peripherals, such as the serial interfaces. But even for those devices, the power control methods appear to be incomplete and were hence not revealing any useful information that would help optimize the power management behavior of a device. Therefore the example given in the ACPI specification v5.1 on page 58 for value-added hardware for placing the device in custom lower power states does not apply for this notebook. Given the fact that the Intel chipset offers many custom lower power states according to Intel's marketing material as well as many technical presentations about Haswell by Intel, we conclude that the ACPI methodology was not followed by Intel in order to control the power states of their devices.

The Haswell architecture does support some Intel specific features for tuning of power and performance, namely configurable thermal design power (cTDP) and low power mode[15] (LPM). These features can only be configured through Intel's proprietary DPTF driver which isn't publicly available for the analyzed Dell laptop.

Our conclusion is that the power management technology resides in the Windows device driver DLLs from Intel and other chip vendors, without exposing power management controls in ACPI, which is the standard supposedly used in Windows for the power management framework to support the OEMs. While ACPI has provisions for declaring all power management details of devices in a description table, the notebook computer we have analyzed contains a very sparse ACPI description table, leaving most power management controls inside the closed source device drivers supplied by the chip vendors.

On the application software side we are missing the ability to see how software components affect the power consumption of the device, especially when it comes to background tasks running while the PC is idle. There's no provision for software designers to declare or announce its resource requirements and it's also very difficult to observe resource usage for such processes, since Windows is running many background tasks as so-called services. These tasks don't get exposed in the system analysis tools provided by Windows or third parties. Instead, up to 16 services get run within a single servicehost process, of which multiple instances are running at any time. It is possible to monitor resource usage for individual processes, but for the servicehost processes there's no way to tell which one of the services is being executed at the time the resource usage is observed.

A second problem on the software side is that Windows doesn't provide any transparency in terms of activity of its background tasks. Services are installed and activated without giving users any choice or information. It's difficult to figure out which application depends on which services, what those services are doing, how often they are being run, and when they are scheduled. The effects can be observed when monitoring resource usage and power consumption of laptop computers in their Short Idle and Long Idle states, which can sometimes be more active than when users are interacting with the device. Whether or not such activity is justified and whether the performed tasks are serving the user's interest can only be determined through expert analysis, while the common user remains completely unaware of these activities and how they impact the power consumption of their device.

---

[15] http://www.intel.com/content/dam/www/public/us/en/documents/datasheets/4th-gen-core-family-mobile-m-h-processor-lines-vol-1-datasheet.pdf - p. 61

The lack of transparency and control on the hardware side as well as on the software side makes it difficult for OEMs to understand the impact of their design decisions on the energy consumption of their device, often leaving many optimization opportunities on the table. Especially newcomers and smaller vendors in this field will have a hard time achieving optimal energy efficiency for their devices, as they often don't have the means to receive the support of the large chip vendors for help in optimizing power and energy efficiency.

Software developers often have no awareness of the energy impacts of their software. Even though some tools, such as Intel's VTune, are available to optimize the energy efficiency of their applications, as of today, little insight is offered for background tasks or services used by their applications[16]. This issue is exacerbated by the fact that software developers feel no pressure from their customers to focus on energy efficiency as users rarely associate poor energy efficiency with individual applications due to the fact that operating systems are lacking the infrastructure to break out energy consumption by software application More transparency and control could lead to more innovation in the field of energy efficiency, both by device manufacturers and software developers.

---

[16] https://software.intel.com/en-us/node/529958

Opening up the power management infrastructure

The results of this project highlights a systematic problem in the computer development chain exists: none of the parties in the design and production chain has the responsibility for the energy management and the comprehensive energy picture of the complete device including the main system on a chip (SoC), memory, components on the device, motherboard and the attached devices. The producers of the notebook SoC provide the reference design and the reference software including the OS drivers for the SoC. They are not responsible for the whole notebook device, as they do not know what kind of components their customers and their OEMs shall use for the subsequent motherboard development. The notebook manufacturers rely on the OS software including the drivers which the SoC provider is delivering. They mostly just take what the SoC manufacturer provides, so can manage the sleep states only through the OS system calls, not the OS or a dedicated energy management kernel what is the prescribed approach heavily used in mobile devices. This leaves the notebook motherboard with a number of power hungry components beyond the SoC which are hard to include in the power management, like the DRAM, flash memory, Ethernet controller, PHY, WiFi, HDD and board level PMIC. Simply put, there is a systematic problem in the design and development of computers for the companies which are not large enough to receive a dedicated support service from the SoC and OS producers.

This problem was supposed to be addressed by the ACPI technology presented in the previous sections. Unfortunately, the ACPI technology has barely changed during the last 18 years despite the rapid advanced in hardware and software, especially in the mobile space. As such, it is less and less likely that it can be of help to the OEMs or new innovative companies who are trying to reach the highest energy efficiency levels for their products. In the past, the BIOS (called also core system software), like Phoenix Technologies, American Megatrends and Insyde Software filled this void between the SoC and OS manufacturers on one side and the OEMs on the other by offering low level software and professional design services, which in particular were in high demand by the OEMs. As the activity of the BIOS was seen as an unnecessary cost factor, most of these companies are now exiting the BIOS business.

In the next sections we shall discuss approaches leading to a more efficient and widely deployable power management.

Energy Proportionality

For successful application and deployment of energy management techniques it is important to carefully analyze the application software activities, identify power state dependencies and objectively determine their realistic energy needs. Very often the lack of alignment between the various software programs and their development teams combined with the insufficient information about the power needs of the software causes all components to be kept at their full power all the time. A good example is the protocol software which provides the security, update, digital right management or communication synchronization between the device and the remote location during the low power modes.

Such software should be executed in an energy proportional manner, which means that there is no reason to wake up the main processing cluster and all the peripherals of the notebook just to send and receive a limited number of internet protocol (IP) packets. In the current notebook designs not only is the energy intensive main processing cluster used for that purpose, but in order to be available at every possible instant for such an infrequent activity the whole device is kept up and prolongs entering into the deeper sleep states. Apple's dark wake concept is a good example of the right approach to power management of such less frequent and background device activities. Energy proportionality[17] should be the governing principle for energy management to guarantee that the work done by the software and the energy consumed by the hardware are all the time kept in the right proportions.

Energy Design Standards

The key for transparency and industry collaboration in power management is the usage of standardized descriptions of the power and energy aspects of the system and the individual components. The industry needs a standard that can capture all of the complexities of today's hardware in terms of power states, dependencies between components, as well as the controls and conditions necessary for entering those power states. At the same time, the standard should also allow to describe the energy needs of the software running on those platforms. Software applications should declare the resource usage of their primary as well as background tasks in a specific manner, including intensity and frequency. Software should also expose controls to change how individual tasks are scheduled and how this affects functionality and quality of service.

Currently, such standards do not exist. The energy design and management flow for electronic devices, including notebooks, is disconnected and lacks the abstractions, formats, interfaces, and automated methodologies long established and standardized in functional design and verification of hardware and software[18]. The main disconnects (or collaboration walls) on energy issues are at the handover points between various design teams shown on **Error! Reference source not found.**Figure 3:

- From hardware designers and component providers to system integrators and power designers at the hardware level;
- From system integrators to firmware developers at the core system software level;
- From device developers and firmware developers to OS integrators at the OS API level;
- From device developers and OS integrators to the application programmers at the application software level;

---

[17] L. Barroso, and U. Hö zle, "The Case for Energy-Proportional Computing," *IEEE Comp. Soc. Press*, vol. 40, issue 12, Dec. 2007.

[18] http://www.eetimes.com/author.asp?section_id=36&doc_id=1320073

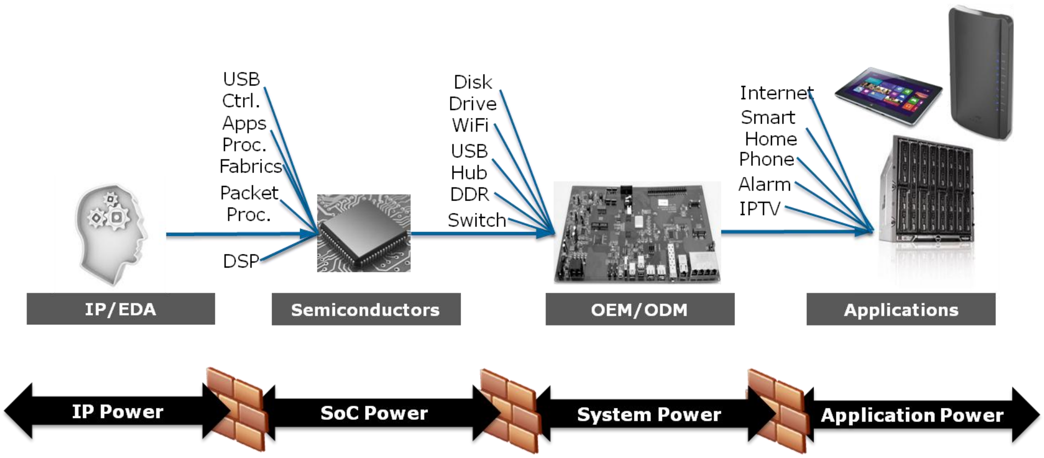

Figure 3: Power and Energy Design Flow

A new standard should address all of these challenges, mainly from an engineering perspective, by automating the design of power management software and hardware. But such an approach has broader benefits, including:

- Lowers the cost and shortens the time-to-market, as companies across equipment classes can share the same power management infrastructure;
- Enables the introduction of formal energy description of equipment and automated generation of the power management software and hardware;
- Fosters energy savings competition among manufacturers of various equipment classes;
- Allows for unified run-time reporting and control of equipment's energy consumption;
- Establishes tighter minimum energy efficiency requirements for higher energy savings and enables faster adoption of energy standards for future equipment classes;
- Allows unified testing, measurement, and standardization procedures;
- Brings together a larger pool of power management experts to focus on the same problem and improves technical education in academia;

In August 2014 the IEEE Standards Association formed a new IEEE standardization project and the IEEE P2415 Working Group (WG) focused on the “Unified Hardware Abstraction and Layer for Energy Proportional Electronic Systems”[19]. The IEEE P2415 WG shall standardize the best energy savings practices from the electronic industry and make them available to the public. The P2415 WG includes a large number of worldwide leading electronic companies who supported the P2415 initiative during the 12-month formation process, including ARM, Atrenta, Broadcom, Cadence, Cisco, Comcast, Doceapower, IBM, Intel, LGE, Mentor, Microsoft, Pace, Qualcomm, Realintent, Synopsys and Xilinx.

---

[19] http://standards.ieee.org/news/2014/ieee_p2415_p2416_wgs.html

Software and Tools for Energy Proportional Design

To fully leverage the energy standard for energy proportional devices, it is advisable to develop software tools acting as the design, visualization, generation and test/measurement cockpit for all energy management and design related activities. The required features of such tools are:

- Create and edit descriptions of the hardware and software details of their systems;
- Export and import a variety of system and power description formats, including ACPI, SFI, DTS, UPF/CPF and IP-XACT;
- Represent the device and component hierarchy in graphical and tabular format;
- Enter and represent the FSMs as state transition tables (STT) or diagrams;
- Verify the validity and consistency of the description;
- Create and manage the IEEE P2415 library of components and devices;
- Generate the energy management driver, user space daemon and kernel for always on cores;
- Generate the FSM for dedicated hardware implementations;
- Simulate power management scenarios on FPGAs or virtual prototypes;

The tool should be integrated with physical power measurement and data acquisition (DAQs) devices for automated measurement, test and verification of the power and energy data collected from multiple voltage rails, supply inputs and battery contacts.

The run-time energy management software should run on the device and complement the standard OS power management. The main tasks of such software are:

- Detect and respond to power management directives by the OS or during OS transitions;
- During OS suspend periods manages the sleep, suspend and dark wake states of the device;
- Coordinate power state changes of components, clusters and subsystems;
- Manage IEE P2415 operating points and scenes;
- Execute control code to retain state information during transitions;
- Change power states by directly accessing hardware components, including clock dividers, PLLs and PMICs;
- Provide run-time power and energy estimates;

The energy management software has to have:

- Fast and precise execution of power state changes;
- Small footprint to fit within local memory of dedicated cores;
- Flexibility to distribute energy management across all the participating cores and the OS(es);
- Support for wide range of processors ranging from application processors to dedicated power management and always-on-cores;

Energy-Driven Component Hierarchy

Energy proportional systems require a shift towards energy proportional components which scale their energy consumption based on the workload. Some components exhibit better energy proportionality than others. The energy consumption of processors can be scaled by scaling the clock frequency, the voltage or just by cutting the voltage supply altogether. The memory can be put into a low power mode (refresh mode) without losing its content or turned off in which case the memory content gets lost and has to be preserved by the software. Communication components like Ethernet or WiFi are harder to scale as their energy consumption is dependent on interactions with other electronic components mostly outside of the motherboard and the device.

The mobile industry is increasingly introducing component hierarchies where the same internal functionality is provided by a number of alternative components with different energy characteristics. Similar to memory hierarchies with the L1 cache, L2 cache and the external memory, where the tradeoff is between speed, area and costs, mobile devices increasingly include processor hierarchies in order to smoothly control the energy consumption. An example of such a processor hierarchy is presented in an Anandtech article[20] (Figure 4).

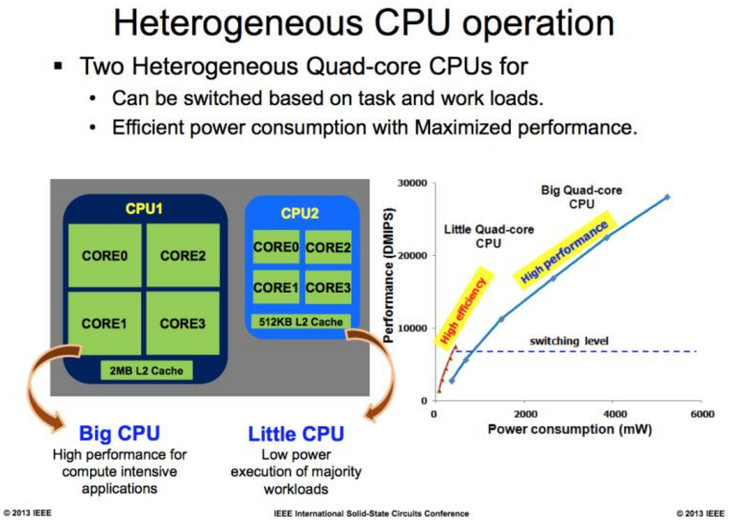

Heterogeneous CPU operation

- Two Heterogeneous Quad-core CPUs for
  - Can be switched based on task and work loads.
  - Efficient power consumption with Maximized performance.

© 2013 IEEE   IEEE International Solid-State Circuits Conference   © 2013 IEEE

Figure 4: Processor Hierarchy for Energy Proportionality

Improved power management in notebooks can be achieved by adding specialized microcontrollers to the device, also called always on cores or power management units. Such cores execute most of their software from private memories and consume between 50 and 100mW, which is one or two orders of magnitude less than what the powerful and energy hungry main processors used in notebooks consume. Although such microcontrollers are part of almost all PC systems on chips found in desktops and notebooks, their use today is limited to the operations within the system on chip and are

---

[20] http://www.anandtech.com/show/6768/samsung-details-exynos-5-octa-architecture-power-at-isscc-13

not used for controlling the energy activities of the motherboard and the whole device. Even more importantly, no provisions exist for the support of such microcontrollers in the operating system, so its use is limited to the system on chip producer and excludes the OEM.

Cost/Benefit Analysis

The hardware and software modifications and the engineering design procedures necessary to reach the savings in short idle energy consumption of 25% and long idle of 50% were estimated based on the qualitative energy analysis using the available data sheets, as well as the results we have achieved on the IPTV set top box project[21]. (In order to provide a detailed quantitative analysis of the proposed improvements for computers we recommend a project to develop a fully functional computer reference design and to conduct detailed measurements, testing and costs/benefits analysis. The estimated duration of such a project is between 12 and 18 months.

The additional costs for the proposed improvements fall into three categories: 1) component and royalty costs per device, 2) annual costs for the design technology and software infrastructure and 3) annual engineering labor costs.

For the costs for the inclusion of an always-on microcontroller on the SoC and to open up the energy management software we estimate at 15 cents per computer. These costs are based on the typical royalty costs for this type of microcontrollers and the accompanying energy management software. The cost for the increase in the silicon area of the system on chip is minimal and was neglected.

For the engineering team developing the improved computer the necessary investment for the software and tools infrastructure, the estimate is at around $1M per year. This includes the term (annual) licenses for various electronic design automation (EDA) software tools for energy design and management and the run-time software. Similar to other software term licenses, it represents the part of the costs for the software which is provided by specialized third parties and shared among the involved companies who use it.

For the engineering R&D labor to implement the hardware and software improvements, we estimate $2M for the increased non-recurring engineering (NRE) costs per year. We assume that these costs shall continue year after year as long as the company is in the computer business and is devoted to produce energy efficient devices. It is predicted that the reuse effect on the amount of work in the subsequent years shall be offset by the increasing complexity of the new components, software and devices, as well as the increased efforts to sustain the energy gains.

On the 59 million new computers (desktops and notebooks) shipped annually in the US and estimating around 10 computer engineering teams dealing with the hardware and software for energy optimizations, the software/tools and labor costs results in an additional NRE cost of approximately 51 cents per new computer produced. Summarizing the above we estimate that the total additional net costs per new computer shall be not more than 70 cents and have a retail price of less than $1 with the

---

[21] Final report of CEC EISG project #57751A/13-14: Unified Power Control for Set Top Box Devices (AGGIOS)

incremental retail markup as per the 2012 estimates of the Department of Energy for battery chargers. These costs are well below the typical CPU price of $40-$164.

Utilizing duty cycles taken from Appendix B and power draw estimates taken from the ENERGY STAR 6.0 QPL (October 2014), the energy benefits are calculated as follows:

- Desktops spend ~20% of their duty cycle in long idle (i.e., 1,753h per year) consuming on average 25.2W and ~20% in short idle consuming on average 26.8W;
- Notebooks spend 10% in long idle (i.e., 876h per year) consuming on average 4.85W and 10% in short idle consuming on average 8.55W;
-

Based on the above numbers, which are likely an underestimation of the market average given that ENERGY STAR products are the most efficient and assuming the current rate of 15 cents per kWh, a typical desktop computer with the suggested improvements would save 34 kWh/year and a typical notebook 4 kWh/year. Assuming a 4 year lifetime for desktops and a 3 year lifetime for notebooks, the net benefits of ~$20 and ~$2, for desktops and notebooks respectively, exceed the incremental costs described above, with lifecycle benefit cost ratios of ~22 and ~2, respectively.

Assuming a 37-63% split between produced desktops and notebooks, the suggested improvements would result in the combined savings for the US electricity ratepayers of $133M in the first year and ramping to $510M in the 4th year. This compares to the total incremental costs of $54M per year.

Conclusion

The qualitative analysis presented in this report proposes the steps necessary to improve computer energy consumption. The report suggests opening up the power management hardware and software in order to improve the collaboration on energy issues between various companies and design teams developing the computers. The largest potential for improvement is in the area of energy aware software development where the improved energy proportionality can lead to major additional savings. To reach such savings we advocate investments in industry standards, software, tools and development of energy proportional components. Our cost/benefit analysis indicates that every dollar invested in such improvements has the potential to return almost ten dollars in energy savings for the ratepayers year after year. In order to provide a detailed quantitative analysis of the proposed improvements we recommend initiating the development of a fully functional computer reference design and conducting detailed measurements, testing and costs/benefits analysis.