DOCKETED

Docket Number:	14-AAER-02
Project Title:	Computer, Computer Monitors, and Electronic Displays
TN #:	205337
Document Title:	Chris Hankin, Information Technology Industry Council Comments: June Deep Dive Mtg, ITI/Technet Computers Innovation Presentation -- Intel Corp
Description:	N/A
Filer:	System
Organization:	Chris Hankin/Information Technology Industry Council
Submitter Role:	Public
Submission Date:	7/13/2015 2:04:20 PM
Docketed Date:	7/13/2015

Comment Received From: Chris Hankin, Information Technology Industry Council
Submitted On: 7/13/2015
Docket Number: 14-AAER-02

June Deep Dive Mtg, ITI/Technet Computers Innovation Presentation -- Intel Corp

submitted in behalf of ITI and Technet

Additional submitted attachment is included below.

PC Power Management Innovation

Barnes Cooper

Chief Power Management Architect (Sr. PE)
Client Computing Group
Intel Corporation

Introduction

- About me
  - 23 years working on PC power management, mobile system architecture, energy efficiency
  - Lead architect for energy efficiency on phones to high end desktop
  - Projects:
    - ACPI (one original architects), Intel SpeedStep™ Technology, multi-core/thread PM
    - Energy Efficiency enhancements to numerous bus standards (USB, eDP, PCIe, SATA, …)
    - Lead architect Intel Power Optimizer
    - 20 year history working with Microsoft on OS/driver power management features/enhancements
  - Own/drive energy efficiency optimizations across the system
- Agenda
  - A look back at the history of PC power management
  - Challenges specific to desktop platforms

Early 90's – first mobile 'laptops'

- Very primitive power and thermal management schemes
  - DOS-based 286 through Pentium® processor designs
  - BIOS and then SMM-based power management
- Hard to figure out basic things
  - Is the user present?
    - Timers/trap on KBD/mouse
  - Is the HDD in use?
    - Timers/trap on disk accesses, set timers
  - Is the system compute busy?
    - Global timer/trap on all system resources
  - Is the CPU idle?
    - I don't know, OS is always doing something!
- 3 hours of battery life was great…

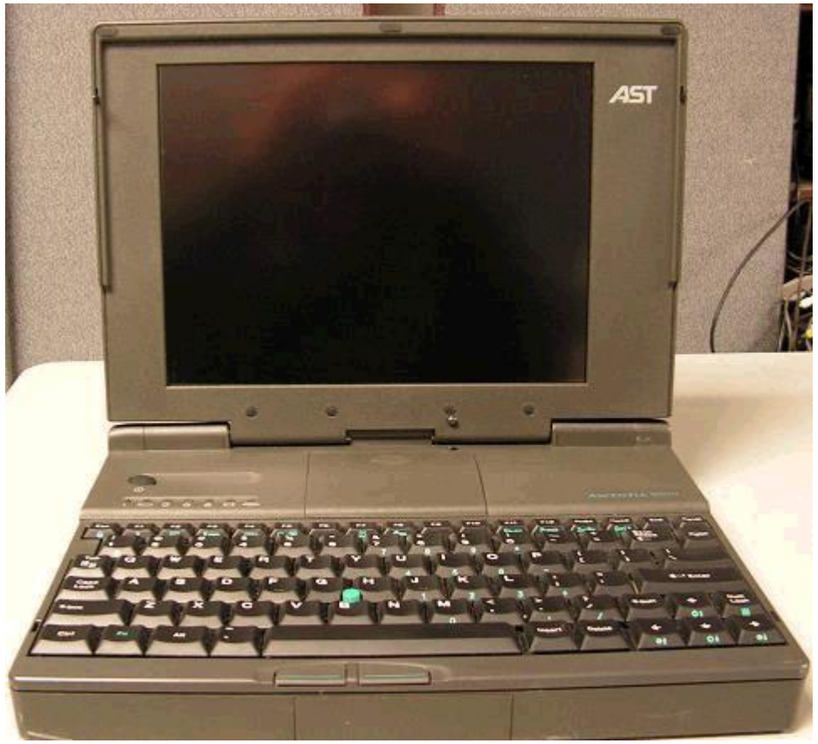

Advanced Power Management (APM)

- APM was the first interface to allow power management cooperation from OS (Windows 3.1 – 98)
  - OS tells BIOS when system resources are not in use
    - Screen, disk, CPU, keyboard
    - Coordinated system idle state and some primitive power policies (AC/battery)
  - BIOS (SMM) did the system specific controls to power down idle elements
- Good step forward, but this approach did not work for Windows NT
  - Windows NT was deemed the future of Windows OSes
- So, we created Advanced Configuration and Power Interface (ACPI)…

ACPI

- In 1995, we started working with Microsoft on ACPI
  - We were told OS knows everything, so that is the right place to do it
  - It was highly controversial as OEM's valued their BIOS-driven PM
- ACPI introduced concepts that allowed for OS-control of PM
  - Interpreted new language for doing configuration and power management
    - ACPI Source Language (aka ASL Code)
  - Invocation of BIOS-supplied control methods (to do system-specific actions)
  - Formal definitions and semantics for CPU, system and device low power states (C, S, D)
- A key dramatic improvement is hardware knows when the processor(s) are idle
  - C1, C2, C3, … C10, etc.
- ACPI also supports performing OEM specific value add power management (and configuration) features in OS-friendly manner
  - ACPI ASL code ships with firmware image (BIOS) and is invoked on events as specified
- ACPI *is* the configuration and power management interface between Microsoft and Linux-based operation systems on PCs

P States (Intel SpeedStep™, Geyserville)

- In the late '90s, the frequency race was on
  - We shipped processors binned at highest possible performance (V/f)
  - Power levels became intractable (e.g. pushing people to 18W and higher thermal solutions)
  - We started looking at dynamically changing CPU frequency and voltage
- At the time, it was a difficult technology to gain confidence in
  - Silicon designed to always pass through a reset once voltage and clocks are stable
  - P states were a new paradigm (stop, well-specified voltage changes, PLL relocks, no reset, etc.)
- Industry followed immediately, and we applied to other SoC elements over time (e.g. graphics RP states)
- But, we had to make OS-vendors comfortable with the concept
  - That was hard, with lots fear/doubt about software-timing loops. However, we never found any and we added P state control to ACPI 2.0
- Net result is silicon runs at the most power efficient state to meet the demand, or thermal/power limit at any instant

Intel Power Optimizer

- After processor and graphics were well power-managed, we started working on reducing the idle power of rest of platform (RoP)
  - Key concept: We need time (latency-tolerance) to do long latency power management optimizations (PLL shutdown, power gating, voltage regulator efficiency changes)
- We needed guaranteed idle and warmup times to stagger activity
- We needed to preserve Intel Architecture-based Personal Computers (IAPC) goodness
  - Timers should fire when scheduled (no skid)
  - Bus mastering devices should not overrun/underrun
  - No end user artifacts (video or audio glitches)
  - Solving all the above was a very big challenge
- We finally launched Power Optimizer on Haswell offering a 25x platform idle power reduction
  - Our biggest leap forward in battery life/energy efficiency as a result of the most extensive re-architecture of the platform
  - Framework requires all devices on platform to participate (cooperative-model)

Continuous Work to Improve Energy Efficiency

- USB - Originally developed as a ISA plug-in card replacement technology, its architecture was inherently inefficient
  - Software polled controller, controller polled devices, devices stayed active all the time (crying babies analogy), all controllers did this asynchronously (unaligned)
  - Developed a multitude of technologies to address through bus standards work (LPM L1), through OS, existing controllers (Enhanced Host Controller - EHCI)
  - USB 3.0 and eXtensible Host Controller (xHCI) were made energy efficient from inception
- Displays – we drive the same image out to display 60 or more times per second (it worked the same way on CRTs and LCDs, image decays)
  - We made a change to eDP for Panel Self Refresh, only send changed frames
- Software (enabling ISVs for energy efficiency) is very challenging
  - Over time, we have made great progress with tools and enabling effort
  - SoCWatch, VTune, Battery Life Analyzer are robust tools for identifying hardware and software inefficiencies
- All new technologies are reviewed at Intel to ensure they do not introduce an energy efficiency issue

Bringing Green Application to Consumers

- Optimized 30+ applications for power hygiene in past 5+ years
- Reduced average battery power consumption by 1.5 Watts through software optimization in top 20 application

- Collection done on same hardware but different software version
- Several application in its own category has different battery life but still significant progress has been made to meet best in class application

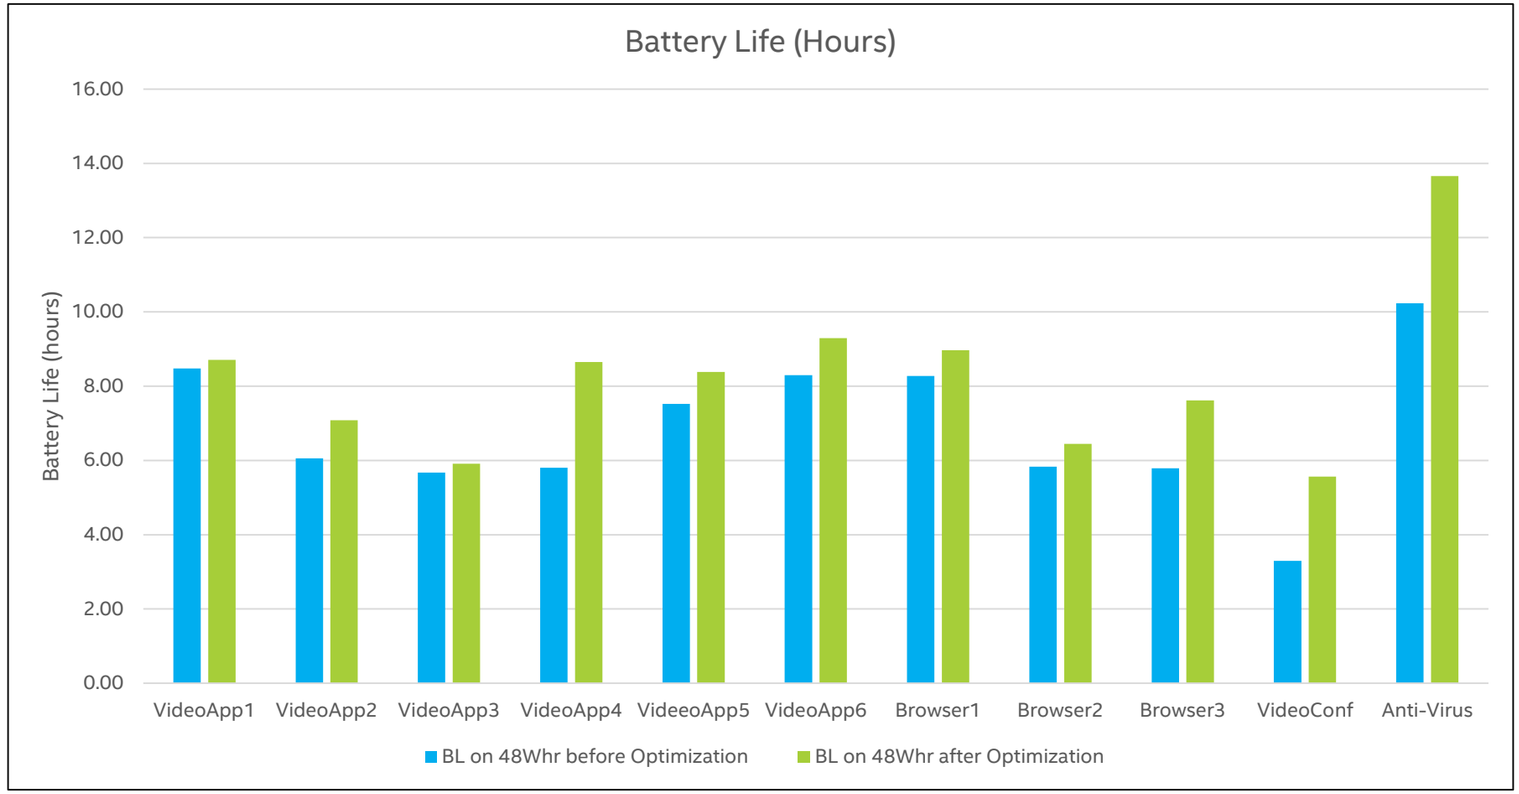

*Power Collection done with NiDAQ on BDW-U Platform with Windows 8.1 x64

Continued PC Power Management Innovation

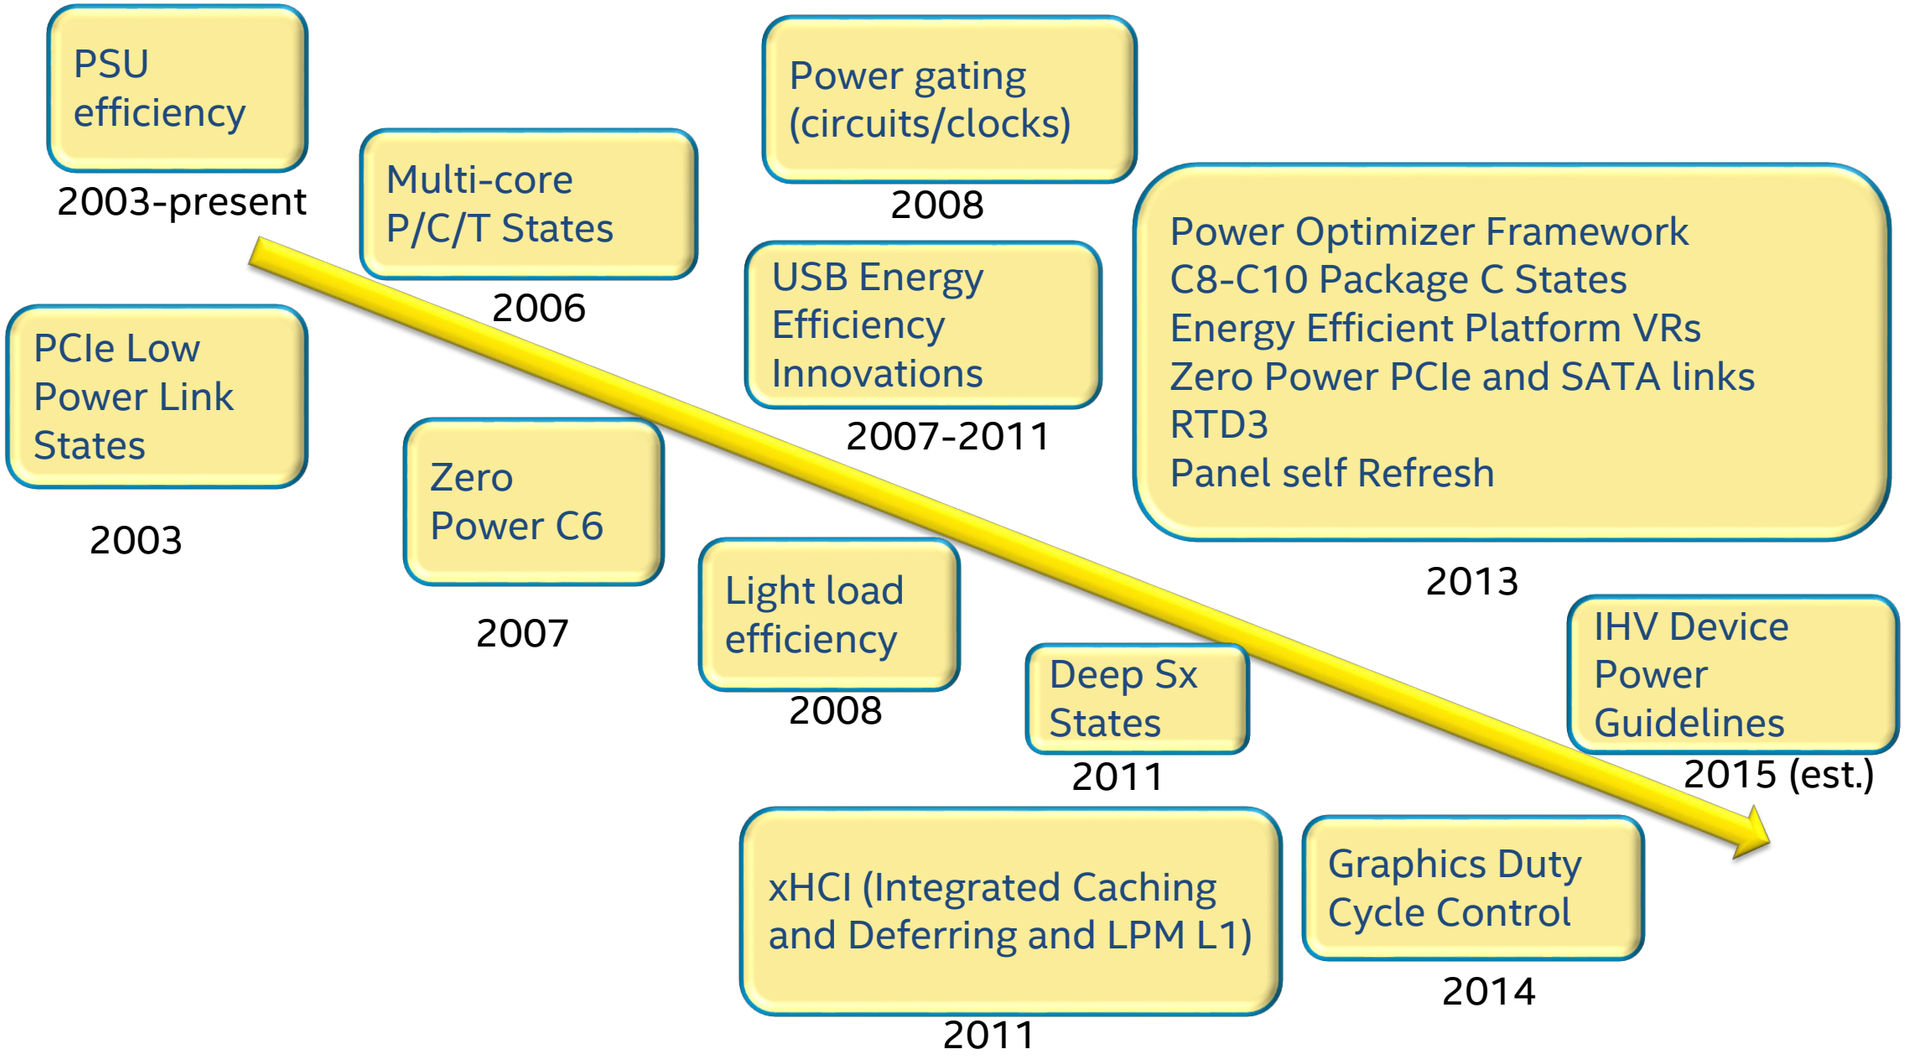

Source: Intel Corp.

Moore's Law drives Energy Savings at the Silicon

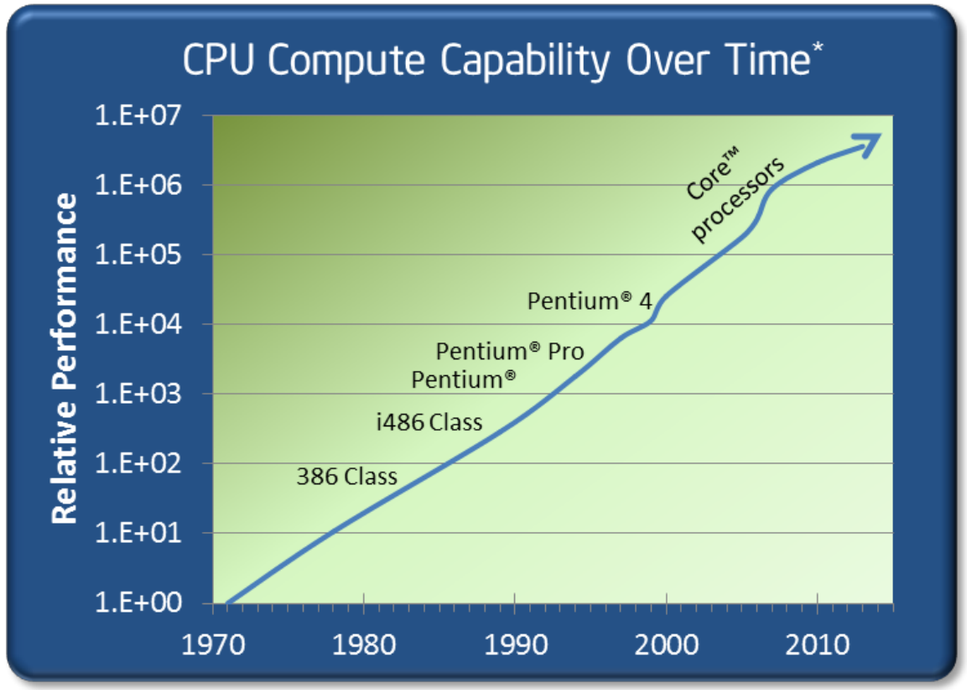

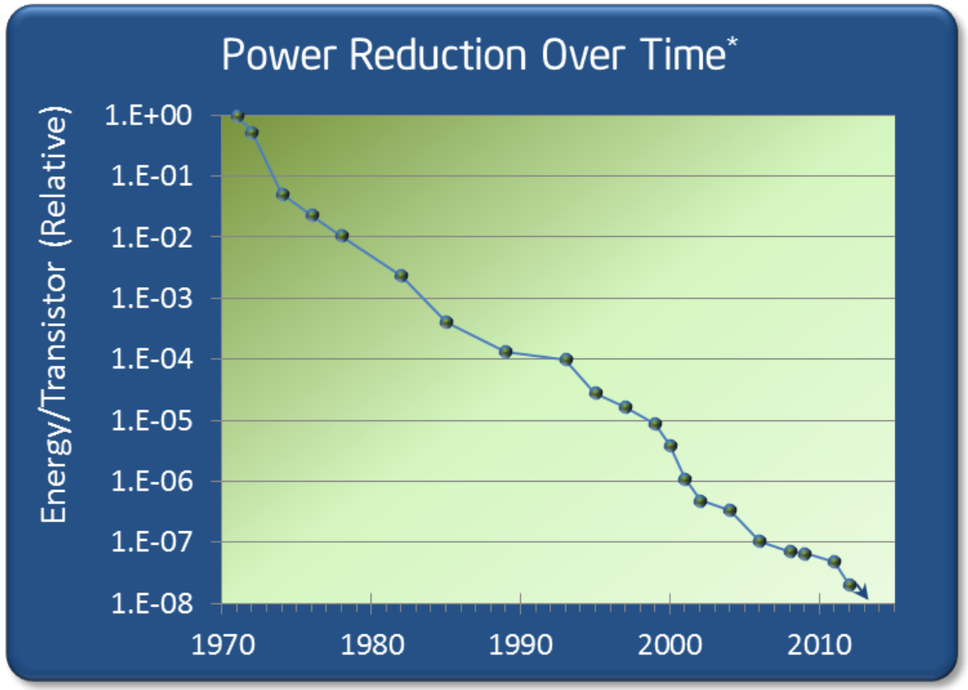

Source: Intel Corporation – January 2014. Relative performance chart estimates based on reported MIPS and SPEC CPU scores over this time period (as configurations and workloads change with time)

- ❑ Delivering great performance within power envelope
- ❑ Compute Energy Efficiency ➔ Positive Impact On Environment

What about desktop power?

- We have been reducing YoY desktop platform power
  - Skylake (2015) lower than Haswell and Broadwell
  - Best case motherboard design shown below (Note: There is large variability between sampled designs)

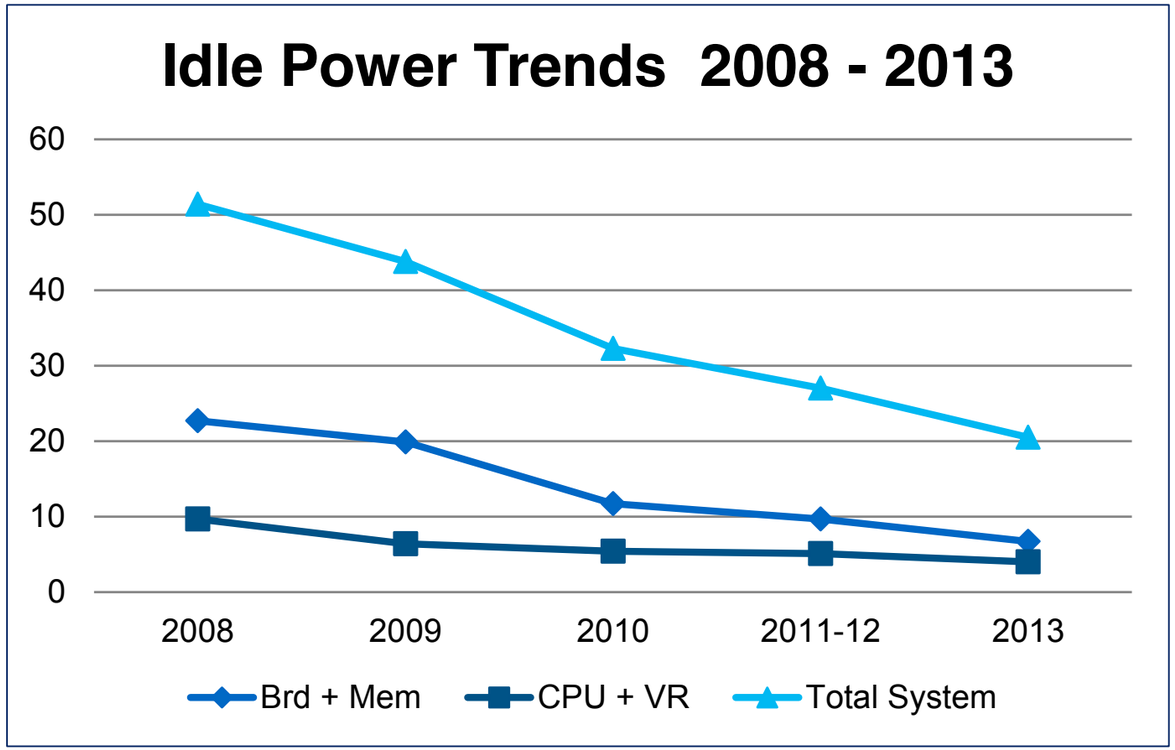

- Why are desktop parts higher power than mobile components at idle?
  - They are higher power as they are binned for performance and frequency, thus, they inherently have higher leakage
  - PC6 turns off the IA and GT cores, but rest of SoC remains powered (uncore, etc.)

Table 25. Maximum Idle Power Specification

Symbol	Parameter	U-Processor 15W with GT3		U-Processor 28W with GT3		Y-Processor 6W SDP / 4.5W SDP with GT2		Unit	Note
		Min	Max	Min	Max	Min	Max		
$P_{PACKAGE(C7)}$	Package power in Package C7 state	–	0.95	–	1.5	–	0.85	W	1, 3
$P_{PACKAGE(C8)}$	Package power in Package C8 state	–	0.12	–	0.18	–	0.1	W	1
$P_{PACKAGE(C9)}$	Package power in Package C9 state	–	0.052	–	0.1	–	0.04	W	1
$P_{PACKAGE(C10)}$	Package power in Package C10 state	–	0.052	–	0.1	–	0.04	W	1

Table 21. Desktop Processor Thermal Specifications

Product	PCG[8]	Max Power Package C1E (W) [1, 2, 5, 9]	Max Power Package C3 (W) [1, 3, 5, 9]	Min Power Package C3 (W)[9]	Max Power Package C6 (W) [1, 4, 5, 9]	Max Power Package C7 (W) [1, 4, 5, 9]	Min Power Package C6/C7 (W)[9]	TTV Thermal Design Power (W) [6, 7, 10]	Min T_{CASE} (°C)	Max TTV T_{CASE} (°C)
Quad Core Processor with Graphics	2013D	26	20	1.0	3.5	3.4	0	84	5	Processor (PCG 2013D) Thermal Profile on page 67

Why not just turn on mobile features on desktop silicon?

Reducing Desktop Power

- To get the power closer to mobile power levels, we would need to enable deep package C states
  - e.g. PC10 @ ~3ms
  - We have not turned on all mobile features in desktop, because it adds cost and introduces significant ecosystem enabling and end user visible challenges
  - What should we do when older card, not supporting Power Optimizer, is plugged into newer system?
- Power supplies inherently inefficient at ultra light load
  - Some even fall out of regulation
    - Likely can be enabled, but will add to system cost
  - Likely need to add VR signaling from SoC/platform to VR as well as potentially a light load rail from power supply to system
    - These are long lead time items (new power supplies, new connectors, new motherboards)
    - This will add to system cost

Other Desktop Power Challenges

- Hard disk drive (HDD) and optical disk drive (ODD) consume considerable power (Watts)
  - ODD use diminishing; will add to system cost to enable zero power ODD
  - Powering HDD off results in poor end user experience when drive access is required (multi-second system stall)
  - Caching or hybrid technologies can mitigate but not guaranteed
  - Bad policy results in end user turning power management features off
- Increased display resolution and performance will require additional power allowances within CEC regulation
  - 8k displays require >50Gbps of bandwidth for static screen refresh, due to OS architecture, the display bandwidth is much higher on the SoC (multiple planes)
  - New bus standards work to support more intelligent display update protocols are required analogous Embedded Display Port's Panel Self Refresh (eDP PSR)
  - Only send frame/portion or frame that have changed, etc.
  - This is a long term effort involving bus standards and large scale ecosystem enabling

Some Other Discussion Points

- Aggios report compares IAPCs to closed systems (set top box)
  - A closed system can be better than an open system, but the value of a PC is in its openness
- Decades of work with OS vendors have resulted in dramatic behavioral changes
  - All OSes are 'tickless', dramatic reductions in OS-idle activity factors (S0ix-based OSes)
  - In-box drivers for display, I/O, storage, audio, etc. are heavily optimized
- ACPI does allow for customer value add feature energy efficiency enhancement
- Add a $0.15 microcontroller to do fine-grained power management
  - Intel platforms have had these since Nehalem (2009 product)
    - http://www.treehugger.com/gadgets/intels-next-cpu-to-include-dedicated-power-control-unit-to-save-power.html
  - Provides fine-grain PM control and sequencing of power gates, VRs, clocks in all our products, but still operates within constraints and guidance of the OS
- No easy ways to dramatically reduce power at this point
  - We can and should pursue, but they will add significant cost and take a considerable amount of time (ecosystem enabling challenges)
  - Key question: *Will customers pay for it and/or turn such features off?*

Summary / Challenges

- As an industry, we have made computing dramatically more efficient continuously, year over year, with some remarkable breakthroughs
  - This continues every generation
- Reducing desktop idle power levels is non-trivial, due to
  - Large enabling ecosystem (a multitude of third party devices, old and new)
  - Power supply dependencies
  - Customer aversion to performance/usability tradeoffs in desktop
    - e.g. waiting 5+ seconds for disk to spin up
  - Desktop platforms/markets segments are extremely cost sensitive
  - Exponential increase in external display resolution and bits per pixel will dramatically increase baseline short idle power; will require additional allowances, and industry momentum to improve power levels

Thank you!!!

Legal Notices and Disclaimers

INFORMATION IN THIS DOCUMENT IS PROVIDED IN CONNECTION WITH INTEL PRODUCTS. NO LICENSE, EXPRESS OR IMPLIED, BY ESTOPPEL OR OTHERWISE, TO ANY INTELLECTUAL PROPERTY RIGHTS IS GRANTED BY THIS DOCUMENT. EXCEPT AS PROVIDED IN INTEL'S TERMS AND CONDITIONS OF SALE FOR SUCH PRODUCTS, INTEL ASSUMES NO LIABILITY WHATSOEVER AND INTEL DISCLAIMS ANY EXPRESS OR IMPLIED WARRANTY, RELATING TO SALE AND/OR USE OF INTEL PRODUCTS INCLUDING LIABILITY OR WARRANTIES RELATING TO FITNESS FOR A PARTICULAR PURPOSE, MERCHANTABILITY, OR INFRINGEMENT OF ANY PATENT, COPYRIGHT OR OTHER INTELLECTUAL PROPERTY RIGHT.

Intel may make changes to specifications and product descriptions at any time, without notice. Designers must not rely on the absence or characteristics of any features or instructions marked "reserved" or "undefined". Intel reserves these for future definition and shall have no responsibility whatsoever for conflicts or incompatibilities arising from future changes to them. Do not finalize a design with this information.

The products described in this document may contain design defects or errors known as errata which may cause the product to deviate from published specifications. Current characterized errata are available on request.

- All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest Intel product specifications and roadmaps.

Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors.

All products, computer systems, dates and figures specified are preliminary based on current expectations, and are subject to change without notice. Results have been estimated or simulated using internal Intel analysis or architecture simulation or modeling, and provided to you for informational purposes.

Intel technologies may require enabled hardware, specific software, or services activation. No computer system can be absolutely secure. Intel does not assume any liability for lost or stolen data or systems or any damages resulting from such losses. Tests document performance of components on a particular test, in specific systems. Differences in hardware, software, or configuration will affect actual performance. Consult other sources of information to evaluate performance as you consider your purchase.
- For more complete information about performance and benchmark results, visit http://www.intel.com/performance

Intel is a sponsor and member of the BenchmarkXPRT Development Community, and was the major developer of the XPRT family of benchmarks. Principled Technologies is the publisher of the XPRT family of benchmarks.

Intel Corporation may have patents or pending patent applications, trademarks, copyrights, or other intellectual property rights that relate to the presented subject matter. The furnishing of documents and other materials and information does not provide any license, express or implied, by estoppel or otherwise, to any such patents, trademarks, copyrights, or other intellectual property rights.

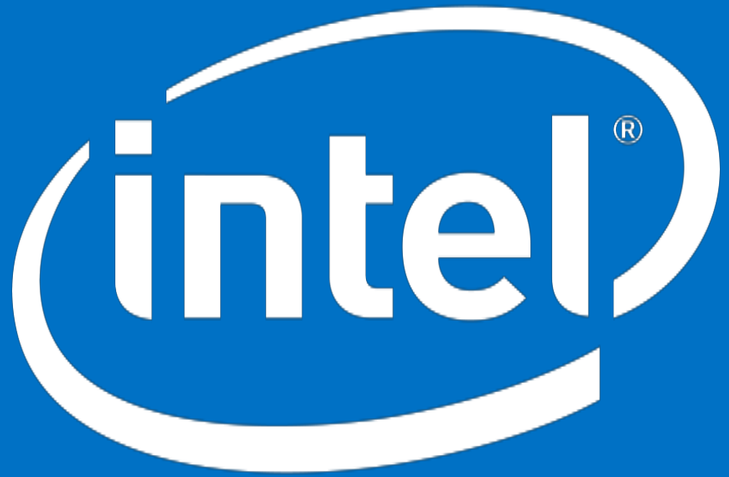
intel