DOCKETED

Docket Number:	14-AAER-02
Project Title:	Computer, Computer Monitors, and Electronic Displays
TN #:	206276
Document Title:	ITI & Technet 9/29 F2F Presentation: EE DT Feasibility
Description:	Energy Efficient Desktop Feasibility Discussion
Filer:	System
Organization:	Chris Hankin, Information Technology Industry Council
Submitter Role:	Public
Submission Date:	10/5/2015 6:25:16 AM
Docketed Date:	10/5/2015

Comment Received From: Chris Hankin, Information Technology Industry Council
Submitted On: 10/5/2015
Docket Number: 14-AAER-02

ITI & Technet 9/29 F2F Presentation: EE DT Feasibility

Additional submitted attachment is included below.

Energy Efficient Desktop Feasibility Discussion

Barnes Cooper
Chief Power Management Architect (Sr. PE)

Vidoot Ponnala
Desktop PM Architect

Client Computing Group
Intel Corporation

Desktop Platform Idle Power Trend 2008 – 2017

- Estimates beyond 2015 are based on trends
  - –Desktop power has dropped significantly since 2008
  - –Year over year power reduction opportunity leveling off

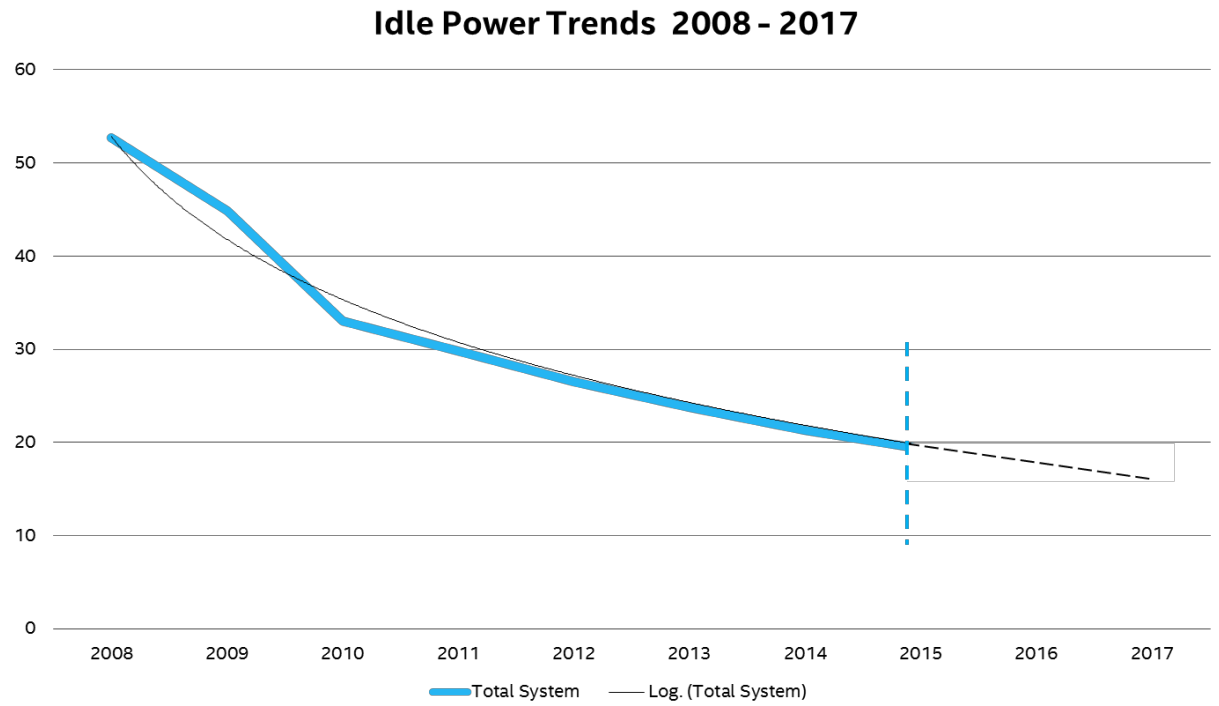

Data Source - Intel

Technical Feasibility Synopsis

- Meeting CEC proposed requirements for desktops may be technically feasible, but present significant industry investment and enabling challenges, that are unlikely to meet CEC's cost effectiveness criteria. Desktop platform focus areas:
- Deep silicon and power supply enhancements
  - Deeper package C states on DT silicon
  - With considerable system ecosystem dependencies and validation costs, as well as usability issues (in field H/W changes)
- Power supply efficiency optimizations (covered as a separate topic)
- Motherboard optimizations and design guidelines
  - 6-7W measured power deltas between best and worst of class motherboards
- 0W ODD and 1W HDD
  - HDD through OS background activity suppression (e.g. Microsoft Modern Standby) as well as integrated or on system non-volatile memory
  - Zero power ODD (commonplace for mobile and present on mobile FF drives), but added cost
- OS
  - OS support of a long idle state (i.e. S0ix) critical
  - Key investments needed to reduce background activity, wired network offloading/filtering, and HDD access reduction
- Longer term, short idle power increases due to bandwidth impact of higher resolution displays will be a huge issue
  - We must work/support new I/O protocols to reduce external display platform power impacts
  - Similar protocols to eDP Panel Self Refresh (send only what has changed, not full frames)

Deeper Desktop C States

- 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.)
- C10 can potentially be enabled for desktop components, but there are many challenges
  - –Next slide…

Table 24. 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, 4

Source: http://www.intel.com/content/www/us/en/processors/core/4th-gen-core-family-mobile-u-y-processor-lines-vol-1-datasheet.html

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 and 2014	26	20	1.0	3.5	3.4	0	84	5	Processor (PCG 2013D and PCG 2014) Thermal Profile on page 68

Source: http://www.intel.com/content/dam/www/public/us/en/documents/datasheets/4th-gen-core-family-desktop-vol-1-datasheet.pdf

Enabling Deep Lower Power States – HW Impacts

- Depending upon the actions we wish to perform while in a low power state, integrated IPs and all motherboard devices will need to be tolerant of the required latency for the action
  - For example, today mobile C10 can incur a 3ms latency
  - All peripherals must be tolerant of this latency
- Latency is communicated to the platform in several ways:
  - Implicit: ascertained by link state (e.g. SATA active/partial/slumber/DEVSLP#)
  - Explicit: communicated by device (e.g. PCIe LTR, USB2 LPM L1)
  - Assumed: OS turns the device off (third party graphics)
- There will be significant validation burden to silicon suppliers as well as our customers building systems (added cost)
  - Beyond cost, there are long lead-time validation and industry enabling challenges in a diverse desktop ecosystem
- There are significant challenges with respect to in field hardware upgrades
  - System purchased with integrated graphics, and populates older external graphics card
  - Device does not support RTD3 and/or LTR
  - We can no longer enter the low power state!
- In short, we need all integrated peripherals to support latency plumbing and/or RTD3, as well as a graceful way to tell user that an 'unfriendly' peripheral has been populated

Enabling Deep Lower Power States – OS Impacts

- To maximize ROI of energy cost of transition to low power states, OS to minimize number of break events
  - For example, Microsoft Connected Standby implementations switch to demand-based tick
- Network traffic suppressed through push infrastructure and programming of network filters, so we only wake on network packets that are targeted to the particular system of interest
- Unused devices should be placed in RTD3 aggressively beforehand
- Depending on the power supply configuration and actions taken there, there is likely more RTD3 dependence than on a traditional mobile system
  - For example, we likely need the graphics to be in RTD3 cold such that it has no dependency on low latency access to system resources
- Need OS infrastructure to inform user of non-participating device
  - e.g. *You have populated a device that has dramatically impacted your energy efficiency*
- Note that installation of poorly behaving software can result in significant power regressions (increases)

Intel Device Power Management Collateral

- Driving power management improvements for all client platform devices with detailed collateral since 2013
  - Audio, Bluetooth, Display, Gfx, LAN, NFC, Touch Screen, WLAN, Storage (HDD, SSD, ODD), USB, WWAN
  - OS agnostic – include OS caveats as appropriate
  - Creation of future generation collateral through 2018 in progress
- Content and updates based on engineering evaluation and iterative customer & vendor feedback
- Collateral available under Intel NDA to OEM/ODM customers and IHV technology vendors
- *"Platform Device Power Targets and Related Recommendations: Device Power Specification"*
  - *2015 Guidelines v1.1, September 2015; 2016 Guidelines v0.7, September 2015; 2017 Guidelines v0.5, September 2015; 2018 Guidelines, WIP.*
- *"Runtime D3 (RTD3) Hardware and Software Recommendations"*
  - *2015 Guidelines v1.0, September 2015; 2016 Guidelines v0.7, September 2015; 2017 Guidelines v0.5, September 2015; 2018 Guidelines, WIP.*

HDD Power Reduction

- Power reduction of hard drives in idle mode is problematic (HDD presentations to follow)
  - Most systems configured to never spin down drives in short idle mode
  - If not, most users disable once they feel the 5-8s spin up time of a desktop HDD (per OEM data)
  - Need a cost effective solution that allows drive to be powered off, and avoid usability issues such that it is actually used (long idle, user away, immediate spin-up when user wakes system)
- Solving this requires a combination of OS as well as hardware support
  - Some amount of non-volatile memory (at least 16GB, likely 32GB)
  - OS S0ix support such as Windows Modern Standby such that:
    - Background app activity suppressed
    - Background disk service activity suppressed (virus scanner, etc.)
- Issues
  - Reliability of drive, rated spindown times
  - Energy cross-over point (amortization), quick spin up/down cycles should be avoided
  - Cost adder could be appreciable to do this properly and have the feature actually be used

Conclusions and Next Steps

- CEC targets will require comprehensive system-wide approach to reduce system power through a combination of:
- Deep silicon low power states and the validation and ecosystem work associated with that
- Power supply efficiency optimizations
- Motherboard optimizations and design guidelines
- 0W ODD and 1W HDD
  - Non-volatile memory and OS support for HDD power down, zero-power ODD (if present)
- OS
  - OS support of a long idle state (i.e. S0ix) is critical to making it all work
  - Key investments needed to ensure/enforce platform quiescence (background activity, wired network filtering, and HDD accesses)
- Longer term, short idle power increases due to bandwidth impact of higher resolution displays will be a huge issue (will discuss later)
  - We must work/support new I/O protocols to reduce external display platform power impacts
  - Similar protocols to eDP Panel Self Refresh (send only what has changed, not

Considerations for future system bandwidth

Gary Chow
Client Planning & Architecture

Client Computing Group
Intel Corporation

System SoC Integration Trends

- Moore's law continues to drive silicon integration of new capabilities and functionality with each new generation of products
- Total SoC BW requirements are determined by evaluating the needs of each subsystem (e.g. CPU, Graphics, Media, Imaging, I/O, etc)
- Each PC segment will undergo integration at differing rates due to performance, power, cost, design complexity trade-offs.
- Forward looking market access requirement must comprehend the "design space" that future products can occupy to ensure future innovation.

Potential Memory topologies

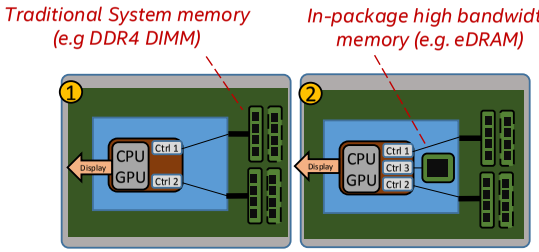

1-Mainstream Integrated Graphics, 2-Performance Integrated Graphics

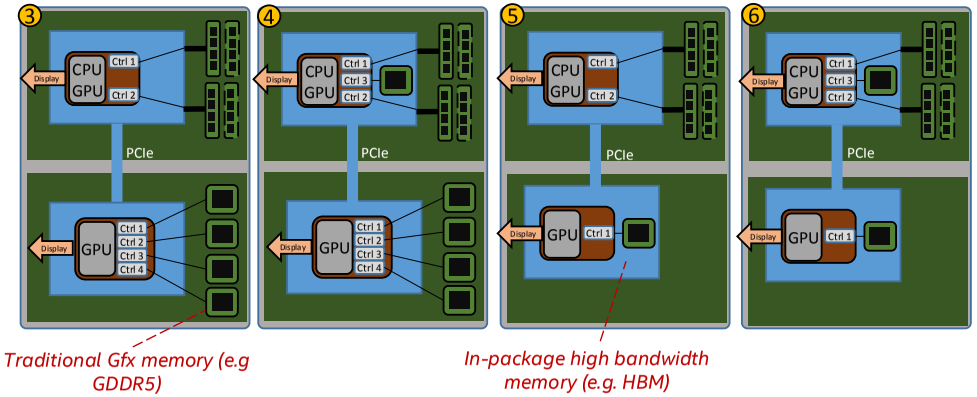

3-Mainstream Discrete Graphics, 4-Performance Integrated & Discrete Switchable Graphics, 5-High Performance Discrete Graphics, 6-High Performance Integrated & Discrete Switchable Graphics,

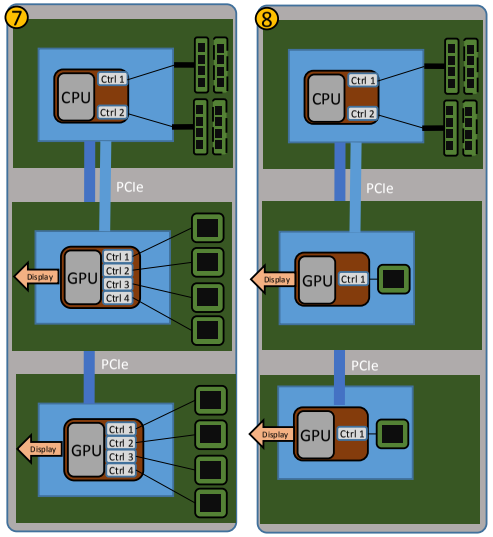

7- Desktop with two mainstream Discrete Graphics, 8- Performance Desktop with two High Performance Discrete Graphics

System Bandwidth Trends

- Different tiers of performance categories have diverging BW trajectories. Higher performance tiers have steeper year-over-year growth.
- As functionality is integrated into silicon, BW requirements takes on a broader range but currently there is no mechanism for additional integrated functionality to expand beyond Base TEC limits.

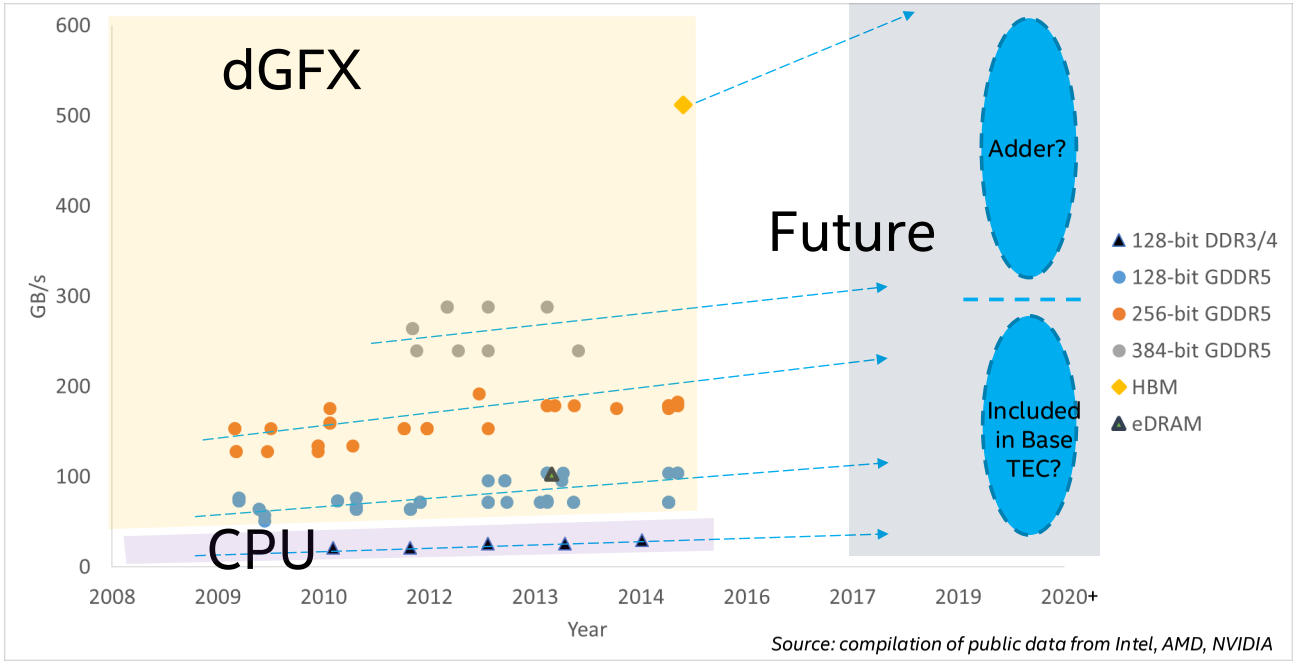

Industry to propose TEC adders based on system memory BW:

- System memory BW serves as a proxy for processor capability
- System BW adders to coexist with other proposed adders (does not replace other proposals)
- Products following historical BW growth trend should be covered by Base TEC while step-function improvement capability should be eligible for an adder.

Legal Notices and Disclaimers

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 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