NINETEENTH ANNUAL

BiTS

Burn-in & Test Strategies Workshop™

March 4 - 7, 2018

Hilton Phoenix / Mesa Hotel
Mesa, Arizona

Archive

COPYRIGHT NOTICE

www.bitsworkshop.org

The Long Road to Energy Efficiency
Are we there yet?

Barnes Cooper

Intel Fellow, Chief Architect, Power Management and Energy Efficiency

Client Computing Group, Intel Corporation

Acknowledgements: Tawfik Arabi, Gary Ghow, Jim Hermerding, John Powell

CLIENT COMPUTING GROUP

Introduction

About me

- 23 years working on PC power management, mobile system architecture, energy efficiency
- Lead architect for energy efficiency (battery life) for all client devices (mobile to high end desktop)
- Projects:
  - ACPI (one of original architects), Intel SpeedStep™ Technology, multi-core/thread PM
  - Energy Efficiency enhancements to numerous bus standards (USB, eDP, PCIe, SATA, …)
  - Lead architect behind biggest reductions in client idle power (mobile and desktop)
  - 20 year history working with Microsoft on OS/driver power management features/enhancements
- Mission: Pursue energy efficiency optimizations anywhere on the platform (software, board, firmware, OS, silicon, I/O standards/implementation, etc.)

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? (timer/trap on KBD/mouse I/O)
- Is the HDD in use? (timer trap on disk I/O)
- Is the system compute busy? (timers based on all relevant I/O and processing)
- Is the CPU idle? (I don't know, OS is always doing something!)

2-3 hours of battery life was great…

CLIENT COMPUTING GROUP

intel 3

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)…

CLIENT COMPUTING GROUP

intel | 4

ACPI

In 1995, we started working with Microsoft on ACPI

- Move power management control/policy to the OS; 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)
- 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

  - Provided a framework for innovation 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 still* the configuration and power management interface between Microsoft and Linux-based operation systems on PCs

ACPI offered and immediate power savings by providing process CPU power control

CLIENT COMPUTING GROUP

intel | 5

Dynamic frequency / voltage control

In the late '90s, the frequency race was on

- We shipped processors binned at highest possible performance (V/f), so did our competitors
- Power levels became intractable (e.g. pushing people to 18W and higher thermal solutions in mobile systems)
- We started looking at dynamically changing CPU frequency and voltage

At the time, it was a difficult technology to gain confidence in (voltage/freq change with no reset)

- 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 (fear of timing loops)

Net result is silicon runs at the most power efficient state to meet the demand, or thermal/power limit at any instant

This mechanism is applied broadly across IPs throughout the platform today known as P states (performance states)

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 launched Power Optimizer on 4^{th}-gen Core Platforms offering significant 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)

Rest of System Power

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, send only update part of frame in the future

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

Power Deliver – spec'd tested, constant innovation and change both internal and external chasing highest possible efficiency at all point of load

All new technologies are reviewed at Intel to ensure they do not introduce an energy efficiency issue

CLIENT COMPUTING GROUP 8

Desktop Platforms

In 2015, California Energy Commission began work on a new energy regulation that had dramatic reductions for next generation platforms

- The whole industry became involved; creative solutions were created
- CEC made reduction law Tier 1 effective date: 1/1/19; Tier 2: 7/1/2021

For Intel platforms, we will introduce solutions fully capable of meeting this standard as well as world-wide embodiments this year

- Through a combination of deep-silicon PM, power supply management, NVM storage, and introduction of S0ix states through OS frameworks such as Windows Modern Standby, we expect to deliver significant idle power reduction (as measured by AC power at the wall)
- Expect some system vendors to offer full Tier 2 CEC compliance in 2018

Product Energy Efficiency

- From 2010 to 2020, compute performance per unit of energy will improve by:
  - 50x for Notebook processors
  - 20x for Desktop processors

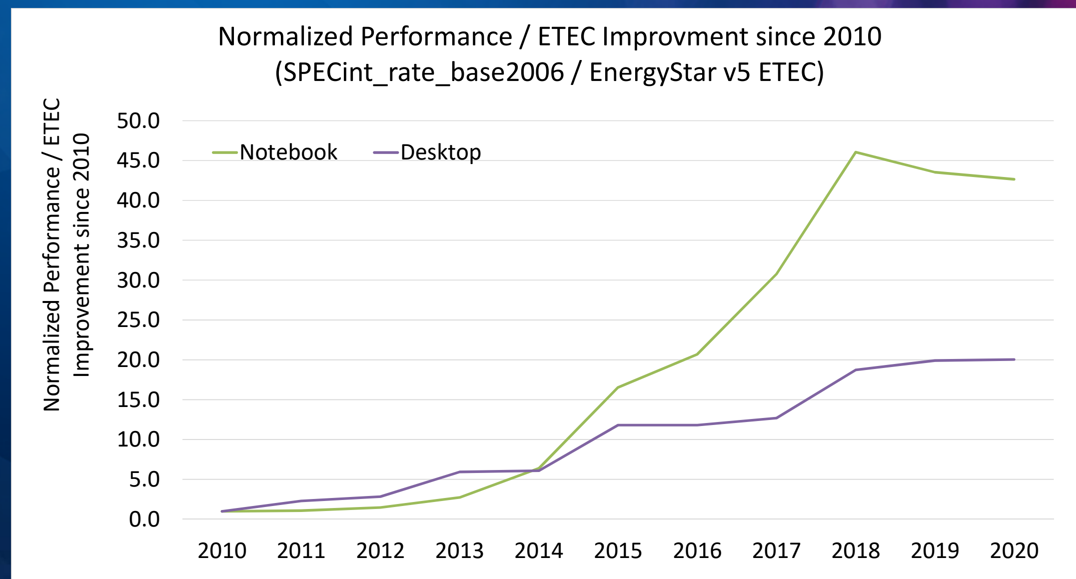

1. Improvement metric based on SPECint_rate2006 divided by ETEC as specified by Energy Star version 5.0
2. Data represents processor models from mainstream Desktop and Notebook product lines.

Benchmark results were obtained prior to implementation of recent software patches and firmware updates intended to address exploits referred to as "Spectre" and "Meltdown". Implementation of these updates may make these results inapplicable to your device or system.

- Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors.
- Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products.
- **For more information go to** http://www.intel.com/benchmarks

CLIENT COMPUTING GROUP

intel | 10

Normalized C02 Emissions

- From 2010 to 2020, compute performance per unit of energy will improve by:
  - >90% reduction for Notebook processors
  - >80% reduction for Desktop processors

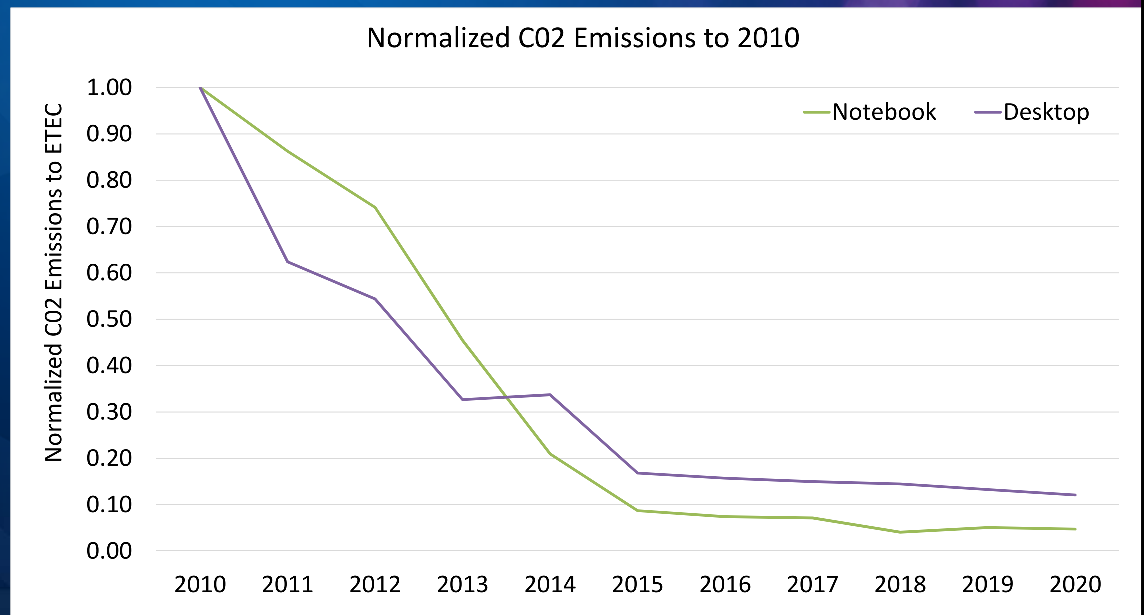

1. Improvement metric based on ETEC as specified by Energy Star version 5.0
2. Data represents processor models from mainstream Desktop and Notebook product lines.

Benchmark results were obtained prior to implementation of recent software patches and firmware updates intended to address exploits referred to as "Spectre" and "Meltdown". Implementation of these updates may make these results inapplicable to your device or system.

- Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors.
- Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products.
- **For more information go to** http://www.intel.com/benchmarks

CLIENT COMPUTING GROUP

intel | 11

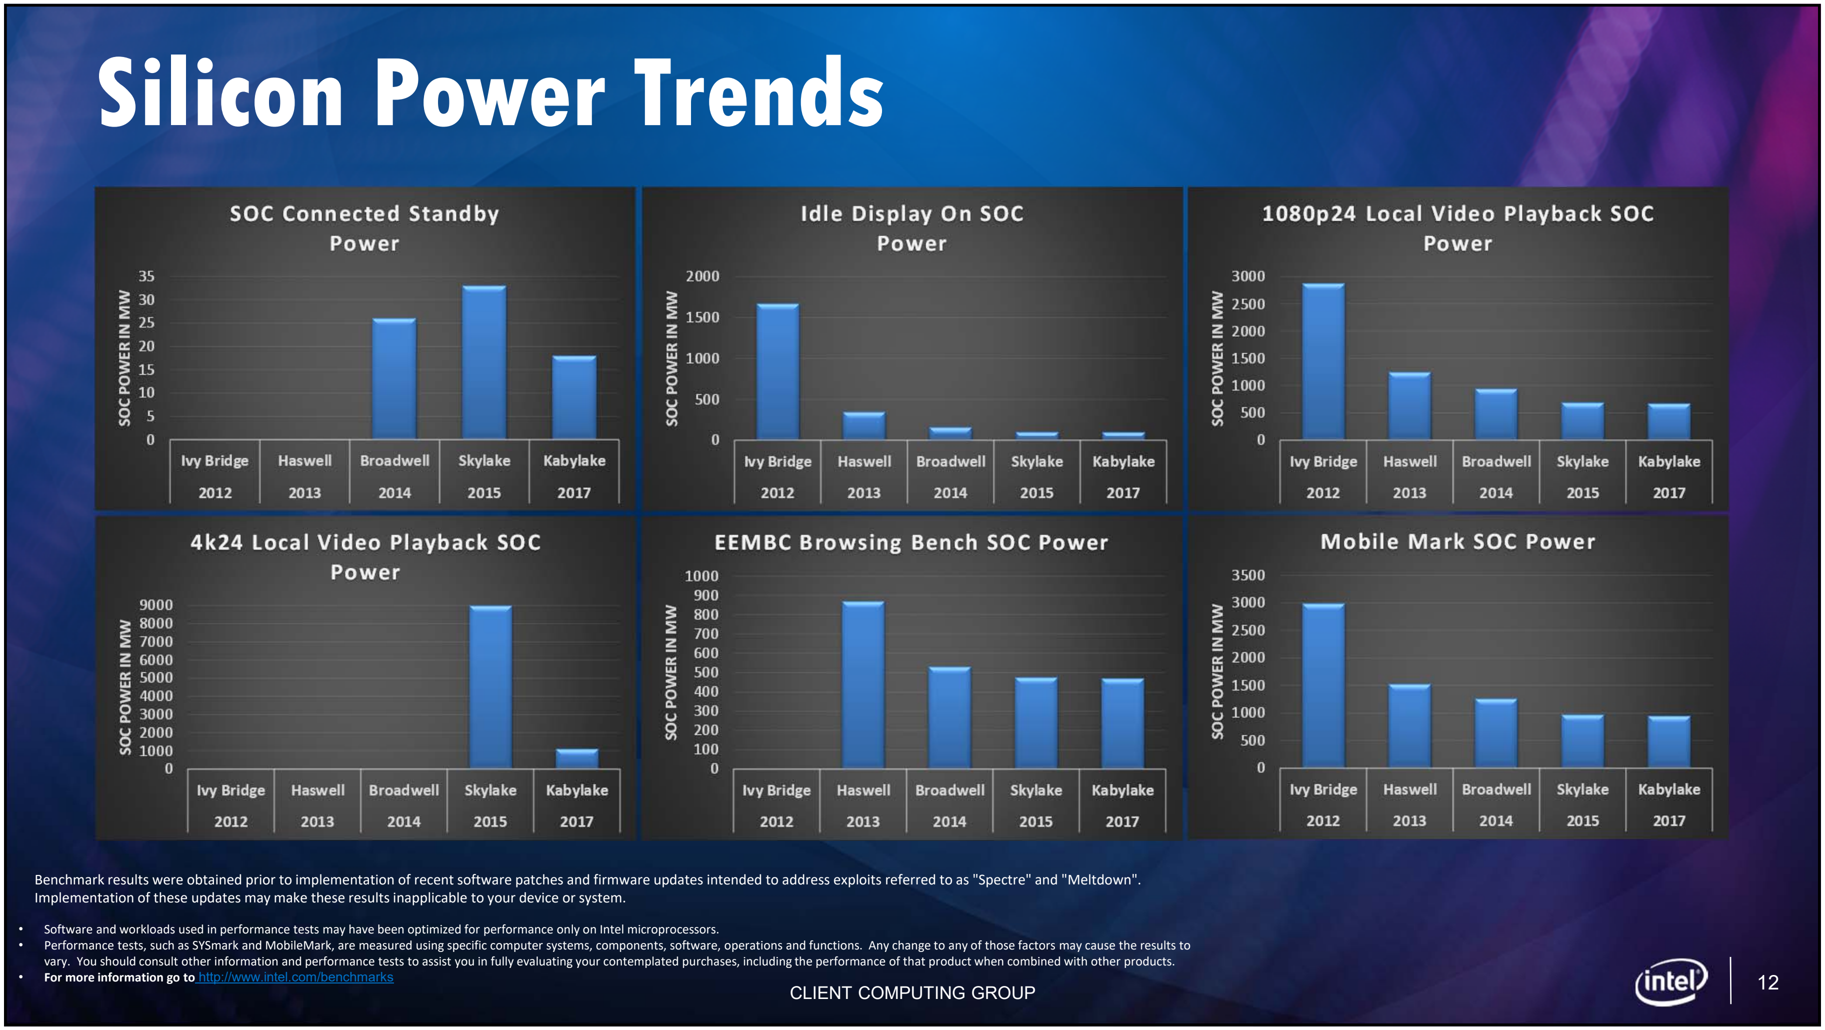
Silicon Power Trends
SOC Connected Standby Power
Idle Display On SOC Power
1080p24 Local Video Playback SOC Power
4k24 Local Video Playback SOC Power
EEMBC Browsing Bench SOC Power
Mobile Mark SOC Power
SOC POWER IN MW
Ivy Bridge 2012
Haswell 2013
Broadwell 2014
Skylake 2015
Kabylake 2017
Benchmark results were obtained prior to implementation of recent software patches and firmware updates intended to address exploits referred to as "Spectre" and "Meltdown". Implementation of these updates may make these results inapplicable to your device or system.
Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors.
Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products.
For more information go to http://www.intel.com/benchmarks
CLIENT COMPUTING GROUP
intel
12

Platform PM Trend

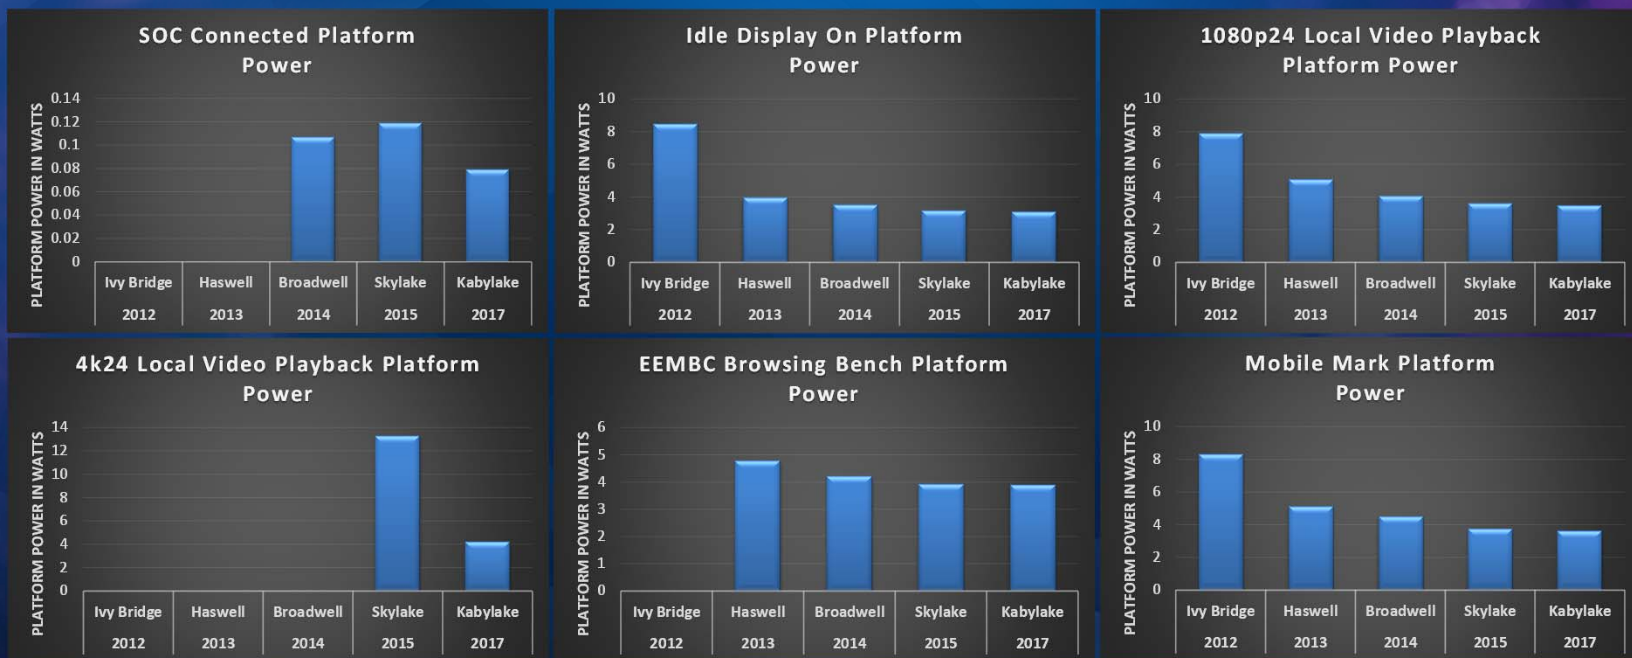

Benchmark results were obtained prior to implementation of recent software patches and firmware updates intended to address exploits referred to as "Spectre" and "Meltdown". Implementation of these updates may make these results inapplicable to your device or system.

- Software and workloads used in performance tests may have been optimized for performance only on Intel microprocessors.
- Performance tests, such as SYSmark and MobileMark, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products.
- **For more information go to** http://www.intel.com/benchmarks

CLIENT COMPUTING GROUP

intel 13

Changing How We Sleep

- The modern standby PC user experience is designed to model that of a cellular phone.
- A large benefit to the modern standby PC experience is instant resume, which is significantly faster then today's designs with S3
- Modern standby also enables PCs to stay fresh and up-to-date by maintaining connectivity and controlling the execution of background tasks
  - But, Modern Standby is also a building block to Modernize PCs and allow innovation of ambient capabilities that are sensitive and responsive to the presence of people and other devices
- Modern PCs are context aware, recognizing you and your situational context, they are tailored to meet your needs and learn to anticipate your desires
- All Intel PCs are moving to Modern and will be Always On, Always Connected, and Always Available in future
- Modern Standby has been extremely difficult to build, enable, deploy, and keep stable
- Microsoft and Intel are investing heavily in making this transition happen, but it still require large industry investment

CLIENT COMPUTING GROUP

intel | 14

Modern Standby Benefits

		Legacy Standby	Disconnected Standby	Connected Standby
Wake Sources	Wake on Voice (WoV)	✗	✓	✓
	Wake on "Miracast to PC"	✗	✗	✓
	Wake on Automatic Wireless Docking	✗	✓	✓
	Wake on Finger Print Reader	✗	✓	✓
Performance & Security	Faster resume from standby	✗	✓	✓
	Consumer Device Encryption	✗	✓	✓
	Windows is up-to-date	✗	✗	✓
Entertainment & Staying Connected	Uninterrupted Casting & Low Power Audio Playback	✗	✓	✓
	Receive Skype Calls & IM's	✗	✗	✓
	Email & Calendar are up-to-date	✗	✗	✓
	Receive push notifications (E.g., Social or IoT)	✗	✗	✓
	Sync with Bluetooth Devices (E.g., Fitness watch)	✗	✓	✓

Modern Standby Benefits

Always On
- Wake-on scenarios
- Notifications
- Background Tasks (i.e. Music/Skype*)

Always Connected
- OS / Apps Downloads
- Fresh Data

Always Available
- 3X faster resume
- Energy Efficient Ambient Compute
- 14+ days on Standby

Intel® Dynamic Tuning

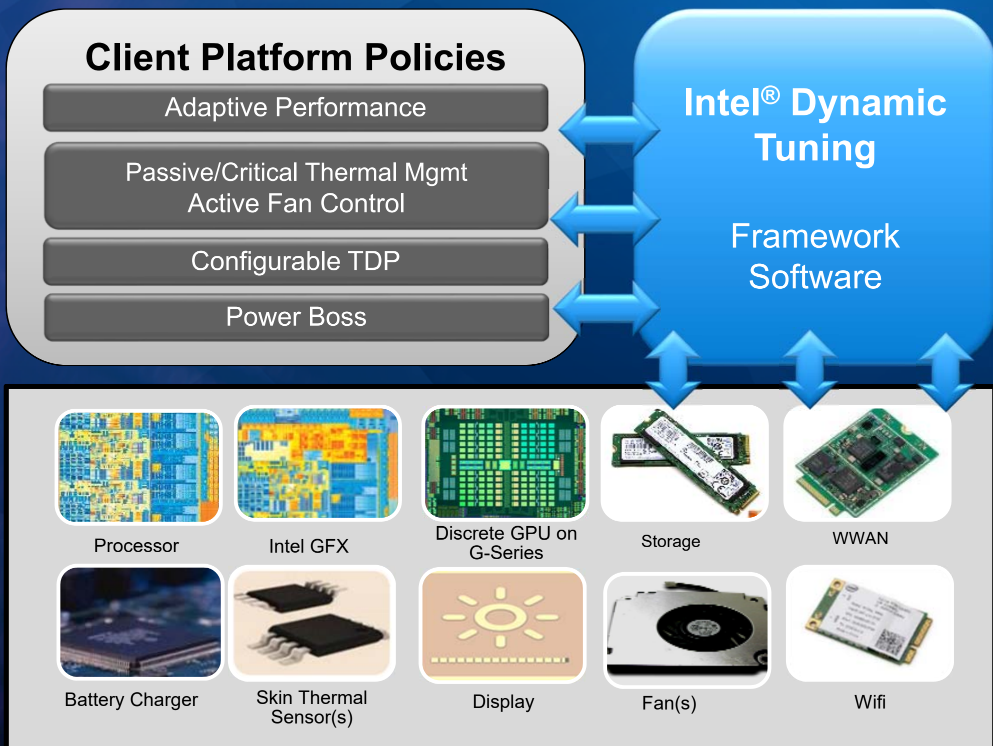

Intel Dynamic Tuning is a uniform thermal management approach across Intel platforms

- ✓ Manage combined Power & Thermal Budget
- ✓ Intelligently balance CPU & GPU power budget based on performance need
- ✓ Power manage multitude of hardware devices, including discrete graphics*
- ✓ Easily tune performance in various system states (e.g. lap, tablet or tent mode etc.)
- ✓ OEM value add customization via configuration/tuning tables

* Intel G-series discrete graphics only

CLIENT COMPUTING GROUP 17 intel 17

Industry Challenges

Power and performance within a given thermal envelope, as integration of high perf IPs continues (and thermal density increases)

Silicon scaling is challenging

- Core performance may be scaled, but surrounding infrastructure is hard to scale
- Hetero cores help somewhat, but often comes with an overhead

Display I/O and graphics compute requirements exponentially increasing

- Raw: 4k@60Hz(24bpp) = ~14Gbps; 8k@60Hz(24bpp) = ~50Gbps
- Martialing around raw data is very power hungry

Modern Standby is still hard to enable

- We are working hard on all the top challenges debug, IHV compliance testing, ecosystem enabling, power delivery, and OS features to ease stability and compliance

Challenges Related to PNP Test

Measuring sub milliwatts Power in a Modern-Standby Sessions, in Validation and HVM Testing

- Many components (SSD, EC, WIFI, SOC, etc.) expected to consume few milli-Watts and in some cases sub-mW
- Well within range of measurement equipment accuracy, especially in automatic test equipment environment

Measuring Real Benchmark Performance in HVM

- Benchmarks are typically 20 minutes benchmarks but need to be measured in milli-seconds in HVM

Running OS-based workloads (for power, performance, and power management) in pre-silicon environment, simulation and emulation

- Emulation speed is less than 1MHZ and RTL simulators are slower – Need 1000x speedup to run OS
- Need proxies that very accurately represent actual benchmarks
- Shorter version of the OS that loads what is only needed to run the benchmark

Temperature and Voltage Control

- Temperature accuracy is very key to performance, especially in burst mode; need to control temperature very accurately in a temperature range From -40C to +125C

Power Delivery and Droop Management

- Need to ensure test environment produces same droops as intended system
- Impact of fast droops on power and performance very difficult to characterize and test in HVM

Summary

As an industry, we have made computing dramatically more efficient continuously, year over year, with some remarkable breakthroughs

- This continues every generation

The challenges that remain are extremely difficult

- Delivering peak performance and best possible battery life with the same silicon
- Driving to extremely low idle powers and testing for them in HVM
- Added cost for energy savings continues to be an end user aversion, unless mandated
  - 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
- Modern Standby will bring mobile phone attributes and capabilities to the PC, but it is a very challenging transition

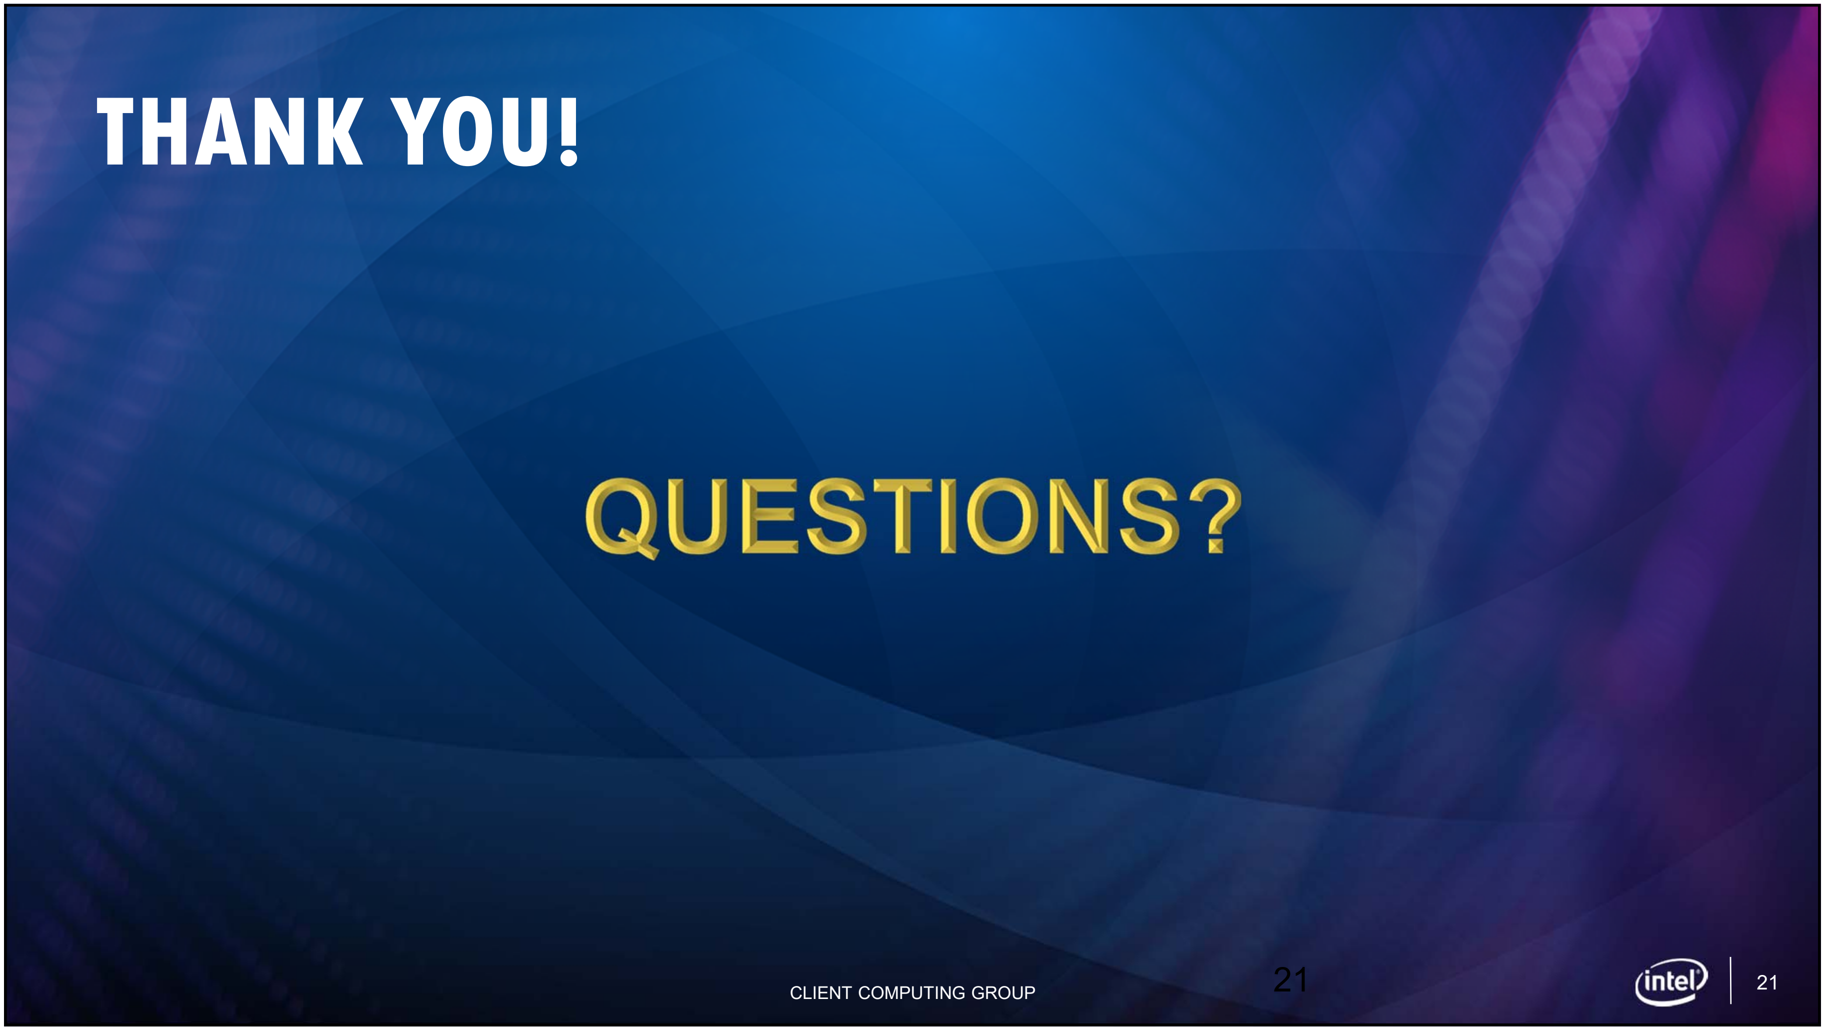
THANK YOU!
QUESTIONS?
CLIENT COMPUTING GROUP
21
intel
21

Legal Disclaimer

Intel technologies may require enabled hardware, specific software, or services activation. Check with your system manufacturer or retailer.

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/benchmarks

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

Performance tests, such as SYSmark* and MobileMark*, are measured using specific computer systems, components, software, operations and functions. Any change to any of those factors may cause the results to vary. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases, including the performance of that product when combined with other products. **For more information go to** http://www.intel.com/benchmarks

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. You should consult other information and performance tests to assist you in fully evaluating your contemplated purchases.

Benchmark results were obtained prior to implementation of recent software patches and firmware updates intended to address exploits referred to as "Spectre" and "Meltdown". Implementation of these updates may make these results inapplicable to your device or system.

(for data marked '(e)'): Results have been estimated or simulated using internal Intel analysis or architecture simulation or modeling, and provided to you for informational purposes. Any differences in your system hardware, software or configuration may affect your actual performance.

Measured Power Data Configurations

	Intel® Core™ i7-3689Y, 3rd Gen ULT i7+ UM77	Intel® Core™ i5-4302Y cTDP Down, PL1=4.5W	Intel® Core™ M-5Y71 QS TDP = 4.5W	Intel® Core™ M7-6Y75, PL1 = 4.5W	KBL-Y 2+2 4.5W
Memory	Memory: DDR3-1600, CL11, 2x4GB	LPDDR3; 2X2GB @1600MHz	LPDDR3; 2X2GB @ 1600MHz	LPDDR3; 2X2GB @ 1600MHz	LPDDR3; 2X2GB @ 1600MHz
Panel	Panel: eDP 13.3" 19x10	eDP 11.6" 19x10	eDP CMI 11.6" 19x10	12.5" 25X14 SHARP* eDP Panel (LQ125T1JX03)	12.5" 25X14 SHARP eDP Panel (LQ125T1JX03)
Storage	Storage: Intel SSD 320series 160GB	Samsung 120GB, NGFF, PM841	Samsung 128GB, NGFF, VS951 PCIe	Intel Temple Star 180GB SSDSCKGW180A4	MZVPV128HDGM_FW-BXW7000Q
WLAN + BT		Broadcom -BCM43241- NGFF - SDIO/UART	PCIe Stone Peak C0	PCIe Snow Field Peak C0	PCIe Snow Field Peak C0
Touch		Synaptics*: I2C+NFC: TM-02399-001; I2C (NO NFC antenna): TM-02400-001	Atmel I2C	Elentec I2C	Elentec I2C
Camera		Chicony* CKFCF01	1)UF USB2 Cam = Chicony CKFDF01 2)WF USB3 Cam = Chicony CKAD802	1)IVCAM USB Camera 2)SKYCAM MIPI CSI2	1)IVCAM USB Camera
Sensors		ST Micro*	ST Micro (STM32L151RBT6)	Integrated Sensor Hub (ISH)	Integrated Sensor Hub (ISH)
TPM		fTPM only	SLB9665T2.0 (TPM 2.0)	SLB9665T2.0 (TPM 2.0)	SLB9665T2.0 (TPM 2.0)
Audio		Realtek*: ALC5505 (External DSP) + ALC 282 (HD codec)	Realtek*: ALC286S	Realtek*: ALC298	Realtek: ALC298
RTD3		Enabled	Enabled	Enabled	Enabled
OS	OS: Win8 Pro Build 9200	Win 8.1 Update	Win 8.1 Update	Win 10 TH2	Win 10 TH2

http://competition.intel.com/Default.aspx?tabid=2255&DMXModule=6079&EntryId=10641&Command=Core_Download
http://competition.intel.com/Default.aspx?tabid=2255&DMXModule=6079&EntryId=12769&Command=Core_Download
http://competition.intel.com/Default.aspx?tabid=2255&DMXModule=6079&EntryId=14016&Command=Core_Download
http://competition.intel.com/Default.aspx?tabid=2254&DMXModule=6070&EntryId=14860&Command=Core_Download