Thin & Light & High Performance Graphics

Srinivas Chennupaty

Intel Corporation, 2018

Notices and Disclaimers

This document contains information on products, services and/or processes in development. All information provided here is subject to change without notice. Contact your Intel representative to obtain the latest forecast, schedule, specifications and roadmaps.

Intel technologies' features and benefits depend on system configuration and may require enabled hardware, software or service activation. Learn more at intel.com, or from the OEM or retailer.

No computer system can be absolutely secure.

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 complete information visit http://www.intel.com/performance.

Optimization Notice: Intel's compilers may or may not optimize to the same degree for non-Intel microprocessors for optimizations that are not unique to Intel microprocessors. These optimizations include SSE2, SSE3, and SSSE3 instruction sets and other optimizations. Intel does not guarantee the availability, functionality, or effectiveness of any optimization on microprocessors not manufactured by Intel. Microprocessor-dependent optimizations in this product are intended for use with Intel microprocessors. Certain optimizations not specific to Intel microarchitecture are reserved for Intel microprocessors. Please refer to the applicable product User and Reference Guides for more information regarding the specific instruction sets covered by this notice.

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.

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.

Cost reduction scenarios described are intended as examples of how a given Intel-based product, in the specified circumstances and configurations, may affect future costs and provide cost savings. Circumstances will vary. Intel does not guarantee any costs or cost reduction.

Intel does not control or audit third-party benchmark data or the web sites referenced in this document. You should visit the referenced web site and confirm whether referenced data are accurate.

Intel, the Intel logo, Intel Optane and Xeon are trademarks or registered trademarks of Intel Corporation or its subsidiaries in the united states and other countries.

Problem Statement

Consumers are faced with a stark choice

- Use CPU Integrated Graphics(IG) offer good base-line performance, features, battery life, form-factor along with a programmer friendly unified memory model, but do not satisfy the needs of high-performance graphics applications

OR

- Use external Graphics(EG) for greater graphics performance but suffer from lower battery life, larger form-factor and programmer un-friendly distributed memory model

How do we enable both?

- Package level integration of IG+EG using high bandwidth coherent interconnect to drive smaller form factor, higher performance/watt, longer battery life , better memory and programming model for developers and overall reduction in graphics solution cost

8th Gen Intel® Core™ with Radeon™ RX Vega M

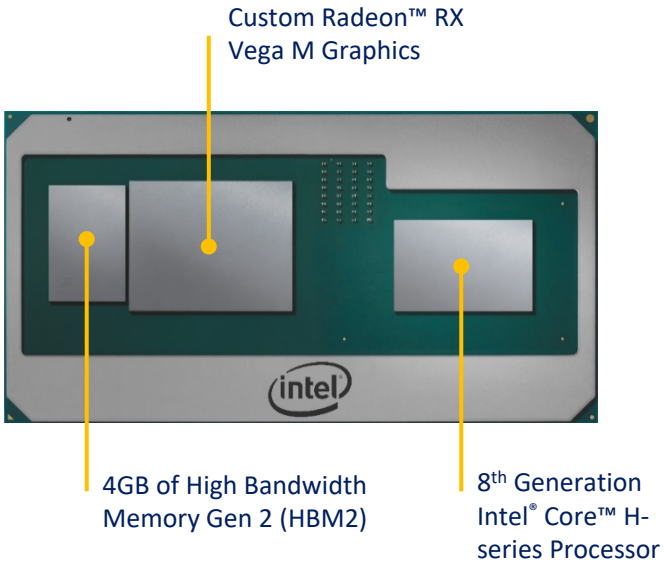

Kaby Lake Processor

- 4-core 8-thread, 3.1GHz base clock, turbo up to 4.2GHz
- 8MB cache w/ 2 channels x DDR4-2400
- Intel HD graphics 630: 24EUs up to 1100MHz
- Overclocking SKU available

Radeon RX Vega M Graphics

- GL SKU - 65W Package TDP
  - 20 compute units
  - Base/boost clock: 931/1011MHz, 2.6TFLOPS
  - ROPs: 32 pix/clk
- GH SKU - 100W Package TDP
  - 24 compute units
  - Base/boost clock: 1063/1190MHz, 3.7TFLOPS
  - ROPs 64 pix/clk

High Bandwidth Memory

- 4GB capacity
- 1.4Gbps (GL) and 1.6Gbps (GH) via 1024-bit interface

Two Graphics Subsystems on One Small Package

High Bandwidth Cache Controller

High Bandwidth Cache
- 4GB @ 180GB/s or 205GB/s
- Low power

Quad Geometry Engines

Vega Pixel Engine
- Up to 16 Render Back Ends
- Up to 64 Pixels/Clock

Compute Units
- Up to 24 Compute Units
- Asynchronous Dispatch
- Per Compute Unit Power Gating
- Vulkan® & DirectX™ 12 Ready
- Supports Radeon Shader Intrinsics

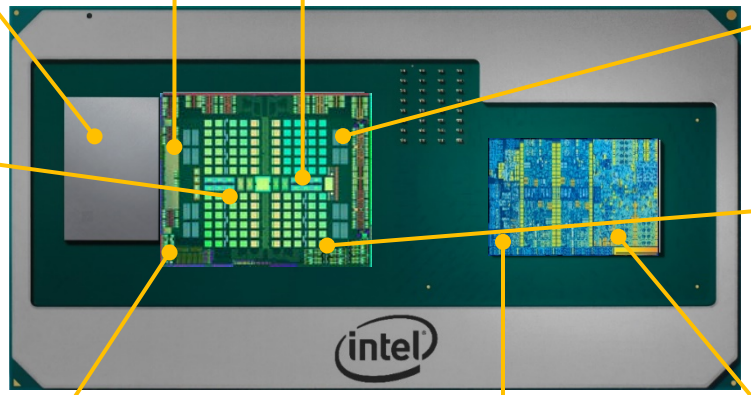

Radeon™ Multimedia Engine
- 4K60 encode / decode with Radeon ReLive
- HEVC, H264 HDR enc/dec

Intel® Quick Sync Video
- VP9 & HEVC 10b HW enc/dec
- H264 HW enc/dec

Displa

Radeon™ Display Engine
- 6 Displays
- Up to 4K resolution
- Display Port 1.4 w/ HDR
- HDMI 2.0b with HDR10 support

Intel® Gfx Display Engine
- 3 Displays
- 4K resolution
- eDP /PSR for long battery life

How did we make this happen?

Key Enabling Intel Technologies

Embedded Multi-Die Interconnect Bridge (EMIB)

- Low Cost high density 2.5D interconnect

Dynamic Platform Thermal Framework

- Platform level thermal mgmt.

Heterogeneous Integration Options

Multi-Chip Package

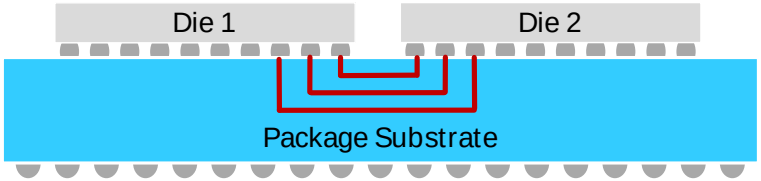

- Poor density of die-package connections
- Poor density of die-die interconnects

Silicon Interposer

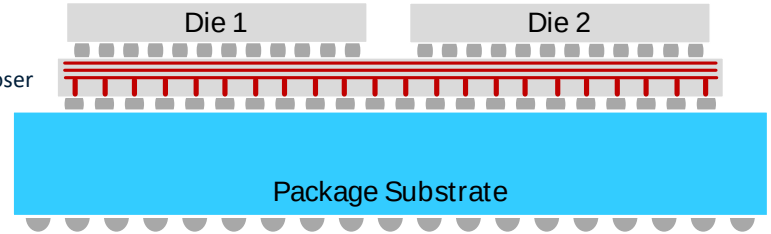

- Good density of die-interposer connections
- Good density of die-die interconnects
- Higher cost of large interposer + thru-silicon vias

Embedded Multi-Die Interconnect Bridge

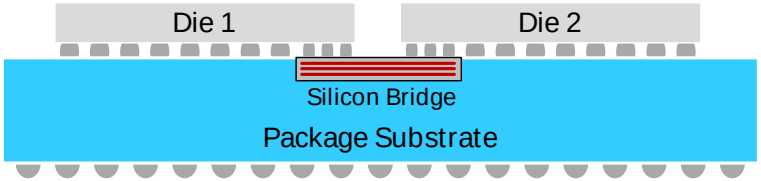

- Good density of die-bridge connections
- Good density of die-die interconnects
- Low cost of small silicon bridges

EMIB technology provides high density, high bandwidth die-die interconnects

Intel EMIB Packaging

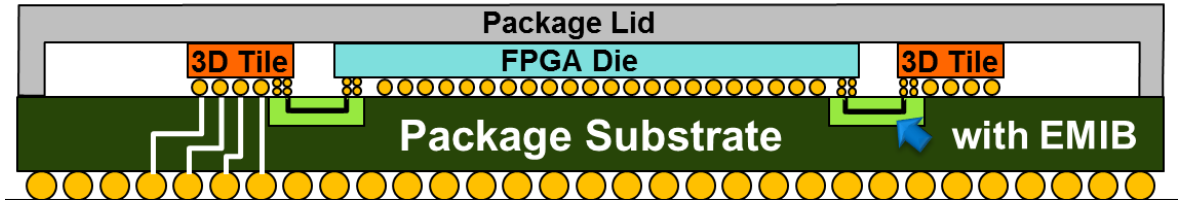

Innovative, Simpler, Higher-Performance Solution
Even for parts from different fabs, process nodes and vendors!

Note: Drawings are conceptual and NOT drawn to scale

Connecting the two components

Highly constrained problem and many unique challenges

PCI-e routing – off package repurposed to on-pkg

Z-height challenges – required custom thinned HBM devices

Engineering and production test flows, while protecting critical IP

Supply chain enabling across **3 GEOs** from FAB to Assembly

Develop common engineering & production test flows, yet protect critical IP in both organizations

Key Enabling Intel Technologies

Embedded Multi-Die Interconnect Bridge (EMIB)

- Low Cost high density 2.5D interconnect

Dynamic Platform Thermal Framework

- Platform level thermal management

Platform Level Power Management

Platform Power Management controls user experience

OEMs design to System Design Point (SDP), not combined TDP

Other factors which affect mobile performance

- Static vs dynamic power allocation
- Skin temperature management
- AC power availability

Low-latency response to workload variations maximizes performance

System Optimized Thermal Management:

Platform Power Sharing for Optimal Performance

Intel® Dynamic Platform and Thermal Framework

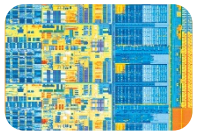

Processor
- Temperature
- Power Control
- P/T States

Processor Graphics
- Temperature
- Power Control
- RP States, EU

PCH
- Temperature
- Power Control

Memory
- Temperature
- Power Control

WLAN, WWAN
- Temperature
- Power Control

Battery Charger
- Charge Rate Control

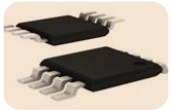

Skin Thermal Sensor(s)
- Temperature

Display
- Brightness Control

System Fan(s)
- Fine Grained Fan Control

Intel® DPTF – Active Skin Temperature Management

Monitor platform constraints

Modulate system and SoC parameters to operate within constraints

- e.g. Monitor skin temperature dynamically adjust PL1/PL2

DPTF can delivery upto 30% performance increase on cold systems

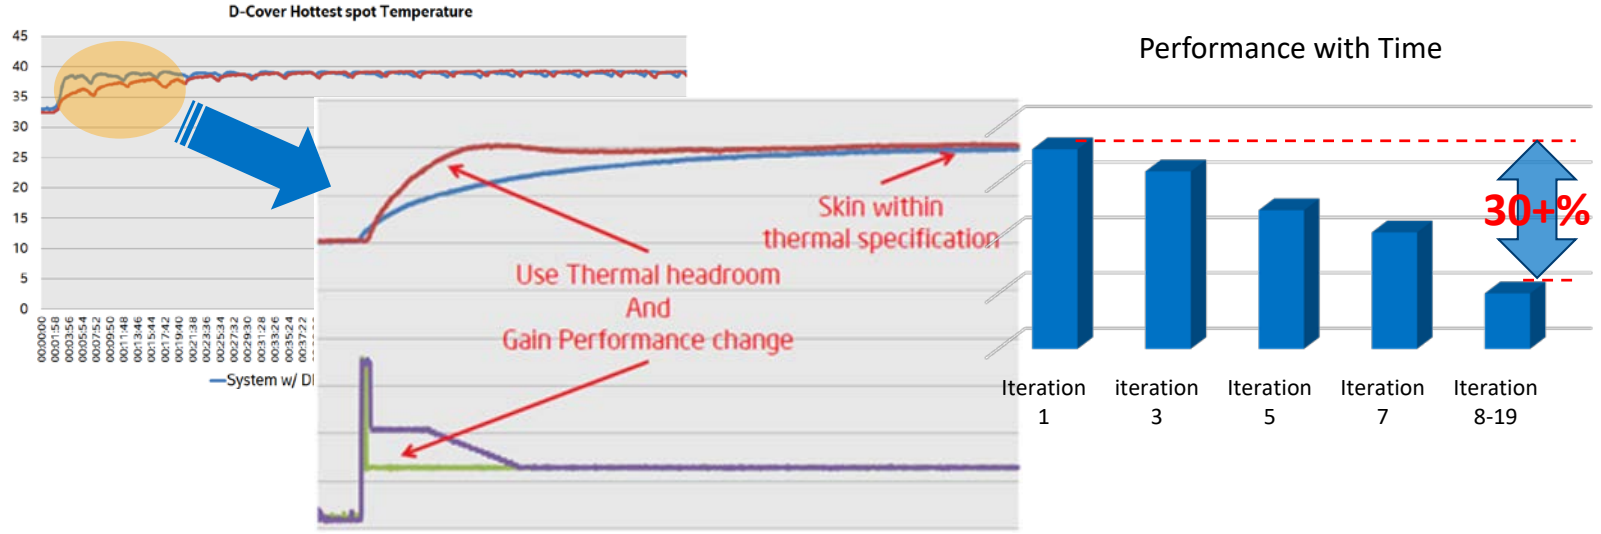

DPTF with discrete graphics

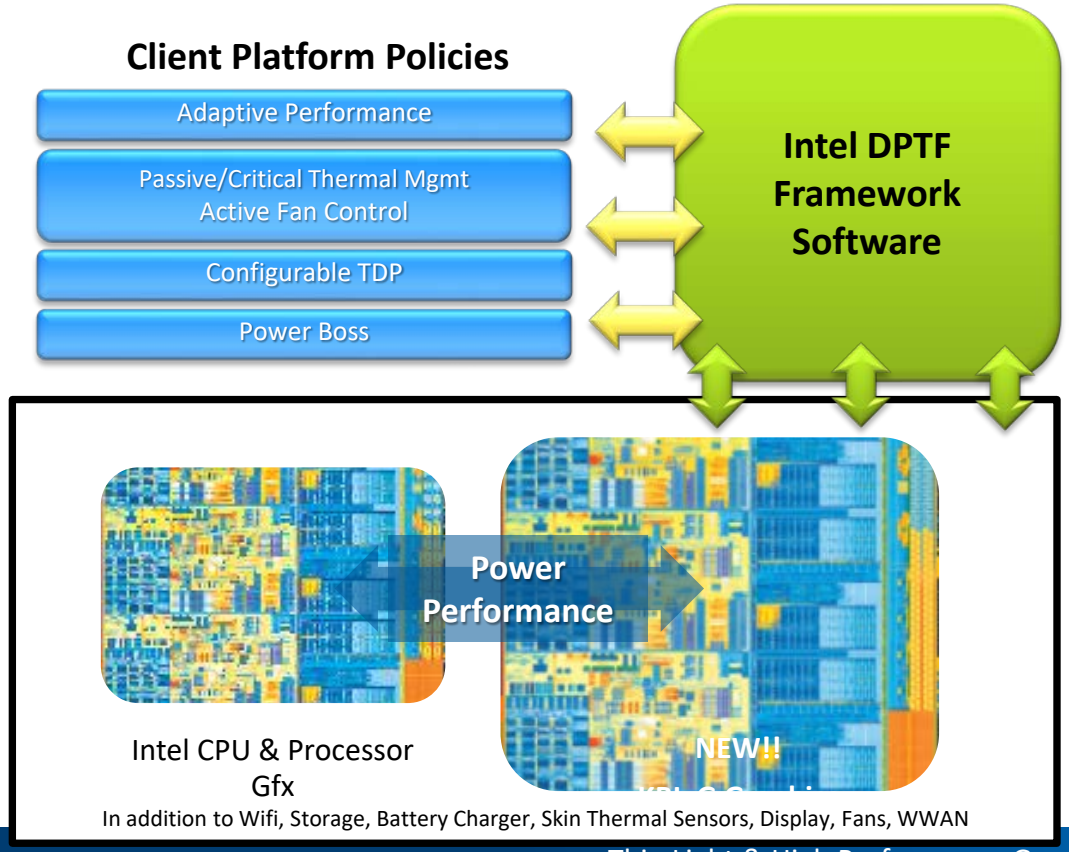

- X Current CPU & dGPU power management is rudimentary and difficult to replicate, tune and update
- ✓ DPTF is a uniform Thermal Mgmt. approach across all platforms
- ✓ Customizable via configuration/tuning tables
- ✓ Manage combined Power & Thermal Budget
- ✓ Intelligently balance CPU & GPU power budget based on performance need.

Power Sharing Control

DPTF w/ Power Sharing Policy – Manages Combined MCP Power

Split “SOC” (KBL-H + dGFX) into two power domains

- CPU participant: KBL-H die
- dGFX participant: dGPU die + HBM

Each domain is provided a power target to meet total MCP Budget

- Participants autonomously manage their individual budgets
- ~100ms control loop bandwidth

Power Sharing algorithm

- PID controller to track and manage overall MCP budget
- Budget available in the PID controller decides MCP power headroom (TDP over next evaluation interval)
- Utilization from each Participant (CPU and dGPU) and from “Platform BIAS” decide how that budget is divided between the 2 participants each polling loop
- Allows “turbo” similar to Intel Turbo Boost 2.0 Technology

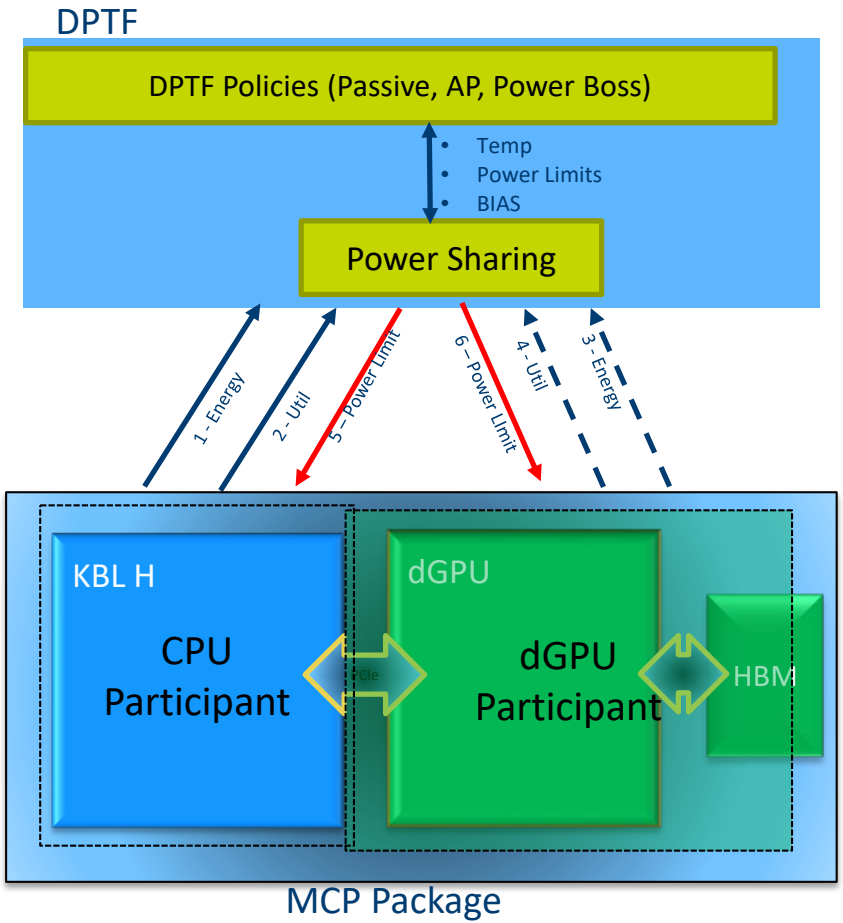

Bringing It together

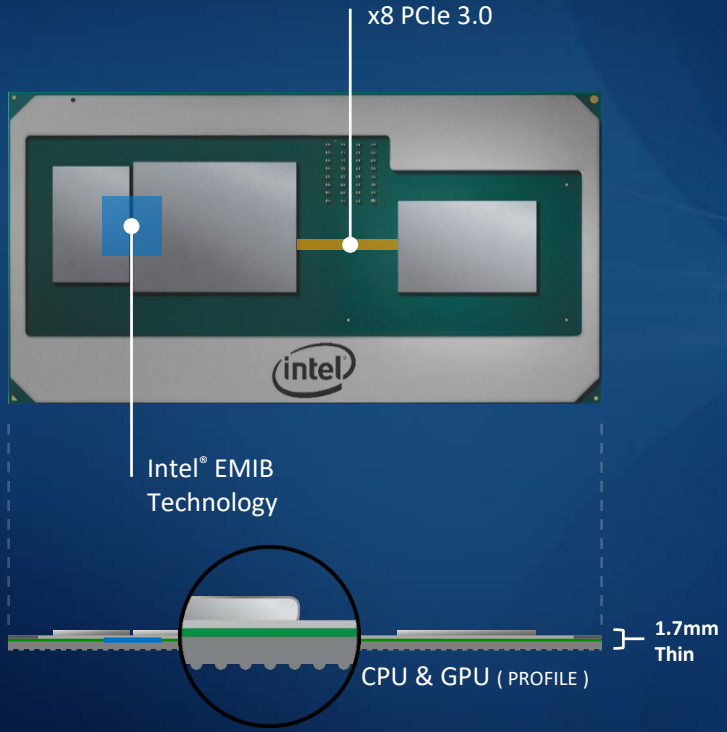

8th Gen Intel® Core™ processors
With Radeon™ RX Vega M Graphics

Smaller, thinner solution through Intel EMIB

- Embedded high speed connector in package
- Reduced silicon footprint over 50%[4]
- Keeps CPU and GPU z-height 1.7mm slim

Enthusiast processor adds needed connectivity

- Eight lanes of PCI Express Gen 3 connecting CPU & GPU
- Provides necessary throughput to feed intense gfx workloads
- Remaining PCIe lanes available for direct CPU access

Hardware Features

- Efficient HBM, up to 80% less power than GDDR5[5]
- Intel® Graphics efficient display and Quick Sync Video capabilities available
- 9 Display outputs available for design flexibility

Design Flexibility through innovation

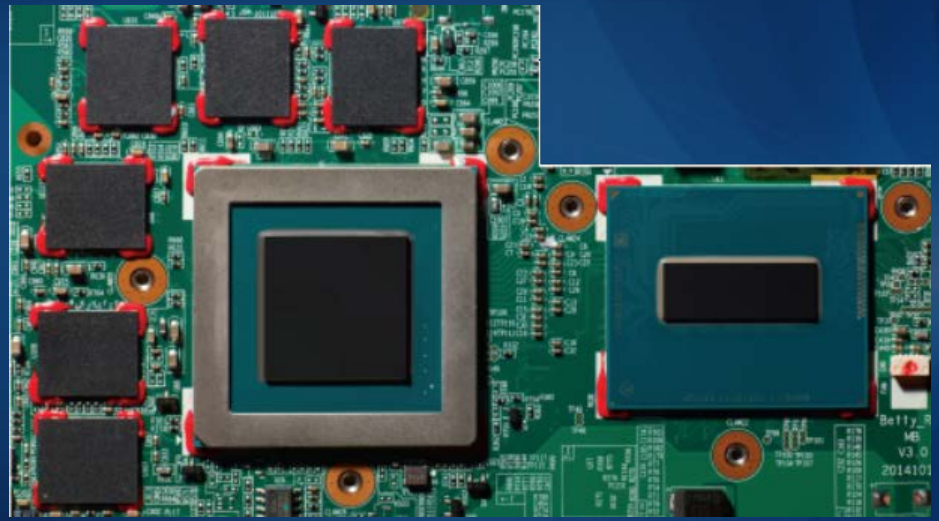

Typical Enthusiast Motherboard Design
CPU + GPU + GDDR5

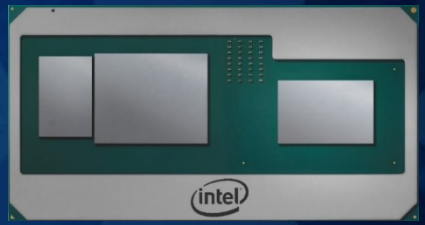

8th Gen Intel® Core™ Processor

Images are shown to scale

1900mm² (3in²) board space savings

Thinner Designs Thru Dynamic Power Sharing

Traditional Platform

OEMs design to System Design Point (SDP), not combined TDP

62.5 W SDP

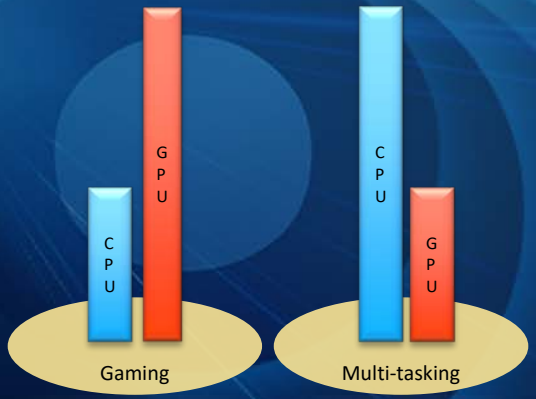

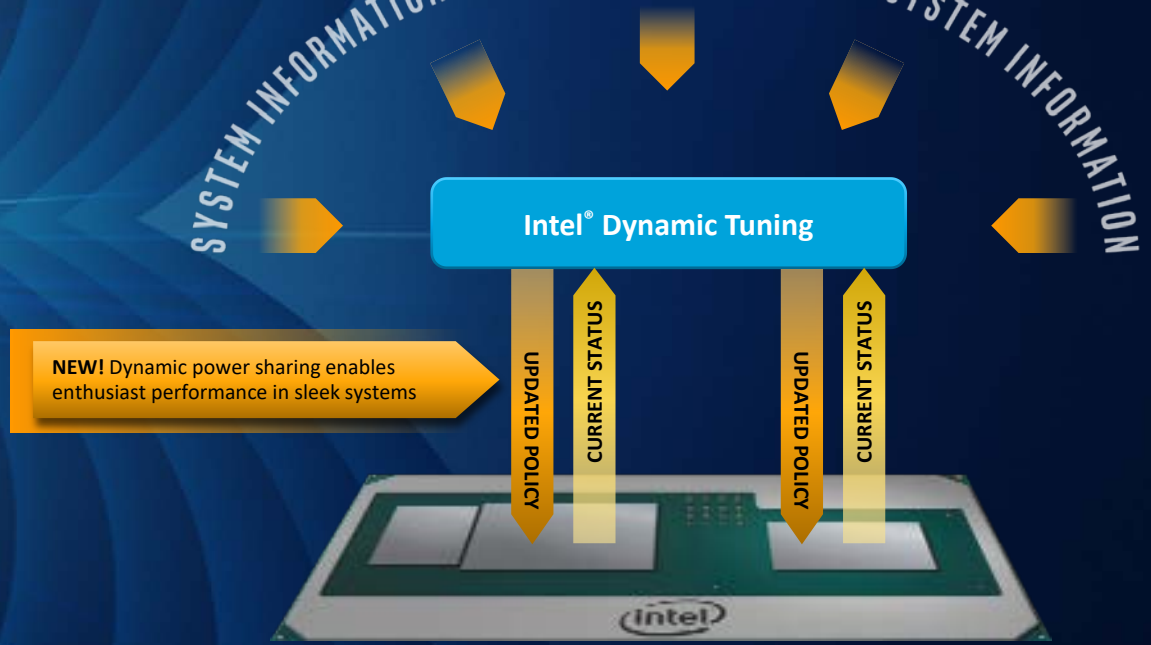

Thinner Designs Thru Dynamic Power Sharing

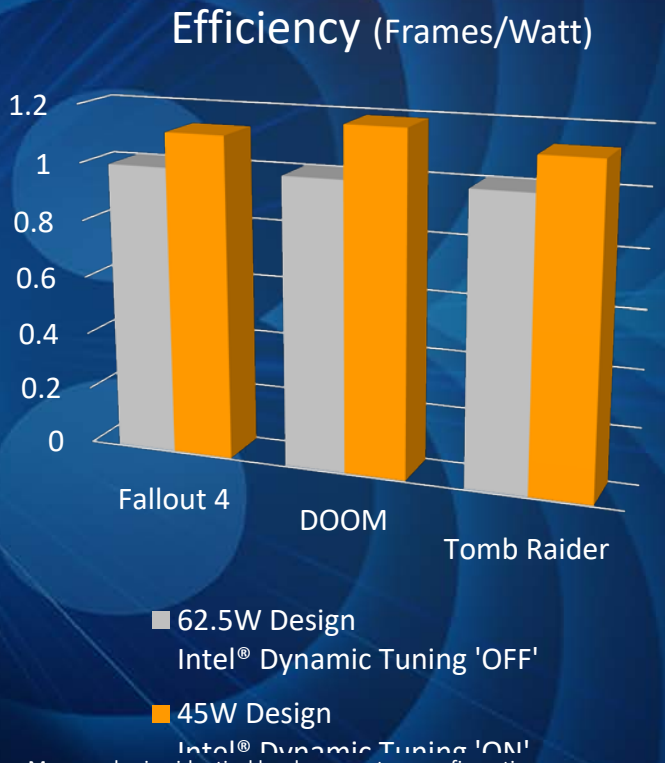

Measured using identical hardware system configuration.

Up front design benefit of 17.5W

Same performance with up to 18% higher efficiency*

* Intel® Dynamic Tuning as measured on Intel Reference Platform: 8th Gen: Intel® Core™ i7-8705G Processor, 4C8T, Turbo up to 4.1GHz, Memory: 16GB, Storage: SSD, Graphics: Radeon* RX Vega M GL, OS: Windows* 10. Power Sharing "ON" at 45W package power. Power Sharing "OFF" at CPU PL1: 45W, GPU 40W TGP

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 complete information visit www.intel.com/benchmarks.

8th gen Intel® Core™ Processor with radeon™ RX Vega M GL Graphics

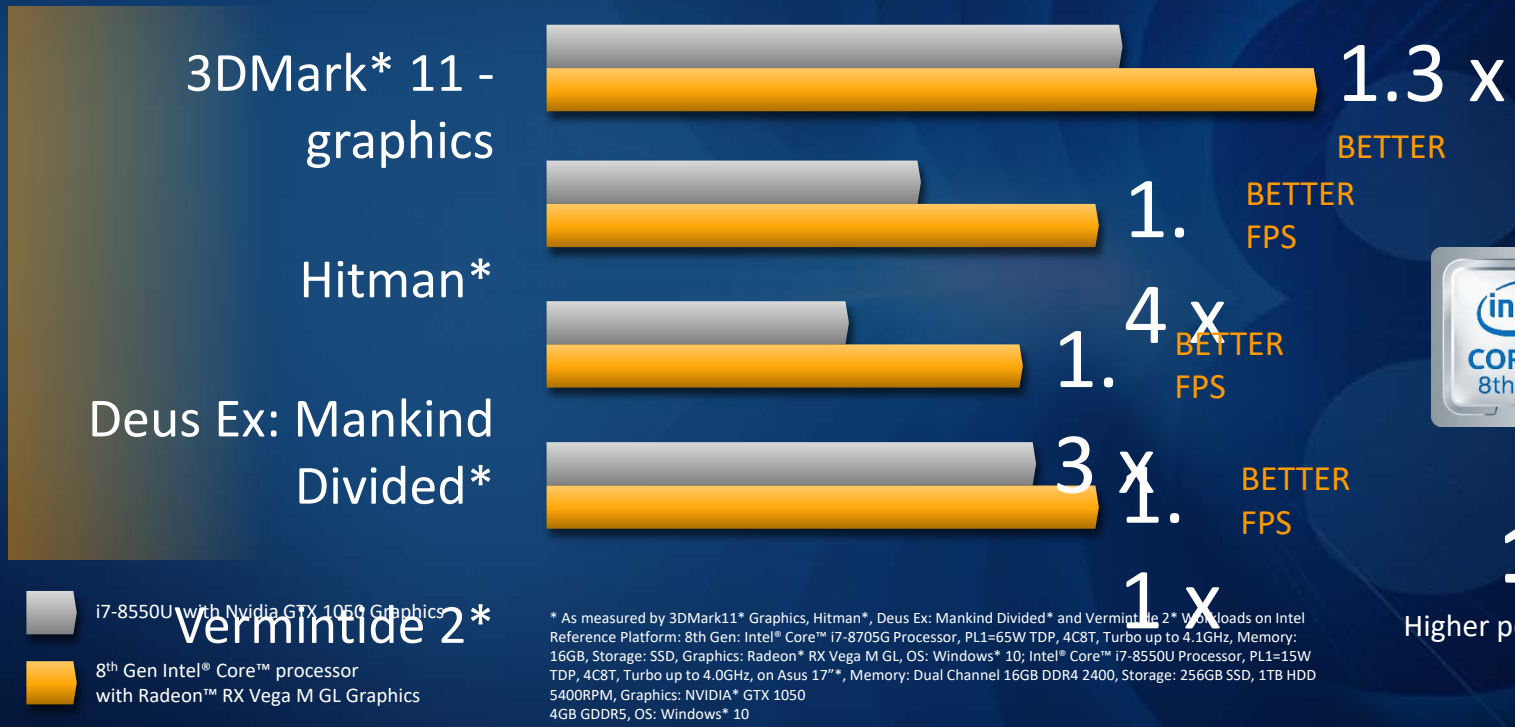

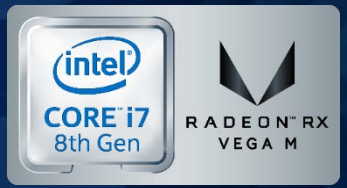

Up to

1.4x*

Higher performance than GTX 1050 (4GB)

* As measured by 3DMark11* Graphics, Hitman*, Deus Ex: Mankind Divided* and Vermintide 2* Workloads on Intel Reference Platform: 8th Gen: Intel® Core™ i7-8705G Processor, PL1=65W TDP, 4C8T, Turbo up to 4.1GHz, Memory: 16GB, Storage: SSD, Graphics: Radeon* RX Vega M GL, OS: Windows* 10; Intel® Core™ i7-8550U Processor, PL1=15W TDP, 4C8T, Turbo up to 4.0GHz, on Asus 17"*, Memory: Dual Channel 16GB DDR4 2400, Storage: 256GB SSD, 1TB HDD 5400RPM, Graphics: NVIDIA* GTX 1050 4GB GDDR5, OS: Windows* 10

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 complete information visit www.intel.com/benchmarks.

Conclusion

Packaging is a critical tool to build interesting new products

Embedded Multi-Die Interconnect Bridge (EMIB) – Flexible way to build heterogeneous products rapidly

Platform level Thermal/Power Management maximizes performance