



# 4<sup>th</sup> Generation Intel® Core™ Processor, codenamed Haswell

Per Hammarlund

Haswell Chief Architect, Intel Fellow

August, 2013



# 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 PRODUCTS ARE NOT INTENDED FOR USE IN MEDICAL, LIFE SAVING, OR LIFE SUSTAINING APPLICATIONS.

Intel may make changes to specifications and product descriptions at any time, without notice.

All products, dates, and figures specified are preliminary based on current expectations, and are subject to change without notice.

Intel, processors, chipsets, and desktop boards 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.

**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 purchase, including the performance of that product when combined with other products.**

Intel, Core i7, Core i5, Core i3, Ultrabook, and the Intel logo are trademarks of Intel Corporation in the United States and other countries.

\* Other names and brands may be claimed as the property of others.

Copyright © 2013 Intel Corporation.



# Sequence

## Family of Innovations!

Power Efficiency and Management

FIVR – Fully Integrated Voltage Regulator

Cache Hierarchy and Interconnects

Gfx/Media

Intel® Microarchitecture (Haswell): Core

ISA

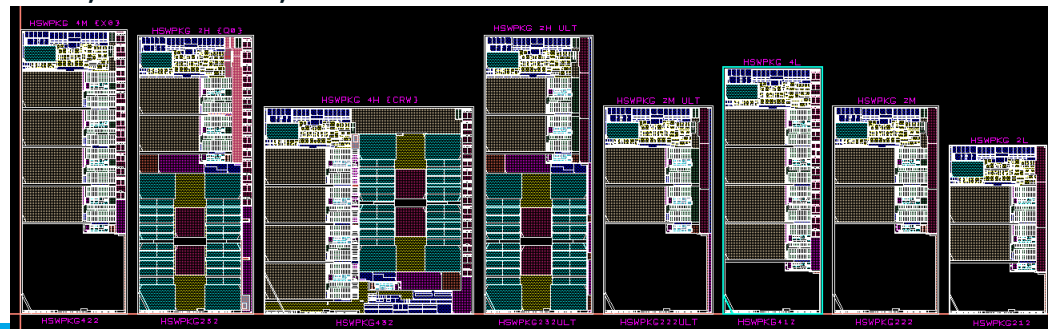
Wrap Up

# Family of Innovations

- **Huge family:** SOC methodology, common architecture
- **Low power platform:** 20x idle power reduction, low power IO (I2C, SDIO, I2S, UART), Link power management (USB, PCIe, SATA)
- **Large eDRAM Cache**
- **Platform:** PSR (Panel Self Refresh)
- **FIVR:** Fully Integrated Voltage Regulator
- **Core:** FMA (Floating-point Multiply Add), 2x Cache BW, TSX (Transaction Synchronization Extension)
- **Graphics:** 2x in Ultrabooks, OpenCL 1.2, DX 11.1, OpenGL 4.0
- **Media:** 5x faster at 0.5x power

## Modularity Options

Value	Range
Core Count	2-4
Graphics	GT1, GT2, GT3
Active Power Level	Tablet to Desktop
Idle Power	Variable
Cache Size	Variable
Interconnects	Variable
Platforms	Traditional, power optimized




# Intel Process 22nm Process Technology and Tick/Tock Development Model

45nm Process Technology		32nm Process Technology		22nm Process Technology	
<b>Nehalem</b>	<b>Westmere</b>	<b>Sandy Bridge</b>	<b>Ivy Bridge</b>	<b>Haswell</b>	
NEW Intel® Microarchitecture (Nehalem)	Intel Microarchitecture (Nehalem)	NEW Intel Microarchitecture (Sandy Bridge)	Intel Microarchitecture (Sandy Bridge)	NEW Intel Microarchitecture (Haswell)	
<i>TOCK</i>	<i>TICK</i>	<i>TOCK</i>	<i>TICK</i>	<i>TOCK</i>	

## Enhanced version of Intel's 22nm Process Technology

- 22nm Tri-Gate transistors enhanced to reduce leakage current 2-3X with the same frequency capability
- Haswell version of 22nm has 11 metal interconnect layers compared to 9 layers on Ivy Bridge to optimize performance, area and cost



Haswell builds on innovations in 2<sup>nd</sup> and 3<sup>rd</sup> Generation Intel® Core™ i3/i5/i7 Processors (Sandy Bridge/Ivy Bridge) with optimized Intel process technology!

# Sequence

Family of Innovations!

Power Efficiency and Management

FIVR – Fully Integrated Voltage Regulator

Cache Hierarchy and Interconnects

Gfx/Media

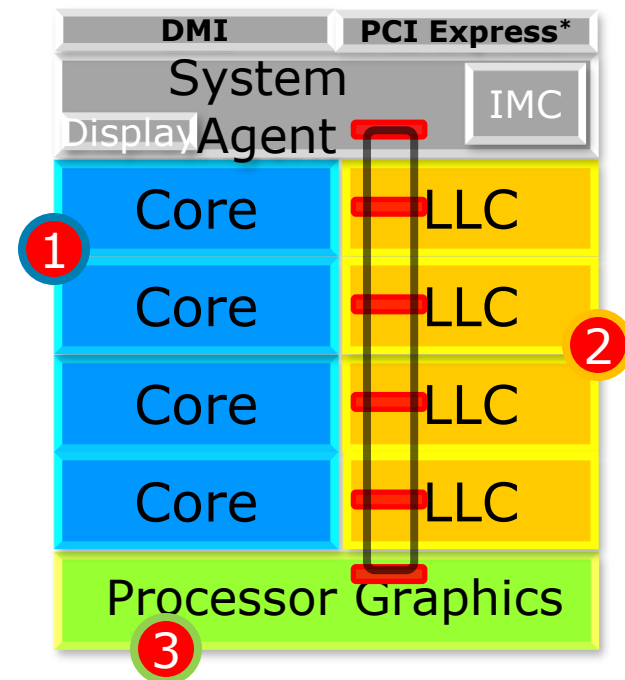
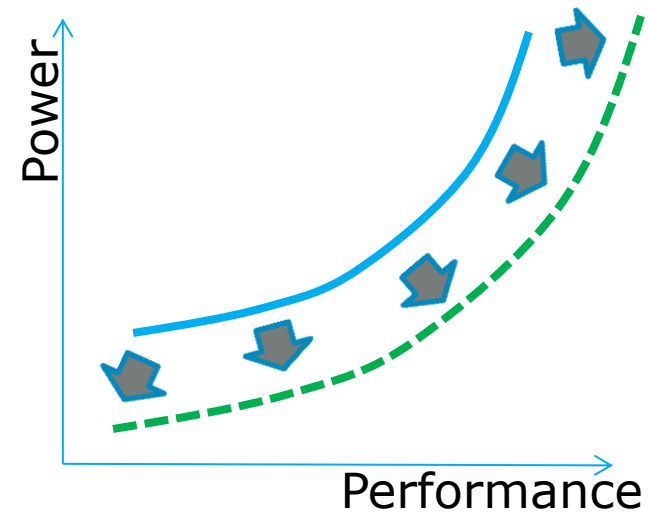
Intel® Microarchitecture (Haswell): Core

ISA

Wrap Up

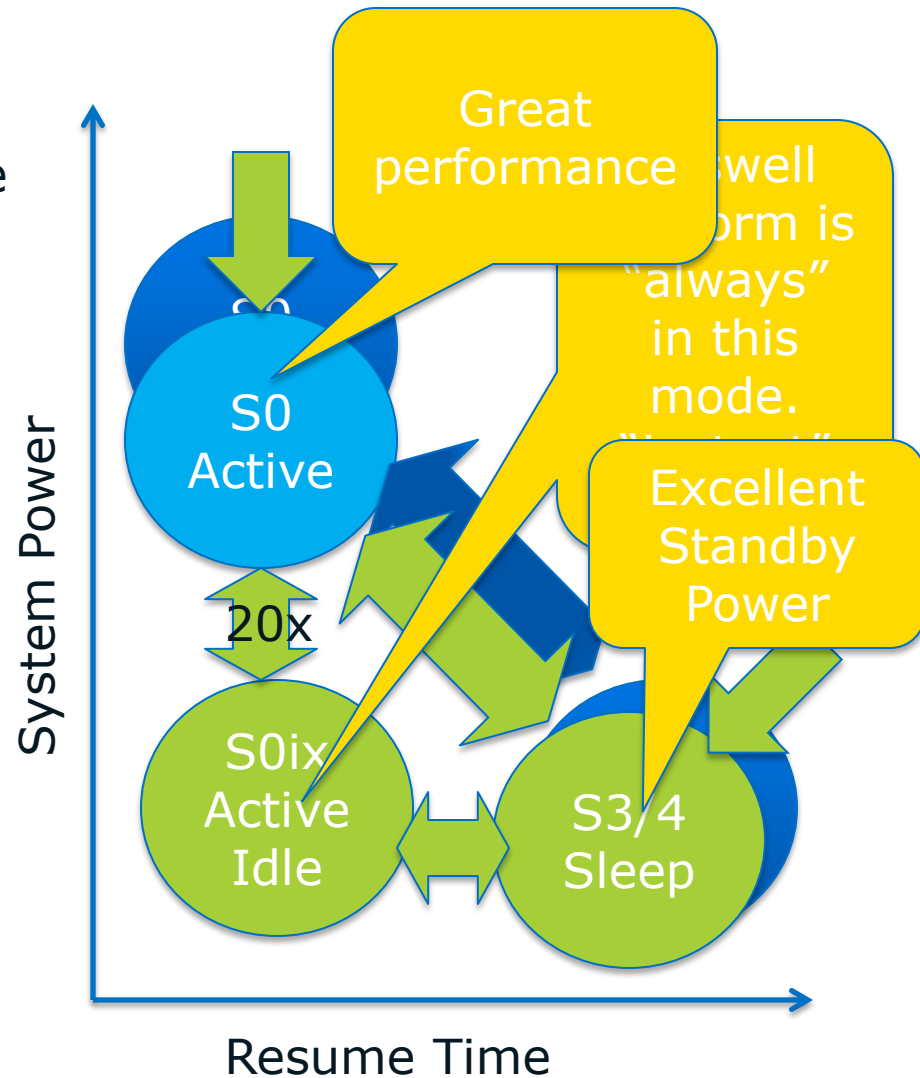
# Power Efficiency: Maximizing Power-Limited Performance

- Extended operating range
  - Increased Turbo
  - New C-states, improved latency
  - Power efficient features: better than voltage / frequency scaling
  - Continued focus on gating unused logic and low-power modes
  - Optimized manufacturing and circuits
- Independent frequency domains
  - Cores separated from LLC+Ring for fine-grained control
  - Power Control Unit dynamically allocates budget when power-limited
  - Prioritization based on run-time characteristics selects domain with the highest performance return



# Haswell Power Management Innovation

- All day experiences
  - Improving power efficiency for active workloads
- Evolutionary improvements
- New extremely low-power active state
  - 20x improvement from prior generation
  - Enables significant improvement in realizable battery life
  - Automatic, continuous, fine-grained, transparent to well written SW
  - Leverages learnings from phone & tablet development



***Everything that is not needed is turned off!***



# Sequence

Family of Innovations!

Power Efficiency and Management

FIVR – Fully Integrated Voltage Regulator

Cache Hierarchy and Interconnects

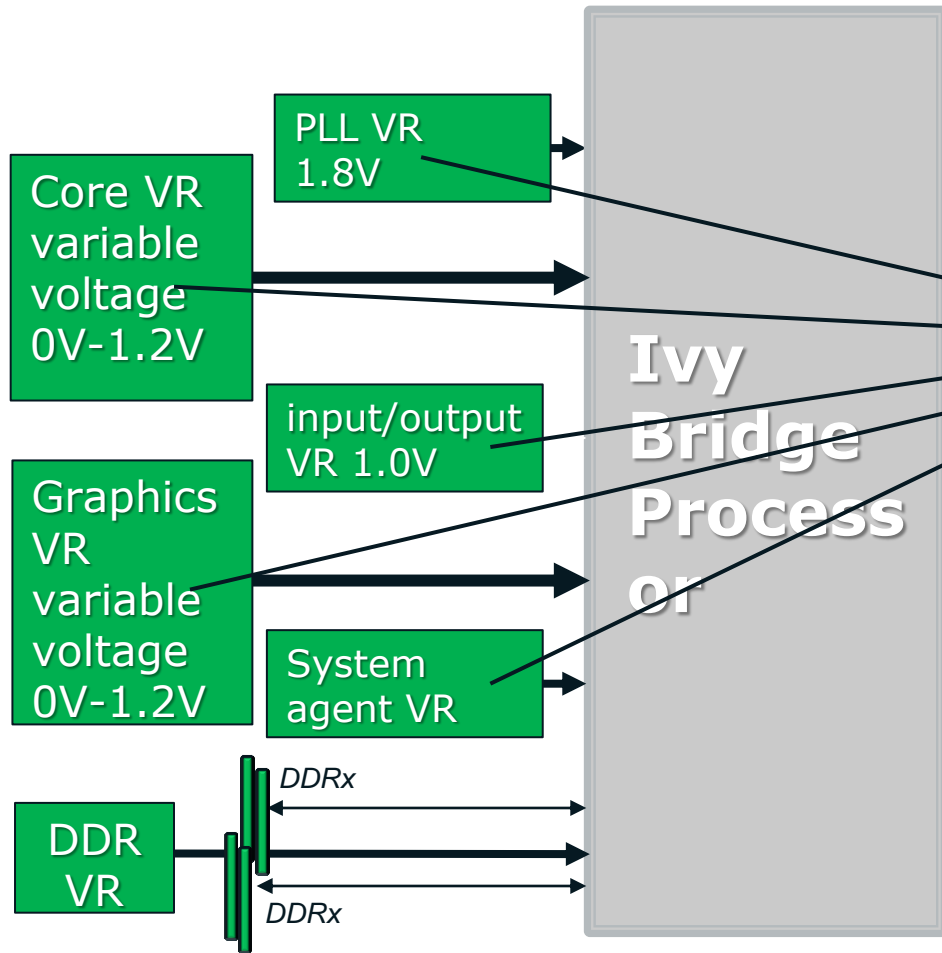
Gfx/Media

Intel® Microarchitecture (Haswell): Core

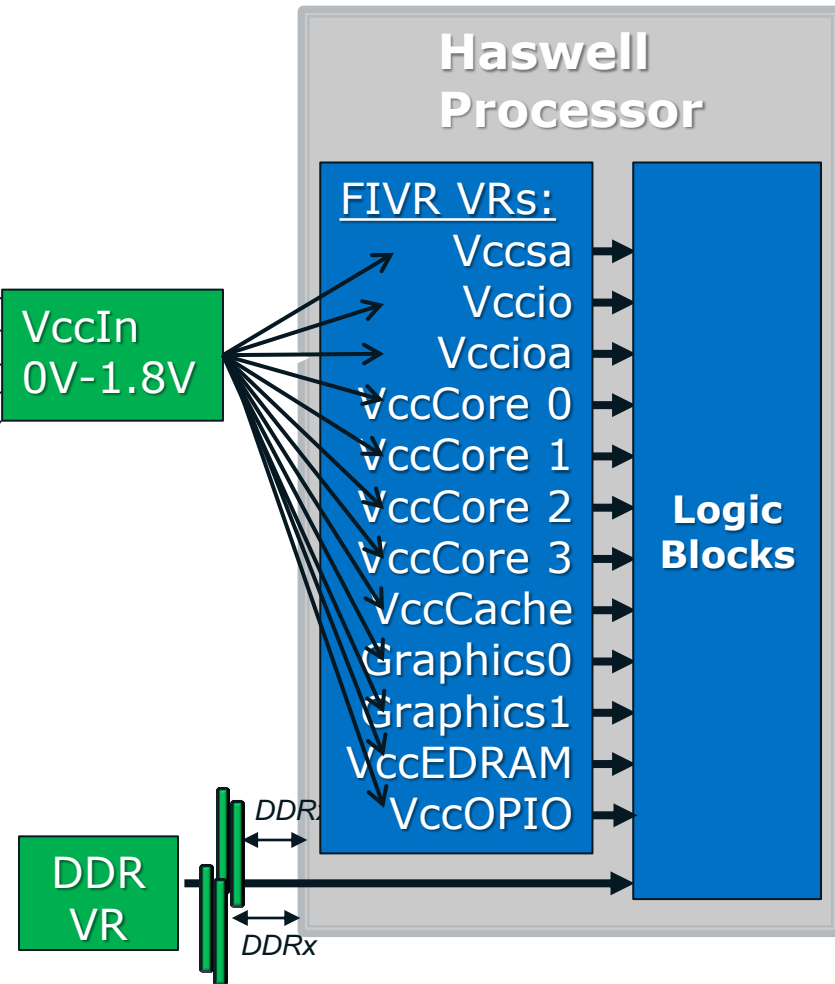
ISA

Wrap Up

# Ivy Bridge Platform



# Haswell Platform



Example voltage planes


# FIVR: Platform Goodness

## Ivy Bridge

- Back is all power
- Large inductors, butterfly mounted through board
- 5.4mm thick

## Haswell

- Backside bare
- Small inductors & caps & 75% fewer
- Space for 10% larger battery
- 3.4mm thick



**2mm thinner; ~\$5 cheaper; space for 10% larger battery**

# Sequence

Family of Innovations!

Power Efficiency and Management

FIVR – Fully Integrated Voltage Regulator

Cache Hierarchy and Interconnects

Gfx/Media


Intel® Microarchitecture (Haswell): Core

ISA

Wrap Up

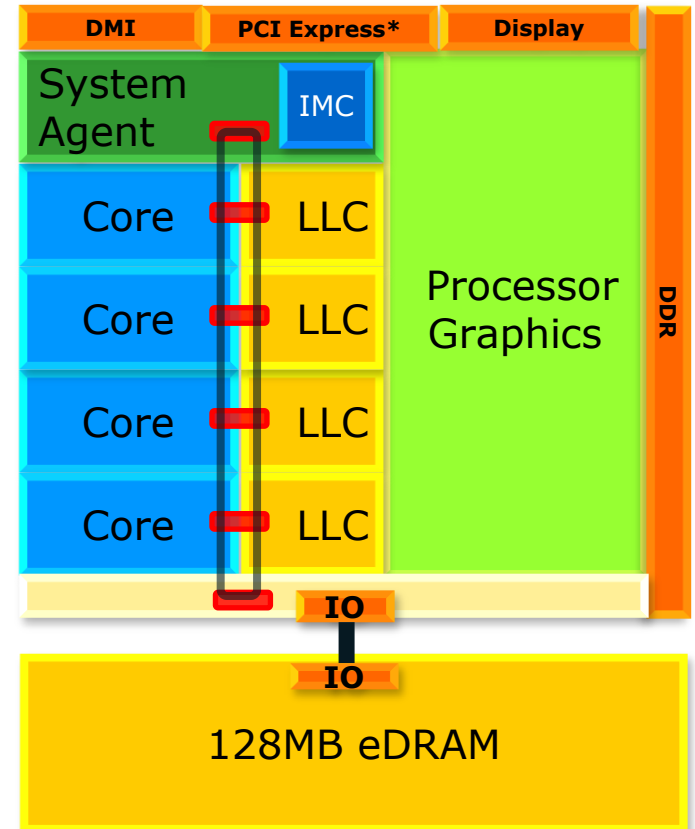
# Cache, Interconnect and System Agent

- More access bandwidth per slice of shared LLC
  - New dedicated pipelines handle data and non-data accesses independently
- Improved load balancing to System Agent
  - Better credit-based management more efficiently shares resources
- Improved DRAM write throughput
  - Deeper pending queues: more decoupling, better scheduling
- Lower power, better efficiency
  - Focused effort to reduce idle and active power (next section)



# Large eDRAM Cache

- Haswell introduces configurations with large graphics & large cache
- Cache attributes
  - High throughput and low latency
  - Flat latency vs. sustained bandwidth curve
  - Fully shared between Graphics, Media, and Cores for very efficient multi-media computing






# Large Caches in Graphics Workloads

- Intra-frame
  - Intra-render pass – capture spatial and temporal locality within a surface
    - Captured in moderate cache sizes (1-8MB LLC). SNB Si shows 20-30% speedup
  - Inter-render pass – capture a full surface from generation to subsequent consumption (shadow maps, render targets)
    - Captured in big cache sizes (16-64+MB LLC). CRW Si shows 20-30% speedup
- Inter-frame
  - Capture texture reuse across frames due to continuity between frames




# Large Cache Performance and Latency

speedup of CRW 128MB vs OMB



CRW Loaded Latency



Small latency sensitivity with load. Sustainable for random traffic

Pre-production system measurements, product measurements may vary.





# Sequence

Family of Innovations!

Power Efficiency and Management

FIVR – Fully Integrated Voltage Regulator

Cache Hierarchy and Interconnects

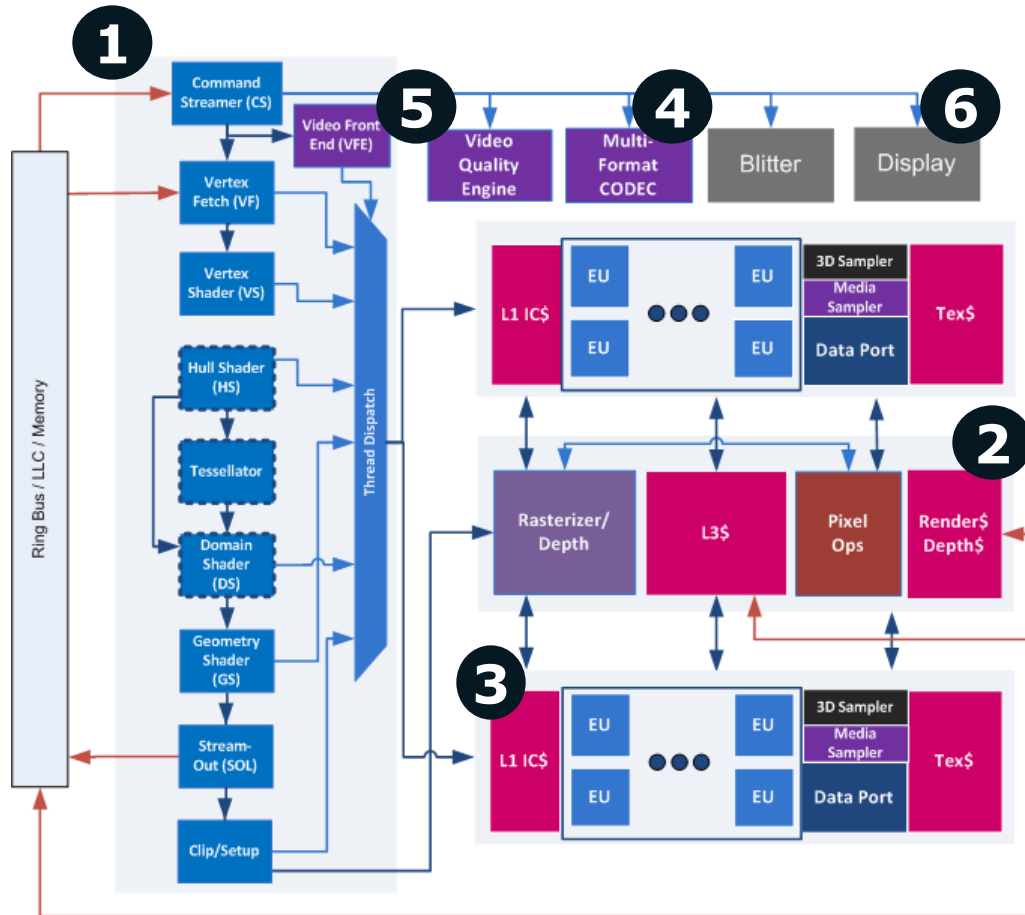
Gfx/Media

Intel® Microarchitecture (Haswell): Core

ISA

Wrap Up

# Haswell Processor Graphics Architecture Building Blocks



## Scalable Architecture partitioned into 6 domains:

1. Global Assets: Geometry Front-end up to Setup
2. Slice Common: Rasterizer, Level 3 Cache (L3\$) and Pixel Back-end
  - Scalable slices for performance and GFlop tuning
3. Sub-Slice: Shaders (EUs), Instruction Caches (IC\$) and Samplers
4. Multi-Format Video CODEC Engine (MFX)
5. Video Quality Enhancement Engine
6. Displays

**Sets the stage for Scale-up!!**

# Video Codec

Introducing hardware-based SVC (Scalable Video Coding) codec

- Allowing single encoded bit-stream for heterogeneous devices
- Key enabler for multi-participant video conferencing

MJPEG (Motion JPEG) hardware decoder

- Enabling low power HD video conferencing for USB2 webcam

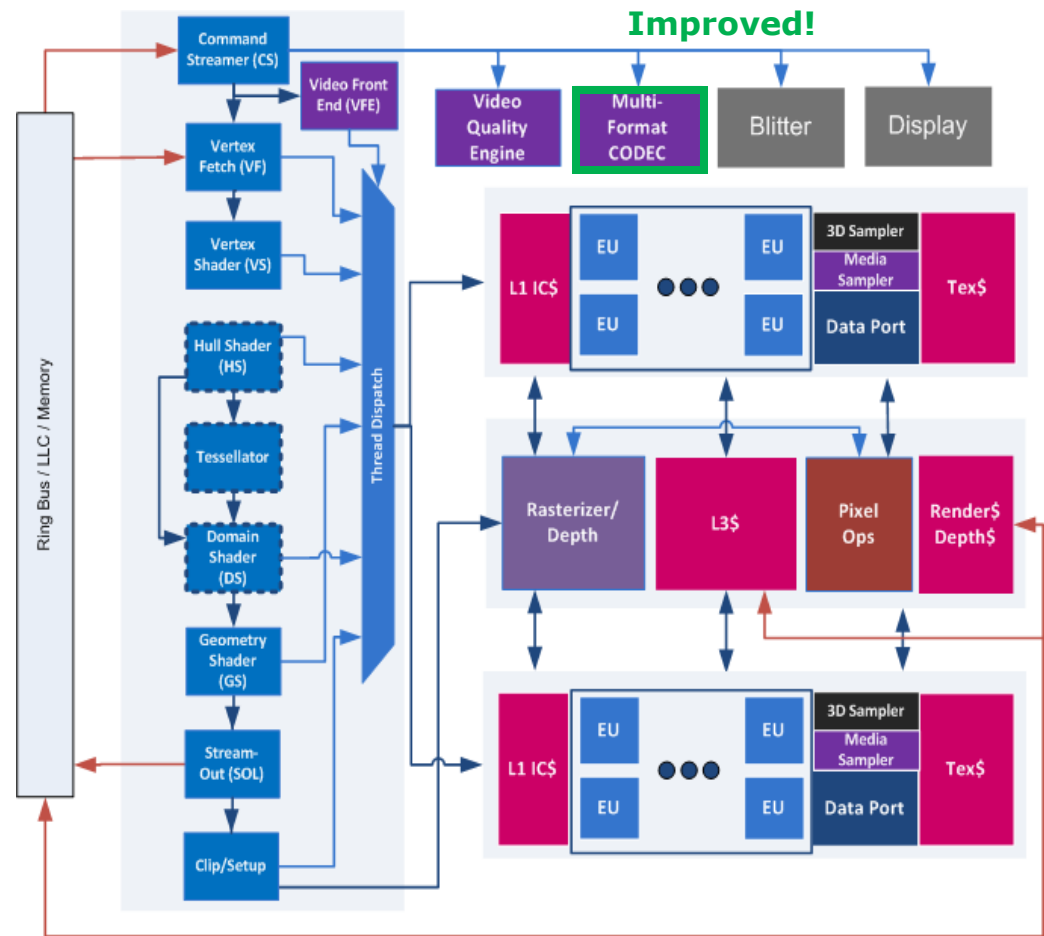
MPEG2 hardware encoder

- DVD creation
- DLNA streaming

4Kx2K video playback

Continue to drive encoder quality

- Introduced through the encoding modes in Media SDK



***Haswell adds newer codec on top existing codecs in 3<sup>rd</sup> Generation Intel® Core™ processors***

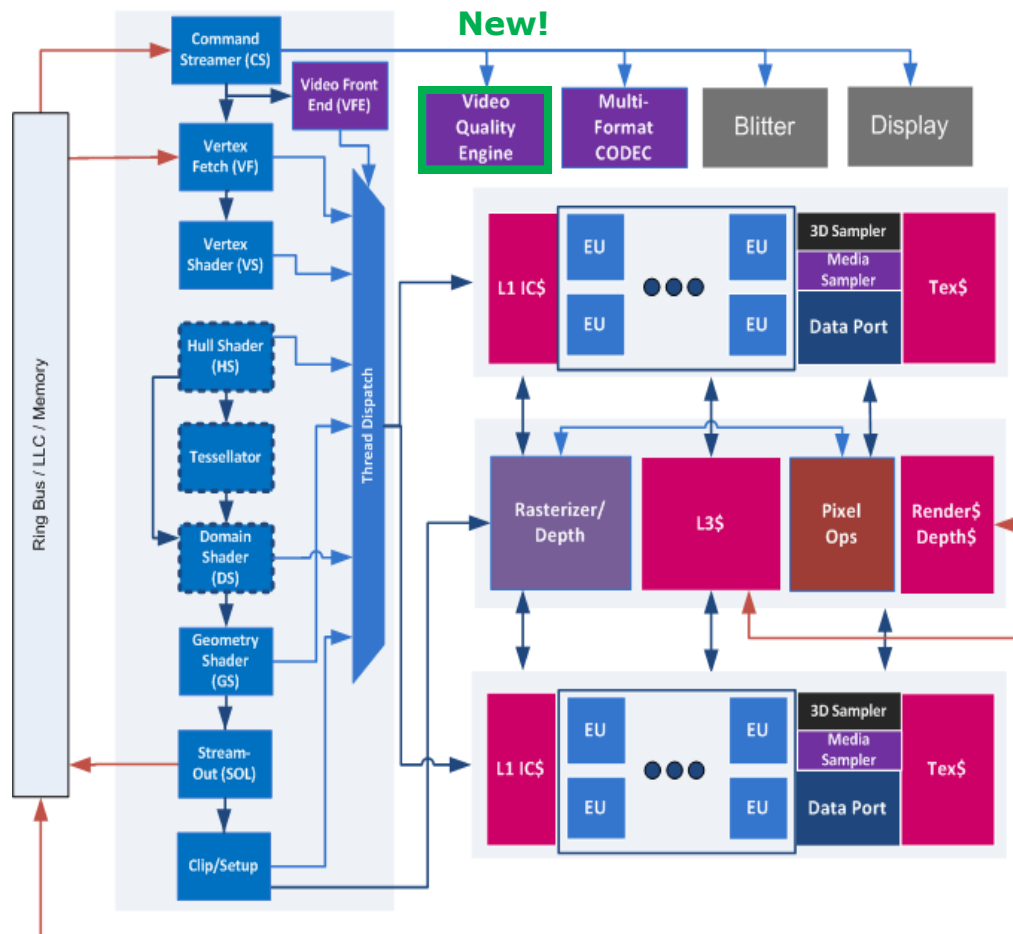
# High Quality Video Processing

Dedicated video processing on newly designed Video Quality Engine (VQE)

Haswell supports an extensive suite of video processing functions including:

- De-Noise (DN)
- De-Interlace (DI)
- Film-mode Detection (FMD)
- Skin Tone Detection (STD)
- Skin Tone Enhancement (STE)
- Total Color Control (TCC)
- Adaptive Contrast Enhancement (ACE)
- Advanced Video Scalar (AVS)
- Gamut Compression (GC)
- Gamut Expansion (GE)<sup>1</sup>
- Skin Tone Tuned Image Enhancement Filter<sup>1</sup>
- Frame Rate Conversion (FRC)<sup>1</sup>
- Image Stabilization (IS)<sup>1</sup>

<sup>1</sup>New on Haswell



**Higher quality video  
at lower power!**

# Media: Quick Sync Video Performance and Power


- 4-12x real-time transcode at various quality modes
- 10-hour video playback time on latest Apple MacBook Air
- Multi-stream 4K decode
- > real-time 4K Encode



HD Mosaic



4K Mosaic



\*Measurements based on Intel Demo Clip in Cyberlink Media Espresso Fast Conversion Mode

# Sequence

Family of Innovations!

Power Efficiency and Management

FIVR – Fully Integrated Voltage Regulator

Cache Hierarchy and Interconnects

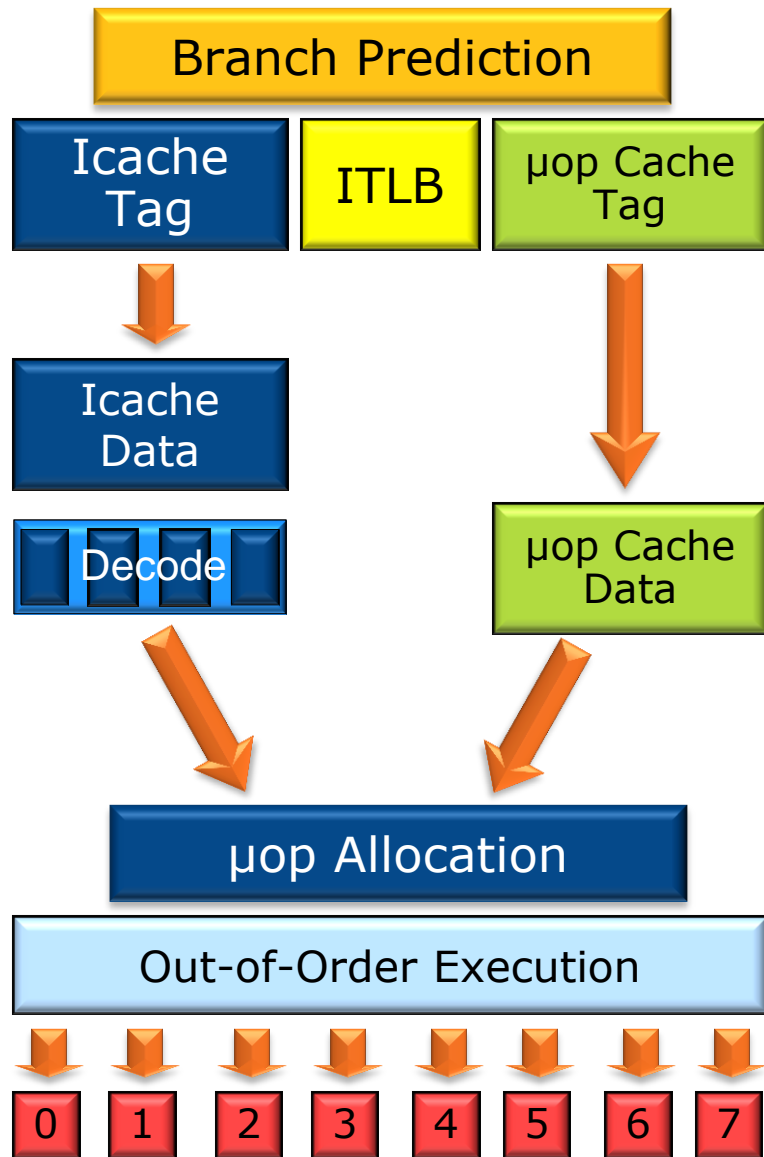
Gfx/Media

Intel® Microarchitecture (Haswell): Core

ISA

Wrap Up

# Haswell Core at a Glance



## Next generation branch prediction

- Improves performance *and* saves wasted work

## Improved front-end

- Initiate TLB and cache misses speculatively
- Handle cache misses in parallel to hide latency
- Leverages improved branch prediction

## Deeper buffers

- Extract more instruction parallelism
- More resources when running a single thread

## More execution units, shorter latencies

- Power down when not in use

## More load/store bandwidth








- Better prefetching, better cache line split latency and throughput, double L2 bandwidth
- New modes save power without losing performance

## No pipeline growth

- Same branch mis-prediction latency
- Same L1/L2 cache latency

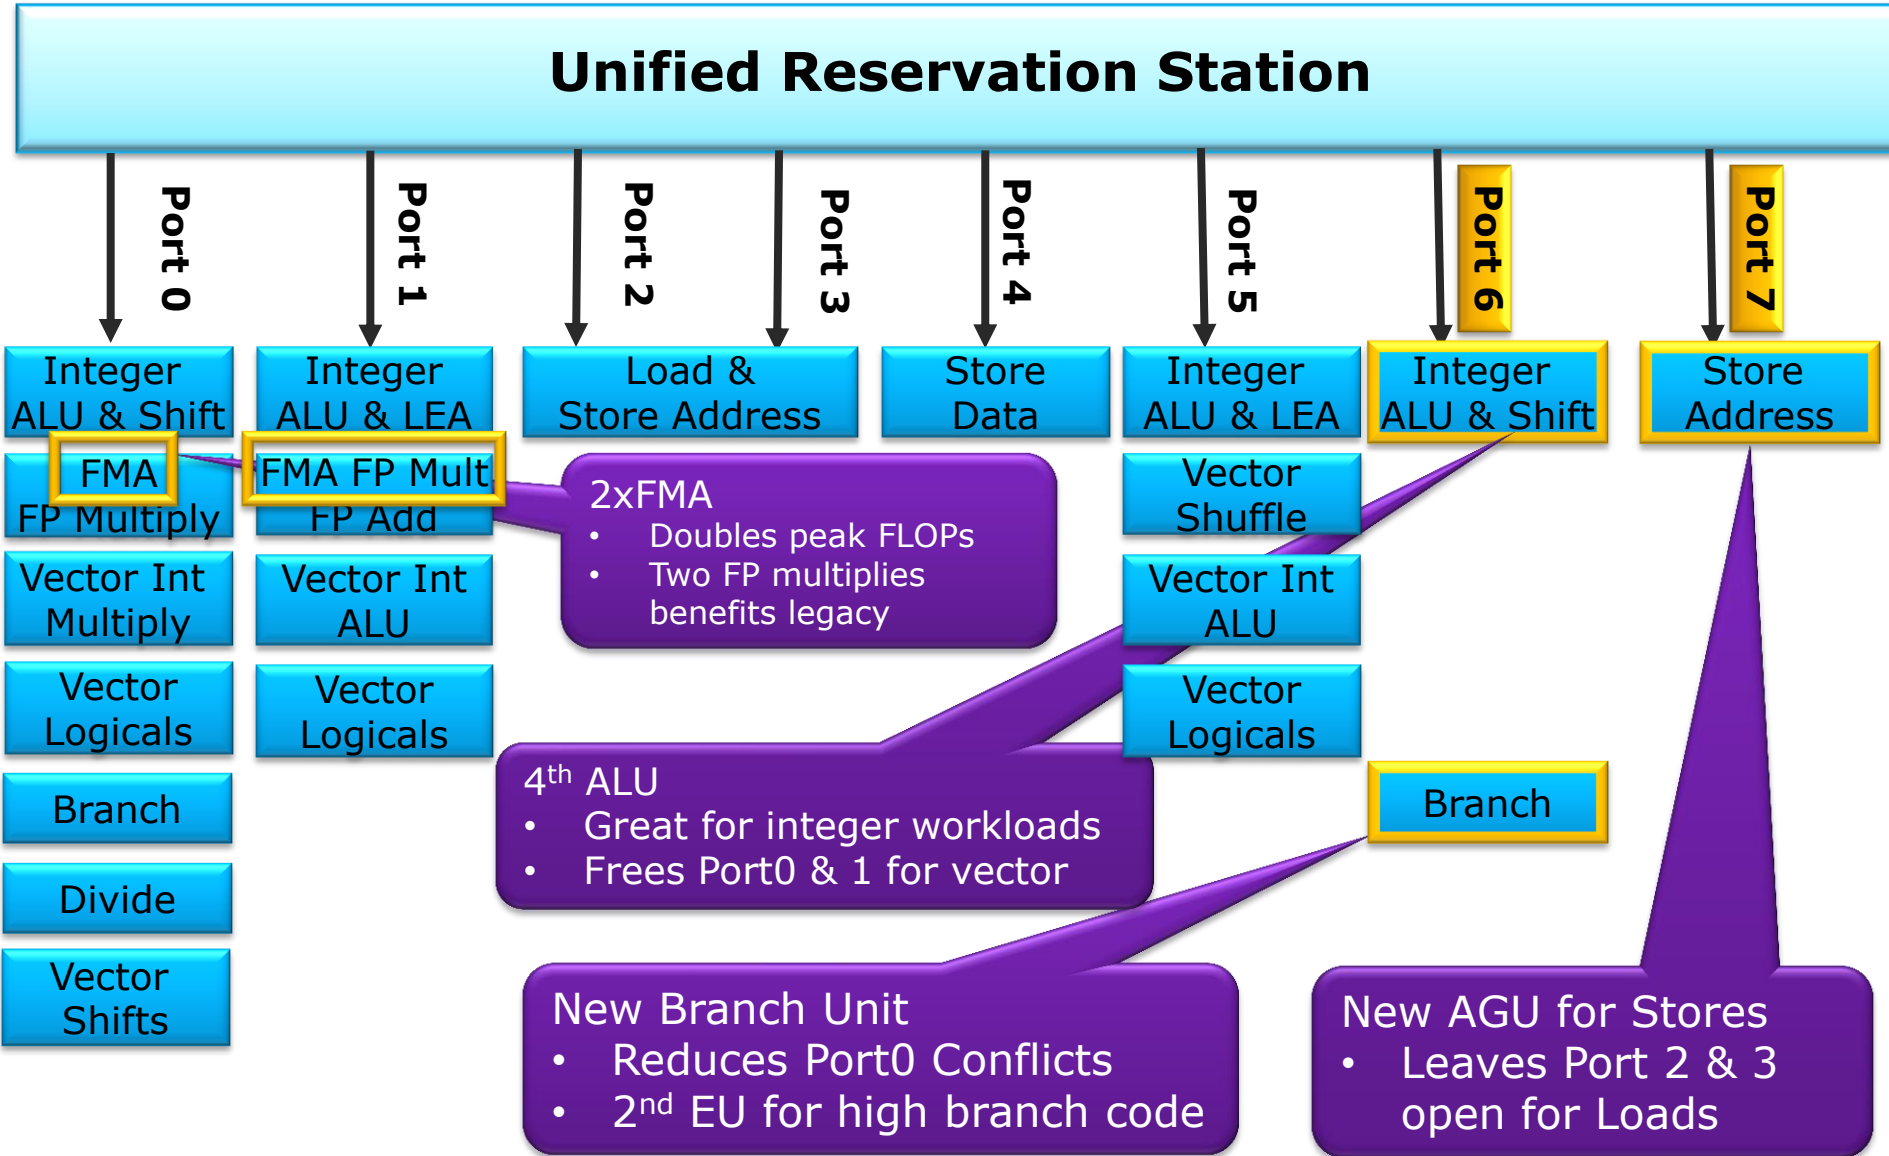
# Haswell Buffer Sizes

Extract more parallelism in every generation


	Nehalem	Sandy Bridge	Haswell	
Out-of-order Window	128	168	192	
In-flight Loads	48	64	72	
In-flight Stores	32	36	42	
Scheduler Entries	36	54	60	
Integer Register File	N/A	160	168	
FP Register File	N/A	144	168	
Allocation Queue	28/thread	28/thread	56	



# Haswell Execution Unit Overview



# FMA (Floatingpoint Multiply Add)



- 2 new FMA units provide 2x peak FLOPs/cycle of previous generation
- 2X cache bandwidth to feed wide vector units
  - 32-byte load/store for L1
  - 2x L2 bandwidth
- 5-cycle FMA latency same as an FP multiply

Latency (clks)	Prior Gen	<b>New</b> Haswell	Ratio
MulPS, PD	5	5	
AddPS, PD	3	3	
Mul+Add /FMA	8	5	<b>1.6</b>

***FMA provides improved accuracy and performance***

# Core Cache Size/Latency/Bandwidth

Metric	Nehalem	Sandy Bridge	Haswell
L1 Instruction Cache	32K, 4-way	32K, 8-way	32K, 8-way
L1 Data Cache	32K, 8-way	32K, 8-way	32K, 8-way
Fastest Load-to-use	4 cycles	4 cycles	4 cycles
Load bandwidth	16 Bytes/cycle	32 Bytes/cycle (banked)	<b>64 Bytes/cycle</b>
Store bandwidth	16 Bytes/cycle	16 Bytes/cycle	<b>32 Bytes/cycle</b>
L2 Unified Cache	256K, 8-way	256K, 8-way	256K, 8-way
Fastest load-to-use	10 cycles	11 cycles	11 cycles
Bandwidth to L1	32 Bytes/cycle	32 Bytes/cycle	<b>64 Bytes/cycle</b>
L1 Instruction TLB	4K: 128, 4-way 2M/4M: 7/thread	4K: 128, 4-way 2M/4M: 8/thread	4K: 128, 4-way 2M/4M: 8/thread
L1 Data TLB	4K: 64, 4-way 2M/4M: 32, 4-way 1G: fractured	4K: 64, 4-way 2M/4M: 32, 4-way 1G: 4, 4-way	4K: 64, 4-way 2M/4M: 32, 4-way 1G: 4, 4-way
L2 Unified TLB	4K: 512, 4-way	4K: 512, 4-way	<b>4K+2M shared: 1024, 8-way</b>

All caches use 64-byte lines

# Sequence

Family of Innovations!

Power Efficiency and Management

FIVR – Fully Integrated Voltage Regulator

Cache Hierarchy and Interconnects

Gfx/Media

Intel® Microarchitecture (Haswell): Core

ISA

Wrap Up

# Haswell New Compute Instructions

## Intel® Advanced Vector Extensions 2 (Intel® AVX2)

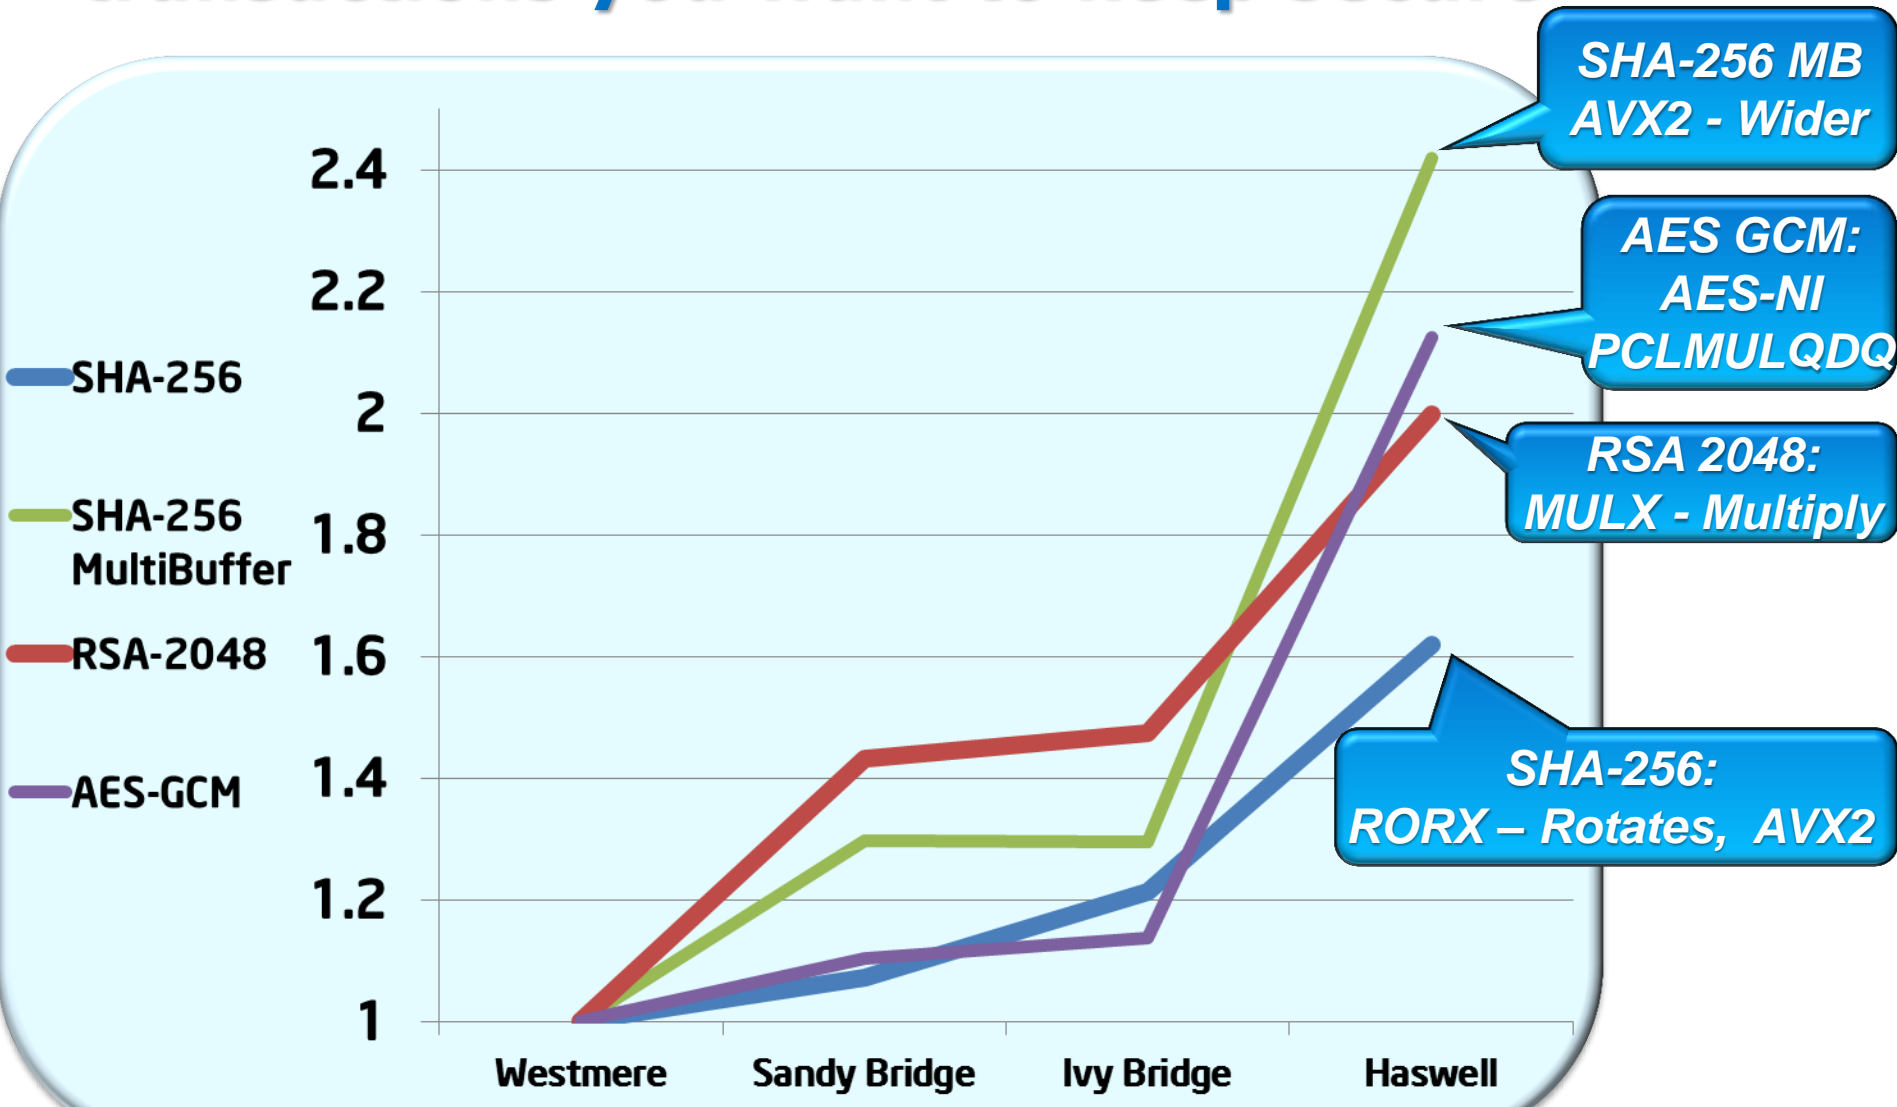
- Includes
  - 256-bit Integer vectors
  - FMA: Fused Multiply-Add
  - Full-width element permutes
  - Gather
- Benefits
  - High performance computing
  - Audio & Video
  - Games
- New Integer Instructions
  - Indexing and hashing
  - Cryptography
  - Endian conversion – MOVBE

	Instruction Set	SP FLOPs per cycle	DP FLOPs per cycle
Nehalem	SSE (128-bits)	8	4
Sandy Bridge	AVX (256-bits)	16	8
Haswell	AVX2 & FMA	32	16

Group	Instructions
Bit Field Pack/Extract	BZHI, SHLX, SHR, SARX, BEXTR
Variable Bit Length Stream Decode	LZCNT, TZCNT, BLSR, BLSMSK, BLSI, ANDN
Bit Gather/Scatter	PDEP, PEXT
Arbitrary Precision Arithmetic & Hashing	MULX, RORX

- Full Instruction Specification Available at: <http://software.intel.com/en-us/avx/>

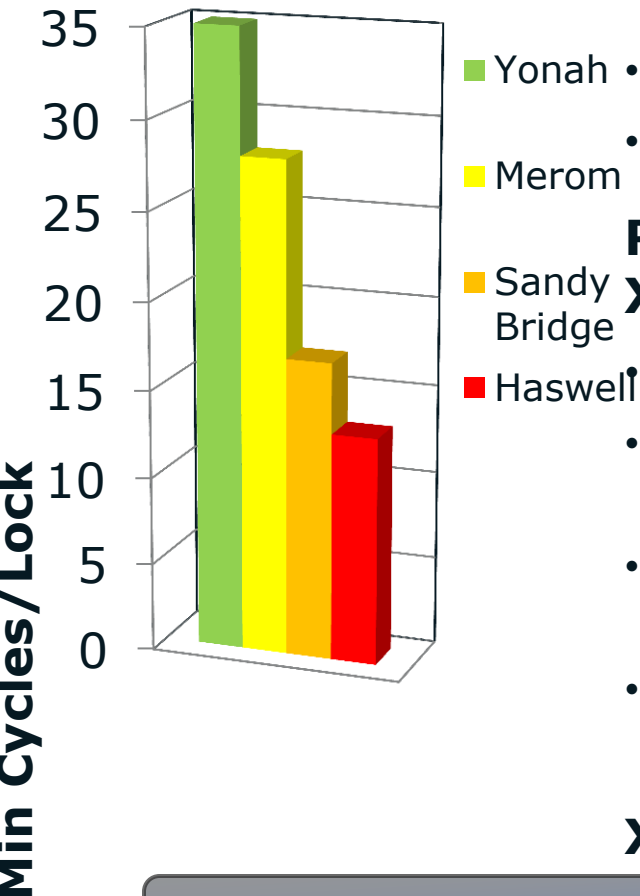
# Cryptography protects nearly all data and transactions you want to keep secure



Haswell's microarchitecture improvements and new instructions enable substantial gains in cryptography

# Locks!

## Cached Lock Performance



## HLE: Hardware Lock Elision – XACQUIRE/XRELEASE

- Software uses legacy compatible hints to identify critical section. Hints ignored on hardware without TSX.
- Hardware support to execute transactionally without acquiring lock
- Abort causes a re-execution without elision
- Hardware manages all architectural state


## RTM: Restricted Transactional Memory – XBEGIN/XEND

- Software uses new instructions to specify critical sections
- Similar to HLE but flexible interface for software to do lock elision
- Abort transfers control to target specified by XBEGIN operand
- Abort information returned in a general purpose register (EAX)

## XTEST and XABORT – Additional instructions

*Bringing Transactional Memory to the Mainstream*

# TSX Evaluation on HPC Workloads



Substitute atomic operations, locks, and non-blocking sync. with RTM

- Average 1.41x speedup with 8 threads

Workloads benefit from RTM by

1. Exploiting concurrency within a critical section (**nufft**)
2. Reducing the synchronization cost (**ssca2, physicsSolver, nufft, histogram**)
3. Replacing complex non-blocking sync. w/ regular memory ops (**canneal**)



# Virtualization on Haswell with Intel® VT

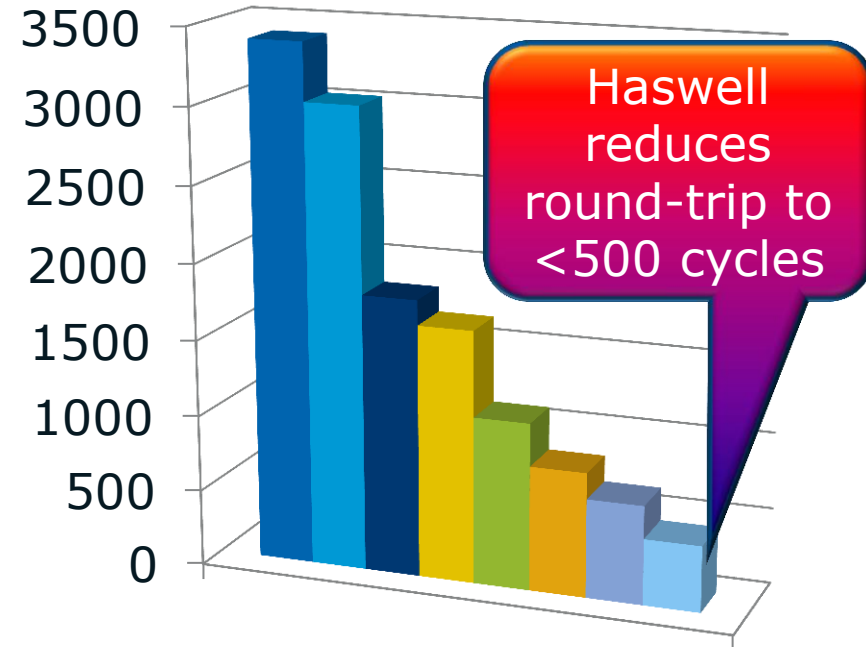
Substantially improved guest/host transition times

New *Accessed* and *Dirty* bits for Extended Page Tables (EPT) eliminates major cause of vmexits

Overhauled TLB invalidations – lower latency, less serialization

New VMFUNC instruction enables hyper-calls without a vmexit

Intel® VT-d adds 4-level page walks to match Intel® VT-x



Intel VT-x Roundtrip over Generations

# Sequence

Family of Innovations!

Power Efficiency and Management

FIVR – Fully Integrated Voltage Regulator

Cache Hierarchy and Interconnects

Gfx/Media

Intel® Microarchitecture (Haswell): Core

ISA

Wrap Up

# Wrap Up!

- **Huge family:** SOC methodology, common architecture
- **Low power platform:** 20x idle power reduction, low power IO (I2C, SDIO, I2S, UART), Link power management (USB, PCIe, SATA)
- **Large eDRAM Cache**
- **Platform:** PSR (Panel Self Refresh)
- **FIVR:** Fully Integrated Voltage Regulator
- **Core:** FMA (Floating-point Multiply Add), 2x Cache BW, TSX (Transaction Synchronization Extension)
- **Graphics:** 2x in Ultrabooks, OpenCL 1.2, DX 11.1, OpenGL 4.0
- **Media:** 5x faster at 0.5x power