Ivy Bridge

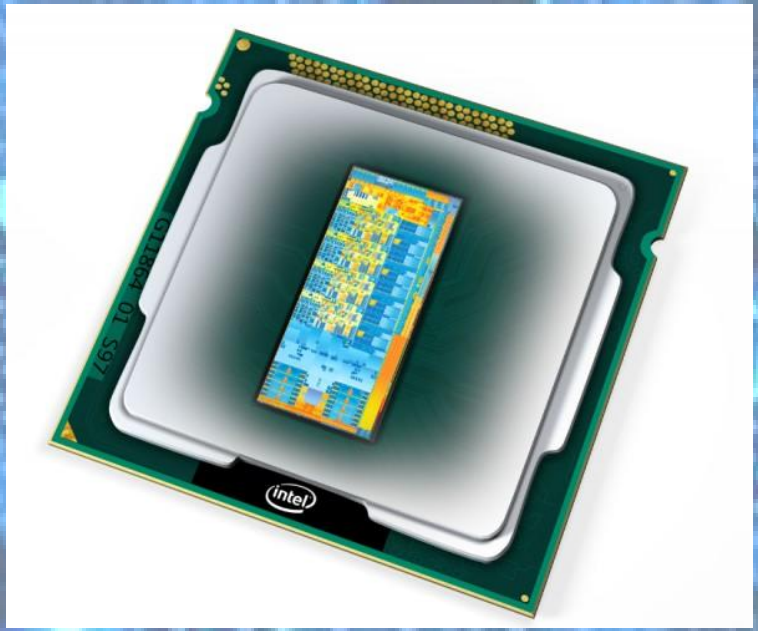

Dan Bower
Ashley Sukhavong

Moore's Law

Moore’s Law

- Intel’s co founder Gordon Moore.
- “The number of transistors on a chip will double approximately every two years.”
- It’s a guiding principle that has been used to produce new technological advances at Intel.
- When you have more transistors on a chip, you increase productivity and performance while cutting down the cost of each transistor.

Moore’s Second Law

- The cost of semiconductor fab also increases exponentially over time

Limits of Moore’s Law

- Eventually miniaturization will reach a limit at atomic level.
- While you can have a growth in transistor density, that does not mean you get greater CPU performance.
- You need to find ways to maintain power dissipation and leakage through using different materials

The Tick-Tock model through the years

Manufacturing process technology | Microarchitectures

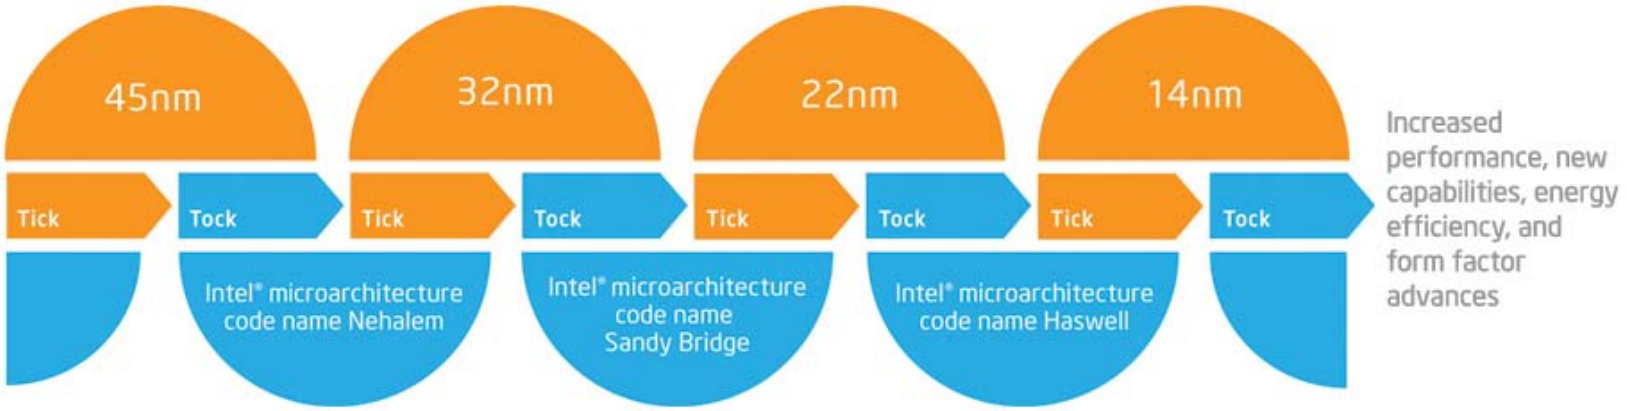

Newest Manufacturing Technology Delivers

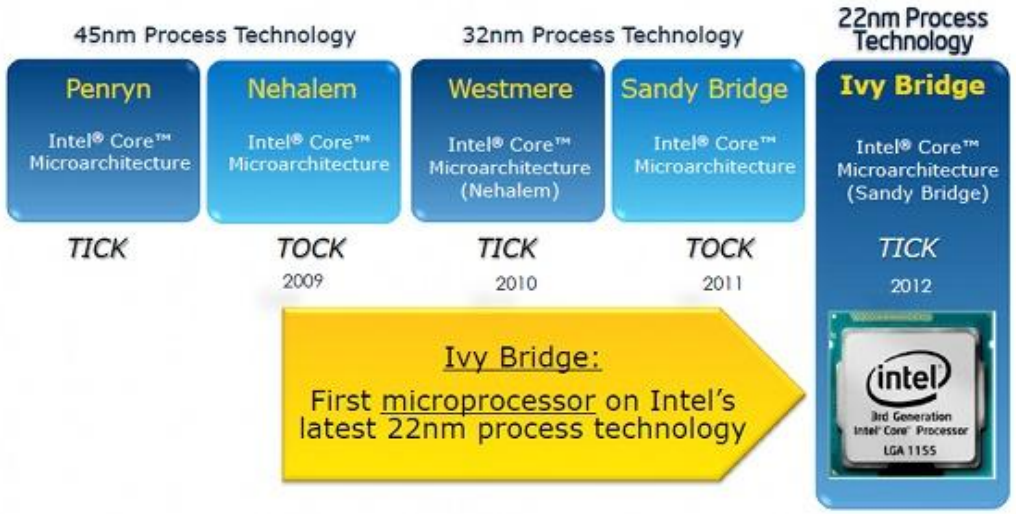

Sandy Bridge 995 million transistors in a 216 m^2

Memory controller, the PCI Express (PCIe) controller, and video functions are all located within the processor die

Features

- Uop Micro-operation Cache
- Improved Branch Prediction
- Advanced Vector Extensions (AVX)
- The last level cache (LLC)
- The System Agent
- Turbo Boost
- Quick Sync

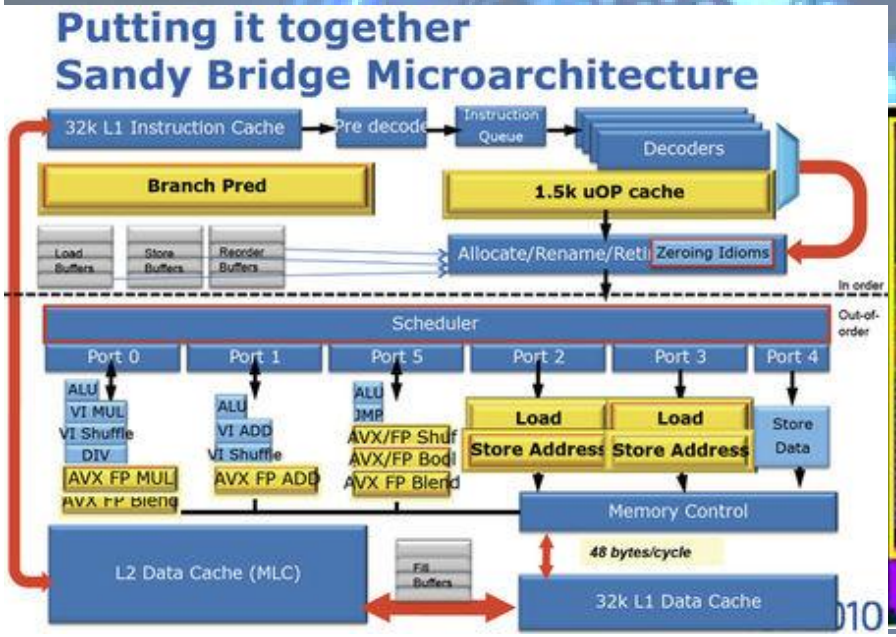

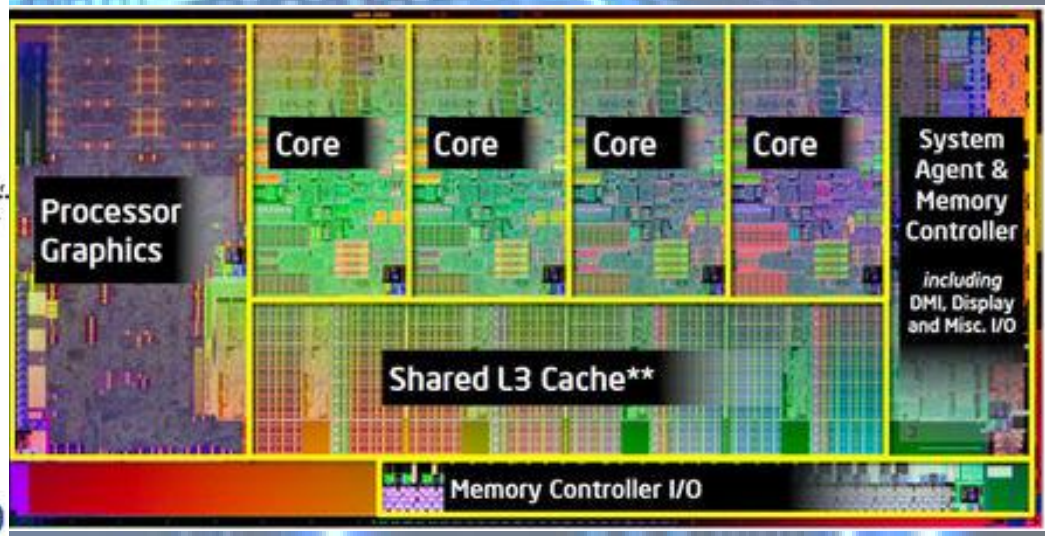

3D Tri-gate Transistors

"The easiest way to imagine a Tri-Gate transistor is to think of it as a normal transistor folded up in the middle"

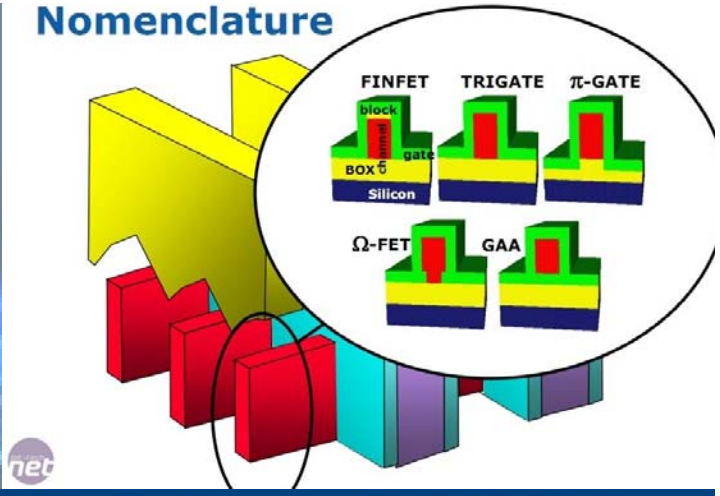

- Ivy bridge's three sides of the channel equate to 22 nm (transistor width = 2H + w)
- 3-D concept Digh Hisamoto MuGFet in 1989
- First tri-gate transistor was made by Ghau and Kayalieros in 2002 and 2006.
- 3-D transistors are known as FinFET, a block material across the top of the channel
- A tri-gate transistor has a channel with three dimensions. The flow of electricity is on all three sides. It reduces transistor size on silicon die to the width of the fin while still being a long enough gate for a good signal.

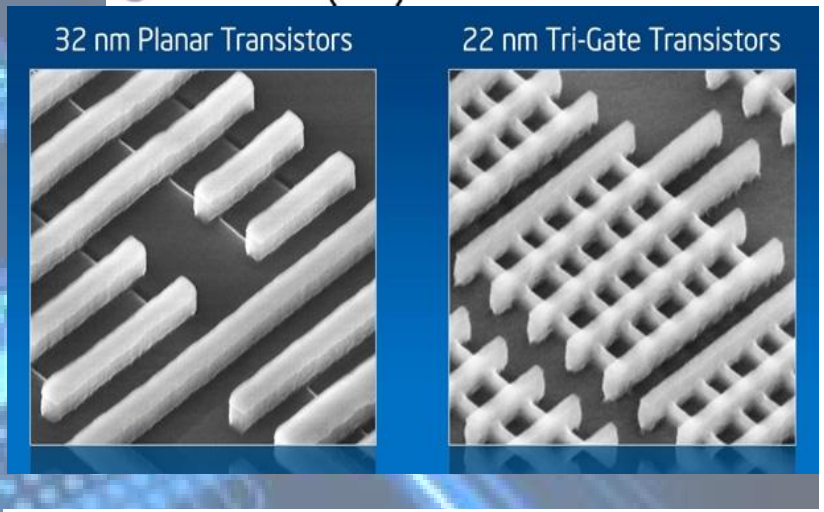

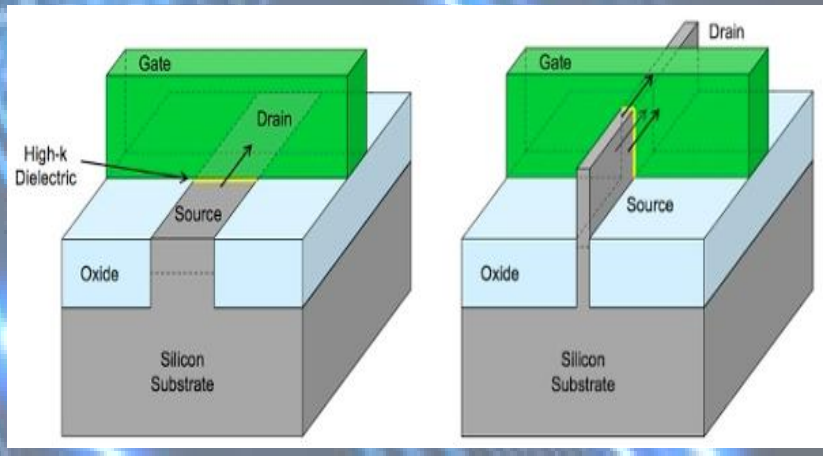

Ivy Bridge Architecture

1.4 Billion Transistors and 160mm^2

- 2-platform partition: CPU and PCH
- 4 cores with shared L3 cache.
- To each side of this central portion are the system agent and the graphics core.
- All these components are bound by a ring-bus that transports data between them.
- Hyper threading: dynamically allocate resources to threads so that if there is only a single thread active, all resources will be dedicated to that thread rather than some going unused
- Instruction set: added support for conversion between 16-bit floating point and 32-bit single precision floating point numbers

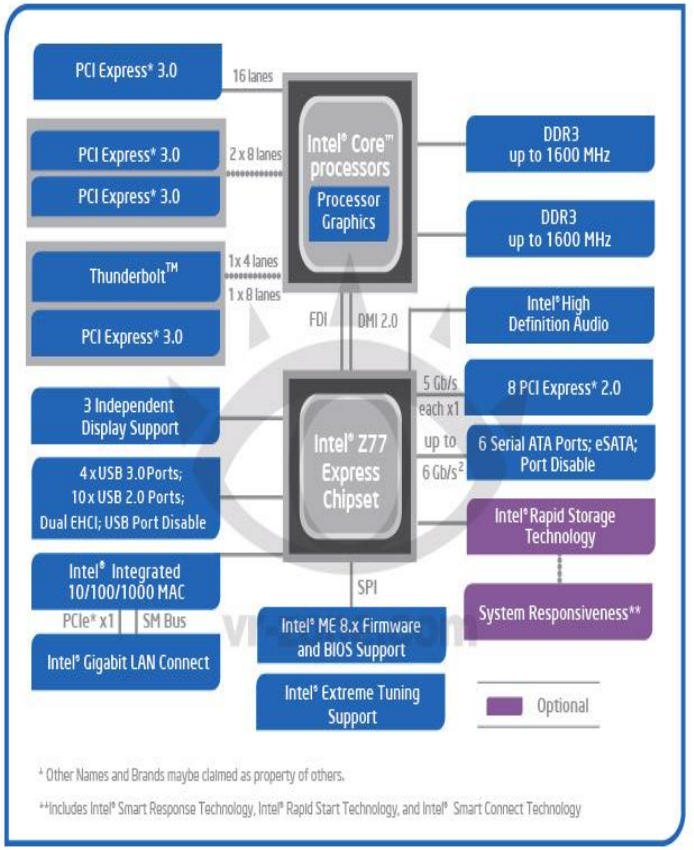

Intel® Z77 Express Chipset Platform Block Diagram

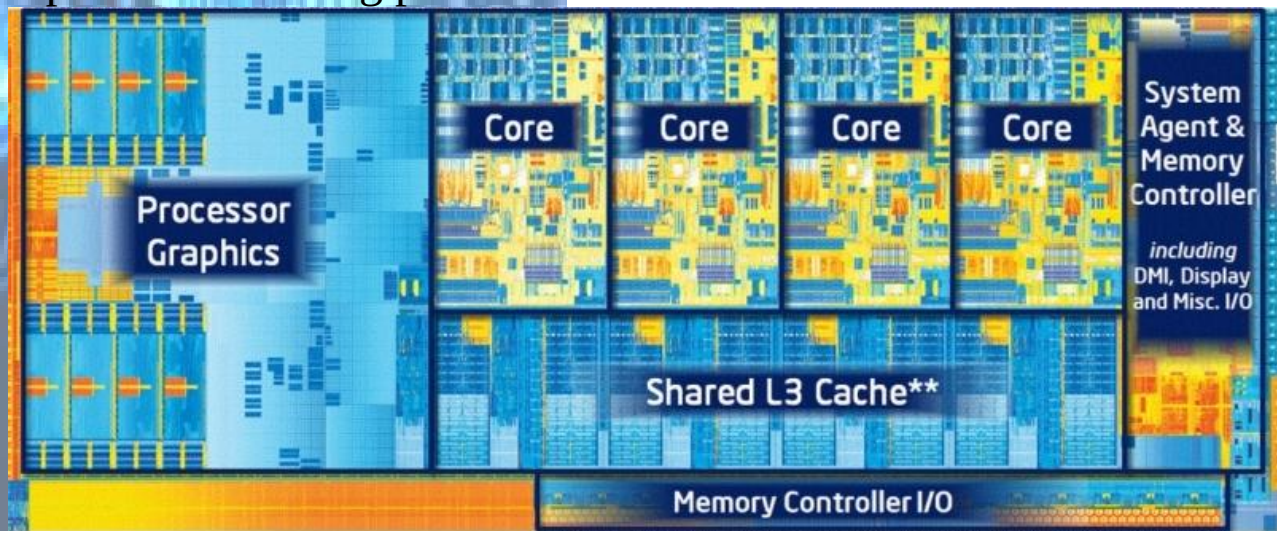

Graphics Architecture

- Ivy Bridge has two GPUs, the high-end HD 4000 and the low-end HD 2500
- OpenCL 1.1, DirectX 11 and OpenGL 3.1 support
- 5 different domains:
  global, slice common, slice, media and display
- The global domain includes the bulk of the graphics and media front-end
- The slice common domain is for the hardware that is shared by the entire shader array, such as the rasterizer, render output pipelines (ROPs) and the brand new L3 cache and unified return buffer (URB).
- A slice contains shader cores and shared resources such as the instruction cache and sampling pipeline.
- The media-processing domain includes Intel's programmable codecs and the video front-end that spawns media threads.

2 levels of concurrency for media encoding

1. programmable media processing and fixed functions stages are now asynchronous. Each stage can concurrently operate on different frames within the same stream.
2. the media and graphics shaders can now share the shader array in time slices

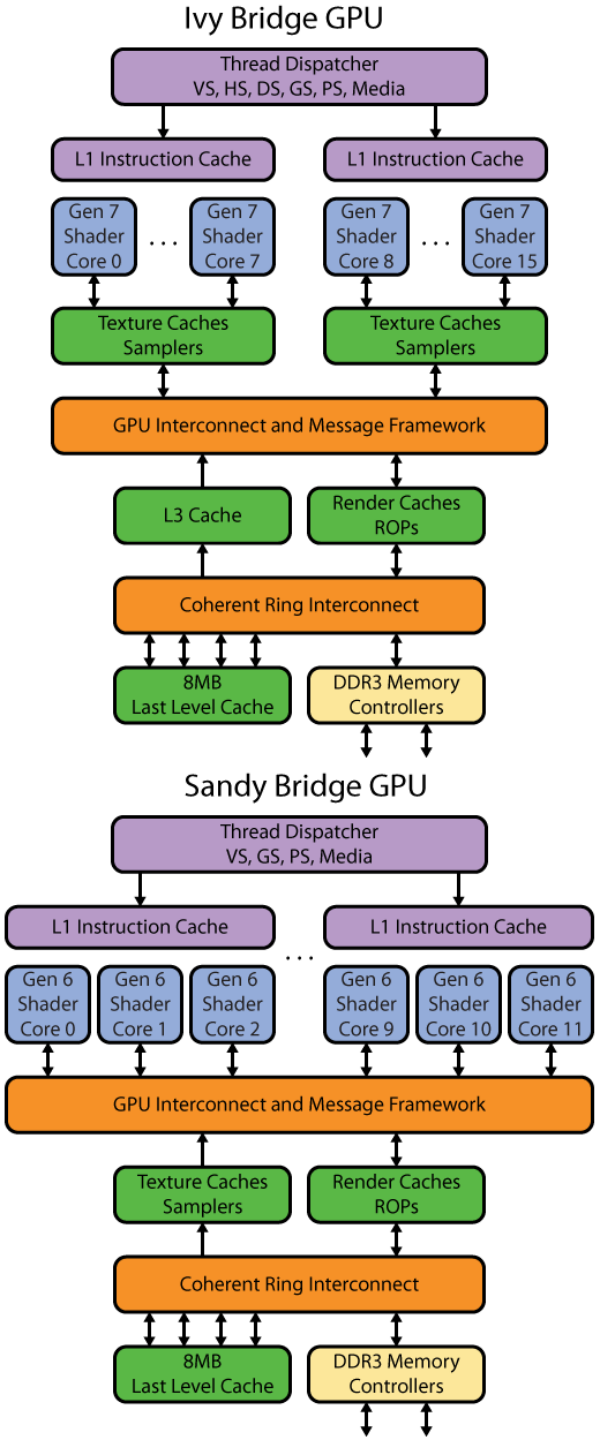

Graphics Architecture

- When texture fetch or memory accesses is needed, an integer address is sent through the messaging network to the shared sampling pipelines.
- Incorporates a high bandwidth L3 cache that is shared by the entire shader array.
- L3 a graphics only cache:backs the L1 and L2 texture caches
- Separate from Ivy Bridge's last level cache (LLC)
- The actual array is physically 512 KB, but it is partitioned into sections shared with other parts of the GPU
- L3 cache takes up half the array which is 256 KB and 32-way associative with 64B lines. It is implemented as 4 banks, each containing 32 sets and delivering a full cache line for an aggregate 256KB/cycle.
- The replacement policy is a pseudo-LRU algorithm and the L3 is way partitioned between data, textures, instructions and constants.
- he main motivation for the L3 cache is to absorb the bandwidth required by the texture pipeline.T
- URB is significantly larger than in Sandy Bridge, and shares the 512KB array for the L3 cache.
  - 256 KB of the array is used for the L3 cache.
  - 256 KB is available for the URB and the OpenCL shared

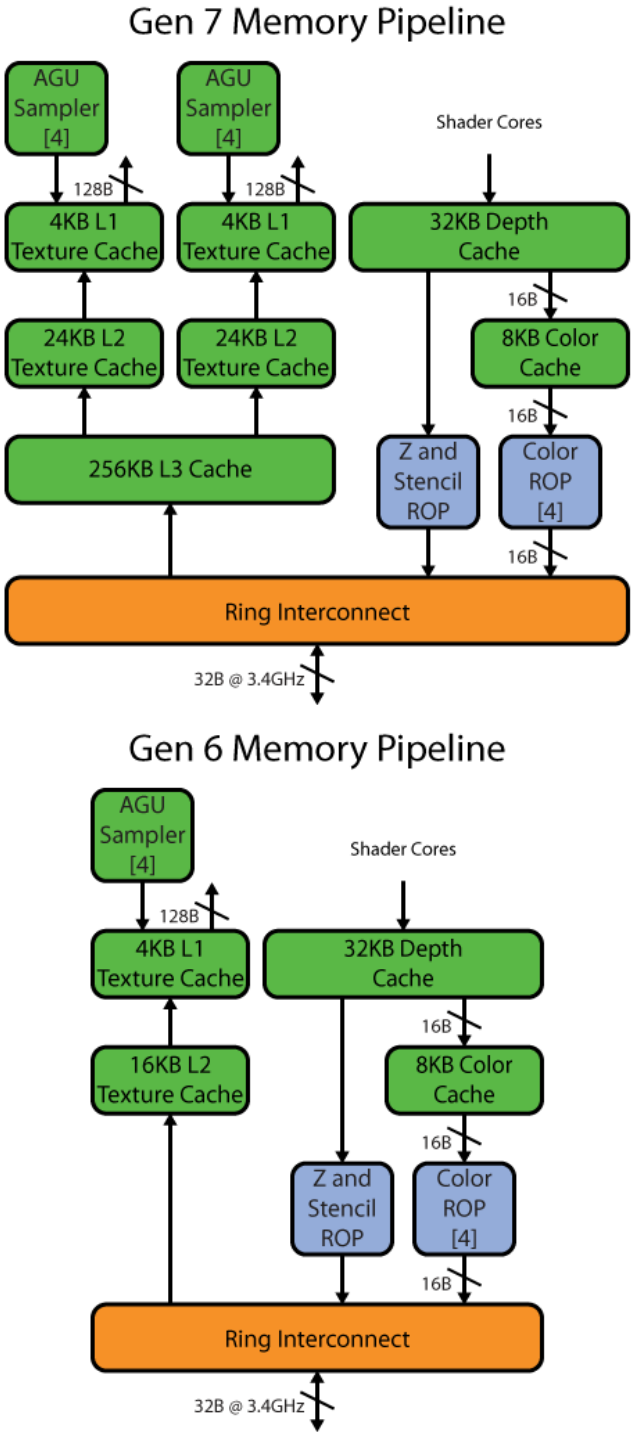

Graphic Architecture

- Sandy Bridge was a scoreboard and could select two instructions per cycle for execution. The instructions needed co-issued and executed in-order from different threads.
- Removed the message register file (MRF).
- Messages can be sent directly from the general-purpose register file (GRF)
  - The GRF has expanded to 32 KB to handle 8 threads.
  - Each thread is allocated 128 general-purpose registers
  - A separate architectural register file (ARF) contains the control information and special purpose registers for each thread which doubled the number of operations per cycle for each shader core.

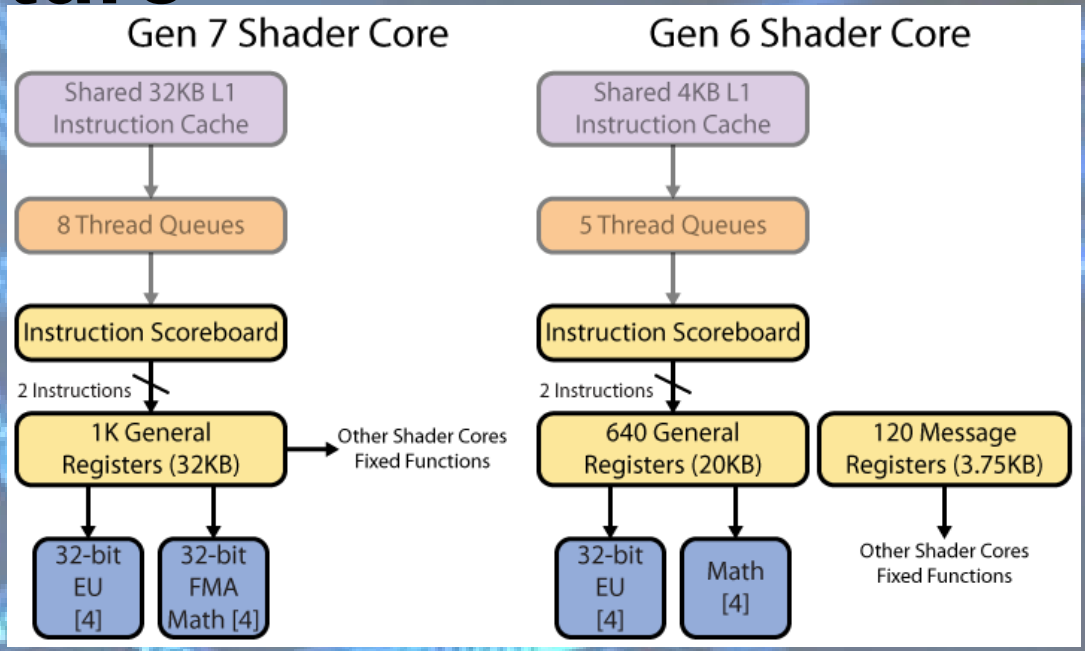

- The execution units gives more opportunities for parallel execution
- The first pipeline for Ivy Bridge has also been enhanced for double precision support.
  - 64-bit floating point instructions are executed at half speed.
  - Denormal numbers which are underflow are handled without any latency or throughput penalties.
- The second pipeline significantly to create more opportunities for co-issuing and attain higher IPC.
  - The second pipeline does not handle 64-bit data.

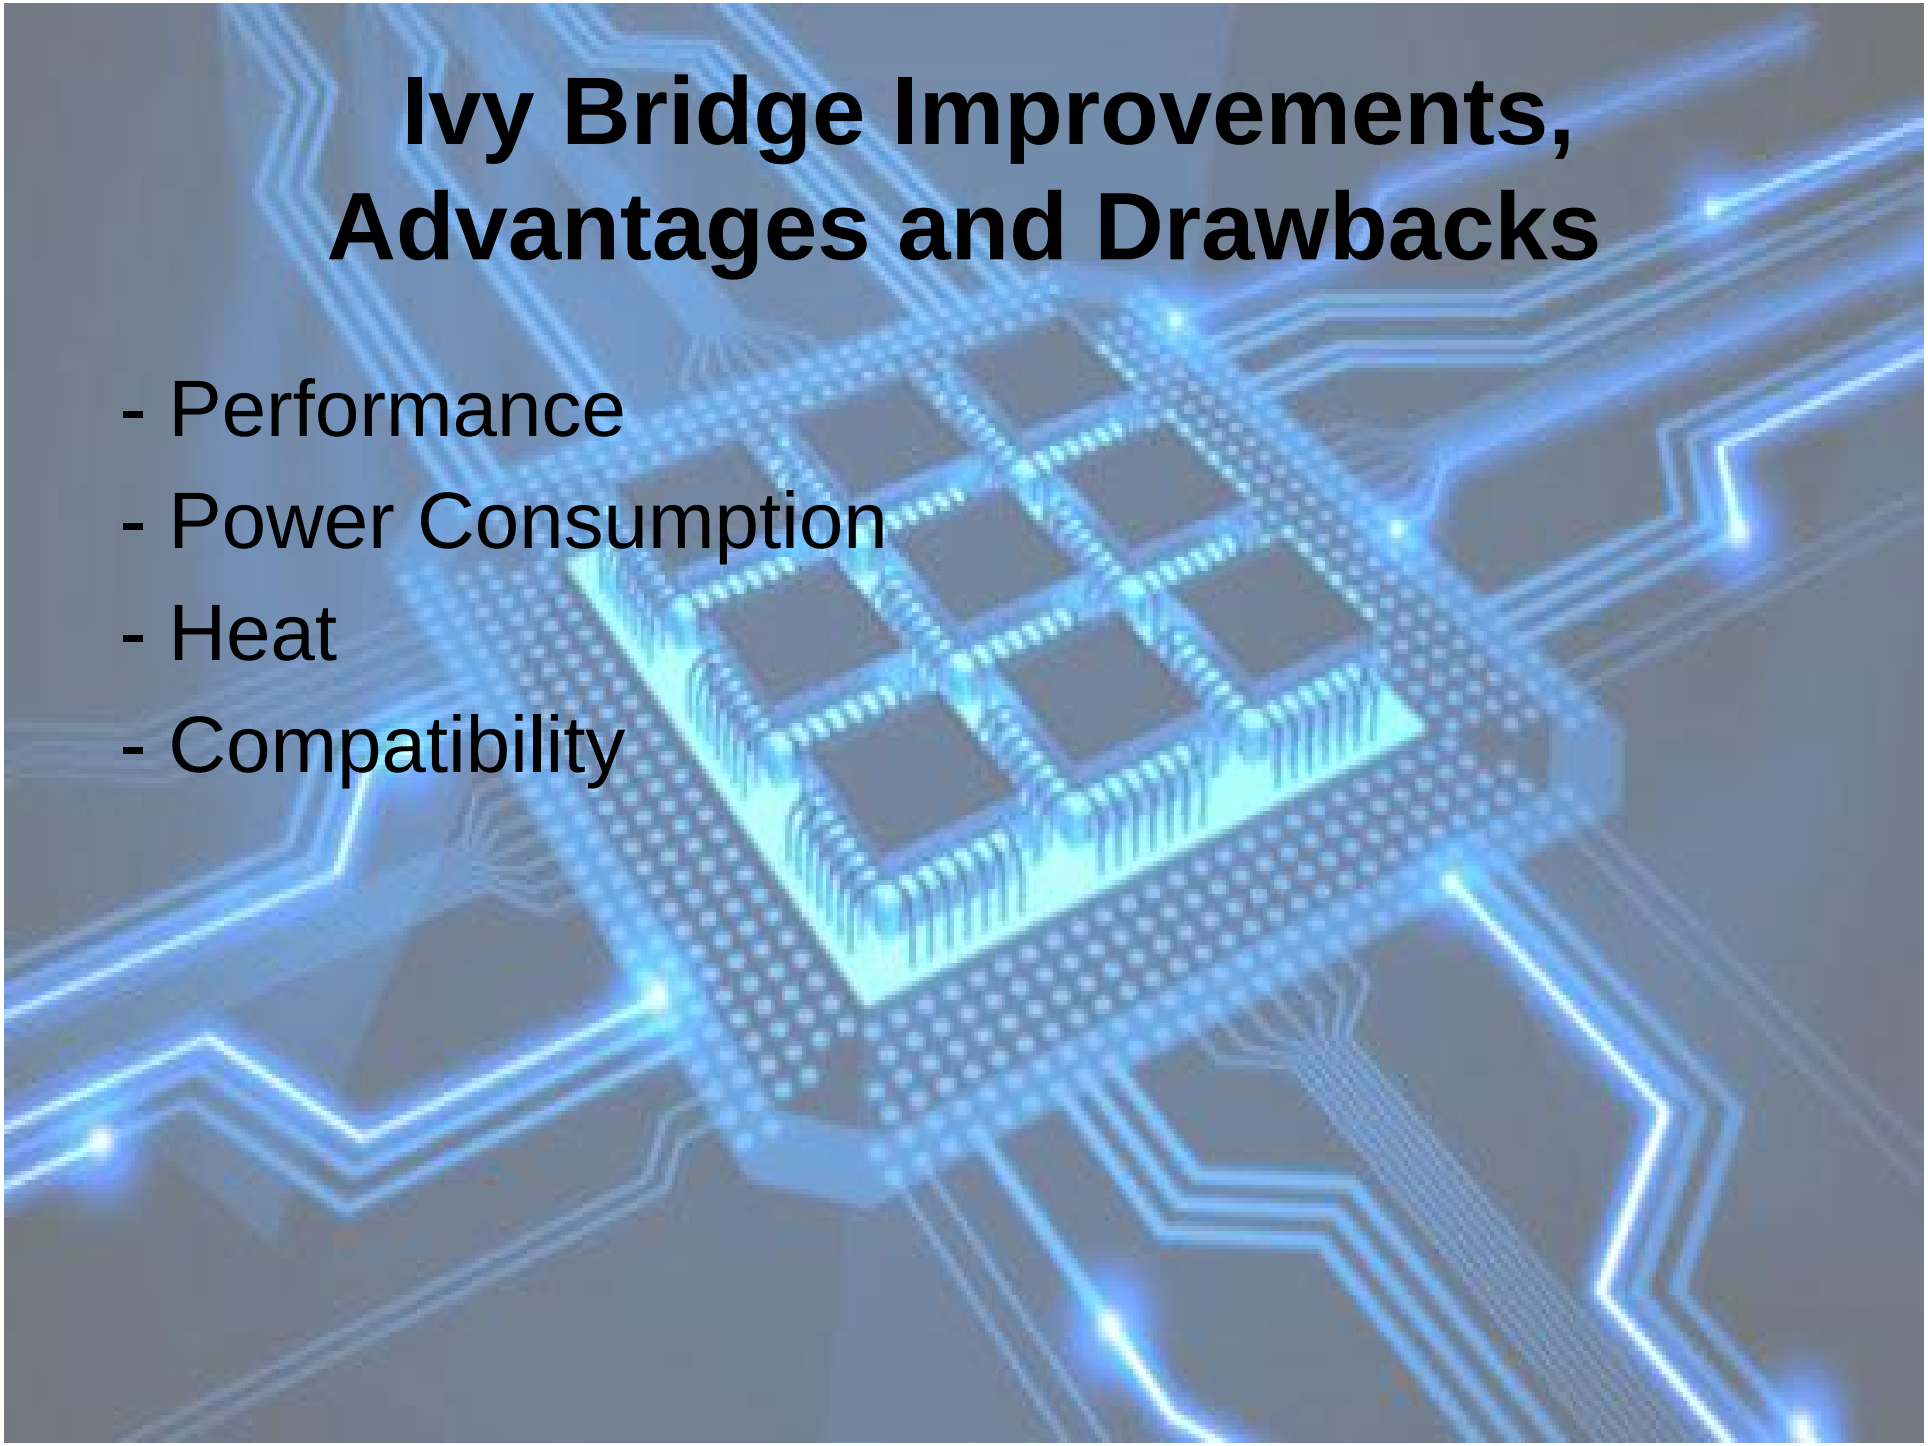

Ivy Bridge Improvements, Advantages and Drawbacks

- Performance
- Power Consumption
- Heat
- Compatibility

Performance

- 5% to 15% increase in CPU performance
- 20% to 50% increase in integrated GPU performance

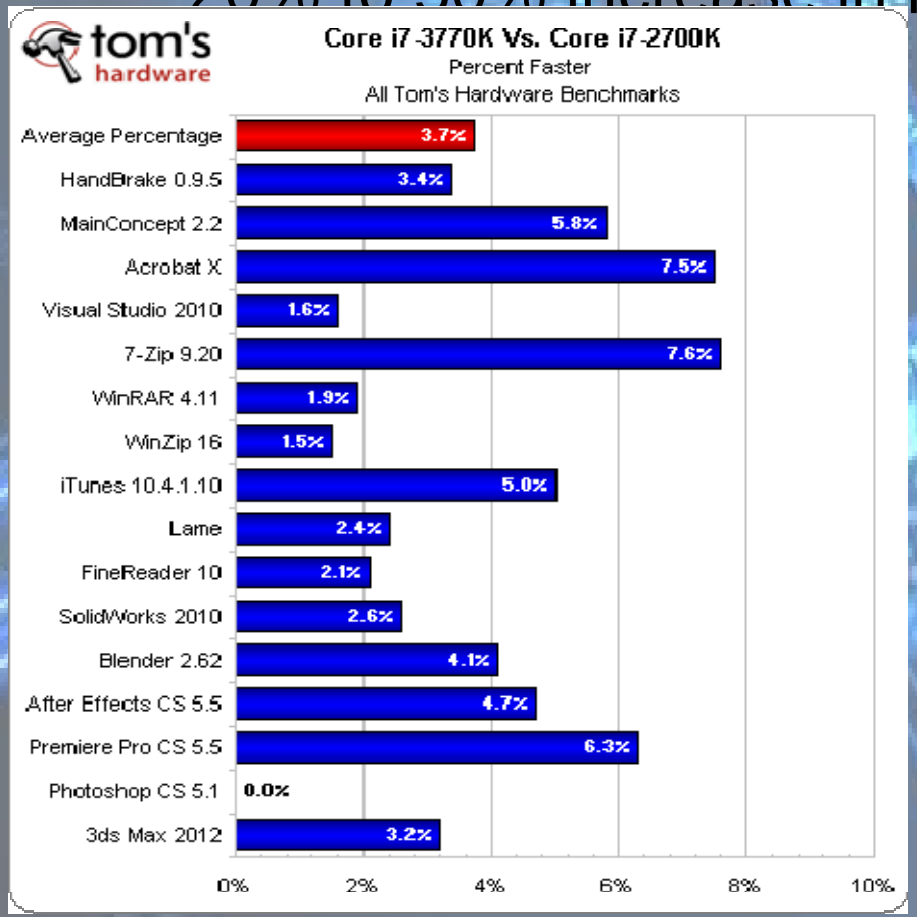

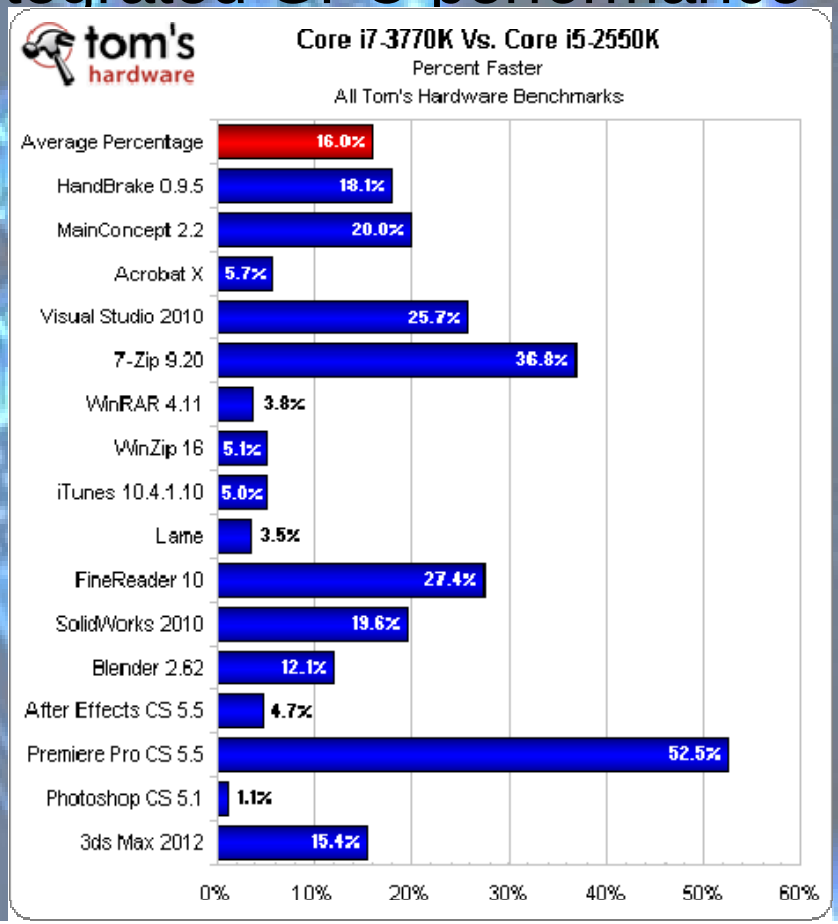

More test bench results

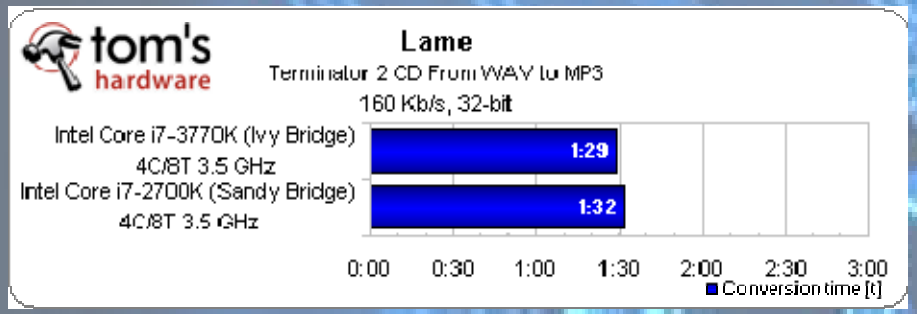

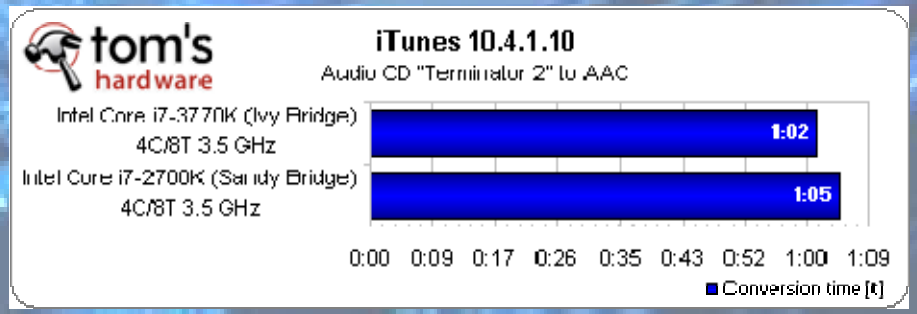

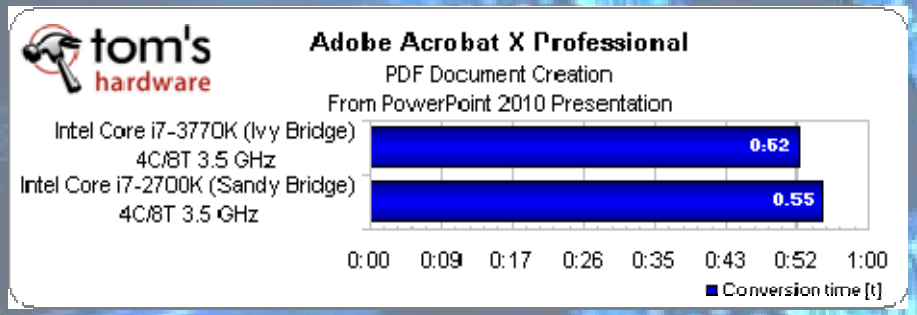

Power Consumption / Heat

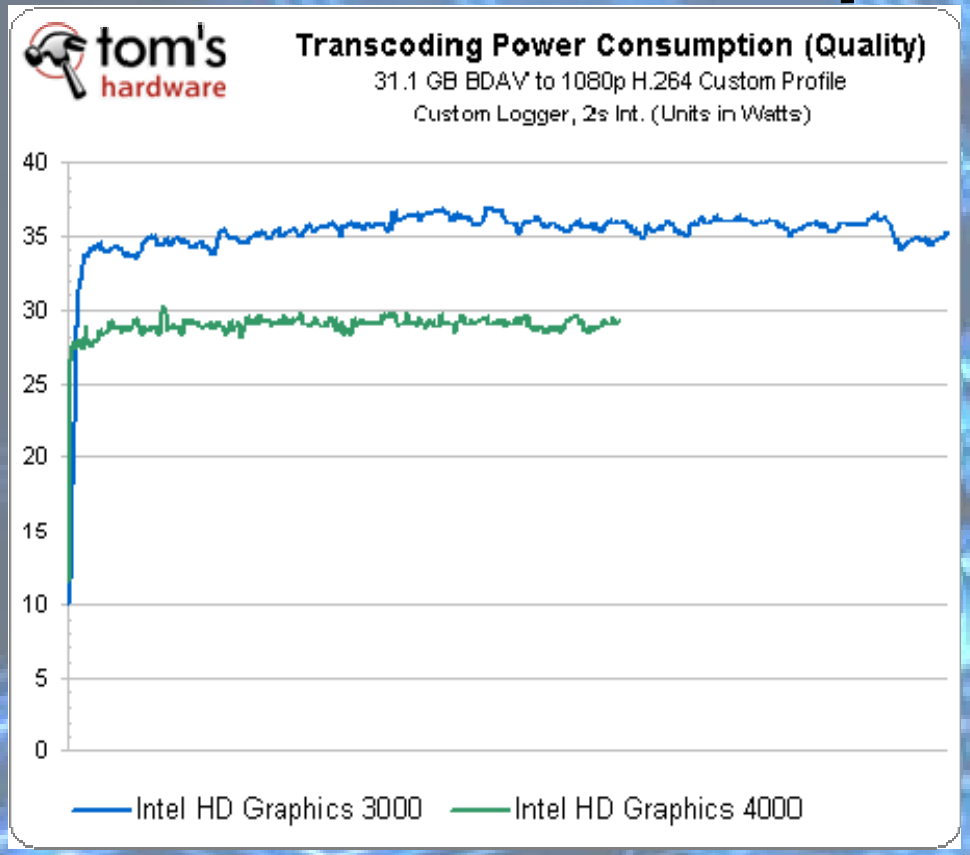

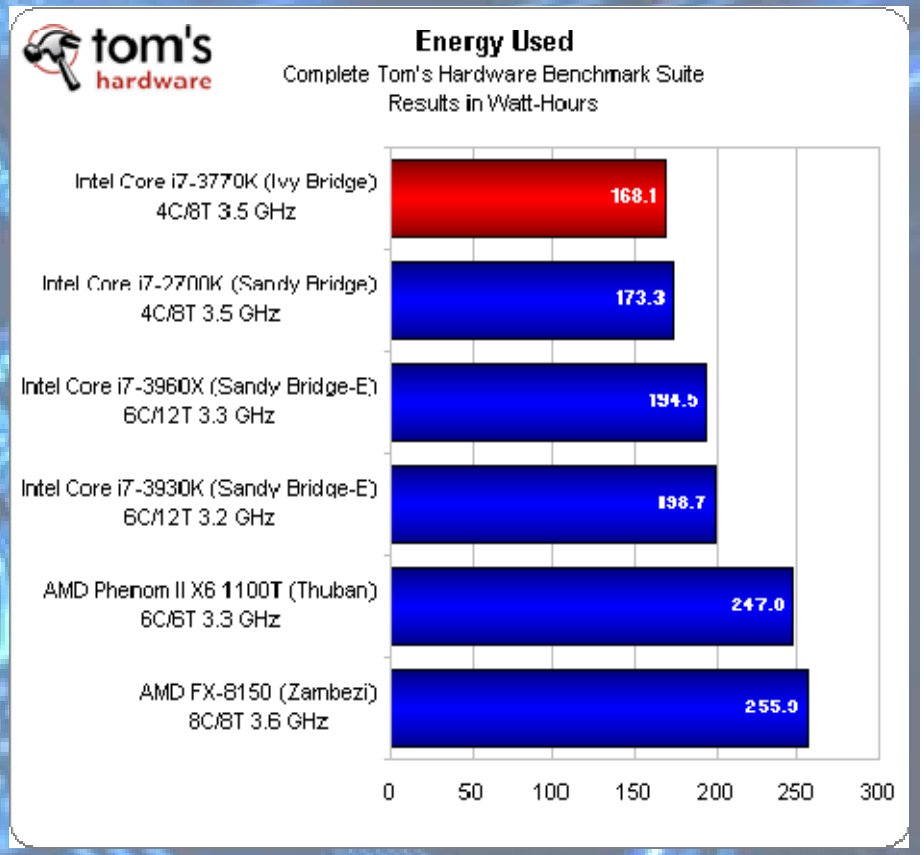

Ivy bridge idle temp: 26-35 C

Sandy bridge idle temp: 29-34C

Heat / Power when overclocked

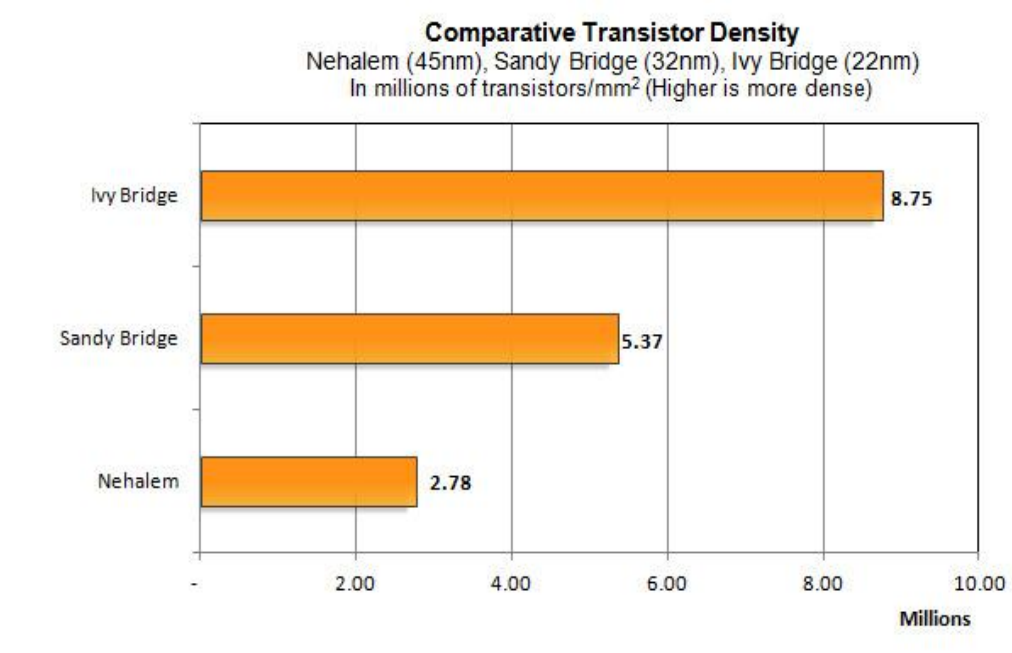

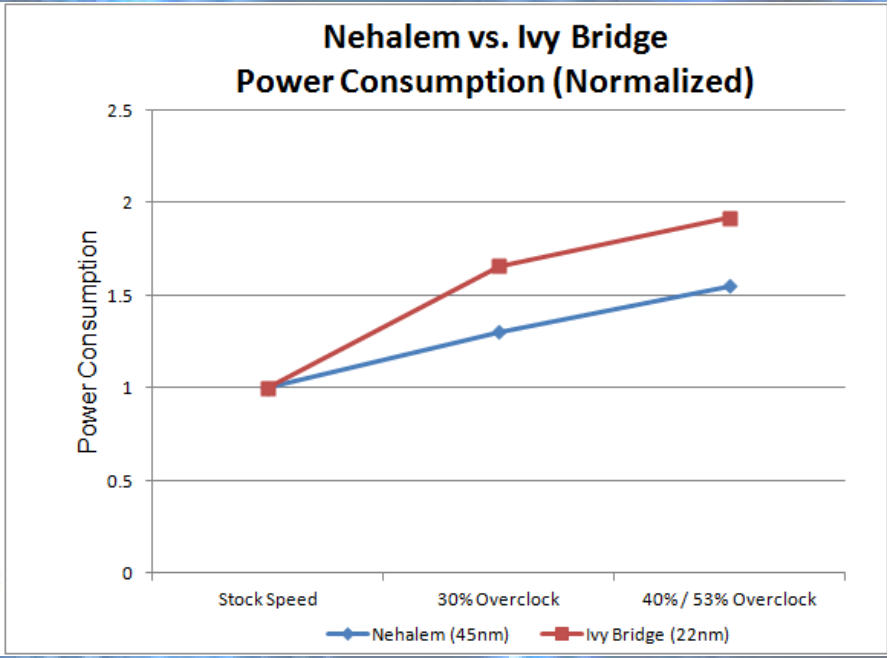

Higher temperatures result from a smaller die size as well as a larger transistor density. Also the increased transistor density results in a larger overall resistance, which means more power is required to overclock the chip.

Compatibility

The ivy bridge was designed to be compatible with the existing SAndy bridge chipsets

with the exception of Q65, Q67, and B65

Compatibility cont - Sandy bridge chipsets

Name	B65[1]	H61	Q67[1]	H67	P67	Z68
Maximum DDR3 slots	4	2	4	4	4	4
Overclocking	GPU	GPU	GPU	GPU	CPU + RAM	CPU + GPU + RAM
Allows using built-in GPU	Yes	Yes	Yes	Yes	No	Yes
RAID	No	No	Yes	Yes	Yes	Yes
Maximum USB 2.0 ports[2]	12	10	14	14	14	14
Maximum SATA 2.0/3.0 ports	4 / 1	4 / 0	4 / 2	4 / 2	4 / 2	4 / 2
Main PCIe Configuration	1 × PCIe 2.0 ×16	1 × PCIe 2.0 ×16	1 x PCIe 2.0 x16	1 × PCIe 2.0 ×16	1 × PCIe 2.0 ×16 or 2 × PCIe 2.0 ×8	1 × PCIe 2.0 ×16 or 2 × PCIe 2.0 ×8
Secondary PCIe	8 × PCIe 2.0	6 × PCIe 2.0	8 x PCIe 2.0	8 × PCIe 2.0	8 × PCIe 2.0	8 × PCIe 2.0
Intel Rapid Storage Technology	No	No	Yes	Yes[5]	Yes	Yes
Intel Active Management Technology	No	No	Yes	No	No	No
Intel Trusted Execution Technology	No	No	Yes	No	No	No
Smart Response Technology	No	No	No	No	No	Yes[6]
Ivy Bridge Processor Support	No	Yes	No	Yes	Yes	Yes
Release Date	February 2011	February 2011	May 2011	January 2011	January 2011	May 2011[6]
Max TDP	6.1 W					
Chipset lithography	65 nm					

Compatibility cont - ivy bridge chipsets

Name	B75	Q75	Q77	H77	Z75	Z77
Overclocking	CPU(Bclk) + GPU	CPU(Bclk) + GPU	CPU(Bclk) + GPU	GPU	CPU + GPU + RAM	CPU + GPU + RAM
Allows using built-in GPU	Yes	Yes	Yes	Yes	Yes	Yes
RAID	Yes	Yes	Yes	Yes	Yes	Yes
Maximum USB 2.0/3.0 ports	8 / 4	10 / 4	10 / 4	10 / 4	10 / 4	10 / 4
Maximum SATA 2.0/3.0 ports	5 / 1	5 / 1	4 / 2	4 / 2	4 / 2	4 / 2
Main PCIe Configuration	1 × PCIe 3.0 ×16	1 × PCIe 3.0 ×16	1 × PCIe 3.0 ×16	1 × PCIe 3.0 ×16	1 × PCIe 3.0 ×16 or 2 × PCIe 3.0 ×8	1 × PCIe 3.0 ×16 or 2 × PCIe 3.0 ×8 or 1 × PCIe 3.0 ×8 + 2 × PCIe 3.0 ×4
Secondary PCIe	8 × PCIe 2.0	8 × PCIe 2.0	8 × PCIe 2.0	8 × PCIe 2.0	8 × PCIe 2.0	8 × PCIe 2.0
PCI	Yes	Yes	Yes	No	No	No
Intel Rapid Storage Technology	No	No	Yes	Yes	Yes	Yes
Smart Response Technology	No	No	Yes	Yes	No	Yes
Release Date	April 2012[8]					
Chipset lithography	65 nm[9]					

Added Features

- Supervisory Mode Execution Protection(SMEP)
- PCI Express 3.0 support.
- Max CPU multiplier of 63 (57 for Sandy Bridge).
- RAM support up to 2800 MT/s in 200 MHz increments.
- A new random number generator and the RdRand instruction, codenamed Bull Mountain.
- DDR3L and Configurable TDP for mobile processors.
- Multiple 4K video playback.
- Intel Quick Sync Video.

Digital random number generation(DRNG)

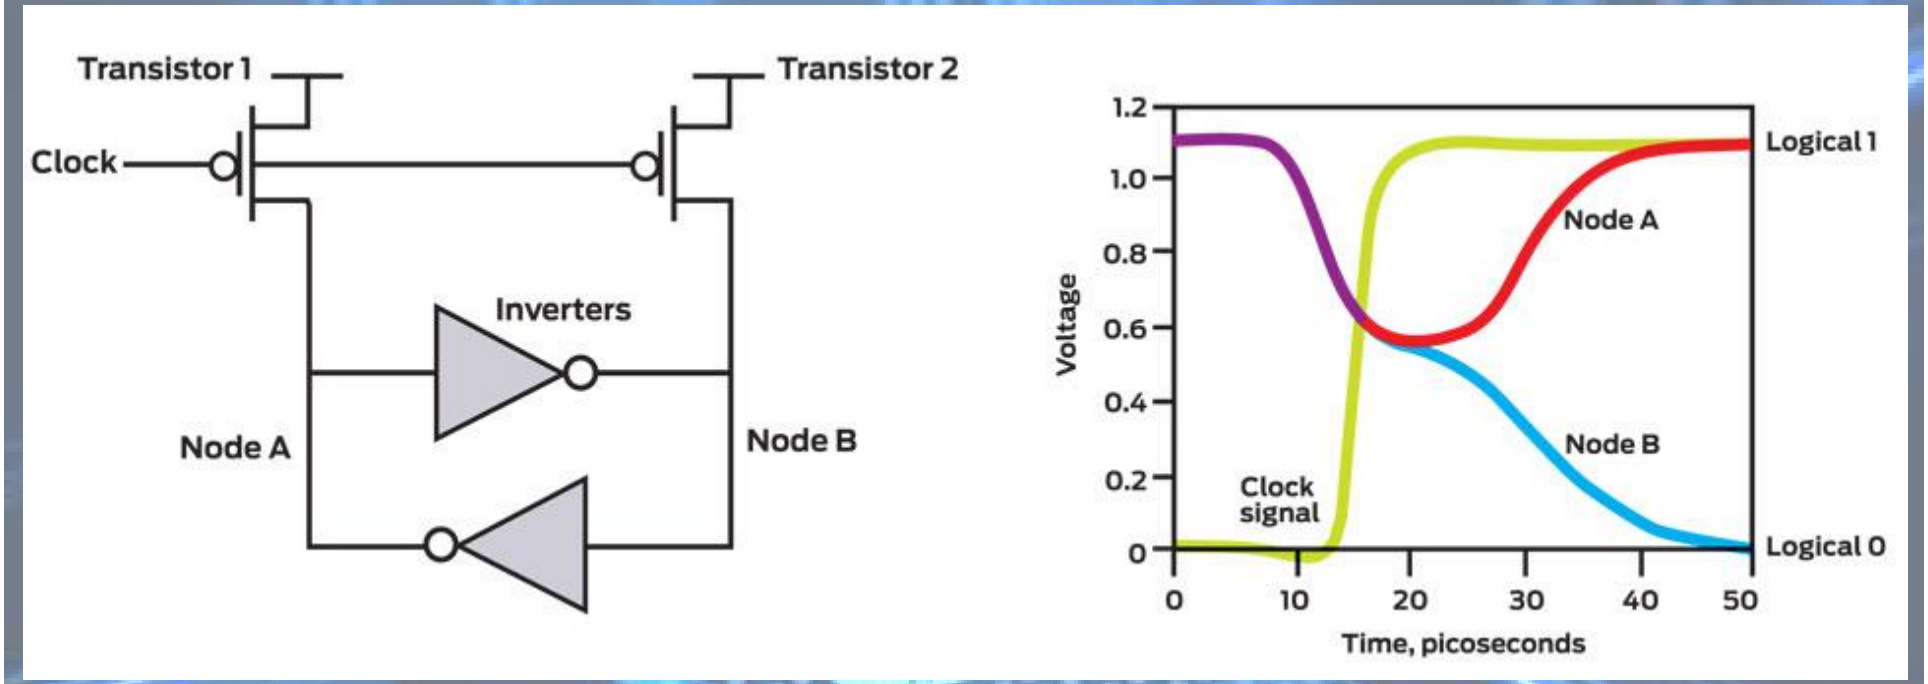

Intel Secure Key consists of a digital random number generator that creates random numbers to strengthen encryption algorithms. Intel OS Guard helps defend against privilege escalation attacks where a hacker remotely takes over another person's system. These two features join existing platform security features such as Intel Identity Protection Technology (Intel IPT) and Intel Anti-Theft technology (Intel AT) to help make Intel platforms some of the most secure in the industry.

Supervisory Mode Execution Protection(SMEP)

Supervisory Mode Execute Protection (SMEP)

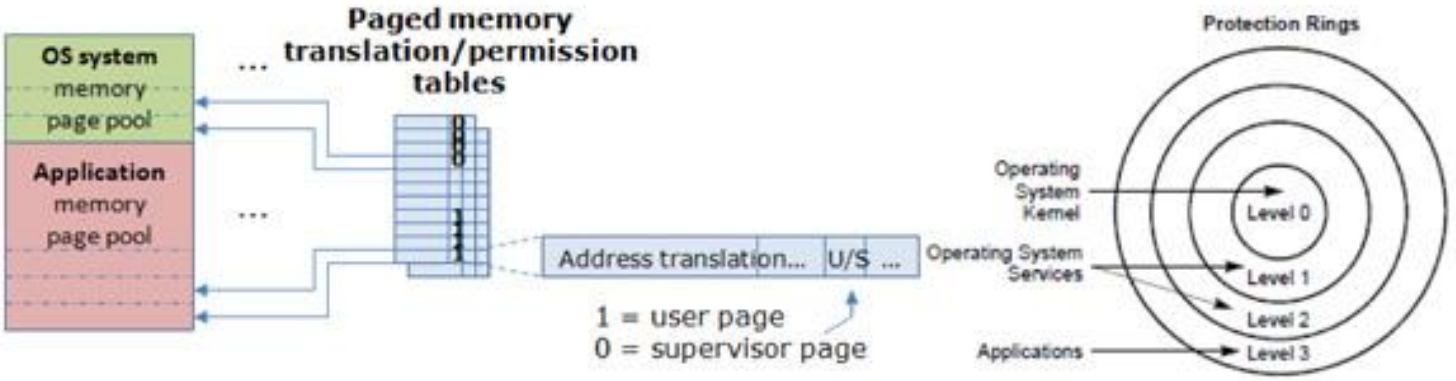

- **Ivy Bridge introduces SMEP to help prevent Escalation of Privilege (EoP) security attacks**
  - Prevents execution out of untrusted application memory while operating at a more privileged level
  - If CR4.SMEP set to 1 and in supervisor mode (CPL<3), instructions may not be executed from a linear address for which the user mode flag is 1
  - Available in both 32- and 64-bit operating modes
  - SMEP is enumerated via CPUID.7.0.EBX[7]

References

http://vr-zone.com/articles/ivy-bridge-should-work-in-h67-and-p67-motherboards/11077.html
http://www.bit-tech.net/hardware/cpus/2011/10/10/all-about-ivy-bridge/6
http://newsroom.intel.com/community/intel_newsroom/blog/2011/05/04/intel-reinvents-transistors-using-new-3-d-structure
http://www.anandtech.com/show/4790/ivy-bridge-overclocking-ratio-changes-without-reboot-more-ratios-and-ddr32800
http://www.anandtech.com/show/4798/ivy-bridge-148b-transistors
http://www.tomshardware.com/news/ivy-bridge-overclocking-high-temp,15512.html
http://www.nordichardware.com/news/69-cpu-chipset/43332-intel-ivy-bridge-gets-variable-tdp-and-thunderbolt.html
http://www.intel.com/content/www/us/en/silicon-innovations/intel-22nm-technology.html
http://www.pcworld.com/article/253992/intels_first_ivy_bridge_chips_wont_be_for_ultrabooks.html
http://www.pcmag.com/article2/0,2817,2403407,00.asp
http://vr-zone.com/articles/ivy-bridge-vs-sandy-bridge--4.8ghz-quad-core-cpu-showdown/15637.html
http://www.techspot.com/guides/502-intel-ivy-bridge/
http://scoop.intel.com/top-5-new-features-ivy-bridge-processors/
http://www.bbc.co.uk/news/technology-17785464
http://lenzfire.com/2011/09/intel-introduces-smep-for-ivy-bridge-a-new-security-feature-80649/
http://en.wikipedia.org/wiki/Privilege_escalation
http://www.hardwaresecrets.com/printpage/Inside-the-Intel-Ivy-Bridge-Microarchitecture/13
http://spectrum.ieee.org/computing/hardware/behind-intels-new-randomnumber-generator/0

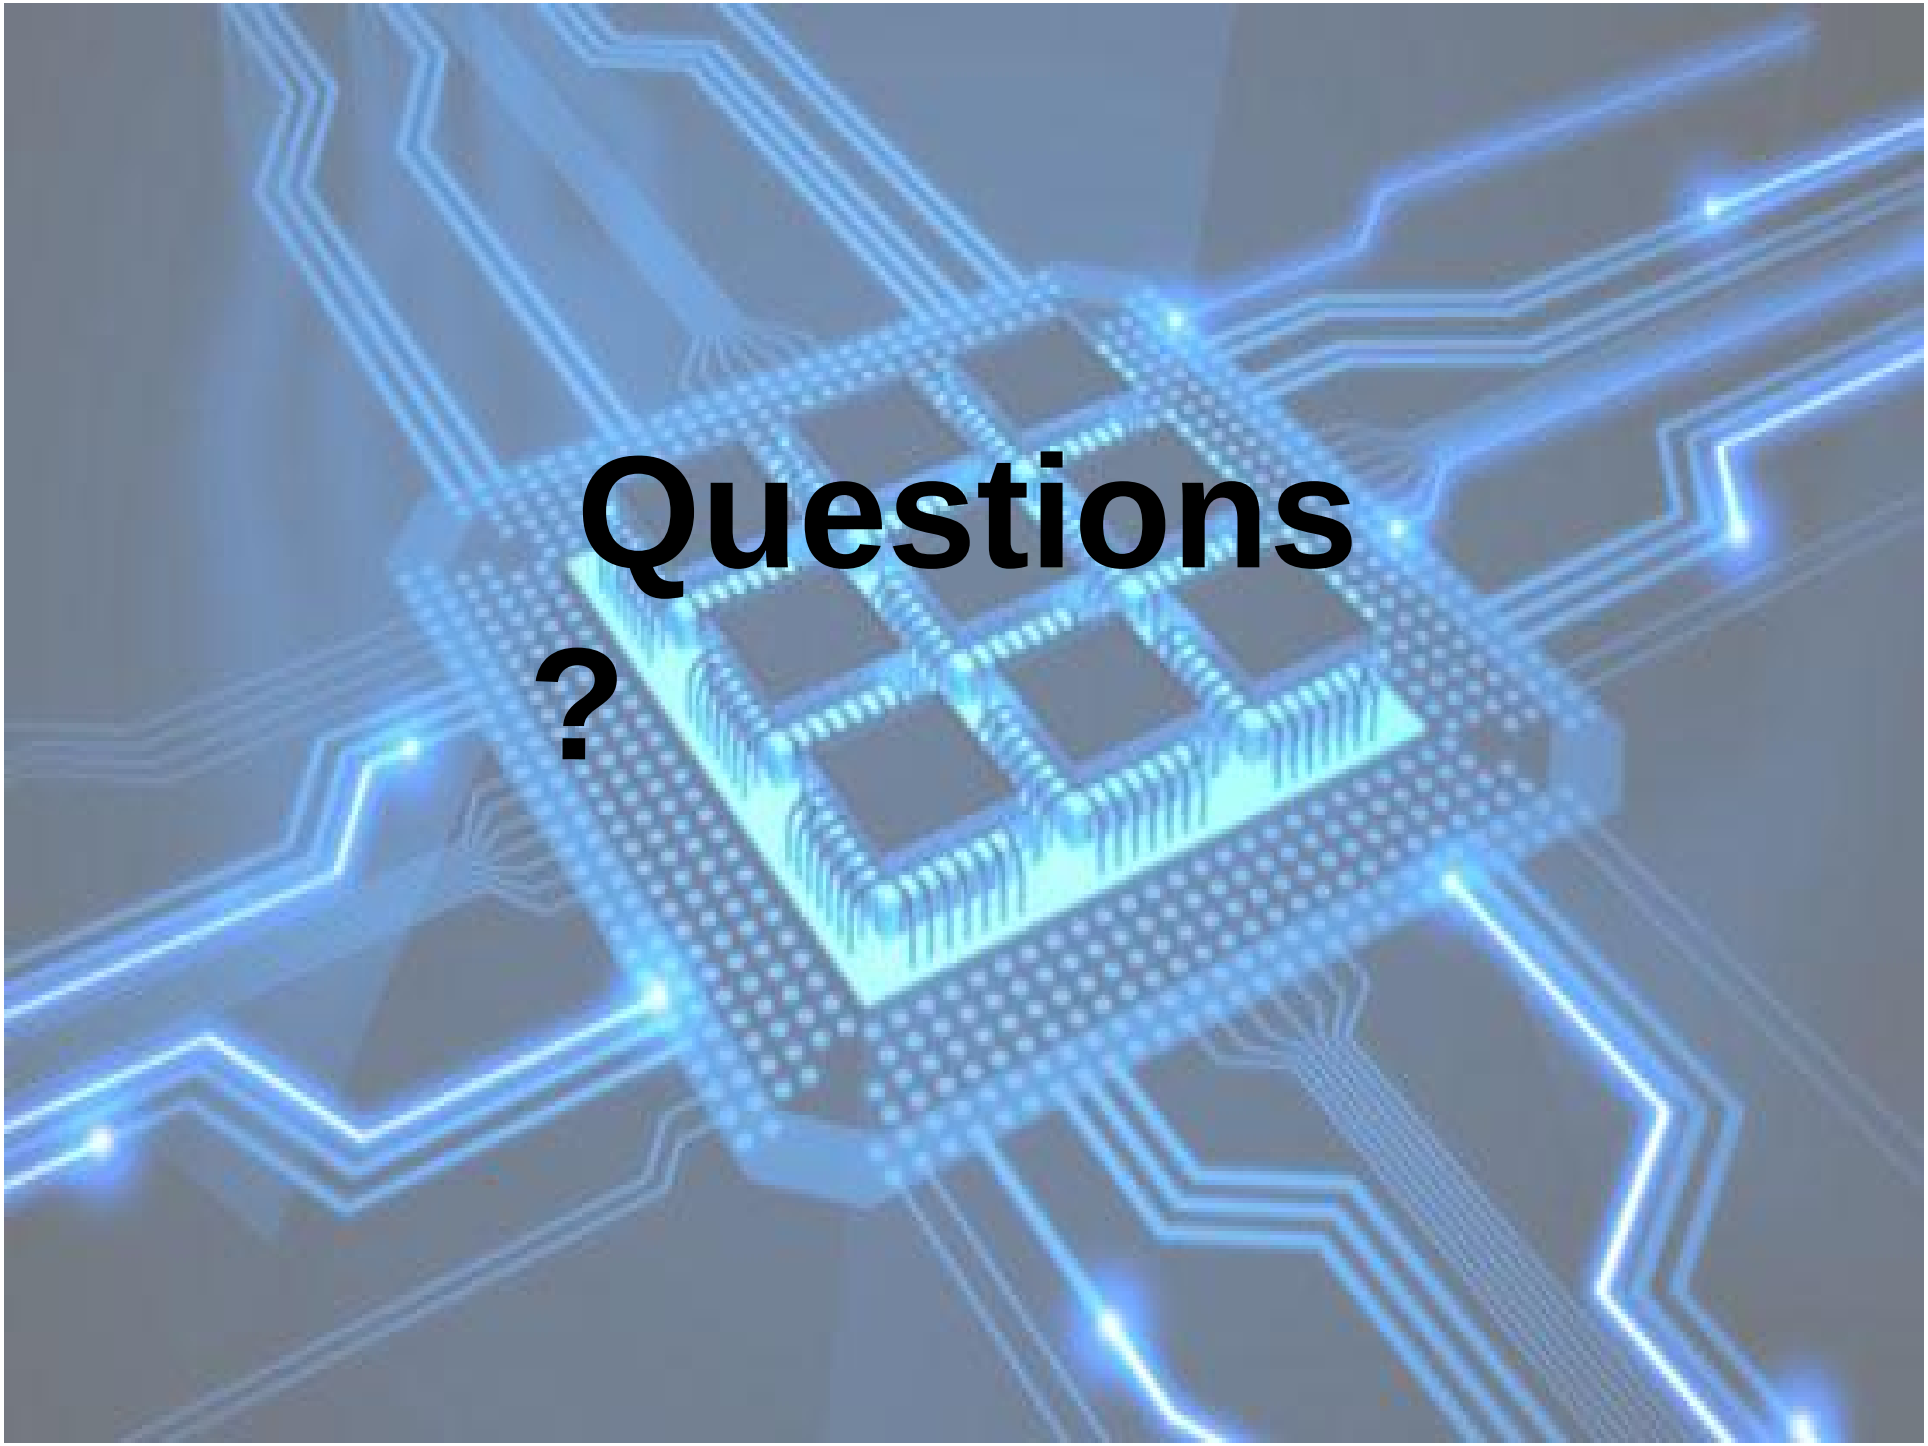
Questions
?