

**Lecture 22:**

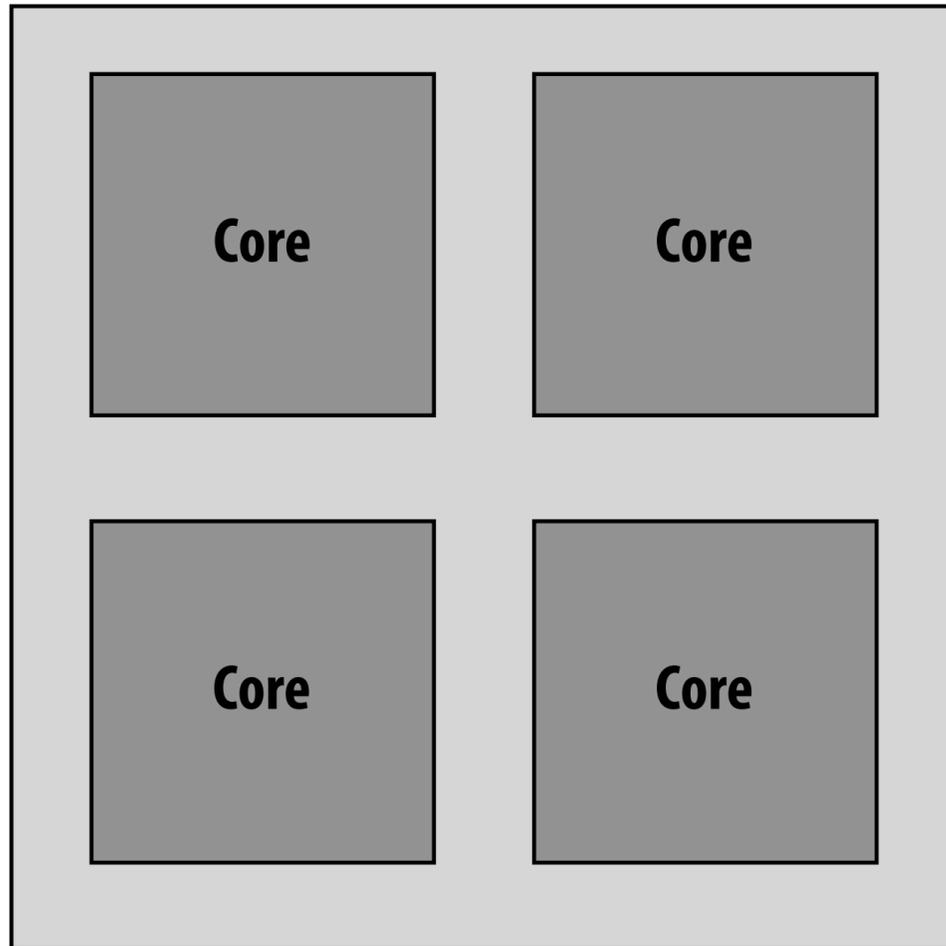
# **Heterogeneous Parallelism and Hardware Specialization**

---

**Parallel Computer Architecture and Programming**

**CMU 15-418, Spring 2013**

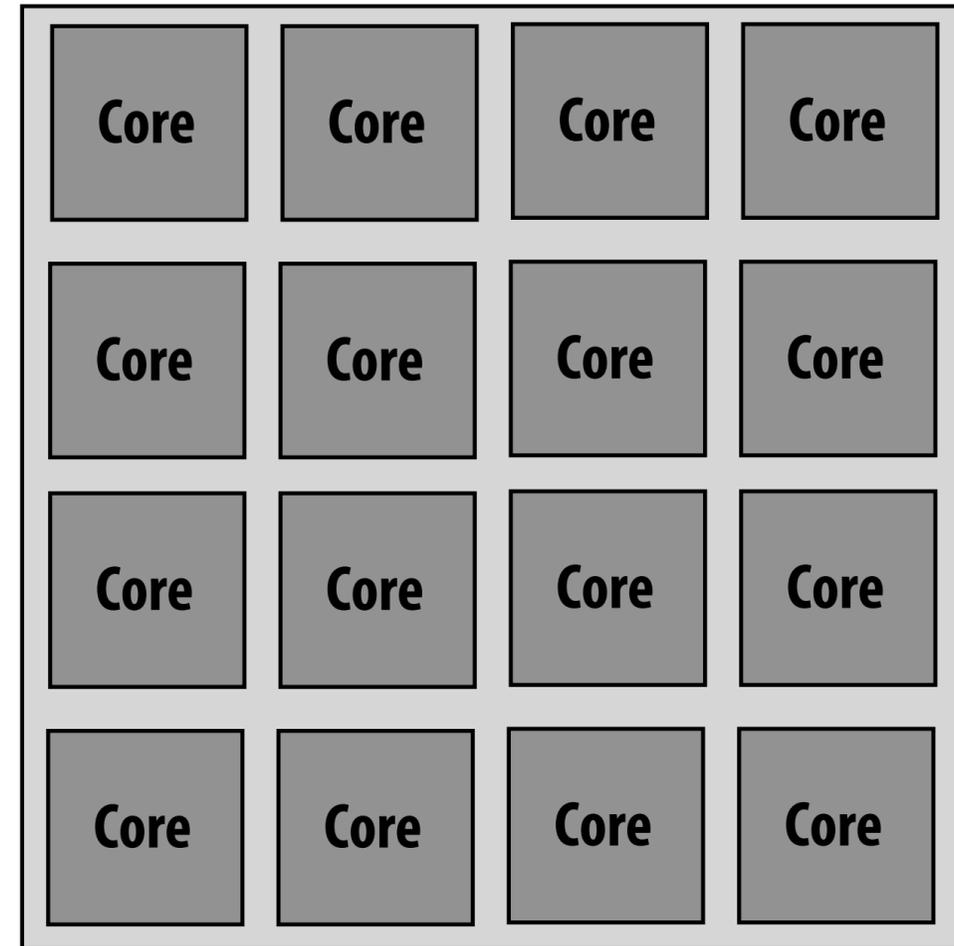
# You need to buy a computer system



**Processor A**

4 cores

Each core has sequential performance  $P$



**Processor B**

16 cores

Each core has sequential performance  $P/2$

**All other components of the system are equal.**

**Which do you pick?**

# Amdahl's law revisited

$$\textit{speedup}(f,n) = \frac{1}{(1-f) + \frac{f}{n}}$$

$f$  = fraction of program that is parallelizable

$n$  = parallel processors

**Assumptions:**

**Parallelizable work distributes perfectly onto  $n$  processors of equal capability**

# Account for resource limits

$$speedup(f, n, r) = \frac{1}{\frac{(1-f)}{perf(r)} + \frac{f \cdot r}{perf(r) \cdot n}}$$

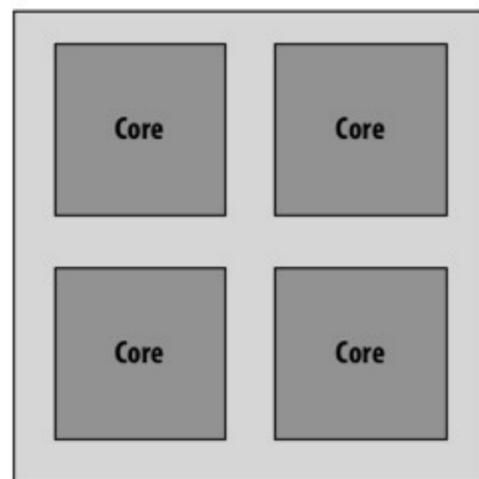
(relative to processor with  
1 unit worth of resources,  $n=1$ )

$f$  = fraction of program that is parallelizable

$n$  = total processing resources (e.g., transistors on a chip)

$r$  = resources dedicated to each processing core,

(each of the  $n/r$  cores has sequential performance  $perf(r)$ )



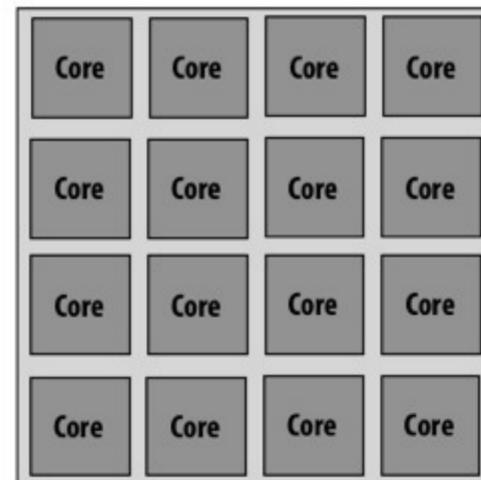
Processor A

**Example:**

Let  $n=16$

$r_A = 4$

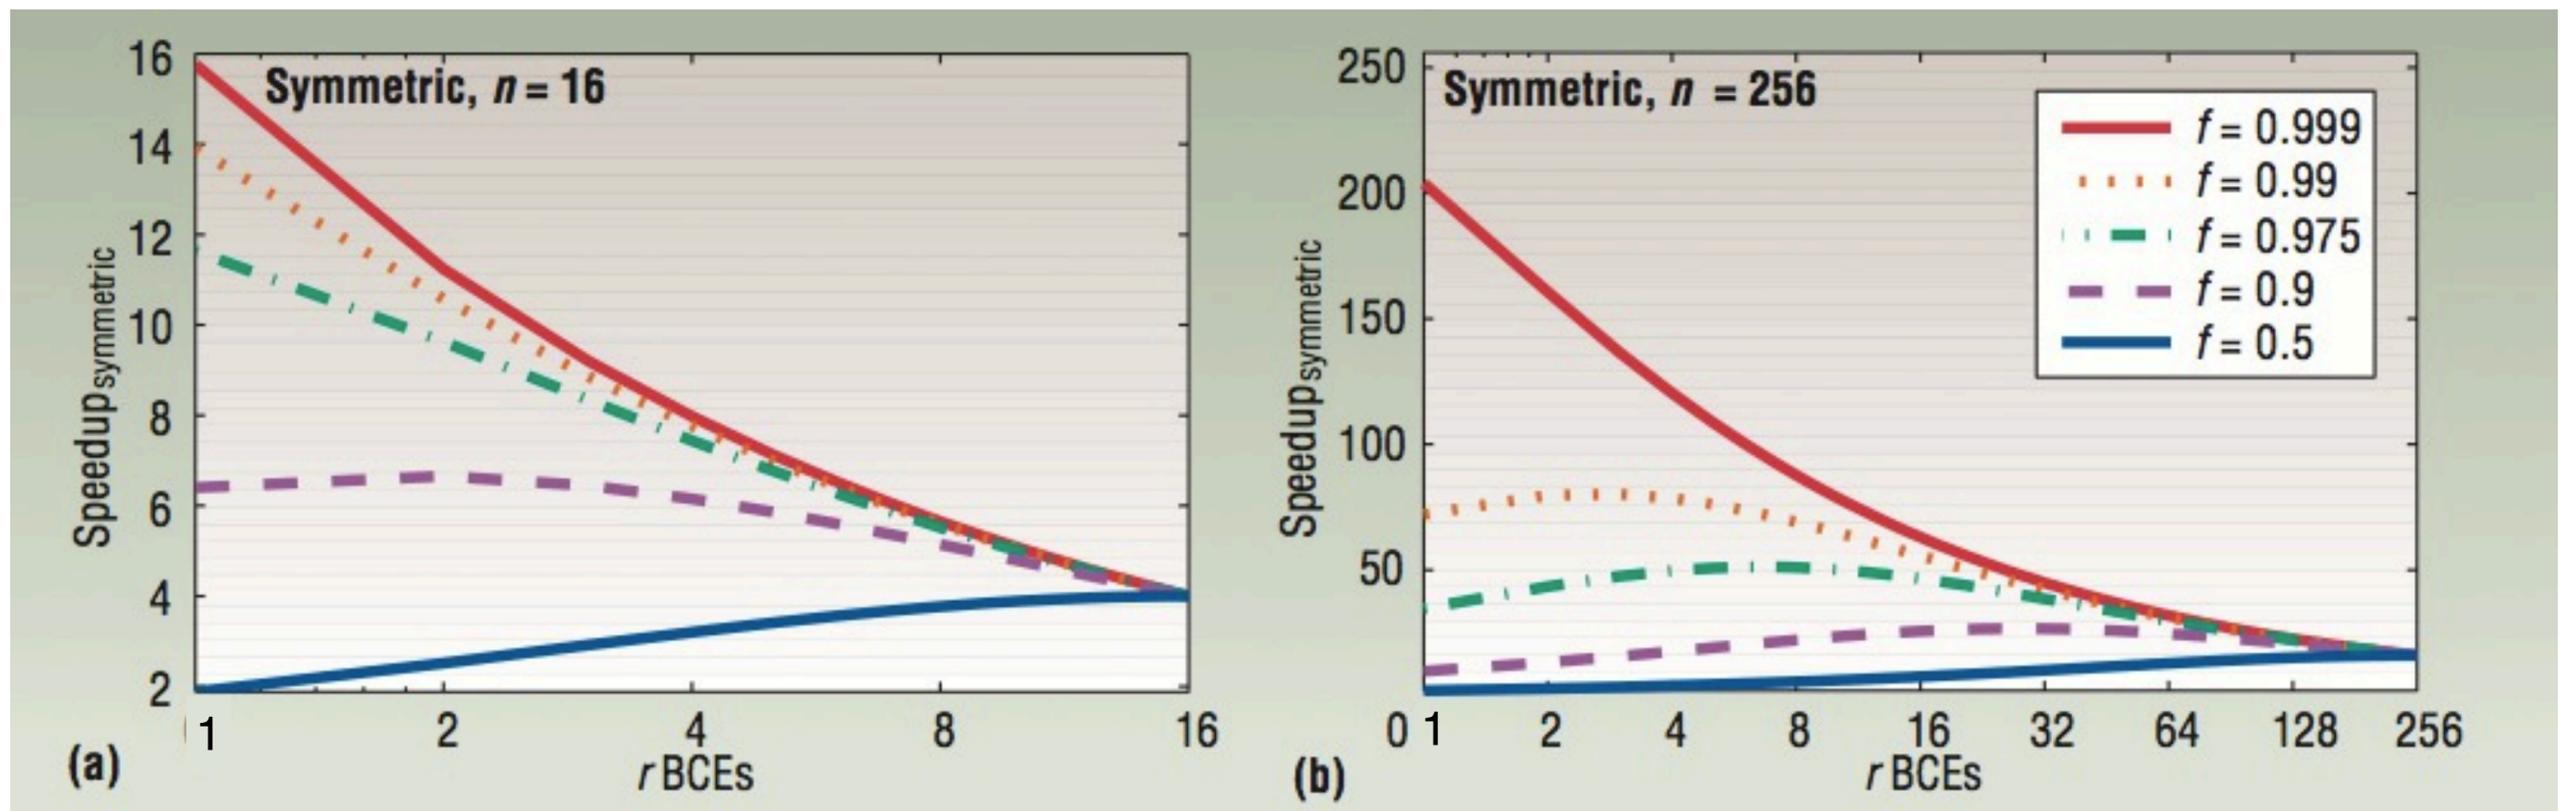
$r_B = 1$



Processor B

# Speedup (relative to n=1)

[Source: Hill and Marty 08]



Up to 16 cores ( $n=16$ )

Up to 256 cores ( $n=256$ )

Text

Each graph is iso-resources

X-axis =  $r$  (many small cores to left, fewer “fatter” cores to right)

$perf(r)$  modeled as  $\sqrt{r}$

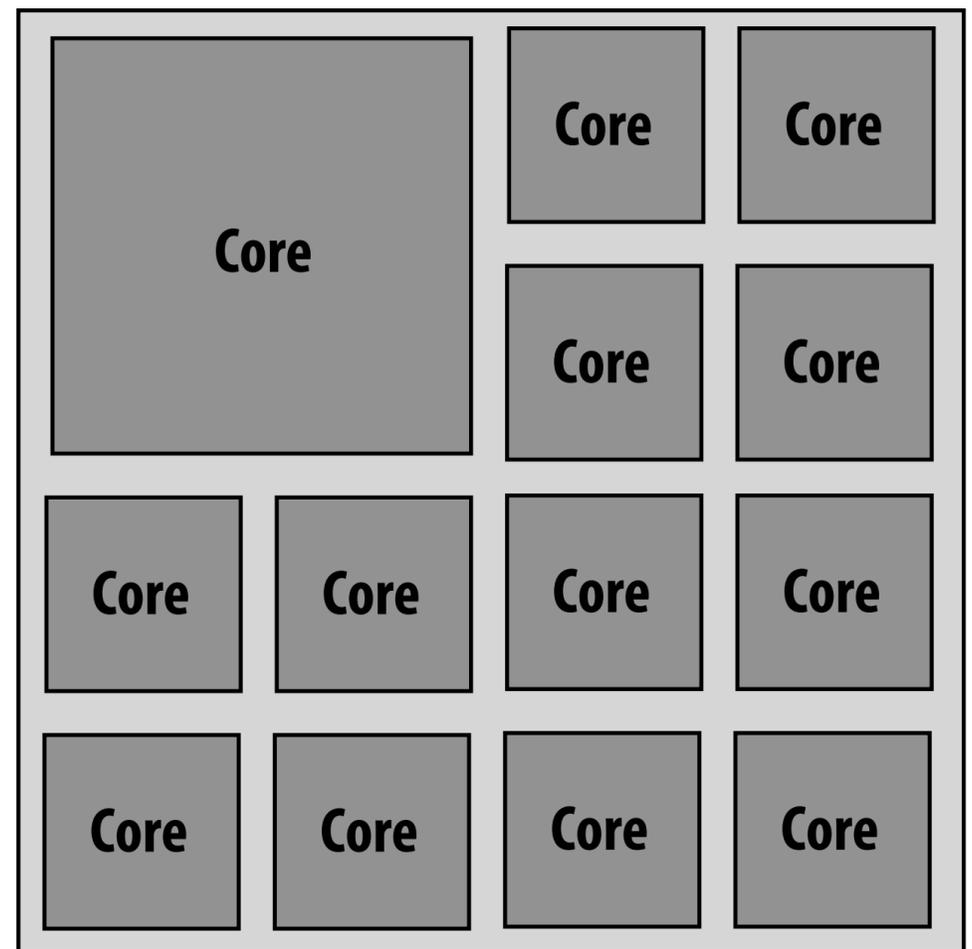
# Asymmetric processing cores

**Example:**

Let  $n=16$

One core:  $r = 4$

12 cores:  $r = 1$



$speedup(f, n, r)$

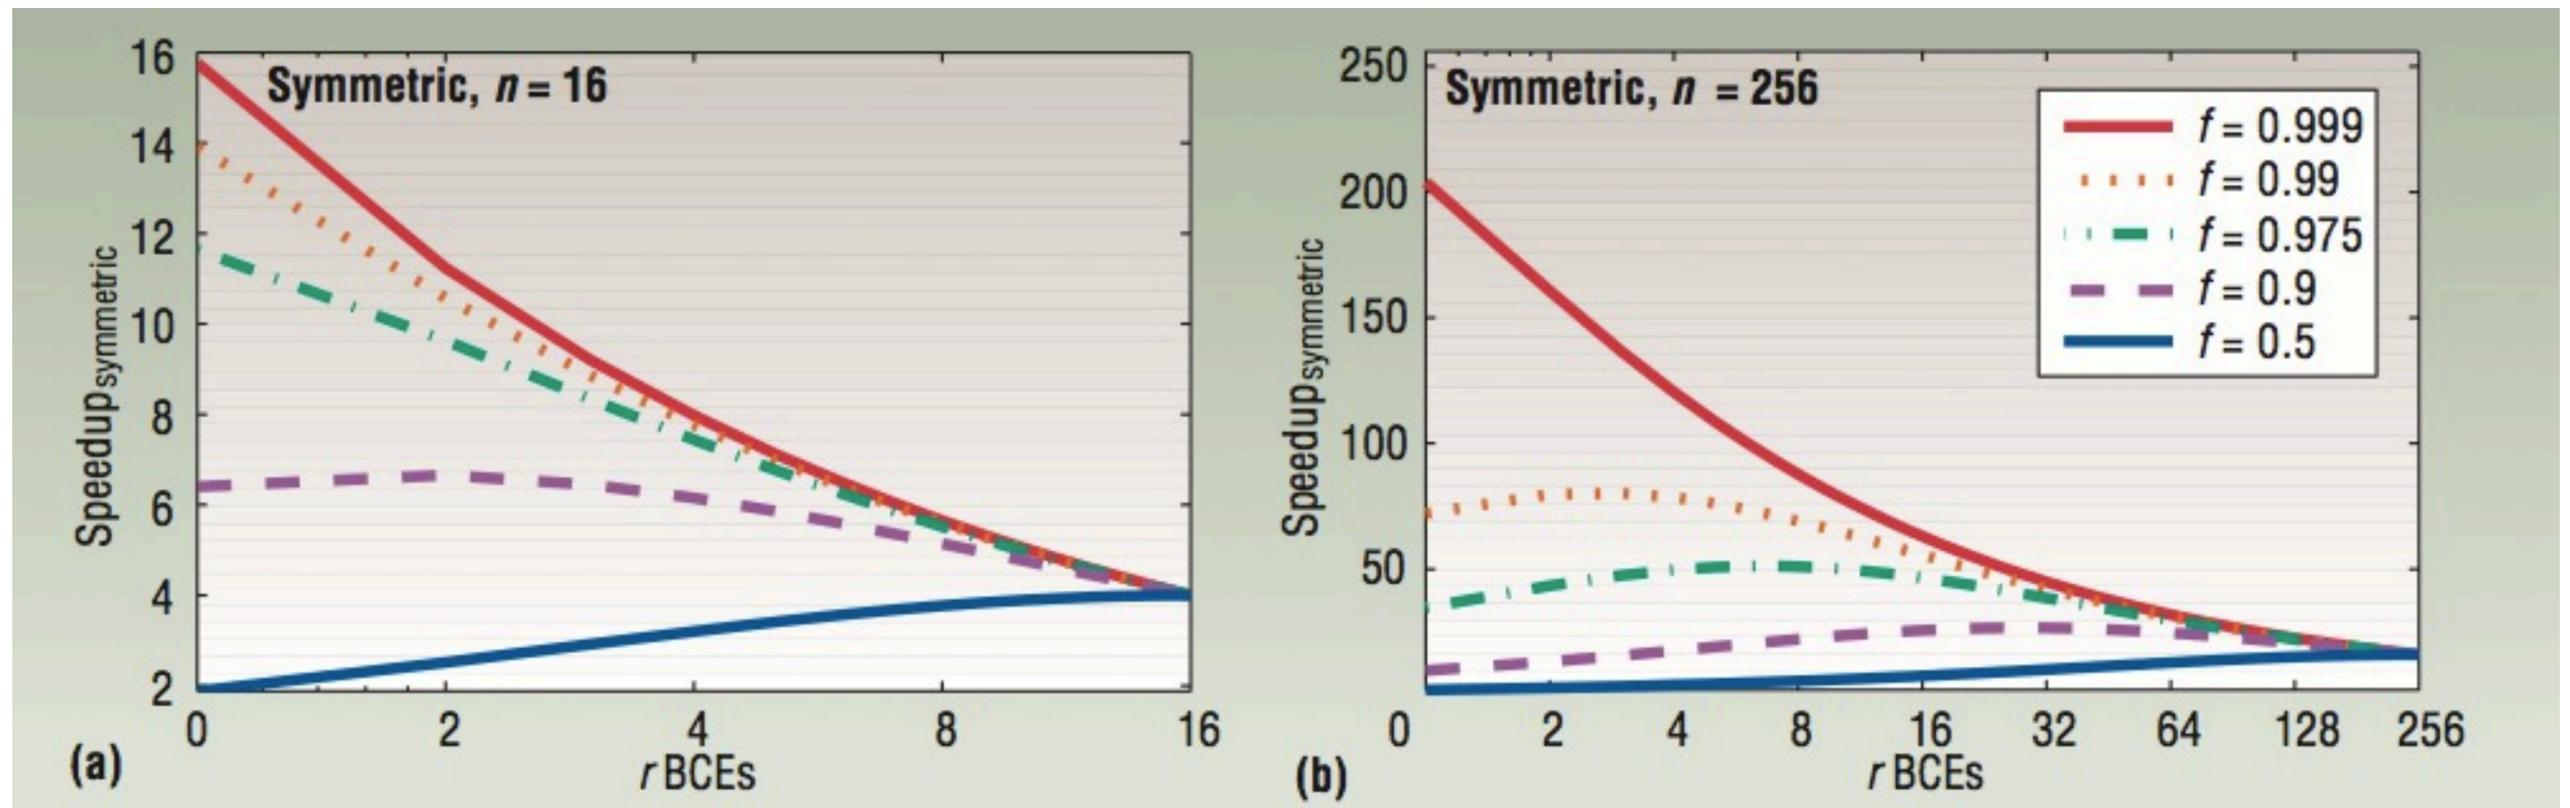
(heterogeneous processor with  $n$  resources relative to uniprocessor with one unit worth of resources,  $n=1$ )

$$= \frac{1}{\frac{(1-f)}{perf(r)} + \frac{f}{perf(r) + (n-r)}}$$

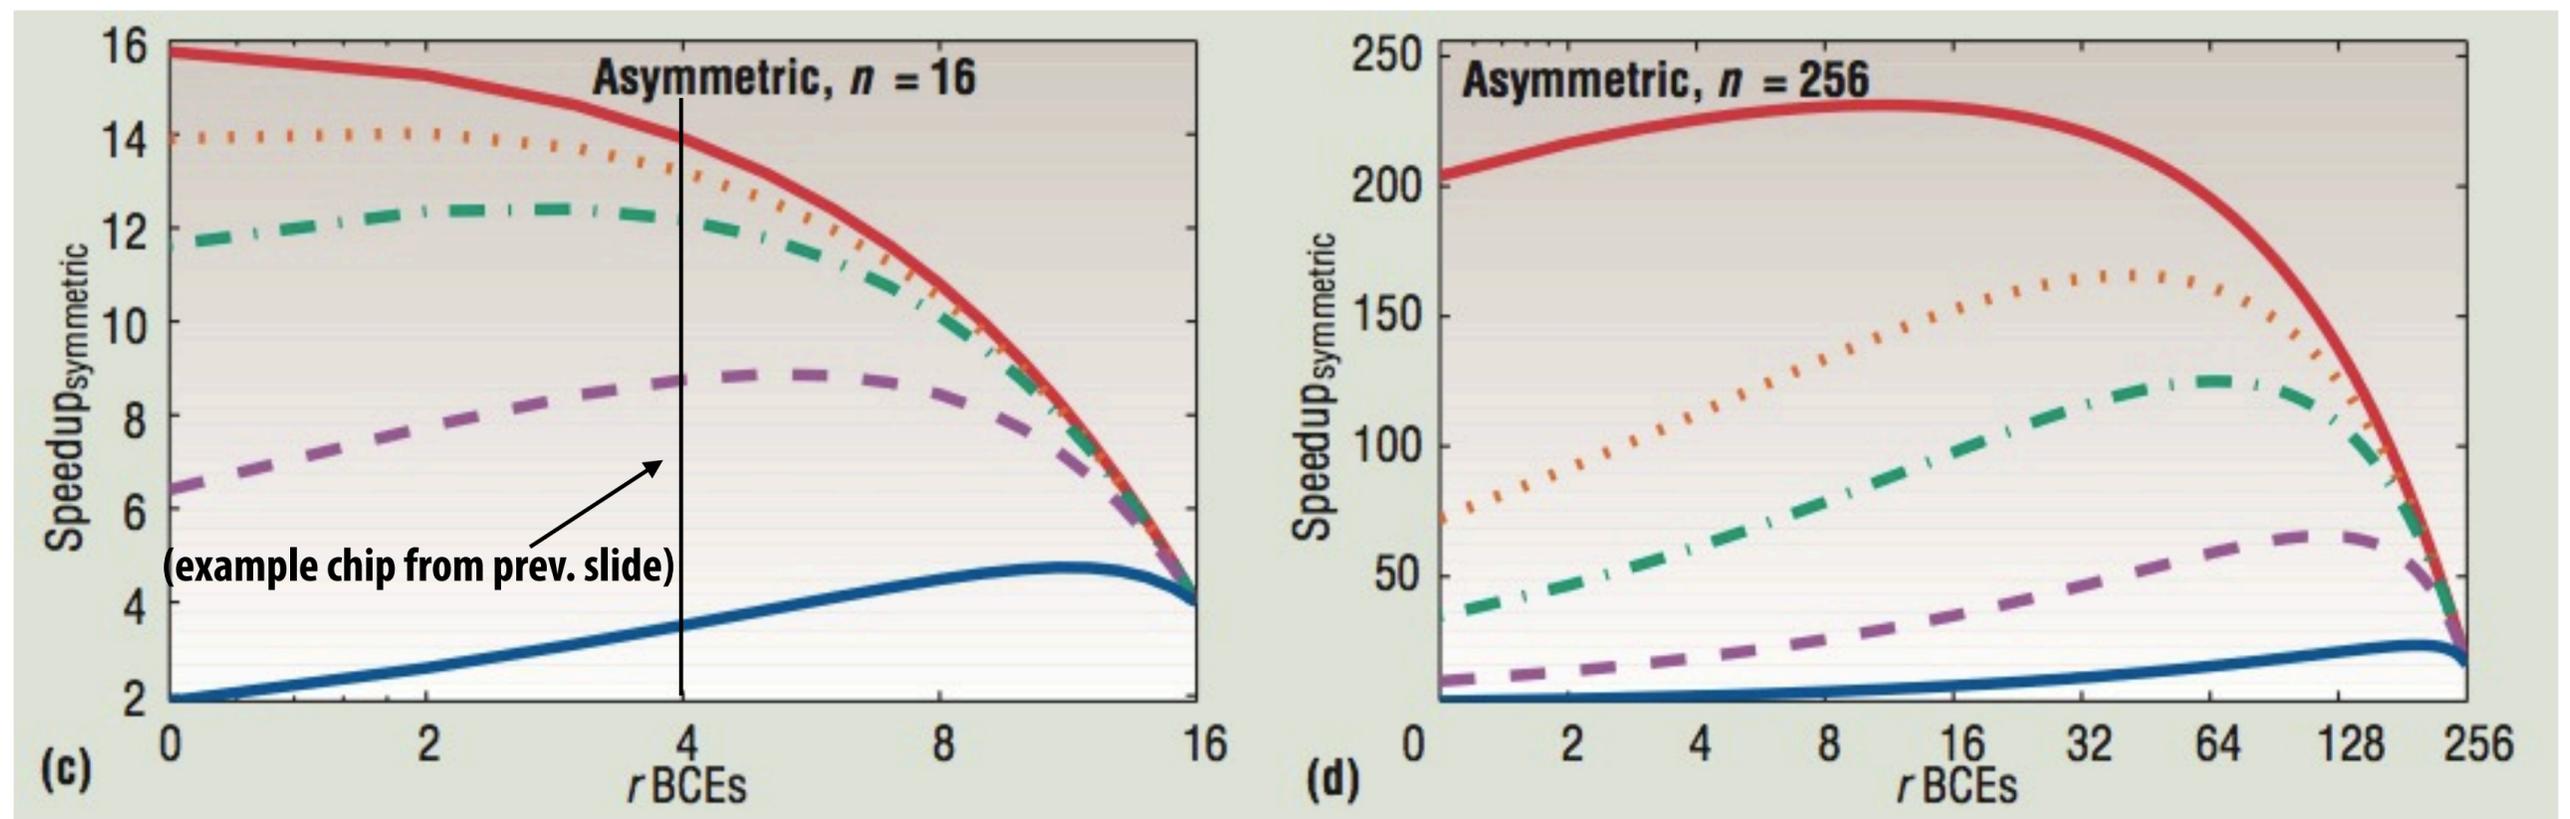
one  $perf(r)$  processor +  $n-r$   $perf=1$  processors

# Speedup (relative to $n=1$ )

[Source: Hill and Marty 08]



X-axis for symmetric architectures gives  $r$  all cores (many small cores to left)



X-axis for asymmetric architectures gives  $r$  for the single "fat" core (rest of cores are  $r=1$ )

# Heterogeneous processing

**Observation: most real applications are complex \*\***

**They have components that can be widely parallelized.**

**And components that are difficult to parallelize.**

**They have components that are amenable to wide SIMD execution.**

**And components that are not. (divergent control flow)**

**They have components with predictable data access**

**And components with unpredictable access, but those accesses might cache well.**

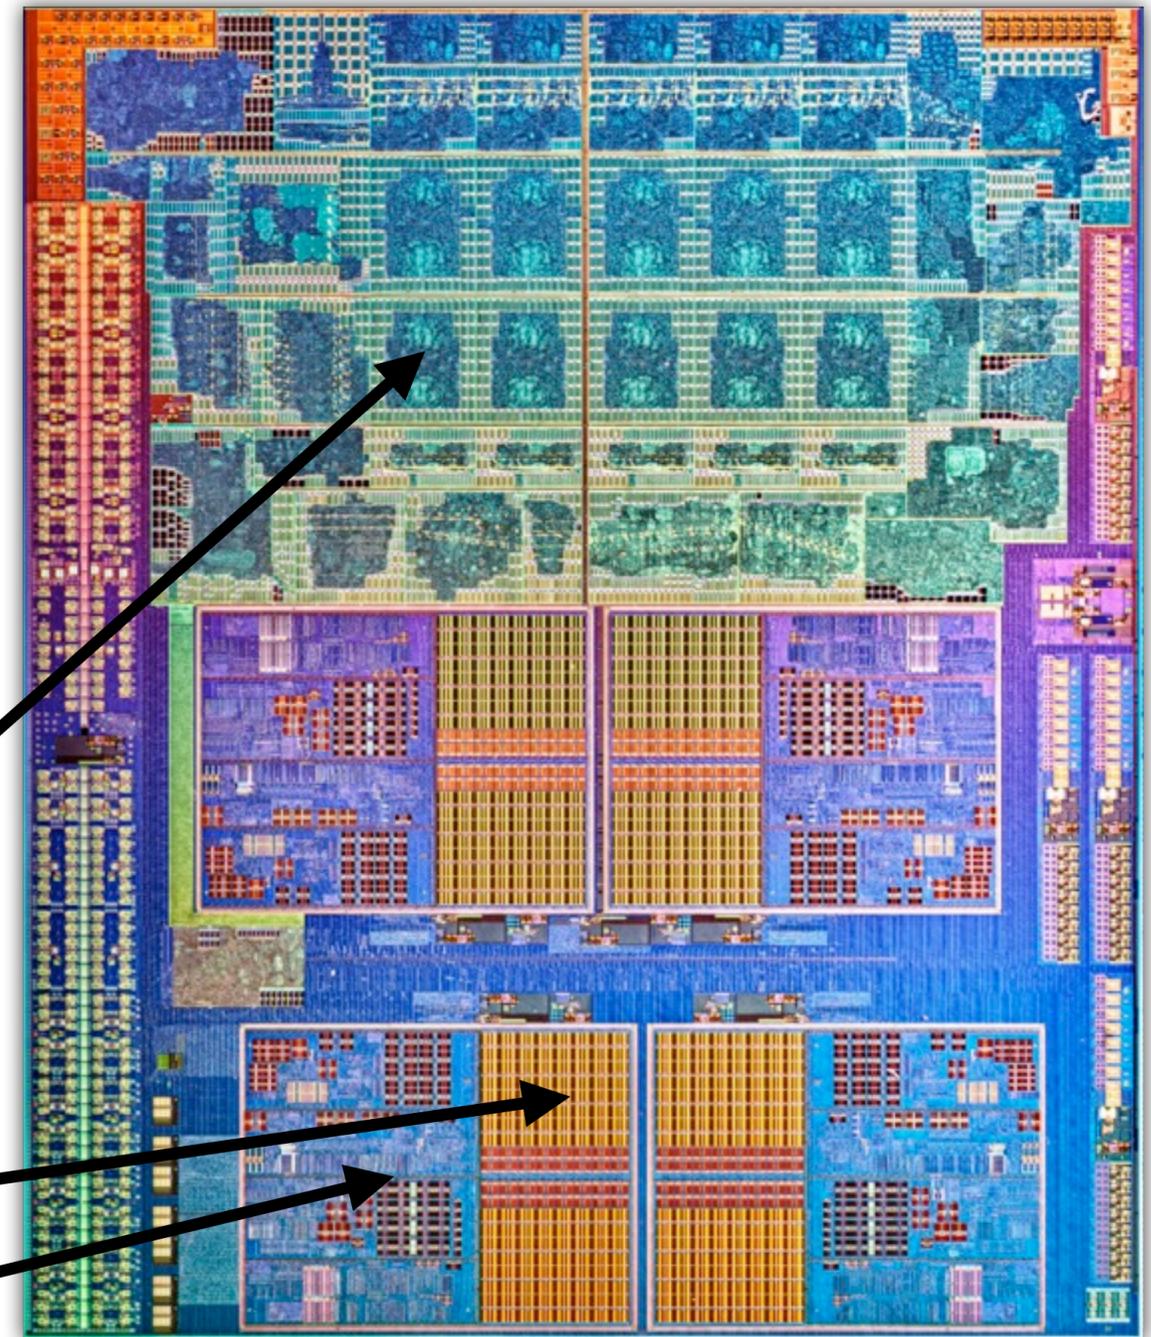
**Most efficient processor is a heterogeneous mixture of resources.  
("use the most efficient tool for the job")**

**\*\* You will likely make this observation during your projects**

# Example: AMD Fusion



- **“APU”**: accelerated processing unit
- **Integrate CPU cores and GPU-style cores on same chip**
- **Share memory system**
  - **Regions of physical memory reserved for graphics (not x86 coherent)**
  - **Rest of memory is x86 coherent**
  - **CPU and graphics memory spaces are not coherent, but at least there is no need to copy data in physical memory (or over PCIe bus) to communicate between CPU and graphics**



Graphics data-parallel (accelerator) core

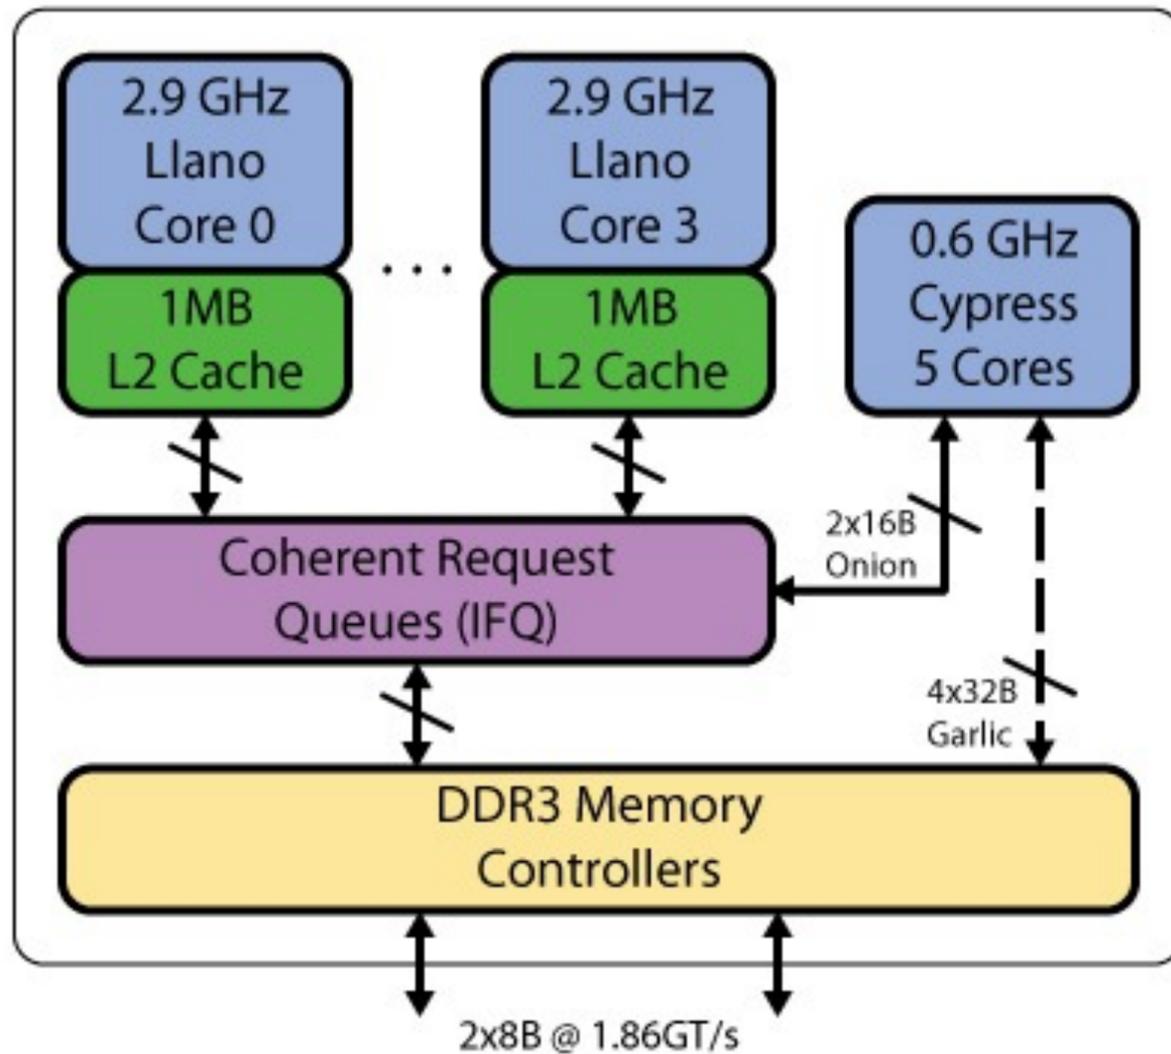
CPU L2

CPU core

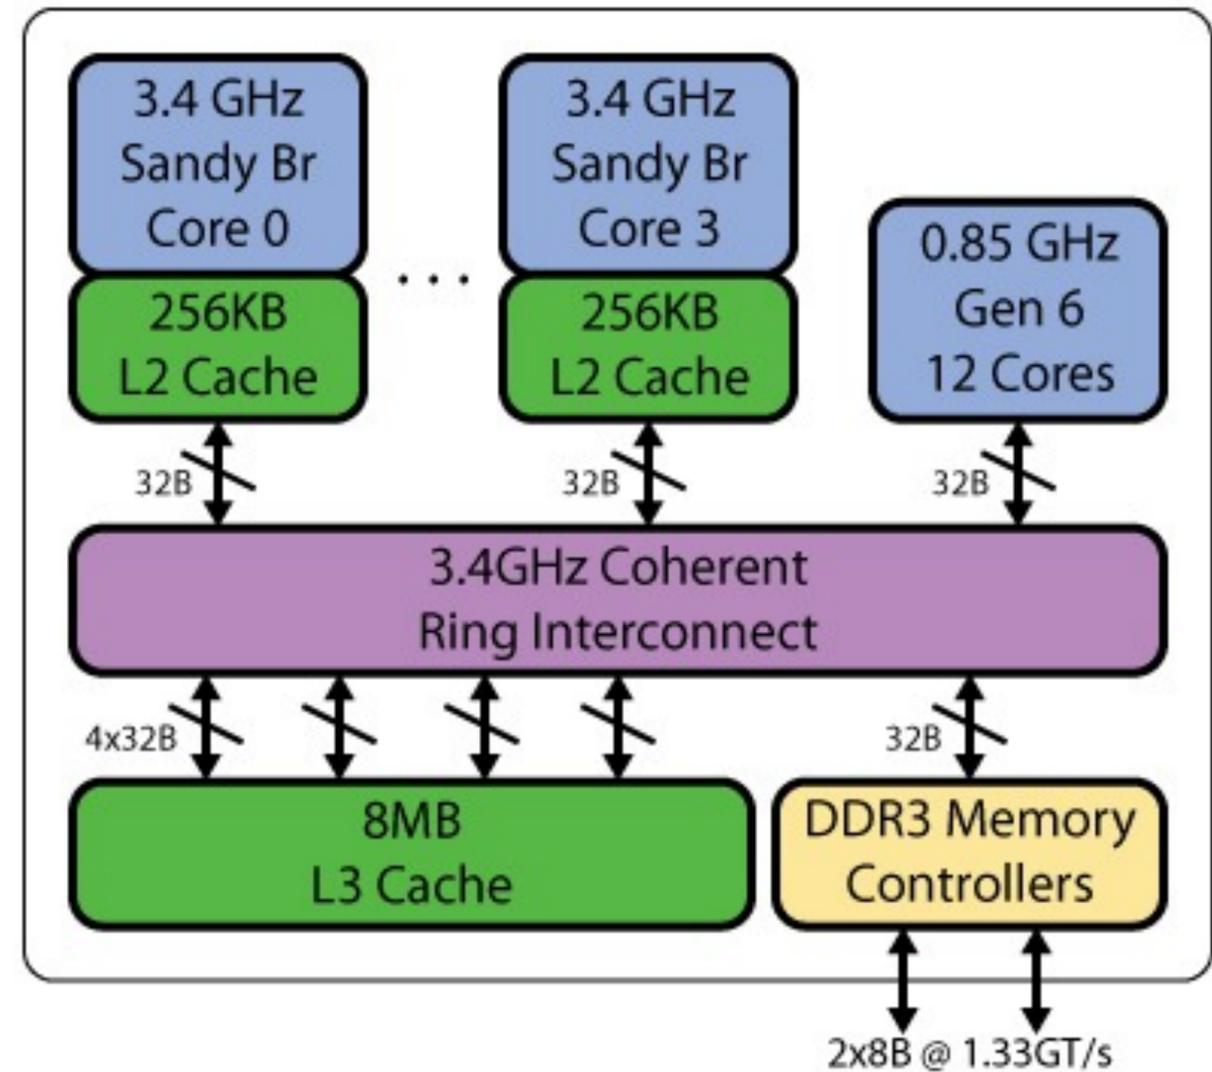
**AMD Llano**  
**(4 CPU cores + integrated GPU cores)**

# Multi-core CPU + integrated graphics

## AMD Llano



## Intel Sandy Bridge



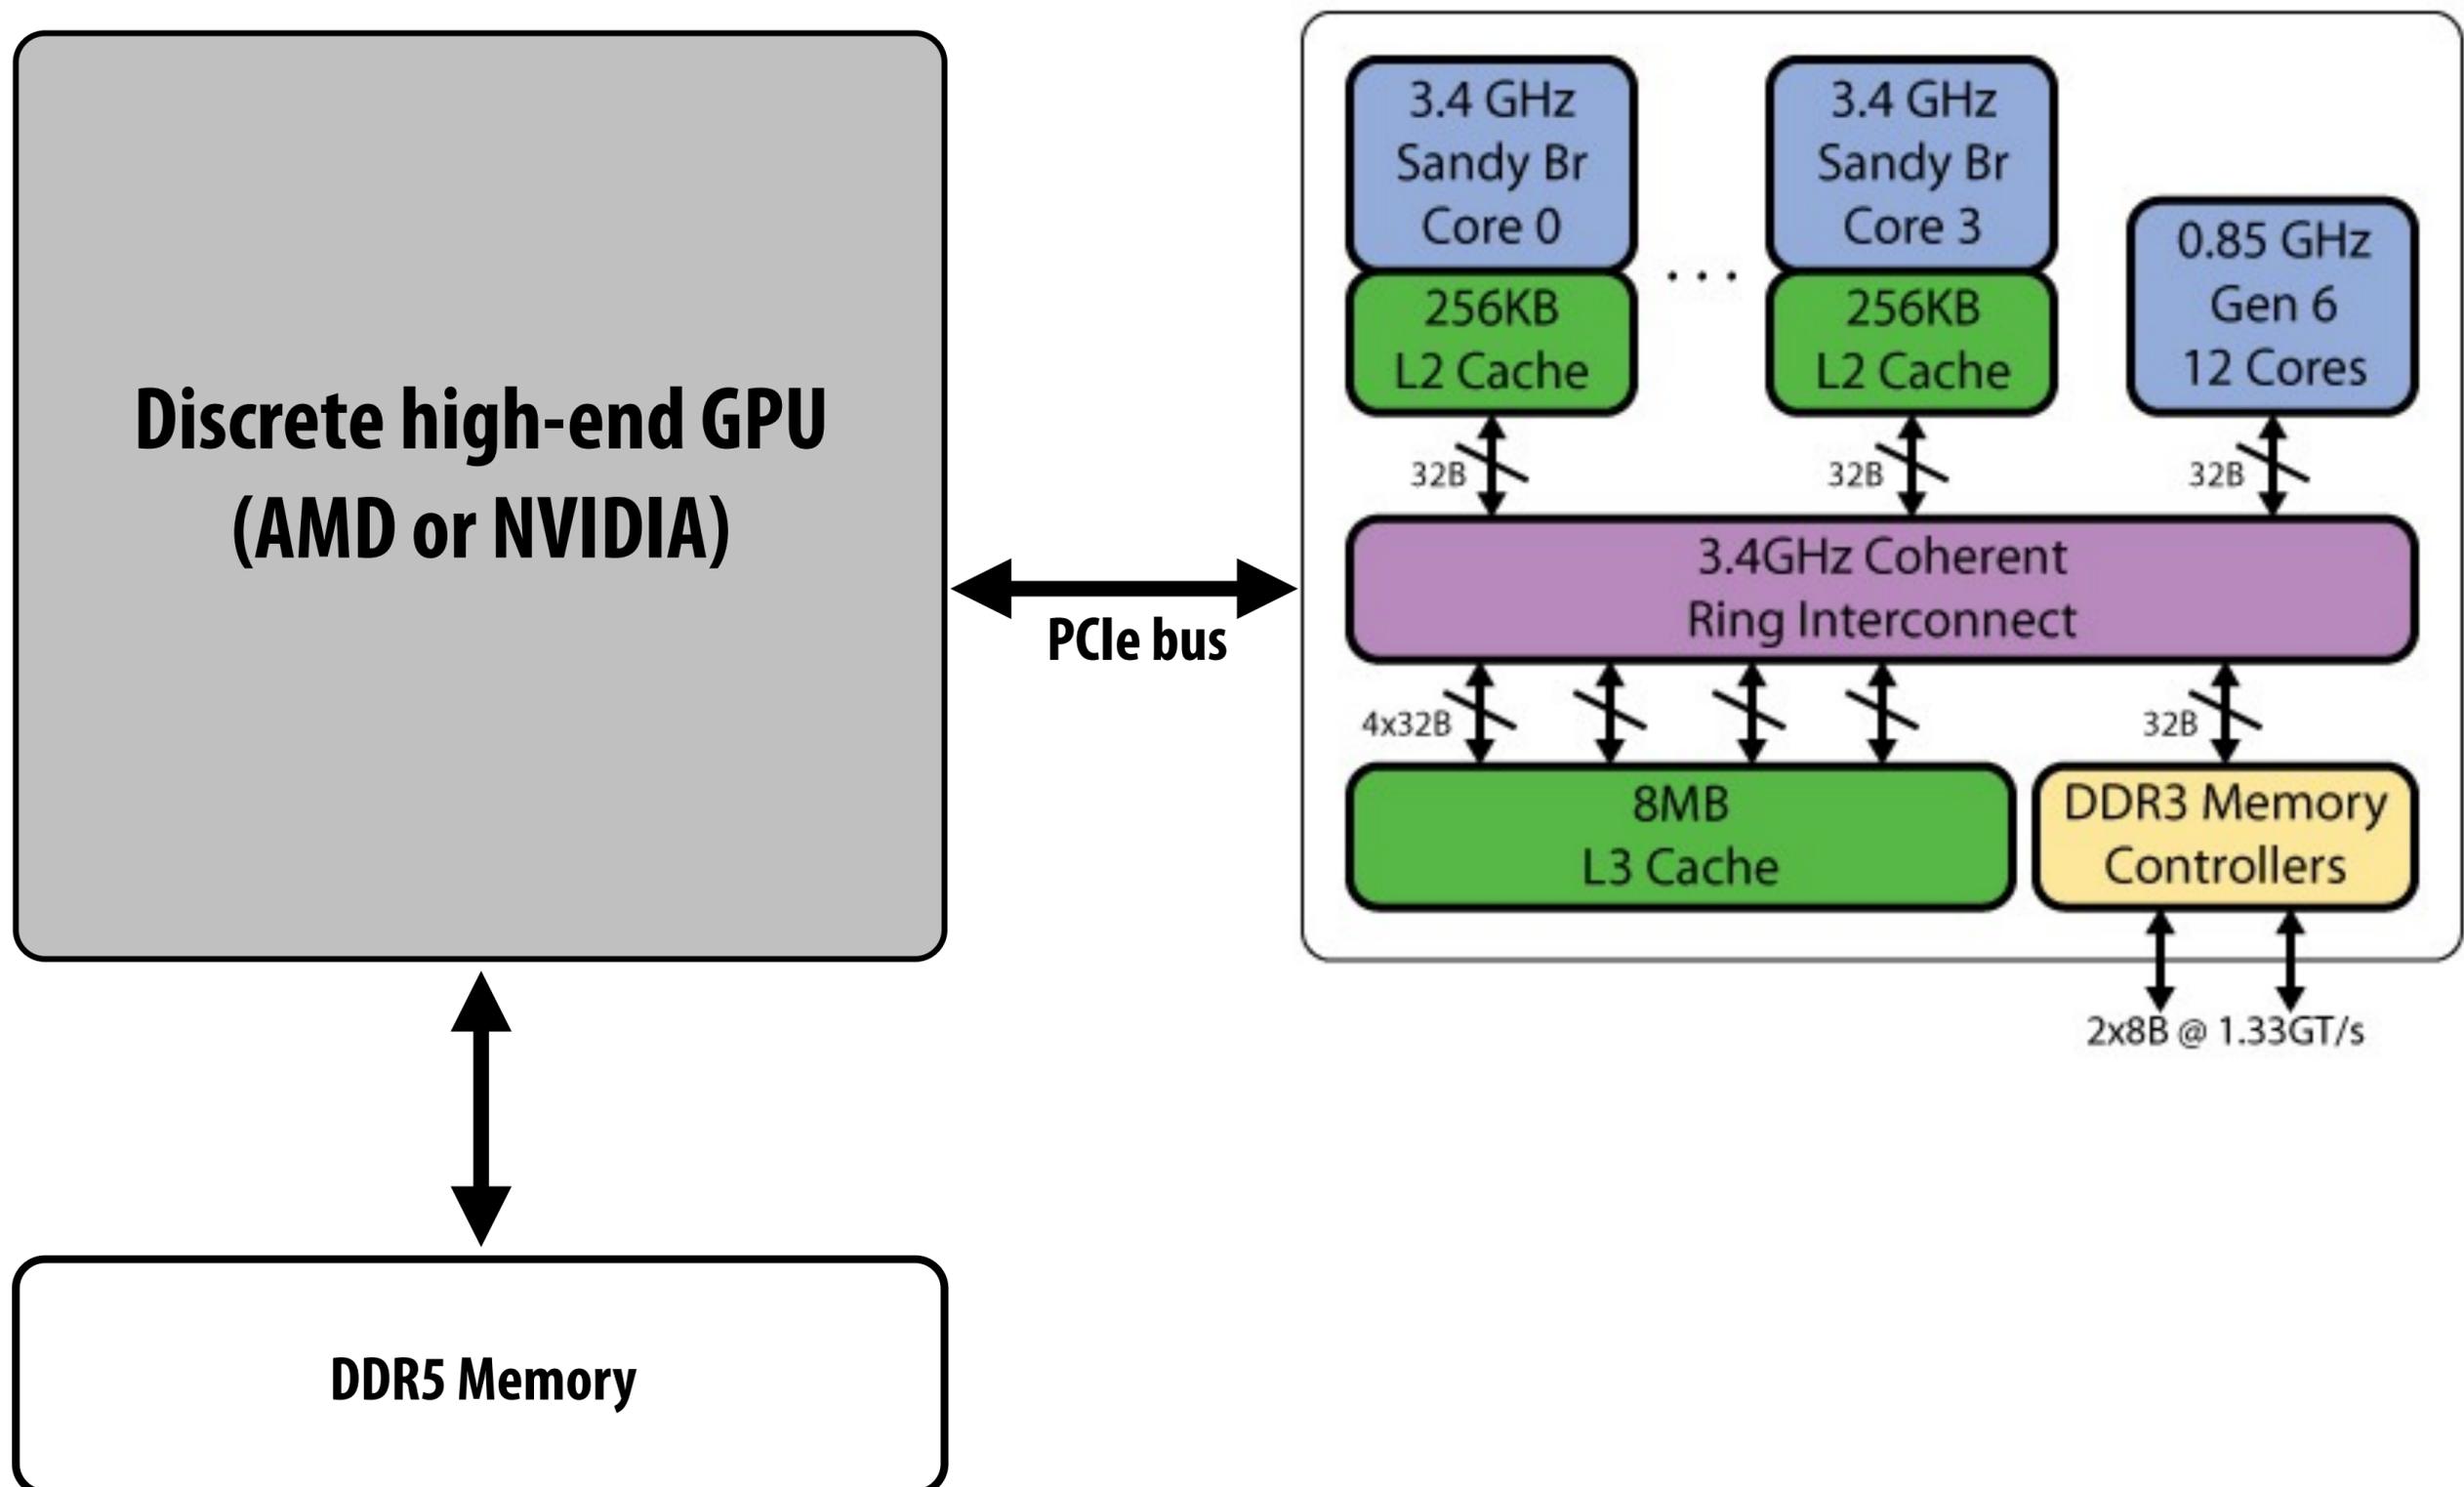
# More heterogeneity: add discrete GPU

Keep discrete (power hungry) GPU unless needed for graphics-intensive applications

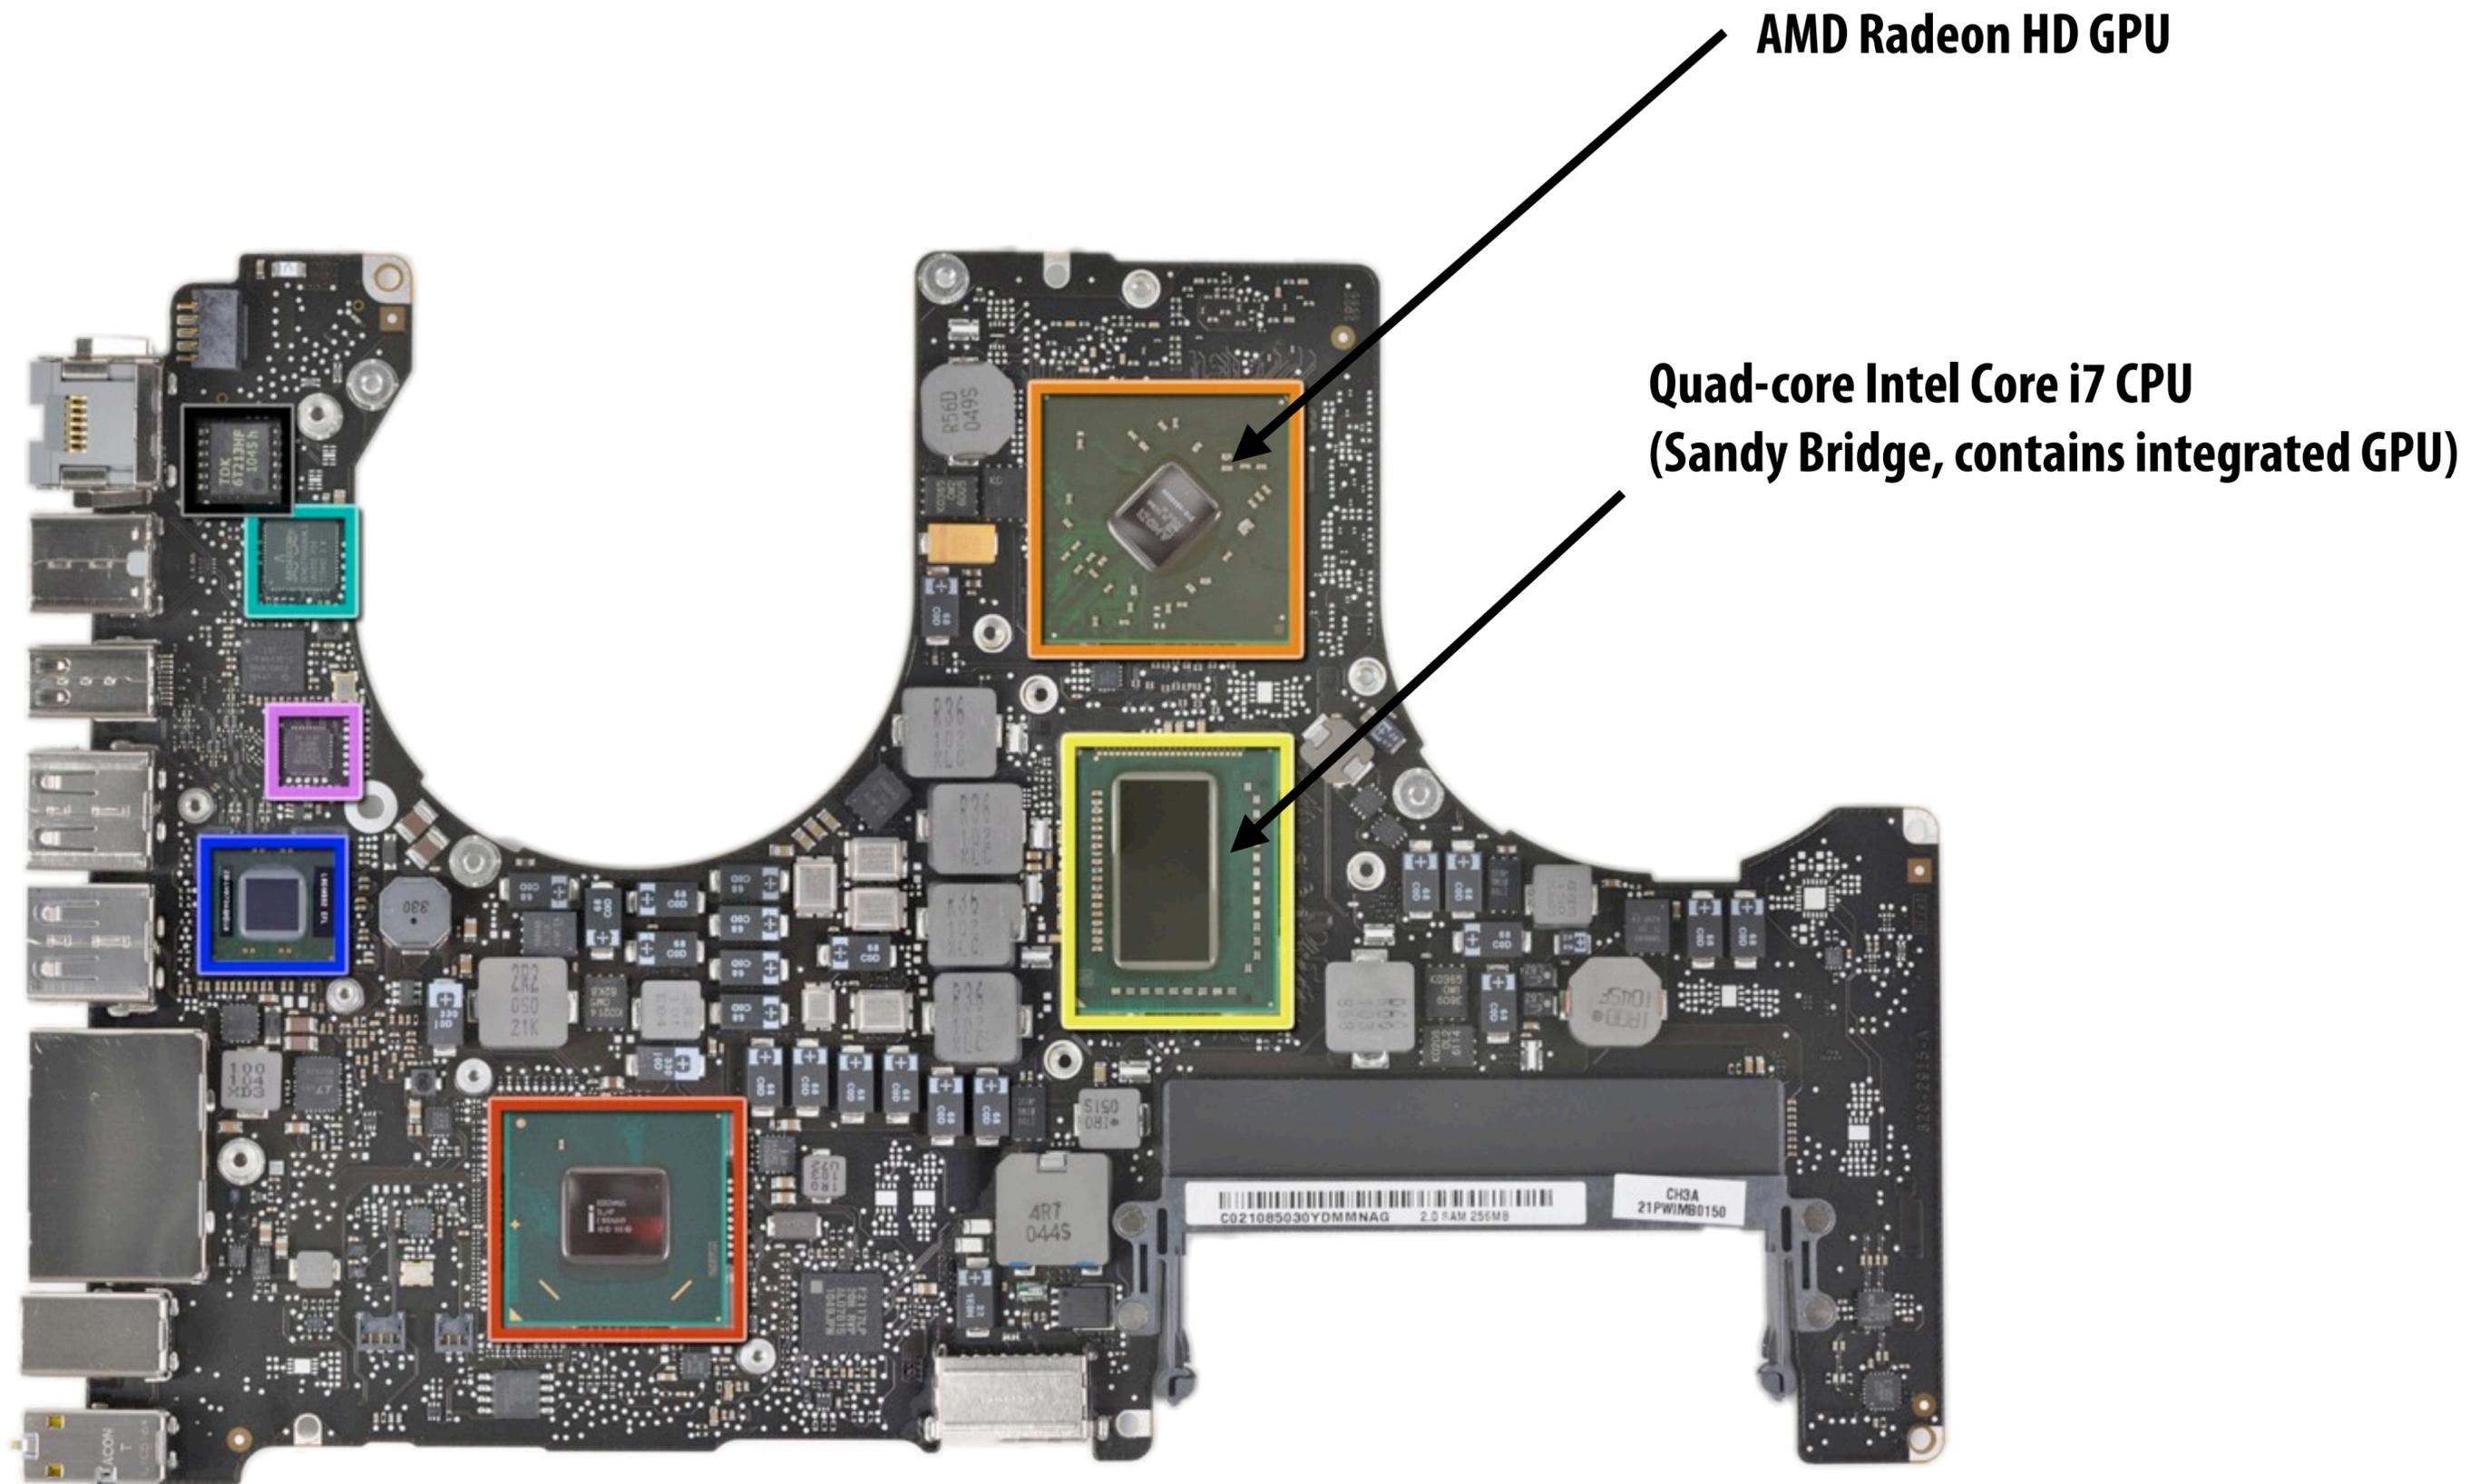
Use integrated, low power graphics for window manager/UI

(Neat: AMD Fusion can parallelize graphics across integrated and discrete GPU)

## Intel Sandy Bridge



# My Macbook Pro 2011 (two GPUs)



# Supercomputers use heterogeneous processing

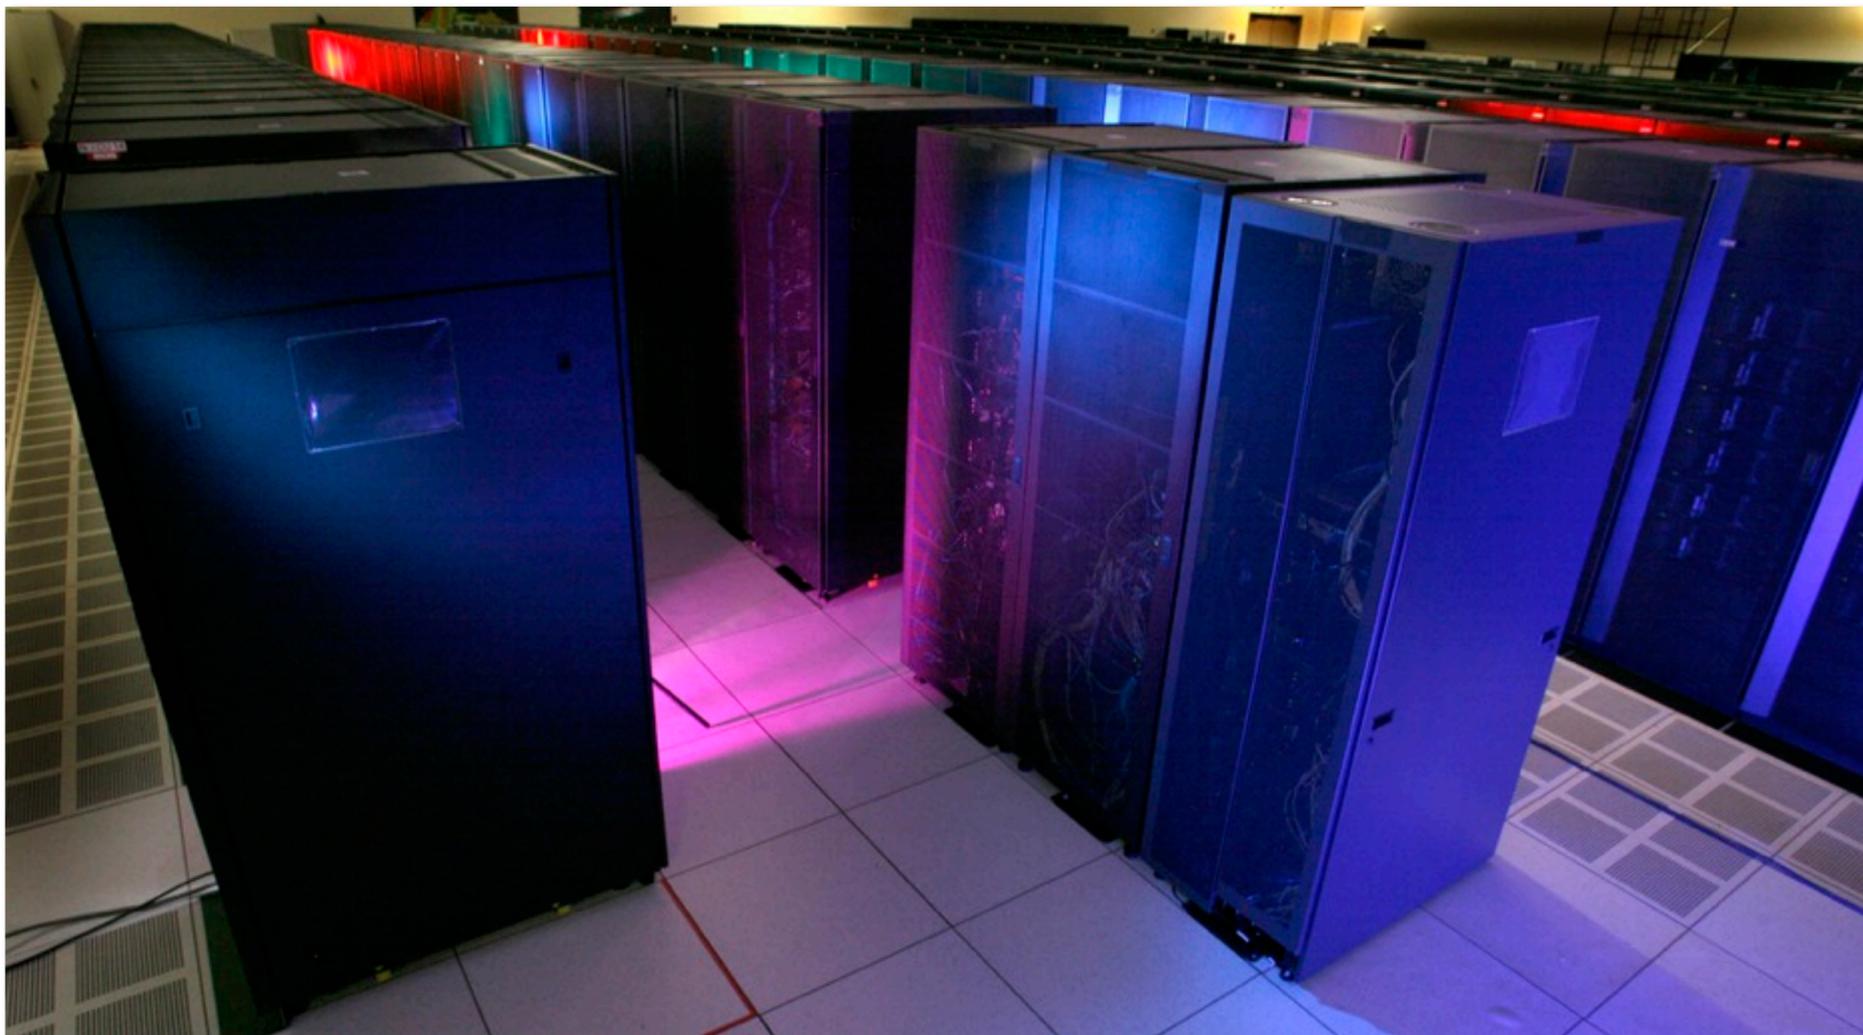
## ■ Los Alamos National Laboratory: Roadrunner

Fastest US supercomputer in 2008, first to break Petaflop barrier: 1.7 PFLOPS

Unique at the time due to use of two types of processing elements

(IBM's Cell processor served as "accelerator" to achieve desired compute density)

- 6,480 AMD Opteron dual-core CPUs (12,960 cores)
- 12,970 IBM Cell Processors (1 CPU + 8 accelerator cores per Cell = 116,640 cores)
- 2.4 MWatt (about 2,400 average US homes)



# Heterogeneous architecture for supercomputing

Source: Top500.org Fall 2012 rankings

Rank	Site	System	Cores	Rmax (TFlop/s)	Rpeak (TFlop/s)	Power (kW)
1	DOE/SC/Oak Ridge National Laboratory United States	<b>Titan</b> - Cray XK7 , Opteron 6274 16C 2.200GHz, Cray Gemini interconnect, NVIDIA K20x Cray Inc. <b>GPU</b>	560640	17590.0	27112.5	8209
2	DOE/NNSA/LLNL United States	<b>Sequoia</b> - BlueGene/Q, Power BQC 16C 1.60 GHz, Custom IBM	1572864	16324.8	20132.7	7890
3	RIKEN Advanced Institute for Computational Science (AICS) Japan	<b>K computer</b> , SPARC64 VIIIfx 2.0GHz, Tofu interconnect Fujitsu	705024	10510.0	11280.4	12660
4	DOE/SC/Argonne National Laboratory United States	<b>Mira</b> - BlueGene/Q, Power BQC 16C 1.60GHz, Custom IBM	786432	8162.4	10066.3	3945
5	Forschungszentrum Juelich (FZJ) Germany	<b>JUQUEEN</b> - BlueGene/Q, Power BQC 16C 1.600GHz, Custom Interconnect IBM	393216	4141.2	5033.2	1970
6	Leibniz Rechenzentrum Germany	<b>SuperMUC</b> - iDataPlex DX360M4, Xeon E5-2680 8C 2.70GHz, Infiniband FDR IBM	147456	2897.0	3185.1	3423
7	Texas Advanced Computing Center/Univ. of Texas United States	<b>Stampede</b> - PowerEdge C8220, Xeon E5-2680 8C 2.700GHz, Infiniband FDR, Intel Xeon Phi Dell <b>two types of CPU</b>	204900	2660.3	3959.0	
8	National Supercomputing Center in Tianjin China	<b>Tianhe-1A</b> - NUDT YH MPP, Xeon X5670 6C 2.93 GHz, NVIDIA 2050 NUDT <b>GPU</b>	186368	2566.0	4701.0	4040

**27 PFLOPS,  
8.2 MWatt**

# GPU-accelerated supercomputing

- **Oak Ridge Titan (world's #1)**
- **18,688 AMD Opteron 16-core CPUs**
- **18,688 NVIDIA Tesla K20X GPUs**
- **710 TB RAM**



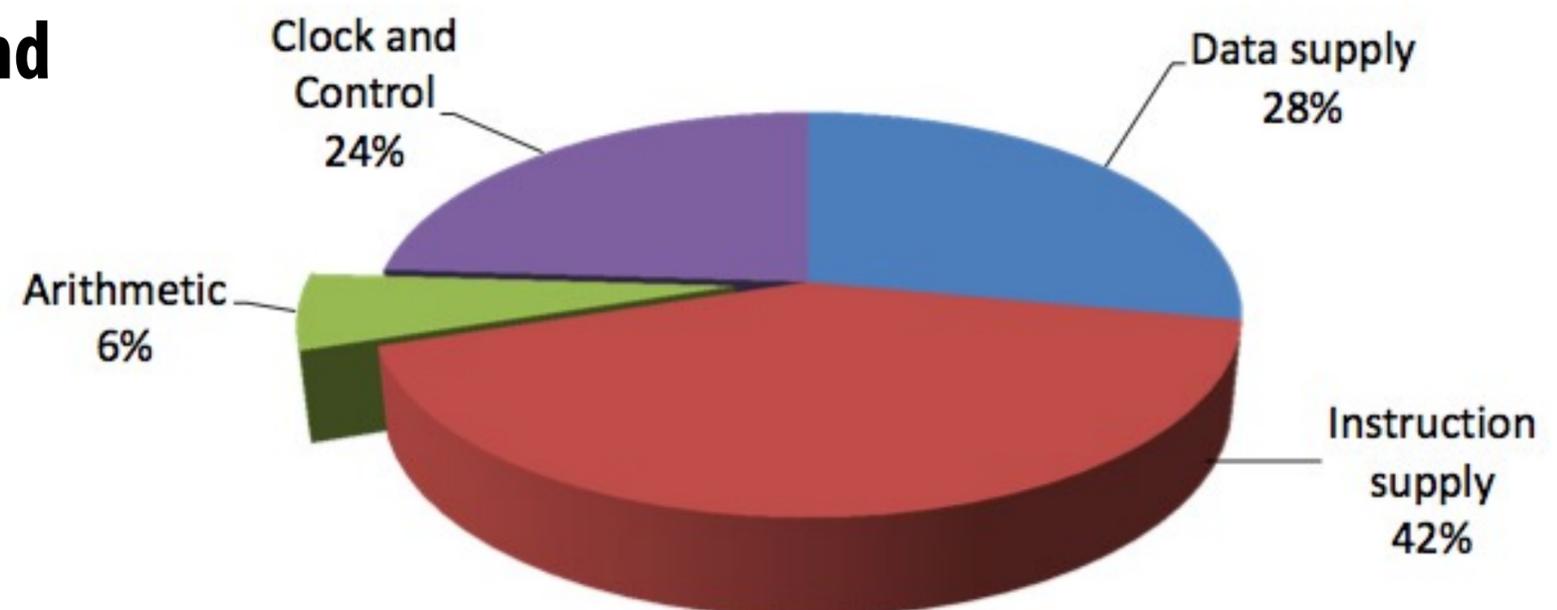
- **Estimated machine cost \$97M**
- **Estimated annual power/operating cost: ~ \$9M \***

# Energy-constrained computing

- **Supercomputers are energy constrained**
  - **Due to sheer scale**
  - **Overall cost to operate (power for machine and for cooling)**
  
- **Mobile devices are energy constrained**
  - **Limited battery life**
  - **Heat dissipation**

# Efficiency benefits of compute specialization

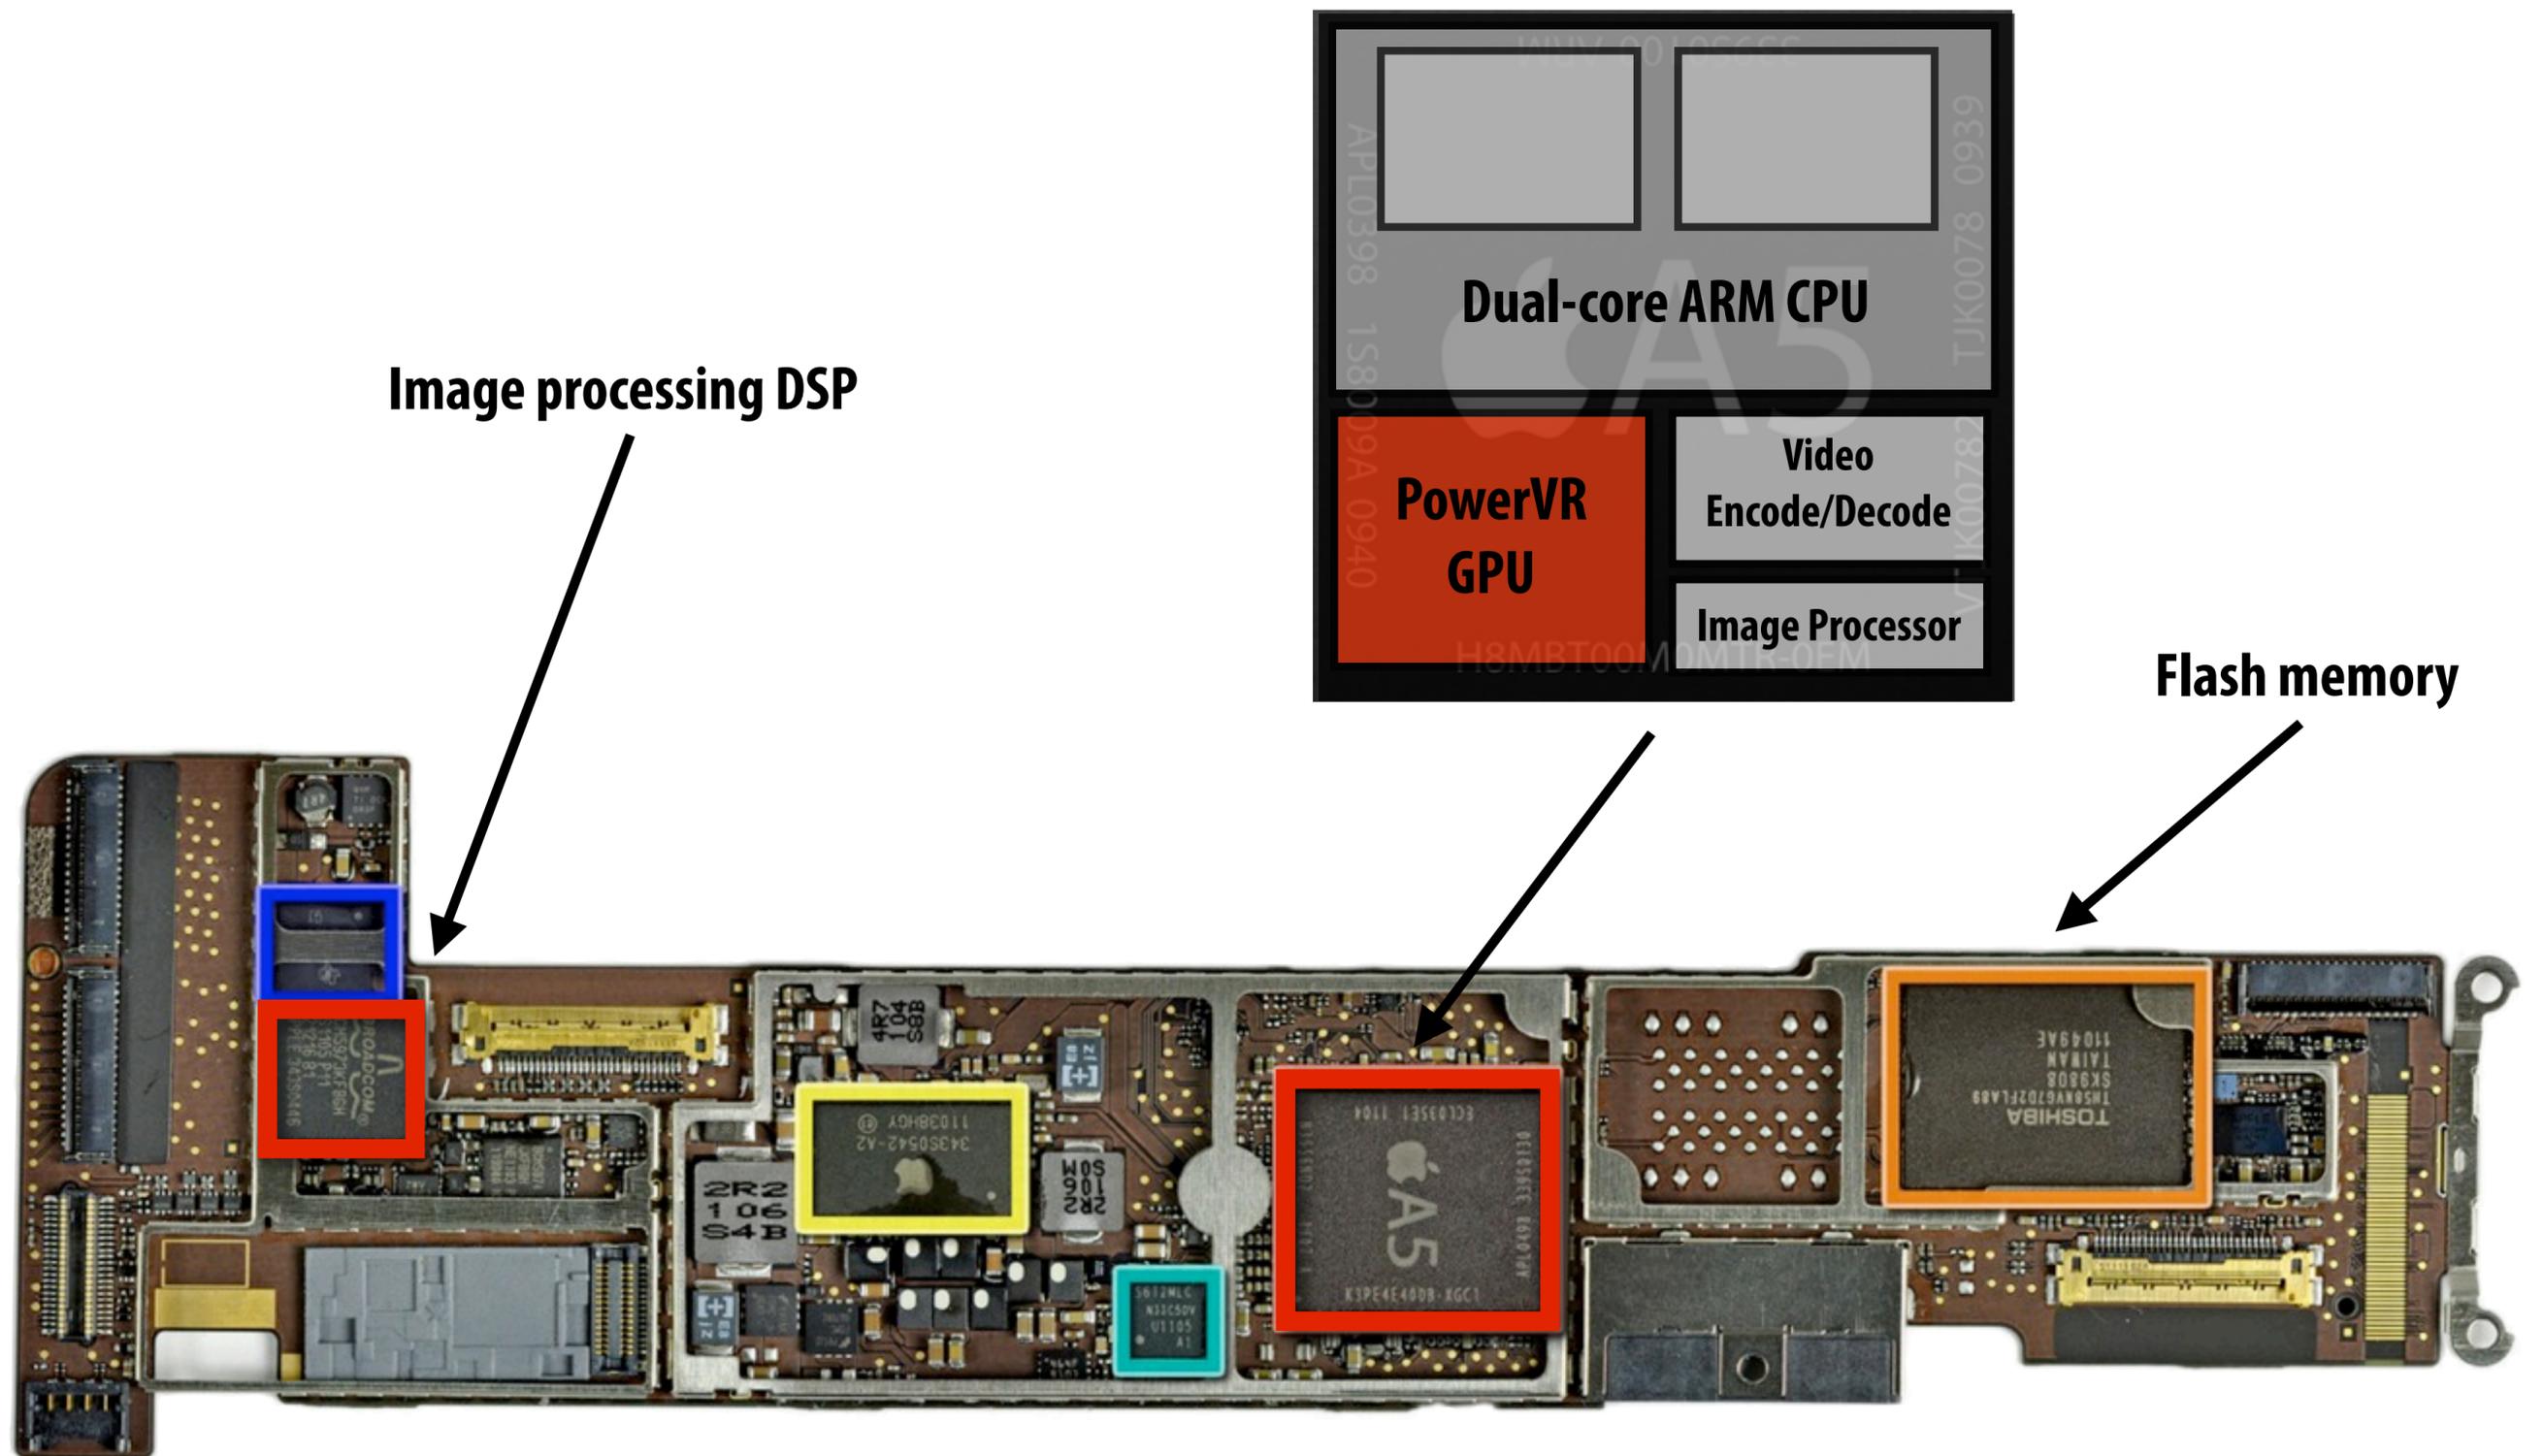
- **Rules of thumb: compared to average-quality C code on CPU...**
- **Throughput-maximized processor architectures: e.g., GPU cores**
  - ~ 10x improvement in perf / watt
  - Assuming code maps well to wide data-parallel execution and is compute bound
- **Fixed-function ASIC (“application-specific integrated circuit”)**
  - ~ 100x or greater improvement in perf/watt
  - Assuming code is compute bound



*Efficient Embedded Computing [Dally et al. 08]*

[Figure credit Eric Chung]

# Example: iPad 2 (2011)



# Original iPhone touchscreen controller

Separate digital signal processor to interpret raw signal from capacitive touch sensor (do not burden main CPU)

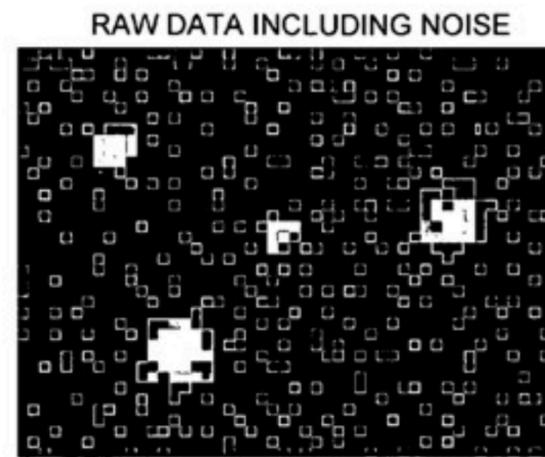
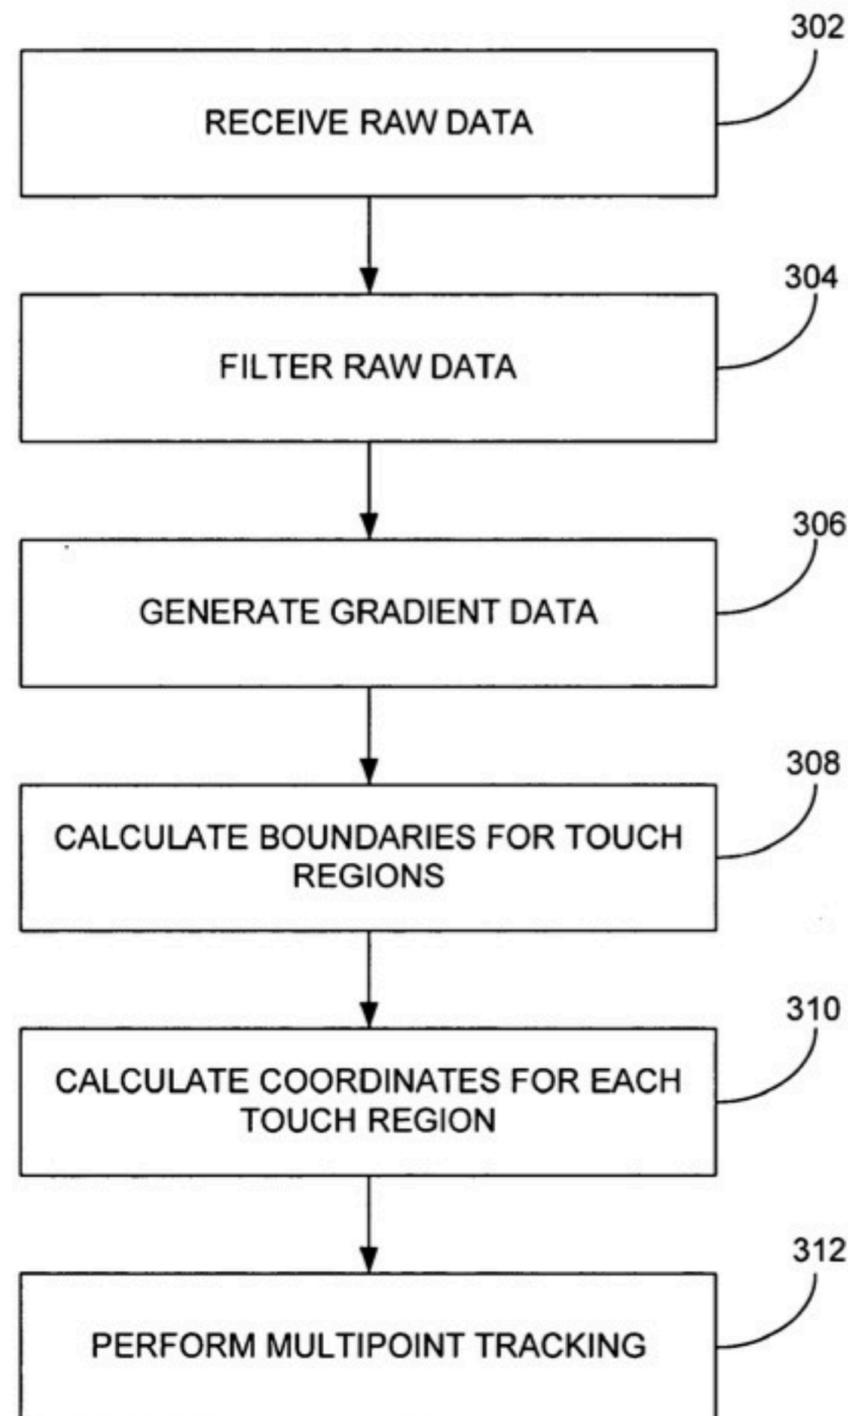


FIG. 17A

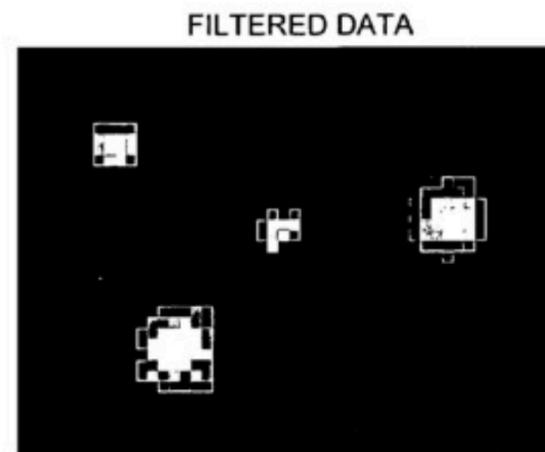


FIG. 17B

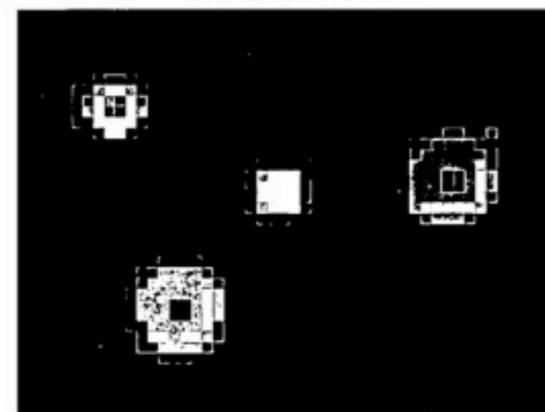


FIG. 17C

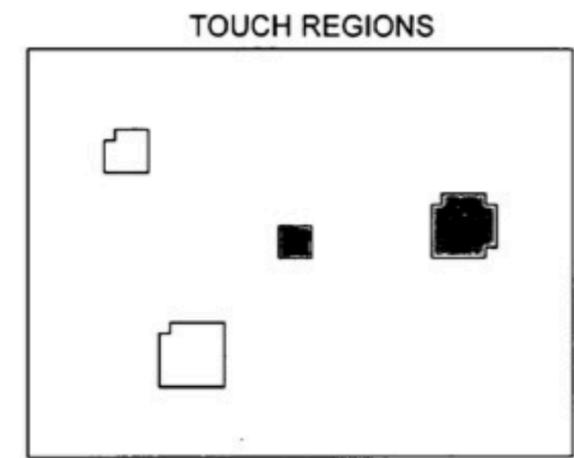


FIG. 17D

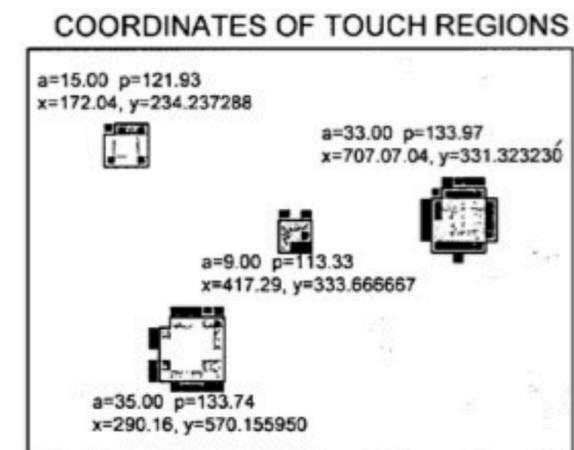
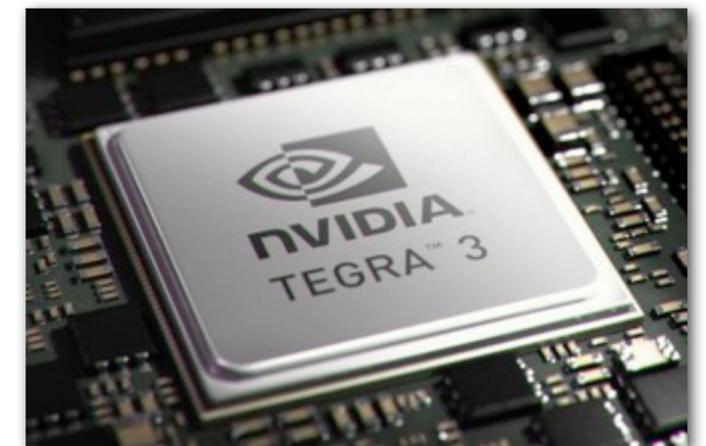
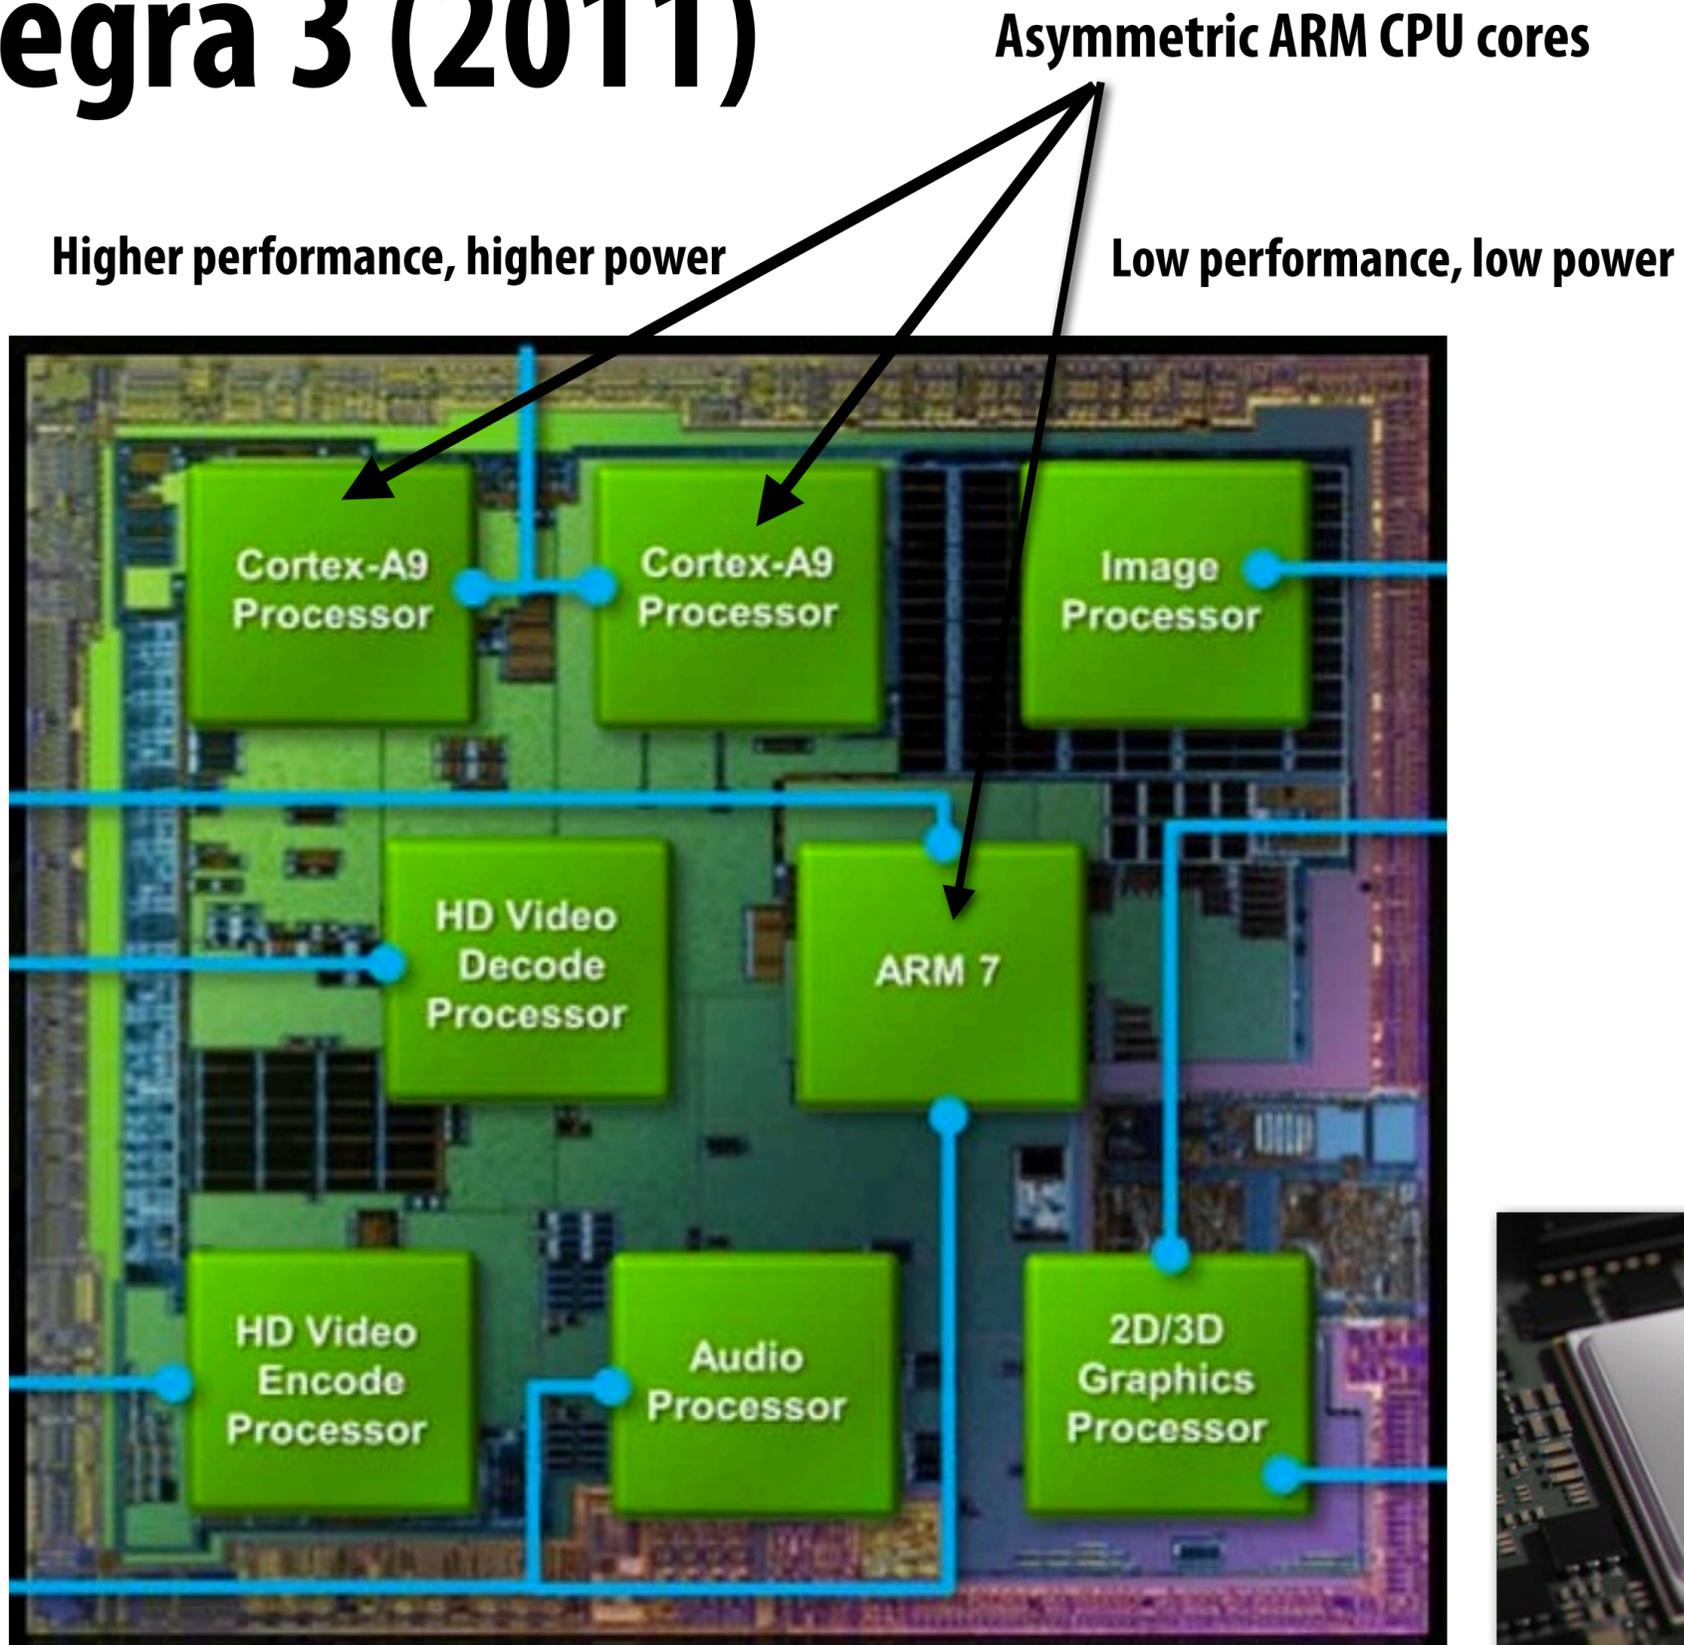
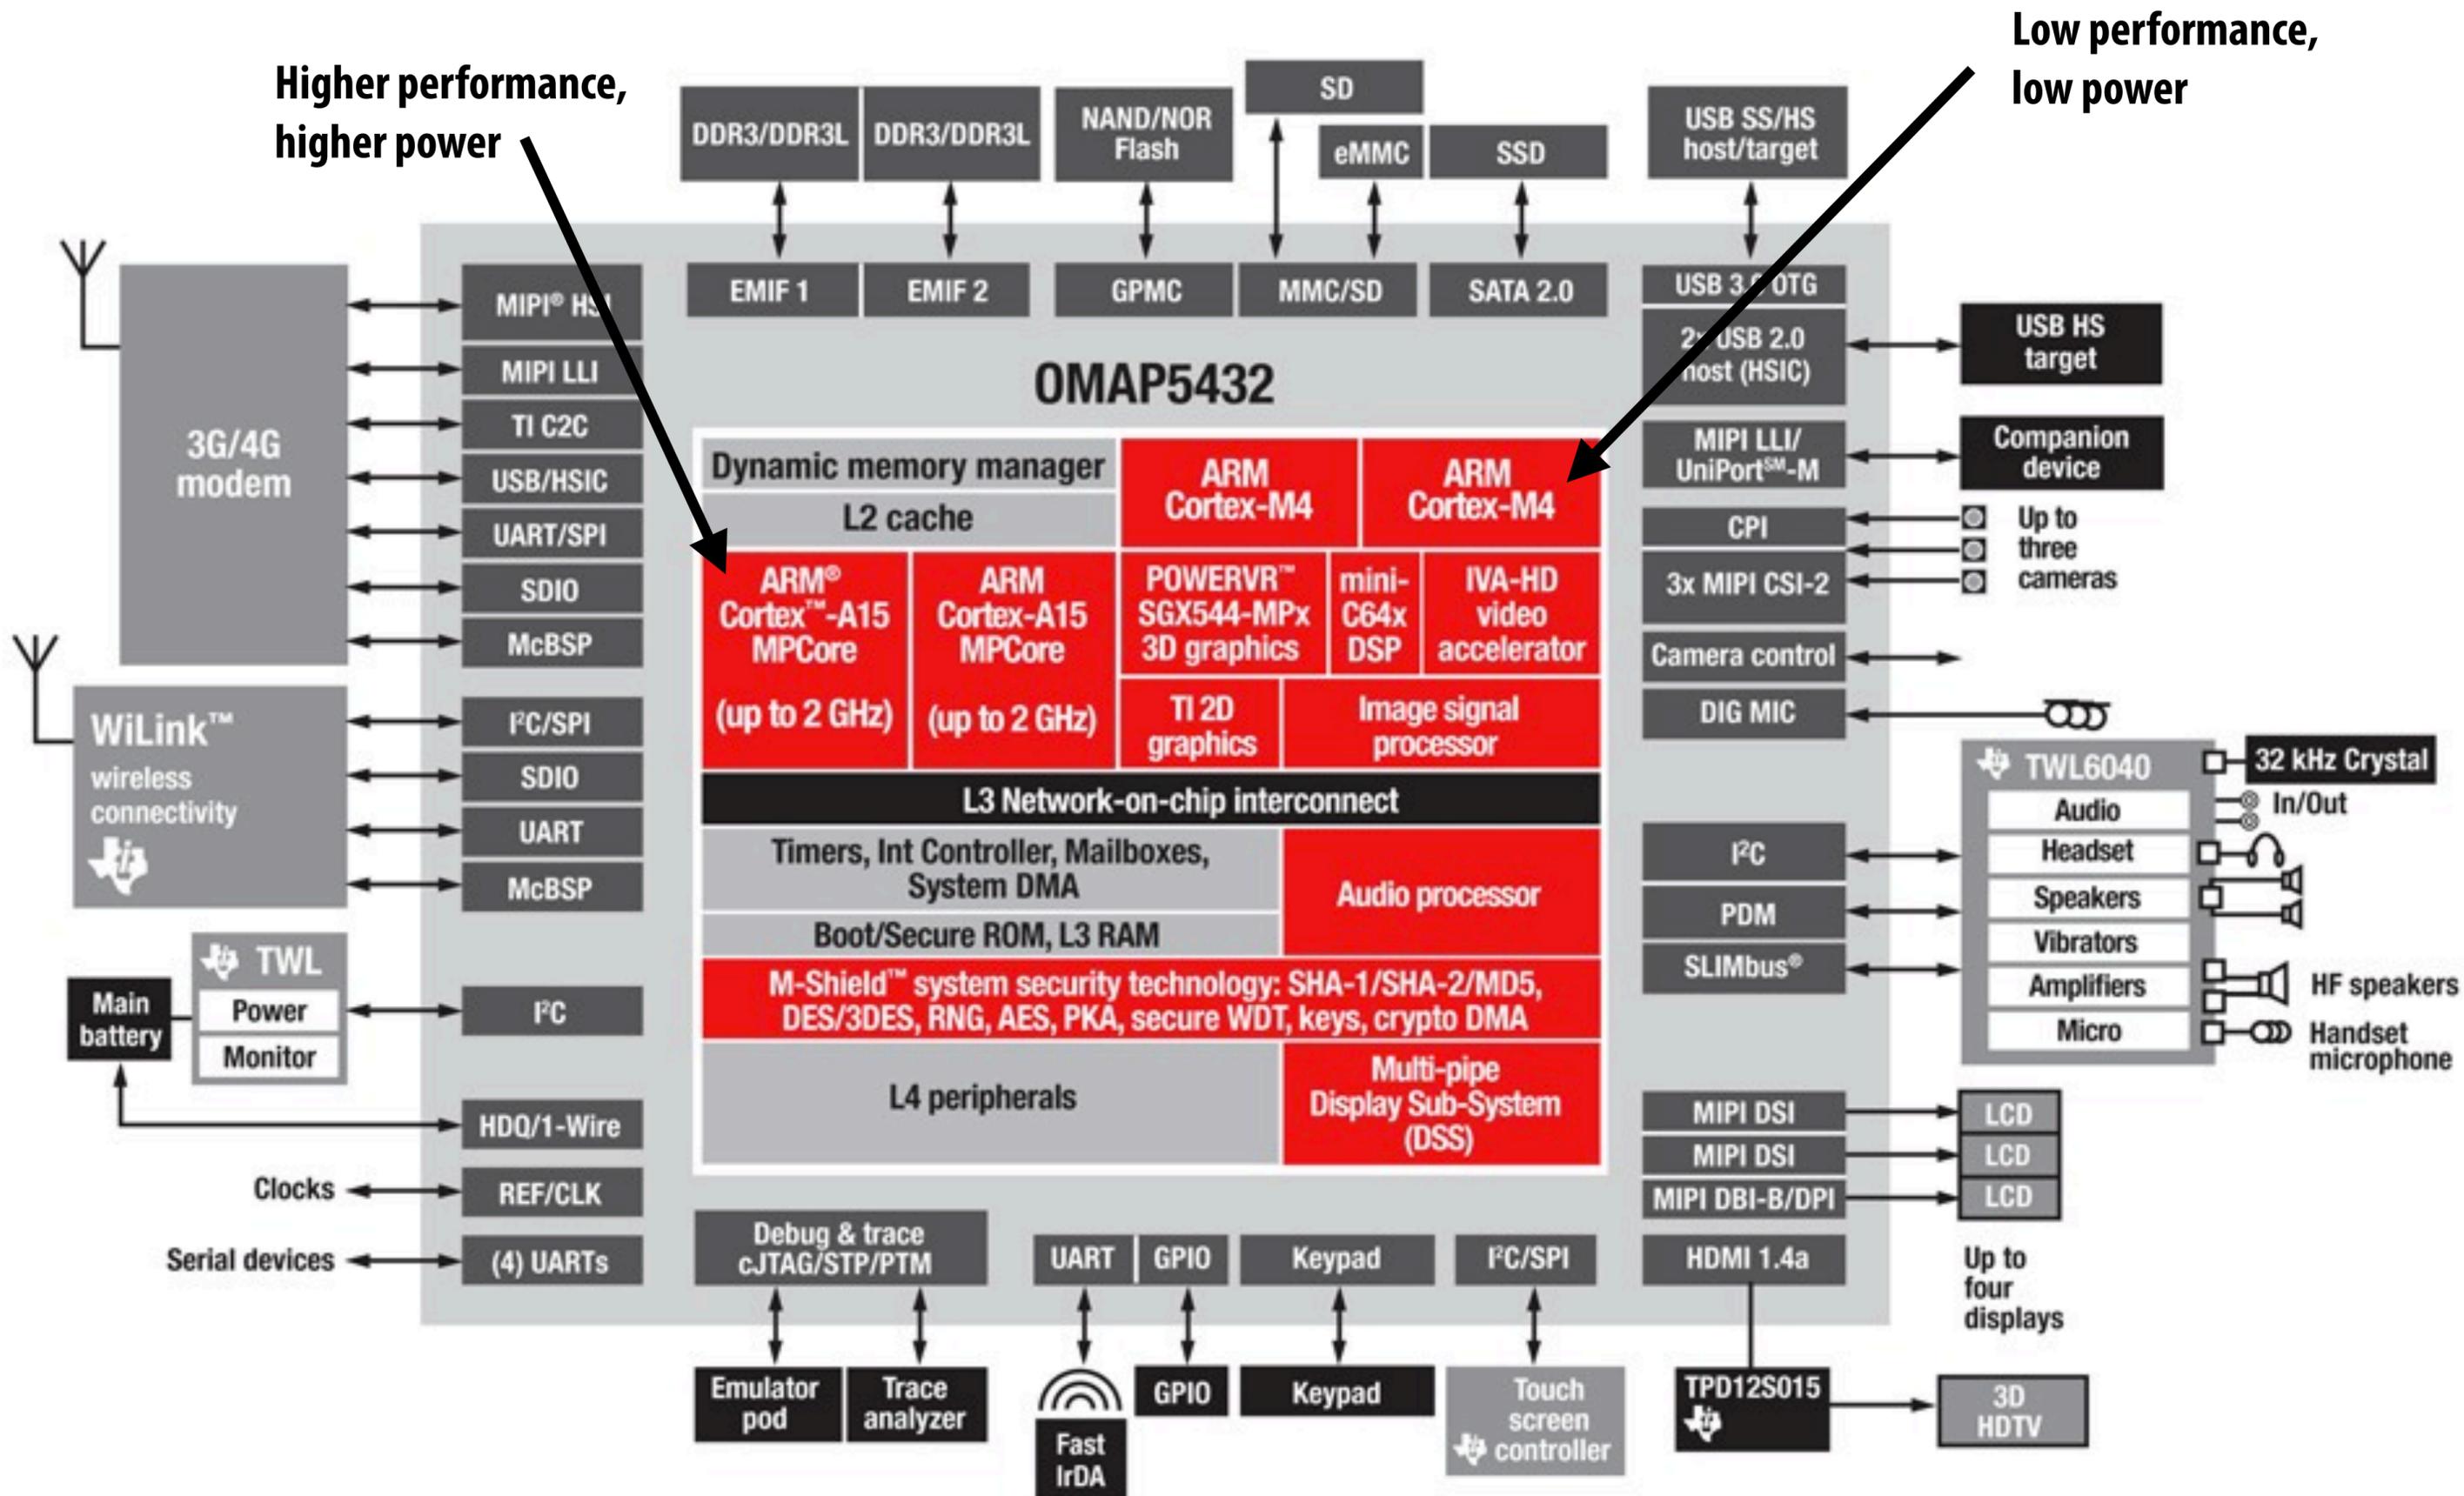


FIG. 17E

# NVIDIA Tegra 3 (2011)



# Texas Instruments OMAP 5 (2012)



# Today: performance often matters more, not less

Fourth, there's battery life.

To achieve long battery life when playing video, mobile devices must decode the video in hardware; decoding it in software uses too much power. Many of the chips used in modern mobile devices contain a decoder called H.264 – an industry standard that is used in every Blu-ray DVD player and has been adopted by Apple, Google (YouTube), Vimeo, Netflix and many other companies.

Although Flash has recently added support for H.264, the video on almost all Flash websites currently requires an older generation decoder that is not implemented in mobile chips and must be run in software. The difference is striking: on an iPhone, for example, H.264 videos play for up to 10 hours, while videos decoded in software play for less than 5 hours before the battery is fully drained.

When websites re-encode their videos using H.264, they can offer them without using Flash at all. They play perfectly in browsers like Apple's Safari and Google's Chrome without any plugins whatsoever, and look great on iPhones, iPods and iPads.

**Steve Jobs' "Thoughts on Flash", 2010**

**<http://www.apple.com/hotnews/thoughts-on-flash/>**

# Demo: image processing on Nikon D7000



**16 MPixel RAW image to JPG image conversion:**

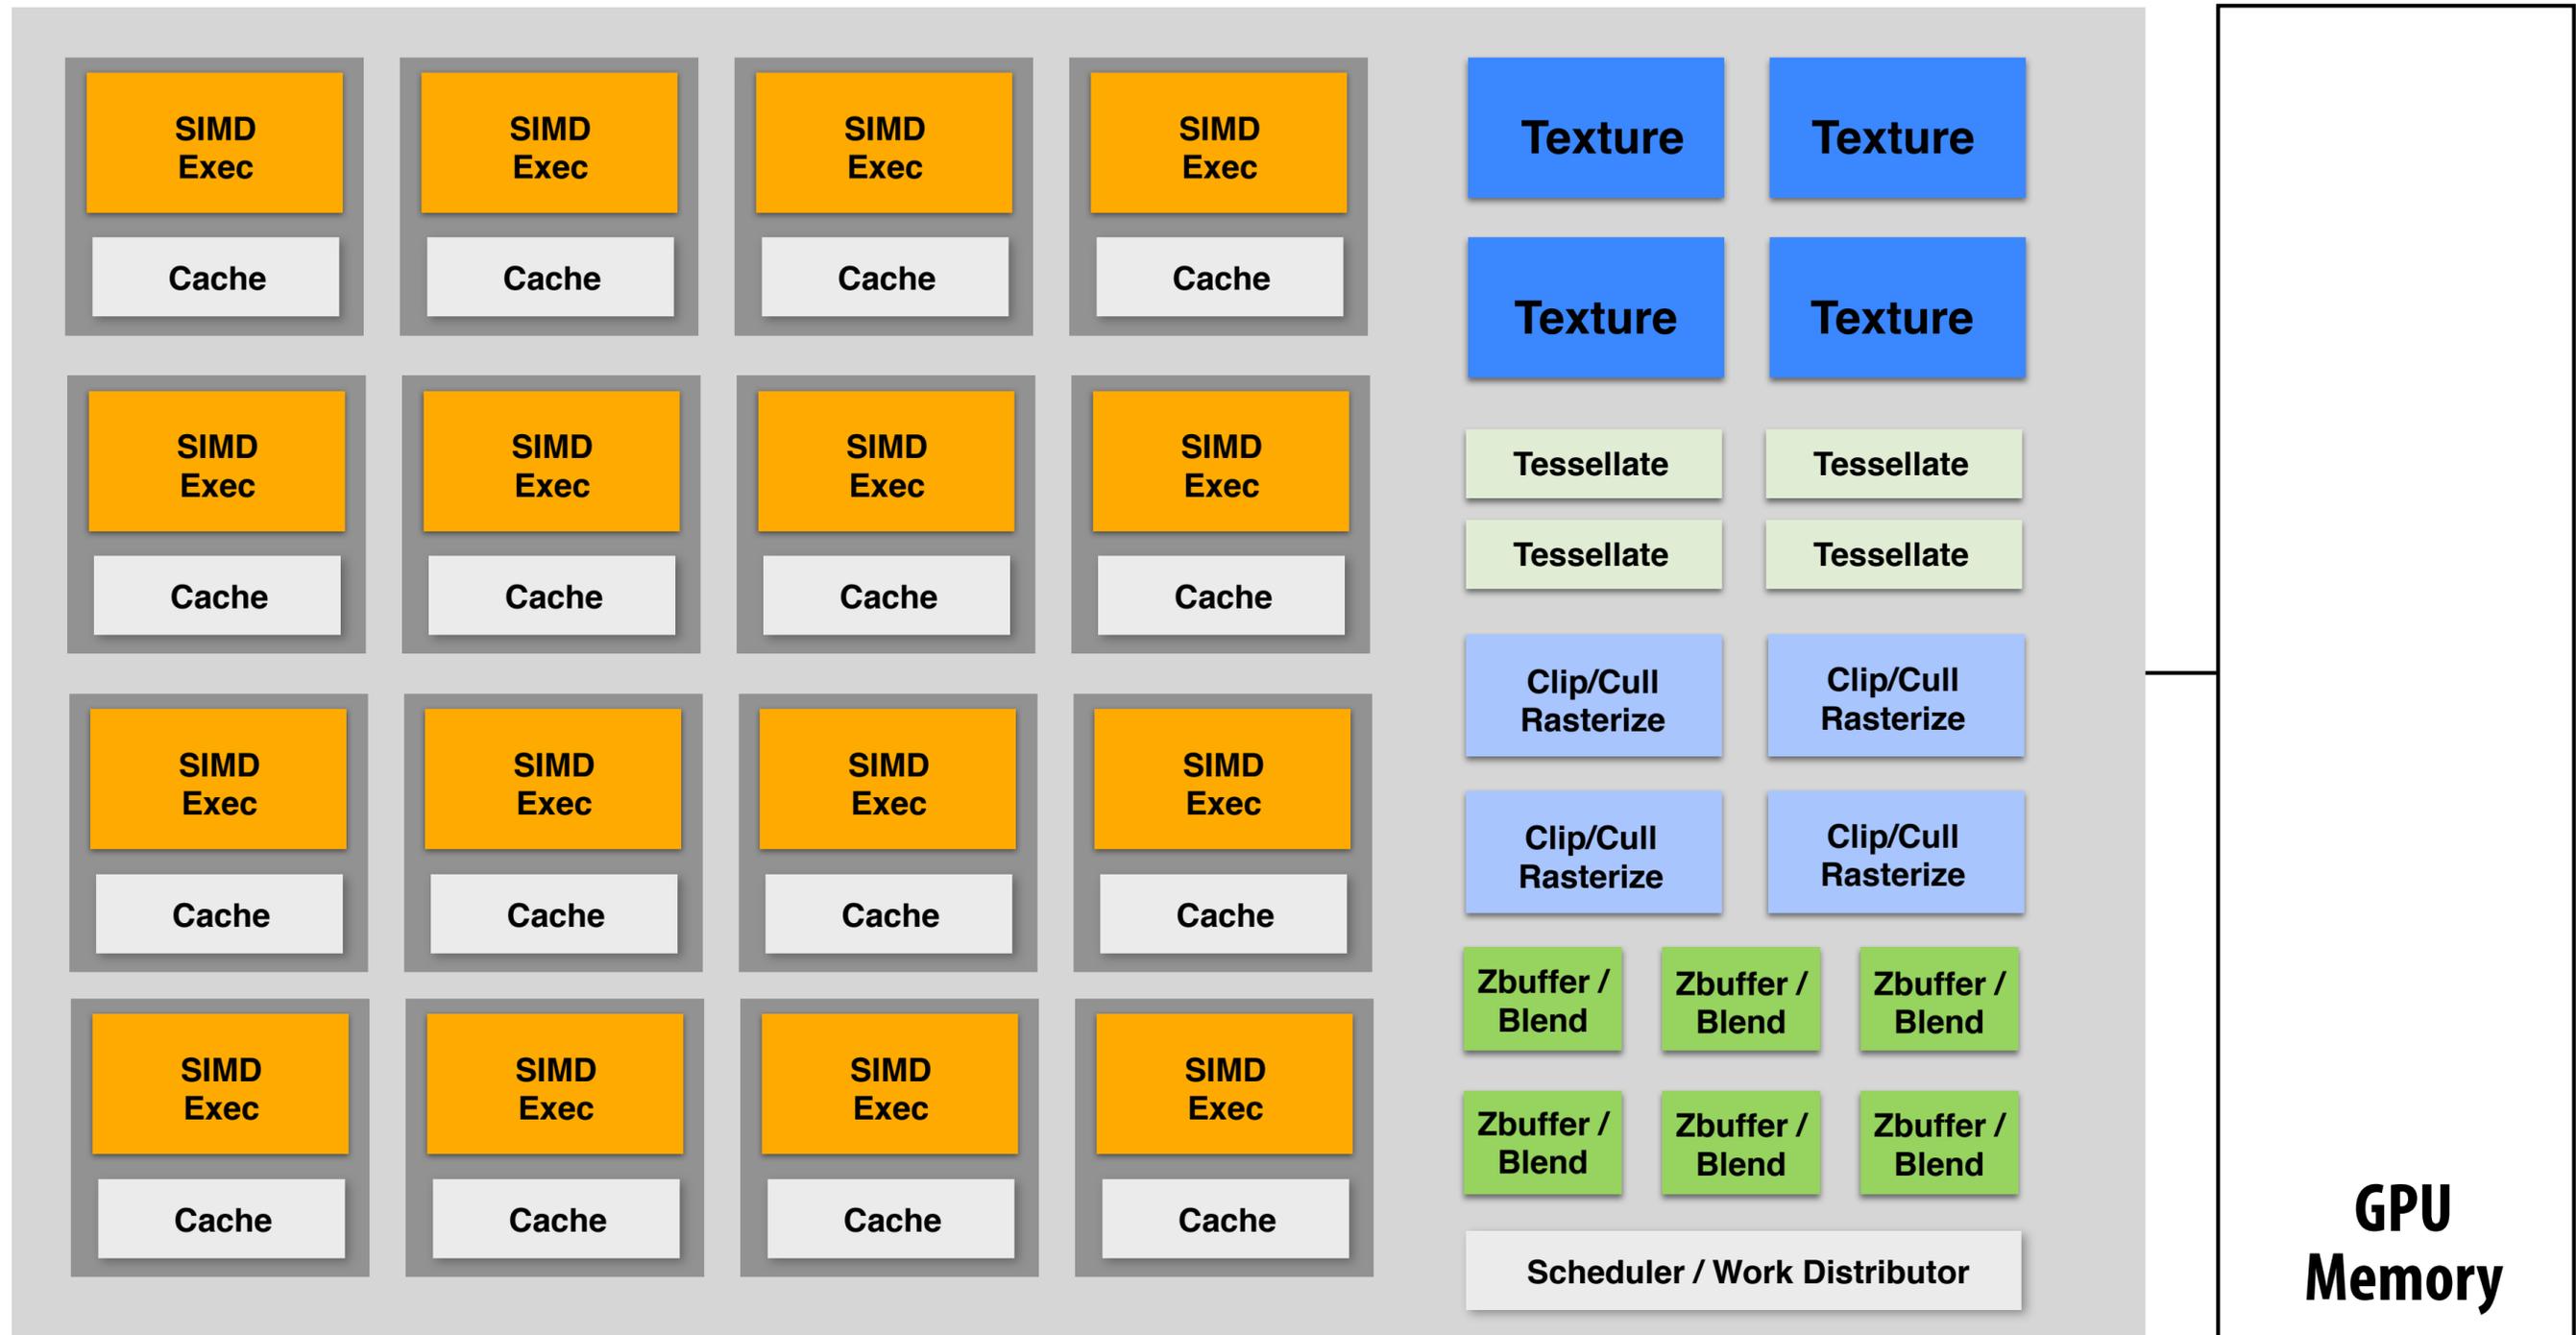
**Adobe Lightroom on my quad-core Macbook Pro laptop: 1-2 sec**

**Camera: ~ 1/6 sec**

# GPU is itself a heterogeneous multi-core processor

Compute resources your CUDA programs used in assignment 2

Graphics-specific, fixed-function compute resources

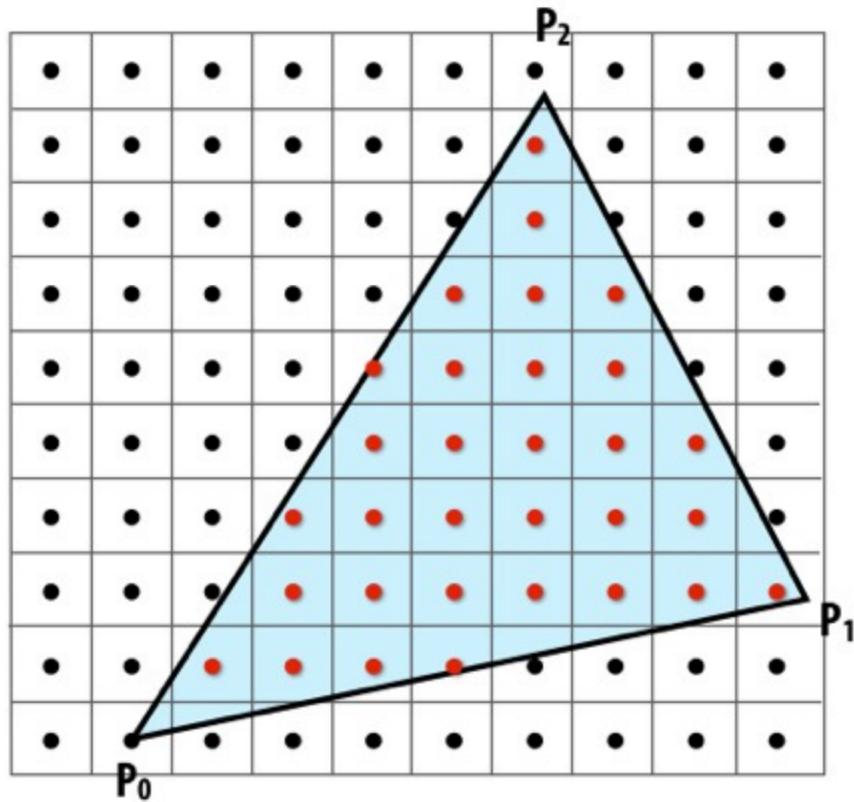


GPU

# Example graphics tasks performed in fixed-function HW

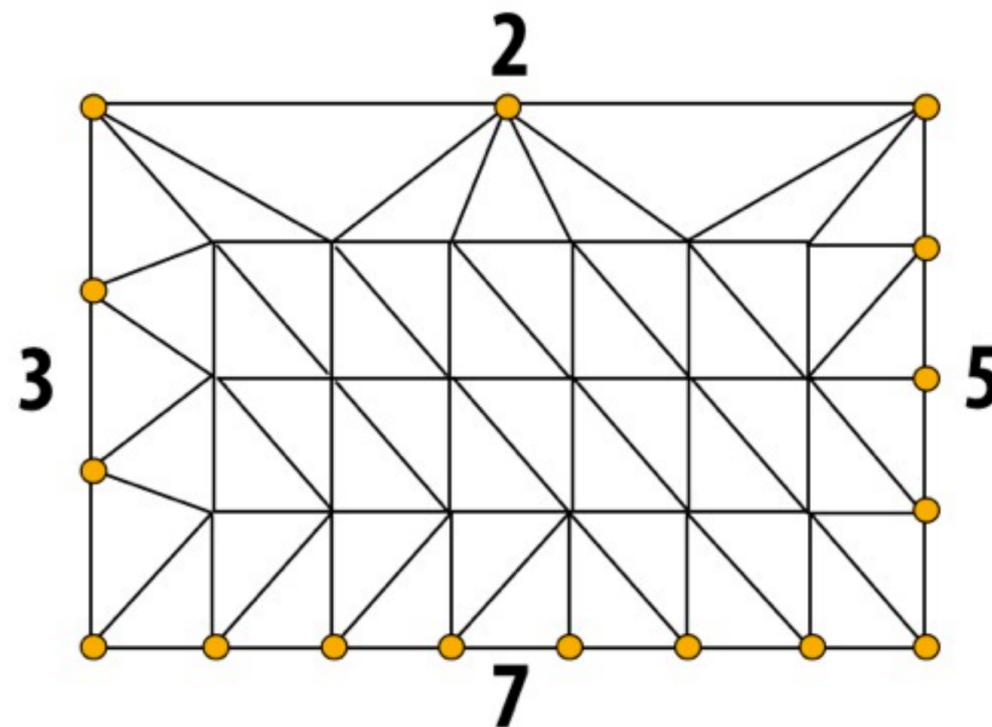
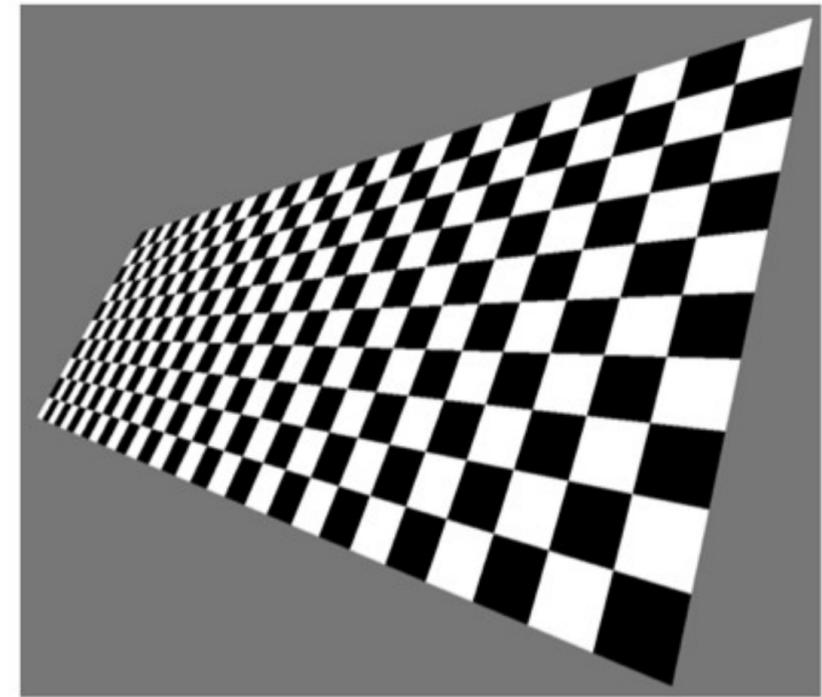
## Rasterization:

Determining what pixels a triangle overlaps



## Texture mapping:

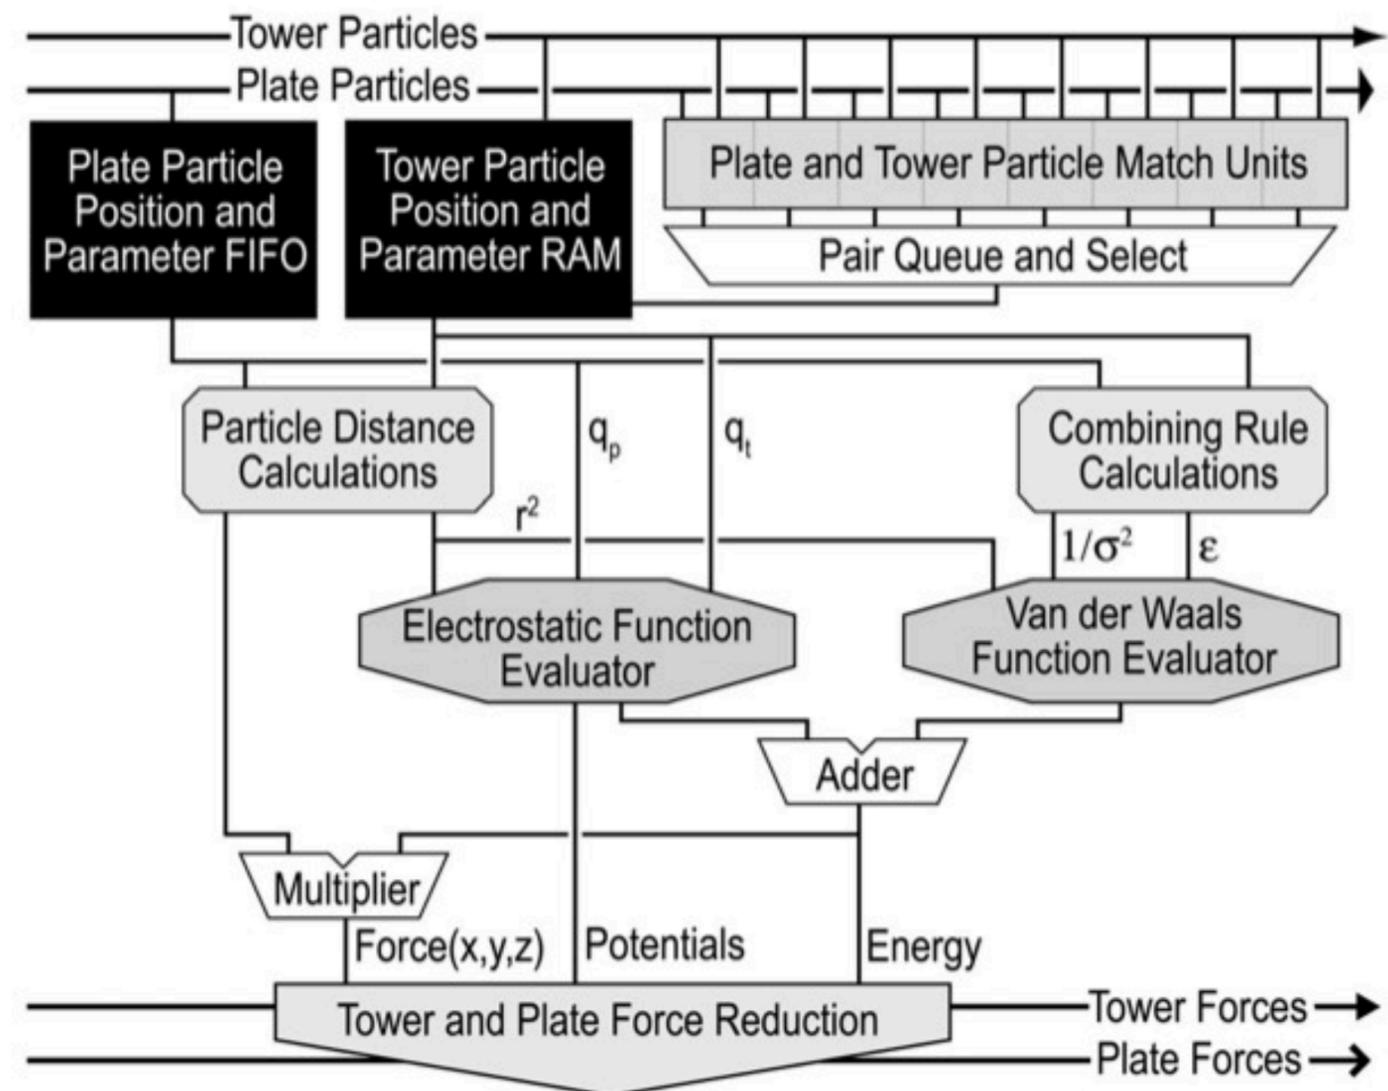
Warping/filtering images to apply detail to surfaces



**Geometric tessellation:**  
computing fine-scale geometry  
from coarse geometry

# DESRES Anton supercomputer

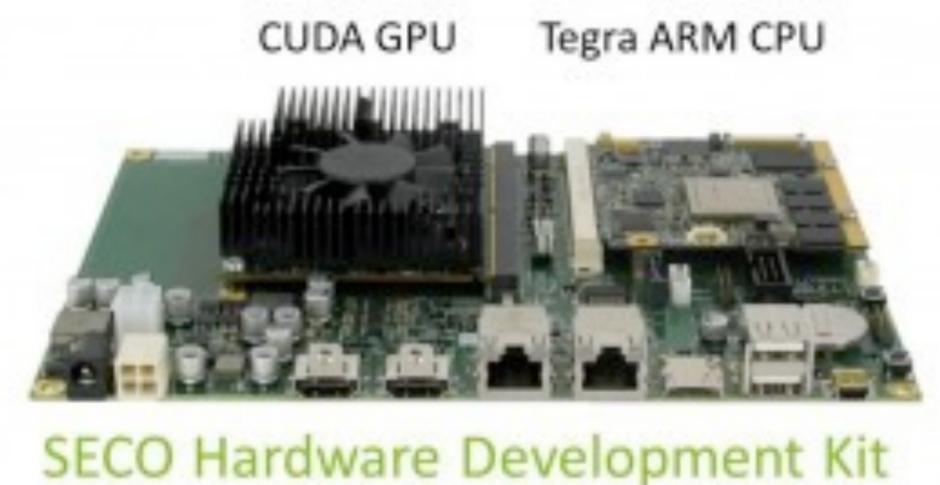
- Supercomputer highly specialized for molecular dynamics
  - Simulate time evolution of proteins
- ASIC for computing particle-particle interactions (512 of them in machine)
- Throughput-oriented subsystem for efficient fast-fourier transforms
- Custom, low-latency communication network designed for communication patterns of N-body simulations



# ARM + GPU Supercomputer

- **Observation: heavy lifting in supercomputing applications is the data-parallel part of workload**
  - **Less need for “beefy” sequential performance cores**
- **Idea: build supercomputer out of power-efficient building blocks**
  - **ARM CPUs (for control) + GPU cores (primary compute engine)**
- **Goal: 7 GFLOPS/Watt efficiency**
- **Project underway at Barcelona Supercomputing Center**

<http://www.montblanc-project.eu>



# Challenges of heterogeneity

## ■ So far in this course:

- **Homogeneous system: every processor can be used for every task**
- **Goal: to get best speedup, keep all processors busy all the time**

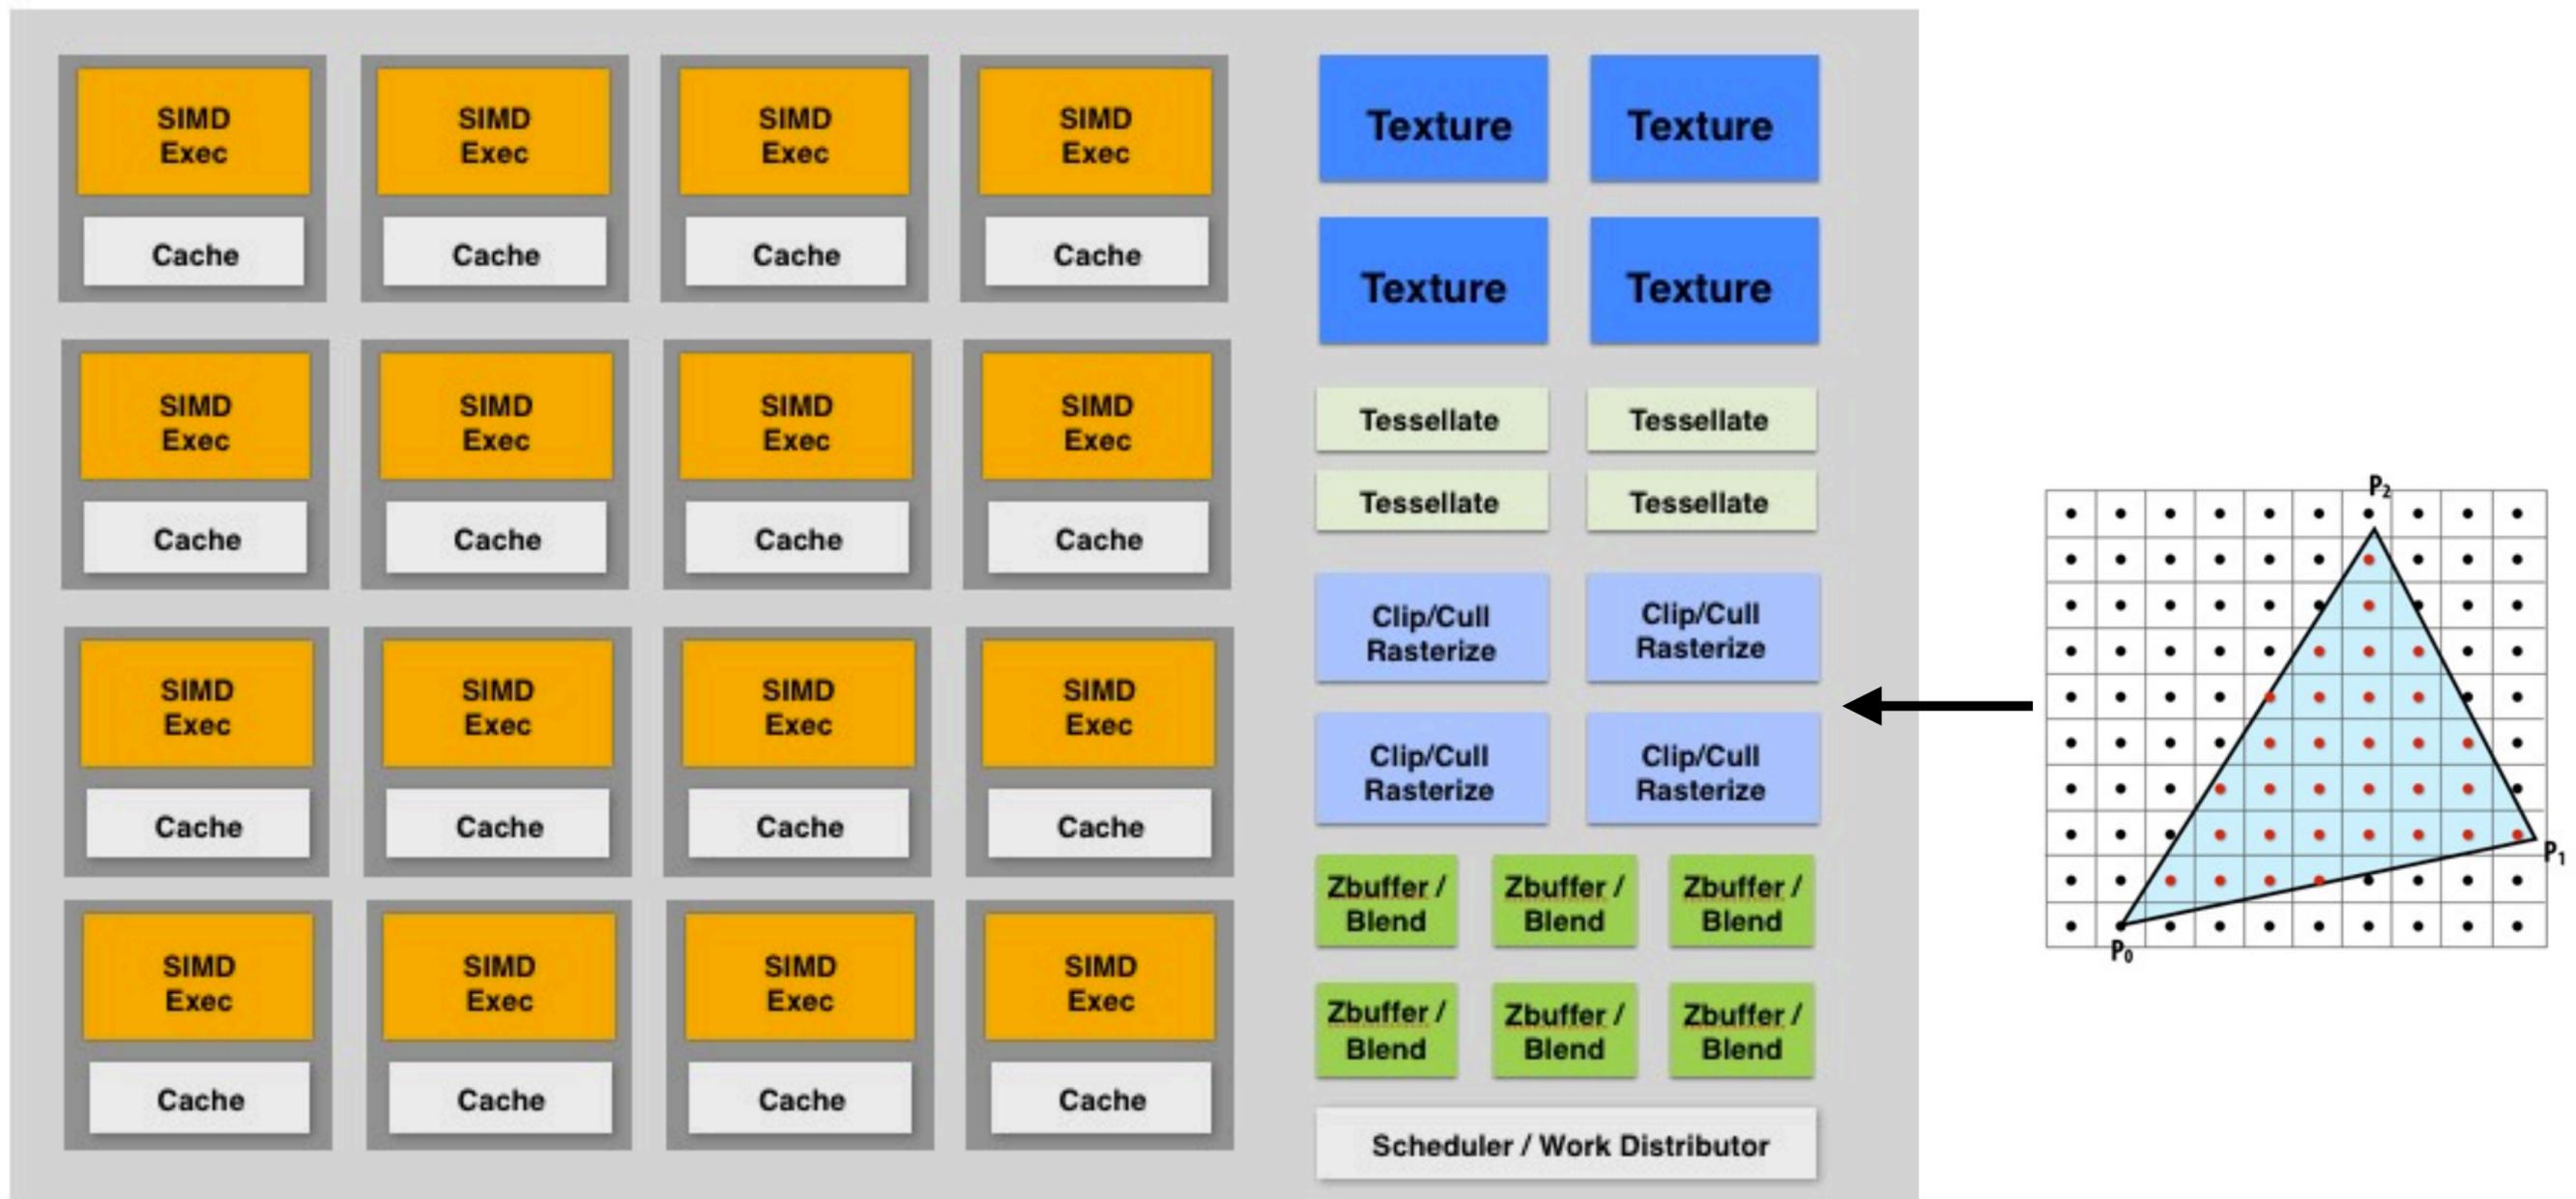
## ■ Heterogeneous system: use preferred processor for each task

- **Challenge for system designer: what is the right mixture of resources to meet performance, cost, and energy goals?**
  - **Too few throughput-oriented resources (lower peak performance/efficiency for parallel workloads -- should have used resources for more throughput cores)**
  - **Too few sequential processing resources (get bitten by Amdahl's Law)**
  - **How much chip area should be dedicated to a specific function, like video? (these resources are taken away from general-purpose processing)**

## ■ Increased pressure to understand workloads accurately at chip design time

# Pitfalls of heterogeneous design

[Molnar 2010]



**Say 10% of the computation is rasterization. (most of graphics workload is computing color of pixels)**

**Let's say you under-provision the fixed-function component of chip for rasterization:**

**(e.g., 1% of chip used for rasterizer, really needed 1.2%)**

**Problem: rasterization is bottleneck, so the expensive programmable processors (99% of chip) are idle waiting on rasterization. So the other 99% of the chip runs at 80% efficiency!**

**Tendency is to be conservative, and over-provision fixed-function components (diminishing their advantage)**

# Challenges of heterogeneity

## ■ Heterogeneous system: preferred processor for each task

- Challenge for system designer: what is the right mixture of resources?
  - Too few throughput oriented resources (lower peak throughput for parallel workloads)
  - Too few sequential processing resources (limited by sequential part of workload)
  - How much chip area should be dedicated to a specific function, like video? (these resources are taken away from general-purpose processing)
  - Work balance must be anticipated at chip design time
    - System cannot adapt to changes in usage over time, new algorithms, etc.
- Challenge to software developer: how to map programs onto a heterogeneous collection of resources?
  - Challenge: design of new algorithms that decompose well into components that each map well to different processing components of the machine
  - Scheduling problem is more complex on a heterogeneous system
  - Available mixture of resources can dictate choice of algorithm
  - Software portability nightmare

# Data movement has high energy cost

- **Rule of thumb in mobile system design: reduce amount of data transferred from memory**
  - Earlier in class we discussed minimizing communication to reduce stalls (poor performance). Now, we wish to reduce communication to reduce energy consumption
- **“Ballpark” numbers** [Sources: Bill Dally (NVIDIA), Tom Olson (ARM)]
  - Integer op: ~ 1 pJ \*
  - Floating point op: ~20 pJ \*
  - Reading 64 bits from small local SRAM (1mm away on chip): ~ 26 pJ
  - Reading 64 bits from low power mobile DRAM (LPDDR): ~1200 pJ
- **Implications**
  - Reading 10 GB/sec from memory: ~1.6 watts
  - Entire power budget for mobile GPU: ~1 watt (remember phone is also running CPU, display, radios, etc.)
  - iPhone 5 battery: ~5.5 watt-hours (note: my Macbook Pro laptop: 77 watt-hour battery)
  - Exploiting locality matters!!!

\* Cost to just perform the logical operation, not counting overhead of instruction decode, load data from registers, etc.

# Trends in energy-focused computing

## ■ Compute less!

- **Computing more costs energy: parallel algorithms that do more work than sequential counterparts may not be desirable even if they run faster**

## ■ Reduce bandwidth requirements

- **Exploit locality (restructure algorithms to reuse on-chip data as much as possible)**
- **Aggressive use of compression: perform extra computation to compress application data before transferring to memory (likely to see fixed-function HW to reduce overhead of general data compression/decompression)**

## ■ Specialize compute units:

- **Heterogeneous programmable processors: CPU-like cores + throughput-optimized cores (GPU-like cores)**
- **Fixed-function units: video decode/encode, image processing, sound processing**
- **Specialized instructions: expanding set of AVX vector instructions, new instructions for accelerating AES encryption (AES-NI)**
- **Increasing use of programmable logic: FPGAs**

# Fun reads about mobile power concerns

## ■ Power concerns in ARM Mali GPUs

- <http://blogs.arm.com/multimedia/780-how-low-can-you-go-building-low-power-low-bandwidth-arm-mali-gpus/>

## ■ Display Backlight measurements

- [http://www.displaymate.com/iPad\\_ShootOut\\_1.htm#Backlight\\_Power](http://www.displaymate.com/iPad_ShootOut_1.htm#Backlight_Power)

# Summary

- **Heterogeneous processing: use a mixture of computing resources that each fit with mixture of needs of target applications**
  - Latency-optimized sequential cores, throughput-optimized parallel cores, domain-specialized fixed-function processors
  - Examples exist throughout modern computing: mobile processors, desktop processors, supercomputers
- **Traditional rule of thumb in system design is to design simple, general-purpose components.**
  - This is not the case with emerging processing systems (perf/watt)
  - Today: want collection of components that meet perf requirement AND minimize energy use
- **Challenge of using these resources effectively is pushed up to the programmer**
  - Current CS research challenge: how to write efficient, portable programs for emerging heterogeneous architectures?