"New" Golden Age of Computer Architecture

Hung-Wei Tseng

Recap: Trends in the Dark Silicon Era

- Aggressive dynamic voltage/frequency scaling
- Throughout oriented slower, but more
 - GPUs
 - Single ISA heterogeneous CMP
- Just let it dark activate part of circuits, but not all
 - Modern GPU microarchitectures
- From general-purpose to domain-specific ASIC



In-Datacenter Performance Analysis of a Tensor Processing Unit

N. P. Jouppi, C. Young, N. Patil, D. Patterson, G. Agrawal, R. Bajwa, S. Bates, S. Bhatia, N. Boden, A. Borchers, R. Boyle, P.-I. Cantin, C. Chao, C. Clark, J. Coriell, M. Daley, M. Dau, J. Dean, B. Gelb, T. V. Ghaemmaghami, R. Gottipati, W. Gulland, R. Hagmann, C. R. Ho, D. Hogberg, J. Hu, R. Hundt, D. Hurt, J. Ibarz, A. Jaffey, A. Jaworski, A. Kaplan, H. Khaitan, D. Killebrew, A. Koch, N. Kumar, S. Lacy, J. Laudon, J. Law, D. Le, C. Leary, Z. Liu, K. Lucke, A. Lundin, G. MacKean, A. Maggiore, M. Mahony, K. Miller, R. Na-garajan, R. Narayanaswami, R. Ni, K. Nix, T. Norrie, M. Omernick, N. Penukonda, A. Phelps, J. Ross, M. Ross, A. Salek, E. Samadiani, C. Severn, G. Sizikov, M. Snelham, J. Souter, D. Steinberg, A. Swing, M. Tan, G. Thorson, B. Tian, H. Toma, E. Tuttle, V. Vasudevan, R. Wal- ter, W. Wang, E. Wilcox, and D. H. Yoon Google Inc.

https://www.pollev.com/hungweitseng close in 1:30 **TPU (Tensor Processing Unit)**

- Regarding TPUs, please identify how many of the following statements are correct.
 - ① TPU is optimized for highly accurate matrix multiplications
 - ② TPU is designed for dense matrices, not for sparse matrices
 - A majority of TPU's area is used by memory buffers 3
 - ④ All TPU instructions are equally long
 - A. 0
 - B. 1
 - C. 2
 - D. 3

TPU

А

С

D

Ε

What TPU looks like



TPU Floorplan



TPU Block diagram



TPU (Tensor Processing Unit)

- Regarding TPUs, please identify how many of the following statements are correct.
 - ① TPU is optimized for highly accurate matrix multiplications TPU is designed for dense matrices, not for sparse matrices A majority of TPU's area is used by memory buffers ④ All TPU instructions are equally long
 - A. 0
 - B. 1
- C. 2 D. 3
 - E. 4



Experimental setup

Model		Die								Benchmarked Servers					
	mm ²	nm	MHz	TDP	Measured		TOPS/s		CP/a	On-Chip	Dias	DRAMSing	מתד	Measured	
					Idle	Busy	8b	FP	GD/S	Memory	Dies	DRAM Size	IDF	Idle	Busy
Haswell E5-2699 v3	662	22	2300	145 W	41W	145W	2.6	1.3	51	51 MiB	2	256 GiB	504W	159W	455W
NVIDIA K80 (2 dies/card)	56 1	28	560	1 50W	25W	98W		2.8	160	8 MiB	8	256 GiB (host) + 12 GiB x 8	1838W	357W	991W
TPU	NA*	28	700	75W	28W	40W	92		34	28 MiB	4	256 GiB (host) + 8 GiB x 4	861W	290W	384W

Performance/Rooflines



TeraOps/sec (log scale)



- TPU Roofline
- K80 Roofline
- HSW Roofline

Tail latencies

Type	Batch	99th% Response	Inf/s (IPS)	% Max IPS
CPU	16	7.2 ms	5,482	42%
CPU	64	21.3 ms	13,194	100%
GPU	16	6.7 ms	13,461	37%
GPU	64	8.3 ms	36,465	100%
TPU	200	7.0 ms	225,000	80%
TPU	250	10.0 ms	280,000	100%

Table 4. 99-th% response time and per die throughput (IPS) for MLP0 as batch size varies for MLP0. The longest allowable latency is 7 ms. For the GPU and TPU, the maximum MLP0 throughput is limited by the host server overhead. Larger batch sizes increase throughput, but as the text explains, their longer response times exceed the limit, so CPUs and GPUs must use less-efficient, smaller batch sizes (16 vs. 200).



Tail latencies



- Why is this bad? Each user request experience tail is much higher
- If 99% of the server's response time is response

 - than 1 seconds.

 Tail Latency == 1 in X servers being slow now needs several servers – Changes of 10ms, but 1% of them take 1 second to

• If the user only needs one, the mean is OK If the user needs 100 partitions from 100 servers, 63% of the requests takes more

Tail latencies

Type	Batch	99th% Response	Inf/s (IPS)	% Max IPS
CPU	16	7.2 ms	5,482	42%
CPU	64	21.3 ms	13,194	100%
GPU	16	6.7 ms	13,461	37%
GPU	64	8.3 ms	36,465	100%
TPU	200	7.0 ms	225,000	80%
TPU	250	10.0 ms	280,000	100%

Table 4. 99-th% response time and per die throughput (IPS) for MLP0 as batch size varies for MLP0. The longest allowable latency is 7 ms. For the GPU and TPU, the maximum MLP0 throughput is limited by the host server overhead. Larger batch sizes increase throughput, but as the text explains, their longer response times exceed the limit, so CPUs and GPUs must use less-efficient, smaller batch sizes (16 vs. 200).



What NVIDIA says

https://blogs.nvidia.com/blog/2017/04/10/ai-drives-rise-accelerated-computing-datacenter/

	K80 2012	TPU 2015	P40 2016	
Inferences/Sec <10ms latency	¹ / ₁₃ X	1X	2X	
Training TOPS	6 FP32	NA	12 FP32	
Inference TOPS	6 FP32	90 INT8	48 INT8	
On-chip Memory	16 MB	24 MB	11 MB	
Power	300W	75W	250W	
Bandwidth	320 GB/S	While Google and NVIDIA chose different development paths, they both our approaches. Specifically:		

- · Al requires accelerated computing. Accelerators provide the significant data processing necessary to keep up with the growing demands of deep learning in an era when Moore's law is slowing.
- Tensor processing is at the core of delivering performance for deep learning training and inference.
- Tensor processing is a major new workload enterprises must consider when building modern data. centers.
- Accelerating tensor processing can dramatically reduce the cost of building modern data centers.

ral themes common to

7.10 Fallacies and Pitfalls

In these early days of both DSAs and DNNs, fallacies abound.

Fallacy It costs \$100 million to design a custom chip.

Figure 7.51 shows a graph from an article that debunks the widely quoted \$100million myth that it was "only" \$50 million, with most of the cost being salaries (Olofsson, 2011). Note that the author's estimate is for sophisticated processors that include features that DSAs by definition omit, so even if there were no improvement to the development process, you would expect the cost of a DSA design to be less.

Why are we more optimistic six years later, when, if anything, mask costs are even higher for the smaller process technologies?

First, software is the largest category, at almost a third of the cost. The availability of applications written in domain-specific languages allows the compilers to do most of the work of porting the applications to your DSA, as we saw for the TPU and Pixel Visual Core. The open RISC-V instruction set will also help reduce the cost of getting system software as well as cut the large IP costs.

Mask and fabrication costs can be saved by having multiple projects share a single reticle. As long as you have a small chip, amazingly enough, for \$30,000 anyone can get 100 untested parts in 28-nm TSMC technology (Patterson and Nikolić, 2015).

Fallacies & Pitfalls

- Fallacy: NN inference applications in data centers value throughput as much as response time.
- Fallacy: The K80 GPU architecture is a good match to NN inference GPU is throughput oriented
- Pitfall: For NN hardware, Inferences Per Second (IPS) is an inaccurate summary performance metric — it's simply the inverse of the complexity of the typical inference in the application (e.g., the number, size, and type of NN layers)
- Fallacy: The K80 GPU results would be much better if Boost mode were enabled Boost mode increased the clock rate by a factor of up to 1.6—from 560 to 875 MHz which increased performance by 1.4X, but it also raised power by 1.3X. The net gain in performance/Watt is 1.1X, and thus Boost mode would have a minor impact on LSTM1 Fallacy: CPU and GPU results would be comparable to the TPU if we used them more
- efficiently or compared to newer versions.

Fallacies & Pitfalls

- Pitfall: Architects have neglected important NN tasks.
 - CNNs constitute only about 5% of the representative NN workload for Google. More attention should be paid to MLPs and LSTMs. Repeating history, it's similar to when many architects concentrated on floating-point performance when most mainstream workloads turned out to be dominated by integer operations.
- Pitfall: Performance counters added as an afterthought for NN hardware.
- Fallacy: After two years of software tuning, the only path left to increase TPU performance is hardware upgrades.
- Pitfall: Being ignorant of architecture history when designing a domain-specific architecture
 - Systolic arrays
 - Decoupled-access/execute
 - CISC instructions

A Cloud-Scale Acceleration Architecture

Adrian Caulfield, Eric Chung, Andrew Putnam, Hari Angepat, Jeremy Fowers, Michael Haselman, Stephen Heil, Matt Humphrey, Puneet Kaur, Joo-Young Kim, Daniel Lo, Todd Massengill, Kalin Ovtcharov, Michael Papamichael, Lisa Woods, Sitaram Lanka, Derek Chiou, **Doug Burger Microsoft**



- Field Programmable Gate Array
 - An array of "Lookup tables (LUTs)"
 - Reconfigurable wires or say interconnects of LUTs
 - Registers
- An LUT
 - Accepts a few inputs
 - Has SRAM memory cells that store all possible outputs
 - Generates outputs according to the given inputs
- As a result, you may use FPGAs to emulate any kind of gates or logic combinations, and create an ASIC-like processor



FPGA

https://www.pollev.com/hungweitseng close in 1:30 MS' "Configurable Clouds"

- Regarding MS' configurable clouds that are powered by FPGAs, please identify how many of the following are correct
 - ① Each FPGA is dedicated to one machine
 - ② Each FPGA is connected through a network that is separated from the data center network
 - ③ FPGA can deliver shorter average latency for AES-CBC-128-SHA1 encryption and decryption than Intel's high-end processors
 - ④ FPGA-accelerated search queries are always faster than a pure software-based datacenter

A.	0	١.	
Β.	1		
C.	2		
D.	3		
E.	4		
			6 - C

catapult

A



С

D

Е

Configurable cloud



Traditional software (CPU) server plane

Gen2 shell

- Foundation for all accelerators
 - Includes PCIe, Networking and DDR IP
 - Common, well tested platform for development
- Lightweight Transport Layer
 - Reliable FPGA-to-FPGA Networking
 - Ack/Nack protocol, retransmit buffers
 - Optimized for lossless network
 - Minimized resource usage



Use cases

- Local: Great service acceleration
- Infrastructure: Fastest cloud network
- Remote: Reconfigurable app fabric (DNNs)

5 day bed-level latency

- Lower & more consistent <u>99.9th tail latency</u>
- In production for years

Even at 2× query load, accelerated ranking has lower latency than software at any load





Query Load (normalized to lowest throughput)



Accelerated networking

- Software defined networking
 - Generic Flow Table (GFT) rule based packet rewriting
 - 10x latency reduction vs software, CPU load now <1 core
 - 25Gb/s throughput at 25µs latency the fastest cloud network
- Capable of 40 Gb line rate encrypt and decrypt
 - On Haswell, AES GCM-128 costs 1.26 cycles/byte[1] (5+ 2.4Ghz cores to sustain 40Gb/s)
 - CBC and other algorithms are more expensive
 - AES CBC-128-SHA1 is 11µs in FPGA vs 4µs on CPU (1500B packet)
 - Higher latency, but significant CPU savings



VMs • 40G NIC GFT 40G Crypto

Shared DNN

- Economics: consolidation
 - Most accelerators have more throughput than a single host requires
 - Share excess capacity, use fewer instances
 - Frees up FPGAs for other use services
- DNN accelerator
 - Sustains 2.5x busy clients in microbenchmark, before queuing delay drives latency up





Oversubscription: # Remote Clients / # FPGAs

MS' "Configurable Clouds"

- Regarding MS' configurable clouds that are powered by FPGAs, please identify how many of the following are correct
 - ① Each FPGA is dedicated to one machine

manager. Instead, this paper focuses first on the hardware architecture necessary to treat remote FPGAs as available resources for global acceleration pools. We describe the com-

- ② Each FPGA is connected through a network that is separated from the data center network
- ③ FPGA can deliver shorter average latency for AES-CBC-128-SHA1 encryption and decryption than Intel's high-end processors
- FPGA-accelerated search queries are always faster than a pure software-based datacenter

A. 0 B. 1 C. 2 D. 3 E. 4

Datacenter networks handle multiple traffic classes and protocols, some of which expect near-lossless behavior. FP-GAs routing traffic between the server's NIC and TOR, as a bump-in-the-wire, must not interfere with the expected behavior of these various traffic classes. To that end, the LTL Protocol Engine shown in Figure 4 allows roles to send and receive packets from the network without affecting-and while supporting-existing datacenter protocols.

Our FPGA implementation supports full 40 Gb/s encryption and decryption. The worst case half-duplex FPGA crypto latency for AES-CBC-128-SHA1 is 11 μ s for a 1500B packet, from first flit to first flit. In software, based on the Intel numbers, it is approximately 4 μ s. AES-CBC-SHA1 is, however,

latencies begin exceeding acceptable thresholds. Because the FPGA is able to process requests while keeping latencies low, it is able to absorb more than twice the offered load, while executing queries at a latency that never exceeds the software datacenter at any load.







- Which of the following is the main reason why Microsoft adopts FPGAs instead of the alternatives chosen by their rivals?
 - A. Cost
 - B. Performance
 - C. Scalability
 - D. Flexibility
 - E. Easier to program

https://www.pollev.com/hungweitseng close in 1:30

FPGA

А

D

С



- Which of the following is the main reason why Microsoft adopts FPGAs instead of the alternatives chosen by their rivals?
 - A. Cost
 - B. Performance
 - C. Scalability
 - D. Flexibility
 - E. Easier to program

Why FPGA?

This model offers significant flexibility. From the local perspective, the FPGA is used as a compute or a network accelerator. From the global perspective, the FPGAs can be managed as a large-scale pool of resources, with acceleration

These programmable architectures allow for hardware homogeneity while allowing fungibility via software for different services. They must be highly *flexible* at the stem even hyperscale infrastructure. The acceleration system we describe is sufficiently flexible to cover three scenarios: local compute acceleration (through PCIe), network acceleration, and global application acceleration, through configuration as pools of emotel, accessible FPGAs. Local acceleration handles high-

This paper described Configurable Clouds, a datacenterscale acceleration architecture, based on FPGAs, that is both scalable and flexible. By putting in FPGA cards both in I/O In addition to architectural requirements that provide sufficient flexibility to justify scale production deployment, there are also physical restrictions in current infrastructures that

Summary: What makes a configurable cloud?

- Local, infrastructure and remote acceleration
 - Gen1 showed significant gains even for complex services (~2x for Bing)
 - Needs to have clear benefit for majority of servers: infrastructure
- Economics must work
 - What works at small scale doesn't always work at hyperscale and vice versa
 - Little tolerance for superfluous costs
 - Minimized complexity and risk in deployment and maintenance
- Must be flexible
 - Support simple, local accelerators and complex, shared accelerators at the same time
 - Rapid deployment of new protocols, algorithms and services across the cloud

Reflection: A New Golden Age for Computer Architecture

Current Challenges for Processor Architecture "If a problem has no solution, it may not be a problem, but a fact—not to be solved, but to be coped with over time." -Shimon Peres

https://www.pollev.com/hungweitseng close in 1:30 **Putting it all together**

- How many of the following would happen given the modern processor microarchitecture?
 - ① The branch predictor will predict not taken for branch A
 - The cache may contain the content of array2[array1[16] * 512]; (2)
 - temp can potentially become the value of array2[array1[16] *(3) 512];
 - The program will raise an exception (4)
 - unsigned int array1_size = 16; A. 0 uint8_t array1[160] = { 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16, 260}; uint8_t array2[256 * 512]; B. 1 void bar(size t x) { if (x < array1_size) { // Branch A: Taken if the statement is not going to be executed. C. 2 temp &= array2[array1[x] * 512]; D. 3 } void foo(size t x) { E. 4 int i = 0, j=0;for(j=0;j<10000;j++) bar(rand()%17); }



Putting it all together

- How many of the following would happen given the modern processor microarchitecture?
 - The branch predictor will predict not taken for branch A very likely The cache may contain the content of array2[array1[16] * 512];temp can potentially become the value of array2[array1[16] *(3) — not really, as x < array1_size 512]; The program will raise an exception - maybe? A. 0

```
B. 1
```

```
D. 3
```

E. 4

```
unsigned int array1_size = 16;
uint8_t array1[160] = { 1, 2, 3, 4, 5, 6, 7, 8, 9, 10, 11, 12, 13, 14, 15, 16, 260};
uint8_t array2[256 * 512];
void bar(size t x) {
  if (x < array1_size) { // Branch A: Taken if the statement is not going to be executed.
    temp &= array2[array1[x] * 512];
}
void foo(size t x) {
```

```
int i = 0, j=0;
    for(j=0;j<10000;j++)
            bar(rand()%17);
}
```



Future Opportunities in Computer Architecture

- End of Moore's Law and Dennard Scaling
- Overlooked Security
- Domain-Specific Languages & Domain-specific architectures.
- Open Architectures, Open-source implementations
- Agile Hardware Development
- Great for architects in academia & industry!


Final words

Conclusion

- Computer architecture is now more important than you could ever imagine
- Being a "programmer" is easy. You need to know architecture a lot to be a "performance programmer"
 - Branch prediction
 - Cache
- Multicore era to get your multithreaded program correct and perform well, you need to take care of coherence and consistency
- We're now in the "dark silicon era"
 - Single-core isn't getting any faster
 - Multi-core doesn't scale anymore
 - We will see more and more ASICs
 - You need to write more "system-level" programs to use these new ASICs.

Thank you all for this great quarter! Let's take a group photo now!

One more thing...



Prefetcher contest

- Top Li, Guiquan 1.22x speedup (avg.)/1.37x(max)
 - Global History Prefetcher
 - Kyle J, Nesbit and James E.Smith, Data Cache Prefetching Using a Global History Buffer
 - Martin Dimitrov, Huiyang Zhou, Combining Local and Global History for High Performance Data, Prefetching
- Runner-up Jung, Boram 1.202x speedup (avg.)/1.41x (max)
 - Next two/four lines prefetcher
- Honorable mention Yu, Hongmiao 1.201x speedup (avg.)/1.32x(max)
 - Local History Buffer(LHB) indexed by instruction PC
 - Dimitrov, Martin, and Huiyang Zhou. "Combining local and global history for high performance data prefetching."

Historical record

- Per-PC stride-based predictor 1.31x (avg.)/2.21x (max)
- Sequential/anti-sequential prefetcher 1.29x(avg.)/2.25x (max.)
- Aggressive stream buffer 1.29x (avg.)/2.09x(max)

Sample Final

Format of the final

- Multiple choices (20 questions)
 - They're like your clicker/midterm multiple choices questions
 - Cumulative, don't forget your midterm and midterm review
- Homework style calculation/operation based questions * 2 problem sets, 10 questions in total
 - They are also MSCS comprehensive exam questions
- Brief discussion/Open-ended * 8
 - Explain your answer using less than 100 words. Some of them must be as short as 30 words
 - May not have a standard answer. You need to understand the concepts to provide a good answer

e of them must be as

eview stions * 2 problem

Multiple choices



How many dependencies do we have?

How many pairs of data dependences are there in the following RISC-V instructions?

ld	X6,	0(X1	10)
add	X7,	Χ6,	X12
sd	X7,	0(X1	10)
addi	X10,X	X10,	8
bne	X10,	X5,	LOOP

- B. 2
- C. 3

D. 4

E. 5



False dependencies

- Consider the following dynamic instructions
 - X12, 0(X20) 1 ld
 - X12, X10, X12 2 add
 - ③ sub X18, X12, X10
 - ④ ld X12, 8(X20)
 - ⑤ add X14, X18, X12
 - X18, X14, X14 add
 - X14, 16(X20) ⊘ sd
 - addi X20, X20, 8

which of the following pair is not a "false dependency"

- A. (1) and (4)
- B. (1) and (8)
- C. (5) and (7)
- D. (4) and (8)
- E. (7) and (8)



What about "linked list"

• For the following C code and it's translation in RISC-V, how many cycles it takes the processor to issue all instructions? Assume the current PC is already at the first instruction and this linked list has only three nodes. This processor can fetch 2 instruction per cycle, with exactly the same register renaming hardware and pipeline as we showed previously.

```
do {
    number of nodes++;
    current = current->next;
} while ( current != NULL )
```

- A. 9
- B. 10
- C. 11
- D. 12
- E. 13

LOOP: 1d X10, 8(X10) addi X7, X7, 1 bne X10, X0, LOOP

CMP advantages

- How many of the following are advantages of CMP over traditional superscalar processor
 - ① CMP can provide better energy-efficiency within the same area
 - ② CMP can deliver better instruction throughput within the same die area (chip size)
 - ③ CMP can achieve better ILP for each running thread
 - ④ CMP can improve the performance of a single-threaded application without modifying code
 - A. 0
 - B. 1
 - C. 2

D. 3

E. 4

How good is SS/OoO/ROB with this code?

Consider the following dynamic instructions

- ld X1, 0(X10) 1
- ② addi X10, X10, 8
- 3 add X20, X20, X1
- bne X10, X2, LOOP 4

Assume a superscalar processor with issue width as 2 & unlimited physical registers that can fetch up to 4 instructions per cycle, 3 cycles to execute a memory instruction and the loop will execute for 10,000 times, what's the average CPI?

- A. 0.5
- B. 0.75
- C. 1
- D. 1.25
- E. 1.5



Amdahl's Law on Multicore Architectures

- Regarding Amdahl's Law on multicore architectures, how many of the following statements is/are correct?
 - ① If we have unlimited parallelism, the performance of each parallel piece does not matter as long as the performance slowdown in each piece is bounded
 - ② With unlimited amount of parallel hardware units, single-core performance does not matter anymore
 - ③ With unlimited amount of parallel hardware units, the maximum speedup will be bounded by the fraction of parallel parts
 - ④ With unlimited amount of parallel hardware units, the effect of scheduling and data exchange overhead is minor
 - A. 0
 - B. 1
 - C. 2
 - D. 3
 - E. 4

Summary of Optimizations

- Regarding the following cache optimizations, how many of them would help improve miss rate?
 - ① Non-blocking/pipelined/multibanked cache
 - ② Critical word first and early restart
 - ③ Prefetching
 - ④ Write buffer
 - A. 0
 - **B**. 1
 - C. 2

D. 3

E. 4



Virtual indexed, physical tagged cache limits the cache size

- If you want to build a virtual indexed, physical tagged cache with 32KB capacity, which of the following configuration is possible? Assume the system use 4K pages.
 - A. 32B blocks, 2-way
 - B. 32B blocks, 4-way
 - C. 64B blocks, 4-way
 - D. 64B blocks, 8-way



Power & Energy

- Regarding power and energy, how many of the following statements are correct?
 - ① Lowering the power consumption helps extending the battery life
 - Lowering the power consumption helps reducing the heat generation (2)
 - ③ Lowering the energy consumption helps reducing the electricity bill
 - ④ A CPU with 10% utilization can still consume 33% of the peak power
 - A. 0
 - B. 1
 - C. 2

D. 3

E. 4

Why is D better than C?

- How many of the following statements explains the main reason why B outperforms C with compiler optimizations
 - ① D has lower dynamic instruction count than C
 - ② D has significantly lower branch mis-prediction rate than C
 - ③ D has significantly fewer branch instructions than C
 - ④ D can incur fewer memory accesses than C

A. 0 B. 1 C. 2 D. 3 E. 4

ပ	<pre>inline int popcount(uint64_t x) { int c = 0; int table[16] = {0, 1, 1, 2, 1, 2, 2, 3, 1, 2, 2, 3, 2, 3, 3, 4}; while(x)</pre>	<pre>inline int o int o int d int</pre>
	67	



t popcount(uint64_t x) { c = 0; table[16] = {0, 1, 1, 2, 1, 1, 2, 2, 3, 2, 3, 3, 4; (uint64_t i = 0; i < 16; i++) += table[(x & $0 \times F$)]; = x >> 4;

rn c;

Demo revisited

- Why the performance is better when option is not "0"
 - ① The amount of dynamic instructions needs to execute is a lot smaller
 - ② The amount of branch instructions to execute is smaller
 - ③ The amount of branch mis-predictions is smaller
 - ④ The amount of data accesses is smaller

}

```
A. 0 if(option)
    std::sort(data, data + arraySize);
B. 1
C. 2 for (unsigned i = 0; i < 100000; ++i) {
    int threshold = std::rand();
D. 3 for (unsigned i = 0; i < arraySize; ++i) {
        if (data[i] >= threshold)
            sum ++;
```

ot "O" ecute is a lot smaller smaller

Why can't we proceed without stalls/no-ops?

- How many of the following statements are true regarding why we have to stall for each branch in the current pipeline processor
 - ① The target address when branch is taken is not available for instruction fetch stage of the next cycle
 - ② The target address when branch is not-taken is not available for instruction fetch stage of the next cycle
 - ③ The branch outcome cannot be decided until the comparison result of ALU is not out
 - ④ The next instruction needs the branch instruction to write back its result
 - A. 0
 - B. 1
 - C. 2
 - D. 3

E. 4

n result of ALU is not out ck its result

What if the code look like this?

- D-L1 Cache configuration of AMD Phenom II
 - Size 64KB, 2-way set associativity, 64B block, LRU policy, write-allocate, write-back, and assuming 64-bit address.

```
int a[16384], b[16384], c[16384];
/* c = 0x10000, a = 0x20000, b = 0x30000 */
for(i = 0; i < 512; i++)
    c[i] = a[i]; //load a and then store to c
for(i = 0; i < 512; i++)
    c[i] += b[i]; //load b, load c, add, and then store to c
```

What's the data cache miss rate for this code?

- A. 6.25%
- B. 56.25%
- C. 66.67%
- 68.75% D.
- E. 100%



What kind(s) of misses can matrix transpose remove?

• By transposing a matrix, the performance of matrix multiplication can be further improved. What kind(s) of cache misses does matrix transpose help to remove?

```
for(i = 0; i < ARRAY_SIZE; i+=(ARRAY_SIZE/n)) {</pre>
  for(j = 0; j < ARRAY_SIZE; j+=(ARRAY_SIZE/n)) {</pre>
    for(k = 0; k < ARRAY_SIZE; k+=(ARRAY_SIZE/n)) {</pre>
      for(ii = i; ii < i+(ARRAY_SIZE/n); ii++)</pre>
         for(jj = j; jj < j+(ARRAY_SIZE/n); jj++)</pre>
           for(kk = k; kk < k+(ARRAY_SIZE/n); kk++)</pre>
             c[ii][jj] += a[ii][kk]*b[kk][jj];
}
```

- A. Compulsory miss
- B. Capacity miss

Block

- C. Conflict miss
- D. Capacity & conflict miss
- E. Compulsory & conflict miss



for(i = 0; i < ARRAY_SIZE; i+=(ARRAY_SIZE/n)) {</pre> for(j = 0; j < ARRAY_SIZE; j+=(ARRAY_SIZE/n)) {</pre>

for(i = 0; i < ARRAY_SIZE; i+=(ARRAY_SIZE/n)) {</pre> for(j = 0; j < ARRAY_SIZE; j+=(ARRAY_SIZE/n)) {</pre> for(k = 0; k < ARRAY_SIZE; k+=(ARRAY_SIZE/n)) {</pre> for(ii = i; ii < i+(ARRAY_SIZE/n); ii++)</pre> for(jj = j; jj < j+(ARRAY_SIZE/n); jj++)</pre> for(kk = k; kk < k+(ARRAY_SIZE/n); kk++</pre> // Compute on b_t c[ii][jj] += a[ii][kk]*b_t[jj][kk];

MS' "Configurable Clouds"

- Regarding MS' configurable clouds that are powered by FPGAs, please identify how many of the following are correct
 - ① Each FPGA is dedicated to one machine
 - ② Each FPGA is connected through a network that is separated from the data center network
 - ③ FPGA can deliver shorter average latency for AES-CBC-128-SHA1 encryption and decryption than Intel's high-end processors
 - ④ FPGA-accelerated search queries are always faster than a pure software-based datacenter
 - A. 0
 - B. 1
 - C. 2
 - D. 3
 - E. 4



Summary of Optimizations

- Regarding the following cache optimizations, how many of them would help improve miss rate?
 - ① Non-blocking/pipelined/multibanked cache
 - ② Critical word first and early restart
 - ③ Prefetching
 - ④ Write buffer
 - A. 0
 - **B**. 1
 - C. 2

D. 3

E. 4



What data structure is performing better

	Array of objects	object	
	<pre>struct grades { int id; double *homework; double average; };</pre>	<pre>struct grades { int *id; double **homework; double *average; };</pre>	
average of each homework	<pre>for(i=0;i<homework_items; (double)total_number_students;="" +="gradesheet[j].homework[i];" =="" for(j="0;j<total_number_students;j++)" gradesheet[total_number_students].homework[i]="" i++)="" pre="" {="" }<=""></homework_items;></pre>	<pre>for(i = 0;i < homework_items; i+ { gradesheet.homework[i][total_n for(j = 0; j <total_number_stu gradesheet.homework[i][j];="" gradesheet.homework[i][tot="" pre="" total_number_students;="" {="" }="" }<=""></total_number_stu></pre>	

- Considering your workload would like to calculate the average score of one of the homework for all students, which data structure would deliver better performance?
 - A. Array of objects
 - B. Object of arrays

of arrays

+)

umber_students] = 0.0; dents; j++)

cal_number_students] +=

al_number_students] /=

3Cs and A, B, C

- Regarding 3Cs: compulsory, conflict and capacity misses and A, B, C: associativity, block size, capacity How many of the following are correct?
 - ① Increasing associativity can reduce conflict misses
 - Increasing associativity can reduce hit time 2
 - Increasing block size can increase the miss penalty 3
 - Increasing block size can reduce compulsory misses (4)
 - A. 0
 - **B**. 1
 - C. 2
 - D. 3

E. 4

Cache coherency

• Assuming that we are running the following code on a CMP with a cache coherency protocol, how many of the following outputs are possible? (a is initialized to 0 as assume we will output more than 10 numbers)

thread 1	
while(1) printf("%d ",a);	while(1) a++;
① 0123456789	
② 1259368101213	
③ 111111164100	
④ 11111111100	
A. 0	
B. 1	
C. 2	
D. 3	
E. 4	

thread 2

Performance comparison

 Comparing implementations of thread_vadd — L and R, please identify which one will be performing better and why Version L

```
void *threaded_vadd(void *thread_id)
 int tid = *(int *)thread_id;
 int i;
  for(i=tid;i<ARRAY_SIZE;i+=NUM_OF_THREADS)</pre>
        c[i] = a[i] + b[i];
  return NULL;
```

```
void *threaded_vadd(void *thread_id)
  int tid = *(int *)thread_id;
  int i;
  for(i=tid*(ARRAY_SIZE/NUM_OF_THREADS);i<(tid+1)*(ARRAY_SIZE/NUM_OF_THREADS);i++)</pre>
      c[i] = a[i] + b[i];
  return NULL;
```

- A. L is better, because the cache miss rate is lower
- B. R is better, because the cache miss rate is lower
- C. L is better, because the instruction count is lower
- D. R is better, because the instruction count is lower
- E. Both are about the same

```
{
 tids[i] = i;
```



Version R



Free-answer questions



Register renaming

Draw the pipeline diagram for the following instructions

1	Loop:LD	F1,0(X3)
2	FADD	F2,F1,F4
3	FMUL	F1,F2,F6
4	FADD	F1,F1,F5
5	FADD	F7,F7,F1
6	ADD	X2,X2,-1
0	BNEZ	X2,Loop
8	ADDI	X6,X6,4
9	LD	F3,0(X6)

- Assume we have a dual-fetch, dual-issue, out-of-order pipeline where
 - INT ALU takes 1 cycle
 - FP ALU takes 3 cycles
 - MEM pipeline: AR-AQ-MEM 3 cycles in total
 - BR takes 1 cycle to resolve
- If the loop is taken twice, how many cycles it takes to issue all instructions?
- If the loop is taken 100 times, what's the average CPI?

Best cache configuration

Consider the following code. Integers and pointers are both 4 bytes.

```
struct List {
   List * next;
   int data;
}
```

```
void foo(List *head) {
   List * cur = head;
   while(cur->next) {
      cur = cur->next;
   }
}
```

• For a given total cache size, what cache line size will provide the best performance for this code? (hint: Your answer should not depend on the number of lines or the associativity of the cache.)



Reverse caching

- Below, we have given you four different sequences of addresses generated by a program running on a processor with a data cache. Cache hit ratio for each sequence is also shown below. Assuming that the cache is initially empty at the beginning of each sequence, find out the following parameters of the processor's data cache (ensure that you sufficiently explain your answer)
 - Associativity (1, 2, or 4 ways)
 - Block size (1, 2, 4, 8, 16, or 32 bytes)
 - Total cache size (256B, or 512B)
 - Replacement policy (LRU or FIFO)
- 1. Address Sequence 1: [0, 2, 4, 8, 16, 32] Hit Ratio: 0.33
- 2. Address Sequence 2: [0, 512, 1024, 1536, 2048, 1536, 1024, 512, 0] Hit Ratio: 0.33
- 3. Address Sequence 3: [0, 64, 128, 256, 512, 256, 128, 64, 0] Hit Ratio: 0.33
- 4. Address Sequence 4: [0, 512, 1024, 0, 1536, 0, 2048, 512] Hit Ratio: 0.25

Open-ended questions



Code and cache miss rate

 Assume my cache has 16KB capacity, 16 byte block size and is 2-way set associative. Integers are 4 bytes. Give the C code for a loop that has a very poor hit rate in this cache but whose hit rate raises to almost 100% if we double the capacity to 32KB.



Branch predictions

- Increasing the size of a branch predictor typically reduces the chances of "aliasing" -- two branches sharing the same predictor. Usually, sharing results in negative interference (decreased prediction) accuracy), but sometimes it can result in positive interference. Assuming a PC-indexed table of 2-bit predictors
 - Give an example of two branches (eg, show the T, N patterns for each, and how they are interleaved) that would result in positive interference (increased overall prediction accuracy).
 - Give an example of two branches that would result in negative interference.
 - Explain why most of the time you would expect to see negative interference with real code.
SMT v.s. CMP

- Both CMP & SMT exploit thread-level or task-level parallelism. Assuming both application X and application Y have similar instruction combination, say 60% ALU, 20% load/store, and 20% branches. Consider two processors:
 - P1: CMP with a 2-issue pipeline on each core. Each core has a private L1 32KB D-cache
 - P2: SMT with a 4-issue pipeline. 64KB L1 D-cache

Which one do you think is better?

- A. P1
- B. P2

Other open-ended questions

- Given the instruction front-end is decoupled from the backend of the pipeline ALUs, do you think ISA still affect performance?
 - Emily Blem, Jaikrishnan Menon, and Karthikeyan Sankaralingam. 2013. Power struggles: Revisiting the RISC vs. CISC debate on contemporary ARM and x86 architectures. In Proceedings of the 2013 IEEE 19th International Symposium on High Performance Computer Architecture (HPCA) (HPCA) '13). https://minds.wisconsin.edu/handle/1793/64923
 - Ashish Venkat and Dean M. Tullsen. 2014. Harnessing ISA diversity: design of a heterogeneous-ISA chip multiprocessor. In Proceeding of the 41st annual international symposium on Computer architecuture (ISCA '14). http://www.cs.virginia.edu/venkat/papers/isca2014.pdf
- What features in modern processor architecture enable the potential of "Meltdown and Spectre" attacks? Should we live without those features? How to solve these security issues?
 - M. D. Hill, J. Masters, P. Ranganathan, P. Turner and J. L. Hennessy, "On the Spectre and Meltdown Processor Security Vulnerabilities," in IEEE Micro, vol. 39, no. 2. http://pages.cs.wisc.edu/~markhill/papers/ieeemicro19_spectre_meltdown_2019_01_30
- What compiler optimizations would not be effective given OoO execution hardware?



Other open-ended questions

- Can you name and briefly describe a few "trends" in the dark silicon era?
- If you're asked to design a machine learning hardware, what will you do?
- If you're asked to build an Xeon Phi type processor where each core also has many-way SMT, are you going to give the processor more cache or better branch predictor?
- Can we focus on improving the throughput of computing instead of latency? Can you give an example on what type of applications will not work well in this way
- Pros and cons for branch prediction using perceptrons?



Announcement — final exam

- Final Exam
 - The final can be opened only once -- if you accidentally close the browser or the browser crashes or your lose Internet connection, you cannot re-initiate it and we WILL NOT help you for these cases. Browsers crash and accidental closing of tabs occur a lot when you have many opened tabs. Please be careful. • Q21 – Q30 are comprehensive exam questions – You must receive at least 60% from Q21–Q25 AND 60%
 - from Q26-Q30 to be considered as PASS
 - This final covers EVERYTHING mentioned/assigned this quarter.
 - This is an open-book, open-note test, but again, the more you open, the higher chance your computer will have issues.
 - We have MANY questions for you, but you only have a total of 180 minutes to finish. Heavily rely on your notes/ book/cheatsheets is not a good idea.
 - Please show your work if appropriate -- we don't give credits to answers only have the final result
 - There is no partial credits for multiple choice questions. Please think thoroughly.
 - Reference online solution, discuss with ANY other human being or digital assistant (e.g, Siri, Google Home, Alexa or whatever you name it) is considered as cheating.
 - We will not automatically submit your test when time is up. If your submission is late by x sec, your grade is max(raw_score * ((100-x)/100),0)



Announcement

- Project due tonight try your best
- Assignment #4 due this Wednesday
- iEVAL, until 12/3
 - Please fill the survey to let us know your opinion!
 - Don't forget to take a screenshot of your submission and submit through iLearn it counts as a **full credit assignment**
 - We will drop your lowest 2 assignment grades
- Final Exam
 - Starting from 12/6 to 12/10 12:00pm, any consecutive 180 minutes you pick
 - Similar to the midterm, but more time and about 1.5x longer
 - Two of the problem sets will be comprehensive exam questions
 - Will release a sample final at the end of the last lecture

Computer Science & Engineering





