The Memory Hierarchy

CSC 211 – December 9, 2020

Datapath Lab Q&A

Pay attention to file names when you submit.

You can use MARS to test assembly programs

- don't set $sp to 0xFF80
- printing to terminal uses the syscall instruction in MARS, not sb to 0xFF

When Caches Fail

Cache Misses

What are the three types of cache misses? They all start with C.

Capacity Miss – the cache runs out of space, so it had to evict something we are now accessing.

Compulsory Miss – this is the first time accessessing a location, so it is not yet in the cache.

Conflict Miss – the value used to be in the cache, but another cached value replaced it. The cache may not have been full at the time.

Array Accesses

```
#define N ___
int arr[N] = {...};

for(int i=0; i<N; i++) {
  printf("%d\n", arr[i]);
}

for(int i=0; i<N; i++) {
  printf("%d\n", arr[i]);
}
```

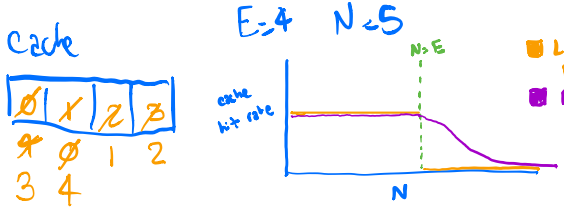

Imagine this code executing on a processor with a fully-associative cache that uses 4-byte cache lines, and has E cache entries.

How does the cache perform when...

N is less than E

- N misses in the first loop (compulsory misses)
- 0 misses in the second loop

N is equal to E

- N misses in the first loop (compulsory)
- 0 misses in the second loop

N is larger than E

- N misses in the first loop (compulsory)
- N misses in the second loop (capacity / conflict)

Array Accesses

```
#define N ___
int arr[N] = {...};

for(int i=0; i<N; i++) {
  printf("%d\n", arr[i]);
}

for(int i=0; i<N; i++) {
  printf("%d\n", arr[i]);
}
```

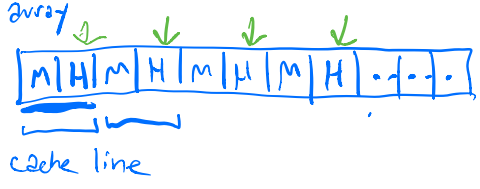

Using LRU replacement

What if we use the same fully-associative cache, but now with 8-byte cache lines?

What values of N lead to the following outcomes? Express N in terms of E.

N/2 cache misses in the first loop, none in the second.

25% miss rate

$N \leq E \cdot 2$

N/2 cache misses in the first loop, N/2 cache misses in the second loop.

50% miss rate

$N > 2 \cdot E$

spatial locality helps, even though we over-fill the cache

Multi-Level Caching

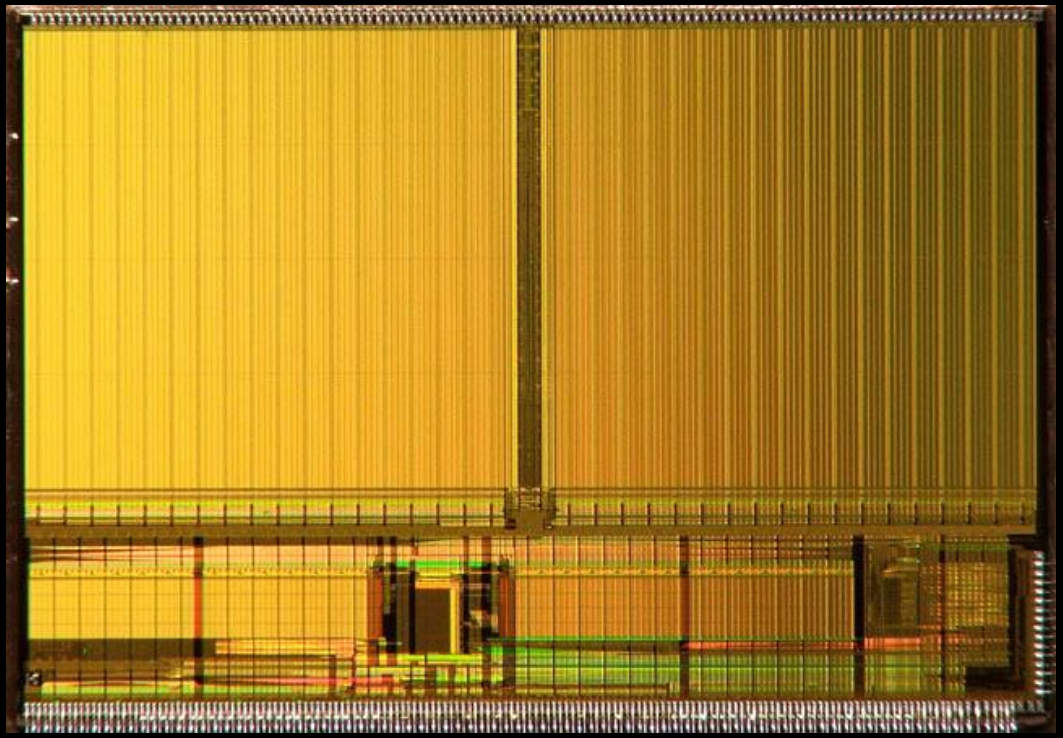

Intel Pentium Pro 256KB Cache (source: CPU-World)

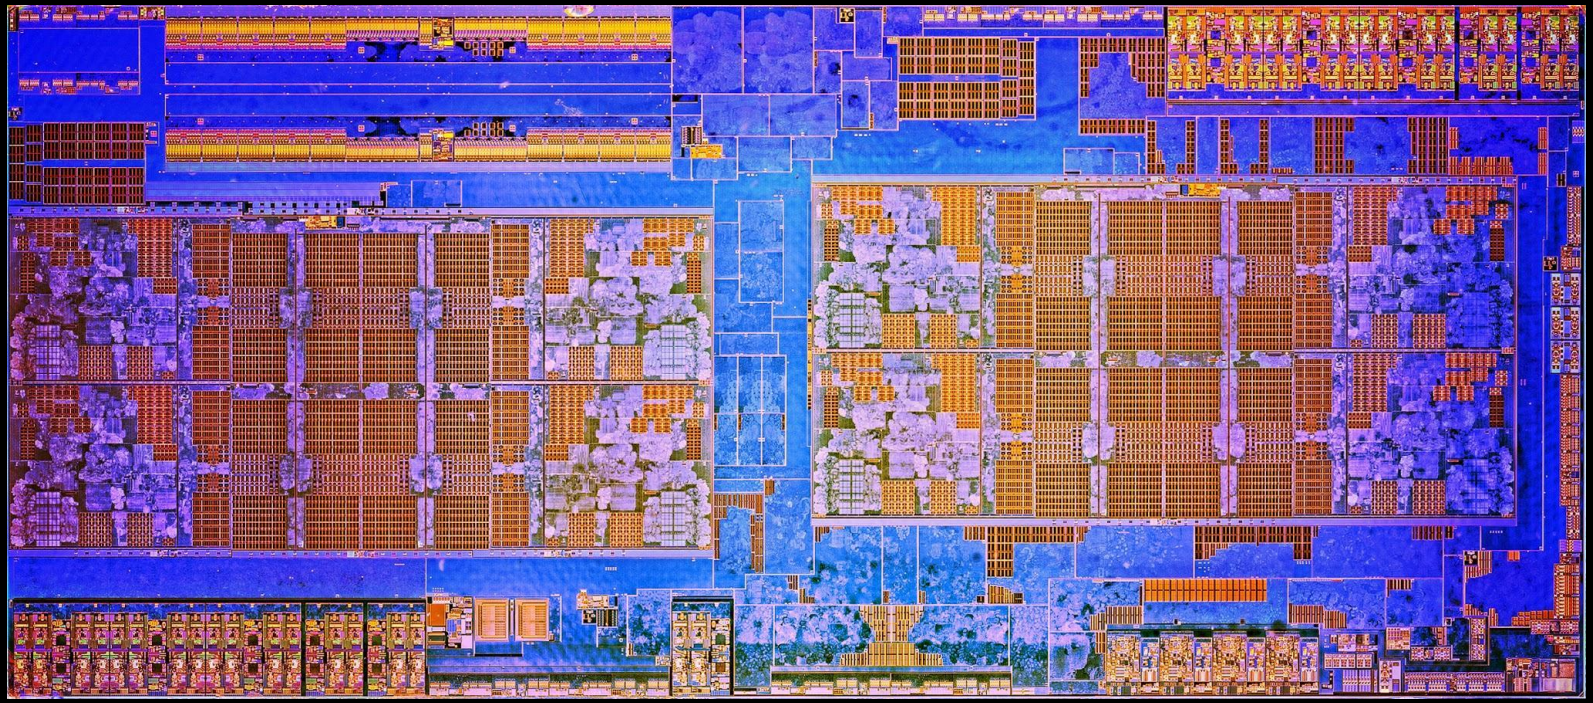

AMD Ryzen 4-Core Unit (source: PCWorld)

Virtual Memory