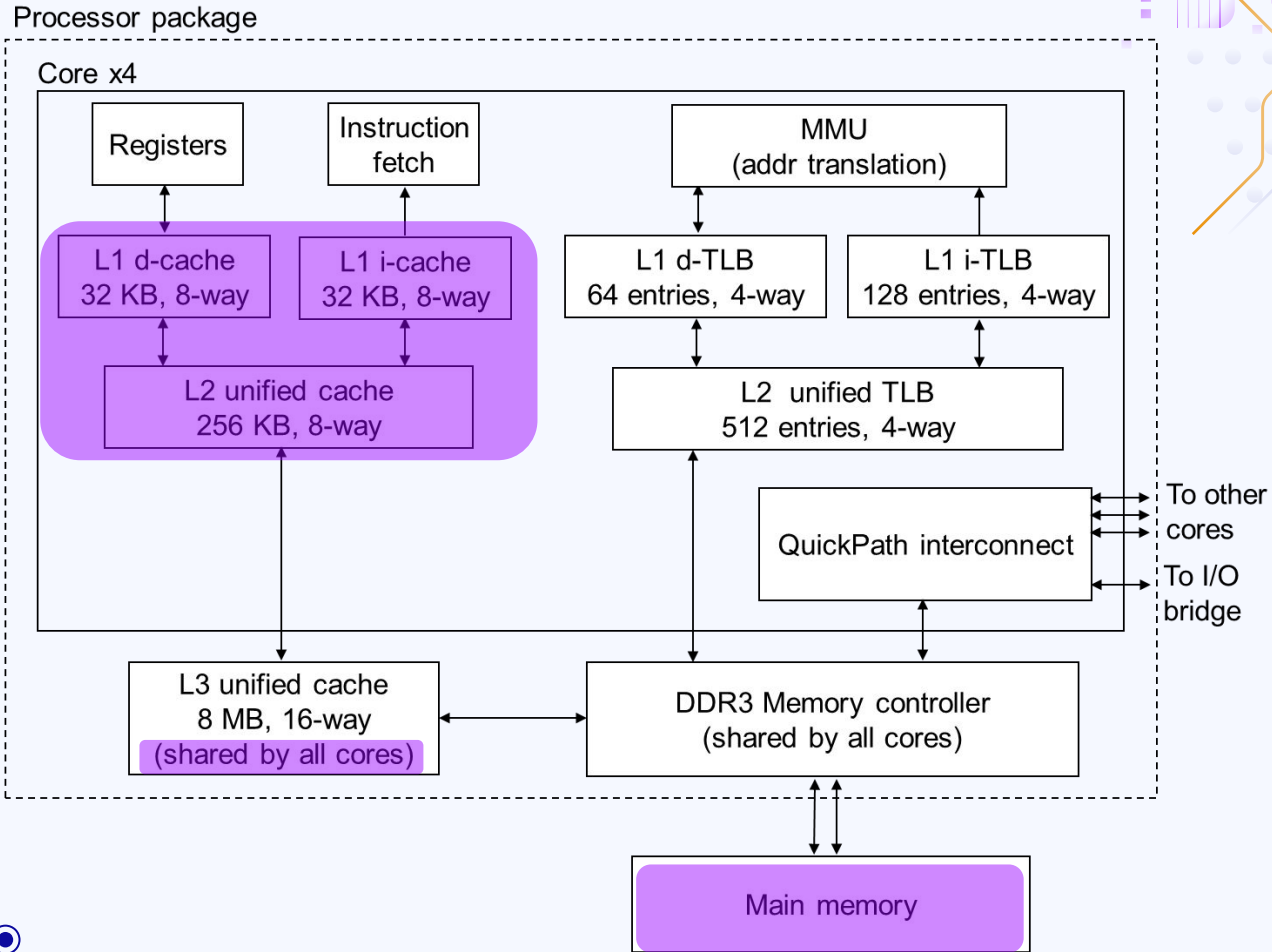


# **The perils of shared and unshared caches (coherence, side channels)**


# Intel i7

Source  
(Bryant & O'Hallaron)




# Cache coherence problem

Core A \$



Core B \$



Memory

$x = D$

# Definitions

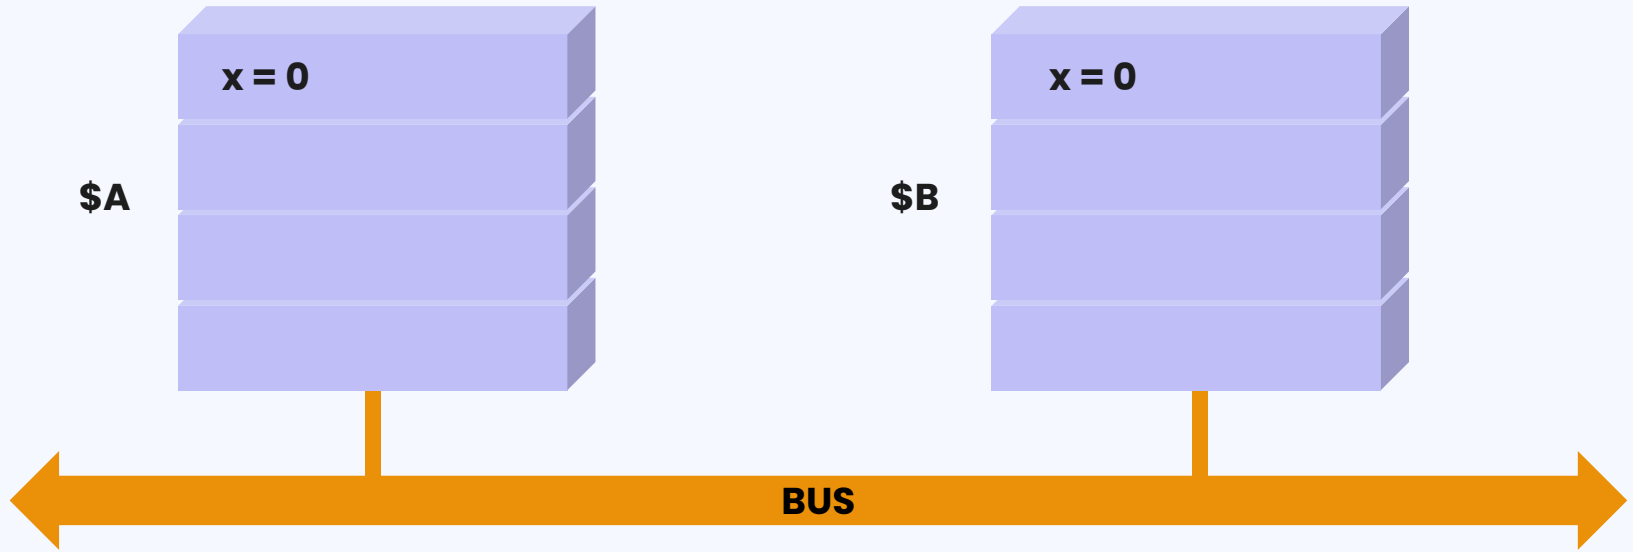
Intuitively: want any read of an item to return most recently written value to item

**Coherence** – what values can be returned by a read?

1. On a uniprocessor: reads after writes return written value
2. On a multiprocessor: reads by B after writes by A return written value when given sufficient time
3. Two writes to same location by one processor are seen in the same order by all processors


# Snooping

1. A reads x
2. B reads x



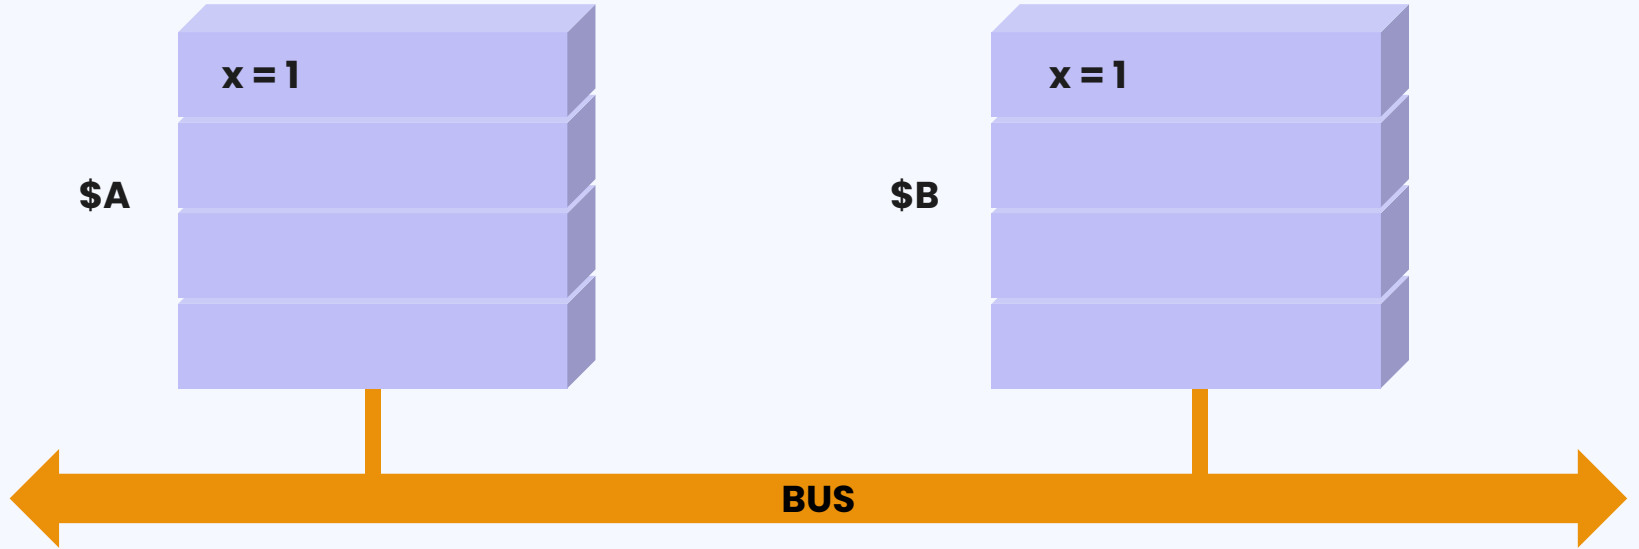
# Snooping

3. A writes x




4. B reads x

# Snooping



Cache needs to keep state for each block based on bus messages (common protocol: MSI)

# What should happen here?







What effects does block size have on  
coherence protocols?





(Pivoting away from coherence)  
How can processor A learn about what  
processor B is doing through the shared  
cache?

# Security and side channels

Security: protection against threats from malicious actor

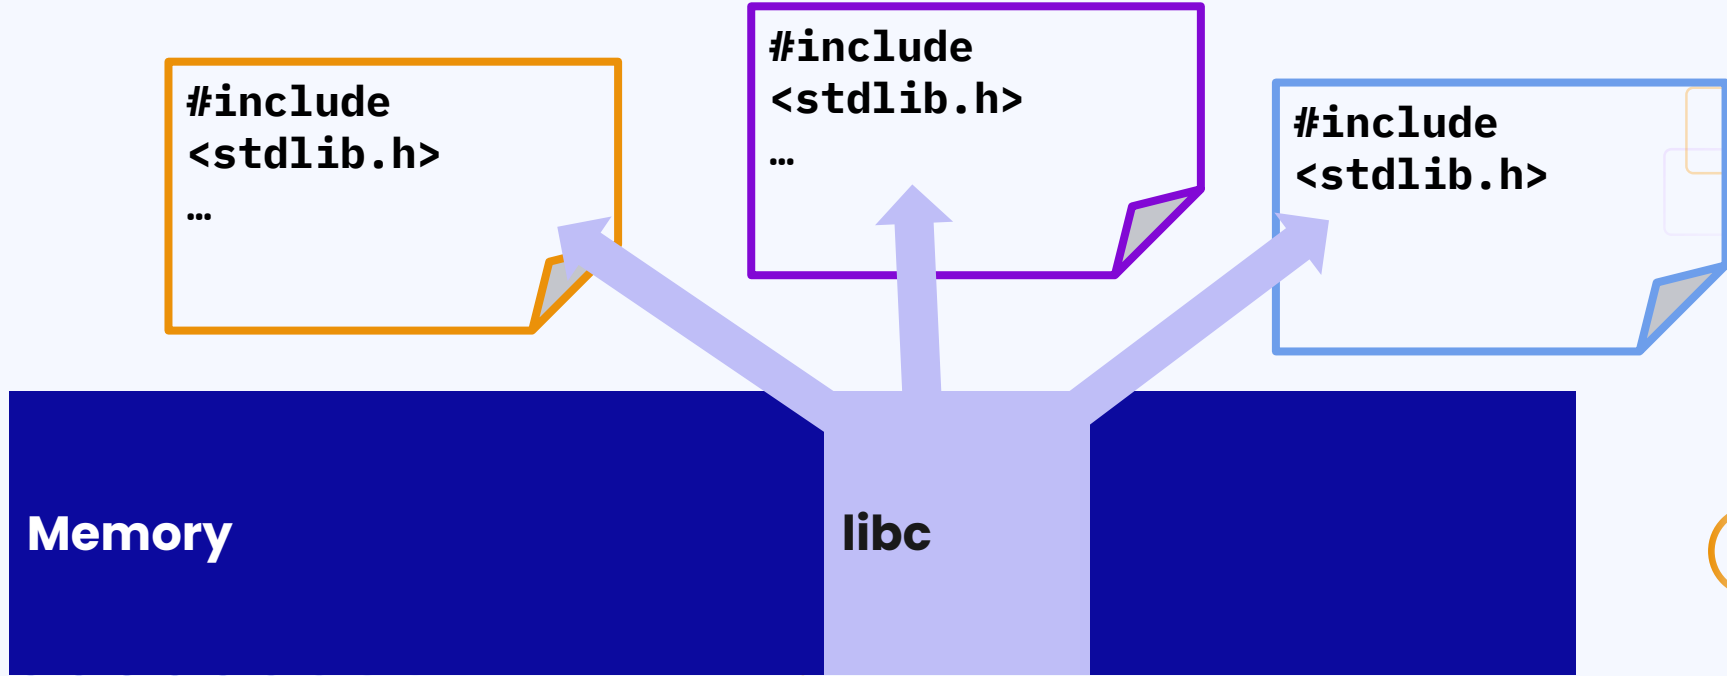
Obviously a **large** field, this is just a taste

Side channels: Incidental information leakage inferred from observing normal execution





# Shared cache side channels

What precisely can CPU B learn about CPU A?



# Flush & Reload attack



```
// flush the block  
cflush 0xLIBCADDR;
```

```
// wait some time  
t1 = time.now();  
while(time.now() - t1  
< 100ns);
```

```
// access block  
t2 = time.now();  
x = * 0xLIBCADDR;  
accessTime =  
time.now() - t2;
```

```
// if slow: unused  
// if fast: used!
```

Shared \$

# Flush & Reload in the wild

So what? An attacker knows I used libc...

```
1 function exponent( $b, e, m$ )
2 begin
3    $x \leftarrow 1$ 
4   for  $i \leftarrow |e| - 1$  downto 0 do
5      $x \leftarrow x^2$ 
6      $x \leftarrow x \bmod m$ 
7     if ( $e_i = 1$ ) then
8        $x \leftarrow xb$ 
9        $x \leftarrow x \bmod m$ 
10    endif
11  done
12  return  $x$ 
13 end
```

Only multiply when bit is 1

Yarom, Yuval, and Katrina Falkner. "{FLUSH+RELOAD}: A high resolution, low noise, l3 cache {Side-Channel} attack." 23rd USENIX security symposium (USENIX security 14). 2014. [link](#)




Figure 6: Time measurements of probes


Figure 2: Exponentiation by Square-and-Multiply



What does flush & reload depend on?



# Prime & Probe

Doesn't require precise timing or access to victim's memory; measures cache contention (less granular)




Fill up the cache!

♪ just executin' some code ♪



Read all addresses  
Which ones were misses?





# Caveats

It takes a *lot* to instrument a side-channel attack

Often can't learn *everything*, but narrow down search space

What can be done to guard against attacks?

- Oblivious RAM (largely theoretical)

- Trusted execution

- Cache partitioning

- Use an abacus