Lecture 16: Virtual Memory

- Today: DRAM innovations, virtual memory (Sections 5.3-5.4)

DRAM Technology Trends

- Improvements in technology (smaller devices) → DRAM capacities double every two years, but latency does not change much
- Power wall: 25-40% of datacenter power can be attributed to the DRAM system
- Will soon hit a density wall; may have to be replaced by other technologies (phase change memory, STT-RAM)
- Interconnects may have to be photonic to overcome the bandwidth limitation imposed by pins on the chip

Latency and Power Wall

- Latency and power can be both improved by employing smaller arrays; incurs a penalty in density and cost
- Latency and power can be both improved by increasing the row buffer hit rate; requires intelligent mapping of data to rows, clever scheduling of requests, etc.
- Power can be reduced by minimizing overfetch – either read fewer chips or read parts of a row; incur penalties in area or bandwidth

Density Wall

- New emerging non-volatile memories that have better scalability; instead of storing data in the form of charge, data encoded in cell resistance (phase change memory) or in electron spin (spin torque transfer STT-RAM)

- Phase change memory: the cell can be either amorphous (high resistance, represents zero) or crystalline (low resistance, represents one)

- Data is written by heating the material and cooling it at different rates (with electrical pulses); short intense pulse → amorphous; long medium pulse → crystalline

Phase Change Memory

- Can also have multi-level cells; each resistance value represents a different encoding; enables scalability
- Each cell can only be written about 10^8 times; need many tricks to improve endurance: write on change, shift bits within a row, re-map data to rows, etc.
- Reads are relatively quick (~50 ns), writes are very slow (~1000 ns)
- Has potential to replace DRAM, disk, or both in at least some classes of computers

Photonics

- A single waveguide carries light that was generated off-chip to multiple “nodes”
- The nodes can act as transmitters or receivers; transmitters can vary the amplitude of a light signal based on the input electrical signal
- Since multiple light wavelengths can be multiplexed on a waveguide and because each wavelength can carry a different signal, photonic interconnects have high bandwidth
- The E→O and O→E conversion overhead means that the photonic signal must travel far enough to out-do an electrical interconnect in terms of latency and power

Virtual Memory

- Processes deal with virtual memory – they have the illusion that a very large address space is available to them
- There is only a limited amount of physical memory that is shared by all processes – a process places part of its virtual memory in this physical memory and the rest is stored on disk
- Thanks to locality, disk access is likely to be uncommon
- The hardware ensures that one process cannot access the memory of a different process

Address Translation

- The virtual and physical memory are broken up into pages

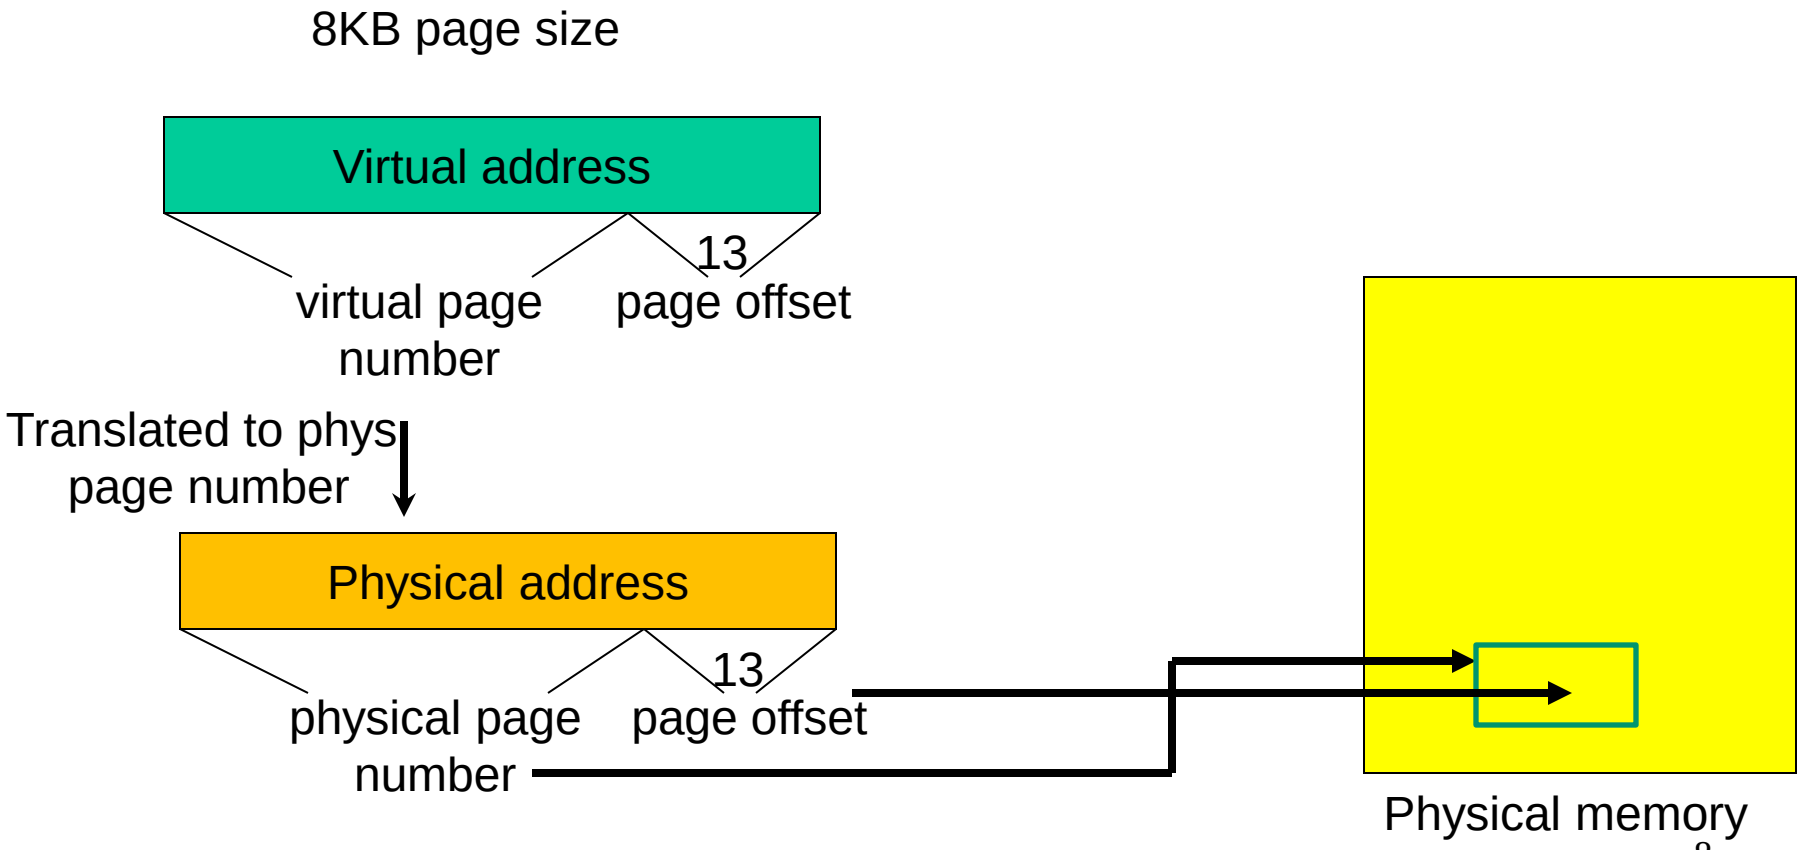

Memory Hierarchy Properties

- A virtual memory page can be placed anywhere in physical memory (fully-associative)
- Replacement is usually LRU (since the miss penalty is huge, we can invest some effort to minimize misses)
- A page table (indexed by virtual page number) is used for translating virtual to physical page number
- The memory-disk hierarchy can be either inclusive or exclusive and the write policy is writeback

TLB

- Since the number of pages is very high, the page table capacity is too large to fit on chip
- A translation lookaside buffer (TLB) caches the virtual to physical page number translation for recent accesses
- A TLB miss requires us to access the page table, which may not even be found in the cache – two expensive memory look-ups to access one word of data!
- A large page size can increase the coverage of the TLB and reduce the capacity of the page table, but also increases memory wastage

TLB and Cache

- Is the cache indexed with virtual or physical address?
  - To index with a physical address, we will have to first look up the TLB, then the cache → longer access time
  - Multiple virtual addresses can map to the same physical address – can we ensure that these different virtual addresses will map to the same location in cache? Else, there will be two different copies of the same physical memory word

- Does the tag array store virtual or physical addresses?
  - Since multiple virtual addresses can map to the same physical address, a virtual tag comparison can flag a miss even if the correct physical memory word is present

Virtually Indexed Caches

- 24-bit virtual address, 4KB page size → 12 bits offset and 12 bits virtual page number
- To handle the example below, the cache must be designed to use only 12 index bits – for example, make the 64KB cache 16-way
- Page coloring can ensure that some bits of virtual and physical address match

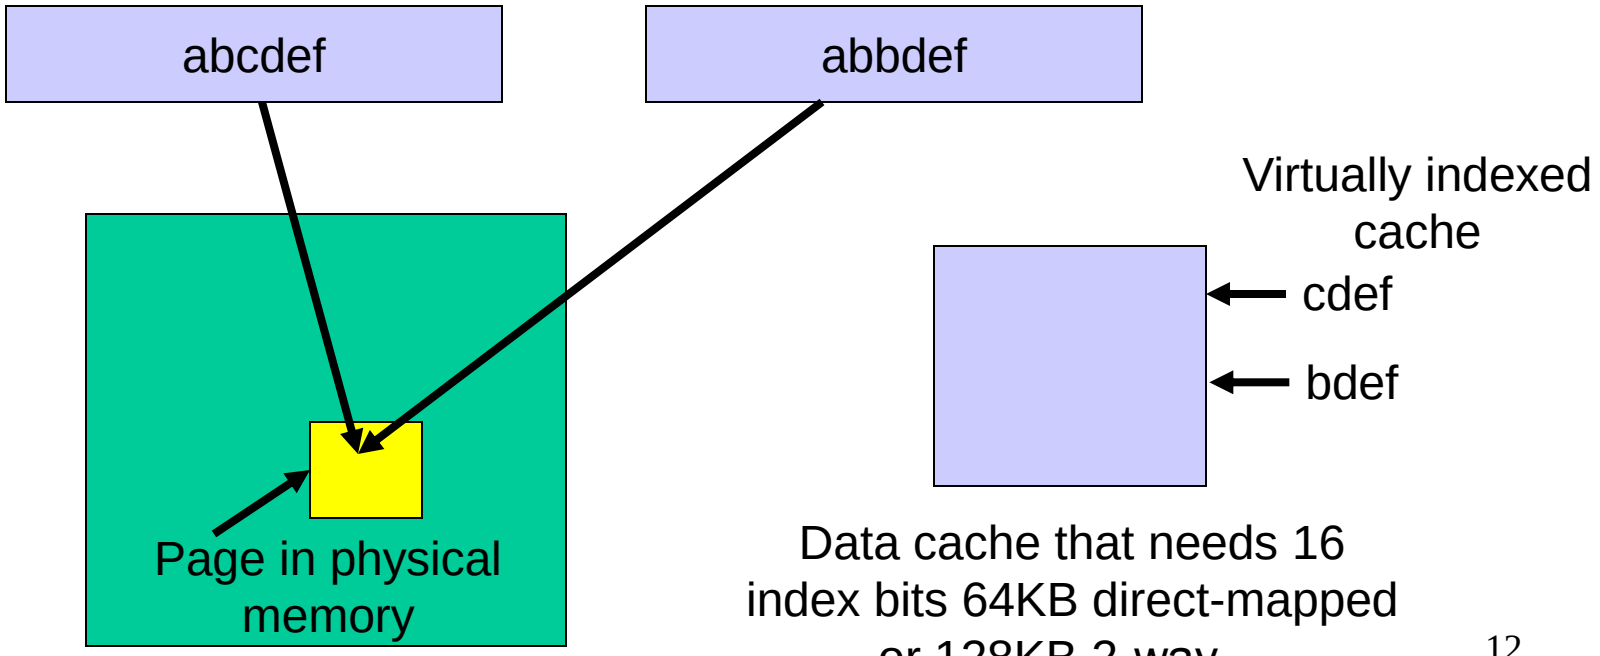

Cache and TLB Pipeline

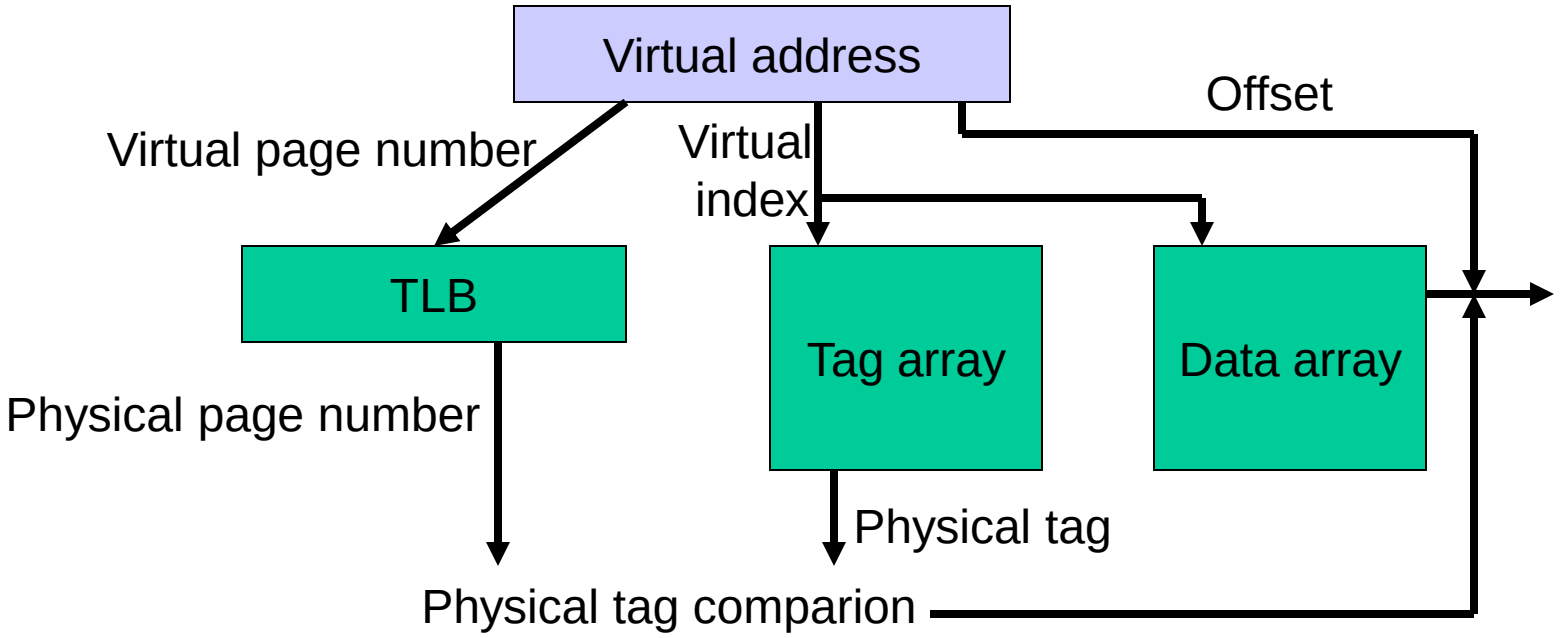

Virtually Indexed; Physically Tagged Cache

Title

- Bullet