

Amortized modification cost (cont.)

Thm: Any sequence of  $m$  INSERTS/DELETES on initially empty  $(a, 2a)$ -tree does  $O(m)$  modifications.

Pf: goal: find  $\Phi$  s.t.  $A_{\text{add}} = O(1)$ ,  $A_{\text{split}} \leq 0$  w.r.t.  $\Phi$ .

$\Phi := \sum_v f(|v|)$  where  $f: \{a-2, a-1, \dots, 2a-1, 2a\} \rightarrow \mathbb{N}$  properly chosen:  $(\sum R_i = \sum A_i - \Delta\Phi)$

$\uparrow$  #keys in  $v$        $\uparrow$  undersized       $\uparrow$  oversized

1)  $|f(i) - f(i+1)| \leq c$  for some const.  $c$

2)  $f(2a) \geq f(a) + f(a-1) + c + 1$  for  $A_{\text{split}} = O(1) + \Delta\Phi \leq 0$

$\uparrow$  oversized       $\uparrow$  new nodes       $\uparrow$  add to parent       $\uparrow$  rescaling to cover real cost

3)  $f(a-2) + f(a-1) \geq f(2a-2) + c + 1$  for  $A_{\text{merge}} = O(1) + \Delta\Phi \leq 0$

$\uparrow$  undersized       $\uparrow$  sibling       $\uparrow$  new node       $\uparrow$  remove from parent       $\uparrow$  rescaling

For  $f$  we can take e.g.

$k$	$a-2$	$a-1$	$\dots$	$2a-1$	$2a$	$\checkmark$
$f(k)$	$2$	$1$	$0 \dots 0$	$2$	$4$	

$\Rightarrow$  INSERT/DELETE  $O(1)$  amortized modification cost

$\Rightarrow$  total real cost  $\sum R_i = \sum A_i - \underbrace{\Phi_m}_{\geq 0} + \underbrace{\Phi_0}_0 = O(m)$ .  $\square$

Top-down  $(a, b)$ -trees

INSERT: if the current node is full ( $b-1$  keys), split it (median to parent)

$\Rightarrow$  parent always has a free space  $\Downarrow$

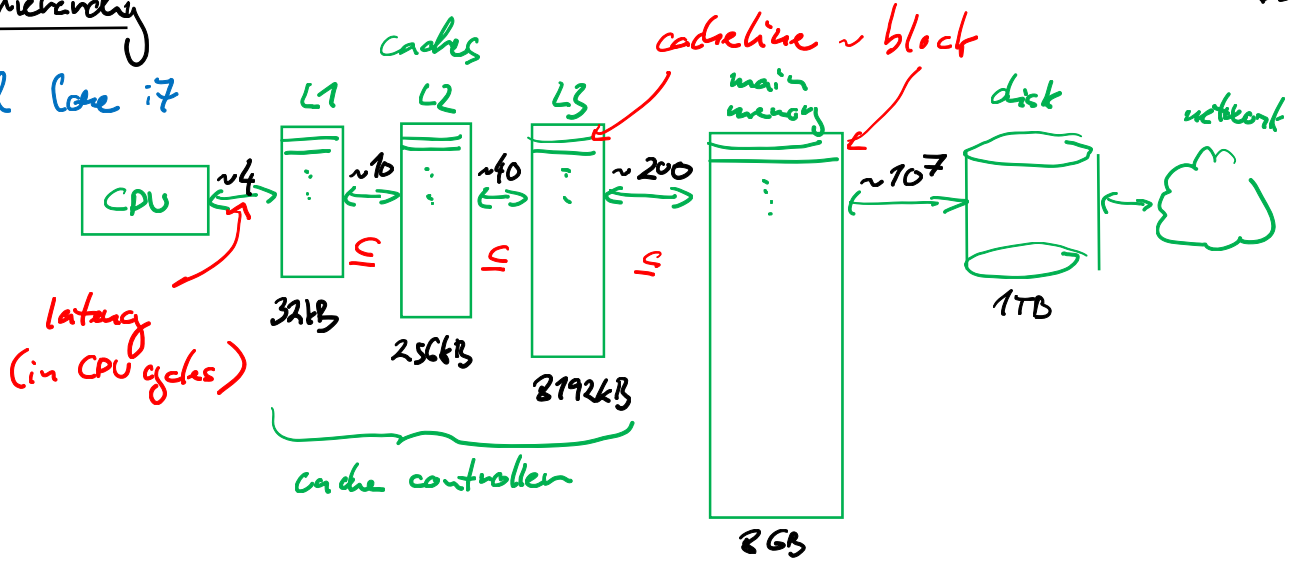
(preemptive splitting) needs  $\lfloor \frac{b-2}{2} \rfloor \geq a-1$

DELETE: similarly, preemptive merging/moving (exercise)  $b \geq 2a$

$\Rightarrow$  useful for parallel  $(a, b)$ -trees: simple locking mechanism (to be discussed later)

# Memory hierarchy

Ex. Intel Core i7




cache line (block) of size  $B \sim 32/64 \text{ B}$  (unit of transfer)


cache hit/miss  $\Rightarrow$  eviction (replacement strategy)

## cache mapping


1) fully associative: every block can be in any cache line




$\Rightarrow$  need many tag comparisons




2) direct mapping: every block can be in only one cache line




$\Rightarrow$  no need for a replacement strategy



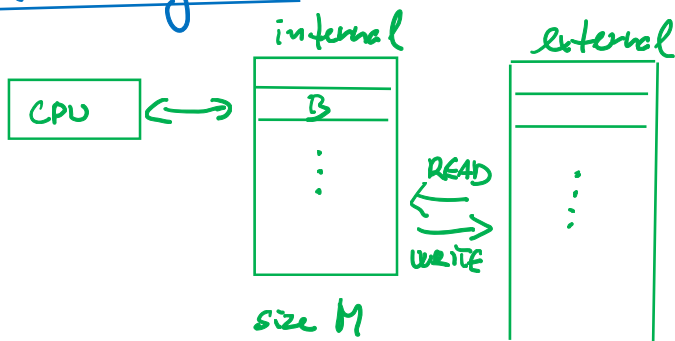
3) set associative: every block can be in a set of  $2^j$  cache lines



(compromise between 1) and 2))



1) external memory model (I/O model)

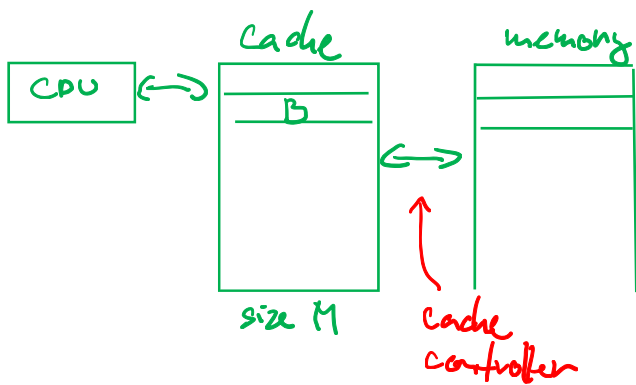


block size  $B$  (# items / Byte / ...)  
 internal memory of size  $M$   
 (alg.) ( $M/B$  blocks)  
 CPU controls READs / WRITEs to the external memory

Def: An algorithm has I/O complexity  $f(n, B, M)$  if for any parameters  $B, M$  any input of size  $n$  the number of I/O operations is  $\leq f(n, B, M)$ .  
 ↑  
 read / write

Note: Typically, # writes =  $O(\#reads)$  (if  $|output| = O(|input|)$ ).  
 $\Rightarrow$  analysis reduces only to # reads.

2) cache-aware model




Alg. **knows** the parameters  **$B, M$** .  
 Alg. has no control of block transfers (it is handled by the controller).

3) cache-oblivious model : Alg. does not know  $B, M$ . (Maybe assumes  $M \in \Omega(B^2)$ ).  
 $\Rightarrow$  needs a strategy that works for any cache  
 $\Rightarrow$  more complex (overhead) but more general (good for cache hierarchies)

Assumptions for analysis:

- 1) cache controller has optimal replacement strategy: (offline, knows future) accesses to memory  
 "evict the cache line whose block is needed furthest"
- 2) fully associative cache


Scan of an array (example)



1) I/O model: aligned array start  
 1 internal block suffices  
 $\Rightarrow$  I/O complexity =  $\lceil N/B \rceil$


2) cache-aware: also aligned (since we know B)  
 I/O complexity = # cache misses in  $\leq$  # cache misses in =  $\lceil N/B \rceil$   
 (reads) OPT strategy "use 1 cache line" strategy  
 maybe  $M/B$  blocks already cached

3) cache oblivious: no alignment guaranteed  $\Rightarrow$  I/O complexity  $\leq \lceil N/B \rceil + 1$   
 (extra read)




$\Rightarrow O(N/B + 1)$  in all models  
 cannot be omitted!

Merge sort (for sorting in external memory)



3x sequential scans  $\Rightarrow$   
 3 blocks / cache lines suffice

$\Rightarrow$  single merge in  $O(u_1 + u_2)$  time  
 I/O complexity  $O((u_1 + u_2)/B + 1)$   
 (in all models)



single pass in  $O(N)$  time  
 $O(N/B + 1)$  block transfers  
 runs are consecutive

$\Rightarrow$  Merge sort in  $\Theta(N \log N)$  time and I/O complexity  
 $O(N/B \log N + \log N)$

1  $\leftarrow$  improved since for  $N \geq B$  we have  $O(N/B + 1)$   
 and for  $N < B$  entire array fits in 1 block  $\Rightarrow O(1)$  block transfers