



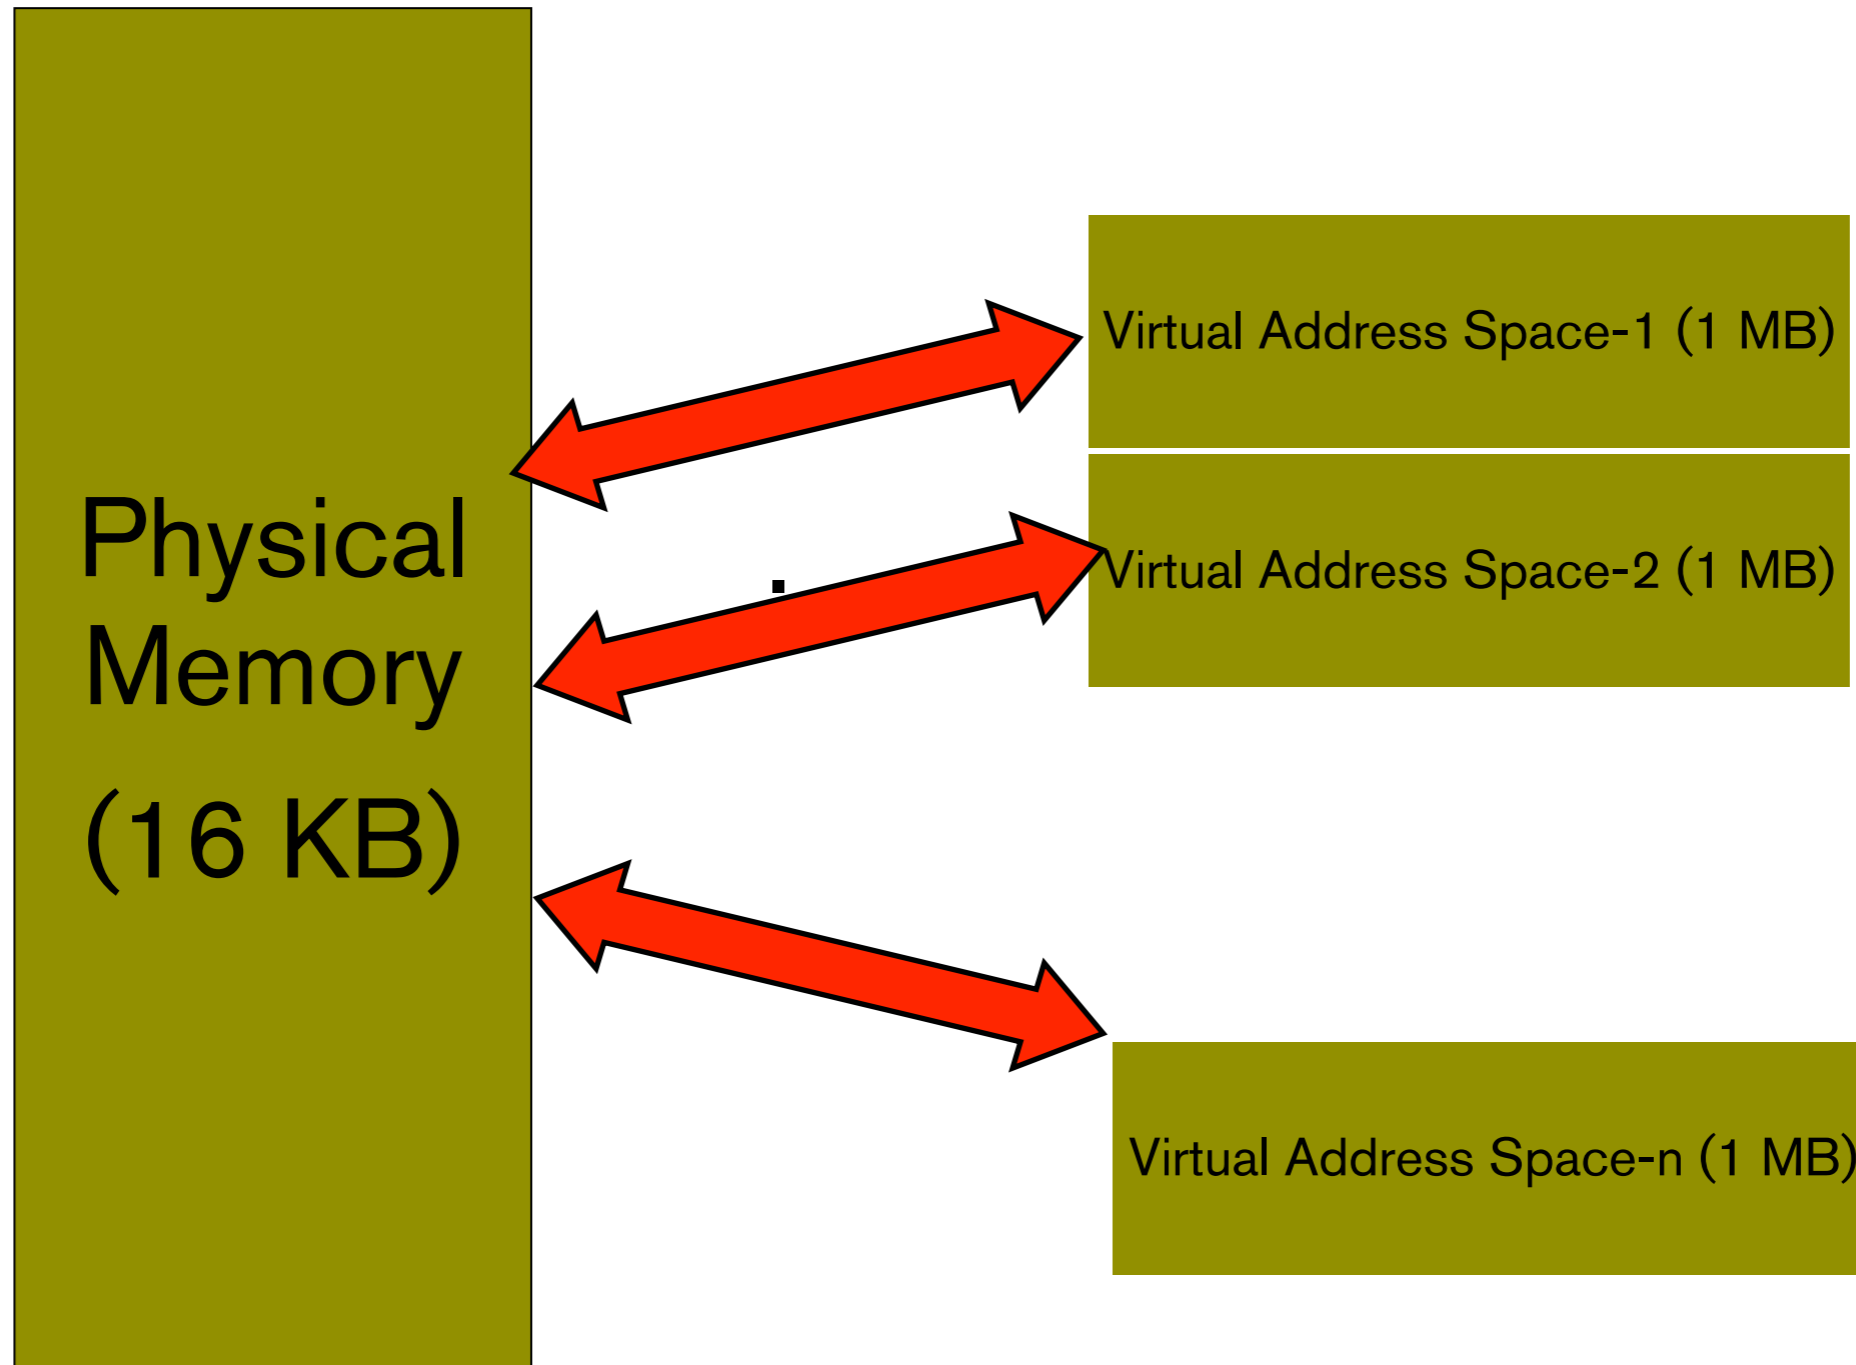
UNIVERSITY OF OREGON

CIS 415:
Operating Systems
Virtual Memory

Spring 2012
Prof. Kevin Butler

- What if programs require more memory than available physical memory?
 - ▶ Use overlays
 - Difficult to program though!
 - ▶ Virtual Memory.
 - Supports programs that are larger than available physical memory.
 - Allows several programs to reside in physical memory (or at-least the relevant portions of them).
 - Allows non-contiguous allocation without making programming difficult.

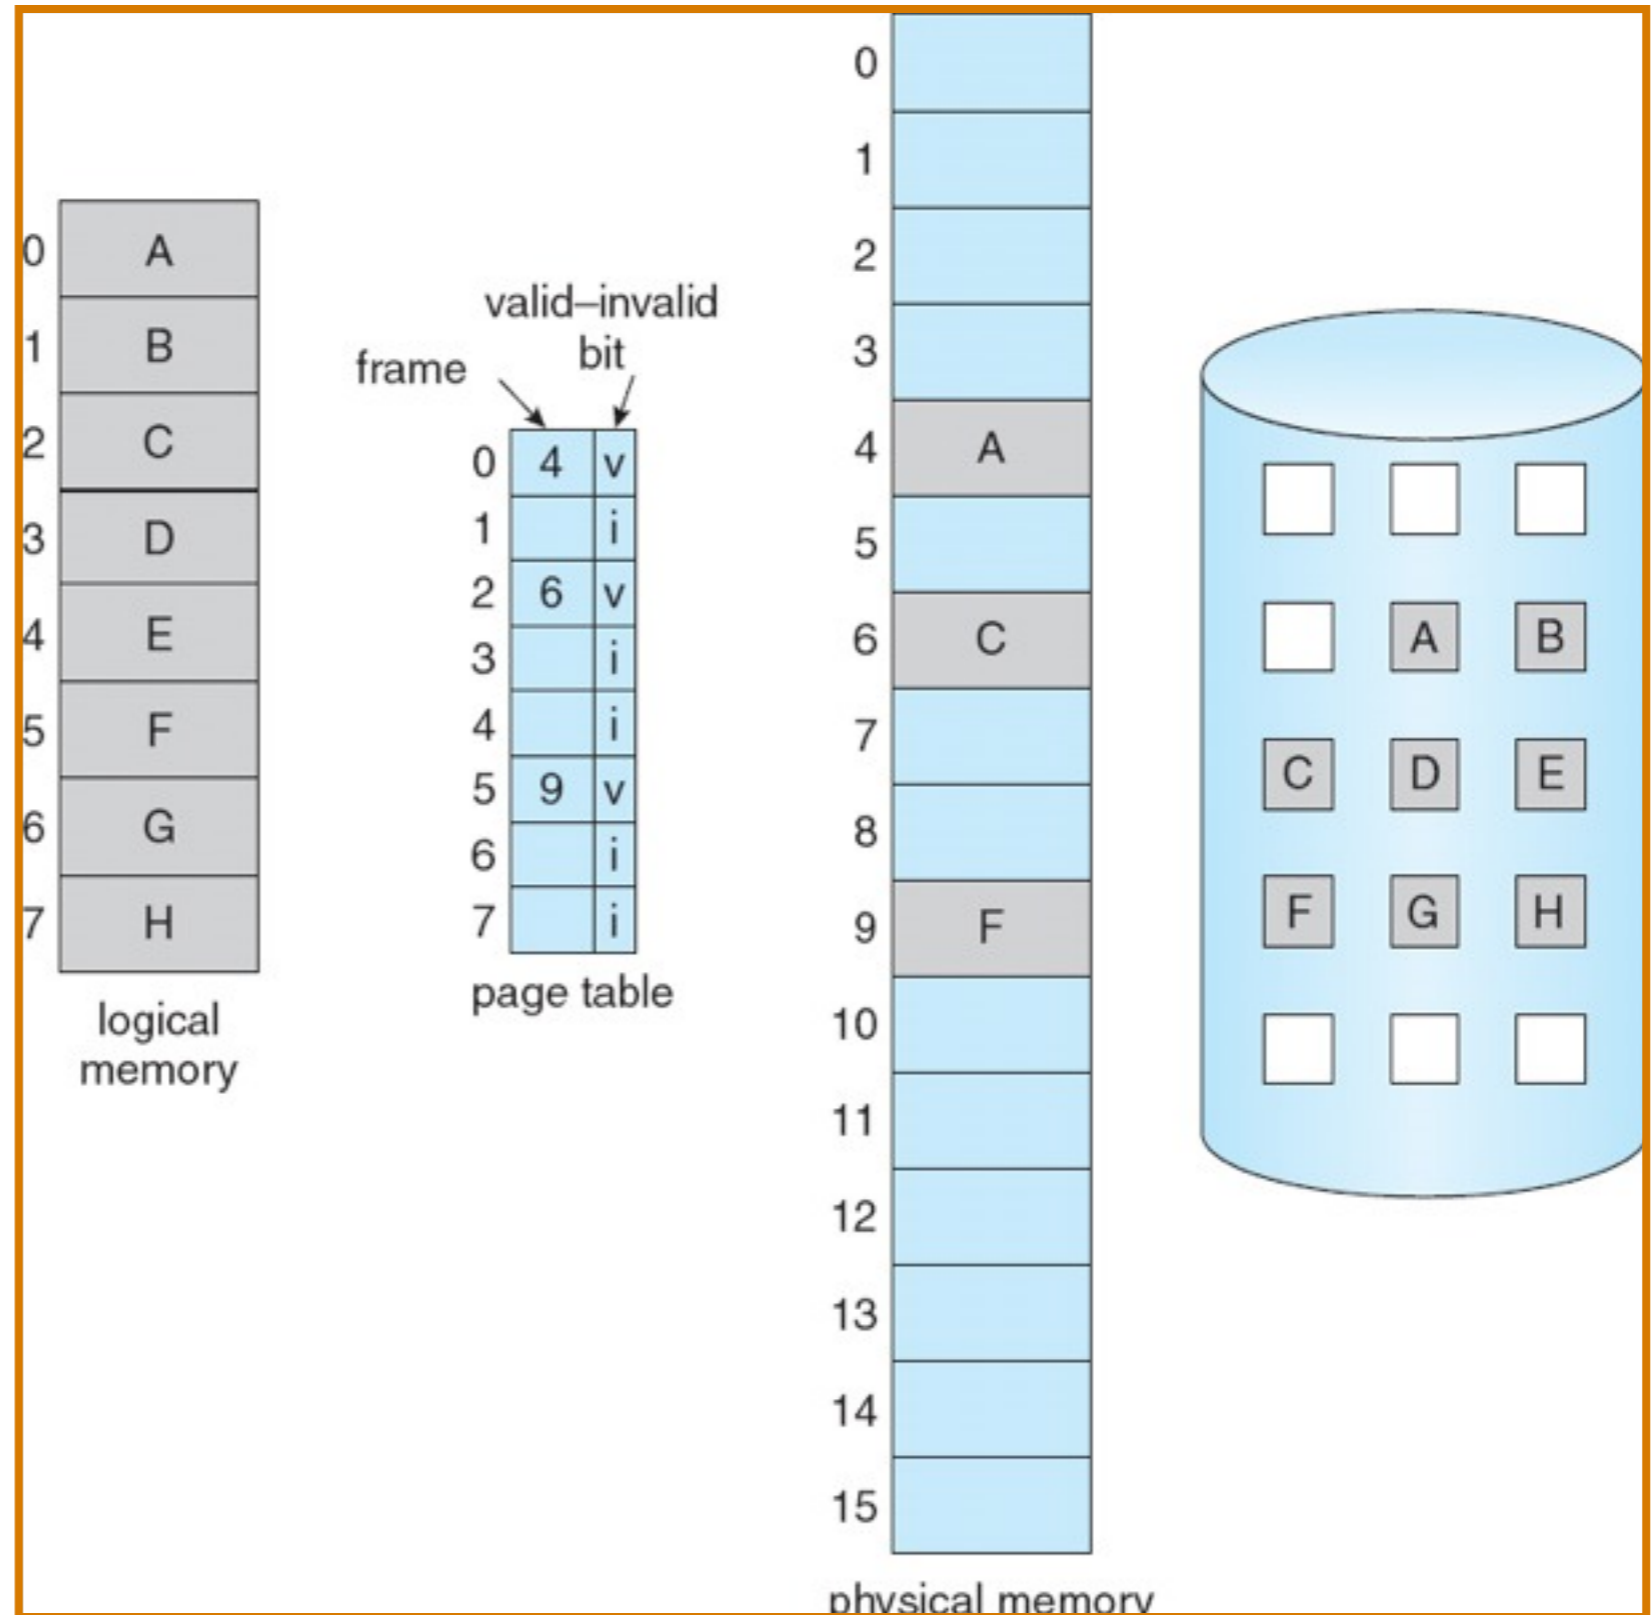
Example



- If a page table mapping indicates an absence of the page in physical memory, hardware raises a *Page Fault*.
- OS traps this fault and the interrupt handler services the fault by initiating a disk read request.
- Once page is brought in from disk to main memory, page table entry is updated and the process which faulted is restarted.
 - ▶ May involve replacing another page and invalidating the corresponding page table entry.

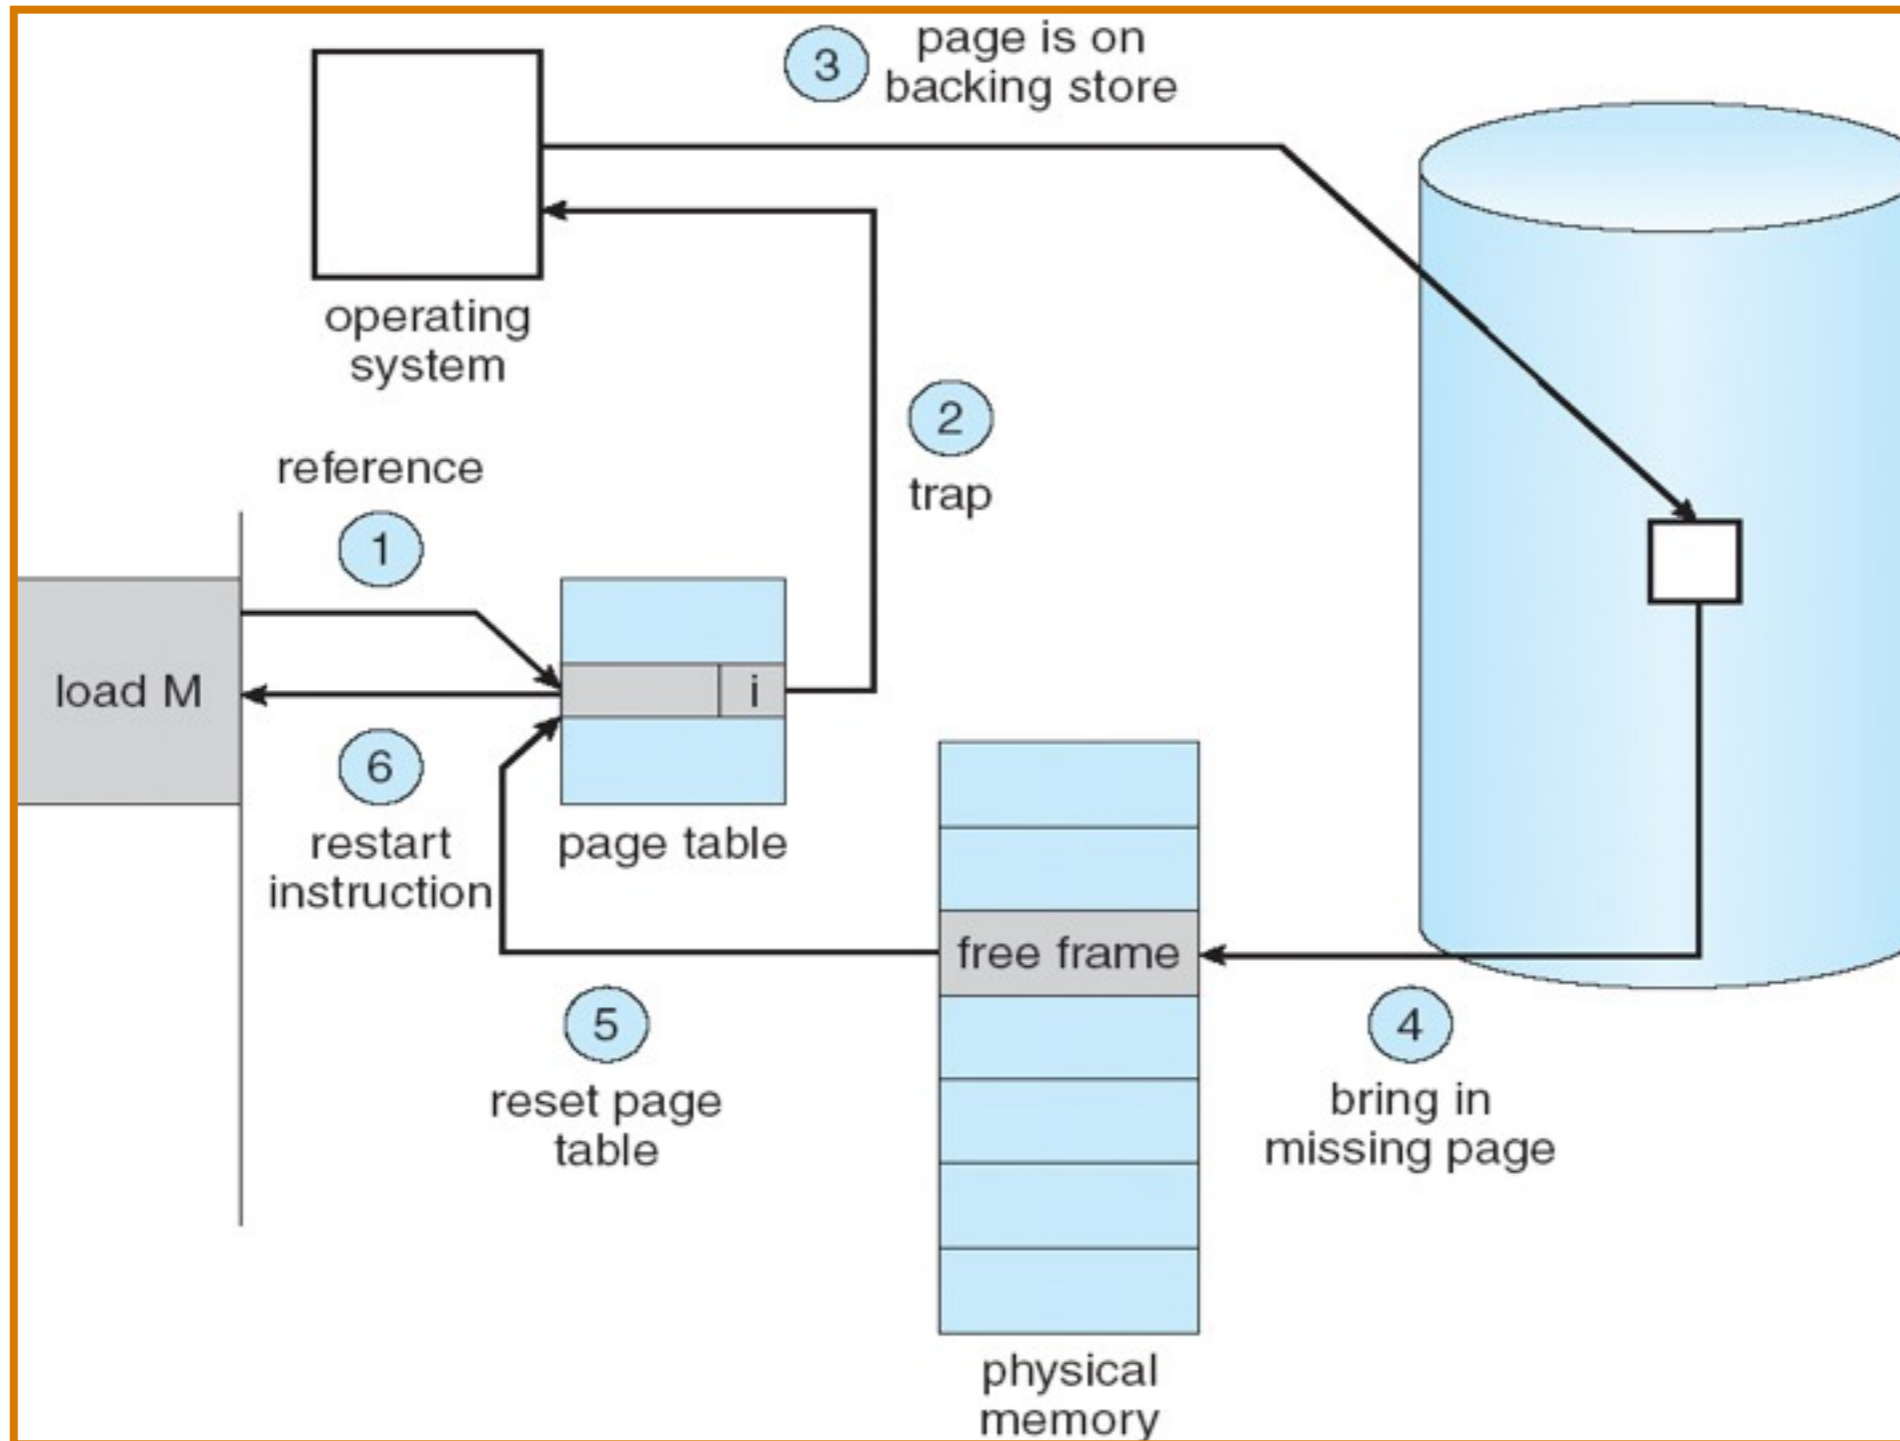
Page Table

Some pages not in main memory



- If there is a reference to a page, first reference to that page will trap to operating system:
 - ▶ page fault
- Operating system looks at another table to decide:
 - ▶ Invalid reference -- abort
 - ▶ Just not in memory
- Get empty frame
- Swap page into frame
- Reset tables
- Set validation bit = v
- Restart the instruction that caused the page fault

Handling a Page Fault



Demand Paging Performance

- Page Fault Rate
 - ▶ $0 \leq p \leq 1.0$
 - ▶ if $p = 0$, no page faults
 - ▶ if $p = 1$, every reference is a fault
- Effective Access Time (EAT)
- $$\text{EAT} = (1 - p) \times \text{memory access} \\ + p \times (\text{page fault overhead} \\ + \text{swap page out} \\ + \text{swap page in} \\ + \text{restart overhead})$$

Demand Paging Example

- Memory access time = 200 nanoseconds
- Average page-fault service time = 8 milliseconds
- $EAT = (1 - p) \times 200 + p (8 \text{ milliseconds})$
 - ▶ $= (1 - p) \times 200 + p \times 8,000,000$
 - ▶ $= 200 + p \times 7,999,800$
- If one access out of 1,000 causes a page fault, then
 - $EAT = 8.2 \text{ microseconds.}$
 - This is a slowdown by a factor of 40!!

- When bringing in a page, something has to be evicted.
- What should we evict? – page replacement algorithm.

- Why optimal?
 - ▶ No other algorithm can have # of page faults lower than this, for a given page reference stream.
- Algorithm:
 - ▶ At any point, amongst the given pages in memory, evict the one whose next reference from now is the furthest in the future.

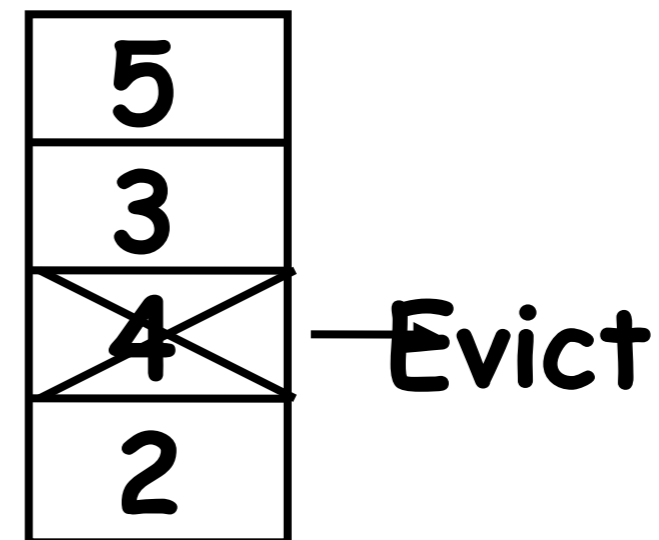
An example of OPT

Reference String

..... 7, 5, 3, 3, 5, 2, **4**, 4, 3, 2,

At this point,
what do we
replace?

Current Physical Memory



Problem with OPT



- Not implementable!
- Requires us to know the future.
- But it has the best page fault behavior
- How do we approach OPT?

First-in, First-out (FIFO)



- Maintain a linked list of pages in the order they were brought into PM.
- On a page fault, evict the one at the head.
- Put the newly brought in page (from disk) at tail of this list.
- Problems:
 - ▶ Reference String: 1,2,3,4,1,1,5,1,1,...
 - ▶ Page fault at (5) would replace (1) !
 - ▶ Need to know what is in recent use!

Not Recently Used



- Referenced bit set on each Read/write by h/w
- Modified set on each write by h/w
- On startup set both R and M bits to 0.
- Periodically (using clock interrupts) the R bit is cleared.

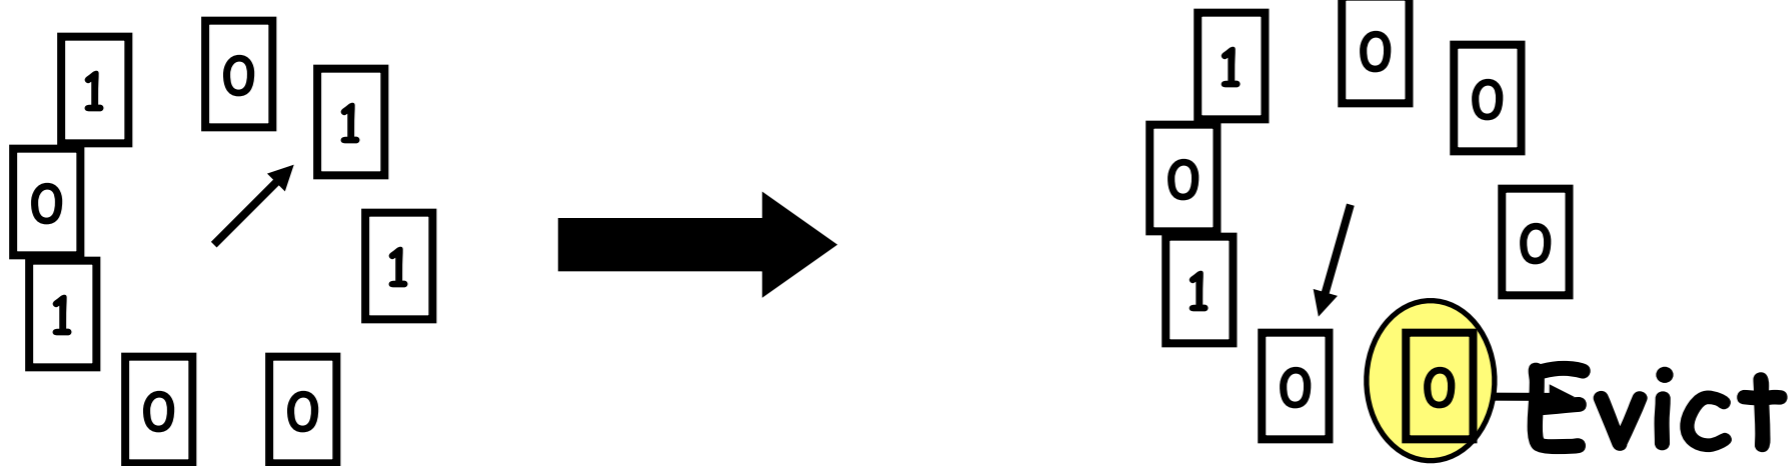
- On a page fault, examine the state of a page
 - ▶ Class 0: R = M = 0
 - ▶ Class 1: R = 0, M = 1
 - ▶ Class 2: R = 1 M = 0
 - ▶ Class 3: R = 1 M = 1

- NRU replaces a page chosen at random from the lowest numbered nonempty class.

Second Chance Replacement



- Aka. clock algorithm (why?)
- Same as FIFO, except you skip over the pages whose reference bit is set, resetting this bit, and moving those pages to end of list.
- Implementation:



Least Recently Used (LRU)



- Order the list of physical memory pages in decreasing order of recency of usage.
- Replace the page at the tail.
- Problem:
 - ▶ This list will need to be updated on each memory reference.
 - ▶ Asking the h/w to do this is ridiculous!
- Solution: Approximate LRU

Approximate LRU (counters)



- Keep a counter for each physical page.
- Initially set to 0.
- At the end of each time interval (interval to be determined), shift the bits right by one position.
- Copy the reference bit to the MSB of counter and reset reference.
- For a page replacement, pick the one with the lowest counter value.
- It is an approximation of LRU because:
 - ▶ we do not differentiate between references that occurred in the same tick.
 - ▶ the history is limited by the size of the counter.

- OPT, FIFO, NRU, second-chance/clock, LRU, approximate LRU
- In practice, OSes use second chance/clock or some variations of it.

Belady's Anomaly



- Normally you expect number of page faults to decrease as you increase physical memory size.
- However, this may not be the case in certain replacement algorithms

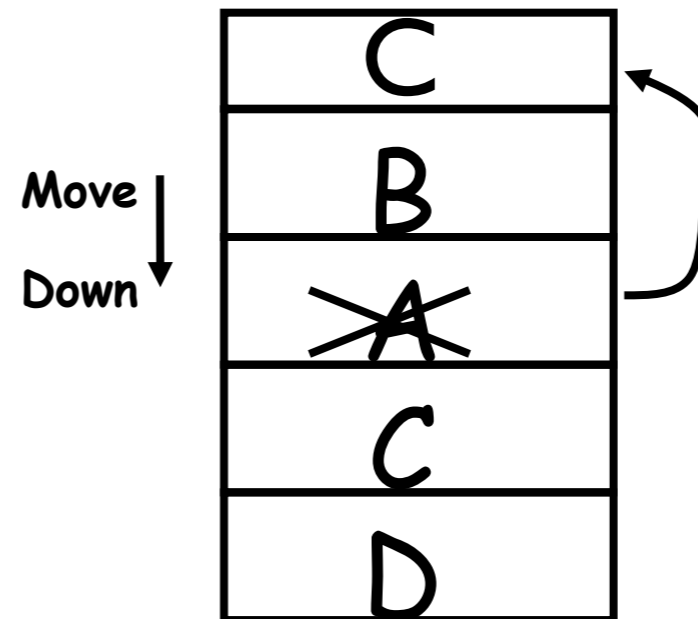
- FIFO replacement Algorithm
- Reference string:
 - ▶ 1 2 3 4 1 2 5 1 2 3 4 5
- 3 physical frames
 - ▶ F F F F F F F - - F F -
 - ▶ # of faults = 9
- 4 physical frames
 - ▶ F F F F - - F F F F F F
 - ▶ # of faults = 10

- Algorithms which do NOT suffer from Belady's anomaly are called *stack algorithms*
- E.g. OPT, LRU.

- Paging behavior characterized by
 - ▶ Reference string
 - ▶ Physical memory size
 - ▶ Replacement algorithm

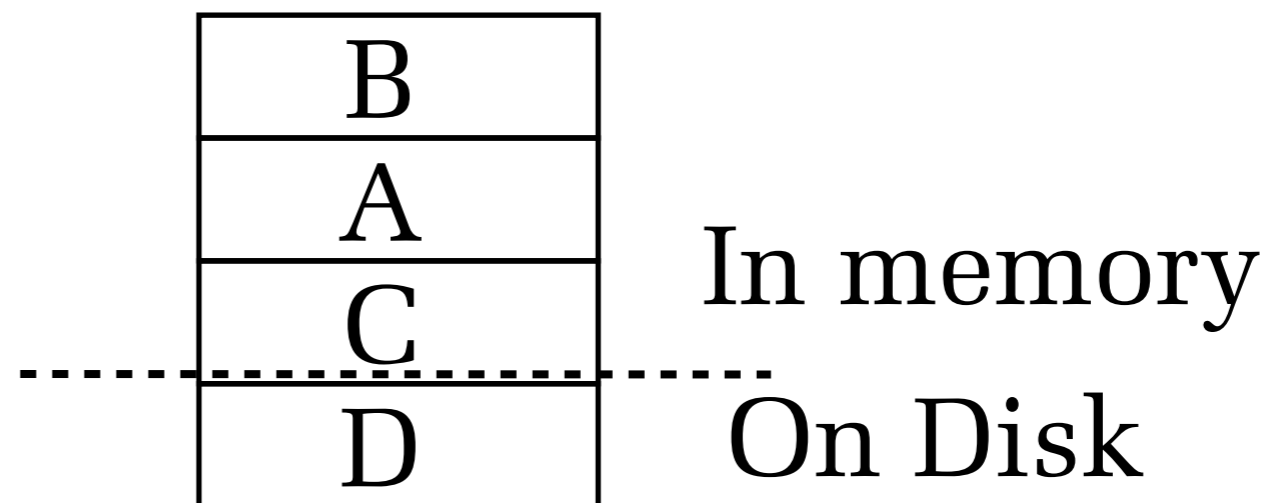
**e.g. A, B, C and D
are virtual pages**

- Visualize it as a stack (say *M*), where a page that is “referenced” is brought to the top of the stack from wherever it is.



When C is referenced ...

- Whatever is in recent use is on the top of M , and the ones that are not in recent use are at the bottom.
- In fact, the top P entries of M represent the pages in physical memory, where P is the # of physical frames.

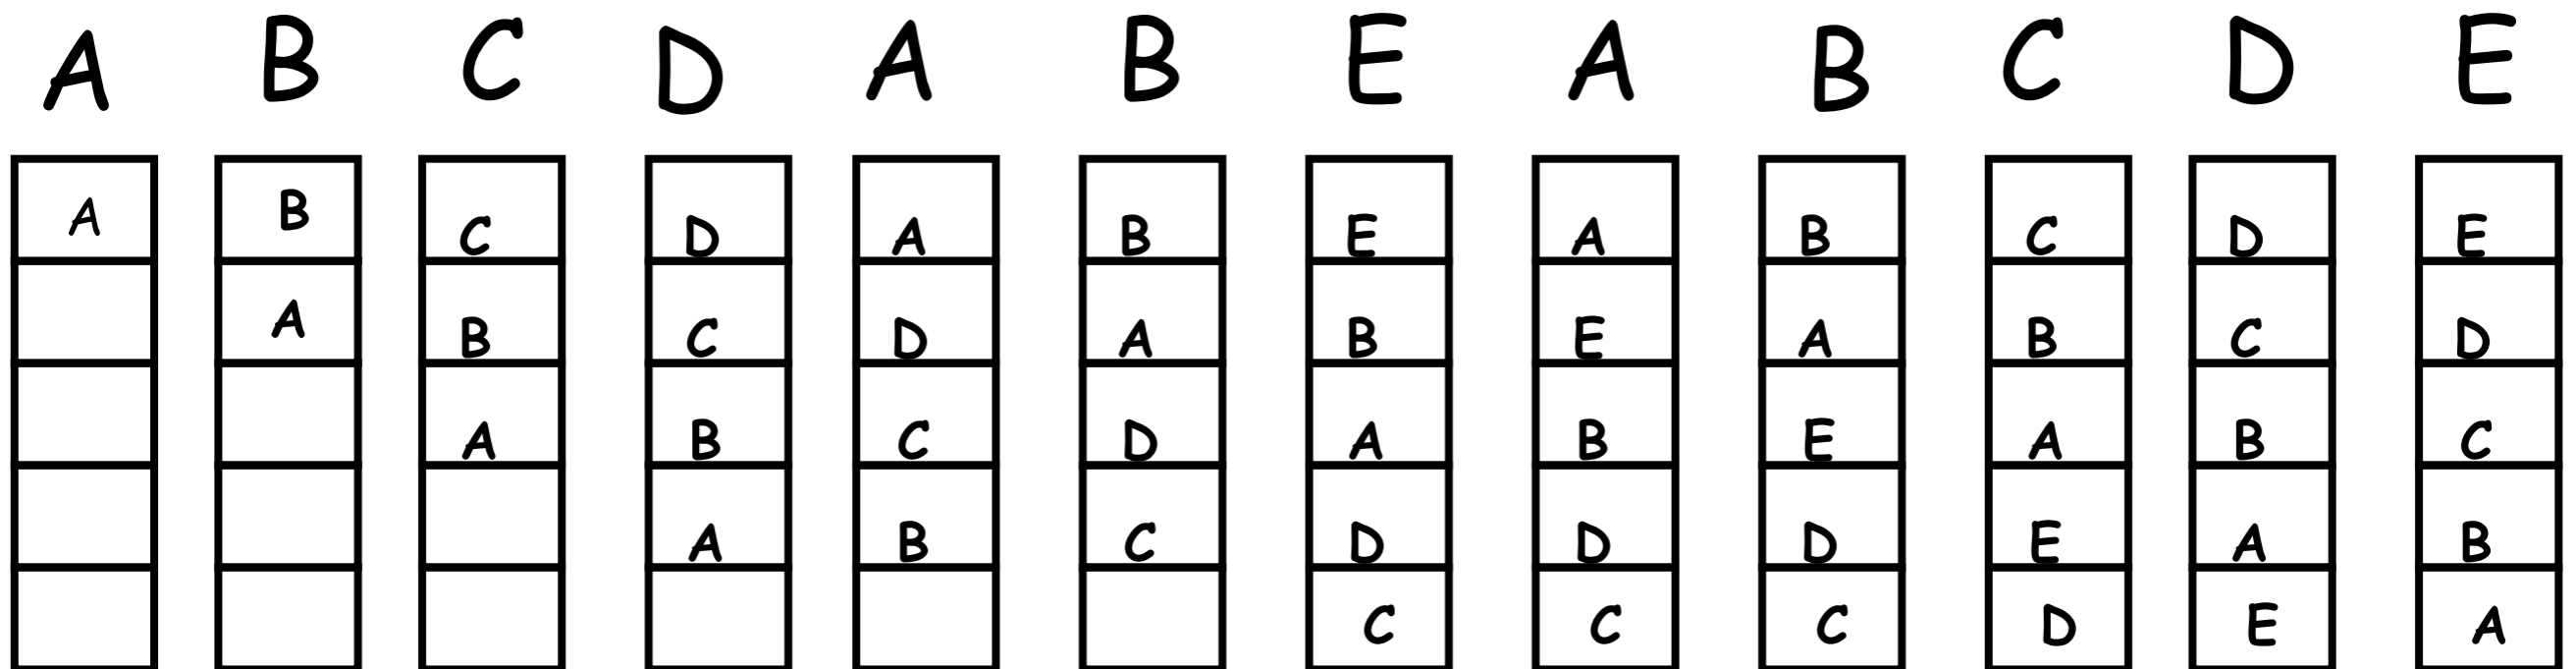


- Distance String:
 - ▶ For each element r of reference string R , this represents the distance of that element from the top of stack in M .

An example of how M changes with 5 virtual pages



Reference String



∞ ∞ ∞ ∞ 3 3 ∞ 2 2 4 4 4

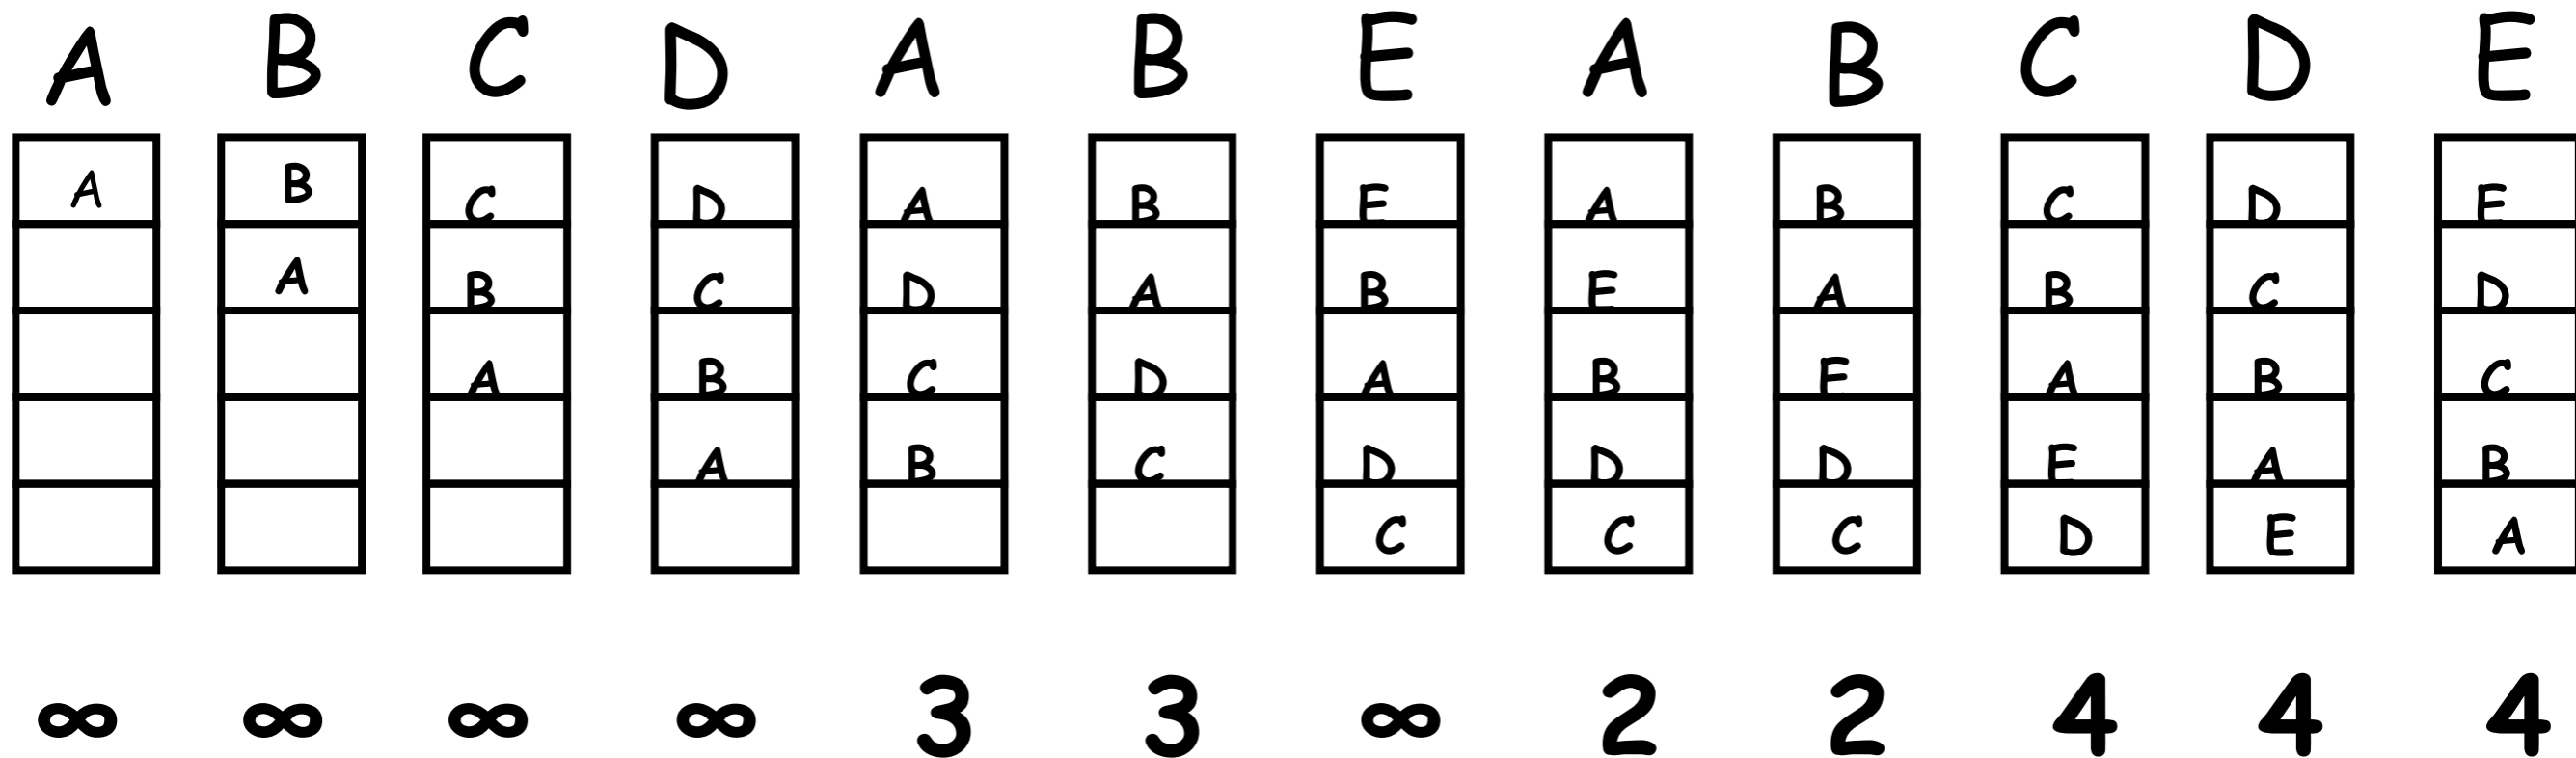
Distance String

Define vector C



- $C[i]$ represents the number of times “ i ” appears in the distance string.

Reference String



Distance String

**C vector: C[0]=0, C[1]=0, C[2]=2,
C[3]=2, C[4]=3, C[5] ... =0
C[∞]=5**

Define Vector F



- $F[j]$ is the number of page faults that will occur for the given reference string R with “ j ” physical frames.
- $F[j] = C[j] + C[j+1] + C[j+2] + C[j+3] \dots + C[\infty]$

- It is now straightforward to prove LRU does not suffer from Belady's anomaly.
 - ▶ The M vector tracks what is in physical memory in the top P slots for LRU.
 - ▶ Note that vector $C[i]$ is independent of physical memory size.
 - ▶ When you go from physical memory with j frames to $(j+x)$ frames, note that the number of C vector terms in the RHS of equation for F decreases \Rightarrow Page faults can only decrease if at all!

$$M(j, R) \subseteq M(j + x, R)$$

- Keep the essentials of what you currently need (*working set*) in physical memory.
- When something you need is not in memory, bring it in from disk:
 - ▶ On demand (demand-paging)
 - ▶ Ahead of need (pre-paging)
- Programs need to exhibit good locality to avoid “thrashing” of pages in memory.
- This usually requires good programming skills!

Fragmentation in paging



- Note that there is only internal fragmentation, and that too only in the last allocated page.
- Smaller the page, smaller the internal fragmentation.
- However, this reduces spatial locality.

Page size trade-offs

- Average process size = s bytes
- Page size = p bytes
- Page Table entry = e bytes

$$\text{Overhead} = s * \frac{e}{p} + \frac{p}{2}$$

- **Page Replacement**
 - ▶ Virtual memory
 - ▶ Page faults
 - ▶ Optimal page replacement not achievable
 - ▶ Variety of algorithms
 - ▶ Anomalies

- Next time: Virtual memory issues