

③ Performance

- caching } today
- concurrency } today
- batching } web.
- speculative I/O
- ~~also~~ scheduling
- (replication/partitioning)

→ flash crowds.

④ Cache (6.004)

$$A = t_1 \times A_R + (1-t_1) D_D$$

$0.99 \times 10 \text{ nsec} + 0.01 \times 100 \mu\text{sec}$   
 $100 \mu\text{sec} \Rightarrow 10000/s$   
 $0.999 \rightarrow 19.9 \mu\text{sec}$

locality of reference

→ Web not exactly like computer (pages) but zip code but not.

How good are we doing?  
 block request on miss!  
 → ~~one~~ multiple threads (show code)

⑥

T<sub>1</sub>: A, B, C, D 6.004  
 T<sub>2</sub>: A, B, C, D 6.033

① Ex

(show slide 1)

⑤ Replacement

LRU → 2 3 4 MRU

1	4
2	3
3	2

1	1
2	2
3	4

used by flash ⇒ 1/verdict

① what is the best page to throw out?

HW: replacement strategy dictated by code  
 structure  
 direct-mapped  
 N-way associative

SW: more flexibly

⑦ Race condition

A<sub>1</sub>, A<sub>2</sub> B<sub>2</sub> B<sub>1</sub>  
 C<sub>1</sub>, C<sub>2</sub> D<sub>2</sub> D<sub>1</sub>

difficult to debug  
 the rac-25  
 mod → passed file  
 preemptive scheduling  
 in our thread package

Simple:

latency: disk limited  
 10 msec

throughput: 100/s  
 ⇒ 8M/day

(not all pages are likely to be accessed)

↓ slide 2.

⑧ acquire(lock);  
 while (lock == 1);  
 lock ← 1

{  
 1: taken  
 0: free

⑩ TSL mem, R  
 R ← [mem]  
 [mem] ← 1  
 (bus arbiter)  
 (show code)

small atomic → bigatomic

⑧ solution, lock obj

acquire(lock)  
 ↓ atomic  
 release(lock)

whenever thread write-share a variable

programmer convention →  
 split modularity  
 show code  
 (still wrong)