











5/31/2018

LATENCY TIMES						
 Design considera SSDs 4x the time HDDs 80x the time 	ations of DRAM e of DRAM					
Action	Latency (ns)	(µs)				
L1 cache reference	0.5ns					
L2 cache reference	7 ns		14x L1 cache			
Mutex lock/unlock	25 ns					
Main memory reference	100 ns		20x L2 cache, 200x L1			
Read 4K randomly from SSD*	150,000 ns	150 µs	~1GB/sec SSD			
Read 1 MB sequentially from memory	250,000 ns	250 µs				
Read 1 MB sequentially from SSD*	1,000,000 ns	1,000 µs	1 ms ~1GB/sec SSD, 4X memory			
Read 1 MB sequentially from disk	20,000,000 ns	20,000 µs	20 ms 80x memory, 20X SSD			
 Latency numbers every progra From: https://gist.github.com 	immer should know /jboner/2841832#f	ile-latency-t	xt			
May 30, 2018 TCSS422: O Institute of	erating Systems [Spring 2018] echnology, University of Washington - Tacoma					











CACHE MANAGEMENT EXAMPLE						
Replacement	policies apply to "any" cache					
Goal is to min	imize the number of misses					
Average mem	ory access time (AMAT) can be estimated:					
	$AMAT = (P_{Hit} * T_M) + (P_{Miss} * T_D)$					
Argument	Meaning					
T _M	The cost of accessing memory (time)					
T _D	The cost of accessing disk (time)					
P _{Hit}	The probability of finding the data item in the cache(a hit)					
P _{Miss}	The probability of not finding the data in the cache(a miss)					
Consider T _M =	100 ns, T _D = 10ms					
For a batch of memory accesses:						
Consider P _{hit} = .9 (90%), P _{miss} = .1						
Consider P _{hit}	Consider P _{hit} = .999 (99.9%), P _{miss} = .001					
May 30, 2018	TCSS422: Operating Systems [Spring 2018] Institute of Technology, University of Washington - Tacoma	L16.13				











































POLLING						
 OS checks if device is <i>READY</i> by repeatedly checking the STATUS register Simple approach CPU cycles are wasted without doing meaningful work Ok if only a few cycles, for rapid devices that are often READY BUT polling, as with "spin locks" we understand is inefficient 						
'waiting 10' 1 : task 1 P : polling						
CPU 1 1 1 1 1 p p p p 1 1 1 1 1 1						
Disk 1 1 1 1 1						
CPU utilization by polling						
May 30, 2018 TCSS422: Operating Systems [Spring 2018] Institute of Technology, University of Washington - Tacoma L16.37						









Transfer Modes						
Mode •	# •	Maximum transfer rate (MB/s)	cycle time ¢			
PIO	0	3.3	600 ns			
	1	5.2	383 ns			
	2	8.3	240 ns			
	3	11.1	180 ns			
	4	16.7	120 ns			
Single-word DMA	0	2.1	960 ns			
	1	4.2	480 ns			
	2	8.3	240 ns			
Multi-word DMA	0	4.2	480 ns			
	1	13.3	150 ns			
	2	16.7	120 ns			
	3[34]	20	100 ns			
	4[34]	25	80 ns			
Ultra DMA	0	16.7	240 ns ÷ 2			
	1	25.0	160 ns ÷ 2			
	2 (Ultra ATA/33)	33.3	120 ns + 2			
	3	44.4	90 ns + 2			
	4 (Ultra ATA/66)	66.7	60 ns ÷ 2			
	5 (Ultra ATA/100)	100	40 ns ÷ 2			
	6 (Ultra ATA/133)	133	30 ns ÷ 2			
	7 (Ultra ATA/167)[35]	167	24 ns + 2			





















Many devices provide special capabilities

- Example: SCSI Error handling
- SCSI devices provide extra detail which are lost to the OS

Buggy device drivers

- 70% of OS code is in device drivers
- Device drivers are required for every device plugged in
 Drivers are often 3rd party, which is not quality controlled at
- the same level as the OS (Linux, Windows, MacOS, etc.)

May 30, 2018 TCSS422: Operating Systems [Spring 2018] Institute of Technology, University of Washington - Tacoma



L16.52