



How to streamline your life (lessons from computer architecture).

COS 116

3/29/2011

Instructor: Sanjeev Arora

Lesson 1: Caching

(and the 80-20 rule)



Customer Rating



4.7 out of 5

XPS 420

Intel® Core™2 Q6600

Quad-Core (8MB L2

cache, 2.4GHz, 1066FSB)

(Aside:


1066 FSB = Front side bus operates at 1066Mhz

(upper limit on the speed of data
communication of CPU with rest of computer)

Quad-core: Processor with 4 separate built-in processors)

The Tired Librarian

Reserves



50 ft



- 1000 checkouts/returns per day
- Distance covered = $50 \times 2 \times 1000 = 100,000$ feet
~ 20 miles
- Please help!!!




80-20 “Rule”

- Pareto [1906]: 20% of the people own 80% of the wealth
- Juran [1930' s]: 20% of the organization does 80% of the work

Better Arrangement

Reserves




“Most popular” shelf:
20% most popular
books



- Distance covered per day?

Even better arrangement?


Reserves



“Most popular” shelf:
20% most popular
books



Top 4%



- Distance covered per day?



Discussion Time

- Is the librarian's problem solved?

How to predict the 20% most popular books for next day? Suggestions?



How to predict the 20% most popular books for next day?

- In general, no easy solution
- In practice, use rules of thumb
 - Example: “Least Recently Used”. When you need to create space on the desk (or shelf), move out the book that was used least recently
 - Many others (LRU is computationally expensive)


Connection to Computer Organization

- Speed vs cost of various memories


	Cost: \$ / GB	Speed: GB/s
Hard drive	0.05	0.3
RAM	50	1.5
On-chip memory for CPU (L2 Cache)	40000	15

Computer ~~Librarian~~ arrangement

Disks
Registers



“Most popular” shelf:
20% most popular
books Memory



Often, today's computers have even more levels of caching

Question



- How does the same program (.exe file) run on different PCs with different memory configurations?
- Answer: “Virtual Memory”
 - All programs live a fiction: allowed to pretend it has 2^{64} bytes of memory
 - Illusion is preserved by hardware

Goodbye Lenin

"The German Democratic Republic lives on – in 79 m²!"

(Die DDR lebt weiter – auf 79 qm!)



Virtual Memory


■ Program's view:

Powerpoint



Memory:

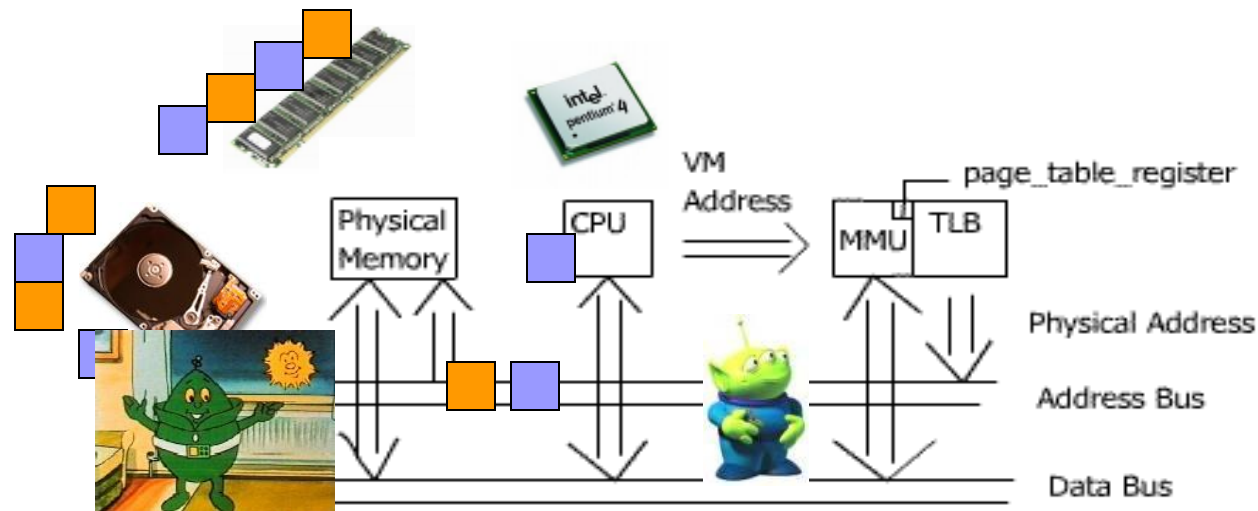
Lec15.ppt P ≠ NP.ppt



Address 0

Address $2^{64} - 1$

■ Underlying truth:





Discussion Time

	Cost: \$ / GB	Speed: GB/s
Hard drive	0.05	0.3
RAM	50	5

Assuming the 80-20 rule, what is the maximum data transfer speed the CPU can expect to see?

How does your answer change if cost is no object?

What determines RAM size?

Divide into groups of 3 and hand in your calculations & thoughts.



Moral

- Performance:

- Speed is close to that of fastest memory (cache)
- Overall capacity is that of largest memory (disk)

Lesson 2: Multitasking

- “The Multitasking Generation”



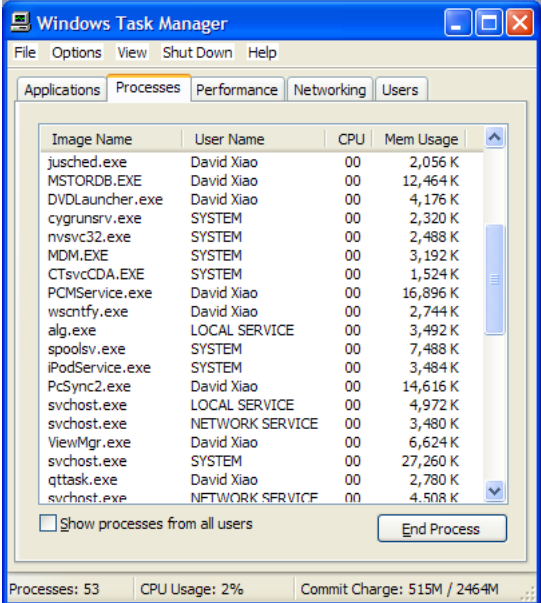
An Evening's Tasks for a Gen-M'er

- Homework
 - Listen to music
 - Instant Messaging
 - Call Mom (goes to bed by 11 PM!)
 - Answer phone
 - Read a bit more of Joyce's *Ulysses*
 - Watch the Daily Show
- How do you do it all? Can the brain do multiple things at a single instant?



Tasks done by my PC last night

- Word processing
- Play CD
- Download news updates
- Download email
- Run clock
- Hidden tasks: handle network traffic, manage disk and RAM traffic, scheduler, etc.



The screenshot shows the Windows Task Manager window with the 'Processes' tab selected. The window title is 'Windows Task Manager' and it has a menu bar with 'File', 'Options', 'View', 'Shut Down', and 'Help'. Below the menu bar are tabs for 'Applications', 'Processes', 'Performance', 'Networking', and 'Users'. The 'Processes' tab is active, displaying a list of running processes with columns for 'Image Name', 'User Name', 'CPU', and 'Mem Usage'. The status bar at the bottom indicates 'Processes: 53', 'CPU Usage: 2%', and 'Commit Charge: 515M / 2464M'. There is also an 'End Process' button at the bottom right of the list area.

Image Name	User Name	CPU	Mem Usage
jusched.exe	David Xiao	00	2,056 K
MSTORDB.EXE	David Xiao	00	12,464 K
DVDLauncher.exe	David Xiao	00	4,176 K
cygrunsvr.exe	SYSTEM	00	2,320 K
nsvsc32.exe	SYSTEM	00	2,488 K
MDM.EXE	SYSTEM	00	3,192 K
CTsvcCDA.EXE	SYSTEM	00	1,524 K
PCMSservice.exe	David Xiao	00	16,896 K
wscntfy.exe	David Xiao	00	2,744 K
alg.exe	LOCAL SERVICE	00	3,492 K
spoolsv.exe	SYSTEM	00	7,488 K
iPodService.exe	SYSTEM	00	3,484 K
PcSync2.exe	David Xiao	00	14,616 K
svchost.exe	LOCAL SERVICE	00	4,972 K
svchost.exe	NETWORK SERVICE	00	3,480 K
ViewMgr.exe	David Xiao	00	6,624 K
svchost.exe	SYSTEM	00	27,260 K
qtask.exe	David Xiao	00	2,780 K
svchost.exe	NETWORK SERVICE	00	4,508 K

Managed by “Operating System”
(WinXP, Linux, MacOS, etc.)

Scheduler's objectives

- Fairness
- Timeliness
- Critical tasks processed promptly
- Low overhead



Discussion Time

How can one achieve these (often conflicting) goals?

Multitasking versus Parallel Processing

Multitasking: A single CPU handles many tasks by switching rapidly among them.

(e.g., all Wintel machines since early 1990s; all Unix machines since the 1970s)

Parallel Processing: Multiple CPUs that do the work of a single CPU. (But, 4 CPUs do not necessarily mean 4x speed.)

XPS 420

Intel® Core™2 Q6600

Quad-Core (8MB L2

cache, 2.4GHz, 1066FSB)

The Legal View....

UNITED STATES DISTRICT COURT
FOR THE DISTRICT OF COLUMBIA

UNITED STATES OF AMERICA,

Plaintiff

v.

MICROSOFT CORPORATION,

Defendant.

Civil Action No. 98-1232 (CKK)

FINAL JUDGMENT
(November 12, 2002)

WHEREAS, plaintiffs United States of America ("United States") and the States of New York, Ohio, Illinois, Kentucky, Louisiana, Maryland, Michigan, North Carolina and Wisconsin and defendant Microsoft Corporation ("Microsoft"), by their respective attorneys, have consented to the entry of this Final Judgment;

One main point studied by the judge:

What is an OS?

Software that always runs on the computer in the background and acts as intermediary between the underlying hardware and all software applications