inst.eecs.berkeley.edu/~cs61c

CS61C : Machine Structures

Lecture #27
Performance II & Summary

2005-12-07

There is one handout today at the front and back of the room!

Lecturer PSOE, new dad Dan Garcia

www.cs.berkeley.edu/~ddgarcia

The ultimate gift under the tree? ⇒

Since it's the holiday season, it's time to to consider what would be the best gift for a CS61C student. Nothing says "I love you" like a $2,300 81-game retro system, available <u>today</u> @ Costco. Xbox360 who?

www.costco.com/Browse/Product.aspx?prodid=11098104

Review

- RAID
  - Motivation: In the 1980s, there were 2 classes of drives: expensive, big for enterprises and small for PCs. They thought "make one big out of many small!"
  - Higher performance with more disk arms/$, adds option for small # of extra disks (the R)
  - Started @ Cal by CS Profs Katz & Patterson
- Latency v. Throughput
- Performance doesn't depend on any single factor: need Instruction Count, Clocks Per Instruction (CPI) and Clock Rate to get valid estimations
- User Time: time user waits for program to execute: depends heavily on how OS switches between tasks
- CPU Time: time spent executing a single program: depends solely on processor design (datapath, pipelining effectiveness, caches, etc.)

$$\text{CPU time} = \frac{\text{Instructions}}{\text{Program}} \times \frac{\text{Cycles}}{\text{Instruction}} \times \frac{\text{Seconds}}{\text{Cycle}}$$

What Programs Measure for Comparison?

- Ideally run typical programs with typical input before purchase, or before even build machine
  - Called a "workload"; For example:
  - Engineer uses compiler, spreadsheet
  - Author uses word processor, drawing program, compression software
- In some situations its hard to do
  - Don't have access to machine to "benchmark" before purchase
  - Don't know workload in future
- Next: benchmarks & PC-Mac showdown!

Benchmarks

- Obviously, apparent speed of processor depends on code used to test it
- Need industry standards so that different processors can be fairly compared
- Companies exist that create these benchmarks: "typical" code used to evaluate systems
- Need to be changed every 2 or 3 years since designers could (and do!) target for these standard benchmarks

Example Standardized Benchmarks (1/2)

- Standard Performance Evaluation Corporation (SPEC) SPEC CPU2000
  - CINT2000 12 integer (gzip, gcc, crafty, perl, ...)
  - CFP2000 14 floating-point (swim, mesa, art, ...)
  - All relative to base machine Sun 300MHz 256Mb-RAM Ultra5_10, which gets score of 100
  - `www.spec.org/osg/cpu2000/`
  - They measure
    - System speed (SPECint2000)
    - System throughput (SPECint_rate2000)

Example Standardized Benchmarks (2/2)

- SPEC
  - Benchmarks distributed in source code
  - Members of consortium select workload
    - 30+ companies, 40+ universities
  - Compiler, machine designers target benchmarks, so try to change every 3 years
  - The last benchmark released was SPEC 2000
    - They are still finalizing SPEC 2005

CINT2000

gzip	C	Compression
vpr	C	FPGA Circuit Placement and Routing
gcc	C	C Programming Language Compiler
mcf	C	Combinatorial Optimization
crafty	C	Game Playing: Chess
parser	C	Word Processing
eon	C++	Computer Visualization
perlbmk	C	PERL Programming Language
gap	C	Group Theory, Interpreter
vortex	C	Object-oriented Database
bzip2	C	Compression
twolf	C	Place and Route Simulator

CFP2000

wupwise	Fortran77	Physics / Quantum Chromodynamics
swim	Fortran77	Shallow Water Modeling
mgrid	Fortran77	Multi-grid Solver: 3D Potential Field
applu	Fortran77	Parabolic / Elliptic Partial Diff Equations
mesa	C	3-D Graphics Library
galgel	Fortran90	Computational Fluid Dynamics
art	C	Image Recognition / Neural Networks
equake	C	Seismic Wave Propagation Simulation
facerec	Fortran90	Image Processing: Face Recognition
ammp	C	Computational Chemistry
lucas	Fortran90	Number Theory / Primality Testing
fma3d	Fortran90	Finite-element Crash Simulation
sixtrack	Fortran77	High Energy Nuclear Physics Accelerator Design
apsi	Fortran77	Meteorology: Pollutant Distribution

Example PC Workload Benchmark

- PCs: Ziff-Davis Benchmark Suite
  - "Business Winstone is a system-level, application-based benchmark that measures a PC's overall performance when running today's top-selling Windows-based 32-bit applications… it doesn't mimic what these packages do; it runs real applications through a series of scripted activities and uses the time a PC takes to complete those activities to produce its performance scores.
  - Also tests for CDs, Content-creation, Audio, 3D graphics, battery life

`http://www.etestinglabs.com/benchmarks/`

Performance Evaluation

- Good products created when have:
  - Good benchmarks
  - Good ways to summarize performance
- Given sales is a function of performance relative to competition, should invest in improving product as reported by performance summary?
- If benchmarks/summary inadequate, then choose between improving product for real programs vs. improving product to get more sales; Sales almost always wins!

Performance Evaluation: The Demo

If we're talking about performance, let's discuss the ways shady salespeople have fooled consumers (so that you don't get taken!)

5. Never let the user touch it
4. Only run the demo through a script
3. Run it on a stock machine in which "no expense was spared"
2. Preprocess all available data
1. Play a movie

PC / PC / Mac Showdown!!! (1/4)

- PC
  - 1 GHz Pentium III
  - 256 Mb RAM
  - 512KB L2 Cache
  - No L3
  - 133 MHz Bus
  - 20 GB Disk
  - 16MB VRAM
- PC 800MHz PIII
- Mac
  - 800 MHz PowerbookG4
  - 1 Gb RAM
    - 2 512Mb SODIMMs
  - 32KB L1Inst, L1Data
  - 256KB L2 Cache
  - 1Mb L3 Cache
  - 133 MHz Bus
  - 40 GB Disk
  - 32MB VRAM

Let's take a look at SPEC2000 and a simulation of a real-world application.

PC / Mac Showdown!!! (2/4)

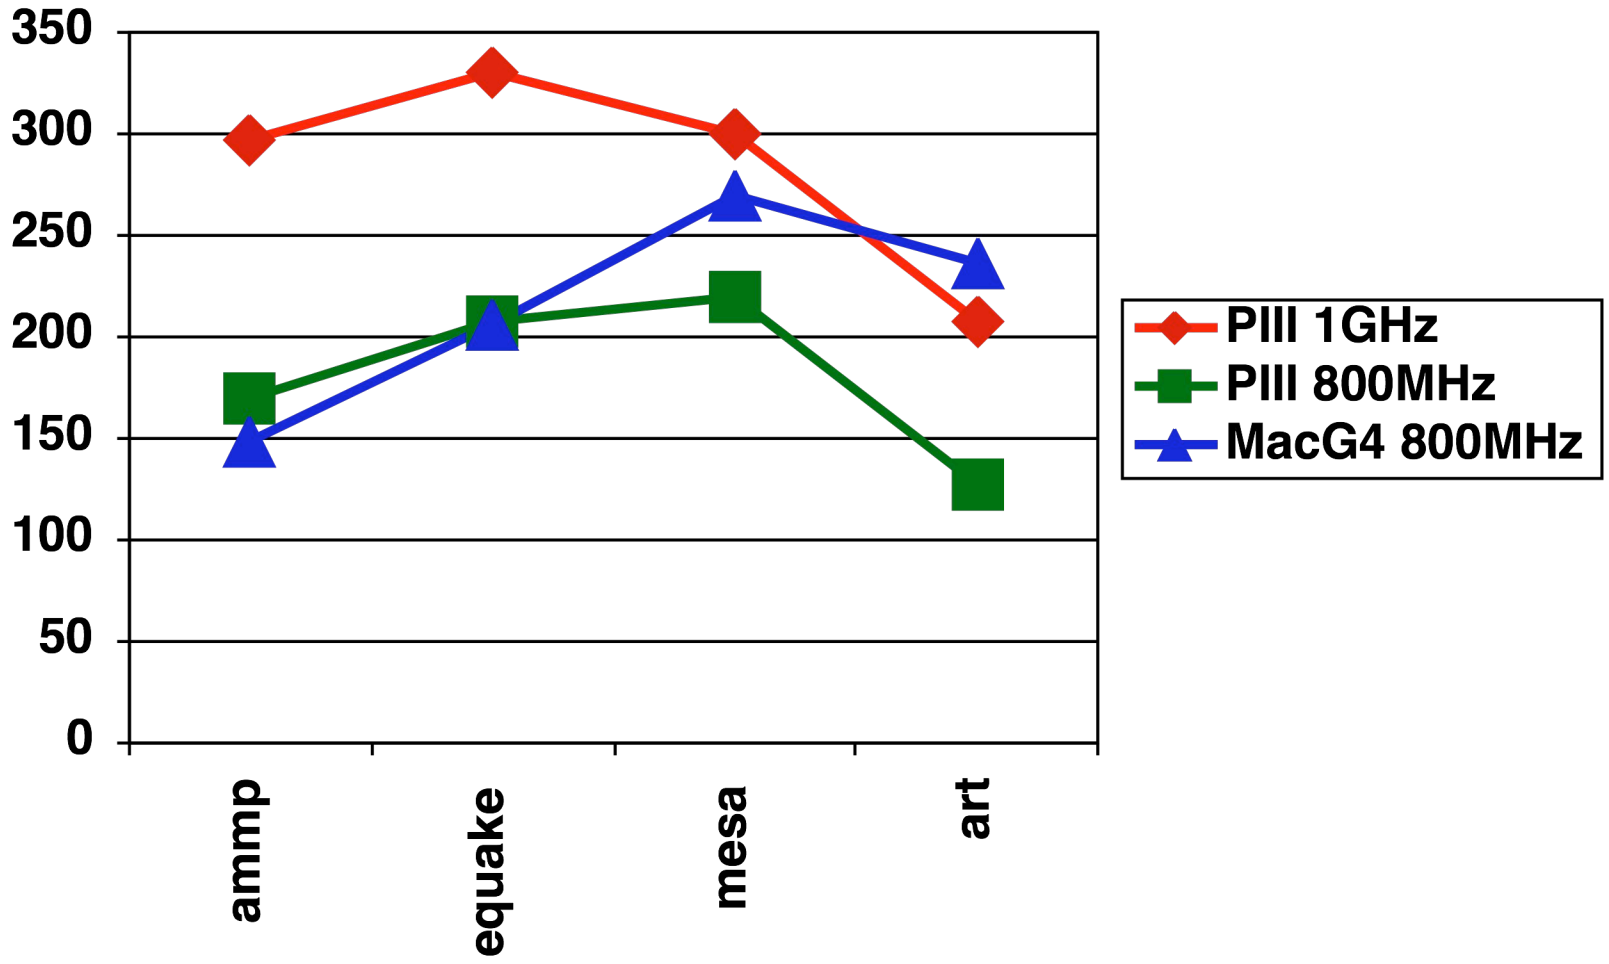

CFP2000 (bigger better)

[left-to-right by G4/PIII 800MHz ratio]

PC / Mac Showdown!!! (3/4)

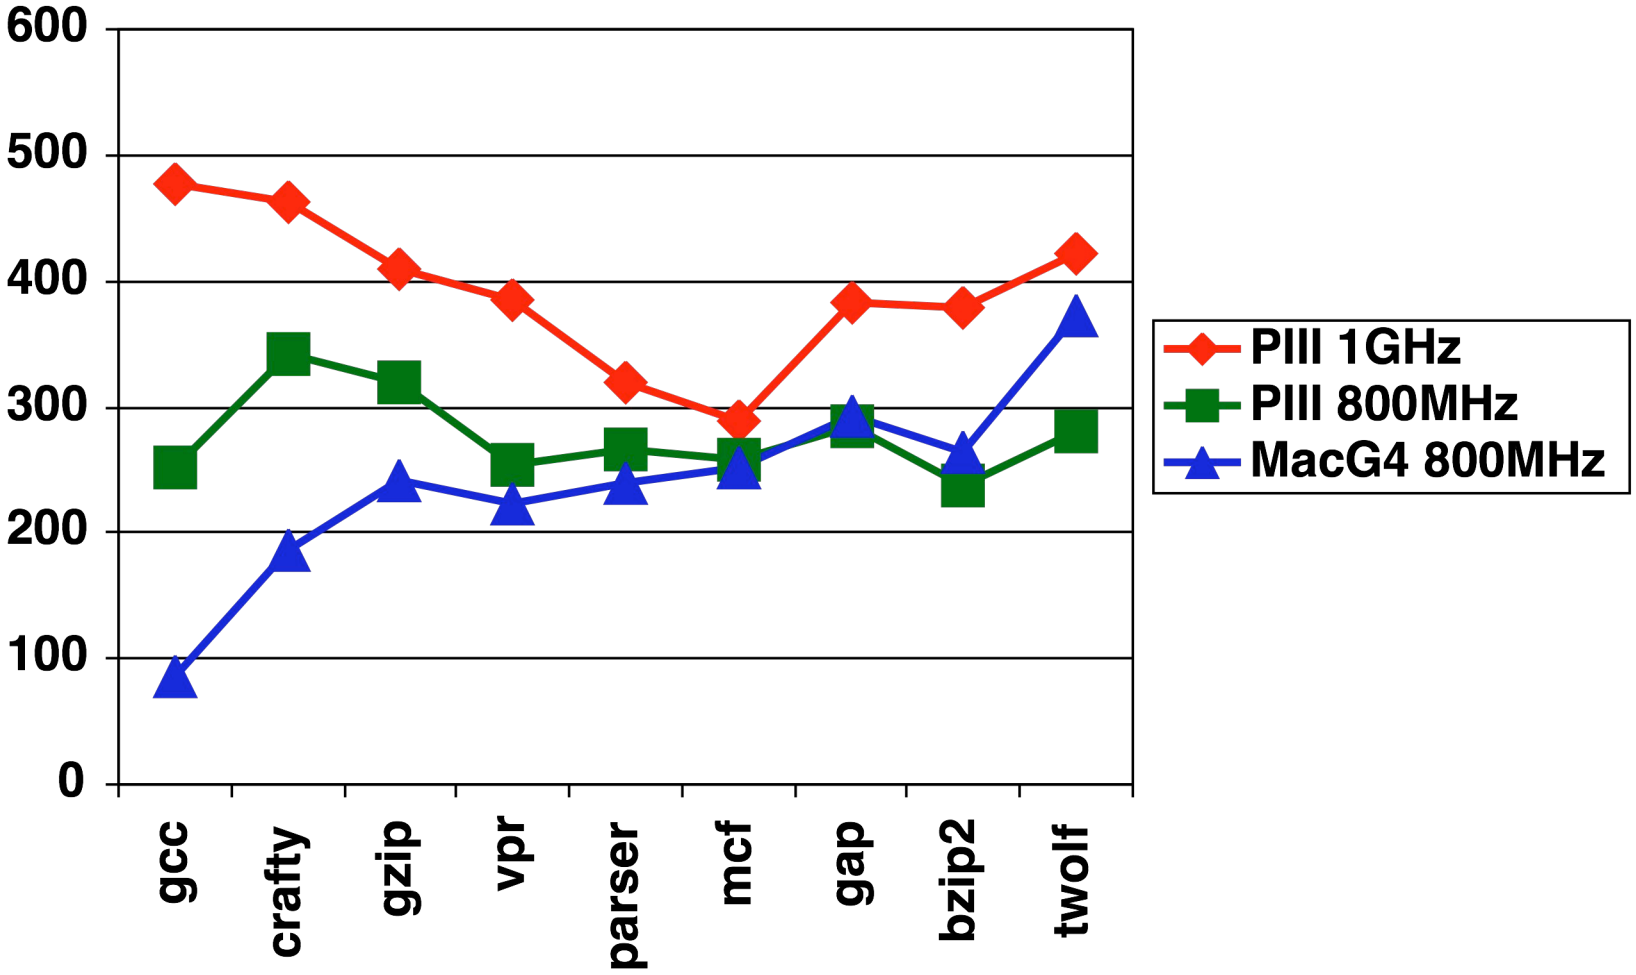

CINT2000 (bigger better)

[left-to-right by G4/PIII 800MHz ratio]

PC / Mac Showdown!!! (4/4)

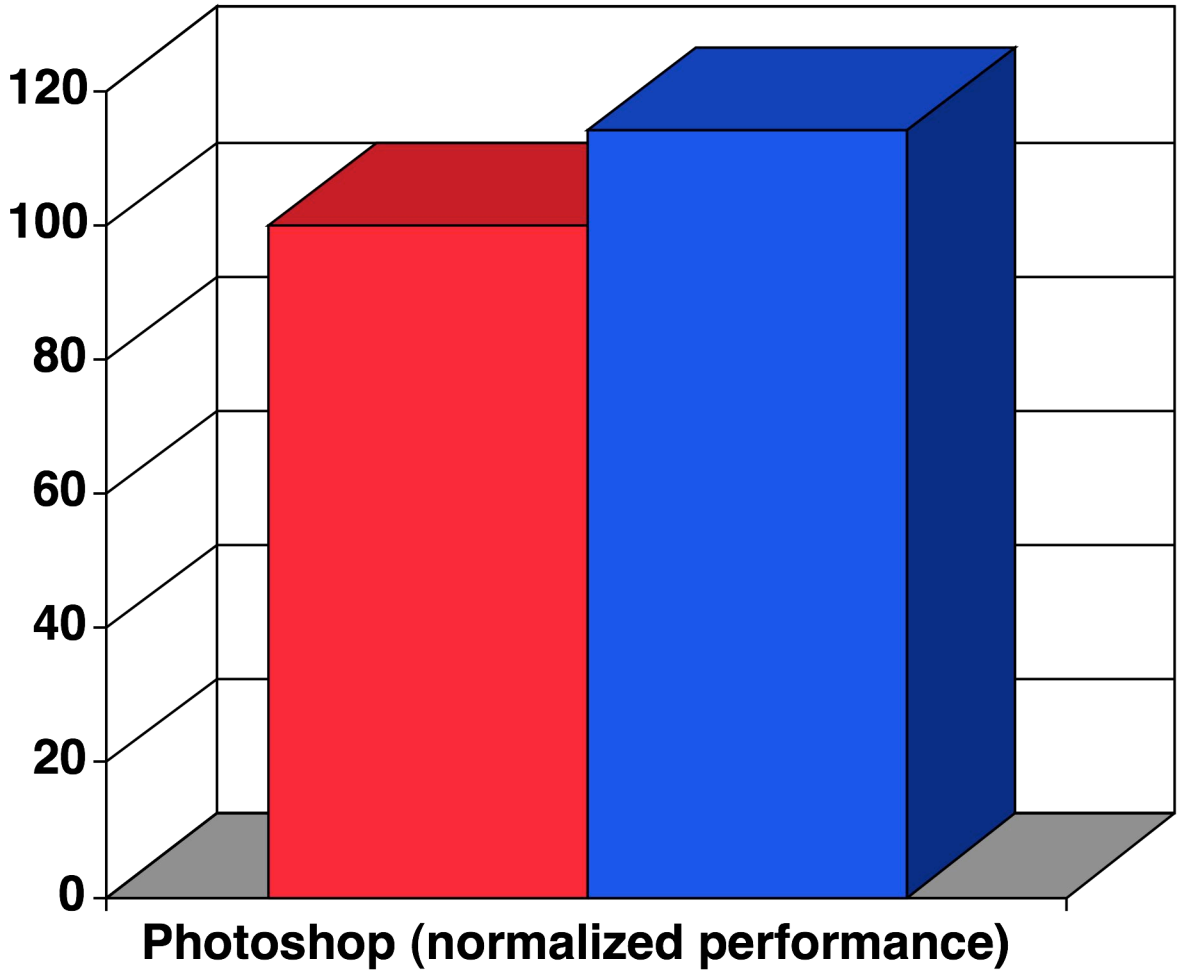

…Apple got in a heap of trouble when claiming the G5 was the "worlds fastest personal computer"

…lies, damn lies, and statistics.

Normalized Photoshop radial blur (bigger better)

[Amt=10,Zoom,Best](PIII = 79sec = "100", G4= 69sec)

Administrivia

- If you did well in CS3 or 61{A,B,C} (A- or above) and want to be on staff?
  - Usual path: Lab assistant ⇒ Reader ⇒ TA
  - Fill in form outside 367 Soda before first week of semester…
  - I strongly encourage anyone who gets above a B+ in the class to follow this path…
- Sp04 Final exam + solutions online!
- Final Review: 2005-12-11 @ 2pm in 10 Evans
- Final: 2005-12-17 @ 12:30pm in 2050 VLSB
  - Only bring pen{,cil}s, two 8.5”x11” handwritten sheets + green. Leave backpacks, books, calculators, cells & pagers home!

Upcoming Calendar

Week #	Mon	Wed	Thu Lab	Sat
#15 Last Week o' Classes	Performance	LAST CLASS Summary, Review, & HKN Evals	I/O Networking & 61C Feedback Survey	
#16 Sun 2pm Review 10 Evans	Performance competition due tonight @ midnight			FINAL EXAM SAT 12-17 @ 12:30pm-3:30pm 2050 VLSB Performance awards

CS61C: So what's in it for me? (1st lecture)

Learn some of the big ideas in CS & engineering:

- 5 Classic components of a Computer
- Principle of abstraction, systems built as layers
- Data can be anything (integers, floating point, characters): a program determines what it is
- Stored program concept: instructions just data
- Compilation v. interpretation thru system layers
- Principle of Locality, exploited via a memory hierarchy (cache)
- Greater performance by exploiting parallelism (pipelining)
- Principles/Pitfalls of Performance Measurement

Conventional Wisdom (CW) in Comp Arch

Thanks to Dave Patterson for these

- Old CW: Power free, Transistors expensive
- New CW: Power expensive, Transistors free
  - Can put more on chip than can afford to turn on
- Old CW: Chips reliable internally, errors at pins
- New CW: ≤ 65 nm ⇒ high error rates
- Old CW: CPU manufacturers minds closed
- New CW: Power wall + Memory gap = Brick wall
  - New idea-receptive environment
- Old CW: Uniprocessor performance 2X / 1.5 yrs
- New CW: 2X CPUs per socket / ~ 2 to 3 years
  - More simpler processors more power efficient

Massively Parallel Socket

- Processor = new transistor?
  - Does it only help power/cost/performance?
- Intel 4004 (1971): 4-bit processor, 2312 transistors, 0.4 MHz, 10 μm PMOS, 11 mm^2 chip

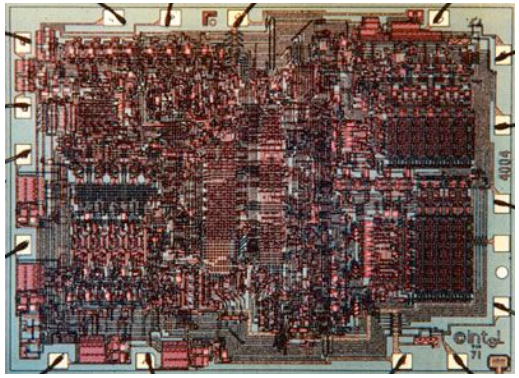

- RISC II (1983): 32-bit, 5 stage pipeline, 40,760 transistors, 3 MHz, 3 μm NMOS, 60 mm^2 chip
  - 4004 shrinks to ~ 1 mm^2 at 3 micron

- 125 mm^2 chip, 65 nm CMOS = 2312 RISC IIs + Icache + Dcache
  - RISC II shrinks to ~ 0.02 mm^2 at 65 nm
  - Caches via DRAM or 1 transistor SRAM (www.t-ram.com)?
  - Proximity Communication at > 1 TB/s ?
  - Ivan Sutherland @ Sun spending time in Berkeley!

20th vs. 21st Century IT Targets

- 20th Century Measure of Success
  - Performance (peak vs. delivered)
  - Cost (purchase cost vs. ownership cost, power)
- 21st Century Measure of Success? "SPUR"
  - Security
  - Privacy
  - Usability
  - Reliability
- Massive parallelism greater chance (this time) if
  - Measure of success is SPUR vs. only cost-perf
  - Uniprocessor performance improvement decelerates

Other Implications

- Need to revisit chronic unsolved problem
  - Parallel programming!!
- Implications for applications:
  - Computing power >>> CDC6600, Cray XMP (choose your favorite) on an economical die inside your watch, cell phone or PDA
    - On your body health monitoring
    - Google + library of congress on your PDA
- As devices continue to shrink…
  - The need for great HCI critical as ever!

Taking advantage of Cal Opportunities

"The Godfather answers all of life's questions"
– Heard in "You've got Mail"

- **Why are we the #2 Univ in the WORLD?**
  So says the 2004 ranking from the "Times Higher Education Supplement"
  - **Research, reseach, research!**
  - **Whether you want to go to grad school or industry, you need someone to vouch for you! (as is the case with the Mob)**
- **Techniques**
  - **Find out what you like, do lots of web research (read published papers), hit OH of Prof, show enthusiasm & initiative**

`http://research.berkeley.edu/`

Dan's CS98/198 Opportunities Spring 2006

- **GamesCrafters (Game Theory R & D)**
  - We are developing SW, analysis on small 2-person games of no chance. (e.g., achi, connect-4, dots-and-boxes, etc.)
  - Req: A- in CS61C, Game Theory Interest
  - `http://GamesCrafters.berkeley.edu`
- **MS-DOS X (Mac Student Developers)**
  - Learn to program Macintoshes. No requirements (other than Mac, interest)
  - `http://msdosx.berkeley.edu`
- **UCBUGG (Recreational Graphics)**
  - Develop computer-generated images and animations.
  - `http://ucbugg.berkeley.edu`

Penultimate slide: Thanks to the staff!

- TAs
  - Head TA
    Jeremy Huddleston
  - Zhangxi Tan
  - Michael Le
  - Navtej Sadhal
- Readers
  - Mario Tanev
  - Mark Whitney

Thanks to Dave Patterson for these CS61C notes…

The Future for Future Cal Alumni

- What's The Future?
- New Millennium
  - Internet, Wireless, Nanotechnology, ...
  - Rapid Changes in Technology
  - World's (2nd) Best Education
  - Never Give Up!

"The best way to predict the future is to invent it" – Alan Kay

The Future is up to you!