

L1: Intro to Computer Systems: Complexity

Robert Morris and Frans Kaashoek

6.033 Spring 2009

<http://web.mit.edu/6.033>

<http://web.mit.edu/6.033>

- Schedule has all assignments
 - Every meeting has preparation/assignment
 - First deliverable is due Tuesday
 - Read the Therac paper for Friday
- Return sign-up sheet at the end of lecture (if you didn't do so yesterday)
 - We will post sections assignment tonight

Monday	Tuesday	Wednesday	Thursday	Friday
feb.2 <i>Registration Day</i>	feb.3 REC 1 Worse is Better <i>First day of classes</i>	feb.4 LEC 1 Intro to systems Preparation: Read 1.1, 1.2, 1.3 (click on link)	feb.5 <small>REC 2</small> The Architecture of Complexity Preparation: Read Simon paper	feb.6 Writing Program Recitation Preparation: Read Therac-25 paper and Writing Assignment Assigned: One-pager #1

What is a system? hidden

- 6.033 is about the design of s/w systems.
- System = [diagram]
 - Inside vs outside
 - achieve specific external behavior
 - many components
- Examples: Bank ATMs, Web
- Much of 6.033 will operate at design level
 - Relationships of components
 - Internals of components that help structure

Problem: Complexity hidden

- Hard to define; symptoms:
 - Large # of components
 - Large # of connections
 - Irregular
 - No short description
 - Many people required to design/maintain
- Technology rarely the limit!
 - Indeed tech opportunity is the problem
 - Limit is usually designers' understanding

6.033 Approach hidden

- Lectures/book: big ideas, technology, examples
- Recitations: papers, discussion
 - Design examples
 - Writing examples: core prob/soln vs detail
 - Learn how to read a paper, skim vs meat
- Design projects: practice designing and writing
 - Design: choose problem, tradeoffs, structure
 - Writing: explain core ideas concisely
- Exams: focus on reasoning about system design
- Ex-6.033 students: papers and recitations

Example 6.033 Readings

Therac-25

UNIX

Ethernet

End-to-End Arguments

System R

Papers hidden

- ONE SENTENCE EACH
- Therac: bad design, at many levels. detailed post-mortem.
- UNIX: successful “New Jersey” design. small, careful choice of problems to solve.
- Ethernet: elegant, beat competition, scaled by 3000x!
- End-to-end: captures a non-obvious and very useful philosophical point.
- System R: huge influence on construction of fault-tolerant systems - recovery from crash at any point.

Problem Types hidden

- Emergent properties
 - surprises
- Propagation of effects
 - Small change -> big effect
- [Incommensurate] scaling
 - Design for small model may not scale

Lessons hidden, running

- Expect surprises
- There is no small change
- 10x increase \Rightarrow perhaps re-design
- Just one more feature!
- Complexity is super-linear
- Performance cost is super-linear


Emergent Property Example: Ethernet

- All computers share single cable
- Goal is reliable delivery
- Listen before send to avoid collisions
- Will listen-while-send detect collisions?
- Maximum cable length!
- Minimum packet size!

Propagation of Effects Example (L. Cole 1969)

- WHO attempted to control malaria in North Borneo
- Sprayed villages with DDT
- Wiped out mosquitoes, but
 - Roaches collected DDT in tissue
 - Lizards ate roaches and became slower
 - Easy target for cats
 - Cats didn't deal with DDT well and died
 - Forest rats moved into villages
 - Rats carried the bacillus for the plague
- WHO replaced malaria with the plague

Galileo in 1638



“To illustrate briefly, I have sketched a bone whose natural length has been increased three times and whose thickness has been multiplied until, for a correspondingly large animal, it would perform the same function which the small bone performs for its small animal. From the figures here shown you can see how out of proportion the enlarged bone appears. Clearly then if one wishes to maintain in a great giant the same proportion of limb as that found in an ordinary man he must either find a harder and stronger material for making the bones, or he must admit a diminution of strength in comparison with men of medium stature; for if his height be increased inordinately he will fall and be crushed under his own weight. Whereas, if the size of a body be diminished, the strength of that body is not diminished in the same proportion; indeed the smaller the body the greater its relative strength. Thus a small dog could probably carry on his back two or three dogs of his own size; but I believe that a horse could not carry even one of his own size.” [Dialog Concerning Two New Sciences, 2nd Day]

Sources of Complexity hidden


- Many goals/requirements
- Interaction of features
- Performance

Example: more goals, more complexity


- 1975 Unix kernel: 10,500 lines of code
- 2008 Linux 2.6.24 line counts:
 - 85,000 processes
 - 430,000 sound drivers
 - 490,000 network protocols
 - 710,000 file systems
 - 1,000,000 different CPU architectures
 - 4,000,000 drivers
 - 7,800,000 Total

Example: interacting features, more complexity

- Call Forwarding
- Call Number Delivery Blocking
- Automatic Call Back
- Itemized Billing



CNDB



ACB + IB



- A calls B, B is busy
- Once B is done, B calls A
- A's number on appears on B's bill

Interacting Features hidden

- Each feature has a spec.
- An interaction is bad if feature X breaks feature Y.
- ...
- The point is not that these bad interactions can't be fixed.
- The point is that there are so many interactions that have to be considered: they are a huge source of complexity.
- Perhaps more than n^2 interactions, e.g. triples.
- Cost of thinking about / fixing interaction gradually grows to dominate s/w costs.
- The point: Complexity is super-linear

Example: single rail hidden

- Performance -> complexity
- One track in a narrow canyon [diagram]
- Base design: alternate trains
 - Low throughput, high delay
 - Worse than two-track, cheaper than blasting
- Lower delay w/ a siding and two trains
 - Precise schedule
 - Risk of collision / signal lights
 - Siding limits train length (a global effect!)
- Point: performance cost super-linear

Coping with Complexity hidden

- Simplifying insights / experience
- Modularity
 - Split up system, consider separately
- Abstraction
 - Interfaces/hiding, e.g. standard size windows
 - Helps avoid propagation of effects
- Hierarchy
 - Reduce connections
 - Divide-and-conquer
- Layering
 - Gradually build up capabilities

Class plan

- Next lecture: computer systems are different
- Naming: gluing modules together
- Client/server: enforced modularity
- Networks: hard boundaries between modules
- Reliability and transactions: handling failures
- Security: handling malicious failures