

CSE503: Software Engineering

David Notkin
University of Washington
Department of Computer Science & Engineering
Spring 2006

1

Very High-Level View

- Requirements define the clients' view
 - What the system is supposed to do
 - Focuses on external behavior
- Design captures the developers' view
 - How the requirements are realized
 - Defines the internal structure of the solution
- But: “What” vs. “How”
- Also, reminiscent of the Brian Cantwell Smith diagram in Jackson's video

2

Today

- Some general background on design
 - Decomposition, composition
 - Managing complexity, anticipating change
 - Etc.
- Then ...
 - Information hiding
 - Layering/ “uses” relation
 - Software architecture, patterns, etc.
 - Later in the quarter – cross-cutting concerns (aspect-oriented approaches) by Miryung

3

Complexity

- “Software entities are more complex for their size than perhaps any other human construct, because no two parts are alike (at least above the statement level). If they are, we make the two similar parts into one... In this respect software systems differ profoundly from computers, buildings, or automobiles, where repeated elements abound.” —Brooks, 1986

4

Continuous & iterative

- High-level (“architectural”) design
 - What pieces?
 - How connected?
- Low-level design
 - Should I use a hash table or binary search tree?
- Very low-level design
 - Variable naming, specific control constructs, etc.
 - About 1000 design decisions at various levels are made in producing a single page of code

5

Multiple design choices

- There are multiple (perhaps unbounded) designs that satisfy (at least the functional) aspects of a given set of requirements
- How does one choose among these alternatives?
 - How does one even identify the alternatives?
 - How does one reject most bad choices quickly?
 - What criteria distinguish good choices from bad choices?

6

What criteria?

- In general, there are three high level answers to this question: and, it is very difficult to answer precisely
 1. Satisfying functional and performance requirements
 - Maybe this is too obvious to include
 - Often not achieved, though
 2. Managing complexity
 3. Accommodating future change

7

1. Managing complexity

- “The technique of mastering complexity has been known since ancient times: *Divide et impera* (Divide and Rule).” —Dijkstra, 1965
- “...as soon as the programmer only needs to consider intellectually manageable programs, the alternatives he is choosing from are much, much easier to cope with.” —Dijkstra, 1972
- The complexity of the software systems we are asked to develop is increasing, yet there are basic limits upon our ability to cope with this complexity. How then do we resolve this predicament?” —Booch, 1991

8

Brooks: accidental vs. inherent complexity

- It’s important to know which you are facing
- That is, which come from the problem itself and which come from your design, your platform, etc.

9

Divide and conquer

- We have to decompose large systems to be able to build them
 - The “modern” problem of composing systems from pieces is equally or more important
 - It’s not modern, though: we’ve had to compose for as long as we have decomposed
 - And closely related to decomposition in many ways
- For software, decomposition techniques are distinct from those used in physical systems
 - Fewer constraints are imposed by the material
 - Shanley principle?

10

Composition

- “Divide and conquer. Separate your concerns. Yes. But sometimes the conquered tribes must be reunited under the conquering ruler, and the separated concerns must be combined to serve a single purpose.” —M. Jackson, 1995
- Jackson’s view of composition as printing with four-color separation
- Composition in programs is not as easy as conjunction in logic

11

Benefits of decomposition

- Decrease size of tasks
- Support independent testing and analysis
- Separate work assignments
- Ease understanding

12

Which decomposition?

- How do we select a decomposition?
 - We determine the desired criteria
 - We select a decomposition (design) that will achieve those criteria
- In theory, that is; in practice, it's hard to
 - Determine the desired criteria with precision
 - Tradeoff among various conflicting criteria
 - Figure out if a design satisfies given criteria
 - Find a better one that satisfies more criteria
- In practice, it's easy to
 - Build something designed pretty much like the last one
 - This has benefits, too: understandability, properties of the pieces, etc.

13

Structure

- The focus of most design approaches is structure
- What are the components and how are they put together?
- Behavior is important, but largely indirectly
 - Satisfying functional and performance requirements

14

So what happens?

- People often buy into a particular approach or methodology
 - Ex: UML, structured analysis and design, object-oriented design, JSD, Hatley-Pirbair, etc.
- “Beware a methodologist who is more interested in his methodology than in your problem.” —M. Jackson

15

Conceptual integrity

- Brooks and others assert that conceptual integrity is a critical criterion in design
 - “It is better to have a system omit certain anomalous features and improvements, but to reflect one set of design ideas, than to have one that contains many good but independent and uncoordinated ideas.” —Brooks, MMM
- Such a design often makes it far easier to decide what is easy and reasonable to do as opposed to what is hard and less reasonable to do
 - This is not always what management wants to hear

16

2. Accommodating change

- “...accept the fact of change as a way of life, rather than an untoward and annoying exception.” —Brooks, 1974
- Software that does not change becomes useless over time. —Belady and Lehman
- Internet time makes the need to accommodate change even more apparent

17

Anticipating change

- It is generally believed that to accommodate change one must anticipate possible changes
 - Counterpoint: Extreme Programming
- By anticipating (and perhaps prioritizing) changes, one defines additional criteria for guiding the design activity
- It is not possible to anticipate all changes

18

Rationalism vs. empiricism: Brooks talk -- The Design of Design

- *rationalism* — the doctrine that knowledge is acquired by reason without resort to experience [WordNet]
- *empiricism* — the doctrine that knowledge derives from experience [WordNet]

19

Examples

- Life
 - Aristotle vs. Galileo
 - France vs. Britain
 - Descartes vs. Hume
 - Roman law vs. Anglo-Saxon law
- Software (Wegner)
 - Prolog vs. Lisp
 - Algol vs. Pascal
 - Dijkstra vs. Knuth
 - Proving programs vs. testing programs

20

Brooks' view

- Brooks says he is a “thoroughgoing, died-in-the-wool empiricist.”
- “Our designs are so complex there is no hope of getting them right first time by pure thought. To expect to is arrogant.”
- “So, we must adopt design-build processes that incorporate evolutionary growth ...
 - “Iteration, and restart if necessary
 - “Early prototyping and testing with *real users*”
- Maybe this is more an issue of requirements and specification, but I think it applies to design, too
 - “Plan to throw one away, you will anyway.”

21

Properties of design

- Cohesion
- Coupling
- Complexity
- Correctness
- Correspondence
- Makes designs “better”, one presumes
- Worth paying attention to

22

Cohesion

- The reason that elements are found together in a module
 - Ex: coincidental, temporal, functional, ...
- The details aren't critical, but the intent is useful
- During maintenance, one of the major structural degradations is in cohesion

23

Coupling

- Strength of interconnection between modules
- Hierarchies are touted as a wonderful coupling structure, limiting interconnections
 - But don't forget about composition, which requires some kind of coupling
- Coupling also degrades over time
 - “I just need one function from that module...”
 - Low coupling vs. no coupling

24

Unnecessary coupling hurts

- Propagates effects of changes more widely
- Harder to understand interfaces (interactions)
- Harder to understand the design
- Complicates managerial tasks
- Complicates or precludes reuse

25

It's easy to...

- ...reduce coupling by calling a system a single module
- ...increase cohesion by calling a system a single module

⇒ No satisfactory measure of coupling
– Either across modules or across a system

26

Complexity

- Well, yeah, simpler designs are better, all else being equal
- But, again, no useful measures of design/program complexity exist
 - Although there are dozens of such measures
 - My understanding is that, to the first order, most of these measures are linearly related to “lines of code”

27


Correctness

- Well, yeah
- Even if you “prove” modules are correct, composing the modules’ behaviors to determine the system’s behavior is hard
- Leveson and others have shown clearly that a system can fail even when each of the pieces work properly
 - Many systems have “emergent” properties

28

Correspondence

- “Problem-program mapping”
- The way in which the design is associated with the requirements
- The idea is that the simpler the mapping, the easier it will be to accommodate change in the design when the requirements change
- M. Jackson: problem frames
 - In the style of Polya



29

Functional decomposition

- Divide-and-conquer based on functions
 - `input;`
 - `compute;`
 - `output`
- Then proceed to decompose compute
- This is stepwise refinement (Wirth, 1971)
- There is an enormous body of work in this area, including many formal calculi to support the approach
 - Closely related to proving programs correct
- More effective in the face of stable requirements

30

Question

- To what degree do you consider your systems
 - as having modules?
 - as consisting of a set of files?
- This is a question of physical vs. logical structure of programs
 - In some languages/environments, they are one and the same
 - Ex: Smalltalk-80

31

Aside: Physical structure

- Almost all the literature focuses on logical structures in design
- But physical structure plays a big role in practice
 - Sharing
 - Separating work assignments
 - Degradation over time
- Why so little attention paid to this?

32

Recap

- High-level, but needed for basic discussion on design

33

Information hiding

- Information hiding is perhaps the most important intellectual tool developed to support software design [Parnas 1972]
 - Makes the anticipation of change a centerpiece in decomposition into modules
- Provides the fundamental motivation for abstract data type (ADT) languages
 - And thus a key idea in the OO world, too

34

Basics of information hiding

- Modularize based on anticipated change
 - Fundamentally different from Brooks' approach in OS/360 (see old and new MMM)
- Separate interfaces from implementations
 - Implementations capture decisions likely to change
 - Interfaces capture decisions unlikely to change
 - Clients know only interface, not implementation
 - Implementations know only interface, not clients
- Modules are also work assignments

35

Anticipated changes

- The most common anticipated change is “change of representation”
 - Anticipating changing the representation of data and associated functions (or just functions)
 - Again, a key notion behind abstract data types
- Ex:
 - Cartesian vs. polar coordinates; stacks as linked lists vs. arrays; packed vs. unpacked strings

36

Claim

- We less frequently change representations than we used to
 - We have significantly more knowledge about data structure design than we did 25 years ago
 - Memory is less often a problem than it was previously, since it's much less expensive
- Therefore, we should think twice about always anticipating that representations will change
 - This is important, since we can't simultaneously anticipate all changes
 - Ex: Changing the representation of null-terminated strings in Unix systems wouldn't be sensible
 - And this doesn't represent a stupid design decision

37

Other anticipated changes?

- Information hiding isn't only ADTs
- Algorithmic changes
 - (These are almost always part and parcel of ADT-based decompositions)
 - Monolithic to incremental algorithms
 - Improvements in algorithms
- Replacement of hardware sensors
 - Ex: better altitude sensors
- More?

38

Ubiquitous computing domain

- Portolano is a UW CSE project on this topic
 - Devices everywhere, handhelds, on-body devices, automated laboratories, etc.
- The set of anticipated changes is significantly different than in many other domains
 - Data is more stable than computations
 - Must accommodate diversity in communication speed, reliability, etc.
- Interesting domain for information hiding

39

Central premise I

- We can effectively anticipate changes
 - Unanticipated changes require changes to interfaces or (more commonly) simultaneous changes to multiple modules
- How accurate is this premise?
 - We have no idea
 - There is essentially no research about whether anticipated changes happen
 - Nor do we have disciplined ways to figure out how to better anticipate changes

40

The A-7 Project

- In the late 1970's, Parnas led a project to redesign the software for the A-7 flight program
 - One key aspect was the use of information hiding
- The project had successes, including a much improved specification of the system and the definition of the SCR requirements language
 - Actually, in one way it was a failure – the idea was to compare the old and new approaches, but the new one was so much better that the old one was dropped quickly, allowing few direct comparisons
- But little data about actual changes was gathered

41

Central premise II

- Changing an implementation is the best change, since it's isolated
- This may not always be true
 - Changing a local implementation may not be easy
 - Some global changes are straightforward
 - Mechanically or systematically
 - VanHilst's work showed an alternative
 - Using parameterized classes with a deferred supertype [ISOTAS, FSE, OOPSLA]
 - Nice followup work by Smaragdakis & Batory [ECOOP 98]
 - Griswold's work on information transparency

42

Central premise III

- The semantics of the module must remain unchanged when implementations are replaced
 - Specifically, the client should not care how the interface is implemented by the module
- But what captures the semantics of the module?
 - The signature of the interface? Performance? What else?

43

Central premise IV

- One implementation can satisfy multiple clients
 - Different clients of the same interface that need different implementations would be counter to the principle of information hiding
 - Clients should not care about implementations, as long as they satisfy the interface
 - Kiczales' work on open implementations (the predecessor of aspect-oriented programming)

44

Information hiding reprise

- It's probably the most important design technique we know
- And it's broadly useful
- It raised consciousness about change
- But one needs to evaluate the premises in specific situations to determine the actual benefits (well, the actual potential benefits)

45

Information Hiding and OO

- Are these the same? Not really
 - OO classes are chosen based on the domain of the problem (in most OO analysis approaches)
 - Not necessarily based on change
- But they are obviously related (separating interface from implementation, e.g.)
- What is the relationship between sub- and super-classes?

46

Layering [Parnas 79]

- A focus on information hiding modules isn't enough
- One may also consider abstract machines
 - In support of program families
 - Systems that have "so much in common that it pays to study their common aspects before looking at the aspects that differentiate them"
- Still focusing on anticipated change

47

The uses relation

- A program A uses a program B if the correctness of A depends on the presence of a correct version of B
- Requires specification and implementation of A and the specification of B
- Again, what is the "specification"? The interface? Implied or informal semantics?
 - Can uses be mechanically computed?

48

uses vs. invokes

- These relations often but do not always coincide
- Invocation without use: name service with cached hints
- Use without invocation: examples?

```
ipAddr := cache(hostName);  
if wrong(ipAddr,hostName) then  
    ipAddr := lookup(hostName)  
endif
```

49

Parnas' observation

- A non-hierarchical uses relation makes it difficult to produce useful subsets of a system
 - It also makes testing difficult
 - (What about upcalls?)
- So, it is important to design the *uses* relation

50

Criteria for *uses* (A, B)

- A is essentially simpler because it uses B
- B is not substantially more complex because it does not use A
- There is a useful subset containing B but not A
- There is no useful subset containing A but not B

51


Layering in THE (Dijkstra's layered OS)

- OK, those of you who took OS
- How was layering used, and how does it relate to this work?
- (For thinking about off-line, or for email discussion)

52

Modules and layers interact?

- Information hiding modules and layers are distinct concepts
- How and where do they overlap in a system?



53

Language support

- We have lots of language support for information hiding modules
 - C++ classes, Ada packages, etc.
- We have essentially no language support for layering
 - Operating systems provide support, primarily for reasons of protection, not abstraction
 - Big performance cost to pay for “just” abstraction

54

Implicit invocation: event-based design

- Implicit invocation
 - Essentially, event-based design

55

Implicit invocation

- Components announce events that other components can choose to respond to
- Components register interest in those events that they want to respond to
- In implicit invocation, the `invokes` relation is the inverse of the `names` relation
- Invocation does not require ability to name
- The central goal is to ease change: the components invoked can be changed without modifying the announcing component

56

Old II mechanisms

- Field [Reiss], DEC FUSE, HP Softbench, etc.
 - Components announce events as ASCII messages
 - Components register interest using regular expressions
 - Centralized multicast message server
- Smalltalk's Model-View-Controller
 - Registering with objects
 - Separating UI views from internal models
 - May request permission to change

57

New II mechanisms: or extensive uses of them

- JDK's
 - Different versions have somewhat different event models
- Java beans, Swing, ...
- CORBA and COM

58

Objective

- Most of you are at least comfortable with using events
 - Probably primarily in the context of existing components and frameworks
- Several issues to cover
 - Thinking of implicit invocation as more than "just" events
 - Identifying some concrete software engineering reasons to use it
 - Identifying some limitations

59

Not just indirection

- There is often confusion between implicit invocation and indirect invocation
 - Calling a virtual function is a good example of indirect invocation
 - The calling function doesn't know the precise callee, but it knows it is there and that there is only one
 - Not true in general in implicit invocation
- An announcing component should not use (in the Parnas sense) any responding components
 - This is extremely difficult to define precisely
 - Roughly, the post-condition of the announcing component should not depend on any computation of the implicitly invoked components

60

Mediators


- One style of using implicit invocation is the use of mediators [Sullivan & Notkin]
 - Also the “mediator” design pattern
- This approach combines events with entity-relationship designs
- The intent is to ease the development and evolution of integrated systems
 - Manage the coupling and isolate behavioral relationships between components

61

Experience

- A radiation treatment planning (RTP) system (Prism) was designed and built using this technique
 - By a radiation oncologist [Kalet]
 - A third generation RTP system
 - In clinical use at UW and several other major research hospitals
 - <http://www.radonc.washington.edu/physics/prism/>
 - See the screenshots on next slides

62



63



64



Example problem

- Given two sets S1 and S2, how do you ensure that they are always consistent?
 - You can independently add or delete elements from either set
- Ideas?

66


A standard solution

- Encapsulate both S1 and S2 in a component that
 - Exports insert-S1, insert-S2, delete-S1, delete-S2 methods
 - Implements these new methods to ensure consistency
- How effective is this solution?

67

Example solution: mediators

- Each set allows insertion and deletion of elements
- Each set also announces an “inserted” event and a “removed” event when the associated operation is performed
- A separate component, a mediator M, registers interest in the events from both sets
- Ex: If an element is inserted into S1, then M receives an “inserted” event; it then can invoke the insert method on S2



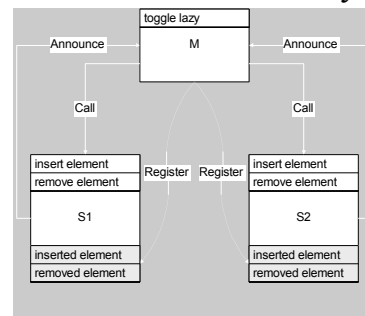
68

Mediator issues

- Must avoid circularity
- Events are first-class elements in interfaces
 - “interface” and “out-erface”
- Makes many changes easier
 - lazy equivalence
 - allow size of the sets to be changed directly
 - ...

69


Mediator: with lazy update



- Only maintain equivalence of sets when a bit is set
- When the bit is moved from off to on, re-establish consistency

70

Mediators: lazy and count



- Track an integer with the count of the set elements
- With or without the lazy modification

71

Assessment

- For some classes of systems and changes, mediator-based designs seem attractive
- Lots of outstanding issues
 - Circularities in relations
 - Ordering of mediators
 - Distributed and concurrent variants
 - Reasoning (even informally) about systems built with implicit invocation
 - Even “just” debugging

72