CSE 331
Software Design and Implementation

Lecture 4
Specifications

Zach Tatlock / Winter 2016

Administrivia

Next assignments posted tonight

- HW2: Written problems on loops

Readings posted as well!

- Quizzes coming soon ☺

2 Goals of Software System Building

- Building the *right system*
  - Does the program meet the user's needs?
  - Determining this is usually called *validation*

- Building the *system right*
  - Does the program meet the specification?
  - Determining this is usually called *verification*

- CSE 331: the second goal is the focus – creating a correctly functioning artifact
  - Surprisingly hard to specify, design, implement, test, and debug even simple programs

Where we are

- We've started to see how to reason about code
- We'll build on those skills in many places:
  - *Specification*: What are we supposed to build?
  - *Design*: How do we decompose the job into manageable pieces? Which designs are "better"?
  - *Implementation*: Building code that meets the specification
  - *Testing*: Systematically finding problems
  - *Debugging*: Systematically fixing problems
  - *Maintenance*: How does the artifact adapt over time?
  - *Documentation*: What do we need to know to do these things? How/where do we write that down?

The challenge of scaling software

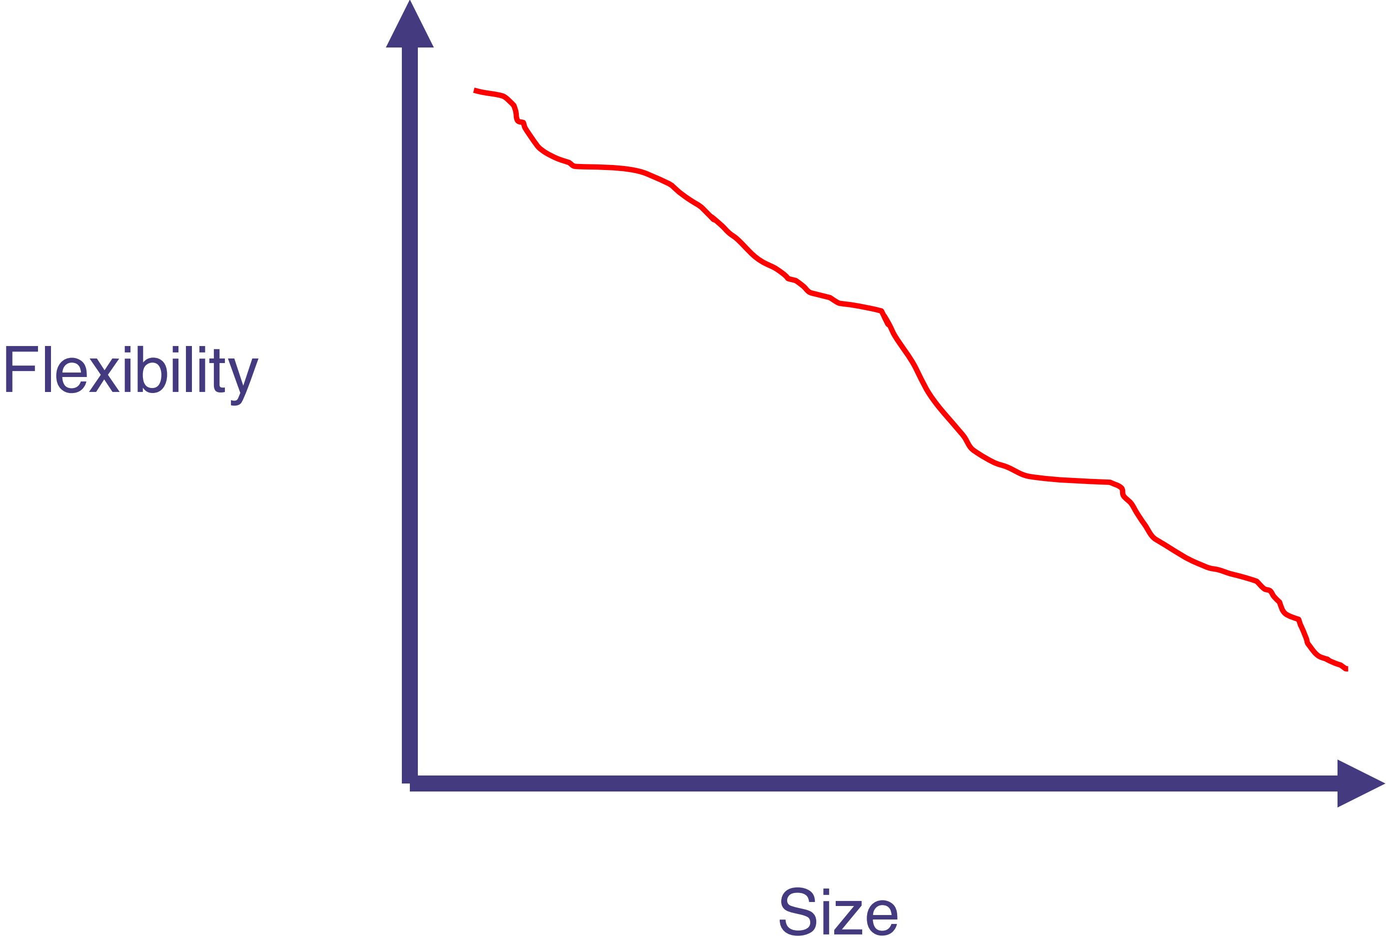

The challenge of scaling software

- Small programs are simple and malleable
  - Easy to write
  - Easy to change

- Big programs are (often) complex and inflexible
  - Hard to write
  - Hard to change

- Why does this happen?
  - Because *interactions* become unmanageable

- How do we keep things simple and malleable?

THE ARCHITECTURE OF COMPLEXITY

HERBERT A. SIMON*

Professor of Administration, Carnegie Institute of Technology

(*Read April 26, 1962*)

A NUMBER of proposals have been advanced in recent years for the development of "general systems theory" which, abstracting from properties peculiar to physical, biological, or social systems, would be applicable to all of them. [1] We might well feel that, while the goal is laudable, systems of such diverse kinds could hardly be expected to have any nontrivial properties in common. Metaphor and analogy can be helpful, or they can be misleading. All depends on whether the similarities the metaphor captures are significant or superficial.

It may not be entirely vain, however, to search for common properties among diverse kinds of complex systems. The ideas that go by the name of cybernetics constitute, if not a theory, at least a point of view that has been proving fruitful over a wide range of applications.[2] It has been useful to look at the behavior of adaptive systems in terms of the concepts of feedback and homeostasis, and to analyze adaptiveness in terms of the theory of selective information.[3] The ideas of feedback and information provide a frame of reference for viewing a wide range of situations, just as do the ideas of evolution, or relativism, of axiomatic method, and of operationalism.

In this essay I should like to report on some things we have been learning about particular kinds of complex systems encountered in the behavioral sciences. The developments I shall discuss arose in the context of specific phenomena, but the theoretical formulations themselves make little reference to details of structure. Instead they refer primarily to the complexity of the systems under view without specifying the exact content of that complexity. Because of their abstractness, the theories may have relevance—application would be too strong a term— to other kinds of complex systems that are observed in the social, biological, and physical sciences.

In recounting these developments, I shall avoid technical detail, which can generally be found elsewhere. I shall describe each theory in the particular context in which it arose. Then, I shall cite some examples of complex systems, from areas of science other than the initial application, to which the theoretical framework appears relevant. In doing so, I shall make reference to areas of knowledge where I am not expert—perhaps not even literate. I feel quite comfortable in doing so before the members of this society, representing as it does the whole span of the scientific and scholarly endeavor. Collectively you will have little difficulty, I am sure, in distinguishing instances based on idle fancy or sheer ignorance from instances that cast some light on the ways in which complexity exhibits itself wherever it is found in nature. I shall leave to you the final judgment of relevance in your respective fields.

I shall not undertake a formal definition of

---

* The ideas in this paper have been the topic of many conversations with my colleague, Allen Newell. George W. Corner suggested important improvements in biological content as well as editorial form. I am also indebted, for valuable comments on the manuscript, to Richard H. Meier, John R. Platt, and Warren Weaver. Some of the conjectures about the nearly decomposable structure of the nucleus-atom-molecule hierarchy were checked against the available quantitative data by Andrew Schoene and William Wise. My work in this area has been supported by a Ford Foundation grant for research in organizations and a Carnegie Corporation grant for research on cognitive processes. To all of the above, my warm thanks, and the usual absolution.

[1] See especially the yearbooks of the Society for General Systems Research. Prominent among the exponents of general systems theory are L. von Bertalanffy, K. Boulding, R. W. Gerard, and J. G. Miller. For a more skeptical view—perhaps too skeptical in the light of the present discussion—see H. A. Simon and A. Newell, Models: their uses and limitations, in L. D. White, ed., *The state of the social sciences*, 66-83, Chicago, Univ. of Chicago Press, 1956.

[2] N. Wiener, *Cybernetics*, New York, Wiley, 1948. For an imaginative forerunner, see A. J. Lotka, *Elements of mathematical biology*, New York, Dover Publications, 1951, first published in 1924 as *Elements of physical biology*.

[3] C. Shannon and W. Weaver, *The mathematical theory of communication*, Urbana, Univ. of Illinois Press, 1949; W. R. Ashby, *Design for a brain*, New York, Wiley, 1952.

A discipline of modularity

- Two ways to view a program:
  - The implementer's view (how to build it)
  - The client's view (how to use it)

- It helps to apply these views to program parts:
  - While implementing one part, consider yourself a client of any other parts it depends on
  - Try *not* to look at those other parts through an implementer's eyes
  - Helps dampen interactions between parts

- Formalized through the idea of a *specification*

A specification is a contract

- A set of requirements agreed to by the user and the manufacturer of the product
  - Describes their expectations of each other

- Facilitates simplicity via *two-way* isolation
  - Isolate client from implementation details
  - Isolate implementer from how the part is used
  - Discourages implicit, unwritten expectations

- Facilitates change
  - Reduces the “Medusa effect”: the specification, rather than the code, gets “turned to stone” by client dependencies

Isn't the interface sufficient?

The interface defines the boundary between implementers and users:

```
public class List<E> {
    public E get(int x) { return null; }
    public void set(int x, E y){}
    public void add(E) {}
    public void add(int, E){}
    …
    public static <T> boolean isSub(List<T>, List<T>){
        return false;
    }
}
```

Interface provides the *syntax and types*

But nothing about the *behavior and effects*

– *Provides too little information to clients*

Note: Code above is right concept but is not (completely) legal Java

– *Parameters need names; no static interface methods before Java 8*

Why not just read code?

```
static <T> boolean sub(List<T> src, List<T> part) {
     int part_index = 0;
     for (T elt : src) {
          if (elt.equals(part.get(part_index))) {
               part_index++;
               if (part_index == part.size()) {
                    return true;
               }
          } else {
               part_index = 0;
          }
     }
     return false;
 }
```

Why are you better off with a specification?

Code is complicated

- Code gives more detail than needed by client
- Understanding or even reading every line of code is an excessive burden
  - Suppose you had to read source code of Java libraries to use them
  - Same applies to developers of different parts of the libraries
- Client cares only about *what* the code does, not *how* it does it

Code is ambiguous

- Code seems unambiguous and concrete
  - But which details of code's behavior are essential, and which are incidental?

- Code invariably gets rewritten
  - Client needs to know what they can rely on
    - What properties will be maintained over time?
    - What properties might be changed by future optimization, improved algorithms, or bug fixes?
  - Implementer needs to know what features the client depends on, and which can be changed

Comments are essential

Most comments convey only an informal, general idea of what that the code does:

```
// This method checks if "part" appears as a
// sub-sequence in "src"
static <T> boolean sub(List<T> src, List<T> part){
  ...
}
```

Problem: ambiguity remains

- What if `src` and `part` are both empty lists?
- When does the function return `true`?

From vague comments to specifications

- Roles of a specification:
  - Client agrees to rely *only* on information in the description in their use of the part
  - Implementer of the part promises to support everything in the description
    - Otherwise is perfectly at liberty
- Sadly, much code lacks a specification
  - Clients often work out what a method/class does in ambiguous cases by running it and depending on the results
  - Leads to bugs and programs with unclear dependencies, reducing simplicity and flexibility

Recall the sublist example

```
static <T> boolean sub(List<T> src, List<T> part) {
    int part_index = 0;
    for (T elt : src) {
        if (elt.equals(part.get(part_index))) {
            part_index++;
            if (part_index == part.size()) {
                return true;
            }
        } else {
            part_index = 0;
        }
    }
    return false;
}
```

A more careful description of **`sub`**

```
// Check whether “part” appears as a sub-sequence in “src”
```

needs to be given some caveats (why?):

```
// * src and part cannot be null
// * If src is empty list, always returns false
// * Results may be unexpected if partial matches
//   can happen right before a real match; e.g.,
//   list (1,2,1,3) will not be identified as a
//   sub sequence of (1,2,1,2,1,3).
```

or replaced with a more detailed description:

```
// This method scans the “src” list from beginning
// to end, building up a match for “part”, and
// resetting that match every time that...
```

A better approach

It's better to simplify than to describe complexity!

Complicated description suggests poor design

- Rewrite **sub** to be more sensible, and easier to describe

```
// returns true iff possibly empty sequences A, B where
//   src = A : part : B
// where ":" is sequence concatenation
static <T> boolean sub(List<T> src, List<T> part) {
```

- Mathematical flavor not always necessary, but often helps avoid ambiguity
- "Declarative" style is important: avoids reciting or depending on operational/implementation details

The benefits of spec #1

- The discipline of writing specifications changes the incentive structure of coding
  - Rewards code that is easy to describe and understand
  - Punishes code that is hard to describe and understand
    - Even if it is shorter or easier to write
- If you find yourself writing complicated specifications, it is an incentive to redesign
  - In **`sub`**, code that does exactly the right thing may be slightly slower than a hack that assumes no partial matches before true matches, but cost of forcing client to understand the details is too high

Writing specifications with Javadoc

- Javadoc
  - Sometimes can be daunting; get used to using it

- Javadoc convention for writing specifications
  - Method signature
  - Text description of method
  - **`@param`**: description of what gets passed in
  - **`@return`**: description of what gets returned
  - **`@throws`**: exceptions that may occur

Example: Javadoc for `String.contains`

```
public boolean contains(CharSequence s)

Returns true if and only if this string contains the
  specified sequence of char values.

Parameters:

  s- the sequence to search for

Returns:

  true if this string contains s, false otherwise

Throws:

  NullPointerException - if s is null

Since:

  1.5
```

CSE 331 specifications

- The *precondition*: constraints that hold before the method is called (if not, all bets are off)
  - **`@requires`**: spells out any obligations on client
- The *postcondition*: constraints that hold after the method is called (if the precondition held)
  - **`@modifies`**: lists objects that may be affected by method; any object not listed is guaranteed to be untouched
  - **`@throws`**: lists possible exceptions and conditions under which they are thrown (Javadoc uses this too)
  - **`@effects`**: gives guarantees on final state of modified objects
  - **`@return`**: describes return value (Javadoc uses this too)

Example 1

static <T> int change(List<T> lst, T oldelt, T newelt)

requires	lst, oldelt, and newelt are non-null. oldelt occurs in lst.
modifies	lst
effects	change the first occurrence of oldelt in lst to newelt & makes no other changes to lst
returns	the position of the element in lst that was oldelt and is now newelt

---

```
static <T> int change(List<T> lst,
                      T oldelt, T newelt) {
   int i = 0;
   for (T curr : lst) {
      if (curr == oldelt) {
         lst.set(newelt, i);
         return i;
      }
      i = i + 1;
   }
   return -1;
}
```

Example 2

static List<Integer> zipSum(List<Integer> lst1, List<Integer> lst2)

requires	lst1 and lst2 are non-null. lst1 and lst2 are the same size.
modifies	none
effects	none
returns	a list of same size where the ith element is the sum of the ith elements of lst1 and lst2

```
static List<Integer> zipSum(List<Integer> lst1
                            List<Integer> lst2) {
  List<Integer> res = new ArrayList<Integer>();
  for(int i = 0; i < lst1.size(); i++) {
    res.add(lst1.get(i) + lst2.get(i));
  }
  return res;
}
```

Example 3

static void listAdd(List<Integer> lst1, List<Integer> lst2)

requires	lst1 and lst2 are non-null. lst1 and lst2 are the same size.
modifies	lst1
effects	ith element of lst2 is added to the ith element of lst1
returns	none

```
static void listAdd(List<Integer> lst1,
                    List<Integer> lst2) {
  for(int i = 0; i < lst1.size(); i++) {
     lst1.set(i, lst1.get(i) + lst2.get(i));
  }
}
```

Example 4 (Watch out for bugs!)

static void uniquify(List<Integer> lst)

requires ???

???

modifies ???

effects ???

returns ???

```
static void uniquify(List<Integer> lst) {
   for (int i=0; i < lst.size()-1; i++)
        if (lst.get(i) == lst.get(i+1))
             lst.remove(i);
}
```

Should requires clause be checked?

If the client calls a method without meeting the precondition, the code is free to do *anything*

- Including pass corrupted data back
- It is polite, nevertheless, to *fail fast*: to provide an immediate error, rather than permitting mysterious bad behavior

Preconditions are common in “helper” methods/classes

- In public libraries, it’s friendlier to deal with all possible input
- *Example: binary search would normally impose a pre-condition rather than simply failing if list is not sorted. Why?*

Rule of thumb: Check if cheap to do so

- *Example: list has to be non-null → check*
- *Example: list has to be sorted → skip*

Satisfaction of a specification

Let M be an implementation and S a specification

M satisfies S if and only if

- Every behavior of M is permitted by S
- "The behavior of M is a subset of S"

The statement "M is correct" is meaningless!

- Though often made!

If M does not satisfy S, either (or both!) could be "wrong"

- *"One person's feature is another person's bug."*
- Usually better to change the program than the spec

The benefits of specs #2

Specification means that client doesn't need to look at implementation

- So the code may not even exist yet!

Write specifications first, make sure system will fit together, and then assign separate implementers to different modules

- Allows teamwork and parallel development
- Also helps with testing (future topic)

Comparing specifications

Occasionally, we need to compare different versions of a specification (*Why?*)

- For that, talk about *weaker* and *stronger* specifications

A weaker specification gives greater freedom to the implementer

- If specification S_1 is weaker than S_2, then for any implementation M,
  - M satisfies S_2 => M satisfies S_1
  - but the opposite implication does not hold in general

Given two specifications, they may be *incomparable*

- Neither is weaker/stronger than the other
- *Some* implementations might still satisfy them both

Why compare specifications?

We wish to relate procedures to specifications

- Does the procedure satisfy the specification?
- Has the implementer succeeded?

We wish to compare specifications to one another

- Which specification (if either) is stronger?
- A procedure satisfying a stronger specification can be used anywhere that a weaker specification is required
  - Substitutability principle
  - Accept at least as many inputs
  - Produce no more outputs

Example 1

```
int find(int[] a, int value) {
    for (int i=0; i<a.length; i++) {
        if (a[i]==value)
          return i;
    }
    return -1;
}
```

Specification A

- requires: value occurs in **`a`**
- returns: **`i`** such that **`a[i]`** = **`value`**

Specification B

- requires: value occurs in **`a`**
- returns: *smallest* **`i`** such that **`a[i]`** = **`value`**

Example 2

```
int find(int[] a, int value) {
  for (int i=0; i<a.length; i++) {
      if (a[i]==value)
        return i;
  }
  return -1;
}
```

Specification A

- requires: value occurs in **a**
- returns: **i** such that **a[i]** = **value**

Specification C

- returns: **i** such that **a[i]** = **value**, or **-1** if value is not in **a**

Stronger and weaker specifications

A stronger specification is

- Harder to satisfy (more constraints on the implementation)
- Easier to use (more guarantees, more predictable, client can make more assumptions)

A weaker specification is

- Easier to satisfy (easier to implement, more implementations satisfy it)
- Harder to use (makes fewer guarantees)

Strengthening a specification

Strengthen a specification by:

- Promising more – any or all of:
  - Effects clause harder to satisfy
  - Returns clause harder to satisfy
  - Fewer objects in modifies clause
  - More specific exceptions (subclasses)
- Asking less of client
  - Requires clause easier to satisfy

Weaken a specification by:

- (Opposite of everything above)

“Strange” case: @throws

[Prior versions of course, including old exams, were clumsy/wrong about this]

Compare:

S1:

```
@throws FooException if x<0
@return x+3
```

S2:

```
@return x+3
```

- These are *incomparable* because they promise different, incomparable things when x<0
- Both are *stronger* than @requires x>=0; @return x+3

Which is better?

Stronger does not always mean better!

Weaker does not always mean better!

Strength of specification trades off:

- Usefulness to client
- Ease of simple, efficient, correct implementation
- Promotion of reuse and modularity
- Clarity of specification itself

"It depends"

More formal stronger/weaker

A specification is a logical formula

- S1 stronger than S2 if S1 implies S2
- From implication all things follow:
  - Example: S1 stronger if requires is weaker
  - Example: S1 stronger if returns is stronger

As in all logic (cf. CSE311), two rigorous ways to check implication

- Convert entire specifications to logical formulas and use logic rules to check implication (e.g., $P1 \wedge P2 \Rightarrow P2$)
- Check every *behavior* described by stronger also described by the other
  - CSE311: truth tables
  - CSE331: *transition relations*

Transition relations

There is a program state before a method call and after

- All memory, values of all parameters/result, whether exception happened, etc.

A specification "means" a set of pairs of program states

- The legal pre/post-states
- This is the transition relation defined by the spec
  - Could be infinite
  - Could be multiple legal outputs for same input

Stronger specification means the transition relation is a subset

Note: Transition relations often are infinite in size