2. Overview and Background

This is a course for computer system designers and builders, and for people who want to really understand how systems work, especially concurrent, distributed, and fault-tolerant systems.

The course teaches you

    how to write precise specifications for any kind of computer system,

    what it means for code to satisfy a specification, and

    how to prove that it does.

It also shows you how to use the same methods less formally, and gives you some suggestions for deciding how much formality is appropriate (less formality means less work, and often a more understandable spec, but also more chance to overlook an important detail).

The course also teaches you a lot about the topics in computer systems that we think are the most important: persistent storage, concurrency, naming, networks, distributed systems, transactions, fault tolerance, and caching. The emphasis is on

    careful specifications of subtle and sometimes complicated things,

    the important ideas behind good code, and

    how to understand what makes them actually work.

We spend most of our time on specific topics, but we use the general techniques throughout. We emphasize the ideas that different kinds of computer system have in common, even when they have different names.

The course uses a formal language called Spec for writing specs and code; you can think of it as a very high level programming language. There is a good deal of written introductory material on Spec (explanations and finger exercises) as well as a reference manual and a formal semantics. We introduce Spec ideas in class as we use them, but we do not devote class time to teaching Spec per se; we expect you to learn it on your own from the handouts. The one to concentrate on is handout 3, which has an informal introduction to the main features and lots of examples. Section 9 of handout 4, the reference manual, should also be useful. The rest of the reference manual is for reference, not for learning. Don't overlook the one page summary at the end of handout 3.

Because we write specs and do proofs, you need to know something about logic. Since many people don't, there is a concise treatment of the logic you will need at the end of this handout.

This is not a course in computer architecture, networks, operating systems, or databases. We will not talk in detail about how to code pipelines, memory interconnects, multiprocessors, routers, data link protocols, network management, virtual memory, scheduling, resource allocation, SQL, relational integrity, or TP monitors, although we will deal with many of the ideas that underlie these mechanisms.

Topics

General

Specifications as state machines.

The Spec language for describing state machines (writing specs and code).

What it means to implement a spec.

Using abstraction functions and invariants to prove that a program implements a spec.

What it means to have a crash.

What every system builder needs to know about performance.

Specific

Disks and file systems.

Practical concurrency using mutexes (locks) and condition variables; deadlock.

Hard concurrency (without locking): models, specs, proofs, and examples.

Transactions: simple, cached, concurrent, distributed.

Naming: principles, specs, and examples.

Distributed systems: communication, fault-tolerance, and autonomy.

Networking: links, switches, reliable messages and connections.

Remote procedure call and network objects.

Fault-tolerance, availability, consensus and replication.

Caching and distributed shared memory.

Previous editions of the course have also covered security (authentication, authorization, encryption, trust) and system management, but this year we are omitting these topics in order to spend more time on concurrency and semantics and to leave room for project presentations.

Prerequisites

There are no formal prerequisites for the course. However, we assume some knowledge both of computer systems and of mathematics. If you have taken 6.033 and 6.042, you should be in good shape. If you are missing some of this knowledge you can pick it up as we go, but if you are missing a lot of it you can expect to have serious trouble. It's also important to have a certain amount of maturity: enough experience with systems and mathematics to feel comfortable with the basic notions and to have some reliable intuition.

If you know the meaning of the following words, you have the necessary background. If a lot of them are unfamiliar, this course is probably not for you.

Systems

Cache, virtual memory, page table, pipeline

Process, scheduler, address space, priority

Thread, mutual exclusion (locking), semaphore, producer-consumer, deadlock

Transaction, commit, availability, relational data base, query, join

File system, directory, path name, striping, RAID

LAN, switch, routing, connection, flow control, congestion

Capability, access control list, principal (subject)

If you have not already studied Lampson's paper on hints for system design, you should do so as background for this course. It is Butler Lampson, Hints for computer system design, *Proceedings of the Ninth ACM Symposium on Operating Systems Principles*, October 1983, pp 33-48. There is a pointer to it on the course Web page.

Programming

Invariant, precondition, weakest precondition, fixed point
Procedure, recursion, stack
Data type, sub-type, type-checking, abstraction, representation
Object, method, inheritance
Data structures: list, hash table, binary search, B-tree, graph

Mathematics

Function, relation, set, transitive closure
Logic: proof, induction, de Morgan's laws, implication, predicate, quantifier
Probability: independent events, sampling, Poisson distribution
State machine, context-free grammar
Computational complexity, unsolvable problem

If you haven't been exposed to formal logic, you should study the summary at the end of this handout.

References

These are places to look when you want more information about some topic covered or alluded to in the course, or when you want to follow current research. You might also wish to consult Prof. Saltzer's bibliography for 6.033, which you can find on the course web page.

Books

Some of these are fat books better suited for reference than for reading cover to cover, especially Cormen, Leiserson, and Rivest, Jain, Mullender, Hennessy and Patterson, and Gray and Reuter. But the last two are pretty easy to read in spite of their encyclopedic character.

Systems programming: Greg Nelson, ed., *Systems Programming with Modula-3*, Prentice-Hall, 1991. Describes the language, which has all the useful features of C++ but is much simpler and less error-prone, and also shows how to use it for concurrency (a version of chapter 4 is a handout in this course), an efficiently customizable I/O streams package, and a window system.

Performance: Jon Bentley, *Writing Efficient Programs*, Prentice-Hall, 1982. Short, concrete, and practical. Raj Jain, *The Art of Computer Systems Performance Analysis*, Wiley, 1991. Tells you much more than you need to know about this subject, but does have a lot of realistic examples.

Algorithms and data structures: Robert Sedgwick, *Algorithms*, Addison-Wesley, 1983. Short, and usually tells you what you need to know. Tom Cormen, Charles Leiserson, and Ron Rivest, *Introduction to Algorithms*, McGraw-Hill, 1989. Comprehensive, and sometimes valuable for that reason, but usually tells you a lot more than you need to know.

Distributed algorithms: Nancy Lynch, *Distributed Algorithms*, Morgan Kaufmann, 1996. The bible for distributed algorithms. Comprehensive, but a much more formal treatment than in this course. The topic is algorithms, not systems.

Computer architecture: John Hennessy and David Patterson, *Computer Architecture: A Quantitative Approach*, 2nd edition, Morgan Kaufmann, 1995. The bible for computer architecture. The second edition has lots of interesting new material, especially on multiprocessor memory systems and interconnection networks. There's also a good appendix on computer arithmetic; it's useful to know where to find this information, though it has nothing to do with this course.

Transactions, data bases, and fault-tolerance: Jim Gray and Andreas Reuter, *Transaction Processing: Concepts and Techniques*, Morgan Kaufmann, 1993. The bible for transaction processing, with much good material on data bases as well; it includes a lot of practical information that doesn't appear elsewhere in the literature.

Networks: Radia Perlman, *Interconnections: Bridges and Routers*, Addison-Wesley, 1992. Not exactly the bible for networking, but tells you nearly everything you might want to know about how packets are actually switched in computer networks.

Distributed systems: Sape Mullender, ed., *Distributed Systems*, 2nd ed., Addison-Wesley, 1993. A compendium by many authors that covers the field fairly well. Some chapters are much more theoretical than this course. Chapters 10 and 11 are handouts in this course. Chapters 1, 2, 8, and 12 are also recommended. Chapters 16 and 17 are the best you can do to learn about real-time computing; unfortunately, that is not saying much.

User interfaces: Alan Cooper, *About Face*, IDG Books, 1995. Principles, lots of examples, and opinionated advice, much of it good, from the original designer of Visual Basic.

Journals

You can find all of these in the LCS reading room. The cryptic strings in brackets are call numbers there. You can also find the last few years of the ACM publications in the ACM digital library at www.acm.org.

For the current literature, the best sources are the proceedings of the following conferences. 'Sig' is short for "Special Interest Group", a subdivision of the ACM that deals with one field of computing. The relevant ones for systems are SigArch for computer architecture, SigPlan for programming languages, SigOps for operating systems, SigComm for communications, SigMod for data bases, and SigMetrics for performance measurement and analysis.

Symposium on Operating Systems Principles (SOSP; published as special issues of ACM SigOps *Operating Systems Review*; fall of odd-numbered years) [P4.35.06]

Operating Systems Design and Implementation (OSDI; Usenix Association, now published as special issues of ACM *SigOps Review*; fall of even-numbered years, except spring 1999 instead of fall 1998) [P4.35.U71]

Architectural Support for Programming Languages and Operating Systems (ASPLOS; published as special issues of ACM SigOps *Operating Systems Review*, SigArch *Computer Architecture News*, or *SigPlan Notices*; fall of even-numbered years) [P6.29.A7]

Applications, Technologies, Architecture, and Protocols for Computer Communication, (SigComm conference; published as special issues of ACM SigComm *Computer Communication Review*; annual) [P6.24.D31]

Principles of Distributed Computing (PODC; ACM; annual) [P4.32.D57]

Very Large Data Bases (VLDB; Morgan Kaufmann; annual) [P4.33.V4]

International Symposium on Computer Architecture (ISCA; published as special issues of ACM SigArch *Computer Architecture News*; annual) [P6.20.C6]

Less up to date, but more selective, are the journals. Often papers in these journals are revised versions of papers from the conferences listed above.

ACM Transactions on Computer Systems

ACM Transactions on Database Systems

ACM Transactions on Programming Languages and Systems

There are often good survey articles in the less technical IEEE journals:

IEEE *Computer*, *Networks*, *Communication*, *Software*

The Internet Requests for Comments (RFC's) can be reached from

    http://www.cis.ohio-state.edu/hypertext/information/rfc.html

Rudiments of logic

Propositional logic

The basic type is `Bool`, which contains two elements `true` and `false`. Expressions in these operators (and the other ones introduced later) are called 'propositions'.

Basic operators. These are ∧ (and), ∨ (or), and ~ (not).[1] The meaning of these operators can be conveniently given by a 'truth table' which lists the value of `a op b` for each possible combination of values of `a` and `b` (the operators on the right are discussed later) along with some popular names for certain expressions and their operands.

		negation	conjunction	disjunction	equality		implication
		not	and	or			implies
a	b	~a	a ∧ b	a ∨ b	a = b	a ≠ b	a ⟹ b
T	T	F	T	T	T	F	T
T	F		F	T	F	T	F
F	T	T	F	T	F	T	T
F	F		F	F	T	F	T
name of a			conjunct	disjunct			antecedent
name of b			conjunct	disjunct			consequent

Note: In Spec we write `==>` instead of the ⟹ that mathematicians use for implication. Logicians write ⊃ for implication, which looks different but is shaped like the > part of ⟹.

Since the table has only four rows, there are only 16 Boolean operators, one for each possible arrangement of T and F in a column. Most of the ones not listed don't have common names, though 'not and' is called 'nand' and 'not or' is called 'nor' by logic designers.

The ∧ and ∨ operators are
    commutative and
    associative and
    distribute over each other.
That is, they are just like * (times) and + (plus) on integers, except that + doesn't distribute over *:
    a + (b * c) ≠ (a + b) * (a + c)
but ∨ does distribute over ∧:
    a ∨ (b ∧ c) = (a ∨ b) ∧ (a ∨ c)
An operator that distributes over ∧ is called 'conjunctive'; one that distributes over ∨ is called 'disjunctive'. Both ∧ and ∨ are both conjunctive and disjunctive. This takes some getting used to.

The relation between these operators and ~ is given by DeMorgan's laws (sometimes called the "bubble rule" by logic designers), which say that you can push ~ inside ∧ or ∨ (or pull it out) by flipping from one to the other:
    ~ (a ∧ b)   =   ~a ∨ ~b

___
[1] It's possible to write all three in terms of the single operator 'nor' or 'nand', but our goal is clarity, not minimality.

```
    ~ (a ∨ b)  =  ~a ∧ ~b
```

To put a complex expression into " disjunctive normal form" replace terms in = and ⇒ with their equivalents in ∧, ∨, and ~ (given below), use DeMorgan's laws to push all the ~'s in past ∧ and ∨ so that they apply to variables, and then distribute ∧ over ∨ so that the result looks like

$$(a_1 \land \sim a_2 \land ...) \lor (\sim b_1 \land b_2 \land ...) \lor ...$$

The disjunctive normal form is unique (up to ordering, since ∧ and ∨ are commutative). Of course, you can also distribute ∨ over ∧ to get a unique "conjunctive normal form".

If you want to find out whether two expressions are equal, one way is to put them both into disjunctive (or conjunctive) normal form, sort the terms, and see whether they are identical. Another way is to list all the possible values of the variables (if there are n variables, there are 2^n of them) and tabulate the values of the expressions for each of them; we saw this 'truth table' for some two-variable expressions above.

Because `Bool` is the result type of relations like =, you can write expressions that mix up relations with other operators in ways that are impossible for any other type. Notably

```
(a = b)  =  ((a ∧ b) ∨ (~a ∧ ~b))
```

Some people feel that the outer = in this expression is somehow different from the inner one, and write it ≡. Experience suggests, however, that this is often a harmful distinction to make.

Implication. We can define an ordering on `Bool` with `false > true`, that is, `false` is greater than `true`. The non-strict version of this ordering is called 'implication' and written ⇒ (rather than ≥ or >= as we do with other types; logicians write it ⊃, which also looks like an ordering symbol). So `(true ⇒ false) = false` (read this as: "`true` is greater than or equal to `false`" is false) but all other combinations are `true`. The expression a ⇒ b is pronounced "a implies b", or "if a then b".[2]

There are lots of rules for manipulating expressions containing ⇒; the most useful ones are given below. If you remember that ⇒ is an ordering you'll find it easy to remember most of the rules, but if you forget the rules or get confused, you can turn the ⇒ into ∨ by the rule

```
(a ⇒ b)  =  ~a ∨ b
```

and then just use the simpler rules for ∧, ∨, and ~. So remember this even if you forget everything else.

The point of implication is that it tells you when one proposition is stronger than another, in the sense that if the first one is true, the second is also true (because if both a and a ⇒ b are true, then b must be `true` since it can't be `false`).[3] So we use implication all the time when reasoning

---

[2] It sometimes seems odd that `false` implies b regardless of what b is, but the "if ... then" form makes it clearer what is going on: if `false` is `true` you can conclude anything, but of course it isn't. A proposition that implies `false` is called 'inconsistent' because it implies anything. Obviously it's bad to think that an inconsistent proposition is true. The most likely way to get into this hole is to think that each of a collection of innocent looking propositions is true when their conjunction turns out to be inconsistent.

[3] It may also seem odd that `false > true` rather than the other way around, since `true` seems better and so should be bigger. But in fact if we want to conclude lots of things, being close to `false` is better because if `false` is true we can conclude anything, but knowing that `true` is true doesn't help at all. Strong propositions are as close to `false` as possible; this is logical brinkmanship. For example, a ∧ b is closer to `false` than a (there are more values of the variables a and b that make it `false`), and clearly we can conclude more things from it than from a alone.

from premises to conclusions. Two more ways to pronounce a ⇒ b are "a is stronger than b" and "b follows from a". The second pronunciation suggests that it's sometimes useful to write the operands in the other order, as b ⇐ a, which can also be pronounced "b is weaker than a" or "b only if a"; this should be no surprise, since we do it with other orderings.

Of course, implication has the properties we expect of an ordering:

Transitive: If a ⇒ b and b ⇒ c then a ⇒ c.[4]

Reflexive: a ⇒ a.

Anti-symmetric: If a ⇒ b and b ⇒ a then a = b.[5]

Furthermore, ~ reverses the sense of implication (this is called the 'contrapositive'):

```
(a ⇒ b)  =  (~b ⇒ ~a)
```

More generally, you can move a disjunct on the right to a conjunct on the left by negating it, or vice versa. Thus

```
(a ⇒ b ∨ c)  =  (a ∧ ~b ⇒ c)
```

As special cases in addition to the contrapositive we have

```
(a ⇒ b)  =  (a ∧ ~b ⇒ false )  =  ~ (a ∧ ~b) ∨ false = ~a ∨ b
(a ⇒ b)  =  (true  ⇒ ~a ∨ b)  =  false   ∨ ~a ∨ b = ~a ∨ b
```

since `false` and `true` are the identities for ∨ and ∧.

We say that an operator `op` is 'monotonic' in an operand if replacing that operand with a stronger (or weaker) one makes the result stronger (or weaker). Precisely, "`op` is monotonic in its first operand" means that if a ⇒ b then (a op c) ⇒ (b op c). Both ∧ and ∨ are monotonic; in fact, any operator that is conjunctive (distributes over ∧) is monotonic, because if a ⇒ b then a = (a ∧ b), so

```
a op c = (a ∧ b) op c = a op c ∧ b op c ⇒ b op c
```

If you know what a lattice is, you will find it useful to know that the set of propositions forms a lattice with ⇒ as its ordering and (remember, think of ⇒ as "greater than or equal"):

top	= `false`	
bottom	= `true`	
meet	= ∧	least upper bound, so (a ∧ b) ⇒ a and (a ∧ b) ⇒ b
join	= ∨	greatest lower bound, so a ⇒ (a ∨ b) and b ⇒ (a ∨ b)

This suggests two more expressions that are equivalent to a ⇒ b:

```
(a ⇒ b) = (a = (a ∧ b))    'and'ing a weaker term makes no difference,
                            because a ⇒ b iff a = least upper bound(a, b).

(a ⇒ b) = (b = (a ∨ b))    'or'ing a stronger term makes no difference,
                            because a ⇒ b iff b = greatest lower bound(a, b).
```

---

[4] We can also write this ((a ⇒ b) ∧ (b ⇒ c)) ⇒ (a ⇒ c).

[5] Thus (a = b) = (a ⇒ b ∧ b ⇒ a), which is why a = b is sometimes pronounced "a if and only if b" and written "a iff b".

Predicate logic

Propositions that have free variables, like $x < 3$ or $x < 3 \Rightarrow x < 5$, demand a little more machinery. You can turn such a proposition into one without a free variable by substituting some value for the variable. Thus if $P(x)$ is $x < 3$ then $P(5)$ is $5 < 3 = \text{false}$. To get rid of the free variable without substituting a value for it, you can take the 'and' or 'or' of the proposition for all the possible values of the free variable. These have special names and notation[6]:

$\forall\ x\ |\ P(x)\ =\ P(x_1)\ \wedge\ P(x_2)\ \wedge\ \ldots$ for all x, $P(x)$. In Spec,
(ALL x | P(x)) or $\wedge$: {x | P(x)}

$\exists\ x\ |\ P(x)\ =\ P(x_1)\ \vee\ P(x_2)\ \vee\ \ldots$ there exists an x such that $P(x)$. In Spec,
(EXISTS x | P(x)) or $\vee$: {x | P(x)}

Here the x_i range over all the possible values of the free variables.[7] The first is called 'universal quantification'; as you can see, it corresponds to conjunction. The second is called 'existential quantification' and corresponds to disjunction. If you remember this you can easily figure out what the quantifiers do with respect to the other operators.

In particular, DeMorgan's laws generalize to quantifiers:
~ ($\forall$ x | P(x)) = ($\exists$ x | ~P(x))
~ ($\exists$ x | P(x)) = ($\forall$ x | ~P(x))

Also, because $\wedge$ and $\vee$ are conjunctive and therefore monotonic, $\forall$ and $\exists$ are conjunctive and therefore monotonic.

It is not true that you can reverse the order of $\forall$ and $\exists$, but it's sometimes useful to know that having $\exists$ first is stronger:
$\exists$ y | $\forall$ x | P(x, y) $\Rightarrow$ $\forall$ x | $\exists$ y | P(x, y)

Intuitively this is clear: a y that works for every x can surely do the job for each particular x.

If we think of P as a relation, the consequent in this formula says that P is total (relates every x to some y). It doesn't tell us anything about how to find a y that is related to x. As computer scientists, we like to be able to compute things, so we prefer to have a function that computes y, or the set of y's, from x. This is called a 'Skolem function'; in Spec you write P.func (or P.setF for the set). P.func is total if P is total. Or, to turn this around, if we have a total function f such that $\forall$ x | P(x, f(x)), then certainly $\forall$ x | $\exists$ y | P(x, y); in fact, y = f(x) will do. Amazing.

Summary of logic

The $\wedge$ and $\vee$ operators are commutative and associative and distribute over each other.

DeMorgan's laws: ~ (a $\wedge$ b) = ~a $\vee$ ~b

---

[6] There is no agreement on what symbol should separate the $\forall$ x or $\exists$ x from the P(x). We use '|' here as Spec does, but other people use '.' or ':' or just a space, or write ($\forall$x) and ($\exists$x). Logicians traditionally write (x) and ($\exists$x).
[7] In general this might not be a countable set, so the conjunction and disjunction are written in a somewhat misleading way, but this complication won't make any difference to us.

~ (a $\vee$ b) = ~a $\wedge$ ~b

Any expression has a unique (up to ordering) disjunctive normal form in which $\vee$ combines terms in which $\vee$ combines (possibly negated) variables: $(a_1 \wedge \sim a_2 \wedge \ldots) \vee (\sim b_1 \wedge b_2 \wedge \ldots) \vee \ldots$

Implication: (a $\Rightarrow$ b) = ~a $\vee$ b

Implication is the ordering in a lattice (a partially ordered set in which every subset has a least upper and a greatest lower bound) with

top	= false		so false $\Rightarrow$ true
bottom	= true		
meet	= $\wedge$	least upper bound,	so (a $\wedge$ b) $\Rightarrow$ a
join	= $\vee$	greatest lower bound, so a $\Rightarrow$ (a $\vee$ b)	

For all x, $P(x)$:
$\forall$ x | P(x) = P(x_1) $\wedge$ P(x_2) $\wedge$ ...

There exists an x such that $P(x)$:
$\exists$ x | P(x) = P(x_1) $\vee$ P(x_2) $\vee$ ...

Index for logic