Massachusetts Institute of Technology
Department of Electrical Engineering and Computer Science

6.826 — Principles of Computer Systems

Handout 14 March 5, 1997

Practical Concurrency

We begin our study of concurrency by describing how to use it in practice; later, in handout 16, we shall study it more formally. First we explain where the concurrency in a system comes from, and discuss the main ways to express concurrency. Then we describe the difference between 'hard' and 'easy' concurrency[1]: the latter is done by locking shared data before you touch it, the former in subtle ways that are so error-prone that simple prudence requires correctness proofs. We give the rules for easy concurrency using locks, and discuss various issues that complicate the easy life: scheduling, locking granularity, and deadlocks.

Sources of concurrency

Before studying concurrency in detail, it seems useful to consider how you might get concurrency in your system. Obviously if you have a multiprocessor or a distributed system you will have concurrency, since in these systems there is more than one CPU executing instructions. Similarly, most hardware has separate parts that can change state simultaneously and independently. But suppose your system consists of a single CPU running a program. Then you can certainly arrange for concurrency by multiplexing that CPU among several tasks, but why would you want to do this? Since the CPU can only execute one instruction at a time, it isn't entirely obvious that there is any advantage to concurrency. Why not get one task done before moving on to the next one?

There are only two possible reasons:

> A task might have to wait for something else to complete before it can proceed, for instance for a disk read. But this means that there is some concurrent task that is going to complete, in the example an I/O device. So we have concurrency in any system that has I/O, even when there is only one CPU.
>
> Something else might have to wait for the result of one task but not for the rest of the computation, for example a human user. But this means that there is some concurrent task that is waiting, in the example the user. Again we have concurrency in any system that has I/O.

[1] I am indebted to Greg Nelson for this taxonomy, and for the object and set example of deadlock avoidance.

In the first case one task must wait for I/O, and we can get more work done by running another task on the CPU, rather than letting it idle during the wait. Thus the concurrency of the I/O system leads to concurrency on the CPU. If the I/O wait is explicit in the program, the programmer can know when other tasks might run; this is often called a ‘non-preemptive’ system, because it has sequential semantics except when the program explicitly allows concurrent activity by waiting. But if the I/O is done at some low level of abstraction, higher levels may be quite unaware of it. The most insidious example of this is I/O caused by the virtual memory system: every instruction can cause a disk read. Such a system is called ‘preemptive’; for practical purposes a task can lose the CPU at any point, since it’s too hard to predict which memory references might cause page faults.

In the second case we have a motivation for true preemption: we want some tasks to have higher priority for the CPU than others. An important special case is interrupts, discussed below.

Ways to package concurrency

In the last section we used the work ‘task’ informally to describe a more-or-less independent, more-or-less sequential part of a computation. Now we shall be less coy about how concurrency shows up in a system.

The most general way to describe a concurrent system is in terms of a set of atomic actions with the property that usually more than one of them can occur (is enabled); we will use this viewpoint in our later study of formal concurrency. In practice, however, we usually think in terms of several ‘threads’ of concurrent execution. Within a single thread only one action is enabled at a time; in general one action may be enabled from each thread, though often some of the threads are waiting or ‘blocked’, that is, have no enabled actions.

The most convenient way to do concurrent programming is in a system that allows each thread to be described as an execution path in an ordinary-looking program with modules, routines, commands, etc., such as Spec, C, or Java. In this scheme more than one thread can execute the code of the same procedure; threads have local state which is the local variables of the procedures they are executing. All the languages mentioned and many others allow you to program in this way.

In fault-tolerant systems there is a conceptual drawback to this thread model. If a failure can occur after each atomic command, it is hard to understand the program by following the sequential flow of control in a thread, because there are so many other paths that result from failure and recovery. In these systems it is often best to reason strictly in terms of independent atomic actions. We will see detailed examples of this when we study reliable messages, consensus, and replication. Applications programmed in a transaction system are another example of this approach: each application runs in response to some input and is a single atomic action.

The biggest drawback of this kind of ‘official’ thread, however, is the costs of representing the local state and call stack of each thread and of a general mechanism for scheduling the threads. There are several alternatives that reduce these costs: interrupts, control blocks, and SIMD computers. They are all based on restricting the freedom of a thread to block, that is, to yield the

processor until some external condition is satisfied, for example, until there is space in a buffer or until a lock is free.

Interrupts

An interrupt routine is not the same as a thread, because it always starts at the same point and cannot wait for another thread. The reason for these restrictions is that the execution context for an interrupt routine is allocated on someone else's stack, which means that the routine must complete before the thread that it interrupted can continue to run. On the other hand, the hardware that schedules an interrupt routine is efficient and takes account of priority within certain limits. In addition, the interrupt routine doesn't pay the cost of its own stack like an ordinary thread.

It's possible to have a hybrid system in which an interrupt routine that needs to wait turns itself into an ordinary thread by copying its state. This is tricky if the wait happens in a subroutine of the main interrupt routine, since the relevant state may be spread across several stack frames. If the copying doesn't happen too often, the interrupt-thread hybrid is efficient. The main drawbacks are that the copying usually has to be done by hand, which is error-prone, and that without compiler and runtime support it's not possible to reconstruct the call stack, which means that the thread has to be structured differently from the interrupt routine.

A simpler strategy that is widely used is to limit the work in the interrupt routine to simple things that don't require waits, and to wake up a thread to do anything more complicated.

Control blocks

Another, related strategy is to package all the permanent state of a thread, including its program counter, in a record (usually called a 'control block') and to explicitly schedule the execution of the threads. When a thread runs, it starts at the saved program counter (usually a procedure entry point) and runs until it explicitly gives up control or 'yields'. During execution it can call procedures, but when it yields its stack must be empty so that there's no need to save it, because all the state has to be in the control block. When it yields, a reference to the control block is saved where some other thread or interrupt routine can find it and queue the thread for execution when it's ready to run, for instance after an I/O operation is complete.

The advantages of this approach are similar to those of interrupts: there are no stacks to manage, and scheduling can be carefully tuned to the application. The main drawback is also similar: a thread must unwind its stack before it can wait. In particular, it cannot wait to acquire a lock at an arbitrary point in the program.

It is very common to implement the I/O system of an operating system using this kind of thread. Most people who are used to this style do not realize that it is a restricted, though efficient, case of general programming with threads.

In 'active messages', a recent variant of this scheme, you break your computation down into non-blocking segments; as the end of a segment, you package the state into an 'active message' and

send it to the agent that can take the next step. Incoming messages are queued until the receiver has finished processing earlier ones.[2]

There are lots of other ways to use the control block idea. In 'scheduler activations', for example, kernel operations are defined to that they always run to completion; if an operation can't do what was requested, it returns intermediate state and can be retried later.[3]

SIMD

This acronym stands for 'single instruction, multiple data', and refers to processors in which several execution units all execute the same sequence of instructions on different data values. In a 'pure' SIMD machine every instruction is executed at the same time by all the processors (except that some of them might be disabled for that instruction). Each processor has its own memory, and the processors can exchange data as part of an instruction. A few such machines were built between 1970 and 1993, but they are now out of favor. The same programming paradigm is still used in many scientific problems however, at a coarser grain. In one step each processor does some computation on its private data. When all of them are done, they exchange some data and then take the next step. The action of detecting that all are done is called 'barrier synchronization'.

The term 'SIMD' has been recycled in the Intel MMX instruction set, and similar designs from several other manufacturers, to describe something much more prosaic: doing 8 8-bit adds in parallel on a 64-bit data path.

Easy concurrency

Concurrency is easy when you program with locks. The rules are simple:

> Every shared variable must be protected by a lock. A variable is shared if it is touched by more than one thread. You can think of data that is private to a thread as being protected by an implicit lock that is always held by the thread.
>
> You must hold the lock for a shared variable before you touch the variable. The essential property of a lock is that two threads can't hold the same lock at the same time. This property is called 'mutual exclusion'; the abbreviation 'mutex' is another name for a lock.
>
> If you want an atomic operation on several shared variables that are protected by different locks, you must not release any locks until you are done. This is called 'two-phase locking', because there is a phase in which you only acquire locks and don't release any, followed by a phase in which you only release locks and don't acquire any.

Then your computation between the point that you acquire a lock and the point that you release it is equivalent to a single atomic action, and therefore you can reason about it sequentially. This atomic part of the computation is called a 'critical section'. To use this method reliably, you should annotate each shared variable with the name of the lock that protects it, and clearly

[2] T. von Eiken et al., Active messages: A mechanism for integrated communication and computation. *Proc. International Symposium on Computer Architecture*, May 1992, pp 256-267.
[3] T. Anderson et al., Scheduler activations: Effective kernel support for the user-level management of parallelism. *ACM Transactions on Computer systems* **10**, 1 (Feb. 1992), pp 54-79.

bracket the regions of your program within which you hold each lock. Then it is a mechanical process to check that you hold the proper lock whenever you touch a shared variable.

Why do locks lead to big atomic actions. Intuitively, the reason is that no other well-behaved thread can touch any shared variable while you hold its lock, because a well-behaved thread won't touch a shared variable without itself holding its lock, and only one thread can hold a lock at a time. We will make this more precise in handout 16 and give a proof of atomicity.

Actually locks give you a bit more atomicity than this. If a well-behaved thread acquires a sequence of locks and then releases them (not necessarily in the same order), the entire computation from the first acquire to the last release is atomic. Once you have done a release, however, you can't do another acquire without losing atomicity.

The simple locks we have been describing are also called 'mutexes'; this is short for "mutual exclusion". As we shall see, more complicated kinds of locks are often useful.

Here is the spec for a mutex. It maintains mutual exclusion by allowing the mutex to be acquired only when no one already holds it. If a thread other than the current holder releases the mutex, the result is undefined. If you try to do an `Acquire` when the mutex is not free, you have to wait, since `Acquire` has no transition from that state because of the `m = nil` guard.

```
MODULE OneMutex EXPORT Acquire, Release =

VAR  m: (Thread + Null) := nil
% A mutex is either nil or the thread holding the mutex.
% The variable SELF is defined to be the thread currently making a transition.

APROC Acquire() = << m = nil  => m := SELF; RET >>
APROC Release() = << m = SELF => m := nil ; RET [*] HAVOC >>

END OneMutex
```

We usually need lots of mutexes, not just one, so we define a `Mutex.M` type which indexes a function that maps `M` to the state of one mutex, and we give the type `acq` and `rel` methods. There's also a `New` procedure for making a new, free mutex; it chooses an unused `M`.

```
MODULE Mutex EXPORT M, New =

TYPE M = Int WITH {acq:=Acquire, rel:=Release}
VAR  s: M -> (Thread + Null) := {}

APROC New() -> M = << VAR m | ~s!m => s(m) := nil; RET m >>

APROC Acquire(m) = << s(m) = nil  => s(m) := SELF; RET >>
APROC Release(m) = << s(m) = SELF => s(m) := nil ; RET [*] HAVOC >>

END Mutex
```

Invariants

In fact things are not so simple, since a computation seldom consists of a single atomic action. A thread should not hold a lock forever (except on private data) because that will prevent any other thread that needs to touch the data from making progress. Furthermore, it often happens that a thread can't make progress until some other thread changes the data protected by a lock. A

simple example of this is a FIFO buffer, in which a consumer thread doing a `Get` on an empty buffer must wait until some other producer thread does a `Put`. In order for the producer to get access to the data, the consumer must release the lock. Atomicity does not apply to code like this that touches a shared variable `x` protected by a mutex `m`:

```
m.acq; touch x; m.rel; private computation; m.acq; touch x; m.rel
```

This code releases a lock and later re-acquires it, and therefore isn't atomic. So we need a different way to think about this situation, and here it is.

> After the `m.acq` the only thing you can assume about `x` is an invariant that holds whenever `m` is unlocked.

As usual, the invariant must be true initially, and the program must establish it before unlocking `m`. While `m` is locked, you can poke around in `x` and discover facts that are not implied by the invariant, but you cannot assume that any of these facts are still true after you unlock `m`.

To use this methodology effectively, of course, you must *write the invariant down*.

Here is a more picturesque way of describing this method. To do easy concurrent programming:

first you put your hand over some shared variables, say `x` and `y`, so that they can't change,

then you do something with them, and

finally you take your hand away.

The reason `x` and `y` can't change is that the rest of the program obeys some conventions; in particular, it acquires locks before touching shared variables.

This viewpoint sheds light on why fault-tolerant programming is hard: `Crash` is no respecter of conventions, and the invariant must be maintained even though a `Crash` may stop an update in mid-flight and reset all or part of the volatile state.

Scheduling: Condition variables

If a thread can't make progress until some condition is established, and therefore has to release a lock so that some other thread can establish the condition, the simplest idiom is

```
m.acq; DO ~ condition involving x => m.rel; m.acq OD; touch x; m.rel
```

This is called "busy waiting", because the thread keeps computing, waiting for the condition to become true; it keeps releasing the lock so that some other thread can make the condition true. This code is correct, but reacquiring the lock immediately makes it more difficult for another thread to get it, and going around the loop while the condition remains false wastes processor cycles. Even if you have your own processor, this isn't a good scheme because of the system-wide cost of repeatedly acquiring the lock.

The way around these problems is an optimization that replaces `m.rel; m.acq` with `c.wait(m)`, where `c` is a 'condition variable'. The `c.wait(m)` releases `m` and then blocks the thread until some other thread does `c.signal`. Then it reacquires `m` and returns. If several threads are waiting, `signal` picks one or more to continue in a fair way. The variation `c.broadcast` continues all the waiting threads.

Here is the spec for condition variables. It says that a `C` is the set of threads waiting on the condition, and it allows for lots of `C`'s just as the `Mutex` spec does. The `wait` method is especially interesting, since it's the first procedure we've seen in a spec that is not atomic. This is because the whole point is that during the `wait` other threads have to run, access the variables protected by the mutex, and signal the condition variable. Note that `wait` takes an extra parameter, the mutex to release and reacquire.

```
MODULE Condition EXPORT C, New =

TYPE C = Int WITH {wait:=Wait, signal:=Signal, broadcast:=Broadcast}
     M = Mutex.M

VAR  s: C -> SET Thread := {}
% Each condition variable is the set of waiting threads.

APROC New() -> C = << VAR c | ~s!c => s(c) := {}; RET c >>

PROC Wait(c, m) =
     << s(c) := s(c) + {SELF}; m.rel >>;              % m.rel=HAVOC unless SELF IN m
     << ~ (SELF IN s(c)) => m.acq >>

APROC Signal(c) = <<
% Remove at least one thread from c.  In practice, usually just one.
     IF   VAR t: SET Thread | t <= s(c) /\ t # {} => s(c) := s(c) - t
     [*]  SKIP
     FI >>

APROC Broadcast(c) = << s(c) := {} >>

END Condition
```

For this scheme to work, a thread that changes `x` so that the condition becomes true must do a `signal` or `broadcast`, in order to allow some waiting thread to continue. A foolproof but inefficient strategy is to have a single condition variable for `x` and to do `broadcast` whenever `x` changes at all. More complicated schemes can be more efficient, but are more likely to omit a `signal` and leave a thread waiting indefinitely. The paper by Birrell in handout 15[4] gives many examples and some good advice.

Note that you are *not* entitled to assume that the condition is true just because `wait` returns. That would be a little more efficient for the waiter, but it would be much more error prone, and it would require a tighter spec for `wait` and `signal` that is often less efficient to implement. You are supposed to think of `c.wait(m)` as just an optimization of `m.rel; m.acq`. This idiom is very robust.

Remember that after `c.wait(m)` you cannot assume anything about `x` beyond its invariant, since the `wait` unlocks `m` and then locks it again. After a `wait`, only the invariant is guaranteed to hold, not anything else that was true about `x` before the wait.

[4] Andrew Birrell, *An Introduction to Programming with Threads*, research report 35, Systems Research Center, Digital Equipment Corporation, January 1989.

Really easy concurrency

An even easier kind of concurrency uses buffers to connect independent modules, each with its own set of variables disjoint from those of any other module. Each module consumes data from some predecessor modules and produces data for some successor modules. In the simplest case the buffers are FIFO, but they might be unordered or use some other ordering rule. A little care is needed to program the buffers' `Put` and `Get` operations, but that's all. This is often called 'pipelining'. The fancier term 'data flow' is used if the modules are connected not linearly but by a more general DAG.

Another really easy kind of concurrency is provided by transaction processing systems, in which an application program accepts some input, reads and updates a shared database, and generates some output. The transaction mechanism makes this entire operation atomic, using techniques that we will describe later. The application programmer does not have to think about concurrency at all.

Hard concurrency

If you don't program according to the rules for locks, then you are doing hard concurrency, and it will be hard. Why bother? There are three reasons:

> You may have to implement mutexes and condition variables on top of something weaker, such as the atomic reads and writes of memory that a basic processor or file system gives you. Of course, only the low-level runtime implementer will be in this position.
>
> It may be cheaper to do the work you need to do using weaker primitives than mutexes. If efficiency is important, hard concurrency may be worth the trouble. But you will pay for it either in bugs or in careful proofs of correctness.
>
> It may be important to avoid waiting for a lock to be released. Even if a critical section is coded carefully so that it doesn't do too much computing, there are still ways for the lock to be held for a long time. If the thread holding the lock can fail independently (for example, if it is in a different address space or on a different machine), then the lock can be held indefinitely. If the thread can get a page fault while holding the lock, then the lock can be held for a disk access time. A concurrent algorithm that prevents one slow (or failed) thread from delaying other thread too much is called 'wait free'.[5]

In fact, the "put out your hand" way of looking at things applies to hard concurrency as well. The difference is that instead of preventing `x` and `y` from changing at all, you do something to ensure that some predicate `p(x, y)` will remain true. The convention that the rest of the program obeys may be quite subtle. A simple example is the careful write solution to keeping track of free space in a file system (handout 7, page 16), in which the predicate is

```
free(da) ==> ~ Reachable(da).
```

The special case of locking maintains the strong predicate `x = x0 /\ y = y0` (unless you change `x` or `y` yourself).

[5] M. Herlihy, Wait-free implementations of concurrent objects. *Proc. ACM Symposium on Principles of Distributed Computing*, Toronto, Aug 1988, pp 276-290.

We postpone a detailed study of hard concurrency to handout 16.

Problems in easy concurrency: Scheduling

If there is a shortage of processor resources, there are various ways in which the simple easy concurrency method can go astray. In this situation we may want some threads to have priority over others, but subject to this constraint we want the processor resources allocated fairly. This means that the amount of time a task takes should be roughly proportional to the amount of work it does; in particular, we don't want short tasks to be blocked by long ones.

Priority inversion

When there are priorities there can be "priority inversion". This happens when a low-priority thread `A` acquires a lock and then loses the CPU, either to a higher-priority thread or to round-robin scheduling. Now a high-priority thread `B` tries to acquire the lock and ends up waiting for `A`. Clearly the priority of `A` should be temporarily increased to that of `B` until `A` completes its critical section, so that `B` can continue. Otherwise `B` may wait for a long time while threads with priorities between `A` and `B` run, which is not what we had in mind when we set up the priority scheme.

Granularity of locks

A different issue is the 'granularity' of the locks: how much data is protected by each lock. A single lock is simple and cheap, but doesn't allow any concurrency. Lots of fine-grained locks allow lots of concurrency, but the program is more complicated, there's more overhead for acquiring locks, and there's more chance for deadlock (discussed in the next section). For example, a file system might have a single global lock, one lock on each directory, one lock on each file, or locks only on byte ranges within a file. The goal is to have fine enough granularity that the queue of threads waiting on a mutex is empty most of the time. More locks than that don't accomplish anything.

It's possible to have an adaptive scheme in which locks start out fine-grained, but when a thread acquires too many locks they are collapsed into fewer coarser ones that cover larger sets of variables. This process is called 'escalation'. It's also possible to go the other way: a process keeps track of the exact variables it needs to lock, but takes out much coarser locks until there is contention. Then the coarse locks are 'de-escalated' to finer ones until the contention disappears.

Closely related to the choice of granularity is the question of how long locks are held. If a lock that protects a lot of data is held for a long time (for instance, across a disk reference or an interaction with the user) concurrency will obviously suffer. Such a lock should protect the minimum amount of data that is in flux during the slow operation.

On the other hand, sometimes you want to minimize the amount of communication needed for acquiring and releasing the same lock repeatedly. To do this, you hold onto the lock for longer than is necessary for correctness. Another thread that wants to acquire the lock must somehow signal the holder to release it. This scheme is commonly used in distributed coherent caches, in which the lock only needs to be held across a single read, write, or test-and-set operation, but one thread may access the same location (or cache line) many times before a different thread touches it.

Lock modes

Another way to get more concurrency at the expense of complexity is to have many lock 'modes'. A mutex has one mode, usually called 'exclusive' since 'mutex' is short for 'mutual exclusion'. A reader/writer lock has two modes, called exclusive and 'shared'. It's possible to have as many modes as there are different kinds of commuting operations. Thus all reads commute and therefore need only shared mode (reader) locks. But a write commutes with nothing and therefore needs an exclusive mode (write) lock. The commutativity of the operations is reflected in a 'conflict relation' on the locks. For reader/writer or shared/exclusive locks this matrix is:

	None	Shared (read)	Exclusive (write)
None	OK	OK	OK
Shared (read)	OK	OK	Conflict
Exclusive (write)	OK	Conflict	Conflict

Just as different granularities bring a need for escalation, different modes bring a need for 'lock conversion', which upgrades a lock to a higher mode, for instance from shared to exclusive, or downgrades it to a lower mode.

Explicit scheduling

In simple situations, queuing for locks is an adequate way to schedule threads. When things are more complicated, however, it's necessary to program the scheduling explicitly because the simple first-come first-served queuing of a lock isn't what you want. A set of printers with different properties, for example, can be optimized across a set of jobs with different priorities, requirements for paper handling, paper sizes, color, etc. There have been many unsuccessful attempts to build general resource allocation systems to handle these problems. They fail because they are too complicated and expensive for simple cases, and not flexible enough for complicated ones. A better strategy is to program the scheduling as part of the application, using as many condition variables as necessary to queue threads that are waiting for resources. Application-specific data structures can keep track of the various resource demands and application-specific code, perhaps written on top of a library, can do the optimization.

Problems in easy concurrency: Deadlock

The biggest problem for easy concurrency is deadlock, in which there is a cycle of the form

> Lock a is held by thread 1.
> Thread 1 is waiting for lock b.
> Lock b is held by thread 2.
> ...
> Lock h is held by thread 8.
> Thread 8 is waiting for lock a.

All the locks and threads are nodes in a lock graph with the edges "lock a is held by thread 1", "thread 1 is waiting for lock b", etc.

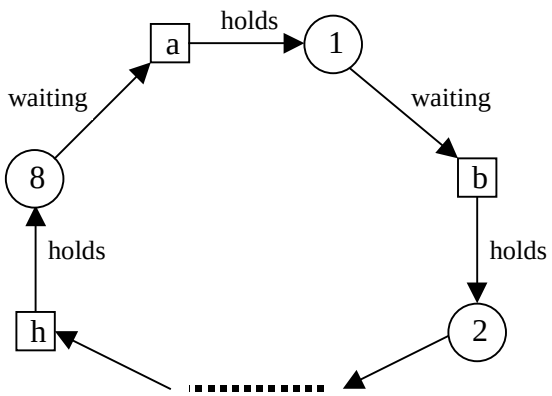

The way to deal with this that is simplest for the application programmer is to *detect* a deadlock[6] and automatically roll back one of the threads, undoing any changes it has made and releasing its locks. Then the rolled-back thread retries; in the meantime, the others can proceed. Unfortunately, this approach is only practical when automatic rollback is possible, that is, when all the changes are done as part of a transaction. We will be talking about this in a few weeks.

The main alternative is to *avoid* deadlocks by defining a partial order on the locks, and abiding by a rule that you only acquire a lock if it is greater than every lock you already hold. This ensures that there can't be any cycles in the graph of threads and locks. Note that there is no requirement to release the locks in order, since a release never has to wait.

To implement this idea you

> annotate each shared variable with its protecting lock (which you are supposed to do anyway when practicing easy concurrency),
>
> state the partial order on the locks, and
>
> annotate each procedure or code block with its 'locking level' `ll`, the maximum lock that can be held when it is entered, like this: `ll <= x`.

Then you always know textually the biggest lock that can be held (by starting at the procedure entry with the annotation, and adding locks that are acquired), and can check whether an `Acquire` is for a bigger lock as required, or not. With a stronger annotation that tells exactly what locks are held, you can subtract those that are released as well. You also have to check when you call a procedure that the current locking level is consistent with the procedure's annotation. This check is very similar to type checking.

Having described the basic method, we look at some examples of how it works and where it runs into difficulties.

If resources are arranged in a tree and the program always traverses the tree down from root to leaves, or up from leaves to root, (in the usual convention, which draws trees upside down) then the tree defines a suitable lock ordering. Examples are a strictly hierarchical file system or a tree of windows. If the program sometimes goes up and sometimes goes down, there are problems;

---

[6] For ways of detecting deadlocks, see Gray and Reuter, pp 481-483 and A. Thomasian, Two phase locking performance and its thrashing behavior. *ACM Trans. on Database Systems* **18**, 4 (Dec. 1993), pp. 579-625

we discuss some solutions shortly. If instead of a tree we have a DAG, it still defines a suitable lock ordering.

Often, as in the file system example, this graph is actually a data structure whose links determine the accessibility of the nodes. In this situation you can choose when to release locks. If the graph is static, it's all right to release locks at any time. If you release each lock before acquiring the next one, there is no danger of deadlock regardless of the structure of the graph, because a flat ordering (everything unordered) is good enough as long as you hold at most one lock at a time. If the graph is dynamic and a node can disappear when it isn't locked, you have to hold on to one lock at least until after you have acquired the next one. This is called 'lock coupling', and a cyclic graph can cause deadlock. We will see an example of this when we study hierarchical file systems.

Here is another common locking pattern. Consider a program that manipulates objects named by handles and maintains a set of these objects. For example, the objects might be buffers, and the set the buffers that are non-empty. One thread works on an object and sometimes needs to mess with the set, for instance when a buffer changes from empty to non-empty. Another thread processes the set and needs to mess with some of the objects, for instance to empty out the buffers at regular intervals. It's natural to have a lock `h.m` on each object and a lock `ms` on the set. How should they be ordered? We work out a solution in which the ordering of locks is every `h.m < ms`.

```
TYPE
    H           =  Int WITH {acq:=(\h|ot(h).m.acq),   % Handle (index in ot)
                             rel:=(\h|ot(h).m.rel),
                             y  :=(\h|ot(h).y ), empty:=...}

VAR
    s           :  SET H                              % Set protected by ms
    ms          :  Mutex.M
    ot          :  H -> [m, y: Any]                   % Object Table. m protects y,
                                                      % which is the object's data
```

Note that each piece of state that is not a mutex is annotated with the lock that protects it: `s` with `ms` and `y` with `m`. The 'object table' `ot` is fixed and therefore doesn't need a lock.

We would like to maintain the invariant "object is non-empty" = "object in set": `~ h.empty = h IN s`. This requires holding both `h.m` and `ms` when the emptiness of an object changes. Actually we maintain `h.m is locked \/ (~ h.empty = h IN s)`, which is just as good. The `Fill` procedure that works on objects is very straightforward; `Add` and `Drain` are functions that compute the new state of the object in some unspecified way, leaving it non-empty and empty respectively. Note that `Fill` only acquires `ms` when it becomes non-empty, and we expect this to happen on only a small fraction of the calls.

```
PROC Fill(h, x: Any) =
% Update the object h using the data x
    h.acq;
    IF h.empty => ms.acq; s := s + {h}; ms.rel [*] SKIP FI;
    ot(h).y := Add(h.y, x);
    h.rel
```

The Demon thread that works on the set is less straightforward, since the lock ordering keeps it from acquiring the locks in the order that is natural for it.

```
THREAD Demon() = DO
    ms.acq;
    IF  VAR h | h IN s  =>
            ms.rel;
            h.acq; ms.acq;                          % acquire locks in order
            IF   h IN s =>                          % is h still in s?
                 s := s - {h}; ot(h).y := Drain(h.y)
            [*]  SKIP
            FI;
            ms.rel; h.rel
    [*]  ms.rel
    FI
  OD
```

Drain itself does no locking, so we don't show its body.

The general idea, for parts of the program that can't acquire locks in the natural order, is to collect the information you need, one mutex at a time. Then reacquire the locks according to the lock ordering, check that things haven't changed (or at least that your conclusions still hold), and do the updates. If it doesn't work out, retry. Version numbers can make the 'didn't change' check cheap. This scheme is closely related to optimistic concurrency control, which we discuss later in connection with concurrent transactions.

It's possible to use a hybrid scheme in which you keep locks as long as you can, rather than preparing to acquire a lock by always releasing any larger locks. This works if you can acquire a lower lock 'cautiously', that is, with a failure indication rather than a wait if you can't get it. If you fail in getting a lower lock, fall back to the conservative scheme of the last paragraph. This doesn't simplify the code (in fact, it makes the code more complicated), but it may be faster.

Nested monitors

A case in which deadlock avoidance by ordering doesn't work is known as "nested monitors". It comes up when there are two levels of abstraction, H and L (for high and low), each with its own lock lH and lL. L has a condition variable cL. The code that deadlocks looks like this, if two threads 1 and 2 are using H, 1 needs to wait on cL, and 2 will signal cL.

```
    H1: lH.lock; call L1
    L1: lL.lock; cL.wait(lL)

    H2: lH.lock; call L2
    L2: lL.lock; cL.signal
```

This will deadlock because the wait in L1 releases lL but not lH, so that H2 can never get past lH.lock to reach L2 and do the signal. This is not a lock-lock deadlock because it involves the condition variable cL, so a straightforward deadlock detector will not find it. The picture below illustrates the point.

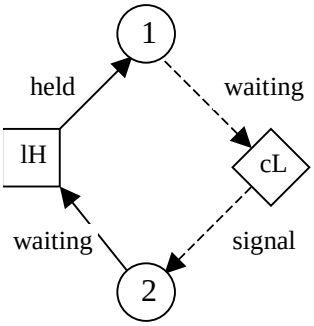

To avoid this deadlock, don't wait on a condition with *any* locks held, unless you know that the `signal` can happen without acquiring any of these locks. The 'don't wait' is simple to check, given the annotations that the methodology requires, but the 'unless' is not simple.

People have proposed to solve this problem by generalizing `wait` so that it takes a set of mutexes to release instead of just one. Why is this a bad idea? Aside from the problems of passing the right mutexes down from `H` to `L`, it means that any call on `L` might release `lH`. The `H` programmer must to be careful not to depend on anything more than the `lH` invariant across any call to `L`. This style of programming is very error-prone.

Simple vs. fancy locks

We have described a number of features that you might want in a locking system:

- multiple modes with conversion, for instance from shared to exclusive;
- multiple granularities with escalation from fine to coarse and de-escalation from coarse to fine;
- deadlock detection.

Database systems typically provide these features. In addition, they acquire locks automatically based on how an application transaction touches data, choosing the mode based on what the operation is, and they can release locks automatically when a transaction commits. For a thorough discussion of database locking see Jim Gray and Andreas Reuter, *Transaction Processing: Concepts and Techniques,* Morgan Kaufmann, 1993, Chapter 8, pages 449-492.

The main reason that database systems have such elaborate locking facilities is that the application programmers are quite naive and can't be expected to understand the subtleties of concurrent programming. Instead, the system does almost everything automatically, and the programmers can safely assume that execution is sequential. Automatic mechanisms that work well across a wide range of applications need to adapt in the ways listed above.

By contrast, a simple mutex has only one mode (exclusive), only one granularity, and no deadlock detection. If these features are needed, the programmer has to provide them using the mutex and condition primitives. We will study one example of this in detail in handout 16: building a reader/writer lock from a simple mutex. Many others are possible.