

# Basic Threads Programming: Standards and Strategy

Mike Dahlin  
Department of Computer Sciences  
University of Texas at Austin  
TR-08-07

## Forward

What follows is a handout I use when teaching undergraduate and graduate students the basics of multithreaded programming.

The approach I take may be controversial; it certainly is opinionated; but I am convinced that it works. For approximately five years before developing this handout and the associated lecture, midterms and finals revealed that a large fraction of my students emerged from my classes unable to write even the simplest concurrent programs. In the five or so years since I have adopted this approach, the vast majority are able to write correct programs, at least for the small programs that one can assign as an examination question or course project.

Advanced programmers may find the rules I outline to be too restrictive, but I strongly believe that the right way to teach this material is to first make sure students have a rock solid grasp of the basics. Once they *understand* the basic rules described here and have used them for a while, they may be in a position to learn to properly apply more advanced techniques.

That said, I suspect that many advanced threads programmers might find their code improved if they followed the rules here more often than they do. Since I began teaching these rules, I have also disciplined myself to follow them unless there is a very good reason not to. I have found exceptions few and far between, and the vast majority of the time when I find myself deviating from the rules and then force myself to rewrite the code to follow them, the code improves. Also, advanced techniques are harder to use than many programmers realize [1].

These rules may come across as opinionated. They are. When I deliver them in course, I literally stand on a soapbox and pronounce them as commandments. My experience is that teaching things this way is effective.

I should note that the rules outlined here are far from novel. For example, in 1980 Lampson and Redell articulate similar advice in a much more measured tone [2].

## References

- [1] D. Bacon, J. Bloch, J. Bogda, C. Click, P. Haahr, D. Lea, T. May, J. Maessen, J. Mitchell, K. Nilsen, B. Pugh, and E. Sirer. The “Double-checked locking is broken” declaration. <http://www.cs.umd.edu/~pugh/java/memoryModel/DoubleCheckedLocking.html>.
- [2] B. W. Lampson and D. D. Redell. Experience with processes and monitors in Mesa. *Communications of the ACM*, 23(2):105–117, 1980.

# 1 Motivation

Some people rebel against coding standards. I don't understand the logic. For concurrent programming in particular, there are a few good solutions that have stood the test of time (and many unhappy people who have departed from these solutions.) For concurrent programming, *debugging won't work*. You must rely on (a) writing correct code and (b) writing code that you and others can read and understand. Following the rules below will help you write correct, readable code.

Rules 2 through 6 below are *required coding standards* for CS372. Answers to homework and exam problems and project code that do not follow these standards are by definition *incorrect*. Section 5 discusses additional restrictions for project code (or exam pseudo-code) in Java. Again, answers that deviate from these *required coding standards* are, by definition, *incorrect* for this course.

If you believe you have a better way to write concurrent programs, that's great! Bring it to us (before you use it on an assignment!) We will examine your approach the same way we hope that a civil engineering manager would examine a proposal by a bright young civil engineer who has a proposal for a better configuration of rebar for reinforcing a new bridge: we will hope you have found a better way, but the burden of proof for explaining the superiority of your approach and proving that there are no hidden pitfalls is on you.

## 2 Coding standards for concurrent programs:

1. Always do things the same way

Even if one way is not better than another, following the same strategy every time (a) frees you to focus on the core problem because the details of the standard approach become a habit and (b) makes it easier for the people who follow you and have to review, maintain, or debug your code understand your code. (Often the person that has to debug the code is you!)

Electricians follow standards for the colors of wire they use for different tasks. White is neutral. Black or red is hot. Copper is ground. An electrician doesn't have to decide "Hm. I have a bit more white on my belt today than black, should I use white or black for my grounds?" When an electrician walks into a room she wired last month, she doesn't have to spend 2 minutes trying to remember which color is which. If an electrician walks into a room she has never seen before, she can immediately figure out what the wiring is doing (without having to trace it back into the switchboard.) Similar advantages apply to following coding standards.

But for concurrency programs, the evidence is that in fact the abstractions we describe *are* better than almost all others. Until you become a *very* experienced concurrent programmer, you should take advantage of the hard-won experience of those that have come before you. Once you are a concurrency guru, you are welcome to invent a better mousetrap.

Sure, you can cut corners and occasionally save a line or two of typing by departing from the standards, but you'll have to spend a few minutes thinking to convince yourself that you are right on a case-by-case basis (and another few minutes typing comments to convince the next person to look at the code that you're right), and a few hours or weeks tracking down bugs when you're wrong. NOT WORTH IT!

2. Always use monitors (condition variables + locks)

Either semaphores or monitors (condition variables + locks) could be used to write concurrent programs. We recommend that you be able to read semaphores (so you can understand legacy code), but that you only write new code using condition variables and locks (e.g., monitors.)

99% of the time monitor code is more clear than the equivalent semaphore code because monitor code is more "self-documenting". If the monitor code is well-structured, it is usually clear what each synchronization action is doing. Sure, occasionally semaphores seem to fit what you are doing perfectly because you can map the invariants onto the internal state of the semaphore exactly (e.g., producer consumer), but what happens when the code changes a bit next month? Will the fit be as good? Rule #1 says that you should choose one of the two styles and stick with it, and my opinion is the right one to pick is monitors.

3. Always hold lock when operating on a condition variable

You signal on a condition variable because you just got done manipulating shared state. You proceed when some condition about a shared state becomes true. Condition variables are useless without shared state and shared state is useless without holding a lock.

Many modern libraries enforce this rule – you cannot call any condition variable methods unless you hold the corresponding lock. But some run-time systems you may see in the future allow sloppiness – don't fall for it.

4. Always grab lock at beginning of procedure and release it right before return

This is mainly an extension of rule #1 - pick one way of doing things and always follow it. (In particular, it is easy to read code and see where the lock is held and where it isn't because things break down on a procedure by procedure basis.)

Also, if there is a logical chunk of code that you can identify as a set of actions that require a lock, then that section should probably be its own procedure - it is a set of logically related actions. If you find yourself wanting to grab a lock in the middle of a procedure, that is usually a red flag that you should break the piece you are considering into a separate procedure. We are all sometimes lazy about creating new procedures when we should. Take advantage of this signal, and you will write clearer code.

5. Always use “`while(predicateOnStateVariables(...) == true/false){ condition->wait()}`”, not “`if...`”

(Where `PredicateOnStateVariables(...)` looks at the state variables of the current object to decide if it is OK to proceed.)

`While` works any time `if` does, and it works in situations when `if` doesn't. By rule 1, you should do things the same way every time.

`If` breaks modularity. One might be tempted to use `if` if one convinced oneself that there will be exactly one signal when one waiting thread should proceed. The problem is that a change in code in one procedure (say, adding a `signal()`) can then cause a bug in another procedure (where the `wait()` is). `While` code is also self-documenting – one can look at the `wait()` and see exactly under what conditions a thread may proceed.

When you always use `while`, you are given incredible freedom about where you put the `signals()`. In fact, `signal()` becomes a hint – you can add more signals to a correct program in arbitrary places and it remains a correct program!

6. (Almost) never `sleep()`

Never use `sleep()` to wait for another thread to do something. The correct way to wait for a condition to become true is to `wait()` on a condition variable.

In general, `sleep()` is only appropriate when there is a particular real-world moment in time when you want to perform some action. If you catch yourself writing `while(some condition){sleep();}`, treat this as a big red flag that you are probably making a mistake.

I'm sure there are valid exceptions to all of the above rules, but they are few and far between. And the benefit you get by occasionally breaking the rules is unlikely to make up for the cost in your effort, extra debugging and maintenance cost, and loss of modularity.

### 3 Problem solving strategy

The following strategy should help you take a systematic approach to organizing your thoughts on a synchronization problem. Unlike the rules above, this is just advice. You are free to take any approach that works for you. Of course, if you write down your thinking at each of these steps, this is a good way to get lots of partial credit on exam questions.

#### 1. Decompose problem into objects

- Identify units of concurrency. Make each a thread with a `go()` method. Write down the actions a thread takes at a high level.
- Identify shared chunks of state. Make each shared *thing* an object. Identify the methods on those objects – the high-level actions made by threads on these objects.
- Write down the high-level main loop of each thread.

Advice: stay high level here. Don't worry about synchronization yet. Let the objects do the work for you.

Now, for each object:

2. Write down the synchronization constraints on the solution. Identify the type of each constraint: *mutual exclusion* or *scheduling*
3. Create a lock or condition variable corresponding to each constraint
4. Write the methods, using locks and condition variables for coordination

## 4 Example

Basic template:

- Object oriented style of programming - encapsulate shared state and synchronization variables inside of objects

(**Hint:** don't manipulate synchronization variables or shared state variables in the code associated with a thread, do it with the code associated with a shared object. Warning: most examples in the book are lazy and talk about "thread 1's code" and "thread 2's code", etc. This is b/c most of the "classic" problems were studied before OO programming was widespread, and the textbooks have not caught up.)

- Always use condition variables for code you write.

Be able to understand code written in semaphores. But the coding standard your manager (me) is enforcing for this group is condition variables for synchronization

```
class Foo{

private:
    // Synchronization variables
    Lock mutex;
    Cond condition1;
    Cond condition2;

    // State variables

public:
    Foo::foo()
    {
    /*
    * 3) Always, grab mutex at start of procedure, release at
    * end (or at any return!!!). Reasoning: if there is a logical
    * set of actions to do when you hold a mutex, that logical
    * set of actions should be expressed as a procedure, right?
```

```

*/
Assert(invariants hold - shared variables in consistent state)
mutex->acquire(){
    /*
    * Invariants may or may not hold; shared variables may be
    * in inconsistent state
    */

    /*
    * 4) always "while" never "if"
    */
while(shared variables in some state){
    assert(invariants hold)

    /*
    * 5) Always hold lock when operating on C.V.
    */
    condition1->wait(&mutex)
    assert(invariants hold);
}

    /*
    * Invariants may or may not hold; shared variables may be
    * in inconsistent state
    */

    /*
    * 5) Always hold lock when operating on C.V.
    */
    condition2->signal(&mutex);
    condition1->signal(&mutex);
}mutex->release()
}
}; // Class

```

- Rule 6) (Almost) never sleep()

Sleep(time) puts the current thread on a waiting queue at the timer - only use it to wait until a specific time, not to wait for an event of a different sort

Hint: sleep should never be in a while()sleep

Problems with using sleep:

1. no atomic release/reacquire lock
2. really inefficient (example - cascading sleeps in Aname)
3. not logical

Warning: on the project and on exams, improper use of sleep will be regarded as strong evidence that you have no idea how to write multi-threaded programs and will affect your grade accordingly. (I make this a point of emphasis b/c this error is so common in past years and easy to avoid.)

As we begin examining more complicated problem, some additional issues will arise in your problem solving strategy (e.g., your deadlock avoidance scheme), but this should be a good start.

## 5 Java

### 5.1 Restrictions for this course

Java is a programming language that has included support for threads from day 1. There are two challenges for using it in a course to teach multi-threaded programming. First, Java includes some language features that should almost never be used (IMHO.) Second, Java includes some language features that are good features for production programming, but that can interfere with teaching/learning good multi-threaded programming practice. As a result, for this course, we require you to follow a heavily restricted programming style.

The following rules are *required coding standards* for all projects. Also, if you use Java-like pseudocode on the exams, you must adhere to these standards as well.

1. Keep shared state classes separate from thread classes. No class that extends `Thread` or implements `Runnable` may include code that manipulates a lock or condition variable.

Threads operate on shared state by accessing shared objects. These shared objects include synchronization. Threads, themselves, just have private per-thread state and do not need synchronization. Since threads are Objects in Java, it is easy to get confused and start mixing these ideas up. This rule tries to make the distinction clear.

2. The *synchronized* keyword is forbidden.
3. You must explicitly allocate locks. You may not use the “default” lock associated with each object.
4. You must explicitly allocate condition variables. You may not use the “default” condition variable associated with each object.

The bottom line for rules 2-4 is that I am forcing you to manually construct monitors using locks and condition variables rather than allowing the Java language to handle most of the details. I am doing this primarily to ensure that you write down your thought process so that we can evaluate whether you understand all of the steps. A secondary advantage is that you will emerge from the course knowing this “advanced” Java technique that you will sometimes need to use (and it will be easy for you to quickly learn the simpler “standard” technique when you leave the course.)

These rules yield the following *preferred* method for writing a synchronized method

```
public class Foo{
    private Lock lock;
    private Condition c1;
    private Condition c2;

    public Foo()
    {
        lock = new SimpleLock();
        c1 = lock.newCondition();
        c2 = lock.newCondition();
        ...
    }

    public void doIt()
    {
        try{
            lock.lock();
            ...
        }
    }
}
```

```

    while(...){
        c1.awaitUninterruptibly();
    }

    ...

    c2.signal();
}
finally{
    lock.unlock();
}
}
}

```

Notice how embedding the `lock.lock()` call in a `try` block and putting the `lock.unlock()` call in the `finally` block ensures that you release the lock as the last action in the method. A slightly less preferred, but acceptable template would be the following:

```

public class Foo{
    private Lock lock;
    private Condition c1;
    private Condition c2;

    public Foo()
    {
        lock = new SimpleLock();
        c1 = lock.newCondition();
        c2 = lock.newCondition();
    }

    public void doIt()
    {
        lock.lock();

        ...

        while(...){
            c1.awaitUninterruptibly();
        }

        ...

        c2.signal();
        lock.unlock();
    }
}

```

## 5.2 Restrictions for production programming

The *synchronized* keyword is fine most of the time, and I encourage you to use it when you write most of your Java code once you leave this course. In particular, you should use synchronized methods and the default object locks and condition variables when one condition variable is all you need for an object. This approach maps well to the rules discussed in Section 1 of this document. And you can still use the “advanced” templates above when multiple condition variables per object are needed.

Note that there are still some language features I would avoid. In particular, Java allows “synchronized blocks” in the middle of a method (not just synchronized method). This feature violates the rule of grabbing locks at the beginning of a method and releasing them at the end, and should be avoided.

Recommendations:

1. Keep shared state classes separate from thread classes. No class that extends `Thread` or implements `Runnable` should include code that manipulates a lock or condition variable.

This is still a good rule for the reasons mentioned above.

2. You should refrain from defining a synchronized block in the middle of a procedure

So, the code such as the following is *forbidden in this course* but *encouraged in the real world*:

```
public class Foo{

    public Foo()
    {
    }

    public synchronized void doIt()
    {

        ...

        while(...){
            this.wait();
        }

        ...

        this.notify();
    }
}
```

You should always avoid code like the following:

```
public class Foo implements Thread{

    public Foo()
    {
    }

    public void run()
    {

        ...

        synchronized(this){

            ...

        }

        ...

    }
}
```



```
}  
}
```

### 5.3 A last word

I realize that these restrictions may annoy some of you. I strongly believe that following these restrictions will help the majority of the course become better at writing multi-threaded code. You have to write a few extra lines of code, but the benefits are well worth it. In particular, I doubt these restrictions will add as many as 50 lines of code that you have to write over the entire semester, so please bear with me.

Good luck.