


Traffic Signal Controller Concurrent Programming Project 1

16th September 2003



The above diagram is an illustration of a particular street intersection that I used to traverse daily. Each circle stands for a set of traffic signals that all behave identically at all times. For example, both northbound traffic lanes get the same signal, but the left and right lanes of eastbound traffic get different signals. Northbound and southbound traffic get different signals. The signals are named: S, controlling southbound traffic (which may also make a left turn); N, controlling northbound traffic (which may also make a left turn); EWR controlling east- and west-bound through traffic in the east-west right lanes; and EWL controlling left turns from the east-west left lanes.

Compatibility of movement: no two signals may be green or yellow at the same time. If they are, crashes will result.

For each of the signals there is a sensor with the same name that indicates the presence or passing of a vehicle. The sensor sends a pulse when a vehicle arrives at it.

The behavior of the signals is dictated by what is detected by the sensors. If there is no traffic, or only EWR traffic, the EWR signal stays green. When traffic arrives for another direction the sensor indicates this to the controller. The controller stops any incompatible movements by first signalling yellow for 6 seconds and then changing to red. When all conflicting movements are stopped, the controller changes the signal to green. Once green, a signal changes to yellow (and then red)

- no sooner than 6 seconds after turning green
- no later than 12 seconds after turning green if there is conflicting traffic
- 2 seconds after the last signal from its sensor provided this is within the 6..12 second period
- In addition to these rules there is a nearby traintrack that has to be accomodated (see below).

If sensors indicate that there are conflicting requests they are to be served in round-robin order, beginning with EWR, N, S, and finally EWL. Except for the EWR signals, lights should only turn green when sensors indicate traffic is present except as these rules are modified to handle the train.

The Train

Not shown on the diagram but nevertheless present, there are train tracks to the south of the intersection, parallelling the east-west roadway. A train-arriving sensor, named TA, and a train-departing sensor, named TD, are on the tracks. When an arriving train is detected the signals must allow northbound traffic to clear the tracks for a 10 second green period, regardless of the information from the N sensor and regardless of the current state. Of course any current incompatible signals must first be changed immediately to yellow for 6 seconds and then to red. (Note that for train arrivals the current green phase may last less than 6 seconds). After that, only ER and WR traffic can proceed until the train departs. You may assume that each train departs before the next train arrives. After the train leaves, the signal should pick up where it left off, either restarting the interrupted phase if it has traffic, or continuing with the next phase in round robin order that does have traffic.

Designing a solution

I encourage you to initially think about your design independently of the details of how we represent the signal lights and sensors in the program. What states have to be made mutually exclusive? How might you ensure mutual exclusion using the synchronization mechanisms that we've talked about?

Suggestion: start with a simpler intersection that has only a single signal controlling north-south traffic and a single signal controlling east-west traffic. If you come up with

good design principles you should be able to extend your solution to the more complex intersection in a straightforward way.

Decide what threads you are going to have and what conditions you need to signal between threads. Decide whether you are going to use monitors or semaphores. How many monitors or semaphores do you need? That is, how do the monitors and semaphores correspond to the principle elements of the problem, the signals and sensors. How will your program know when to take action based on the time specifications in the problem description?

Implementing signals and sensors

Your program will interact with the signal lights and sensors through its standard input and output.

Sensor input

Sensor input comes in the form of lines read from standard input. A line consists of the identity of the sensor (N, S, EWR, EWL, TA, TD) followed by a newline character (i.e. you can run your program in a terminal session and type sensor input to it). Example input:

```
N
S
EWR
EWL
TA
TD
```

Timing of the input clearly matters to the behavior of the signal.

Signal and sensor output

Output from your program will be a series of lines, each line containing a *timestamp* and either a *light command* to a single signal or an echo of a sensor input. *Timestamps* are integers representing the number of milliseconds since the start of the run. A *light command* line consists of 4 tokens: a timestamp, the letter L, the signal name (EWR, EWL, N, S) and a color indicator: G, Y or R. The command indicates that the named signal is to show the indicated color.

Sensor inputs must be echoed to the output as soon as they arrive. These lines consist of three tokens separated by whitespace: a timestamp, as above, the letter S, and the sensor name.

The output format allows a program that I will write to check whether the controller is operating properly. The output might also be used to drive a graphical display of the state of the intersection.

If you wish to interleave debugging or other information in output lines, feel free. Just make the first token a timestamp and the second token, which must be present, different from L and from S.

Example output

```
0 L N R
0 L S R
0 L EWL R
0 L EWR G
2250 S N
5500 S EWR
7500 L EWR Y
13500 L EWR R
13500 L N G
19500 L N Y
25500 L N R
25500 L EWR G
...
```

Assignment details

You may write your program in Java, C, or C++. If using C or C++ you will use the pthreads package. If using Java you'll start from the built-in synchronization facilities. If you want to program this for Windows please see me. We'll have to negotiate how to make the program testable.

Preliminary design document due Tuesday, Sept. 23: this should lay out your general approach to the problem, including what you've determined to be the basic synchronization requirements, areas that you see as problematic, what threads you are going to use, etc. (10%)

Preliminary code due Friday, Oct. 3 (under my door): this should illustrate that you have mastered the input and output requirements, the synchronization primitives, and the compiling and running steps for your choice of concurrent language. Working code for a 2-way intersection with sensors would represent good progress. (10%)

Final turn-in: Thursday, Oct 16th in class for printed material including code and a revised design document describing in detail what you built. Tell me what is the most important thing for me to look at in your design and code. Tell me what changed between the preliminary design and what you ended up with. Tell me how well it works. Does it meet the specification? I will also provide an on-line submission mechanism for your source code and executables. Your executable will be tested with programs that supply sensor input and check the output to make sure the signals work correctly.

Suggested plan

I suggest building this project incrementally: build working code for a series of simpler problems leading up to the final program. One possible series of increasingly complicated intersections would be:

1. A two-way intersection, call them NS and EW, using only timed lights – no sensors.
2. A two-way intersection that stays in say EW mode until traffic is sensed in the NS direction, stays in NS until traffic is sensed in EW. Try to get the minimum and maximum times working at this stage.
3. Split N and S.
4. Split EW into EWR and EWL corresponding to through traffic and left turns on EW road.
5. Add the train.

I think you can manage with just two threads: one to manage the signals, and the other to read the input.

You will find it helpful to describe the intersection in data structures that are interpreted by the code for the signals. This will vastly reduce the amount of code you have to write. You do not want to write, and I do not want to read, 150 lines of code for each different signal in this problem! As a point of comparison, my java code corresponding roughly to a combination of steps 4 above contains 30 lines to build the datastructure describing the intersection, 175 lines to describe the behavior of the signals including producing the output, and 20 lines to read the sensor input. I estimate that adding the train will cost less than 20 lines of code.

Do not be fooled by the small size of the code for this project: there is a great deal going on in this problem. You will have to spend considerable time thinking about the problem and about your plan for solution in order to even get started. You cannot put this off until the week before it is due and expect to succeed. The preliminary due dates are intended to both get you moving on the project, but also to give me a chance to provide feedback to keep you on the right track.