Assignment 3: Switched Networks

Due the week of October 1 & 2, 2015

This week, we will consider another approach to constructing a data network: switching. This will mark the beginning of our consideration of networks designs appropriate for wide area networks. The contention based protocols for Ethernet and wireless networks we considered last week are suitable mainly for local area networks. Switching techniques, however, can be used in networks that span much larger distances. At the same time, switching techniques are employed in local area networks. In particular, our introduction to switching will include its application in local area networks in the form of transparent bridges/switches. Modern Ethernets make heavy use of such switches. In a typical modern Ethernet each computer is attached to a two-way segment of cable with the other end connected to a bridge/switch. Because the cable is two-way, transmissions by the computer and switch do not collide with one another. In such a configuration, an Ethernet is no longer a contention-based broadcast network. It is instead an example of a switched network.

A significant issue addressed in switched networks is the problem of finding routes to use when delivering data. We will discuss several algorithms for choosing routes this week.

Please read Chapter §3.1 and §3.2 from Peterson and Davie. Note that for now we are skipping §3.2. This is a little odd because Chapter 3 is named "Internetworking" and the sections I am telling you to skip is the only one that is actually about internetworking. I promise, however, that we will come back and read §3.2 later. To save you a little time, while reading §3.1 you can skip the section on ATM networks (if you want to know why, reading this paper: http://dx.doi.org/10.1145/774749.774751).

In addition, I would like you to read a paper that discusses the details of some of the improvements made to the routing algorithms used in the ARPANET discussed in the text. (the ARPANET eventually morphed into the Internet).

- Khanna, Atul and John Zinky, "The Revised ARPANET Routing Metric," *Proceedings of the ACM SIGCOMM '89 Symposium on Communication Architectures and Protocols*, pp. 45-56.

The last two questions refer to this paper. You should read the questions before reading the paper! They are this week's open-ended questions.

The first few of this week's questions are short questions from the text. Yes, this is a good sign! In our meetings, we will use these questions to briefly discuss the basic operations of the protocols/algorithms described in the text. We will discuss the remaining questions in more depth.

In our organizational meeting, I urged you to take some time to think about how you would present you answers to the weekly exercises orally. During the process of grading your written work this week, I realized I need to urge you to do the same thing when completing your written work. Your answers to the weekly questions do not need to be long, but they should present a clear, logical explanation of your answer/opinion. Some of you did this for the first assignment, but others provided answers that failed to communicate your train of thought. Doing so is important.

Exercises

1. Complete parts a, b, and c of 3.1.
2. Complete problem 3.3.
3. Complete problem 3.13. Explain the process used, not just the final results.

4. Complete problem 3.17.

5. Complete problem 3.48.

6. Consider the problem of designing the network infrastructure for an office building. The building has four floors and a basement. There are computers located on all five levels. The organization using the building has assigned offices in such a way that members of the same department/group tend to have offices on the same level. So, they hope to maximize network performance by installing a separate LAN on each floor and interconnecting the LANs with bridges and a backbone network. The backbone network will support a higher data rate than the networks on individual floors. Therefore, if possible, all traffic between stations on distinct floors should be routed through the backbone.

   They want to build redundancy into the design. At the least, if any single bridge fails, they want all the LANs to remain interconnected.

   (a) A possible design for this network is shown below. As in the diagrams in the text, the little circles represent bridges and the straight lines are Ethernet segments. The goal is to use the bridges that connect networks on adjacent floors (i.e. the bridges on the right side of the diagram) only as backups in case one of the connections to the backbone fails. What id numbers should be assigned to the bridges in the network to provide the desired behavior?

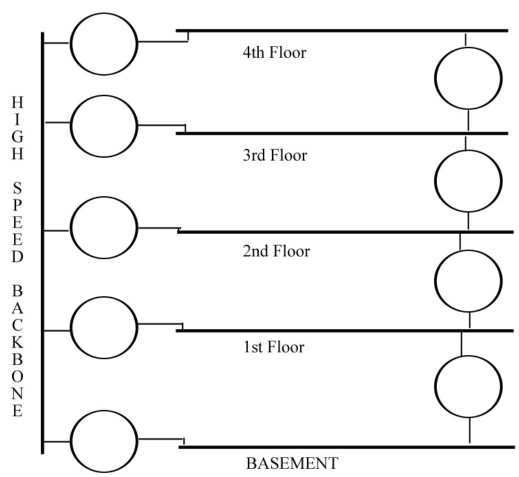

   (b) The preceding design only has one backbone. So, if one of the routers connecting a floor to the backbone were to fail, the network would have to route traffic from that floor to other floors through either the floor above or below where the failure occurred. To make the network even more tolerant of failure, we might consider adding a backup backbone. That is, we will add a second backbone network and connect it to the floor networks with bridges as shown below.

   The hope would be that if one of the bridges connected to the primary backbone failed, the traffic would be shifted to the added backbone. It would also be nice, of course, if during normal operation the two backbones could share the load. Explain why neither of these goals is attainable using the transparent bridge spanning tree algorithm.

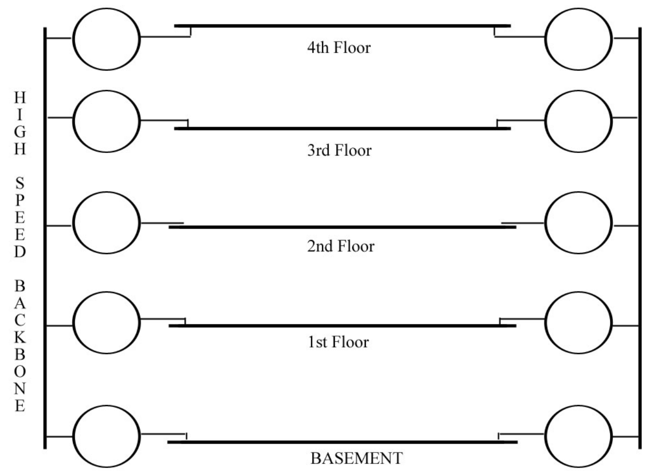

(c) Suggest a way to use one or more backbone networks together with bridges to interconnect the floor networks so that all inter-floor traffic will use a backbone network even in the event that a single bridge fails. That is, all the packets seen on the network serving a given floor should either come from or be destined to one of the computers on that floor.

Show both the topology and the bridge id number assignments.

7. For this question, I would like you to compare the routes used by a local network built using transparent bridges to a switched network with the same topology in which shortest path routing is used. To do this, consider the network used as an example in problem 3.13 from the text.

   (a) First, redraw the network as a graph in which:
      - each of the network segments shown in Figure 3.48 (i.e. A, B, C...) is a node, and
      - there is an edge between two nodes exactly when the corresponding network segments are attached to a common bridge.

   (b) Simulate Dijkstra's algorithm on the resulting graph to find the best routes from B to all other network segments. Assume all edge weights are 1.

   (c) Compare the route that would be used in this scenario to transmit packets from B to A and B to D to the routes determined by the spanning tree algorithm used by transparent bridges (i.e. the answer to the original question from the book). Which appears to be more efficient?

   (d) When the transparent bridge spanning tree approach is used, bridges that are not forwarding do not forward any packets. Their potential bandwidth is completely unused. The spanning tree you computed using Dijkstra's algorithm for part (b) does not use every bridge. However, this tree is only used when sending from B. Other nodes compute their own spanning trees and use those trees to determine their routes. When all the spanning trees used by all the nodes are considered, will any bridges go completely unused? Would this change if the weights on the edges were not all 1s?

8. For this week's last exercise, I would like you to analyze one of two aspects of the "Revised ARPANET Routing Metric" paper assigned as a reading. Two possibilities are described below. You get to pick the one you find most interesting.

   Here are your options:

(a) One nice consequence of the fact that the "Revised ARPANET Routing Metric" paper describes routing in a network that has vanished into the mists of history, is that it makes it clear that the fact that time marches on should be an inherent part of the design process of any network algorithm/protocol. What was good enough when the network was small (100s of computers) and slow (56Kbps lines), might work less well as the network grew larger (1000s of computer) and faster (10Mbps lines) and might be unworkable in today's network. Basically, good protocols must scale.

One concrete place to look for scaling issues is in the magic numeric constants that appear here and there in the description of a network protocol. For example, in the "Revised ARPANET Routing Metric" paper, the authors state that the reported cost for a 56 kb/s link ranges from 30 units to 90 units. Why 30 to 90? Why not 0 to 100? Are 30 and 90 universal constants of the network universe, or are they likely to require adjustment as the network grows larger or faster?

What I would like you to do for this option is make an annotated list of examples of such constants that appear in this paper. For each example, I would ask you to:

- identify the section of the paper in which the parameter is discussed,
- specify the value being used, and
- "briefly" explain the role the parameter plays in the protocol/algorithm.

For most of the examples you can find, that is all I ask. For two of the examples (any two you like), however, I would like you to explore whether the parameter values chosen by the authors is likely to require adjustment as the network grows. In particular, I would like you to discuss what if any changes you would expect to be necessary in the values used as the network grows larger and faster. For both examples you choose, you should justify your prediction. It would be particularly nice if you could suggest a procedure for identifying good values to use either analytically or experimentally.

(b) The Khanna and Zinky paper is a classic paper. It has been very widely cited. You will find in on innumerable course reading lists. The impact of the work it describes on the ARPANet was quite significant.

At the same time, it is possible to criticize the paper as an example of a bad scientific experiment. In a well designed experiment, one carefully varies one variable while keeping other variables fixed and measures the results of the changes. In their work, however, Khanna and Zinky modified many aspects of the existing routing mechanisms and evaluated the net effect of all the changes. For this analysis, I would like you to identify the potentially independent changes they made. That is, I would like you to give a brief list of changes to the routing process they proposed that each could have been made independently of the other changes they proposed. Continuing to follow the mantra of being "scientific", for each of the independent changes you identify I would like you to state your hypothesis of how you think network behavior might be impacted if this change were made independent of the other changes proposed in the paper. For example, one change might reduce the amount of traffic required to distribute routing updates without improving the actual routing of data packets while another might actually improve both aspects of routing. Briefly justify each of your hypotheses.