CS 740, Fall 1997
Assignment 5:
Part II: Constructing a Nonblocking Network

Assigned: Wed., Nov. 5
Due: Mon., Nov. 17

Policy

You may work in a group of up to 3 people in solving the problems for this assignment. You should turn in a single report for your entire group, identifying all of the group members.

Logistics

Any clarifications and revisions to the assignment will be posted on the class bboard and Web page. There is no electronic handin for this assignment.

Nonblocking networks

In this part of the assignment we are going to prove the existence of a multistage nonblocking network with N input ports, N output ports, $O(N \log N)$ switches, and $O(N \log N)$ edges.

In order to get started, we need to establish some definitions. A *multistage* network is one in which every switch can be labeled with an integer, called its *level*, such that a switch with level i is connected only to switches with levels $i-1$ and $i+1$ (if these levels exist). Without loss of generality, let us assume that the smallest level number assigned to a switch is 0, and that the largest level number assigned is L. In this case we say that the network has L levels of edges. We call the switches on level 0 the *input ports* and the switches on level L the *output ports*.

An example of a multistage network is shown below. This network is called a butterfly network. In this example, there are $N = 8$ input ports and also $N = 8$ output ports. The network has $L = \log_2 N = 3$ levels of edges, for a total of $2N \log_2 N$ edges.

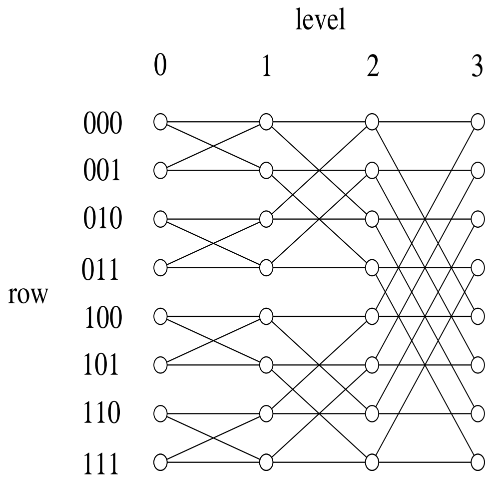

The goal throughout this assignment will be to route *switch-disjoint* paths between the input ports and the output ports. A set of paths is switch disjoint if no two paths pass through the same switch. We say that an input port has already been *used* if a path has been established from that input port to some output port, and similarly for an output port. A network is said to be *nonblocking* if no matter what paths have already been established, it is possible to establish a new path, disjoint from all previously established paths, between any unused input port and any unused output port.

Warming up

Begin by proving that the butterfly network is *not* a nonblocking network. In particular, show that there exists a butterfly network with $N = 2^k$ input ports, for some $k > 0$, such that there is a permutation π from the input ports to the output ports that cannot be routed with switch-disjoint paths. **Hint**: You won't need a very large value of k to find a counterexample.

Establishing the optimality of the contruction

Before we begin with the construction of the nonblocking network, you will first demonstrate that the size of the network, in terms of the number of edges, is optimal. In particular, you will show that any nonblocking network with N input ports and N output ports must have at least $\Omega(N \log N)$ edges. (This is actually an old theorem due to Shannon.)

Here is a rough outline of the proof. Your job is to fill in the details.

1. Count the number of different possible permutations from the N input ports to the N output ports.
2. Suppose that permutation π can be routed using a subset $E(\pi)$ of the edges of the network. (There may be other ways to route π as well.) Argue that if $E(\pi) = E(\pi')$, then $\pi = \pi'$. In other words, argue that it's not possible to route two different permutations using the exact same subset of edges. Recall that the paths must be switch disjoint, so that at most one incoming and one outgoing edge can be used at any switch.
3. Count the number of different subsets of edges in a network with $|E|$ edges.
4. Conclude that in order to route all possible permutations, $|E| = \Omega(N \log N)$.

Majority networks

Rather than constructing the nonblocking network in one shot, you will first show how to construct a *majority* network. A network with N input ports and $M > N$ output ports is called a majority network if no matter what paths have already been established between input ports and output ports, it is still possible to route a path from any unused input port to at least $(M/2) + 1$ different unused output ports. In other words, it is possible to route a path from any unused input port to a majority of the output ports. Note that a majority network is not necessarily a nonblocking network, because there may be some unused output ports that cannot be reached from some unused input ports.

Prove that a nonblocking network can be constructed using two back-to-back majority networks. This should be a short proof.

Bipartite expander graphs

The construction of majority networks will be based on a graph-theoretic construction called a *bipartite expander graph*. A graph is *bipartite* if its nodes can be partitioned into two sets U and V such that every edge has one endpoint in U and the other in V. Suppose that we think of the edges as being directed from U to V. If $|U| = n$ and $|V| = m$, then we call the graph an $n \times m$ bipartite graph.

For the purposes of this assignment, an $n \times m$ bipartite graph is called an α-*expander* if every subset of $n/2$ nodes of U has at least $(\alpha n/2) + 1$ distinct neighbors in V. The bipartite graphs used in this construction will have $m = 2n$ and $2 < \alpha < 4$, and the degree of each node in U will be exactly d, where d is a fixed constant.

To show that these bipartite expander graphs actually exist (which is not obvious), you will use a nonconstructive proof technique called the probabilistic method. In particular, you are going to show that if you choose the edges of the graph at random, then with nonzero probability you get the expander that you are after.

Prove that for any n, and $m = 2n$, and any fixed constant α between 2 and 4, there is a fixed integer constant $d > 0$ such that if every node in U has d randomly chosen neighbors in V, the probability that the resulting graph is an α-expander is nonzero. To help you with this proof, I've provided some hints.

1. Rather than showing that the probabiltiy of success is greater than 0, show that the probability of failure is less than 1.

2. Even a gross overcounting of the probability of failure will work in this argument. It suffices to sum the probability that a set of $n/2$ nodes in U has $\alpha n/2$ or fewer neighbors over every possible choice of $n/2$ nodes in U. Furthermore, for any particular set of $n/2$ nodes in U, it suffices to sum the probability that all of the edges out of these nodes have endpoints in V in a particular set $\alpha n/2$ nodes in V, over all possible sets of $\alpha n/2$ nodes in V.

3. Each node in U should choose d neighbors in V. Each neighbor should be chosen randomly, uniformly from among all of the nodes in V, and independently of the other choices made by the node or by any other node.

4. To simplify your calculations, use the handy inequality
$$\binom{a}{b} \leq \left(\frac{ae}{b}\right)^b,$$
where a and b are integers, $0 < b < a$, and e is the base of the natural logarithm.

Recursive construction of a majority network

Now that you've proved the existence of α-expander graphs, you're in a position to prove that majority networks exist, from which it will folow that nonblocking networks exist.

The construction of a majority network will be recursive. Recall that the network has N input ports and M output ports, where $M > N$. Let us assume that $N = \epsilon M$, where ϵ is a fixed constant between 0 and 1, e.g., something like $1/10$.

Your job is to fill in the details in the following proof.

1. Begin with the base case. Show how to contruct a majority network with $N = 1$ and $M = 1/\epsilon$. How many edges does your network have? What makes this case so easy?

2. For the inductive step, assume that you know how to construct a majority network with $N/2$ input ports and $M/2$ ouput ports. You construct a majority network with N input ports and M output ports by gluing two of the $N/2$-input-port networks together with an expander graph. In particular, suppose that an $N/2$-input-port network can be constructed using l levels of edges. You place two copies of this network in parallel, to form a new network with N input ports, and l levels of edges. Note that these two subnetworks are entirely disjoint, and not connected in any way. To connect them, you add a final level of M switches to the network. These will be the M output ports of the network. The edges in the final level are specified as follows. Insert edges from the $M/2$ output ports of each of the two subnetworks to the M output ports on the last level so as to form a bipartite α-expander graph.

   Prove that, for the right values of α and ϵ, this construction gives a majority network. This step is the crux of the entire proof. You'll have to use the fact that the last level of edges forms a pair of α-expanders, that $M = N/\epsilon$, and that at most one path can originate at any input port.

Bounding the size of the nonblocking network

Finish off the proof by analyzing the number of edges, the number of switches, and the number of levels in the nonblocking network. State your bounds using big-O notation.

Extra credit: Bounding the in-degree of each node

One downside to the construction of majority networks that you have just been through is that although you bounded the number of edges leaving each switch by d (except possibly at the first level), you did not bound the number of edges entering each switch. Each switch randomly and independently chooses its neighbors at the next level, and if you're unlucky, many of them might choose the same neighbor. It turns out that the randomized construction can be refined to limit the in-degree of each node as well. For the case of $m = 2n$, the idea is to randomly choose a set of edges subject to the constraint that the out-degree of each switch is d, and the in-degree is $d/2$. Because of this contraint, the choices of the neighbors for different switches are no longer independent, which complicates the counting argument. For extra credit, prove that for sufficiently large d the resulting graph will be an α-expander.