Lecture 5

1 Overview

We will give a couple of networks that can be used for routing messages. We will illustrate their use in implementing a BSP computer that was described in the previous lecture.

2 A couple of networks.

Today, we will talk about a particular parallel network; the butterfly and its brother the hypercube. The hypercube in n dimensions consists of a set of nodes labeled with the binary n-bit strings, thus it has 2^n nodes. Each node is connected to any node whose label differs in exactly one bit. (Would you advise drawing your own figure here.)

To route a message from node $a = a_1, \ldots, a_n$ to $b = b_1, \ldots, b_n$ proceeds by "correcting" the bits in some order, say from bit 1 to bit n. Of course, there are many (shortest) paths depending on the order of bit corrections, and of course many many non shortest paths.

A butterfly is the network that "embodies" this routing in the hypercube. Here, the nodes are labeled with a pair (i, x) where i is an integer in $\{0, \ldots n\}$ and x is a n bit string. One can view the first index as the "row" and the second as the "column". The edges basically form a line on the row, i.e., (i, x) is connected to $(i + 1, x)$. We also have "hypercube" edges. That is $(i, x_1, \ldots x_n)$ is connected to $(i, x_1, \ldots \overline{x_i} \ldots x_n)$. That is, at the ith place you can flip the i bit of your row. (Here a figure would be drawn in class.)

Typically, for the butterfly network, the first node in a row is viewed as the input and the last node is viewed as an output. One can also collapse the rows, and obtain the hypercube. Routing from an input $(1, x)$ to (n, y) by correcting bits in order corresponds to the canonical shortest path routing algorithm above.

3 What is totally excellent?

Are these good routing networks?

- Small degree. Easy to build.
- It has short paths. Any graph with degree at most k, has diameter at least $\log_k N$. The hypercube diameter is close to optimal, i.e., $\log N$ versus $\log_{\log n} N = \log N / \log \log N$. (De Bruijn networks have optimal diameter.) The butterfly has very near optimal diameter, $O(\log N)$.

A binary tree has both properties. What else? Perhaps, you want two properties. Hard to cut? No, the binary tree is easy to cut. When you route many messages one would like

not to congest any link too much. Is this true? Not for the binary tree, routing messages across the root is very bad.

How about the butterfly. (For now, we leave the hypercube.) Hard to cut? We will address this later.

First, we restrict the routings a bit. Say a single message starts at each input and is destined for each output. That is, no output is the destination of more than one message, and no input is the destination of more than one message.

Let's consider the following pattern? Route a message from each input $(1, x0, 1y)$ to $(n, y0, 1x)$ where x is the first half (actually $n-1/2$) of the bits and y is the second (actually $n-1/2$) of the bits. This pattern is one to one as required.

THEOREM 1
The transpose pattern gives congestion $\Omega(\sqrt{N})$ on an N-input butterfly, using "greedy routing."

This theorem follows by cutting the butterfly at level $n/2$. This leaves $2^{n-1/2}$ separate $22^{n-1/2}$ input butterflies on the left and on the right, each of size $2^{n-1}/2$. The ith sub butterfly on the left have a common $(n-1)/2$ bit prefix representing i for the row identifier, while the jth sub butterfly on the right has a common suffix j. There is a single edge from the ith sub butterfly on the left to the jth sub butterfly on the right. But, all the messages from the ith sub butterfly on the left are destined for the ith sub butterfly on the right? (Note actually there is a $i, 0$ sub butterfly on the left and a $i, 1$ on the left, and similarly for the right.)

This causes a congestion of $2^{n-1/2}$ on a single edge. Yet, most edges in the middle layer are empty? Can we do better?

4 Benes Routing

A Benes network is a butterfly network connected to a reveres ed butterfly network. Or one, can view it as a network that goes up the butterfly and back down. This network gives many natural shortest paths from input. One for each of N intermediate outputs.

This choice allows us to overcome the problem from the last section.

THEOREM 2
Any (partial) permutation can be routed on a Benes network, with no congestion.

PROOF:

Notice the following recursive nature of the Benes network. If one removes all the inputs and outputs, that one obtains two separate Benes networks. We note that for a trivial Benes network (i.e., one node) it is easy to route a permutation. Do nothing!!

So, now we reduce the problem of routing a permutation on an N input Benes network to one of routing on a permutation on a $N/2$ input Benes network. What to do?

Well, at the "first" step, we must choose to route message at input i (and destined for say j) to the "top" sub-Benes or to the bottom. This decision then forces us to route the packet to a particular output of the sub-Benes based on its destination j.

We wish to coordinate the decisions so that

- Each input of a sub-Benes is chosen by a single input of the original.
- Each output of a sub-Benes is forced to be used by a single message.

Either is easy to enforce. For example, only two inputs of the original can choose any particular input of a sub-Benes and there are two sub-Benes networks. Same for outputs. But the choice at the inputs *forces* the choice at the outputs.

We encode these dependencies in a bipartite graph each side of which contains $N/2$ nodes. For each message, which starts at $i = i_1 \dots i_n$ and is destined for $j = j_1 \dots j_n$, we add an edge from node $i_2 \dots i_n$ to node $j_1 \dots j_{n-1}$. This encodes the fact, that either original input $0i_2 \dots i_n$ or $1i_2 \dots i_n$ can choose input $i_2 \dots i_n$ of a sub-Benes network. A similar statement hold for the outputs.

The degree of this graph is at most two. If we can partition the edges (messages) into two sets each of which induces degree at most 1 in the graph, we can assign each of the two sets of messages to the two sub-Benes networks in a manner where the induced message source, destinations is a permutation.

Clearly, since the degree of the graph is two, the graph consists of a set of cycles. Moreover, the length of the cycles is even. Thus, we can partition the cycles into two sets to meet this degree condition.

□

Question 1: Prove that separating any k of the inputs from any k outputs in the Benes network requires one to cut at least k edges. (That is, it is much "harder to cut" than a binary tree, or for that matter planar graphs.)

4.1 Some comments, or perhaps the main point.

This along with much of modern engineering science was invented at Bell Labs in the fifties. It is the basis of routing in the telephone network switches called Clos networks (which are essentially high degree butterflies.)

This particular observation was the basis of the parallel computer built by IBM in the eighties to do quantum chromodynamics computations. This computation proceeded as did the relaxation algorithms of the previous lecture. The pattern of communication in that lecture as well as in QCD computations is fixed in every step. Thus, one could implement the BSP model by precomputing the switch settings in the network.

For example, with very fast switches (say as fast as the processors) and p computers, this network essentially implemented the BSP model with g being one. And L being $\log p$.

One can also view the old-style CRAY supercomputers in this light. They had tremendous processor to memory networks. We will look at this more later.

5 On-line solutions: Valiant/Brebner routing.

Benes routing in some sense says that the hypercube/butterfly are great networks. On the other hand, the algorithm is not very "parallel", i.e., for each pattern of messages one needs to figure out how to route them.

Here we consider a really parallel algorithm for routing permutations in the Benes network. Route to a random intermediate destination in the first butterfly. Then route greedily to the final destination. This method was proposed by Valiant and Brebner for the hypercube. (Valiant has done many truly amazing things. This being merely a comment of wide reaching consequences.)

THEOREM 3
The maximum number of packets that use any edge is $O(\log n)$ with high probability.

What is the expected number of messages that use any edge? Consider the ith level in the first butterfly. The number of inputs that could use it are 2^i.

That is the load can be expressed as the sum of 2^i $0-1$ random variables. The number of outputs it can get to is 2^{n-i} (within a factor of 2). The probability of any of the inputs choosing this output is thus $2^{n-i}/2^n = 2^{-i}$ Thus, the expected number of messages is 1.

Question 2: Show that with high probability that max load is $O(\log n)$. (Hint: use a chernoff bound.)

Question 3: (Optional) Show that with high probability that max load is $O(\log n/\log\log n)$.