Finding the bisection width of a graph is NP-complete. But for nice graphs such as hypercubes, multidimensional arrays and several others, we can compute it exactly. The upper bounds are usually given by inspection. For some graphs such as paths, or the complete binary tree, or the complete graph, lower bounds are easily argued. But in general lower bounds are trickier. It may be possible to make an *exchange argument*, i.e. consider a candidate minimum bisection, argue that getting a new bisection by swapping a pair of vertices should only increase the cost, and using this to get a characterization of the optimal bisection. This argument is often messy, even for simple graphs such as say $P_n \Box P_n$. It turns out that it is possible to get lower bounds by looking at the eigenvalues of certain matricess associated with the graph, which we will study later.

A very elegant strategy is based on estimating ratios of bisection widths of graphs.

Theorem 1 *Suppose a graph G on n vertices is embedded in a graph H with load 1, and congestion C. Suppose B_G, B_H denote the bisection widths of G, H respectively. Then $B_H \geq B_G/C$.*

Proof: Let E' denote a set of edges bisecting H into H_1, H_2 such that $|V(H_1)| = \lfloor n/2 \rfloor$, and $|V(H_2)| = \lceil n/2 \rceil$, and $|E'| = B_H$. Consider the subgraphs G_1, G_2 of G induced by the vertices embedded into H_1, H_2 respectively. This is a bisection of G, though not necessarily a minimum bisection. There must be at least B_G edges with one endpoint in G_1 and another in G_2. All such edges must be embedded into paths that pass through E' at least once. Thus the total congestion in E' must be at least B_G. But the congestion of any edge is at most C, hence the total congestion is $C|E'| = CB_H$. Hence $CB_H \geq B_G$, and hence the result follows. ∎

Suppose we somehow know the bisection width of G. Then by embedding G into H, we get a lower bound on B_H using the above theorem! The most suitable candidate for embedding is the complete directed on n nodes. The reason for choosing the complete directed graph rather than the complete graph will become clear later.

Corollary 1 *Let H be a graph on n nodes, and B its bisection width. Suppose the complete directed graph on n nodes is embedded into H with load 1, such that the (unidirectional) congestion in each edge is at most C. Then $B \geq \frac{1}{C} \left\lfloor \frac{n}{2} \right\rfloor \left\lceil \frac{n}{2} \right\rceil$.*

Proof: The proof of Theorem 1 applies with G the complete directed graph, B_G the size of the bisection counting each edge only in one direction. Let $G_1, G_2, H, H_1, H_2, E', B_H$ be as before. Then we have $B_H \geq B_G/C$. Noting $B_G = \lfloor n/2 \rfloor \lceil n/2 \rceil$ the bound follows. ∎

1 Hypercube Bisection Width

We begin by embedding a n node complete directed graph into the k dimensional hypercube Q_k, where $n = 2^k$. For this we use canonical paths, i.e. the ones obtained by correcting bits lsb to msb.

Let u, v be two nodes in Q_k. Let $u = u_{k-1} \dots u_0$ denote the bits of u, and $v = v_{k-1} \dots v_0$. Consider the sequence $P(u, v)$ obtained as we change the bits of u to the bits of v, one bit at a time from lsb to msb, i.e.

$$P(u, v) = \{u_{k-1} \dots u_0, u_{k-1} \dots u_1 v_0, \dots, u_{k-1} v_{k-2} \dots v_0, v_{k-1} \dots v_0\}$$

Assume that the elements of $P(u, v)$ are numbered 0 through $k - 1$. Then the ith element of $P(u, v)$ is either identical to the $i + 1$th element, or differs in bit i only. Thus the elements of $P(u, v)$ define a path from u to v. This is the so called canonical path. This is how the directed edge in K_n from u to v is embedded.

Example: for $u = 1101$ and $v = 0110$, we have $P(u, v) = 1101, 1100, 1110, 1110, 0110$. From this if we get the path $1101, 1100, 1110$, 0110.

So let us estimate the congestion of an edge from $w = w_{k-1} \dots w_0$ to its neighbour z across dimension i. Clearly $z = w_{k-1} \dots w_{i+1} \overline{w_i} w_{i-1} \dots w_0$

Suppose a path from u to v uses this edge. Then we know that w, z must occur consecutively in $P(u, v)$. Further, w, z differ only in the ith bit. But the only possible consecutive vertices in $P(u, v)$ that can differ in the ith bit are the ith and $i + 1$th. Thus these must equal w, z respectively! Thus we know that

1. $w = u_{k-1} \dots u_{i+1} u_i v_{i-1} \dots v_0$
2. $z = u_{k-1} \dots u_{i+1} v_i v_{i-1} \dots v_0$

Or alternatively, $w_{k-1} \dots w_i = u_{k-1} \dots u_i$, and $w_{i-1} \dots w_0 = v_{i-1} \dots v_0$. Note further that the ith bit of z is $\overline{w_i}$. Thus $v_i = \overline{w_i}$.

Thus the fact that $P(u, v)$ goes through w, z constrains how u, v can be chosen. The most significant $k - i$ bits of u are required to agree with those of w. Thus there are only 2^i ways to choose the remaining bits for u, and those are the different possible choices for u. The least significant i bits of w, v must agree and the ith least bit must differ, thus there are $k - i - 1$ bits which can be chosen arbitrarily to decide v. Thus there are 2^{k-i-1} ways in which v can be chosen. Thus the number of possible pairs u, v can be chosen in $2^i \cdot 2^{k-i-1} = 2^{k-1}$ ways. Thus the edge (w, z) has congestion 2^{k-1}. But this applies to any edge. Thus the unidirectional congestion is uniformly 2^{k-1} in all edges.

Thus the bisection width is $n^2/4C = 2^{2k}/(4 \cdot 2^{k-1}) = 2^{k-1} = n/2$. This is precisely the number of edges along any dimension, and hence there is a matching upper bound as well.

This argument is adequate to give good lower bounds on bisection widths of many, many networks.

Exercises

1. Show that $H = P_r \Box P_c$ where $r \leq c$ has bisection width r without embedding the complete directed graph. *Hint:* Suppose H_1, H_2 is an optimal bisection. Start by arguing that H_1, H_2 need not contain non-consecutive vertices in any column or row.

2. Show that $P_r \Box P_c$ where $r \leq c$ has bisection width r by embedding the complete directed graph. Compare this proof with the preceding one.

3. Show that the converse of Theorem 1 is not true.

4. *(I dont know the answer to this.)* Are there classes of graphs for which the converse might be true? Say perhaps vertex transitive graphs? Say not the exact converse, but something like the converse?

5. Consider the graph obtained by attaching $P_{n^2/3}$ to the center of the longer side of $P_n \Box P_{2n/3}$. Give an upper bound on the bisection width of this graph. Get a lower bound by embedding a complete graph as above. Get a lower bound by embedding a graph G consisting of a $K_{5n^2/6}$ to which is attached $P_{n^2/6}$. Argue a bound on the bisection width of G from first principles.

6. Suppose I want remove minimum number of edges to partition Q_n into one subgraph having 2^k vertices, $k < n$, and the rest. Show that it is possible to do this by removing $2^k(n-k)$ edges and that at least 2^k edges must be removed. The lower bound is also based on embedding the complete directed graph; perhaps it can be improved by embedding some other graph.

7. Suppose each vertex u in the 2^n node hypercube Q_n sends a message to a vertex $\pi(u)$ where π is a permutation, i.e. $\pi(u) = \pi(v)$ only if $u = v$. Suppose canonical paths are used for sending the messages. Show that there exists π such that some edge will have congestion $\Omega(2^{n/2})$. The term congestion in this context means the number of messages going through the edge.