1.4 Communication Networks Considered

1.4.1 Hypercubes

Definition

In general, a hypercube has $p = 2^d$ nodes, where d is the dimension of the cube. We introduce hypercubes by showing how we can build large hypercubes from small ones, in Figure 1.1.

An immediate consequence of the indexing scheme shown in the figure is that in a hypercube two nodes are neighbors if and only if the binary representation of their indices differ in exactly one bit.

We now give a definition of a hypercube network:

Definition 1 *The communication network is a hypercube has the property that two nodes $\mathbf{P}_i$ and $\mathbf{P}_j$ are nearest neighbors if and only if the binary representations of i and j differ in at most one bit.*

Properties

Given two nodes $\mathbf{P}_i$ and $\mathbf{P}_j$, it is relatively easy to generate a path between them: Consider the binary representations of i and j:

$$\begin{array}{llll} i \rightarrow & i_{d-1} & \cdots & i_1 \quad i_0 \\ j \rightarrow & j_{d-1} & \cdots & j_1 \quad j_0 \end{array}$$

One algorithm for finding a path between the two nodes is

do $k = 0, \ldots, d-1$
  if $i_k \neq i_j$
    send message from
      $i_{d-1} \quad \cdots \quad i_k \quad j_{k-1} \quad \cdots \quad j_0$
    to
      $i_{d-1} \quad \cdots \quad i_{k+1} \quad j_k \quad \cdots \quad j_0$
  fi
od

In most hypercubes architectures this routing is done in hardware, using pipelining as discussed previously.

It is not hard to see that a shortest path will only consider each dimension of the hypercube once, i.e. each bit is toggled at most once in a routing algorithm. Notice that there are up to $d!$ different shortest paths between two nodes. If the order in which the bits are considered is always fixed, the resulting routing algorithm is often called dimensional routing.

The *network diameter*, defined as maximum over all node pairs of length of the shortest path between them is only the dimension of the hypercube.

1.4.2 Linear Arrays and Rings

Definition

A network forms a linear array if there exists a link between $\mathbf{P}_i$ and $\mathbf{P}_j$ if and only if $i = j + 1$ or $j = i + 1$. If in addition there is a link between $\mathbf{P}_{p-1}$ and $\mathbf{P}_0$, then the linear array becomes a ring. If each link in the ring is only unidirectional, e.g. there is a unidirectional link from $\mathbf{P}_i$ to $\mathbf{P}_{(i+1) \bmod p}$, then the ring is said to be a unidirectional ring.

dim. d	# of nodes p	comment	
0	1	A single node.	
1	2	Duplicate the 0 dimensional cube, and connect corresponding nodes with a link.	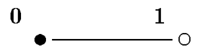
2	2^2	Duplicate the 1 dimensional cube, and connect corresponding nodes with a link.	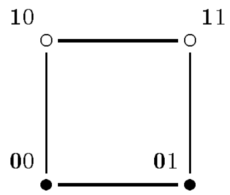
3	2^3	Duplicate the 2 dimensional cube, and connect corresponding nodes with a link.	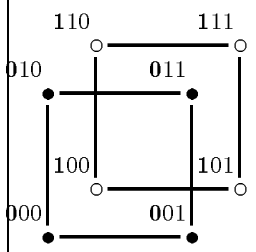
$\vdots$	$\vdots$	$\vdots$	$\vdots$
d	2^d	Duplicate the $d-1$ dimensional cube, and connect corresponding nodes with a link.	$\vdots$

Figure 1.1: Construction of hypercubes

Properties

Notice that under our assumptions that messages traveling in opposite directions on a link do not conflict, all linear arrays are also unidirectional rings.

The network diameter for a linear array is $p-1$, and for a ring is $\lfloor p/2 \rfloor$. **check this!**

1.4.3 Meshes and Tori

A natural extention to linear arrays and rings are higher dimensional meshes and tori.

Definition

A network is said to be a two dimensional mesh if there is an indexing of the nodes $\mathbf{P}_{i,j}$, $0 \leq i < r$ and $0 \leq j < c$ such that P_{i_1,j_1} is a neighbor of P_{i_2,j_2} if and only if one of the following is true

$$
\begin{aligned}
i_1 = i_2 &\quad \text{and} \quad j_2 = j_1 + 1, \quad \text{or} && (1.3)\\
i_1 = i_2 &\quad \text{and} \quad j_1 = j_2 + 1, \quad \text{or} && (1.4)\\
j_1 = j_2 &\quad \text{and} \quad i_2 = i_1 + 1, \quad \text{or} && (1.5)\\
j_1 = j_2 &\quad \text{and} \quad i_1 = i_2 + 1, \quad \text{or} && (1.6)
\end{aligned}
$$

Parameters r and c will be referred to as the row and column dimensions of the mesh, with $p = r \times c$. i and j are the row and column indices, respectively. If in addition, there are "wrap-around" connections like for the ring, the mesh becomes a torus. A unidirectional torus is a torus with unidirectional links.

The order in which message are routed is often fixed. We will assume that all messages take the shortest path between two nodes, traversing the x-direction first, followed by the y-directions. In other words, a message between two nodes will travel from the source within the row in which the source lies to the column in which the destination lies, after which it is routed to the destination within that column.

This definition naturally generalizes to higher dimensions.

Properties

Notice that a $2 \times 2 \times \cdots \times 2$ mesh is a hypercube. Under our model, any mesh can be viewed as a unidirectional torus.

1.5 Viewing Topologies as Higher Logical Meshes

1.6 MPI