Advanced Computer Architecture (0630561)

Lecture 17

Static Interconnection Networks

Prof. Kasim M. Al-Aubidy
Computer Eng. Dept.

INs Taxonomy:

- An IN could be either static or dynamic. Connections in a static network are fixed links, while connections in a dynamic network are established on the fly as needed.
- Static networks can be further classified according to their interconnection pattern as one-dimension (1D), two-dimension (2D), or hypercube (HC).
- Dynamic networks can be classified based on interconnection scheme as bus-based versus switch-based.

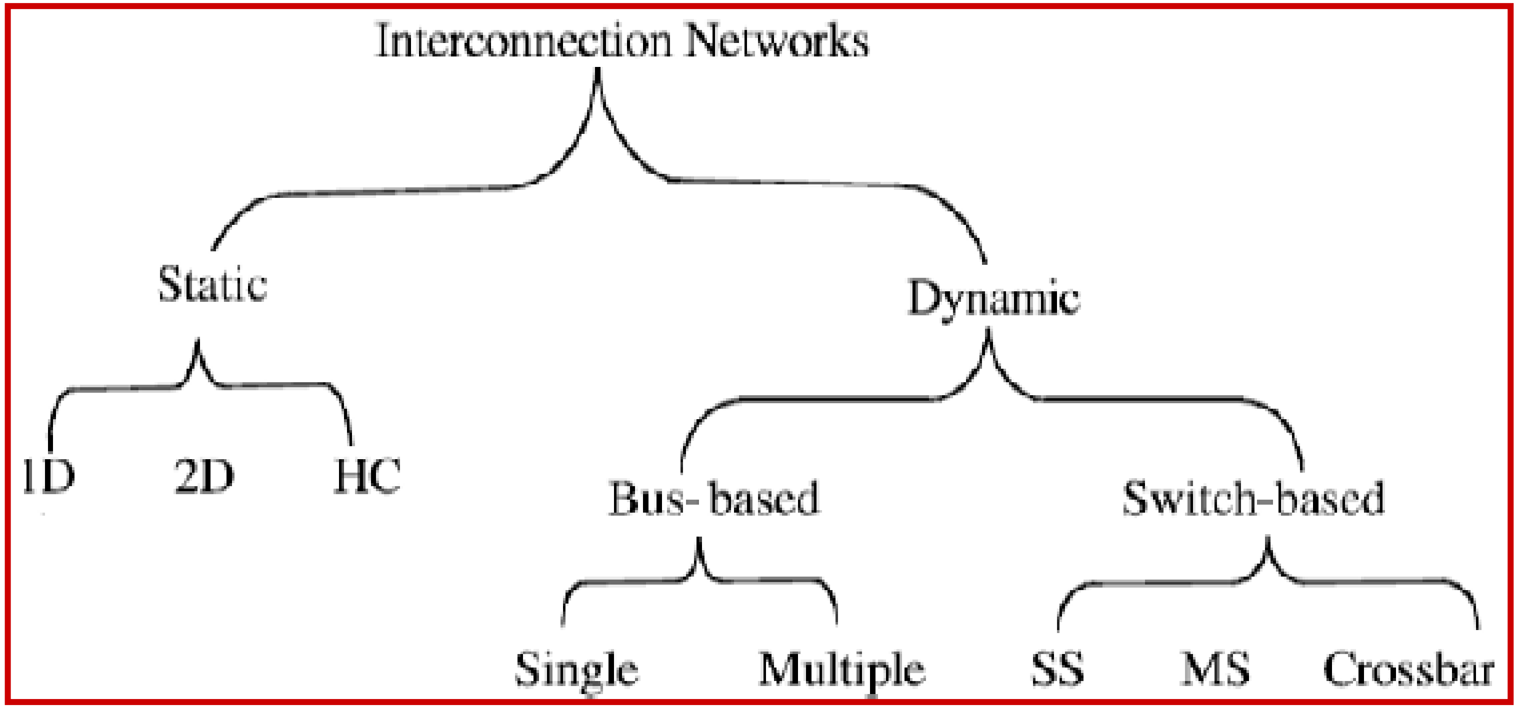

Static Interconnection Networks:

- Static (fixed) interconnection networks are characterized by having fixed paths, unidirectional or bidirectional, between processors. Two types of static networks can be identified;
  - Completely Connected Networks (CCNs),
  - Limited Connection Networks (LCNs).

1. Completely Connected Networks:

- In a completely connected network each node is connected to all other nodes in the network.
- CCNs guarantee fast delivery of messages from any source node to any destination node (only one link has to be traversed).
- Every node is connected to every other node in the network, therefore routing of messages between nodes becomes a straightforward task.

CCN Characteristics:

- CCNs guarantee fast delivery of messages from any source node to any destination node. Only one link has to be traversed.
- CCNs are expensive in terms of the number of links needed
- for their construction. This disadvantage becomes more and more apparent for higher values of N.
- The number of links in a CCN is given by N(N-1)/2.
- The delay complexity of CCNs, measured in terms of the number of links traversed as messages are routed from any source to any destination is constant, that is, O(1).

Example:

CCN having N=6 nodes. A total of 15 links are required in order to satisfy the complete interconnectivity of the network.

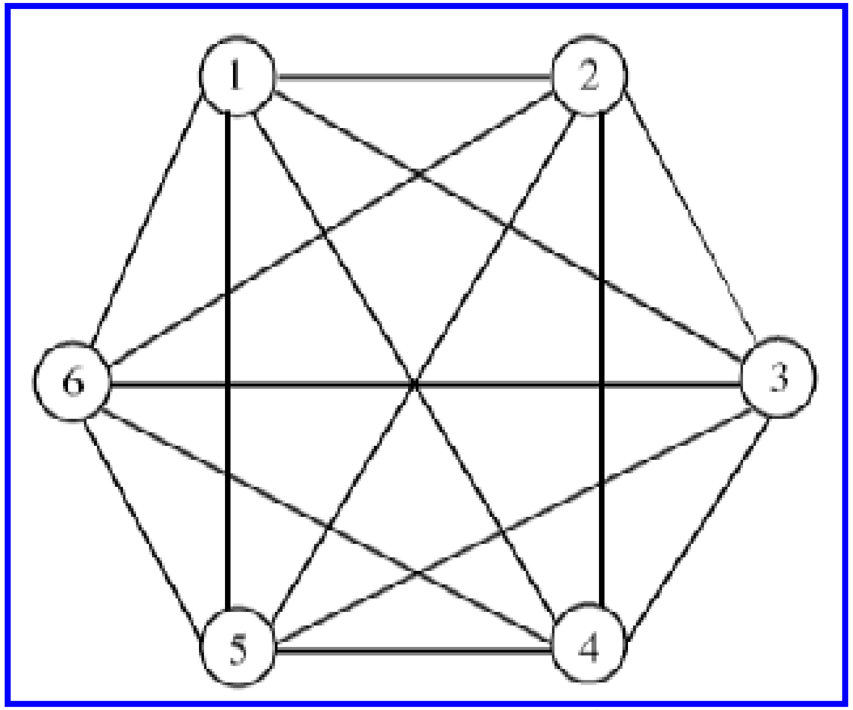

Limited Connection Networks:

- Limited connection networks (LCNs) do not provide a direct link from every node to every other node in the network. Instead, communications between some nodes have to be routed through other nodes in the network.
- The length of the path between nodes, measured in terms of the number of links that have to be traversed, is expected to be longer compared to the case of CCNs.
- Two other conditions seem to have been imposed by the existence of limited interconnectivity in LCNs, these are:
  - the need for a pattern of interconnection among nodes, and
  - the need for a mechanism for routing messages around the network until they reach their destinations.
- A number of regular interconnection patterns for LCNs include:
  - . linear arrays;
  - . ring (loop) networks;
  - . two-dimensional arrays (nearest-neighbor mesh);
  - . tree networks; and
  - . cube networks.

Examples of Static LCNs

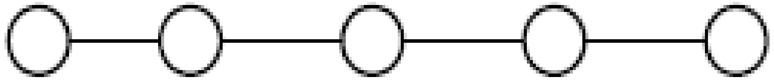

linear array network

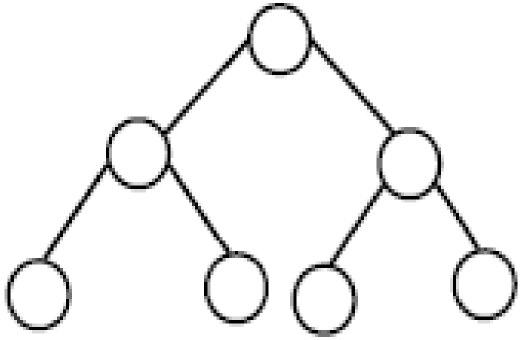

tree network:

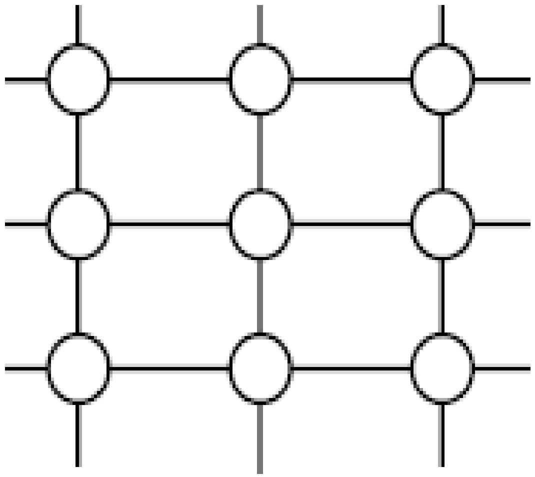

two-dimensional
array (mesh) network

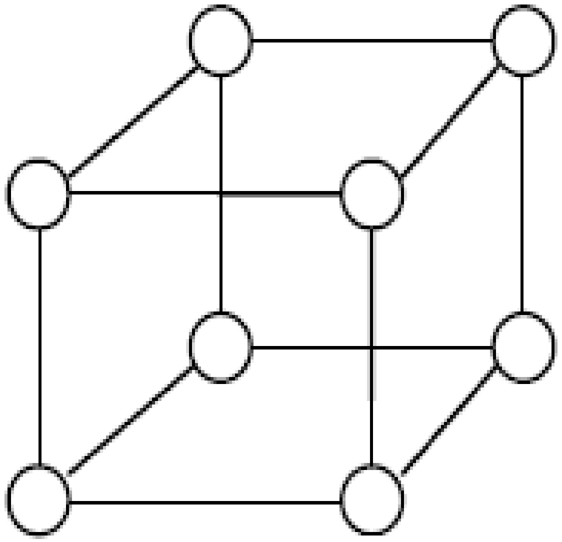

three-cube network

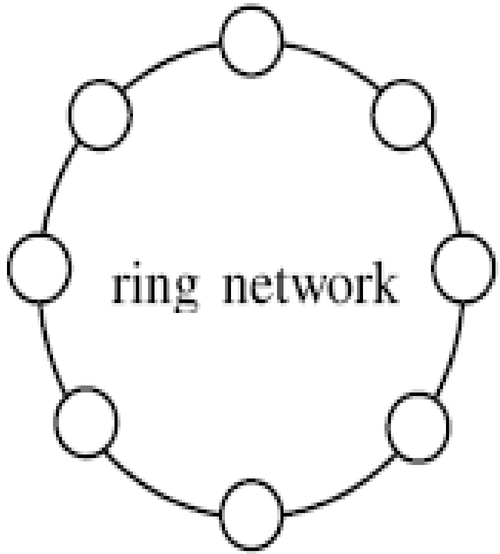

Linear Array Static LCN:

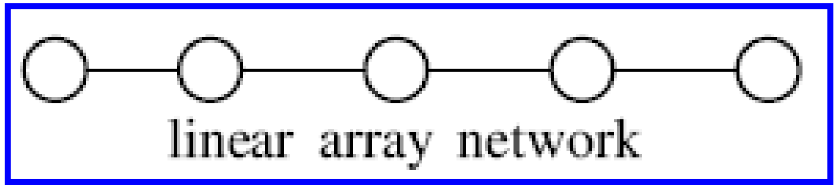

- Each node is connected to its two immediate neighboring nodes.
- If node i needs to communicate with node j, j>i, then the message from node i has to traverse nodes i+1, i+2, . . . , j-i.
- In the worst possible case, when node 1 has to send a message to node N, the message has to traverse a total of N-1nodes before it can reach its destination.
- Linear arrays are simple in their architecture and have simple routing mechanisms.
- Linear arrays are slow. This is particularly true when the number of nodes N is large.
- The network complexity of the linear array is O(N) and its time complexity is O(N).
- If the two nodes at the extreme ends of a linear array network are connected, then the resultant network has ring (loop) architecture.

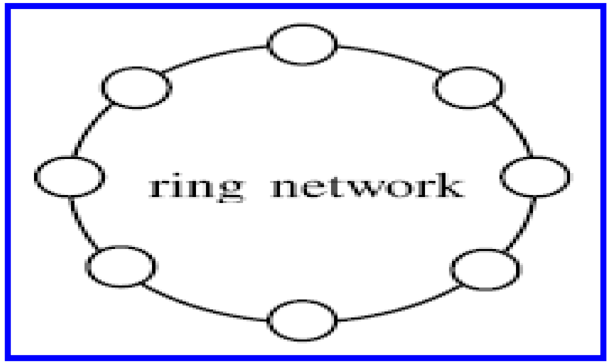

Tree Network Static LCN:

- If a node at level i needs to communicate with a node at level j, where i>j and the destination node belongs to the same root's child subtree, then it will have to send its message up the tree traversing nodes at levels i-1, i-2, . . . , j+1 until it reaches the destination node.
- If a node at level i needs to communicate with another node at the same level i (or with node at level j≠i where the destination node belongs to a different root's child subtree), it will have to send its message up the tree until the message reaches the root node at level 0. The message will have to be then sent down from the root nodes until it reaches its destination.
- The number of nodes (processors) in a binary tree system having k levels can be calculated as:

$$\boxed{\begin{aligned} N(k) &= 2^0 + 2^1 + 2^2 + \cdots + 2^k \\ &= \frac{(2^k - 1)}{2 - 1} = 2^k - 1 \end{aligned}}$$

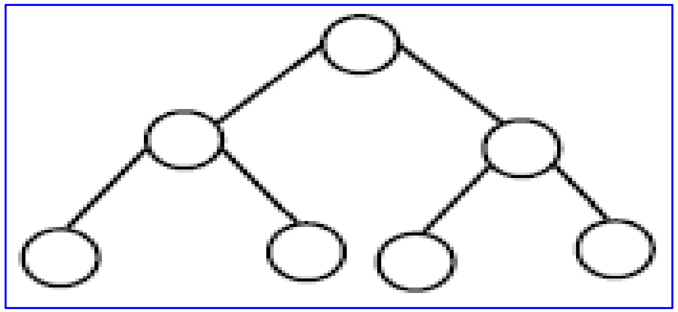

- The maximum depth of a binary tree system is $[\log_2 N]$, where N is the number of nodes (processors) in the network.
- The network complexity is O(2↑k) and the time complexity is $O(\log_2 N)$.

Cube-Connected Networks:

- An n-cube (hypercube of order n) is defined as an undirected graph having (2↑n) vertices labeled 0 to (2↑n) -1 such that there is an edge between a given pair of vertices if and only if the binary representation of their addresses differs by one and only one bit.
- In a cube-based multiprocessor system, processing elements are positioned at the vertices of the graph. Edges of the graph represent the point-to-point communication links between processors.
- Each processor in a 4-cube is connected to four other processors.

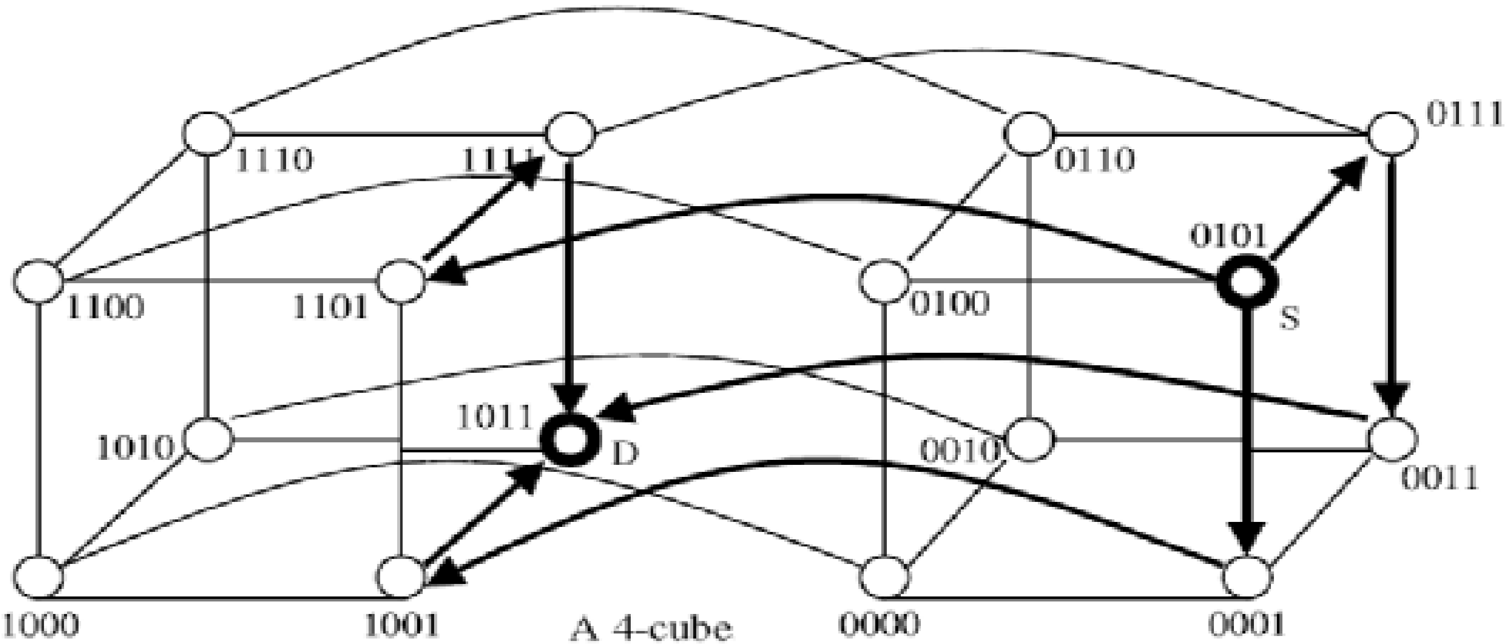

A 4-cube

- The route of a message originating at node i and destined for node j can be found by XOR-ing the binary address representation of i and j.
- If the XOR-ing operation results in a 1 in a given bit position, then the message has to be sent along the link that spans the corresponding dimension.

Example: If a message is sent from source (S) node 0101 to destination (D) node 1011, then the XOR operation results in 1110.

  - The message will be sent only along dimensions 2, 3, and 4 (counting from right to left) in order to arrive at the destination.
  - Once the message traverses the three dimensions in any order it will reach its destination.

- The hypercube is referred to as a logarithmic architecture. This is because the maximum number of links a message has to traverse in order to reach its destination in an n-cube containing N= (2↑n) nodes is log2N=n links.

Mesh-Connected Networks:

Example: 3*3*2 mesh network.

- A node whose position is (i, j, k) is connected to its neighbors at dimensions i+1, j+1, and k+1.
- A number of routing mechanisms have been used to route messages around meshes. One such routing mechanism is known as the dimension-ordering routing.
- A message is routed in one given dimension at a time, arriving at the proper coordinate in each dimension before proceeding to the next dimension.

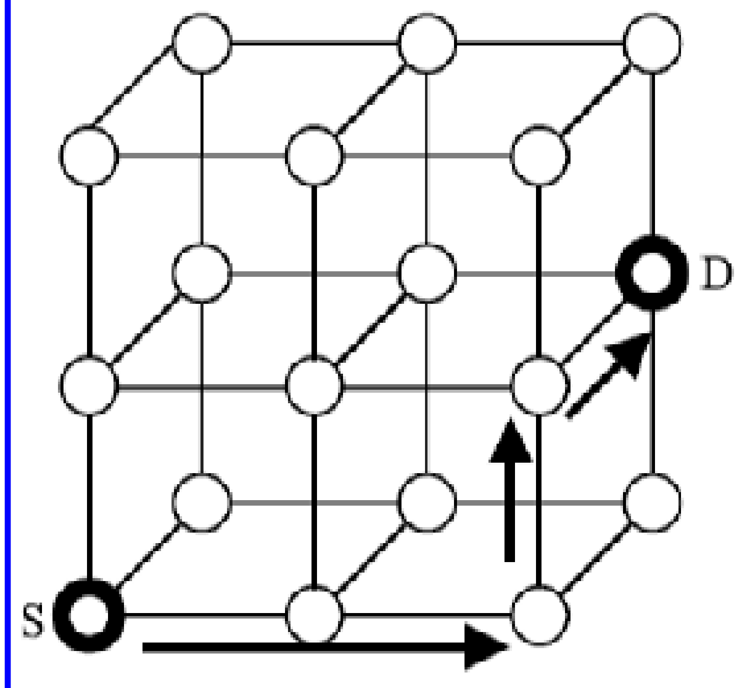

A 3 × 3 × 2 mesh network

- For a 3D mesh, messages are first sent along the i dimension, then along the j dimension, and finally along the k dimension.
- The route of a message sent from node S at position (0, 0, 0) to node D at position (2, 1, 1).
- For a mesh interconnection network with N nodes, the longest distance traveled between any two arbitrary nodes is O(√N).