CMSC416: Introduction to Parallel Computing

Topic: Parallel Networks and Filesystems
Date: April 18, 2024

High-speed interconnection networks

- A high-speed interconnection network is important to supercomputers and HPC (high-performance computer system) clusters. The network provides low latency and high bandwidth connections.
- Bandwidth is the amount of data that can be transferred per second.

- The different methods of connecting the nodes form different topologies.

- The longest and shortest distance determines network diameter. Distance or hop(s) is the number of cable(s) a switch must travel to reach another switch.
- Ex: If a router, A, is directly linked by a cable to another router, B, then the shortest distance to travel from A to B is one hop.

- Radix is the amount of ports on a router/switch.

Network components

- A network interface controller (NIC) or network interface card provides connectivity capability between devices in a network.

- Both terms, switch and router, can be used interchangeably.
- It connects multiple networks together; it can be connected to other routers.
- Switches are connected by network cables.

- There are two types of wires/cables: copper and optical.
- Copper cables cannot transfer data over long distances compared to optical cables (due to signal loss). However, copper cables are cheaper to produce.
- The length of the cable has no impact on the performance.

N-dimensional network

- Each switch has a few nodes connected to it (usually one).
- In an N-dimensional network, each switch is directly linked to 2n switches.

- **Trous** networks are networks that contain wrapped-around links.
- “An ideal torus network uses the same number of nodes along each dimension to maintain balance” (https://clusterdesign.org/torus/).
- Examples from https://en.wikipedia.org/wiki/Torus_interconnect :

1. 1D Torus: In one dimension, nodes are connected in a closed loop, forming a simple circle. Communication can occur in two directions: +x and −x. Essentially, a 1D torus is equivalent to a ring interconnection.
2. 2D Torus: In two dimensions, nodes are laid out in a rectangular lattice with n rows and n columns. Each node is connected to its four nearest neighbors (+x, −x, +y, and −y). The total number of nodes in a 2D torus is n^2.
3. 3D Torus: In three dimensions, nodes form a three-dimensional lattice resembling a rectangular prism. Each node is connected to its six neighbors (+x, −x, +y, −y, +z, and −z). Each edge of a 3D torus consists of n nodes.

Fat-tree network

- It optimizes communication efficiency.
- A Fat-tree network has a radix of k, and each router has k/2 nodes.
- A pod is a group of k/2 switches (at each level), and the max number of pods = k.