CS 498
Hot Topics in High Performance Computing

Networks and Fault Tolerance

8. Advanced Network Topologies

Intro

- What did we learn in the last lecture
  - Introduction to network topologies
  - Parallel sorting
- What will we learn today
  - Topology metrics (diameter, bisection bandwidth, cost, …)
  - Common network topologies

Network Topology Graph

- We now introduce a graph-theoretic view of network topology
- A network is a directed (unidirectional) or undirected (bidirectional) graph G=(V,E)
  - Direct networks have PEs at each vertex
  - Indirect networks have two types of vertices: pure forwarding vertices and vertices with PEs
  - Edges denote the physical or virtual links between nodes in the network, might be weighted

Diameter of Graphs

- Definition: "The maximum length of the shortest path between any two vertices in G"
  - Only defined for strongly connected digraphs or connected undirected graphs
- Class Question: What are the diameters of: bidirectional ring, unidirectional ring, 2d array, 3d array?

Properties of the Diameter

- Diameter indicates:
  - Maximum transmission delay
  - Maximum power consumption to transmit a packet
  - Rough cost of the interconnection network
- Average Distance matters too
  - Average path length for all node pairs
  - Average delay and power consumption
- We focus on diameter for simplicity

Minimum Diameter Directed

- Switches have fixed out-degree d
- Theorem: Let G with n=|V| be a strongly connected digraph with fixed d, then the diameter

$$d(G) = \begin{cases} = n - 1: & d = 1 \\ \geq \lceil \log_d(n(d-1)+1) \rceil - 1: & \text{otherwise.} \end{cases}$$

- Proof:
  - From any vertex, at most d vertices can be reached at distance 1 and for i≥1, at most d^i vertices can be reached at distance i

Minimum Diameter Directed

– Let d(G) = k, then, n ≤ 1+d+d^2+…+d^{k-1}+d^k

$$= \begin{cases} = k+1: & d = 1 \\ \frac{d^{k+1}-1}{d-1} & \text{otherwise.} \end{cases}$$

– for d=1, n ≤ k+1, i.e., d(G)=k ≥ n-1 and d(G) ≤ n-1

– for d>1, (d-1)n ≤ d^{k+1} – 1, i.e.,

$$d(G) = k \geq \lceil \log_d(n(d-1)+1) \rceil - 1$$

– Q.e.d.

Minimum Diameter Undirected

- Theorem: Let G be a connected undirected Graph with n=|V|>2 and d>1, then

$$= \begin{cases} = 2k-1: & d=2 \\ \frac{d(d-1)^k-2}{d-2}: & d>2 \end{cases}$$

- Proof (Moore bound):
  - Let d(G)=k
  - At most d vertices can be reached at distance 1 and at most $d(d-1)^{i-1}$ can be reached at distance i

Minimum Diameter Undirected

- It follows $n \leq 1+d+d(d-1)+\ldots+d(d-1)^{k-1}$

$$= \begin{cases} = 2k-1: & d=2 \\ \frac{d(d-1)^k-2}{d-2}: & d>2 \end{cases}$$

  - For d=2, $n \leq 2k+1$, i.e., $d(G) = k = \lfloor \frac{n}{2} \rfloor$
  - For d>2, $(d-2)n \leq d(d-1)^k-2$ which implies

$$d(G) \geq \left\lceil \log_{d-1} \frac{n(d-2)+2}{d} \right\rceil$$

- Q.e.d.

Degree-Diameter Problem

- Show table of graphs
- Examples:
  - Peterson Graph
    - degree=3
    - d=2
    - P=10
  - Hoffman-Singleton Graph
    - degree=7
    - d=2
    - P=50

Bisection (Band)width

- Definition: “The bisection with is the minimum number of links that need to be cut to bisect the graph into two equal partitions”
- Class Question: What are the bisection widths of: 1d array, ring, 2d array, 3d array?

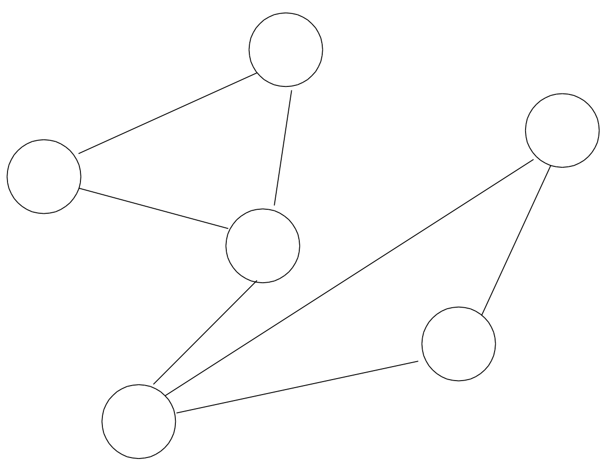

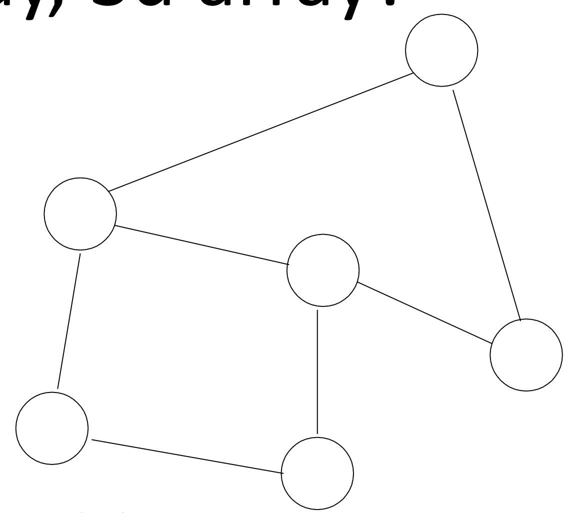

Computing the Bisection Width

- Minimum bisection is NP-hard
  - Proof in Garey&Johnson MINCUT (omitted here)
- Approximation algorithms are available (Kernighan-Lin, also METIS or SCOTCH)
- Most regular network structures can be derived analytically
  - We will discuss this soon for some topologies

More Metrics

- **Degree**: maximum degree of any vertex in the network
- **Cost**: number of total links in the network
- **Connectivity**: minimum number of edges that need to be removed to disconnect the network
  – Class Question: What does this metric indicate?

Fully Connected

- Examples: many SMP systems (POWER7, most modern x86 multicores)
- Metrics
  - Construction: K(P)
  - Diameter: 1
  - Degree: P-1
  - Bisection width: P^2
  - Connectivity: P-1
  - Cost: P*(P-1)/2

Arrays and Meshes

- 1-d, 2-d, 3-d arrays (with and without wrap-around) – direct network!
  - Not necessarily cubic but we will limit our analysis to cubic networks for simplicity!
  - Parameters: n-dimensions, k-PEs in each dim.
  - Examples: 1-d array with P=4, 2-d array with P=4, 3-d array with P=8!
  - Class Question: diameter, degree, bisection width, connectivity, cost (all exact with regards to $P=k^d$!)

Arrays and Meshes Metrics

- d-dimensional arrays with P Pes, $P=k^d$
- Metrics
  - Diameter: $d \cdot \sqrt[d]{P}$
  - Degree: $2d$
  - Bisection width: $k^{d-1} = \sqrt[d]{P}^{d-1} = P^{\frac{d-1}{d}}$
  - Connectivity: d
  - Cost: $d \cdot P - d \cdot k^{d-1}$

Hypercubes

- Examples: Cosmic Cube from Caltec, iPSC/2 from Intel, Connection Machine
  - Generally used as direct network
- Recursive construction: Q(d) – d-dimensional Hypercube
  - Q(0) = single PE, Q(1) = line, Q(2) = array, …
  - Show construction!
  - Class Question: diameter, degree, bisection width, connectivity, cost (all exact with regards to $P=2^d$!)

Hypercubes

- Q(n) → $P=2^d$
- Metrics
  - Diameter: $\log_2(P)$
  - Degree: $\log_2(P)$
  - Bisection width: P/2
  - Connectivity: $\log_2(P)$
  - Cost: $P*\log_2(P)/2$

Generalizing to k-ary n-cubes

- Examples: Cell B.E., Intel's SCC, Cray supercomputers, BG/L, BG/P
  - Direct network
- "square" n-dimensional with k PEs in each dimension
  - Examples: 2-ary 2-cube, 2-ary 3-cube, 3-ary 2-cube
  - Class Question: What other name do you know for 2-ary n-cubes, n-ary 2-cubes, n-ary 3-cubes?

Generalizing to k-ary n-cubes

- Examples: Cell B.E., Intel’s SCC, Cray supercomputers, BG/L, BG/P
  - Direct network
- “square” n-dimensional with k PEs in each dimension
  - Examples: 2-ary 2-cube, 2-ary 3-cube, 3-ary 2-cube
  - Class Question: What other name do you know for 2-ary n-cubes, n-ary 2-cubes, n-ary 3-cubes?
  - Yes, Hypercube and 2-D Torus, 3D Torus

Kautz Graphs

- Minimum diameter (dir.): $k \geq \lceil \log_d(n(d-1)+1) \rceil - 1$
  - Class Question: What is the minimum diameter topology that you could build right now?

Kautz Graph

- Minimum diameter (dir.): $k \geq \lceil \log_d(n(d-1)+1) \rceil - 1$
  - Class Question: What is the minimum diameter topology that you could build right now?
  - Yes, 2-ary n-cubes aka. Hypercubes with $k = \log_2(n)$
- Kautz graphs can fix the disparity!
  - Reaches smallest directed diameter possible!
    - Can be constructed easily (cf. degree-diameter graphs)
  - Degree-k Kautz graph is k-connected
  - Omitted definition for brevity

Kautz Graph Example

- Diameter: 3
- Degree: 2 (2x2 unidirectional)
- P=12

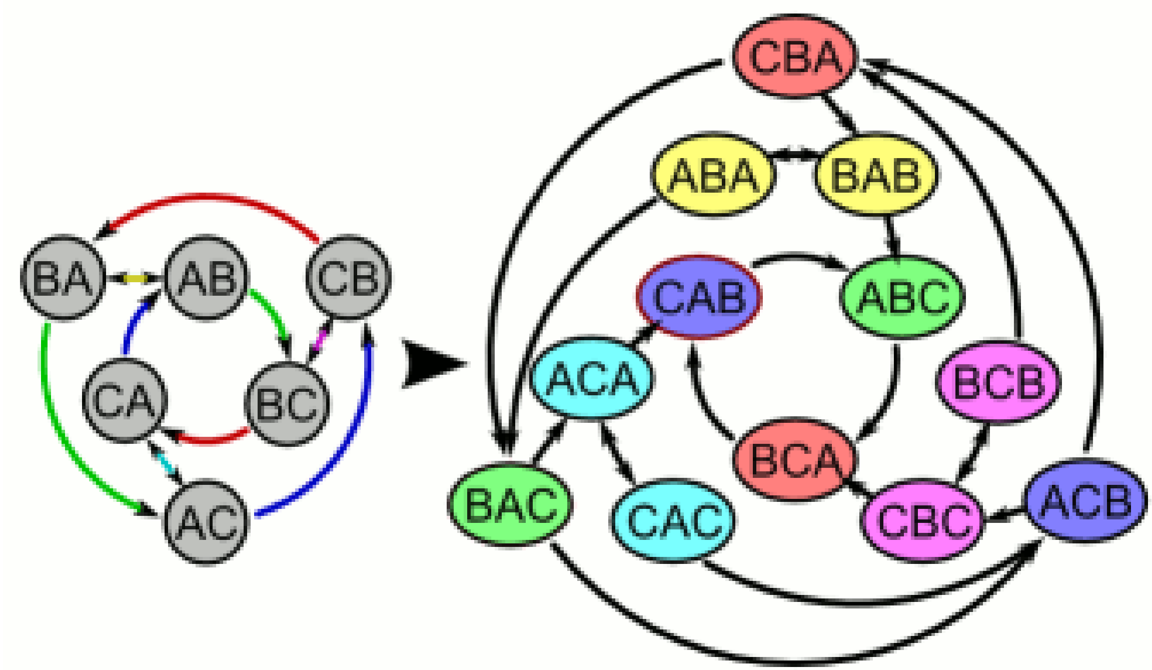

Butterfly Network

- Examples: CM5
  - It is a directed indirect network!
- k-ary n-fly: each switch has k inputs and k outputs and the network has n stages
  - Show example 2-ary 3-fly!
  - Class Question: What is the total number of PEs in a k-ary n-fly?

Butterfly Network

- Examples: CM5
  - It is a directed indirect network!
- k-ary n-fly: each switch has k inputs and k outputs and the network has n stages
  - Show example 2-ary 3-fly!
  - $P=k^{n-1}$
  - Class Question: How many paths are there from one process to one other process?

Butterfly Network

- Examples: CM5
  – It is a directed indirect network!
- k-ary n-fly: each switch has k inputs and k outputs and the network has n stages
  – Show example 2-ary 3-fly!
  – $P=k^{n-1}$
  – There is one path between two processes!
  – Class Question: diameter, degree, bisection width, connectivity, cost (all exact!)

Butterfly Metrics

- Class Question: diameter, degree, bisection width, connectivity, cost (all exact!)
  - $P=k^{n-1}$
  - Diameter: $\log_2 P$
  - Degree: 2+2
  - Bisection width: P/2
  - Connectivity: 2
  - Cost: $2*P*\log_2 P$

Benes Networks

- Two 2-ary butterflies back-to-back
  - Increases path-diversity
- Class Question: What are the metrics?

Benes Networks

- Two 2-ary butterflies back-to-back
  - Increases path-diversity
- Metrics:
  - Diameter: $2*\log_2 P$
  - Degree: 2+2
  - Bisection width: P/2
  - Connectivity: 2
  - Cost: $4*P*\log_2 P$

Folded Butterfly (similar to Clos)

- Fold Benes network in the middle
  - Utilizes bidirectional links
  - Retains path diversity
- Class Question: What are the metrics?

Folded Butterfly (similar to Clos)

- Fold Benes network in the middle
  - Utilizes bidirectional links
  - Retains path diversity
- Metrics:
  - Diameter: $2*\log_2 P$
  - Degree: 4
  - Bisection width: P/2
  - Connectivity: 2
  - Cost: $2*P*\log_2 P$ (bidirectional)