“A Case for Random Shortcut Topologies for HPC Interconnects” by Koibuchi et al. (koibuchi12)

An Trieu, Rosie Vuong

Introduction

- Problem: scaling parallel application increases communication latencies on performance
- Solution to lower latency: a topology of switches that has low diameter and low average shortest path length

  → Random shortcut topology can do this

- Random shortcuts: add random edges in network topologies for HPC systems

Related work

- Distributed Loop Networks (DLN)

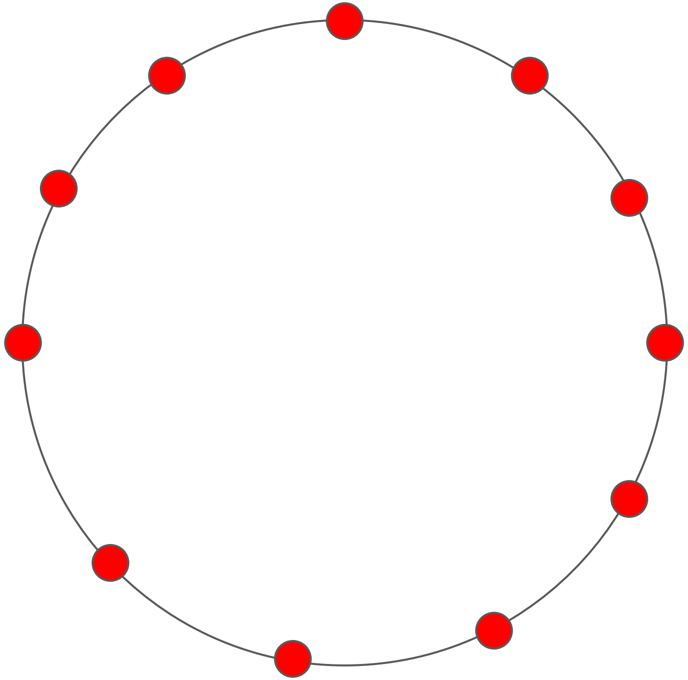

Random Shortcut Topology

- Reduces both diameter and average shortest path length, compared with non-random topologies with the same degree
- Achieve comparable throughput and have lower latency compared to non-random topologies (hypercubes, tori)
- Lead to robustness to random edge removals due to a small-world effect
- Good way to generate: add random shortcuts to simplest base topology (ring) using a simple uniform distribution

Scalability

- Difference between non-random topologies and DLN-2-y is positive

→ DLN-2-y (random shortcut topology) has advantage as the topology gets larger

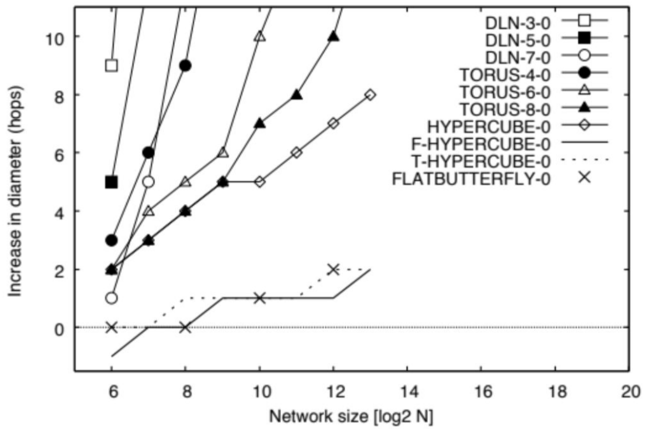

Figure 4. Diameter increase over DLN-2-y when using non-random topologies, where y is chosen so that comparisons are between topologies of the same degree, vs. N.

Scalability

- DLN-2-y also has advantage for average shortest path length when the topology gets larger
- Lower diameter leads to lower end-to-end path length → latencies between pairs are uniform → better for task allocation

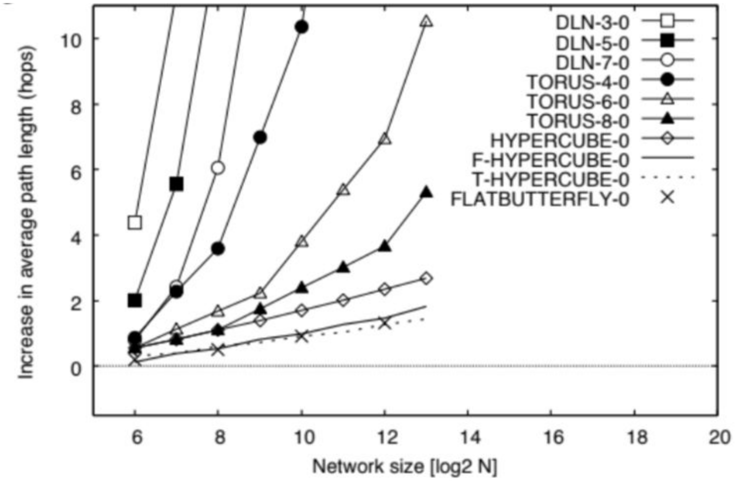

Figure 5. Average shortest path length increase over DLN-2-y when using non-random topologies, where y is chosen so that comparisons are between topologies of the same degree, vs. N.

Fault Tolerance

- DLN good fault tolerance result: diameter increase by 2, 30% edges are removed is sufficient in practice
- Non-random topologies: use custom deadlock-free routing algorithms, not robust in failed links

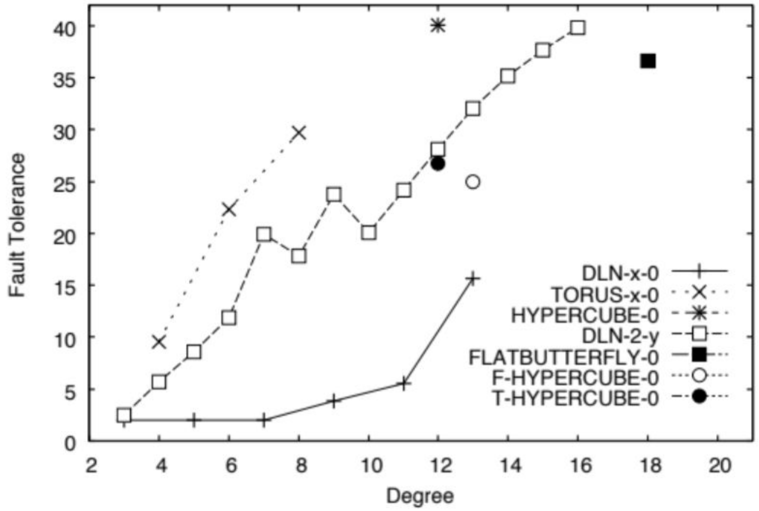

Figure 6. Fault tolerance vs. degree ($N = 2^{12}$ vertices).

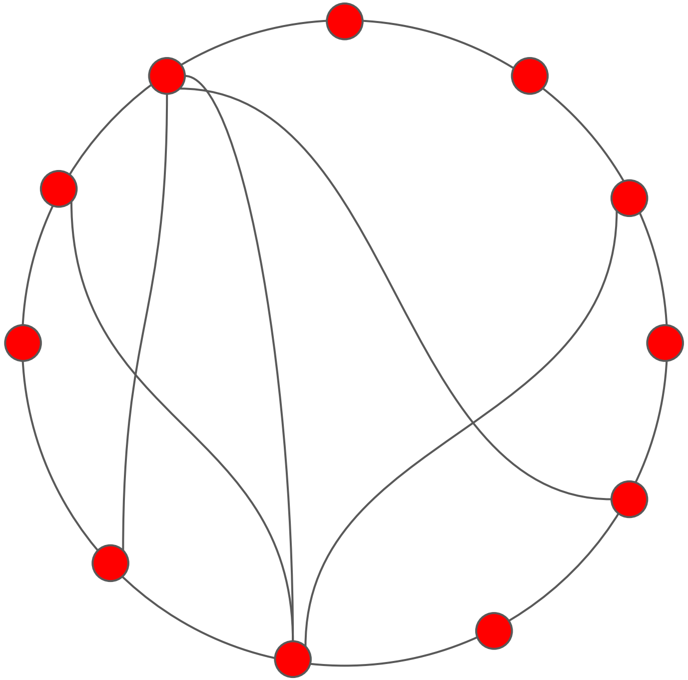

Discussion Of Limitations

- Case Study: Random Shortcut Links on Myrinet-Clos
    - Well distributed paths
    - Relatively low diameter.

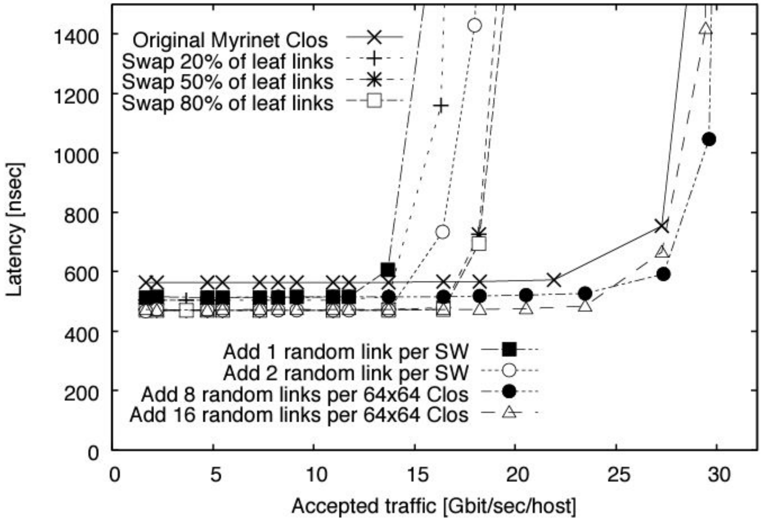

Figure 26. Latency vs. accepted traffic for random shortcut patterns on Myrinet Clos (80 switches, 256 hosts).

Discussion Of Limitations

- Routing Scalability
  - No Scheme → does not have a simple structure.
  - The scale of random shortcut topologies can be limited by routing table size at eachs witch.
- Physical Cable Length and Maintainability