Real World Evaluation of the Slim Fly Network

Edward Brees

Bet You Haven't Heard of This Before

- Network Topologies
 - The need for efficient data transfer and communication in high-performance computing
 - Latency, cost, and power consumption
- Slim Fly (SF)
 - Proposed solution with theoretical (key word) benefits of low-diameter
 - This implies low latency and cost savings



From Theoretical to Reality

- Slim Fly's benefits have not been empirically tested in real-world environments.
- Research paper objective
 - To deploy an at-scale SF network and rigorously evaluate its performance, cost, and power efficiency.
- Approach
 - Implementation of open-source cabling and physical layout routines to facilitate practical deployments.



Deployment Challenges and Implementation

- Challenges
 - Ensuring accurate cabling to maintain SF's low-diameter benefits
 - Managing cost and power efficiency while keeping practical scalability
- Implementation
 - Physical layout was optimized to reflect Slimfly setup.
 - $\circ \qquad {\sf Developed \ open-source \ cabling \ tools \ to \ assist \ in \ error-free \ installation.}$



Proposed Multipathing

- FatPaths as a basis
 - Utilizes multiple paths to avoid congestion and evenly distribute traffic.
- Improvements made upon it (Layering)
 - Achieves improved load balancing through layered routing approaches.
 - Incorporates adaptive routing algorithms that are aware of network state.

Ь	aput : Network topology $G = (V, E)$, number of layers $ L $
R	esult : A set of L routing layers
/	/ $W \in \mathbb{R}^{N_r \times N_r}$ contains weights of links; p is a priority gueue, with entries being pairs of nodes
1 W	= init link weight matrix() // Set all matrix entries to 0
2 p	= init_p_queue(G) // Each node pair gets the same priority = [F] // Laver 0 contains all the links (F)
5 L	$= \{L_{j} / j \}$ Layer 0 concarns are the rinks (L_{j})
4 10	r t = 1 to $ L - 1$ do
5	<pre>imit_layer(l) // Initialize the next layer as empty</pre>
6	$node_pairs = copy_pairs(p)$
7	while <i>node_pairs</i> $\neq \emptyset$ do
8	$pair = node \ pairs.dequeue()$
9	path = find path(G, W, pair, l)
0	if valid(nath) then
	undate priorities(nath n)
	update_priorities(pain, p)
2	update_weights(path, W)
3	add_path_to_layer(path, G, l)
4	end
5	end
~	chu



Implementing Physical Multipathing

- Final tweaks and adjustments
 - Accounted for practical problems such as link failures, deadlocks and traffic variations
 - Tailored FatPaths routing for optimized performance with SlimFly characteristics
- Physical Implementation
 - Load balancing is achieved by distributing traffic over multiple equal-cost paths.
 - Systematic approach to physical and logical network design enabled efficient routing.
 - Validated the FatPaths approach through real-world measurements and experiments.

Theoretical Analysis of Paths

- RUES and p%
 - RUES (Random Uniform Edge Selection) is used to measure theoretical equitable traffic handling
 - p% indicates the fairness in path usage at equilibrium
 - 'This Work' is unadjusted physical observed data
- Path Lengths
 - Theoretical shortest path lengths remain consistent, facilitating predictability in performance.
- Path Distribution
 - Analyses show a favorable distribution of paths, which is essential for load balancing.
- Path Diversity
 - High path diversity is inherent in SF topology, contributing to robustness against failures and congestion.





Results and Insights From Analysis

- Throughput
 - \circ \qquad Total amount of data transmitted through the network per unit time.
- SF's max throughput outperforms standard FatPaths in most cases, indicating potential for high efficiency in data handling.



Empirical Evaluation

- Comparisons
 - Slimfly was shown to surpass traditional 2-Level 0 Fat-Trees in both latency and throughput.
 - Slimfly was tested across diverse HPC and deep learning Ο workloads.
- Microbenchmarks were used to present SF's raw performance on basic network tests.









Empirical Results and Conclusions Made

- Comparison Benchmarks
 - SF showcases general improvements in common HPC tasks and scientific simulations over other methods
- Deep Learning Uses
 - Showed significant speedups in distributed deep neural network training over traditional topologies.
- Scalability & Cost Analysis
 - Offers generally higher scalability and cost-effectiveness, especially in larger installations, but becomes disproportionate as size and scale increase.

Questions?