Gigaflow: Pipeline-Aware Sub-Traversal Caching for Modern SmartNICs

Annus Zulfiqar, Ali Imran, Venkat Kunaparaju[1], Ben Pfaff[2], Gianni Antichi[3], Muhammad Shahbaz

University of Michigan, [1]Purdue University, [2]Feldera, [3]Politecnico di Milano

Learn More

Abstract and Motivation

- Virtual switches optimize performance by caching multi-table lookup traversals to single-table Megaflow cache, which SmartNICs offload directly to hardware
- We present **Gigaflow: a multi-table sub-traversal cache for SmartNICs**, designed to capture a much larger rule space using the same cache size

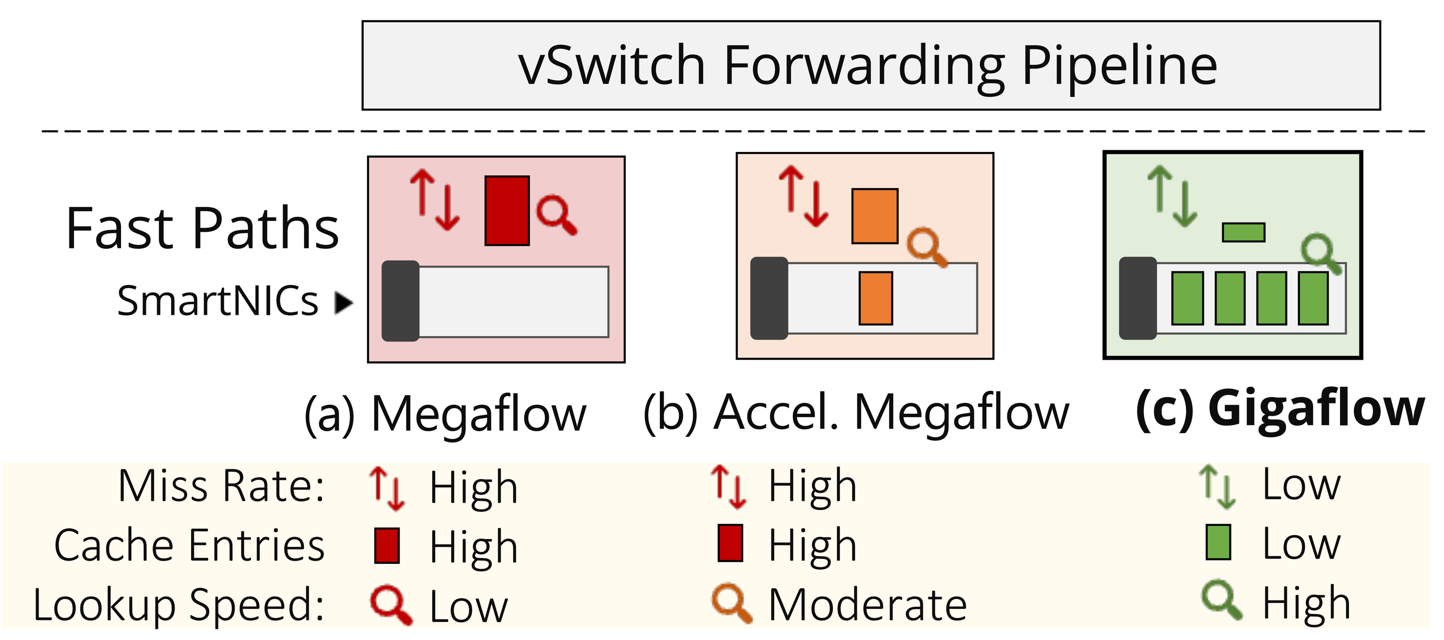

Figure 1: Comparison of OVS cache miss rate, entries and lookup speed

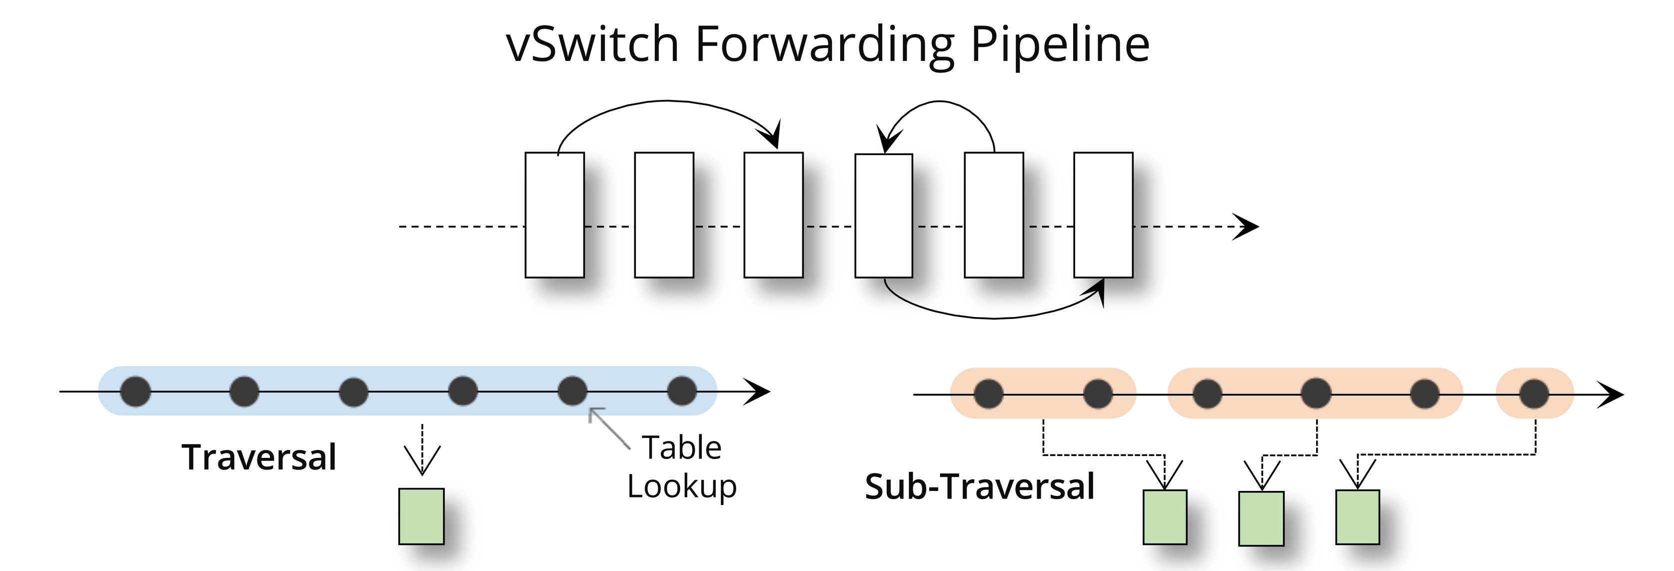

Figure 2: A traversal of vSwitch pipeline yields a Megaflow rule

- Open vSwitch caches traversals into Megaflow and can't share sub-traversals among flows, making the captured rule space proportional to cache size
- By caching sub-traversals into a multi-table cache, we can capture **3 orders of magnitude more rule space**, attain **51% higher cache hit rate**, and **30% lower end-to-end packet latency** with manageable processing overhead

Design: A Pipeline of TCAM Tables to Cache Sub-Traversals

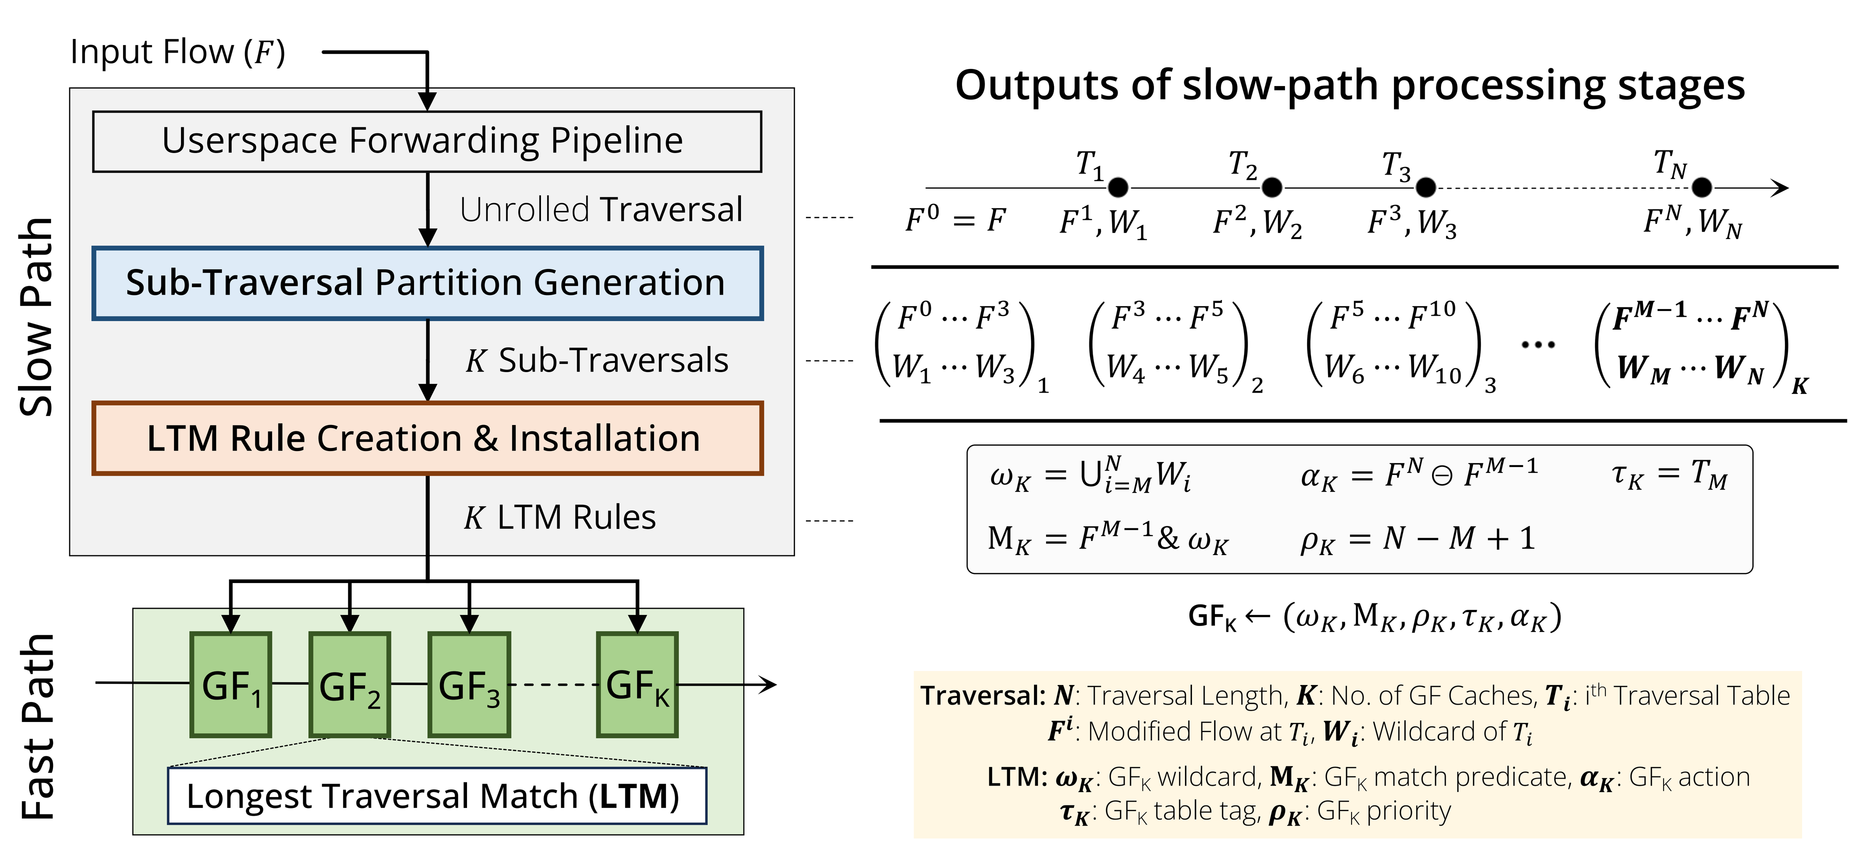

Figure 3: A high-level view of slow-path processing for cache misses in Gigaflow

The vSwitch processes cache misses through its userspace forwarding pipeline and **unrolls the traversal**. A **sub-traversal partitioning algorithm** explores its possible partitions to **maximize disjointedness in sub-traversal matching fields**, which maximizes the captured cross-product rule space. We convert these sub-traversals into Gigaflow cache entries.

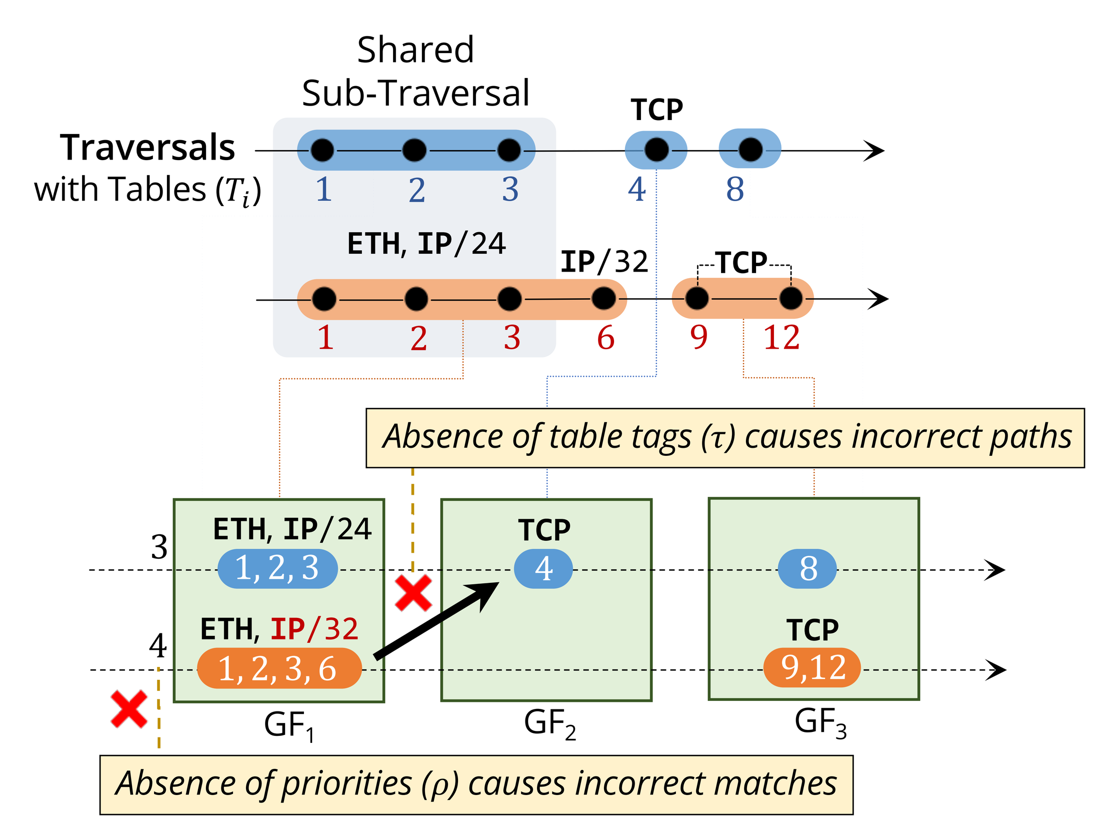

Figure 4: An example Gigaflow fast-path with cache entries

The multi-table cache maps nicely to RMT architecture, where individual tables contribute sub-traversal-level cache hits. To ensure lookup correctness, **Gigaflow** uses priorities (ρ) and table tags (τ) for **Longest Traversal Matching (LTM)**.

Evaluation and Results

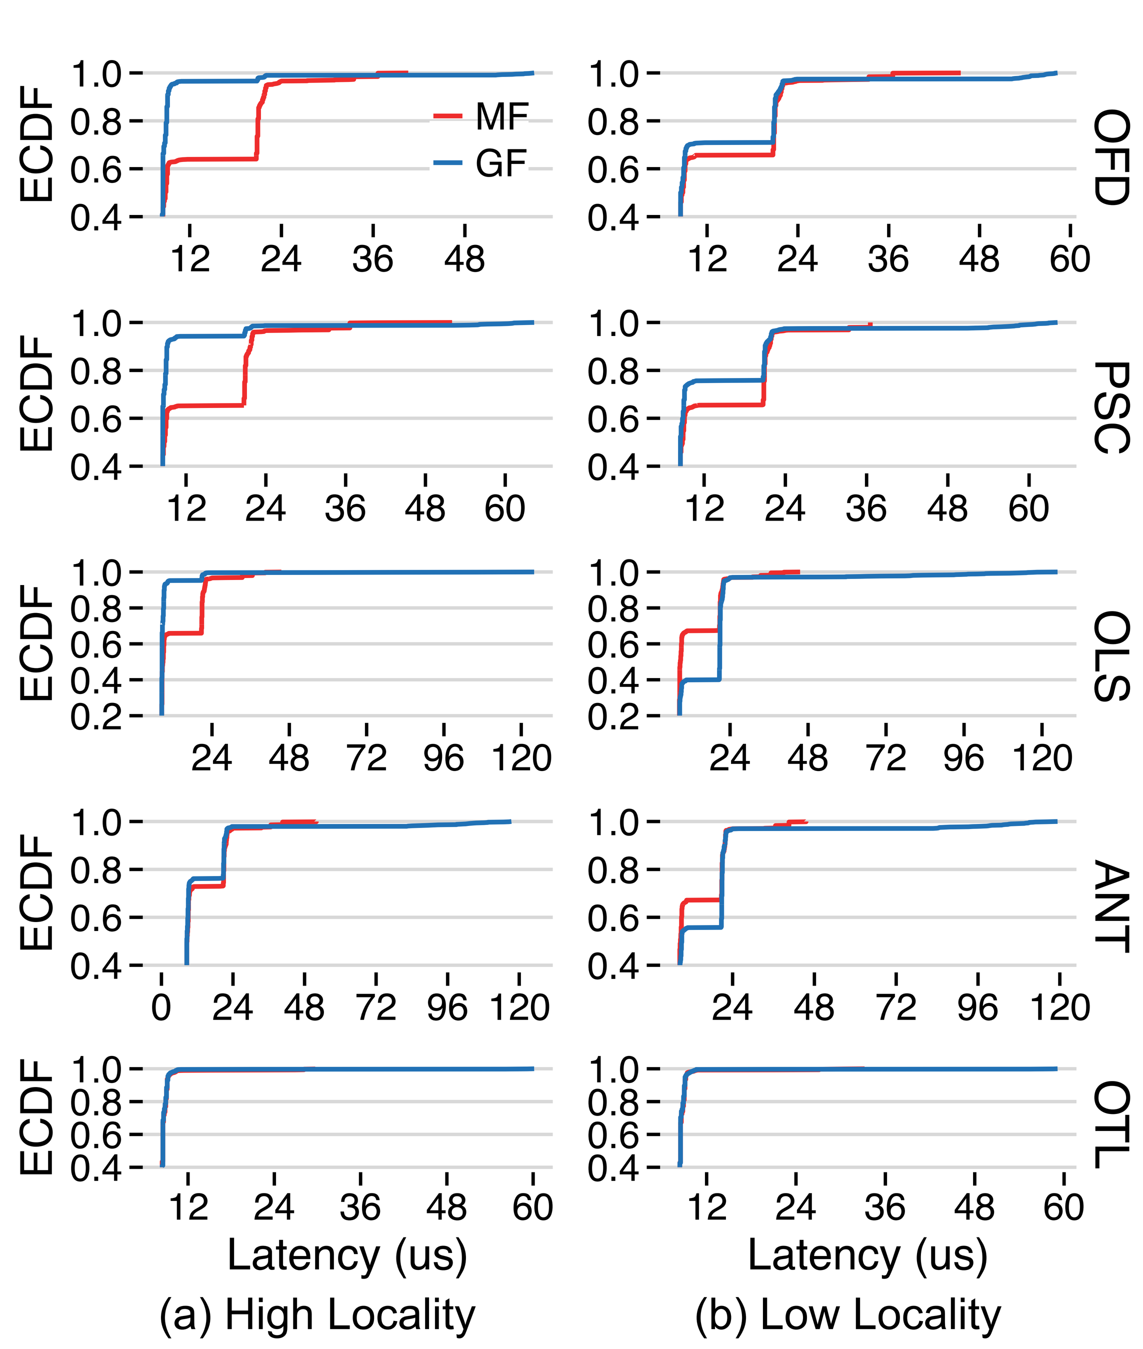

(a) End-to-end latency with Megaflow (MF) vs Gigaflow (GF) cache in high/low locality

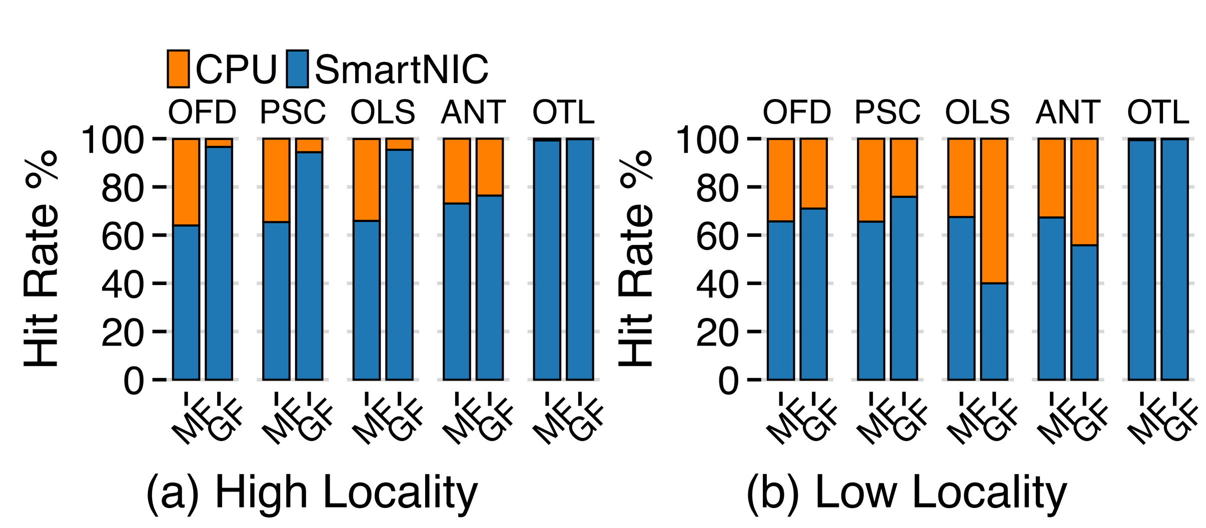

(b) Cache hit rates for real-world vSwitch pipelines with Megaflow (MF) vs Gigaflow (GF) cache on the SmartNIC

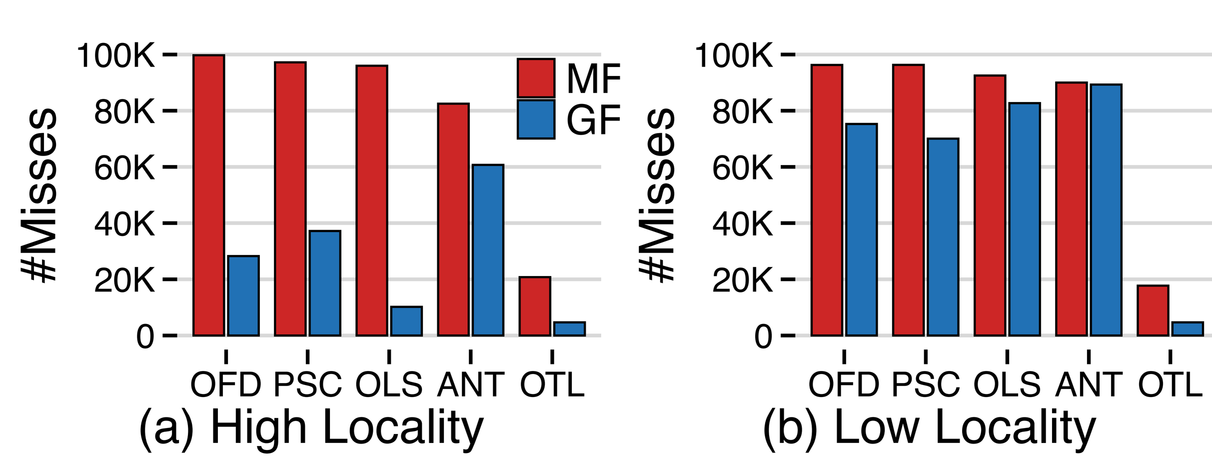

(c) Number of Cache Misses with Megaflow (MF) vs Gigaflow (GF) cache on the SmartNIC

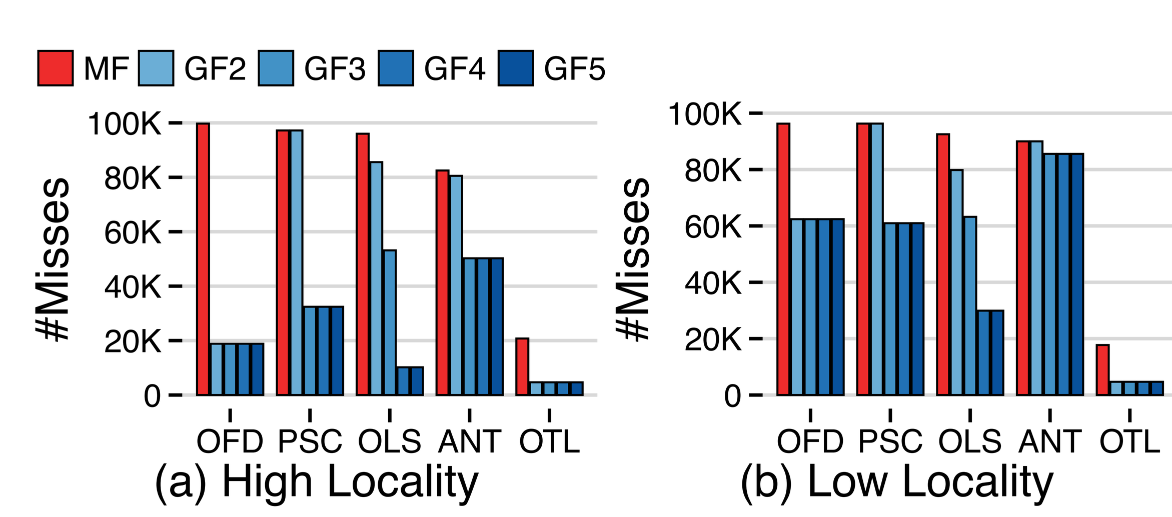

(d) Cache Misses vs. Number of Gigaflow tables

	OFD	PSC	OLS	ANT	OTL
Megaflow	32K	32K	32K	32K	32K
Gigaflow	14.7M	4.9M	10.8M	1.3M	48K

(e) Flow space capacity of Megaflow and Gigaflow cache with 32K entries

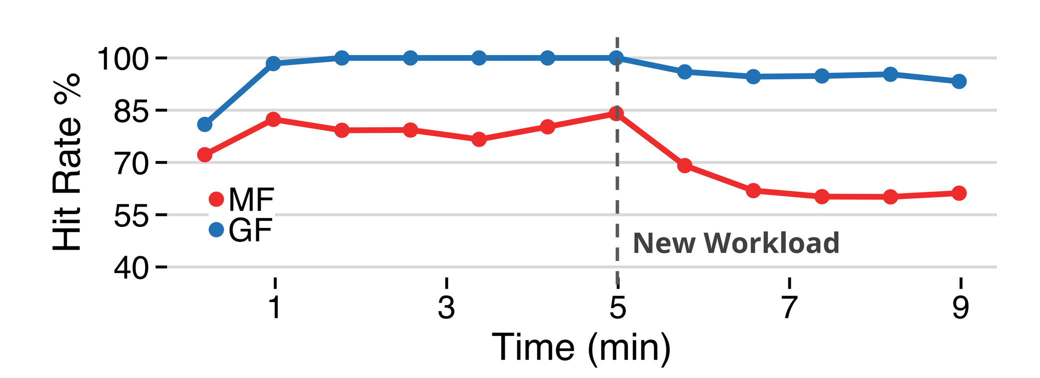

(f) End-to-end cache hit rate with Megaflow (MF) vs Gigaflow (GF) cache with dynamic workload

Contact us: zulfiqaa@umich.edu