

Scalable, High Performance Ethernet Forwarding Lookup on Commodity x86 Platform

Dong Zhou, Bin Fan, Hyeontaek Lim, David G. Andersen, Michael Kaminsky* (CMU, *Intel)

Goal

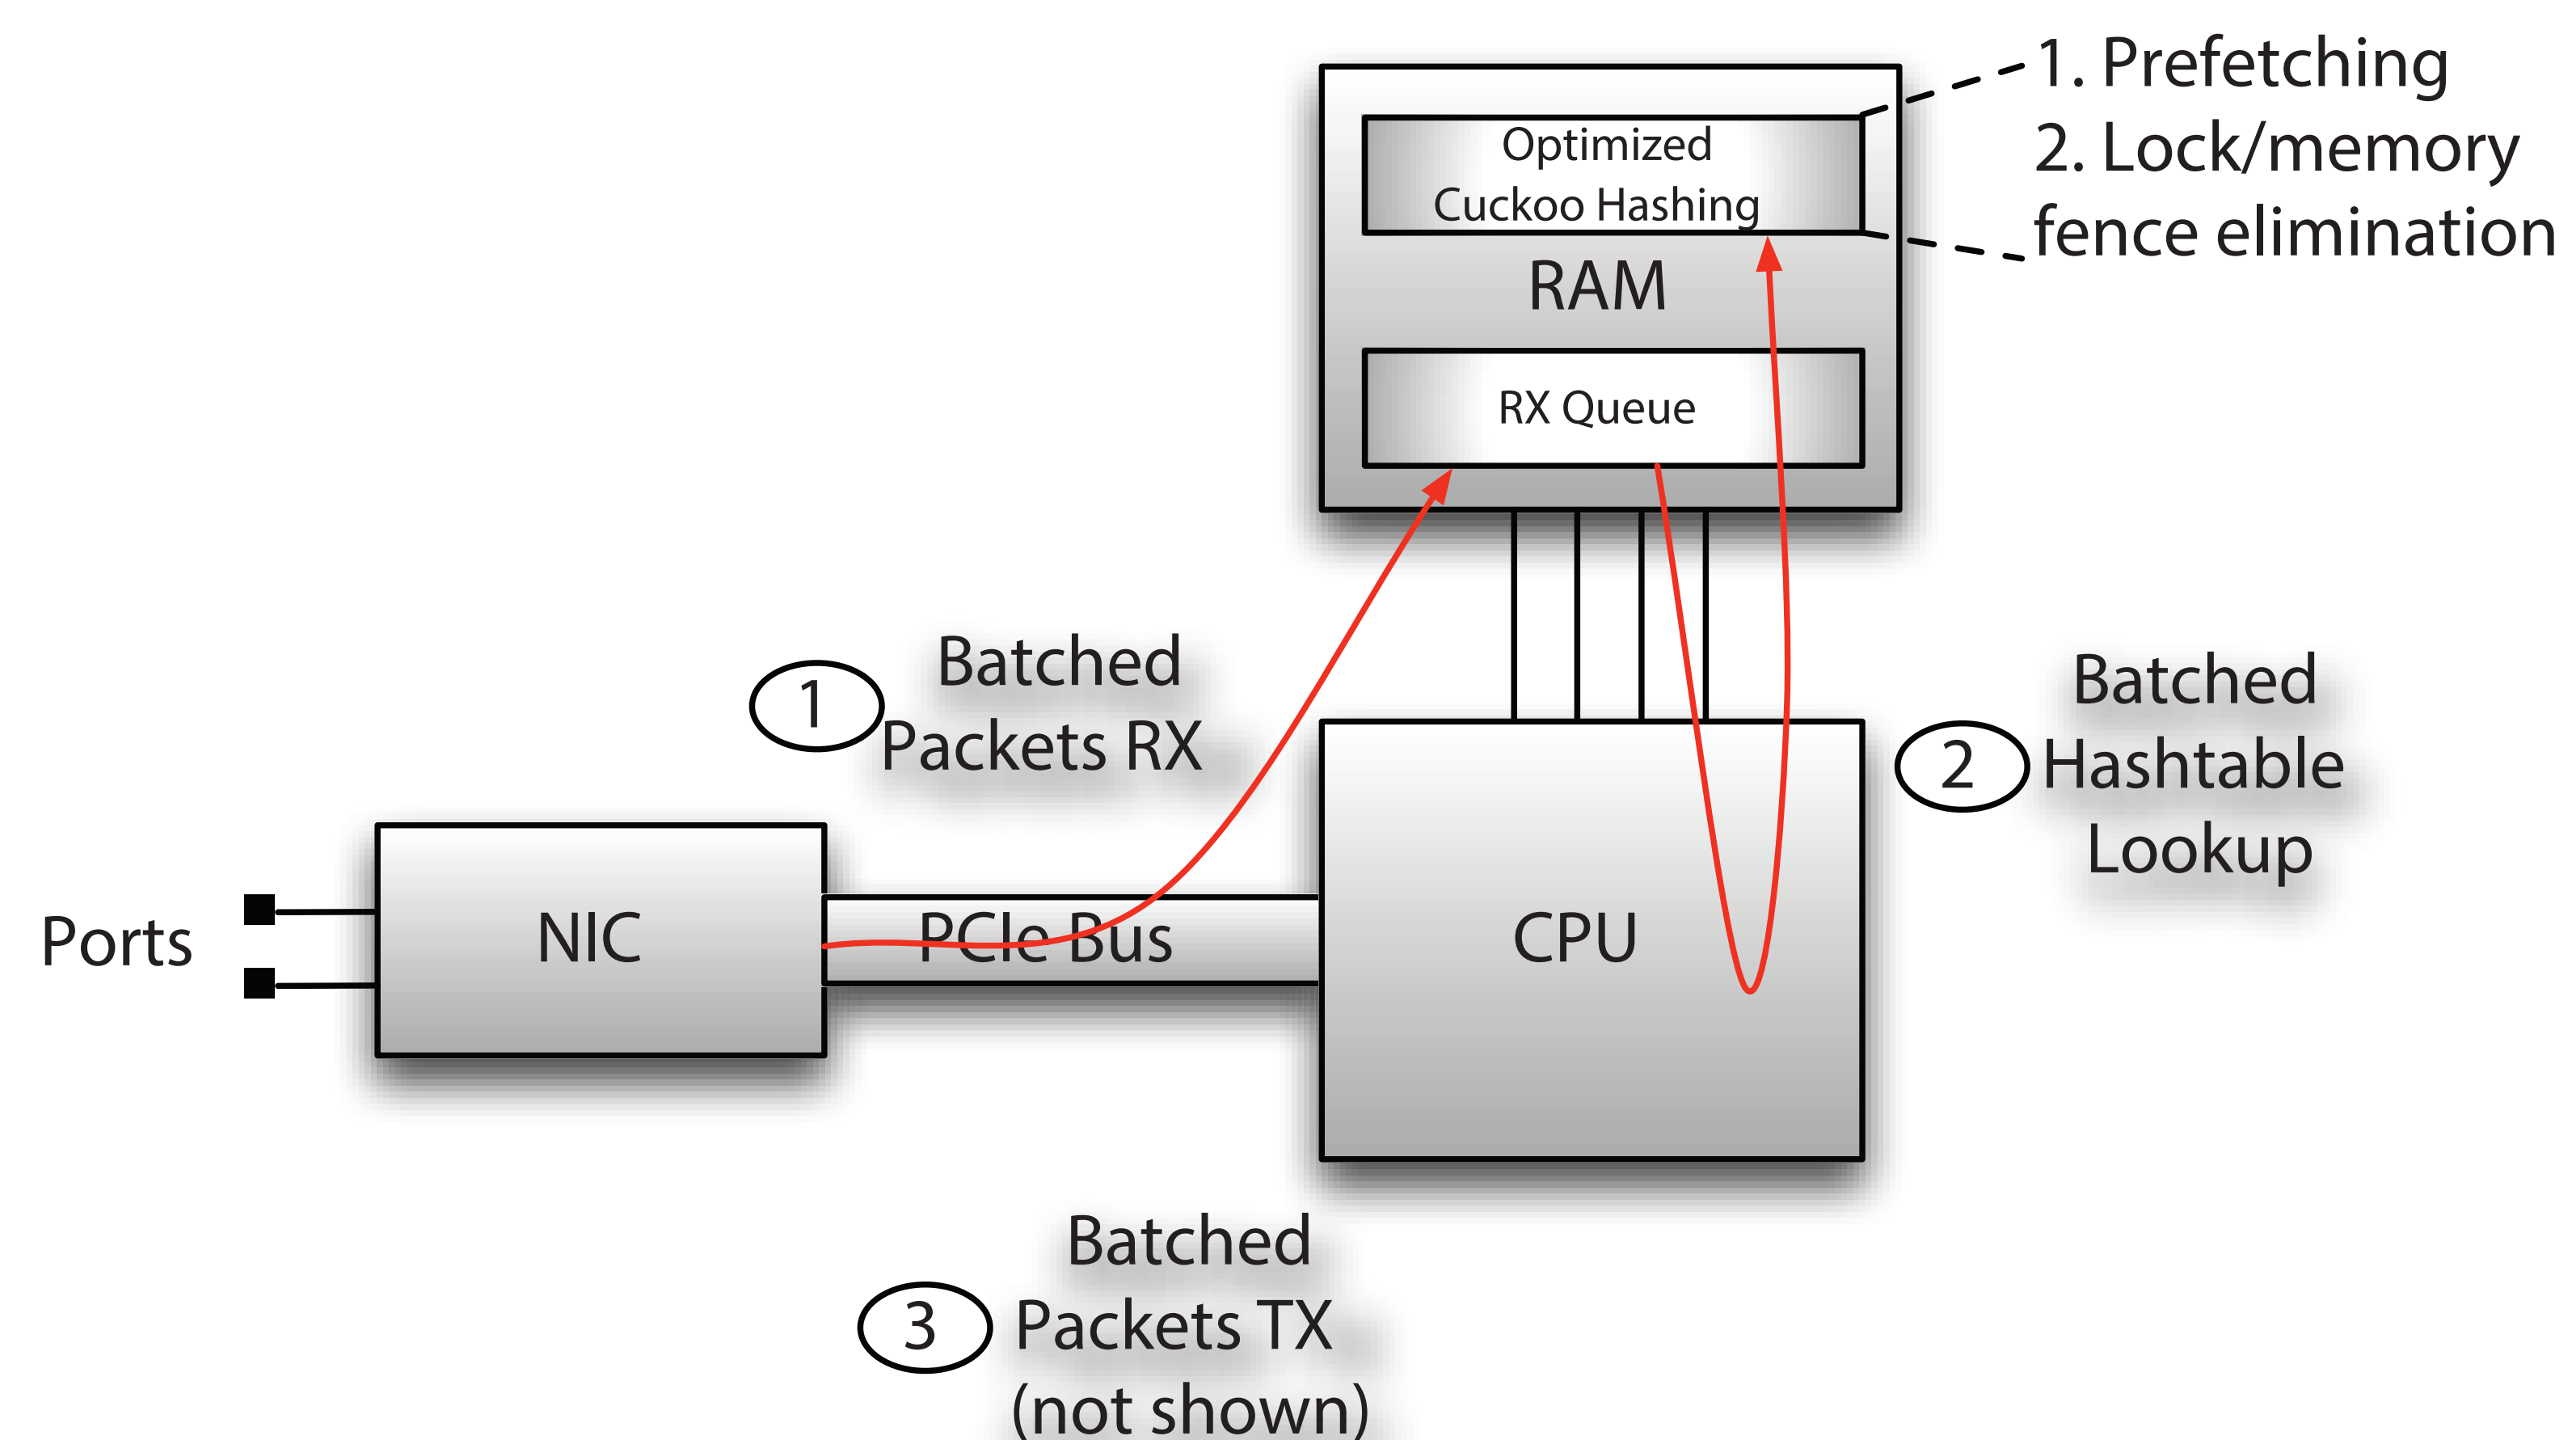
- Single-node software switch, with
 - One billion entries in the forwarding table
 - 10Gbps line speed
 - 4-8 ports support
- Motivation
 - Content-based networking
 - Ever-larger layer-2 networks enabled by recent research efforts
 - Pushing the limit of the scalability of forwarding table

Why Current Technologies Won't Do

- High-speed memories (e.g. TCAM)
 - Small size severely limits scalability
- Memory efficient but approximate solutions
 - Induce path stretch
- Prior hashing schemes
 - Either *memory inefficient* and/or have unacceptable lookup performance to handle collisions

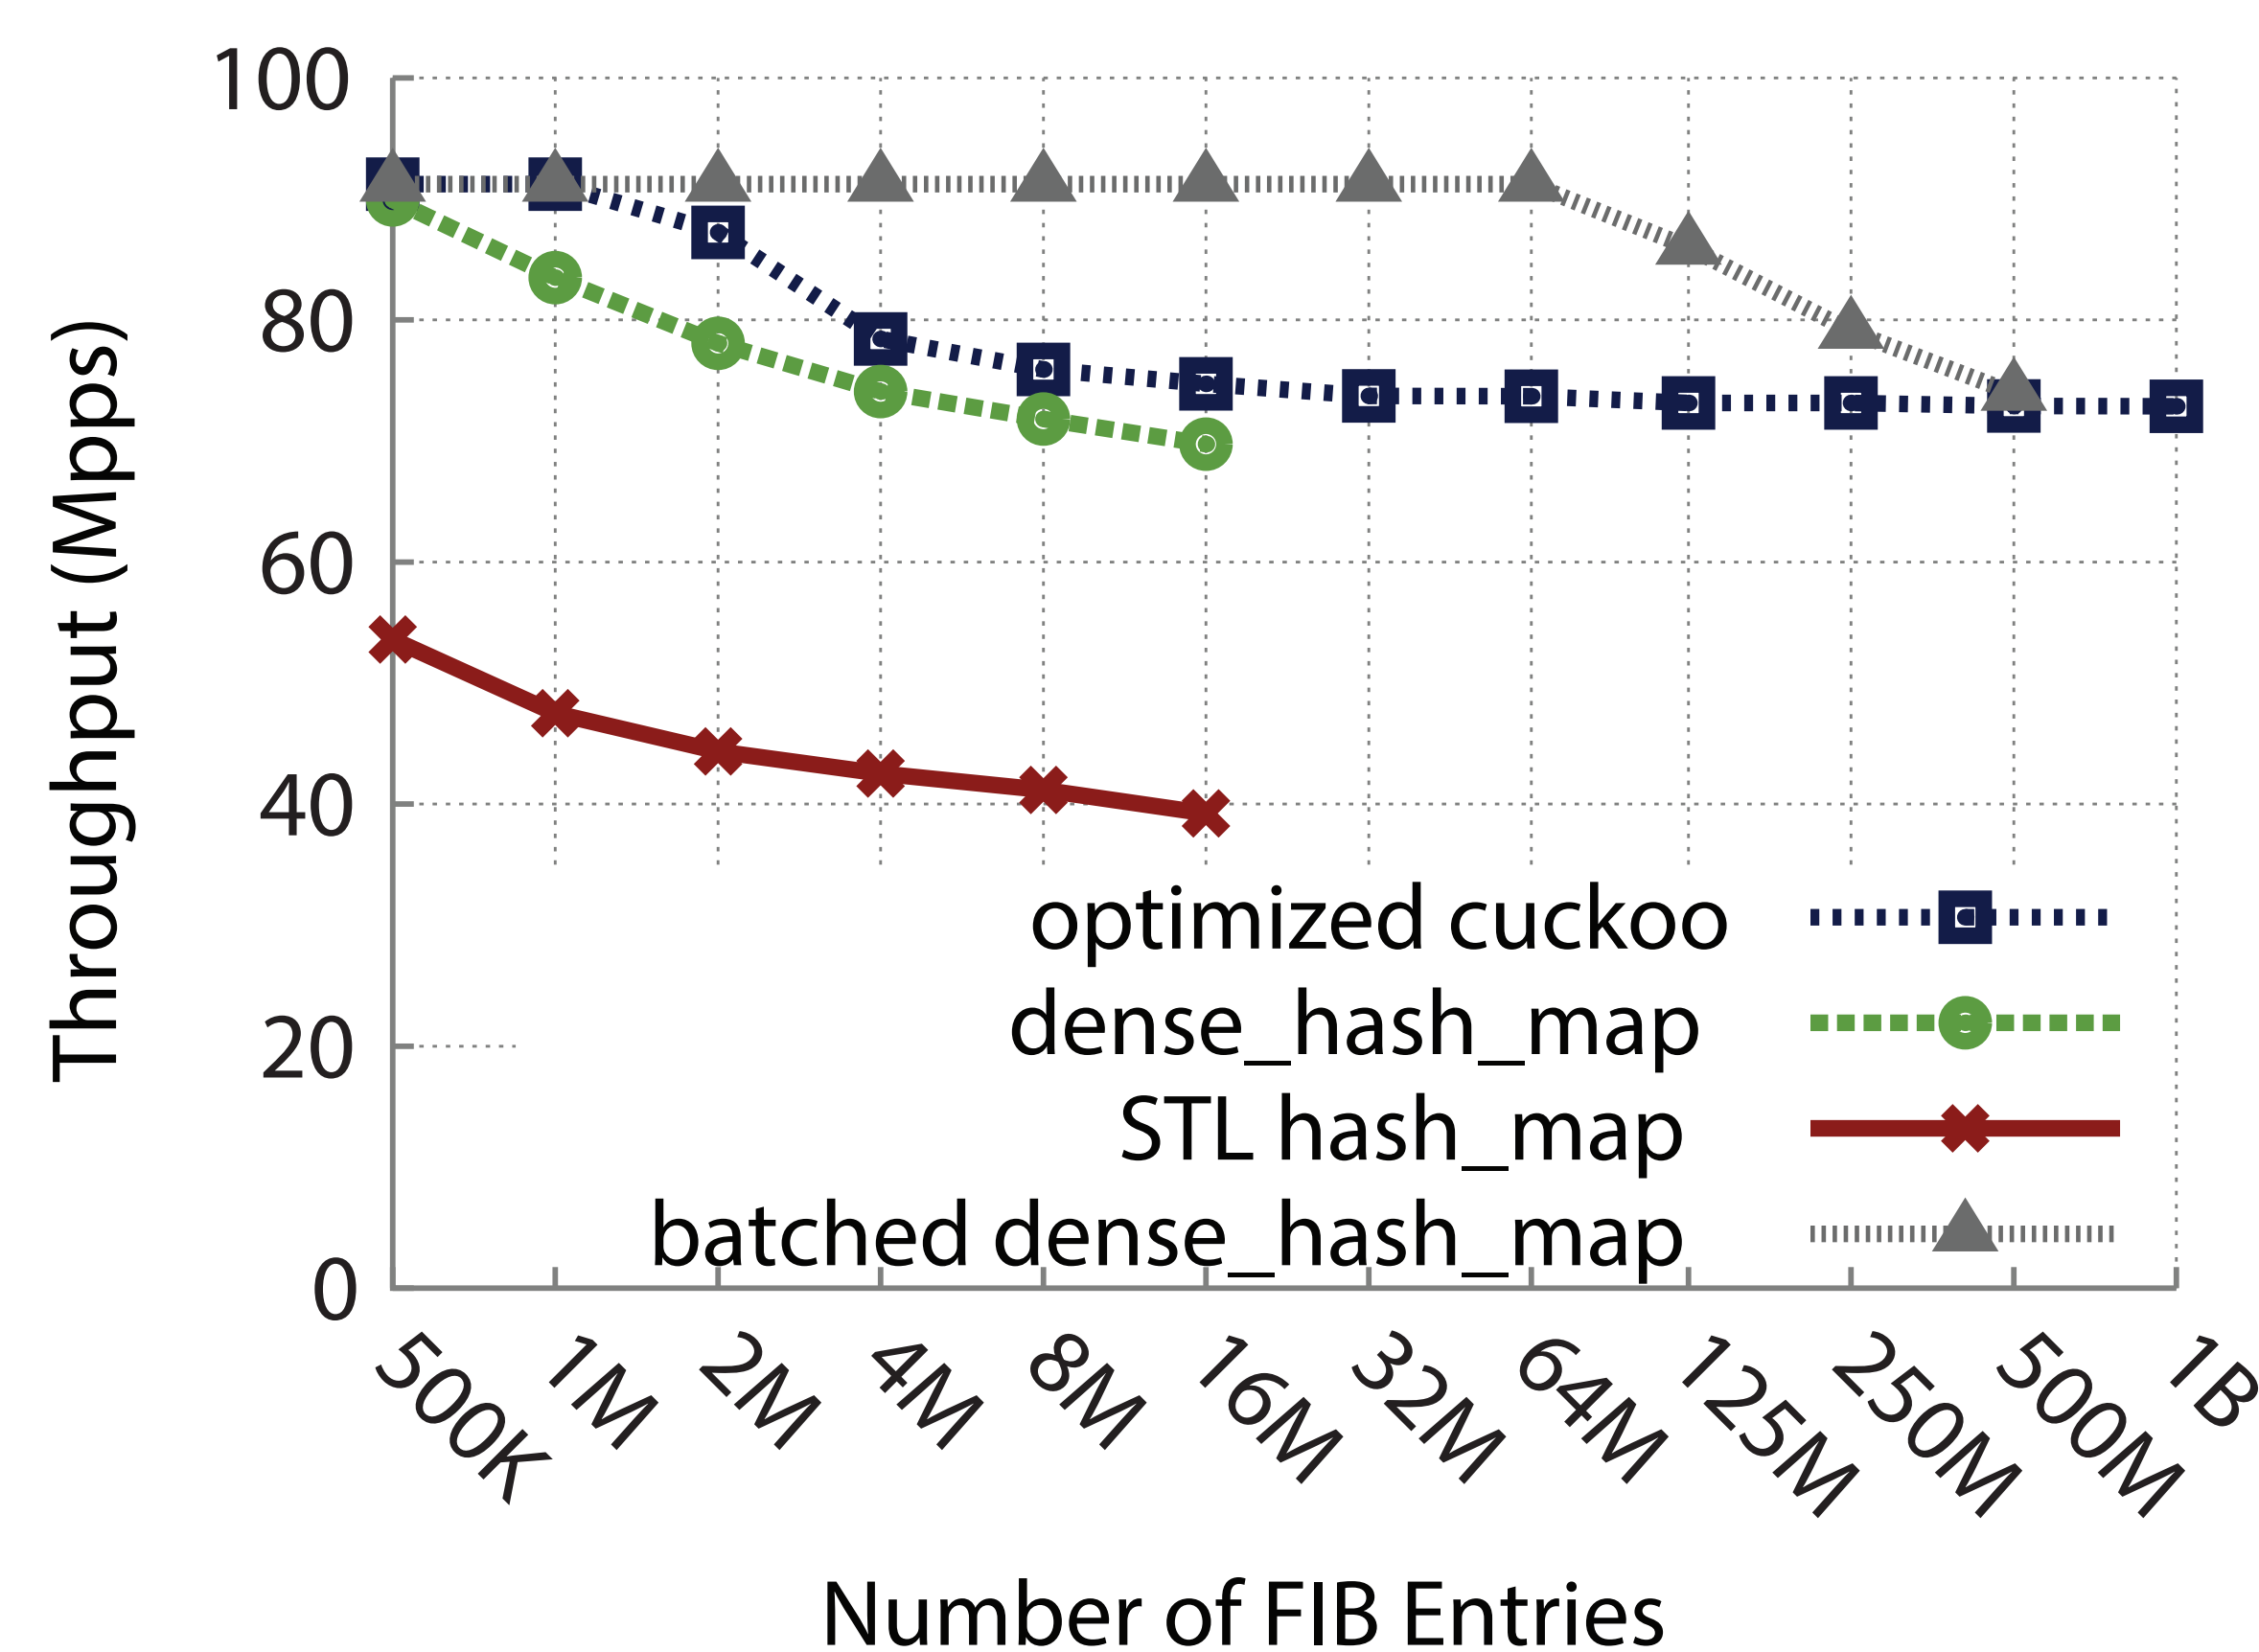
Solution: Intel DPDK + Optimized Concurrent Cuckoo Hashing

- Intel DPDK: high-throughput packet I/O to user-space
 - NUMA-aware memory management
 - Polling mode user-space driver
 - Batching for efficiency
- Optimized Concurrent Cuckoo Hashing
 - Fast lookups: two parallel cacheline-sized reads
 - Highly Compact
 - Concurrent reads with in-place writes



Evaluation

Minimum (64B) Packet



Throughput vs. Update Rate

