

# Scalable, High Performance Ethernet Forwarding on x86 Platforms

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

## GOAL

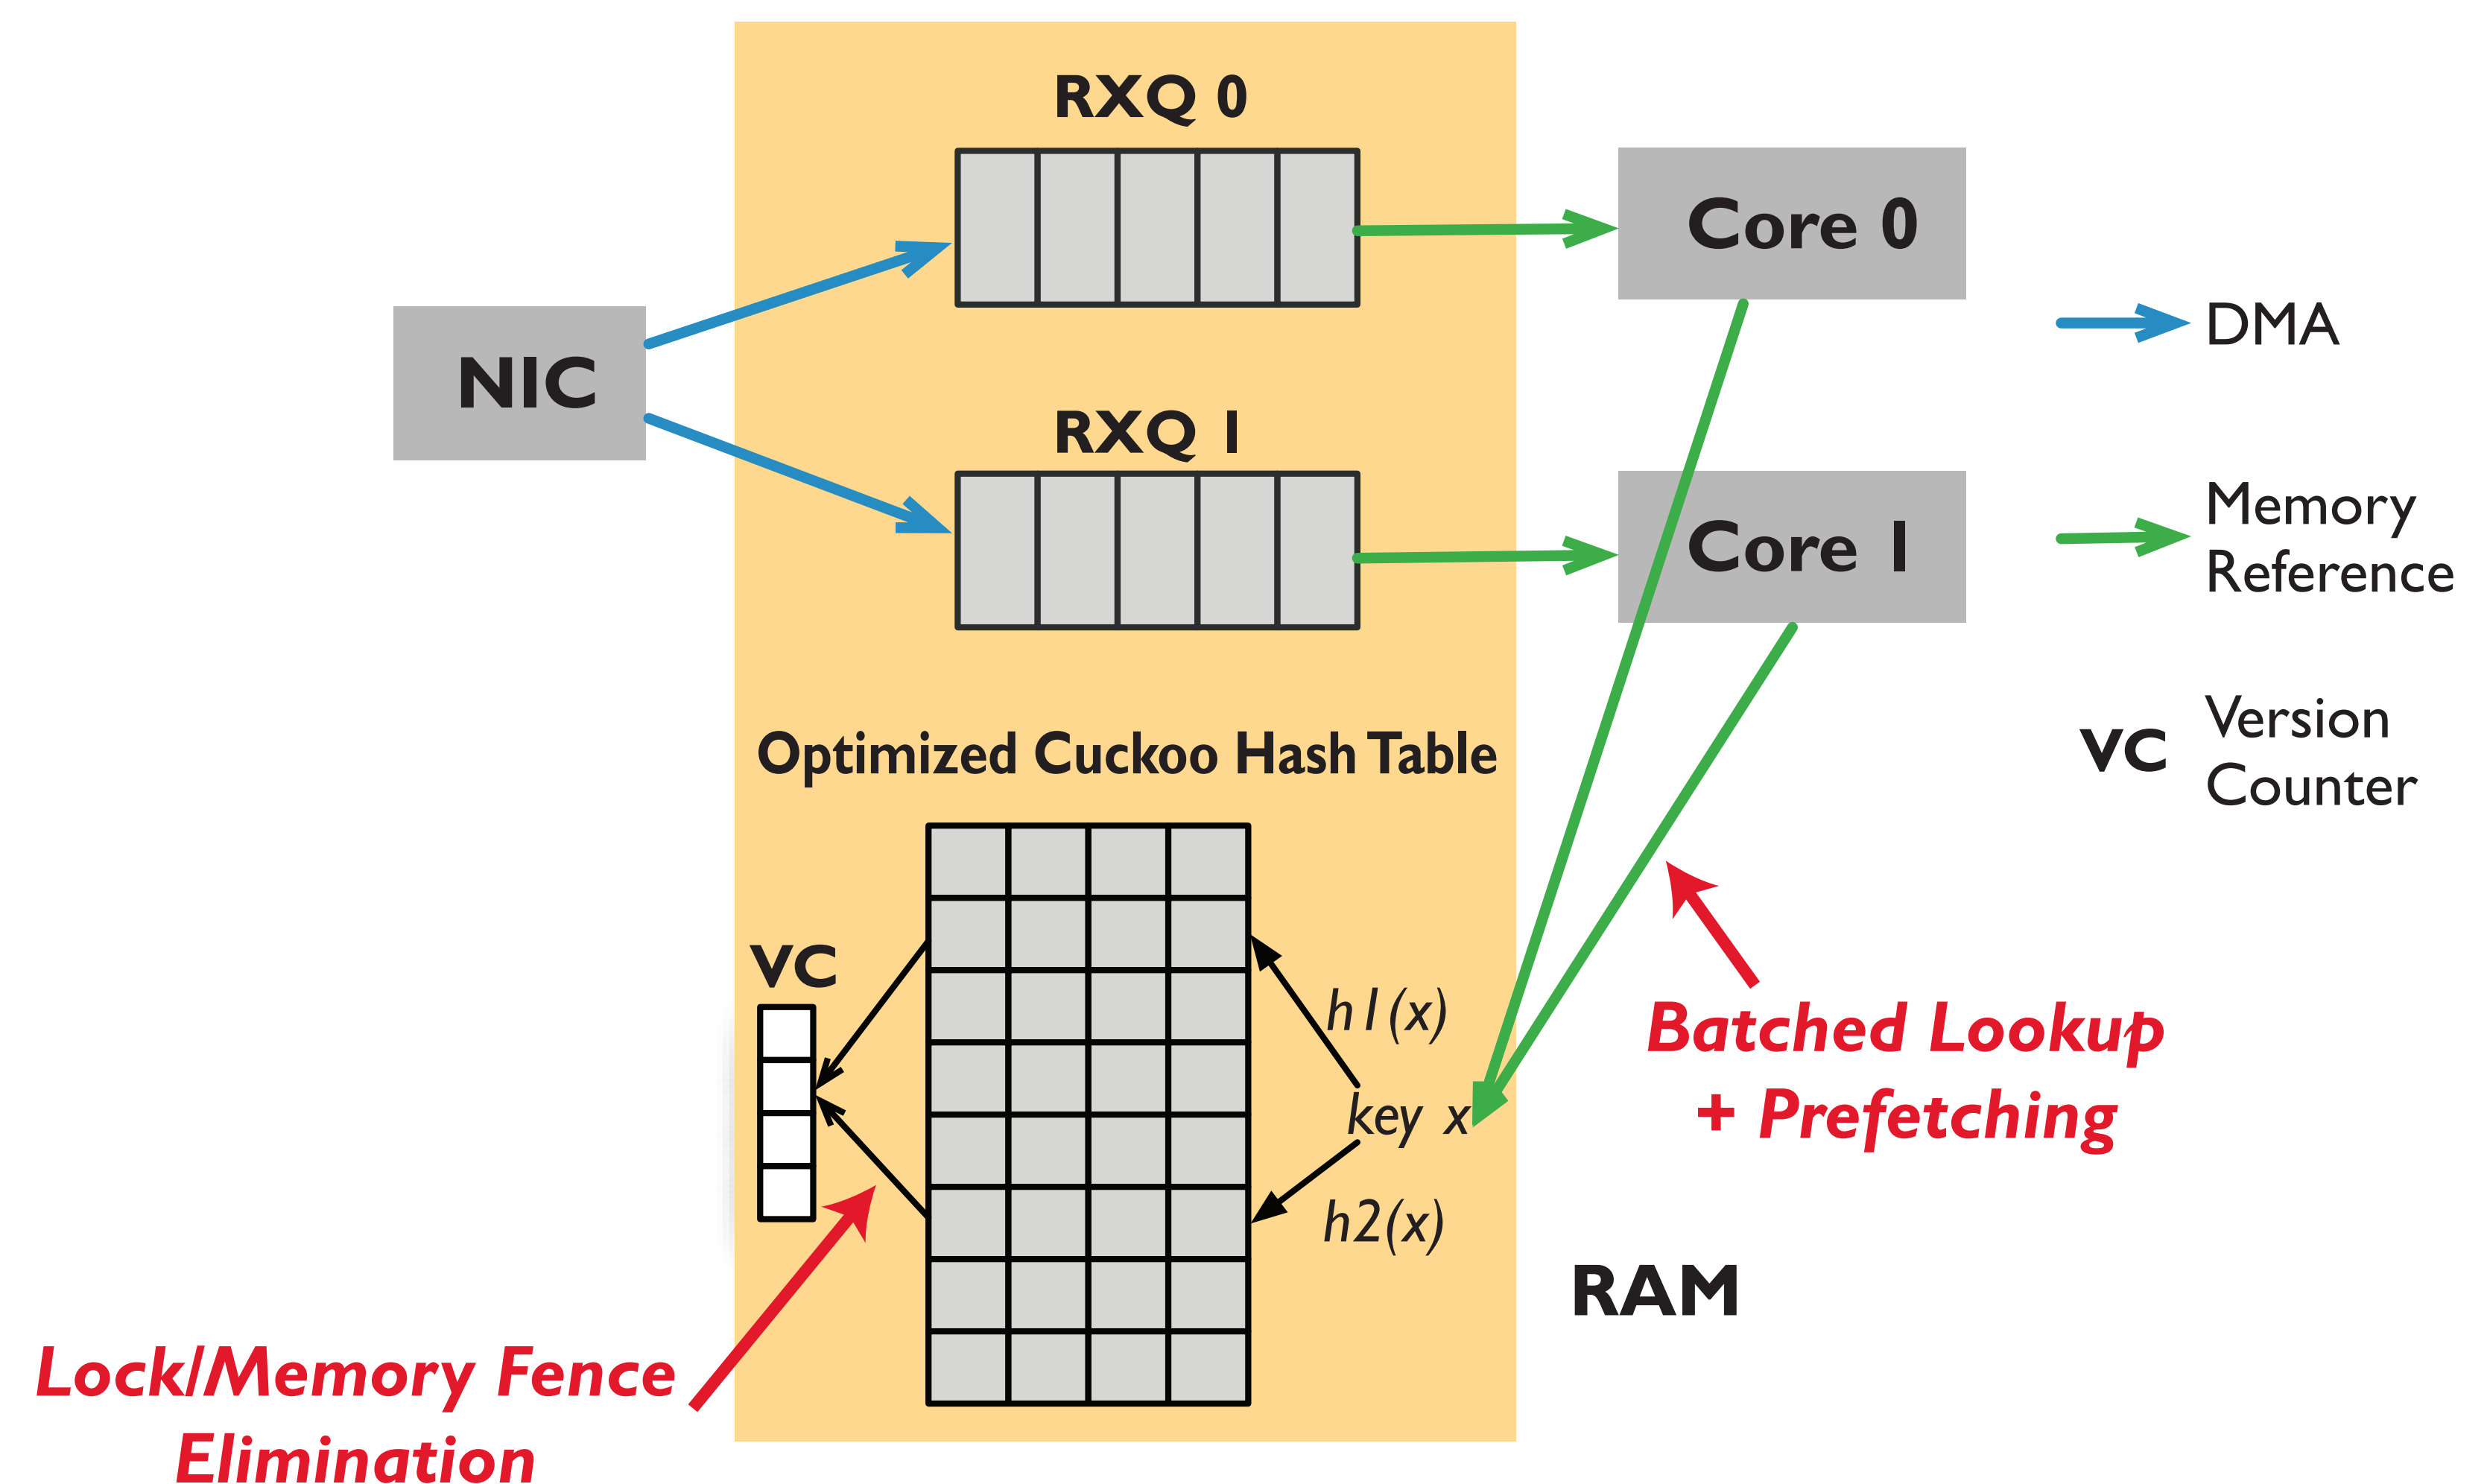
- Single-node software switch, with
  - One billion entries in the forwarding table
  - 10Gbps line speed
  - 4-8 ports support
- Motivation
  - Content-centric networks
  - Ever-larger layer-2 networks
  - Push the limits of hardware

## CURRENT TECHNIQUES

- 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 collision


## SOLUTION: INTEL DPDK + OPTIMIZED CUCKOO HASHING

- Intel Data Plane Development Kit (DPDK): high-throughput packet I/O to user-space
  - NUMA-aware memory management
  - Polling mode user-space driver
  - Batching for efficiency
- Optimized Cuckoo Hash Table
  - Fast lookup: ~350 Mops for small key/value pairs
  - Compact: 95% occupancy
  - Concurrent reads with in-place writes



## EVALUATION

### END-TO-END PERFORMANCE



### FACTOR ANALYSIS

