

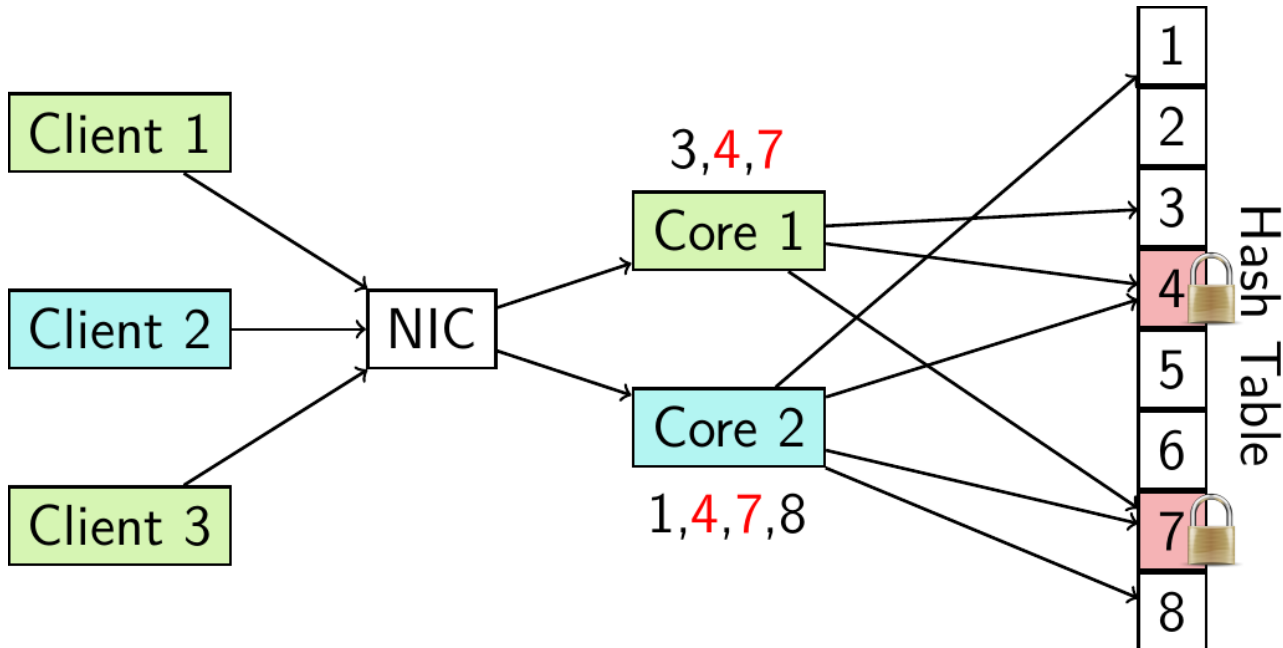
Key-value store cluster with programmability in NICs and switches

Antoine Kaufmann, Naveen Kr. Sharma

Scaling Up a Key-Value Store

- Implementing load-balancing today:
 - Client chooses destination explicitly
 - Load balancer (server, or dedicated core)
- Problem:
 - Client based: dynamic load balancing is hard
 - Load balancer: does not scale with bandwidth
- Key Idea: Use flexible network hardware to steer packets
 - NICs: FlexNIC, Intel RRC, Netronome
 - Switches: RMTs, Intel Flexpipe

Packet Steering: Status Quo



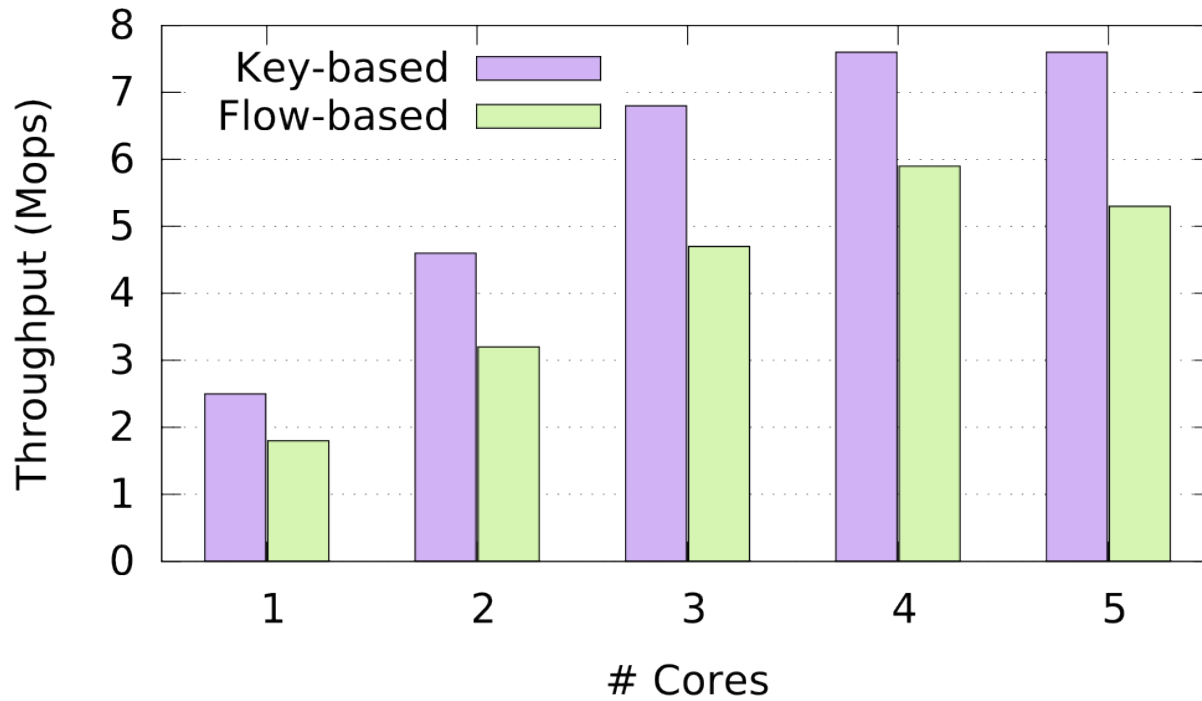
- Problem:

- Lock contention
- Poor cache utilization

Scaling to Multiple Cores

- Idea: NIC steers packets to cores based on hash over key
- Implementation:
 - Place key in IPv6 src/dst address fields
 - Use receive-side scaling in NIC to steer packets to cores
 - NIC also provides RSS hash to Software, use for hash table lookup

Resulting Throughput



Scaling to Multiple Servers

- Idea: Switch steers packets to servers based on hash over key
- Implementation:
 - Again key in IPv6 src/dst address fields
 - Use Link Aggregation over server ports on switch

# Servers	1	2	3
Requests/s	320,017	633,712	950,914
Scale-up	1x	1.98x	2.97x