

# Is it a SmartNIC or a Key-Value Store? Both!

G. Siracusano\*, R. Bifulco†  
University of Rome Tor Vergata\*, NEC Laboratories Europe†

## Motivation

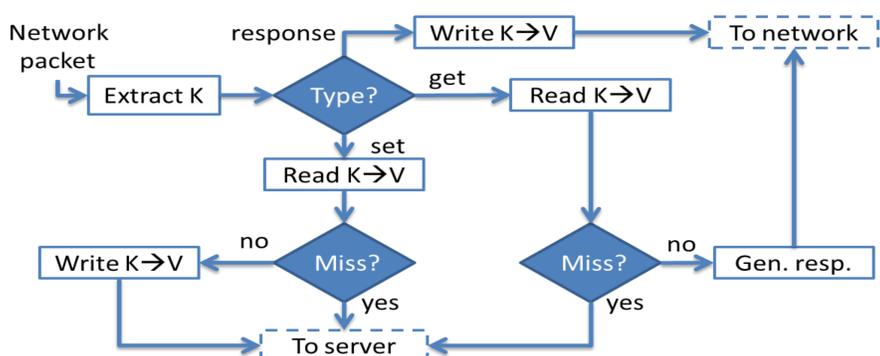
In-memory Key-Value stores (KVSs) are critical components of modern web services.

Given their key role in website performance, KVS are carefully tuned to maximize throughput and minimize response time.

### Related work

- KVSs have been optimized to take advantage of modern server's hardware [1].
- Purpose-built accelerators have been proposed to improve performance [2].

## NICached



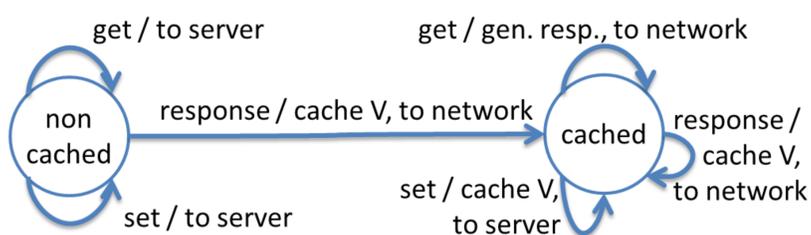
**Idea:** combine the two approaches in a single implementation that runs both on general purpose servers and on SmartNICs.

### Solution:

- **NICached**, a general caching system for KVSs, which can be supported by upstream Linux Kernels and amenable to be offloaded to different types of SmartNICs.
- A cache that sits at the earliest packet processing point of a modern server's network processing stack.

**Caching strategy:** Store the most recently requested Key/Value entries.

## Design

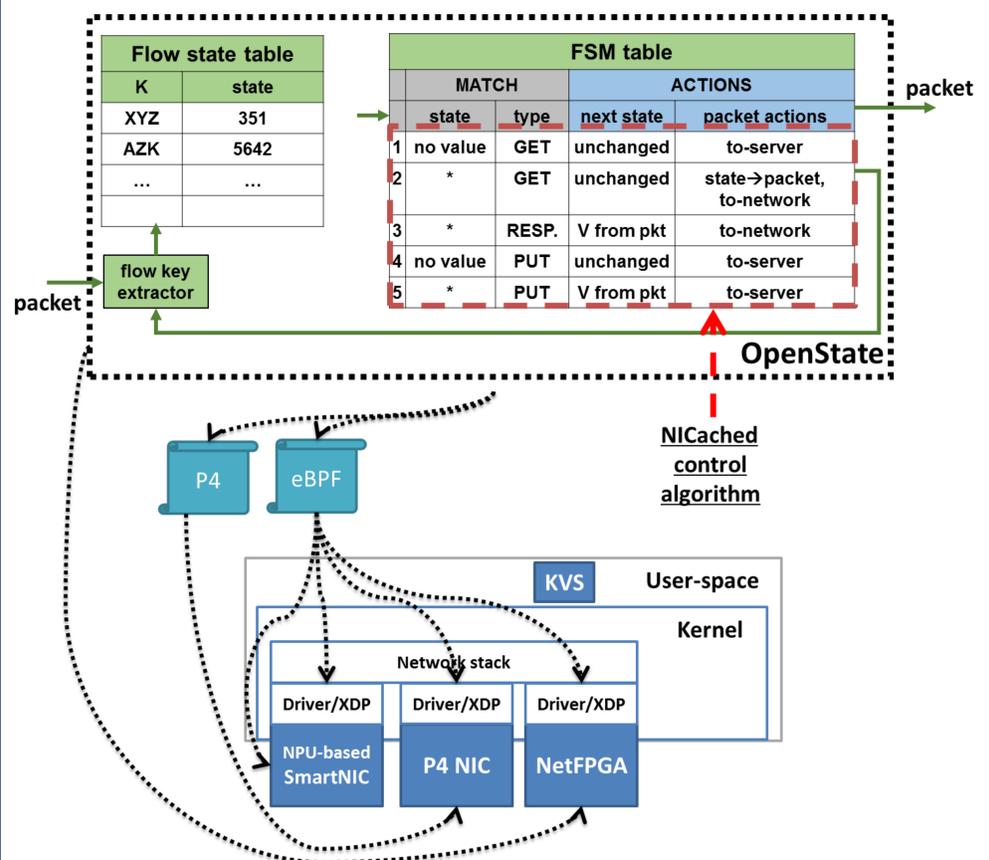


Several proposals in programmable network data planes suggest FSM abstraction as programming model [3,4].

NICached algorithm can be easily expressed using a Finite State Machine (FSM).

- we describe NICached using the OpenState [3] abstraction.

## PoC Implementation



NICached implemented using the OpenState abstraction could be deployed to several targets.

- NetFPGA.
- P4 based NICs.
- eBPF for Linux or for NPU-based SmartNIC.

NICached PoC deployed using the eBPF target.

- eXpress Data Path (XDP) hook is used.
- XDP\_TX action to transmit packets.

Testbed:

- Two machine with Intel Xeon E5-1630 CPUs (4 cores @3.70GHz).
- Two Mellanox ConnectX-3 (40Gbps) Ethernet cards.
- XDP\_TX baseline 6.9 million packets per second (Mpps).

Software micro-benchmarks:

- Cache composed by 3.2 million of entries, Key/Value size 8 bytes.
- Small entries represent the critical workload for KVS[1,5].
- memcached throughput ~ 0.9 million of request per second.
- NICached throughput ~ 5.4 million of request per second, using a single core.

[1] Sheng Li et al.2015. Architecting to achieve a billion requests per second through-put on a single key-value store server platform. In ACM SIGARCH Comput Architect News. ACM.  
[2] Michaela Blott et al.2013. Achieving 10Gbps Line-rate Key-value Stores withFPGAs.. InHotCloud  
[3] Giuseppe Bianchi et al.2014. OpenState: programming platform-independent stateful openflow applications inside the switch. ACM SIGCOMM CCR.  
[4]Masoud Moshref et al.2014. Flow-level state transition as a new switch primitivefor SDN. In ACM SIGCOMM HotSDN.  
[5]Berk Atikoglu et al.2012. Workload analysis of a large-scale key-value store. In ACM SIGMETRICS Performance Eval. Rev.

### ACKNOWLEDGMENTS

This work has received funding from the European Union's Horizon 2020 research and innovation programme under grant agreement No. 671648 ("VirtuWind"). This paper reflects only the authors' views and the European Commission is not responsible for any use that may be made of the information it contains