

StRoM: Smart Remote Memory

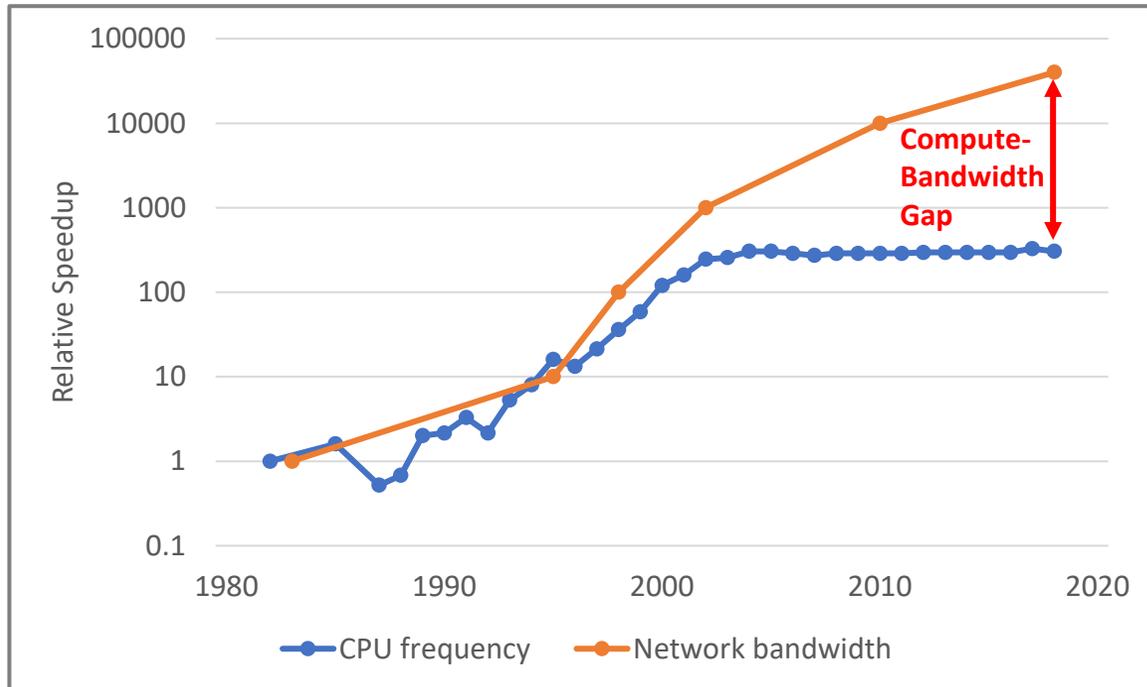
David Sidler^{*‡}, Zeke Wang^{†‡}, Monica Chiosa[‡], Amit Kulkarni[‡], Gustavo Alonso[‡]

* Microsoft Corporation

† Collaborative Innovation Center of Artificial Intelligence, Zhejiang University

‡ Systems Group, Department of Computer Science, ETH Zürich

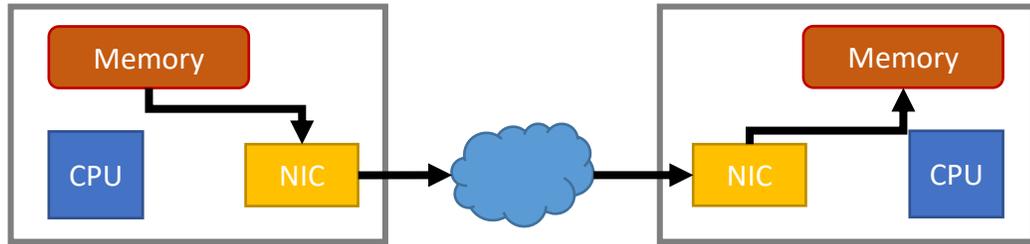
Increasing Compute-Bandwidth Gap



- Increase in CPU cycles allocated towards network processing
- Context switches between OS network stack and application amplify the issue

RDMA (Remote Direct Memory Access)

RDMA (Remote Direct Memory Access)



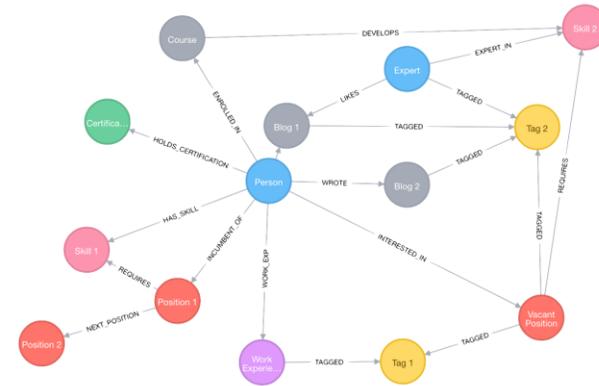
Complete Hardware offload
=> **Bypasses OS and CPU**



Distributed key-value stores[1,2]



Parallel database systems



Distributed graph computation[3]

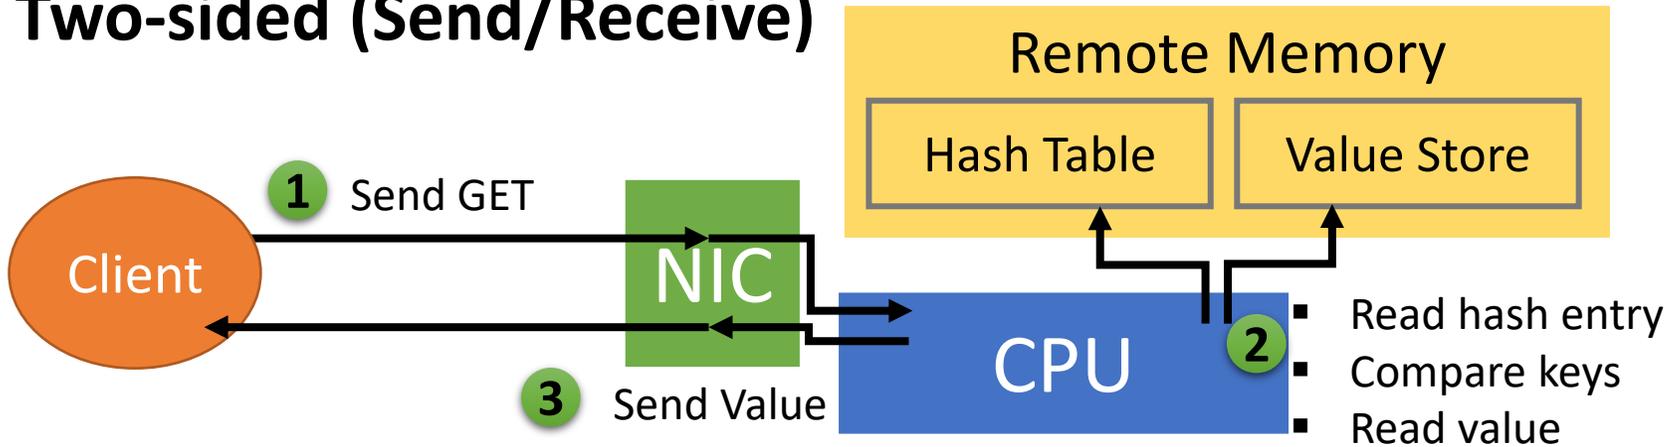
[1] C. Mitchell, et al., *Using One-sided RDMA Reads to build a fast, CPU-efficient key-value store*, ATC'13

[2] A. Dragojevic, et al., *FaRM: Fast Remote Memory*, NSDI'14

[3] M. Wu, et al., *GRAM: Scaling graph computation to the trillions*, SoCC'15

Get over RDMA: Two-sided vs One-sided

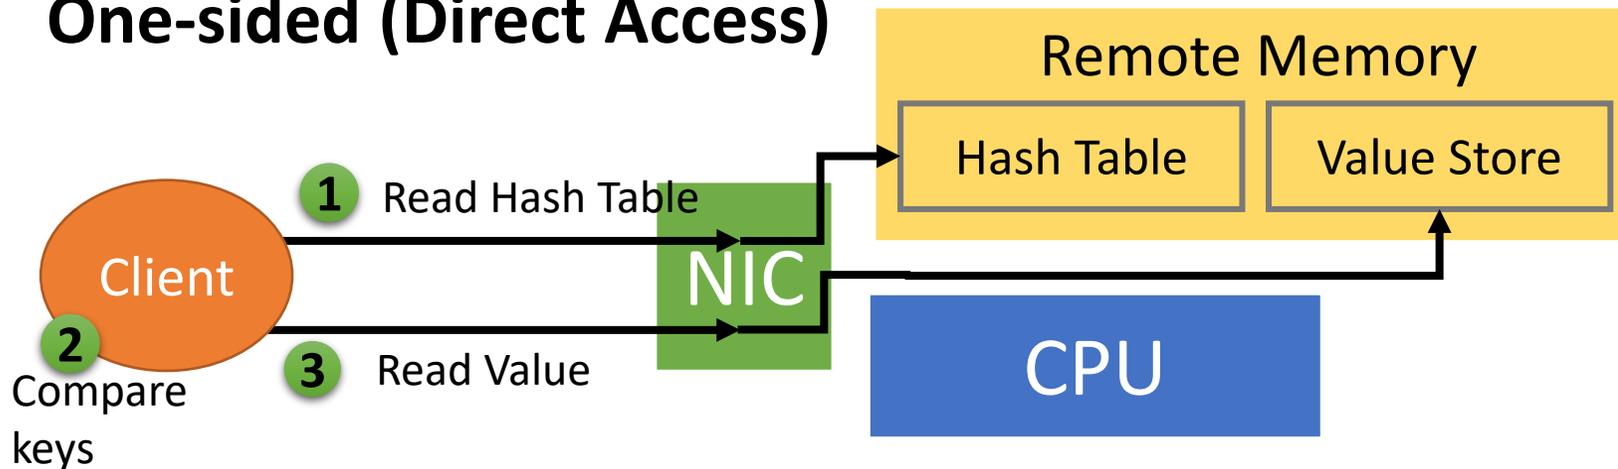
Two-sided (Send/Receive)



- Single round trip
- Simple client-server model
- Remote CPU involved

No clear winner

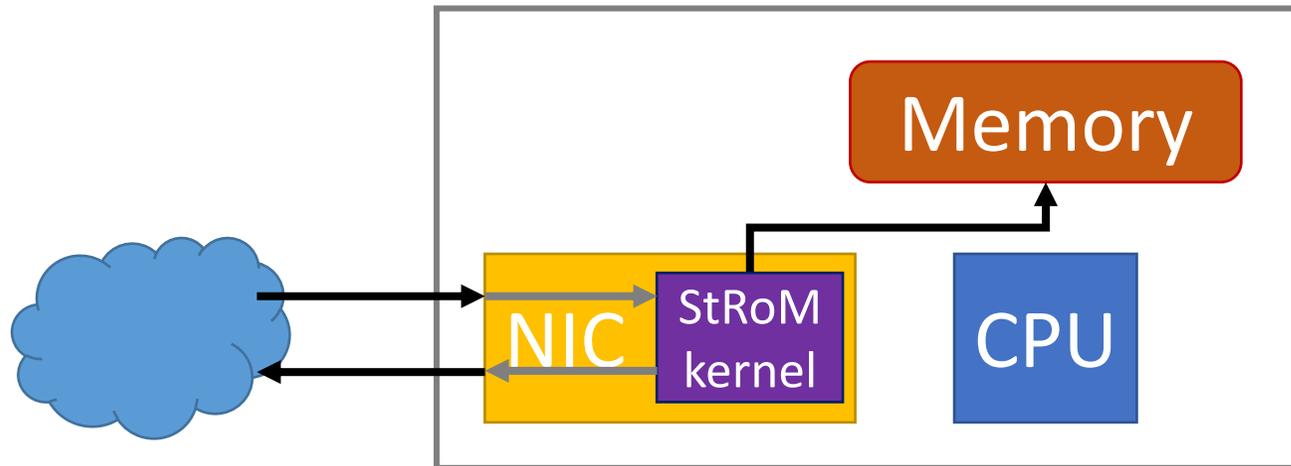
One-sided (Direct Access)



- Remote CPU not involved
- At least 2 RTs necessary
- Handling of misses costly

StRoM: Smart Remote Memory

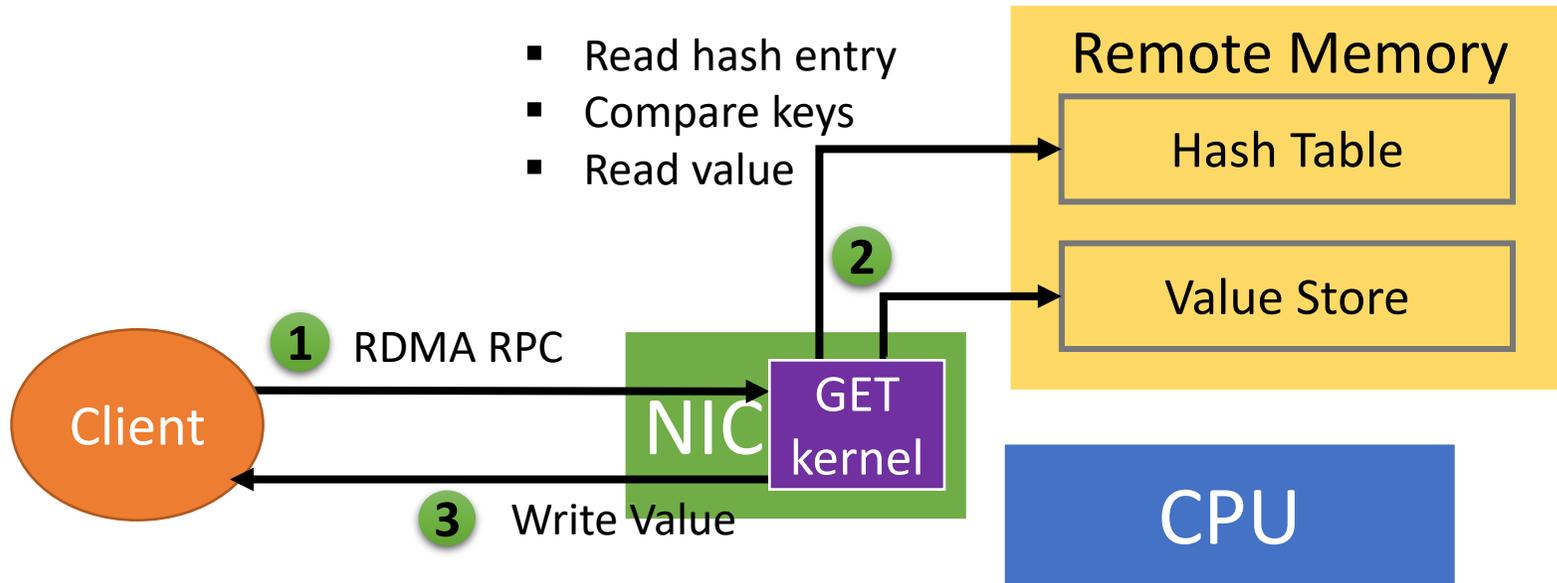
StRoM: Deployment of Acceleration kernels on the NIC



StRoM kernel

- Direct access to host memory
- Able to receive/transmit data over RDMA

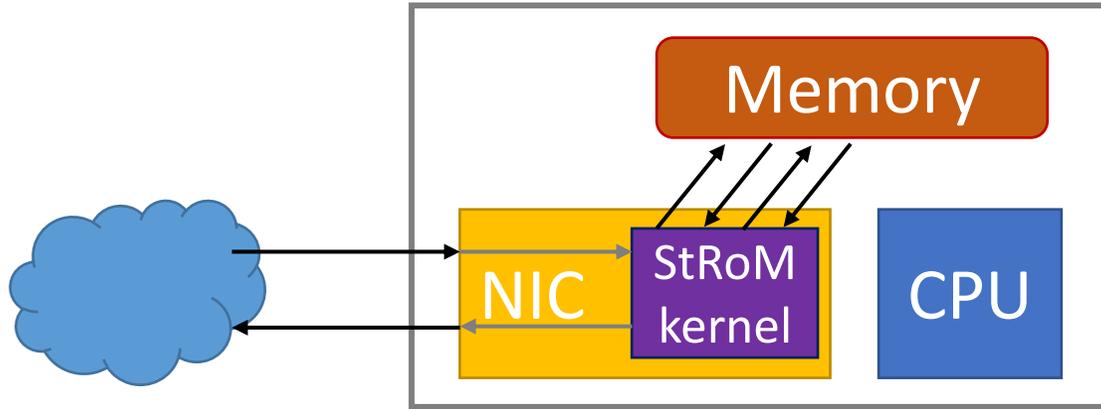
GET as StRoM Kernel



- **Single round trip**
- **Remote CPU not involved**

Acceleration Capabilities

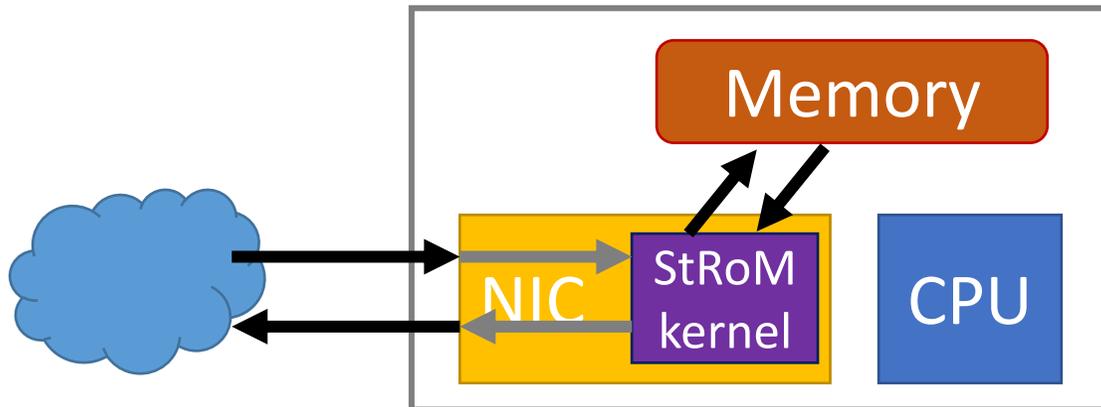
Accelerating Data Access



Invoke one-sided RPCs on the remote NIC

- Traversal of remote data structures
- Verification of data objects
- Manipulation of simple data structures

Accelerating Data Processing



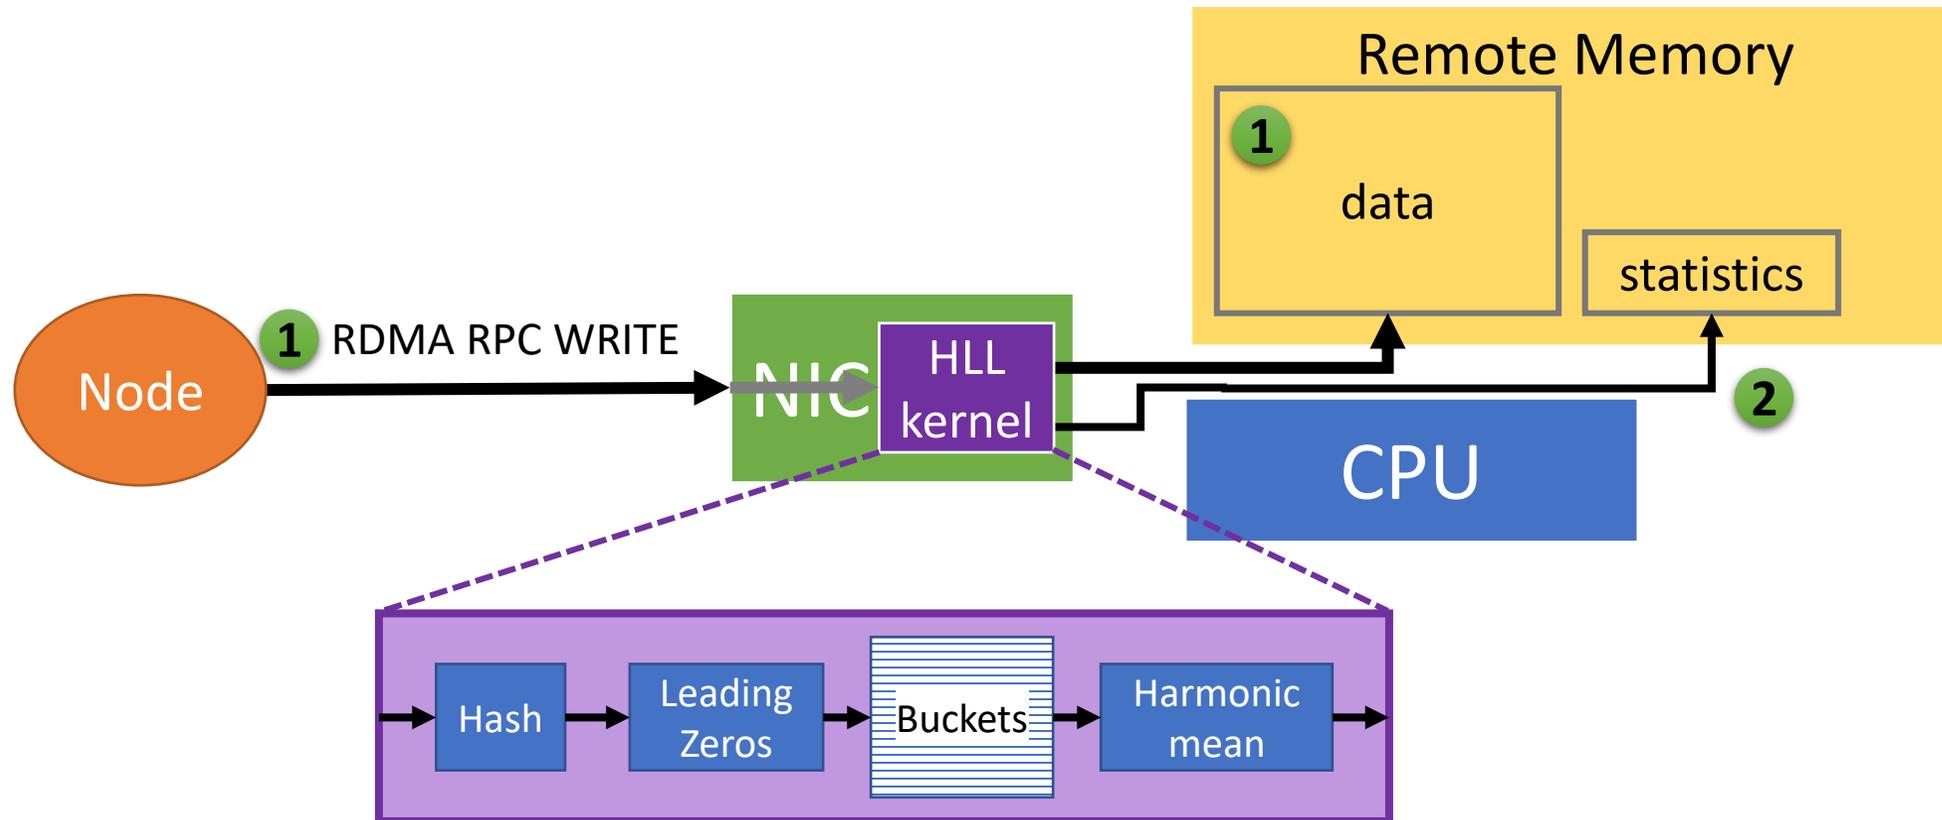
On-the-fly data processing when transmitting/receiving

- Data shuffling
- Filtering
- Pattern/event detection
- Aggregation
- Compression
- Statistics gathering

Use Case: Gathering Statistics

HyperLogLog (HLL) kernel to estimate cardinality of a data set

- Bump-in-the-wire kernel
- Cardinality estimation can augment the optimizer in data processing systems



Evaluation – StRoM NIC

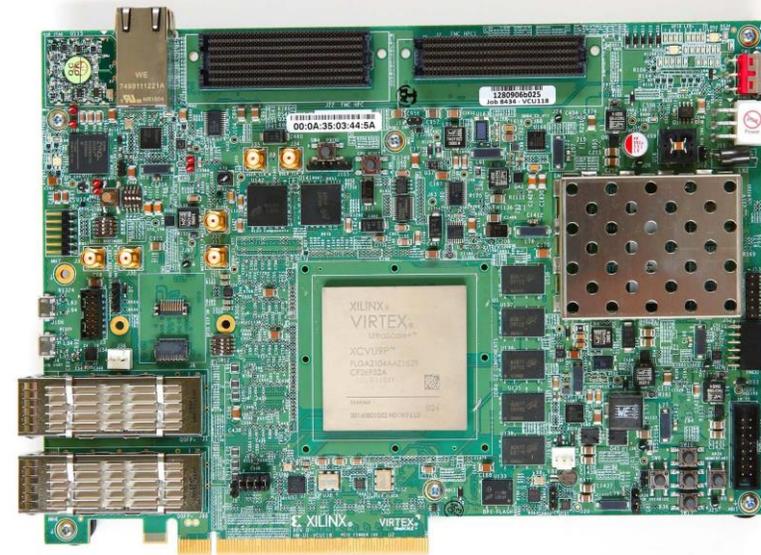
- FPGA-based prototype RDMA NIC
- Extended RoCEv2 implementation with support for StRoM

StRoM at 10G



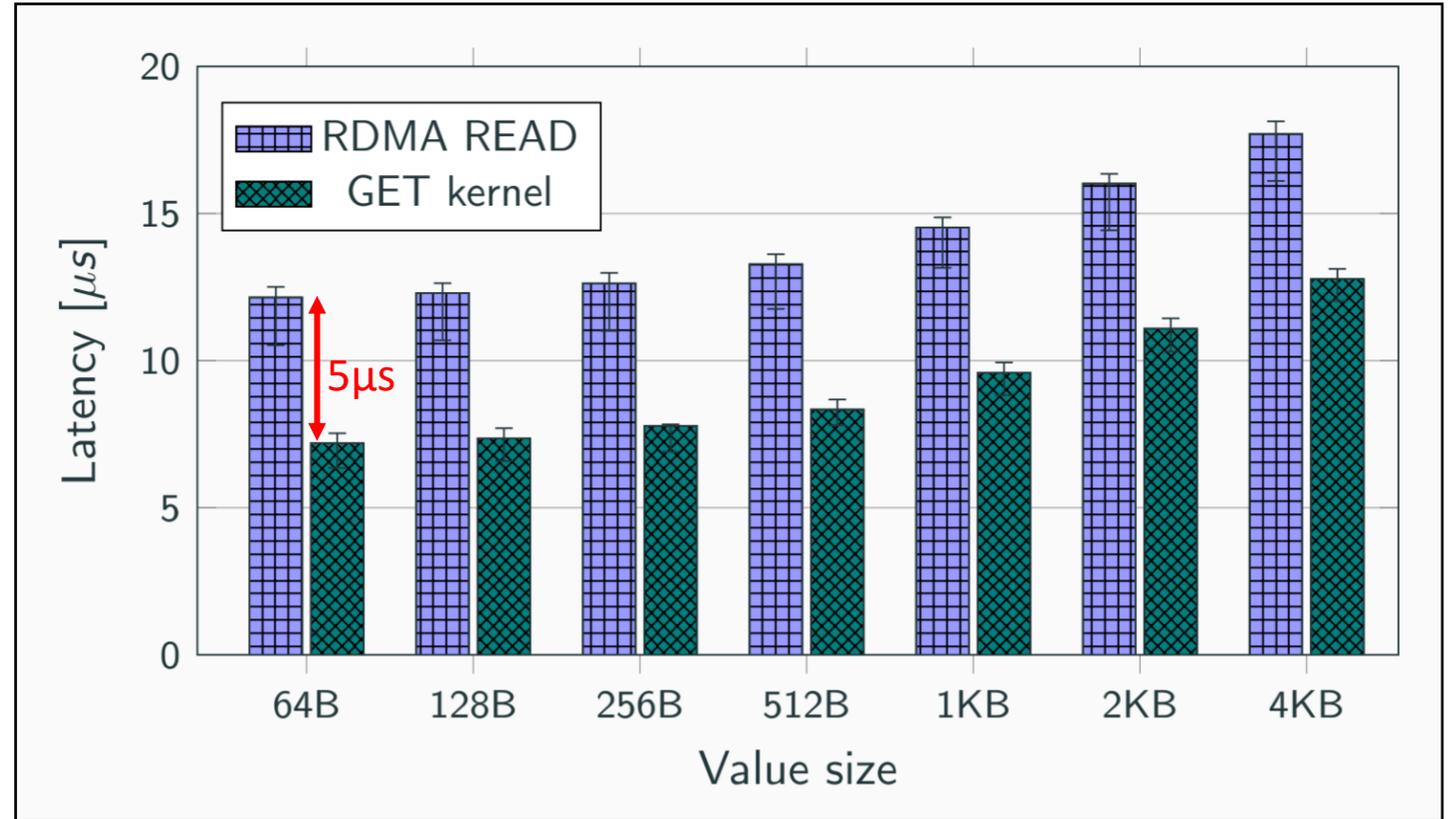
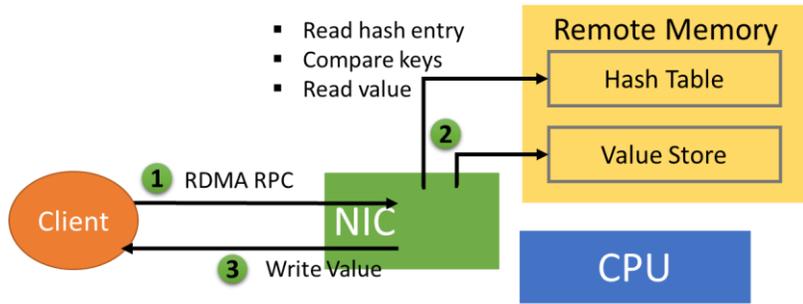
Alpha Data ADM-PCIE-7V3

StRoM at 100G

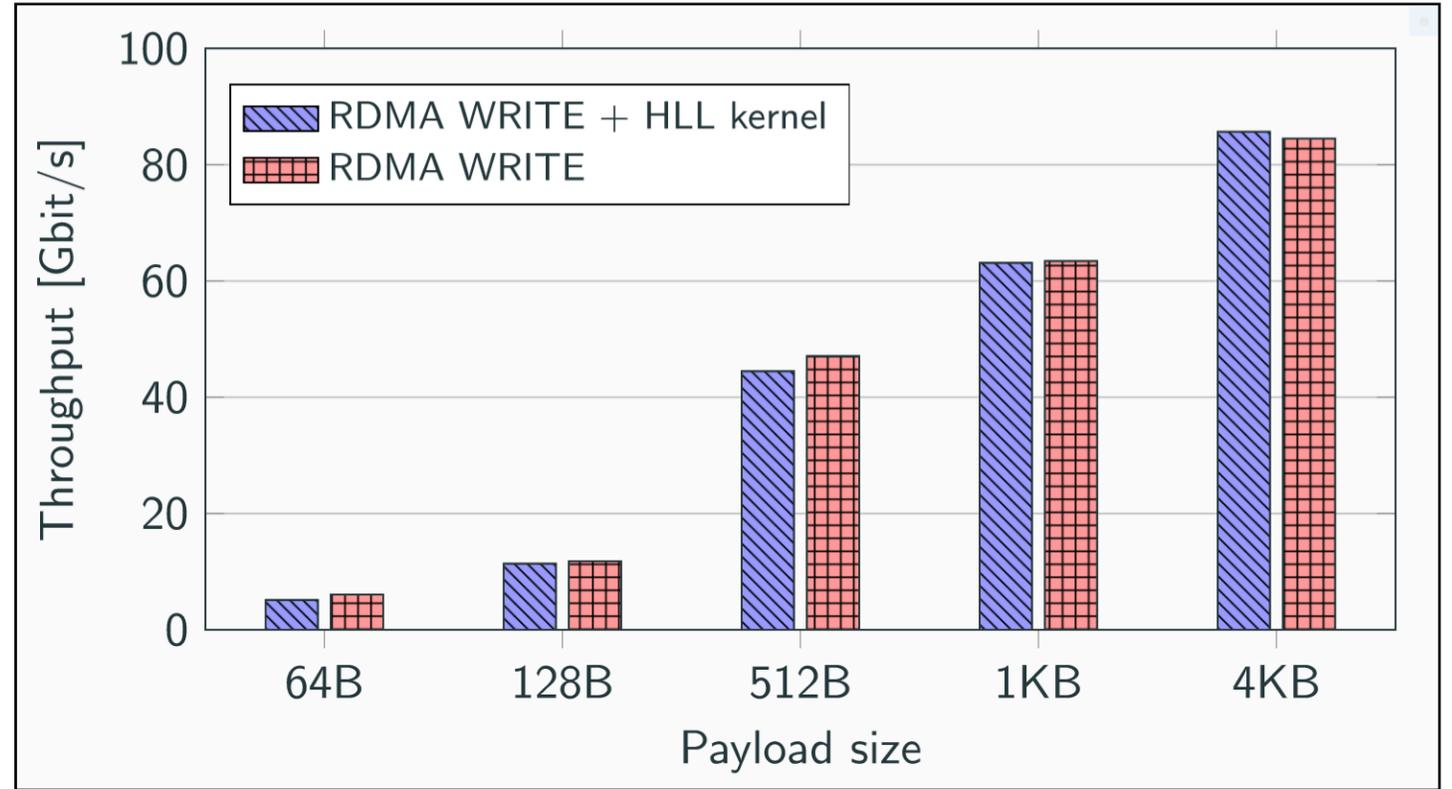
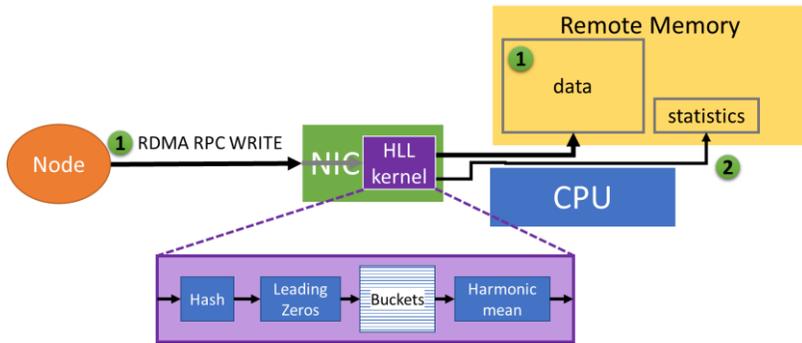


Xilinx VCU118

Evaluation – GET kernel



Evaluation – HLL kernel



Conclusion

StRoM: Smart Remote Memory

- Deployment of acceleration kernels on the NIC
- Acceleration of data access and data processing at up to 100G
- Research platform



Open source at github.com/fpgasystems/fpga-network-stack