

ALGORITHMS IN LOGIC



[HTTP://ALGO-LOGIC.COM](http://ALGO-LOGIC.COM)

*Comparison of Key/Value Store (KVS)
in Software and Programmable Hardware*

John W. Lockwood, CEO: Algo-Logic Systems, Inc.

<http://Algo-Logic.com> • Solutions@Algo-Logic.com • (408) 707-3740 • 2255-D Martin Ave., Santa Clara, CA 95050

Why Share Data by Name (Key) Instead of Address?



Examples:

Key

Value

Directory

Company

Algo-Logic

Phone #

(408) 707-3740

Forwarding
Tables

IP Address

204.2.34.5

Interface : MAC Address

Eth6 : 02:33:29:F2:AB:CC

Data De-
duplication

Content Hash

XYZ

Storage Block ID

948830038411

Stock Trading

Order ID

ATY11217911101

Symbol, Side, Price

AAPL, B, 126.75

Graph Search

Virtex

v140

Edge List

v201, v206, v225

- **Key/Value Store (KVS)**

- Simplifies implementation of large-scale distributed computation algorithms
- Data Center Servers exchanges data over standard Ethernet

- **Challenges**

- Operating System delays packets and limits throughput
- Per-core processing inefficient at high-speed packet processing

- **Solutions**

- Bypass kernel bypass with DPDK
- Offload of packet processing with FPGA



Why the Move to Programmable Hardware?

“There are large challenges in scaling the performance of software now. The question is: ‘What’s next?’ We took a bet on programmable hardware.”
- Doug Burger, Microsoft

- **Driving Metrics in the Data Center**

- Latency:

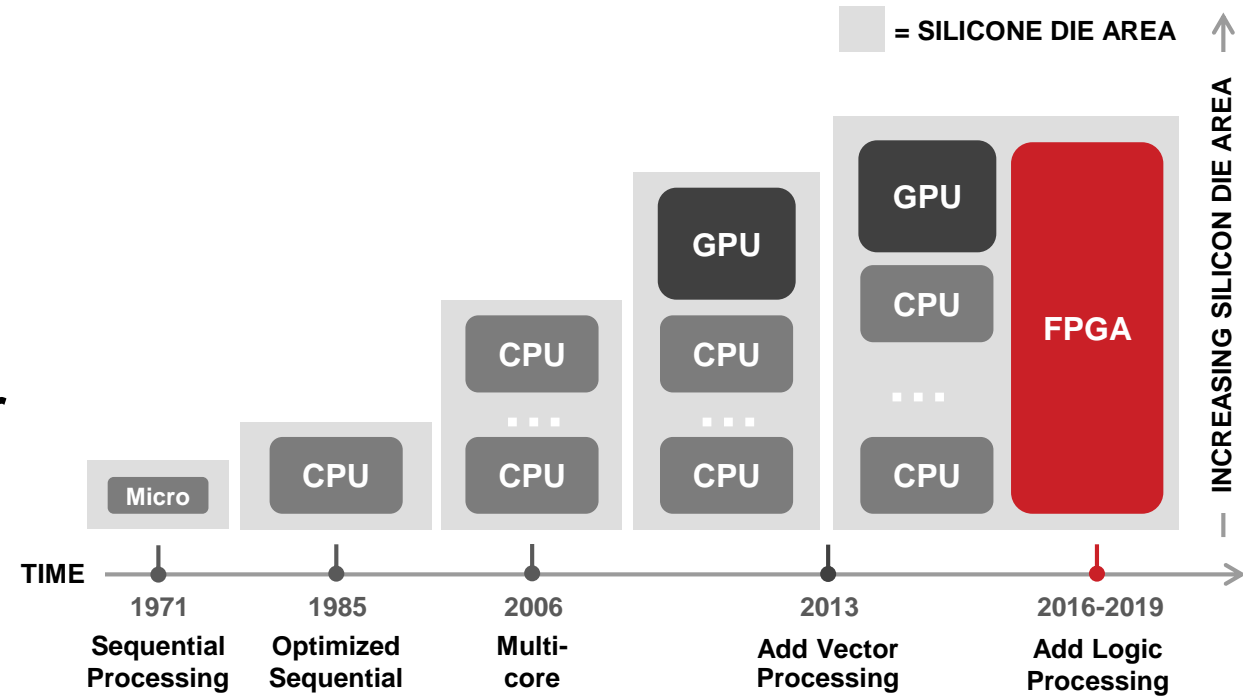
- Reduce delay
 - Avoid jitter

- Throughput

- Processing packets at line rate
 - Handle 10G, 25G, 40G, and 100G

- Power:



- Driving cost of OpEx



- Field Programmable Gate Array (FPGA) logic moves into the CPU
- Microsoft accelerates BING search with FPGA
- Intel acquires Altera

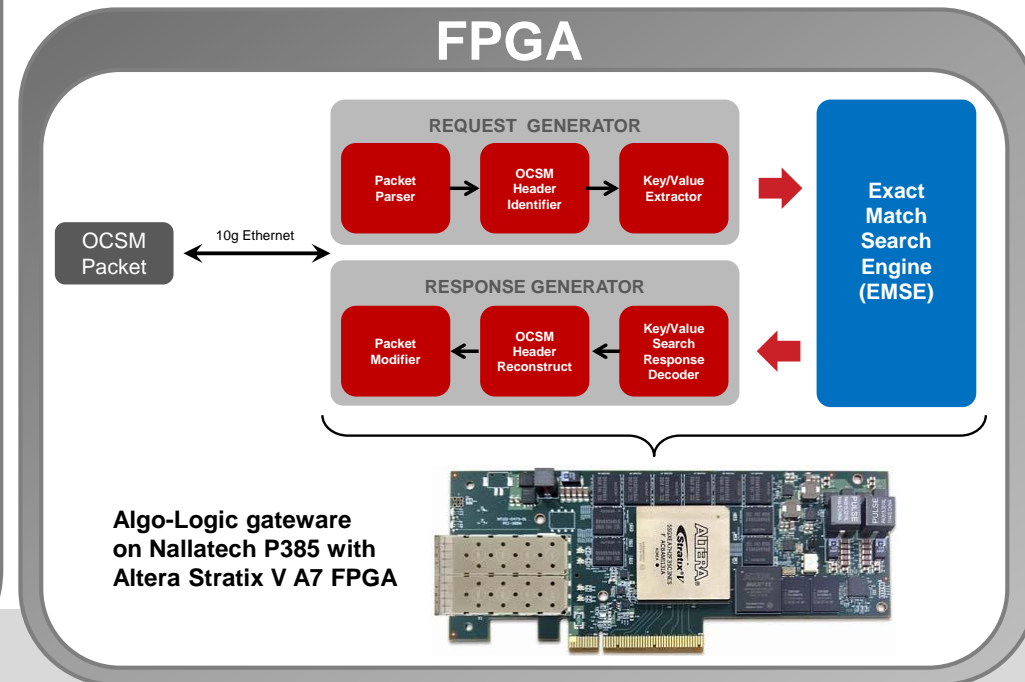
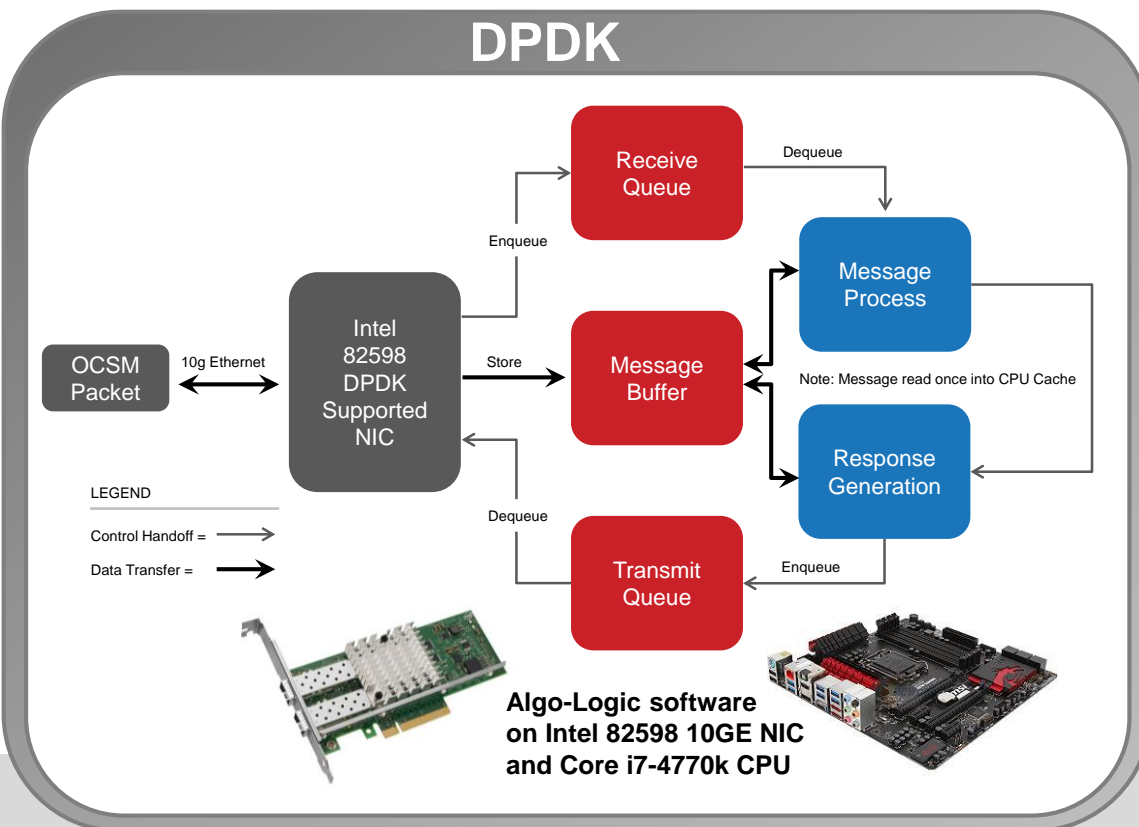
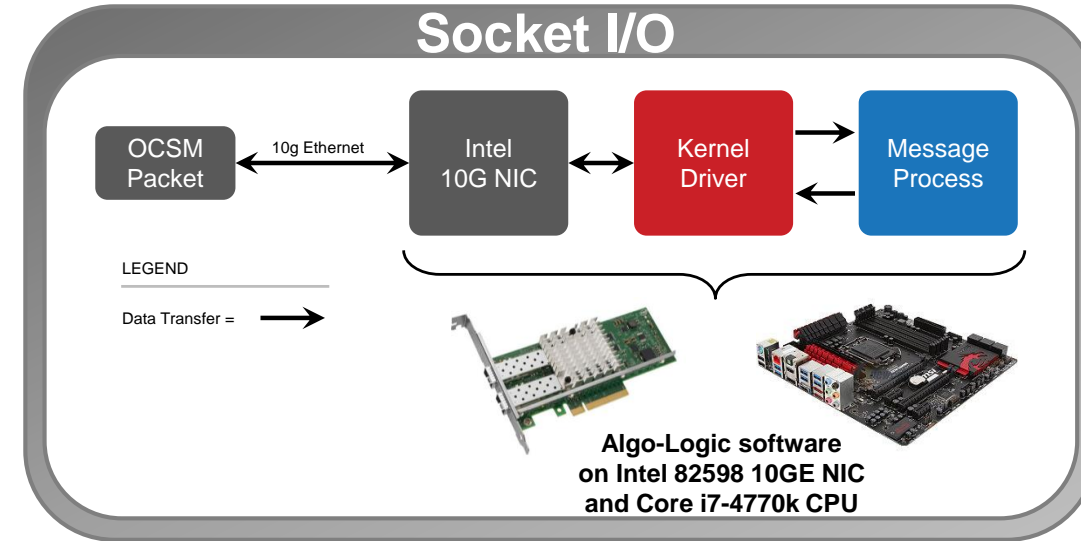
Servers Accelerated with FPGA Gateway

- **FPGA Augments Existing Servers**
 - Can run on an expansion card (same size as a GPU)
 - Or may be integrated into the CPU socket
- **GDN Applications run on FPGA**
 - Implements low-latency, low-power, high-throughput data processing

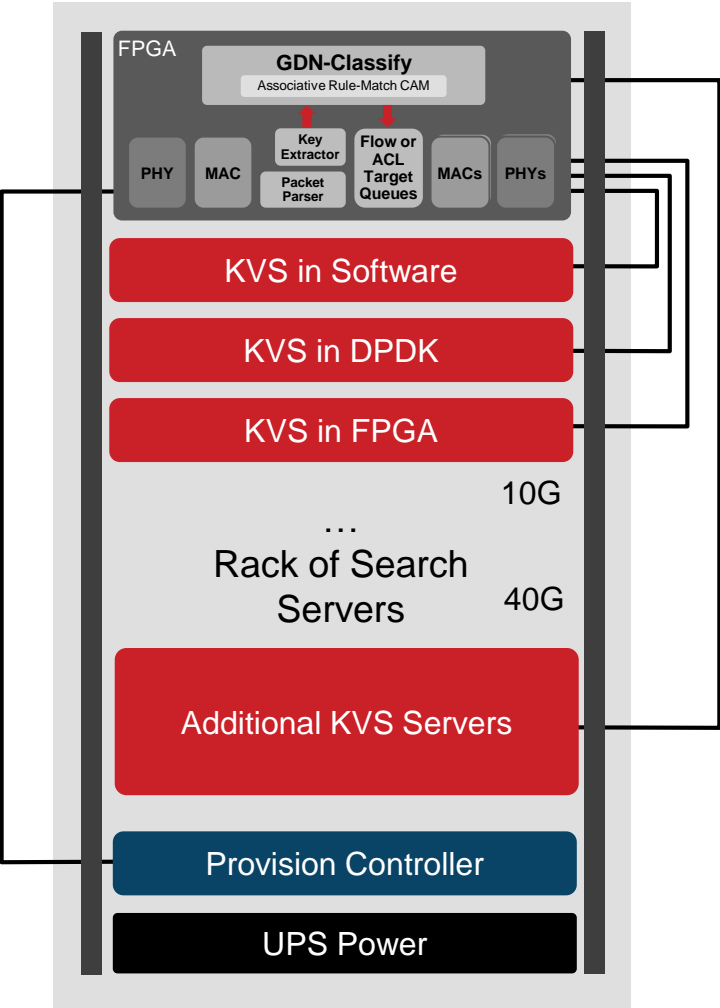


Implementation of KVS with Socket I/O, DPDK, and FPGA

- **Benchmark same application**
 - Key/Value Store (KVS)
- **Running on the same PC**
 - Intel i7-4770k CPU, 82598 NIC, and Altera Stratix V A7 FPGA
- **With three different implementations**
 - Socket I/O, DPDK, FPGA



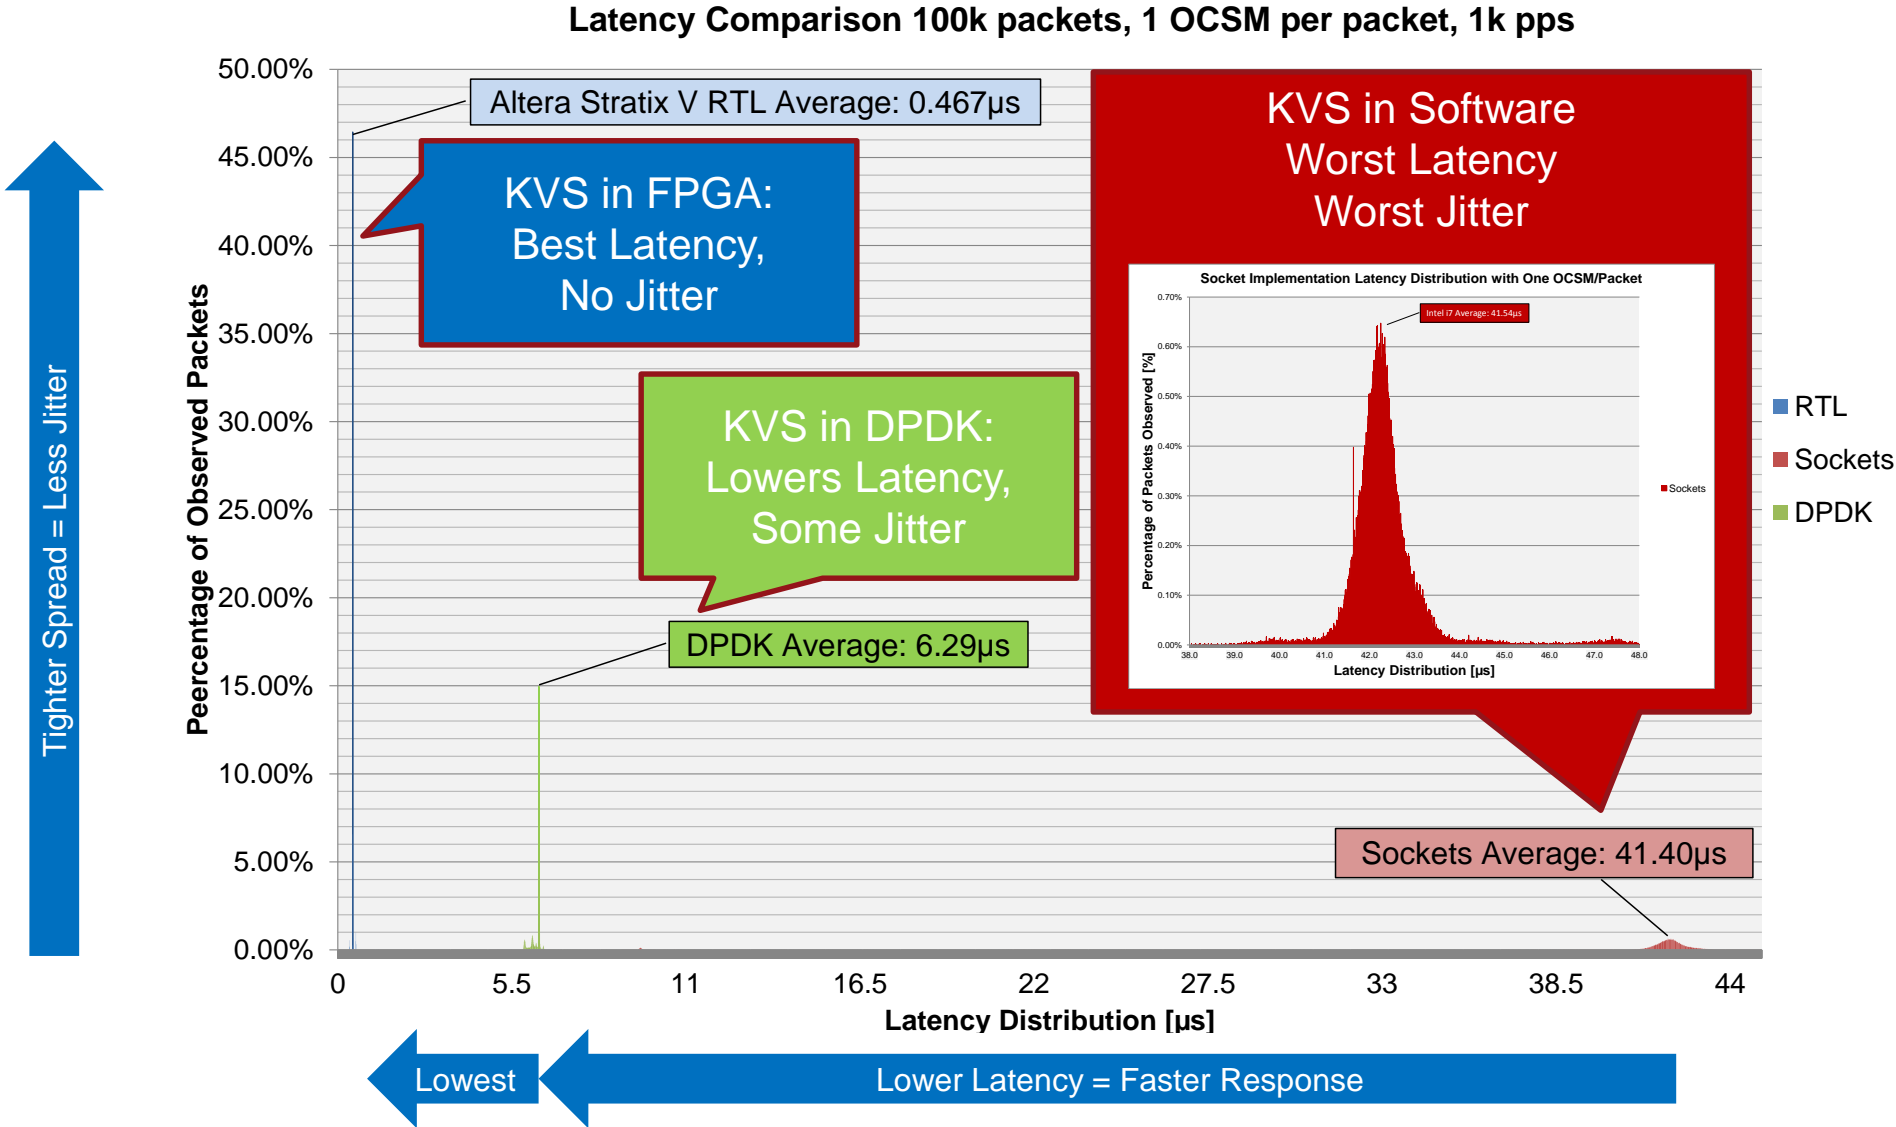
Measured Latency, Throughput, and Power Results



All Datapaths Summary	Latency (μseconds)	Tested Throughput (CSMs/sec)	Power (μJoules/CSM)
Sockets	41.54	4.0	11
DPDK	6.434	16	6.6
RTL	0.467	15	0.52

All Datapaths Summary	Latency (μseconds)	Maximum Throughput (CSMs/sec)	Power (μJoules/CSM)
GDN vs. Sockets	88x less	13x	21x less
GDN vs. DPDK	14x less	3.2x	13x less

KVS Latency in FPGA, DPDK, and Sockets



Conclusions: Key/Value Store in Programmable Hardware

- **Lowers Latency**
 - 88x faster than Linux networking sockets
 - 14x faster than optimized DPDK (kernel bypass)
- **Increases Throughput (IOPs)**
 - 3x to 13x improvement in throughput
 - Lowers Capital Expenditures (CapEx)
- **Reduces Power**
 - 13x to 21x reduction in power
 - Reduces Operating Expenditures (OpEx)

Gateway Defined Networking® dramatically reduces latency and power and improves throughput in the data center