



An FPGA Memcached Appliance

Sai Rahul Chalamalasetti^{†*}, Kevin Lim[‡], Mitch Wright, Alvin Auyoung[‡],
Parthasarathy Ranganathan[‡], Martin Margala^{*}

February 13th, 2013

[†] HP, Houston, TX

[‡] HP Labs, Palo Alto, CA

^{*} University of Massachusetts Lowell, Lowell, MA

© Copyright 2013 Hewlett-Packard Development Company, L.P. The information contained herein is subject to change without notice.

Background

The Web is driving an increase in structured and unstructured data

Different applications to access and analyze the unprecedented amount of data

Memcached has become de facto standard for low latency access to data

Key innovation is all **in-memory storage**

Major organizations with entire tiers of memcached servers

Facebook, Zynga, Twitter are among top users

Other in-memory Database applications

MonetDB, VoltDB, and HANA

Contribution

Implemented **base memcached appliance** on a **standalone FPGA board**

Our design vs. CPU based systems

- Performance **on-par** with baseline servers
- Consumes **9%** of the power of the baseline
- **3.2X to 10.9X** Energy efficiency improvement
- **Significant total-cost-of-ownership** improvements at Data Center Scale



Outline

**Memcached
Application
Overview**

**CPU
Challenges
vs. FPGA
Opportunities**

**Memcached
Appliance on
FPGA**

Results

Summary

Outline

**Memcached
Application
Overview**

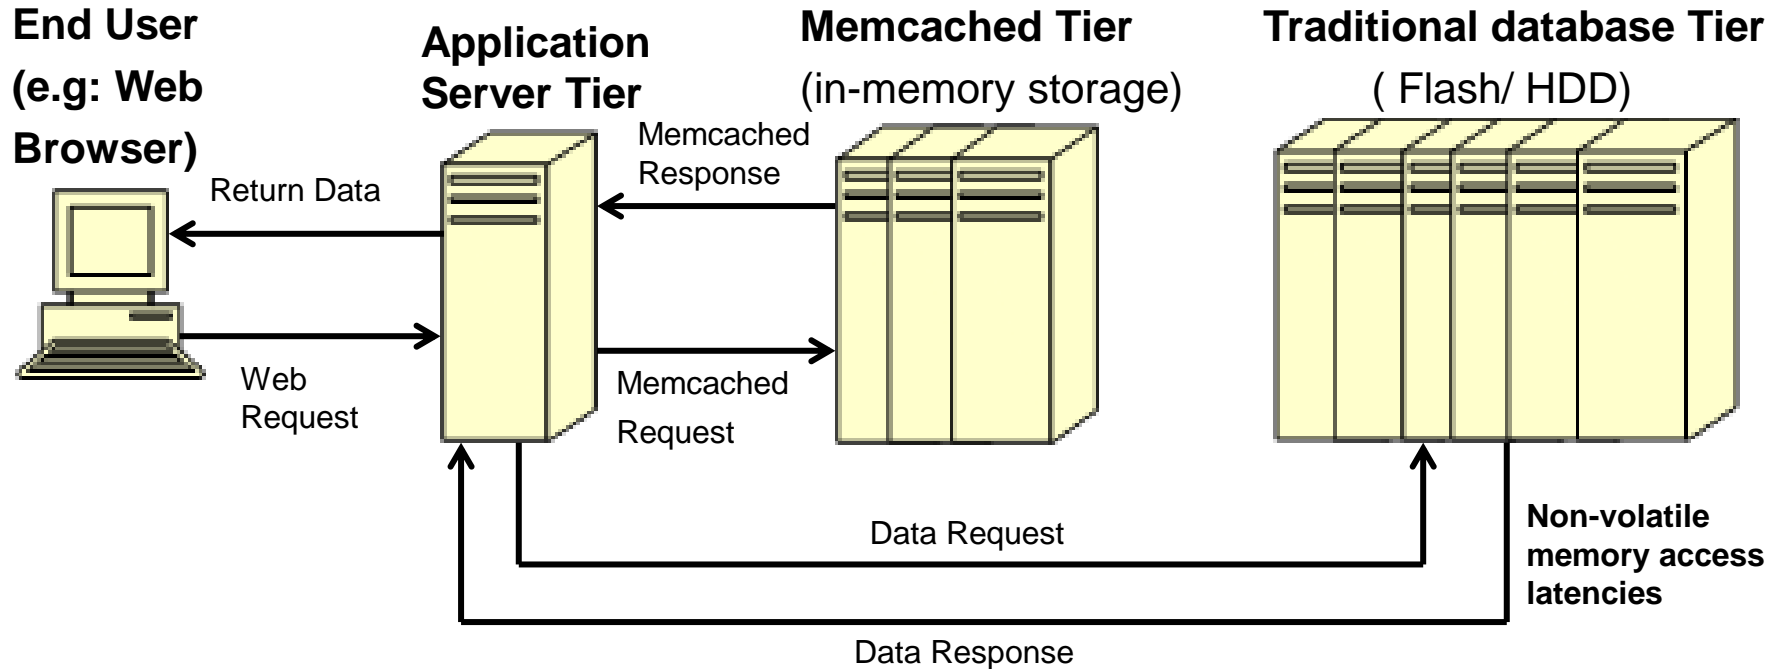
**CPU
Challenges
vs. FPGA
Opportunities**

**Memcached
Appliance on
FPGA**

Results

Summary

Memcached Application



Memcached Application

- **Distributed in-memory** caching system
- Used to speed up **dynamic-driven websites** through caching smaller data and objects into RAM
- Uses **distributed hash table** to map the data on multiple machines
- Originally developed by **Danga Interactive** for Live Journal
- **Two main operations** are: **Set**(write), and **Get**(read)
- Numerous transactions from web makes the memcached application **network intensive**

Outline

**Memcached
Application
Overview**

**CPU
Challenges
vs. FPGA
Opportunities**

**Memcached
Appliance on
FPGA**

Results

Summary

CPU challenges vs. FPGAs opportunities

Network Processing on CPUs

- Round trip packet latency of NIC is 60 μs ¹
- Linux Kernel is 30 μs ¹
- Memcached appliance only take 30 μs on CPU

Power Consumption on CPUs

- Intel Xeon CPUs
 - 258W for 12 core(two socket) CPU with 64GB of DRAM
 - 190W TDP is for two sockets

Network Processing on FPGAs

- FPGAs are suitable for applications that require processing data on wire speed
 - Network processing
 - Matured IP(TCP/UDP OE) base to map networking blocks

Power Consumption on FPGAs

- Altera FPGAs
 - 18.5W for FPGA board with 8GB RAM
- Programmable Gate Arrays
 - Application can be customized to utilize necessary hardware resources
 - Lower frequency of operation

Ref1 Low Latency RPC in RAMCloud, Rosenblum .et al



Outline

**Memcached
Application
Overview**

**CPU
Challenges
vs. FPGA
Opportunities**

**Memcached
Appliance on
FPGA**

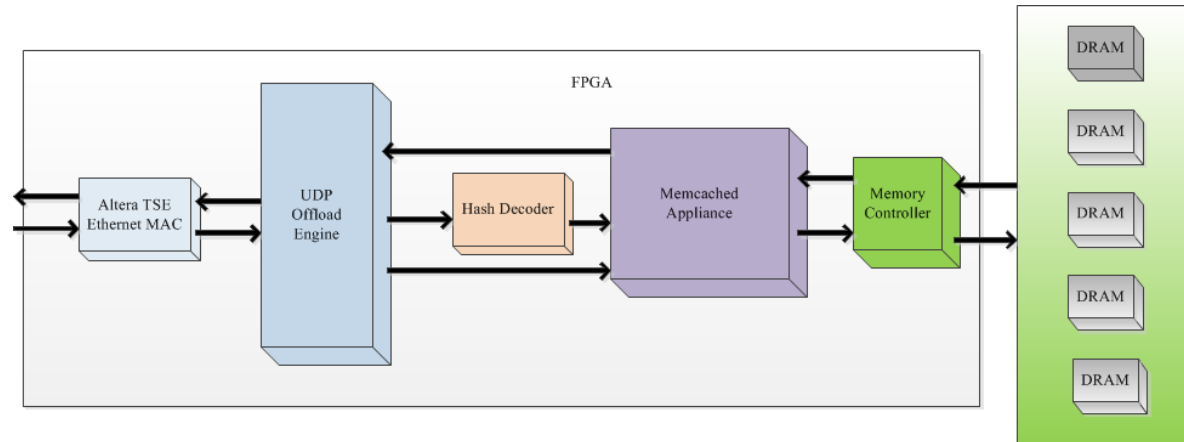
Results

Summary

Memcached Appliance on FPGA

Network Processing

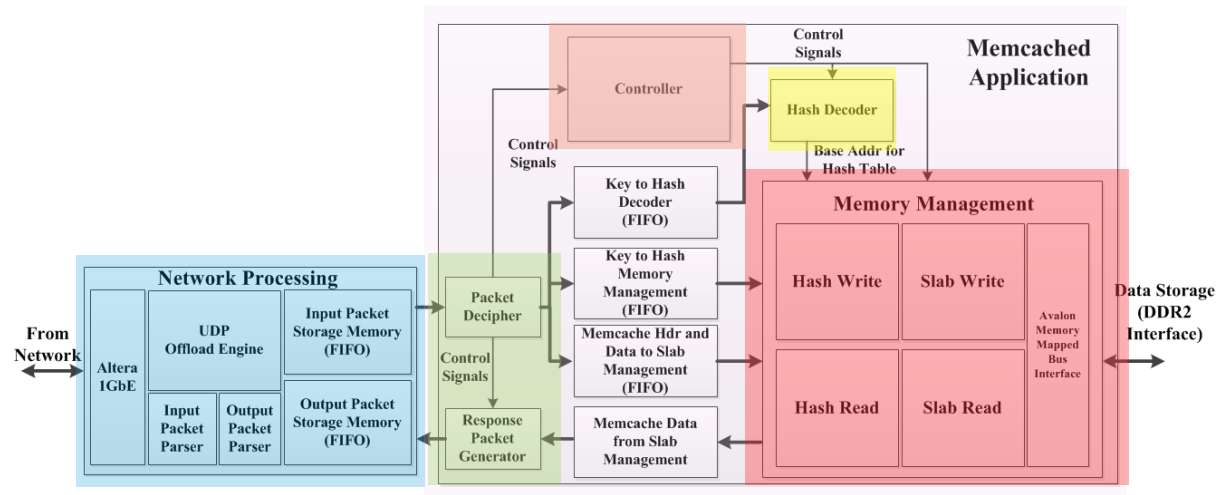
- 1GbE Altera IP
- UDP Offload Engine
- Multi-request packet data storage



Standalone FPGA Memcached Appliance

Memcached Application

- Packet Decipher
- Response Packet Generator
- Controller
- Hash Decoder
- Memory Management
 - Slab Memory
 - Hash Memory

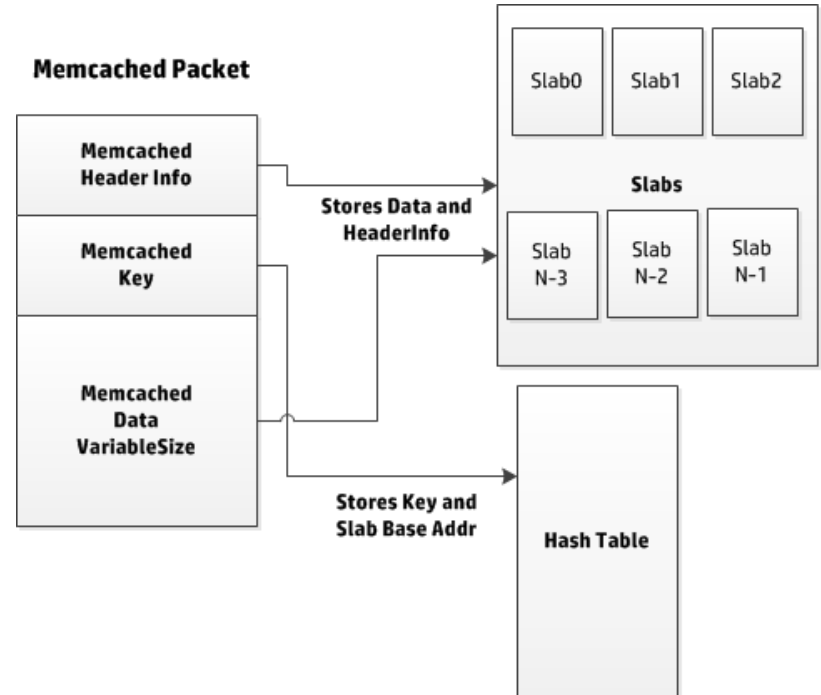


Block Diagram of Memcached Appliance on FPGA

Memcached Memory Management

Two different data structures

- Hash Memory
 - Stores basic memcached header information
 - **Key**
 - **Data** size
 - Used to find key-value pairs
- Slab Memory
 - Stores **user data**
 - Variable size data storage



Replacement Policy

Least Recently Used (LRU) is used in software

Round Robin (RR) is used in hardware

Outline

**Memcached
Application
Overview**

**CPU
Challenges
vs. FPGA
Opportunities**

**Memcached
Appliance on
FPGA**

Results



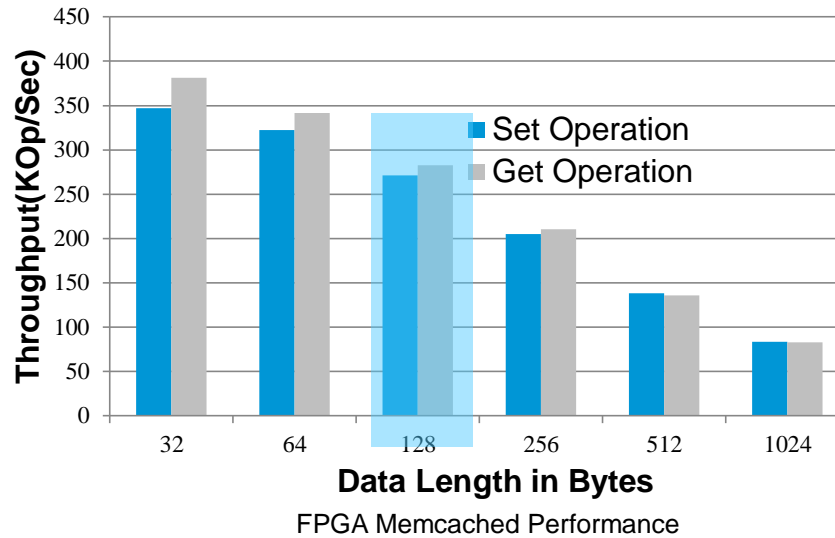
Summary

Testing Methodology and Utilization Results



- Micro Benchmark to generate set and get operations
- Java Program to send UDP packets to FPGA IP Address
 - Read trace file, and pass the packets on to network
 - Wait till acknowledgment packet from FPGA memcached appliance
- Measure performance through counters inserted for each hardware block
- **Board Specifications**
 - Altera DE-4 Development Board
 - Stratix IV 530 FPGA
 - Interface to 8GB DDR2 memory
- **The proposed design utilizes**
 - Single Memcached Instance : **7%** of total Stratix IV FPGA resources
 - Dual Memcached Instance : **13%** of total Stratix IV FPGA resources

Performance Results



Energy Efficiency and TCO

System	Power Consumption (W)	Performance (KOps/Sec)	Performance/Watt (KOps/Sec/W)
CPU ₁	258 [†]	1000	4.95
CPU ₂	258 [†]	300	1.49
FPGA ₁	17.4	283	16.26
FPGA ₂	18.84	566	30.04

Power Consumption and Energy Efficiency

Total Cost of Ownership (TCO)

- **Two CPU systems with 64GB**
 - CPU1 with Optimistic Performance : 1000 KOps/Sec
 - CPU2 with Realistic Performance: 300 KOps/Sec
- **Four FPGA system configurations**
 - Cost of FPGA: \$1000, \$3000
 - Memory density on the board: 16GB, 32GB
- **A best case FPGA cost and memory density board vs. CPU1**
 - **Performance per dollar** improvement of **3.5X**

[†] Power Consumption estimated by using various device and component data sheets



Outline

**Memcached
Application
Overview**

**CPU
Challenges
vs. FPGA
Opportunities**

**Memcached
Appliance on
FPGA**

Results

Summary

Summary

Implemented an **FPGA-based standalone base memcached** system

The proposed design achieves about **280KOps/sec**

Low resource utilization enable multiple(two) memcached appliances on one FPGA

- Achieves performance of about **565KOps/sec**
- Only increases overall board power consumption by **3.9%**

Our system achieves energy efficiency improvements of **3.2X to 10.9X**

A significant improvement in Total-cost-of-ownership at Data Center Level

Accelerators offer promise in the era of data-centric computing



Thank you

