

Low-Latency TCP/IP Stack for Data Center Applications

David Sidler, Zsolt István, Gustavo Alonso Systems Group, Dept. of Computer Science, ETH Zürich

Systems Group, Dept. of Computer Science, ETH Zürich

Original Architecture[1]

Scalable 10 Gbps TCP/IP Stack Architecture for Reconfigurable Hardware

David Sidler, Gastaro Alonso	Michaela Blott, Kimon Kamus, Kees Vosers	Reymond Carley
Suptoma Group	Xilinx Research	Dept. of Electrical
Dept. of Computer Science	Dublin, holand	& Computer Engineering
ETH Zarish	& San Jess, CA	Camegic Middon Universit
[dasidler, alonse]/@inf.ethit.ch	(rabbott, kimonk, kees, vissers) @xillins.com	scarley/iPece.cmm.edu

As one -1537 % for production or constrained on product Consequently, ICFFF effects is known is strained as the strained strained product the strained st

stark. This paper initializes a source excitation for a 100 kpc binary of the stark of the transmission of the stark of t

1. INTRODUCTION

TCP/IP is the consensione of modern network communitions with its support for reliable data transfer including ov control, congestion avoidance, duplicate data suppression

or constraints application [3]. As such as the second probability integrated in terms of the second probability of the se

time of sections. To institute the stack's applicables, it was essential to creat a fasible solution that allows to efficiently and easily adapt the design to different competition avoidness schemes, prioritisk ent of order processing, etc., while using a maintainal reasons.

energy to universe congenerative proteines proteines, protein and of-order processing, etc., while using a minimal resource frontprint. To achieve this, we adopted a Cu+-based design flux using high-level synthesis (HLS) that simplifies the design prairies it more faculties and matter to constraine towards receils

- 10 Gbps bandwidth TCP/IP stack
- Supporting thousands of concurrent connections
- Generic implementation as close to specification as possible
- Enables seamless integration of FPGA-based applications into existing networks

[1] Sidler et al., Scalable 10 Gbps TCP/IP Stack Architecture for Reconfigurable Hardware, FCCM'15, http://github.com/dsidler/fpga-network-stack

Systems Group, Dept. of Computer Science, ETH Zürich

FPL'16, Lausanne | August 30, 2016 | 2 / 6

Application Integration



Application Integration



• Memory bandwidth is shared among multiple modules \rightarrow potential bottleneck

Systems Group, Dept. of Computer Science, ETH Zürich

FPL'16, Lausanne | August 30, 2016 | 3 / 6

Application Integration



 \blacksquare Memory bandwidth is shared among multiple modules \rightarrow potential bottleneck

 \blacksquare Distributed systems rely on very low latency \rightarrow to guarantee latency bounds to clients

Systems Group, Dept. of Computer Science, ETH Zürich

FPL'16, Lausanne | August 30, 2016 | 3 / 6

Assumptions

- Application
 - Client requests fit into an MTU (maximum transfer unit)
 - Synchronous clients
 - Application logic consumes data at line-rate

Assumptions

- Application
 - Client requests fit into an MTU (maximum transfer unit)
 - Synchronous clients
 - Application logic consumes data at line-rate
- Data center network
 - High reliability and structured topology
 - \blacksquare Data loss less common \rightarrow fewer retransmission
 - Packets are rarely reordered

Optimizations for Data Center Applications



Systems Group, Dept. of Computer Science, ETH Zürich

FPL'16, Lausanne | August 30, 2016 | 5 / 6

Optimizations for Data Center Applications



Systems Group, Dept. of Computer Science, ETH Zürich

Optimizations for Data Center Applications



Results

Systems Group, Dept. of Computer Science, ETH Zürich

FPL'16, Lausanne | August 30, 2016 | 6 / 6

Results



Results



Results



	Mem. allocated	Mem. bandwidth
TCP org.	1,300 MB	40 Gbps
TCP opt.	650 MB	10 Gbps
Diff	-50%	-75%

Systems Group, Dept. of Computer Science, ETH Zürich

Results



These results enabled a consistent distributed key-value store [2]

	Mem. allocated	Mem. bandwidth
TCP org.	1,300 MB	40 Gbps
TCP opt.	650 MB	10 Gbps
Diff	-50%	-75%

[2] István et al., Consensus in a Box: Inexpensive Coordination in Hardware, NSDI16 Systems Group, Dept. of Computer Science, ETH Zürich FPL'16, Lausanne | August 30, 2016 | 6 / 6

Results



	Mem. allocated	Mem. bandwidth
TCP org.	1,300 MB	40 Gbps
TCP opt.	650 MB	10 Gbps
Diff	-50%	-75%

[2] István et al., Consensus in a Box: Inexpensive Coordination in Hardware, NSDI16 Systems Group, Dept. of Computer Science, ETH Zürich FPL'16, Lausanne | August 30, 2016 | 6 / 6