

Exploring RapidIO technology within a DAQ system event building network

Simaolhoda Baymani, Konstantinos Alexopoulos, Sébastien Valat
CERN

Introduction

RapidIO is a high-performance, low pin count, packet switched, system level interconnect architecture.

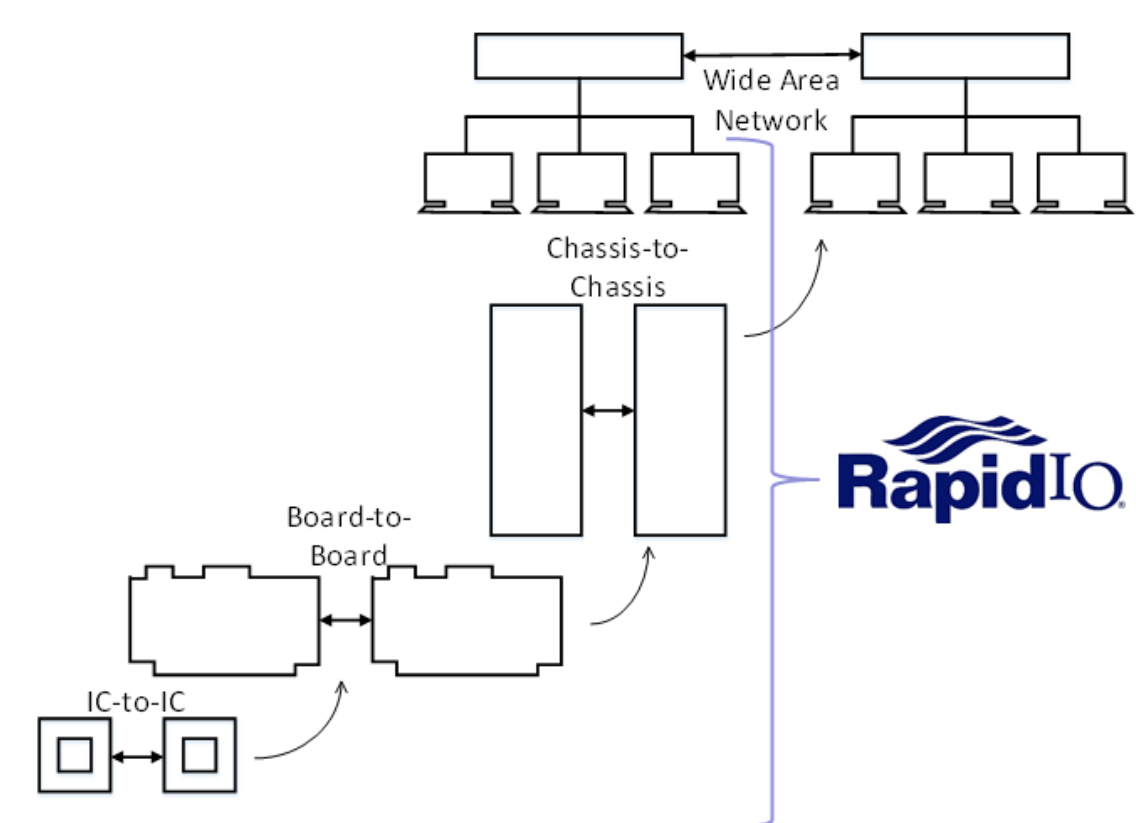
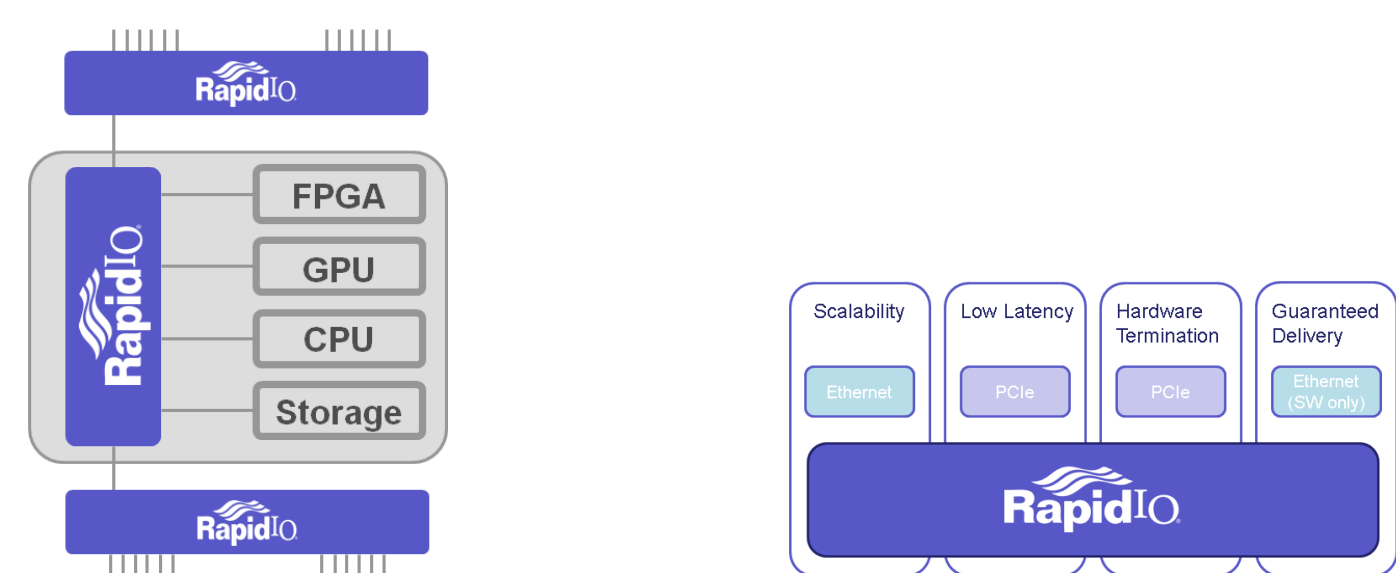


Figure: Interconnect application domains

Characteristics:

- Promises the combined strengths of PCI Express and Ethernet
- Supports heterogeneous systems (e.g. DSP and FPGA on the same fabric)
- Achieves minimal overhead through hardware implementation
- Supports error handling at the physical level



(a) Heterogeneous systems (b) Combined strengths

Data propagation can be achieved by utilizing:

- Channelized Messaging
- Direct Memory Access

Setup

- **Nodes:** 4 x Intel Xeon L5640 @ 2.27GHz, 48GB of RAM
- **NICs:** 4 x IDT PCIe to Rapidio Bridge, 14.5 Gbps
- **Switch:** 38-port RapidIO Generation 2 switch box, 20Gbps
- **Cables:** 4 x QSFP+ cables
- **OS:** CERN CentOS release 7.2.1511
- **RapidIO:** IDT Linux kernel drivers and libraries

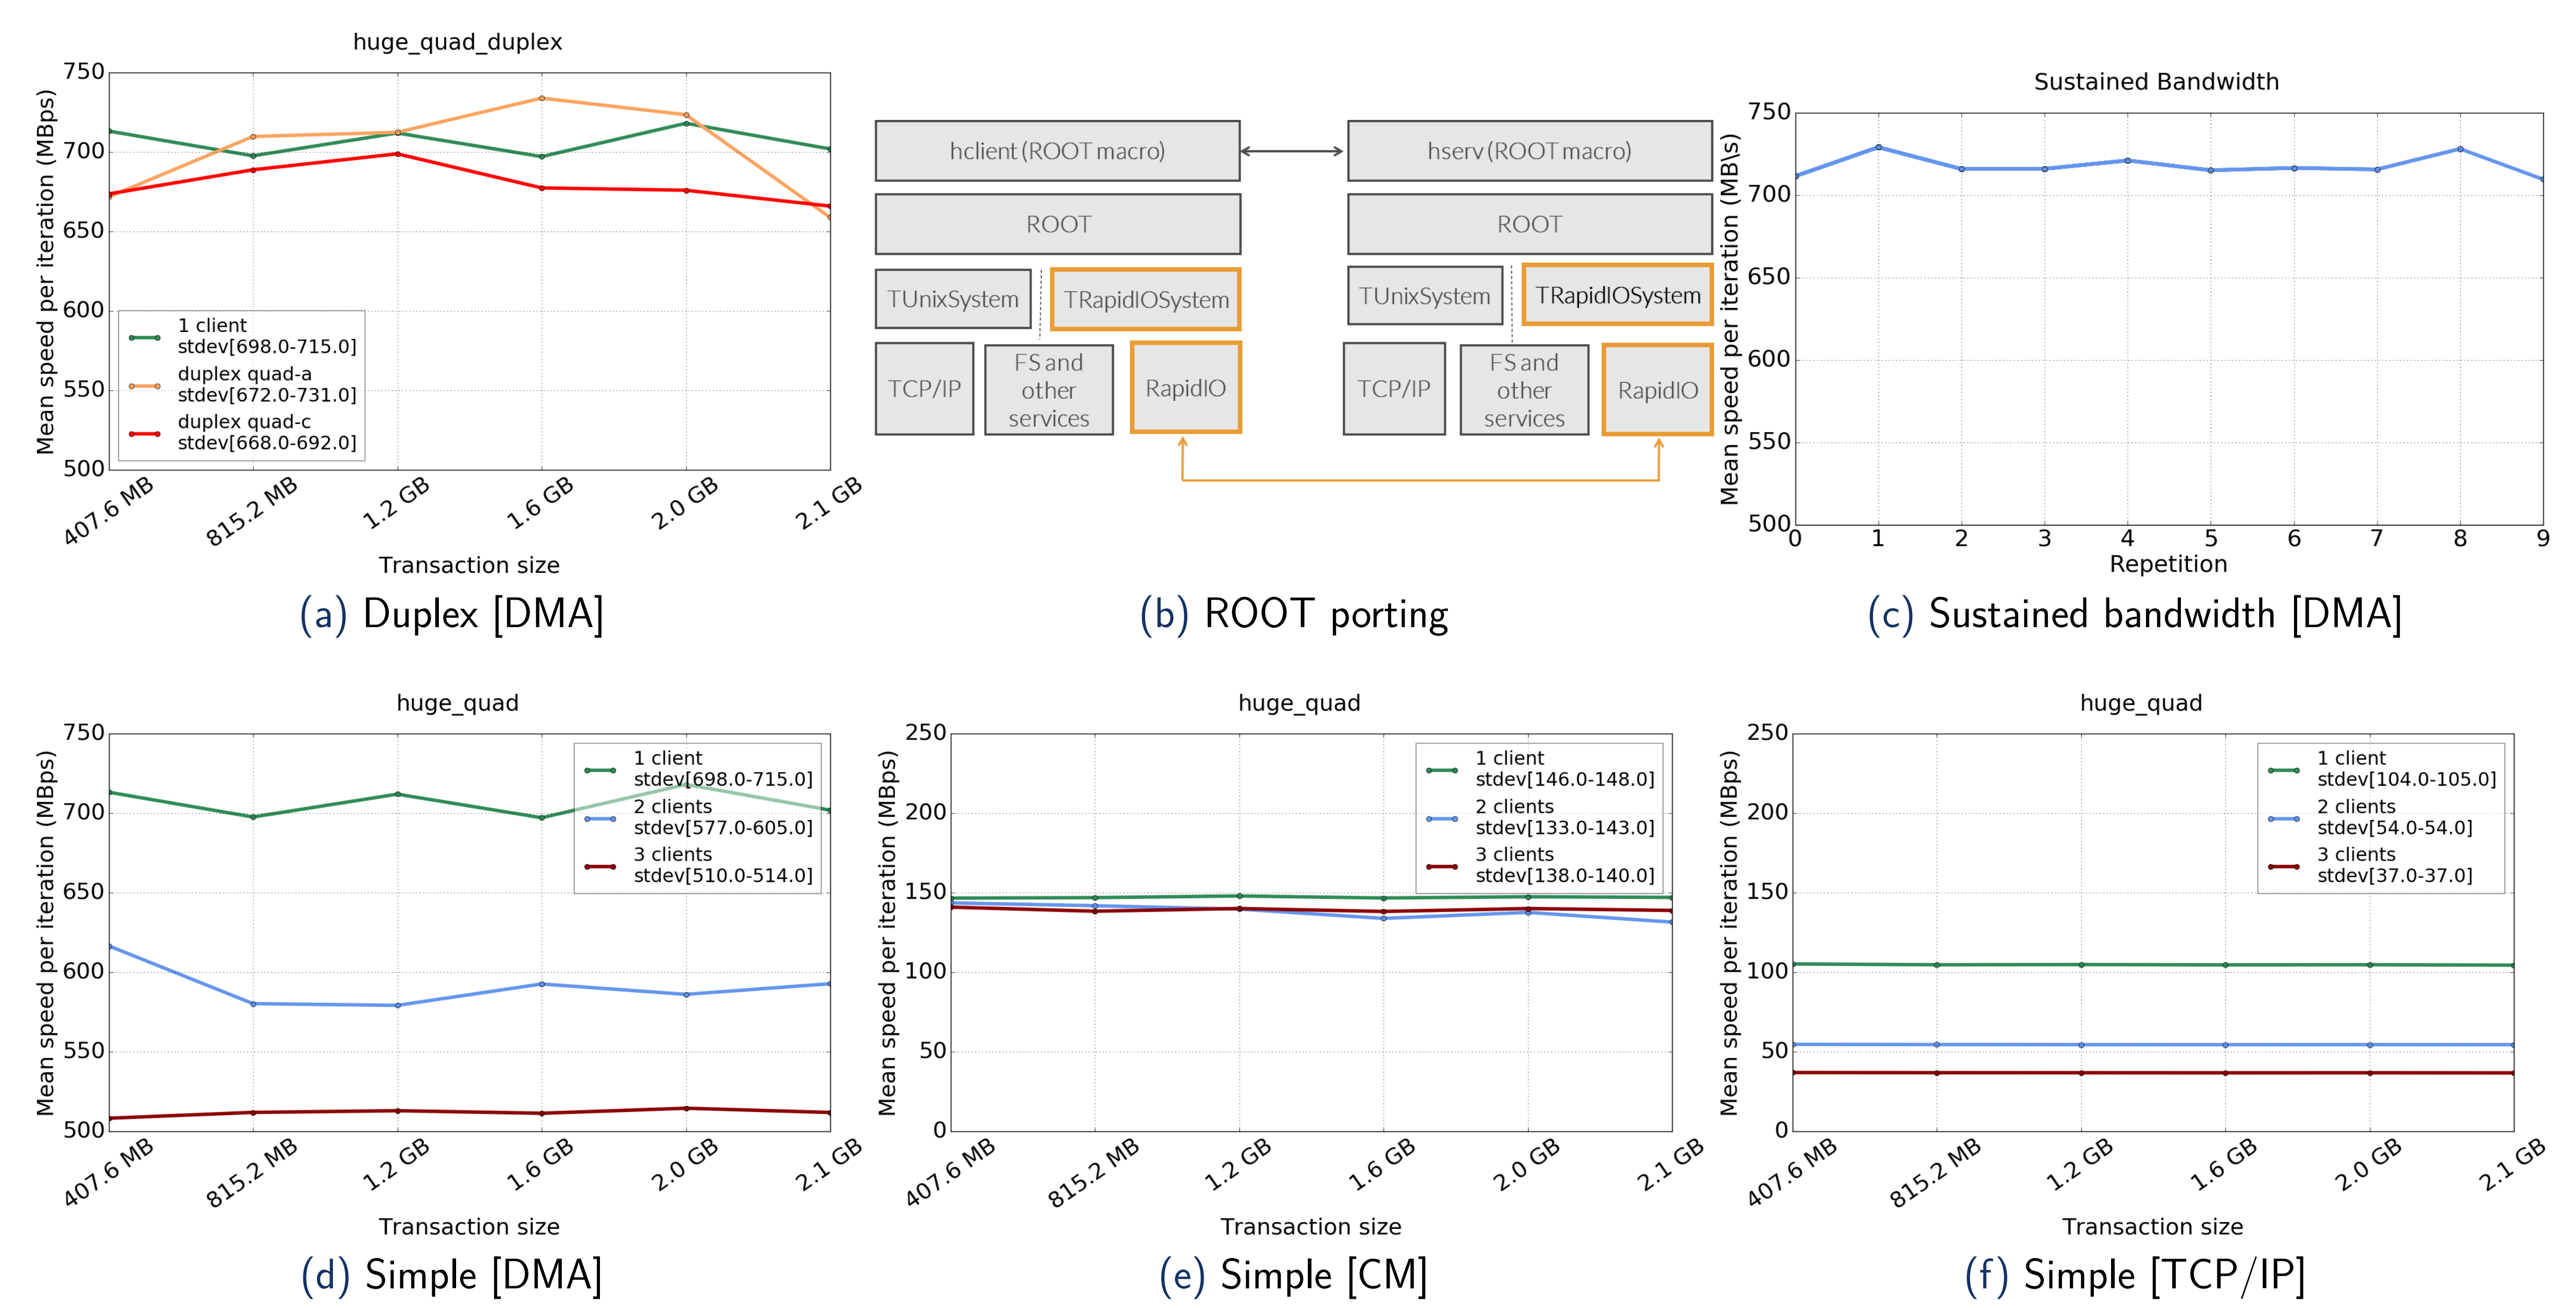
Acknowledgments

We would like to thank IDT for its membership and technical collaboration with CERN openlab, as well as for providing the hardware and software needed. We would also like to thank our colleagues at CERN for their input and support.



ROOT

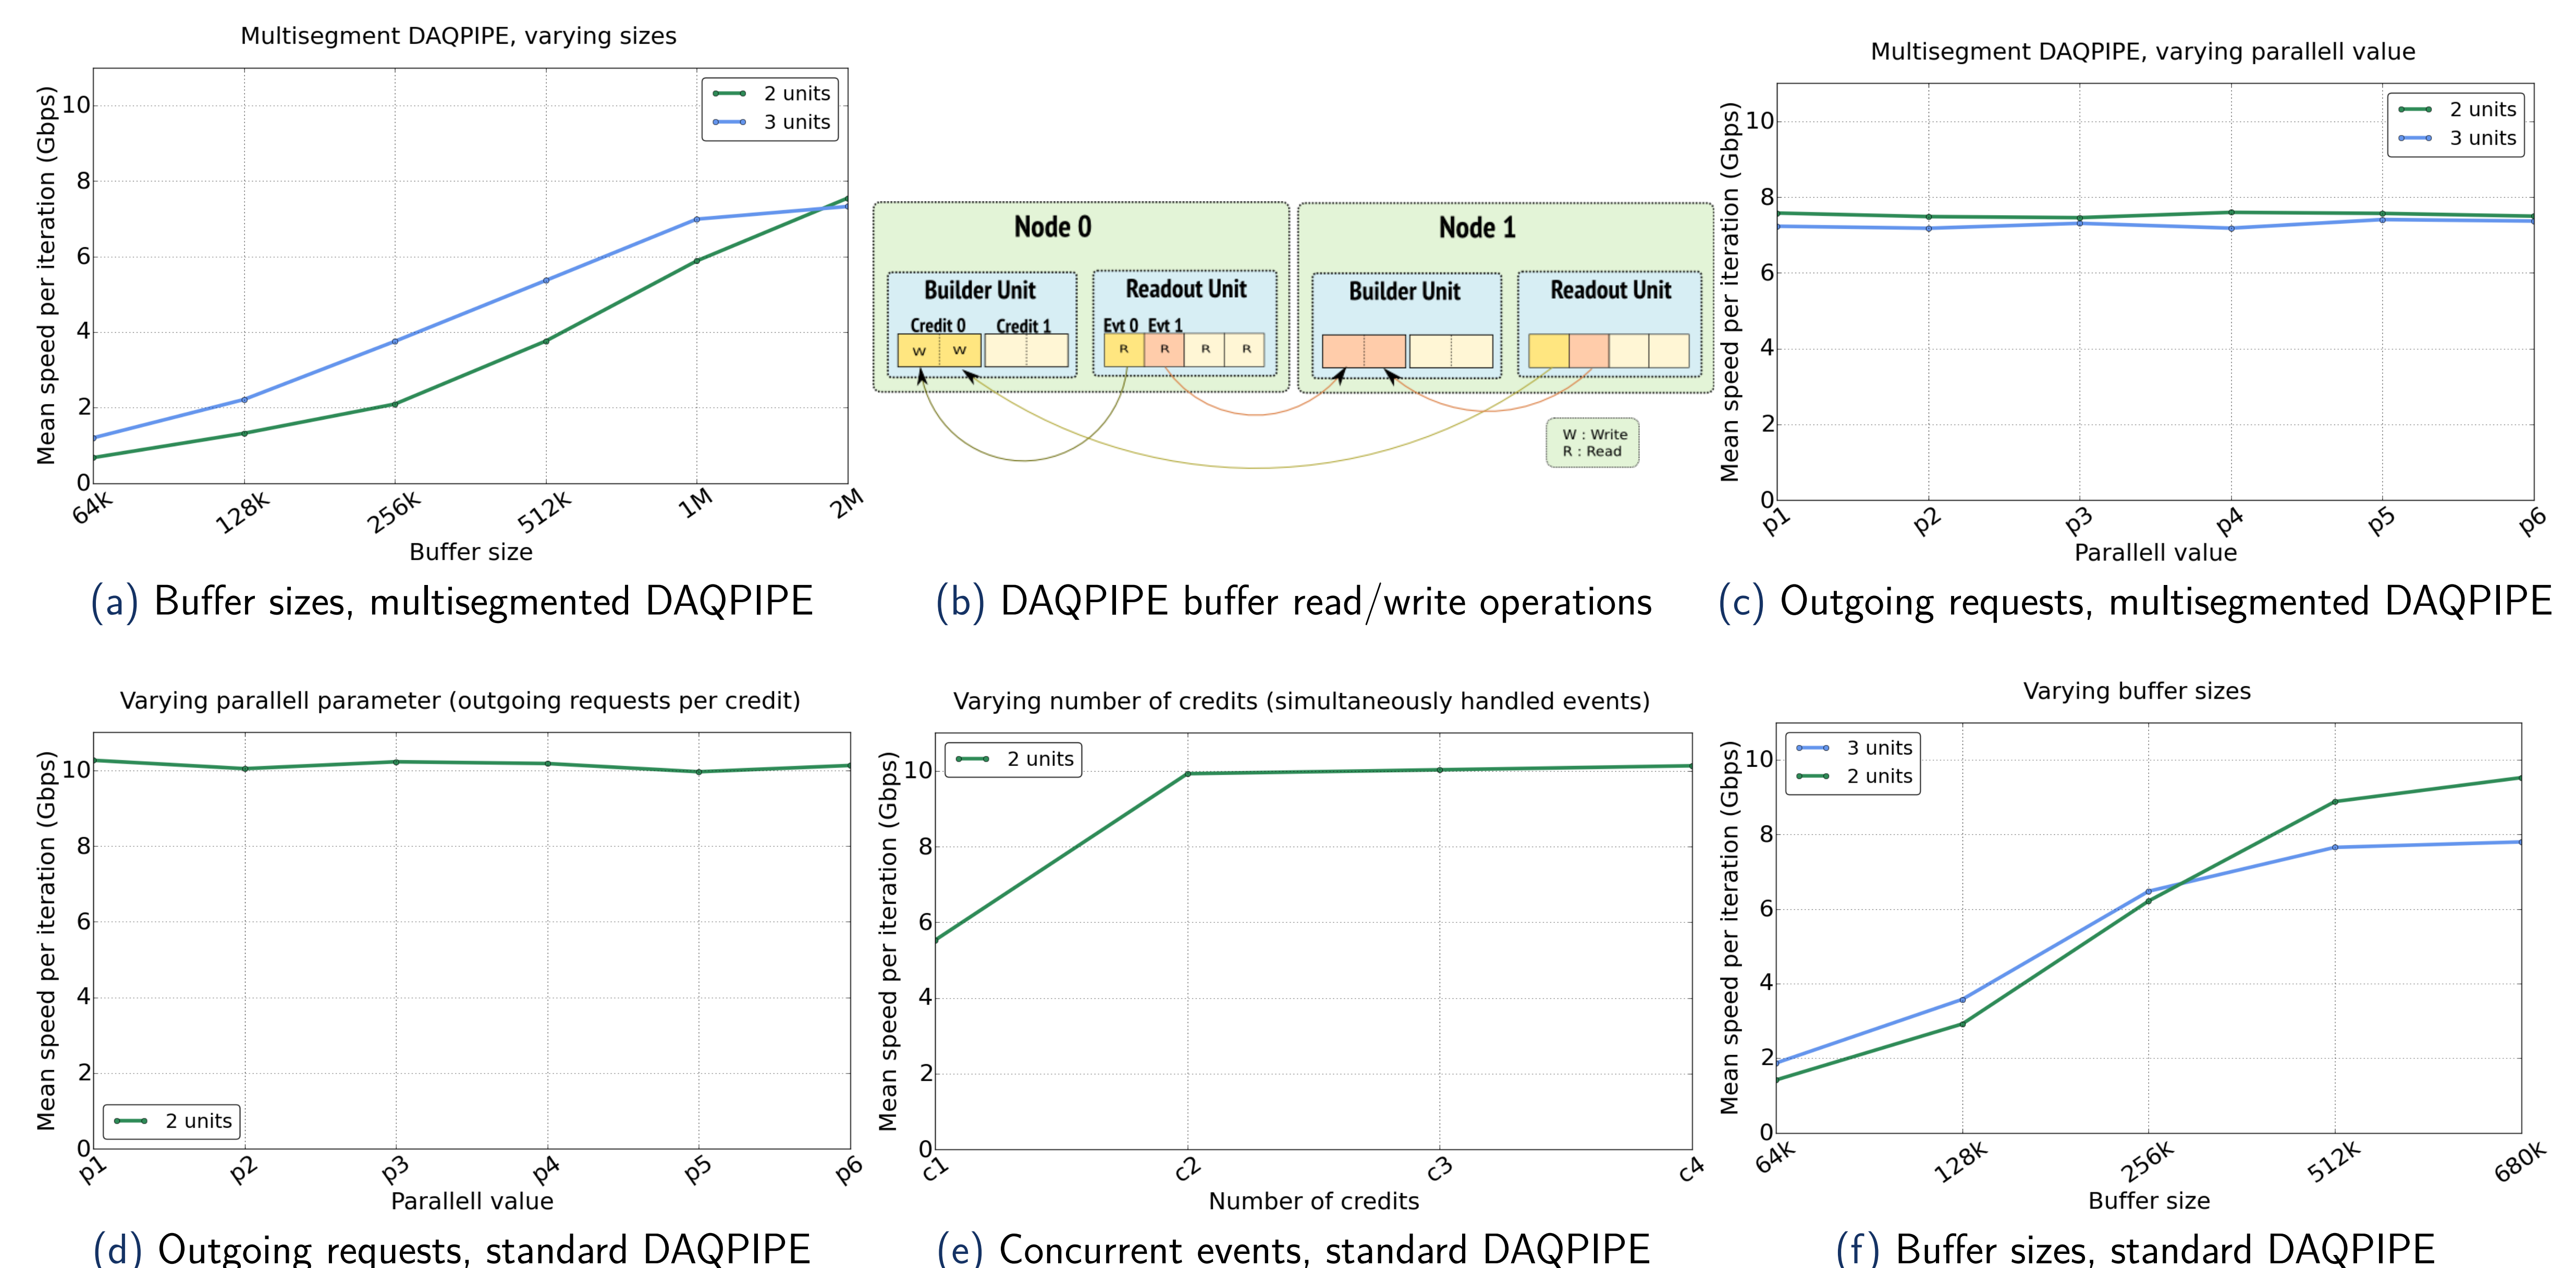
- ROOT is a data processing framework developed at CERN
- Normally operates over TCP/IP
- Ported to use the RapidIO protocol
- Benchmarking data transactions between two or more instances
- Channelized Messaging Implementation
- DMA Implementation



- Duplex communication does not hinder performance
- DMA faster than CM (expected)
- CM not suitable for orchestration
- ROOT not a networking benchmark
- TCP 1Gbps as point of reference
- Bandwidth sustained over time

DAQPIPE

- DAQPIPE is a benchmark application for network fabrics
- Emulates the future LHCb local area network event builder
- Protocol, topology and interconnect agnostic
- Already ported to several network technologies
- Channelized Messages for commands and orchestration
- DMA for memory storage



- Larger buffers give higher speeds
- Expected performance decrease when scaling
- Concurrency (events, requests) become relevant at scale
- Good match in terms of paradigms used
- Orchestration through CM affects overall speed
- DAQPIPE's internal orchestration more efficient