

Towards Energy-Efficient Storage Servers

Vinícius Machado

Natália Rampon

Francieli Boito

Edson Padoin

Julien Diaz

Amanda Braga

Jean Bez

Rodrigo Kassick

Jean-François Méhaut

Philippe Navaux

Federal University of Rio Grande do Sul (UFRGS), Brazil

April 6th, 2017



High Performance Computing

- Increasing performance requirements
- Escalating large scale parallel platforms
- This processing power requires a high power demand

Energy Concern

- HPC's next goal: reach **Exascale** computation
- A 20 MW **limit** was suggested for future systems by the USDOD
- Processors' energy efficiency must increase by two orders of magnitude
- Power demand is a primary concern

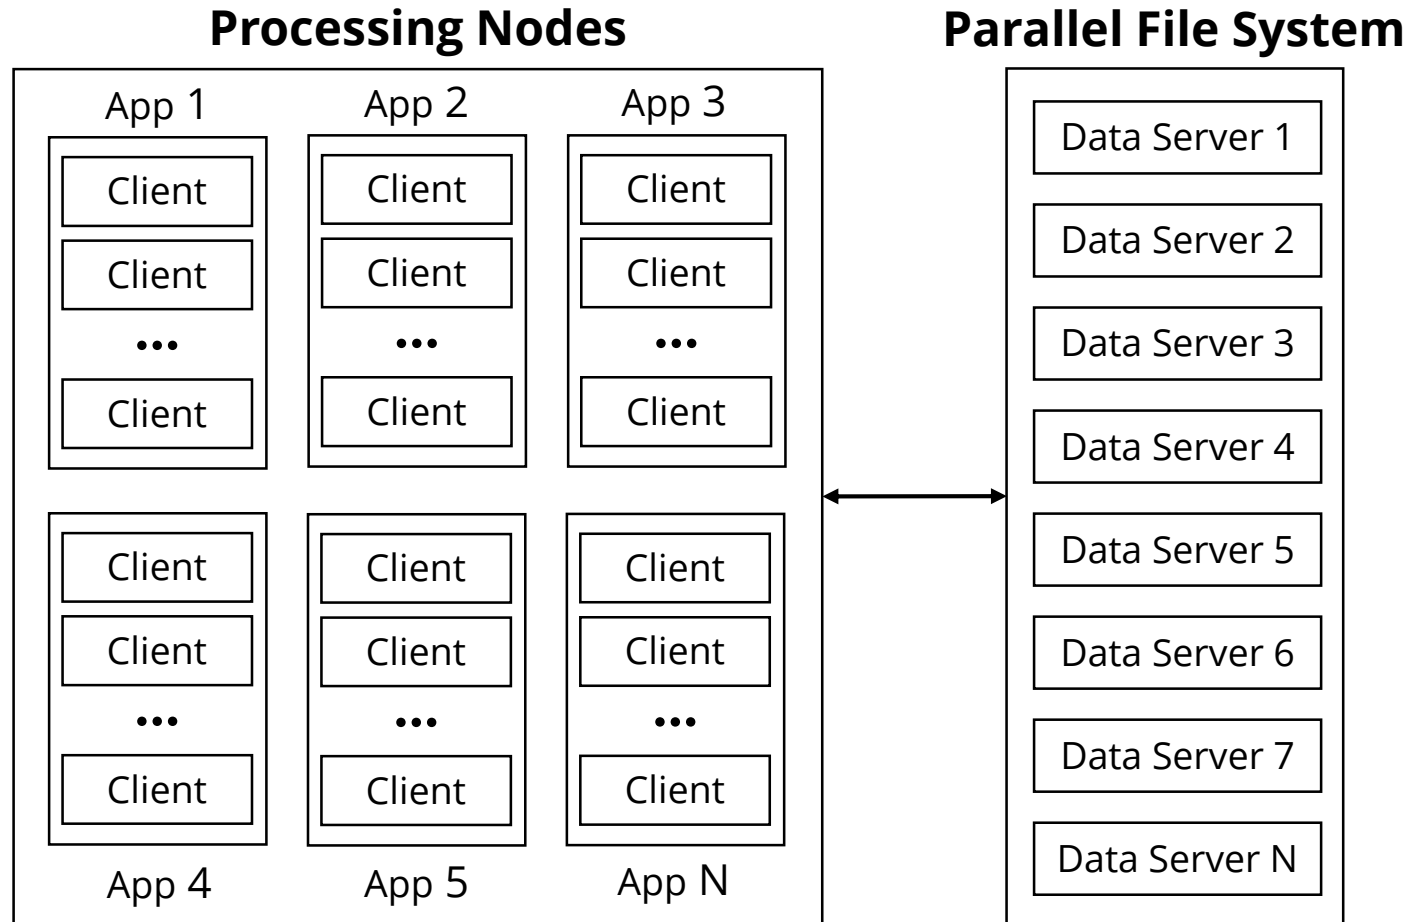
Alternatives

- Research efforts aimed at reducing systems' power demand
- Common strategy: Use **ARM** processors
- Improve energy efficiency at a performance **cost**

I/O

- Processing units demand most of the power
- Historical **gap** between processing and data access speeds
- I/O is a **bottleneck** for an increasing amount of applications
- Better **I/O** energy efficiency is also needed

Parallel File Systems



- Data servers **don't** fully utilize their processing power

Contribution

- Data servers could be **equipped** with **ARM processors**
- Similar research has been done in the past, **mostly on Data Centers** (Paypal, Microsoft)
- None of them evaluate the alternative or measure its performance
- Our goal is to **evaluate energy efficiency and performance** of a ARM-based PFS data server

Summary

I. Introduction

II. **Methodology**

III. Results

IV. Conclusion & Future Work

Towards Energy-Efficient Storage Servers

Tests' Environment

- Low-power machine: **Two Cubietrucks**
 - Dual-core ARM A20 Processor
- Regular machine: **"PC"**
 - Intel i5-4460 Processor
- Power Meter: P4460 Kill-a-Watt Ez
- Storage devices: **Two Samsung 840 SSDs**

Synthetic Experiments

- Cubietrucks use a customized Linux kernel
- We were unable to install typical parallel file systems
- An **emulator** using MPI was developed
 - Only data servers are emulated
 - Isolate effects of the network and of processing nodes
- Access patterns:
 - **Small** (64 KB) and **large** (4 MB) requests
 - Read and write operations
 - **Contiguous** and **1D-strided** spatialities

Hou10ni

- Wave propagation simulator
- Tests were executed on the **Edel** cluster from **Grid'5000**
- Tests' traces were provided to the emulator as a workload to compare with our synthetic tests' results

Statistics

- Each test was repeated 5 times in both architectures
- Results do not follow a normal distribution
- **Medians** were used instead of means
- **No error bars** are presented
- Dunn test was used to compare medians

Summary

I. Introduction


II. Methodology

III. **Results**


IV. Conclusion & Future Work

Towards Energy-Efficient Storage Servers

Power Demand




Power Demand



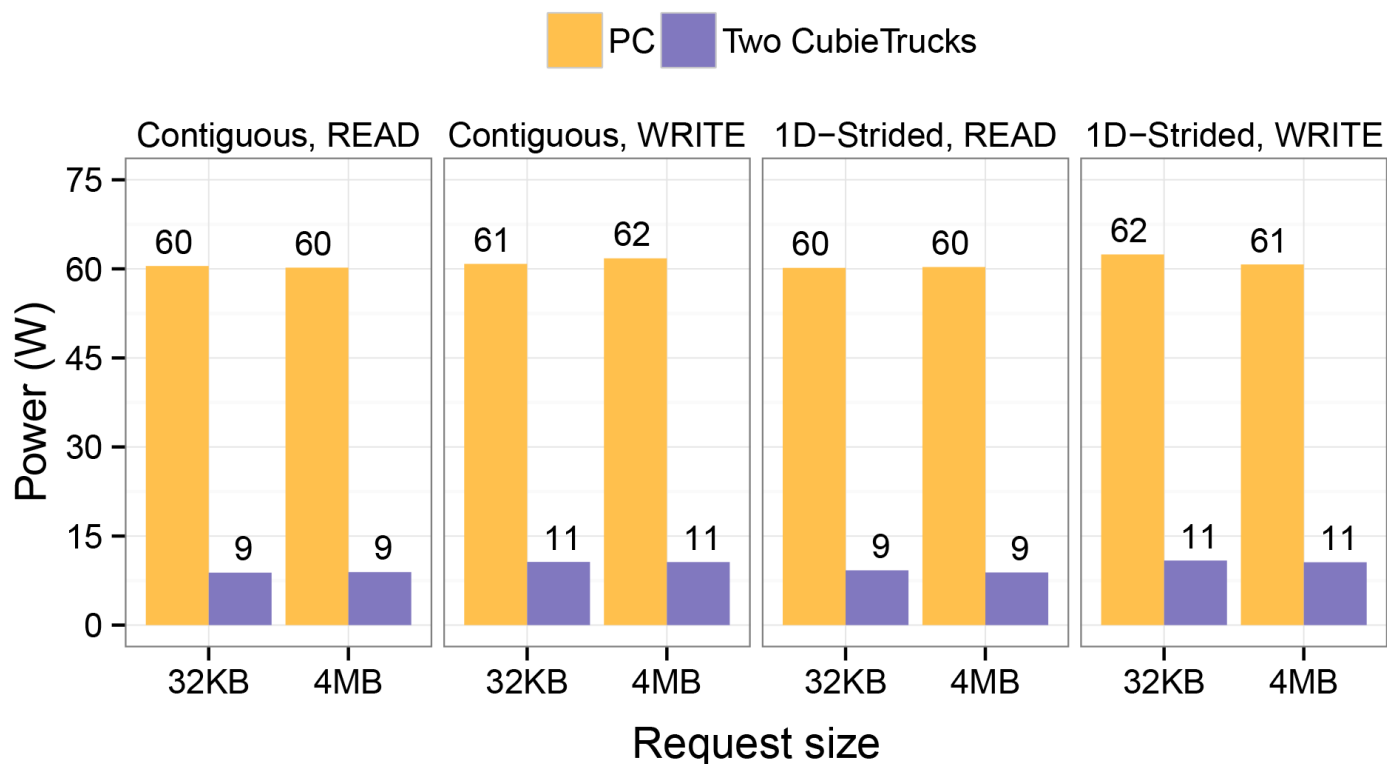
- Request size has no significant impact

Power Demand




- Spatiality has no significant impact

Power Demand




- Write is more power demanding than read in all cases

Power Demand




- Write is 7% more power demanding in the PC

Power Demand




- Write is 23% more power demanding in the Cubietrucks

Power Demand




- Cubietrucks are significantly less power demanding

Power Demand




- Read: 85% lower power demand in the Cubietrucks

Power Demand




- Write: 82% lower power demand in the Cubietrucks

Performance




Performance




- Request size has no significant impact

Performance




- Spatiality has no significant impact

Performance




- Read bandwidth is 240% **higher** in the Cubietrucks

Performance




- Write bandwidth is 37% lower in the Cubietrucks

Energy Efficiency



Energy Efficiency



- Cubietrucks offer higher energy efficiency
- **1437%** for read tests
- **287%** for write tests

Discussion: Read

- A single Cubietruck can achieve similar read bandwidth to a common processor with higher energy efficiency
- Replacing a common server with two Cubietrucks
 - Doubles the bandwidth
 - Decreases the power demand by 85%


Discussion: Write

- Replacing a common server with **two** Cubietrucks
 - Increases energy efficiency by 287%
 - Decreases bandwidth by 37%
- By extrapolating our results, we found that replacing a common server with **four** Cubietrucks
 - Decreases power demand by 64%
 - Increases bandwidth by 25%

Discussion: Cost

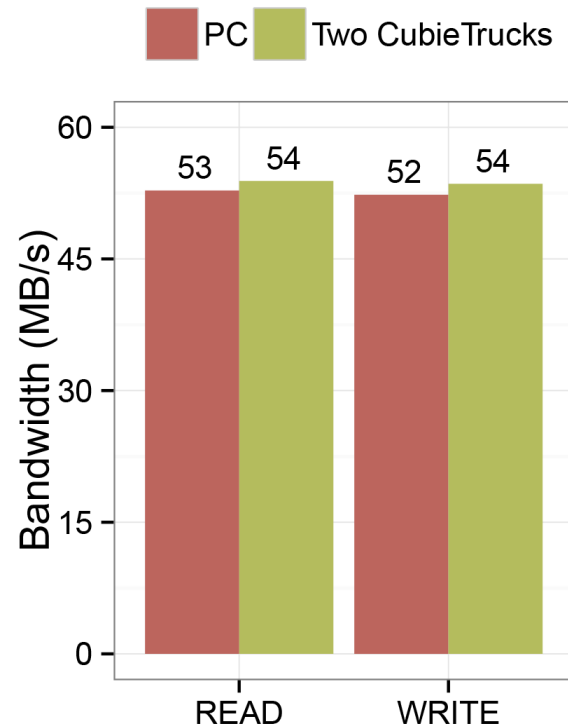
- All costs as of September, 2016
- Each Cubietruck costs 100 USD
- A similar PC costs 500 USD
- SSDs cost 330 USD each
- The replacement is **financially viable** and would have future savings in the electricity bills

Hou10ni: Power Demand



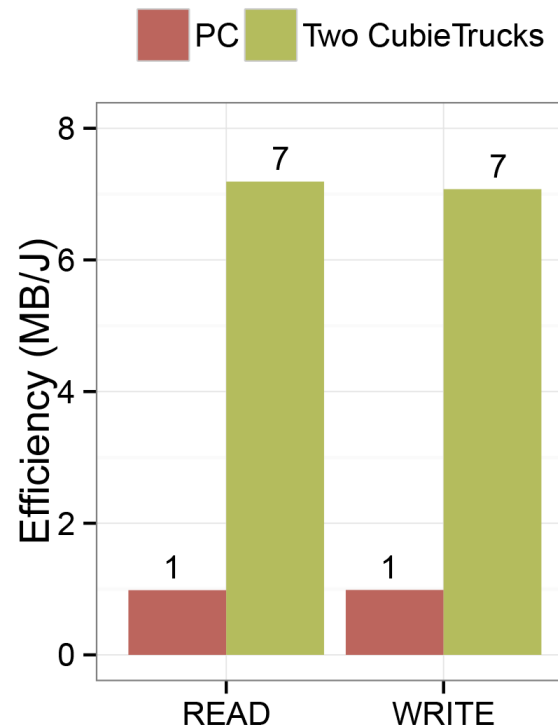
- Cubietrucks demanded 86% less power than the PC

Hou10ni: Performance



- Both options achieved the same performance
- Result due to a less intensive workload

Hou10ni: Energy Efficiency



- Cubietrucks were 7x more energy efficient than the PC

Summary

I. Introduction

II. Methodology

III. Results

IV. **Conclusion** & Future Work

Towards Energy-Efficient Storage Servers

Conclusion

- We presented an evaluation of the energy efficiency and performance of an ARM-based PFS data server
- It is viable to replace one server with two Cubietrucks for read workloads
- More ARM servers could also provide good trade-offs for write workloads
- The alternatives are also financially viable

Future Work

- Include the network and the processing nodes on the emulator
- Evaluate more access patterns
- Explore other low-power architectures

Acknowledgments

This research received funding from the European EU H2020 Programme and from MCTI/RNP-Brasil in the context of the HPC4E Project, number 689772

It was also partially funded by CNPq



Thank You!

Vinícius Machado

<http://inf.ufrgs.br/~vrmachado>

vrmachado@inf.ufrgs.br

