Energy Consumption and Performance Analysis Between SSD and HDD

Pablo J. Pavan, Vinícius R. Machado, Jean L. Bez, Edson L. Padoin, Francieli Z. Boito, Philippe O. A. Navaux, Jean-François Méhaut

WSPPD - 2017

Summary

- Introduction
- Goal
- Methodology
- Results Analysis
- Conclusion and Future Work

Introduction

- **Energy consumption** is a major limitation in the construction Exascale Systems
  - DARPA limits the consumption to **20 MWatt**
- The processors represent a **significant percentage** o the power demand of HPC systems
- **File systems** also have an impact on power demand

Component	*Avg Power Consumption*
CPU [21]	80 W
Memory [20][23]	36 W
Motherboard [16]	25 W
Disk subsystem [27]	12 W
Fans [16]	10 W
Network Interface [44]	2 W

PALMIERI, Francesco et al. Energy-oriented denial of service attacks: an emerging menace for large cloud infrastructures. **The Journal of Supercomputing**, v. 71, n. 5, p. 1620-1641, 2015.

Introduction

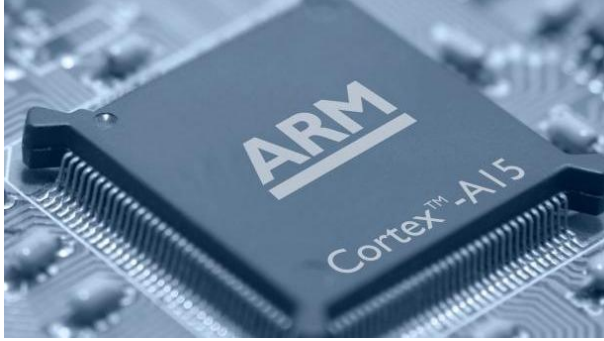

- **Alternatives** that respect the given **power limit**
  - Use Advanced RISC Machine (**ARM**) processors
    - ARM processors focus on low energy consumption but yet present **good energy efficiency** results
  - Replace traditional HDD by Solid State Drive (**SSD**)
    - No Moving Parts
    - Speed
    - **Energy Efficiency**

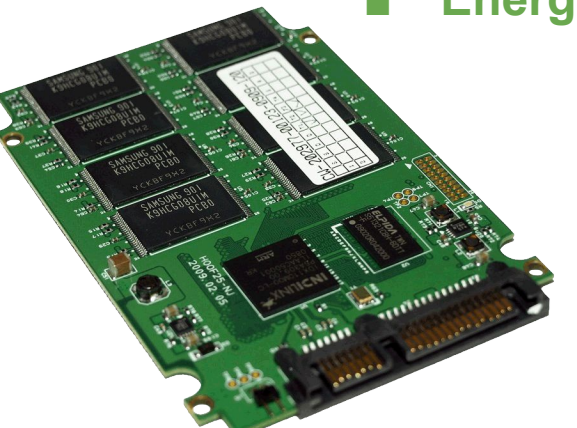

Goal

To analyze the viability of **replacing** conventional servers by **low-power alternatives** to overcome the **need to build exascale systems**

Methodology - Equipments

ARM

VS

intel

	CPU	MPSoC
Processor	Intel Core-I7	ARM Cortex A7
Processor model	4790	AllWinnerTech SoC A20
Technique of Manufacture (nm)	22	40
Clock frequency	3.6GHz	960MHz
Cores/Processor (#)	4 (with Hyper-Threading	2
Memory (GB)	16 DDR3	2 LP DDR3

Methodology - Storage Devices

Type	Manufacturer	Capacity(GB)	RPM
HDD1	Seagate	1000	5400
HDD2	Seagate	60	7200
SSD1	Samsung	240	-
SSD2	Kingston	120	-

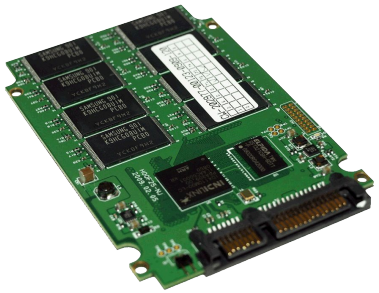

Methodology - Benchmark

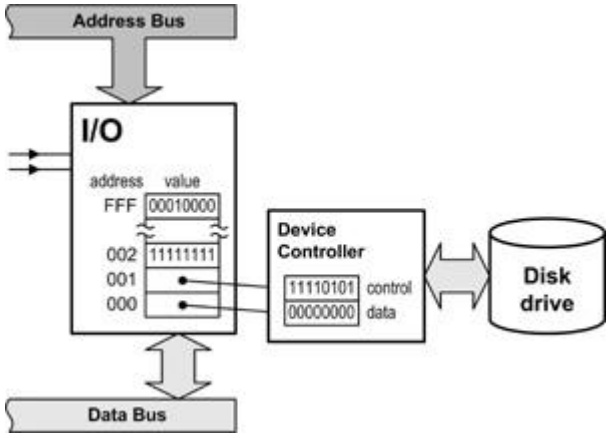

- **FIO**
  - The experiments
    - **with** and **without** the usage of the **buffer cache**
  - **Four access patterns**
    - sequential write, random write, sequential read, random read
  - **Two request** sizes
    - 32 KB and 4 MB
  - Data size **20 GB**
  - Time limit of **60 seconds**
- Total of 96 experiments, each one of them was repeated **10 times**
- A minimum **20-seconds** delay is guaranteed **between tests**

Methodology - Measure

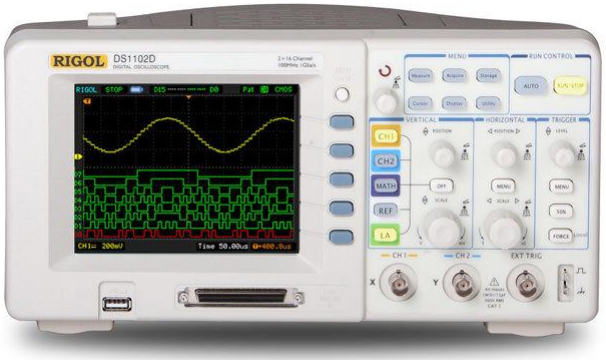

- To **measure the power demand**, we employed an Agilent oscilloscope model DSO6014A
  - A power tip model 1146A was used to **measure the current** for the **entire equipment**
  - **Current** for the **storage devices** was measured from the Hall effect, with an Allegro solution model ACS712T connected to the oscilloscope
- Instantaneous **voltage and current** measurements
  - **500 ms**
- The oscilloscope was connected via USB to a computer, where the BenchVue software logs captured data

Results Analysis - Performance VS Power Demand

- The results have shown that all devices **suffer in performance** when used in the **MPSoC**
- Access patterns **impact on performance**, but **not power demand**
- **MPSoC** has a **lower power** demand than the **PC**

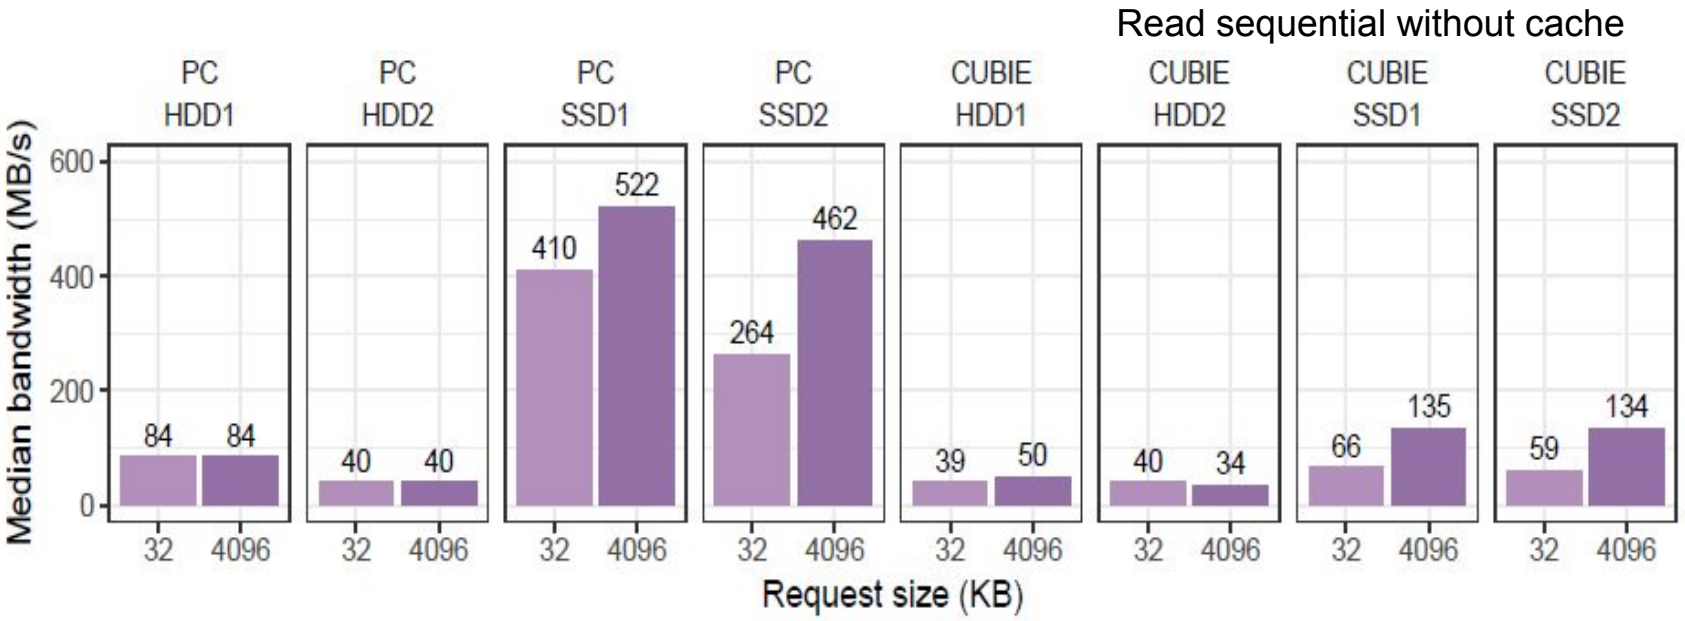

Results Analysis - Performance VS Power Demand

- In the **PC**
  - **write** performance was up to **1062%** higher
  - **read** performance was up to **522%** higher
- The **power demand** show **SSDs** do **not** demand as much power in the **MPSoC** than in the **PC**

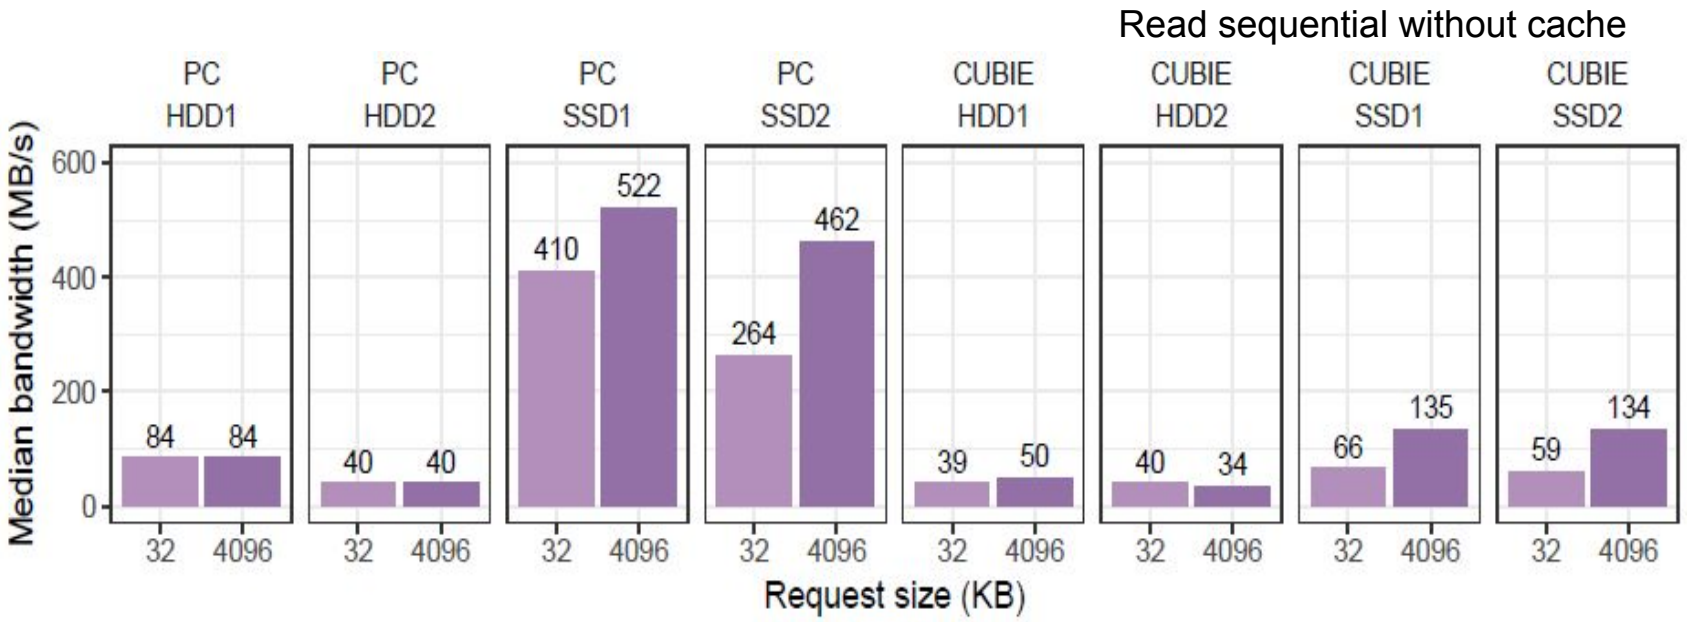

Results Analysis - Energy efficiency

- Using **SSDs** leads to up to **6675% higher energy efficiency** than using **HDDs**
- When using **SSDs** **energy efficiency** is **higher** in the **PC** than in the **MPSoC**
  - up to **196%** for **write** workloads and **564%** for **read** workloads
- Using **HDD2** in the **MPSoC** results in **higher energy efficiency**
  - up to **166%**

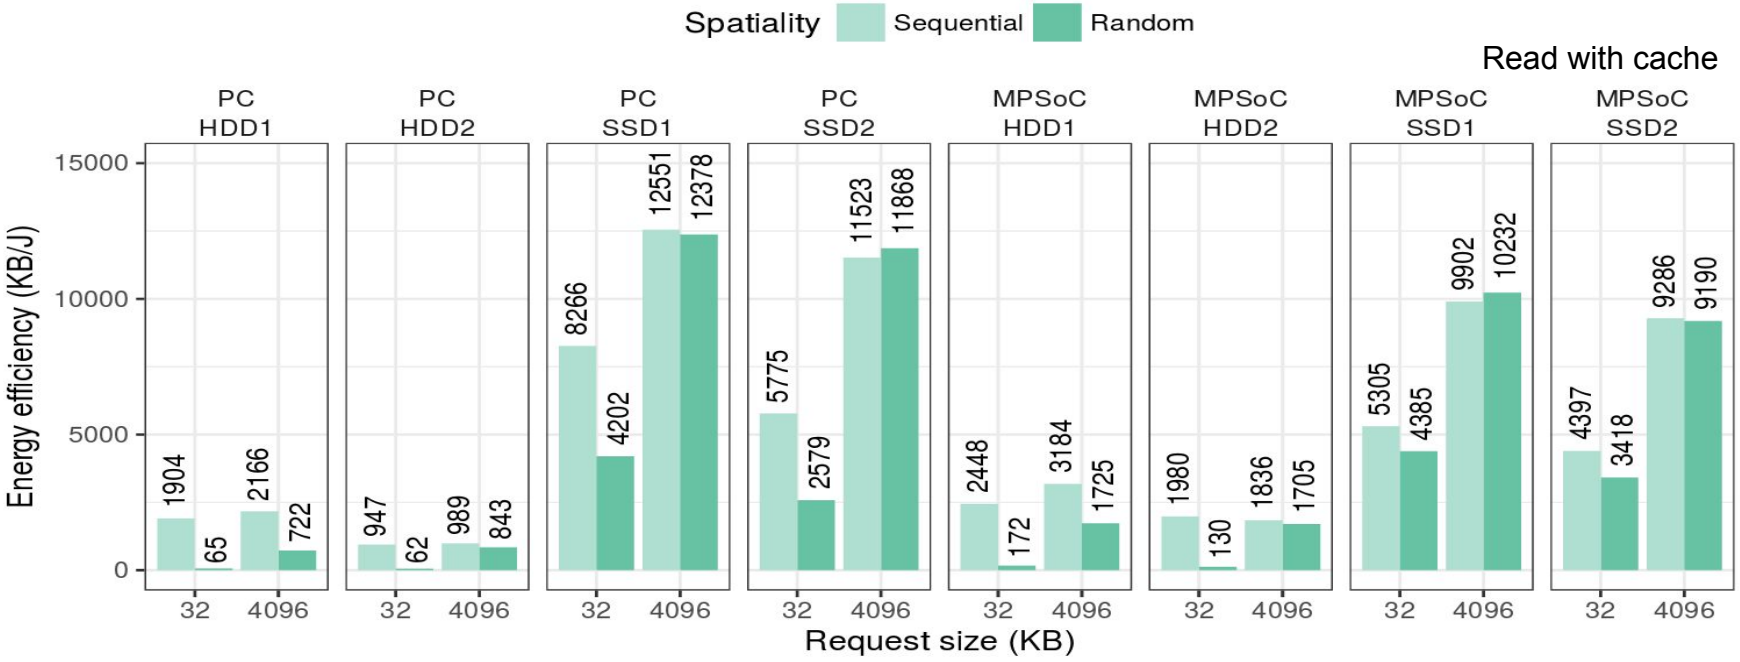

Conclusion

- **Replacing** the traditional server by multiple low-power ones only results in higher **energy efficiency** if the **PC uses HDDs** for storage, and the **MPSoC uses SSDs** to **write workloads**
  - **8% lower power demand** by replacing the PC by **3** MPSoC
    - **Without harm** to **sequential write** bandwidth
    - **Increasing** **random write** bandwidth in up to **40%**

Conclusion

- **Replacing** the traditional server by multiple low-power ones only results in higher **energy efficiency** if the **PC uses HDDs** for storage, and the **MPSoC uses SSDs** to **write workloads**
  - **8% lower power demand** by replacing the PC by **3** MPSoC
    - **Without harm** to **sequential write** bandwidth
    - **Increasing** **random write** bandwidth in up to **40%**

- **Read workloads** this **replacement** could be by **1.4** MPSoC servers
  - **To keep the same** **sequential read** bandwidth with **small requests**
  - **Increasing** **sequential read** bandwidth with large requests in up to **61%**
  - **Increasing** **random read** bandwidth in up to **294%**

Conclusion

- The **replacement** of traditional servers by low-power ones also makes sense **if both use HDDs** for **read workload**
  - **2.2** MPSoC with **HDD1**, resulting in **20% lower power demand**
  - **1.2** MPSoC with **HDD2**, demanding **42% less power**
  - These replacements would **keep** the same **sequential read** bandwidth, and **increase** the **random read** bandwidth in up to **120%**
- **Nonetheless**, **write workloads** would observe **lower performance**

Future Work

- Future work include expanding the investigation to other processors and storage devices.
- Furthermore, we would like to consider the traditional and low-power alternatives as storage servers, receiving requests through the network and processing them.

Energy Consumption and Performance Analysis Between SSD and HDD

Pablo J. Pavan, Vinícius R. Machado, Jean L. Bez, Edson L. Padoin, Francieli Z. Boito, Philippe O. A. Navaux, Jean-François Méhaut

Thanks!

WSPPD - 2017