OPEN POSSIBILITIES.

OMI, The Path to High Bandwidth, High Capacity Near Memory

OCP GLOBAL SUMMIT
NOVEMBER 9-10, 2021

OMI, The Path to High Bandwidth, High Capacity Near Memory

Tom Coughlin, Coughlin Associates, https://tomcoughlin.com
Jim Handy, Objective Analysis, https://objective-analysis.com

OCP GLOBAL SUMMIT
NOVEMBER 9-10, 2021

Preface: Near Memory vs. Far Memory

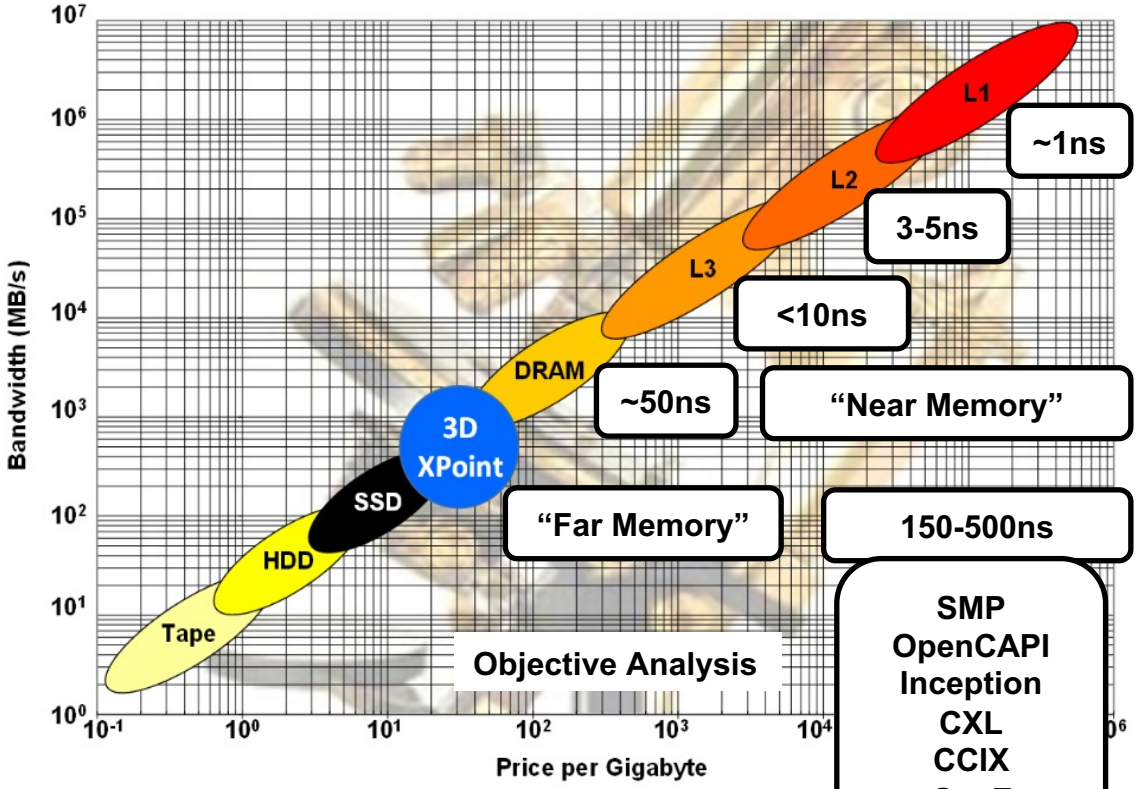

SERVER

HIGH PERFORMANCE COMPUTING

OPEN POSSIBILITIES.

Preface: Near Memory vs. Far Memory

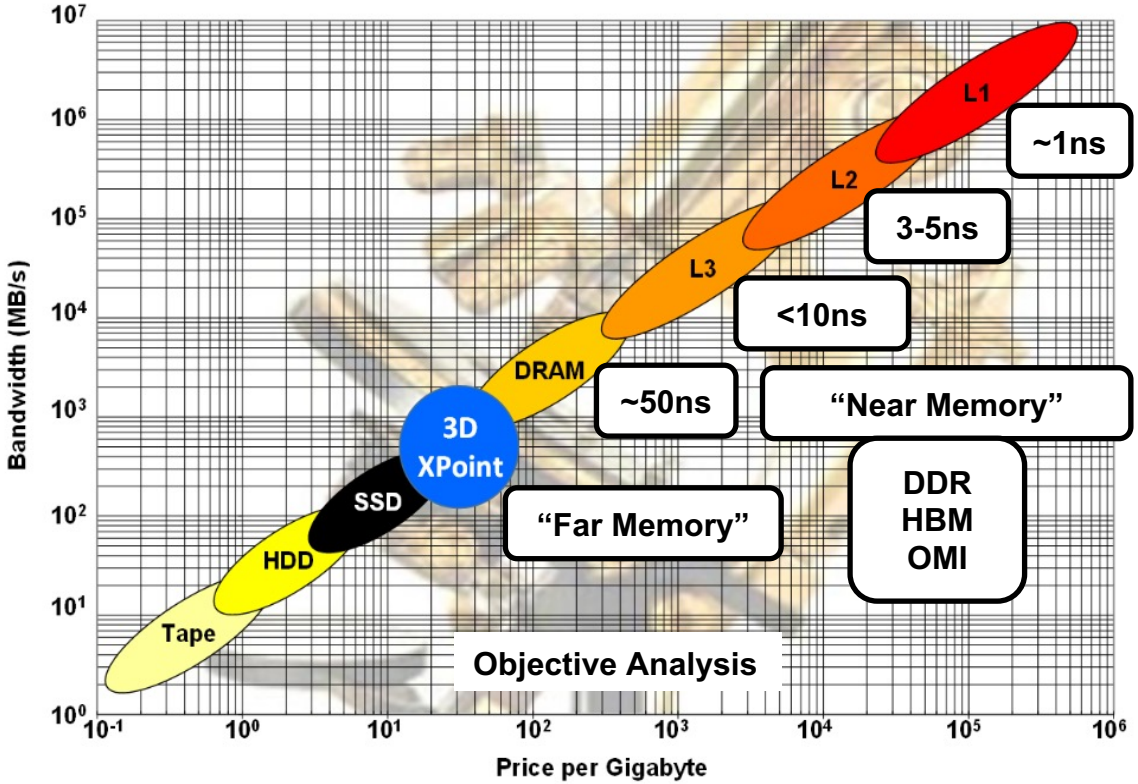

SERVER

HIGH PERFORMANCE COMPUTING

OPEN POSSIBILITIES.

OCP GLOBAL SUMMIT
NOVEMBER 9-10, 2021

Outline

- More Cores Need More Near Memory Bandwidth
- Today's Problems Need Larger and Faster Near Memories
- Three Alternatives DDR, HBM, & OMI
- What Is a DDIMM?
- Supporting a Mix of Near Memories
- Conclusions

SERVER

HIGH PERFORMANCE COMPUTING

Processor Trends Drive Near Memory Demands

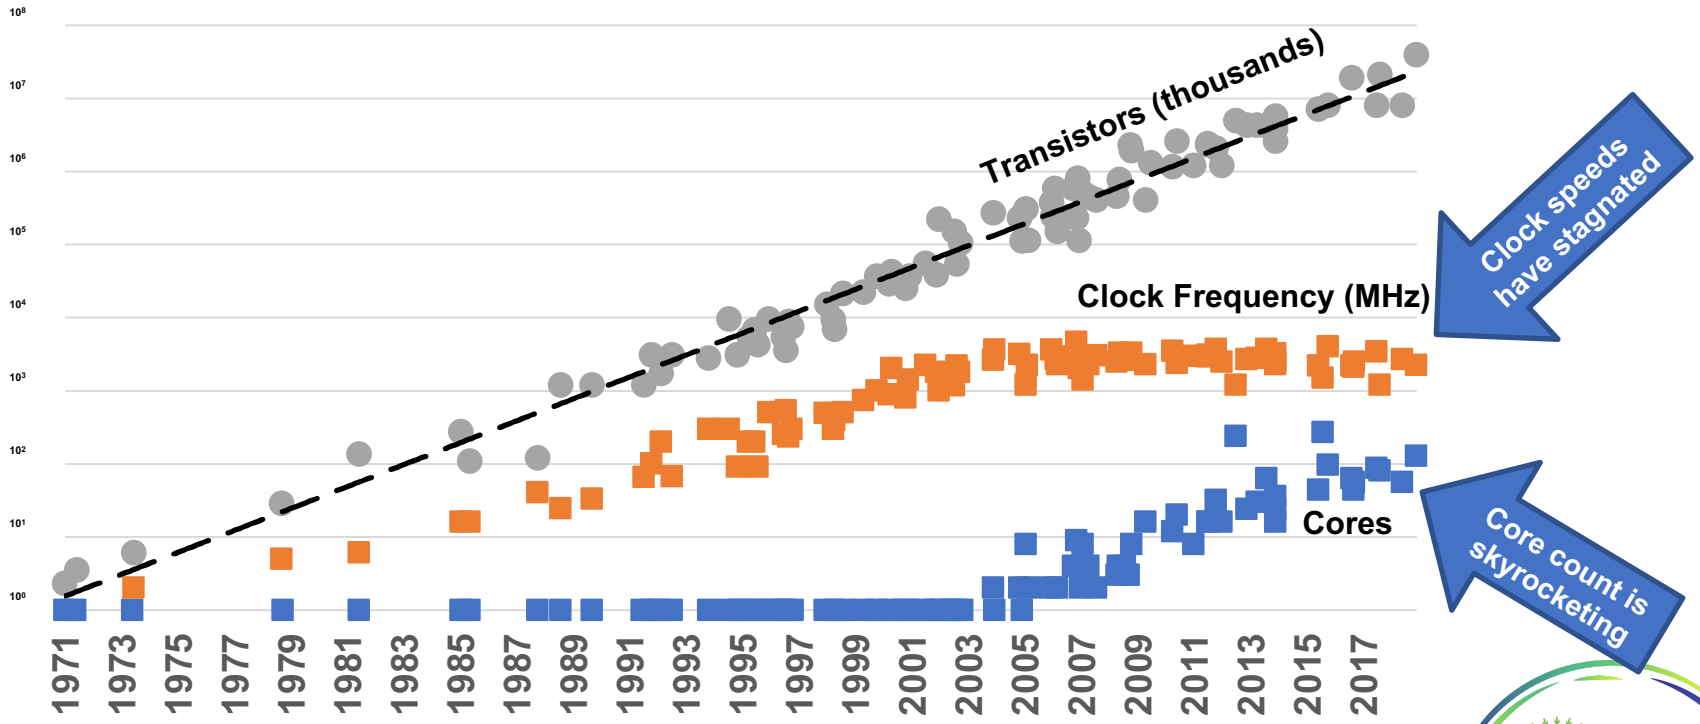

Data Source: Karl Rupp - https://www.karlrupp.net/2015/06/40-years-of-microprocessor-trend-data/

OPEN POSSIBILITIES.

SERVER

HIGH PERFORMANCE COMPUTING

OCP GLOBAL SUMMIT

NOVEMBER 9-10, 2021

Processor Trends Drive Near Memory Demands

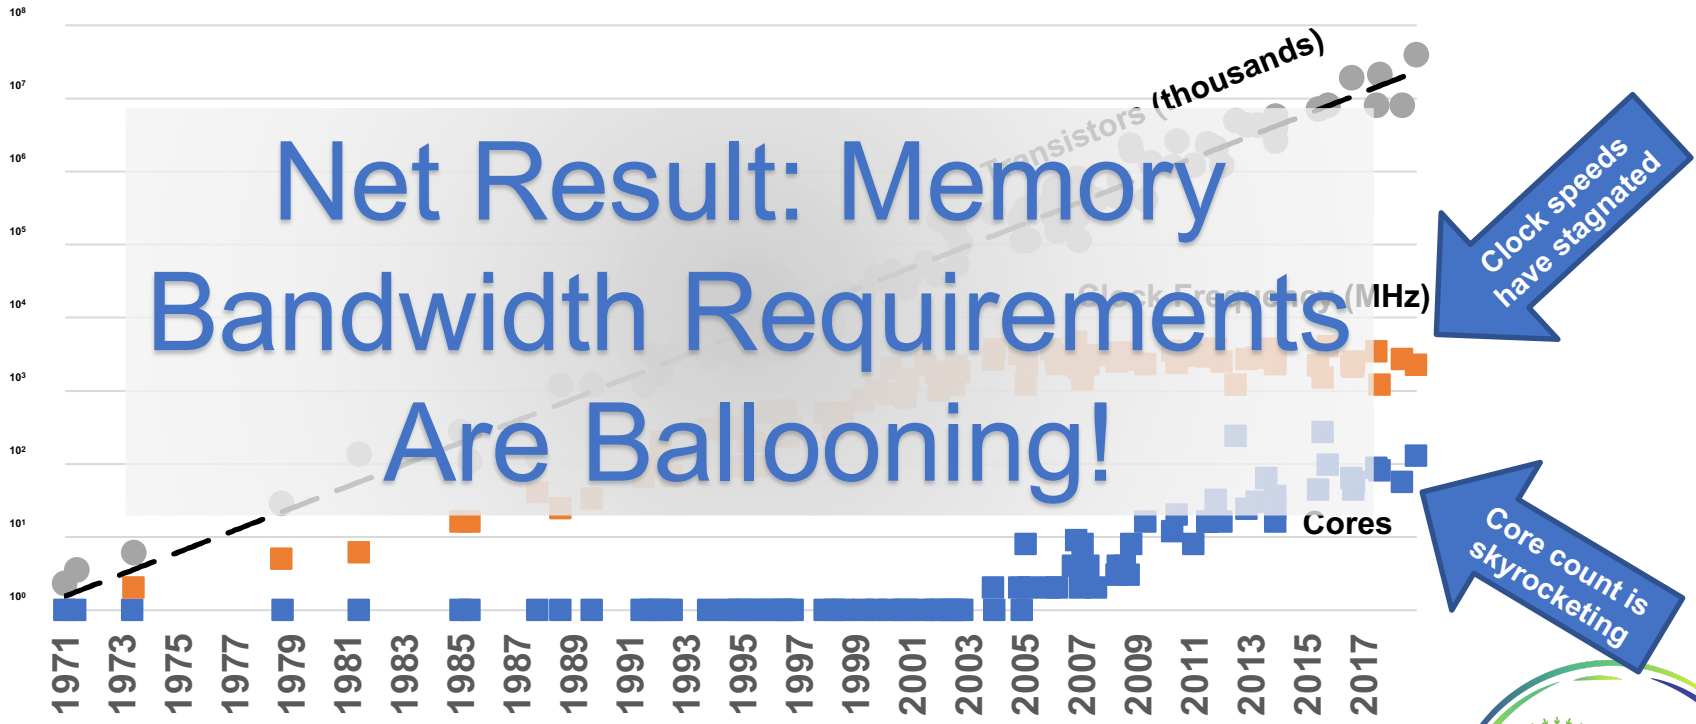

Data Source: Karl Rupp - https://www.karlrupp.net/2015/06/40-years-of-microprocessor-trend-data/

OPEN POSSIBILITIES.

SERVER

HIGH PERFORMANCE COMPUTING

OCP GLOBAL SUMMIT

NOVEMBER 9-10, 2021

Increasing Near Memory Bandwidth

- Improve the memory interface
  - DDR, DDR2, DDR3, DDR4, DDR5
- Increase the bus clock speed
  - 1600, 1866, 2133, 2400, 2666, 2933, 3200
  - But latency doesn't change as clock speed increases
  - Higher clock speeds reduce DIMMs per channel
- Add Memory Channels

SERVER

HIGH PERFORMANCE COMPUTING

More Memory Channels Means More Processor Die Cost & Power

- Each LRDIMM channel requires 152 processor pins
  - 72 data bits + 18 address bits + 36 strobes + 26 command/other bits
  - 6 Channels = 912 pins
  - 8 Channels = 1,216 pins
- The most expensive real estate in the system
- Fast I/O is a power hog
  - Capacitance is the culprit!

SERVER

HIGH PERFORMANCE COMPUTING

Three Near Memory Options

- DDR4
  - Mainstream & inexpensive
- HBM2E
  - Costly, but fast
    - 1,000-bit bus!
- OMI
  - Introduces SerDes into data path
  - Supports high bandwidth and large memory capacities

SERVER

HIGH PERFORMANCE COMPUTING

OMI & IBM's POWER10 CPU

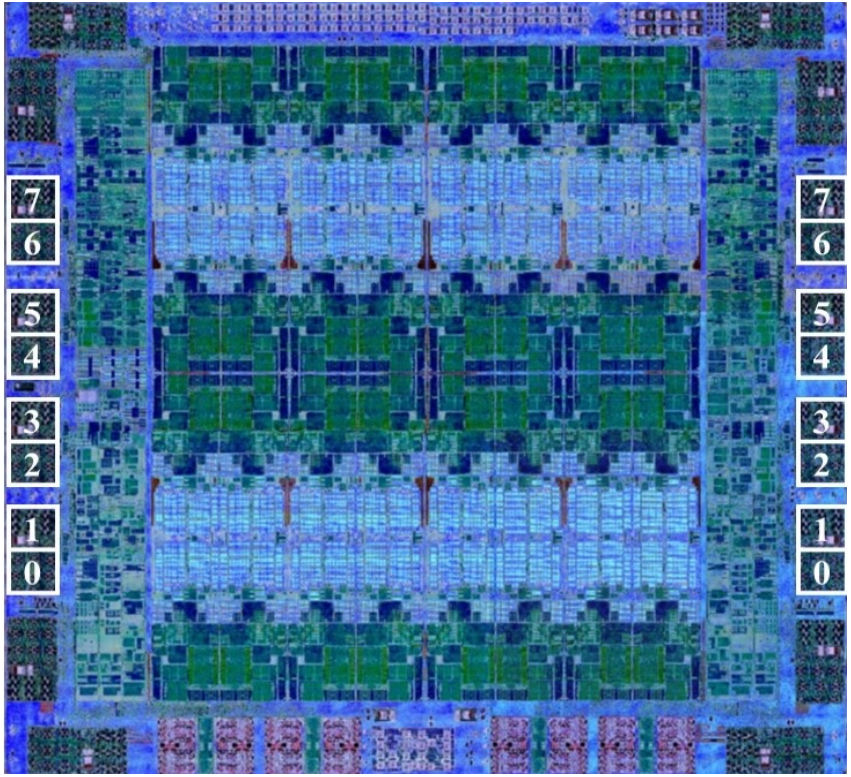

SERVER

HIGH PERFORMANCE COMPUTING

OPEN POSSIBILITIES.

POWER10 Stats

- Sixteen OMI channels
  - 256 GB capacity per channel
  - 1TB/s total memory bandwidth (read + write)
  - 64GB/s per OMI channel
  - 2.2mm² die area per channel
  - 29.6GB/s/mm²
    - Nearly matches HBM

SERVER

HIGH PERFORMANCE COMPUTING

OMI DDIMM

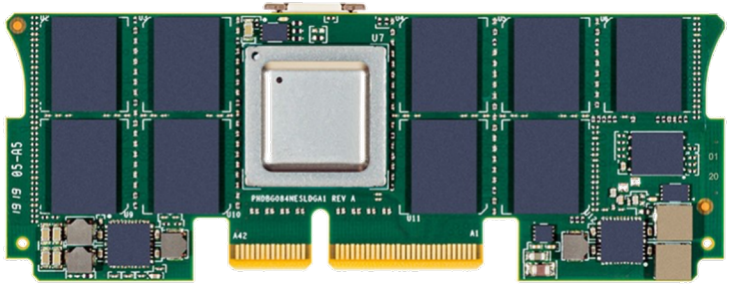

SERVER

- Serial differential signaling
  - ¼ the pins of DDR
  - HBM-like bandwidth
- Moves DRAM signaling away from main board, onto DDIMM
  - Reduces processor's I/O power
- Low latency penalty <4ns
- Requires a controller
  - Replaces LRDIMM buffers
  - Processor no longer tied to one interface (DDR4) or one memory type (DRAM)
  - 30mm² die area
  - Can do on-DIMM processing (ECC, encryption…)

HIGH PERFORMANCE COMPUTING

OPEN POSSIBILITIES.

OCP GLOBAL SUMMIT

NOVEMBER 9-10, 2021

Other Possibilities

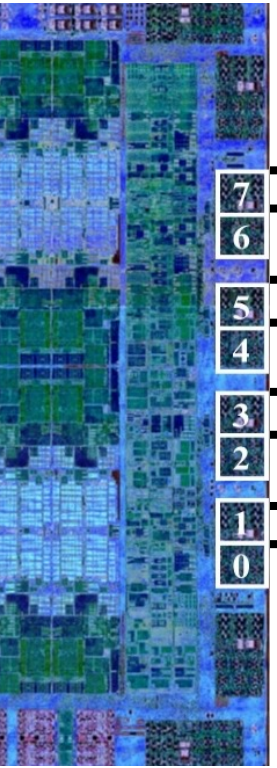

DDR4 | DDR5 | LPDDR | XPoint

MRAM | ReRAM | FRAM | NRAM

SERVER

HIGH PERFORMANCE COMPUTING

OPEN POSSIBILITIES.

Near HBM Bandwidth at >DDR Capacities

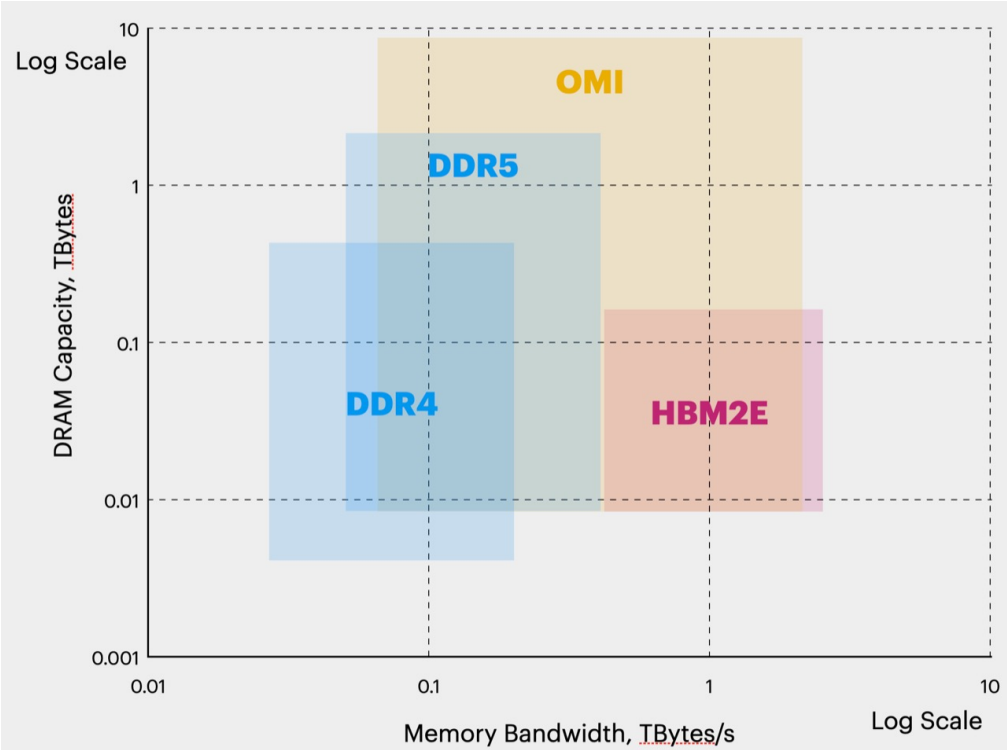

SERVER

HIGH PERFORMANCE COMPUTING

OPEN POSSIBILITIES.

OCP GLOBAL SUMMIT
NOVEMBER 9-10, 2021

Conclusions

Near Memory needs more bandwidth

Near Memory needs larger capacities

More Memory→DDR, Faster Speeds→HBM, Both→OMI

Near-HBM speeds at larger-than-DDR capacities

OMI can support a mix of Near Memory types

SERVER

HIGH PERFORMANCE COMPUTING

Call to Action

Visit the OMI Booth on the OCP Exhibit Floor

Consider adopting OMI as your near memory interface during your next processor design to get a low latency, high bandwidth and high capacity solution

SERVER

HIGH PERFORMANCE COMPUTING

OPEN POSSIBILITIES.

References

CXL Consortium, Introduction to Compute Express Link (CXL), 2019 Hot Interconnects.

Patrick Knebel, et. Al, Gen-Z Chipset for Exascale Fabrics, 2020 Hot Interconnects.

Mazyar Razzaz, DDR Basics, Register Configurations & Pitfalls, Freescale Technology Forum, July 2009.

Low Power Double Data Rate 4 (LPDDR4), JESD209-4D, JEDEC, 6/2021.

High Bandwidth Memory (HBM) DRAM, JESD235D, JEDEC, 3/2021.

Karl Rupp - https://www.karlrupp.net/2015/06/40-years-of-microprocessor-trend-data/
https://github.com/karlrupp/microprocessor-trend-data/blob/master/LICENSE.txt

W. Starke and B. Thompto, IBM's POWER10 Processor, 2020 Hot Interconnects.

T. Coughlin and J. Handy, The Future of Low Latency Memory, White Paper, https://Objective-Analysis.com/reports/#Emerging, July 2021.

SAP Standard Application Benchmarks, https://www.sap.com/dmc/exp/2018-benchmark-directory/#/sd

SERVER

HIGH PERFORMANCE COMPUTING

OPEN POSSIBILITIES.

Thank you!