Funded by

Deutsche Forschungsgemeinschaft

German Research Foundation

DRESDEN concept

André Berthold,
Constantin Fürst, Antonia Obersteiner, Lennart Schmidt, Dirk Habich, Wolfgang Lehner, Horst Schirmeier

Dresden University of Technology, Faculty of Computer Science, Institute of Systems Architecture

Demystifying Intel Data Streaming Accelerator for In-Memory Data Processing

Austin, Texas, DIMES 2nd Workshop on Disruptive Memory Systems // November 3, 2024

Motivation

NUMA-Node 0

NUMA-Node 1

Funded by DFG
DRESDEN concept

Motivation

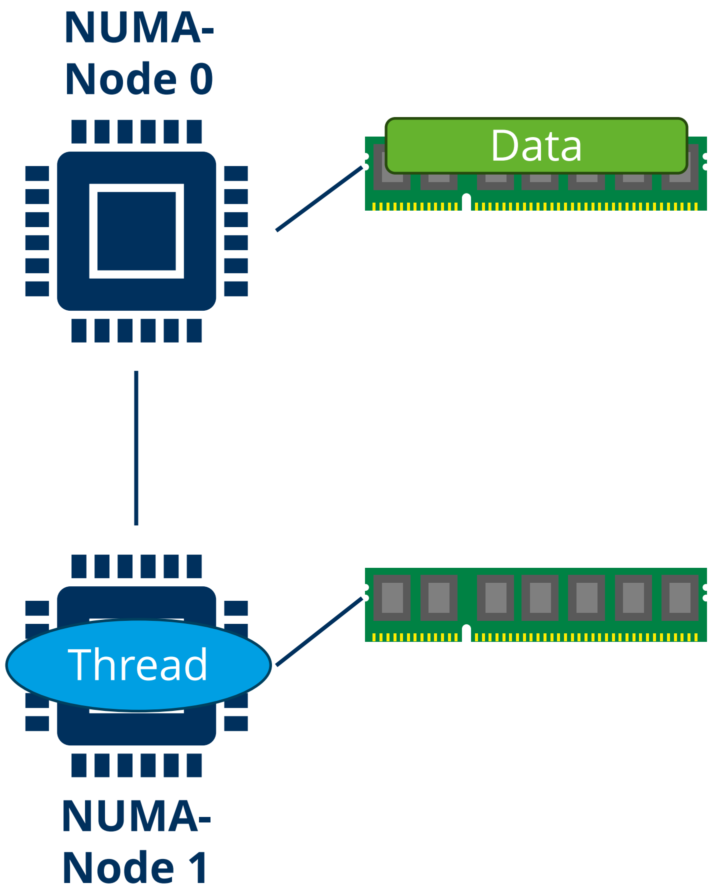

Funded by
DFG
DRESDEN concept

Motivation

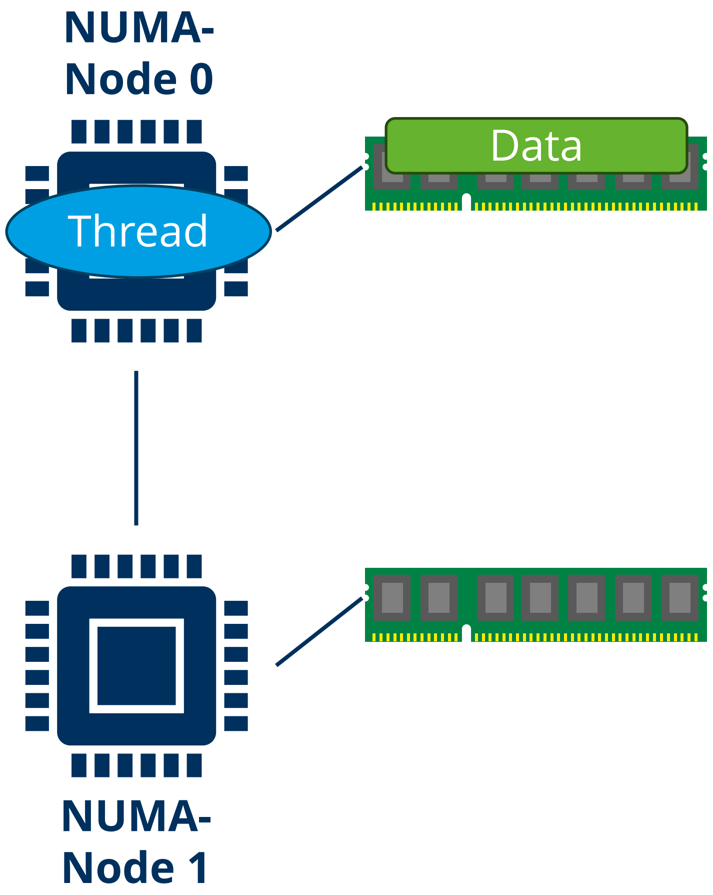

Funded by
DFG
DRESDEN concept

Motivation

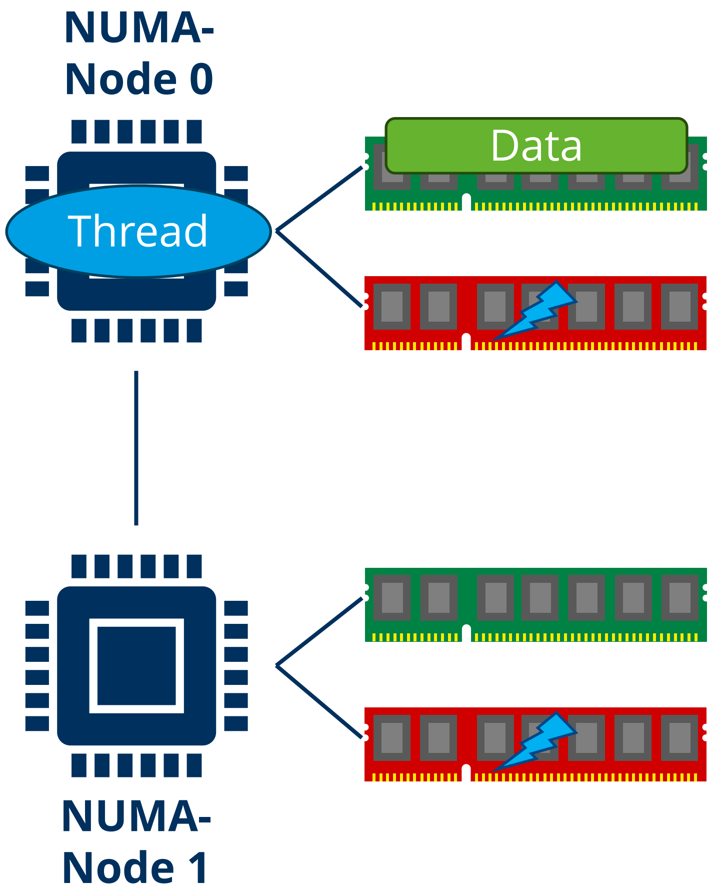

Funded by

Motivation

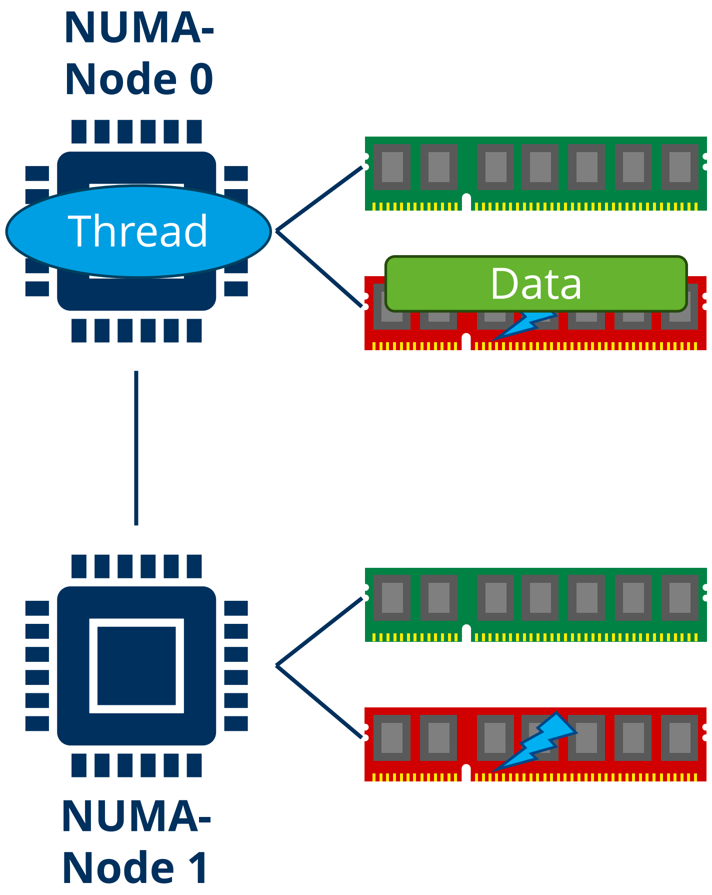

Funded by

Motivation

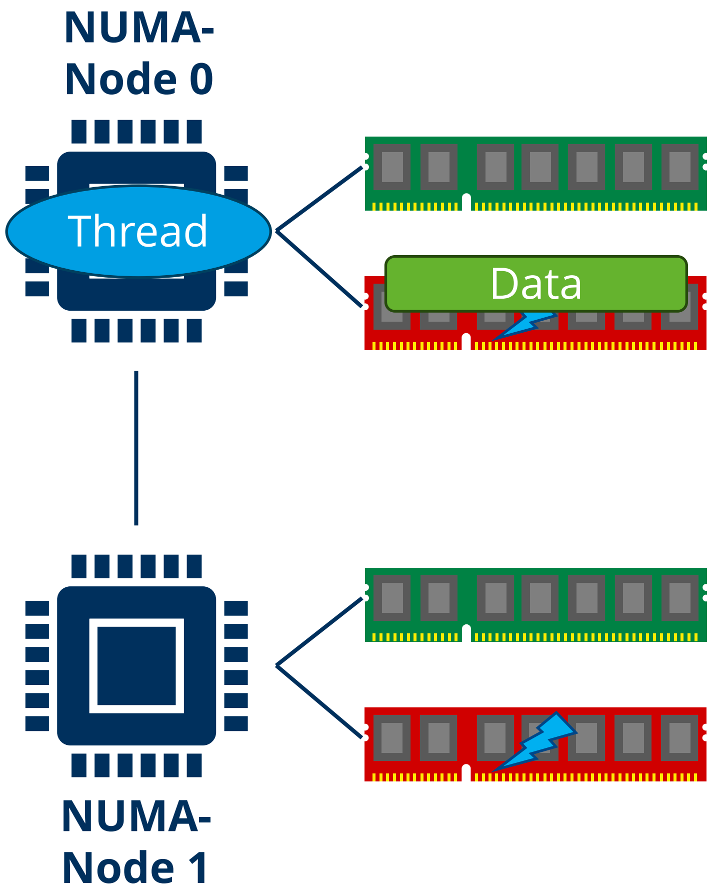

std::memcpy(a, b, 2^{30});

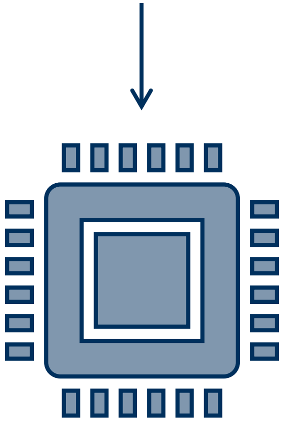

Motivation

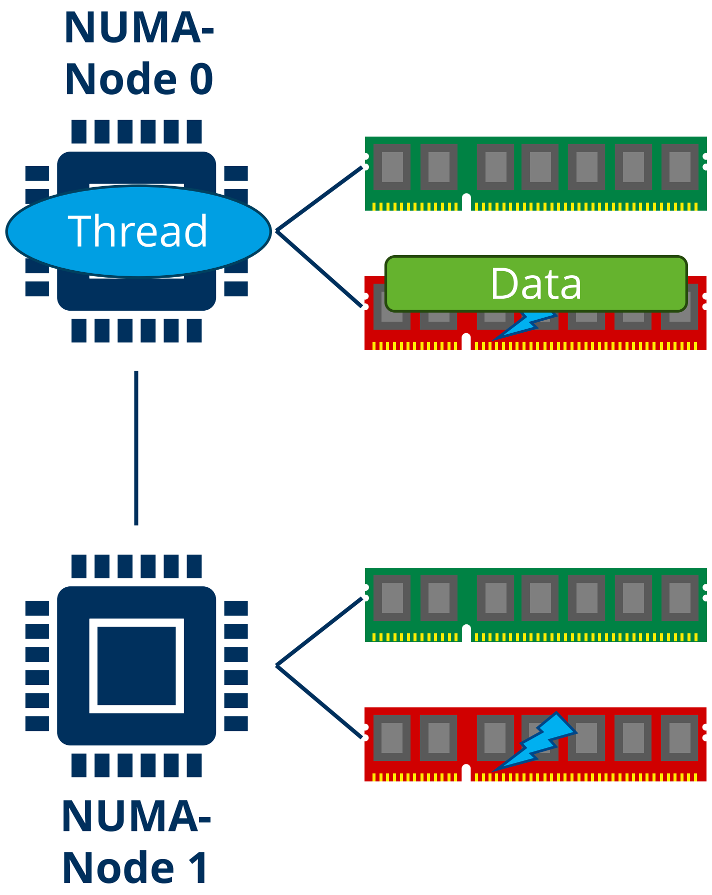

std::memcpy(a, b, 2^{30});

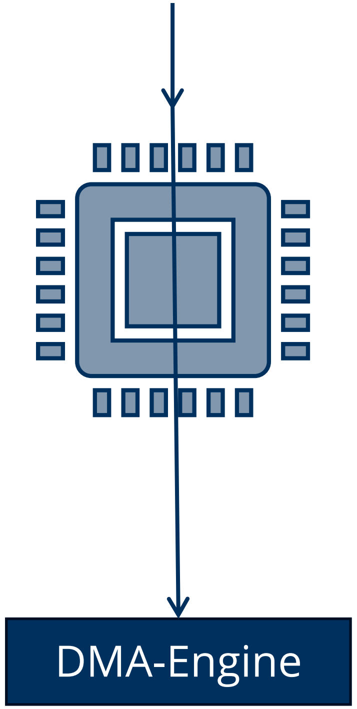

Funded by DFG
DRESDEN concept

Motivation

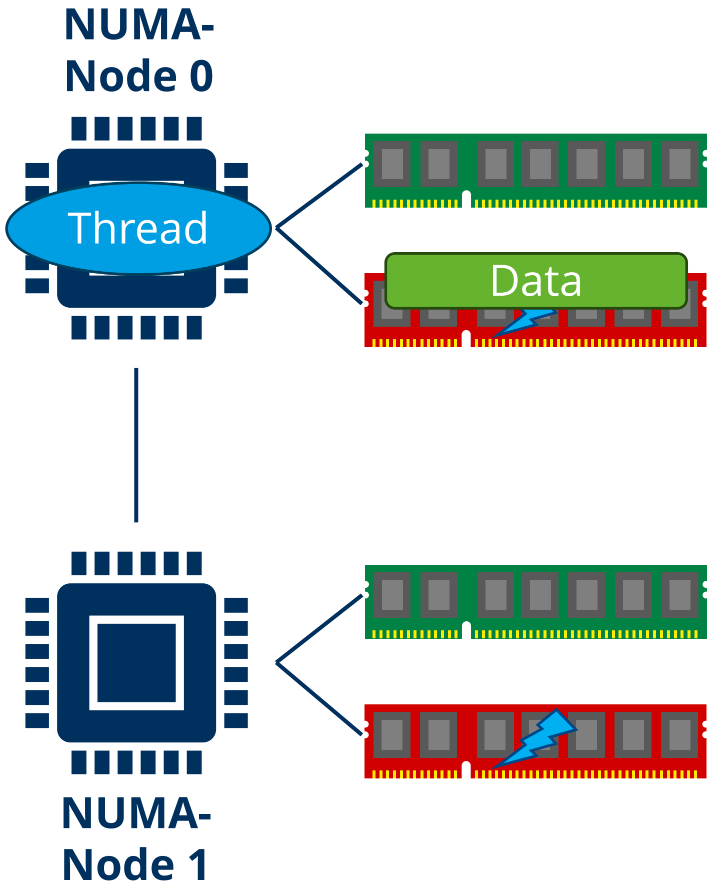

std::memcpy(a, b, 2^{30});

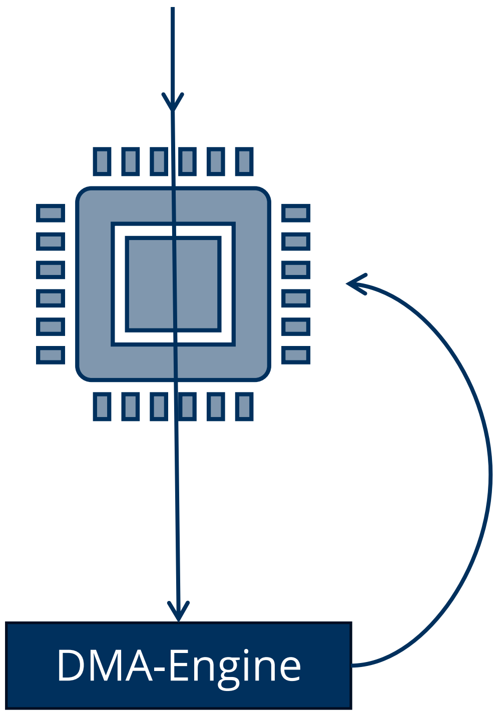

Funded by DFG
DRESDEN concept

Motivation

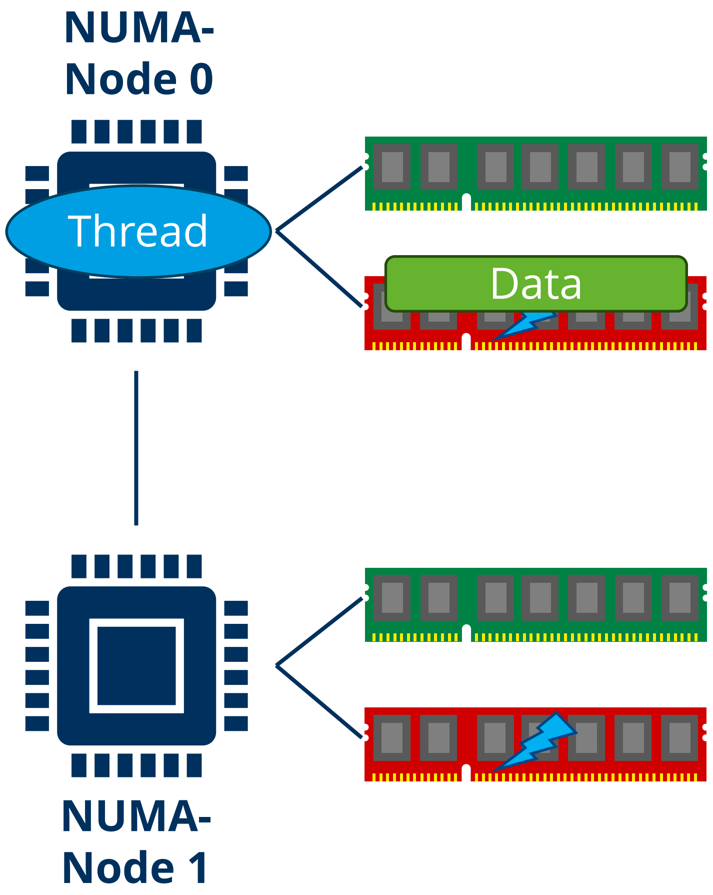

std::memcpy(a, b, 2^{30});

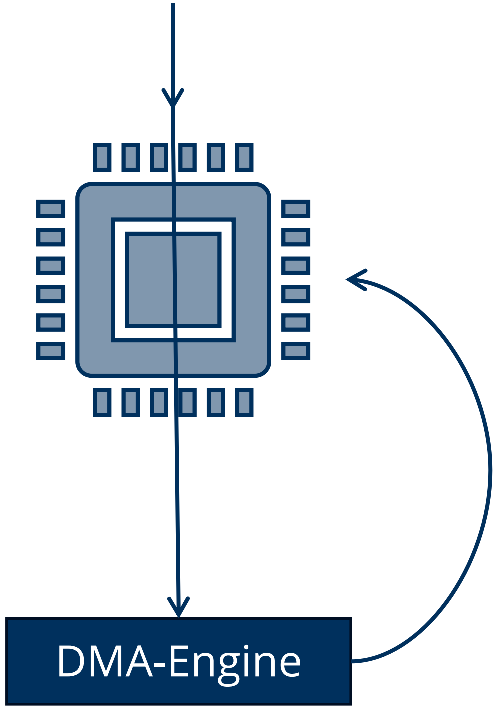

Sapphire Rapids
Microarchitecture

Funded by DFG

Motivation

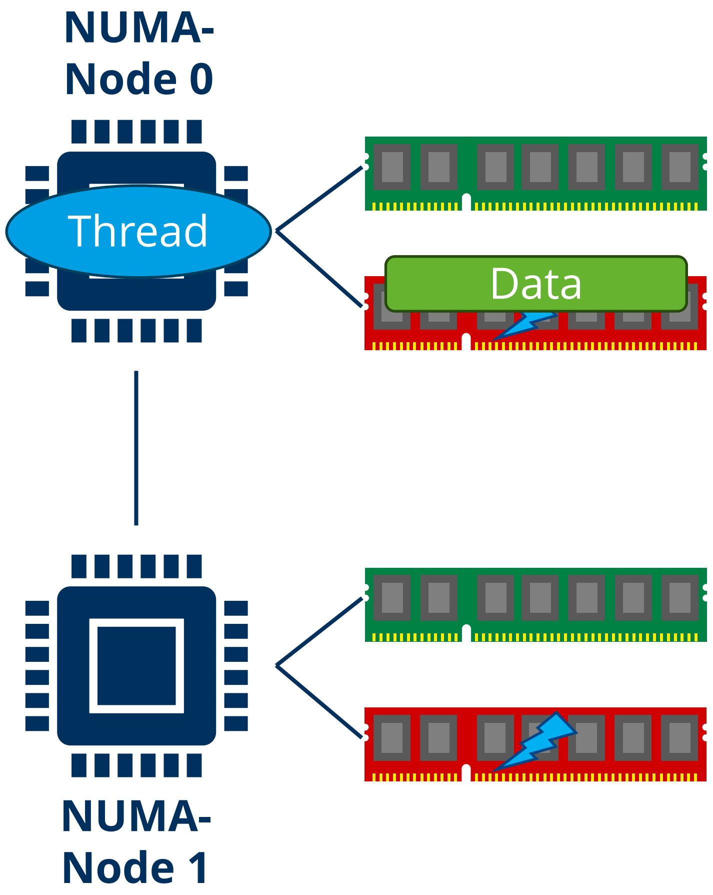

std::memcpy(a, b, 2^{30});

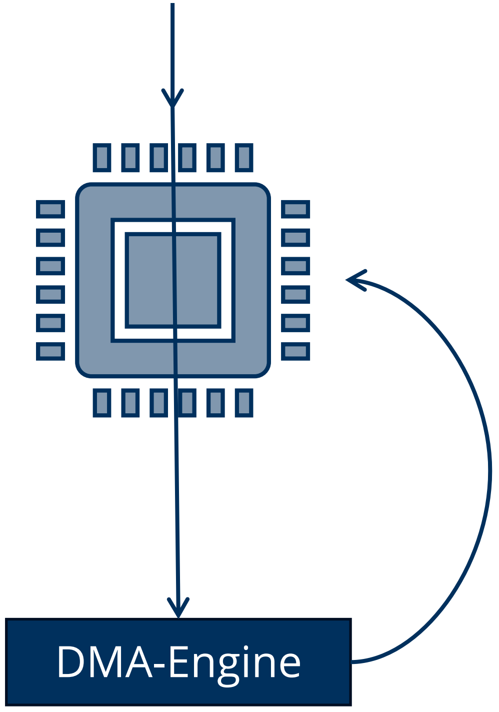

Sapphire Rapids
Microarchitecture

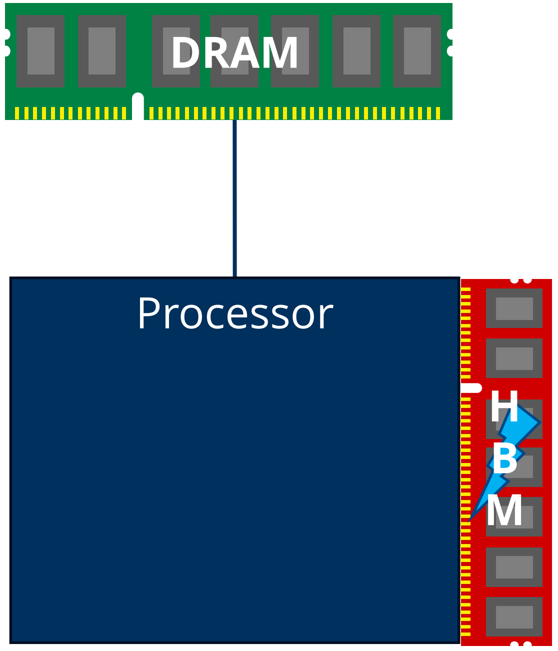

Funded by DFG
DRESDEN concept

Motivation

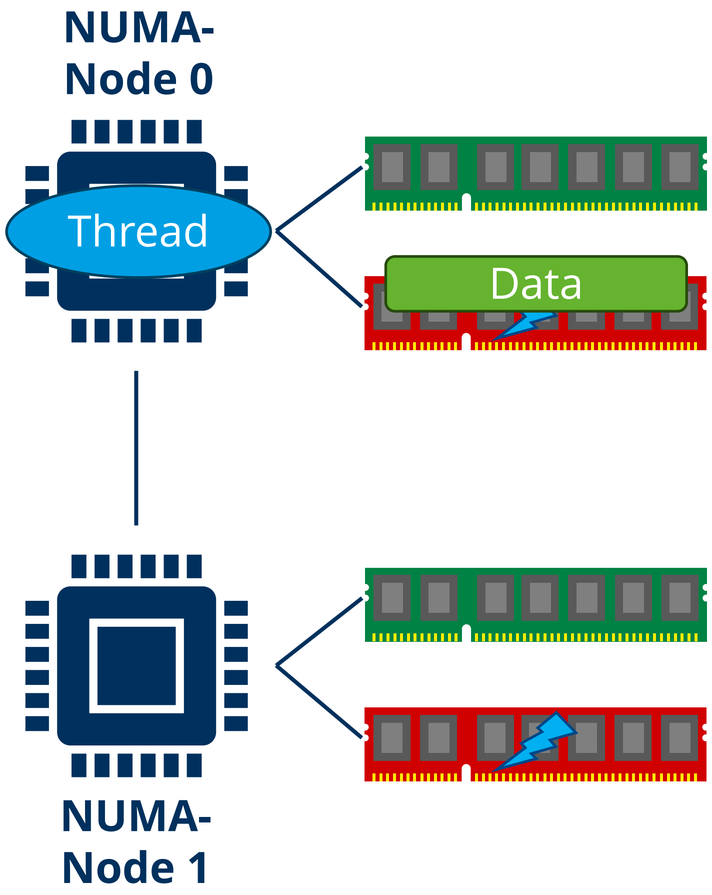

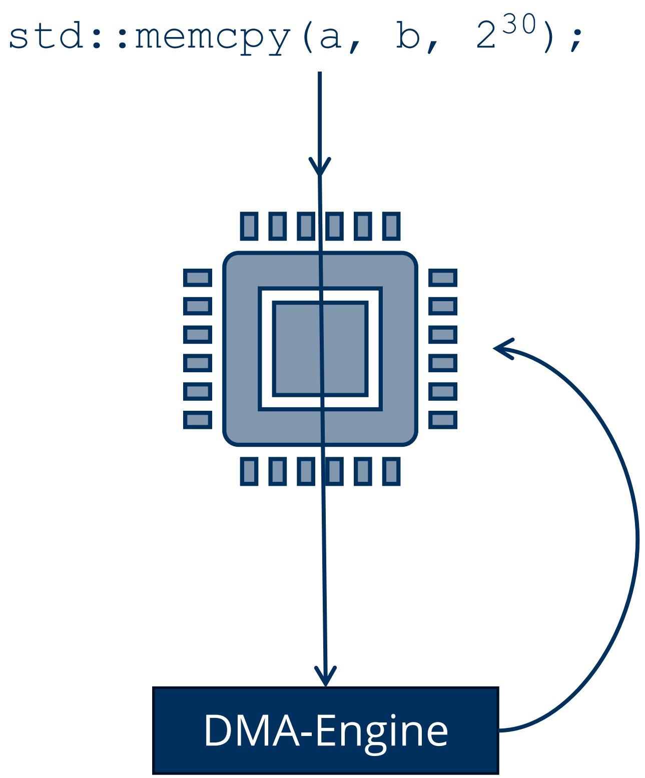
std::memcpy(a, b, 2^{30});

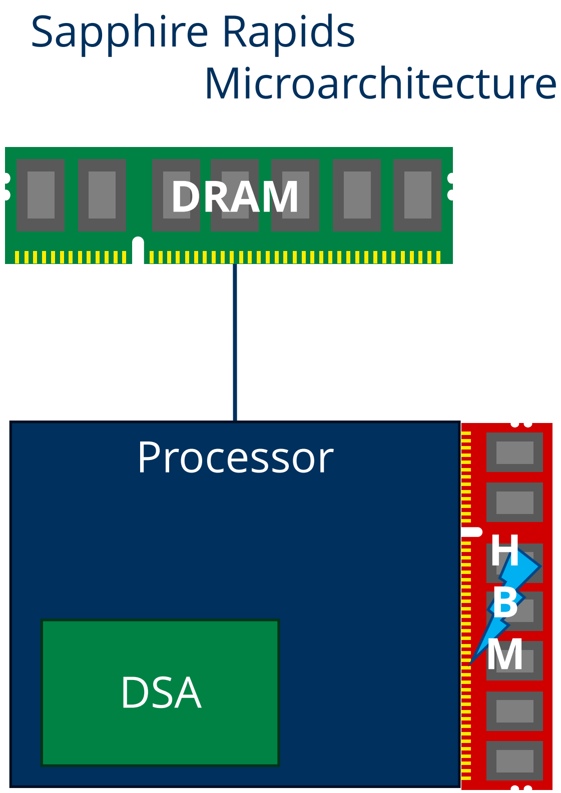

Funded by DFG
DRESDEN concept

DSA - Intel Data Streaming Accelerator

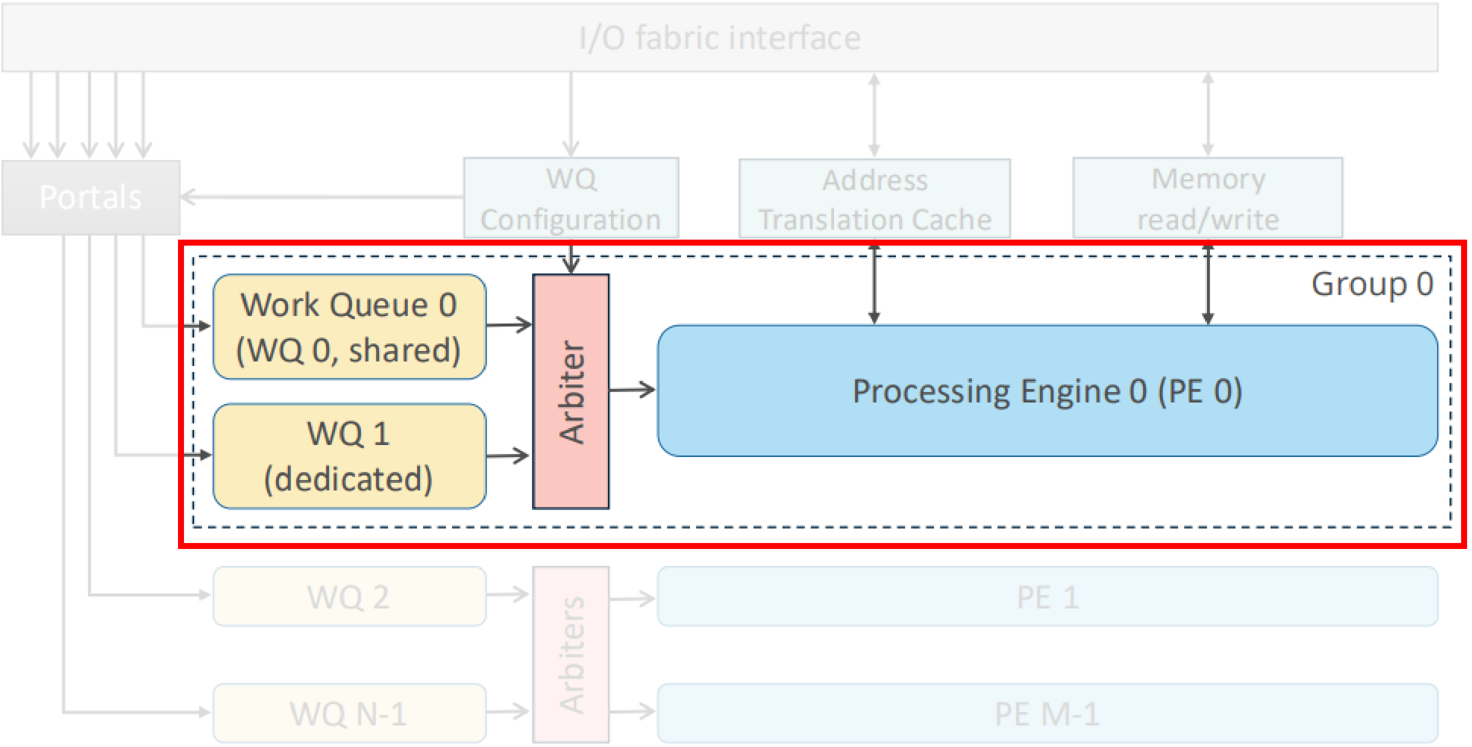

Reese Kuper, et al. 2024. A Quantitative Analysis and Guidelines of Data Streaming Accelerator in Modern Intel Xeon Scalable Processors. In *ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS '24)*, 37–54.

The DSA supports:
- Memory Copy
- Memory Fill
- Memory Compare

⋮

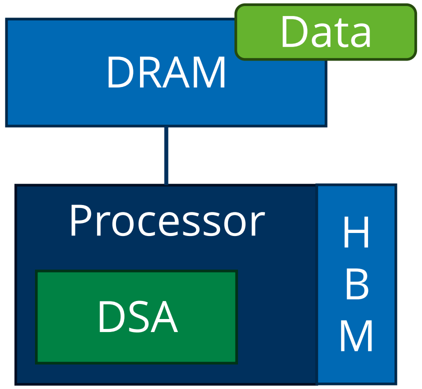

Funded by

DSA - Intel Data Streaming Accelerator

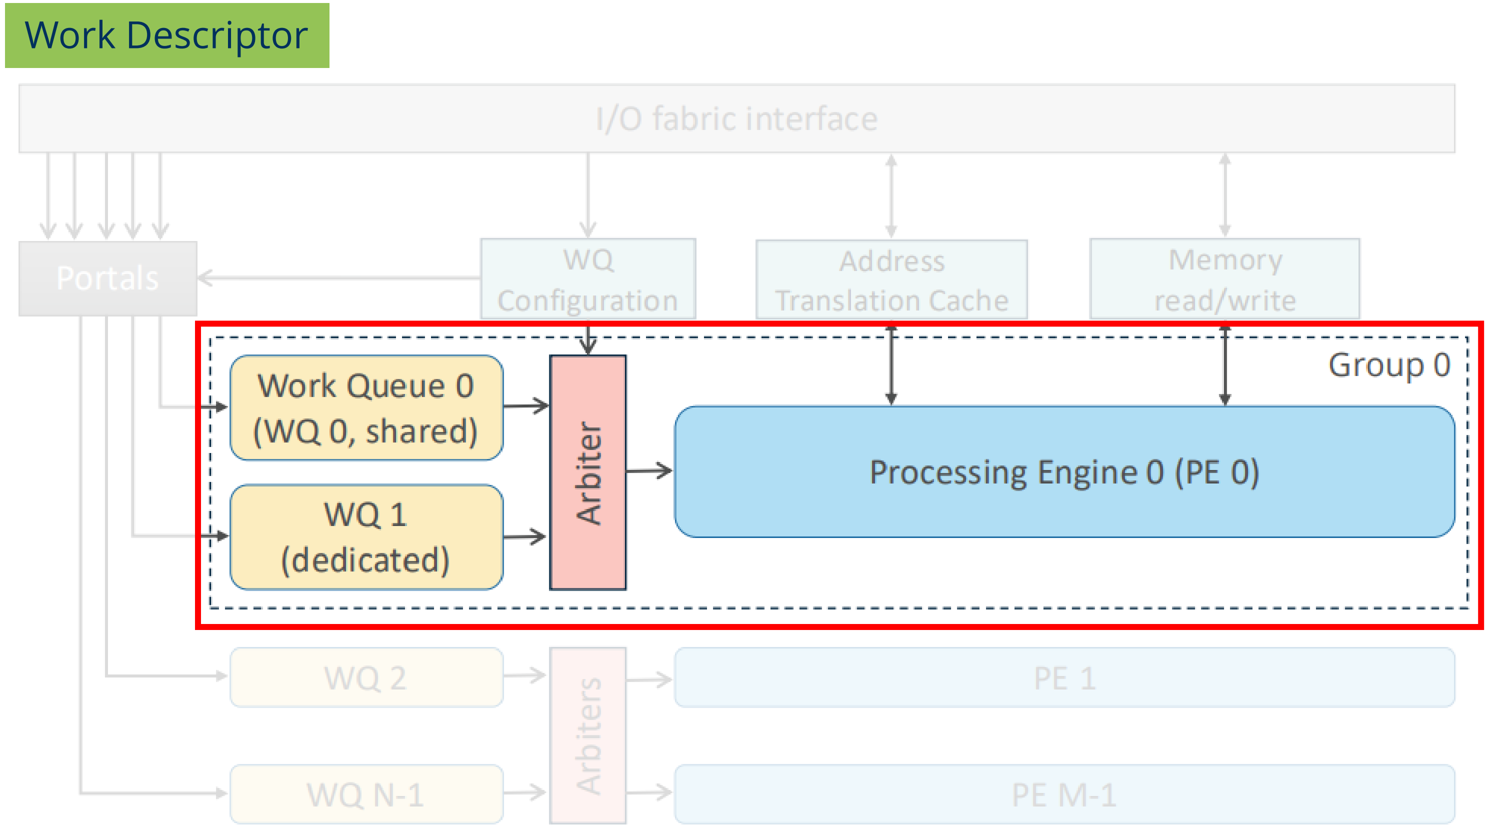

The DSA supports:

- Memory Copy
- Memory Fill
- Memory Compare

⋮

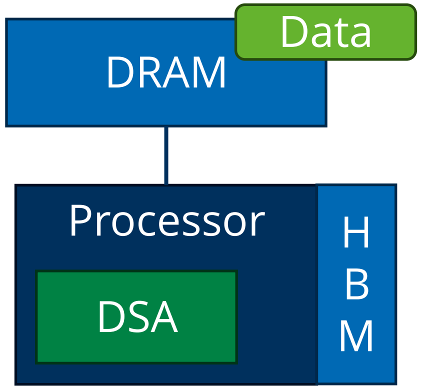

Reese Kuper, et al. 2024. A Quantitative Analysis and Guidelines of Data Streaming Accelerator in Modern Intel Xeon Scalable Processors. In *ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS '24)*, 37–54.

Funded by DFG
DRESDEN concept

DSA - Intel Data Streaming Accelerator

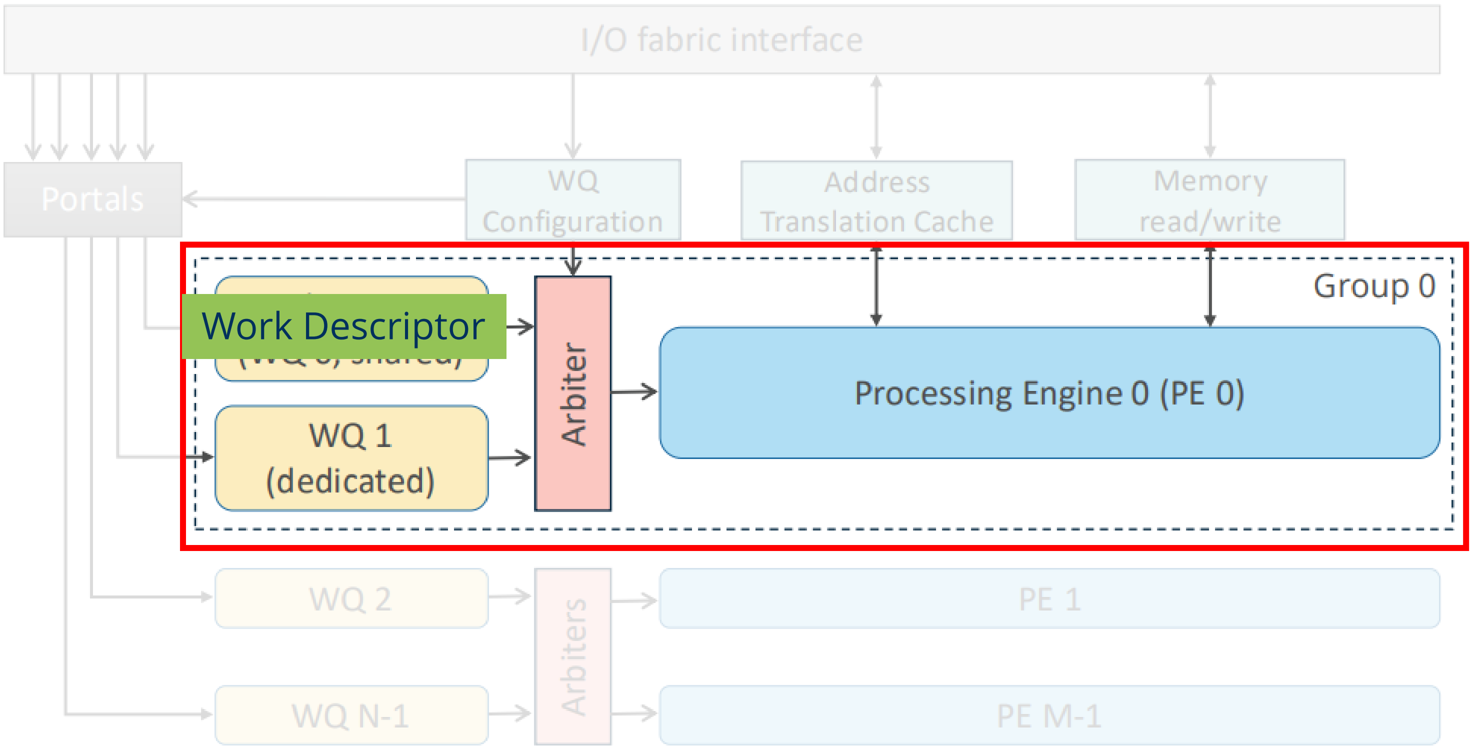

Reese Kuper, et al. 2024. A Quantitative Analysis and Guidelines of Data Streaming Accelerator in Modern Intel Xeon Scalable Processors. In *ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS '24)*, 37–54.

The DSA supports:

- Memory Copy
- Memory Fill
- Memory Compare

⋮

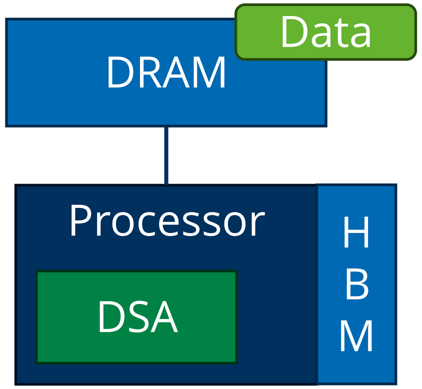

Funded by DFG

DRESDEN concept

DSA - Intel Data Streaming Accelerator

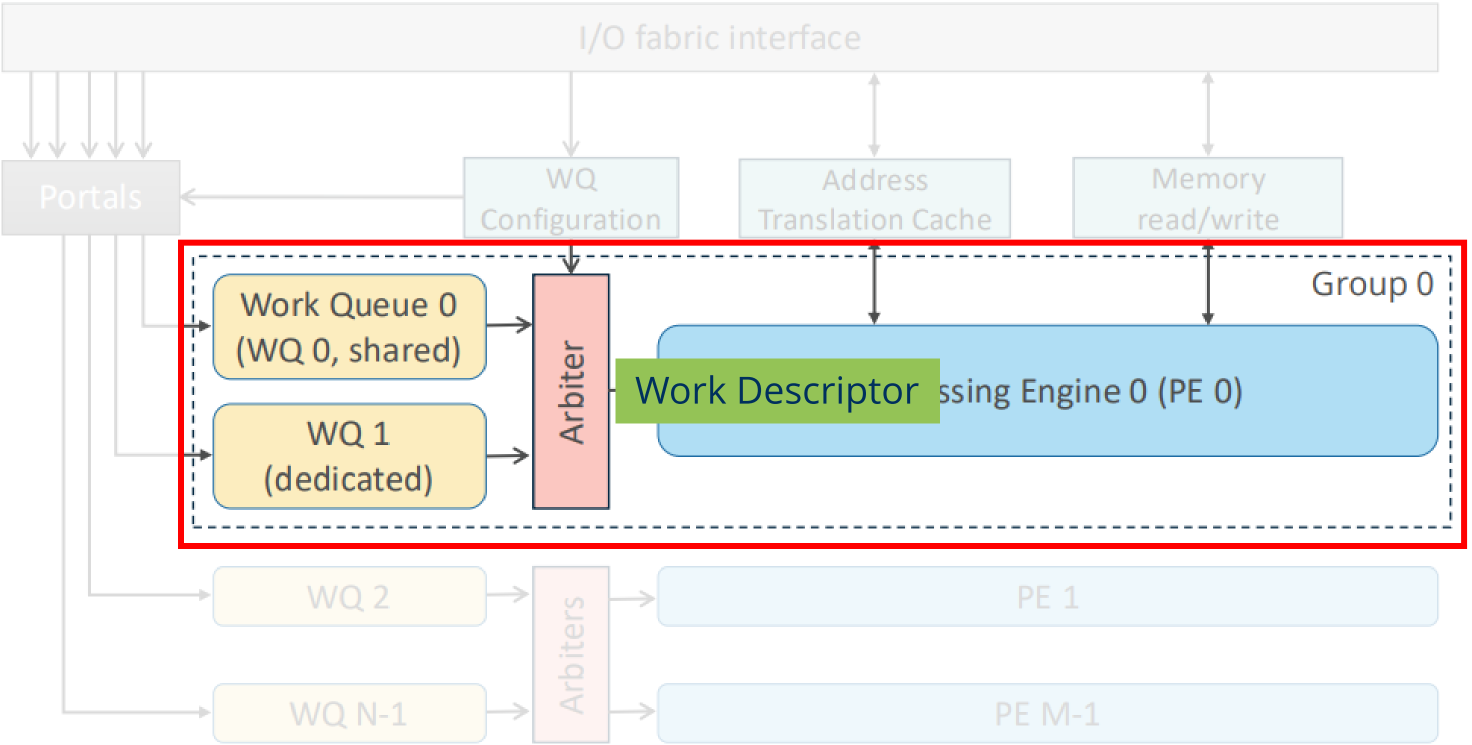

The DSA supports:

- Memory Copy
- Memory Fill
- Memory Compare

⋮

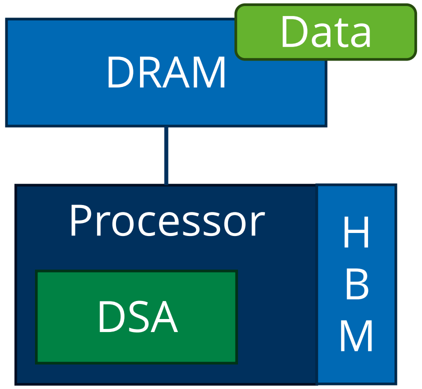

Reese Kuper, et al. 2024. A Quantitative Analysis and Guidelines of Data Streaming Accelerator in Modern Intel Xeon Scalable Processors. In *ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS '24)*, 37–54.

Funded by DFG
DRESDEN concept

DSA - Intel Data Streaming Accelerator

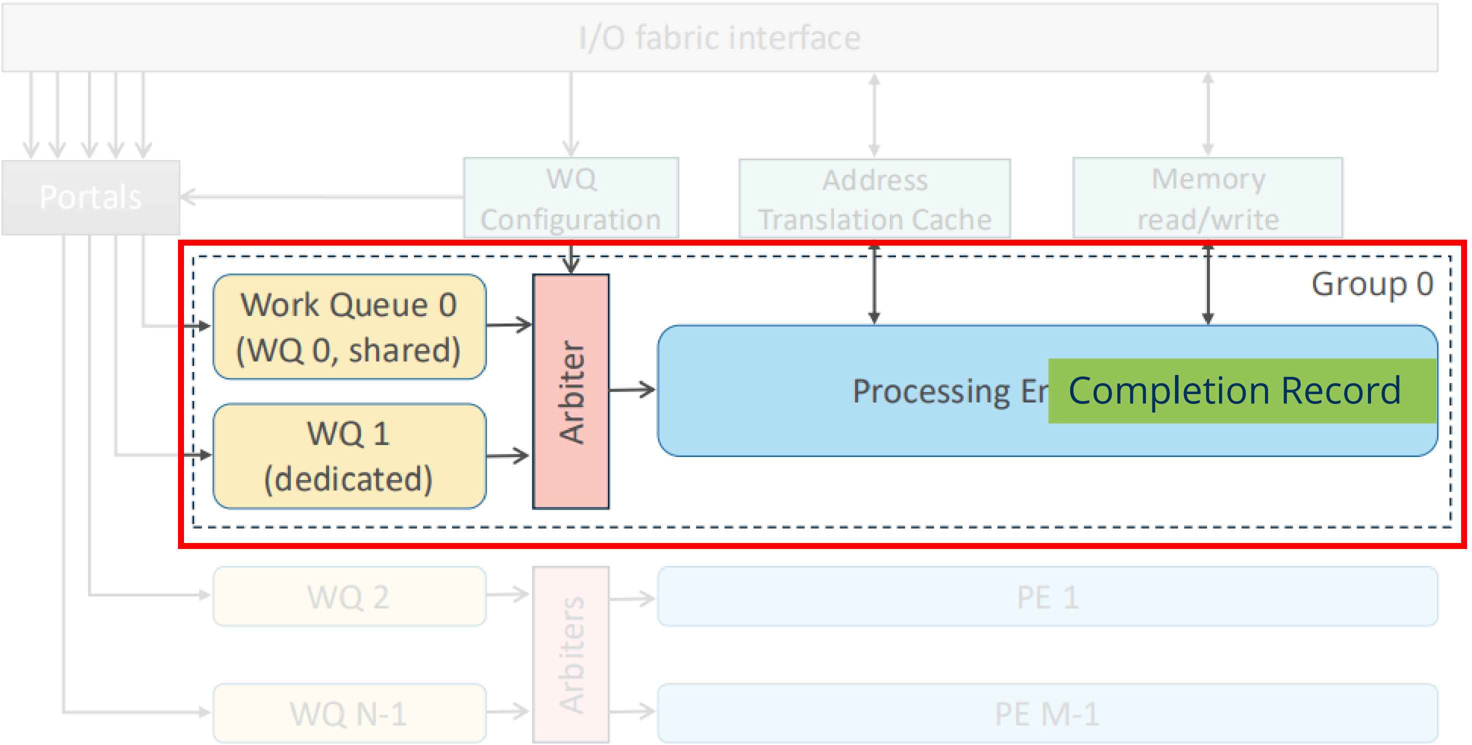

Reese Kuper, et al. 2024. A Quantitative Analysis and Guidelines of Data Streaming Accelerator in Modern Intel Xeon Scalable Processors. In *ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS '24)*, 37–54.

The DSA supports:

- Memory Copy
- Memory Fill
- Memory Compare

⋮

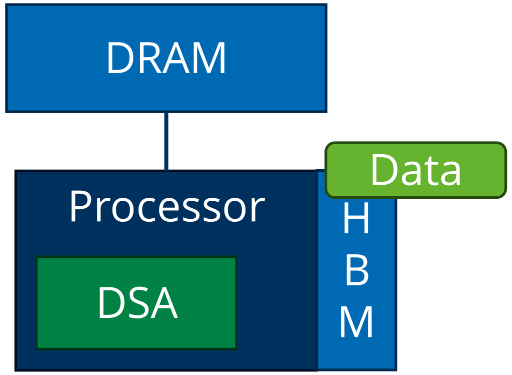

TECHNISCHE UNIVERSITÄT DRESDEN
Funded by DFG
DRESDEN concept

DSA - Intel Data Streaming Accelerator

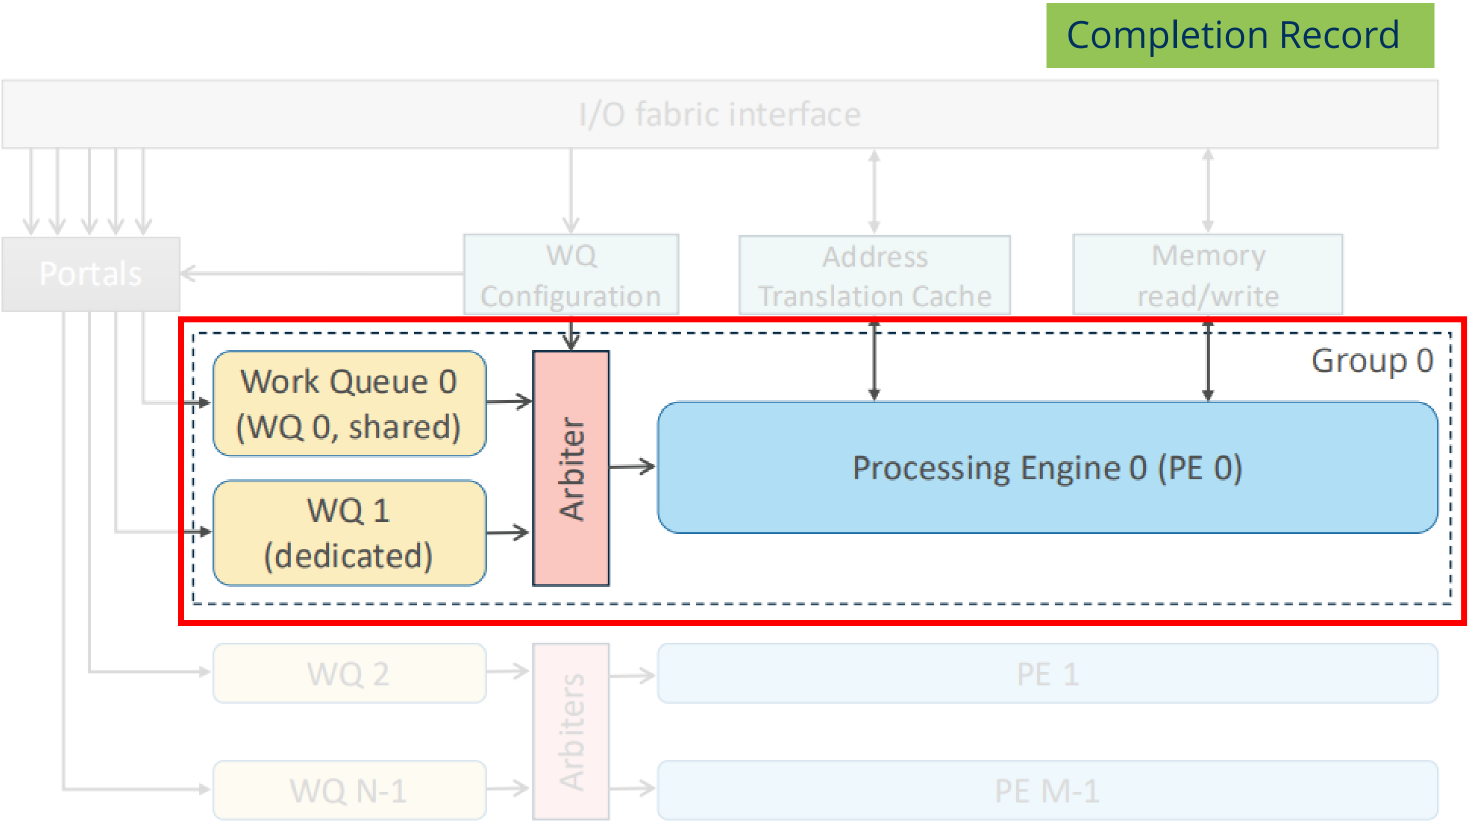

Reese Kuper, et al. 2024. A Quantitative Analysis and Guidelines of Data Streaming Accelerator in Modern Intel Xeon Scalable Processors. In *ACM International Conference on Architectural Support for Programming Languages and Operating Systems (ASPLOS '24)*, 37–54.

The DSA supports:

- Memory Copy
- Memory Fill
- Memory Compare

⋮

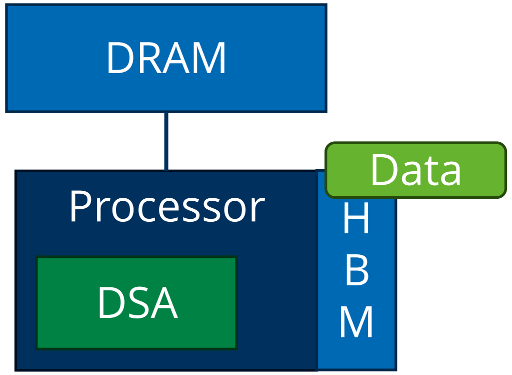

Funded by DFG
DRESDEN concept

Benchmarking the DSA

Funded by
DFG
DRESDEN concept

Benchmarking the DSA

Setup

System under Test: Intel Xeon CPU Max 9468

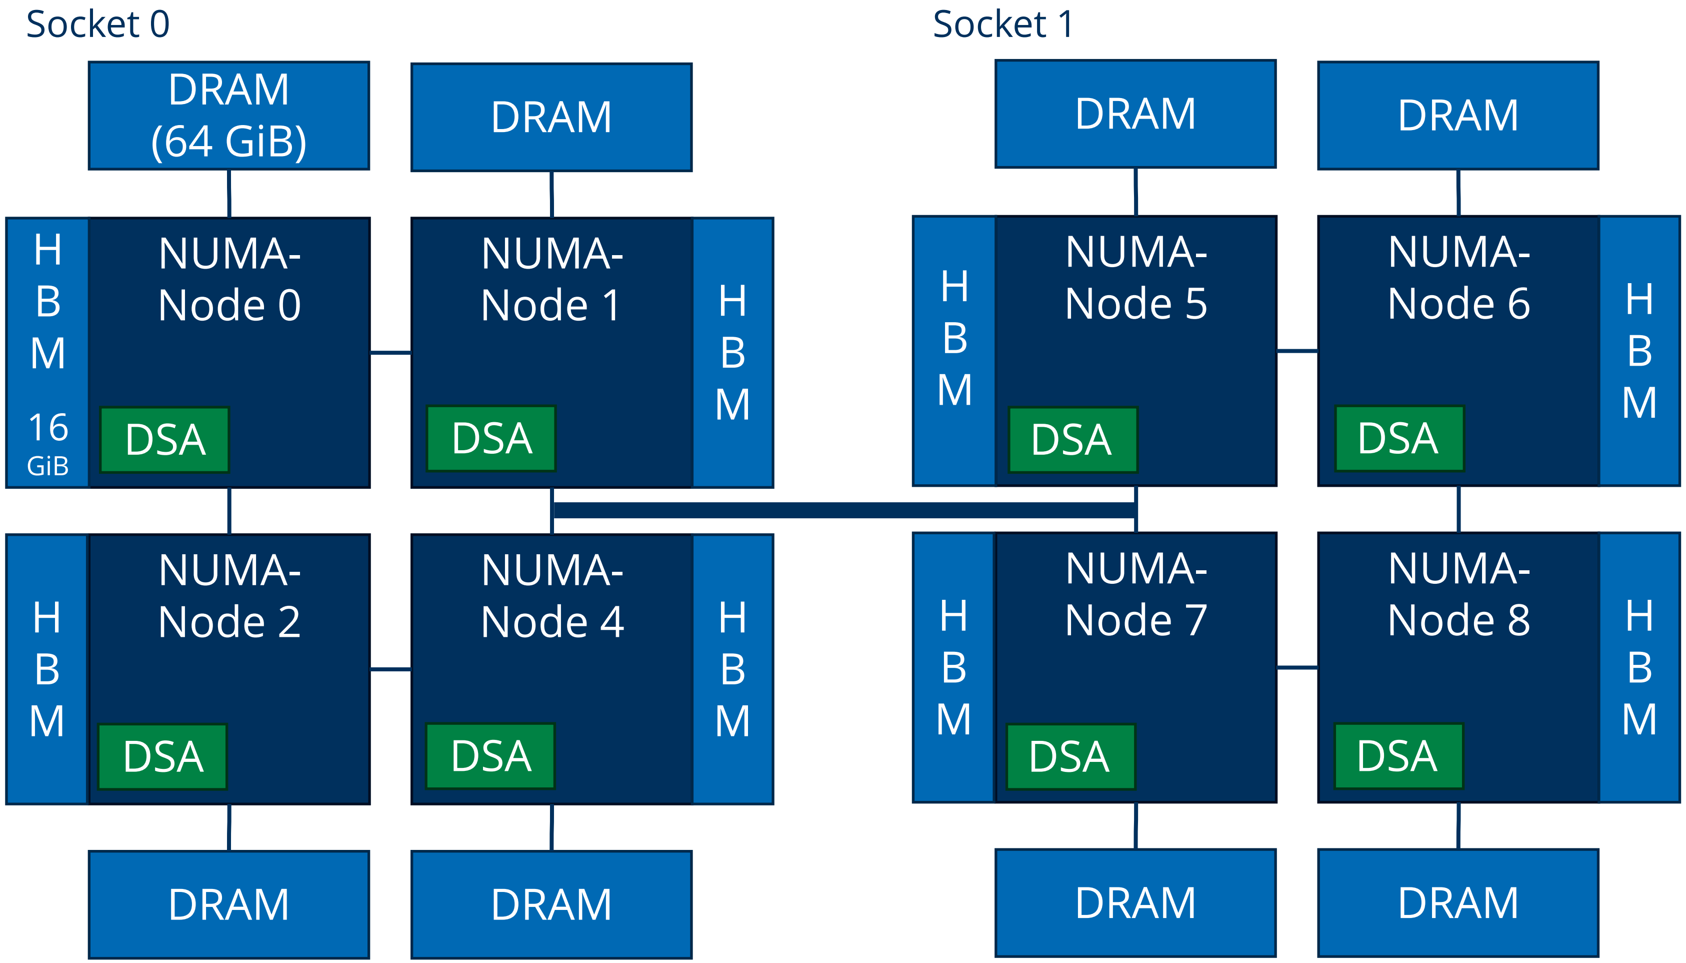

Funded by

Benchmarking the DSA
Setup

1) CPU vs. DSA

System under Test: Intel Xeon CPU Max 9468

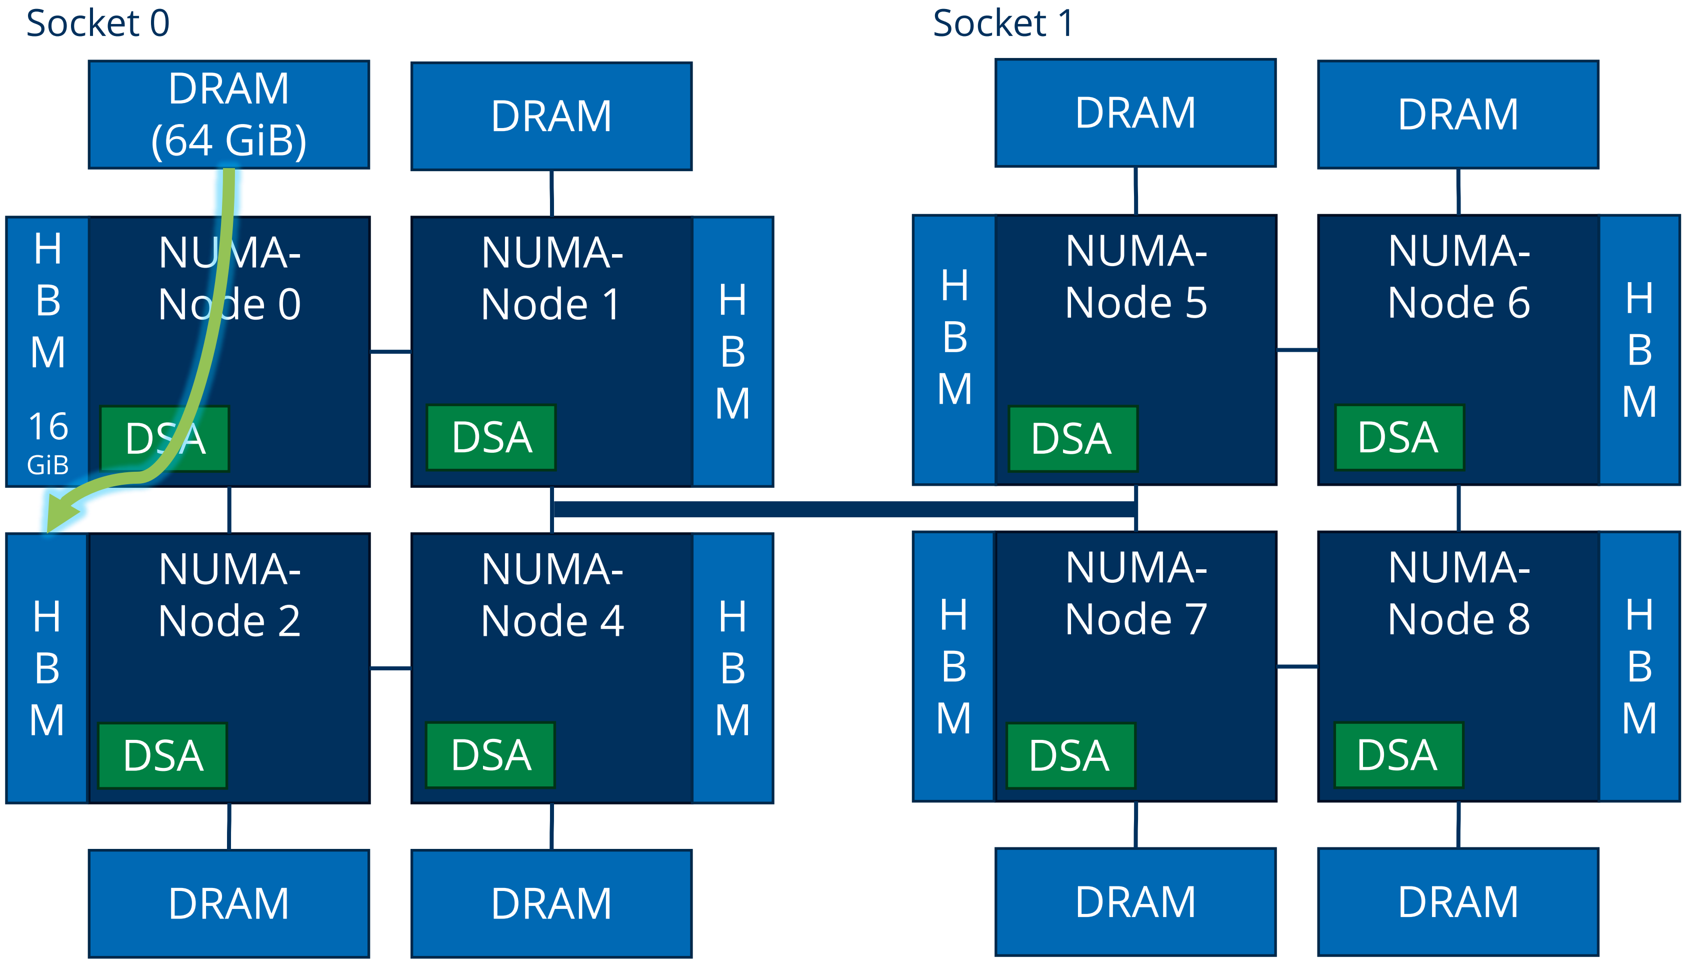

Funded by DFG
DRESDEN concept

Benchmarking the DSA
Setup

1) CPU vs. DSA
2) Interference between DSA and CPU

System under Test: Intel Xeon CPU Max 9468

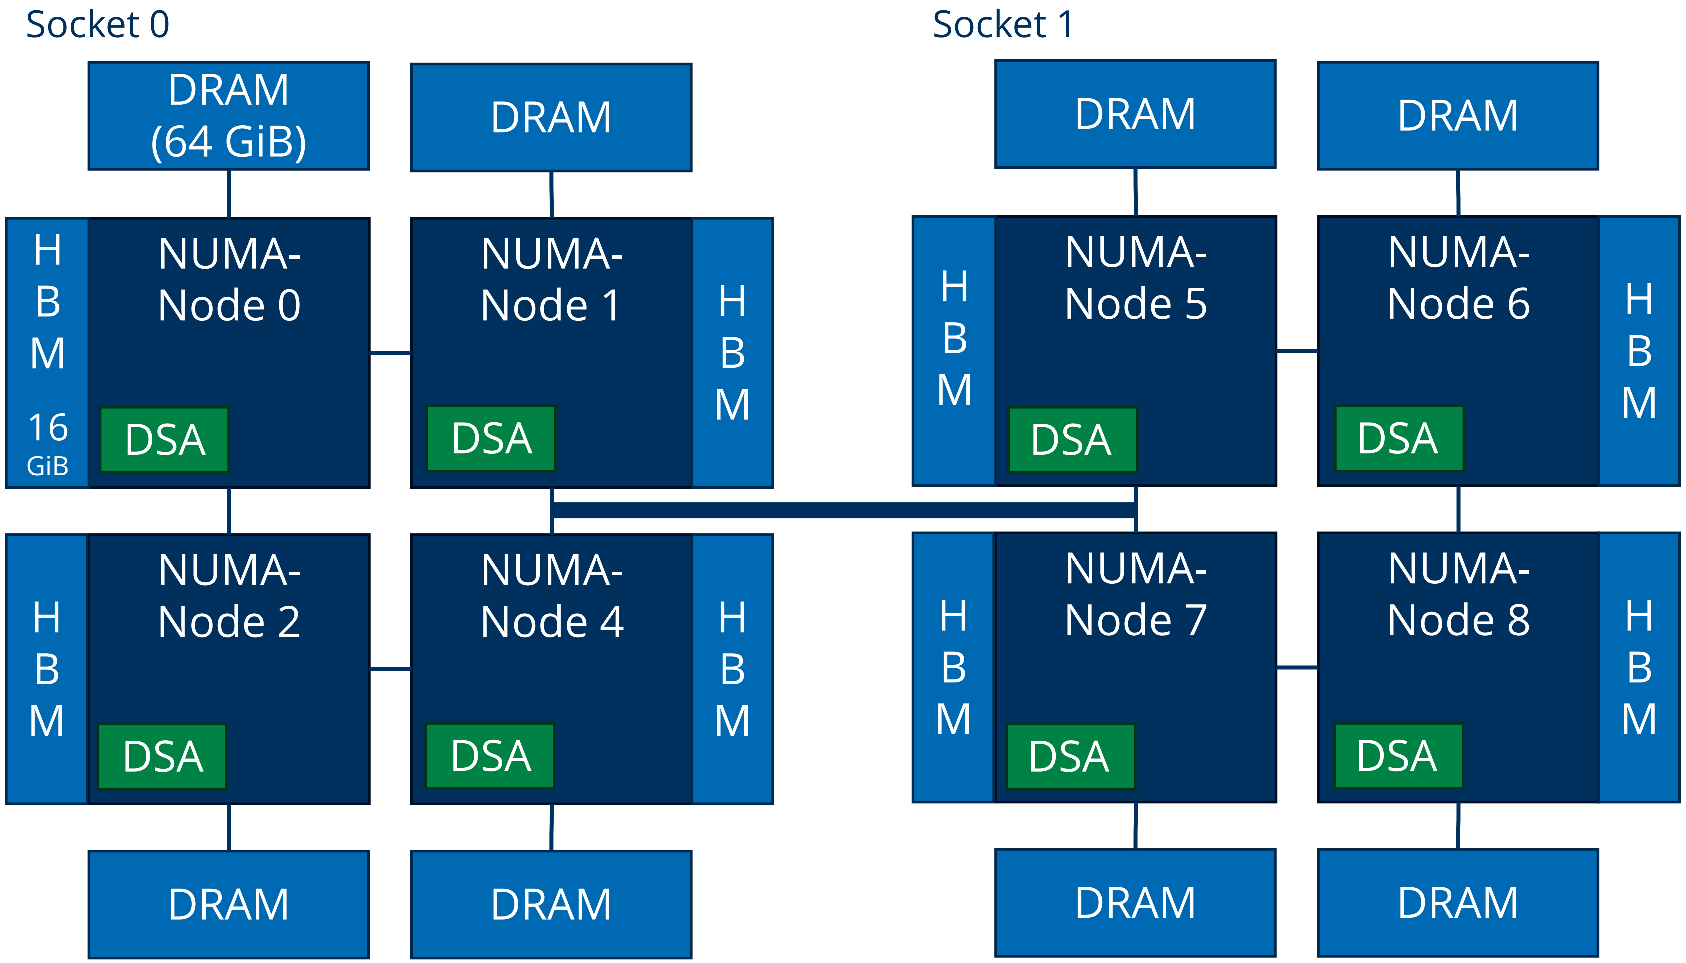

Funded by DFG
DRESDEN concept

Benchmarking the DSA
Setup

1) CPU vs. DSA
2) Interference between DSA and CPU
   1) Compute-Intensive CPU Threads

System under Test: Intel Xeon CPU Max 9468

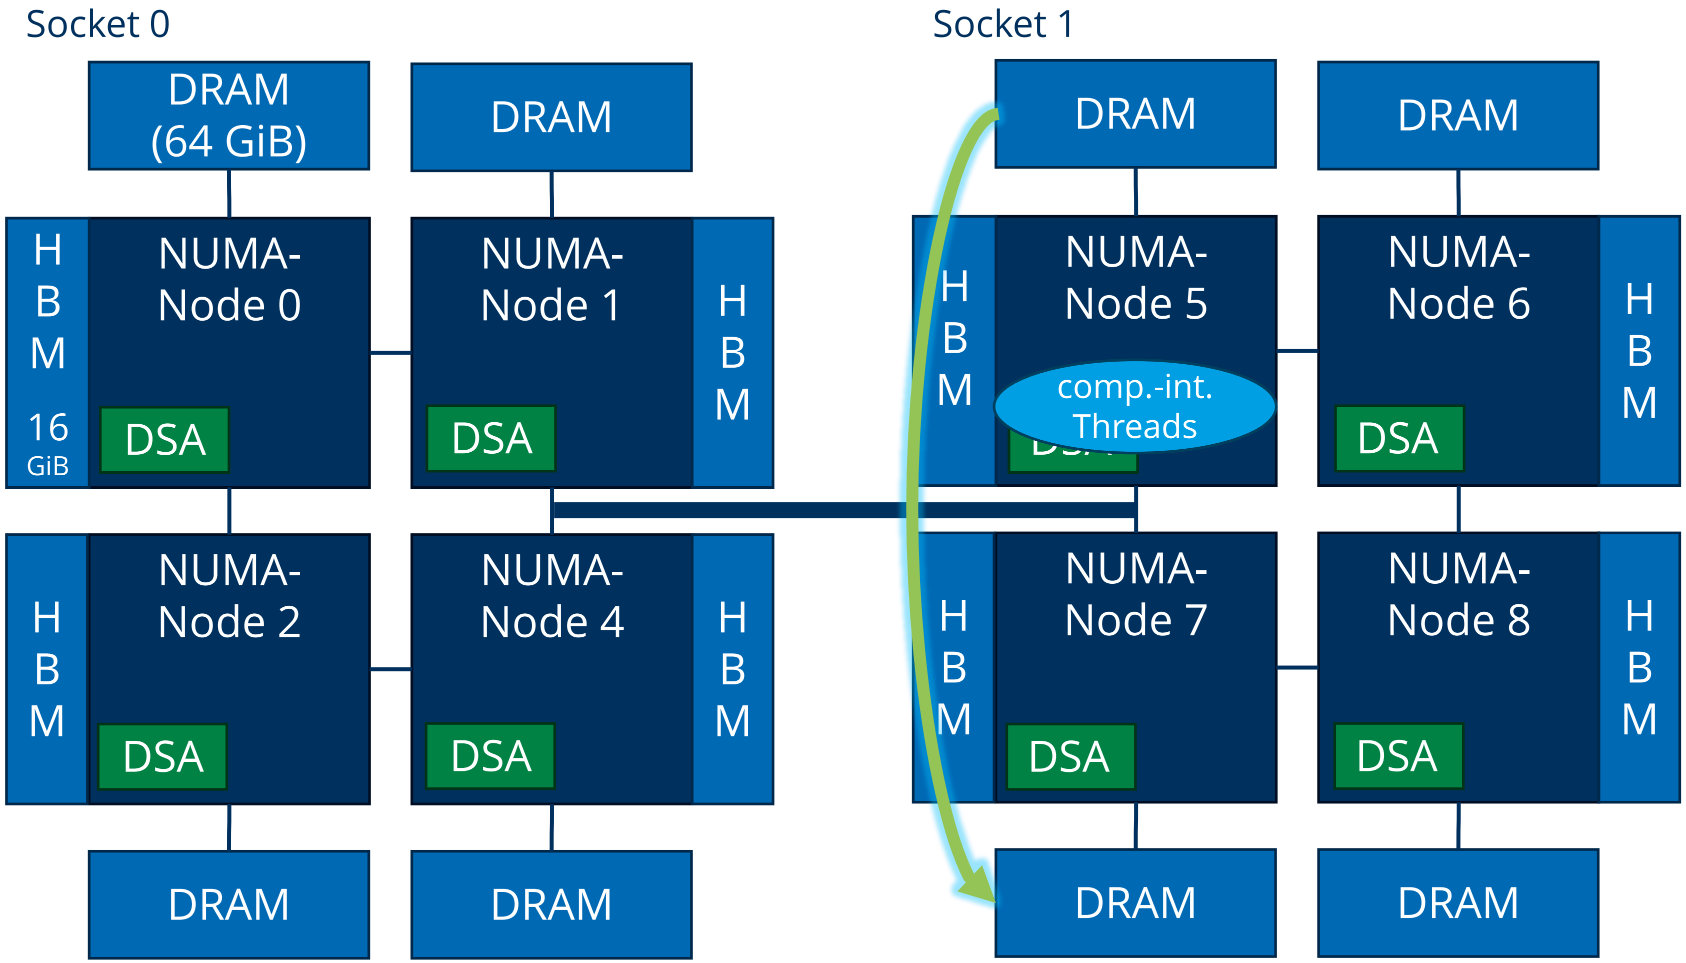

Funded by DFG
DRESDEN concept

Benchmarking the DSA
Setup

1) CPU vs. DSA
2) Interference between DSA and CPU
   1) Compute-Intensive CPU Threads
   2) Data-Intensive CPU Threads

System under Test: Intel Xeon CPU Max 9468

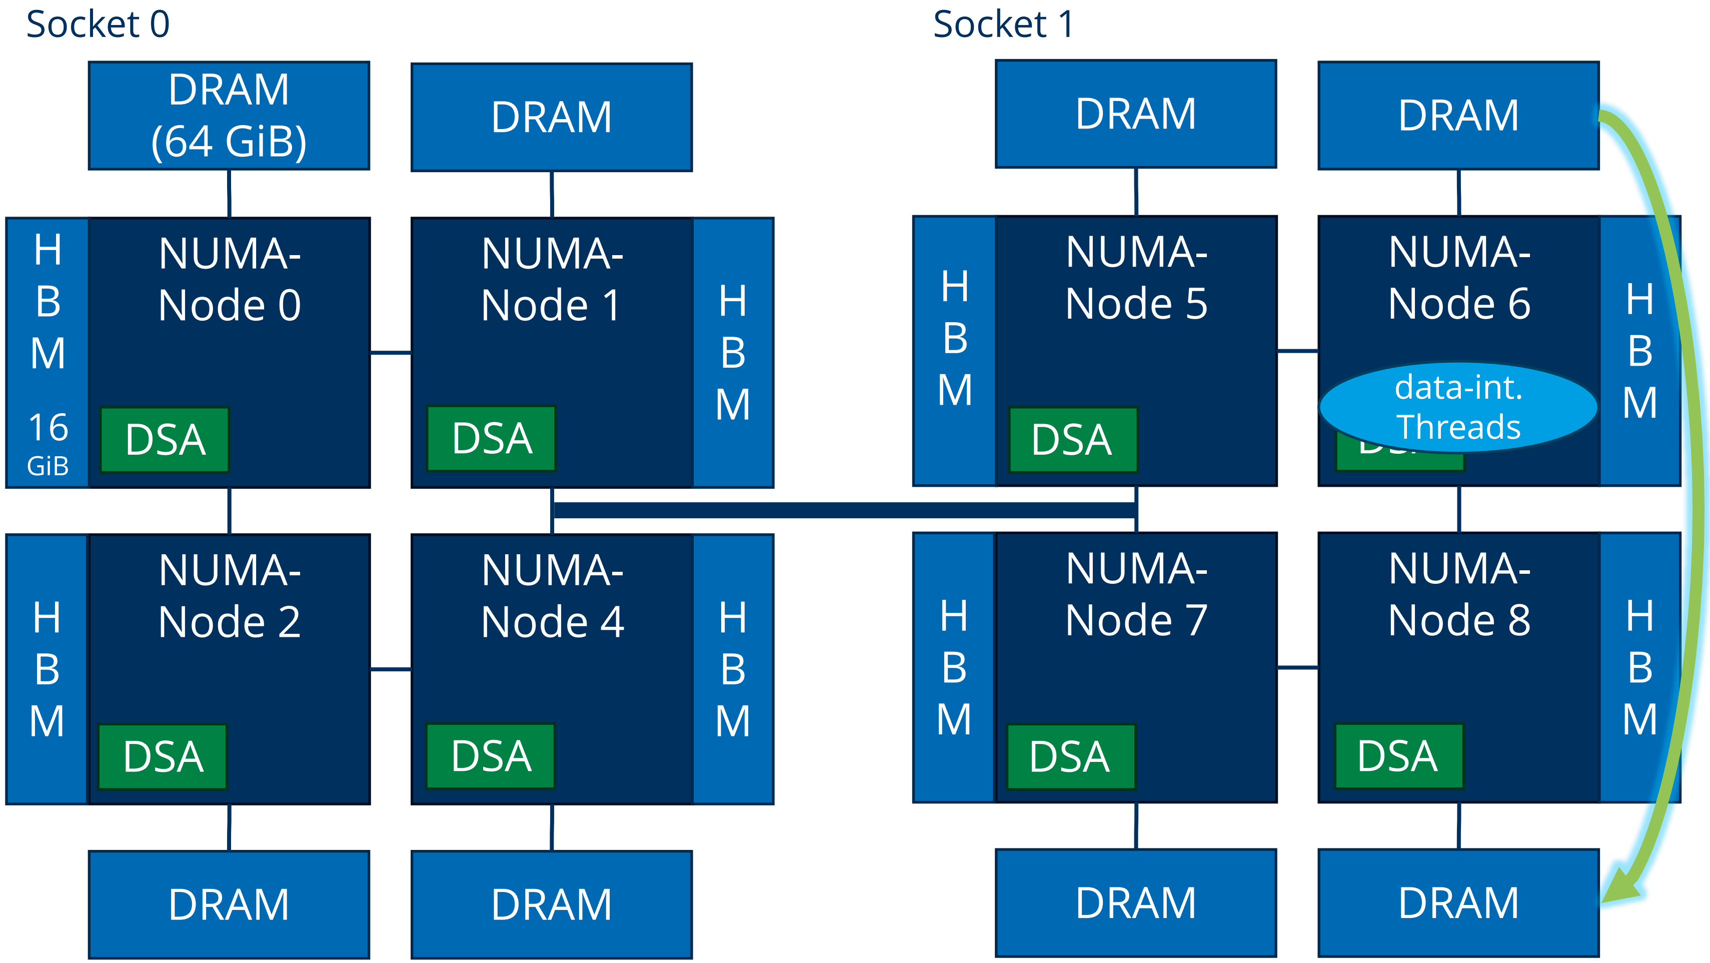

Funded by DFG

Benchmarking the DSA
Setup

1) CPU vs. DSA
2) Interference between DSA and CPU
    1) Compute-Intensive CPU Threads
    2) Data-Intensive CPU Threads
    3) Data-Intensive CPU Threads on HBM

System under Test: Intel Xeon CPU Max 9468

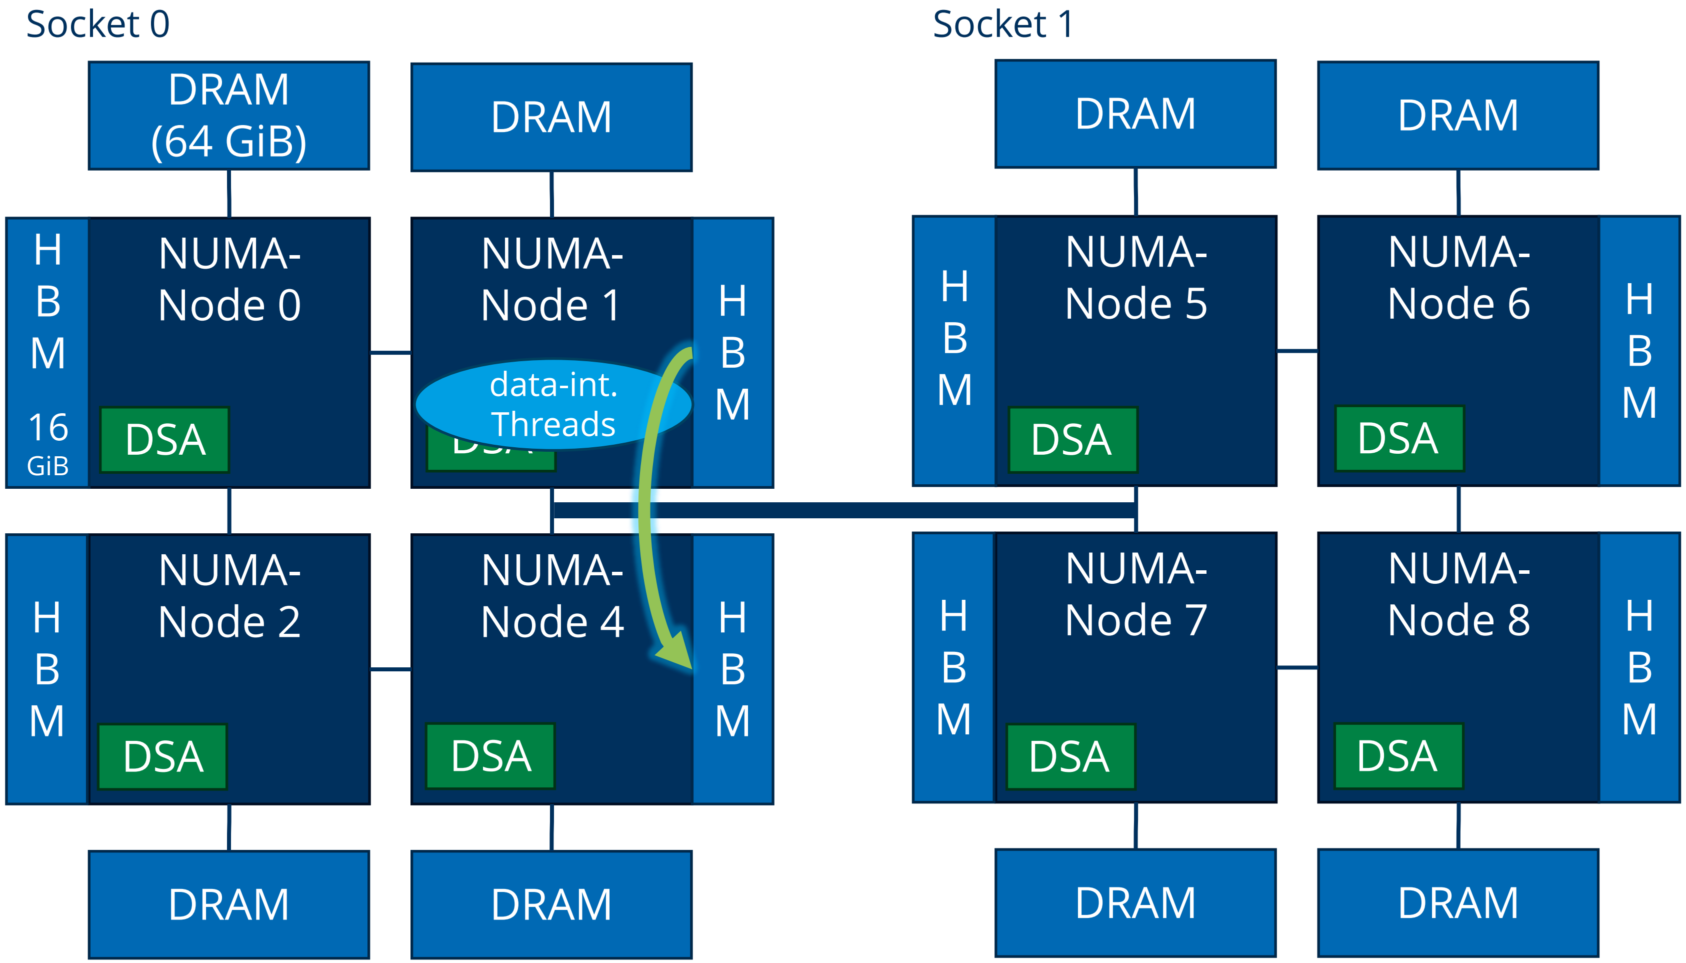

Funded by DFG
DRESDEN concept

Benchmarking the DSA
Setup

1) CPU vs. DSA
2) Interference between DSA and CPU
    1) Compute-Intensive CPU Threads
    2) Data-Intensive CPU Threads
    3) Data-Intensive CPU Threads on HBM
3) Inter-Socket Data Transfer with DSA

System under Test: Intel Xeon CPU Max 9468

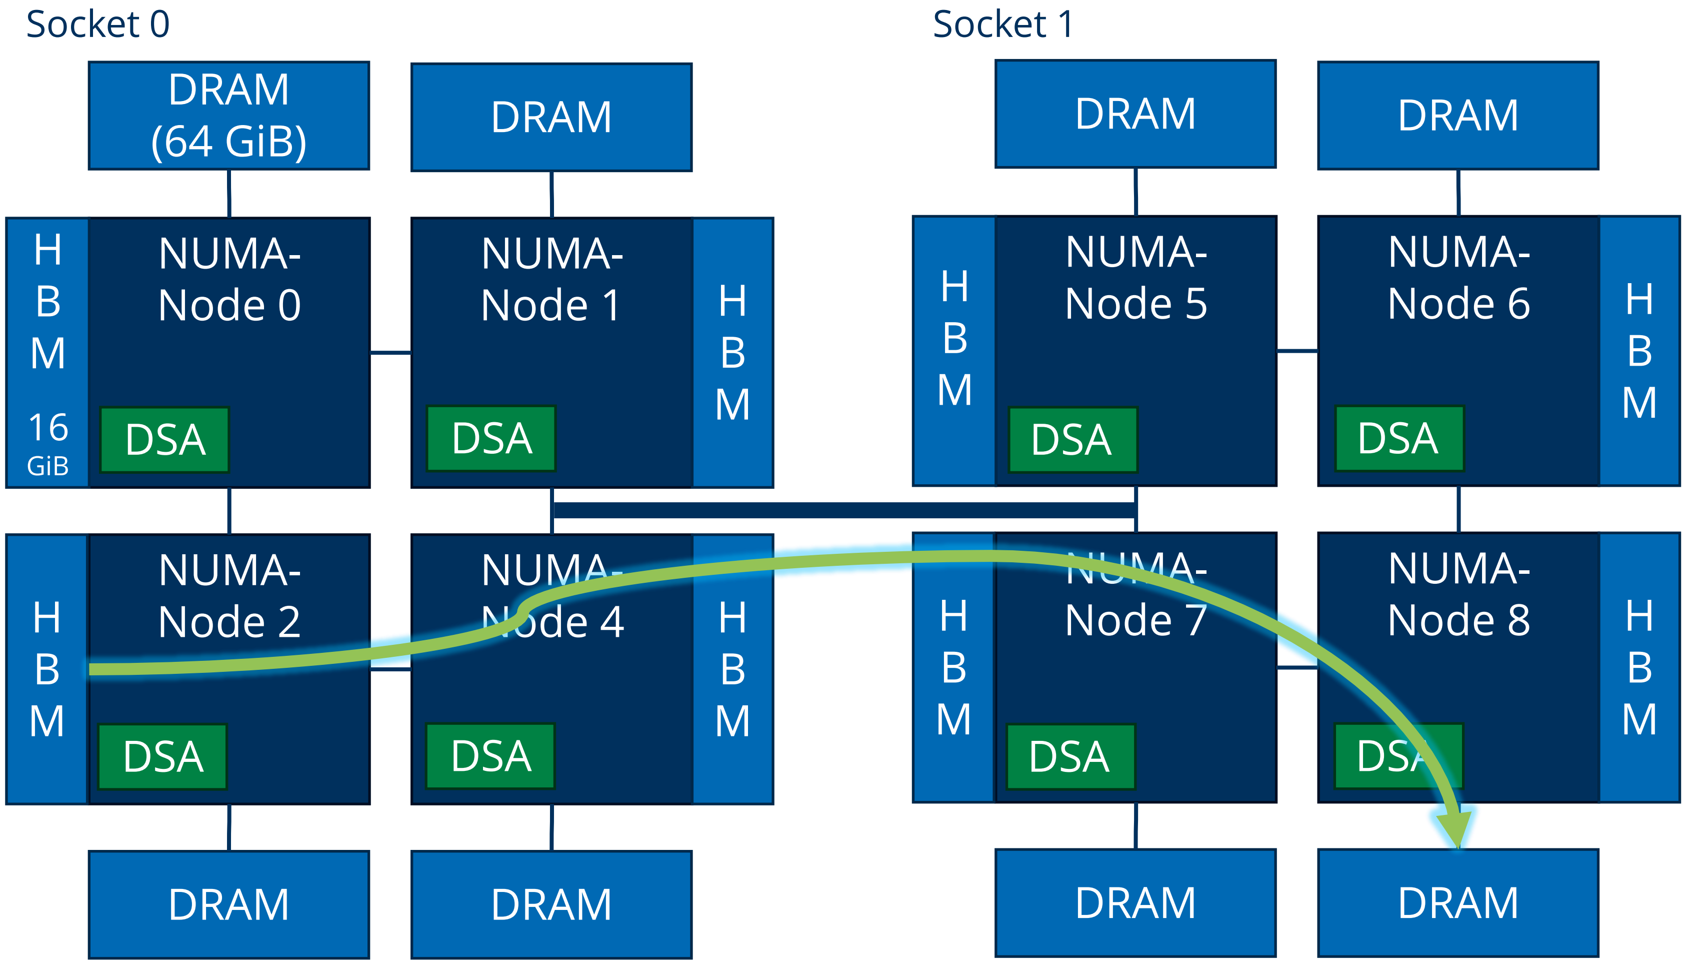

Funded by DFG
DRESDEN concept

Benchmark – CPU vs. DSA

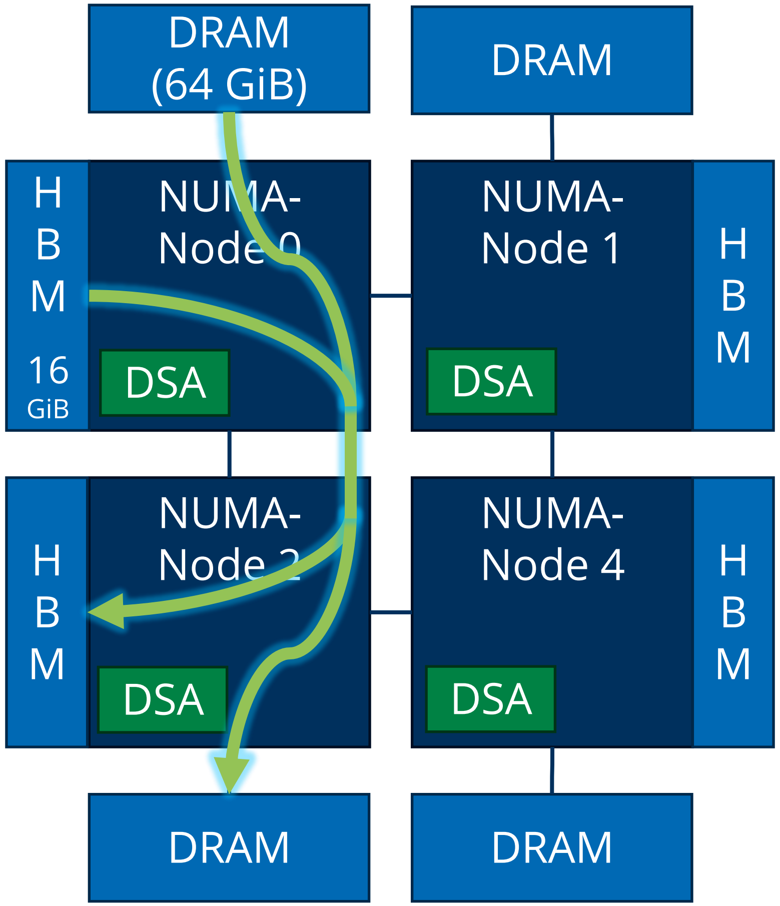

Funded by

Benchmark – CPU vs. DSA

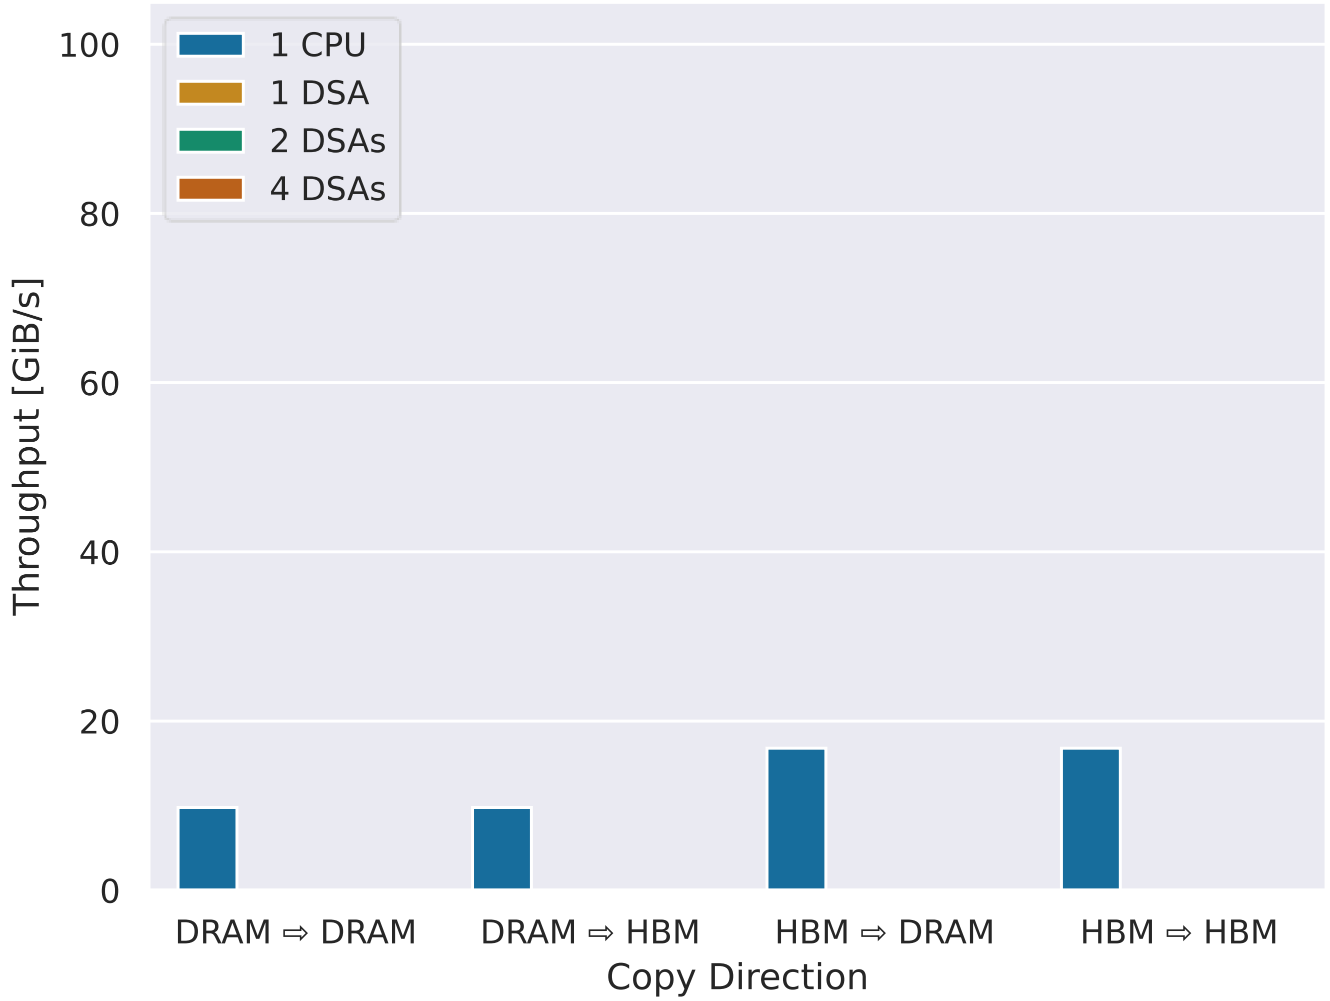

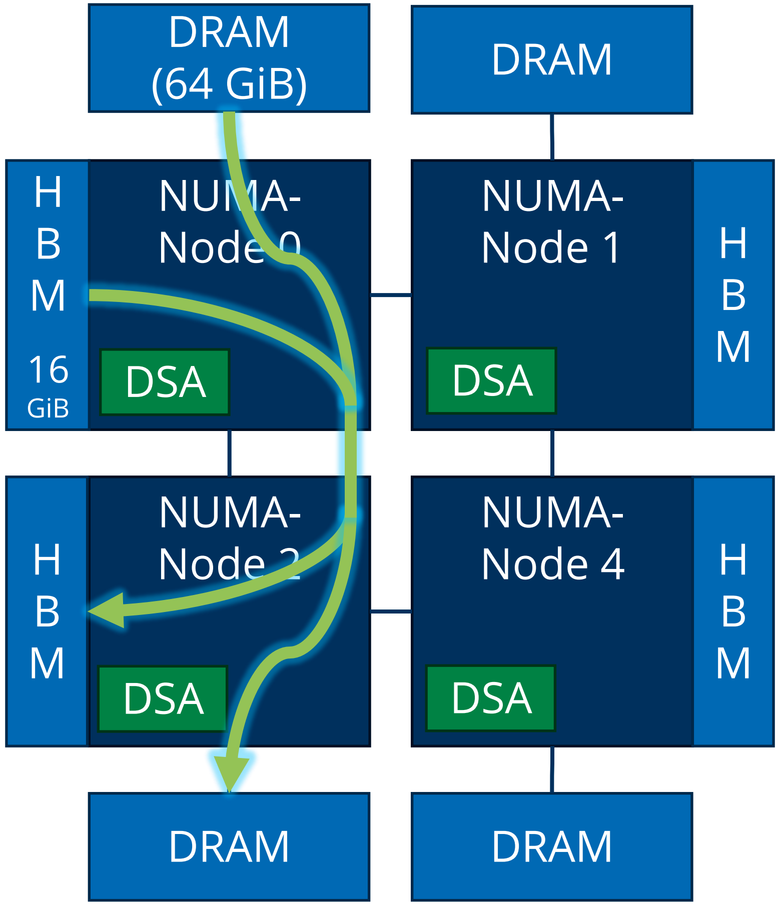

Funded by DFG
DRESDEN concept

Benchmark – CPU vs. DSA

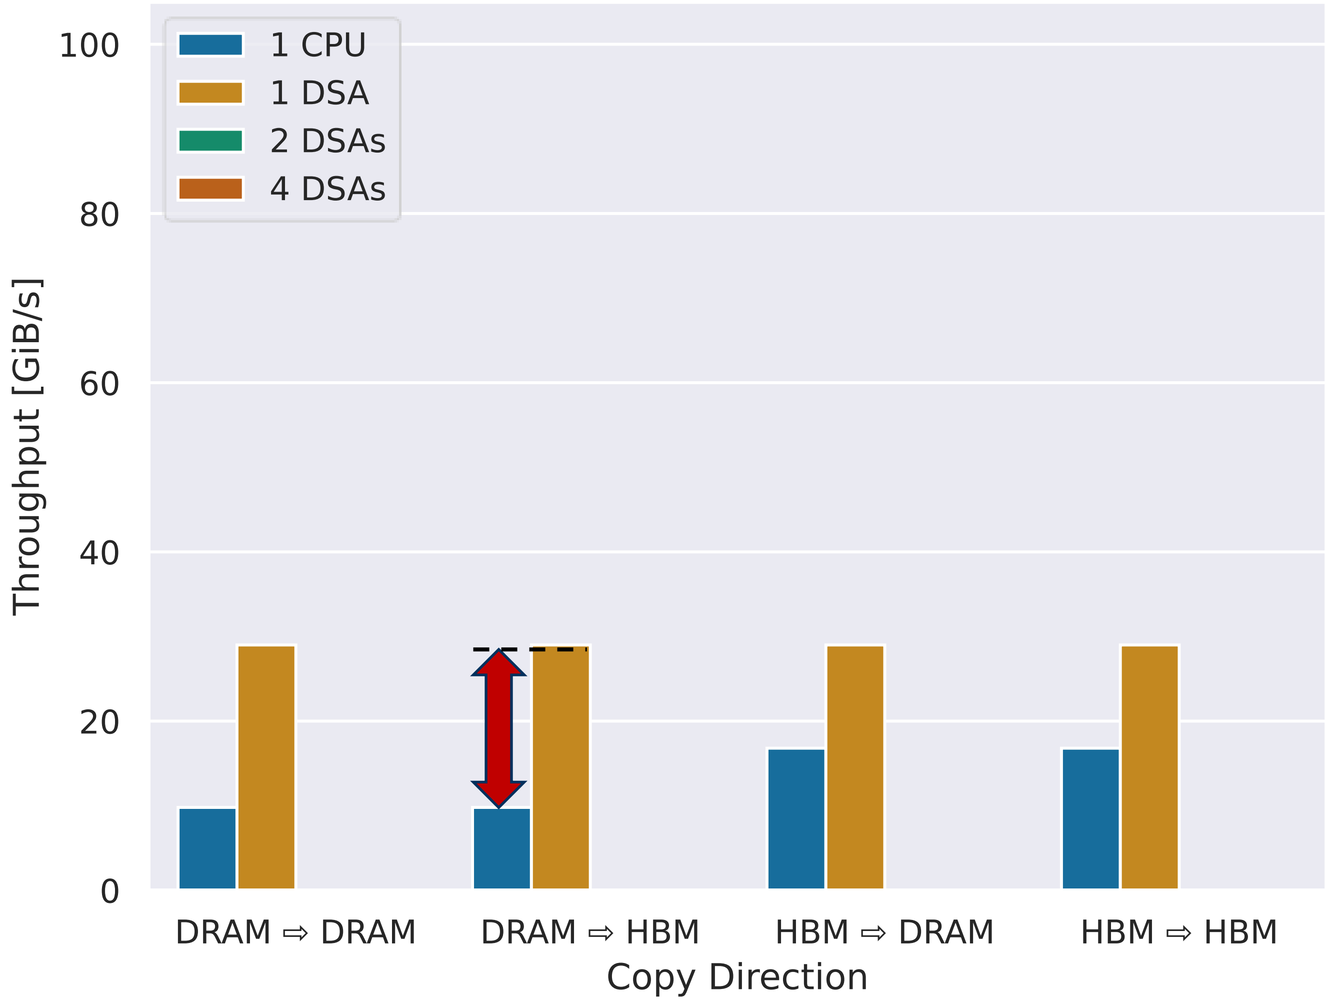

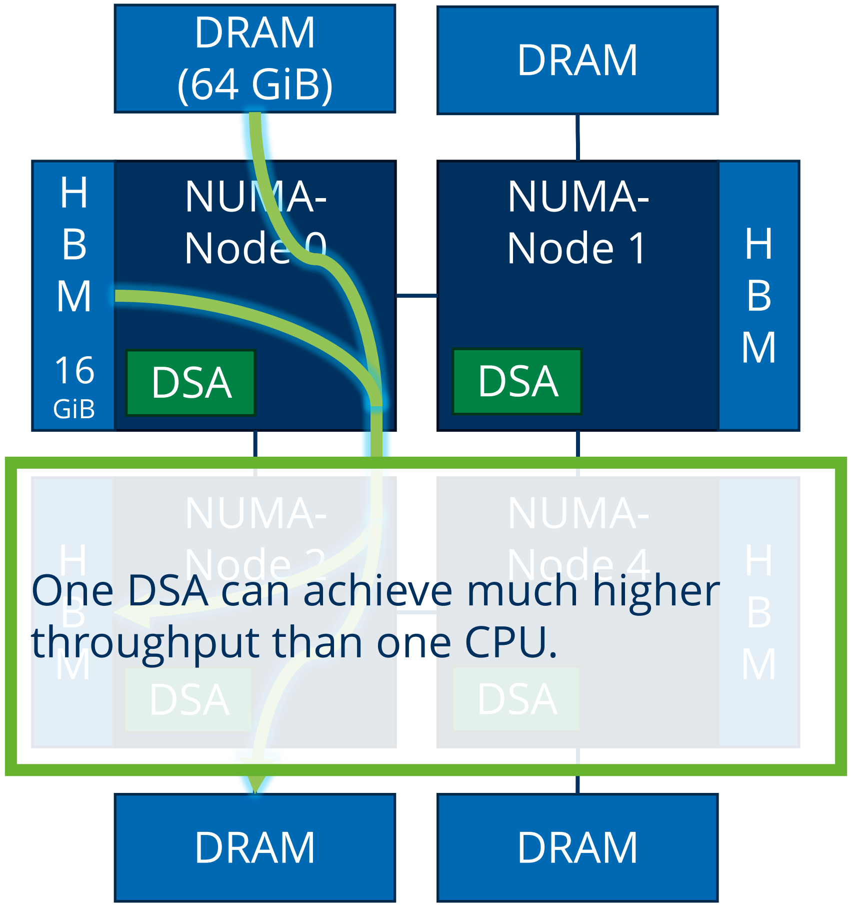

One DSA can achieve much higher throughput than one CPU.

Funded by DFG
DRESDEN concept

Benchmark – CPU vs. DSA

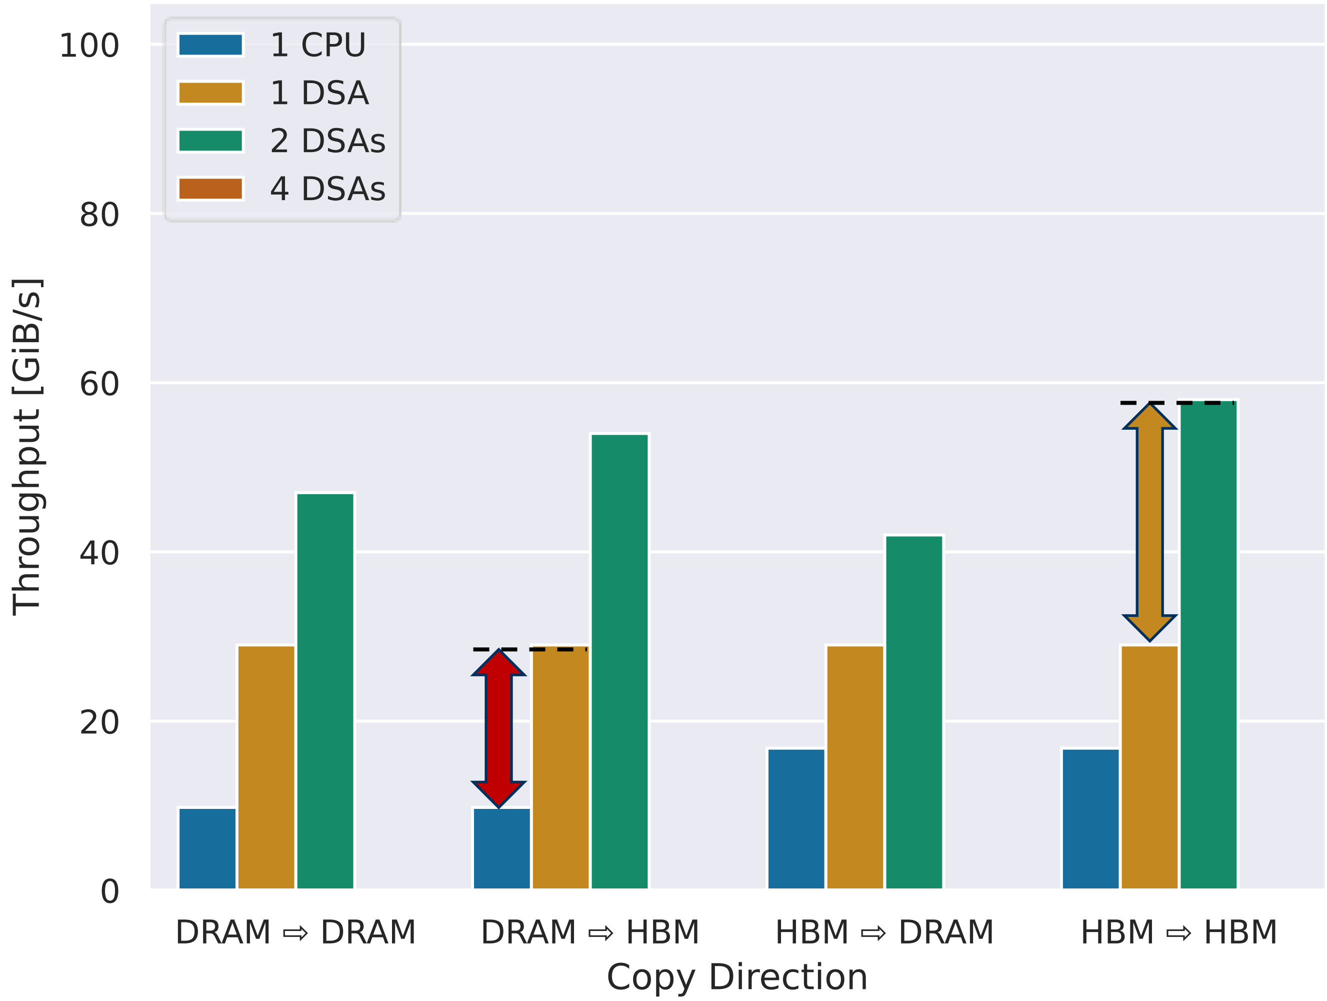

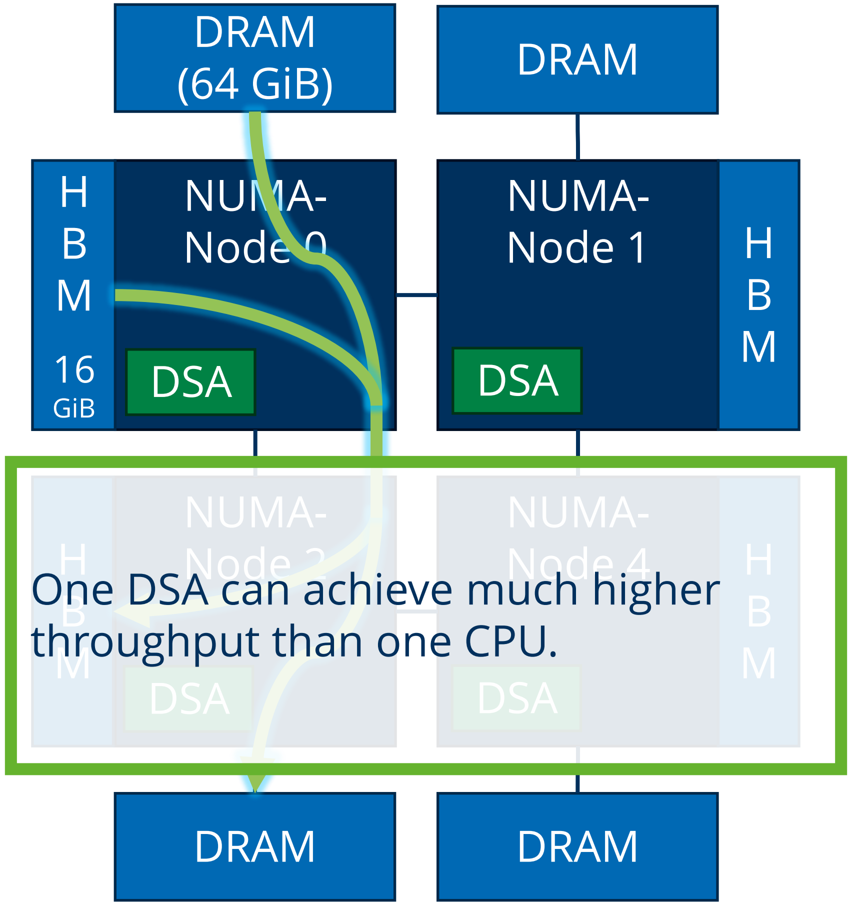

One DSA can achieve much higher throughput than one CPU.

Funded by DFG

DRESDEN concept

Benchmark – CPU vs. DSA

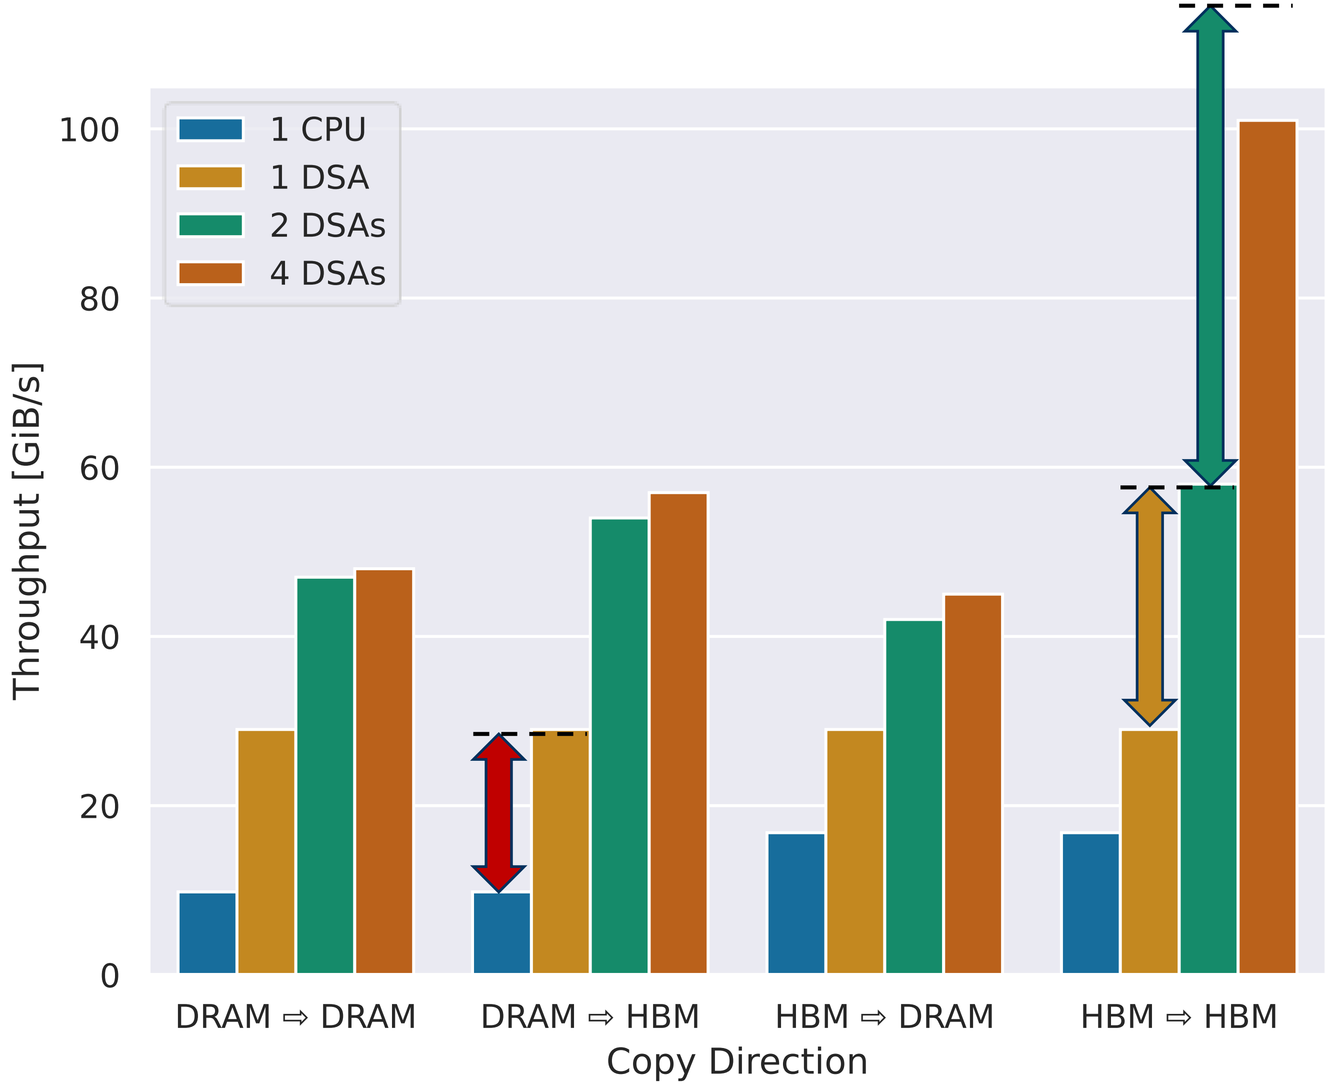

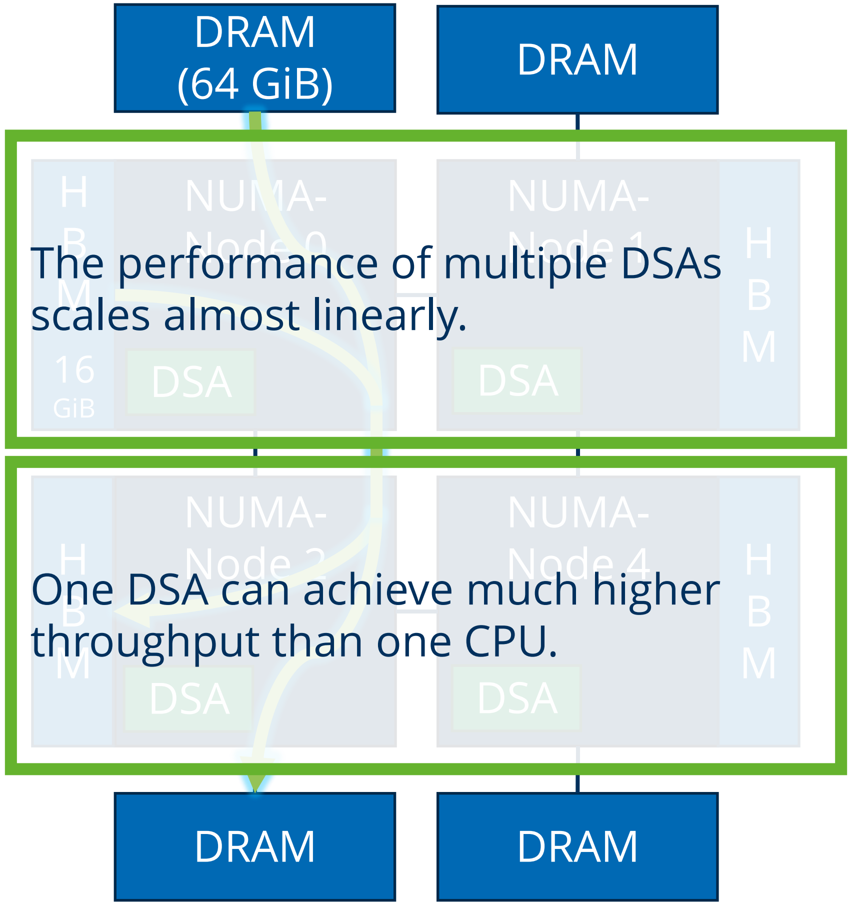

The performance of multiple DSAs scales almost linearly.

One DSA can achieve much higher throughput than one CPU.

Funded by DFG

DRESDEN concept

Benchmark – Interference between DSA and CPU

Compute-Intensive CPU Threads

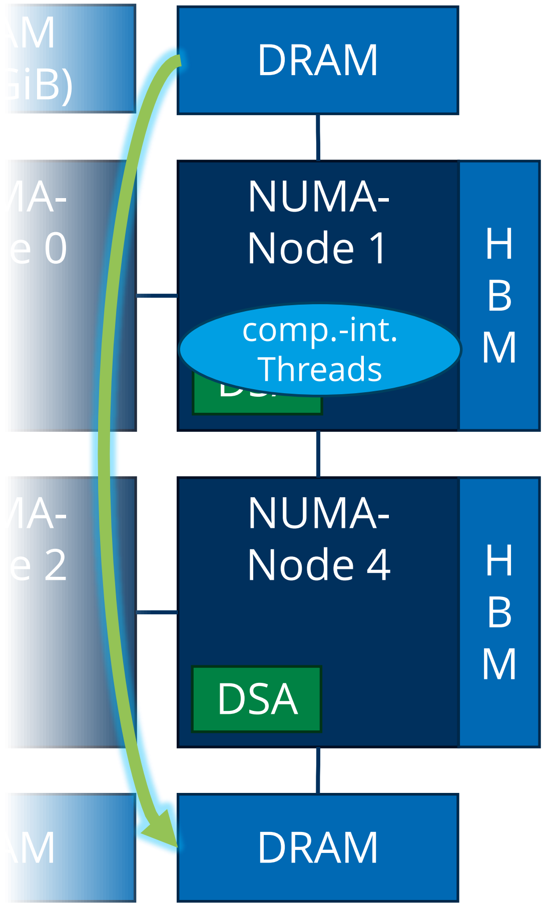

Funded by DFG

Benchmark – Interference between DSA and CPU

Compute-Intensive CPU Threads

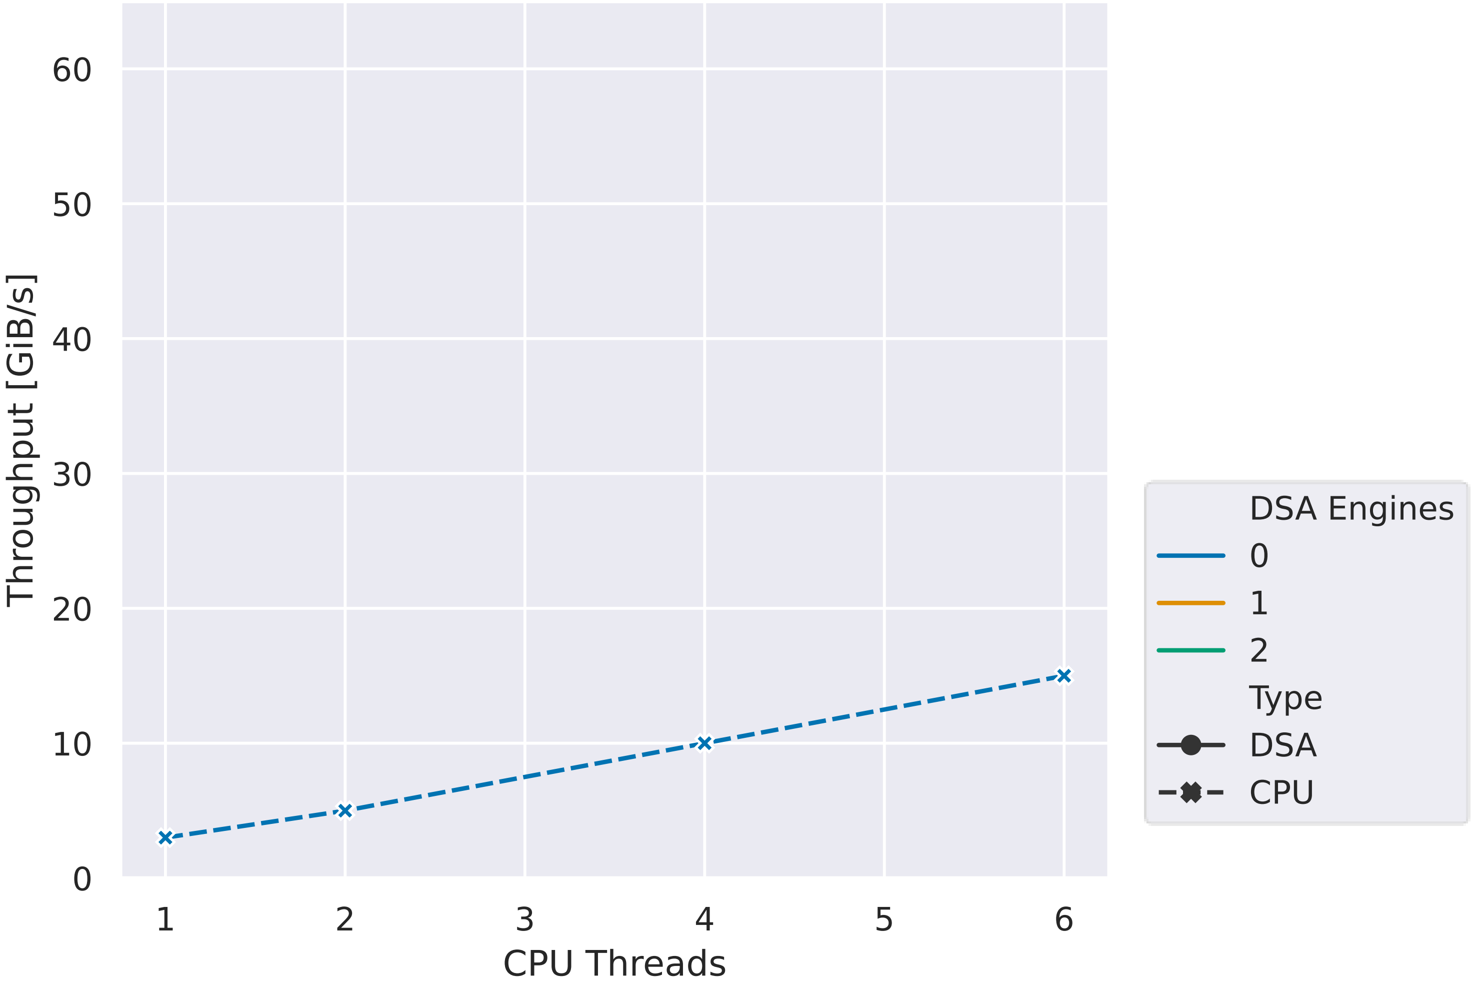

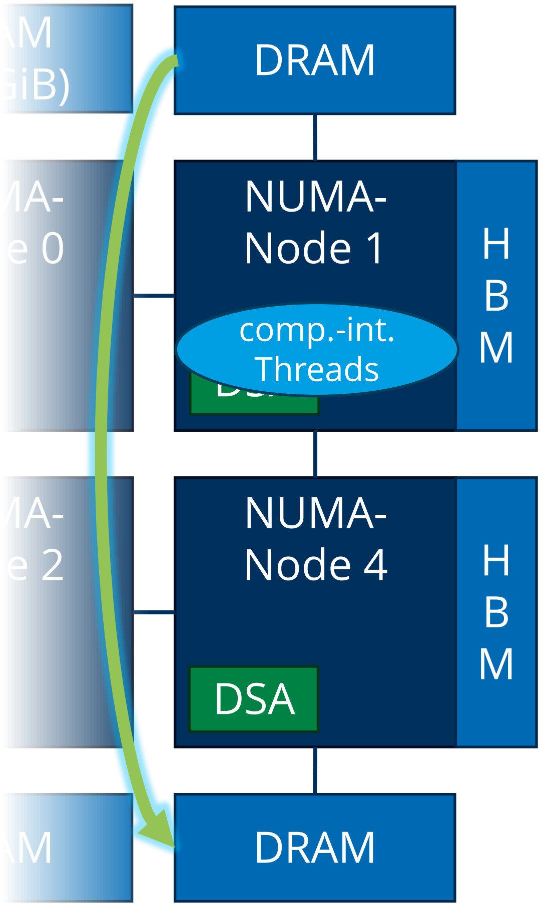

Funded by DFG

DRESDEN concept

Benchmark – Interference between DSA and CPU

Compute-Intensive CPU Threads

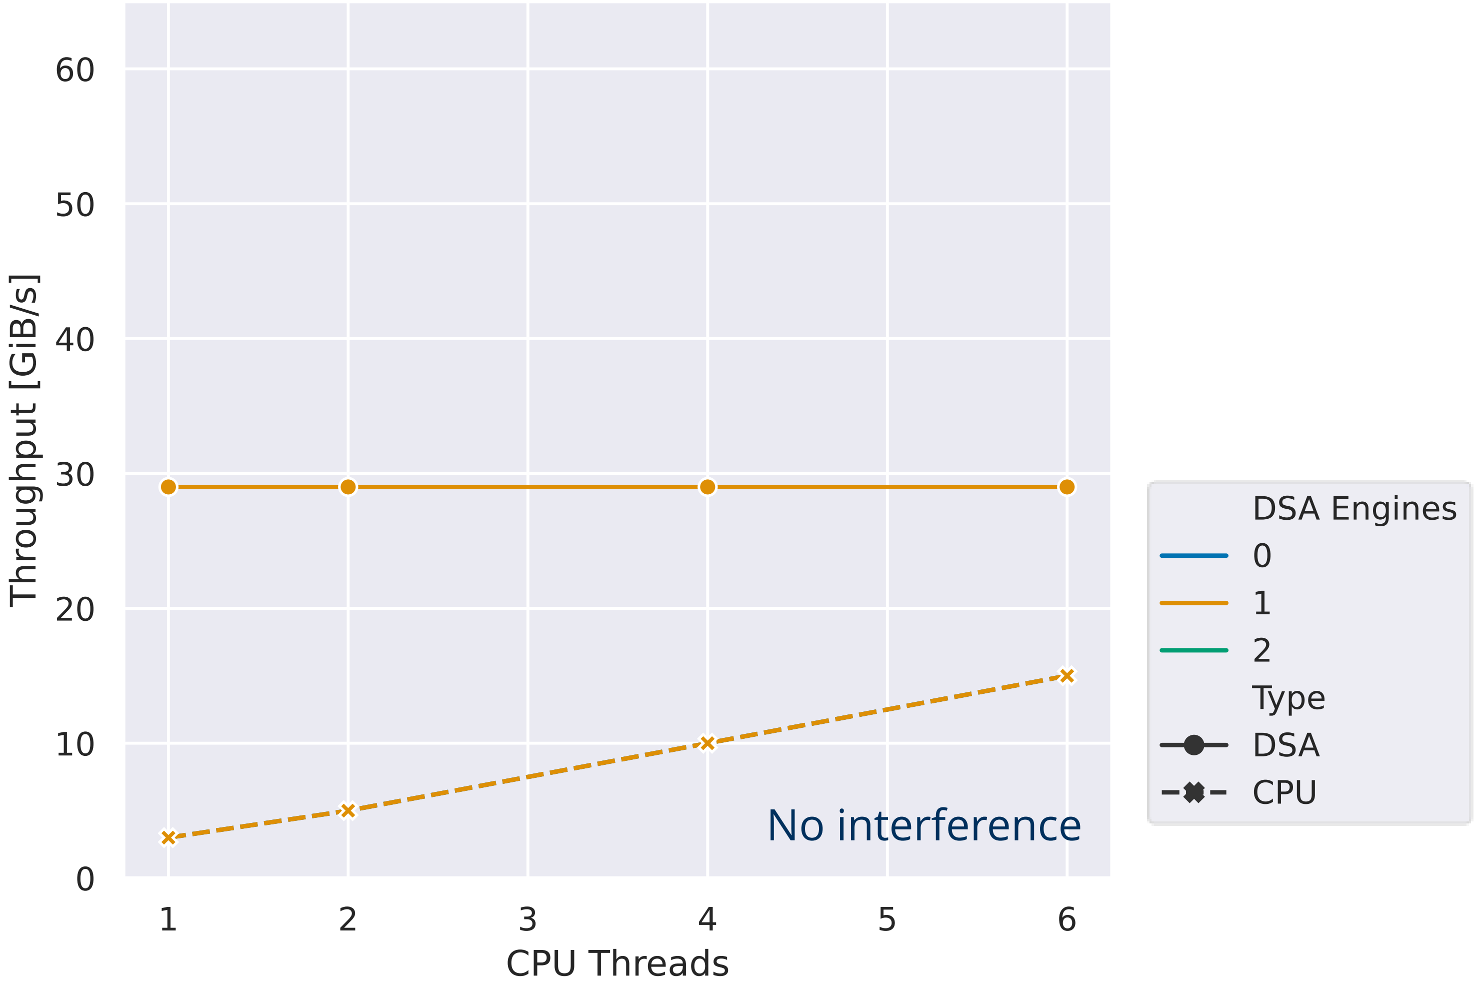

No interference

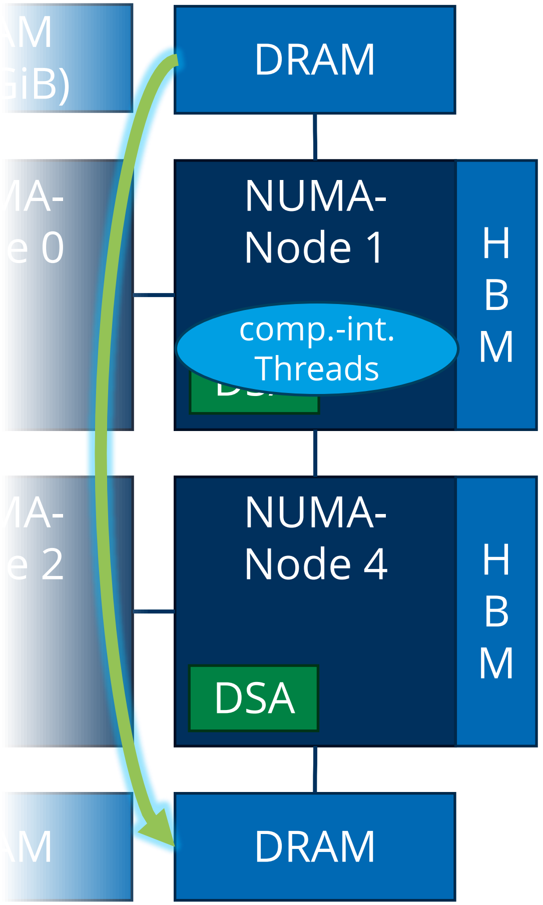

Funded by DFG

DRESDEN concept

Benchmark – Interference between DSA and CPU

Compute-Intensive CPU Threads

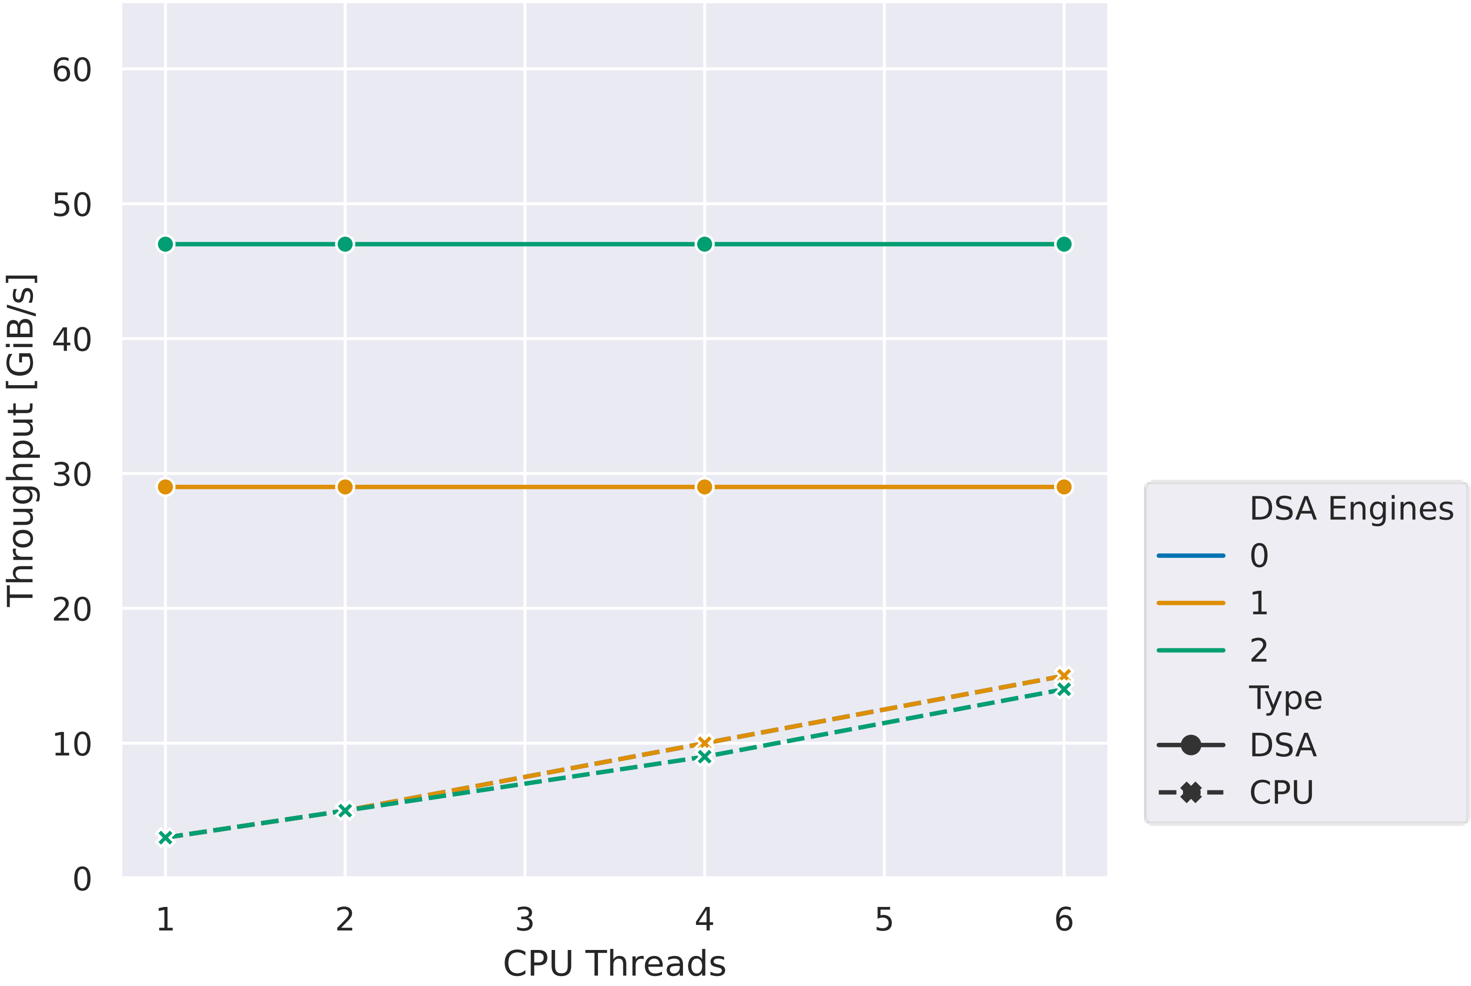

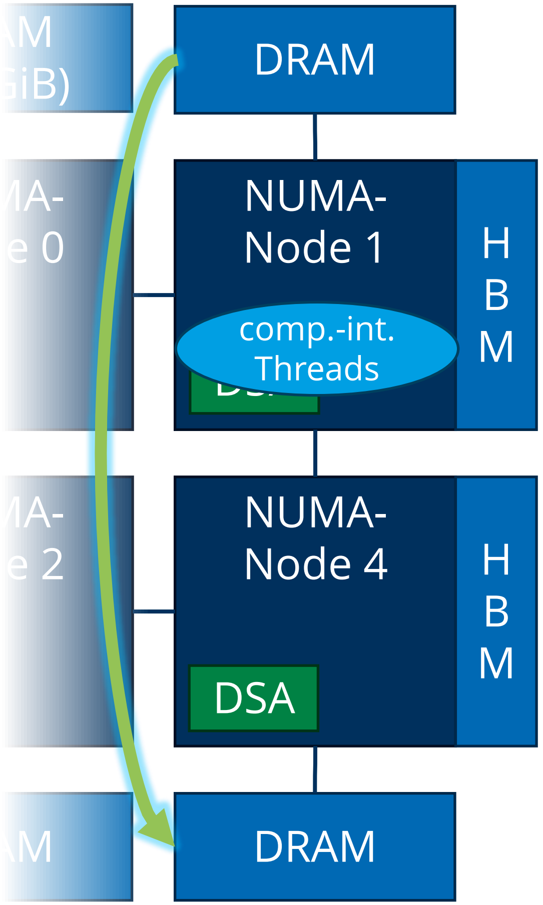

Funded by DFG
DRESDEN concept

Benchmark – Interference between DSA and CPU

Compute-Intensive CPU Threads

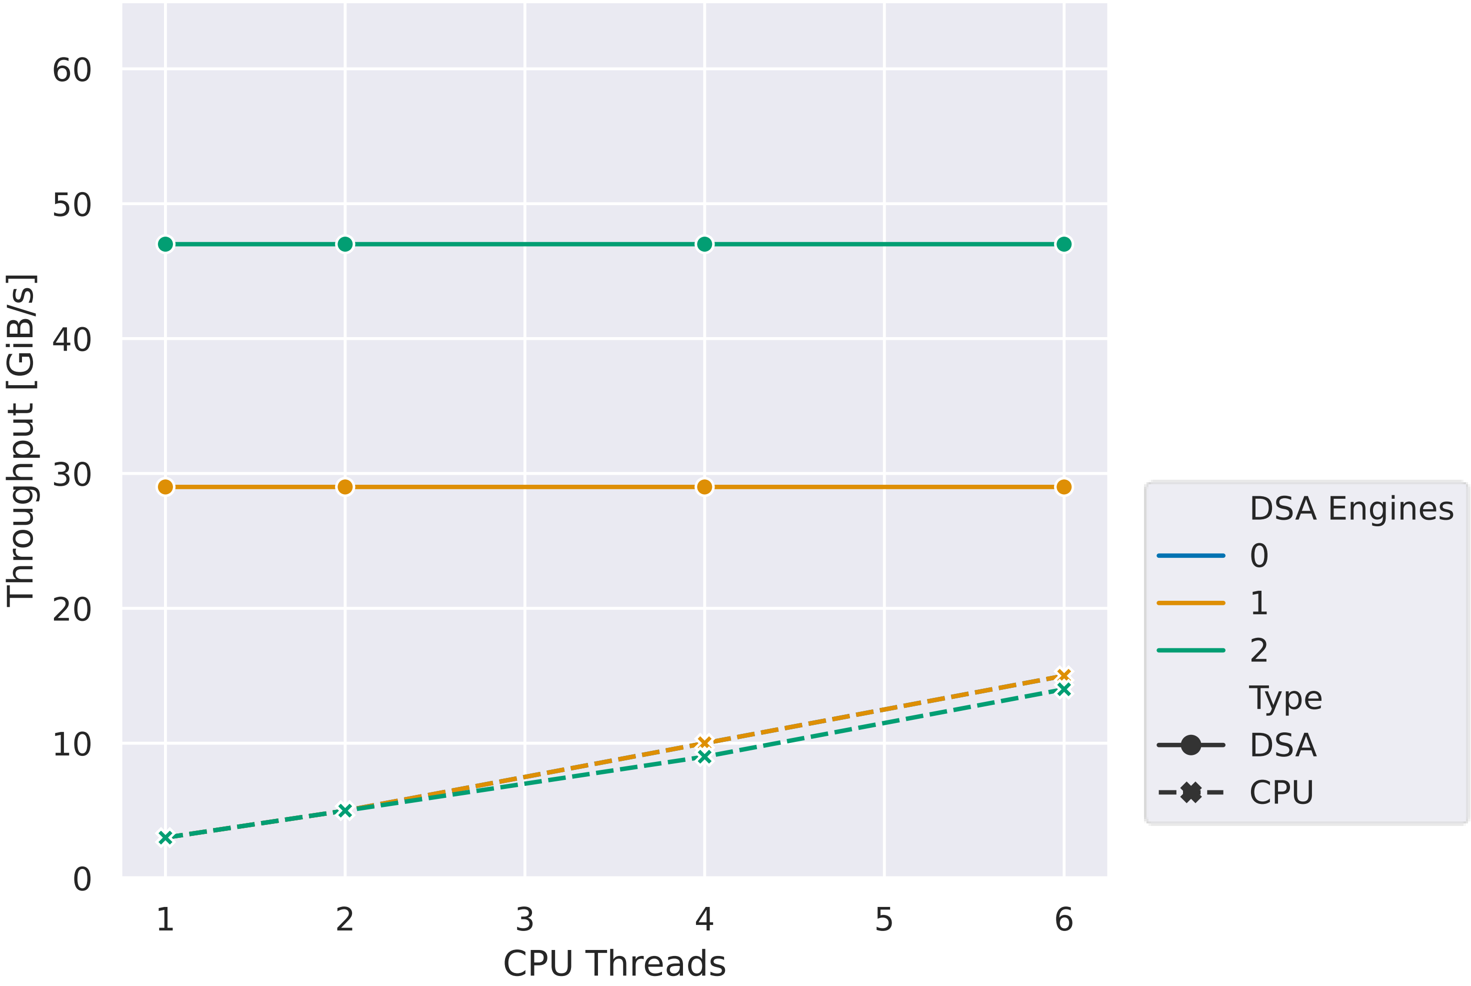

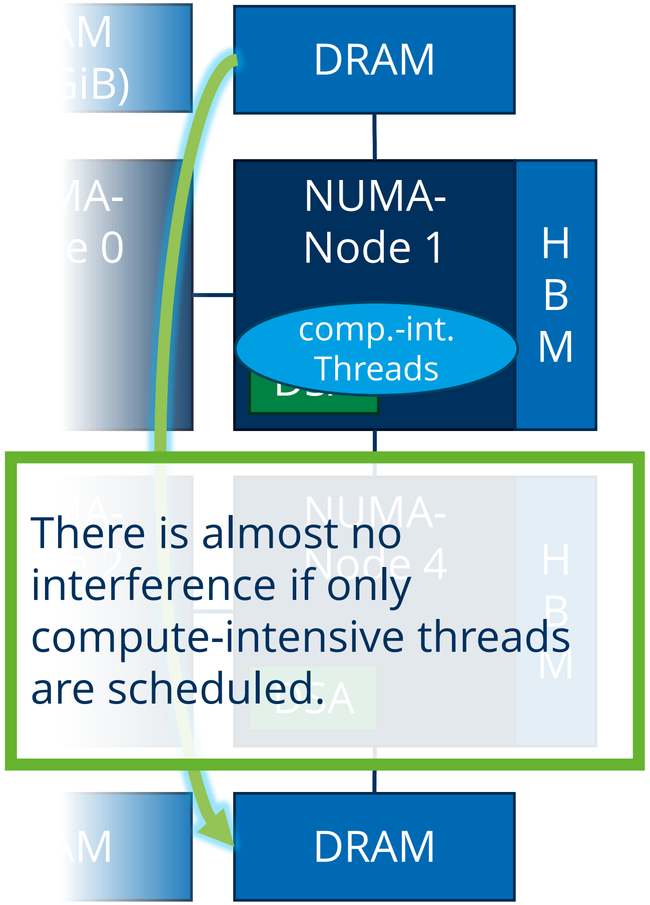

There is almost no interference if only compute-intensive threads are scheduled.

Funded by DFG
DRESDEN concept

Benchmark – Interference between DSA and CPU

Data-Intensive CPU Threads (Vectorized, AVX-512)

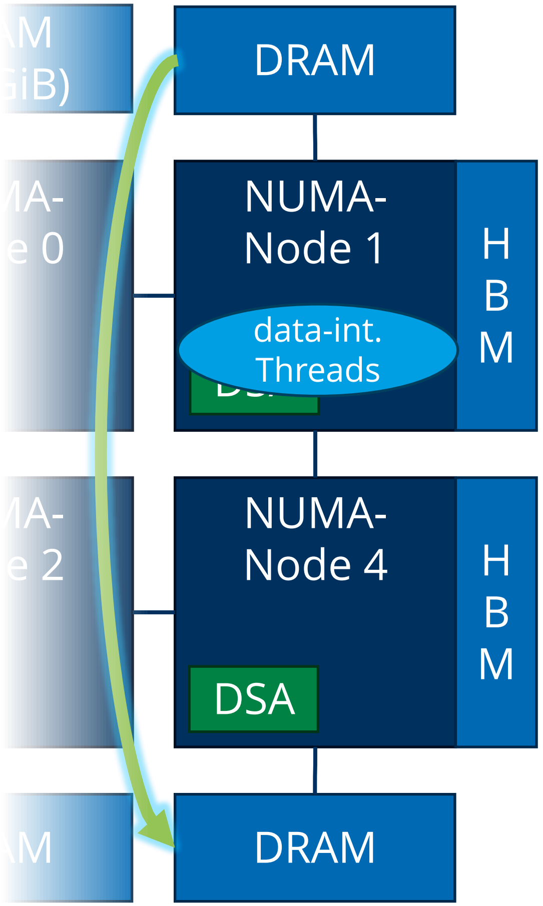

Funded by DFG

Benchmark – Interference between DSA and CPU

Data-Intensive CPU Threads (Vectorized, AVX-512)

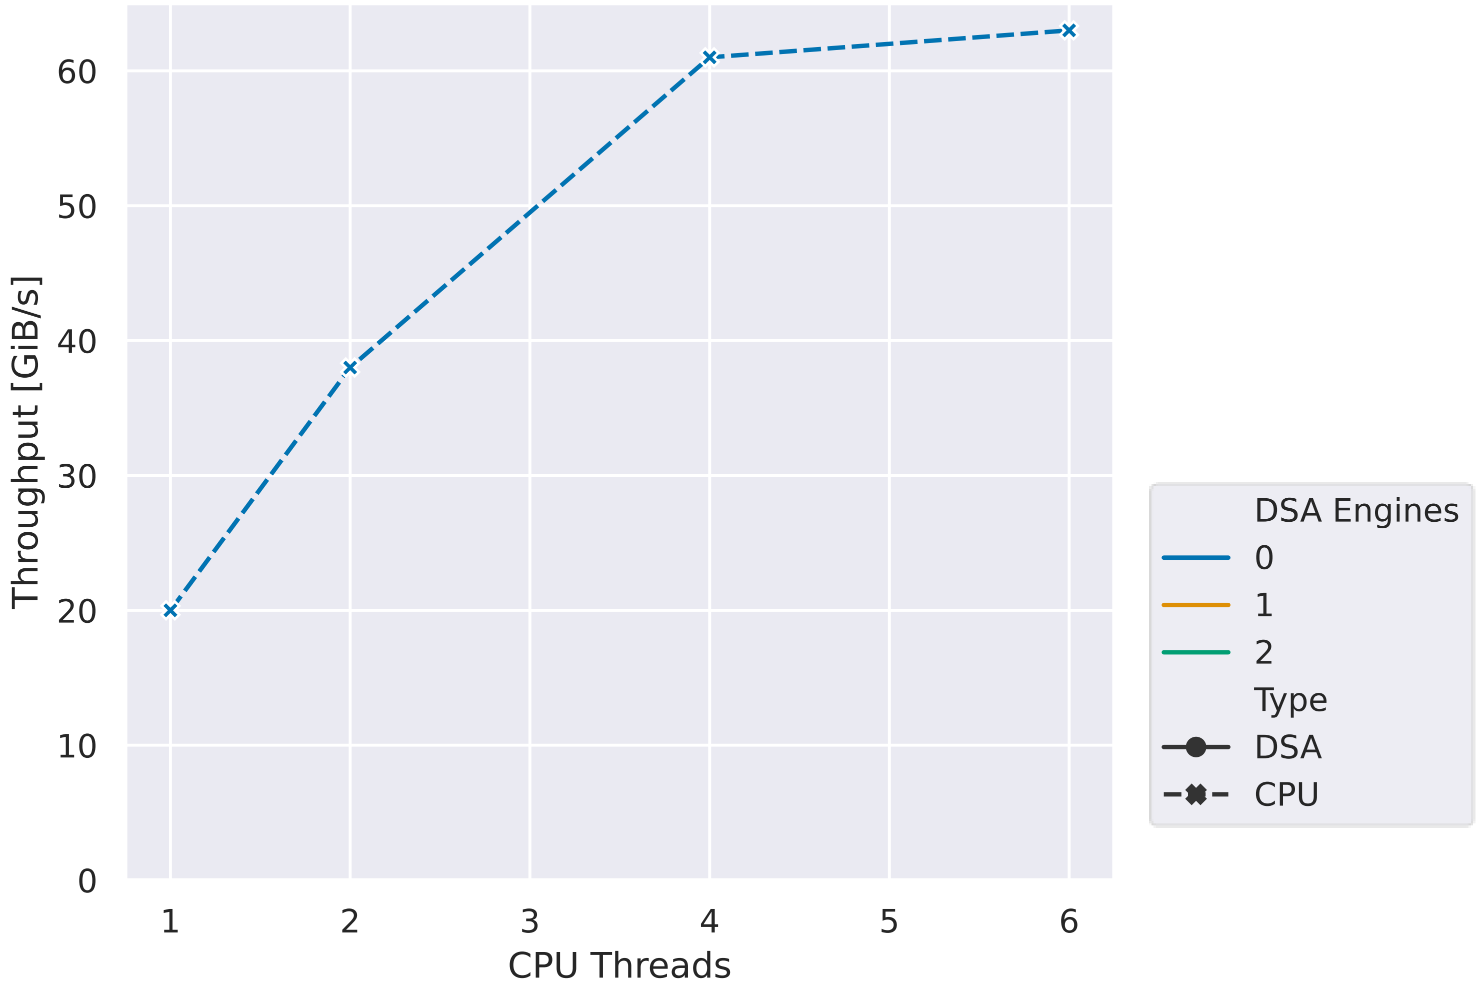

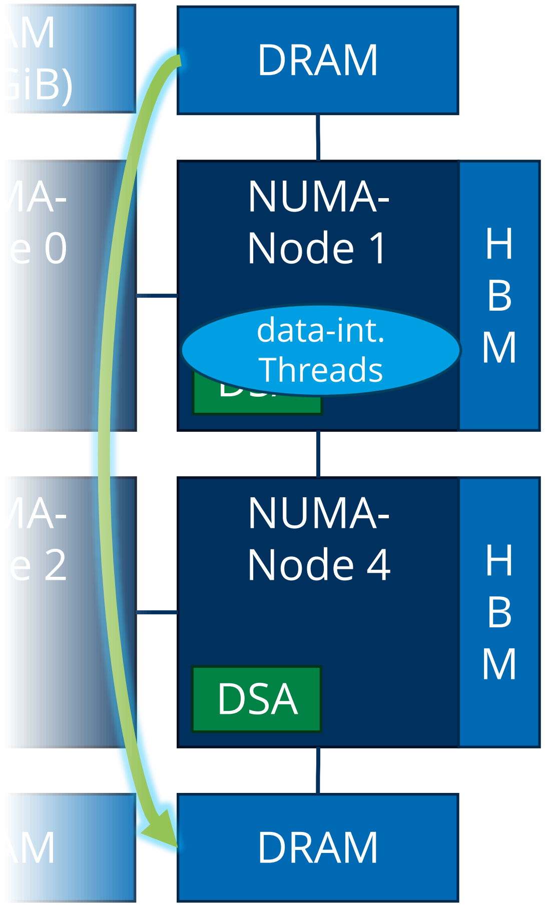

Funded by

Benchmark – Interference between DSA and CPU

Data-Intensive CPU Threads (Vectorized, AVX-512)

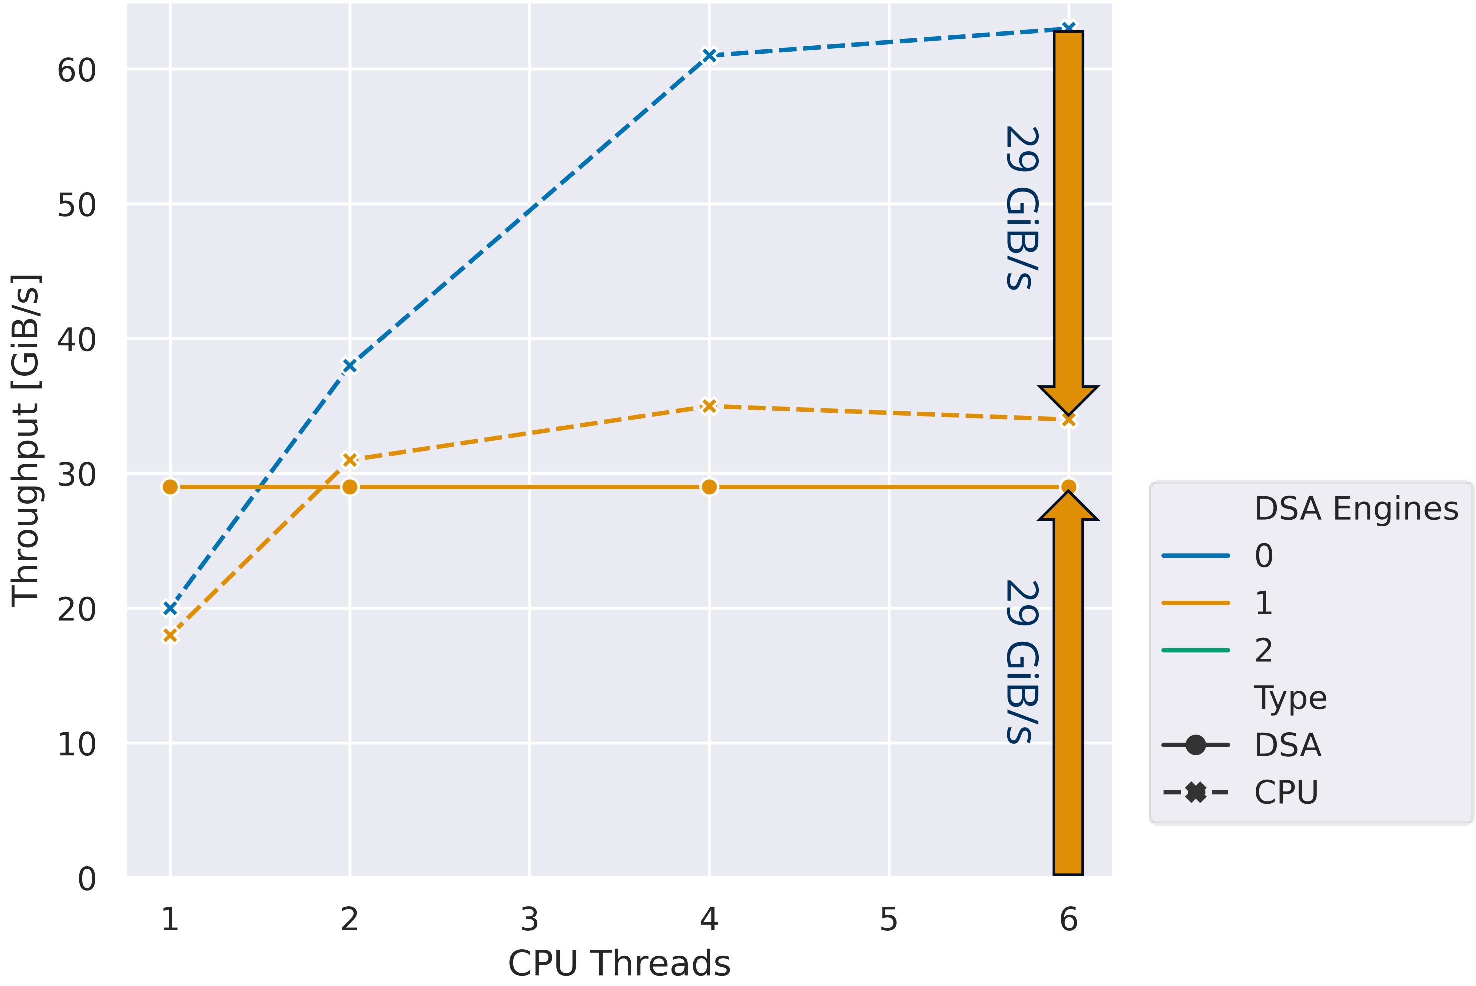

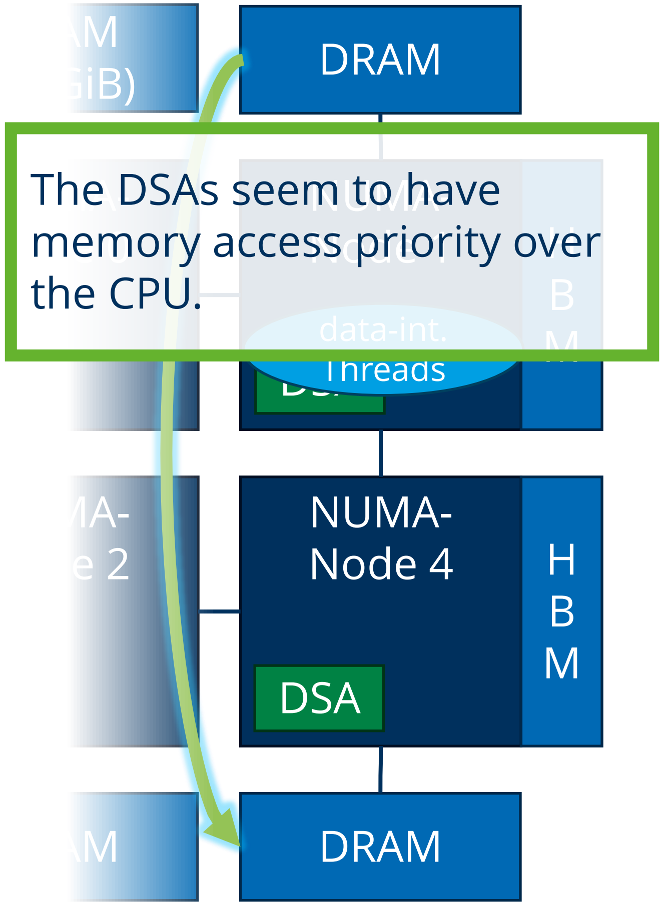

The DSAs seem to have memory access priority over the CPU.

Funded by

Benchmark – Interference between DSA and CPU

Data-Intensive CPU Threads (Vectorized, AVX-512)

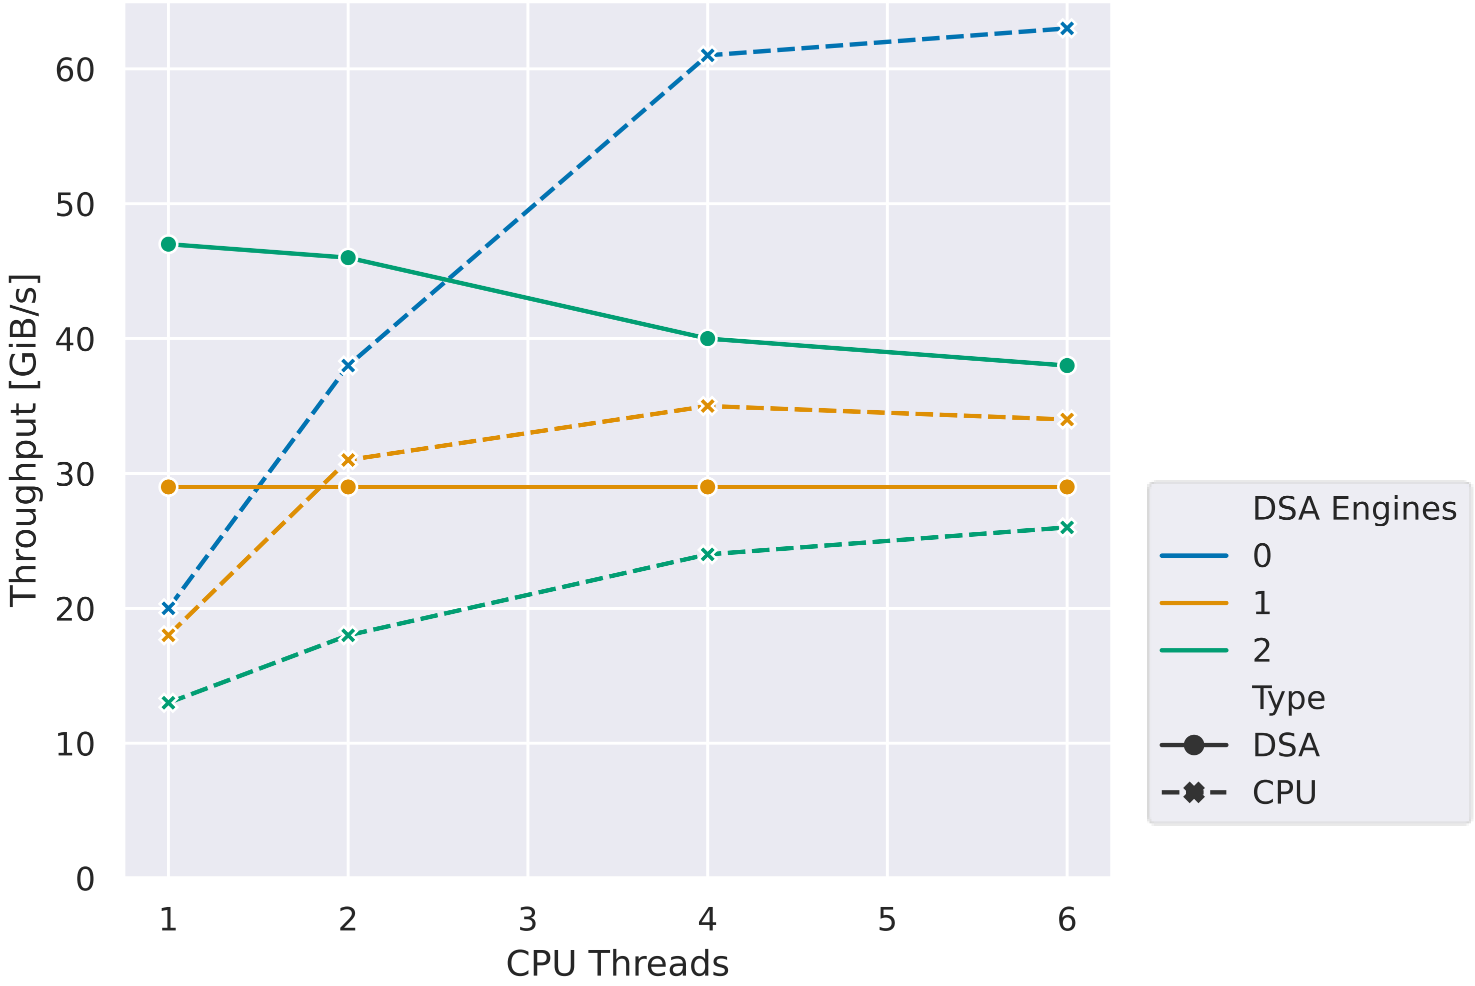

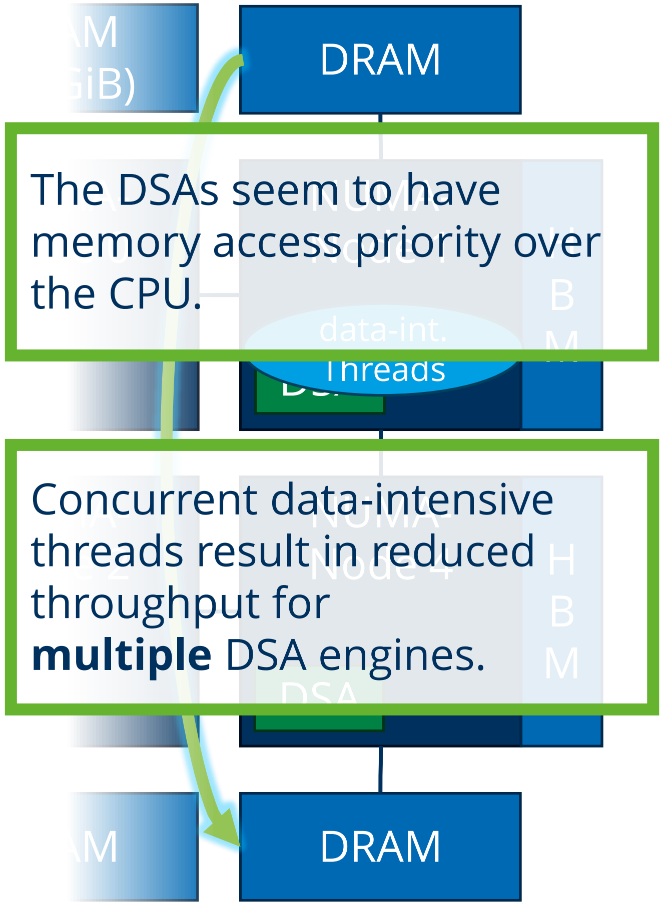

The DSAs seem to have memory access priority over the CPU.

Concurrent data-intensive threads result in reduced throughput for **multiple** DSA engines.

Funded by DFG
DRESDEN concept

Benchmark – Interference between DSA and CPU

Data-Intensive CPU Threads on HBM

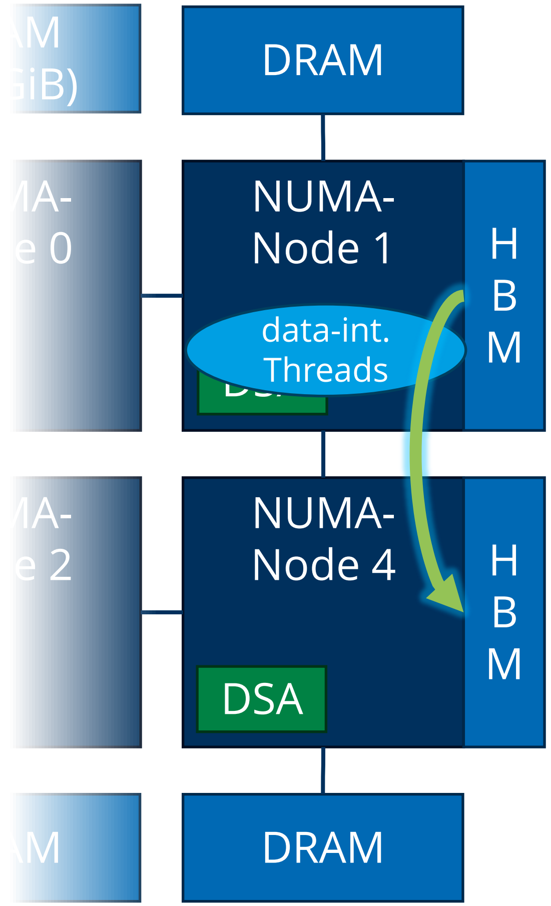

Benchmark – Interference between DSA and CPU

Data-Intensive CPU Threads on HBM

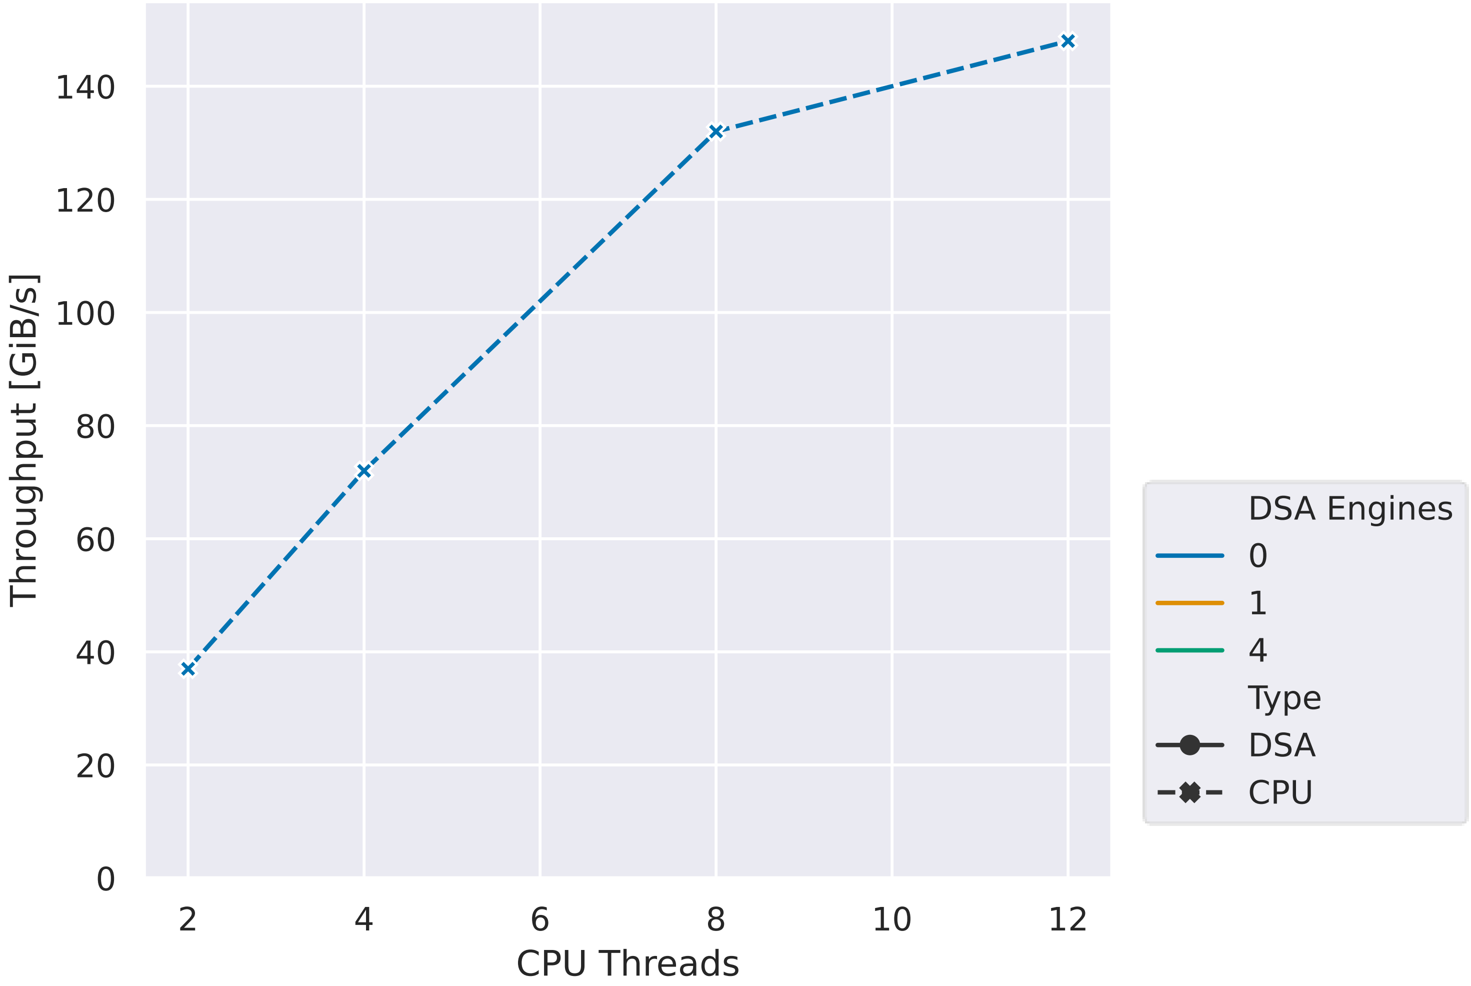

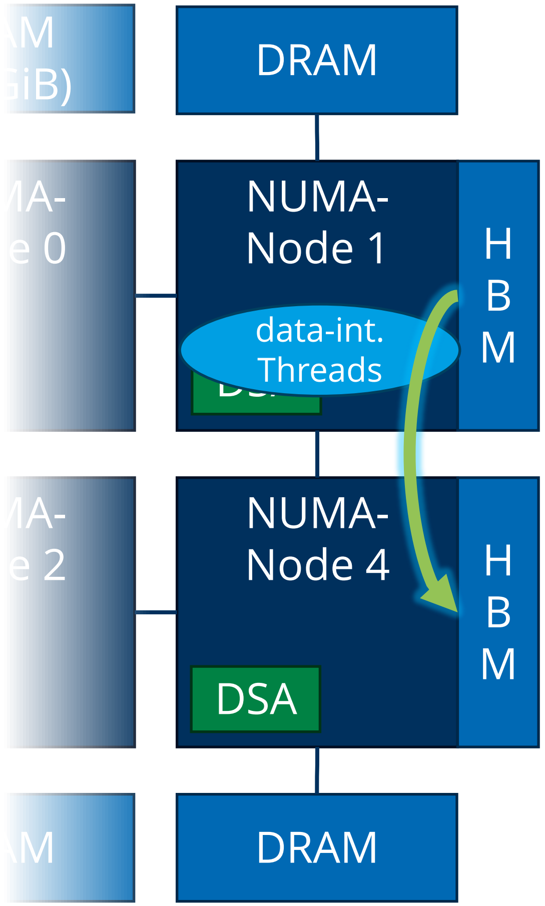

Funded by DFG
DRESDEN concept

Benchmark – Interference between DSA and CPU

Data-Intensive CPU Threads on HBM

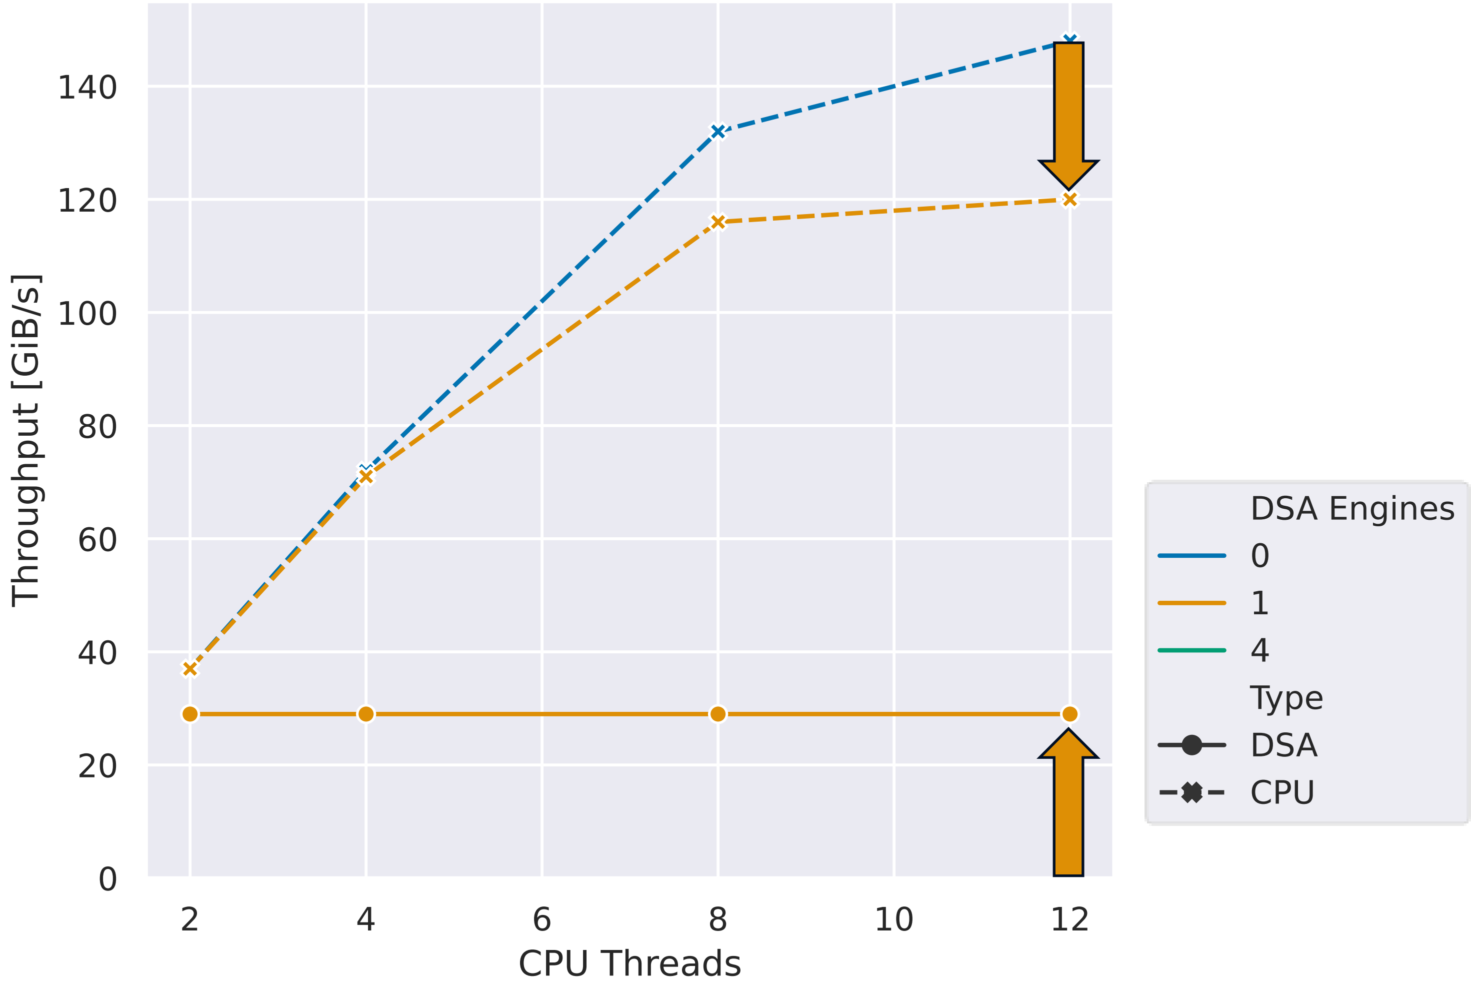

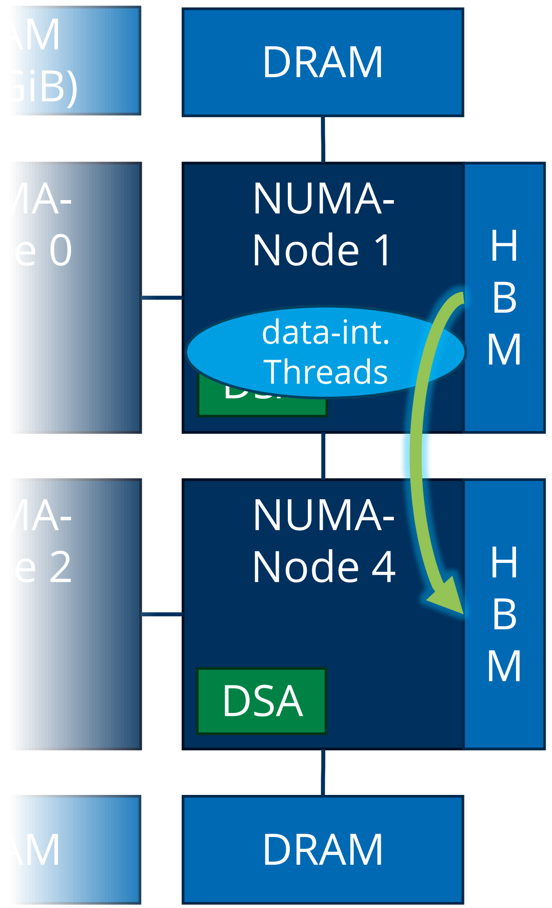

Funded by DFG
DRESDEN concept

Benchmark – Interference between DSA and CPU

Data-Intensive CPU Threads on HBM

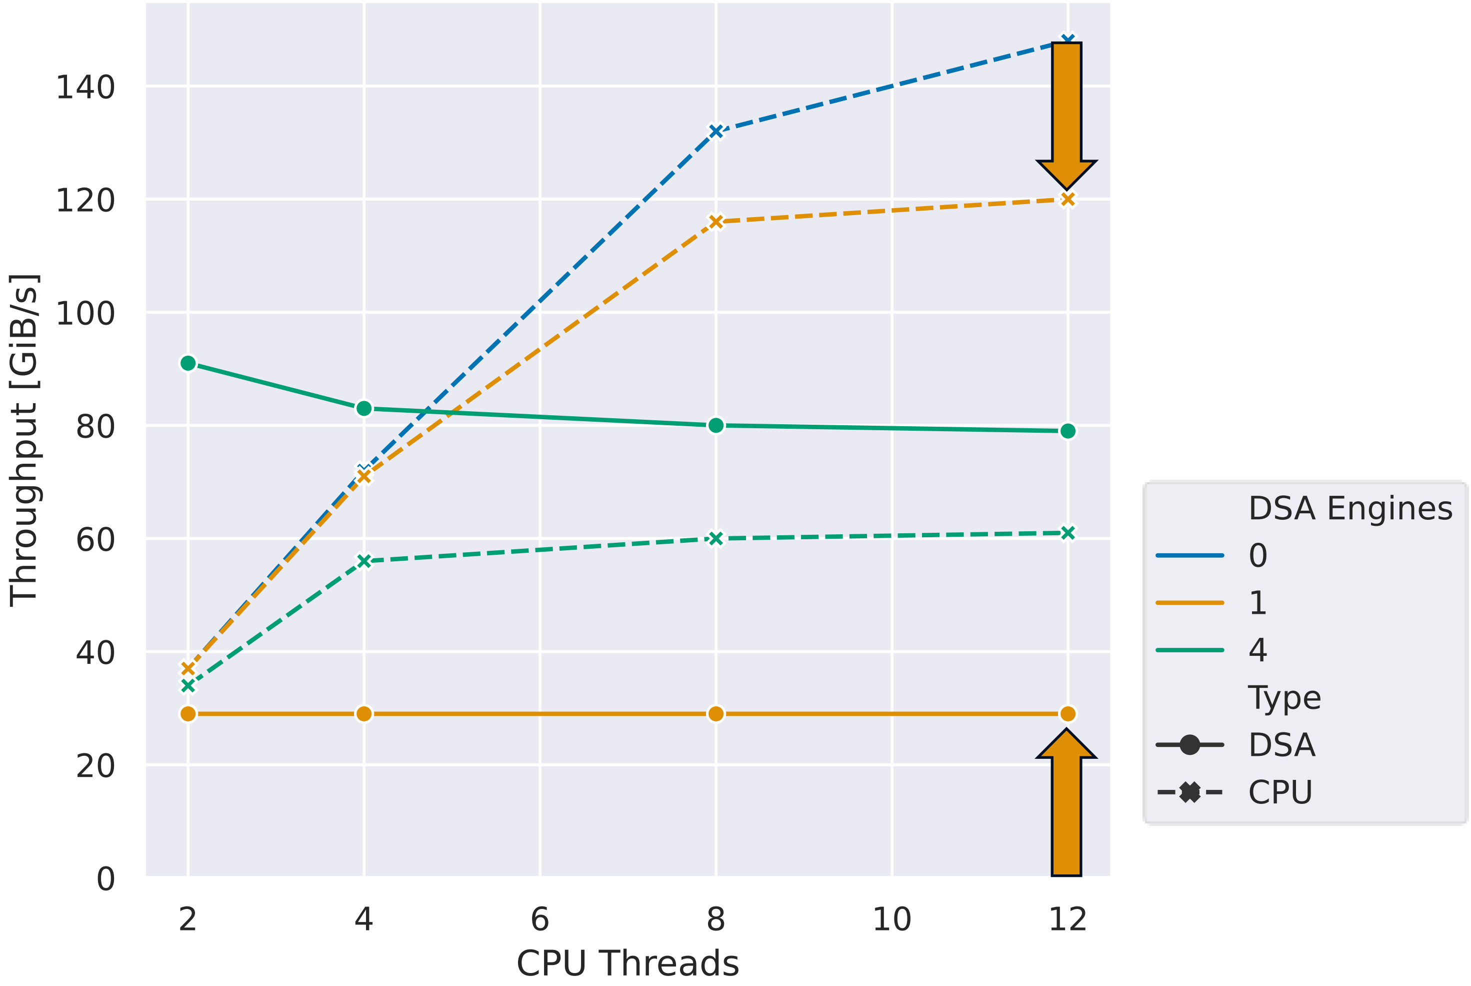

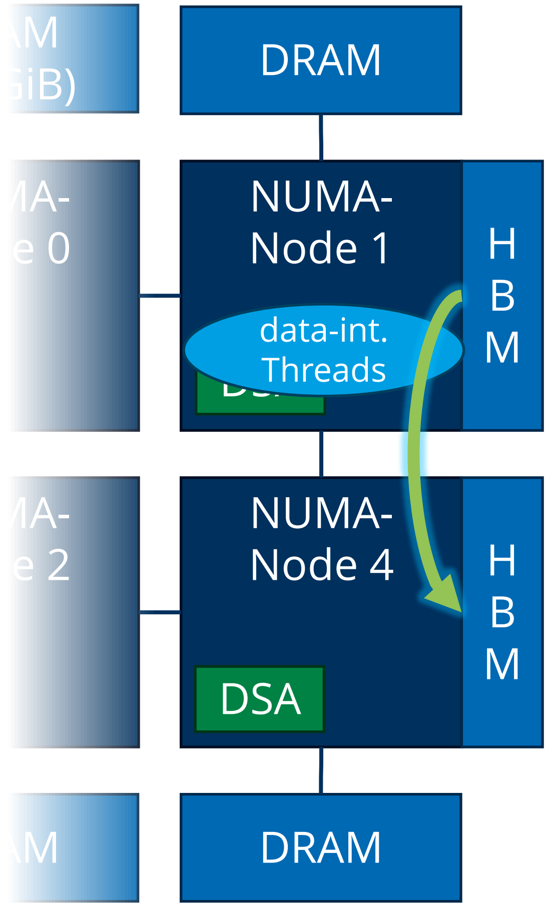

Funded by
DFG
DRESDEN concept

Benchmark – Interference between DSA and CPU

Data-Intensive CPU Threads on HBM

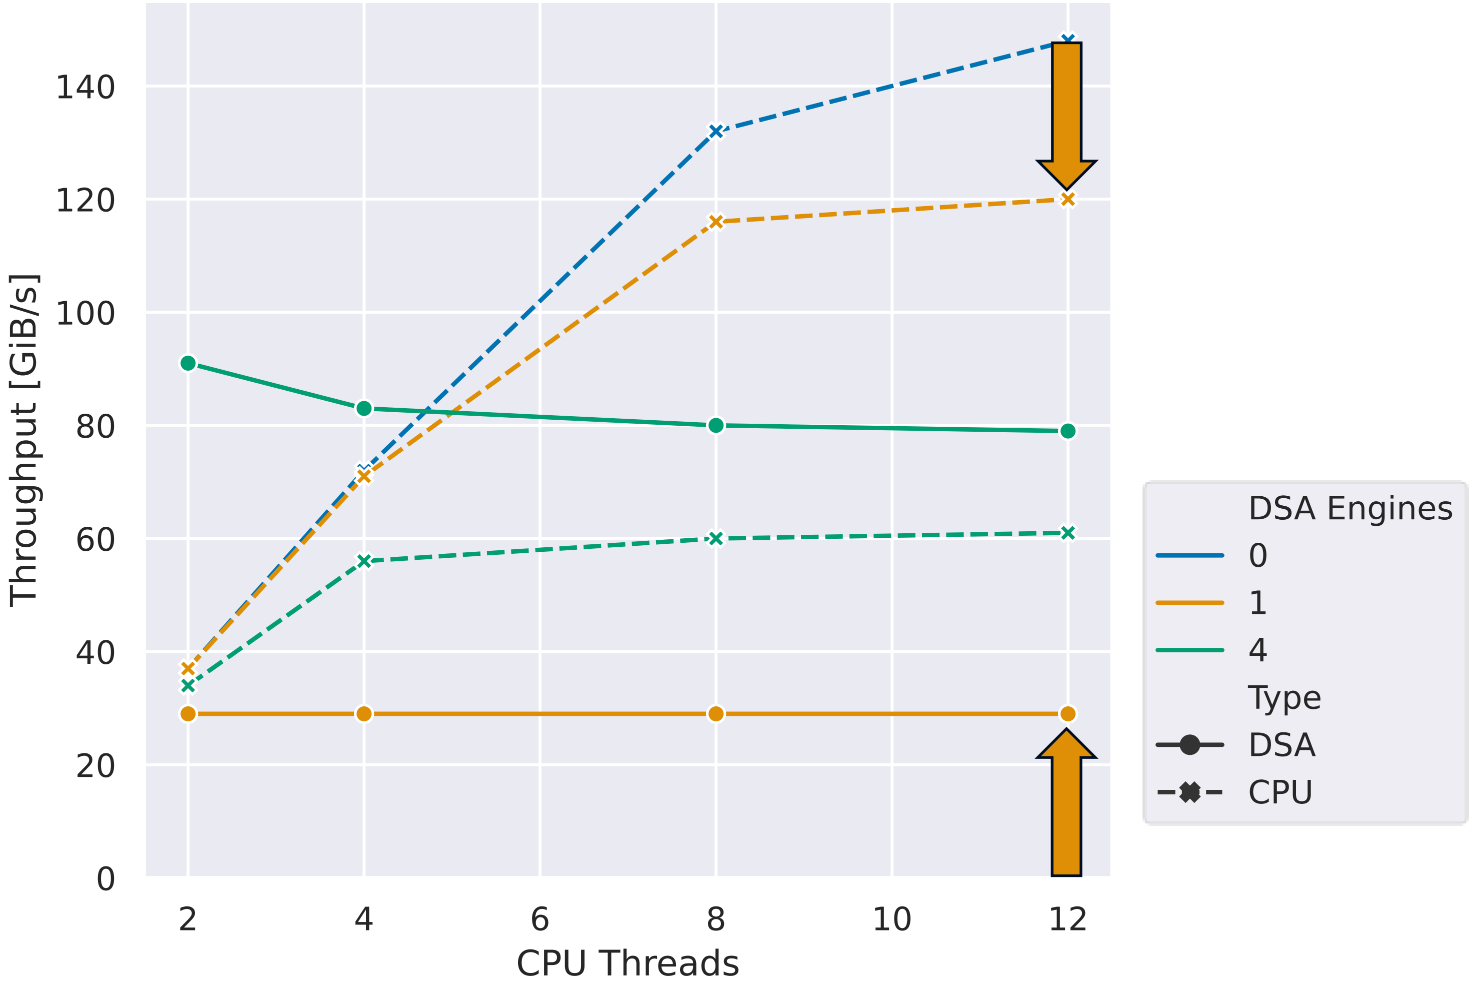

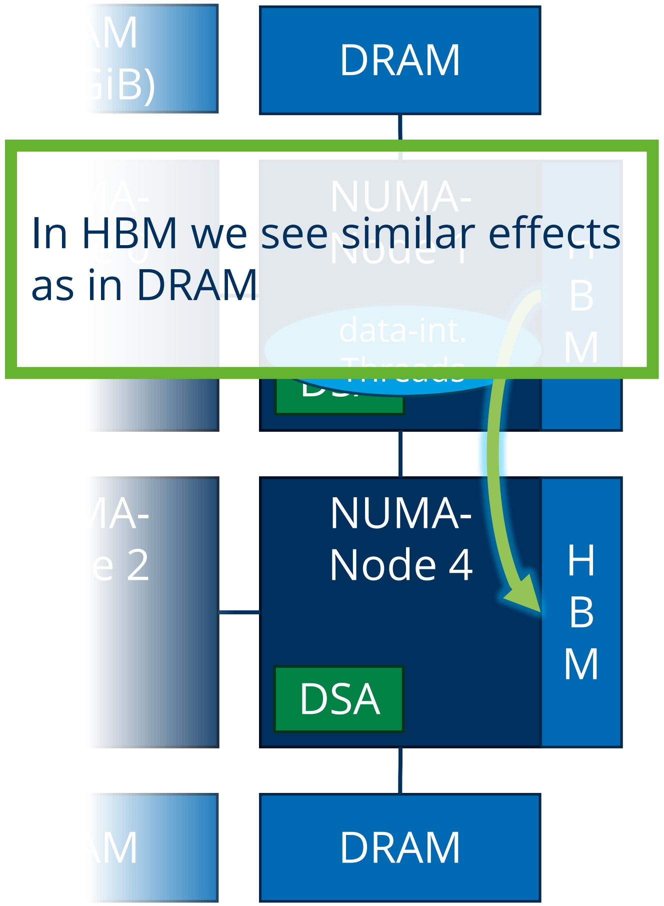

In HBM we see similar effects as in DRAM

Funded by DFG
DRESDEN concept

Benchmark – Inter-Socket Data Transfer with DSA

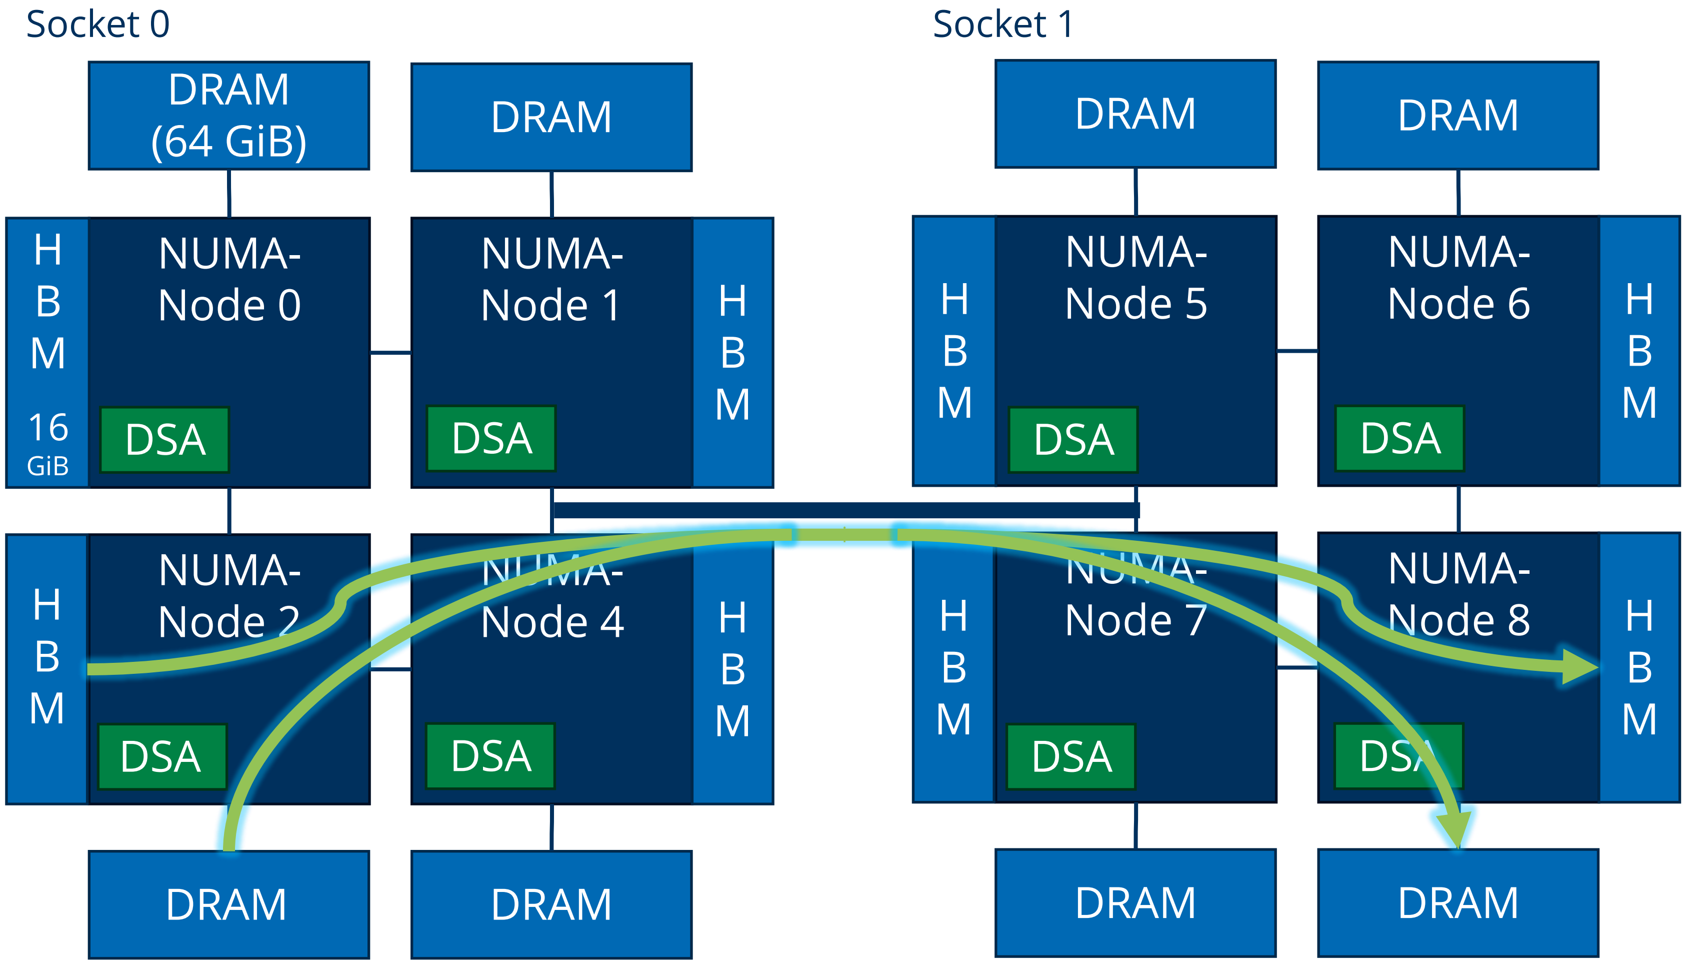

Funded by DFG

Benchmark – Inter-Socket Data Transfer with DSA

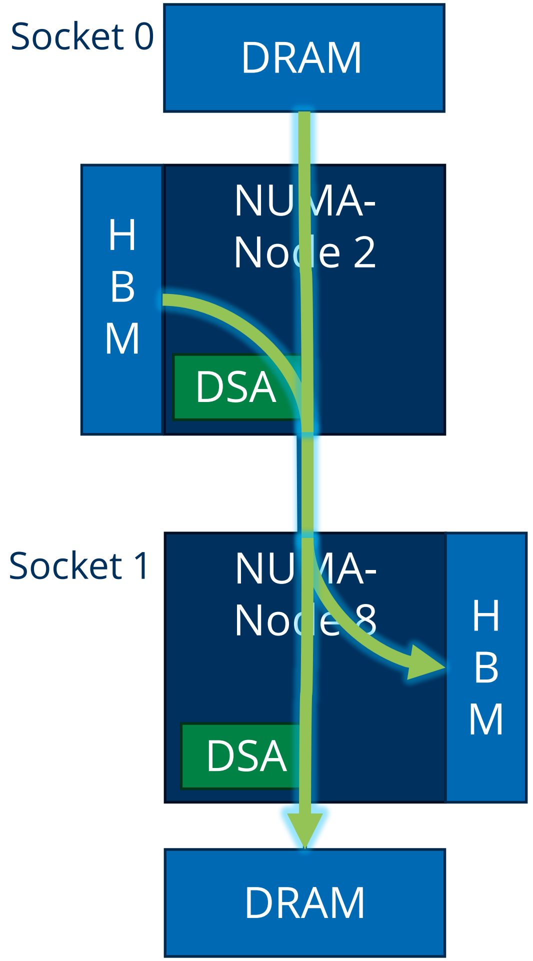

Funded by

Benchmark – Inter-Socket Data Transfer with DSA

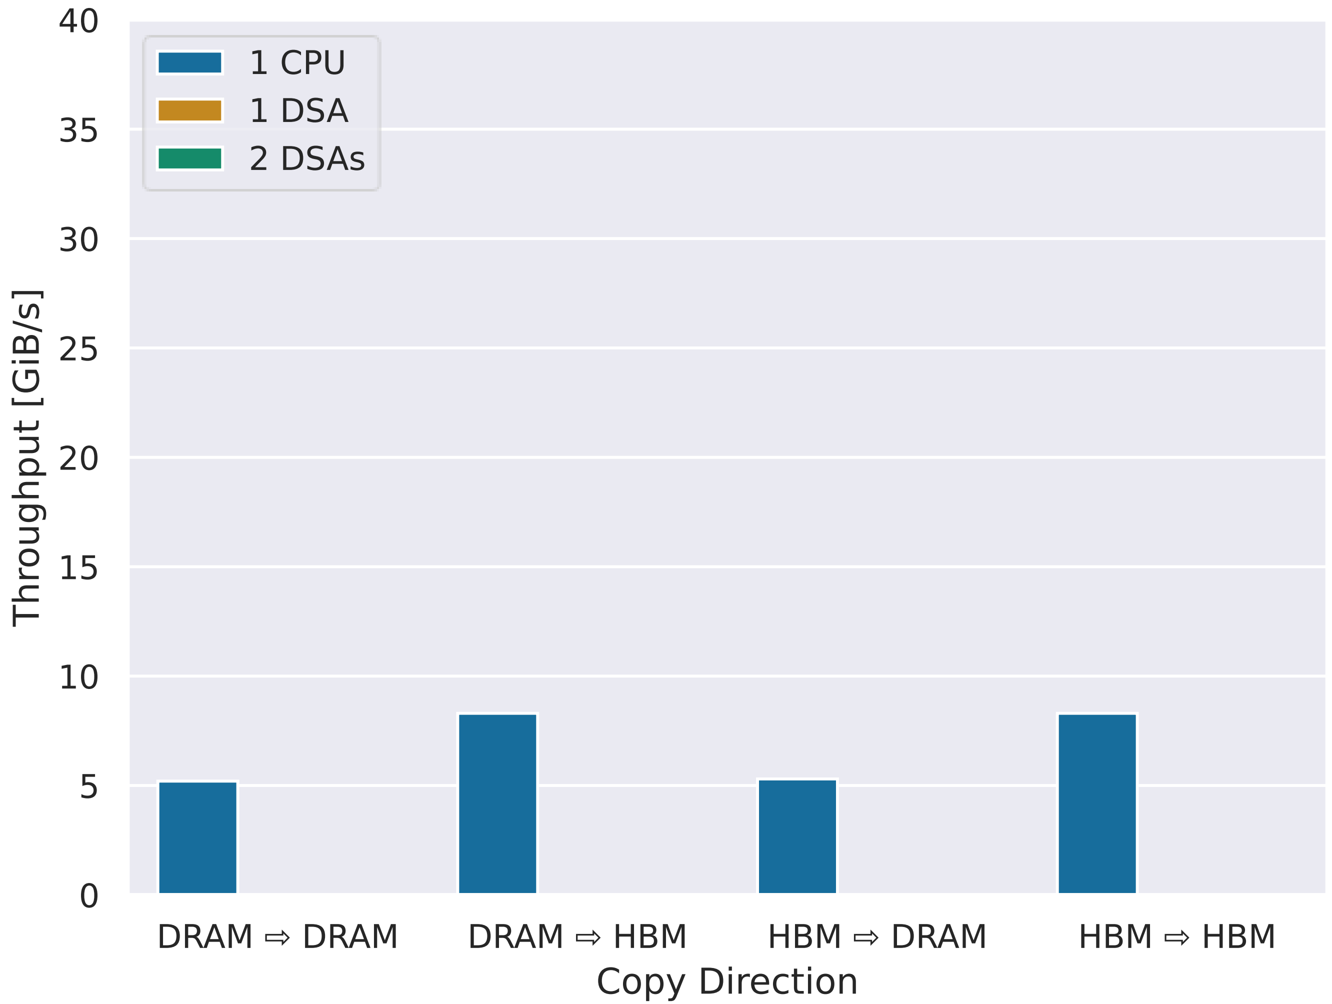

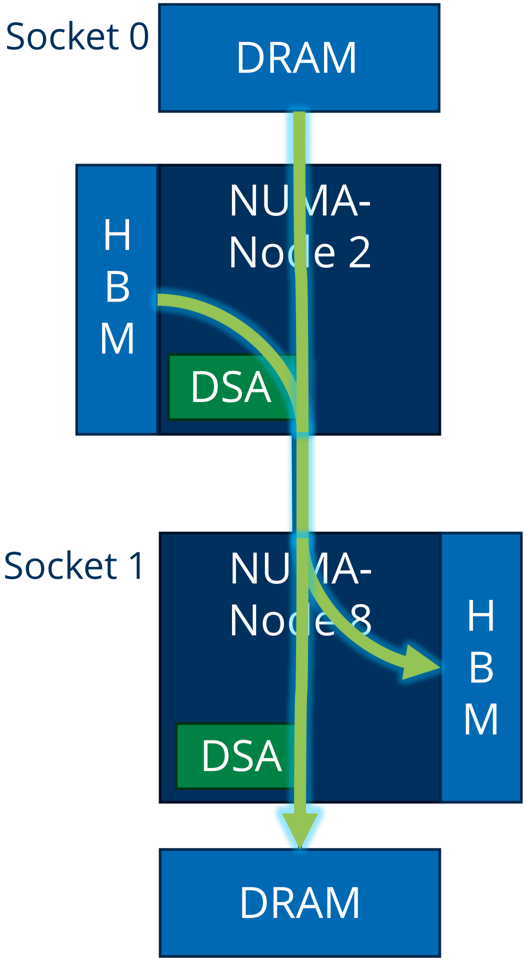

Benchmark – Inter-Socket Data Transfer with DSA

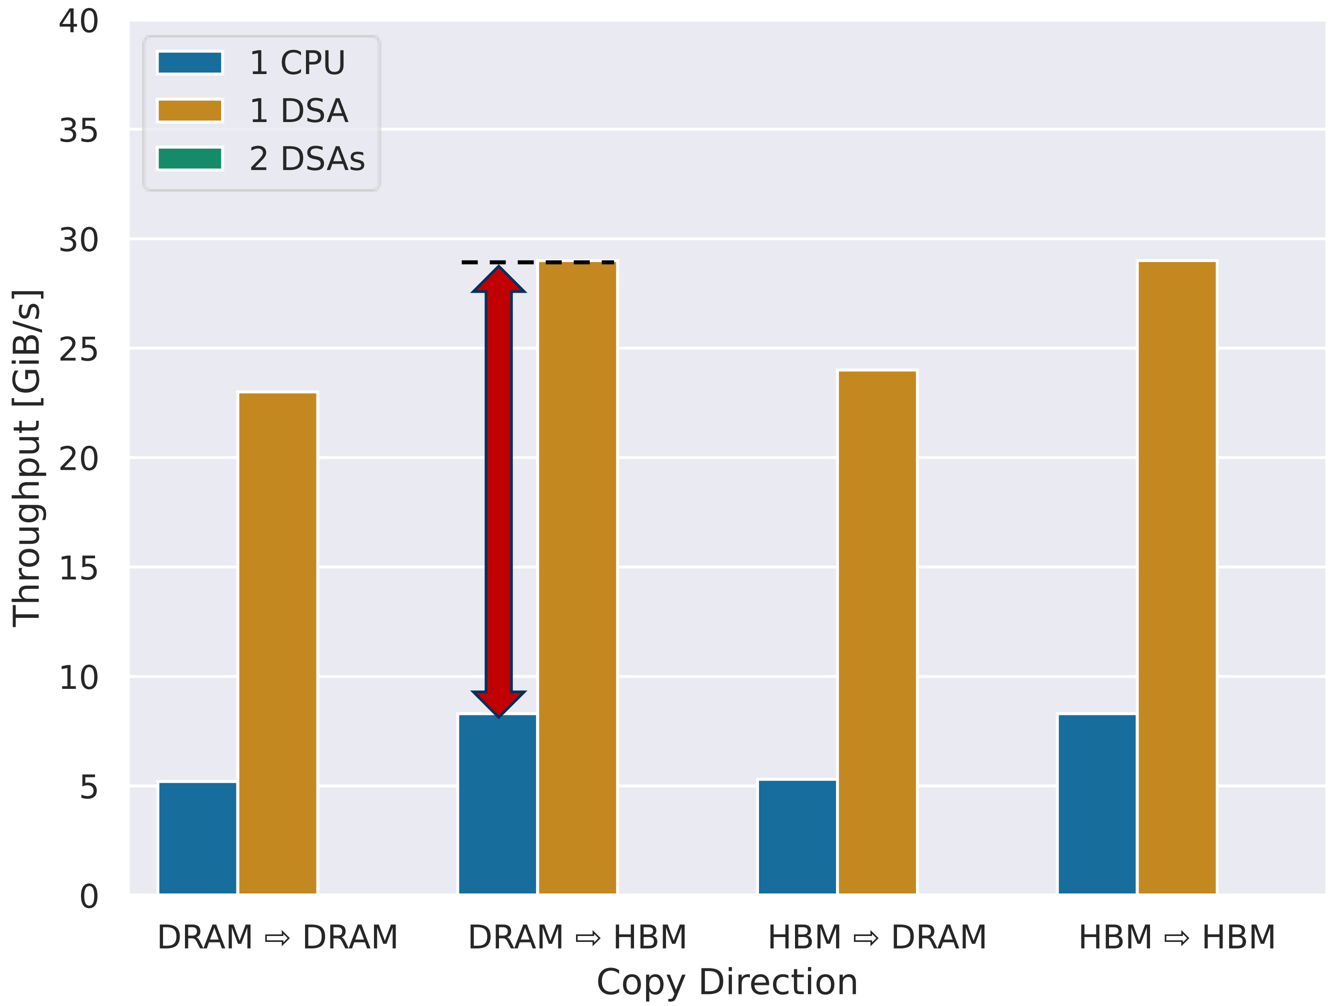

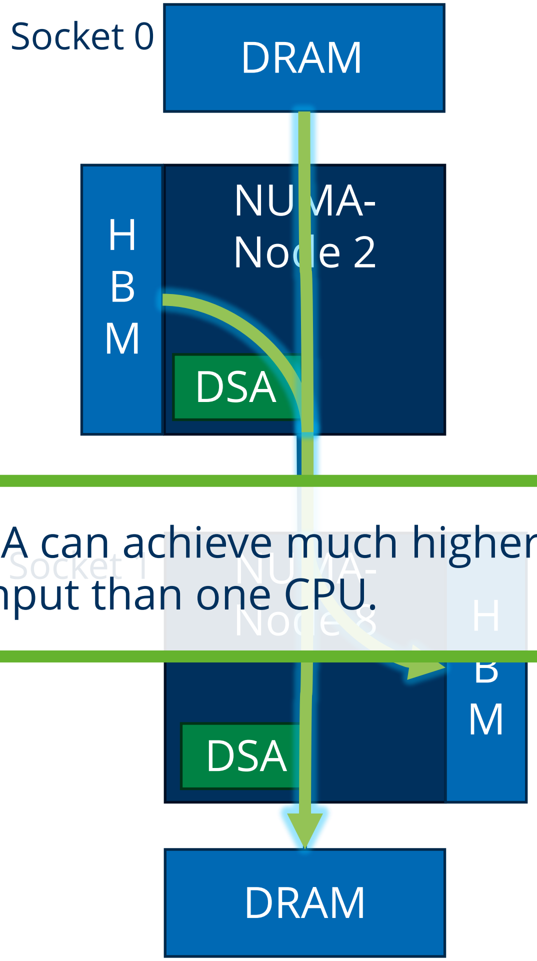

One DSA can achieve much higher throughput than one CPU.

Funded by DFG
DRESDEN concept

Benchmark – Inter-Socket Data Transfer with DSA

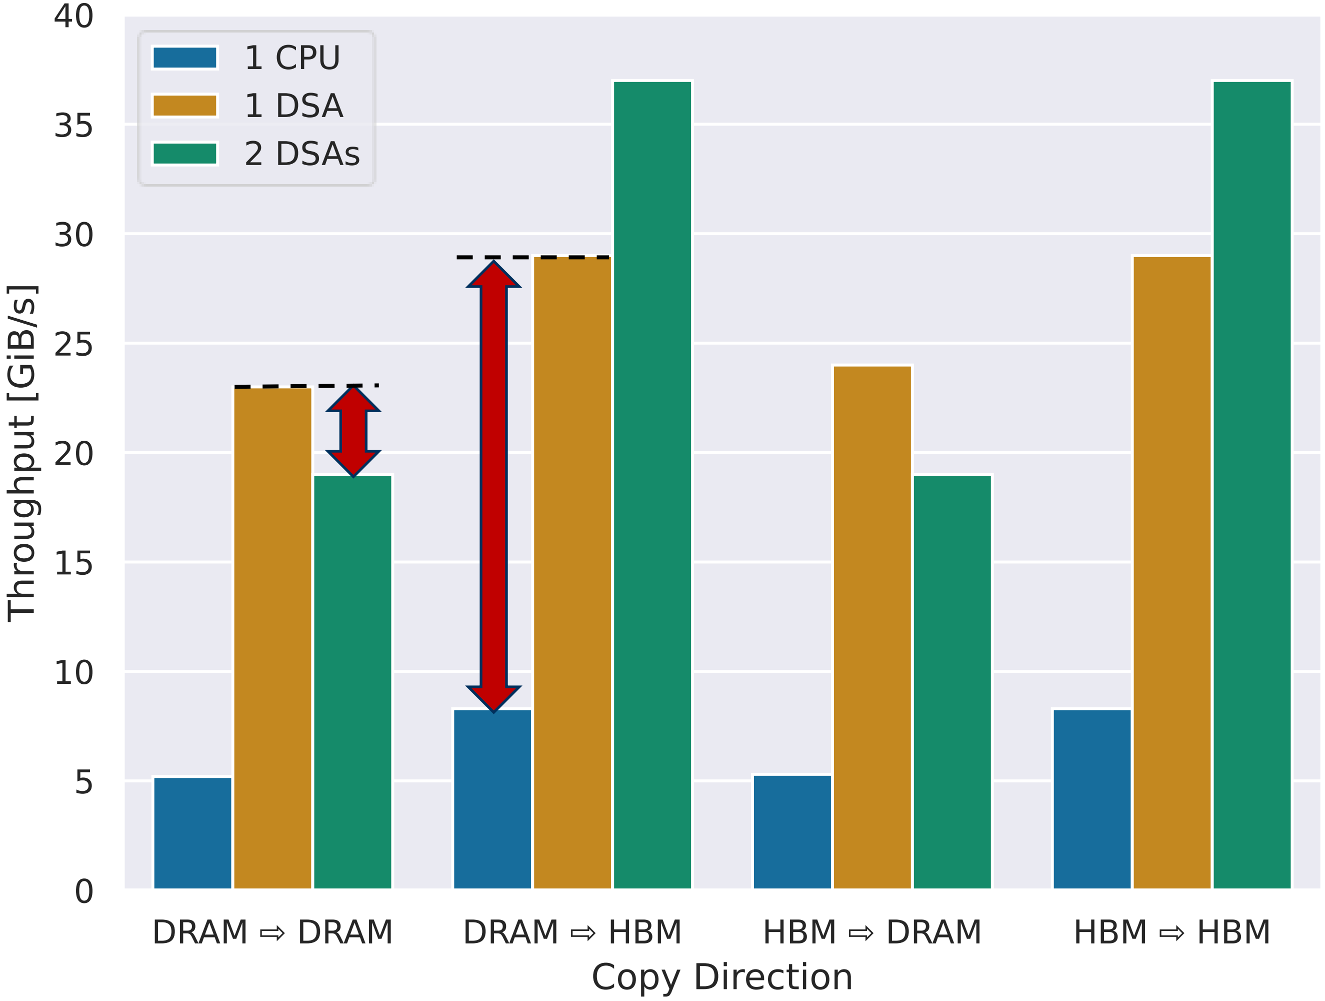

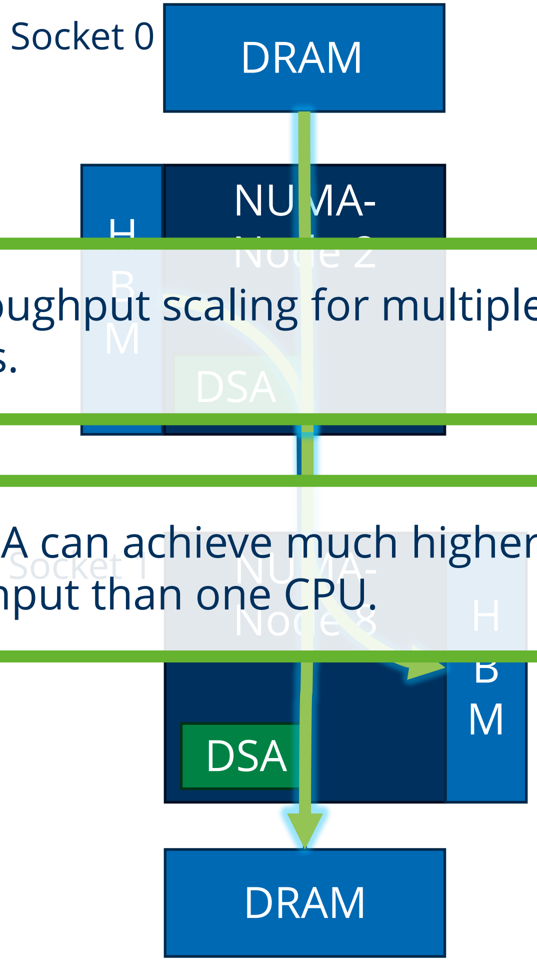

No throughput scaling for multiple DSA engines.

One DSA can achieve much higher throughput than one CPU.

Funded by DFG

DRESDEN concept

Benchmark – Inter-Socket Data Transfer with DSA

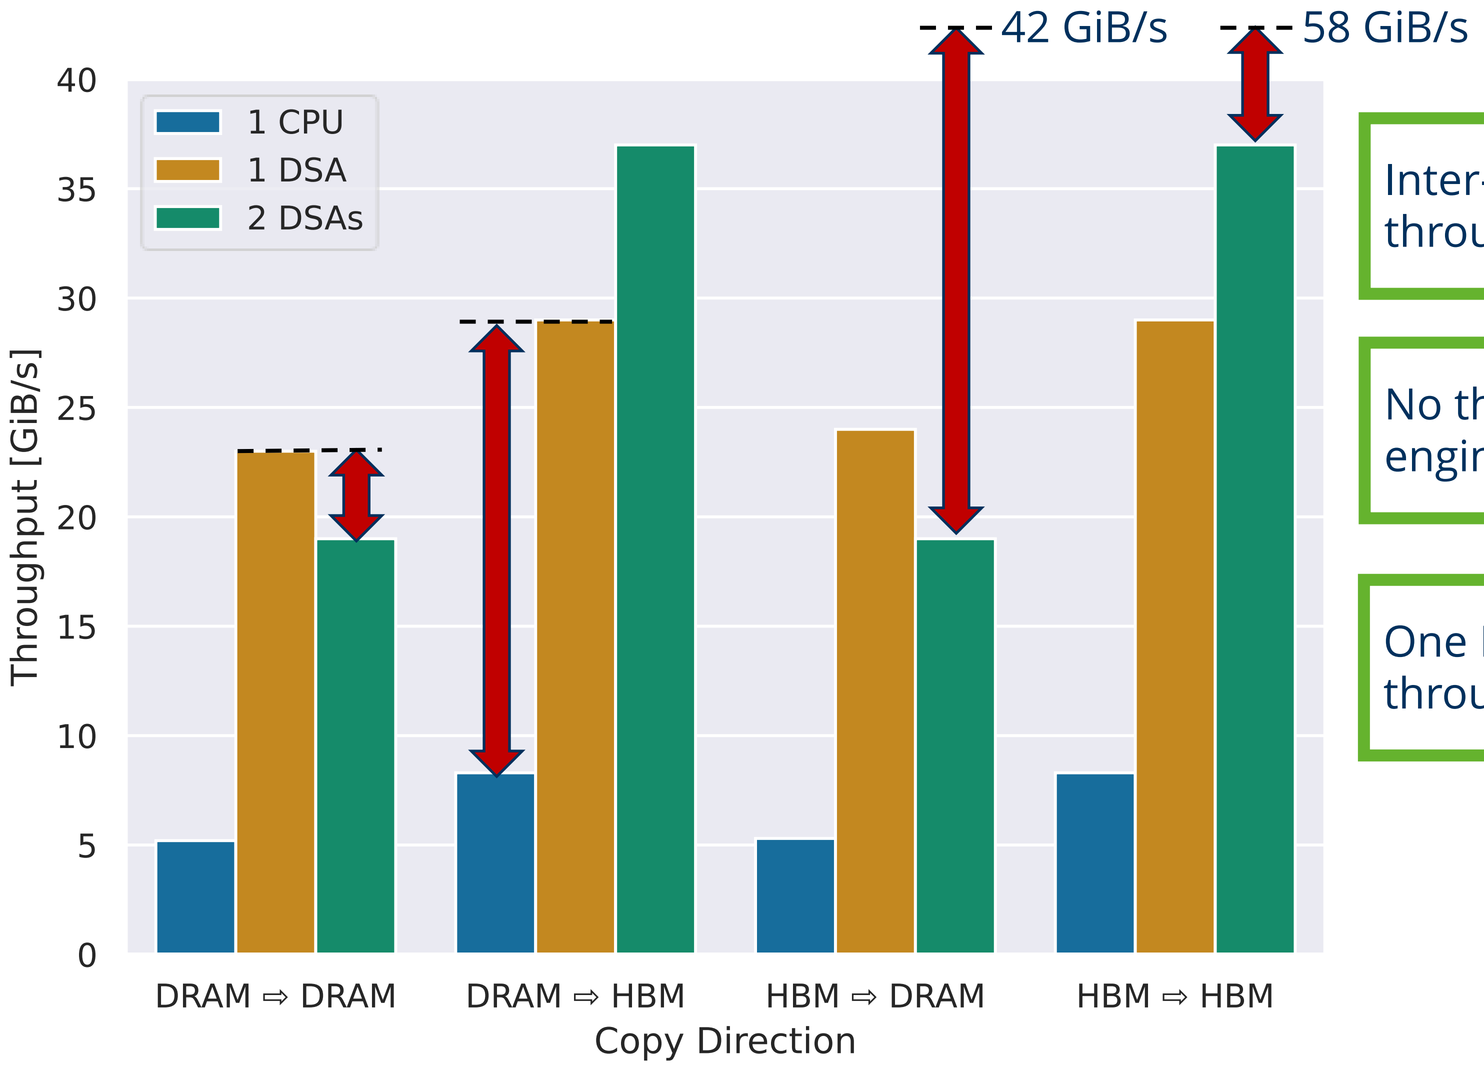

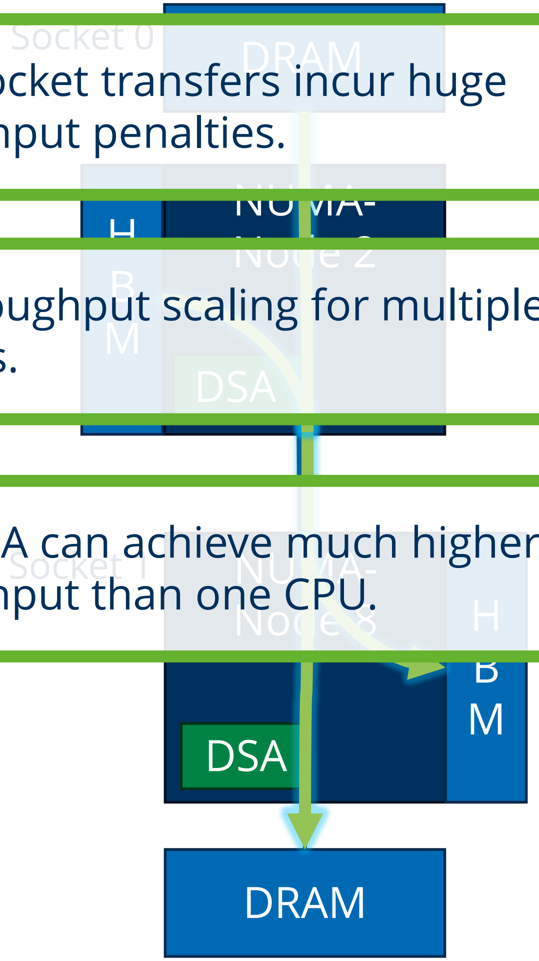

Inter-socket transfers incur huge throughput penalties.

No throughput scaling for multiple DSA engines.

One DSA can achieve much higher throughput than one CPU.

Funded by
DFG
DRESDEN concept

On-the-fly Data Distribution

André Berthold, Lennart Schmidt, Anton Obersteiner, Dirk Habich, Wolfgang Lehner, Horst Schirmeier. 2024. On-The-Fly Data Distribution to Accelerate Query Processing in Heterogeneous Memory Systems. In *28th European Conference on Advances in Databases and Information Systems (ADBIS ’24).* Springer Nature Switzerland, Cham, 170–183.

Funded by DFG
DRESDEN concept

On-the-fly Data Distribution

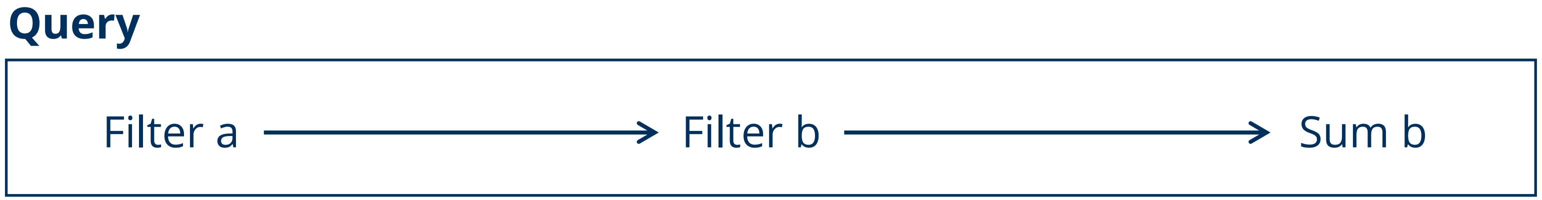

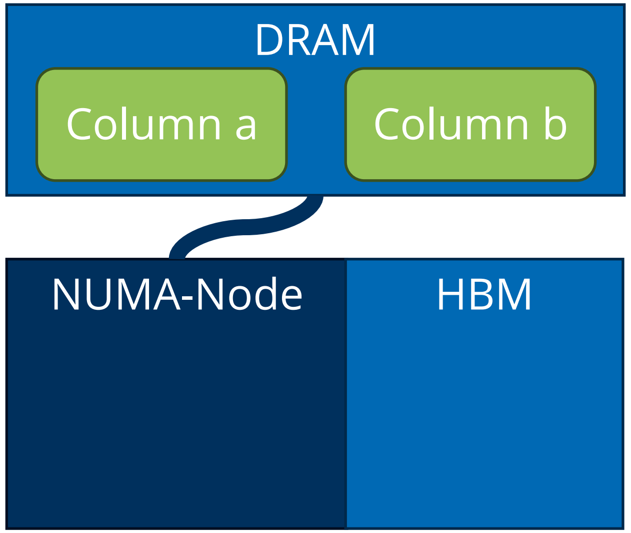

Funded by

On-the-fly Data Distribution

Baseline Execution Time

Query

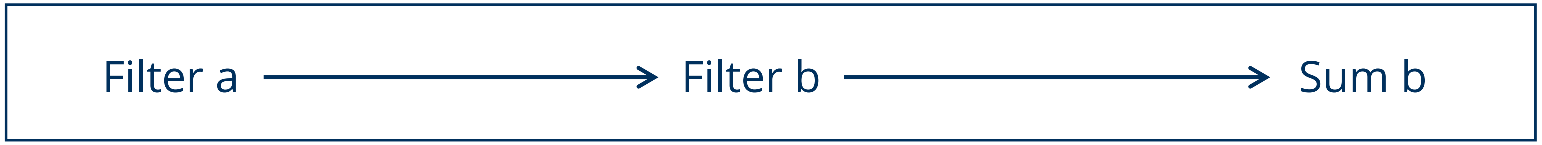

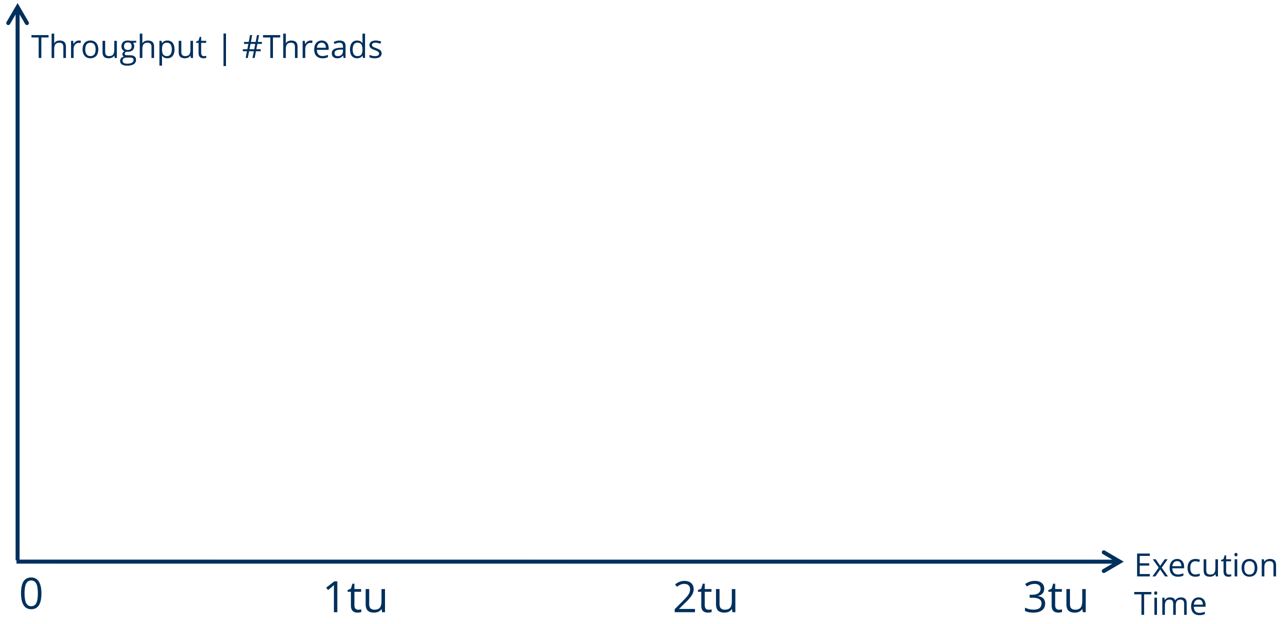

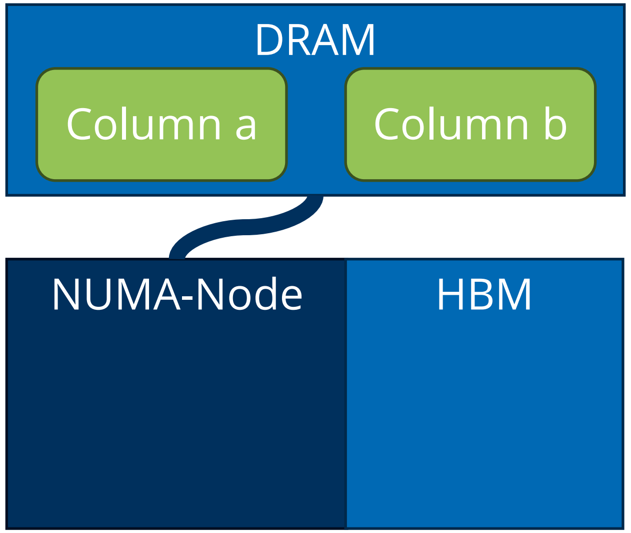

Funded by DFG
DRESDEN concept

On-the-fly Data Distribution

Baseline Execution Time

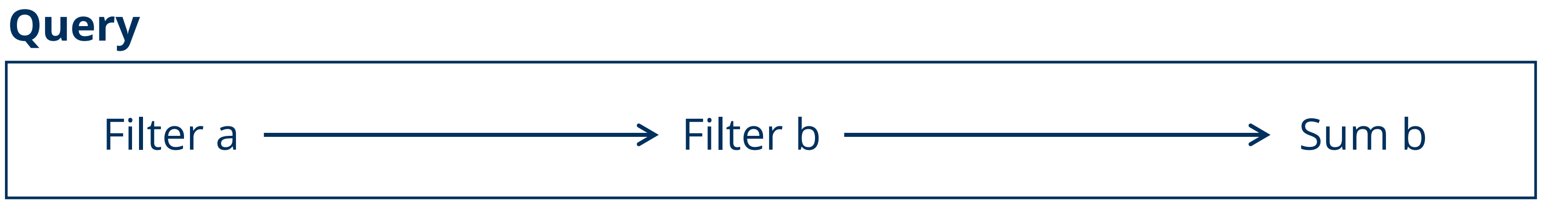

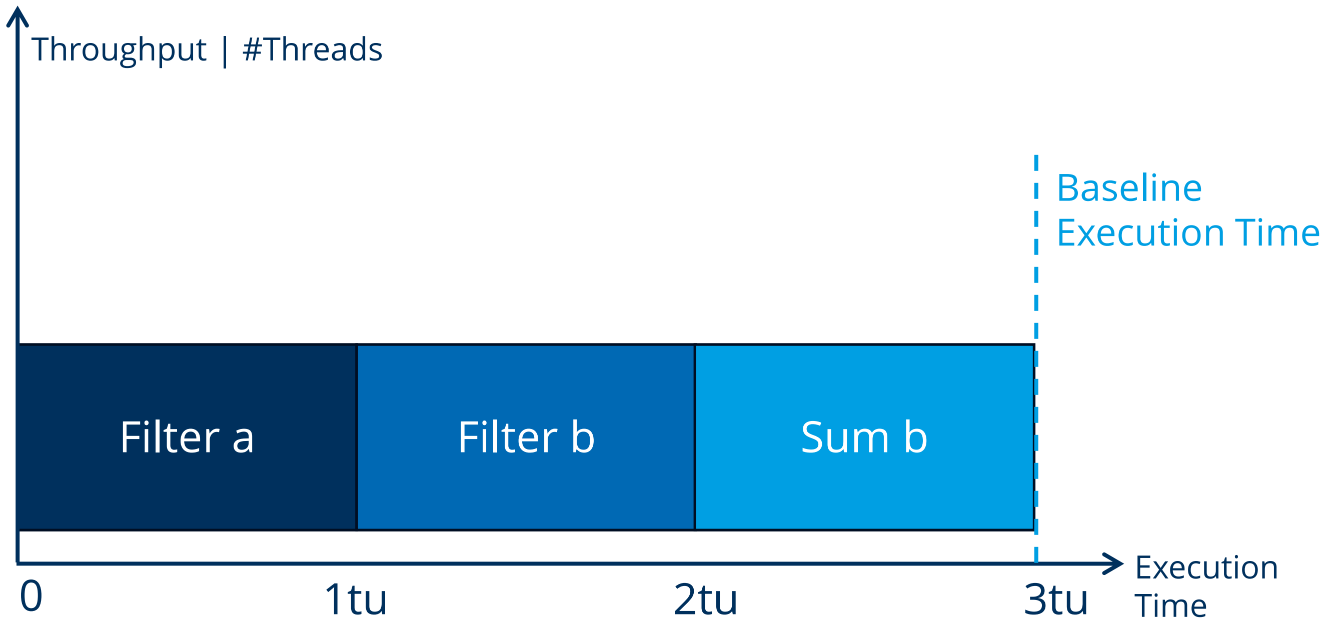

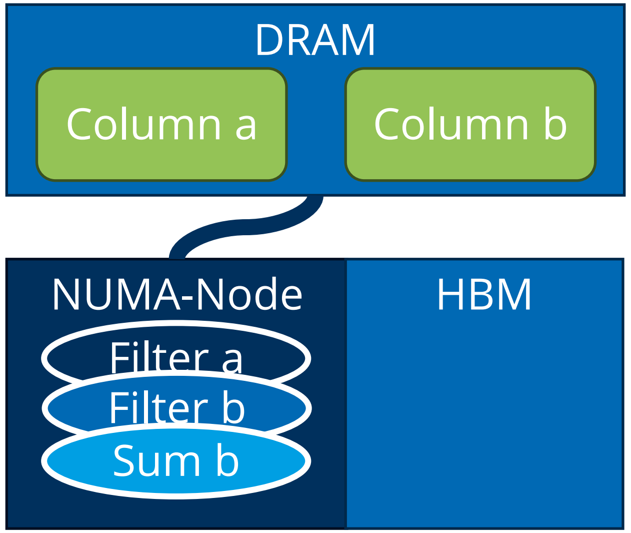

Funded by
DFG
DRESDEN concept

On-the-fly Data Distribution
Optimized Execution Time

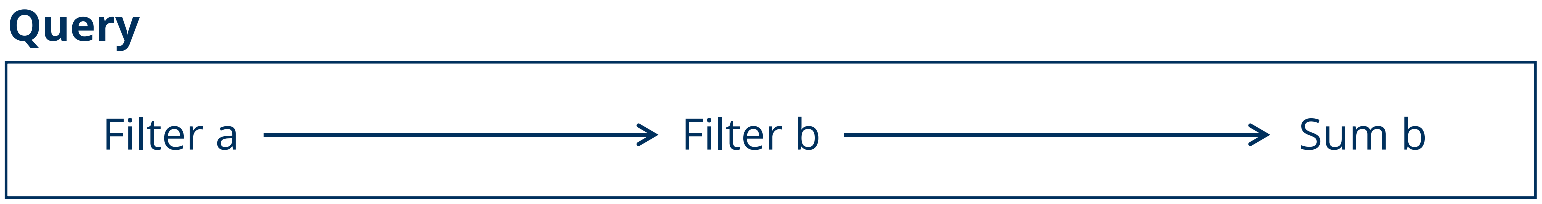

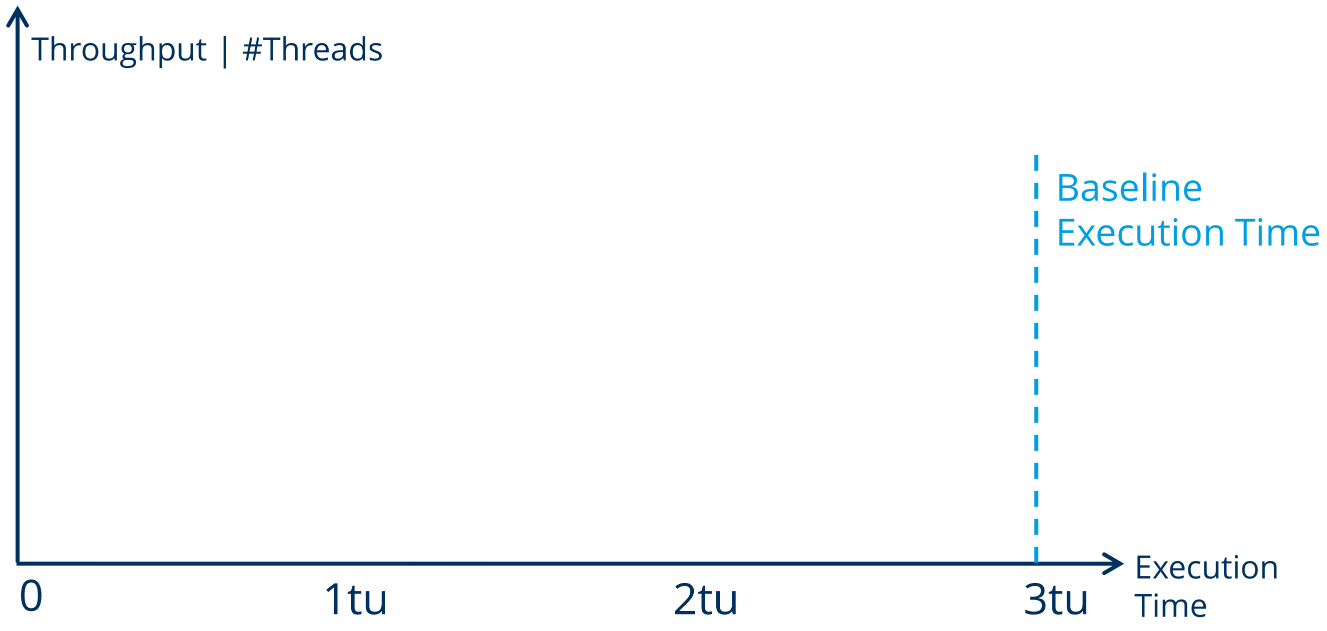

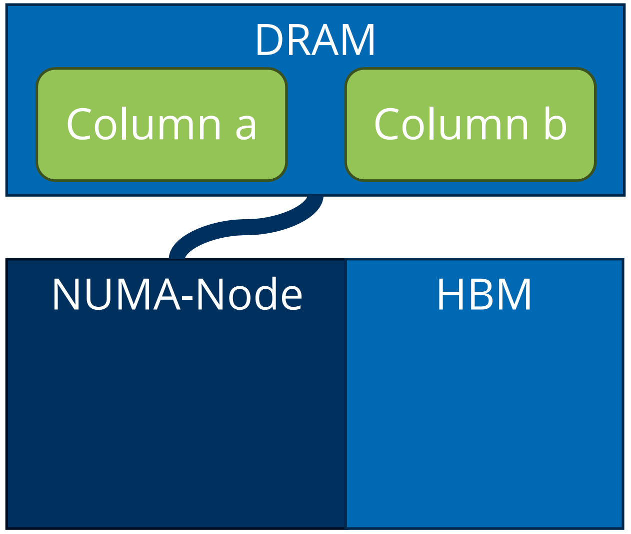

Funded by
DFG
DRESDEN concept

On-the-fly Data Distribution

Optimized Execution Time

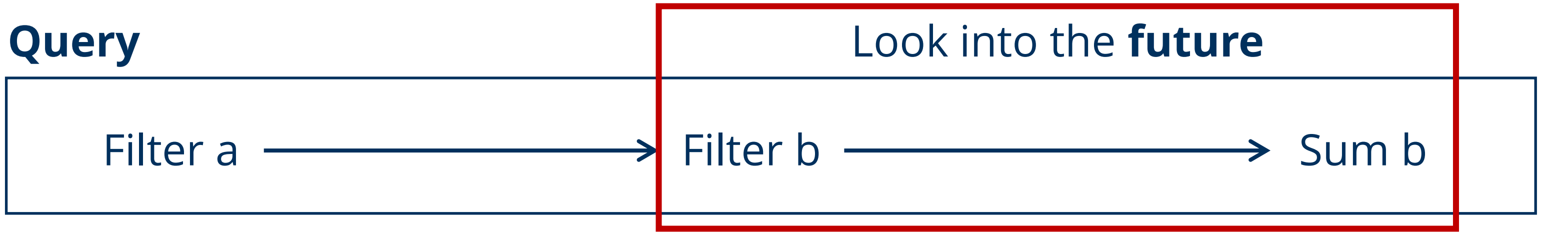

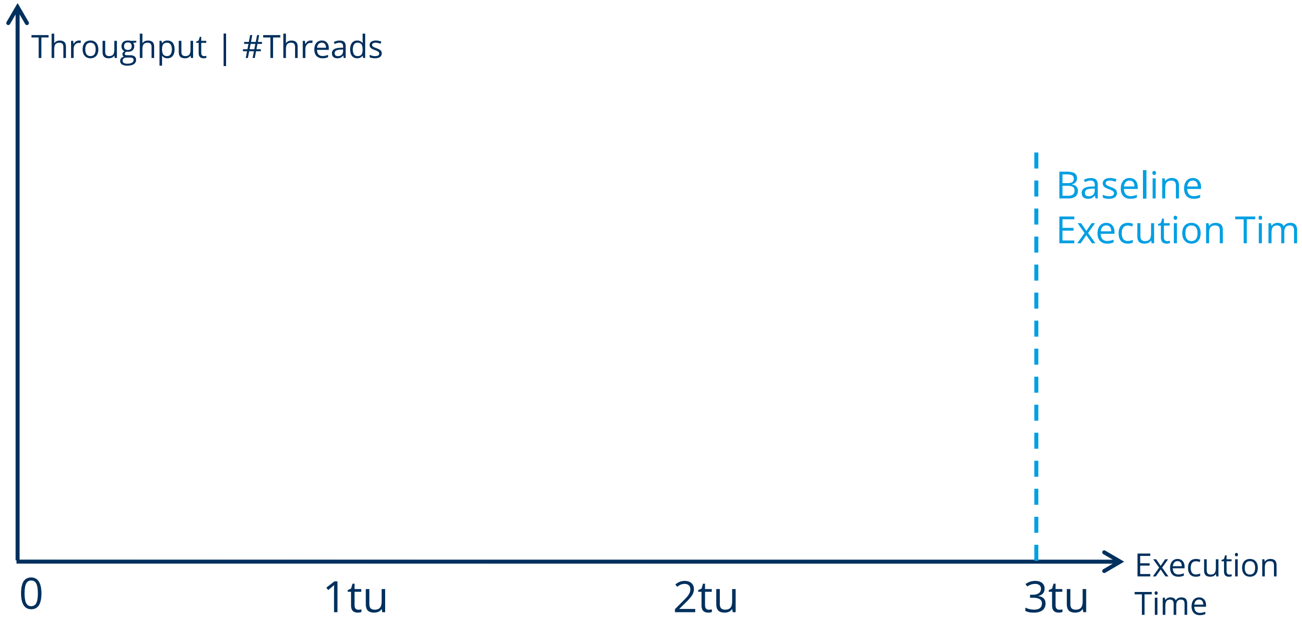

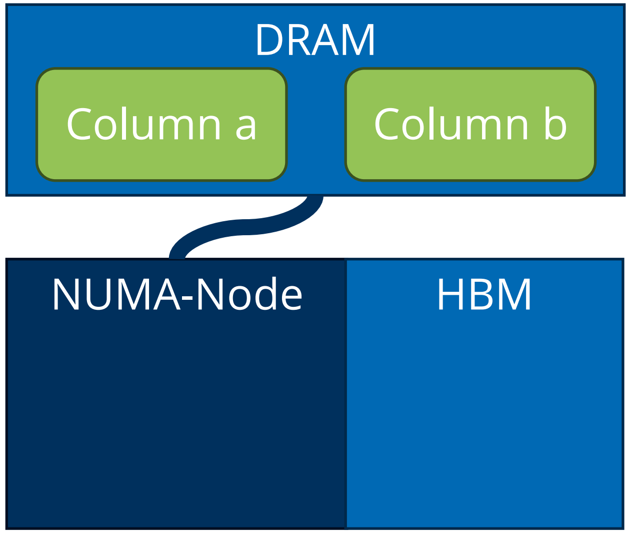

Funded by
DFG
DRESDEN concept

On-the-fly Data Distribution

Optimized Execution Time

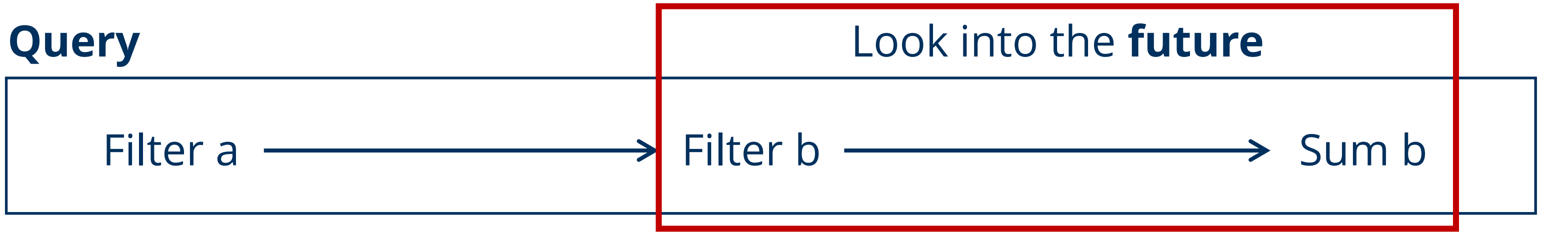

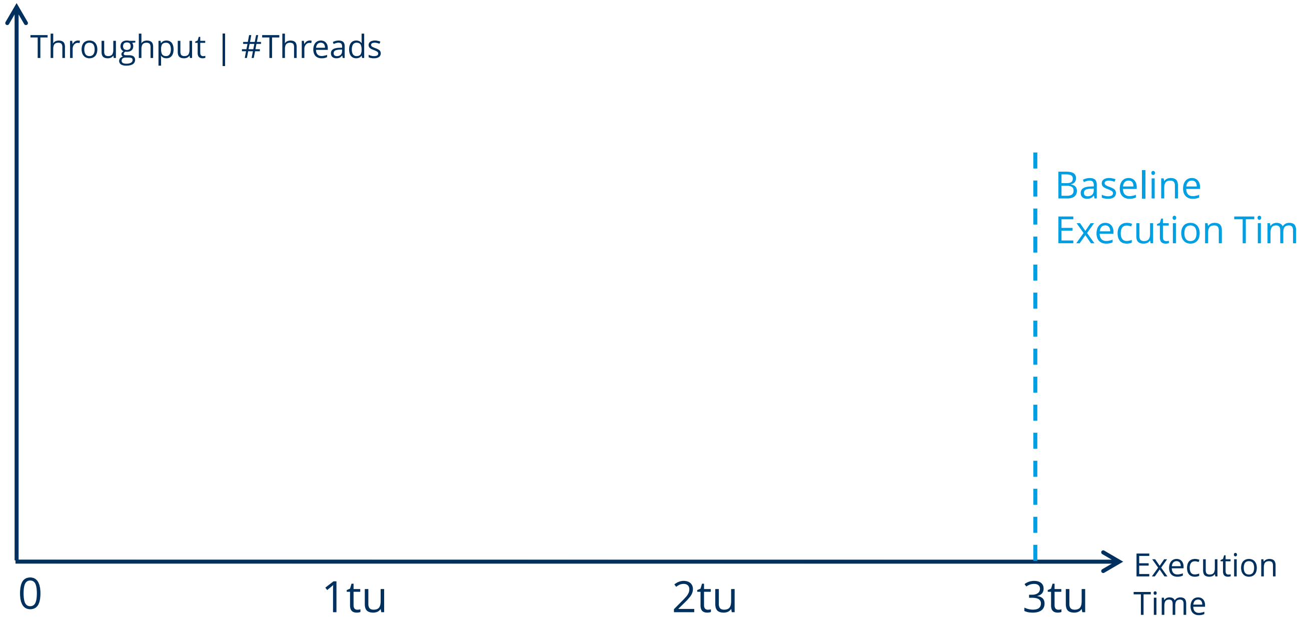

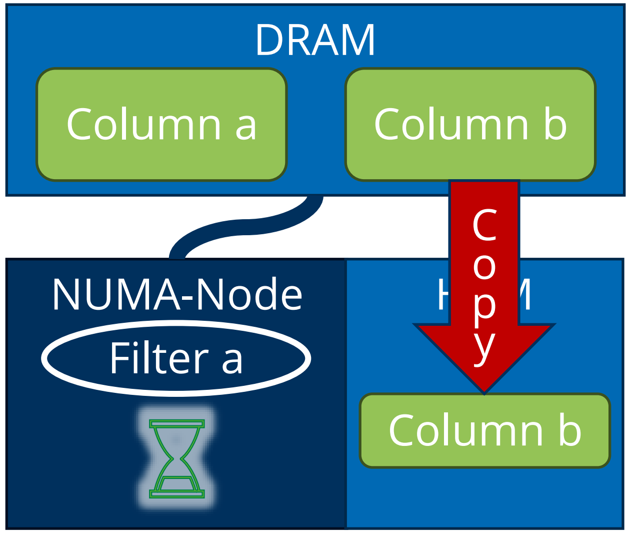

Funded by DFG DRESDEN concept

On-the-fly Data Distribution

Optimized Execution Time

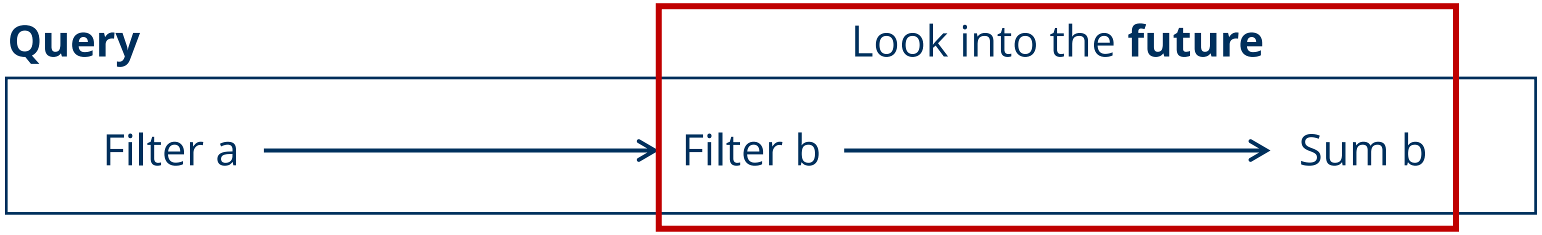

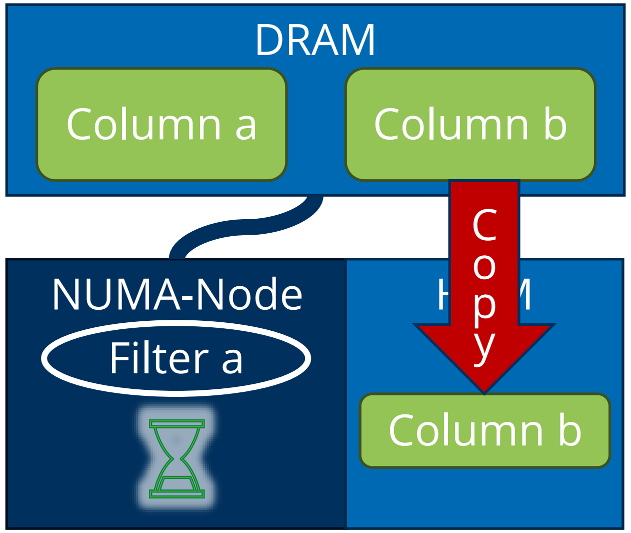

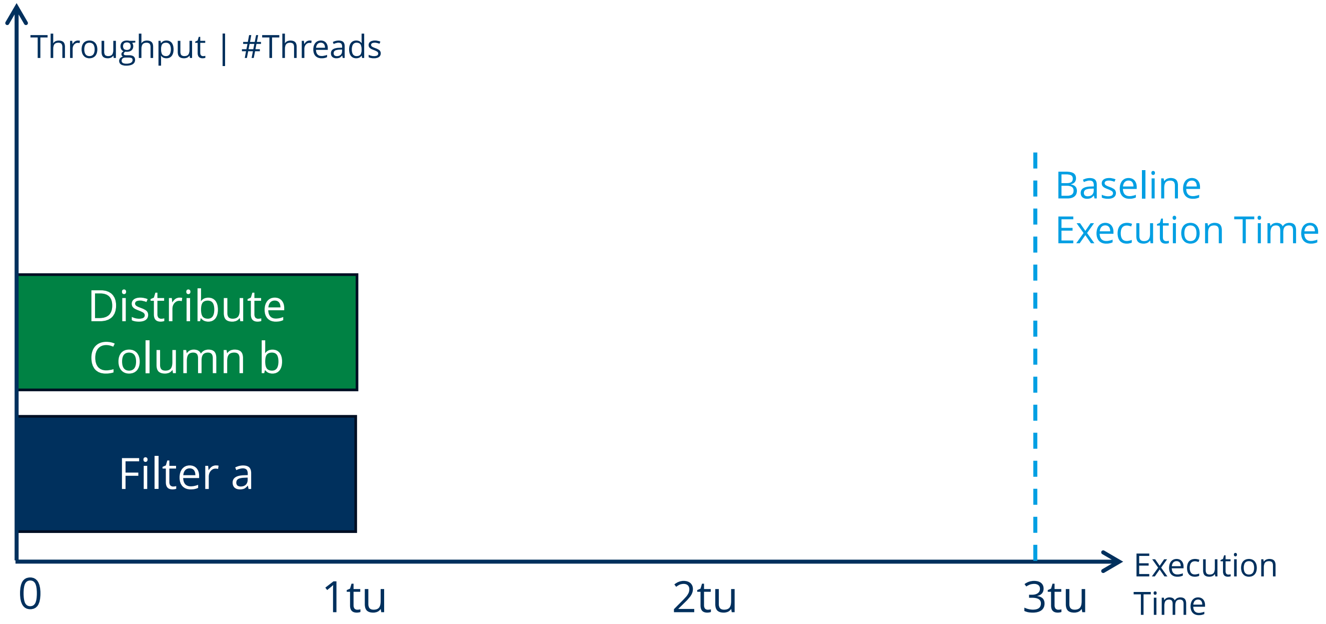

Funded by

On-the-fly Data Distribution
Optimized Execution Time

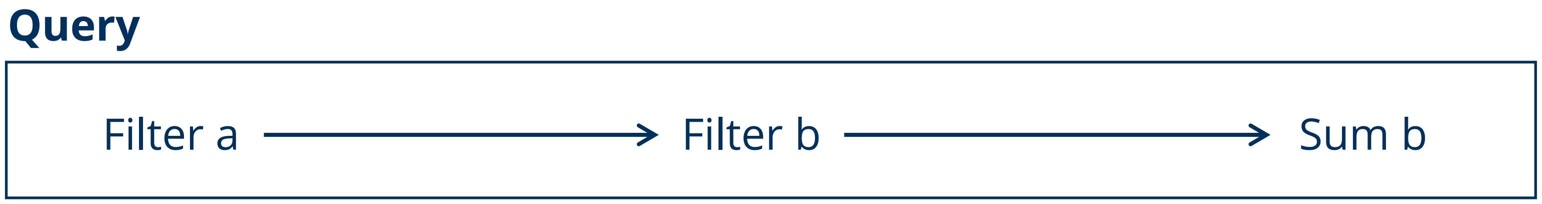

Query

Filter a → Filter b → Sum b

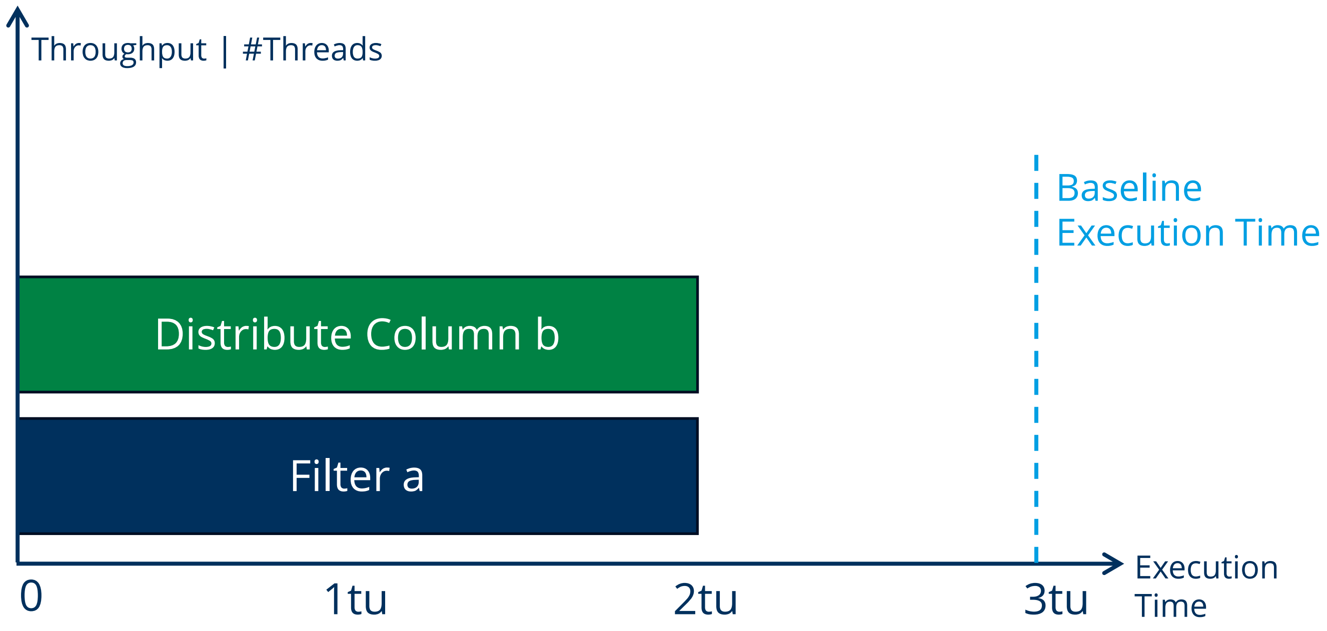

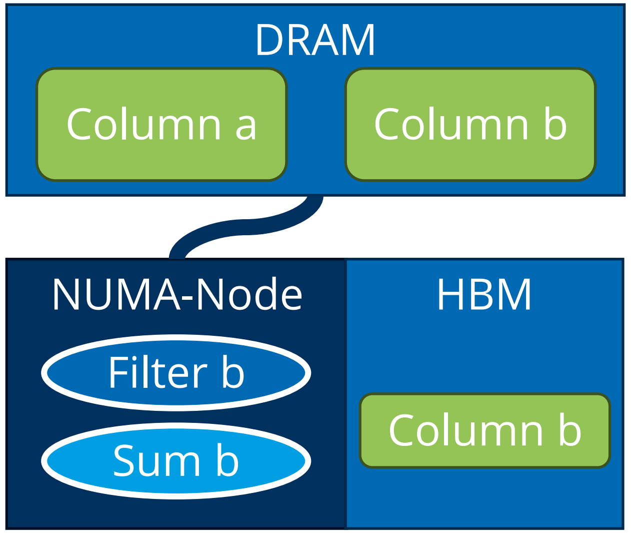

Funded by DFG DRESDEN concept

On-the-fly Data Distribution
Optimized Execution Time

Query

Filter a → Filter b → Sum b

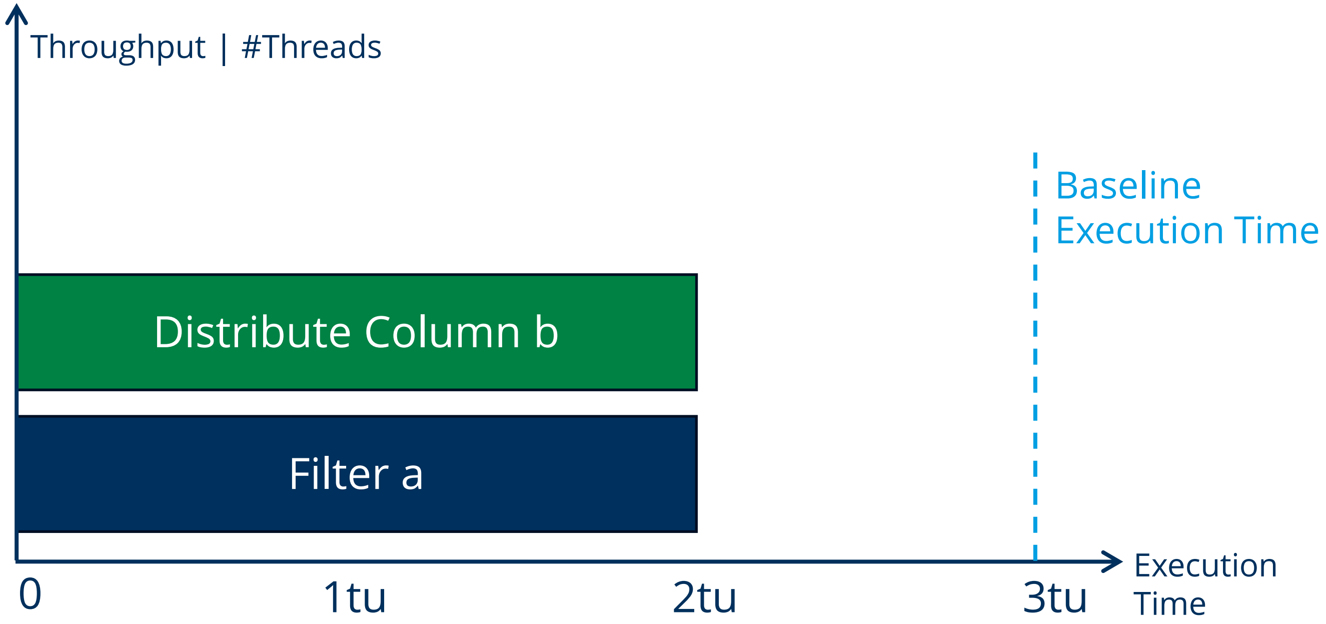

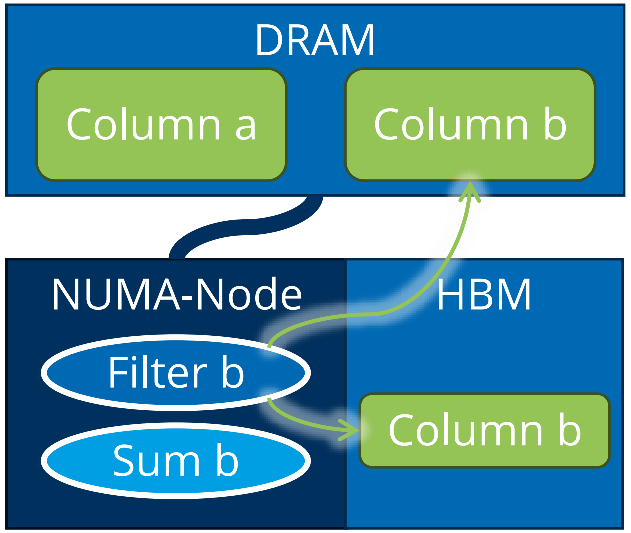

Funded by
DFG
DRESDEN concept

On-the-fly Data Distribution

Optimized Execution Time

Query

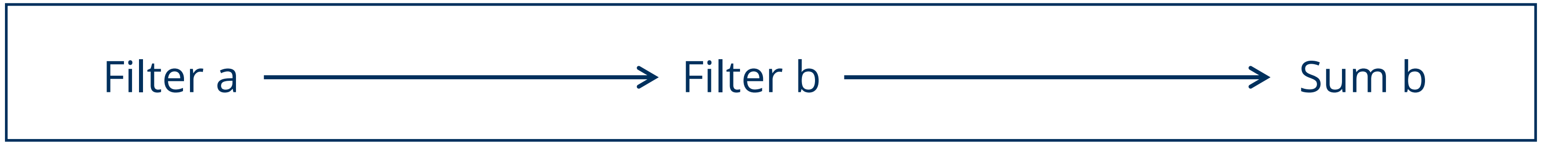

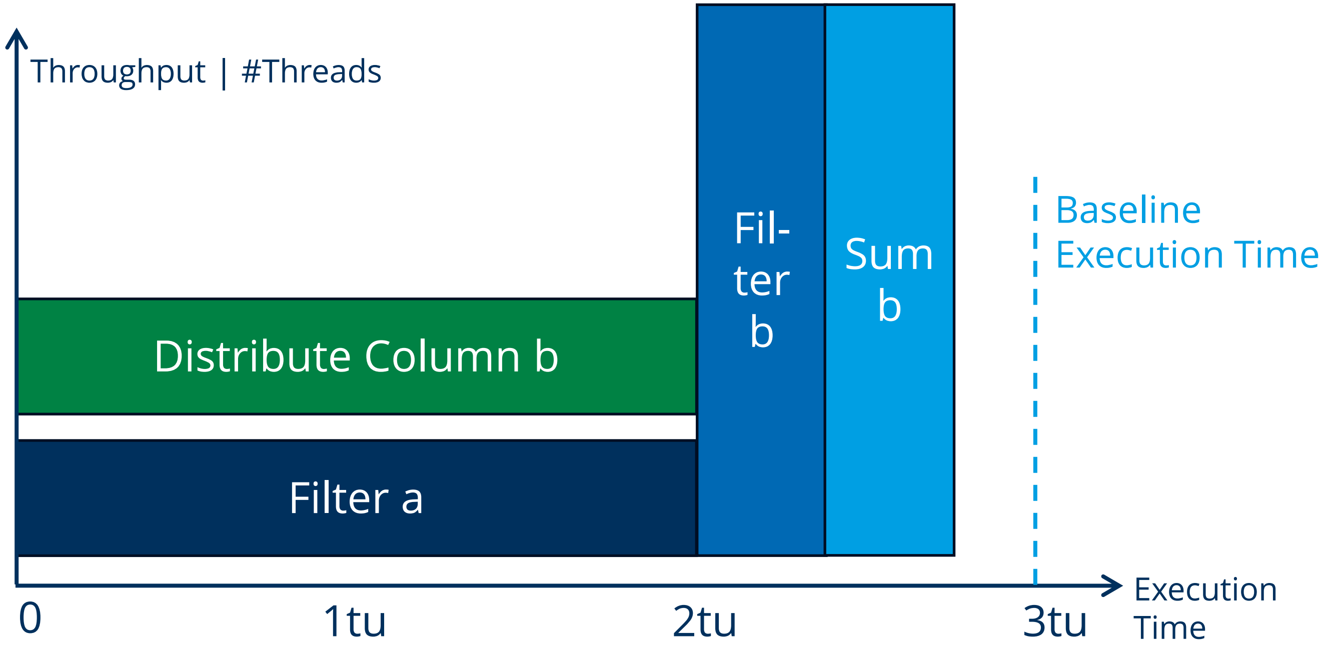

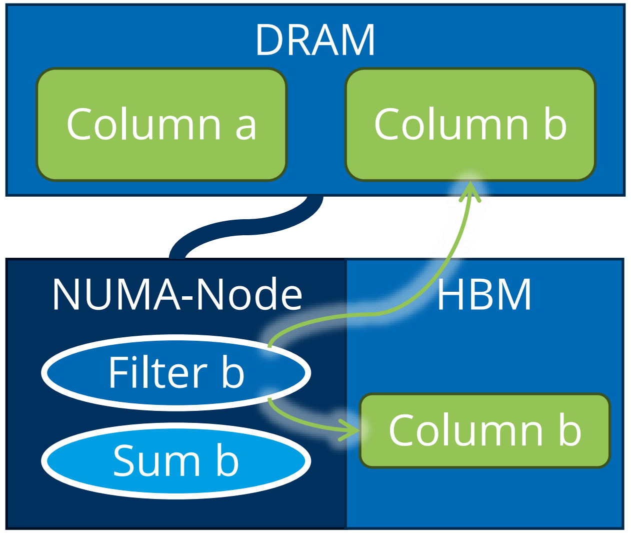

Funded by

On-the-fly Data Distribution
Optimized Execution Time

Query

Filter a → Filter b → Sum b

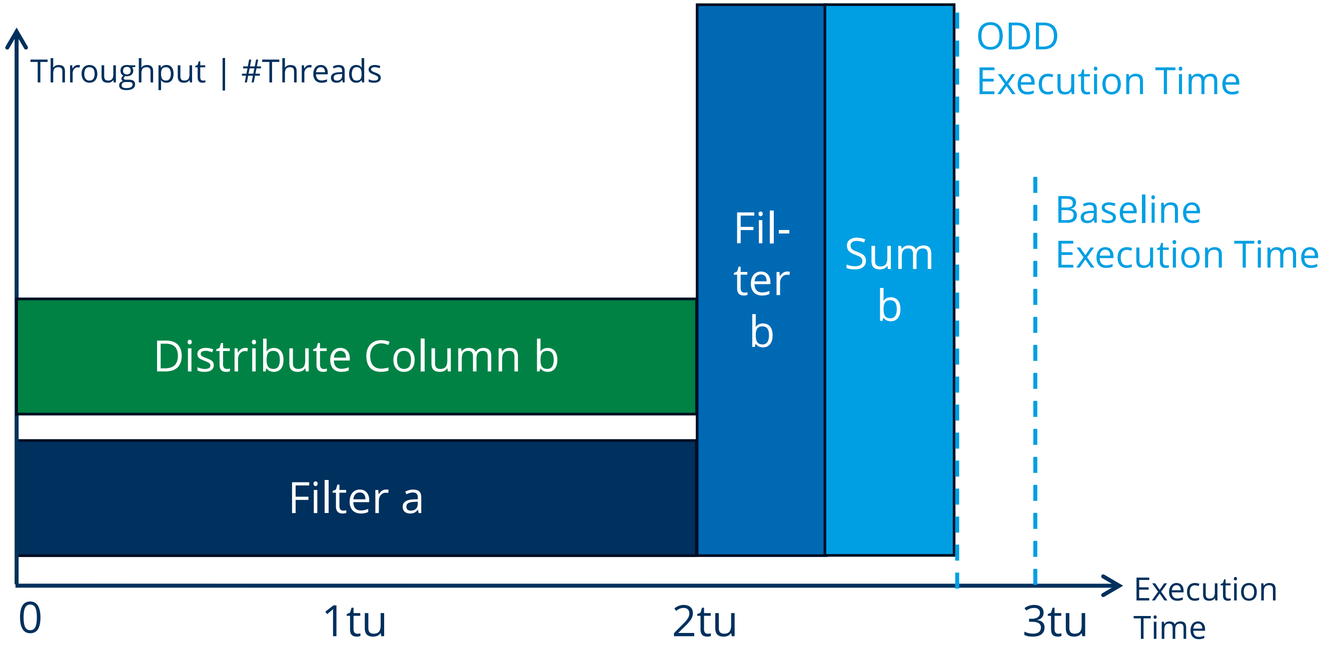

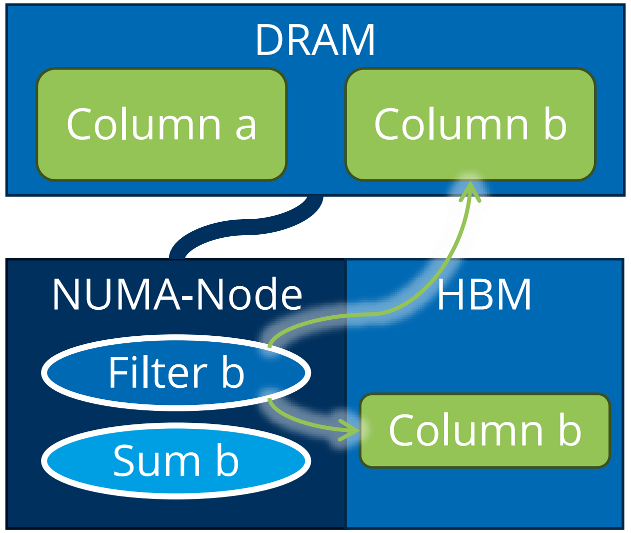

Funded by DFG
DRESDEN concept

On-the-fly Data Distribution

Optimized Execution Time

Query

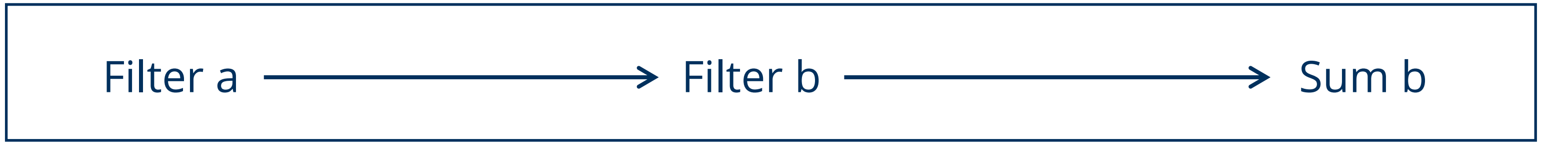

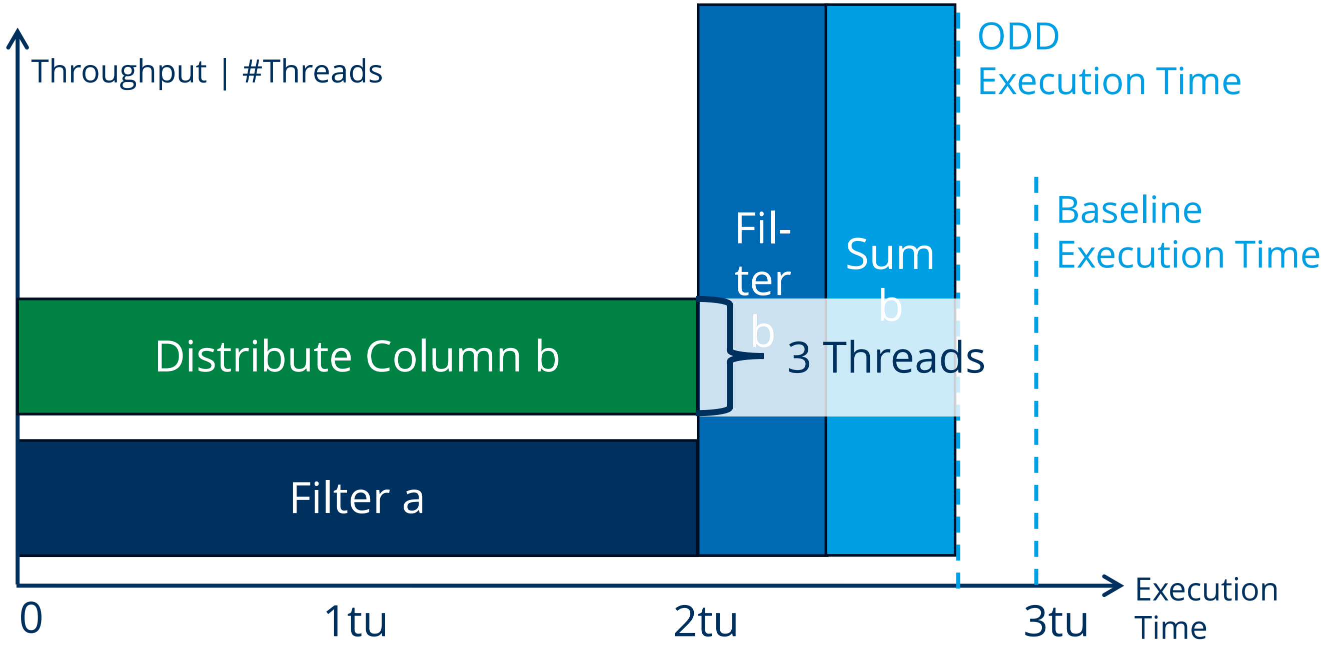

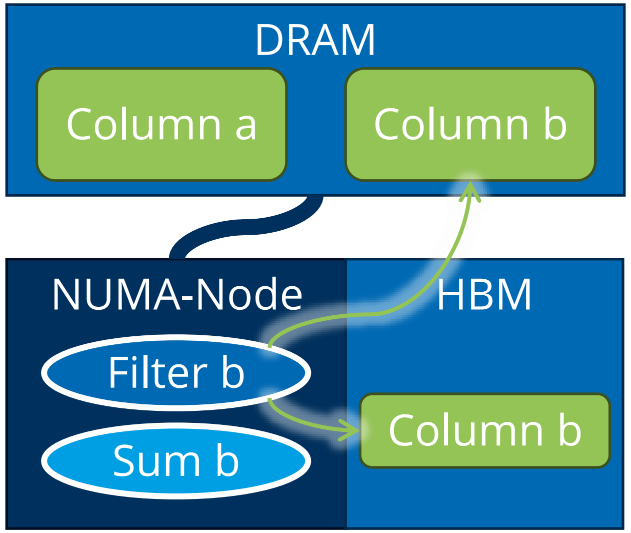

Funded by

On-the-fly Data Distribution

Optimized using DSA

Query

Filter a → Filter b → Sum b

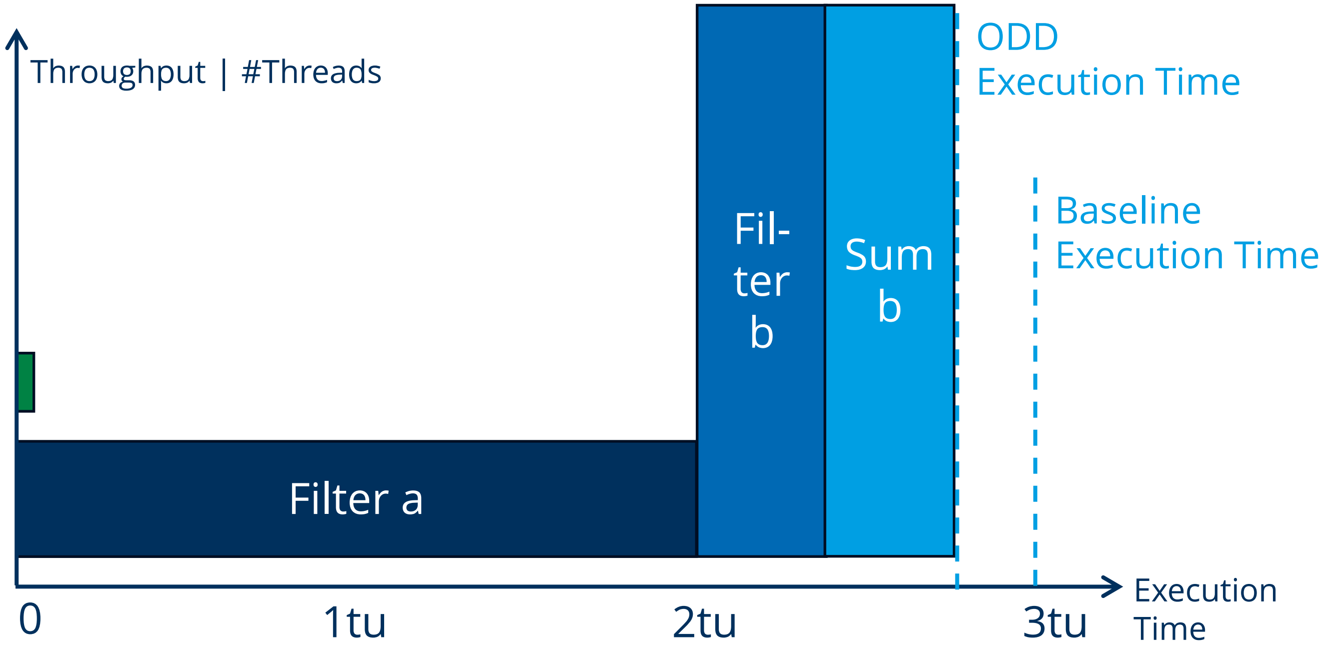

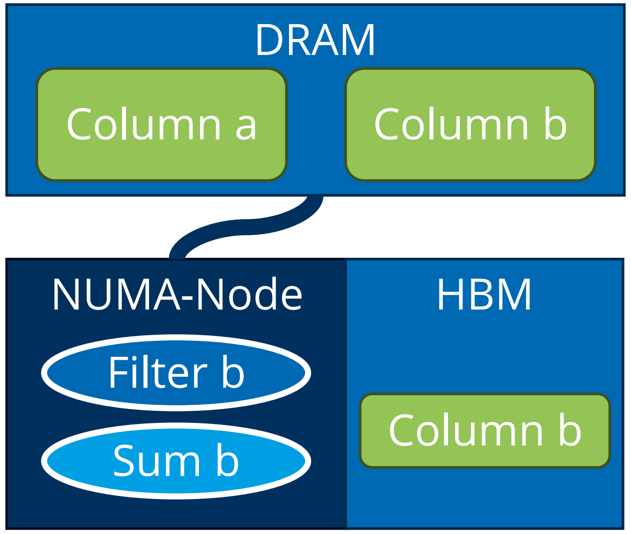

Funded by

On-the-fly Data Distribution

Optimized using DSA

Query

Filter a → Filter b → Sum b

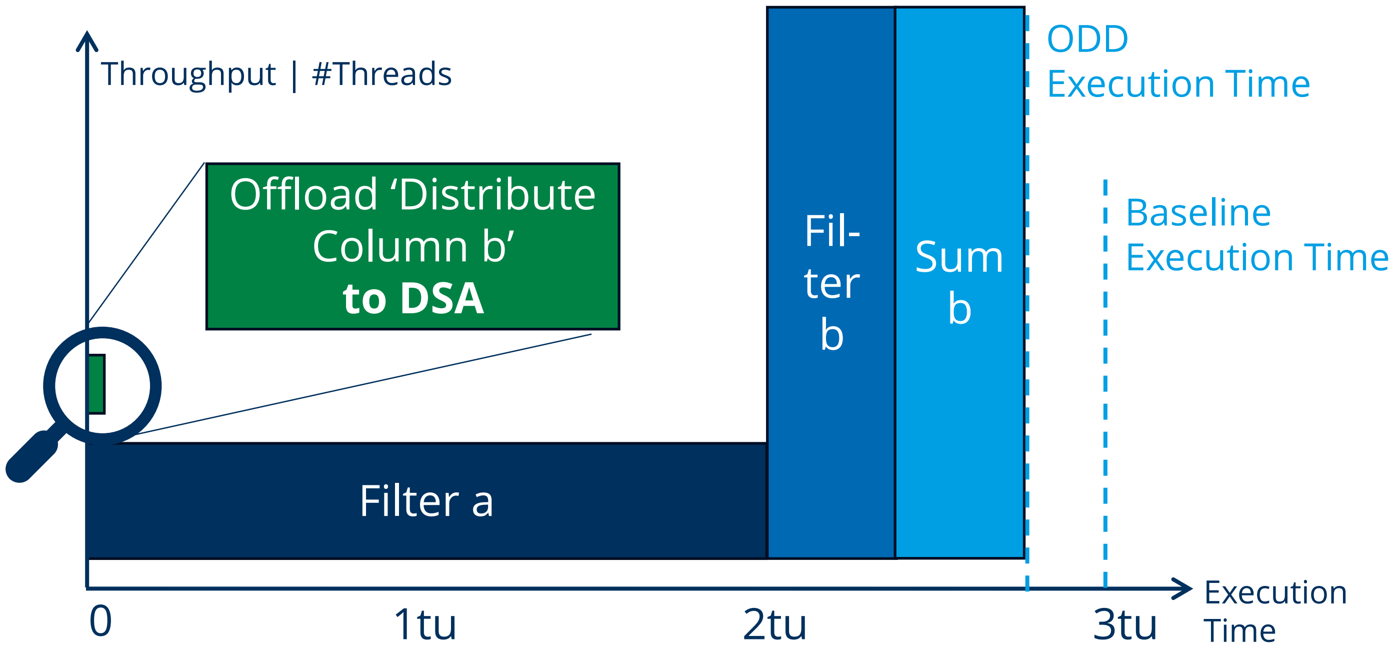

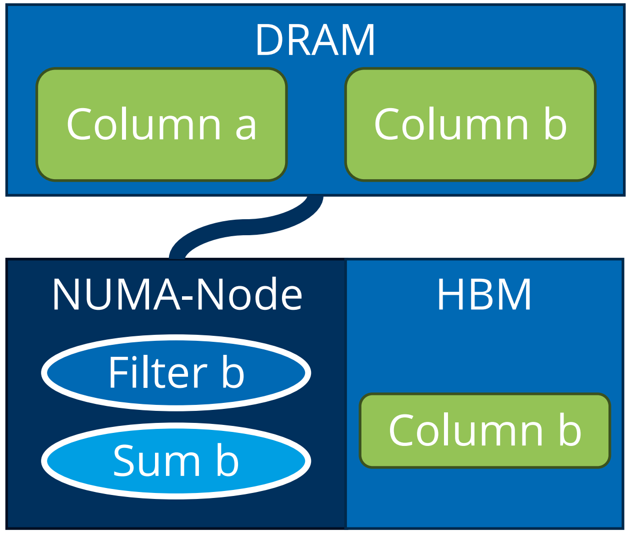

Funded by DFG
DRESDEN concept

On-the-fly Data Distribution

Optimized using DSA

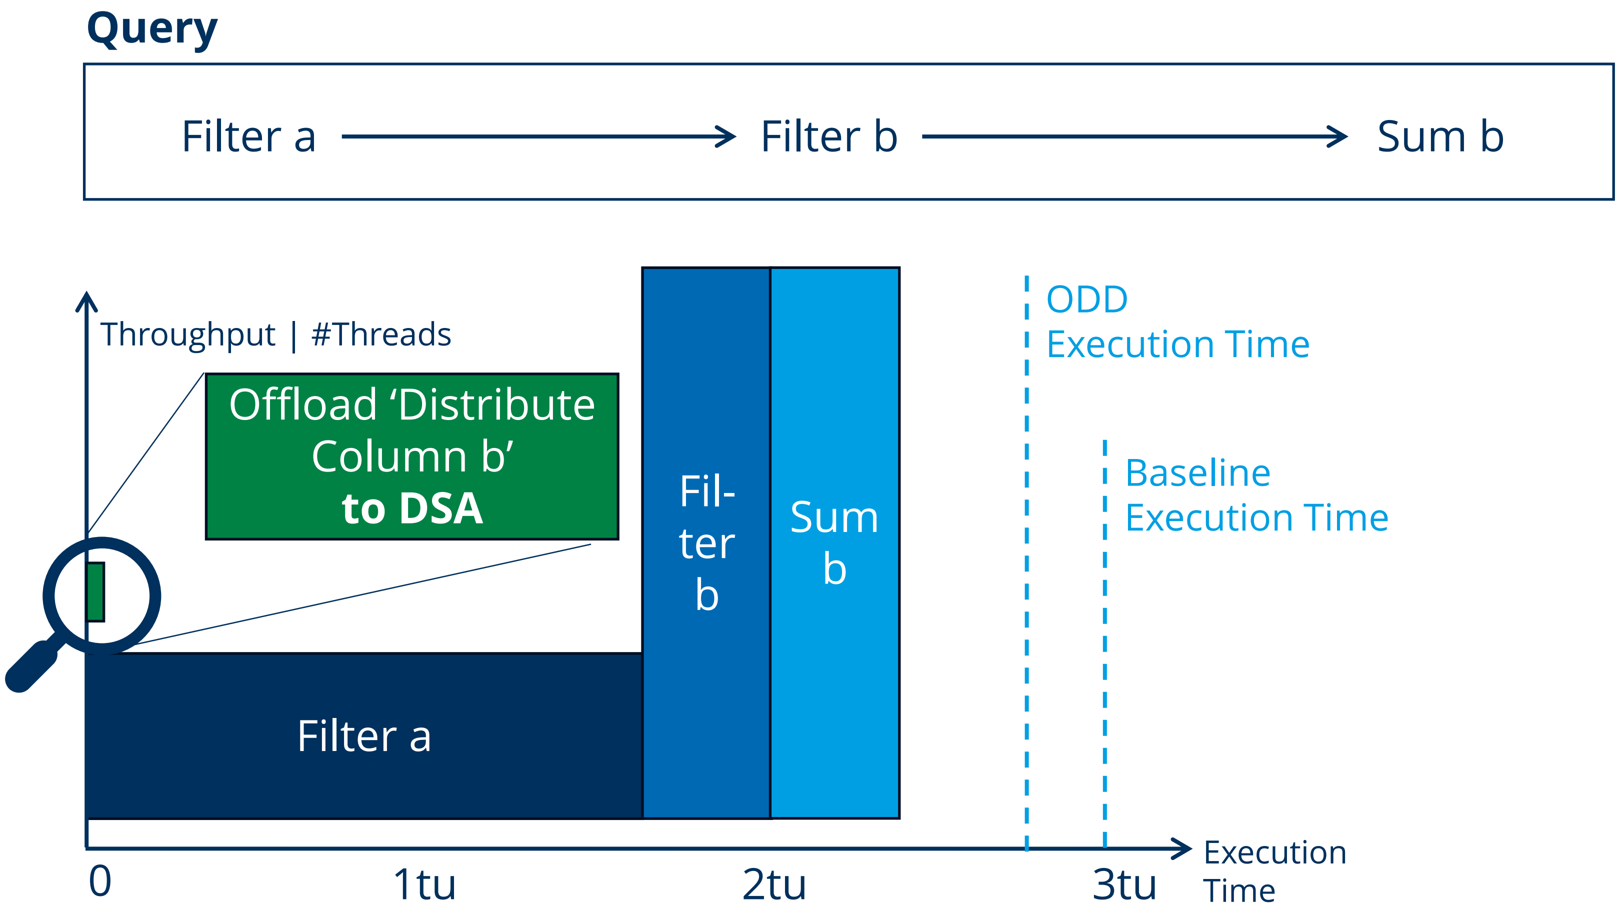

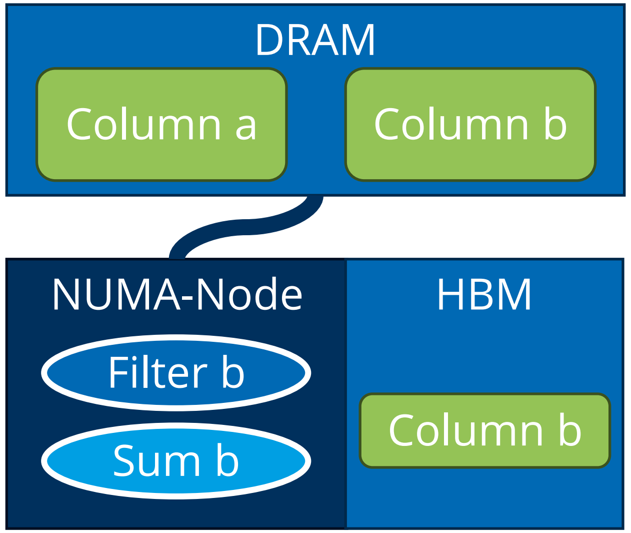

Funded by DFG

DRESDEN concept

On-the-fly Data Distribution with DSA - Benchmark

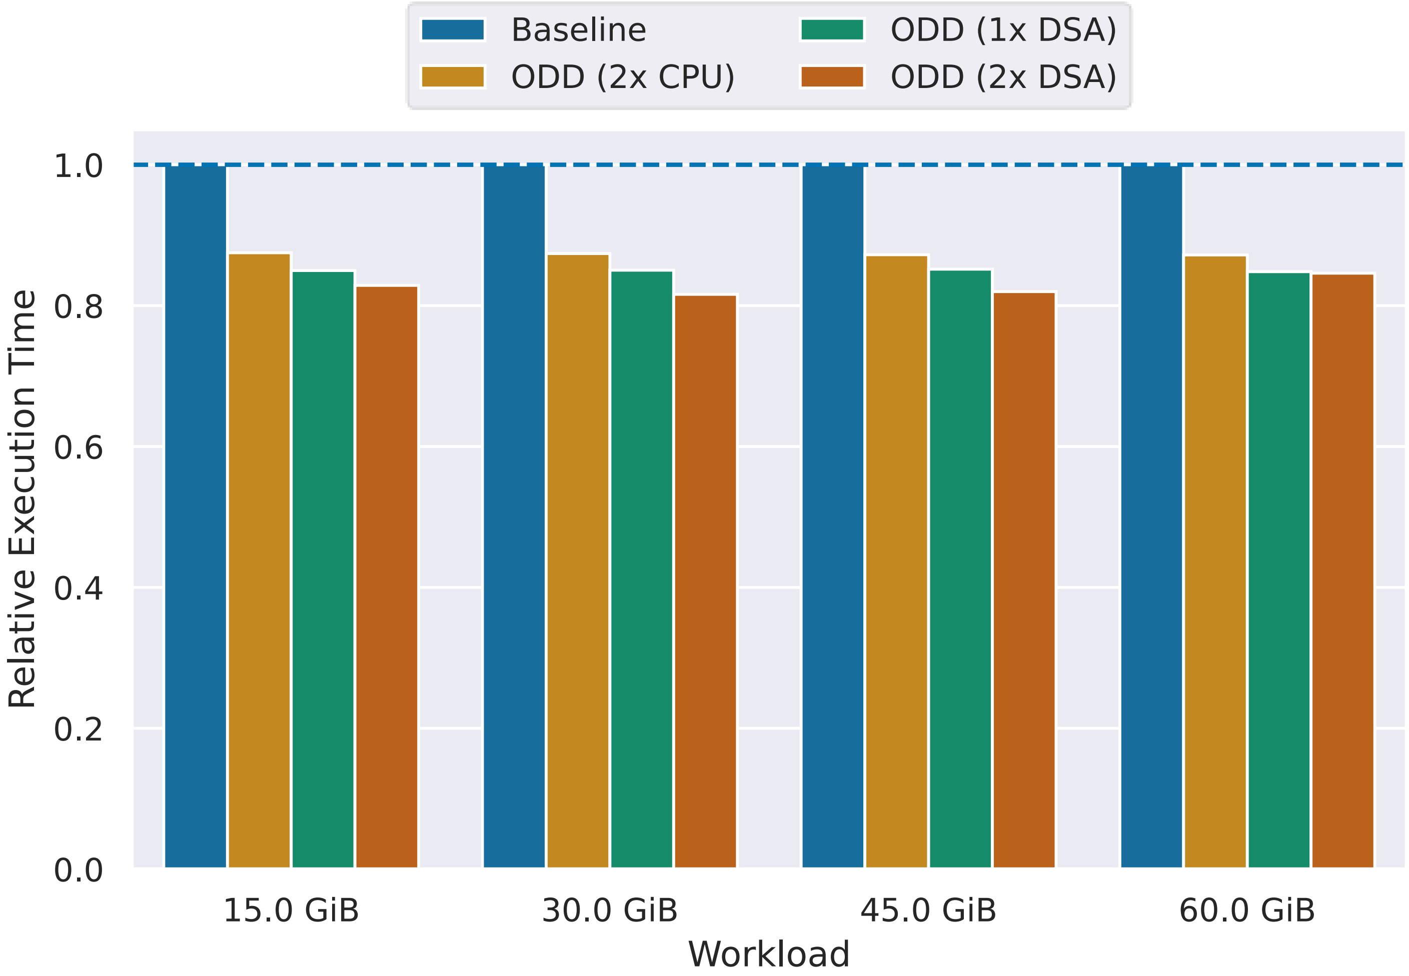

Funded by
DFG
DRESDEN concept

On-the-fly Data Distribution with DSA - Benchmark

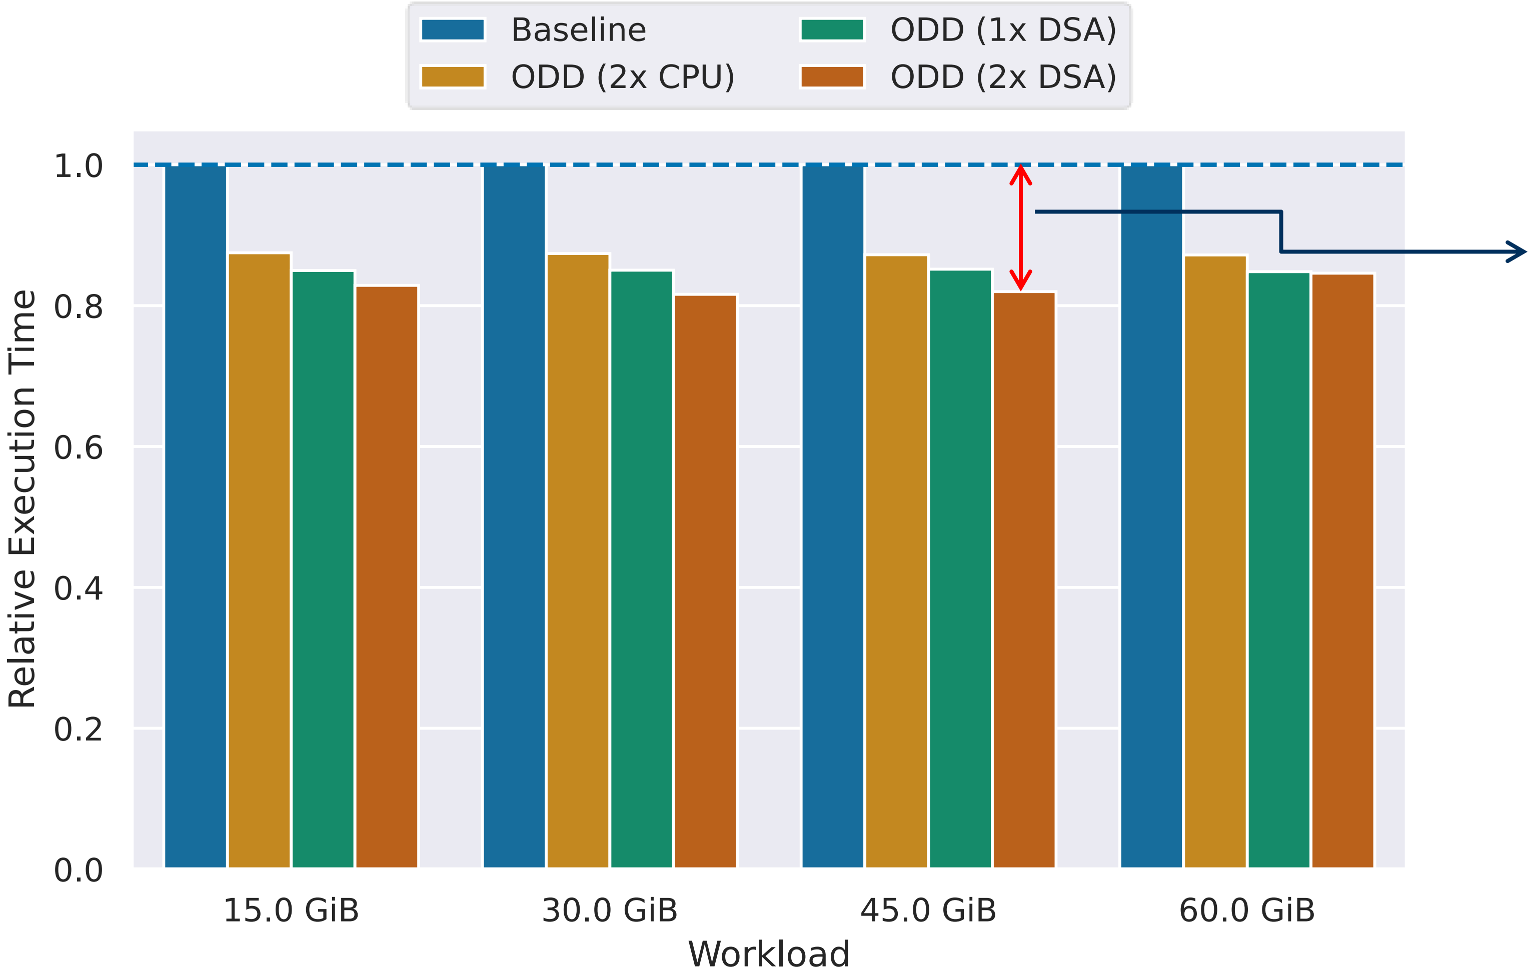

18% performance improvement compared to the **baseline**

Funded by DFG
DRESDEN concept

On-the-fly Data Distribution with DSA - Benchmark

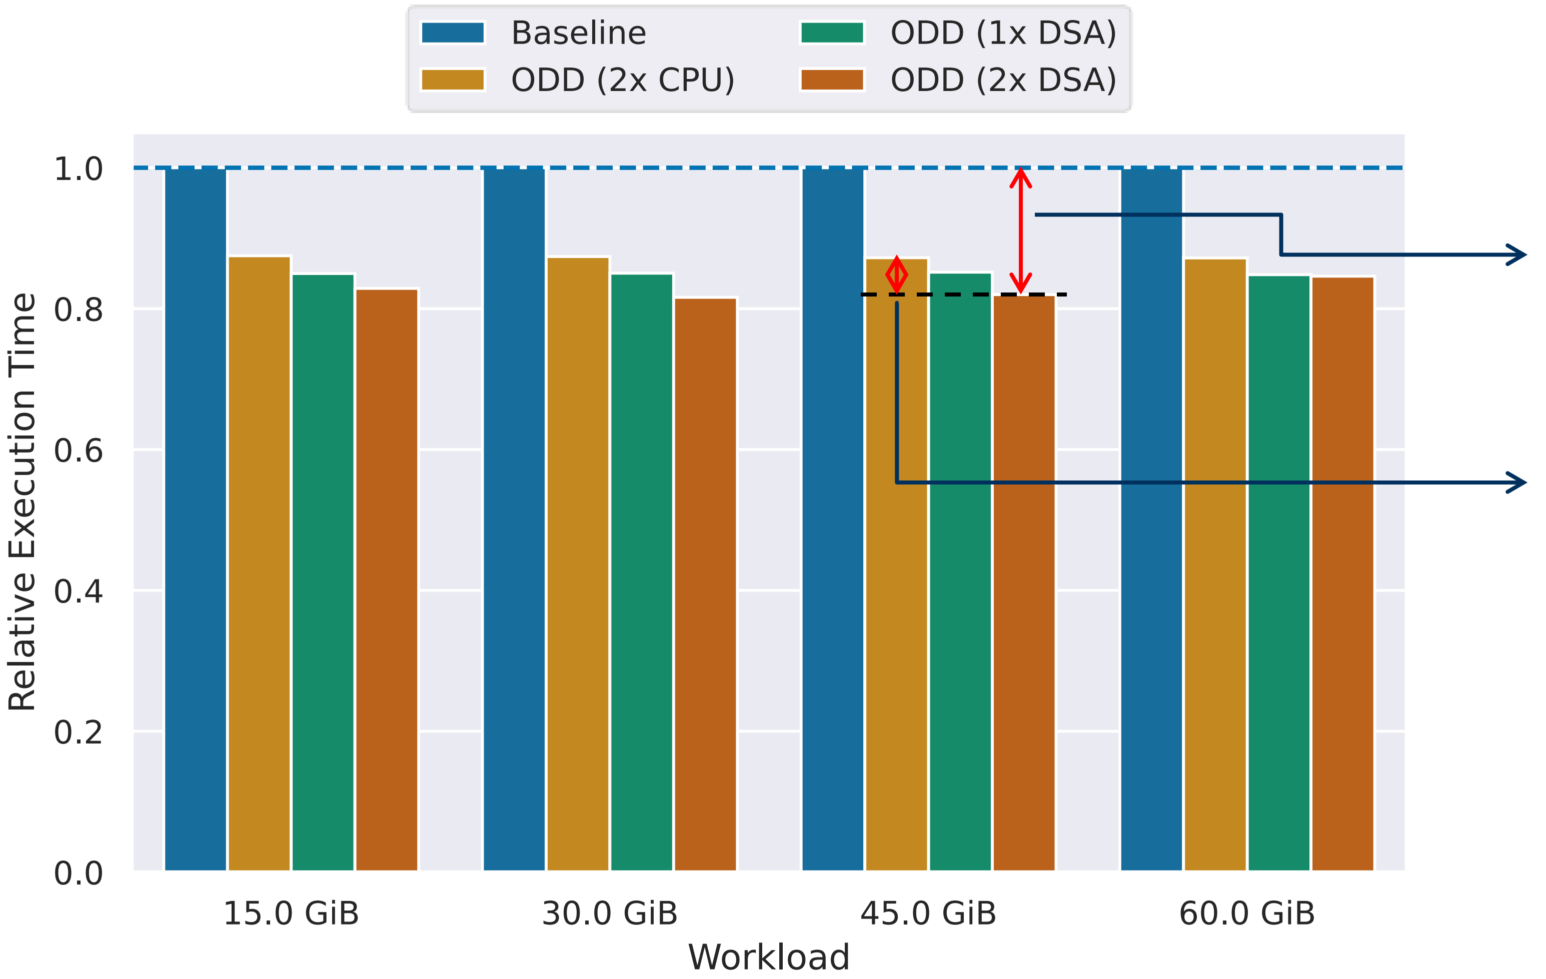

18% performance improvement compared to the **baseline**

6% performance improvement compared to **CPU-based ODD**

Funded by DFG
DRESDEN concept

Summary

DSA allows memory operation offloading …

Funded by
DFG
DRESDEN concept

Summary

DSA allows memory operation offloading …

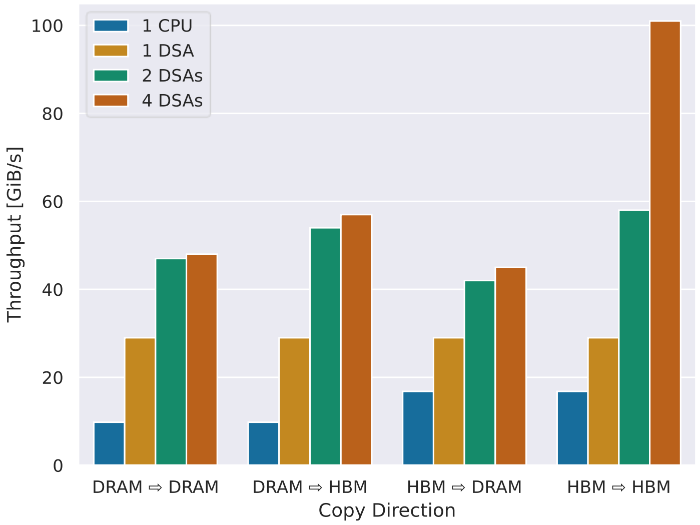

… with **higher throughput** compared to one CPU.

Funded by DFG

DRESDEN concept

Summary

DSA allows memory operation offloading …

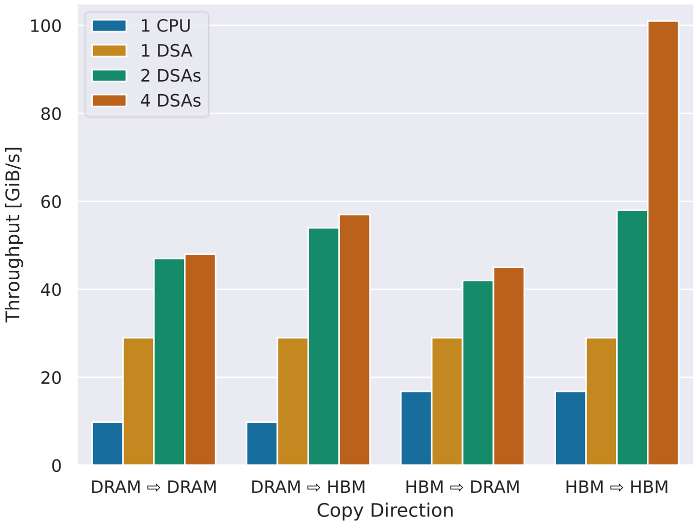

… with **higher throughput** compared to one CPU.

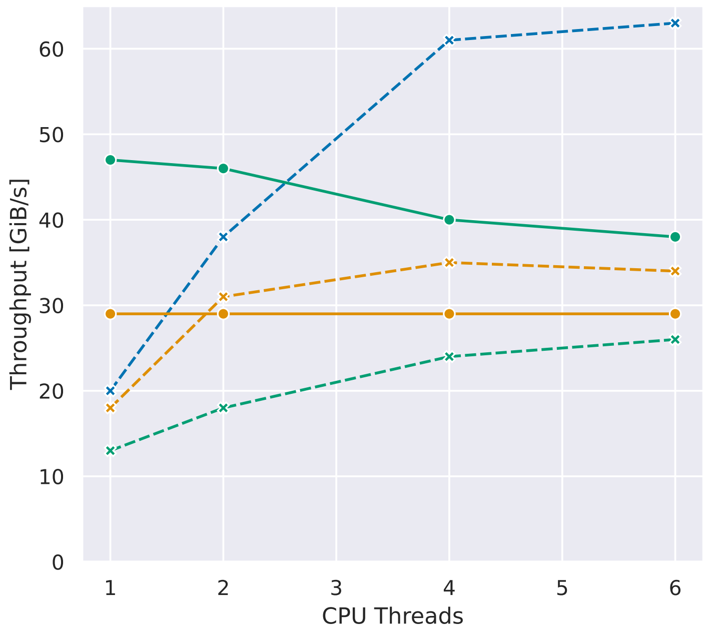

… with **high-priority memory access**

Funded by DFG
DRESDEN concept

Summary

DSA allows memory operation offloading …

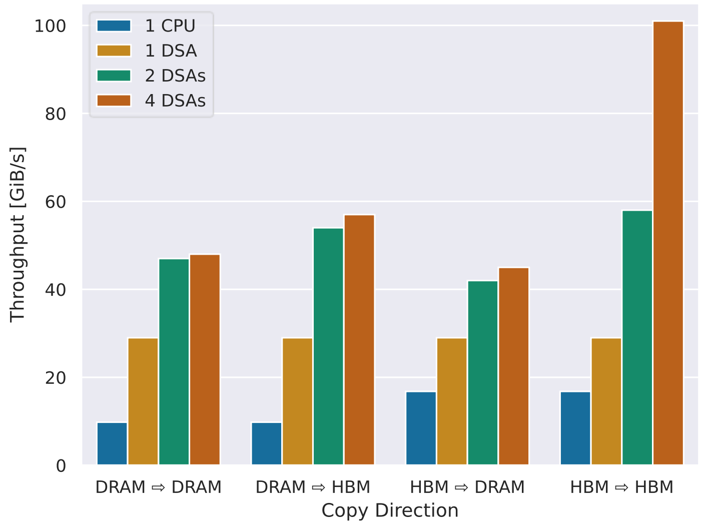

… with **higher throughput** compared to one CPU.

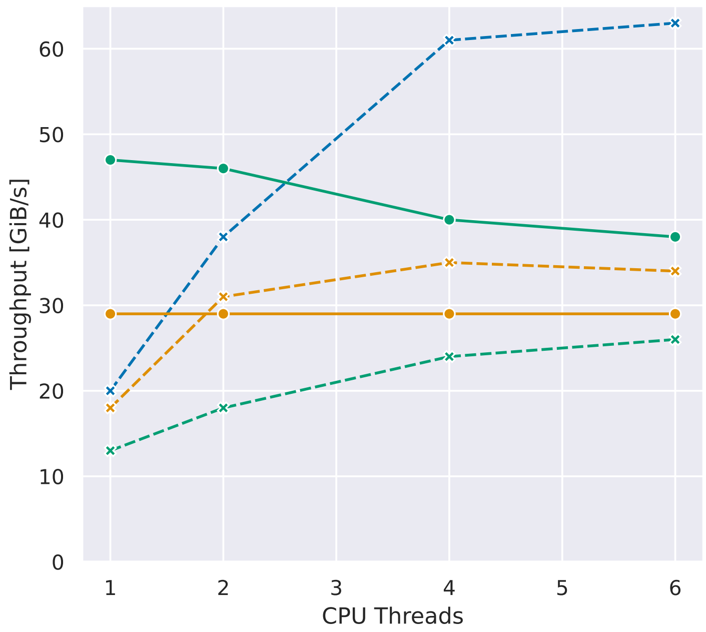

… with **high-priority memory access**

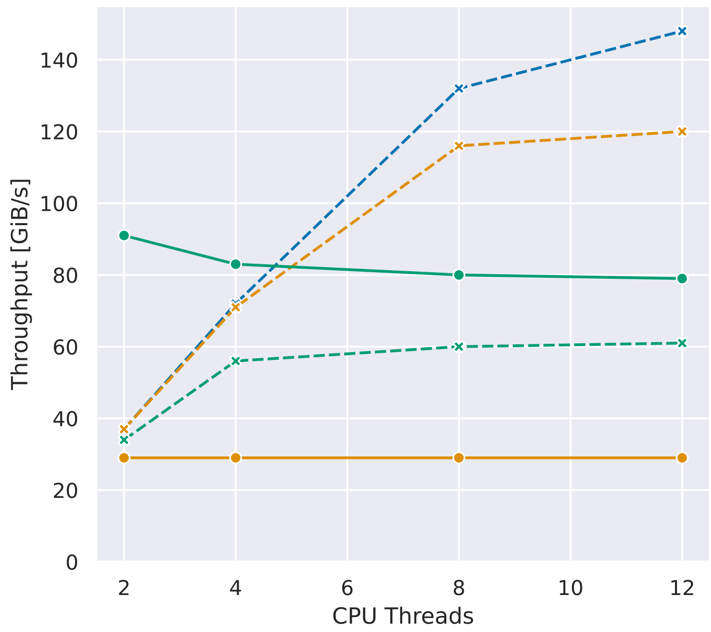

… where DSA interferes with CPU,

but CPU has **minor influence on DSAs** throughput.

Funded by DFG DRESDEN concept

Funded by

Deutsche Forschungsgemeinschaft

German Research Foundation

André Berthold,
Constantin Fürst, Antonia Obersteiner, Lennart Schmidt, Dirk Habich, Wolfgang Lehner, Horst Schirmeier

Dresden University of Technology, Faculty of Computer Science, Institute of Systems Architecture

Demystifying Intel Data Streaming Accelerator for In-Memory Data Processing

Austin, Texas, DIMES 2nd Workshop on Disruptive Memory Systems // November 3, 2024