

PERFORMANCE ANALYSIS OF ARRAY DATABASE SYSTEMS IN NON-UNIFORM MEMORY ARCHITECTURE

EDUARDO C. DE ALMEIDA¹, JORGE A. MEIRA², MARCO A. Z. ALVES¹, SIMONE DOMINICO¹

FEDERAL UNIVERSITY OF PARANÁ, BRAZIL¹, UNIVERSITY OF LUXEMBOURG²



AGENDA

- ▶ **Introduction**

- ▶ Array database systems

- ▶ NUMA architecture

- ▶ Array Database Systems in NUMA architecture

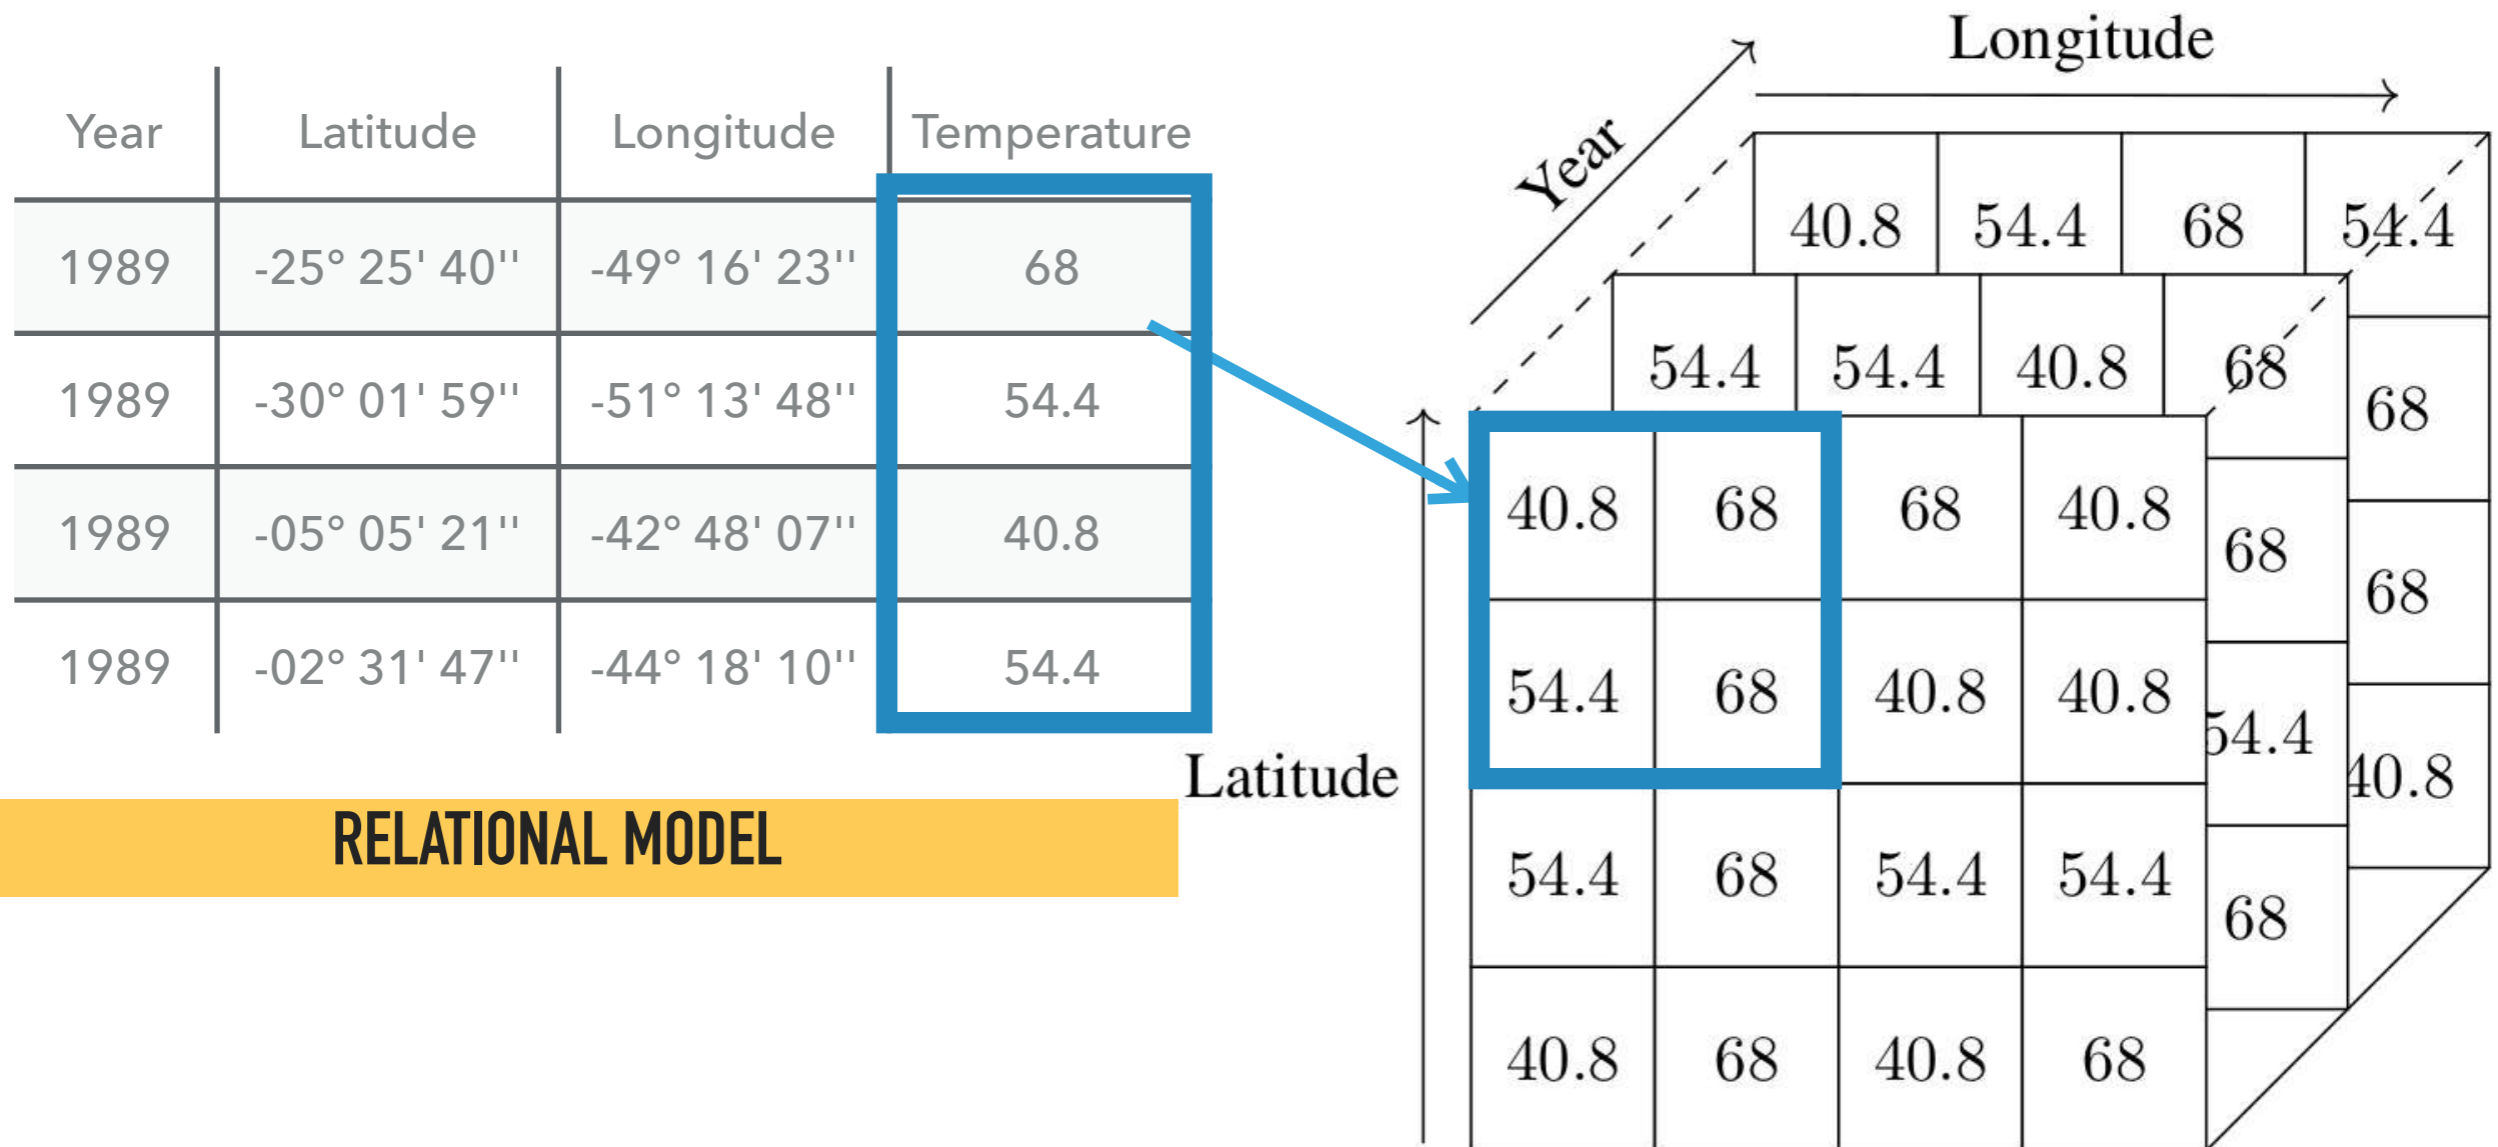
- ▶ Methodology

- ▶ Thread pinning Strategies


- ▶ Experiments

- ▶ Conclusion

ARRAY DATABASE AND RELATIONAL MODEL



ARRAY DATABASE SYSTEMS



AGENDA

- ▶ **Introduction**

- ▶ Array database systems
- ▶ NUMA architecture
- ▶ Array Database Systems in NUMA architecture

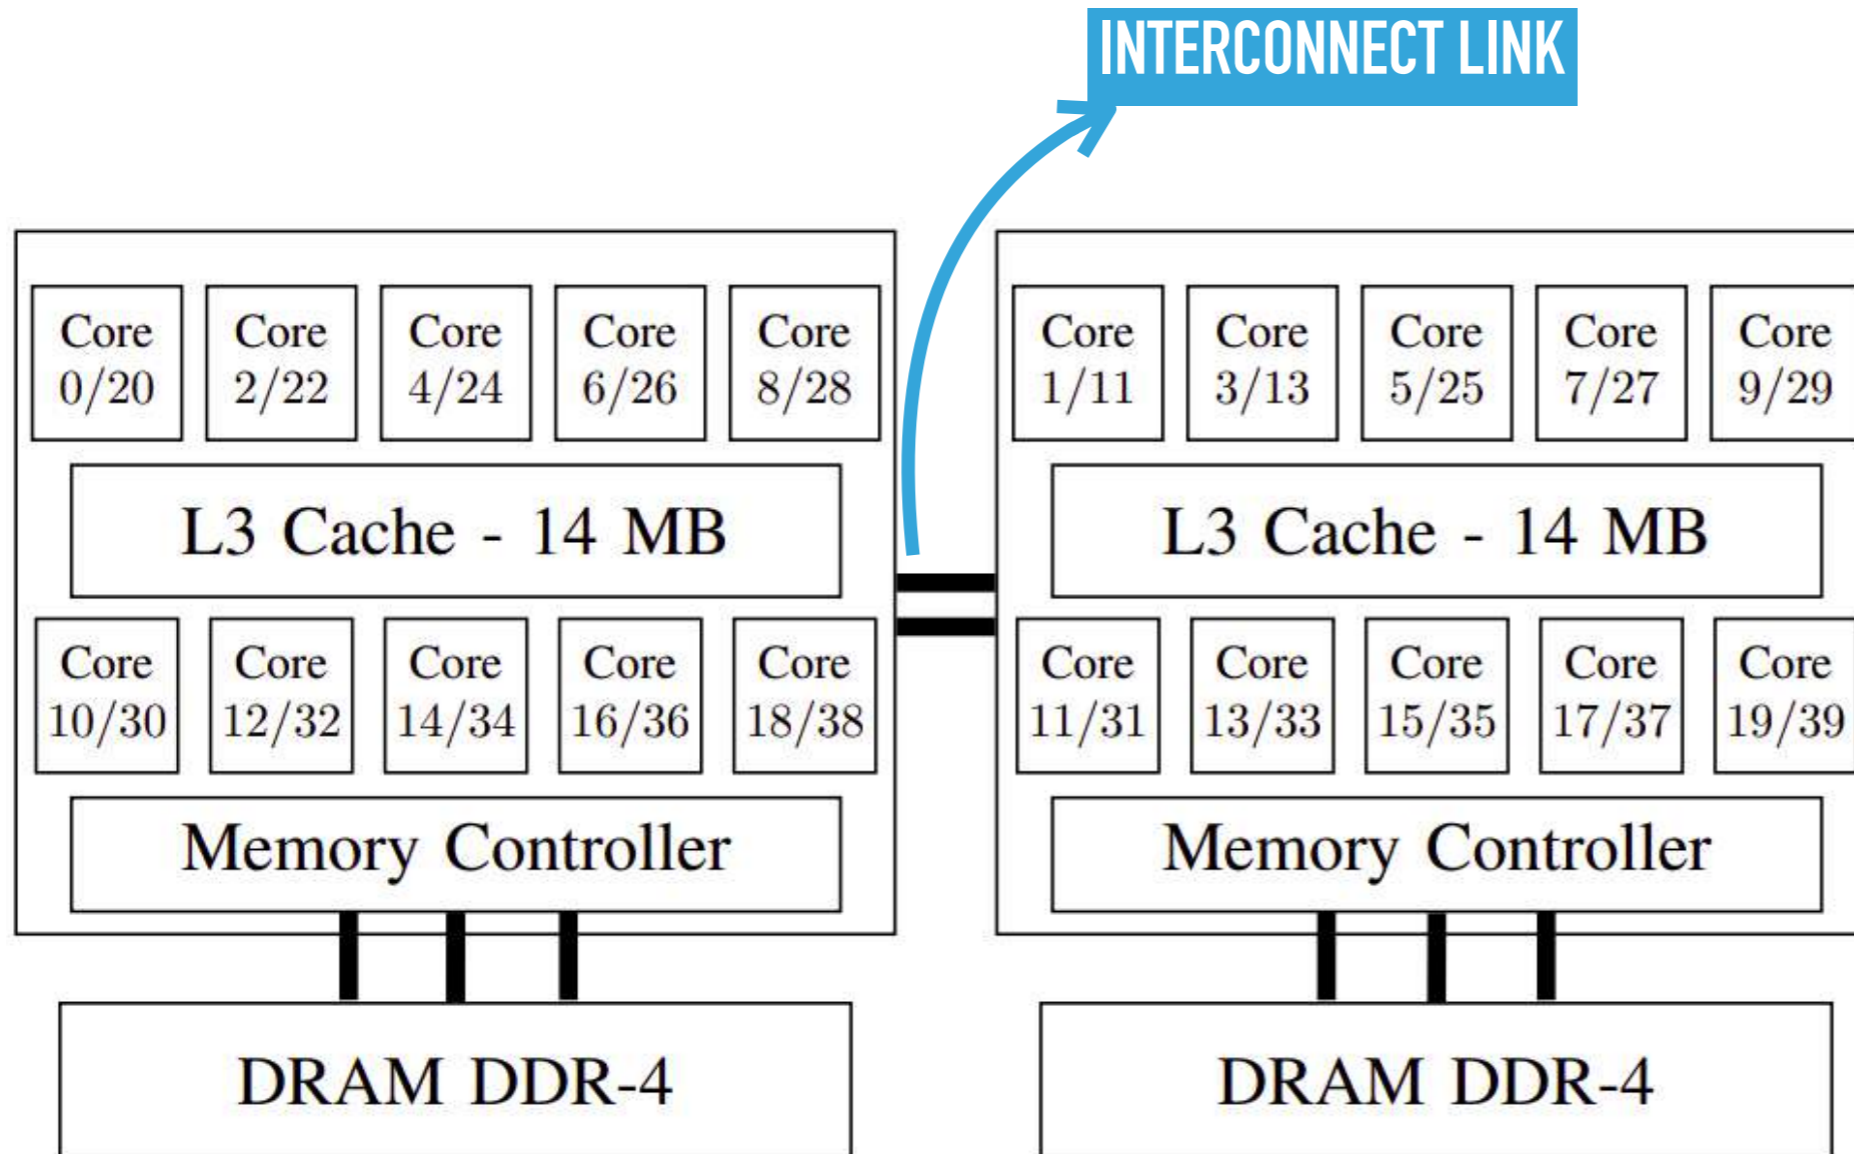
- ▶ Methodology

- ▶ Thread pinning Strategies

- ▶ Experiments

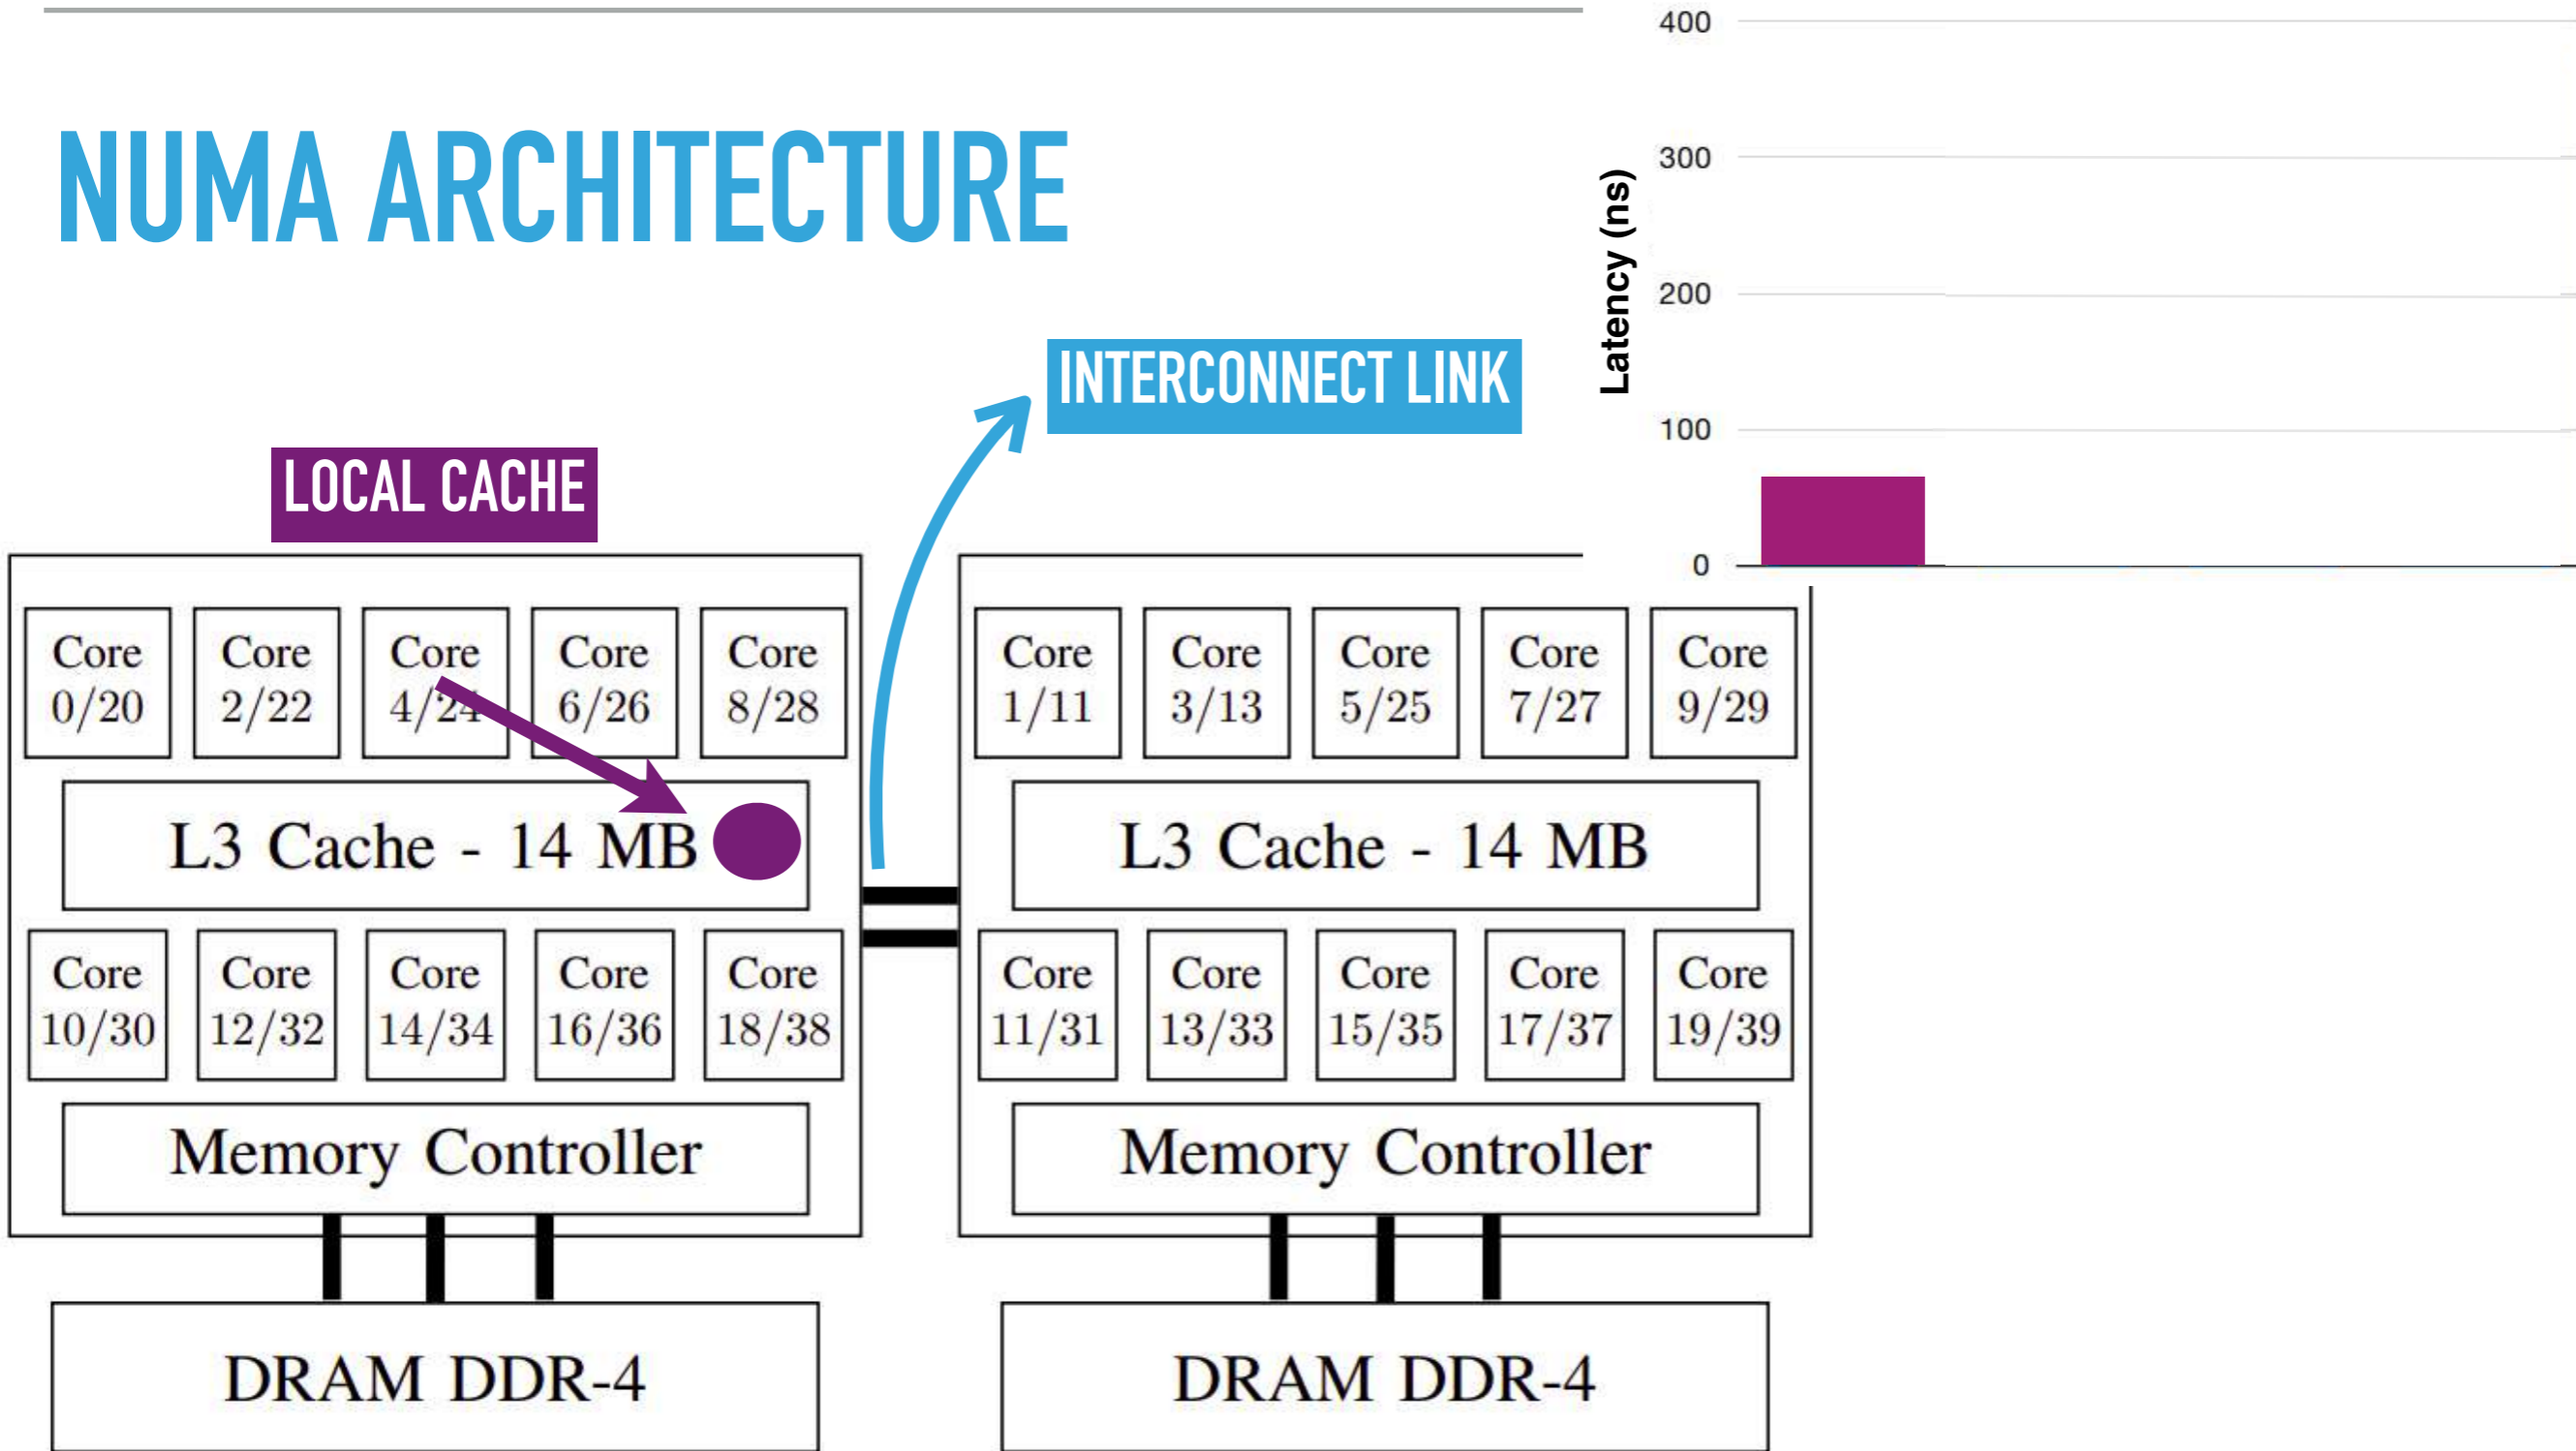
- ▶ Conclusion

NUMA ARCHITECTURE




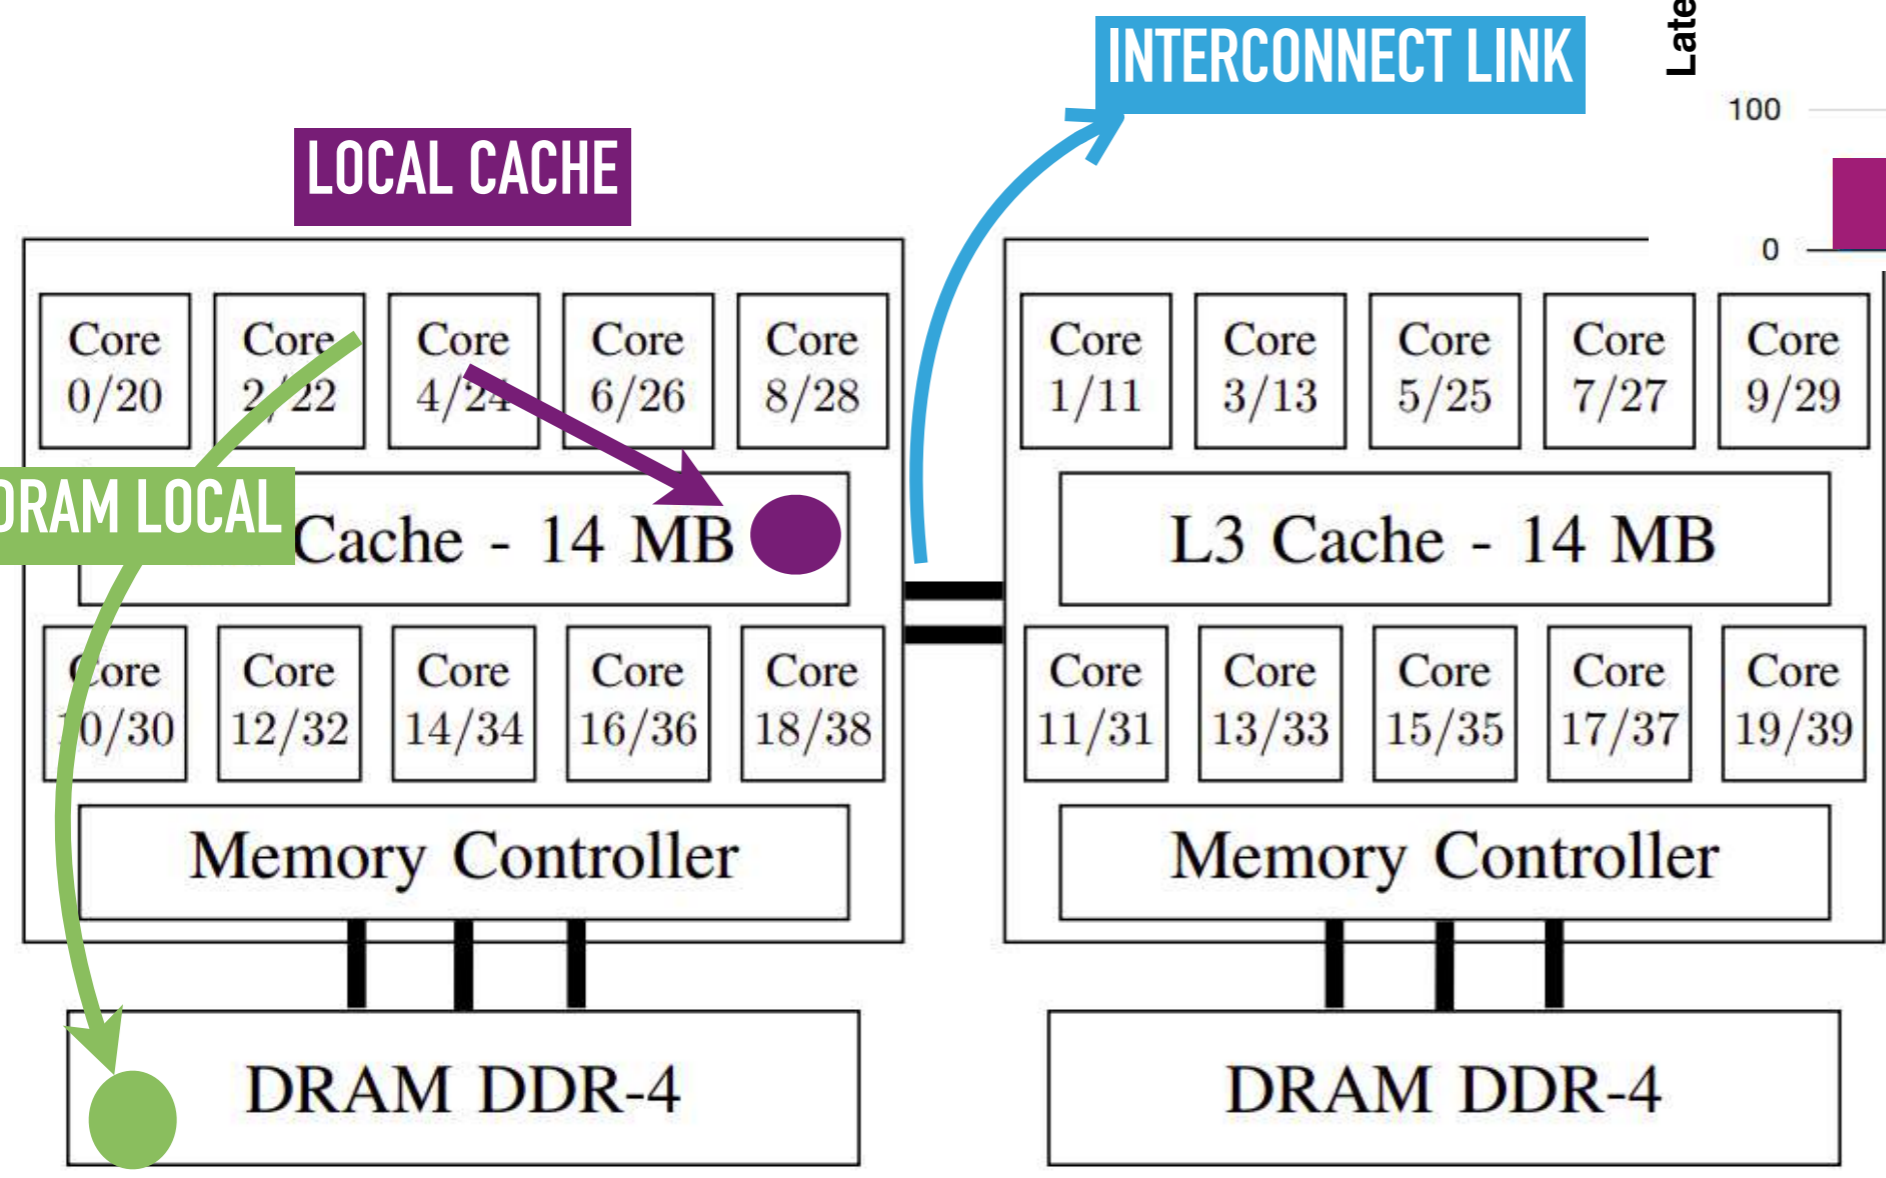
EXAMPLE OF A NUMA ARCHITECTURE WITH 2-NODES INSPIRED ON INTEL XEON SILVER 4114.

NUMA ARCHITECTURE




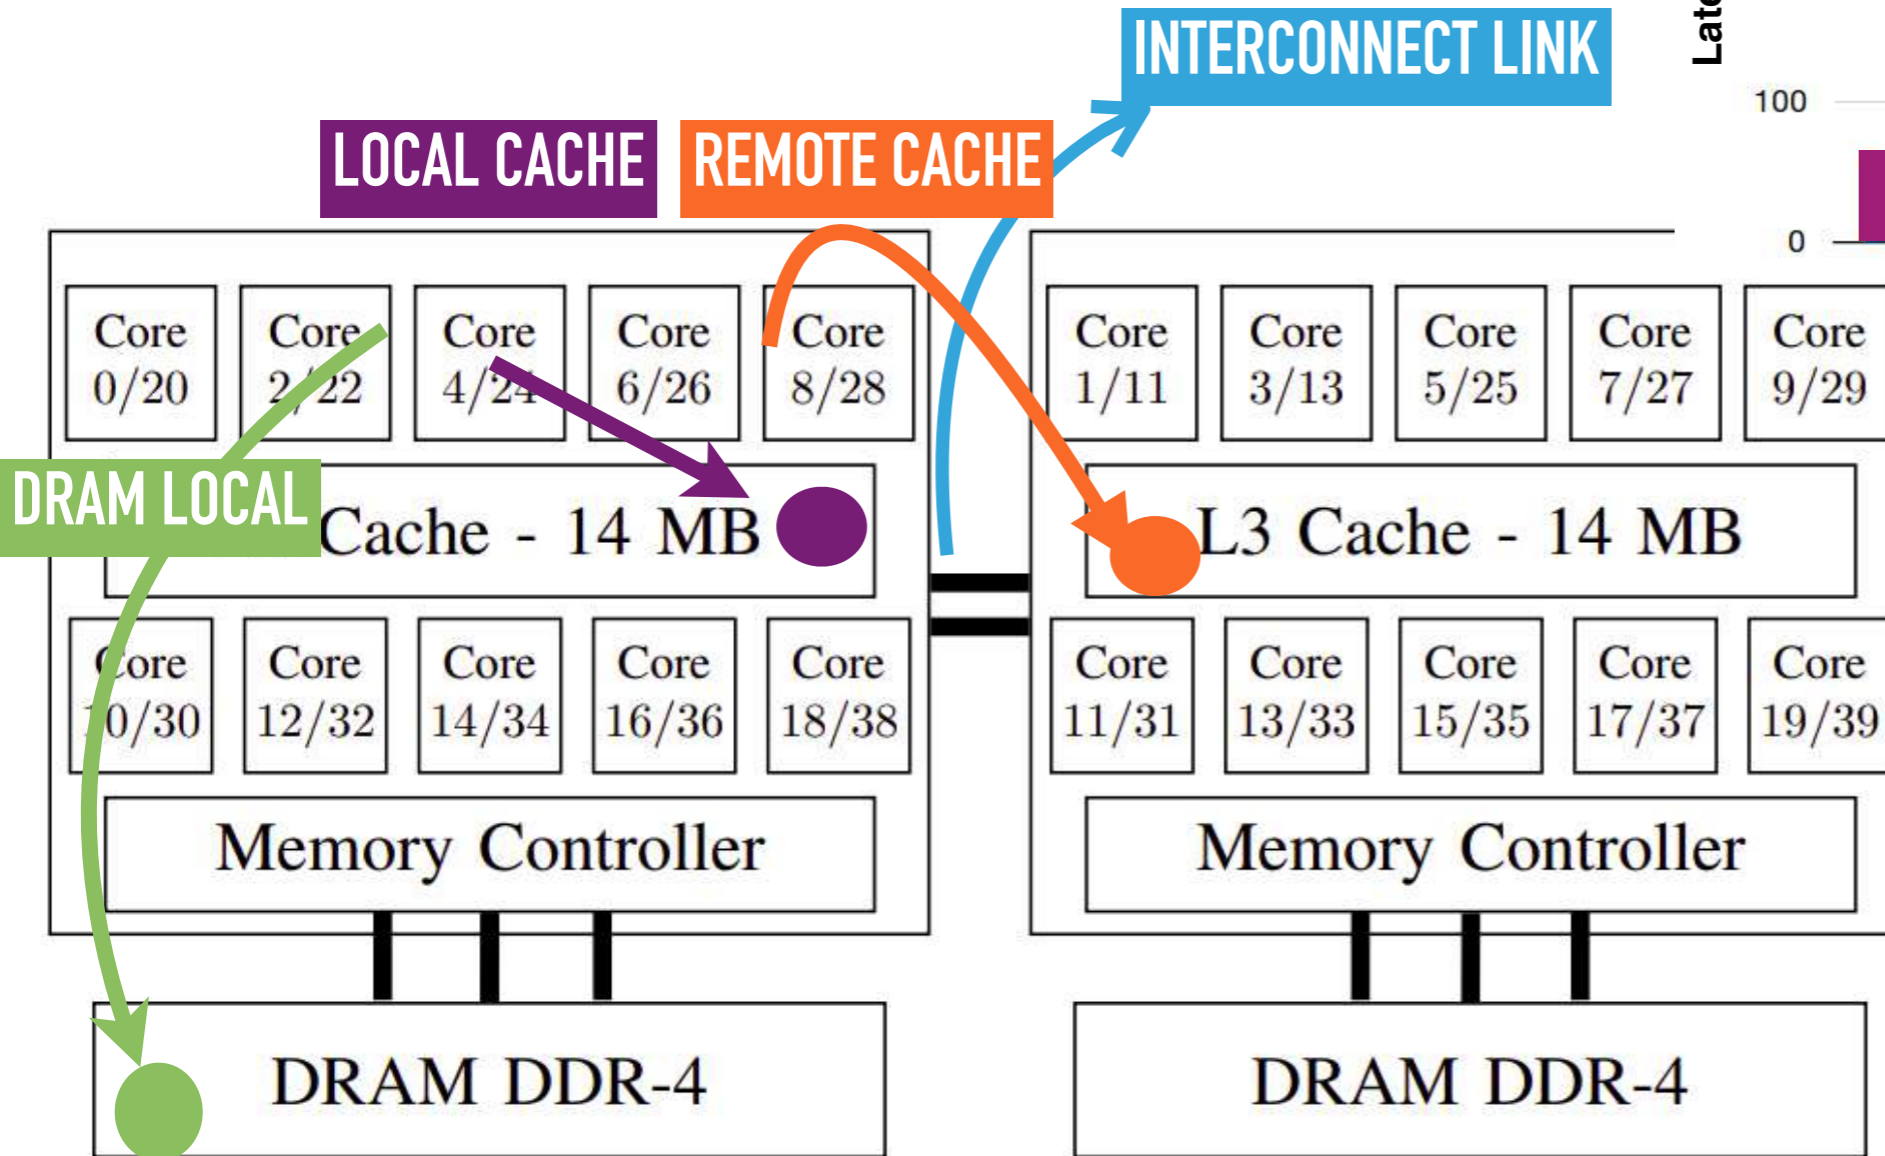
EXAMPLE OF A NUMA ARCHITECTURE WITH 2-NODES INSPIRED ON INTEL XEON SILVER 4114.

NUMA ARCHITECTURE




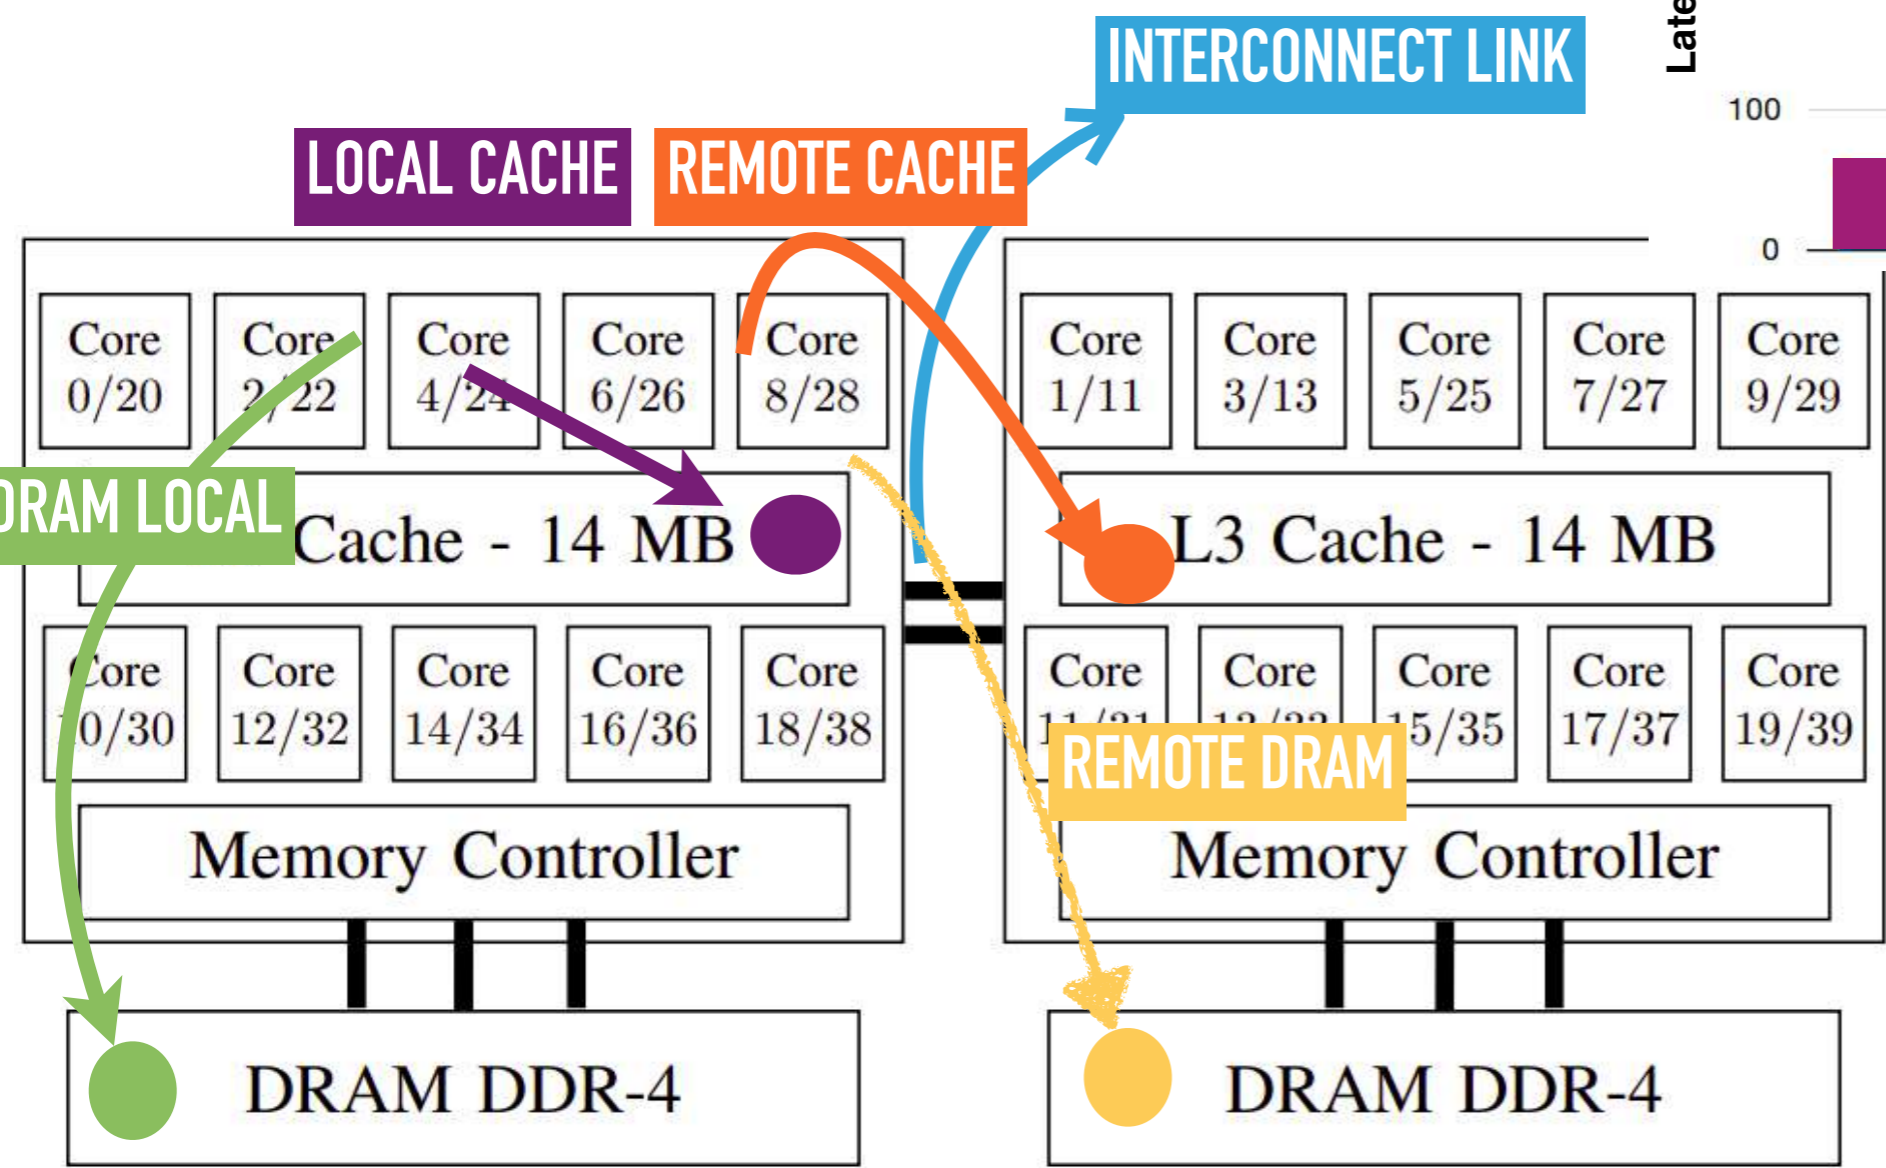
EXAMPLE OF A NUMA ARCHITECTURE WITH 2-NODES INSPIRED ON INTEL XEON SILVER 4114.

NUMA ARCHITECTURE



EXAMPLE OF A NUMA ARCHITECTURE WITH 2-NODES INSPIRED ON INTEL XEON SILVER 4114.

NUMA ARCHITECTURE



EXEMPLO DE ARQUITETURA NUMA COM 2-SOCKETS INSPIRADA NO INTEL XEON SILVER 4114

AGENDA

- ▶ **Introduction**

- ▶ Array database systems

- ▶ NUMA architecture

- ▶ Array Database Systems in NUMA architecture

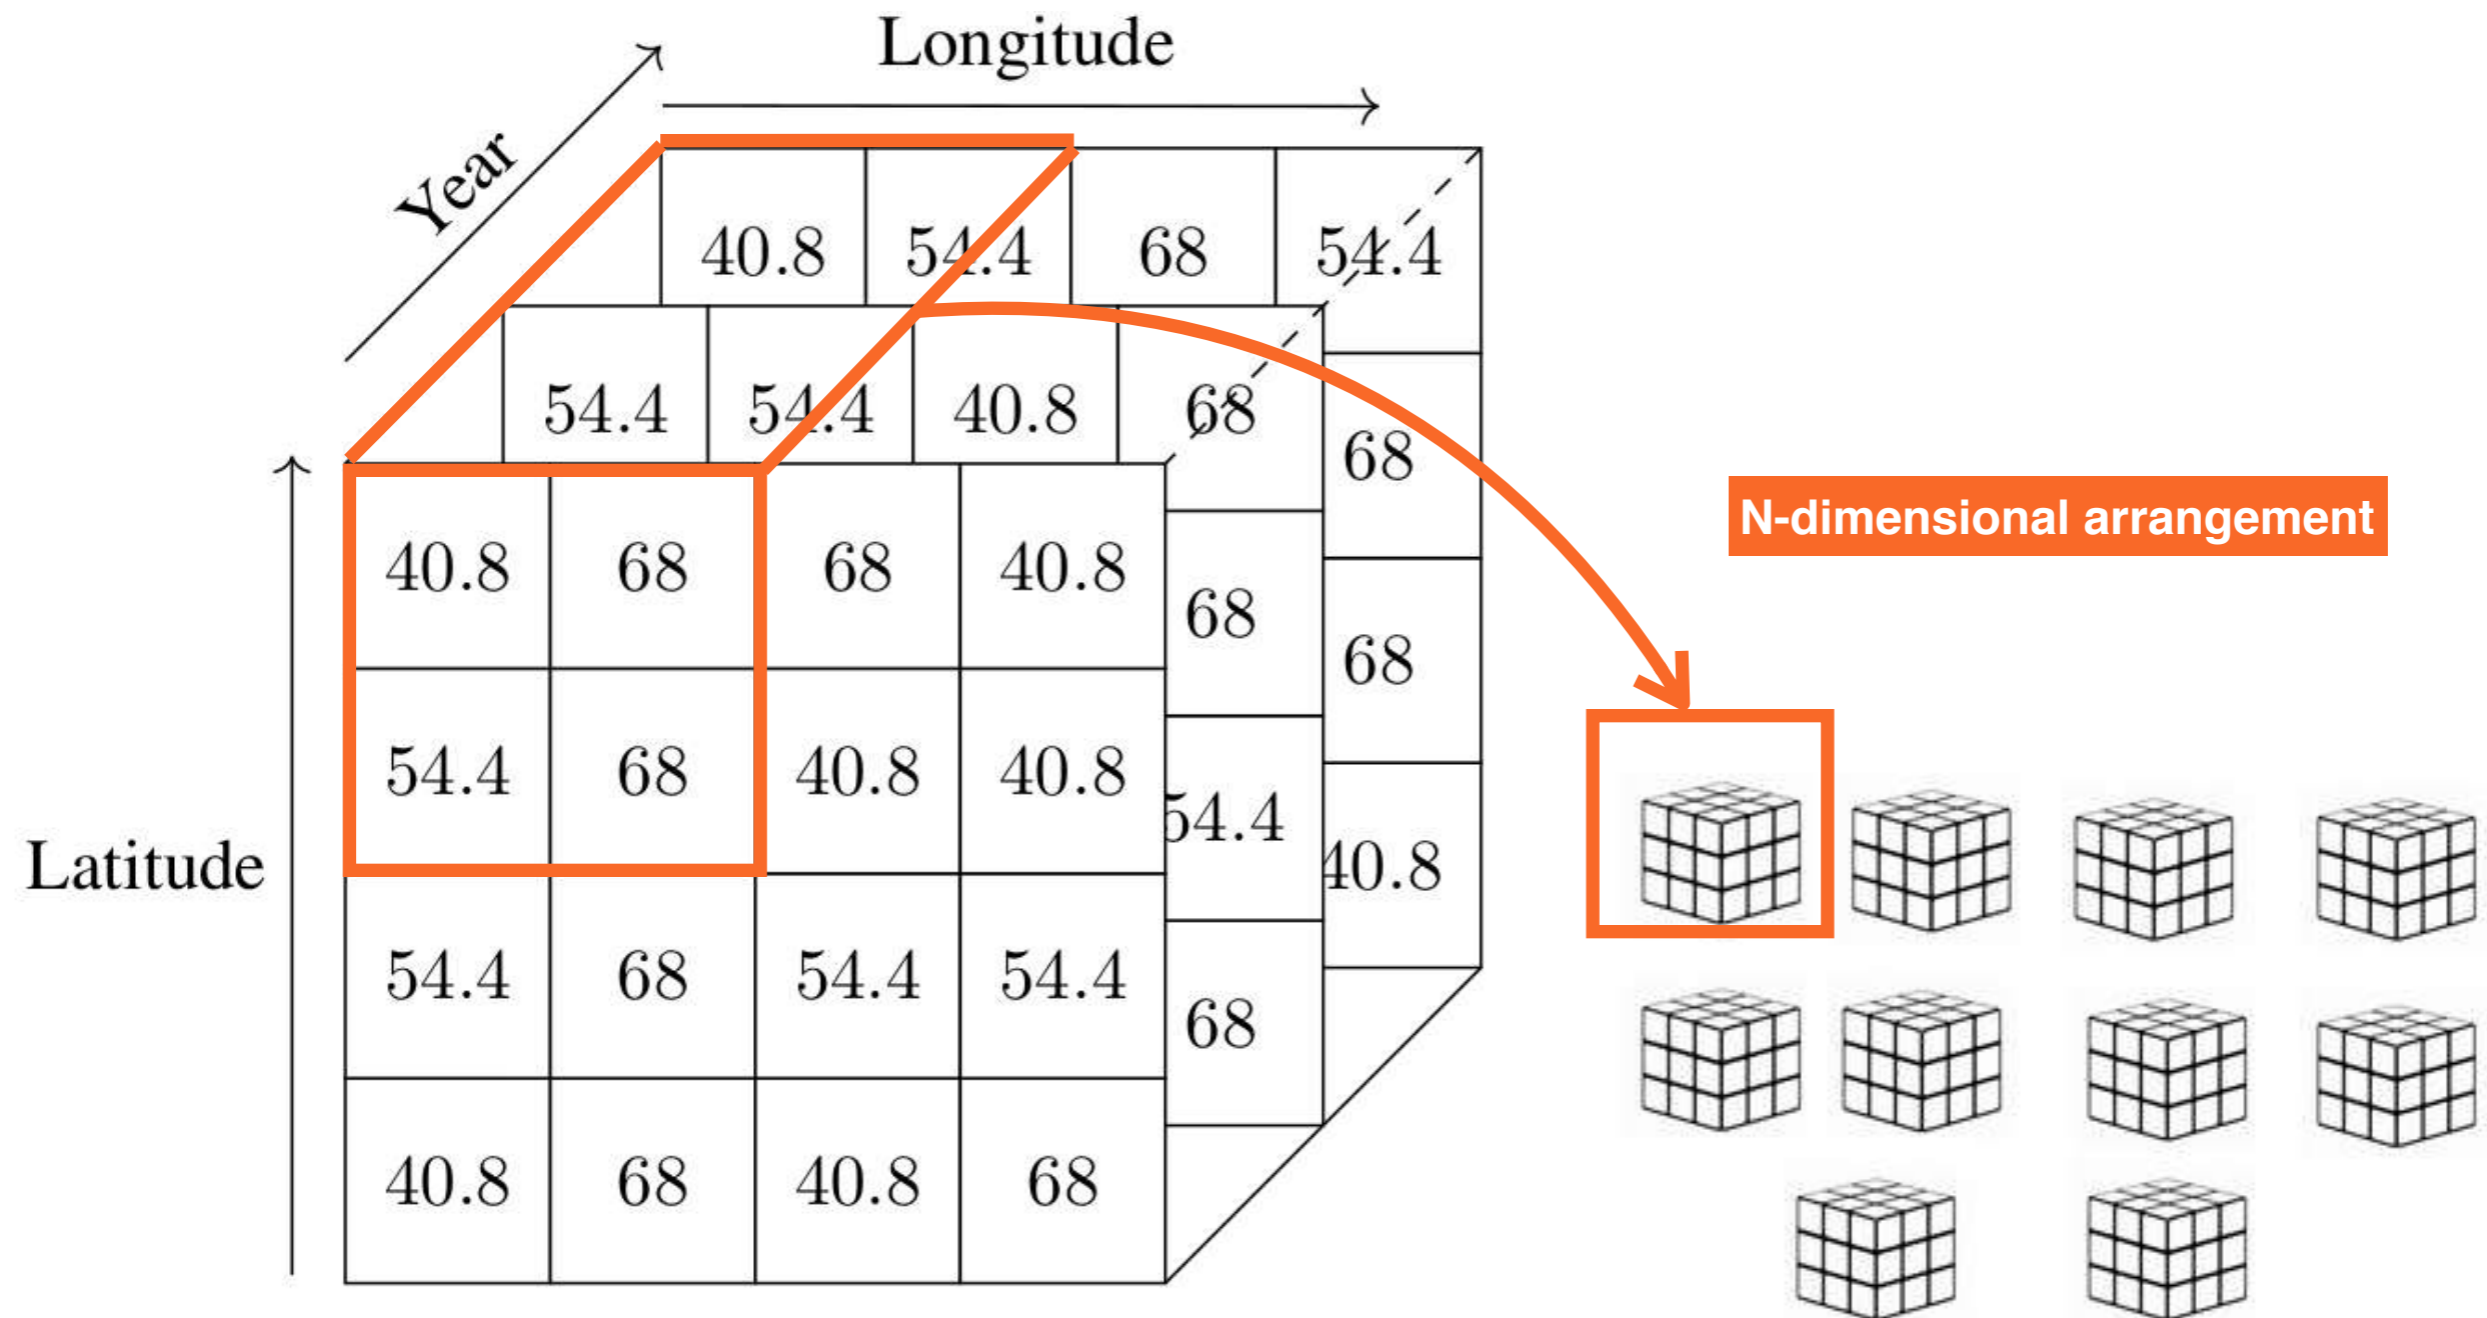
- ▶ Methodology

- ▶ Thread pinning Strategies

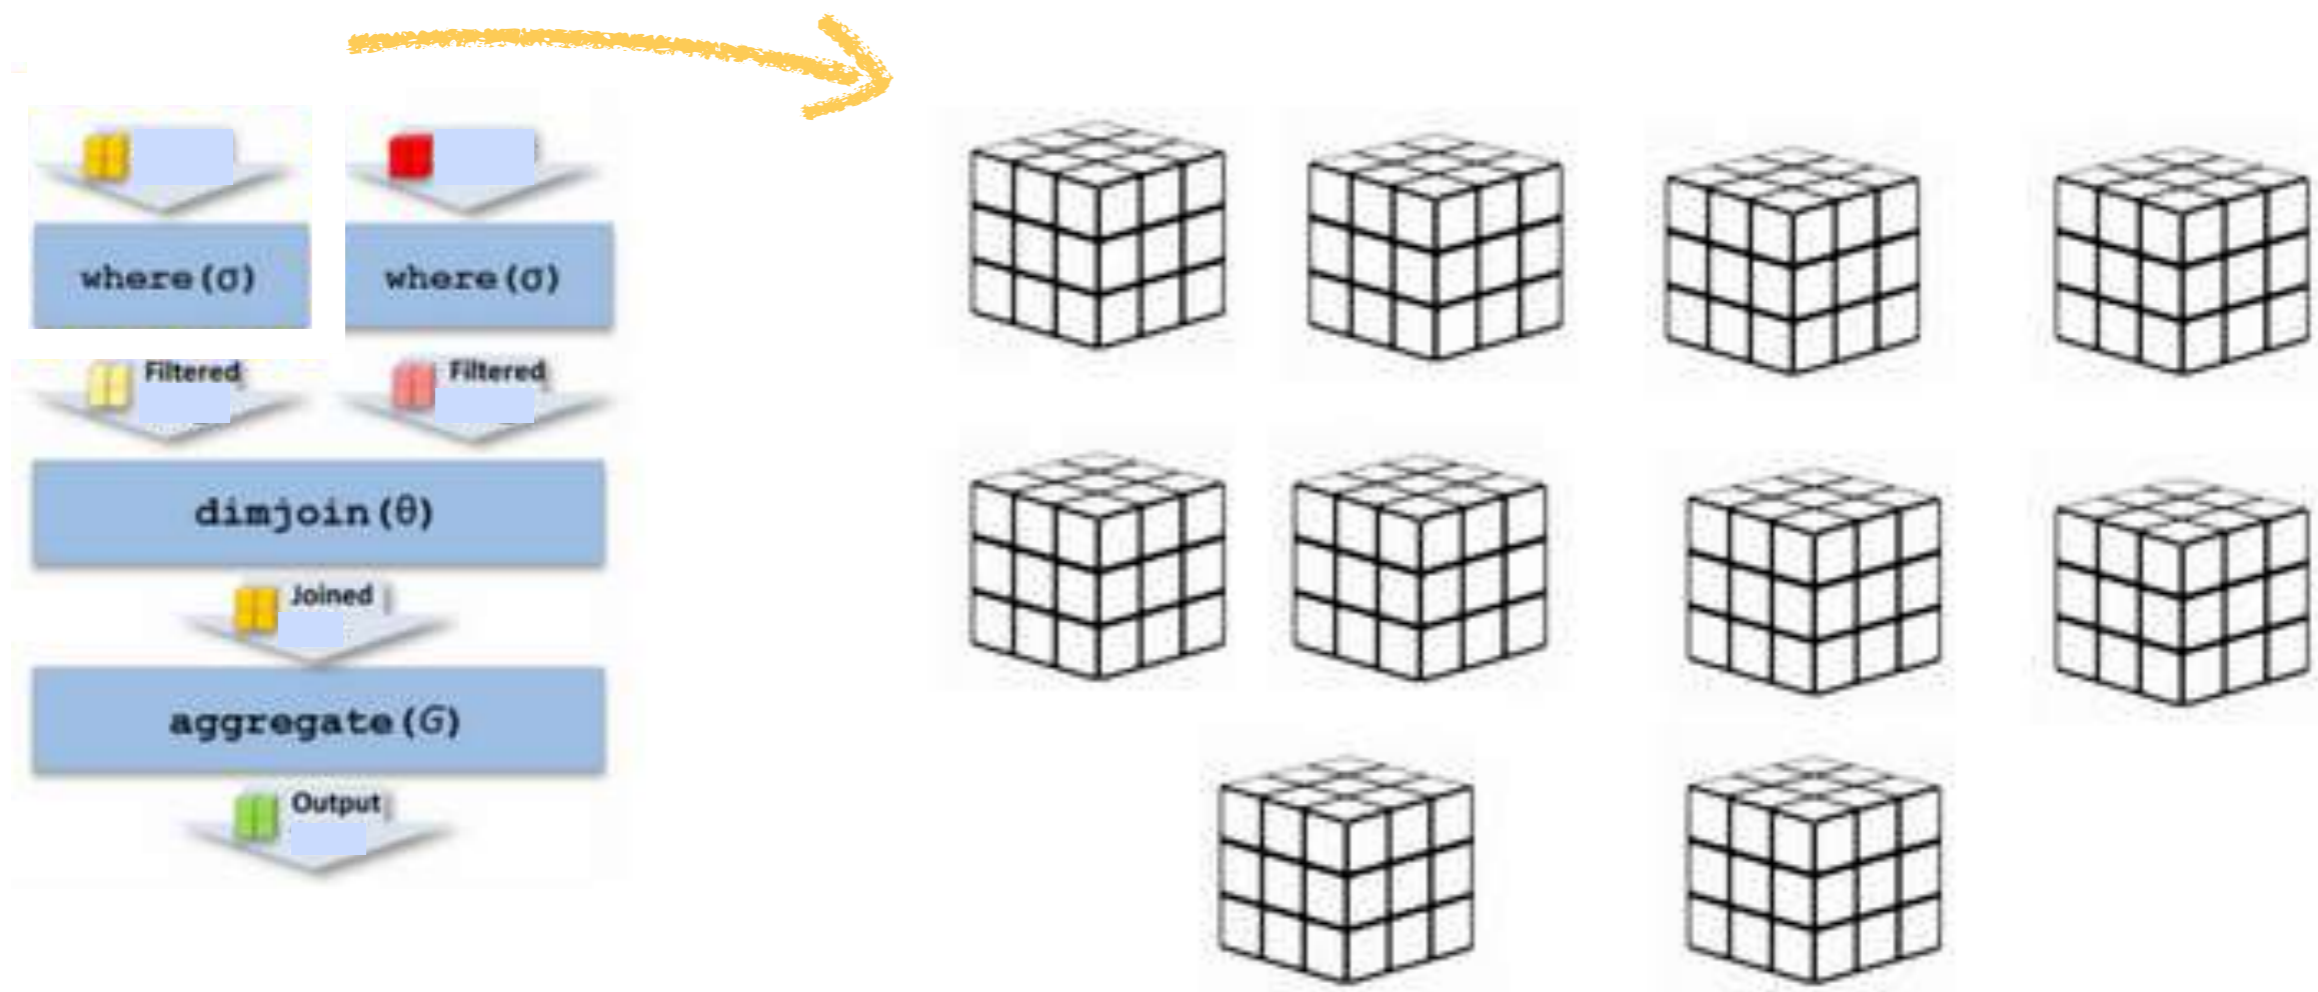
- ▶ Experiments

- ▶ Conclusion

ARRAY STORED AS A SUBARRAY (CHUNK)



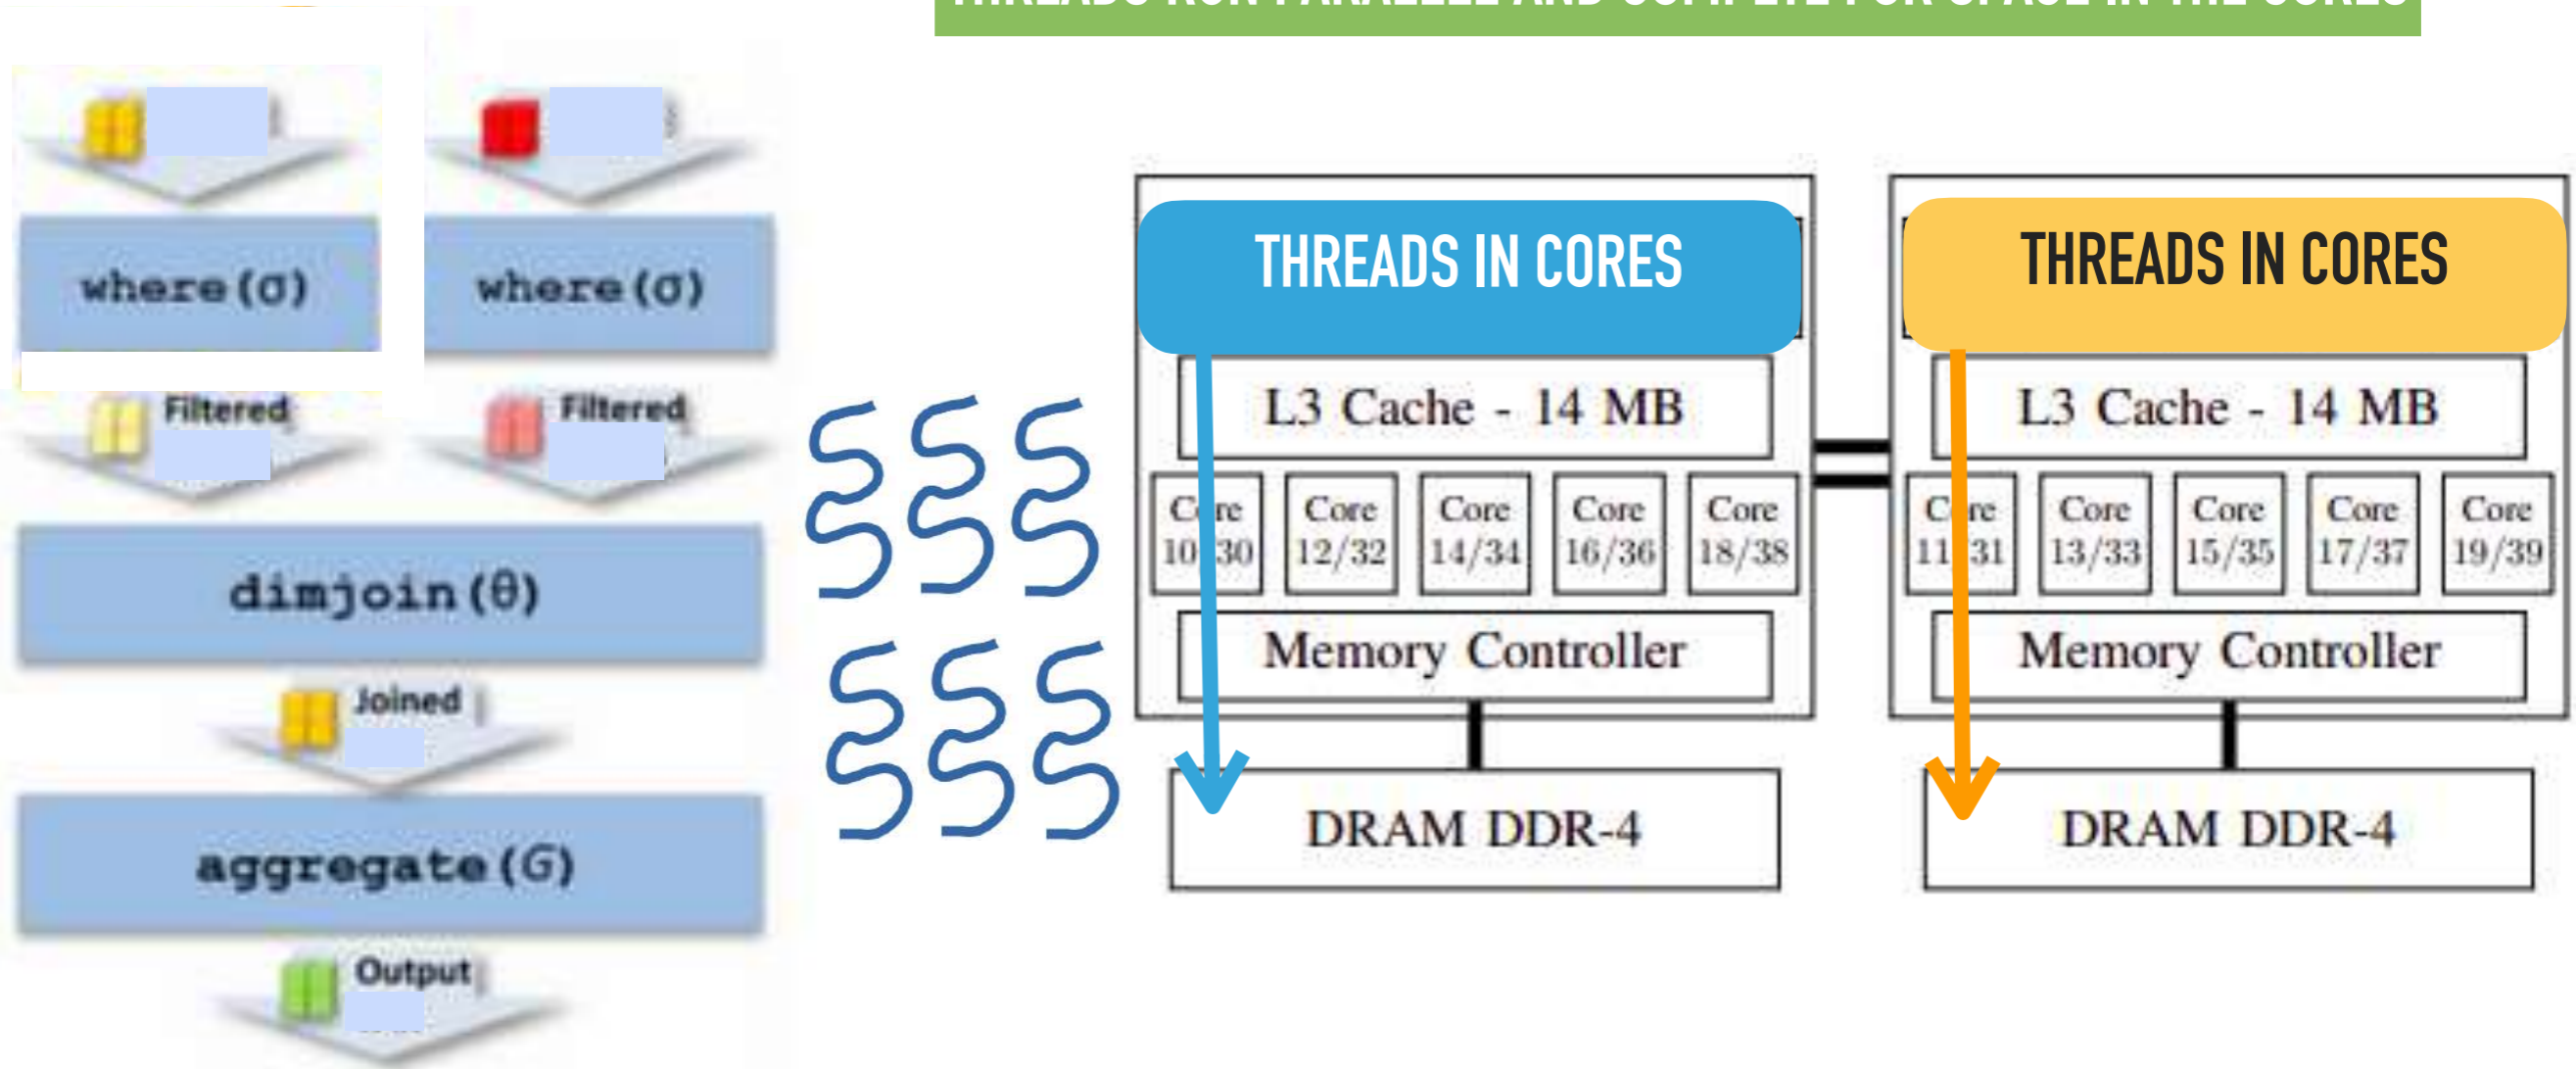
PROCESSING MODEL



EACH DATABASE OPERATION PROCESSES SUBARRAY WITH MULTIPLE THREADS

THREAD POSITIONING - OPERATION SYSTEM

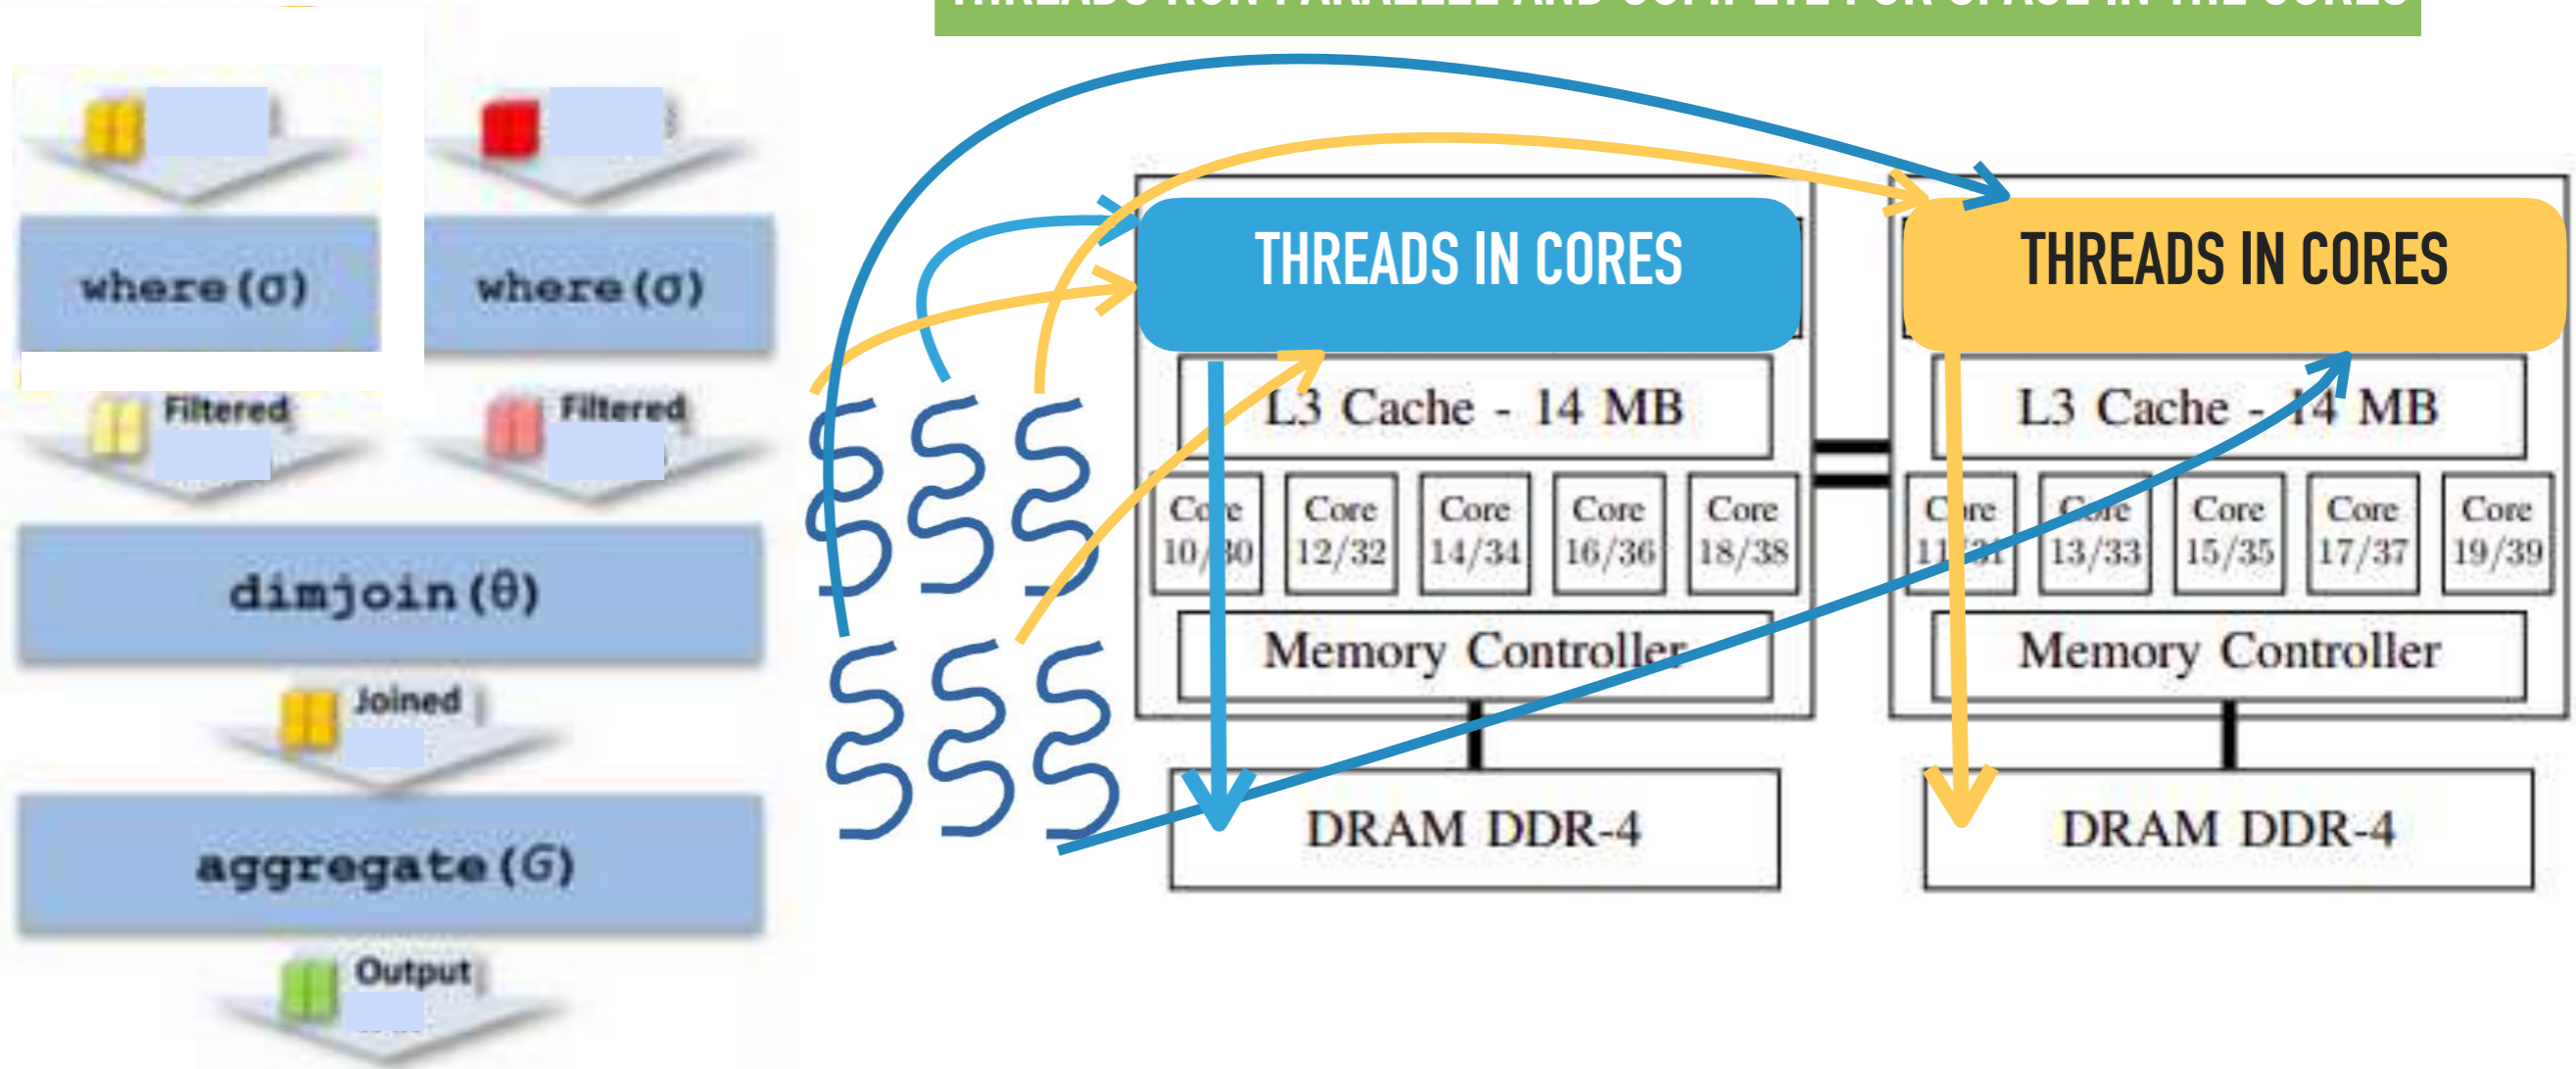
THREADS RUN PARALLEL AND COMPETE FOR SPACE IN THE CORES



MULTIPLE THREADS PROCESS A QUERY OPERATION

THREAD POSITIONING - OPERATION SYSTEM

THREADS RUN PARALLEL AND COMPETE FOR SPACE IN THE CORES



MULTIPLE THREADS PROCESS A QUERY OPERATION

AGENDA

- ▶ Introduction
 - ▶ Array database systems
 - ▶ NUMA architecture
 - ▶ Array Database Systems in NUMA architecture
- ▶ Methodology
 - ▶ Thread pinning Strategies
- ▶ Experiments
- ▶ Conclusion

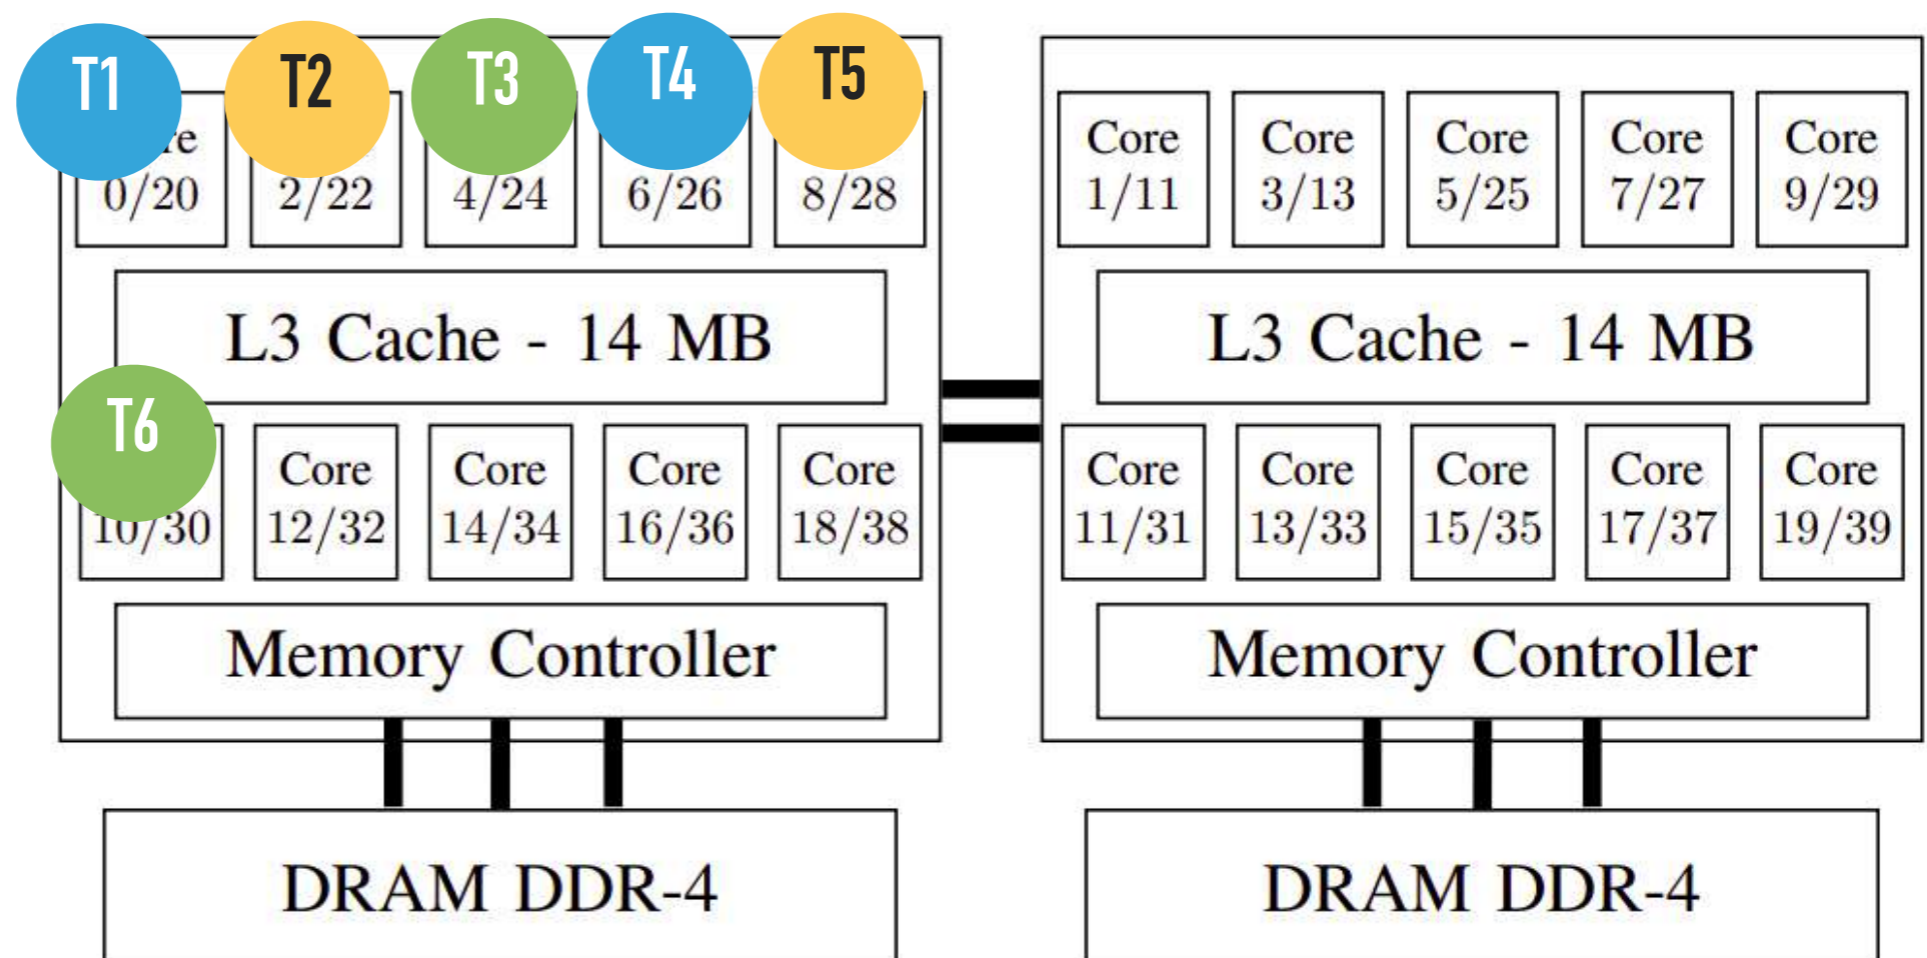
THREAD PINNING STRATEGIES

- ▶ state-of-the-art strategies
- ▶ Petri net
- ▶ Random allocation

THREAD PINNING STRATEGIES

▶ COMPACT

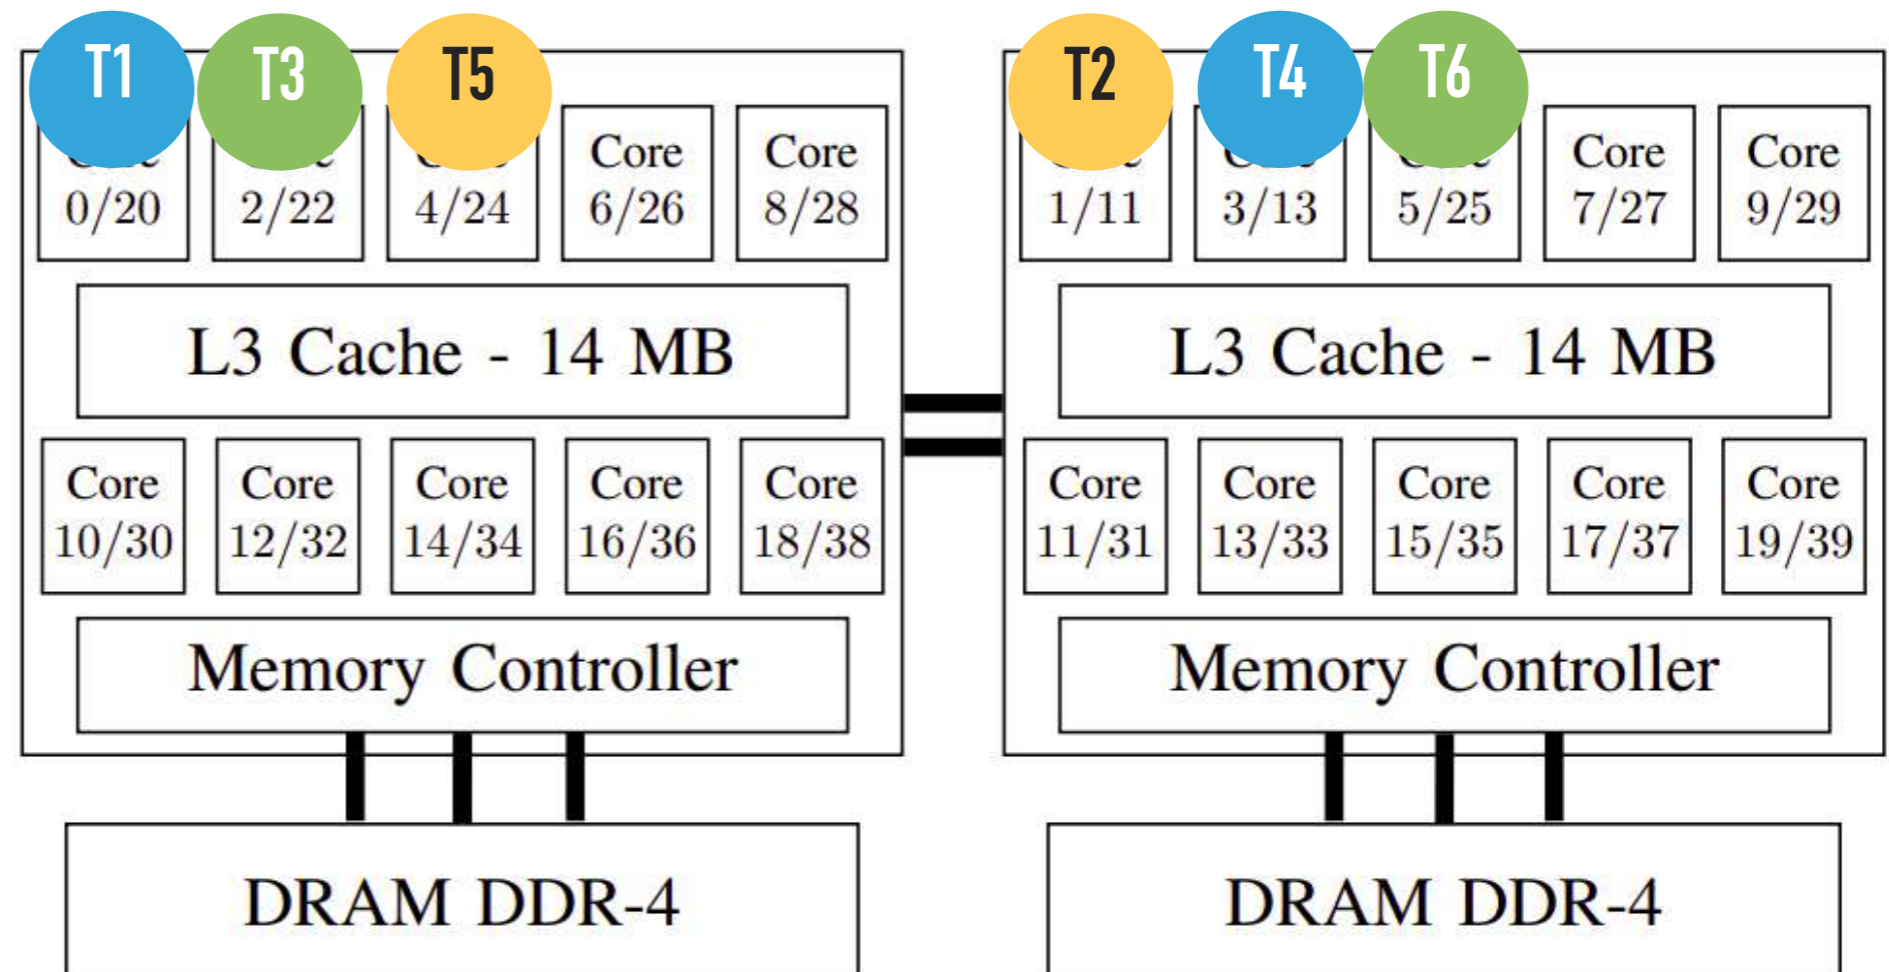
▶ Same Socket



THREAD PINNING STRATEGIES

▶ SPARSE

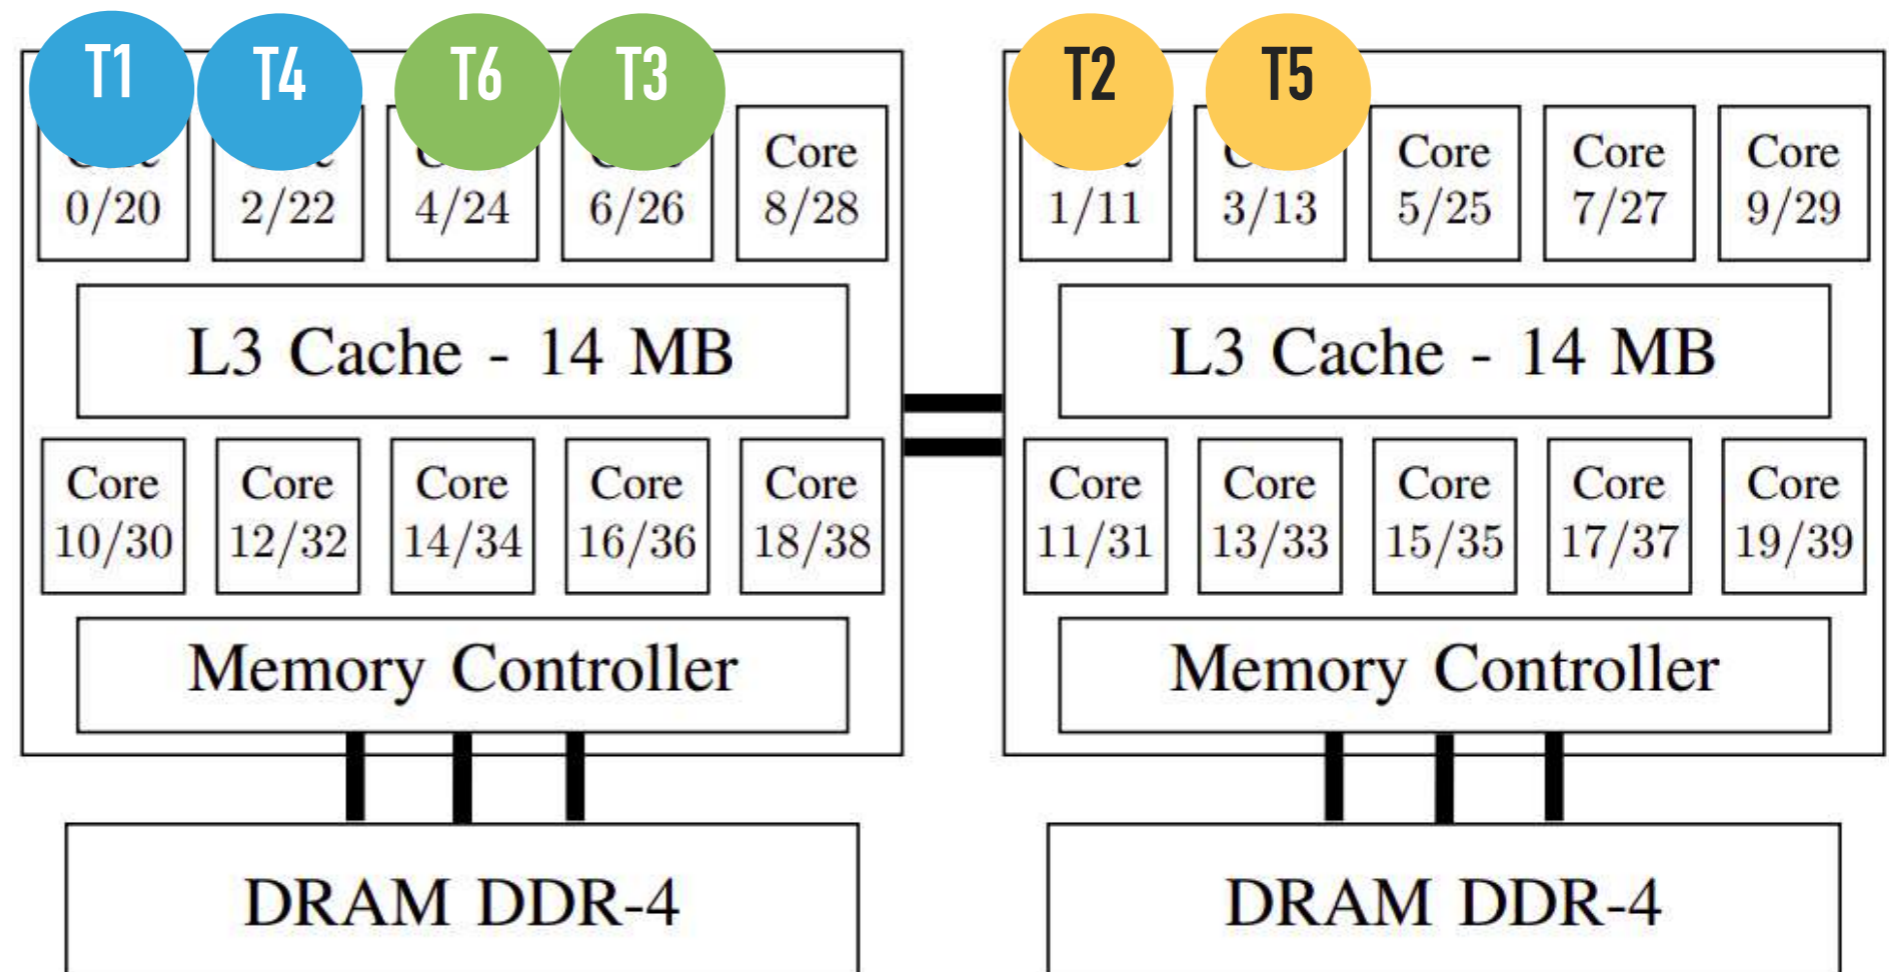
- ▶ One thread each socket



THREAD PINNING STRATEGIES

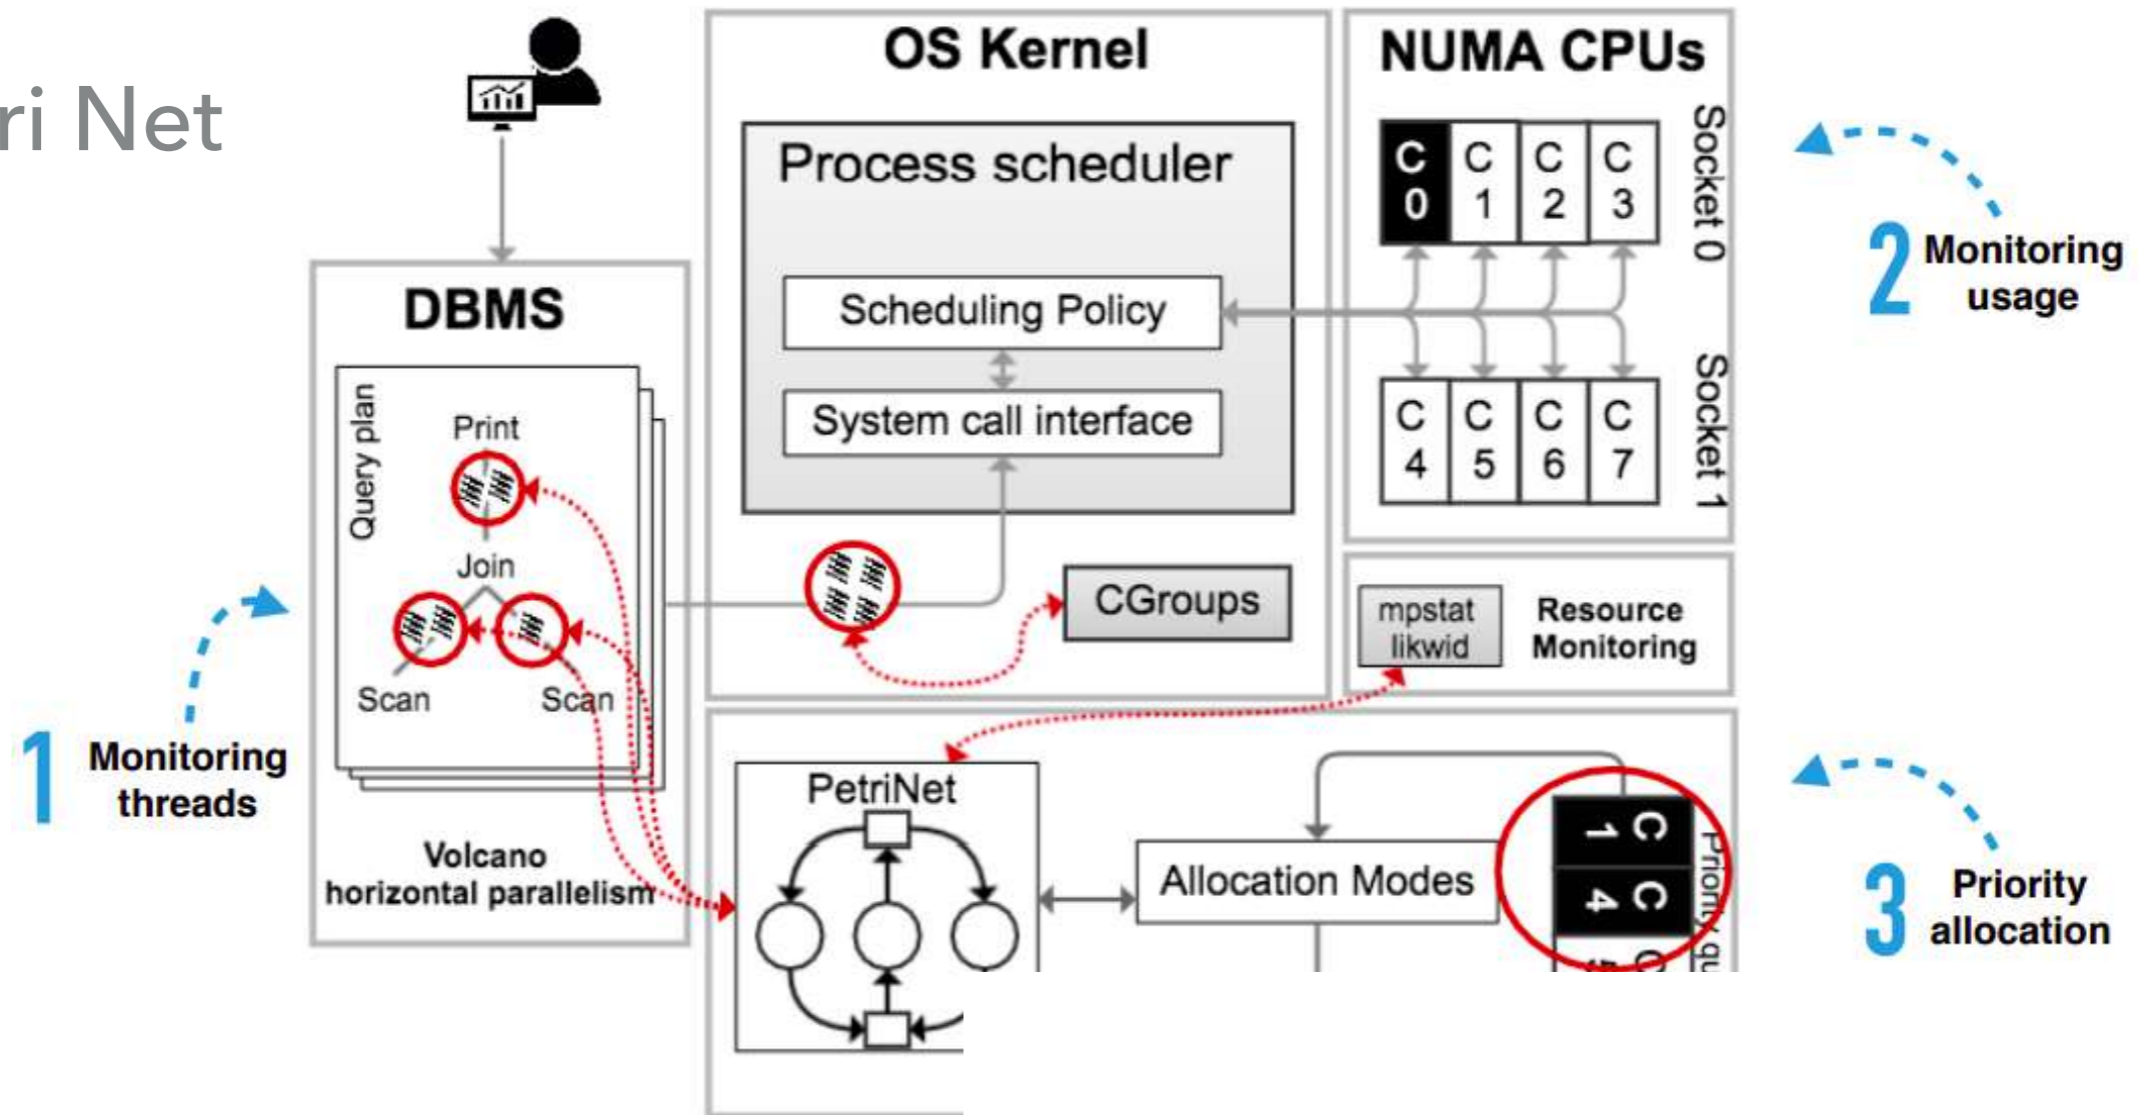
▶ SHARED

- ▶ Threads that share data on the same node.



THREAD PINNING STRATEGIES

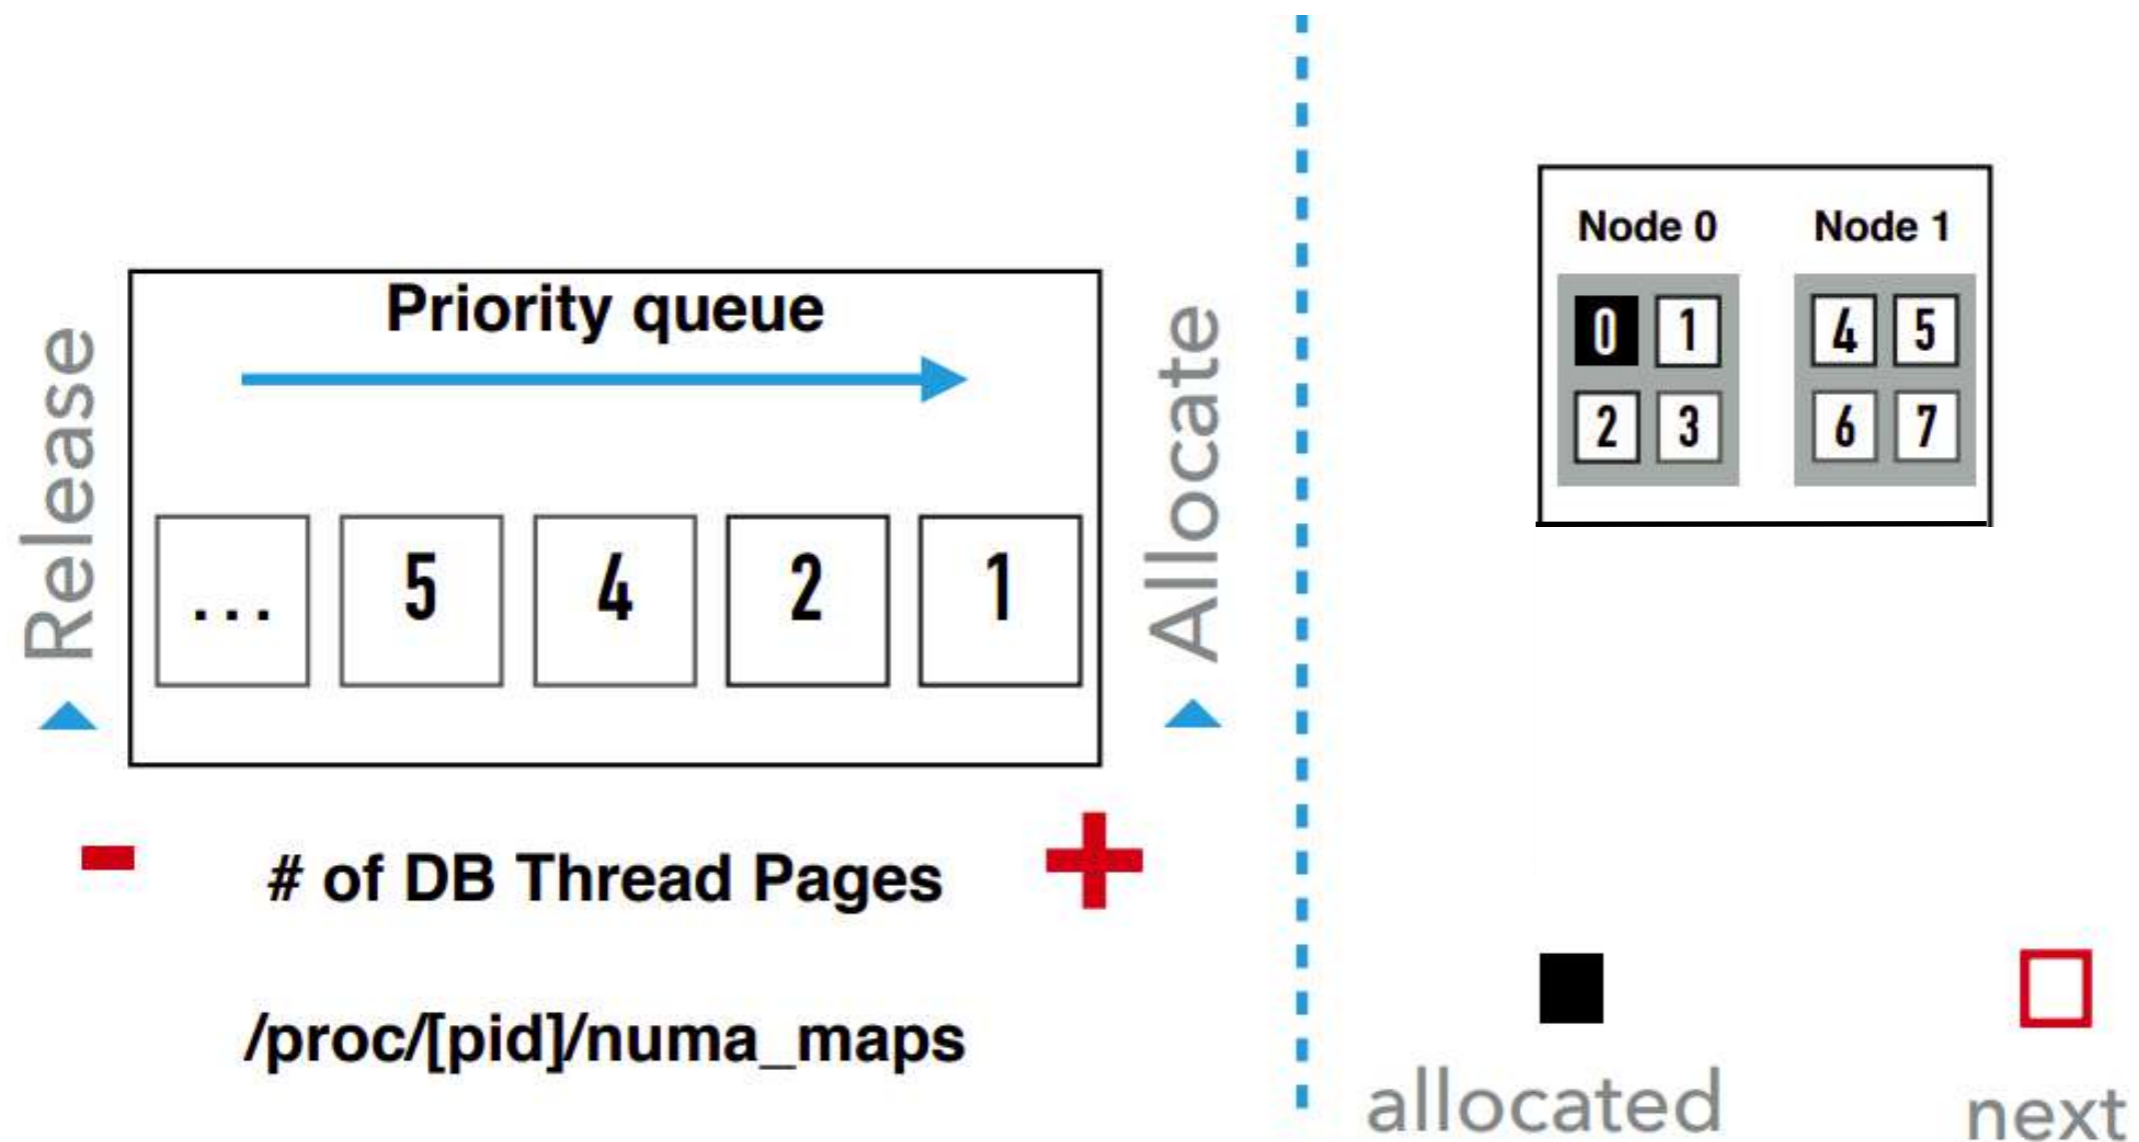
► Petri Net



S. Dominico, E. C. de Almeida, J. A. Meira, and M. A. Z. Alves, "An elastic multi-core allocation mechanism for database systems," in IEEE 34th Int. Conf. on Data Eng. (ICDE), 2018, pp. 473–484.

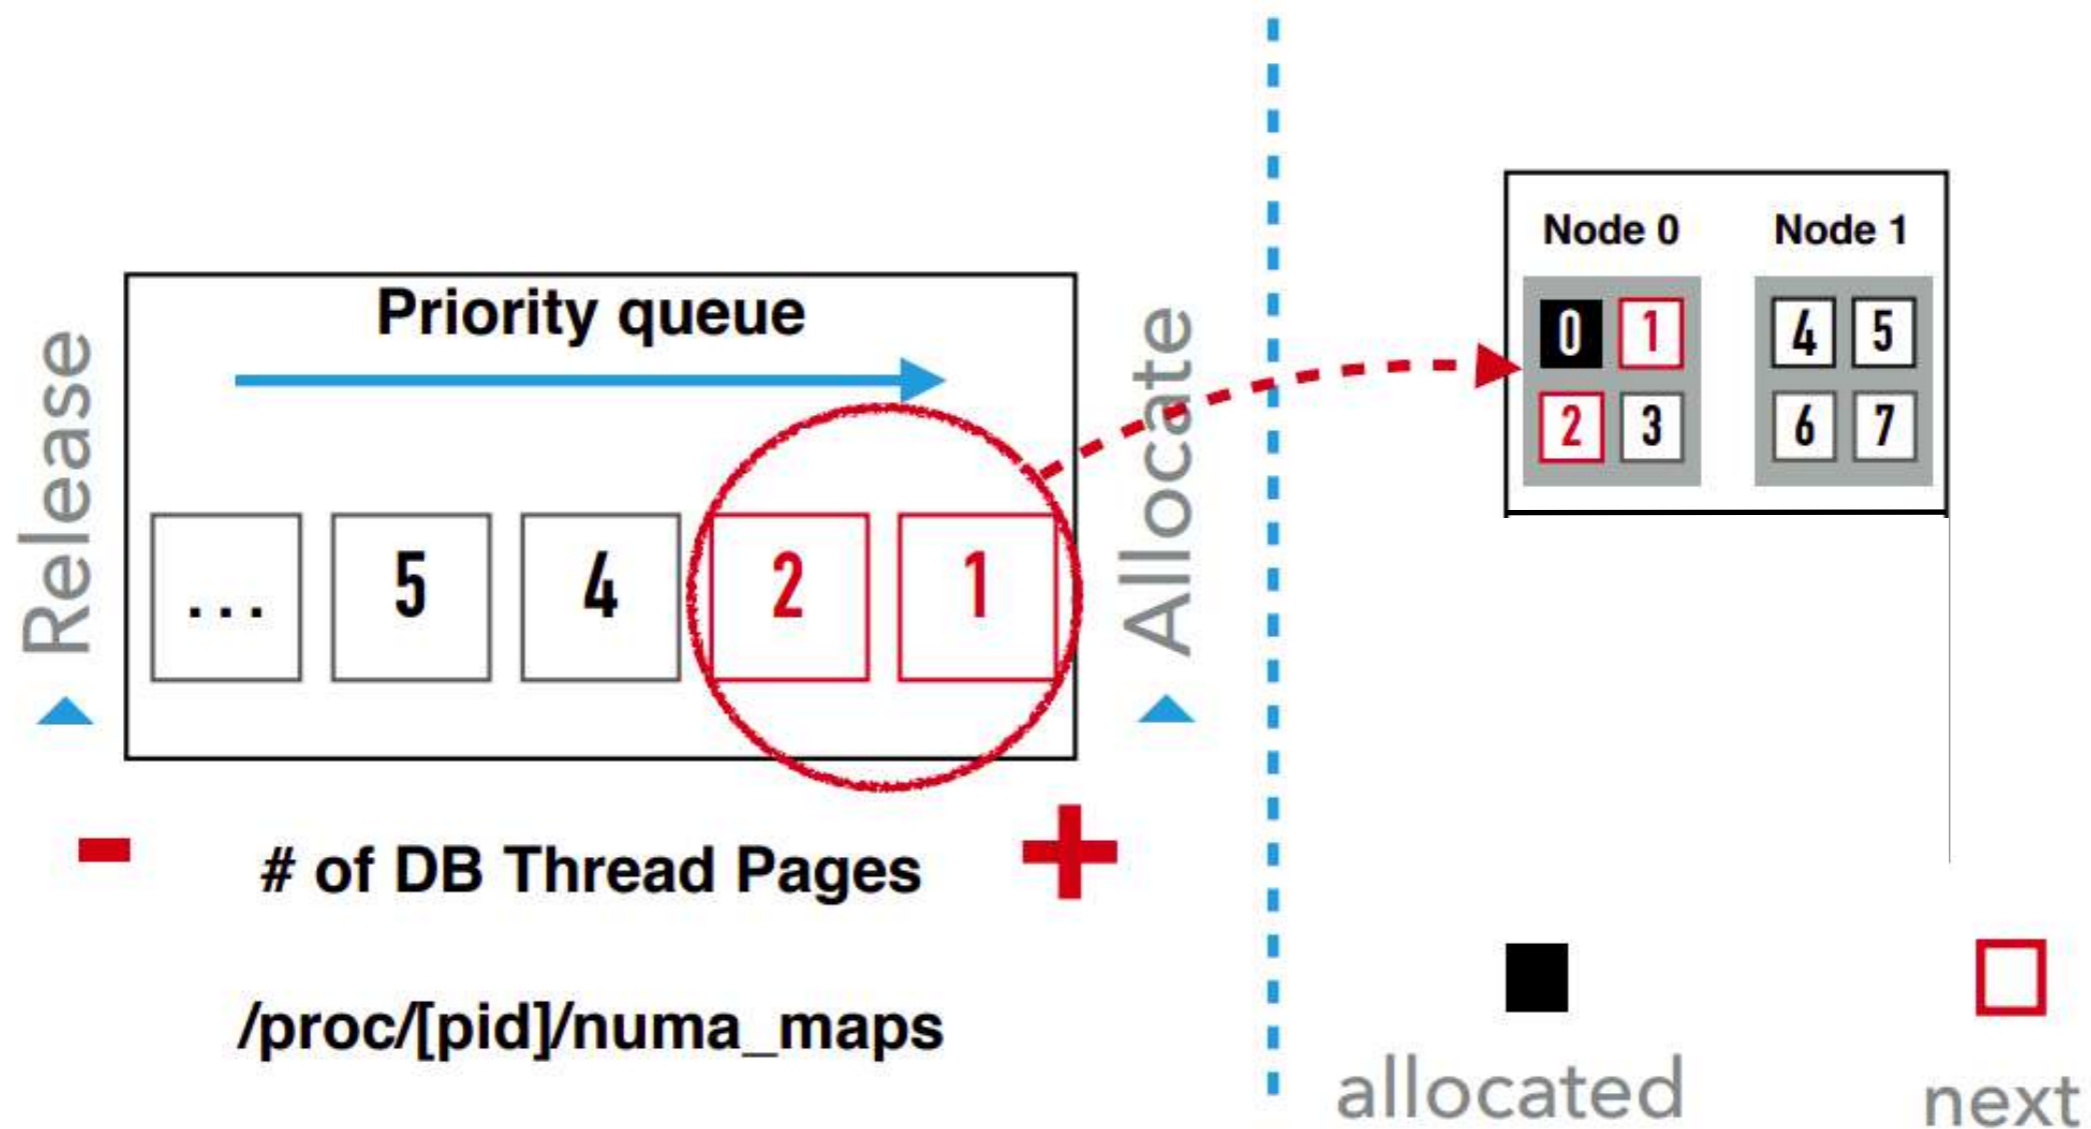
THREAD PINNING STRATEGIES

▶ Petri Net



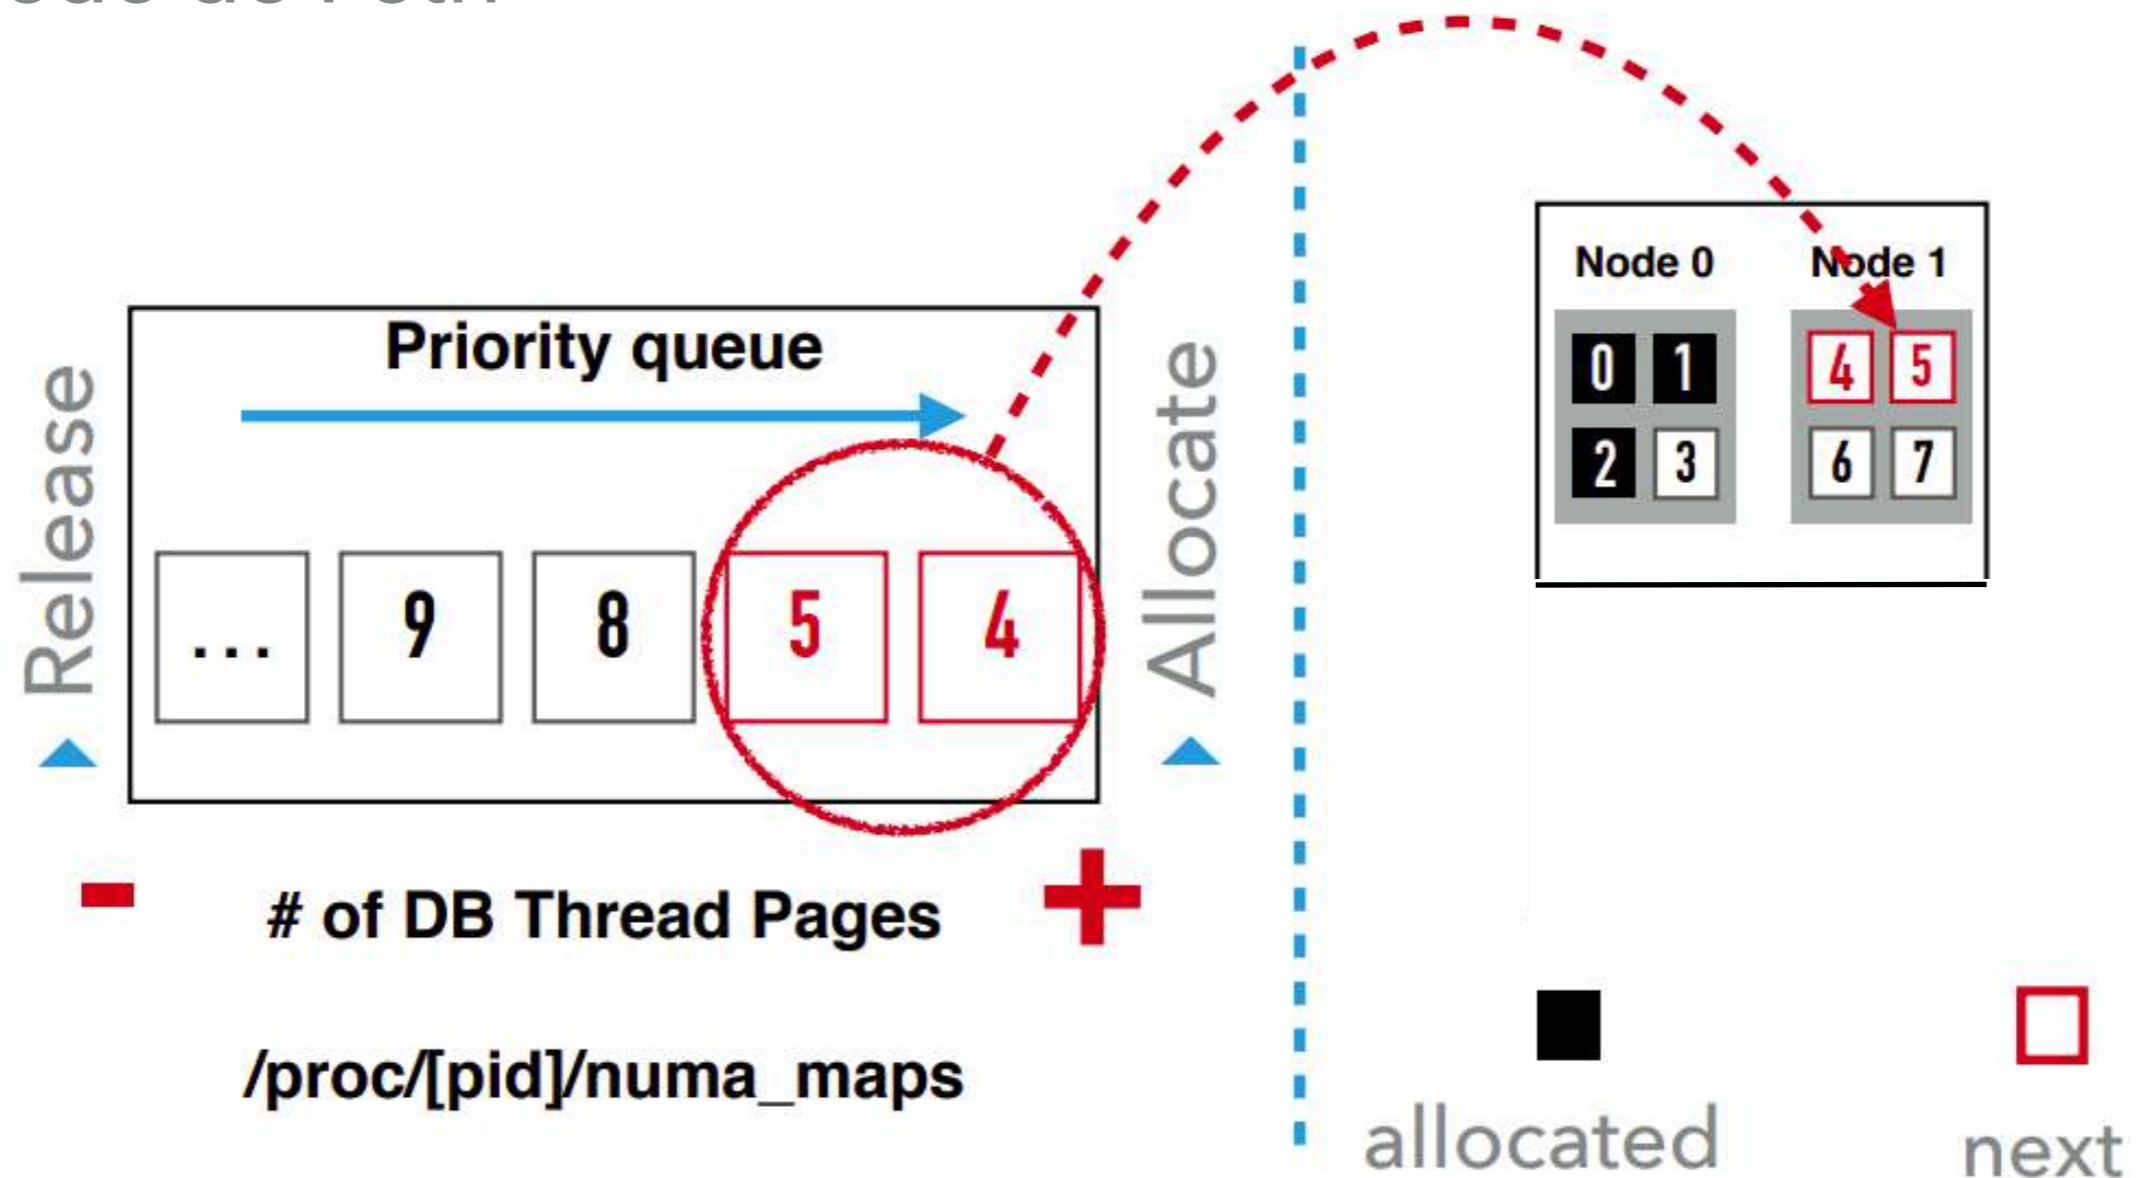
THREAD PINNING STRATEGIES

▶ Petri Net



THREAD PINNING STRATEGIES

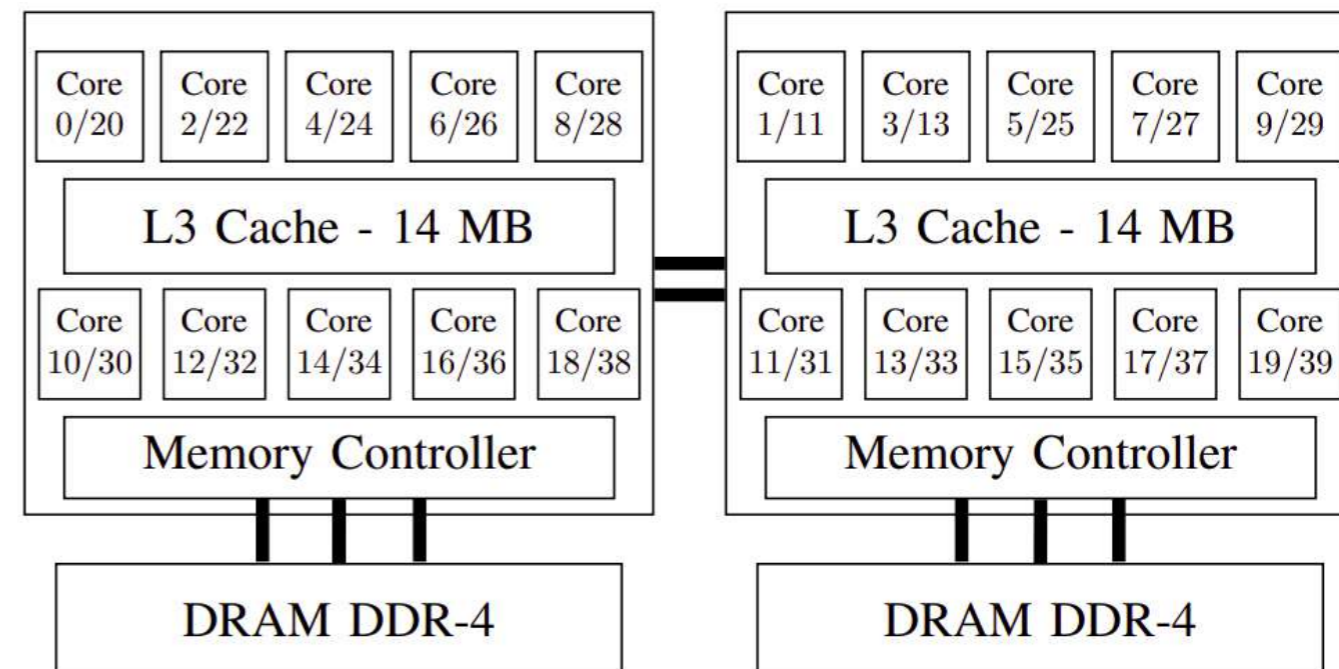
► Rede de Petri



THREAD PINNING STRATEGIES

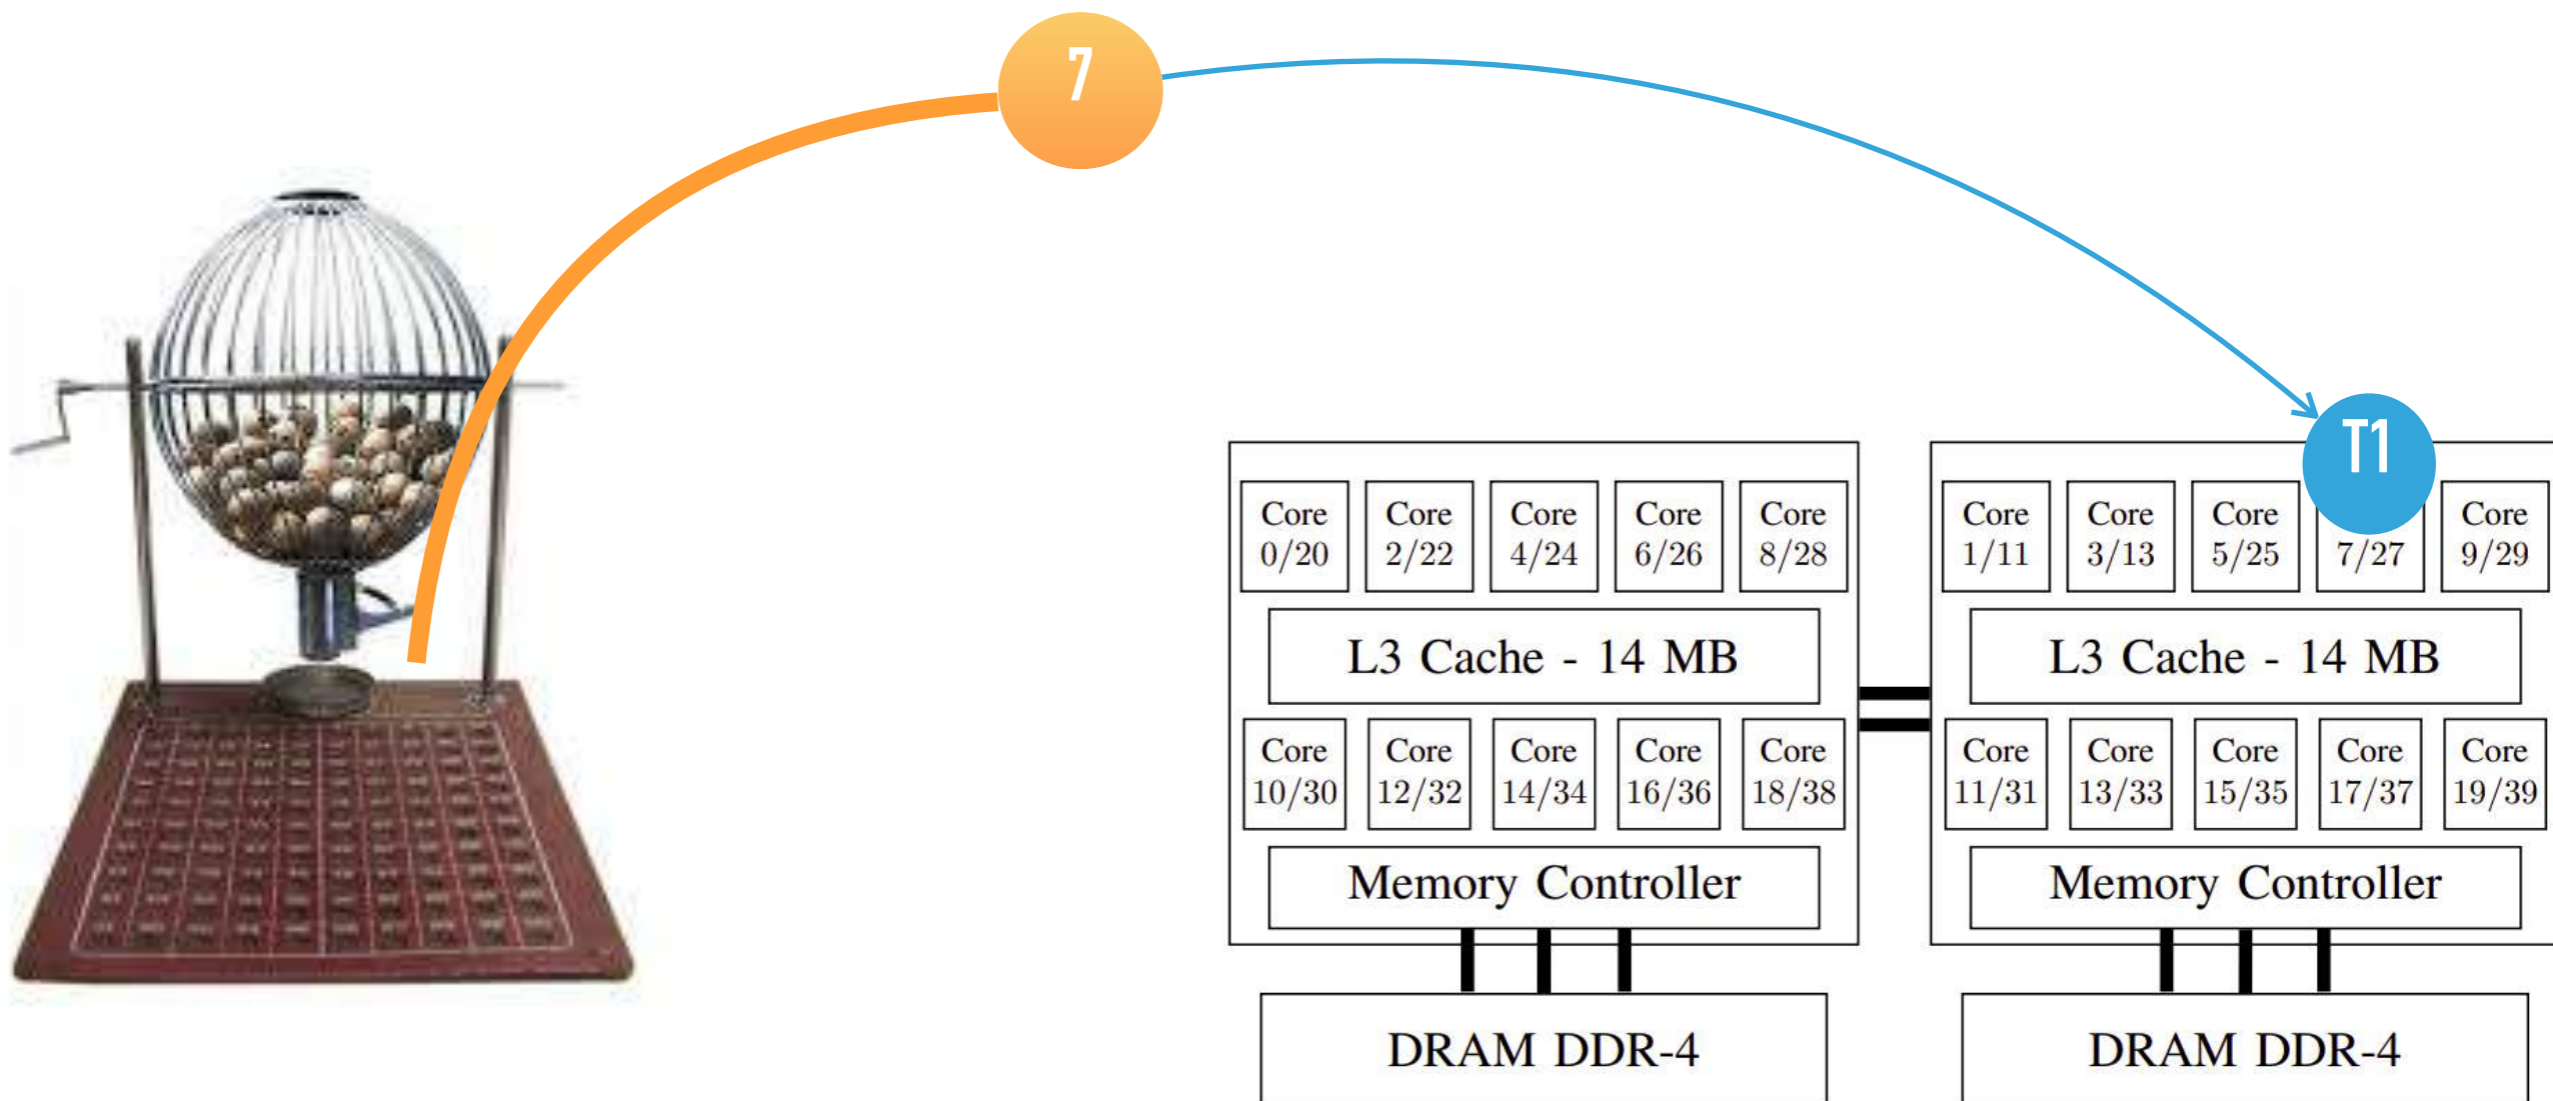
▶ RANDOM

7



THREAD PINNING STRATEGIES

▶ RANDOM



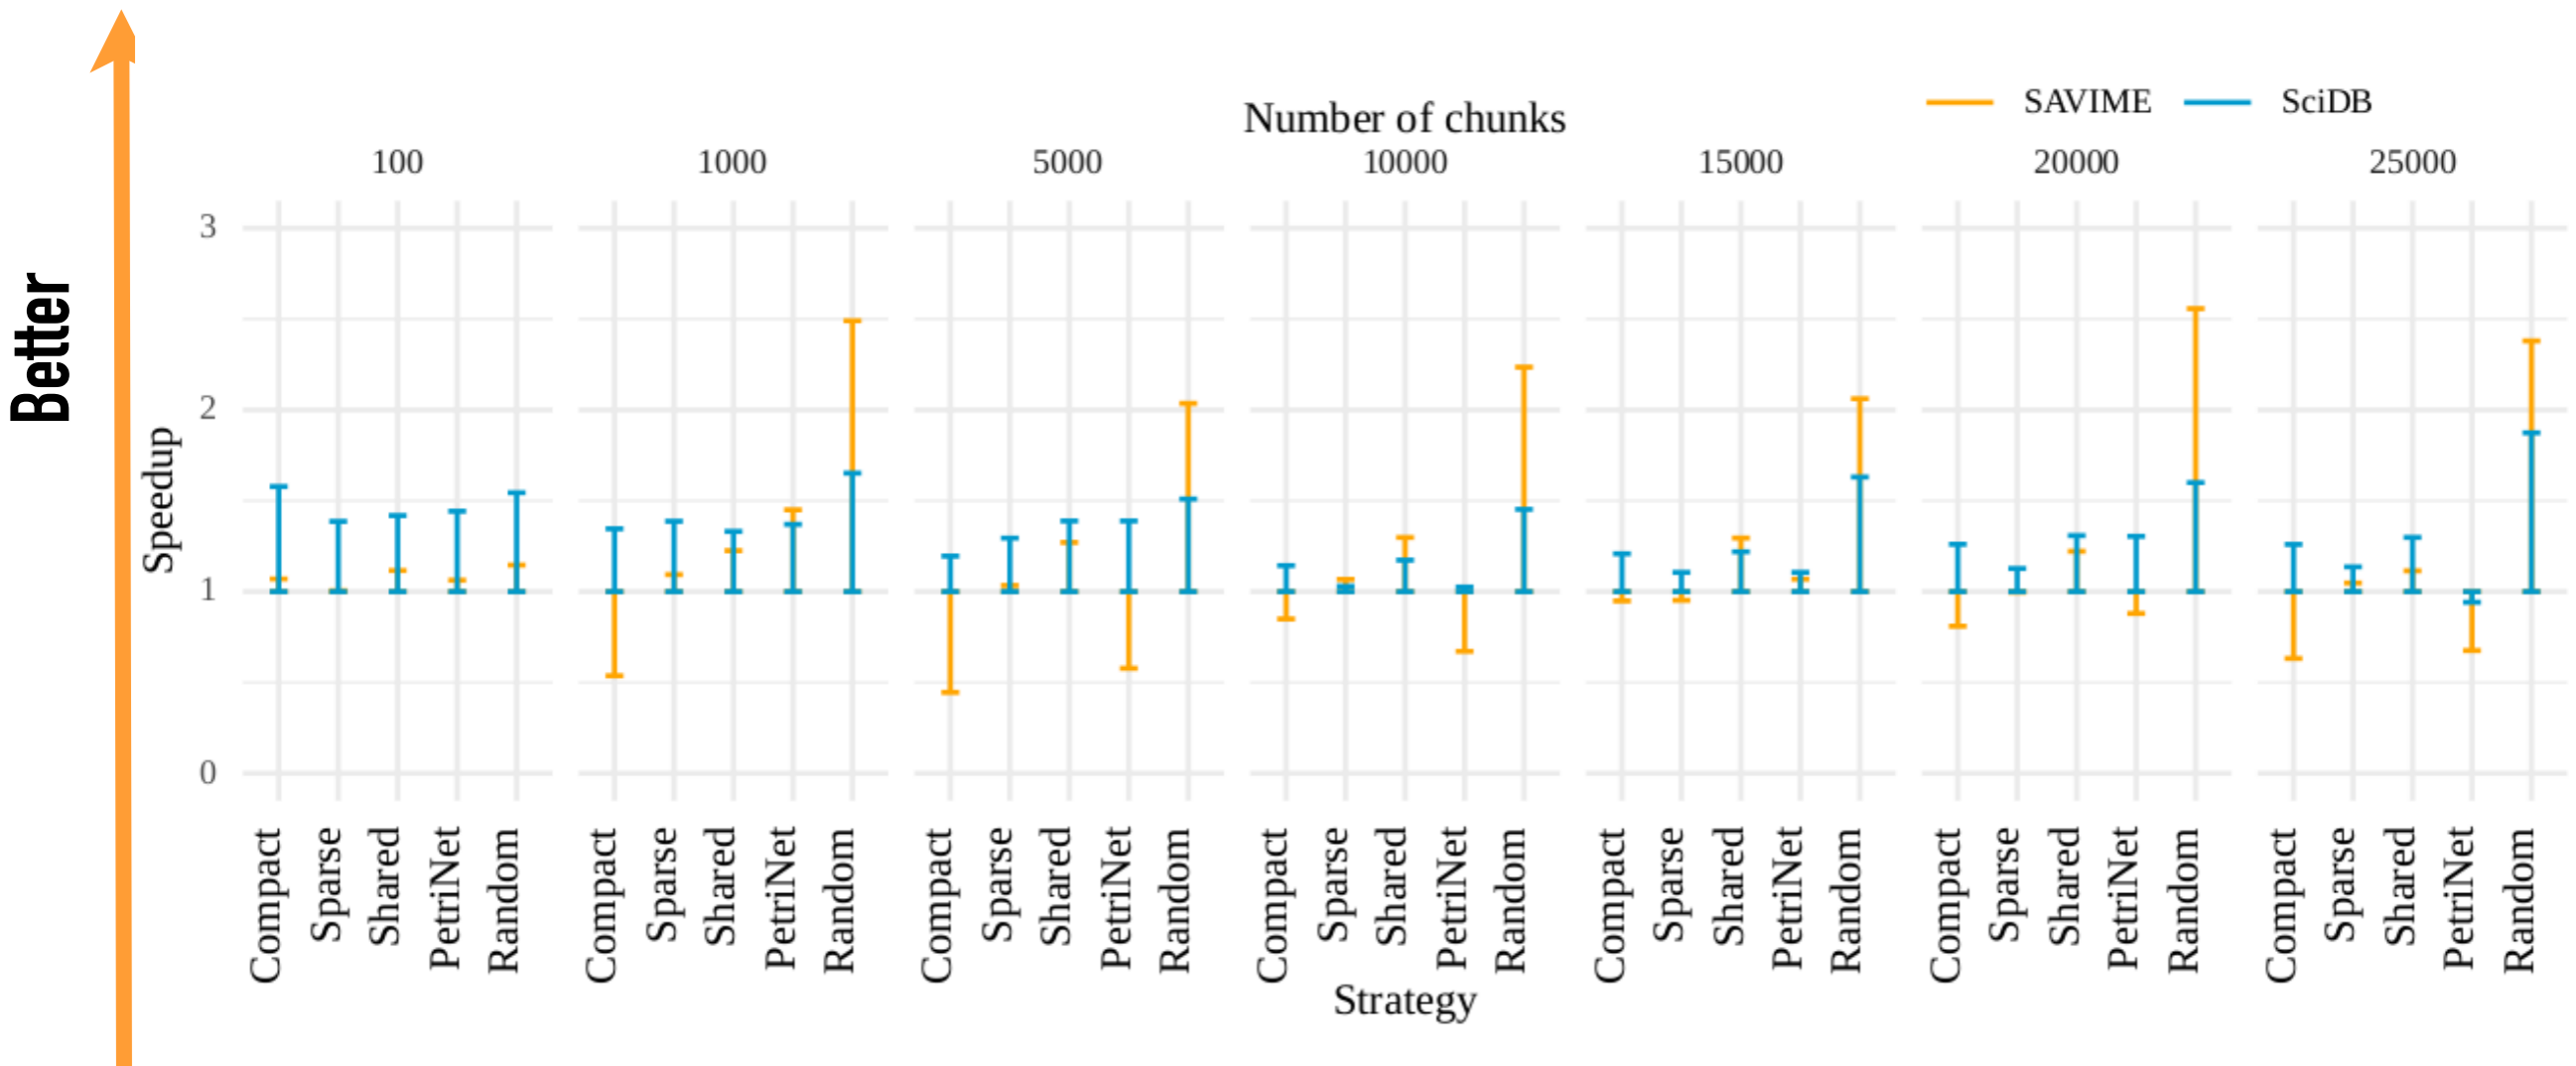
AGENDA

- ▶ Introduction
 - ▶ Array database systems
 - ▶ NUMA architecture
 - ▶ Array Database Systems in NUMA architecture
- ▶ Methodology
 - ▶ Thread pinning Strategies
- ▶ Experiments
- ▶ Conclusion

SETUP

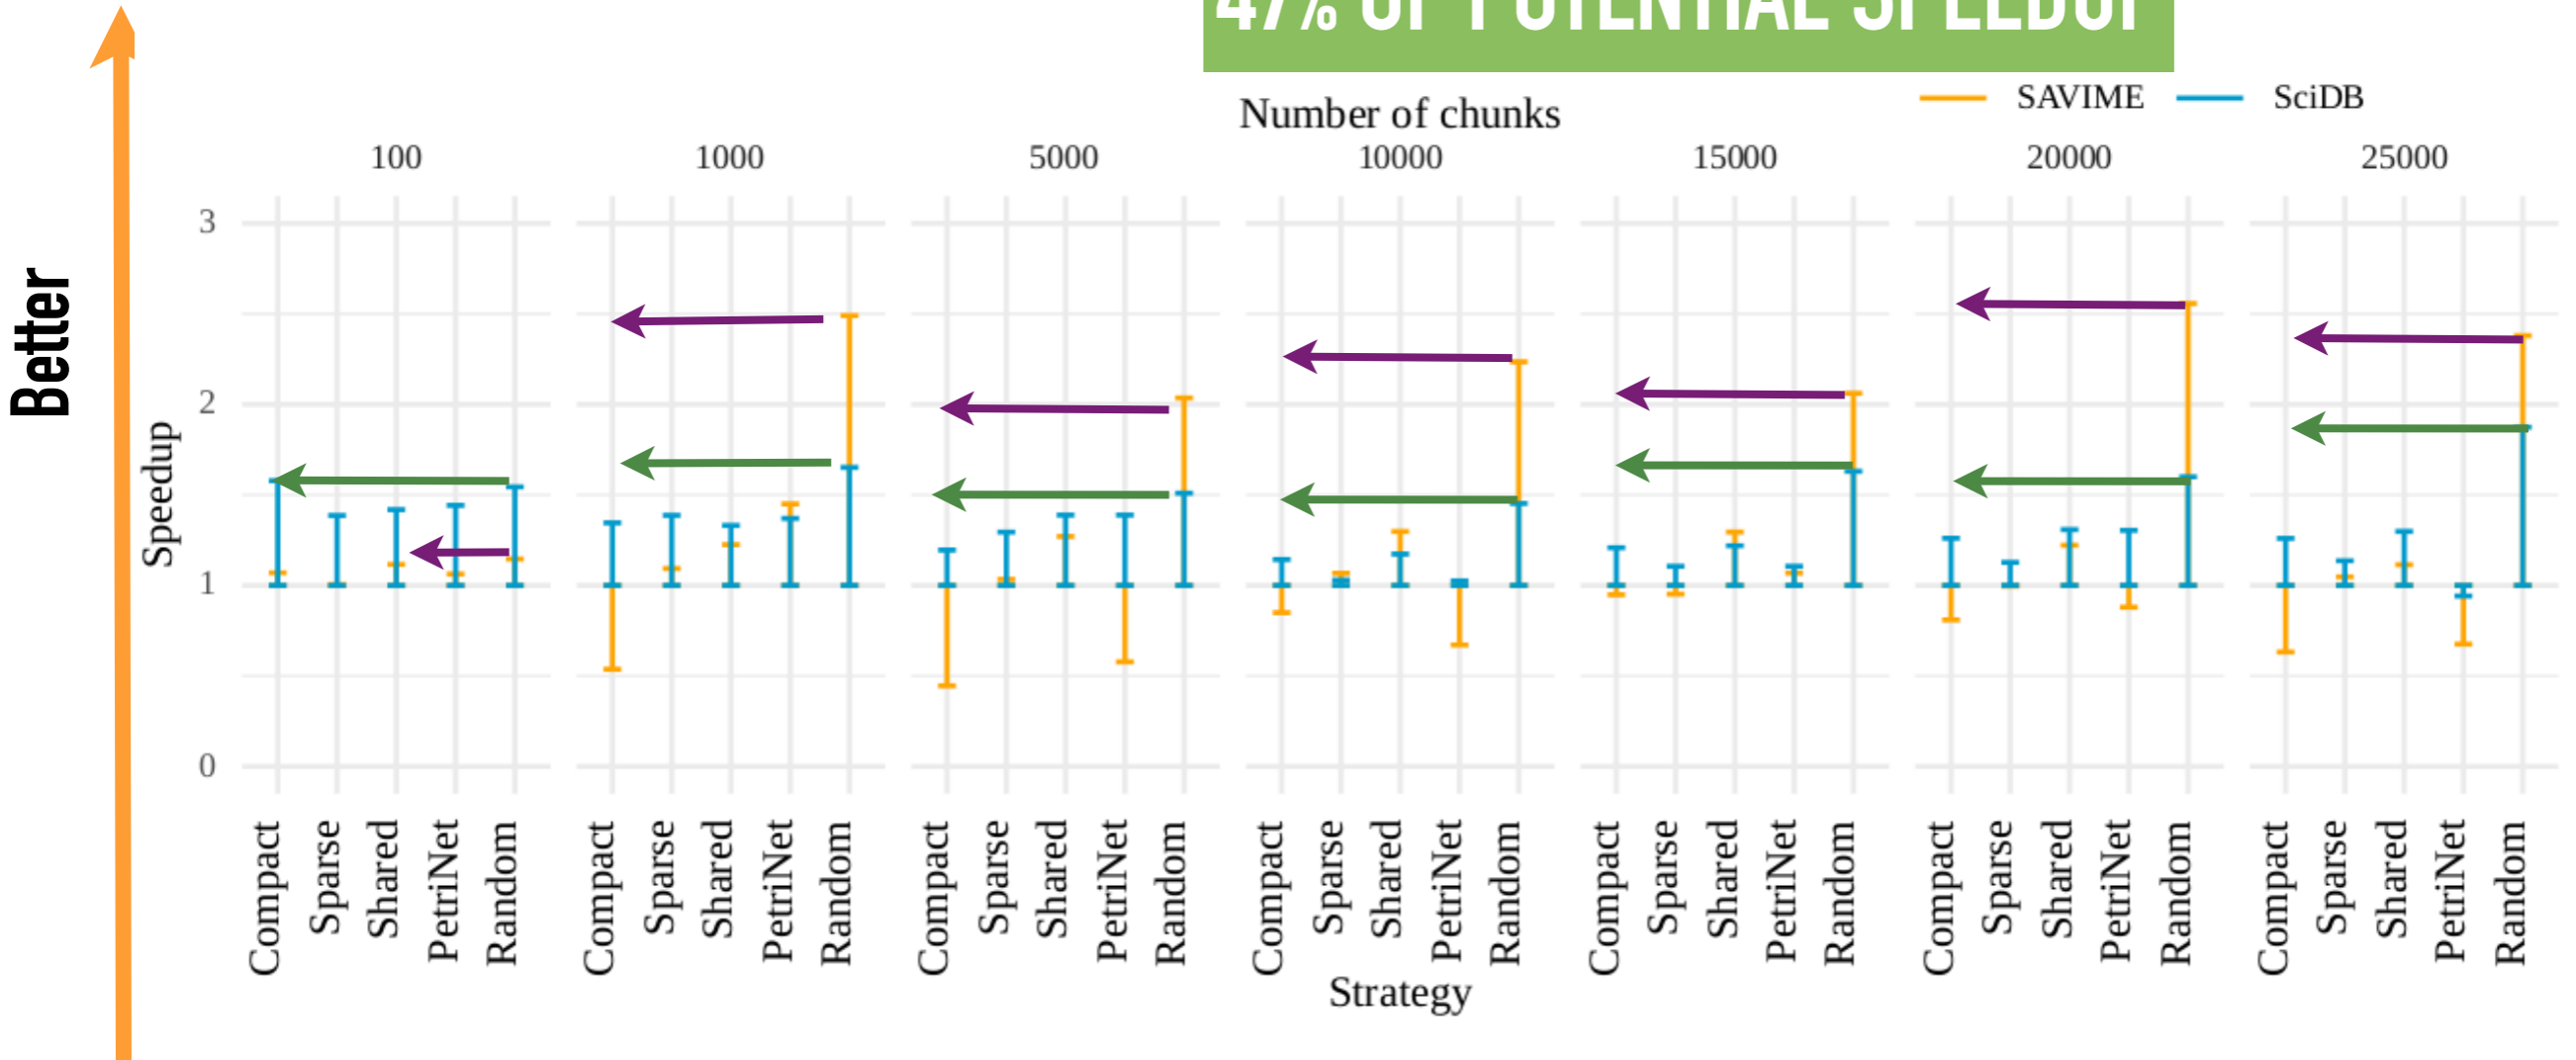
- ▶ Subarray: Savime and SciDB
- ▶ Server with 2-Sockets Intel Xeon Silver 4114 (20 cores)
- ▶ Ubuntu 18.04.01 LTS "Bionic Beaver" (Savime) and Ubuntu 14.04.6 LTS "Trusty Tahr";
- ▶ Threads: 20
- ▶ 1 GB: different numbers of chunks
- ▶ 50 GB: Selectivity

PERFORMANCE COMPARISON OF SUBARRAY OPERATOR IN 1 GB DATABASE - SPEEDUP




PERFORMANCE COMPARISON OF SUBARRAY OPERATOR IN 1 GB DATABASE - SPEEDUP

47% OF POTENTIAL SPEEDUP



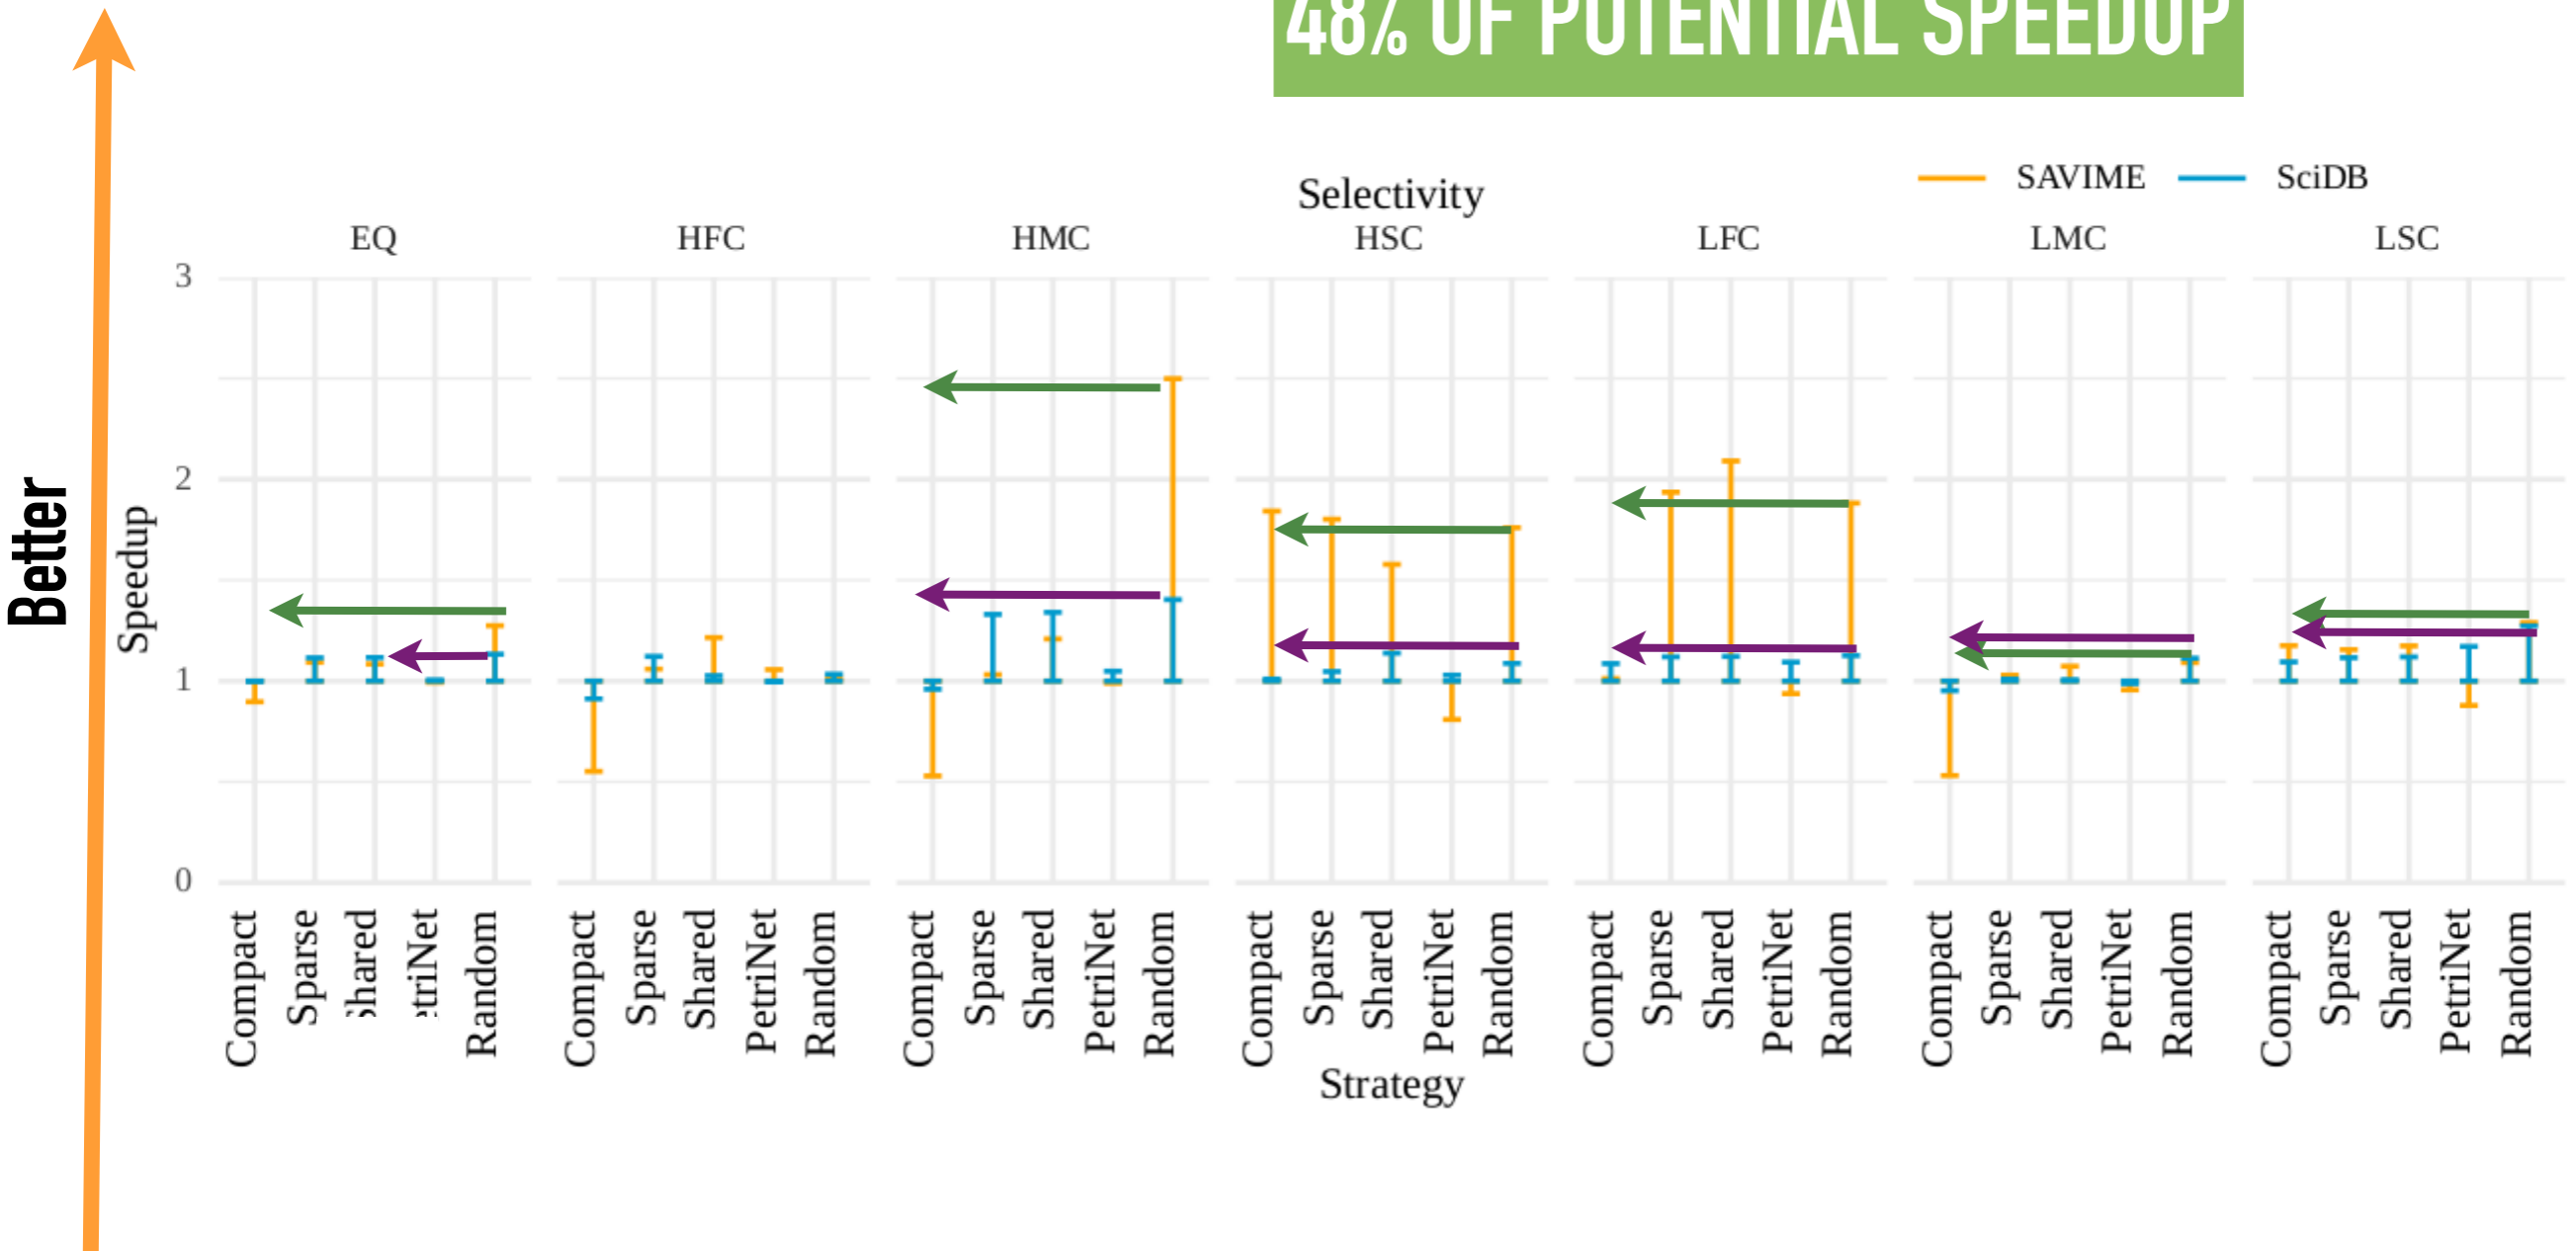
PERFORMANCE COMPARISON OF SUBARRAY OPERATOR IN 1 GB DATABASE - SPEEDUP

PREVIOUS WORK ARE NOT EXTRACTING THE MAXIMUM PERFORMANCE




PERFORMANCE COMPARISON OF SUBARRAY OPERATOR IN A 50 GB DATABASE - SPEEDUP

48% OF POTENTIAL SPEEDUP



PERFORMANCE COMPARISON OF SUBARRAY OPERATOR IN A 50 GB DATABASE- MEMORY ACCESS



Evaluation of remote and local memory access in *subarray* operation in a 50 GB varying the operator selectivity in the Array database.

REMOTE ACCESS - 4.1X

PERFORMANCE COMPARISON OF SUBARRAY OPERATOR IN A 50 GB DATABASE- ENERGY COMSUPTION

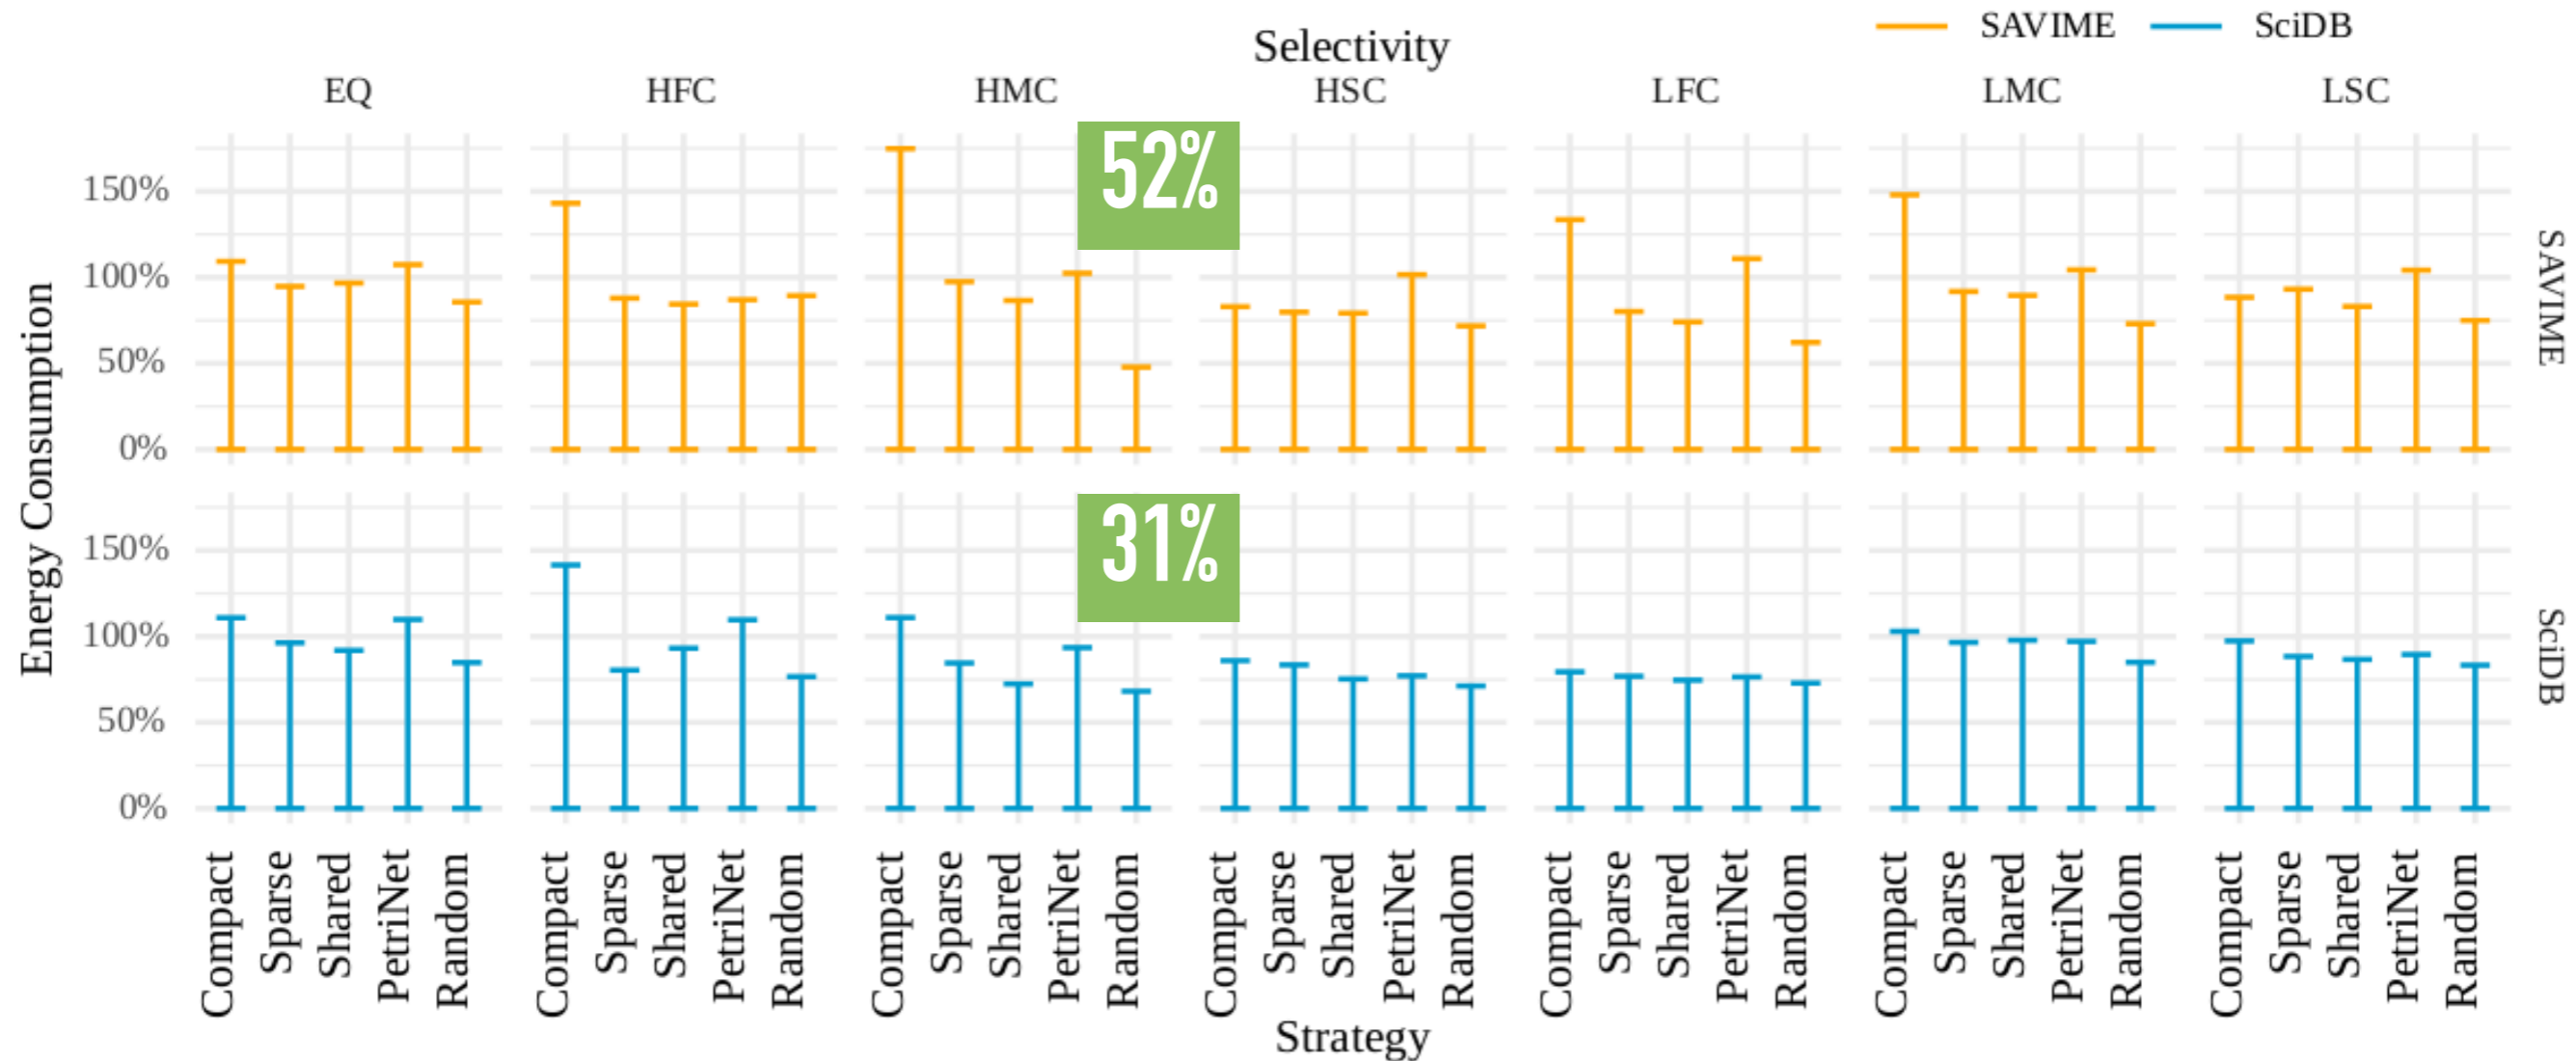


Fig. 8: Energy consumption in DRAM with 50 GB database using different operator selectivity in the Array database.

AGENDA

- ▶ Motivation
 - ▶ Array database systems
 - ▶ NUMA architecture
 - ▶ Array Database Systems in NUMA architecture
- ▶ Methodology
 - ▶ Thread pinning Strategies
- ▶ Experiments
- ▶ Conclusion

CONCLUSION AND FUTURE WORK

- ▶ NUMA architecture affects in performance of subarray operations in array database systems.
- ▶ Traditional techniques are still far from the maximum possible gains.
- ▶ Static strategies only yield 48% from the potential speedup (and 26% of the energy reduction)
- ▶ Opening a new research topic
- ▶ Understanding the NUMA effects in other array database operators and designing an array database scheduler that finds the best thread pinning.

THANKS!

SDOMINICO@INF.UFPR.BR

ACKNOWLEDGMENTS: CAPES, CNPQ AND SERRAPILHEIRA INSTITUTE.