NEAR DATA FILTERS: TAKING ANOTHER BRICK FROM THE MEMORY WALL

Diego G. Tomé*
CWI, The Netherlands
diego.tome@cwi.nl

Tiago R. Kepe
UFPR and IFPR, Brazil
trkepe@inf.ufpr.br

Marco A. Z. Alves
UFPR, Brazil
mazalves@inf.ufpr.br

Eduardo C. de Almeida
UFPR, Brazil
eduardo@inf.ufpr.br

NEAR DATA FILTERS: TAKING ANOTHER BRICK FROM THE MEMORY WALL

INTRODUCTION

MEMORY WALL PROBLEM

- The disparity between CPU performance and main memory latency has grown tightly.

MEMORY WALL PROBLEM

- The disparity between CPU performance and main memory latency has grown tightly.

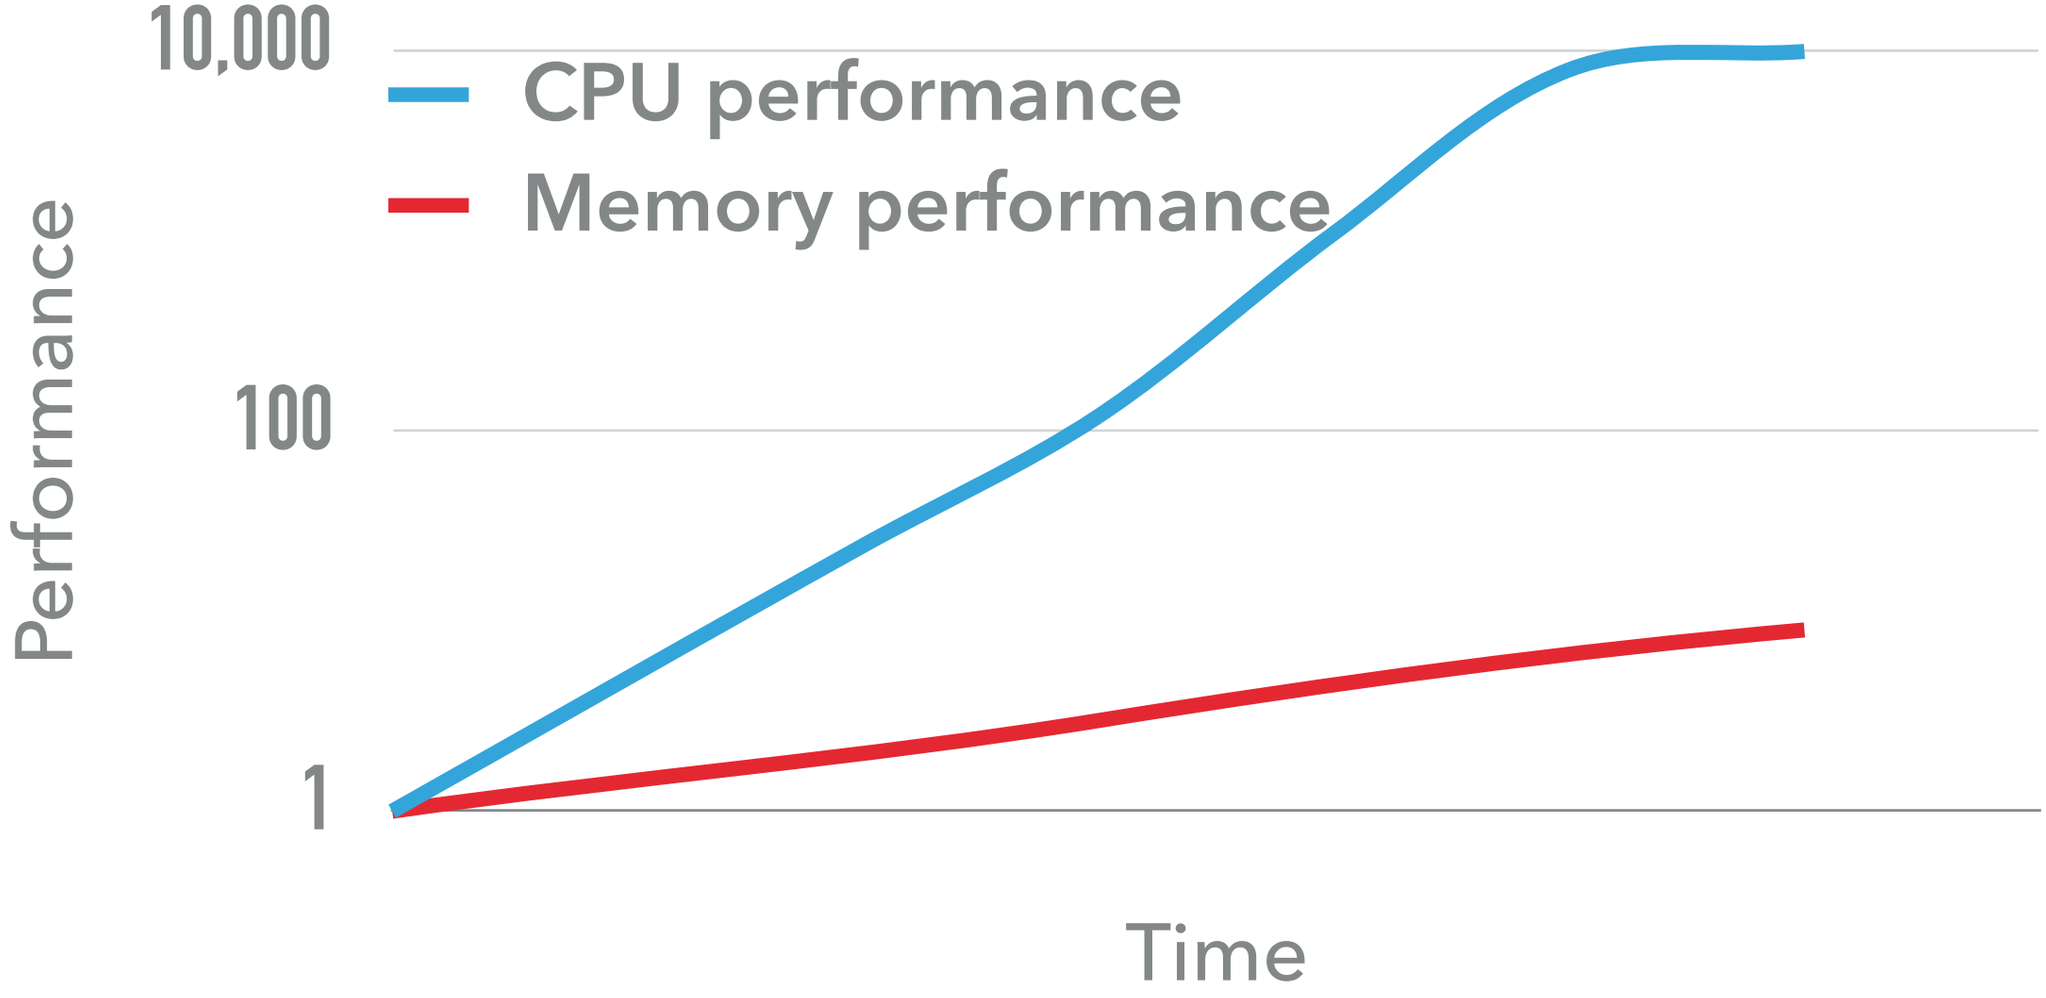

MEMORY WALL PROBLEM

- The disparity between CPU performance and main memory latency has grown tightly.

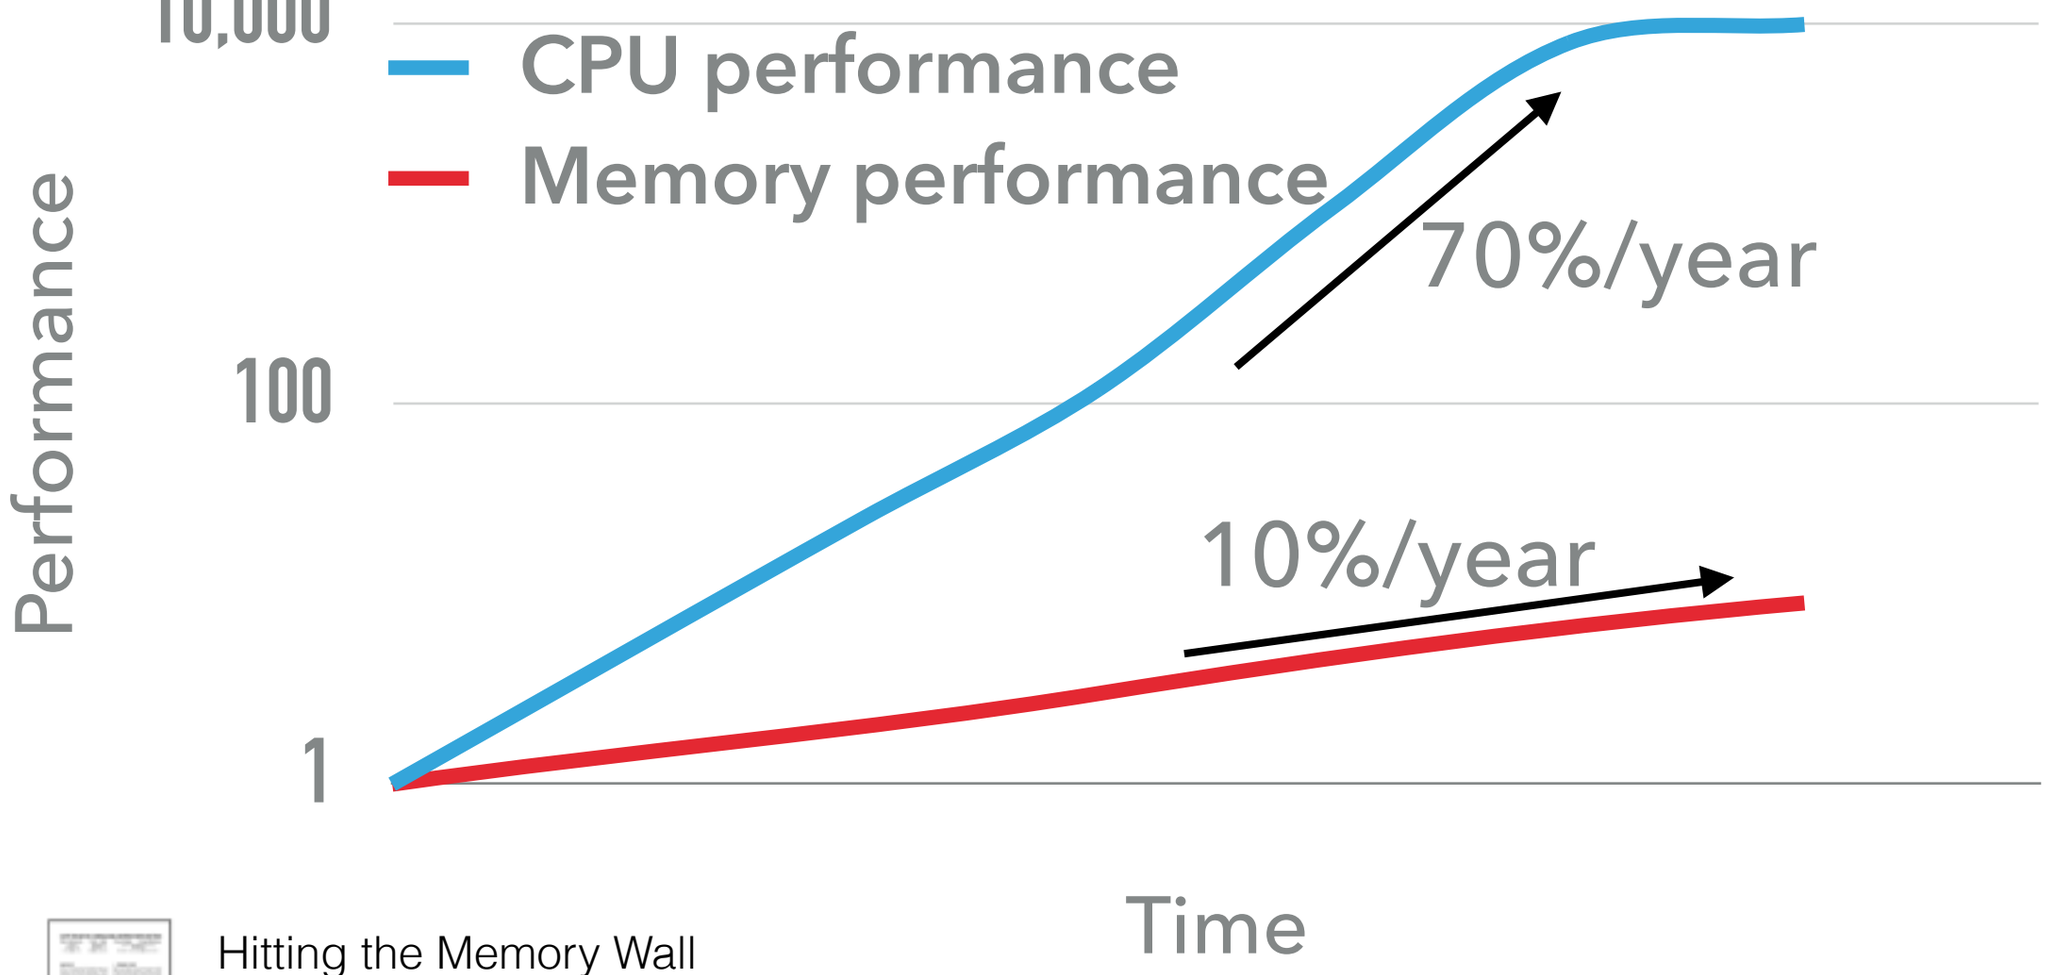

Hitting the Memory Wall
W. Wulf and McKee S.A.
ACM Computer Architecture News, 1994

MEMORY WALL PROBLEM

- The disparity between CPU performance and main memory has grown tightly.

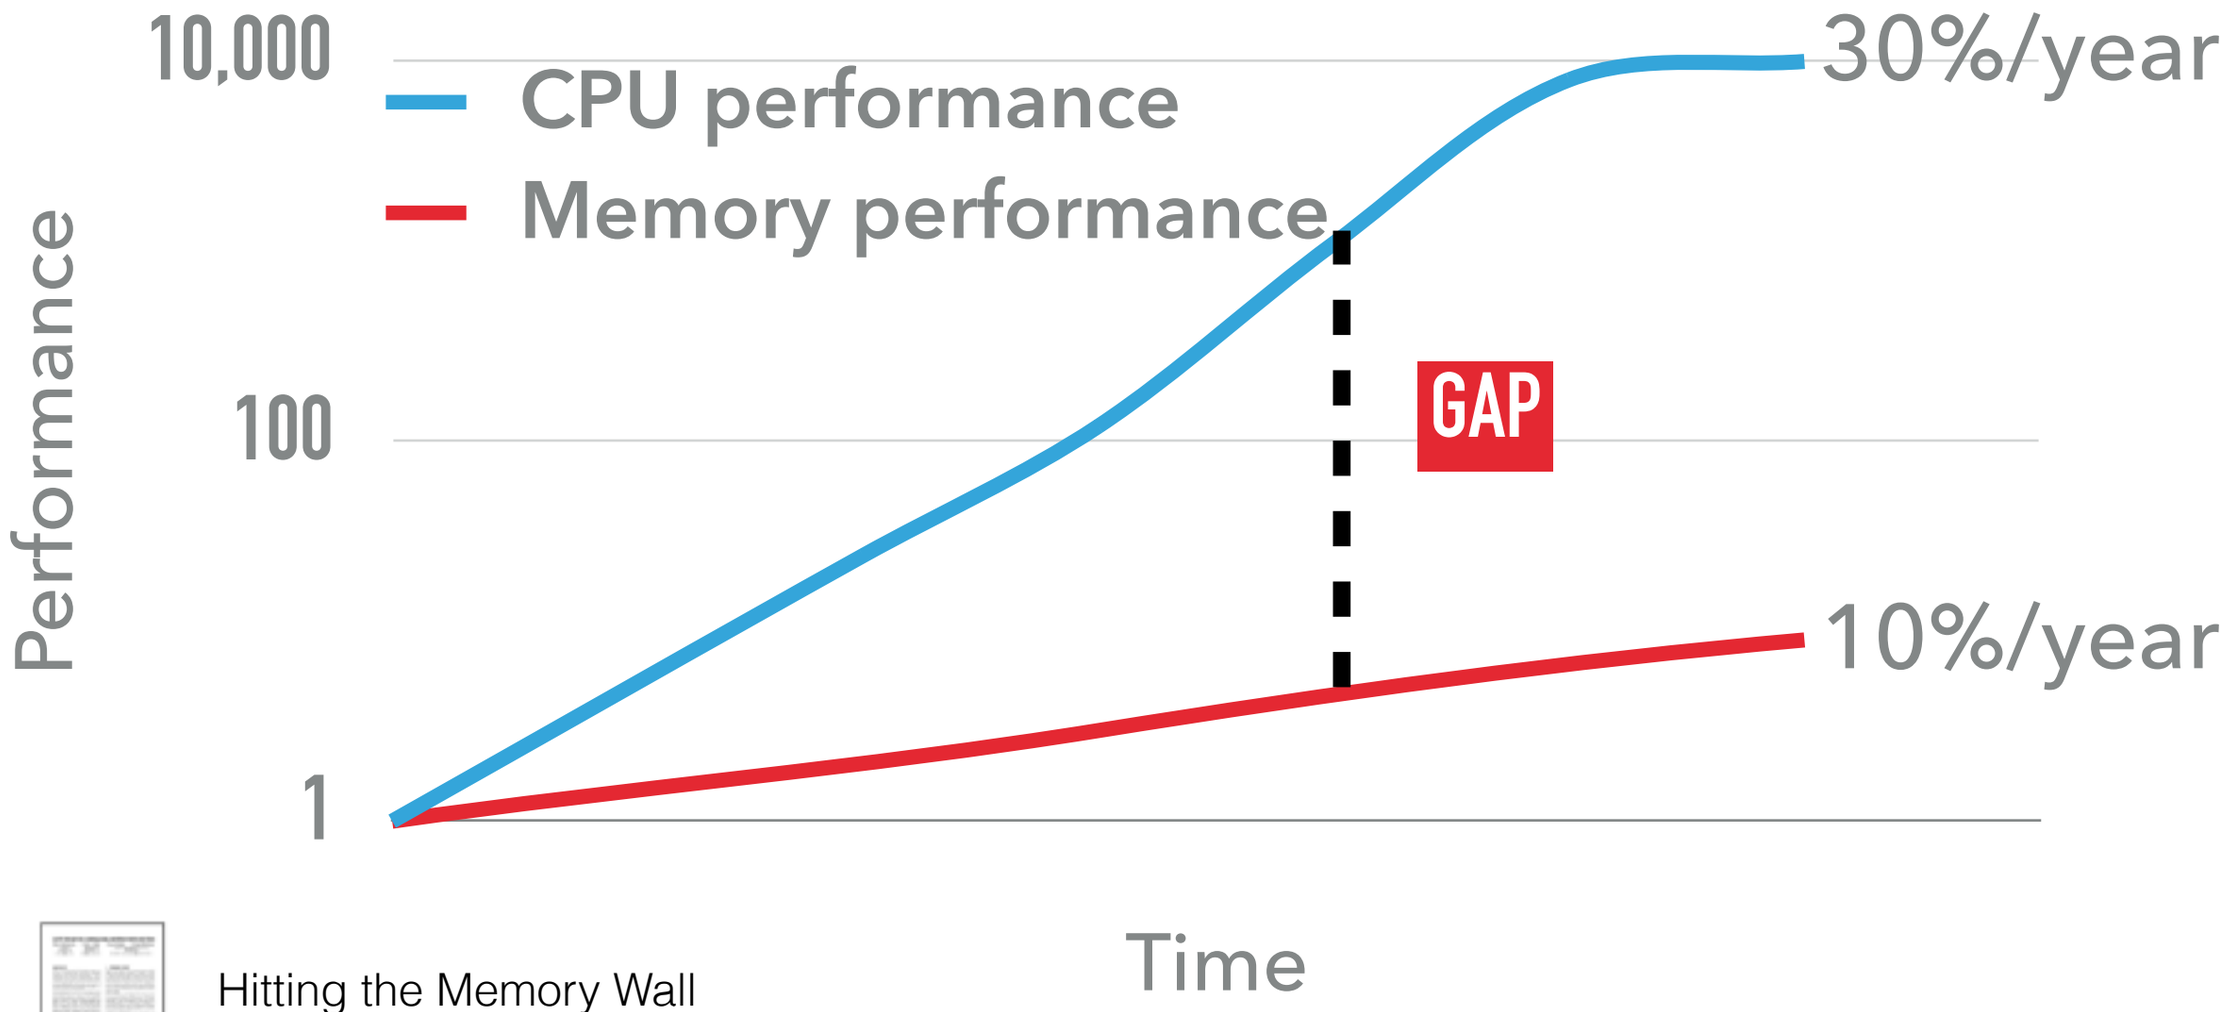

Hitting the Memory Wall
W. Wulf and McKee S.A.
ACM Computer Architecture News, 1994

MEMORY WALL PROBLEM

- The disparity between CPU performance and main memory has grown tightly.

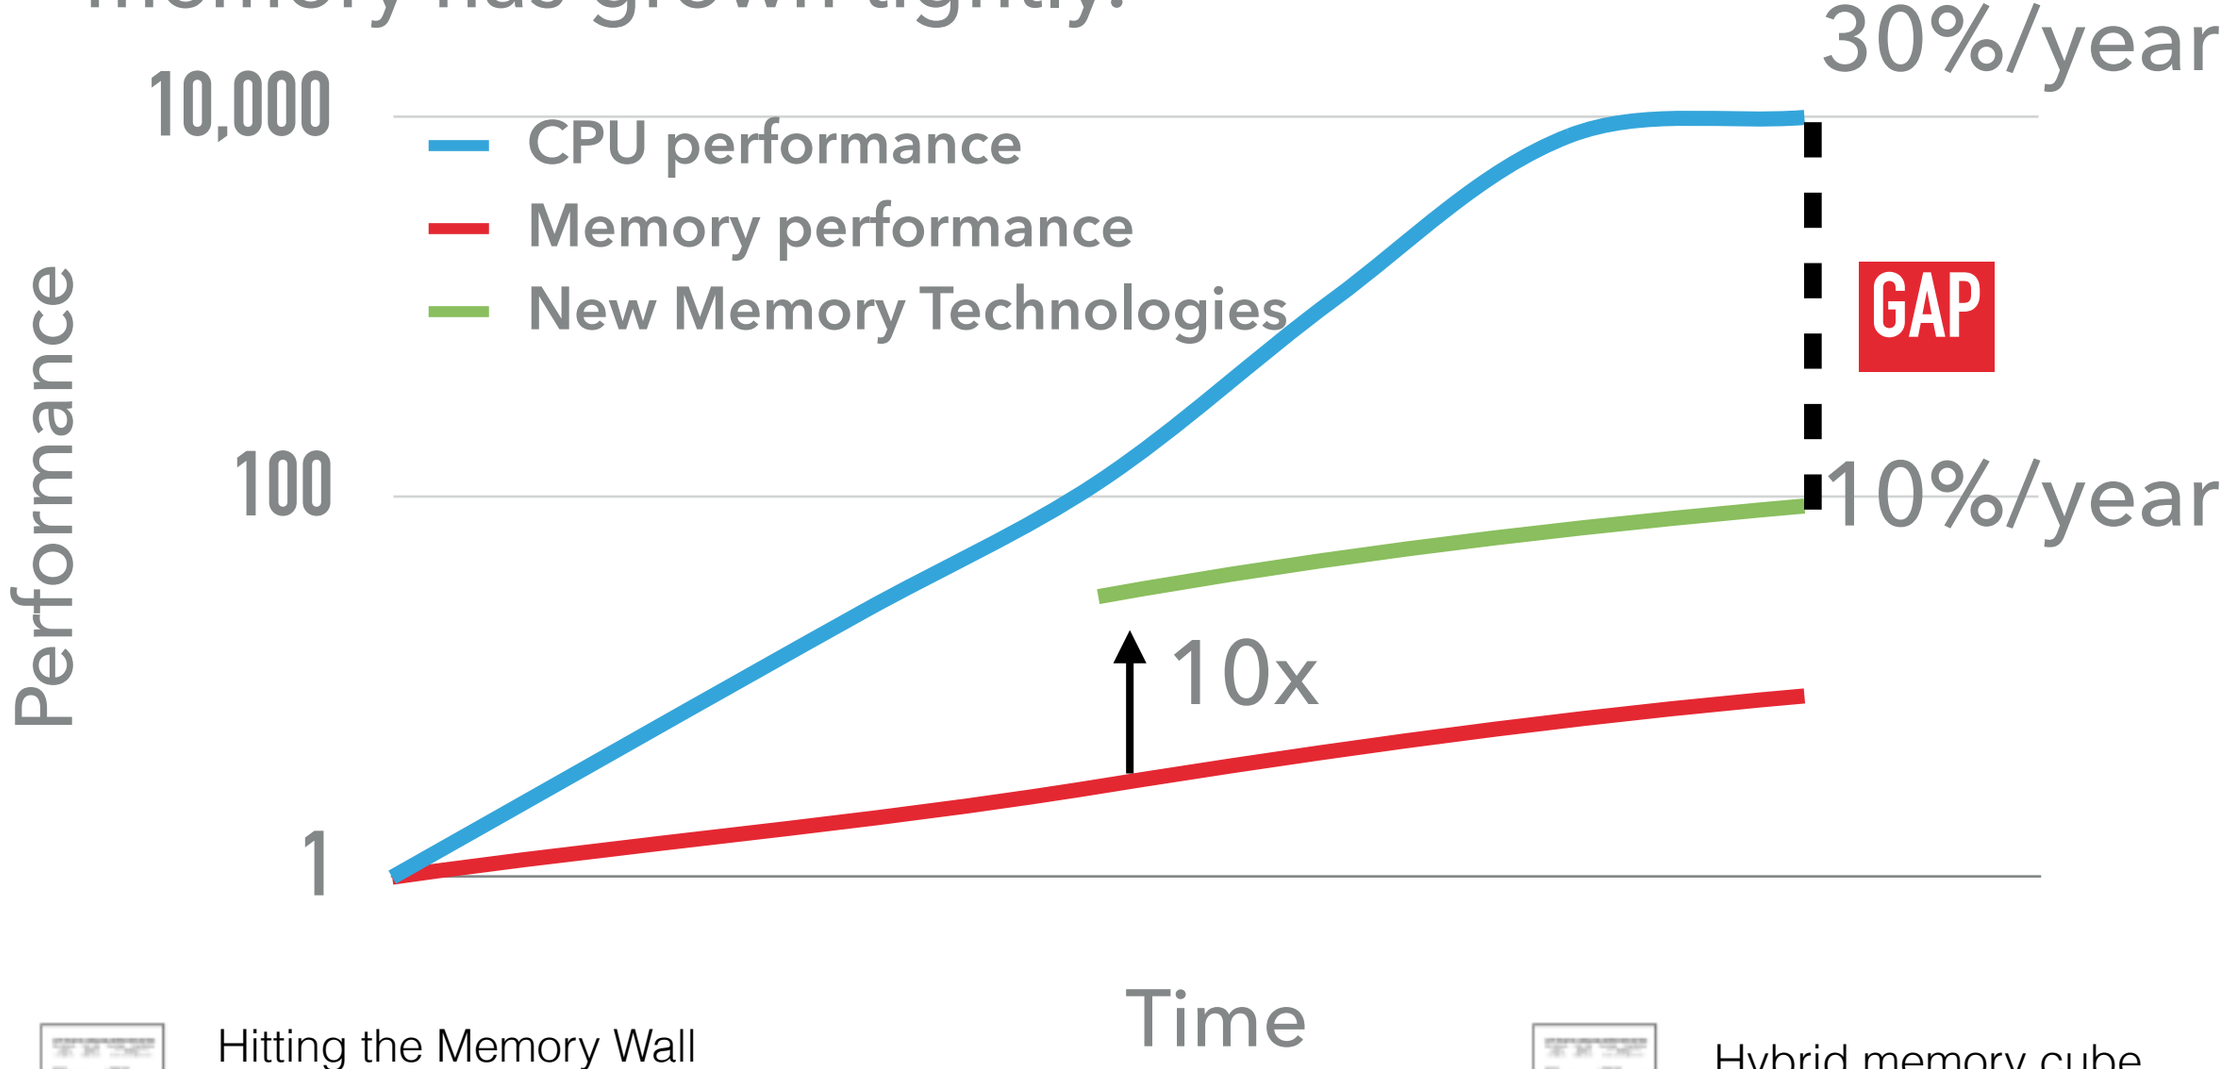

Hitting the Memory Wall
W. Wulf and McKee S.A.
ACM Computer Architecture News, 1994

Hybrid memory cube
J. Jeddeloh and B. Keet.
Symposium on VLSI (VLSIT), 2012

QUERY PROCESSING: WHERE DOES TIME GO?

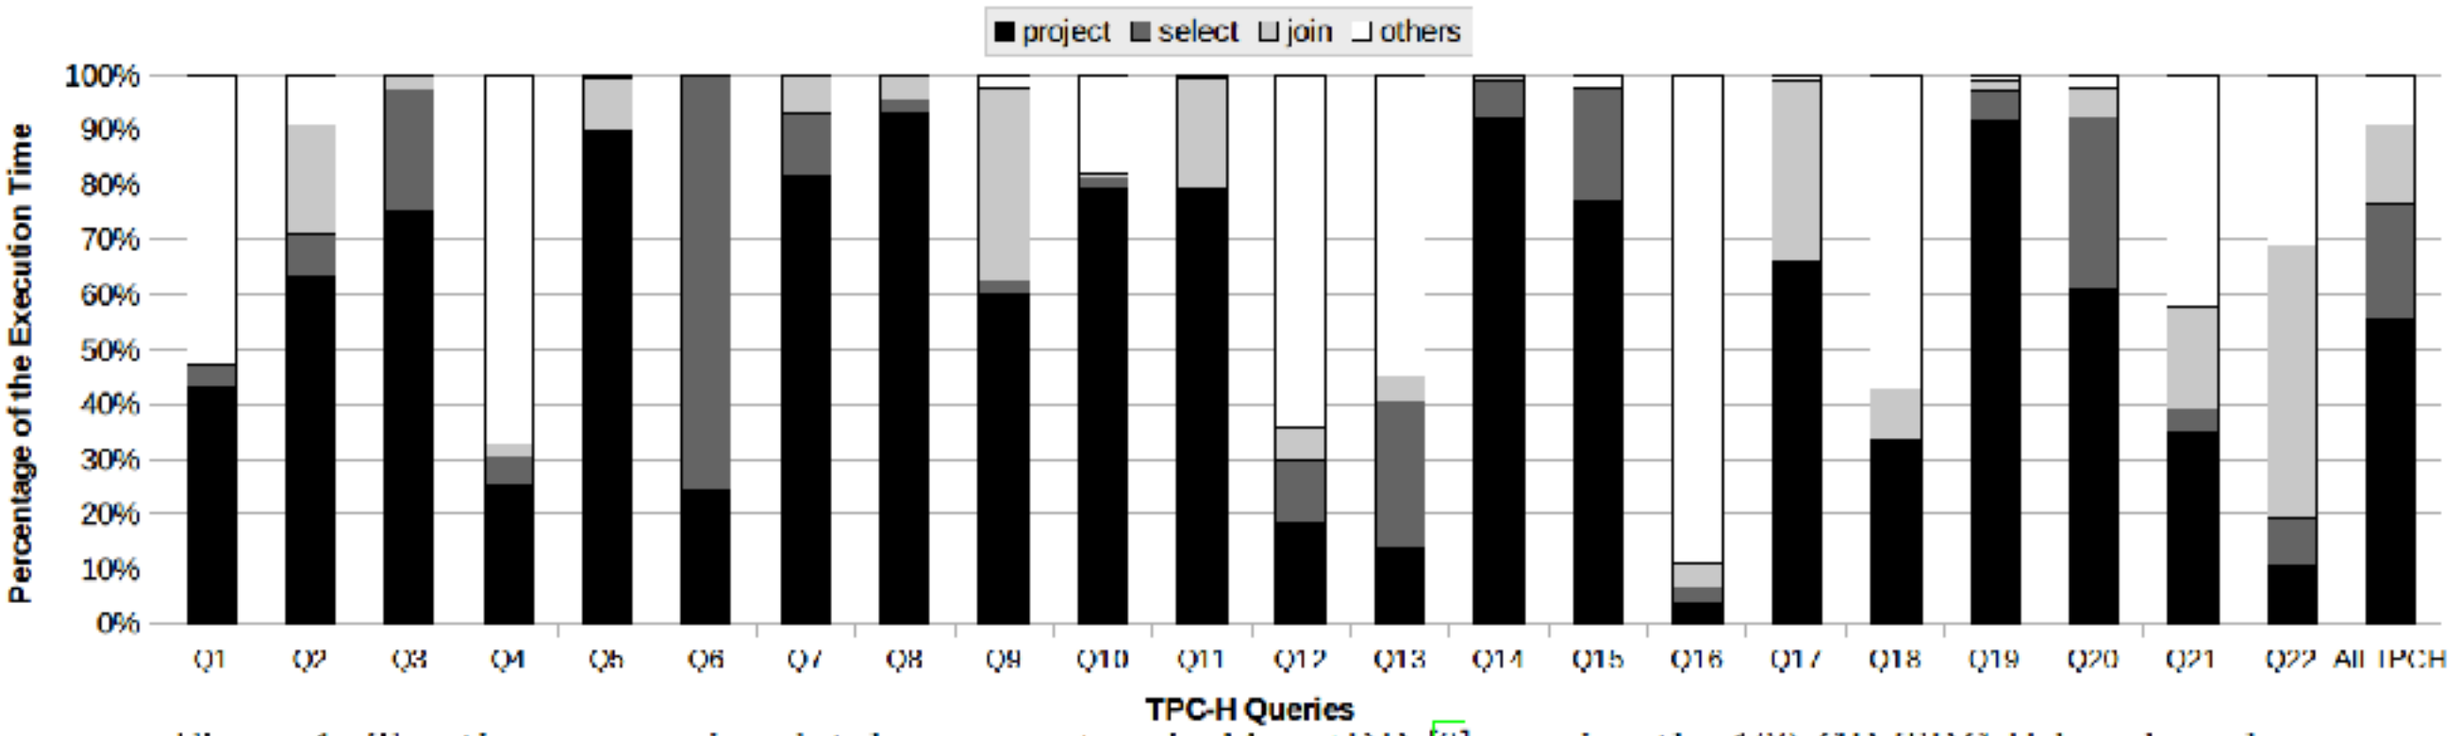

Figure 1: Top time consuming database operators in MonetDB [9] running the 100 GB TPC-H benchmark.

TPC-H Queries - operators breakdown

Monetdb: Two decades of research in column-oriented database architectures
Idreos, S. and Groffen, F and Nes, N. and Manegold, S. and Mullender, S. and Kersten, M.
IEEE Data Eng. Bull, 2012

QUERY PROCESSING: WHERE DOES TIME GO?

Figure 1: Top time consuming database operators in MonetDB [9] running the 100 GB TPC-H benchmark.

76% IS SPENT IN SELECTION AND PROJECTION

TPC-H Queries - operations breakdown

Monetdb: Two decades of research in column-oriented database architectures
Idreos, S. and Groffen, F and Nes, N. and Manegold, S. and Mullender, S. and Kersten, M.
IEEE Data Eng. Bull, 2012

DATA MOVEMENT IN DATABASE SYSTEMS

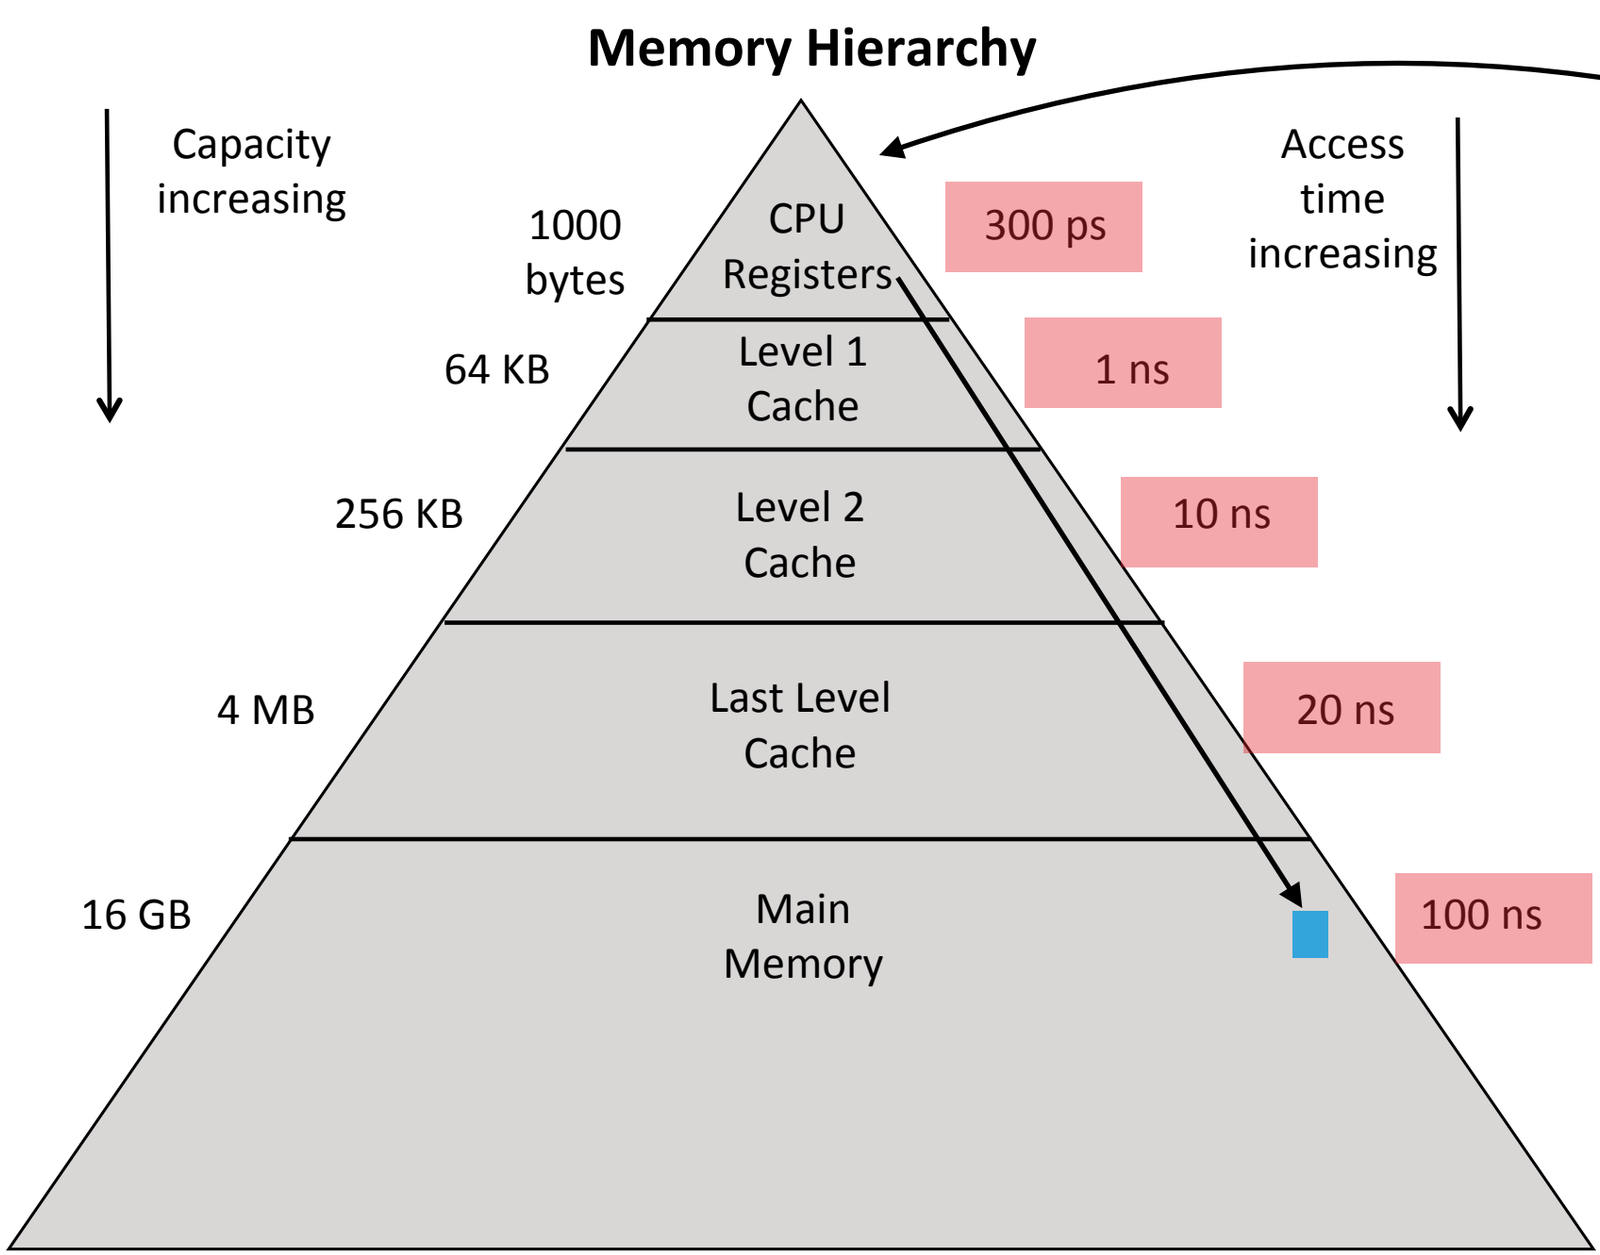

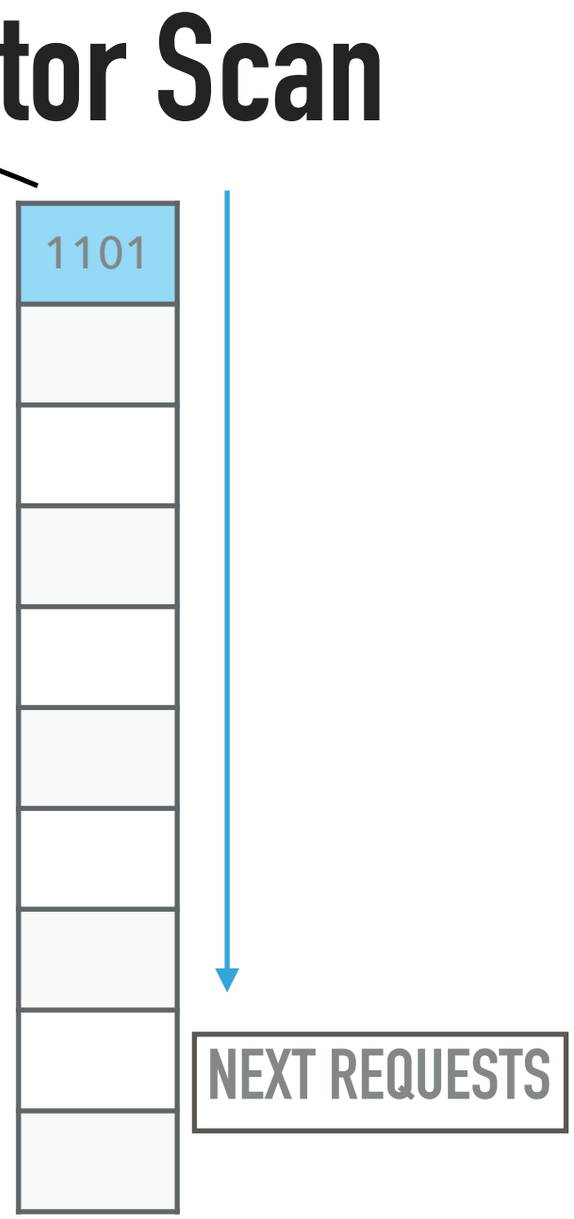

DATA MOVEMENT IN DATABASE SYSTEMS

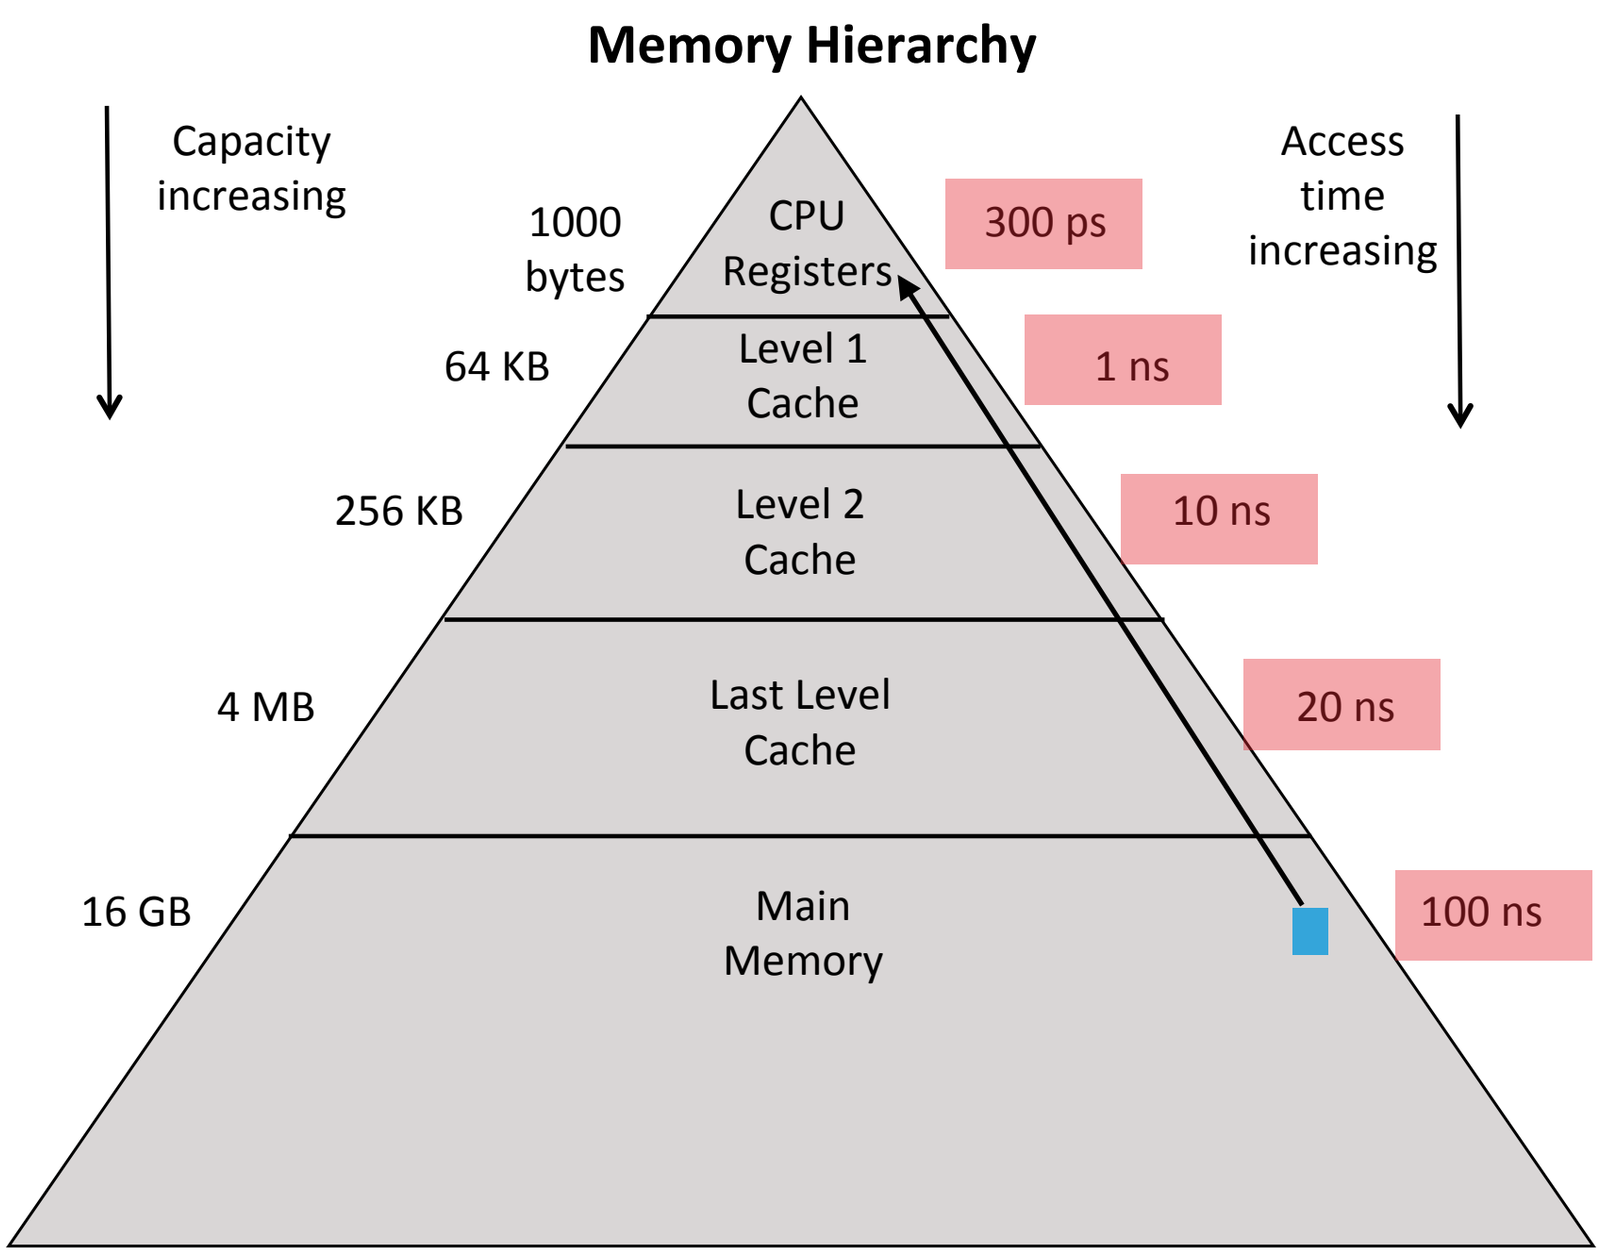

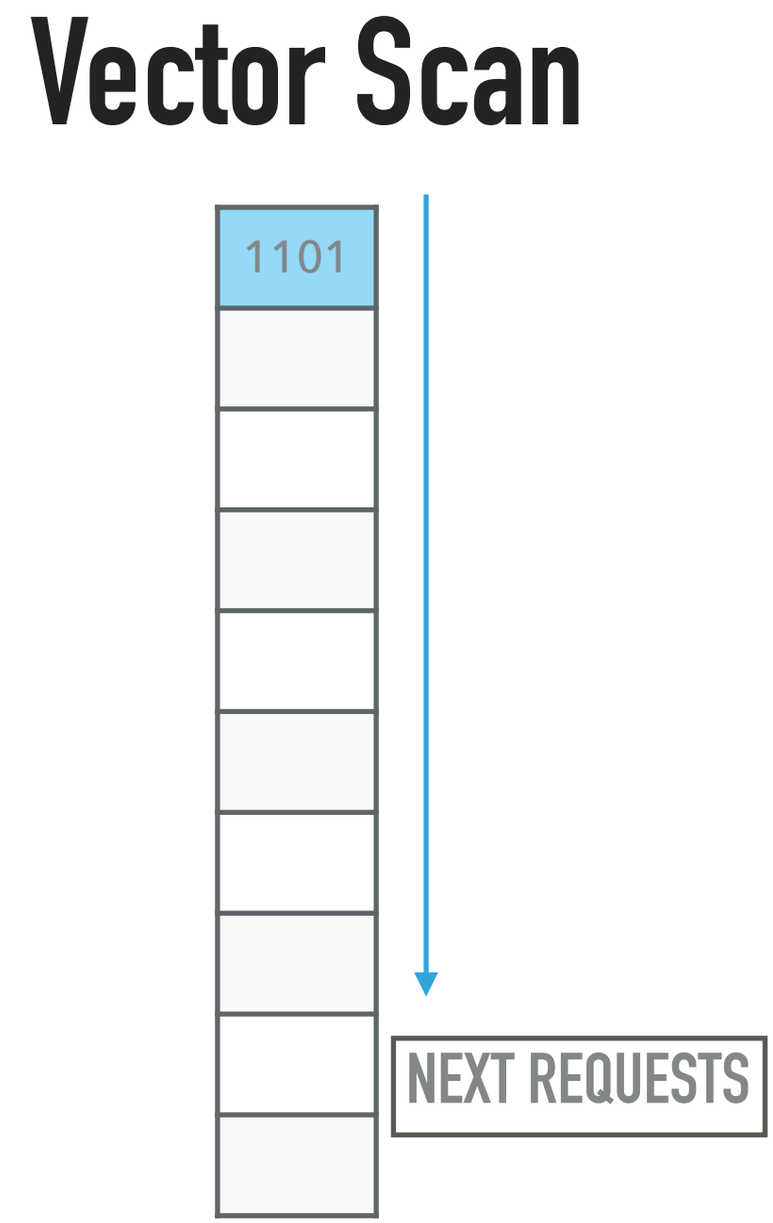

NEAR-DATA PROCESSING (NDP)

- Cost of moving data is significant
- Performs the data processing in the main memory
- Reduces energy consumption (data movement)
- Provides faster response time (no cache latency)
- With HMC the execution inside memory became tangible

Near-Data Processing: Insights from a MICRO-46 Workshop
R. Balasubramonian, J. Chang, T. Manning, J. H. Moreno, R. Murphy, R. Nair, and S. Swanson.
IEEE Micro, 34(4):36–42, 2014

Hybrid memory cube
J. Jeddeloh and B. Keet.
Symposium on VLSI (VLSIT), 2012

HYBRID MEMORY CUBE (HMC)

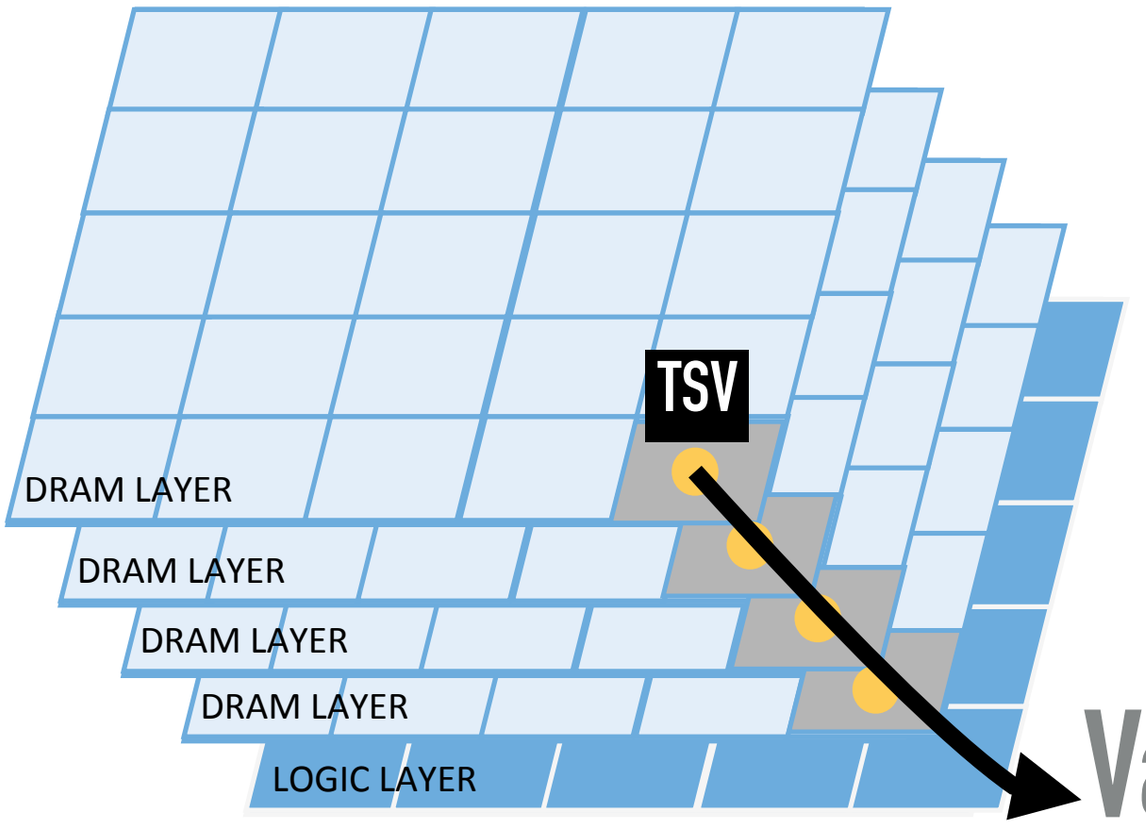

Hybrid memory cube
J. Jeddeloh and B. Keet.
Symposium on VLSI (VLSIT), 2012

TSV - THROUGH-SILICON VIAS

HYBRID MEMORY CUBE (HMC)

- Composed of 32 vaults
- **Bandwidth up to 320GB/s**
- **Instruction execution (16 bytes at a time)**
- 256 Bytes of row buffer

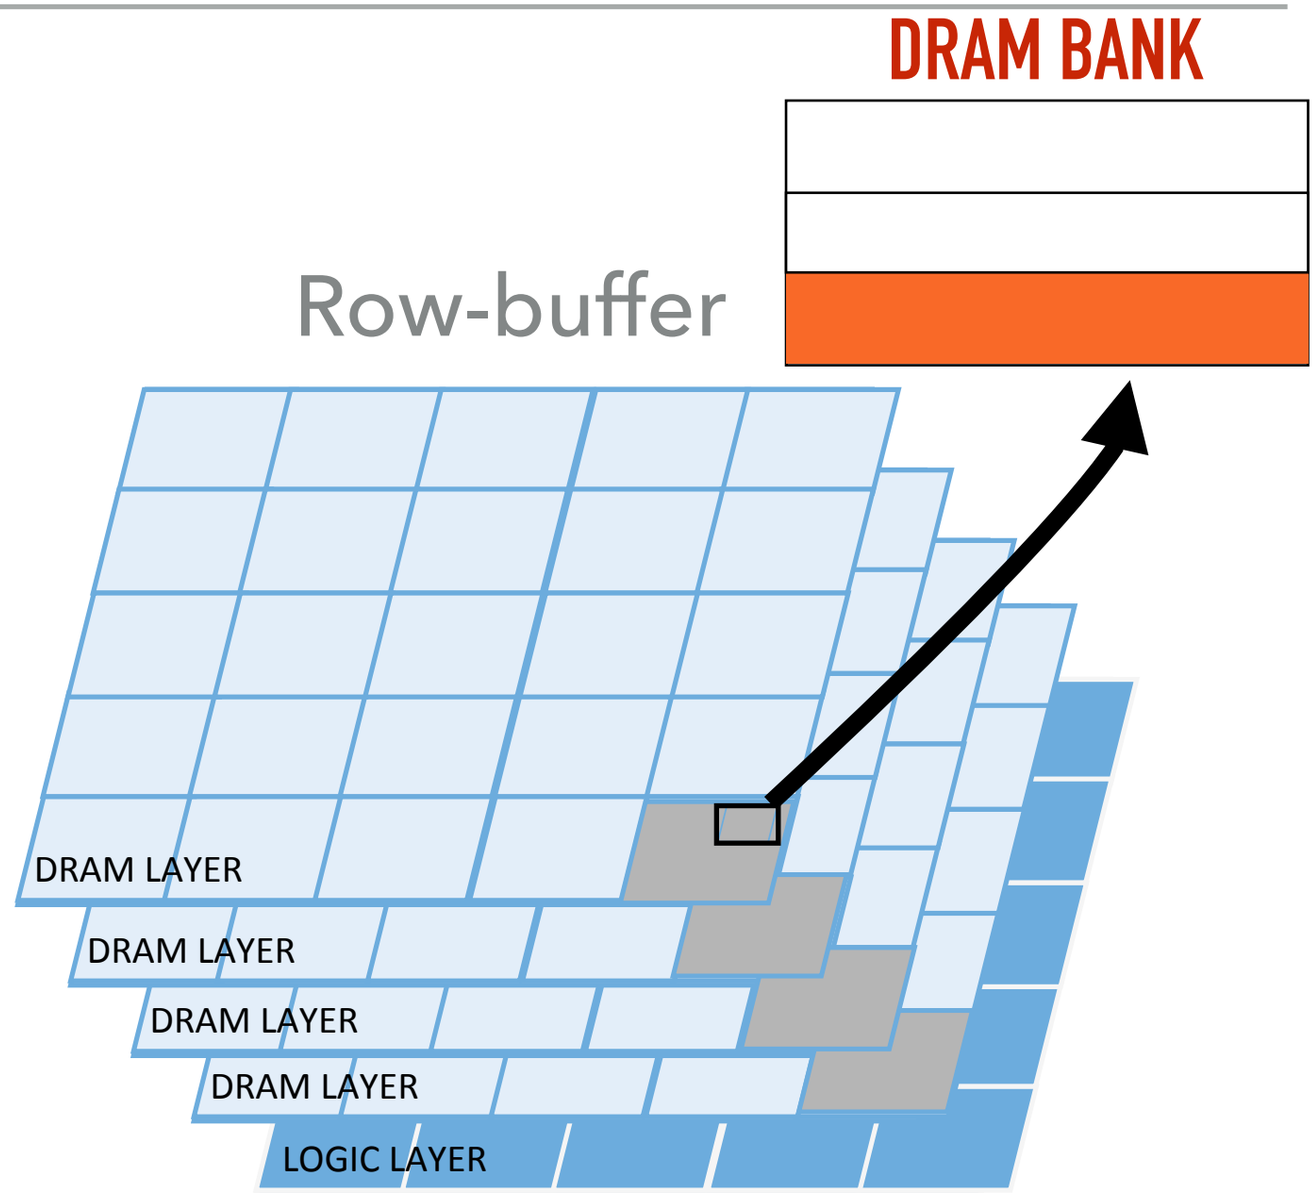

Hybrid memory cube
J. Jeddeloh and B. Keet.
Symposium on VLSI (VLSIT), 2012

1. What happens when database systems run the select scan over the current x86 architecture using the HMC as ordinary DRAM?

Research Question

DATABASE SYSTEMS PREDICATE PROCESSING

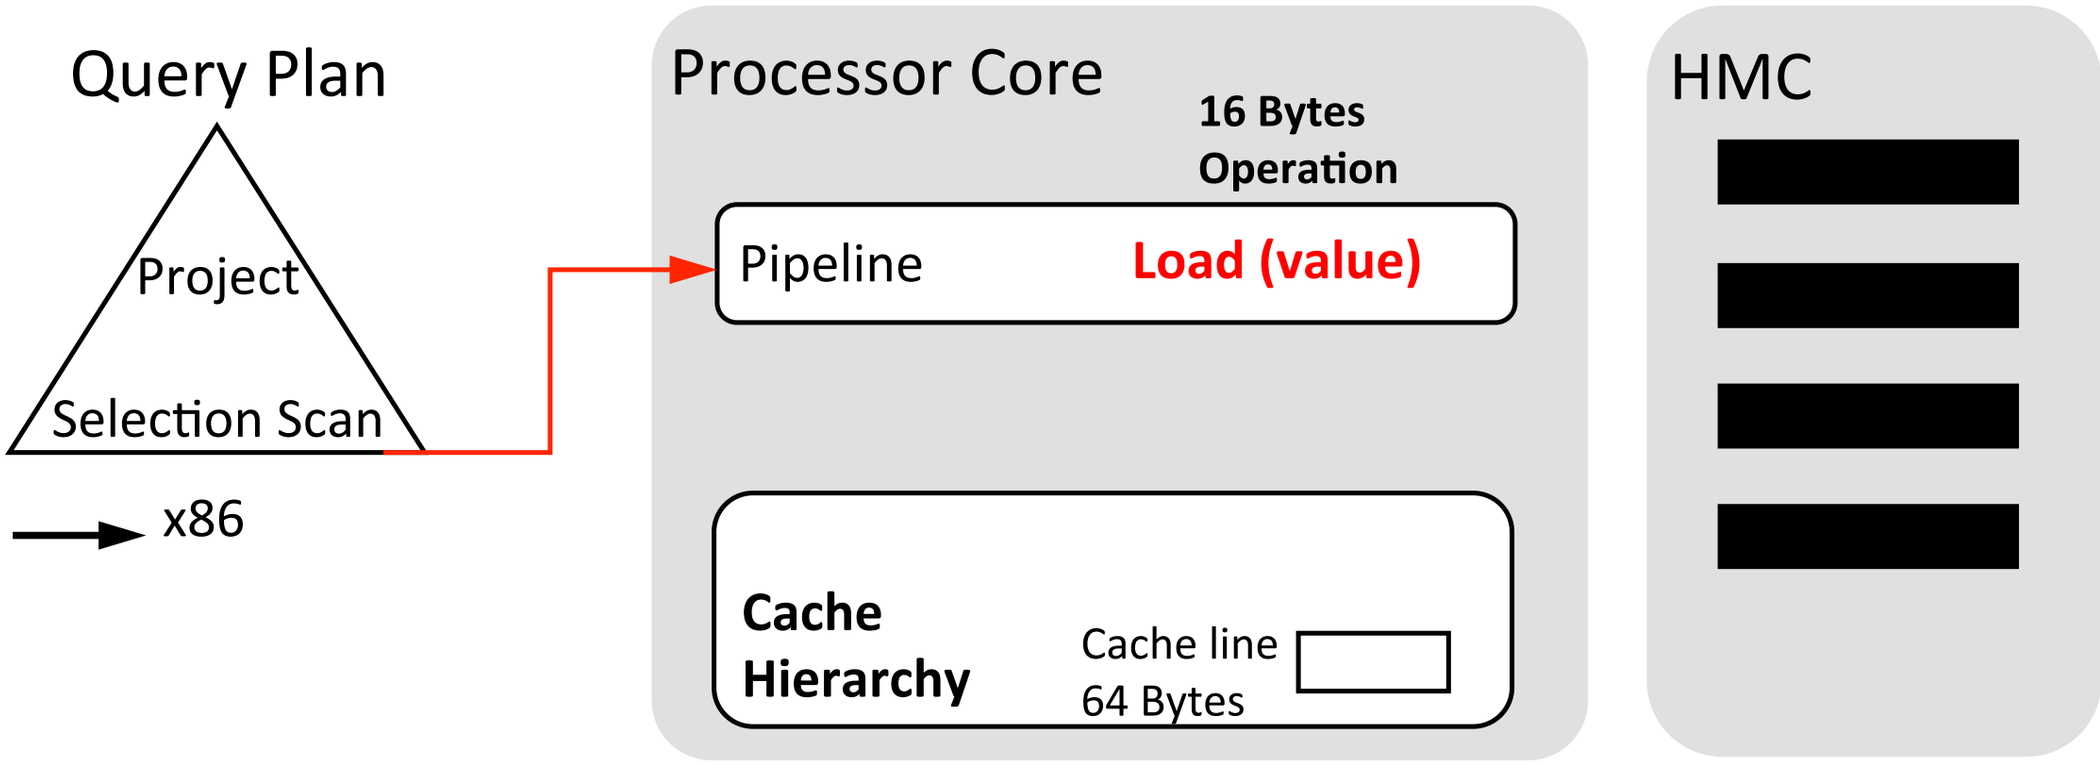

"SELECT<COLUMNS> FROM <TABLES> **WHERE**<CONDITION> AND/OR <CONDITION>"

DATABASE SYSTEMS PREDICATE PROCESSING

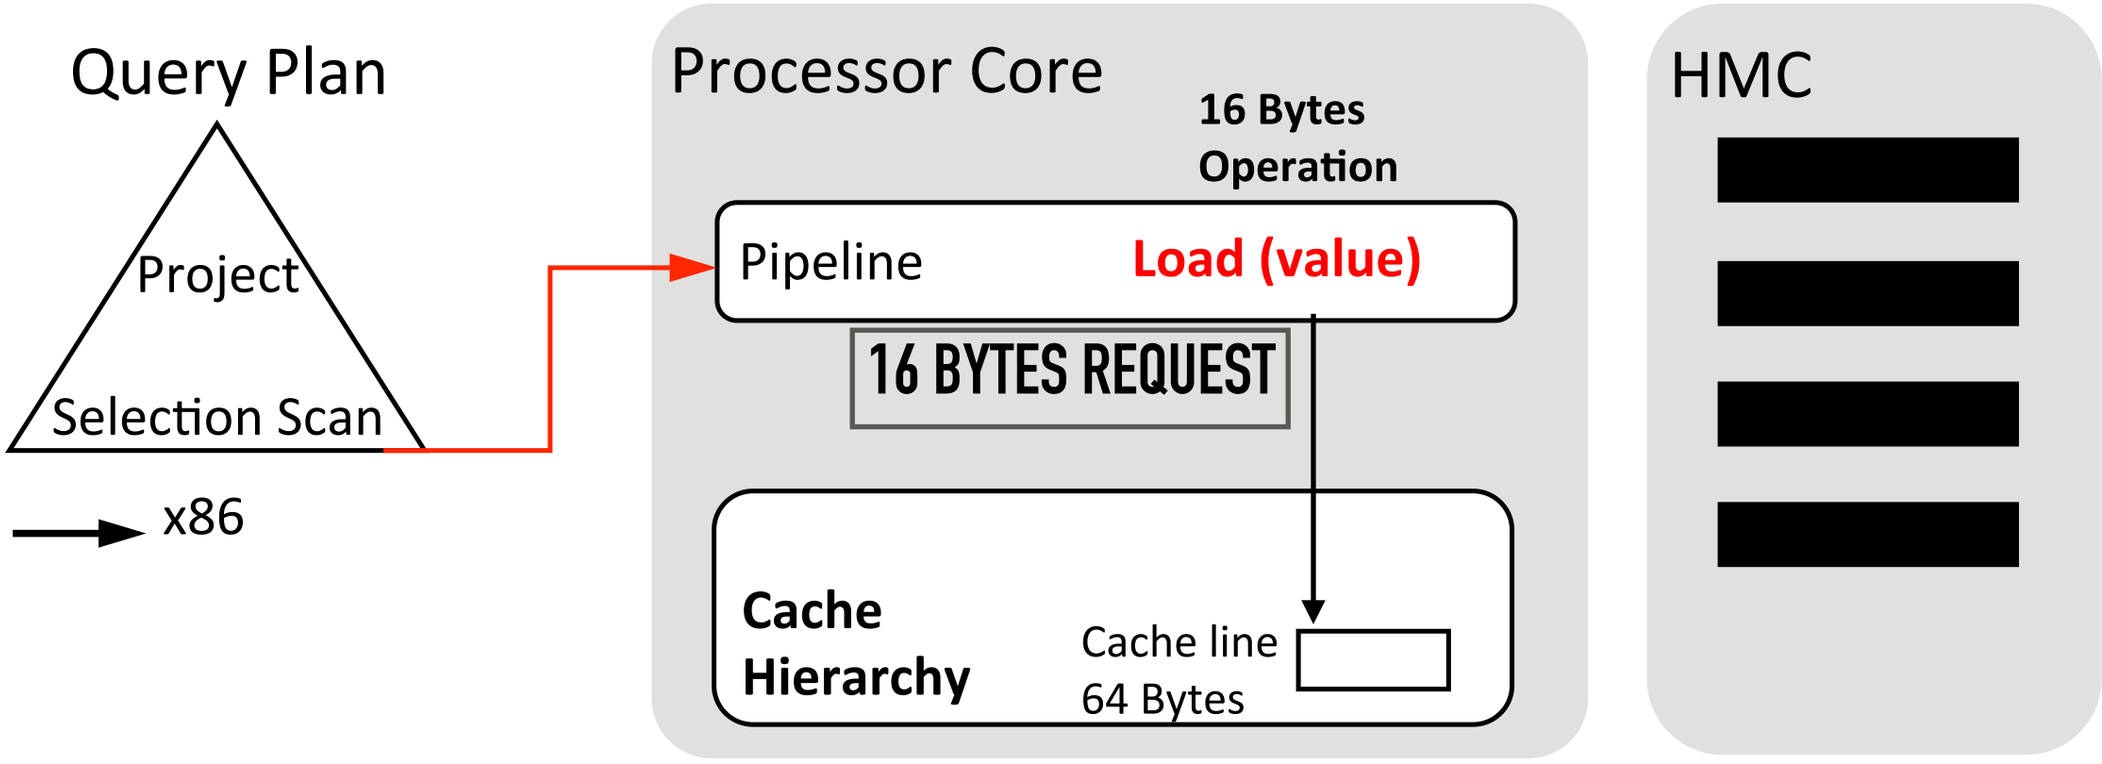

DATABASE SYSTEMS PREDICATE PROCESSING

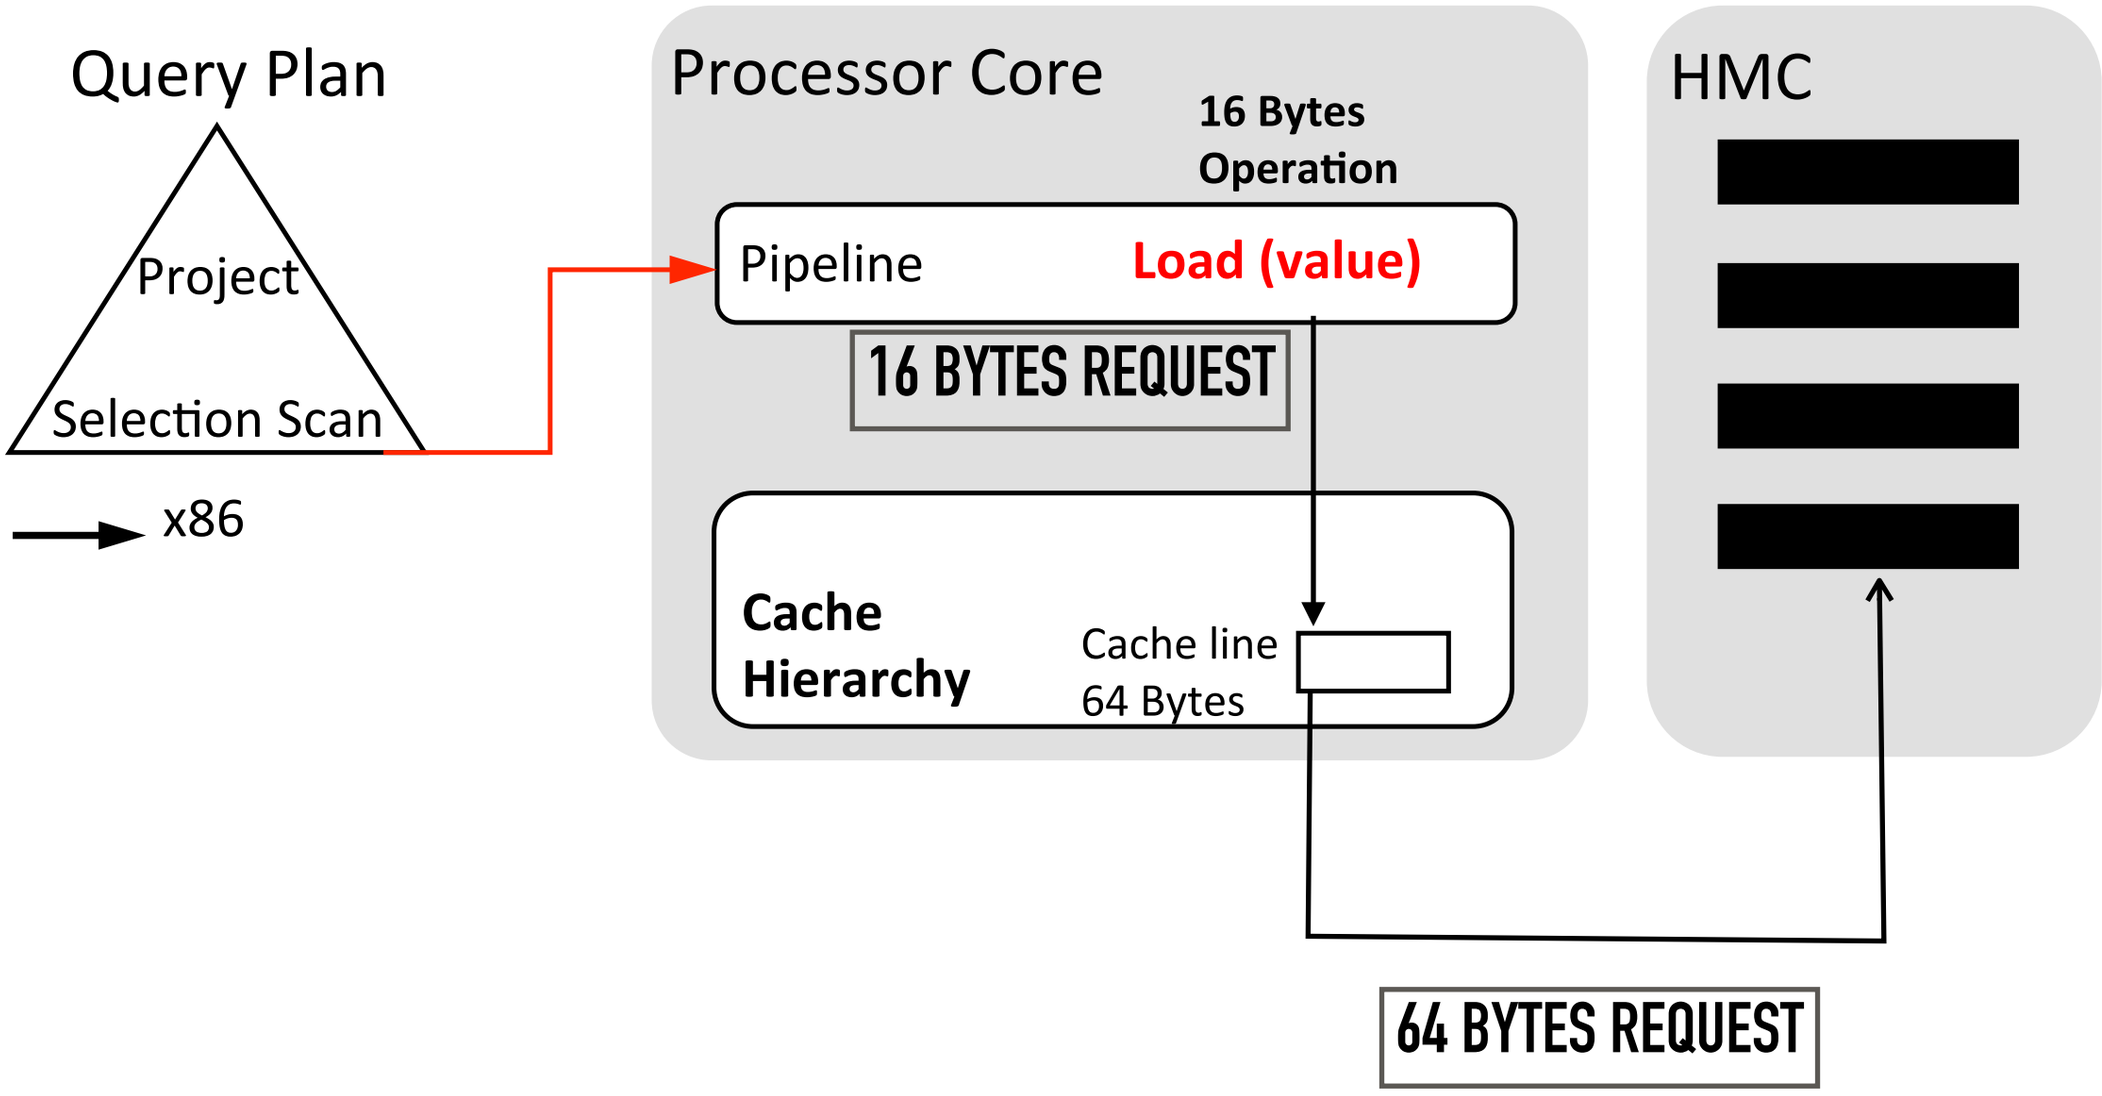

DATABASE SYSTEMS PREDICATE PROCESSING

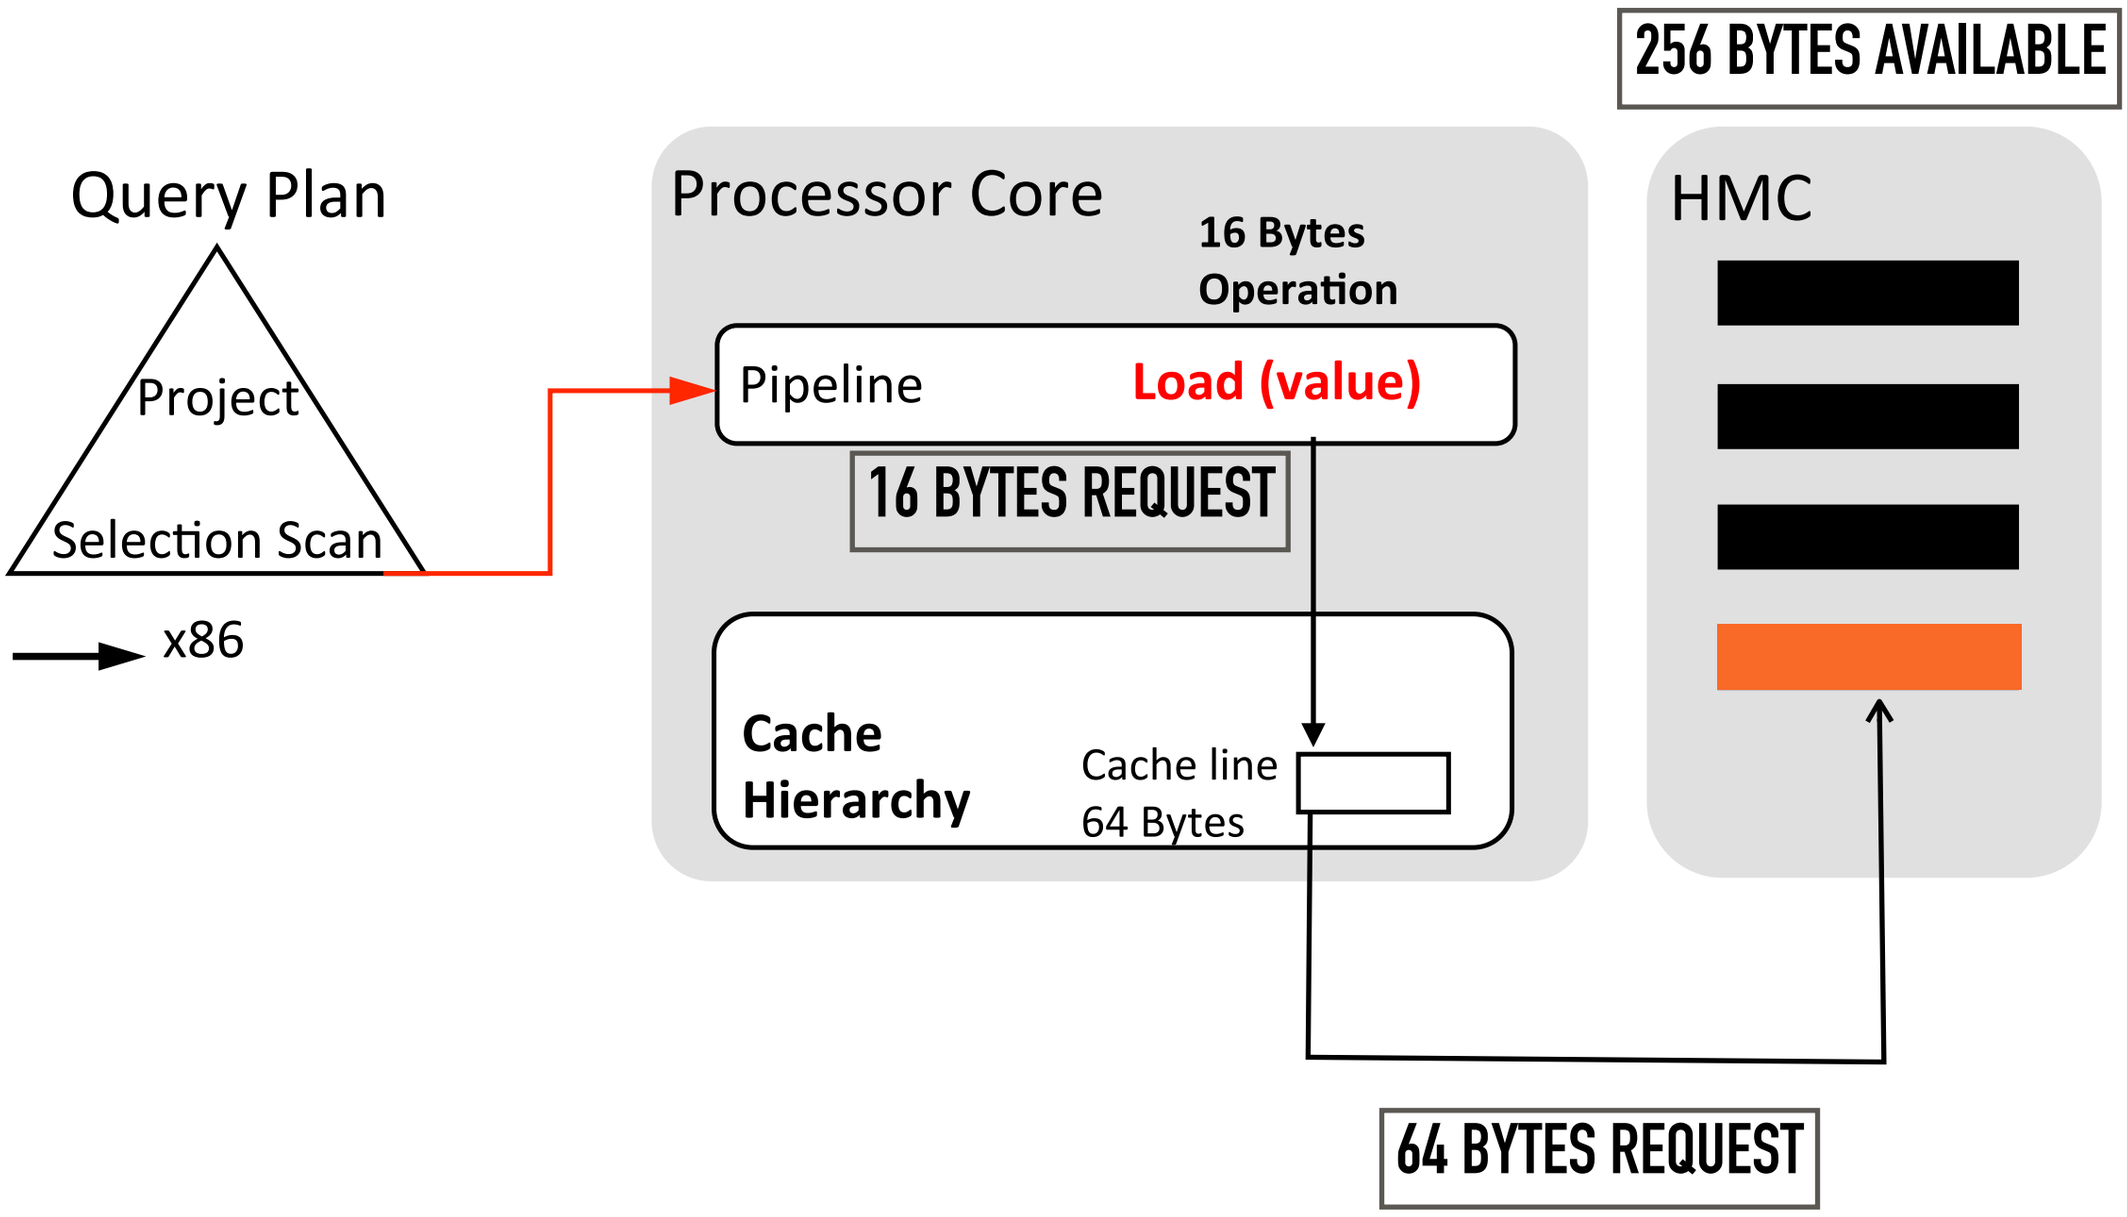

DATABASE SYSTEMS PREDICATE PROCESSING

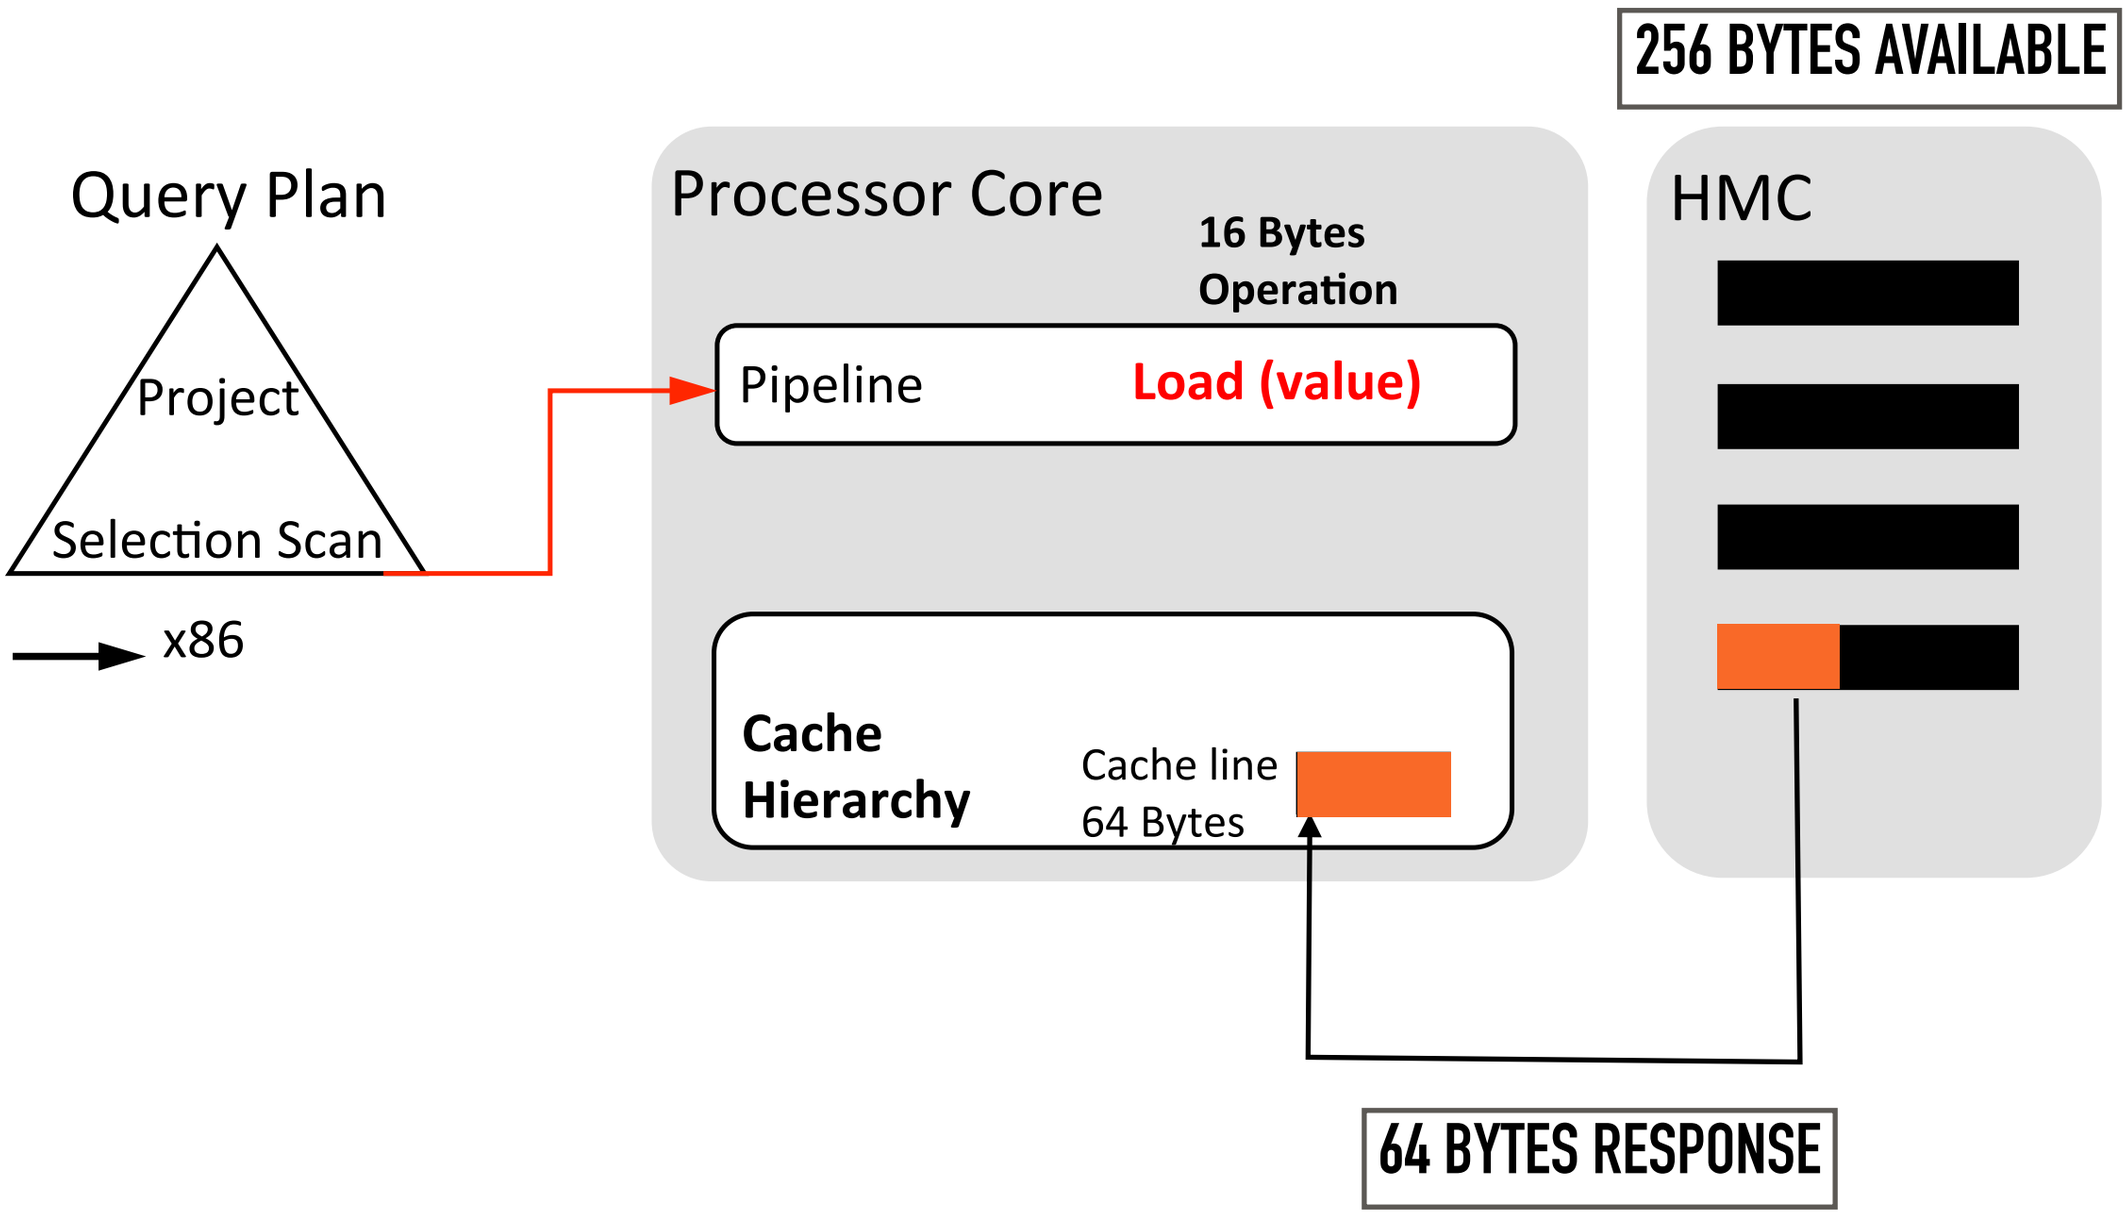

DATABASE SYSTEMS PREDICATE PROCESSING

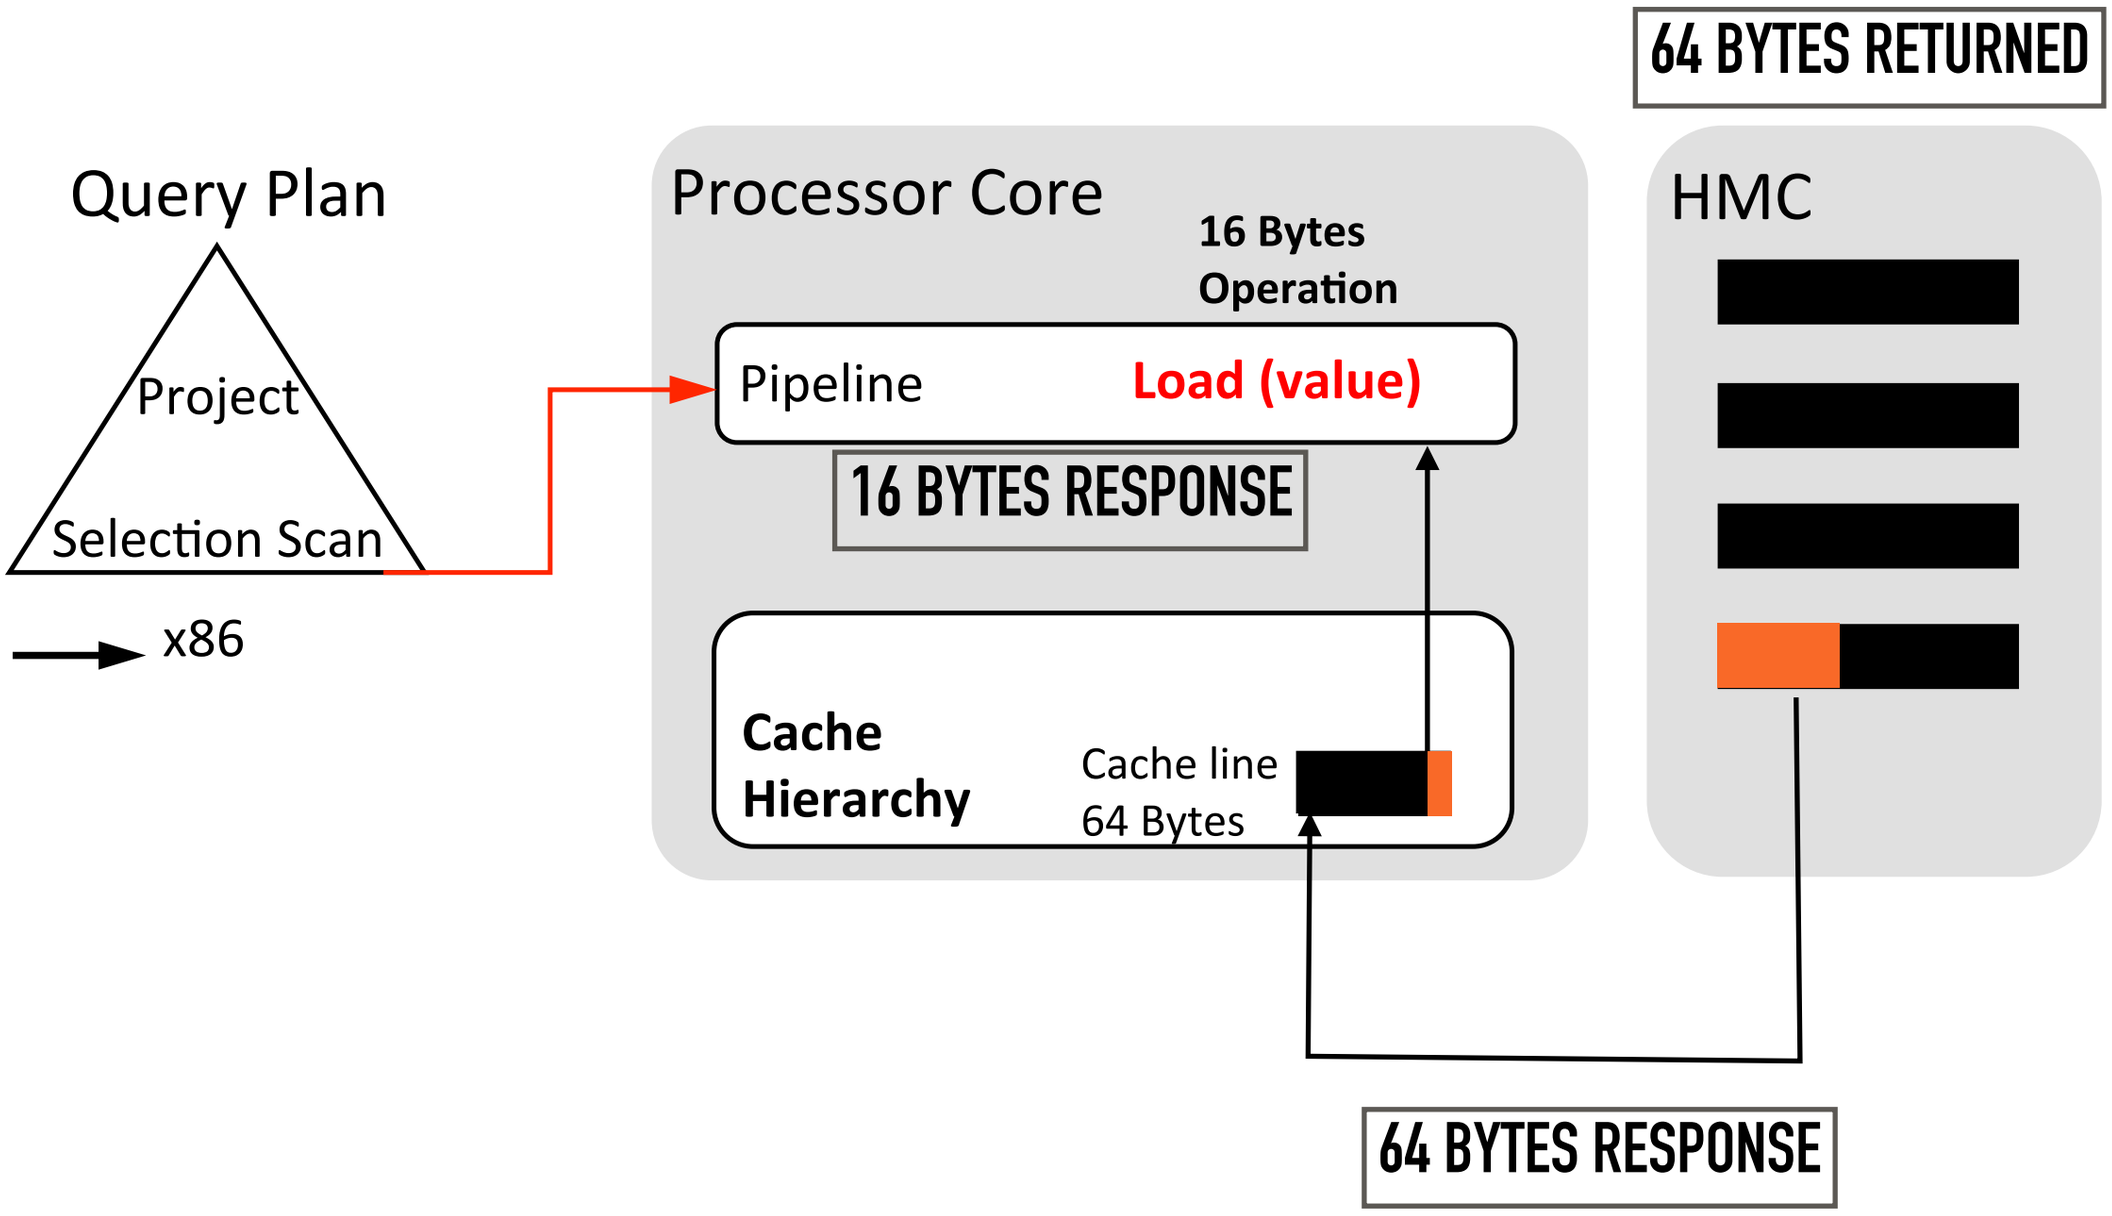

DATABASE SYSTEMS PREDICATE PROCESSING

DATA-MOVEMENT

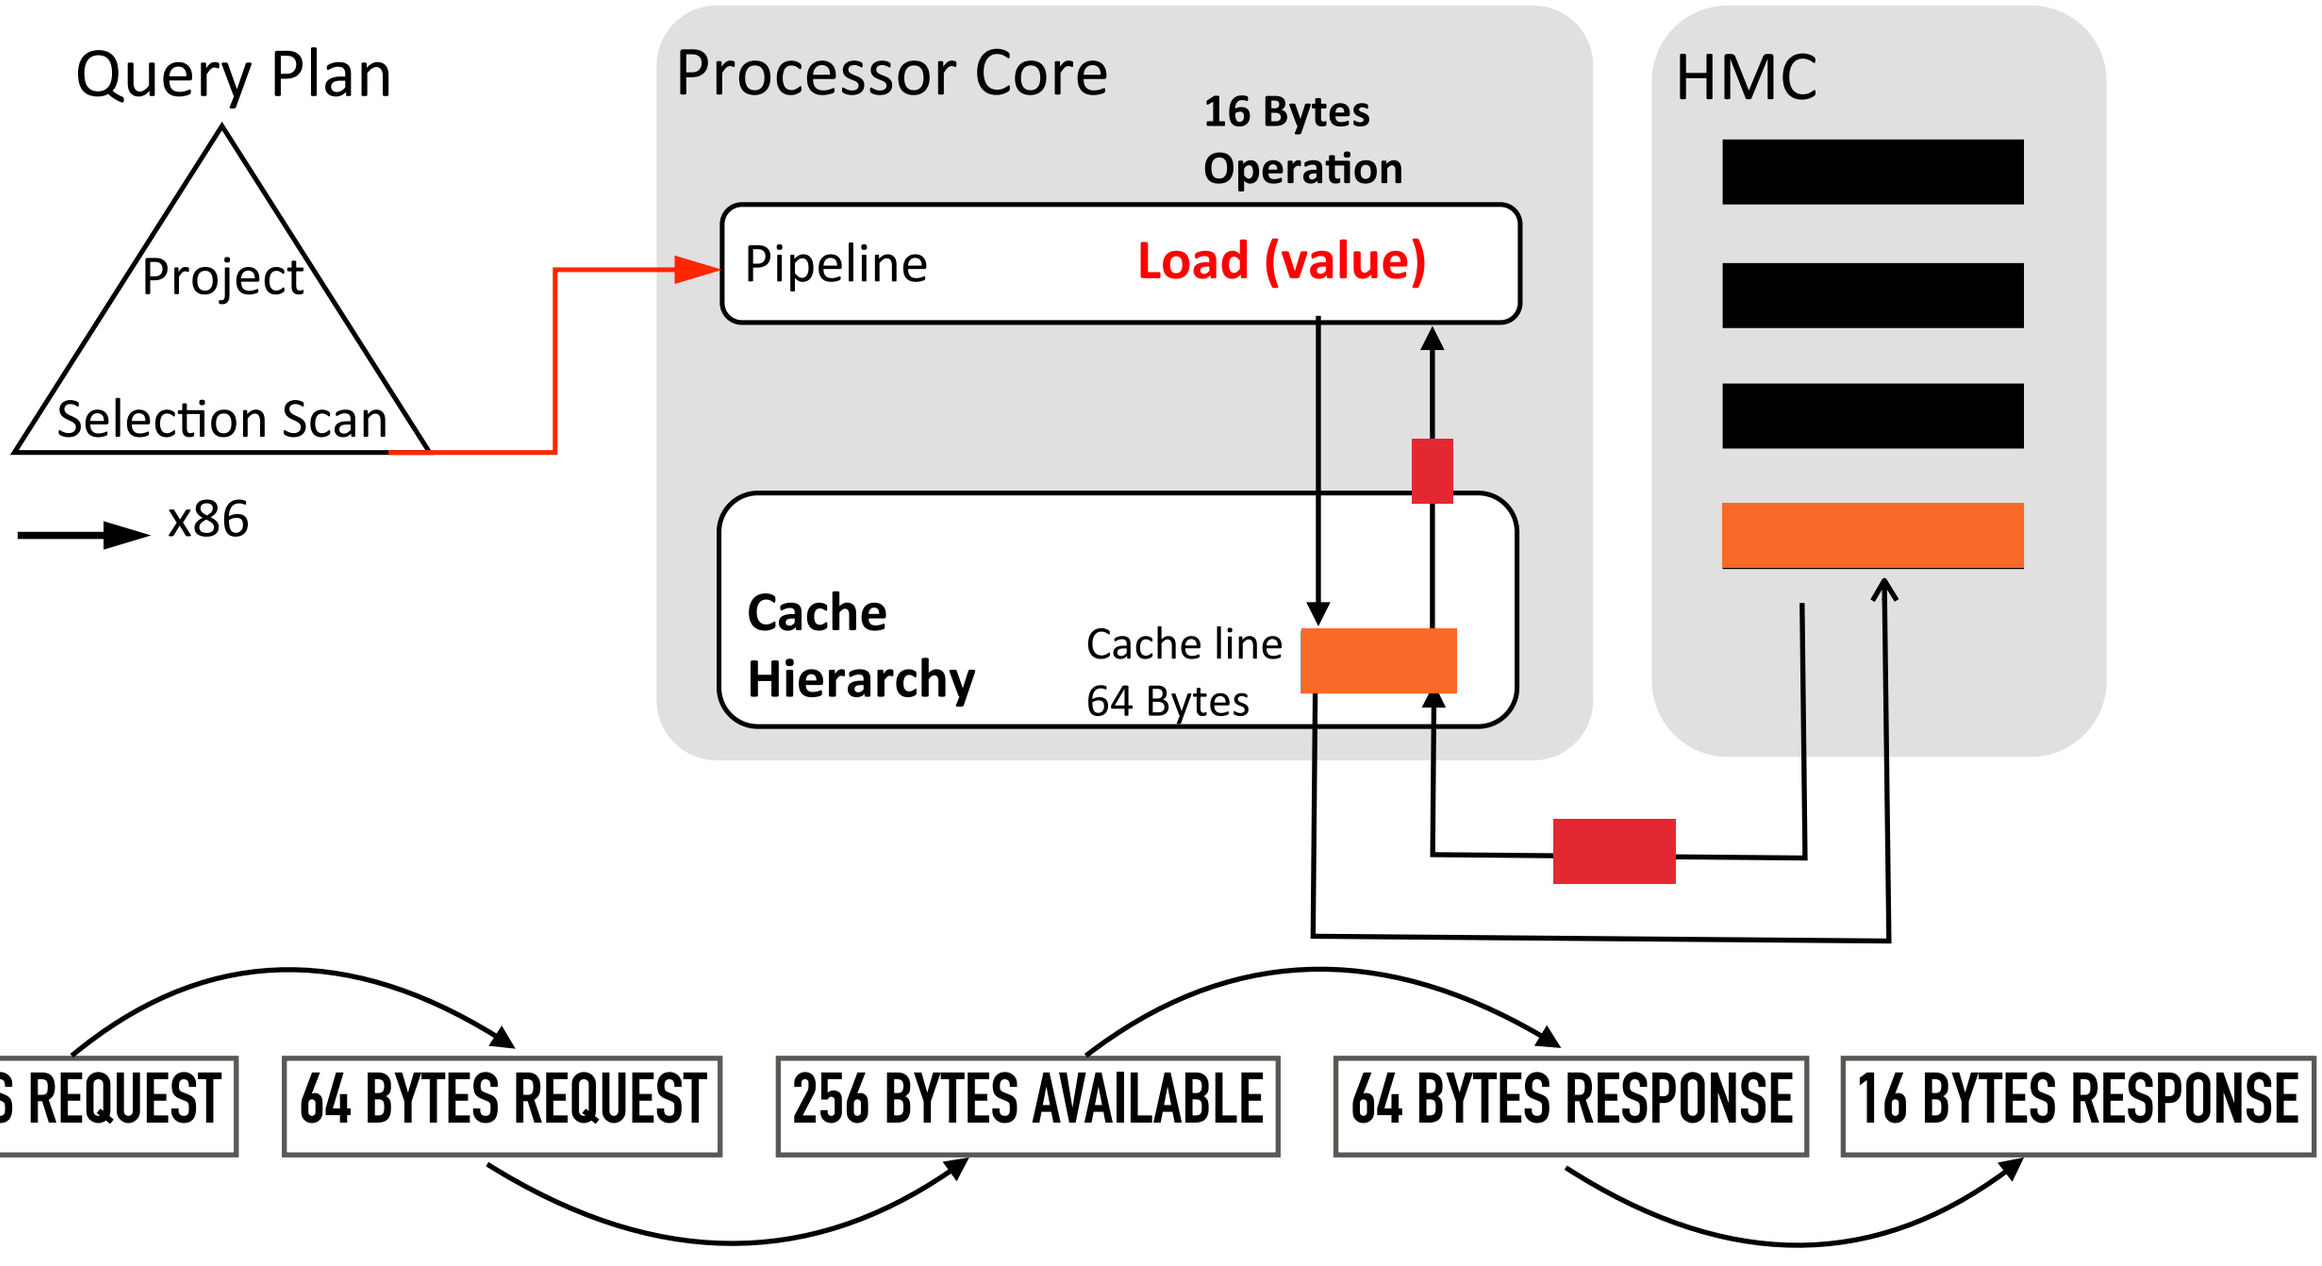

DATABASE SYSTEMS PREDICATE PROCESSING
DO MITIGATE DATA-MOVEMENT

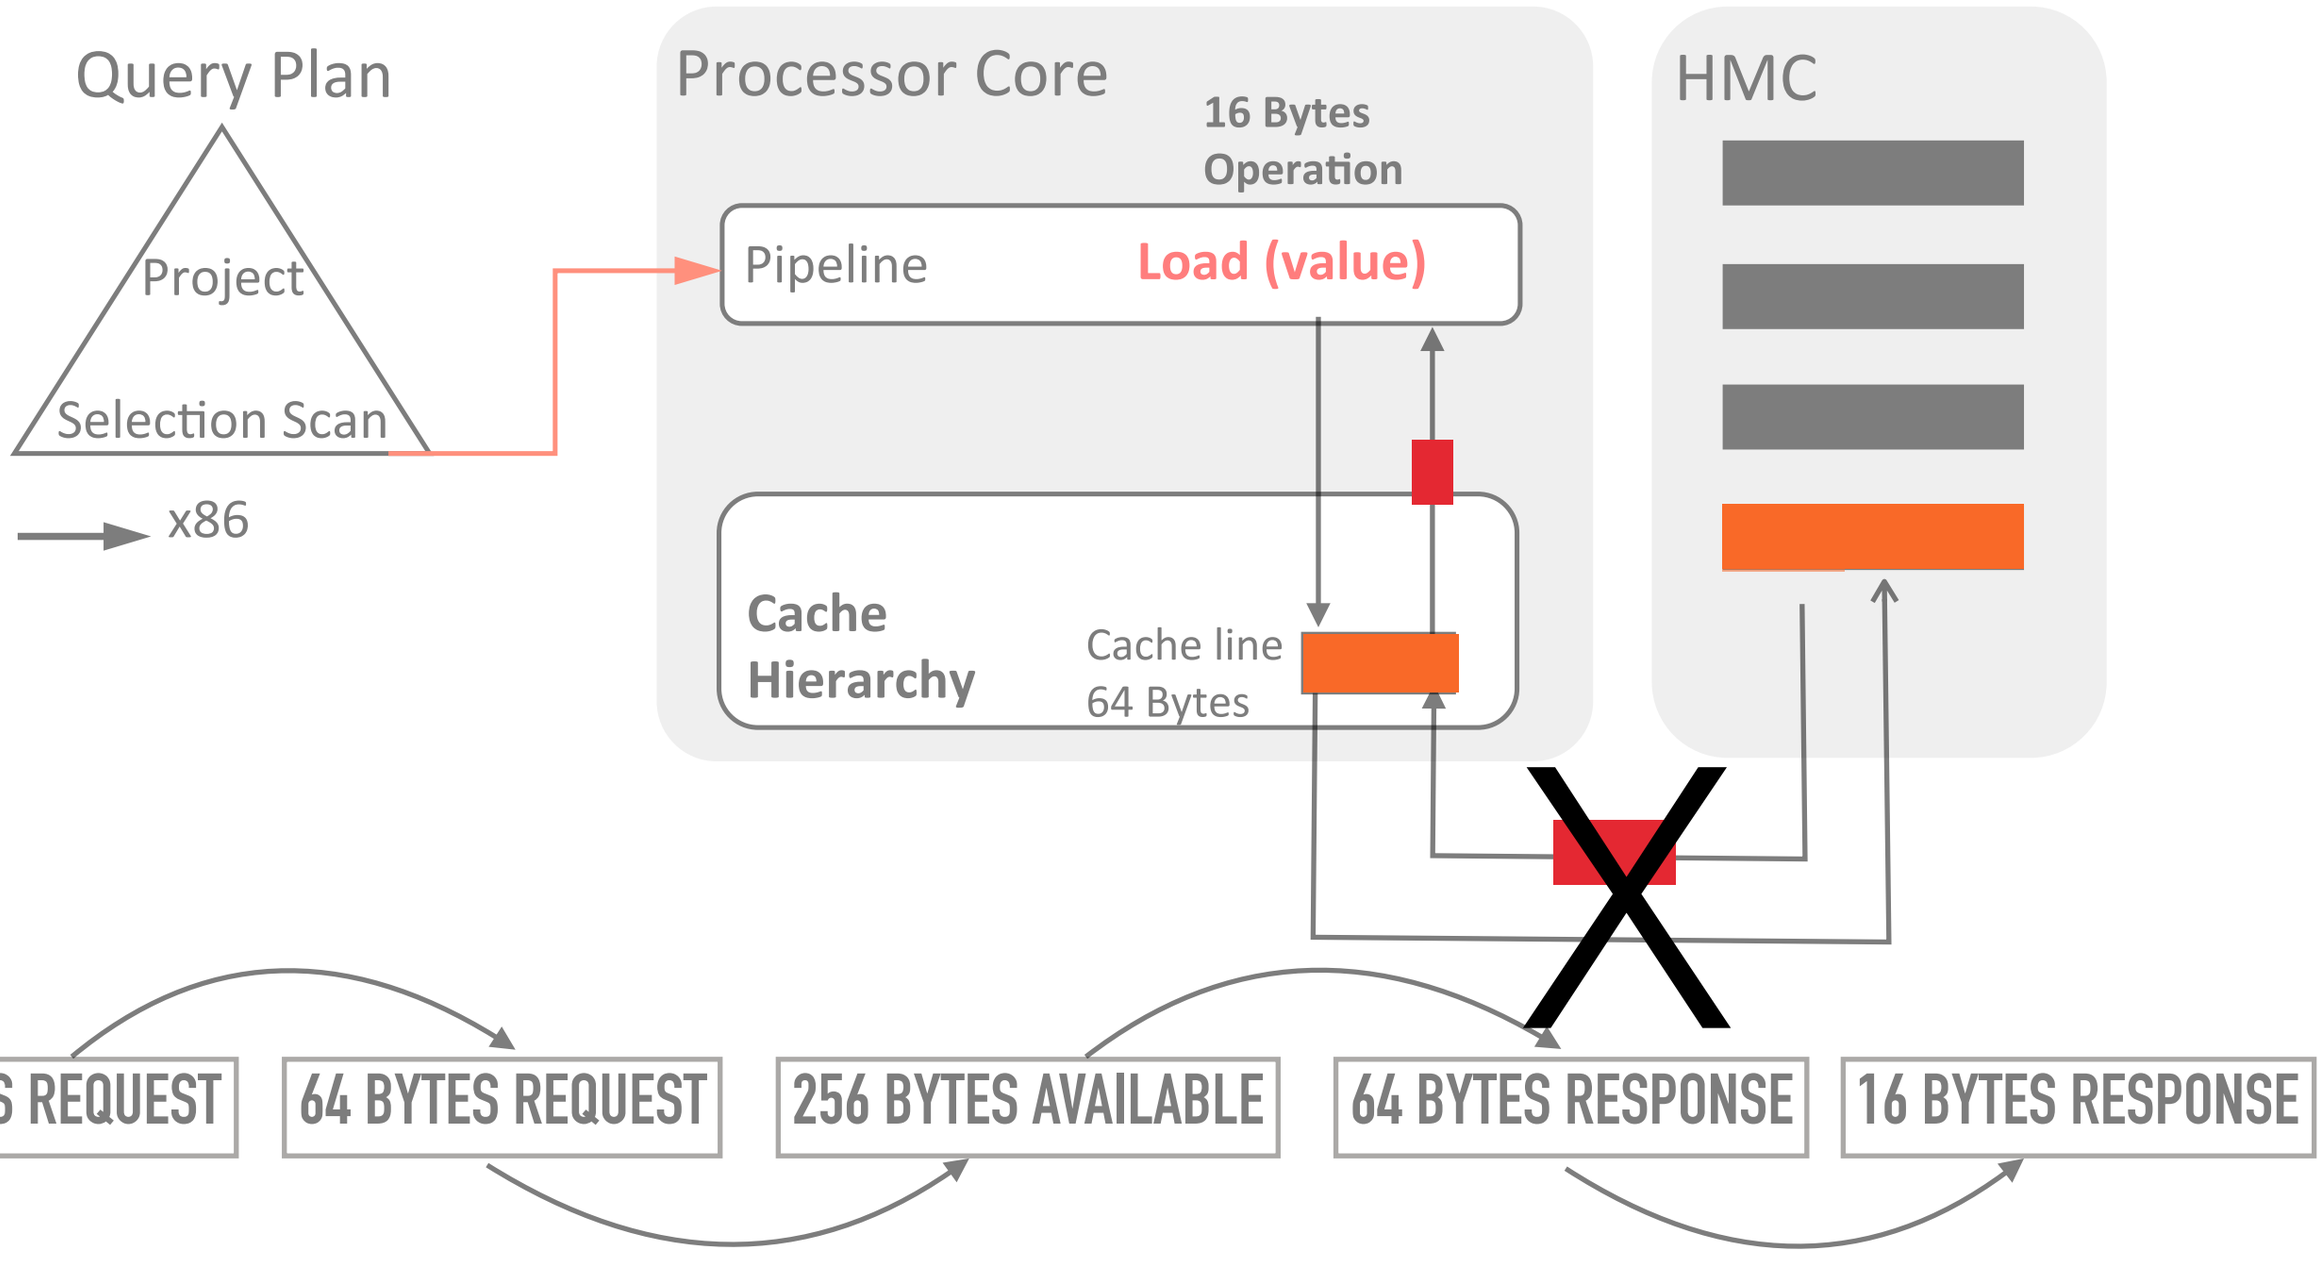

NEAR-DATA FILTERS FOR DATABASE SYSTEMS

- Allow the execution of filtering inside the HMC
- Take advantage of the HMC instruction execution and high bandwidth capacity to mitigate the “memory wall”
- Increase the parallelism of the filtering execution
- With extensions for the HMC ISA filtering is done by branch-less decisions

Near Data Filters: Taking Another Brick from the Memory Wall
Diego G. Tomé, Tiago R. Kepe, Marco A. Z. Alves, and Eduardo C. Almeida.
ADMS@VLDB, 2018

Hipe: Hmc Instruction Predication Extension Applied on Database Processing
Diego G. Tomé, Marco A. Z. Alves, and Eduardo C. Almeida.
DATE, 2018

1. What happens when database systems run the select scan over the current x86 architecture using the HMC as ordinary DRAM? Memory wall is still a problem.

2. Can we use the current HMC Instruction Set Architecture (ISA) to implement the near-data filter?

Research Question

HMC-SCAN LIMITATIONS

```
SELECT COLUMN
FROM TABLE
WHERE COLUMN < 10
```

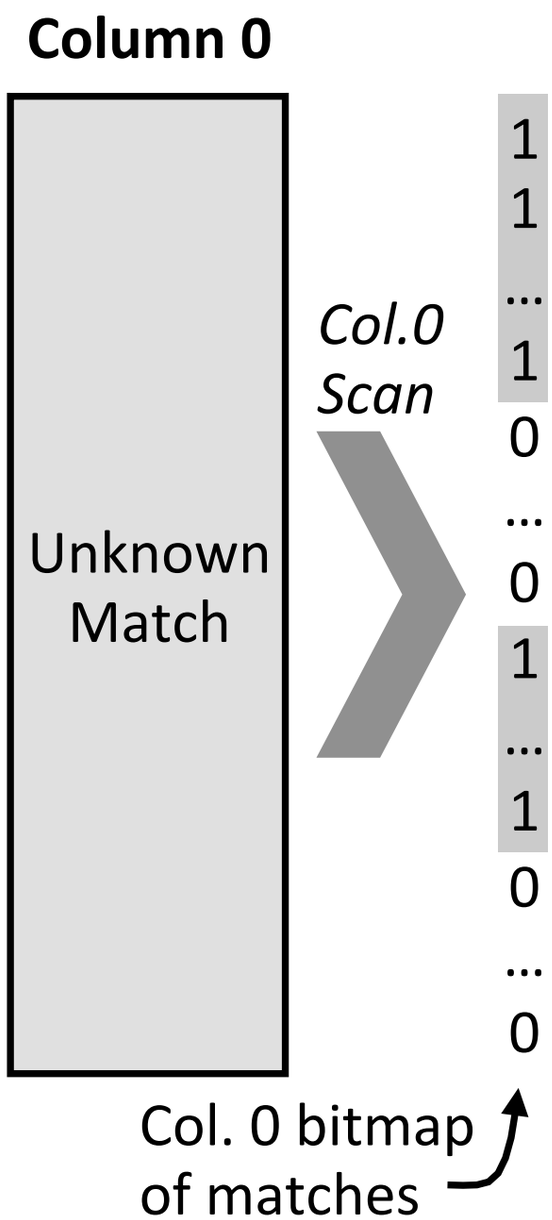

HMC-SCAN LIMITATIONS

```
SELECT COLUMN
FROM TABLE
WHERE COLUMN < 10
AND COLUMN 1 < 10
```

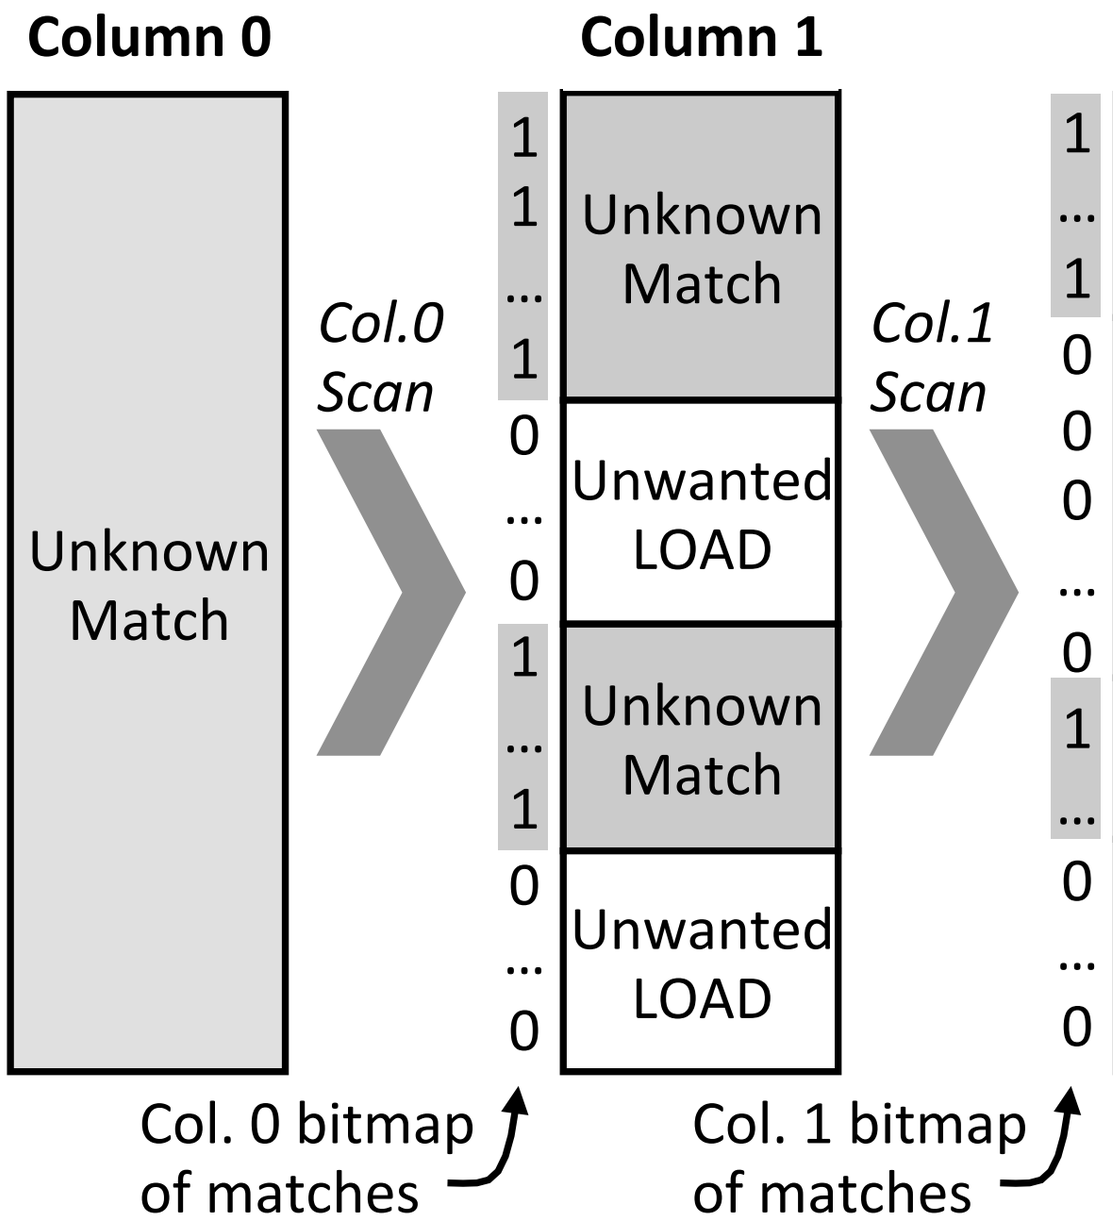

HMC-SCAN LIMITATIONS

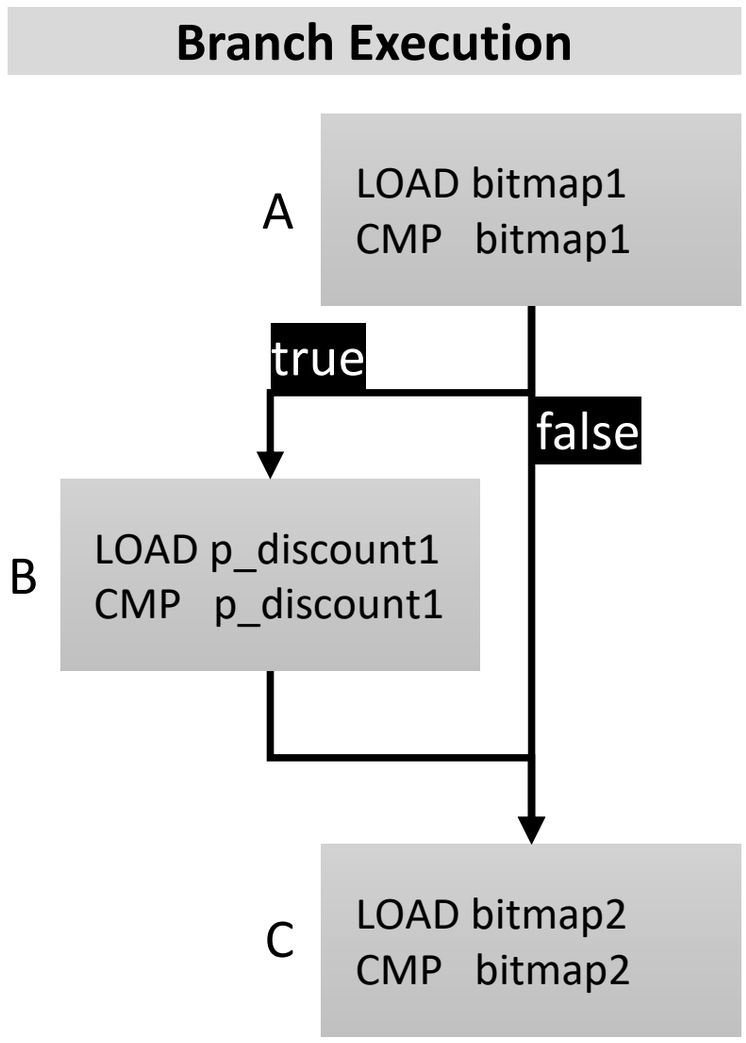

HMC-SCAN LIMITATIONS

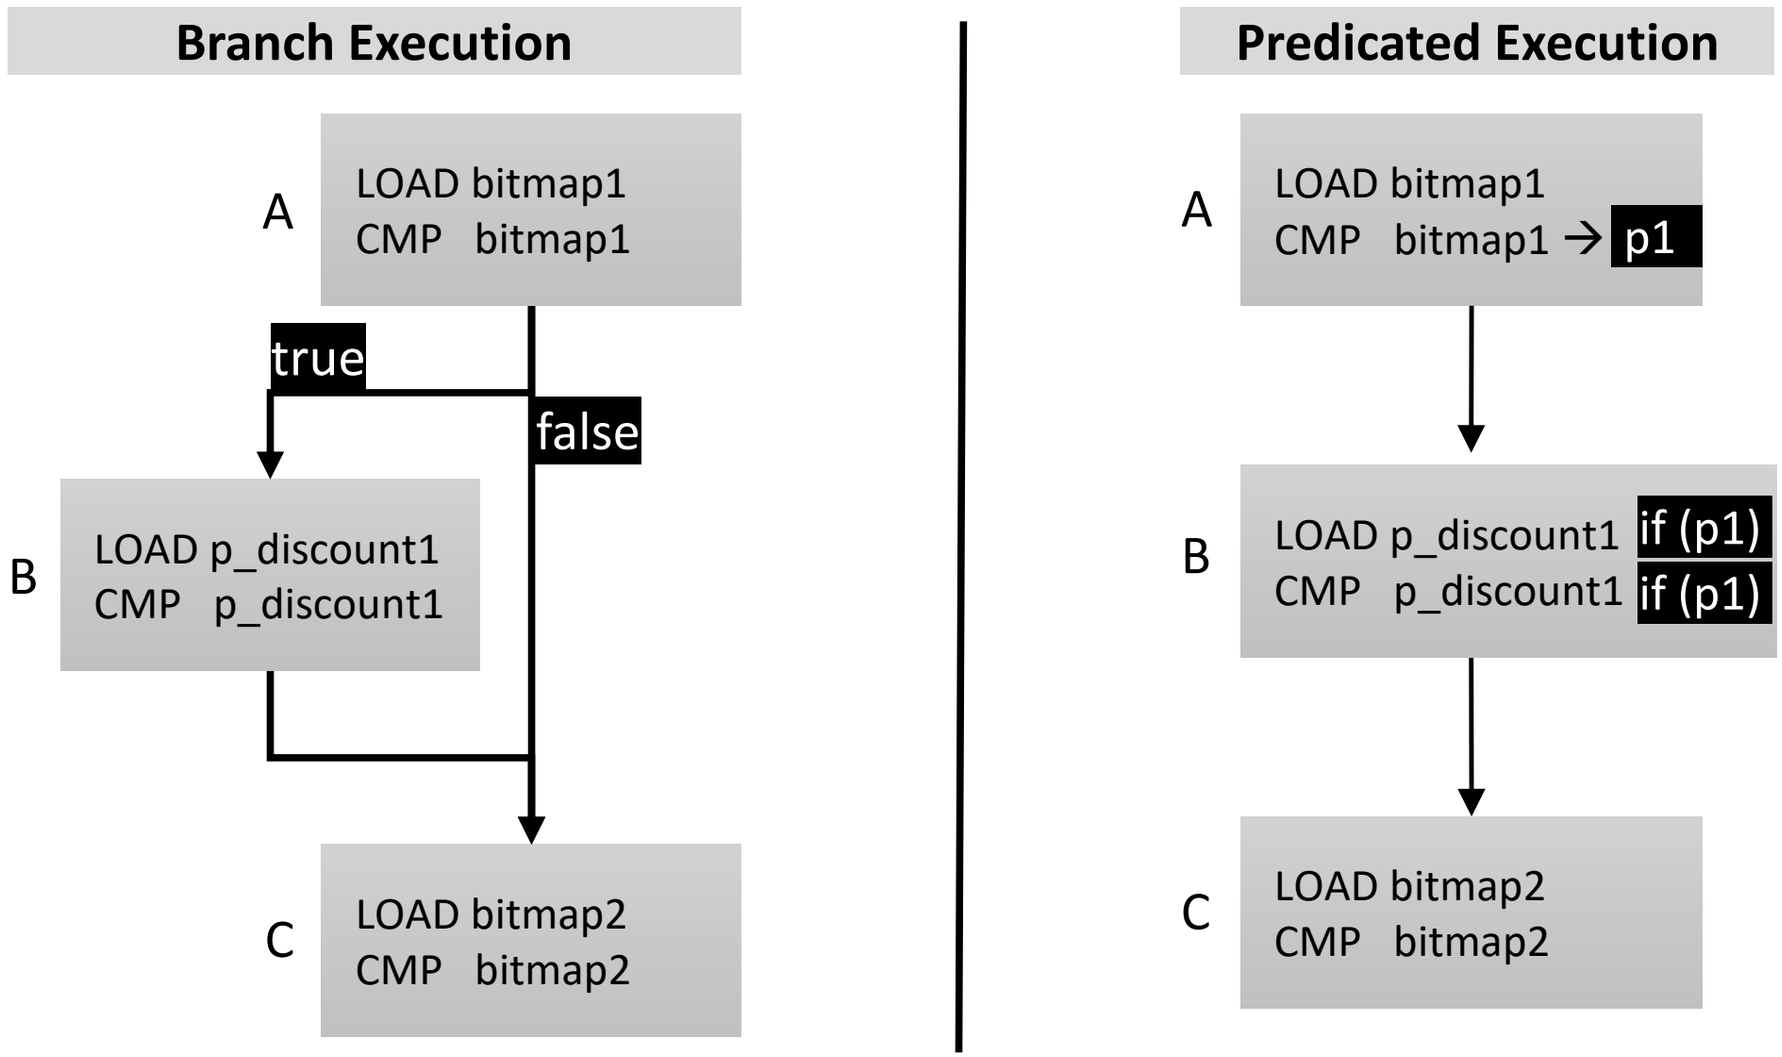

3.What are the extensions to the HMC ISA to leverage the full potential of the hardware and reduce the interleaving with x86 instructions?”

Research Question

HIPE-FILTER: HMC INSTRUCTION PREDICATION EXTENSION

- A comparison instruction between address and immediate
- 32 Vector Functional Units, one per vault
- Bigger operations, with up to 256 bytes
- 36 Registers 256 Bytes wide
- Interleaves x86 with HMC instructions

Hipe: Hmc Instruction Predication Extension Applied on Database
Diego G. Tomé, Marco A. Z. Alves, and Eduardo C. Almeida.
DATE, 2018

DATABASE SYSTEMS PREDICATE PROCESSING

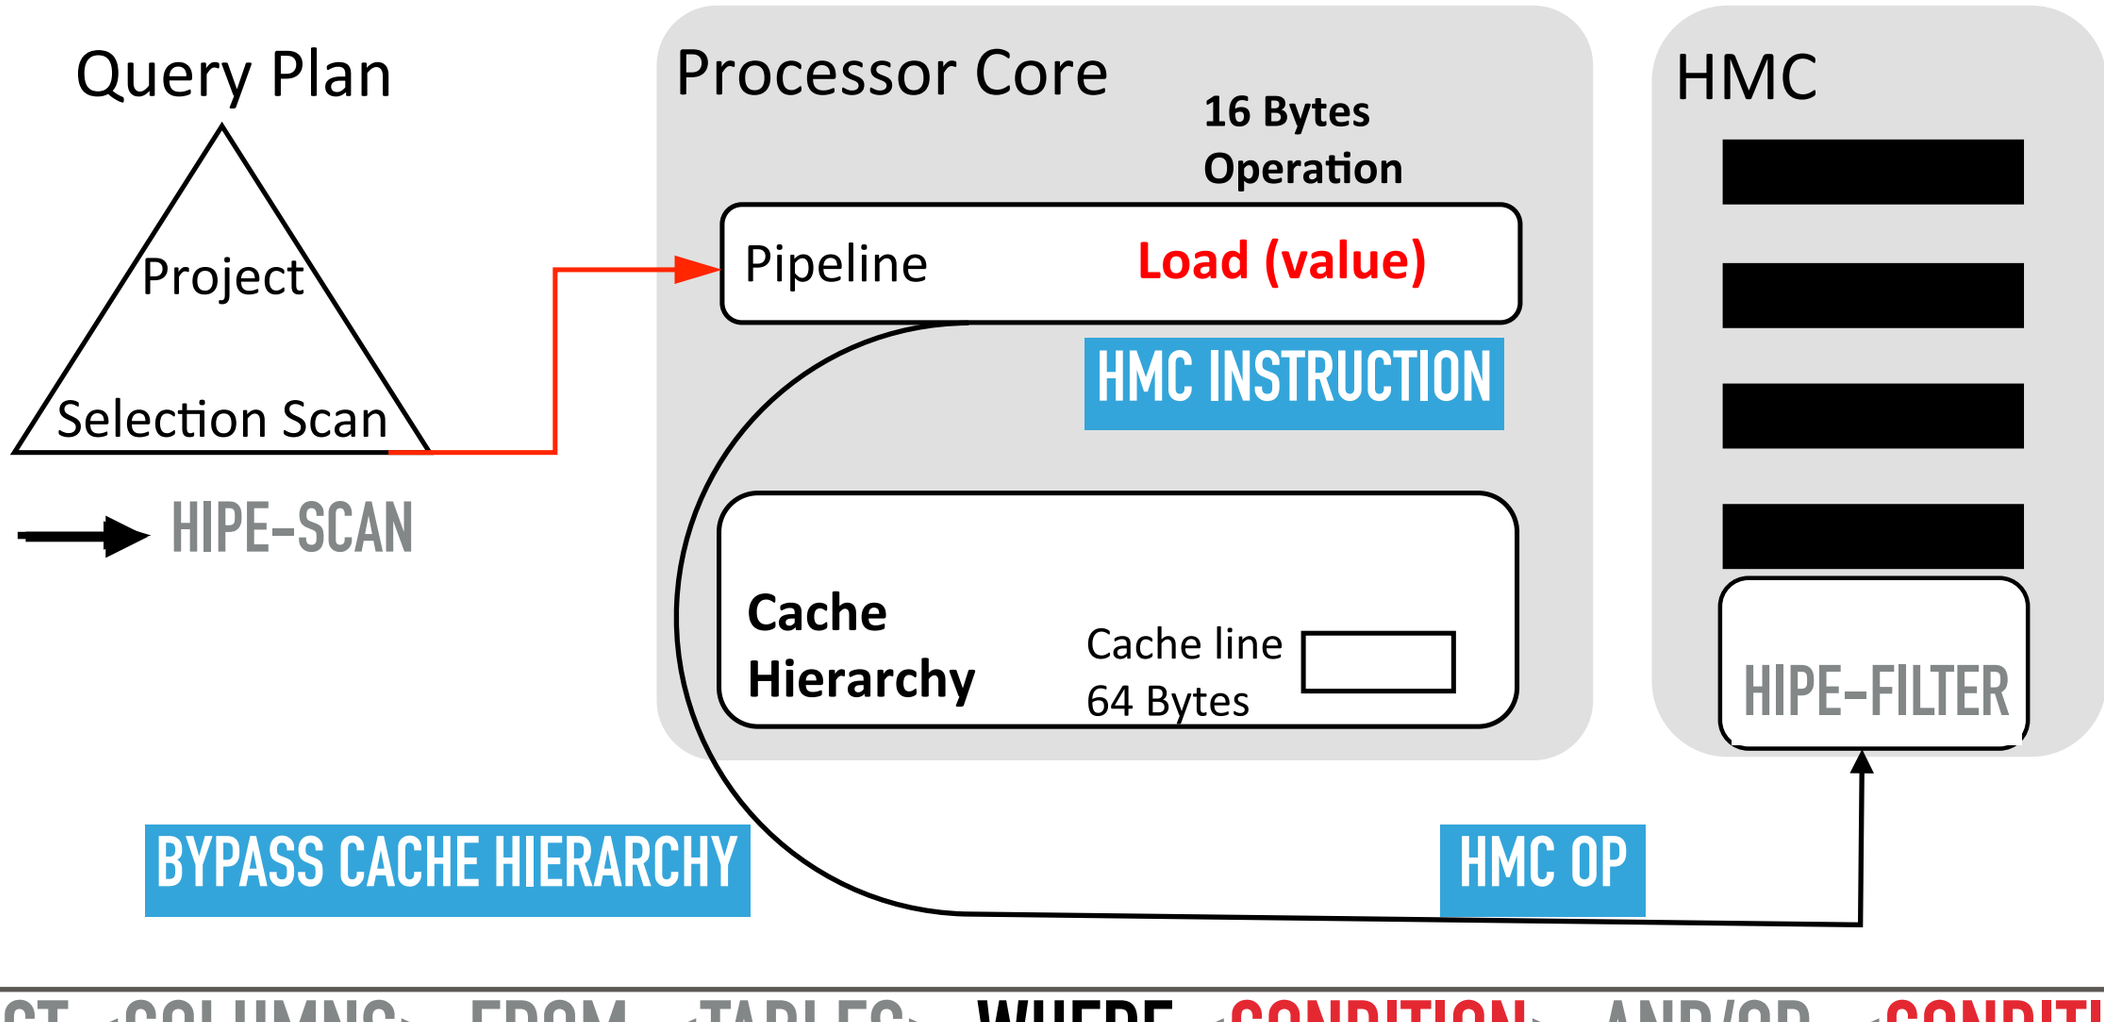

"SELECT<COLUMNS> FROM <TABLES> **WHERE**<CONDITION> AND/OR <CONDITION>"

DATABASE SYSTEMS PREDICATE PROCESSING

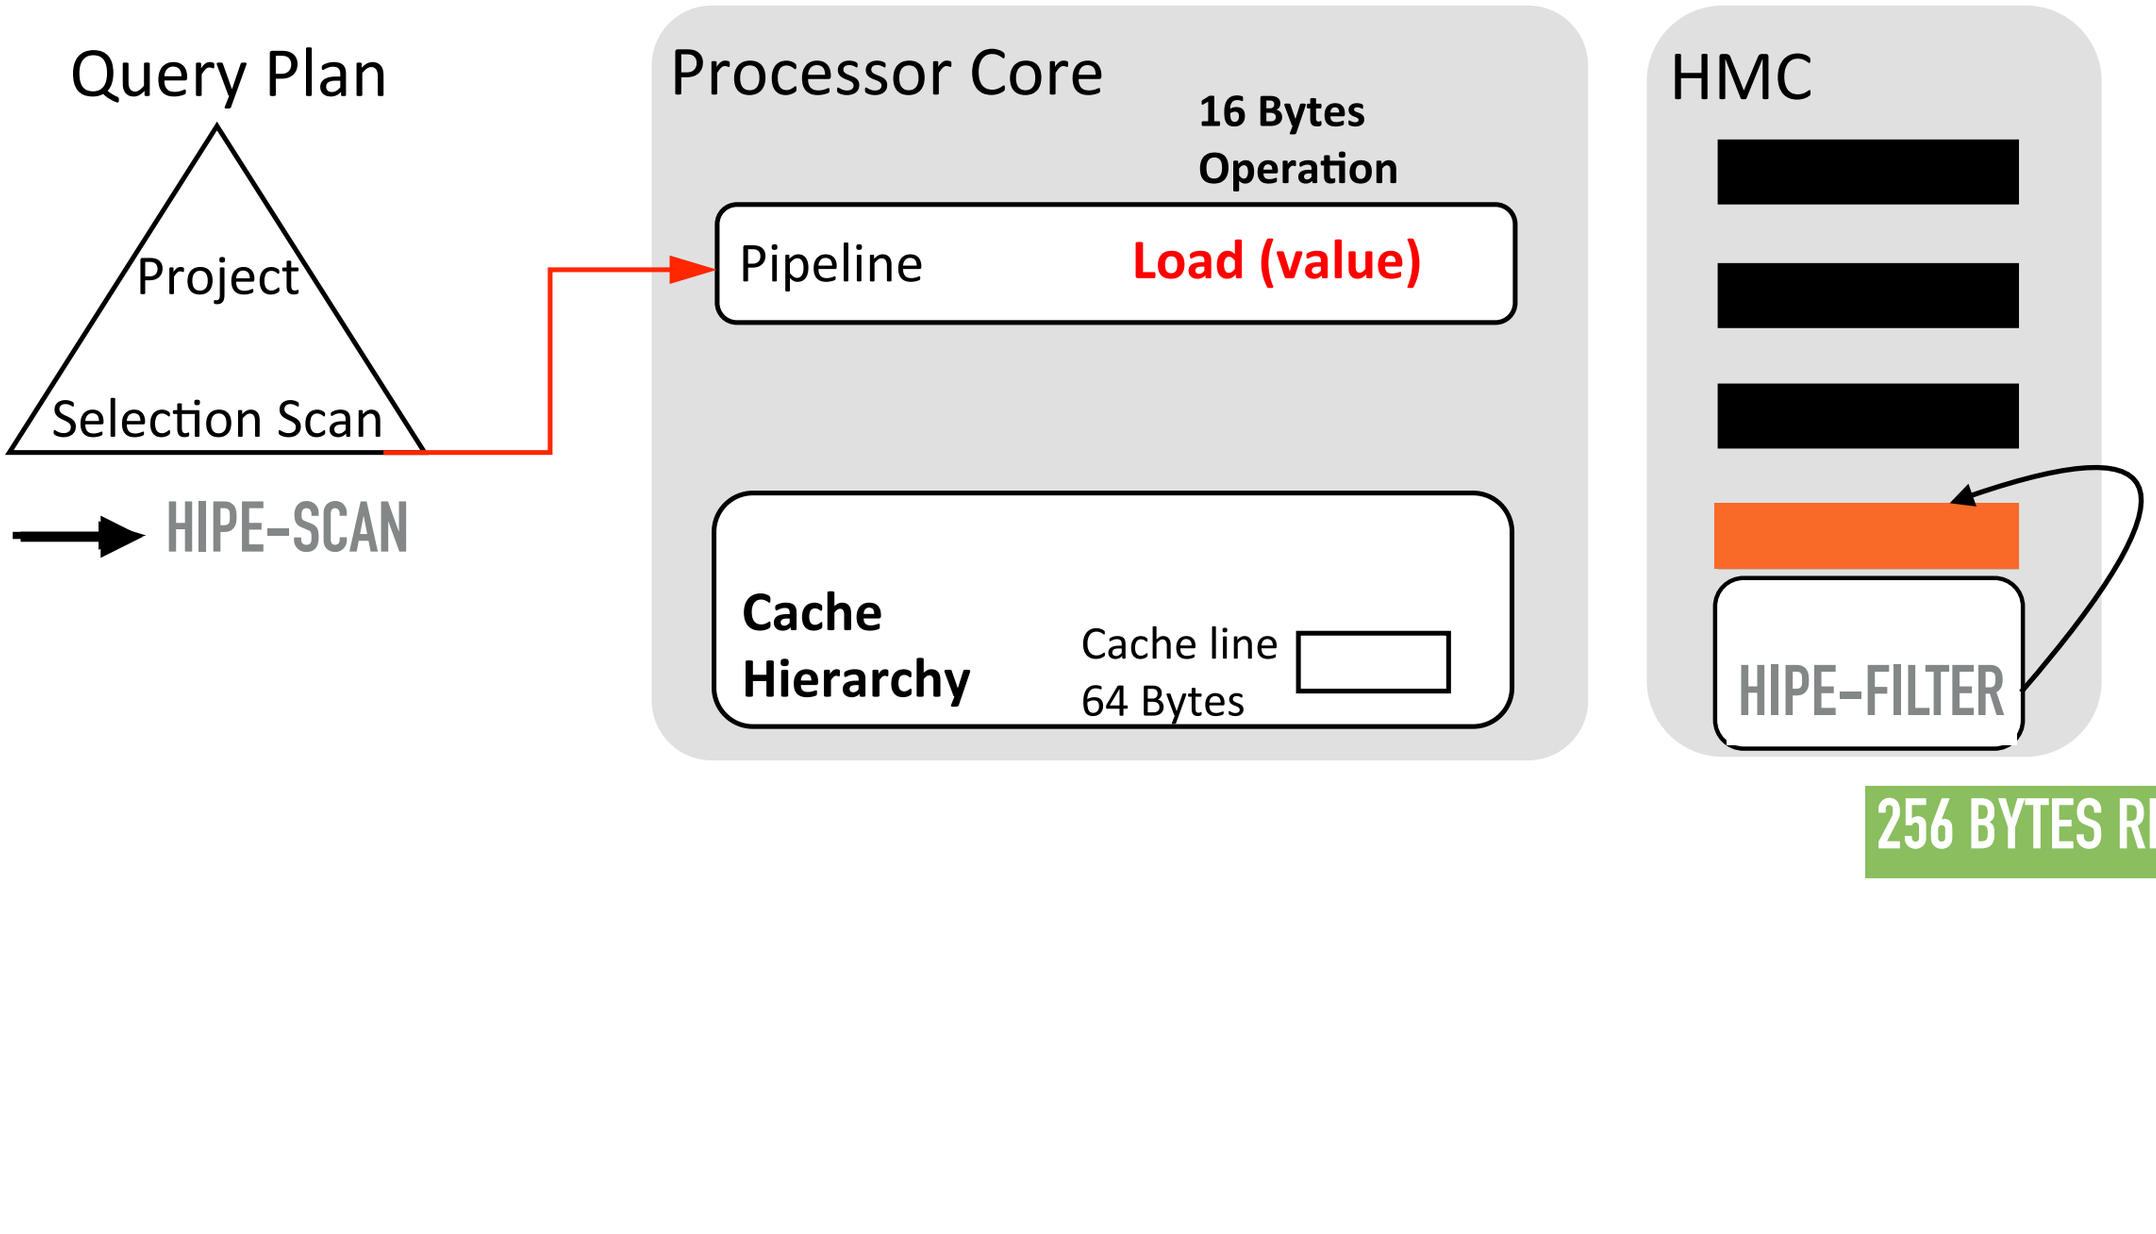

DATABASE SYSTEMS PREDICATE PROCESSING

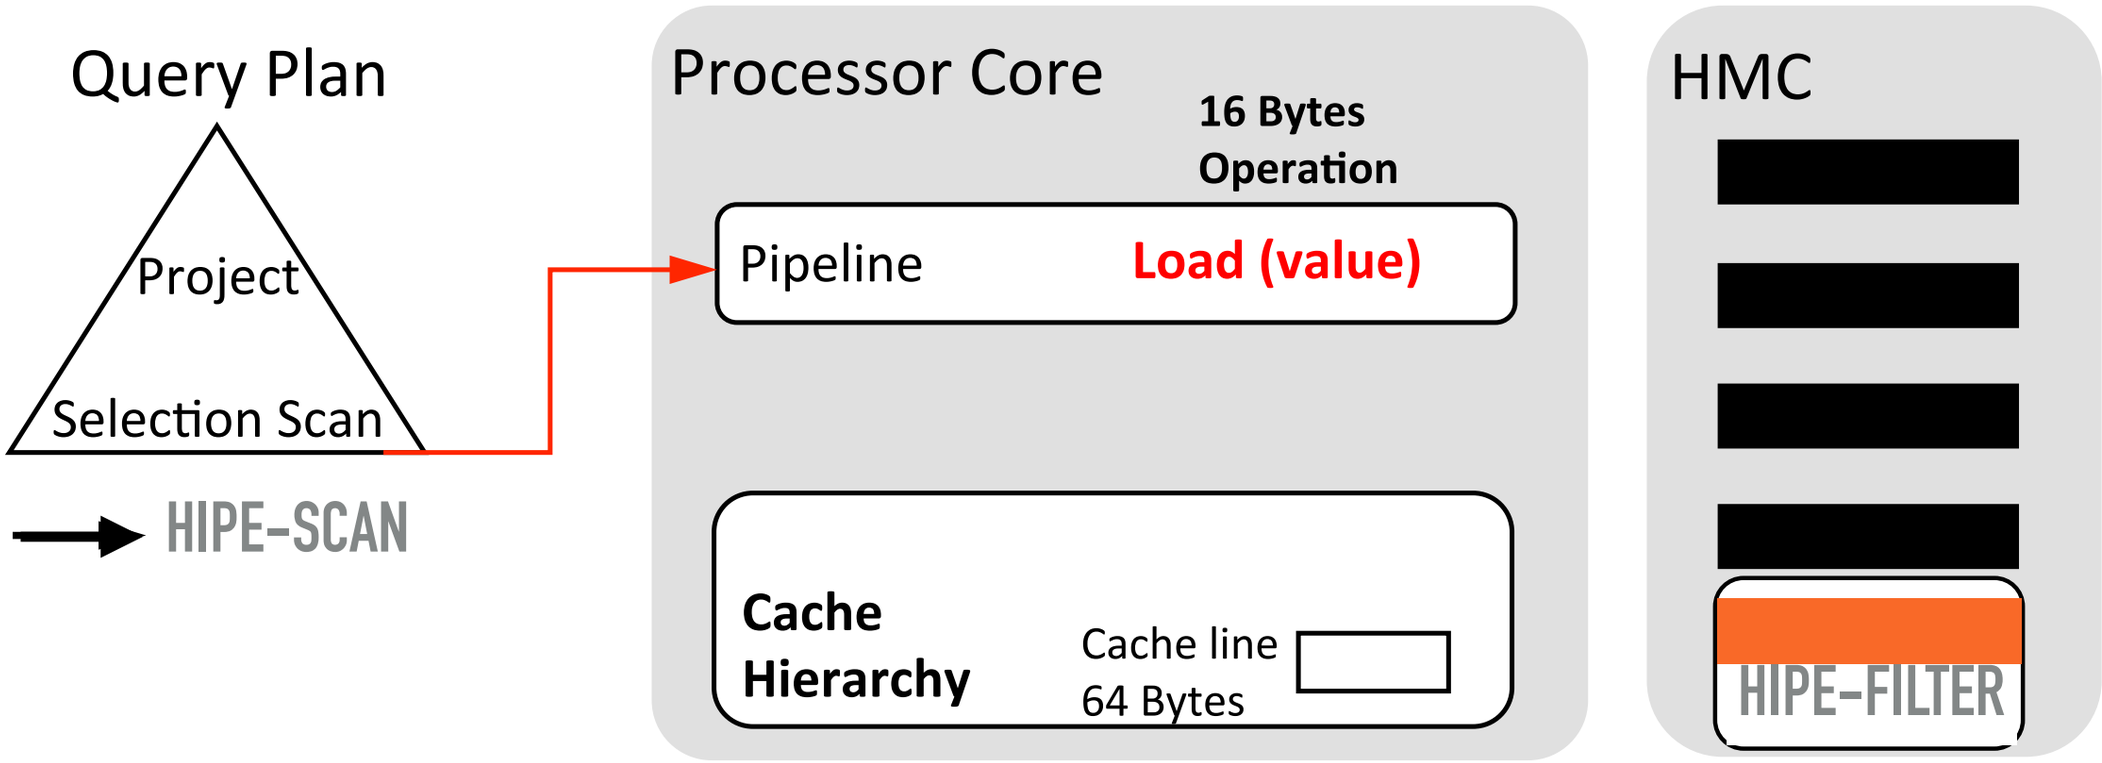

256 BYTES OPERATION

DATABASE SYSTEMS PREDICATE PROCESSING

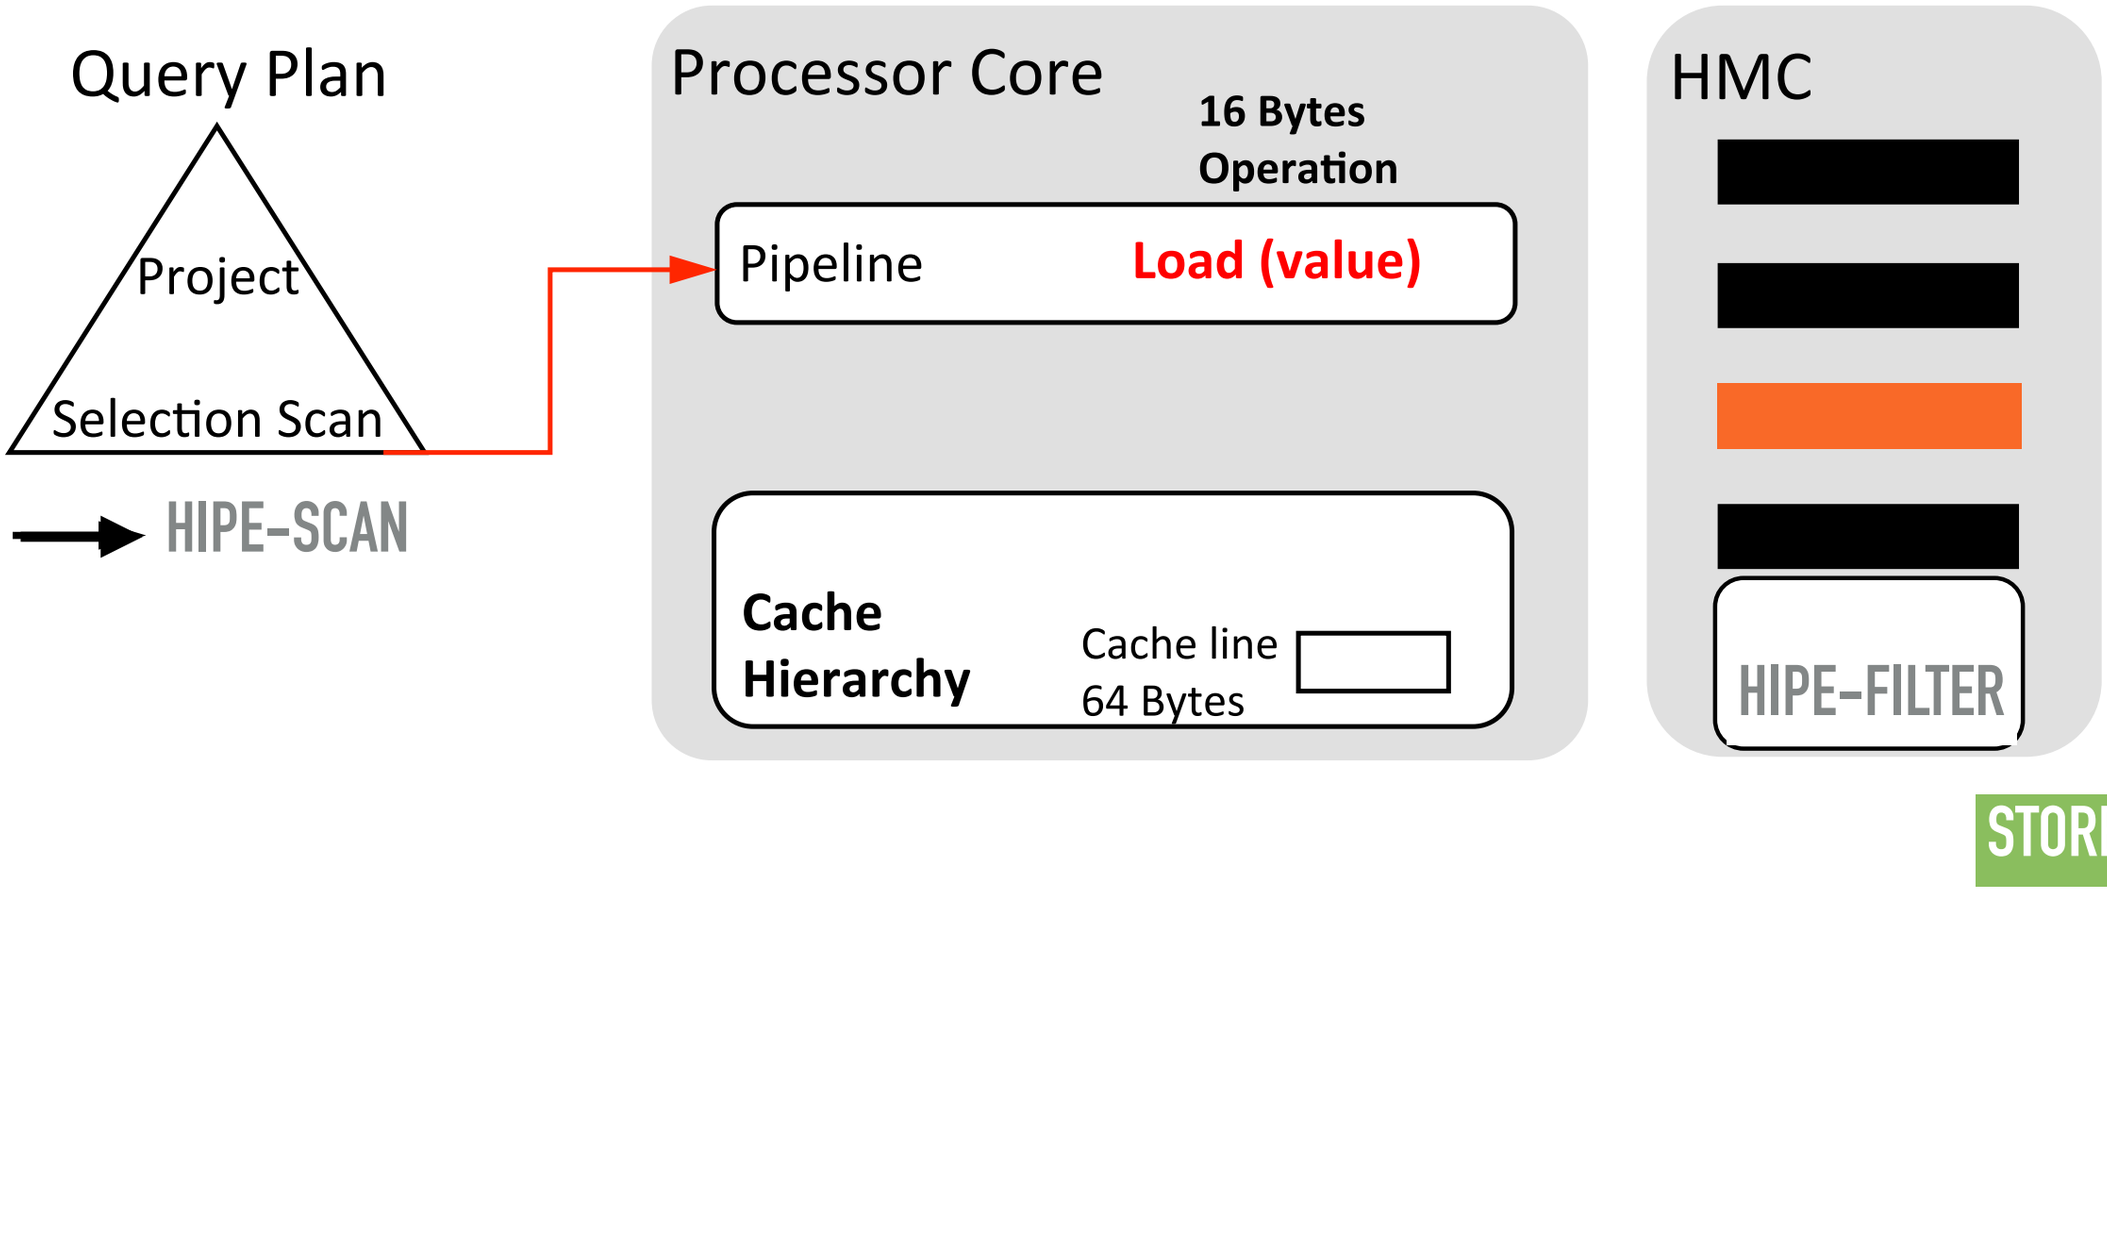

DATABASE SYSTEMS PREDICATE PROCESSING

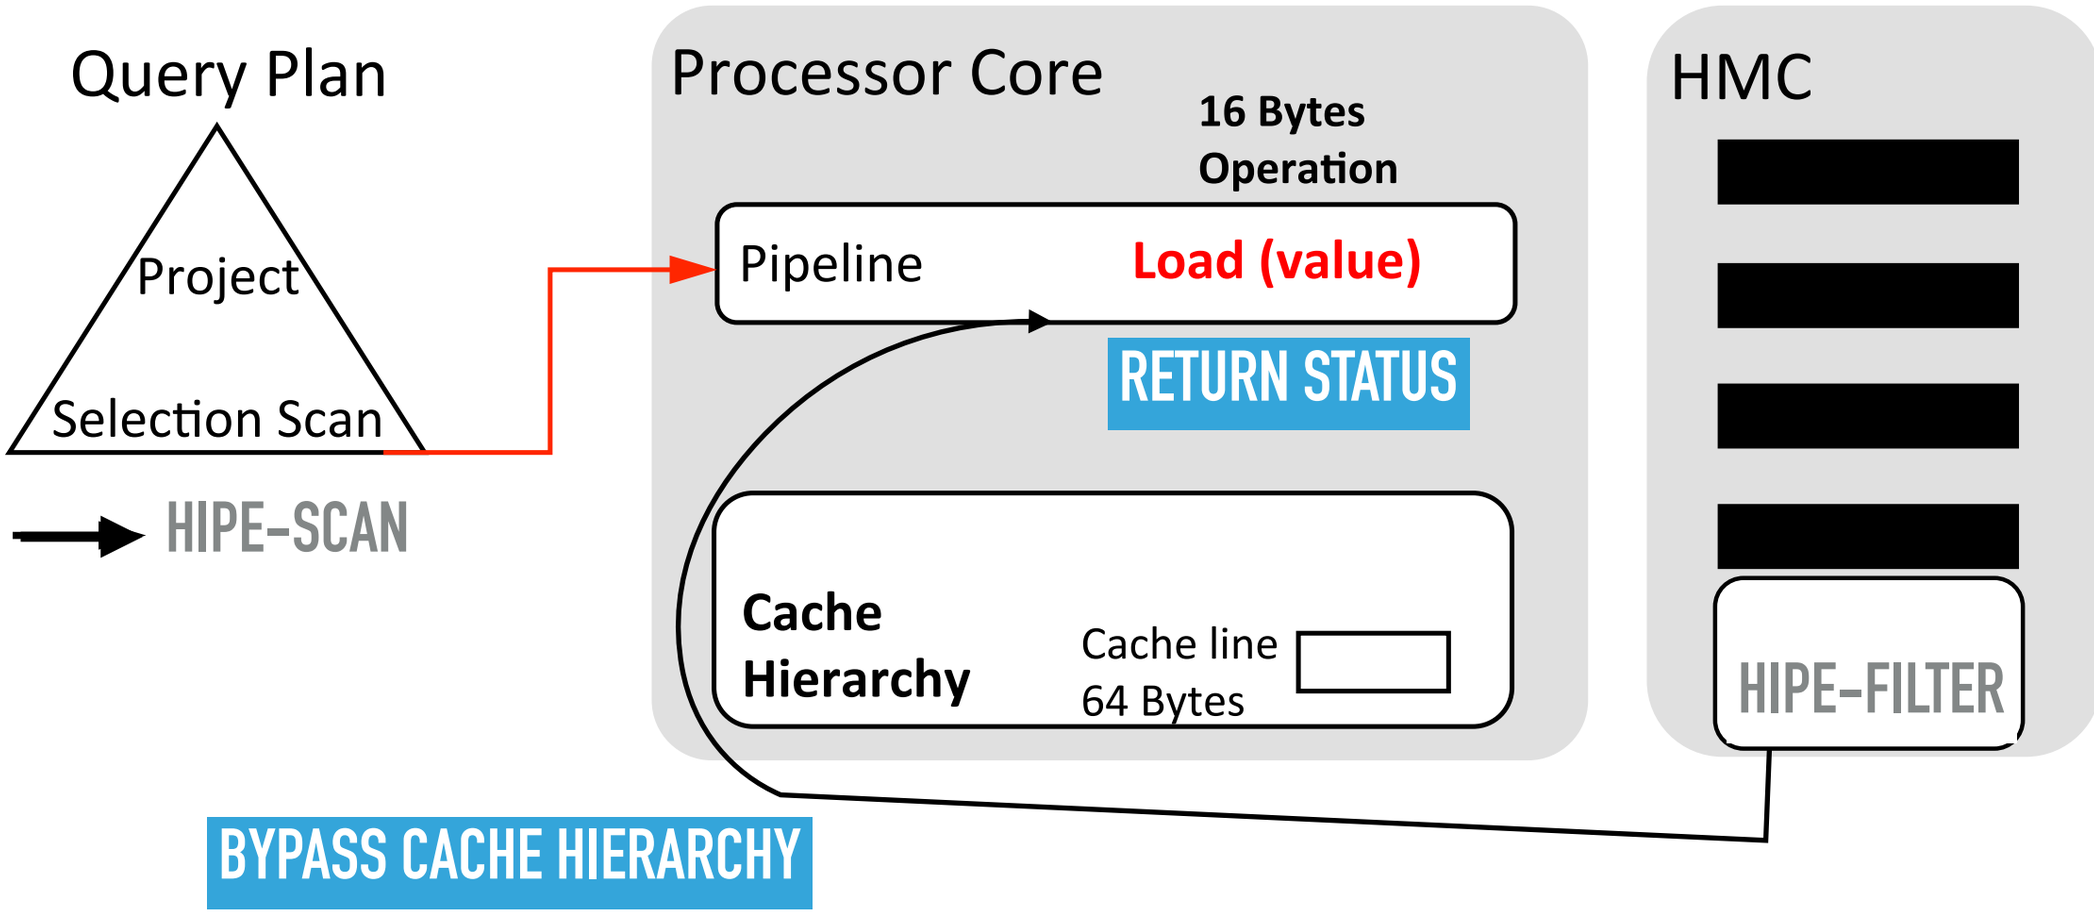

DATABASE SYSTEMS PREDICATE PROCESSING

DATA-MOVEMENT

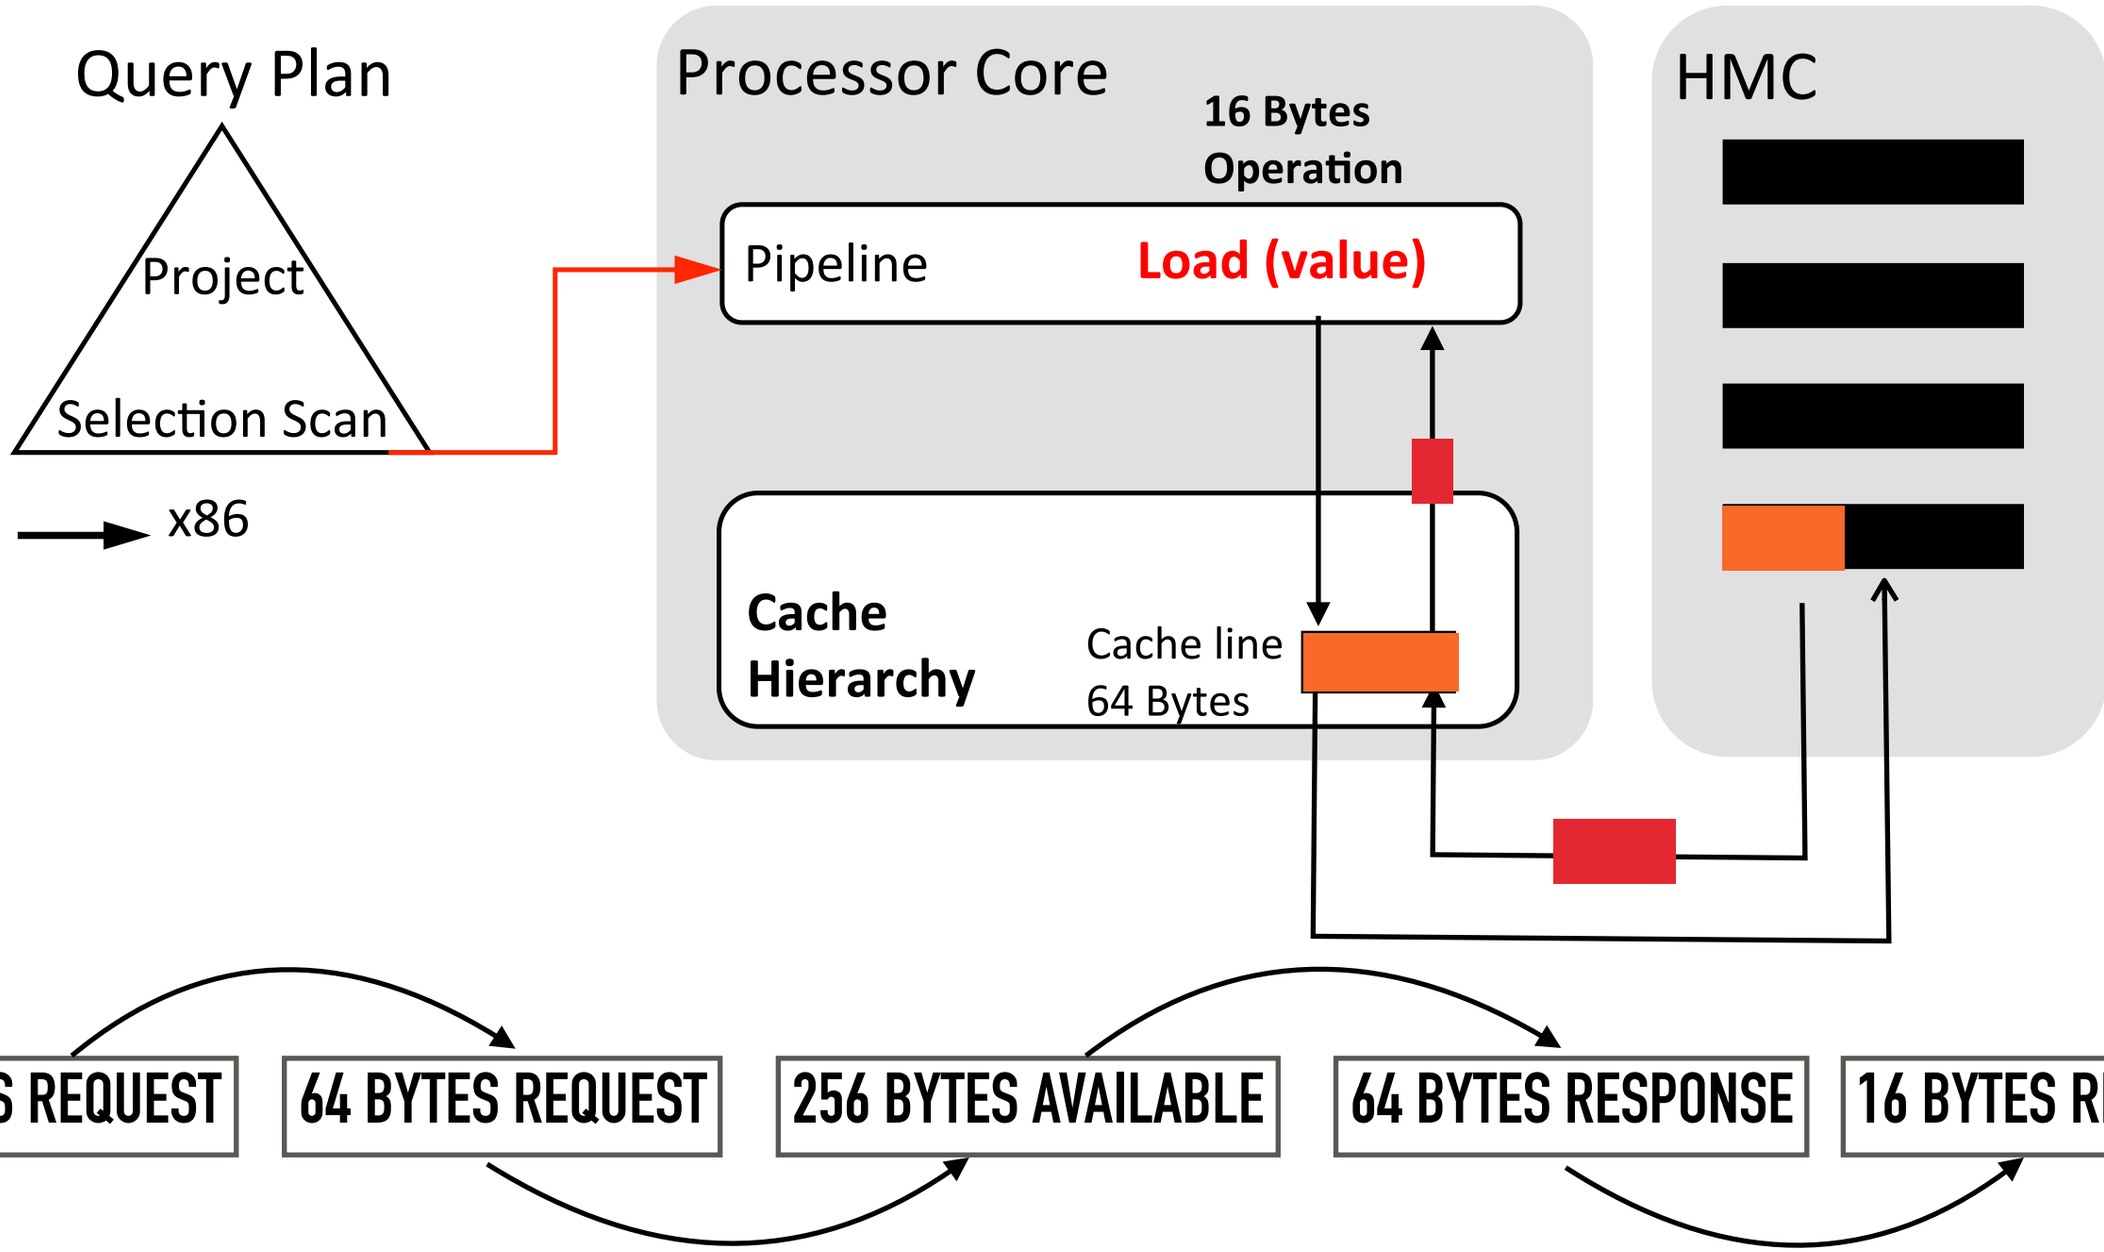

16 BYTES REQUEST → 64 BYTES REQUEST → 256 BYTES AVAILABLE → 64 BYTES RESPONSE → 16 BYTES RESPONSE

DATABASE SYSTEMS PREDICATE PROCESSING

NEAR-MEMORY OPERATIONS

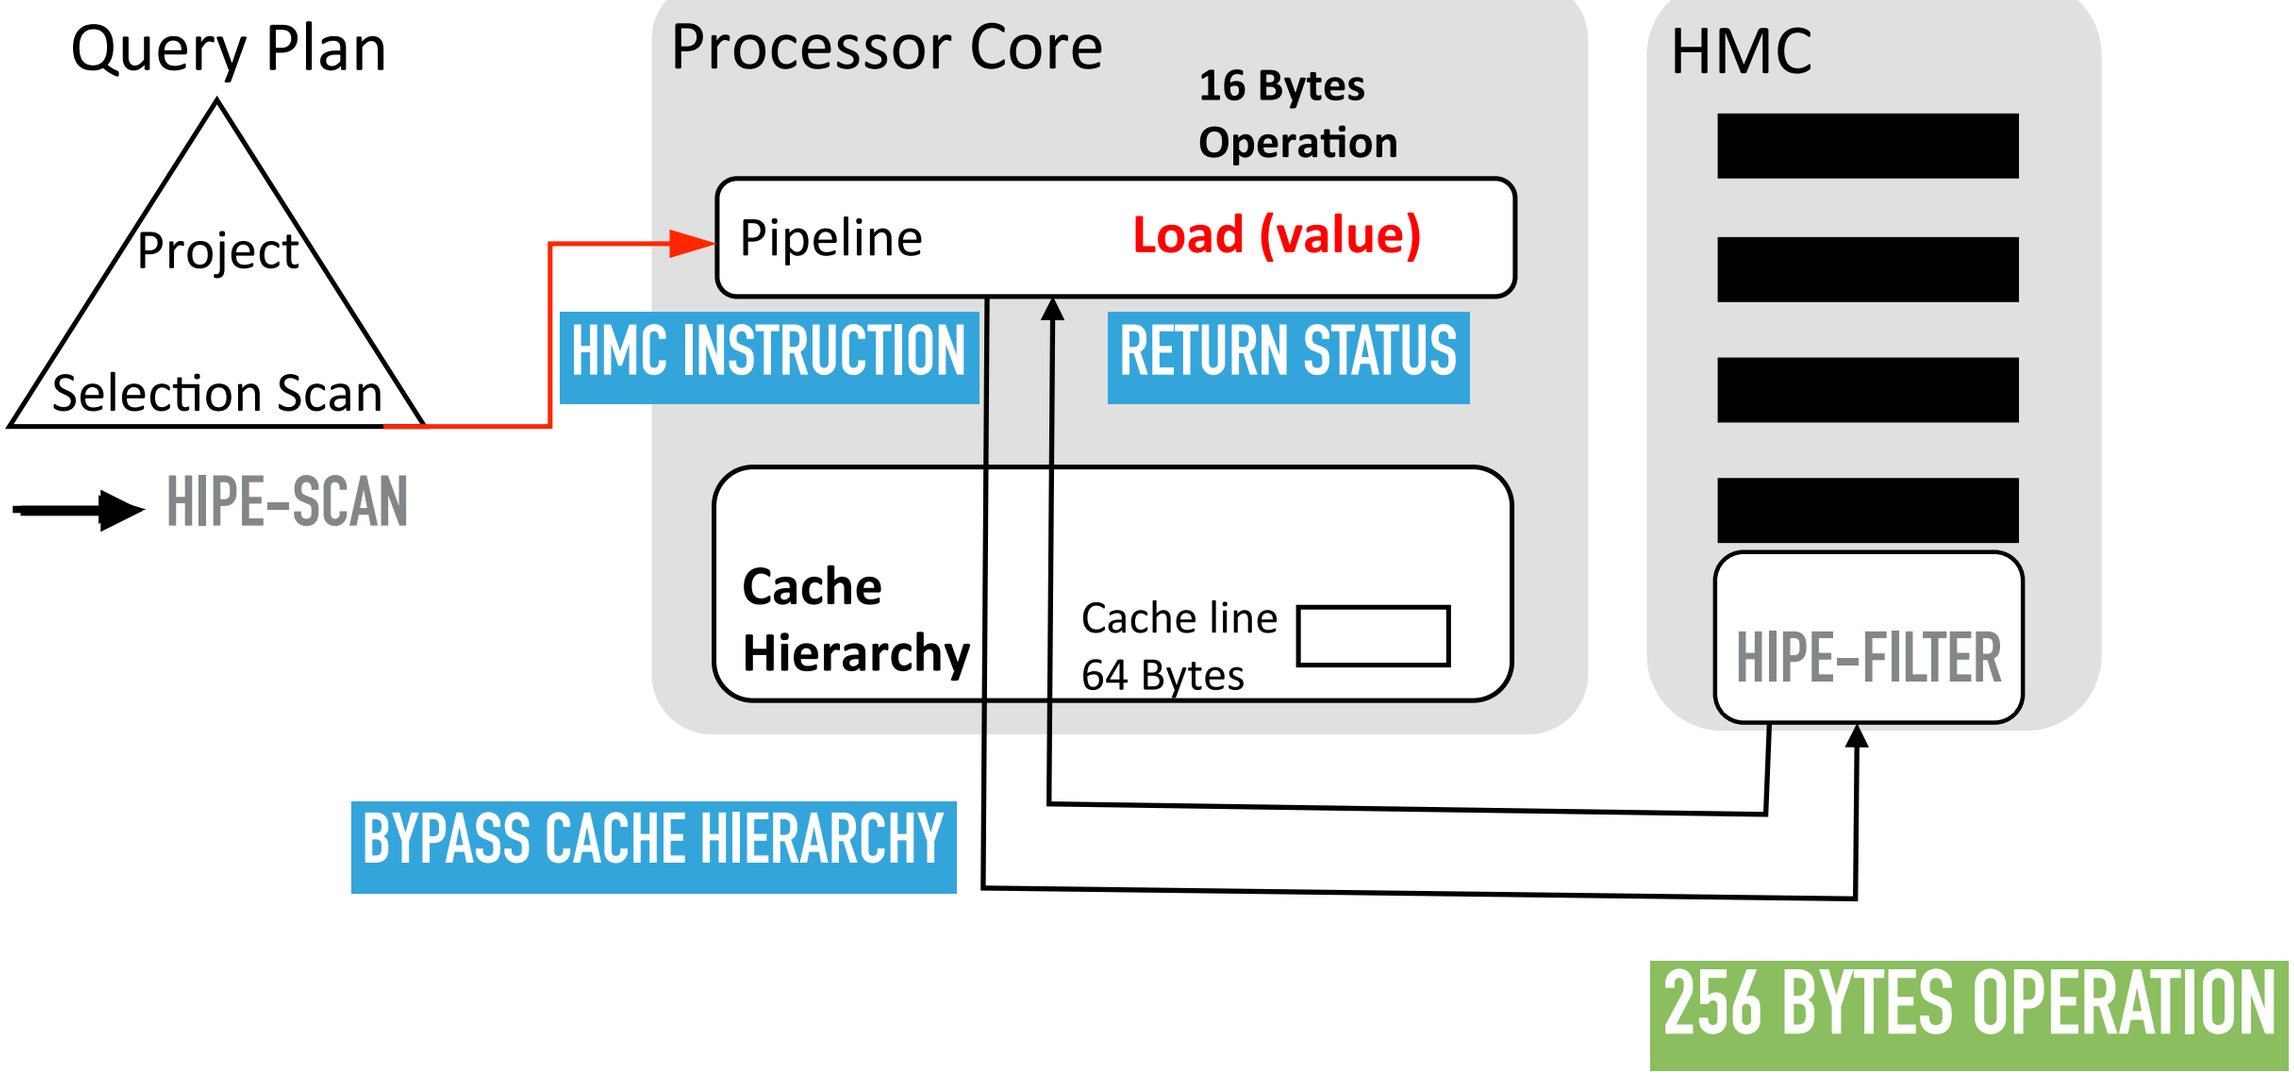

NEAR-DATA FILTERS BENEFITS

- Reduce cache pollution
- No Cache latency
- Mitigate DRAM operations (signals)

EXPERIMENTS

METHODOLOGY AND SETUP

- Using a cycle-accurate simulator to evaluate
- Dataset 1GB from TPC-H for OLAP scenario
- Micro-benchmark TPC-H Query 06

Sinuca: A validated micro-architecture simulator.
M. A. Z. Alves, C. Villavieja, M. Diener, F. B. Moreira, and P. O. A. Navaux.
IEEE 17th International Conference on High Performance Computing and Communications, 2015

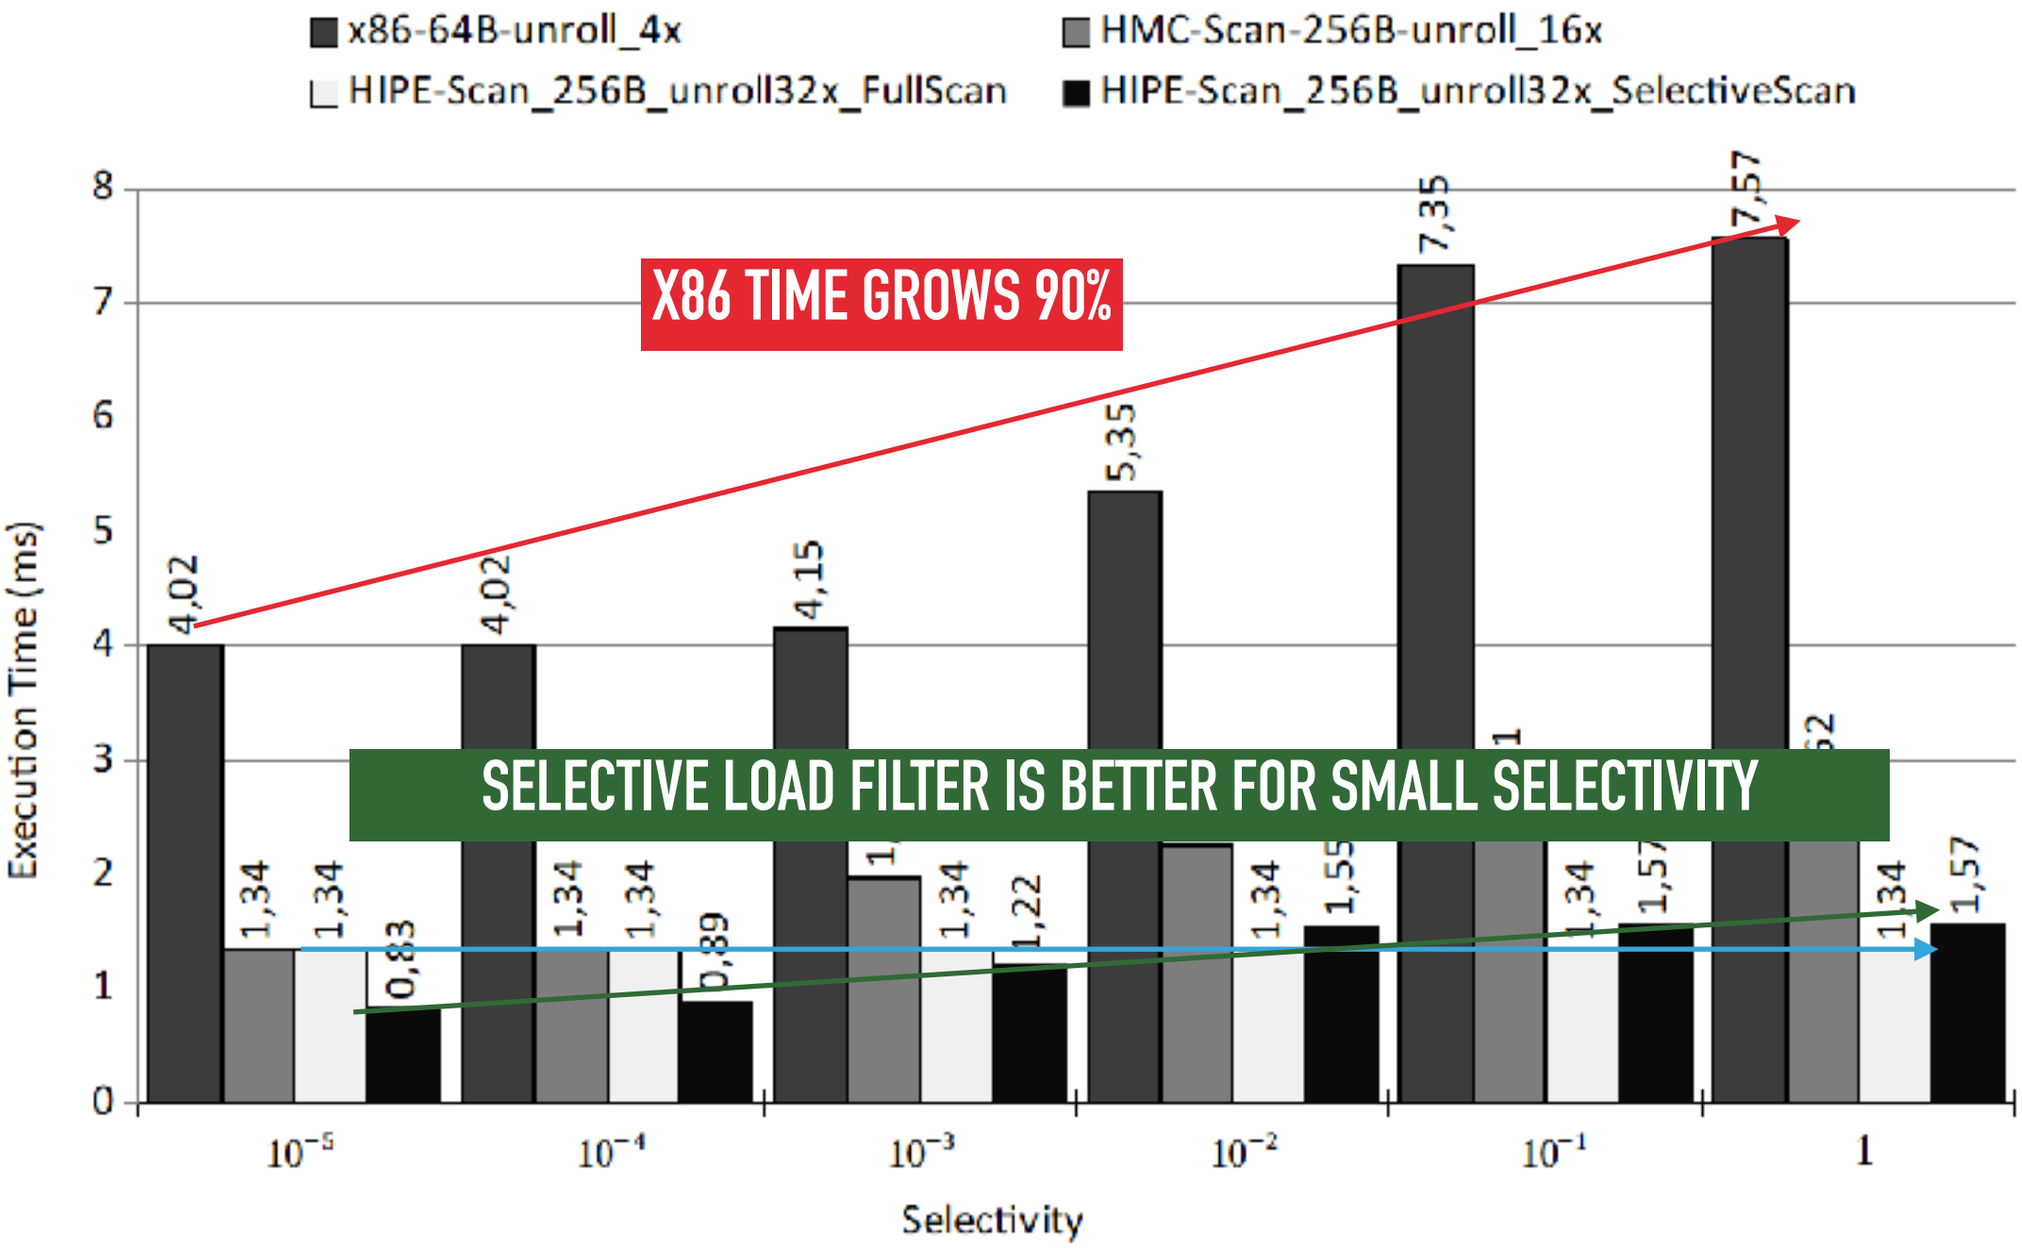

Figure 10: Evaluating execution time of TPC-H Q6 varying the selectivity factor in the different hardware architectures.

TPC-H query 06 Select scan - Performance

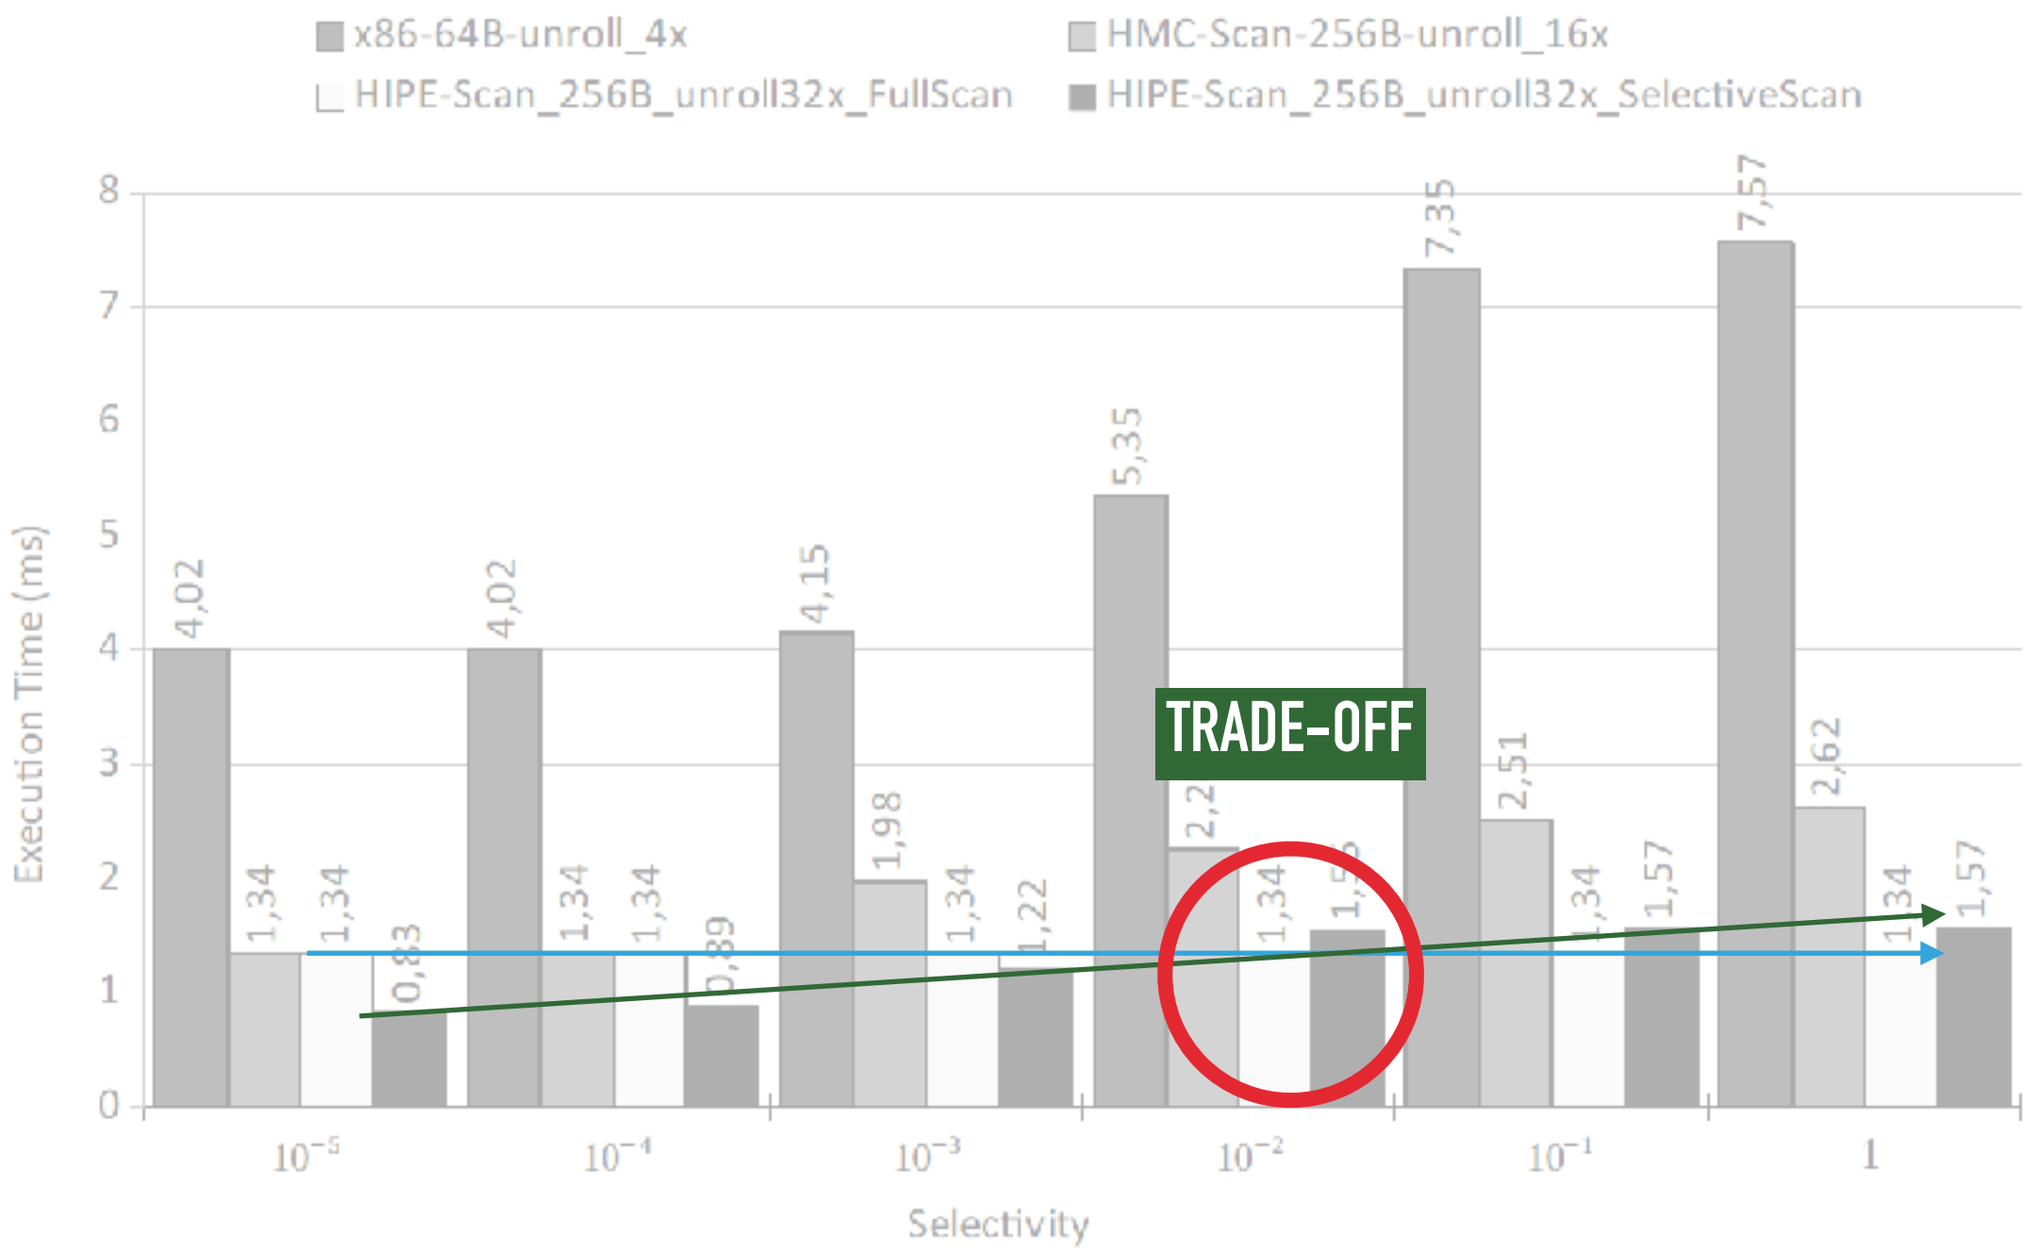

Figure 10: Evaluating execution time of TPC-H Q6 varying the selectivity factor in the different hardware architectures.

TPC-H query 06 Select scan - Performance

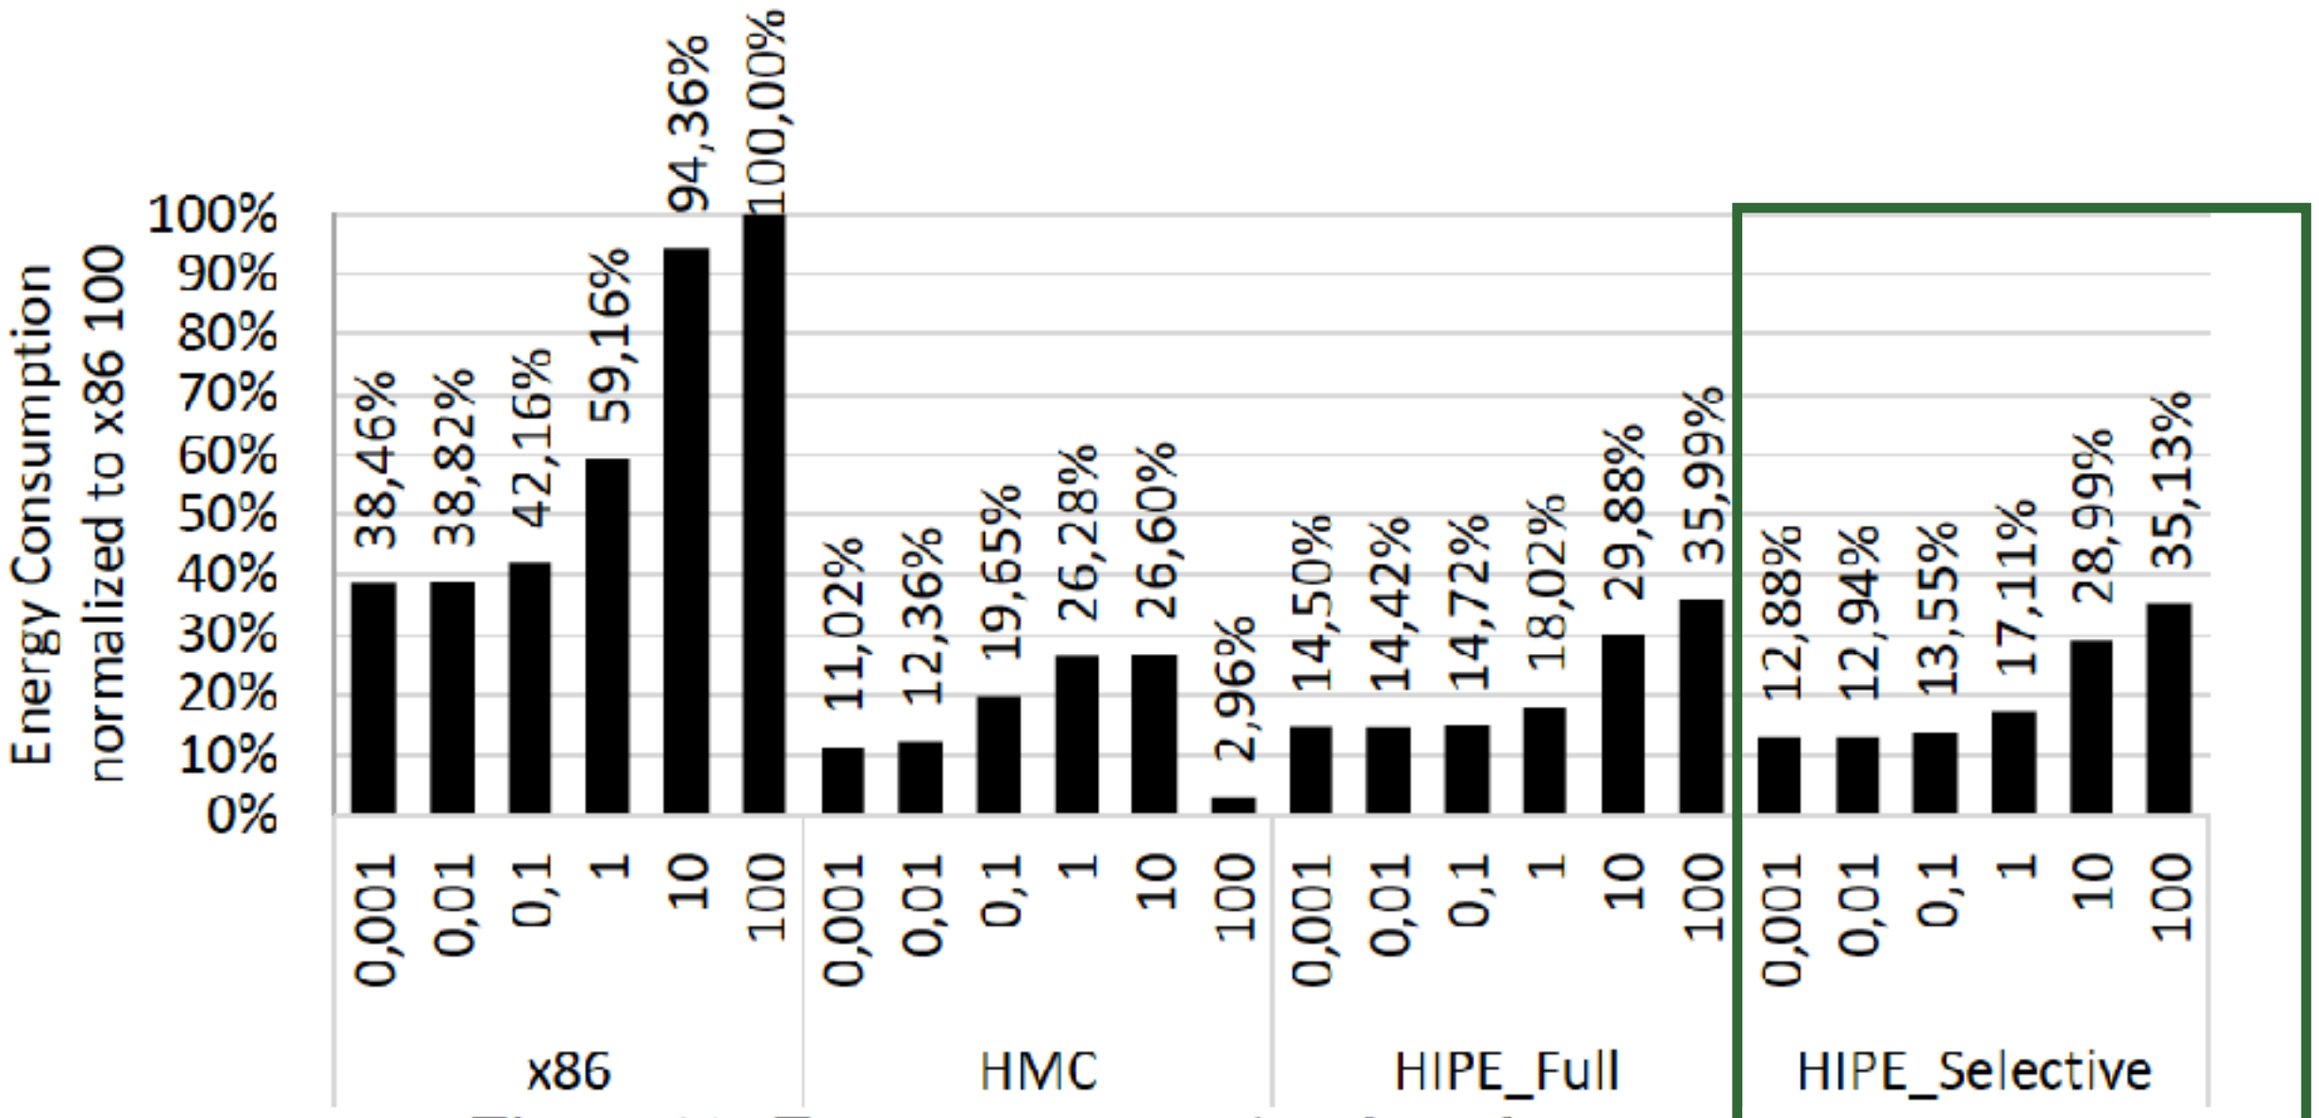

Figure 11: Energy consumption for select scan.

SELECTIVE LOAD FILTER IS 3X MORE EFFICIENT IN ENERGY CONSUMPTION

TPC-H query 06 Select scan - Energy

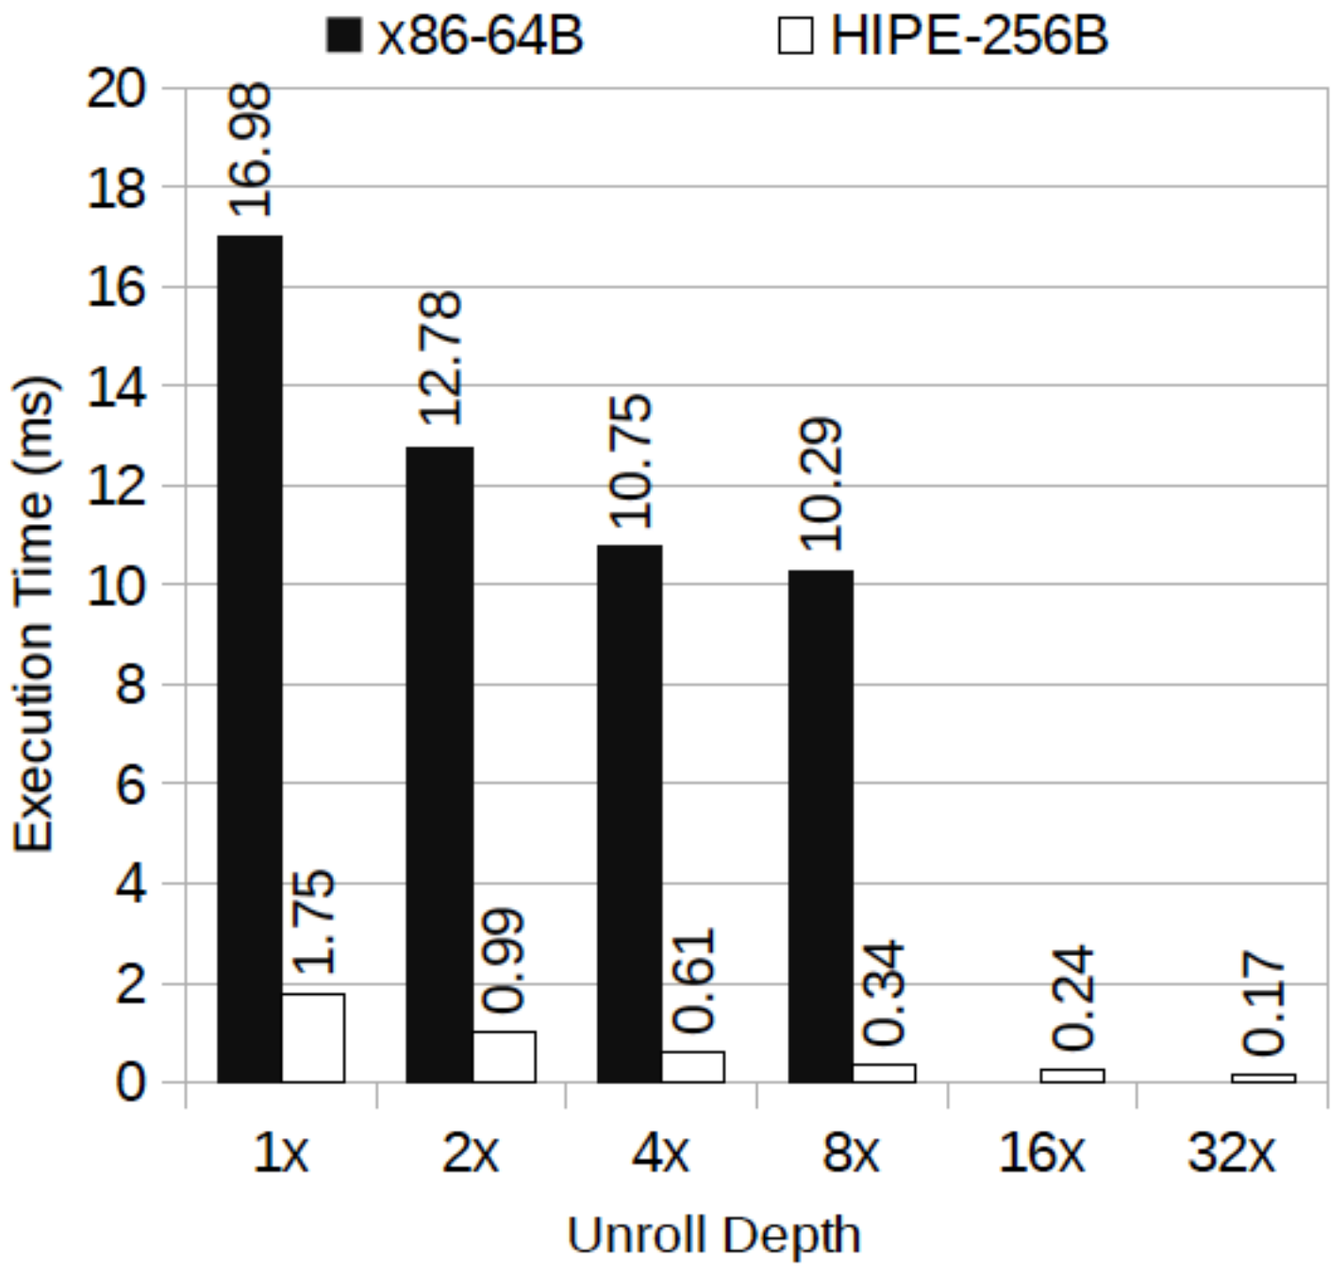

Figure 13: Execution time using the HIPE-Filters in the column projection varying the loop unroll depth.

TPC-H query 06 Projection

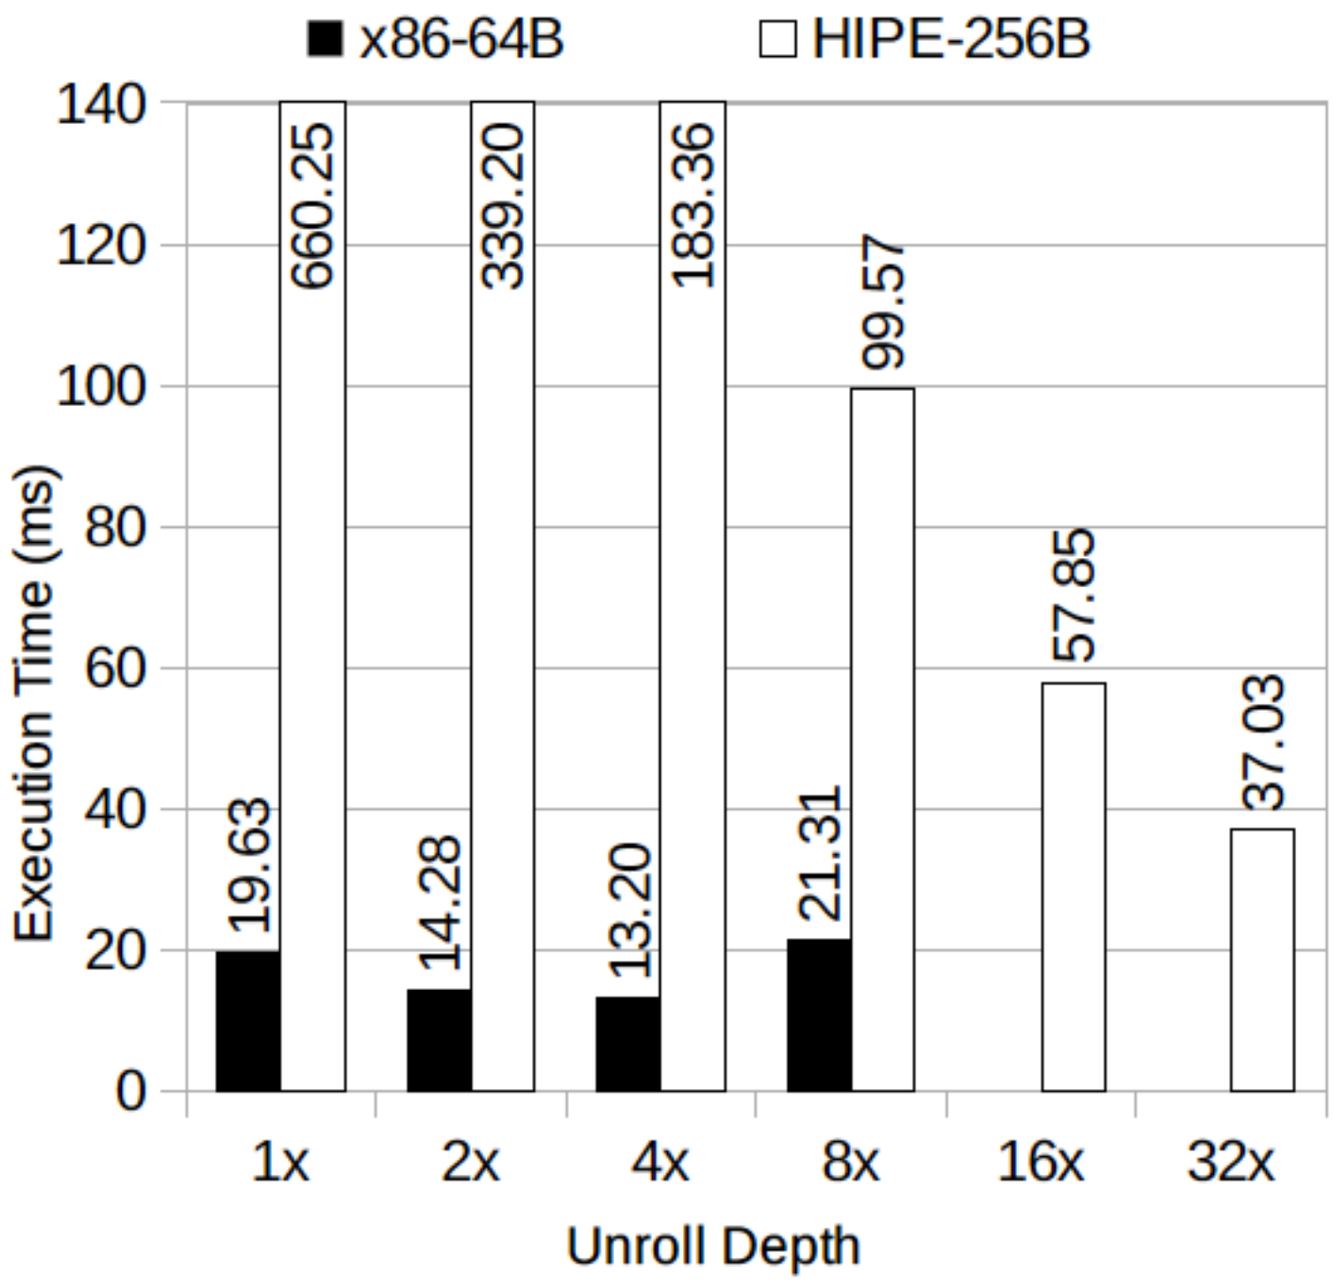

Figure 12: Execution time using the HIPE-Filters in the Nested-Loop join varying the loop unroll depth.

Nested-Loop Join

1. What happens when database systems run the select scan over the current x86 architecture using the HMC as ordinary DRAM? The memory wall is still a problem

2. Can we use the current HMC Instruction Set Architecture (ISA) to implement the near-data select scan? Yes, but it is not suitable to mitigate the memory wall.

3. What are the extensions to the HMC ISA to leverage the full potential of the hardware to mitigate the memory wall and reduce the interleaving with x86 instructions? Branch-less near data filters

Research Questions

CONCLUSION

CONCLUSION

- With the usage of the HMC as ordinary DRAM the memory wall is still a problem
- Performing filtering operations inside the HMC improves performance for OLAP
- Near-data processing is suitable for OLAP mitigating the memory wall for DBMSs
- The correct filter algorithm on different selectivity factors may have great impact on performance.