An Application-Specific Instruction-Set Processor for Database Operators

TensilicaDay 2017

Sebastian Haas
Emil Matúš
Gerhard Fettweis

16.02.2017

Outline

- Development of database accelerator core based on a Tensilica processor
- Deploying instruction set extensions (TIE) for database operators used in query processing
- Analysis of query optimization techniques
- Chip tape-out of database accelerator core
- Evaluation: performance and power measurements of database operators and query optimization

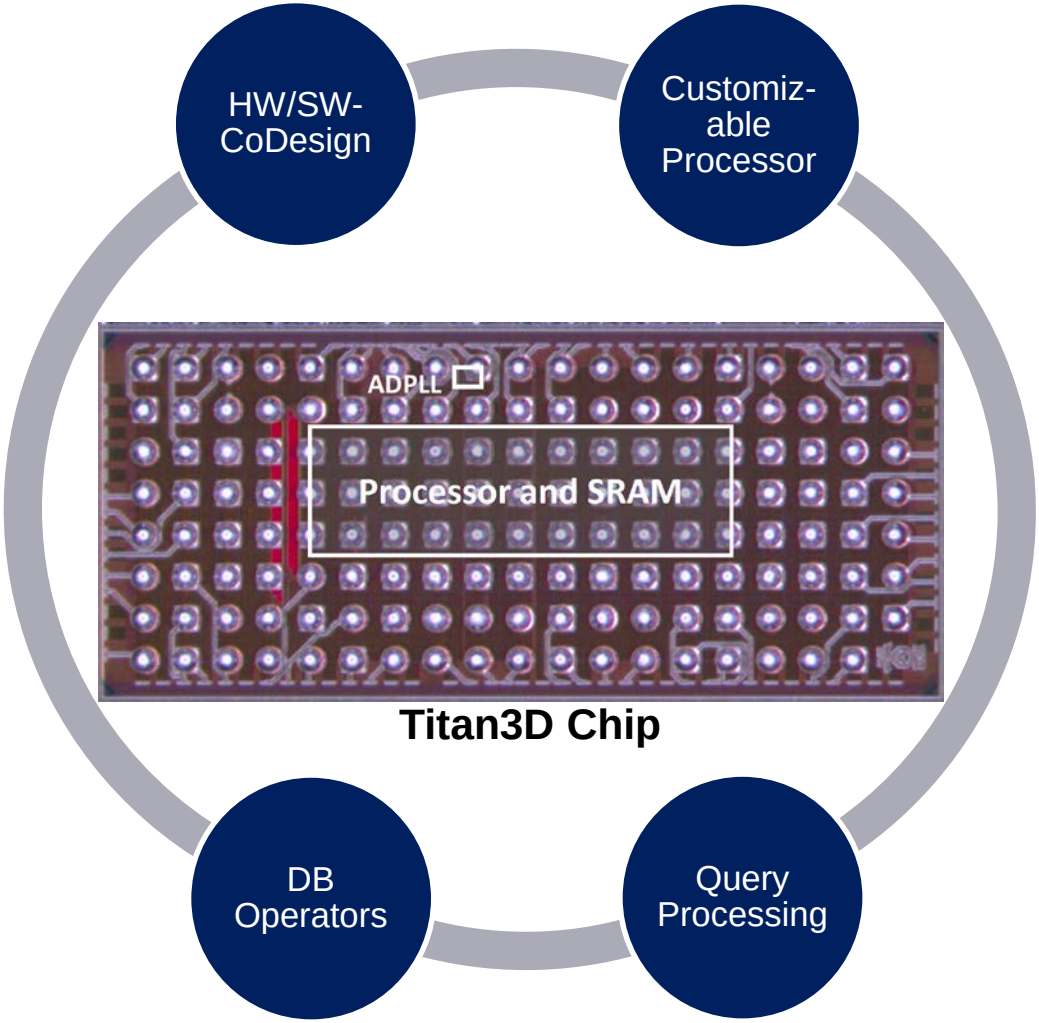

INTRODUCTION

Introduction

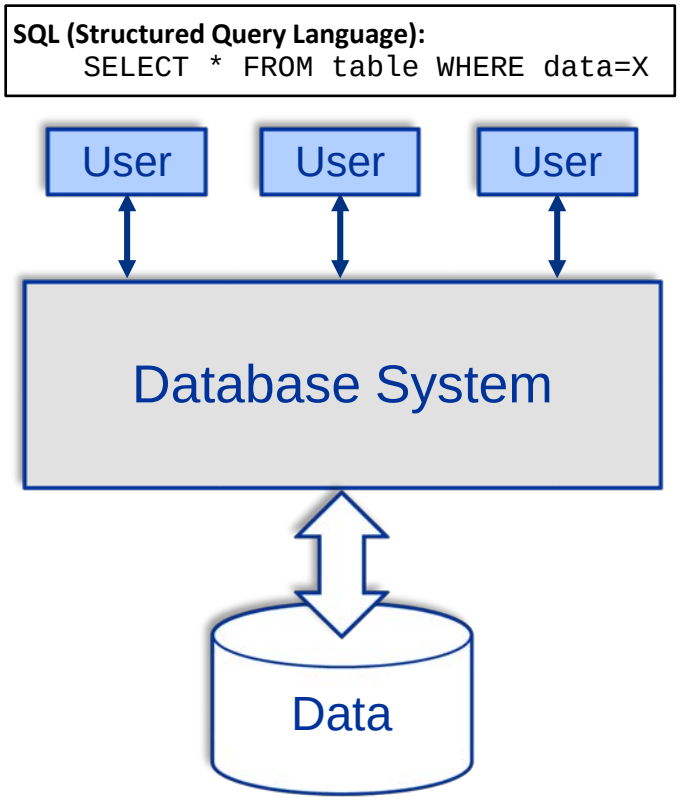

Query Processing

- Demand for
  - Low query latency (fast answer)
  - High throughput
  - Big data
  - Advanced analytics
- Growing demand, data, work, and complexity
- No sufficient improvement in energy efficiency of general-purpose processors

➔ **Custom-made processor**

DATABASE ACCELERATOR

Processing Element (PE)

Tensilica Xtensa LX5 RISC

- Instruction set extensions (TIE) for database operators
- 2 load/store units
- 2x 128-bit SIMD
- 64-bit VLIW

Core Wrapper

- Clock generation (ADPLL): 83-2000 MHz
- Power management controller (PMC): 0.8-1.1 V
- 96 kB single-port SRAM
- JTAG

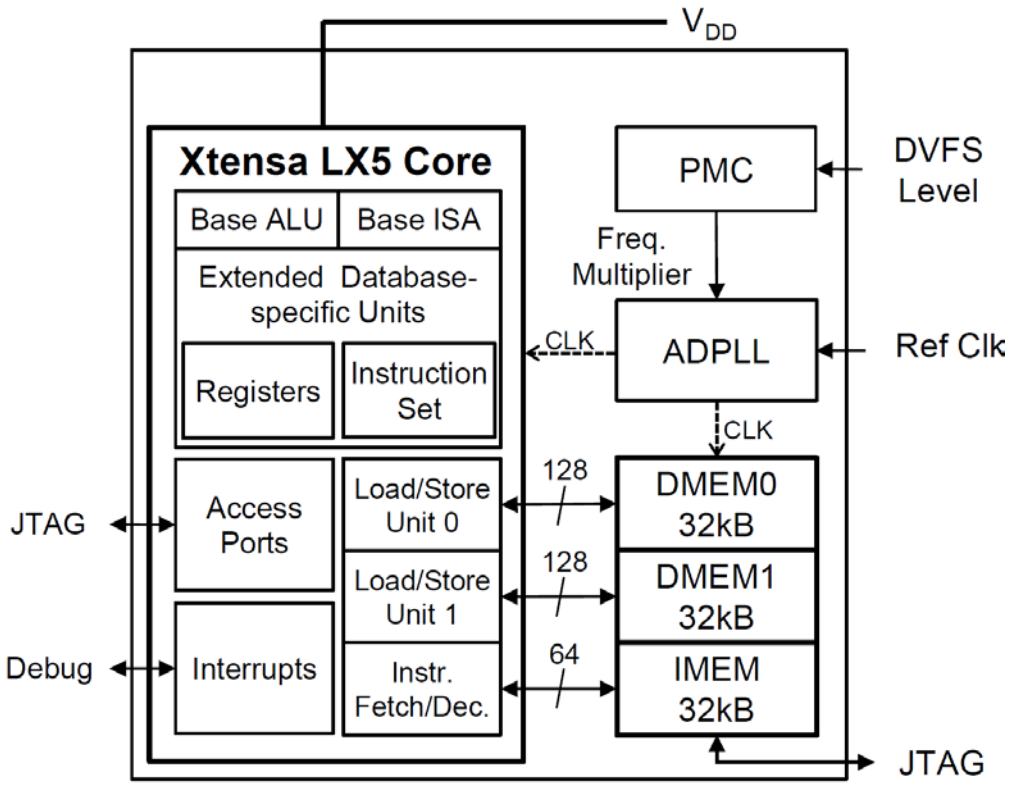

Titan3D Chip

Technology	28 nm SLP CMOS Globalfoundries
Size (pre-shrink)	Chip: 3.3x1.5 mm², PE: 1.6x0.46 mm²
Gate count PE	1.11 M NAND2 gate equivalents
PE area Core Memory	 total: 0.286 mm², basic core (without ISE): 0.09 mm² 0.232 mm² (96 kB)
Max. throughput of arithmetic-logic operations	50 GOPS
Max. clock frequency	250 MHz at V_{DD} = 1.1 V
Avg. power consumption	23.3 mW at 250 MHz, V_{DD} = 1.1 V

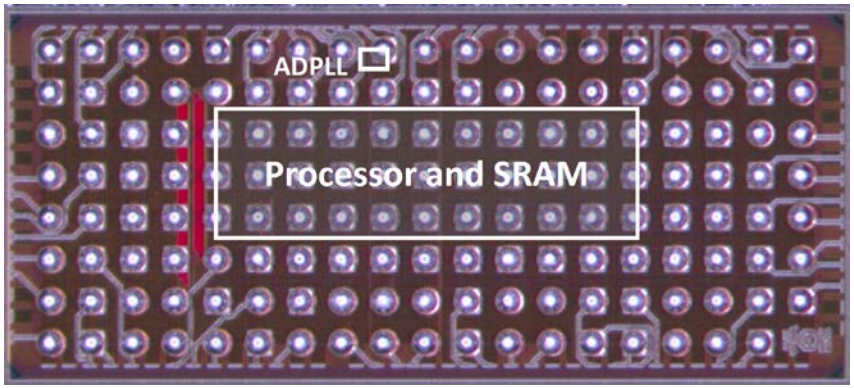

Database Operators

Application-specific units containing load, store, and processing instructions

- Units support SQL queries
  - `INTERSECT`
  - `ORDER BY`
  - `JOIN`
  - `GROUP BY`
- Data types:
  - Sort and intersection of RID sets: 32-bit integers
  - Join and aggregation of relational tables with multiple attributes: 64-bit key-payload pairs (integer tuples)

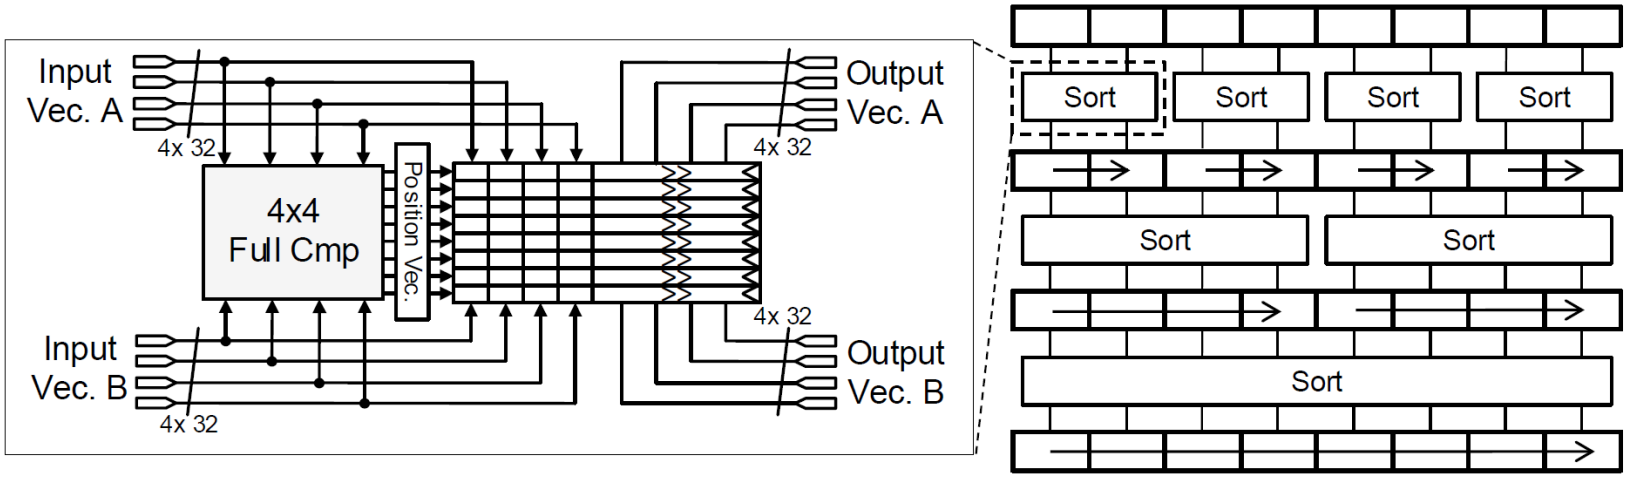

Merge Sort algorithm by using the dedicated hardware unit for a 4x4 merge

EXPERIMENTS

Measurement Set-Up

- Set-up
  - Chip module
  - Power supply board
  - Host-PC

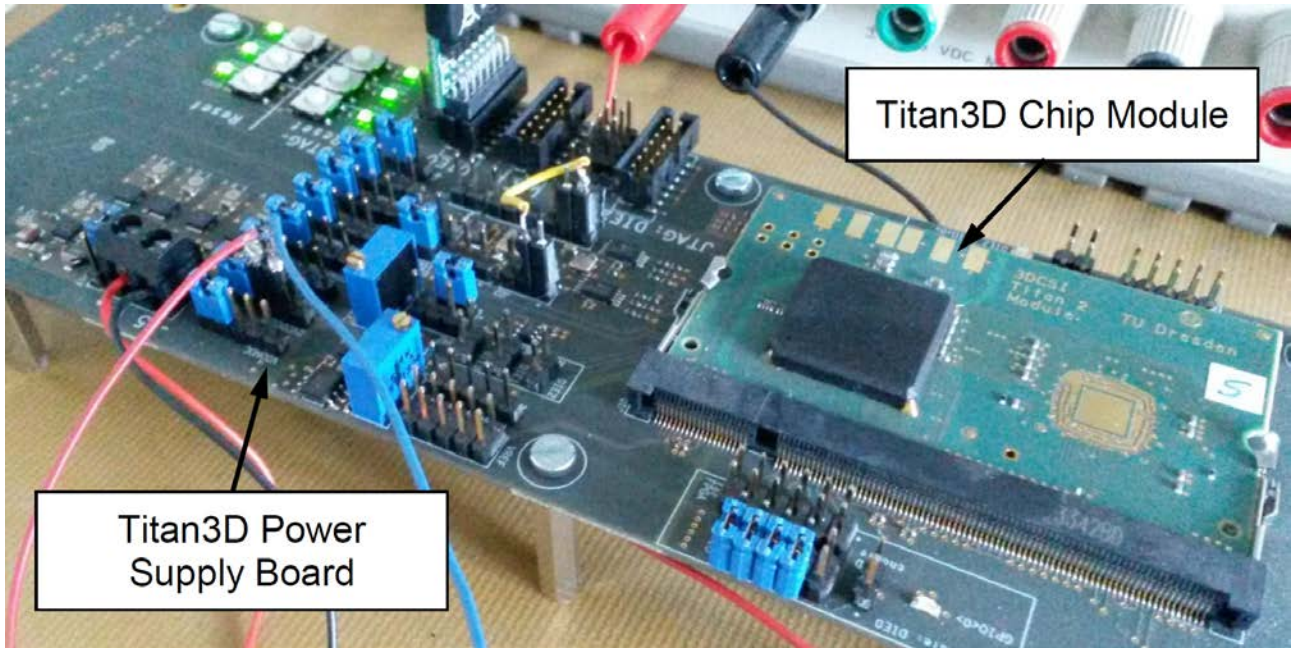

DB Operators: Throughput and Power

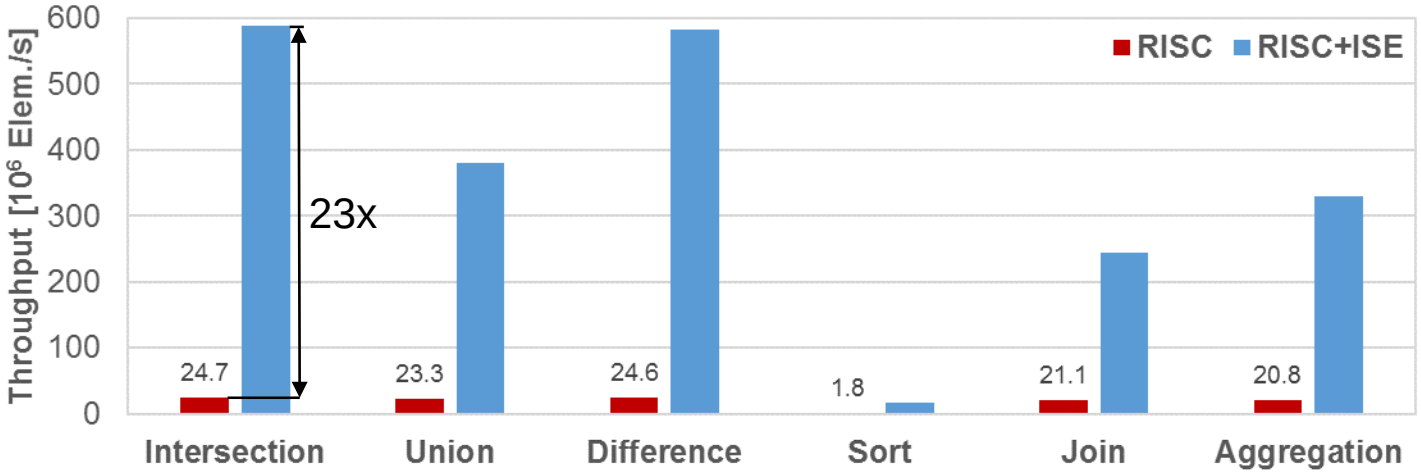

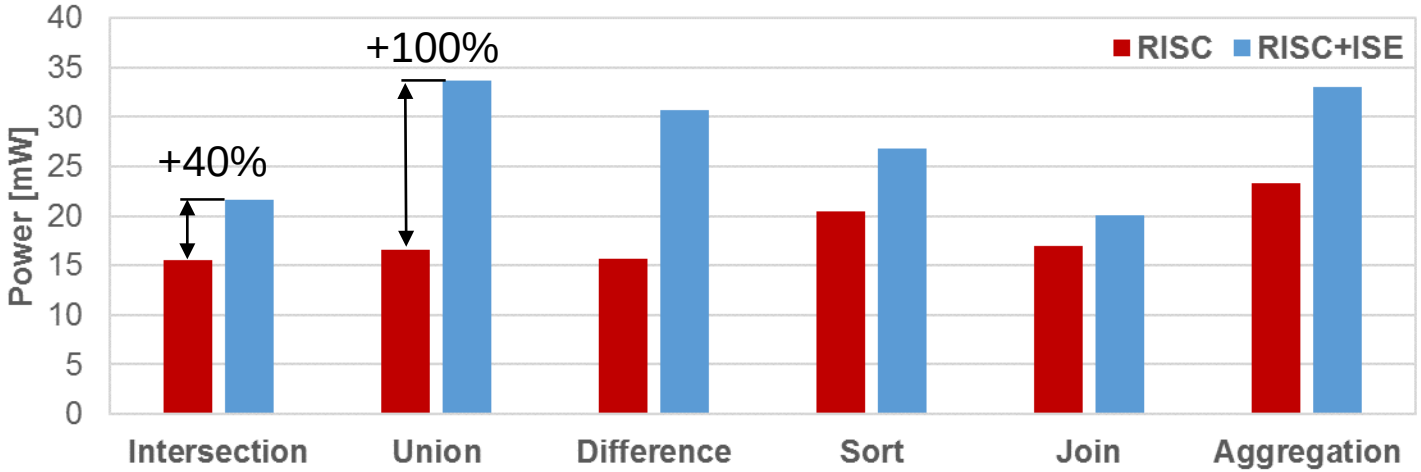

DB Operators: Energy

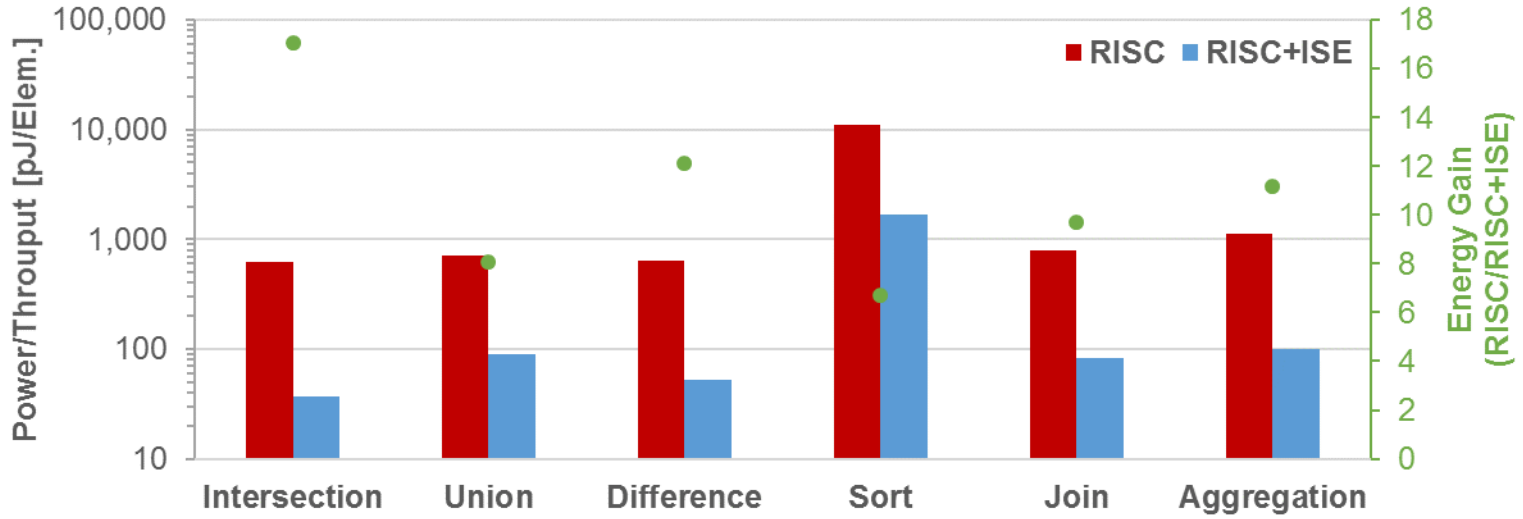

In contrast to pure RISC execution, instruction set extensions lead to

- up to 23x performance improvement
- 55% increase of power consumption
- 8x to 17x energy gain

Query Optimization: Example

- Given **3 tables** with different row count each containing **2 columns** of 32-bit integer values

table_A

id	data

table_B

id	data

table_C

id	data

- **Query:** Intersection of all tables

```
SELECT table_A.id, table_A.data
FROM table_A
INTERSECT
SELECT table_B.id, table_B.data
FROM table_B
INTERSECT
SELECT table_C.id, table_C.data
FROM table_C
```

Query order?

Query Optimization: Example

- Optimized execution by first intersecting the two tables with lowest selectivity
  → Second intersection operates on minimum number of tuples
- Run accelerated 32-bit intersection on six single columns
  → Identify selectivities between tables
  → Choose optimal query order

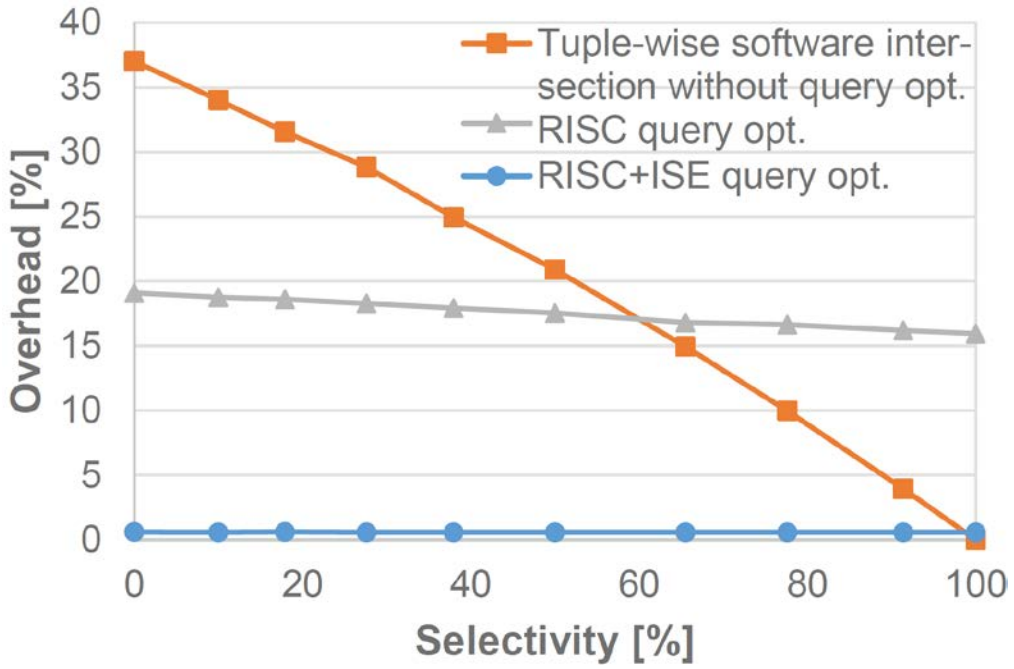

Energy Efficiency

Query optimization: Power measurements at different supply voltages

- Chip runs at 100 MHz
- Voltage reduction from 1.1V to 0.8V

  → Saving almost 50% energy

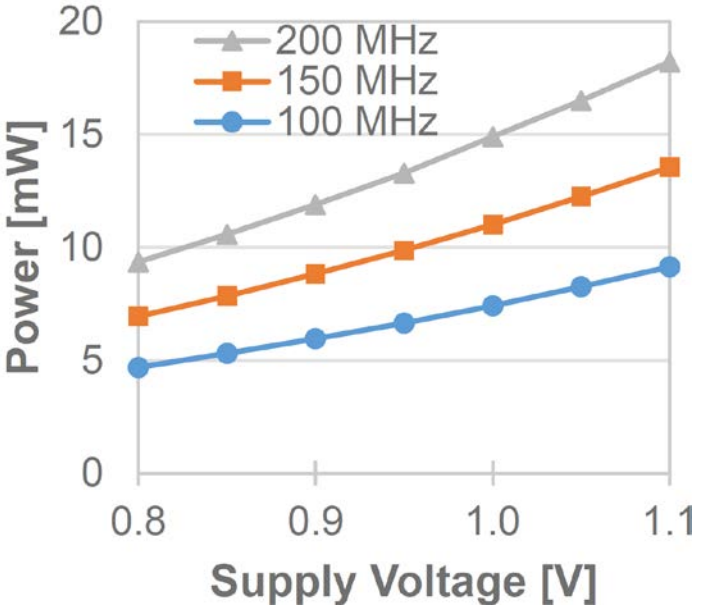

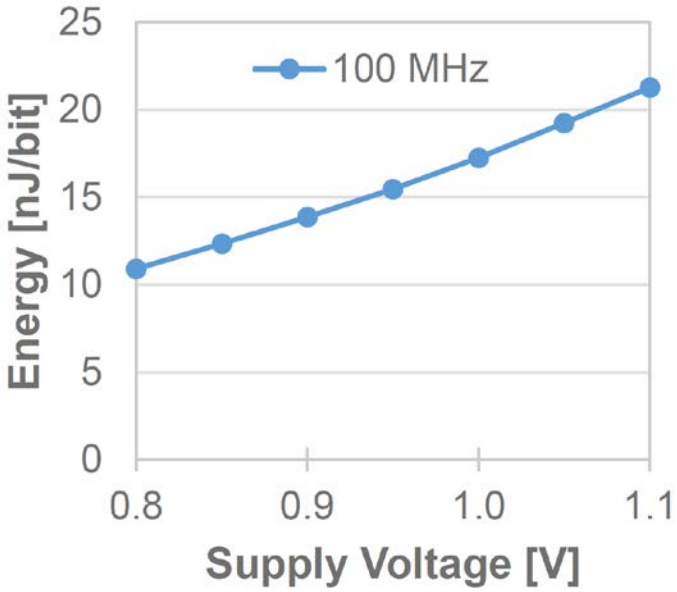

State-of-the-art Comparisons

System Overview	Max. frequency [MHz]	Memory Bandwidth [Gbit/s]	Join Throughput [Gbit/s]	Power [mW]	Energy [pJ/bit]
This work 4.95 mm² chip, 28 nm	**250**	**62.5**	**15.24**	**20.1**	**1.3**
ZYNQ [1]: Artix-7 FPGA, 85k logic cells, 1 GB DDR DRAM	200	68.0	0.55	723	1314
ZYNQ [1]: ARM dual-core SoC, 28 nm, 1 GB DDR3 DRAM	667	68.0	0.04	803	20075
FPGA [2]: 4x Virtex-6 FPGA, each 475k logic cells, each with 24 GB DDR DRAM	200	1228.8	80	19154 [3]	239

References:

[1] E. S. Chung, J. D. Davis, and J. Lee, “LINQits: Big Data on Little Clients,” in Proceedings of the 40th Annual International Symposium on Computer Architecture, ser. ISCA’13. ACM, 2013, pp. 261–272.

[2] J. Casper and K. Olukotun, “Hardware Acceleration of Database Operations,” in Proceedings of the 2014 ACM/SIGDA International Symposium on Field-programmable Gate Arrays, ser. FPGA’14, 2014, pp. 151–160.

[3] Xilinx, Virtex-6 FPGA Data Sheet: DC and Switching Characteristics, March 2014, v3.6.

Summary

- Database accelerator core
  - Tensilica RISC core with instruction set extensions
  - Database operators for energy-efficient query processing
- Hardware/Software co-design approach with sophisticated tool flow
  - Cycle accurate simulator, processor generator
  - RTL code verification and simulation
  - Backend design, PCB development
- Titan3D Chip
  - Fabricated chip with performance and power measurements
  - Improvements on energy efficiency by three orders of magnitude compared to state-of-the-art systems

Finally

The presented results are published in the following paper:

S. Haas, O. Arnold, S. Scholze, S. Höppner, G. Ellguth, A. Dixius,
A. Ungethüm, E. Mier, B. Nöthen, E. Matus, S. Schiefer, L. Cederstroem,
F. Pilz, C. Mayr, R. Schüffny, W. Lehner, and G. P. Fettweis,
"A Database Accelerator for Energy-Efficient Query Processing and Optimization",
in 2016 IEEE Nordic Circuits and Systems Conference
(NORCAS), Nov 2016, pp. 1–5.

Acknowledgements:
We would like to thank Cadence and Tensilica for providing software tools an IP, the Database Group for the fruitful ideas as well as the Chair for *Highly-Parallel VLSI Systems and Neuro-Microelectronics* of TU Dresden for backend design and PCB development of the Titan3D chip.

Thank you!