TECHNISCHE UNIVERSITÄT DRESDEN

Dresden Database Systems Group

Database Operations on top of Complex System Design

Nusrat Jahan Lisa

eBISS 2019, Berlin, Germany, July 5, 2019

AALBORG UNIVERSITET

IT4BI

Outline

Dresden Database Systems Group

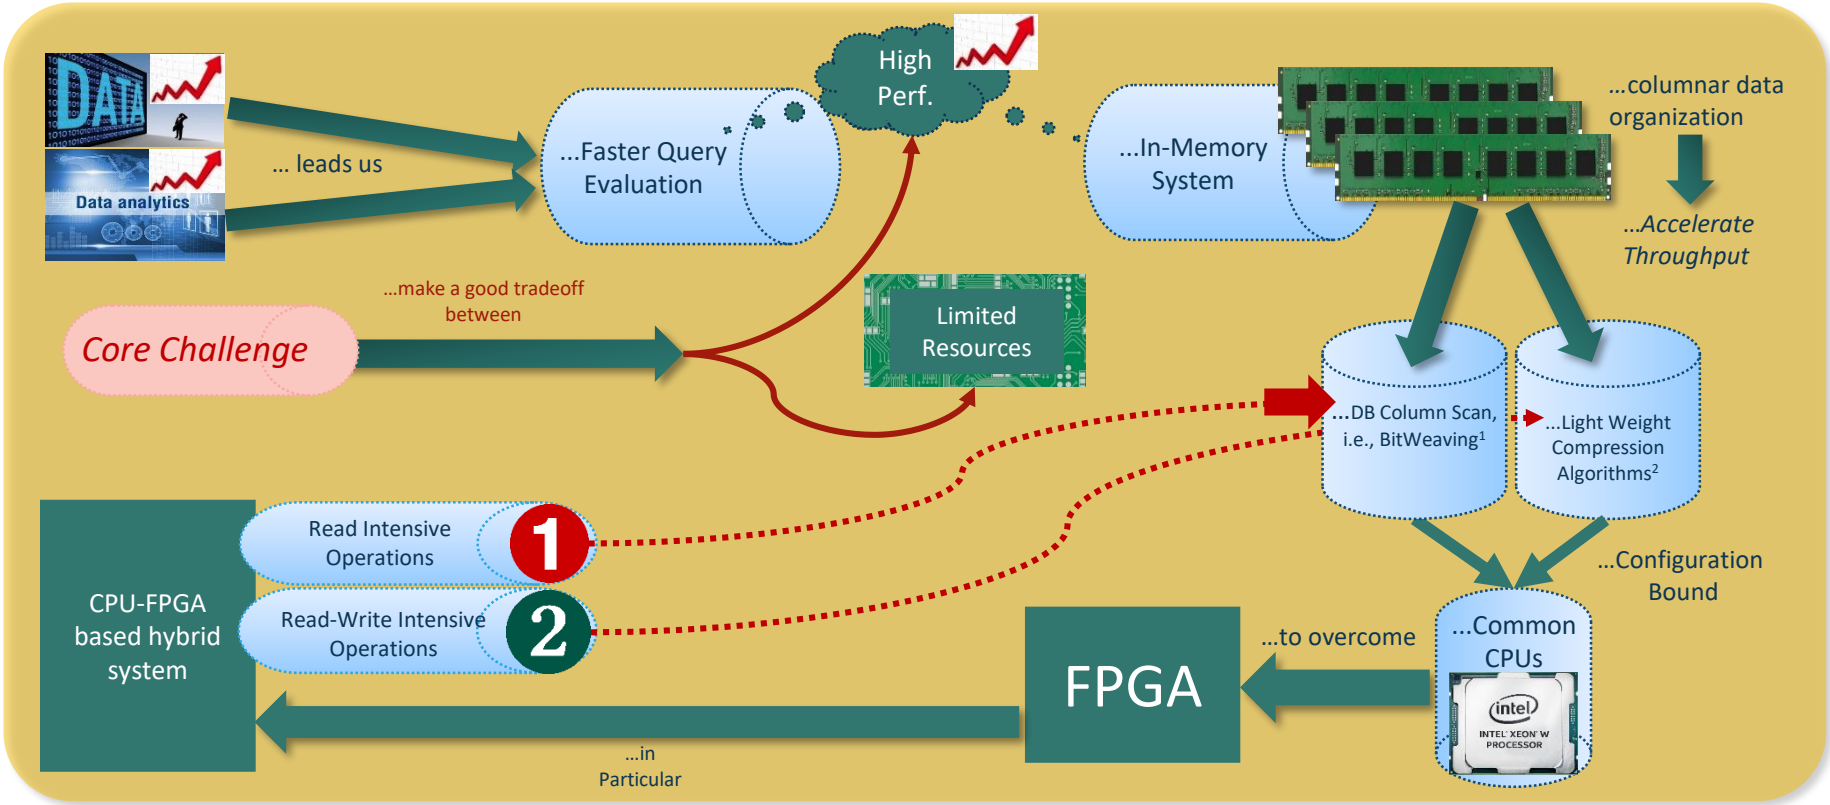

[1]Yinan Li and Jignesh M. Patel, BitWeaving: fast scans for main memory data processing. SIGMOD, pp.289-300, 2013.
[2]D. Lemire and L. Boytsov. Decoding billions of integers per second through vectorization. Softw., Pract. Exper., 45:1–29, 2015.

DB Column Scan

Dresden Database Systems Group

Recent Scan Approach for Compressed Columnar Data

- exploits the intra-instruction parallelism at the bit-level of modern processors.
- multiple compressed data are packed horizontally/vertically into processor words.

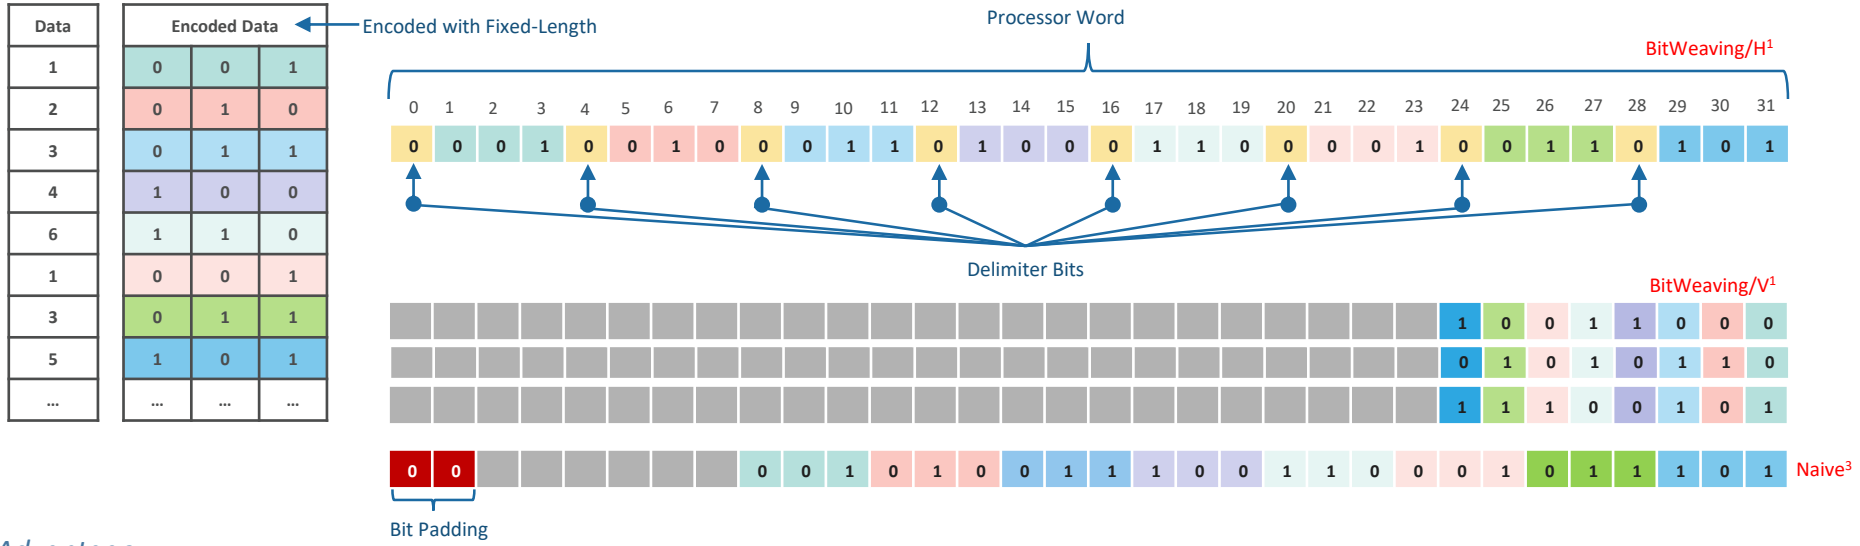

Advantage

- Evaluate any kind of predicates (Equality, Greater than, etc.) directly on the encoded data.

[1]Yinan Li and Jignesh M. Patel, BitWeaving: fast scans for main memory data processing. SIGMOD, pp.289-300, 2013.
[3]Lamport, L.: Multiple byte processing with full-word instructions. Communications of the ACM 18(8), 471{475 (1975).

Evaluating Predicates

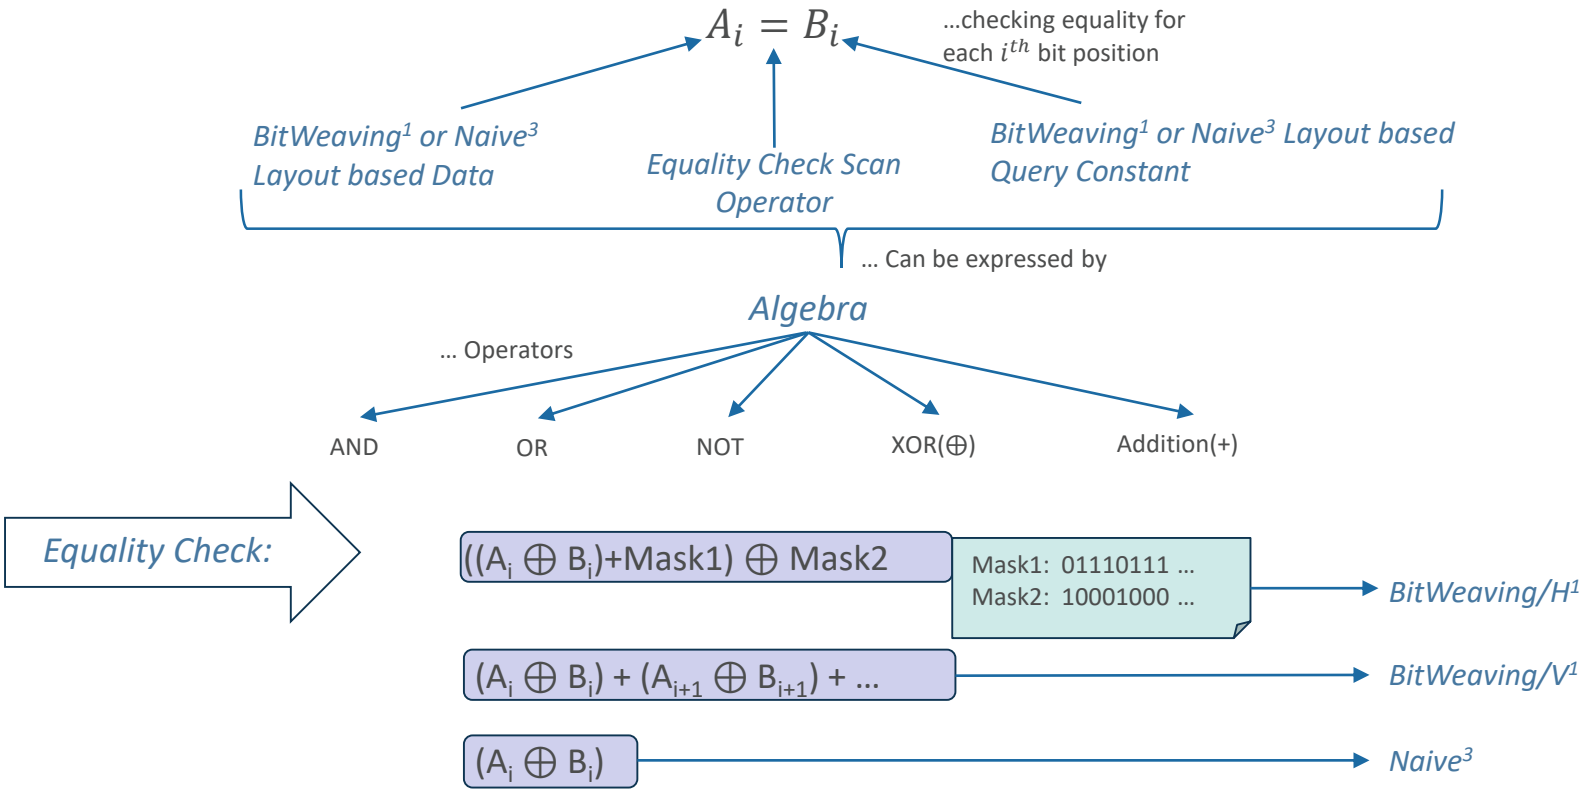

[1]Yinan Li and Jignesh M. Patel, BitWeaving: fast scans for main memory data processing. SIGMOD, pp.289-300, 2013.
[3]Lamport, L.: Multiple byte processing with full-word instructions. Communications of the ACM 18(8), 471{475 (1975).

Read Intensive Operations

Dresden Database Systems Group

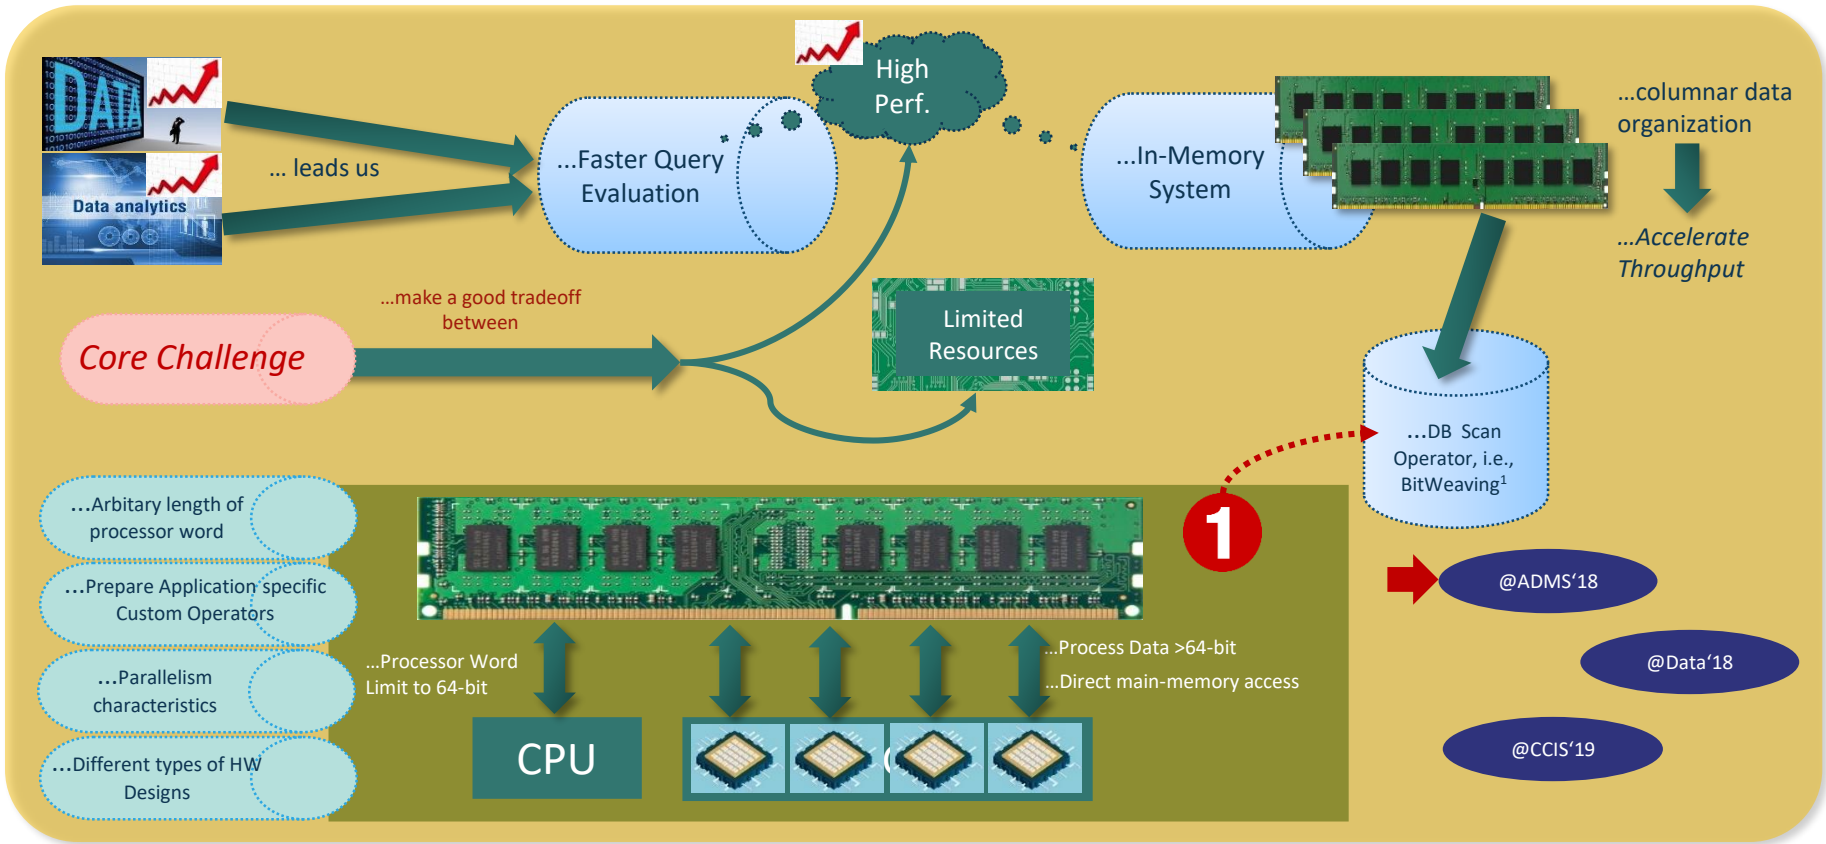

[1]Yinan Li and Jignesh M. Patel, BitWeaving: fast scans for main memory data processing. SIGMOD, pp.289-300, 2013.

Basic Design

@ADMS'18

Dresden Database Systems Group

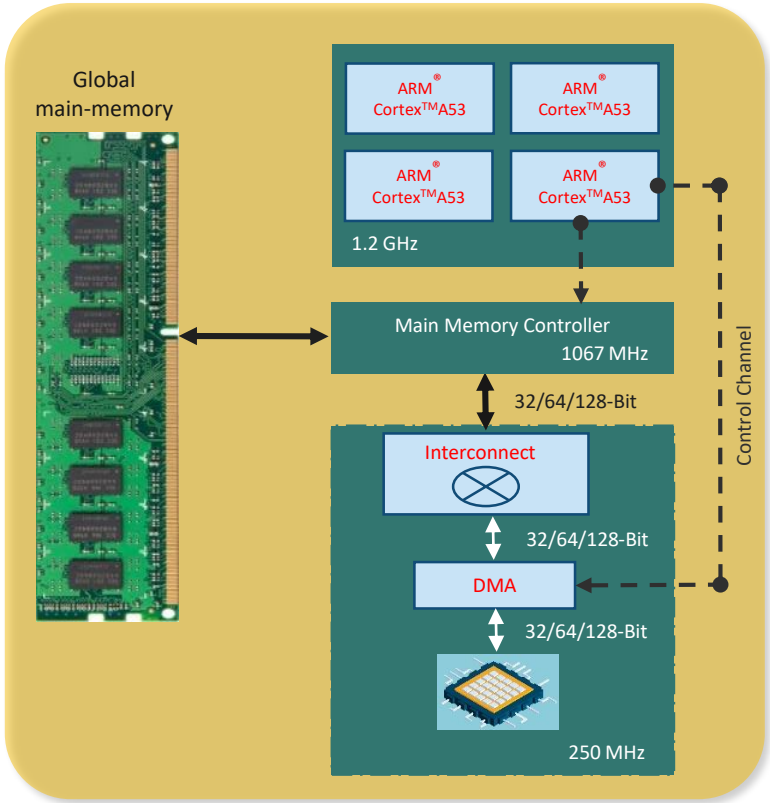

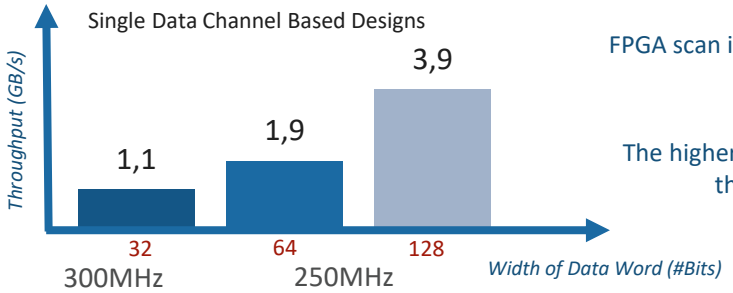

FPGA scan is quite efficient

The higher the „data width" – the better !?

Design Configurations

@ADMS'18

Dresden Database Systems Group

Target CPU-FPGA System: Zynq UltraScale+

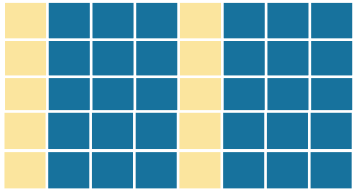

Keep BitWeaving/H Layout based data in main-memory …

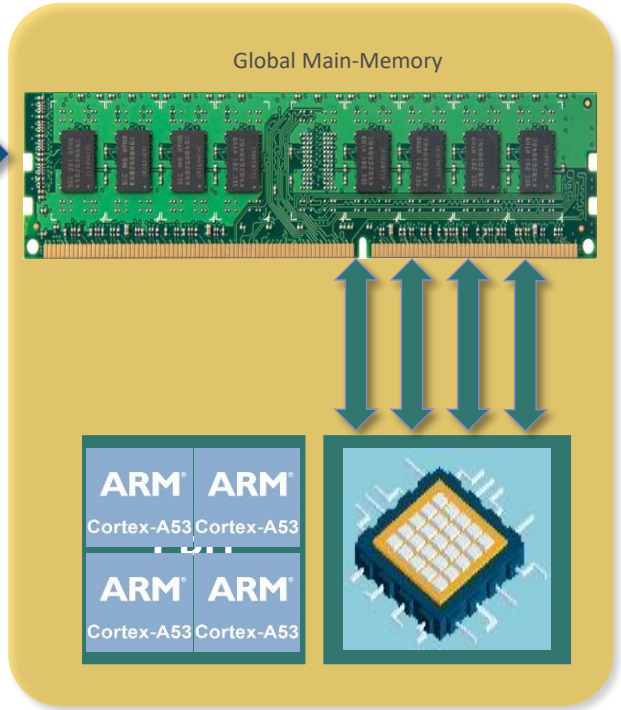

Multiple Data Channels Based Designs

…in order to

extend data width

- *Basic Design*
- *Combined Design*
- *Independent Design*
- *Hybrid Design*
- *Hybrid² Design*

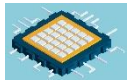

Scan Processing Element

ZYNQ UltraSCALE+

Combined Design

@ADMS'18

Dresden Database Systems Group

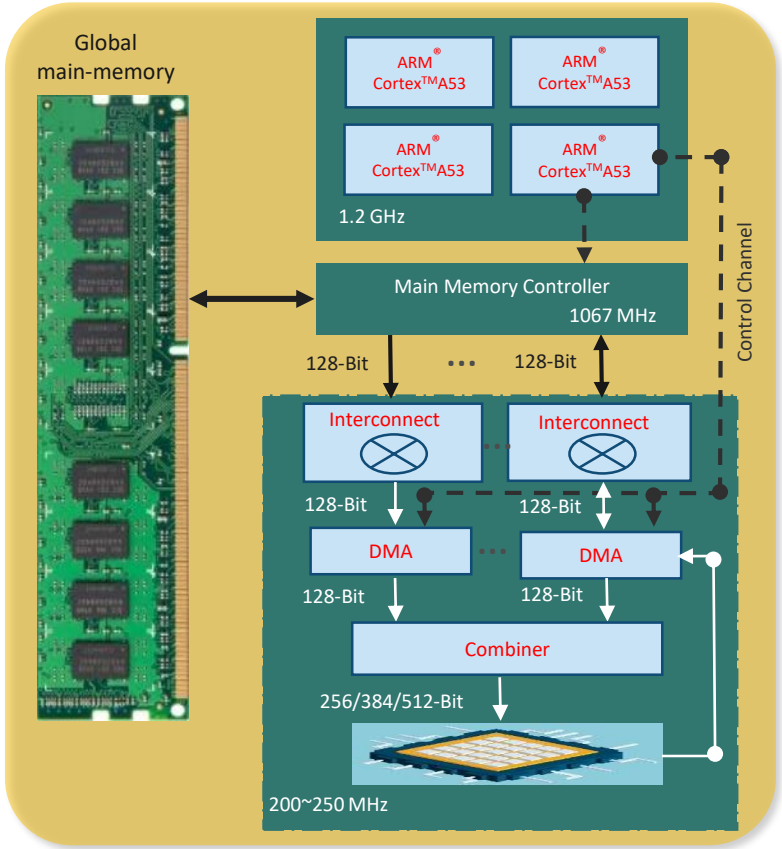

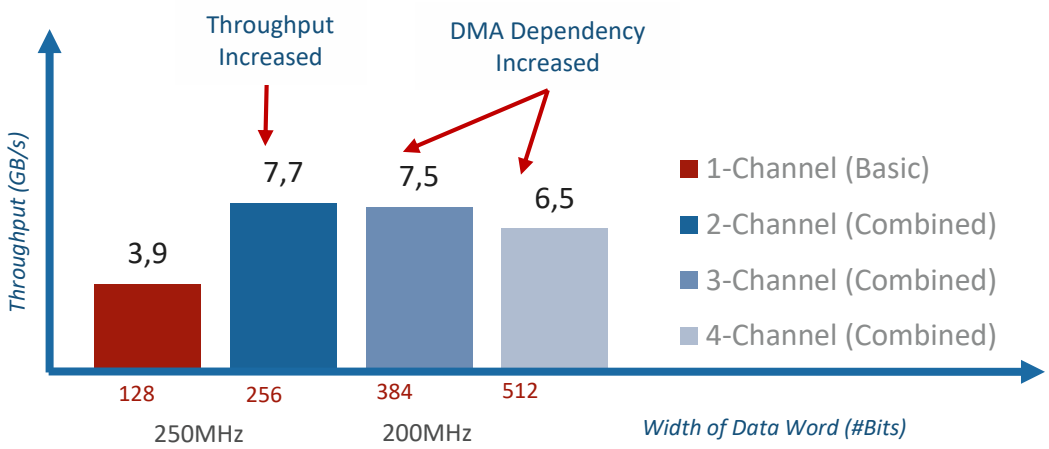

Next Step

Process multiple 128-bit width based data independently by using more than one processing element instead of one.

→ *Independent Design*

Independent Design

@ADMS'18

Dresden Database Systems Group

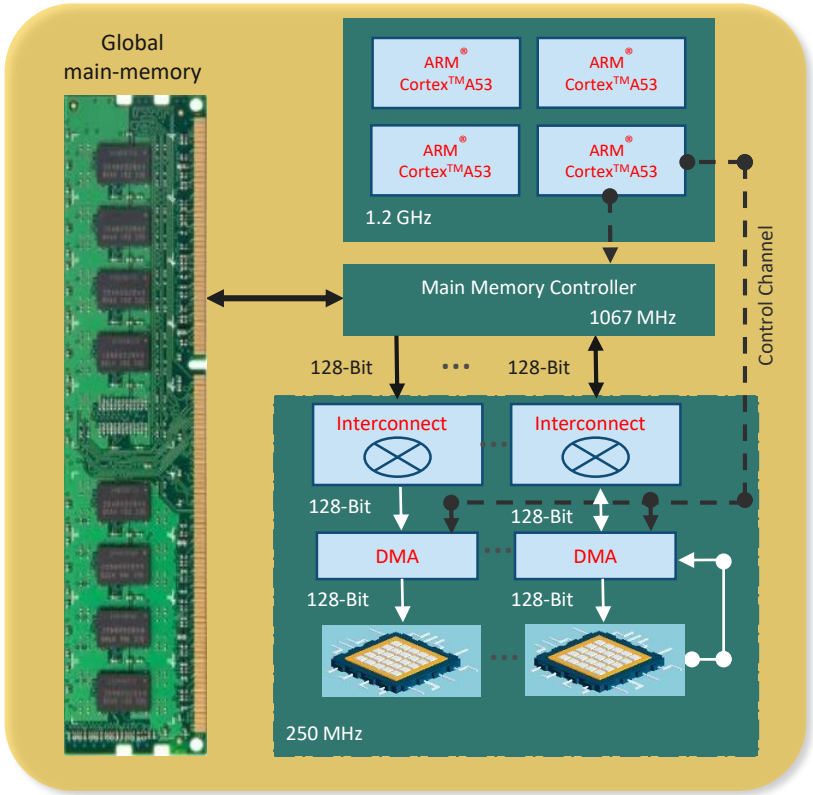

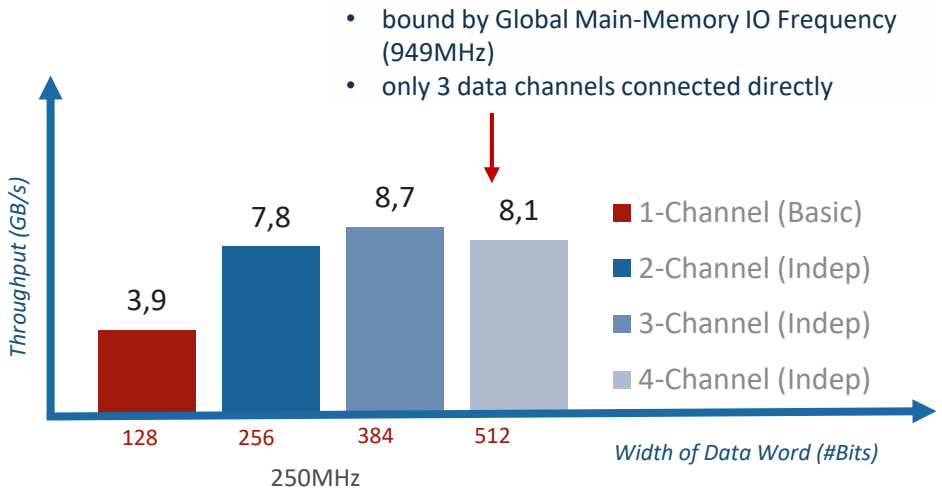

Design Configurations

@ADMS'18

Dresden Database Systems Group

Target CPU-FPGA System: Zynq UltraScale+

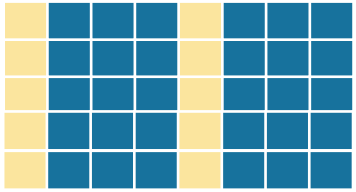

Keep BitWeaving/H Layout based data in main-memory …

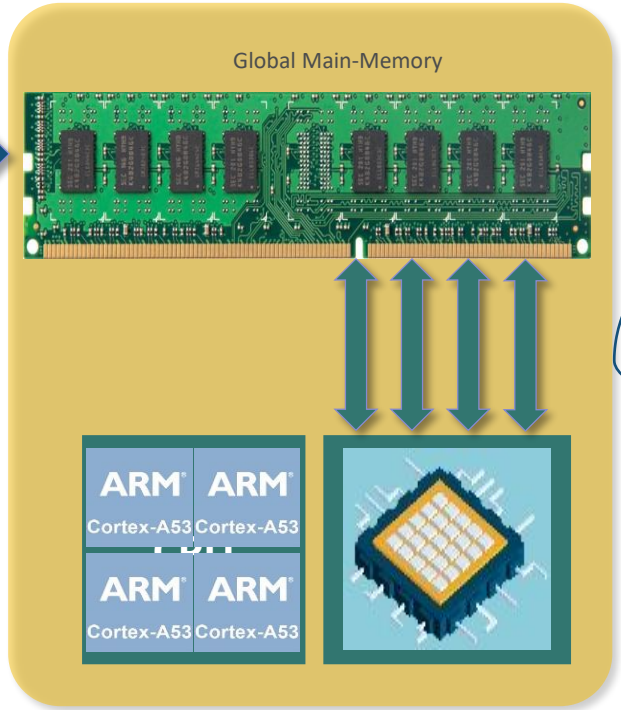

Basic Design

256-bit Width based Data Processing is beneficial ← *Combined Design*

Multiple instances beneficial ← *Independent Design*

…Integrating → *Hybrid Design*

Hybrid² Design

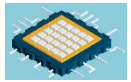

Scan Processing Element

ZYNQ UltraSCALE+

Hybrid Design

@ADMS'18

Dresden Database Systems Group

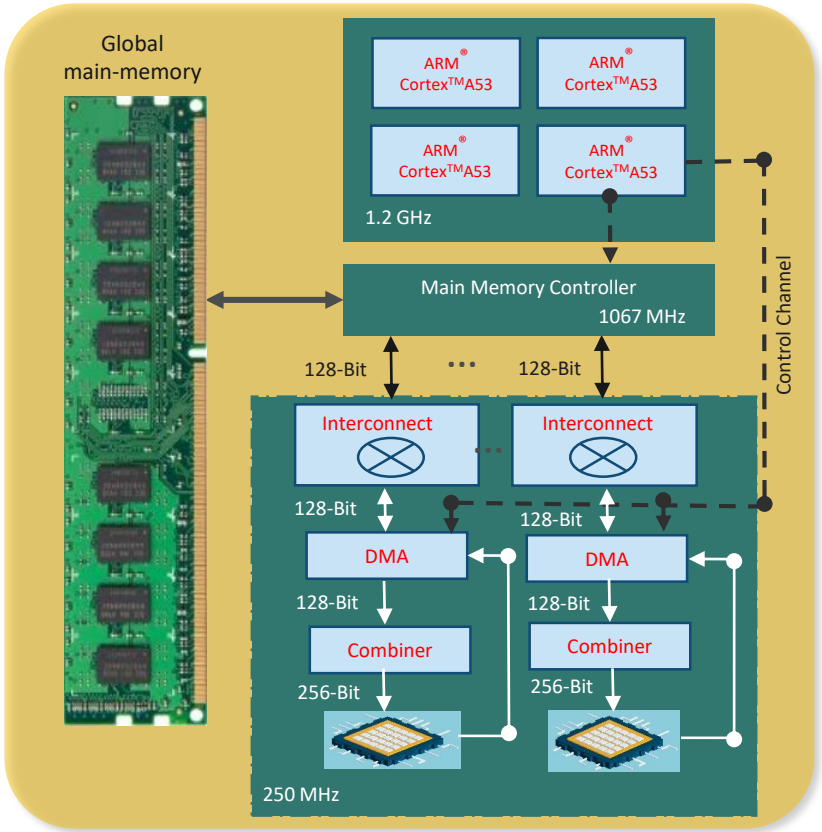

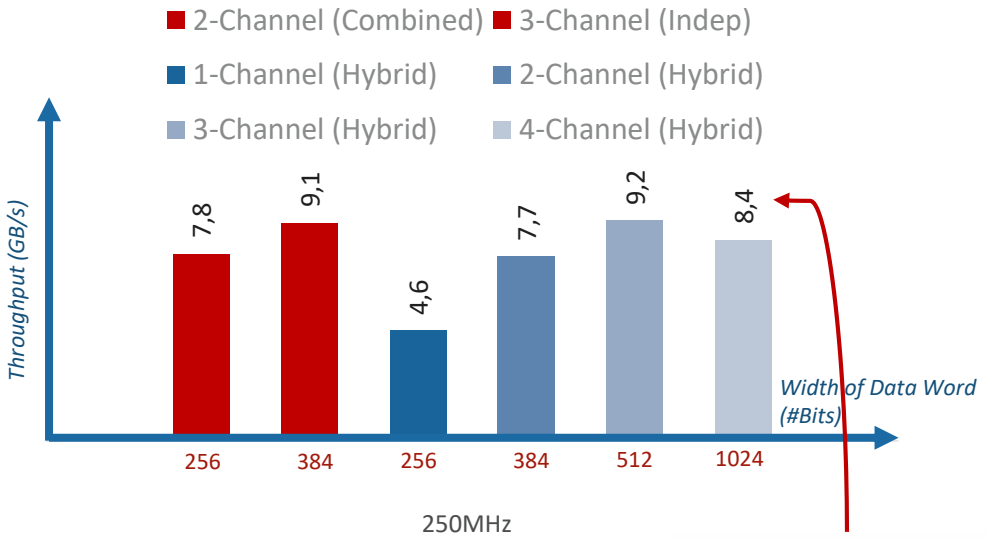

Bound with Global Main-Memory IO Frequency

Design Configurations

@ADMS'18

Dresden Database Systems Group

Target CPU-FPGA System: Zynq UltraScale+

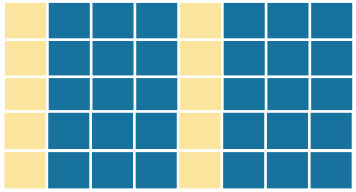

Keep BitWeaving/H Layout based data in main-memory …

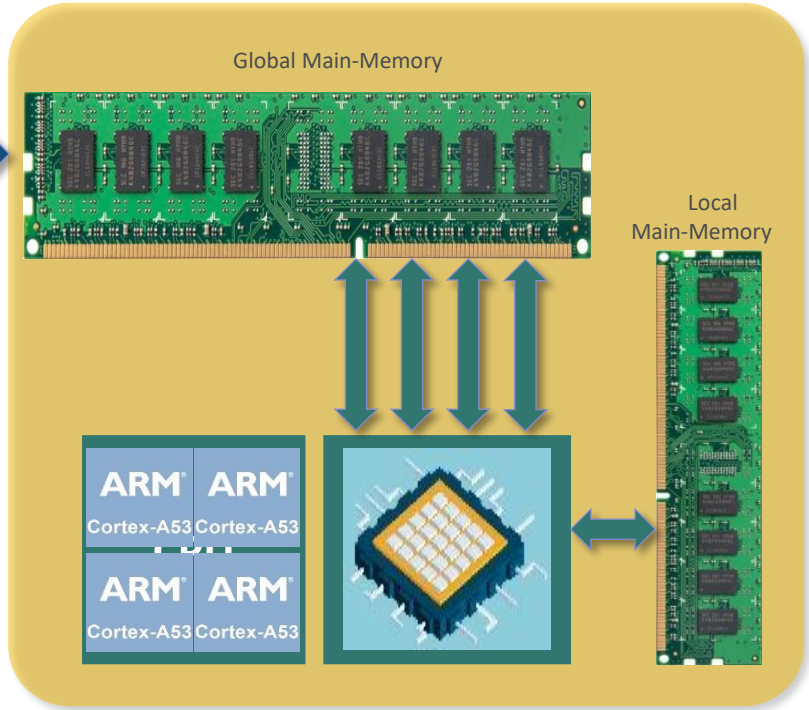

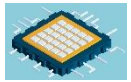

Scan Processing Element

ZYNQ UltraSCALE+

Basic Design

Combined Design

Independent Design

Hybrid Design

Multiple Main-Memory ← *Hybrid² Design*

Hybrid² Design

@ADMS'18

Dresden Database Systems Group

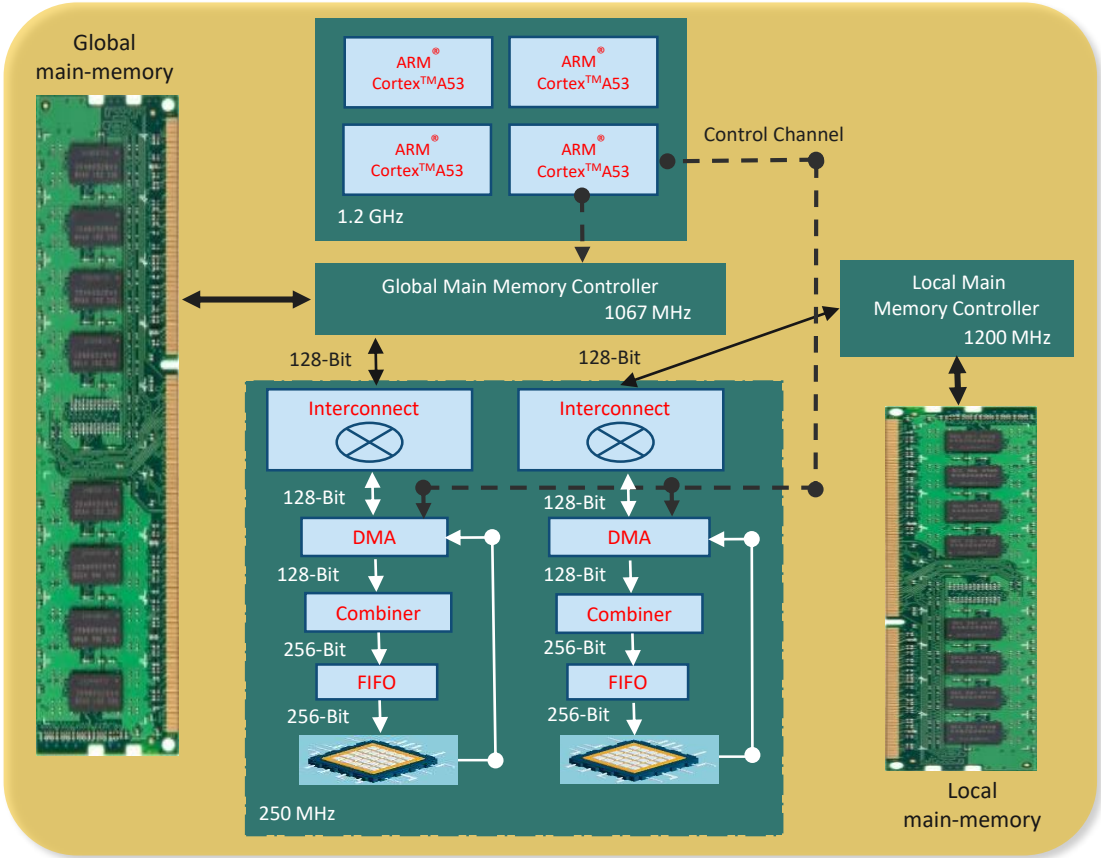

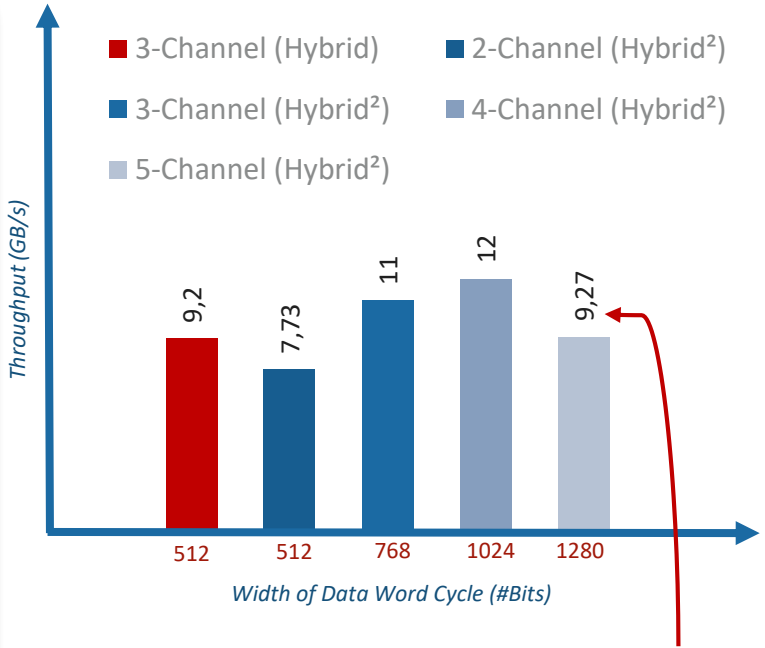

Bound with Global Main-Memory IO Frequency

Summary

@ADMS'18

Dresden Database Systems Group

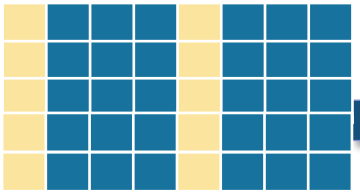

Keep BitWeaving/H Layout based data in main-memory …

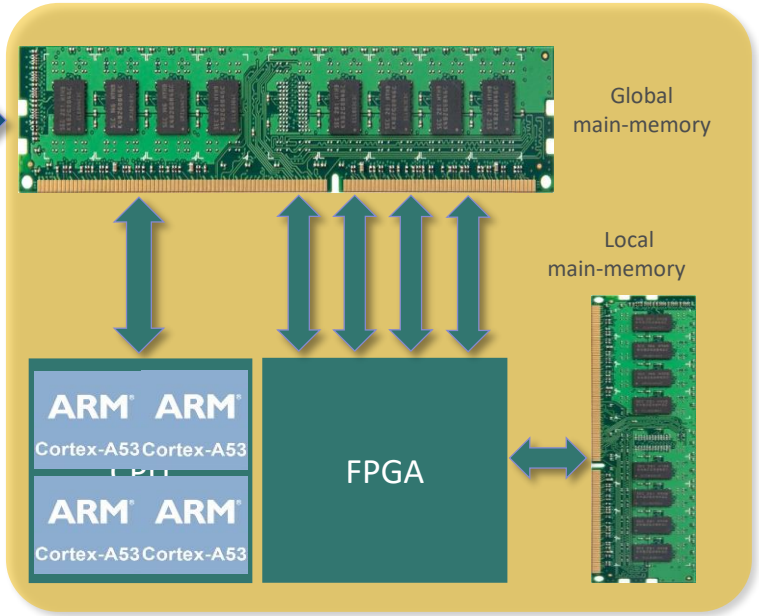

4-Core Processing **4.8 GB/s**

3-Core Processing **4.1 GB/s**

2-Core Processing **3 GB/s**

1-Core Processing **1.9 GB/s**

Basic Design
3.9 GB/s

Combined Design
~7 GB/s

Independent Design
~9 GB/s

Hybrid Design
~9.2 GB/s

Hybrid² Design
12 GB/s

- FPGAs may be extremely useful for off-loading column-scan operations
- "Optimal" design is tradeoff between max performance and system complexity

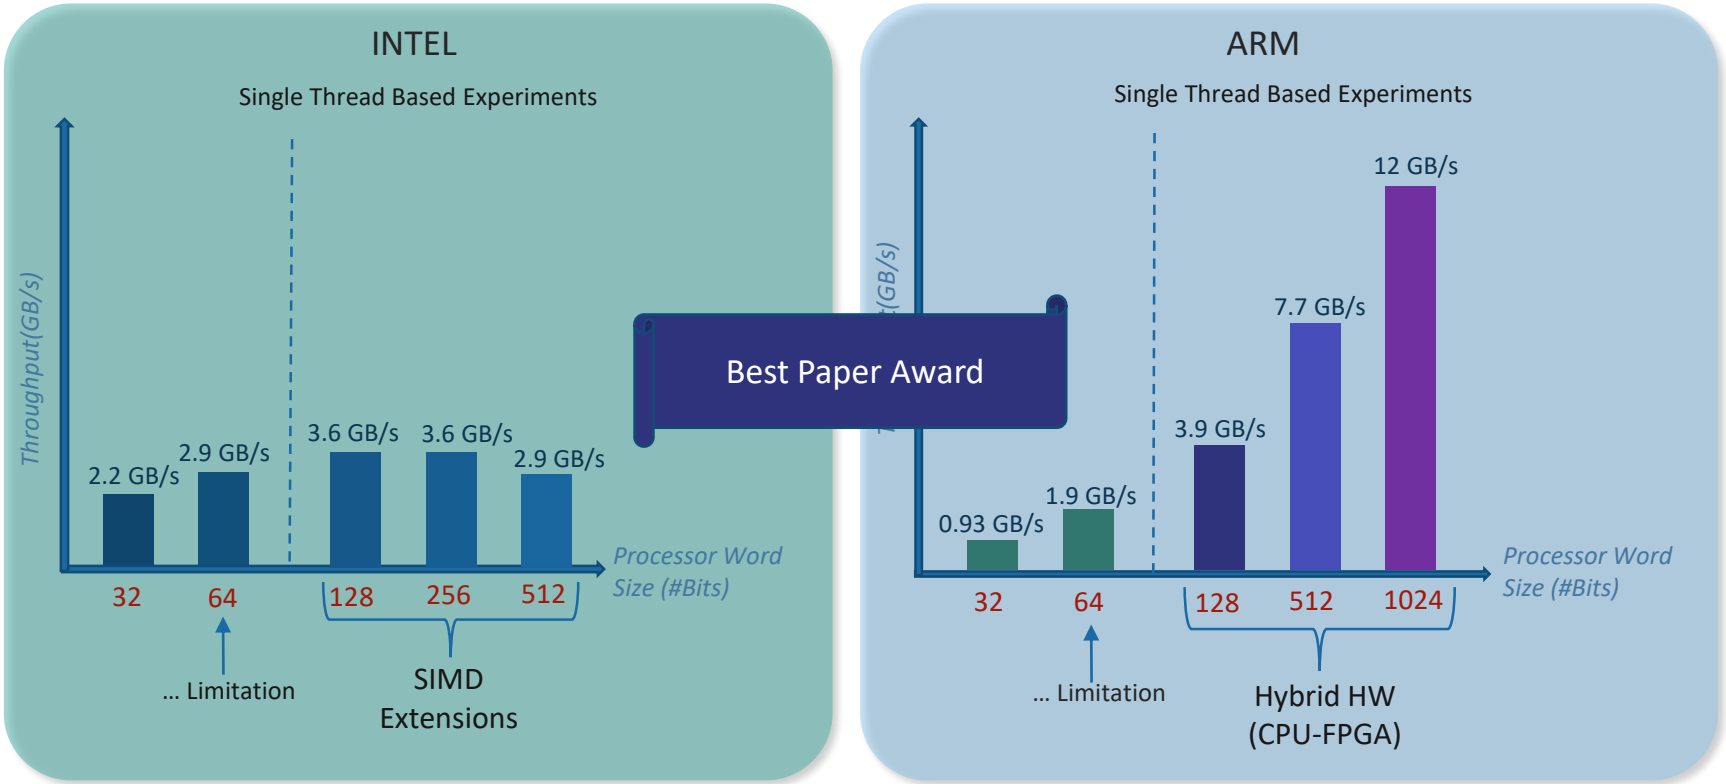
INTEL
Single Thread Based Experiments
Throughput(GB/s)
2.2 GB/s
2.9 GB/s
3.6 GB/s
3.6 GB/s
2.9 GB/s
32
64
128
256
512
Processor Word Size (#Bits)
... Limitation
SIMD Extensions
Best Paper Award
ARM
Single Thread Based Experiments
0.93 GB/s
1.9 GB/s
3.9 GB/s
7.7 GB/s
12 GB/s
32
64
128
512
1024
Processor Word Size (#Bits)
... Limitation
Hybrid HW (CPU-FPGA)

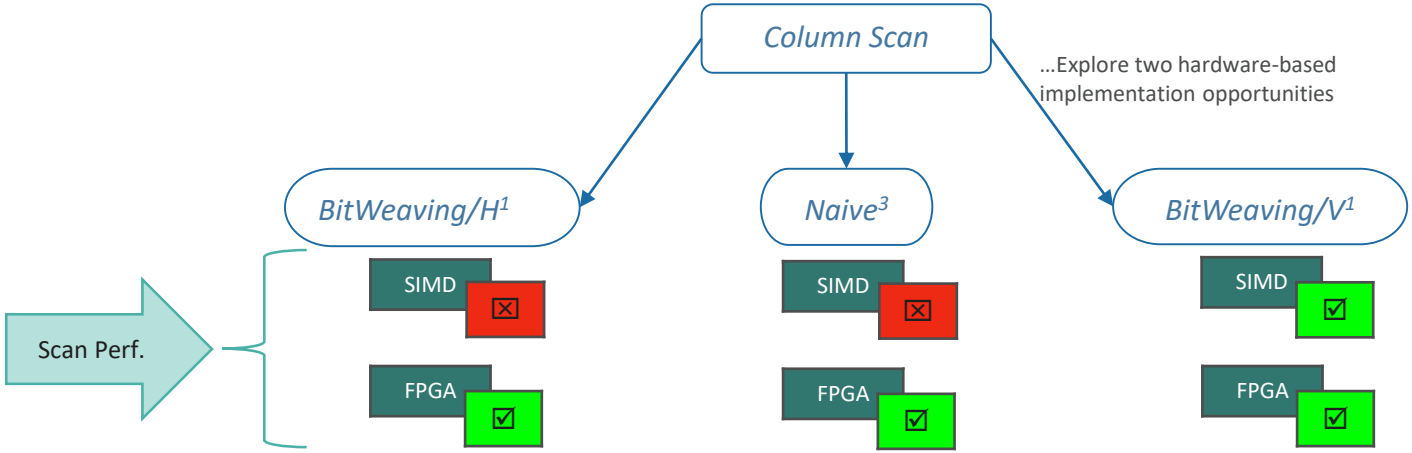

[1]Yinan Li and Jignesh M. Patel, BitWeaving: fast scans for main memory data processing. SIGMOD, pp.289-300, 2013.
[3]Lamport, L.: Multiple byte processing with full-word instructions. Communications of the ACM 18(8), 471{475 (1975).

FPGA Architecture

Dresden Database Systems Group

A Field Programmble Gate Array (FPGA) is a programmble logic device which is capable to implement any type of user defined logic function/circuit.

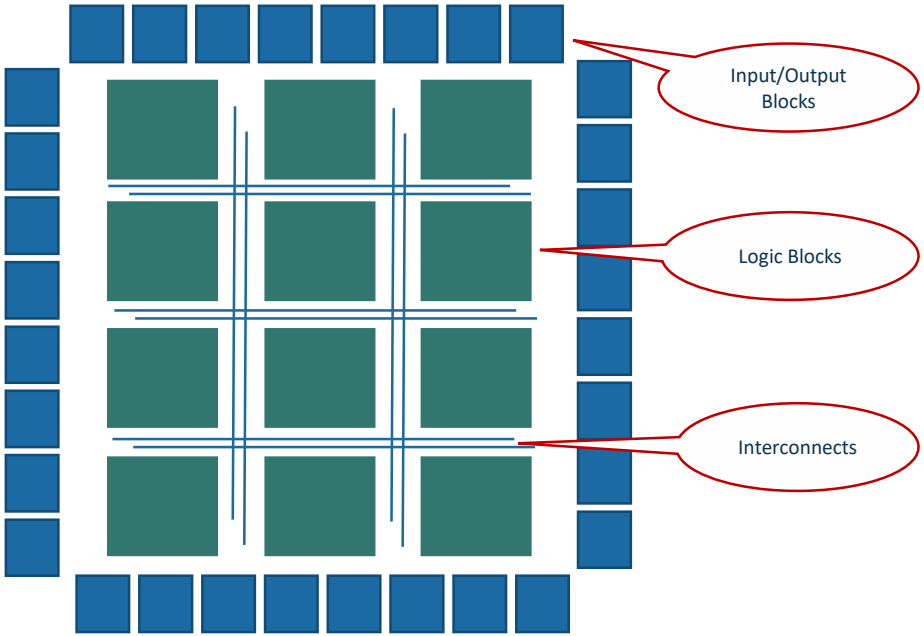

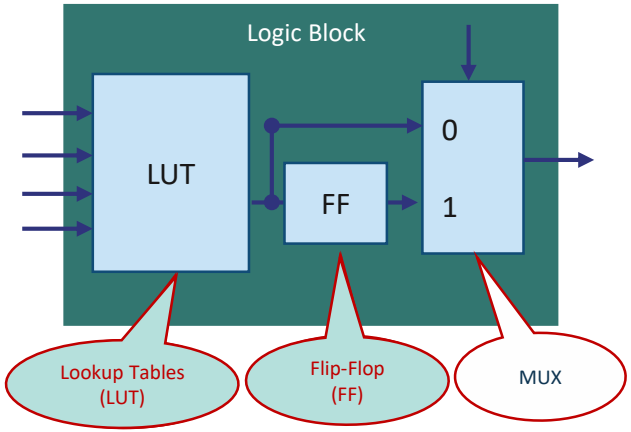

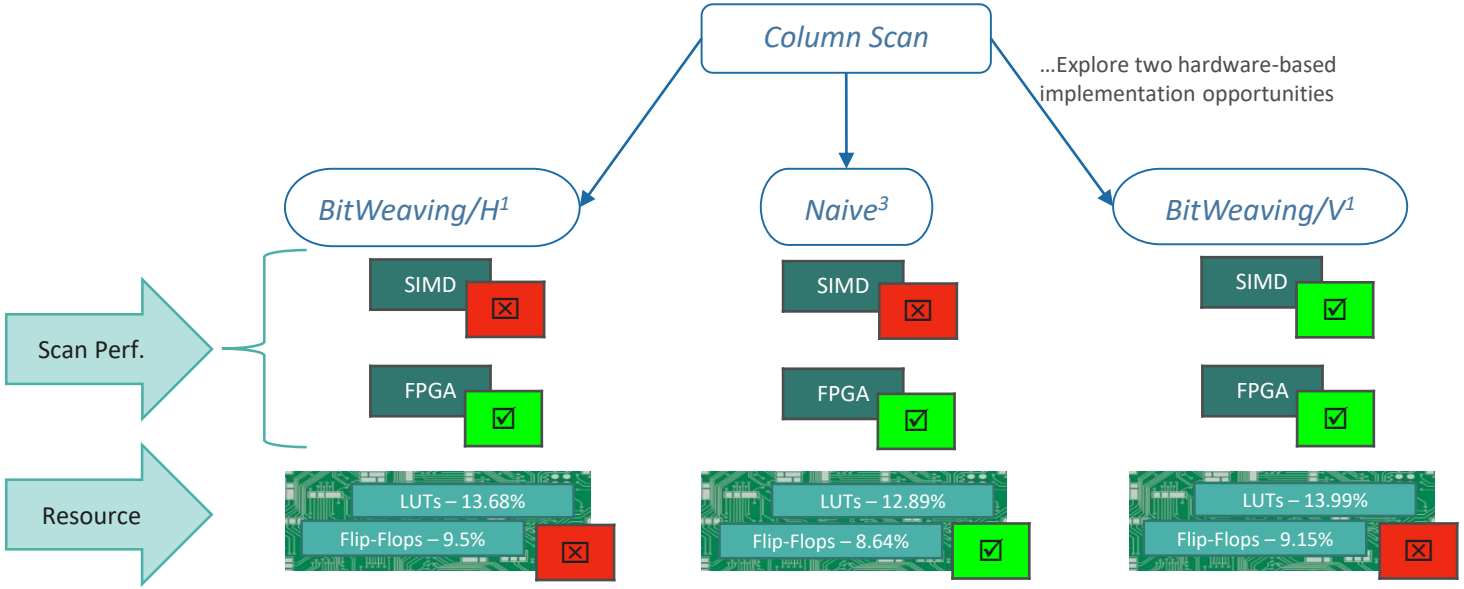

- FPGA is best for Naive technique and BitWeaving is perfect for SIMD.
- To improve scan performance through FPGA do not require any fancy scan mechanism as BitWeaving due to its high parallelism criteria and flexibility to configure hardware as per requirements.

[1]Yinan Li and Jignesh M. Patel, BitWeaving: fast scans for main memory data processing. SIGMOD, pp.289-300, 2013.
[3]Lamport, L.: Multiple byte processing with full-word instructions. Communications of the ACM 18(8), 471{475 (1975).

Read Intensive Operations: Summary

Dresden Database Systems Group

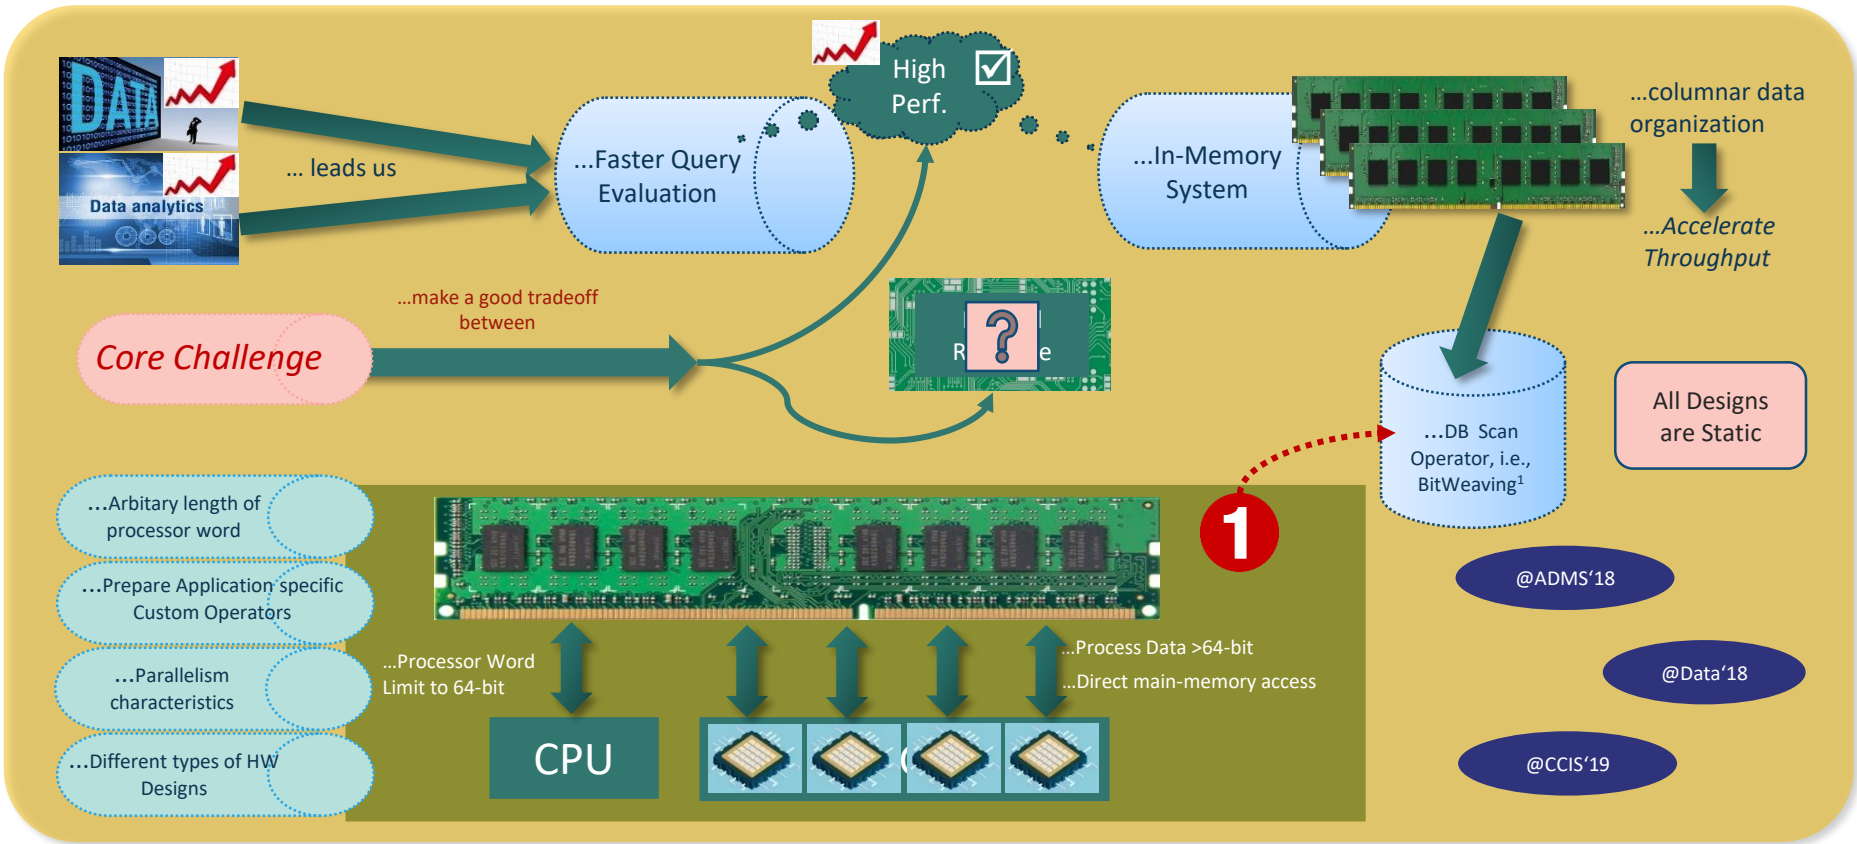

[1]Yinan Li and Jignesh M. Patel, BitWeaving: fast scans for main memory data processing. SIGMOD, pp.289-300, 2013.

Read-Write Intensive Operations

Dresden Database Systems Group

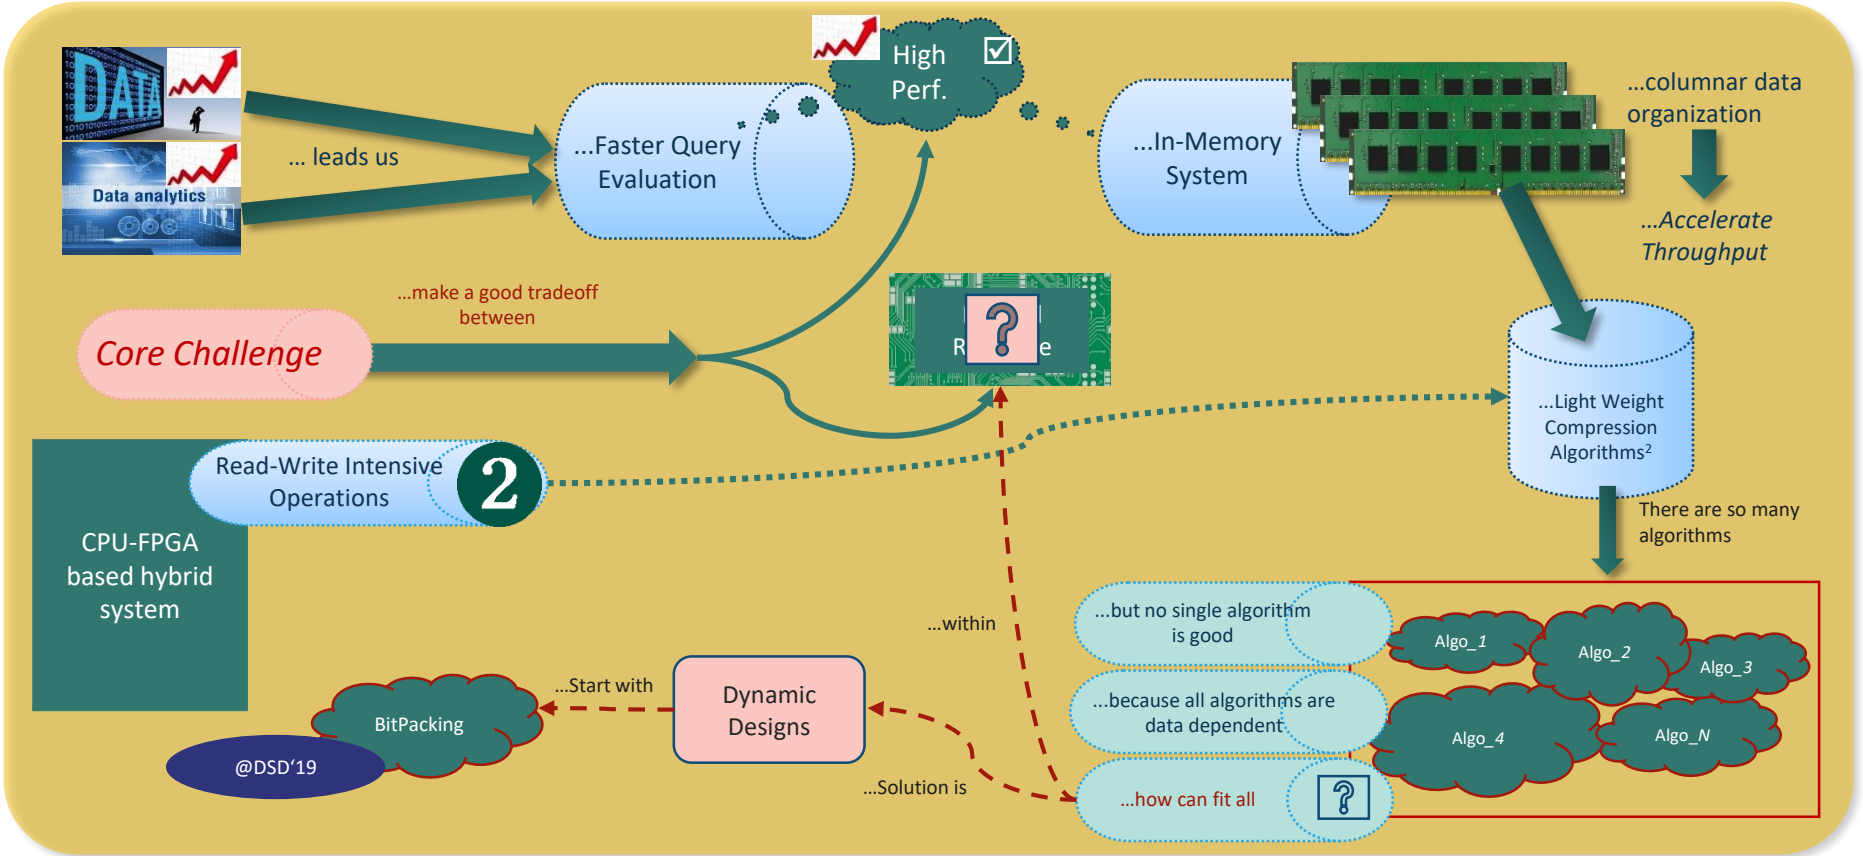

[2]D. Lemire and L. Boytsov. Decoding billions of integers per second through vectorization. Softw., Pract. Exper., 45:1–29, 2015.

BitPacking (BP) Compression

Dresden Database Systems Group

- BitPacking is one of the most applicable compression scheme in this domain showing a very good---not always optimal---behavior for different data properties

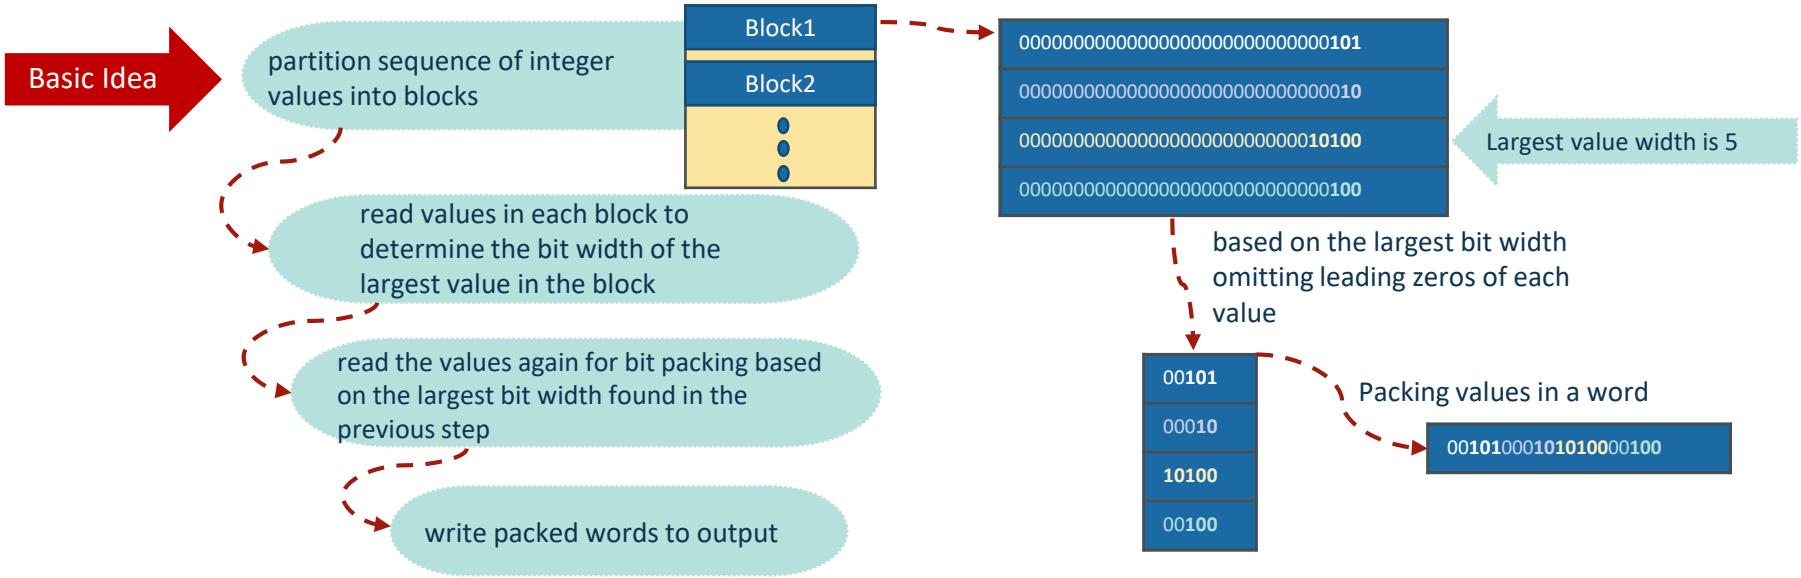

Pipeline-based BP Implementation

@DSD'19

Dresden Database Systems Group

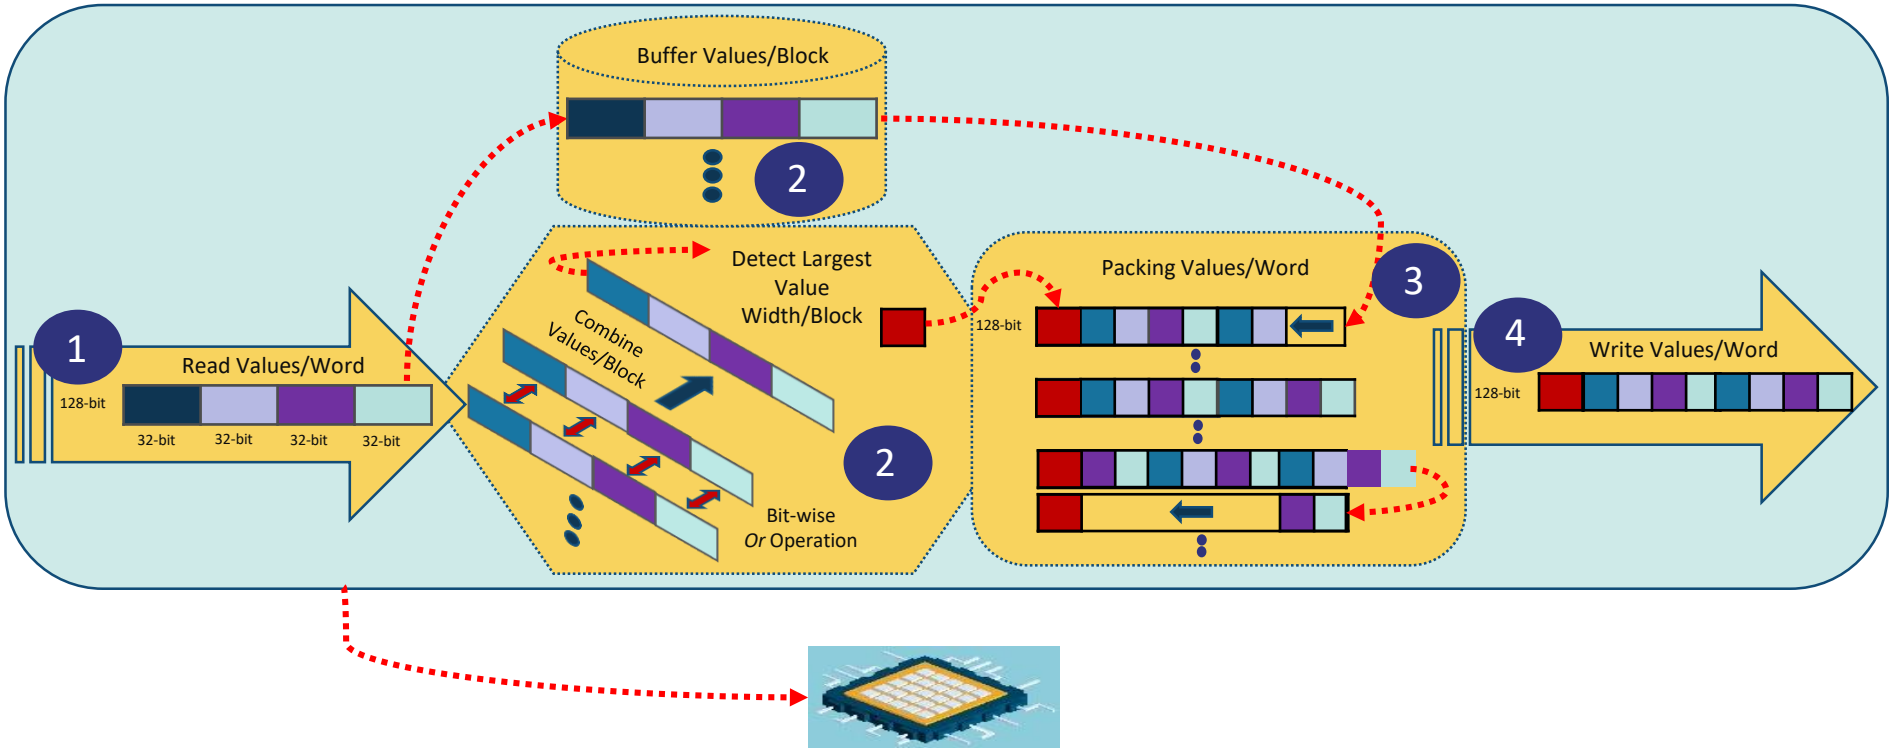

Custom BitPacking Hardware

BitPacking Static Design

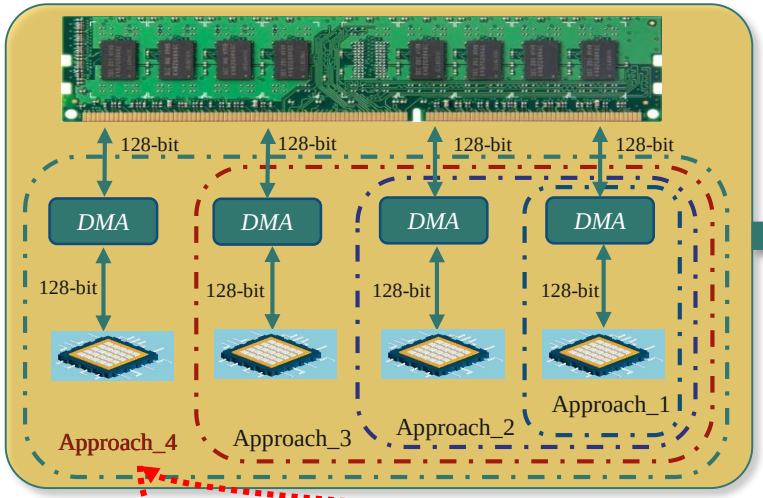

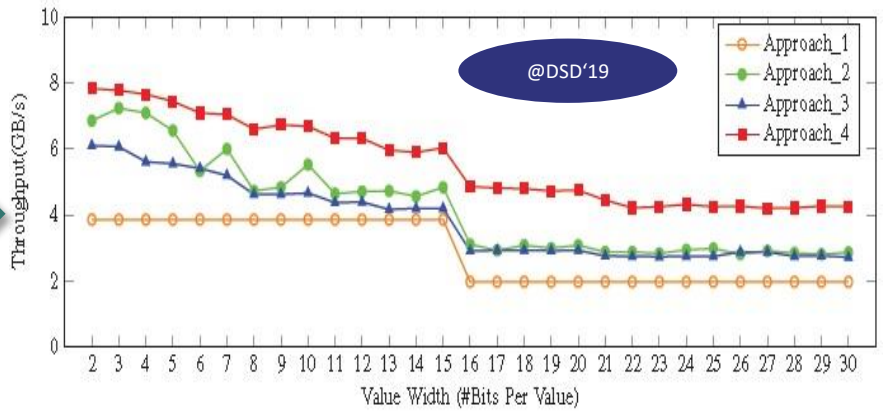

Best Static Design

...require

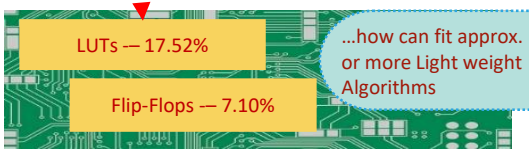

...how can fit approx. 100 or more Light weight Algorithms

...explore

Dynamic Design

...need to

Find out sub-design parts of a design

...so that

Load any specific sub-design during run time

...in order to

Optimize resource utilization

...for that

Include two Preprocessing Algo as Sub-design parts in BitPacking

...Apply Successive Deduction in a Block — Delta

...Find the smallest value; Perform deduction between each value and smallest value in a Block — FOR

BitPacking Dynamic Design

Dresden Database Systems Group

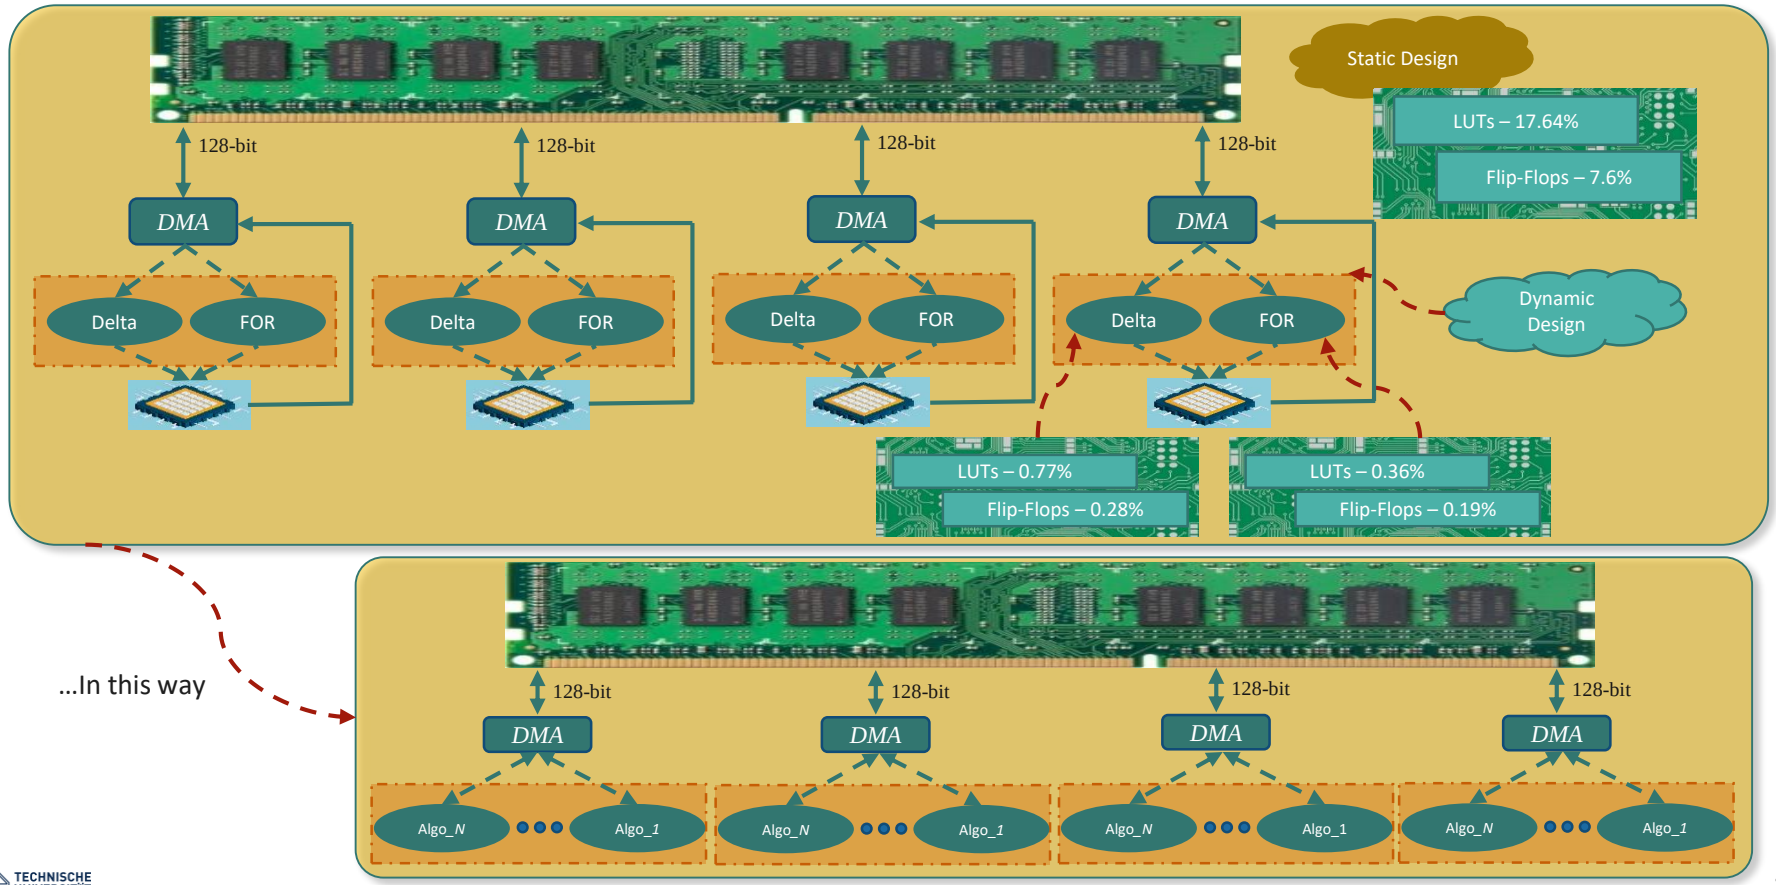

Read-Write Intensive Operations: Summary

Dresden Database Systems Group

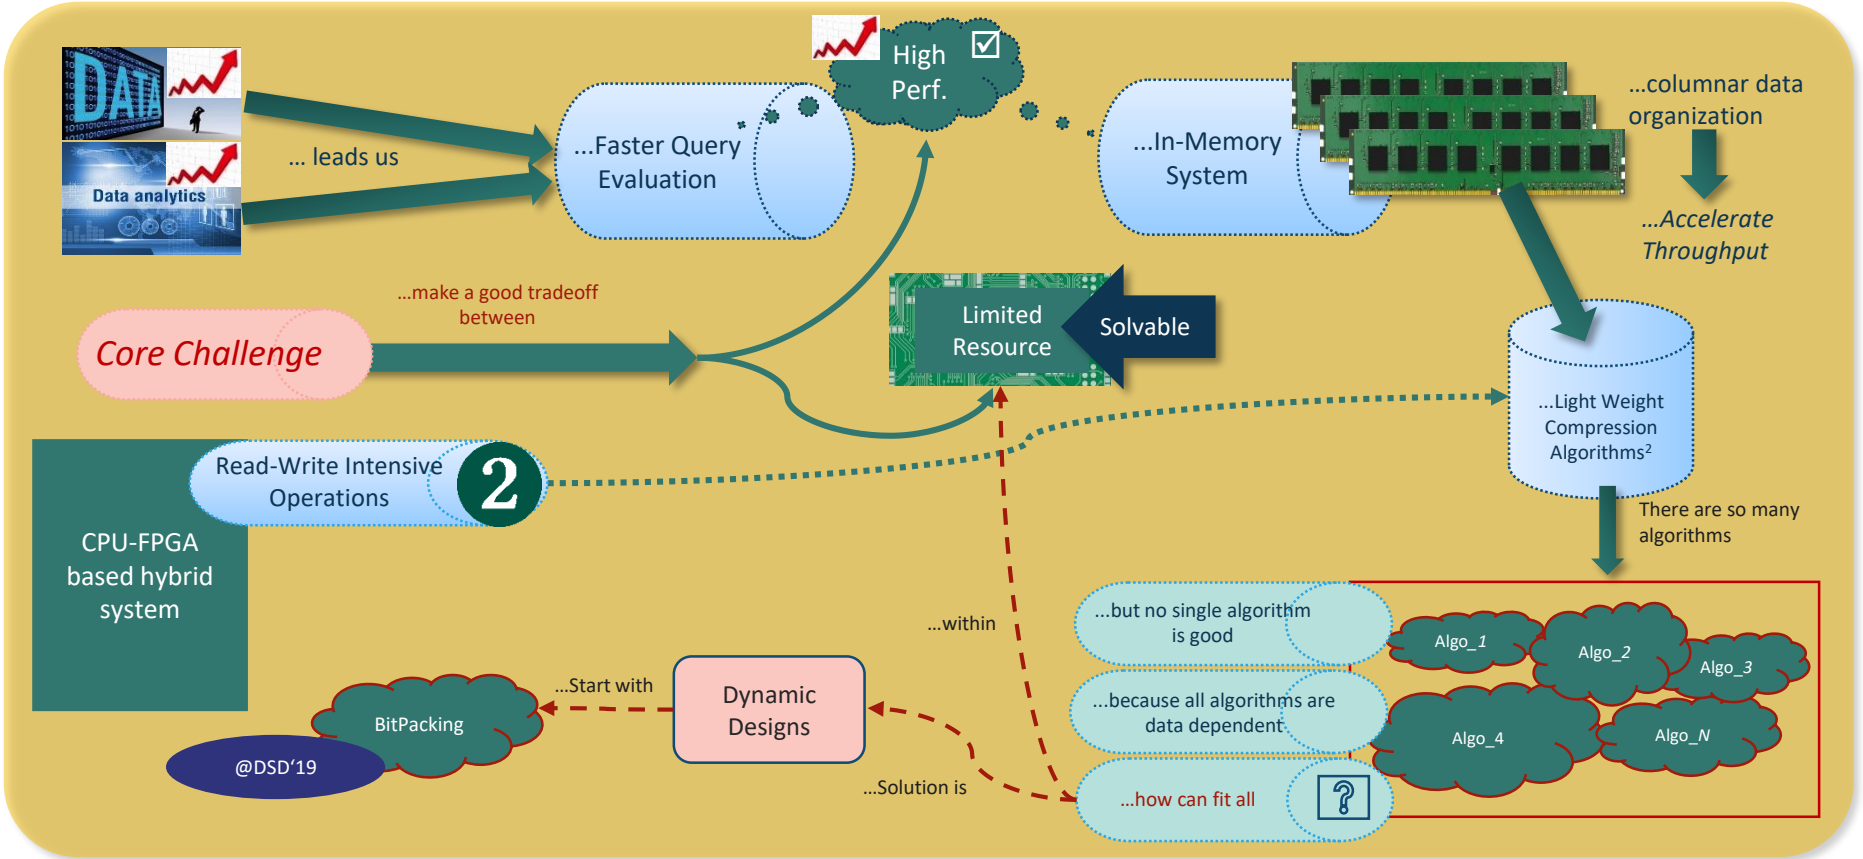

[2]D. Lemire and L. Boytsov. Decoding billions of integers per second through vectorization. Softw., Pract. Exper., 45:1–29, 2015.

Summary

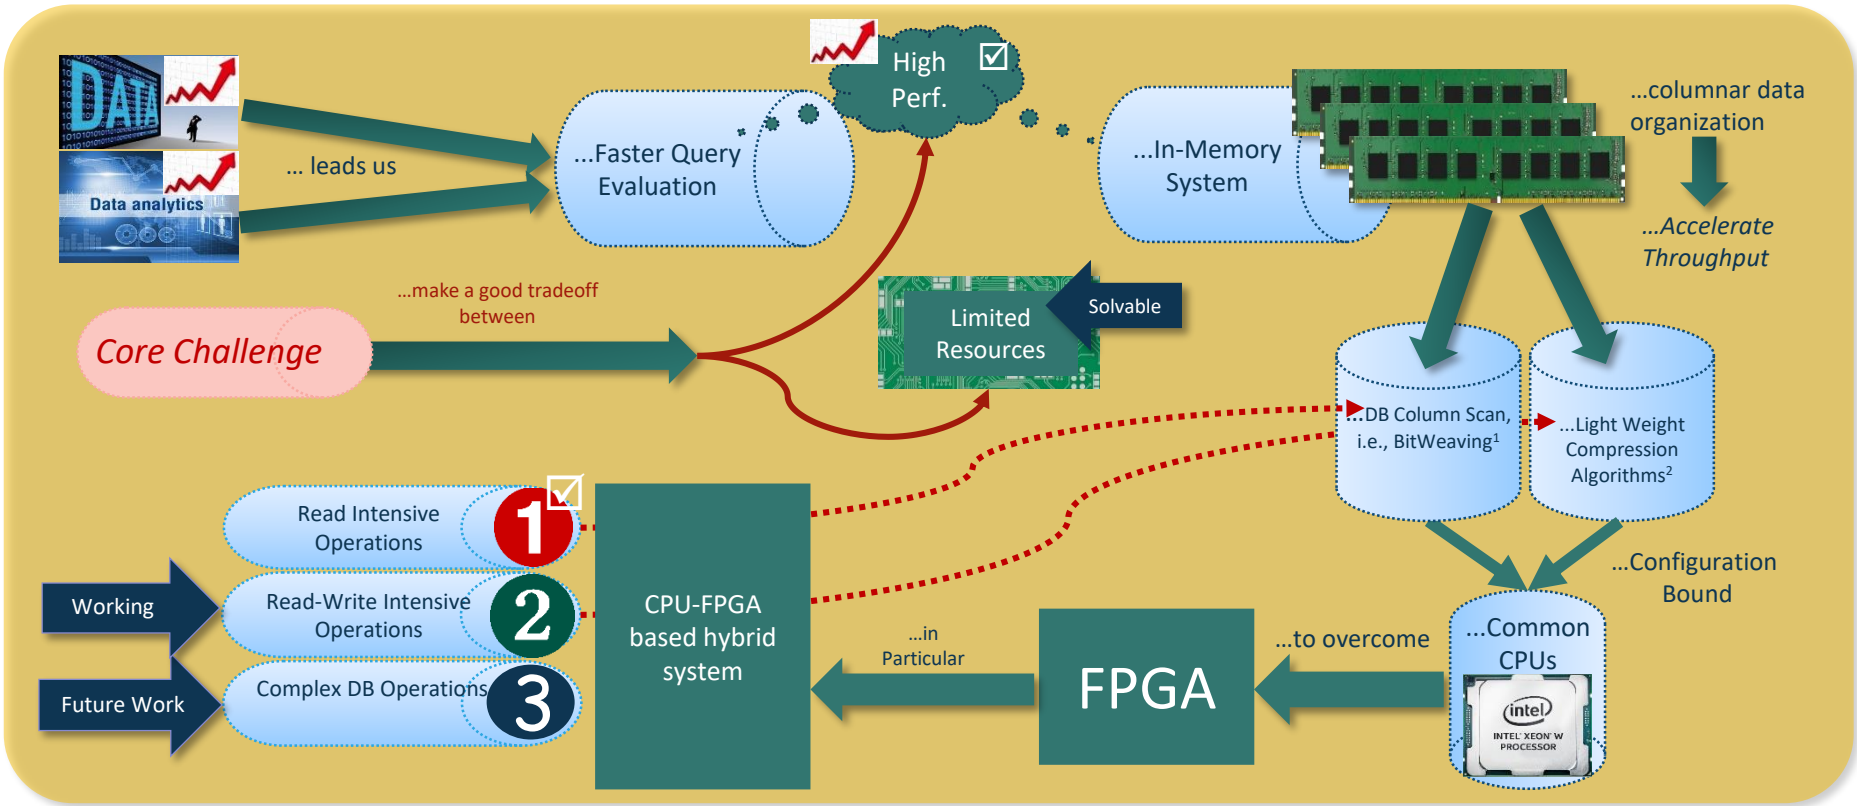

[1]Yinan Li and Jignesh M. Patel, BitWeaving: fast scans for main memory data processing. SIGMOD, pp.289-300, 2013.
[2]D. Lemire and L. Boytsov. Decoding billions of integers per second through vectorization. Softw., Pract. Exper., 45:1–29, 2015.

TimeLine

Dresden Database Systems Group

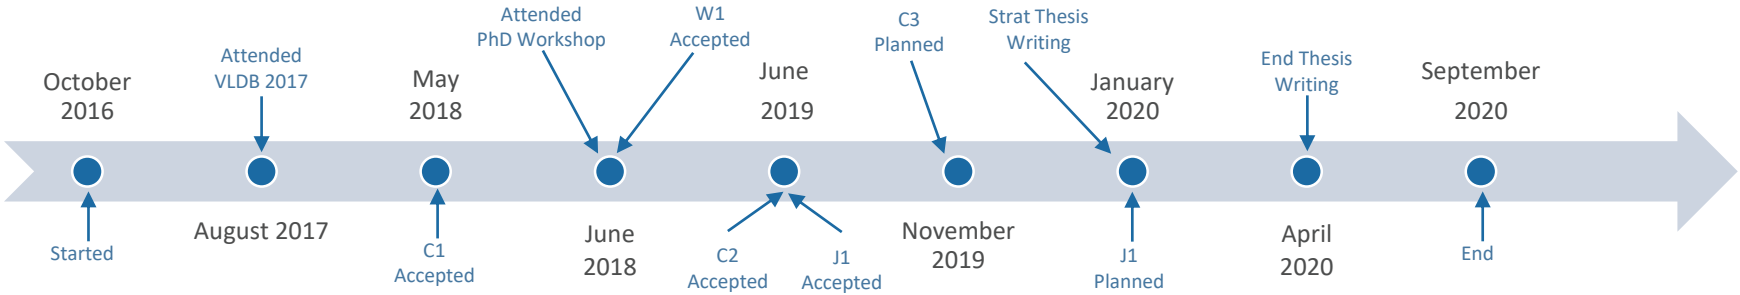

Publishing of papers

Papers	Conference	Workshop	Journal	Status
C1: Column scan optimization by increasing intra-instruction parallelism.	DATA			☑
W1: Column Scan Acceleration in Hybrid CPU-FPGA Systems		ADMS		☑
J1: FPGA vs. SIMD: Comparison for Main Memory-based Fast Column Scan.			CCIS published by Springer	☑
C2: High-Throughput BitPacking Compression	DSD			☑
C3: Hardware-Software Co-Design Architecture for Lightweight Compression Algorithms.	VLDB			●
J2: An Overview of Hardware-Software Co-Design Architecture for Lightweight Compression Algorithms.			VLDBJ	●

Planned ● Accepted ☑

Doctoral Courses Plan

Dresden Database Systems Group

Courses	Place/Organized by	ECTS	General/ Project course	Status
Foreign Language (German-1)	TUD	2.5	General	☑
Transactional Information System	TUD	6	General	☑
Writing and Reviewing of Scientific Papers	TUD	1	General	☑
Introduction to the PhD Study	AAU	1	General	☑
Foreign Language (German-2)	TUD	1.5	General	⊙
Big Data Management on Modern Hardware	AAU	2	Project	☑
Deep Memory Technology for Modern Database Systems	AAU	2	Project	☑
eBISS 2017 (Participate as a Presenter)	IT4BI-DC, Brussels, Belgium	2	Project	☑
eBISS 2018 (Participate as a External Presenter)	IT4BI-DC, Netherland	1	Project	☑
eBISS 2019: IT4BI-DC Doctoral Colloquium	IT4BI-DC, Berlin, Germany	3	Project	☑
VLDB Conference, 2017 (Participate)	Munich, Germany	1	Project	☑
SPP PhD Workshop, 2018	Ilmenau, Germany	1	Project	☑
DATA Conference, 2018 (Participate as a Paper Presenter)	Porto, Portugal	1	Project	☑
ADMS Workshop, 2018	Brazil	1	Project	☑
DB Retreat, TUD, 2018 (Participate as a Presenter)	Meissen, Germany	1	Project	☑
DSD Conference, 2019 (Participate as a Presenter)	Greece	1	Project	☑
Conference/Journal/Workshop/Seminer	TBD	2	Project	⊙
Total ECTS for General Courses = 12 , Total ECTS for Project Courses = 18, Grand Total = 30				

	Completed/ Attended	Planned	Total
General	10.5	1.5	12
Project	16	2	18

Completed/Attended ☑

Planned ⊙

Published Works

Dresden Database Systems Group

@ADMS'18 — Nusrat Jahan Lisa, Annett Ungethüm, Dirk Habich, Wolfgang Lehner, Tuan D. A. Nguyen, Akash Kumar. Column Scan Acceleration in Hybrid CPU-FPGA Systems. ADMS@VLDB 2018: 22-33.

@Data'18 — Nusrat Jahan Lisa, Annett Ungethüm, Dirk Habich, Tuan D. A. Nguyen, Akash Kumar, Wolfgang Lehner. Column Scan Optimization by Increasing Intra-Instruction Parallelism. DATA 2018: 344-353.

@CCIS'19 — Nusrat Jahan Lisa, Annett Ungethüm, Dirk Habich, Wolfgang Lehner, Tuan D. A. Nguyen, Akash Kumar. FPGA vs. SIMD: Comparison for Main Memory-based Fast Column Scan. CCIS 2019, published by Springer (Accepted and on Process).

@DSD'19 — Nusrat Jahan Lisa, Tuan D. A. Nguyen, Dirk Habich, Akash Kumar, Wolfgang Lehner. High-Throughput BitPacking Compression. DSD 2019 (Accepted).