High-throughput conversion of Apache Parquet files to Apache Arrow in-memory format using FPGAs

Lars van Leeuwen (l.t.j.vanleeuwen@student.tudelft.nl)

Johan Peltenburg, Jian Fang, Zaid Al-Ars, Peter Hofstee

Quantum & Computer Engineering

TU Delft

Analyzing some tweets

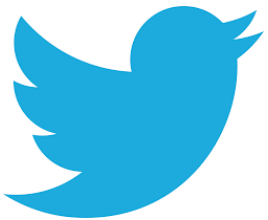

Decompression
Deserialization

Compute

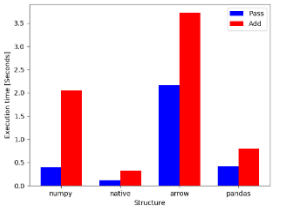

Storage → CPU → Memory → CPU → Memory

Analyzing some tweets

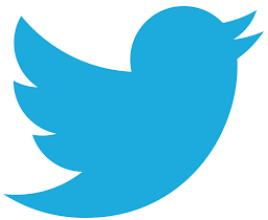

Decompression
Deserialization

Compute

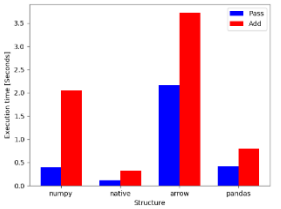

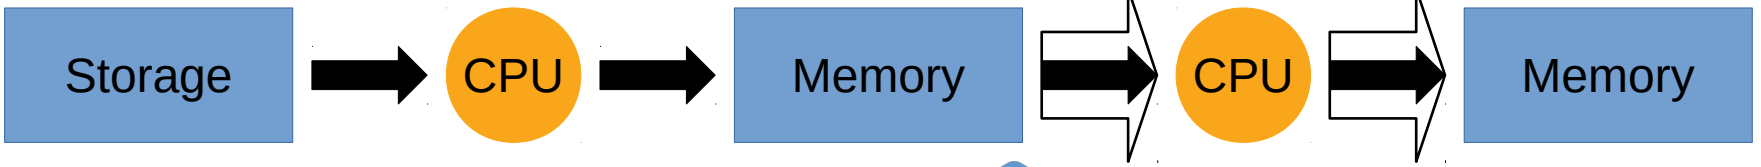

Trends in bandwidth

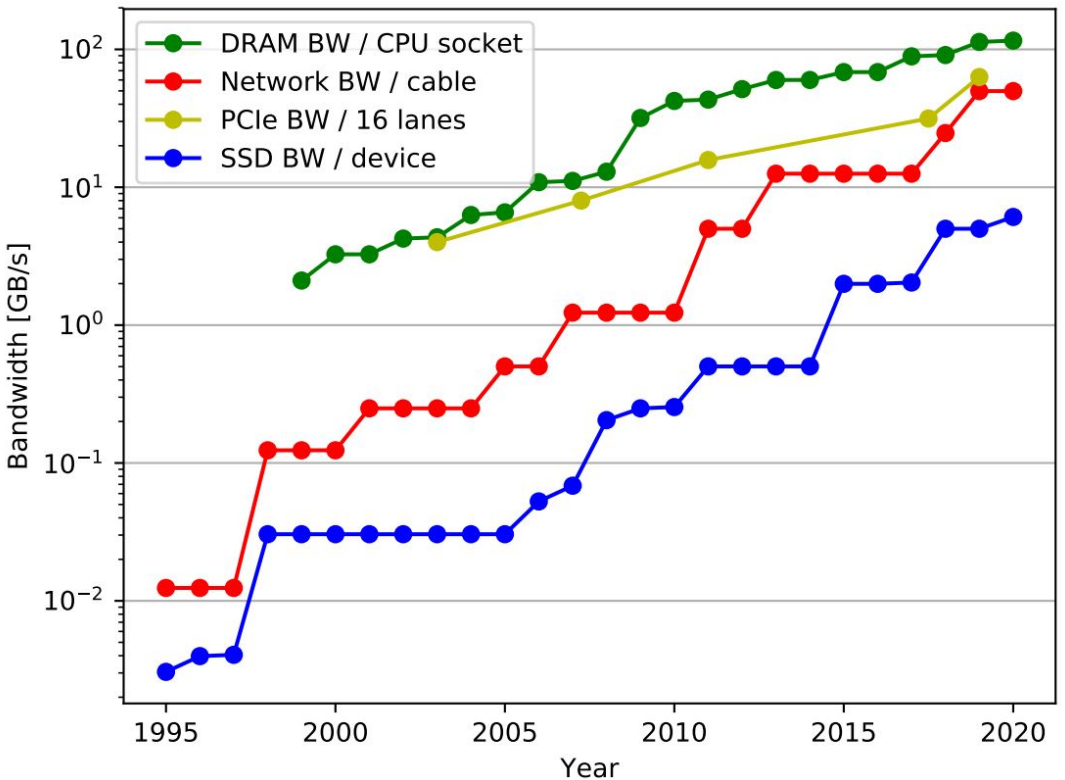

Source: The Consequences of Infinite Storage Bandwidth, SanDisk, Allen Samuels

Reading data into Pandas

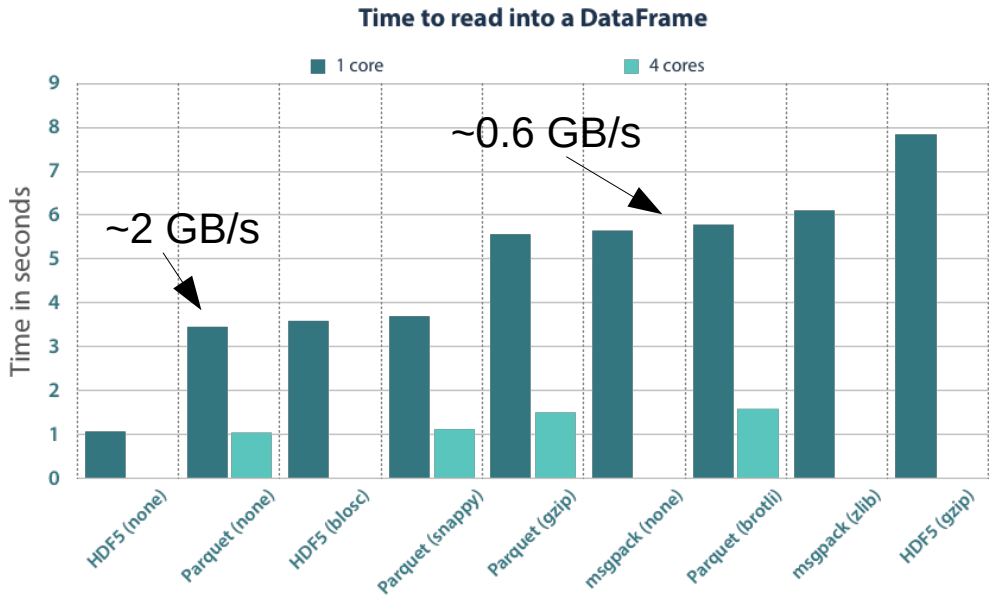

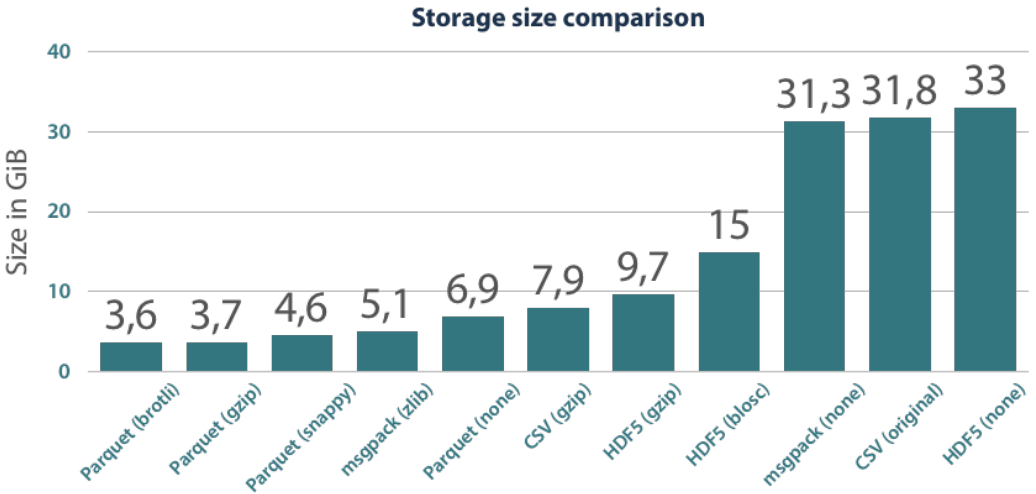

Source: Efficient DataFrame Storage with Apache Parquet, JDA Tech Blog, Uwe Korn

Analyzing some tweets

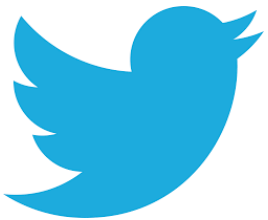

Decompression
Deserialization

Compute

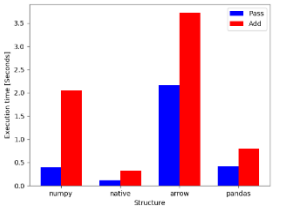

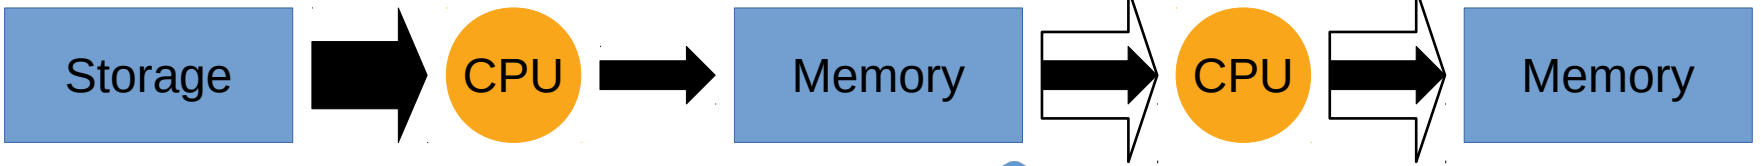

Analyzing some tweets

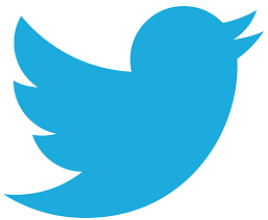

Decompression
Deserialization

Compute

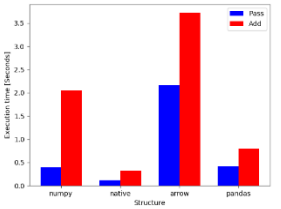

Storage → FPGA → Memory → CPU → Memory

Field Programmable Gate Array?

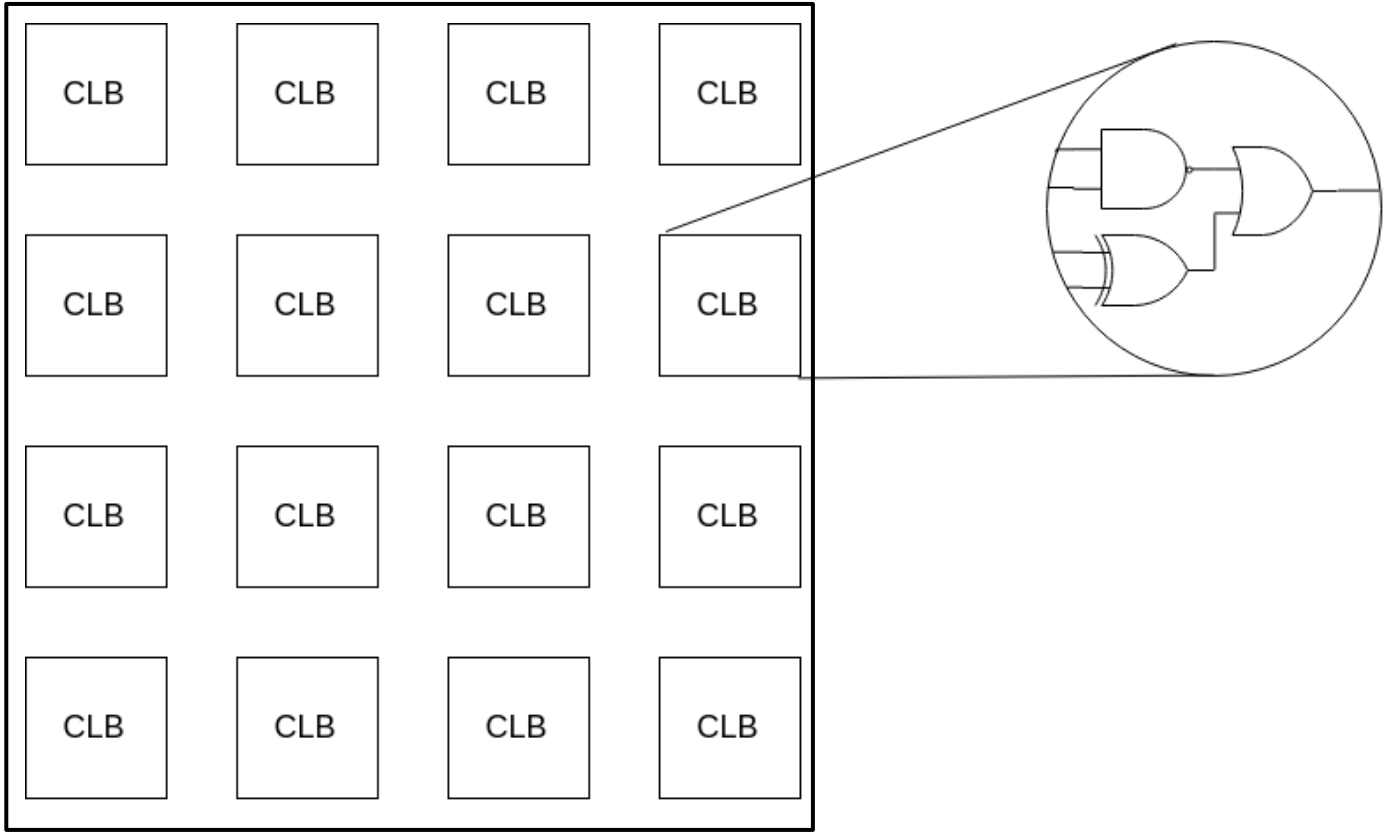

Analyzing some tweets

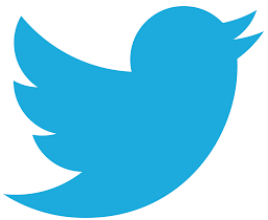

Decompression
Deserialization

Compute

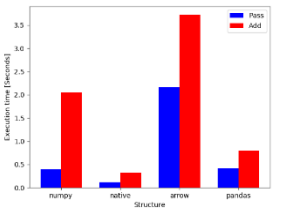

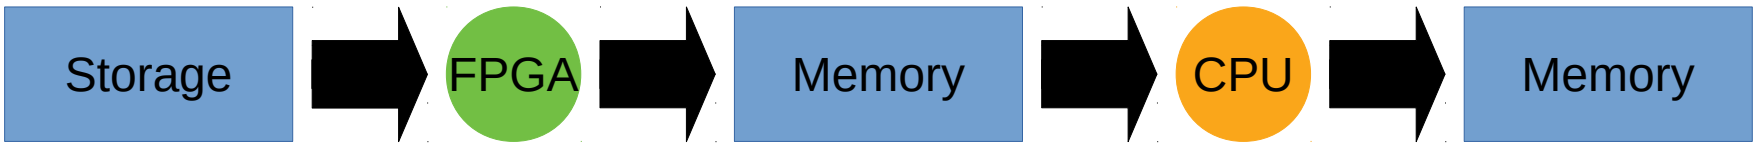

Analyzing some tweets

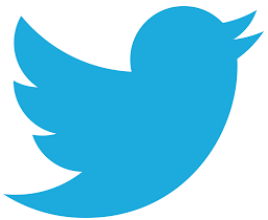

Decompression
Deserialization

Compute

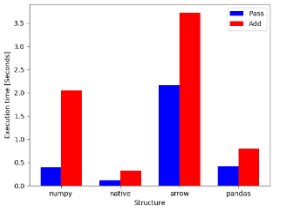

Analyzing some tweets

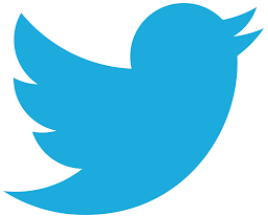

Decompression
Deserialization
Compute

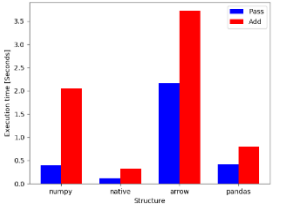

Scoping the project

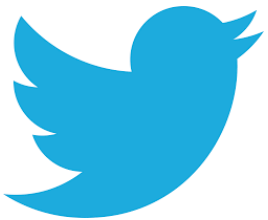

Decompression
Deserialization

Compute

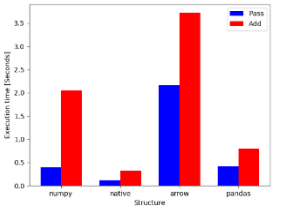

Storage → FPGA → Memory → CPU → Memory

Scoping the project

1) Both Parquet and Arrow are columnar formats
2) Parquet files can be split into parts
3) Arrow data is accessible to multiple language runtimes

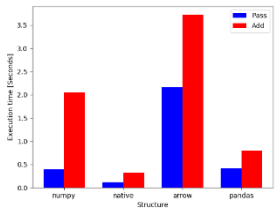

Decompression
Deserialization

Compute

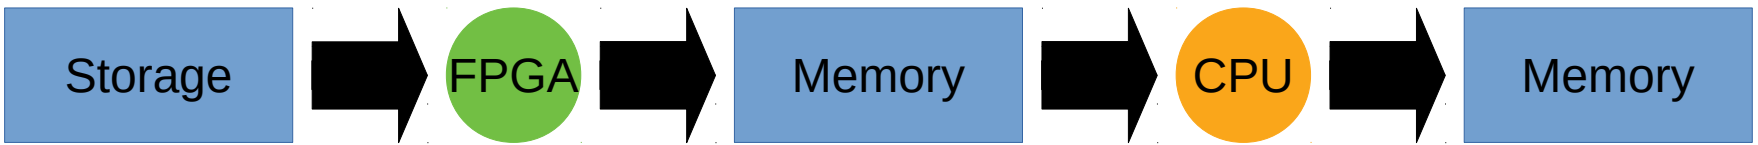

Apache Parquet

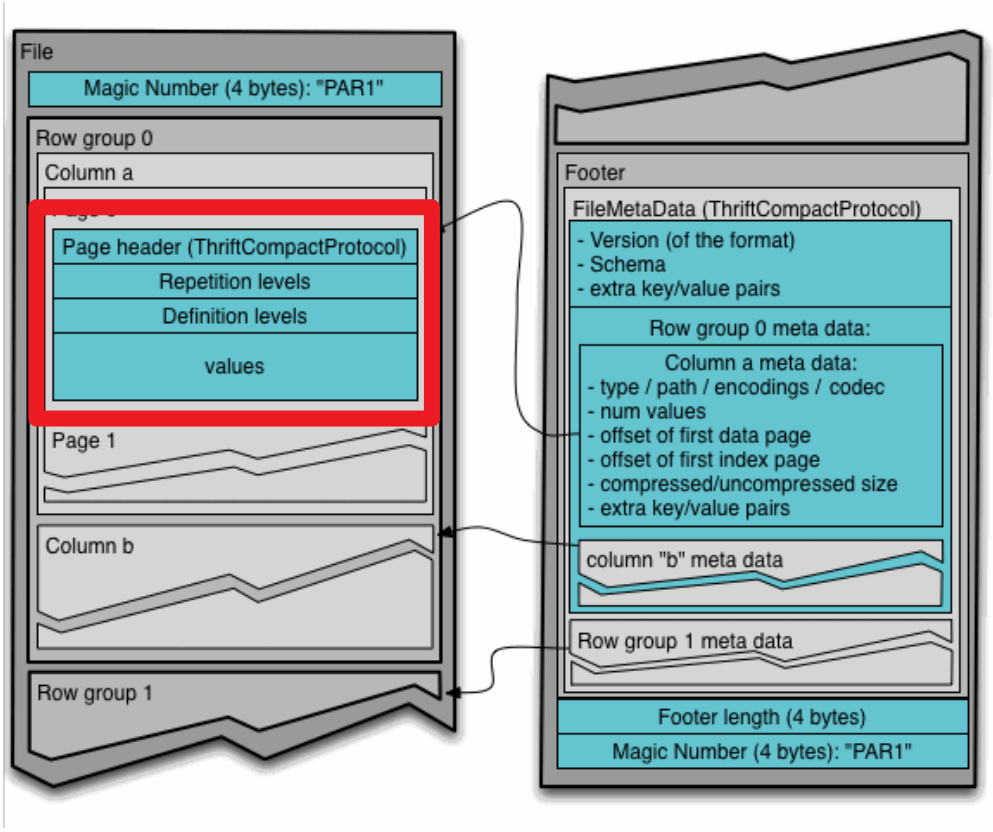

Source: parquet.apache.org

High-level architecture

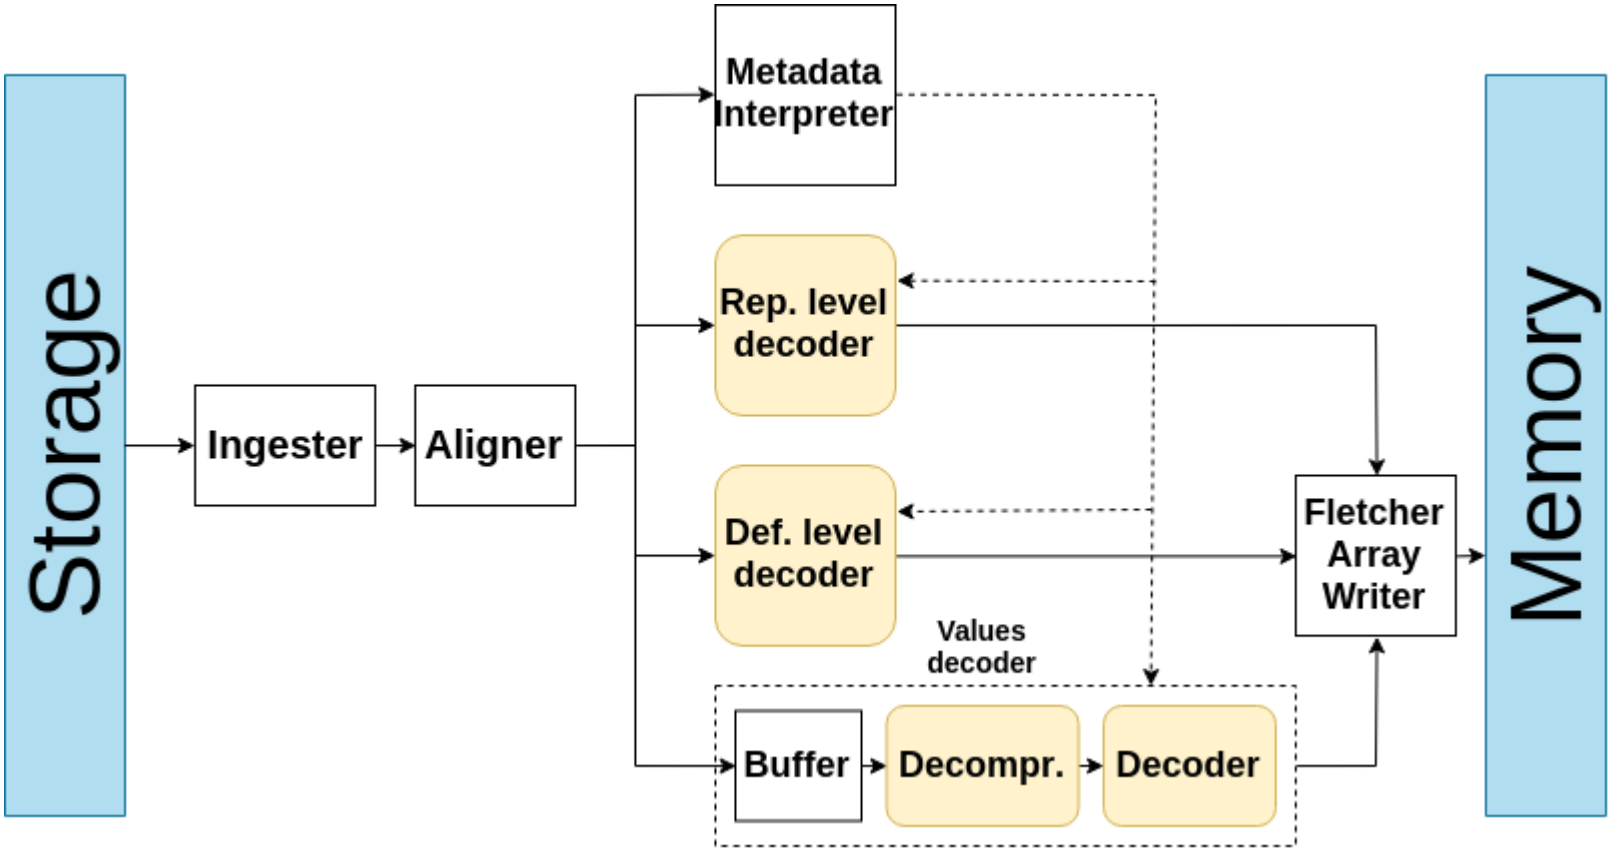

Performance for plain encoding (uncompressed)

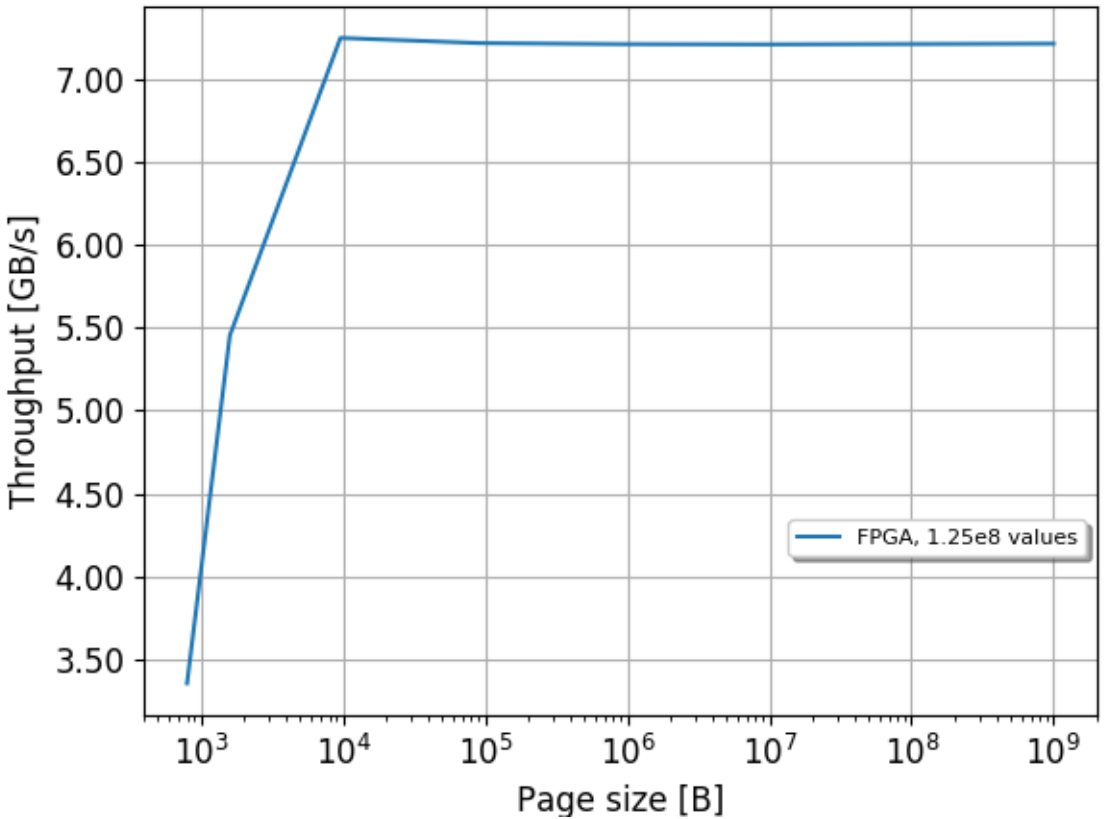

Single ParquetReader
Implemented on XCVU9P

- Area utilization
  - CLB 3.14%
  - BRAM 1.78%
- Estimated power usage
  - ~1 W

Giving the FPGA something to do

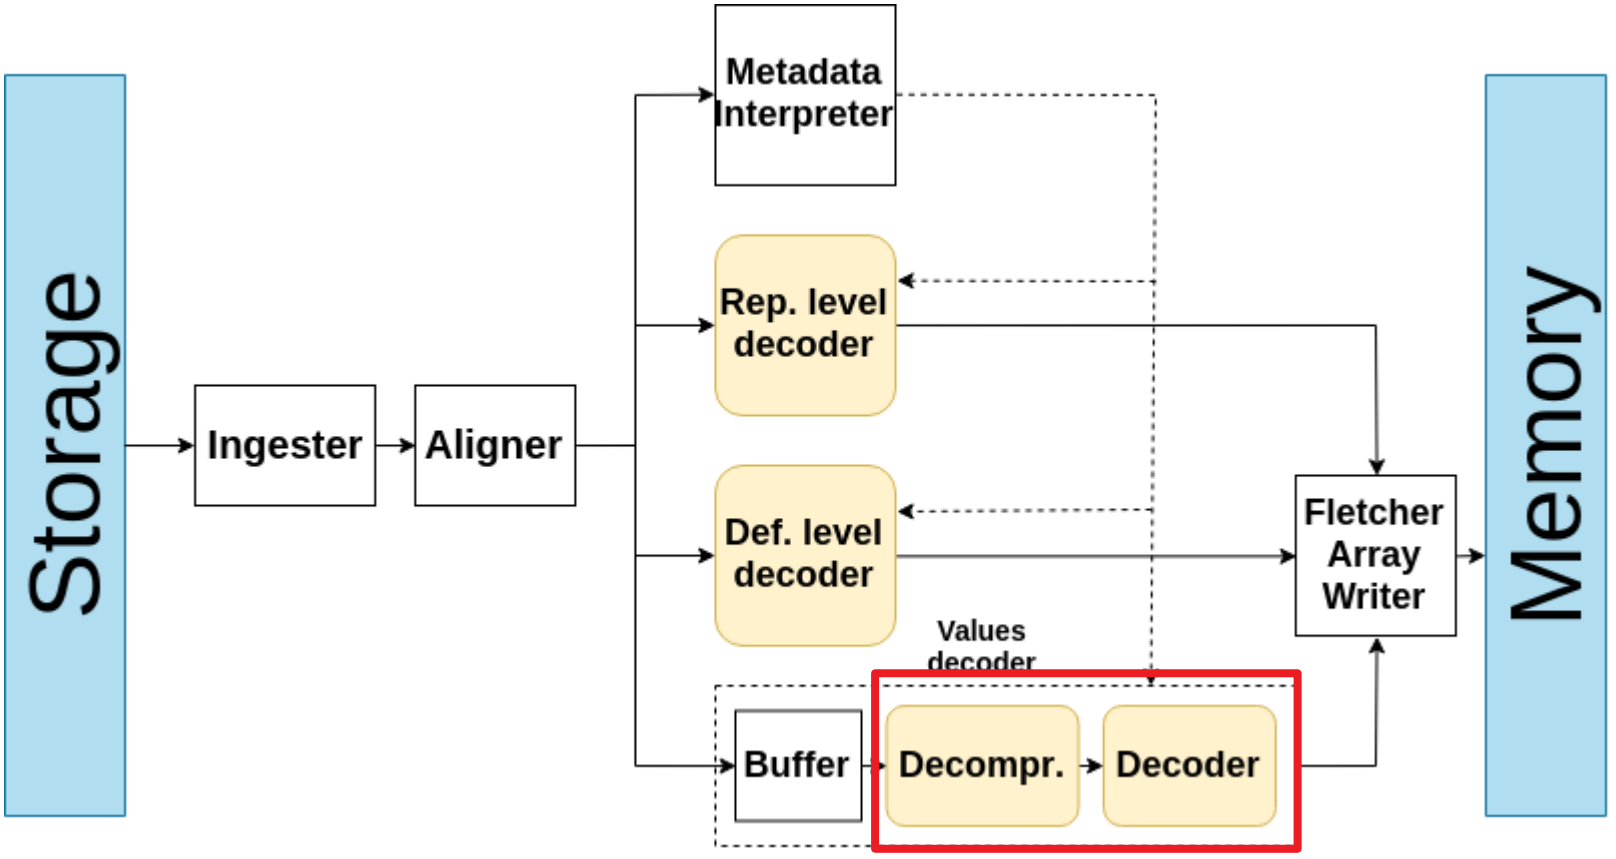

Delta decoding

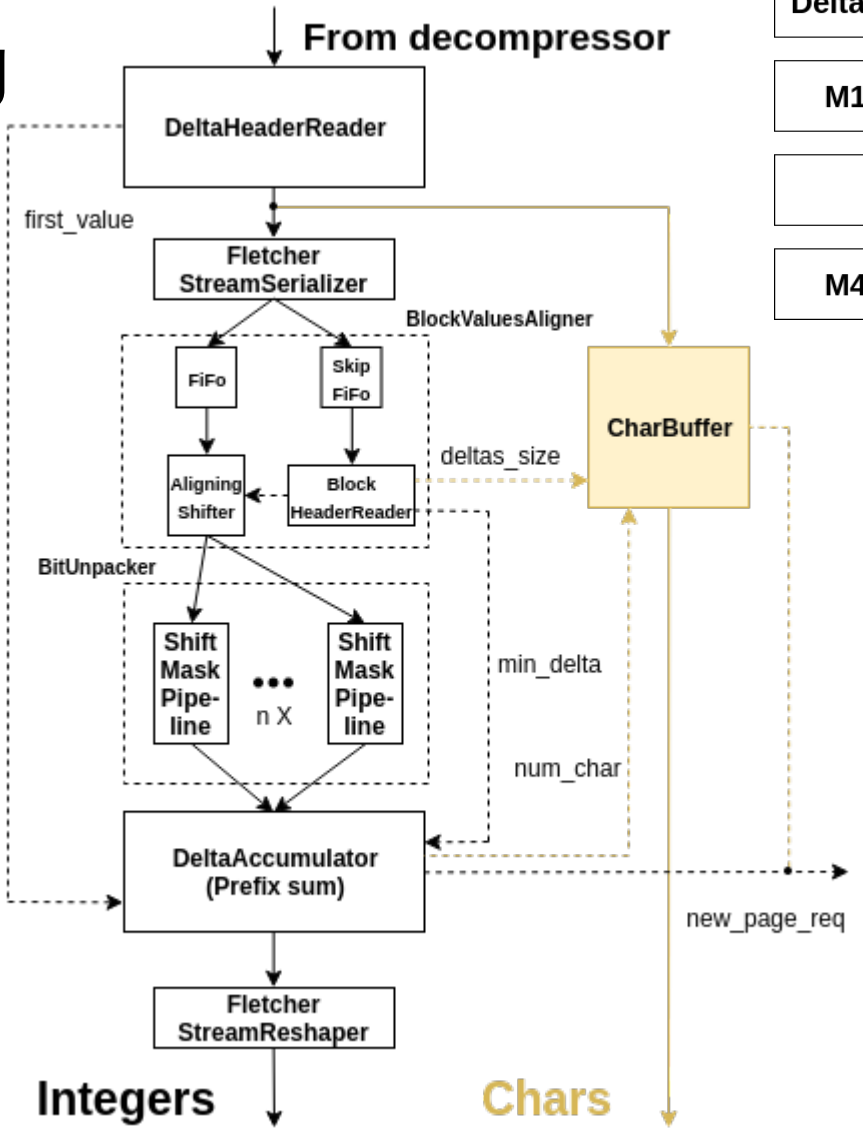

DeltaHeader	BlockHeader	M1

M1	M2	M3

M3	

M4	BlockHeader	M1

Performance with delta decoding

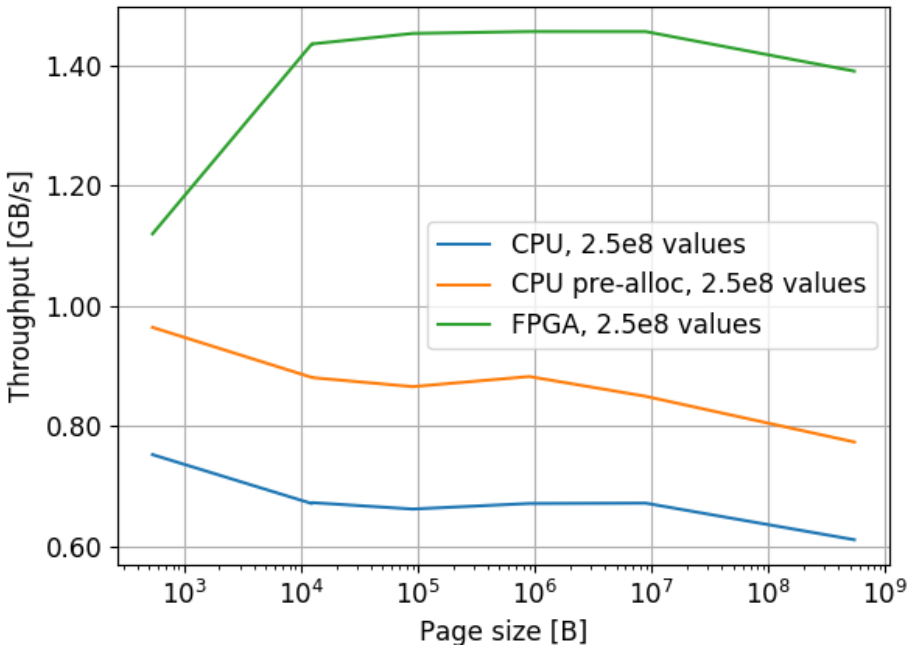

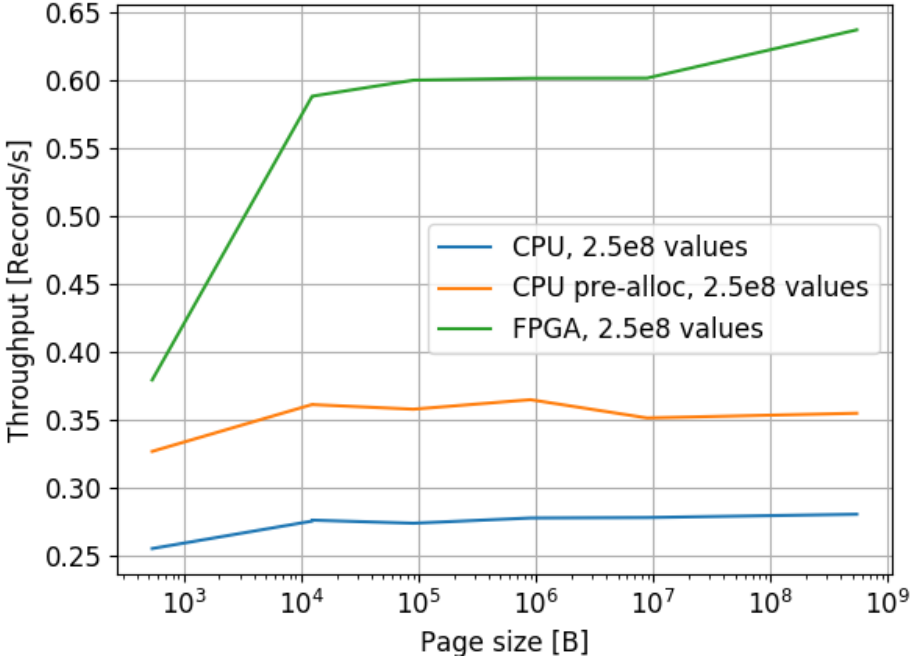

Single ParquetReader implemented on XCVU9P

- Area utilization
  - CLB 3.20%
  - BRAM 2.85%
- Estimated power usage
  - ~1 W

The future

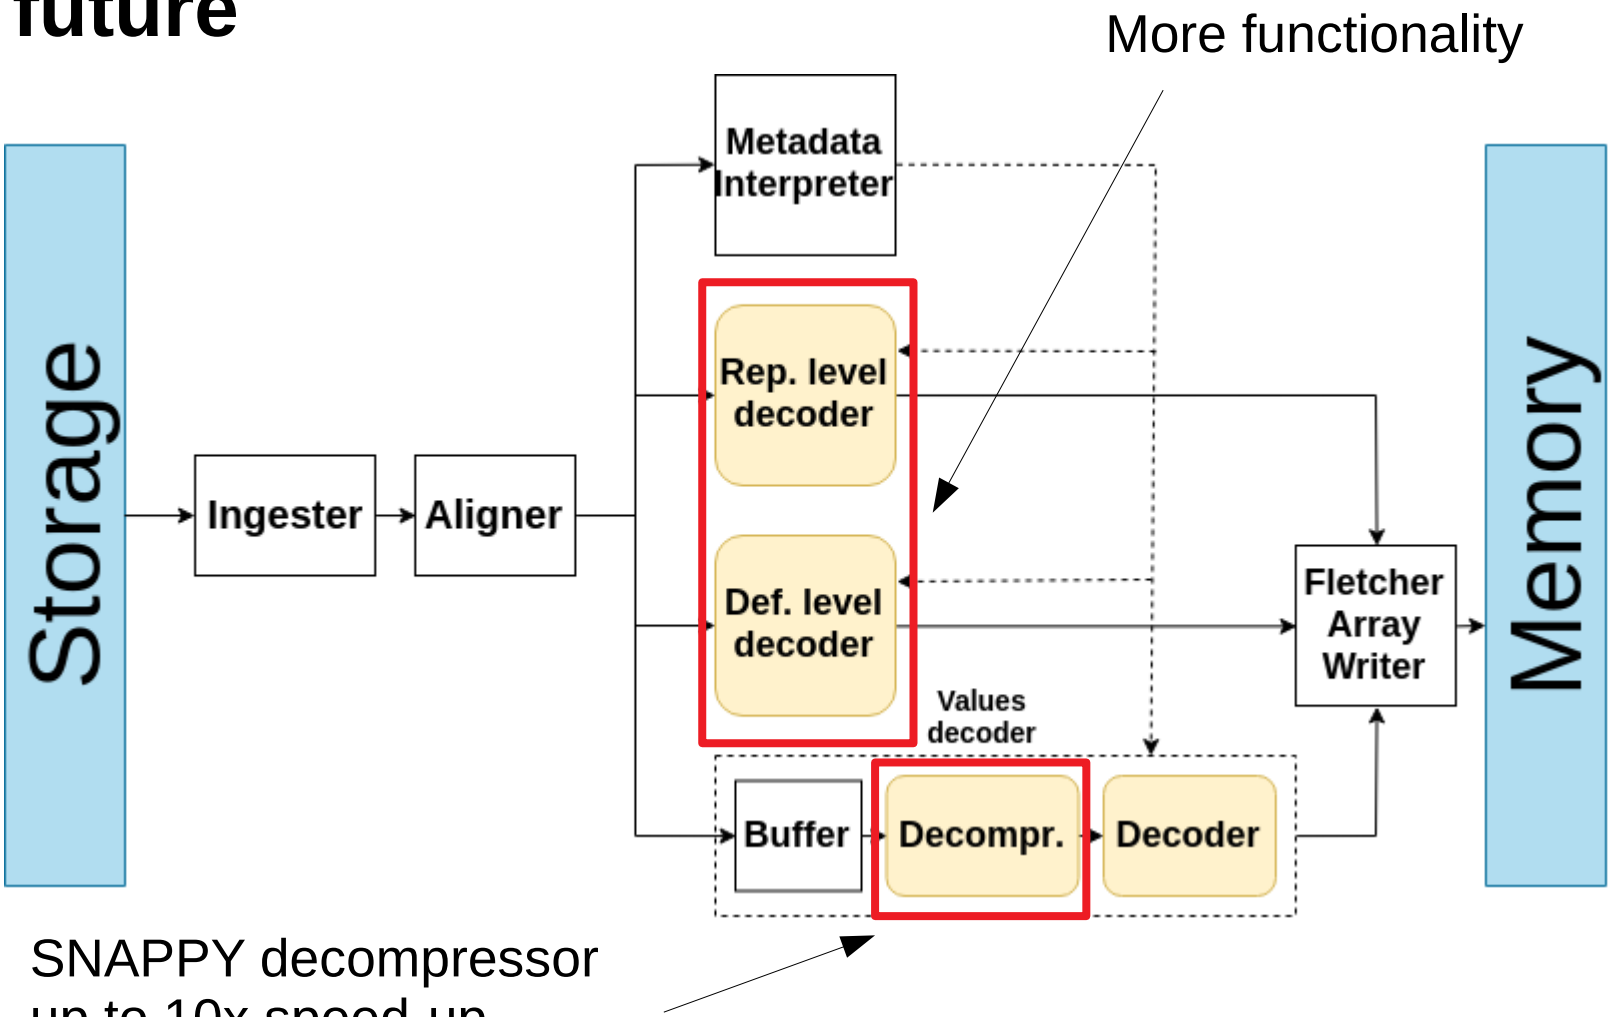
More functionality
Storage
Ingester
Aligner
Metadata Interpreter
Rep. level decoder
Def. level decoder
Values decoder
Buffer
Decompr.
Decoder
Fletcher Array Writer
Memory
SNAPPY decompressor
up to 10x speed-up

Conclusion

- 1.75x speed up shown for delta encoding
- Single ParquetReader uses only 3.2% of CLBs
- Promising future with decompression hardware

High-throughput conversion of Apache Parquet files to Apache Arrow in-memory format using FPGAs

Lars van Leeuwen (l.t.j.vanleeuwen@student.tudelft.nl)

Johan Peltenburg, Jian Fang, Zaid Al-Ars, Peter Hofstee

Quantum & Computer Engineering