Trinity: End-to-End In-Database Near-Data Machine Learning Acceleration Platform for Advanced Data Analytics

Ji-Hoon Kim[1)], **Seunghee Han**[1)],
Kwanghyun Park[2)], **Soo-Young Ji**[3)] **and Joo-Young Kim**[1)]

1) KAIST CastLab 2) Microsoft 3) SAMSUNG

In-DBMS Data Analytics

- **Three Important yet Independent Technology Trends**

ML-based Advanced Data Analytics

Google BigQuery ML

Amazon Redshift

MADlib

Spark MLlib

- Enterprise-level DBMS
- In-DBMS ML support

Database HW Acceleration

blazingSQL

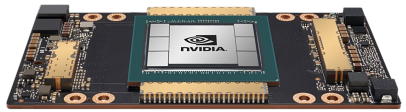

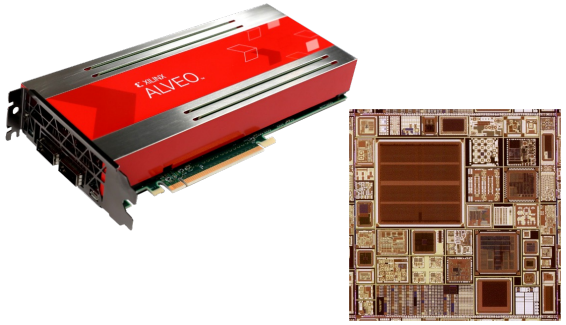

- GPU-based DBMS
- ASIC/FPGA/GPU

Near-Data/In-Storage Processing

- Data-intensive application
- SmartSSD/SmartNIC/HMC

Trinity: In-Database, In-Storage Platform

- **Full Stack System for Near-Data-based In-DBMS ML Inference**

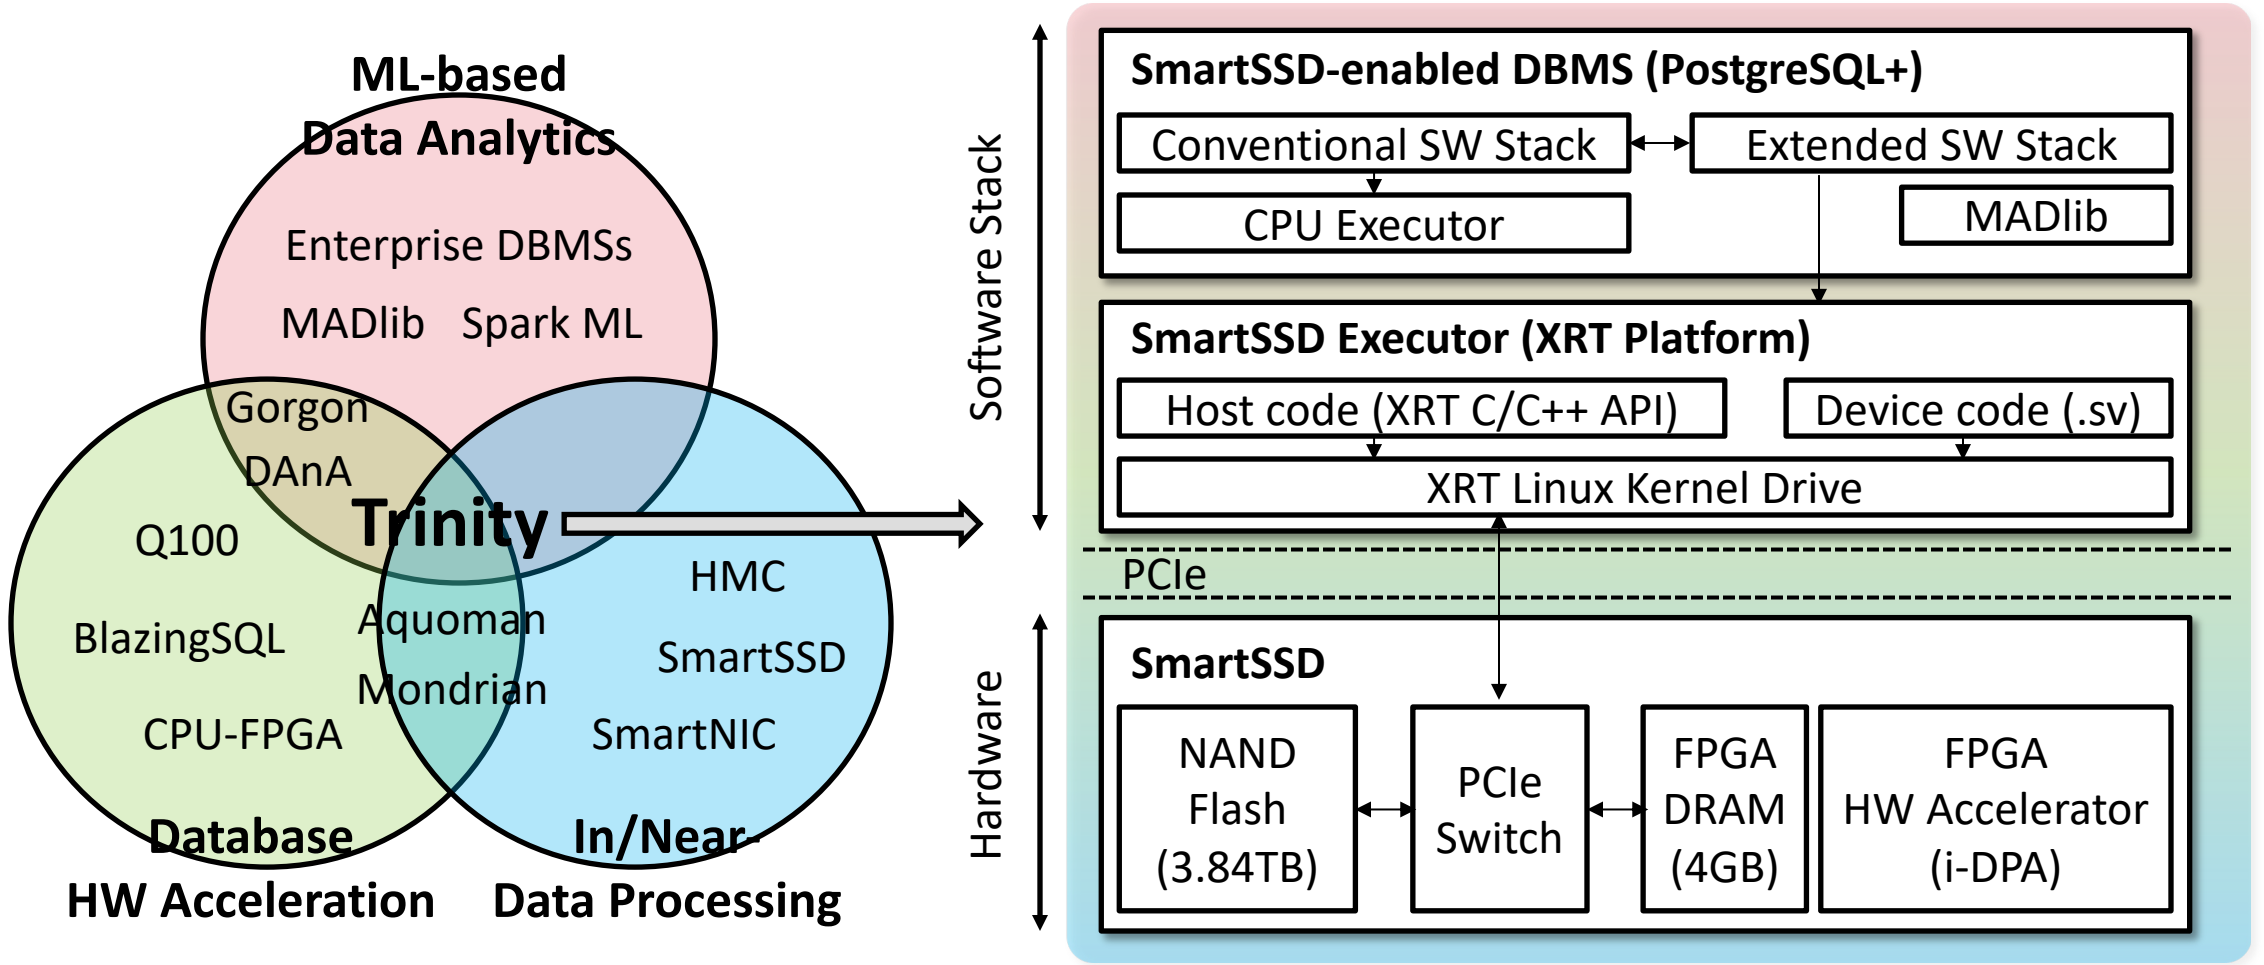

Trinity: In-Database, In-Storage Platform

- **Full Stack System for Near-Data-based In-DBMS ML Inference**

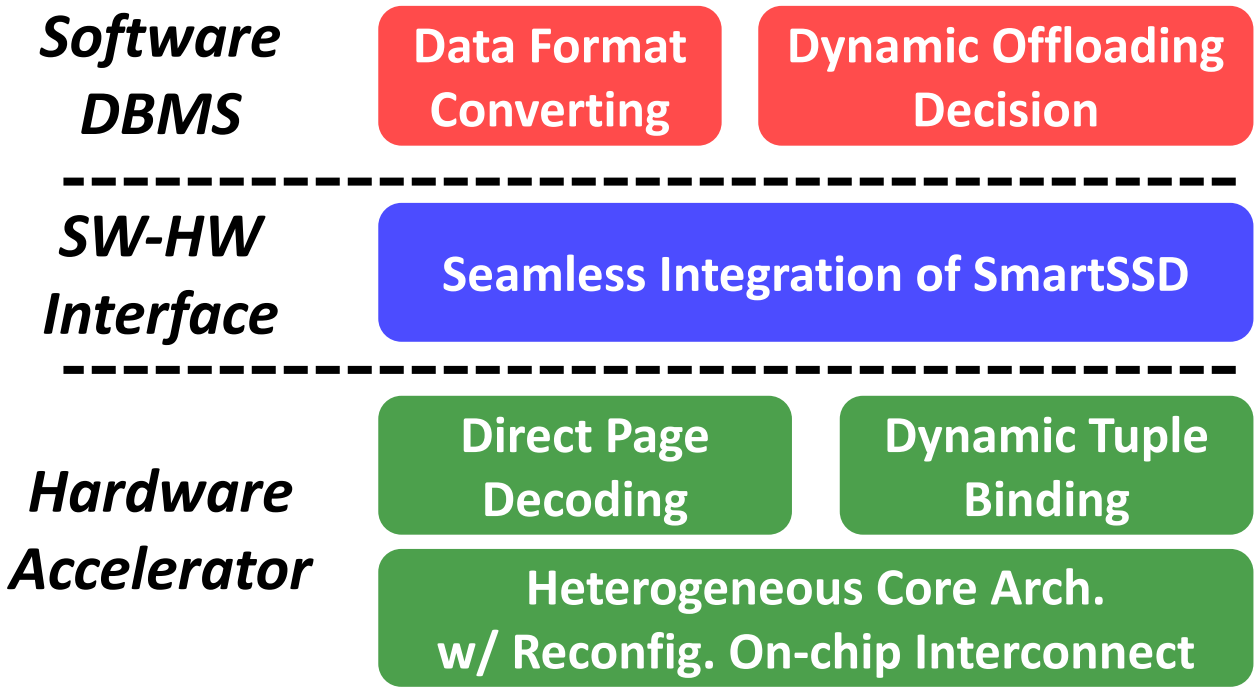

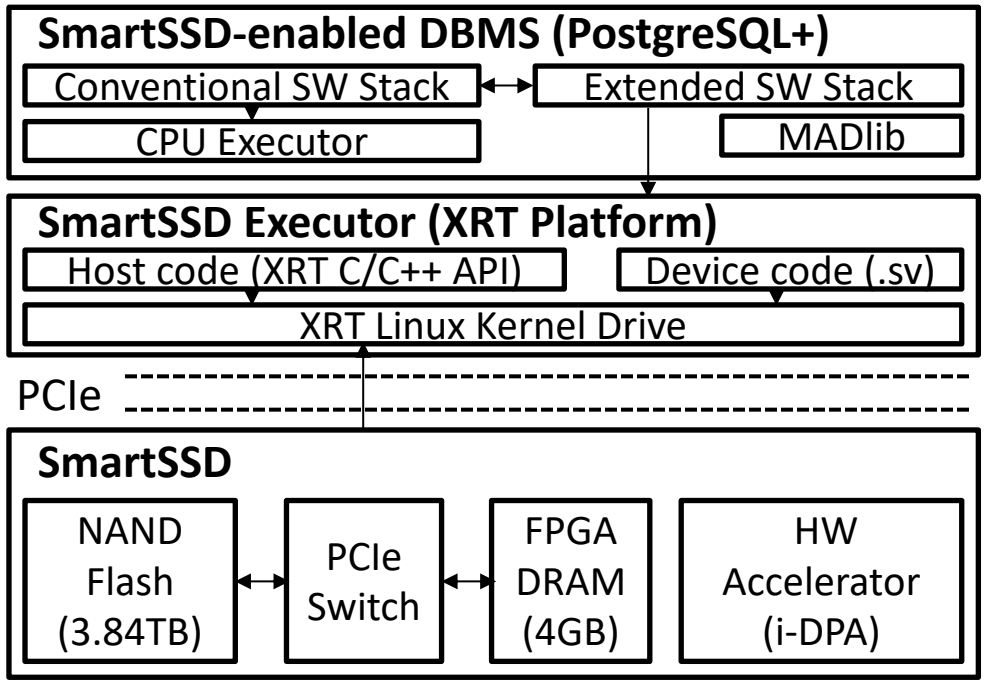

Trinity shows up to 57.18x faster query processing speed than CPU-based DBMS

Computational Storage Device

- **New Hardware Backend: <u>Samsung's SmartSSD</u>**
  - Xilinx Kintex UltraScale+ FPGA, 4GB DRAM and 3.84TB NAND flash
  - Direct FPGA-to-SSD data access using internal PCIe switch

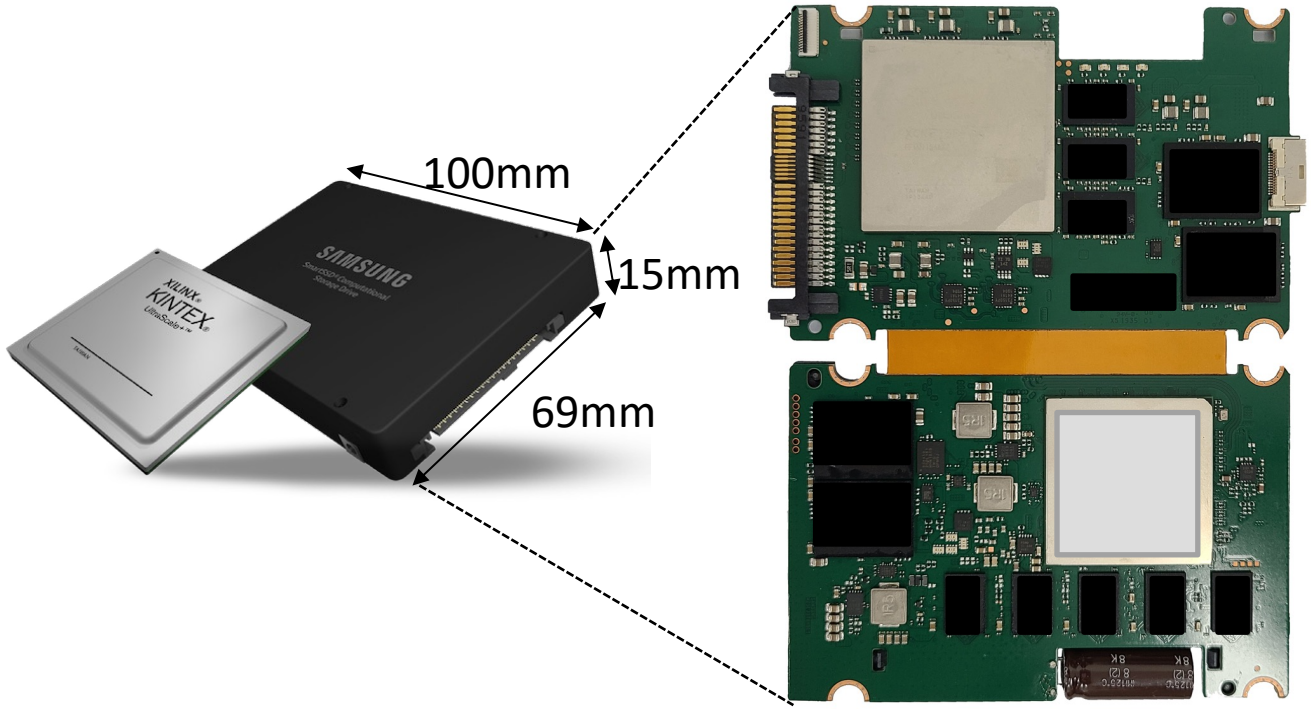

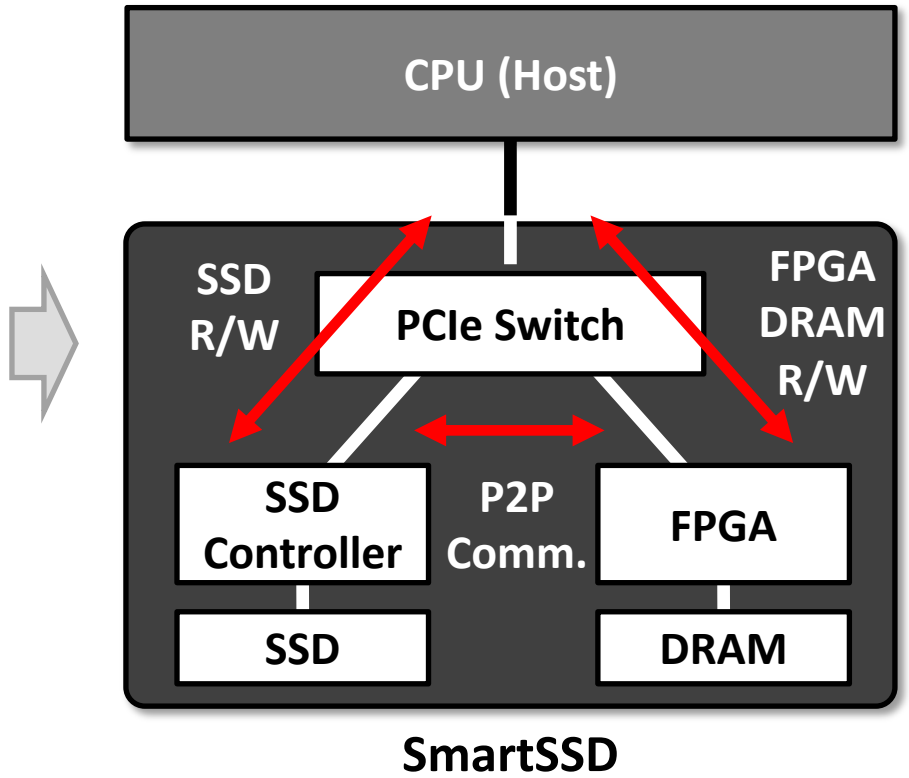

Software Stack for Trinity

- **Seamless Integration of SmartSSD in DBMS**
  - Extended SW stack ➔ extended analyzer[1)] + optimizer[2)]

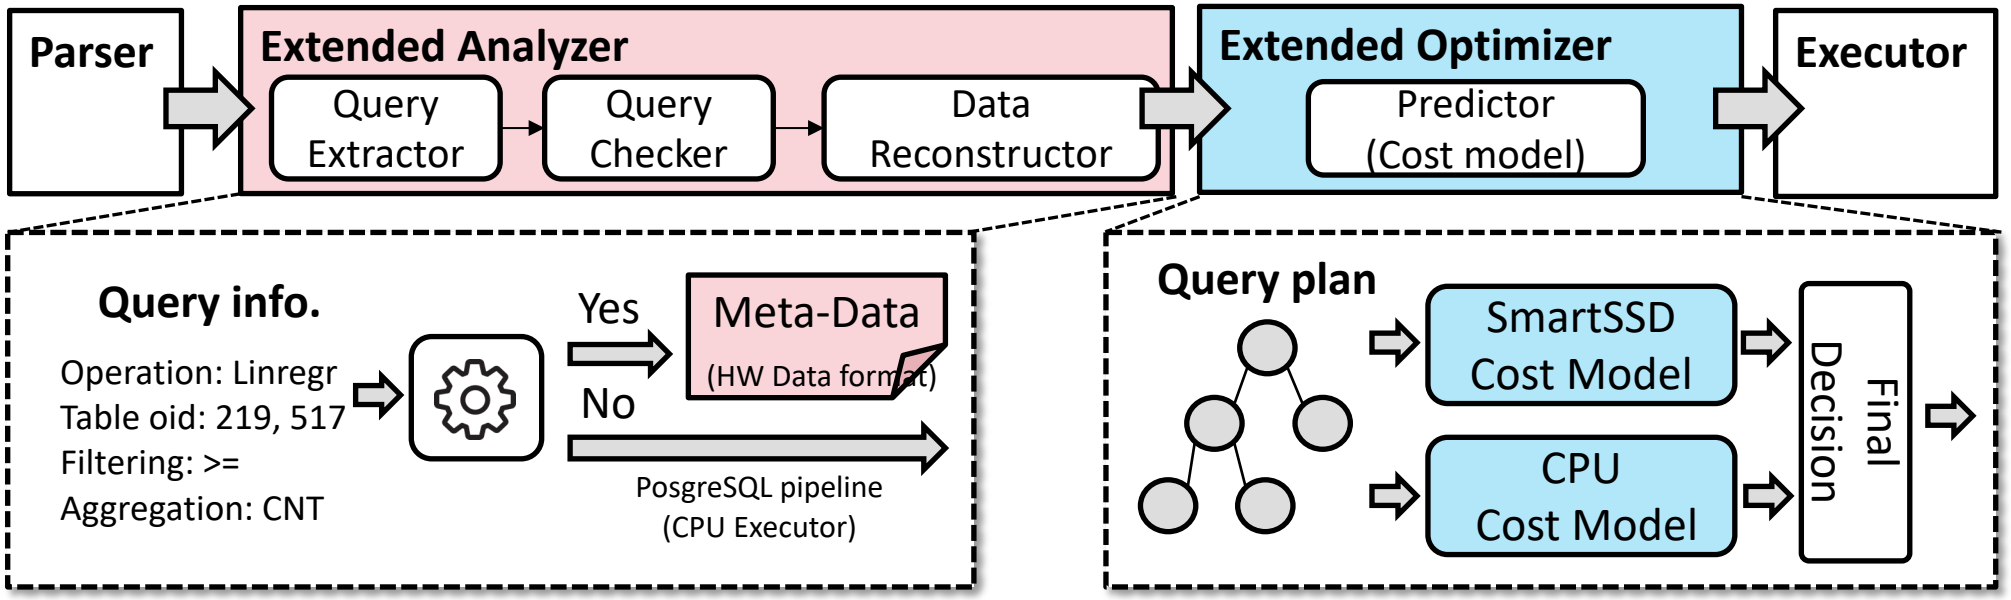

1) Converting query information to the HW data format
2) Making **runtime offloading decision** to select an optimal HW backend

Extended Query Optimizer

- **Performance Cost Model**
  - Determining optimal hardware backend (SmartSSD vs CPU)
  - Showing **5.3%** & **12.97%** average error + **96% offloading accuracy**

SmartSSD Cost Model

Equation-based performance model

$$T_{smartssd} = T_{host} + T_{trasnfer} + T_{kernel}$$

$$T_{buffer} + T_{others} \qquad T_{ssd2fpga} + T_{fpga2host}$$

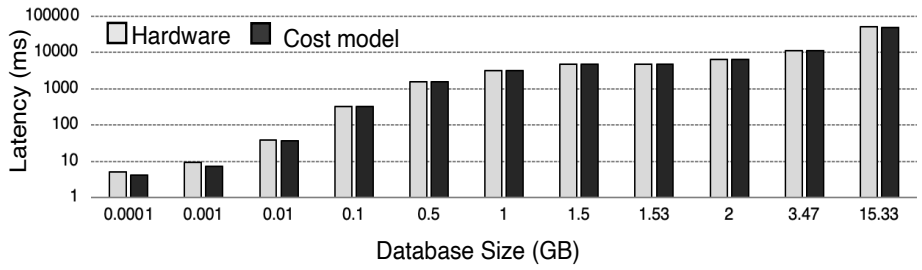

CPU Cost Model

Regression-based performance model

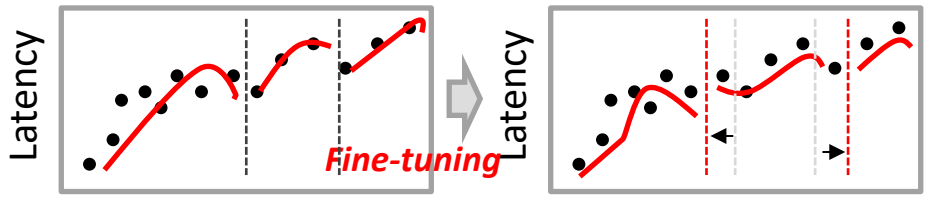

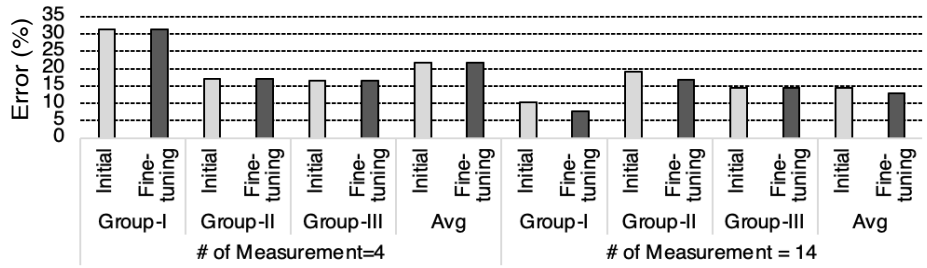

Overall Architecture of i-DPA*

1. **Database Page Decoder**
   - *Page & data processing unit*
   - *Page-level parallelism*
2. **Database Tuple Binder**
   - *Dynamic tuple binding*
   - *Tuple-level parallelism*
3. **Heterogeneous Core Arch.**
   - *Reconfig. on-chip interconnect*
   - *Task-level parallelism*

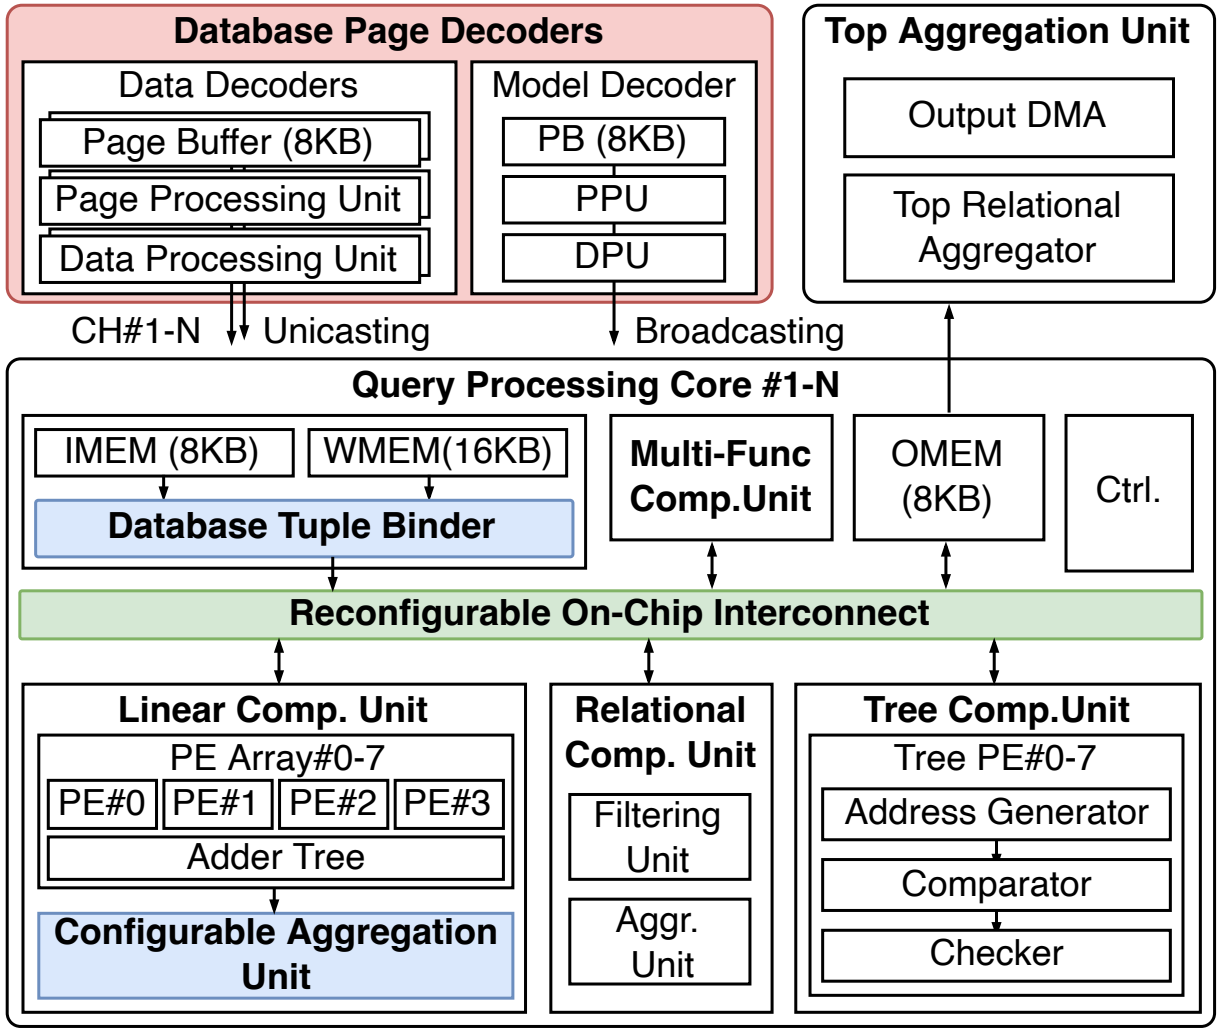

*** i-DPA = in-Database Processing Accelerator**

Database Page Decoder

- **Direct Tuple Extraction from Database Page**
  - Removing host interaction ➔ Keep BW benefit of in-storage processing
  - Faster page decoding w/ **page-level parallelism** ☺

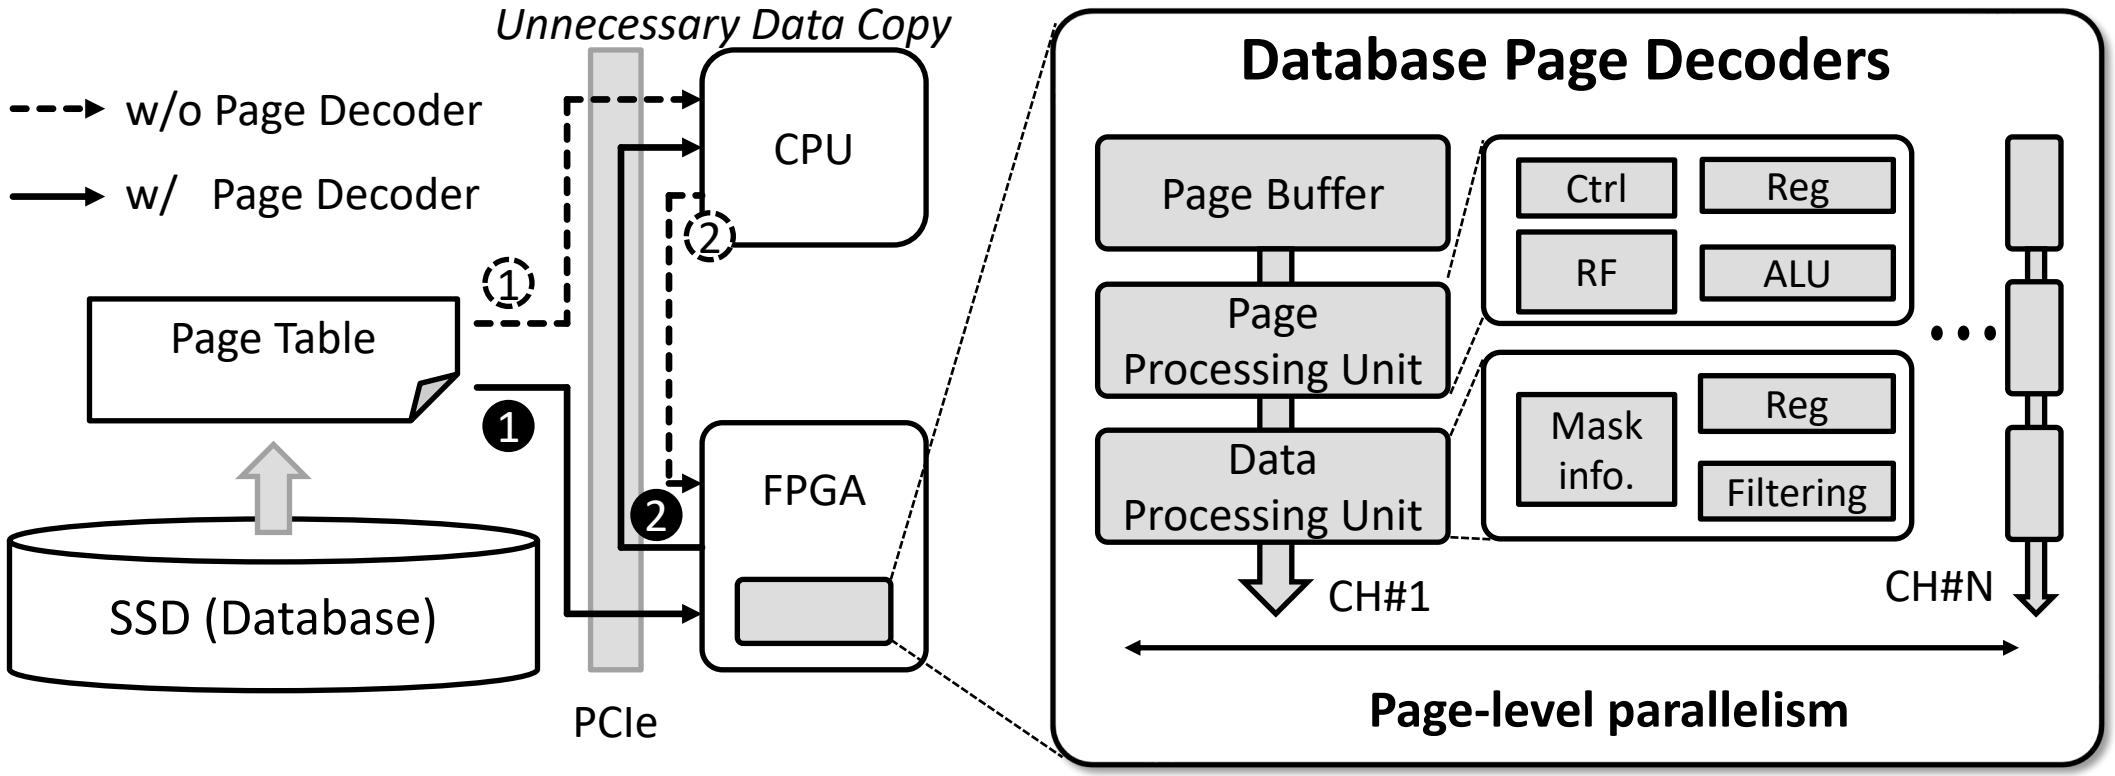

Database Tuple Binder

- **Dynamic Tuple Binding**
  - Dynamically varying tuple packing density according to the tuple size
  - **Tuple-level parallelism** ▲ & hardware utilization ▲ ☺

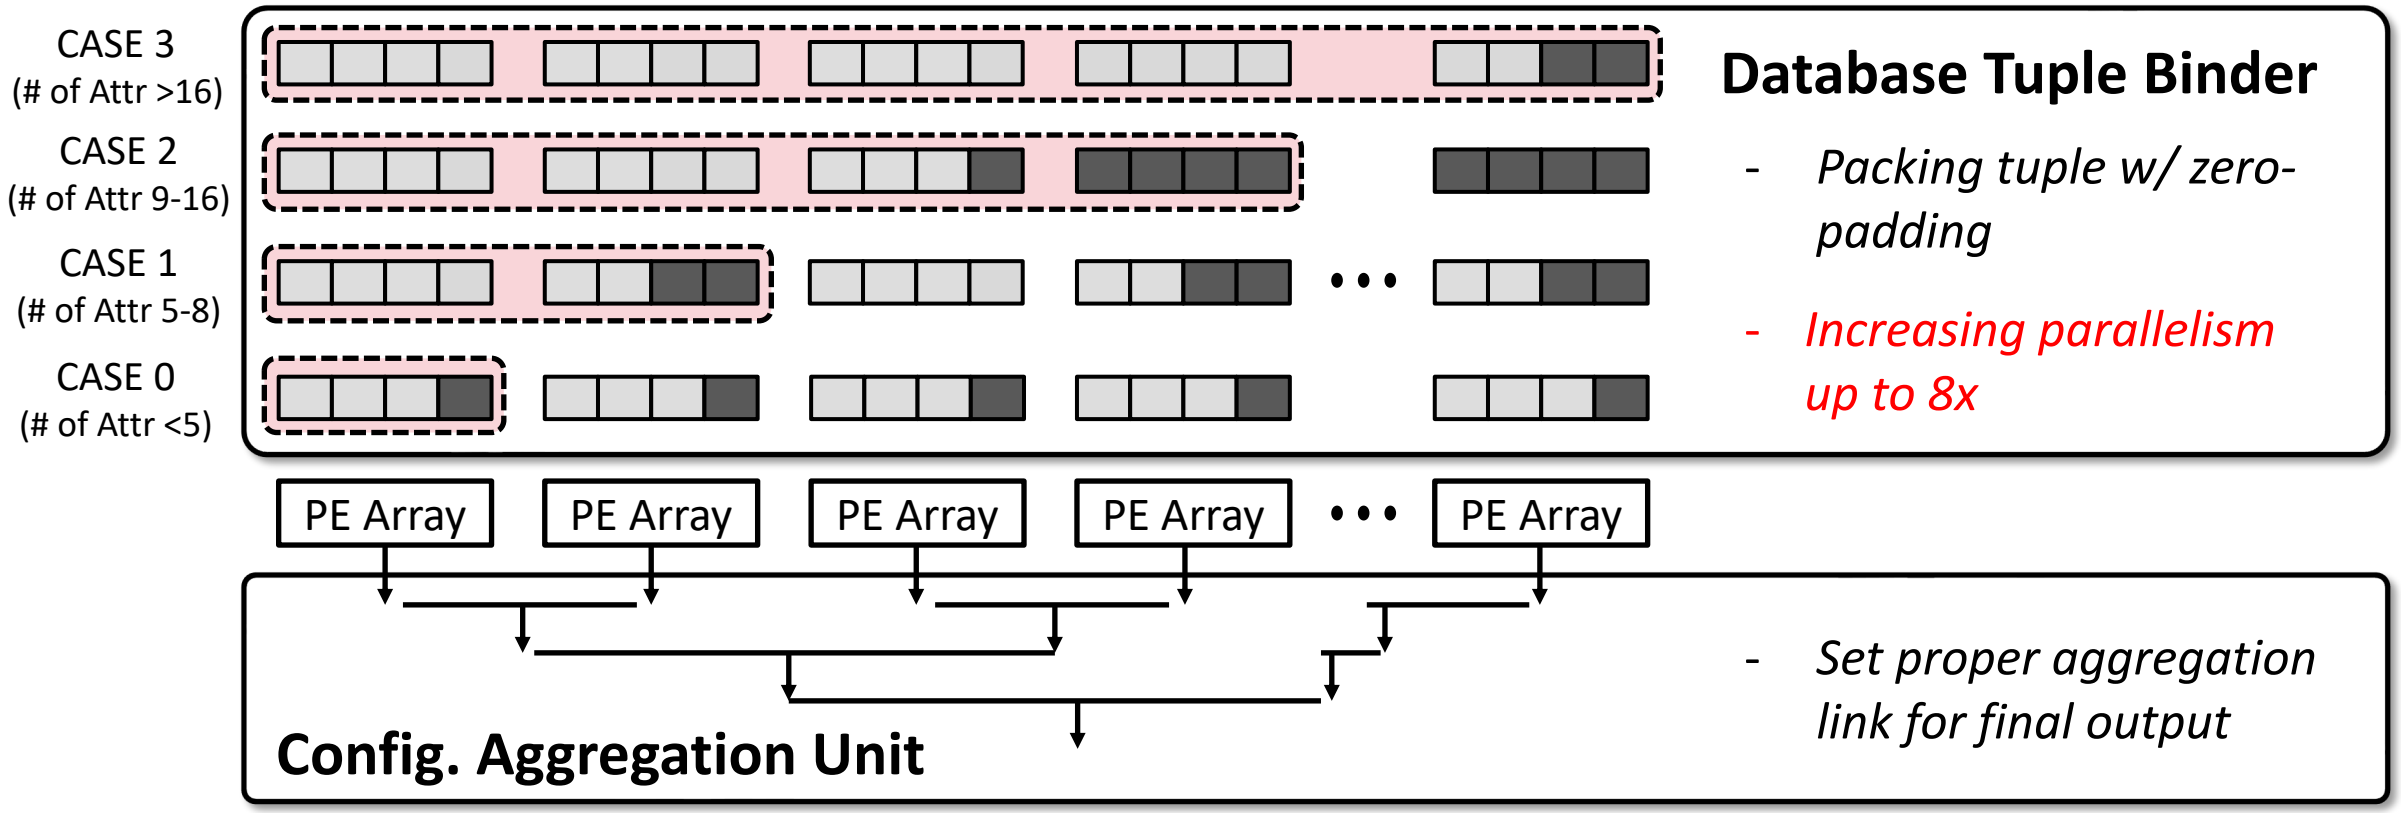

Query Processing Core

- **Heterogeneous Core Architecture**
  - Reconfigurable on-chip interconnect ➔ Enabling flexible data streaming ☺
  - **Task-level parallelism** across the computing units

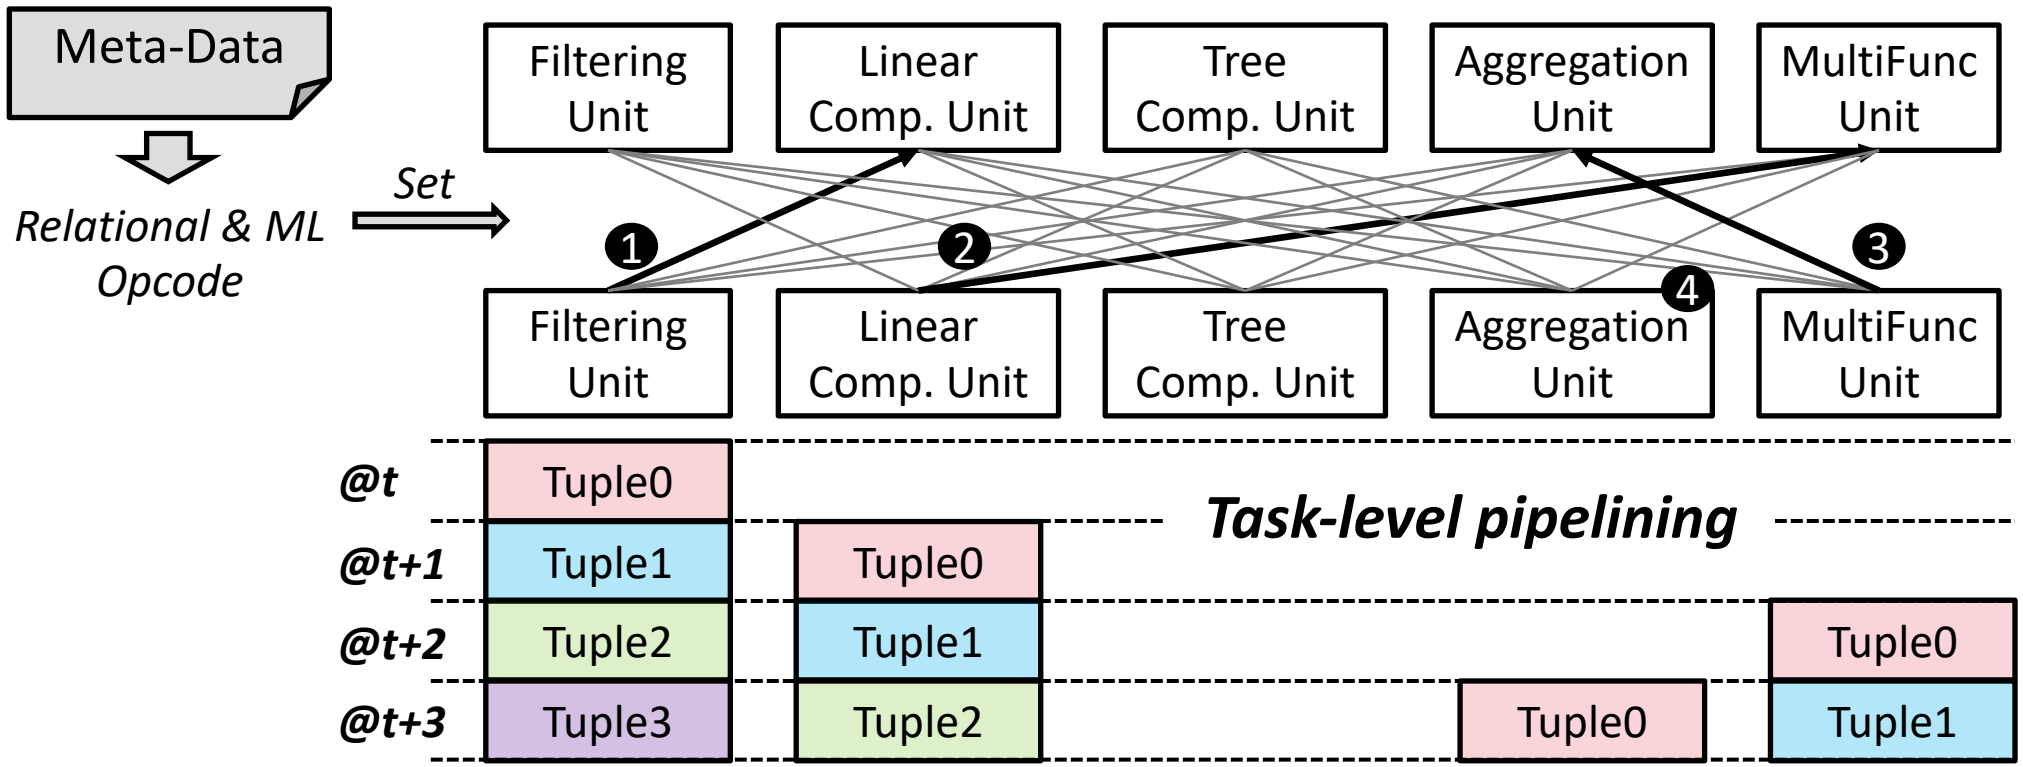

FPGA Implementation Result

- **System Setup & FPGA Implementation Result**

- 2 Intel Xeon Silver 4210 CPUs
  - PostgreSQL v12.6
  - MADlib v1.17.0
- 156GB DIMM
- 3.84TB SmartSSD

Computational Storage Server

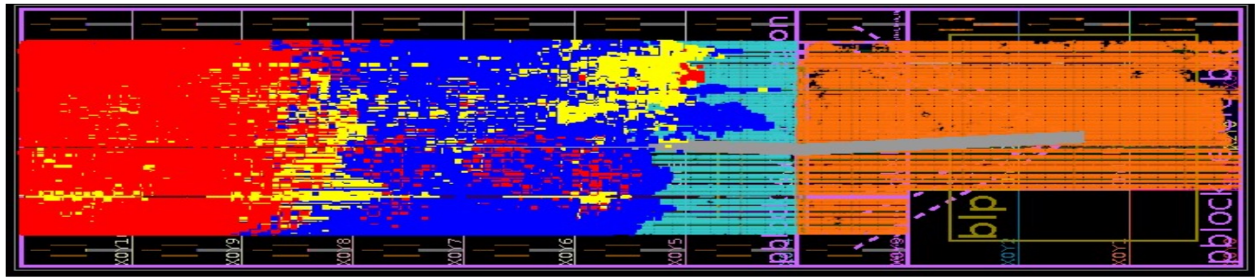

	Specifications				
FPGA	Kintex Ultrascale+				
Freq.	170MHz				
	Resource Utilization				
	LUT	**FF**	**BRAM**	**URAM**	**DSP**
Interface	128665	183077	311.5	8	9
Core0	102672	92142	15	32	342
Core1	110092	92437	15	32	342
Others	10740	9757	22	0	10
Utilization	67.4%	36.1%	46.9%	56.3%	35.7%

End-to-End Trinity Evaluation

- **Evaluate Against CPU-based DBMS Platform**
  - **0.85x – 57.18x** faster query processing than CPU-based DBMS
  - **15.21x** faster than CPU-based DBMS on average

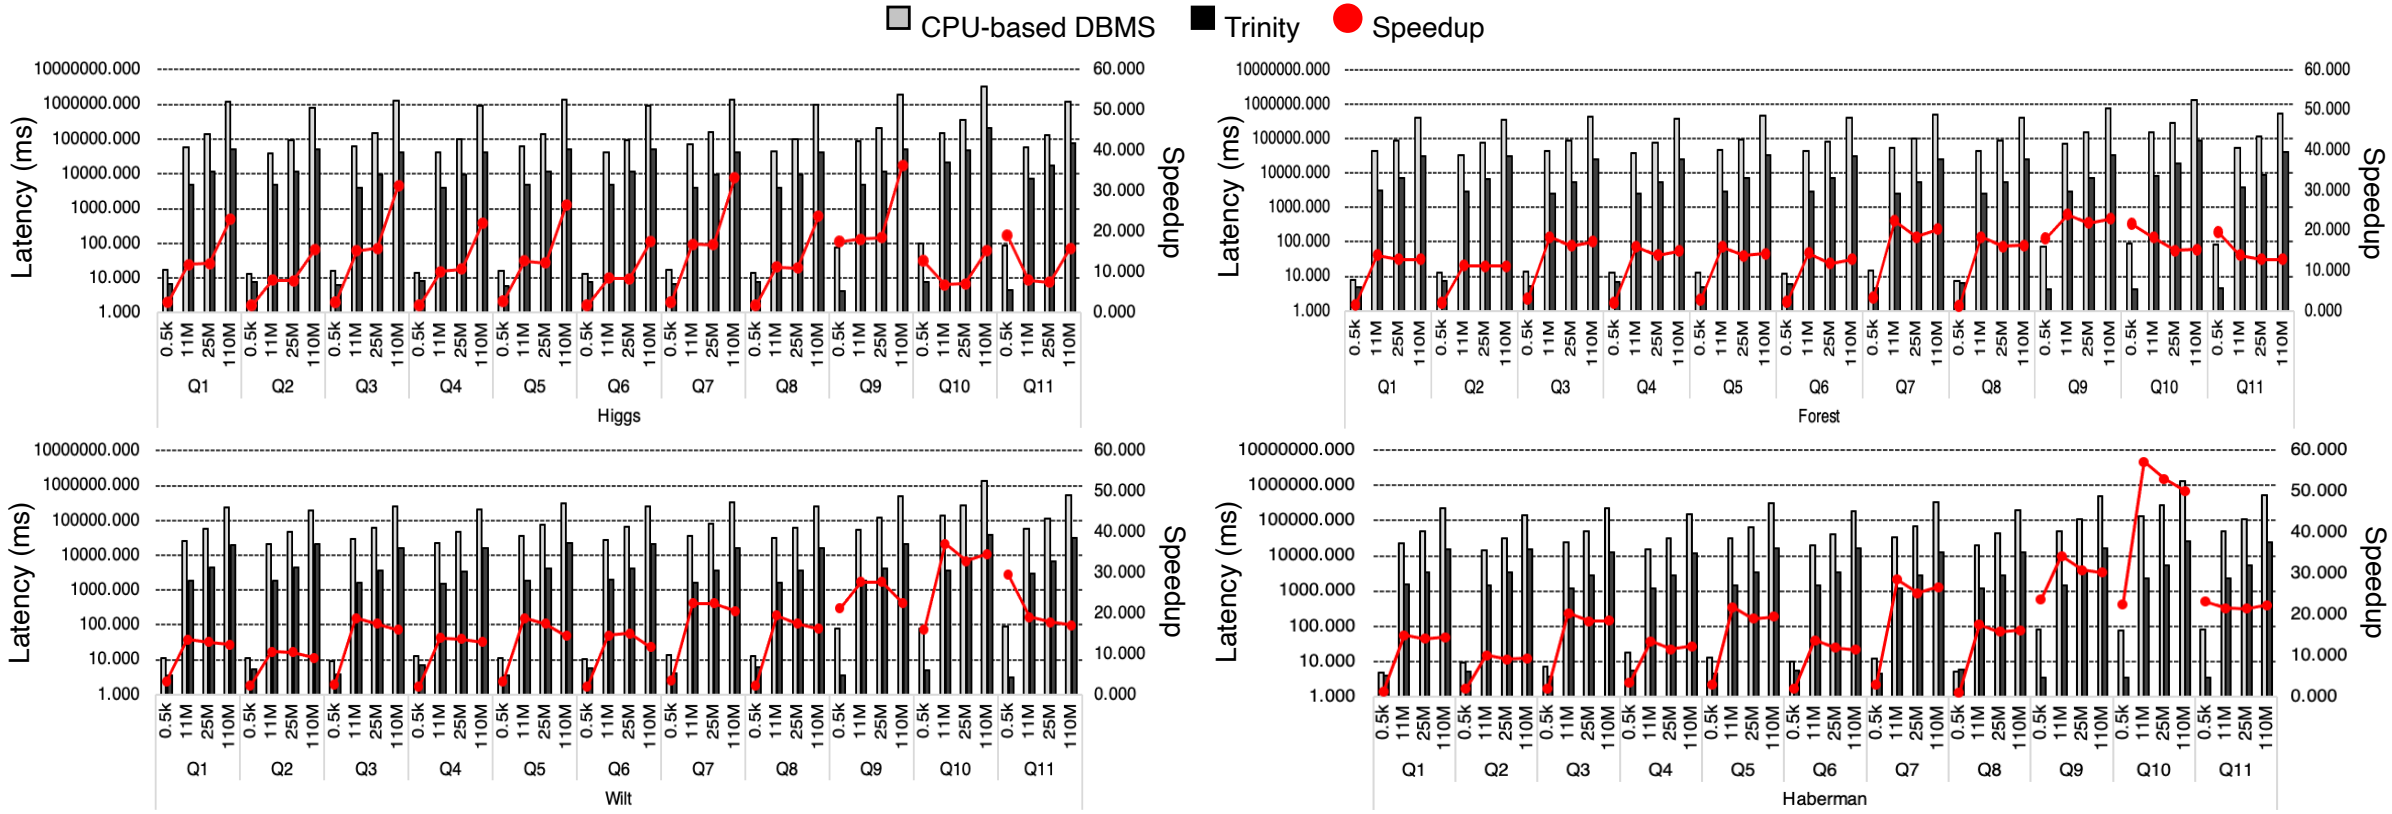

Scaling-up with Multiple SmartSSDs

- Scale-up the Overall System
  - SmartSSD: easy to scale-out the number of devices w/ U.2 form factor
  - Deploying 4 SmartSSDs ➔ **200x faster** than CPU-based DBMS

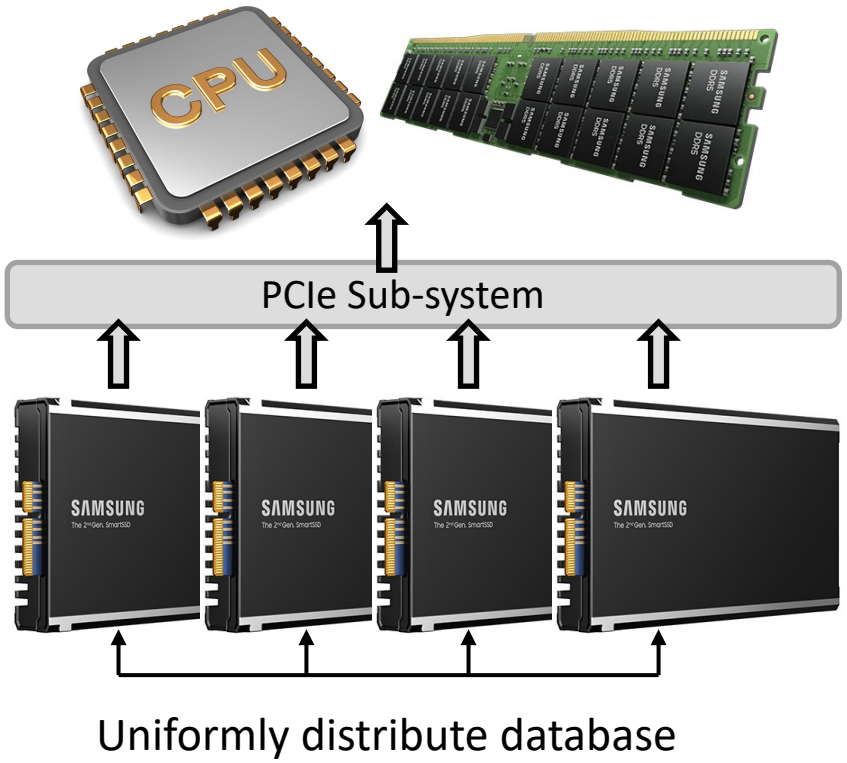

Uniformly distribute database

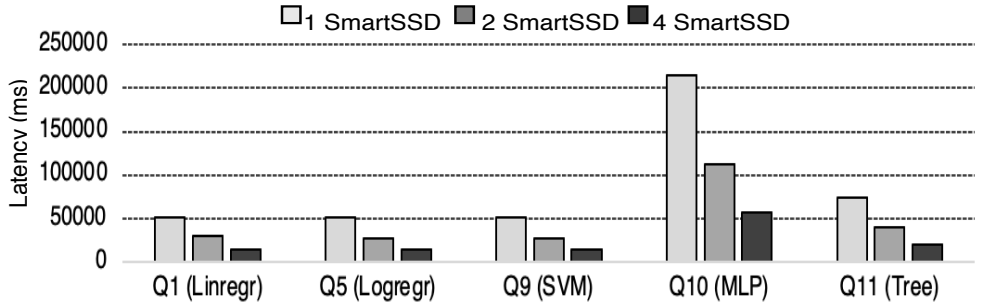

- With 2 SmartSSD ➔ **1.85x performance gain**
- With 4 SmartSSDs ➔ **3.66x performance gain**

Linear Performance Improvement
➔ Overall 200x Speedup Achieved

Conclusion

1. **Full Stack System for In-DBMS Advanced Data Analytics**
   - 57.18x faster query processing than CPU-based DBMS
2. **Software Stack (PostgreSQL+) for SmartSSD Integration**
   - Dynamic offloading decision ➔ 96% accuracy
3. **Near-Storage based Hardware Accelerator (i-DPA)**
   - Direct data page decoding & abundant parallel processing (3-levels)

Trinity: In-Database, Near-Data Machine Learning Acceleration Platform for advanced Data Analytics

Thank You!

- **Questions? Feel Free to Contact Me!**
  - E-mail: jihoon0708@kaist.ac.kr
  - LinkedIn: https://www.linkedin.com/in/ji-hoon-kim-860429225/
  - Slack: Trinity: In-Database Near-Data ML Acceleration platform