Flash Memory Summit

XILINX
ALL PROGRAMMABLE™

Database Acceleration Solution Using FPGAs and Integrated Flash Storage

HK Verma, Xilinx Inc.

FPGA Analytics in Flash Storage System

Flash Memory Summit

- In-memory or Flash storage based DB reduce disk access penalty, but do not ease CPU “compute" bottleneck
- FPGA Analytics accelerators overcome the CPU bottleneck, enable new workloads
- FPGAs bring compute closer to flash storage devices

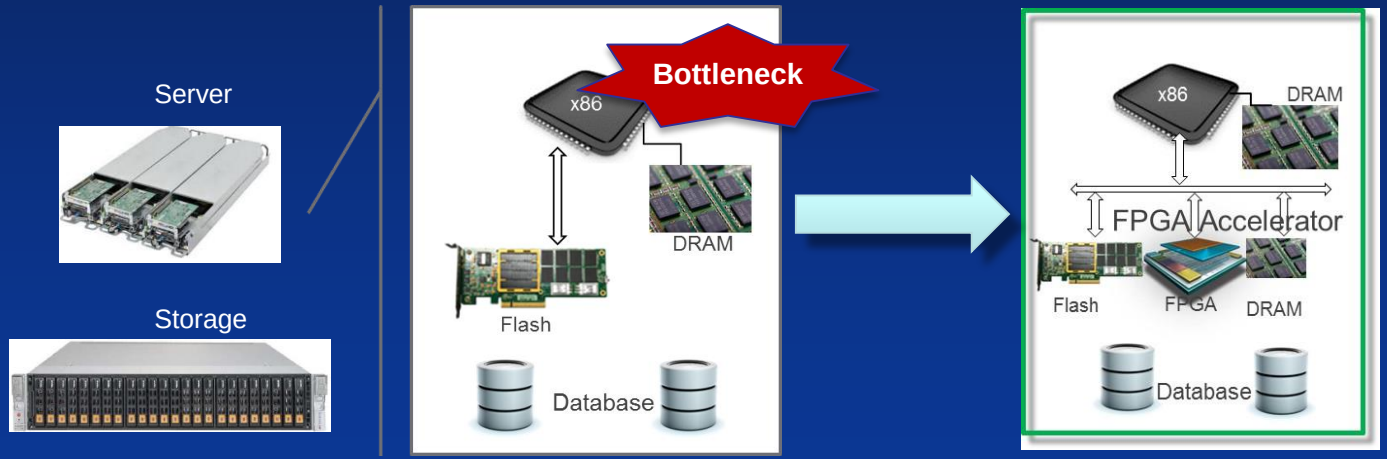

In-memory system does not ease the CPU Bottleneck, FPGA reduces the bottleneck

FPGA Accelerator between Server and Storage

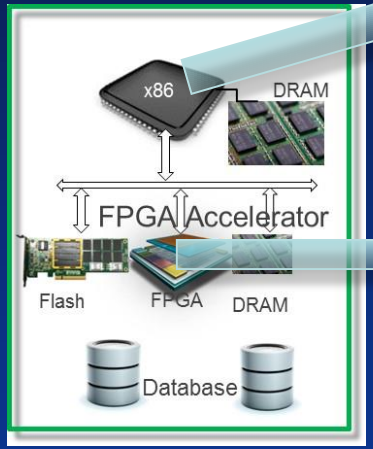

FPGA connected to SSD using PCIe
Integrated FPGA & SSD board

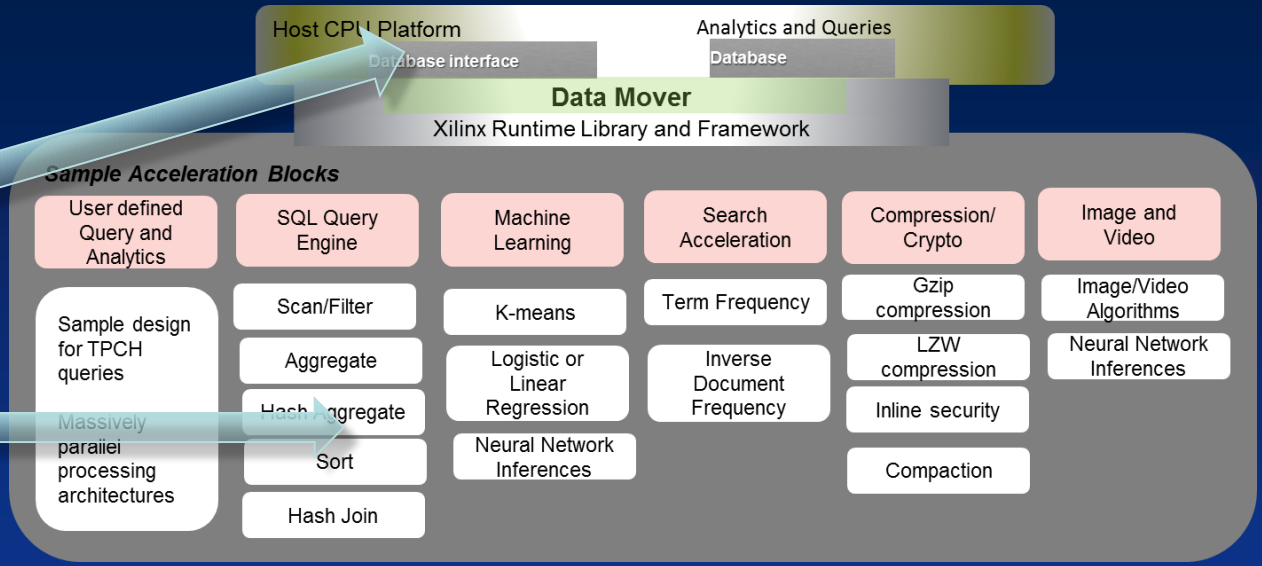

- Tools, methodology, platforms, acceleration design blocks available for customers to implement target acceleration logic on FPGA
- Tighter integration with storage reduces CPU overheads, adds new functionality in storage
- Acceleration solution available on Amazon AWS as F1 instance to try and deploy

Reduce CPU Usage, add new functionalities in storage solutions by using FPGA Acceleration

Database acceleration offload to FPGAs

Big Data Application
10s-100s 0f GigaBytes TeraBytes

Stored in Database or Text files
- RDBMS or NoSQL such as PostgreSQL
- Text Files like csv
- Key-value Database like RocksDB

Application Interface
- Prepare Data Blocks
- Use instruction as metadata of the Data Blocks to program for target functionality
- Stream into FPGA using Data Mover
- Gather the results from FPGA and produce the final result

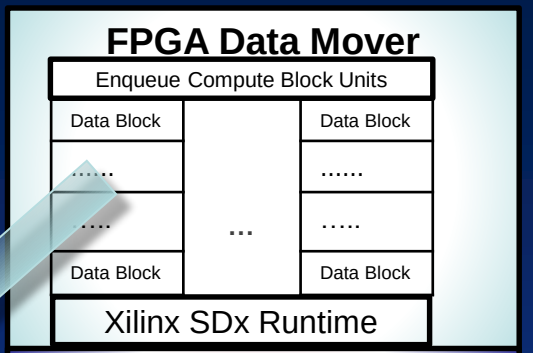

- Identify acceleration components from the big data applications
- Use Xilinx tools, methodology to build an offload engine
- Release and try it on cloud such as Amazon AWS F1

- Standard API for an efficient data mover
- Common across different types of application
- Data DMA, Kernel executions synchronized using OpenCL events to ensure maximal sustained data rate over PCIe

Integrated Flash Storage and Xilinx FPGAs solutions for compute offloads closer to FPGAs

- Discrete or FPGA integrated controller solution
- Direct FPGA connection or connected through PCIe

FPGA and flash storage integrated solutions offering application specific offloads

FPGA Implementation Methodology for Database Acceleration

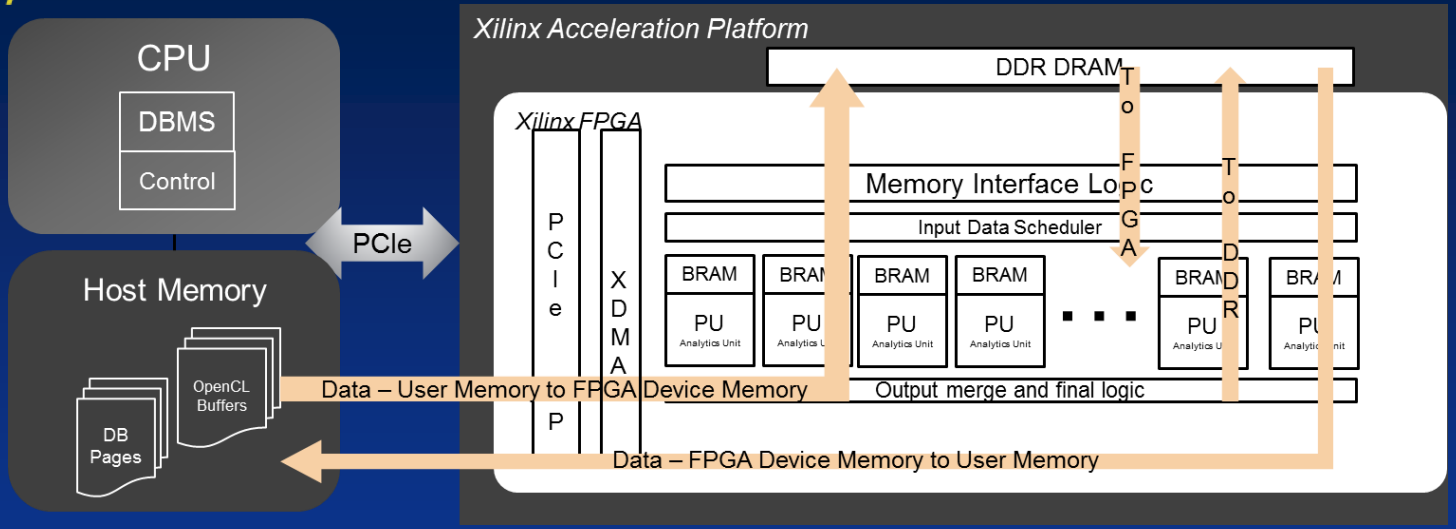

- Efficient implementation of massively parallel processing units (PU)
- Provides 10-25x performance improvement over CPUs across many applications
- Accelerated functionality uses less than half of Xeon power

Innovative technology offering 10-50x compute efficiency improvement over CPUs

Database TPCH Query 6 & 1 Acceleration

Flash Memory Summit

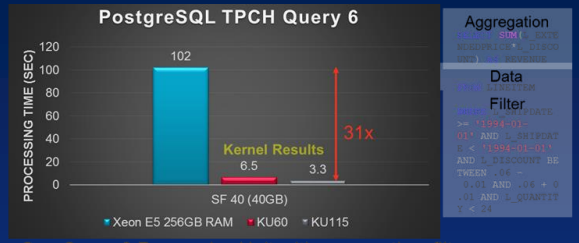

One Query 6 Processing Unit with aggregation, filter

HLS Results	Block Size	PU Block Size	LUT Total	FFs	RAMB18	DSP48E	Freq. (MHz)
Query 6 PU	2MB	32KB	3936	2596	35	22	250

Query 6 - KU60 can fit up to 64 Units, KU115 128 Units

- Application pulls data from PostgreSQL, and sends to FPGA platform accelerating Query 6 and Query 1 of TPCH benchmark processing 40GB database size
- Query implemented on FPGA device using C HLS and massively parallel instantiations of Query Units
- Code and methodology available now
- CPU query implementation is from default PostgreSQL results on Xeon E5 platform

Processing Units in C HLS

- Native support of Date, Decimal SQL types
- Complete query logic on FPGA
- Combine many parallel PU results in HLS for a Compute Unit Kernels

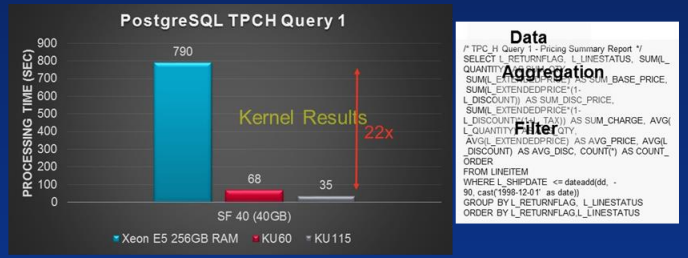

One Query 1 Processing Unit with aggregation, filter

HLS Results	Block Size	PU BlockSize	LUT Total	FFs	DSP48E	Freq (MHz)
Query 1 PU	2MB	128KB	6054	2988	38	250

Query 1 - KU60 can fit up to 16 Units, KU115 32 Units

TPCH Query demo with PostgreSQL, text data available on Nimbix cloud

SQL Query Engine on Xilinx FPGA

Flash Memory Summit

PostgreSQL interface

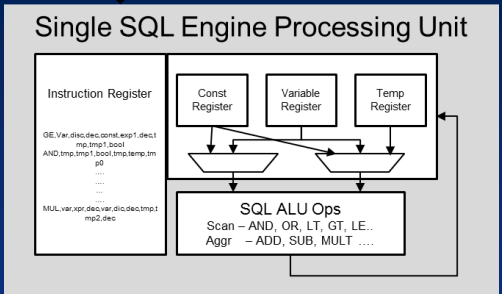

- SQL engine with SQL processing opcodes that processes Postgres storage blocks
- Customers can offload existing SQL query on FPGA
- Suitable for large data with scan, filter, aggregate, hash aggregate processing

PostgreSQL UDF (User Defined Function)

- PostgreSQL UDF exposed as a SQL command
- User calls this query for qualified scanning and aggregation FPGA offload on the PostgreSQL data

PostgreSQL Generic SQL

Amazon EC2 F1 Instances

Run Customizable FPGAs in the AWS Cloud

- Client side plugin to insert FPGA offload computation into a Postgres SQL query plan
- Automatic generation of SQL Engine opcodes for the given SQL query
- Process any user SQL query without reprogramming FPGA
- Provides a solution to offload existing customer SQL commands to an FPGA

PostgreSQL Amazon AWS F1 image (AFI) using Xilinx FPGAs releasing shortly !!

Summary

- A system architecture integrating FPGAs and flash storage proposed and available on cloud for customers to try it out
  - Hardware programming in FPGAs enable tighter integration with flash storage devices
- FPGAs overcome CPU compute bottleneck
  - FPGAs have superior energy efficiency and performance
- Proposed efficient method to build massively parallel Processing Unit architecture on FPGAs
  - Custom hardware integration of parallel units reduces inefficiencies with parallelization
- Standard API based data mover for Xilinx platforms with SDx runtime libraries to move data from Application user memory to FPGA device memory
- Application designers need to build efficient application interface