On-chip Database Acceleration with FPGA

Geetesh More, Suprio Ray, Kenneth Kent
Faculty of Computer Science, University of New Brunswick
{gmore, sray, ken}@unb.ca

PROBLEM STATEMENT

As technology advances, data volume increases. This results in the need for faster computation. While the performance increase of general-purpose CPUs is not able to keep up with the advances in their periphery, reconfigurable logic found in FPGAs along with their highly parallelized architecture enables them as most sought-after alternative to CPU for database acceleration.

WHAT IS FPGA?

FPGAs are the matrix of interconnected and reconfigurable logic block such as memory (BRAM), flip flop (FF), LUTs and IOs. They can be re-programmed to operate as a custom hardware solution. Their ability to operate at low clock frequencies allows to significantly reduce power consumption.

WHAT IS LEARNED INDEX?

Learned index data structures such as RadixSpline, PGM employes certain ML based technique to lookup specific key on sorted data set.

CONTRIBUTIONS

- Development of on-chip database acceleration with FPGA for key search on the sorted data.
- Methodology for predicting selective compute intensive operation.

GOALS

- Acceleration of selective operation with FPGA
- Maintain area utilization under scalable margin
- Attain higher throughput on FPGA as compared to CPU.

APPROACHES

- Accelerated the entire learned index on FPGA.
- Bottom-up approach to predict compute intensive operation
- Selected operation acceleration on FPGA
- Code optimization for better performance..

RESULT

- Accelerating the entire learned index:
  - FPGA resource utilization exceeds the maximum limit.

ARCHITECTURE OVERVIEW

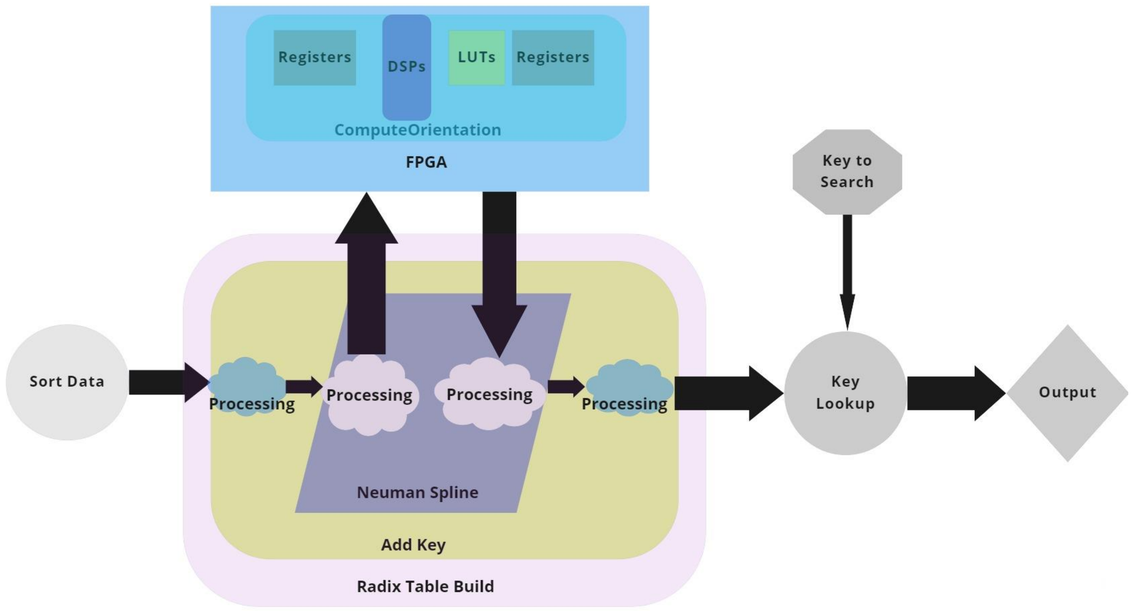

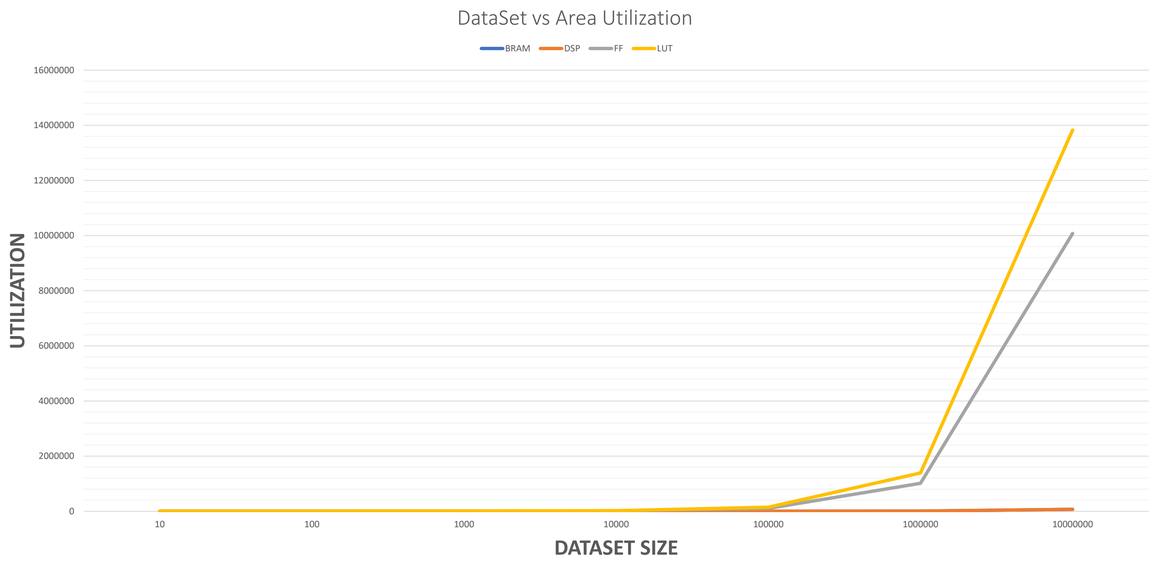

- Accelerating only compute intensive operation:
  - 50% Runtime improvement and
  - Area utilization of FPGA resources is less that 1%.

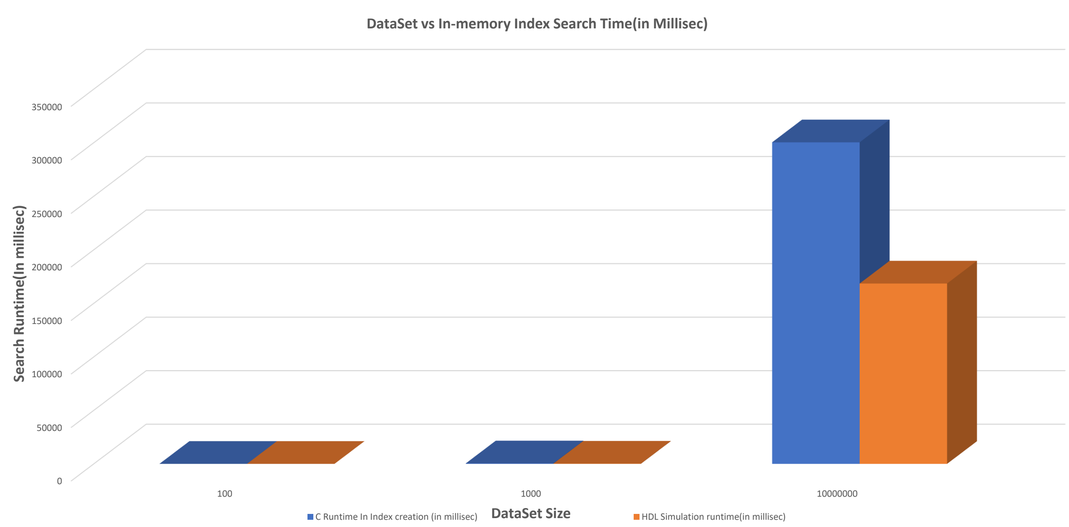

CAS–Atlantic