Concurrent execution of an analytical workload on a POWER8 server with K40 GPUs

A Technology Demonstration

Sina Meraji sinamera@ca.ibm.com

Berni Schiefer schiefer@ca.ibm.com

Tuesday March 17th at 12:00 PM, Hall 1-2

Please Note

- IBM's statements regarding its plans, directions, and intent are subject to change or withdrawal without notice at IBM's sole discretion.
- Information regarding potential future products is intended to outline our general product direction and it should not be relied on in making a purchasing decision.
- The information mentioned regarding potential future products is not a commitment, promise, or legal obligation to deliver any material, code or functionality. Information about potential future products may not be incorporated into any contract.
- The development, release, and timing of any future features or functionality described for our products remains at our sole discretion.

Performance is based on measurements and projections using standard IBM benchmarks in a controlled environment. The actual throughput or performance that any user will experience will vary depending upon many factors, including considerations such as the amount of multiprogramming in the user's job stream, the I/O configuration, the storage configuration, and the workload processed. Therefore, no assurance can be given that an individual user will achieve results similar to those stated here.

Outline

- DB2 BLU Acceleration
- Hardware Acceleration
- Nvidia GPU
- Key Analytic Database Operators
- Our Acceleration Design
- Live Technology Demonstration

DB2 with BLU Acceleration

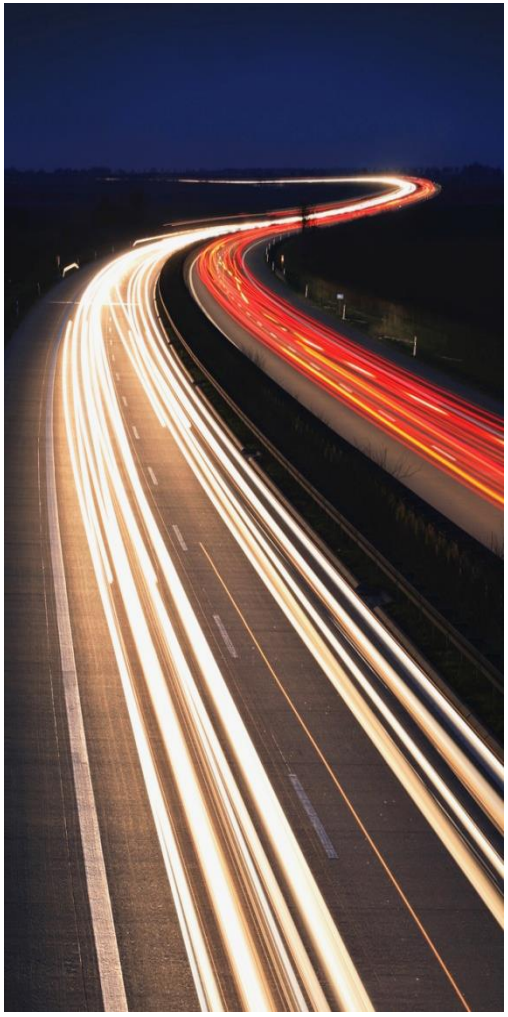

Next generation database

- Super Fast (query performance)
- Super Simple (load-and-go)
- Super Small (storage savings)

Seamlessly integrated

- Built seamlessly into DB2
- Consistent SQL, language interfaces, administration
- Dramatic simplification

Hardware optimized

- Memory optimized
- CPU-optimized
- I/O optimized

OpenPOWER™

Risk system injects 1/2 TB per night from 25 different source systems.
"Impressive load times."

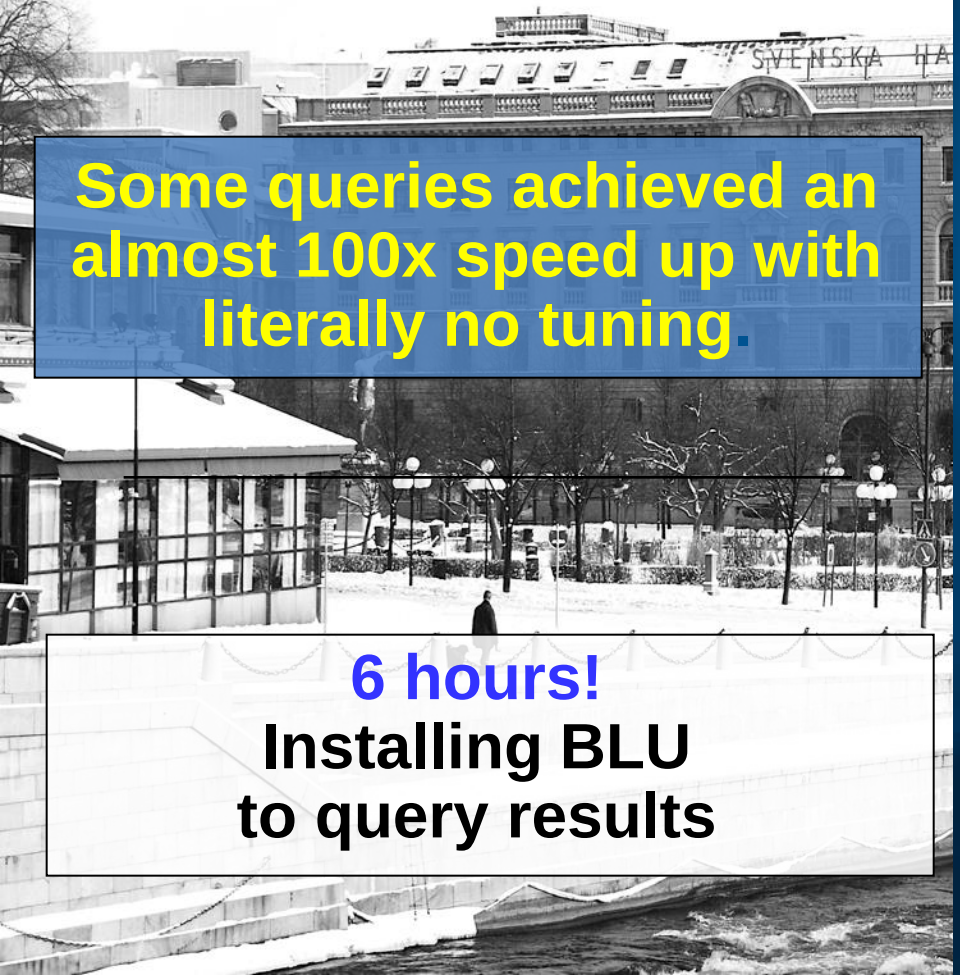

Some queries achieved an almost 100x speed up with literally no tuning.

6 hours!
Installing BLU to query results

One of the world's most profitable and secure rated banks.

DB2 with BLU Acceleration: The 7 Big Ideas

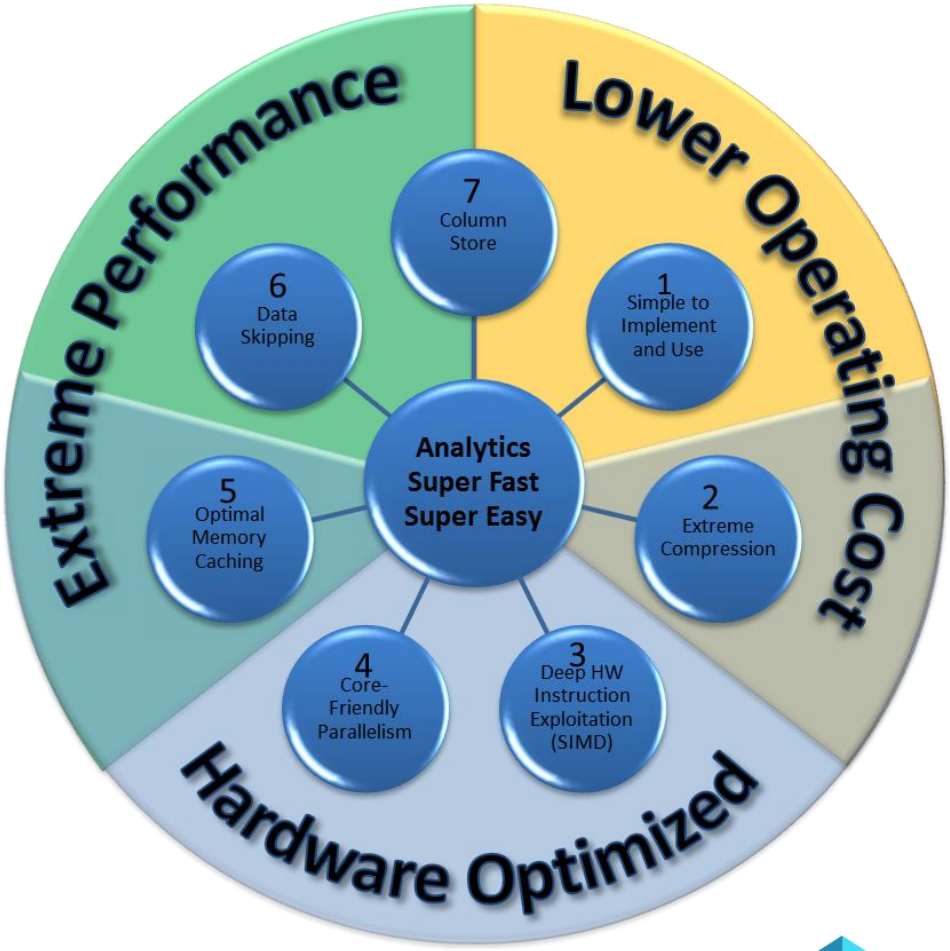

OpenPOWER™

Hardware Acceleration

- Use specific hardware to execute software functions faster
- Popular accelerator technology
  - SIMD
    - Present in every CPU
  - GPUs
    - Easy to program
  - FPGA
    - Hard to program

Nvidia GPU

- NVIDIA Tesla K40
  - Kepler technology
  - Peak double precession performance: 1.66 TFLOPs
  - Peak single precession performance: 5 TFLOPs
  - High Memory Bandwidth: up to 288 GB/Sec
  - Memory Size: 12GB
  - Number of cores: 2880

Key Analytic Database Operators

- GROUP BY / Aggregation
  - SELECT column_name, aggregate_function(column_name)
    FROM table_name
    WHERE column_name operator value
    GROUP BY column_name;
- Join
  - SELECT column_name(s)
    FROM table1
    JOIN table2
    ON table1.column_name=table2.column_name;
- Sort
  - SELECT column_name
    FROM table_name
    ORDER BY column_name;

Hardware Configuration

- POWER8 S824L
  - 2 sockets, 12 cores per socket, SMT-8, 512GB
  - Ubuntu LE 14.04.02 LTS
- GPU:
  - 2 NVIDIA Tesla K40

Our Acceleration Design

- Use parallel POWER8 threads for reading/pre-processing data
- Transfer data to GPU
- Have the GPU to process the query
- Transfer the result back to host machine

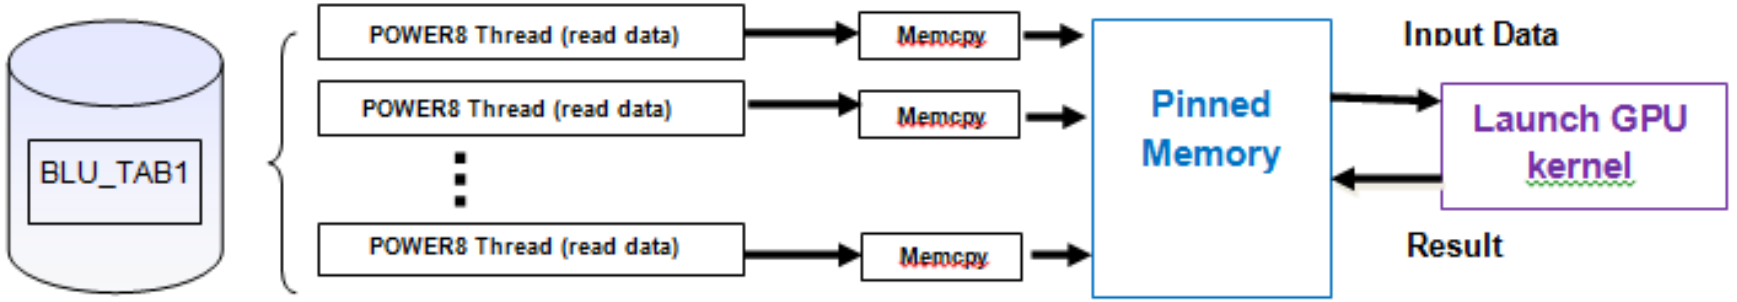

Hybrid Design: Use Both POWER8 and GPU for Query Processing

- Decide where to execute the query dynamically at runtime
  - Use GPU only
  - Use CPU only
  - Use both

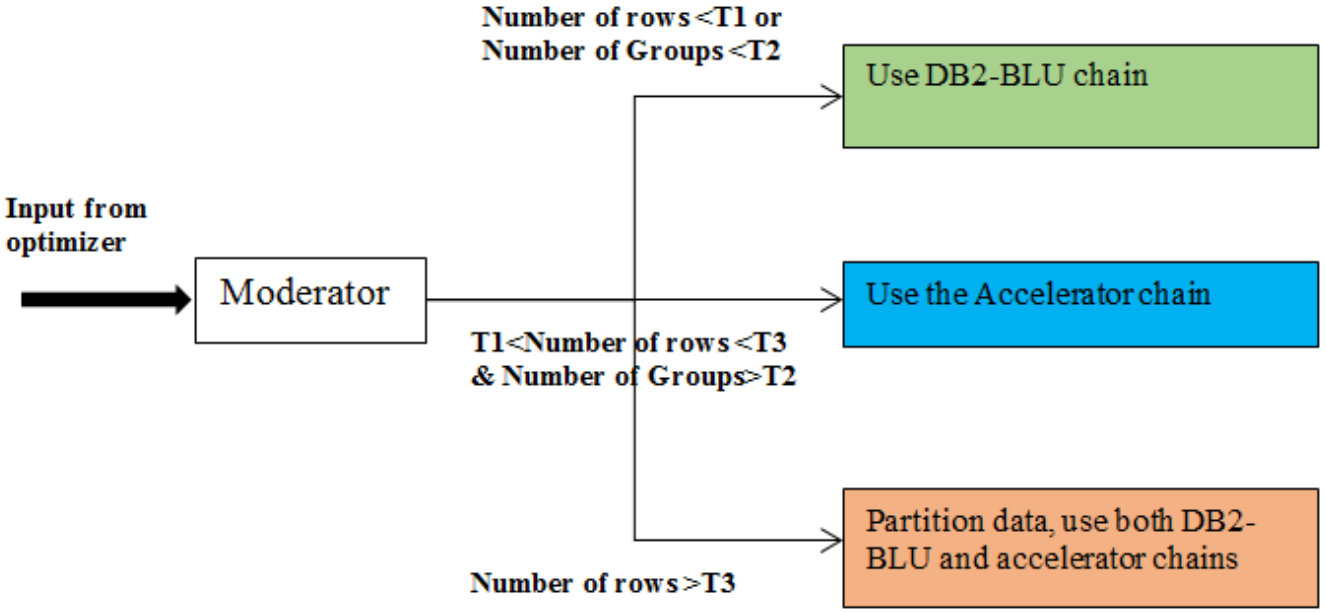

GPU Kernels

- Design and develop our own GPU runtime
- Developed fast kernels
  - e.g. GROUP BY, aggregation
- Use Nvidia CUDA calls
  - e.g. Atomic operations
- Use Nvidia fast kernels
  - e.g. sort

Hash-Based Group By/Aggregate

- SELECT C1, SUM(C2) FROM Simple_Table GROUP BY C1

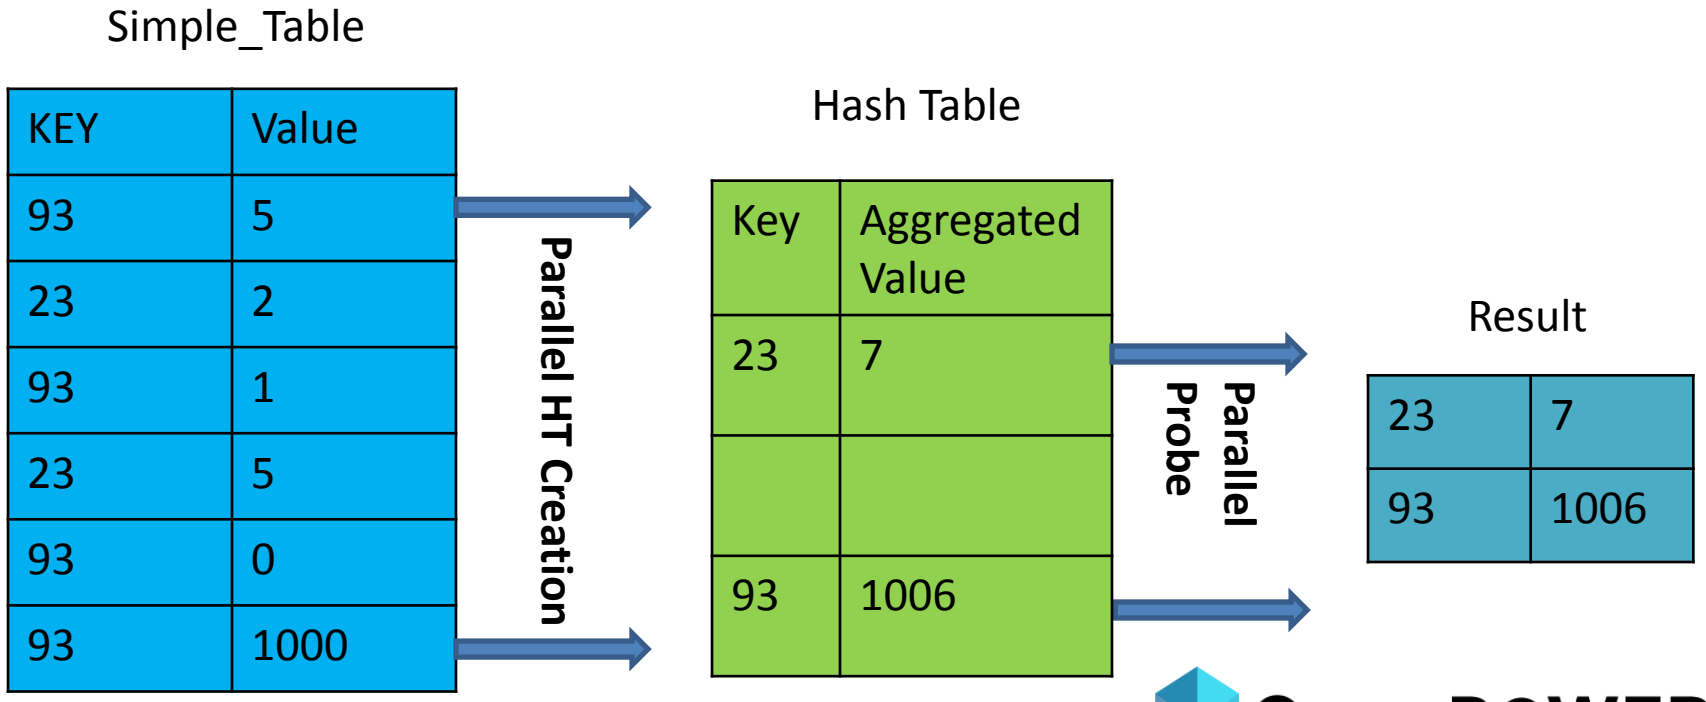

Acceleration Demonstration

- Accelerating DB2 BLU Query Processing with Nvidia GPUs on POWER8 Servers
  - A Hardware/Software Innovation Preview

- Compare query acceleration of DB2 BLU with GPU vs. non- GPU baseline
- Show CPU offload by demonstrating increased multi-user throughput with DB2 BLU with GPU

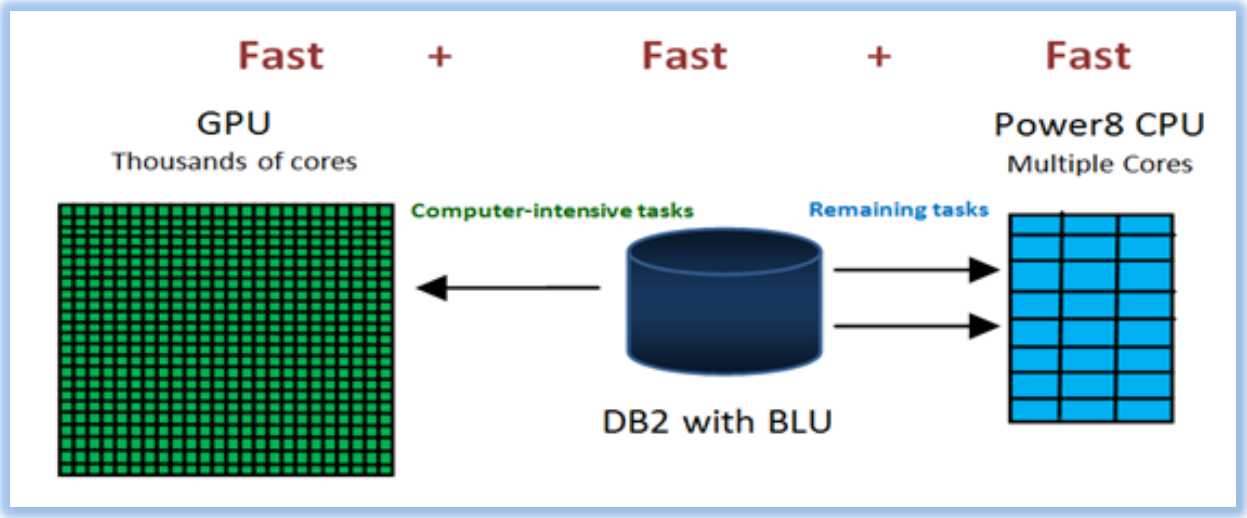

BLU Analytic Workload

- A set of Queries from existing BLU Analytic workloads
  - TPC-DS database schema
    - Based on a retail database with in-store, on-line, and catalog sales of merchandise
  - 15% of queries use GPU heavily
  - 50% of queries use GPU moderately
  - 35% of queries do not use GPU at all
- Benchmark Configuration
  - 100 GB (raw) Data set
  - 10 concurrent users

Performance Result

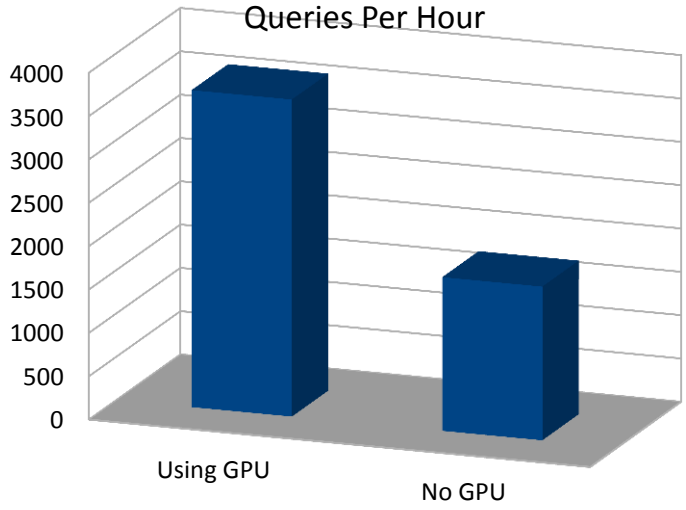

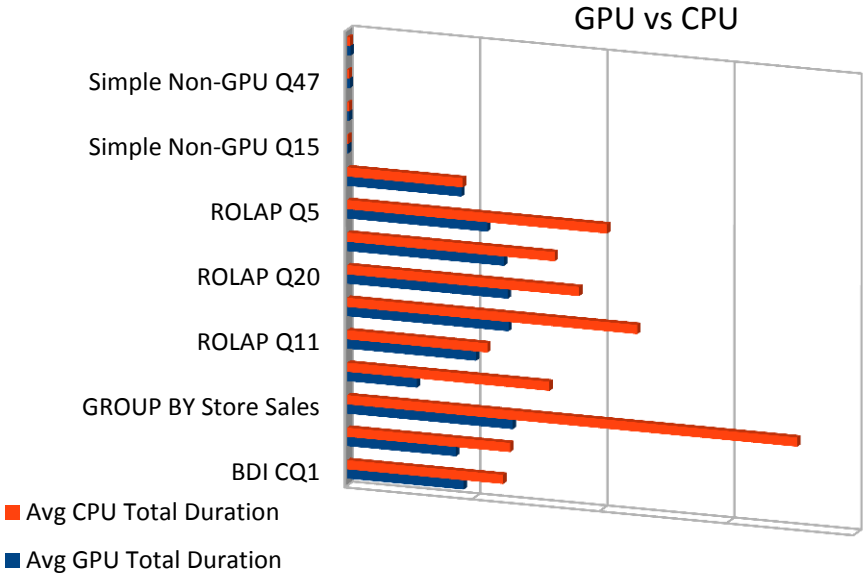

- ~2x improvement in workload throughput
- CPU Offload + improved query runtimes are the main factors
- Most individual queries improve in end-to-end run time

GPU Utilization

The DB2 BLU GPU demo technology will attempt to balance GPU operations across the available GPU devices

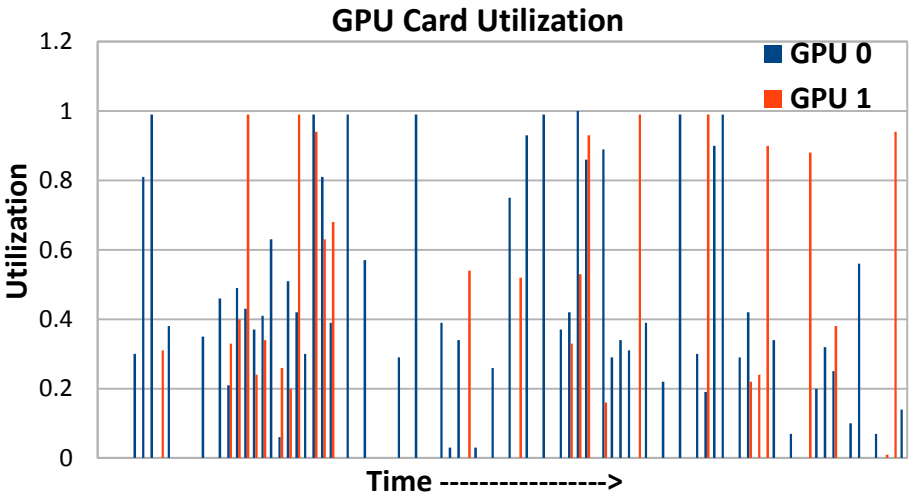

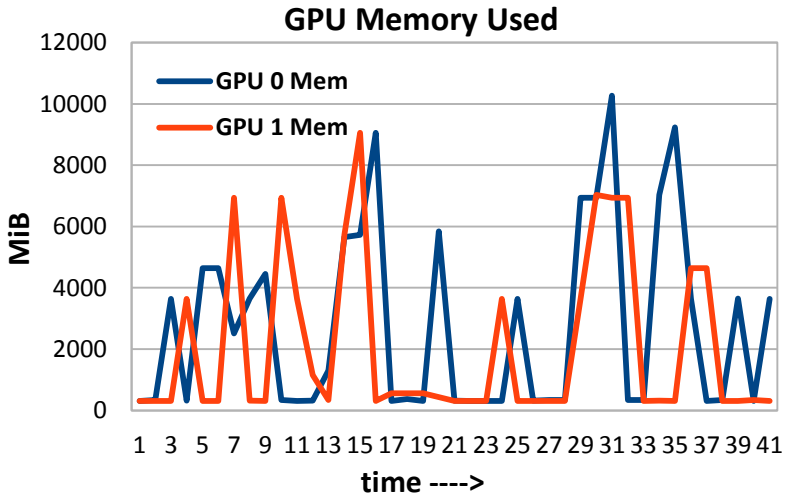

These measurements are taken from the Demo Workload running in continuous mode.

Summary

- Hardware/Software Innovation Preview demonstrated GPU Acceleration
- Improved DB2 BLU query throughput
  - Use both POWER8 processor and Nvidia GPUs
  - Design and develop fast GPU kernels
  - Use Nvidia kernels, function calls, etc
- Hardware Acceleration shows potential for
  - Faster execution time
  - CPU off-loading