

Accelerating Sql Queries Using GPU

Krutika Paradkar¹, Rucha Borane², Pallavee Deshmukh³, Akshay Kale⁴

^{1, 2, 3, 4} Dept of Computer

^{1, 2, 3, 4} Marathwada Mitra Mandal's College of Engineering, Karvenagar, Pune

Abstract- Day by day there is a rapid growth in the generation of data. So, to handle such a huge amount of data there is a need of some effective techniques which avail to store, retrieve and analyze such massive data. Graphics Processing Unit (GPU) were conventionally utilized to optimize image filtering and video processing but it additionally avails to handle critical business applications where the computational power of GPU can provide significant benefits. In this paper, we propose the approach of analyzing the speed up of query execution. In GPU, the data is divided into smaller parts of data with help of threads of the GPU, SQL query is fired parallelly. The proposed approach is implemented on a SQLite database using the CUDA framework. The execution time of both CPU and GPU is analyzed.

Keywords- Parallel Database; Parallel Programming; Distributed Database; Database Management

I. INTRODUCTION

GPU has proved to be an efficient co-processor in the field of conventional computer. In this advancing world of computation the demand for high speed processing is increasing day by day. Web based application and system which uses database on back end required high speed processing to access the data over databases. Most of the top supercomputers in the world contain GPUs to increase performance. Research in applying GPUs to solve parallel problems has been done on numerous applications. The acceleration of SQL queries would enable programmers to do little or no change to their source code.

In database system the computational power of GPU can provide significant benefit to critical business applications for instance Bank's Transaction Process. A database system performs a significant amount of repeated calculations on different data. Traditional databases used to perform the steps of a query sequentially using the CPU. CPU has less but more powerful cores and GPU has more but less powerful cores. Though parallelism is possible in CPUs by using multiple cores, but attempting to achieve greater parallelism requires the addition of costly additional CPUs, and communication between large numbers of CPUs is problematic. A GPU database is feasible due to the considerable inbuilt parallelism in the way databases operate. A database query consists of a

set of operations that are performed on the rows of one or more tables. Each of these operations is repeated, potentially once for each row in the table. Within the query, the execution of these operations is largely data independent; the execution of operations on one row usually does not affect the operations being executed on another row. This introduces to a large amount of data parallelism. Each set of operations could be executed simultaneously on every row.

The proposed approach has been conducted on an NVIDIA GeForce GTX TITAN X which is an ultimate graphics card. It is a combination of the latest technologies and performance of the new NVIDIA Maxwell™ architecture. It has 3072 CUDA cores inside and 12 GB memory.

II. LITERATURE SURVEY

There has been a lot of development in the past few years in the field of general-purpose graphics processing unit databases. In 2010, Peter Bakkum implemented SQLite database management system on a GPU using CUDA framework. In his work he demonstrated core principles of GPU database implementation with the avail of a database capable of performing SELECT queries on integer values. Also, it uses GPU-mapped memory with efficient caching. Consequently, it can compute a larger size of data which exceeds GPU physical memory size.

Glen's explored new ways of handling complex data types and managing data and workloads. Their work focused on three critical areas for developing a database system on a GPU: data caching to manage the data on the GPU, processing table joins and managing the resultant workload on the GPU, and handling irregular and complex data types.

In 2015, Patta's included a query parser which accepts SQL queries in the form of input and the output will be in the form of intermediate code or opcode. They used Virtual Machine to execute the opcode program, to process data records and give a result set as output. Also, they used the Table Functionality for managing groups of tablets in tables. The speed comparison between the GPU and CPU was done with the help of a SQL profiler.

In 2017, Shehab et al proposed a new hybrid query processing technique that makes use of capabilities of CPUs and GPUs where they proposed approach which breaks down each SQL statement into smaller parts during the parsing process. Then in which it automatically manages the distribution of different query parts to be executed either on the CPU only on the GPU co-processor only or on both CPU and GPU. Their main focus was on how to automatically decide which part of the query should be executed on CPUs and which others should be executed on the GPU co-processor for that they demonstrated how to control and dispatch the query parts making a decision about which query part should be executed on the CPU and which should be executed on the GPU co-processor.

III. PROPOSED SYSTEM

The proposed approach used in this paper is used to analyze the speed up of query execution. In this system, the large data is taken for processing. The SQL query is then executed first on CPU. The execution time is calculated for the CPU. The same data is used with the GPU. In GPU, the data is divided into small parts of data. On each part of data with the help of threads of the GPU, SQL query is fired parallelly. The result of each thread is stored in a file. The execution time of both CPU and GPU is analyzed.

In this proposed approach we present the implementation of SQL operations using CUDA framework on the GPU platform. A programming language like CUDA helps to directly interact with the GPU. We propose the optimization of selection queries and aggregation functions with where-clause operations. Using the CUDA programming the data is stored on the GPU in row-column form.

The NVIDIA GeForce GTX TITAN X card is used which is an ultimate graphics card which combines the latest technologies and performance of the new NVIDIA Maxwell™ architecture. It has 3072 CUDA cores and 12 GB memory with memory bandwidth of 336.5 GB/sec and memory type of GDDR5 SDRAM placed on the card.

ADVANTAGES

The main objective is to retrieve data quickly. Most of the databases are built and filled with data due to which leads to slow speed. Sometimes the wait time could range from minutes, to hours, and days as well in worst cases. Query

Optimization is used in database administration. The Queries allow the user to extract significant information from a database. Instead of scanning the whole table, they can

predefine the categories of information which will be retrieved.

IV. CONCLUSION

The SQLite database was used as a platform for executing the queries. SQL parsing mechanism and switching between CPU and GPU execution was useful and enhancing. Execution on the GPU was supported by re-implementing the SQLite virtual machine as a CUDA kernel. In this paper database queries are offloaded on the GPU to achieve effective acceleration of database operations. In the upcoming years the size of NVIDIA GPUs will increase and it will be able to handle large amountof data and performance will also increase.




Fig 1. Overview of system

Sr no	Query	Row effected	CPU time	GPU time	Speed up
1.	select * from account;	1048576	2.01 sec	0.85 sec	42x
2.	select Beat from account where Beat >750	668296	0.65 sec	0.39 sec	60x
3.	select Primary_type from account where Primary_type="ASSAULT"	71218	0.61 sec	0.34 sec	55x

```

10123448 | NARCOTICS | true | false | 923 | 63 | 2015 |
10123449 | THEFT | false | false | 1832 | 8 | 2015 |
10123450 | THEFT | false | false | 832 | 64 | 2015 |
10123451 | CRIMINAL_DAMAGE | false | false | 735 | 67 | 2015 |
10123452 | ASSAULT | false | false | 944 | 61 | 2015 |
10123453 | ASSAULT | false | false | 421 | 43 | 2015 |
10123454 | RECEIPTIVE_PRACTICE | false | false | 133 | 35 | 2015 |
10123455 | BATTERY | false | true | 1131 | 26 | 2015 |
10123456 | CRIMINAL_DAMAGE | false | true | 923 | 63 | 2015 |
10123457 | MOTOR_VEHICLE_THEFT | false | true | 1924 | 6 | 2015 |
10123458 | BATTERY | false | true | 831 | 66 | 2015 |
10123459 | BURGLARY | false | false | 832 | 66 | 2015 |
10123460 | CRIMINAL_DAMAGE | false | false | 822 | 63 | 2015 |
10123461 | THEFT | false | false | 1834 | 8 | 2015 |
10123462 | CRIMINAL_DAMAGE | false | false | 364 | 76 | 2015 |
10123463 | THEFT | false | false | 111 | 32 | 2015 |
10123464 | BATTERY | true | false | 835 | 76 | 2015 |
10123465 | NARCOTICS | true | false | 823 | 66 | 2015 |
10123466 | PROSTITUTION | true | false | 814 | 56 | 2015 |
10123467 | THEFT | false | false | 1123 | 27 | 2015 |
10123468 | OFFENSE_INVOLVING_CHILDREN | false | false | 443 | 46 | 2015 |
10123469 | CRIMINAL_DAMAGE | false | false | 1424 | 24 | 2015 |
10123470 | OTHER_OFFENSE | false | false | 2833 | 3 | 2015 |
10123471 | THEFT | false | false | 1742 | 15 | 2015 |
10123472 | RECEIPTIVE_PRACTICE | false | false | 831 | 66 | 2015 |
10123473 | BATTERY | false | false | 835 | 57 | 2015 |
10123474 | THEFT | false | false | 925 | 61 | 2015 |
10123475 | BURGLARY | false | false | 213 | 38 | 2015 |
10123476 | THEFT | false | false | 921 | 58 | 2015 |
10123477 | CRIMINAL_TRESPASS | false | false | 1932 | 7 | 2015 |
10123478 | NARCOTICS | true | false | 923 | 63 | 2015 |
10123479 | THEFT | false | false | 824 | 66 | 2015 |
10123480 | THEFT | false | false | 1832 | 8 | 2015 |
10123481 | THEFT | false | false | 831 | 66 | 2015 |
10123482 | CRIMINAL_DAMAGE | false | false | 735 | 67 | 2015 |
10123483 | ASSAULT | false | false | 944 | 61 | 2015 |
10123484 | ASSAULT | false | false | 421 | 43 | 2015 |
10123485 | RECEIPTIVE_PRACTICE | false | false | 133 | 35 | 2015 |
10123486 | BATTERY | false | true | 1131 | 26 | 2015 |
10123487 | CRIMINAL_DAMAGE | false | true | 923 | 63 | 2015 |
.....
1048576 rows in set (2.01 sec)
    
```

CPU execution

```

10123457 BATTERY false true 1131 26 2015
10123463 CRIMINAL DAMAGE false true 933 61 2015
10123261 MOTOR VEHICLE THEFT false false 1924 6 2015
10123266 BATTERY false true 831 66 2015
10123267 BURGLARY false false 832 66 2015
10123270 CRIMINAL DAMAGE false false 822 63 2015
10123277 THEFT false false 1894 8 2015
10123285 CRIMINAL DAMAGE false false 1614 76 2015
10123286 THEFT false false 111 32 2015
10123299 BATTERY true false 835 70 2015
10123329 NARCOTICS true false 823 66 2015
10123331 PROSTITUTION true false 814 56 2015
10123337 THEFT false false 1123 27 2015
10123362 OFFENSE INVOLVING CHILDREN false false 413 48 2015
10123379 CRIMINAL DAMAGE false false 1424 24 2015
10123388 OTHER OFFENSE false false 2833 3 2015
10123389 THEFT false false 1722 15 2015
10123399 DECEPTIVE PRACTICE false false 831 66 2015
10123426 BATTERY false false 815 57 2015
10123427 THEFT false false 925 61 2015
10123438 BURGLARY false false 213 38 2015
10123440 THEFT false false 921 58 2015
10123441 CRIMINAL TRESPASS false false 1932 7 2015
10123448 NARCOTICS true false 923 63 2015
10123449 THEFT false false 624 69 2015
10123450 THEFT false false 1832 8 2015
10123451 THEFT false false 812 64 2015
10123452 CRIMINAL DAMAGE false false 735 67 2015
10123454 ASSAULT false false 934 61 2015
10123455 ASSAULT false false 421 43 2015
10123456 DECEPTIVE PRACTICE false false 133 35 2015
10123457 BATTERY false true 1131 26 2015
10123463 CRIMINAL DAMAGE false true 933 61 2015
-----
1048576 rows in set (0.85 sec)
mysql>
    
```

GPU execution

REFERENCES

[1] Esraa Shehab, Alsayed Algergaway, Amany Sarhan "Accelerating relational database operations using both CPU and GPU co-processor

[2] Rajendra A. patta, Anuraj R. Kurup, Sandip M.Walunj "Enhancing Speed of SQL Database Operations using GPU" International Conference on Pervasive Computing:IEEE-978-1-4799-6272-3

[3] Kalle Karkkainen "Parallelism in Database Operation" Helsinki University October 2, 2012

[4] Peter Bakkum and Kevin Skadron "Accelerating SQL Database Operation on GPU with CUDA" Department of Computer Science University of Virginia,Charlottesville,VA22904, March 14,2010