Thesis Portfolio

Benchmarking GPU-Accelerated Databases Against Traditional Databases

(Technical Report)

A Sociotechnical Analysis of High Frequency Trading

(STS Research Paper)

An Undergraduate Thesis

Presented to the Faculty of the School of Engineering and Applied Science
University of Virginia • Charlottesville, Virginia

In Fulfillment of the Requirements for the Degree
Bachelor of Science, School of Engineering

Paul Lee
Spring, 2020

Department of Computer Science

Table of Contents

Sociotechnical Synthesis

Benchmarking GPU-Accelerated Databases Against Traditional Databases

A Sociotechnical Analysis of High Frequency Trading

Thesis Prospectus

Sociotechnical Synthesis

Everyday billions of trades are made on financial exchanges around the world. With the breakthroughs in computational technology in the 20th century, computers have made it possible to conduct millions of trades per day. These technological improvements gave birth to high frequency trading (HFT), the use of fast computers and complex algorithms to make trades. These firms have brought in unprecedented liquidity and efficiency to markets while providing profitable opportunities for HFT firms. Yet, as with any new and empowering technology, the potential to misuse HFT to manipulate markets became apparent. With the sheer volume of trading activity that takes place on a daily basis, the ability to detect and regulate harmful trading activity became difficult to achieve for regulators.

In my research, I propose a novel technology, GPU-accelerated databases, as a potential tool regulators could use to monitor trading activity. Databases, which are software tools used to store and process information, are essential in recording trading activity. However, the sheer volume of trading activity makes most standard databases incapable of processing and storing all the trading data that is generated every day. As such, leveraging a piece of hardware called graphical processing units (GPUs) accelerate database performance and make it possible to out-performing against other databases. I created a testing framework that benchmarks the performance of a GPU-accelerated database, Kinetica, against a standard database, MySQL. Various aspects of database performance are measured against each other, such as data retrieval speed, data storage speed, and data integrity. This testing framework is conducted through the TPC-DS benchmark, an industry-standard framework used to compare databases against each other. Results show that Kinetica has superior processing speed for both data ingestion and data

retrieval speeds. However, hardware costs will have to be considered, as Kinetica runs on more expensive equipment.

In my STS paper, I analyze the sociotechnical dynamics that drive the skepticism around HFT, despite their fundamental market-enhancing role. Through the SCOT framework, I identify the following stakeholders: HFT traders, market participants, and exchanges. Essentially, HFT traders pay money to exchanges to gain fast access to exchange servers, which allow them to find profitable opportunities. Market participants use these exchanges to find counterparties to trade with. HFT activity provides liquidity to these exchanges by making it possible to rapidly match buy and sell orders with each other. This increased liquidity lowers trading fees and subsequently, makes exchanges more efficient. Despite the beneficial impact of HFT, these traders have become infamous in the financial sector. First, there is fear that HFT can be used to manipulate prices, which harm market participants. However, identifying a certain trade as manipulative is an ambiguous question for regulators. Second, the indirect relationship between HFT and improved markets makes it difficult to comprehend the fundamental benefit of HFT. As such, HFT has developed a negative perception in society due to the difficulty of regulating it and the unwarranted emphasis on HFT's ability to artificially move prices. However, HFT is not unique in this regard; the onset of any new technology generally brings upon an initial wave of distrust. Properly informing the public on the role of HFT and overcoming the technical challenge of monitoring and identifying disruptive HFT trades are natural barriers that must be overcome. Nevertheless, societal acceptance of technology is rarely a smooth process.

By working on both the technical and STS components of this project, I have been able to gain a holistic view of the sociotechnical landscape of HFT. To elaborate, the STS portion of this project helped me contextualize my technical project in terms of societal and historical

technological development. By seeing how my technical portion is merely one component of an entire sociotechnical phenomena, I can appreciate the complexities society faces when disruptive technologies arise. By leveraging the SCOT framework, I was able to better understand the different viewpoints behind each stakeholder and why certain perceptions of HFT came to rise. The technical portion of the paper shows me how solutions to various problems is not only a technological obstacle that must be overcome but also a complex problem of weighing costs and benefits. For example, although GPU-accelerated databases show promise as a superior database for monitoring HFT activity, the additional costs of purchasing specialized hardware can potentially be a shortcoming. In addition, by doing hands-on work with certain technologies, I am able to first-hand see how identifying a solution is merely a first step and that implementation is another challenge of its own.