Accelerate Spark

Bigstream Hyper-Acceleration Layer

SOLUTION BRIEF

INTRODUCTION

The field of analytics is experiencing a constant growth in both the amount of data being processed, as well as sophistication of analyses needed to reap business value.
The infrastructure investment growth is driven by three factors:

- Business objective to analyze every aspect of their business
- Adoption of more data sources with larger data sets
- Faster and more frequent iteration of analytics and/or machine learning algorithms for near real-time insight

Today, there is no automated way for analytics platforms to leverage advanced field programmable hardware.

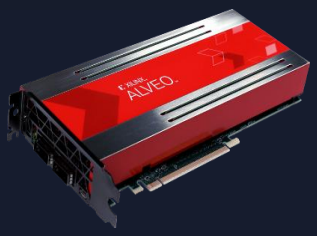

- Zero code change
- Heterogeneous hardware support
- Up to 10X acceleration

KEY BENEFITS

The Bigstream Hyper-Acceleration Layer and the Alveo™ U250 data center accelerator card:

- Does not require any code changes
- Automated programming of applications onto the Alveo U250 data center accelerator card. The result is lower risk and faster time to production
- Accelerates every step of an analytics pipeline, achieving the highest levels of end to end speedup
- Leverages both CPU and FPGA hardware simultaneously to provide optimal performance

SOLUTION OVERVIEW

Bigstream Hyper-acceleration layer automates the process of acceleration for the users of big data platforms. It is comprised of compiler technology for software acceleration via native C++, and FPGA acceleration templates. The automatic compilation of user code for acceleration yields up to a 10X factor in performance.

The technology is available for a number of Xilinx® platforms, including now on U250 accelerators in the Microsoft Azure cloud.

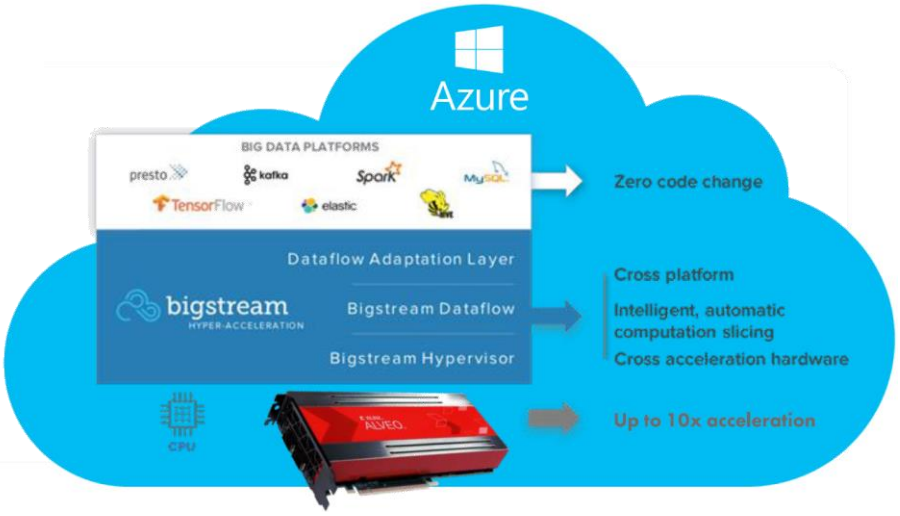

Accelerate Spark

Bigstream Hyper-Acceleration Layer

SOLUTION BENEFITS

The Bigstream Hyper-Acceleration Layer provides:

- **Time-to-Insight:** For many use cases, customers are limited in their ability to iterate their analytics more frequently to manage risk, optimize revenue, or improve customer engagement. Bigstream's acceleration technology enables customers to run their models up to an order of magnitude faster to achieve their business objectives. This is especially important in delay sensitive applications that require a much smaller time scale.
- **TCO Savings:** Doing more with less. Companies want to expand their analytics applications, run them more frequently and leverage larger data sets. But each of these translate to major increase in infrastructure cost. This cost involves charges by cloud providers, or in the case of on premises deployment, involves cost of hardware, power, cooling and additional personnel. Bigstream acceleration can result in savings of 30 70%.
- **Faster pace of innovation:** The scarcity and cost of employing Data Scientists limits the pace of innovation in enterprises. Acceleration of workloads during development and operation frees up Data Scientists to take on more innovative initiatives, driving competitive advantage.

RESULTS

3.9x Faster

- Create Heatmap of number of annoying departure flights on US map in last 5 years

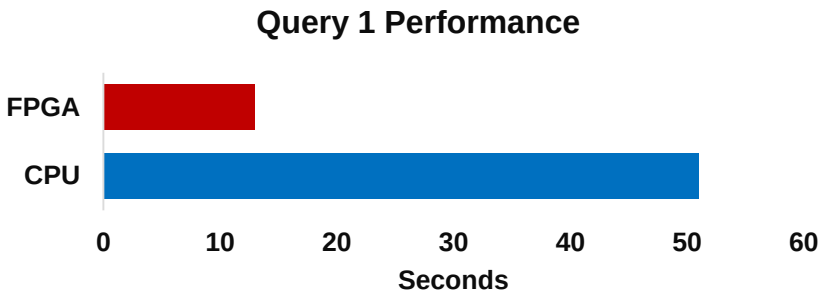

4.4x Faster

- Create Heatmap of number of annoying departure flights in Bay Area since 2000

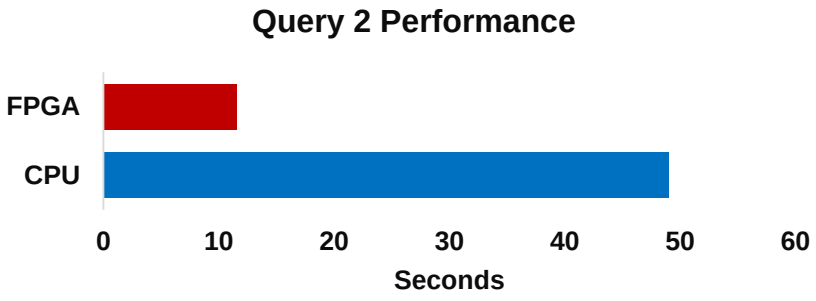

(*) Benchmark results for example on Xilinx Alveo U250 + Intel Xeon Platinum 8171 CPU (management) vs. CPU-only architecture based on Intel Xeon Platinum 8171 CPU

TAKE THE NEXT STEP

Learn more about Xilinx Alveo accelerator cards
Learn more about Bigstream insight.bigstream.co
Reach out to sales@bigstream.co