

Innovation in Investment Banking Technology

Using Graphic Processing Units (GPUs) in Pricing and Risk

A Graphic Processing Unit (GPU) is a **high-speed processor originally designed for processor-intensive commercial gaming applications, which we can use to deliver exceptional performance to pricing and risk models.** This piece of technology allows us to calculate complex factors very quickly, giving traders information they can use to hedge trades appropriately far faster than if conventional CPUs were used.

Power and Versatility

- More than 480 processors and a Teraflop of processing power per card with lower energy consumption than equivalent CPU processing.
- Can accelerate performance by between 10 and 100x across a range of mathematical models using cheap commodity hardware already in use in PCs.
- Powerful software development support and integrated into all the popular development platforms on both Windows and Linux.
- Can be program in open standard language (OpenCL) and already taught in a number of leading Universities.

Business Challenges

- Conversion of existing sequential algorithm into a highly parallel implementation of the GPU.
- Ensure that the final result matches the original algorithm exactly (<0.1 per cent difference).
- Engineering challenges associated with using these technologies (e.g. robustness by enhancement job monitoring, task scheduling, disaster recovery).

Key Benefits (Business/Clients)

- Provide price and risk at very high speed (x10 - x100).
- This capability allows us to achieve more accuracy in pricing by using more factors in the modeling process.
- Provide forward looking view by simulation of the portfolio using various scenarios.
- Manage the risk exposure better using faster and more accurate hedging.

Technology Overview

- Leverage powerful commodity "gamer" chips for pricing of exotics structure trades
- Uses the same GPU (nVidia Tesla M2050) as the world fastest computer Tianhe-1A
- Acceleration through use of coarse-grained parallelism
- Software coded in C using OpenCL and CUDA to keep support cost low
- Provision as part of the J.P. Morgan's Compute BackBone

LOB/Function(s) Impacted

- Credit & interest rates
- Equities
- Commodities
- Risk management & VaR

Industry/External Recognition

- Innovation in Banking Technology Awards 2011 Winner "J.P. Morgan - GPU Deployment for Risk Computation"
- "Developing Effective Systems to calculate Real Time Risk" - TradeTech Architecture 2010
- Computational Fluid Dynamics, Imaging
- Weather, Atmospheric and Space Sciences

Functionality Overview

Introduction of this technology will accelerate (by 10-100x) some of the most complex pricing algorithms (for example Monte Carlo and Finite Difference based algorithms) in the J.P. Morgan Quantitative Financial Library by leveraging the GPU. This involves re-coding some of the routines using nVidia CUDA environment. This allows the GPU to use thousands of parallel threads to achieve the acceleration and performance required. Some of these trades that used to take up to 30 minutes to price can now be completed in < 20s.

Why is the technology innovative?

The tremendous computational power of the GPU was originally developed for the gaming industry. Widespread adoption of this technology to deliver high quality graphics and gameplay has commoditized the hardware and brought exceptional computational speed to a highly affordable price point. At J.P. Morgan, we are harnessing this power to help us to leapfrog our competitors in the speed and accuracy with which we can calculate pricing and risk, giving us a speed advantage in the marketplace.

Development/Delivery

Timeline

- Development start Q1, 2009
- Production: Q4, 2010

Partners

- nVidia: manufacturer of GPU and CUDA
- Mellanox Technologies: Provider of GPU Direct capability through Infiniband interconnect
- Imperial College, London: Ph.D. Research Support

The GPU



NVIDIA GeForce GTX Fermi Products

