

This is a summary of a presentation at Global Derivatives Trading and Risk Management, Paris, May 2010.

About the Authors

Oskar Mencer is CEO at Maxeler Technologies.
Stephen Weston is a Managing Director at J.P. Morgan.

Summary

Derivatives modelers are facing a number of simultaneous challenges: trading and risk management receive risk and PnL the next day, at best; businesses are constantly finding out where they have been rather than looking forward; compute time and cost scale linearly with volume and model complexity; there is no feedback loop between maths, algorithms and computation.

Large scale simulations and higher-order risk analysis using intraday market and trade data will allow businesses to look forward. Increasing the scalability of compute resources, and computing more cheaply, in power, space and time, will remove barriers to business growth.

JPMorgan has assessed a number of solutions to these problems, including GPU and FPGA accelerators. These offer a new approach to such problems and both are being evaluated by JPMorgan. Choosing the right technology for the task is a key factor in acceleration and each technology is suitable for different tasks. Single survival curve bootstrapping for example, is suited to a latency-based (processor) accelerator such as a GPU, because of its recursive nature as a search. Bespoke tranche valuation, with large amounts of regular control and data, is a perfect fit for throughput-based technologies such as FPGAs.

JPMorgan engaged Maxeler in September 2008 to investigate accelerating interest and credit derivative calculations using FPGA accelerators. FPGA technology can achieve high performance in such applications through three key factors: pipelined implementation of the arithmetic functions; multiple, parallel compute pipelines per FPGA device; and multiple devices per compute node.

The project has been structured in three phases. The first phase (Sept 2008 - Dec 2008) was a proof of concept with a reduced problem size that required an order of magnitude improvement in




Figure 1: Example volatility surface.

computation for evaluation of fair value and first order risk.

The second phase (Jan 2009 - Nov 2009) was a detailed research and use case evaluation, which involved a small machine purchase to evaluate the FPGA technology. This phase included fair value and first order risk, developing cross-gamma calculations, limited book-level simulation and migrated multi-factor Monte Carlo interest rates model.

The third phase (Jan 2010 - date) included purchase of a production-sized machine to run large scale simulations and calculate higher order risk measures.

One of the interest rate hybrid models accelerated is SRM3, a multi-variate Monte Carlo model with multiple payoffs which handles combinations of interest rates, FX, credit and equities. The SRM3 application iterates over $T = [100 \rightarrow 1000]$ timesteps, $A = [\sim 2 \rightarrow 4]$ assets / underliers and $P = [10000 \rightarrow 50000]$ paths, creating a 3D calculation space. The FPGA-accelerated solution for a full portfolio of 480 trades provides a 284x speedup over a single-core CPU implementation for fair value calculation.

For a credit hybrid model, calculating fair value and pointwise credit spread risk for a portfolio of 2,925 bespoke tranches showed an improvement of around 240x for an FPGA-accelerated node compared to a single-core CPU implementation. As a result of this project, current computations can run in much less time and operational costs resulting from given computations are dramatically reduced.