

GPU Technology Conference, May 14-17, 2012
McEnery Convention Center, San Jose, California
www.gputechconf.com

Sessions on **Finance** (subject to change)

IMPORTANT: Visit <http://www.gputechconf.com/page/sessions.html> for the most up-to-date schedule.

S0418 - High Productivity Computational Finance on GPUs

Aamir Mohammad (Aon Benfield Securities), Peter Phillips (Aon Benfield Securities)

Day: Tuesday, 05/15 | **Time:** 2:00 pm - 2:50 pm

Topic Areas: Finance; Application Design & Porting Techniques; Parallel Programming Languages & Compilers

Session Level: Beginner

Learn how Aon Benfield helps clients use GPUs to develop and accelerate Monte Carlo derivatives pricing models. We will present our PathWise software tools used by actuaries and quants in order to rapidly develop and deploy production quality, GPU grid enabled, Monte Carlo models, using only high-level languages and tools without requiring any knowledge of CUDA or C/C++. We will describe our approach of using Code Generation, Visual Programming, Domain Specific Languages and scripting languages to create a High Productivity Computing software stack for financial services applications.

S0034 - Real-Time Risk Simulation: The GPU Revolution In Profit Margin Analysis

Gilles Civario (ICHEC), Renato Miceli (ICHEC)

Day: Tuesday, 05/15 | **Time:** 3:00 pm - 3:50 pm

Topic Areas: Finance; Application Design & Porting Techniques; Algorithms & Numerical Techniques

Session Level: Intermediate

Discover how ICHEC helped a world leading company in its sector, to dramatically speed-up and improve the quality of its real-time risk management tool chain. In this session, we present the method used for porting the core-part of the simulation engines to GPUs using CUDA. This porting was realized on two very different simulation algorithms and resulted in speed-ups of 2 to 3 orders of magnitude, allowing much greater accuracy of the results in a real-time environment.

S0435 - Leveraging GPGPU Technology for Valuation of Complex Insurance Products

Chris Stiefeling (Oliver Wyman Financial Services)

Day: Tuesday, 05/15 | **Time:** 4:00 pm - 4:50 pm

Topic Areas: Finance

Session Level: Intermediate

We share our experiences moving a mature, large scale insurance application from a CPU to GPU environment. This session explores the nuances of porting a C++ application when 'blank sheet' re-architecture is not an option. This session will cover: Insurance differences from other financial products (and the implications for the GPU), Considerations when moving an existing, fully featured C++ system to a GPGPU platform, Supporting CPU and GPU implementations from a single code base, Supporting user defined code extensions on the GPU, CUDA 4.0 C++ extensions: experiences, challenges and limitations and Performance case study.

S0219 - Efficient Top-Down Planning in Business Intelligence**Tobias Lauer (Jedox AG), Alexander Haberstroh (Jedox AG)****Day:** Tuesday, 05/15 | **Time:** 5:00 pm - 5:25 pm**Topic Areas:** Databases, Data Mining, Business Intelligence; Finance; Algorithms & Numerical Techniques**Session Level:** Intermediate

In business intelligence, tasks like corporate planning or what-if analysis complement traditional reporting and analysis. One main difference is that while the latter only read data, the former require the change of possibly large numbers of existing and creation of new data records in the business model, preferably in real time. In this session, we describe the extension of an existing BI tool, Jedox OLAP, by GPU-based parallel algorithms for interactive planning scenarios. Compared to sequential in-memory algorithms, our CUDA approach yields tremendous speedups and can also cope with large amounts of data by using multiple GPUs.

S0427 - Intra-Day Risk-Management with Parallelized Algorithms on GPUs**Partha Sen (Fuzzy Logix)****Day:** Tuesday, 05/15 | **Time:** 5:00 pm - 5:50 pm**Topic Areas:** Databases, Data Mining, Business Intelligence; Finance; Algorithms & Numerical Techniques; Supercomputing**Session Level:** Advanced

The challenge with intra-day risk management is that a very large number of calculations are required to be performed in a very short amount of time. Typically, we may be interested in calculating VaR for 100 to 1000 securities per second based on 100 million potential scenarios. The magnitude of these calculations is not Utopian but it reflects the reality of modern financial institutions and exchanges. In this presentation, we outline how the complex problem of intra-day risk management can be solved using parallelized algorithms on GPUs. The methodology has been proven in a POC at 2 financial institutions.

S0383 - Speedup Derivatives and Structured Products Pricing, Reduce TCO Using GPUs**Steve Karmesin (Numerix)****Day:** Wednesday, 05/16 | **Time:** 9:00 am - 9:50 am**Topic Areas:** Finance; Algorithms & Numerical Techniques**Session Level:** Intermediate

Numerix will share its experience using GPU to significantly reduce its customers' Total Cost of Ownership (TCO) and accelerate forward Monte Carlo pricing methods and hybrid models of complex financial structured products and variable annuities. Numerix will describe how it combines complex financial and actuarial modeling with user scripting to drive GPU execution from a script interpreted at run time. This architecture is well suited to financial services firms with portfolios of many different types of structured products where deals are represented independently from the models used to price them.

S0250 - From GPU Computing Toward Full HPC in Finance with GPUs**Pierre Spatz (Murex SAS)****Day:** Wednesday, 05/16 | **Time:** 10:00 am - 10:50 am**Topic Areas:** Finance**Session Level:** Intermediate

During the previous GTC Murex has shown how the company had adapted their generic Monte-Carlo & PDE codes compatible with a payoff language. With one more year of experience with GPUs and OpenCL Murex will show how the company has broadened the usage of GPUs for other subjects like vanilla screening or model calibration and focus on their new challenge 'use as many GPUs as possible' for one single computation.

S0369 - Running Risk on GPUs**Tim Wood (ING Bank nv)****Day:** Wednesday, 05/16 | **Time:** 2:00 pm - 2:50 pm**Topic Areas:** Finance**Session Level:** Intermediate

A key component of Basel III is the Credit Value Adjustment (CVA) which is in essence the value of counter-party credit risk. Quantifying the CVA on simple products already poses considerable computational challenges and considering many banks have hundreds of thousands of positions it becomes clear that the computational challenges of CVA are massive. Calculating CVA sensitivities for hedging only add to this burden. In this talk we will discuss real world applications of GPUs in risk management and show how, using CUDA, GPU computing is an enabling technology to address the computational challenges of an evolving regulatory environment.

S0271 - Fast Adaptive Sampling Technique for Multi-Dimensional Integral Estimation Using GPUs**Pradeep Rao (Infosys Technologies Ltd), Srinivasa Prasanna (Internation Institute of Information Technology Bangalore)****Day:** Wednesday, 05/16 | **Time:** 2:30 pm - 2:55 pm**Topic Areas:** Algorithms & Numerical Techniques; Finance**Session Level:** Intermediate

Evaluating multi-dimensional integrals is a commonly encountered problem in many areas of science including Physics and Volume estimation of convex bodies. One of the widely used techniques for integral evaluation in large dimensions is the Monte Carlo method. Vanilla Monte Carlo methods of Integral Estimation use uniform sampling techniques. Variance of such uniform sampling reduces as $1/\sqrt{\text{Sample-size}}$, which is too slow for most real life applications. In this study, we discuss about an adaptive sampling technique called VEGAS which reduces the variance at a much faster rate than uniform sampling. We present a new parallel implementation for VEGAS based on CUDA that can significantly reduce the computation time of multi-dimensional integrals. We show that our GPU based implementation of VEGAS achieves up to a 45x speed up over an equivalent CPU based implementation.

S0405 - New Generation GPU Accelerated Financial Quant Libraries**Daniel Egloff (QuantAlea GmbH)****Day:** Wednesday, 05/16 | **Time:** 3:00 pm - 3:50 pm**Topic Areas:** Finance; Application Design & Porting Techniques; Algorithms & Numerical Techniques; Cloud Computing**Session Level:** Advanced

Learn from industry experts how new generation GPU accelerated solutions for derivative pricing, hedging, and risk management can be build more efficiently with modern technology and functional programming languages like F# on .NET or Scala on the Java VM. As a concrete example we report from a large derivative pricing project developed in F# on .NET. We will introduce the key design concepts and parallelization strategies, which lead to an efficient and transparent GPU acceleration. Several examples will illustrate the benefit of the functional as compared to the classical object oriented approach.

S0206 - Monte-Carlo Pricing under a Hybrid Local Volatility Model**Sebastien Gurrieri (Mizuho International)****Day:** Wednesday, 05/16 | **Time:** 4:00 pm - 4:25 pm**Topic Areas:** Finance; Algorithms & Numerical Techniques**Session Level:** Intermediate

This session shows how to calculate the prices of several financial products, vanilla and exotic, under Dupire's Local Volatility model. We start with vanilla options on the foreign exchange rate and explain how to rescale the Local Volatility matrix in order to take advantage of the fast texture memory interpolation. We then extend this framework to two factors by including stochastic interest rates following Hull-White model, and show how to price Power-Reverse Dual Coupon swaps with an exotic TARN feature. We provide details of the algorithms and compare accuracy and speed with typical performances of single-core production implementations.

S0377 - C++ Data Marshalling Best Practices**Cliff Woolley (NVIDIA)****Day:** Wednesday, 05/16 | **Time:** 4:30 pm - 4:55 pm**Topic Areas:** Finance; Application Design & Porting Techniques**Session Level:** Intermediate

When integrating CUDA C++ kernels into existing C++ applications, it is at times desirable to migrate a C++ object instance from the host to the device or vice versa. Given variations among host compilers regarding structure layout, accomplishing this data marshalling in a manner that is reliable, simple, and efficient is a complex issue. `cudaMemcpy` is our primary means to transfer data to the GPU, but `memcpy`-style operations are more readily amenable to C-style structures and arrays than to C++ objects or collections of objects. In this session, we will cover the caveats and best practices for marshalling C++ data.

S0100 - Mathematica as a Practical Platform for GPU-Accelerated Finance**Dylan Roeh (Wolfram Research Inc), Abdul Dakkak (Wolfram Research Inc.)****Day:** Wednesday, 05/16 | **Time:** 5:00 pm - 5:25 pm**Topic Areas:** Finance; Development Tools & Libraries**Session Level:** Intermediate

With the introduction of GPU support in version 8, Mathematica has become an excellent environment for integrating CUDA with high level code for interpretation or visualization. In this presentation, we will show the usefulness of Mathematica in the venue of computational finance. In addition to demonstrating the GPU-accelerated financial computations which can be readily performed within Mathematica, we will show that these calculations can easily be integrated with third-party data sources including Microsoft Excel and databases. Furthermore, we will cover the UnRisk Mathematica package written by MathConsult, which seamlessly adds GPU-accelerated complex model calibration algorithms to Mathematica's repertoire.

S0280 - MATLAB and GPU: American Exercise Options Monte Carlo 3 Ways**John Ashley (NVIDIA)****Day:** Wednesday, 05/16 | **Time:** 5:30 pm - 5:55 pm**Topic Areas:** Finance; Development Tools & Libraries**Session Level:** Beginner

The same Longstaff-Schwartz Monte Carlo implementation done 3 ways -- using `GPUArray`, `arrayFun`, and custom kernels -- to illustrate the trade-offs between ease of use and performance. Side by side examination of segments of code and their relative performance will help build intuition about the different ways to leverage GPU acceleration of projects using MATLAB.