

INTRODUCING NVIDIA TESLA GPUs
FOR COMPUTATIONAL FINANCE

NVIDIA® TESLA® ACCELERATE FINANCIAL SERVICES APPLICATIONS WITH GPU SOLUTIONS

“ This is game changing technology. We are now able to calculate fair-value and risks for long-dated books in a matter of minutes during the day, rather than overnight. This is giving us market advantage and driving new business.”

Mark Mullahy,
Co-Head of equity structured trading desk,
JPMorgan

NVIDIA® TESLA® GPUs enable tremendous speed-ups of parallel applications, such as Monte Carlo simulations. These algorithms are widely deployed in financial services firms for pricing and risk analytics.

GPUs enable you to calculate pricing and risk for complex options and OTC derivatives in seconds rather than minutes or hours. This speedup allows you to run more simulations that increase the quality of your results. With more confidence in your data, you are able to offer tighter spreads and gain competitiveness. GPUs make it possible to run complex models that were previously impossible.

With GPUs you can run more complex market or counter party risk calculations and obtain results intra-day rather than overnight. This speed and quality of information provides you with deeper insight into your exposures, enabling you to rapidly adjust positions and reduce risk.

You can dramatically increase the performance of your applications using GPU-accelerated libraries and programming languages, such as the easy-to-use and widely adopted NVIDIA® CUDA®. Accelerate the most performance critical areas in your application using the development tools you already know.

NVIDIA can help you evaluate and realize the benefits of market-changing GPU technology.

GPUs MAKE THE DIFFERENCE IN FINANCE:

- Faster Pricing** → **More Revenue**
- More Modeling** → **Less Risk**
- Maximize Resources** → **Increase Efficiency**

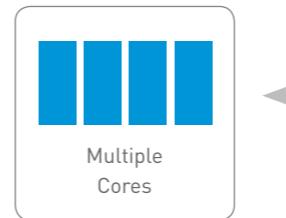
“ The availability of affordable raw computing power is no longer an issue.”

Tim Wood,
Quantitative analyst,
ING Bank

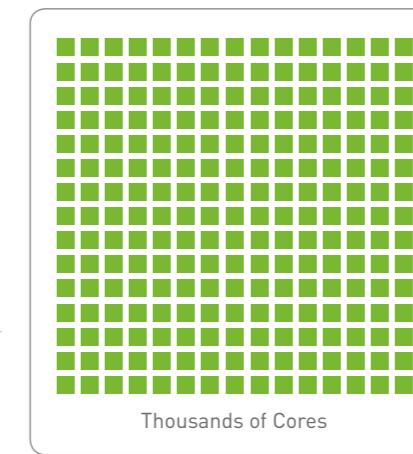
WHY GPU COMPUTING?

With the ever-increasing demand for more computing performance, the HPC industry is moving toward a hybrid computing model, where GPUs and CPUs work together to perform general purpose computing tasks. As parallel processors, GPUs excel at tackling large amounts of similar data because the problem can be split into thousands of pieces and calculated simultaneously. As sequential

www.nvidia.eu/tesla

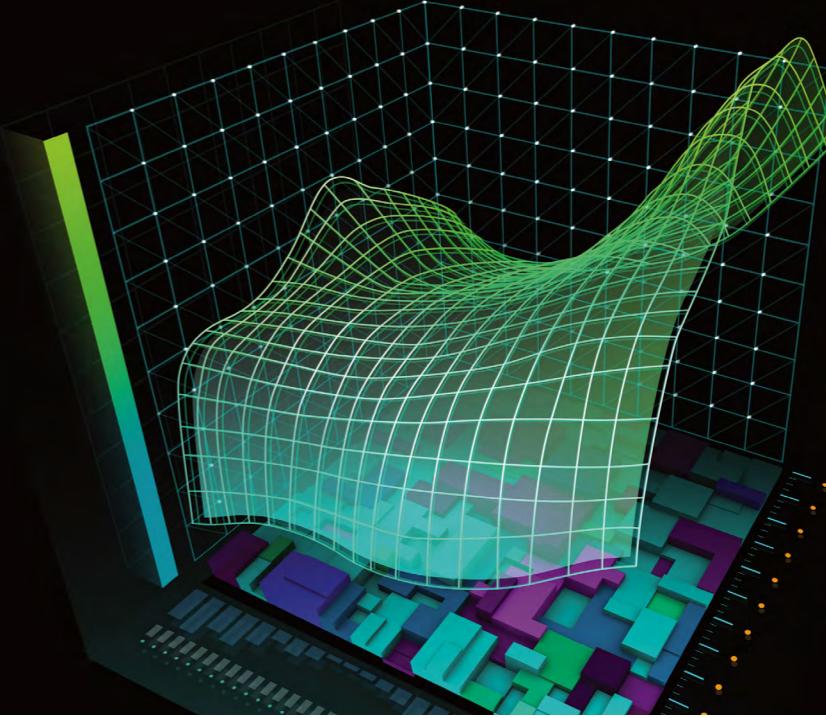
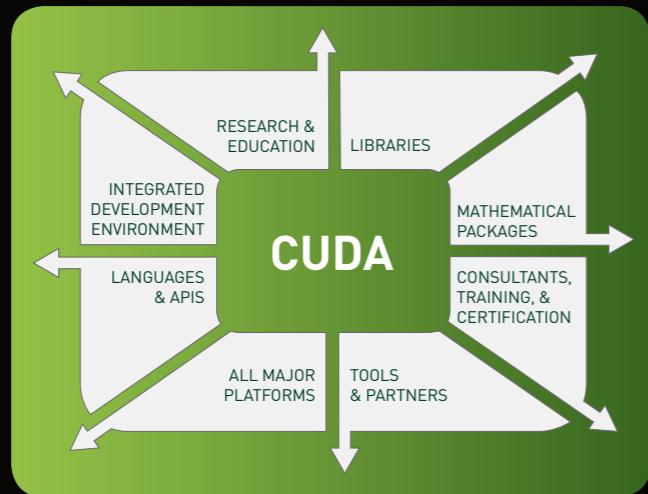


CPU



GPU

processors, CPUs are not designed for this type of computation, but they are adept at more serial based tasks such as running operating systems and organizing data. NVIDIA believes in applying the most relevant processor to the specific task in hand.



PROGRAMMING ENVIRONMENT FOR GPU COMPUTING

NVIDIA's CUDA

CUDA architecture has the industry's most robust language and API support for GPU computing developers, including C, C++, OpenCL, DirectCompute, and Fortran. NVIDIA's Parallel Nsight, a fully integrated development environment for Microsoft Visual Studio is also available. Used by more than six million developers worldwide, Visual Studio is one of the world's most popular development environments for Windows-based applications and services. Adding functionality specifically for GPU computing developers, Parallel Nsight makes the power of the GPU more accessible than ever before.

NVIDIA® NSIGHT™ DEVELOPMENT ENVIRONMENT FOR WINDOWS, LINUX, MAC OS

NVIDIA Nsight™ is the ultimate development platform for heterogeneous computing. Its powerful debugging and profiling tools enables full optimization performance of the GPU.

Nsight comes in two different versions:

- Visual Studio Edition for Windows.
- Eclipse Edition for Linux and Mac OS.

For more information, visit:
developer.nvidia.com/object/nsight.html

NVIDIA TESLA GPUs USED BY J.P. MORGAN TO RUN RISK CALCULATIONS IN MINUTES, NOT HOURS



CHALLENGE

Risk management has been, and continues to be, a huge and increasingly costly focus for the financial services industry. By 2008, JPMorgan's equity derivatives business alone had found itself spending more than \$30M annually on servers for market, credit, scenario and regulatory risk computation. A cornerstone of J.P. Morgan's cost reduction plan to cut the cost of risk calculation involves accelerating its risk library. At the same time, JPMorgan has an aggressive plan to enable more accurate and frequent risk calculations.

SOLUTION

JPMorgan looked to move risk computation from conventional central processing units (CPUs) to graphics processing units (GPUs). The Bank's Equity Derivatives Group added NVIDIA® Tesla® GPUs to its datacenters to meet the demand of the increase in computational modeling. More than half the equity derivative-focused risk computations run by the bank have been moved to running on hybrid GPU/CPU-based systems.

NVIDIA Tesla GPUs were deployed in multiple data centers across the bank's global offices. Now, J.P. Morgan can seamlessly share the GPUs between dozens of global applications.

IMPACT

Utilizing GPUs has accelerated application performance by 40X and delivered over 80 percent savings, enabling greener data centers that deliver higher performance for the same power. For J.P. Morgan, this is game-changing technology, enabling the bank to calculate risk across a range of products in a matter of minutes rather than overnight.

GPU-ACCELERATED LIBRARIES

Take advantage of the massively parallel computing power of the GPU by using the GPU-accelerated versions of your existing libraries. Some examples include:

NVIDIA Math Libraries

A collection of GPU-accelerated libraries—including FFT, BLAS, sparse matrix operations, RNG, performance primitives for image/signal processing, and the Thrust C++ library of high-performance templated algorithms—that all deliver significant speedups when compared to CPU-only libraries. These highly optimized libraries are free of charge in the CUDA® Toolkit available at www.nvidia.com/getcuda

EM Photonics CULA

A GPU-accelerated linear algebra (LA) library that dramatically improves the performance of sophisticated mathematics.

ArrayFire

An easy-to-use GPU library of common functions for C, C++, Python, and Fortran applications. Includes hundreds of functions for math, signal/image processing, statistics, and more. Supports the popular "GFOR" for-loops running on GPUs.

RogueWave IMSL

A comprehensive set of mathematical and statistical functions for Fortran applications to take advantage of GPU acceleration.

MAGMA

MAGMA is a collection of next generation linear algebra (LA) GPU accelerated libraries designed and implemented by the team that developed LAPACK and ScaLAPACK. MAGMA supports interfaces to current LA packages and standards, e.g., LAPACK and BLAS, to allow computational scientists to effortlessly port any LA-relying software components.

GPUs TAKE MONTE CARLO OPTIONS PRICING REAL-TIME

In finance, an option or derivative is a contract that gives a firm the right, under certain conditions, to buy or sell an underlying asset or instrument. Financial firms use options to reduce the risks associated with investing. While options can add balance to a portfolio and limit exposure to potential threats, they are complex securities that must be approached with due diligence. Accurately calculating risk and pricing is a critical part of this strategy, and GPU co-processors have increasingly been commandeered for this computationally intensive task.

There are a number of computational algorithms common to the financial world. In finance pricing, Monte Carlo simulation is the dominant technique for American-style options, which may be exercised at any time before the expiration date. The Monte Carlo method provides the decision maker with a range of possible outcomes and the probabilities they will occur for any choice of action.

Julien Demouth, who works on the developer technology team at Nvidia, describes an implementation of American option pricing using Monte Carlo Simulation with a GPU-optimized implementation of the Longstaff Schwarz algorithm. This setup, developed in collaboration with partners IBM and the Securities Technology Analysis Center (STAC), was used to calculate a risk management benchmark in real-time on a single workstation with Tesla GPUs.

Nvidia worked with IBM and STAC to implement the STAC-A2 benchmark on two Nvidia Tesla K20X GPUs. The system calculated the risk and pricing of this particular complex option in less than a second, showing that risk management benchmarks, like STAC-2, can be operated in real-time, enabling them to be used ahead of actual trades.

STAC-A2 is a suite of benchmarks based on options Greeks, standard financial tools that measure the sensitivity of the price of the option based on a number of variables, such as changes in interest rates.

A key component of the Monte Carlo simulation for pricing American-style options is the Longstaff-Schwartz algorithm, a calculation that determines at which point the option should be exercised. Nvidia's STAC-A2 implementation initially used a hybrid method that relied on the CPU to perform the linear regression, but since then, Nvidia developers have come up with a new implementation that runs entirely on the GPU. To maximize the performance of the linear regression, Nvidia developers figured out a way to reduce the amount of data transferred between the GPU memory and the system main memory attached to the compute cores.

On a Tesla K40 GPU coprocessor, announced last November by Nvidia, the GPU-optimized algorithm prices an American option on 32,000 paths and 100 time steps in under 3ms. Complete time, including the paths generation, is less than 5.5ms.

Nvidia also explains the programming techniques used to obtain a very efficient code for the Andersen Quadratic Exponential (QE) path discretization, used in quantifying the random nature of the stock price and its volatility. Although Andersen-QE presents a challenge for efficient parallel implementations, Nvidia improved performance by finely tuning each of the branches of the code and moving as much computation outside of the block as possible.

According to results published by STAC, STAC-A2 benchmarks running on Nvidia Tesla K20x GPUs showed nearly an order of magnitude speed up compared to traditional X86 CPUs. For a test machine,



STAC used an IBM System x iDataPlex server with two eight-core Intel Xeon E5-2660 processors running at 2.20 GHz processors and two Nvidia K20Xm GPUs. The software stack was coded by Nvidia using the CUDA 5.5 toolkit. The system delivered over six times the average speed of the fastest publicly benchmarked system without GPUs. Furthermore, adding one or two GPUs to a system offers speedups just north of 5X and 9X, respectively, compared to the same system without GPUs.

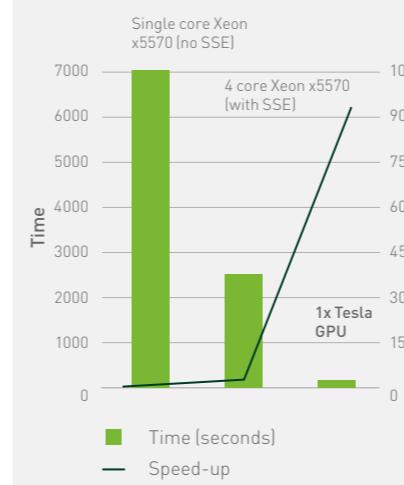
"Our STAC-A2 benchmark code and this implementation of the Longstaff-Schwartz algorithm both illustrate how Nvidia GPUs can make option pricing and risk analysis much faster," notes Demouth.

Julien Demouth,
NVIDIA Developer Technology Engineer

MUREX ANALYTICS ON TESLA GPUs

For twenty five years Murex has focused exclusively on developing enterprise wide-software platforms and technologies for the capital markets.

MUREX — SPEED-UP WITH TESLA GPU



www.nvidia.com/object/tesla-murex-analytics

The Murex trading platform enables trading and risk management across multiple asset classes for over 36,000 users at over 200 institutions in 65 countries.

MUREX MACS

Murex's clients are increasingly demanding massive acceleration for fast and accurate calculation of pricing and risk sensitivities. In response Murex has invested heavily as early as 2008 in targeting NVIDIA GPUs for high-performance computing using the Murex MACS analytics library. The MACS analytics library provides valuation models and greeks for complex structured securities along with VAR, PFE & XVA valuation for exotics & vanillas.

THE VALUE OF GPUs

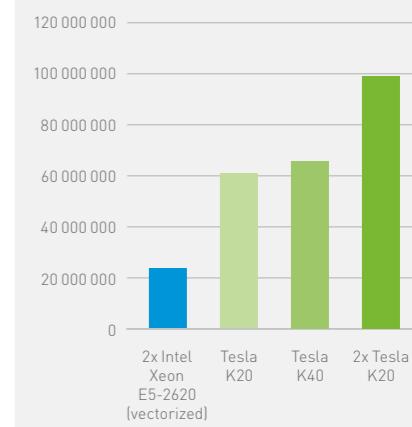
Murex was the first system of its kind in production with NVIDIA Tesla GPUs. Using NVIDIA Tesla GPUs, Murex has achieved



MUREX, SUNGARD,

SUNGARD®

VALUATIONS PER SECOND



EFFORTLESS GPU ACCELERATION FROM SUNGARD'S ADAPTIV ANALYTICS

SunGard's Adaptiv Analytics delivers incredible speed for financial risk and risk-based pricing calculations across the broadest range of financial asset classes and derivatives. Exceptionally engineered, Analytics grows with an organization's risk requirements — whether its volumes are increasing, or its appetite for sophisticated risk calculations.

For SunGard, the key to optimal performance is delivering greater calculation capabilities at a lower cost for our customers, by getting the best from the latest hardware platforms. A single server powered by NVIDIA Tesla GPUs can now deliver 120 million valuations per second when calculating credit risk — more than 100 times faster than pre-vectorized code.

www.sungard.com/adaptiv

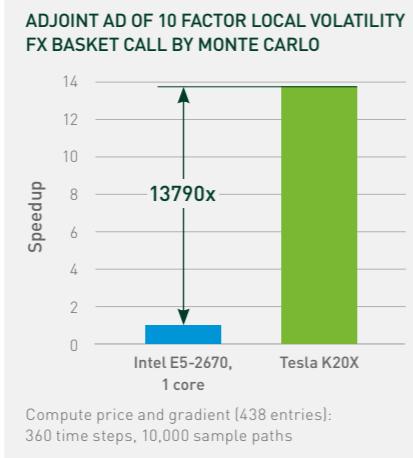
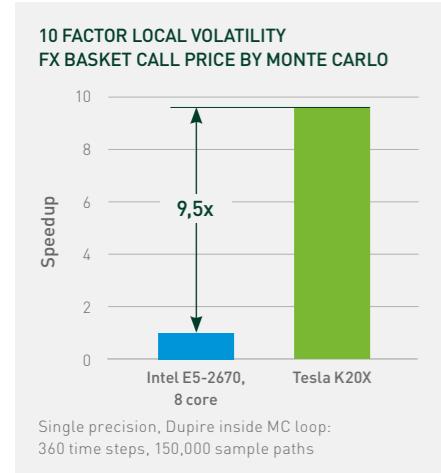
HANWECK ACCELERATES ANALYTICS AS A SERVICE WITH NVIDIA GPUs

The Hanweck Volera™ product line delivers real-time low latency analytics for trading and risk management. Using GPU-based high performance computing technology, Volera systems and services offer significant performance and cost advantage relative to traditional grid computing.

www.hanweckassoc.com

EXPERTS IN NUMERICAL SOFTWARE, SERVICES AND HIGH PERFORMANCE COMPUTING

NAG are world leaders in numerical software, services and high performance computing, having served academia and industry for over 40 years. NAG's numerical software underpins applications where cutting edge numerics are needed, and NAG services such as developing and optimising client numerical code, and the parallelisation of client serial codes help organisations keep pace with the latest technology advances.



www.nag.com



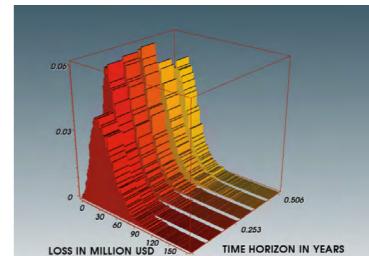
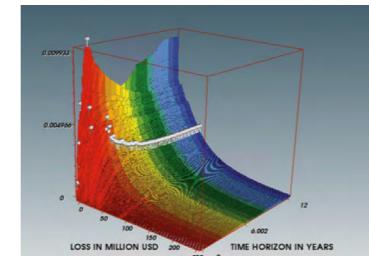
VoleraFEED™ delivers implied volatilities and risk parameters on the entire OPRA universe of equities and index options as a streaming, multicast datafeed. It features real-time implied volatilities, Greeks, and interpolated surfaces.

VoleraRISK™ computes real-time portfolio margin calculations, including Option Clearing Corporation's (OCC) TIMS profit and loss vectors, other risk arrays, and custom scenarios for pre and post-trade risk and margin surveillance.

GVL ESTHER WITH GPU ACCELERATION

Esther is GVL's flagship product. A general purpose valuation and simulation engine for large OTC bank portfolios, GVL Esther is packaged as a software-hardware-data solution and delivered as an appliance capable of running megamodels for CVA/DVA/FVA in addition to running margin and capital simulations.

Designed to tackle the single most massive and demanding computational tasks in financial engineering, GVL Esther is an in-memory application that executes on large heterogeneous boards with acceleration. Tailored to Supermicro equipment and accelerated by NVIDIA K10 GPUs, GVL Esther achieves a sustained performance of about 10 TF/sec on each board. A handful of servers in a compact cabinet enclosure is capable of processing workloads far in excess of what can be achieved by traditional grid computing solutions.



BENEFITS

- Consistent and realistic models provide more reliable results
- Numerical approach eliminates constraints of models with analytic solutions
- Comprehensive modeling of wrong way risk
- New hardware enables fast turn-around times and increased number of MC



HANWECK, NAG, GLOBAL VALUATION, MATLAB



ALGORITHMIC DIFFERENTIATION (AD) FOR GPUs

Algorithmic Differentiation is a sophisticated technique for obtaining accurate first- and higher-order mathematical derivatives of any computer program. It can be orders of magnitude faster than finite differences. NAG provides tools, expert help, support, training and coding to deliver robust and efficient solutions using this powerful technique.

NUMERICAL CODE FOR GPUs

The NAG Numerical Routines for GPUs contain a suite of Monte Carlo components such as Random Number Generators, Linear Algebra and a Brownian Bridge constructor. The routines are high performance, parallelised GPU kernels which give significant performance gains compared to traditional CPU systems.

MATLAB ACCELERATION ON TESLA GPUs

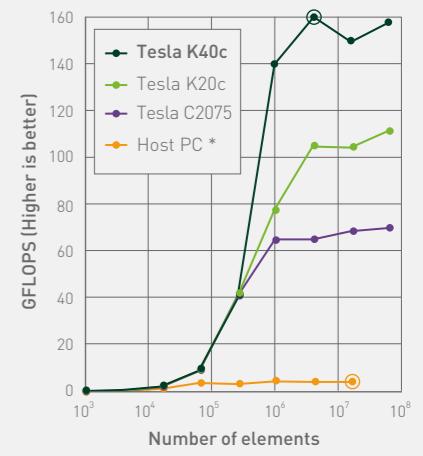
NVIDIA® and MathWorks have collaborated to deliver the power of GPU computing for MATLAB users. Available with the latest release of MATLAB, GPU acceleration enables faster results for users of the Parallel Computing Toolbox and MATLAB Distributed Computing Server. GPU support is also available in several domain-specific MATLAB toolboxes (including Signal Processing, Communications System, Phased Array, and Neural Network toolboxes). MATLAB supports CUDA® — enabled GPUs with compute capability version 1.3 or higher, such as K-series GPUs.

The latest release of Parallel Computing Toolbox and MATLAB Distributed Computing Server takes advantage of the CUDA parallel computing platform to provide users the ability to

- Manipulate data on NVIDIA GPUs
- Perform GPU accelerated MATLAB operations
- Integrate users own CUDA kernels into MATLAB applications
- Compute across multiple NVIDIA GPUs by running multiple MATLAB workers with Parallel Computing Toolbox on the desktop and MATLAB Distributed Computing Server on a compute cluster.



FFT (DOUBLE PRECISION)



www.mathworks.com

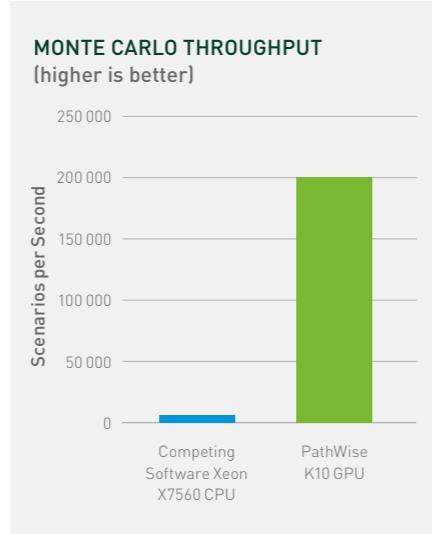
PATHWISE: HIGH PERFORMANCE INSURANCE MODELING AND RISK SOLUTIONS



“ We had ready access to risk information on a regular basis before PathWise, but now the information is refreshed frequently and we’re able to make more timely decisions. Many calculations that we would have done in hours or days are now done every few minutes — our information is up-to-date on a real-time basis and we’re not taking decisions on information that is stale.”

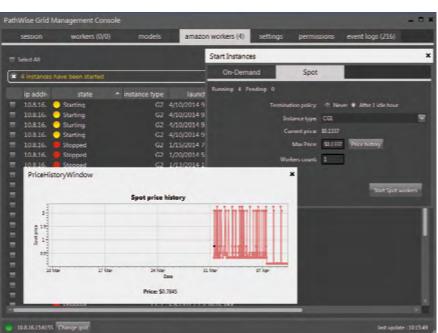
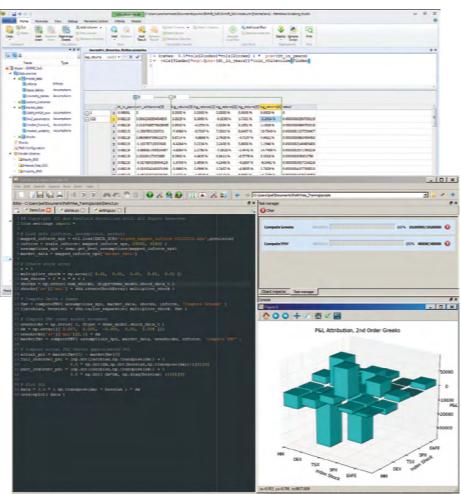
Martin Ettles
Senior Actuary,
Finance and Risk management,
Standard Life

PathWise delivers unprecedented performance and productivity to our clients in the life insurance industry who are increasingly providers of wealth management and retirement investment financial products. Actuarial computational workloads are growing rapidly, due to increasing need for better and faster modeling and forecasting of Capital, Reserves, Managed Volatility and Hedging Strategies for long-dated, embedded, financial guarantees in life insurance products. PathWise is used to develop, automate and run production processes for our clients in 5 different countries and using hundreds of NVIDIA GPUs, 24 hours per day, to support valuation, regulatory reporting, enterprise risk management and hedging activities.



www.aon.com

Learn more about Aon Benfield Analytics
bit.ly/1kPvlZG



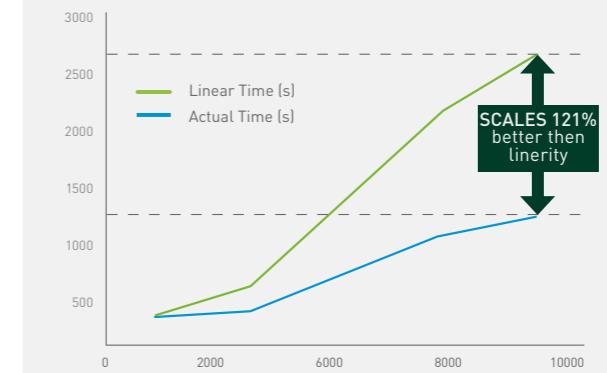
ULTRA-FAST PRICING AND DATA AGGREGATION WITH MISYS GLOBAL RISK TECHNOLOGY



Risk simulation and data aggregation are becoming increasingly time critical and embedded in many of daily risk management tasks. To stay competitive and alert to market changes, risk figures need to keep up more and more with the timeliness required in the front office. General-purpose processing on Graphical Processing Units (GPU's) are today's technical answer to the pricing challenge as it provides a large number of cores and fast memory access on a single card that can be added to a standard PC at low purchase and deployment costs. Together with NVIDIA, Misys developed Misys Parallel Technology to achieve the full potential of GPU pricing.

www.misys.com

AVERAGE PRICING TIME (SEC.) VS. NUMBER OF SCENARIOS, SCALABILITY



The chart below reflect super linearity of Misys Parallel Technology performance: the computation time with scenario numbers from 1000 up to 10000, everything else being equal. Further performance improvements can be achieved with more GPU's.

JEDOX: SELF-SERVICE BUSINESS INTELLIGENCE



Jedox enables business users to get rapid, ad-hoc insights into the economic and operational position of their company. Use Jedox for portfolio analysis, management consolidation, liquidity controlling, cash flow statements, profit center accounting, treasury management, customer value analysis and many more financial applications, all accessible in a powerful web application or your familiar Excel environment.

At the heart of the Jedox solution is a high-performance in-memory database engine, Jedox OLAP, boosted by NVIDIA Tesla GPUs. This database holds all relevant data in GPU memory and is thus an ideal application to utilize the Tesla K40's 12 GB on-board RAM. Scale that up with multiple GPUs and keep close to 100 GB of compressed data in GPU memory on a single server system for fast analysis, reporting and planning.

www.jedox.com

JEDOX GPU PERFORMANCE



Comparison of response times for a large analytical query on a CPU/GPU system with 2x Intel Xeon E5-2643, 256 GB RAM, 2x NVIDIA Tesla K20c, and 2x NVIDIA Tesla K40c. The printed numbers are speedup factors of the GPU variants over the CPU calculation ("GPU Boost" was tested in the 2x K40c variant).

These key technologies in the Jedox architecture enable Jedox users to quickly adapt and vary views, perspectives and

analytical dimensions. For its innovative GPU technology, Jedox was identified as a "Cool Vendor" in 2013 by Gartner.

REDUCE TIME TO MARKET AND DEVELOPMENT COST FOR YOUR NEXT GENERATION GPU APPLICATION

QuantAlea, a member of InCube Group, is a GPU pioneer in the financial service industry. QuantAlea's newest product Alea.cuBase is designed as a high productivity platform for GPU application development. It seamlessly combines the CUDA programming model with the new functional language F# and is unique in several ways:

- Alea.cuBase generates and compiles CUDA code at run-time, opening a whole range of new application scenarios
- Compiler and linker are fully integrated into the F# language and exposed as an API

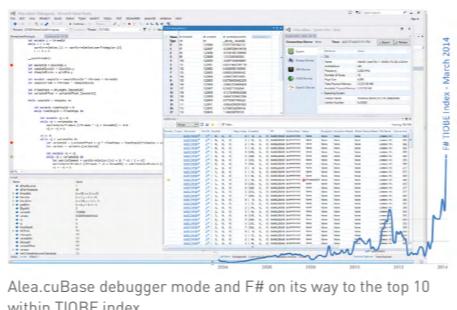
- A rapid prototyping environment with scripting and full GPU debugging support is provided
- Alea.cuBase is extensible allowing to build new domain specific languages on top of it

With Alea.cuBase and the outstanding tooling for F#, development becomes faster and less error-prone. Yet with all the benefits of the F# runtime environment, there is no performance penalty: the generated GPU code runs as fast as CUDA C/C++.

Clients in the financial service industry are building complete new pricing



libraries with GPU acceleration including on the fly payoff scripting, modern risk systems or sophisticated data analysis applications based on a single language, and runtime technology in projects with ambitious time-to-market constraints.



www.quantalea.net/products

OTHERS PROMISE NO PROGRAMMING. SCIFINANCE FROM SCICOMP DELIVERS.

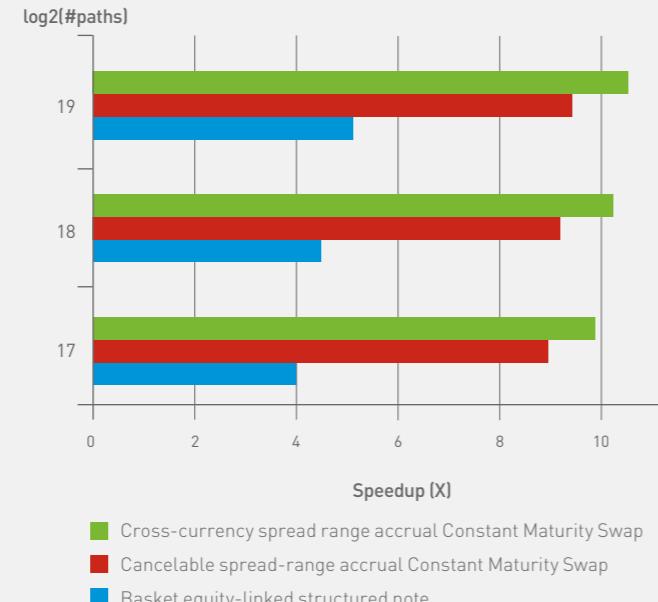
The SciFinance® paradigm, unique in the industry, eliminates programming by automatically translating derivatives model specifications into fully documented C-family source code in minutes. The SciFinance knowledge base has been extended to automatically generate NVIDIA CUDA-enabled source code for any Monte Carlo pricing model. PDE pricing models can be CUDA-enabled on a consulting basis. No hand-coding, CUDA programming or parallel computing expertise is required.

Not a toolkit, SciFinance provides infinitely customizable pricing models, complete model transparency and no run-time licenses. Specifications are written in an intuitive, finance-specific language that exactly describes the financial contracts and numerical methods.



TYPICAL CUDA ACCELERATION FOR SCIFINANCE PRICING MODELS USING QUASI MONTE CARLO SIMULATION

NVIDIA Tesla K20 vs. Dual CPU (Intel Xeon E5-2687W)



www.nvidia.com/object/tesla-scicomp

THE XCELERIT SDK: FAST ANALYTICS MADE EASY



“ With the Xcelerit SDK we can get the speedups we require, maintain a single codebase and run the executable on any of our in-house facilities. This is a very efficient way to cut infrastructure and development costs.”

Eurico Covas
Head of QRVG Development and Hedge Accounting Systems, HSBC



The award-winning Xcelerit SDK is a software toolkit to boost the performance of compute-intensive applications while preserving programmers' productivity.

From a modified sequential source code — free from any parallel constructs, low-level compiler directives, etc. — the Xcelerit SDK can generate highly optimised code for a

variety of target processors. These include multi-core CPUs, GPUs, and combinations of these in a grid. The programming interface is intuitive and suitable for many classes of algorithms, e.g. Monte-Carlo, Finite-difference, or Lattice-based methods.

The Xcelerit SDK is targeted at Quantitative Analysts and Quantitative

www.xcelerit.com/xcelerit-sdk

HYBRIDIZER: DEVELOP IN .NET OR JAVA DEBUG AND EXECUTE ON GPU



The Hybridizer is your in-house GPU expert. This productivity tool generates vectorized C++ source code [AVX] and CUDA-C source code from .NET assemblies (MSIL) or Java archives (java bytecode).

In managed development environments, use freely virtual functions and generics, yet make use of most of the hardware compute capabilities, with 80% usage of peak performance of processors and memory.

From a single version of the source code, debug and execute on both AVX CPU and CUDA GPU, within your favorite development environment stepping into original source code (.NET or Java). Profile your application using state of the art solutions (VTUNE / NSIGHT) pointing to locations in your original source code.

Whether you are an IRD trading firm, an investment bank or a clearing house, these numbers mean that you

can substantially reduce your IT costs while still coping with increasing demand from the regulators. Observed

dramatic performance improvements, x69 in the best case, also apply to more complex use cases.

CASH FLOWS DISCOUNTING FOR RISK PROFILE

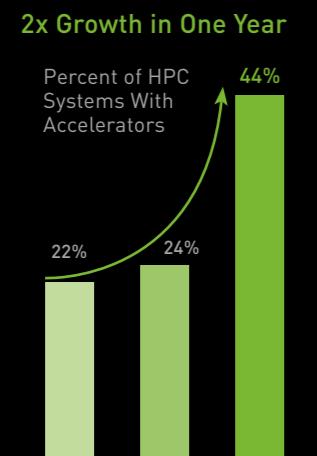
Linear Interpolated Stochastic Rate



Single version of source code x4 with Hybridizer-AVX—same hardware, x16 to x64 with Hybridizer-CUDA on NVIDIA Kepler GPU.
GPU is NVIDIA K20c, CPU is Intel Core i7-4770S @ 3.1 GHz

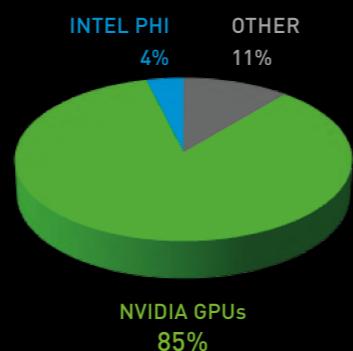
www.altimesh.com
contact@altimesh.com

ACCELERATED COMPUTING GROWING FAST



Intersect360 Research
HPC User Site Census: Systems, July 2013

NVIDIA GPU is
Accelerator of Choice



2
USEFUL LINKS

NVIDIA TESLA GPUs EMPOWER THE DATACENTER

INCREASE OVERALL SYSTEM PERFORMANCE AT LOWER COST AND REDUCE POWER

From the largest supercomputing centers to leading enterprises worldwide, IT managers are adding NVIDIA® Tesla® GPUs to their datacenters to cost-effectively increase performance.

GPUs ALLEVIATE PRESSURES IN THE DATACENTER

The increasing demand from business units for more compute power to perform pricing and risk analytics is posing real challenges to IT resources, with static or declining budgets, maxed out datacenters, and limited power availability.

- GPU computing, where CPUs and GPUs are combined in a server node or blade form factor, delivers massive compute power per dollar & per watt.
- GPUs accelerate large computationally intense applications that are

"embarrassingly parallel", such as Monte Carlo simulation and large matrix multiplication algorithms. Speedups of 200x – 300x can be achieved with GPU/CPU computing, as compared to a single x86 core.

- When parallel applications are migrated to GPUs, the load on the overall datacenter is reduced, resulting in lower power and cooling costs, lower system maintenance charges, smaller footprint, and lower capital cost.
- NVIDIA® Tesla® GPUs protect all internal and external memories with ECC and support IEEE 754-2008 floating point standard.

NVIDIA TESLA AND CUDA LINKS

LINKS:

NVIDIA GPU computing for Finance:
www.nvidia.co.uk/finance

NVIDIA Tesla solutions:
www.nvidia.eu/tesla

Request a CUDA Course Program
www.parallel-computing.pro/downloads/courses.pdf

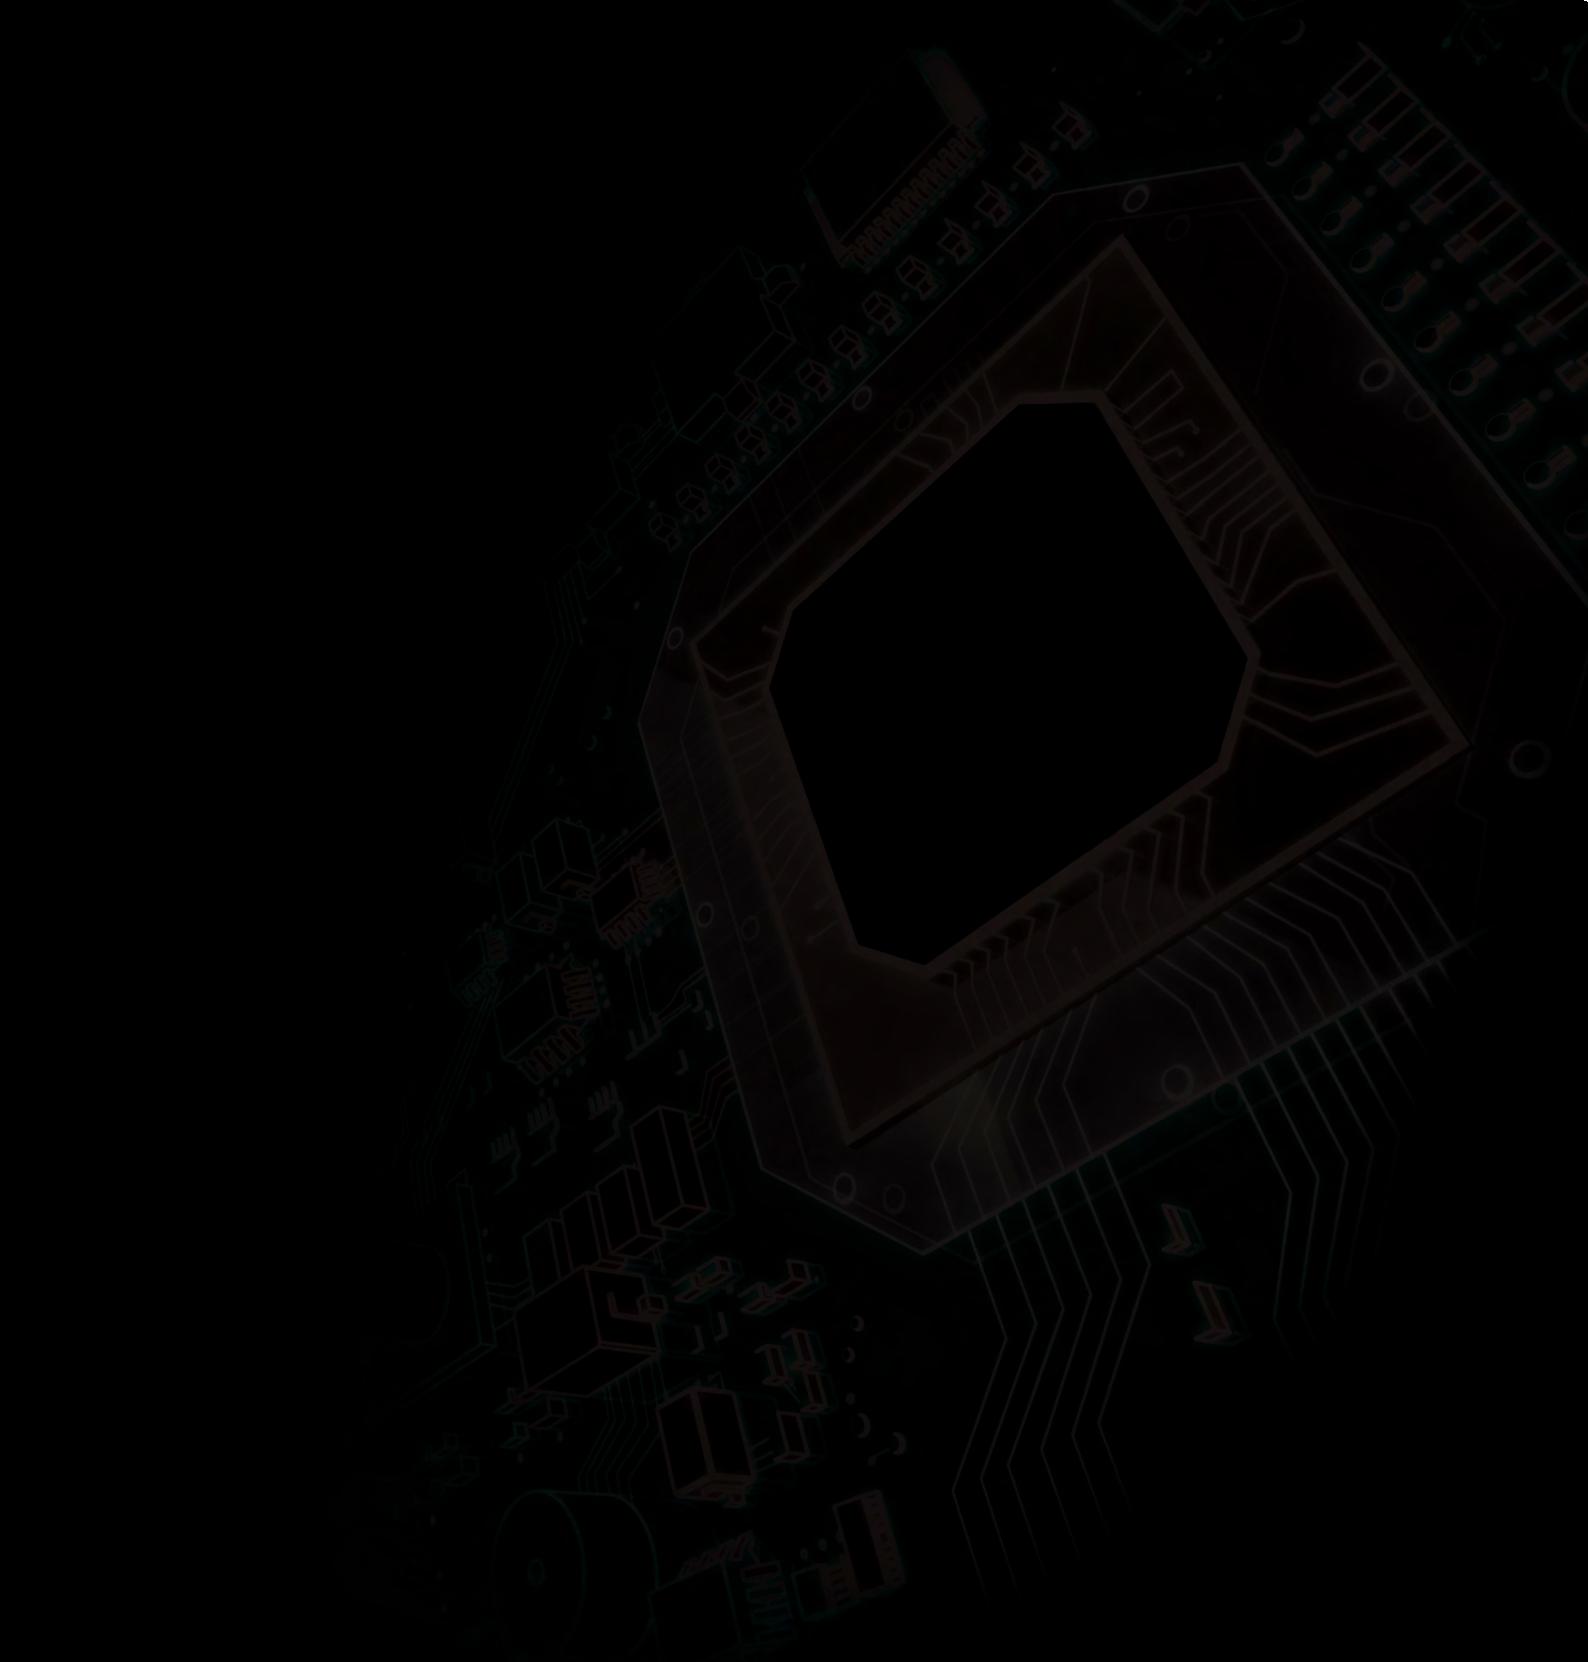
CONTACTS:

Alastair Houston
NVIDIA Compute FSI Sales Manager
Tel: +44 7802 213914
Email: ahouston@nvidia.com

Jeff Sporn
Business Development Manager
Financial Services, North America
Tel: 908 281 9180
Email: jsporn@nvidia.com

Simon See
Director/Chief Solutions Architect, APAC
Tel: 659 619 1816
Email: ssee@nvidia.com

John Ashley
Solutions Architect,
North America and EMEA
Tel: +44 07788 242104
Email: jashley@nvidia.com



© 2014 NVIDIA Corporation. All rights reserved. NVIDIA, the NVIDIA logo, Tesla, Kepler, and CUDA are trademarks and/or registered trademarks of NVIDIA Corporation. All company and product names are trademarks or registered trademarks of the respective owners with which they are associated. Features, pricing, availability, and specifications are all subject to change without notice.

