



STAC[®] Summit

June 4, 2014

Doors open: 9:30am

Meeting starts: 10:00am


Hyatt Andaz Hotel

Great Eastern Room

40 Liverpool Street

London

Platinum Sponsor:



Gold Sponsors:



DataDirect
NETWORKS



AGENDA

Note: Times are approximate.

STAC Update for STAC-A2 (risk computation) [\[slides\]](#)

- *Peter Lankford, Founder & Director, STAC*

The STAC-A2 Benchmark suite is the industry standard for testing technology stacks used for compute-intensive analytic workloads involved in pricing and risk management. Peter will provide a brief update on the latest learnings and the latest activities of the STAC-A2 Working Group.

Advances in C++ Parallelization [\[slides\]](#)

- *Robert Geva, Principal Engineer, Intel*

Parallelization continues to be a focal point for C++ programmers in finance. Some have not yet achieved parallelism in their code, while others debate the best ways to do so. As the leader of Intel's team that helps quants improve their application performance, Robert understands the challenges, benefits, and tradeoffs of parallelization. In this talk, Robert will provide new data on the benefits of parallelism in financial applications, as well as recent insights on best practices in parallelizing specific financial computations.

Implementing PDE solvers on GPUs [\[slides\]](#)

- *Mike Giles, Professor of Scientific Computing, Oxford University*

Noting that GPUs have already been shown to be very efficient for Monte Carlo simulations, Mike will describe research examining GPUs in PDE option pricing models. This work considered both 1-factor and 3-factor models, and both explicit and implicit time-marching numerics. The 1-factor implementation was compute-limited, achieving up to 85% of the GFlop peak, whereas the performance of the 3-factor version was more limited by the bandwidth between the GPU and graphics memory, running at about 50% of the peak bandwidth on average, and up to 40% of the GFlop peak. The key to the implicit solver implementation was the use of a hybrid Thomas / Parallel Cyclic Reduction algorithm for solving tridiagonal systems, and care in the transfer of data to/from the graphics memory to maximise the achieved bandwidth.

Innovation Roundup Presentation

"Static code analysis: Finding flaws as they're typed" [\[slides\]](#)

Scott Lasica, VP, Field Technical Services, Rogue Wave Software

Julia: Why all the buzz? [\[slides\]](#)

- *Avik Sengupta, Head of Algorithmic Solutions, Decimal Point Analytics*

The Julia programming language seems to have come out of nowhere and attracted great interest in the computational finance community. And why wouldn't it? It's MIT creators say they sought to provide an open source language with the best features of Matlab, R, Hadoop, and shell scripts, with the performance of C. "We want to write A*B and launch a thousand computations on a thousand machines, calculating a vast matrix product together," they said. Those are no small ambitions. What are the realities? How does Julia work, and what are its current strengths and limitations? How does its performance really compare to Matlab, R, Python, and C? Who are the main players in the ecosystem, and what is the roadmap? As an enthusiast and major contributor to Julia, as well as a quant programmer experienced with other languages, Avik will take us on a quick tour of Julia and touch on each of these questions.

OpenSPL - A new way of thinking about FPGA [\[slides\]](#)

- *Ryan Eavy, Executive Director, Architecture, CME Group*

The benefit of FPGAs for certain tasks in financial markets is now widely accepted. But as we've discussed at many STAC Summits, one of the challenges slowing FPGA adoption is programmer productivity. A consortium of firms led by the CME, Chevron, and Juniper have set out to change that with a programming framework called the Open Spatial Programming Language, or OpenSPL. OpenSPL is an effort to standardize a "spatial programming" approach for FPGAs and other computing technologies. Whereas traditional programming techniques focus on execution in just one dimension (time), a spatial programmer thinks of execution in two dimensions: time and space. In this talk, Ryan will provide a brief overview of OpenSPL, illustrating how it eases programming of devices like FPGA.

Networking Luncheon

STAC Update: Network research [\[slides\]](#)

- Peter Lankford, Founder & Director, STAC

Peter will review the latest STAC activities related to low-latency networking.

Not Just Server Consolidation: Virtualization for Big and Fast Workloads [\[slides\]](#)

- David Riddoch, Chief Architect, Solarflare

While virtualization dominates the enterprise data center, it has been unrealistic for Tier 1 financial applications due to performance limitations. That may be about to change. In this talk, David will examine some emerging cases in which applications achieve near-native performance by leveraging the latest virtualization technologies such as SR-IOV, multiple PFs and VFs, PCI passthrough, and a preview of Linux containerization that clears a path toward more elegant integration of 10/40GbE, Linux and KVM into OpenStack Clouds. As part of this, he will show how running IEEE 1588 (PTP) in a VM and leveraging ingress and egress time stamping enables real-time latency measurements without external equipment.

Innovation Roundup Presentations

"Achieving ultra low latency with the Exablaze ExaNIC network card family" [slides]	Joshua Rose, Director of Commercialisation, Exablaze
"EMC ScaleIO for Financial Data Management Applications" [slides]	Simon Stevens, Advisory Systems Engineer, ScaleIO UK&I, EMC Advanced Software Division
"Getting the most out of multi-year and multi-source trading history" [slides]	Glenn Wright, Systems Architect, DataDirect Networks

STAC Update - Tick data and backtesting [\[slides\]](#)

- Peter Lankford, Founder & Director, STAC

Peter will summarize the latest STAC Benchmark Council activities in areas relating to backtesting and tick data, including STAC-M3, STAC-A3, and research on flash storage.

Coffee Break

Big data use cases in investment and retail banking [\[slides\]](#)

- Peter Lankford, Founder & Director, STAC
- Introduction by Gordon Hughes, Global Sales Director, Financial Services Alliances, Intel

Investment and retail banking often appear near the top of the list of industries investing in "big data" technology. Yet information about how banks are using that technology is sparse. With the help of seed funding from Intel, STAC recently worked with several of the global banks in the STAC Big Data Special Interest Group to identify concrete use cases that pose big data challenges. By interviewing staff with direct knowledge of these cases, we were able to characterize the workloads involved and understand the business problems that arise with traditional technologies. We were also able to learn about the advantages and challenges of the new approaches that banks were taking to these workloads. After an introduction from Gordon, Peter will review the use cases and the major themes that have emerged so far. He will also discuss how these cases and themes are influencing potential STAC Benchmarks in big data domains.

Big data for Big Time Workloads [\[slides\]](#)

- *Amy McNee, VP for Global Customer Operations, DataStax*

Most transactions in finance today are automated, requiring computerization straight through—from the customer interface to the transaction handling systems to the audit trails and analytics. As financial institutions grapple with growing data volumes, tougher analytic challenges, and stubborn cost pressures, many of them are turning to open source "big data" software for data management. However, in Amy's view, most of these new products are not designed for the continuous uptime required by institutional and retail transaction flows. This marginalizes big data technologies, restricting them to side activities where failures can be tolerated but the payoffs to new software are not as great. In this session, Amy will discuss the principles of a data architecture that can marry the scalability and simplicity of NoSQL with the uptime and security that banks and exchanges take for granted in their legacy systems. In the process, she will present case studies of organizations that have seen the benefits of such a design.

Non-volatile RAM: An insider perspective [\[slides\]](#)

- *David Byrne, EMEA SSD Business Manager, Intel*

The non-volatile memory industry is changing rapidly. The cost of NAND flash memory continues to drop. Vendors are packaging flash in a wide range of form factors and interfaces, from rackable arrays to PCIe cards and even DIMMs. New standards based programming models are emerging, such as NVMe and the NVM Programming Model. New scale-out file systems are replacing traditional SAN and NAS architectures in enterprise IT. Meanwhile, new technologies on the horizon have some pundits predicting a convergence of memory and storage that will overturn the very Von Neumann computing model itself. At Intel, David Byrne runs Intel's non-volatile memory solutions business in EMEA. He'll sort out the space for us and provide some pointers for architects and application developers. What are the key locations on the cost-vs-performance graph? How can performance-critical applications best take advantage of flash today (memory channel? kernel bypass? NoFTL?)? How can solution designers best prepare for the future?

You Snooze you Lose: Lessons from Real-Time Bidding [\[slides\]](#)

- *Brian Bulkowski, Founder & CTO, Aerospike*

Much like financial trading, the online ad market is now dominated by automated decision making. Every time a user opens a web page with ad space, computers from multiple firms engage in an enormous amount of communication, analysis, and bidding behind the scene to decide which company's ad will make it onto the page—all within milliseconds. Similar to many processes in a bank, realtime bidding platforms require a data architecture that can handle terabytes of data and hundreds of thousands of transactions per second. How do the leading platforms achieve this and turn insights into action? Brian will describe his experiences building and operating modern ad-tech platforms, combining an in-memory NoSQL database for predictable high performance with Hadoop based analytics platforms in the back, an architecture that is rapidly becoming the gold standard in the Age of Context.

Apache Spark as Cross-over Hit for Data Science [\[slides\]](#)

- *Sean Owen, Director of Data Science, Cloudera*

Apache Spark is a compelling multi-purpose platform for use cases that span both investigative and operational analytics. This talk will briefly introduce Spark and how it fits into the world of existing data science tools, and how it can speed up computations dramatically with effective use of in-memory caching.

Networking Reception

Speaker Biographies – Feature Sessions



Brian Bulkowski, Founder & CTO, Aerospike. Brian is a founder of Aerospike, CTO & Product, networking whiz, innovator and high performance expert. 'My family has a long and varied history in science and tech, so I wound up shipping code in high school. One of the great things about software is you can build something with it. You don't need plywood or welding. Computers are an easy way to start creating stuff. My first taste of networking was in 1989. I knew there was a whole world out there waiting. A computer that's not connected to a network is kind of dull.' Brian became a Lead Engineer at Novell, and then Chief Architect of Cable Solutions at Liberate – where he built a high-performance, embedded networking stack, as well as the high scale broadcast server infrastructure. As Director of Performance at Aggregate Knowledge, Brian had direct experience with the scaling limitations of sharded

MySQL systems. 'It wasn't hard to see that there was a huge need for a new distributed database, because they all sucked. Everyone was struggling with what was available. That led to the idea for Citrusleaf – which then became Aerospike.' When he's not busy creating stuff without plywood or a welding torch, Brian plays cello in a band called Rosincoven. He also writes about cuisine for the San Jose Metro."




David Byrne, EMEA SSD Business Manager, Intel. David is the Business Manager for Intel's Solid State Drive products in Europe Middle East and Africa. In this role he is responsible for growing the sales of Intel Solid State Drive products in the EMEA market through Intel extensive network of partners and channels to market. He joined Intel in 1995 and has held a wide variety of positions in the sales and marketing organisation. Including managing the marketing of the Intel CPU business to local manufacturers and managing a software as a service product line. 5/2014



Ryan Eavy, Executive Director, Enterprise Architecture, CME Group. Ryan leads the Emerging Technology function as part of the Enterprise Architecture team at the CME Group. He is responsible for ensuring the company's technology is aligned with business strategy through application research and development, technical solution assessment and systems roadmap definition. Among Ryan's many other accomplishments, he has recently led the design and implementation of the Exchange's ultra-low-latency middleware strategy and common abstract middleware layer. Prior to this, he most recently served as a Senior Application and Security Architect for the Chicago Board of Trade. Before joining the CBOT in 2003, Ryan worked as a Senior Applications Developer at Quotes Plus where he designed and

developed trading analysis software and as a Senior Applications Engineer at Dewpoint. Ryan earned a bachelor's in computer science from the University of Michigan and is currently pursuing a master's in Computational Finance from DePaul University.



Robert Geva, Principal Engineer, Intel. Robert is a principal engineer at Intel's software and services group. His current role is the manager of the financial services engineering group, where he leads a team of SW engineers working with customers in FSI. Robert joined Intel in 1991 and has since developed an expertise in compilers and performance analysis and tuning for microarchitectures. Robert has worked on compiler optimizations for a variety of Intel microprocessor based systems, including the 80486, the Pentium Processor, the Pentium Pro Processor, Itanium, the Pentium 4 and Pentium M and core II Duo. Most recently, Robert was an architect in the development products division responsible for driving language extensions and programming models for parallel and heterogeneous programming. Robert has been involved with the development of Intel Cilk™ Plus and the offloading model for Intel® Xeon®

Phi™. Robert has BA and MSc from the Technion, Israel institute of technology.



Mike Giles, Professor of Scientific Computing, Oxford University's Mathematical Institute. Mike Giles is Professor of Scientific Computing in Oxford University's Mathematical Institute. He is also a founding member of the Oxford-Man Institute of Quantitative Finance, and an Associate Director of the Oxford e-Research Centre where he leads the CUDA Centre of Excellence. Half of his research is in the development and numerical analysis of improved Monte Carlo methods for a range of stochastic applications in finance, engineering and science. The other half is in high performance computing with an emphasis on many-core GPU computing, developing numerical libraries and high-level abstractions with code generators to enable others to obtain the benefits more easily.



Gordon Hughes, Global Sales Director, Financial Services Alliances, Intel. Gordon is responsible for Intel's worldwide strategy for eco-system engagements in the company's focus areas in banking and capital markets. Gordon has been with Intel for over 8 years and has a pedigree in Alliance Management with both software and hardware vendors, including Tandem Computers and Informatica. He holds a B.Sc (Hons) in Applied Mathematics from the University College of North Wales.



Peter Lankford, Founder & Director, Securities Technology Analysis Center. Peter leads STAC®, which provides hands-on technology research and testing tools to the finance industry and facilitates the STAC Benchmark Council™, a group of leading financial institutions and vendors that engages in technical dialog and specifies standard ways to assess technologies used in finance. Prior to STAC, Peter was SVP of the \$240M market data technology business at Reuters. Prior to Reuters, Peter held management positions at Citibank, First Chicago Corp., and operating-system maker IGC. Peter has an MBA, Masters in International Relations, and Bachelors in Chemistry from the University of Chicago.



Amy McNee, VP for Global Customer Operations, DataStax. Amy is a twenty-year technology veteran and has worked for companies like GE, SAP, IBM, Cognos and most recently, DataStax. Amy has held a variety of roles, including application development, consulting, presales and strategy. Currently, Amy is the VP for Global Customer Operations for DataStax leading a global team across presales, consulting, training and customer advocacy functions.



Sean Owen, Director of Data Science, Cloudera. Sean is Director of Data Science at Cloudera, based in London. Before Cloudera, he founded Myrrix Ltd (now part of Cloudera), a company commercializing large-scale real-time recommender systems on Apache Hadoop. He has been a primary committer and VP for Apache Mahout, and co-author of Mahout in Action. Previously, Sean was a senior engineer at Google. He holds an MBA from the London Business School and a BA in Computer Science from Harvard.



David Riddoch, Chief Architect, Solarflare. David co-founded Level 5 Networks in July 2002 and joined Solarflare when it merged with Level 5 in April 2006. David is the architect and lead developer of Solarflare's market leading OpenOnload network acceleration middleware. David's mission is to deliver absolutely the best possible performance without asking users to abandon the standard network stack: Sockets, TCP, UDP and Ethernet.



Avik Sengupta, Head of Algorithmic Solutions, Decimal Point Analytics. Avik has developed applications for investment banks in the areas of pricing, risk and trading for over a decade. He is currently a co-founder at AlgoCircle, a firm engaged in building machine learning and NLP solutions for the financial sector. Two years ago he discovered Julia, and hasn't looked back since. He is a contributor to the core language and the maintainer of many Julia packages.
