

AGENDA

Note: Times are approximate.

8:30 am Morning welcome

- *Peter Lankford, Founder & Director, STAC*

8:35 am IIT Mini-Symposium on Algo Trading Management

The Center for Financial Markets at the IIT Stuart School of Business conducts research into important aspects of high-frequency trading by combining technical and quantitative expertise with foundational concepts of modern business. In this mini-symposium, IIT researchers will present three new studies of interest to quants, technologists, and managers in the financial markets as well as other researchers.

- *Implementing Real Time Control in Algorithmic Trading, by Ben Van Vliet, Assistant Professor.* In principle, the obligation is on the trading firm to ensure its trading systems run in control at all times. But what does it mean for a trading system to run in control? We are looking to define signals that indicate there is a probability that the trading strategy is not running according to expectations. We review the ways in which reference distributions of expected performance of critical characteristics can be defined. We develop a simple model for real time statistical process control that allows for market state-varying performance. If supervisory mechanisms audit compliance with a principle-based standard that mandates control, our model should satisfy both market safety and institutional flexibility. Both regulators and the markets will be better off by harmonizing regulation around such an international standard of control.

- *Does HFT benefit LFT?, by Ernie Li, Adjunct Faculty, and Rick Cooper, Assistant Professor.* We examine how high-frequency trading (HFT) affects low-frequency trades (LFT) and the fundamental information they transmit to prices. We find that HFT enhances the flow rate of fundamental information by increasing the size and decreasing the waiting time of LFT limit orders, and increasing the size and frequency of LFT market orders. HFT serves as an intermediary by transporting liquidity to LFT orders, thus replacing market makers as the primary source of liquidity. By examining the interaction effects between market and limit orders, we find that low-frequency market orders take liquidity and lower the flow rate of low-frequency limit orders, and vice versa. This interaction effect shows that the major source of liquidity loss is not high frequency trading but rather low-frequency trading.

- *Performance Attribution in Automated Trading, by Rick Cooper.* Performance attribution is an important component of monitoring trading strategies. We present an intuitive absolute-return-based attribution framework for the volatility and information ratio. The results are valid for strategies at high frequencies given the appropriate selection of measurement and decision frequencies. Derivatives positions are integrated into the attribution primarily by breaking their return into effects due to movements in the underlying equities and arbitrage effects. For derivative positions held longer, a time effect is also introduced.

10:00 am Innovation Roundup – Round 1

“Static code analysis: Finding flaws as they’re typed”

Scott Lasica, VP, Field Technical Services, Rogue Wave

“Getting the most out of multi-year and multi-source trading history”

Michael Chazot, Director, Strategic Accounts, DataDirect Networks

10:20 am STAC Update - Tick data and backtesting

- *Peter Lankford, Founder & Director, STAC*

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

10:30 am Point of View: A faster way to get strategies to market

- *Bryan Lewis, Chief Data Scientist, Paradigm4*

Trading firms today require faster cycle times for development and backtesting of trading strategies. Bryan believes that a key to this is finding more agile and less code-intensive approaches using high-productivity languages like R and Python. He believes the other key is to use shared-nothing architectures on the back end that exploit hundreds or thousands of commodity hardware nodes in a cloud or local cluster. The crucial design imperative for large data is to minimize data movement by moving computation to the data. Many SQL and Hadoop solutions fall short when trying to perform strategy development and backtesting directly in the database. Even good distributed SQL databases often end up largely requiring ETL to an analytics environment like R or Python because of lack of in-database analytics. And Hadoop solutions typically sacrifice powerful data management concepts like join. As a consequence, developers are compelled to learn and use multiple toolsets, losing considerable time moving data among them. Taking the open source SciDB database as an example, Bryan will argue that it's possible to use an array language like R with inexpensive, scale-out hardware to create a single environment for tick storage, strategy development and backtesting.

10:50 am Point of View: Putting the Big Data ecosystem to work for tick data

- *Michel Debiche, Financial Services Practice Lead, ThinkBig Analytics*

As a quantitative trader for over 20 years, Michel built one backtesting platform after another for a series of trading firms. Each time, the requirements were essentially the same. And each time, the platform had to be built largely from scratch. With the advent of "Big Data" technologies like Hadoop, he suddenly felt liberated. In his view, combing through huge volumes of historical data to search for signals and backtest trading strategies is a natural match for these new platforms. Not only because the map-and-reduce paradigm is a close fit to requirements, but also because these tools make it easy to integrate a huge range of content, from market data and news to fundamental data, research reports, and more. He believes that such scalable, open source tools also open up analytic opportunities for groups outside the front office. Risk, compliance, technology, and security groups are among those who must test assumptions and scenarios, and then use the intelligence gleaned from historical research to effectively monitor developments in real time across many dimensions. In this talk, Michel will start from first principles to explain how the "Big Data" ecosystem maps to requirements. He will then address some of the key questions for quants and engineers: what's it like for the developer? how does it scale for different sorts of backtests? when does it make sense to combine "Big Data" with "Cloud Computing" in financial services?

11:10 am Non-volatile RAM: An insider perspective

- *James Myers, Director of SSD Solutions Architecture and Engineering, 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, James runs a team responsible for making non-volatile memory solutions work for customers. In this talk, he will 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?

11:30 am Networking Luncheon

12:30 pm *Financial technologists: Be proud. Be proactive.*

- *Peter Nabicht, Senior Advisor, Modern Markets Initiative*

HFT and other forms of automated trading have become favorite whipping boys for the media. Most trading firms, quiet by design, avoid the debate. Today the Modern Markets Initiative has taken up the task of explaining to the public the crucial role that computerized decisions play in providing fair prices to investors. Peter will argue that automated traders-as well as the technologists who enable them-should be proud of the important functions they perform. But he will also argue that everyone in a trading firm has a role to play in how the industry is perceived and how it evolves. After reviewing key issues in our complex market structure and how tools like MIDAS may shape the way that regulators view those issues, Peter will suggest things that technologists can do to promote a bright future for automated trading.

1:15 pm *STAC Update - Compute-bound workloads*

- *Peter Lankford, Founder & Director, STAC*

Peter will review the latest STAC activities in areas such as options pricing.

1:20 pm *Practical Parallelization in C++*

- *Arch Robison, Senior 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 original architect of Intel's Threading Building Block (TBB) library, Arch understands the challenges, benefits, and tradeoffs of parallelization. In this talk, Arch 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.

1:40 pm *Accelerating Stochastic Volatility Model Calibration in R Using GPUs*

- *Matthew Dixon, Assistant Professor of Analytics, University of San Francisco*

The R statistical software environment is popular with quants partly because it facilitates application design space exploration. However, despite offering a rich set of native numerical and statistical functionality that is useful for options-related analytics, a typical R based implementation of a stochastic volatility model calibration on a CPU does not meet the performance requirements for sub-minute level trading, i.e. mid- to high-frequency trading. Calculations rely on native optimization packages that are difficult to precisely replicate outside of R but are bottlenecked on computationally intensive kernels that can easily be replicated. In this talk, Matthew will present a new R library that that dramatically accelerates options-related calculations by offloading the most computationally intensive part of the volatility model calibration to a GPU via an easy-to-use map-reduce interface. He will provide evidence that the new R-based implementation performs comparably to GPU-based calibration code written in C/C++.

2:00 pm *STAC Update - Low latency research*

- *Peter Lankford, Founder & Director, STAC*

Peter will review the latest STAC activities in areas such as low-latency networking and order routing.

2:05 pm Innovation Roundup – Round 2

“Start Fast & Stay Fast: Priming Java for Market Open with ReadyNow!”	<i>Gil Tene, CTO, Azul Systems</i>
“MetaConnect: Proven performance.”	<i>Dave Snowdon, Founder, co-CTO, Metamako</i>
“Trading Without a Switch/Trading Within a Switch”	<i>Robert Walker, CTO, xCelor</i>
“Achieving ultra low latency with the Exablaze ExaNIC network card family”	<i>Dr. Matthew Chapman, CTO, Exablaze</i>

2:30 pm Not Just Server Consolidation: Virtualization for Big and Fast Workloads

<ul style="list-style-type: none"> <i>Matthew Knight, Marketing Director Financial Services, Solarflare</i> <p>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, Matthew 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.</p>

2:50 pm Innovation Roundup – Round 3

“New Developments in Exchange Latency Management”	<i>Tony Pettipiece, Global Head of Sales and Marketing, Cape City Command</i>
“Where global business and growth connect”	<i>Henrique Hablitschek, Business Development Manager, Equinix</i>
“Tackling 'big data' capture challenges”	<i>Mohammad Darwish, VP Sales, Fiberblaze</i>

3:15 pm OpenSPL - A new way of thinking about FPGA

<ul style="list-style-type: none"> <i>Ari Studnitzer, Managing Director, Platform Development, CME Group</i> <p>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, Ari will provide a brief overview of OpenSPL, illustrating how it eases programming of devices like FPGA.</p>

Networking Reception