

Navigating the Future of Smart Order Routing



Sponsored by:

smartTrade  TECHNOLOGIES

SUNGARD[®]

In this report

- S2 Editor's Letter: Smart Order Routing Grows Up
- S3 Expanding the Frontiers of Smart Order Routing
- S4 Roundtable: A Peek at the Next Generation
- S6 Smart Order Routing in the News

Smart Order Routing Grows Up

It's hard to believe that smart order routing (SOR) is approximately a decade old and, like anyone approaching its teenage years, is about to go through a major growth spurt in several directions. Original routers only had to contend with the displayed U.S. cash equities market and settle and clear trades with National Securities Clearing Corp. (NSCC). Today, the challenge has grown in complexity as smart order routing has moved into the European markets, which trade in multiple currencies and clear and settle in a host of different clearinghouses. Toss in the intelligence needed to estimate where non-displayed liquidity can be found and it becomes a daunting task to keep the smart order router at peak performance.

Looking not too far down the road, all of these complexities will grow exponentially as support for multi-asset trading gets rolled into the routing algorithm as traders look for the best hedge as well as the best execution.

In this report, we speak with ITG's Mark Wright, who discusses where he sees the evolution of SOR is taking his firm from both a product and a technology perspective. *DWT* also quizzes Harry Gozlan, founder and CEO of Smart Trade Technologies; and Vincent Burzynski, chief product officer with SunGard Global Trading, on what the next generation of smart order routers will look like.



Rob Daly
Editor

Expanding the Frontiers of Smart Order Routing

Once alternative trading systems (ATSes) started to pop up in U.S. cash equities markets in the late 1990s with the introduction of the U.S. Securities and Exchange Commission's (SEC's) Regulation ATS and changes to its order handling rules, smart order routing (SOR) has been a fact of life for global agency broker Investment Technology Group (ITG).

Once there were better quotes away from the primary exchanges, ITG traders needed to access them, and deploying a smart order router provided that access, according to Mark Wright, global head of product management with ITG.

As liquidity fragmentation began occurring in other equity markets, such as in Europe under the Markets in Financial Instruments Directive (Mifid), the agency broker managed to leverage its years of SOR experience when deploying its smart order router into the European market.

"There was a whole host of plumbing to be done, such as new data feeds, exchange connections, and, certainly, the economics of trading in Europe are very different," says Wright. "That has to be factored in for best execution because of the multiple clearinghouses and the complex settlement process."

ITG also expanded its SOR capabilities within the Canadian markets as recently as the second quarter of 2009 when it introduced SOR for its direct market access (DMA) clients wishing to trade in the Canadian markets. Previously, the firm was using its SOR capabilities mostly to power

its local trading algorithms, according to Wright.

Transplanting its SOR technology to the Canadian markets wasn't that difficult, says Wright, since the Canadian market structure isn't radically different from that found in the U.S. It does, however, share some characteristics with the European markets—for example, most of the liquidity is still found on the primary exchanges, he adds.

"A lot of analysis and best execution comes in when looking at the alternative liquidity, categorizing it, benchmarking it and so on."

Mark Wright, global head of product management, ITG

Although the Australian Securities and Investments Commission (ASIC) doesn't plan to open the Australian markets to ATS providers until the latter part of 2010, introducing SOR capabilities into a new market should be a relatively quick process for ITG once viable alternative quotes are available, according to Wright

The Next Big Thing

Officials at ITG expect that the next step in SOR's evolution will be trading dual-listed securities across national borders. The firm has had its Best Market Server, which examines prices of dual-listed stocks on the U.S. and Canadian markets using a real-time foreign exchange (FX) price calculation and makes the currency conversion if the cross-border quote is better, in production for a number of years.

Other potential markets for the cross-border router may include the European markets where trading is occurring in

euros, sterling, kroners and other local currencies.

"It is something that we have evaluated, but do not have in production," says Wright. "We are certainly investigating it."

Yet there is a limit in terms of how far the platform can go, according to Wright, who says he sees it operating at the edges of the order routing business. "It really depends on how much cross listing happens

across national borders, which is always the upper boundary," he explains. "You can consider U.S. and European American Depositary Receipt (ADR) trades, but you'll have shorter market hours to work with," he adds.

Growing Pains

Over the next two to three years, ITG officials are planning for a rapid increase in the required infrastructure to manage the firm's smart order routers as new venues continue to appear and as each venue tries to differentiate itself through novel trading capabilities.

The firm expects to scale its routers horizontally to meet the growing resource demand rather than deploy a major overhaul of its router architecture in the near future, says Wright.

The current platform is written in C++ and runs on high-performance servers running the Microsoft Windows operating system. Although ITG

has complex event processing (CEP) technology deployed in several places within the firm, the firm does not use CEP technology within its smart order routing platform due to latency issues, says Wright. ITG's smart order router isn't the lowest latency application within ITG, but is certainly one of the lowest, he adds.

The brokerage also avoids what it sees as other custom low-latency architecture, such as networking appliances and custom application-specific integrated circuits (ASICs).

"It's always a mix of moving fast at a feature/function level versus latency, but we have not [chosen to adopt them]," says Wright, who adds that the firm will re-examine these technologies as they mature.

Where changes are being made within the smart order router is where it interacts with non-displayed liquidity.

"When you go to the displayed market, there is not much to discuss: You are accessing a quote," explains Wright. "A lot of analysis and best execution comes in when looking at the alternative liquidity, categorizing it, benchmarking it and so on."

ITG is building in functionality to the router that will process various market signals, which can be translated into estimates where liquidity can be found based on the success that ITG has had in the past with filling orders and examining the fill rate in order to make a judgment on the size of the liquidity present in the venue.

Wright says ITG already has capabilities such as this in place, but is working on refining them substantially.

A Peek at the Next Generation

As more trading venues come online and differentiate themselves in terms of fee structures and order types, smart order routers need more intelligence in order to exploit these market structure changes. *DWT* sits down with Harry Gozlan, founder and CEO of Smart Trade Technologies, and Vincent Burzynski, chief product officer with SunGard Global Trading, to discuss the future of smart order routing (SOR).

How would you describe the current generation of smart order routing (SOR)?

Harry Gozlan, Smart Trade: Smarter, more proficient and flexible SOR implementation is spreading throughout the U.S., Canada, Europe and soon, Asia. But let me pinpoint a concern. Current smart order routers that aren't repackaged complex event processing (CEP) engines have two specific, real-world drawbacks. They are nearly always single-asset based, preventing them from managing several asset classes in the same routing rule. It's obviously a limiting factor for an institutional investor, regardless of size and assets traded. As a result, these engines usually tend to embed hard-coded rules and order types requiring recompilation of the code, if not a new release, to support a new rule or new order type. Unfortunately, clients are limited in controlling implementation of their own rules, adding cost and longer time-to-market.

The second weakness is that they manage a parent-child relationship on the orders they receive and send back, creating limitations in the entire state management of the workflow. With some implications in the failover mechanism, this makes it very difficult to recover the exact state of all orders received and sent to the different venues, limiting the possibility of orchestrating different order books across different ven-



ues, internal crossing engines and order management systems (OMSs). Although CEP engines can be useful tools to handle SOR actions, they do require increased programming because they weren't originally based on market concepts.

Vincent Burzynski, SunGard: First-generation smart routers, particularly in Europe, have focused on maximizing the efficiency of their trading across "lit" public-book venues—the established exchanges and major multilateral trading facilities (MTFs). Identification of trading opportunities is based primarily on real-time analysis of the prices streamed from these venues, including consideration of size and market depth. Routing strategies—designed as they must be under Markets in Financial Instruments Directive (Mifid) to obtain the best possible trade prices from a broker's connected venues—take into account other considerations such as trading platform latency and the fee costs of trading at each

venue. These strategies may include significant amounts of "passive" trading—i.e., placing orders on the markets as a price maker in anticipation of being hit by a taker on the opposite side of the trade. The MTFs' maker-taker pricing models particularly encourage this tactic. Optimizing the results for both broker and client in such cases is complex, and consultation is often necessary.

In concept, the algorithmic approach is usually quite simple, but it has already yielded two important benefits for user firms:

Prices available on the MTFs can in many cases beat those on the "home" exchanges. Aggressive users of smart routing regularly report average price improvements for trades routed to the MTFs in the order of eight to 10 basis points (bps) and often send 30 to 50 percent of their trades this way as the major markets have fragmented more heavily.

The maker-taker pricing models of the MTFs, and the pricing of their associated clearers, are so advantageous that, even with strict adherence

to best-execution policies, brokers can obtain very significant fee cost savings.

Another significant characteristic of current smart order routers is "non-integration." Although the router can be considered "just another algo" contributing to the increasing fragmentation of institutional orders into many smaller transactions, it is normally quite separate from the mainstream algorithm suites of brokers. Simply put, one can say that the algorithmic suite determines when a trade should be done, and the smart router then decides where it will be done.

What do you see as the characteristics of the next generation of SOR? When and where do you see it being available?

Burzynski: The next generation is already appearing. We won't see a disconnect: Additional parameters and strategies are, in most cases, added to existing installations. The main issues being addressed include the following: SOR strategies are now moving beyond pure real-time analysis. As sufficient large-volume liquidity history becomes available for a given security across several venues, statistical analysis is being used to identify what kind of orders to send to a specific venue and when.

Smart routers are increasingly routing to dark pools of liquidity. The statistical analysis

mentioned above is vital here, as in the absence of real-time prices it is the only valid basis on which to decide how much of an order to send to the pool and for how long, usually before progressing the sweep to another pool.

At the same time, algorithms—while still typically separate from the smart router—are taking real account of multi-listed execution. For instance: Where more than 30 percent of liquidity is away from the reference exchange, percentage volume algorithms should incorporate volumes on the MTFs.

Client demand is also driving improved reporting from smart routers, including detailed transaction cost analysis (TCA).

Gozlan: The next-generation SOR will correct the drawbacks I pinpointed earlier with five additional improvements, starting with clarity between where logic is embedded, controlled by the client, and where the execution factory, including failover management, connectivity, state management and “plumbing,” is located.

There’s also a need for greater integration with other liquidity management components—aggregation engines, connectors and crossing/internal matching engine/market data middleware—whether built by the same firm or external IT providers.

Third, SOR will be offered as a hosted service, with the advantage of gaining pre-packaged connection to all the desired venues.

Fourth, it will support very high throughputs.

Finally, it will include multi-leg, multi-asset-class manage-

SUNGARD®**Vincent Burzynski**Chief Product Officer
SunGard Global Trading
www.sungard.com

“Aggressive users of smart routing regularly report average price improvements for trades routed to the MTFs in the order of eight to 10 basis points.”

Vincent Burzynski, chief product officer, SunGard Global Trading

ment, meaning the ability to orchestrate several instruments at once, on one or several different asset classes.

Are we seeing current smart order routers scale to meet the introduction of new trading venues and order types?

Gozlan: Good question. As I said earlier, with the design of current smart order routers, any major change, such as the introduction of a new execution condition, order type or trading protocol, requires considerable effort. But it is possible to avoid this costly, time-consuming if not painful drawback if the design has been properly achieved and fine-tuned over time. However, the rapid, accelerating increase in the number of trading venues with specific order types attached to them means that the same vendor cannot necessarily handle all asset class changes in a reasonable time-frame, and that one solution is to offer, in a very cleanly isolated container, a flexible way for clients to define their own additional order types and, if needed, connections to new

venues. In our case, purely borne through experience, we keep development of new connectors at our level, and give clients more leeway to change rules and order types in a simple manner. It makes a difference. It matters.

Burzynski: Scaling is an important issue for an application where real-time performance is vital, especially when more calculation intensity—e.g., statistical analysis—is being demanded at the same time as new venues are being added. There is more data to watch and more complex decisions to be made. The optimization and performance tuning of router algorithms is, therefore, getting much attention at present, and, of course, hardware power is also being thrown at the problem.

Arguably, this problem is approaching a peak as the number of public-book trading venues for most stocks still continues to increase, whereas expectations are that we will subsequently see some degree of consolidation among the MTFs—though experience

in the U.S. perhaps indicates otherwise. The need to sweep more dark pools will also keep the pressure on.

Is there an average lag time between the introduction of venues and order types before they are incorporated into smart order routers? Is that time period shrinking?

Burzynski: There is by definition a time lag between the introduction of venues and their availability for trading via a smart order router. Clearly, the specifics of the new venue—order types supported and so on—have to be handled, and testing has to be carried out. Routing algorithms may need to be adapted, particularly if the venue offers new instrument coverage, onward routing, or peculiarities in the commercial offer.

The fact that MTFs tend to follow fairly standardized approaches in many areas makes new introductions progressively easier to handle, so time lags are shrinking. Lags can also be reduced where specifications are available well ahead of start dates, and also where introduction of instruments is phased, as is usually the case.

Current initiatives should contribute further to the trend of market simplification: For example, the common-symbology initiative being driven by several MTFs removes the need for smart routers to manage some “trans-coding” features.

Gozlan: Adding a new venue becomes faster not because of the smart order routers but because of the adoption of

continued on page s6

A Peek at the Next Generation

continued from page s5

message standards such as FIX to communicate to new venues. We tend to see that the type of orders mostly used against most of the venues are always the same four or five, so complexity of message standards tends to shrink over time, offsetting the additional workload due to the increase in number of venues to connect to. New order types, if they are proprietary to the client, are closer to an execution algorithm, and this, as explained before, is easier if the router has split the function of logic management versus execution management.

Brokers' smart order routers are well integrated with their respective dark pools of liquidity. How can third-party smart order routers compete to access dark liquidity?

Gozlan: If third-party smart order routers are pure software engines, such as what we offer, the connection to dark

“There’s a need for greater integration with other liquidity management components—aggregation engines, connectors and crossing/internal matching engine/market data middleware.”

Harry Gozlan, founder and CEO, Smart Trade Technologies

pools or venues is always done on behalf of our client. Here, when the demand to connect to a specific liquidity pool is validated based on the credit that our client presents, we act only as a technology arm. If the third-party router is offered as a packaged service, then it is true to say that they may compete with brokers' smart order routers, who may not want to see the access to their clients be bypassed through an intermediary company. Today, the question of controlling flows is very strategic and we see how this impacts the control of the major exchanges and venues, owned sometimes by dealers. The control of smart order routers is of the same

magnitude, and this is why we always elected to be a pure software provider, running no execution service based on our technology.

How can smart order routing co-exist with low-latency execution? How much latency do smart order routers add to trade execution?

Burzynski: The low-latency objective is of course the reason why the scaling and performance issues discussed above are so important: Smart order routers have to be as fast as they are smart; otherwise their decisions, based on real-time prices, will be useless. Most routers, therefore, aim for sub-millisecond latency, and this

can be achieved where algorithms are not too real-time intensive. Clearly there is a trade-off of speed against complexity, with tuned algorithms and faster hardware deployed to maximize latency as far as possible.

Gozlan: My sense is that if a smart order router adds less than 1 millisecond of latency, it is perfectly acceptable in most of the market cases, whatever the asset class. However, the real trap is in the coding of the rule or execution strategy, because complex strategies may add some loops in their execution, or iterations, which, when reproduced at the scale of entire trading workflow, and especially in fast markets, may be costly in terms of CPUs and thus latencies eventually. For ultra-low latency, not only the smart order router has to be taken into account, but also collocation, transport, network, and so on. They have to work hand in hand.

Smart Order Routing in the News

BOFA-ML TAKES ORC EXNET

In a bid to expand access to its U.S. options algorithmic trading platform globally, Bank of America-Merrill Lynch has linked to trading platform and connectivity provider Orc Software's Orc ExNet connectivity network. The move will allow Orc users to trade U.S. options using the bank's Clean Sweep, Hidden Sweep, Iceberg and NBBO Sweep algorithms, according to vendor officials.

CM-CIC SECURITIES TAPS QUANTHOUSE

Officials from trading technology and data vendor QuantHouse have announced that the investment banking arm of CM4-

CIC, CM-CIC Securities, has selected the vendor's ultra-low-latency market data-feed to feed the firm's smart order router and algorithmic trading engine. According to a firm spokesperson, CM-CIC started testing the data feed in the second quarter of 2009 before going into production the subsequent quarter.

SOR MARKET TO EXCEED \$1 BILLION

According to a recent report published by industry analyst firm Aite Group, the demand for smart order routing will exceed \$1 billion in 2010 as liquidity fragmentation drives further adoption of multi-asset and multi-geographic trading strategies among investors. The report author also expects

that 95 percent of firms will leverage smart order routers from multiple brokers.

PROGRESS APAMA, CQG UPGRADE AUTOEX

Complex event processing (CEP) engine vendor Progress Apama and hosted multi-asset trading and charting platform provider CQG have added price capture, error detection and risk management capabilities to the latest release of Apama's AutoEx for CQG Integrated Client (CQG IC). Unlike the prior release, the platform detects trading opportunities in real time and can immediately execute resting orders that achieve tighter price control and best execution, according to officials.

HOSTED BY
FISD, INSIDE MARKET DATA & INSIDE REFERENCE DATA

ASIA PACIFIC FINANCIAL INFORMATION CONFERENCE

APFIC2009.net

November 3-4, 2009
JW Marriott Hotel • Hong Kong



Bringing together market data, reference data and data management executives from leading financial institutions across the Pac-Rim to examine solutions to urgent challenges facing their business. Program topics will address key issues affecting the Pacific Rim financial information industry and serve as a platform for sharing ideas and solutions on market data and best practices.

Hear from the leading practitioners, including:

- **Robin Leung**, Managing Director, Head of Global Information Technology Division, Bank of China
- **Dennis Mwansa**, Head of Trading Infrastructure Services, Asia Pacific, Royal Bank of Scotland
- **Christopher Butler**, Project Director, Data Management, Basel II Data & Reporting, National Australia Bank
- **Makoto Shibata**, Principal Analyst, e-Business and IT Initiatives Division, The Bank of Tokyo Mitsubishi UFJ
- **Rose Marie E. Sotelo**, Head of Treasury Risk, Land Bank of the Philippines
- **Madhu Gayer**, Head of Investment Risk and Analytical Services for Asset Servicing Clients, Northern Trust
- **Barnaby Nelson**, Head of Business Development, BNP Paribas
- **Michael Mollemans**, SVP, Head of Electronic Execution Sales, Daiwa Securities

Conference highlights include:

- Keynote addresses by **Terry Roche**, Global Head of Information Management Solutions, Thomson Reuters and **Roger Sargeant**, Managing Director, International, Interactive Data
- End-user Panel: Getting the best value for money in a tough economy
- The data demands of advanced trading strategies
- EDM: Integrating reference data sources to manage risk
- Obtaining reliable data for OTC derivatives, complex products and new asset classes
- Vendor Panel: Data providers respond to end-users' questions
- Next-generation technologies: The shape of things to come

Plus superior networking opportunities:

- **Welcome reception** on Monday, November 2
- Breakfast, lunch and coffee breaks
- Exhibition hall with the **leading data and technology providers** in the industry
- **Reception & dinner** at The Verandah on Repulse Bay on Tuesday, November 3

FREE attendance for qualified delegates from financial institutions

CO-HOST SPONSOR



HKE
香港交易所



SSE INFONET LTD
A Shanghai Stock Exchange Company

VISIBILITY SPONSORS



EQUINIX



GOLDENSOURCE



STANDARD
& POOR'S

EXHIBITOR

Capital IQ

A Division of Standard & Poor's

Sponsor / Exhibit:

There are a limited number of opportunities remaining. To leverage this conference to advance your marketplace objectives, we suggest you contact us soon.

Jo Garvey on +1 212.457.7745 or at jo.garvey@incisivemedia.com
Lee Hartt on +44 (0)20 7484 9907 or at lee.hartt@incisivemedia.com

Event updates / enquiries:

To register or for further information, please visit APFIC2009.net or contact Pablo Cot Gonzalez on +852 3411 4888 or at pabloctg@incisivemedia.com

APFIC2009.net