March to the algorithm

Gaining edge in automated trading: ISVs and exchanges level the playing field. **Tom Haldes** and **Diane Saucier** from **Trading Technologies** report

Recent years have seen a surge in the growth of algorithmic trading, with more trades executed through automated systems each year. The availability of global order routing platforms is stimulating an increase in automated cross-market trading. Although average fill sizes are getting smaller, the number of filled orders is increasing dramatically. The highly active global markets continue to attract more participants, including hedge funds, which have embraced and are continuing to use automated trading at a growing rate. Algorithmic trading allows hedge funds to model, test and deploy highly complex quantitative strategies that would be impossible to execute manually.

A large portion of these funds have drawn on the capabilities of third party trading platforms as they move to more high frequency trading models in order to capitalise on new opportunities and new markets. In response to these growing demands, many third party platforms have become more advanced, supporting multi-market access, customised application development and complex strategies, empowering the buyside and further driving liquidity and volume.

Many hedge funds are moving from the traditional buyside position trading model into more high frequency strategies centered on statistical arbitrage and automated market making. Consequently, they are devising increasingly complex trading algorithms that seek to identify temporarily mispriced assets, taking short term positions to capitalise on the price anomaly and reversing the trade when they feel the price has reverted back to its intrinsic value. In some ways this resembles traditional scalping, except that the entry and exit points are determined at the sub-millisecond level, and many markets may be monitored and traded simultaneously. In addition, as new exchanges emerge offering additional sources of liquidity and volatility, market participants, including hedge funds, develop spreading and trading strategies that involve more markets around the world and contribute to increased order flow in the global electronic markets.

As an increasing number of trading firms move toward high frequency, automated trading, more order flow is generated and the increased flow by definition increases liquidity. This in turn stimulates more orders from other firms, and so on. Additionally, as their automated trading strategies have begun to incorporate emerging and often volatile markets, such as the rapidly growing commodities markets, price movements on one exchange may generate transactions across other exchanges around the world, again adding to volume.

A number of consequences flow from these trends. The first is that market data rates are skyrocketing as a result of automated quotation and execution. Ten years ago market data vendors could aggregate real-time market data for all major US exchanges comfortably on a T1 circuit. Since then, market data rates have grown by roughly two orders of magnitude, requiring concommitant upgrades of data networks and computing systems. Yet the truly astounding fact is that this growth is continuing, and today's data infrastructures will require continual expansion to keep up with ever increasing data volume.

But as data volumes soar, the ratio of quotes to trades increases as well. In some equities markets this ratio exceeds 1000:1. While the ratios are not as high on the derivatives side, all exchanges have seen a dramatic

increase in their quote to trade ratio. Eurex, for example, received eight times as many quotes per trade in 2006 as in 2001. And the ICE Futures exchange has experienced a 275% increase in traded volume over the past two years while the transaction volume – representing all message types - has grown 2700% over the same period.

To understand what is behind the market data "tsunami," as it is often called, consider a high frequency trading system that is auto-quoting options. A single price movement in the underlying security may cause the system's logic to re-quote every strike for every month in the options series. Now assume the underlying security is also being quoted by high frequency black boxes, such as index or statistical arbitrage systems, with each new quote from the black box applications potentially resulting in re-quotes of hundreds of strikes in the options series, and you can begin to understand why data is becoming such an issue for exchanges, vendors and trading firms alike.

Technologies devised for high frequency trading are also enabling markets to become more interdependent. As bigger data pipes and faster computers are increasingly deployed by buyside firms to monitor real-time prices across multiple markets, the window of time required to capture inter-market arbitrage opportunities has diminished. High performance trading systems are no longer strictly in the domain of the world's largest banks and funds and instead this technology is becoming a basic necessity for buy side funds.

To address these ever-changing conditions, firms must devise increasingly sophisticated algorithms. At the most basic level, simple synthetic order types can arguably be called "automated". Many commercially available trading platforms support such synthetic orders, where behaviour is determined by market conditions. But most practitioners think of algorithmic trading in terms of grey and black box concepts. Grey box algorithms combine software-based algorithms with human interaction and allow the user to set, monitor and alter the parameters. Traders may run several concurrent grey box applications, monitoring them visually and/or audibly, and stepping in to take manual control when necessary or advantageous.

Black box systems take away much of the human interaction, and typically are designed to run a proprietary strategy developed internally by the quantitative research group of a bank or hedge fund. Because black box trading involves highly proprietary trading models, many quantitative funds employ skilled programmers to develop these systems in-house. This is one of the most secretive activities of hedge funds.

Whereas grey boxes are often deployed as desktop applications, black box systems are typically implemented to run on high performance computational servers. Once developed, black box algorithms are usually back tested against months or years of historical data. If the historical data suggests the model will be profitable, the algorithm is placed into limited production to be scaled up if empirical results are positive. Because black boxes require limited user intervention, they will sometimes include "circuit breaker" logic to stop or unwind trading automatically if things go wrong.

A fund wishing to deploy an algorithmic trading model faces a critical technology decision before the first line of code is written: to build or to buy?

A fund wishing to deploy an algorithmic trading model faces a critical technology decision before the first line of code is written: to build or to buy. A few very large hedge funds insist on writing their own exchange interfaces, or "gateways," to handle market data and order submission, the argument being the firm then does not have to generalise any part of the gateway for use outside of its own intended purpose. This in turn may be expected to yield small but worthwhile performance advantages over commercial products. The downside to this is that the firm must hire and retain a staff of expert programmers to build, update and maintain complex exchange interfaces.

Exchange protocols vary widely, both for market data and order submission, and exchanges are constantly changing and improving their interfaces in order to support performance and functionality enhancements. The common solution for isolating automated trading applications from these complexities is to create a "normalisation layer" that provides black box developers with a common set of programming commands to receive data and route orders across all exchanges. This approach allows modifications to the underlying exchange protocols to happen behind the scenes without requiring changes to trading applications, such as algorithmic black boxes.

But it is the creation of this additional layer, according to some, that eliminates the performance benefit of a proprietary gateway over a vendor's product. To preserve this perceived edge, their implementation may eschew exchange normalisation and ultimately present exchange-specific nuances to their users, namely the algorithm developers. These black box developers now must add additional complexity to their systems to handle the non-normalised exchange interfaces – a time consuming process. In addition, hedge funds and other automated trading firms must develop, test and maintain an increasing number of exchange gateways.

In reality the gateways offered by independent software vendors (ISVs) are capable of receiving market data and routing orders in fractions of a millisecond. Gateway design and order routing is core to the business model of the ISV, and through years of continuous development and tuning, these systems perform faster than many in-house developed gateways. Yet these platforms also provide the substantial benefits of exchange normalisation and independence from exchange interface changes. This means that a black box can trade new products and markets, and take advantage of exchange protocol enhancements, with no changes to the black box application.

Most commercial order routing platforms provide a proprietary application programming interface (API) and/or the Financial Information eXchange (FIX) Protocol for routing orders and receiving market data. Hedge funds in particular are increasingly relying on FIX to integrate their black box trading systems with com-

mercial trading platforms, gaining virtually immediate access to electronic markets worldwide.

Because the data is normalised across all exchanges to FIX, clients only have to integrate one time – from their FIX engine to the ISV's FIX adapter – to be capable of trading thousands of products worldwide. It is useful to think of such a platform, in conjunction with the ISV's FIX interface, as a single, super-exchange through which customer developed algorithms can trade an ever increasing array of global products at sub-millisecond speeds. This provides enormous "business agility" to hedge funds looking to start trading a new product or exchange immediately, since the high speed infrastructure is already in place, and their algorithms require no modification to trade the new product or market. These algorithms gain immediate additional benefits as well, including pre-trade risk checking and real-time position and P&L monitoring, with no additional programming.

In addition to offering open API and FIX connectivity to their platforms, some ISVs employ a dedicated team of programmers to support external developers writing to their platforms.

Whether a fund is using an ISV's grey box products for algorithmic trading, developing black box algorithms in-house or relying on some combination of the two, several industry trends are clear and should be taken into consideration by hedge funds expanding their automation capabilities.

The first is that automated trading is increasing as a percentage of the total traded volume across all electronic markets. This trend will continue as both the buy side and sell side increase their reliance on semi- and fully-automated systems. Another is that funds will increase their desire to rapidly identify and trade more multi-exchange opportunities. This trend is being facilitated by vendors that provide a single interface through which firms can integrate their manual and automated systems one time in order to trade dozens of exchanges worldwide.

The acceptance of FIX will continue, especially as a single interface to multi-exchange order routing platforms, and work is underway in the global FIX Protocol Limited (FPL) organisation to enhance FIX's support of algorithmic trading. All of this will create additional order flow worldwide, increasing the capacity requirements on data circuits, computers, software applications and local network infrastructures.

As these trends continue and data volumes soar, exchanges are making large scale changes to their matching engines and technology infrastructures. For example, many exchanges are increasing the transactional capacity of their matching hosts and expanding the bandwidth of their data circuits in order to accommodate increased quote volume and new product launches.

"Funds will increase their desire to rapidly identify and trade more multi-exchange opportunities"
Diane Saucier

On Liffe, while the transaction volume has increased by over 1000% in the last four years, enhancements to the matching host during this same period have enabled a reduction in the Liffe host calculation time from 250 milliseconds to 1 millisecond.

At ICE Futures, during the most recent two year period that saw inbound message traffic multiply 27 times, average data centre time was reduced from 600 milliseconds to 28 milliseconds, while the actual time in the matching engine has dropped to around one millisecond. As the exchange rolls out a new trading engine this year, data centre times are projected to continue to drop to an estimated seven milliseconds.

According to the Chicago Mercantile Exchange (CME), while transaction volume has grown by a factor of 30 from its 2004 level, average round trip time has fallen 80% during that same period.

While, on Eurex, where average quotes per day rose from 10 million in the beginning of 2005 to 145 million in May of 2006 and more than 233.5 million in July of 2007, plans have already been announced to increase quote transaction capability to 1.6 billion messages per day by 2009. During the same period in 2005-2006, host enhancements on Eurex enabled the exchange to reduce response time from 80 to 100 milliseconds down to 40 milliseconds.

In a nutshell, exchanges are taking advantage of advances in hardware technology, horizontal scaling and clustering, higher bandwidth connections and improved market data dissemination methods in order to continually improve capacity while reducing latency.

Exchanges are also beginning to offer or enhance their collocation facilities through which clearing organisations and their customers can physically install time sensitive components, such as the servers that run high frequency trading algorithms. This way, a hedge fund in Texas can reduce order submission latency from their black box to an exchange's matching engine from over 20 milliseconds to less than one millisecond. In today's high-frequency environment, that difference is monumental and often determines whether or not the order gets filled.

In addition to providing substantial improvements in order latency, buy side firms are beginning to take advantage of hosting and collocation services that allow firms to outsource most maintenance of their electronic trading infrastructure. This allows the firm to focus on its core trading business while systems administration experts at the hosting and collocation centers manage the firm's technology capital.

Electronic trading has opened the door to automated trading. This in turn has brought about a massive increase in data and infrastructure requirements for trading firms, software vendors and exchanges alike. Trading automation has changed the landscape

Tom Haldes, Senior Product Manager, Automated Trading & Diane Saucier, Vice President, Market Development, both with Trading Technologies International