

Me and My Machine: The Automated Trader Interview

Miles Kumaresan, Head of Quantitative Trading at TransMarket Group Ltd, talks to Automated Trader's Editor, William Essex, about the development of the company's High Frequency Quantitative Trading Group.

Do you make a distinction between automated and algorithmic trading?

Not really – we tend to use algorithmic trading as an umbrella term for all aspects of quantitative trading. So while I look after quantitative trading, where I'm building up a new group, I also handle the algorithmic execution of the group's trades.

Do you use your own in house execution algorithms?

Absolutely; I cannot envisage a situation where we would use a third party's execution algorithms.

TransMarket has a strong reputation as a proprietary futures trading operation. Does the creation of your trading group presage an expansion beyond that?

I can't speak for the rest of the business, but certainly the quantitative group isn't constrained by asset class. We trade across various market complexes including cash equities, futures, FX etc.

How does your group fit within the other TransMarket business units?

I was hired in April of this year to build up a high frequency quantitative trading operation from scratch. The company views this new division as a key revenue generator going forward and one that we intend to build up across the regions. The quantitative trading business already has an operation in Chicago as well as London. However, we soon expect to be operating on a 24 hour basis.

Our quantitative business is entirely separate, so we have little day to day involvement with the other TransMarket business units and regional offices. Our focus is 100% on price action and volume based quantitative models, so factors such as news and local market knowledge are not relevant for us.

While we anticipate that the quantitative business will continue to expand globally and tap into existing regional resources, there is always the question of finding somebody suitable to manage those overseas



Miles Kumaresan (right), William Essex (left)

“Ultra high frequency liquidity taking strategies are nice, but they have limited capacity.”

operations. At TransMarket there is a strong awareness of personality and cultural matters and a realisation that just because somebody is a good trader this doesn't automatically qualify them to run a trading operation.

Despite the business autonomy of the quantitative unit, is it included in any internal TransMarket netting programs?

At present it isn't, but we will eventually have internal netting, because we will have consolidated routing of orders, although that is probably another twelve months away. We are still in the early stages of building up the business and the potential benefits of internalisation are small in comparison with the P & L potentially derived from trading strategies. Once the high frequency quantitative trading group reaches a degree of maturity in terms of its core business, it will then justify putting resources into internalisation.

In general terms, what types of quantitative strategies does your group undertake?

I am risk averse, so most of our trading is intraday with positions held overnight in high Sharpe Ratio strategies only. We certainly don't undertake long-term directional trading.

I have a natural bias towards provision of liquidity and I believe there is a lot more opportunity in that direction. Ultra high frequency liquidity taking strategies are nice, but they have limited capacity. If you really want to scale up your business, you must have strategies that allow large capacity, and liquidity providing strategies definitely comply with that requirement. They are also less sensitive to overall market volatility, so even if volatility dries up we can still remain active and generate P & L. Having said that, we will be busy in the liquidity taking spectrum as well.

Are there any particular markets you are already active on as regards liquidity provision?

We are currently trading a number of liquidity making and taking strategies. We are about to go live on Hotspot and intend to add other major FX venues



“...100 millisecond latency for an institution-level algorithmic execution environment is more than sufficient”

as well. However, the whole process of conformance testing, completing the paperwork etc. takes quite a while.

Are other TransMarket business units already active in liquidity provision?

Not really. Some of them are certainly providing liquidity, but they are also takers; so on balance, they are probably not substantial net providers.

Inevitably latency has been a regular topic of conversation in auto/algo trading circles in recent years. What's your perspective on that in the context of trading strategies?

I have to say we have been rather spoilt in this respect. Two years ago TransMarket hired Anthony Stanfield, the head of technology of Citadel Europe as CIO. He has spent twelve months creating a high-performance infrastructure and technology platform, in the process reducing our internal network latency to 45 microseconds. A true partnership of trading and technology...

On the algorithmic execution side, I have always held the view that latency is not of enormous importance. It is nice to have, but 100 millisecond latency for an

institution-level algorithmic execution environment is more than sufficient. Given that limit orders typically take minutes to mature, micro and milliseconds don't mean a great deal.

However, with regard to proprietary and quantitative trading, any latency advantage is very nice to have – but it is not the be all and end all. You cannot realistically put all your eggs in one basket with a strategy that attempts to capture small scale mispricing in the sub millisecond or single digit millisecond time frame. The problem is that this is a 'winner takes all' type of strategy, so if you come to work one morning and find that a competitor is now a millisecond faster, your entire P & L has vanished.

Therefore the objective with the high frequency quantitative trading group at TransMarket is to build a far more diversified and stable business that is not solely dependent on that sort of extremely latency-sensitive strategy. It will certainly be an important part of a diversified strategy mix, but the cost of maintaining the necessary speed edge is extremely high and its sensitivity to being overtaken means that you cannot project long-term investment in an operation with that as the sole revenue generator.

On that point, do you feel that many market participants really exploit all the possibilities of automated trading to the full, in terms of diversification by model, parameter and timeframe?

My sense is probably not. Over the last few years it has been very hard making money in quantitative trading, particularly when volatility dropped off to low levels. Yet despite that, I have the impression that relatively few people have responded to these circumstances by exploiting diversification to the full in order to uncover opportunities.

There has been a great deal of interest in high frequency trading in recent years, with many market participants claiming to be active there. Do you think that this is really the case?

There are capacity restrictions at the very shortest timeframes so I cannot imagine that very many people really are active there. The sophistication and technology investment required to achieve and maintain an edge at the high frequency microstructure level are extremely high, so I suspect that far fewer organisations are active in this spectrum than claim to be.

A relatively common complaint by Automated Trader readers is that they don't often see any real benefits from partnerships between vendors. What's your view on this?

To date I don't believe I have seen a partnership on the system/hardware side that makes a particularly radical difference to performance. We see various scenarios where we are told that running a particular app on a particular hardware platform will enhance performance, but such gains that may arise are typically unspectacular; a 20% performance increase isn't really going to transform things from our perspective.

Having said that, I'm hopeful that the relationship between Intel and Kx Systems (who are a key vendor for us) will deliver a significant performance increase. Based upon my prior experience of Kx and their mindset, I am optimistic that something substantive will result.

What modelling tools do you use and how do they integrate into your development and trading platforms?

Our in-depth data analysis and prototyping is done using MATLAB, which I regard as an excellent tool that I've used for many years. From our perspective, one of its greatest strengths lies in its support of the statistical functions we require to easily analyse data and the optimisation capability is also good.

Once model prototyping is complete, we use an in-house trading and simulation platform that I have developed over the last decade or so and regularly revise approximately every two years. It is a form of complex event processing system and most importantly provides an environment where there is no barrier between simulation and real trading. New trading models follow a plug-and-play philosophy without requiring a single line of code change.

During simulation the platform takes into consideration all the various environmental factors that affect performance, such as latency, the order book queue, probability of fills etc. Therefore when simulation is complete and we are ready to go into production, we do not need to change any of the code. If any statistical/mathematical routine is too complex to write natively for the platform, we will simply make a call to MATLAB.


How do you handle the application latency implications of making that call?

Production usage of MATLAB DLLs here at TransMarket is very limited. In my previous employment however, as head of algorithmic trading, most of the execution algos we developed were written in MATLAB and deployed into production. There

is a twenty millisecond overhead involved in calling MATLAB functions, so we won't make the call for each tick. The platform's design was structured to accommodate this sort of situation. Its order manager component is written in Java, and a lot of the data manipulation and some of the analytics are written in Kx, but the actual execution algorithms are written in MATLAB.

Data feeds come directly into Kx and to circumvent the twenty millisecond overhead involved in calling MATLAB we use buffering. Instead of always making a calculation call for every single stock on every tick we will buffer the data for perhaps fifty milliseconds ▶





“...we have been surprised at how often other participants don't disguise their strategies and how easily discernible they are by analysing the tick data.”

and then make a single call to MATLAB for all the securities in one hit. Over time, that process of mass calculation has also proved to be extremely robust in a production environment; in some two years of operation neither MATLAB nor Kx have ever crashed.

Do you regard the ability to move directly from simulation to live trading within one environment as essential?

Yes, for several reasons. On the one hand, you enjoy a significant time to market advantage in terms of getting models into production. By contrast, handing prototype code to an IT team for production recoding is a far slower process and of course also runs the risk of inadvertent translation errors being introduced. Finally, there is an obvious security benefit in terms of intellectual property (IP); if a trading or execution

model is never disseminated outside the original development group then the chances of IP leakage are minimised. Being able to create everything in a 'safe module' that can be plugged straight into production is therefore beneficial.

Automated Trader's readers seem to have an almost religious fervour in their attitude to the relative merits of C/C++ versus Java for automated trading. What is your view on this language debate?

I have spent a lot of time working on low-level programming and C++. However, for the type of work we do, after conducting a benchmarking comparison of C++ and Java a few years ago, I have mostly opted to use Java. There isn't a performance penalty in using Java for mathematical routines. On the contrary, in some cases Java has better performance



performance is down to individual usage. You see people frantically debating the relative virtues of Java versus C/C++, and yet the same people still happily incur the performance penalty of using hash tables. That penny wise, pound foolish approach rather defeats the object.

Furthermore, the number of programmers with specialist expertise in coding for high frequency trading is small. Therefore a language such as Java that helps to minimise bugs and memory leaks is invaluable in an environment where such flaws simply cannot be tolerated. I am not alone in this view of Java, as two of the top very high frequency hedge funds in London also code in Java.

Do you feel that the growing focus on quantitative strategies and the availability of sophisticated modelling tools is reducing the life span of profitable trading strategies?

That very much depends upon the type of strategy and timeframe. As I mentioned earlier, there is limited capacity in the very high frequency spectrum for price taking strategies and profitability may genuinely hinge upon a sub-millisecond advantage.



than pre-compiled code because of its runtime optimisation technology. Java's on-the-fly code optimisation capabilities (such as its analysis and refinement of 'hot spots') effectively compensate for any more generic performance discrepancies.

Since my primary aim is to make money, the choice of programming language or technology in general for that matter, is driven by what is going to enhance my P & L – it is as simple as that.

As a result, from our perspective there is little to choose between the two in terms of performance. However, Java has very substantial advantages over C++ in terms of code reliability, maintenance and manageability. Furthermore, so much of the relative

However, elsewhere I am not so sure that it applies. A lot depends upon how well you disguise your strategy. If well-implemented, I think a good liquidity provision strategy can be difficult to detect and therefore resilient in the longer term. Having said that, we have been surprised at how often other participants don't disguise their strategies and how easily discernible they are by analysing the tick data. In general, liquidity taking high frequency strategies are easier to detect. ▶

Nevertheless, I think a far more significant issue is the mobility of personnel. As people move on to other firms, there is inevitably a dissipation of trading ideas. That is far more likely to affect model profitability than the efficiency of tools in spotting similar strategies in the market or inadvertently arriving at similar trading ideas. History provides plenty of examples of this, such as the widespread adoption of pairs trading after the break up of Nunzio Tartaglia's quant team at Morgan Stanley in the 1980s.

Do you feel that obfuscation of a trading strategy to prevent its identification is something that should be part of its core business logic or something that should be part of its algorithmic execution?

Certainly the former, and in the case of large orders the latter as well. For example, particularly with a less liquid stock, you need something that trades in a way that mimics the localised noise in the stock's price and volume. Ideally the two time series with and without your trades included should be effectively indistinguishable.

This is particularly important with liquidity provision strategies. If you wish to represent 20% of the bid you can't just place that percentage of orders there en masse. (Well you can, but the consequences will be unfortunate.) You have to break up your flow so that it

mimics the liquidity arrival pattern for that price/depth level and time of day, and appears as if it were natural liquidity coming in. Furthermore, when the price moves up/down, you don't immediately cancel and replace all your orders at the next level simultaneously.

"...ultimately you need a mentality that nags you to find, model and exploit trading opportunities."

Astonishingly, you do occasionally observe some very large participants doing this. From our perspective it is highly convenient when they do, as it makes it possible to estimate and capture short term momentum. Also, as there is a lot of spoofing in the market, we use a bank of anti-gaming filters that validate new price levels as being genuine (or not) before we respond.

How do you handle the recruitment process for your group?

With considerable effort! The group has no limit on headcount, yet it is extremely difficult recruiting suitable personnel. I spend a great deal of time looking at CVs and find that the majority of them end up being discarded within a matter of minutes. A common issue is people claiming to be what they patently are not. There are not many people who are good intraday traders - never mind good high-frequency ones.

After initially screening CVs, I will spend perhaps ten minutes talking to possible candidates on the phone to try and gain an intuition as to whether they might be suitable. A few of these may be asked for formal interview, but so far our actual rate of recruitment has been low – barely a few percent of those who initially approach us with CVs.

Through our relationship with recruitment agencies we have managed to reduce the number of clearly unsuitable CVs that we receive, though obviously that runs the risk that somebody eminently suitable may be inadvertently filtered out. However, as a counter to that, you tend to hear quickly about potential recruits from successful trading groups on the grapevine.

What skills are you looking for?

In essence, we look for entrepreneurial candidates with natural drive and enthusiasm to succeed. From our perspective, being a quant is not necessarily a prerequisite for a quant trader. Where somebody has studied and whether they have a first class honours degree or higher qualification isn't really important. You certainly need a good mathematical head, but some PhDs have considerable difficulty in getting their head around the concept of a trading edge. In particular we are looking for people with a balance of solid programming skills, a good head for mathematics and a strong natural instinct for alpha capture. This last point is critical, because ultimately you need a mentality that nags you to find, model and exploit trading opportunities.

