

ATP001 JP Morgan

Announcer: Voices in Business presents The Algorithmic Trading Podcast. Today's guest, Carl Carrie, JP Morgan.

[music]

Announcer: Hello, and welcome to the first interview in our series on algorithmic trading, sponsored by Sybase. During this series, our interviewers will be asking deeper questions, from a variety of thought leaders on alternative execution that are not covered by traditional media outlets. This is your podcast, and we strongly encourage you to submit your own questions, either via email or in the comments section of the website, www.algotradingpodcast.com

In today's interview, Sinan Baskan, senior product manager at Sybase, and Greg Grimer, editor of Voices in Business, speak to Carl Carrie, the head of product development in the Electronic Client Solutions Group at JP Morgan, about some of the technical challenges around growing volumes of data, the integrated architecture, and the cross-fertilization of skills and experience, which are needed.

Greg Grimer: Welcome to The Algorithmic Trading Podcast, gentlemen. Can I start, Carl, by asking the first question, which is... Recently, I think, I saw in the "Financial Times," a statement that 50 percent of the deals in the London market would be generated by computerized trading systems. Arguably, we're now at the point where more trades are being produced, in certain markets, by computers than by people, you know, generated initially.

Could you give us a feeling of where things stand today, for our listeners who are not, perhaps, familiar, on a day-to-day basis, with algorithmic trading?

Carl Carrie: Clearly, algorithmic trading is a huge factor. High-frequency trading for Arbitrage's, indexes, ADRs, pairs, ETFs, interlisted trading, as well as automation around auto-working, have all been factors contributing to the growth of algorithmic trading and trading on exchanges.

The exchanges themselves have also been contributing factors. They've invested heavily in capacity and throughput. And the allocations of assets to European equities has also been a minor factor.

Combined, that, and the reduction in transaction costs, which is probably a feedback loop that has resulted in even more flow going to the exchanges, because you can transact them at lower cost because you're using algorithms, and using more algorithms because you can transact them at a lower costs. That loop has continued, in the same way that the trend has continued unabated in the States.

Sinan Baskan: In light of what you just stated, do you see the actual bottlenecks, in terms of just data capacity, on the exchanges' side? Or do you think there is also a potential problem in distributing data inside the institutions and capturing the data for algorithmic trading?

Carl: Yeah. The gaps exist in many areas. There is certainly data capture, and we're constantly riding up against laws of existing infrastructure, physics, and just overall science, in handling the data capture. The amount of tick data that we store is increasing. The amount of derived data that we want to store is also increasing. So, implied volatility, implied correlations, implied choppiness or blockiness of particular stocks, implied co-variances between stocks, are all things that we're starting to store, but not storing to the extent that we would ultimately like.

Terabytes are turning into petabytes. Data storage costs are quickly becoming dis-economic, as you look to expand out. So, there's always a balancing act with the data that you want to store and the data they you have to store for regulatory and just marginal competitive efficiency reasons. But, beyond data capture, the amount of technology that's required to process the data and analyze the data. And there, we're being challenged by the availability of the kind of software that can be applied.

We've explored complex event processing. We've explored time series databases. We've implemented both of those technologies in our infrastructure and have achieved some level of improved efficiency, but the science and the implementation kind of reality is there's still a gap between what's desired and what's currently feasible.

Beyond the capturing, the other key technological area that we're spending a lot of time and a lot of energy is really on the productization side for our institutional clients, and delivering that data to our clients, and the presentation format.

So, visualization, Web 2.0, Ajax, making those technologies seamless and efficient, security technologies - newer technologies that are less invasive, so you don't have to go through layers and layers of passwords and layers of encryption technology, which slows the data down and so on - are all contributing to a better experience, ultimately, that's what we're trying to achieve, but there remain challenges nonetheless.

Greg: Is that technological architecture so complex that if one area was improved tomorrow, it wouldn't make a dramatic difference? Or is there an area you can think of where, if there was a huge leap, there would be a big difference in terms of what was possible?

Carl: That's a great question. One of the things that... I have many friends in the consumer technology space, and we talk about new devices, like the iPhone, and the seismic-shift, the tectonic plates shift, and you see that evolving and you get really excited about the potential for new technologies to really be a disruptive force.

The same sort of effect is far more rare within institutional trading, and because it's more

around the innovations that are really successful and more around the integration of newer technologies than any particular technology. But, notwithstanding that comment, there are areas, like complex event processing, that we're looking for advances. The better integration and better vertical kind of support is really what we're after.

Sinan: If you were to pinpoint a possible breakthrough technology for the next 6 to 12 months to help you integrate better vertically and horizontally, what would you pick?

Carl: Certainly, there are still opportunities for complex event processing, particularly as it integrates better into our traditional data repositories, into our additional streams of information. I think, the next wave is really around weaving complex event processing into analytics and into algorithms, and the tighter integration of the three of those.

If you look at algorithmic trading, in effect, that is complex event processing in a nutshell. If you look at what drives complex event processing within algorithms, it's really the analytics. So, I look at CEP as kind of the bridge between the two.

But, the ultimate advance is less about the infrastructure that those three things represent and more about the user experience. So, institutional traders using algorithmic trading to implement their alpha and beta ideas directly, rather than having to go through intermediate steps and costly delays in implementation. That's really where I think the street is going, and where the more sophisticated hedge funds are making more investments, is tying in the alpha to the algorithm and fixing the connectivity, the synaptic layers, in-between.

Sinan: Is this the kind of technology that would diminish the difference between large institutions and the smaller players in the algorithmic trading space? You would get smaller guys become a more robust competitor because of this?

Carl: It's always interesting watching scale merchants, like the investment banks, compete against the smaller, nimble, very agile, technologically savvy, smaller firms. There are always opportunities for small players to excel and become larger players in the space, and there are opportunities for scale player to continue to compete on scale. You traditionally can't have more than a handful of scale players who are real winners.

And in the electronic trading space, you're seeing the beginnings of a fallout, and you're seeing larger scale players, some of them become clear winners. Not that they can permanently sustain their competitive advantage, but for a period of time, there is an economic advantage in being the preeminent, top scale player, and probably the next two rungs below.

The more interesting game is happening on the periphery - in some cases, being waged by the small players who are looking to find their way in. And it's not really competing mano-a-mano with the large investment banks, but it's really creating and developing new services, and then perhaps being acquired, or developing JVs with the larger players.

Some new products, new ideas, new business process work flows, in the electronic space, are all...

Greg: So, you're essentially adding value to the scale players, as a smaller player.

Carl: Exactly. And you see some... You see federation of EMSs and OMSs. At one level, they can provide differentiated value. And then you see complex event processors trying to provide infrastructure into that space. And you're seeing some of the larger data merchants looking to provide newer products - news, algorithmically tailored, for example.

So, you develop an ecosystem that feeds on itself. And that's an unhealthy evolution, and that kind of ecosystem really needs to develop for many players to survive and do well. And I think, it's important for the innovations to continue, not just by the scale players in looking to improve the margins, and by the small players in looking for new value-drivers, but in having them work together in creative ways as well.

Greg: Can you give us a feeling to where the US is now? Because you spent a lot of time in Asia. You're always backwards and forwards.

Carl: Just came back from Hong Kong.

Greg: So, can you give us an idea of where Asia is now, where Europe is now, and where the US is now, if we break it into those three trading areas, or banking areas, as it were? Because, obviously, Europe's got MiFID coming up.

Carl: Sure.

Greg: Can you give us a feeling for where they are and where the growth is? And does the growth differ from, say, Asia to where the US is now? Will it grow in a different way?

Carl: I think, there are enormous opportunities in both Europe and Asia to grow the market overall. We haven't come close to tapping the potential of both markets. The amount of algorithmic trading, even though it's grown in London in particular, it hasn't grown to the emerging Europe, and it hasn't been adopted as a top-two tool for most second-tier buy-side firms. So, there's a lot of growth potential in Europe.

In Asia, it's an even greener field. And, of course, in the major markets, like Japan, algorithmic trading is already very vibrant and very productive. It's the emerging markets. It's Indonesia and Korea and India, to a large extent. As the exchanges on those markets, and as the alternative markets develop - say, develop some fragmentation and increasing need for electronic trading - you'll see those markets grow.

Roughly, if you look at Europe, it's probably about six months to a year behind the US. And if you look at Asia, it's probably about two years. But, I would expect, a year from now, that we'd see that compress dramatically. So, I don't think that would remain at the

same pace of evolution. I think, they'd jump far ahead at that point. Not far ahead of the US necessarily, but much closer in terms of the penetration of electronic trading.

You have some unique advantages in Europe and Asia as well. You don't have the tradition and history of the US markets burdening as well, so I think, the adoption rates will be quicker. You see a lot of advanced North American hedge funds that are trading into both of those regions that are very heavy users of the electronic arsenals, and so they'll drive growth.

And you see a lot of creative, new players, in each of those respective regions, that are exploring tightly coupling their algorithms with their alpha generation, and that is exciting to watch. And we haven't seen that quite evolve on the scale that seems to be evolving in Asia.

Sinan: As a result of this, do you see significant pricing power shift to the buy side of it to the sell side?

Carl: The inexorable trend of pricing power to the buy side continues unabated. The pricing power on the transaction side is one aspect of the pricing power. I think that the sell side sees that both as an opportunity and as a threat, to be honest.

The opportunity is in terms of looking to service their clients better and better with more lower-touch products. But, it's also in how we wrap our products together, that I think that we're spending more time in looking for the packaging as a driver for maintaining the lower cost of the transaction side, but at the same time, making us appear as rich a partner to work with.

So, our partnership extends now into technology, into customization around algorithms, into looking at things like RAP business, into looking at unbundling, and so, servicing our clients in ways that would be unexpected and would be appreciated.

Sinan: In light of this, where do you think the proper place in the trade cycle is for risk pricing? Traditionally, we used to take risk pricing after the trade execution, sort of as a middle office thing. Is that changing significantly? Or will it have to change?

Carl: Oh, it's changed. It's moved deeply into the front office, and also very deeply into the product space as well.

So, the traditional kind of use of credit, on a single-name basis by a sales trader who's covering an account and who's looking at a client on a net profitability basis, and opportunistically as well, and the use of it in our portfolio trading, to manage a risk for our clients, particularly when there's an aspect of risk that they may not be able to manage with the flows that they have versus the flows that we see, those continue to evolve. And there's always a market for firms that can be aggressive in terms of the risk pricing.

An evolving market in terms of risk is the more dynamic, fluid market, as it relates to dark pools. And we're seeing the emergence of a new class of liquidity merchants that are providing, in effect, risk-based pricing, typically in small lots. And they do that to stimulate liquidity and develop a portfolio that they can subsequently hedge out of. But, they provide value as well, in providing more continuous liquidity for those firms that really need to execute.

So, the initial consumers of those liquidity merchants are sales-side algorithms. But, that's evolving. And I expect that we'll see that growth really drive trading in general, over the next two to three years.

So, more liberal use of risk, more use of risk in terms of smaller quantities, more use of risk in terms of small caps, particularly where you can hedge that. And the key driver is the use of internalization engines, and algorithmic access to those internalization engines to provide more broker to broker kind of risk trading as well as broker to buy-side.

Sinan: I just want to shift to the implications of MiFID and RegNMS, which are helping us, other than the obvious requirement of handling more data and turning around faster, what other impositions it has on algorithmic trading infrastructure.

Carl: I'll start with the data side, partially because it's easier, but I think, it's also less appreciated than most would argue. On the data side, we're seeing billions of ticks now. We're seeing massive amounts of tick data around the open and the close as well, so the amount of bursty traffic has jumped through the roof.

So, the amount of processing, just in capturing the data for our smart order routers, for our algorithms, has been talked about, but you really are fighting the laws of physics and the laws of what's achievable with current technologies. And distributing the problem is one tactic that we can all employ, but we're not maximizing our ability to contend with the rapid growth of data.

Then, as you work downstream, you see the challenges then turn into how you utilize the data. Can you develop algorithms that look at every bit of Cisco data, every level? If you've got 40 different price levels of Cisco, that's more than, on some days, the entire Italian marketplace, in terms of market data. Using all of that and looking for opportunities to find size deeper in the price infrastructure becomes part of the challenge.

That's a solvable problem. What's less tractable are the analytics that are really important to determine when data should appear next, or when prices may expect to come down. So, the predictive analytics - the analytics that count choppiness and count clustering of data - becomes a second-order problem that's still vexing most of the street.

With respect to MiFID, it's less clear that the rules are embedded systematically in a way that defines precisely how you need to react. So, because of the looser separation of the rules from the implementation, it gives the street a large swathe of capability in which to look to implement.

So, I think, over the next two years, we'll see some of the same evolution. We'll see rapid growth of market data in Europe, which will have this cascading problem downstream across other architectures and infrastructures. You'll see rapid growth of the adoption of newer technologies, like complex event processing.

But more importantly, you'll look to see, what are the new products that are going to drive value to our clients? Is it going to be regulatory-driven, in terms of the products? I don't think so. I think, it's more about how can we improve our clients' trading behaviors? How can we help them trade better than average against their benchmarks? What kinds of analytical byproducts do they need to see to feed their own algorithms, to feed their own alpha generation? And I think, it's that kind of space that we're looking to operate.

Sinan: What is, essentially, the initiatives that you need to take to retain a competitive advantage in developing algorithms?

Carl: I like to joke with our product team that they all have to think a little bit like Steve Jobs.

Sinan: [laughs]

Carl: It's not just about applying millions of dollars of capital behind our products and our infrastructure to be competitive. It's about coming up with new, killer products in new areas that we can synthesize value for our clients. So, part of it is how we think and how we collaborate with our clients and the kind of synergy effect that we look to achieve.

So, the kinds of meetings that we're excited about, with some of our clients recently, have been around how we can help them achieve better economies of scale as well. They're not sure, as a general theme, how to achieve better than the performance they've already achieved with their trading areas and how much should they expect to achieve. Is that an area they can maintain competitive advantage as their scale goes up, or is that an area they should look to achieve a cost advantage by outsourcing?

We're also looking to help our clients, not just with capital, but intelligent use of capital that reflects where we believe we have competitive advantages and synergies and where our clients can develop new capabilities. So, we're spending a lot of time in small caps, as an example, in the US markets.

Sinan: How competitive the emerging markets are right now? Is it the first-comer advantage's last?

Carl: First, second, and third-comers all have competitive advantages. I think, capital calendars are also hugely important and will remain important in emerging markets.

Risk trading is also a key element in emerging markets. The need and the ability to provide a risk-based price is important where... sometimes, your alpha forecast is for 400

percent of the average daily volume of a stock, or 500 percent, or 1000 percent. And if you look at the kinds of names that are exciting, they're typically second-tier names as well, where there's just insufficient liquidity to move ahead.

So, ways of synthesizing exposures, to use the VTFs, and creative ways of developing exposure to a particular name without quite as much idiosyncratic risk are some of the newer opportunities we're looking at as well for our clients. We have a strong derivative pedigree as well. And we're looking at ways of harmonizing our strengths in derivatives with our newer strengths in cash trading.

Sinan: In the emerging markets, or especially because of the emerging markets, have you identified new risk types?

Carl: There are new risk types. I think, it used to be about timing cost and market impact. Those were two twin pillars that most algorithmic trading has been based on.

And I think, if you look at what's happened recently in the credit markets, it hasn't opened our eyes to liquidity risk, but liquidity cost and liquidity risk is perhaps a different animal. It's not just about price volatility. It's about volume volatility. It's about timing of that volume volatility. It may be there today, and when you want to get out of your position, it may not be there tomorrow. And how do you reflect that into your own trading and into, not just your alpha generation, but on the risk side of the alpha generation? Most risk models don't really take into consideration the kinds of anomalies that we may see on a yearly basis.

It's not a Six Sigma event, typically, that happens when we have a liquidity crisis. And a liquidity crisis very easily moves across from one market, as a class, to another. So, you've got this contagion correlation effect that's massive. So, I think, it's important for all of us to develop new science and new tactics to really deal with that. And particularly, as you talk about emerging markets, there's no sphere that is as liquidity-sensitive as emerging markets is.

Greg: What's required, in terms of inbound data, to improve the identification of dark pools?

Carl: Certainly, tagging is huge. Inbound data to improve the identification of new dark pools of liquidity is not just about dark pools, but when you look at data - corporate actions with market data - they're not necessarily aligned, not necessarily provided by the same vendor. They're not necessarily oriented as well in time. We need them, not on a T plus zero basis. We need them at T.

Greg: [laughs]

Sinan: [laughs]

Carl: Not T plus 10 seconds. Actually, you need them before T. You need to plan your thinking and plan your trading.

So, tagging, or combining of data by, particularly, the larger data vendors, if they can provide more value there, that's a big plus. As more and more algorithmic trading happens in second and third-tier names, where there's less liquidity, less information, more errors, you don't want to have a situation where you're trading the wrong lot size because the information wasn't carried along, or the wrong kind of pattern develops because it went through a recent event and so its volume profiles are no longer representative of the way it trades currently.

So, tagging is meta-important. Combining data sets so that you can interpret them with the same systems is becoming more important. Establishing new sorts of data in combining news, for example, and news tagging into real-time streams of market pricing would be very valuable as well.

Sinan: We used to think of, when we look at front office, middle office, we would sometimes look at it as T plus zero events and T plus one events and so forth. Is that disappearing? Or has it already disappeared, that they're all T plus zero, basically, across the trading cycle?

Carl: I think, even in some markets where we're T plus three or T plus one, our middle offices on the streets are all looking to converge to T plus zero. There're cost economies when you know exactly where you are inside your own house, in being that way.

Unfortunately, you have a dependency on so many other institutions to get there. That trend also continues unabated. There's a deep economic desire, for most of the street and most of the buy side, to ever improve their costs. We'll see that trend really grow in Europe this year, as a result of MiFID, but certainly in Asia as well.

Sinan: What is the best architectural paradigm to distribute data in real time?

Carl: If you need tight synchronization of data, then having a shared memory space enables that. And it's simple logic, and if you need to do something very quickly, that represents the easiest opportunity to kind of develop something.

Complex event processing requires an enormous investment, in time and resources, in order to leverage the benefits of an environment. You've got to use it at the data sourcing side all the way through each element of extraction of it and transformation of the data. And so, it represents kind of a Holy Grail of opportunity, but you still are contending with performance bottlenecks in components of the architectures.

Messaging represents, still, a very important part of most fabrics, whether it's complex event processing, or where you've got distributed systems leveraging a common, shared memory space, you typically still need messaging.

The challenges on architecture are around security and around messaging. Especially, if you're trading globally, typical encryption on high-frequency messages is not the norm.

Strong standards exist, at least in a technical sense, but on an implementation sense, are typically rare.

Coupling that with extraction and the front-end technologies, I think, is really the bigger architectural challenge. What kinds of products are you developing? Are they web-based? Are they front-end, traditional-based? Are they thin client? Are they extremely thin? Or is there a processing layer in-between?

And in reality, what we're finding with many of our clients is, we've got Web 2.0 technologies implemented in more than one of our products. But, in some cases, some of those same clients that are using those Web 2.0. As a firm, a trader may want to use a Web 2.0 product, but their analyst may want to have a more complex interaction. So, now we're integrating at a services layer between firms, for certain kinds of analytics. And in some cases, we're integrating on a message-based layer as well.

So, the harsh reality is that all of those technologies are important, and it really depends on the problem you're looking to solve. The complex event processing represents the biggest, most cataclysmic jump, perhaps, of that space. But, we'll be seeing implementations of each of those technologies, along with asynchronous services and a few other newer technologies, also enter the fold.

Sinan: Do you feel that you need to have controls on data, that a certain people are being restricted from seeing certain data sets when data is distributed? And is there an established way of figuring that out or categorizing that, or it's still a work in progress?

Carl: We have a very rigorous data management protocol that exists within the firm. There are people who are trusted guardians - and they're called data guardians - of data, and they're responsible for managing the use and the security profile for any data sets that they've got that mandate for.

Then there's IRM policies [Information Risk Management Policies] that we all adhere to, and there are audits around. And for businesses like ours that extend from internal trading all the way to client products, there are lots of walls and lots of checks and balances and lots of protocols that we all follow.

And I think, it is continuously work in progress. We're always looking to refine our processes, both to continue to be nimble but, at the same time, try to be ahead of would-be miscreants in a marketplace that would look to take advantage of weaknesses in our protocols.

Greg: I guess, both with your internal people and with the external technologies, they have to start understanding this to a deeper degree than they've understood it already, to start architecting systems so that they can achieve the goals you're after.

Carl: One of my deep frustrations - and you've hit on it, you hit on a nerve - is the lack of real market structure knowledge that exists within technology. If you're a strong

architect and you're well-schooled in architectural principles and agile methodologies, and you've got a pure sense of how to solve problems from a technology perspective, you still may be facing an enormous gap in terms of translating those skills into solving the kinds of business problems that we're facing now.

We just don't need architectural skills. We need the creative problem-solving skills as well. And you really can't use your creativity unless you become, in a Zen-like way, somewhat immersed into the fabric of the problem. And I think, that's a challenge sometimes. And sometimes, I find myself pulling people over the fence to force immersion, and sometimes I'm going over the fence to teach, or looking for ways to kind of make the two sides meet.

There are many technologists on the business side, which helps, within our world, and there are a few business analysts on the technology side as well. And I think, that's part of the solution to the problem.

Greg: Thank you very much for your time today, Carl.

Carl: Thanks a lot. Thank you.

Greg: Very enjoyable talking to you.

Carl: Thank you.

Greg: Thanks. Bye-bye.

Announcer: That was an interview with Carl Carrie, the head of product development in the Electronic Client Solutions Group at JP Morgan. The Algorithmic Trading Podcast, sponsored by Sybase, was brought to you today by Voices in Business, helping our sponsors to achieve thought leadership in their business sector. To find out more about Voices in Business, visit www.voicesinbusiness.com

And as we mentioned at the beginning of this podcast, we strongly encourage feedback; so please submit your own questions, either via email or in the comments section of the website, at www.algotradingpodcast.com

Thanks for listening. Good bye.

[music]