

## WHITE PAPER SERIES

### *Performance: What is real-time?*

by Tim Dodd, head of product management, SunGard Front Arena

## Performance: What is real-time?



Every firm faces financial pressures right now but few are willing to sacrifice performance in their trading. Tim Dodd, Head of Product Management for Front Arena, demonstrates how the latest technology can be used in a strategic way to extract the right level of performance without exceeding budgets.

In an industry where speed is everything, a lot is made of real-time. The ability to run calculations, to receive data and, ultimately, to trade in real-time is often cited as a necessity in today's competitive markets. In this paper, I explore the balance between the benefits of speed and the cost of hardware and software to achieve that speed.

### What benefits does speed bring?

Without doubt, being first to react correctly in financial markets gives you the edge:

- if you are first to provide the best price on some derivative pricing in an OTC competitive auction you will be first to trade
- if you are the first to spot and take out a mispriced instrument in an electronic market place, you gain the mispricing
- if you are first to re-hedge to capture advantageous and but fleeting movements in the underlying instruments of a derivatives portfolio, you will be the one to realise that profit

And so the list goes on.

In an ideal world, you would aim to have a system that gives you first mover advantage with questions answered instantaneously and automatically. However, we live in the real-world where hardware cost rises as you try to achieve real-time and portfolio complexity only amplifies this effect with more complexity implying higher cost for the same performance.

This is true even though hardware costs have fallen (a powerful stationary workstation can cost as little as 200 Euros) since the software, networking, implementation cost as well as maintenance all need to be considered when planning a trading solution.

But here's a thought: in many ways the industry has misunderstood what 'real-time' actually means. Real-time functionality is associated with a constant stream of

updated market data that enables traders to revalue their portfolios with immediate effect.

For certain instruments and markets this quality of data is available: when market-making simple derivatives on electronic exchanges access, you can have prices of derivatives and their underlying instruments readily available and you can calculate and monitor theoretical values in real-time.

However, when it comes to other instruments such as interest rate derivatives, that availability of real-time market data is not as timely. Prices are more illiquid and deducing current implied volatility from brokers is impossible. Even in the most liquid of these markets or currencies, price information is updated every 15 minutes at best and even then prices are only indicative.

Consequently, the most expensive and sophisticated software and technology will not alter the simple fact that the supply of market data is too slow to give a complete 'real-time' picture. Investing in real-time, full re-pricing technology for trading in a market where assumptions might only be updated every 15 minutes can be a waste of resources.

This is not to say that trading systems should simply lay dormant. Instead, high quality and high frequency approximations can be made using a series of 'Greeks' which allows the simulation of movements of your portfolio values based on sensitivities to underlying prices. Specifically, first order sensitivities to prices are calculated using mathematical models and then multiplied by changes in those prices. For small

movements, this best assessment of changes in value can be calculated rapidly (in linear time) but with sufficient accuracy to trade.

For larger movements or when precise market data is available, such approximations can be updated by recalculating implied market data, recalibrating models and running full calculations.

Of course, for other trading activities, many institutions do not need real-time. If your investment horizon is longer (you run “buy-and-hold” strategies) or have very small sensitivity to underlying prices then there is no point in making the investment.

However, today’s markets are once again more volatile, encouraging increased in-and-out trading and trading in instruments with “optionality” and a high sensitivity to underlying prices:

- With the rise of the internet, we have seen a rise in small ticket size stock trading and a consequent need for fast market making systems to consolidate net positions.
- Short maturity options or barrier options trading near their barriers amplify market movements significantly so we have seen greater demand for systems that can calculate second order effects rapidly to be able to realise profits that appear albeit fleetingly.

It is important that system purchasers realise there are some factors – such as the delivery of market data – that are beyond the control of technology vendors. A trading system that works in real-time across all asset classes is not a realistic demand from your traders.

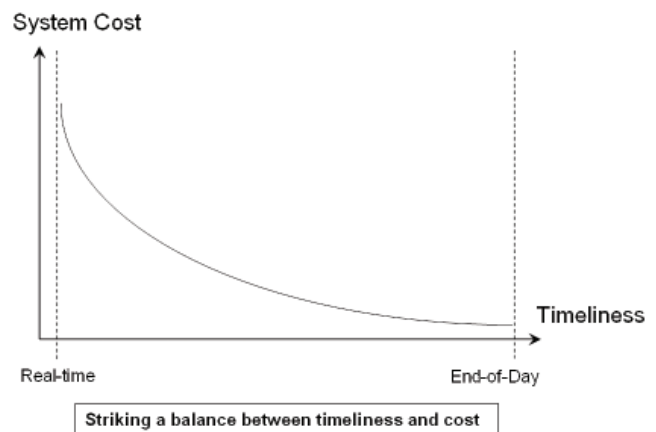
A balance must be struck between cost of system for a given type of trading and the improvements this will bring in trading profitability.

This is particularly true in the current climate, where budgets must be carefully constructed and the maximum value must be extracted from every dollar spent.

### The best performance money can buy?

With many firms operating in an environment where forensic attention to cost is evident, the importance of a trading system that employs all of the latest techniques and strategies to improve performance through other means than simply increasing the amount of computing power, cannot be overstated.

A choice needs to be made as to where on the spectrum of timeliness against cost a system should be deployed.



So what are the most recent technology advances that can you turn to maximise return on investment? Let’s look at the latest development in techniques to see how most recent changes should influence performance strategy.

### How can grid computing help you today?

Grid computing, where processing is spread over a number of machines and a number of networks, has been one of the most important developments in the technology world over the last decade and has been integral in the efforts to increase performance levels without having to employ excessive computing power.

However, as with real-time technology, grid computing is not applicable to every trading scenario and its use has to be justified. For example, when splitting up a transaction or calculation into several components to be spread over a ‘grid’ of computers, the size of the process has to be large enough to justify the procedure. Equally important is the timeliness of the transaction, particularly if the time it takes to chop up the components and rejoining them outweighs the time in which the calculations must be made.

Gridding has become cheaper as costs and performance of hardware have come down but ready availability of multi-core machines is really improving the ease with which calculations can be split out as there is no worrying about network latencies and guaranteed messaging between machines.

### How can multi-threading help you today?

The comparative strategy to grid computing is multi-threading, where a series of smaller calculations can be spread across a number of processor cores but all within the same machine. How many cores does your own computer have today?

The decision on whether to use gridding or multi-threading, or to stick with one processor, is often dependant on the instruments or asset classes involved, their complexity and the quantity. Vanilla products should be ably catered for by a single processor. More exotic instruments require more complex pricing (using Monte Carlo simulations or similar) and are ideal for either gridding or multi-threading.

Of course, all this will not help you if your software can not take advantage of multiple threads: a single threaded product will run at the same speed on a dual core machine as a single core machine.

To take advantage of threading, software writers need to ensure that code is thread safe (i.e. that only one core can process a piece of code and associated variables at any one time) and ensure presented totals make sense (i.e. all calculations dependent on a specific price will only be totalled if the price is consistent across the calculations). Although new programming languages are coming out that can deal with many processors more easily, this is no small task in the world of derivatives.

Processor manufactures will only continue to push this aspect of their products and we will continue to see subsequent improvements in software: single processors cannot run much faster than today as they are restricted by the speed of light so parallel processing is the only way to go.

### Dealing with volume

The quantity of trading activity also carries performance-related implications by virtue of the number of calculations it is necessary to carry out. A hedge fund may invest in straight-forward products but such is the frequency of their trading that millions of simple calculations must be executed.

To put these quantities into some kind of perspective, a typical large trading system has to cater for approximately 600,000 trades a day, a figure that amounts to 17 transactions a second. Moreover, the average trading day is anything but that with

unpredictable volume spikes of activity and troughs of inactivity appearing throughout the day. Systems have to deal with this unpredictability and pull data in from all of the various available sources and make them visible to traders.

Dynamically accumulating many previous trades in a single instrument into a single position can make handling these volumes possible. Systems can then conduct calculations on the position as apposed to once on every trade. Of course, some smarts are required to unwind accumulations to deal with any bad data on a single trade, but the benefits of reducing millions of trades to several thousand accumulations are obvious.

Emphasis in this area has only increased as more algorithmic trading is conducted by the market and software must continue to evolve to meet higher volumes. Today, in the US more than 60% of transactions are initiated by a computer algorithm and number of transactions has increased substantially but size per transaction has fallen. Exchanges have certainly invested in infra-structure and systems to deal with more volumes of price, order and trade data.

Of course, there is another dimension to dealing with volume and that is complexity: a structured product may be a single instrument but such is the nature of its mathematic model that thousands of pricing calculations must be carried out to determine value and sensitivities.

Some convertible bonds need millions of likely future market simulations to price and this can take several hundred milliseconds to compute. To run "what if" analyses, even more calculations are required.

Often these trades have long maturities and are bought and held and so systems need to offer increased computing power to deal with this dimension of volume.

### Performance tools

Bottlenecks can appear anywhere in the technology topology of an organisation. If you spend millions on the best systems they are only as good as their weakest link. Consider spending something on real-time monitoring.

In order to get the maximum performance from a trading system, it helps to have a series of tools that can monitor and measure performance levels making it easier for developers and system managers to resolve any memory or CPU-related issues.

A real-time profile that graphically and dynamically displays all calculations along with how much CPU time is needed to run them can help users to identify short-term problems and devise longer-term performance improvement strategies.

Additionally the provision of cross system monitors allow purchasers to plug in real-time updates to a single consolidated dashboard.

### A new paradigm?

In the constant quest for greater performance, the vendor community has to look beyond its own industry to find the most cutting edge techniques and apply them to the trading technology market. One such example can be found in the increasing adoption of graphic cards in trading systems. These cards are commonly used in the gaming industry and have thousands of very simple processing cores on one powerful card.

The technology is still at a fledgling stage with limited deployment within the trading world so it remains an expensive form of parallelisation. But for large exotic basket option market makers charged with the task of

providing real-time exchange prices growing number of clients, such technology may prove to be invaluable in the near future.

### Pushing technology to its extremities

When it comes to exchange trading, exchanges and traders have pushed their software to the limits. Co-location of the technology to allow exchange members to quote on exchanges with minimum latency by allowing members to locate their machines in exchange premises shows that the speed of light is now an issue.

Interestingly, this puts “best execution” commitments required by regulators in an intriguing light. If you have an instrument traded on two markets there is no way for a best comparison can be made on the prices on the two markets as a message much be sent to some consolidating spot from each exchange. This takes time (over and above the message to actually execute) and so strict “best execution” opportunities can only be missed. Of course, in reality, regulators are happy with transparency and best efforts as this physical location problem cannot be overcome.

---

## Conclusion

The trading arena has always been known as a fiercely competitive landscape and its participants are constantly looking to push technology to its maximum. For those traders engaged in high-frequency trading, arbitrage-based strategies or complex structured products, their decision-making is based on the ability to see all of the information they need in a timely manner and in one easily accessible destination – such as a desktop.

They want to know where they are so that they can navigate their way through the market.

This is a great position to be in but it can come at a price. The technology is out there. Things can be seen in real-time but firms have to consider how necessary this is. Conversely, firms cannot afford to be too casual in their calculations and risk being left behind by the market. It is ultimately about achieving a workable balance.

Even when making allowances for the current climate of cost-related conservatism, it is only to be expected that traders will push technology to the boundaries of what is possible and affordable but firms have to make strong decisions and be loyal to the performance strategy they have chosen. This does not mean picking some arbitrary point between the brutal expense of technology’s cutting edge and the disaster of being left behind by your better performing competitors.

Instead firms should put their performance criteria in the context of both their competitive ambitions and their particular strategies. This will, of course, be different for every trading participant. A market maker will have different priorities to a trading desk manager who will again have different priorities to a frontline trader. But the underlying principle is the same – that technology and its performance should be used to help firms control their positions, the more real-time the higher the cost.

## **About Front Arena**

Front Arena provides straight-through processing by integrating sales and distribution functions, trading capabilities, risk management, as well as settlement and accounting support. Market makers, brokers, risk managers and, institutional asset managers use Front Arena to trade equities, fixed-income, interest rate derivatives, and credit. Front Arena's components can be configured to meet a variety of trading needs and scaled to support operations ranging from local to global.

To find out more about Front Arena visit [www.sungard.com/frontarena](http://www.sungard.com/frontarena) or call an expert +44 (0)208 081 2779

## **About SunGard**

SunGard is one of the world's leading software and IT services companies. SunGard serves more than 25,000 customers in more than 70 countries, including the world's 25 largest financial services companies.

SunGard provides software and processing solutions for financial services, higher education and the public sector. SunGard also provides disaster recovery services, managed IT services, information availability consulting services and business continuity management software.

With annual revenue exceeding \$5 billion, SunGard is ranked 472 on the Fortune 500 and is the largest privately held business software and services company on the Forbes list of private businesses. Based on information compiled by Datamonitor\*, SunGard is the third largest provider of business applications software after Oracle and SAP. Continuity, Insurance & Risk has recognized SunGard as service provider of the year an unprecedented five times.

For more information, please visit SunGard at [www.sungard.com](http://www.sungard.com).