ONETICK®

Enabling the Success of Quantitative Research and Trading

In the financial industry we are in a race, not to the bottom or to the top or the white tape of a finish line. It is not a race to zero or an arms race as some would have you believe. It is a race to succeed. A trading firm's final goal is to win, to be successful in outpacing the market. Firms operate in a fiercely competitive industry where success is measured by profit. In some sense it is a race against time and the road is filled with obstacles – firms need the tools and technology to trade smarter, dig deeper into data and ultimately outmaneuver the competition. In an industry where the thinnest of margins separates success and failure, being smart about technology decisions is a means to be smarter about trading. The race is on; a long running marathon made up of sprints, hurdles and the occasional steeplechase. Being equipped to run it can mean the difference between a thriving business and a fool's errand.

What Fosters Ingenuity

Quants apply an empirically-tested and rules-based approach to exploit market inefficiencies manifested by human behavior, geo-political events and market structure. Just consider that the U.S. market structure is the aggregate of public exchanges, broker-sponsored (dark) execution venues and alternative trading systems (ATS) with no one destination executing more than ~25% of the total traded volume. It is a complex fragmented web to navigate for liquidity, controlling costs, trading patterns and behavior.

With tighter spreads, thinner margins and lower risk appetite, quantitative traders are exploring more cross asset trading models and cross asset hedging. Consequently, the quest for new and revised models is never ending. The side effect of this is increasing demands for deep data over longer time periods across a multiplicity of markets - equities, futures, options and of course cross-border currencies. This data dump is the fuel feeding automation technology, quant's research and strategy modeling.

The challenge lies in managing this increasing volume of data to ensure its overall quality and accuracy for trade-related decision management including econometric trade modeling, risk and transaction cost analysis.

The need for high quality of data cuts across all aspects of the trade lifecycle including model back-testing, portfolio mark-to-market and compliance. They all depend on a high-degree of data quality, where accuracy is vital to determining outcomes.

As we see more volume and market data, we can expect that additional flow will lead to greater message traffic entering the market place as orders and executions. The increased throughput will burden older technology in an attempt to handle the load. Firms will realize it is not simply a matter of buying new versions of old technology to keep up. And the focus will shift towards analysis – the ability to link disparate data sets under some common thread to tease out intelligible answers, a diamond from a mountain of coal.

What Is Time Series Data; CEP, Tick Databases and Applied Analytics

Time series refers to data that has an associative time sequence, a natural ordering to its content. For financial data, it is identified by symbol and type including rates, prices, curves, dividend schedules, index compositions and baskets. The temporal ordering allows for distinct data analysis revealing unique observations and patterns and the possibility for predicting future values. Extracting meaningful statistics such as a time-weighted average or linear regression are unique to and dependent on the data's temporal nature.

Analyzing historical time series data is to use observations from the past to characterize relationships, the interplay between trades and quotes for example. If we want to use those historical relationships to explain future developments we have to assume that the future behaves like the past. This is the basis for back-testing algorithmic trading models whether for determining the profitability of a statistical arbitrage pair's strategy or modeling correlations of the Dow 30 Stocks over the past 4 years.

Fundamentally time series are finite or infinite sequences of data items, where each item has an associated timestamp and the sequence of timestamps is non-decreasing. Time series are often called data streams; the two terms are used interchangeably to represent real-time data, such as intra-day streaming trades. This is a common notation used in Complex Event Processing (CEP). A time series 'tick' database on the other hand is optimized for handling time series data for archiving history.

Historical time series data provides a context to the analysis associated with real-time complex event processing. Consider the scenario where there is a need to understand historic price volatility to determine statistical thresholds of future price movements. This is helpful for both trade models and transaction cost analysis.

The ideal case for CEP analysis is to view historical time series and real-time streaming data as a single time continuum. What happened yesterday, last week or last month is simply as extension of what is occurring today and what may occur in the future. An example may involve comparing current market volumes to historic volumes and prices for trade execution logic. Or the need to act upon live market prices may involve comparisons to benchmarks that include sector and index movements, whose intra-day and historic trends gauge volatility and smooth outliers.

The OneTick Solution

OneMarketData provides market data storage and analytical solutions that enable financial institutions to outmaneuver the competition. Our OneTick product is the only solution that combines complex event processing and time series tick database spanning both historical and real-time in a common paradigm.

OneTick can capture and store market data from any asset class including foreign exchange, equities, fixed income, futures and options. It can also store orders and executions for creating fully customized transaction cost analysis in a multi-asset or multi-bank environment. OneTick can store, aggregate and consolidate price data… trades, quotes and order book depth from numerous sources providing the basis to the price discovery process – vital for understanding liquidity patterns and market behavior.

Data storage can span from a few months to multiple years with no loss of performance.

The product comes supplied with over 130 built-in high performance functions for data analysis. OneTick also embeds the mathematical libraries of R and MATLAB. In analyzing data streams, they can transform the market data dump into actionable trading solutions.

OneTick inherently understands financial data including complex symbologies and order book data structures with nanosecond precision… which means hard problems have easy solutions such as creating specialized FX liquidity models, re-pricing solutions and low-latency arbitrage. OneTick's built-in functions are at your design fingertips for converting currencies, pricing portfolios; manage continuous futures contracts along with dozens of others.

Quants, algo developers and savvy traders, the users of OneTick access this library in a visual query modeling tool. This is a means to construct the semantic logic of a real-time CEP or historical analytical query. Those queries are run as close to the metal as possible, the OneTick server is both a storage engine and the analytics engine, executing data retrieval from archive storage, in-memory and real-time sources. Query analytics are immediately applied; only final results are passed to requesting clients.

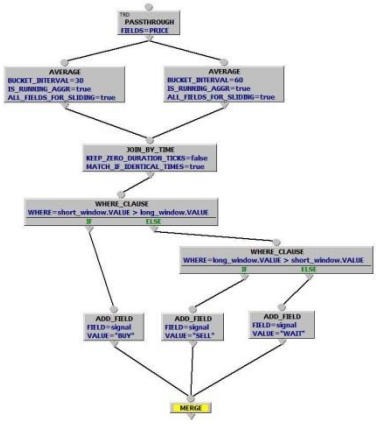

Those clients can access OneTick queries through APIs which are available in all common programming languages. And they include the familiar interfaces of SQL and ODBC. Visual tools such as Panopticon, Microsoft Excel and Spotfire can easily consume query results.

The advantages of OneTick mean you can confidently be first to market with profitable, back-tested strategies constructed in an intuitive visual way. Philippe Lüdi, VP at AlphaSimplex, a research-driven asset management firm "… *wanted a superior data management system that would allow us to develop and execute intra-day trading strategies for today's fast moving global markets and OneTick really impressed us*". Just as AlphaSimplex discovered, the OneTick solution provides superior technology and low-cost of ownership in the highly competitive race to succeed.

OneTick, the Details

Architecture

OneTick is a specialized database optimized for capture, storage, retrieval and analysis of extremely high-frequency time series (tick) data. It is built to specifically handle data in the financial markets domain, such as instrument prices, trades, quotes and order book information. OneTick has pre-defined constructs called "tick types" to support these financial market-specific concepts. The OneTick database is a proprietary technology which includes an in-memory database holding intra-day data and a file-based archive store. It can be configured for row or column storage.

Modeling data schemas only applies to the time series data notion. Unlike in a typical relational storage engine, there is no need to wrestle with a vast amount of decision points for the design of tables, schemas, relationships, indexing strategies, etc.

Managing Order Books

An order book is an organized collection of market depth. OneTick time series as a consumer of exchange order book data represents that market depth as a series of price level modifications. Updates which can be full book snapshots or incremental changes can modify an entry's price & size or remove an entry. This architecture creates a flexible model for reconstructing books from history and consolidating books from multiple sources. User queries can request order book data starting from any time of day for any level of depth. Reconstruction is fast and efficient.

Reference Data

A complement to time series 'tick' data is reference data. Some common examples of reference data include alternative security identifiers, holiday and event calendars and corporate action information. Reference data is vital for elevating the accuracy and quality of the corresponding tick history. OneTick provides a complete reference data architecture directly integrated into the database for exchange calendars, corporate actions, cross symbology maps and rolling symbologies (i.e. futures contracts).

Data Compression

To efficiently manage all storage subsystem types OneTick provides not only a means to compress data, but also compression tuning. The granularity of compression starts at the individual field. There are multiple compression types resulting in archives between 3-5 times smaller than source data.

Sophisticated User Tools, Services and APIs

In addition to the core database engine itself, the OneTick software includes numerous other functional components: data capture components for both live markets and batch loading and fully integrated complex event processing (CEP) and analytical engine to complement historical analysis. A wide variety of programming APIs are available in C++, Java, C#, Perl and Python. These allow for custom data capture components, data query clients and custom analytics.

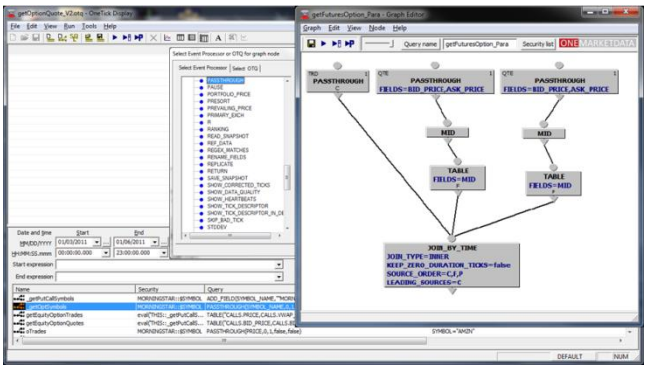

Lastly is a user interface component , a fat-client GUI tool to enable convenient querying of the data stored in OneTick's databases (i.e. in memory/archive store)and real-time feeds. It includes a visual browser that allows users to navigate to the desired dataset, as well as query building tools which allow users to build queries by visually linking together a sequence of functional building blocks known as event processors (EPs) into a acyclic directed graph. When executed in the front end GUI, query results can be shown in table form, chart and graphs, as well as support for exporting results into external formats such as MS Excel.

Analytics

OneTick provides a large collection of built-in analytical functions. These functions referred to as Event Processors (EPs) are semantically assembled in a query and ultimately define the logical, time series result set of a query. Event Processors include filters, transformers, joins & unions, statistical and numerical finance-specific functions, order book management, sorting and ranking, and input and output functions. Together these allow time series tick streams originating from any of the OneTick storage sources (archive, in-memory or real time) to be filtered, reduced and/or enriched into the business logic supporting a wide variety of use cases.

The Business Impact of Latency and Data Quality

The need for high quality from captured market data, orders and executions cuts across all aspects of the trade lifecycle from quantitative research, algo development and deployment, risk, post trade clearing, settlement and TCA. Those trade-related solutions depend on a high-degree of data quality.

Market data comes in many shapes, sizes and encodings. There are 13 major exchanges in the U.S. alone. Across Asia are 15 different markets with different technologies, cost structures, regulations and cultures.

That's a natural barrier to algorithmic trading and it creates unique challenges for efficiently trading across them - recognizing their differences in both market microstructure and regulation. Yet, adoption of algorithmic execution as a percentage of total U.S. equities trading volume has increased from approximately 28% in 2004 to just over 50% in 2010, causing all participates to chase an ever tightening pot. The determinants of price discovery, volume, and trading patterns define the structure unique to a market, an asset class and geography influenced by participants and current regulation. The collection and analysis of data across the wide geography demands sophisticated tools and systems as firms look to discover the unique determinants of transaction costs, prices, and trading behavior to devise smarter trading models.

The challenge of achieving timely data quality in this data dump is dealing with the vagaries of multiple data sources and managing a sea of reference data, such as corporate action stock dividends or ingesting cancelations and corrections. Any and all of these factors are vital to data quality and the science of quantitative trade modeling.

Invesco, an Asset Manager found this out as they began to branch out, expanding their business to be multi-asset, multi-broker. They needed to improve their business visibility and manage trade costs. Invesco looked to OneTick to build a customized TCA solution to accurately measure costs for the purpose of executing orders in a more efficient way. "*We wanted to increase both the level of sophistication of our analysis and experiment with new approaches. OneTick empowered us to take control of our TCA*", said Kevin Cronin, Invesco's Global Head of Equity Trading. Profiling trader and broker behavior, looking at intra-day execution efficiencies and highlighting the impact of delays on strategy performance made all the difference in Invesco's algo profitability.

Latency management is the ante to the game in the trading business, it is evolutionary not revolutionary. The evolution in technology goes hand-in-hand with trade innovation. Firms have to closely monitor latency; it is a tax everyone must pay to play. Good models and a deep understanding of the market through deep data research, coupled with infrastructure that does not constrain is the goal.

To derive true business significance, it's important to understand it centers on time-sensitive data quality. The scope of that deals with extracting effective value from large data sets and real-time feeds.

Tyler Capital, a London-based technology-driven proprietary trading group focused on global futures, recognized in order to increase profitability their business demanded effective data management only achievable through OneTick. To achieve alpha and discover new investment opportunities demands sophistication in tooling, analytics and data management. OneTick provides Tyler with all those capabilities to take their business to the next level.

The quality of data, its timeliness and accuracy ultimately defines an end game, that Holy Grail for sustaining the profitability of the business. Firms should not lose sight of that.

Coastal Management, an investment firms focused on systematic quantitative trading began using OneTick in production at the end of December. Coastal's managing director says OneTick meets the needs of a broad range of quants, "When traders come on board, their first questions are always, 'What is your historical data like? How rich is it? How easy is it to access?' -- and those aren't a concern anymore with OneTick."

Optiver, one of the leading liquidity providers in Europe, the US and Asia Pacific will employ OneTick's powerful analytics tools for their quantitative research and strategy development. OneTick is the first vendor solution of its kind used by the firm. Generating trading signals based on both historical and real-time tick data is a key requirement, and Optiver will leverage OneTick's CEP and time-series tick database for this purpose. As mentioned by the firm's project manager - "OneTick's comprehensive functionality, high performance and intuitive design combined with its ability to draw on real-time and historical data will help us to identify and act on unique trading opportunities throughout Asia".

Managing the Race to Success

Across the world's financial centers countless firms ranging from marketplace data providers like SIX Financial, sell side institutions like Société Générale, Invesco Asset Management and quantitative trading firms like Tyler Capital have recognized the broad and in-depth value of OneTick. Whether researching trade performance, devising new quantitative trade models or working to understand behaviors and patterns in the different aspects of high-frequency, stochastic and fragmented markets, OneTick is the complete solution for optimizing quantitative decisions to devise smarter trading models to enable and sustain the success of your business.

OneMarketData, LLC
2 Hudson Place
Hoboken, New Jersey, USA
+01 201.710.5977 (New York/New Jersey)
+44 (0)7956 370 340 (London)
+81 (90) 9826 4136
+852 28977798 (Hong Kong)

http://www.onetick.com