

Solution Brief

Market Sensing - Discovering important relationships in financial data to reduce risk and maximize gains

The Business Problem

Financial markets are rarely predictable. What moves a price one day might have no effect the next, or it might be felt several steps away from where it's expected. That's where market sensing plays a role.

Broadly defined, market sensing is the ability to bring as much relevant information as possible to bear on trading and risk decision-making. Going beyond traditional sources of information such as price feeds and price-based analytics, enhanced market sensing allows traders and risk managers to consider the impact of news and its ability to move markets.

The first step to market sensing is to provide participants the ability to observe a filtered set of news items that could potentially impact their existing positions or trading outlook. To do so requires the ability to analyze news feeds, extract entities of interest and constantly update a model that relates these entities to existing positions directly or indirectly – in other words, a dependency graph. It also requires the ability to identify and analyze hitherto unknown, interesting relationships among entities mentioned in news that could provide the basis of new trade ideas or risk responses.

The Technical Challenge

Risk and trade decision-making in financial services combines ever-decreasing windows of time and opportunity with ever-increasing volumes and types of data. Market sensing at scale is no different – it is only valuable if done in a timely, repeatable manner with immediate impact. Existing database technologies, being fundamentally disk I/O based are simply incapable of the responsiveness required for market sensing.

While Complex Event Processors exist and seem promising, CEP engines are typically proprietary solutions that work well on very specific problems where the underlying data structures are known in advance. Market sensing requires the ability to constantly adjust the domain data model as new relationships are discovered.



The Urika™ Solution

The Urika graph analytics appliance with its massive parallelism and extremely large shared memory is ideally suited for scalable market sensing. Urika can be the destination for a semantic extraction pipeline that extracts the dependency structure from news items. Simultaneously, existing trade portfolios can be expressed as a graph as well, hence providing a unified graph model of both existing market positions and related entities that can be updated as news arrives.

Urika's speed and thoroughness is directly related to its huge, globally shared memory of up to 512TB which holds the ever historical and current database resource in-memory and, using the SeaStar™ chip which can handle up to 350TB of I/O per hour, can quickly update relationships dynamically as new data becomes available.

The massively multi-threaded architecture of the Threadstorm™ (128 independent threads) is specially designed for analyzing graphs and allows threads to continue executing even if some are waiting for data to be returned from memory. Plus, the extreme flexibility of the Urika system architecture lets I/O or compute nodes be independently added as the application demands. This design facilitates relationship exploration. With scalable market sensing, analysts can go beyond retrospective analytics based on traditional market data to create, validate and modify hypotheses for trade/risk decisions during the course of a single trading day. This allows for unprecedented responsiveness to changing market conditions and provides a powerful predictive complement to existing front-office trade execution and risk analytics capabilities.

About Urika

YarcData's Urika is a big data appliance for graph analytics. Urika helps enterprises gain business insight by discovering relationships in big data. It's highly-scalable, real-time graph analytics warehouse supports ad hoc queries, pattern-based searches, inferencing and deduction. Urika complements an existing data warehouse or Hadoop cluster by offloading graph workloads and interoperating within the existing analytics workflow. Subscription pricing or on-premise deployment of the appliance eases Urika adoption into existing IT environments.

About YarcData

YarcData, a Cray company, delivers business-focused real-time graph analytics for enterprises to gain insight by discovering unknown relationships in big data. Adopters include the Institute of Systems Biology, the Mayo Clinic, Noblis, Sandia National Labs, as well as multiple deployments in the US government. YarcData is based in the San Francisco bay area and more information is at www.yarcdata.com.

