About Kx, a division of First Derivatives plc

kx

Corporate Information

- Founded in 1996
- Listed on London Stock Exchange (FDP)
- Proven in capital markets over 20+ years
- Global coverage with 2,700 employees
- 16 offices globally

Known For

- World's fastest time-series database
- Rich analytics on real-time, streaming and historical data
- Low latency technology stack
- Vertically and horizontally scalable solutions

High Frequency Trading

kx

The world's fastest real-time database leveraged to optimise trading strategies for one of world's largest investment banks

Business Problem

The bank needed to calculate trading profitability per client, analyse liquidity in electronic versus manual trading and evaluate trade decay over pre and post-trade movements.


Kx Solution

Provided fast and reliable technology for capturing and storing trade quotes and orders, running algorithms and providing users with a browser based UI to interact with the data:

1. Captured price quote/tick updates for orders, trades and rejections rates from external and internal sources.
2. Analytics and browser based dashboards with the ability to back test and fine tune parameters.
3. Both on-shore and off-shore development and support was provided.

Results

Improved visibility into rate volatilities
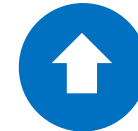
Significantly increased quote and trade info

Improved client flow analysis and profitability

Winning on the Track

Aston Martin Red Bull Racing needed a competitive edge for F1 racing. Selected kdb+ as their data and analytics engine.

Business Problem

Red Bull needed to optimize the performance of their race cars while performing wind tunnel experiments.

https://youtu.be/oycxhxKH0fA

Kx Solution

Provided superfast ingestion and analytic capabilities:

1. Ingested 2000 different sensors up to 500Hz frequency. 270GB/day rates using kdb+.
2. Utilized our programming language, q, to calculate realtime responses and corrections to the race cars while testing in the wind tunnel.
3. Plans to take the data ingestion and analytics to the race track for realtime monitoring and corrections.

Results
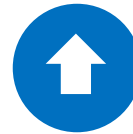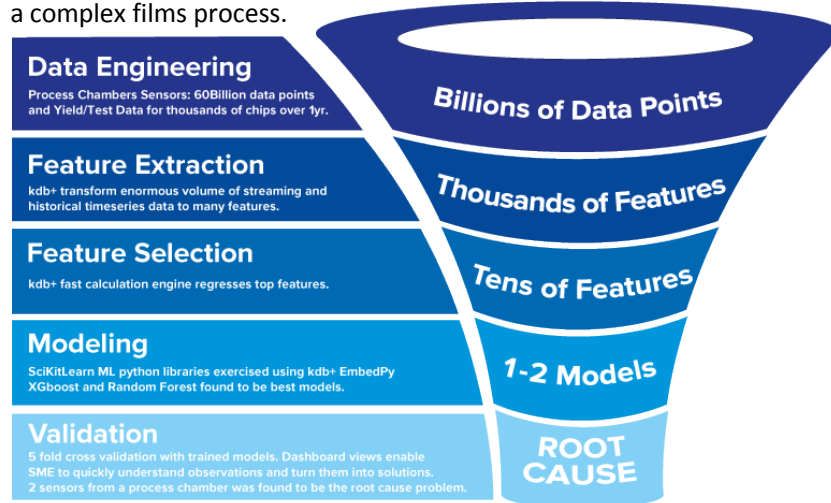
Super-Fast cycles of Learning

Visuals to Interpret Results

Selected as an Innovation Partner

Improving Root Cause Analysis

kx

kdb+ and Kx engineers using open source ML libraries on a single server enabled 2 year project to be solved in weeks.

Business Problem

A semiconductor fab had a chronic yield problem originating from a complex films process.

Data Engineering
Process Chambers Sensors: 60Billion data points and Yield/Test Data for thousands of chips over 1yr.

Feature Extraction
kdb+ transform enormous volume of streaming and historical timeseries data to many features.

Feature Selection
kdb+ fast calculation engine regresses top features.

Modeling
SciKitLearn ML python libraries exercised using kdb+ EmbedPy XGboost and Random Forest found to be best models.

Validation
5 fold cross validation with trained models. Dashboard views enable SME to quickly understand observations and turn them into solutions. 2 sensors from a process chamber was found to be the root cause problem.

Billions of Data Points

Thousands of Features

Tens of Features

1-2 Models

ROOT CAUSE

kdb+ and Kx engineers using open source ML libraries on a single server enabled 2 year project to be solved in weeks

Kx Solution

1. Setup a kdb+ cloud environment with small team of Kx data scientists and customer engineers.
2. Performed extensive problem identification and data management using kdb+ and Kx dashboards.
3. Executed data management and supervised machine learning at 340x faster than Hadoop with 1/12th of the hardware.

Results

Improved Yield

Fast Cycles Of Learning

Visuals to Interpret Results

Space research with NASA FDL

Leveraging best of breed machine learning in Python, with Kx unique ability to process and analyse mass volumes of real-time and historical data

The Problems

- Tess Exoplanet detection - classification of TCEs as real planets or false positives and predict with a level of confidence
- Space weather predictions - predict scintillation events, using historical GNSS data and live data to improve the window of detection
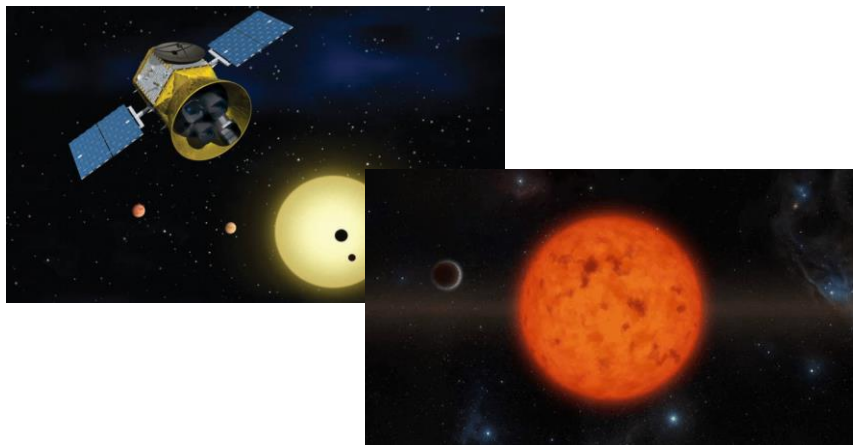

Kx Solution

- For a long time exoplanets were identified by humans looking at the light curves and deciding if the drops in brightness detected in the curves were caused by the transit of a planet. This process was very slow and involved many resources devoted to it.
- The solution achieves a improved detecting capability. Also, and more importantly, it achieves a high precision too, which saves money and time since further analysis of false detections is avoided. In addition, since confidence in predictions is also provided this allows researches to focus on what's important

Results

Improved visibility and detection on less infrastructure
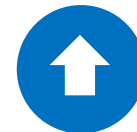
Fast historical and access for training

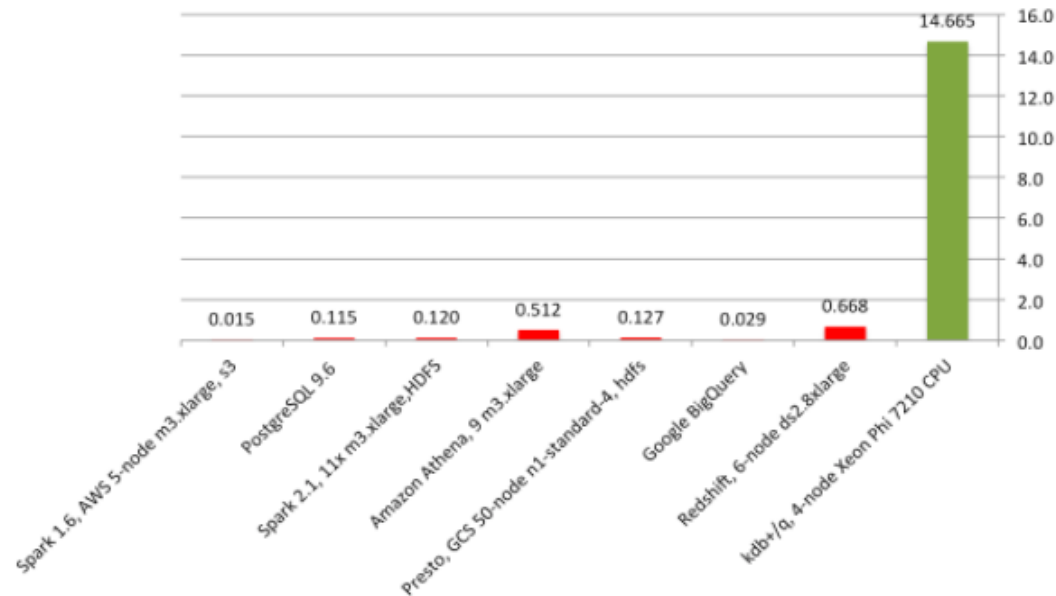Improved flow analysis and reduction of false positives

We are a green technology –

We are impacting the climate with green sustainable technology. Why run large clusters of hardware and damage the environment? You can make the move to a more energy effective software platform today!!

Streaming	Process and analyze 30 million events per second per core
Scan	Search in-memory tables at 4 Billion records per second per core
Batch	Bulk Ingest data at 10 million records per second per core
Store	Accumulate 10 trillion data points (3 PB) of NYSE data
Usage	Trusted by the world's top investment banks
Volume	Daily volumes of 10 TB's of streaming data per day
Scale	From Raspberry Pi, edge devices, to 20,000 cores on AWS Cloud

Energy Efficiency:
Number of Queries Per Watt (more is better)

Platform	Queries per Watt
Spark 1.6, AWS 5-node m3.xlarge, s3	0.015
PostgreSQL 9.6	0.115
Spark 2.1, 11x m3.xlarge HDFS	0.120
Amazon Athena, 9 m3.xlarge	0.512
Presto, GCS 50-node n1-standard-4, hdfs	0.127
Google BigQuery	0.029
Redshift, 6-node ds2.8xlarge	0.668
kdb+/q, 4-node Xeon Phi 7210 CPU	14.665

Thank you

kx

Robert Hill

Head of Space

+447929278501

rhill1@kx.com

Gary Connolly

Business Development

+447885969260

gconnolly@kx.com