# IPRO-313 Fall 2007
# Ultra-High-Speed Market Data System


Sponsor: Townsend Analytics

Manager: Carl Ververs

Illinois Institute of Technology

Advisors: Wai Gen Yee

Ben Van Vliet

| Students: | Design Team: | Hardware Team: | Software Team: |
|---|---|---|---|
| | Philip Pannenko | Michael Lenzen | Jongyon Kim |
| | Devaraj Ramsamy | Yunseok Song | Young Cho |
| | Kenneth Buddell | Noh Hyup Kwak | Usman Jafarey |
| | | Jong Min Lim | Jesus Allan C Tugade |
| | | | Jong Su Yoon |

# INDEX

# OBJECTIVES

IPRO 313's objective is to develop a data ticker plant for our sponsor Townsend Analytics which needs to meet or exceed certain performance requirements. The data ticker plant has to have a sustained optimal throughput of three million price quotes per second and minimize latency while maintain specific constraints. The ticker plant aggregates streaming data for numerous global financial markets and disseminates the data to thousands of users in real time. The data is used in Townsend Analytics' RealTick® Execution Management System (EMS), its flagship institutional product for the financial services industry. Thus, timely and accurate data delivery is a critical component to Townsend's product and competitive position. Through research and development, the group will have a concept of design, prototype development and benchmark testing. Additionally, the groundwork will be laid for future development of the ticker plant and additional trading-platform components.

Our group prioritized objectives as follows:

- Research low latency discussions and reports
    - Fully understand what a ticker plant is and does
    - Learn about new methods for ticker plants and explore their advantages and disadvantages
    - Understand market use
- Explore competitors' solutions
    - Know what is currently on the market
    - Better understand implementation of ticker plant architecture
    - Understand what works and what does not
- Develop a functioning ticker plant system
    - Analyze ticker protocols used
    - Design ticker plant architecture
    - Code a working small system
- Determine hardware requirements
    - Test off-the-shelf hardware for system
    - Design custom hardware configuration
    - Compare each solution
- Benchmarks & Prototype
    - Integrate hardware and software designs
    - Prepare benchmarks
    - Document technical user manual

# BACKGROUND

From the creation of the first stock ticker on the floor of the New York Stock Exchange (NYSE) in 1867, one thing was set in stone: speed meant money. Now if we fast forward to the 21st century we find out the same is true; worse, after the implementation of the computers and automatic trading algorithms, so much information is on the nets now that, it is all becoming unmanageable! [1]

Figure 1: Predicted Increase in Messages Per Day

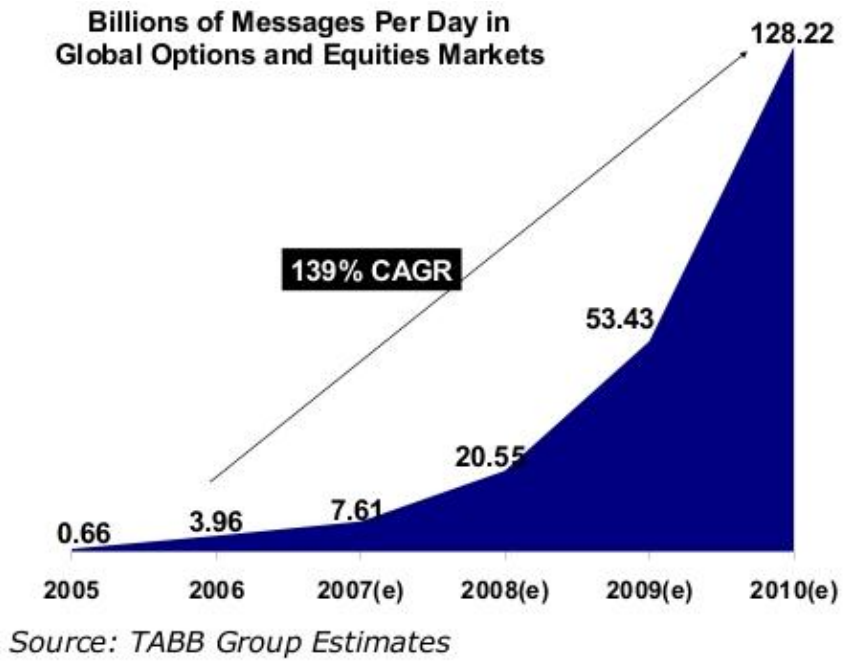


(Source: Trading at Light Speed)

Townsend Analytics (TAL), a direct-access trading-system vendor, provides connectivity to a multitude of electronic-communications networks and stock exchanges:

> Servicing the global capital markets for over 20 years, Townsend Analytics Trading Services offers world class trading and trade order management solutions to the institutional- portfolio management and broker dealer- marketplace. In today's competitive institutional trading environment, portfolio managers and traders are under increasing pressure to access a wider variety of liquidity sources, employ more aggressive yet cost efficient trading strategies and achieve best execution faster than ever. Achieving this demands the right service bureau partner, powerful order management/reporting tools and direct access to a variety of electronic markets across multiple asset classes. [2]

TAL has already developed a means for FH which has captured both the Unites States market as well as the European market:

Powered by RealTick®, Townsend Analytics' flagship institutional product, RealTick EMS (Execution Management System) is the institutional financial service industry's leading multi-asset, multi-broker, multi-routing and multi-regional market data, analytics and direct market access (DMA) trading platform. [2]

Currently there are a few firms that address this same problem; Exegy, Wombat & RMDS are examples of such. These companies however lack an effective system that will aggregate the amount of expected data in the next few years. Exegy currently claims to be the leading provider of aggregation of data at two million messages per second. They achieve this by using specific hardware and software configurations.

If TAL hopes to increase data flow by doubling processing power to their current solution, though bottlenecks will simply appear in different areas in their design. This approach leads to diminishing returns and is only a short-term solution to the problem.

IPRO 313 will research current approaches by TAL and their competitors and evaluate the cost benefits of each approach. Then, according to this research, the team will develop a prototype which will demonstrate:

Figure 1:Example of Real Time Stream Processing

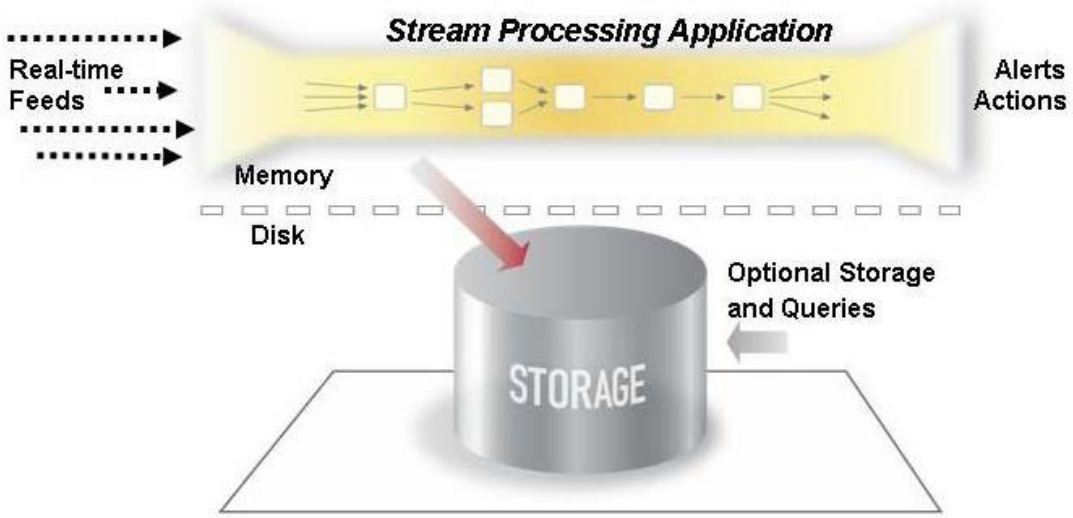


(Source: The 8 Requirements of Real Time Stream Processing)

- Stream Processing Application:
  - Translation of ticker data
  - Archiving complete ticker information
- Storage & Queries:

    - Normalization of incomplete data
    - Accessing cache values
    - Distributing tickers according to subscriptions
- Reduce Bottlenecks:
    - Optimize architecture to streamline data delivery performance
    - Reduce latency

By the end of this IPRO, the team will provide a proof of concept, complete with benchmarks, and a technical and user manual. Also, annotated research papers, competitors' system reviews & comparison charts will be included as extra appendices.

# METHODOLOGY

The IPRO group is broken down into three teams: Design, Hardware and Software. Each team is responsible for researching and developing solutions within its specific area, but also informs and collaborates with the other teams of their findings. Each team will draft its own reports, schedule and presentations. Additionally, each team will report to the group weekly, so that the information is presented in a timely and consistent manner. With a weekly schedule if presentations and reports, the flow of information is communicated to everyone on the project and the oversight of work is being maintained. Documents generated from each group are to support the IPRO deliverables, so that they can be streamlined into a comprehensive deliverable.

The Design team is responsible for all IPRO deliverables and deadlines and for managing the work flow of the entire project. It is also responsible for communicating weekly with the sponsor on the progress of work done for the past week and for upcoming work. The design team, aside from taking a managerial role, will conduct research on the financial industry. Research topics include competitors, regulations, protocols, and overall market conditions now and expected future market conditions.

The Hardware team will determine whether or not standard off-the-shelf hardware will meet the requirements both now and in the future. We will justify our decision with research into the growth in hardware capabilities and analysis of different current hardware configurations through experimentation. If the case is that standard servers will not suffice we will come up with a proposal for a solution that will work and justify why we think that it will work the best. The hardware design encompasses more than just the design of each machine, but also the overall design of the system. We will also analyze competing products, determining what they are using and why.
Work throughout the semester will include:

- Creation of tests and benchmarks to simulate demand on the system
- Test simultaneous network I/O by testing the maximum throughput while varying message sizes and ratio's of input to output
- Testing of data access times and data updates
- Coordinate work with software team in development of the macro-design

The software team will focus on Data Streaming and Last Value Cache in the processing system over this semester with methods via research into relevant paper from academic side and benchmark other systems from industrial side. We will propose reliable solutions and implement programs with statistical data. Major deliverables for the software groups will be:

- Design (Ticker Plant System)
- Codes for Normalizing/Caching Components
- Data Generator
- Benchmarks / Test Results
- Documentation
- Technical Manual
- User Manual

- Actual executable system to show the inputs and outputs ( integrating into the H/W configuration)

# BUDGET

| **Name** | **Description** | **Price** |
|---|---|---|
| **IPRO Items** | | |
| Paper | Brochure, Abstract, Summary Sheets | $20 |
| Poster | Poster board/printing | $30 |
| Printing | Color printing | $40 |
| | | |
| **TAL Meetings** | Lunches with sponsor | $100 |
| | | |
| **Hardware** | | |
| | TBD | TBD |
| | | |
| | Total | $190 |

# SCHEDULE OF TASKS & MILESTONE EVENTS

| **Task Name** | **Time** | **Start** | **Finish** | **Resources** |
|---|---|---|---|---|
| **Research & Developing** | | | | |
| Last Value Cache | 2 Weeks | 1-Oct | 13-Oct | Software |
| Stream Processor | 2 Weeks | 1-Oct | 13-Oct | Software |
| Competitors Documentation | 10 Days | 16-Sept | 26-Sept | Hardware |
| Annotating White Papers | 1 Week | 6-Oct | 13-Oct | Design |
| Market Terms/Value | 2 Week | 1-Oct | 13-Oct | Design |
| **IPRO Midterm Report** | | | | |
| Midterm Report | 1 Week | 10-Oct | 26-Oct | Design |
| Ethics Paper | 1 Week | 10-Oct | 17-Oct | Design |
| Oral Report | 1 Week | 10-Oct | 17-Oct | Design |
| **Test** | | | | |
| Determining Test Case Format | 1 Week | 8-Oct | 15-Oct | Design |
| Module Performance | 2 Weeks | 15-Oct | 27-Oct | Software |
| Modify Code | 2 Weeks | 15-Oct | 27-Oct | Software |
| Simultaneous Network I/O | 2 Weeks | 27-Sept | 9-Oct | Hardware |
| Data Access Times & Updates | 2 Weeks | 27-Sept | 9-Oct | Hardware |
| **Analysis & Improvements** | | | | |
| Designing Guidelines For Next Phase | 1 Week | 27-Oct | 5-Nov | Software |
| Proposing of New Features | 1 Week | 27-Oct | 5-Nov | Design |

| | | | | |
|---|---|---|---|---|
| 64 Bits vs. 32 Bits | 2 Week | 11-Oct | 29-Oct | Hardware |
| Multiple Network Cards | 2 Week | 11-Oct | 29-Oct | Hardware |
| **Integration** | | | | |
| Components Tested As Whole | 3 Weeks | 29-Oct | 17-Nov | ALL |
| Error Debugging & Correction | 3 Weeks | 29-Oct | 17-Nov | ALL |
| **IPRO Deliverables** | | | | |
| Poster | 2 Weeks | 17-Nov | 30-Nov | Design |
| Final Report | 2 Weeks | 17-Nov | 30-Nov | Design |
| Presentation | 2 Weeks | 17-Nov | 30-Nov | ALL |
| Technical Write Up | 2 Weeks | 17-Nov | 30-Nov | Hardware |
| User Manual | 2 Weeks | 17-Nov | 30-Nov | Software |
| Industry Report | 2 Weeks | 17-Nov | 30-Nov | Design |
| Benchmarks | 2 Weeks | 17-Nov | 30-Nov | Hardware |

# INDIVIDUAL TEAM MEMBER ASSIGNMENTS

| Group Members | Major | Skill & Strengths | Experience | Roles |
|---|---|---|---|---|
| Design Team | | | | |
| Philip Pannenko | Computer Science | Management & Communication | CBOE employment | IPRO Team Leader; Agenda Maker |
| Devaraj Ramsamy | Business | Project Controls | Bechtel SAIC LLC, DOE Yucca Mountain Project | iGROUPS & Deliverable Management |
| Kenneth Buddell | Business | | Previous IPRO experience | Timesheet Summarizer<br><br>Industry Research |
| Software Team | | | | |
| Jongyon Kim | Business | | | IPRO Sub Team Leader; Minute Taker; Timesheet Summarizer |
| Young Cho | Applied Math | | | |
| Usman Jafarey | Computer Science | | | |
| Jesus Allan C Tugade | Computer Science | | | |
| Jong su Toon | Computer Science | C, C++, Java, Oracle, PHP | Online Game Company | Software Coding |
| Hardware Team | | | | |
| Michael Lenzen | Applied Math | Java, JavaScript, PHP, Perl | Web start-up company | IPRO Sub Team Leader; Timesheet Summarizer |
| Yunseok Song | Electrical Engineering | C, Matlab, Circuit Analysis | | Research |
| Noh Hyup Kwak | Electrical Engineering | Java, Verilog, Circuit Analysis | CLSI Design<br><br>Automation Lab | Research |

| | | | | |
|---|---|---|---|---|
| Jong Min Lim | Electrical Engineering | Statistical Analysis, Circuit Design | Power Analysis ComEd | Network Analyst |

# RESULTS

## HARDWARE RESULTS

**1. Research Results**

(a) Current Technology

1) FPGA (Field-Programmable Gate Array)

FPGA is a semiconductor device containing logic blocks. Logic blocks are programmed to perform a function. FPGAs are increasingly used in conventional high performance computing applications, which includes the ticker plant we want to develop.

- Programmed with hardware description language (HDL), such as Verilog or VHDL

- Ability to re-program in the field to fix bugs

- Lower non-recurring engineering costs

- More cost effective than ASIC design

At the entry level, external USB-connected FPGA boards start at around $1000, while a top of the range PCI-X based FPGA board might cost $10,000. Obviously in both cases we would still have to factor in the cost of a relatively lowly specified computer to which the FPGA board would be connected. Programmable toolkits for FPGAs usually start in the $40-50,000 range.

2) ASIC (Applied-Specific Integrated Circuit)

ASIC is an integrated circuit customized for a particular use, rather than intended for general-purpose use.

- Programmed with hardware description language (HDL), such as Verilog or VHDL

- Not cost effective

- Non-recurring engineering cost can run into hundreds of thousands of dollars

3) Infiniband

- Switched fabric communications link

- A point-to-point bidirectional serial link

- Signaling rate is 2.5 (Gbit/s) in each direction per connection

- Use 8B/10B encoding - every 10 bits sent carry 8bits of data

- Latency ranges from 1-2.5 microseconds

4) Parallel Processing

Parallel processing makes a program run faster because there are more engines (CPUs) running it. In practice, it is often difficult to divide a program in such a way that separate CPUs can execute different portions without interfering with each other. With single-CPU computers, it is possible to perform parallel processing by connecting the computers in a network. However, this type of parallel processing requires very sophisticated software called distributed processing software.

(b) Competitor benchmarks

1) EXEGY

- 2 Million exchange messages per second

- Less than 100 microseconds of latency

The main key of EXEGY's system is extreme parallelism. They extended computer technology with reconfigurable hardware logic. By processing in parallel, one instruction per cycle can be performed in instructionless manner. EXEGY uses FPGAs to maximize this performance.

2) REUTERS (RMDS)

RMDS is short for Reuters Market Data System. The main two components of their system are Source Distributor and P2PS.

Source Distributor

- implements source specific features

- optionally supports field and update filtering

- simultaneously support independent source applications

P2PS (Point-to-Point Server)

- combines data cache

- provides point-to-point access to all the information available on the Market Data Hub

- reconnects easily and quickly

- automatically converts Market Feed data to RDM/RDF

RDM: Reuter Data Module, RDF: Reuter Data Feed

**Performance of RMDS**

Highest Source Distributor throughput to date on a single 4-socket or 2-socket server

> 2,800,000 updates per second

Highest Point-to-Point Server throughput to date on a single 4-socket or 2-socket server

> 2,200,000 updates per second through a single Point-to-Point Server machine

**2. Hardware Results**

Limitations of Off-the-Shelf Hardware

Faster data feeds lead to more transactions to capitalize off of small market changes which in turn create more of a demand for even faster data. The limitations of off the shelf computers are reached by the compounding of increases in data volume and required speed. At some point the limit of a single processor will be reached and it is not as simple as adding another processor. The fact that multiple processors would not scale linearly is negligible compared to the increased overhead generated.

FPGA

Advantages

- Easy to connect in parallel (Flexible Structure)

- Reconfigurable Hardware

- Cost Effective (Lower non-recurring engineering costs)

Disadvantages

- Exists unnecessary components (gates) due to reconfigurable

- Longer delay time than ASIC

ASIC

Advantages

- There are no unnecessary components

- which can make faster than FPGA

Disadvantages

- It still needs FPGA to verify its function

We need further research and discussions to decide whether we will use FPGA or ASIC to build system. However, since we need to use FPGA anyway, we will focus on FPGA first.

**3. Sources**

White Papers

1) Exegy Ticker Plant 2.0

2) STAC Report - Exegy Ticker Plant

3) STAC Report - RMDS

Websites

1) http://en.wikipedia.org/wiki/Fpga

2) http://en.wikipedia.org/wiki/Asic

3) https://customers.reuters.com/developer/rmdsandtools

# SOURCES

[1] Johnson, Jeromee. Trading At Light Speed: Analyzing Low Latency Data Market Data Infrastructure. New York City: TABB Group, 2007. 1-22. 30 Aug. 2007 <www.tabbgroup.com>.

[2] "Corporate Overview." Townsend Analytics. 2007. 19 Sept. 2007 <http://www.townsendanalytics.com/>.