

CloudEx: Experiments with Building a Financial Exchange in the Cloud

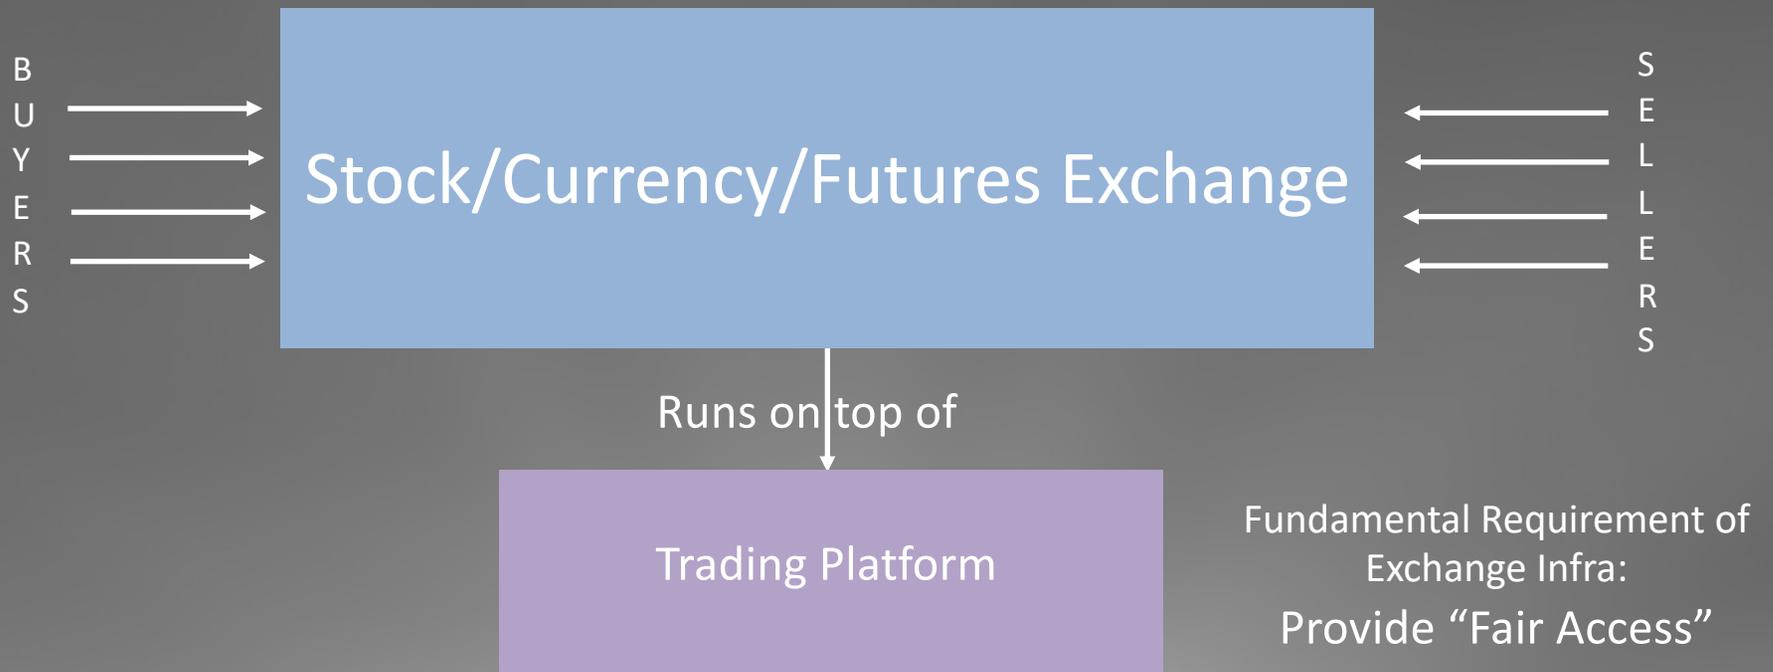
Balaji Prabhakar

Professor of EE and CS, Stanford University

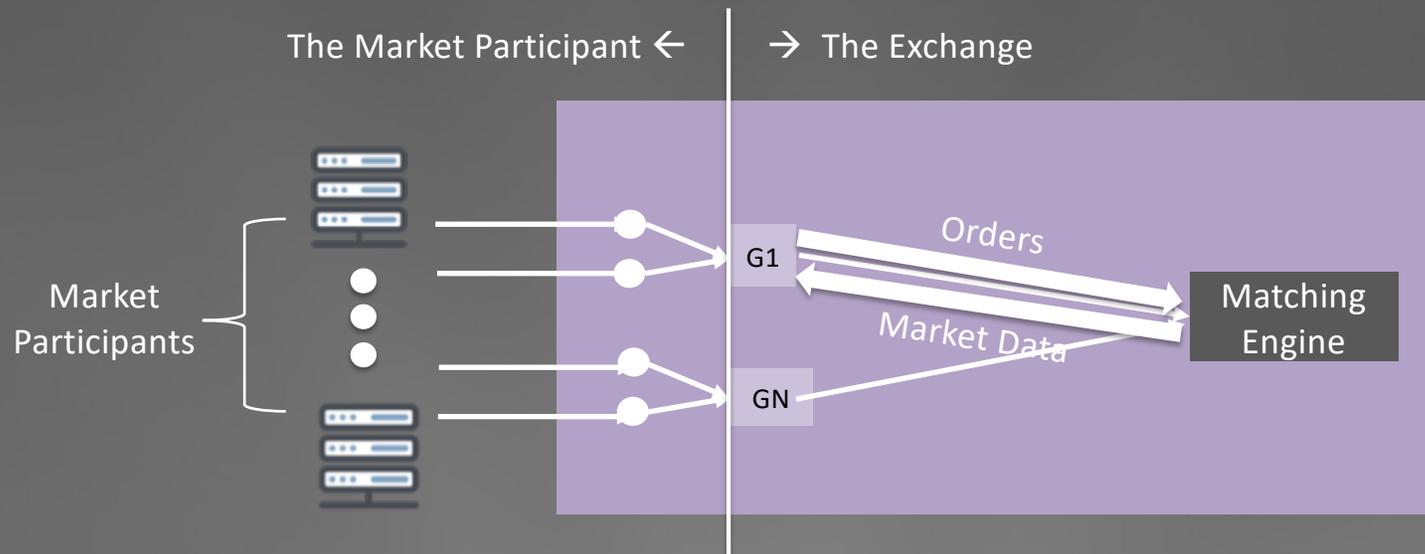
Co-Founder, Tick Tock Networks

Exchanges and Trading Platforms

Fundamental Operation Performed by Exchange:
Price Discovery!



The Exchange Loop Must Be **Fair**

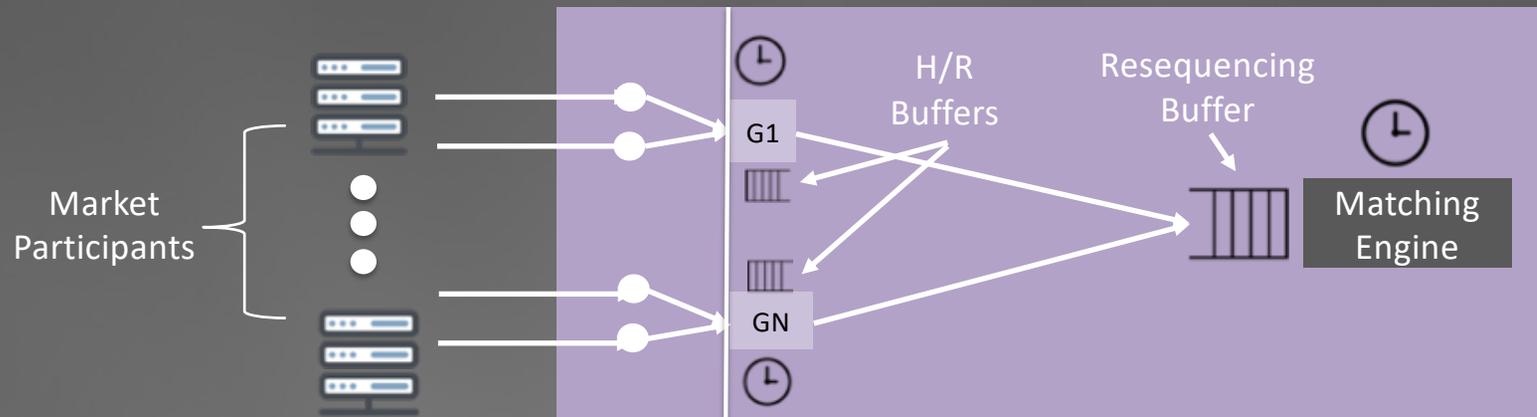


Definition of Fairness

Inbound Fairness: *Orders must be processed in a globally FIFO manner, regardless of GW through which they enter the exchange*

Outbound Fairness: *The Market Data shall reach all MPs (GWs) at exactly the same time*

Providing Fairness using Accurate Clocks



Inbound Fairness (recall: GW and ME clocks are synchronized)

Timestamp orders when they arrive at the gateways to establish order of arrival

Resequence orders according to their gateway timestamps before executing them at the ME

Outbound Fairness

Timestamp market data at the ME

Hold-and-release market data at the gateways using the H/R buffers to ensure simultaneous delivery

→ See STAC talk from Oct 2019: <https://stacresearch.com/STAC-Summit-17-Oct-2019-tick-tock>

"All the News That's Fit to Print"

The New York Times

VOL. CLXVII... No. 58,009 NEW YORK, SATURDAY, JUNE 30, 2018 \$3.00

KREMLIN'S LINKS TO BREXIT PUSH GET A NEW LOOK

COURTING BANKROLLERS

Parallels to U.S. Inquiry as Investigators Focus on 2016 Meeting

By DAVID H. BRINKSTEIN and MATTHEW ROSENBERG

LONDON — Aron Banko, a British financier who bankrolled the campaign for Britain to leave the European Union, has long bragged about his "close relationship" with the Russian ambassador.

Since she wondered about Mr. Banko's Russian ties and her own campaign, Mrs. Clinton, a U.S. senator, after the British withdrawal agreement, Mr. Banko said he had not met the Russian ambassador.

Mr. Banko's Russian ties and her own campaign, Mrs. Clinton, a U.S. senator, after the British withdrawal agreement, Mr. Banko said he had not met the Russian ambassador.

Mr. Banko's Russian ties and her own campaign, Mrs. Clinton, a U.S. senator, after the British withdrawal agreement, Mr. Banko said he had not met the Russian ambassador.



Battling Tanks, Troops and Twisted Tactics
American soldiers in Eastern Europe holding training exercises to combat Russia's military might and manipulation tactics. Page A7.

Once Deemed Too Radical for Mexico, He Now May Be President

By ADAM ARMED

CANADIANA, Mexico — Andrés Manuel López Obrador, a three-time candidate for president who once reflected a nationwide anti- and anti-establishment mood, is now being seen as a deep connection with the president-elect.

Mr. López Obrador, a three-time candidate for president who once reflected a nationwide anti- and anti-establishment mood, is now being seen as a deep connection with the president-elect.

Mr. López Obrador, a three-time candidate for president who once reflected a nationwide anti- and anti-establishment mood, is now being seen as a deep connection with the president-elect.

his left-leaning platform and connect the dots of Mexico's most important political figures.

Mr. López Obrador, a three-time candidate for president who once reflected a nationwide anti- and anti-establishment mood, is now being seen as a deep connection with the president-elect.

Mr. López Obrador, a three-time candidate for president who once reflected a nationwide anti- and anti-establishment mood, is now being seen as a deep connection with the president-elect.

Mr. López Obrador, a three-time candidate for president who once reflected a nationwide anti- and anti-establishment mood, is now being seen as a deep connection with the president-elect.

Mr. López Obrador, a three-time candidate for president who once reflected a nationwide anti- and anti-establishment mood, is now being seen as a deep connection with the president-elect.

Mr. López Obrador, a three-time candidate for president who once reflected a nationwide anti- and anti-establishment mood, is now being seen as a deep connection with the president-elect.

Mr. López Obrador, a three-time candidate for president who once reflected a nationwide anti- and anti-establishment mood, is now being seen as a deep connection with the president-elect.

Mr. López Obrador, a three-time candidate for president who once reflected a nationwide anti- and anti-establishment mood, is now being seen as a deep connection with the president-elect.

G.M. TELLS TRUMP TARIFF PLAN PUTS U.S. JOBS AT RISK

CAUGHT IN TRADE WAR

Warns Auto Prices Could Rise by Thousands, Sapping Demand

By TIFFANY HUI

General Motors warned Friday that if President Trump pushed ahead with another round of tariffs, the more costly hardware, leading to "less investment, fewer jobs and lower wages" for its employees.

General Motors warned Friday that if President Trump pushed ahead with another round of tariffs, the more costly hardware, leading to "less investment, fewer jobs and lower wages" for its employees.

General Motors warned Friday that if President Trump pushed ahead with another round of tariffs, the more costly hardware, leading to "less investment, fewer jobs and lower wages" for its employees.

General Motors warned Friday that if President Trump pushed ahead with another round of tariffs, the more costly hardware, leading to "less investment, fewer jobs and lower wages" for its employees.

Whittling a New York Minute To 100 Billionths of a Second

By JOHN MARKOFF

SAN FRANCISCO — Computer scientists at Stanford University and Google have created technology that can track time down to 100 billionths of a second. It could be just what Wall Street is looking for.

second. Ultimately, this is about money. With stock trading now dominated by computers that make buying and selling decisions and execute them with blazing speed, keeping that order also means protecting profits. So-called high-frequency trading firms place trades in a fraction of a second, sometimes in a bet that they can move faster than bigger competitors.

System engineers at Nasdaq, the New York-based stock exchange, recently began testing an algorithm and software that they hope can synchronize a giant network of computers with that nano-second precision. They say they have built a prototype, and are in the process of deploying a bigger version.

The pressure to manage these high-speed trades grows when the stock market becomes more volatile, as it has been in recent months, in part to prevent the fastest traders from taking unfair advantage of slower firms. High-frequency traders typically ac-

For an exchange like Nasdaq, such refinement is essential to accurately order the millions of stock trades that are placed on their computer systems every

Continued on Page A17

Firing of Comey Left Rosenstein Feeling 'Shaken' and Exploited

WASHINGTON — In the days after the F.B.I. director James B. Comey was fired last year, the deputy attorney general, Rod J. Rosenstein, reportedly expressed anger about how the firing was carried out. He said he felt "shaken" and "exploited" by the move.

Whittling a New York Minute To 100 Billionths of a Second

SAN FRANCISCO — Computer scientists at Stanford University and Google have created technology that can track time down to 100 billionths of a second. It could be just what Wall Street is looking for.

In Gunman's Rage, Something Eerily Familiar for Newsrooms

Years ago, the Kaitia was out to dinner or a nice restaurant with his wife, celebrating their second anniversary. Mr. Kaitia, then the editor of the New York Times, was in the kitchen that led to the center of the island, a friend of the district attorney showed up and asked the Kaitia to go outside.

F.U. Reaches Migration Deal

DOUGHTY making an agreement, the agreement seemed to be enough to save Chancellor Angela Merkel's U.S. ally from appearing lonely.

Nationalist Aid

A ruling in a Kentucky case may affect other states' efforts to restrict abortion to work to health care.

China Curtails Lending Spree

A tight credit policy aimed at curbing excessive lending by banks in China grows cautious.

9/11 Safety Chief Laid to Rest

Chief Ronald B. Spidolow, who kept rescue workers safe at ground zero of World Trade Center at 93.

Riding an Outdoor Wave

The state's program to fund the use for the city's elite high schools could spend in elite middle schools.

The N.E.A.'s Decline

Led by James, James might join a new union in Cleveland, or he might stay in Cleveland.

Tough Math in School Plan

The state's program to fund the use for the city's elite high schools could spend in elite middle schools.

China Curtails Lending Spree

A tight credit policy aimed at curbing excessive lending by banks in China grows cautious.

The N.E.A.'s Decline

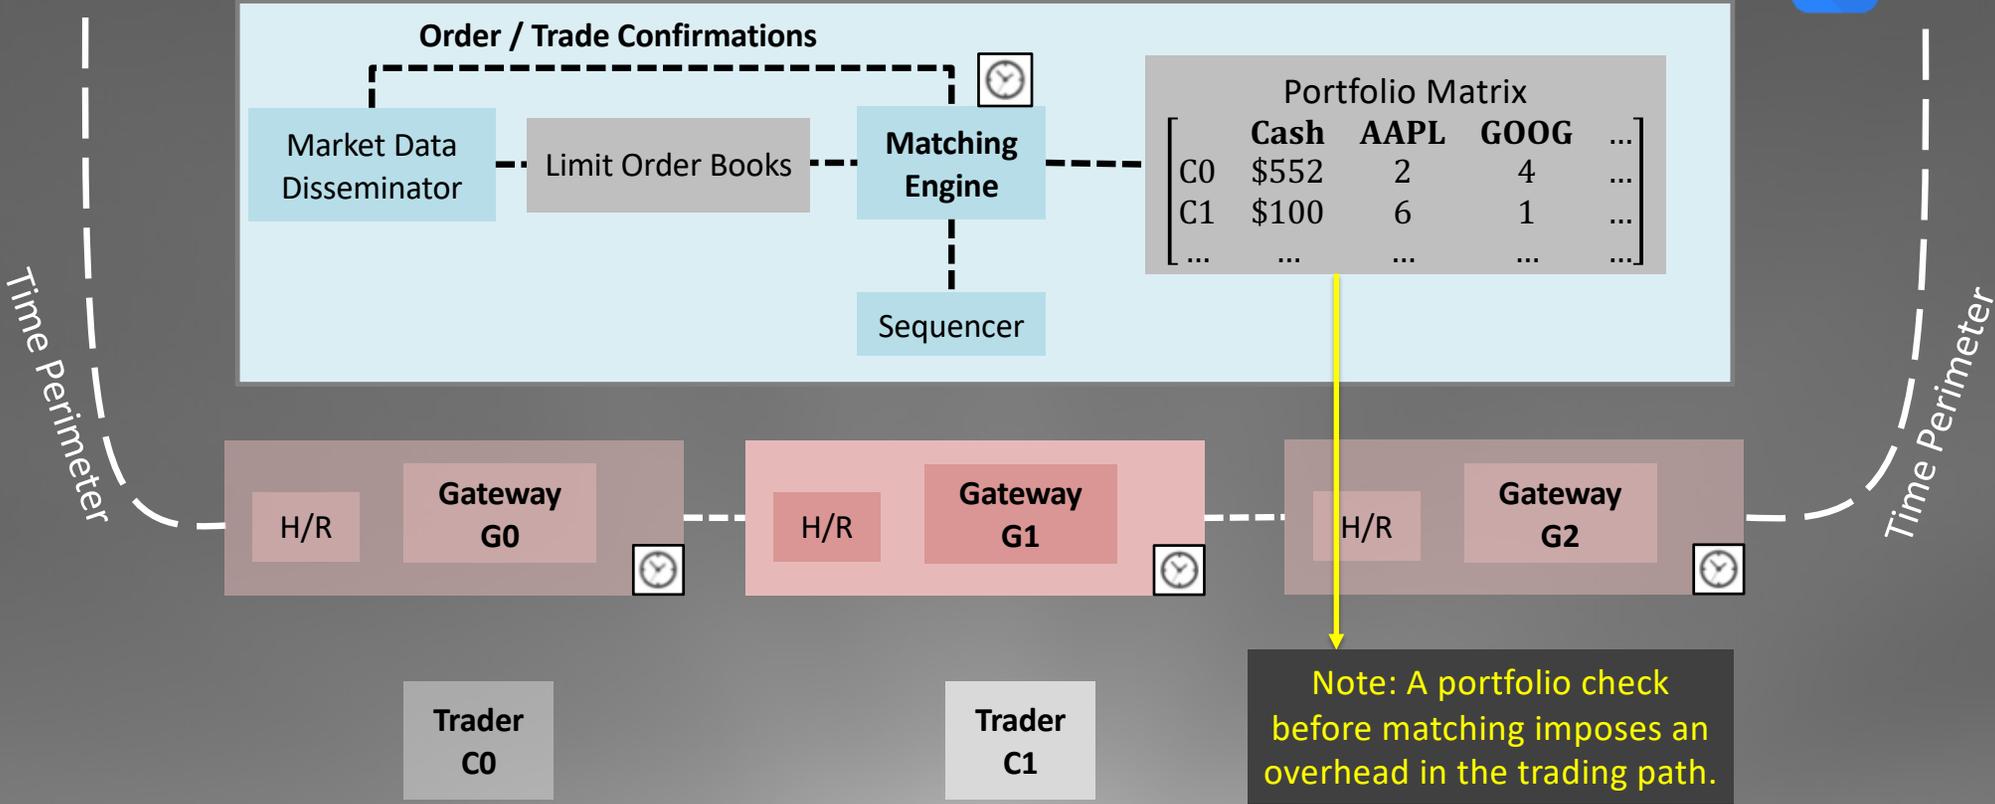
Led by James, James might join a new union in Cleveland, or he might stay in Cleveland.

Motivation for CloudEx

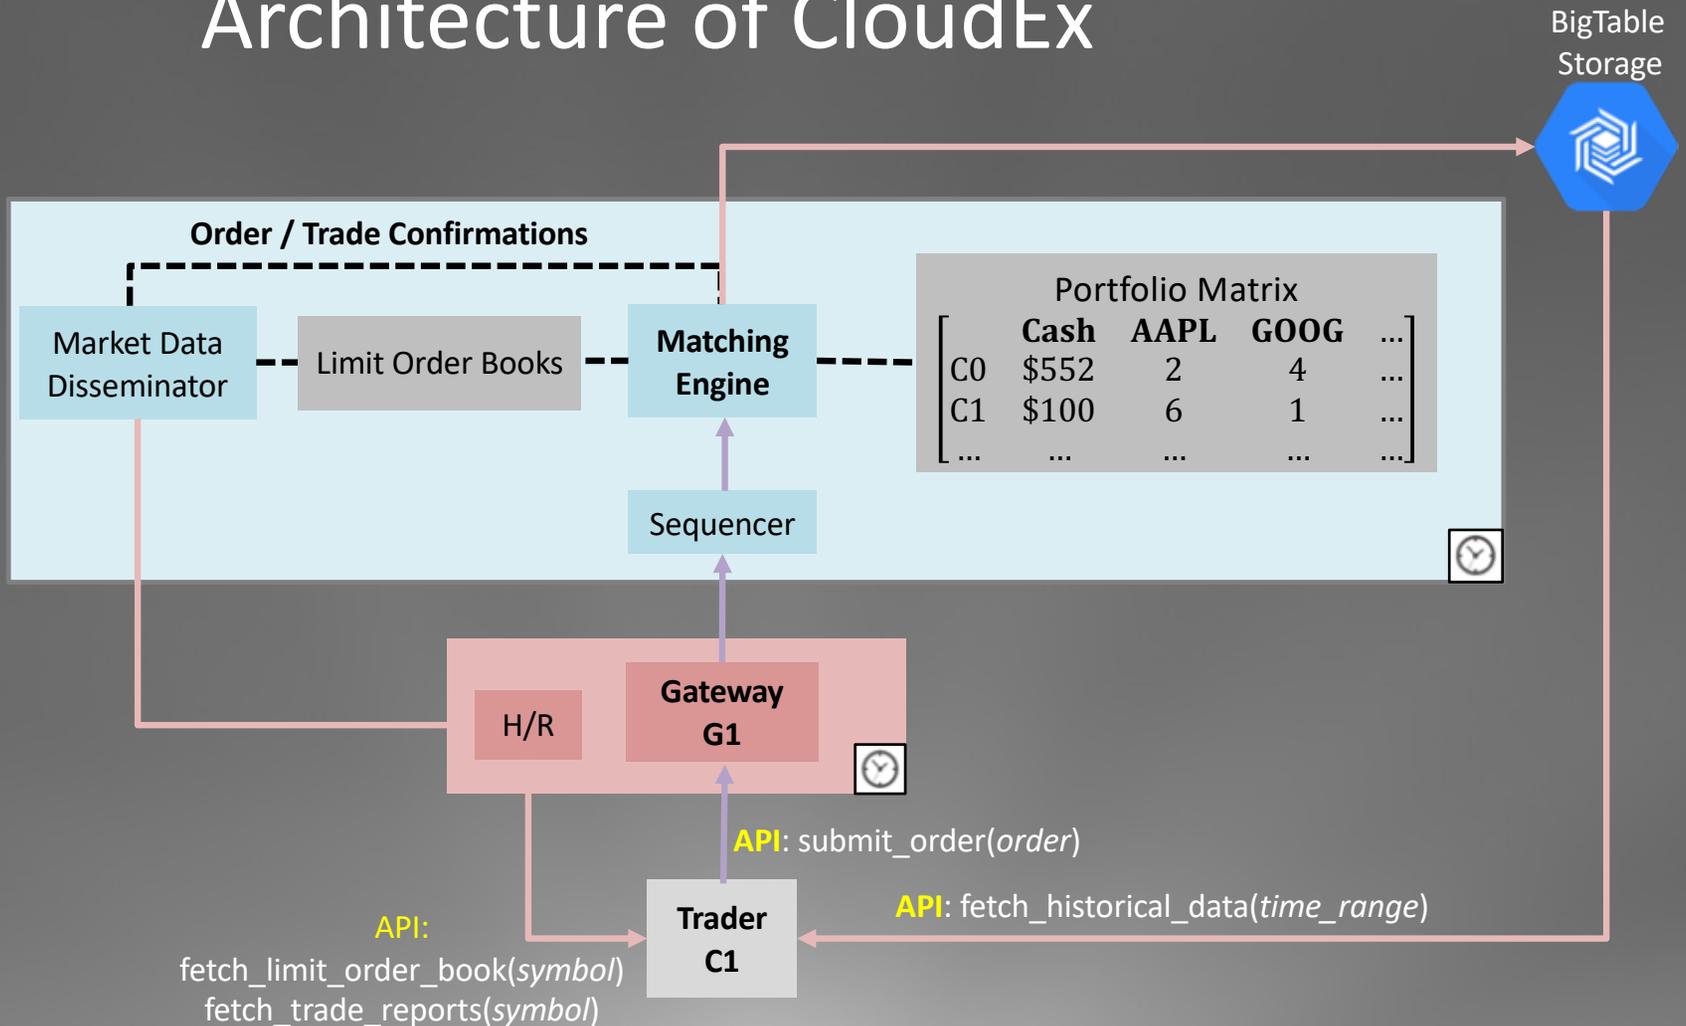
- Study feasibility of building fair and highly performant financial exchanges in the public cloud
- Research tool
 - Fully configurable end-to-end exchange; can experiment with infra changes, matching engine algos, etc
- Teaching tool
 - Prototype system for learning about (i) infrastructure, (ii) trading algorithms, and (iii) matching algorithms
 - Used in a course at Stanford this Autumn; student teams will compete as HF traders
- Joint work with
 - Vinay Sriram, Jinkun Geng, Ahmad Ghalayini, Vig Sachidananda, Yilong Geng, Mendel Rosenblum and Anirudh Sivaraman

Architecture of CloudEx

BigTable
Storage



Architecture of CloudEx



Order Messages

Currently restricted to Market and Limit Orders

Appended by Gateway

Format

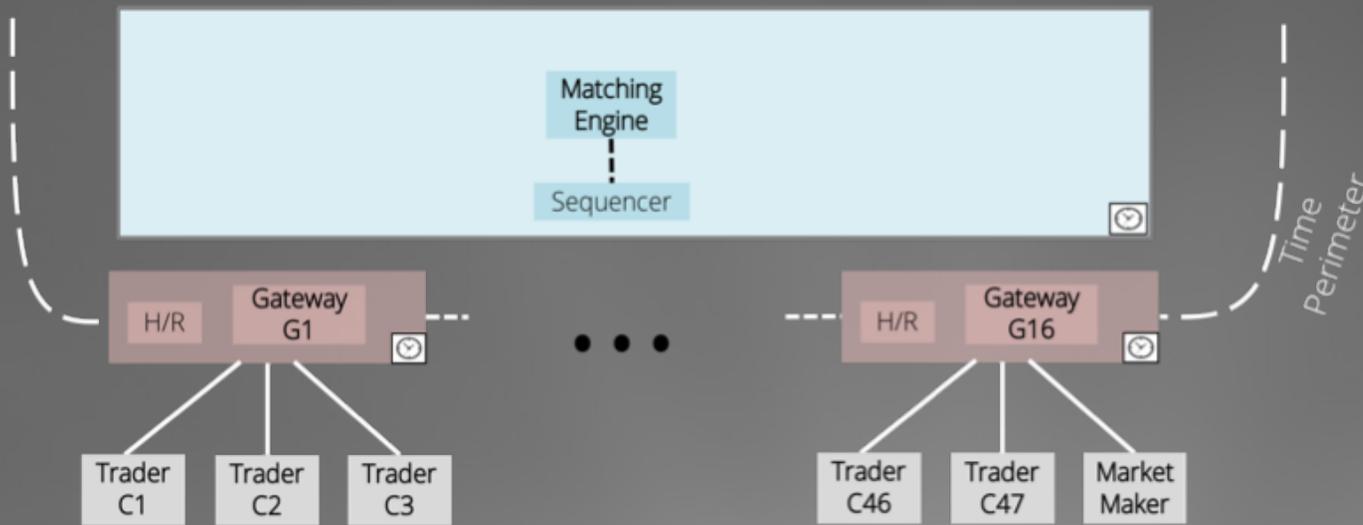
Symbol	Action	Shares	Type	Bid/Ask Price	Timestamp	Client ID	Order ID
--------	--------	--------	------	---------------	-----------	-----------	----------

Examples

AAPL	Buy	100	Limit	\$210	1569290045000	C1	SAD651GH
GOOG	Buy	50	Market	\$1200	1569290048121	A2	3AS5SDF2
MSFT	Sell	325	Limit	\$140	1569290053007	B2	A22S134H

CloudEx Testbed Setup and Numbers

Testbed Setup



- 1 Matching Engine
- 48 Traders, 3 per Gateway
- 10 symbols to trade
- Each trader starts with 5M cash and 10k shares per symbol

Traders

1 Market Maker

Trader
C48

- Has a large amount of resources that allows it to transact in high volume and drive the price of a symbol towards some predetermined price

43 Random Traders

Trader
C21

- Models traders that perform trades without reacting to instantaneous market data

4 Algorithmic Traders

Trader
C1

- Uses a **mean reversion strategy** to profit on volatility in certain symbols
- Also support other strategies such as Pairs Trading and Momentum Trading

CloudEx Testbed Specs

- VMs rented in Google Cloud; heterogeneous capabilities, as follows:
 - Matching Engine:
 - n1-standard-64 (64 vCPUs, 240 GB memory, 32Gbps Bandwidth Maximum)
 - Market Maker:
 - n1-standard-32 (32 vCPUs, 32 GB memory, 32 Gbps Bandwidth Maximum)
 - Gateway:
 - n1-standard-8 (8 vCPUs, 30 GB memory, 16 Gbps Bandwidth Maximum)
 - Trader:
 - n1-standard-2 (2 vCPUs, 7.5 GB memory, 10 Gbps Bandwidth Maximum)

Dashboard Plots

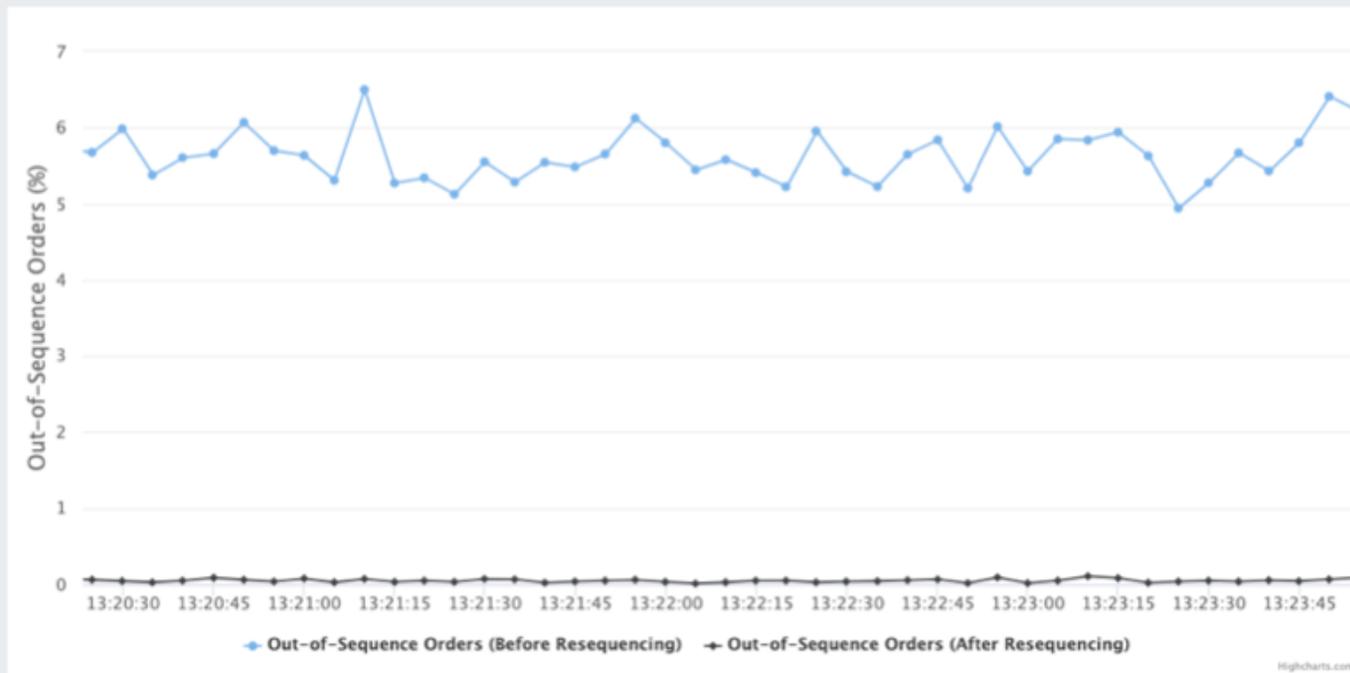
Inbound Fairness

Queue Length

System Delays

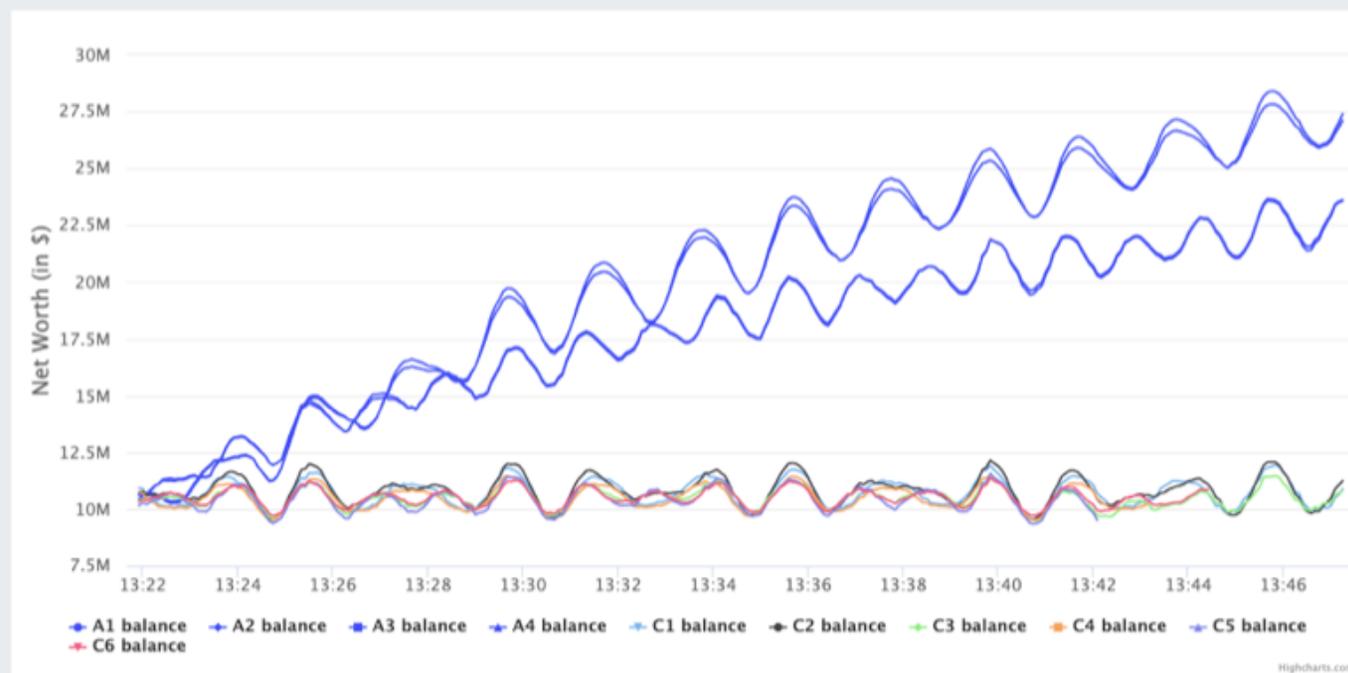
Inbound Fairness

Out-of-Sequence Orders.



Portfolio Values

Evolution of each trader's portfolio value over time.



CloudEx in Numbers

(Not a STAC benchmark)

- **Performance**

- **Throughput**

- 1 ME Shard: Throughput ~ 30K/second
- 2 ME Shards: Throughput ~ 63K/second
- ~30 μ s processing latency for each order per shard

- **Latency: Inbound (order placement, 2 Hops)**

- Trader to Gateway: 164 μ s
- Gateway to Matching Engine 156 μ s

- **Latency: Outbound (market data, 2 Hops)**

- Matching Engine to Gateway: 217 μ s
- Gateway to Trader: 199 μ s

- **Work In-Progress**

- Working with VM timestamps currently: 100s nanosecs accuracy, large stack latencies
 - NIC timestamps + using SR-IOV will lower latencies
- Replace network communication library (ZMQ) with high-performance bypass framework (e.g., DPDK)
 - Significantly decrease latency
- Currently have verification overhead (Portfolio Matrix); may move verification out of trading path
 - Significantly increase throughput