Observing news propagation

Stéphane Tyč and Imad-Eddine Srairi

Quincy Data

revised Oct 09, 2015

Abstract

Large and sudden changes in prices in broad market indices reflect the release of macro news. The detailed observation of the times of these moves can provide lower bounds for the time of transmission of information between data centres. We have studied one macro events and provide our findings. The propagation of news on land in the United States is well explained by optimised microwave networks. The propagation of news to Europe is consistent with the fastest known transatlantic fibre cables at the time.

Introduction

On September 18, 2013 the Federal Open Market Committee (FOMC) released a statement saying it would continue injecting money in the economy – *the Committee decided to continue purchasing additional agency mortgage-backed securities at a pace of $40 billion per month and longer-term Treasury securities at a pace of $45 billion per month.* This announcement was unexpected and triggered large price moves in the markets. In the wake of the large moves, questions were raised on the mechanism used for releasing information[1]. It was surprisingly difficult to obtain a simple answer. A Federal Reserve spokesman even refused to answer questions raised by Eamon Javers of CNBC who said:

> So what exactly were those Federal Reserve lockup rules? Were organisations allowed to transmit information out of the room before 2 p.m. or not? The Federal Reserve won't say? A Fed spokesman declined to answer that question from CNBC.[2]

We studied the impact of the news in several data centres and were able to conclude that the news was simultaneously released in Chicago and in New Jersey under a mechanism called *embargo*[3]. Embargo means that a piece of news can be transported before it is released but that it can only be released in situ after a very precise time. After this controversy, the release rules were changed to another mode called *lockup*. In lockup mode, the piece of news is released in a single point – the lockup – and then it is propagated using the fastest possible transport mechanisms and used to trade on various markets. The lockup release can be very helpful to study the speed of the best networks.

In this article we study another news release and try to infer the propagation latency of the best networks. We propose a simple formula for estimating the latency of optimal microwave networks on land and find it to be consistent with observed data.

[1] http://www.nanex.net/aqck2/4436.html

[2] http://www.cnbc.com/id/101062081

[3] The interested reader may refer to:

- http://www.quincy-data.com/fed-robbery-revisited/
- http://www.quincy-data.com/fed-robbery-new-evidence/
- http://www.quincy-data.com/great-fed-robbery-2013/

Data sources and news impact measurements

Identifying a news trade

The impact of news can be seen in the price moves they trigger. Unfortunately, exchanges do not ask participants to log their thought process when submitting an order. We have to resort to finding a convincing signature of news trading. We have not defined any scientific measure to unambiguously detect news trading. We have studied only a very limited set of events which can be traced to a single announcement and are as unambiguous as possible.

There are two typical signatures of trading motivated by a news release. The first is usually a sudden increase in volume, accompanied by a rapid change in price. However, this alone is not sufficient to identify the impact of a macro news. A macro news will trigger a large volume and a rapid change in price *in more than one asset at the same time.* In the normal course of trading, two different assets may have spikes in volume that are causally related but not co occurring. For instance, the S&P 500 futures moves on the CME and the equities of the underlying index move accordingly about 4 ms later in New Jersey. When two different assets move exactly at the same time, the cause of their move must be external. In the news release studied below, the bonds traded on eSpeed move 2 μs later than the equities on the same venue. It is impossible for one of the assets to trigger the move of the other because the market data of the first move is published after the second asset has traded. This is a clear signature that both moves had an external cause.

The identification of the market data which are related to the news release is discussed in the annexes. Each market is discussed and the actual packets identified as the first signature of the news trade are identified.

Finding the exact time of arrival in a data centre

Another difficulty is that we only have access to the market data disseminated by the exchange. This does not provide the time of arrival of the news in a data centre, it only provides the time of publication inside the data centre of the impact of the news.

Let's introduce -in Table (1)- some definitions for the various times in the chain of events that we are considering.

Time	Definition
t_0	News release in a lockup
t_1	News arrival in data centre, after processing at end point
t_2	Order arrival at exchange gateway, after its submission
t_3	Trade execution by exchange matching engine
t_4	Timestamp of trade in market data published by exchange

Table 1: Definition of various times from news release to trade publication

Our goal is to infer the exact location from which the news was released, which requires to measure the difference $t_1 - t_0$. However, our observables are t_2, t_3, and, or t_4 depending on the exchange. We do not even know the actual value of the release time t_0 although, at least, we can assume that it is unique when a lock-up mechanism is used. We only have a theoretical value for t_0, which is the officially announced news release time, but the release mechanism in the lockup does not guarantee a precise synchronisation to the Coordinated Universal Time (UTC). Stated like this, the problem seems impossible to solve.

In order to achieve this goal, we will have to estimate the delay between t_1 and the smallest available times amongst t_2, t_3 and t_4, which is t_2 for EUREX, and t_3 for all other exchanges. This delay depends on the particular exchange.

Estimating the difference between order matching time found in market data and the time of arrival in a data centre There are two components in this difference $t_3 - t_1$. The first one, $t_2 - t_1$, is the time it takes for a market participant from the reception of the news to the submission of an order on the market. This time presumably very standard across data centres, it only depends on the technology deployed by the market participant. This time is also the subject of a vigorous competition between trading firms and only the smallest time is of interest in our study since the signature of the news is detected through the orders submitted by the fastest participant. This time is known in the industry as "tick to trade". A simple Google search will show numerous responses and many claims of low latency tick to trade numbers. The lowest claims are under 1 μs and most claims are under 10 μs. Since the fastest trading firm wins, it is safe to assume that this number is comprised between 0 and 5 μs. This is very small and well below the expected error bars of our study, we will take a fixed latency of 5 μ s for the purpose of this analysis.

The second component, $t_3 - t_2$, is the time between order sending and order execution on a given matching engine. This time is more difficult to know and is also more variable. We present a short discussion of this time for different markets in the appendix.

Outline of the analysis

We obtain signature times of a news release and we need to convert these times in actual distances on the map. We assume a crude model for the speed of highly optimised networks between Washington DC and the CME and Nasdaq respectively.

We verify that the model is consistent with a news released in a government lockup and made available in the Coresite data centre on K Street, Washington DC and find the likely value of t_0 in the data centre. The actual release of the news in the lockup is not a relevant number for us. Only the time when the news is provided to trading firms in DC is relevant and starts the clock for rapid news transmission.

We invert the question and find the locus of points where the news could have been released. This method could define the point of release if we had three different times of arrivals at data centres. It is done just for illustration in our case.

At this point in the reasoning we are confident that the news was released on K Street, we have the time of release with a very good approximation and we have the time of passage of the news in the Nasdaq data centre in Carteret, NJ. We can proceed to the second question which is the study of the travel time from the US to London and Frankfurt. We try to find if the news propagation time is consistent with the fastest known cable of the time. The latency of the fastest known undersea cable and the simple model for optimised microwave networks on land are consistent with the observations.

The example of the Job report release on February 6th 2015

The announcement of 257,000 new nonfarm payroll jobs was a surprise and triggered large moves in the market [4]. The European markets were open at the time and news propagation can be studied across the Atlantic as well as in the US.

It is ironic, but beside the point, to note that this number was then revised[5] and that the revised number was much more in line with expectations. Nevertheless, the markets moved after this release.

Data centre	First Available Time (ms) (measured wrt 13:30:00 UTC)
NASDAQ	$t_3 = 1,305.796061$
ICE	$t_3 = 1,308$
CME	$t_3 = 1,308.270776$
LIFFE	$t_3 = 1,336$
EUREX	$t_2 = 1,338.754921$

Table 2: News arrival time ranges on 06FEB2015 for various exchanges

New arrival times in various data centres

Transforming times in distances

The Department of Labour releases its data under lockup rule in Washington DC. Only accredited news agencies have access to this data and those news agencies sell the news in machine readable format in a data centre located on K Street in Washington DC. The first task is to translate propagation times into distances. To do this we have to make some assumptions about the transport mechanism.

On land, the fastest known transport mechanism is via microwave networks. The signal is modulated at one end of a series of towers and is repeated with a very small delay all the way to the other end. The total transport time is the sum of end point latencies which include fibre tails and network equipment and transport latencies which, to a first approximation, is proportional to distance. The best networks are very straight and follow the geodesic but they cannot be perfect and suffer from some indirection which can be described as a percentage of the total distance. The time is takes to repeat the signal is small but non zero, it is of the order of 100 to 500 nanosecond per tower and scales with the number of towers. Those two effects can be compounded in a single fudge factor ϵ and we can write the transport latency between two points as

$$T = d/c \times (1 + \epsilon) + T_{\text{endpoint}} \tag{1}$$

Carrying the news from K Street to the trading venues is one of the most competitive latency races and therefore it is highly optimised. The FCC database can be used to reconstruct the various microwave networks between K Street and the Aurora data centre where the CME is located, or Carteret where Nasdaq is located[6]. For the purpose of the discussion, we will take $\epsilon = 0.005$, half a percent, and we will take the end point latency to be $15\mu s$. In practice the end point latency depends on the various fibre lengths and equipment at the end points but we will not take this into account and assume it can be neglected.

Equation (1) can provide the distance travelled for a given time. We are now equipped to study where the news was released.

Is the data consistent with a release on K Street?

Writing Equation (1) for two data centres and subtracting both equations leads to

$$\Delta T = (1 + \epsilon)\frac{\Delta d}{c}$$

where ΔT is the difference between the news arrival times at both data centres. Notice that this relationship does no longer depend on the value of T_{endpoint}. As news arrival times t_1 are not

[4] http://www.bls.gov/news.release/archives/empsit_02062015.htm

[5] This number was revised twice, a first time on March, 6th: http://www.bls.gov/news.release/archives/empsit_03062015.htm and then on April, 3rd: http://www.bls.gov/news.release/archives/empsit_04032015.htm.

[6] For a discussion of these please see http://www.quincy-data.com/transmitting-data-via-microwave/.

observable, it proves practical to express this difference in terms of other, observable times, as the trade execution times by the matching engines t_3. Writing

$$t_3 = t_1 + \underbrace{(\ t2\ -\ t_1\ )}_{\substack{\text{tick-to-trade} \\ \text{supposed to be} \\ \text{the same everywhere}}} + \underbrace{(\ t_3\ -\ t_2\ )}_{\substack{\tau \\ \equiv \\ \text{gateway-to-exec} \\ \text{delay}}}$$

we get hence

$$\Delta T \equiv \Delta t_1 = \Delta t_3 - \Delta\tau = (1+\epsilon)\frac{\Delta d}{c}$$

Let's see what we get for CME and Nasdaq. The distance between CoreSite K Street and CME is 1004.5 km, and Nasdaq is 303.1 km far away from CoreSite K Street. The right-hand side of previous equation is worth 2.350 ms. On the other hand, using times from Table (2), we have $\Delta t_3 = 2.475$ ms. For the equality above to hold, we infer from this that $\Delta\tau$ should be equal to 125 μs. In other terms, we have to assume that the delay between order arrival at the exchange gateway and execution is 125 μs higher at CME than at Nasdaq.
This difference in speed sounds plausible / realistic. *Should we discuss this result more, considering for instance the particular time these events occurred (compare with markets opening time)?*

Now, if we attempt to determine the actual time when the news was released, our assumption above on the value of T_{endpoint} and on the tick-to-trade are not sufficient. On top of them, we have to adopt a value for $\tau = t_3 - t_1$ for either markets and then solve for t_0 in Equation (1)

$$\begin{aligned} T \equiv t_1 - t_0 &= (t_2 - \text{tick-to-trade}) - t_0 \\ &= ((t_3 - \tau) - \text{tick-to-trade}) - t_0 \\ &= d/c \times (1+\epsilon) + T_{\text{endpoint}} \end{aligned}$$

Adopting $\tau = 20$ μs for Nasdaq, we now have a good estimate of the actual release time :

$$t_0 = \underbrace{\underbrace{\underbrace{t_3(NASDAQ) - \tau(NASDAQ)}_{t_2\ :\ \text{exchange gateway arrival time}} - \text{tick-to-trade}}_{t_1\ :\ \text{arrival in Carteret, including local processing}} - T_{\text{endpoint}}}_{\text{arrival in Carteret, before local processing}} - \underbrace{d/c \times (1+\epsilon)}_{\text{"on air" time}} = 1,304,741\ \mu s$$

after 13:30:00 UTC. Of course, we would have gotten the same result if we had considered CME rather than NASDAQ, except that we would have used a value of τ consistent with our findings above i.e. $\tau(CME) = \Delta\tau + \tau(NASDAQ)$ that is $125 + 20 = 145$ μs.

Can we find the release site by observing the data?

Now, let's stick with the same example where we are detecting -though indirectly, through times t_3 instead of t_1- the arrival of the news at Nasdaq and at CME and let's assume that we do not know where the data is released. Remember that the *actual* time of release t_0 is unknown at first, we only have a theoretical, announced value (08:30:00am Eastern Standard Time corresponding to 13:30:00 UTC). Had we had an actual, accurate, reliable (i.e synchronized with markets) value for t_0, we would have converted delays $t_1 - t_0$ into distances using the speed of light in the air, and could have located the place of release as the intersection of two circles, the radii of which are the previously computed distances. Since we do not know the *actual* value of t_0, we can only rely on the difference Δt_1 between arrival times and are left with the following equation

$$\begin{aligned} d(S, \text{Aurora}) - d(S, \text{Carteret}) &= \frac{c}{1+\epsilon} \times (t_1(\text{CME}) - t_1(\text{Nasdaq})) \\ &= \frac{c}{1+\epsilon} \times (t_3(\text{CME}) - t_3(\text{Nasdaq}) - \Delta\tau) \qquad (2) \end{aligned}$$

S being the unknown location of the source of the news and d the distance between two points at the surface of the Earth. Can we still find the place where the data was disseminated based on the observed timestamps?

This question is very similar to the one addressed by the so-called *hyperbolic navigation systems* such as LORAN or DECCA: a navigator on a vessel is measuring the relative time of arrival of two micro-wave pulses which have left two separate transmitters, and attempts to infer from this her position or *fix*. If there is no difference between received times, then the vessel is at the same distance from both transmitters and could be anywhere on a straight line intersecting the transmitters axis at its middle point and perpendicular to it. If this difference is non-zero, then the locus of points exhibiting the same delay on a flat Earth is a branch of a hyperbola with the transmitters at the focal points, explaining why these systems are termed as *hyperbolic*. If there are at least three signals received from three different transmitters, then the navigator can infer her position from the intersection of these hyperbolic lines. Of course, the Earth is not flat, but we can replace hyperbolas by spherical hyperbolas, and everything works the same.

In our case, we have only one transmitter and several receivers instead of several transmitters and one receiver (aboard the vessel), but the physics and the equations are exactly the same. With three receivers, we could have located the place from which the news had been released. Since we only have two receivers and hence only one difference between received times, we can only draw the line on which it is sufficient for the source S to lie in order to verify previous equation.

Results[7] are shown on figures (1) and (2). We have used five different values for $\Delta\tau$ ranging from 75 to 150 μs. We notice that the data centre CoreSite on K Street lies exactly on the line corresponding to 125 μs.

With this calculation we have derived an adjustment parameter, $\Delta\tau$, which is the "relative slowness of CME with respect to Nasdaq". How accurate is our determination of this number? The sources of uncertainty come from our simple model of microwave network latency (Equation (1)) and on the sensitivity of the position of the pseudo-hyperbola to this parameter.

We can derive an expression for it when the source of the news S lies exactly on the geodesic line (that is: the *great circle*) joining Carteret and Aurora. In this case, we can write

$$d(S,\text{Aurora}) + d(S,\text{Carteret}) \quad = \quad d(\text{Aurora},\text{Carteret})$$

Substituting in Equation (2) yields

$$d(S,\text{Carteret}) = \frac{1}{2}\underbrace{d(\text{Aurora},\text{Carteret})}_{\text{distance between the focii}} - \underbrace{\frac{c}{2(1+\epsilon)}(\Delta t_3 - \Delta\tau)}_{\text{length of the semi-major axis}}$$

so that a $\delta\Delta\tau = 10$ μs increase of the value of $\Delta\tau$ results in a move of

$$\delta\Delta\tau \times \frac{\partial d(S,Carteret)}{\partial\Delta\tau} = +10^{-6} \times \frac{c}{2(1+\epsilon)} \approx 1.5 \text{ km}$$

towards Aurora and the East on the hyperbola axis i.e. the line joining its focal points Aurora and Carteret.

How did the data get to EUREX?

We have determined with very good precision the time of arrival of the news at the Nasdaq data centre and we have inferred the time of release of the news on K Street. We have also

[7] These results were obtained by numerically solving with a great accuracy the constant-delay equation on the reference oblate spheroid used in WGS84. We made no "Euclidian, planar-inspired" approximation.

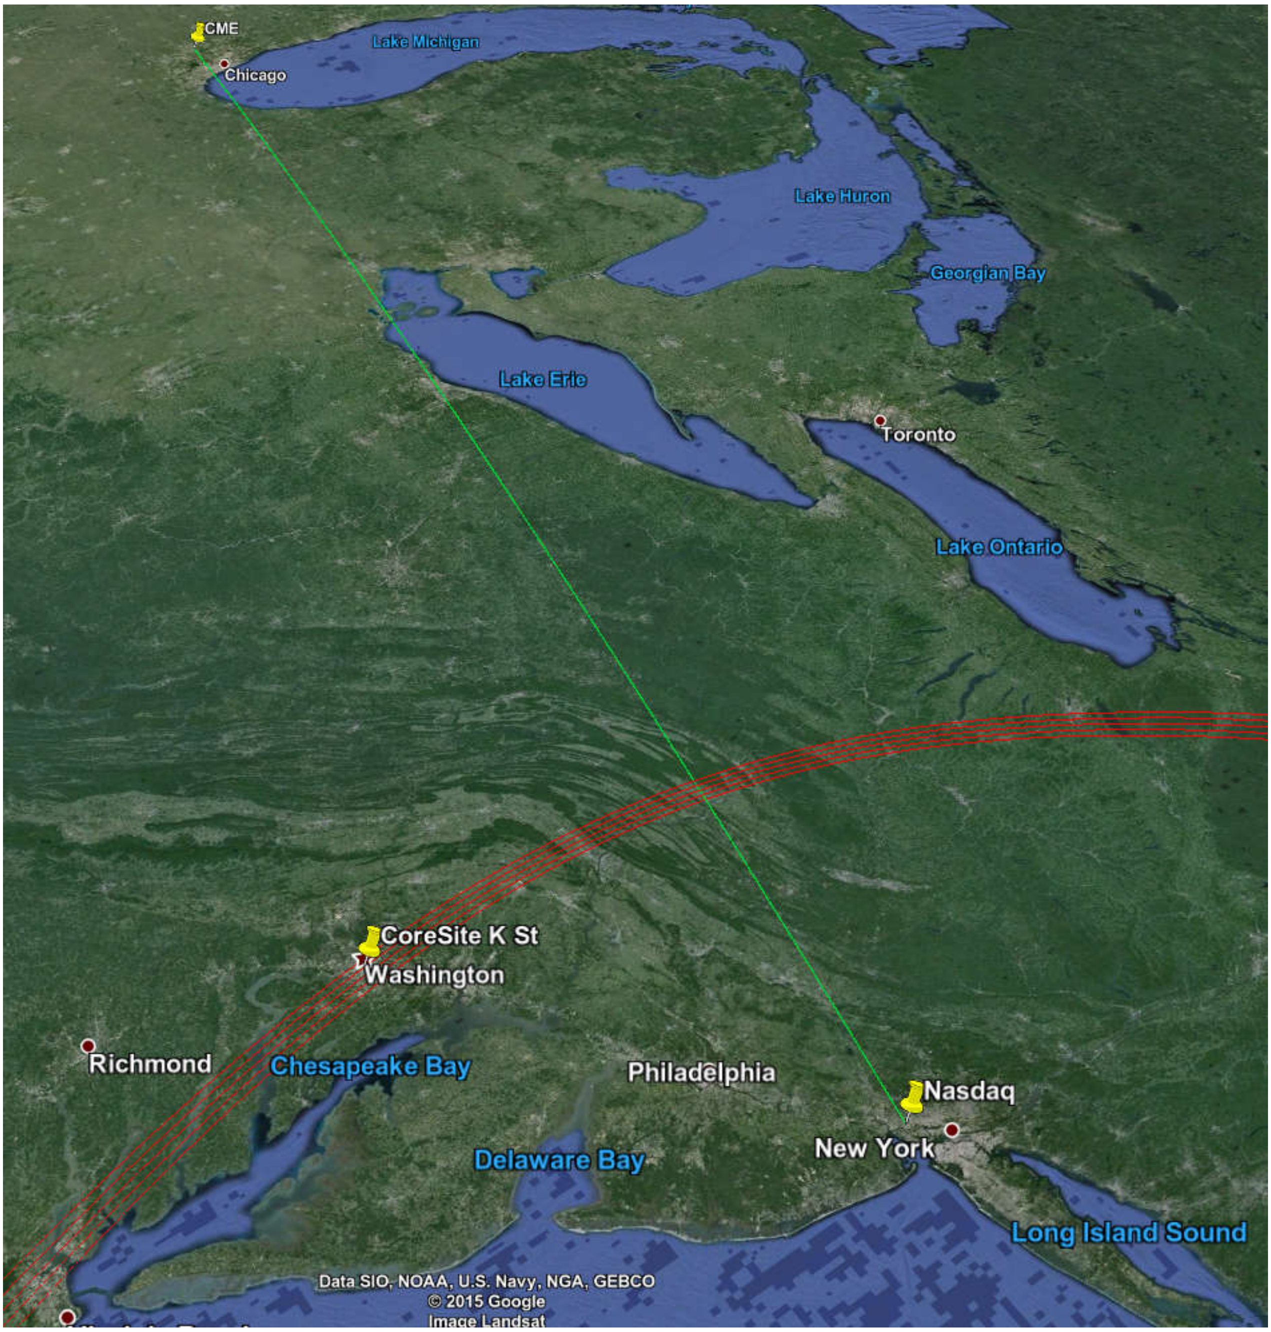

Figure 1: Position lines of the news release point for various assumptions on $\Delta\tau$

determined with high accuracy the arrival time of the news at EUREX in Frankfurt. So we have an upper bound on the latency of the best network from K Street to Frankfurt. The question that we would like to explore is the means of transport of the news to Frankfurt. Can the transmission time be explained by optimal microwave networks and the best transatlantic undersea cable? If the transmission time is lower than this by more than the error margin, then we have to speculate on the transport mechanism and wonder if HF radio signal might have been used.

In February 2015, the fastest transatlantic cable was AC1, owned by Level 3. There are a few public references that provide the latency of the wet portion of the cable. The latency from Brookhaven in Long Island to Whitesands Bay in Cornwall is 58.2 ms round

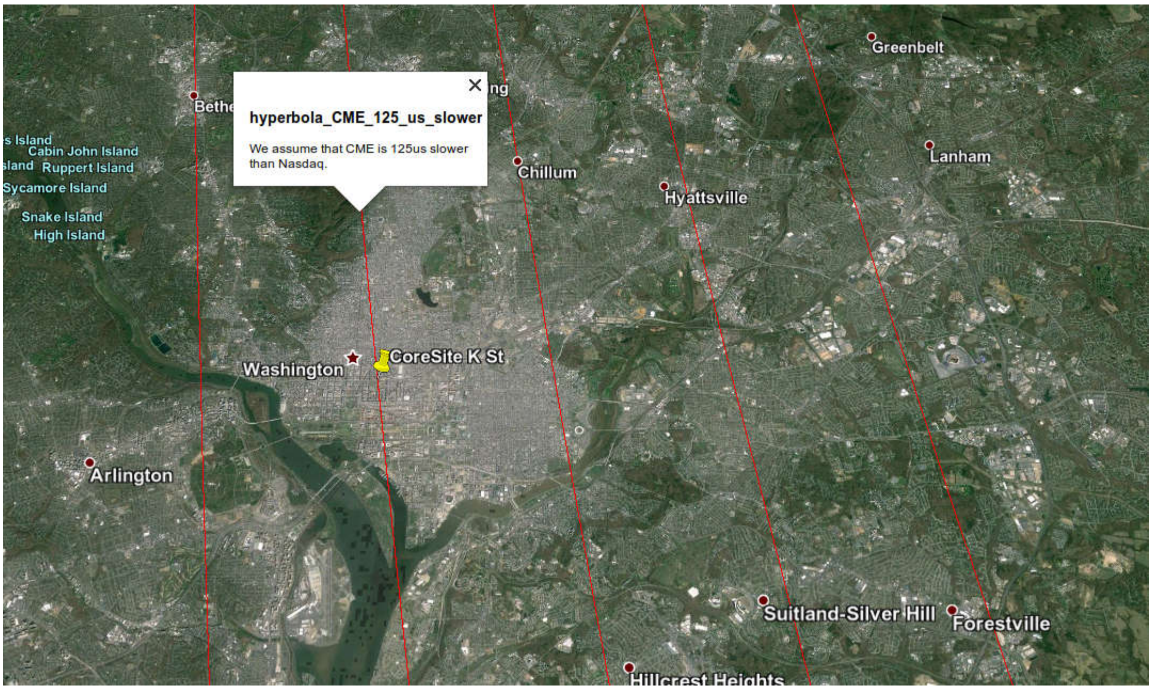

Figure 2: Position lines of the news release point for various assumptions on $\Delta\tau$

trip[8,9,10].

The fastest way to get to Europe from K Street should be a direct microwave network to the AC1 landing station in Brookhaven, the AC1 cable to the UK landing station in Whitesands bay and a direct microwave network from Whitesands bay to the various data centres in Europe. Estimating the latency of the networks on land is easy because microwave networks are straight and optimised for latency. Figures (3) and (4) show our assumptions for optimised paths on land. The latency associated with these path can be computed with Equation (1), which yields respectively 1391 μs and 3497 μs. Adding those two estimates to the latency of the undersea cable yields $t_{opt} = 33988$ μs. Now we have an estimation for the fastest transmission possible to EUREX. Namely the arrival time with optimal wireless and AC1 should be $t_1(EUREX) = t_0 + t_{opt} = 1338729$ μs after 13:30:00 UTC.

The actual arrival time of the news on EUREX is less than 21 μs after this ideal time. This shows that the data transport is consistent with a microwave transport and that these networks are indeed optimised.

Conclusion

We have shown that the study of market data can be used to find the real life latencies of optimised networks and that a simple model for the latencies of these networks was a reasonable approximation.

[8] https://www.nanog.org/meetings/nanog51/presentations/Tuesday/C_FREEDMAN-miami-welcome-to-europe.pdf

[9] http://www.telecomramblings.com/2010/10/industry-spotlight-global-crossings-neil-barua/

[10] We thank an anonymous source for pointing us to this material

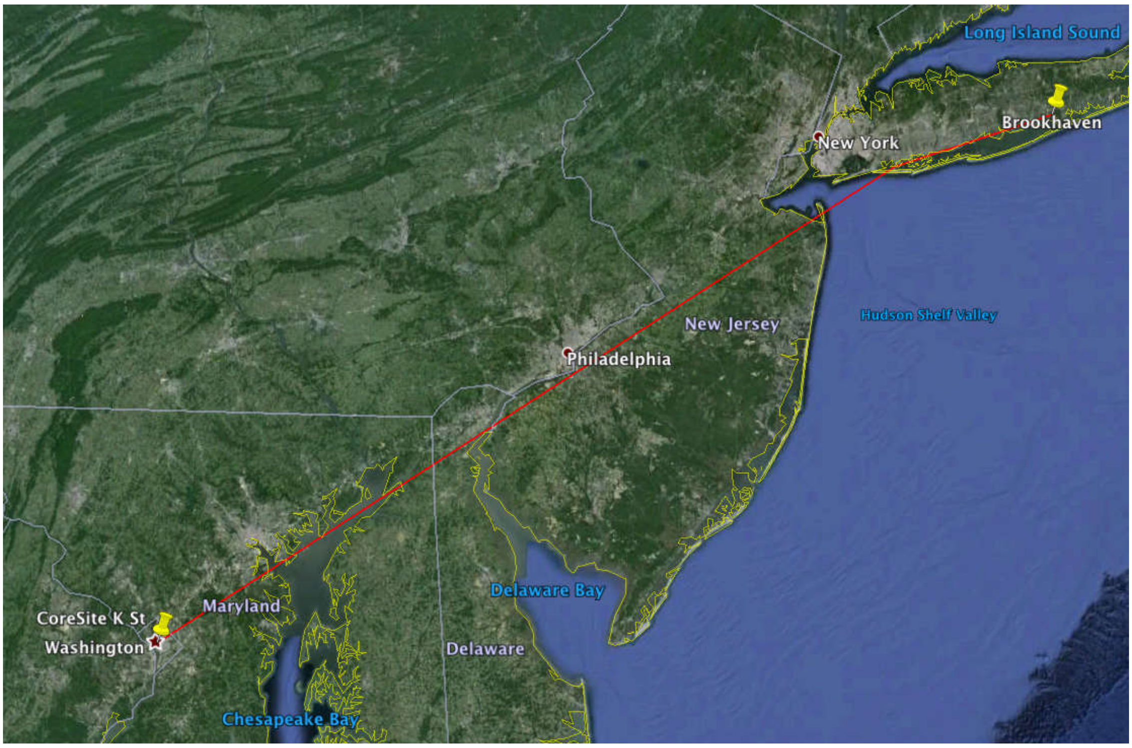

Figure 3: Ideal path from K Street to Brookhaven : 410.8km

Figure 4: Ideal path from Whitesands to EUREX : 1039.4km

Appendix

Nasdaq

Nasdaq publishes its market data with a nanosecond time stamp. However, the round trip latencies are not published, the Nasdaq website only says that[11]:

> [Their] global platform can handle more than one million messages per second at sub-40 microsecond speeds.

The time between order sending and order matching will be estimated at 20 μs, and this is $t_2 - t_1$ with definitions of Table (1).
The uncertainty is small because the orders are sent at a time when Nasdaq is quiet, just before the news is released (see Figure 5).

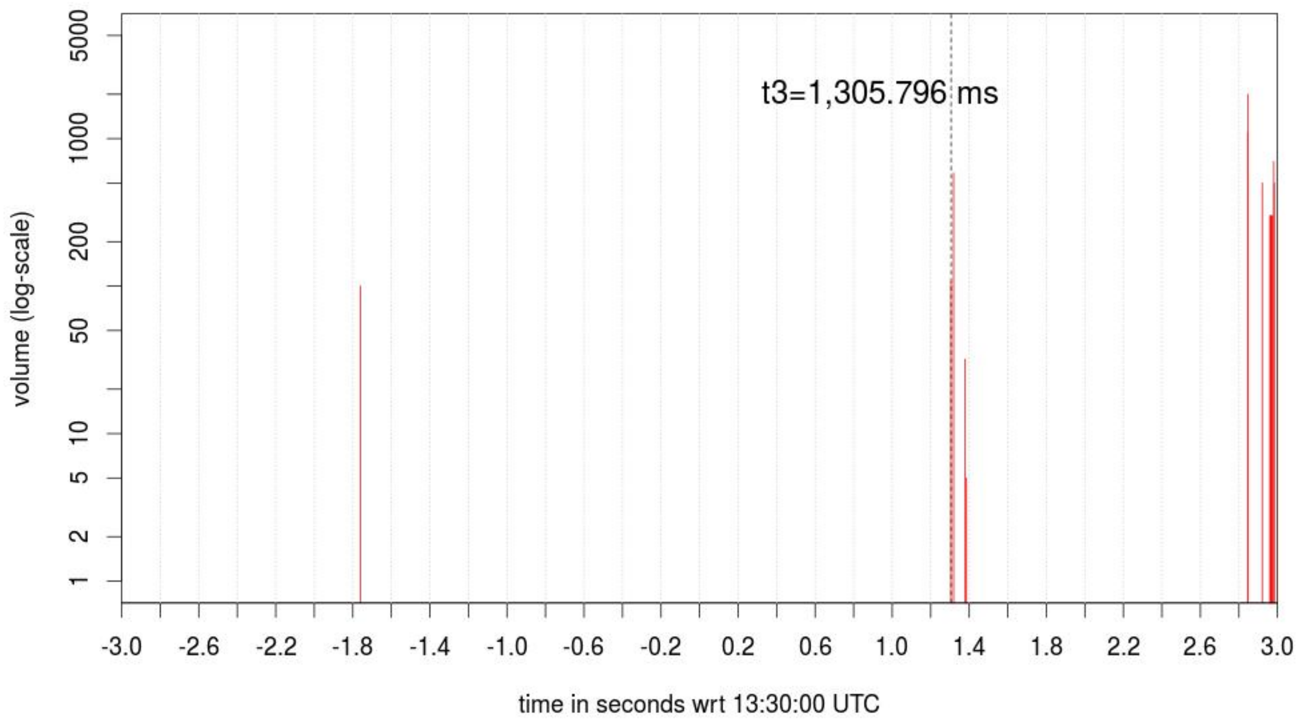

Figure 5: NASDAQ trades on 06FEB2015 around 13:30:00 UTC

There is no trade at all around 13:30:00 UTC until large ones, shown on the market data exerpt below, occur on the SPDR Gold ETF (NSDQ.GLD):

```
Mold Header 20
|- Session                = 000006940B
|- seqNum                 = 2211744
|- msgNum                 = 3

Mkt stamp = 08:30:01.305796061 <=> 30601305796061
Order Executed 31
|- message Type           = E
|- stockLocate            = 3111
|- trackingNumber         = 2
```

[11]Retrieved from http://www.nasdaqtrader.com/Trader.aspx?id=Latencystats on September 7th 2015

```
|- timestamp            = 30601305796061
|- order Ref Number     = 2442738
|- executed Shares      = 100
|- match Number         = 24159

Mkt stamp = 08:30:01.305796061 <=> 30601305796061
Order Executed 31
|- message Type         = E
|- stockLocate          = 3111
|- trackingNumber       = 4
|- timestamp            = 30601305796061
|- order Ref Number     = 735760
|- executed Shares      = 5
|- match Number         = 24160

Mkt stamp = 08:30:01.305796061 <=> 30601305796061
Order Executed 31
|- message Type         = E
|- stockLocate          = 3111
|- trackingNumber       = 6
|- timestamp            = 30601305796061
|- order Ref Number     = 753093
|- executed Shares      = 2
|- match Number         = 24161
```

CME

On the CME we find that the duration between the matching engine time stamp and the publishing time stamp is 2.2 ms. This is $t_4 - t_3$ by the definitions of Table (1). It is probably slowed because of the high activity caused by the news. We have no data to estimate the delay between the sending of an order and its execution within the matching engine, i.e. $t_3 - t_2$.

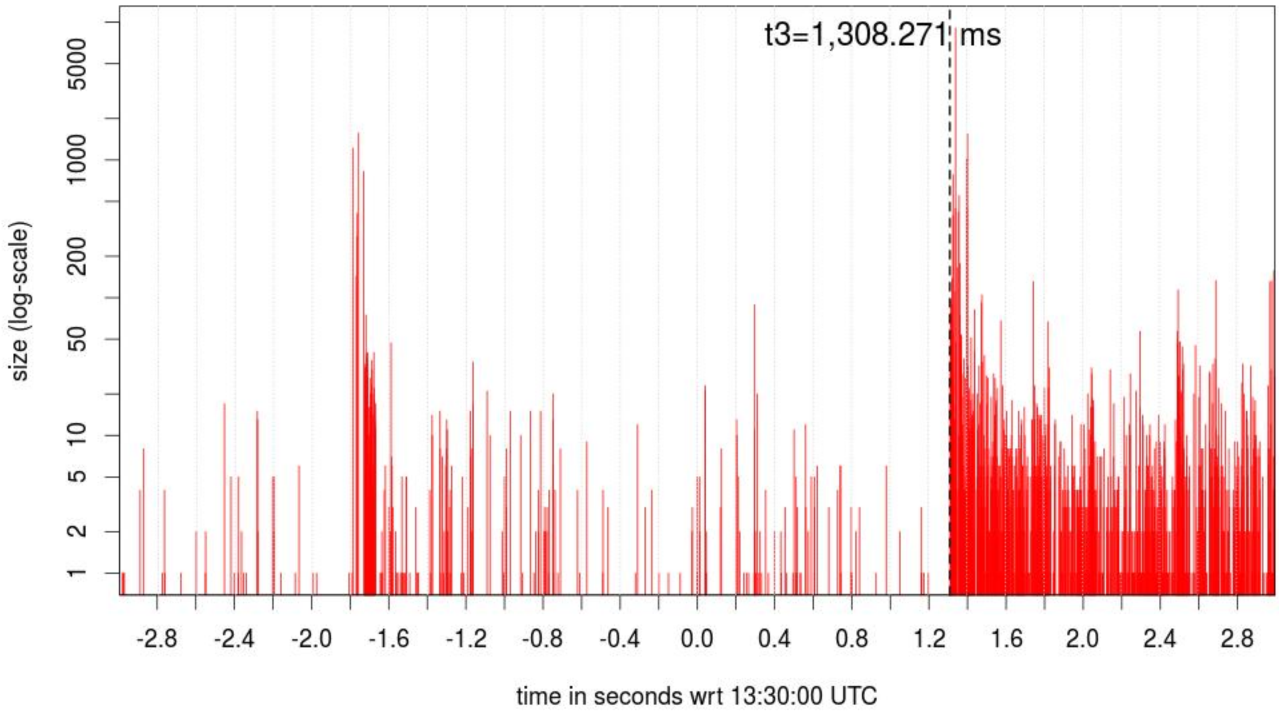

Figure 6: CME trades on 06FEB2015 around 13:30:00 UTC

All trades found in the CME raw feed around 13:30:00 UTC are displayed on Figure (6). Several trades on various maturities of *Light Sweet Crude Oil* (WTI) future occur starting 1,308.270776 milliseconds after 13:30:00 UTC (matching engine time), triggering a sustained trading sequence which first extends to JPY/USD futures.
First data relating to these trades are shown below:

```
MDP Packet Header 12
|-            msgSeqNum = 222028438
|-          sendingTime = 1423229401310473540 => 2015-02-06 13:30:01.310473540
MDP Message Header 10
|- msgLen               = 896
|- blockLength          = 11
|- templateId           = 42
|- schemaId             = 1
|- version              = 5
MDIncrementalRefreshTradeSummary 40
|- timestamp            = 1423229401308270776
|- matchEventInd        = 0
|- Number of actions    = 13
  |- 1/13 MDIncrementalTradeSummaryElmt 40
  |- 1/13 |- mdPrice              = 51870000000 * 1e-7 => 5187
```

```
|- 1/13 |- mdEntrySize          = 6
|- 1/13 |- securityID           = 250614
|- 1/13 |- rptSeq               = 22082533
|- 1/13 |- aggressorSide        = 1
|- 1/13 |- mdUpdateAction       = 0
|- 1/13 |- mdEntryType          = 2
|- 1/13 |- numberOfOrders       = 2
|- 2/13 MDIncrementalTradeSummaryElmt 40
|- 2/13 |- mdPrice              = -850000000 * 1e-7 => -85
|- 2/13 |- mdEntrySize          = 1
|- 2/13 |- securityID           = 108373
|- 2/13 |- rptSeq               = 7316447
|- 2/13 |- aggressorSide        = 0
|- 2/13 |- mdUpdateAction       = 0
|- 2/13 |- mdEntryType          = 2
|- 2/13 |- numberOfOrders       = 2
|- 3/13 MDIncrementalTradeSummaryElmt 40
|- 3/13 |- mdPrice              = 52720000000 * 1e-7 => 5272
|- 3/13 |- mdEntrySize          = 1
|- 3/13 |- securityID           = 301158
|- 3/13 |- rptSeq               = 11822345
|- 3/13 |- aggressorSide        = 0
|- 3/13 |- mdUpdateAction       = 0
|- 3/13 |- mdEntryType          = 2
|- 3/13 |- numberOfOrders       = 2
|- 4/13 MDIncrementalTradeSummaryElmt 40
|- 4/13 |- mdPrice              = 51880000000 * 1e-7 => 5188
|- 4/13 |- mdEntrySize          = 3
|- 4/13 |- securityID           = 250614
|- 4/13 |- rptSeq               = 22082534
|- 4/13 |- aggressorSide        = 1
|- 4/13 |- mdUpdateAction       = 0
|- 4/13 |- mdEntryType          = 2
|- 4/13 |- numberOfOrders       = 2
|- 5/13 MDIncrementalTradeSummaryElmt 40
|- 5/13 |- mdPrice              = 51890000000 * 1e-7 => 5189
|- 5/13 |- mdEntrySize          = 3
|- 5/13 |- securityID           = 250614
|- 5/13 |- rptSeq               = 22082535
|- 5/13 |- aggressorSide        = 1
|- 5/13 |- mdUpdateAction       = 0
|- 5/13 |- mdEntryType          = 2
|- 5/13 |- numberOfOrders       = 2
```

ICE

ICE publishes its time stamps with a millisecond accuracy. The time published is consistent with the time of trades on the CME but does not provide extra information.

LIFFE

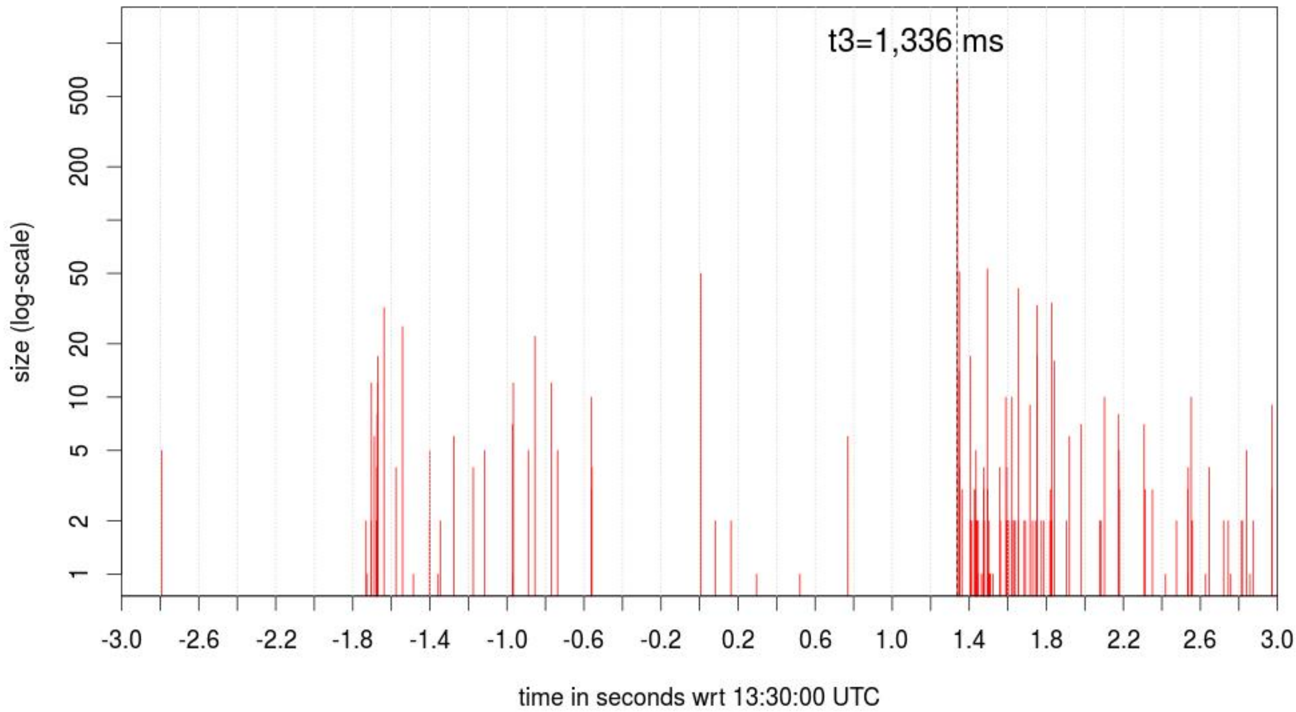

Figure 7: LIFFE trades on 06FEB2015 around 13:30:00 GMT

Looking at the trades which occurred around 13:30:00 UTC (figure 7) and based on the data below, we estimate that the first trade triggered by the news publication occurred 1,336 milliseconds after 13:30:00 UTC (matching engine time, i.e. t_3).

```
TradeMessage:
 |- messageType                    : G
 |- bodyLength                     : 42
 |- marketID                       : 5089543
 |- TradeID                        : 4242284
 |- IsSystemPricedLeg              : N
 |- Price                          : 12198
 |- Quantity                       : 1
 |- OffMarketTradeType             :
 |- TransacDateTime                : 1423229401336
 |- SystemPricedLegType            :
 |- IsImpliedSpreadAtMarketOpen    : N
 |- IsAdjustedTrade                : N
 |- AggressorSide                  : 2
 |- ExtraFlags                     : 0
TradeMessage:
 |- messageType                    : G
```

```
|- bodyLength                      : 42
|- marketID                        : 5089543
|- TradeID                         : 4242254
|- IsSystemPricedLeg               : N
|- Price                           : 12198
|- Quantity                        : 1
|- OffMarketTradeType              :
|- TransacDateTime                 : 1423229401336
|- SystemPricedLegType             :
|- IsImpliedSpreadAtMarketOpen     : N
|- IsAdjustedTrade                 : N
|- AggressorSide                   : 2
|- ExtraFlags                      : 0
```

EUREX

In a presentation entitled *Insights into trading system dynamics*[12], EUREX has published a very detailed analysis of its internal latencies. Based on these numbers, the minimum time between order sending and market data publication is about 200 μs with a variation between 120 and more than 1000. We are interested in the lower number and this is fortunately the more precise one. If the number is higher it would simply mean that the news propagation is faster but we are trying to establish a lower bound and, therefore, the fastest possible time between order sending and market data publication is the relevant time.

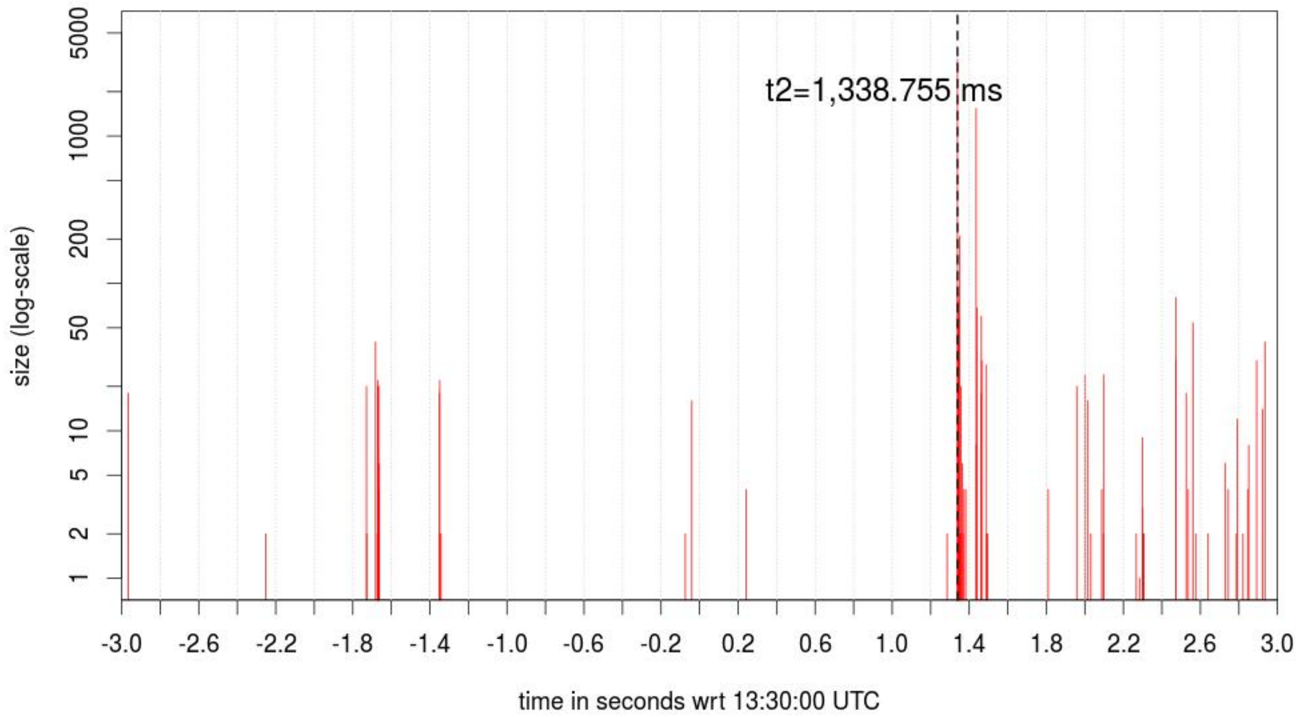

Figure 8: EUREX trades on 06FEB2015 around 13:30:00 UTC

Large trades on the BOBL and then the Bund futures occur about 1,339 milliseconds after 13:30:00 UTC (see figure 8). Since EUREX is providing the gateway arrival times in their EOBI feed (see `requestTime` in excerpt below), we can precisely date the arrival of the first of these trades to 1,338.754921 milliseconds after 13:30:00 GMT.

```
PacketHeaderMessage:
 |- bodyLen                            : 32
 |- templateID                         : 13003
 |- msgSeqNum                          : 4294967295
 |- applSeqNum                         : 159154
 |- marketSegmentID                    : 689
 |- partitionID                        : 4
 |- completionIndicator                : 0
 |- applSeqResetIndicator              : 0
 |- pad5                               :^@^@^@^@
```

[12] http://www.eurexchange.com/blob/238346/6ba9ee0e49239954278105f1c6d44068/data/presentation_insights-into-trading-system-dynamics_en.pdf

```
 |- transactTime                    : 1423229401339494744
ExecutionSummaryMessage:
 |- bodyLen                         : 64
 |- templateID                      : 13202
 |- msgSeqNum                       : 169318
 |- securityID                      : 847252
 |- aggressorTimestamp              : 1423229401338828547
 |- requestTime                     : 1423229401338754921
 |- execID                          : 1423229401338844689
 |- lastQty                         : 766
 |- aggressorSide                   : 2
 |- tradeCondition                  : 255
 |- pad2                            : ^@^@
 |- lastPx                          : 13087000000
 |- restingHiddenQty                : 0
 |- pad4                            : ^@^@^@^@
FullOrderExecutionMessage:
 |- bodyLen                         : 56
 |- templateID                      : 13104
 |- msgSeqNum                       : 169319
 |- side                            : 1
 |- pad7                            : ^@^@^@^@^@^@^@
 |- price                           : 13091000000
 |- trdRegTSTimePriority            : 1423229390373001007
 |- securityID                      : 847252
 |- trdMatchID                      : 3679
 |- lastQty                         : 10
 |- lastPx                          : 13091000000
```