



Bitstamp's New Matching Engine:

How Implementing Nasdaq Technology
Has Improved Trading Frequency and
Order Book Efficiency

In partnership with

Bitstamp

Bitstamp's New Matching Engine: How Implementing Nasdaq Technology Has Improved Trading Frequency and Order Book Efficiency

Sacha Ghebali¹ and Olivier Mammet

Introduction

The matching engine of an exchange is the cornerstone of its infrastructure. Its crucial role is to electronically match the bids and asks sent by traders in a fair and reliable way. In times of intense trading activity, the arrival rate of market orders can increase by orders of magnitude relative to calmer market states. This imposes stringent requirements on the matching engine performance to allow efficient price adjustments and avoid a throttling of the market. The need for more efficient engines has been growing in the cryptocurrency market due to increasing trading activity. In response, Bitstamp recently upgraded their matching engine with Nasdaq's technology. This article attempts to identify and quantify the improvements brought by the new matching engine on trades and order book efficiency.

Key Insights

- The upgrade to the new matching engine imparted a clear change in the capacity to handle a large number of trades over a short amount of time.
- For the most traded pairs, the change is so sharp that the hour of the upgrade can be identified in the data.
- Following the upgrade, a reduction in the autocorrelation of the 6-second returns occurred between two-month windows placed before/after the upgrade. This points to an improved capacity for market participants to quickly adjust prices.
- Relative to the rest of the market, order-book amounts placed close to the best bid/ask prices on Bitstamp appear to have mildly increased from February to May. In addition, bid-ask spreads have improved relative to the rest of the market.

¹ sacha@kaiko.com

Contents

Introduction.....	2
Key Insights.....	2
Content.....	3
Data Used in This Study.....	3
Upgrade of the Matching Engine	4
Impact on Trading Activity	4
Impact on Market Liquidity	7
Conclusion.....	10
Acknowledgements.....	11
More Information	11

Data Used in This Study

The underlying data used in this analysis were collected and normalized by Kaiko and consist of tick-by-tick trades as well as order book snapshots from November 2019 until three months past the upgrade date, i.e. 5 May 2020.

The instruments considered in this analysis consist of BTC/USD, BTC/EUR, XRP/USD, XRP/EUR, LTC/USD, and LTC/EUR on Bitstamp, Bitfinex, BitFlyer Bittrex, Coinbase, Gemini, Itbit and Kraken.

Upgrade of the Matching Engine

Bitstamp's matching engine was upgraded progressively and cautiously across pairs, starting by assets with lower trading activity and finishing with BTC pairs on February 5th. Table 1 shows how the upgrade was rolled out in early 2020 for BTC, XRP, and LTC pairs quoted in USD and EUR. All times are expressed in UTC time zone.

Pairs	Date	Hour
BTC/USD and BTC/EUR	05 Feb 2020	10 a.m.
XRP/USD and XRP/EUR	29 Jan 2020	10 a.m.
LTC/USD and LTC/EUR	22 Jan 2020	10 a.m.

Table 1: Upgrade schedule for BTC, XRP and LTC pairs quoted in USD and EUR.

Throughout this article, the pre-upgrade timeframe will be shaded in light gray in order to easily identify the date at which the new matching engine was introduced.

Impact on Trading Activity

Trades coming into the matching engine are akin to items going through a cash register on a conveyor belt. By making the cash register more efficient, items can flow faster and impatient customers can be done with their shopping faster. Imagine the performance of the register solely depends on the performance of the employee on duty. Alice is the young and multiple-times employee of the month; she can process a whopping 60 articles per minute on a good day. Her colleague Bob should have already retired but still needs to make ends meet, so he does his best to keep up with the pace but he is only able to process 15 articles per minute. On a quiet day, when there is virtually nobody in the shop and articles come sporadically, Bob is doing a fine job and Alice couldn't have done better. However, in the heat of holiday sales, having Alice as a cashier will make an important improvement.

Upgrading the matching engine is similar to shifting from Bob to Alice: the effects can only be seen in periods of high intensity.

Back to the trades, Figures 1 to 3 show the evolution of the number of trades Bitstamp processed over the course of one second during extreme market activity on a daily and quarter-hourly basis (left and right, respectively). **The introduction of the new matching engine is evident for all pairs**

on the daily timeframe with more than a ten-fold increase for BTC pairs, from around 15 trades/s before the upgrade up to about 200 trades/s after the upgrade.

Furthermore, it is remarkable that even the hour at which the upgrade was rolled out can be identified for BTC/USD on the right-hand side of Figure 1.




Figure 1: Maximum number of trades processed in one second over the span of 1 day (left) or 15 minutes (right). The shaded area highlights the time period before the upgrade. Asset: **BTC**




Figure 2: Same description as Figure 1. Asset: **XRP**




Figure 3: Same description as Figure 1. Asset: **LTC**

Next, the autocorrelation of returns for BTC/USD were studied at short timescales before and after the upgrade. Studying the autocorrelation of returns provides information on the nature of the time series: the returns of a price series following a pure random walk would not exhibit any serial correlation, but in practice returns are far from being normally distributed and this is the case too for cryptocurrencies time series. Given a time series of returns, the auto-correlation function quantifies the similarity between the returns and the same series of returns lagged by a fixed amount of time. It can shed light on temporal patterns that are hidden in the noise when simply looking at the returns.

The results are shown in Figure 4. The specific time of 6-second log returns was chosen so that a lag of 10 time intervals would represent one minute. The mean reversion character seen in a negative autocorrelation for small lags is consistent with traditional markets. Interestingly, peaks can be observed at regular intervals, roughly every 30 seconds, with larger peaks at 2 and 4 minutes. This may be the footprint of automated order-splitting strategies. Although other factors or events (such as for example Crypto Black Thursday²) might explain the differences between the two windows analyzed in Figure 4, a significant reduction in the autocorrelation level can be observed at most lags up to 5 minutes. This reduction is particularly striking for short lags between 10 and 30 seconds, a sign of improved market efficiency where market participants are able to execute at a faster pace to adjust prices.




Figure 4: Autocorrelation of 6-second returns, calculated from the trade data for BTC/USD on Bitstamp for the two months preceding/following the upgrade on February 5th.

² Cf. our past analysis *Crypto 'Black Thursday' Under the Microscope* available on our blog at: <https://blog.kaiko.com/crypto-black-thursday-under-the-microscope-a86770df5c29>

Impact on Market Liquidity

Next, the changes in the order book structure are investigated. Figure 5 shows the total amount available across limit orders placed on the order book within 10 basis points away from the best bid/ask prices. In order to put those amounts in perspective, an equally weighted average was assembled from a basket of other major exchanges (dash-dotted line in Figure 5). Comparing the evolution of the market depth to this average across a basket of other crypto-currency exchanges highlights the systemic effects that dominate the market depth evolution: the two curves shown in Figure 5 are highly correlated with a sharp dip following the March 12th crash.




Figure 5: Average daily market depth located within 10 basis points of the best bid/ask price on Bitstamp (orange solid line) and average market depth across Bitfinex, BitFlyer Bittrex, Coinbase, Gemini, Itbit and Kraken order books (equally weighted daily average, black dash-dotted line).

Given the predominance of the impact of market conditions on liquidity, the systemic effects (the tendency of market depth to move in sync across exchanges depending on global market conditions) need to be removed or at least attenuated in order to look for changes in the order book structure following the matching-engine upgrade. The method employed here to reduce those systemic effects is to introduce a relative measure of depth. The latter was constructed by measuring the ratio between the market depth on Bitstamp and the average depth on the set of selected exchanges (using a simple uniformly weighted average – cf. orange and black dash-dotted lines in Figure 5), and shifting this ratio by its value on the day of the upgrade:

$$Relative\ depth(t) = \frac{Bitstamp\ depth(t)}{Basket\ depth(t)} - \frac{Bitstamp\ depth(t_{upgrade})}{Basket\ depth(t_{upgrade})}$$

Where depth is defined as the cumulative sum of limit order amounts placed within a given range around the best bid/ask prices, the upgrade time is shown for each pair in Table 1, and the basket consists of the average depth across Bitfinex, BitFlyer, Bittrex, Coinbase, Gemini, Itbit and Kraken markets.

The above quantity is shown for 5, 10, 20 and 30 basis points away from the best bid/ask prices on Figure 6. Despite the noise present in the signal, a relative increase in market depth between 5 and 10 basis points was observed from January until early May 2020.

Going hand-in-hand with the relative increase in market depth, the bid–ask spread also saw some improvement (i.e. a reduction in spreads) relative to the rest of the market between January and May 2020 as shown in Figure 7. In particular, Bitstamp had a noticeably tight market spread relative to the market in the midst of Crypto Black Thursday, indicating a stronger robustness of the liquidity levels relative to the rest of the market.

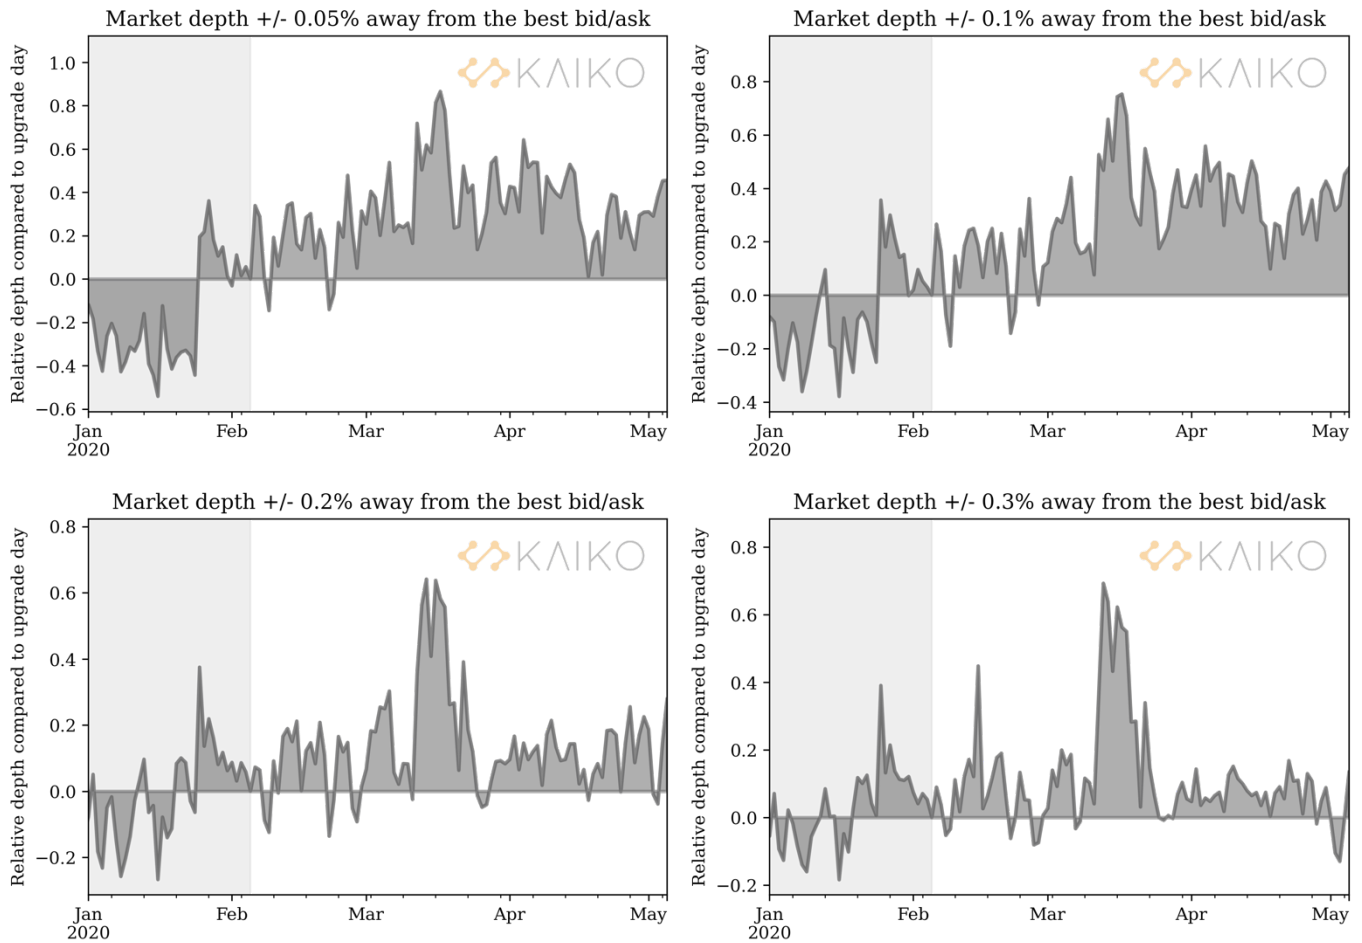


Figure 6: Relative measure of the cumulative market depth (daily average) at various price distances from the best bid/ask (5, 10, 20 and 30 basis points). The shaded rectangular area highlights the time period before the upgrade.




Figure 7: Average daily bid–ask spread on Bitstamp relative to the average spread across Bitfinex, BitFlyer Bittrex, Coinbase, Gemini, Itbit and Kraken.

Conclusion

The matching engine upgrade performed by Bitstamp had a noticeable impact on the observed trading frequency. Almost instantly after the matching-engine upgrade, Bitstamp was able to cope with many more trades in a given second which convincingly demonstrates the effectiveness of the upgrade.

Quantifying the effects on liquidity and trading behavior is a much harder and more subtle task because of the strong systemic fluctuations that largely outweigh these effects. Yet, weaker autocorrelation on short timescales was brought to light on the BTC/USD market when comparing the two months that preceded and followed the upgrade. A reduction of the autocorrelation function was observed on 6-second returns at all lags studied (ranging between 6 seconds and 5 minutes). This points to an enhanced ability for market participants to quickly readjust prices, and to overall greater market efficiency.

A relative measure of market depth and spread was introduced in order to adjust for changing market conditions and allow a comparison of liquidity levels relative to the global market at any given instant. Over the period January – May 2020, a relative improvement in bid–ask spreads and market depth within 10 bps of the best bid/ask prices was observed. These relative measures highlighted a robust liquidity response on Bitstamp’s BTC/USD market relative to the market average during and after the collapse of Crypto Black Thursday.

Acknowledgements

The authors would like to acknowledge interesting discussions with Anastasia Melachrinou and Colin Scanes.

More Information

For more information about our products, please reach out to hello@kaiko.com.

For more of this content, visit our blog at <https://blog.kaiko.com/> or visit our [interactive factsheet](#) to explore liquidity metrics, volatility, asset correlations and more!

[Subscribe to our research newsletter](#) to receive weekly reports on the evolution of the crypto currency market structure.