Better, faster, stronger?

Eurex has poured plenty of development resource into its T7 platform and seems determined to push for more improvements. In the previous edition of *Automated Trader,* Ion Asset Architecture's Dennis Lohfert looked at how the new architecture was performing and the results were impressive. He follows up here with a look at the impact of the latest upgrade and considers whether Eurex can maintain its momentum.

In a previous article, we examined the performance of the new Eurex T7 architecture from Eurex Exchange. By that time, we had only seen the first upgrade iteration, which was to Version 1.1 and which hit rack shelves on July 1. There were not enough data to meaningfully analyse the impact, so we are due for a catch up. (Readers are encouraged to refer to the last issue in order to reference some of the concepts, particularly as far as the meaning of various measurements is concerned.)

Version 1.1 was primarily a performance-tuning update, introducing a few enhancements on the functionality side, but most importantly increasing throughput and decreasing latency.

Eurex's engineering team has been maintaining an aggressive update schedule. Version 1.2 was deployed on September 9 and Version 2.0 is slated for release at the end of November. This is an unusually fast update schedule for a derivatives exchange and is an example of just how much development work is being poured into the platform. It is definitely paying off, as we will see in a moment.

We can analyse and visualise the improvements in different ways. Perhaps the most relevant and most obvious way is to look at how transit times are affected. We collected statistics for 2.5 million DAX futures orders that we sent between the time T7 became the new platform for interesting futures contracts and the end of September, a period spanning about 150 calendar days. As an aside, 99.99% of these orders were passive limit orders; in other words, they were intended to provide liquidity at the time of sending.

In numbers

In order to determine the effect of the various optimisations we calculate the ratio of various transit times before and after the upgrade. In the table below we can see that all numbers are improved, some dramatically more so than others. This is because there are two performance factors that have been improved simultaneously: base latency and throughput. Different parts of the curve are impacted to different degrees. For example, the 99th percentile improvement is primarily a reflection of improved throughput capacity (almost 100% more throughput, which cuts the duration of the 99th percentile in half).

	Core Transit Time				Gateway Transit Time			
	Avg	Min	50%	99%	Avg	Min	50%	99%
T7 Version 1.0	134	60	84	842	350	210	293	1202
T7 Version 1.1	97	37	68	454	275	153	244	707
Reduction	28%	38%	19%	46%	21%	27%	18%	41%

timings are in microseconds

If we look at the official numbers across the entire exchange below, we see numbers that closely match ours, except for the 99th percentile. This makes sense because our order set is much more tied to market activity. As such, the timings for our orders are affected more by the throughput part of the equation in many cases (as opposed to just hitting the "steady state" at random).

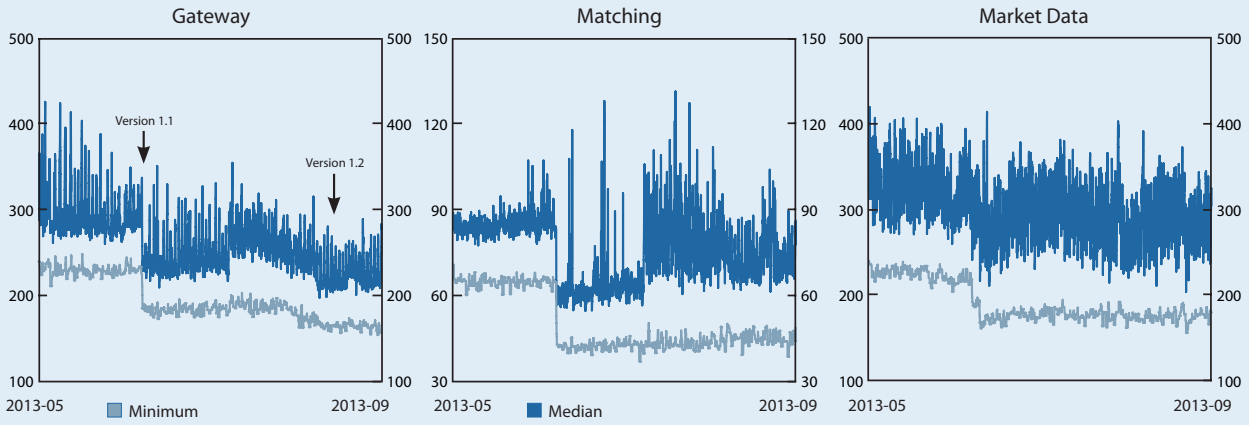
	Core Transit Time			Gateway Transit Time		
	Avg	50%	99%	Avg	50%	99%
Reduction	27%	25%	34%	21%	19%	24%

In any case, the improvements rolled out in July point to a 25% reduction in response time and roughly a doubling of throughput numbers. Both are respectable numbers and we see Eurex being able to make additional improvements. As mentioned earlier, they are maintaining an impressive update schedule. Historically, their engineers have managed time and time again to incrementally squeeze more and more out of their infrastructure.

Evolution over time

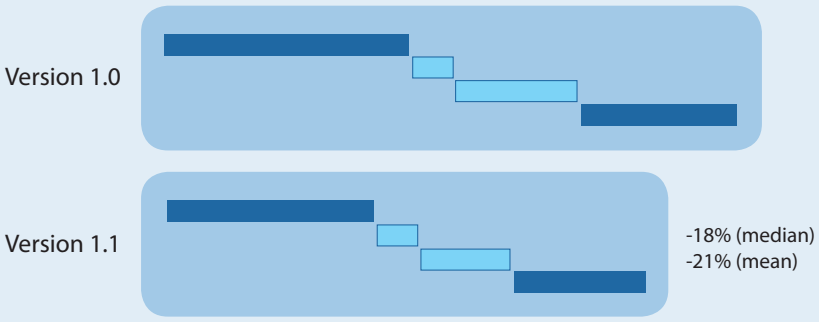
In the figure below, we plot two time series for each measurement of transit time: the minimum over a window of 100 orders and the corresponding median. Minimum transit times, as we noted in the previous article, have special significance, as they tend to represent the "steady state". It is not what you get most of the time, but it does show what the system is capable of and is much less sensitive to noise (caused by market activity), thus making performance improvements more clearly visible. To re-emphasise: this is not what most of your orders will be seeing, it is mostly a tool to keep tabs on the steady state of the exchange system. ▶

Figure 01: Transit Times over Time



Looking at these graphs, we can easily spot the upgrade to Version 1.1 which brought significant improvements across the entire spectrum of the order chain: gateways, matching engine and market data publishers (there are a few more pieces in the infrastructure but this is what we care about).

Figure 02: Improvement by Stage



Where to from here?

The public eye has focused on the supposed detrimental effects of high frequency trading in recent years. As a result, many of the engineering efforts that went into creating new, reliable, deterministic and fast systems have gone largely unnoticed by those not directly involved in trading on these exchanges. These are systems that facilitate the distribution of risk and disseminate timely and accurate pricing data. Some of the engineering has been so successful that it has made fast exchange connectivity essentially a commoditised product. Commoditisation tends to be the end-result of good, broadly based engineering efforts.

There is ever-increasing regulatory transparency on the trading process. As a consequence of the German High Frequency Trading Act, which came into effect in May, Eurex Exchange will introduce new features. Version 2.0 due in November will for the first time require explicit tagging of all algorithmically generated orders for the benefit of regulators (BaFin primarily, but within the context of ESMA this undoubtedly can be shared with every other regulator).

Additionally, we will see a revised implementation of the order-to-trade ratio and excessive system usage charges. The notion here is that the regulators feel that excessive quotation activity without any resulting trades needs to be curtailed. I don't necessarily agree with this notion, but one of the hallmarks of good trading is adapting to new environments and changing circumstances.