

Trading Real-Time Risk and Liquidity

Citi deploys Intel field programmable gate arrays (FPGAs) to accelerate FX options pricing, providing traders with real-time risk management and access to liquidity to hedge their positions.

Authors

Ben Nasatyr

Head of FX Markets Quantitative
Analysis, Citigroup

Graham McKenzie

Field Applications Engineer, Intel
Programmable Solutions Group

Any trading bank needs to manage risk appropriately and make the best use of available liquidity. By working closely with Intel, Citi has been able to integrate field programmable gate arrays (FPGAs) into its high-performance compute and low-latency electronic trading applications to accelerate the pricing of FX options. As a result, Citi FX traders can now assess risk and access liquidity on a real-time basis, in fast-moving markets.

Challenge

- The Citi FX Markets Quantitative Analysis (FX MQA) team is tasked with providing Citi FX traders with real-time valuations of FX options.
- Armed with this information, FX traders can hedge their position giving Citi a competitive edge in the FX market segment.
- To further improve risk calculation time, Citi's FX MQA team needed a technical solution to accelerate the compute-intensive calculations involved in pricing FX options.

Solution

- Citi and Intel collaborated on an initial proof of concept (PoC) to evaluate the performance impact of integrating field programmable gate arrays (FPGAs) into two of Citi's primary FX pricing engines.
- The PoC showed that FPGA-based acceleration provides the performance needed to enable Citi traders to assess risk across large, complex option portfolios and react quickly to changes in market conditions.
- Through high-level design tools and languages, such as OpenCL, the FX MQA team can deploy algorithms on an FPGA without having to employ specialist FPGA developers.

Results

- The FPGA-based system is now in production enabling Citi FX traders to compute risks instantly and access scarce liquidity to hedge their positions on FX options.
- The use of standard servers, processors, acceleration cards, and tools provide a platform that is scalable and extensible to meet growing compute needs.
- Citi's improved risk application allows it to continue to sharpen its competitive edge in the FX market segment.

Real-Time Risk Application Demands Higher Compute Performance

The Citi FX Markets Quantitative Analysis (FX MQA) team supports the bank's FX Options trading desk in calculating the valuations and risks of FX options. Understanding the risks of a large portfolio of complex options enables Citi's FX traders to hedge their position. As such, more accurate and up-to-date calculation of risk leads to a competitive advantage for the bank.

To further strengthen Citi's competitive position in the FX market segment and improve real-time risk application, the FX MQA team sought a technical solution to accelerate the compute-intensive calculations involved in pricing FX options.

Solution Details

The FX MQA team wanted to deploy field programmable gate arrays (FPGAs) in a way that utilized standard hardware in the form of servers and plug-in accelerator cards. It also wanted to standardize programming languages to ease integration.

Citi partnered with Intel to evaluate two of its primary pricing engines on FPGAs. These engines were a Partial Differential Equation (PDE) solver used to price barrier options, and a Vanilla Pricer used to price vanilla options.

For the PDE solver, Citi developed a specification, which was then implemented by Intel. This part of the project was implemented over two phases. The first comprised a standard implementation of a PDE solver and the second added bespoke features for Citi. To learn the development flow and to be sure that it could become self-sufficient in creating and maintaining its own algorithms on the FPGA, the FX MQA team developed the Vanilla Pricer itself.

The use of high-level programming for FPGAs was critical for Citi to be effective at using the technology. OpenCL was used as the programming language, which provides a C/C++ based entry method that is more familiar to software developers than traditional register transfer level (RTL) methods. Coupled with the Intel® FPGA Acceleration Stack, this high-level programming provided a flow that abstracts all the lower-level FPGA details, such as PCIe and memory interfacing, assignments, and timing closure. This flow enables the team to successfully implement the algorithms with limited prior knowledge of FPGA development.

Daan Posthuma from the FX MQA team commented: "The Intel® FPGA development tools are easy to use. As OpenCL is so familiar to C++, it was straightforward to port our code to FPGAs and gain significant speed-ups."

In parallel with the programming efforts of the FX MQA team, which took place in London, the Low Latency team at Citi assessed the full acceleration stack in the Citi Technology Infrastructure organization in New York.

The hardware solution comprised a Dell R740 dual-socket server fitted with Intel® Xeon® Gold 6254 processors and three Intel® Programmable Acceleration Cards with Intel®

Arria® 10 FPGAs (Intel® PACs with Intel® Arria® 10 FPGAs). The team was focused on an acceleration solution that could be utilized by both high-performance compute and low-latency electronic trading applications. The acceleration stack was selected to provide high CPU frequency speed and number of cores, with an optimized BIOS and operating system (OS) to reduce latency and jitter.

Jeremy York, Head of Low Latency Infrastructure at Citi said: "This allowed Citi to reduce footprint and allow applications to run in the same host as the acceleration cards, where the computation part of the algorithm is offloaded from the CPU."

Efficient Management of Risk and Liquidity

The PoC proved to be successful and is now running in production, with both the PDE and Vanilla Pricer exceeding performance goals. The use of FPGAs has enabled Citi to compute its risks instantly and, in fast-moving markets, access scarce liquidity to hedge its positions. By working with Intel, Citi has integrated the Intel PACs and associated frameworks into its existing production environments and software platforms. The FX MQA team has proven that it can efficiently develop algorithms in FPGAs using high-level languages and tool flows.

Ben Nasatyr, head of the FX MQA team, commented: "Intel® FPGAs are a key part of this work, but their performance benefits would have remained inaccessible to my team without a high-level language like OpenCL. The support and partnership we have received from Intel throughout has been the key to making this project a success and has given us confidence that, in the future, we can easily apply the same acceleration techniques to other models and algorithms."

With the current pricers now in deployment, Citi is considering which other algorithms could gain from FPGA acceleration. An obvious next step for Citi is Monte Carlo simulation. The team is now investigating this model with some very promising early results showing a significant acceleration for path generation compared with previous implementations.

"During periods of extreme volatility, such as we saw in the first half of 2020, having FPGA-powered real-time risk proved to be invaluable to efficiently manage risk and liquidity, and to extract significant additional value from our portfolio dynamics."

—Rob Finn
Head of FX Options Trading, Citi.

Spotlight on Citi

Citi, the leading global bank, has approximately 200 million customer accounts and does business in more than 160 countries and jurisdictions. Citi provides consumers, corporations, governments, and institutions with a broad range of financial products and services, including consumer banking and credit, corporate and investment banking, securities brokerage, transaction services, and wealth management. For further information, visit www.citigroup.com

Learn More

You may find the following resources helpful:

- [Intel® Field Programmable Gate Arrays \(Intel® FPGAs\) and Programmable Devices](#)
- [Intel FPGA Products](#)

Find the solution that's right for your organization. Contact your Intel representative or visit:

- <https://www.intel.com/content/www/us/en/programmable/solutions/acceleration-hub/solutions.html#financial>
- <https://www.intel.com/content/www/us/en/programmable/solutions/acceleration-hub/overview.html>



Intel does not control or audit third-party data. You should consult other sources to evaluate accuracy.

Intel technologies may require enabled hardware, software or service activation.

Your costs and results may vary.

© Intel Corporation. Intel, the Intel logo, and other Intel marks are trademarks of Intel Corporation or its subsidiaries. Other names and brands may be claimed as the property of others. 0322/TVL/CAT/PDF Please Recycle 346568 -001EN