Rapid Computation of Value and Risk for Derivative Porfolios

MAXELER Technologies

This is a summary of a special issue paper published online in Wiley Online Library in April 2011.

About the Authors

James Spooner is Head of Acceleration (Finance), Sebastien Racaneire is Senior Acceleration Architect, Stephen Weston is Chief Development Officer and Oskar Mencer is CEO at Maxeler Technologies.

Summary

The use of complex mathematical models for valuing and managing risk of derivatives and asset classes has grown to the point where thousands of CPUs are used daily this computation. Increasing computational and power costs have led financial institutions to pursue ways to add greater computational power while reducing operating costs.

Maxeler Technologies employs a four stage systematic process to accelerate computation, as shown by Figure 1. The analysis stage for collateralized default obligation (CDO) pricing entails mapping the relationship between computation and the input and output data in order to determine the main areas of computation to accelerate. These two areas are calculation of conditional survival probabilities and probability distribution (copula evaluation and convolution). Monte Carlo analysis revealed that payoff scripts accounted for 20% of CPU runtime and thus were a target area for acceleration. The transformation stage examined how loops and control are structured within the C++ code and identified data layout transformations necessary to accelerate the performance of the algorithm. The partitioning stage created a connected block of operations that is both malleable to accelerate and achieve the maximum possible runtime coverage when CPU performance and data input and output are taken into consideration. The implementation stage consisted of kernels that provided multiple compute pipelines. This made it possible for the copula and convoluter to be executed in separate kernels, allowing the two computations to be run independently of each other.

The dataflow engine (DFE) design comprised of kernels that handled the arithmetic computation and managers which manipulated the data input-output of these kernels. The separation of computation and communication allows for highly parallel pipelined kernels. CDO results were obtained using the J.P. Morgan MaxRack equipped with MaxNode-1821 compute nodes and Monte Carlo results were acquired using the J.P. Morgan MaxRack fitted with MaxNode-1834 compute nodes. Maxeler provides a programming environment called MaxCompiler which raises the level of abstraction of DFE design to address the complex programming task.

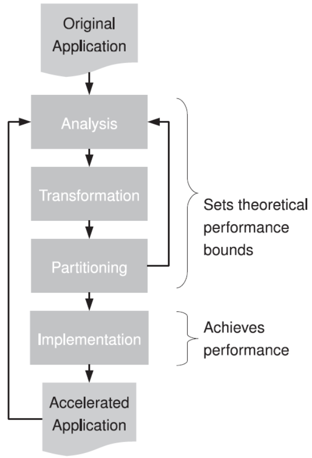

Figure 1: Iterative Maxeler process for accelerating software.

Implementation of CDO pricing targeted a balance of arithmetic optimizations, desired accuracy, power consumption and reliability. In order to meet these specifications, two designs were selected; a full precision design and a reduced precision design. Monte Carlo implementation included random numbers loaded into the main application that were read in more than one process. As a result, it became possible to overlap the software and run three simultaneous processes so the DFE was permanently busy.

The acceleration of this computational application was achieved by optimizing the computer based on the program, rather than the program based on the computer. When assessing performance, 29,250 CDO tranches were used. Utilizing the MaxNode-1821, a speedup over an eight Core Xeon server of 31x was achieved using the full precision design and 37x using the reduced precision design.