

This is a summary of a paper presented at the IEEE Workshop on High Performance Computational Finance, New Orleans, November 2010.

About the Authors

Stephen Weston and Jean-Tristan Marin are with the Applied Analytics group at J.P. Morgan. Oliver Pell is VP of Engineering, James Spooner is Head of Acceleration (Finance) and Oskar Mencer is CEO at Maxeler Technologies.

Summary

Huge growth in the trading and complexity of credit derivative instruments over the past five years has driven the need for ever more computationally-demanding mathematical models. This has led to massive growth in data center compute capacity, power and cooling requirements. We report the results of an on-going joint project between J.P. Morgan and specialist acceleration solutions provider Maxeler Technologies to improve the price-performance for calculating the value and risk of a large complex credit derivatives portfolio. Our results show that valuing tranches of Collateralized Default Obligations (CDOs) on Maxeler accelerated systems is over 30 times faster per cubic foot and per Watt than solutions using standard multi-core Intel Xeon processors.

In this paper we use the stochastic recovery model which has recently become popular for pricing CDO tranches, using the following algorithm:

1. The loss distribution is discretized and computed using convolution, given the conditional survival probabilities and losses from the copula model.
2. The two closest bins are discretized using the standard method with a weighting such that the expected loss is conserved.
3. The final loss distribution is computed using a weighted sum over all of the market factors evaluated, using the copula model.




Figure 1: Tranched credit exposure

Starting with the program which runs on a CPU, we transformed the program into the spatial domain, creating a structure on the FPGA that matched the data-flow structure of the program. Thus, we *optimized the computer based on the program*, rather than optimizing the program based on the computer.

Precision	Speedup
Full Precision	31x
Reduced Precision	37x

Table 2: MaxNode-1821 vs. Eight Core Xeon Server speedup

For the results in this paper we used the J.P. Morgan 10-node MaxRack configured with MaxNode-1821 compute nodes. Each node has eight Intel Xeon cores and two Xilinx FPGAs connected to the CPU via PCI-Express. Comparison was made between two hardware configurations, the difference being the addition of the MAX2-4412C card with two Xilinx FPGAs and 24GB additional DDR DRAM. We saw that by using the FPGAs, the power usage per node is decreased by 6% (see Table 1), even with a 31x increase in computational performance (see Table 2). It follows that the speedup per Watt when computing is greater than the speedup per cubic foot for this application.

Platform	Idle	Processing
Dual Xeon L5430 2.66GHz Quad Core 48GB DDR DRAM	185W	255W
(as above) with MAX2-4412C Dual Xilinx SX240T, 24GB DDR DRAM	210W	240W

Table 1: Power usage for 1U compute nodes when idle and while processing