Purpose

To provide a fast and practical solution to the dose calculation problem, two implementations of the collapsed cone convolution/superposition (CCCS) algorithm based on different technologies: Graphics Processing Unit (GPU) and Field Programming Gate Array (FPGA), are developed and evaluated.

Method and Materials

GPU-based approach and FPGA-based approach have been two most favorable design choices for application acceleration. However, the answer to the question of which one is better is strongly dependent on the target application domain and the problem's computation characteristics. To address this issue, we have implemented both approaches and evaluated them under the same metrics. Both solutions have been thoughoutly optimized to ensure the results reach the limit of hardware.

Results

Both implementations are compared with a commercial multi-threaded implementation running on an Intel quad-core computer with 2.4GHz frequency and 4GB memory. The performance data are collected for different phantom sizes and field sizes. Although the memory technology of the FPGA board is several generations behind the GPU, it still provides same level of speedup. For all test cases, the FPGA board showed a speedup in the range of 21.37-24.26X and the GPU solution (based on GTX260) demonstrates 12.90-16.33X speedup. Some other important design issues, such as cost, system specification, and design efforts are also compared. The anticipated speedups for both platforms are also provided based on the current technology.

Conclusion

Our results have shown that both implementations achieved significant speedup over a multi-threaded software implementation on a Quad-core system. Although FPGA still outperforms GPU in terms of performance, the merits of GPU such as low-cost and off-the-shelf availability make it a preferred solution for many scenarios. Ultra-high-performance scenarios would prefer the FPGA solution as it provides a compact and powerful computation engine within a reasonable budget.