Investigating Design Space of Five Final SHA-3 Candidates in High-Performance FPGAs

Kris Gaj, Ekawat Homsirikamol, and Marcin Rogawski
George Mason University

Abstract

In this paper, we focus on comparing hardware performance of the remaining five final candidates in the SHA-3 contest organized by NIST in the period from 2007 to 2012. The unique feature of our approach is the investigation of multiple hardware architectures of each algorithm. Our goal is analyze the entire performance space in terms of the throughput to area trade-offs, for all Round 3 SHA-3 candidates, as well as the current standard, SHA-2. We perform our investigation using four high-performance FPGA families from two major vendors: Virtex 5 and Virtex 6 from Xilinx, and Stratix III and Stratix IV from Altera. All algorithms have been implemented based on their updated Round 3 specifications, published in January 2011.

Three major performance metrics used in our study are throughput, area, and throughput to area ratio. Throughput is understood as the throughput for long messages, or cumulative throughput for multiple small messages (where processing and input/output functions overlap in time). Area is understood as the number of reconfigurable logic resources, as the use of embedded logic resources is set to zero in this study.

A starting point for our exploration of various architectures of hash functions is the basic iterative architecture, in which each round/step of a hash function is implemented in one clock cycle. Each algorithm is then implemented using multiple architectures based on the concepts of folding, unrolling, and pipelining. Trade-offs between speed and area are investigated, and the best architecture from the point of view of the throughput to area ratio is identified. Finally, all algorithms are ranked based on their overall performance, and the characteristic features of each algorithm important from the point of view of its implementation in hardware are identified.

In case of four out of five candidates (all except JH), the most efficient architecture appeared to be a pipelined architecture. The optimum number of pipeline stages was specific to the algorithm, and was equal to two for Keccak and Groestl, and four for BLAKE. The optimum pipelined architecture for Skein was the architecture with four rounds unrolled, and n pipeline stages, where the optimum value of n was equal to two for Xilinx high-performance FPGAs, and five for Altera high-performance FPGAs.

The results for all investigated functions, and the most successful architectures have been then summarized on the comprehensive throughput vs. area graphs. These graphs have revealed that Keccak is the only candidate that consistently outperforms SHA-2 for all considered FPGA families and two hash function variants (with 256-bit and 512-bit output). JH performed better than SHA-2 in three out of four investigated scenarios, namely it was outperformed by SHA-2 only for the 256-bit function variants implemented using Altera FPGAs. The remaining three Round 3 candidates (BLAKE, Groestl, and Skein) lagged in performance compared to the current standard, SHA-2, for majority of investigated FPGA families and hash function variants.