C vs. VHDL: Comparing Performance of CAESAR Candidates Using High-Level Synthesis on Xilinx and Altera FPGAs

Ekawat Homsirikamol and Kris Gaj
George Mason University

Abstract

The growing number of candidates competing in the cryptographic contests, such as CAESAR, makes the hardware performance evaluation extremely time consuming and tedious, especially at the early stages of the competitions. The main difficulties include the long time necessary to develop and verify Register Transfer Level (RTL), hardware description language (HDL) codes of all candidates, and the need of developing (or at least tweaking) codes for multiple variants and architectures of each algorithm. High-level synthesis (HLS), based on the newly developed Xilinx Vivado HLS tool, offers a potential solution to the aforementioned problems. In order to verify a potential validity of this approach, we have applied both the traditional RTL methodology and the newly proposed HLS-based methodology to the comparison of AES-GCM and ten arbitrarily selected Round 1 CAESAR candidates.

The reference C source codes for the HLS-based approach were based on the submission packages for the respective candidates. Each reference C implementation was then manually modified to create an optimum HLS-ready C code, and verified in software. The C code was then passed as an input to Vivado HLS, and the corresponding VHDL code was automatically generated. If the number of clock cycles was too high, the HLS-ready code was optimized, and the entire process repeated. The VHDL code obtained after the last round of revisions was then simulated for functional correctness, and the number of clock cycles required to process a block of data was verified. If the HDL code performed as expected, this code was benchmarked using ATHENa, and the final netlist verified using timing simulation. Our implementations targeted four modern, high-performance FPGA families: Stratix IV and Stratix V from Altera, and Virtex 6 and Virtex 7 from Xilinx.

Our case study has demonstrated quite substantial (but still far from perfect) correlation in terms of the algorithm rankings according to three basic performance metrics: frequency, throughput, and area. The HLS approach clearly identified two fastest candidates, ICEPOLE and Keyak, and demonstrated that none of the remaining eight CAESAR candidates investigated in this study consistently outperformed the current standard, AES-GCM, in Xilinx or Altera FPGAs, in terms of the throughput to area ratio. The only two other candidates that came relatively close to beating AES-GCM were CLOC and PAEQ.