Fair and Comprehensive Benchmarking of 29 Round 2 CAESAR Candidates in Hardware: Preliminary Results

Ekawat Homsirikamol, William Diehl, Ahmed Ferozpuri,
Farnoud Farahmand, and Kris Gaj
George Mason University

Abstract

Cryptographic contests have emerged as a commonly accepted way of developing cryptographic standards. This process has appeared to work particularly well in case of Advanced Encryption Standard (AES), developed in the period 1997-2001, and Secure Hash Algorithm 3 (SHA-3), developed in the period 2007-2012. In 2013, a new contest, called CAESAR - Competition for Authenticated Encryption: Security, Applicability, and Robustness - has been announced. This contest is currently reaching the end of Round 2, with 29 candidates remaining.

Performance of candidates in hardware has always been a very important evaluation factor, especially at the final stages of the competitions, when all remaining algorithms have been found to have adequate security strength. In CAESAR, for the first time, an attempt has been made to conduct hardware benchmarking of candidates at the very early stages of the contest, when the number of competing algorithms is still very large, namely there are still 29 authenticated cipher families remaining, with multiple variants for some of them (such as PRIMATEs).

This early hardware evaluation has become possible in CAESAR because of the two novel approaches. First, the design teams have been asked to submit their own Verilog/VHDL code before the end of Round 2. Secondly, High-Level Synthesis, based on on the newly developed Xilinx Vivado HLS tool, has been applied to transform reference C implementations of CAESAR candidates to the corresponding efficient Register Transfer Level (RTL), hardware description language (HDL) code.

Our group has contributed to this effort in multiple ways: First, we have proposed a universal hardware Application Programming Interface (API) for authenticated ciphers. The major parts of our proposal include: minimum compliance criteria, interface, communication protocol, and timing characteristics supported by the cores. All of these parts have been defined with the goals of guaranteeing (a) compatibility among implementations of the same algorithm by different designers, and (b) fair benchmarking of authenticated ciphers in hardware. This API is about to be adopted by the CAESAR Committee as a recommended way of implementing all remaining candidates in Verilog and VHDL. Second, we have developed a comprehensive package of VHDL and Python code supporting the development of implementations compliant with our API. This package includes in particular: a) VHDL code of generic pre-processing and post-processing units, common for all modern authenticated ciphers, b) a universal testbench

to verify the functionality of any CAESAR candidate implemented using our hardware API, c) a Python application used to automatically generate test vectors for this testbench, d) VHDL wrappers used to determine the maximum clock frequency and the resource utilization of all implementations, e) RTL VHDL source code of high-speed implementations of AES and the Keccak Permutation F, which may be used as building blocks in implementations of related ciphers, and f) several reference high-speed implementations of Dummy authenticated ciphers, implemented using our API and design methodology. All these resources have been made available in the public domain and documented using a comprehensive Implementer's Guide. Third, we have developed RTL implementations of 16 CAESAR candidates and the current NIST standard, AES-GCM, following the recommended API. Fourth, we have implemented 27 out of 29 Round 2 CAESAR candidates in hardware using the HLS-based approach. The only ones omitted were the two-pass algorithms, AEZ and HS1-SIV. For the 17 ciphers implemented using both RTL and HLS approaches, we have demonstrated that the ranking of candidates remains almost identical, independently of which approach is used.

In this talk, we will present preliminary results of hardware benchmarking for all 29 Round 2 CAESAR candidates, and compare their implementations with the results for the current standard, AES-GCM. For the comparison using the RTL approach, we will present results for all Verilog/VHDL code submitted by the CAESAR design teams, as well as our own implementations. In case multiple implementations are available, only the best results obtained for each candidate will be taken into account. For the comparison using HLS approach, we will rely exclusively on the results obtained using HLS-ready C code generated by our group, based on the reference C implementations, submitted by design teams.

The post place & route results will be generated for multiple FPGA families, and optimized using either ATHENa or 25 default optimization strategies available in Xilinx Vivado. All results will be presented in a graphical and easy to interpret form. Conclusions regarding the suitability of all Round 2 CAESAR candidates for high-speed, non-pipelined hardware implementations will be drawn, and the corresponding recommendations, regarding the advancement to Round 3, made to the CAESAR Committee.

The differences between the RTL and HLS rankings will be identified and analyzed, leading to potential improvements in both kinds of implementations. The feasibility of applying HLS approach alone in the early stages of any future cryptographic contests will be investigated, and the appropriate conclusions drawn.