A 1 Gbit/s Partially Unrolled Architecture of Hash Functions SHA-1 and SHA-512

Roar Lien, Tim Grembowski, and Kris Gaj

ECE Department, George Mason University, 4400 University Drive, Fairfax, VA 22030
kgaj@gmu.edu

Abstract. Hash functions are among the most widespread cryptographic primitives, and are currently used in multiple cryptographic schemes and security protocols, such as IPSec and SSL. In this paper, we investigate a new hardware architecture for a family of dedicated hash functions, including American standards SHA-1 and SHA-512. Our architecture is based on unrolling several message digest steps and executing them in one clock cycle. This modification permits implementing majority of dedicated hash functions with the throughput exceeding 1 Gbit/s using medium-size Xilinx Virtex FPGAs. In particular, our new architecture has enabled us to speed up the implementation of SHA-1 compared to the basic iterative architecture from 544 Mbit/s to 1 Gbit/s using Xilinx XCV1000. The implementation of SHA-512 has been sped up from 717 to 929 Mbit/s for Virtex FPGAs, and exceeded 1 Gbit/s for Virtex-E Xilinx FPGAs.

1 Introduction

Hash functions are very common and important cryptographic primitives. Their primary application is their use for message authentication, integrity, and non-repudiation as a part of the Message Authentication Codes (MACs) and digital signatures [1].

The current American federal standard, FIPS 180-2, recommends the use of one of the four hash functions developed by National Security Agency (NSA) and approved by NIST. By far the most widely used of these four functions is SHA-1 (Secure Hash Algorithm-1), a revised version of the standard algorithm introduced in 1993. The best attack against this algorithm is in the range of 2^{80} operations, which makes its security equivalent to the security of Skipjack and the Digital Signature Standard (DSS). After introducing a new secret-key encryption standard, AES (Advanced Encryption Standard), with three key sizes, 128, 192, and 256 bits, the security of SHA-1 did not any longer match the security guaranteed by the encryption standard. Therefore, an effort was initiated by NSA to develop three new hash functions, with the security equivalent to the security of AES with 128, 192, and 256 bit key respectively. This effort resulted in the development and standardization of three new hash functions referred to as SHA-256, SHA-384, and SHA-512 [1].

T. Okamoto (Ed.): CT-RSA 2004, LNCS 2964, pp. 324–338, 2004.

All four standardized algorithms have a similar internal structure and operation. All of them are based on sequential processing of consecutive blocks of data, and therefore cannot be easily sped up by using pipelining or parallel processing (at least when only one stream of data is being processed).

The majority of reported implementations of SHA-1 based on the current generation of FPGA devices, such as Virtex [2], can only reach the throughputs up to 500 Mbit/s [3-9]. The higher speeds can only be accomplished by using more expensive FPGA devices, such as Virtex-E or Virtex II (see Table 1). Similarly, the FPGA implementations of SHA-512 based on the medium cost Virtex devices reach the speeds in the range of 700 Mbit/s [3, 4].

Significantly higher speeds might be required for applications such as High Definition Television (HDTV), videoconferencing, Virtual Private Networks, etc. [10]. Our goal was to propose, implement, and verify a new architecture of standard hash functions that would allow them to be executed with the throughputs in the range of 1 Gbit/s using medium cost FPGA devices, such as Xilinx Virtex 1000.

2 Hardware Architectures of Hash Functions

A general block diagram common for all four SHA standards and many other dedicated hash functions is shown in Fig. 1. An input message passes first through the preprocessing unit which performs padding and forms message blocks of the fixed length, 512 or 1024 bits, depending on the hash function. The preprocessing unit passes message blocks to the message scheduler unit. Message scheduler unit generates message dependent words, W_t, for each step of the message digest. The message digest unit performs actual hashing. In each step, it processes a new word generated by the message scheduler unit. The message digest is the most critical part of the implementation, as it determines both the speed and area of the circuit.

The most straightforward implementation of the message digest, most often used in practice is shown in Fig. 2a. It is called the basic iterative architecture (or just basic architecture). In this architecture, registers R and H are first both initialized with a value of the constant initialization vector, IV. Subsequently, the architecture executes one step of the message digest per one clock period. In each step t, the message digest accepts a different message dependent word, W_t, and a different step dependent constant, K_t. After executing all steps, the result of the last step, stored in the register R, is added to the previous value of the register H. Then, the processing of the message di-

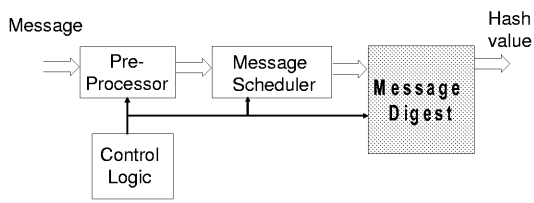

Fig. 1. General block diagram of the hardware implementation of a dedicated hash function, such as SHA-1 and SHA-512

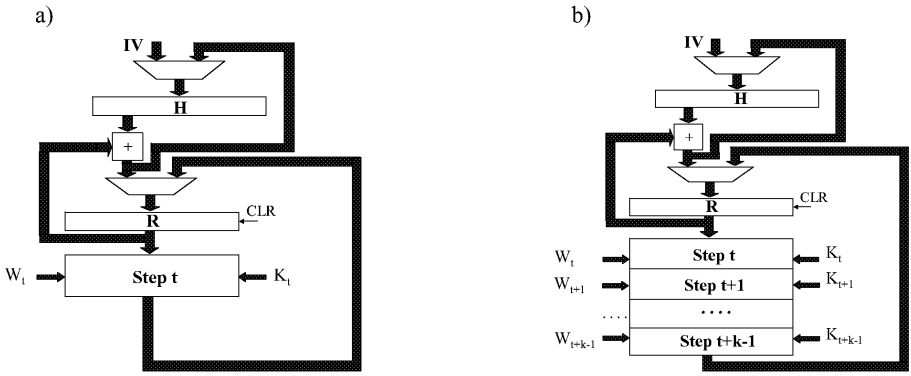

Fig. 2. General diagrams of the message digest units for a) basic architecture, b) partially unrolled architecture with k steps unrolled

gest resumes for a new set of the message dependent words, W_t, corresponding to the new block of the message.

Two straightforward ways of speeding up hardware implementations of hash functions (and any other logic functions) are parallel processing using multiple instantiations of the basic architecture, and pipelining. Out of these two methods, pipelining is more attractive because of the smaller area penalty. Nevertheless, both of these architectures are able to improve an average circuit throughput only under the assumption that multiple independent streams of data are processed simultaneously. If a single long message needs to be hashed, none of these architectures offers *any* improvement in terms of the execution time.

A new architecture of the dedicated hash functions investigated in this paper is shown in Fig. 2b. It is called *partially unrolled architecture*. In this architecture, k steps have been "unrolled" and are executed in the same clock cycle. As a result, the total number of clock cycles necessary to compute one iteration of the message digest has been reduced by a factor of k. At the same time, the critical path through k steps is likely to be significantly shorter than k times the path through a single step. This is because in hash functions, the critical path through a step of the message digest is different for each word of the step input (see Fig. 3).

3 Previous Work

Fully and partially unrolled architectures of dedicated hash functions have been investigated by several authors in the past, but no definite conclusions have been made. In [11] a fully unrolled architecture of MD5 has been compared with a basic iterative architecture. Unrolling of all 64 rounds resulted in a throughput increase by a factor of 2.1, while at the same time the circuit area increased by a factor of 5.4. In [12] a partially unrolled architecture of SHA-1, with the number of rounds unrolled k=5, has been investigated. A high level architecture presented in this paper was very similar to the one proposed in this paper. Nevertheless, the reported results were rather discouraging, with only 11% gain in the circuit throughput and a 43% penalty in the

aging, with only 11% gain in the circuit throughput and a 43% penalty in the circuit area for the partially unrolled architecture over the basic iterative architecture.

All other hardware implementations of dedicated hash functions reported in the literature [9, 13, 14] or available as commercial IP cores [3-8] have followed the basic iterative architecture with only one step of hash function executed in each clock cycle.

4 Details of the Hardware Architectures

4.1 Internal Structure of the Message Digests of SHA-1 and SHA-512

Internal structures of the message digests for SHA-1 and SHA-512 are shown in Fig. 3. In both functions, input registers are initialized with the constant initialization vector, and are updated with the new value in each round. In SHA-1, four out of five words (A, B, C, and D) remain almost unchanged by a single round. These words are only shifted by one position down. The last word, E, undergoes a complicated transformation equivalent to multioperand addition modulo 2^{32}, with five 32-bit operands dependent on all input words, the round-dependent constant K_t, and the message dependent word W_t. The internal structure of the message digest of SHA-512 is similar. The primary differences are as follows: The number of words processed by each round is 8, each word is 64 bits long, and the longest path is equivalent to addition of seven 64-bit operands modulo 2^{64}. These operands depend on seven out of eight input words (all except D), the round-dependent constant K_t, and a message dependent word W_t. Six out of eight input words remain unchanged by a single round.

4.2 Basic Architecture of SHA-1

From Fig. 3a, the critical path of a single SHA-1 round involves the calculation of the chaining variable A at the moment t+1, given by the following formula:

$$A_{t+1} = A_t<<<5 + f_t(B_t, C_t, D_t) + E_t + K_t + W_t + HA'_t$$

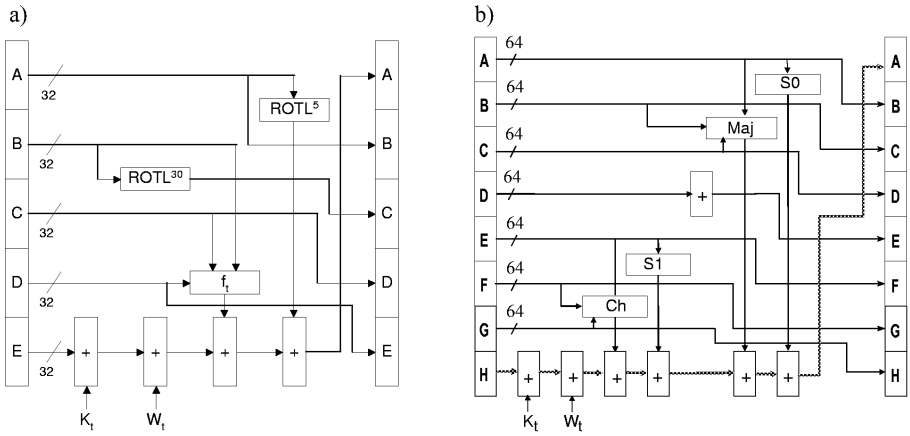

Fig. 3. Internal structure of a single message digest round of a) SHA-1, b) SHA-512

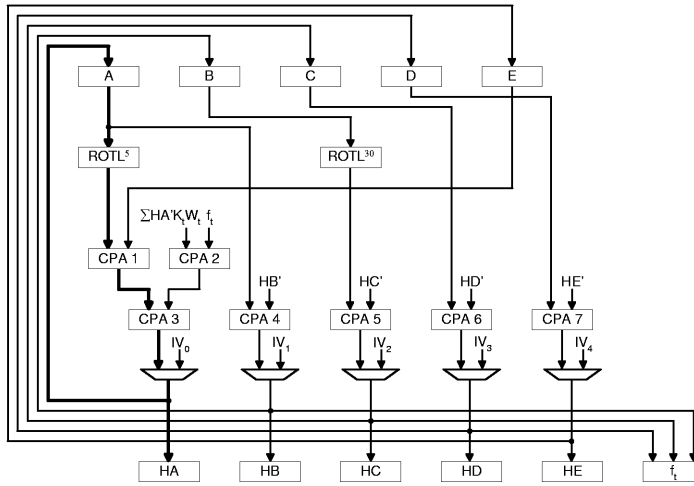

Fig. 4. Our implementation of the message digest unit of SHA-1 in the basic iterative architecture

where X_t is a value of the variable X in the step t, and $HA'_t = HA$ when t=79, otherwise 0. HA is a word A of the register H in Fig. 2a.

Additionally, we know that

$$B_t, = A_{t-1}, \quad C_t = B_{t-1} <<< 30, \quad D_t = C_{t-1}.$$

None of these operations involve any logic, consequently, the expression

$$f_t (B_t, C_t, D_t) = f_t (A_{t-1}, B_{t-1} <<< 30, C_{t-1})$$

can be precomputed in the previous clock cycle, t-1, and will not contribute to the critical path. Similarly, the sum

$$\sum_{HA' Kt Wt} = K_t + W_t + HA'_t$$

can be precomputed by the message scheduler unit, because all values are known already in the previous clock cycle.

As a result, the critical path reduces to the addition of four operands

$$A_{t+1} = A_t <<< 5 + E_t + \sum_{HA' Kt Wt} + f_t (A_{t-1}, B_{t-1} <<< 30, C_{t-1}).$$

All aforementioned optimizations lead to the schematic of the basic architecture of SHA-1 shown in Fig. 4. The lowest level multiplexers choose initialization vectors IV_0 to IV_4 only in the first clock cycle of computations for any new message. The variables HB'.. HE' are equal to HB..HE only in the last step of the message digest computations for a given message block, i.e., only when t=79; otherwise, they are equal to zero.

4.3 Partially Unrolled Architecture of SHA-1

The optimization of the unrolled message digest is relatively straightforward. The general technique employed is to precalculate sums at the earliest possible stage using either regular carry propagate adders (CPAs) or carry save adders (CSAs) (see Fig. 5). The calculations in the critical path follow a sequence of computations described by the equations below:

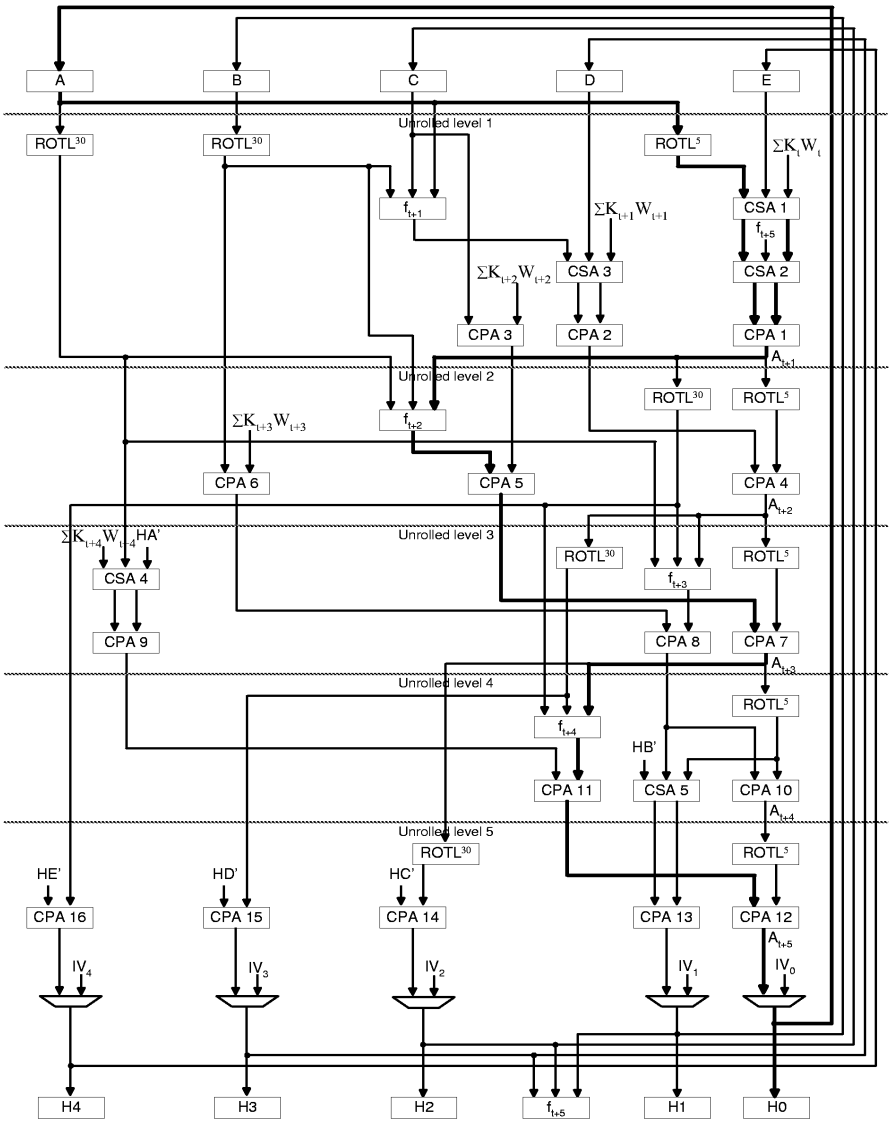

Fig. 5. Our implementation of the message digest unit of SHA-1 in the partially unrolled architecture with 5 steps unrolled

$A_{t+1} = A_t<<<5 + f_t(B_t, C_t, D_t) + E_t + K_t + W_t = A_t<<<5 + f_t(B_t, C_t, D_t) + E_t + \sum K_t W_t$
$A_{t+2} = A_{t+1}<<<5 + f_{t+1}(B_{t+1}, C_{t+1}, D_{t+1}) + E_{t+1} + K_{t+1} + W_{t+1} =$
$= A_{t+1}<<<5 + [f_{t+1}(A_t, B_t<<<30, C_t) + D_t + \sum K_{t+1} W_{t+1}]$
$A_{t+3} = A_{t+2}<<<5 + [f_{t+2}(A_{t+1}, A_t<<<30, B_t<<<30) + [C_t + \sum K_{t+2} W_{t+2}]]$
$A_{t+4} = A_{t+3}<<<5 + [f_{t+3}(A_{t+2}, A_{t+1}<<<30, A_t<<<30) + [B_t<<<30 + \sum K_{t+3} W_{t+3}]]$
$A_{t+5} = A_{t+4}<<<5 + [f_{t+4}(A_{t+3}, A_{t+2}<<<30, A_{t+1}<<<30) + [A_t<<<30 + \sum K_{t+4} W_{t+4} + HA'_{t+4}]].$

At each stage two paths are critical. One is a calculation of the new value of A_{t+i} (i=1..5), which involves rotation by five positions and a single addition. The second is the precalculation of the value of $[f_{t+i} + [E_{t+i} + \sum K_{t+i}W_{t+i}]]$ to be used in the next stage. This precalculation involves the calculation of f_{t+i} and a single addition of a precalculated value $[E_{t+i} + \sum K_{t+i}W_{t+i}]$.

In the first stage of computations (computing A_{t+1}), precalculated values do not exist, so the computations must be performed from scratch. In every second stage starting from stage two, the precomputation of the sum $[f_{t+i} + [E_{t+i} + \sum K_{t+i}W_{t+i}]]$ is the most time consuming operation. Finally, in every second stage starting from stage three, the only contribution to the critical path is a single addition.

4.4 Basic Architecture of SHA-512

From Fig. 3b, the critical path of a single SHA-512 round involves the calculation of the chaining variable A at the moment t+1, given by the following formula:

$$A_{t+1} = S0(A_t) + Maj(A_t, B_t, C_t) + S1(E_t) + Ch(E_t, F_t, G_t) + K_t + W_t + H_t + HA_t'$$

where X_t is a value of the variable X in the step t; S0, Maj, S1, Ch are the logic functions defined in the SHA-512 standard, and $HA'_t = HA$ when t=79, otherwise 0.

Additionally, we know that

$$H_t = G_{t-1}.$$

The functions S0 and Maj execute in parallel in approximately the same amount of time. The same holds true for functions S1 and Ch.

The sum

$$KWHA_t = K_t + W_t + G_{t-1} + HA'_t$$

can be precomputed in the previous clock cycle, t-1.

As a result, the critical path reduces to the addition of five operands

$$A_{t+1} = S0(A_t) + Maj(A_t, B_t, C_t) + S1(E_t) + Ch(E_t, F_t, G_t) + KWHA_t.$$

All aforementioned optimizations lead to the schematic of the basic architecture of SHA-512 shown in Fig. 6. The registers HA-HH are set to the initialization vectors IV_0 to IV_7 only in the first clock cycle of computations for any new message. The multiplexers selecting between HB and '0', HC and '0', etc. choose non-zero values only in the last step of the message digest computations for a given message block, i.e., only when t=79.

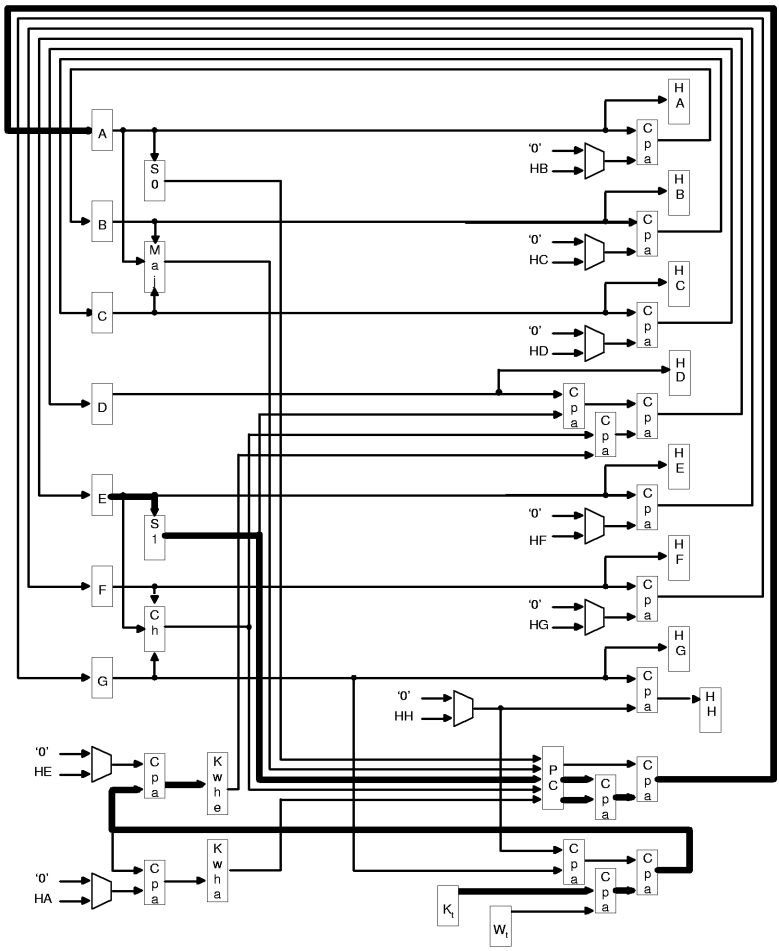

Fig. 6. Our implementation of the message digest unit of SHA-512 in the basic iterative architecture (PC – a 5-to-3 parallel counter, see [9])

4.5 Unrolled Architecture of SHA-512

The unrolled architecture of SHA-512 is shown in Fig.7. Because of the dependence of E_{t+1} on E_t, and A_{t+1} on A_t and E_t (see Fig. 3b), three major critical paths (A0 to A0, E0 to A0 and E0 to E0) exist in the circuit. These paths are marked in Fig. 7 with thicker lines. Values of variables A_{t+i}, and E_{t+i} are denoted as "Ai" and "Ei" respectively, e.g., "E2" denotes E_{t+2}. Precomputations in the previous clock cycle are used to reduce the number of operands in the first four stages of the unrolled architecture. Recall that in the basic architecture, the $KWHA_t$ sum is computed based on the equation $H_t = G_{t-1}$. In the unrolled architecture with k=5, t changes by 5 every clock cycle As a result, $H_t = G_{t-1} = F_{t-2} = E_{t-3} = E_{t+2-5}$ = "E2" in the previous clock cycle.

On the far left side of Fig. 7, "E2" is used to precompute KWH0 (notation for $KWHA_{t+0}$) for the next clock cycle.

$$KWH0 = KWHA_t = K_t + W_t + H_t + HA'_t$$

This method is repeated in stages two to four in order to compute $KWHA_{t+i}$ (denoted in Fig. 7 as KWHi, i=1..3). In stage 5, $H_{t+4} = E_{t+1}$ = "E1", so this value is computed in the same clock cycle, and as a result is not included in the earlier precomputed KWH4 = $KWHA_{t+4}$, which reduces to $KWHA_{t+4} = K_{t+4} + W_{t+4}$. Please, note that in Fig. 7, the sum $K_{t+i} + W_{t+i}$ is denoted as KWi.

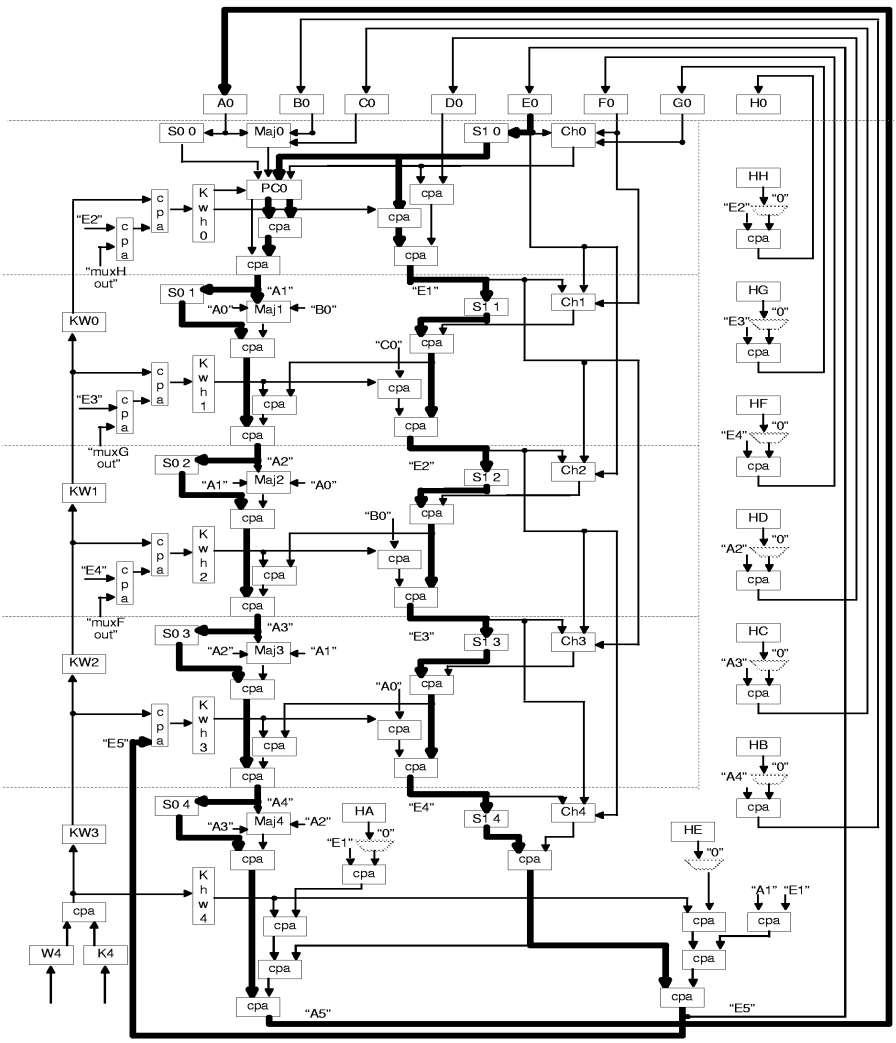

Fig. 7. Our implementation of the message digest unit of SHA-512 in the partially unrolled architecture with 5 steps unrolled

Further reductions in critical paths were accomplished in each stage by adding values of logic functions S1 and Ch as early as possible, reusing values of S1 + Ch, and by selective routing to balance the number of slices in various critical paths.

5 Design Methodology and Results

Our target FPGA device was the Xilinx Virtex XCV1000-6. This device is composed of 12,288 basic logic cells referred to as CLB (Configurable Logic Block) slices, includes 32 4-kbit blocks of synchronous dual-ported RAM, and can achieve synchronous system clock rates up to 200 MHz [2]. XCV1000 was chosen because of the availability of a general purpose PCI board, SLAAC-1V, based on three FPGA devices of this type [10]. Additionally, a new family of Virtex-E Xilinx devices was targeted as well.

All hardware architectures were first described in VHDL, and their operation verified through functional simulation using Active HDL, from Aldec, Inc. Test vectors and intermediate results from the reference software implementations based on the Crypto++ library [15] were used for debugging and verification of VHDL codes. The revised VHDL code became an input to logic synthesis performed using FPGA Compiler II from Synopsys. Tools from Xilinx ISE 4.2 were used for mapping, placing, and routing. These tools generated reports describing area and speed of implementation, a netlist used for timing simulation, and a bitstream used to configure an actual FPGA device. All designs were fully verified through behavioral, post-synthesis, and timing simulations.

The experimental testing of our cryptographic modules was performed using the SLAAC-1V hardware accelerator board, including three Virtex 1000 FPGAs as the primary processing elements. Only one of the three FPGA devices was used to implement hash core.

Test program written in C used the SLAAC-1V APIs and the SLAAC-1V driver to communicate with the board. Our testing procedure is composed of three groups of tests. The first group verifies the circuit functionality at a single clock frequency. The goal of the second group is to determine the maximum clock frequency at which the circuit operates correctly. Finally, the purpose of the third group is to determine the limit on the maximum encryption and decryption throughput, taking into account the limitations of the PCI interface.

In Fig. 8, the minimum clock periods of SHA-1 and SHA-512 obtained using static timing analysis and the experiment are given. For the unrolled architecture, the effective clock period is the minimum time necessary for the data signals to pass the critical path. Since in both our unrolled designs, the data signal is traveling through the critical path over multiple clock periods, the effective clock period is a multiple of the actual clock period. In case of the unrolled architecture for SHA-1 the multiplication factor is 2, in case of the SHA-512 architecture, the multiplication factor is 5.

Based on the knowledge of the minimum clock period, the maximum data throughput has been computed according to the equation:

Throughput=Message_block_size / (Effective_clock_period * Number_of_rounds/k)

The maximum throughput values calculated based on the minimum clock periods obtained using static timing analysis and experiment are shown in Fig. 9. In the same figure, these results are compared with the experimentally measured data throughputs that take into account the delay contributions and the bandwidth limit of the PCI interface. This comparison demonstrates that the PCI interface is capable of operating with a constant uninterrupted data flow up to about 960-990 Mbit/s, and has a negligible influence on the data throughput below this communication rate.

The number of CLB Slices used by our implementations of SHA-1 and SHA-512 are shown in Tables 1 and 2. In SHA-512, four 4 kbit block RAMs are used to store 80 64-bit constants K_t.

Out of the two analyzed hash standards, SHA-1 offers much better potential for loop unrolling. As a result of loop unrolling, the throughput of SHA-1 increased by a factor of almost two (1.9 times), while at the same time its area grew only by a factor of three. SHA-512 is much less suitable for loop unrolling, as its observed speed-up was only 30%, and the area increase 48%.

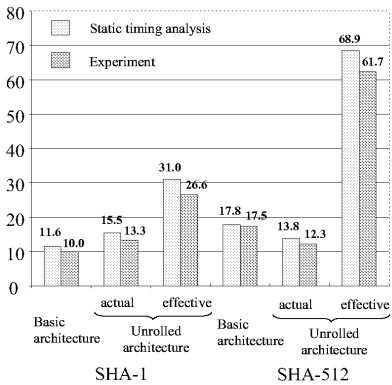

Fig. 8. Minimum clock periods of SHA-1 and SHA-512 in the basic iterative architecture and partially unrolled architecture

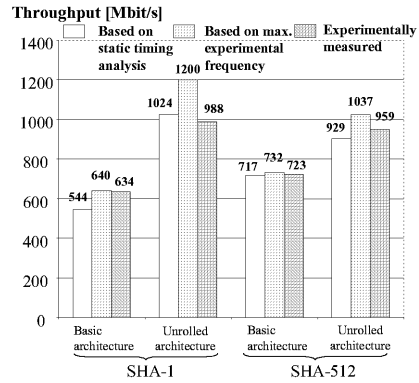

Fig. 9. Maximum throughputs of SHA-1 and SHA-512 in the basic iterative architecture and partially unrolled architecture

6 Comparison with Other Hash Cores

There exist multiple commercial IP cores implementing SHA-1 [3-8]. In Table 2, we present the comparison of our designs for SHA-1 with the most representative IP cores with equivalent functionality. For the Xilinx Virtex family of FPGA devices, our core for SHA-1 in the basic iterative architecture outperforms the second best core (from Helion Technology Ltd) by 13%, using 30% less CLB slices. Our core for the partially unrolled architecture of SHA-1 with 5 rounds unrolled, outperforms all reported Virtex cores by a factor of at least two in terms of throughput, and uses about two times more area. The similar advantages exist for the implementations using

Virtex-E devices, where our core for the unrolled architecture approaches the throughput of 1.2 Gbit/s.

Table 1. Comparison of our designs for SHA-1 with the representative commercial IP cores with equivalent functionality

Source	Clock frequency [MHz]	Throughput [Mbit/s]	Area [CLB Slices]
Xilinx Virtex			
Our, basic	**85**	**544**	**480**
Our, unrolled (k=5)	**64**[1]	**1024**	**1480**
ALMA Technologies	70	442	686
Helion Technology Ltd.	76	480	689
Ocean Logic Pty Ltd	56	352	612
Xilinx Virtex-E			
Our, basic	**103**	**659**	**484**
Our, unrolled (k=5)	**72.5**	**1160**	**1484**
ALMA Technologies	87	549	686
Bisquare Systems Private Limited	66	422	579
Helion Technology Ltd.	95	600	689
Intron, Ltd.	71	449	716
Ocean Logic Pty Ltd	71.5	452	612
Xilinx Virtex-II			
ALMA Technologies	102	644	686
Amphion Semiconductor	99	626	854
Helion Technology Ltd.	103.5	654	569
Ocean Logic Pty Ltd	79	498	612

Table 2. Comparison of our designs for SHA-512 with the representative commercial IP cores with equivalent functionality

Source	Clock frequency [MHz]	Throughput [Mbit/s]	Area [3] [CLB Slices]
Xilinx Virtex			
Our, basic	**56**	**717**	**2384 Slices**
Our, unrolled (k=5)	**67**[2]	**929**	**3521 Slices**
ALMA Technologies	56	707	2690 Slices
Xilinx Virtex-E			
Our, unrolled (k=5)	**72**[2]	**1034**	**3517 Slices**
ALMA Technologies	68	859	2690 Slices
Xilinx Virtex-II			
ALMA Technologies	72	910	2507 Slices
Amphion Semiconductor	50	626	2403 Slices

[1] multi-cycle clock used in the critical path, critical path $\leq 2\ T_{CLK} = 2/f_{CLK}$, 5 steps executed in 2 clock cycles

[2] multi-cycle clock used in the critical path, critical path $\leq 5\ T_{CLK} = 5/f_{CLK}$, 5 steps executed in 5 clock cycles;

[3] each circuit contains additionally 4 Block RAMs

At this point, there are relatively few cores available for the new standard, SHA-512 (see Table 2) [3, 4]. Our implementation of the basic iterative architecture slightly outperforms the equivalent core from ALMA Technologies in terms of throughput, using a smaller amount of FPGA resources. Our partially unrolled architecture is the fastest core for the Virtex family of FPGA devices outperforming the second best core by 30% at the cost of only 31% increase in the circuit area. For the Virtex-E family of FPGA devices our core is the only currently available SHA-512 core that exceeds the throughput of 1 Gbit/s.

7 Comparison with Software Implementations

Efficient software implementations of hash functions have been extensively studied in the literature [17-20]. In [17], basic recommendations on developing an efficient and portable implementation of SHA-1 in C have been formulated. In [18], a close to optimum implementations of dedicated hash functions using Pentium's superscalar architecture have been presented. In [19], software parallelism of all major dedicated hash functions have been studied. Finally, in [20], optimizations targeting Pentium III have been investigated. These optimizations made use of MMX registers and instructions available in Pentium III.

In this paper, we used for comparison, software implementations of SHA-1 and SHA-512, available as a part of the Crypto++ library [15]. Although Crypto++ is not the fastest of the reported software implementations, the reason for using this library was its portability, availability in public domain, and wide practical deployment.

A PC with 2.2 GHz clock, 1 GByte RAM, and cache size 512KB, running Windows XP was used in our measurements. The Crypto++ implementation of hash functions written in C++ was compiled using MS Visual Studio with Service Pack 5. The obtained throughput was 40.5 Mbit/s for SHA-1 and 30.4 Mbit/s for SHA-512. These throughputs were respectively 25 times and 31 times smaller than the throughputs of our partially unrolled hardware implementations of SHA-1 and SHA-512 for Xilinx Virtex 1000-6 FPGAs.

8 Summary

A new *partially unrolled* architecture has been proposed for a family of dedicated hash functions, including four American standard algorithms SHA-1, SHA-256, SHA-384, and SHA-512. The unrolled architecture has been designed, optimized, and experimentally verified for the most widely used hash algorithm, SHA-1, and one of the new hash standard algorithms SHA-512. For the purpose of comparison, the basic iterative architecture has been implemented for both functions as well.

The new architecture appeared to be particularly suitable for the implementation of SHA-1. For the number of rounds unrolled equal to k=5, it allowed to almost double the throughput of SHA-1 compared to the basic iterative architecture, at the cost of increasing circuit area by a factor of three. The similar design for SHA-512 appeared to have much less benefit; the increase in the circuit throughput was only 30%, and the area of the circuit increased by 48%.

This different behavior of two hash algorithms could be easily explained by analyzing the structure of both algorithms. In the unrolled architecture of SHA-1, many message digest steps could be substantially sped up by preprocessing partial results of a given step in the previous steps. The same optimization was not possible in SHA-512 due to sequential dependencies present in the algorithm.

Our partially unrolled implementation of SHA-1 reached the target throughput of 1 Gbit/s in Virtex XCV1000, and outperformed all known to the authors commercial IP cores with equivalent functionality by at least a factor of two. Our implementation of SHA-512 also compared favorably with commercial IP cores, and reached a target throughput of 1 Gbit/s using Virtex-E family of Xilinx FPGAs. To our best knowledge, our implementations of SHA-1 and SHA-512 are the only FPGA implementations of these hash functions available to date that can sustain a throughput over 1 Gbit/s for a single stream of data.

References

1. NIST Cryptographic Toolkit, available at http://csrc.nist.gov/CryptoToolkit/
2. Xilinx, Inc.: Virtex 2.5 V Field Programmable Gate Arrays, available at www.xilinx.com
3. ALMA Technologies web page, available at http://www.alma-tech.com
4. Amphion Semiconductor web page, available at http://www.amphion.com
5. Bisquare Systems Private Limited web page, available at http://www.bisquare.com
6. Helion Technology Limited web page, available at http://www.heliontech.com
7. Intron, Ltd. Web page, available at http://www.intron.lviv.ua
8. Ocean Logic Pty Ltd web page, available at http://www.ocean-logic.com
9. Grembowski T., Lien R., Gaj K., Nguyen N., Bellows P., Flidr J., Lehman T., Schott B.: Comparative Analysis of the Hardware Implementations of Hash Functions SHA-1 and SHA-512, LNCS 2433, 5th International Conference, ISC 2002, Sao Paulo, Brazil, Sep./Oct. 2002, 75–89
10. Bellows P., Flidr J., Gharai L., Perkins C., Chodowiec P., and Gaj K.: IPsec-Protected Transport of HDTV over IP, LNCS 2778, 13th International Conference on Field Programmable Logic and Applications, FPL 2003, Lisbon, Portugal, Sep. 2003, 869–879
11. Deepakumara J., Heys H.M., and Venkatesan R.: FPGA Implementation of MD5 Hash Algorithm, Proc. IEEE Canadian Conference on Electrical and Computer Engineering (CCECE 2001), Toronto, Ontario, May 2001, available at http://www.engr.mun.ca/~howard/PAPERS/ccece_2001.pdf
12. Hoare R., Menon P., and Ramos M.: 427 Mbits/sec Hardware Implementation of the SHA-1 Algorithm in an FPGA, International Association of Science and Technology for Development (IASTED) Journal 2002
13. Ting K.K., Yuen S.C.L., Lee K.H., and Leong P.H.W.: An FPGA Based SHA-256 Processor, Proc. 12th International Conference, FPL 2002, Montpellier, France September 2–4, 2002
14. Kang K.Y., Kim D.W., Kwon T.W., and Choi J.R.: Hash Function Processor Using Resource Sharing for IPSec, Proc. 2002 International Technical Conference On Circuit/Systems, Computers and Communications
15. Crypto++, free C++ class library of cryptographic schemes, available at http://www.eskimo.com/~weidai/cryptlib.html
16. Digital Signature Standard Validation System (DSSVS) User's Guide available at http://csrc.nist.gov/cryptval/shs.html
17. McCurley K.S.: A Fast Portable Implementation of the Secure Hash Algorithm, Sandia National Laboratories Technical Report SAND93–2591

18. Bosselaers A., Govaerts R. and Vandewalle J.: Fast Hashing on the Pentium, in N. Koblitz (Ed.): Advances in Cryptology–CRYPT0 '96, LNCS 1109, Springer-Verlag Berlin Heidelberg 1996, 298–312
19. Bosselaers A., Govaerts R. and Vandewalle J.: SHA: A Design for Parallel Architectures?, in W. Fumy (Ed.): Advances in Cryptology–EUROCRYPT '97, LNCS 1233, Springer-Verlag Berlin Heidelberg 1997, 348–362
20. Nakajima J. and Matsui M.: Performance Analysis and Parallel Implementation of Dedicated Hash Functions, in L.R. Knudsen (Ed.): EUROCRYPT 2002, LNCS 2332, Springer-Berlin Heidelberg 2002, 165–180