

Opera suspicionată (OS)		Opera autentică (OA)
Suspicious work		Authentic work
OS	Mang I., "Considerations on hardware implementations of encryption algorithms", <i>Visnik Sumskovo derjavnovo universitetu, Seria Technichni nauki</i> 12(71), p.41-50, 2004 Disponibil la: http://essuir.sumdu.edu.ua/handle/123456789/10536 .	
OA	Elbirt, A.J., Yip,W., Chetwynd, B., Paar, C., „An FPGA Implementation and Performance Evaluation of the AES Block Cipher Candidate Algorithm Finalists”, The Third Advance Encryption Standard (AES3) Candidate Conference, New York, April 13-14, 2000. Disponibil la: http://www.ei.rub.de/media/crypto/veroeffentlichungen/2011/01/21/aelbirtetalaes3.pdf	

Incidența minimă a suspiciunii / Minimum incidence of suspicion

p.41:6 – p.41:17	p.1:24 – p.2:7
p.41:31 – p.42:36	p.3:5 – p3:28
p.43:Table 1	p.5:Table 1

Fișa întocmită pentru includerea suspiciunii în Indexul Operelor Plagiate în România de la www.plagiate.ro

An FPGA Implementation Performance Evaluation of the AES Block Cipher Counter Algorithm Finalists *

AJ Elbirt¹, W Yi¹, B Chetwyn², C Paar¹
Electrical and Computer Engineering Department
Worcester Polytechnic Institute
100 Institute Road, Worcester, MA 01609, USA

¹ Email: aelbirt, waihyi, christf@ece.wpi.edu

² Email: sungsunge@alum.wpi.edu

The Third Advanced Encryption Standard (AES3) Candidate Conference,
April 13-14, 2000, New York, USA.

Abstract

The technical analysis used in determining which of the Advanced Encryption Standard candidates will be selected as the Advanced Encryption Algorithm includes efficiency testing of both hardware and software implementations of candidate algorithms. Reconfigurable devices such as Field Programmable Gate Arrays (FPGAs) are highly attractive candidates for encryption algorithm agility, physical security, and potentially much higher performance than software solutions. This contribution investigates the significance of FPGA implementations for the Advanced Encryption Standard candidate algorithm finalists. Multiple architectural implementations of each algorithm are explored for each algorithm. A strong focus is placed on high throughput implementations, which are required to support security for current and future high bandwidth applications. The implementations of each algorithm will be compared in an effort to determine the most suitable candidate hardware implementation within commercially available FPGAs.

Keywords: cryptography, algorithm-agility, FPGA, block cipher, VHDL

1 Introduction

The National Institute of Standards and Technology (NIST) has initiated a process to develop a Federal Information Processing Standard (FIPS) for the Advanced Encryption Standard (AES), specifying an Advanced Encryption Algorithm to replace the Data Encryption Standard (DES) which expired in 1998 [1]. NIST has selected candidate algorithms from among AES, resulting in fifteen official candidate algorithms of which five have been selected as finalists. Unlike DES, which was designed specifically for hardware implementation, one of the design criteria for AES candidate algorithms is that they can be efficiently implemented in both hardware and software. Thus, NIST has announced that hardware and software performance measurements will be included in their efficiency testing. So far, however, virtually all performance measures have been restricted to software implementations on various platforms [2].

*This research was supported in part through NSF CAREER award #CCR-9733246.

The advantages of software implementation include ease of use, ease of upgrade, portability, and flexibility. However, software implementation offers only limited physical security, especially with respect to key storage [3] [4]. Conversely, cryptographic algorithms (and their associated keys) that are implemented in hardware are, by nature, more physically secure as they cannot easily be read or modified by an outside attacker [4]. The main disadvantages of hardware implementation are the lack of flexibility with respect to algorithm parameter switch. A promising alternative for implementation lock-in is reconfigurable hardware devices such as Field Programmable Gate Arrays (FPGAs). FPGAs are hardware devices whose functionality is not fixed and which can be programmed in-system. The potential advantages of encryption algorithms implemented in FPGAs include:

Algorithm Agility This term refers to the switching of cryptographic algorithms during operation. The majority of modern security protocols, such as TLS IPsec, allow for multiple encryption algorithms. The encryption algorithm is negotiate during a session; e.g., IPsec allows using DES, 3DES, Blowfish, CAST, IDEA, RC4 and RC6 as algorithms, and future extensions are possible. Whereas algorithm agility is costly with traditional hardware, FPGAs can reprogram on-the-fly.

Algorithm Update It is conceivable that field devices are upgraded with a new encryption algorithm which does not exist (or was not standardize!) at design time. In particular, it is very attractive for modern security protocols to upgrade for use of AE once the selection process is over. Assuming there is some kind of temporary connection to a network such as the Internet, FPGAs equipped for encryption devices can update the new configuration code.

Algorithm Modification There are applications which require modification of a standard algorithm, e.g., by using rotary - axes permutations. Such modifications are easily made with reconfigurable hardware. Similarly, a standard algorithm can be swapped with a rotary one. Also, modifications for certain can be easily changed.

Architecture Efficiency In certain cases, a hardware architecture can be much more efficient if it is designed for a specific set of parameters; e.g., constant multiplication (of integers in Galois fields) is far more efficient than general multiplication. With FPGAs it is possible to design and optimize an architecture for a specific parameter set.

Throughput Although generally slower than an ASIC implementations, FPGA implementations have the potential of running substantially faster than software implementations.

Cost Efficiency The time and costs for developing an FPGA implementation of a given algorithm are much lower than for an ASIC implementation. (However, for high-volume applications, ASIC solutions usually become the more cost-efficient choice.)

Note that algorithm agility remains an open research issue in regards to security, physical security, and the cost associated with current high-end FPGA devices. However, we believe that cost is not a long-term limiting factor, as will be discussed in section 3.3. For these reasons, this paper explores a throughput comparison between the AE finalists algorithms RC6, Rijndael, Twofish, and Blowfish with respect to implementation in state-of-the-art FPGAs. One aspect that seems particularly relevant is the investigation of achievable encryption rates for FPGA-based implementations. We demonstrate that FPGA solutions encrypt at rates in the Giga bit range for all four algorithms investigated, which is at least twice the magnitude faster than most software implementations [1].

What follows is an investigation of the AE finalists to determine the nature of their underlying components. The characterization of the algorithms' components will lead to a discussion of the hardware architectures best suited for implementation of the AE finalists. A performance metric to measure the hardware cost for the throughput achieved by each algorithm's implementations will be evaluated and a

target FPGA will each select the best implementation that are optimized for high-throughput operation within the commercially available device. Finally, multiple architectures' times for the algorithms within the target FPGA will be discussed and the overall performance of the implementations will be evaluated versus typical software implementations.

2 Previous Work

As seen in custom hardware implementations, little work exists in the area of clock circuit implementation within existing FPGAs. DE, the most common clock circuit implementation targeted at FPGAs, has been shown to operate at speeds up to 400 MHz/s [6]. We believe that this performance can be greatly enhanced using today's technology. These speeds are significantly faster than the best software implementations of DE [7] [8] [9], which typically have throughputs of up to 100 MHz/s, although a 137 MHz/s implementation has been reported as well [7]. This performance differential is an extreme result of DE having been designed in the 1970s with hardware implementations in mind.

Other clock circuits have been implemented in FPGAs with varying degrees of success. A typical example is the IDEA clock circuit which has been implemented at speeds ranging from 2.8 MHz/s [10] to 28 MHz/s [11]. Note that while the 28 MHz throughput was achieved in a fully pipeline architecture, the implementation requires four Xilinx XC4000 FPGAs.

The FPGA implementations throughputs for the AE candidates have been shown to be far slower than their software counterparts. Hardware throughput for a DE [12] [13] have been achieved for CA T-26. However, software implementations have resulted in throughputs of 37.8 MHz/s for CA T-26 on a 200 MHz Pentium Pro PC [12], a factor of three faster than FPGA implementations. When scaled to a current 600 MHz Pentium Pro PC, it is expected that the same software implementation would achieve a throughput of 37.8 MHz/s [13], confirming that custom hardware higher data rates are achievable.

When examining the AE finalists, it is important to note that they do not necessarily exhibit similar behavior to DE when comparing hardware and software implementations. One reason for this is that the AE finalists have been designed with efficient software implementations in mind. Additionally, software implementations may execute in parallel, operating at frequencies as high as 800 MHz while typical implementations that target FPGAs reach a maximum clock frequency of 0 MHz.

3 Methodology

3.1 Design Methodology

There are two basic hardware design methodologies currently available: language-based (high level) design and schematic-based (low level) design. Language-based design relies upon synthesis tools to implement the desired hardware. While synthesis tools continue to improve, they rarely achieve the most optimization in terms of the area and speed when compared to a schematic implementation. As a result, synthesized designs tend to be (slightly) larger and slower than their schematic counterparts. Additionally, implementation results can greatly vary depending on the synthesis tool as well as the design being synthesized, leading potentially increase variances in the synthesis results when comparing synthesis tools. This situation is not entirely different from a software implementation of an algorithm in a high-level language such as C, which is also dependent on coding style and compiler quality. As shown in [14], schematic-based design methodologies are no longer feasible for supporting the increase in architectural complexity experienced by modern FPGAs. As a result, a language-based design methodology was chosen as the implementation framework for the AE finalists with VHDL being the specific language chosen.

3.2 Implementations — General Considerations

Each AE finalist was implemented in VHDL using a timing analysis test methodology. The same hardware interface was used for each of the implementations. In an effort to achieve the maximum efficiency possible, note that key scheduling and encryption were not implemented for each of the AE finalists. Because FPGAs may reconfigure in-system, the FPGA may reconfigure for key scheduling and then later reconfigure for either encryption or decryption. This timing is a major advantage of FPGAs implementations over classical ASIC implementations. Run keys for encryption are loaded from the external key bus and are stored in internal registers and all keys must be loaded before encryption may begin. Key loading is idle until encryption is complete. Each implementation was simulated for functional correctness using the test vectors provided in the AE submission package [1] [16] [17] [18]. After verifying the functionality of the implementations, the VHDL code was synthesized, placed and routed, and re-simulated with annotated timing using the same test vectors, verifying that the implementations were successful.

3.3 Selection of Target FPGA

When examining the AE finalists for hardware implementation within an FPGA, a number of key aspects emerge. First, it is visible that the implementation will require a large amount of I/O pins to fully support the 128-bit data stream at high speeds where bus multiplexing is not an option. It is desirable to use the 128-bit input and output data streams to allow for a fully pipeline architecture. Since the run keys cannot change during the encryption process, they may be loaded via a separate key input bus instead of the start of encryption. Additionally, to implement a fully pipeline architecture requires 128-bit wide pipeline stages, resulting in the need for a register-rich architecture to achieve a fast, synchronous implementation. Moreover, it is desirable to have as many register bits as possible for each of the FPGA's configuration units to allow for a regular layout of design elements as well as to minimize the routing requirements between configuration units. Finally, it is critical that fast carry-chaining be provided between the FPGA's configuration units to maximize the performance of AE finalists that utilize arithmetic operations [13] [12].

In addition to architectural requirements, scalability and cost must be considered. We believe that the chosen FPGA should be the most highly available, capable of supporting the largest amount of hardware resources as well as being highly flexible as to yield and performance. Unfortunately, the cost associated with current high-end FPGAs is relatively high (several hundred US dollars per device). However, it is important to note that the FPGA market has historically evolved at an extremely rapid pace, with larger and faster devices being released to industry at a constant rate. This evolution has resulted in FPGA cost-curves that decrease sharply over relatively short periods of time. Hence, selecting a high-end device provides the closest model for the typical FPGA that will be available over the expected lifespan of an AE.

Based on the aforementioned considerations, the Xilinx Virtex XCV1000BG-60-4 FPGA was chosen as the target device. The XCV1000 has 128K bits of embedded RAM, including thirty-two RAM blocks that are separate from the main memory of the FPGA. The 60-pin ball grid array package contains 12 usable I/O pins. The XCV1000 is composed of a 64×6 array of logic blocks (CLBs), each of which acts as a 4-bit element composed of two 2-bit slices for a total of 12288 CLB slices [1]. This type of configuration results in a highly flexible architecture that will accommodate the run functions' use of various functions. Note that the XCV1000 is a first-generation representative of modern FPGAs and that devices from other vendors are not fundamentally different. It is thus hoped that our results carry over, within limits, to other devices.

3.4 Design Tools

FPGA Express, Inc. and Synopsys, Inc. were used to synthesize the VHDL implementations of the AE finalists. As this study focuses on high throughput implementations,

the synthesis tools were set to optimize for speed. As will be discussed in section 6, the resultant implementations exhibit the best possible throughput with the associated cost being an increase in the area required in the FPGA for each of the implementations. Similarly, if the synthesis tools were set to optimize for area, the resultant implementations would exhibit reduced area requirements at the cost of decreased throughput.

XACTste 2.1 by Xilinx, Inc. was used to place and route the synthesized implementations. For the pipeline architectures, a 40 MHz timing constraint was used in the synthesis and place-and-route processes as it resulted in significantly higher system clock frequencies. However, the 40 MHz timing constraint was found to have little effect on the other architectures, resulting in nearly identical system clock frequencies to those achieved without the timing constraint.

Finally, see wave Viewlogic systems, Inc. and Active-HDLTM by ALDEC, Inc. were used to perform behavioral timing simulations for the implementations of the AE finalists. The simulations verified the functionality and the ability to operate at the designated clock frequencies for the implementations.

4 Architecture details in the AES Finalists

Before attempting to implement the AE finalists in hardware, it is important to understand the nature of each algorithm as well as the hardware architectures most suitable for their implementation. What follows is an investigation into the key components of the AE finalists. Based on this breakdown, a discussion is presented on the hardware architectures most suitable for implementation of the AE finalists.

4.1 Comparisons of the AES Finalist Algorithms

Algorithm	XOR	$M^{2^{32}}$ Adder	$M^{2^{32}}$ Subtractor	Fixed Shift	Variable Rotations	$M^{2^{32}}$ Multiplication	$GF(2^8)$ Multiplication	LUT
MAR	•	•	•	•	•	•		•
RC6	•	•		•	•	•		
Rijndael	•			•			•	•
PRESENT	•			•				•
Twofish	•	•		•			•	•

Table 1: AE finalists comparisons [20]

Modern FPGAs have a structure consisting of a two-dimensional array of configuration units interconnected via horizontal and vertical routing channels. Configuration function units are typically configured from look-up-tables and flip-flops. Look-up-tables may be configured as either combinational logic or memory elements. Additionally, many modern FPGAs provide variable-size RAM blocks that may be used as either memory elements or look-up-tables [21].

In terms of complexity, the entries listed in Table 1 that require the most hardware resources as well as computation time are the multi 2^{32} multiplication and the variable rotation entries [20]. Implementing with multipliers in hardware is an inherently difficult task that requires significant hardware resources. Additionally, algorithms that employ large variable rotations require a moderate amount of multiplexing hardware if carefully designed (see section 4.1 for further discussion). Boxes may be implemented in either combinational logic or RAM—the advantages of each of these techniques are discussed in section 4.2. Fast multipliers such as bit-wise XOR, multi 2^{32} addition, subtraction, and fixed value shifting are constructed from simple hardware elements. Additionally, the Galois field multipliers require in Rijndael and Twofish can also be implemented very efficiently in hardware as they are multipliers by a constant. Galois field constant multipliers require far less resources than general multipliers [22].

Based on our evaluation of the AE finalists, the MAR algorithm appears to be the most resource intensive due to its use of large boxes, an $M^{2^{32}}$ multiplier. As a result, it was conjecture

that the MAR algorithm will exhibit lesser performance when compared to the other AE finalists. Due to this evaluation and a lack of evaluation resources, the MAR algorithm was omitted from this study.

4.2 Hardware Architectures

The AE finalists are all composed of a basic building structure (some from either Feistel or substitution-permutation network) where data is iteratively processed through various functions. Based on this building structure, the following architectures listed were investigated as to yield minimize implementations:

- Iterative Looping
- Full Unrolling
- Partial Pipelining
- Partial Pipelining with Sub-Pipelining

Iterative looping over a cipher's round structure is an effective method for minimizing the hardware requirements when implementing an iterative architecture. When only one round is implemented, an n -round cipher must iterate n times to perform an encryption. This architecture has a low register-to-register delay but requires a large number of clock cycles to perform an encryption. This architecture also minimizes in general the hardware requirements for round function implementation but can easily be extended to the hardware requirements for round key and -Box multileveling. Iterative looping is a subset of full unrolling in that only one round is unrolled whereas a full unrolling architecture allows for the unrolling of multilevel rounds, until the total number of rounds required by the cipher. As seen in an iterative looping architecture, a full unrolling architecture where all n rounds are unrolled and implemented as a single computation block maximizes the hardware requirements for round function implementation while the hardware requirements for round key and -Box multileveling is completely eliminated. However, while this architecture minimizes the number of clock cycles required to perform an encryption, it maximizes the worst case register-to-register delay for the system, resulting in an extremely slow system clock.

A partially pipeline architecture offers the advantage of high throughput rates by increasing the number of blocks of data that are being simultaneously processed. This is achieved by relocating the round function hardware and registering the intermediate data between rounds. Moreover, in the case of a full-length pipeline (as opposed from a partial pipeline), the system will wait until a 128-bit block of ciphertext at each clock cycle once the latency of the pipeline has been met. However, an architecture of this form requires significantly more hardware resources as compared to a full unrolling architecture. In a partially pipeline architecture, each round is implemented as the pipeline's atomic unit and are separate by the registers that form the actual pipeline. However, many of the AE finalists cannot be implemented using a full-length pipeline due to the large size of their associated round functions and -Boxes, which must be located n times for an n -round cipher. As such, these algorithms must be implemented as partial pipelines. Additionally, a pipeline architecture can only execute in modes of parallelism which do not require feedback of the encrypted data, such as Electronic Code Book Counter Mode [3, section 4]. When operating in feedback modes such as Cipher Block Chaining Feedback Mode, the ciphertext from one clock must be available before the next clock can be encrypted. As a result, multiple blocks of plaintext cannot be encrypted in a pipeline fashion when operating in feedback modes. For the remainder of our discussion, feedback modes will be abbreviated as FB and non-feedback modes will be abbreviated as NFB.

Sub-pipeline (partially) pipeline architecture is a vantageous when the round function of the pipeline architecture is multilevel, resulting in a large delay between pipeline stages. By adding sub-pipeline stages, the atomic function of each pipeline stage is subdivided into smaller functional blocks. This results in a decrease in the pipeline's delay between stages. However, each subdivision of the atomic function increases the number of clock cycles required to perform an encryption by a factor equal to the number of

su - ivisi ns. At the same time, the num er f lcks f ata that may e erate u n in NFB m e is increase y a fact r equal t the num er f su - ivisi ns. Theref re, f r this technique t e effective, the w rst case elay etween stages will e ecrease y a fact r f m where m is the num er f a e su - ivisi ns. H wever, if the at mic functi n f the artially i eline architecture has a small stage elay, su - ivi ing the stage will achieve n significant ecrease in the w rst case stage elay. In this case, su - i elining w ul result in n significant increase in the system's cl ck frequency ut w ul increase the l gic res urces an el ck cycles require t erf rm an encry ti n, resulting in re use thr ugh ut.

Many FPGAs r vi e em e e RAM which may e use t re lace the r un key an -B x multilexing har ware. By st ring the keys within the RAM lcks, the a r riate key may ea resse ase n the current r un . H wever, ue t the limite num er f RAM lcks, as well as their restricte it wi th, this meth l gy is n t feasi le f r architectures with many i eline stages r unr lle l s. Th se architectures require m re RAM lcks than are ty ically availa le. A iti nally, the switching time f r the RAM is m re than a fact r f three l nger than that f a stan ar CLB slice element, resulting in the RAM element having a lesser s ee -u effect nthe verall im lementati n. Theref re, the use f em e e RAM is n t c nsi ere f r this stu y t maintain c nsistency etween architectural im lementati ns.

5 Architectur 1 Im lement ti n An lysis

F r each f the AE finalists, the fur architecture ti ns escri e in ecti n 4.2 were im mente in VHDL using a tt m-u esign an test meth l gy. The same har ware interface was use f r each f the im lementati ns. R un keys are st re in internal registers an all keys must el a e ef re encry ti n may begin. Key la ing is isale until encry ti n is c m lete . These im lementati ns yiel e a great eal f kn wle ge in regar st the FPGA suita lity f each AE finalist. What f ll ws is a discussi n f the kn wle ge gained regar ing each alg rithm when im mente using the fur architecture ty es.

5.1 Architectur 1 Im lement ti n An lysis — RC6

When im lementing the RC6 alg rithm, it was first determine that the RC6 m ul 2^{32} multi licati n was the minant element f the r un functi n in terms f require l gic res urces. Each RC6 r un requires tw c ies f the m ul 2^{32} multi lier. H wever, it was fun that the RC6 r un functi n es n t require a general m ul 2^{32} multi lier. The RC6 multi liers im mente the functi n $A(2A +)$ which may e im mente as $2A^2 + A$. Theref re, the multi licati n erati n was re lace with an array squarer with summe artial r ucts, requiring fewer har ware res urces an resulting in a faster im lementati n. The remaining c m nents f the RC6 r un functi n — fixe an varia le shifting, it-wise X R, an m ul 2^{32} a iti n — were fun t e sim le in structure, resulting in these elements f the r un functi n requiring few har ware res urces. While varia le shifting erati ns have the tentiat t require c nsi era le har ware res urces, the - it varia le shifting require y the RC6 r un functi n require few har ware res urces. Instead f im lementing a 32-t -1 multi lex r f r each f the thirty-tw r tati n ut ut its (c ntr ille y the five shifting its), a five-level multi lexing a r ach was use . The varia le r tati n is r ken int five stages, each f which is c ntr ille y ne f the five shifting its. F r each r tati n ut ut it f a given stage, a 2-t -1 multi lex r c ntr ille y the stage's shifting it is use . This im lementati n requires a t tal f 160 2-t -1 multi lex rs as se t the thirty-tw 32-t -1 multi lex rs require f r a ne-stage im lementati n. H wever, using 2-t -1 multi lex rs t f rm the five-stage arrel-shifter results in an verall im lementati n that is smaller an faster when c m are t the ne-stage arrel-shifter im lementati n as escri e in [18, ecti n 3.4]. Finally, it was fun that the synthesis t ls c ul n t minimize the verall size f a RC6 r un sufficiently t all w f r a fully unr lle r fully i eline im lementati n f the entire twenty r un s f the alg rithm within the target FPGA.

As discusse in ecti n 4.2, im lementing a single r un f the RC6 alg rithm r vi es the greatest area- timize s luti n. Further l unrling r vi e nly min r thr ugh ut increases as the ecrease in

the number of cycles required to encode was offset by the relatively decreasing system clock frequency. 2-stage partial interleaving was found to yield the highest throughput when operating in FB mode, but performing the single run iterative ring implementation achieving a significantly higher system clock frequency.

When operating in NFB mode, a partially interleaved architecture with two additional sub-interleaved stages was found to offer the advantage of extremely high throughput rates once the latency of the pipeline was met, with the 10-stage partial pipeline implementation laying the best throughput among results. Based on the delay analysis of the partial pipeline implementations, it was determined that nearly two-thirds of the run function's associated delay was attributed to the multi 2^{32} -bit multiplier. Therefore, two additional pipeline stages were implemented as part of the multiplier instead of smaller blocks, resulting in a total of three pipeline stages per run function. As a result, an increase by a factor of more than 2 was seen in the system's clock frequency, resulting in a similar increase in throughput when operating in NFB mode. Furthermore, sub-interleaving was not implemented as this would require sub-interleaving the address registers to sum the partial products (a non-trivial task) to balance the delay between sub-interleaved stages.

5.2 Architecture Implementation Analysis — Rijnael

When implementing the Rijnael algorithm, it was first determined that the Rijnael -B boxes were the dominant element of the run function in terms of required logic resources. Each Rijnael round requires sixteen cycles of the -B boxes, each of which is an 8-bit \times 8-bit \times 8-bit multiplier, requiring significant hardware resources. However, the remaining components of the Rijnael round function — byte swapping, constant Galois field multiplication, and key addition — were found to be simple in structure, resulting in these elements of the run function requiring few hardware resources. Additionally, it was found that the synthesis tool could minimize the overall size of a Rijnael round sufficiently to allow for a fully unrolled fully interleaved implementation of the entire ten rounds of the algorithm within the target FPGA.

Surprisingly, a one-round partially interleaved implementation with one sub-interleaved stage achieved the most area-optimized solution. As compared to a one-stage implementation with no sub-interleaving, the addition of a sub-interleaved stage afforded the synthesizer greater flexibility in its optimizations, resulting in a more area-efficient implementation. While 2-stage unrolling was found to yield the highest throughput when operating in FB mode, the measured throughput was within 10% of the single stage implementation. Due to the parallelistic nature of the lace-and-route algorithms, one can expect a variance in performance based on differences in the starting point of the process. When performing this process multiple times, known as multi-pass lace-and-route, it is likely that the single round implementation would achieve a throughput similar to that of the 2-stage unrolled implementation.

When operating in NFB mode, partial interleaving was found to offer the advantage of extremely high throughput rates once the pipeline latency was met, with the 10-stage partial pipeline implementation laying the best throughput results. While Rijnael cannot be implemented using a fully interleaved architecture due to the large size of the run function, significant throughput increases were seen as compared to the full unrolling architecture.

Sub-interleaving of the partially interleaved architectures was implemented by inserting a pipeline sub-stage within the Rijnael round function. Based on the delay analysis of the partial pipeline implementations, it was determined that nearly half of the run function's associated delay was attributed to the -B boxes substitutions. Therefore, the additional pipeline stage was implemented as to separate the -B boxes from the rest of the run function. As a result, an increase by a factor of nearly 2 was seen in the system's clock frequency, resulting in a similar increase in throughput when operating in NFB mode. Furthermore, sub-interleaving was not implemented as this would require sub-interleaving the -B boxes (a non-trivial task) to balance the delay between sub-interleaved stages.

5.3 Architectural Implementation Analysis — Serial

When implementing the serial algorithm, it was first determined that since the serial -Buses are relatively small (4-bit to 4-bit), it is sufficient to implement them using conventional logic as separate memory elements. Additionally, the -Buses map extremely well to the Xilinx CLB slice, which is common to all 4-bit-to-4-bit slices, allowing nearly all serial elements to be implemented in a total of two CLB slices, yielding a simple implementation which minimizes routing between CLB slices. Finally, the components of the serial round function — key masking, -Bus substitution, and linear transformation — were found to be simple in structure, resulting in the round function requiring few hardware resources.

Implementing a single round of the serial algorithm requires the greatest area-time product. However, a significant performance improvement was achieved by performing 8-rounds in parallel, reducing the need for 4-bit multi-lexing hardware as needed for each serial -Bus grouping is now included within one of the eight rounds. This amounts to parallelizing achieving a significant performance increase with little increase in hardware resources due to the compact nature of the serial round function. As expected, parallelizing thirty-two rounds of the serial algorithm results in a lesser performance when compared to the eight-round implementation. Implementing the thirty-two rounds of the algorithm in conventional logic severely hampered the overall clock frequency of the system, verifying the performance increase caused by the removal of the multi-lexing hardware required to switch between keys.

When operating in NFB mode, a full-length pipeline architecture was found to offer the advantage of extremely high throughput rates once the latency of the pipeline was met, but performing smaller partially-pipeline implementations. In the fully-pipeline architecture, all of the elements of a given round function are implemented as conventional logic. Other AE finalists cannot implement using a fully-pipeline architecture due to the larger round functions. However, due to the small size of the serial -Buses (4-bit to 4-bit slices), the cost of -Bus replacement is minimal in terms of the required hardware.

Finally, serializing of the partially-pipeline architectures was determined to yield in throughput increase. Because the round functions components are all simple in structure, there is little performance to gain by serializing them with registers in an attempt to reduce the delay between stages. As a result, the increase in the system's clock frequency would not outweigh the increase in the number of clock cycles required to perform an encryption, resulting in a performance degradation.

5.4 Architectural Implementation Analysis — Twofish

When implementing the Twofish algorithm, it was first determined that the synthesis tools were unable to minimize the Twofish -Buses to the extent of the AE finalist algorithms due to the -Buses being key-dependent. Therefore, the overall size of a Twofish round was too large to allow for a fully parallel pipeline implementation of the algorithm within the target FPGA. Moreover, the key-dependent -Buses were found to require nearly half of the delay associated with the Twofish round function.

As expected, implementing a single round of the Twofish algorithm requires the greatest area-time product in terms of total CLB slices required for the implementation. Additionally, parallelizing round increases as the increase in the number of cycles per encryption clock was offset by the rapidly increasing system clock frequency. However, single stage partial pipelining with no serial pipeline stage was found to yield the best throughput when operating in feedback mode. With a small increase in the required hardware resources, the serial-pipeline architecture was able to reach a significantly faster system clock frequency as compared to the parallel-pipeline implementation.

When operating in NFB mode, a partially-pipeline architecture was found to offer the advantage of extremely high throughput rates once the latency of the pipeline was met, with the 8-stage partial pipeline implementation laying the best throughput results. While Twofish cannot be implemented using a fully-pipeline architecture due to the large size of the round function, significant throughput increases were seen as compared to the parallel-pipeline architecture.

Finally, substituting of the partially pipeline architectures was implemented by inserting a pipeline stage within the two fish run function. Based on the delay analysis of the partial pipeline implementations, it was determined that nearly half of the run function's associated delay was attributed to the -B boxes from the rest of the run function. As a result, an increase of a factor of nearly 2 was seen in the system's clock frequency, resulting in a similar increase in throughput when operating in NFB mode. Furthermore, pipelining was not implemented as this would require substituting the -B boxes (a non-trivial task) to balance the delay between pipeline stages.

6 Performance Evaluation

Tables 2 and 3 detail the throughput measurements for the implementations of the three architecture types for each of the AE finalists for the NFB and FB mode. The architecture types — 1-unit timing (LU), full partial pipelining (PP), and partial pipelining with pipeline (P) — are listed along with the number of stages and (if necessary) pipeline stages in the associated implementation; e.g., LU-4 implements a 1-unit timing architecture with four runs, while P-2-1 implements a partially pipeline architecture with two stages and one pipeline stage per pipeline stage. As a result, the P-2-1 architecture implements two runs of the given cipher with a total of two stages per run. Throughput is calculated as:

$$\text{Throughput} := (128 \text{ Bits} * \text{Clock Frequency}) / (\text{Cycles Per Encryption Block})$$

Note that the implementation of a single stage partial pipeline architecture, an iterative timing architecture, and a non-run 1-unit timing architecture are all equivalent and are therefore not listed separately. Also note that the compute throughput for implementations that employ any form of hardware pipelining (as discussed in section 4) are made assuming that the pipeline latency has been met.

The number of CLBs required as well as the maximum operating frequency for each implementation was taken from the Xilinx report files. Note that the Xilinx tools assume the absolute worst-case operating conditions — highest possible operating temperature, lowest supply voltage, and worst-case fabrication tolerance for the specific grade of the FPGA [23]. As a result, it is common for actual implementations to achieve slightly better performance results than those specified in the Xilinx report files.

While this study focuses on high throughput implementations, the hardware resources required to achieve this throughput is also a critical parameter. No standard metric exists to measure the hardware resources consumed associated with the measured throughput of an FPGA implementation. Two area measurements of FPGA utilization are reasonably apparent — logic gates and CLB slices. It is important to note that the logic gate count does not yield a true measure of actual FPGA utilization. Hardware resources within CLB slices may not be fully utilized by the logic-and-route resources as they relieve routing congestion. This results in an increase in the number of CLB slices with a corresponding increase in logic gates. To achieve a more accurate measure of chip utilization, CLB slice count was chosen as the most reliable area measurement. Therefore, to measure the hardware resource consumption with an implementation's resultant throughput, the Throughput Per Slice (TPS) metric is used. We define TPS as:

$$TPS := (\text{Encryption Rate}) / (\# \text{ CLB slices Used})$$

Therefore, the optimal implementation will display the highest throughput and have the largest TPS. Note that the TPS metric behaves inversely to the classical time-area (TA) metric.

Alg.	Arch.	Thr ugh ut (G it/s)	Slices	TPS
RC6	P-10-2	2.40	108 6	220881
Rijndael	P- -1	1. 4	10 2	1762 7
er ent	PP-32	4.86	004	3 778
Tw fish	P-8-1	1.5	34	16 63

Table 4: AE finalist performance evaluation — n n-fee ack m e s ee - timize imlementati ns

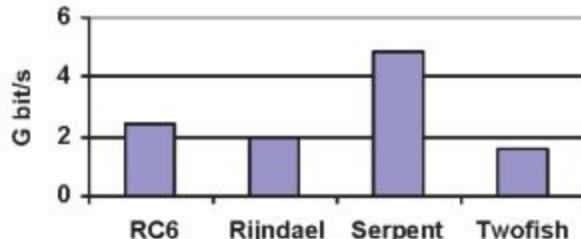


Figure 1: Best throughput — n n-fee ack m e

Alg.	Arch.	Thr ugh ut (M it/s)	Slices	TPS
RC6	PP-2	126.5	318	3 662
Rijndael	LU-2	300.1	302	660
er ent	LU-8	444.2	7 64	771
Twofish	P-1-1	11 .6	30 3	3 16

Table 5: AE finalist performance evaluation — fee ack m e s ee - timize imlementati ns

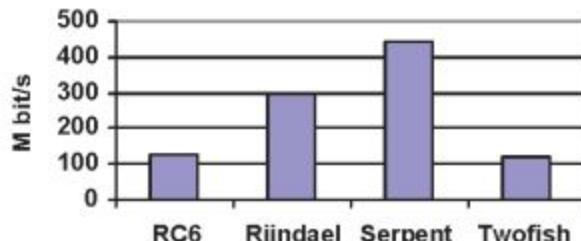


Figure 2: Best throughput — fee ack m e

Tables 4 and 5 detail the timing implementations of the AE finalists in both FB and NFB modes. Additionally, TPS is also shown for each of the implementations. It is critical to note that for the current set of this study, the timing implementation of an AE finalist is fine-tuned to yield the highest throughput. As previously discussed, the synthesis system was specifically optimized for the target architecture to guarantee that the highest throughput would be achieved for each implementation. However, it should be noted that the optimal implementation is not always the fastest. In fact, the optimal implementation is often slower than the fastest implementation. This is because the synthesis system must be programmed with the target system's timing constraints. While an area-efficiency analysis of the AE finalists warrants investigation, it is beyond the scope of this study.

Based on the data shown in Tables 4 and 5, the entitlement algorithm clearly outperforms the other AE finalists in both modes of operation. As can be seen in its nearest competitor, entitlement exhibits a throughput increase of a factor of 2.2 in NFB mode and a factor of 1.1 in FB mode. Interestingly, RC6, Rijndael, and Twofish

all exhibited similar performance results in NFB mode. However, Rijnael exhibits significantly improved performance in FB mode as compared to RC6 and Twofish, although it is still 0% slower than serpent.

None of the main findings from our investigation, namely that serpent appears to be especially well suited for an FPGA implementation for performance reasons, seems especially interesting considering that serpent is clearly not the fastest algorithm with respect to most software candidates [1]. Another major result of our study is that all four algorithms can easily achieve Gigabit encryption rates with standard commercially available FPGAs. The algorithms are at least nearly as fast as magnitudes faster than the best reported software realizations. These software solutions are essentially achieved by parallelizing (parallelizing and unrolling) the structure and using register bypassing (e.g., bypassing of 128 bits in one clock cycle), which are not feasible in current processors. We would like to stress that the pipeline architectures cannot reuse their maximum availability if they are for encryption which require feedback (CFB, FB, etc.) However we believe that for many applications which require high encryption rates, non-feedback modes (including feedback modes such as interleaved CFB [3, section 12]) will be the modes of choice. Note that the Counter Mode grew out of the need for high speed encryption of ATM networks which require parallelization of the encryption algorithm.

7 Conclusions

The importance of the Advanced Encryption Standard and the significance of high throughput implementations of the AE finalists has been examined. A design methodology was established which in turn led to the architectural requirements for a target FPGA. The characteristics of the AE finalists were identified and multiple architecture types were discussed. The implementation of each architecture type for each of the AE finalists was analyzed to determine their suitability for hardware implementation. Based on the implementation results, the best software-optimized implementations were identified for each AE finalist in both non-feedback and feedback modes. Using benchmarks, it was determined that the serpent algorithm yields the best performance in both modes, where serpent performance was defined strictly as the highest throughput. The serpent algorithm outperforms its nearest competitor by a factor of 2.2 in non-feedback mode and by a factor of 1.1 in feedback mode.

8 Acknowledgment

We would like to thank Pawel Chwiecek and Kris Gaj from George Mason University for their helpful discussions on the VHDL code modules that were provided to assist in the implementation of some of the AE finalists. We would also like to thank Alan Martell from the University of Pittsburgh for his public-domain VHDL code module that was used in implementation of the AE finalists.

References

- [1] D. Tuck, *Cryptography, Theory and Practice*. Boca Raton, FL: CRC Press, 1995.
- [2] National Institute of Standards and Technology (NIST), *Second Advanced Encryption Standard (AES) Candidate Algorithm Presentations*, (Rome, Italy), March 1999.
- [3] B. Schneier, *Applied Cryptography*. John Wiley & Sons Inc., 2nd ed., 1996.
- [4] R. Duan, "Hardware Cryptographic Solutions Based VPN," *EETimings*, pp. 7-14, April 1999.
- [5] B. Glaman, "Implementation Experience with AE Candidates Algorithms," in *Proceedings: Second AES Candidates Presentations (AES 2)*, (Rome, Italy), March 1999.