A Scalable and Unified Multiplier Architecture for Finite Fields $GF(p)$ and $GF(2^m)$*

Erkay Savaş, Alexandre F. Tenca, and Çetin K. Koç

Electrical & Computer Engineering
Oregon State University, Corvallis, Oregon 97331
{savas,tenca,koc}@ece.orst.edu

Abstract. We describe a scalable and unified architecture for a Montgomery multiplication module which operates in both types of finite fields $GF(p)$ and $GF(2^m)$. The unified architecture requires only slightly more area than that of the multiplier architecture for the field $GF(p)$. The multiplier is scalable, which means that a fixed-area multiplication module can handle operands of any size, and also, the wordsize can be selected based on the area and performance requirements. We utilize the concurrency in the Montgomery multiplication operation by employing a pipelining design methodology. The upper limit on the precision of the scalable and unified Montgomery multiplier is dictated only by the available memory to store the operands and internal results, and the module is capable of performing infinite-precision Montgomery multiplication in both types of finite fields.
Keywords: Prime fields, binary extension fields, multiplication, Montgomery multiplication, scalability, hardware implementation.

1 Introduction

The basic arithmetic operations (i.e., addition, multiplication, and inversion) in prime and binary extension fields, $GF(p)$ and $GF(2^m)$, have several applications in cryptography, such as decipherment operation of RSA algorithm [17], Diffie-Hellman key exchange algorithm [3], elliptic curve cryptography [7,12], and the Digital Signature Standard including the Elliptic Curve Digital Signature Algorithm [15]. The most important of these three arithmetic operations is the field multiplication operation since it is the core operation in many cryptographic functions.

The Montgomery multiplication algorithm [13] is an efficient method for doing modular multiplication with an odd modulus. The Montgomery multiplication algorithm is very useful for obtaining fast software implementations of the multiplication operation in prime fields $GF(p)$. The algorithm replaces division operation with simple shifts, which are particularly suitable for implementation on both general-purpose computers and application specific hardware.

* Readers should note that Oregon State University filed a patent application containing this work to the US Patent and Trademark Office.

Ç.K. Koç and C. Paar (Eds.): CHES 2000, LNCS 1965, pp. 277–292, 2000.

The Montgomery multiplication operation has been extended to the finite field $GF(2^k)$ in [9]. Efficient software implementations of the multiplication operation in $GF(2^k)$ can be obtained using this algorithm, particularly when the irreducible polynomial generating the field is chosen arbitrarily. The main idea of the architecture proposed in this paper is based on the observation that the Montgomery multiplication algorithm for both fields $GF(p)$ and $GF(2^k)$ are essentially the same algorithm. The proposed unified architecture performs the Montgomery multiplication in the field $GF(p)$ generated by an arbitrary prime p and in the field $GF(2^m)$ generated by an arbitrary irreducible polynomial $p(x)$. We show that a unified multiplier performing the Montgomery multiplication operation in the fields $GF(p)$ and $GF(2^k)$ can be designed at a cost only slightly higher than the multiplier for the field $GF(p)$, providing significant savings when both types of multipliers are needed.

Several variants of the Montgomery multiplication algorithm [16,10,2] have been proposed to obtain more efficient software implementations on specific processors. Various hardware implementations of the Montgomery multiplication algorithm for limited precision operands are also reported [2,16,4]. On the other hand, implementations utilizing high-radix modular multipliers have also been proposed [16,11,18]. Advantages and disadvantages of using high-radix representation have been discussed in [21,20]. Because high-radix Montgomery multiplication designs introduce longer critical paths and more complex circuitry, these designs are less attractive for hardware implementations.

A scalable Montgomery multiplier design methodology for $GF(p)$ was introduced in [20] in order to obtain hardware implementations. This design methodology allows to use a fixed-area modular multiplication circuit for performing multiplication of unlimited precision operands. The design tradeoffs for best performance in a limited chip area were also analyzed in [20]. We use the design approach as in [20] to obtain a scalable hardware module. Furthermore, the scalable multiplier described in this paper is capable of performing multiplication in both types of finite fields $GF(p)$ and $GF(2^k)$, i.e., it is a scalable and unified multiplier.

The main contributions of this paper are summarized below.

- We show that a unified architecture for multiplication module which operates both in $GF(p)$ and $GF(2^m)$ can be designed easily without compromising scalability, time and area efficiency.
- We analyze the design tradeoffs such as the effect of word length, the number of the pipeline stages, and the chip area by supplying implementation results obtained by Mentor graphics synthesis tools.

We start with a short discussion of scalability in §2 and explain the main idea behind the unified multiplier architecture in §3. We then present the methodology to perform the Montgomery multiplication operation in both types of finite fields using the unified architecture. We give the original and modified definitions of Montgomery algorithm for $GF(p)$ and $GF(2^m)$ in §4. We discuss concurrency in the Montgomery multiplication and show the methodology to design a pipeline module utilizing the concurrency in §5. We present the processing unit and the

modifications needed to make the unit operate in prime and binary extension fields in §6. In §7, we discuss the area/time tradeoffs and suitable choices for word lengths, the number of pipeline stages, and typical chip area requirements. Finally, we summarize our conclusions in §8.

2 Scalable Multiplier Architecture

An arithmetic unit is called scalable if it can be reused or replicated in order to generate long-precision results independently of the data path precision for which the unit was originally designed. To speed up the multiplication operation, various dedicated multiplier modules were developed in [18,1,14]. These designs operate over a fixed finite field. For example, the multiplier designed for 155 bits [1] cannot be used for any other field of higher degree. When a need for a multiplication of larger precision arises, a new multiplier must be designed. Another way to avoid redesigning the module is to use software implementations and fixed precision multipliers. However, software implementations are inefficient in utilizing inherent concurrency of the multiplication because of the inconvenient pipeline structure of the microprocessors being used. Furthermore, software implementations on fixed digit multipliers are more complex and require excessive amount of effort in coding. Therefore, a scalable hardware module specifically tailored to take advantage of the concurrency of the Montgomery multiplication algorithm becomes extremely attractive.

3 Unified Multiplier Architecture

Even though prime and binary extension fields, $GF(p)$ and $GF(2^m)$, have dissimilar properties, the elements of either field are represented using almost the same data structures inside the computer. In addition, the algorithms for basic arithmetic operations in both fields have structural similarities allowing a unified module design methodology. For example, the steps of the Montgomery multiplication algorithm for binary extension field $GF(2^m)$ given in [9] only slightly differs from those of the integer Montgomery multiplication algorithm [13,10]. Therefore, a scalable arithmetic module, which can be adjusted to operate in both types of fields, is feasible, provided that this extra functionality does not lead to an excessive increase in area or a dramatic decrease in speed. In addition, designing such a module must require only a small amount of extra effort and no major modification in control logic of the circuit.

Considering the amount of time, money and effort that must be invested in designing a multiplier module or more generally speaking a cryptographic coprocessor, a scalable and unified architecture which can perform arithmetic in two commonly used algebraic fields is definitely beneficial. In this paper, we show the method to design a Montgomery multiplier that can be used for both types of fields following the design methodology presented in [20]. The proposed unified architecture is obtained from the scalable architecture given in [20] after

minor modifications. The propagation time is unaffected and the increase in chip area is insignificant.

4 Montgomery Multiplication

Given two integers A and B, and the prime modulus p, the Montgomery multiplication algorithm computes

$$C = \mathsf{MonMul}(A, B) = A \cdot B \cdot R^{-1} \pmod{p} , \tag{1}$$

where $R = 2^m$ and $A, B < p < R$, and p is an m-bit number. The original algorithm works for any modulus n provided that $\gcd(n, R) = 1$. In this paper, we assume that the modulus is a prime number, thus, we perform multiplication in the field defined by this prime number. This issue is also relevant when the algorithm is defined for the binary extension fields.

The Montgomery multiplication algorithm relies on a different representation of the finite field elements. The field element $A \in GF(p)$ is transformed into another element $\bar{A} \in GF(p)$ using the formula $\bar{A} = A \cdot R \pmod{p}$. The number $\bar{A}$ is called Montgomery image of the element, or $\bar{A}$ is said to be in the Montgomery domain. Given two elements in the Montgomery domain $\bar{A}$ and $\bar{B}$, the Montgomery multiplication computes

$$\bar{C} = \bar{A}\cdot\bar{B}\cdot R^{-1} \pmod{p} = (A\cdot R)\cdot(B\cdot R)\cdot R^{-1} \pmod{p} = C\cdot R \pmod{p} , \tag{2}$$

where $\bar{C}$ is again in the Montgomery domain. The transformation operations between the two domains can also be performed using the MonMul function as

$$\begin{aligned}
\bar{A} &= \mathsf{MonMul}(A, R^2) = A \cdot R^2 \cdot R^{-1} = A \cdot R \pmod{p} , \\
\bar{B} &= \mathsf{MonMul}(B, R^2) = B \cdot R^2 \cdot R^{-1} = B \cdot R \pmod{p} , \\
C &= \mathsf{MonMul}(\bar{C}, 1) = C \cdot R \cdot R^{-1} = C \pmod{p} .
\end{aligned}$$

Provided that $R^2 \pmod{p}$ is precomputed and saved, we need only a single MonMul operation to carry out each of these transformations. However, because of these transformation operations, performing a single modular multiplication using MonMul might not be advantageous. The advantage of the Montgomery multiplication becomes much more apparent in applications requiring multiplication-intensive calculations, e.g., modular exponentiation or elliptic curve point operations. In order to exploit this advantage, all arithmetic operations are performed in the Montgomery domain, including the inversion operation [6,19].

Below, we give bitwise Montgomery multiplication algorithm for obtaining $C := ABR^{-1} \pmod{p}$, where $A = (a_{m-1}, \ldots, a_1, a_0)$ and $C = (c_m, \ldots, c_1, c_0)$.

Input:	$A, B \in GF(p)$ and $m = \lceil \log_2 p \rceil$
Output:	$C \in GF(p)$
Step 1:	$C := 0$

Step 2: for $i = 0$ to $m - 1$
Step 3: $C := C + a_i B$
Step 4: $C := C + c_0 p$
Step 5: $C := C/2$
Step 6: if $C \geq p$ then $C := C - p$
Step 7: return C

In the case of $GF(2^m)$, the definitions and the algorithms are slightly different since we use polynomials of degree at most $m - 1$ with coefficients from the binary field $GF(2)$ to represent the field elements. Given two polynomials

$$A(x) = a_{m-1}x^{m-1} + a_{m-2}x^{m-2} + \cdots + a_1 x + a_0$$
$$B(x) = b_{m-1}x^{m-1} + b_{m-2}x^{m-2} + \cdots + b_1 x + b_0 \ ,$$

and the irreducible monic degree-m polynomial

$$p(x) = x^m + p_{m-1}x^{m-1} + p_{m-2}x^{m-2} + \cdots + p_1 x + p_0$$

generating the field $GF(2^m)$, the Montgomery multiplication of $A(x)$ and $B(x)$ is defined as the field element $C(x)$ which is given as

$$C(x) = A(x) \cdot B(x) \cdot R(x)^{-m} \pmod{p(x)} \ . \tag{3}$$

We note that, as compared to Equation 1, $R(x) = x^m$ replaces $R = 2^m$. The representation of x^m in the computer is exactly the same as the representation of 2^m, i.e., a single 1 followed by 2^m zeros. Furthermore, the elements of $GF(p)$ and $GF(2^m)$ are represented using the same data structures. Only the arithmetic operations acting on the field elements differ. The Montgomery image of a polynomial $A(x)$ is given as $\bar{A}(x) = A(x) \cdot x^m \pmod{p(x)}$. Similarly, before performing Montgomery multiplication, the operands must be transformed into the Montgomery domain and the result must be transformed back. These transformations are accomplished using the precomputed variable $R^2(x) = x^{2m} \pmod{p(x)}$ as follows:

$$\bar{A}(x) = \mathsf{MonMul}(A, R^2) = A(x) \cdot R^2(x) \cdot R^{-1}(x) = A(x) \cdot R(x) \pmod{p(x)} \ ,$$
$$\bar{B}(x) = \mathsf{MonMul}(B, R^2) = B(x) \cdot R^2(x) \cdot R^{-1}(x) = B(x) \cdot R(x) \pmod{p(x)} \ ,$$
$$C(x) = \mathsf{MonMul}(\bar{C}, 1) = C(x) \cdot R(x) \cdot R^{-1}(x) = C(x) \pmod{p(x)} \ .$$

The bit-level Montgomery multiplication algorithm for the field $GF(2^m)$ is given below:

Input: $A(x), B(x) \in GF(2^m)$, $p(x)$, and m
Output: $C(x)$
Step 1: $C(x) := 0$
Step 2: for $i = 0$ to $m - 1$
Step 3: $C(x) := C(x) + a_i B(x)$
Step 4: $C(x) := C(x) + c_0 p(x)$
Step 5: $C(x) := C(x)/x$
Step 6: return $C(x)$

We note that the extra subtraction operation in Step 6 of the previous algorithm is not required in the case of $GF(2^m)$, as proven in [9]. Also, the addition operations are different. While addition in binary field is just bitwise mod 2 addition, the addition in $GF(p)$ requires carry propagation.

Our basic observation is that it is possible to design a unified Montgomery multiplier which can perform multiplication in both types of fields if an adder module, equipped with the property of performing addition with or without carry, is available. The design of an adder with this property is provided in the following sections.

The algorithms presented in this section require that the operations be performed using full precision arithmetic modules, thus, limiting the designs to a fixed degree. In order to design a scalable architecture, we need modules with the scalability property. The scalable algorithms are word-level algorithms, which we give in the following sections.

4.1 The Multiple-Word Montgomery Multiplication Algorithm for $GF(p)$

The use of fixed precision words alleviates the broadcast problem in the circuit implementation. Furthermore, a word-oriented algorithm allows design of a scalable unit. For a modulus of m-bit precision, $e = \lceil m+1/w \rceil$ words (each of which is w bits) are required. Note that one extra bit is used for all the variables in the actual implementation in order to take care of partial sum in the Montgomery algorithm, which can reach $(m+1)$-bit precision. The algorithm proposed in [20] scans the operand B (multiplicand) word-by-word, and the operand A (multiplier) bit-by-bit. The vectors involved in multiplication operations are expressed as

$$\begin{aligned} B &= (B^{(e-1)}, \ldots, B^{(1)}, B^{(0)}) \ , \\ A &= (a_{m-1}, \ldots, a_1, a_0) \ , \\ p &= (p^{(e-1)}, \ldots, p^{(1)}, p^{(0)}) \ , \end{aligned}$$

where the words are marked with superscripts and the bits are marked with subscripts. For example, the ith bit of the kth word of B is represented as $B_i^{(k)}$. A particular range of bits in a vector B from position i to j where $j > i$ is represented as $B_{j..i}$. $(x|y)$ represents the concatenation of two bit sequence. Finally, 0^m stands for an all-zero vector of m bits. The algorithm is given below:

Input: $A, B \in GF(p)$ and p
Output: $C \in GF(p)$
Step 1: $T := 0^{m+1}$
Step 2: for $i = 0$ to $m - 1$
Step 3: $(Carry|T^{(0)}) := a_i \cdot B^{(0)} + T^{(0)}$
Step 4: $Parity := T_0^{(0)}$
Step 5: $(Carry|T^{(0)}) := Parity \cdot p^{(0)} + (Carry|T^{(0)})$
Step 6: for $j = 1$ to $e - 1$

Step 7: $(Carry|T^{(j)}) := a_i \cdot B^{(j)} + Carry + T^{(j)} + Parity * p^{(j)}$
Step 8: $T^{(j-1)} := (T_0^{(j)}|T_{w-1..1}^{(j-1)})$
Step 9: $T^{(e-1)} := (Carry|T_{w-1..1}^{(e-1)})$
Step 10: $C := T$
Step 11: if $C > p$ then $C := C - p$
Step 12: return C

Note that the variable $Carry$ must be capable of accumulating more than one single bit. As suggested in [20], we use the Carry-Save form for the partial sum T, thus $T = (TC, TS)$ where TC and TS are carry and sum part of T, respectively.

4.2 Multiple-Word Montgomery Multiplication Algorithm for $GF(2^m)$

The Montgomery multiplication algorithm for $GF(2^m)$ is given below. Since there is no carry computation in $GF(2^m)$ arithmetic, the intermediate addition operations are replaced by bitwise XOR operations, which are represented below using the symbol $\oplus$.

Input: $A, B \in GF(2^m)$ and $p(x)$
Output: $C \in GF(2^m)$
Step 1: $T := 0^{m+1}$
Step 2: for $i = 0$ to m
Step 3: $T^{(0)} := a_i B^{(0)} \oplus T^{(0)}$
Step 4: $Parity := T_0^{(0)}$
Step 5: $T^{(0)} := Parity \cdot p^{(0)} \oplus T^{(0)}$
Step 6: for $j = 1$ to $e - 1$
Step 7: $T^{(j)} := a_i B^{(j)} \oplus TS^{(j)} \oplus Parity \cdot p^{(j)}$
Step 8: $T^{(j-1)} := (T_0^{(j)}|T_{w-1..1}^{(j-1)})$
Step 9: $T^{(e-1)} := (0|T_{w-1..1}^{(e-1)})$
Step 10: $C := T$
Step 11: return C

Notice that in the outer loop the index i runs from 0 to m. Since $(m + 1)$ bits are required to represent the irreducible polynomial of $GF(2^m)$, we prefer to allocate $(m + 1)$ bits to express the field elements.

5 Concurrency in Montgomery Multiplication

In this section, we analyze the concurrency in Montgomery multiplication algorithms as given in the subsections §4.1 and §4.2. In order to accomplish this task, we need to determine the inherent data dependencies in the algorithm and describe a scheme to allow the Montgomery multiplication to be computed on an array of processing units organized in a pipeline.

We prefer to accomplish concurrent computation of the Montgomery multiplication by exploiting the parallelism among the instructions across the different iterations of i-loop of the algorithms, as proposed in [20]. We scan the multiplier one bit at a time, and after the first words of the intermediate variables (TC, TS) are fully determined, which takes two clock cycles, the computation for the second bit of A can start. In other words, after the inner loop finishes the execution for $j = 0$ and $j = 1$ in ith iteration of the outer loop, the $(i+1)th$ iteration of outer loop starts its execution immediately. The dependency graph shown in Figure 1 illustrates these computations.

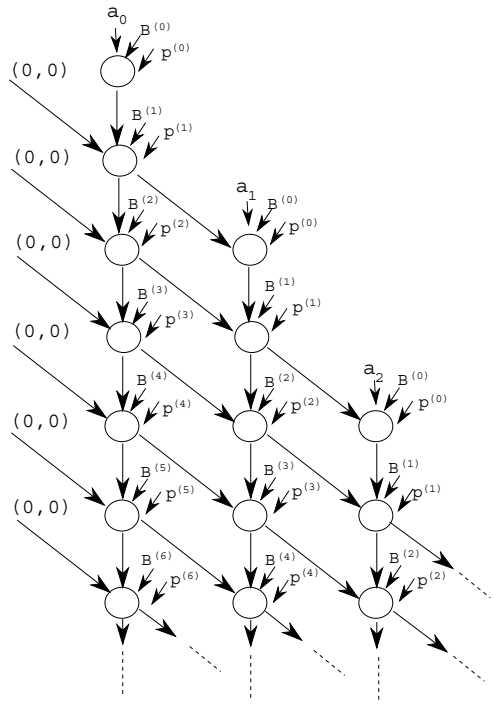

Figure 1: The Dependency Graph of the MonMul Algorithm.

Each circle in the graph represents an elementary computation performed in each iteration of the j-loop. We observe from this graph that these computations are very suitable for pipelining. Each column in the graph represents operations that can be performed by separate processing units (PU) organized as a pipeline. Each PU takes only one bit from multiplier A and operates on each word of multiplicand, B, each cycle. Starting from the second clock cycle, a PU generates one word of partial sum $T = (TC, TS)$ in the Carry-Save form at each cycle, and communicates it to the next PU which adds its contribution to the partial sum, when its turn comes. After $e + 1$ clock cycles, the PU finishes its portion of work, and becomes available for further computation. In case there is no

available PU and there is work to do, the pipeline must stall and wait for the working PUs to finish their jobs. Since the PU at the end of the pipeline has no way of communicating its result to another PU, we need to provide extra buffers for it. In the worst case, which happens when there is only one PU, there must be $2e$ extra buffers of w length to hold these partial sum words. In the last clock cycle of each column, the The PU responsible for this column must receive $p^{(e)} = B^{(e)} = 0$. Elementary computations represented by circles in Figure 1 are performed on the same hardware module. Local control module in the PU must be able to extract $T_0^{(0)}$ and keep this value for the entire operand scanning. Each PU, in other words, has to obtain this value and use it to decide whether to add the modulus p to the partial sum. This value is determined in the first clock cycle of each stage.

An example of the computation for 6-bit operands is shown in Figure 2 for the word size $w = 1$ provided that there are sufficient number of PUs preventing the pipeline to stall. Note that there is a delay of 2 clock cycles between the stage for x_i and the stage for x_{i+1}. The total execution time for the computation takes 20 clock cycles in this example.

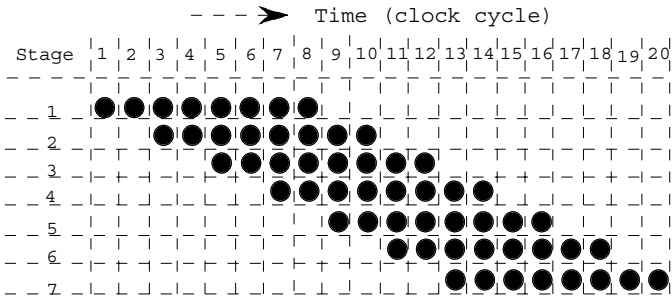

Figure 2: An Example of Pipeline Computation for 6-Bit Operands, where $w = 1$.

If there are at least $\lceil (e+1)/2 \rceil$ PUs in the pipeline organization the pipeline stalls do not take place. The total computation time, CC (clock cycles), is slightly different from the one in [20] and is given as

$$CC = \begin{cases} (\lceil \frac{m+1}{k} \rceil - 1)2k + e + 1 + 2(k-1) & \text{if } (e+1) < 2k \ , \\ (\lceil \frac{m+1}{k} \rceil)(e+1) + 2(k-1) & \text{otherwise} \ , \end{cases}$$

where k is the number of PUs in the pipeline. Notice that the first line of the formula gives the execution time in clock cycles when there are sufficiently many PUs while the second line corresponds to the case when there are stalls in the pipeline.

6 Scalable Architecture

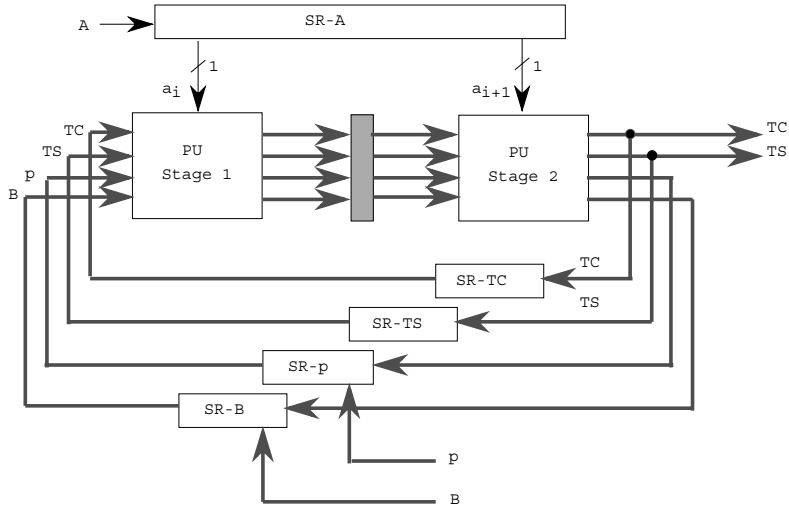

Figure 3: Pipeline Organization with 2 PUs.

An example of pipeline organization with 2 PUs is shown in Figure 3. An important aspect of this organization is the register file design. The bits of multiplier a_i are given serially to the PUs, and are not used again in later stages and can be discarded immediately. Therefore, a simple shift register would be sufficient for the multiplier. The registers for the modulus p and multiplicand B can also be shift registers. When there is no pipeline stall, the latches between PUs forward the modulus and multiplicand to next PU in the pipeline. However, if pipeline stalls occur, the modulus and multiplicand words generated at the end of the pipeline enter the $SR - p$ and $SR - B$ registers. The length of these shift registers are of crucial importance and determined by the number of pipeline stages (k) and the number of words (e) in the modulus. By considering that $SR - p$ and $SR - B$ values require one extra register to store the all-zero word needed for the last clock cycle in every stage (recall that $p^{(e)} = B^{(e)} = 0$) the length of these registers can be given as

$$L_1 = \begin{cases} e - 2 \cdot (k-1) \text{ if } (e+1) > 2k \ , \\ 0 \qquad\qquad\quad \text{otherwise.} \end{cases} \tag{4}$$

The width of the shift registers is equal to w, the wordsize. Once the partial sum (TC, TS) is generated, it is transmitted to the next stage without any delay. However, we need two shift registers, $SR-TC$ and $SR-TS$, to hold the partial sums from the last stage until the job in the first stage is completed. The length (L_2) of the registers TC and TS is equal to L_1.

The registers for TC, TS, B, and p must have loading capability which can complicate the local control circuit by introducing several multiplexers (MUX).

The delay imposed by these MUXes will not create a critical path in the final circuit. The global control block was not mentioned since its function can be inferred from the dependency graph and the algorithms.

6.1 Processing Unit

The processing unit (PU) consists of two layers of adder blocks, which we call *dual-field adders*. A dual-field adder is basically a full adder which is capable of performing addition both with carry and without carry. Addition with carry corresponds to the addition operation in the field $GF(p)$ while addition without carry corresponds to the addition operation in the field $GF(2^m)$. We give the details about the dual-field adder in the next subsection. The block diagram of a processing unit (PU) for $w = 3$ is shown in Figure 4.

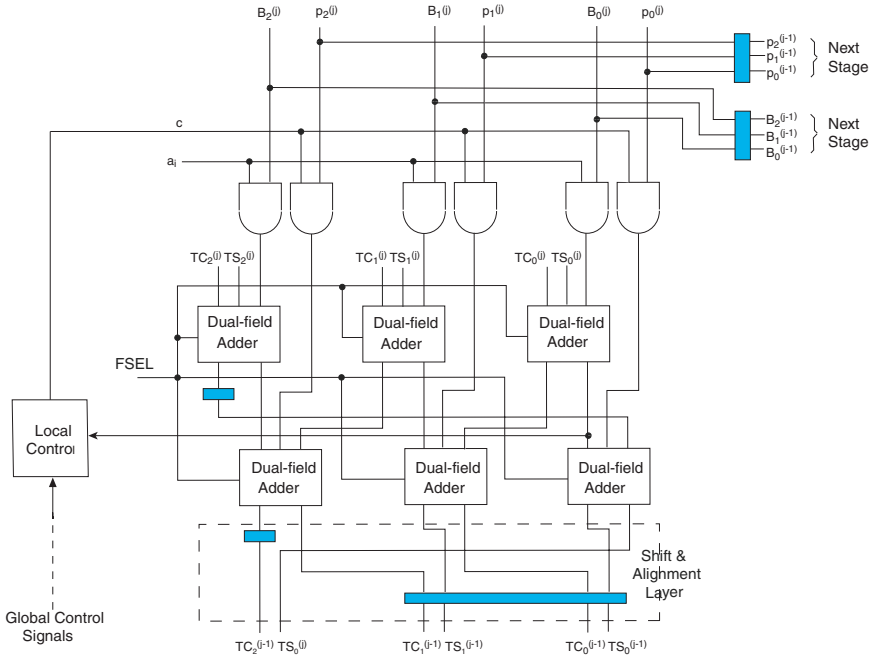

Figure 4: Processing Unit (PU) with $w = 3$.

The unit receives the inputs from the previous stage and/or from the registers $SR - A$, $SR - B$ and $SR - p$, and computes the partial sum words. It delays p and B for the first cycle, then, it transmits them to the next stage along with the first partial sum word (which is ready at the second clock cycle) if there is an available PU. The data path for partial sum $T = (TC, TS)$ (which is expressed in the redundant Carry-Save form) is $2w$ bits long while it is w bits long for p and B and 1 bit long for a_i. At the first cycle, the decision to add the modulus to the partial sum is determined, and this information is kept during the following e clock cycles by the local control. FSEL selects between $GF(p)$ and $GF(2^m)$ fields.

6.2 Dual-Field Adder

The dual-field adder (DFA) shown in Figure 5a, as mentioned before, is basically a full-adder equipped with the capability of doing bit addition both with and without carry. It has an input called $FSEL$ (field select) that enables this functionality. When $FSEL = 1$, the DFA performs the bit-wise addition with carry, which enables the multiplier to do arithmetic in the field $GF(p)$. When $FSEL = 0$, on the other hand, the output $Cout$ is forced to 0 regardless of the values of the inputs. The output S produces the result of bitwise modulo-2 addition of three input values. At most 2 of 3 input values of dual-field adder can have nonzero values while in the $GF(2^m)$ mode.

An important aspect of designing the dual-field adder is not to increase the critical path of the circuit compared to the full-adder, which can have an effect on the clock speed which this would be against our design goal. However, a small amount of extra area can be sacrificed. We show in the following section that this extra area is very insignificant. Figure 5b shows the actual circuit synthesized by Mentor Graphics tools using the $1.2\mu m$ CMOS technology.

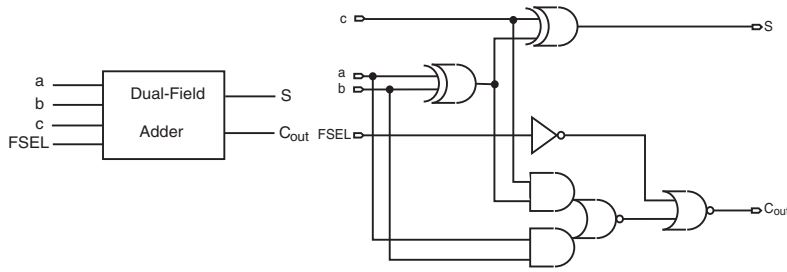

Figure 5: The Dual-Field Adder Circuit.

In the circuit, the two XOR gates are dominant in terms of both area and propagation time. As in the standard full-adder circuit, the dual-field adder has two XOR gates connected serially. Thus, propagation time of the dual-field adder is not larger than that of full adder. Their areas differ slightly.

7 Design Considerations

In [20], an analysis of the area and time tradeoffs is given for the scalable multiplier. The architecture allows designs with different word lengths and different pipeline organizations for varying values of operand precision. In addition, the area can be treated as a design constraint. Thus, one can adjust the design according to the given area, and choose appropriate values for the word length and the number of pipeline stages, in accordance. We give a similar analysis for the scalable and unified architecture. We are targeting two different classes of ranges for operand precision:

- *High precision range* which includes 512, 768 and 1024, is intended for applications requiring the exponentiation operation.
- *Moderate precision range* which includes 160, 192, 224, and 256, is typical for elliptic curve cryptography.

The propagation delay of the PU is independent of the wordsize w when w is relatively small, and thus all comparisons among different designs can be made under the assumption that the clock cycle is the same for all cases. The area consumed by the registers for the partial sum, the operands, and modulus is also the same for all designs, and we are not treating them as parts of the multiplier module.

The proposed scheme yields the worst performance for the case $w = m$ in the high precision range, since some extra cycles are introduced by the PU in order to allow word-serial computation, when compared to other full-precision conventional designs. On the other hand, using many pipeline stages with small wordsize values brings about no advantage after a certain point. Therefore, the performance evaluation reduces into finding an optimum organization for the circuit.

In order to determine the optimum selection for our organization, we obtain implementation results by synthesizing the circuit with Mentor Graphics tools using $1.2\mu m$ CMOS technology. The cell area for a given word size w is obtained as

$$A_{cell}(w) = 48.5w \tag{5}$$

units, and is slightly different from the one found in [20], where the multiplication factor in the formula is the area cost provided by the synthesis tool for a single bit slice. Note that a 2-input NAND gate takes up 0.94 units of area. In the pipelined organization, the area of the inter-stage latches is important, which was measured as

$$A_{latch}(w) = 8.32w \tag{6}$$

units. Thus, the area of a pipeline with k processing elements is given as

$$A_{pipe}(k, w) = (k-1)A_{latch}(w) + kA_{cell}(w) = 56,82kw - 8.32w \tag{7}$$

units. For a given area, we are able to evaluate different organizations and select the most suitable one for our application. The graphs given in Figure 6 allow to make such evaluations for a fixed area of 15,000 gates.

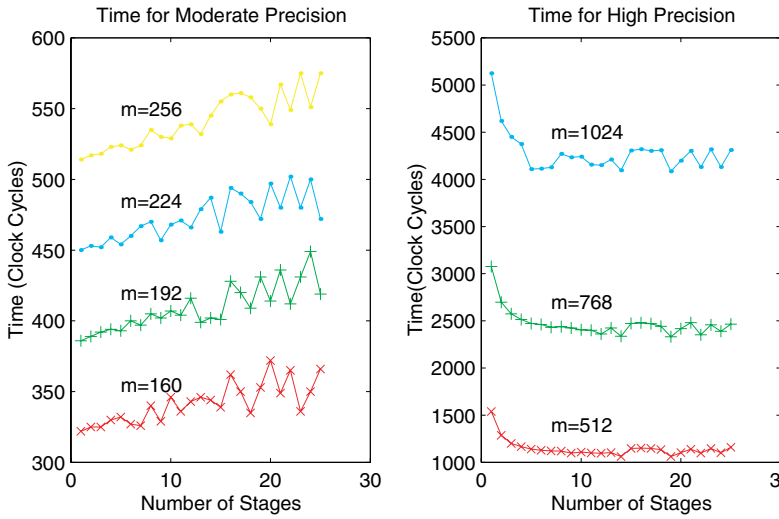

Figure 6: Time Efficiency for Different Configurations with a Fixed Area of 15,000 Gates.

For both moderate and high precision ranges, the number of stages between 5 and 10 are likely to give the best performance. For the high precision cases, fewer than 5 stages yields very poor performance since the fixed area becomes insufficient for large wordsizes and the performance degradation due to pipeline stalls becomes a major problem. The small number of stages with very long word sizes seem to provide a reasonable performance in the moderate range, however, because of the incompatibility issues about using very long word sizes and inefficiency when the precision increases, using fewer than 5 stages is not advised. We avoid using many stages for two reasons:

- high utilization of the PUs will be possible only for very high precision, and
- the execution time may have undesirable oscillations.

The behavior mentioned in the latter category is the result of the facts that

- extra stages at the end of the computations, and
- there is not a good match between the number of words e and the number of stages k, causing a underutilization of stages in the pipeline.

From the synthesis tool we obtained a minimum clock cycle time of 11 nanoseconds, which allows to use a clock frequency of up to 90MHz with $1.2\mu m$ CMOS. Using the CMOS technology with smaller feature size, we can attain much faster clock speeds. It is very important to know how fast this hardware organization really is when comparing it to a software implementation. The answer to this would determine whether it is worth to design a hardware module. In general, it is difficult to compare hardware and software implementations. In order to obtain realistic comparisons, a processor which uses similar clock cycles and technology must be chosen. We selected an ARM microprocessor [5] with 80

MHz clock which has a very simple pipeline. We compare the $GF(p)$ multiplication timing on this processor against that of our hardware module. We use the same clock frequency 80 MHz for the module of the pipeline organization with $w = 32$ and $k = 7$ for the hardware module. On the other hand, the Montgomery multiplication algorithm is written in the ARM assembly language by using all known optimization techniques [8,10]. Table 1 shows the multiplication timings and the speedup.

Table 1: The Execution Times of Hardware and Software Implementations of the $GF(p)$ Multiplication.

precision	Hardware (μs) (80 MHz, $w = 32$, $k = 7$)	Software (μs) (on ARM with Assembly)	speedup
160	4.1	18.3	4.46
192	5.0	25.1	5.02
224	5.9	33.2	5.63
256	6.6	42.3	6.41
1024	61	570	9.34

8 Conclusion

Using the design methodology proposed in [20], we obtained a scalable field multiplier for $GF(p)$ and $GF(2^m)$ in unified hardware module. The methodology can also be used to design separate modules for $GF(p)$ and $GF(2^m)$ which are fast, scalable and area-efficient. The fundamental contribution of this research is to show that it is possible to design a dual-field arithmetic unit without compromising scalability, the time performance and area efficiency. Our analysis shows that a pipeline consisting of several stages is adequate and more efficient than a single unit processing very long words. Working with relatively short words diminishes data paths in the final circuit, reducing the required bandwidth.

The proposed multiplier was synthesized using the Mentor tools, and a circuit capable of working with clock frequencies up to 90 MHz is obtained. Except for the upper limit on the precision which is dictated only by the availability of memory to store the operands and internal results, the module is capable of performing infinite-precision Montgomery multiplication in $GF(2^m)$ and $GF(p)$.

References

1. G. B. Agnew, R. C. Mullin, and S. A. Vanstone. An implementation of elliptic curve cryptosystems over F_{2155}. *IEEE Journal on Selected Areas in Communications*, 11(5):804–813, June 1993.
2. A. Bernal and A. Guyot. Design of a modular multiplier based on Montgomery's algorithm. In *13th Conference on Design of Circuits and Integrated Systems*, pages 680–685, Madrid, Spain, November 17–20 1998.
3. W. Diffie and M. E. Hellman. New directions in cryptography. *IEEE Transactions on Information Theory*, 22:644–654, November 1976.

4. S. E. Eldridge and C. D. Walter. Hardware implementation of Montgomery's modular multiplication algorithm. *IEEE Transactions on Computers*, 42(6):693–699, June 1993.
5. Steve Furber. *ARM System Architecture*. Addison-Wesley, Reading, MA, 1997.
6. B. S. Kaliski Jr. The Montgomery inverse and its applications. *IEEE Transactions on Computers*, 44(8):1064–1065, August 1995.
7. N. Koblitz. Elliptic curve cryptosystems. *Mathematics of Computation*, 48(177):203–209, January 1987.
8. Ç. K. Koç. High-Speed RSA Implementation. Technical Report TR 201, RSA Laboratories, 73 pages, November 1994.
9. Ç. K. Koç and T. Acar. Montgomery multiplication in GF(2^k). *Designs, Codes and Cryptography*, 14(1):57–69, April 1998.
10. Ç. K. Koç, T. Acar, and B. S. Kaliski Jr. Analyzing and comparing Montgomery multiplication algorithms. *IEEE Micro*, 16(3):26–33, June 1996.
11. P. Kornerup. High-radix modular multiplication for cryptosystems. In E. Swartzlander, Jr., M. J. Irwin, and G. Jullien, editors, *Proceedings, 11th Symposium on Computer Arithmetic*, pages 277–283, Windsor, Ontario, June 29 – July 2 1993. IEEE Computer Society Press, Los Alamitos, CA.
12. A. J. Menezes. *Elliptic Curve Public Key Cryptosystems*. Kluwer Academic Publishers, Boston, MA, 1993.
13. P. L. Montgomery. Modular multiplication without trial division. *Mathematics of Computation*, 44(170):519–521, April 1985.
14. D. Naccache and D. M'Raïhi. Cryptographic smart cards. *IEEE Micro*, 16(3):14–24, June 1996.
15. National Institute for Standards and Technology. Digital Signature Standard (DSS). FIPS PUB 186-2, January 2000.
16. H. Orup. Simplifying quotient determination in high-radix modular multiplication. In S. Knowles and W. H. McAllister, editors, *Proceedings, 12th Symposium on Computer Arithmetic*, pages 193–199, Bath, England, July 19–21 1995. IEEE Computer Society Press, Los Alamitos, CA.
17. J.-J. Quisquater and C. Couvreur. Fast decipherment algorithm for RSA publickey cryptosystem. *Electronics Letters*, 18(21):905–907, October 1982.
18. A. Royo, J. Moran, and J. C. Lopez. Design and implementation of a coprocessor for cryptography applications. In *European Design and Test Conference*, pages 213–217, Paris, France, March 17-20 1997.
19. E. Savaş and Ç. K. Koç. The Montgomery modular inverse - revisited. *IEEE Transactions on Computers*, 49(8), July 2000. To appear.
20. A. F. Tenca and Ç. K. Koç. A scalable architecture for Montgomery multiplication. In Ç. K. Koç and C. Paar, editors, *Cryptographic Hardware and Embedded Systems*, Lecture Notes in Computer Science, No. 1717, pages 94–108. Springer, Berlin, Germany, 1999.
21. C. D. Walter. Space/Time trade-offs for higher radix modular multiplication using repeated addition. *IEEE Transactions on Computers*, 46(2):139–141, February 1997.