

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
20 November 2003 (20.11.2003)

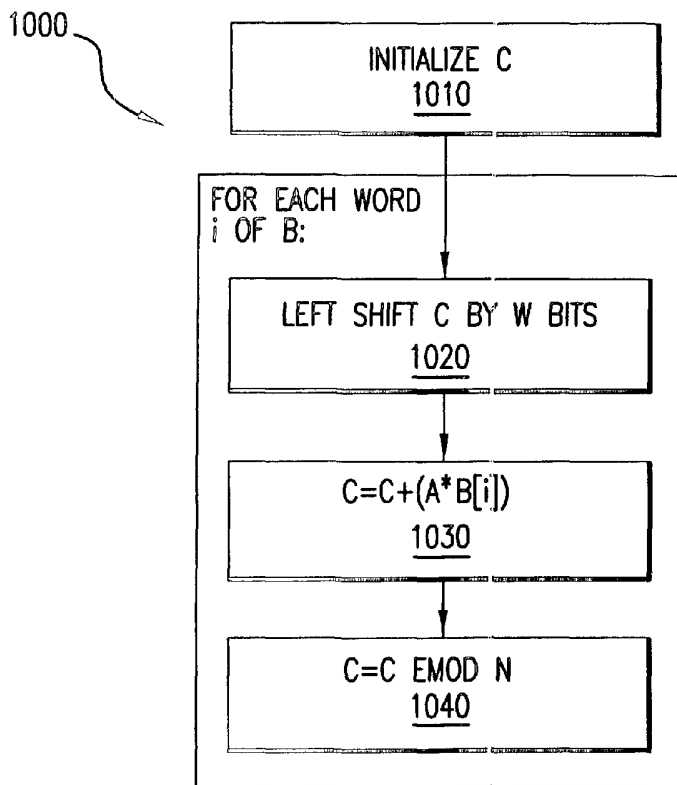
PCT

(10) International Publication Number
WO 03/096182 A1

- (51) International Patent Classification⁷: **G06F 7/72**
 - (21) International Application Number: PCT/US03/13807
 - (22) International Filing Date: 2 May 2003 (02.05.2003)
 - (25) Filing Language: English
 - (26) Publication Language: English
 - (30) Priority Data:
10/140,885 9 May 2002 (09.05.2002) US
 - (71) Applicant: **INTEL CORPORATION** [US/US]; 2200 Mission College Boulevard, Santa Clara, CA 95052 (US).
 - (72) Inventor: **HOEJSTED, Erik**; Bogholder Alle 488, 1 TH, DK-2720 Vanldese (DK).
 - (74) Agents: **O'DOWD, Shawn, W.** et al.; Kenyon & Kenyon, 1500 K Street, N.W., Suite 700, Washington, DC 20005 (US).
 - (81) Designated States (*national*): AE, AG, AL, AM, AT, AU, AZ, BA, BB, BG, BR, BY, BZ, CA, CH, CN, CO, CR, CU, CZ, DE, DK, DM, DZ, EC, EE, ES, FI, GB, GD, GE, GH, GM, HR, HU, ID, IL, IN, IS, JP, KE, KG, KP, KR, KZ, LC, LK, LR, LS, LT, LU, LV, MA, MD, MG, MK, MN, MW, MX, MZ, NI, NO, NZ, OM, PH, PL, PT, RO, RU, SC, SD, SE, SG, SK, SL, TJ, TM, TN, TR, TT, TZ, UA, UG, UZ, VC, VN, YU, ZA, ZM, ZW.
 - (84) Designated States (*regional*): ARIPO patent (GH, GM, KE, LS, MW, MZ, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian patent (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European patent (AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HU, IE, IT, LU, MC, NL, PT, RO, SE, SI, SK, TR), OAPI patent (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).
- Published:**
— with international search report

[Continued on next page]

(54) Title: "EMOD" A FAST MODULUS CALCULATION FOR COMPUTER SYSTEMS



(57) Abstract: An emod operation is a computational substitute for a traditional modulus operation, one that is computationally less expensive but also less precise. Where a modulus operation may be defined for some base number n, the emod operation determines a modulus of an operand using a "phantom modulus," one that is an integer multiple of n. The phantom modulus is chosen to make emod calculations computationally inexpensive when compared to a modulus operation. Thus, the emod operation is particularly useful for multiplications or exponential operations using very large operands. Upon conclusion of interstitial processing associated with the multiplications or exponential operations, a single, traditional modulus operation may be used to obtain a final result.



WO 03/096182 A1



— *before the expiration of the time limit for amending the claims and to be republished in the event of receipt of amendments* *For two-letter codes and other abbreviations, refer to the "Guidance Notes on Codes and Abbreviations" appearing at the beginning of each regular issue of the PCT Gazette.*

"EMOD" A FAST MODULUS CALCULATION FOR COMPUTER SYSTEMS

BACKGROUND

The present invention relates to modulus calculations. In particular, it relates to modulus calculations that may be performed with high degrees of efficiency.

5 A modulus calculation (colloquially, a "mod" calculation) determines the remainder of a division operation. Thus, the expression $A \text{ mod } N$ determines a result that is the remainder obtained by dividing the number A by N . Example: 17 divided by 3 is 5 with a remainder of 2. "17 mod 3" yields a result having value 2.

10 Mod calculations are performed in many computing applications including key negotiation conducted between two parties before engaging in encrypted communication. In the key negotiation context, evaluation of equations having the form $(A^B) \text{ mod } n$ is performed at two terminals. Often, the A and B values may be quite large – from 1024 to 2048 bits long. Of course, when two operands having length l are multiplied, the result may have a length of up to $2l$. With such large operands, it is impractical to build result
15 registers in a processor that have the full width of the multiplication result. Instead, as multiplication results are generated, they typically are truncated by applying the mod calculation to each product. Because the mod calculation divides each product by a modulus having its own length (say, j), the result always has a length less than j .

20 Assuming operands of length l and an equal number of 0s and 1s therein, evaluation of $A^B \text{ mod } n$ may require l multiplications and l mod operations. This involves considerable computational expense. The expense associated with such computations becomes particularly severe in high-load environments such as computer servers where it can be expected that several thousand key negotiation requests (maybe more) would be received per hour.

25 Accordingly, there is a need in the art for a fast, computationally inexpensive technique for resolving mod operations with large operands.

BRIEF DESCRIPTION OF THE DRAWINGS

FIG. 1 illustrates a method according to an embodiment of the present invention.

FIG. 2 illustrates a method according to an embodiment of the present invention.

FIG. 3 illustrates another method according to an embodiment of the present invention.

FIG. 4 is a block diagram illustrating an interstitial product generator according to an embodiment of the present invention.

5 FIG. 5 is a block diagram illustrating an IPG according to an alternate embodiment of the present invention.

FIG. 6 illustrates a multiplier according to an embodiment of the present invention.

FIG. 7 illustrates a multiplier circuit according to an embodiment of the present invention.

10 FIG. 8 illustrates a multiplier circuit according to another embodiment of the present invention.

FIG. 9 illustrates a multiplier circuit according to an embodiment of the present invention.

15 FIG. 10 illustrates a multiplier circuit according to another embodiment of the present invention.

FIG. 11 illustrates another method according to an embodiment of the present invention.

DETAILED DESCRIPTION

Embodiments of the present invention introduce an "emod" operation for use in
20 mod calculations. The emod is a computational substitute for a traditional mod operation, one that is computationally less expensive but also less precise. The emod operation may be used in connection with interstitial multiplications that may be generated during evaluation of an $A^B \bmod n$ calculation or an $(A \cdot B) \bmod n$ calculation. At the end, when a final product is available, a traditional mod operation may be performed to obtain a final
25 result. In this way, the embodiments of the present invention avoid the computational expense of perhaps thousands of mod operations that might otherwise be performed at interstitial stages of operation.

Although computers perform arithmetic operations having binary values (base 2), the advantages of the emod operation might best be understood with an example using

traditional decimal numbers (base 10). To evaluate the operation $23754 \bmod 3331$, it would be conventional to divide 23754 by the modulus 3331 to obtain the remainder 437. However, such division is computationally expensive. It would be far easier to use some multiple of the modulus that is closer to 10^k for some arbitrary number k . Using such a multiplier (say 9993 instead of 3331), one may employ a series of subtractions instead of a division operation. After twice subtracting 9993 from 23754, one is left with a residual of 3768 (which includes the correct remainder 437 plus 3331). This residual is sufficient for use with the interstitial products obtained at intermediate stage of computation. When a final result is obtained and a true mod operation is employed, the correct remainder will be isolated from any multiples of the modulus that may have been carried over from the intermediate stages. Use of this "phantom modulus," however, improves processing speed.

The example illustrated above also works in a binary scheme. In the base 2 domain, for some modulus n , a multiple is chosen that closely approximates some 2^k for some arbitrary k . Just as the decimal example above included a consecutive series of 9s in the most significant bit positions, in the binary example, the phantom modulus will include a consecutive series of 1s in the most significant bit positions. This property simplifies the subtraction that takes place when reducing the operand A by the phantom modulus.

Given a modulus n , a phantom modulus mn may be chosen such that $m \cdot n = 2^k - d$, where $d < n$. Then the emod operation may be employed as a recursive subtraction in which mn is subtracted from the source operand until the residual is less than mn . These two parameters, d and m , control the emod operation.

Evaluation of $(A \cdot B) \bmod n$

According to an embodiment, evaluation of:

$$c = (A \cdot B) \bmod n \quad (1)$$

may be performed by parsing the operand B into multiple words of w bits each. Thus:

$$\begin{aligned} B &= B[M-1]B[M-2] \dots B[0] \\ &= \sum_{i=0}^{M-1} 2^{w \cdot i} B[i] \end{aligned} \quad (2)$$

Equation 1, then becomes:

$$\begin{aligned}
 c &= \left(A \cdot \sum_{i=0}^{M-1} 2^{w \cdot i} B[i] \right) \text{emod } n \\
 &= \left(\dots \left(A \cdot B[M-1] \cdot 2^w + A \cdot B[M-2] \right) \cdot 2^w + \dots \right) \cdot 2^w + A \cdot B[0] \text{emod } n
 \end{aligned}
 \tag{3}$$

The emod operation is distributive and may be replicated within the parenthetical. This property leads to the method illustrated in FIG. 1.

FIG. 1 illustrates a method 1000 according to an embodiment of the present invention. According to the method 1000, a variable *c* may be initialized to be zero (box 1010). Thereafter, the method 1000 iteratively may consider each word of the multiplicand *B*, starting with the word corresponding to the most significant bit position of *B* and working toward the word corresponding to the least significant bit position. During each iteration, the method 1000 may shift left the *c* value from a prior iteration by the length of a word (box 1020). The value of *A* multiplied by the new *B* word (labeled, “*B*[*i*]”) may be added to the shifted value of *c* (box 1030). Thereafter, the emod operation may be performed on the result obtained from box 1030 (box 1040). The result of the emod may be used as the initial value *c* of a subsequent iteration. Following the last iteration, the *c* value obtained is the result of the calculation.

In one embodiment, the method of FIG. 1 may be implemented in software according to the following pseudocode.

```

c := 0
for i := (M-1) downto 0 do
    c := (a • b[i] + shiftleft(w,c)) emod n
    
```

The Basic Multiply/Modulo algorithm

Where *shiftleft(w,c)* merely shifts left the *c* operand by *w* bits. This is equivalent to a multiplication by 2^w in binary.

The emod operator operates based on a phantom modulus $mn = m \cdot n$, yielding a precision factor $d = 2^k - mn$. If the operand *c* were split into two parts, a quotient *q* and a remainder *r*, so that:

$$c = q \cdot 2^k + r \tag{4}$$

then the *emod* operation may be defined as:

$$\begin{aligned}
 (c \text{ emod } n) &= \left((q \cdot 2^k + r) \text{ emod } n \right) \\
 &= \left(q \cdot 2^k + r - q \cdot mn \right) \\
 &= (r + q \cdot d)
 \end{aligned}
 \tag{5}$$

This *emod* function may be integrated into the method of FIG. 1 as shown in the embodiment of FIG. 2.

FIG. 2 illustrates a method 1100 according to an embodiment of the present invention. According to the method 1100, a dummy variable *c* may be initialized to be zero (box 1110). Thereafter, the method iteratively may consider each word of the multiplicand *B*, starting with the word corresponding to the most significant bit position of *B* and working to the word corresponding to the least significant bit position. The method may calculate a quotient value *q* and a remainder value *r* from the value *c* obtained from a prior iteration (boxes 1120, 1130). The quotient *q* may be taken as a span of bits from *c* extending from the most significant bit position to the k^{th} bit position, shifted left by *w* bits. The remainder *r* may be taken as the remaining bits of *c*, extending from the $k-1^{\text{th}}$ bit position to the 0^{th} bit position, shifted left by *w* bits. Thereafter, the *c* value may be evaluated as:

$$c = A \cdot B[i] + r + d \cdot q \quad (6)$$

(box 1140). The *c* value obtained at the last iteration may be taken as the result of the *emod* function.

In one embodiment, the method of FIG. 2 may be implemented in software according to the following pseudocode.

```

c := 0
for i := (M-1) downto 0 do
  q := selectbits(max_bits downto k, c)
  r := selectbits((k-1) downto 0, c)
  c := a * b[i] + shiftleft(w, r) + shiftleft(w, d * q)

```

A Multiply/Modulo algorithm

This implementation requires that the product $d \cdot q$ is available immediately. In practice, since this product may take some time to generate, the method effectively becomes stalled until the product becomes available.

In an alternate embodiment, the method may complete a current iteration without having the $d \cdot q$ product available. Instead, it may advance to the next iteration of *i* and integrate the $d \cdot q$ product from a previous iteration. FIG. 3 illustrates this embodiment.

FIG. 3 illustrates a method 1200 according to an embodiment of the present invention. According to the method 1200, a variable *c* may be initialized to be zero (box 1210). Thereafter, the method iteratively may consider each word of the multiplicand *B*,

starting with the word corresponding to the most significant bit position of B and working to the word corresponding to the least significant bit position. At each iteration, the method 1200 may calculate a quotient value $q[i]$ and a remainder value $r[i]$ from the value c obtained from a prior iteration (boxes 1120, 1130). The quotient $q[i]$ may be taken as a span of bits from c extending from the most significant bit position to the k^{th} bit position, shifted left by w bits. The remainder $r[i]$ may be taken as the remaining bits of c , extending from the $k-1^{\text{th}}$ bit position to the least significant bit position, shifted left by w bits. The quotient $q[i-1]$, obtained from a prior iteration, also may be shifted left by w bits (box 1240). Thereafter, the c value may be evaluated as:

$$c = A \cdot B[i] + r[i] + d \cdot q[i-1], \quad (7)$$

where the $q[i-1]$ value is the shifted value obtained in box 1240 (box 1250).

Following the final iteration, the quotient from the final iteration may be added to c (box 1260). The value obtained from this operation may be taken as the result from the emod operation.

In one embodiment, the method of FIG. 3 may be implemented in software according to the following pseudocode.

```

c := 0
for i := (M-1) downto 0 do
  q1 := q
  q := selectbits(max_bits downto k, c)
  r := selectbits((k-1) downto 0, c)
  c := a • b[i] + shiftright(w, r) + shiftright(2•w, d • q1)
q1 := q
c := shiftright(w, c) + shiftright(2•w, d • q1)
c := shiftright(w, c)

```

The Multiply/Modulo algorithm with one cycle q-delay

In this embodiment, the $d \cdot q$ product from a prior iteration (reabeled as $d \cdot q1$) is shifted left to account for positional differences between the two words.

As noted above, the embodiment of FIG. 3 need not be stalled while waiting for evaluation of a $d \cdot q$ operation. This embodiment may find application in a high-load application where avoidance of computational latencies may be at a premium.

Multiplication of Large Numbers with Small Numbers

As described above, the multiplicand B may be parsed into a plurality of smaller words $B[w]$, $w=0$ to $M-1$, and the words may be used as a basis on which to perform the multiplication with the multiplier A. A discussion of a circuit implementation for this embodiment follows.

FIG. 4 is a block diagram illustrating an interstitial product generator ("IPG") 100 according to an embodiment of the present invention. The IPG 100 generates an interstitial product from a multiplicand A. It may include a multiplicand register (called, an "A register" herein) 110, a pair of shifters 120, 130 (labeled "shift 1" and "shift 2" respectively) and a "3A register" 140. The A and 3A registers are illustrated in phantom because they may (but need not) be placed within the IPG 100 itself; alternatively, they may be provided in some other circuit but their contents may be provided as an input to the IPG 100. The IPG 100 also may include a pair of multiplexers (colloquially, "MUXes") 150, 160 and an inverter 170.

The shifters 120, 130 each present values representing the value stored in the A register shifted by a predetermined number of bit positions. The first shifter 120 may present an A value shifted one bit position toward the most significant bit position. It is labeled "shift 1." The second shifter 130 may present the A value having been shifted two bit positions toward the most significant bit position, labeled "shift 2." In binary data systems, a single or double bit shift causes a two-fold or four-fold multiplication of a source data value respectively.

The shifters 120, 130 may be provided as any number of embodiments. Perhaps the simplest embodiment is to provide a shifter as a wired interconnection between the A register 110 and the MUX 150. For example, each bit position i in the A register 110 may be connected to a position $i+1$ of the MUX 150 to constitute the "shift 1" shifter 120. Similarly, each bit position i in the A register may be connected to a position $i+2$ of the MUX 150 to satisfy the "shift 2" shifter 130. The least significant bit position of the shift 1 shifter may be grounded. The two least significant bit position of the shift 2 shift inputs to the MUX 150 also may be grounded. This architecture provides the desired shift functions with the least implementation cost in terms of area or control hierarchy.

Alternatively, the shifters 120, 130 could be provisioned as formal shift registers complete with storage cells (not shown) for storage of shifted values. Although there is no performance advantage for this alternative, it may find use in applications where such shift registers are employed for other purposes.

5 The 3A register 140, as its name implies, is a register to store a value representing three times the value in the A register. The values stored in this register may be obtain from a straightforward addition of the values from the A register 110 and the shift 1 register 120 or, alternatively, from a subtraction of the values in the shift 1 register 120 from the value in the shift 2 register 140. Circuitry to implement these
10 functions is straightforward and omitted from the illustration of FIG. 1 to keep the illustration simple. In an embodiment, the 3A register 140 also may be provided in a location external to the IPG 100; it is illustrated in phantom accordingly.

 Outputs from the A register 110, the two shifters 120 and 130 and the 3A register 140 may be input to the first MUX 150. An output from the first MUX 150 may be input to
15 both the second MUX 160 and the inverter 170. An output from the inverter 170 may be provided as a second input to the second MUX 160. The inverter 170 may generate a two's complement inversion of the multibit output from the first MUX 150. The second MUX 160 may have a third input coupled directly to a zero value "Ø." Alternatively, the zero value could be input to the first MUX 150. Thus, given an input value A, the IPG 100
20 may generate any of the following outputs: A, \bar{A} , 2A, $\bar{2A}$ 3A, $\bar{3A}$, 4A, $\bar{4A}$ and Ø.

 The IPG 100 may include a controller 180 that governs operation of the two MUXes 150, 160. As discussed below, given an input "segment," the controller 180 may generate a control signal (labeled c_i) that causes the MUXes 150, 160 to output a selected one of the possible outputs on each cycle of a driving clock (not shown).

25 FIG. 5 is a block diagram illustrating an IPG 200 according to an alternate embodiment of the present invention. According to an embodiment, the IPG 200 may include a plurality of inverters 210, 220, 230-1, 230-2, a 3X multiplier 240, a pair of shifters 250-1, 250-2 and a multiplexer 260. In this embodiment, the IPG 200 is illustrated as connected to an external multiplicand register rather than including the
30 multiplicand register as part of it. The multiplicand may be input to the IPG 200 on a first terminal 270 thereof. One of the inverters 210 may be coupled to the first terminal 270, to invert a multiplicand when presented.

The 3X multiplier, as its name implies, may generate a value that is three times a multiplicand when presented at the input terminal. A second inverter 220 may be coupled to the 3x multiplier 240 to invert the output thereof.

5 The shifters 250-1, 250-2 provided shifted versions of the multiplicand as in the embodiment of FIG. 2. One of them (say, shifter 250-1) shifts the input multiplicand by a single bit position; the other 250-2 shifts the multiplicand by two bit positions. Inverters 230-1, 230-2 from the respective shifters 250-1, 250-2 may generate inverted shifted values of the multiplicand. The shifters 250-1, 250-2 may be provided according to any of the embodiments described above.

10 Outputs from the inverters 210, 220, 230-1, 230-1, the 3X multiplier 240 and the shifters 250-1, 250-2 may be input to the multiplexer 260. The multiplexer 260 also may be controlled to output none of the inputs from the IPG 200. In this condition, the multiplexer 260 causes the IPG 200 to generate a zero output therefrom.

15 According to an embodiment, when it is desired to perform a multiplication based on a long multiplicand A and a shorter multiplier B, the multiplicand A may be input to the IPG 200. Values of A, 3A, \bar{A} and $\overline{3A}$ will be available to the multiplexer 260 after a short initialization period. Similarly, shifted values of A and \bar{A} also will be available to the multiplexer 260. Once these values are available, they may be retrieved from the IPG and forwarded to a remainder of a multiplication circuit (not shown in FIG. 5) based on
20 values of multiplier segments.

The IPG may include a controller 290 responsive to these multiplier segments to cause the multiplexer 260 to retrieve one of the previously stored values from within the IPG. As is known, many multiplication circuits include controllers for other purposes. The controller 290 may be integrated into these known controllers or may be provided as a
25 separate element as may be desired. FIG. 5 illustrates the controller 290 as being separate from the IPG 200 for convenience only.

In an embodiment, the multiplier B may be parsed into several four bit segments s_i . Each segment s_i includes the bits B_{3i+2} - B_{3i-1} from the multiplier B. From these segments, a control signal CTRL may be generated to determine which value from within
30 the IPG should be output from the multiplexer. In an embodiment, the IPG may generate outputs according to the scheme shown in Table 1 below.

Input Pattern of Segment s_i	CTRL	Interstitial Product
0000	0	0
0001	1	A
0010	1	A
0011	2	$A \ll 1$
0100	2	$A \ll 1$
0101	3	3A
0110	3	3A
0111	4	$A \ll 2$
1000	-4	$\overline{A} \ll 2$
1001	-3	$\overline{3A}$
1010	-3	$\overline{3A}$
1011	-2	$\overline{A} \ll 1$
1100	-2	$\overline{A} \ll 1$
1101	-1	\overline{A}
1110	-1	\overline{A}
1111	0	0

Table 1

Where \overline{A} is the two's complement of an input A. The control value CTRL may be related to the four bits input pattern by:

$$5 \quad \text{CTRL} = -4s_{i3} + 2s_{i2} + s_{i1} + s_{i0} \quad (8)$$

where s_{ij} represents the j^{th} bit position of segment s_i . For the segment s_0 , the zeroth bit position, which would be a fictional bit position "B₋₁," may be set to 0 to render the response of the controller to segment s_0 consistent with Table 1.

10 It may be observed from FIG. 6 that a multiplier B will not fill all segments completely unless the length of the multiplier is a multiple of 3. In an embodiment, when the length of the multiplier is not a multiple of 3, it may be sign extended in sufficient length to fill an otherwise unused portion of the last segment. This typically involves copying the sign bit, the most significant bit B_{MSB} , to fill fictional bit positions beyond the most significant bit.

15 The IPG embodiments described above may be applied to multiplier circuits of a variety of architectures. In each application, use of an IPG permits the multiplier circuit to

achieve faster operation essentially by permitting additions to occur once every three bit positions rather than once per bit position as is conventional.

FIG. 7 is a diagram of an application of an IPG in a combinatorial multiplier 300 according to an embodiment of the present invention. The combinatorial multiplier 300 may include first and second registers 310 for storage of a multiplicand A and a multiplier B respectively. It may include an IPG 330, a controller 340, a plurality of interstitial product registers 350.1-350.L, an adder tree 360 a product register 370.

During operation, the IPG 330 may be initialized to create the values A, \bar{A} , $3A$ and $3\bar{A}$ and shifted values of A and \bar{A} . The controller 340 may parse the multiplier B into segments and, responsive to the bit pattern in each segment, cause the IPG 330 to load one of the values into a corresponding interstitial product (say, 350.1). In an embodiment, the number of interstitial products 350.1-350.L may be based upon the number of segments supported by the multiplier B. Thus, the number of interstitial products 350.1-350.L may be tied to the length of the multiplier register 320.

Once values are loaded in each of the interstitial product registers 350.1-350.L, the combinatorial multiplier 300 may cause a final product to be stored in the product register 370 by summing across all the interstitial product registers 350.1-350.L. An adder tree 360 accepts the interstitial product values from each product registers and sums them in a manner that is cognizant of the respective bit offsets among the interstitial product registers. Adder trees are well known per se and may be modified for this purpose.

In an embodiment for a multiplicand A of length m and a multiplier B of length n, the product register 370 may have a length n+m just as in the traditional combinatorial multiplier. Interstitial product registers 350.1-350.L may have a length m+2 whereas in traditional combinatorial multipliers, they would have a length m.

Traditional combinatorial multipliers include one interstitial product register for each bit position of a multiplier B. In the foregoing embodiment, there need be only one interstitial product register (say, 350.1) for every three bit positions of the multiplier B. Thus, because the combinatorial multiplier 300 of the foregoing embodiments include approximately one-third the number of interstitial registers than conventional counterparts, the adder tree include one-third the number of adders to compute the final summation. Products are expected to be generated faster in the present embodiment

because the final addition is carried out across fewer interstitial values and, therefore, is available more quickly. This embodiment, therefore, yields higher throughput with less logic.

FIG. 8 is a block diagram of an application of an IPG integrated with a shift-add multiplier circuit 400 according to an embodiment of the present invention. The multiplier circuit 400 may include a pair of registers 410, 420 for storage of the multiplicand A and the multiplier B. The shift-add multiplier circuit 400 also may include an IPG 430, a controller, a carry save adder 450 and a product register 460.

During operation, the multiplier circuit 400 may be initialized. In this embodiment, the product register 460 may be cleared to zero and the interstitial product generator 430 may be initialized with the value of the multiplicand A. Thereafter, during operation, the controller 440 may shift each segment out of the multiplier register 420 and, responsive to the new segment, may cause a selected value to be output from the IPG 430 to a first input of the carry save adder 450. The most significant bits from the product register 460 may be shifted 3 places and input to a second input of the carry save adder. The carry save adder 450 may add the values presented on each of its two inputs together and write the value back to the product register 460. This process may repeat in an incremental fashion for as many segments as are supported by the multiplier B.

As in the embodiment of FIG. 4, the embodiment of FIG. 8 provides for improved performance over traditional shift-add multiplier circuits. The traditional circuits perform an addition for each bit position of a multiplier B. By contrast, the embodiment shown in FIG. 8, provides an addition only once for every three bit positions of the multiplier B. Again, a fewer number of additions permit the shift-add multiplier of FIG. 8 to generate multiplication products in a shorter amount of time than would be available from traditional circuits.

In another embodiment, a multiplier circuit may omit use of a multiplier register (such as the multiplier 420 of FIG. 8). FIG. 9 illustrates a multiplier 500 populated by a multiplicand register 510, an IPG 520, a controller 530, a carry save adder 540 and a product register 550. A multiplicand value (A) may be input to the IPG 520 as an initialization step. Additionally, a multiplier value (B) may be loaded into the least significant bit positions of the product register 550.

On each clock cycle, the contents of the product register 550 may be downshifted by three bit positions. When the least significant bits of the product register are shifted out of the product register, they may be input to the controller 530 as a new segment. In response to these three bits (and one bit from the shift of a prior clock cycle), the controller 530 may cause the IPG 520 to generate an output as shown above in Table 1. The IPG output may be provided to a first input of the carry save adder 540. The downshifted value from the product register 550 may be provided to a second input of the carry save adder 540. The carry save adder 540 may add the two input values and store them to the product register in the most significant bit positions. Again, this embodiment provides improved performance over other shift-add multipliers that operate a single bit shift at a time. As compared to the embodiment of FIG. 4, this embodiment also provides better register utilization because the register 420 (FIG. 4) may be omitted.

FIG. 10 illustrates a multiplication circuit 600 according to another embodiment of the present invention. This embodiment may be used to evaluate equation 7 above. In this embodiment, the multiplier 600 may include a pair of IPGs 610, 620. A first IPG 610 may evaluate the $A \cdot B[i]$ term of equation of equation 7, the second IPG 620 may evaluate the $d \cdot q[i-1]$ term. The first IPG 610 may accept as inputs values of A and 3A from respective registers 630, 640. The first IPG 610 may be controlled by a controller 650 which in turn may be controlled by the respective word of B (input not shown).

The second IPG 620 may accept values of d and 3d from respective sources 660, 670. The second IPG 620 may be controlled by a q value 680.

The multiplication circuit 600 may include a carry save adder 690 that accepts, as inputs, outputs from the two IPGs 610, 620 and a result value from a product register 700. An output of the carry save adder 690 may be input to the product register 700. The product register 700 may be provided as a shift register to implement the shifts described in connection with FIG. 3. A span of bits from the least significant bit positions may output to the carry save adder 690 as the $r[i]$ value.

As is known, a carry save adder generates a result in the so-called "redundant form." Carry save adders are faster than other types of adders because they generate results of addition operations without performing a traditional carry propagation (a time consuming operation). Instead, the addition results are stored using multiple bits per "bit position." Multiple additions can be performed in redundant form. After a final addition, a single carry propagation may be performed to obtain a result in non-redundant form.

According to an embodiment, a portion of the product register 700 corresponding to the quotient q may be input to a second adder 710 such as a carry lookahead adder. The second adder 710 may generate a non-redundant result which may be fed back to the second IPG 610 as the q value.

5 The foregoing embodiments have been presented in connection with an evaluation of a $(A \cdot B) \bmod n$ operation. If done in connection with an $(A \cdot B) \bmod n$ operation, the methods illustrated in FIGS. 1-3 may be followed by a single mod operation to complete the evaluation. These embodiments may be used advantageously, however, for the evaluation of other operations, such as an $A^B \bmod n$ operation, in which case the
 10 mod n operation may be deferred.

Evaluation of $A^B \bmod n$

Embodiments of the emod operation are described below in the context of an $A^B \bmod n$ calculation. Resolution of the calculation may proceed as a nested loop process, constituting an outer loop and an inner loop.

15 The outer loop may scan the bits of the exponent and to control the multiplications. Each pass of the outer loop may include squaring operation:

$$c = (c \cdot c) \bmod n, \quad (9)$$

and, depending on the current bit in the exponent, an additional multiplication:

$$20 \quad c = (c \cdot a) \bmod n, \quad (10)$$

This embodiment is illustrated in FIG. 11.

FIG. 11 illustrates another method 1300 according to an embodiment of the present invention. According to the method 1200, a dummy variable c may be initialized to be zero (box 1210). Thereafter, the method iteratively considers each bit i of the
 25 exponent B . During each iteration, the method may evaluate equation 9 using an inner loop (box 1320). Thereafter, the method 1300 may determine whether the i^{th} bit of the exponent B ($B[i]$) is 1. If so, the method 1300 may evaluate equation 10 using an inner loop (box 1330). When the operation of box 1330 concludes or if the i^{th} bit was not a 1, the operation may advance to a next successive iteration.

30 Upon conclusion of the last iteration, the method 1300 may invoke a traditional mod operation upon the value c (box 1340). This result obtains the result of $A^B \bmod n$.

The inner loop of the operation may be performed using any of the methods described hereinabove with respect to FIGS. 1-3. During processing of the inner loop, both cases may be implemented using the same hardware. There need be no special handling to accommodate the squaring case. Typically the operands c and mn are of the same size and a is somewhat smaller.

Determining d and m.

To determine the parameters d and m, the number of significant bits in n, size(n), may be determined so that:

$$2^{\text{size}(n)-1} \leq n < 2^{\text{size}(n)}. \tag{11}$$

Also, k may be defined as:

$$k = \text{size}(n) + \text{precision} \tag{12}$$

for some desired precision value. Further, mn may be defined to be m•n, an integer multiple of n, so that:

$$d = 2^k - mn < n \tag{13}$$

The parameter mn, when viewed in binary format, has a number of leading 1s equaling or exceeding the precision value.

To find m and d from n and 2k, note that:

$$mn = m \cdot n = \text{int}\left(\frac{2^k}{n}\right) \cdot n \tag{14}$$

where int(x) finds the largest integer of some value x. To find m, one may use the

Newton method:

$$u_{i+1} = (u_i \cdot (2^{k+1} - u_i \cdot n)) / 2^k. \tag{15}$$

Equation 15 has a property that:

$$u_i \Rightarrow (2^k / n) \text{ as } i \Rightarrow \infty \tag{16}$$

As it is desirable to find only the integer part of (2^k/n), this iteration converges quite fast.

To keep the number of bits in u_i to a minimum, it is possible also to take the integer part after each iteration and throw away the fractional part:

$$u_{i+1} = (u_i \cdot (2^{k+1} - u_i \cdot n)) / 2^k \tag{17}$$

This implementation at times can lead to a problem where the iteration stalls before reaching the correct number. Modifying the final iteration slightly however cures the problem.

$$u_{\text{Final}} = ((u_{\text{Final-1}}+1) \cdot (2^{k+1} - u_{\text{Final-1}} \cdot n) - 2^k) \text{ div } 2^k$$

In another embodiment, the calculation of $(2^{k+1} - u_i \cdot n)$ where $(u_i \cdot n) < 2^{k+1}$ can be done simply by setting the leading ones of the negative number $-(u_i \cdot n)$ to zero from position $(k+1)$ and up. That is equivalent to keeping the bits from position k down to 0.

- 5 In one embodiment, the calculation of m and d may be implemented in software according to the following pseudocode.

```

u := 1.5 • 2^(precision-1)
for i := 1 to (final_iteration-1) do
  x := -n•u
  x := selectbits(k downto 0, x)           // 2^(k+1) - n•u
  x := u•x
  u := selectbits(max_bits downto k, x)    // (...) div 2^k
x := -n•u
x := selectbits(k downto 0, x)
u := u + 1
x := u•x
x := x - 2^k
m := selectbits(max_bits downto k, x)     // (...) div 2^k
x := -n•m
d := selectbits((k-1) downto 0, x)        // 2^k - m•n

```

An algorithm for calculating $m = (2^k \text{ div } n)$ and $d = (2^k - mn)$

where the function `selectbits(max downto min, c)` returns the bits in c from and including positions max and down to min .

- 10 The embodiments presented hereinabove provide a computational substitute for a mod operation, labeled "emod," that incurs much less computational expense at the cost of lost precision. It is useful when performing a mod operation in connection with multiplications or exponential operations. By applying an emod operation to interstitial products, the length operands may be maintained to be within some predetermined
- 15 length window. When a final product is obtained, a traditional mod operation may be applied to obtain a final result. This scheme obtains a final result with much less processing than would be possible using only the mod operation.

Several embodiments of the present invention are specifically illustrated and described herein. However, it will be appreciated that modifications and variations of the

present invention are covered by the above teachings and within the purview of the appended claims without departing from the spirit and intended scope of the invention.

WE CLAIM:

1. A method of performing a modulus operation upon an operand that is represented by mathematical function to be performed iteratively, the modulus taken with respect to a value n , comprising:
 - 5 upon conclusion of each of a plurality of interstitial iterations, determining a modulus of a result thereof using a phantom modulus that is a multiple of n , the phantom modulus approximating 2^k for some arbitrary k , and
 - upon conclusion of a final iteration, determining a modulus of a result thereof using a true modulus n .

- 10 2. A method of resolving an $(A \cdot B) \bmod n$ operation, comprising,
 - iteratively, for each of several words of B :
 - shifting a value c by a word length,
 - adding to the shifted value c , a value $A \cdot B[i]$, where $B[i]$ is the i th word of B ,
 - performing an emod operation on a result of the addition, using a phantom
 - 15 modulus mn that is a multiple of n , and
 - following a last iteration, performing a modulus operation on the result of the addition from the last iteration.

3. A method of resolving an $(A \cdot B) \bmod n$, operation comprising:
 - iteratively, for each of several words of B :
 - 20 generating a quotient q from a higher order portion of an earlier valuation of c ,
 - generating a remainder r from a lower order portion of the value c ,
 - revaluating c to be $A \cdot B[i] + r + d \cdot q$, where d is a precision value of a phantom modulus, where $B[i]$ is the i th word of B ,
 - 25 following a last iteration, performing a modulus operation on the result of the revaluating from the last iteration.

4. A method of resolving an $(A \cdot B) \bmod n$ operation, comprising:
 - iteratively, for each of several words of B :
 - generating a quotient $q[i]$ from a higher order portion of an earlier valuation
 - 30 of c , where i represents the current iteration,
 - generating a remainder $r[i]$ from a lower order portion of the value c ,
 - shifting left a quotient $q[i-1]$ by a word width

revaluing c to be $A \cdot B[i] + r[i] + d \cdot q[i-1]$, where d is a precision value of a phantom modulus, where $B[i]$ is the i th word of B and the revaluing uses the shifted quotient,

5 following a last iteration, performing a modulus operation on the result of the revaluing from the last iteration.

5. A method of resolving an $A^B \bmod n$ operation, comprising:

iteratively, for each bit position i of B :

performing a $c = (c \cdot c) \bmod n$ operation, and

if the i th bit position of B is a 1, performing a $c = (c \cdot A) \bmod n$ operation;

10 and

following a last iteration, performing $A \bmod N$ operation of c obtained from the last operation,

wherein the $\bmod n$ operation calculates a modulus result from a phantom modulus mn , where mn is an integer multiple of n .




FIG. 1

FIG. 2

2/8

1200




FIG.3




FIG. 4




FIG.5




FIG.6




FIG. 7

6/8




FIG. 8




FIG. 9




FIG.10

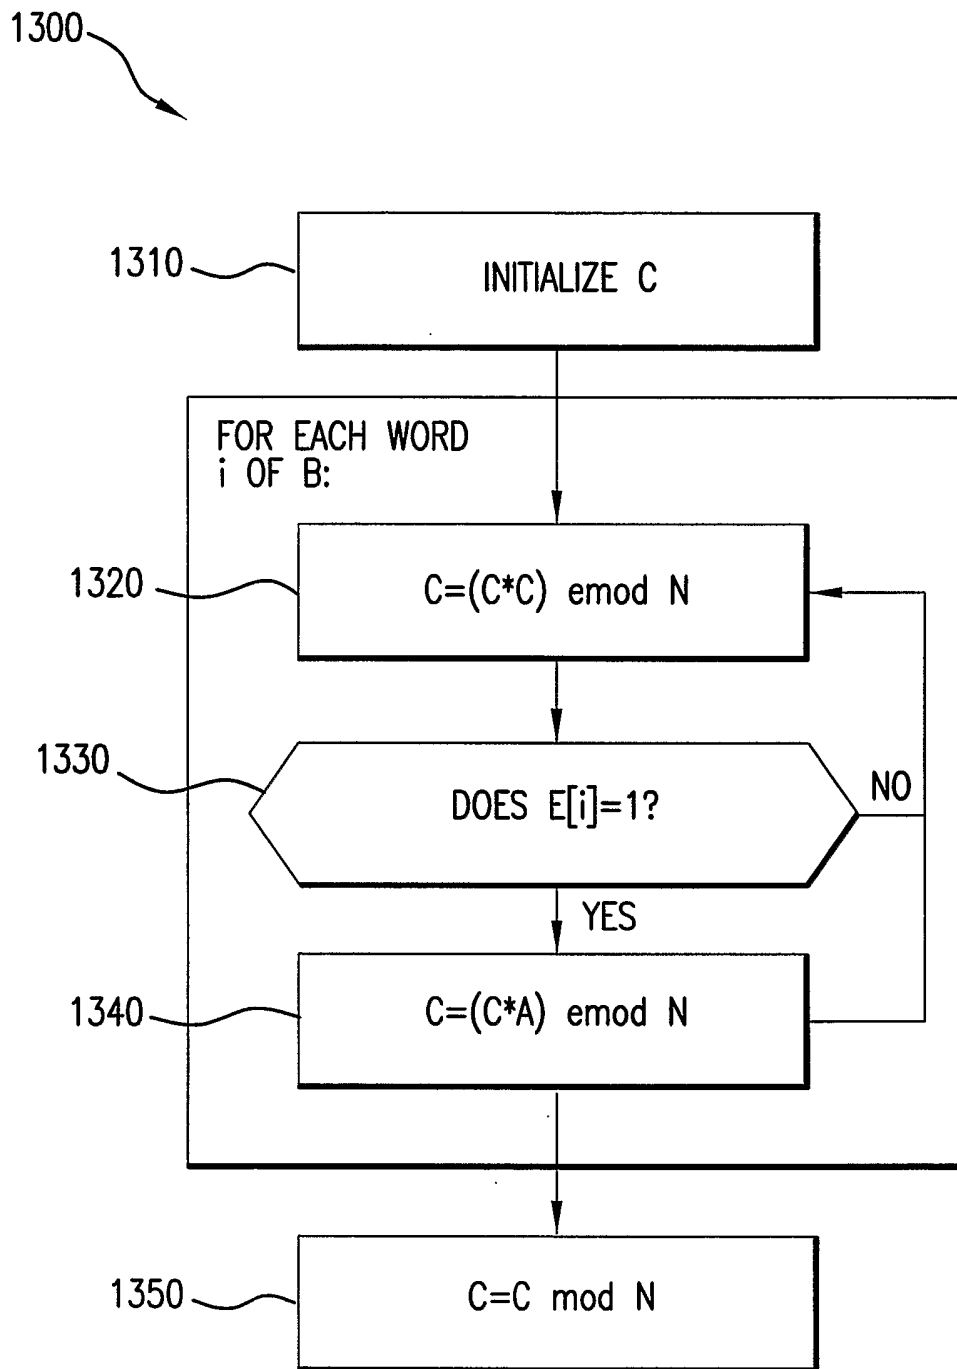


FIG.11

INTERNATIONAL SEARCH REPORT

International Application No

PCT/US 03/13807

A. CLASSIFICATION OF SUBJECT MATTER

IPC 7 G06F7/72

According to International Patent Classification (IPC) or to both national classification and IPC

B. FIELDS SEARCHED

Minimum documentation searched (classification system followed by classification symbols)

IPC 7 G06F

Documentation searched other than minimum documentation to the extent that such documents are included in the fields searched

Electronic data base consulted during the international search (name of data base and, where practical, search terms used)

EPO-Internal, WPI Data, PAJ

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category °	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
X	US 5 166 978 A (QUISQUATER JEAN-JACQUES) 24 November 1992 (1992-11-24)	3,4
A	column 2, line 31 -column 9, line 2; figures 1-3	1,2,5
A	---- MAYERWIESER W ET AL: "TESTING A HIGH-SPEED DATA PATH THE DESIGN OF THE RSABETA CRYPTO CHIP" JOURNAL OF UNIVERSAL COMPUTER SCIENCE, SPRINGER VERLAG, DE, vol. 1, no. 11, 1995, pages 728-743, XP001111443 paragraph '0002! ---- -/--	1-5

 Further documents are listed in the continuation of box C. Patent family members are listed in annex.

° Special categories of cited documents:

A document defining the general state of the art which is not considered to be of particular relevance

E earlier document but published on or after the international filing date

L document which may throw doubts on priority claim(s) or which is cited to establish the publication date of another citation or other special reason (as specified)

O document referring to an oral disclosure, use, exhibition or other means

P document published prior to the international filing date but later than the priority date claimed

T later document published after the international filing date or priority date and not in conflict with the application but cited to understand the principle or theory underlying the invention

X document of particular relevance; the claimed invention cannot be considered novel or cannot be considered to involve an inventive step when the document is taken alone

Y document of particular relevance; the claimed invention cannot be considered to involve an inventive step when the document is combined with one or more other such documents, such combination being obvious to a person skilled in the art.

& document member of the same patent family

Date of the actual completion of the international search

1 September 2003

Date of mailing of the international search report

11/09/2003

Name and mailing address of the ISA

European Patent Office, P.B. 5818 Patentlaan 2
NL - 2280 HV Rijswijk
Tel. (+31-70) 340-2040, Tx. 31 651 epo nl,
Fax: (+31-70) 340-3016

Authorized officer

Prins, L

INTERNATIONAL SEARCH REPORT

International Application No
PCT/US 03/13807

C.(Continuation) DOCUMENTS CONSIDERED TO BE RELEVANT		
Category °	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	<p>POSCH K C ET AL: "APPROACHING ENCRYPTION AT ISDN SPEED USING PARTIAL PARALLEL MODULUS MULTIPLICATION" MICROPROCESSING AND MICROPROGRAMMING, ELSEVIER SCIENCE PUBLISHERS, BV., AMSTERDAM, NL, vol. 29, no. 3, October 1990 (1990-10), pages 177-184, XP000151455 ISSN: 0165-6074 paragraphs '0001!, '0002! -----</p>	1-5

INTERNATIONAL SEARCH REPORT

Information on patent family members

International Application No
PCT/US 03/13807

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 5166978	A	24-11-1992	FR 2658932 A1
			DE 69130581 D1
			DE 69130581 T2
			EP 0443679 A1
			JP 4216588 A
			JP 2003029627 A
