



- (51) International Patent Classification:
G06F 9/46 (2006.01) *G06F 9/38* (2006.01)
- (21) International Application Number:
PCT/US2013/045465
- (22) International Filing Date:
12 June 2013 (12.06.2013)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data:
13/631,165 28 September 2012 (28.09.2012) US
- (71) Applicant: INTEL CORPORATION [US/US]; 2200 Mission College Boulevard, MS: RNB-4-150, Santa Clara, California 95054 (US).
- (72) Inventors: WOLRICH, Gilbert M.; 1 Macomber Lane, Framingham, Massachusetts 01702 (US). YAP, Kirk S.; 224 Belknap Road, Framingham, Massachusetts 01701 (US). GUILFORD, James D.; 17 Mashpee Circle, Northborough, Massachusetts 01532 (US). GOPAL, Vinodh; 15 West End Avenue, Westborough, Massachusetts 01581 (US). GULLEY, Sean M.; 146 Marlborough Street, Boston, Massachusetts 02116 (US).
- (74) Agents: VINCENT, Lester J. et al.; Blakely Sokoloff Taylor & Zafman, 1279 Oakmead Parkway, Sunnyvale, California 94085 (US).

(81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BN, BR, BW, BY, BZ, CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PA, PE, PG, PH, PL, PT, QA, RO, RS, RU, RW, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, RW, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, RS, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, KM, ML, MR, NE, SN, TD, TG).

Published:

— with international search report (Art. 21(3))

(54) Title: INSTRUCTION SET FOR MESSAGE SCHEDULING OF SHA256 ALGORITHM




FIG. 4

(57) Abstract: A processor includes a first execution unit to receive and execute a first instruction to process a first part of secure hash algorithm 256 (SHA256) message scheduling operations, the first instruction having a first operand associated with a first storage location to store a first set of message inputs and a second operand associated with a second storage location to store a second set of message inputs. The processor further includes a second execution unit to receive and execute a second instruction to process a second part of the SHA256 message scheduling operations, the second instruction having a third operand associated with a third storage location to store an intermediate result of the first part and a third set of message inputs and a fourth operand associated with a fourth storage location to store a fourth set of message inputs.

INSTRUCTION SET FOR MESSAGE SCHEDULING OF SHA256 ALGORITHMTECHNICAL FIELD

Embodiments of the present invention relate generally to instruction processing
 5 apparatuses. More particularly, embodiments of the invention relate to instruction processing
 apparatus to process SHA-256 secure hashing algorithms.

BACKGROUND ART

SHA stands for Secure Hash Algorithm. It consists of five hash functions designed by the
 National Security Agency (NSA) and published by the National Institute of Standards and
 10 Technology (NIST). One of them is SHA-2. SHA-2 is a set of secure hash functions including
 SHA 224, SHA 256, SHA 384, and SHA 512 developed by the NSA intended to provide a
 higher level of security than the SHA-1 algorithm. SHA 224 and SHA 256 are similar
 algorithms based on a 32 bit word length producing digests of 224 and 256 bits. SHA 384 and
 SHA 512 are based on 64 bit words and produce digests of 384 and 512 bits.

15 The SHA-2 algorithm is computationally more complex the SHA 1, relying on carry
 propagate additions as well as logical operations and rotates. A critical path for a round of SHA-
 2 operations consists of four consecutive propagate additions with adder inputs being determined
 by complex logical and rotation functions. Figure 1 depicts details of the SHA-2 algorithm. A,
 B, C, D, E, F, G, and H represent the 8 words of state (32 bits for SHA 224/256 and 64 bits for
 20 SHA384/512). Following operations are performed for each iteration:

$$\begin{aligned} \text{Ch}(E, F, G) &= (E \wedge F) \oplus (\neg E \wedge G) \\ \text{Ma}(A, B, C) &= (A \wedge B) \oplus (A \wedge C) \oplus (B \wedge C) \\ \Sigma_0(A) &= (A \ggg 2) \oplus (A \ggg 13) \oplus (A \ggg 22) \\ \Sigma_1(E) &= (E \ggg 6) \oplus (E \ggg 11) \oplus (E \ggg 25) \end{aligned}$$

25 The bitwise rotation uses different constants for SHA-512. In this example, the given numbers
 are for SHA-256. Constant K plus W_i message input addition can be performed ahead of the
 round critical path. The message scheduling function for the SHA-2 algorithm is also more
 complex than SHA-1 relying on rotated copies of previous message inputs to form message
 inputs:

30 for i from 16 to 63

$$\begin{aligned} s0 &:= (w[i-15] \text{ ROTR } 7) \text{ XOR } (w[i-15] \text{ ROTR } 18) \text{ XOR } (w[i-15] \text{ SHR } 3) \\ s1 &:= (w[i-2] \text{ ROTR } 17) \text{ XOR } (w[i-2] \text{ ROTR } 19) \text{ XOR } (w[i-2] \text{ SHR } 10) \\ w[i] &:= w[i-16] + s0 + w[i-7] + s1 \end{aligned}$$

where ROTR (also used as ">>>") denotes a bitwise right-rotate operator; SHR denotes a bitwise
 35 right-shift operator; and XOR denotes a bitwise exclusive-OR operator.

For SHA-256, each iteration is performed as follows:

```

5       $\Sigma_0 := (a \text{ ROTR } 2) \text{ XOR } (a \text{ ROTR } 13) \text{ XOR } (a \text{ ROTR } 22)$ 
       $\text{maj} := (a \text{ AND } b) \text{ XOR } (a \text{ AND } c) \text{ XOR } (b \text{ AND } c)$ 
       $t2 := \Sigma_0 + \text{maj}$ 
       $\Sigma_1 := (e \text{ ROTR } 6) \text{ XOR } (e \text{ ROTR } 11) \text{ XOR } (e \text{ ROTR } 25)$ 
       $\text{ch} := (e \text{ AND } f) \text{ XOR } ((\text{NOT } e) \text{ AND } g)$ 
       $t1 := h + \Sigma_1 + \text{ch} + k[i] + w[i]$ 

10      $h := g$ 
       $g := f$ 
       $f := e$ 
       $e := d + t1$ 
       $d := c$ 
       $c := b$ 
15      $b := a$ 
       $a := t1 + t2$ 

```

Message input $w[i]$ for rounds 1 to 16 is the 32bit x 16 = 512 bit block of data. $W[i]$ for rounds 17 to 64 must be derived. Constant K is specified for each round, the $W[i] + K[i]$ value for each round can be calculated ahead of the actual round iteration. Further detailed information concerning the SHA-2 specification can be found in Secure Hash Standard published by Federal Information Processing Standard Publication (FIPS PUB 180-3, published October, 2008).

The round processing requires 8 (32-bit) state variables A through H. It is possible to split these across two 128-bit registers. However, to be able to compute a whole round in the data-path, we also need the $w[i] + k[i]$ inputs. Even if these can be added earlier, it introduces at least another source operand which we cannot use in a 2-source processor. An option would be to add the WK value to H prior to each round, and have the instruction process one round. This would add instructions that would limit the throughput, and more importantly directly add to the latency of around. The next round cannot start before the completion of a previous round. If we have 1 and 3-cycle pipelines in the single instruction multiple data (SIMD) units, then we would be limited to 1+3 cycles per round at best. In a previous design we proposed an instruction set on the 128-bit register set with two operands, that is able to achieve 3 cycles per round, by careful partitioning of the state variables, injection of the WK values and computing multiple rounds per instruction. There has been a lack of efficient ways to perform the above operations.

BRIEF DESCRIPTION OF THE DRAWINGS

Embodiments of the invention are illustrated by way of example and not limitation in the figures of the accompanying drawings in which like references indicate similar elements.

Figure 1 depicts details of the SHA-2 algorithm.

5 Figure 2 is a block diagram of an execution pipeline of a processor or processor core according to one embodiment of the invention.

Figure 3 is a block diagram illustrating a SHA-256 round operation according to one embodiment.

10 Figure 4 is a block diagram illustrating a process for SHA-256 operations according to one embodiment.

Figure 5 is a flow diagram illustrating a method for performing SHA-2 message scheduling operations according to one embodiment.

Figure 6 is a flow diagram illustrating a method for performing SHA-2 message scheduling operations according to another embodiment.

15 Figure 7 is pseudocode illustrating a process of SHA256 operations according to one embodiment.

Figures 8A and 8B show alternative instructions for SHA512 operations according to certain embodiments.

20 Figure 9A illustrates an exemplary advanced vector extensions (AVX) instruction format according to one embodiment of the invention.

Figure 9B illustrates an exemplary advanced vector extensions (AVX) instruction format according to another embodiment of the invention.

Figure 9C illustrates an exemplary advanced vector extensions (AVX) instruction format according to another embodiment of the invention.

25 Figure 10A is a block diagram illustrating a generic vector friendly instruction format and class A instruction templates thereof according to embodiments of the invention.

Figure 10B is a block diagram illustrating the generic vector friendly instruction format and class B instruction templates thereof according to embodiments of the invention.

30 Figure 11A is a block diagram illustrating an exemplary specific vector friendly instruction format according to one embodiment of the invention.

Figure 11B is a block diagram illustrating a generic vector friendly instruction format according to another embodiment of the invention.

Figure 11C is a block diagram illustrating a generic vector friendly instruction format according to another embodiment of the invention.

Figure 11D is a block diagram illustrating a generic vector friendly instruction format according to another embodiment of the invention.

Figure 12 is a block diagram of register architecture according to one embodiment of the invention.

5 Figure 13A is a block diagram illustrating both an exemplary in-order pipeline and an exemplary register renaming, out-of-order issue/execution pipeline according to embodiments of the invention.

10 Figure 13B is a block diagram illustrating both an exemplary embodiment of an in-order architecture core and an exemplary register renaming, out-of-order issue/execution architecture core to be included in a processor according to embodiments of the invention.

Figure 14A is a block diagram of a processor core according to one embodiment of the invention.

Figure 14B is a block diagram of a processor core according to another embodiment of the invention.

15 Figure 15 is a block diagram of a processor according to embodiments of the invention.

Figure 16 is a block diagram of a system in accordance with one embodiment of the invention.

Figure 17 is a block diagram of a more specific exemplary system in accordance with an embodiment of the invention.

20 Figure 18 is a block diagram of a more specific exemplary system in accordance with another embodiment of the invention.

Figure 19 is a block diagram of a SoC in accordance with an embodiment of the invention.

25 Figure 20 is a block diagram contrasting the use of a software instruction converter to convert binary instructions in a source instruction set to binary instructions in a target instruction set according to embodiments of the invention.

DESCRIPTION OF THE EMBODIMENTS

Various embodiments and aspects of the inventions will be described with reference to details discussed below, and the accompanying drawings will illustrate the various embodiments. The following description and drawings are illustrative of the invention and are not to be
30 construed as limiting the invention. Numerous specific details are described to provide a thorough understanding of various embodiments of the present invention. However, in certain instances, well-known or conventional details are not described in order to provide a concise discussion of embodiments of the present inventions.

Reference in the specification to “one embodiment” or “an embodiment” means that a

particular feature, structure, or characteristic described in conjunction with the embodiment can be included in at least one embodiment of the invention. The appearances of the phrase “in one embodiment” in various places in the specification do not necessarily all refer to the same embodiment.

5 According to some embodiments, a new instruction set architecture (ISA) is utilized to perform one or more rounds of SHA-2 operations described above in response to a single instruction (e.g., a single instruction multiple data or SIMD instruction) to improve the efficiency of the SHA-2 computation. New instructions for SHA256 round processing using 128-bit registers (e.g., XMM registers) with two source limitations are utilized, which can be performed
10 in a 3-cycle pipeline processor or processor core. In addition, new instructions for SHA256 message scheduling are defined which utilize 128-bit SIMD with two source operand constraints.

Figure 2 is a block diagram of an execution pipeline of a processor or processor core according to one embodiment of the invention. Referring to Figure 2, processor 100 may represent any kind of instruction processing apparatuses. For example, processor 100 may be a
15 general-purpose processor. Processor 100 may be any of various complex instruction set computing (CISC) processors, various reduced instruction set computing (RISC) processors, various very long instruction word (VLIW) processors, various hybrids thereof, or other types of processors entirely. Processor 100 may also represent one or more processor cores.

Processor cores may be implemented in different ways, for different purposes, and in
20 different processors. For instance, implementations of such cores may include: 1) a general purpose in-order core intended for general-purpose computing; 2) a high performance general purpose out-of-order core intended for general-purpose computing; 3) a special purpose core intended primarily for graphics and/or scientific (throughput) computing. Implementations of different processors may include: 1) a central processing unit (CPU) including one or more
25 general purpose in-order cores intended for general-purpose computing and/or one or more general purpose out-of-order cores intended for general-purpose computing; and 2) a coprocessor including one or more special purpose cores intended primarily for graphics and/or scientific (throughput). Such different processors lead to different computer system architectures, which may include: 1) the coprocessor on a separate chip from the CPU; 2) the coprocessor on a
30 separate die in the same package as a CPU; 3) the coprocessor on the same die as a CPU (in which case, such a coprocessor is sometimes referred to as special purpose logic, such as integrated graphics and/or scientific (throughput) logic, or as special purpose cores); and 4) a system on a chip that may include on the same die the described CPU (sometimes referred to as the application core(s) or application processor(s)), the above described coprocessor, and

additional functionality. Exemplary core architectures are described next, followed by descriptions of exemplary processors and computer architectures.

In one embodiment, processor 100 includes, but is not limited to, instruction decoder 101 and one or more execution units 102. Instruction decoder 101 is to receive and decode instructions 103 from an instruction fetch unit (not shown). Instruction decoder 102 may generate and output one or more micro-operations, micro-code, entry points, microinstructions, other instructions, or other control signals, which reflect, or are derived from, the instructions. Instruction decoder 102 may be implemented using various different mechanisms. Examples of suitable mechanisms include, but are not limited to, microcode read only memories (ROMs), look-up tables, hardware implementations, programmable logic arrays (PLAs), and the like.

Execution units 102, which may include an arithmetic logic unit, or another type of logic unit capable of performing operations based on instructions. As a result of instruction decoder 102 decoding the instructions, execution unit 102 may receive one or more micro-operations, micro-code entry points, microinstructions, other instructions, or other control signals, which reflect, or are derived from, the instructions. Execution unit 102 may be operable as a result of instructions indicating one or more source operands (SRC) and to store a result in one or more destination operands (DEST) of a register set indicated by the instructions. Execution unit 102 may include circuitry or other execution logic (e.g., software combined with hardware and/or firmware) operable to execute instructions or other control signals derived from the instructions and perform an operation accordingly. Execution unit 102 may represent any kinds of execution units such as logic units, arithmetic logic units (ALUs), arithmetic units, integer units, etc.

Some or all of the source and destination operands may be stored in registers of a register set or memory. The register set may be part of a register file, along with potentially other registers, such as status registers, flag registers, etc. A register may be a storage location or device that may be used to store data. The register set may often be physically located on die with the execution unit(s). The registers may be visible from the outside of the processor or from a programmer's perspective. For example, instructions may specify operands stored in the registers. Various different types of registers are suitable, as long as they are capable of storing and providing data as described herein. The registers may or may not be renamed. Examples of suitable registers include, but are not limited to, dedicated physical registers, dynamically allocated physical registers using register renaming, combinations of dedicated and dynamically allocated physical registers, etc. Alternatively, one or more of the source and destination operands may be stored in a storage location other than a register, such as, for example, a location in system memory.

Referring back to Figure 2, according to one embodiment, execution unit 102 includes one or more SHA-2 units 106 to, in response to a first instruction received and provided by instruction decoder 101, to perform one or more rounds of SHA-2 round operations using data 110 such as SHA-2 states A to H, message inputs (w), and corresponding constants Wt and Kt stored in one or more registers 104. The one or more rounds of SHA-2 round operations are performed in response to one or more SIMD instructions. The number of rounds of SHA-2 round operations depends on the specific design or configuration (e.g., pipeline latency requirement) of the processor pipeline, which may be configured to an appropriate number that will optimize the overall performance of the processor pipeline. For the purpose of illustration, it is assumed two rounds of SHA-2 round operations are performed. It will be appreciated, more or fewer rounds of SHA-2 round operations can also be performed, as long as the required resources such as registers or memory with proper sizes are available.

According to one embodiment, two rounds of SHA-2 round operations are performed using multiple SIMD instructions in a 128-bit data path: 1) SHA2_IN; 2) SHA2_LO; and 3) SHA2_HI instructions. Execution of these three instructions completes at least two rounds of SHA256 operations with a critical path six cycles for two rounds: a SHA2_IN 3-cycle pipeline followed by a SHA2_LO 3-cycle pipeline. SHA2_HI and message scheduling (MSG1SHA256 and MSG2SHA256) operations are hidden inside execution of the SHA2_IN and SHA2_LO instructions (e.g., executed in parallel). All of these instructions utilize operands with at least 128 bits, such as XMM registers. Throughout this application, XMM registers are utilized as examples of 128-bit registers.

In one embodiment, instruction SHA_IN may be defined as follows:

SHA2_IN XMM0, XMM1

SHA2_IN instruction performs partial calculation for next two rounds, injecting messages plus constants for two rounds. In this instruction, XMM0 is a source/destination register. As an input, XMM0 stores SHA2 states E, F, G, and H. XMM1 stores constants WK2 and WK1 for two rounds. Specifically, as an input according to one embodiment, the SHA2 states are stored in XMM0 register as follows:

XMM0 [127:96] = E

XMM0 [95:64] = F

XMM0 [63:32] = G

XMM0 [31:0] = H

The output of the operations is stored in XMM0 as a destination register, including the states E and F, and the intermediate results X and Y. According to one embodiment, as an

output, the destination register XMM0 is defined as follows:

$$\text{XMM0 [127:96]} = E$$

$$\text{XMM0 [95:64]} = F$$

$$\text{XMM0 [63:32]} = X = \text{WK2} + G$$

$$5 \quad \text{XMM0 [31:0]} = Y = \text{WK1} + H + \text{Ch}(E, F, G) + \sum(E)$$

In one embodiment, instruction SHA_LO may be defined as follows:

SHA2_LO XMM0, XMM1

SHA2_LO partially calculates two rounds of SHA256 states E, F, G, and H for current round, referred to as updated E, F, G, and H, from the input [A, B, C, D] and [E, F, G, H]. In this instruction, XMM0 is a source/destination register. As an input, XMM0 stores SHA2 states E, F, X, and Y that are calculated from the above SHA2_IN instruction, as follows:

$$\text{XMM0 [127:96]} = E$$

$$\text{XMM0 [95:64]} = F$$

$$\text{XMM0 [63:32]} = X$$

$$15 \quad \text{XMM0 [31:0]} = Y$$

XMM1 stores input states A, B, C, and D. Specifically, according to one embodiment, as an input, the SHA2 states are stored in XMM1 register as follows:

$$\text{XMM1 [127:96]} = A$$

$$\text{XMM1 [95:64]} = B$$

$$20 \quad \text{XMM1 [63:32]} = C$$

$$\text{XMM1 [31:0]} = D$$

After the execution of instruction SHA_LO, as an output according to one embodiment, XMM0 is defined as follows:

$$\text{XMM0 [127:96]} = \text{updated E}$$

$$25 \quad \text{XMM0 [95:64]} = \text{updated F}$$

$$\text{XMM0 [63:32]} = \text{updated G}$$

$$\text{XMM0 [31:0]} = \text{updated H}$$

In one embodiment, instruction SHA_HI may be defined as follows:

SHA2_HI XMM0, XMM1

SHA2_HI partially calculates two rounds of SHA256 states A, B, C, and D for current round, referred to as updated A, B, C, and D, from the input [A, B, C, D] and [E, F, X, Y]. In this instruction, XMM0 is a source/destination register. As an input according to one embodiment, XMM0 stores SHA2 original states A, B, C, and D as follows:

$$\text{XMM0 [127:96]} = A$$

XMM0 [95:64] = B

XMM0 [63:32] = C

XMM0 [31:0] = D

As an input according to one embodiment, XMM1 stores E, F, X, and Y (which is
5 calculated by SHA2_IN) as follows:

XMM1 [127:96] = E

XMM1 [95:64] = F

XMM1 [63:32] = X

XMM1 [31:0] = Y

10 As an output according to one embodiment, XMM0 is defined as follows:

XMM0 [127:96] = updated A

XMM0 [95:64] = updated B

XMM0 [63:32] = updated C

XMM0 [31:0] = updated D

15 Figure 3 is a block diagram illustrating an example of an execution pipeline for SHA2 round operations according to one embodiment. System 300 is configured to process the instructions SHA2_IN, SHA2_LO, and SHA2_HI, where microcodes 301-303 may be executed in different pipeline stages.

According to another embodiment, in response to an instruction, the SHA-2 unit 106 is
20 configured to perform SHA-2 message scheduling operations to produce multiple message inputs for the next cycle (e.g., one or more rounds of SHA-2 round operations in a next SIMD instruction cycle). In one embodiment, there may be two instructions needed to prepare the message inputs for the next SIMD cycle dependent upon the specific design or configuration of the processor pipeline. In supporting at least two rounds of SHA-2 round operations, the first
25 instruction includes two operands to store at least 8 previous message inputs and after the first instruction is executed, an intermediate result is generated and returned in a register specified in one of the operands. The second instruction takes the intermediate result generated from the first instruction, as well as at least 4 previous message inputs, as an input in one operand. Another operand of the second instruction specifies at least 4 other previous message inputs. The final
30 result represents 4 message inputs for the next SIMD cycle. In one embodiment, registers involved in the message scheduling operations have at least 128 bits.

To keep pace with the above round instructions, according to one embodiment, two specialized instructions performing the SHA-256 message schedule are provided, referred to herein as instructions MSG1SHA256 and MSG2SHA256. MSG1SHA256 and MSG2SHA256

calculate words $w(i)$ for rounds 16 to 63 as follows:

$$s0 = (w(i-15) \gg 7) \text{ XOR } (w(i-15) \gg 18) \text{ XOR } (w(i-15) \gg 3)$$

$$s1 = (w(i-2) \gg 17) \text{ XOR } (w(i-2) \gg 19) \text{ XOR } (w(i-2) \gg 10)$$

$$w[i] = w(i-16) + s0 + w(i+7) + s1$$

5

$$\text{msg1} = w(i-16) + s0$$

$$\text{msg1} = w(i-16) + ((w(i-15) \gg 7) \text{ XOR } (w(i-15) \gg 18) \text{ XOR } (w(i-15) \gg 3))$$

$$w[i] = \text{msg1}(i-16) + w(i-7) + s1$$

In one embodiment, instruction MSG1SHA256 may be defined as follows:

10

MSG1SHA256 XMM0, XMM1

where registers XMM0 and XMM1 represent registers having at least 128 bits such as those XMM registers available from the Intel AVX processor family. XMM0 is a source/destination register while XMM1 is a source register. Instruction MSG1SHA256 performs an intermediate calculation for the next four SHA256 message inputs.

15 In one embodiment, input XMM0 represents messages $w(i-13)$, $w(i-14)$, $w(i-15)$, and $w(i-16)$. An embodiment of the format of XMM0 can be defined as follows:

$$\text{XMM0}[127:96] = w(i-13)$$

$$\text{XMM0}[95:64] = w(i-14)$$

$$\text{XMM0}[63:32] = w(i-15)$$

20

$$\text{XMM0}[31:00] = w(i-16)$$

Input XMM1 represents messages $w(i-9)$, $w(i-10)$, $w(i-11)$, and $w(i-12)$. An embodiment of the format of XMM1 can be defined as follows:

$$\text{XMM1}[127:96] = w(i-9)$$

$$\text{XMM1}[95:64] = w(i-10)$$

25

$$\text{XMM1}[63:32] = w(i-11)$$

$$\text{XMM1}[31:00] = w(i-12)$$

Output XMM0 represents $w(i-13) + s0(w(i-12))$, $w(i-14) + s0(w(i-13))$, $w(i-15) + s0(w(i-14))$, and $w(i-16) + s0(w(i-15))$. An embodiment of the format of XMM0 can be defined as follows:

30

$$\text{XMM0}[127:96] = \text{msg1}(i-13) = w(i-13) + s0(w[i-12])$$

$$\text{XMM0}[95:64] = \text{msg1}(i-14) = w(i-14) + s0(w[i-13])$$

$$\text{XMM0}[63:32] = \text{msg1}(i-15) = w(i-15) + s0(w[i-14])$$

$$\text{XMM0}[31:00] = \text{msg1}(i-16) = w(i-16) + s0(w[i-15])$$

where $\text{XMM0}[127:96]$ represents $\text{Msg1}(i-13)$ for determining $w(i+3)$; $\text{XMM0}[95:64]$ represents

Msg1(i-14) for determining $w(i+2)$; XMM0[63:32] represents Msg1(i-15) for determining $w(i+1)$; and XMM0[31:00] represents Msg1(i-16) for determining $w(i)$.

In one embodiment, instruction MSG2SHA256 may be defined as follows:

MSG2SHA256 XMM0, XMM1

5 where registers XMM0 and XMM1 represent registers having at least 128 bits such as those XMM registers available from Intel AVX processor family. XMM0 register is a source/destination register while XMM1 is a source register. Instruction MSG2SHA256 performs calculation for the next four SHA-256 message inputs containing 1) the msg1 values for $w(i-13)$ to $w(i-16)$ and message inputs $w(i-5)$ to $w(i-7)$; and 2) message inputs $w(i-1)$ to $w(i-4)$.

In one embodiment, input XMM0 is defined as follows:

XMM0 [127:96] = msg1(i-13)

XMM0 [95:64] = msg1 (i-14) + $w(i-5)$

XMM0 [63:32] = msg1 (i-15) + $w(i-6)$

15 XMM0 [31:00] = msg1 (i-16) + $w(i-7)$

In one embodiment, for a current round i of SHA-256 round operations, input XMM1 represents messages $w(i-1)$, $w(i-2)$, $w(i-3)$, and $w(i-4)$ as follows:

XMM1[127:96] = $w(i-1)$

XMM1[95:64] = $w(i-2)$

20 XMM1[63:32] = $w(i-3)$

XMM1[31:00] = $w(i-4)$

In one embodiment, for a current round i of SHA-256 round operations, output XMM0 represents new messages $w(i+3)$, $w(i+2)$, $w(i+1)$, and $w(i)$ as follows:

XMM0[127:96] = $w(i+3)$ = msg1(i-13) + $w(i-4)$ + s1($w(i+1)$)

25 XMM0[95:64] = $w(i+2)$ = msg1(i-14) + $w(i-5)$ + s1($w(i)$)

XMM0[63:32] = $w(i+1)$ = msg1(i-15) + $w(i-6)$ + s1($w(i-1)$)

XMM0[31:00] = $w(i)$ = msg1(i-16) + $w(i-7)$ + s1($w(i-2)$)

Note that input XMM0 to this instruction is not the direct output of msg1, but a modified version with some w words added in 3 of the 4 lanes. This can be achieved by a new instruction or a combination of two existing instructions such as PERM, ADD instructions.

Figure 4 is a block diagram illustrating a process for SHA-256 round operations according to one embodiment. Referring to Figure 4, an embodiment of the process can be defined with at least three pipeline stages 401-403 of a processor or processor core. Note that pipeline stages 401-403 may or may not be consecutive pipeline stages dependent upon the specific design or

configuration of the processor. MSG1SHA256 microcode 404, MSG2SHA256 microcode 405, and SHA256 round microcode 406 may be implemented as part of SHA-2 unit 106 of Figure 2. In one embodiment, SHA256 round microcode 406 may be implemented as part of logic 300 of Figure 3.

5 According to one embodiment, in response to instruction MSG1SHA256, MSG1SHA256 microcode 404 is to perform a first part of SHA-256 message scheduling operations as described above. For a given round i of SHA-256 round operations, during pipeline stage 401, MSG1SHA256 microcode 404 is to perform the first part of the message scheduling operations on previously generated messages 407-408 and to generate intermediate message 409.

10 According to one embodiment, in response to instruction MSG2SHA256, MSG2SHA256 microcode 405 is to perform a second part of SHA-256 message scheduling operations as described above. For a given round i of SHA-256 round operations, during pipeline stage 402, MSG2SHA256 microcode 405 is to perform the second part of the message scheduling operations on previously generated message 410 and intermediate message 409 and message
15 inputs 411 (which may be combined using another instruction such as PERM and/or ADD), and to generate message 412.

 According to one embodiment, in response to instruction SHA256RND2, SHA256RND2 microcode 406 is to perform a round of SHA-256 round operations as described above. For a given round i of SHA-256 round operations, during pipeline stage 403,
20 SHA256RND2 microcode 406 is to perform a round operation on messages 412 and current SHA-256 states 413, and to generate SHA-256 states 414 for next round or iteration. Note that in this example, each of $w(i+3)$, $w(i+2)$, $w(i+1)$, and $w(i)$ stored in register 412 has 32 bits. It can be used to perform at least two rounds of SHA-256 round operations. If register 412 can store more message inputs, more rounds of SHA-256 round operations can also be performed, as
25 long as the pipeline latency requirement can be satisfied.

 Figure 5 is a flow diagram illustrating a method for performing SHA-2 message scheduling operations according to one embodiment. Method 500 may be performed by SHA-2 unit 106 of Figure 1. Referring to Figure 5, at block 501, an instruction (e.g., MSG1SHA256) is received by an execution unit of a processor, where the instruction includes a first operand (e.g., XMM0) and
30 a second operand (e.g., XMM1) storing message inputs for performing a first part of SHA256 message scheduling operations. At block 502, multiple message inputs (e.g., at least 8 message inputs) are extracted from storage locations indicated by the first and second operands. At block 503, a first part of the SHA256 message scheduling operations is performed based on the extracted message inputs. At block 504, an intermediate result of the first part of the SHA256

message scheduling operations is stored in a storage location indicated by the first operand (e.g., XMM0).

Figure 6 is a flow diagram illustrating a method for performing SHA-2 message scheduling operations according to one embodiment. Method 600 may be performed by SHA-2 unit 106 of Figure 1 as a second part of the SHA-2 message scheduling operations. Referring to Figure 6, at block 601, an instruction (e.g., MSG2SHA256) is received, where the instruction includes two operands (e.g., XMM0, XMM1), each identifying a storage location with at least 128 bits. The first operand refers to a first storage location storing an intermediate result of the first part of message scheduling operations and a first set of message inputs. The second operand refers to a second storage location storing a second set of message inputs. At block 602, the intermediate result and the first set of message input are extracted from the storage location associated with the first operand and a second set of message input is extracted from the storage location associated with the second operand. At block 603, a second part of SHA-2 message scheduling operations is performed based on the intermediate result and the message inputs. At block 604, a final result is generated and stored in a register associated with the first operand (e.g., XMM0). Figure 7 is pseudocode illustrating a process of SHA256 operations according to one embodiment.

According to another embodiment, a new round instruction SHA2_2RND may be defined that essentially performs the SHA2_LO and SHA2_HI operations. According to one embodiment, instruction SHA2_2RND may be defined as follows:

SHA2_2RND XMM0, XMM1

where both XMM0 and XMM1 are source/destination registers. As inputs, XMM0 stores states A, B, C, and D, while XMM1 stores states E and F, and intermediate result X and Y. The sample sequence of operations of SHA2_2RND may be as follows:

Tmp = XMM1
SHA2_LO XMM1, XMM0
SHA2_HI XMM0, Tmp

After two rounds, as outputs, XMM0 stores updated A, B, C, and D, while XMM1 stores updated E, F, G, and H.

Although the embodiments have been described for SHA256 operations, however, the above techniques can also be applied to other wider data paths such as SHA512 operations, as shown in Figures 8A and 8B. Execution of SHA512IN1, SHA512IN2, SHA512LO, and SHA512HI completes two rounds of SHA512 with a critical path of 9 cycles for two rounds: SHA512IN1 (3-cycle pipeline) followed by SHA512IN2 (3-cycle pipeline) and followed by

SHA512HI (3-cycle pipeline). SHA512LO and the message scheduling instructions are hidden inside the execution of SHA512IN1, SHA512IN2, and SHA512HI.

An instruction set, or instruction set architecture (ISA), is the part of the computer architecture related to programming, and may include the native data types, instructions, register architecture, addressing modes, memory architecture, interrupt and exception handling, and external input and output (I/O). The term instruction generally refers herein to macro-instructions – that is instructions that are provided to the processor (or instruction converter that translates (e.g., using static binary translation, dynamic binary translation including dynamic compilation), morphs, emulates, or otherwise converts an instruction to one or more other instructions to be processed by the processor) for execution – as opposed to micro-instructions or micro-operations (micro-ops) – that is the result of a processor's decoder decoding macro-instructions.

The ISA is distinguished from the microarchitecture, which is the internal design of the processor implementing the instruction set. Processors with different microarchitectures can share a common instruction set. For example, Intel® Pentium 4 processors, Intel® Core™ processors, and processors from Advanced Micro Devices, Inc. of Sunnyvale CA implement nearly identical versions of the x86 instruction set (with some extensions that have been added with newer versions), but have different internal designs. For example, the same register architecture of the ISA may be implemented in different ways in different microarchitectures using well-known techniques, including dedicated physical registers, one or more dynamically allocated physical registers using a register renaming mechanism (e.g., the use of a Register Alias Table (RAT), a Reorder Buffer (ROB), and a retirement register file; the use of multiple maps and a pool of registers), etc. Unless otherwise specified, the phrases register architecture, register file, and register are used herein to refer to that which is visible to the software/programmer and the manner in which instructions specify registers. Where a specificity is desired, the adjective logical, architectural, or software visible will be used to indicate registers/files in the register architecture, while different adjectives will be used to designation registers in a given microarchitecture (e.g., physical register, reorder buffer, retirement register, register pool).

An instruction set includes one or more instruction formats. A given instruction format defines various fields (number of bits, location of bits) to specify, among other things, the operation to be performed (opcode) and the operand(s) on which that operation is to be performed. Some instruction formats are further broken down though the definition of instruction templates (or subformats). For example, the instruction templates of a given

instruction format may be defined to have different subsets of the instruction format's fields (the included fields are typically in the same order, but at least some have different bit positions because there are less fields included) and/or defined to have a given field interpreted differently. Thus, each instruction of an ISA is expressed using a given instruction format (and, if defined, in a given one of the instruction templates of that instruction format) and includes fields for specifying the operation and the operands. For example, an exemplary ADD instruction has a specific opcode and an instruction format that includes an opcode field to specify that opcode and operand fields to select operands (source1/destination and source2); and an occurrence of this ADD instruction in an instruction stream will have specific contents in the operand fields that select specific operands.

Scientific, financial, auto-vectorized general purpose, RMS (recognition, mining, and synthesis), and visual and multimedia applications (e.g., 2D/3D graphics, image processing, video compression/decompression, voice recognition algorithms and audio manipulation) often require the same operation to be performed on a large number of data items (referred to as "data parallelism"). Single Instruction Multiple Data (SIMD) refers to a type of instruction that causes a processor to perform an operation on multiple data items. SIMD technology is especially suited to processors that can logically divide the bits in a register into a number of fixed-sized data elements, each of which represents a separate value. For example, the bits in a 256-bit register may be specified as a source operand to be operated on as four separate 64-bit packed data elements (quad-word (Q) size data elements), eight separate 32-bit packed data elements (double word (D) size data elements), sixteen separate 16-bit packed data elements (word (W) size data elements), or thirty-two separate 8-bit data elements (byte (B) size data elements). This type of data is referred to as packed data type or vector data type, and operands of this data type are referred to as packed data operands or vector operands. In other words, a packed data item or vector refers to a sequence of packed data elements, and a packed data operand or a vector operand is a source or destination operand of a SIMD instruction (also known as a packed data instruction or a vector instruction).

By way of example, one type of SIMD instruction specifies a single vector operation to be performed on two source vector operands in a vertical fashion to generate a destination vector operand (also referred to as a result vector operand) of the same size, with the same number of data elements, and in the same data element order. The data elements in the source vector operands are referred to as source data elements, while the data elements in the destination vector operand are referred to a destination or result data elements. These source vector operands are of the same size and contain data elements of the same width, and thus they contain

the same number of data elements. The source data elements in the same bit positions in the two source vector operands form pairs of data elements (also referred to as corresponding data elements; that is, the data element in data element position 0 of each source operand correspond, the data element in data element position 1 of each source operand correspond, and so on). The operation specified by that SIMD instruction is performed separately on each of these pairs of source data elements to generate a matching number of result data elements, and thus each pair of source data elements has a corresponding result data element. Since the operation is vertical and since the result vector operand is the same size, has the same number of data elements, and the result data elements are stored in the same data element order as the source vector operands, the result data elements are in the same bit positions of the result vector operand as their corresponding pair of source data elements in the source vector operands. In addition to this exemplary type of SIMD instruction, there are a variety of other types of SIMD instructions (e.g., that has only one or has more than two source vector operands, that operate in a horizontal fashion, that generates a result vector operand that is of a different size, that has a different size data elements, and/or that has a different data element order). It should be understood that the term destination vector operand (or destination operand) is defined as the direct result of performing the operation specified by an instruction, including the storage of that destination operand at a location (be it a register or at a memory address specified by that instruction) so that it may be accessed as a source operand by another instruction (by specification of that same location by the another instruction).

The SIMD technology, such as that employed by the Intel® Core™ processors having an instruction set including x86, MMX™, Streaming SIMD Extensions (SSE), SSE2, SSE3, SSE4.1, and SSE4.2 instructions, has enabled a significant improvement in application performance. An additional set of SIMD extensions, referred to the Advanced Vector Extensions (AVX) (AVX1 and AVX2) and using the Vector Extensions (VEX) coding scheme, has been , has been released and/or published (e.g., see Intel® 64 and IA-32 Architectures Software Developers Manual, October 2011; and see Intel® Advanced Vector Extensions Programming Reference, June 2011).

Embodiments of the instruction(s) described herein may be embodied in different formats. Additionally, exemplary systems, architectures, and pipelines are detailed below. Embodiments of the instruction(s) may be executed on such systems, architectures, and pipelines, but are not limited to those detailed.

VEX encoding allows instructions to have more than two operands, and allows SIMD vector registers to be longer than 128 bits. The use of a VEX prefix provides for three-operand

(or more) syntax. For example, previous two-operand instructions performed operations such as $A = A + B$, which overwrites a source operand. The use of a VEX prefix enables operands to perform nondestructive operations such as $A = B + C$.

Figure 9A illustrates an exemplary AVX instruction format including a VEX prefix 2102, real opcode field 2130, Mod R/M byte 2140, SIB byte 2150, displacement field 2162, and IMM8 2172. Figure 9B illustrates which fields from Figure 9A make up a full opcode field 2174 and a base operation field 2142. Figure 9C illustrates which fields from Figure 9A make up a register index field 2144.

VEX Prefix (Bytes 0-2) 2102 is encoded in a three-byte form. The first byte is the Format Field 2140 (VEX Byte 0, bits [7:0]), which contains an explicit C4 byte value (the unique value used for distinguishing the C4 instruction format). The second-third bytes (VEX Bytes 1-2) include a number of bit fields providing specific capability. Specifically, REX field 2105 (VEX Byte 1, bits [7-5]) consists of a VEX.R bit field (VEX Byte 1, bit [7] – R), VEX.X bit field (VEX byte 1, bit [6] – X), and VEX.B bit field (VEX byte 1, bit[5] – B). Other fields of the instructions encode the lower three bits of the register indexes as is known in the art (rrr, xxx, and bbb), so that Rrrr, Xxxx, and Bbbb may be formed by adding VEX.R, VEX.X, and VEX.B. Opcode map field 2115 (VEX byte 1, bits [4:0] – mmmmm) includes content to encode an implied leading opcode byte. W Field 2164 (VEX byte 2, bit [7] – W) - is represented by the notation VEX.W, and provides different functions depending on the instruction. The role of VEX.vvvv 2120 (VEX Byte 2, bits [6:3]-vvvv) may include the following: 1) VEX.vvvv encodes the first source register operand, specified in inverted (1s complement) form and is valid for instructions with 2 or more source operands; 2) VEX.vvvv encodes the destination register operand, specified in 1s complement form for certain vector shifts; or 3) VEX.vvvv does not encode any operand, the field is reserved and should contain 1111b. If VEX.L 2168 Size field (VEX byte 2, bit [2]-L) = 0, it indicates 128 bit vector; if VEX.L = 1, it indicates 256 bit vector. Prefix encoding field 2125 (VEX byte 2, bits [1:0]-pp) provides additional bits for the base operation field.

Real Opcode Field 2130 (Byte 3) is also known as the opcode byte. Part of the opcode is specified in this field. MOD R/M Field 2140 (Byte 4) includes MOD field 2142 (bits [7-6]), Reg field 2144 (bits [5-3]), and R/M field 2146 (bits [2-0]). The role of Reg field 2144 may include the following: encoding either the destination register operand or a source register operand (the rrr of Rrrr), or be treated as an opcode extension and not used to encode any instruction operand. The role of R/M field 2146 may include the following: encoding the instruction operand that references a memory address, or encoding either the destination register operand or a source

register operand.

Scale, Index, Base (SIB) - The content of Scale field 2150 (Byte 5) includes SS2152 (bits [7-6]), which is used for memory address generation. The contents of SIB.xxx 2154 (bits [5-3]) and SIB.bbb 2156 (bits [2-0]) have been previously referred to with regard to the register indexes
 5 Xxxx and Bbbb. The Displacement Field 2162 and the immediate field (IMM8) 2172 contain address data.

A vector friendly instruction format is an instruction format that is suited for vector instructions (e.g., there are certain fields specific to vector operations). While embodiments are described in which both vector and scalar operations are supported through the vector friendly
 10 instruction format, alternative embodiments use only vector operations the vector friendly instruction format.

Figure 10A, Figure 10B, and Figure 10C are block diagrams illustrating a generic vector friendly instruction format and instruction templates thereof according to embodiments of the invention. Figure 10A is a block diagram illustrating a generic vector friendly instruction format
 15 and class A instruction templates thereof according to embodiments of the invention; while Figure 10B is a block diagram illustrating the generic vector friendly instruction format and class B instruction templates thereof according to embodiments of the invention. Specifically, a generic vector friendly instruction format 2200 for which are defined class A and class B instruction templates, both of which include no memory access 2205 instruction templates and
 20 memory access 2220 instruction templates. The term generic in the context of the vector friendly instruction format refers to the instruction format not being tied to any specific instruction set.

While embodiments of the invention will be described in which the vector friendly instruction format supports the following: a 64 byte vector operand length (or size) with 32 bit (4
 25 byte) or 64 bit (8 byte) data element widths (or sizes) (and thus, a 64 byte vector consists of either 16 doubleword-size elements or alternatively, 8 quadword-size elements); a 64 byte vector operand length (or size) with 16 bit (2 byte) or 8 bit (1 byte) data element widths (or sizes); a 32 byte vector operand length (or size) with 32 bit (4 byte), 64 bit (8 byte), 16 bit (2 byte), or 8 bit (1 byte) data element widths (or sizes); and a 16 byte vector operand length (or size) with 32 bit
 30 (4 byte), 64 bit (8 byte), 16 bit (2 byte), or 8 bit (1 byte) data element widths (or sizes); alternative embodiments may support more, less and/or different vector operand sizes (e.g., 256 byte vector operands) with more, less, or different data element widths (e.g., 128 bit (16 byte) data element widths).

The class A instruction templates in Figure 10A include: 1) within the no memory access

2205 instruction templates there is shown a no memory access, full round control type operation 2210 instruction template and a no memory access, data transform type operation 2215 instruction template; and 2) within the memory access 2220 instruction templates there is shown a memory access, temporal 2225 instruction template and a memory access, non-temporal 2230 instruction template. The class B instruction templates in Figure 10B include: 1) within the no memory access 2205 instruction templates there is shown a no memory access, write mask control, partial round control type operation 2212 instruction template and a no memory access, write mask control, vsize type operation 2217 instruction template; and 2) within the memory access 2220 instruction templates there is shown a memory access, write mask control 2227 instruction template.

The generic vector friendly instruction format 2200 includes the following fields listed below in the order illustrated in Figure 10A and Figure 10B. Format field 2240 – a specific value (an instruction format identifier value) in this field uniquely identifies the vector friendly instruction format, and thus occurrences of instructions in the vector friendly instruction format in instruction streams. As such, this field is optional in the sense that it is not needed for an instruction set that has only the generic vector friendly instruction format. Base operation field 2242 – its content distinguishes different base operations.

Register index field 2244 – its content, directly or through address generation, specifies the locations of the source and destination operands, be they in registers or in memory. These include a sufficient number of bits to select N registers from a PxQ (e.g. 32x512, 16x128, 32x1024, 64x1024) register file. While in one embodiment N may be up to three sources and one destination register, alternative embodiments may support more or less sources and destination registers (e.g., may support up to two sources where one of these sources also acts as the destination, may support up to three sources where one of these sources also acts as the destination, may support up to two sources and one destination).

Modifier field 2246 – its content distinguishes occurrences of instructions in the generic vector instruction format that specify memory access from those that do not; that is, between no memory access 2205 instruction templates and memory access 2220 instruction templates. Memory access operations read and/or write to the memory hierarchy (in some cases specifying the source and/or destination addresses using values in registers), while non-memory access operations do not (e.g., the source and destinations are registers). While in one embodiment this field also selects between three different ways to perform memory address calculations, alternative embodiments may support more, less, or different ways to perform memory address calculations.

Augmentation operation field 2250 – its content distinguishes which one of a variety of different operations to be performed in addition to the base operation. This field is context specific. In one embodiment of the invention, this field is divided into a class field 2268, an alpha field 2252, and a beta field 2254. The augmentation operation field 2250 allows common groups of operations to be performed in a single instruction rather than 2, 3, or 4 instructions. Scale field 2260 – its content allows for the scaling of the index field's content for memory address generation (e.g., for address generation that uses $2scale * index + base$).

Displacement Field 2262A – its content is used as part of memory address generation (e.g., for address generation that uses $2scale * index + base + displacement$). Displacement Factor Field 2262B (note that the juxtaposition of displacement field 2262A directly over displacement factor field 2262B indicates one or the other is used) – its content is used as part of address generation; it specifies a displacement factor that is to be scaled by the size of a memory access (N) – where N is the number of bytes in the memory access (e.g., for address generation that uses $2scale * index + base + scaled\ displacement$). Redundant low-order bits are ignored and hence, the displacement factor field's content is multiplied by the memory operands total size (N) in order to generate the final displacement to be used in calculating an effective address. The value of N is determined by the processor hardware at runtime based on the full opcode field 2274 (described later herein) and the data manipulation field 2254C. The displacement field 2262A and the displacement factor field 2262B are optional in the sense that they are not used for the no memory access 2205 instruction templates and/or different embodiments may implement only one or none of the two.

Data element width field 2264 – its content distinguishes which one of a number of data element widths is to be used (in some embodiments for all instructions; in other embodiments for only some of the instructions). This field is optional in the sense that it is not needed if only one data element width is supported and/or data element widths are supported using some aspect of the opcodes.

Write mask field 2270 – its content controls, on a per data element position basis, whether that data element position in the destination vector operand reflects the result of the base operation and augmentation operation. Class A instruction templates support merging-writemasking, while class B instruction templates support both merging- and zeroing-writemasking. When merging, vector masks allow any set of elements in the destination to be protected from updates during the execution of any operation (specified by the base operation and the augmentation operation); in other one embodiment, preserving the old value of each element of the destination where the corresponding mask bit has a 0. In contrast, when zeroing

vector masks allow any set of elements in the destination to be zeroed during the execution of any operation (specified by the base operation and the augmentation operation); in one embodiment, an element of the destination is set to 0 when the corresponding mask bit has a 0 value. A subset of this functionality is the ability to control the vector length of the operation being performed (that is, the span of elements being modified, from the first to the last one); however, it is not necessary that the elements that are modified be consecutive. Thus, the write mask field 2270 allows for partial vector operations, including loads, stores, arithmetic, logical, etc. While embodiments of the invention are described in which the write mask field's 2270 content selects one of a number of write mask registers that contains the write mask to be used (and thus the write mask field's 2270 content indirectly identifies that masking to be performed), alternative embodiments instead or additionally allow the mask write field's 2270 content to directly specify the masking to be performed.

Immediate field 2272 – its content allows for the specification of an immediate. This field is optional in the sense that it is not present in an implementation of the generic vector friendly format that does not support immediate and it is not present in instructions that do not use an immediate. Class field 2268 – its content distinguishes between different classes of instructions. With reference to Figure 10A and Figure 10B, the contents of this field select between class A and class B instructions. In Figure 10A and Figure 10B, rounded corner squares are used to indicate a specific value is present in a field (e.g., class A 2268A and class B 2268B for the class field 2268 respectively in Figure 10A and Figure 10B).

In the case of the non-memory access 2205 instruction templates of class A, the alpha field 2252 is interpreted as an RS field 2252A, whose content distinguishes which one of the different augmentation operation types are to be performed (e.g., round 2252A.1 and data transform 2252A.2 are respectively specified for the no memory access, round type operation 2210 and the no memory access, data transform type operation 2215 instruction templates), while the beta field 2254 distinguishes which of the operations of the specified type is to be performed. In the no memory access 2205 instruction templates, the scale field 2260, the displacement field 2262A, and the displacement scale field 2262B are not present.

In the no memory access full round control type operation 2210 instruction template, the beta field 2254 is interpreted as a round control field 2254A, whose content(s) provide static rounding. While in the described embodiments of the invention the round control field 2254A includes a suppress all floating point exceptions (SAE) field 2256 and a round operation control field 2258, alternative embodiments may support may encode both these concepts into the same field or only have one or the other of these concepts/fields (e.g., may have only the round

operation control field 2258).

SAE field 2256 – its content distinguishes whether or not to disable the exception event reporting; when the SAE field's 2256 content indicates suppression is enabled, a given instruction does not report any kind of floating-point exception flag and does not raise any

5 floating point exception handler.

Round operation control field 2258 – its content distinguishes which one of a group of rounding operations to perform (e.g., Round-up, Round-down, Round-towards-zero and Round-to-nearest). Thus, the round operation control field 2258 allows for the changing of the rounding mode on a per instruction basis. In one embodiment of the invention where a processor includes

10 a control register for specifying rounding modes, the round operation control field's 2250 content overrides that register value.

In the no memory access data transform type operation 2215 instruction template, the beta field 2254 is interpreted as a data transform field 2254B, whose content distinguishes which one of a number of data transforms is to be performed (e.g., no data transform, swizzle, broadcast).

15 In the case of a memory access 2220 instruction template of class A, the alpha field 2252 is interpreted as an eviction hint field 2252B, whose content distinguishes which one of the eviction hints is to be used (in Figure 10A, temporal 2252B.1 and non-temporal 2252B.2 are respectively specified for the memory access, temporal 2225 instruction template and the memory access, non-temporal 2230 instruction template), while the beta field 2254 is interpreted

20 as a data manipulation field 2254C, whose content distinguishes which one of a number of data manipulation operations (also known as primitives) is to be performed (e.g., no manipulation; broadcast; up conversion of a source; and down conversion of a destination). The memory access 2220 instruction templates include the scale field 2260, and optionally the displacement field 2262A or the displacement scale field 2262B.

25 Vector memory instructions perform vector loads from and vector stores to memory, with conversion support. As with regular vector instructions, vector memory instructions transfer data from/to memory in a data element-wise fashion, with the elements that are actually transferred is dictated by the contents of the vector mask that is selected as the write mask.

Temporal data is data likely to be reused soon enough to benefit from caching. This is,

30 however, a hint, and different processors may implement it in different ways, including ignoring the hint entirely. Non-temporal data is data unlikely to be reused soon enough to benefit from caching in the 1st-level cache and should be given priority for eviction. This is, however, a hint, and different processors may implement it in different ways, including ignoring the hint entirely.

In the case of the instruction templates of class B, the alpha field 2252 is interpreted as a

write mask control (Z) field 2252C, whose content distinguishes whether the write masking controlled by the write mask field 2270 should be a merging or a zeroing.

In the case of the non-memory access 2205 instruction templates of class B, part of the beta field 2254 is interpreted as an RL field 2257A, whose content distinguishes which one of the different augmentation operation types are to be performed (e.g., round 2257A.1 and vector length (VSIZE) 2257A.2 are respectively specified for the no memory access, write mask control, partial round control type operation 2212 instruction template and the no memory access, write mask control, VSIZE type operation 2217 instruction template), while the rest of the beta field 2254 distinguishes which of the operations of the specified type is to be performed.

In the no memory access 2205 instruction templates, the scale field 2260, the displacement field 2262A, and the displacement scale field 2262B are not present.

In the no memory access, write mask control, partial round control type operation 2210 instruction template, the rest of the beta field 2254 is interpreted as a round operation field 2259A and exception event reporting is disabled (a given instruction does not report any kind of floating-point exception flag and does not raise any floating point exception handler).

Round operation control field 2259A – just as round operation control field 2258, its content distinguishes which one of a group of rounding operations to perform (e.g., Round-up, Round-down, Round-towards-zero and Round-to-nearest). Thus, the round operation control field 2259A allows for the changing of the rounding mode on a per instruction basis. In one embodiment of the invention where a processor includes a control register for specifying rounding modes, the round operation control field's 2250 content overrides that register value.

In the no memory access, write mask control, VSIZE type operation 2217 instruction template, the rest of the beta field 2254 is interpreted as a vector length field 2259B, whose content distinguishes which one of a number of data vector lengths is to be performed on (e.g., 128, 256, or 512 byte).

In the case of a memory access 2220 instruction template of class B, part of the beta field 2254 is interpreted as a broadcast field 2257B, whose content distinguishes whether or not the broadcast type data manipulation operation is to be performed, while the rest of the beta field 2254 is interpreted the vector length field 2259B. The memory access 2220 instruction templates include the scale field 2260, and optionally the displacement field 2262A or the displacement scale field 2262B.

With regard to the generic vector friendly instruction format 2200, a full opcode field 2274 is shown including the format field 2240, the base operation field 2242, and the data element width field 2264. While one embodiment is shown where the full opcode field 2274 includes all

of these fields, the full opcode field 2274 includes less than all of these fields in embodiments that do not support all of them. The full opcode field 2274 provides the operation code (opcode).

The augmentation operation field 2250, the data element width field 2264, and the write mask field 2270 allow these features to be specified on a per instruction basis in the generic vector friendly instruction format. The combination of write mask field and data element width field create typed instructions in that they allow the mask to be applied based on different data element widths.

The various instruction templates found within class A and class B are beneficial in different situations. In some embodiments of the invention, different processors or different cores within a processor may support only class A, only class B, or both classes. For instance, a high performance general purpose out-of-order core intended for general-purpose computing may support only class B, a core intended primarily for graphics and/or scientific (throughput) computing may support only class A, and a core intended for both may support both (of course, a core that has some mix of templates and instructions from both classes but not all templates and instructions from both classes is within the purview of the invention). Also, a single processor may include multiple cores, all of which support the same class or in which different cores support different class. For instance, in a processor with separate graphics and general purpose cores, one of the graphics cores intended primarily for graphics and/or scientific computing may support only class A, while one or more of the general purpose cores may be high performance general purpose cores with out of order execution and register renaming intended for general-purpose computing that support only class B. Another processor that does not have a separate graphics core, may include one more general purpose in-order or out-of-order cores that support both class A and class B. Of course, features from one class may also be implemented in the other class in different embodiments of the invention. Programs written in a high level language would be put (e.g., just in time compiled or statically compiled) into an variety of different executable forms, including: 1) a form having only instructions of the class(es) supported by the target processor for execution; or 2) a form having alternative routines written using different combinations of the instructions of all classes and having control flow code that selects the routines to execute based on the instructions supported by the processor which is currently executing the code.

Figure 11 is a block diagram illustrating an exemplary specific vector friendly instruction format according to embodiments of the invention. Figure 11 shows a specific vector friendly instruction format 2300 that is specific in the sense that it specifies the location, size, interpretation, and order of the fields, as well as values for some of those fields. The specific

vector friendly instruction format 2300 may be used to extend the x86 instruction set, and thus some of the fields are similar or the same as those used in the existing x86 instruction set and extension thereof (e.g., AVX). This format remains consistent with the prefix encoding field, real opcode byte field, MOD R/M field, SIB field, displacement field, and immediate fields of the existing x86 instruction set with extensions. The fields from Figure 10 into which the fields from Figure 11 map are illustrated.

It should be understood that, although embodiments of the invention are described with reference to the specific vector friendly instruction format 2300 in the context of the generic vector friendly instruction format 2200 for illustrative purposes, the invention is not limited to the specific vector friendly instruction format 2300 except where claimed. For example, the generic vector friendly instruction format 2200 contemplates a variety of possible sizes for the various fields, while the specific vector friendly instruction format 2300 is shown as having fields of specific sizes. By way of specific example, while the data element width field 2264 is illustrated as a one bit field in the specific vector friendly instruction format 2300, the invention is not so limited (that is, the generic vector friendly instruction format 2200 contemplates other sizes of the data element width field 2264).

The generic vector friendly instruction format 2200 includes the following fields listed below in the order illustrated in Figure 11A. EVEX Prefix (Bytes 0-3) 2302 - is encoded in a four-byte form. Format Field 2240 (EVEX Byte 0, bits [7:0]) - the first byte (EVEX Byte 0) is the format field 2240 and it contains 0x62 (the unique value used for distinguishing the vector friendly instruction format in one embodiment of the invention). The second-fourth bytes (EVEX Bytes 1-3) include a number of bit fields providing specific capability.

REX field 2305 (EVEX Byte 1, bits [7-5]) – consists of a EVEX.R bit field (EVEX Byte 1, bit [7] – R), EVEX.X bit field (EVEX byte 1, bit [6] – X), and 2257BEX byte 1, bit[5] – B). The EVEX.R, EVEX.X, and EVEX.B bit fields provide the same functionality as the corresponding VEX bit fields, and are encoded using 1s complement form, i.e. ZMM0 is encoded as 1111B, ZMM15 is encoded as 0000B. Other fields of the instructions encode the lower three bits of the register indexes as is known in the art (rrr, xxx, and bbb), so that Rrrr, Xxxx, and Bbbb may be formed by adding EVEX.R, EVEX.X, and EVEX.B.

REX' field 2210 – this is the first part of the REX' field 2210 and is the EVEX.R' bit field (EVEX Byte 1, bit [4] - R') that is used to encode either the upper 16 or lower 16 of the extended 32 register set. In one embodiment of the invention, this bit, along with others as indicated below, is stored in bit inverted format to distinguish (in the well-known x86 32-bit mode) from the BOUND instruction, whose real opcode byte is 62, but does not accept in the MOD R/M

field (described below) the value of 11 in the MOD field; alternative embodiments of the invention do not store this and the other indicated bits below in the inverted format. A value of 1 is used to encode the lower 16 registers. In other words, R'Rrrr is formed by combining EVEX.R', EVEX.R, and the other RRR from other fields.

- 5 Opcode map field 2315 (EVEX byte 1, bits [3:0] – mmmm) – its content encodes an implied leading opcode byte (0F, 0F 38, or 0F 3). Data element width field 2264 (EVEX byte 2, bit [7] – W) - is represented by the notation EVEX.W. EVEX.W is used to define the granularity (size) of the datatype (either 32-bit data elements or 64-bit data elements). EVEX.vvvv 2320 (EVEX Byte 2, bits [6:3]-vvvv)- the role of EVEX.vvvv may include the following: 1) EVEX.vvvv encodes the first source register operand, specified in inverted (1s complement) form and is valid for instructions with 2 or more source operands; 2) EVEX.vvvv encodes the destination register operand, specified in 1s complement form for certain vector shifts; or 3) EVEX.vvvv does not encode any operand, the field is reserved and should contain 1111b. Thus, EVEX.vvvv field 2320 encodes the 4 low-order bits of the first source register specifier stored in inverted (1s complement) form. Depending on the instruction, an extra different EVEX bit field is used to extend the specifier size to 32 registers. EVEX.U 2268 Class field (EVEX byte 2, bit [2]-U) – If EVEX.U = 0, it indicates class A or EVEX.U0; if EVEX.U = 1, it indicates class B or EVEX.U1.

- Prefix encoding field 2325 (EVEX byte 2, bits [1:0]-pp) – provides additional bits for the base operation field. In addition to providing support for the legacy SSE instructions in the EVEX prefix format, this also has the benefit of compacting the SIMD prefix (rather than requiring a byte to express the SIMD prefix, the EVEX prefix requires only 2 bits). In one embodiment, to support legacy SSE instructions that use a SIMD prefix (66H, F2H, F3H) in both the legacy format and in the EVEX prefix format, these legacy SIMD prefixes are encoded into the SIMD prefix encoding field; and at runtime are expanded into the legacy SIMD prefix prior to being provided to the decoder's PLA (so the PLA can execute both the legacy and EVEX format of these legacy instructions without modification). Although newer instructions could use the EVEX prefix encoding field's content directly as an opcode extension, certain embodiments expand in a similar fashion for consistency but allow for different meanings to be specified by these legacy SIMD prefixes. An alternative embodiment may redesign the PLA to support the 2 bit SIMD prefix encodings, and thus not require the expansion.

Alpha field 2252 (EVEX byte 3, bit [7] – EH; also known as EVEX.EH, EVEX.rs, EVEX.RL, EVEX.write mask control, and EVEX.N; also illustrated with α) – as previously described, this field is context specific. Beta field 2254 (EVEX byte 3, bits [6:4]-SSS, also

known as EVEX.s2-0, EVEX.r2-0, EVEX.rr1, EVEX.LL0, EVEX.LLB; also illustrated with $\beta\beta\beta$) – as previously described, this field is context specific.

REX' field 2210 – this is the remainder of the REX' field and is the EVEX.V' bit field (EVEX Byte 3, bit [3] - V') that may be used to encode either the upper 16 or lower 16 of the extended 32 register set. This bit is stored in bit inverted format. A value of 1 is used to encode the lower 16 registers. In other words, V'VVVV is formed by combining EVEX.V', EVEX.vvvv.

Write mask field 2270 (EVEX byte 3, bits [2:0]-kkk) – its content specifies the index of a register in the write mask registers as previously described. In one embodiment of the invention, the specific value EVEX.kkk=000 has a special behavior implying no write mask is used for the particular instruction (this may be implemented in a variety of ways including the use of a write mask hardwired to all ones or hardware that bypasses the masking hardware).

Real Opcode Field 2330 (Byte 4) is also known as the opcode byte. Part of the opcode is specified in this field. MOD R/M Field 2340 (Byte 5) includes MOD field 2342, Reg field 2344, and R/M field 2346. As previously described, the MOD field's 2342 content distinguishes between memory access and non-memory access operations. The role of Reg field 2344 can be summarized to two situations: encoding either the destination register operand or a source register operand, or be treated as an opcode extension and not used to encode any instruction operand. The role of R/M field 2346 may include the following: encoding the instruction operand that references a memory address, or encoding either the destination register operand or a source register operand.

Scale, Index, Base (SIB) Byte (Byte 6) - As previously described, the scale field's 2250 content is used for memory address generation. SIB.xxx 2354 and SIB.bbb 2356 – the contents of these fields have been previously referred to with regard to the register indexes Xxxx and Bbbb. Displacement field 2262A (Bytes 7-10) – when MOD field 2342 contains 10, bytes 7-10 are the displacement field 2262A, and it works the same as the legacy 32-bit displacement (disp32) and works at byte granularity.

Displacement factor field 2262B (Byte 7) – when MOD field 2342 contains 01, byte 7 is the displacement factor field 2262B. The location of this field is that same as that of the legacy x86 instruction set 8-bit displacement (disp8), which works at byte granularity. Since disp8 is sign extended, it can only address between -128 and 127 bytes offsets; in terms of 64 byte cache lines, disp8 uses 8 bits that can be set to only four really useful values -128, -64, 0, and 64; since a greater range is often needed, disp32 is used; however, disp32 requires 4 bytes. In contrast to disp8 and disp32, the displacement factor field 2262B is a reinterpretation of disp8; when using

displacement factor field 2262B, the actual displacement is determined by the content of the displacement factor field multiplied by the size of the memory operand access (N). This type of displacement is referred to as $\text{disp8} \times N$. This reduces the average instruction length (a single byte of used for the displacement but with a much greater range). Such compressed displacement is based on the assumption that the effective displacement is multiple of the granularity of the memory access, and hence, the redundant low-order bits of the address offset do not need to be encoded. In other words, the displacement factor field 2262B substitutes the legacy x86 instruction set 8-bit displacement. Thus, the displacement factor field 2262B is encoded the same way as an x86 instruction set 8-bit displacement (so no changes in the ModRM/SIB encoding rules) with the only exception that disp8 is overloaded to $\text{disp8} \times N$. In other words, there are no changes in the encoding rules or encoding lengths but only in the interpretation of the displacement value by hardware (which needs to scale the displacement by the size of the memory operand to obtain a byte-wise address offset). Immediate field 2272 operates as previously described.

Figure 11B is a block diagram illustrating the fields of the specific vector friendly instruction format 2300 that make up the full opcode field 2274 according to one embodiment of the invention. Specifically, the full opcode field 2274 includes the format field 2240, the base operation field 2242, and the data element width (W) field 2264. The base operation field 2242 includes the prefix encoding field 2325, the opcode map field 2315, and the real opcode field 2330.

Figure 11C is a block diagram illustrating the fields of the specific vector friendly instruction format 2300 that make up the register index field 2244 according to one embodiment of the invention. Specifically, the register index field 2244 includes the REX field 2305, the REX' field 2310, the MODR/M.reg field 2344, the MODR/M.r/m field 2346, the VVVV field 2320, xxx field 2354, and the bbb field 2356.

Figure 11D is a block diagram illustrating the fields of the specific vector friendly instruction format 2300 that make up the augmentation operation field 2250 according to one embodiment of the invention. When the class (U) field 2268 contains 0, it signifies EVEX.U0 (class A 2268A); when it contains 1, it signifies EVEX.U1 (class B 2268B). When $U=0$ and the MOD field 2342 contains 11 (signifying a no memory access operation), the alpha field 2252 (EVEX byte 3, bit [7] – EH) is interpreted as the rs field 2252A. When the rs field 2252A contains a 1 (round 2252A.1), the beta field 2254 (EVEX byte 3, bits [6:4] – SSS) is interpreted as the round control field 2254A. The round control field 2254A includes a one bit SAE field 2256 and a two bit round operation field 2258. When the rs field 2252A contains a 0 (data

transform 2252A.2), the beta field 2254 (EVEX byte 3, bits [6:4]- SSS) is interpreted as a three bit data transform field 2254B. When U=0 and the MOD field 2342 contains 00, 01, or 10 (signifying a memory access operation), the alpha field 2252 (EVEX byte 3, bit [7] – EH) is interpreted as the eviction hint (EH) field 2252B and the beta field 2254 (EVEX byte 3, bits [6:4]- SSS) is interpreted as a three bit data manipulation field 2254C.

When U=1, the alpha field 2252 (EVEX byte 3, bit [7] – EH) is interpreted as the write mask control (Z) field 2252C. When U=1 and the MOD field 2342 contains 11 (signifying a no memory access operation), part of the beta field 2254 (EVEX byte 3, bit [4]- S0) is interpreted as the RL field 2257A; when it contains a 1 (round 2257A.1) the rest of the beta field 2254 (EVEX byte 3, bit [6-5]- S2-1) is interpreted as the round operation field 2259A, while when the RL field 2257A contains a 0 (VSIZE 2257.A2) the rest of the beta field 2254 (EVEX byte 3, bit [6-5]- S2-1) is interpreted as the vector length field 2259B (EVEX byte 3, bit [6-5]- L1-0). When U=1 and the MOD field 2342 contains 00, 01, or 10 (signifying a memory access operation), the beta field 2254 (EVEX byte 3, bits [6:4]- SSS) is interpreted as the vector length field 2259B (EVEX byte 3, bit [6-5]- L1-0) and the broadcast field 2257B (EVEX byte 3, bit [4]- B).

Figure 12 is a block diagram of a register architecture 2400 according to one embodiment of the invention. In the embodiment illustrated, there are 32 vector registers 2410 that are 512 bits wide; these registers are referenced as zmm0 through zmm31. The lower order 256 bits of the lower 16 zmm registers are overlaid on registers ymm0-16. The lower order 128 bits of the lower 16 zmm registers (the lower order 128 bits of the ymm registers) are overlaid on registers xmm0-15. The specific vector friendly instruction format 2300 operates on these overlaid register file as illustrated in the below tables.

Adjustable Vector Length	Class	Operations	Registers
Instruction Templates that do not include the vector length field 2259B	A (Figure 10A; U=0)	2210, 2215, 2225, 2230	zmm registers (the vector length is 64 byte)
	B (Figure 10B; U=1)	2212	zmm registers (the vector length is 64 byte)

Instruction Templates that do include the vector length field 2259B	B (Figure 10B; U=1)	2217, 2227	zmm, ymm, or xmm registers (the vector length is 64 byte, 32 byte, or 16 byte) depending on the vector length field 2259B
---	------------------------	------------	---

In other words, the vector length field 2259B selects between a maximum length and one or more other shorter lengths, where each such shorter length is half the length of the preceding length; and instructions templates without the vector length field 2259B operate on the maximum vector length. Further, in one embodiment, the class B instruction templates of the specific vector friendly instruction format 2300 operate on packed or scalar single/double-precision floating point data and packed or scalar integer data. Scalar operations are operations performed on the lowest order data element position in an zmm/ymm/xmm register; the higher order data element positions are either left the same as they were prior to the instruction or zeroed depending on the embodiment.

Write mask registers 2415 - in the embodiment illustrated, there are 8 write mask registers (k0 through k7), each 64 bits in size. In an alternate embodiment, the write mask registers 2415 are 16 bits in size. As previously described, in one embodiment of the invention, the vector mask register k0 cannot be used as a write mask; when the encoding that would normally indicate k0 is used for a write mask, it selects a hardwired write mask of 0xFFFF, effectively disabling write masking for that instruction.

General-purpose registers 2425 - in the embodiment illustrated, there are sixteen 64-bit general-purpose registers that are used along with the existing x86 addressing modes to address memory operands. These registers are referenced by the names RAX, RBX, RCX, RDX, RBP, RSI, RDI, RSP, and R8 through R15.

Scalar floating point stack register file (x87 stack) 2445, on which is aliased the MMX packed integer flat register file 2450 - in the embodiment illustrated, the x87 stack is an eight-element stack used to perform scalar floating-point operations on 32/64/80-bit floating point data using the x87 instruction set extension; while the MMX registers are used to perform operations on 64-bit packed integer data, as well as to hold operands for some operations performed

between the MMX and XMM registers.

Alternative embodiments of the invention may use wider or narrower registers. Additionally, alternative embodiments of the invention may use more, less, or different register files and registers.

5 Processor cores may be implemented in different ways, for different purposes, and in different processors. For instance, implementations of such cores may include: 1) a general purpose in-order core intended for general-purpose computing; 2) a high performance general purpose out-of-order core intended for general-purpose computing; 3) a special purpose core intended primarily for graphics and/or scientific (throughput) computing. Implementations of
10 different processors may include: 1) a CPU including one or more general purpose in-order cores intended for general-purpose computing and/or one or more general purpose out-of-order cores intended for general-purpose computing; and 2) a coprocessor including one or more special purpose cores intended primarily for graphics and/or scientific (throughput). Such different processors lead to different computer system architectures, which may include: 1) the
15 coprocessor on a separate chip from the CPU; 2) the coprocessor on a separate die in the same package as a CPU; 3) the coprocessor on the same die as a CPU (in which case, such a coprocessor is sometimes referred to as special purpose logic, such as integrated graphics and/or scientific (throughput) logic, or as special purpose cores); and 4) a system on a chip that may include on the same die the described CPU (sometimes referred to as the application core(s) or
20 application processor(s)), the above described coprocessor, and additional functionality. Exemplary core architectures are described next, followed by descriptions of exemplary processors and computer architectures.

Figure 13A is a block diagram illustrating both an exemplary in-order pipeline and an exemplary register renaming, out-of-order issue/execution pipeline according to embodiments of
25 the invention. Figure 13B is a block diagram illustrating both an exemplary embodiment of an in-order architecture core and an exemplary register renaming, out-of-order issue/execution architecture core to be included in a processor according to embodiments of the invention. The solid lined boxes illustrate the in-order pipeline and in-order core, while the optional addition of the dashed lined boxes illustrates the register renaming, out-of-order issue/execution pipeline and
30 core. Given that the in-order aspect is a subset of the out-of-order aspect, the out-of-order aspect will be described.

In Figure 13A, a processor pipeline 2500 includes a fetch stage 2502, a length decode stage 2504, a decode stage 2506, an allocation stage 2508, a renaming stage 2510, a scheduling (also known as a dispatch or issue) stage 2512, a register read/memory read stage 2514, an execute

stage 2516, a write back/memory write stage 2518, an exception handling stage 2522, and a commit stage 2524.

Figure 13B shows processor core 2590 including a front end unit 2530 coupled to an execution engine unit 2550, and both are coupled to a memory unit 2570. The core 2590 may be a reduced instruction set computing (RISC) core, a complex instruction set computing (CISC) core, a very long instruction word (VLIW) core, or a hybrid or alternative core type. As yet another option, the core 2590 may be a special-purpose core, such as, for example, a network or communication core, compression engine, coprocessor core, general purpose computing graphics processing unit (GPGPU) core, graphics core, or the like.

The front end unit 2530 includes a branch prediction unit 2532 coupled to an instruction cache unit 2534, which is coupled to an instruction translation lookaside buffer (TLB) 2536, which is coupled to an instruction fetch unit 2538, which is coupled to a decode unit 2540. The decode unit 2540 (or decoder) may decode instructions, and generate as an output one or more micro-operations, micro-code entry points, microinstructions, other instructions, or other control signals, which are decoded from, or which otherwise reflect, or are derived from, the original instructions. The decode unit 2540 may be implemented using various different mechanisms. Examples of suitable mechanisms include, but are not limited to, look-up tables, hardware implementations, programmable logic arrays (PLAs), microcode read only memories (ROMs), etc. In one embodiment, the core 2590 includes a microcode ROM or other medium that stores microcode for certain macroinstructions (e.g., in decode unit 2540 or otherwise within the front end unit 2530). The decode unit 2540 is coupled to a rename/allocator unit 2552 in the execution engine unit 2550.

The execution engine unit 2550 includes the rename/allocator unit 2552 coupled to a retirement unit 2554 and a set of one or more scheduler unit(s) 2556. The scheduler unit(s) 2556 represents any number of different schedulers, including reservations stations, central instruction window, etc. The scheduler unit(s) 2556 is coupled to the physical register file(s) unit(s) 2558. Each of the physical register file(s) units 2558 represents one or more physical register files, different ones of which store one or more different data types, such as scalar integer, scalar floating point, packed integer, packed floating point, vector integer, vector floating point, status (e.g., an instruction pointer that is the address of the next instruction to be executed), etc.

In one embodiment, the physical register file(s) unit 2558 comprises a vector registers unit, a write mask registers unit, and a scalar registers unit. These register units may provide architectural vector registers, vector mask registers, and general purpose registers. The physical register file(s) unit(s) 2558 is overlapped by the retirement unit 2554 to illustrate various ways in

which register renaming and out-of-order execution may be implemented (e.g., using a reorder buffer(s) and a retirement register file(s); using a future file(s), a history buffer(s), and a retirement register file(s); using a register maps and a pool of registers; etc.). The retirement unit 2554 and the physical register file(s) unit(s) 2558 are coupled to the execution cluster(s) 2560.

5 The execution cluster(s) 2560 includes a set of one or more execution units 2562 and a set of one or more memory access units 2564. The execution units 2562 may perform various operations (e.g., shifts, addition, subtraction, multiplication) and on various types of data (e.g., scalar floating point, packed integer, packed floating point, vector integer, vector floating point). While some embodiments may include a number of execution units dedicated to specific
10 functions or sets of functions, other embodiments may include only one execution unit or multiple execution units that all perform all functions.

The scheduler unit(s) 2556, physical register file(s) unit(s) 2558, and execution cluster(s) 2560 are shown as being possibly plural because certain embodiments create separate pipelines for certain types of data/operations (e.g., a scalar integer pipeline, a scalar floating point/packed
15 integer/packed floating point/vector integer/vector floating point pipeline, and/or a memory access pipeline that each have their own scheduler unit, physical register file(s) unit, and/or execution cluster – and in the case of a separate memory access pipeline, certain embodiments are implemented in which only the execution cluster of this pipeline has the memory access unit(s) 2564). It should also be understood that where separate pipelines are used, one or more
20 of these pipelines may be out-of-order issue/execution and the rest in-order.

The set of memory access units 2564 is coupled to the memory unit 2570, which includes a data TLB unit 2572 coupled to a data cache unit 2574 coupled to a level 2 (L2) cache unit 2576. In one exemplary embodiment, the memory access units 2564 may include a load unit, a store address unit, and a store data unit, each of which is coupled to the data TLB unit 2572 in the
25 memory unit 2570. The instruction cache unit 2534 is further coupled to a level 2 (L2) cache unit 2576 in the memory unit 2570. The L2 cache unit 2576 is coupled to one or more other levels of cache and eventually to a main memory.

By way of example, the exemplary register renaming, out-of-order issue/execution core architecture may implement the pipeline 2500 as follows: 1) the instruction fetch 2538 performs
30 the fetch and length decoding stages 2502 and 2504; 2) the decode unit 2540 performs the decode stage 2506; 3) the rename/allocator unit 2552 performs the allocation stage 2508 and renaming stage 2510; 4) the scheduler unit(s) 2556 performs the schedule stage 2512; 5) the physical register file(s) unit(s) 2558 and the memory unit 2570 perform the register read/memory read stage 2514; the execution cluster 2560 perform the execute stage 2516; 6) the memory unit

2570 and the physical register file(s) unit(s) 2558 perform the write back/memory write stage 2518; 7) various units may be involved in the exception handling stage 2522; and 8) the retirement unit 2554 and the physical register file(s) unit(s) 2558 perform the commit stage 2524.

5 The core 2590 may support one or more instructions sets (e.g., the x86 instruction set (with some extensions that have been added with newer versions); the MIPS instruction set of MIPS Technologies of Sunnyvale, CA; the ARM instruction set (with optional additional extensions such as NEON) of ARM Holdings of Sunnyvale, CA), including the instruction(s) described herein. In one embodiment, the core 2590 includes logic to support a packed data instruction set extension (e.g., AVX1, AVX2, and/or some form of the generic vector friendly instruction
10 format (U=0 and/or U=1) previously described), thereby allowing the operations used by many multimedia applications to be performed using packed data.

It should be understood that the core may support multithreading (executing two or more parallel sets of operations or threads), and may do so in a variety of ways including time sliced multithreading, simultaneous multithreading (where a single physical core provides a logical
15 core for each of the threads that physical core is simultaneously multithreading), or a combination thereof (e.g., time sliced fetching and decoding and simultaneous multithreading thereafter such as in the Intel® Hyperthreading technology).

While register renaming is described in the context of out-of-order execution, it should be understood that register renaming may be used in an in-order architecture. While the illustrated
20 embodiment of the processor also includes separate instruction and data cache units 2534/2574 and a shared L2 cache unit 2576, alternative embodiments may have a single internal cache for both instructions and data, such as, for example, a Level 1 (L1) internal cache, or multiple levels of internal cache. In some embodiments, the system may include a combination of an internal cache and an external cache that is external to the core and/or the processor. Alternatively, all of
25 the cache may be external to the core and/or the processor.

Figure 14A and Figure 14B illustrate a block diagram of a more specific exemplary in-order core architecture, which core would be one of several logic blocks (including other cores of the same type and/or different types) in a chip. The logic blocks communicate through a high-bandwidth interconnect network (e.g., a ring network) with some fixed function logic, memory
30 I/O interfaces, and other necessary I/O logic, depending on the application.

Figure 14A is a block diagram of a single processor core, along with its connection to the on-die interconnect network 2602 and with its local subset of the Level 2 (L2) cache 2604, according to embodiments of the invention. In one embodiment, an instruction decoder 2600 supports the x86 instruction set with a packed data instruction set extension. An L1 cache 2606

allows low-latency accesses to cache memory into the scalar and vector units. While in one embodiment (to simplify the design), a scalar unit 2608 and a vector unit 2610 use separate register sets (respectively, scalar registers 2612 and vector registers 2614) and data transferred between them is written to memory and then read back in from a level 1 (L1) cache 2606,

5 alternative embodiments of the invention may use a different approach (e.g., use a single register set or include a communication path that allow data to be transferred between the two register files without being written and read back).

The local subset of the L2 cache 2604 is part of a global L2 cache that is divided into separate local subsets, one per processor core. Each processor core has a direct access path to its
10 own local subset of the L2 cache 2604. Data read by a processor core is stored in its L2 cache subset 2604 and can be accessed quickly, in parallel with other processor cores accessing their own local L2 cache subsets. Data written by a processor core is stored in its own L2 cache subset 2604 and is flushed from other subsets, if necessary. The ring network ensures coherency for shared data. The ring network is bi-directional to allow agents such as processor cores, L2
15 caches and other logic blocks to communicate with each other within the chip. Each ring data-path is 1012-bits wide per direction.

Figure 14B is an expanded view of part of the processor core in Figure 14A according to embodiments of the invention. Figure 14B includes an L1 data cache 2606A part of the L1 cache 2604, as well as more detail regarding the vector unit 2610 and the vector registers 2614.
20 Specifically, the vector unit 2610 is a 16-wide vector processing unit (VPU) (see the 16-wide ALU 2628), which executes one or more of integer, single-precision float, and double-precision float instructions. The VPU supports swizzling the register inputs with swizzle unit 2620, numeric conversion with numeric convert units 2622A-B, and replication with replication unit 2624 on the memory input. Write mask registers 2626 allow predicating resulting vector writes.

Figure 15 is a block diagram of a processor 2700 that may have more than one core, may have an integrated memory controller, and may have integrated graphics according to embodiments of the invention. The solid lined boxes in Figure 15 illustrate a processor 2700 with a single core 2702A, a system agent 2710, a set of one or more bus controller units 2716, while the optional addition of the dashed lined boxes illustrates an alternative processor 2700
30 with multiple cores 2702A-N, a set of one or more integrated memory controller unit(s) 2714 in the system agent unit 2710, and special purpose logic 2708.

Thus, different implementations of the processor 2700 may include: 1) a CPU with the special purpose logic 2708 being integrated graphics and/or scientific (throughput) logic (which may include one or more cores), and the cores 2702A-N being one or more general purpose

cores (e.g., general purpose in-order cores, general purpose out-of-order cores, a combination of the two); 2) a coprocessor with the cores 2702A-N being a large number of special purpose cores intended primarily for graphics and/or scientific (throughput); and 3) a coprocessor with the cores 2702A-N being a large number of general purpose in-order cores. Thus, the processor 2700
5 may be a general-purpose processor, coprocessor or special-purpose processor, such as, for example, a network or communication processor, compression engine, graphics processor, GPGPU (general purpose graphics processing unit), a high-throughput many integrated core (MIC) coprocessor (including 30 or more cores), embedded processor, or the like. The processor may be implemented on one or more chips. The processor 2700 may be a part of and/or may be
10 implemented on one or more substrates using any of a number of process technologies, such as, for example, BiCMOS, CMOS, or NMOS.

The memory hierarchy includes one or more levels of cache within the cores, a set or one or more shared cache units 2706, and external memory (not shown) coupled to the set of integrated memory controller units 2714. The set of shared cache units 2706 may include one or
15 more mid-level caches, such as level 2 (L2), level 3 (L3), level 4 (L4), or other levels of cache, a last level cache (LLC), and/or combinations thereof. While in one embodiment a ring based interconnect unit 2712 interconnects the integrated graphics logic 2708, the set of shared cache units 2706, and the system agent unit 2710/integrated memory controller unit(s) 2714, alternative embodiments may use any number of well-known techniques for interconnecting such units. In
20 one embodiment, coherency is maintained between one or more cache units 2706 and cores 2702-A-N.

In some embodiments, one or more of the cores 2702A-N are capable of multi-threading. The system agent 2710 includes those components coordinating and operating cores 2702A-N. The system agent unit 2710 may include for example a power control unit (PCU) and a display
25 unit. The PCU may be or include logic and components needed for regulating the power state of the cores 2702A-N and the integrated graphics logic 2708. The display unit is for driving one or more externally connected displays.

The cores 2702A-N may be homogenous or heterogeneous in terms of architecture instruction set; that is, two or more of the cores 2702A-N may be capable of execution the same
30 instruction set, while others may be capable of executing only a subset of that instruction set or a different instruction set.

Figure 16 to Figure 20 are block diagrams of exemplary computer architectures. Other system designs and configurations known in the arts for laptops, desktops, handheld PCs, personal digital assistants, engineering workstations, servers, network devices, network hubs,

switches, embedded processors, digital signal processors (DSPs), graphics devices, video game devices, set-top boxes, micro controllers, cell phones, portable media players, hand held devices, and various other electronic devices, are also suitable. In general, a huge variety of systems or electronic devices capable of incorporating a processor and/or other execution logic as disclosed

5 herein are generally suitable.

Referring now to Figure 16, shown is a block diagram of a system 2800 in accordance with one embodiment of the present invention. The system 2800 may include one or more processors 2810, 2815, which are coupled to a controller hub 2820. In one embodiment the controller hub 2820 includes a graphics memory controller hub (GMCH) 2890 and an Input/Output Hub (IOH) 2850 (which may be on separate chips); the GMCH 2890 includes memory and graphics controllers to which are coupled memory 2840 and a coprocessor 2845; the IOH 2850 is couples input/output (I/O) devices 2860 to the GMCH 2890. Alternatively, one or both of the memory and graphics controllers are integrated within the processor (as described herein), the memory 2840 and the coprocessor 2845 are coupled directly to the processor 2810, and the controller hub 2820 in a single chip with the IOH 2850.

The optional nature of additional processors 2815 is denoted in Figure 16 with broken lines. Each processor 2810, 2815 may include one or more of the processing cores described herein and may be some version of the processor 2700.

The memory 2840 may be, for example, dynamic random access memory (DRAM), phase change memory (PCM), or a combination of the two. For at least one embodiment, the controller hub 2820 communicates with the processor(s) 2810, 2815 via a multi-drop bus, such as a frontside bus (FSB), point-to-point interface such as QuickPath Interconnect (QPI), or similar connection 2895.

In one embodiment, the coprocessor 2845 is a special-purpose processor, such as, for example, a high-throughput MIC processor, a network or communication processor, compression engine, graphics processor, GPGPU, embedded processor, or the like. In one embodiment, controller hub 2820 may include an integrated graphics accelerator.

There can be a variety of differences between the physical resources 2810, 2815 in terms of a spectrum of metrics of merit including architectural, microarchitectural, thermal, power consumption characteristics, and the like.

In one embodiment, the processor 2810 executes instructions that control data processing operations of a general type. Embedded within the instructions may be coprocessor instructions. The processor 2810 recognizes these coprocessor instructions as being of a type that should be executed by the attached coprocessor 2845. Accordingly, the processor 2810 issues these

coprocessor instructions (or control signals representing coprocessor instructions) on a coprocessor bus or other interconnect, to coprocessor 2845. Coprocessor(s) 2845 accept and execute the received coprocessor instructions.

Referring now to Figure 17, shown is a block diagram of a first more specific exemplary system 2900 in accordance with an embodiment of the present invention. As shown in Figure 17, multiprocessor system 2900 is a point-to-point interconnect system, and includes a first processor 2970 and a second processor 2980 coupled via a point-to-point interconnect 2950. Each of processors 2970 and 2980 may be some version of the processor 2700. In one embodiment of the invention, processors 2970 and 2980 are respectively processors 2810 and 2815, while coprocessor 2938 is coprocessor 2845. In another embodiment, processors 2970 and 2980 are respectively processor 2810 coprocessor 2845.

Processors 2970 and 2980 are shown including integrated memory controller (IMC) units 2972 and 2982, respectively. Processor 2970 also includes as part of its bus controller units point-to-point (P-P) interfaces 2976 and 2978; similarly, second processor 2980 includes P-P interfaces 2986 and 2988. Processors 2970, 2980 may exchange information via a point-to-point (P-P) interface 2950 using P-P interface circuits 2978, 2988. As shown in Figure 17, IMCs 2972 and 2982 couple the processors to respective memories, namely a memory 2932 and a memory 2934, which may be portions of main memory locally attached to the respective processors.

Processors 2970, 2980 may each exchange information with a chipset 2990 via individual P-P interfaces 2952, 2954 using point to point interface circuits 2976, 2994, 2986, 2998. Chipset 2990 may optionally exchange information with the coprocessor 2938 via a high-performance interface 2939. In one embodiment, the coprocessor 2938 is a special-purpose processor, such as, for example, a high-throughput MIC processor, a network or communication processor, compression engine, graphics processor, GPGPU, embedded processor, or the like.

A shared cache (not shown) may be included in either processor or outside of both processors, yet connected with the processors via P-P interconnect, such that either or both processors' local cache information may be stored in the shared cache if a processor is placed into a low power mode. Chipset 2990 may be coupled to a first bus 2916 via an interface 2996. In one embodiment, first bus 2916 may be a Peripheral Component Interconnect (PCI) bus, or a bus such as a PCI Express bus or another third generation I/O interconnect bus, although the scope of the present invention is not so limited.

As shown in Figure 17, various I/O devices 2914 may be coupled to first bus 2916, along with a bus bridge 2918 which couples first bus 2916 to a second bus 2920. In one embodiment, one or more additional processor(s) 2915, such as coprocessors, high-throughput MIC

processors, GPGPU's, accelerators (such as, e.g., graphics accelerators or digital signal processing (DSP) units), field programmable gate arrays, or any other processor, are coupled to first bus 2916. In one embodiment, second bus 2920 may be a low pin count (LPC) bus.

Various devices may be coupled to a second bus 2920 including, for example, a keyboard and/or mouse 2922, communication devices 2927 and a storage unit 2928 such as a disk drive or other mass storage device which may include instructions/code and data 2930, in one embodiment. Further, an audio I/O 2924 may be coupled to the second bus 2920. Note that other architectures are possible. For example, instead of the point-to-point architecture of Figure 17, a system may implement a multi-drop bus or other such architecture.

Referring now to Figure 18, shown is a block diagram of a second more specific exemplary system 3000 in accordance with an embodiment of the present invention. Like elements in Figure 18 and Figure 19 bear like reference numerals, and certain aspects of Figure 17 have been omitted from Figure 18 in order to avoid obscuring other aspects of Figure 18. Figure 18 illustrates that the processors 2970, 2980 may include integrated memory and I/O control logic ("CL") 2972 and 2982, respectively. Thus, the CL 2972, 2982 include integrated memory controller units and include I/O control logic. Figure 18 illustrates that not only are the memories 2932, 2934 coupled to the CL 2972, 2982, but also that I/O devices 3014 are also coupled to the control logic 2972, 2982. Legacy I/O devices 3015 are coupled to the chipset 2990.

Referring now to Figure 19, shown is a block diagram of a SoC 3100 in accordance with an embodiment of the present invention. Similar elements in Figure 15 bear like reference numerals. Also, dashed lined boxes are optional features on more advanced SoCs. In Figure 19, an interconnect unit(s) 3102 is coupled to: an application processor 3110 which includes a set of one or more cores 202A-N and shared cache unit(s) 2706; a system agent unit 2710; a bus controller unit(s) 2716; an integrated memory controller unit(s) 2714; a set or one or more coprocessors 3120 which may include integrated graphics logic, an image processor, an audio processor, and a video processor; an static random access memory (SRAM) unit 3130; a direct memory access (DMA) unit 3132; and a display unit 3140 for coupling to one or more external displays. In one embodiment, the coprocessor(s) 3120 include a special-purpose processor, such as, for example, a network or communication processor, compression engine, GPGPU, a high-throughput MIC processor, embedded processor, or the like.

Embodiments of the mechanisms disclosed herein may be implemented in hardware, software, firmware, or a combination of such implementation approaches. Embodiments of the invention may be implemented as computer programs or program code executing on

programmable systems comprising at least one processor, a storage system (including volatile and non-volatile memory and/or storage elements), at least one input device, and at least one output device.

Program code, such as code 2930 illustrated in Figure 17, may be applied to input instructions to perform the functions described herein and generate output information. The output information may be applied to one or more output devices, in known fashion. For purposes of this application, a processing system includes any system that has a processor, such as, for example; a digital signal processor (DSP), a microcontroller, an application specific integrated circuit (ASIC), or a microprocessor.

The program code may be implemented in a high level procedural or object oriented programming language to communicate with a processing system. The program code may also be implemented in assembly or machine language, if desired. In fact, the mechanisms described herein are not limited in scope to any particular programming language. In any case, the language may be a compiled or interpreted language.

One or more aspects of at least one embodiment may be implemented by representative instructions stored on a machine-readable medium which represents various logic within the processor, which when read by a machine causes the machine to fabricate logic to perform the techniques described herein. Such representations, known as “IP cores” may be stored on a tangible, machine readable medium and supplied to various customers or manufacturing facilities to load into the fabrication machines that actually make the logic or processor.

Such machine-readable storage media may include, without limitation, non-transitory, tangible arrangements of articles manufactured or formed by a machine or device, including storage media such as hard disks, any other type of disk including floppy disks, optical disks, compact disk read-only memories (CD-ROMs), compact disk rewritable’s (CD-RWs), and magneto-optical disks, semiconductor devices such as read-only memories (ROMs), random access memories (RAMs) such as dynamic random access memories (DRAMs), static random access memories (SRAMs), erasable programmable read-only memories (EPROMs), flash memories, electrically erasable programmable read-only memories (EEPROMs), phase change memory (PCM), magnetic or optical cards, or any other type of media suitable for storing electronic instructions.

Accordingly, embodiments of the invention also include non-transitory, tangible machine-readable media containing instructions or containing design data, such as Hardware Description Language (HDL), which defines structures, circuits, apparatuses, processors and/or system features described herein. Such embodiments may also be referred to as program products.

In some cases, an instruction converter may be used to convert an instruction from a source instruction set to a target instruction set. For example, the instruction converter may translate (e.g., using static binary translation, dynamic binary translation including dynamic compilation), morph, emulate, or otherwise convert an instruction to one or more other instructions to be
5 processed by the core. The instruction converter may be implemented in software, hardware, firmware, or a combination thereof. The instruction converter may be on processor, off processor, or part on and part off processor.

Figure 20 is a block diagram contrasting the use of a software instruction converter to convert binary instructions in a source instruction set to binary instructions in a target instruction
10 set according to embodiments of the invention. In the illustrated embodiment, the instruction converter is a software instruction converter, although alternatively the instruction converter may be implemented in software, firmware, hardware, or various combinations thereof. Figure 20 shows a program in a high level language 3202 may be compiled using an x86 compiler 3204 to generate x86 binary code 3206 that may be natively executed by a processor with at least one
15 x86 instruction set core 3216. The processor with at least one x86 instruction set core 3216 represents any processor that can perform substantially the same functions as an Intel processor with at least one x86 instruction set core by compatibly executing or otherwise processing (1) a substantial portion of the instruction set of the Intel x86 instruction set core or (2) object code versions of applications or other software targeted to run on an Intel processor with at least one
20 x86 instruction set core, in order to achieve substantially the same result as an Intel processor with at least one x86 instruction set core. The x86 compiler 3204 represents a compiler that is operable to generate x86 binary code 3206 (e.g., object code) that can, with or without additional linkage processing, be executed on the processor with at least one x86 instruction set core 3216. Similarly, Figure 20 shows the program in the high level language 3202 may be compiled using
25 an alternative instruction set compiler 3208 to generate alternative instruction set binary code 3210 that may be natively executed by a processor without at least one x86 instruction set core 3214 (e.g., a processor with cores that execute the MIPS instruction set of MIPS Technologies of Sunnyvale, CA and/or that execute the ARM instruction set of ARM Holdings of Sunnyvale, CA). The instruction converter 3212 is used to convert the x86 binary code 3206 into code that
30 may be natively executed by the processor without an x86 instruction set core 3214. This converted code is not likely to be the same as the alternative instruction set binary code 3210 because an instruction converter capable of this is difficult to make; however, the converted code will accomplish the general operation and be made up of instructions from the alternative instruction set. Thus, the instruction converter 3212 represents software, firmware, hardware, or

a combination thereof that, through emulation, simulation or any other process, allows a processor or other electronic device that does not have an x86 instruction set processor or core to execute the x86 binary code 3206.

According to some embodiments, a processor includes a first execution unit to receive a first instruction to process a first part of secure hash algorithm 256 (SHA256) message scheduling operations, the first instruction having a first operand associated with a first storage location to store a first set of message inputs and a second operand associated with a second storage location to store a second set of message inputs, and to execute the first instruction based on the first and second set of message inputs to produce an intermediate result; and a second execution unit to receive a second instruction to process a second part of the SHA256 message scheduling operations, the second instruction having a third operand associated with a third storage location to store the intermediate result and a third set of message inputs and a fourth operand associated with a fourth storage location to store a fourth set of message inputs, and to execute the second instruction based on the intermediate result and the third and fourth sets of message inputs to produce message inputs for next iteration of SHA256 round operations. The first, second, third, and fourth operands refer to storage locations each having at least 128 bits. The intermediate result is stored in the first storage location associated with the first operand. The message inputs for the next iteration of SHA256 round operations are stored in the third storage location associated with the third operand. For a current iteration i of the SHA256 message scheduling operations, the first storage location stores message inputs $w(i-13)$, $w(i-14)$, $w(i-15)$, and $w(i-16)$, and wherein the second storage location stores message inputs $w(i-9)$, $w(i-10)$, $w(i-11)$, and $w(i-12)$. For a current iteration i of the SHA256 message scheduling operations, the intermediate result represents $w(i-13)+s_0(w(i-12))$, $w(i-14)+s_0(w(i-13))$, $w(i-15)+s_0(w(i-14))$, and $w(i-16)+s_0(w(i-15))$. For a current iteration i of the SHA256 message scheduling operations, the third storage location stores a combination of the intermediate result and message inputs $w(i-5)$, $w(i-6)$, and $w(i-7)$. For a current iteration i of the SHA256 message scheduling operations, the message inputs for the next iteration represent $w(i+3)$, $w(i+2)$, $w(i+1)$, and $w(i)$.

Some portions of the preceding detailed descriptions have been presented in terms of algorithms and symbolic representations of operations on data bits within a computer memory. These algorithmic descriptions and representations are the ways used by those skilled in the data processing arts to most effectively convey the substance of their work to others skilled in the art. An algorithm is here, and generally, conceived to be a self-consistent sequence of operations leading to a desired result. The operations are those requiring physical manipulations of physical

quantities.

It should be borne in mind, however, that all of these and similar terms are to be associated with the appropriate physical quantities and are merely convenient labels applied to these quantities. Unless specifically stated otherwise as apparent from the above discussion, it is appreciated that throughout the description, discussions utilizing terms such as those set forth in the claims below, refer to the action and processes of a computer system, or similar electronic computing device, that manipulates and transforms data represented as physical (electronic) quantities within the computer system's registers and memories into other data similarly represented as physical quantities within the computer system memories or registers or other such information storage, transmission or display devices.

The techniques shown in the figures can be implemented using code and data stored and executed on one or more electronic devices. Such electronic devices store and communicate (internally and/or with other electronic devices over a network) code and data using computer-readable media, such as non-transitory computer-readable storage media (e.g., magnetic disks; optical disks; random access memory; read only memory; flash memory devices; phase-change memory) and transitory computer-readable transmission media (e.g., electrical, optical, acoustical or other form of propagated signals – such as carrier waves, infrared signals, digital signals).

The processes or methods depicted in the preceding figures may be performed by processing logic that comprises hardware (e.g. circuitry, dedicated logic, etc.), firmware, software (e.g., embodied on a non-transitory computer readable medium), or a combination of both. Although the processes or methods are described above in terms of some sequential operations, it should be appreciated that some of the operations described may be performed in a different order. Moreover, some operations may be performed in parallel rather than sequentially.

In the foregoing specification, embodiments of the invention have been described with reference to specific exemplary embodiments thereof. It will be evident that various modifications may be made thereto without departing from the broader spirit and scope of the invention as set forth in the following claims. The specification and drawings are, accordingly, to be regarded in an illustrative sense rather than a restrictive sense.

CLAIMS

What is claimed is:

- 5 1. A processor, comprising:
a first execution unit to receive a first instruction to process a first part of secure hash
algorithm 256 (SHA256) message scheduling operations, the first instruction
having a first operand associated with a first storage location to store a first set of
message inputs and a second operand associated with a second storage location to
10 store a second set of message inputs, and to execute the first instruction based on
the first and second set of message inputs to produce an intermediate result; and
a second execution unit to receive a second instruction to process a second part of the
SHA256 message scheduling operations, the second instruction having a third
operand associated with a third storage location to store the intermediate result
15 and a third set of message inputs and a fourth operand associated with a fourth
storage location to store a fourth set of message inputs, and to execute the second
instruction based on the intermediate result and the third and fourth sets of
message inputs to produce message inputs for next iteration of SHA256 round
operations.
20
2. The processor of claim 1, wherein the first, second, third, and fourth operands refer to
storage locations each having at least 128 bits.
3. The processor of claim 1, wherein the intermediate result is stored in the first storage
25 location associated with the first operand.
4. The processor of claim 1, wherein the message inputs for the next iteration of SHA256
round operations are stored in the third storage location associated with the third operand.
- 30 5. The processor of claim 1, wherein for a current iteration i of the SHA256 message
scheduling operations, the first storage location stores message inputs $w(i-13)$, $w(i-14)$, $w(i-15)$,
and $w(i-16)$, and wherein the second storage location stores message inputs $w(i-9)$, $w(i-10)$, $w(i-11)$,
and $w(i-12)$.

6. The processor of claim 1, wherein for a current iteration i of the SHA256 message scheduling operations, the intermediate result represents $w(i-13)+s0(w(i-12))$, $w(i-14)+s0(w(i-13))$, $w(i-15)+s0(w(i-14))$, and $w(i-16)+s0(w(i-15))$.

5 7. The processor of claim 1, wherein for a current iteration i of the SHA256 message scheduling operations, the third storage location stores a combination of the intermediate result and message inputs $w(i-5)$, $w(i-6)$, and $w(i-7)$.

8. The processor of claim 1, wherein for a current iteration i of the SHA256 message
10 scheduling operations, the message inputs for the next iteration represent $w(i+3)$, $w(i+2)$, $w(i+1)$, and $w(i)$.

9. A method, comprising:

15 receiving, by a first execution unit of a processor, a first instruction to process a first part of secure hash algorithm 256 (SHA256) message scheduling operations, the first instruction having a first operand associated with a first storage location to store a first set of message inputs and a second operand associated with a second storage location to store a second set of message inputs;

20 executing, by the first execution unit, the first instruction based on the first and second set of message inputs to produce an intermediate result;

receiving, by a second execution unit of the processor, a second instruction to process a second part of the SHA256 message scheduling operations, the second instruction having a third operand associated with a third storage location to store the
25 intermediate result and a third set of message inputs and a fourth operand associated with a fourth storage location to store a fourth set of message inputs;
and

executing, by the second execution unit, the second instruction based on the intermediate result and the third and fourth sets of message inputs to produce message inputs
30 for next iteration of SHA256 round operations.

10. The method of claim 9, wherein the first, second, third, and fourth operands refer to storage locations each having at least 128 bits.

11. The method of claim 9, wherein the intermediate result is stored in the first storage location associated with the first operand.

12. The method of claim 9, wherein the message inputs for the next iteration of SHA256 round operations are stored in the third storage location associated with the third operand.

13. The method of claim 9, wherein for a current iteration i of the SHA256 message scheduling operations, the first storage location stores message inputs $w(i-13)$, $w(i-14)$, $w(i-15)$, and $w(i-16)$, and wherein the second storage location stores message inputs $w(i-9)$, $w(i-10)$, $w(i-11)$, and $w(i-12)$.

14. The method of claim 9, wherein for a current iteration i of the SHA256 message scheduling operations, the intermediate result represents $w(i-13)+s0(w(i-12))$, $w(i-14)+s0(w(i-13))$, $w(i-15)+s0(w(i-14))$, and $w(i-16)+s0(w(i-15))$.

15. The method of claim 9, wherein for a current iteration i of the SHA256 message scheduling operations, the third storage location stores a combination of the intermediate result and message inputs $w(i-5)$, $w(i-6)$, and $w(i-7)$.

16. The method of claim 9, wherein for a current iteration i of the SHA256 message scheduling operations, the message inputs for the next iteration represent $w(i+3)$, $w(i+2)$, $w(i+1)$, and $w(i)$.

17. A data processing system, comprising:

an interconnect;

a dynamic random access memory (DRAM) coupled to the interconnect; and

a processor coupled the interconnect, the processor including

a first execution unit to receive a first instruction to process a first part of secure

hash algorithm 256 (SHA256) message scheduling operations, the first

instruction having a first operand associated with a first storage location to

store a first set of message inputs and a second operand associated with a

second storage location to store a second set of message inputs, and to

execute the first instruction based on the first and second set of message

inputs to produce an intermediate result, and

a second execution unit to receive a second instruction to process a second part of the SHA256 message scheduling operations, the second instruction having a third operand associated with a third storage location to store the intermediate result and a third set of message inputs and a fourth operand associated with a fourth storage location to store a fourth set of message inputs, and to execute the second instruction based on the intermediate result and the third and fourth sets of message inputs to produce message inputs for next iteration of SHA256 round operations.

18. The system of claim 17, wherein the first, second, third, and fourth operands refer to storage locations each having at least 128 bits.

19. The system of claim 17, wherein the intermediate result is stored in the first storage location associated with the first operand.

20. The system of claim 17, wherein the message inputs for the next iteration of SHA256 round operations are stored in the third storage location associated with the third operand.

21. The system of claim 17, wherein for a current iteration i of the SHA256 message scheduling operations, the first storage location stores message inputs $w(i-13)$, $w(i-14)$, $w(i-15)$, and $w(i-16)$, and wherein the second storage location stores message inputs $w(i-9)$, $w(i-10)$, $w(i-11)$, and $w(i-12)$.

22. The system of claim 17, wherein for a current iteration i of the SHA256 message scheduling operations, the intermediate result represents $w(i-13)+s0(w(i-12))$, $w(i-14)+s0(w(i-13))$, $w(i-15)+s0(w(i-14))$, and $w(i-16)+s0(w(i-15))$.

23. The system of claim 17, wherein for a current iteration i of the SHA256 message scheduling operations, the third storage location stores a combination of the intermediate result and message inputs $w(i-5)$, $w(i-6)$, and $w(i-7)$.

24. The system of claim 17, wherein for a current iteration i of the SHA256 message scheduling operations, the message inputs for the next iteration represent $w(i+3)$, $w(i+2)$, $w(i+1)$, and $w(i)$.




FIG. 1
(Prior Art)




FIG. 2





FIG. 3






FIG. 5




FIG. 6

7/23

```

MOV          XMM0, {3, 2, 1, 0}
VADD  XMM11, XMM0, MEMK3-0
    SHA2_in   XMM14, XMM11          /* XMM14 = EFX Y for rounds 0, 1 */
    SHA2_2rnd XMM15, XMM14          /* ABCD, EFGH after 2 rounds */
VPERM  XMM12, XMM11                /* move wK3, wK2 to low words */
    SHA2_in   XMM14, XMM12          /* XMM14 = EFX Y for rounds 2, 3 */
    SHA2_2rnd XMM15, XMM14          /* ABCD, EFGH after 4 rounds */

MOV          XMM1, {7, 6, 5, 4}
MSG1SHA256 XMM0, XMM1
VADD  XMM11, XMM1, MEMK7-4
    SHA2_in   XMM14, XMM11          /* XMM14 = EFX Y for rounds 4, 5 */
    SHA2_2rnd XMM15, XMM14          /* ABCD, EFGH after 6 rounds */
VPERM  XMM12, XMM11                /* move wK7, wK6 to low words */
    SHA2_in   XMM14, XMM12          /* XMM14 = EFX Y for rounds 6, 7 */
    SHA2_2rnd XMM15, XMM14          /* ABCD, EFGH after 8 rounds */

MOV          XMM2, {11, 10, 9, 8}
VPERM XMM5, XMM2                    /* x, 11, 10, 9 from 11, 10, 9, 8 */
VADD XMM0, XMM5
MSG1SHA256 XMM1, XMM2
VADD  XMM11, XMM2, MEMK11-8
    SHA2_in   XMM14, XMM11
    SHA2_2rnd XMM15, XMM14
VPERM  XMM12, XMM11
    SHA2_in   XMM14, XMM12
    SHA2_2rnd XMM15, XMM14          /* ABCD, EFGH after 12 rounds */

MOV          XMM3, {15, 14, 13, 12}
VPERM XMM5, XMM3                    /* x, 15, 14, 13 from 15, 14, 13, 12 */
VADD XMM1, XMM5
MSG1SHA256 XMM2, XMM3
VADD  XMM11, XMM3, MEMK15-12
    SHA2_in   XMM14, XMM11
    SHA2_2rnd XMM15, XMM14
VPERM  XMM12, XMM11
    SHA2_in   XMM14, XMM12
    SHA2_2rnd XMM15, XMM14          /* ABCD, EFGH after 16 rounds */

MSG2SHA256 XMM0, XMM3    {19, 18, 17, 16}
VPERM XMM5, XMM0          /* x, 19, 18, 17 from 19, 18, 17, 16 */
VADD XMM2, XMM5
MSG1SHA256 XMM3, XMM0
VADD  XMM11, XMM0, MEMK19-16
    SHA2_in   XMM14, XMM11
    SHA2_2rnd XMM15, XMM14
VPERM  XMM12, XMM11
    SHA2_in   XMM14, XMM12
    SHA2_2rnd XMM15, XMM14          /* ABCD, EFGH after 20 rounds */
...

```

FIG. 7

```

SHA512IN1 XMM0,XMM1,XMM2
/* partial calculation for next 2 rounds, inject message + constants for 2 rounds */
XMM0 has WK3: WK1
XMM1 source state words E, F
    XMM1[127:63]=E, XMM1[63:0]= F
XMM2 source state words G, H
    XMM2[127:63]=G, XMM2[63:0]= H
XMM0 destination, Intermediate calculations X, Y
    XMM0[127:64] = X = WK1 + G
    XMM0[63: 0] = Y = WK1 + H + Ch(E,F,G) + s1(E)

SHA512IN2 XMM0,XMM1,XMM2
/* partial calculation for next 2 rounds, inject message + constants for 2 rounds */
XMM0 has X: Y, as defined earlier
XMM1 source state words E, F
    XMM1[127:63]=E, XMM1[63:0]= F
XMM2 source state words C, D
    XMM2[127:63]=C, XMM2[63:0]= D
XMM0 destination, Intermediate calculations X, Y
    XMM0[127:64] = Z = Ch(Y+D,E,F) + s1(Y+D) + X
    XMM0[63: 0] = Y

```

FIG. 8A

```

SHA512LO XMM0, XMM1, XMM2
/* calculate 2 rounds of SHA512 E, F from state [C, D], [Y and Z] */
XMM1 is source Z, Y
    XMM1[127:64]=Z, XMM1[63:0]= Y
XMM2 is source of state words C, D
    XMM2[127:64]=C, XMM2[63:0]= D
XMM0 is destination is updated E, F after two rounds of SHA512
    XMM0[127:64]=E, XMM0[63:0]= F

SHA512HI XMM0, XMM1, XMM2
/* calculate 2 rounds of SHA512 A, B from state [A,B,C, D], [Y and Z] */
XMM0 is source of state words C, D
    XMM0[127:64]=C, XMM0[63:0]= D
XMM1 is source of state words A, B
    XMM1[127:64]=A, XMM1[63:0]= B
XMM2 is source Z, Y
    XMM2[127:64]=Z, XMM2[63:0]= Y
XMM0 is destination is updated A, B after two rounds of SHA512
    XMM0[127:64]=A, XMM0[63:0]= B

```

FIG. 8B

FIG. 9A





FIG. 9B





FIG. 9C



12/23




FIG. 11A




FIG. 11B





FIG. 11C







FIG. 11D







FIG. 13B




FIG. 14B




FIG. 14A




FIG. 15




FIG. 16




FIG. 17




FIG. 18




FIG. 19




FIG. 20

A. CLASSIFICATION OF SUBJECT MATTER**G06F 9/46(2006.01)i, G06F 9/38(2006.01)i**

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)

G06F 9/46; H04K 1/00; H04L 9/28; H04L 9/06; G06F 9/38

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

Korean utility models and applications for utility models

Japanese utility models and applications for utility models

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

eKOMPASS(KIPO internal) & Keywords: secure hash algorithm, SHA, message scheduling, intermediate result

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	KR 10-2010-0047592 A (ELECTRONICS AND TELECOMMUNICATIONS RESEARCH INSTITUTE) 10 May 2010 See abstract, paragraphs [0031]-[0052],[0074]-[0085], claims 1-8, and figures 1,2,4.	1-24
A	US 2010-0278331 A1 (WALKER DOUGLAS WAYNE et al.) 04 November 2010 See abstract, paragraphs [0032]-[0038], claims 1-5, and figures 4,5,8.	1-24
A	KR 10-2010-0065721 A (ELECTRONICS AND TELECOMMUNICATIONS RESEARCH INSTITUTE) 17 June 2010 See abstract, paragraphs [0053]-[0088], claims 1-7, and figures 3-5.	1-24
A	US 2005-0089160 A1 (THOMAS CRISPIN et al.) 28 April 2005 See abstract, paragraphs [0043]-[0047],[0064]-[0067], claims 1-5, and figures 3,12,13.	1-24



Further documents are listed in the continuation of Box C.



See patent family 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 application or patent 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 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

02 September 2013 (02.09.2013)

Date of mailing of the international search report

02 September 2013 (02.09.2013)

Name and mailing address of the ISA/KR


Korean Intellectual Property Office
189 Cheongsa-ro, Seo-gu, Daejeon Metropolitan City,
302-701, Republic of Korea

Facsimile No. +82-42-472-7140

Authorized officer

PARK In Hwa

Telephone No. +82-42-481-5722



INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2013/045465

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
KR 10-2010-0047592 A	10/05/2010	JP 2010-107947A US 2010-0104098 A1	13/05/2010 29/04/2010
US 2010-0278331 A1	04/11/2010	EP 2425580 A1 EP 2425580 B1 WO 2010-127075 A1	07/03/2012 17/04/2013 04/11/2010
KR 10-2010-0065721 A	17/06/2010	JP 2010-134465 A JP 5143817 B2 US 2010-0146296 A1 US 8275126 B2	17/06/2010 13/02/2013 10/06/2010 25/09/2012
US 2005-0089160 A1	28/04/2005	CN 100463392 C0 CN 100527664 C CN 100539495 C CN 100573441 C CN 102594547 A CN 1538656 A CN 1538656 C0 CN 1558591 A CN 1558591 C0 CN 1592189 A CN 1592189 C0 CN 1607763 A CN 1649296 A CN 1649296 C0 CN 1652163 A CN 1652163 B CN 1652163 C0 CN 1655496 A CN 1655496 B CN 1655496 C0 CN 1658548 A CN 1658548 B CN 1658548 C0 CN 1658550 A CN 1658550 C0 CN 1661958 A CN 1661958 B CN 1661958 C CN 1661958 C0 CN 1684408 A CN 1684408 B CN 1684408 C0 CN 1684409 A CN 1684409 C0 CN 1684412 A CN 1684412 B CN 1684412 C0	18/02/2009 12/08/2009 09/09/2009 23/12/2009 18/07/2012 20/10/2004 15/08/2007 29/12/2004 28/05/2008 09/03/2005 09/03/2005 20/04/2005 03/08/2005 03/08/2005 10/08/2005 11/07/2012 10/08/2005 17/08/2005 14/07/2010 17/08/2005 24/08/2005 15/09/2010 24/08/2005 24/08/2005 24/08/2005 31/08/2005 28/04/2010 28/04/2010 31/08/2005 19/10/2005 23/06/2010 19/10/2005 19/10/2005 19/10/2005 13/10/2010 19/10/2005

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2013/045465

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
		CN 1834898 A0	20/09/2006
		CN 1838140 A0	27/09/2006
		EP 1469371 A2	20/10/2004
		EP 1469371 A3	07/06/2006
		EP 1469371 B1	02/07/2008
		EP 1496421 A2	12/01/2005
		EP 1496421 A3	05/04/2006
		EP 1496421 B1	28/10/2009
		EP 1519509 A2	30/03/2005
		EP 1519509 A3	21/03/2007
		EP 1538510 A1	08/06/2005
		EP 1538510 B1	13/08/2008
		EP 1586971 A2	19/10/2005
		EP 1586971 A3	27/09/2006
		EP 1596281 A2	16/11/2005
		EP 1596281 A3	02/01/2008
		EP 1596530 A1	16/11/2005
		EP 1596530 B1	02/07/2008
		EP 1724675 A1	22/11/2006
		TW 268686 A	11/12/2006
		TW 268689 A	11/12/2006
		TW 272815 A	01/02/2007
		TW 272815 B	01/02/2007
		TW 274280 A	21/02/2007
		TW 274280 B	21/02/2007
		TW 274281 A	21/02/2007
		TW 274281 B	21/02/2007
		TW 282230 A	01/06/2007
		TW 282230 B	01/06/2007
		TW I272815 B	01/02/2007
		TW I274280 B	21/02/2007
		TW I274281 B	21/02/2007
		TW I282230 B	01/06/2007
		TW I303936 A	01/12/2008
		TW I303936 B	01/12/2008
		TW I336858 A	01/02/2011
		TW I336858 B	01/02/2011
		TW I351864 I	01/11/2011
		US 2004-0208072 A1	21/10/2004
		US 2004-0208318 A1	21/10/2004
		US 2004-0223610 A1	11/11/2004
		US 2004-0228479 A1	18/11/2004
		US 2004-0228481 A1	18/11/2004
		US 2004-0228483 A1	18/11/2004
		US 2004-0250090 A1	09/12/2004
		US 2004-0250091 A1	09/12/2004
		US 2004-0252841 A1	16/12/2004
		US 2004-0252842 A1	16/12/2004
		US 2004-0255129 A1	16/12/2004
		US 2004-0255130 A1	16/12/2004

INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2013/045465

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
		US 2005-0160279 A1	21/07/2005
		US 2005-0188216 A1	25/08/2005
		US 2005-0256920 A1	17/11/2005
		US 2011-0142228 A1	16/06/2011
		US 2011-0142229 A1	16/06/2011
		US 2011-202775 A1	18/08/2011
		US 7321910 B2	22/01/2008
		US 7392400 B2	24/06/2008
		US 7502943 B2	10/03/2009
		US 7519833 B2	14/04/2009
		US 7529367 B2	05/05/2009
		US 7529368 B2	05/05/2009
		US 7532722 B2	12/05/2009
		US 7536560 B2	19/05/2009
		US 7539876 B2	26/05/2009
		US 7542566 B2	02/06/2009
		US 7664810 B2	16/02/2010
		US 7844053 B2	30/11/2010
		US 7900055 B2	01/03/2011
		US 7921300 B2	05/04/2011
		US 7925891 B2	12/04/2011
		US 8060755 B2	15/11/2011
		US 8132022 B2	06/03/2012
		US 8132023 B2	06/03/2012
		US 8255703 B2	28/08/2012