



(51) International Patent Classification:

G06F 9/30 (2006.01) G06F 21/22 (2006.01)
G06F 9/305 (2006.01) H04L 9/28 (2006.01)

(21) International Application Number:

PCT/US201 1/064640

(22) International Filing Date:

13 December 2011 (13.12.2011)

(25) Filing Language:

English

(26) Publication Language:

English

(71) Applicant (for all designated States except US): INTEL CORPORATION [US/US]; 2200 Mission College Boulevard, MS: RNB-4-150, Santa Clara, California 95052 (US).

(72) Inventors; and

(75) Inventors/Applicants (for US only): YAP, Kirk S. [US/US]; 224 Belknap Road, Framingham, Massachusetts 01701 (US). WOLRICH, Gilbert M. [US/US]; 1 Macomber Lane, Framingham, Massachusetts 01702 (US). GUILFORD, James D. [US/US]; 17 Mashpee Circle, Northborough, Massachusetts 01532 (US). GOPAL, Vinodh [US/US]; 15 West End Avenue, Westborough, Massachusetts 01581 (US). OZTURK, Erdinc [TR/US]; 71 Christie Way, Apt. 93C, Marlborough, Massachusetts 01752 (US). GULLEY, Sean M. [US/US]; 146 Marlborough Street, Boston, Massachusetts 02116 (US). FEGHALI, Wajdi K. [CA/US]; 199 Massachusetts Ave., Apt. 206, Boston, Massachusetts 02115 (US). DIXON, Martin G. [US/US]; 4005 NE Hazelfern Place, Portland, Oregon 97232 (US).

(74) Agents: SHAO, Kevin G. et al; Blakely Sokoloff Taylor & Zafman LLP, 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, 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, KM, 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, 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.

[Continued on nextpage]

(54) Title: METHOD AND APPARATUS TO PROCESS KECCAK SECURE HASHING ALGORITHM




FIG. 2

(57) Abstract: A processor includes a plurality of registers, an instruction decoder to receive an instruction to process a KECCAK state cube of data representing a KECCAK state of a KECCAK hash algorithm, to partition the KECCAK state cube into a plurality of subcubes, and to store the subcubes in the plurality of registers, respectively, and an execution unit coupled to the instruction decoder to perform the KECCAK hash algorithm on the plurality of subcubes respectively stored in the plurality of registers in a vector manner.



(84) **Designated States** (unless otherwise indicated, for every *Mnd 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, MD, 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, ML, MR, NE, SN, TD, TG).

Declarations under Rule 4.17:

— as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(Hi))

Published:

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

METHOD AND APPARATUS TO PROCESS KECCAK SECURE HASHING ALGORITHMFIELD OF THE INVENTION

5

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

10 BACKGROUND

KECCAK is a new secure hashing function that maintains state in an array of bits arranged with a 5x5 edge surface and depth of 2^L where $L = 0$ to 6. KECCAK is under consideration by the National Institute of Standards and Technology (NIST) as an
 15 algorithm for selection as the SHA-3 standard. The performance of KECCAK in hardware implementations exceeds that of other secure hash algorithms competing for the standard due to the simple logic functions required and the parallelism that can be utilized. The performance of KECCAK for current software implementations is constrained by the high number of logic operations that require individual integer instructions.

20 A KECCAK state can be viewed as a three dimensional array of elements (bits) with a 5x5 element edge termed a "slice" and a depth (z direction) of w bits where the depth is a power of 2, i.e. $w = 2^1$ for $1=0$ to 6, as shown in Figure 1A. KECCAK uses a "sponge" construction where r bits are input to the KECCAK state with an XOR of the "first" r bits of state, followed by the KECCAK-f state update function. KECCAK-1600 ($1=64$) is the target function providing
 25 the highest capacity for message authentication. Mapping the lanes of the state, i.e., the one-dimensional sub-arrays in the direction of the z axis, onto 64 bit processor words, results in simple and efficient software implementation for the step mappings. For $1=6$, KECCAK-1600 ($5 \times 5 \times 64$) is the state update function consisting of n_r rounds of five steps/permutations, Θ , ρ , π , χ , ι , as shown in Figure 1B.

30 The Θ function is performed as illustrated in Figure 1C and the ρ function rotates the lanes/registers by specified offsets, requiring 24 register rotates (one offset is zero) as shown in Figure 1D. The ι function is a transposition of the lanes. This transposition changes the usage and grouping of every 64 bit section as shown in Figure 1E. The transpose has a period of 24

rounds before lanes return to their original position. During the χ step, each row is transformed by neighboring elements of that row:

$$A[x, y] = a[x, y] \oplus ((\text{NOT } a[x + 1, y]) \text{ AND } a[x + 2, y])$$

5 A row of lanes are calculated together, with 5 NOTs, 5 ANDs, and 5 XORs needed for a total of $15 \times 5 = 75$ operations. The /add round constant is applied to a single register/lane requiring 1 XOR instruction as shown in Figure 1F. The operations per round are 55 for Θ , 24 for ρ , zero for π , 75 for χ , 1 for ι . There are a total of 155 operations per round. The number of rounds n_r is $12 + 21 = 24$ for $l=6$, i.e. 64 bit registers. For 24 rounds KECCAK requires $155 * 24$
10 $= 3720$ operations. On a four execution unit processor if each operation requires an instruction, a minimum of 930 cycles are required.

The above operations cannot be performed efficiently using parallel execution of the functions using vector instructions, such as streaming single instruction multiple data (SIMD) extensions (SSE) or advanced vector extensions (AVX) instructions from Intel®
15 Corporation of Santa Clara, California, due to the /rfunction since the location of the lanes and corresponding words of the cube are scrambled each round.

BRIEF DESCRIPTION OF THE DRAWINGS

20 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.

Figures 1A-1F are diagrams illustrating a process of a typical KECCAK hash algorithm.

Figure 2 is a block diagram illustrating an example of a processor according one embodiment.

25 Figure 3 is a block diagram illustrating partitions of a KECCAK state cube according to one embodiment.

Figure 4 is a block diagram illustrating a register layout used in KECCAK operations according to one embodiment.

Figure 5 is pseudocode to perform KECCAK hash algorithm according to one
30 embodiment.

Figure 6 is a flow diagram illustrating a method to perform KECCAK hash algorithm according to one embodiment.

Figure 7 is a block diagram illustrating an example of a data processing system according to one embodiment.

Figure 8 is a block diagram illustrating an example of a data processing system according to another embodiment.

DETAILED DESCRIPTION

5

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 construed as limiting the invention. Numerous specific details are described to provide
10 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
15 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.

According to some embodiments, a vector instruction and a data path are utilized to reduce instructions of the KECCAK function from more than one hundred instructions per
20 round to approximately eight instructions per KECCAK round. In one embodiment, in response to an instruction to perform KECCAK operations received at a processor, several registers having at least 450 bits are utilized to store data representing a KECCAK state cube as shown in Figure 1A. As described above, a KECCAK state cube (e.g., 64 slices) requires 1600 storage bits. In one embodiment, the KECCAK state cube is partitioned into
25 at least four subcubes, where each subcube represents at least 16 slices or a $\frac{1}{4}$ state cube. The four subcubes are stored in at least four registers that have at least 450 bits, such as ZMM compatible registers (e.g., 512-bit wide registers) available from Intel processors, where each register stores a $5 \times 5 \times 16$ or 400 bit portion of the KECCAK state cube (e.g., $\frac{1}{4}$ state cube). The 400 bits stored in each register are partitioned as 25 16-bit sections,
30 where each 16-bit section represents $\frac{1}{4}$ of a lane of the full KECCAK state cube. The rest of 50 bits may be utilized to store intermediate results (e.g., additional two slices) during the operations as set forth below. Thereafter, at least some of the KECCAK functions set forth above can be performed by a vector capable processor in a vector manner or parallel.

Figure 2 is a block diagram illustrating an example of a processor according one embodiment. Referring to Figure 2, processor 100 may represent any kind of instruction processing apparatuses. For example, processor 101 may be a 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. In one embodiment, processor 100 includes instruction decoder 101 to receive and decode instruction 106. Instruction decoder 101 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, instruction 106. Instruction decoder 101 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.

Processor 100 further includes one or more execution units 102, which may include an arithmetic logic unit, or another type of logic unit capable of performing operations based on instruction 106. As a result of instruction decoder 101 decoding instruction 106, 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, instruction 106. Execution unit 102 may be operable as a result of instruction 106 indicating one or more source operands (SRC) 108 and to store a result in one or more destination operands (DEST) 109 of register set 105 indicated by instruction 106. 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 instruction 106 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.

In one embodiment, instruction 106 may implicitly indicate and/or explicitly specify (e.g., through one or more dedicated fields or sets of bits) the source and destination operands. Examples of suitable sources and/or destination of the operands include registers, memory, immediate of the instruction, and combinations thereof. In various embodiments, the source and destination operands may be 8-bit, 16-bit, 32-bit, or 64-bit operands, although this is not required.

In one embodiment, some or all of the source and destination operands may be stored in registers of a register set or memory 105. 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.

According to one embodiment, a vector instruction and data path for performing KECCAK operations are utilized to reduce instructions of the KECCAK function from more than one hundred instructions per round to approximately eight instructions per KECCAK round. In one embodiment, in response to instruction 106 to perform KECCAK operations received at instruction decoder 101, instruction decoder 101 is configured to arrange registers 104 having at least 400 bits to store data 110 representing a KECCAK state cube and additional at least 50 bits to store intermediate results. Note that the registers 104 may be part of register set 105 (e.g., source and/or destination registers) and KECCAK state cube 110 may be retrieved from the memory. As described above, a KECCAK state cube (e.g., 64 slices) requires 1600 storage bits.

In one embodiment, in response to instruction 106, instruction decoder 101 and/or execution unit 102 may partition KECCAK state cube 110 into four subcubes, where each subcube represents 16 slices or a $\frac{1}{4}$ state cube, as shown in Figure 3. Referring to Figure 3, KECCAK state cube 110 is partitioned into four subcubes 301-304. Four subcubes 301-304 are then stored in at least four registers 305-308, respectively. In one embodiment, registers 305-308 have at least 450 bits, preferably 512-bits such as ZMM compatible registers (e.g., 512-bit wide registers) available from Intel processors. Each of registers 305-308 stores a $5 \times 5 \times 16$ or 400 bit portion of the KECCAK state cube (e.g., $\frac{1}{4}$ state cube). The 400 bits stored in each of registers 305-308 are partitioned as 25 16-bit sections, as shown in Figure 4.

Referring to Figure 4, register 400 may represent any one of registers 305-308 of Figure 3. In one embodiment, register 400 is partitioned into 25 16-bit sections 401-403, where each of 16-bit sections 401-403 represents a $\frac{1}{4}$ lane of the full KECCAK state cube 110. Referring back to Figure 2, once KECCAK state cube 110 has been partitioned and

properly stored in registers 305-308 of Figure 3, at least some of the KECCAK functions are performed by KECCAK unit(s) 106 of execution unit(s) 102 in a vector manner or parallel.

According to some embodiments, the π , χ , and ρ functions can be performed for each $\frac{1}{4}$ cube independent of data from the other $\frac{1}{4}$ cube sections. To calculate the Θ function of a quadrant (e.g., subcube or $\frac{1}{4}$ state cube), data from the "next" quadrant will also be needed. Calculating the ρ function for the first 15 slices only need data contained within the quadrant being operated on. However, to operate on the 16th slice, the 17th slice that is stored in the "next" quadrant is needed. The ρ rotate function requires state from each of the $\frac{1}{4}$ cubes to perform the lane rotate functions, which are identical for each lane of all the $\frac{1}{4}$ cubes.

According to one embodiment, at least two instructions or phases are used to process a round of KECCAK operations on each $\frac{1}{4}$ cube for a total of eight instructions per round. The first instruction or phase is referred to as KECCAK_THETA and can be defined as follows according to one embodiment:

15 KECCAK_THETA Dst/Src 1, Src2, Src3

During the KECCAK_THETA phase, the processor is configured to perform the Θ function and the 1st part of the ρ (rotate) function on each $\frac{1}{4}$ slice:

Src1 = The State Quadrant being processed, $z = (z_0 \text{ to } z_0+15) \bmod 64$
 Src2 = The State Quadrant next to Src1, $z = (z_0+16 \text{ to } z_0+31) \bmod 64$
 20 Src3 = The State Quadrant next to Src2, $z = (z_0+32 \text{ to } z_0+47) \bmod 64$

In one embodiment, the first part of the instruction or phase KECCAK_THETA performs the Θ function on Src1, Src2 and Src3. Note that the Θ of the last slice of Src3 cannot be computed since it needs a slice from the 4th quadrant. The second part of KECCAK_THETA performs a portion of the ρ function. Specifically, for each lane, the processing logic extracts the post Θ bits (e.g., bits 404 of Figure 4) for bits in the src1 quadrant whose bits will come from the bits in quadrants 1, 2, and 3. The last part of the instruction or phase KECCAK_THETA is to extract the 2 unprocessed slices from Src1 and Src3 (e.g., the first slice of Src1 and the last slice of Src3) and place them in the 17th slice of a destination register Dst, Dst[449:400], i.e. Dst[424:400] = Src1[0, 16, 32, ..., 384] and the 18th slice of Dst, Dst[449:425] = Src3[15, 31, ..., 399].

The second instruction or phase is referred to as KECCAK_ROUND and can be defined as follows according to one embodiment:

KECCAK_ROUND Dst/Src 1, Src2, Src3

During the KECCAK_ROUND phase, processor is to complete p rotate function and perform π, χ , /functions on slices in Src1 quadrant:

Src1 = Result of KECCAK_THETA instruction

Src2 = Unprocessed Quadrant furthest from the Src1 Quadrant

5 Src3 = KECCAK round constant to be used for /function

According to one embodiment, the first part of the instruction KECCAK_ROUND is to perform the Θ function on the last slice of the "third" quadrant that is stored in Src1 [449:425] and the Src2 quadrant. Note that in order to calculate Θ of the last slice of Src2, it needs the unprocessed first slice of Src1 quadrant, which is stored in Src1 [424:400]. The second part of
 10 KECCAK_ROUND is to extract the bits from the result of the first part that will be needed to complete the p function for the quadrant of Src1. This is then followed by performing the π, χ , /functions on the completed p transformed quadrant corresponding to Src1.

Figure 5 is an example of a pseudocode that performs a round of KECCAK according to one embodiment. Pseudocode 500 may be generated by a compiler and executed by processor
 15 100 of Figure 2 in a vector manner. Referring to Figure 5, the state corresponding to the four quadrants is stored in registers ZMM1, ZMM2, ZMM3, and ZMM4, respectively. Here a ZMM compatible register is at least 450-bit wide, preferably at least 512-bit wide. During initialization 501, a copy of subcubes (e.g., a $\frac{1}{4}$ KECCAK state cube) stored in registers ZMM1-ZMM4 are saved in temporary registers ZMMTMP1- ZMMTMP4, respectively. During phase 502, which
 20 is the KECCAK_THETA phase, the Θ function and the first part of the p function are performed on each of the subcubes stored in registers ZMM1-ZMM4.

During phase 503, which is the KECCAK_ROUND phase, the second part of p function, as well as the π, χ , /functions are performed. Note that register ZMMRC herein contains the round constant used for the /function. In one embodiment, the second SIMD port of the
 25 processor is used to execute the MOV instructions that are used to retain the unprocessed $\frac{1}{4}$ cube state in order to complete the rotations. In one embodiment, the data path for KECCAK_THETA and KECCAK_ROUND is implemented with a three-cycle data path for a throughput of 10 cycles per round. The 24 rounds of KECCAK complete in 240 cycles or an approximately 4x improvement over the single operation per instruction integer code.

30 Figure 6 is a flow diagram illustrating a method for performing KECCAK operations in a vector manner according to one embodiment. Method 600 may be performed by processor 100 of Figure 2. Referring to Figure 6, at block 601, a processor receives an instruction to perform KECCAK operations on data representing a KECCAK state cube having 64 slices. At block 602, the KECCAK state cube is partitioned into four subcubes (e.g., $\frac{1}{4}$ state cubes), each having

16 slices. At block 603, the subcubes are stored in four registers, respectively, where each register is at least 450-bit wide, preferably 512-bit wide. For each round of the KECCAK operations, at block 604, the KECCAK_THETA phase is performed on the subcubes, including performing a Θ function and a first part of a ρ function on the subcubes. At block 605, the
5 KECCAK_ROUND phase is performed, including performing a second part of the ρ function and π, χ, γ functions on the subcubes.

Figure 7 is a block diagram illustrating an example of a data processing system according to one embodiment of the invention. System 900 may represent any of the systems described above. For example, processor 901 may represent processor 100 of Figure 2. System 900 may
10 represent a desktop, a laptop, a tablet, a server, a mobile phone (e.g., Smartphone), a media player, a personal digital assistant (PDA), a personal communicator, a gaming device, a network router or hub, a wireless access point or repeater, a set-top box, or a combination thereof. Note that while Figure 7 illustrates various components of a data processing system, it is not intended to represent any particular architecture or manner of interconnecting the components; as such
15 details are not germane to embodiments of the present invention. It will also be appreciated that network computers, handheld computers, mobile phones, and other data processing systems which have fewer components or perhaps more components may also be used with embodiments of the present invention.

Referring to Figure 7, in one embodiment, system 900 includes processor 901 and chipset
20 902 to couple various components to processor 901 including memory 905 and devices 903-904 via a bus or an interconnect. Processor 901 may represent a single processor or multiple processors with a single processor core or multiple processor cores 909 included therein. Processor 901 may represent one or more general-purpose processors such as a microprocessor, a central processing unit (CPU), or the like. More particularly, processor 901 may be a complex
25 instruction set computing (CISC) microprocessor, reduced instruction set computing (RISC) microprocessor, very long instruction word (VLIW) microprocessor, or processor implementing other instruction sets, or processors implementing a combination of instruction sets. Processor 901 may also be one or more special-purpose processors such as an application specific integrated circuit (ASIC), a field programmable gate array (FPGA), a digital signal processor
30 (DSP), a network processor, a graphics processor, a network processor, a communications processor, a cryptographic processor, a co-processor, an embedded processor, or any other type of logic capable of processing instructions. For example, processor 901 may be a Pentium® 4, Pentium® Dual-Core, Core™ 2 Duo and Quad, Xeon™, Itanium™, XScale™, Core™ i7, Core™ i5, Celeron®, or StrongARM™ microprocessor available from Intel Corporation of Santa

Clara, California. Processor 901 is configured to execute instructions for performing the operations and steps discussed herein.

Processor 901 may include an instruction decoder, which may receive and decode a variety of instructions. The decoder 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, an original input instruction. The decoder 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.

The decoder may not be a required component of processor 901. In one or more other embodiments, processor 901 may instead have an instruction emulator, an instruction translator, an instruction morpher, an instruction interpreter, or other instruction conversion logic. Various different types of instruction emulators, instruction morphers, instruction translators, and the like, are known in the arts. The instruction conversion logic may receive the bit range isolation instruction, emulate, translate, morph, interpret, or otherwise convert the bit range isolation instruction, and output one or more instructions or control signals corresponding to the original bit range isolation instruction. The instruction conversion logic may be implemented in software, hardware, firmware, or a combination thereof. In some cases, some or all of the instruction conversion logic may be located off-die with the rest of the instruction processing apparatus, such as a separate die or in a system memory. In some cases, the instruction processing apparatus may have both the decoder and the instruction conversion logic.

Processor 901 and/or cores 909 may further include one or more execution units coupled with, or otherwise in communication with, an output of the decoder. The term "coupled" may mean that two or more elements are in direct electrical contact or connection. However, "coupled" may also mean that two or more elements are not in direct connection with each other, but yet still co-operate or interact or communicate with each other (e.g., through an intervening component). As one example, the decoder and the execution unit may be coupled with one another through an intervening optional buffer or other component(s) known in the arts to possibly be coupled between a decoder and an execution unit. Processor 901 and/or cores 909 may further include multiple different types of execution units, such as, for example, arithmetic units, arithmetic logic units (ALUs), integer units, etc.

Processor 901 may further include one or more register files including, but are not limited to, integer registers, floating point registers, vector or extended registers, status registers, and an instruction pointer register, etc. The term "registers" is used herein to refer to the on-board

processor storage locations that are used as part of macro-instructions to identify operands. In other words, the registers referred to herein are those that are visible from the outside of the processor (from a programmer's perspective). However, the registers should not be limited in meaning to a particular type of circuit. Rather, a register need only be capable of storing and providing data, and performing the functions described herein. The registers described herein can be implemented by circuitry within a processor using any number of different techniques, such as dedicated physical registers, dynamically allocated physical registers using register renaming, combinations of dedicated and dynamically allocated physical registers, etc. In one embodiment, integer registers store 32-bit or 64-bit integer data. A register file may contain extended multimedia SIMD registers (e.g., XMM) for packed data. Such registers may include 128 bits wide XMM registers and 256 bits wide registers (which may incorporate the XMM registers in their low order bits) relating to SSE2, SSE3, SSE4, GSSE, and beyond (referred to generically as "SSEx") technology to hold such packed data operands.

Processor 901 and/or cores 909 may also optionally include one or more other well-known components. For example, processor 901 may optionally include instruction fetch logic, pre-decode logic, scheduling logic, re-order buffers, branch prediction logic, retirement logic, register renaming logic, and the like, or some combination thereof. These components may be implemented conventionally, or with minor adaptations that would be apparent to those skilled in the art based on the present disclosure. Further description of these components is not needed in order to understand the embodiments herein, although further description is readily available, if desired, in the public literature. There are literally numerous different combinations and configurations of such components known in the arts. The scope is not limited to any known such combination or configuration. Embodiments may be implemented either with or without such additional components.

Chipset 902 may include memory control hub (MCH) 910 and input output control hub (ICH) 911. MCH 910 may include a memory controller (not shown) that communicates with a memory 905. MCH 910 may also include a graphics interface that communicates with graphics device 912. In one embodiment of the invention, the graphics interface may communicate with graphics device 912 via an accelerated graphics port (AGP), a peripheral component interconnect (PCI) express bus, or other types of interconnects. ICH 911 may provide an interface to I/O devices such as devices 903-904. Any of devices 903-904 may be a storage device (e.g., a hard drive, flash memory device), universal serial bus (USB) port(s), a keyboard, a mouse, parallel port(s), serial port(s), a printer, a network interface (wired or wireless), a

wireless transceiver (e.g., WiFi, Bluetooth, or cellular transceiver), a media device (e.g., audio/video codec or controller), a bus bridge (e.g., a PCI-PCI bridge), or a combination thereof.

MCH 910 is sometimes referred to as a Northbridge and ICH 911 is sometimes referred to as a Southbridge, although some people make a technical distinction between them. As used
5 herein, the terms MCH, ICH, Northbridge and Southbridge are intended to be interpreted broadly to cover various chips whose functions include passing interrupt signals toward a processor. In some embodiments, MCH 910 may be integrated with processor 901. In such a configuration, chipset 902 operates as an interface chip performing some functions of MCH 910 and ICH 911, as shown in Figure 8. Furthermore, graphics accelerator 912 may be integrated
10 within MCH 910 or processor 901.

Memory 905 may store data including sequences of instructions that are executed by processor 901, or any other device. For example, executable code 913 and/or data 914 of a variety of operating systems, device drivers, firmware (e.g., input output basic system or BIOS), and/or applications can be loaded in memory 905 and executed by processor 901. An operating
15 system can be any kind of operating systems, such as, for example, Windows® operating system from Microsoft®, Mac OS®/iOS® from Apple, Android® from Google®, Linux®, Unix®, or other real-time operating systems. In one embodiment, memory 905 may include one or more volatile storage (or memory) devices such as random access memory (RAM), dynamic RAM (DRAM), synchronous DRAM (SDRAM), static RAM (SRAM), or other types of storage devices.

Nonvolatile memory may also be utilized such as a hard disk or a flash storage device. Front
20 side bus (FSB) 906 may be a multi-drop or point-to-point interconnect. The term FSB is intended to cover various types of interconnects to processor 901. Chipset 902 may communicate with other devices such as devices 903-904 via point-to-point interfaces. Bus 906 may be implemented as a variety of buses or interconnects, such as, for example, a quick path interconnect (QPI), a hyper transport interconnect, or a bus compatible with advanced
25 microcontroller bus architecture (AMBA) such as an AMBA high-performance bus (AHB).

Cache 908 may be any kind of processor cache, such as level-1 (L1) cache, L2 cache, L3 cache, L4 cache, last-level cache (LLC), or a combination thereof. Cache 908 may be shared with processor cores 909 of processor 901. Cache 908 may be embedded within processor 901
30 and/or external to processor 901. Cache 908 may be shared amongst cores 909. Alternatively, at least one of cores 909 further includes its own local cache embedded therein. At least one of cores 909 may utilize both the local cache and the cache shared with another one of cores 909. Processor 901 may further include a direct cache access (DCA) logic to enable other devices such as devices 903-904 to directly access cache 908. Processor 901 and/or chipset 902 may

further include an interrupt controller, such as an advanced programmable interrupt controller (APIC), to handle interrupts such as message signaled interrupts.

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.

5 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.

10 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
15 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
20 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,
25 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
30 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, 5 to be regarded in an illustrative sense rather than a restrictive sense.

CLAIMS

What is claimed is:

1. A processor, comprising:
a plurality of registers;
an instruction decoder to receive an instruction to process a KECCAK state cube of data
5 representing a KECCAK state of a KECCAK hash algorithm, to partition the
KECCAK state cube into a plurality of subcubes, and to store the subcubes in the
plurality of registers, respectively; and
an execution unit coupled to the instruction decoder to perform the KECCAK hash
algorithm on the plurality of subcubes respectively stored in the plurality of
10 registers in a vector manner.
2. The processor of claim 1, wherein the KECCAK state cube includes 64 slices partitioned
into 4 subcubes, wherein each subcube contains 16 slices.
3. The processor of claim 2, wherein the plurality of registers include 4 registers, each
having at least 450 bits.
4. The processor of claim 1, wherein, for each round of the KECCAK algorithm, the
execution unit is configured to perform KECCAK_THETA operations, including
performing a Θ function of the KECCAK algorithm on the subcubes stored in the
15 registers in parallel, and
performing a first portion of a ρ function of the KECCAK algorithm on the subcubes in
parallel.
5. The processor of claim 4, wherein the execution unit is further configured to perform
KECCAK_ROUND operations, including
performing a second portion of the ρ function of the KECCAK algorithm on the
20 subcubes in parallel,
performing a π function of the KECCAK algorithm on the subcubes in parallel,
performing a χ function of the KECCAK algorithm on the subcubes in parallel, and
performing an ι function of the KECCAK algorithm on the subcubes in parallel.

6. The processor of claim 5, wherein the instruction comprises a first instruction followed by a second instruction, wherein in response to the first instruction, the execution unit is configured to perform the KECCAK_THETA operations of the KECCAK algorithm, and wherein in response to the second instruction, the execution unit is configured to perform the KECCAK_ROUND operations of the KECCAK algorithm.

7. The processor of claim 5, wherein each of the plurality of registers contains at least 512 bits, wherein at least 400 bits are utilized to store a subcube, and wherein a remainder of the at least 512 bits is utilized to store intermediate results of the KECCAK_THETA operations and KECCAK_ROUND operations.

8. A method, comprising:

receiving, at an instruction decoder of a processor, an instruction to process a KECCAK state cube of data representing a KECCAK state of a KECCAK hash algorithm, to partition the KECCAK state cube into a plurality of subcubes, and to store the subcubes in a plurality of registers, respectively; and
performing, by an execution unit of the processor, the KECCAK hash algorithm on the plurality of subcubes respectively stored in the plurality of registers in a vector manner.

9. The method of claim 8, wherein the KECCAK state cube includes 64 slices partitioned into 4 subcubes, wherein each subcube contains 16 slices.

10. The method of claim 9, wherein the plurality of registers include 4 registers, each having at least 450 bits.

11. The method of claim 8, wherein, for each round of the KECCAK algorithm, the execution unit is configured to perform KECCAK_THETA operations, including performing a Θ function of the KECCAK algorithm on the subcubes stored in the registers in parallel, and performing a first portion of a ρ function of the KECCAK algorithm on the subcubes in parallel.

12. The method of claim 11, wherein the execution unit is further configured to perform KECCAK_ROUND operations, including

performing a second portion of the p function of the KECCAK algorithm on the subcubes in parallel,

performing a π function of the KECCAK algorithm on the subcubes in parallel,

performing a χ function of the KECCAK algorithm on the subcubes in parallel, and

5 performing an ι function of the KECCAK algorithm on the subcubes in parallel.

13. The method of claim 12, wherein the instruction comprises a first instruction followed by a second instruction, wherein in response to the first instruction, the execution unit is configured to perform the KECCAK_THETA operations of the KECCAK algorithm, and wherein in response to the second instruction, the execution unit is configured to perform the KECCAK_ROUND operations of the KECCAK algorithm.

14. The method of claim 12, wherein each of the plurality of registers contains at least 512 bits, wherein at least 400 bits are utilized to store a subcube, and wherein a remainder of the at least 512 bits is utilized to store intermediate results of the KECCAK_THETA operations and KECCAK_ROUND operations.

15. A data processing system, comprising:
an interconnect;

10 a processor coupled to the interconnect to receive an instruction to process a KECCAK state cube of data representing a KECCAK state of a KECCAK hash algorithm, to partition the KECCAK state cube into a plurality of subcubes, to store the subcubes in a plurality of registers, respectively, and to perform the KECCAK hash algorithm on the plurality of subcubes respectively stored in the plurality of registers in a vector manner; and

15 a dynamic random access memory (DRAM) coupled to the interconnect.

16. The system of claim 15, wherein the KECCAK state cube includes 64 slices partitioned into 4 subcubes, wherein each subcube contains 16 slices.

17. The system of claim 16, wherein the plurality of registers include 4 registers, each having at least 450 bits.

18. The system of claim 1, wherein, for each round of the KECCAK algorithm, the processor is configured to perform KECCAK_THETA operations, including

performing a Θ function of the KECCAK algorithm on the subcubes stored in the registers in parallel, and

performing a first portion of a ρ function of the KECCAK algorithm on the subcubes in parallel.

19. The system of claim 4, wherein the processor is further configured to perform KECCAK_ROUND operations, including

performing a second portion of the ρ function of the KECCAK algorithm on the subcubes in parallel,

performing a π function of the KECCAK algorithm on the subcubes in parallel,

performing a χ function of the KECCAK algorithm on the subcubes in parallel, and

performing an ι function of the KECCAK algorithm on the subcubes in parallel.

20. The system of claim 5, wherein the instruction comprises a first instruction followed by a second instruction, wherein in response to the first instruction, the processor is configured to perform the KECCAK_THETA operations of the KECCAK algorithm, and wherein in response to the second instruction, the processor is configured to perform the KECCAK_ROUND operations of the KECCAK algorithm.

1/11




FIG. 1A
(Prior Art)

$$\begin{aligned} \theta: \quad & a[x][y][z] \quad \leftarrow a[x][y][z] + \sum_{y'=0}^4 a[x-1][y'][z] + \sum_{y'=0}^4 a[x+1][y'][z-1], \\ \rho: \quad & a[x][y][c] \quad \leftarrow a[x][y][z - (t+1)(t+2)/2], \\ & \text{with } t \text{ satisfying } 0 \leq t < 24 \text{ and } \begin{pmatrix} 0 & 1 \\ 2 & 3 \end{pmatrix}^t \begin{pmatrix} 1 \\ 0 \end{pmatrix} = \begin{pmatrix} x \\ y \end{pmatrix} \text{ in } \text{GF}(5)^{2 \times 2}, \\ & \text{or } t = -1 \text{ if } x = y = 0, \\ \pi: \quad & a[x][y] \quad \leftarrow a[x'][y'], \text{ with } \begin{pmatrix} x \\ y \end{pmatrix} = \begin{pmatrix} 0 & 1 \\ 2 & 3 \end{pmatrix} \begin{pmatrix} x' \\ y' \end{pmatrix}, \\ \chi: \quad & a[x] \quad \leftarrow a[x] + (a[x+1] + 1)a[x+2], \\ \iota: \quad & a \quad \leftarrow a + \text{RC}[i_t]. \end{aligned}$$

FIG. 1B
(Prior Art)




FIG. 1C
(Prior Art)




FIG. 1D
(Prior Art)




FIG. 1F
(Prior Art)




FIG. 1E
(Prior Art)




FIG. 2




FIG. 3

400




FIG. 4

500

501 Initialization: Save a copy of the unprocessed State

Mov ZMMTMP1, ZMM1

Mov ZMMTMP2, ZMM2

Mov ZMMTMP3, ZMM3

Mov ZMMTMP4, ZMM4

502 KECCAK_THETA: Perform θ function and 1st part of ρ rotate function.

Keccak_Theta ZMM1, ZMM2, ZMM3

Keccak_Theta ZMM2, ZMM3, ZMM4

Keccak_Theta ZMM3, ZMM1, ZMMTMP1

Keccak_Theta ZMM4, ZMMTMP1, ZMMTMP2

503 KECCAK_ROUND: Complete ρ rotate function and perform π, χ, ι Functions

Keccak_Rnd ZMM1, ZMMTMP4, ZMMRC

Keccak_Rnd ZMM2, ZMMTMP1, ZMMRC

Keccak_Rnd ZMM3, ZMMTMP2, ZMMRC

Keccak_Rnd ZMM4, ZMMTMP3, ZMMRC

FIG. 5




FIG. 6




FIG. 7




FIG. 8

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2011/064640**A. CLASSIFICATION OF SUBJECT MATTER****G06F 9/30(2006.01)i, G06F 9/305(2006.01)1, G06F 21/22(2006.01)1, H04L 9/28(2006.01)1**

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 9/30; G06F 15/16; G06F 21/24; H04L 9/28; H04L 9/32; G06F 7/04; G06F 21/00

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: hash, KECCAK, partition, subcube, plurality of registers

C. DOCUMENTS CONSIDERED TO BE RELEVANT

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
A	US 2011-0283110 A1 (DAPKUS PETER et al.) 17 November 2011 See abstract; paragraphs [0122]-[0134] ; figures 4-7.	1-20
A	US 2011-0138192 A1 (KÖCHER PAUL c. et al.) 09 June 2011 See abstract; paragraphs [0048]-[0053] ; figures 4-6.	1-20
A	US 2011-0040977 A1 (FARRUGIA AUGUSTIN J. et al.) 17 February 2011 See abstract; paragraphs [0028]-[0031] ; figures 2, 3.	1-20
A	US 2008-0279369 A1 (PALMER TOMAS EARL) 13 November 2008 See abstract; paragraphs [0021]-[0069] ; figures 5-10.	1-20



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

25 JULY 2012 (25.07.2012)

Date of mailing of the international search report

27 JULY 2012 (27.07.2012)

Name and mailing address of the ISA/KR


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

Facsimile No. 82-42-472-7140

Authorized officer

Ko, Yeon Hwa

Telephone No. 82-42-481-8569



INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2011/064640

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2011-0283110 A1	17.11.2011	US 2011-282908 A1 US 2011-283356 A1	17.11.2011 17.11.2011
US 2011-0138192 A1	09.06.2011	WO 2011-068996 A1	09.06.2011
US 2011-0040977 A1	17.02.2011	None	
US 2008-0279369 A1	13.11.2008	None	