



(51) International Patent Classification:
G06F 9/30 (2006.01)

(21) International Application Number:
PCT/US2017/035191

(22) International Filing Date:
31 May 2017 (31.05.2017)

(25) Filing Language: English

(26) Publication Language: English

(30) Priority Data:
15/196,712 29 June 2016 (29.06.2016) US

(71) Applicant: **INTEL CORPORATION** [US/US]; 2200 Mission College Boulevard, Santa Clara, California 95054 (US).

(72) Inventors: **SURESH, Vikram B.**; 19962 NW Gloucester Way, Hillsboro, Oregon 97124 (US). **SATPATHY, Sudhir K.**; 2111 NE 25th Avenue, Hillsboro, Oregon 97124 (US). **MATHEW, Sanu K.**; 2620 NE Katie Drive, Hillsboro, Oregon 97124 (US).

(74) Agent: **PORTNOVA, Marina** et al.; Lowenstein Sandler LLP, 65 Livingston Avenue, Roseland, New Jersey 07068 (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, DJ, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IR, IS, JP, KE, KG, KH, KN, KP, KR, KW, KZ, LA, LC, LK, LR, LS, 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, SA, 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, ST, 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,

(54) Title: A LOW CLOCK-ENERGY 3-PHASE LATCH-BASED CLOCKING SCHEME




Figure 1

100

(57) Abstract: A processing system includes a processor and a hardware accelerator, communicatively coupled to the processor, comprising a plurality of circuits to perform a plurality rounds of computation, wherein the plurality of circuits comprise a first set of level-sensitive latches enabled by a first clock signal to store data associated with a first round of the plurality of rounds of computation and a second set of level-sensitive latches enabled by a second clock signal to store data associated with a second round of the plurality of rounds of computation, and wherein a duty cycle of the first clock signal and a duty cycle of the second clock signal are non-overlapping.

TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW,
KM, ML, MR, NE, SN, TD, TG).

Declarations under Rule 4.17:

- *as to applicant's entitlement to apply for and be granted a patent (Rule 4.17(ii))*
- *as to the applicant's entitlement to claim the priority of the earlier application (Rule 4.17(iii))*

Published:

- *with international search report (Art. 21(3))*

A LOW CLOCK-ENERGY 3-PHASE LATCH-BASED CLOCKING SCHEME

TECHNICAL FIELD

[0001] The present disclosure relates to hardware accelerators and, more specifically, to a processing system including a processor employing energy-efficient hardware accelerators with a low energy-consumption clock.

BACKGROUND

[0002] Bitcoin is a type of digital currency used in peer-to-peer transactions. The use of Bitcoin in transactions may eliminate the need for intermediate financial institutes because Bitcoin may enforce authenticity and user anonymity by employing digital signatures. Bitcoin resolves the “double spending” problem (namely, using the same Bitcoin more than once by a same entity in different transactions) using block chaining, whereas a public ledger records all the transactions that occur within the Bitcoin currency system. Every block added to the block chain validates a new set of transactions by compressing a 1024-bit message which includes a cryptographic root (e.g., the Merkle root) of the transaction along with bits representing other information such as, for example, a time stamp associated with the transaction, a version number, a target, the hash value of the last block in the block chain and a nonce. The process of validating transactions and generating new blocks of the block chain is commonly referred to as Bitcoin mining.

BRIEF DESCRIPTION OF THE DRAWINGS

[0003] The disclosure will be understood more fully from the detailed description given below and from the accompanying drawings of various embodiments of the disclosure. The drawings, however, should not be taken to limit the disclosure to the specific embodiments, but are for explanation and understanding only.

[0004] **Figure 1** illustrates a processing system to perform Bitcoin mining by employing energy-efficient hardware accelerators according to an embodiment of the present disclosure.

[0005] **Figure 2** illustrates a process to hash a 1024-bit message into a hash value using three stages of SHA-256 hash in Bitcoin mining.

[0006] **Figure 3A** illustrates three consecutive rounds of compression calculation using three generated non-overlapping clocks to enable level-sensitive latches according to an embodiment of the present disclosure.

[0007] **Figure 3B** illustrates the three generated clock signals with respect to the system clock according to an embodiment of the present disclosure.

[0008] **Figure 3C** illustrates a clock generator circuit that may produce three generated clocks (CLK1, CLK2, CLK3) according to an embodiment of the present disclosure.

[0009] **Figure 4** is a block diagram of a method to use 3-phase clock to perform SHA-256 in Bitcoin mining according to an embodiment of the present disclosure.

[0010] **Figure 5A** is a block diagram illustrating a micro-architecture for a processor including heterogeneous core in which one embodiment of the disclosure may be used.

[0011] **Figure 5B** is a block diagram illustrating an in-order pipeline and a register renaming stage, out-of-order issue/execution pipeline implemented according to at least one embodiment of the disclosure.

[0012] **Figure 6** illustrates a block diagram of the micro-architecture for a processor that includes logic in accordance with one embodiment of the disclosure.

[0013] **Figure 7** is a block diagram illustrating a system in which an embodiment of the disclosure may be used.

[0014] **Figure 8** is a block diagram of a system in which an embodiment of the disclosure may operate.

[0015] **Figure 9** is a block diagram of a system in which an embodiment of the disclosure may operate.

[0016] **Figure 10** is a block diagram of a System-on-a-Chip (SoC) in accordance with an embodiment of the present disclosure

[0017] **Figure 11** is a block diagram of an embodiment of an SoC design in accordance with the present disclosure.

[0018] **Figure 12** illustrates a block diagram of one embodiment of a computer system.

DETAILED DESCRIPTION

[0019] The reward for a successful Bitcoin mining is the generation of a certain number of new Bitcoins (e.g., 25 Bitcoins) and the service fee associated with the transactions validated during the mining process. Each Bitcoin may be exchanged for currencies in circulation (e.g., U.S. dollars) or used in transactions with merchants that accept Bitcoins. Bitcoin mining may be associated with certain costs such as, for example, the computing resources consumed to perform Bitcoin mining operations. The most expensive operation in Bitcoin mining involves the computationally-intensive task of determining the validity of a 32-bit nonce. The nonce is a number or a string of bits that is used only once. A 32-bit nonce is a number (or a string of bits) that is represented by 32 bits. The 32-bit nonce may be part of a 1024-bit input message that may also include the Merkle root, the hash of the last chain block, and other parameters. The 1024-bit message may be hashed using three stages of a secure hash algorithm (e.g., SHA-256) to produce a 256-bit hash value that may be compared to a target value also contained in the input message to determine the validity of the nonce. The operations to calculate the hash value are commonly performed on hardware accelerators (e.g., the SHA-256 hash may be performed on application-specific integrated circuits (ASICs)) and may consume a lot of power. The power consumption by the hardware accelerators is the recurring cost for the Bitcoin mining. Embodiments of the present disclosure provide technical solutions including hardware accelerators to perform energy-efficient Bitcoin mining using energy-efficient clock system.

[0020] Embodiments of the present disclosure may include a clock generator that generates three non-overlapping clock signals to enable level-sensitive latches that are used to store intermediate hash values, messages, and parameters of compression functions used in hardware circuits implementing the SHA-256 hash, whereas a level-sensitive latch is a circuit that follows its input as long as the latch is enabled. By replacing commonly used flip-flops (which are sensitive to edges of enable signal) enabled by a single system clock with the level-sensitive latches enabled by the duty cycles of the three generated clock signals, the energy consumption relating to state switching may be improved by 33%.

[0021] Bitcoin mining operations include operations to generate a 256-bit hash value from a 1024-bit message. The operations are part of cryptographic hash that is one-way (very hard to reverse) and collision-resistant. The hash operations may include two stages (stage-0 and stage-1) of SHA-256 hash to compress a 1024-bit input message into intermediate results, followed by another round (stage-2) of SHA-256 hash applied to the intermediate

results generated by the first two stages of SHA-256 hash. The 1024-bit input message to the three stages of SHA-256 hash contains header information, a 32-bit nonce, and padding bits. The padding bits may include 1s and 0s that are generated using a padding generation formulae. The 32-bit nonce is incremented every cycle of the Bitcoin mining process to generate an updated input message, where each cycle takes approximate 10 minutes. A valid nonce is identified if the final hash value contains a certain number of leading zeros. A miner may use the valid nonce as a proof of a successful Bitcoin mining.

[0022] The software application of Bitcoin mining may be implemented on a processing system including processors executing Bitcoin mining applications and dedicated hardware accelerators such as, for examples, ASICs containing clusters of SHA engines that run in parallel to deliver high-performance SHA-256 hash operations. The clusters of SHA engines may consume a lot of powers (e.g., at a rate of greater than 200W). Embodiments of the present disclosure include energy-efficient ASIC-based SHA engines that consume less power for Bitcoin mining operations.

[0023] **Figure 1** illustrates a processing system 100 to perform Bitcoin mining by employing energy-efficient hardware accelerators including SHA-256 engines according to an embodiment of the present disclosure. As shown in **Figure 1**, processing system 100 (e.g., a system-on-a-chip (SOC)) may include a processor 102 and ASICs 104 communicatively coupled to processor 102 via a bus 106. Processor 102 may be a hardware processing device such as, for example, a central processing unit (CPU) or a graphic processing unit (GPU) that includes one or more processing cores (not shown) to execute software applications. Processor 102 may execute a Bitcoin mining application 108 which may include operations to employ multi-stage of SHA-256 hash to compress a 1024-bit input message. For example, Bitcoin mining application 108 may delegate the calculation of the three stages of SHA-256 hash to hardware accelerators such as, for example, SHA-256 engines 110 to perform stage-0 hash, SHA-256 engines 112 to perform stage-1 hash, and SHA-256 engines 114 to perform stage-2 hash. These SHA-256 engines are implemented on one or more ASICs 104. Each one of ASICs 104 may contain multiple SHA-256 engines (e.g., > 1000) that run in parallel. Embodiments of the present disclosure may take advantage of characteristics of different stages of SHA-256 hash to implement them in energy efficient manners to save power consumption in Bitcoin mining.

[0024] The three stages of SHA-256 hash engines 110, 112, 114 are used to convert a 1024 input message into a 256-bit hash output that is compared to a 256-bit target value to determine whether a 32-bit nonce in the input message is a valid proof of successful Bitcoin

mining. Each one of the SHA-256 hash engines 110, 112, 114 may receive a 512-bit input and include 64 rounds of calculation which uses the 512-bit input as a key to compress eight 32-bit state (A, B, C, D, E, F, G, H) stored in eight registers (a, b, c, d, e, f, g, h). Each round of the compression is achieved by applying compression functions to the eight states.

[0025] In some implementations, each one of the 32-bit registers is implemented using 32 flip-flops that are enabled by rising edges (or falling edges) of a system clock 116. A flip-flop is a circuit that includes an output of two stable states (0, 1) that switch at the rising edge (or the falling edge) of an enable signal. The enable signal is the system clock (CLK) in current implementations of SHA-256 engines of Bitcoin mining. The flip-flop is used to store one bit information. Thus, in a hardware SHA-256 engine, 256 flip-flops are employed to store the eight 32-bit states (A, B, C, D, E, F, G, H), and 512 flip-flops are employed to store the 512-bit input to the SHA-256 engine. SHA-256 engine may include circuits to perform up to 64 rounds of compression. In some implementations, the compression computation of SHA-256 engine may be implemented as a pipeline, wherein each round of compression computation may employ $256 + 512$ flip-flops and the SHA-256 engine employs $64 * (256 + 512)$ flip-flops. The output states of these flip-flops are changed at the rising edges (or falling edges) of the clock signal (CLK). Between two consecutive rounds of compression calculation, the states (A, B, C, D, E, F, G, H) stored in eight registers (a, b, c, d, e, f, g, h) may need to be shifted between registers. Thus, there may be up to 672 bits that need to be shifted per round of compression computation, which create 672 potential minimum delay paths due to direct flop-to-flop shift operations. Since the shift operation may cause the new data arrive before the stored data are shifted out due to clock skews (i.e., imperfect edges in clock signals), the input data to the registers may need to be held stable for a hold time. Thus, due to clock skew between consecutive rounds, minimum delay buffers need to be inserted on these minimum delay paths to provide a guard band for any potential hold time violations. These minimum delay buffers increase the total power consumption by the system.

[0026] Instead of using flip-flops directly enabled by system clock 116, embodiments of the present disclosure employ level-sensitive latches to implement registers that store states (A, B, C, D, E, F, G, H) and the 512-bit input values associated with the three stages of SHA-256 engines 110, 112, 114. A level-sensitive latch includes an output that follows the input signal when the latch is enabled. The processing system 100 may include a clock generator 118 that may generate three non-overlapping clock signals CLK1, CLK2, CLK3 that are provided to these level-sensitive latches, as the enable signals, employed in different

rounds of compression calculation. Non-overlapping clock signals are those that have no overlapping duty cycles at any given time. The combined energy consumption by these generated clocks (CLK1, CLK2, CLK3) is significantly reduced compared to using flip-flops driven by edges of system clock (CLK).

[0027] **Figure 2** illustrates a process 200 to hash a 1024-bit message into a 256-bit hash value using three stages of SHA-256 hash during Bitcoin mining. In SHA-256 hash, the hash value may be stored in eight state registers (*a, b, c, d, e, f, g, h*) associated with each SHA-256 engine, where each of the state registers is a hardware register that stores a 32-bit word referred to as a state (represented by A, B, C, D, E, F, G, H). The initial values of these states can be 32-bit constants. Alternatively, the state registers may initially store a hash value calculated from a previous iteration of the hashing process. The states (A, B, C, D, E, F, G, H) are updated during SHA-256 hash calculation to generate a 256-bit hash value as the output. SHA-256 hash consumes a block of 512-bit message and compresses it into a 256-bit hash (*A – H*) stored in state registers (*a – h*). The Bitcoin mining process employs three stages of SHA-256 hash to convert the 1024-bit input message to a 256-bit hash value that may be compared to a target value to determine whether a Bitcoin has been identified.

[0028] The SHA-256 hash may include 64 rounds (identified as round 0, 1, . . . , 63) of applications of compression functions to the states stored in state registers. The compression function employs a 512-bit input value to manipulate the contents stored in registers (*a – h*). Table 1 illustrates the 64 rounds of the SHA-256 operations as applied to the states stored in registers (*a – h*) to generate a hash value that can be used to determine if a valid nonce is found as a proof of the identification of a Bitcoin.

Table 1

• Apply the SHA-256 compression function to update registers $a, b, \dots, h$ •

For $j = 0$ to 63

{

Compute $Ch(e, f, g)$, $Maj(a, b, c)$, $\Sigma_0(a)$, $\Sigma_1(e)$, and W_j (see Definitions below)

$T_1 \leftarrow h \div \Sigma_1(e) \div Ch(e, f, g) \div K_j \div W_j$

$T_2 \leftarrow \Sigma_0(a) \div Maj(a, b, c)$

$h \leftarrow g$

$g \leftarrow f$

$f \leftarrow e$

$e \leftarrow d \div T_1$

$d \leftarrow c$

$c \leftarrow b$

$b \leftarrow a$

$a \leftarrow T_1 \div T_2$

}

where logic functions $Ch(x, y, z)$, $Maj(x, y, z)$, $\Sigma_0 x$, $\Sigma_1 x$ are compression functions that are defined according the SHA-256 specification, and each registers ($a - h$) is initiated with a 32-bit initial values, and $W_j, j = 0, \dots, 63$, are 32-bit values derived from a 512-bit message which can be part of the 1024-bit input message of the Bitcoin mining.

[0029] As shown in **Figure 2**, the process of the Bitcoin mining 200 starts with a 1024-bit message 218. The 1024-bit input message 218 may be composed of header information, a nonce 212, and padding bits 214 that make the input message 218 to the length of 1024 bits. The header information may include a 32-bit version number 202, a 256-bit hash value 204 generated by the immediate preceding block in the block chain of Bitcoin public ledger, a 256-bit Merkle root 206 of the transaction, a 32-bit time stamp 208, and a 256-bit target value 210. Version number 202 is an identifier associated with the version of the block chain. Hash value 204 is the hashing result from the immediate preceding block in the block chain recorded in the public ledger. Merkle root 206 is 256-bit hash based on all of the transactions in the block. Time stamp 208 represents the current time when the Bitcoin mining process starts. Target value 210 represents a threshold value that the resulting hash value generated by the Bitcoin mining is compared to. If the resulting hash value (“hash out”) is smaller than the target value 210, the nonce 212 in the input message 218 is identified as a valid nonce that can be used as the proof of the identification of a Bitcoin. If the final result is no less than the target value 210, the nonce 212 is determined to be invalid, or the

Bitcoin mining failed to find a Bitcoin. The value of nonce 212 may be updated (e.g., incremented by one), and the Bitcoin mining process is repeated to determine the validity of the updated nonce.

[0030] In one embodiment, instead of comparing the final hashing result with the target value, Bitcoin mining application may determine whether the hash out has a minimum number of leading zeros. The minimum number of leading zeros may ensure that the final hashing value is smaller than the target value. The target value (or the number of leading zeros) may be changed to adjust the complexity of Bitcoin mining: decreasing the target value decreases the probability of finding a valid nonce and hence increases the overall search space to generate a new block in the block chain. By modifying the target value 210, the complexity of the Bitcoin mining is adjusted to ensure that the time used to find a valid nonce is relative constant (approximately 10 minutes). For a given header, the Bitcoin mining application may sweep through the search space of 2^{32} possibilities to find a valid nonce. The Bitcoin mining process includes a series of mining iterations to sweeping through these possibilities of valid nonce. The header information is kept the same through these mining iterations while the nonce 212 is incremented by one.

[0031] Each Bitcoin mining calculation to find a valid nonce may include three stages (stage-0 – stage-2) of SHA-256 hash calculations. Referring to **Figure 2**, at stage-0 SHA-256 hash, the state (A, B, C, D, E, F, G, H) stored in state registers (a, b, c, d, e, f, g, h) may be initiated with eight 32-bit constants. Stage-0 SHA-256 hash may receive a 512-bit input message including the 32-bit version number 202, 256-bit hash value 204 from the last block in the block chain, and a portion (the first 224 bits) of Merkle root 206. Stage-0 SHA-256 hash may produce a first 256-bit intermediate hash value. The first intermediate hash value is then employed to initiate the state registers A – H of the stage-1 SHA-256 hash. The 512-bit input message to the stage-1 SHA-256 hash may include the rest portion (32 bits) of the Merkle root 206, 32-bit time stamp 208, 256-bit target value 210, 32-bit nonce 212, and 128 padding bits 214. Stage-1 SHA-256 hash may produce a second 256-bit intermediate hash value.

[0032] At the stage-2 SHA-256 hash, the state registers (a, b, c, d, e, f, g, h) of the stage-2 SHA-256 hash may be set with the 256-bit constant same as the constant input of stage-0 SHA-256 hash. The 512-bit input message to the stage-2 SHA-256 hash may include the second 256-bit intermediate hash result (from the stage-1 SHA-256 hash output) combined with 256 padding bits to make a 512-bit input message to the stage-2 SHA-256 hash. The stage-2 SHA-256 hash may produce a third 256-bit hash value as the hash out for

the three stages of SHA-256 hash. The Bitcoin mining application may then determine whether the hash out is smaller than the target value 210. If the hash out is smaller than the target value 210, the nonce 212 in the input message is identified as a valid nonce. If the hash out is no less than the target value 210, the nonce 212 is an invalid nonce. After the determination, nonce 212 is incremented to repeat the process to determine the validity of the updated nonce 212 using the process as shown in **Figure 2**.

[0033] As shown in Table 1, SHA-256 includes 64 rounds of compression calculation that includes applying compression functions to data stored in state registers (a, b, c, d, e, f, g, h) and registers that store the 512-bit input. The compression functions $Ch(x, y, z)$, $Maj(x, y, z)$, $\sum_0 x$, $\sum_1 x$ along with register shifts in a round of compression calculation are implemented in message digest circuits. Thus, during a round of compression calculation (e.g., round i , where i is an integer index), the state (A, B, C, D, E, F, G, H) stored in registers (a, b, c, d, e, f, g, h) associated with round i may be provided to a message digest circuit associated with the same round. The calculated results from the message digest for round i are then provided, in a pipeline manner, to registers (a, b, c, d, e, f, g, h) associated with round $i+1$. The states (A, B, C, D, E, F, G, H) stored in registers (a, b, c, d, e, f, g, h) associated with round $i+1$ may then be processed by message digest circuit associated round $i+1$.

[0034] **Figure 3A** illustrates three consecutive rounds of compression calculation using three generated non-overlapping clocks to enable level-sensitive latches according to an embodiment of the present disclosure. A SHA-256 engine may include hardware circuits to perform 64 rounds of compression calculation. **Figure 3A** shows three consecutive rounds 302, 304, 306 of the 64 rounds. In one embodiment, the circuit to perform a round (e.g., round i , $i+1$, or $i+2$) of compression calculation may include 256 one-bit level-sensitive latches 308, 310, 312 (referred to as 256-bit latch) that are employed to implement the eight registers (a, b, c, d, e, f, g, h) for storing states (A, B, C, D, E, F, G, H) of the round. In one embodiment, eight registers (a, b, c, d, e, f, g, h) are used for all 64 rounds of compression calculation in a SHA-256 engine. Thus, the SHA-256 engine generates a hash value in every 64 system clock cycles. In another embodiment, each round is associated with respective eight registers (a, b, c, d, e, f, g, h) to allow the 64 rounds of SHA-256 hash performed as a pipeline. Thus, the SHA-256 engine may generate a hash value in alternative system clock cycles.

[0035] For example, as shown in **Figure 3A**, the SHA-256 engine is implemented as a pipeline, and each one of rounds 302, 304, 306 may be associated with a respective latch

308, 310, 312. The state A stored in latches implementing register a associated with round i is an input to latches implementing register b for storing state B associated with round $i+1$. Similarly, states (B, C, D, E, F, G) stored in latches implementing registers (b, c, d, e, f, g) associated with round i is the input to latches implementing registers (c, d, e, f, g, h) for storing state (C, D, E, F, G, H) associated with round $i+1$. Thus, the output of latches associated with round $i+1$ are updated to the input (or the states of round i) when the latches associated with round $i+1$ are enabled.

[0036] The circuit for the round may also include message digest circuit 314, 316 which may also include level-sensitive latches to store parameters of compression functions. Thus, the hardware implementation of each compression round of SHA-256 hash may include multiple level-sensitive latches to store data (e.g., states and parameters of compression functions) that may be updated when these latches are enabled. Because the circuit of a level-sensitive latch is about half of the size of a flip-flop circuit, the hardware implementation using level-sensitive latches may reduce the overall circuit area, switching capacitance, and clock load.

[0037] To achieve the same results as the hardware implementation using flip-flops using the system clock as the enable signal, three generated clock signals (CLK1, CLK2, CLK3) are used to enable three consecutive rounds of compression calculation. **Figure 3B** illustrates the three generated clock signals (CLK1, CLK2, CLK3) with respect to the system clock (CLK) according to an embodiment of the present disclosure. A clock generator 116 (as shown in **Figure 1**) may receive the system clock (CLK) as an input and generate three clock signals (CLK1, CLK2, CLK3).

[0038] As shown in **Figure 3B**, the system clock (CLK) may have a period which is the time duration (T) of one clock cycle. Within each period, the system clock (CLK) may include a half period (0.5T) of duty cycle (CLK is high) and half period of non-duty cycle (CLK is low). The three generated clocks (CLK1, CLK2, CLK3) each may have a period (T') which is at a ratio (e.g., one and half time) of the period of the system clock (e.g., $T' = 1.5T$). Within a period (T'), the generated clocks (CLK1, CLK2, CLK3) may include a duty cycle with a duration of at a ratio (e.g., one quarter) of the system clock (CLK) period (e.g., 0.25T) and a non-duty cycle with a duration at a ratio (e.g., one and one quarter) of the system clock (CLK) period (e.g., 1.25T). The duty cycles of the generated clocks (CLK1, CLK2, CLK3) are non-overlapping – namely, when one of the generated clocks (e.g., CLK1) is in the duty cycle (at high state), the other two generated clocks (e.g., CLK2, CLK3) are in the non-duty cycle (at low state). In one embodiment, as shown in **Figure 3B**, the duty

cycles between any two of the generated clocks (e.g., CLK1 and CLK2, CLK1 and CLK3, CLK2 and CLK3) are separated by time gaps (e.g., at least $0.25T$).

[0039] Clock generator 118 may be implemented as a logic circuit on ASICs 104.

Figure 3C illustrates a clock generator circuit 340 that may produce three generated clocks (CLK1, CLK2, CLK3) according to an embodiment of the present disclosure. As shown in **Figure 3C**, the frequency of a system clock (CLK) may be doubled to generate an intermediate clock (CLK_{2x}) having a period of $0.5T$, where T is the period of the system clock (CLK). The intermediate clock (CLK_{2x}) is provided to the enabling pins of three D latches 320, 322, 324 which are connected in a master-slave manner and are combined with an XOR gate 338 to generate second intermediate clocks C_1 , C_2 , C_3 which have a period of $1.5T$. The inverse of the intermediate clock (CLK_{2x}) may be provided to the enabling pins of another three D latches 332, 334, 336 which are also connected in a master-slave manner to generate third intermediate clocks (C'_1 , C'_2 , C'_3) which have a period of $1.5T$. The pairs of (C_1 , C'_1), (C_2 , C'_2), and (C_3 , C'_3) are respectively provided to AND gates 332, 334, 336 to generate clocks CLK1, CLK2, CLK3 that have a period of $1.25T$ and non-overlapping duty cycles over time.

[0040] Using three non-overlapping clocks (CLK1, CLK2, CLK3) allow the latches in each round to be enabled once every $T' = 1.5T$, thereby reducing the throughput by about 33%. However, since the switching capacitance in the sequential logic decreases by about 50% due to employing level-sensitive latches (vs. flip-flops), the performance per watt, as a whole, improves about 33%. The gaps among the duty cycles of the generated clocks (CLK1, CLK2, CLK3) may avoid simultaneous raise and fall edges on the clocks of two consecutive rounds and provide a $0.25T$ margin on the minimum delay paths that may guarantee reliable operation as long as the clock skew is less than $0.25T$. This timing margin may eliminate the need for using minimum delay buffers to guarantee the hold time for data stored in registers.

[0041] Although **Figures 1-3** illustrate embodiments using 3-Phase generated clocks in the context of Bitcoin mining application, aspects of the disclosure may be applicable to any logic circuits to perform sequential operations (e.g., SHA used in applications other than the Bitcoin mining applications) where N -phase clock ($N \geq 2$) can be generated to enable the sequential operations. The N -phase clock may have a period at a ratio of the period (T) of the system clock. The duty circles of the N -phase clock may not overlap over the time axis.

[0042] **Figure 4** is a block diagram of a method 400 to use 3-phase clock to perform SHA-256 in Bitcoin mining according to an embodiment of the present disclosure. Method

400 may be performed by processing logic that may include hardware (e.g., circuitry, dedicated logic, programmable logic, microcode, etc.), software (such as instructions run on a processing device, a general purpose computer system, or a dedicated machine), firmware, or a combination thereof. In one embodiment, method 400 may be performed, in part, by processing logics of processor 102 and ASIC 104 as shown in **Figure 1**.

[0043] For simplicity of explanation, the method 400 is depicted and described as a series of acts. However, acts in accordance with this disclosure can occur in various orders and/or concurrently and with other acts not presented and described herein. Furthermore, not all illustrated acts may be performed to implement the method 400 in accordance with the disclosed subject matter. In addition, those skilled in the art will understand and appreciate that the method 400 could alternatively be represented as a series of interrelated states via a state diagram or events.

[0044] Referring to **Figure 4**, processor 102 may be communicatively coupled to ASICs 104 which may include clusters of SHA-256 engines to perform stage-0, stage-1, and stage-2 SHA-256 hash for the Bitcoin mining application including the speculative nonce pre-selection. Each one of the stage-0, stage-1, or stage-2 SHA-256 hash may include up to 64 rounds of compression calculation that may be implemented using the 3-phase generated clocks and level-sensitive latches.

[0045] At 402, a SHA-256 engine implemented on a hardware accelerator may perform a first round of compression computation of a SHA-256 hash by employing a first set of level-sensitive latches enabled by a first clock signal.

[0046] At 404, the SHA-256 engine may perform a second round of the compression computation of the SHA-256 hash by employing a second set of level-sensitive latches enabled by a second clock signal.

[0047] At 406, the SHA-256 engine may perform a third round of the compression computation of the SHA-256 hash by employing a third set of level-sensitive latches enabled by a third clock signal, in which the first clock signal, the second clock signal, and the third clock signal are generated from a system clock.

[0048] **Figure 5A** is a block diagram illustrating a micro-architecture for a processor 500 that implements the processing device including heterogeneous cores in accordance with one embodiment of the disclosure. Specifically, processor 500 depicts an in-order architecture core and a register renaming logic, out-of-order issue/execution logic to be included in a processor according to at least one embodiment of the disclosure.

[0049] Processor 500 includes a front end unit 530 coupled to an execution engine unit 550, and both are coupled to a memory unit 570. The processor 500 may include 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, processor 500 may include a special-purpose core, such as, for example, a network or communication core, compression engine, graphics core, or the like. In one embodiment, processor 500 may be a multi-core processor or may part of a multi-processor system.

[0050] The front end unit 530 includes a branch prediction unit 532 coupled to an instruction cache unit 534, which is coupled to an instruction translation lookaside buffer (TLB) 536, which is coupled to an instruction fetch unit 538, which is coupled to a decode unit 540. The decode unit 540 (also known as a 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 decoder 540 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. The instruction cache unit 534 is further coupled to the memory unit 570. The decode unit 540 is coupled to a rename/allocator unit 552 in the execution engine unit 550.

[0051] The execution engine unit 550 includes the rename/allocator unit 552 coupled to a retirement unit 554 and a set of one or more scheduler unit(s) 556. The scheduler unit(s) 556 represents any number of different schedulers, including reservations stations (RS), central instruction window, etc. The scheduler unit(s) 556 is coupled to the physical register file(s) unit(s) 558. Each of the physical register file(s) units 558 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, etc., status (e.g., an instruction pointer that is the address of the next instruction to be executed), etc. The physical register file(s) unit(s) 558 is overlapped by the retirement unit 554 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.).

[0052] In one implementation, processor 500 may be the same as processor 102 described with respect to **Figure 1**.

[0053] Generally, the architectural registers are visible from the outside of the processor or from a programmer's perspective. The registers are not limited to any known particular type of circuit. Various different types of registers are suitable as long as they are capable of storing and providing data as described herein. 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. The retirement unit 554 and the physical register file(s) unit(s) 558 are coupled to the execution cluster(s) 560. The execution cluster(s) 560 includes a set of one or more execution units 562 and a set of one or more memory access units 564. The execution units 562 may perform various operations (e.g., shifts, addition, subtraction, multiplication) and operate on various types of data (e.g., scalar floating point, packed integer, packed floating point, vector integer, vector floating point).

[0054] While some embodiments may include a number of execution units dedicated to specific 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) 556, physical register file(s) unit(s) 558, and execution cluster(s) 560 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 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) 564). It should also be understood that where separate pipelines are used, one or more of these pipelines may be out-of-order issue/execution and the rest in-order.

[0055] The set of memory access units 564 is coupled to the memory unit 570, which may include a data prefetcher 580, a data TLB unit 572, a data cache unit (DCU) 574, and a level 2 (L2) cache unit 576, to name a few examples. In some embodiments DCU 574 is also known as a first level data cache (L1 cache). The DCU 574 may handle multiple outstanding cache misses and continue to service incoming stores and loads. It also supports maintaining cache coherency. The data TLB unit 572 is a cache used to improve virtual address translation speed by mapping virtual and physical address spaces. In one exemplary embodiment, the memory access units 564 may include a load unit, a store address unit, and a

store data unit, each of which is coupled to the data TLB unit 572 in the memory unit 570. The L2 cache unit 576 may be coupled to one or more other levels of cache and eventually to a main memory.

[0056] In one embodiment, the data prefetcher 580 speculatively loads/prefetches data to the DCU 574 by automatically predicting which data a program is about to consume. Prefetching may refer to transferring data stored in one memory location of a memory hierarchy (e.g., lower level caches or memory) to a higher-level memory location that is closer (e.g., yields lower access latency) to the processor before the data is actually demanded by the processor. More specifically, prefetching may refer to the early retrieval of data from one of the lower level caches/memory to a data cache and/or prefetch buffer before the processor issues a demand for the specific data being returned.

[0057] The processor 500 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).

[0058] 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 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).

[0059] 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 embodiment of the processor also includes a separate instruction and data cache units and a shared L2 cache unit, 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 the cache may be external to the core and/or the processor.

[0060] **Figure 5B** is a block diagram illustrating an in-order pipeline and a register renaming stage, out-of-order issue/execution pipeline implemented by processing device 500 of **Figure 5A** according to some embodiments of the disclosure. The solid lined boxes in **Figure 5B** illustrate an in-order pipeline, while the dashed lined boxes illustrates a register renaming, out-of-order issue/execution pipeline. In **Figure 5B**, a processor pipeline 500

includes a fetch stage 502, a length decode stage 504, a decode stage 506, an allocation stage 508, a renaming stage 510, a scheduling (also known as a dispatch or issue) stage 512, a register read/memory read stage 514, an execute stage 516, a write back/memory write stage 518, an exception handling stage 522, and a commit stage 524. In some embodiments, the ordering of stages 502-524 may be different than illustrated and are not limited to the specific ordering shown in **Figure 5B**.

[0061] **Figure 6** illustrates a block diagram of the micro-architecture for a processor 600 that includes hybrid cores in accordance with one embodiment of the disclosure. In some embodiments, an instruction in accordance with one embodiment can be implemented to operate on data elements having sizes of byte, word, doubleword, quadword, etc., as well as datatypes, such as single and double precision integer and floating point datatypes. In one embodiment the in-order front end 601 is the part of the processor 600 that fetches instructions to be executed and prepares them to be used later in the processor pipeline.

[0062] The front end 601 may include several units. In one embodiment, the instruction prefetcher 626 fetches instructions from memory and feeds them to an instruction decoder 628 which in turn decodes or interprets them. For example, in one embodiment, the decoder decodes a received instruction into one or more operations called “micro-instructions” or “micro-operations” (also called micro op or uops) that the machine can execute. In other embodiments, the decoder parses the instruction into an opcode and corresponding data and control fields that are used by the micro-architecture to perform operations in accordance with one embodiment. In one embodiment, the trace cache 630 takes decoded uops and assembles them into program ordered sequences or traces in the uop queue 634 for execution. When the trace cache 630 encounters a complex instruction, the microcode ROM 632 provides the uops needed to complete the operation.

[0063] Some instructions are converted into a single micro-op, whereas others need several micro-ops to complete the full operation. In one embodiment, if more than four micro-ops are needed to complete an instruction, the decoder 628 accesses the microcode ROM 632 to do the instruction. For one embodiment, an instruction can be decoded into a small number of micro ops for processing at the instruction decoder 628. In another embodiment, an instruction can be stored within the microcode ROM 632 should a number of micro-ops be needed to accomplish the operation. The trace cache 630 refers to an entry point programmable logic array (PLA) to determine a correct micro-instruction pointer for reading the micro-code sequences to complete one or more instructions in accordance with one embodiment from the micro-code ROM 632. After the microcode ROM 632 finishes

sequencing micro-ops for an instruction, the front end 601 of the machine resumes fetching micro-ops from the trace cache 630.

[0064] The out-of-order execution engine 603 is where the instructions are prepared for execution. The out-of-order execution logic has a number of buffers to smooth out and re-order the flow of instructions to optimize performance as they go down the pipeline and get scheduled for execution. The allocator logic allocates the machine buffers and resources that each uop needs in order to execute. The register renaming logic renames logic registers onto entries in a register file. The allocator also allocates an entry for each uop in one of the two uop queues, one for memory operations and one for non-memory operations, in front of the instruction schedulers: memory scheduler, fast scheduler 602, slow/general floating point scheduler 604, and simple floating point scheduler 606. The uop schedulers 602, 604, 606, determine when a uop is ready to execute based on the readiness of their dependent input register operand sources and the availability of the execution resources the uops need to complete their operation. The fast scheduler 602 of one embodiment can schedule on each half of the main clock cycle while the other schedulers can only schedule once per main processor clock cycle. The schedulers arbitrate for the dispatch ports to schedule uops for execution.

[0065] Register files 608, 610, sit between the schedulers 602, 604, 606, and the execution units 612, 614, 616, 618, 620, 622, 624 in the execution block 611. There is a separate register file 608, 610, for integer and floating point operations, respectively. Each register file 608, 610, of one embodiment also includes a bypass network that can bypass or forward just completed results that have not yet been written into the register file to new dependent uops. The integer register file 608 and the floating point register file 610 are also capable of communicating data with the other. For one embodiment, the integer register file 608 is split into two separate register files, one register file for the low order 32 bits of data and a second register file for the high order 32 bits of data. The floating point register file 610 of one embodiment has 128 bit wide entries because floating point instructions typically have operands from 64 to 128 bits in width.

[0066] The execution block 611 contains the execution units 612, 614, 616, 618, 620, 622, 624, where the instructions are actually executed. This section includes the register files 608, 610, that store the integer and floating point data operand values that the micro-instructions need to execute. The processor 600 of one embodiment is comprised of a number of execution units: address generation unit (AGU) 612, AGU 614, fast ALU 616, fast ALU 618, slow ALU 620, floating point ALU 622, floating point move unit 624. For one

embodiment, the floating point execution blocks 622, 624, execute floating point, MMX, SIMD, and SSE, or other operations. The floating point ALU 622 of one embodiment includes a 64 bit by 64 bit floating point divider to execute divide, square root, and remainder micro-ops. For embodiments of the present disclosure, instructions involving a floating point value may be handled with the floating point hardware.

[0067] In one embodiment, the ALU operations go to the high-speed ALU execution units 616, 618. The fast ALUs 616, 618, of one embodiment can execute fast operations with an effective latency of half a clock cycle. For one embodiment, most complex integer operations go to the slow ALU 620 as the slow ALU 620 includes integer execution hardware for long latency type of operations, such as a multiplier, shifts, flag logic, and branch processing. Memory load/store operations are executed by the AGUs 612, 614. For one embodiment, the integer ALUs 616, 618, 620, are described in the context of performing integer operations on 64 bit data operands. In alternative embodiments, the ALUs 616, 618, 620, can be implemented to support a variety of data bits including 16, 32, 128, 256, etc. Similarly, the floating point units 622, 624, can be implemented to support a range of operands having bits of various widths. For one embodiment, the floating point units 622, 624, can operate on 128 bits wide packed data operands in conjunction with SIMD and multimedia instructions.

[0068] In one embodiment, the uops schedulers 602, 604, 606, dispatch dependent operations before the parent load has finished executing. As uops are speculatively scheduled and executed in processor 600, the processor 600 also includes logic to handle memory misses. If a data load misses in the data cache, there can be dependent operations in flight in the pipeline that have left the scheduler with temporarily incorrect data. A replay mechanism tracks and re-executes instructions that use incorrect data. Only the dependent operations need to be replayed and the independent ones are allowed to complete. The schedulers and replay mechanism of one embodiment of a processor are also designed to catch instruction sequences for text string comparison operations.

[0069] The processor 600 also includes logic to implement store address prediction for memory disambiguation according to embodiments of the disclosure. In one embodiment, the execution block 611 of processor 600 may include a store address predictor (not shown) for implementing store address prediction for memory disambiguation.

[0070] The term “registers” may refer to the on-board processor storage locations that are used as part of instructions to identify operands. In other words, registers may be those that are usable from the outside of the processor (from a programmer’s perspective).

However, the registers of an embodiment should not be limited in meaning to a particular type of circuit. Rather, a register of an embodiment is 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 thirty-two bit integer data. A register file of one embodiment also contains eight multimedia SIMD registers for packed data.

[0071] For the discussions below, the registers are understood to be data registers designed to hold packed data, such as 64 bits wide MMXTM registers (also referred to as ‘mm’ registers in some instances) in microprocessors enabled with MMX technology from Intel Corporation of Santa Clara, California. These MMX registers, available in both integer and floating point forms, can operate with packed data elements that accompany SIMD and SSE instructions. Similarly, 128 bits wide XMM registers relating to SSE2, SSE3, SSE4, or beyond (referred to generically as “SSEx”) technology can also be used to hold such packed data operands. In one embodiment, in storing packed data and integer data, the registers do not need to differentiate between the two data types. In one embodiment, integer and floating point are either contained in the same register file or different register files. Furthermore, in one embodiment, floating point and integer data may be stored in different registers or the same registers.

[0072] Referring now to **Figure 7**, shown is a block diagram illustrating a system 700 in which an embodiment of the disclosure may be used. As shown in **Figure 7**, multiprocessor system 700 is a point-to-point interconnect system, and includes a first processor 770 and a second processor 780 coupled via a point-to-point interconnect 750. While shown with only two processors 770, 780, it is to be understood that the scope of embodiments of the disclosure is not so limited. In other embodiments, one or more additional processors may be present in a given processor. In one embodiment, the multiprocessor system 700 may implement hybrid cores as described herein.

[0073] Processors 770 and 780 are shown including integrated memory controller units 772 and 782, respectively. Processor 770 also includes as part of its bus controller units point-to-point (P-P) interfaces 776 and 778; similarly, second processor 780 includes P-P interfaces 786 and 788. Processors 770, 780 may exchange information via a point-to-point (P-P) interface 750 using P-P interface circuits 778, 788. As shown in **Figure 7**, IMCs 772

and 782 couple the processors to respective memories, namely a memory 732 and a memory 734, which may be portions of main memory locally attached to the respective processors.

[0074] Processors 770, 780 may each exchange information with a chipset 790 via individual P-P interfaces 752, 754 using point to point interface circuits 776, 794, 786, 798. Chipset 790 may also exchange information with a high-performance graphics circuit 738 via a high-performance graphics interface 739.

[0075] 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.

[0076] Chipset 790 may be coupled to a first bus 716 via an interface 796. In one embodiment, first bus 716 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 disclosure is not so limited.

[0077] As shown in **Figure 7**, various I/O devices 714 may be coupled to first bus 716, along with a bus bridge 718 which couples first bus 716 to a second bus 720. In one embodiment, second bus 720 may be a low pin count (LPC) bus. Various devices may be coupled to second bus 720 including, for example, a keyboard and/or mouse 722, communication devices 727 and a storage unit 728 such as a disk drive or other mass storage device which may include instructions/code and data 730, in one embodiment. Further, an audio I/O 724 may be coupled to second bus 720. Note that other architectures are possible. For example, instead of the point-to-point architecture of **Figure 7**, a system may implement a multi-drop bus or other such architecture.

[0078] Referring now to **Figure 8**, shown is a block diagram of a system 800 in which one embodiment of the disclosure may operate. The system 800 may include one or more processors 810, 815, which are coupled to graphics memory controller hub (GMCH) 820. The optional nature of additional processors 815 is denoted in **Figure 8** with broken lines. In one embodiment, processors 810, 815 implement hybrid cores according to embodiments of the disclosure.

[0079] Each processor 810, 815 may be some version of the circuit, integrated circuit, processor, and/or silicon integrated circuit as described above. However, it should be noted that it is unlikely that integrated graphics logic and integrated memory control units would exist in the processors 810, 815. **Figure 8** illustrates that the GMCH 820 may be coupled to a memory 840

that may be, for example, a dynamic random access memory (DRAM). The DRAM may, for at least one embodiment, be associated with a non-volatile cache.

[0080] The GMCH 820 may be a chipset, or a portion of a chipset. The GMCH 820 may communicate with the processor(s) 810, 815 and control interaction between the processor(s) 810, 815 and memory 840. The GMCH 820 may also act as an accelerated bus interface between the processor(s) 810, 815 and other elements of the system 800. For at least one embodiment, the GMCH 820 communicates with the processor(s) 810, 815 via a multi-drop bus, such as a frontside bus (FSB) 895.

[0081] Furthermore, GMCH 820 is coupled to a display 845 (such as a flat panel or touchscreen display). GMCH 820 may include an integrated graphics accelerator. GMCH 820 is further coupled to an input/output (I/O) controller hub (ICH) 850, which may be used to couple various peripheral devices to system 800. Shown for example in the embodiment of **Figure 8** is an external graphics device 860, which may be a discrete graphics device, coupled to ICH 850, along with another peripheral device 870.

[0082] Alternatively, additional or different processors may also be present in the system 800. For example, additional processor(s) 815 may include additional processor(s) that are the same as processor 810, additional processor(s) that are heterogeneous or asymmetric to processor 810, accelerators (such as, e.g., graphics accelerators or digital signal processing (DSP) units), field programmable gate arrays, or any other processor. There can be a variety of differences between the processor(s) 810, 815 in terms of a spectrum of metrics of merit including architectural, micro-architectural, thermal, power consumption characteristics, and the like. These differences may effectively manifest themselves as asymmetry and heterogeneity amongst the processors 810, 815. For at least one embodiment, the various processors 810, 815 may reside in the same die package.

[0083] Referring now to **Figure 9**, shown is a block diagram of a system 900 in which an embodiment of the disclosure may operate. **Figure 9** illustrates processors 970, 980. In one embodiment, processors 970, 980 may implement hybrid cores as described above. Processors 970, 980 may include integrated memory and I/O control logic ("CL") 972 and 982, respectively and intercommunicate with each other via point-to-point interconnect 950 between point-to-point (P-P) interfaces 978 and 988 respectively. Processors 970, 980 each communicate with chipset 990 via point-to-point interconnects 952 and 954 through the respective P-P interfaces 976 to 994 and 986 to 998 as shown. For at least one embodiment, the CL 972, 982 may include integrated memory controller units. CLs 972, 982 may include I/O control logic. As depicted, memories 932, 934 coupled to CLs 972, 982 and I/O devices 914 are also coupled to the control logic 972, 982. Legacy I/O devices 915 are coupled to the chipset 990 via interface 996.

[0084] Embodiments may be implemented in many different system types. **Figure 10** is a block diagram of a SoC 1000 in accordance with an embodiment of the present disclosure. Dashed lined boxes are optional features on more advanced SoCs. In **Figure 10**, an interconnect unit(s) 1012 is coupled to: an application processor 1020 which includes a set of one or more cores 1002A-N and shared cache unit(s) 1006; a system agent unit 1010; a bus controller unit(s) 1016; an integrated memory controller unit(s) 1014; a set or one or more media processors 1018 which may include integrated graphics logic 1008, an image processor 1024 for providing still and/or video camera functionality, an audio processor 1026 for providing hardware audio acceleration, and a video processor 1028 for providing video encode/decode acceleration; an static random access memory (SRAM) unit 1030; a direct memory access (DMA) unit 1032; and a display unit 1040 for coupling to one or more external displays. In one embodiment, a memory module may be included in the integrated memory controller unit(s) 1014. In another embodiment, the memory module may be included in one or more other components of the SoC 1000 that may be used to access and/or control a memory. The application processor 1020 may include a store address predictor for implementing hybrid cores as described in embodiments herein.

[0085] The memory hierarchy includes one or more levels of cache within the cores, a set or one or more shared cache units 1006, and external memory (not shown) coupled to the set of integrated memory controller units 1014. The set of shared cache units 1006 may include one or 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.

[0086] In some embodiments, one or more of the cores 1002A-N are capable of multi-threading. The system agent 1010 includes those components coordinating and operating cores 1002A-N. The system agent unit 1010 may include for example a power control unit (PCU) and a display unit. The PCU may be or include logic and components needed for regulating the power state of the cores 1002A-N and the integrated graphics logic 1008. The display unit is for driving one or more externally connected displays.

[0087] The cores 1002A-N may be homogenous or heterogeneous in terms of architecture and/or instruction set. For example, some of the cores 1002A-N may be in order while others are out-of-order. As another example, two or more of the cores 1002A-N may be capable of execution the same instruction set, while others may be capable of executing only a subset of that instruction set or a different instruction set.

[0088] The application processor 1020 may be a general-purpose processor, such as a Core™ i3, i5, i7, 2 Duo and Quad, Xeon™, Itanium™, Atom™ or Quark™ processor, which

are available from Intel™ Corporation, of Santa Clara, Calif. Alternatively, the application processor 1020 may be from another company, such as ARM Holdings™, Ltd, MIPS™, etc. The application processor 1020 may be a special-purpose processor, such as, for example, a network or communication processor, compression engine, graphics processor, co-processor, embedded processor, or the like. The application processor 1020 may be implemented on one or more chips. The application processor 1020 may be a part of and/or may be implemented on one or more substrates using any of a number of process technologies, such as, for example, BiCMOS, CMOS, or NMOS.

[0089] **Figure 11** is a block diagram of an embodiment of a system on-chip (SoC) design in accordance with the present disclosure. As a specific illustrative example, SoC 1100 is included in user equipment (UE). In one embodiment, UE refers to any device to be used by an end-user to communicate, such as a hand-held phone, smartphone, tablet, ultra-thin notebook, notebook with broadband adapter, or any other similar communication device. Often a UE connects to a base station or node, which potentially corresponds in nature to a mobile station (MS) in a GSM network.

[0090] Here, SOC 1100 includes 2 cores—1106 and 1107. Cores 1106 and 1107 may conform to an Instruction Set Architecture, such as an Intel® Architecture Core™-based processor, an Advanced Micro Devices, Inc. (AMD) processor, a MIPS-based processor, an ARM-based processor design, or a customer thereof, as well as their licensees or adopters. Cores 1106 and 1107 are coupled to cache control 1108 that is associated with bus interface unit 1109 and L2 cache 1110 to communicate with other parts of system 1100. Interconnect 1110 includes an on-chip interconnect, such as an IOSF, AMBA, or other interconnect discussed above, which potentially implements one or more aspects of the described disclosure. In one embodiment, cores 1106, 1107 may implement hybrid cores as described in embodiments herein.

[0091] Interconnect 1110 provides communication channels to the other components, such as a Subscriber Identity Module (SIM) 1130 to interface with a SIM card, a boot ROM 1135 to hold boot code for execution by cores 1106 and 1107 to initialize and boot SoC 1100, a SDRAM controller 1140 to interface with external memory (e.g. DRAM 1160), a flash controller 1145 to interface with non-volatile memory (e.g. Flash 1165), a peripheral control 1150 (e.g. Serial Peripheral Interface) to interface with peripherals, video codecs 1120 and Video interface 1125 to display and receive input (e.g. touch enabled input), GPU 1115 to perform graphics related computations, etc. Any of these interfaces may incorporate aspects of the disclosure described herein. In addition, the system 1100 illustrates peripherals for

communication, such as a Bluetooth module 1170, 3G modem 1175, GPS 1180, and Wi-Fi 1185.

[0092] **Figure 12** illustrates a diagrammatic representation of a machine in the example form of a computer system 1200 within which a set of instructions, for causing the machine to perform any one or more of the methodologies discussed herein, may be executed. In alternative embodiments, the machine may be connected (e.g., networked) to other machines in a LAN, an intranet, an extranet, or the Internet. The machine may operate in the capacity of a server or a client device in a client-server network environment, or as a peer machine in a peer-to-peer (or distributed) network environment. The machine may be a personal computer (PC), a tablet PC, a set-top box (STB), a Personal Digital Assistant (PDA), a cellular telephone, a web appliance, a server, a network router, switch or bridge, or any machine capable of executing a set of instructions (sequential or otherwise) that specify actions to be taken by that machine. Further, while only a single machine is illustrated, the term “machine” shall also be taken to include any collection of machines that individually or jointly execute a set (or multiple sets) of instructions to perform any one or more of the methodologies discussed herein.

[0093] The computer system 1200 includes a processing device 1202, a main memory 1204 (e.g., read-only memory (ROM), flash memory, dynamic random access memory (DRAM) (such as synchronous DRAM (SDRAM) or DRAM (RDRAM), etc.), a static memory 1206 (e.g., flash memory, static random access memory (SRAM), etc.), and a data storage device 1218, which communicate with each other via a bus 1230.

[0094] Processing device 1202 represents one or more general-purpose processing devices such as a microprocessor, central processing unit, or the like. More particularly, the processing device may be complex instruction set computing (CISC) microprocessor, reduced instruction set computer (RISC) microprocessor, very long instruction word (VLIW) microprocessor, or processor implementing other instruction sets, or processors implementing a combination of instruction sets. Processing device 1202 may also be one or more special-purpose processing devices such as an application specific integrated circuit (ASIC), a field programmable gate array (FPGA), a digital signal processor (DSP), network processor, or the like. In one embodiment, processing device 1202 may include one or processing cores. The processing device 1202 is configured to execute the processing logic 1226 for performing the operations and steps discussed herein. In one embodiment, processing device 1202 is the same as processor architecture 100 described with respect to **Figure 1** as described herein with embodiments of the disclosure.

[0095] The computer system 1200 may further include a network interface device 1208 communicably coupled to a network 1220. The computer system 1200 also may include a video display unit 1210 (e.g., a liquid crystal display (LCD) or a cathode ray tube (CRT)), an alphanumeric input device 1212 (e.g., a keyboard), a cursor control device 1214 (e.g., a mouse), and a signal generation device 1216 (e.g., a speaker). Furthermore, computer system 1200 may include a graphics processing unit 1222, a video processing unit 1228, and an audio processing unit 1232.

[0096] The data storage device 1218 may include a machine-accessible storage medium 1224 on which is stored software 1226 implementing any one or more of the methodologies of functions described herein, such as implementing store address prediction for memory disambiguation as described above. The software 1226 may also reside, completely or at least partially, within the main memory 1204 as instructions 1226 and/or within the processing device 1202 as processing logic 1226 during execution thereof by the computer system 1200; the main memory 1204 and the processing device 1202 also constituting machine-accessible storage media.

[0097] The machine-readable storage medium 1224 may also be used to store instructions 1226 implementing store address prediction for hybrid cores such as described according to embodiments of the disclosure. While the machine-accessible storage medium 1128 is shown in an example embodiment to be a single medium, the term “machine-accessible storage medium” should be taken to include a single medium or multiple media (e.g., a centralized or distributed database, and/or associated caches and servers) that store the one or more sets of instructions. The term “machine-accessible storage medium” shall also be taken to include any medium that is capable of storing, encoding or carrying a set of instruction for execution by the machine and that cause the machine to perform any one or more of the methodologies of the present disclosure. The term “machine-accessible storage medium” shall accordingly be taken to include, but not be limited to, solid-state memories, and optical and magnetic media.

[0098] The following examples pertain to further embodiments. Example 1 is a processing system including a processor and a hardware accelerator, communicatively coupled to the processor, comprising a plurality of circuits to perform a plurality rounds of computation, wherein the plurality of circuits comprise a first set of level-sensitive latches enabled by a first clock signal to store data associated with a first round of the plurality of rounds of computation and a second set of level-sensitive latches enabled by a second clock signal to store data associated with a second round of the plurality of rounds of computation,

and wherein a duty cycle of the first clock signal and a duty cycle of the second clock signal are non-overlapping.

[0099] In Example 2, the subject matter of Example 1 further provides that the hardware accelerator comprises a clock generator to generate the first clock signal and the second clock signal based on a system clock, wherein the clock generator is further to generate a third clock signal based on the system clock, wherein a duty cycle of third clock signal is non-overlapping with the duty cycle of the second clock signal and non-overlapping with the duty cycle of the third clock signal.

[00100] In Example 3, the subject matter of any of Examples 1 and 2 further provides that the hardware accelerator is to provide the third clock signal to a third set of level-sensitive latches employed by a third round of the plurality of rounds of computation.

[00101] In Example 4, the subject matter of any of Examples 1 and 2 further provides that the system clock is associated with a first time period, wherein the first clock signal, the second clock signal, and the third clock signal are associated with a second time period, and wherein the second time period is at a pre-defined ratio of the first time period.

[00102] In Example 5, the subject matter of Example 4 further provides that the duty cycle of the system clock occupies a portion of the period of the system clock, and wherein the duty cycle of each one of the first clock signal, the second clock signal, or the third clock signal is at a second pre-defined ratio of the period of the system clock.

[00103] In Example 6, the subject matter of Example 4 further provides that the duty cycle of the first clock signal is separated from the duty cycle of the second clock signal by a time duration, the duty cycle of the first clock signal is separated from the duty cycle of the third clock signal by the time duration, and the duty cycle of the second clock signal is separated from the duty cycle of the third clock signal by the time duration.

[00104] In Example 7, the subject matter of any of Examples 1 and 2 further provides that an input of one of the second set of latches is communicatively coupled to an output of one of the first set of latches.

[00105] In Example 8, the subject matter of any of Examples 1 and 2 further provides that an input of one of the third set of latches is communicatively coupled to an output of one of the second set of latches.

[00106] In Example 9, the subject matter of Example 1 further provides that the plurality of rounds of computation is to implement one of stage-0 secure hash algorithm (SHA) hash, stage-1 SHA hash, or stage-2 SHA hash of Bitcoin mining.

[00107] Example 10 is an application specific integrated circuit (ASIC) comprising a plurality of circuits to perform a plurality rounds of computation, wherein the plurality of circuits comprise a first set of level-sensitive latches enabled by a first clock signal to store data associated with a first round of the plurality of rounds of computation and a second set of level-sensitive latches enabled by a second clock signal to store data associated with a second round of the plurality of rounds of computation, and wherein a duty cycle of the first clock signal and a duty cycle of the second clock signal are non-overlapping.

[00108] In Example 11, the subject matter of Example 10 further provides that a clock generator is to generate the first clock signal and the second clock signal, wherein the clock generator is further to generate a third clock signal based on the system clock, wherein a duty cycle of third clock signal is non-overlapping with the duty cycle of the second clock signal and non-overlapping with the duty cycle of the third clock signal.

[00109] In Example 12, the subject matter of any of Examples 10 and 11 further provides that the ASIC is to provide the third clock signal to a third set of level-sensitive latches employed by a third round of the plurality of rounds of computation.

[00110] In Example 13, the subject matter of any of Examples 10 and 11 further provides that the system clock is associated with a first time period, wherein the first clock signal, the second clock signal, and the third clock signal are associated with a second time period, and wherein the second time period is at a pre-defined ratio of the first time period.

[00111] In Example 14, the subject matter of Example 13 further provides that the duty cycle of the system clock occupies a portion of the period of the system clock, and wherein the duty cycle of each one of the first clock signal, the second clock signal, or the third clock signal is at a second pre-defined ratio of the period of the system clock.

[00112] In Example 15, the subject matter of Example 13 further provides that the duty cycle of the first clock signal is separated from the duty cycle of the second clock signal by a time duration, the duty cycle of the first clock signal is separated from the duty cycle of the third clock signal by the time duration, and the duty cycle of the second clock signal is separated from the duty cycle of the third clock signal by the time duration.

[00113] In Example 16, the subject matter of any of Examples 10 and 11 further provides that an input of one of the second set of latches is communicatively coupled to an output of one of the first set of latches.

[00114] In Example 17, the subject matter of any of Examples 10 and 11 further provides that an input of one of the third set of latches is communicatively coupled to an output of one of the second set of latches.

[00115] In Example 18, the subject matter of Example 10 further provides that the plurality of rounds of computation is to implement one of stage-0 secure hash algorithm (SHA) hash, stage-1 SHA hash, or stage-2 SHA hash of Bitcoin mining.

[00116] Example 19 is a method comprising performing a first round of a plurality of rounds of computation by employing a first set of level-sensitive latches enabled by a first clock signal, performing a second round of the plurality of rounds of computation by employing a second set of level-sensitive latches enabled by a second clock signal, and performing a third round of the plurality of rounds of computation by employing a third set of level-sensitive latches enabled by a third clock signal, wherein the first clock signal, the second clock signal, and the third clock signal are generated from a system clock.

[00117] In Example 20, the subject matter of Example 19 further provides that the system clock is associated with a first time period, wherein the first clock signal, the second clock signal, and the third clock signal are associated with a second time period, and wherein the second time period is at a pre-defined ratio of the first time period.

[00118] In Example 21, the subject matter of any of Examples 19 and 20 further provides that a duty cycle of the first clock signal and a duty cycle of the second clock signal are non-overlapping, and wherein a duty cycle of third clock signal is non-overlapping with the duty cycle of the second clock signal and non-overlapping with the duty cycle of the third clock signal.

[00119] Example 22 is an apparatus comprising: means for performing the method of any of Examples 19 to 21.

[00120] Example 23 is a machine-readable non-transitory medium having stored thereon program code that, when executed by a processor, perform operations comprising performing a first round of a plurality of rounds of computation by employing a first set of level-sensitive latches enabled by a first clock signal, performing a second round of the plurality of rounds of computation by employing a second set of level-sensitive latches enabled by a second clock signal, and performing a third round of the plurality of rounds of computation by employing a third set of level-sensitive latches enabled by a third clock signal, wherein the first clock signal, the second clock signal, and the third clock signal are generated from a system clock.

[00121] In Example 24, the subject matter of Example 23 further provides that the system clock is associated with a first time period, wherein the first clock signal, the second clock signal, and the third clock signal are associated with a second time period, and wherein the second time period is at a pre-defined ratio of the first time period.

[00122] In Example 35, the subject matter of any of Examples 23 and 24 further provides that a duty cycle of the first clock signal and a duty cycle of the second clock signal are non-overlapping, and wherein a duty cycle of third clock signal is non-overlapping with the duty cycle of the second clock signal and non-overlapping with the duty cycle of the third clock signal.

[00123] While the disclosure has been described with respect to a limited number of embodiments, those skilled in the art will appreciate numerous modifications and variations there from. It is intended that the appended claims cover all such modifications and variations as fall within the true spirit and scope of this disclosure.

[00124] A design may go through various stages, from creation to simulation to fabrication. Data representing a design may represent the design in a number of manners. First, as is useful in simulations, the hardware may be represented using a hardware description language or another functional description language. Additionally, a circuit level model with logic and/or transistor gates may be produced at some stages of the design process. Furthermore, most designs, at some stage, reach a level of data representing the physical placement of various devices in the hardware model. In the case where conventional semiconductor fabrication techniques are used, the data representing the hardware model may be the data specifying the presence or absence of various features on different mask layers for masks used to produce the integrated circuit. In any representation of the design, the data may be stored in any form of a machine readable medium. A memory or a magnetic or optical storage such as a disc may be the machine readable medium to store information transmitted via optical or electrical wave modulated or otherwise generated to transmit such information. When an electrical carrier wave indicating or carrying the code or design is transmitted, to the extent that copying, buffering, or re-transmission of the electrical signal is performed, a new copy is made. Thus, a communication provider or a network provider may store on a tangible, machine-readable medium, at least temporarily, an article, such as information encoded into a carrier wave, embodying techniques of embodiments of the present disclosure.

[00125] A module as used herein refers to any combination of hardware, software, and/or firmware. As an example, a module includes hardware, such as a micro-controller, associated with a non-transitory medium to store code adapted to be executed by the micro-controller. Therefore, reference to a module, in one embodiment, refers to the hardware, which is specifically configured to recognize and/or execute the code to be held on a non-transitory medium. Furthermore, in another embodiment, use of a module refers to the non-

transitory medium including the code, which is specifically adapted to be executed by the microcontroller to perform predetermined operations. And as can be inferred, in yet another embodiment, the term module (in this example) may refer to the combination of the microcontroller and the non-transitory medium. Often module boundaries that are illustrated as separate commonly vary and potentially overlap. For example, a first and a second module may share hardware, software, firmware, or a combination thereof, while potentially retaining some independent hardware, software, or firmware. In one embodiment, use of the term logic includes hardware, such as transistors, registers, or other hardware, such as programmable logic devices.

[00126] Use of the phrase ‘configured to,’ in one embodiment, refers to arranging, putting together, manufacturing, offering to sell, importing and/or designing an apparatus, hardware, logic, or element to perform a designated or determined task. In this example, an apparatus or element thereof that is not operating is still ‘configured to’ perform a designated task if it is designed, coupled, and/or interconnected to perform said designated task. As a purely illustrative example, a logic gate may provide a 0 or a 1 during operation. But a logic gate ‘configured to’ provide an enable signal to a clock does not include every potential logic gate that may provide a 1 or 0. Instead, the logic gate is one coupled in some manner that during operation the 1 or 0 output is to enable the clock. Note once again that use of the term ‘configured to’ does not require operation, but instead focus on the latent state of an apparatus, hardware, and/or element, where in the latent state the apparatus, hardware, and/or element is designed to perform a particular task when the apparatus, hardware, and/or element is operating.

[00127] Furthermore, use of the phrases ‘to,’ ‘capable of/to,’ and or ‘operable to,’ in one embodiment, refers to some apparatus, logic, hardware, and/or element designed in such a way to enable use of the apparatus, logic, hardware, and/or element in a specified manner. Note as above that use of to, capable to, or operable to, in one embodiment, refers to the latent state of an apparatus, logic, hardware, and/or element, where the apparatus, logic, hardware, and/or element is not operating but is designed in such a manner to enable use of an apparatus in a specified manner.

[00128] A value, as used herein, includes any known representation of a number, a state, a logical state, or a binary logical state. Often, the use of logic levels, logic values, or logical values is also referred to as 1’s and 0’s, which simply represents binary logic states. For example, a 1 refers to a high logic level and 0 refers to a low logic level. In one embodiment, a storage cell, such as a transistor or flash cell, may be capable of holding a

single logical value or multiple logical values. However, other representations of values in computer systems have been used. For example the decimal number ten may also be represented as a binary value of 910 and a hexadecimal letter A. Therefore, a value includes any representation of information capable of being held in a computer system.

[00129] Moreover, states may be represented by values or portions of values. As an example, a first value, such as a logical one, may represent a default or initial state, while a second value, such as a logical zero, may represent a non-default state. In addition, the terms reset and set, in one embodiment, refer to a default and an updated value or state, respectively. For example, a default value potentially includes a high logical value, i.e. reset, while an updated value potentially includes a low logical value, i.e. set. Note that any combination of values may be utilized to represent any number of states.

[00130] The embodiments of methods, hardware, software, firmware or code set forth above may be implemented via instructions or code stored on a machine-accessible, machine readable, computer accessible, or computer readable medium which are executable by a processing element. A non-transitory machine-accessible/readable medium includes any mechanism that provides (i.e., stores and/or transmits) information in a form readable by a machine, such as a computer or electronic system. For example, a non-transitory machine-accessible medium includes random-access memory (RAM), such as static RAM (SRAM) or dynamic RAM (DRAM); ROM; magnetic or optical storage medium; flash memory devices; electrical storage devices; optical storage devices; acoustical storage devices; other form of storage devices for holding information received from transitory (propagated) signals (e.g., carrier waves, infrared signals, digital signals); etc., which are to be distinguished from the non-transitory mediums that may receive information there from.

[00131] Instructions used to program logic to perform embodiments of the disclosure may be stored within a memory in the system, such as DRAM, cache, flash memory, or other storage. Furthermore, the instructions can be distributed via a network or by way of other computer readable media. Thus a machine-readable medium may include any mechanism for storing or transmitting information in a form readable by a machine (e.g., a computer), but is not limited to, floppy diskettes, optical disks, Compact Disc, Read-Only Memory (CD-ROMs), and magneto-optical disks, Read-Only Memory (ROMs), Random Access Memory (RAM), Erasable Programmable Read-Only Memory (EPROM), Electrically Erasable Programmable Read-Only Memory (EEPROM), magnetic or optical cards, flash memory, or a tangible, machine-readable storage used in the transmission of information over the Internet via electrical, optical, acoustical or other forms of propagated signals (e.g., carrier waves,

infrared signals, digital signals, etc.). Accordingly, the computer-readable medium includes any type of tangible machine-readable medium suitable for storing or transmitting electronic instructions or information in a form readable by a machine (e.g., a computer).

[00132] Reference throughout this specification to “one embodiment” or “an embodiment” means that a particular feature, structure, or characteristic described in connection with the embodiment is included in at least one embodiment of the present disclosure. Thus, the appearances of the phrases “in one embodiment” or “in an embodiment” in various places throughout this specification are not necessarily all referring to the same embodiment. Furthermore, the particular features, structures, or characteristics may be combined in any suitable manner in one or more embodiments.

[00133] In the foregoing specification, a detailed description has been given with reference to specific exemplary embodiments. It will, however, be evident that various modifications and changes may be made thereto without departing from the broader spirit and scope of the disclosure as set forth in the appended claims. The specification and drawings are, accordingly, to be regarded in an illustrative sense rather than a restrictive sense. Furthermore, the foregoing use of embodiment and other exemplarily language does not necessarily refer to the same embodiment or the same example, but may refer to different and distinct embodiments, as well as potentially the same embodiment.

CLAIMS

What is claimed is:

1. A processing system comprising:
a processor; and
a hardware accelerator, communicatively coupled to the processor, comprising a plurality of circuits to perform a plurality rounds of computation, wherein the plurality of circuits comprise a first set of level-sensitive latches enabled by a first clock signal to store data associated with a first round of the plurality of rounds of computation and a second set of level-sensitive latches enabled by a second clock signal to store data associated with a second round of the plurality of rounds of computation, and wherein a duty cycle of the first clock signal and a duty cycle of the second clock signal are non-overlapping.
2. The processing system of claim 1, wherein the hardware accelerator comprises:
a clock generator to generate the first clock signal and the second clock signal based on a system clock, wherein the clock generator is further to generate a third clock signal based on the system clock, wherein a duty cycle of third clock signal is non-overlapping with the duty cycle of the second clock signal and non-overlapping with the duty cycle of the third clock signal.
3. The processing system of any of claims 1 and 2, wherein the hardware accelerator is to provide the third clock signal to a third set of level-sensitive latches employed by a third round of the plurality of rounds of computation.
4. The processing system of any of claims 1 and 2, wherein the system clock is associated with a first time period, wherein the first clock signal, the second clock signal, and the third clock signal are associated with a second time period, and wherein the second time period is at a pre-defined ratio of the first time period.
5. The processing system of claim 4, wherein the duty cycle of the system clock occupies a portion of the period of the system clock, and wherein the duty cycle of each one of the first clock signal, the second clock signal, or the third clock signal is at a second pre-defined ratio of the period of the system clock.
6. The processing system of claim 4, wherein the duty cycle of the first clock signal is separated from the duty cycle of the second clock signal by a time duration, the duty cycle of

the first clock signal is separated from the duty cycle of the third clock signal by the time duration, and the duty cycle of the second clock signal is separated from the duty cycle of the third clock signal by the time duration.

7. The processing system of any of claims 1 and 2, wherein an input of one of the second set of latches is communicatively coupled to an output of one of the first set of latches.

8. The processing system of any of claims 1 and 2, wherein an input of one of the third set of latches is communicatively coupled to an output of one of the second set of latches.

9. The processing system of claim 1, wherein the plurality of rounds of computation is to implement one of stage-0 secure hash algorithm (SHA) hash, stage-1 SHA hash, or stage-2 SHA hash of Bitcoin mining.

10. An application specific integrated circuit (ASIC) comprising:
a plurality of circuits to perform a plurality rounds of computation, wherein the plurality of circuits comprise a first set of level-sensitive latches enabled by a first clock signal to store data associated with a first round of the plurality of rounds of computation and a second set of level-sensitive latches enabled by a second clock signal to store data associated with a second round of the plurality of rounds of computation, and wherein a duty cycle of the first clock signal and a duty cycle of the second clock signal are non-overlapping.


11. The ASIC of claim 10, wherein a clock generator is to generate the first clock signal and the second clock signal, wherein the clock generator is further to generate a third clock signal based on the system clock, wherein a duty cycle of third clock signal is non-overlapping with the duty cycle of the second clock signal and non-overlapping with the duty cycle of the third clock signal.

12. The ASIC of any of claims 10 and 11, wherein the ASIC is to provide the third clock signal to a third set of level-sensitive latches employed by a third round of the plurality of rounds of computation.

13. The ASIC of any of claims 10 and 11, wherein the system clock is associated with a first time period, wherein the first clock signal, the second clock signal, and the third clock signal are associated with a second time period, and wherein the second time period is at a pre-defined ratio of the first time period.

14. The ASIC of claim 13, wherein the duty cycle of the system clock occupies a portion of the period of the system clock, and wherein the duty cycle of each one of the first clock signal, the second clock signal, or the third clock signal is at a second pre-defined ratio of the period of the system clock.
15. The ASIC of claim 13, wherein the duty cycle of the first clock signal is separated from the duty cycle of the second clock signal by a time duration, the duty cycle of the first clock signal is separated from the duty cycle of the third clock signal by the time duration, and the duty cycle of the second clock signal is separated from the duty cycle of the third clock signal by the time duration.
16. The ASIC of any of claims 10 and 11, wherein an input of one of the second set of latches is communicatively coupled to an output of one of the first set of latches.
17. The ASIC of any of claims 10 and 11, wherein an input of one of the third set of latches is communicatively coupled to an output of one of the second set of latches.
18. The ASIC of claim 10, wherein the plurality of rounds of computation is to implement one of stage-0 secure hash algorithm (SHA) hash, stage-1 SHA hash, or stage-2 SHA hash of Bitcoin mining.
19. A method comprising:
performing a first round of a plurality of rounds of computation by employing a first set of level-sensitive latches enabled by a first clock signal;
performing a second round of the plurality of rounds of computation by employing a second set of level-sensitive latches enabled by a second clock signal; and
performing a third round of the plurality of rounds of computation by employing a third set of level-sensitive latches enabled by a third clock signal, wherein the first clock signal, the second clock signal, and the third clock signal are generated from a system clock.
20. The method of claim 19, wherein the system clock is associated with a first time period, wherein the first clock signal, the second clock signal, and the third clock signal are associated with a second time period, and wherein the second time period is at a pre-defined ratio of the first time period.

21. The method of any of claims 19 and 20, wherein a duty cycle of the first clock signal and a duty cycle of the second clock signal are non-overlapping, and wherein a duty cycle of third clock signal is non-overlapping with the duty cycle of the second clock signal and non-overlapping with the duty cycle of the third clock signal.
22. An apparatus comprising: means for performing the method of any of claims 19 to 21.
23. A machine-readable non-transitory medium having stored thereon program code that, when executed by a processor, perform operations comprising:
- performing a first round of a plurality of rounds of computation by employing a first set of level-sensitive latches enabled by a first clock signal;
 - performing a second round of the plurality of rounds of computation by employing a second set of level-sensitive latches enabled by a second clock signal; and
 - performing a third round of the plurality of rounds of computation by employing a third set of level-sensitive latches enabled by a third clock signal, wherein the first clock signal, the second clock signal, and the third clock signal are generated from a system clock.
24. The machine-readable non-transitory medium of claim 23, wherein the system clock is associated with a first time period, wherein the first clock signal, the second clock signal, and the third clock signal are associated with a second time period, and wherein the second time period is at a pre-defined ratio of the first time period.
25. The machine-readable non-transitory medium of any of claims 23 and 24, wherein a duty cycle of the first clock signal and a duty cycle of the second clock signal are non-overlapping, and wherein a duty cycle of third clock signal is non-overlapping with the duty cycle of the second clock signal and non-overlapping with the duty cycle of the third clock signal.



100

Figure 1




Figure 2

3/14




Figure 3 A




Figure 3B





Figure 3C

6/14

400

*Figure 4*




FIG. 5A




FIG. 5B




FIG. 6





FIG. 7

10/14

**FIG. 8**




FIG. 9




FIG. 10




FIG. 11




FIG. 12

INTERNATIONAL SEARCH REPORT

International application No.
PCT/US2017/035191**A. CLASSIFICATION OF SUBJECT MATTER****G06F 9/30(2006.01)i**

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

B. FIELDS SEARCHEDMinimum documentation searched (classification system followed by classification symbols)
G06F 9/30; H03L 7/00; H04K 1/00; G05B 19/045; H04L 9/28; H03K 3/356; H03L 7/06; H04L 9/06Documentation 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 modelsElectronic 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), rounds of computation, level-sensitive latches,
duty cycle, non-overlapping, Bitcoin mining, clock signal**C. DOCUMENTS CONSIDERED TO BE RELEVANT**

Category*	Citation of document, with indication, where appropriate, of the relevant passages	Relevant to claim No.
Y	US 2016-0109870 A1 (21, INC.) 21 April 2016 See paragraphs [0011]–[0014], [0032], [0048]; claims 1-7; and figures 6-7.	1-25
Y	US 2009-0058478 A1 (FENG LIN) 05 March 2009 See paragraphs [0026], [0032]; and figures 1, 4.	1-25
A	WO 2016-046821 A1 (SPONDOOLIES TECH LTD.) 31 March 2016 See paragraphs [0032]–[0044]; and figures 1-6.	1-25
A	US 2010-0086127 A1 (MIKHAIL GRINCHUK et al.) 08 April 2010 See paragraphs [0059]–[0073]; and figures 1-4.	1-25
A	US 2008-0107259 A1 (HIROSHI SATOU) 08 May 2008 See paragraphs [0024]–[0059]; and figures 1A-5.	1-25



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 another citation or other special reason (as specified)

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

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

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

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

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

"&" document member of the same patent family

Date of the actual completion of the international search

08 September 2017 (08.09.2017)

Date of mailing of the international search report

08 September 2017 (08.09.2017)

Name and mailing address of the ISA/KR

International Application Division
Korean Intellectual Property Office
189 Cheongsu-ro, Seo-gu, Daejeon, 35208, Republic of Korea

Facsimile No. +82-42-481-8578

Authorized officer

Lee, Eun Kyu

Telephone No. +82-42-481-3580



INTERNATIONAL SEARCH REPORT

Information on patent family members

International application No.

PCT/US2017/035191

Patent document cited in search report	Publication date	Patent family member(s)	Publication date
US 2016-0109870 A1	21/04/2016	TW 201626717 A WO 2016-060930 A1	16/07/2016 21/04/2016
US 2009-0058478 A1	05/03/2009	US 2007-0033464 A1 US 7454646 B2 US 8001410 B2	08/02/2007 18/11/2008 16/08/2011
WO 2016-046821 A1	31/03/2016	EP 3198783 A1	02/08/2017
US 2010-0086127 A1	08/04/2010	US 8160242 B2	17/04/2012
US 2008-0107259 A1	08/05/2008	CN 101178861 A CN 101178861 B JP 04308840 B2 JP 2008-116856 A KR 10-1389034 B1 KR 10-2008-0041990 A US 7962753 B2	14/05/2008 25/05/2011 05/08/2009 22/05/2008 28/04/2014 14/05/2008 14/06/2011