



US011770258B2

(12) **United States Patent**
Suresh et al.

(10) **Patent No.:** **US 11,770,258 B2**
(45) **Date of Patent:** ***Sep. 26, 2023**

(54) **ACCELERATORS FOR POST-QUANTUM CRYPTOGRAPHY SECURE HASH-BASED SIGNING AND VERIFICATION**

(58) **Field of Classification Search**
CPC ... H04L 9/3239; H04L 9/0869; H04L 9/3247; H04L 9/50
See application file for complete search history.

(71) Applicant: **Intel Corporation**, Santa Clara, CA (US)

(56) **References Cited**

(72) Inventors: **Vikram Suresh**, Portland, OR (US);
Sanu Mathew, Portland, OR (US);
Manoj Sastry, Portland, OR (US);
Santosh Ghosh, Hillsboro, OR (US);
Raghavan Kumar, Hillsboro, OR (US);
Rafael Misoczki, Hillsboro, OR (US)

U.S. PATENT DOCUMENTS

6,944,313 B1 9/2005 Donescu
2004/0228360 A1 11/2004 Bae

(Continued)

FOREIGN PATENT DOCUMENTS

CN 112152786 A1 12/2020
CN 112152788 A 12/2020

(Continued)

OTHER PUBLICATIONS

Notice of Allowance; U.S. Appl. No. 16/455,967, dated Oct. 28, 2022, 11 pages.

(Continued)

(21) Appl. No.: **17/562,461**

(22) Filed: **Dec. 27, 2021**

Primary Examiner — Ghodrat Jamshidi

(65) **Prior Publication Data**

US 2022/0123943 A1 Apr. 21, 2022

(74) *Attorney, Agent, or Firm* — JAFFERY WATSON MENDONSA & HAMILTON LLP

Related U.S. Application Data

(63) Continuation of application No. 16/455,908, filed on Jun. 28, 2019, now Pat. No. 11,218,320.

(51) **Int. Cl.**
H04L 9/40 (2022.01)
H04L 9/32 (2006.01)

(Continued)


(57) **ABSTRACT**

In one example an apparatus comprises a computer readable memory, hash logic to generate a message hash value based on an input message, signature logic to generate a signature to be transmitted in association with the message, the signature logic to apply a hash-based signature scheme to a private key to generate the signature comprising a public key, and accelerator logic to pre-compute at least one set of inputs to the signature logic. Other examples may be described.


(52) **U.S. Cl.**
CPC **H04L 9/3239** (2013.01); **H04L 9/0869** (2013.01); **H04L 9/3247** (2013.01); **H04L 9/50** (2022.05)

20 Claims, 18 Drawing Sheets


One-Time Hash-Based Signatures



1 signing key



1 signature



1 verification key

A private key must only sign a single message

- (51) **Int. Cl.**
H04L 9/08 (2006.01)
H04L 9/00 (2022.01)

(56) **References Cited**

U.S. PATENT DOCUMENTS

2005/0198468	A1*	9/2005	Rubin	G06F 8/443 712/22
2007/0104329	A1*	5/2007	England	G06F 21/57 380/255
2007/0130470	A1	6/2007	Blom		
2007/0143826	A1	6/2007	Sastry		
2010/0235588	A1*	9/2010	Maeda	H04L 9/0891 711/147
2011/0083015	A1*	4/2011	Meier	H04L 9/3236 713/176
2012/0008775	A1	1/2012	Natarajan		
2012/0246625	A1*	9/2012	Puthuff	G06F 11/28 717/130
2013/0322621	A1*	12/2013	Yoon	H04L 9/0816 380/44
2014/0019759	A1	1/2014	Burmester		
2014/0037089	A1	2/2014	Itoh		
2014/0270162	A1*	9/2014	Greiner	G06F 9/3001 380/46
2015/0100785	A1	4/2015	Joye		
2017/0331631	A1	11/2017	Peeters		
2018/0091309	A1*	3/2018	Misoczki	H04L 9/3252
2020/0014535	A1	1/2020	Baskaran		
2020/0177389	A1	6/2020	Tsuchida		
2021/0167945	A1	6/2021	Yamakawa		
2021/0194705	A1	6/2021	Takazoe		

FOREIGN PATENT DOCUMENTS

EP		3758278	A1	12/2020
EP		3758281	A1	12/2020

OTHER PUBLICATIONS

Communication Pursuant to Article 94(3) in EP Application No. 20162601.7, dated Feb. 3, 2023, 10 pages.

Glabb et al: “Multi-mode operator for SHA-2 hash functions”, Journal of Systems Architecture, Elsevier BV, NL, vol. 53, No. 2-3, Jan. 20, 2007 (Jan. 20, 2007), pp. 127-138, XP005855099, ISSN: 1383-7621, DOI: 10.1016/J.SYSARC.2006.09.006.

Communication—Extended European Search Report for EP Application No. 20162601.7, dated Sep. 16, 2020. 9 pages.

Wen Wang et al: “XMSS and Embedded Systems—XMSS Hardware Accelerators for RISC-V”, IACR, International Association for Cryptologic Research, vol. 20181230:125359, Dec. 21, 2018 (Dec. 21, 2018), pp. 1-23, XP061027415.

Aysu Aydin et al.: “Precomputation Methods for Hash-Based Signatures on Energy-Harvesting Platforms”, IEEE Transactions on Computers, IEEE, USA, vol. 65, No. 9, Sep. 1, 2016 (Sep. 1, 2016), pp. 2925-2931, XP011618562.

Santosh Ghosh et al: “Lightweight Post-Quantum-Secure Digital Signature Approach for IoT Motes”, IACR, International Association for Cryptologic Research, vol. 20190213:033538, Feb. 7, 2019 (Feb. 7, 2019), pp. 1-23, XP061031800.

Yavux Attila Altay, et al., “Real-Time Digital Signatures for Time-Critical Networks”, IEEE Transactions on Information Forensics and Security, IEEE, Piscataway NJ, US, vol. 12, No. 11, Nov. 1, 2017, pp. 2627-2639, XP011660596.

Communication—Extended European Search Report for EP Application 20161899.8, dated Aug. 20, 2020, 8 pages.

Huelsing et al. “XMSS: extended Merkle Signature Scheme” [online] Internet Research Task Force, May 2018 [retrieved Mar. 26, 2021]. Retrieved from the Internet: URL: <https://tools.ietf.org/pdf/rfc8391.pdf> (Year: 2018).

Perin et al. “Tuning the Winternitz hash-based digital signature scheme” [online] IEEE, Jun. 2018 [retrieved Mar. 26, 2021]. Retrieved from the Internet: URL: <https://ieeexplore.ieee.org/document/8538642> (Year: 2018).

Non-Final Office Action issued in U.S. Appl. No. 16/455,967 dated Apr. 2, 2021, 24 pages.

Daniel J. Bernstein & Tanja Lang, “Post-quantum cryptography”, Sep. 14, 2017, Nature, vol. 549, pp. 1-7 (Year: 2017).

Final Office Action issued in U.S. Appl. No. 16/455,967 dated Oct. 1, 2021, 27 pages.

Communication—Extended European Search Report for EP Application No. 20161899.8, dated Mar. 16, 2021, 7 pages.

* cited by examiner

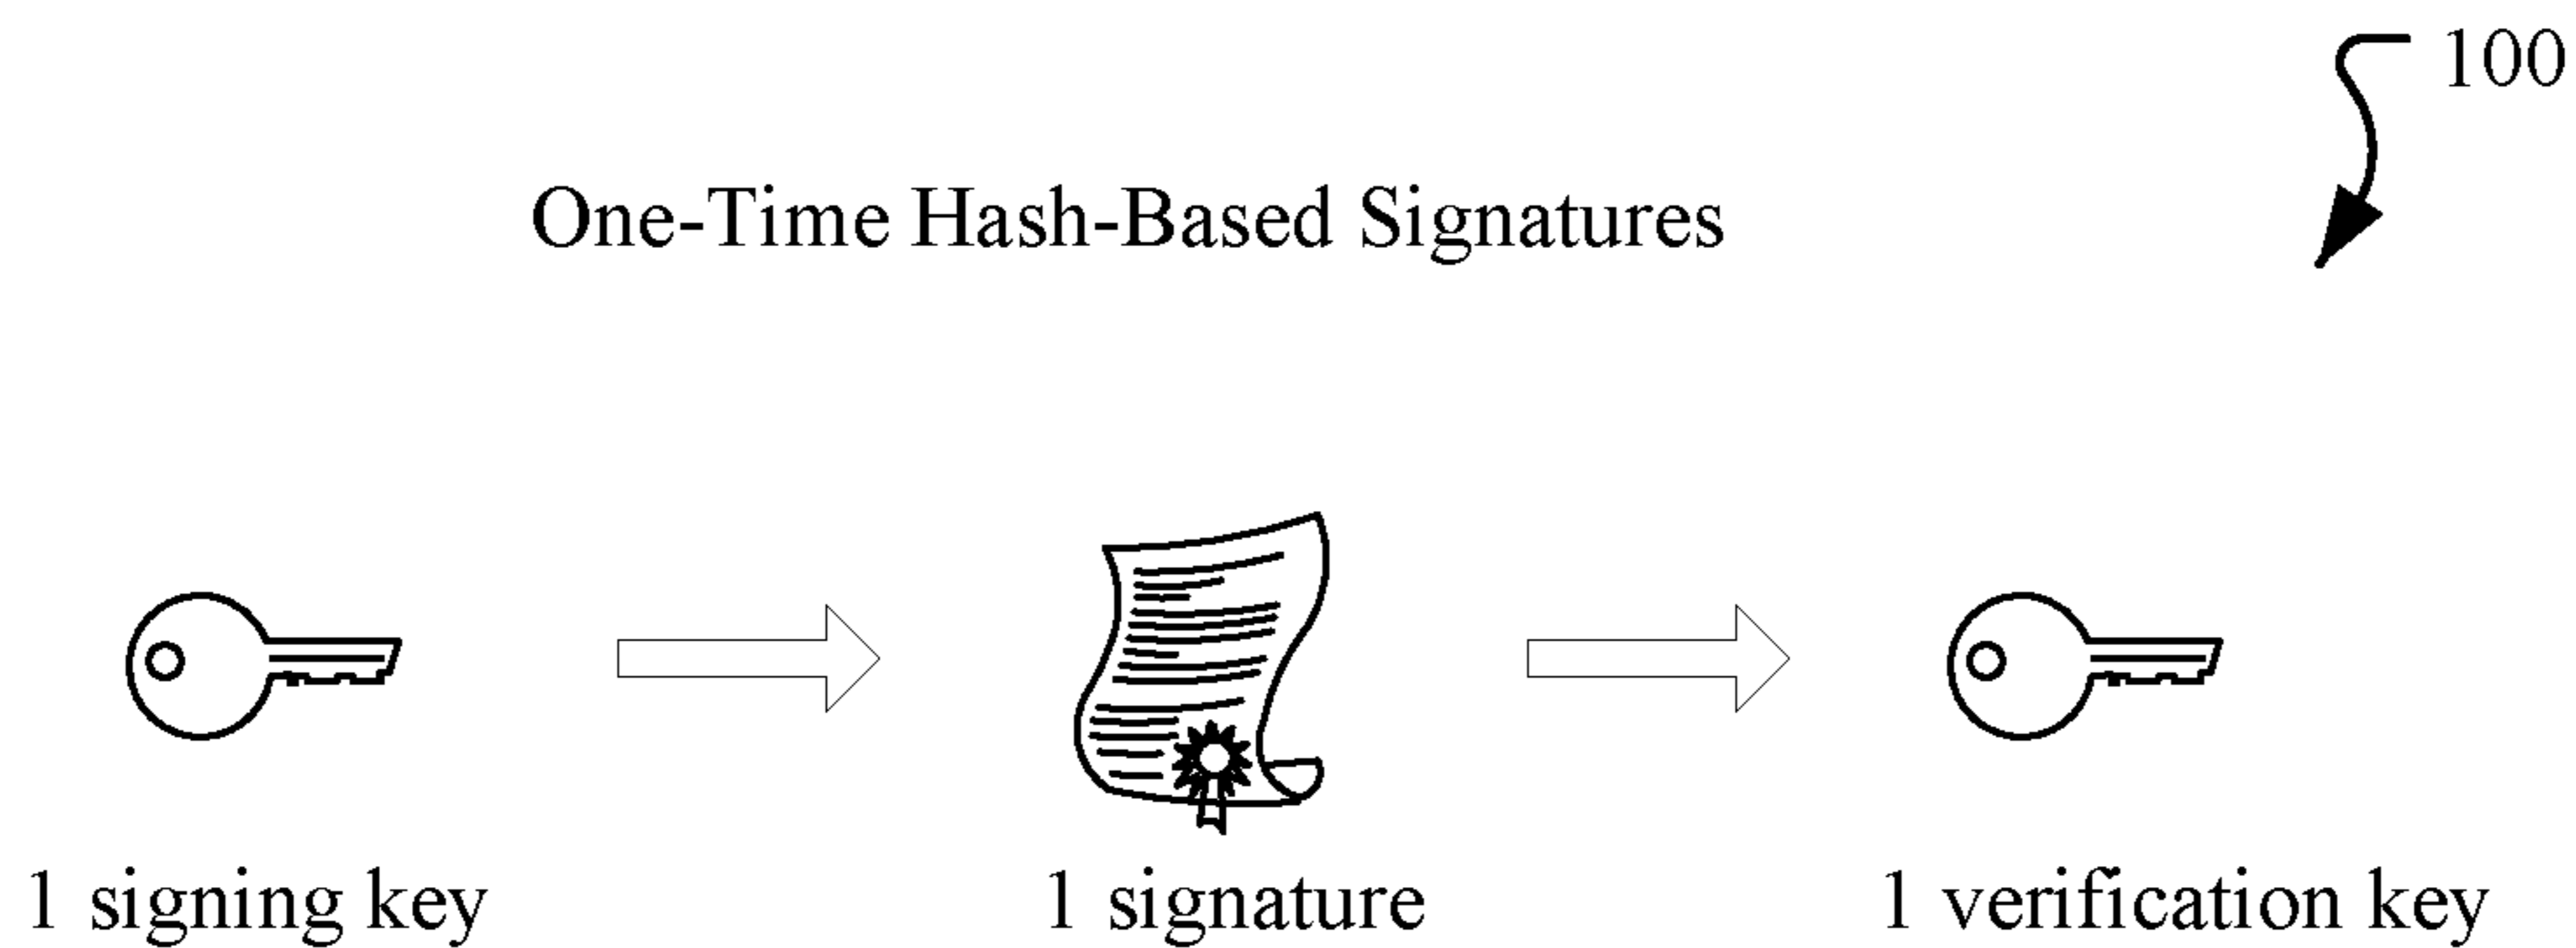


FIG. 1A




FIG. 1B

200

WOTS One-Time Signature Scheme

Message is hashed and parsed into M:67 integers between [0,1,2,...15]

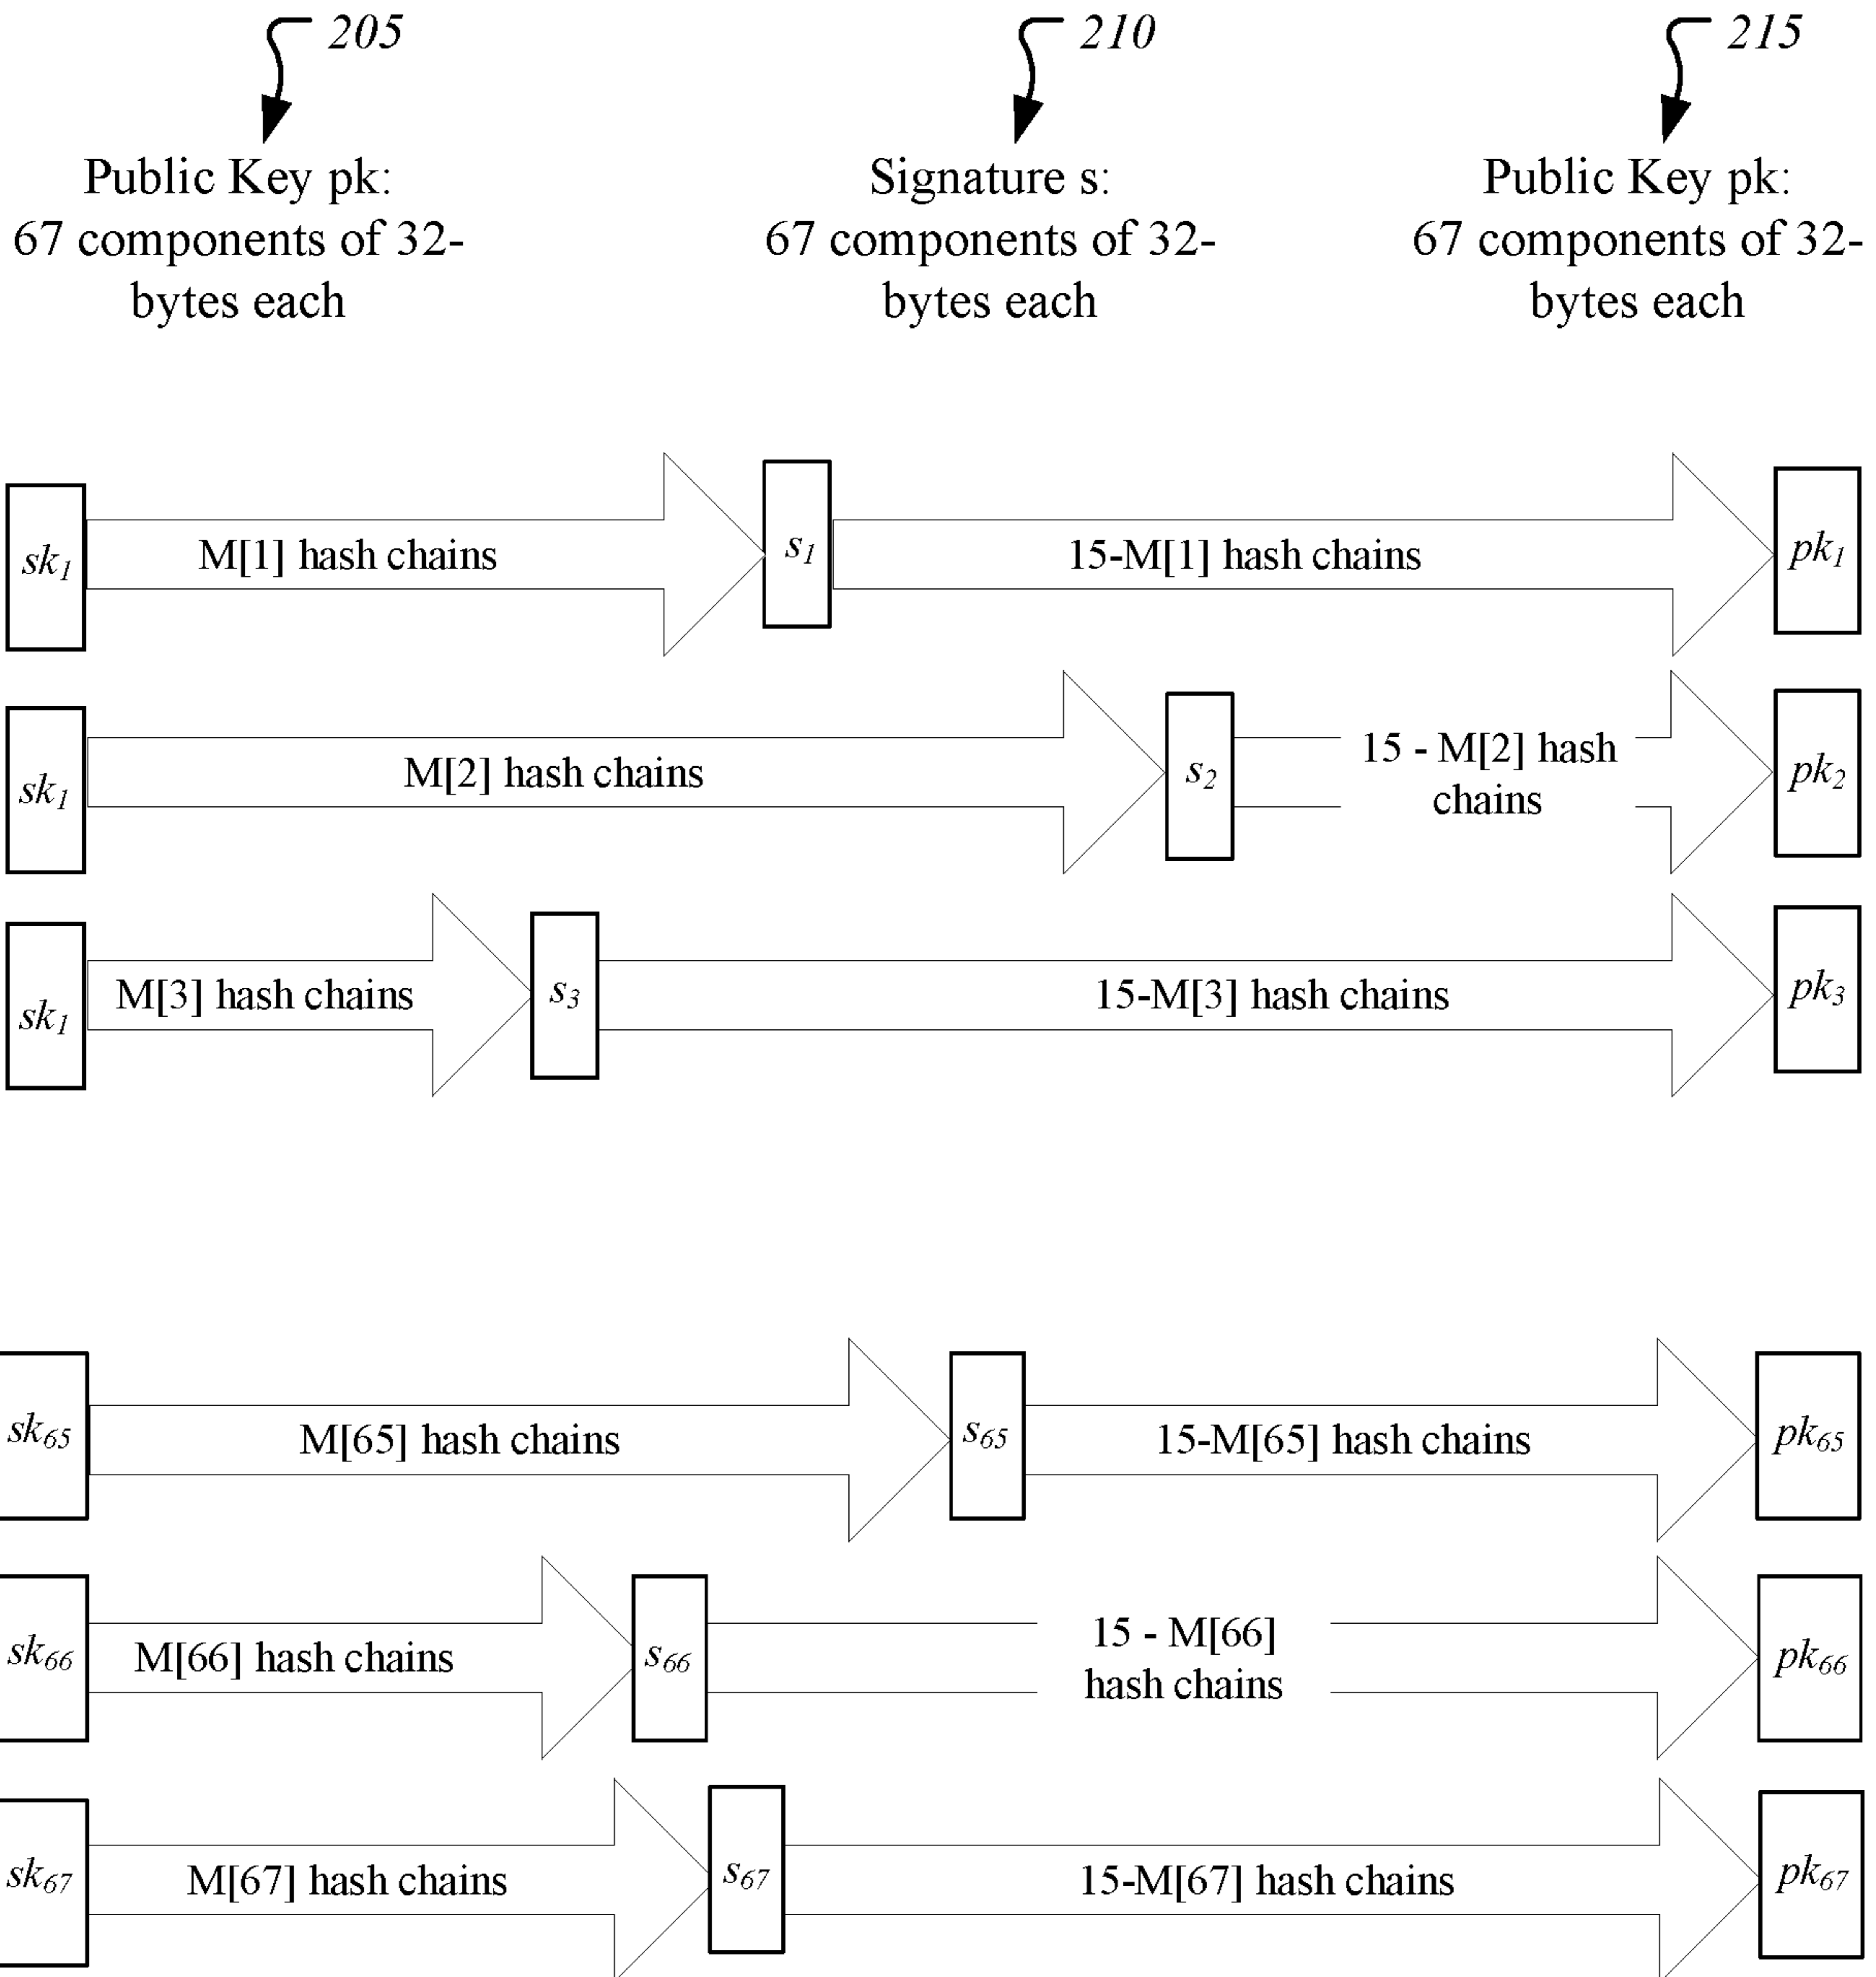


FIG. 2A

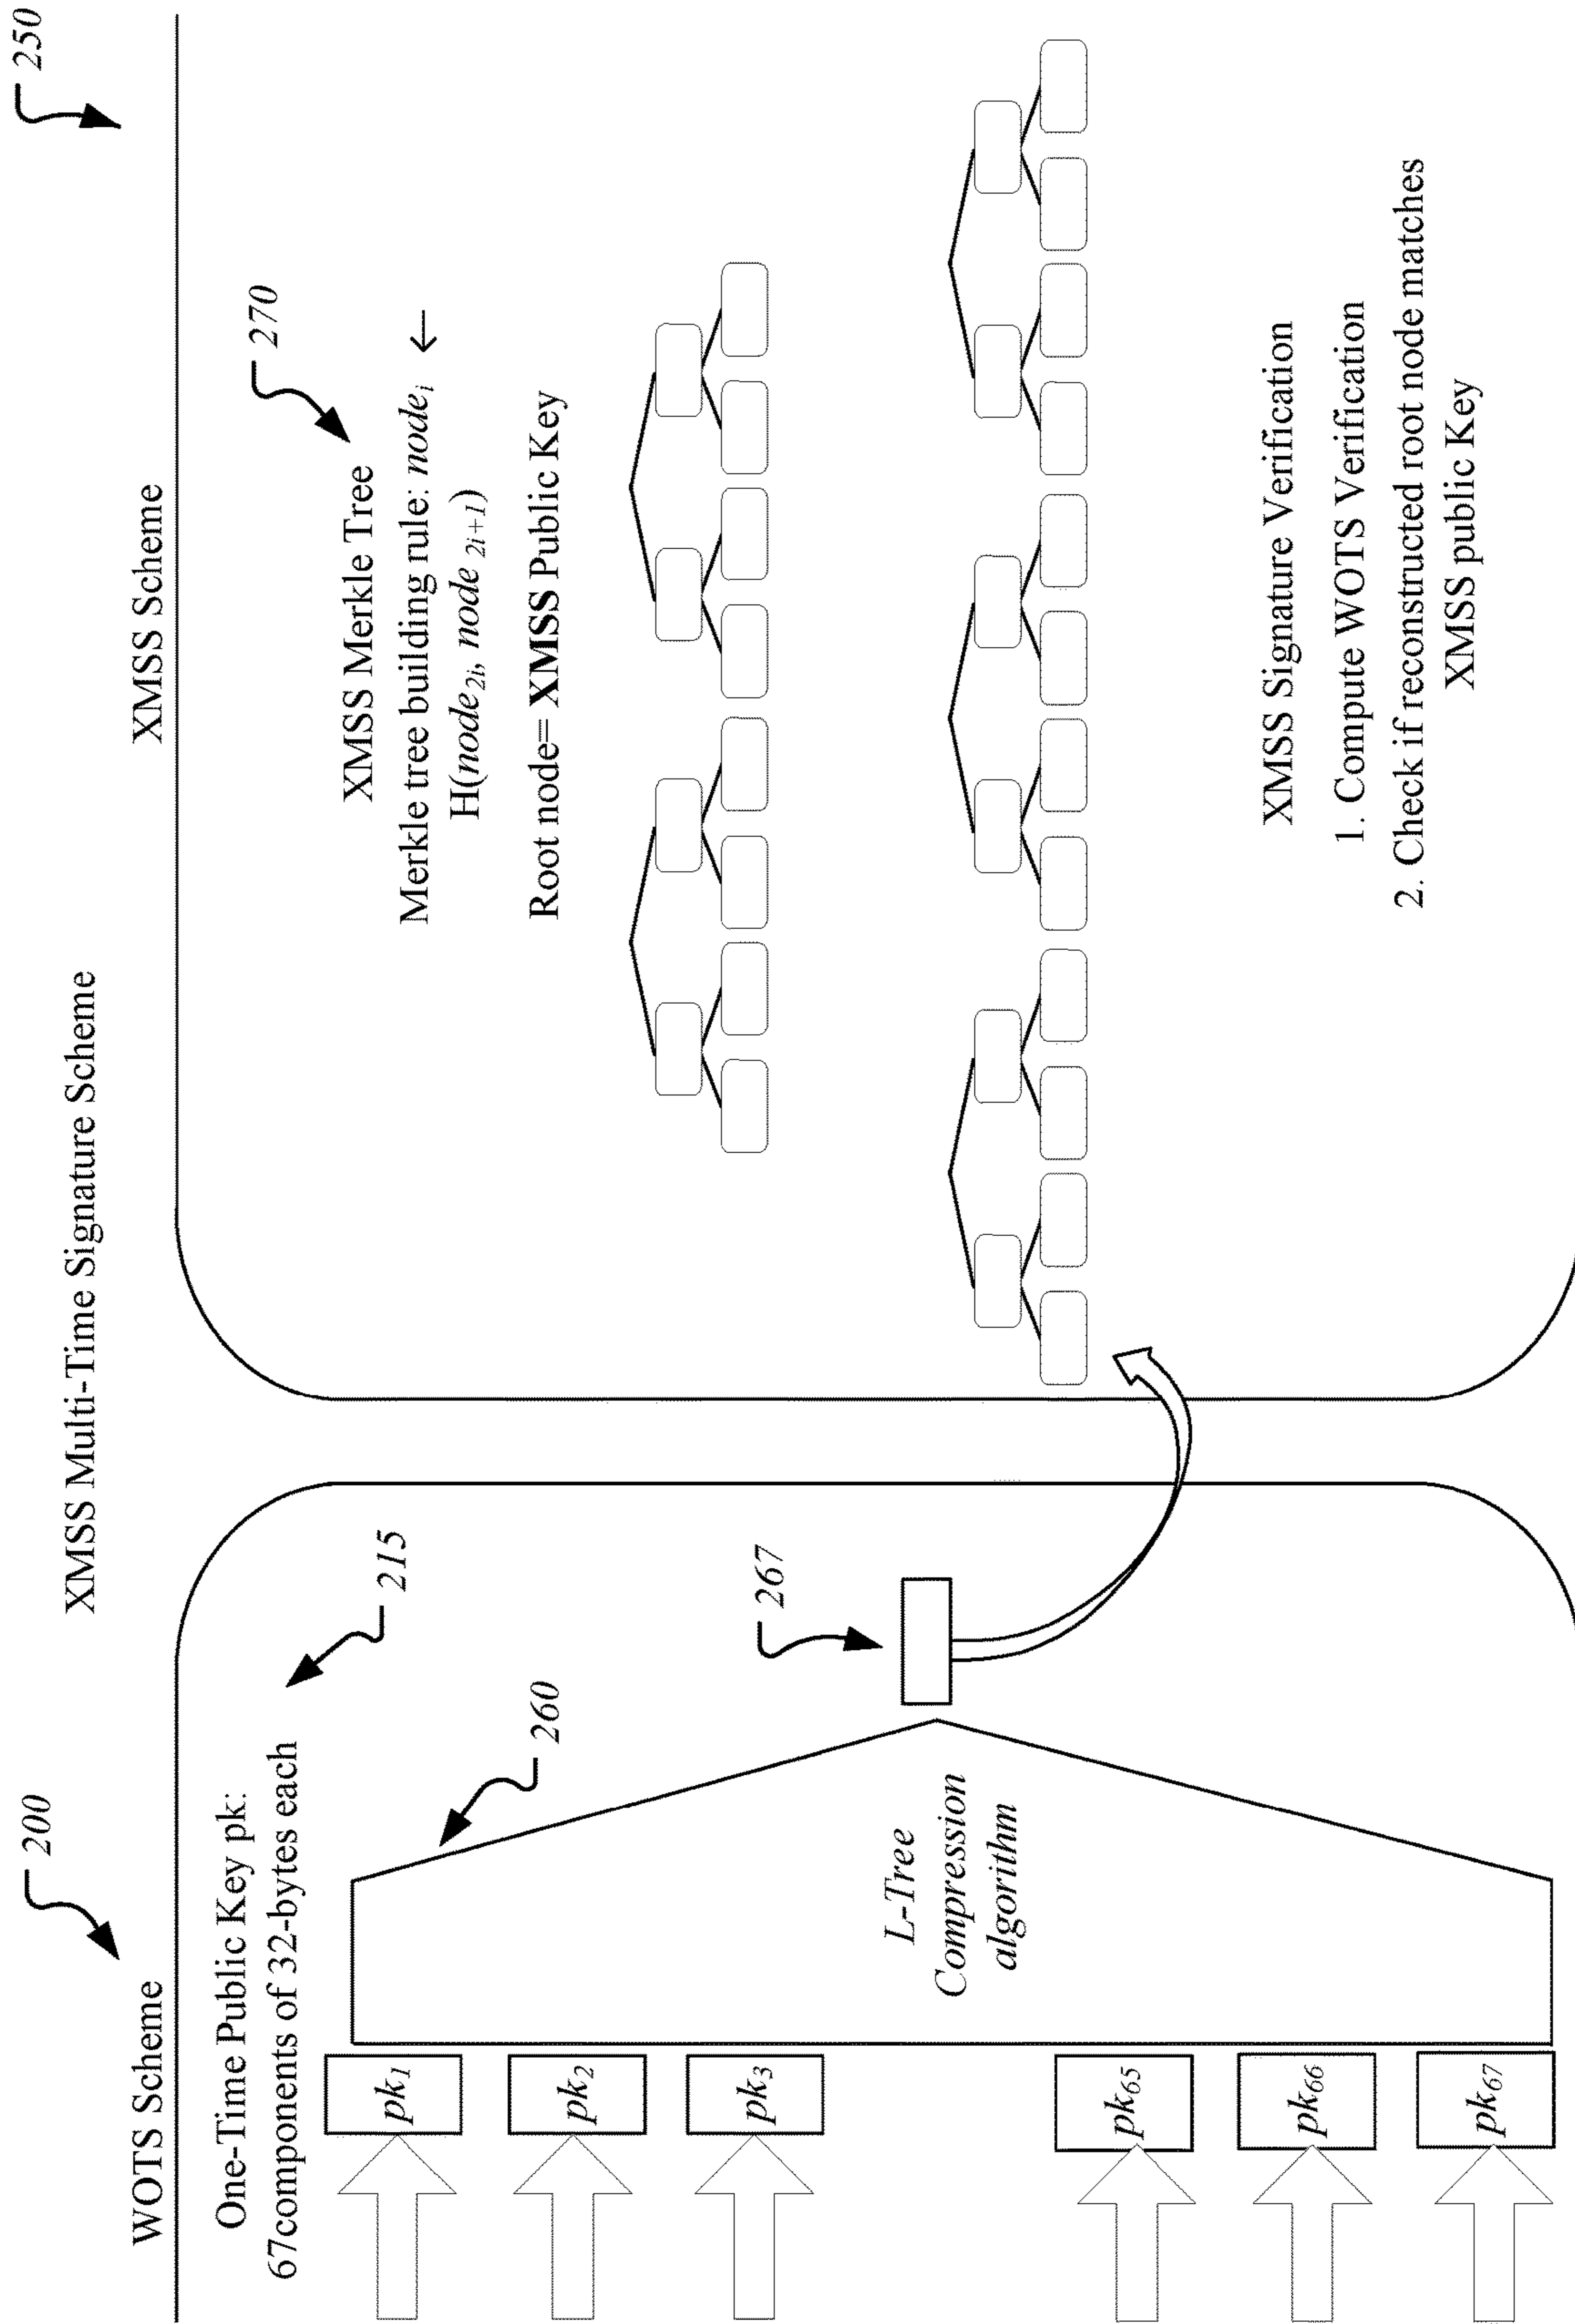


FIG. 2B

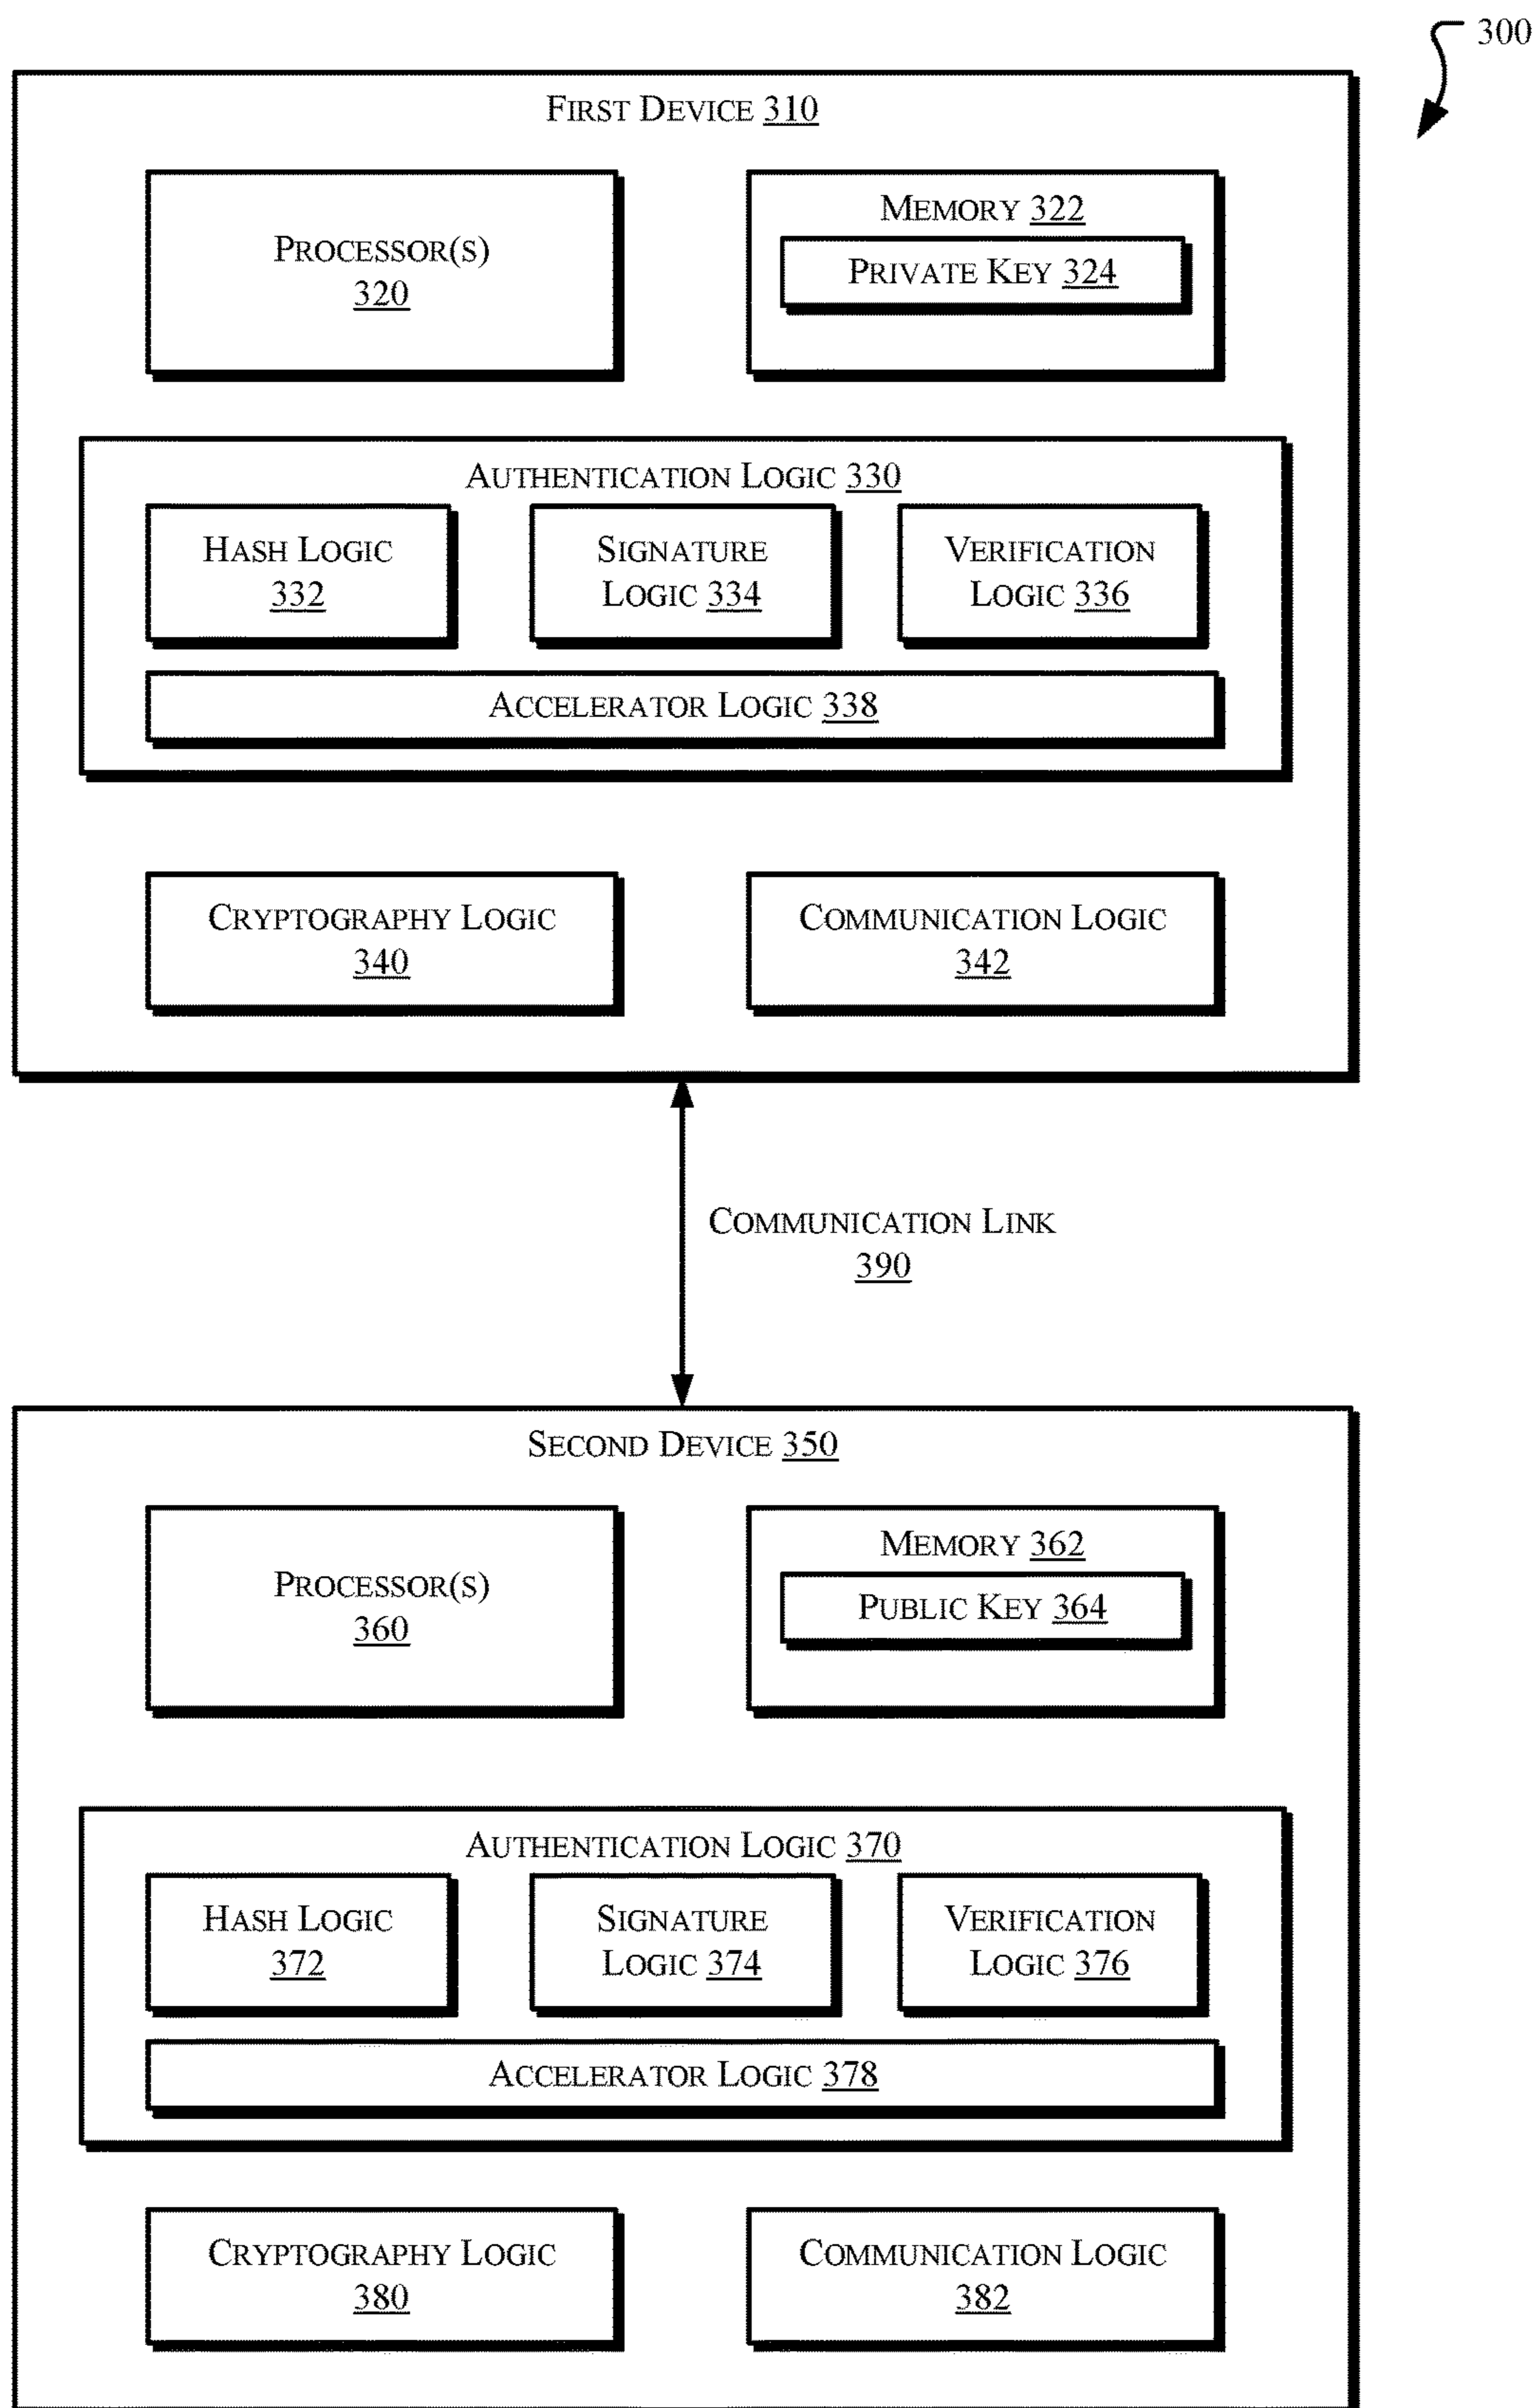


FIG. 3




FIG. 4A

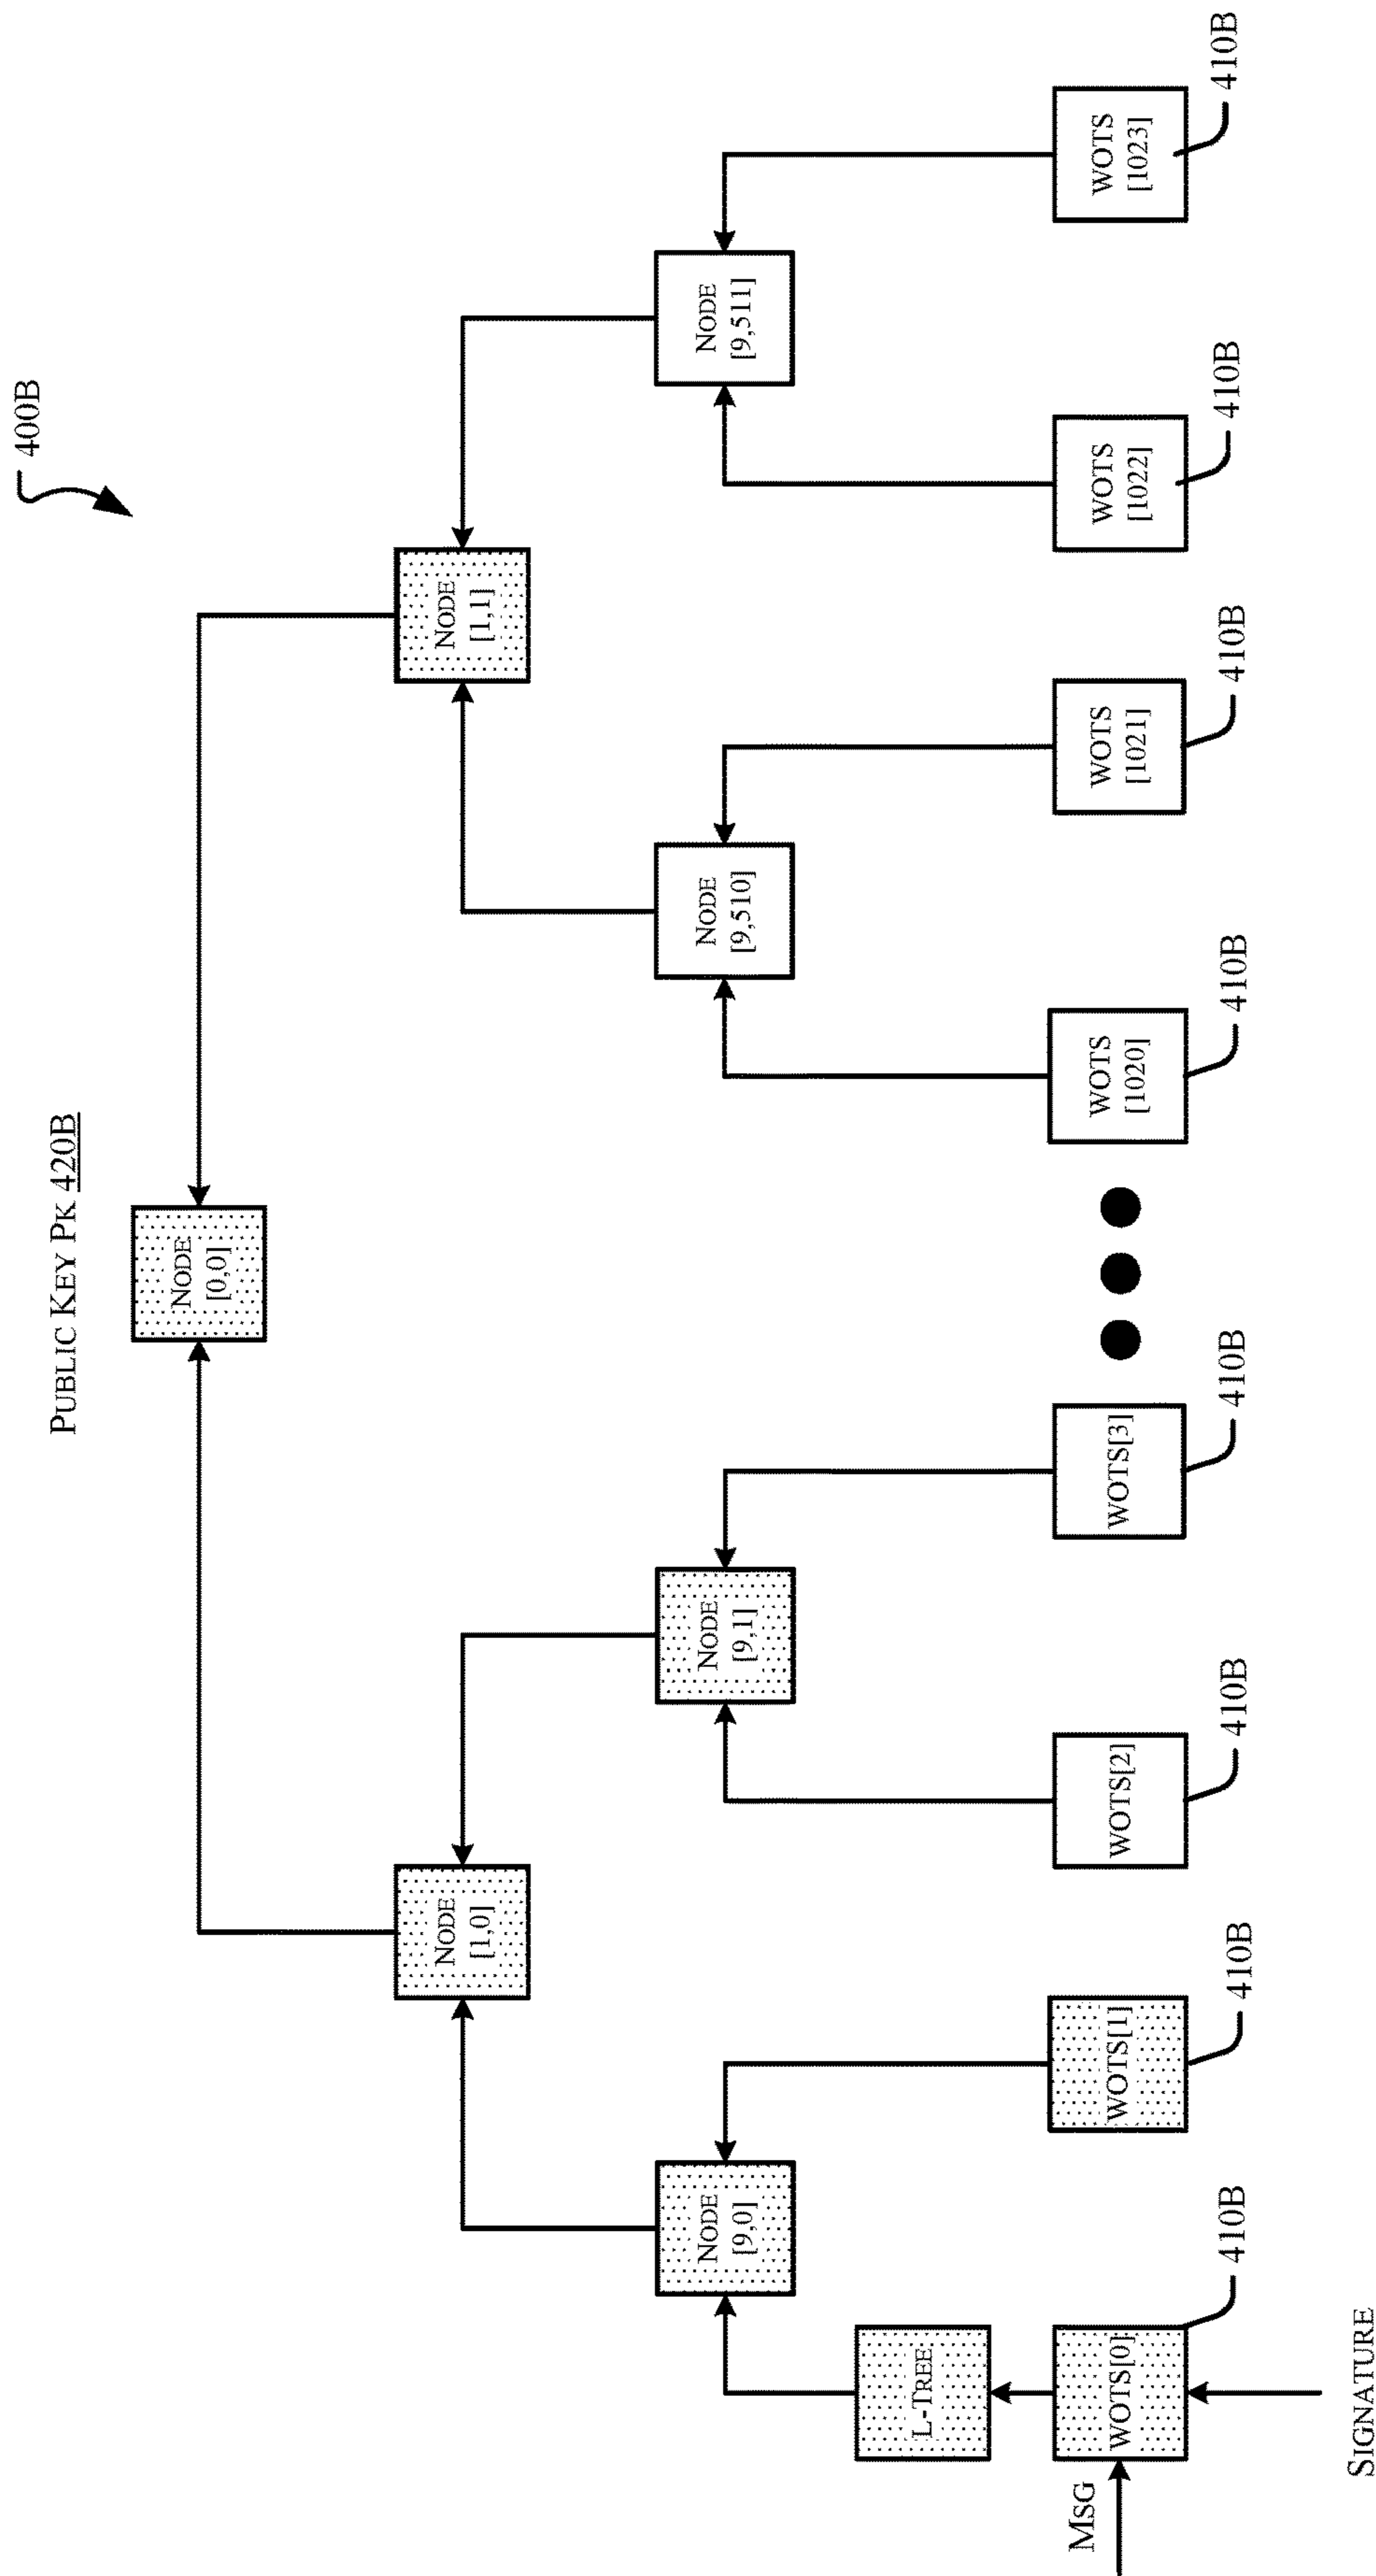


FIG. 4B

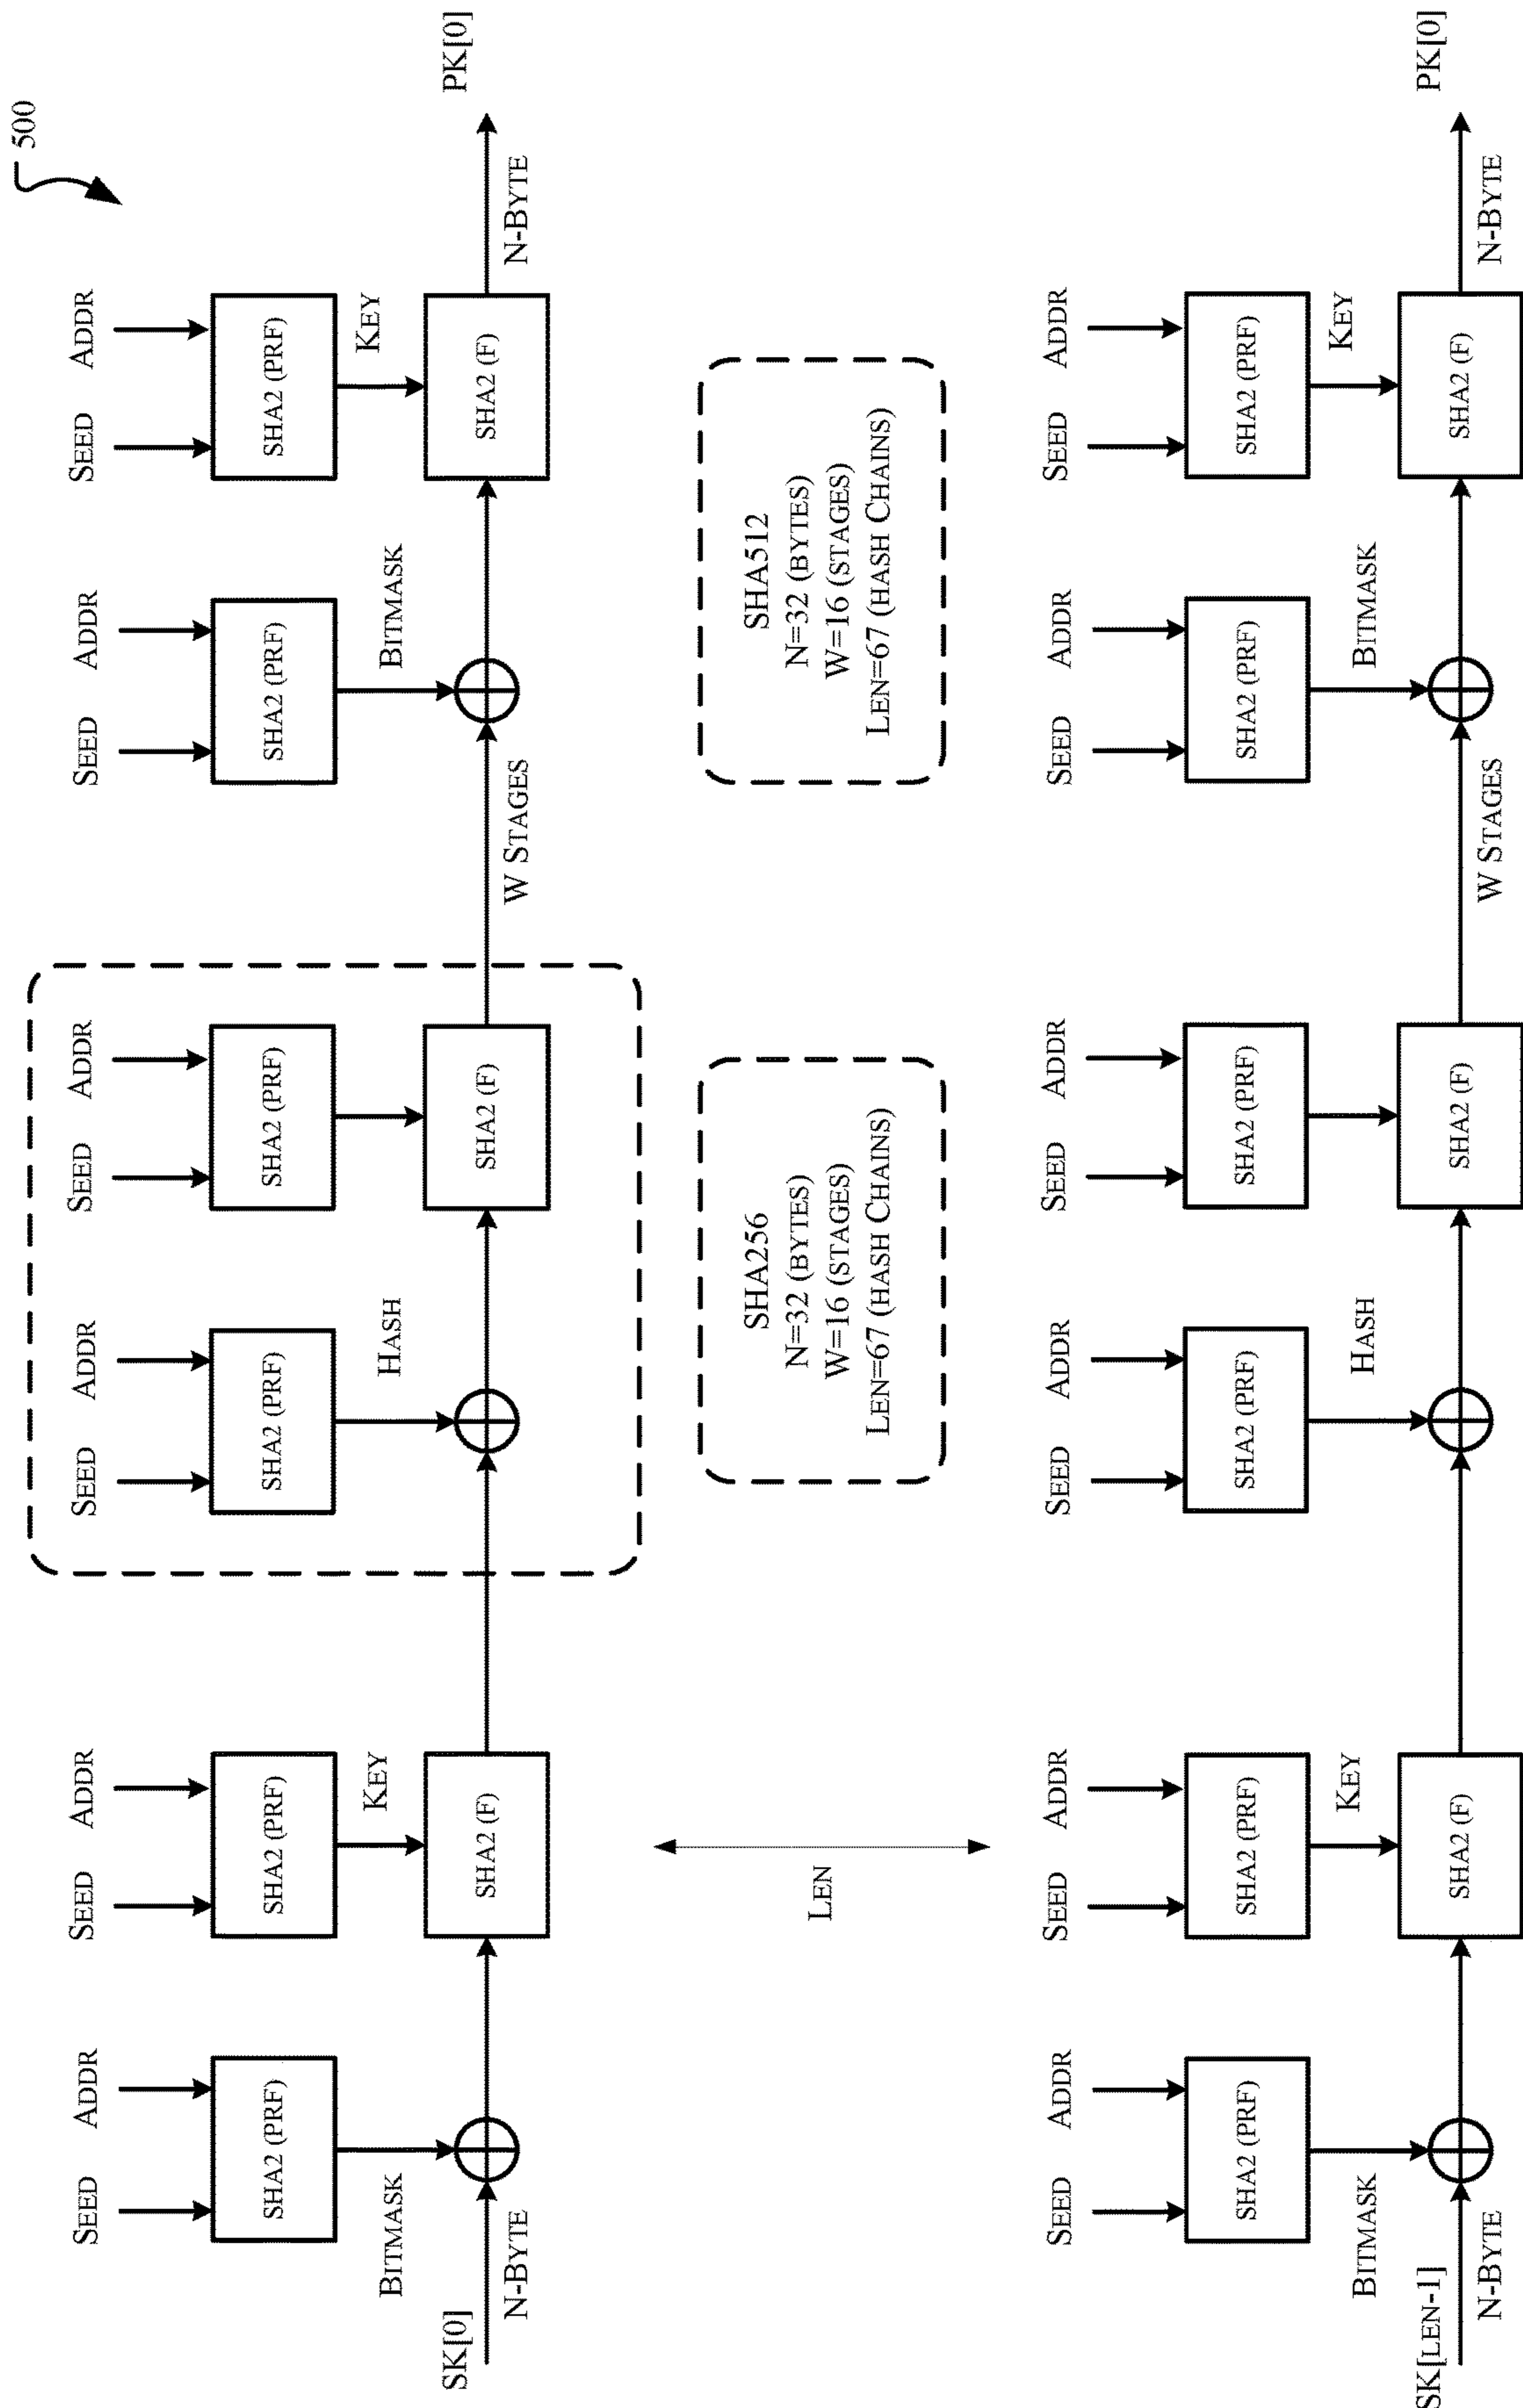


FIG. 5

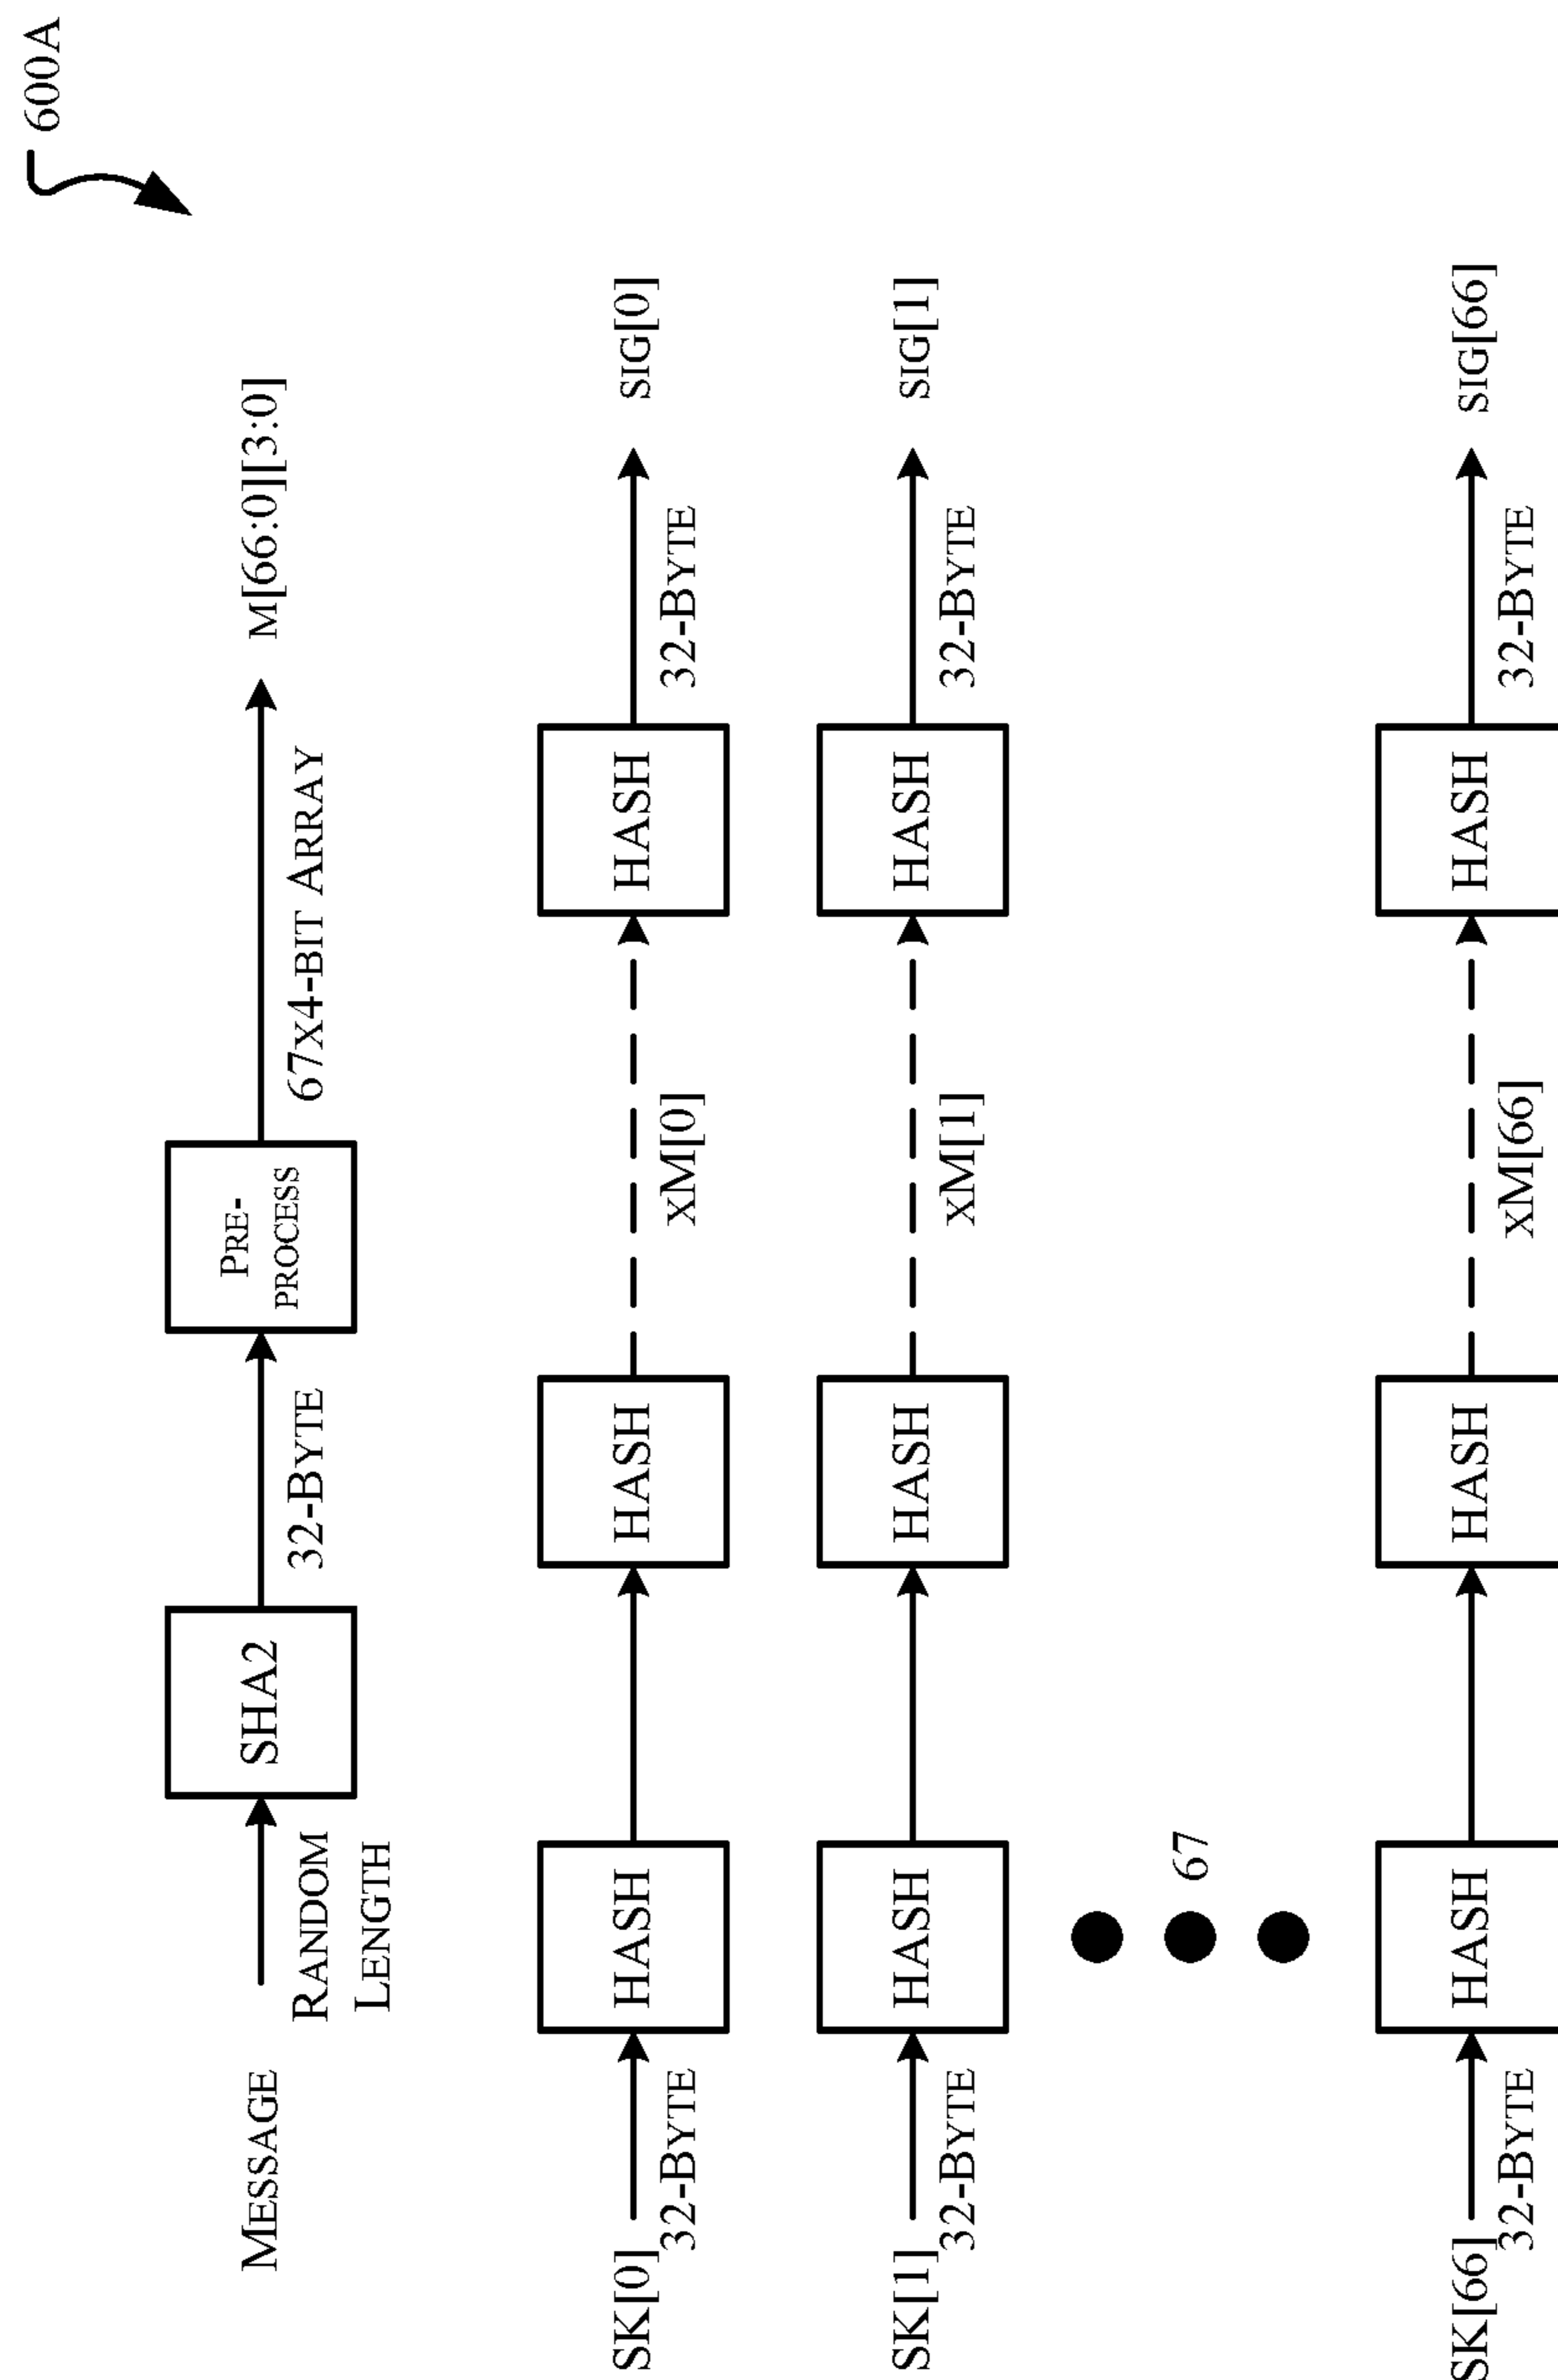


FIG. 6A

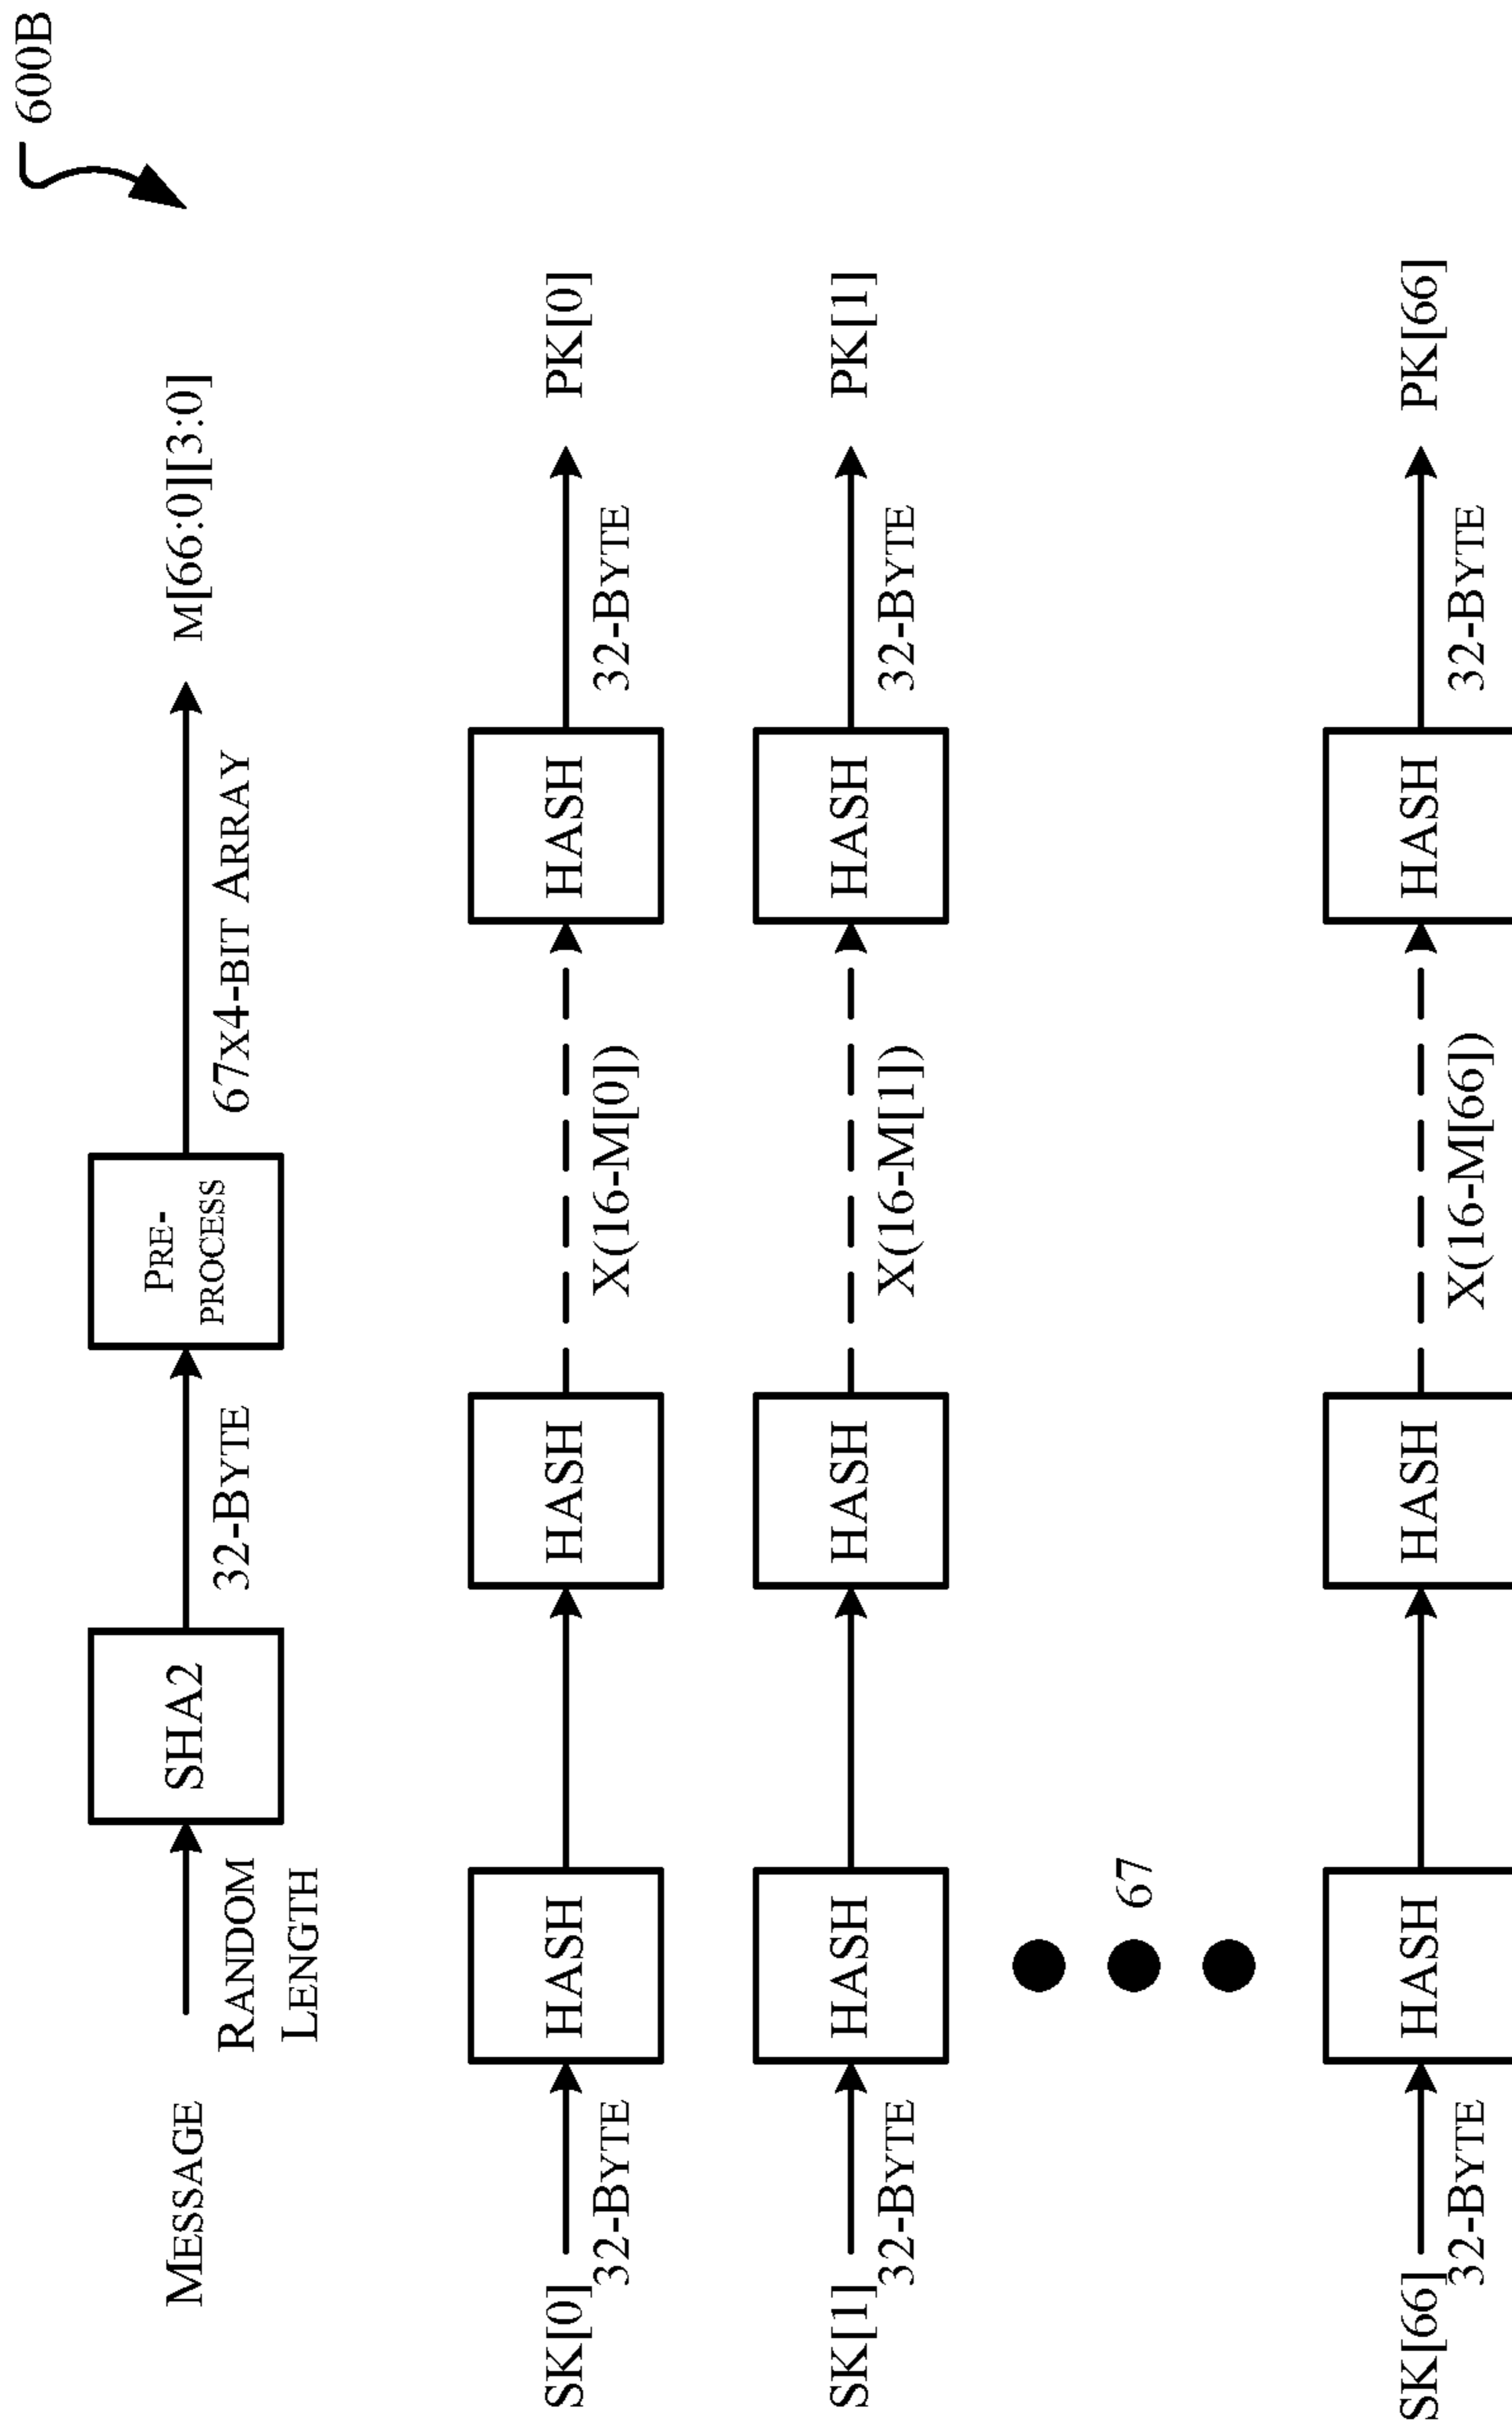


FIG. 6B

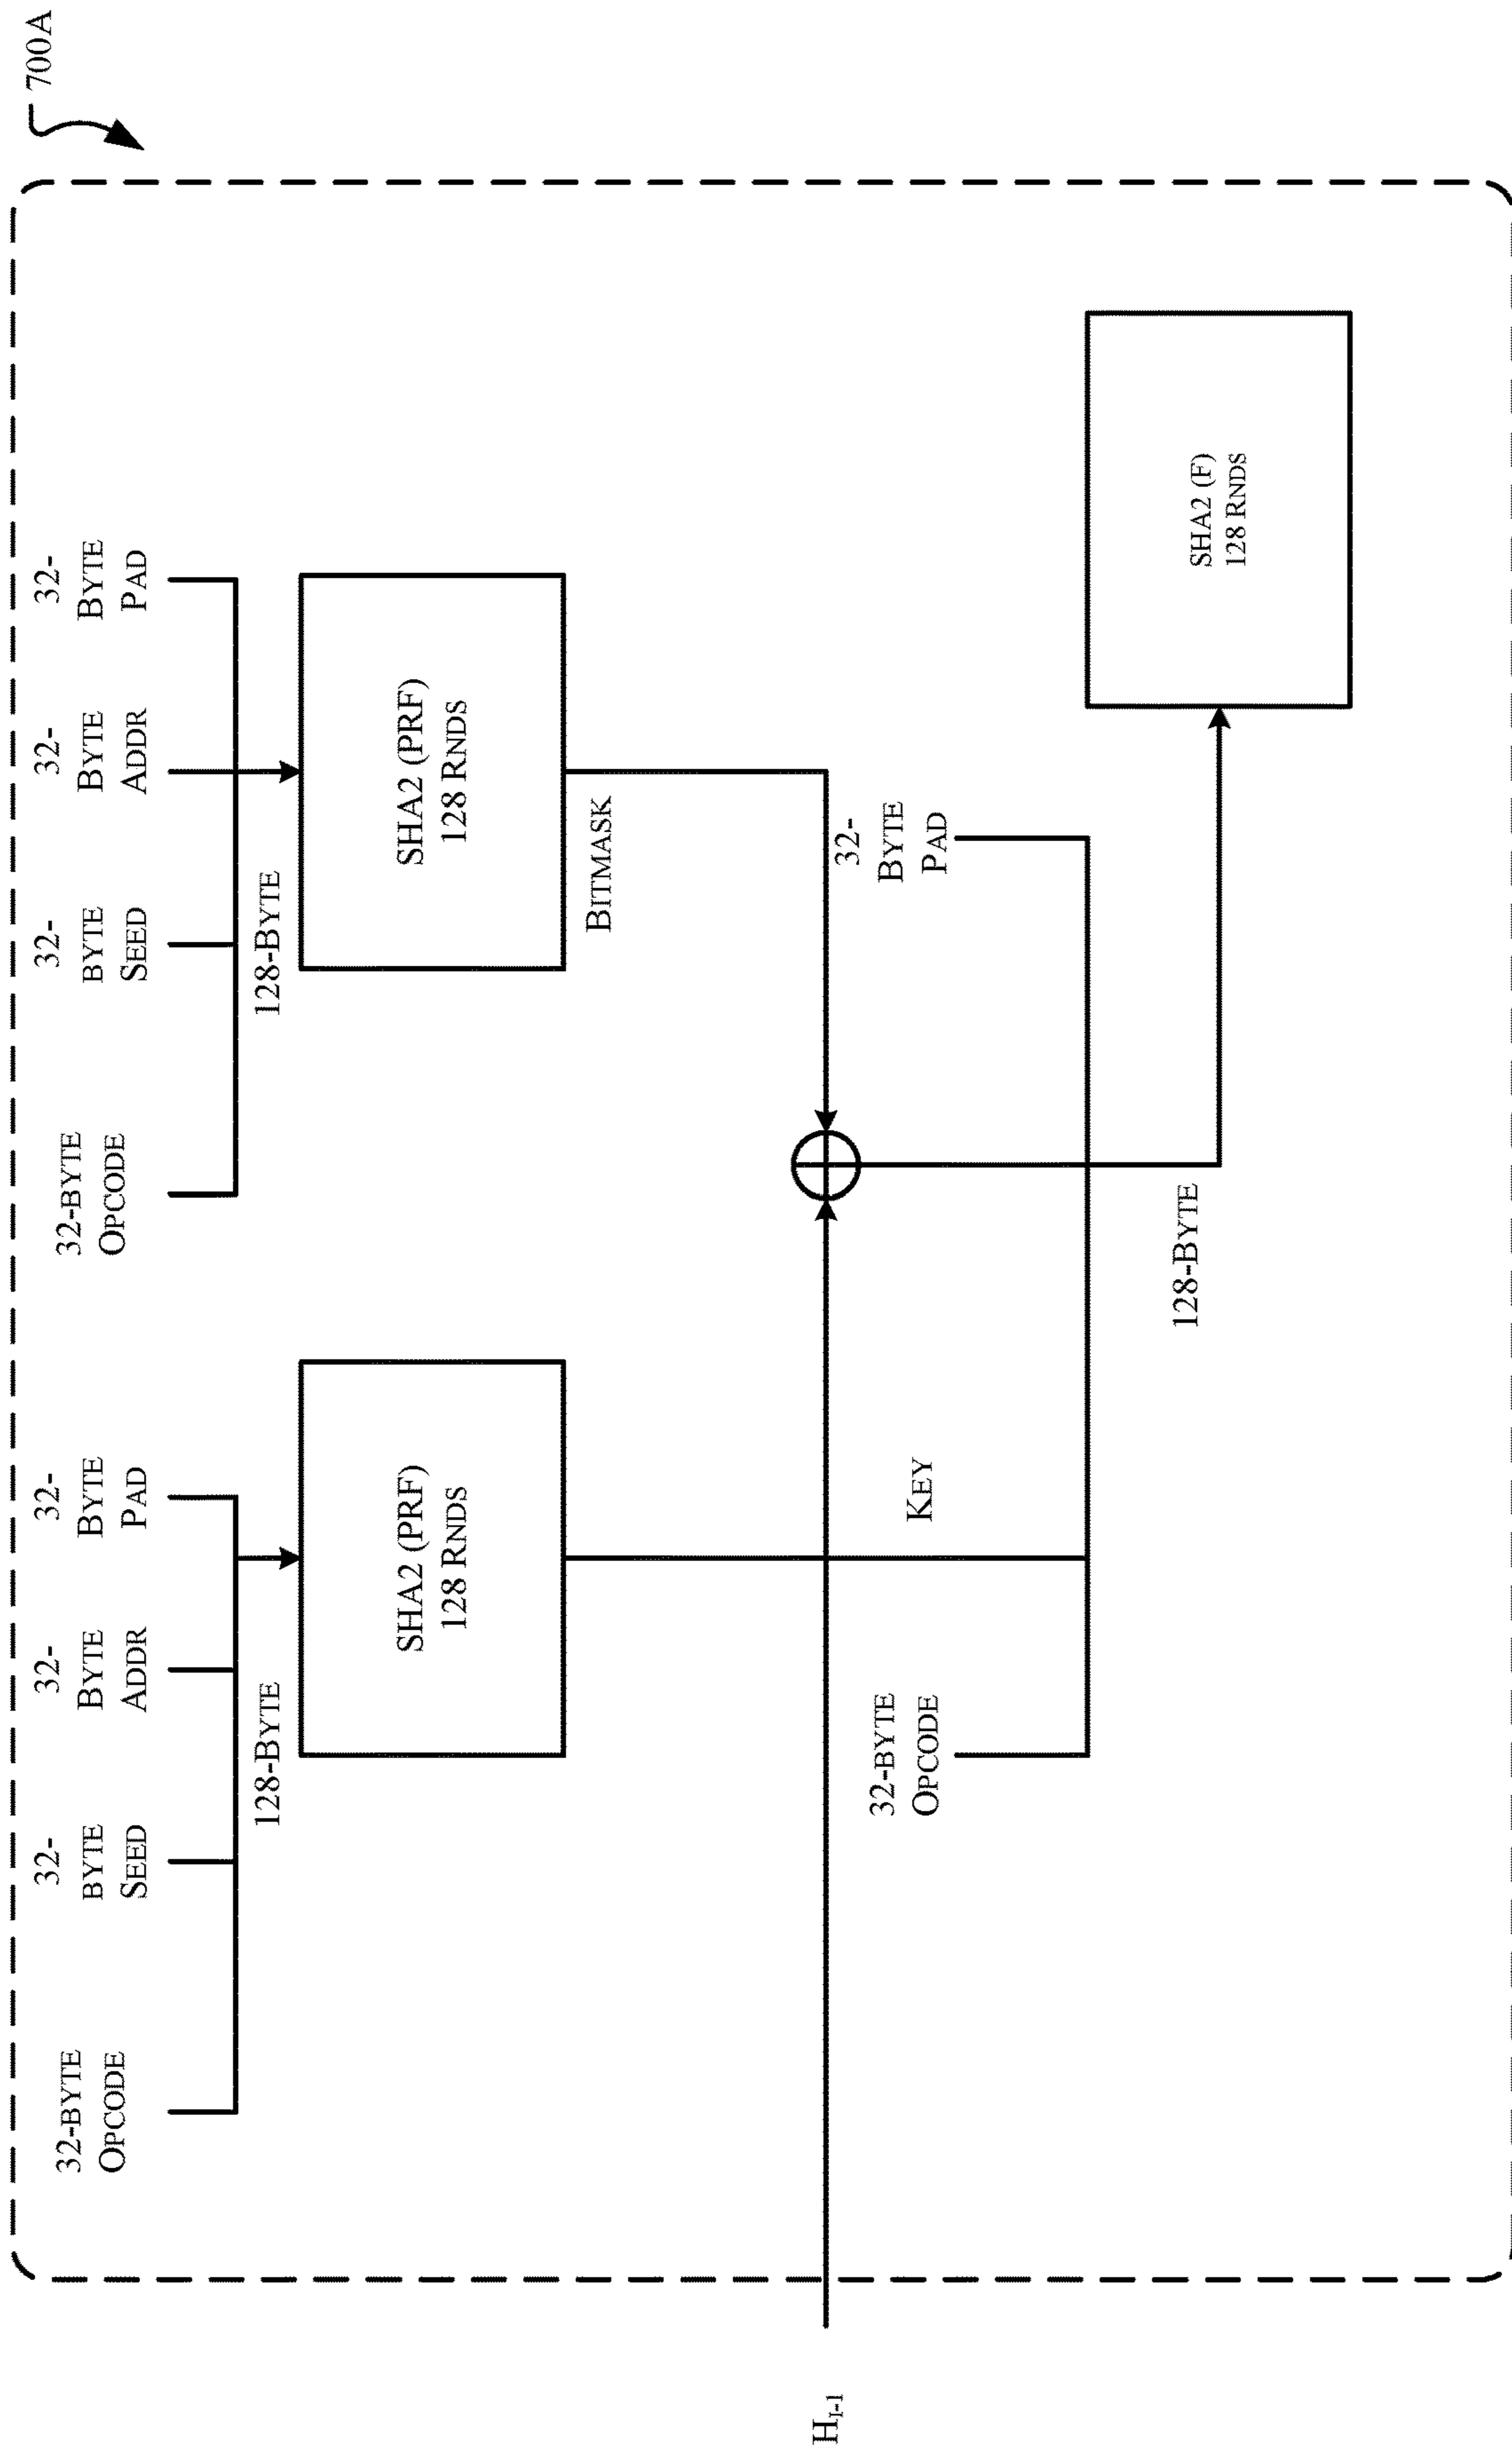


FIG. 7A

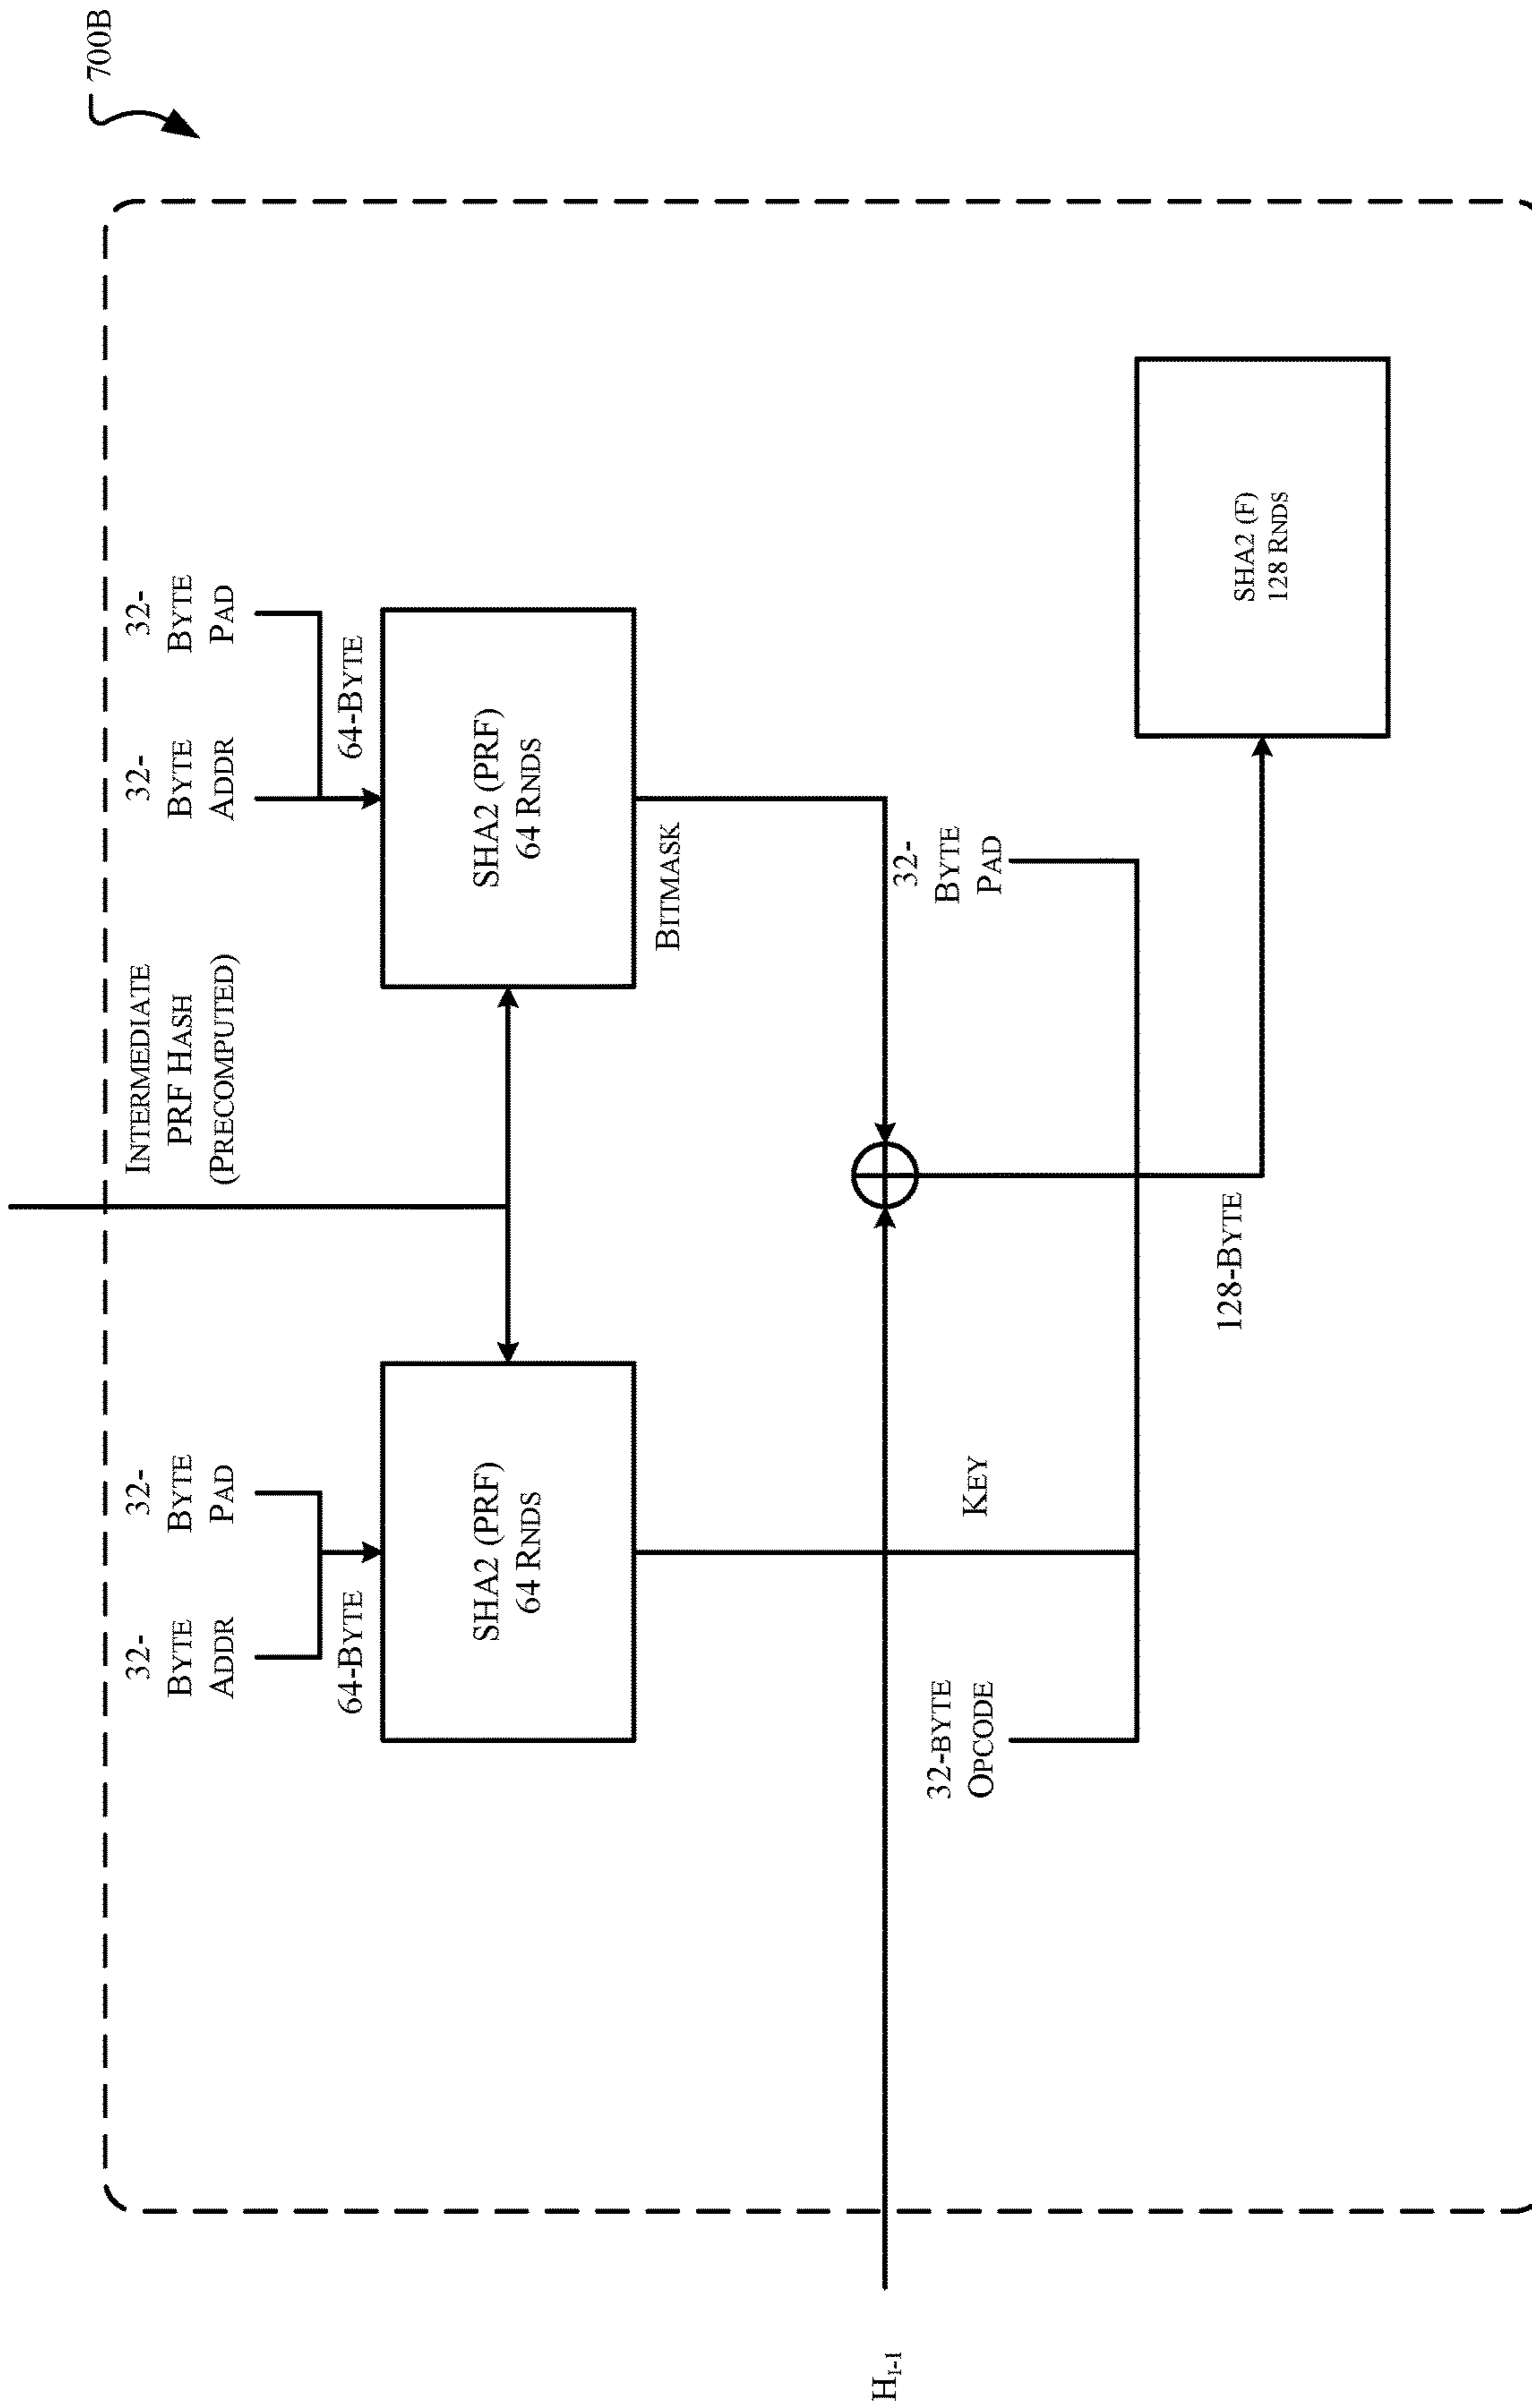


FIG. 7B

700B

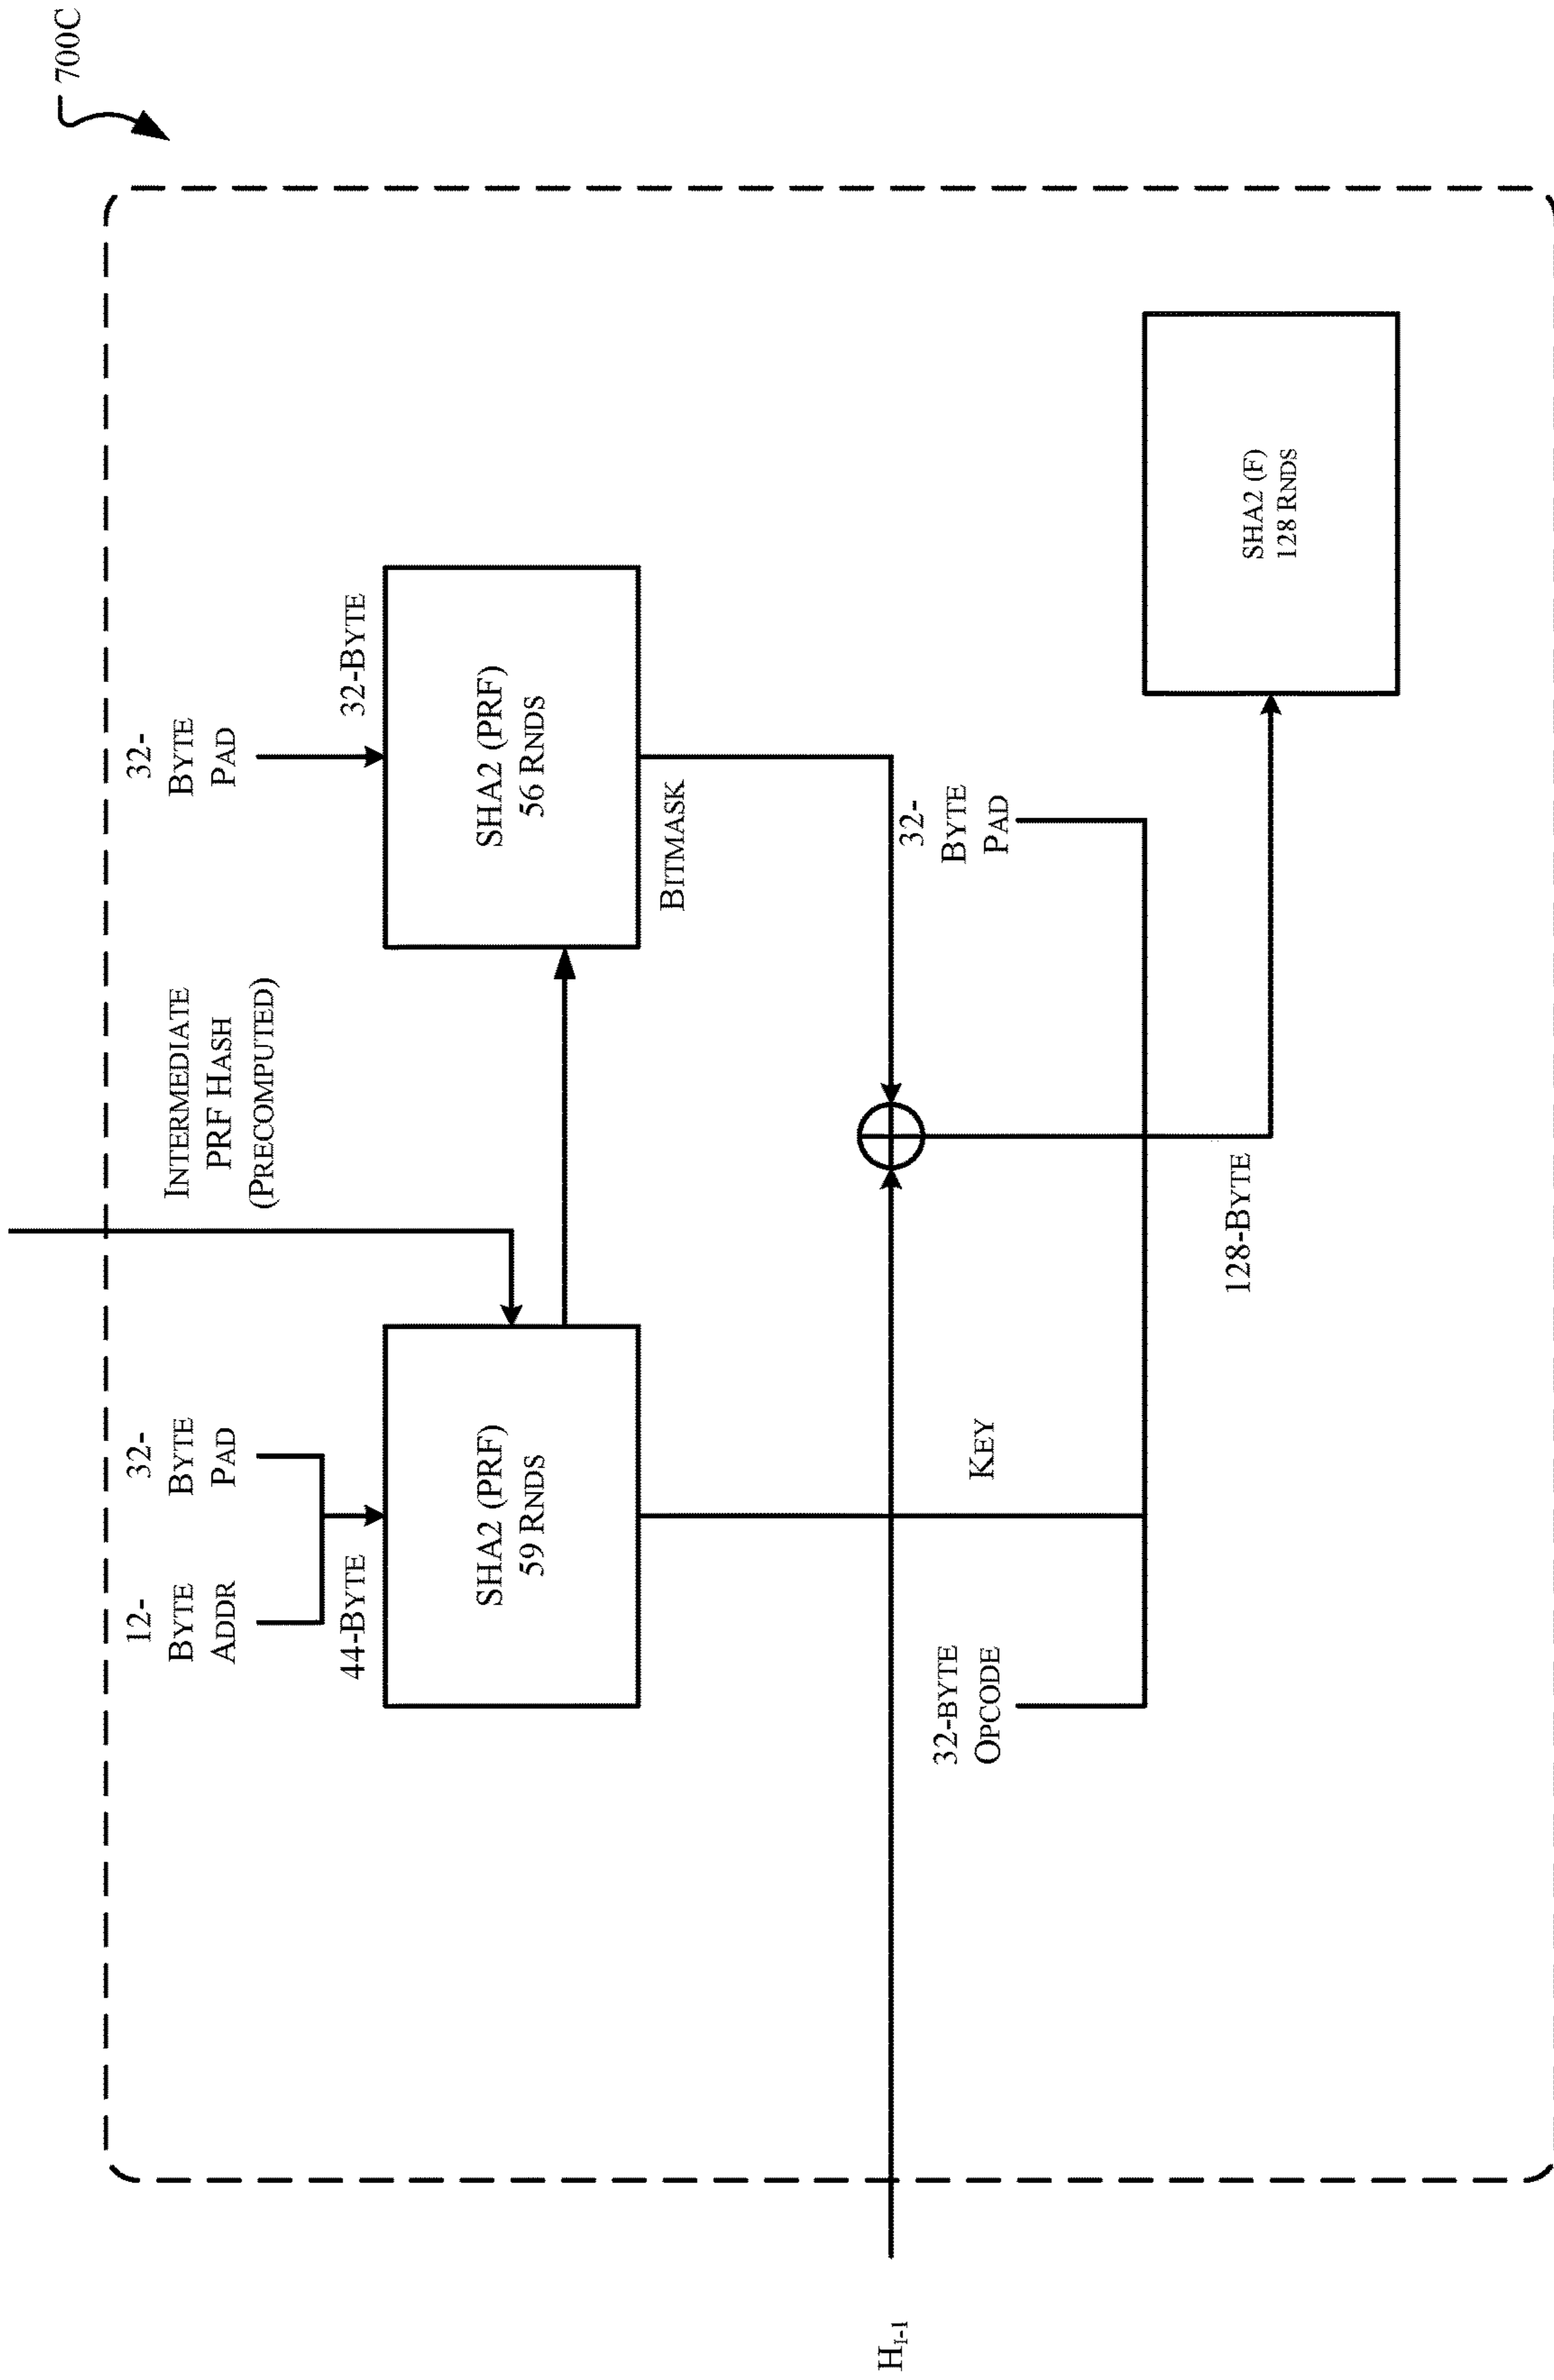


FIG. 7C

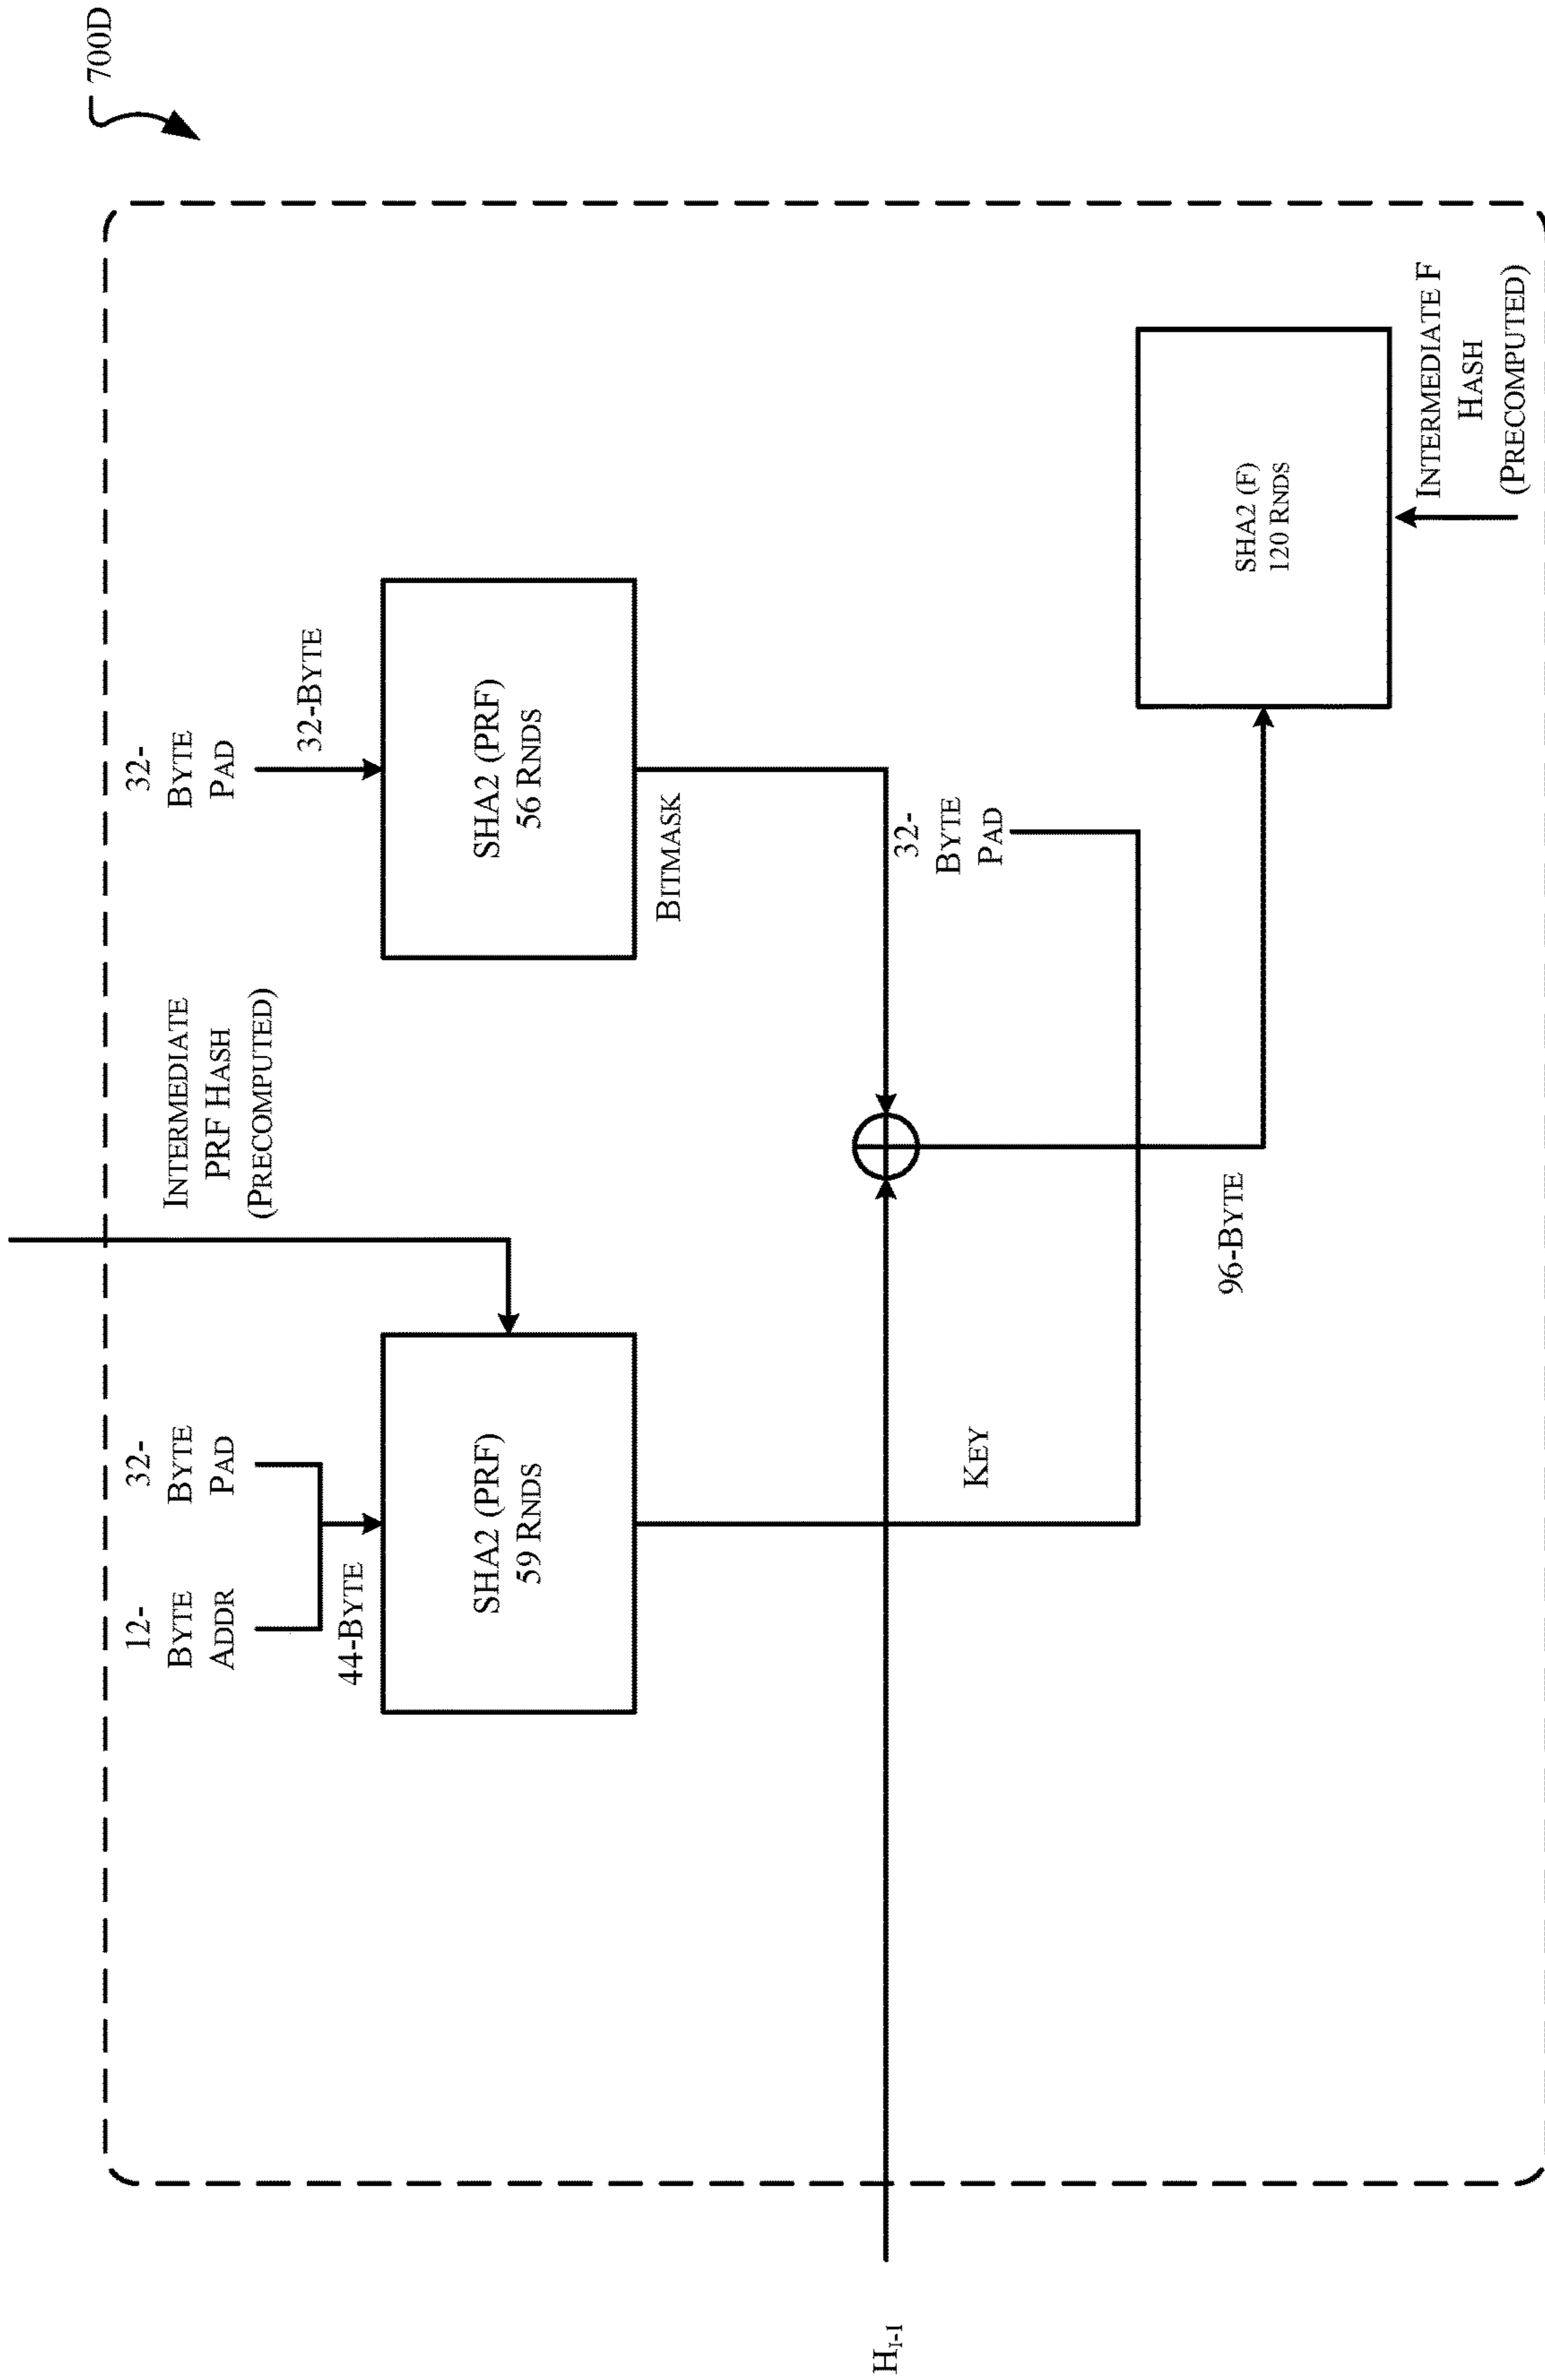


FIG. 7D

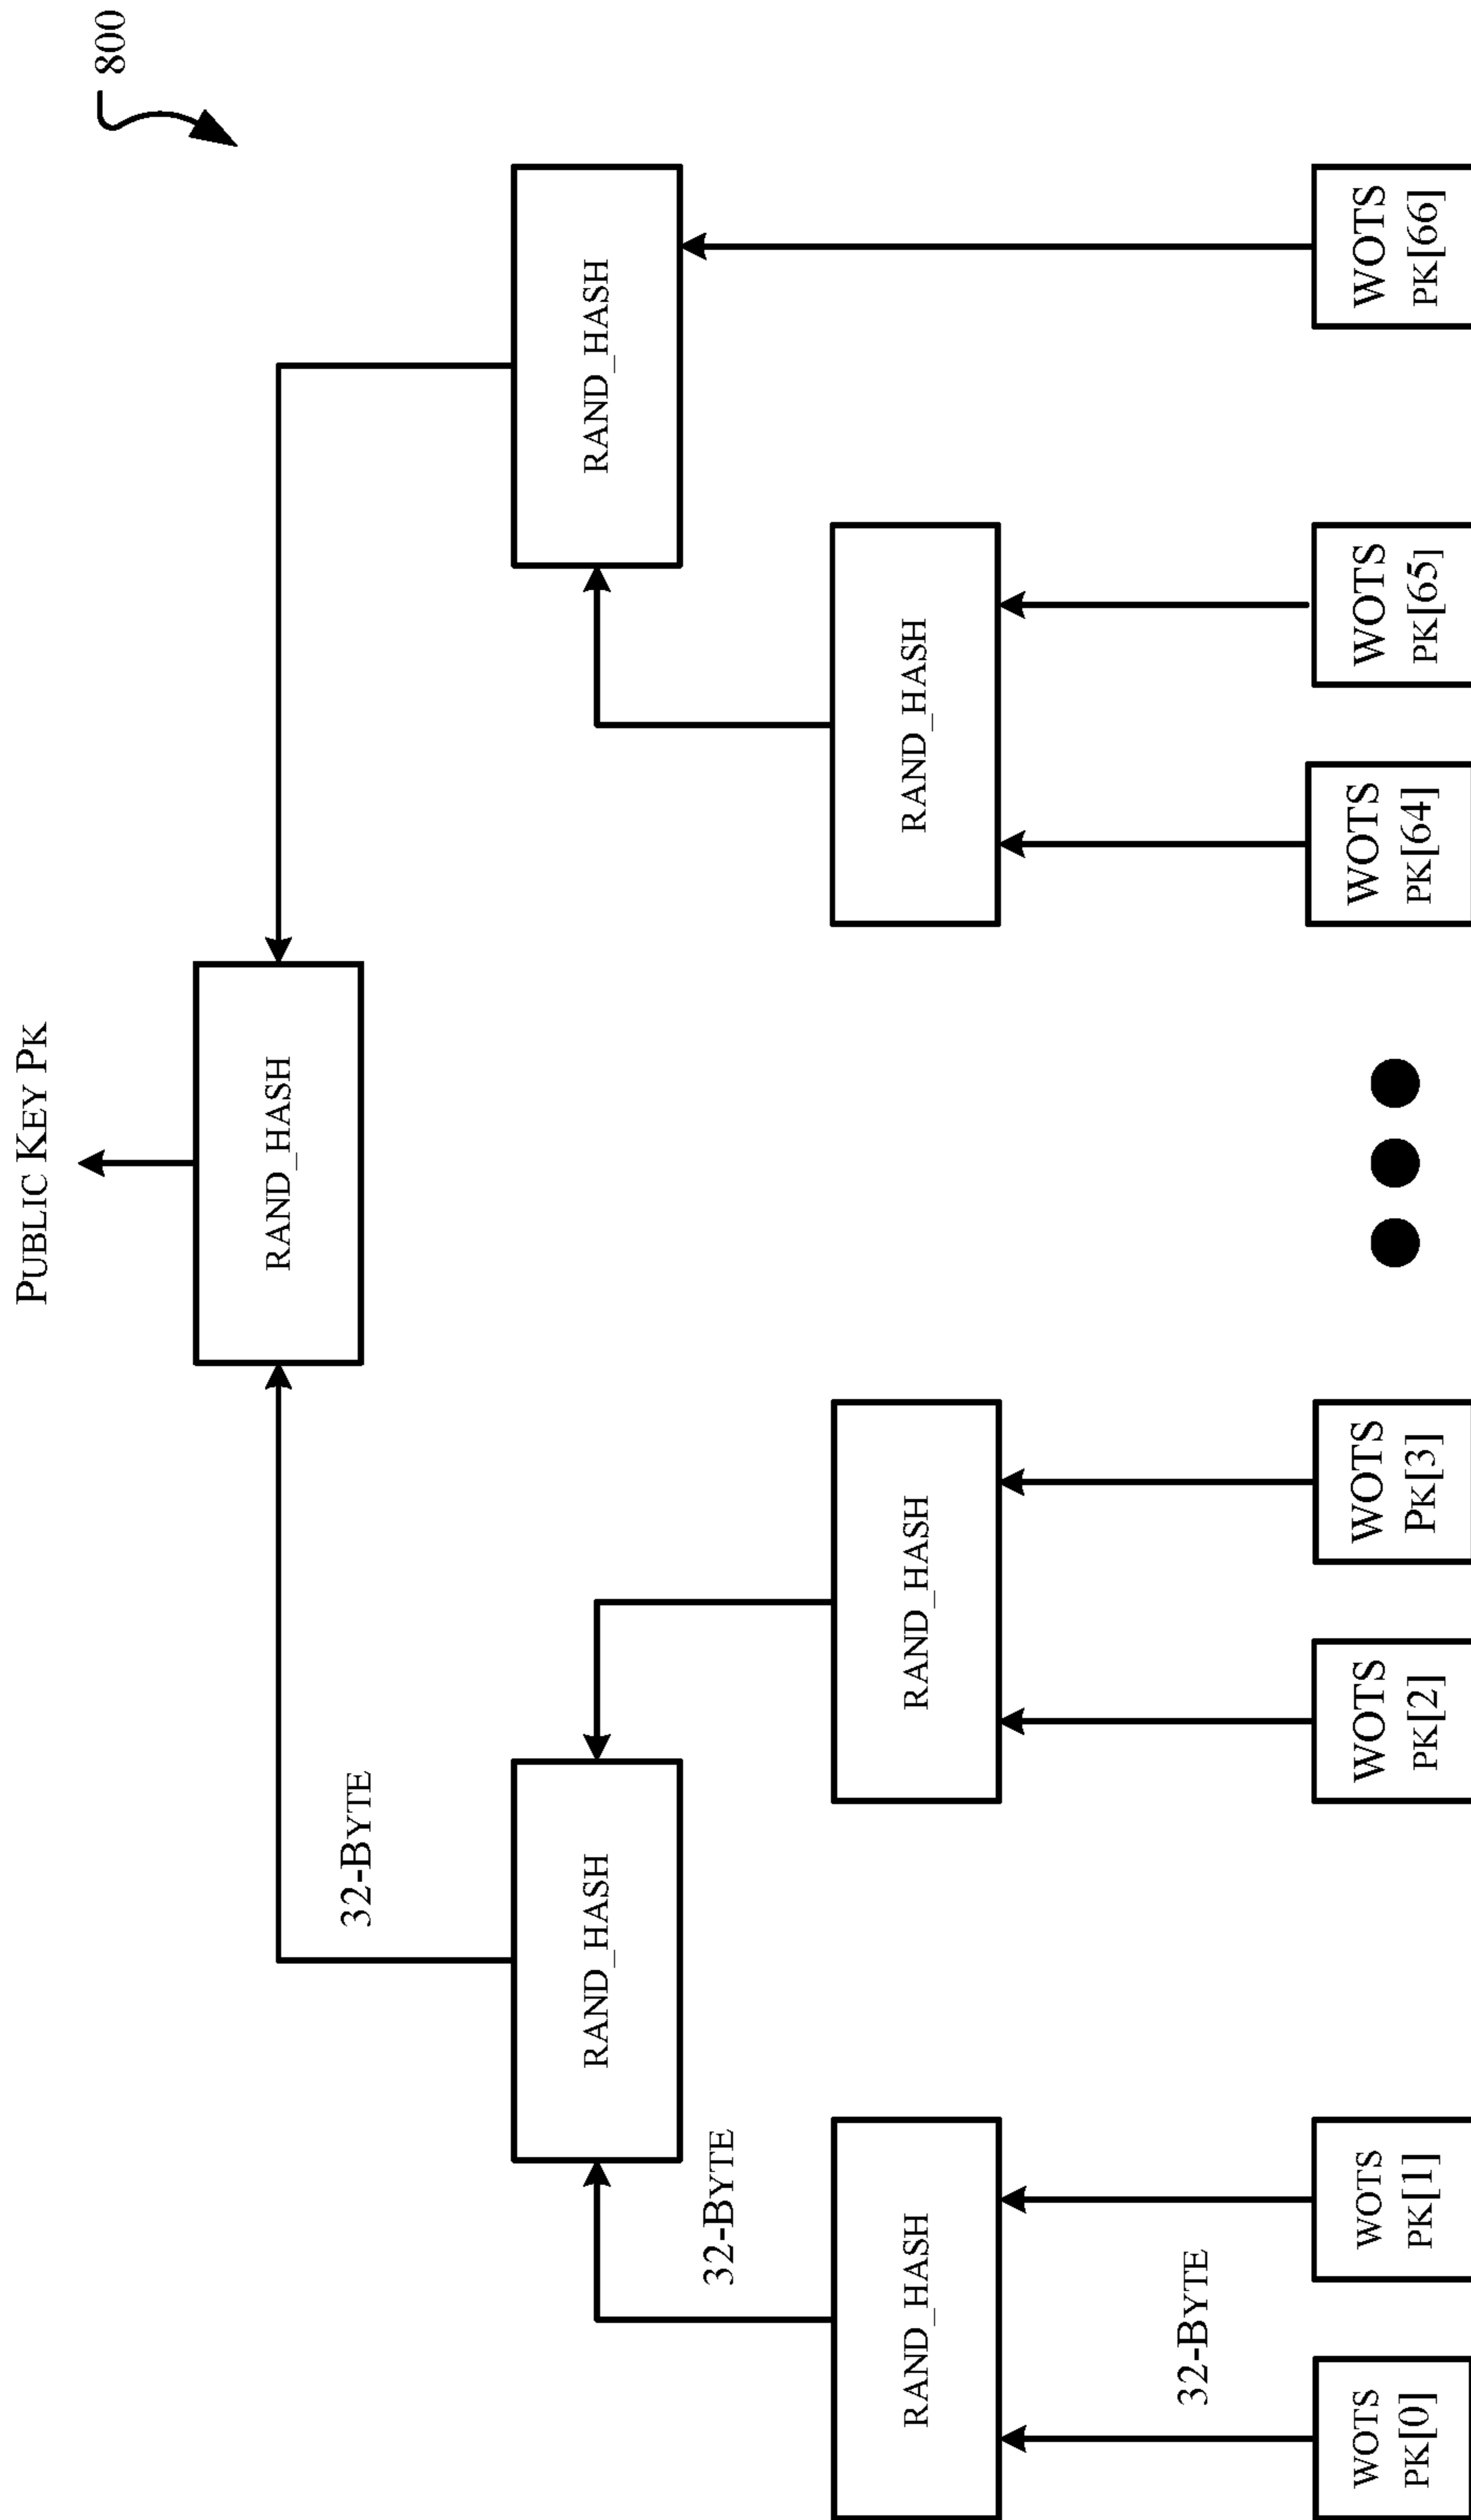


FIG. 8

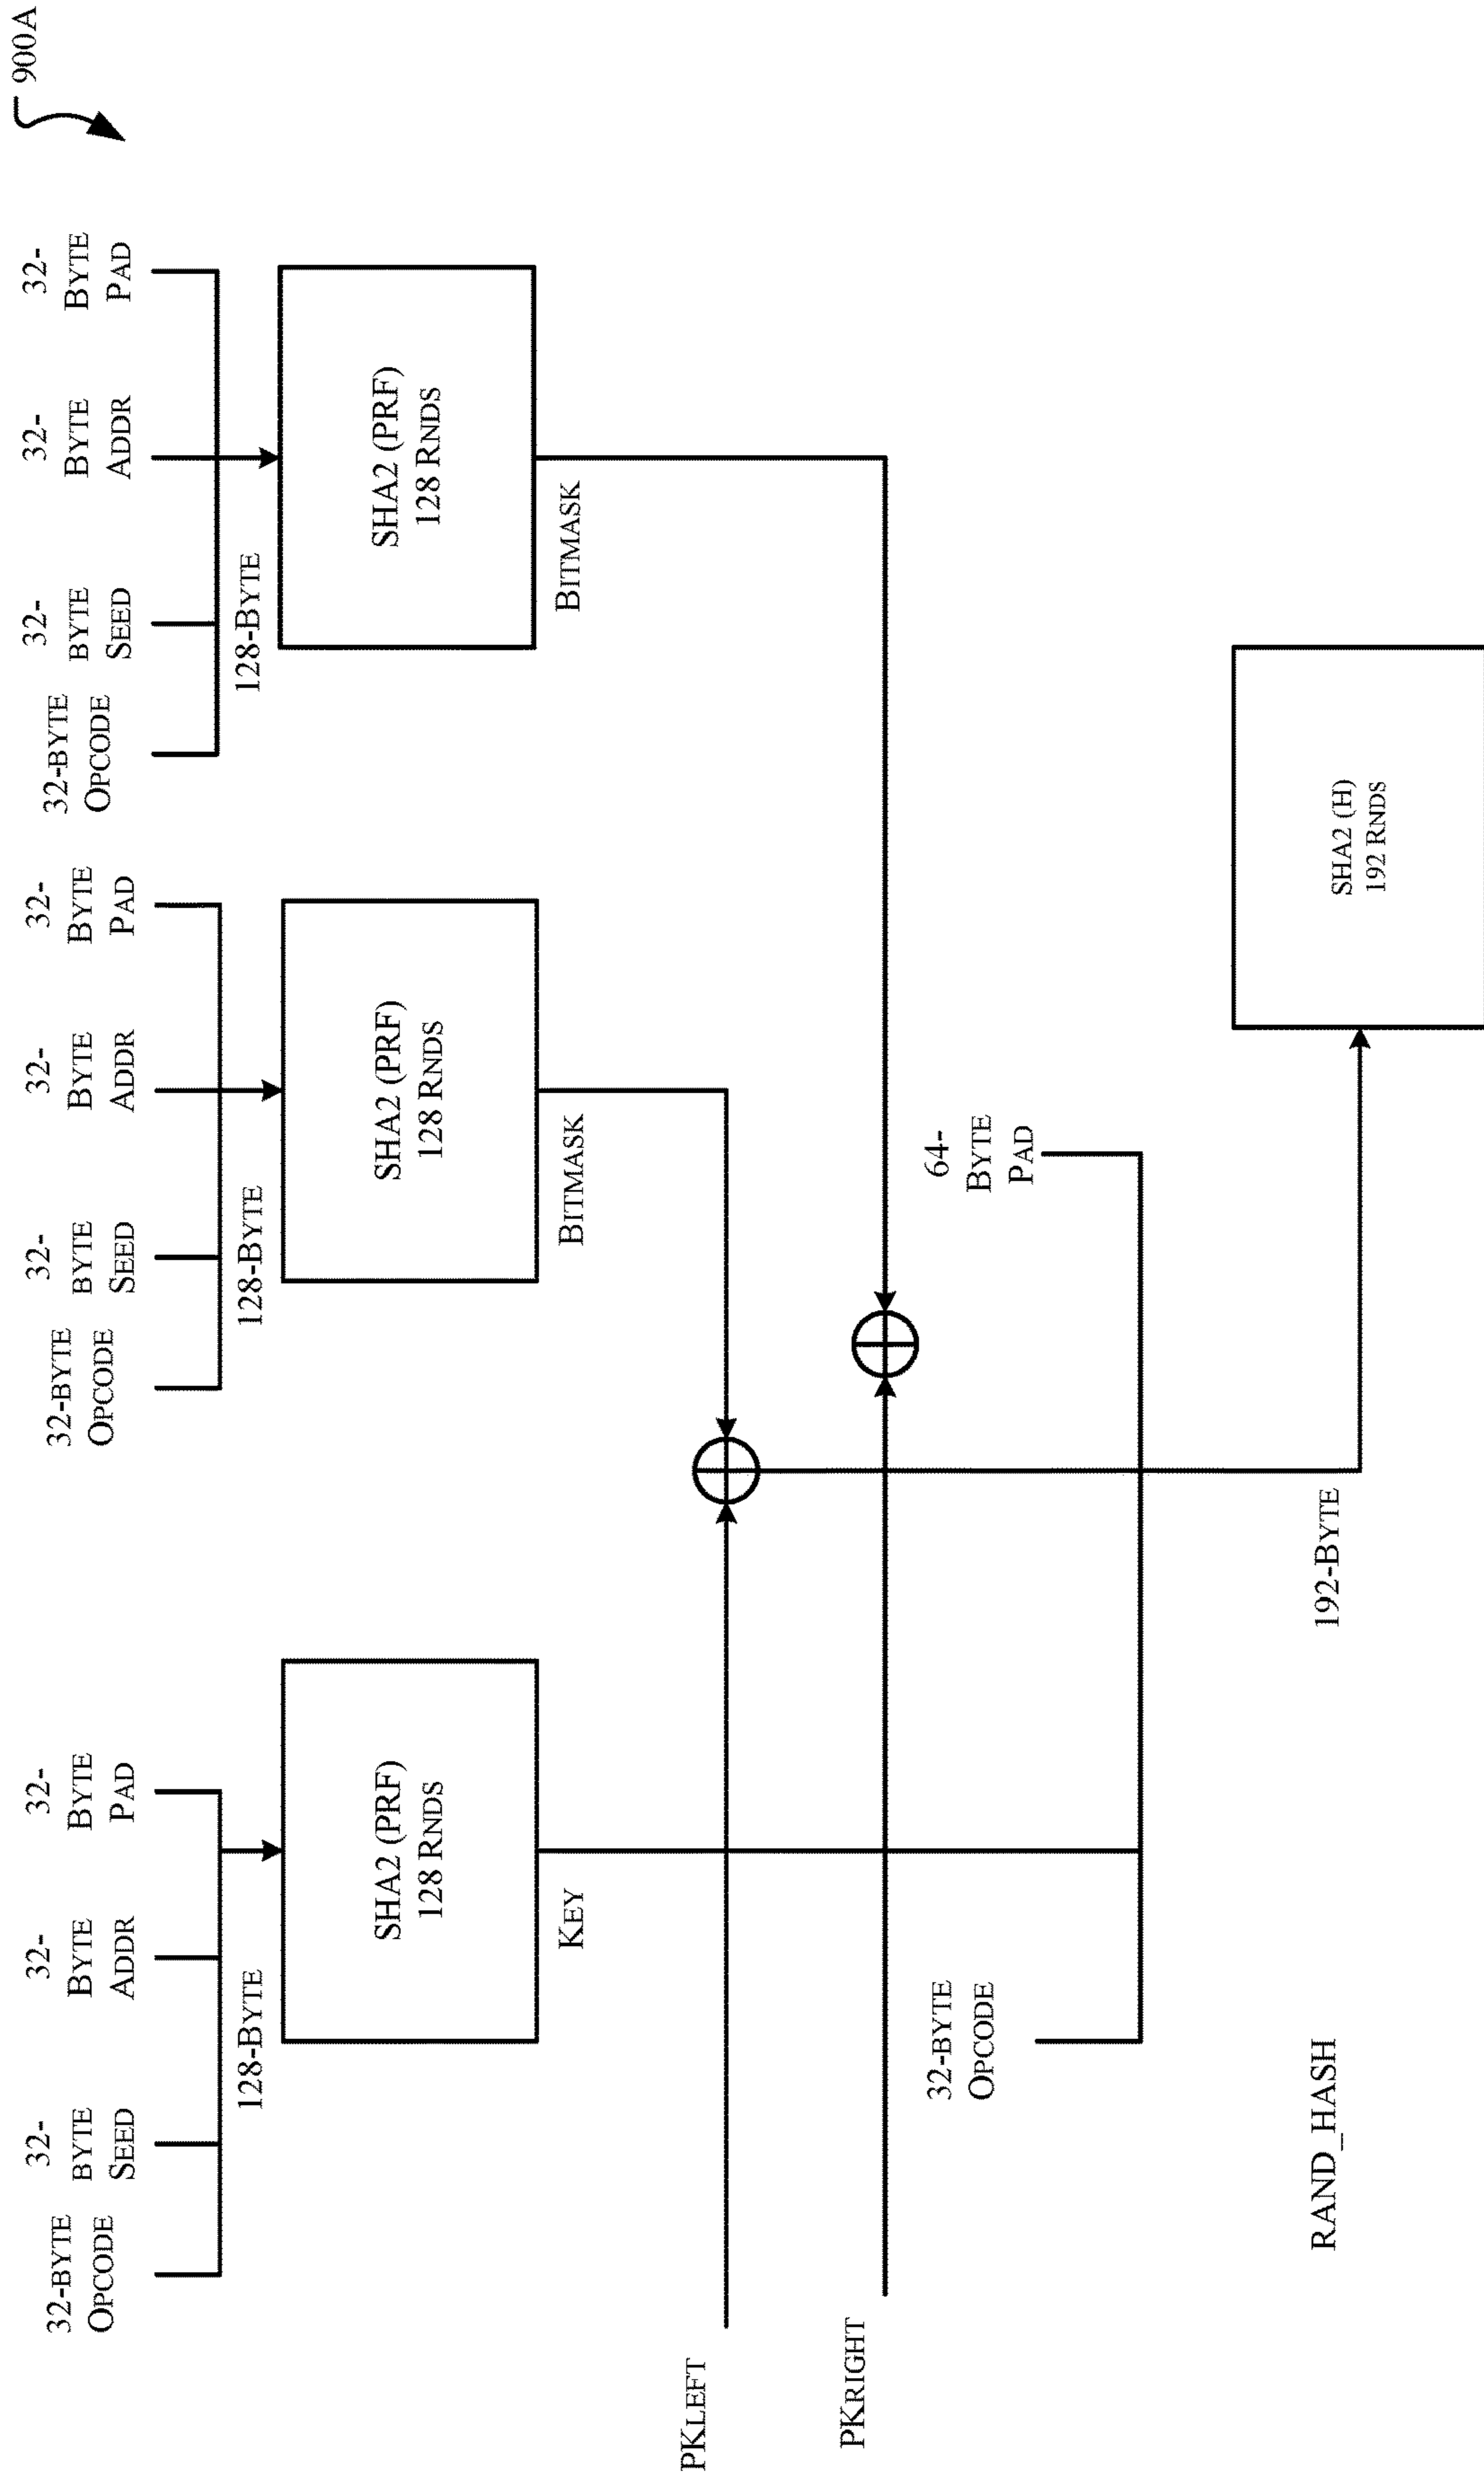


FIG. 9A

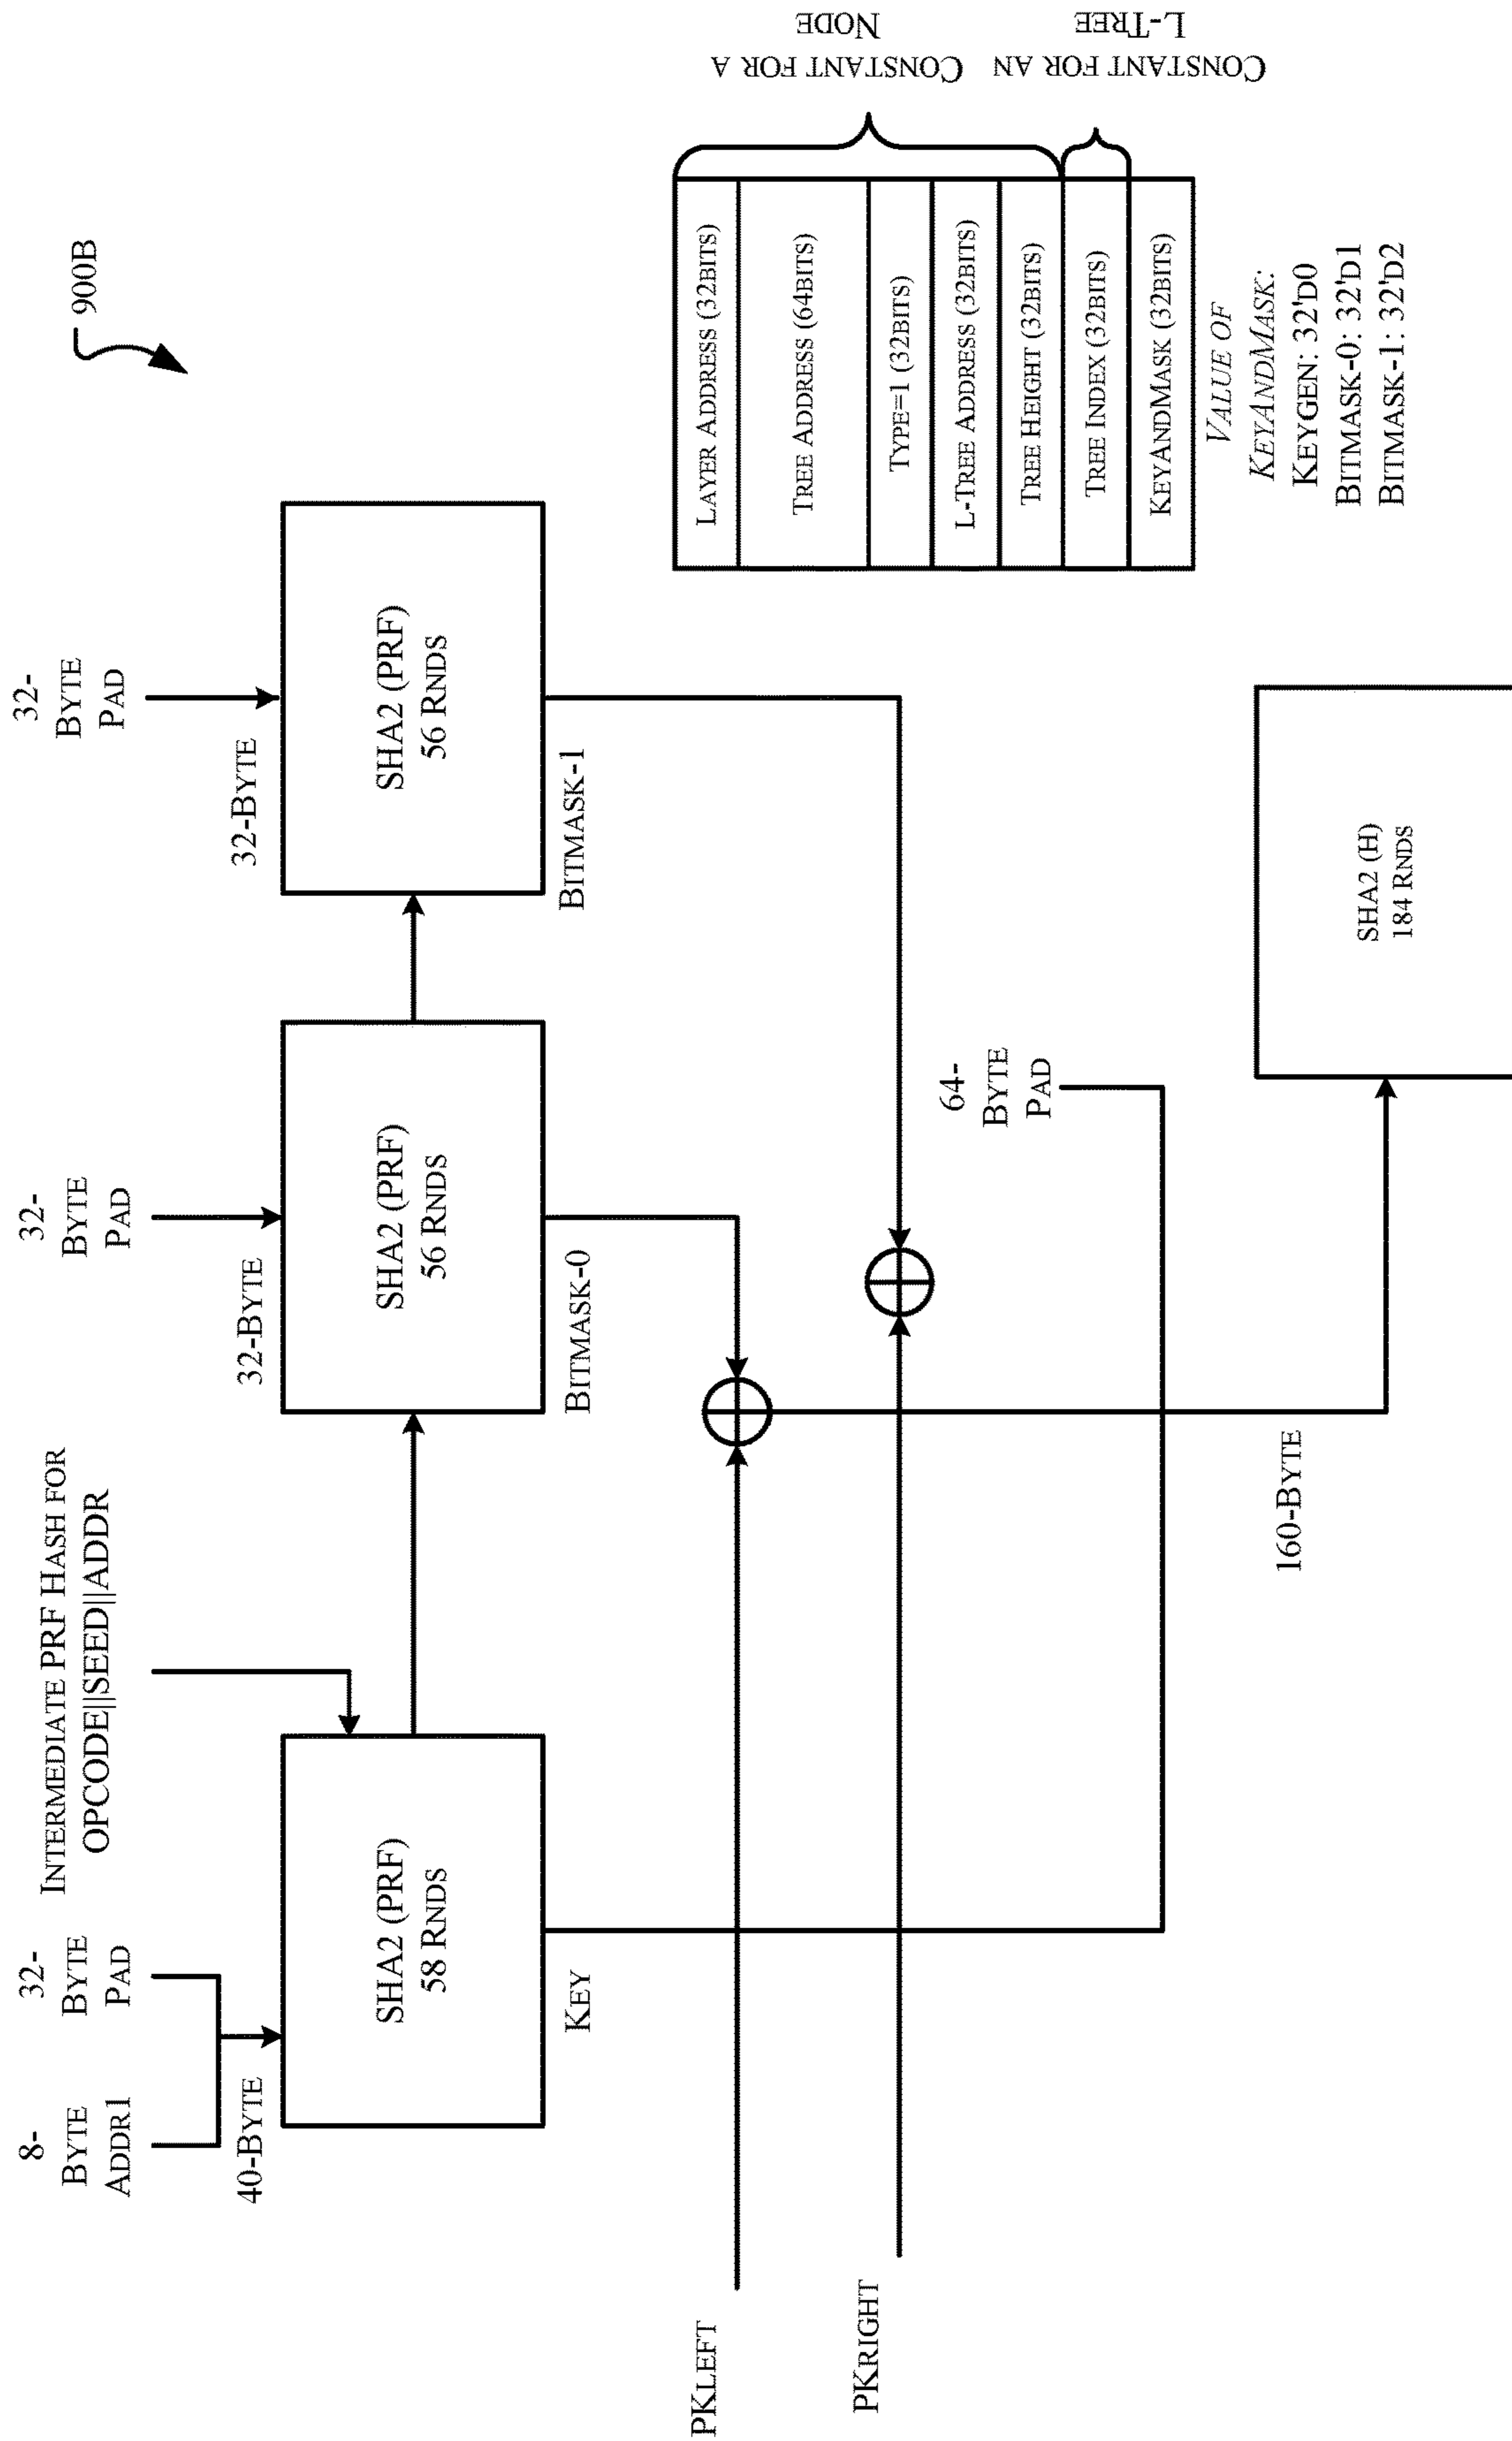


FIG. 9B

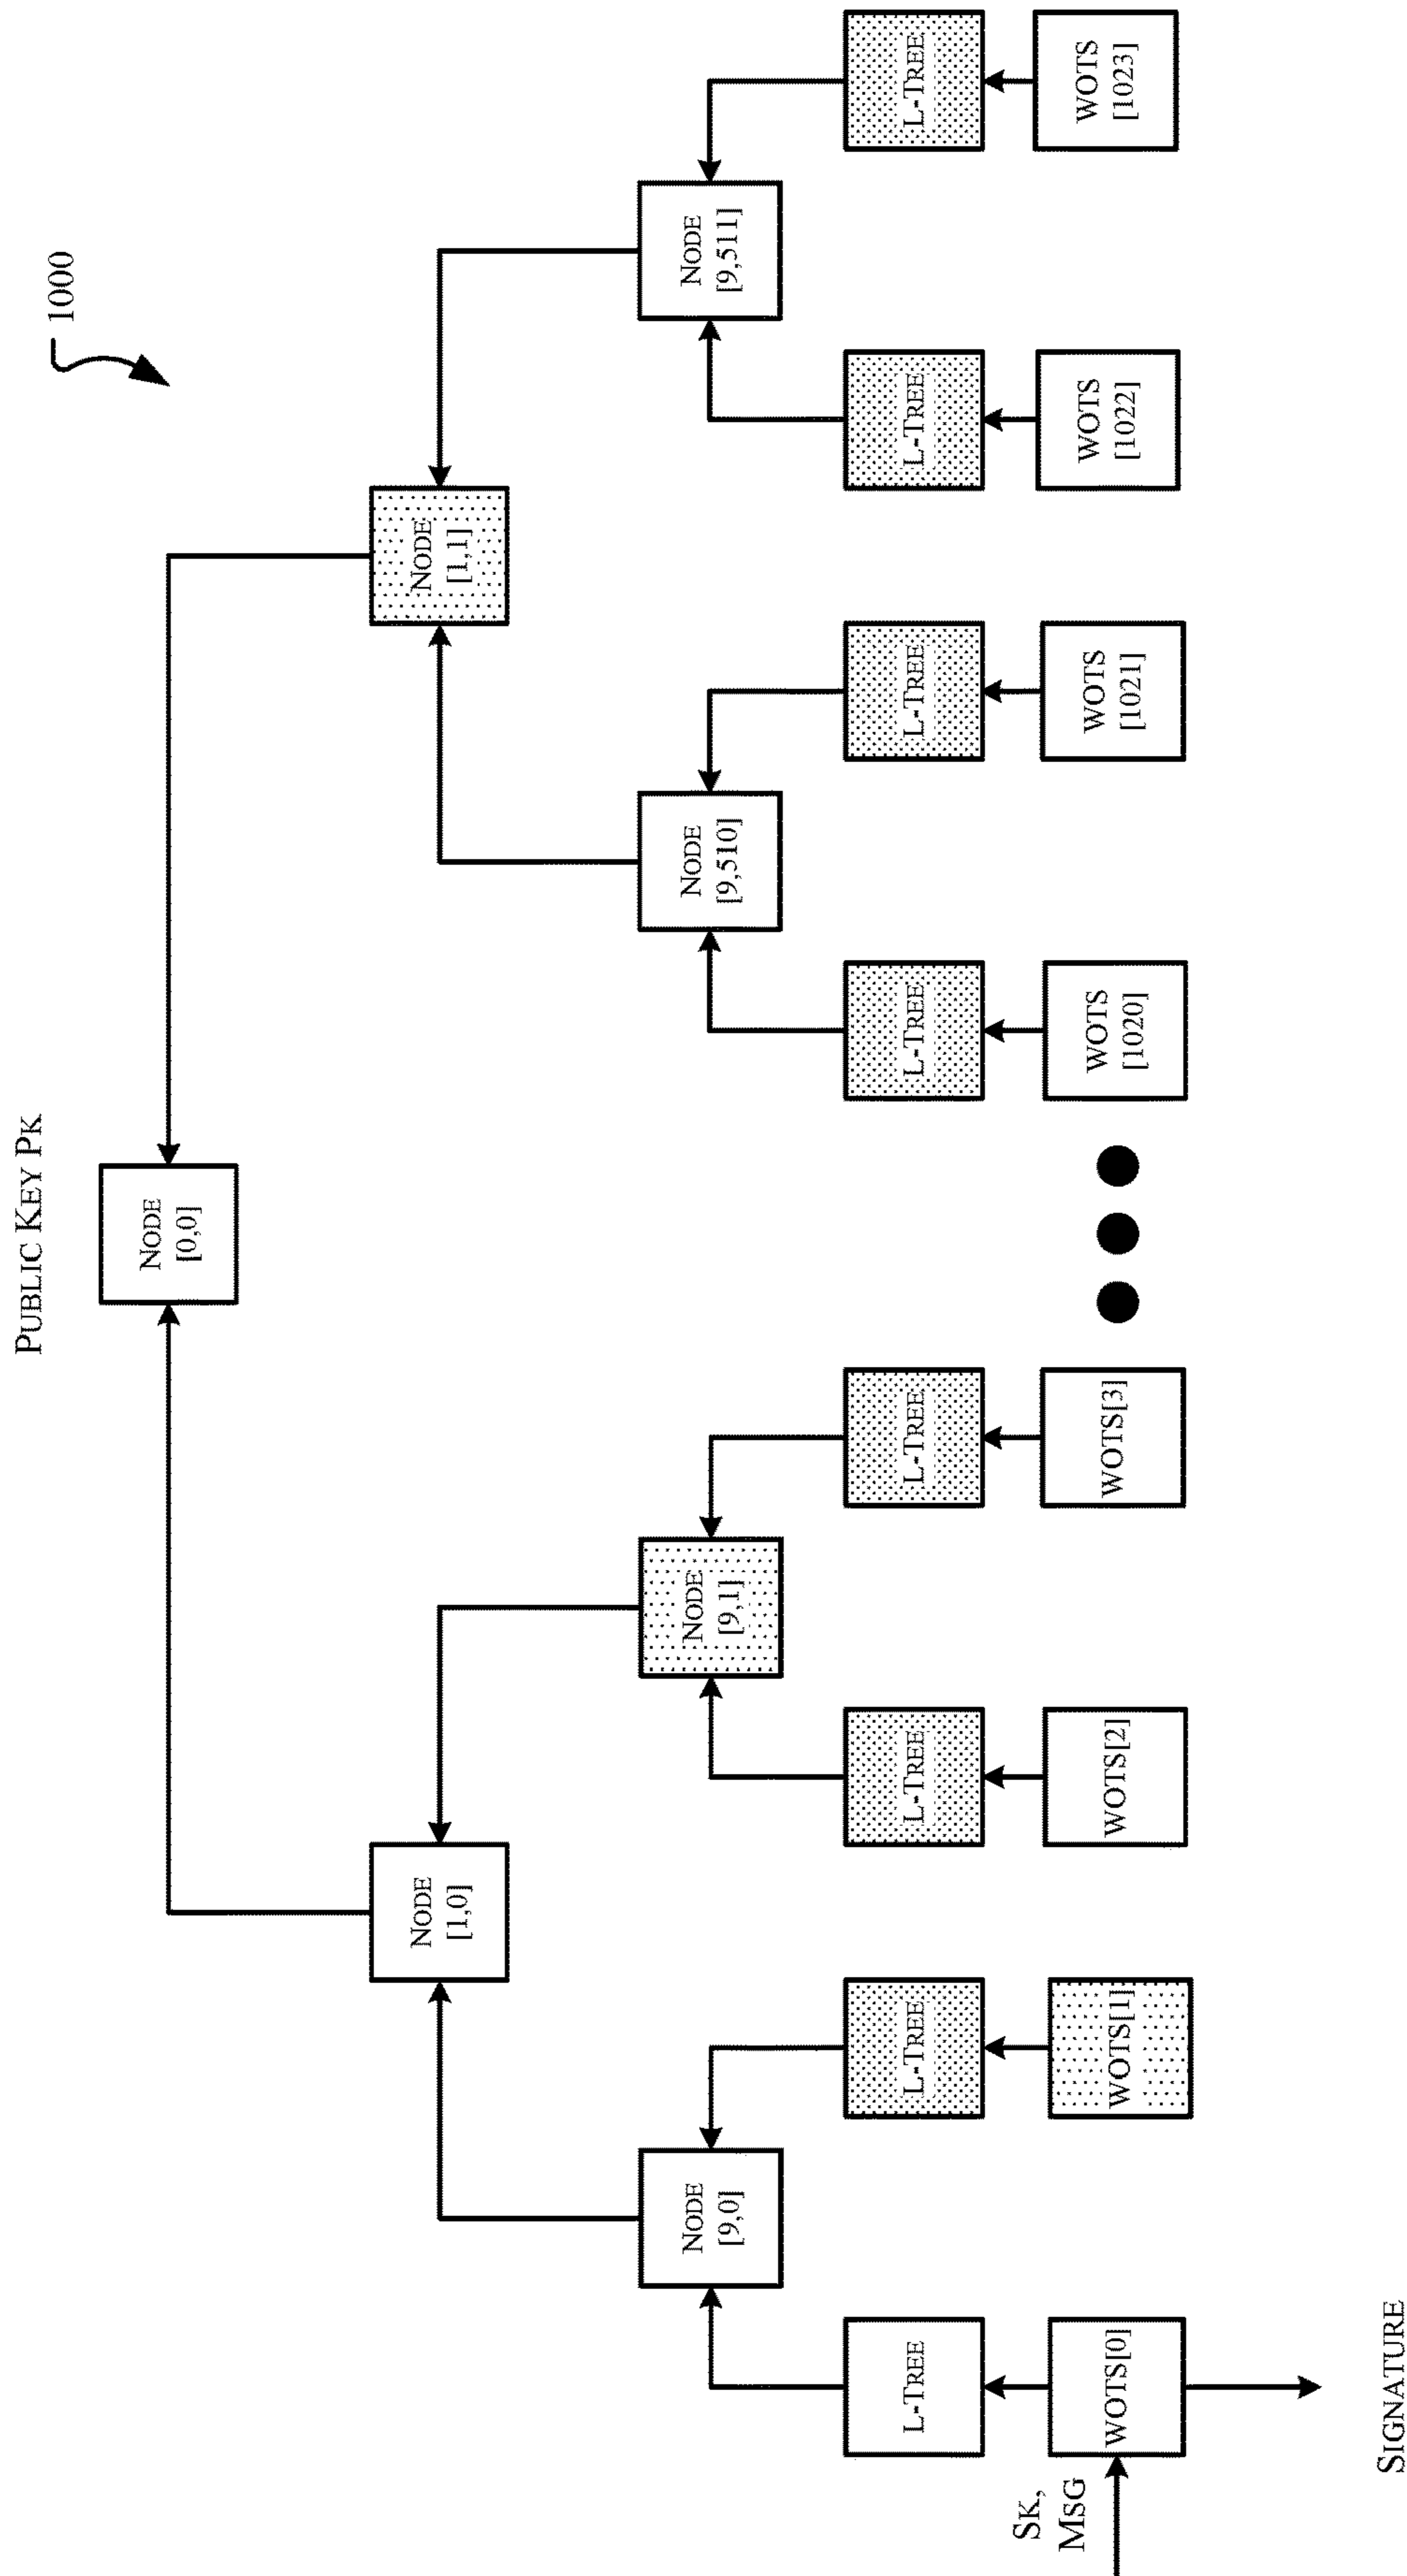


FIG. 10

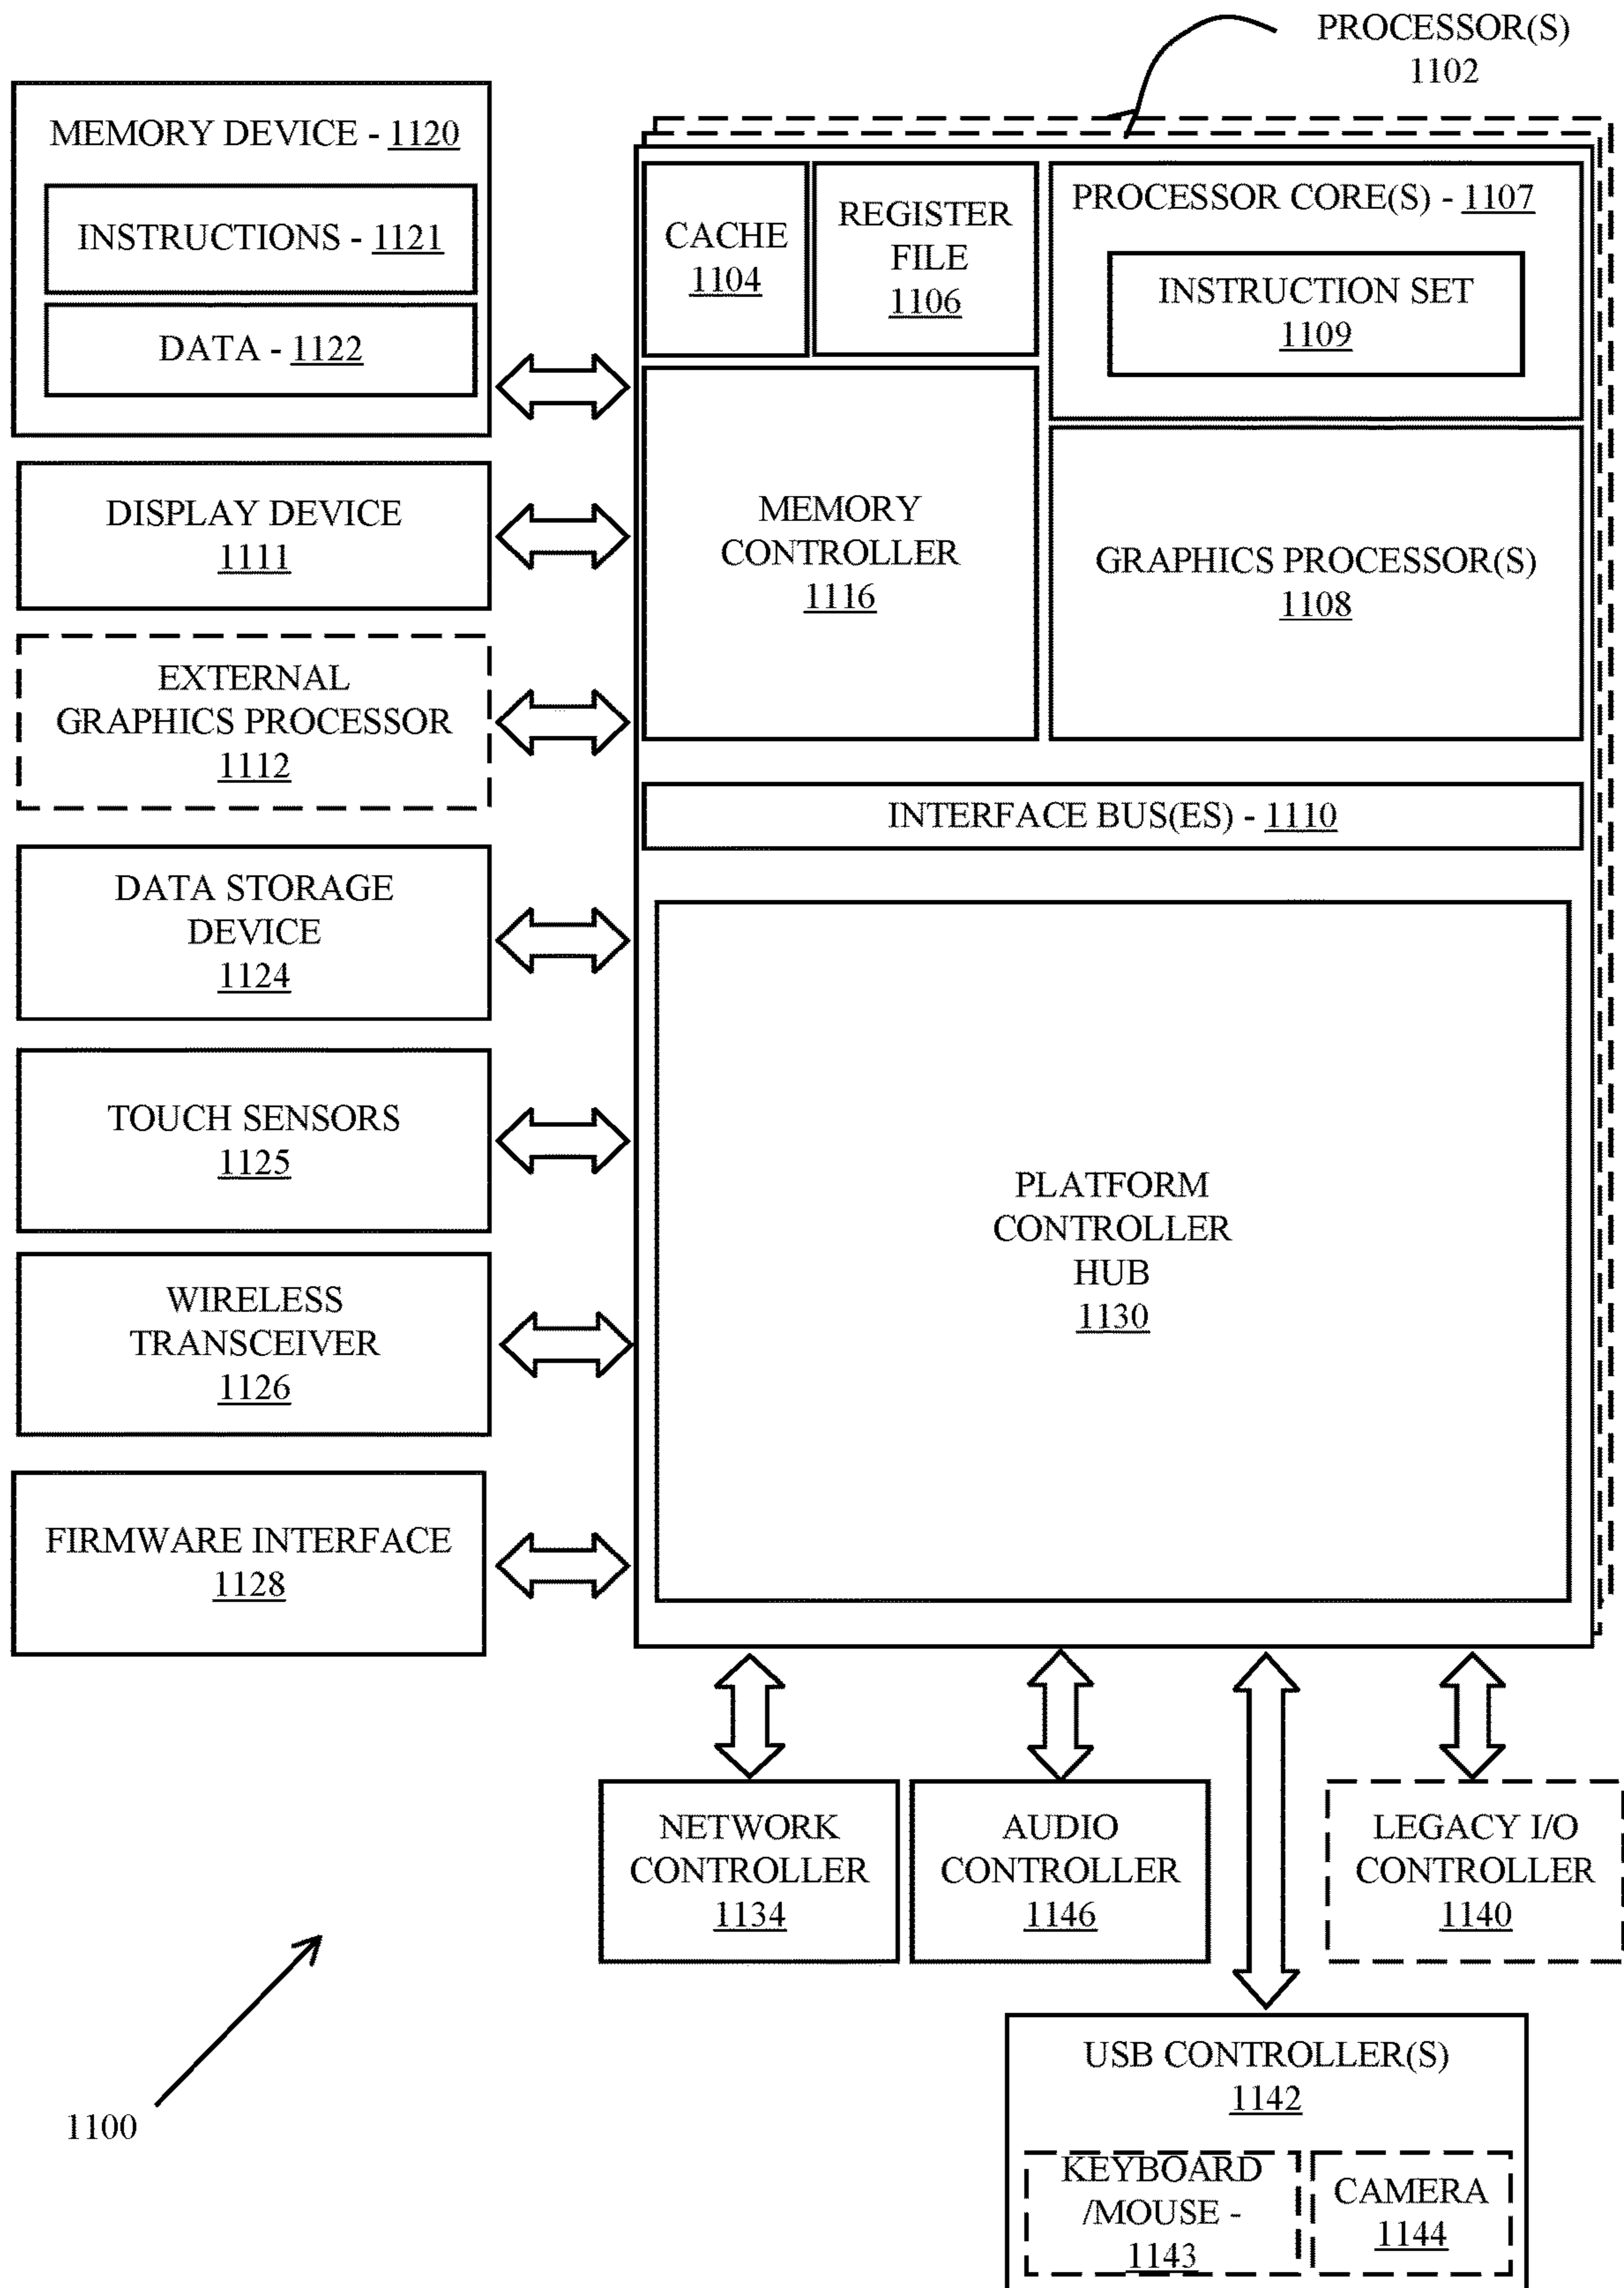


FIG. 11

**ACCELERATORS FOR POST-QUANTUM
CRYPTOGRAPHY SECURE HASH-BASED
SIGNING AND VERIFICATION**

CLAIM TO PRIORITY

This Application is a continuation of and claims the benefit of and priority to U.S. application Ser. No. 16/455,908, entitled ACCELERATORS FOR POST-QUANTUM CRYPTOGRAPHY SECURE HASH-BASED SIGNING AND VERIFICATION, by Vikram Suresh, et al., filed Jun. 28, 2019, the entire contents of which are incorporated herein by reference.

BACKGROUND

Subject matter described herein relates generally to the field of computer security and more particularly to accelerators for post-quantum cryptography secure hash-based signing and verification.

Existing public-key digital signature algorithms such as Rivest-Shamir-Adleman (RSA) and Elliptic Curve Digital Signature Algorithm (ECDSA) are anticipated not to be secure against brute-force attacks based on algorithms such as Shor's algorithm using quantum computers. As a result, there are efforts underway in the cryptography research community and in various standards bodies to define new standards for algorithms that are secure against quantum computers.

Accordingly, techniques to accelerate signature and verification schemes such as XMSS and LMS may find utility, e.g., in computer-based communication systems and methods.

BRIEF DESCRIPTION OF THE DRAWINGS

The detailed description is described with reference to the accompanying figures.

FIGS. 1A and 1B are schematic illustrations of a one-time hash-based signatures scheme and a multi-time hash-based signatures scheme, respectively.

FIGS. 2A-2B are schematic illustrations of a one-time signature scheme and a multi-time signature scheme, respectively.

FIG. 3 is a schematic illustration of a signing device and a verifying device, in accordance with some examples.

FIG. 4A is a schematic illustration of a Merkle tree structure, in accordance with some examples.

FIG. 4B is a schematic illustration of a Merkle tree structure, in accordance with some examples.

FIG. 5 is a schematic illustration of a compute blocks in an architecture to implement a signature algorithm, in accordance with some examples.

FIG. 6A is a schematic illustration of a compute blocks in an architecture to implement signature generation in a signature algorithm, in accordance with some examples.

FIG. 6B is a schematic illustration of a compute blocks in an architecture to implement signature verification in a verification algorithm, in accordance with some examples.

FIG. 7A is a schematic illustration of compute blocks in an architecture to implement a HASH algorithm, in accordance with some examples.

FIG. 7B is a schematic illustration of compute blocks in an architecture to implement a HASH algorithm with a pseudo-random function (PRF) precomputed opcode/key, in accordance with some examples.

FIG. 7C is a schematic illustration of compute blocks in an architecture to implement a HASH algorithm with a pseudo-random function (PRF) precomputed message, in accordance with some examples.

FIG. 7D is a schematic illustration of compute blocks in an architecture to implement a HASH algorithm with a function-F opcode precompute, in accordance with some examples.

FIG. 8 is a schematic illustration of an L-tree structure, in accordance with some examples.

FIG. 9A is a schematic illustration of a compute blocks in an architecture to implement a RAND_HASH algorithm, in accordance with some examples.

FIG. 9B is a schematic illustration of a compute blocks in an architecture to implement a RAND_HASH algorithm with pre-computations, in accordance with some examples.

FIG. 10 is a schematic illustration of an L-tree computation for an authentication path generation during signing, in accordance with some examples.

FIG. 11 is a schematic illustration of a computing architecture which may be adapted to implement adversarial training of neural networks using information about activation path differentials in accordance with some examples.

DETAILED DESCRIPTION

Described herein are exemplary systems and methods to implement accelerators for post-quantum cryptography secure hash-based signature algorithms. In the following description, numerous specific details are set forth to provide a thorough understanding of various examples. However, it will be understood by those skilled in the art that the various examples may be practiced without the specific details. In other instances, well-known methods, procedures, components, and circuits have not been illustrated or described in detail so as not to obscure the examples.

As described briefly above, existing public-key digital signature algorithms such as Rivest-Shamir-Adleman (RSA) and Elliptic Curve Digital Signature Algorithm (ECDSA) are anticipated not to be secure against brute-force attacks based on algorithms such as Shor's algorithm using quantum computers. The eXtended Merkle signature scheme (XMSS) and/or an eXtended Merkle many time signature scheme (XMSS-MT) are hash-based signature schemes that can protect against attacks by quantum computers. As used herein, the term XMSS shall refer to both the XMSS scheme and the XMSS-MT scheme.

An XMSS signature process implements a hash-based signature scheme using a one-time signature scheme such as a Winternitz one-time signature (WOTS) or a derivative thereof (e.g., WOTS+) in combination with a secure hash algorithm (SHA) such as SHA2-256 as the primary underlying hash function. In some examples the XMSS signature/verification scheme may also use one or more of SHA2-512, SHA3-SHAKE-256 or SHA3-SHAKE-512 as secure hash functions. XMSS-specific hash functions include a Pseudo-Random Function (PRF), a chain hash (F), a tree hash (H) and message hash function (H_{msg}). As used herein, the term WOTS shall refer to the WOTS signature scheme and or a derivative scheme such as WOTS+.

The Leighton/Micali signature (LMS) scheme is another hash-based signature scheme that uses Leighton/Micali one-time signatures (LM-OTS) as the one-time signature building block. LMS signatures are based on a SHA2-256 hash function.

An XMSS signature process comprises three major operations. The first major operation receives an input message

(M) and a private key (sk) and utilizes a one-time signature algorithm (e.g., WOTS+) to generate a message representative (M') that encodes a public key (pk). In a 128-bit post quantum security implementation the input message M is subjected to a hash function and then divided into 67 message components (n bytes each), each of which are subjected to a hash chain function to generate the a corresponding 67 components of the digital signature. Each chain function invokes a series of underlying secure hash algorithms (SHA).

The second major operation is an L-Tree computation, which combines WOTS+ (or WOTS) public key components (n-bytes each) and produces a single n-byte value. For example, in the 128-bit post-quantum security there are 67 public key components, each of which invokes an underlying secure hash algorithm (SHA) that is performed on an input block.

The third major operation is a tree-hash operation, which constructs a Merkle tree. In an XMSS verification, an authentication path that is provided as part of the signature and the output of L-tree operation is processed by a tree-hash operation to generate the root node of the Merkle tree, which should correspond to the XMSS public key. For XMSS verification with 128-bit post-quantum security, traversing the Merkle tree comprises executing secure hash operations. In an XMSS verification, the output of the Tree-hash operation is compared with the known public key. If they match then the signature is accepted. By contrast, if they do not match then the signature is rejected.

The XMSS signature process is computationally expensive. An XMSS signature process invokes hundreds, or even thousands, of cycles of hash computations. Subject matter described herein addresses these and other issues by providing systems and methods to implement accelerators for post-quantum cryptography secure XMSS and LMS hash-based signing and verification.

Post-Quantum Cryptography Overview

Post-Quantum Cryptography (also referred to as “quantum-proof”, “quantum-safe”, “quantum-resistant”, or simply “PQC”) takes a futuristic and realistic approach to cryptography. It prepares those responsible for cryptography as well as end-users to know the cryptography is outdated; rather, it needs to evolve to be able to successfully address the evolving computing devices into quantum computing and post-quantum computing.

It is well-understood that cryptography allows for protection of data that is communicated online between individuals and entities and stored using various networks. This communication of data can range from sending and receiving of emails, purchasing of goods or services online, accessing banking or other personal information using websites, etc.

Conventional cryptography and its typical factoring and calculating of difficult mathematical scenarios may not matter when dealing with quantum computing. These mathematical problems, such as discrete logarithm, integer factorization, and elliptic-curve discrete logarithm, etc., are not capable of withstanding an attack from a powerful quantum computer. Shor, P. W. “Algorithms for quantum computation: discrete logarithms and factoring”. Proceedings 35th Annual Symposium on Foundations of Computer Science. IEEE Comput. Soc. Press. doi:10.1109/sfcs.1994.365700. ISBN 0818665807. Although any post-quantum cryptography could be built on the current cryptography, the novel

approach would need to be intelligent, fast, and precise enough to resist and defeat any attacks by quantum computers.

Today's PQC is mostly focused on the following approaches: 1) hash-based cryptography based on Merkle's hash tree public-key signature system of 1979, which is built upon a one-message-signature idea of Lamport and Diffie; 2) code-based cryptography, such as McEliece's hidden-Goppa-code public-key encryption system; 3) lattice-based cryptography based on Hoffstein-Pipher-Silverman public-key-encryption system of 1998; 4) multivariate-quadratic equations cryptography based on Patarin's Hidden Field Equation (HFE) public-key-signature system of 1996 that is further based on the Matumoto-Imai proposal; 5) supersingular elliptical curve isogeny cryptography that relies on supersingular elliptic curves and supersingular isogeny graphs; and 6) symmetric key quantum resistance, such as HBS.

FIGS. 1A and 1B illustrate a one-time hash-based signatures scheme and a multi-time hash-based signatures scheme, respectively. As aforesaid, hash-based cryptography is based on cryptographic systems like Lamport signatures, Merkle Signatures, extended Merkle signature scheme (XMSS), SPHINCS scheme, SPHINCS+ scheme, etc. With the advent of quantum computing and in anticipation of its growth, there have been concerns about various challenges that quantum computing could pose and what could be done to counter such challenges using the area of cryptography.

One area that is being explored to counter quantum computing challenges is hash-based signatures (HBS) since these schemes have been around for a long while and possess the necessary basic ingredients, such as relying on symmetric cryptography building blocks (e.g., hash functions), to counter the quantum counting and post-quantum computing challenges. HBS schemes are regarded as fast signature algorithms working with fast platform secured-boot, which is regarded as the most resistant to quantum attacks.

For example, as illustrated with respect to FIG. 1A, a scheme of HBS is shown that uses Merkle trees along with one-time signature (OTS) scheme **100**, such as using a private key to sign a message and a corresponding public key to verify the OTS message, where a private key only signs a single message.

Similarly, as illustrated with respect to FIG. 1B, another HBS scheme is shown, where this one relates to multi-time signatures (MTS) scheme **150**, where a private key can sign multiple messages.

FIGS. 2A and 2B illustrate a one-time signature scheme and a multi-time signature scheme, respectively. Continuing with HBS-based OTS scheme **100** of FIG. 1A and MTS scheme **150** of FIG. 1B, FIG. 2A illustrates Winternitz OTS (WOTS) scheme **200**, which was offered by Robert Winternitz of Stanford Mathematics Department, while FIG. 2B illustrates XMSS MTS scheme **250**, respectively.

For example, WOTS scheme **200** of FIG. 2A provides for hashing and parsing of messages into M, with 67 integers between [0, 1, 2, . . . , 15], such as private key, sk, **205**, signature, s, **210**, and public key, pk, **215**, with each having 67 components of 32 bytes each.

Now, for example, FIG. 2B illustrates XMSS MTS scheme **250** that allows for a combination of WOTS scheme **200** of FIG. 2A and XMSS scheme **255** having XMSS Merkle tree **270**. As discussed previously with respect to FIG. 2A, WOTS scheme **200** is based on a one-time public key, pk, **215**, having 67 components of 32 bytes each, that is then put through L-Tree compression algorithm **260** to

offer WOTS compressed pk **267** to take a place in the XMSS Merkle tree **270** of XMSS scheme **255**. It is contemplated that XMSS signature verification may include computing WOTS verification and checking to determine whether a reconstructed root node matches the XMSS public key, such as root node=XMSS public key.

Accelerators for Post-Quantum Cryptography

FIG. **3** is a schematic illustration of a high-level architecture of a secure environment **300** that includes a first device **310** and a second device **350**, in accordance with some examples. Referring to FIG. **3**, each of the first device **310** and the second device **350** may be embodied as any type of computing device capable of performing the functions described herein. For example, in some embodiments, each of the first device **310** and the second device **350** may be embodied as a laptop computer, tablet computer, notebook, netbook, Ultrabook™, a smartphone, cellular phone, wearable computing device, personal digital assistant, mobile Internet device, desktop computer, router, server, workstation, and/or any other computing/communication device.

First device **310** includes one or more processor(s) **320** and a memory **322** to store a private key **324**. The processor(s) **320** may be embodied as any type of processor capable of performing the functions described herein. For example, the processor(s) **320** may be embodied as a single or multi-core processor(s), digital signal processor, microcontroller, or other processor or processing/controlling circuit. Similarly, the memory **322** may be embodied as any type of volatile or non-volatile memory or data storage capable of performing the functions described herein. In operation, the memory **322** may store various data and software used during operation of the first device **310** such as operating systems, applications, programs, libraries, and drivers. The memory **322** is communicatively coupled to the processor(s) **320**. In some examples the private key **324** may reside in a secure memory that may be part memory **322** or may be separate from memory **322**.

First device **310** further comprises authentication logic **330** which includes memory **332**, signature logic, and verification logic **336**. Hash logic **332** is configured to hash (i.e., to apply a hash function to) a message (M) to generate a hash value (m') of the message M. Hash functions may include, but are not limited to, a secure hash function, e.g., secure hash algorithms SHA2-256 and/or SHA3-256, etc. SHA2-256 may comply and/or be compatible with Federal Information Processing Standards (FIPS) Publication 180-4, titled: "Secure Hash Standard (SHS)", published by National Institute of Standards and Technology (NIST) in March 2012, and/or later and/or related versions of this standard. SHA3-256 may comply and/or be compatible with FIPS Publication 202, titled: "SHA-3 Standard: Permutation-Based Hash and Extendable-Output Functions", published by NIST in August 2015, and/or later and/or related versions of this standard.

Signature logic **332** may be configured to generate a signature to be transmitted, i.e., a transmitted signature. In instances in which the first device **310** is the signing device, the transmitted signature may include a number, L, of transmitted signature elements with each transmitted signature element corresponding to a respective message element. For example, for each message element, m_i , signature logic **332** may be configured to perform a selected signature operation on each private key element, sk_i , of the private key, sk , a respective number of times related to a value of each message element, m_i , included in the message representative

m' . For example, signature logic **332** may be configured to apply a selected hash function to a corresponding private key element, sk_i , m_i times. In another example, signature logic **332** may be configured to apply a selected chain function (that contains a hash function) to a corresponding private key element, sk_i , m_i times. The selected signature operations may, thus, correspond to a selected hash-based signature scheme.

As described above, hash-based signature schemes may include, but are not limited to, a Winternitz (W) one time signature (OTS) scheme, an enhanced Winternitz OTS scheme (e.g., WOTS+), a Merkle many time signature scheme, an extended Merkle signature scheme (XMSS) and/or an extended Merkle multiple tree signature scheme (XMSS-MT), etc. Hash functions may include, but are not limited to SHA2-256 and/or SHA3-256, etc. For example, XMSS and/or XMSS-MT may comply or be compatible with one or more Internet Engineering Task Force (IETF®) informational draft Internet notes, e.g., "XMSS: Extended Hash-Based Signatures, released May, 2018, by the Internet Research Task Force (IRTF), Crypto Forum Research Group.

A WOTS signature algorithm may be used to generate a signature and to verify a received signature utilizing a hash function. WOTS is further configured to use the private key and, thus, each private key element, sk_i , one time. For example, WOTS may be configured to apply a hash function to each private key element, m_i , or $N-m_i$ times to generate a signature and to apply the hash function to each received message element $N-m_i$, or m_i , times to generate a corresponding verification signature element. The Merkle many time signature scheme is a hash-based signature scheme that utilizes an OTS and may use a public key more than one time. For example, the Merkle signature scheme may utilize Winternitz OTS as the one-time signature scheme. WOTS+ is configured to utilize a family of hash functions and a chain function.

XMSS, WOTS+ and XMSS-MT are examples of hash-based signature schemes that utilize chain functions. Each chain function is configured to encapsulate a number of calls to a hash function and may further perform additional operations. In some examples, the number of calls to the hash function included in the chain function may be fixed. Chain functions may improve security of an associated hash-based signature scheme.

Cryptography logic **340** is configured to perform various cryptographic and/or security functions on behalf of the signing device **310**. In some embodiments, the cryptography logic **340** may be embodied as a cryptographic engine, an independent security co-processor of the signing device **310**, a cryptographic accelerator incorporated into the processor(s) **320**, or a standalone software/firmware. In some embodiments, the cryptography logic **340** may generate and/or utilize various cryptographic keys (e.g., symmetric/asymmetric cryptographic keys) to facilitate encryption, decryption, signing, and/or signature verification. Additionally, in some embodiments, the cryptography logic **340** may facilitate to establish a secure connection with remote devices over communication link. It should further be appreciated that, in some embodiments, the cryptography module **340** and/or another module of the first device **310** may establish a trusted execution environment or secure enclave within which a portion of the data described herein may be stored and/or a number of the functions described herein may be performed.

After the signature is generated as described above, the message, M, and signature may then be sent by first device

310, e.g., via communication logic 342, to second device 350 via network communication link 390. In an embodiment, the message, M, may not be encrypted prior to transmission. In another embodiment, the message, M, may be encrypted prior to transmission. For example, the message, M, may be encrypted by cryptography logic 340 to produce an encrypted message.

Second device 350 may also include one or more processors 360 and a memory 362 to store a public key 364. As described above, the processor(s) 360 may be embodied as any type of processor capable of performing the functions described herein. For example, the processor(s) 360 may be embodied as a single or multi-core processor(s), digital signal processor, microcontroller, or other processor or processing/controlling circuit. Similarly, the memory 362 may be embodied as any type of volatile or non-volatile memory or data storage capable of performing the functions described herein. In operation, the memory 362 may store various data and software used during operation of the second device 350 such as operating systems, applications, programs, libraries, and drivers. The memory 362 is communicatively coupled to the processor(s) 360.

In some examples the public key 364 may be provided to second device 350 in a previous exchange. The public key, pk, is configured to contain a number L of public key elements, i.e., $pk = [pk_1, \dots, pk_L]$. The public key 364 may be stored, for example, to memory 362.

Second device 350 further comprises authentication logic 370 which includes hash logic 372, signature logic, and verification logic 376. As described above, hash logic 372 is configured to hash (i.e., to apply a hash function to) a message (M) to generate a hash message (m'). Hash functions may include, but are not limited to, a secure hash function, e.g., secure hash algorithms SHA2-256 and/or SHA3-256, etc. SHA2-256 may comply and/or be compatible with Federal Information Processing Standards (FIPS) Publication 180-4, titled: "Secure Hash Standard (SHS)", published by National Institute of Standards and Technology (NIST) in March 2012, and/or later and/or related versions of this standard. SHA3-256 may comply and/or be compatible with FIPS Publication 202, titled: "SHA-3 Standard: Permutation-Based Hash and Extendable-Output Functions", published by NIST in August 2015, and/or later and/or related versions of this standard.

In instances in which the second device is the verifying device, authentication logic 370 is configured to generate a verification signature based, at least in part, on the signature received from the first device and based, at least in part, on the received message representative (m'). For example, authentication logic 370 may be configured to perform the same signature operations, i.e., apply the same hash function or chain function as applied by hash logic 332 of authentication logic 330, to each received message element a number, $N - m_i$, (or m_i), times to yield a verification message element. Whether a verification signature, i.e., each of the L verification message elements, corresponds to a corresponding public key element, pk_i , may then be determined. For example, verification logic 370 may be configured to compare each verification message element to the corresponding public key element, pk_i . If each of the verification message element matches the corresponding public key element, pk_i , then the verification corresponds to success. In other words, if all of the verification message elements match the public key elements, pk_1, \dots, pk_L , then the verification corresponds to success. If any verification message element does not match the corresponding public key element, pk_i , then the verification corresponds to failure.

As described in greater detail below, in some examples the authentication logic 330 of the first device 310 includes one or more accelerators 338 that cooperate with the hash logic 332, signature logic 334 and/or verification logic 336 to accelerate authentication operations. Similarly, in some examples the authentication logic 370 of the second device 310 includes one or more accelerators 378 that cooperate with the hash logic 372, signature logic 374 and/or verification logic 376 to accelerate authentication operations. Examples of accelerators are described in the following paragraphs and with reference to the accompanying drawings.

The various modules of the environment 300 may be embodied as hardware, software, firmware, or a combination thereof. For example, the various modules, logic, and other components of the environment 300 may form a portion of, or otherwise be established by, the processor(s) 320 of first device 310 or processor(s) 360 of second device 350, or other hardware components of the devices. As such, in some embodiments, one or more of the modules of the environment 300 may be embodied as circuitry or collection of electrical devices (e.g., an authentication circuitry, a cryptography circuitry, a communication circuitry, a signature circuitry, and/or a verification circuitry). Additionally, in some embodiments, one or more of the illustrative modules may form a portion of another module and/or one or more of the illustrative modules may be independent of one another.

FIG. 4A is a schematic illustration of a Merkle tree structure illustrating signing operations, in accordance with some examples. Referring to FIG. 4A, an XMSS signing operation requires the construction of a Merkle tree 400A using the local public key from each leaf WOTS node 410 to generate a global public key (PK) 420. In some examples the authentication path and the root node value can be computed off-line such that these operations do not limit performance. Each WOTS node 410 has a unique secret key, "sk" which is used to sign a message only once. The XMSS signature consists of a signature generated for the input message and an authentication path of intermediate tree nodes to construct the root of the Merkle tree.

FIG. 4B is a schematic illustration of a Merkle tree structure 400B during verification, in accordance with some examples. During verification, the input message and signature are used to compute the local public key 420B of the WOTS node, which is further used to compute the tree root value using the authentication path. A successful verification will match the computed tree root value to the public key PK shared by the signing entity. The WOTS and L-Tree operations constitute a significant portion of XMSS sign/verify latency respectively, thus defining the overall performance of the authentication system. Described herein are various pre-computation techniques which may be implemented to speed-up WOTS and L-Tree operations, thereby improving XMSS performance. The techniques are applicable to the other hash options and scale well for both software and hardware implementations.

FIG. 5 is a schematic illustration of a compute blocks in an architecture 500 to implement a signature algorithm, in accordance with some examples. Referring to FIG. 5, the WOTS+ operation involves 67 parallel chains of 16 SHA2-256 HASH functions, each with the secret key $sk[66:0]$ as input. Each HASH operation in the chain consists of 2 pseudo-random functions (PRF) using SHA2-256 to generate a bitmask and a key. The bitmask is XOR-ed with the previous hash and concatenated with the key as input message to a 3rd SHA2-256 hash operation. The 67x32-byte

WOTS public key $pk[66:0]$ is generated by hashing secret key sk across the 67 hash chains.

FIG. 6A is a schematic illustration of a compute blocks in an architecture 600A to implement signature generation in a signature algorithm, in accordance with some examples. As illustrated in FIG. 6A, for message signing, the input message is hashed and pre-processed to compute a 67×4 -bit value, which is used as an index to choose an intermediate hash value in each operation of the chain function.

FIG. 6B is a schematic illustration of a compute blocks in an architecture 600B to implement signature verification in a verification algorithm, in accordance with some examples. Referring to FIG. 6B, during verification, the message is again hashed to compute the signature indices and compute the remaining HASH operations in each chain to compute the WOTS public key pk . This value and the authentication path are used to compute the root of the Merkle tree and compare with the shared public key PK to verify the message.

FIG. 7A is a schematic illustration of compute blocks in an architecture 700A to implement a HASH algorithm, in accordance with some examples. By way of background, SHA256 processes 64 bytes of input at a time. These 64 bytes of input are divided into 16 values of 32-bits each, denoted as W_1, W_2, \dots, W_{16} . From these 16 values, SHA256 derives other 48 values of 32-bits each, denoted as $W_{17}, W_{18}, \dots, W_{64}$. The SHA256 function has 64 rounds and it uses one W_i for each one of the rounds, $1 \leq i \leq 64$. This means that if the first 16 W s are known, it is possible to compute all 64 rounds. If fewer than the first 16 W s are known, say 15, then only 15 rounds can be pre-computed.

Referring to FIG. 7A, the inputs to the three SHA2-256 function calls in the HASH operation (i.e., 2 PRFs and one function F) consists of a 128-byte message obtained by concatenating 32-bytes of OPCODE, KEY, MESSAGE and PAD. In subject matter described herein, common features of the input message may leveraged in different pre-compute techniques to speed-up the hashing operation in WOTS. Each HASH operation works on 3×128 -bytes of message, thus requiring 384 hashing rounds for SHA2-256 which consumes message at a rate of 1 byte/round. The overall WOTS public key generation requires $67 \times 16 \times 384 = 411,648$ hash rounds. The message dependent sign/verify operations will require on average half the number of rounds equating to 205,824.

FIG. 7B is a schematic illustration of compute blocks in an architecture 700B to implement a HASH algorithm with a pseudo-random function (PRF) precomputed opcode/key, in accordance with some examples. Referring to FIG. 7B, in some examples the opcode for PRF function across the entire WOTS operation may be set to a constant value 32-byte value fixed to, for example, '3'. Further, the KEY part of the message in PRF is a pseudo-randomly generated non-secret 32-byte value identified as the SEED, which is common for all PRF functions in each WOTS operation. Since SHA2-256 consumes messages in blocks of 64-bytes and operates for 64 rounds per block, the intermediate hash for the first 64-bytes of PRF functions (i.e., the OPCODE and the SEED) can be computed once and used across all 1072 HASH operations in WOTS to accelerate computation of the HASH algorithm. In some examples, this modification reduces the number of iterative rounds for PRF from 128 to 64, thus enabling WOTS key generation in 274,432 hash cycles and signing/verification in approximately 137,216 cycles, resulting in improvement in WOTS performance.

FIG. 7C is a schematic illustration of compute blocks in an architecture 700C to implement a HASH algorithm with

a pseudo-random function (PRF) precomputed message, in accordance with some examples. Referring to FIG. 7C, in some examples the message part of the input to PRF consists of the 32-byte OTS address, divided into 6×4 -byte and 1×8 -byte address fields. Since the first 20-bytes of the OTS address is a constant for each WOTS operation, the hash value for these can be pre-computed to eliminate 5-rounds in each PRF function. Further, the remaining 12-bytes of OTS address differ by only a single bit between key and bitmask generation. This may be leveraged by computing the intermediate hash for the remaining 12-bytes once during key generation and then re-using the computed value during bitmask generation to eliminate an additional 3 rounds from bitmask PRF. This technique reduces the number of iterative rounds for PRF from 128 to 59/56 respectively, enabling WOTS key generation in approximately 260,496 hash cycles and signature/verification in approximately 130,248 cycles, resulting in improvement in WOTS performance.

FIG. 7D is a schematic illustration of compute blocks in an architecture 700D to implement a HASH algorithm with a function- F opcode precompute, in accordance with some examples. Referring to FIG. 7D, in some examples the hash function "F" has a 128-byte message comprising of a 32-byte opcode, 32-byte PRF generated key, 32-byte XOR of previous hash and the bitmask and a 32-byte padding. The 32-byte opcode is a constant and set to the value '0' and hence can be pre-computed and shared across all HASH operations in WOTS. This reduces the number of iterative rounds for function- F from 128 to 56, enabling WOTS key generation in approximately 251,920 hash cycles and signature/verification in approximately 125,960 cycles, resulting in improvement in WOTS performance.

These pre-computation optimizations can be extended to other XMSS hash functions as well and are scalable across both software and hardware implementations.

L-Tree Pre-Hash Improvements

Another contributor to XMSS latency is the L-Tree operation. L-Tree generates a binary tree to compress the 67×32 -byte WOTS public key into a single 32-byte leaf node value using a keyed-hash function called RAND_HASH.

FIG. 8 is a schematic illustration of an L-tree structure 800, in accordance with some examples. Referring to FIG. 8, in some examples the RAND_HASH function internally generates a local key ($pk[n]$) and two bitmask values using the PRF function. The key and bitmasks xor-ed with left/right child nodes are concatenated and hashed to compute a new node value of L-Tree. All hash operations in L-Tree use SHA256 in the default recommended mode, requiring 3×128 -rounds for PRF and 192-rounds for function-H operations.

FIG. 9A is a schematic illustration of a compute blocks in an architecture 900A to implement a RAND_HASH algorithm. Referring to FIG. 9A, in some examples the OPCODE and SEED for the PRF function are constant for an XMSS operation and can be pre-computed in some instances.

FIG. 9B is a schematic illustration of a compute blocks in an architecture 900B to implement a RAND_HASH algorithm with pre-computations, in accordance with some examples. As indicated in FIG. 9B, the 64-bytes of OPCODE||SEED can be pre-hashed to reduce the latency of all three PRF functions to 64-rounds. Further, 24-bytes of the address (ADDR*) field are constant for an L-Tree operation and can be pre-hashed. The remaining 2×4 -byte fields of the address can be hashed once during key genera-

11

tion and re-used for bitmask generation by incrementing the pre-hashed digest value. This reduces the number of rounds in key generation to 58 and bitmask generation to 56. Similar to PRF, the OPCODE input to function-H is also a constant. The OPCODE value is pre-hashed and re-used during all function-F call, reducing the number of rounds to 184. These pre-hash techniques can reduce overall L-Tree latency significantly.

FIG. 10 is a schematic illustration of an L-tree computation 1000 for an authentication path generation during

12

Table 2 shows the core part of the LM-OTS signature generation (on the left) and signature verification algorithms that process each one of these chunks. These algorithms are similar to WOTS algorithms. At first, a message representative Q is generated from the message. Then a for-loop iterates over each one of the $p=67$ chunks. For each one of them, the hash function is called a number of times that depend on the coefficient a extracted from the message representative Q. In the signature verification algorithm, there is a final step that compresses all $p=67$ chunks into a single 32-bytes value Kc.

TABLE 2

1. $Q = H(\mathbf{I} \parallel \mathbf{q} \parallel D_MSG \parallel C \parallel \text{message})$	1. $Q = H(\mathbf{I} \parallel \mathbf{q} \parallel D_MSG \parallel C \parallel \text{message})$
2. for ($i = 0; i < p; i = i + 1$) {	2. for ($i = 0; i < p; i = i + 1$) {
3. $a = \text{coef}(Q \parallel \text{Cksm}(Q), i, w)$	3. $a = \text{coef}(Q \parallel \text{Cksm}(Q), i, w)$
4. $\text{tmp} = x[i]$	4. $\text{tmp} = y[i]$
5. for ($j = 0; j < a; j = j + 1$) {	5. for ($j = a; j < 2^w - 1; j = j + 1$) {
6. $\text{tmp} = H(\mathbf{I} \parallel \mathbf{q} \parallel \mathbf{i} \parallel \mathbf{j} \parallel \text{tmp})$	6. $\text{tmp} = H(\mathbf{I} \parallel \mathbf{q} \parallel \mathbf{i} \parallel \mathbf{j} \parallel \text{tmp})$
7. }	7. }
8. $y[i] = \text{tmp}$	8. $z[i] = \text{tmp}$
9. }	9. }
	10. $Kc = H(\mathbf{I} \parallel \mathbf{q} \parallel D_PBLC \parallel z[0] \parallel z[1] \parallel \dots \parallel z[p-1])$

signing, in accordance with some examples. The L-Tree pre-hash pre-computations have a significant impact during XMSS signing, since an L-Tree operation is performed during each signature generation to compute the authentication path. Since the SEED and ADDR fields remain constant for an XMSS operation, the pre-computed values can be re-used across all signatures and thereby shared across 1024 signatures in a 10-level XMSS tree. For taller trees (e.g. a height of 20) the number of possible signature is 2^{20} , thus the savings achieved by this technique are even greater than in the 10-level case. The WOTS and L-Tree pre-computations improve the overall XMSS sign/verify latency significantly, thereby enabling a high-performance, energy-efficient digital signature implementation. Table 1 illustrates performance improvements achieved by various techniques described herein.

TABLE 1

	Average No. of SHA256 rounds for XMSS Verify
Baseline	249,600
WOTS PRF Opcode/Seed Pre-hash	180,991
WOTS PRF Address Pre-hash	173,488
WOTS Function-F Opcode Pre-hash	169,200
L-Tree PRF/Function-F Pre-hash	154,548 (-38%)

Application of Techniques to LMS Signature Scheme

The LMS signature and verification scheme is a competing proposal to XMSS scheme. LMS does not require the functions F, PRF, H, RAND_HASH as needed in XMSS. Instead, LMS scheme allows direct calls to the underlying hash function SHA-256. Similar to WOTS in XMSS, the LM-OTS one-time signature building block also operates over $p=67$ chunks of the private-key/signature to generate and verify signatures. Each one of these chunks requires a number of hash calls that depends on the message that is being signed/verified.

Table 2 illustrates the central part of the LM-OTS signature generation (left side) and signature verification (right side) algorithms. In both algorithms, the inner hash call in Line 6 of the algorithm (in bold) is called on average $p \times (2^w - 1) / 2$ times. One useful parameter set is $p=67$ and $w=4$, which leads to approximately 427 hash calls per signature generation/verification. The inputs to this hash call are the following:

- I: a 16-byte identifier of the LMS public/private keypair
- q: an integer that is expanded to 32-bytes
- i: an integer that is expanded to 16-bytes
- j: an integer that is expanded to 8-bytes
- tmp: a 32-bytes value

Acceleration techniques for LM-OTS signature/verification may leverage two facts. The first is that the certain input values are fixed for all hash calls (or change after a given number of hash calls). The second fact is that SHA256 consumes 64 input bytes at a time, and it runs 64 rounds for each of these 64-bytes input blocks. Thus acceleration techniques may be based on identifying sequences of fixed inputs to save SHA256 rounds computations.

The input size of hash call in line 6 of algorithms given in Table 1 is 104 bytes, which means that the input will be treated as two different blocks (each of 64 bytes, the last one padded with 24 bytes), thus leading to $2 \times 64 = 128$ SHA256 rounds to be computed. In total, this would lead to $67 \times (2^4 - 1) / 2 \times 128 = 64,320$ SHA256 rounds per signature generation/verification in a conventional design.

In some examples an acceleration technique may leverage the fact that I and q are always the same, which means that rounds processing 48 bytes input can be pre-computed. This would allow pre-computing 12 rounds of SHA256 for the first block. However, we point out that the value i only changes after $(2^w - 1) / 2 \approx 7$ hash calls on average. This means that in only $p=67$ times out of $67 \times 7 = 469$ it will be able to pre-compute only 12 rounds, while in all other occasions it will be able to pre-compute all 64 rounds. Thus, in some examples the number of hash calls in line 6 using such a pre-computation technique is $67 \times (128 - 12) + 67 \times ((2^4 - 1) / 2 - 1) \times (128 - 64) = 7,772 + 27,872 = 35,644$ SHA256 rounds. This represents a significant gain in the speed of the chain procedure of LMS. Since this is the most expensive step in LMS signature generation/verification (i.e., about 80-90% of

the total cost), this represents a significant gain in LMS signature generation/verification.

FIG. 11 illustrates an embodiment of an exemplary computing architecture that may be suitable for implementing various embodiments as previously described. In various embodiments, the computing architecture 1100 may comprise or be implemented as part of an electronic device. In some embodiments, the computing architecture 1100 may be representative, for example of a computer system that implements one or more components of the operating environments described above. In some embodiments, computing architecture 1100 may be representative of one or more portions or components of a DNN training system that implement one or more techniques described herein. The embodiments are not limited in this context.

As used in this application, the terms “system” and “component” and “module” are intended to refer to a computer-related entity, either hardware, a combination of hardware and software, software, or software in execution, examples of which are provided by the exemplary computing architecture 1100. For example, a component can be, but is not limited to being, a process running on a processor, a processor, a hard disk drive, multiple storage drives (of optical and/or magnetic storage medium), an object, an executable, a thread of execution, a program, and/or a computer. By way of illustration, both an application running on a server and the server can be a component. One or more components can reside within a process and/or thread of execution, and a component can be localized on one computer and/or distributed between two or more computers. Further, components may be communicatively coupled to each other by various types of communications media to coordinate operations. The coordination may involve the uni-directional or bi-directional exchange of information. For instance, the components may communicate information in the form of signals communicated over the communications media. The information can be implemented as signals allocated to various signal lines. In such allocations, each message is a signal. Further embodiments, however, may alternatively employ data messages. Such data messages may be sent across various connections. Exemplary connections include parallel interfaces, serial interfaces, and bus interfaces.

The computing architecture 1100 includes various common computing elements, such as one or more processors, multi-core processors, co-processors, memory units, chipsets, controllers, peripherals, interfaces, oscillators, timing devices, video cards, audio cards, multimedia input/output (I/O) components, power supplies, and so forth. The embodiments, however, are not limited to implementation by the computing architecture 1100.

As shown in FIG. 11, the computing architecture 1100 includes one or more processors 1102 and one or more graphics processors 1108, and may be a single processor desktop system, a multiprocessor workstation system, or a server system having a large number of processors 1102 or processor cores 1107. In one embodiment, the system 1100 is a processing platform incorporated within a system-on-a-chip (SoC or SOC) integrated circuit for use in mobile, handheld, or embedded devices.

An embodiment of system 1100 can include, or be incorporated within a server-based gaming platform, a game console, including a game and media console, a mobile gaming console, a handheld game console, or an online game console. In some embodiments system 1100 is a mobile phone, smart phone, tablet computing device or mobile Internet device. Data processing system 1100 can

also include, couple with, or be integrated within a wearable device, such as a smart watch wearable device, smart eyewear device, augmented reality device, or virtual reality device. In some embodiments, data processing system 1100 is a television or set top box device having one or more processors 1102 and a graphical interface generated by one or more graphics processors 1108.

In some embodiments, the one or more processors 1102 each include one or more processor cores 1107 to process instructions which, when executed, perform operations for system and user software. In some embodiments, each of the one or more processor cores 1107 is configured to process a specific instruction set 1109. In some embodiments, instruction set 1109 may facilitate Complex Instruction Set Computing (CISC), Reduced Instruction Set Computing (RISC), or computing via a Very Long Instruction Word (VLIW). Multiple processor cores 1107 may each process a different instruction set 1109, which may include instructions to facilitate the emulation of other instruction sets. Processor core 1107 may also include other processing devices, such as a Digital Signal Processor (DSP).

In some embodiments, the processor 1102 includes cache memory 1104. Depending on the architecture, the processor 1102 can have a single internal cache or multiple levels of internal cache. In some embodiments, the cache memory is shared among various components of the processor 1102. In some embodiments, the processor 1102 also uses an external cache (e.g., a Level-3 (L3) cache or Last Level Cache (LLC)) (not shown), which may be shared among processor cores 1107 using known cache coherency techniques. A register file 1106 is additionally included in processor 1102 which may include different types of registers for storing different types of data (e.g., integer registers, floating point registers, status registers, and an instruction pointer register). Some registers may be general-purpose registers, while other registers may be specific to the design of the processor 1102.

In some embodiments, one or more processor(s) 1102 are coupled with one or more interface bus(es) 1110 to transmit communication signals such as address, data, or control signals between processor 1102 and other components in the system. The interface bus 1110, in one embodiment, can be a processor bus, such as a version of the Direct Media Interface (DMI) bus. However, processor busses are not limited to the DMI bus, and may include one or more Peripheral Component Interconnect buses (e.g., PCI, PCI Express), memory busses, or other types of interface busses. In one embodiment the processor(s) 1102 include an integrated memory controller 1116 and a platform controller hub 1130. The memory controller 1116 facilitates communication between a memory device and other components of the system 1100, while the platform controller hub (PCH) 1130 provides connections to I/O devices via a local I/O bus.

Memory device 1120 can be a dynamic random-access memory (DRAM) device, a static random-access memory (SRAM) device, flash memory device, phase-change memory device, or some other memory device having suitable performance to serve as process memory. In one embodiment the memory device 1120 can operate as system memory for the system 1100, to store data 1122 and instructions 1121 for use when the one or more processors 1102 executes an application or process. Memory controller hub 1116 also couples with an optional external graphics processor 1112, which may communicate with the one or more graphics processors 1108 in processors 1102 to perform graphics and media operations. In some embodiments a display device 1111 can connect to the processor(s) 1102.

The display device **1111** can be one or more of an internal display device, as in a mobile electronic device or a laptop device or an external display device attached via a display interface (e.g., DisplayPort, etc.). In one embodiment the display device **1111** can be a head mounted display (HMD) such as a stereoscopic display device for use in virtual reality (VR) applications or augmented reality (AR) applications.

In some embodiments the platform controller hub **1130** enables peripherals to connect to memory device **1120** and processor **1102** via a high-speed I/O bus. The I/O peripherals include, but are not limited to, an audio controller **1146**, a network controller **1134**, a firmware interface **1128**, a wireless transceiver **1126**, touch sensors **1125**, a data storage device **1124** (e.g., hard disk drive, flash memory, etc.). The data storage device **1124** can connect via a storage interface (e.g., SATA) or via a peripheral bus, such as a Peripheral Component Interconnect bus (e.g., PCI, PCI Express). The touch sensors **1125** can include touch screen sensors, pressure sensors, or fingerprint sensors. The wireless transceiver **1126** can be a Wi-Fi transceiver, a Bluetooth transceiver, or a mobile network transceiver such as a 3G, 4G, or Long Term Evolution (LTE) transceiver. The firmware interface **1128** enables communication with system firmware, and can be, for example, a unified extensible firmware interface (UEFI). The network controller **1134** can enable a network connection to a wired network. In some embodiments, a high-performance network controller (not shown) couples with the interface bus **1110**. The audio controller **1146**, in one embodiment, is a multi-channel high definition audio controller. In one embodiment the system **1100** includes an optional legacy I/O controller **1140** for coupling legacy (e.g., Personal System 2 (PS/2)) devices to the system. The platform controller hub **1130** can also connect to one or more Universal Serial Bus (USB) controllers **1142** connect input devices, such as keyboard and mouse **1143** combinations, a camera **1144**, or other USB input devices.

The following pertains to further examples.

Example 1 is an apparatus comprising a computer readable memory; hash logic to generate a message hash value based on an input message; signature logic to generate a signature to be transmitted in association with the message, the signature logic to apply a hash-based signature scheme to a private key to generate the signature comprising a public key; and accelerator logic to pre-compute at least one set of inputs to the signature logic.

In Example 2, the subject matter of Example 1 can optionally include an arrangement in which the hash logic is to compute the message hash of the input message using a secure hash algorithm (SHA2) hash function; and process the message hash to generate an array of values.

In Example 3, the subject matter of any one of Examples 1-2 can optionally include an arrangement in which the signature logic is to apply a chain function to the private key, the chain function comprising a chain of SHA2 hash operations, wherein hash operations in the chain of SHA2 hash operations comprises a first pseudo-random function (PRF) to generate a bitmask and a second PRF to generate a key; an XOR logic to XOR the bitmask with a hash value from a previous hash function to generate an XOR output; and a SHA2 hash function to receive a concatenation of the XOR output and the bitmask as an input and to generate a hash output.

In Example 4, the subject matter of any one of Examples 1-3 can optionally include an arrangement in which the accelerator logic is to precompute a first intermediate hash value based on an OPCODE input and a SEED input; and

apply the first intermediate hash value to the first PRF and the second PRF in the chain of SHA2 hash operations in the signature logic.

In Example 5, the subject matter of any one of Examples 1-4 can optionally include an arrangement in which the accelerator logic is to precompute a second intermediate hash value based on a portion of an address input; and apply the second intermediate hash value to the first PRF and the second PRF in the chain of SHA2 hash operations in the signature logic.

In Example 6, the subject matter of any one of Examples 1-5 can optionally include an arrangement in which the accelerator logic is to precompute a third intermediate hash value based on a portion of an address input; and apply the third intermediate hash value to the first SHA2 hash function in the chain of SHA2 hash operations in the signature logic.

In Example 7, the subject matter of any one of Examples 1-6 can optionally include an arrangement in which the signature logic is to apply a L-tree operation which generates a binary tree to compress the public key into a single leaf node value using a series of keyed hash operations, wherein hash operations in the series of keyed hash operations comprises a first pseudo-random function (PRF) to generate a first bitmask; a second PRF to generate a second bitmask; a third PRF to generate a key; a first XOR logic to XOR the first bitmask with a first public key value from to generate a first XOR output; a second XOR logic to XOR the second bitmask with a second public key value from to generate a second XOR output and a SHA2 hash function to receive a concatenation of the first XOR output, the second XOR output, an OPCODE and PAD as an input and to generate a hash output.

In Example 8, the subject matter of any one of Examples 1-7 can optionally include an arrangement in which the accelerator logic is to precompute a first intermediate hash value based on an OPCODE input and a SEED input; and apply the first intermediate hash value to the first PRF, the second PRF, and the third PRF in the series of SHA2 hash operations in the L-Tree operation.

In Example 9, the subject matter of any one of Examples 1-8 can optionally include an arrangement in which the accelerator logic is to precompute a second intermediate hash value based on a portion of an address input; and apply the second intermediate hash value to the first PRF, the second PRF, and the third PRF in the series of SHA2 hash operations in the L-Tree operation.

In Example 10, the subject matter of any one of Examples 1-9 can optionally include an arrangement in which the accelerator logic is to precompute a third intermediate hash value based on a portion of an address input; and apply the third intermediate hash value to the SHA2 hash function in the series of SHA2 hash operations in the L-Tree operation.

Example 11 is a computer-implemented method, comprising generating a message hash value based on an input message; generating a signature to be transmitted in association with the message, the signature logic to apply a hash-based signature scheme to a private key to generate the signature comprising a public key; and pre-computing at least one set of inputs to the signature logic.

In Example 12, the subject matter of Example 11 can optionally further include computing the message hash of the input message using a secure hash algorithm (SHA2) hash function; and process the message hash to generate an array of values.

In Example 13, the subject matter of any one of Examples 11-12 can optionally include applying a chain function to the private key, the chain function comprising a chain of SHA2

hash operations, wherein hash operations in the chain of SHA2 hash operations comprises a first pseudo-random function (PRF) to generate a bitmask and a second PRF to generate a key; an XOR logic to XOR the bitmask with a hash value from a previous hash function to generate an XOR output; and a SHA2 hash function to receive a concatenation of the XOR output and the bitmask as an input and to generate a hash output.

In Example 14, the subject matter of any one of Examples 11-13 can optionally include precomputing a first intermediate hash value based on an OPCODE input and a SEED input; and applying the first intermediate hash value to the first PRF and the second PRF in the chain of SHA2 hash operations in the signature logic.

In Example 15, the subject matter of any one of Examples 11-14 can optionally include precomputing a second intermediate hash value based on a portion of an address input; and applying the second intermediate hash value to the first PRF and the second PRF in the chain of SHA2 hash operations in the signature logic.

In Example 16, the subject matter of any one of Examples 11-15 can optionally include precomputing a third intermediate hash value based on a portion of an address input; and applying the third intermediate hash value to the first SHA2 hash function in the chain of SHA2 hash operations in the signature logic.

In Example 17, the subject matter of any one of Examples 11-16 can optionally include applying a L-tree operation which generates a binary tree to compress the public key into a single leaf node value using a series of keyed hash operations, wherein hash operations in the series of keyed hash operations comprises a first pseudo-random function (PRF) to generate a first bitmask; a second PRF to generate a second bitmask; a third PRF to generate a key; a first XOR logic to XOR the first bitmask with a first public key value from to generate a first XOR output; a second XOR logic to XOR the second bitmask with a second public key value from to generate a second XOR output and a SHA2 hash function to receive a concatenation of the first XOR output, the second XOR output, an OPCODE and PAD as an input and to generate a hash output.

In Example 18, the subject matter of any one of Examples 11-17 can optionally include precomputing a first intermediate hash value based on an OPCODE input and a SEED input; and applying the first intermediate hash value to the first PRF, the second PRF, and the third PRF in the series of SHA2 hash operations in the L-Tree operation.

In Example 19, the subject matter of any one of Examples 11-18 can optionally include precomputing a second intermediate hash value based on a portion of an address input; and applying the second intermediate hash value to the first PRF, the second PRF, and the third PRF in the series of SHA2 hash operations in the L-Tree operation.

In Example 20, the subject matter of any one of Examples 11-19 can optionally include precompute a third intermediate hash value based on a portion of an address input; and applying the third intermediate hash value to the SHA2 hash function in the series of SHA2 hash operations in the L-Tree operation.

The above Detailed Description includes references to the accompanying drawings, which form a part of the Detailed Description. The drawings show, by way of illustration, specific embodiments that may be practiced. These embodiments are also referred to herein as "examples." Such examples may include elements in addition to those shown or described. However, also contemplated are examples that include the elements shown or described. Moreover, also

contemplated are examples using any combination or permutation of those elements shown or described (or one or more aspects thereof), either with respect to a particular example (or one or more aspects thereof), or with respect to other examples (or one or more aspects thereof) shown or described herein.

Publications, patents, and patent documents referred to in this document are incorporated by reference herein in their entirety, as though individually incorporated by reference. In the event of inconsistent usages between this document and those documents so incorporated by reference, the usage in the incorporated reference(s) are supplementary to that of this document; for irreconcilable inconsistencies, the usage in this document controls.

In this document, the terms "a" or "an" are used, as is common in patent documents, to include one or more than one, independent of any other instances or usages of "at least one" or "one or more." In addition "a set of" includes one or more elements. In this document, the term "or" is used to refer to a nonexclusive or, such that "A or B" includes "A but not B," "B but not A," and "A and B," unless otherwise indicated. In the appended claims, the terms "including" and "in which" are used as the plain-English equivalents of the respective terms "comprising" and "wherein." Also, in the following claims, the terms "including" and "comprising" are open-ended; that is, a system, device, article, or process that includes elements in addition to those listed after such a term in a claim are still deemed to fall within the scope of that claim. Moreover, in the following claims, the terms "first," "second," "third," etc. are used merely as labels, and are not intended to suggest a numerical order for their objects.

The terms "logic instructions" as referred to herein relates to expressions which may be understood by one or more machines for performing one or more logical operations. For example, logic instructions may comprise instructions which are interpretable by a processor compiler for executing one or more operations on one or more data objects. However, this is merely an example of machine-readable instructions and examples are not limited in this respect.

The terms "computer readable medium" as referred to herein relates to media capable of maintaining expressions which are perceivable by one or more machines. For example, a computer readable medium may comprise one or more storage devices for storing computer readable instructions or data. Such storage devices may comprise storage media such as, for example, optical, magnetic or semiconductor storage media. However, this is merely an example of a computer readable medium and examples are not limited in this respect.

The term "logic" as referred to herein relates to structure for performing one or more logical operations. For example, logic may comprise circuitry which provides one or more output signals based upon one or more input signals. Such circuitry may comprise a finite state machine which receives a digital input and provides a digital output, or circuitry which provides one or more analog output signals in response to one or more analog input signals. Such circuitry may be provided in an application specific integrated circuit (ASIC) or field programmable gate array (FPGA). Also, logic may comprise machine-readable instructions stored in a memory in combination with processing circuitry to execute such machine-readable instructions. However, these are merely examples of structures which may provide logic and examples are not limited in this respect.

Some of the methods described herein may be embodied as logic instructions on a computer-readable medium. When

executed on a processor, the logic instructions cause a processor to be programmed as a special-purpose machine that implements the described methods. The processor, when configured by the logic instructions to execute the methods described herein, constitutes structure for performing the described methods. Alternatively, the methods described herein may be reduced to logic on, e.g., a field programmable gate array (FPGA), an application specific integrated circuit (ASIC) or the like.

In the description and claims, the terms coupled and connected, along with their derivatives, may be used. In particular examples, connected may be used to indicate that two or more elements are in direct physical or electrical contact with each other. Coupled may mean that two or more elements are in direct physical or electrical contact. However, coupled may also mean that two or more elements may not be in direct contact with each other, but yet may still cooperate or interact with each other.

Reference in the specification to “one example” or “some examples” means that a particular feature, structure, or characteristic described in connection with the example is included in at least an implementation. The appearances of the phrase “in one example” in various places in the specification may or may not be all referring to the same example.

The above description is intended to be illustrative, and not restrictive. For example, the above-described examples (or one or more aspects thereof) may be used in combination with others. Other embodiments may be used, such as by one of ordinary skill in the art upon reviewing the above description. The Abstract is to allow the reader to quickly ascertain the nature of the technical disclosure. It is submitted with the understanding that it will not be used to interpret or limit the scope or meaning of the claims. Also, in the above Detailed Description, various features may be grouped together to streamline the disclosure. However, the claims may not set forth every feature disclosed herein as embodiments may feature a subset of said features. Further, embodiments may include fewer features than those disclosed in a particular example. Thus, the following claims are hereby incorporated into the Detailed Description, with each claim standing on its own as a separate embodiment. The scope of the embodiments disclosed herein is to be determined with reference to the appended claims, along with the full scope of equivalents to which such claims are entitled.

Although examples have been described in language specific to structural features and/or methodological acts, it is to be understood that claimed subject matter may not be limited to the specific features or acts described. Rather, the specific features and acts are disclosed as sample forms of implementing the claimed subject matter.

What is claimed is:

1. An apparatus, comprising a hardware processor to:
 - receive an input message;
 - precompute a first intermediate hash value based on an OPCODE input and a SEED input; and
 - apply the first intermediate hash value to the first pseudo-random function (PRF) to generate a bitmask and the second pseudo-random function (PRF) to generate a private key;
 - apply a chain function to the private key, the chain function comprising a chain of secure hash algorithm (SHA)2 hash operations;
 - generate a message hash value based at least in part on the input message;

generate a signature to be transmitted in association with the message, the signature comprising a public key generated by applying a hash-based signature scheme to the private key.

2. The apparatus of claim 1, the hardware processor to:
 - compute the message hash of the input message using a secure hash algorithm (SHA2) hash function; and
 - process the message hash to generate an array of values.

3. The apparatus of claim 2, wherein hash operations in the chain of SHA2 hash operations comprises:
 - an XOR logic to XOR the bitmask with a hash value from a previous hash function to generate an XOR output; and
 - a SHA2 hash function to receive a concatenation of the XOR output and the bitmask as an input and to generate a hash output.

4. The apparatus of claim 3, the hardware processor to:
 - apply the first intermediate hash value to the first PRF and the second PRF in the chain of SHA2 hash operations in the signature logic.

5. The apparatus of claim 4, the hardware processor to:
 - precompute a second intermediate hash value based on a portion of an address input; and
 - apply the second intermediate hash value to the first PRF and the second PRF in the chain of SHA2 hash operations.

6. The apparatus of claim 5, the hardware processor to:
 - precompute a third intermediate hash value based on a portion of an address input; and
 - apply the third intermediate hash value to the first SHA2 hash function in the chain of SHA2 hash operations.

7. The apparatus of claim 2, the hardware processor to:
 - apply a L-tree operation which generates a binary tree to compress the public key into a single leaf node value using a series of keyed hash operations, wherein hash operations in the series of keyed hash operations comprises:
 - a first pseudo-random function (PRF) to generate a first bitmask;
 - a second PRF to generate a second bitmask;
 - a third PRF to generate a key;
 - a first XOR logic to XOR the first bitmask with a first public key value from to generate a first XOR output;
 - a second XOR logic to XOR the second bitmask with a second public key value from to generate a second XOR output and
 - a SHA2 hash function to receive a concatenation of the first XOR output, the second XOR output, an OPCODE and PAD as an input and to generate a hash output.

8. The apparatus of claim 7, the hardware processor to:
 - precompute a first intermediate hash value based on an OPCODE input and a SEED input; and
 - apply the first intermediate hash value to the first PRF, the second PRF, and the third PRF in the series of SHA2 hash operations in the L-Tree operation.

9. The apparatus of claim 8, the hardware processor to:
 - precompute a second intermediate hash value based on a portion of an address input; and
 - apply the second intermediate hash value to the first PRF, the second PRF, and the third PRF in the series of SHA2 hash operations in the L-Tree operation.

10. The apparatus of claim 9, the hardware processor to:
 - precompute a third intermediate hash value based on a portion of an address input; and

11. The apparatus of claim 10, the hardware processor to:
 - compute the message hash of the input message using a secure hash algorithm (SHA2) hash function; and
 - process the message hash to generate an array of values.

12. The apparatus of claim 11, wherein hash operations in the chain of SHA2 hash operations comprises:
 - an XOR logic to XOR the bitmask with a hash value from a previous hash function to generate an XOR output; and
 - a SHA2 hash function to receive a concatenation of the XOR output and the bitmask as an input and to generate a hash output.

13. The apparatus of claim 12, the hardware processor to:
 - apply the first intermediate hash value to the first PRF and the second PRF in the chain of SHA2 hash operations in the signature logic.

14. The apparatus of claim 13, the hardware processor to:
 - precompute a second intermediate hash value based on a portion of an address input; and
 - apply the second intermediate hash value to the first PRF and the second PRF in the chain of SHA2 hash operations.

15. The apparatus of claim 14, the hardware processor to:
 - precompute a third intermediate hash value based on a portion of an address input; and
 - apply the third intermediate hash value to the first SHA2 hash function in the chain of SHA2 hash operations.

16. The apparatus of claim 11, the hardware processor to:
 - apply a L-tree operation which generates a binary tree to compress the public key into a single leaf node value using a series of keyed hash operations, wherein hash operations in the series of keyed hash operations comprises:
 - a first pseudo-random function (PRF) to generate a first bitmask;
 - a second PRF to generate a second bitmask;
 - a third PRF to generate a key;
 - a first XOR logic to XOR the first bitmask with a first public key value from to generate a first XOR output;
 - a second XOR logic to XOR the second bitmask with a second public key value from to generate a second XOR output and
 - a SHA2 hash function to receive a concatenation of the first XOR output, the second XOR output, an OPCODE and PAD as an input and to generate a hash output.

21

apply the third intermediate hash value to the SHA2 hash function in the series of SHA2 hash operations in the L-Tree operation.

11. A computer-implemented method, comprising:
 receiving an input message;
 precomputing a first intermediate hash value based on an
 OPCODE input and a SEED input; and
 applying the first intermediate hash value to the first
 pseudo-random function (PRF) to generate a bitmask
 and the second pseudo-random function (PRF) to generate
 a private key;
 applying a chain function to the private key, the chain
 function comprising a chain of secure hash algorithm
 (SHA)2 hash operations;
 generating a message hash value based at least in part on
 the input message;
 generating a signature to be transmitted in association
 with the message, the signature comprising a public
 key generated by applying a hash-based signature
 scheme to the private key.

12. The method of claim 11, further comprising:
 computing the message hash of the input message using
 a secure hash algorithm (SHA2) hash function; and
 processing the message hash to generate an array of
 values.

13. The method of claim 12, wherein hash operations in
 the chain of SHA2 hash operations comprises:
 an XOR logic to XOR the bitmask with a hash value from
 a previous hash function to generate an XOR output;
 and
 a SHA2 hash function to receive a concatenation of the
 XOR output and the bitmask as an input and to generate
 a hash output.

14. The method of claim 13, further comprising:
 precomputing a first intermediate hash value based on an
 OPCODE input and a SEED input; and
 applying the first intermediate hash value to the first PRF
 and the second PRF in the chain of SHA2 hash operations.

15. The method of claim 14, further comprising:
 precomputing a second intermediate hash value based on
 a portion of an address input; and
 applying the second intermediate hash value to the first
 PRF and the second PRF in the chain of SHA2 hash
 operations.

22

16. The method of claim 15, further comprising:
 precomputing a third intermediate hash value based on a
 portion of an address input; and
 applying the third intermediate hash value to the first
 SHA2 hash function in the chain of SHA2 hash operations.

17. The method of claim 12, further comprising:
 apply a L-tree operation which generates a binary tree to
 compress the public key into a single leaf node value
 using a series of keyed hash operations, wherein hash
 operations in the series of keyed hash operations comprises:
 a first pseudo-random function (PRF) to generate a first
 bitmask;
 a second PRF to generate a second bitmask;
 a third PRF to generate a key;
 a first XOR logic to XOR the first bitmask with a first
 public key value from to generate a first XOR output;
 a second XOR logic to XOR the second bitmask with
 a second public key value from to generate a second
 XOR output and
 a SHA2 hash function to receive a concatenation of the
 first XOR output, the second XOR output, an
 OPCODE and PAD as an input and to generate a
 hash output.

18. The method of claim 17, further comprising:
 precomputing a first intermediate hash value based on an
 OPCODE input and a SEED input; and
 applying the first intermediate hash value to the first PRF,
 the second PRF, and the third PRF in the series of
 SHA2 hash operations in the L-Tree operation.

19. The method of claim 18, further comprising:
 precomputing a second intermediate hash value based on
 a portion of an address input; and
 applying the second intermediate hash value to the first
 PRF, the second PRF, and the third PRF in the series of
 SHA2 hash operations in the L-Tree operation.

20. The method of claim 19, further comprising:
 precomputing a third intermediate hash value based on a
 portion of an address input; and
 applying the third intermediate hash value to the SHA2
 hash function in the series of SHA2 hash operations in
 the L-Tree operation.

* * * * *