

BLOCKCHAIN DATA STORAGE BASED ON ERROR CORRECTION CODE FOR PERMISSIONED BLOCKCHAIN NETWORK

TECHNICAL FIELD

[0001] This specification relates to storage of blockchain data based on error correction code in permissioned blockchain networks.

BACKGROUND

[0002] Distributed ledger systems (DLSs), which can also be referred to as consensus networks, and/or blockchain networks, enable participating entities to securely and immutably store data. DLSs are commonly referred to as blockchain networks without referencing any particular user case. Examples of types of blockchain networks can include public blockchain networks, private blockchain networks, and consortium blockchain networks. A consortium blockchain network is provided for a select group of entities, which control the consensus process, and includes an access control layer.

[0003] Blockchain-based programs can be executed by a distributed computing platform. For example, the distributed computing platform can include a virtual machine that provides the runtime environment for executing smart contracts. A blockchain computing platform can be viewed as a transaction-based state machine. State data in the platform can be assembled to a global shared-state referred to as a world state. The world state includes a mapping between account addresses and account states. The world state can be stored in data structures such as the Merkle Patricia tree (MPT).

[0004] Besides state data, blockchain networks can also store other types of data such as block data and index data. Block data can include block header and block body. The block header can include identity information of a particular block and the block body can include transactions that are confirmed with the block. As transactions are increasingly entered into the blockchain, state data and block data can grow very large in size. In some DLSs, every node stores an entire copy of the blockchain, which can take large amount of storage space. This is because all block data and state data are stored going back to the first transaction recorded to the blockchain. In some DLSs, a few shared nodes store the entire copy of the blockchain and share blockchain data with other blockchain nodes which can create “data

inequality.” That is, when data are unevenly distributed across different nodes, the risk of data security can be high when nodes that store majority of data are at fault.

[0005] In some consortium blockchain networks or private blockchain networks, some nodes do not participate in consensus. It would be desirable to shift some storage burden of the consensus nodes by enabling a data storage scheme among the non-consensus nodes. It would also be desirable to shift the computational burden of recovering blockchain data to the non-consensus nodes while maintaining blockchain data equality and integrity.

SUMMARY

[0006] Described embodiments of the subject matter can include one or more features, alone or in combination.

[0007] For example, in one embodiment, a computer-implemented method for blockchain data storage performed by a non-validating node (NVN) of a permissioned blockchain network is disclosed. The NVN can receive transactions issued by a client node communicably coupled to the permissioned blockchain network, send the transactions to a plurality of validating nodes for validating the transactions and running consensus, receive one or more blocks generated based on consensus of at least a portion of the transactions by the plurality of validating nodes, encode the one or more blocks using ECC to generate one or more encoded blocks. For each encoded block of the one or more encoded blocks, the NVN can divide the encoded block into a plurality of datasets based on a data storage scheme associated with the one or more blocks, and store at least one of the plurality of datasets based on the assignments provided in the data storage scheme.

[0008] In some embodiments, these general and specific aspects may be implemented using a system, a method, or a computer program, or any combination of systems, methods, and computer programs. The foregoing and other described embodiments can each, optionally, include one or more of the following features:

[0009] In some embodiments, the NVN further receives confirmations from a remainder of the plurality of NVNs other than the NVN that at least one of the plurality of datasets are stored by the corresponding NVN based on the assignments provided in the data storage scheme, and sends a notification to the plurality of validating nodes to delete the one or more blocks.

[0010] In some embodiments, the NVN stores blockchain data in a tiered storage that a newer block is stored in a same or faster tier of the tiered storage than an older block.

[0011] In some embodiments, the NVN further receives a request from the client node for verifying a transaction included in the one or more blocks, retrieves, based on the index, datasets associated with the one or more blocks from the remainder of the plurality of NVNs, decodes the one or more blocks based on the datasets, and verifies that the transaction exists if the transaction is included in the one or more blocks.

[0012] In some embodiments, the NVN further receives a request from a validating node for retrieving a transaction included in the one or more blocks, retrieve, based on the index, datasets associated with the one of the one or more blocks from the remainder of the plurality of NVNs, and decodes the one of the one or more blocks based on the datasets, and sends the transaction to the validating node.

[0013] In some embodiments, the ECC is performed when utilization rate of computational resource of the NVN is less than or equal to a predetermined value.

[0014] In some embodiments, for each encoded block of the plurality of encoded blocks, the NVN further hashes a remainder of the plurality of datasets other than the at least one of the plurality of datasets to generate hash values, stores the hash values, and discards the one or more blocks and the remainder of the plurality of datasets.

[0015] In some embodiments, the ECC is performed when usage of storage space of the NVN is greater than or equal to a predetermined percentage.

[0016] In some embodiments, the ECC is erasure coding performed by adding redundant bits to the plurality of blocks.

[0017] It is appreciated that methods in accordance with this specification may include any combination of the aspects and features described herein. That is, methods in accordance with this specification are not limited to the combinations of aspects and features specifically described herein, but also include any combination of the aspects and features provided.

[0018] The details of one or more embodiments of this specification are set forth in the accompanying drawings and the description below. Other features and advantages of this specification will be apparent from the description and drawings, and from the claims.

BRIEF DESCRIPTION OF THE DRAWINGS

[0019] FIG. 1 depicts an example of an environment that can be used to execute embodiments of this specification.

[0020] FIG. 2 depicts an example of an architecture in accordance with embodiments of this specification.

[0021] FIG. 3 depicts an example of a block data encoding and hashing process in accordance with embodiments of this specification.

[0022] FIG. 4 depicts an example of a data storage scheme in accordance with embodiments of this specification.

[0023] FIG. 5 depicts another example of a block data encoding and hashing process in accordance with embodiments of this specification.

[0024] FIG. 6 depicts an example of a process of shared blockchain data storage based on non-validating nodes (NVNs) in accordance with embodiments of this specification.

[0025] FIG. 7 depicts an example of a process that can be executed in accordance with embodiments of this specification.

[0026] FIG. 8 depicts examples of modules of an apparatus in accordance with embodiments of this specification.

[0027] Like reference numbers and designations in the various drawings indicate like elements.

DETAILED DESCRIPTION

[0028] This specification describes technologies for storing blockchain data based on error correction code (ECC) in a permissioned blockchain network. These technologies generally involve receiving transactions issued by a client node communicably coupled to a permissioned blockchain network, sending the transactions to a plurality of validating nodes for validating the transactions and running consensus, receiving one or more blocks generated based on consensus of at least a portion of the transactions by the plurality of validating nodes, encoding the one or more blocks using ECC to generate one or more encoded blocks, and for each encoded block of the one or more encoded blocks: dividing the encoded block into a plurality of datasets based on a data storage scheme associated with the one or more blocks.

[0029] As described herein, blockchain networks can store different types of data such as state data, block data, and index data. Block data includes all transactions in the blockchain network, which can take a large amount of storage space as new blocks are constantly adding to the blockchain. It can be inefficient for the blockchain nodes to each store all the block data, especially for data of infrequently accessed blocks (e.g., blocks added to the blockchain long time ago).

[0030] In permissioned blockchain networks, such as private blockchain networks or consortium networks, the distributed ledger protocols can be run by validating nodes (VNs) and non-validating nodes (NVNs). The VNs and NVNs are also known as peers. The VNs can be consensus nodes responsible for running consensus, validating transactions, and maintaining the distributed ledger. The NVNs do not participate in consensus or execute transactions. In some embodiments, instead of storing blocks generated through consensus, the VNs can send the blocks (except for the most recent ones) to the NVNs for storage. Moreover, after the NVNs receive the blocks, they can perform ECC to encode the blocks and divide the ECC encoded blocks to datasets to be separately stored by each NVN. Accordingly, some embodiments of this specification provide that in a permissioned blockchain network that includes VNs and NVNs, the VNs store only recently created blocks and the NVNs store a portion of ECC encoded blocks to lower the storage cost of both VNs and NVNs.

[0031] The techniques described in this specification produce several technical effects. For example, embodiments of the subject matter reduce the burden on storage resources of blockchain networks, while maintaining computational efficiency and data equality of the blockchain nodes. For infrequently accessed blocks (e.g., older blocks), storage resources of blockchain nodes can be conserved by saving only a portion of ECC encoded blocks (also referred to herein as encoded blocks) on each blockchain node and retrieving the remainder of the encoded blocks from other blockchain nodes when needed.

[0032] As another example, when the blockchain network is a permissioned blockchain network that includes VNs and NVNs. The VNs can reach consensus to generate blocks and send the blocks to the NVNs for storage. After the NVNs receive the blocks, they can perform ECC to encode the blocks and divide the encoded blocks to datasets to be separately stored by each NVN to save storage space. As such, some embodiments of the subject matter

reduce the storage consumption of both the VNs and NVNs. Moreover, because NVNs do not participate in running consensus of the transactions, performing ECC by the NVNs can shift some computational burden from the VNs to the NVNs, while achieving cost savings of storage space. Similarly, when transactions in old blocks need to be accessed by the VNs, the VNs can send a request to the NVNs to recover the corresponding block that contains the transactions. The computational resource consumption of the VNs for decoding the ECC encoded blocks can also be saved.

[0033] To provide further context for embodiments of this specification, and as introduced above, distributed ledger systems (DLSs), which can also be referred to as consensus networks (e.g., made up of peer-to-peer nodes), and blockchain networks, enable participating entities to securely, and immutably conduct transactions, and store data. Although the term blockchain is generally associated with particular networks, and/or use cases, blockchain is used herein to generally refer to a DLS without reference to any particular use case.

[0034] A blockchain is a data structure that stores transactions in a way that the transactions are immutable. Thus, transactions recorded on a blockchain are reliable and trustworthy. A blockchain includes one or more blocks. Each block in the chain is linked to a previous block immediately before it in the chain by including a cryptographic hash of the previous block. Each block also includes a timestamp, its own cryptographic hash, and one or more transactions. The transactions, which have already been verified by the nodes of the blockchain network, are hashed and encoded into a Merkle tree. A Merkle tree is a data structure in which data at the leaf nodes of the tree is hashed, and all hashes in each branch of the tree are concatenated at the root of the branch. This process continues up the tree to the root of the entire tree, which stores a hash that is representative of all data in the tree. A hash purporting to be of a transaction stored in the tree can be quickly verified by determining whether it is consistent with the structure of the tree.

[0035] Whereas a blockchain is a decentralized or at least partially decentralized data structure for storing transactions, a blockchain network is a network of computing nodes that manage, update, and maintain one or more blockchains by broadcasting, verifying and validating transactions, etc. As introduced above, a blockchain network can be provided as a public blockchain network, a private blockchain network, or a consortium blockchain

network. Embodiments of this specification are described in further detail herein with reference to a consortium blockchain network. It is contemplated, however, that embodiments of this specification can be realized in any appropriate type of blockchain network.

[0036] In general, a consortium blockchain network is private among the participating entities. In a consortium blockchain network, the consensus process is controlled by an authorized set of nodes, which can be referred to as consensus nodes, one or more consensus nodes being operated by a respective entity (e.g., a financial institution, insurance company). For example, a consortium of ten (10) entities (e.g., financial institutions, insurance companies) can operate a consortium blockchain network, each of which operates at least one node in the consortium blockchain network.

[0037] In some examples, within a consortium blockchain network, a global blockchain is provided as a blockchain that is replicated across all nodes. That is, all consensus nodes are in perfect state consensus with respect to the global blockchain. To achieve consensus (e.g., agreement to the addition of a block to a blockchain), a consensus protocol is implemented within the consortium blockchain network. For example, the consortium blockchain network can implement a practical Byzantine fault tolerance (PBFT) consensus, described in further detail below.

[0038] FIG. 1 is a diagram illustrating an example of an environment 100 that can be used to execute embodiments of this specification. In some examples, the environment 100 enables entities to participate in a consortium blockchain network 102. The environment 100 includes computing devices 106, 108, and a network 110. In some examples, the network 110 includes a local area network (LAN), wide area network (WAN), the Internet, or a combination thereof, and connects web sites, user devices (e.g., computing devices), and back-end systems. In some examples, the network 110 can be accessed over a wired and/or a wireless communications link. In some examples, the network 110 enables communication with, and within the consortium blockchain network 102. In general, the network 110 represents one or more communication networks. In some cases, the computing devices 106, 108 can be nodes of a cloud computing system (not shown), or each computing device 106, 108 can be a separate cloud computing system including a number of computers interconnected by a network and functioning as a distributed processing system.

[0039] In the depicted example, the computing systems 106, 108 can each include any appropriate computing system that enables participation as a node in the consortium blockchain network 102. Examples of computing devices include, without limitation, a server, a desktop computer, a laptop computer, a tablet computing device, and a smartphone. In some examples, the computing systems 106, 108 host one or more computer-implemented services for interacting with the consortium blockchain network 102. For example, the computing system 106 can host computer-implemented services of a first entity (e.g., user A), such as a transaction management system that the first entity uses to manage its transactions with one or more other entities (e.g., other users). The computing system 108 can host computer-implemented services of a second entity (e.g., user B), such as a transaction management system that the second entity uses to manage its transactions with one or more other entities (e.g., other users). In the example of FIG. 1, the consortium blockchain network 102 is represented as a peer-to-peer network of nodes, and the computing systems 106, 108 provide nodes of the first entity, and second entity respectively, which participate in the consortium blockchain network 102.

[0040] FIG. 2 depicts an example of an architecture 200 in accordance with embodiments of this specification. The example conceptual architecture 200 includes participant systems 202, 204, 206 that correspond to Participant A, Participant B, and Participant C, respectively. Each participant (e.g., user, enterprise) participates in a blockchain network 212 provided as a peer-to-peer network including a plurality of nodes 214, at least some of which immutably record information in a blockchain 216. Although a single blockchain 216 is schematically depicted within the blockchain network 212, multiple copies of the blockchain 216 are provided, and are maintained across the blockchain network 212, as described in further detail herein.

[0041] In the depicted example, each participant system 202, 204, 206 is provided by, or on behalf of Participant A, Participant B, and Participant C, respectively, and functions as a respective node 214 within the blockchain network. As used herein, a node generally refers to an individual system (e.g., computer, server) that is connected to the blockchain network 212, and enables a respective participant to participate in the blockchain network. In the example of FIG. 2, a participant corresponds to each node 214. It is contemplated, however, that a participant can operate multiple nodes 214 within the blockchain network 212, and/or

multiple participants can share a node 214. In some examples, the participant systems 202, 204, 206 communicate with, or through the blockchain network 212 using a protocol (e.g., hypertext transfer protocol secure (HTTPS)), and/or using remote procedure calls (RPCs).

[0042] Nodes 214 can have varying degrees of participation within the blockchain network 212. For example, some nodes 214 can participate in the consensus process (e.g., as miner nodes that add blocks to the blockchain 216), while other nodes 214 do not participate in the consensus process. As another example, some nodes 214 store a complete copy of the blockchain 216, while other nodes 214 only store copies of portions of the blockchain 216. For example, data access privileges can limit the blockchain data that a respective participant stores within its respective system. In the example of FIG. 2, the participant systems 202, 204, and 206 store respective, complete copies 216', 216'', and 216''' of the blockchain 216.

[0043] A blockchain (e.g., the blockchain 216 of FIG. 2) is made up of a chain of blocks, each block storing data. Examples of data include transaction data representative of a transaction between two or more participants. While transactions are used herein by way of non-limiting example, it is contemplated that any appropriate data can be stored in a blockchain (e.g., documents, images, videos, audio). Examples of a transaction can include, without limitation, exchanges of something of value (e.g., assets, products, services, currency). The transaction data is immutably stored within the blockchain. That is, the transaction data cannot be changed.

[0044] Before storing in a block, the transaction data is hashed. Hashing is a process of transforming the transaction data (provided as string data) into a fixed-length hash value (also provided as string data). It is not possible to un-hash the hash value to obtain the transaction data. Hashing ensures that even a slight change in the transaction data results in a completely different hash value. Further, and as noted above, the hash value is of fixed length. That is, no matter the size of the transaction data the length of the hash value is fixed. Hashing includes processing the transaction data through a hash function to generate the hash value. An example of a hash function includes, without limitation, the secure hash algorithm (SHA)-256, which outputs 256-bit hash values.

[0045] Transaction data of multiple transactions are hashed and stored in a block. For example, hash values of two transactions are provided, and are themselves hashed to provide another hash. This process is repeated until, for all transactions to be stored in a block, a

single hash value is provided. This hash value is referred to as a Merkle root hash, and is stored in a header of the block. A change in any of the transactions will result in change in its hash value, and ultimately, a change in the Merkle root hash.

[0046] Blocks are added to the blockchain through a consensus protocol. Multiple nodes within the blockchain network participate in the consensus protocol, and perform work to have a block added to the blockchain. Such nodes are referred to as consensus nodes. PBFT, introduced above, is used as a non-limiting example of a consensus protocol. The consensus nodes execute the consensus protocol to add transactions to the blockchain, and update the overall state of the blockchain network.

[0047] In further detail, the consensus node generates a block header, hashes all of the transactions in the block, and combines the hash value in pairs to generate further hash values until a single hash value is provided for all transactions in the block (the Merkle root hash). This hash is added to the block header. The consensus node also determines the hash value of the most recent block in the blockchain (i.e., the last block added to the blockchain). The consensus node also adds a nonce value, and a timestamp to the block header.

[0048] In general, PBFT provides a practical Byzantine state machine replication that tolerates Byzantine faults (e.g., malfunctioning nodes, malicious nodes). This is achieved in PBFT by assuming that faults will occur (e.g., assuming the existence of independent node failures, and/or manipulated messages sent by consensus nodes). In PBFT, the consensus nodes are provided in a sequence that includes a primary consensus node, and backup consensus nodes. The primary consensus node is periodically changed. Transactions are added to the blockchain by all consensus nodes within the blockchain network reaching an agreement as to the world state of the blockchain network. In this process, messages are transmitted between consensus nodes, and each consensus nodes proves that a message is received from a specified peer node, and verifies that the message was not modified during transmission.

[0049] In PBFT, the consensus protocol is provided in multiple phases with all consensus nodes beginning in the same state. To begin, a client sends a request to the primary consensus node to invoke a service operation (e.g., execute a transaction within the blockchain network). In response to receiving the request, the primary consensus node multicasts the request to the backup consensus nodes. The backup consensus nodes execute the request, and each sends a

reply to the client. The client waits until a threshold number of replies are received. In some examples, the client waits for $f+1$ replies to be received, where f is the maximum number of faulty consensus nodes that can be tolerated within the blockchain network. In PBFT, the maximum number of faulty consensus nodes needs to be less than $1/3$ of the total number of consensus nodes. The final result is that a sufficient number of consensus nodes come to an agreement on the order of the record that is to be added to the blockchain, and the record is either accepted, or rejected.

[0050] In some blockchain networks, cryptography is implemented to maintain privacy of transactions. For example, if two nodes want to keep a transaction private, such that other nodes in the blockchain network cannot discern details of the transaction, the nodes can encrypt the transaction data. An example of cryptography includes, without limitation, symmetric encryption, and asymmetric encryption. Symmetric encryption refers to an encryption process that uses a single key for both encryption (generating ciphertext from plaintext), and decryption (generating plaintext from ciphertext). In symmetric encryption, the same key is available to multiple nodes, so each node can en-/de-crypt transaction data.

[0051] Asymmetric encryption uses keys pairs that each include a private key, and a public key, the private key being known only to a respective node, and the public key being known to any or all other nodes in the blockchain network. A node can use the public key of another node to encrypt data, and the encrypted data can be decrypted using other node's private key. For example, and referring again to FIG. 2, Participant A can use Participant B's public key to encrypt data, and send the encrypted data to Participant B. Participant B can use its private key to decrypt the encrypted data (ciphertext) and extract the original data (plaintext). Messages encrypted with a node's public key can only be decrypted using the node's private key.

[0052] Asymmetric encryption is used to provide digital signatures, which enables participants in a transaction to confirm other participants in the transaction, as well as the validity of the transaction. For example, a node can digitally sign a message, and another node can confirm that the message was sent by the node based on the digital signature of Participant A. Digital signatures can also be used to ensure that messages are not tampered with in transit. For example, and again referencing FIG. 2, Participant A is to send a message to Participant B. Participant A generates a hash of the message, and then, using its private

key, encrypts the hash to provide a digital signature as the encrypted hash. Participant A appends the digital signature to the message, and sends the message with digital signature to Participant B. Participant B decrypts the digital signature using the public key of Participant A, and extracts the hash. Participant B hashes the message and compares the hashes. If the hashes are same, Participant B can confirm that the message was indeed from Participant A, and was not tampered with.

[0053] FIG. 3 depicts an example of a block data encoding and hashing process 300 in accordance with embodiments of this specification. In this example, a blockchain network of four blockchain nodes is depicted, which are blockchain nodes 302, 304, 306, and 308. Using blockchain node 302 as an example to illustrate the encoding and hashing process 300, the blockchain node 302 can store block data of the blockchain network to block body of a block 312. In the illustrated example, the block data is stored in block 100. Afterwards, the blockchain node 302 can engage in a consensus process with other blockchain nodes 304, 306, and 308. During the consensus process, the blockchain node 302 can perform a consensus algorithm, such as proof of work (PoW) or proof of stake (PoS) to create a corresponding block on the blockchain.

[0054] In some embodiments, the blockchain node 302 can identify one or more infrequently accessed blocks. In practice, the longer a block has been created, the less likely the corresponding block data is needed for operations such as executing smart contracts. The blockchain node 302 can determine that locally stored blocks are infrequently accessed when they are historical blocks that have been created on the blockchain for a predetermined amount of time. For example, the predetermined amount of time can be one or two times of the average time a block is created. In some examples, a block can also be determined as infrequently accessed when no block data in the block is retrieved for the predetermined amount of time to execute smart contracts.

[0055] After identifying infrequently accessed blocks, the blockchain node 302 can perform ECC 314 of block data in the block body of each of the infrequently accessed blocks. ECC can be used for controlling errors or losses of data over unreliable transmissions by adding redundant bits (also referred to as redundancy) to the data. Redundant bits can be a complex function of many original information bits. The redundancy can allow errors or losses of data to be corrected without retransmission of the data. The original information

may or may not appear literally in the encoded output. ECC codes that include the unmodified original information in the encoded output are referred to as systematic ECC codes, while those that do not are referred to as non-systematic ECC codes. The maximum fractions of errors or of missing bits that can be corrected by ECC is determined by the design of the ECC code. Therefore, different error correction codes are suitable for different conditions. In general, a stronger ECC code induces more redundancy, which increases storage consumption of the code and reduces communication efficiency if the encoded information is to be transmitted.

[0056] One example ECC can be the erasure coding. Using the erasure coding, a message of k symbols can be encoded to a codeword with n symbols, where k and n are natural numbers, and $k < n$. The message can be recovered from a subset of the n -symbol codeword. The fraction $r = k/n$ is the code rate of the erasure code.

[0057] By using ECC, each of the blockchain nodes can store a portion of the encoded block data and retrieve the rest of the encoded block data from other blockchain nodes when needed. In some embodiments, the ECC can be performed when utilization rate of computational resource of the blockchain node 302 is lower than a predetermined value (e.g., 40%). As such, the interference with other computational operations on the blockchain node 302 can be reduced. In some embodiments, ECC can be performed when the usage of storage space of the blockchain node 302 is greater than or equal to a predetermined percentage, such that after ECC, some portions of the encoded block data can be deleted to free up storage space.

[0058] Again, using block 100 as an example, assuming that the blockchain node 302 determines the block 100 as an infrequently accessed block and performs ECC 314, the ECC encoded data can be divided into a plurality of datasets based on a data storage scheme. A data storage scheme can be provided as a set of computer-executable instructions that define where and/or how data is to be stored within the blockchain network. In some examples, the data storage scheme can be provided by a trusted node with proof of authority and agreed to by the blockchain nodes. In some examples, the data storage scheme can be agreed to by the blockchain nodes through consensus. Generally, the data storage scheme can include one or more predetermined rules for dividing the encoded data to a plurality of datasets based on the number of blockchain nodes in a blockchain network. The data storage scheme can also

include assignments of one or more datasets of the plurality of datasets to be stored or hashed by each of the blockchain nodes. To ensure data equality, the data storage scheme can include an assignment of at least one dataset to be stored by each blockchain node of the blockchain network.

[0059] In the example shown in FIG. 3, the encoded block data of block 100 is divided into four datasets, which are Data1, Data2, Data3, and Vdata1, each to be stored by one of the blockchain nodes 302, 304, 306, and 308. Vdata1 can represent the redundant bits of the ECC for error correction. Data1 is selected to be stored by the blockchain node 302 according to the data storage scheme. Data2, Data3, and Vdata1 are selected to be separately hashed 316 to generate hash values Dhash2, Dhash3, and Vhash1, respectively. In accordance with embodiments of this specification, the encoded data can be divided to more than four datasets when the blockchain network has more than four nodes. In some examples, each of the blockchain nodes can store more than one dataset and hash the rest of the datasets assigned to be stored by other nodes.

[0060] Referring now to FIG. 4, FIG. 4 depicts an example of a data storage scheme 400 in accordance with embodiments of this specification. As discussed earlier, Data1 is selected to be stored by the blockchain node 302 according to the data storage scheme 400. Based on the data storage scheme 400, blockchain node 304 stores Data2 and separately hashes Data1, Data3, and Vdata1 to generate hash values Dhash1, Dhash3, and Vhash1, respectively. Blockchain node 306 stores Data3 and separately hashes Data1, Data2, and Vdata1 to generate hash values Dhash1, Dhash2 and Vhash1, respectively. Blockchain node 308 stores Vdata1 and separately hashes Data1, Data2, and Vdata3 to generate hash values Dhash1, Dhash2 and Dhash3, respectively.

[0061] Referring back to FIG. 3, because the hash values correspond to encoded datasets of the same block, they can be indexed by a block ID of the block. For example, the blockchain node 302 can index Data1, Dhash1, Dhash2, and Vhash1 associated with block 100 with a block ID 100. As such, the blockchain node 302 can use the indexed block ID to map the hash values to their corresponding blocks. A more detailed example of indexing the datasets and hash values is discussed in the description of FIG. 6.

[0062] It is to be understood that other data storage schemes can be made for the blockchain nodes 302, 304, 306, and 308, according to the data storage scheme. In some

examples, the encoded block data of block 100 can be divided to more than four datasets. It is to be understood that other data storage schemes can be made for the blockchain nodes 502, 504, 506, and 508, according to the data storage scheme.

[0063] After generating and storing Dhash2, Dhash3, and Vhash1, the blockchain node 302 can delete Data2, Data3, and Vdata1 from storage to save storage space. As such, for each block, the blockchain node 302 only stores one ECC encoded dataset (i.e., Data1) and three hash values (i.e., Dhash2, Dhash3, and Vhash1), instead of the entire block. As such, storage space can be significantly reduced. Similar to block 100, the encoding and hashing process can be performed for other infrequently accessed blocks that are stored by the blockchain nodes 304, 306, and 308.

[0064] When the blockchain node 302 determines that block data of the block 100 is needed for executing a smart contract, it can retrieve Data2, Data3, and Vdata1 from blockchain nodes 304, 306, and 308, respectively, according to the data storage scheme. To retrieve datasets from other blockchain nodes 304, 306, and 308, blockchain node 302 can send hash values corresponding to the datasets to be retrieved according to the data storage scheme.

[0065] For example, to retrieve Data2, the blockchain node 302 can send Dhash2 to the blockchain node 304. If the blockchain node 304 has Data2 stored, it can send the Data2 back to the blockchain node 302 in response to receiving the Dhash2. After receiving the Data2 from the blockchain node 304, the blockchain node 302 can hash the received dataset and compare the hash value with Dhash2. If the hash value is the same as Dhash2, the blockchain node 302 can determine that the received dataset is authentic. Otherwise, the received dataset is determined to be unauthentic. When the received dataset is determined as unauthentic, the blockchain node 302 can report the blockchain node 304 as a faulty node (or a Byzantine node). If the percentage of unauthentic data received by the blockchain node 302 is less than or equal to the maximum fraction of erroneous or missing bits that can be corrected by the ECC, block 100 can be recovered from the locally stored and received datasets.

[0066] As described earlier, blockchain networks can store different types of data such as state data, block data, and index data. State data are often stored as a content-addressed state tree, such as the MPT or the fixed depth Merkle tree (FDMT). Content-addressed state trees are incremental in nature. That is, changes of account states are reflected by adding new tree

structures instead of only updating values of the existing state tree. Therefore, the content-addressed state trees can grow very large in size when blocks are continuously added to the blockchain. Under the FDMT storage scheme, state data can be separated into current state data associated with the current block and historic state data associated with all blocks of the blockchain. Most data in the FDMT are infrequently used historic state data. Storing all historic state data in every consensus node can be quite inefficient in terms of storage resource usage.

[0067] In some embodiments, similar to encoding and sharing block data, ECC such as erasure coding can be used to encode the historic state data. Each consensus node in the blockchain network stores only a portion of the historic state data and retrieves the rest of the historic state data from other nodes to reduce storage consumption. By sharing ECC encoded historic state data instead of the original historic state data, even if unauthentic data exists or data loss occurs, the original historic state data can be recovered, as long as the percentage of unauthentic data or data loss is less than or equal to the maximum fraction of erroneous or missing bits that can be corrected by the ECC.

[0068] FIG. 5 depicts another example of a block data encoding and hashing process 500 in accordance with embodiments of this specification. In this example, a blockchain network of four blockchain nodes is depicted, which are blockchain nodes 502, 504, 506, and 508. Using blockchain node 502 as an example to illustrate the encoding and hashing process 500, when new block data are added to the block 512, the blockchain node 502 can perform ECC 514 to encode the block data. As compared to the encoding and hashing process 300 discussed in the description of FIG. 3, the blockchain node 502 performs ECC on the block data as they are written to a block. As such, the blockchain node 502 does not need to store the entire block, but can instead, store a selected portion of the ECC encoded block data and hash values corresponding to the rest of the encoded block data based on the data storage scheme. This encoding and hashing process 500 can be especially suitable for scenarios when blockchain node 502 has low disk space.

[0069] In some embodiments, instead of storing data as blocks, the blockchain node 502 can store a write-ahead log (WAL) file or other similar roll-forward journal files. The WAL file can record block data that have been committed but not yet stored by the blockchain node 502. Using the WAL file, the original blockchain data can be preserved in the database file,

while changes of the blockchain data can be written into a separate WAL file. A commit to roll-forward with the changes can happen without ever writing to the original blockchain data. This arrangement allows continued operations of the blockchain data while changes are committed into the WAL file. By using the WAL file to store changes made through the encoding and hashing process 500, the blockchain node 502 can indicate that it has the block data for consensus, while performing the ECC in the background when appropriate. As such, the ECC can be performed when utilization rate of computational resource of the blockchain node 302 is low, in order to reduce the impact on computational efficiency or latency of the consensus process.

[0070] In some embodiments, the blockchain node 502 can store the block data in a buffer. The blockchain node 502 can perform ECC to the block data stored in the buffer when the size of the data is greater than a predetermined threshold or when the buffer is full. After performing ECC, the blockchain node 502 can follow the encoding and hashing process 500 to store encoded block data and hash values, as discussed in the description below.

[0071] Using block 100 as an example again, after performing the ECC, the encoded block data can be divided into a plurality of datasets based on the data storage scheme. Similar to the example discussed in the description of FIG. 3, the encoded block data of block 100 can be divided into four datasets, which are Data1, Data2, Data3, and Vdata1, each to be stored by one of the blockchain nodes 502, 504, 506, and 508. Vdata1 can represent the redundant bits of the ECC. Data1 is selected to be stored by the blockchain node 502 according to the data storage scheme. Data2, Data3, and Vdata1 are selected to be separately hashed 516 to generate hash values Dhash2, Dhash3, and Vhash1, respectively.

[0072] The hash values can be indexed by a block ID of a corresponding block of the hash values. For example, the blockchain node 502 can index Data1, Dhash1, Dhash2, and Vhash1 associated with block 100 with a block ID 100. As such, the blockchain node 502 can use the indexed block ID to map the hash values to their corresponding blocks. A more detailed example of indexing the datasets and hash values is discussed in the description of FIG. 6.

[0073] It is to be understood that other data storage schemes can be made for the one or more blockchain nodes 502, 504, 506, and 508, according to the data storage scheme. For

example, the encoded block data of block 100 can be divided into more than four datasets. Each of the blockchain nodes 502, 504, 506, and 508 can store more than one dataset and hash the rest of the datasets stored by other nodes.

[0074] After generating Dhash2, Dhash3, and Vhash1, the blockchain node 502 can store Data1, Dhash2, Dhash3, and Vhash1 and delete Data2, Data3, and Vdata1 from storage to save storage space. As such, for each block of the blockchain, the blockchain node 502 only stores one dataset (i.e., Data1) and three hash values (i.e., Dhash2, Dhash3, and Vhash1) of the ECC encoded block data instead of the original block data to save on storage space. When the blockchain node 502 determines that block data of the block 100 is needed for executing a smart contract, it can retrieve Data2, Data3, and Vdata1 from blockchain nodes 504, 506, and 508, respectively, according to the data storage scheme.

[0075] To retrieve datasets from other blockchain nodes 504, 506, and 508, blockchain node 502 can send hash values corresponding to the datasets to be retrieved according to the data storage scheme. For example, to retrieve Data2, the blockchain node 502 can send Dhash2 to the blockchain node 504. If the blockchain node 504 has Data2 stored, it can send the Data2 back to the blockchain node 502 in response to receiving the Dhash2. After receiving the Data2 from the blockchain node 504, the blockchain node 502 can hash the received dataset and compare the hash value with Dhash2. If the hash value is the same as Dhash2, the blockchain node 502 can determine that the received dataset is authentic. Otherwise, the received dataset can be determined as unauthentic. When the received dataset is determined as unauthentic, the blockchain node 502 can report the blockchain node 504 as a faulty node (or a Byzantine node). If the percentage of unauthentic data received by the blockchain node 502 is less than or equal to the maximum fraction of erroneous or missing bits that can be corrected by the ECC, block 100 can be recovered from the locally stored and received datasets.

[0076] As discussed earlier, by performing the encoding and hashing process, blockchain data can be ECC encoded and divided into a plurality of datasets. To save on storage space, each blockchain node can store one or more of the plurality of datasets and hash values of rest of the datasets based on a data storage scheme. The stored datasets and hash values can be indexed with Block IDs in order for a blockchain node to retrieve datasets from other nodes to recover original data.

[0077] FIG. 6 depicts an example of a process 600 of shared blockchain data storage based on NVNs in accordance with embodiments of this specification. VNs (referred to here as validating peers) and NVNs (referred to here as non-validating peers) are two types of peers that run distributed ledger protocol for a permissioned blockchain platform, such as a Hyperledger fabric. A permissioned blockchain network can be a private or consortium blockchain network, where every member node maintains an identity in the network. End user or client node needs to be authorized and authenticated before using the network. In other words, the permissioned blockchain controls who can participate in validation and the distributed ledger protocol. The VNs can be responsible for running consensus, validating transactions, and maintaining the distributed ledger. The NVNs can function as a proxy to connect client nodes (e.g., that issue transactions) to VNs. The NVNs do not execute transactions, but may verify them.

[0078] In the illustrated example, it is assumed that the blockchain network 610 is a permissioned blockchain network that includes VNs 602, 604, 606 and 608. The blockchain network 610 also includes an NVN cluster 610 formed by NVNs 612, 614, 616, 618. In some embodiments, the transactions can be generated from a client node 630. The client node 630 can be a piece of computer hardware or software. In some embodiments, the client node 310 can send the transactions 632 to the NVN cluster 610 or one of the NVNs 612, 614, 616, 618. After receiving the transactions, the NVN cluster 610 can forward the transactions 622 to the VNs 602, 604, 606, and 608 for running consensus. After receiving the transactions, the VNs 602, 604, 606, and 608 can validate the transactions by engaging in a consensus process (e.g., PBFT) to generate a block (e.g., block 100) based on at least a subset of the received transactions and append it to the blockchain.

[0079] After block 100 is generated and appended to the blockchain, a VN 602 can send the block 100 to the NVN cluster 610 for storage. In some embodiments, an NVN can also automatically pull the blocks from a VN. The blocks can be pulled periodically or after a predetermined number of blocks (e.g., 50 blocks) are generated. After receiving one or more blocks from the VN 602, the NVNs 612, 614, 616, and 618 can then engage in an encoding and hashing process as discussed in the descriptions of FIGS. 3 and 5. Each NVN can encode the one or more blocks using ECC 626 (e.g., erasure coding) to generate one or more ECC encoded blocks. In some embodiments, the ECC can be performed when a utilization rate of

the computational resource of the NVN is less than or equal to a predetermined value (e.g., 50%) to maintain the computational efficiency of the NVN. In some embodiments, the ECC can be performed when usage of storage space of the NVN is greater than or equal to a predetermined percentage (e.g., 80%).

[0080] After generating the ECC encoded block, an NVN (here, NVN 612) can divide the ECC encoded block into four datasets Data1, Data2, Data3, and Vdata1, based on a data storage scheme 400 as discussed in the description of FIG. 4. The NVN 612 can then store Data1 based on the assignments provided in the data storage scheme, and can perform hashing 628 to the remainder of the datasets to generate hash values Dhash2, Dhash3, and Vhash1. After storing the dataset (Data1) and hash values (Dhash2, Dhash3, and Vhash1), the NVN 612 can discard the ECC encoded block and the remainder of the datasets, Data2, Data3, and Vdata1 to save storage space.

[0081] In some embodiments, the NVN 612 can confirm with other NVNs 614, 616, and 618 whether they have stored their assigned datasets and hash values according to the data storage scheme. After receiving confirmations from other NVNs 614, 616, and 618, the NVN 612 can send a notification to the plurality of VNs 602, 604, 606, and 608 that the NVNs 612, 614, 616, and 618 have completed the encoding and hashing process. It is then safe for the VNs 602, 604, 606, and 608 to delete the block 100 to save storage space.

[0082] In some embodiments, a client node 630 or a VN can send a request to the NVN 612 for verifying whether a transaction exists on the blockchain. The NVN 612 receiving the request can retrieve datasets stored by other NVNs 614, 616, and 618 based on the index and decode the datasets to recover the corresponding blocks. After the blocks are recovered, the NVN 612 can verify whether the transaction is included in the blocks.

[0083] In some embodiments, a VN can send a request to the NVN 612 for retrieving transactions to execute smart contracts. The NVN 612 receiving the request can determine that the transactions are included in one or more blocks saved as datasets by the NVN cluster 610. The NVN 612 can then retrieve datasets stored by other NVNs 614, 616, and 618 based on the index and decode the datasets to recover the corresponding one or more blocks. Afterwards, the NVN node 612 can locate the transactions and send the transactions to the VN.

[0084] In some embodiments, the NVN can store blockchain data in a tiered storage. With tiered storage, data can be placed in a hierarchy of tiers according to its relative business value. For example, the business value of data can be determined based on how often it is accessed. Generally, more valuable data can be saved in faster or more expensive storage media (i.e., tiers). Because blockchain data in newer blocks are more likely to be accessed than older blocks, the older blocks can be stored in a slower storage medium as compared to newer blocks to reduce the cost of storage.

[0085] FIG. 7 depicts an example of a process 700 that can be executed in accordance with embodiments of this specification. For convenience, the process 700 will be described as being performed by a NVN. The NVN can be a computer or a system of one or more computers, located in one or more locations, and programmed appropriately in accordance with this specification. For example, the NVN can be a computing device in a computing system, e.g., the computing system 106, 108 of FIG. 1, appropriately programmed, can perform the process 700.

[0086] At 702, the NVN receives transactions issued by a client node communicably coupled to a permissioned blockchain network that includes the NVN. At 704, the NVN sends the transactions to a plurality of validating nodes for validating the transactions and running consensus.

[0087] At 706, the NVN receives one or more blocks generated based on consensus of at least a portion of the transactions by the plurality of validating nodes. At 708, the NVN encodes the one or more blocks using ECC to generate one or more encoded blocks. In some cases, the ECC is performed when utilization rate of computational resource of the NVN is less than or equal to a predetermined value. In some cases, the ECC is performed when usage of storage space of the NVN is greater than or equal to a predetermined percentage. In some cases, the ECC is erasure coding performed by adding redundant bits to the plurality of blocks.

[0088] At 710, for each encoded block of the one or more encoded blocks, the NVN divides the encoded block into a plurality of datasets based on a data storage scheme associated with the one or more blocks, wherein the data storage scheme provides assignments of the plurality of datasets to a plurality of NVNs. In some cases, the NVN stores blockchain data in a tiered storage that a newer block is stored in a same or faster tier

of the tiered storage than an older block. At 712, the NVN stores at least one of the plurality of datasets based on the assignments provided in the data storage scheme.

[0089] In some cases, the NVN further performs: receiving confirmations from a remainder of the plurality of NVNs other than the NVN that at least one of the plurality of datasets are stored by the corresponding NVN based on the assignments provided in the data storage scheme; and sending a notification to the plurality of validating nodes to delete the one or more blocks.

[0090] In some cases, the NVN further performs: receiving a request from the client node for verifying a transaction included in the one or more blocks; retrieving, based on the index, datasets associated with the one or more blocks from the remainder of the plurality of NVNs; decoding the one or more blocks based on the datasets; and verifying that the transaction exists if the transaction is included in the one or more blocks.

[0091] In some cases, the NVN further performs: receiving a request from a validating node for retrieving a transaction included in the one or more blocks; retrieving, based on the index, datasets associated with the one of the one or more blocks from the remainder of the plurality of NVNs; decoding the one of the one or more blocks based on the datasets; and sending the transaction to the validating node.

[0092] In some cases, for each encoded block of the plurality of encoded blocks, the NVN further performs: hashing a remainder of the plurality of datasets other than the at least one of the plurality of datasets to generate hash values; storing the hash values; and discard the one or more blocks and the remainder of the plurality of datasets.

[0093] FIG. 8 is a diagram of an example of modules of an apparatus 800 in accordance with embodiments of this specification. The apparatus 800 can be an example of an embodiment of an NVN configured to storing and processing blockchain data. The apparatus 800 can correspond to the embodiments described above, and the apparatus 800 includes the following: a receiving module 802 that receives transactions issued by a client node communicably coupled to a permissioned blockchain network; a sending module 804 that sends the transactions to a plurality of validating nodes for validating the transactions and running consensus; the receiving module 802 that receives one or more blocks generated based on consensus of at least a portion of the transactions by the plurality of validating nodes; an encoding module 806 that encodes the one or more blocks using ECC to generate

one or more encoded blocks; and for each encoded block of the one or more encoded blocks: a dividing module 808 that divides the encoded block into a plurality of datasets based on a data storage scheme associated with the one or more blocks, wherein the data storage scheme provides assignments of the plurality of datasets to a plurality of NVNs; a storing module 810 that stores at least one of the plurality of datasets based on the assignments provided in the data storage scheme.

[0094] In some embodiments, the receiving module 802 further receives confirmations from a remainder of the plurality of NVNs other than the NVN that at least one of the plurality of datasets are stored by the corresponding NVN based on the assignments provided in the data storage scheme, and the sending module 804 further sends a notification to the plurality of validating nodes to delete the one or more blocks.

[0095] In some embodiments, the NVN stores blockchain data in a tiered storage that a newer block is stored in a same or faster tier of the tiered storage than an older block.

[0096] In some embodiments, the receiving module 802 further receives a request from the client node for verifying a transaction included in the one or more blocks. The apparatus further includes a retrieving submodule to retrieve, based on the index, datasets associated with the one or more blocks from the remainder of the plurality of NVNs, a decoding submodule to decode the one or more blocks based on the datasets, and a verification submodule to verify that the transaction exists if the transaction is included in the one or more blocks.

[0097] In some embodiments, the receiving module 802 further receives a request from a validating node for retrieving a transaction included in the one or more blocks. The retrieving submodule further retrieves, based on the index, datasets associated with the one of the one or more blocks from the remainder of the plurality of NVNs. The decoding submodule further decodes the one of the one or more blocks based on the datasets, and the sending module 804 further sends the transaction to the validating node.

[0098] In some embodiments, the ECC is performed when utilization rate of computational resource of the NVN is less than or equal to a predetermined value.

[0099] In some embodiments, the apparatus 800 further includes a hashing submodule. For each encoded block of the plurality of encoded blocks, the hashing submodule hashes a remainder of the plurality of datasets other than the at least one of the plurality of datasets to

generate hash values. The storing module 812 further stores the hash values and discards the one or more blocks and the remainder of the plurality of datasets.

[0100] In some embodiments, the ECC is performed when usage of storage space of the NVN is greater than or equal to a predetermined percentage.

[0101] In some embodiments, the ECC is erasure coding performed by adding redundant bits to the plurality of blocks.

[0102] The system, apparatus, module, or unit illustrated in the previous embodiments can be implemented by using a computer chip or an entity, or can be implemented by using a product having a certain function. A typical embodiment device is a computer, and the computer can be a personal computer, a laptop computer, a cellular phone, a camera phone, a smartphone, a personal digital assistant, a media player, a navigation device, an email receiving and sending device, a game console, a tablet computer, a wearable device, or any combination of these devices.

[0103] For an embodiment process of functions and roles of each module in the apparatus, references can be made to an embodiment process of corresponding steps in the previous method. Details are omitted here for simplicity.

[0104] Because an apparatus embodiment basically corresponds to a method embodiment, for related parts, references can be made to related descriptions in the method embodiment. The previously described apparatus embodiment is merely an example. The modules described as separate parts may or may not be physically separate, and parts displayed as modules may or may not be physical modules, may be located in one position, or may be distributed on a number of network modules. Some or all of the modules can be selected based on actual demands to achieve the objectives of the solutions of the specification. A person of ordinary skill in the art can understand and implement the embodiments of the present application without creative efforts.

[0105] Referring again to FIG. 8, it can be interpreted as illustrating an internal functional module and a structure of a blockchain node. An execution body in essence can be an electronic device, and the electronic device includes the following: one or more processors; and one or more computer-readable memories configured to store an executable instruction of the one or more processors. In some embodiments, the one or more computer-readable memories are coupled to the one or more processors and have programming

instructions stored thereon that are executable by the one or more processors to perform algorithms, methods, functions, processes, flows, and procedures as described in this specification. This specification also provides one or more non-transitory computer-readable storage media coupled to one or more processors and having instructions stored thereon which, when executed by the one or more processors, cause the one or more processors to perform operations in accordance with embodiments of the methods provided herein.

[0106] This specification further provides a system for implementing the methods provided herein. The system includes one or more processors, and a computer-readable storage medium coupled to the one or more processors having instructions stored thereon which, when executed by the one or more processors, cause the one or more processors to perform operations in accordance with embodiments of the methods provided herein.

[0107] Embodiments of the subject matter and the actions and operations described in this specification can be implemented in digital electronic circuitry, in tangibly-embodied computer software or firmware, in computer hardware, including the structures disclosed in this specification and their structural equivalents, or in combinations of one or more of them. Embodiments of the subject matter described in this specification can be implemented as one or more computer programs, e.g., one or more modules of computer program instructions, encoded on a computer program carrier, for execution by, or to control the operation of, data processing apparatus. For example, a computer program carrier can include one or more computer-readable storage media that have instructions encoded or stored thereon. The carrier may be a tangible non-transitory computer-readable medium, such as a magnetic, magneto optical, or optical disk, a solid state drive, a random access memory (RAM), a read-only memory (ROM), or other types of media. Alternatively, or in addition, the carrier may be an artificially generated propagated signal, e.g., a machine-generated electrical, optical, or electromagnetic signal that is generated to encode information for transmission to suitable receiver apparatus for execution by a data processing apparatus. The computer storage medium can be or be part of a machine-readable storage device, a machine-readable storage substrate, a random or serial access memory device, or a combination of one or more of them. A computer storage medium is not a propagated signal.

[0108] A computer program, which may also be referred to or described as a program, software, a software application, an app, a module, a software module, an engine, a script, or

code, can be written in any form of programming language, including compiled or interpreted languages, or declarative or procedural languages; and it can be deployed in any form, including as a stand-alone program or as a module, component, engine, subroutine, or other unit suitable for executing in a computing environment, which environment may include one or more computers interconnected by a data communication network in one or more locations.

[0109] A computer program may, but need not, correspond to a file in a file system. A computer program can be stored in a portion of a file that holds other programs or data, e.g., one or more scripts stored in a markup language document, in a single file dedicated to the program in question, or in multiple coordinated files, e.g., files that store one or more modules, sub programs, or portions of code.

[0110] Processors for execution of a computer program include, by way of example, both general- and special-purpose microprocessors, and any one or more processors of any kind of digital computer. Generally, a processor will receive the instructions of the computer program for execution as well as data from a non-transitory computer-readable medium coupled to the processor.

[0111] The term “data processing apparatus” encompasses all kinds of apparatuses, devices, and machines for processing data, including by way of example a programmable processor, a computer, or multiple processors or computers. Data processing apparatus can include special-purpose logic circuitry, e.g., an FPGA (field programmable gate array), an ASIC (application specific integrated circuit), or a GPU (graphics processing unit). The apparatus can also include, in addition to hardware, code that creates an execution environment for computer programs, e.g., code that constitutes processor firmware, a protocol stack, a database management system, an operating system, or a combination of one or more of them.

[0112] The processes and logic flows described in this specification can be performed by one or more computers or processors executing one or more computer programs to perform operations by operating on input data and generating output. The processes and logic flows can also be performed by special-purpose logic circuitry, e.g., an FPGA, an ASIC, or a GPU, or by a combination of special-purpose logic circuitry and one or more programmed computers.

[0113] Computers suitable for the execution of a computer program can be based on general or special-purpose microprocessors or both, or any other kind of central processing unit. Generally, a central processing unit will receive instructions and data from a read only memory or a random access memory or both. Elements of a computer can include a central processing unit for executing instructions and one or more memory devices for storing instructions and data. The central processing unit and the memory can be supplemented by, or incorporated in, special-purpose logic circuitry.

[0114] Generally, a computer will also include, or be operatively coupled to receive data from or transfer data to one or more storage devices. The storage devices can be, for example, magnetic, magneto optical, or optical disks, solid state drives, or any other type of non-transitory, computer-readable media. However, a computer need not have such devices. Thus, a computer may be coupled to one or more storage devices, such as, one or more memories, that are local and/or remote. For example, a computer can include one or more local memories that are integral components of the computer, or the computer can be coupled to one or more remote memories that are in a cloud network. Moreover, a computer can be embedded in another device, e.g., a mobile telephone, a personal digital assistant (PDA), a mobile audio or video player, a game console, a Global Positioning System (GPS) receiver, or a portable storage device, e.g., a universal serial bus (USB) flash drive, to name just a few.

[0115] Components can be “coupled to” each other by being commutatively such as electrically or optically connected to one another, either directly or via one or more intermediate components. Components can also be “coupled to” each other if one of the components is integrated into the other. For example, a storage component that is integrated into a processor (e.g., an L2 cache component) is “coupled to” the processor.

[0116] To provide for interaction with a user, embodiments of the subject matter described in this specification can be implemented on, or configured to communicate with, a computer having a display device, e.g., a LCD (liquid crystal display) monitor, for displaying information to the user, and an input device by which the user can provide input to the computer, e.g., a keyboard and a pointing device, e.g., a mouse, a trackball or touchpad. Other kinds of devices can be used to provide for interaction with a user as well; for example, feedback provided to the user can be any form of sensory feedback, e.g., visual feedback, auditory feedback, or tactile feedback; and input from the user can be received in any form,

including acoustic, speech, or tactile input. In addition, a computer can interact with a user by sending documents to and receiving documents from a device that is used by the user; for example, by sending web pages to a web browser on a user's device in response to requests received from the web browser, or by interacting with an app running on a user device, e.g., a smartphone or electronic tablet. Also, a computer can interact with a user by sending text messages or other forms of message to a personal device, e.g., a smartphone that is running a messaging application, and receiving responsive messages from the user in return.

[0117] This specification uses the term “configured to” in connection with systems, apparatus, and computer program components. For a system of one or more computers to be configured to perform particular operations or actions means that the system has installed on it software, firmware, hardware, or a combination of them that in operation cause the system to perform the operations or actions. For one or more computer programs to be configured to perform particular operations or actions means that the one or more programs include instructions that, when executed by data processing apparatus, cause the apparatus to perform the operations or actions. For special-purpose logic circuitry to be configured to perform particular operations or actions means that the circuitry has electronic logic that performs the operations or actions.

[0118] While this specification contains many specific embodiment details, these should not be construed as limitations on the scope of what is being claimed, which is defined by the claims themselves, but rather as descriptions of features that may be specific to particular embodiments. Certain features that are described in this specification in the context of separate embodiments can also be realized in combination in a single embodiment. Conversely, various features that are described in the context of a single embodiments can also be realized in multiple embodiments separately or in any suitable subcombination. Moreover, although features may be described above as acting in certain combinations and even initially be claimed as such, one or more features from a claimed combination can in some cases be excised from the combination, and the claim may be directed to a subcombination or variation of a subcombination.

[0119] Similarly, while operations are depicted in the drawings and recited in the claims in a particular order, this should not be understood as requiring that such operations be performed in the particular order shown or in sequential order, or that all illustrated

operations be performed, to achieve desirable results. In certain circumstances, multitasking and parallel processing may be advantageous. Moreover, the separation of various system modules and components in the embodiments described above should not be understood as requiring such separation in all embodiments, and it should be understood that the described program components and systems can generally be integrated together in a single software product or packaged into multiple software products.

[0120] Particular embodiments of the subject matter have been described. Other embodiments are within the scope of the following claims. For example, the actions recited in the claims can be performed in a different order and still achieve desirable results. As one example, the processes depicted in the accompanying figures do not necessarily require the particular order shown, or sequential order, to achieve desirable results. In some cases, multitasking and parallel processing may be advantageous.

CLAIMS

1. A computer-implemented method for processing blockchain data by a non-validating node (NVN) of a permissioned blockchain network, the method comprising:

receiving transactions issued by a client node communicably coupled to a permissioned blockchain network;

sending the transactions to a plurality of validating nodes for validating the transactions and running consensus;

receiving one or more blocks generated based on consensus of at least a portion of the transactions by the plurality of validating nodes;

encoding the one or more blocks using error correction coding (ECC) to generate one or more encoded blocks; and

for each encoded block of the one or more encoded blocks:

dividing the encoded block into a plurality of datasets based on a data storage scheme associated with the one or more blocks, wherein the data storage scheme provides assignments of the plurality of datasets to a plurality of non-validating nodes (NVNs); and

storing at least one of the plurality of datasets based on the assignments provided in the data storage scheme.

2. The method of claim 1, further comprising:

receiving confirmations from a remainder of the plurality of NVNs other than the NVN that at least one of the plurality of datasets are stored by the corresponding NVN based on the assignments provided in the data storage scheme; and

sending a notification to the plurality of validating nodes to delete the one or more blocks.

3. The method of claim 1, wherein the NVN stores blockchain data in a tiered storage that a newer block is stored in a same or faster tier of the tiered storage than an older block.

4. The method of claim 1, further comprising:

receiving a request from the client node for verifying a transaction included in the one or more blocks;

retrieving, based on the index, datasets associated with the one or more blocks from the remainder of the plurality of NVNs;

decoding the one or more blocks based on the datasets; and

verifying that the transaction exists if the transaction is included in the one or more blocks.

5. The method of claim 1, further comprising:

receiving a request from a validating node for retrieving a transaction included in the one or more blocks;

retrieving, based on the index, datasets associated with the one of the one or more blocks from the remainder of the plurality of NVNs;

decoding the one of the one or more blocks based on the datasets; and

sending the transaction to the validating node.

6. The method of any previous claim, wherein the ECC is performed when utilization rate of computational resource of the NVN is less than or equal to a predetermined value.

7. The method of any previous claim, further comprising:

for each encoded block of the plurality of encoded blocks:

hashing a remainder of the plurality of datasets other than the at least one of the plurality of datasets to generate hash values;

storing the hash values; and

discard the one or more blocks and the remainder of the plurality of datasets.

8. The method of any previous claim, wherein the ECC is performed when usage of storage space of the NVN is greater than or equal to a predetermined percentage.

9. The method of any previous claim, wherein the ECC is erasure coding performed by adding redundant bits to the plurality of blocks.

10. A system for processing blockchain data, comprising:
 - one or more processors; and
 - one or more computer-readable memories coupled to the one or more processors and having instructions stored thereon that are executable by the one or more processors to perform the method of any of claims 1 to 9.
11. An apparatus for processing blockchain data, the apparatus comprising a plurality of modules for performing the method of any of claims 1 to 9.