

INDEX STRUCTURE FOR BLOCKCHAIN LEDGER

TECHNICAL FIELD

[0001] This application generally relates to storing data via a blockchain, and more particularly, to an index structure which stores key values from the blockchain ledger and which can be accessed for non-critical queries to reduce phantom data items.

BACKGROUND

[0002] A centralized database stores and maintains data in a single database (e.g., a database server) at one location. This location is often a central computer, for example, a desktop central processing unit (CPU), a server CPU, or a mainframe computer. Information stored on a centralized database is typically accessible from multiple different points. Multiple users or client workstations can work simultaneously on the centralized database, for example, based on a client/server configuration. A centralized database is easy to manage, maintain, and control, especially for purposes of security because of its single location. Within a centralized database, data redundancy is minimized as a single storing place of all data also implies that a given set of data only has one primary record.

[0003] Meanwhile, blockchain systems store data on an immutable ledger, provide distributed and decentralized access to the immutable ledger through non-trusting participants, establish consensus requirements for agreement between the non-trusting participants such that no one entity can change the immutable ledger without agreement from others, invoke smart contracts, and the like. A blockchain is formed by a network of participants which agree to add a block (with data stored therein) to the immutable ledger. Before being added, the block is linked to a previous block on the immutable ledger thereby forming a chain. This immutable and incorruptible nature of blockchain makes it safe from falsified information and hacks. The decentralized nature also gives it the unique quality of being trustless, in that parties do not need to establish trust before they can transact safely. Through blockchain, new and improved opportunities for data storage, communication, and security are emerging.

SUMMARY

[0004] One example embodiment provides an apparatus that includes one or more of a storage configured to store an index structure that comprises an index of keys from a blockchain ledger, where the keys are stored as nodes in the index structure, and a processor configured to one or more of receive a blockchain request for data stored on the blockchain ledger, read a set of keys of a non-critical query included in the blockchain request from the nodes in the index structure, and generate and store a read set for the blockchain request which does not include values for the set of keys of the non-critical query

[0005] Preferably, the present invention provides an apparatus wherein the index structure comprises a binary tree in which the keys are arranged as nodes in a hierarchical manner.

[0006] Preferably, the present invention provides an apparatus wherein the processor is further configured to build the index structure based on key values that are added to the blockchain ledger.

[0007] Preferably, the present invention provides an apparatus wherein the processor is configured to read the set of keys one-by-one from the nodes in the index structure via a next function.

[0008] Preferably, the present invention provides an apparatus wherein the processor is further configured to identify a query of the blockchain request as the non-critical query when the query does not impact a validity of the blockchain request.

[0009] Preferably, the present invention provides an apparatus wherein the processor is further configured to validate the blockchain request based on a previous read set for the blockchain request which does not include the values for the set of keys read from the index structure.

[0010] Preferably, the present invention provides an apparatus wherein the index structure is implemented within logic of chaincode of a blockchain peer node.

[0011] Preferably, the present invention provides an apparatus wherein the keys in the index structure correspond to unique values stored on a blockchain of the blockchain ledger.

[0012] Another example embodiment provides a method that includes one or more of storing an index structure that comprises an index of keys from a blockchain ledger, where the keys are stored as nodes in the index structure, receiving a blockchain request for data stored on the blockchain ledger, reading a set of keys of a non-critical query included in the blockchain request from the nodes in the index structure, and generating and storing a read set for the blockchain request which does not include values for the set of keys of the non-critical query.

[0013] Preferably, the present invention provides a method wherein the index structure comprises a binary tree in which the keys are arranged as nodes in a hierarchical manner.

[0014] Preferably, the present invention provides a method further comprising building the index structure based on key values that are added to the blockchain ledger.

[0015] Preferably, the present invention provides a method wherein the reading comprises reading the set of keys one-by-one from the nodes in the index structure via a next function.

[0016] Preferably, the present invention provides a method further comprising identifying a query of the blockchain request as the non-critical query when the query does not impact a validity of the blockchain request.

[0017] Preferably, the present invention provides a method further comprising validating the blockchain request based on a previous read set for the blockchain request which does not include the values for the set of keys read from the index structure.

[0018] Preferably, the present invention provides a method wherein the index structure is implemented within logic of chaincode of a blockchain peer node.

[0019] Preferably, the present invention provides a method wherein the keys in the index structure correspond to unique values stored on a blockchain of the blockchain ledger.

[0020] Viewed from another aspect, the present invention provides a non-transitory computer readable medium comprising instructions, that when read by a processor, cause the processor to perform a method comprising: storing an index structure that comprises an index of keys from a blockchain ledger, where the keys are stored as nodes in the index structure; receiving a blockchain request for data stored on the blockchain ledger; reading a set of keys of a non-critical query included in the blockchain request from the nodes in the index structure; and generating and storing a read set for the blockchain request which does not include values for the set of keys of the non-critical query.

[0021] Preferably, the present invention provides a non-transitory computer readable medium wherein the index structure comprises a binary tree in which the keys are arranged as nodes in a hierarchical manner.

[0022] Preferably, the present invention provides a non-transitory computer readable medium wherein the method further comprises building the index structure based on key values that are added to the blockchain ledger.

[0023] Preferably, the present invention provides a non-transitory computer readable medium wherein the reading comprises reading the set of keys one-by-one from the nodes in the index structure via a next function.

BRIEF DESCRIPTION OF THE DRAWINGS

[0024] FIG. 1A is a diagram illustrating a blockchain network implementing a binary index structure according to example embodiments.

- [0025]** FIG. 1B is a diagram illustrating an example of the binary index structure of FIG. 1A, according to example embodiments.
- [0026]** FIG. 2A is a diagram illustrating an example blockchain architecture configuration, according to example embodiments.
- [0027]** FIG. 2B is a diagram illustrating a blockchain transactional flow among nodes, according to example embodiments.
- [0028]** FIG. 3A is a diagram illustrating a permissioned network, according to example embodiments.
- [0029]** FIG. 3B is a diagram illustrating another permissioned network, according to example embodiments.
- [0030]** FIG. 3C is a diagram illustrating a permissionless network, according to example embodiments.
- [0031]** FIG. 4A is a diagram illustrating a process of querying a blockchain ledger via a range query according to example embodiments.
- [0032]** FIG. 4B is a diagram illustrating a process of querying an index structure on top of a blockchain ledger according to example embodiments.
- [0033]** FIG. 5 is a diagram illustrating a method of accessing an index structure on top of a blockchain ledger according to example embodiments.
- [0034]** FIG. 6A is a diagram illustrating an example system configured to perform one or more operations described herein, according to example embodiments.
- [0035]** FIG. 6B is a diagram illustrating another example system configured to perform one or more operations described herein, according to example embodiments.
- [0036]** FIG. 6C is a diagram illustrating a further example system configured to utilize a smart contract, according to example embodiments.
- [0037]** FIG. 6D is a diagram illustrating yet another example system configured to utilize a blockchain, according to example embodiments.
- [0038]** FIG. 7A is a diagram illustrating a process of a new block being added to a distributed ledger, according to example embodiments.
- [0039]** FIG. 7B is a diagram illustrating data contents of a new data block, according to example embodiments.
- [0040]** FIG. 7C is a diagram illustrating a blockchain for digital content, according to example embodiments.
- [0041]** FIG. 7D is a diagram illustrating a block which may represent the structure of blocks in the blockchain, according to example embodiments.
- [0042]** FIG. 8A is a diagram illustrating an example blockchain which stores machine learning (artificial intelligence) data, according to example embodiments.
- [0043]** FIG. 8B is a diagram illustrating an example quantum-secure blockchain, according to example embodiments.
- [0044]** FIG. 9 is a diagram illustrating an example system that supports one or more of the example embodiments.

DETAILED DESCRIPTION

[0045] It will be readily understood that the instant components, as generally described and illustrated in the figures herein, may be arranged and designed in a wide variety of different configurations. Thus, the following detailed description of the embodiments of at least one of a method, apparatus, non-transitory computer readable medium and system, as represented in the attached figures, is not intended to limit the scope of the application as claimed but is merely representative of selected embodiments.

[0046] The instant features, structures, or characteristics as described throughout this specification may be combined or removed in any suitable manner in one or more embodiments. For example, the usage of the phrases "example embodiments", "some embodiments", or other similar language, throughout this specification refers to the fact that a particular feature, structure, or characteristic described in connection with the embodiment may be included in at least one embodiment. Thus, appearances of the phrases "example embodiments", "in some embodiments", "in other embodiments", or other similar language, throughout this specification do not necessarily all refer to the same group of embodiments, and the described features, structures, or characteristics may be combined or removed in any suitable manner in one or more embodiments. Further, in the diagrams, any connection between elements can permit one-way and/or two-way communication even if the depicted connection is a one-way or two-way arrow. Also, any device depicted in the drawings can be a different device. For example, if a mobile device is shown sending information, a wired device could also be used to send the information.

[0047] In addition, while the term "message" may have been used in the description of embodiments, the application may be applied to many types of networks and data. Furthermore, while certain types of connections, messages, and signaling may be depicted in exemplary embodiments, the application is not limited to a certain type of connection, message, and signaling.

[0048] Example embodiments provide methods, systems, components, non-transitory computer readable media, devices, and/or networks, which implement an index structure that is logically on top of a blockchain ledger and which prevents non-critical phantom items from occurring during a blockchain transaction.

[0049] In one embodiment, the application utilizes a decentralized database (such as a blockchain) that is a distributed storage system, which includes multiple nodes that communicate with each other. The decentralized database includes an append-only immutable data structure resembling a distributed ledger capable of maintaining records between mutually untrusted parties. The untrusted parties are referred to herein as peers or peer nodes. Each peer maintains a copy of the database records and no single peer can modify the database records without a consensus being reached among the distributed peers. For example, the peers may execute a consensus protocol to validate blockchain storage transactions, group the storage transactions into blocks, and build a hash chain over the

blocks. This process forms the ledger by ordering the storage transactions, as is necessary, for consistency. In various embodiments, a permissioned and/or a permissionless blockchain can be used. In a public or permission-less blockchain, anyone can participate without a specific identity. Public blockchains can involve native cryptocurrency and use consensus based on various protocols such as Proof of Work (PoW). On the other hand, a permissioned blockchain database provides secure interactions among a group of entities which share a common goal but which do not fully trust one another, such as businesses that exchange funds, goods, information, and the like.

[0050] The application can utilize a blockchain that operates arbitrary, programmable logic, tailored to a decentralized storage scheme and referred to as “smart contracts” or “chaincodes.” In some cases, specialized chaincodes may exist for management functions and parameters which are referred to as system chaincode. The application can further utilize smart contracts that are trusted distributed applications which leverage tamper-proof properties of the blockchain database and an underlying agreement between nodes, which is referred to as an endorsement or endorsement policy. Blockchain transactions associated with this application can be “endorsed” before being committed to the blockchain while transactions, which are not endorsed, are disregarded. An endorsement policy allows chaincode to specify endorsers for a transaction in the form of a set of peer nodes that are necessary for endorsement. When a client sends the transaction to the peers specified in the endorsement policy, the transaction is executed to validate the transaction. After validation, the transactions enter an ordering phase in which a consensus protocol is used to produce an ordered sequence of endorsed transactions grouped into blocks.

[0051] This application can utilize nodes that are the communication entities of the blockchain system. A “node” may perform a logical function in the sense that multiple nodes of different types can run on the same physical server. Nodes are grouped in trust domains and are associated with logical entities that control them in various ways. Nodes may include different types, such as a client or submitting-client node which submits a transaction-invocation to an endorser (e.g., peer), and broadcasts transaction-proposals to an ordering service (e.g., ordering node). Another type of node is a peer node which can receive client submitted transactions, commit the transactions and maintain a state and a copy of the ledger of blockchain transactions. Peers can also have the role of an endorser, although it is not a requirement. An ordering-service-node or orderer is a node running the communication service for all nodes, and which implements a delivery guarantee, such as a broadcast to each of the peer nodes in the system when committing transactions and modifying a world state of the blockchain, which is another name for the initial blockchain transaction which normally includes control and setup information.

[0052] This application can utilize a ledger that is a sequenced, tamper-resistant record of all state transitions of a blockchain. State transitions may result from chaincode invocations (i.e., transactions) submitted by participating parties (e.g., client nodes, ordering nodes, endorser nodes, peer nodes, etc.). Each participating party (such as a peer node) can maintain a copy of the ledger. A transaction may result in a set of asset key-value pairs being committed to the ledger as one or more operands, such as creates, updates, deletes, and the like. The ledger includes a blockchain

(also referred to as a chain) which is used to store an immutable, sequenced record in blocks. The ledger also includes a state database which maintains a current state of the blockchain.

[0053] This application can utilize a chain that is a transaction log which is structured as hash-linked blocks, and each block contains a sequence of N transactions where N is equal to or greater than one. The block header includes a hash of the block's transactions, as well as a hash of the prior block's header. In this way, all transactions on the ledger may be sequenced and cryptographically linked together. Accordingly, it is not possible to tamper with the ledger data without breaking the hash links. A hash of a most recently added blockchain block represents every transaction on the chain that has come before it, making it possible to ensure that all peer nodes are in a consistent and trusted state. The chain may be stored on a peer node file system (i.e., local, attached storage, cloud, etc.), efficiently supporting the append-only nature of the blockchain workload.

[0054] The current state of the immutable ledger may be stored in a state database, and may represent the latest values for all keys that are included in the chain transaction log. Since the current state represents the latest key values known to a channel, it is sometimes referred to as a world state. Chaincode invocations execute transactions against the current state data of the ledger. To make these chaincode interactions efficient, the latest values of the keys may be stored in a state database. The state database may be simply an indexed view into the chain's transaction log, it can therefore be regenerated from the chain at any time. The state database may automatically be recovered (or generated if needed) upon peer node startup, and before transactions are accepted.

[0055] Before data may be added to a distributed database such as blockchain, a consensus protocol must be performed to ensure the data is agreed upon. This form of decentralized agreement enables blockchain to be managed without a central/controlling entity. In Hyperledger Fabric, the consensus protocol has various phases including endorsement, ordering, and validation/commitment. During an endorsement phase, a subset of blockchain peers simulate a transaction to ensure the transaction can be processed based on the current state of the blockchain (stored on a blockchain ledger). However, no data is added to the blockchain at this point. If endorsed, the transaction is sent to an orderer node which aggregates endorsed transactions into a data block. The orderer node then cuts the data block when it is full and broadcasts the data block to blockchain peers on that channel. The blockchain peers then re-execute the transactions based on the current state of the blockchain to ensure that data is still valid. If approved, the transactions are used to update the state of the ledger and the data block is stored on the blockchain.

[0056] To provide proof of the endorsement phase, endorser nodes will generate a read set that identifies the key values that are read from the blockchain ledger and a write set that identifies values to be written to the blockchain ledger. When the transactions are subsequently committed by a blockchain peer during the validation/commitment phase, the read set is re-run at commit time. If the result is different from proposal/endorsement time, a phantom read error is thrown. This check ensures that the state and environment has not been changed from proposal time to commit

time. However, the read set is only evaluated during range queries and is the result of the range query performed by Hyperledger Fabric. A phantom read error from non-critical range queries can be a huge obstacle for performance of smart contract logic because phantom read errors cause a transaction to be denied.

[0057] In many cases, a transaction can run multiple range queries and not all range queries are critical to evaluate endorsement. In the cases where a range query is not critical to the evaluation, there is no need for a peer to check for phantom read errors. However, the current framework of blockchain solutions such as Hyperledger Fabric do not provide a way to avoid execution of range queries. Due to this limitation, some transactions fail even though they should succeed. Such failures eliminate the possibility of running concurrent transactions that make the same non-critical range queries.

[0058] The example embodiments overcome the deficiencies in related blockchain solutions by building an index structure on top of the blockchain ledger and enabling blockchain peers to query the index structure instead of range-querying the blockchain ledger. In some embodiments, the index structure is a binary tree where keys are stored as nodes that are connected to each other via links. The index structure may include a hierarchical format and nodes / keys may be connected to each other based on any desired business rules. When a blockchain transaction is received, a blockchain peer may endorse the transaction by reading keys for non-critical queries from the index structure instead of the blockchain ledger. In doing so, when the endorsing peer generates a read set for the transaction, the blockchain peer can omit the data read from the index structure. Accordingly, when a committing peer re-runs the queries at commit time, the committing peer may also read the keys for non-critical queries from the index structure, and avoid performing a comparison of the non-critical data. Thus, phantom read errors for non-critical items can be avoided.

[0059] The index may be structured as a binary search tree that is logically on top of the blockchain ledger and which allows transactions to bypass phantom error checks for non-critical range queries. The benefits of the example embodiments include significant improvement in chaincode performance by allowing developers to implement concurrent smart contract logic. Furthermore, transaction throughput is also improved because transactions are no longer denied due to phantom items of non-critical transactional data. The system still allows for critical range queries to be processed off of the ledger. Therefore, phantom read errors can be checked for critical range queries.

[0060] FIG. 1A illustrates a blockchain network 100 implementing an index structure 130 on top of a blockchain ledger 120, according to example embodiments. Referring to FIG. 1A, a plurality of blockchain peers 121-125 share in management of a blockchain (not shown). The blockchain may include a chain of blocks that are hash-linked together on the blockchain ledger 120. The blockchain ledger 120 may also store a state database (referred to herein as a world state) which includes the current values of all keys (unique data items) stored on the blockchain. In a typical

transaction from a client 110, a blockchain peer would endorse the transaction based on key-values obtained from the state database of the blockchain ledger 120.

[0061] Meanwhile, in the example embodiments, the blockchain peers 121-125 may build and maintain an index structure 130 on top of the blockchain ledger 120. The index structure 130 may include an index of keys from the blockchain ledger 130 which are ordered/connected based on business logic defined by the system. When a blockchain request (such as a transaction) is received from client 110, non-critical queries of the blockchain request may be processed using the index structure 130 while critical queries can still be processed from key values on the blockchain ledger 120.

[0062] In FIG. 1A, the client 110 has chosen blockchain peers 123 and 124 to be endorsing peers for a transaction. Each of these blockchain peers 123 and 124 may simulate the transaction by invoking chaincode (deployed / packaged smart contract logic) which reads data from the blockchain ledger 120 to verify whether the transaction will be successful. To endorse the transaction, the blockchain peer 123 may run a plurality of range queries which retrieve data from the blockchain ledger 120. According to various embodiments, the blockchain peer 123 may identify one or more range queries that are not critical to the endorsement (*i.e.*, they do not affect the transactions validity) and query the index structure 130 for the keys of such query. When the blockchain peer 123 generates a read set for the endorsement, the blockchain peer 123 may include all of the keys that are read from the blockchain ledger 120 (critical range queries) and omit the keys that are read from the index structure 130 (non-critical queries).

[0063] The resulting read set may be returned to the client 110 and forward to an ordering node (not shown) for inclusion in the blockchain. When a data block including the transaction (and its read set) is cut and sent to the blockchain peers 121-125, each of the blockchain peers 121-125 may again perform the queries to validate the transaction before committing the block to its local copy of the blockchain ledger 120. Here, the blockchain peers 121-125 may query the blockchain ledger 120 for critical range queries, and compare the key values read from the blockchain ledger 120 to the key values stored in the read set of the transaction. However, for non-critical queries, the blockchain peers 121-125 may query their own copy of the index structure 130 without having to perform a validation of the non-critical queries. As a result, phantom read errors for non-critical data/queries can be avoided.

[0064] FIG. 1B illustrates an example of the index structure 130 of FIG. 1A, according to example embodiments. Referring to FIG. 1B, chaincode running on a blockchain peer may build the index structure 130 on top of the blockchain ledger 120 to create a custom index for less error-prone queries. In this example, the index structure 130 is a binary tree. The index structure 130 may be part of the smart contract logic that is packaged into chaincode that is deployed and executed by the blockchain peers 121-125 shown in FIG. 1A.

[0065] A range query uses a range iterator to go over a set of ledger keys in the blockchain ledger 120. The iterator can be used to iterate over all keys between a start key and an end key within the range. Each key has a value associated with it. The peers use range queries to get a range of values. A phantom Read error occurs if two query executions render different results. For example, if a range query executed during an endorsement of a transaction and the same range query executed during a validation/commitment of the transaction result in different keys/values being returned, a phantom error is detected. A phantom error can be caused by a concurrent transaction modifying a key(s) within the range of records in between the endorsement and the validation/commitment.

[0066] In the example embodiments, the index structure 130 persists between transactions. It is an alternative way to store ledger keys rather than simply storing keys in the blockchain ledger 120 database. In FIG. 1B, the keys are stored as nodes 132 in the index structure. Furthermore, the keys/nodes 132 are linked with links 134 in the binary tree. Each node 132 in the index structure 130 corresponds to a key value stored on the blockchain ledger 120, as denoted by the arrows in FIG. 1B. The blockchain ledger 120 may store key-value pairs which include a key (e.g., system value) and a timestamp at which the key was added to the system. Each ledger key corresponds to a node in the index structure 130. Furthermore, each key node's next property points to another key. The keys may be sorted in the index structure 130 at insertion based on business logic within the smart contract of the blockchain peer.

[0067] To query the index structure 130, the blockchain peer may read the keys one by one as nodes in the tree instead of querying them by range on the blockchain ledger 120. In the example of a binary search tree as the index structure 130, the blockchain peer can identify the next node 132 in the index structure 130 by calling a Next(). This function identifies an order among the keys within the index structure 130. The blockchain peer will not generate a phantom read error by simply querying the next node. Some phantom read errors are valid as certain data should not be modified within the same read-write set. These data items may be still be queried from the blockchain ledger 120. However, other phantom read errors are due to changes in data that is not critical to the transaction's validity. The example embodiments provide an innovative way to avoid phantom read errors that are not critical by querying the index structure 130 and not adding the results to a read set.

[0068] Without the custom index structure built by the example embodiments, running range queries (such as by using Fabric's GetStateByRange function) with an iterator is the only way for the blockchain peer to understand what the next key is. So if any key within the range is added, updated, or deleted between proposal and commit time, GetStateByRange would return different results and the phantom read error would occur. In contrast, instead of reading keys in a range query, the example embodiments read individual elements in the range one by one using an index structure such as a binary search tree. With a binary search tree, the blockchain peer can identify the next key is by calling Next(). By eliminating the need for the ledger's native range query function, the embodiments create a way to eliminate non-critical phantom read checks to allow better chaincode performance.

[0069] Blockchain, i.e. Hyperledger Fabric, has the following functions (or similar) to save and query ledger data. Ledger data is typically saved as a key-value pair. For example, PutState(k,v) is a function to store value v in the ledger with key k. GetState(k) is a function to get a value stored in the ledger for the key k. This query is not added to the read set of transaction during a phantom read error check. GetStateByRange(startKey, endKey) is a function to query ledger data between start key (inclusive) and end key (exclusive). This range query is added to the read set of transaction during a phantom read error check.

[0070] Meanwhile, the index structure (e.g., a binary search tree, red-black tree, etc.) provides the following functions to add a node or find a node in the binary tree. In this example, the node values in the binary search tree represent ledger keys. For example, Insert(k) may insert the key k to the binary search tree and rebalance if needed. Find(k) may find a node in the binary search tree whose key is greater than or equal to input key value k. Next(k) may receive the current node k in the binary search tree, and return the next key k+1 in the binary search tree. The order of the keys in the index structure may be based on business logic implemented within the smart contract.

[0071] Furthermore, the example embodiments enable a peer to build an index of the ledger keys using a binary search tree on top of the ledger. When a piece of data (key-value pair) is added to the blockchain ledger, the blockchain peer may use the following logic instead of using blockchain's native PutState() function.

```
PutStateBST (k, v) {
    Insert(k)
    PutState(k, v)
}
```

[0072] In this example, before adding the data to the ledger, the peer may add the key to the binary search tree in order to build an index of ledger keys. In order to bypass a phantom read error check for non-critical range queries, the blockchain peer may use the binary search tree index to query a range of data using the following logic instead of Blockchain's native GetStateByRange() function. The following function uses the binary search tree index in order to find the next key to read, and then reads each value from the ledger individually using the native GetState() function. This avoids adding a read set for phantom read error checking.

```
RangeQueryBST(sk, ek) {
    result = []
    n = Find(sk)
    while ( n < ek ) {
        v = GetState(n)
        result = append(result, v)
        n = Next(n)
    }
}
```

```

    return result
}

```

[0073] FIG. 2A illustrates a blockchain architecture configuration 200, according to example embodiments. Referring to FIG. 2A, the blockchain architecture 200 may include certain blockchain elements, for example, a group of blockchain nodes 202. The blockchain nodes 202 may include one or more nodes 204-210 (these four nodes are depicted by example only). These nodes participate in a number of activities, such as blockchain transaction addition and validation process (consensus). One or more of the blockchain nodes 204-210 may endorse transactions based on endorsement policy and may provide an ordering service for all blockchain nodes in the architecture 200. A blockchain node may initiate a blockchain authentication and seek to write to a blockchain immutable ledger stored in blockchain layer 216, a copy of which may also be stored on the underpinning physical infrastructure 214. The blockchain configuration may include one or more applications 224 which are linked to application programming interfaces (APIs) 222 to access and execute stored program/application code 220 (e.g., chaincode, smart contracts, etc.) which can be created according to a customized configuration sought by participants and can maintain their own state, control their own assets, and receive external information. This can be deployed as a transaction and installed, via appending to the distributed ledger, on all blockchain nodes 204-210.

[0074] The blockchain base or platform 212 may include various layers of blockchain data, services (e.g., cryptographic trust services, virtual execution environment, etc.), and underpinning physical computer infrastructure that may be used to receive and store new transactions and provide access to auditors which are seeking to access data entries. The blockchain layer 216 may expose an interface that provides access to the virtual execution environment necessary to process the program code and engage the physical infrastructure 214. Cryptographic trust services 218 may be used to verify transactions such as asset exchange transactions and keep information private.

[0075] The blockchain architecture configuration of FIG. 2A may process and execute program/application code 220 via one or more interfaces exposed, and services provided, by blockchain platform 212. The code 220 may control blockchain assets. For example, the code 220 can store and transfer data, and may be executed by nodes 204-210 in the form of a smart contract and associated chaincode with conditions or other code elements subject to its execution. As a non-limiting example, smart contracts may be created to execute reminders, updates, and/or other notifications subject to the changes, updates, etc. The smart contracts can themselves be used to identify rules associated with authorization and access requirements and usage of the ledger. For example, the smart contract (or chaincode executing the logic of the smart contract) may read blockchain data 226 from an index structure. The blockchain data may be processed by one or more processing entities (e.g., virtual machines) included in the blockchain layer 216 to generate results 228 of the chaincode invocation. The results may be determined by the business logic within the smart contract. The physical infrastructure 214 may be utilized to retrieve any of the data or information described herein.

[0076] A smart contract may be created via a high-level application and programming language, and then written to a block in the blockchain. The smart contract may include executable code which is registered, stored, and/or replicated with a blockchain (e.g., distributed network of blockchain peers). To be deployed, the smart contract may be packaged into chaincode (with other smart contracts) and deployed onto a blockchain peer. A transaction is an execution of the smart contract code which can be performed in response to conditions associated with the smart contract being satisfied. The executing of the smart contract may trigger a trusted modification(s) to a state of a digital blockchain ledger. The modification(s) to the blockchain ledger caused by the smart contract execution may be automatically replicated throughout the distributed network of blockchain peers through one or more consensus protocols.

[0077] The smart contract may write data to the blockchain in the format of key-value pairs. Furthermore, the smart contract code can read the values stored in a blockchain and use them in application operations. The smart contract code can write the output of various logic operations into one or more blocks within the blockchain. The code may be used to create a temporary data structure in a virtual machine or other computing platform. Data written to the blockchain can be public and/or can be encrypted and maintained as private. The temporary data that is used/generated by the smart contract is held in memory by the supplied execution environment, then deleted once the data needed for the blockchain is identified.

[0078] A chaincode may include the code interpretation of a smart contract, with additional features. Here, the chaincode may include a packaged and deployable version of the logic within the smart contract. As described herein, the chaincode may be program code deployed on a computing network, where it is executed and validated by chain validators together during a consensus process. The chaincode receives a hash and retrieves from the blockchain a hash associated with the data template created by use of a previously stored feature extractor. If the hashes of the hash identifier and the hash created from the stored identifier template data match, then the chaincode sends an authorization key to the requested service. The chaincode may write to the blockchain data associated with the cryptographic details.

[0079] FIG. 2B illustrates an example of a blockchain transactional flow 250 between nodes of the blockchain in accordance with an example embodiment. Referring to FIG. 2B, the transaction flow may include a client node 260 transmitting a transaction proposal 291 to an endorsing peer node 281. The endorsing peer 281 may verify the client signature and execute a chaincode function to initiate the transaction. The output may include the chaincode results, a set of key/value versions that were read in the chaincode (read set), and the set of keys/values that were written in chaincode (write set). Here, the endorsing peer 281 may determine whether or not to endorse the transaction proposal. The proposal response 292 is sent back to the client 260 along with an endorsement signature, if approved. The client 260 assembles the endorsements into a transaction payload 293 and broadcasts it to an ordering service node 284. The ordering service node 284 then delivers ordered transactions as blocks to all peers 281-283 on a

channel. Before committal to the blockchain, each peer 281-283 may validate the transaction. For example, the peers may check the endorsement policy to ensure that the correct allotment of the specified peers have signed the results and authenticated the signatures against the transaction payload 293.

[0080] Referring again to FIG. 2B, the client node initiates the transaction 291 by constructing and sending a request to the peer node 281, which is an endorser. The client 260 may include an application leveraging a supported software development kit (SDK), which utilizes an available API to generate a transaction proposal. The proposal is a request to invoke a chaincode function so that data can be read and/or written to the ledger (i.e., write new key value pairs for the assets). The SDK may serve as a shim to package the transaction proposal into a properly architected format (e.g., protocol buffer over a remote procedure call (RPC)) and take the client's cryptographic credentials to produce a unique signature for the transaction proposal.

[0081] In response, the endorsing peer node 281 may verify (a) that the transaction proposal is well formed, (b) the transaction has not been submitted already in the past (replay-attack protection), (c) the signature is valid, and (d) that the submitter (client 260, in the example) is properly authorized to perform the proposed operation on that channel. The endorsing peer node 281 may take the transaction proposal inputs as arguments to the invoked chaincode function. The chaincode is then executed against a current state database to produce transaction results including a response value, read set, and write set. However, no updates are made to the ledger at this point. In 292, the set of values, along with the endorsing peer node's 281 signature is passed back as a proposal response 292 to the SDK of the client 260 which parses the payload for the application to consume.

[0082] In response, the application of the client 260 inspects/verifies the endorsing peers signatures and compares the proposal responses to determine if the proposal response is the same. If the chaincode only queried the ledger, the application would inspect the query response and would typically not submit the transaction to the ordering node service 284. If the client application intends to submit the transaction to the ordering node service 284 to update the ledger, the application determines if the specified endorsement policy has been fulfilled before submitting (i.e., did all peer nodes necessary for the transaction endorse the transaction). Here, the client may include only one of multiple parties to the transaction. In this case, each client may have their own endorsing node, and each endorsing node will need to endorse the transaction. The architecture is such that even if an application selects not to inspect responses or otherwise forwards an unendorsed transaction, the endorsement policy will still be enforced by peers and upheld at the commit validation phase.

[0083] After successful inspection, in step 293 the client 260 assembles endorsements into a transaction proposal and broadcasts the transaction proposal and response within a transaction message to the ordering node 284. The transaction may contain the read/write sets, the endorsing peers signatures and a channel ID. The ordering node 284 does not need to inspect the entire content of a transaction in order to perform its operation, instead the

ordering node 284 may simply receive transactions from all channels in the network, order them chronologically by channel, and create blocks of transactions per channel.

[0084] The blocks are delivered from the ordering node 284 to all peer nodes 281-283 on the channel. The data section within the block may be validated to ensure an endorsement policy is fulfilled and to ensure that there have been no changes to ledger state for read set variables since the read set was generated by the transaction execution. Furthermore, in step 295 each peer node 281-283 appends the block to the channel's chain, and for each valid transaction the write sets are committed to current state database. An event may be emitted, to notify the client application that the transaction (invocation) has been immutably appended to the chain, as well as to notify whether the transaction was validated or invalidated.

[0085] FIG. 3A illustrates an example of a permissioned blockchain network 300, which features a distributed, decentralized peer-to-peer architecture. In this example, a blockchain user 302 may initiate a transaction to the permissioned blockchain 304. In this example, the transaction can be a deploy, invoke, or query, and may be issued through a client-side application leveraging an SDK, directly through an API, etc. Networks may provide access to a regulator 306, such as an auditor. A blockchain network operator 308 manages member permissions, such as enrolling the regulator 306 as an "auditor" and the blockchain user 302 as a "client". An auditor could be restricted only to querying the ledger whereas a client could be authorized to deploy, invoke, and query certain types of chaincode.

[0086] A blockchain developer 310 can write chaincode and client-side applications. The blockchain developer 310 can deploy chaincode directly to the network through an interface. To include credentials from a traditional data source 312 in chaincode, the developer 310 could use an out-of-band connection to access the data. In this example, the blockchain user 302 connects to the permissioned blockchain 304 through a peer node 314. Before proceeding with any transactions, the peer node 314 retrieves the user's enrollment and transaction certificates from a certificate authority 316, which manages user roles and permissions. In some cases, blockchain users must possess these digital certificates in order to transact on the permissioned blockchain 304. Meanwhile, a user attempting to utilize chaincode may be required to verify their credentials on the traditional data source 312. To confirm the user's authorization, chaincode can use an out-of-band connection to this data through a traditional processing platform 318.

[0087] FIG. 3B illustrates another example of a permissioned blockchain network 320, which features a distributed, decentralized peer-to-peer architecture. In this example, a blockchain user 322 may submit a transaction to the permissioned blockchain 324. In this example, the transaction can be a deploy, invoke, or query, and may be issued through a client-side application leveraging an SDK, directly through an API, etc. Networks may provide access to a regulator 326, such as an auditor. A blockchain network operator 328 manages member permissions, such as enrolling the regulator 326 as an "auditor" and the blockchain user 322 as a "client". An auditor could be restricted only to querying the ledger whereas a client could be authorized to deploy, invoke, and query certain types of chaincode.

[0088] A blockchain developer 330 writes chaincode and client-side applications. The blockchain developer 330 can deploy chaincode directly to the network through an interface. To include credentials from a traditional data source 332 in chaincode, the developer 330 could use an out-of-band connection to access the data. In this example, the blockchain user 322 connects to the network through a peer node 334. Before proceeding with any transactions, the peer node 334 retrieves the user's enrollment and transaction certificates from the certificate authority 336. In some cases, blockchain users must possess these digital certificates in order to transact on the permissioned blockchain 324. Meanwhile, a user attempting to utilize chaincode may be required to verify their credentials on the traditional data source 332. To confirm the user's authorization, chaincode can use an out-of-band connection to this data through a traditional processing platform 338.

[0089] In some embodiments, the blockchain herein may be a permissionless blockchain. In contrast with permissioned blockchains which require permission to join, anyone can join a permissionless blockchain. For example, to join a permissionless blockchain a user may create a personal address and begin interacting with the network, by submitting transactions, and hence adding entries to the ledger. Additionally, all parties have the choice of running a node on the system and employing the mining protocols to help verify transactions.

[0090] FIG. 3C illustrates a process 350 of a transaction being processed by a permissionless blockchain 352 including a plurality of nodes 354. A sender 356 desires to send payment or some other form of value (e.g., a deed, medical records, a contract, a good, a service, or any other asset that can be encapsulated in a digital record) to a recipient 358 via the permissionless blockchain 352. In one embodiment, each of the sender device 356 and the recipient device 358 may have digital wallets (associated with the blockchain 352) that provide user interface controls and a display of transaction parameters. In response, the transaction is broadcast throughout the blockchain 352 to the nodes 354. Depending on the blockchain's 352 network parameters the nodes verify 360 the transaction based on rules (which may be pre-defined or dynamically allocated) established by the permissionless blockchain 352 creators. For example, this may include verifying identities of the parties involved, etc. The transaction may be verified immediately or it may be placed in a queue with other transactions and the nodes 354 determine if the transactions are valid based on a set of network rules.

[0091] In structure 362, valid transactions are formed into a block and sealed with a lock (hash). This process may be performed by mining nodes among the nodes 354. Mining nodes may utilize additional software specifically for mining and creating blocks for the permissionless blockchain 352. Each block may be identified by a hash (e.g., 256 bit number, etc.) created using an algorithm agreed upon by the network. Each block may include a header, a pointer or reference to a hash of a previous block's header in the chain, and a group of valid transactions. The reference to the previous block's hash is associated with the creation of the secure independent chain of blocks.

[0092] Before blocks can be added to the blockchain, the blocks must be validated. Validation for the permissionless blockchain 352 may include a proof-of-work (PoW) which is a solution to a puzzle derived from the block's header. Although not shown in the example of FIG. 3C, another process for validating a block is proof-of-stake. Unlike the proof-of-work, where the algorithm rewards miners who solve mathematical problems, with the proof of stake, a creator of a new block is chosen in a deterministic way, depending on its wealth, also defined as "stake." Then, a similar proof is performed by the selected/chosen node.

[0093] With mining 364, nodes try to solve the block by making incremental changes to one variable until the solution satisfies a network-wide target. This creates the PoW thereby ensuring correct answers. In other words, a potential solution must prove that computing resources were drained in solving the problem. In some types of permissionless blockchains, miners may be rewarded with value (e.g., coins, etc.) for correctly mining a block.

[0094] Here, the PoW process, alongside the chaining of blocks, makes modifications of the blockchain extremely difficult, as an attacker must modify all subsequent blocks in order for the modifications of one block to be accepted. Furthermore, as new blocks are mined, the difficulty of modifying a block increases, and the number of subsequent blocks increases. With distribution 366, the successfully validated block is distributed through the permissionless blockchain 352 and all nodes 354 add the block to a majority chain which is the permissionless blockchain's 352 auditable ledger. Furthermore, the value in the transaction submitted by the sender 356 is deposited or otherwise transferred to the digital wallet of the recipient device 358.

[0095] FIG. 4A illustrates a process 400A of querying a blockchain ledger via a range query according to example embodiments. Referring to FIG. 4A, a smart contract 444 running on a blockchain peer includes a state database 446 (stored on a blockchain ledger). In this example, a logging system 442 tracks users that are logged onto an enterprise network and records data about the user activities. The data may be stored on a blockchain managed by smart contract 444 including keys. The current value of all keys may be stored by the smart contract 444 within the state database 446.

[0096] In this example, the smart contract 444 has logic to count a number of user/employees that are logged into logging system 442 and log the results on the blockchain ledger (state DB 446) every second of time. FIG. 4A illustrates an example of a communication sequence of the smart contract 444 storing transactions to the state database 446, and the occurrence of a phantom read error.

[0097] Because of network overhead, the time taken between proposal of a transaction and commit of that transaction can vary. In 401, a logging system starts a first transaction at a time of 12:00:00. In 402, the smart contract 444 counts the number of employees logged into the system by a range query. Currently 6 employees are logged into the system. The range query is then added to read set to check a phantom read error.

[0098] Meanwhile, in 403 a new user 440 logs onto the system. Now 7 employees are logged onto the system. Furthermore, the smart contract 444 updates the blockchain ledger 446 with the new user information. This happens concurrently between the endorsement and the validation of the first transaction received in 401. In 405, the logging system 442 triggers the start of a next count transaction (transaction 2) at 12:00:01. The smart contract 444 counts the number of employees logged in by a range query, in 406. The total count is 7. The range query added to read set for transaction 2 to check a phantom read error. Here, a phantom error check for the second transaction is performed in 407 and it succeeds. Thus, in 408, the second transaction is committed to the blockchain and the state DB 446, and a success notification is sent in 409.

[0099] Meanwhile, the first transaction was delayed due to overhead. Here, in 411, the smart contract 444 performs a phantom error check based on the current state of the ledger (state DB 446). In this example, the current number of users logged onto the system is recorded as 7, but the read set of the first transaction indicates the number of users is 6. Therefore, the phantom error check fails. Accordingly, in 412, the smart contract 444 outputs a failure notification.

[00100] FIG. 4B illustrates a process 400B of querying an index structure 448 on top of a blockchain ledger 446 according to example embodiments. In this example, the index structure 448 stores current keys of the blockchain ledger in an index that logically resides above the state DB 446 of the blockchain ledger. In 421, the logging system 442 triggers a first transaction at 12:00:00. In 422, in response, the smart contract 444 counts the number of employees logged into the system by querying the index structure 448 without issuing a range query to the state DB 446. Currently 6 employees are logged into the system. This read is not added to the read set for transaction 1 since no range query was used. Instead, the read set is created with only critical queries (not shown in this example).

[00101] In 423, a new user 440 logs into the system. Now 7 employees are logged in. In this example, both the state DB 446 and the index structure 448 are updated with the new key for the user 440.

[00102] Meanwhile, as in the case of FIG. 4A, the logging system 442 triggers the start of a second transaction in 425. In this case, transaction 2 starts at 12:00:01. In 426, the smart contract 444 counts the number of employees logged in by querying the index structure 448. The total count is 7. This read, however, is not added to the read set for transaction 2 since no range query was used. In 427, the smart contract 444 performs a phantom error check based on the read set of the second transaction and the check succeeds. In 428, the second transaction is added to the blockchain ledger and a success message is sent to the logging system 442.

[00103] Meanwhile, just as in the example of FIG. 4A, the first transaction is delayed. However, in 430, a phantom error check is performed and succeeds because the non-critical transaction (employee/user count) was not

added to the read set. Therefore, there is no need to perform a phantom error check on the employee count value read in 422. Instead, the first transaction is committed in 431, and a success message is provided in 432.

[00104] FIG. 5 illustrates a method 500 of accessing an index structure on top of a blockchain ledger according to example embodiments. As a non-limiting example, the method 500 may be performed by a blockchain peer, chaincode that is running on a blockchain peer, and the like, however embodiments are not limited thereto. Referring to FIG. 5, in 510, the method may include storing an index structure that comprises an index of keys from a blockchain ledger, where the keys are stored as nodes in the index structure. As an example, the index structure may be a binary tree, also referred to as a binary search tree. The tree may include nodes which represent keys stored in the blockchain.

[00105] In some embodiments, the blockchain peer node may also build the index structure over time, as keys are added to the blockchain. For example, the index structure may be created and implemented within a smart contract (deployed as chaincode) running on the blockchain peer node. The nodes may be ordered or otherwise connected to each other in any manner desired. For example, the smart contract may include business rules, methods, procedures, etc. for connecting the nodes/keys in the index structure. When the index structure is a tree, such as a binary tree, the keys may be arranged as nodes in a hierarchical manner.

[00106] In 520, the method may include receiving a blockchain request for data stored on the blockchain ledger. For example, the blockchain request may include a blockchain transaction which invokes a smart contract on a blockchain peer node. The blockchain request may be processed using one or more queries to the blockchain ledger. In a traditional blockchain, these requests may be processed using range queries. However, in 530, the method may include reading a set of keys of a non-critical query included in the blockchain request from the nodes in the index structure instead of executing a range query on the blockchain ledger. Here, the blockchain request may include / need both critical and non-critical queries to the ledger. The blockchain peer may process the critical queries on the ledger. However, for non-critical queries, the blockchain peer node may instead read the data from the index structure.

[00107] Furthermore, in 540, the method may include generating and storing a read set for the blockchain request which does not include values for the set of keys of the non-critical query. Thus, a phantom read error can be avoided by processing non-critical queries on the index structure instead of the blockchain ledger. Furthermore, the peer can avoid putting the keys read from the index structure within a read set of the smart contract.

[00108] In some embodiments, the reading may include reading the set of keys one-by-one from the nodes in the index structure via a next function. Here, the next function may move to a next key/node in the index structure based on how the nodes are ordered by the smart contract when the smart contract is building the index structure. In some embodiments, the method may further include identifying a query of the blockchain request as the non-critical query

when the query does not impact a validity of the blockchain request. In some embodiments, the method may include validating the blockchain request based on a previous read set for the blockchain request which does not include the values for the set of keys read from the index structure. In some embodiments, the index structure may be implemented within logic of chaincode of a blockchain peer node. In these examples, the keys in the index structure correspond to unique values stored on a blockchain of the blockchain ledger.

[00109] FIG. 6A illustrates an example system 600 that includes a physical infrastructure 610 configured to perform various operations according to example embodiments. Referring to FIG. 6A, the physical infrastructure 610 includes a module 612 and a module 614. The module 614 includes a blockchain 620 and a smart contract 630 (which may reside on the blockchain 620), that may execute any of the operational steps 608 (in module 612) included in any of the example embodiments. The steps/operations 608 may include one or more of the embodiments described or depicted and may represent output or written information that is written or read from one or more smart contracts 630 and/or blockchains 620. The physical infrastructure 610, the module 612, and the module 614 may include one or more computers, servers, processors, memories, and/or wireless communication devices. Further, the module 612 and the module 614 may be a same module.

[00110] FIG. 6B illustrates another example system 640 configured to perform various operations according to example embodiments. Referring to FIG. 6B, the system 640 includes a module 612 and a module 614. The module 614 includes a blockchain 620 and a smart contract 630 (which may reside on the blockchain 620), that may execute any of the operational steps 608 (in module 612) included in any of the example embodiments. The steps/operations 608 may include one or more of the embodiments described or depicted and may represent output or written information that is written or read from one or more smart contracts 630 and/or blockchains 620. The physical infrastructure 610, the module 612, and the module 614 may include one or more computers, servers, processors, memories, and/or wireless communication devices. Further, the module 612 and the module 614 may be a same module.

[00111] FIG. 6C illustrates an example system configured to utilize a smart contract configuration among contracting parties and a mediating server configured to enforce the smart contract terms on the blockchain according to example embodiments. Referring to FIG. 6C, the configuration 650 may represent a communication session, an asset transfer session or a process or procedure that is driven by a smart contract 630 which explicitly identifies one or more user devices 652 and/or 656. The execution, operations and results of the smart contract execution may be managed by a server 654. Content of the smart contract 630 may require digital signatures by one or more of the entities 652 and 656 which are parties to the smart contract transaction. The results of the smart contract execution may be written to a blockchain 620 as a blockchain transaction. The smart contract 630 resides on the blockchain 620 which may reside on one or more computers, servers, processors, memories, and/or wireless communication devices.

[00112] FIG. 6D illustrates a system 660 including a blockchain, according to example embodiments. Referring to the example of FIG. 6D, an application programming interface (API) gateway 662 provides a common interface for accessing blockchain logic (e.g., smart contract 630 or other chaincode) and data (e.g., distributed ledger, etc.). In this example, the API gateway 662 is a common interface for performing transactions (invoke, queries, etc.) on the blockchain by connecting one or more entities 652 and 656 to a blockchain peer (i.e., server 654). Here, the server 654 is a blockchain network peer component that holds a copy of the world state and a distributed ledger allowing clients 652 and 656 to query data on the world state as well as submit transactions into the blockchain network where, depending on the smart contract 630 and endorsement policy, endorsing peers will run the smart contracts 630.

[00113] The above embodiments may be implemented in hardware, in a computer program executed by a processor, in firmware, or in a combination of the above. A computer program may be embodied on a computer readable medium, such as a storage medium. For example, a computer program may reside in random access memory ("RAM"), flash memory, read-only memory ("ROM"), erasable programmable read-only memory ("EPROM"), electrically erasable programmable read-only memory ("EEPROM"), registers, hard disk, a removable disk, a compact disk read-only memory ("CD-ROM"), or any other form of storage medium known in the art.

[00114] An exemplary storage medium may be coupled to the processor such that the processor may read information from, and write information to, the storage medium. In the alternative, the storage medium may be integral to the processor. The processor and the storage medium may reside in an application specific integrated circuit ("ASIC"). In the alternative, the processor and the storage medium may reside as discrete components.

[00115] FIG. 7A illustrates a process 700 of a new block being added to a distributed ledger 720, according to example embodiments, and FIG. 7B illustrates contents of a new data block structure 730 for blockchain, according to example embodiments. Referring to FIG. 7A, clients (not shown) may submit transactions to blockchain nodes 711, 712, and/or 713. Clients may be instructions received from any source to enact activity on the blockchain 720. As an example, clients may be applications that act on behalf of a requester, such as a device, person or entity to propose transactions for the blockchain. The plurality of blockchain peers (e.g., blockchain nodes 711, 712, and 713) may maintain a state of the blockchain network and a copy of the distributed ledger 720. Different types of blockchain nodes/peers may be present in the blockchain network including endorsing peers which simulate and endorse transactions proposed by clients and committing peers which verify endorsements, validate transactions, and commit transactions to the distributed ledger 720. In this example, the blockchain nodes 711, 712, and 713 may perform the role of endorser node, committer node, or both.

[00116] The distributed ledger 720 includes a blockchain which stores immutable, sequenced records in blocks, and a state database 724 (current world state) maintaining a current state of the blockchain 722. One distributed ledger 720 may exist per channel and each peer maintains its own copy of the distributed ledger 720 for each channel of

which they are a member. The blockchain 722 is a transaction log, structured as hash-linked blocks where each block contains a sequence of N transactions. Blocks may include various components such as shown in FIG. 7B. The linking of the blocks (shown by arrows in FIG. 7A) may be generated by adding a hash of a prior block's header within a block header of a current block. In this way, all transactions on the blockchain 722 are sequenced and cryptographically linked together preventing tampering with blockchain data without breaking the hash links. Furthermore, because of the links, the latest block in the blockchain 722 represents every transaction that has come before it. The blockchain 722 may be stored on a peer file system (local or attached storage), which supports an append-only blockchain workload.

[00117] The current state of the blockchain 722 and the distributed ledger 722 may be stored in the state database 724. Here, the current state data represents the latest values for all keys ever included in the chain transaction log of the blockchain 722. Chaincode invocations execute transactions against the current state in the state database 724. To make these chaincode interactions extremely efficient, the latest values of all keys are stored in the state database 724. The state database 724 may include an indexed view into the transaction log of the blockchain 722, it can therefore be regenerated from the chain at any time. The state database 724 may automatically get recovered (or generated if needed) upon peer startup, before transactions are accepted.

[00118] Endorsing nodes receive transactions from clients and endorse the transaction based on simulated results. Endorsing nodes hold smart contracts which simulate the transaction proposals. When an endorsing node endorses a transaction, the endorsing nodes creates a transaction endorsement which is a signed response from the endorsing node to the client application indicating the endorsement of the simulated transaction. The method of endorsing a transaction depends on an endorsement policy which may be specified within chaincode. An example of an endorsement policy is "the majority of endorsing peers must endorse the transaction". Different channels may have different endorsement policies. Endorsed transactions are forward by the client application to ordering service 710.

[00119] The ordering service 710 accepts endorsed transactions, orders them into a block, and delivers the blocks to the committing peers. For example, the ordering service 710 may initiate a new block when a threshold of transactions has been reached, a timer times out, or another condition. In the example of FIG. 7A, blockchain node 712 is a committing peer that has received a new data new data block 730 for storage on blockchain 720. The first block in the blockchain may be referred to as a genesis block which includes information about the blockchain, its members, the data stored therein, etc.

[00120] The ordering service 710 may be made up of a cluster of orderers. The ordering service 710 does not process transactions, smart contracts, or maintain the shared ledger. Rather, the ordering service 710 may accept the endorsed transactions and specifies the order in which those transactions are committed to the distributed ledger 720. The architecture of the blockchain network may be designed such that the specific implementation of 'ordering' (e.g., Solo, Kafka, BFT, etc.) becomes a pluggable component.

[00121] Transactions are written to the distributed ledger 720 in a consistent order. The order of transactions is established to ensure that the updates to the state database 724 are valid when they are committed to the network. Unlike a cryptocurrency blockchain system (e.g., Bitcoin, etc.) where ordering occurs through the solving of a cryptographic puzzle, or mining, in this example the parties of the distributed ledger 720 may choose the ordering mechanism that best suits that network.

[00122] When the ordering service 710 initializes a new data block 730, the new data block 730 may be broadcast to committing peers (e.g., blockchain nodes 711, 712, and 713). In response, each committing peer validates the transaction within the new data block 730 by checking to make sure that the read set and the write set still match the current world state in the state database 724. Specifically, the committing peer can determine whether the read data that existed when the endorsers simulated the transaction is identical to the current world state in the state database 724. When the committing peer validates the transaction, the transaction is written to the blockchain 722 on the distributed ledger 720, and the state database 724 is updated with the write data from the read-write set. If a transaction fails, that is, if the committing peer finds that the read-write set does not match the current world state in the state database 724, the transaction ordered into a block will still be included in that block, but it will be marked as invalid, and the state database 724 will not be updated.

[00123] Referring to FIG. 7B, a new data block 730 (also referred to as a data block) that is stored on the blockchain 722 of the distributed ledger 720 may include multiple data segments such as a block header 740, block data 750 (block data section), and block metadata 760. It should be appreciated that the various depicted blocks and their contents, such as new data block 730 and its contents, shown in FIG. 7B are merely examples and are not meant to limit the scope of the example embodiments. In a conventional block, the data section may store transactional information of N transaction(s) (e.g., 1, 10, 100, 500, 1000, 2000, 3000, etc.) within the block data 750.

[00124] The new data block 730 may also include a link to a previous block (e.g., on the blockchain 722 in FIG. 7A) within the block header 740. In particular, the block header 740 may include a hash of a previous block's header. The block header 740 may also include a unique block number, a hash of the block data 750 of the new data block 730, and the like. The block number of the new data block 730 may be unique and assigned in various orders, such as an incremental/sequential order starting from zero.

[00125] The block metadata 760 may store multiple fields of metadata (e.g., as a byte array, etc.). Metadata fields may include signature on block creation, a reference to a last configuration block, a transaction filter identifying valid and invalid transactions within the block, last offset persisted of an ordering service that ordered the block, and the like. The signature, the last configuration block, and the orderer metadata may be added by the ordering service 710. Meanwhile, a committer of the block (such as blockchain node 712) may add validity/invalidity information based on an endorsement policy, verification of read/write sets, and the like. The transaction filter may include a byte array of a size

equal to the number of transactions that are included in the block data 750 and a validation code identifying whether a transaction was valid/invalid.

[00126] FIG. 7C illustrates an embodiment of a blockchain 770 for digital content in accordance with the embodiments described herein. The digital content may include one or more files and associated information. The files may include media, images, video, audio, text, links, graphics, animations, web pages, documents, or other forms of digital content. The immutable, append-only aspects of the blockchain serve as a safeguard to protect the integrity, validity, and authenticity of the digital content, making it suitable use in legal proceedings where admissibility rules apply or other settings where evidence is taken in to consideration or where the presentation and use of digital information is otherwise of interest. In this case, the digital content may be referred to as digital evidence.

[00127] The blockchain may be formed in various ways. In one embodiment, the digital content may be included in and accessed from the blockchain itself. For example, each block of the blockchain may store a hash value of reference information (e.g., header, value, etc.) along the associated digital content. The hash value and associated digital content may then be encrypted together. Thus, the digital content of each block may be accessed by decrypting each block in the blockchain, and the hash value of each block may be used as a basis to reference a previous block. This may be illustrated as follows:

<u>Block 1</u>	<u>Block 2</u>		<u>Block N</u>
Hash Value 1	Hash Value 2		Hash Value N
Digital Content 1	Digital Content 2		Digital Content N

[00128] In one embodiment, the digital content may be not included in the blockchain. For example, the blockchain may store the encrypted hashes of the content of each block without any of the digital content. The digital content may be stored in another storage area or memory address in association with the hash value of the original file. The other storage area may be the same storage device used to store the blockchain or may be a different storage area or even a separate relational database. The digital content of each block may be referenced or accessed by obtaining or querying the hash value of a block of interest and then looking up that hash value in the storage area, which is stored in correspondence with the actual digital content. This operation may be performed, for example, a database gatekeeper. This may be illustrated as follows:

<u>Blockchain</u>	<u>Storage Area</u>
Block 1 Hash Value	Block 1 Hash Value ... Content
.	.
.	.
.	.

Block N Hash Value

Block N Hash Value ... Content

[00129] In the example embodiment of FIG. 7C, the blockchain 770 includes a number of blocks 778₁, 778₂, ... 778_N cryptographically linked in an ordered sequence, where $N \geq 1$. The encryption used to link the blocks 778₁, 778₂, ... 778_N may be any of a number of keyed or un-keyed Hash functions. In one embodiment, the blocks 778₁, 778₂, ... 778_N are subject to a hash function which produces n-bit alphanumeric outputs (where n is 256 or another number) from inputs that are based on information in the blocks. Examples of such a hash function include, but are not limited to, a SHA-type (SHA stands for Secured Hash Algorithm) algorithm, Merkle-Damgard algorithm, HAIFA algorithm, Merkle-tree algorithm, nonce-based algorithm, and a non-collision-resistant PRF algorithm. In another embodiment, the blocks 778₁, 778₂, ..., 778_N may be cryptographically linked by a function that is different from a hash function. For purposes of illustration, the following description is made with reference to a hash function, e.g., SHA-2.

[00130] Each of the blocks 778₁, 778₂, ..., 778_N in the blockchain includes a header, a version of the file, and a value. The header and the value are different for each block as a result of hashing in the blockchain. In one embodiment, the value may be included in the header. As described in greater detail below, the version of the file may be the original file or a different version of the original file.

[00131] The first block 778₁ in the blockchain is referred to as the genesis block and includes the header 772₁, original file 774₁, and an initial value 776₁. The hashing scheme used for the genesis block, and indeed in all subsequent blocks, may vary. For example, all the information in the first block 778₁ may be hashed together and at one time, or each or a portion of the information in the first block 778₁ may be separately hashed and then a hash of the separately hashed portions may be performed.

[00132] The header 772₁ may include one or more initial parameters, which, for example, may include a version number, timestamp, nonce, root information, difficulty level, consensus protocol, duration, media format, source, descriptive keywords, and/or other information associated with original file 774₁ and/or the blockchain. The header 772₁ may be generated automatically (e.g., by blockchain network managing software) or manually by a blockchain participant. Unlike the header in other blocks 778₂ to 778_N in the blockchain, the header 772₁ in the genesis block does not reference a previous block, simply because there is no previous block.

[00133] The original file 774₁ in the genesis block may be, for example, data as captured by a device with or without processing prior to its inclusion in the blockchain. The original file 774₁ is received through the interface of the system from the device, media source, or node. The original file 774₁ is associated with metadata, which, for example, may be generated by a user, the device, and/or the system processor, either manually or automatically. The metadata may be included in the first block 778₁ in association with the original file 774₁.

[00134] The value 776₁ in the genesis block is an initial value generated based on one or more unique attributes of the original file 774₁. In one embodiment, the one or more unique attributes may include the hash value for the original file 774₁, metadata for the original file 774₁, and other information associated with the file. In one implementation, the initial value 776₁ may be based on the following unique attributes:

- 1) SHA-2 computed hash value for the original file
- 2) originating device ID
- 3) starting timestamp for the original file
- 4) initial storage location of the original file
- 5) blockchain network member ID for software to currently control the original file and associated

metadata

[00135] The other blocks 778₂ to 778_N in the blockchain also have headers, files, and values. However, unlike the first block 772₁, each of the headers 772₂ to 772_N in the other blocks includes the hash value of an immediately preceding block. The hash value of the immediately preceding block may be just the hash of the header of the previous block or may be the hash value of the entire previous block. By including the hash value of a preceding block in each of the remaining blocks, a trace can be performed from the Nth block back to the genesis block (and the associated original file) on a block-by-block basis, as indicated by arrows 780, to establish an auditable and immutable chain-of-custody.

[00136] Each of the header 772₂ to 772_N in the other blocks may also include other information, e.g., version number, timestamp, nonce, root information, difficulty level, consensus protocol, and/or other parameters or information associated with the corresponding files and/or the blockchain in general.

[00137] The files 774₂ to 774_N in the other blocks may be equal to the original file or may be a modified version of the original file in the genesis block depending, for example, on the type of processing performed. The type of processing performed may vary from block to block. The processing may involve, for example, any modification of a file in a preceding block, such as redacting information or otherwise changing the content of, taking information away from, or adding or appending information to the files.

[00138] Additionally, or alternatively, the processing may involve merely copying the file from a preceding block, changing a storage location of the file, analyzing the file from one or more preceding blocks, moving the file from one storage or memory location to another, or performing action relative to the file of the blockchain and/or its associated metadata. Processing which involves analyzing a file may include, for example, appending, including, or otherwise associating various analytics, statistics, or other information associated with the file.

[00139] The values in each of the other blocks 776_2 to 776_N in the other blocks are unique values and are all different as a result of the processing performed. For example, the value in any one block corresponds to an updated version of the value in the previous block. The update is reflected in the hash of the block to which the value is assigned. The values of the blocks therefore provide an indication of what processing was performed in the blocks and also permit a tracing through the blockchain back to the original file. This tracking confirms the chain-of-custody of the file throughout the entire blockchain.

[00140] For example, consider the case where portions of the file in a previous block are redacted, blocked out, or pixelated in order to protect the identity of a person shown in the file. In this case, the block including the redacted file will include metadata associated with the redacted file, e.g., how the redaction was performed, who performed the redaction, timestamps where the redaction(s) occurred, etc. The metadata may be hashed to form the value. Because the metadata for the block is different from the information that was hashed to form the value in the previous block, the values are different from one another and may be recovered when decrypted.

[00141] In one embodiment, the value of a previous block may be updated (e.g., a new hash value computed) to form the value of a current block when any one or more of the following occurs. The new hash value may be computed by hashing all or a portion of the information noted below, in this example embodiment.

- a) new SHA-2 computed hash value if the file has been processed in any way (e.g., if the file was redacted, copied, altered, accessed, or some other action was taken)
- b) new storage location for the file
- c) new metadata identified associated with the file
- d) transfer of access or control of the file from one blockchain participant to another blockchain participant

[00142] FIG. 7D illustrates an embodiment of a block which may represent the structure of the blocks in the blockchain 790 in accordance with one embodiment. The block, Block_i , includes a header 772_i , a file 774_i , and a value 776_i .

[00143] The header 772_i includes a hash value of a previous block Block_{i-1} and additional reference information, which, for example, may be any of the types of information (e.g., header information including references, characteristics, parameters, etc.) discussed herein. All blocks reference the hash of a previous block except, of course, the genesis block. The hash value of the previous block may be just a hash of the header in the previous block or a hash of all or a portion of the information in the previous block, including the file and metadata.

[00144] The file 774_i includes a plurality of data, such as Data 1, Data 2, ..., Data N in sequence. The data are tagged with metadata Metadata 1, Metadata 2, ..., Metadata N which describe the content and/or characteristics

associated with the data. For example, the metadata for each data may include information to indicate a timestamp for the data, process the data, keywords indicating the persons or other content depicted in the data, and/or other features that may be helpful to establish the validity and content of the file as a whole, and particularly its use as a digital evidence, for example, as described in connection with an embodiment discussed below. In addition to the metadata, each data may be tagged with reference REF₁, REF₂, ..., REF_N to a previous data to prevent tampering, gaps in the file, and sequential reference through the file.

[00145] Once the metadata is assigned to the data (e.g., through a smart contract), the metadata cannot be altered without the hash changing, which can easily be identified for invalidation. The metadata, thus, creates a data log of information that may be accessed for use by participants in the blockchain.

[00146] The value 776_i is a hash value or other value computed based on any of the types of information previously discussed. For example, for any given block Block_i, the value for that block may be updated to reflect the processing that was performed for that block, e.g., new hash value, new storage location, new metadata for the associated file, transfer of control or access, identifier, or other action or information to be added. Although the value in each block is shown to be separate from the metadata for the data of the file and header, the value may be based, in part or whole, on this metadata in another embodiment.

[00147] Once the blockchain 770 is formed, at any point in time, the immutable chain-of-custody for the file may be obtained by querying the blockchain for the transaction history of the values across the blocks. This query, or tracking procedure, may begin with decrypting the value of the block that is most currently included (e.g., the last (Nth) block), and then continuing to decrypt the value of the other blocks until the genesis block is reached and the original file is recovered. The decryption may involve decrypting the headers and files and associated metadata at each block, as well.

[00148] Decryption is performed based on the type of encryption that took place in each block. This may involve the use of private keys, public keys, or a public key-private key pair. For example, when asymmetric encryption is used, blockchain participants or a processor in the network may generate a public key and private key pair using a predetermined algorithm. The public key and private key are associated with each other through some mathematical relationship. The public key may be distributed publicly to serve as an address to receive messages from other users, e.g., an IP address or home address. The private key is kept secret and used to digitally sign messages sent to other blockchain participants. The signature is included in the message so that the recipient can verify using the public key of the sender. This way, the recipient can be sure that only the sender could have sent this message.

[00149] Generating a key pair may be analogous to creating an account on the blockchain, but without having to actually register anywhere. Also, every transaction that is executed on the blockchain is digitally signed by the sender.

using their private key. This signature ensures that only the owner of the account can track and process (if within the scope of permission determined by a smart contract) the file of the blockchain.

[00150] FIGS. 8A and 8B illustrate additional examples of use cases for blockchain which may be incorporated and used herein. In particular, FIG. 8A illustrates an example 800 of a blockchain 810 which stores machine learning (artificial intelligence) data. Machine learning relies on vast quantities of historical data (or training data) to build predictive models for accurate prediction on new data. Machine learning software (e.g., neural networks, etc.) can often sift through millions of records to unearth non-intuitive patterns.

[00151] In the example of FIG. 8A, a host platform 820 builds and deploys a machine learning model for predictive monitoring of assets 830. Here, the host platform 820 may be a cloud platform, an industrial server, a web server, a personal computer, a user device, and the like. Assets 830 can be any type of asset (e.g., machine or equipment, etc.) such as an aircraft, locomotive, turbine, medical machinery and equipment, oil and gas equipment, boats, ships, vehicles, and the like. As another example, assets 830 may be non-tangible assets such as stocks, currency, digital coins, insurance, or the like.

[00152] The blockchain 810 can be used to significantly improve both a training process 802 of the machine learning model and a predictive process 804 based on a trained machine learning model. For example, in 802, rather than requiring a data scientist / engineer or other user to collect the data, historical data may be stored by the assets 830 themselves (or through an intermediary, not shown) on the blockchain 810. This can significantly reduce the collection time needed by the host platform 820 when performing predictive model training. For example, using smart contracts, data can be directly and reliably transferred straight from its place of origin to the blockchain 810. By using the blockchain 810 to ensure the security and ownership of the collected data, smart contracts may directly send the data from the assets to the individuals that use the data for building a machine learning model. This allows for sharing of data among the assets 830.

[00153] The collected data may be stored in the blockchain 810 based on a consensus mechanism. The consensus mechanism pulls in (permissioned nodes) to ensure that the data being recorded is verified and accurate. The data recorded is time-stamped, cryptographically signed, and immutable. It is therefore auditable, transparent, and secure. Adding IoT devices which write directly to the blockchain can, in certain cases (i.e. supply chain, healthcare, logistics, etc.), increase both the frequency and accuracy of the data being recorded.

[00154] Furthermore, training of the machine learning model on the collected data may take rounds of refinement and testing by the host platform 820. Each round may be based on additional data or data that was not previously considered to help expand the knowledge of the machine learning model. In 802, the different training and testing steps (and the data associated therewith) may be stored on the blockchain 810 by the host platform 820. Each

refinement of the machine learning model (e.g., changes in variables, weights, etc.) may be stored on the blockchain 810. This provides verifiable proof of how the model was trained and what data was used to train the model.

Furthermore, when the host platform 820 has achieved a finally trained model, the resulting model may be stored on the blockchain 810.

[00155] After the model has been trained, it may be deployed to a live environment where it can make predictions / decisions based on the execution of the final trained machine learning model. For example, in 804, the machine learning model may be used for condition-based maintenance (CBM) for an asset such as an aircraft, a wind turbine, a healthcare machine, and the like. In this example, data fed back from the asset 830 may be input the machine learning model and used to make event predictions such as failure events, error codes, and the like. Determinations made by the execution of the machine learning model at the host platform 820 may be stored on the blockchain 810 to provide auditable / verifiable proof. As one non-limiting example, the machine learning model may predict a future breakdown/failure to a part of the asset 830 and create alert or a notification to replace the part. The data behind this decision may be stored by the host platform 820 on the blockchain 810. In one embodiment the features and/or the actions described and/or depicted herein can occur on or with respect to the blockchain 810.

[00156] New transactions for a blockchain can be gathered together into a new block and added to an existing hash value. This is then encrypted to create a new hash for the new block. This is added to the next list of transactions when they are encrypted, and so on. The result is a chain of blocks that each contain the hash values of all preceding blocks. Computers that store these blocks regularly compare their hash values to ensure that they are all in agreement. Any computer that does not agree, discards the records that are causing the problem. This approach is good for ensuring tamper-resistance of the blockchain, but it is not perfect.

[00157] One way to game this system is for a dishonest user to change the list of transactions in their favor, but in a way that leaves the hash unchanged. This can be done by brute force, in other words by changing a record, encrypting the result, and seeing whether the hash value is the same. And if not, trying again and again and again until it finds a hash that matches. The security of blockchains is based on the belief that ordinary computers can only perform this kind of brute force attack over time scales that are entirely impractical, such as the age of the universe. By contrast, quantum computers are much faster (1000s of times faster) and consequently pose a much greater threat.

[00158] FIG. 8B illustrates an example 850 of a quantum-secure blockchain 852 which implements quantum key distribution (QKD) to protect against a quantum computing attack. In this example, blockchain users can verify each other's identities using QKD. This sends information using quantum particles such as photons, which cannot be copied by an eavesdropper without destroying them. In this way, a sender and a receiver through the blockchain can be sure of each other's identity.

[00159] In the example of FIG. 8B, four users are present 854, 856, 858, and 860. Each of pair of users may share a secret key 862 (*i.e.*, a QKD) between themselves. Since there are four nodes in this example, six pairs of nodes exists, and therefore six different secret keys 862 are used including QKD_{AB}, QKD_{AC}, QKD_{AD}, QKD_{BC}, QKD_{BD}, and QKD_{CD}. Each pair can create a QKD by sending information using quantum particles such as photons, which cannot be copied by an eavesdropper without destroying them. In this way, a pair of users can be sure of each other's identity.

[00160] The operation of the blockchain 852 is based on two procedures (i) creation of transactions, and (ii) construction of blocks that aggregate the new transactions. New transactions may be created similar to a traditional blockchain network. Each transaction may contain information about a sender, a receiver, a time of creation, an amount (or value) to be transferred, a list of reference transactions that justifies the sender has funds for the operation, and the like. This transaction record is then sent to all other nodes where it is entered into a pool of unconfirmed transactions. Here, two parties (*i.e.*, a pair of users from among 854-860) authenticate the transaction by providing their shared secret key 862 (QKD). This quantum signature can be attached to every transaction making it exceedingly difficult to tamper with. Each node checks their entries with respect to a local copy of the blockchain 852 to verify that each transaction has sufficient funds. However, the transactions are not yet confirmed.

[00161] Rather than perform a traditional mining process on the blocks, the blocks may be created in a decentralized manner using a broadcast protocol. At a predetermined period of time (*e.g.*, seconds, minutes, hours, etc.) the network may apply the broadcast protocol to any unconfirmed transaction thereby to achieve a Byzantine agreement (consensus) regarding a correct version of the transaction. For example, each node may possess a private value (transaction data of that particular node). In a first round, nodes transmit their private values to each other. In subsequent rounds, nodes communicate the information they received in the previous round from other nodes. Here, honest nodes are able to create a complete set of transactions within a new block. This new block can be added to the blockchain 852. In one embodiment the features and/or the actions described and/or depicted herein can occur on or with respect to the blockchain 852.

[00162] FIG. 9 illustrates an example system 900 that supports one or more of the example embodiments described and/or depicted herein. The system 900 comprises a computer system/server 902, which is operational with numerous other general purpose or special purpose computing system environments or configurations. Examples of well-known computing systems, environments, and/or configurations that may be suitable for use with computer system/server 902 include, but are not limited to, personal computer systems, server computer systems, thin clients, thick clients, hand-held or laptop devices, multiprocessor systems, microprocessor-based systems, set top boxes, programmable consumer electronics, network PCs, minicomputer systems, mainframe computer systems, and distributed cloud computing environments that include any of the above systems or devices, and the like.

[00163] Computer system/server 902 may be described in the general context of computer system-executable instructions, such as program modules, being executed by a computer system. Generally, program modules may include routines, programs, objects, components, logic, data structures, and so on that perform particular tasks or implement particular abstract data types. Computer system/server 902 may be practiced in distributed cloud computing environments where tasks are performed by remote processing devices that are linked through a communications network. In a distributed cloud computing environment, program modules may be located in both local and remote computer system storage media including memory storage devices.

[00164] As shown in FIG. 9, computer system/server 902 in cloud computing node 900 is shown in the form of a general-purpose computing device. The components of computer system/server 902 may include, but are not limited to, one or more processors or processing units 904, a system memory 906, and a bus that couples various system components including system memory 906 to processor 904.

[00165] The bus represents one or more of any of several types of bus structures, including a memory bus or memory controller, a peripheral bus, an accelerated graphics port, and a processor or local bus using any of a variety of bus architectures. By way of example, and not limitation, such architectures include Industry Standard Architecture (ISA) bus, Micro Channel Architecture (MCA) bus, Enhanced ISA (EISA) bus, Video Electronics Standards Association (VESA) local bus, and Peripheral Component Interconnects (PCI) bus.

[00166] Computer system/server 902 typically includes a variety of computer system readable media. Such media may be any available media that is accessible by computer system/server 902, and it includes both volatile and non-volatile media, removable and non-removable media. System memory 906, in one embodiment, implements the flow diagrams of the other figures. The system memory 906 can include computer system readable media in the form of volatile memory, such as random-access memory (RAM) 910 and/or cache memory 912. Computer system/server 902 may further include other removable/non-removable, volatile/non-volatile computer system storage media. By way of example only, storage system 914 can be provided for reading from and writing to a non-removable, non-volatile magnetic media (not shown and typically called a "hard drive"). Although not shown, a magnetic disk drive for reading from and writing to a removable, non-volatile magnetic disk (e.g., a "floppy disk"), and an optical disk drive for reading from or writing to a removable, non-volatile optical disk such as a CD-ROM, DVD-ROM or other optical media can be provided. In such instances, each can be connected to the bus by one or more data media interfaces. As will be further depicted and described below, memory 906 may include at least one program product having a set (e.g., at least one) of program modules that are configured to carry out the functions of various embodiments of the application.

[00167] Program/utility 916, having a set (at least one) of program modules 918, may be stored in memory 906 by way of example, and not limitation, as well as an operating system, one or more application programs, other program modules, and program data. Each of the operating system, one or more application programs, other program modules,

and program data or some combination thereof, may include an implementation of a networking environment. Program modules 918 generally carry out the functions and/or methodologies of various embodiments of the application as described herein.

[00168] As will be appreciated by one skilled in the art, aspects of the present application may be embodied as a system, method, or computer program product. Accordingly, aspects of the present application may take the form of an entirely hardware embodiment, an entirely software embodiment (including firmware, resident software, micro-code, etc.) or an embodiment combining software and hardware aspects that may all generally be referred to herein as a "circuit," "module" or "system." Furthermore, aspects of the present application may take the form of a computer program product embodied in one or more computer readable medium(s) having computer readable program code embodied thereon.

[00169] Computer system/server 902 may also communicate with one or more external devices 920 such as a keyboard, a pointing device, a display 922, etc.; one or more devices that enable a user to interact with computer system/server 902; and/or any devices (e.g., network card, modem, etc.) that enable computer system/server 902 to communicate with one or more other computing devices. Such communication can occur via I/O interfaces 924. Still yet, computer system/server 902 can communicate with one or more networks such as a local area network (LAN), a general wide area network (WAN), and/or a public network (e.g., the Internet) via network adapter 926. As depicted, network adapter 926 communicates with the other components of computer system/server 902 via a bus. It should be understood that although not shown, other hardware and/or software components could be used in conjunction with computer system/server 902. Examples, include, but are not limited to: microcode, device drivers, redundant processing units, external disk drive arrays, RAID systems, tape drives, and data archival storage systems, etc.

[00170] Although an exemplary embodiment of at least one of a system, method, and non-transitory computer readable medium has been illustrated in the accompanied drawings and described in the foregoing detailed description, it will be understood that the application is not limited to the embodiments disclosed, but is capable of numerous rearrangements, modifications, and substitutions as set forth and defined by the following claims. For example, the capabilities of the system of the various figures can be performed by one or more of the modules or components described herein or in a distributed architecture and may include a transmitter, receiver or pair of both. For example, all or part of the functionality performed by the individual modules, may be performed by one or more of these modules. Further, the functionality described herein may be performed at various times and in relation to various events, internal or external to the modules or components. Also, the information sent between various modules can be sent between the modules via at least one of: a data network, the Internet, a voice network, an Internet Protocol network, a wireless device, a wired device and/or via plurality of protocols. Also, the messages sent or received by any of the modules may be sent or received directly and/or via one or more of the other modules.

[00171] One skilled in the art will appreciate that a "system" could be embodied as a personal computer, a server, a console, a personal digital assistant (PDA), a cell phone, a tablet computing device, a smartphone or any other suitable computing device, or combination of devices. Presenting the above-described functions as being performed by a "system" is not intended to limit the scope of the present application in any way but is intended to provide one example of many embodiments. Indeed, methods, systems and apparatuses disclosed herein may be implemented in localized and distributed forms consistent with computing technology.

[00172] It should be noted that some of the system features described in this specification have been presented as modules, in order to more particularly emphasize their implementation independence. For example, a module may be implemented as a hardware circuit comprising custom very large-scale integration (VLSI) circuits or gate arrays, off-the-shelf semiconductors such as logic chips, transistors, or other discrete components. A module may also be implemented in programmable hardware devices such as field programmable gate arrays, programmable array logic, programmable logic devices, graphics processing units, or the like.

[00173] A module may also be at least partially implemented in software for execution by various types of processors. An identified unit of executable code may, for instance, comprise one or more physical or logical blocks of computer instructions that may, for instance, be organized as an object, procedure, or function. Nevertheless, the executables of an identified module need not be physically located together but may comprise disparate instructions stored in different locations which, when joined logically together, comprise the module and achieve the stated purpose for the module. Further, modules may be stored on a computer-readable medium, which may be, for instance, a hard disk drive, flash device, random access memory (RAM), tape, or any other such medium used to store data.

[00174] Indeed, a module of executable code could be a single instruction, or many instructions, and may even be distributed over several different code segments, among different programs, and across several memory devices. Similarly, operational data may be identified and illustrated herein within modules and may be embodied in any suitable form and organized within any suitable type of data structure. The operational data may be collected as a single data set or may be distributed over different locations including over different storage devices, and may exist, at least partially, merely as electronic signals on a system or network.

[00175] It will be readily understood that the components of the application, as generally described and illustrated in the figures herein, may be arranged and designed in a wide variety of different configurations. Thus, the detailed description of the embodiments is not intended to limit the scope of the application as claimed but is merely representative of selected embodiments of the application.

[00176] One having ordinary skill in the art will readily understand that the above may be practiced with steps in a different order, and/or with hardware elements in configurations that are different than those which are disclosed.

Therefore, although the application has been described based upon these preferred embodiments, it would be apparent to those of skill in the art that certain modifications, variations, and alternative constructions would be apparent.

[00177] While preferred embodiments of the present application have been described, it is to be understood that the embodiments described are illustrative only and the scope of the application is to be defined solely by the appended claims when considered with a full range of equivalents and modifications (e.g., protocols, hardware devices, software platforms etc.) thereto.

CLAIMS

1. An apparatus comprising:

a storage configured to store an index structure that comprises an index of keys from a blockchain ledger, where the keys are stored as nodes in the index structure; and

a processor configured to receive a blockchain transaction comprising a range query for data stored on the blockchain ledger, determine that a validity of the blockchain transaction is not affected by the range query, read a set of keys for the range query from the nodes in the index structure, generate a read set for the blockchain transaction which omits the set of keys read for the range query, and store the read set in a blockchain block.

2. The apparatus of claim 1, wherein the index structure comprises a binary tree in which the keys are arranged as nodes in a hierarchical manner.

3. The apparatus of claim 1, wherein the processor is further configured to build the index structure based on key values that are added to the blockchain ledger.

4. The apparatus of claim 1, wherein the processor is configured to read the set of keys one-by-one from the nodes in the index structure via a next function.

5. The apparatus of claim 1, wherein the processor is further configured to validate the blockchain transaction based on a previous read set for the blockchain transaction.

6. The apparatus of claim 1, wherein the index structure is implemented within logic of chaincode of a blockchain peer node.

7. The apparatus of claim 1, wherein the keys in the index structure correspond to unique values stored on a blockchain of the blockchain ledger.

8. A method comprising:

storing an index structure that comprises an index of keys from a blockchain ledger, where the keys are stored as nodes in the index structure;

receiving a blockchain transaction comprising a range query for data stored on the blockchain ledger;

determining that a validity of the blockchain transaction is not affected by the range query;

reading a set of keys for the range query from the nodes in the index structure;

generating a read set for the blockchain transaction which omits the set of keys read for the range query; and

storing the read set in a blockchain block.

9. The method of claim 8, wherein the index structure comprises a binary tree in which the keys are arranged as nodes in a hierarchical manner.

10. The method of claim 8, further comprising building the index structure based on key values that are added to the blockchain ledger.

11. The method of claim 8, wherein the reading comprises reading the set of keys one-by-one from the nodes in the index structure via a next function.

12. The method of claim 8, further comprising validating the blockchain transaction based on a previous read set for the blockchain transaction.

13. The method of claim 8, wherein the index structure is implemented within logic of chaincode of a blockchain peer node.

14. The method of claim 8, wherein the keys in the index structure correspond to unique values stored on a blockchain of the blockchain ledger.

15. A non-transitory computer readable medium comprising instructions, that when read by a processor, cause the processor to perform a method comprising:

storing an index structure that comprises an index of keys from a blockchain ledger, where the keys are stored as nodes in the index structure;

receiving a blockchain transaction comprising a range query for data stored on the blockchain ledger;

determining that a validity of the blockchain transaction is not affected by the range query;

reading a set of keys for the range query from the nodes in the index structure;

generating a read set for the blockchain transaction which omits the set of keys read for the range query; and

storing the read set in a blockchain block.

16. The non-transitory computer readable medium of claim 15, wherein the index structure comprises a binary tree in which the keys are arranged as nodes in a hierarchical manner.

17. The non-transitory computer readable medium of claim 15, wherein the method further comprises building the index structure based on key values that are added to the blockchain ledger.

18. The non-transitory computer readable medium of claim 15, wherein the reading comprises reading the set of keys one-by-one from the nodes in the index structure via a next function.