

SPLITTABLE SECURITY TOKEN

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This application claims the benefit of United States Provisional Patent Application Serial No. 62/717,575 (Attorney Docket 270.023USPR) filed on August 10, 2018, entitled “SELF-ENFORCING SECURITY TOKEN IMPLEMENTING SMART-CONTRACT-BASED COMPLIANCE RULES CONSULTING SMART-CONTRACT-BASED GLOBAL REGISTRY OF INVESTORS”; United States Provisional Patent Application Serial No. 62/871,543 (Attorney Docket 270.032USPR) filed on July 8, 2019, entitled “SPLITTABLE SECURITY TOKEN”; and United States Provisional Patent Application Serial No. 62/881,121 (Attorney Docket 270.033USPR) filed on July 31, 2019, entitled “UPGRADEABLE SECURITY TOKEN”; all of which are hereby incorporated herein by reference.

[0002] This application is related to the following co-pending United States patent applications, which are hereby incorporated herein by reference:

[0003] United States Patent Application Serial No. _____ (attorney docket number 270.023US01) entitled “SELF-ENFORCING SECURITY TOKEN IMPLEMENTING SMART-CONTRACT-BASED COMPLIANCE RULES CONSULTING SMART-CONTRACT-BASED GLOBAL REGISTRY OF INVESTORS” and filed on even date herewith, which is hereby incorporated herein by reference; and

[0004] United States Patent Application Serial No. _____ (attorney docket number 270.033US01) entitled “UPGRADEABLE SECURITY TOKEN” and filed on even date herewith, which is hereby incorporated herein by reference.

BACKGROUND

[0005] Cryptographic tokens are used in connection with various aspects of blockchains. In examples, smart contracts implemented on an Ethereum blockchain enable the creation and issuance of tokens with complex behaviors attached to them. Ethereum Request for Comments 20 (ERC20) is a technical standard used for smart contracts (on the Ethereum blockchain) implementing tokens. The ERC20 standard is incorporated by reference in its entirety. Additionally, other methods may be

implemented in order to satisfy Title 8 of the Delaware Code Relating to the General Corporation Law.

SUMMARY

[0006] A network node includes at least one processor, at least one memory, and at least one network interface. The network node is configured to be within a plurality of network nodes communicatively coupled in a peer-to-peer network of network nodes implementing a distributed ledger. The network node is configured to be communicatively coupled to at least one remotely located computing device through the at least one network interface. The at least one processor is configured to determine a split ratio for a new split of security tokens held in at least one address of the distributed ledger. The at least one processor is also configured to store an indication of the split ratio without updating token balances for all of the at least one address in a list of addresses holding at least one of the security tokens.

DRAWINGS

[0007] Understanding that the drawings depict only exemplary embodiments and are not therefore to be considered limiting in scope, the exemplary embodiments will be described with additional specificity and detail through the use of the accompanying drawings, in which:

[0008] Figure 1 is a block diagram of an example system using smart-contract-based security tokens that are splittable according to the present systems and methods;

[0009] Figure 2A is a block diagram illustrating a data structure used in the system illustrated in Figure 1;

[0010] Figure 2B is a block diagram illustrating a hierarchy of entities in the system illustrated in Figure 1;

[0011] Figure 3 is a flow diagram illustrating a method for token splitting without formally updating token balances;

[0012] Figure 4 is a flow diagram illustrating a method for transferring a security token following a token split;

[0013] Figure 5A is a block diagram illustrating a table of balances and a splits data structure before a new token split is declared;

[0014] Figure 5B is a block diagram illustrating a table of balances and a splits data structure after the new token split is declared;

[0015] Figure 6A is a block diagram illustrating a table of balances and a splits data structure before a token transaction is performed;

[0016] Figure 6B is a block diagram illustrating a table of balances and a splits data structure after the token transaction is performed;

[0017] Figure 7 is a block diagram illustrating an example computer system with which embodiments of the present disclosure may be utilized; and

[0018] Figure 8 is a block diagram illustrating another example computing device with which embodiments of the present disclosure may be utilized.

[0019] In accordance with common practice, the various described features are not drawn to scale but are drawn to emphasize specific features relevant to the exemplary embodiments.

DETAILED DESCRIPTION

[0020] In the following detailed description, reference is made to the accompanying drawings that form a part hereof, and in which is shown by way of illustration specific illustrative embodiments. However, it is to be understood that other embodiments may be utilized and that logical, mechanical, and electrical changes may be made. Furthermore, the method presented in the drawing figures and the specification is not to be construed as limiting the order in which the individual steps may be performed. The following detailed description is, therefore, not to be taken in a limiting sense.

[0021] A security token is a cryptographic token that represents a security. A security may be any fungible, negotiable financial instrument that holds some type of monetary value. A security may represent an ownership position, a creditor relationship, or rights to ownership as represented by an option. Examples of securities include, without limitation, a piece of real property, at least one commodity, a piece of personal property, at least one bond, at least one derivative, at least one future, at least one fund, at least one currency fund, at least one exchange traded fund,

at least one mutual fund, at least one index fund, at least one bond fund, at least one commodity fund, and/or at least one real estate fund. The security token may be implemented as a smart contract on (e.g., stored on) a distributed ledger (e.g., a blockchain). Various aspects of a given security token (e.g., issuance, transfers, etc.) may be regulated by the Securities and Exchange Commission (SEC) or other regulatory body.

[0022] In some configurations, each security token represents a single share of a company. Accordingly, a share of stock and a security token representing a share of stock may be referred to interchangeably. In examples, a forward stock split is a decision by a company's board of directors to increase the number of shares that are outstanding by issuing more shares to current shareholders (where the increased number of shares are proportionally less valuable). In examples, a forward split might be used to keep a share/token price remains under a certain value, e.g., to make it more attractive to prospective investors. In examples, stock splits may also occur in the opposite way, when a company does a reverse stock split. In examples, a reverse stock split is a type of corporate action that consolidates the number of existing shares of stock into fewer, proportionally more valuable, shares. In examples, a reverse split might be used to ensure that a share/token price remains above a certain value, e.g., to prevent the share/token from being de-listed from an exchange.

[0023] As security tokens become more widely used, the question of how to process stock splits in a blockchain ecosystem has started to emerge. As used herein, the term "split" is used to refer to a forward stock split (or a security token split where each security token represents a single share of a company) or a reverse stock split (or reverse security token split where each security token represents a single share of a company). The terms "security token" and "token" may be used interchangeably herein. The terms "security token holder," "token holder," and "shareholder" may be used interchangeably herein.

[0024] A record of various security token transactions may be stored in a distributed ledger. Security token transactions may require the transfer of some small quantity of cryptocurrency. In examples, if a security token is implemented on an Ethereum blockchain, transfers (e.g., purchases and sales) of the security token may include the transfer of one or more Ethereum tokens to record the transfer on the Ethereum blockchain. Therefore, manually changing the existing number of shares for all

shareholders, in response to a split, can be expensive when there are many shareholders of a company's security tokens. Such expense would typically be borne by the issuing company and is undesirable.

[0025] Instead, the present systems and methods enable security tokens to be "split" in a way that minimizes costs. Specifically, as described below, the actual token balances of each shareholder may not be updated in response to a split. Instead, split values may be tracked and stored for each declared split, which indicate a multiplier that should be applied to the token balance for the shareholder. In order to reduce blockchain-related fees, the actual balances may not be updated in response to a split. Rather, the balances may be updated when security tokens are actually transferred following a split. Otherwise, the added split values can be used to determine split-updated balances without having to commit security token transactions to the distributed ledger.

[0026] More description of security tokens, including issuance, transfer, and regulations pertaining thereto, can be found in U.S. Provisional Application No. 62/717,575 (Attorney Docket 270.023USPR) filed on August 10, 2018, entitled "SELF-ENFORCING SECURITY TOKEN IMPLEMENTING SMART-CONTRACT-BASED COMPLIANCE RULES CONSULTING SMART-CONTRACT-BASED GLOBAL REGISTRY OF INVESTORS", the entirety of which is incorporated by reference herein.

[0027] The examples described herein may be compatible with a custodial model where custodians hold security tokens on behalf of investors and/or broker dealers. However, the systems herein are also compatible with non-custodial models, e.g., where investors and/or broker dealers hold their security tokens directly with fewer or no intermediaries.

[0028] Various method names may be used herein as examples. It should be appreciated that any method names are illustrative in purpose, and not limiting. Additionally, optional devices and method steps are indicated with dashed lines in the Figures.

[0029] Figure 1 is a block diagram of an example system 100 using smart-contract-based security tokens 102 that are splittable according to the present systems and methods. The system 100 may include one or more security tokens 102, a system

owner 112, an issuer 114, and at least one network node 140 implementing a distributed ledger 122. The system 100 may also include various optional entities and/or devices, as described below.

[0030] An investor 120 may be a person or entity that has previously, or is considering, receiving or transferring a security token 102. In examples, an investor 120 may be a customer of a broker dealer 118 (or directly of the custodian 116), where the investor 120 may send orders to the broker dealer 118 (or directly to the custodian 116) relating to the purchase and/or sale of one or more security tokens 102. Alternatively, an investor 120 may be an “external investor” that does not need to connect to a broker dealer 118 or a custodian 116. Instead, an external investor may call the relevant methods described herein to initiate security token transactions, e.g., methods in the token 102 itself and/or the global registry 108, described below.

[0031] The term “security token transaction” or similar refers to any transfer of a security token 102, and may be used to describe an issuance, offering, gift, purchase, sale, and/or employee compensation carried out with the transfer of the security token 102. A security token 102 transaction may transfer a quantity of security tokens 102 from one or more input addresses (belonging to one or more investors 120, broker dealers 118, and/or custodians 116) to at least one output address (belonging to one or more investors 120, broker dealers 118, and/or custodians 116). As used herein, the terms “buyer” and “seller” refer to investors 120 that are transferring and receiving a security token 102, respectively. Security token transaction may be performed on-exchange or off-exchange.

[0032] The broker dealer 118 in the system 100 may be a person or entity that purchases or sells security tokens 102 for its own account and/or on behalf of its customers. During onboarding, a broker dealer 118 may create an account for a new customer, including collection of information, e.g., a broker dealer 118 may onboard an investor 120 as a new customer. This may include performing anti-money laundering and/or know-your-customer (AML/KYC) checks on its customers, e.g., provided by an identity services provider 126.

[0033] The issuer 114 may be a person or entity that instructs the owner 112 to issue security tokens 102, i.e., the owner 112 may own one or more originating smart contracts 113 that deploy/issue security tokens 102. The system 100 may include

many issuers 114. In examples, the issuer 114 may be a company where each security token 102 represents a share of the company.

[0034] The broker dealer 118 may have an account with a custodian 116, i.e., a person or entity that holds custody, possession, and/or ownership of security tokens 102 for one or more broker dealers 118. The custodian 116 may be a person or entity that holds custody, possession, and/or ownership of security tokens 102 on behalf of many broker dealers 118 (which broker dealers 118 may have many investors 120 as its customers). The system 100 may include many custodians 116 and/or many broker dealers 118.

[0035] An owner 112 may be a person or entity that owns, deploys, re-deploys, and/or transfers smart contracts in the system 100. In examples, the owner 112 may use a method call (e.g., that is executed by a virtual machine) to the originating smart contract 113 to deploy/issue security tokens 102. The owner 112 may also be referred to as the administrator of the system 100. The system 100 may only have a one or more owners 112. Optionally, the owner 112 may provide data services describing the “market” for a security token 102 based on the security token 102 transactions executed in the system 100, e.g., a feed indicating volume of security token 102 transactions, transaction price for the most recent security token 102 transaction, etc.

[0036] As described above, the terms owner 112, issuer 114, custodian 116, broker dealer 118, and investor 120 may refer to a person or entity filling the various roles in the system 100. Alternatively, these terms may refer to a computing device used by the owner 112, issuer 114, custodian 116, broker dealer 118, and investor 120, respectively. When referring to a computing device, each of the owner 112, issuer 114, custodian 116, broker dealer 118, and investor 120 may be implemented with one or more processors that execute instructions in a memory on the respective computing device. Each respective computing device may be a mobile device, such as a cell phone or tablet, personal computer, or a server configured to send and receive instructions and/or other data to other computing devices, e.g., via a network 138, such as the Internet.

[0037] A security token exchange 124 may be a marketplace or a business entity that operates the marketplace (or one or more computing devices operated by a security token exchange 124) in which security tokens 102, commodities, derivatives and/or

other financial instruments are traded. In examples, the security token exchange 124 may record successfully executed transactions on the distributed ledger 122. The security token exchange 124 may be custodial, i.e., where custodians 116 hold and transact security tokens 102 on behalf of broker dealers 118 and/or investors 120.

[0038] The network node 140 may be a computing device implemented using one or more processors that execute instructions stored in memory to implement functionality described herein. In examples, a network node 140 may run a virtual machine (e.g., Ethereum Virtual Machine) that executes any of the smart contracts described herein.

[0039] The term “distributed ledger” refers to an electronic ledger that is distributed across multiple interconnected network nodes 140 (i.e., computing devices executing instructions stored in memory), where more than one of the network nodes 140 stores a copy of the distributed ledger 122. In examples, one or more network nodes 140 in a peer-to-peer network may implement the distributed ledger 122.

[0040] The distributed ledger 122 may implement one or more blockchains to validate the data stored within the distributed ledger 122. A blockchain is a verifiable permanent ledger constructed one block at a time with a proof-of-work seal (such as a hash) affixed to each block that validates that block. In a blockchain, the hash of the previous block is included in the current block, and therefore by recursion the current hash also validates all previous blocks back to the original genesis block. Inserting a hash into a blockchain permanently records that hash and acts as a notary verifying the time stamped proof-of-existence of the hashed data at the moment in time that block is added to the chain. Any future blocks add a layer of protection from manipulation of the data stored in the chain or a chain re-organization and therefore provide additional certainty that no changes can be made to blocks earlier in the chain. A blockchain is an implementation of a distributed ledger 122, and may be public (e.g., viewable by anyone) or private (e.g., viewable only to authorized users or entities). Exemplary blockchains include, but are not limited to, the Bitcoin blockchain, the Ethereum blockchain, BigchainDB, Billon, Chain, Corda, Credits, Elements, Monax, Fabric, HydraChain, Hyperledger, Multichain, Openchain, Quorum, Ravencoin, Sawtooth, and Stellar. Preferably, the distributed ledger 122 is the public Ethereum blockchain, however, other implementations are possible. If a private blockchain is used, a hash of the private blockchain may be periodically

committed to (i.e., recorded on) a public blockchain, e.g., the Ethereum blockchain or Bitcoin blockchain.

[0041] Additionally, or alternatively, the distributed ledger 122 may implement a directed acyclic graph (DAG), e.g., IOTA or Hashgraph that uses a gossip protocol to share information between network nodes 140 of the system 100. Furthermore, consensus may be reached in the distributed ledger 122 without proof-of-work and can instead use proof-of-stake. Furthermore, any Merkle tree (or hash tree) that connects different sets of data using cryptographic hashes may be implemented by the distributed ledger 122.

[0042] The term “wallet” refers to a software program, digital file, and/or memory used to store and/or manage digital assets, such as security tokens 102. Although the present systems and methods are described herein using security tokens 102, they are also compatible with any type of digital asset. In examples, a wallet may be defined by one or more private keys, one or more public keys derived from one or more private keys, and/or one or more addresses derived from one or more private keys and/or one or more public keys. In examples, a wallet may be defined by one or more private account keys (and optional corresponding public account key(s)), each of which may have one or more child and/or grandchild transaction keys. The term “account” may be used to refer to an address on the distributed ledger 122, e.g., an Ethereum address in a hierarchical deterministic (HD) wallet.

[0043] The term “smart contract” refers to a set of conditional logic that may be implemented in software, e.g., one or more sequential steps that are executed in response to one or more relevant conditions being satisfied. A smart contract may be stored at an address of a distributed ledger 122. In examples, smart contracts may be programmed in the Solidity programming language. Smart contracts may be executed by a processor on a network node 140, e.g., that is running a virtual machine, such as the Ethereum Virtual Machine (EVM).

[0044] Security tokens 102 may be issued privately, e.g., as part of a pre-initial public offering (pre-IPO) to a dark pool of investors that qualify under the relevant exemptions to the SEC regulations. Alternatively, security tokens 102 may be issued as part of an initial public offering (IPO) under SEC regulations. The owner 112 may deploy/issue one or more security tokens 102, e.g., with a security token offering

(STO). Optionally, the issuer 114 may instruct the owner 112 to deploy/issue one or more security tokens 102, e.g., with a security token offering (STO). An STO may include deploying/issuing or selling a security where the security is represented by a security token 102. Optionally, an STO may include allocating them to one or more wallets, e.g., belonging to the owner 112 or issuer 114. In examples, all deployed security tokens 102 may be initially allocated to a wallet belonging to the owner 112 or issuer 114 before allocation to individual investors 120. Alternatively, the deployed security tokens 102 may be directly allocated to wallets owned by the individual investors 120. An STO may include one or more originating smart contracts 113 issuing one or more security tokens 102.

[0045] Following initial deployment and allocation, the security token 102 may be used to transfer the asset that the security token 102 represents from one party to another, e.g., on a security token exchange (STX) 124 or in a peer to peer (P2P) transaction (if permitted by the security token 102). Records of issuance and/or subsequent transactions of the security token 102 may be committed to (i.e., recorded in) the security token 102 itself and/or the distributed ledger 122.

[0046] The system 100 may be used to deploy, allocate and/or subsequently transfer various different types of security tokens 102. In examples, company A may deploy X security tokens 102 that each represent a share of company A and have a symbol AAAX, company B may deploy Y security tokens 102 that each represent a share of company B and have a symbol BBBX, and company C may deploy Z security tokens 102 that each represent a share of company C and have a symbol CCCX. Furthermore, security tokens 102 AAAX, BBBX, and CCCX may all be deployed, allocated, and/or tradeable in the same system 100.

[0047] Ethereum Request for Comments 20 (ERC20) is a standard defining a set of methods. To be compliant with ERC20, a cryptographic token (e.g., utility, cryptocurrency, or security) must implement these methods, although the owner of the token can implement the methods in their own way. One of the advantages of ERC20 is that each application, wallet, exchange, or interface doesn't have to be tailored to each specific token 102. Rather, if a project supports the ERC20 standard, it may support many different ERC20 tokens 102.

[0048] In order to comply with ERC20, the security token 102 may include token logic and/or parameters 104 that implements six required methods, including totalSupply(), balanceOf(), transfer(), transferFrom(), approve(), and allowance(). These required methods describe how tokens can be transferred and how token-related data can be accessed. The token logic/parameters 104 may also implement various events to comply with ERC20, e.g., Transfer() and Approval(). These events describe formatting guidelines for transfers and approvals.

[0049] The transfer() method/function 109 may be called when a security token 102 transaction is requested, and the transfer() method 109 may optionally call method(s) in one or more other smart contracts, e.g., in the token logic/parameters 104 and/or the global registry 108. The transfer() method 109 may be a particular implementation that accounts for split(s) in the security token 102. However, it is understood that other implementations of the transfer() method 109 are possible, which do not account for split(s) in the security token 102.

[0050] The token logic/parameters 104 may also implement one or more optional methods to comply with ERC20, such as name(), symbol(), and/or decimals(). These optional methods can be used to assign the security token 102 a name and a symbol, as well as define the number of decimals the security token 102 uses, respectively. These methods and events may be executed by a virtual machine (e.g., the Ethereum Virtual Machine) on a network node 140, and are further described in the ERC20 standard.

[0051] The token logic and/or parameters 104 may also include data that indicates the name of the security token 102, how many security tokens 102 of the same type were deployed or issued, and/or a symbol for the security token 102. In examples, the symbol of a security token 102 may be a multi-character (e.g., four or five) identifier for the security token 102 similar to a stock ticker symbol. The symbol of a security token 102 may be registered with Financial Industry Regulatory Authority (FINRA). The security token 102 symbol issued by a company may preferably be an extended version of the company's ticker symbol, e.g., by concatenating a particular character (e.g., Q, X, or Z) to the end of the company's ticker symbol.

[0052] The security token 102 may also include a table of balances 105, e.g., stored on a distributed ledger 122. The table of balances 105 may indicate balances of all

holders of a particular type of security tokens 102, e.g., the table of balances 105 may indicate what addresses (e.g., Ethereum addresses) hold what quantity of a security token 102. Alternatively, instead of indexing by investor 120 (or custodian 116 or broker dealer 118) address, the table of balances 105 may be indexed in other ways, e.g., a hash of the investor's name, etc. The table of balances 105 for a security token 102 may be public, semi-public, or private, and may be anonymous such that balances of security tokens 102 held by a particular investor 120, custodian 116, or broker dealer 118 cannot be derived purely from publicly-available information.

[0053] The table of balances 105 may account for all security tokens 102 that have been deployed or issued. When a transfer of the security token 102 occurs (i.e., balances change), the table of balances 105 may be updated and distributed to each copy of the security token 102 stored on the different network nodes 140. In examples, the transfer of a security token 102 may include updating the table of balances 105 in the security token 102. Table 1 is an example table of balances 105 that includes a key/value pair for each investor address:

Investor, Custodian, or Broker Dealer Address	Security Token Balance
0x7d2a3d9f938e13cd947eb85a924bb34df8dd866	525
0x1d2a3d9f991e13c5da7ec05a1c7fe734df8dd84a	50
0x1d2a3da66b8f13c5da7e9d5afc7fe734df8dd826	15000
0x1d2a3d9f938e13c4fa7ec0577c7fe734df8dd55e	1100

Table 1

In some examples, a table of balances 105 may be stored as a multi-object data structure (e.g., array) that is indexed by address (illustrated in the left column in Table 1) with the value at each index being a security token balance (illustrated in the right column in Table 1) associated with the address.

[0054] However, security tokens 102 (unlike other types of tokens) may be regulated by the Securities and Exchange Commission (SEC) in the United States. Accordingly, it may be desirable to satisfy Title 8 of the Delaware Code Relating to the General Corporation Law, which is important because many publicly traded companies are incorporated in Delaware. Additional methods satisfying Title 8 are referred to as Title 8 methods herein and may be implemented in a variety of ways, e.g., with or

without Ethereum Request for Comments 884 (ERC884). Specifically, the Title 8 methods allows for the creation of tokens where each token represents a single share issued by a Delaware corporation, i.e., such tokens may be designed to represent equity issued by any Delaware corporation, whether private or public.

[0055] Accordingly, the token logic and/or parameters 104 may implement one or more Title 8 methods (beyond the six required by ERC20): including addVerified(), removeVerified(), updateVerified(), cancelAndReissue(), isVerified(), isHolder(), holderCount(), holderAt(), hasHash(), isSuperseded(), and/or getCurrentFor(). The token logic/parameters 104 may also implement one or more events (beyond the six required for ERC20 compatibility), e.g., verifiedAddressAdded(), VerifiedAddressRemoved(), VerifiedAddressUpdate(), and/or VerifiedAddressSuperseded(). These Title 8 methods and events may be executed by a virtual machine on a network node 140.

[0056] After a company has issued shares (in the form of security tokens), it may decide to change the number of existing shares using a split. A simple way to perform a split action on issued security tokens 102 would be to traverse through the list of addresses holding at least one security token 102 (e.g., in a table of balances 105) and update their token balances one by one. With a security token 102 running on Ethereum, this is not practical because a small blockchain-related fee (also referred to as “gas”) is required to conduct a transaction or execute a method in a smart contract on a blockchain platform, e.g., “gas” is the pricing value required to successfully conduct a transaction or execute a smart contract on the Ethereum blockchain platform. Furthermore, separate transactions (instead of in batches) may be required to update the balance of each individual token holder, due to upper gas limits. This process could be expensive and take a long time to complete. Instead, the present systems and methods may simulate splits until an actual transfer of security token(s) 102 is made, thereby saving gas on unnecessary transactions. In other words, the functionality described herein can implement splits at a lower cost than updating the token balances by manually iterating through the shareholder list (e.g., in a table of balances 105).

[0057] In order to enable split functionality, the system 100 may implement functionality beyond ERC20 and the Title 8 methods. In examples, in addition to implementing methods required by ERC20 and/or the Title 8 methods, the security

token 102 may implement the following split-related methods 152: (1) the transfer() method 109 (mentioned above) that transfers a quantity of security tokens 102 from an input address to at least one output address, while accounting for splits of the security token 102 that have taken place, but are not yet accounted for in the token balances; (2) an update split method 142 that initiates a split of a token (updateSplit()); (3) a balance check method 144 that determines a security token balance for a token holder that accounts for one or more splits of the token (balanceOf()), without actually updating token balances in the list addresses holding tokens (e.g., the table of balances 105); (4) an update balances method 146 to update token balances of one or more token holders based on one or more splits of the token (updateBalances()); (5) an update total tokens method 148 to update the total amount of outstanding tokens following one or more splits of the token (splitTotalSupply()); and/or (6) a remove shareholder method 150 to remove fractional shareholders following a split (removeShareholder()). The split-related methods 152 may be executed by a processor on a network node 140, e.g., that is running a virtual machine, such as the Ethereum Virtual Machine (EVM).

[0058] Alternatively, instead of the security token 102 implementing the split-related methods 152, the token logic and/or parameters 104 in the security token 102 may include pointer(s) (not shown) that indicate the address(es) of one or more split-related methods 152 implemented outside of the security token 102, e.g., the actual split-related methods 152 may be implemented in the optional global registry 108 described below. In this alternative configuration, the split-related methods 152 may be updated, if necessary, without modifying the security token 102 itself.

[0059] In some configurations, as described below, a split-related method 152 may itself call one or more other split-related method 152. For example, the transfer() method 109 might call the update balances method 146, which calls the balance check method 144 for the input and output addresses. Furthermore, the split-related methods 152 may be executed on network node 140 implementing the distributed ledger 122.

[0060] When determining the token balance for a shareholder, a quick calculation may be performed (as part of the balanceOf() method 144) to obtain the split-updated balance (the balance after the split(s) is/are applied, if applicable), which is also called newBalance herein. In examples, this method 144 is free to execute on the blockchain because it is read-only, i.e., it does not require any gas (blockchain-related fee) to

execute. Then, when a transfer is requested (e.g., the transfer() method 109 in the security token 102 is called), the actual token balance is updated for both sides of the trade (input address and output address(es)), and the desired number of tokens is transferred, thus bringing the shareholders' token balances up-to-date with the current split. A record of the security token transaction may be recorded on the distributed ledger 122 during the transfer() method 109.

[0061] In addition to split-related methods 152, the token logic and/or parameters 104 (in the security token 102 itself) may also include one or more split-related parameters 154. Specifically, when a stock goes through a split (including reverse splits), an event is recorded. Any subsequent splits are also tracked/recorded in the order that they occurred. The split events may be recorded in a data structure (called splits herein) using the updateSplit() method, where the data structure (splits) has an object for each split in the token's history. For each shareholder/token holder, the last split that is reflected in their token balance is also tracked. The last split may be recorded in a data structure called holdersSplit. On a reverse split, an audit of the shareholder list may be performed to determine the number of tokens (cachedOutTokenQuantity) that will be eliminated due to partial tokens not being allowed (according to regulatory and/or technical limitations). The split-related parameters 154 may also include an index (currentSplitIndex) indicating the most recent split for the token 102 (without respect to any specific token holder). It should be noted that the specific split-related parameters 154 may be implemented using any suitable data structures, e.g., a struct, array, linked list, integer, etc. Furthermore, one or more of the split-related parameters 154 may not be used in some implementations, e.g., if the functionality described herein can be accomplished without (or with a different) parameter.

[0062] The split-related parameters 154 are implemented in the security token 102, e.g., in each security token 102 on each network node 140 implementing the distributed ledger 122. In some configurations, where a split-related parameter 154 is being updated, stored, adjusted, removed, changed, added, or any variants herein, it may be happening on a distributed ledger 122, e.g., an Ethereum blockchain. In some configurations, updating a split-related parameter 154 may include executing a transaction on the distributed ledger 122.

[0063] The system 100 may also include an optional global registry 108, which includes one or more smart contracts, such as a data storage smart contract 128 that stores data. The data stored in the data storage smart contract 128 may include a mapping of addresses (in the table of balances 105) to investors 120, personally identifiable information (PII) of an investor 120, and/or any other investor information that may be used. Optionally, pointer(s) in the token logic and/or parameters 104 reference split-related methods 152 implemented in the optional global registry 108. Additionally, in some configurations, one or more split-related parameters 154 may be stored in the data storage smart contract 128.

[0064] The smart contracts 128 in the global registry 108 may be stored on a public distributed ledger 122. Alternatively, the global registry 108 could be a traditional database that stores attributes of the security tokens 102, custodians 116, broker dealers 118, and/or investors 120 in the system 100, e.g., stored on a single computing device. Alternatively, the global registry 108 may be a database that stores attributes of the security tokens 102, custodians 116, broker dealers 118, and/or investors 120 on the distributed ledger 122 without utilizing smart contracts 128.

[0065] A data structure 110 in the data storage smart contract 128 may store at least one investor element, each with information about a particular investor 120. Additionally, each investor element may include an address (e.g., Ethereum address) of the investor's broker dealer 118, and/or various regulatory accreditation information (such as whether AML/KYC has been performed for the investor 120, the investor 120 country of origin, length of time a particular security token 102 has been owned by the investor 120, etc.). The attributes in an investor element (e.g., a value in a slot in an investor element) may be stored as comma-separated values, a hash table, a mapping, a dictionary, and/or key/value pairs, etc., e.g., as illustrated in Figure 2A. Each investor element may be created by an owner 112, issuer 114, broker dealer 118, custodian 116, and/or an investor 120 (e.g., if the investor 120 is an external investor).

[0066] In examples, the data storage smart contract 128 may store at least one investor element for each investor 120. The data storage smart contract 128 may store multiple investor elements for a particular investor 120 if the investor 120 is a customer of multiple different broker dealers 118. In examples, only the broker dealer 118 associated with the investor element (or their agent) may update the information for that investor element. Different investor elements for the same investor 120 (at

different broker dealers 118) may be stored in the same data structure 110.

Alternatively, other configurations in which the data storage smart contract 128 stores a single investor element (for an investor 120 having accounts with multiple broker dealers 118) are possible. Optionally, an interface smart contract (not shown) may be used to add, remove, and/or update investor elements in one or more data structures 110.

[0067] The attributes in an investor element may include information about the investor 120, e.g., collected by the broker dealer 118 during onboarding. Such information may include: customer name; tax identification number (e.g., Social Security number); address; telephone number; email address; date of birth; driver's license, passport information or information from other government-issued identification; employment status and occupation; whether the customer is employed by a brokerage firm; annual income; net worth; account investment objectives; and/or any other data that can be used to distinguish individual identity. The investor elements in the data storage smart contract 128 may be updated, as appropriate. For example, when the investor 120 closes their account with a broker dealer 118, the corresponding investor element may be removed from the data storage smart contract 128. Additionally or alternatively, the investor 120 may be responsible for creating, updating, and/or removing the attributes in the investor element 223 associated with itself, e.g., if an investor 120 is an external investor.

[0068] As mentioned above, the data storage smart contract 128 may include an investor element for each instance of the investor 120 per broker dealer 118 that the investor 120 has an account with, i.e., the data storage smart contract 128 may store multiple investor elements for an investor 120 with accounts at multiple broker dealers 118. In examples, the investor element(s) for an investor 120 may only be created, updated, and/or removed by a broker dealer 118 that the investor 120 is associated with, however, other configurations are possible. Accordingly, each broker dealer 118 can preferably create, modify, or remove an investor element that is associated with investors 120 that are customers of the respective broker dealer 118.

[0069] Optionally, one or more other types of elements may be stored in the data storage smart contract 128 that include attributes about various actors in the system 100. In examples, the other elements (i.e., other than investor elements) may include at least one custodian element, at least one broker dealer element, and/or at least one

security token element. In examples, the broker dealer element for a broker dealer 118 may only be created, updated, and/or removed by a custodian 116 that the broker dealer 118 is associated with, however, other configurations are possible. Optionally, an issuer element and/or security token exchange element may be stored for every issuer 114 and/or security token exchange 124 in the system. The attributes in one of the other elements (e.g., a value in a slot in an element) may be stored as comma-separated values, a hash table, a mapping, a dictionary, and/or key/value pairs, etc.

[0070] Figure 2A is a block diagram illustrating a data structure 110 used in the system 100 illustrated in Figure 1. One or more data structures 110 may be stored in the data storage smart contract 128. Each data structure 110 may include one or more elements 223A-M, where each element 223 corresponds to (i.e., includes information about) a particular custodian 116, broker dealer 118, or investor 120. Optionally, security token elements 223 may also be stored in one or more data structures 110.

[0071] In some configurations, all elements 223 in a particular data structure 110 correspond to the same type of entity, e.g., all elements 223 correspond to investors 120 in the system 100. In this configuration, the data storage smart contract 128 may include a first data structure 110 for all investors 120, a second data structure 110 for all custodians 116, a third data structure 110 for all broker dealers 118, and/or a fourth data structure 110 for all security tokens 102 in the system 100. Alternatively, a particular data structure 110 may include elements 223 corresponding to different types of entities, i.e., investor element(s) 223, custodian element(s) 223, and broker dealer element(s) 223 may be included in a single data structure 110.

[0072] In some configurations, each element 223 may include two storage slots 215A-N where each storage slot is a fixed number of bytes. For example, a storage slot may be 32 bytes long (or any suitable length), e.g., each storage slot 215 may be slot in the Solidity programming language. Although each element 223 is illustrated in Figure 2A as including two storage slots 215 (i.e., $N=2M$), an element 223 may include only a single storage slot 215 (i.e., $N=M$) or more than two storage slots 215 (i.e., $N>2M$).

[0073] Each storage slot 215 may include a key 219A-N and a value 221A-N, i.e., a key/value pair. Each key 219 may be a nested structure with a first level indicating the type of entity the element 223 corresponds to (e.g., custodian 116, broker dealer 118,

investor 120, etc.) and a second level that indicates an address (e.g., Ethereum address) owned by the particular entity. In this way, the key 215 may be encoded (with the type of entity). Instead of indexing by address, the second level of the key 219 may be indexed in other ways, e.g., a hash of the entity/investor/security token name, etc. Optionally, the keys 219 may include more than two levels. If the element 223 includes two slots 215, the key 219 in the first storage slot 215 of the element 223 may be the same or different than the key 219 in the second slot 215 of the element 223.

[0074] The information included in the values 221 may vary depending on the type of entity the element 223 corresponds to. Each value 221 in the data structure 110 may be a nested structure with a first level indicating the type of entity the element 223 corresponds to (e.g., custodian 116, broker dealer 118, investor 120, etc.) and one or more additional levels indicating additional information. In this way, the value 221 may be encoded (with the type of entity) so that a computing device (e.g., a network node 140) will know how to interpret the data in the value 221. In other words, a computing device (e.g., a network node 140) will interpret the data in the value 221 differently based on the encoding, i.e., the type indicated in the value 221. Multiple parameters may be bit-shifted (and/or concatenated) so they fit in the same value 221.

[0075] For example, the value 221 in a first storage slot 215 in a particular investor element 223 may include one or more bits indicating “investor” in the first level and the PII hash for the investor 120 in the second level. Furthermore, the value 221 in the second storage slot 215 in the investor element 223 may include one or more bits indicating other attributes of the investor 120.

[0076] Examples of other data structures 110 include custodian elements 223, broker dealer elements 223, and/or optional security token elements 223. Each custodian element 223 may include attribute(s) about a particular custodian 116. Optionally, a custodian element 223 may reference an address (e.g., Ethereum address) of one or more investor elements 223, broker dealer elements 223, and/or security token elements 223 that the custodian element 223 is associated with. For example, a custodian element 223 may include an address (e.g., Ethereum address) of one or more broker dealers 118 that the custodian 116 is associated with.

[0077] Each broker dealer element 223 may include attribute(s) about a particular broker dealer 118. The value(s) 221 in a particular broker dealer element 223 may include one or more bits indicating “broker dealer” in the first level and one or more bits indicating attribute(s) about the broker dealer 118 in a second level. Optionally, a broker dealer element 223 may reference an address (e.g., Ethereum address) of one or more investor elements 223, custodian element 223, and/or security token elements 223 that the broker dealer element 223 is associated with. For example, a broker dealer element 223 may include an address (e.g., Ethereum address) of one or more custodians 116 that the broker dealer 118 is associated with.

[0078] Each security token element 223 may include attribute(s) about a particular security token 102. The value(s) 221 in an optional security token element 223 may include one or more bits indicating “security token” in the first level and one or more bits indicating attribute(s) about the security token 102 in a second level. Various types of information about the security token 102 may be included in the value(s) a security token element 223.

[0079] Each optional issuer element 223 may include attribute(s) about a particular issuer 114. The value(s) 221 in a particular optional issuer element 223 may include one or more bits indicating “issuer” in the first level and one or more bits indicating attribute(s) about the issuer 114 in a second level. Optionally, an issuer element 223 may reference addresses (e.g., Ethereum addresses) of one or more investors 120, custodians 116, broker dealers 118 and/or security token 102 that the issuer 114 is associated with.

[0080] Figure 2B is a block diagram illustrating a hierarchy of entities in the system 100 illustrated in Figure 1. Specifically, Figure 2B shows relationships between one or more investors 120, optional one or more broker dealers 118, optional one or more custodians 116, and/or one or more owners 112 of the system 100.

[0081] In examples, the hierarchy in Figure 2B may define data permissions in the system 100. The hierarchy may be stored in any suitable form and/or may be implicit in the data stored in various elements 223.

[0082] A data structure 110 may store multiple investor elements 223 for an investor 120 that has accounts at multiple broker dealers 118. In examples, a first investor element 223A may describe a particular investor 120, according to data collected by a

first broker dealer 118A, e.g., during onboarding performed by the first broker dealer 118A. Similarly, a different investor element 223B may describe the same particular investor 120, according to data collected by a different broker dealer 118B, e.g., during onboarding performed by the different broker dealer 118B. Additionally or alternatively, an investor element 120I may be stored for an external investor that is not a customer of a broker dealer 118 or custodian 116.

[0083] Any smart contracts implementing the global registry 108 may be executed by a virtual machine (e.g., the Ethereum Virtual Machine) running on a network node 140. Each entity in the system 100 may have different data permissions to the data storage smart contract 128. In examples, the owner 112 may add, remove, and/or change elements for custodians 116, broker dealers 118 and/or investors 120 in its own hierarchy. A custodian 116 may add, remove, and/or change elements for broker dealers 118 and/or investors 120 in its own hierarchy. A broker dealer 118 may add, remove, and/or change elements for investors 120 in its own hierarchy. An owner 112 may add, remove, and/or change elements for custodians 116, broker dealers 118 and/or investors 120 in its own hierarchy. In some configurations, an investor 120 may have permission to call the transfer() method 109 that transfers funds and/or security tokens 102 between two addresses.

[0084] Figure 3 is a flow diagram illustrating a method 300 for token splitting without formally updating token balances. The method 300 may be executed by at least one processor running a virtual machine in the system 100, e.g., at least one processor in a network node 140. The method 300 may include executing one or more split-related methods 152 implemented in a security token 102, e.g., stored in a distributed ledger 122. Additionally, the method 300 may utilize one or more split-related parameters 154 in a security token 102, e.g., stored in a distributed ledger 122. The method 300 may incur low (or no) blockchain-related fees (gas).

[0085] The method 300 may begin at step 302 where the at least one processor determines a split ratio for a new split of security tokens 102 held in at least one address of a distributed ledger 122, e.g., an Ethereum blockchain. This may include an issuer 114 (or owner 112) setting a numerator and a denominator, which together form the split ratio. The new split may be a forward split or a reverse split. For example, a numerator of 2 with a denominator of 1 (considered a forward split since numerator > denominator) means that the number of existing security tokens 102 (of

the specified type) is meant to double after the new split, with each resulting security token 102 having half the value compared to before the split. As another example, a numerator of 1 with a denominator of 2 (considered a reverse split since numerator < denominator) means that the number of existing security tokens (of the specified type) is meant to half after the new split, with each resulting security token having twice the value compared to before the split.

[0086] As discussed above, the issuer 114 may be a company (or a computing device owned and/or controlled by the company) that issues security tokens 102, each representing a share of the company. The split ratio may be determined by the company (e.g., the board of directors of the company) when deciding to initiate a split on security tokens 102 representing the company's shares. In some configurations, the issuer 114 may call the update split method 142 (implemented in the security token 102) and pass the split ratio as one of the parameters. For example, the issuer may invoke the following function call: `updateSplit(numerator, denominator)`.

[0087] The method 300 may proceed at optional step 304 where the at least one processor determines a number of the security tokens that will be paid out in response to the new split. Specifically, the system 100 may not allow for fractional shares. Therefore, token holders will be paid out for their resulting fractional tokens (e.g., half tokens) following a split. This is referred to as "burning" these fractional tokens 102, which reduces the total number of outstanding security tokens (of the specified type). The number of tokens that will be paid out/burned is referred to as `cachedOutTokenQuantity`, and may be stored as a split-related parameter 154 in the security token 102. Fractional tokens typically result from reverse splits (where the total number of tokens/shares is reduced, each resulting share having a proportionally greater value post-split), but can also result from a "forward" split with a non-integer (e.g., $3/2$) split ratio.

[0088] Optional step 304 may include an audit of a list of addresses holding token(s) (e.g., in a table of balances 105 stored in the security token 102) to determine how many security tokens 102 will be burned. In examples, the at least one processor may iteratively step through all addresses in the list of addresses and, for each address, determine the fractional (non-integer) tokens in the resulting post-split token balance associated with the address. The fractions of tokens identified across all addresses in

the list of addresses may be summed together to determine the number of security tokens 102 that will be paid out/burned (cachedOutTokenQuantity).

[0089] As a simple example of the audit in optional step 304, suppose that there are three token holders: A, B, and C of a particular token 102. If the split ratio of a new split is 1/2, A's address possessing 100 pre-split tokens 102 will have 50 split-updated tokens (with no fractional tokens), B's address possessing 75 pre-split tokens 102 will have 37 1/2 post-split tokens (i.e., 1/2 fractional tokens), and C's address possessing 55 pre-split tokens 102 will have 27 1/2 post-split tokens (i.e., 1/2 fractional tokens). The amount of all fractional tokens to be paid out/burned (cachedOutTokenQuantity) following the new split will be 1 token (1/2 from B and 1/2 from C). In examples, the split-updated token balance for each address may be determined by calling the balance check method 144, passing a different address in each invocation of the balance check method 144.

[0090] The method 300 may proceed at optional step 306 where the at least one processor initiates at least one transaction to compensate each security token holder, which will possess a fractional security token following the new split, for their respective fractional security token 102. The transaction may be a cryptocurrency transaction that is recorded on the distributed ledger 122. For example, a cryptocurrency transaction may be initiated for each fraction token holder, where the fractional token holder's address is the output address of a transaction, and where the amount of cryptocurrency transferred has a value equal (or very nearly equal to) their fractional token being paid out/burned in response to the new split. The issuer 114 may be the payor in these transaction(s), so the issuer's address may be the input address in the transaction(s). In some configurations, more than one of the (e.g., all) fractional token holders may be compensated for their respective fractional token in a single transaction, e.g., the issuer's address is the input address, and all the addresses of the fractional token holders are the output addresses.

[0091] In some configurations, optional step 306 may include the at least one processor updating the token balances (e.g., in the table of balances 105) for the address(es) that will have fractional tokens paid out/burned following the new split. This may include calling the update balances method 146 (updateBalances()). For example, all addresses with a fractional token balance of less than one (following the new split) may have their balances updated. Alternatively, all addresses with a

fractional token balance following the new split may have their balances updated. Alternatively, none of the addresses have their token balances updated following the new split (unless a subsequent transfer is initiated). In some configurations, paying out/burning tokens may include invoking the transfer() method 109.

[0092] The method 300 may proceed at optional step 308 where the at least one processor updates the total number of outstanding security tokens 102 following the new split. The post-split total number of outstanding security tokens may be determined as the pre-split total number of outstanding security tokens minus the number tokens paid out/burned in response to the new split (cachedOutTokenQuantity). Accordingly, optional step 308 may include the at least one processor calling an update total tokens (splitTotalSupply()) method 148, which takes cachedOutTokenQuantity as a parameter.

[0093] The method 300 may proceed at optional step 310 where the at least one processor removes each security token holder, which will possess less than one security token 102 following the new split, from a list of addresses holding at least one of the security tokens 102, e.g., in a table of balances 105 stored in the security token 102. In other words, the address(es) of token holder(s) that will no longer possess any tokens (after the fractional tokens are paid out in response to the new split) will no longer be considered token holders (of the specified type of security token 102) following the new split.

[0094] Optional step 310 may include the at least one processor calling a remove shareholder function 150, e.g., implemented in the security token 102. A single address (of a token holder possessing less than one security token 102 following the new split) may be passed as a parameter to each invocation of the remove shareholder function 150. Alternatively, multiple addresses of at least one token holder may be passed to a single invocation of the remove shareholder function 150.

[0095] The method may proceed at step 312 where an indication of the split ratio is stored without updating the token balances for all addresses in the list of addresses holding at least one of the security tokens (e.g., the table of balances 105). The stored split ratio may be used to determine split-updated token balances, e.g., as part of a subsequent token transfer described in Figure 4. Performing the method 300 without actually changing the token balances for all addresses in the list of addresses holding

at least one of the security tokens (e.g., the table of balances 105 in the security token 102) may reduce blockchain-related fees (gas) incurred when a token is split.

[0096] It should be noted that the method 300 may be performed each time a split is initiated on the security token 102. A history of splits may be stored in a single data structure (e.g., in an array called splits[] having an object for each split, each object having a numerator and a denominator of a particular split).

[0097] Figure 4 is a flow diagram illustrating a method 400 for transferring at least one security token 102 following a token split. The method 400 may be executed by at least one processor running a virtual machine in the system 100, e.g., at least one processor in a network node 140. The method 400 may include executing one or more split-related methods 152 implemented in a security token 102, e.g., stored in a distributed ledger 122. Additionally, the method 400 may utilize one or more split-related parameters 154 in a security token 102, e.g., stored in a distributed ledger 122.

[0098] The method 400 of Figure 4 may be performed after one or more iterations of the method 300 of Figure 3. In examples, if a particular token holder holds a quantity of tokens for a long period of time (across multiple splits of the security token 102), the method 400 may have to account for multiple splits (e.g., according to the method 300 of Figure 3) as part of the transfer of the at least one security token 102.

Alternatively, if the particular token holders involved in the transfer have not yet acquired tokens or only recently acquired tokens (and no intervening splits have been declared for the security token 102), the method 400 may not have to account for any splits (e.g., according to the method 300 of Figure 3).

[0099] As noted above, the token balances in the list of addresses (e.g., in the table of balances 105) may not actually be updated when a token is split (e.g., according to the method 300 of Figure 3). Instead, an indicator of the split ratio may be stored, which can be accounted for during future executions of the check balance method 144 and/or the transfer() method 109. In contrast, the method 400 of Figure 4 may actually update the list of addresses (e.g., in the table of balances 105).

[0100] The method 400 may begin at step 402 where the at least one processor determines that a new transfer of a quantity of the security tokens 102, from a first address to a second address, has been requested. In examples, step 402 may include a transfer() method 109 being called. The transfer() method 109 may be implemented in

the security token 102 and may be called by a wallet (e.g., implemented on an investor 120 smartphone or in a browser using Metamask) or any other computing device (e.g., a hosted Ethereum provider, such as Infura) or via direct private network node 140 access. In examples, the following parameters may be passed as part of the transfer() method 109: the second address (of the transferee token holder) as the output address, and the quantity (of security tokens being transferred). In examples, the at least one processor uses the address invoking the transaction (the first address of the transferring/transferor token holder, in this example) as the input address from which the at least one security token 102 will be transferred, i.e., so the first address does not get sent as an explicit parameter.

[0101] The method 400 may proceed at optional step 404 where the at least one processor adds the second address to the list of addresses holding at least one of the security tokens 102 (e.g., in a table of balances 105 in the security token 102). If the second address (received as a parameter in the transfer() method 109 call) is already in the list of addresses, optional step 404 may be skipped. If the second address is not already in the list of addresses, it may be added with an initial token balance of 0.

[0102] The method 400 may proceed at step 406 where the at least one processor determines split-updated token balances (newBalance) for the first address and the second address based on at least one stored indication of at least one split ratio. Step 406 may include the transfer() method 109 calling a update balances method 146 (updateBalances()), which calls a balance check method 144 (balanceOf()), all being implemented in the security token 102. In some configurations, the balance check method 144 may be called twice: once for the first address and once for the second address.

[0103] During the balance check method 144 for the first address, the at least one processor may iterate through all the splits for which the token balance associated with the first address (of the token transferor) has not been adjusted, but should be. For the first unaccounted-for split, the at least one processor may determine a new token balance as the respective pre-split token balance multiplied by a stored indication of a first split ratio, e.g., $\text{newBalance} = \text{pre-split balance} * 1\text{st split ratio}$. If multiple splits need to be accounted for, the resulting new token balance is then multiplied by the stored indications of all remaining unaccounted-for splits, until all splits have been accounted for, e.g., $\text{newBalance} = \text{pre-split balance} * 1\text{st split ratio} *$

2nd split ratio If all splits have been accounted for (including for new token holders purchasing their first quantity of the security token; or if no splits have been declared for the security token 102 since the first address has held tokens), no adjustments are necessary for the first address and the second address, i.e., $\text{newBalance} = \text{current token balance}$.

[0104] A similar balance check method 144 may be performed for the second address (to determine the newBalance for the second address). In some configurations, the token balance associated with the first address will be adjusted to account for a different number of splits than the token balance associated with the second address. In other words, the newBalance for the first address may be determined by multiplying the current token balance by a different number of split ratios than the newBalance for the second address. For example, if a first token holder (owning the first address) has held their security tokens 102 for ten years, during which time the token has been split two times, while the second token holder (owning the second address) has held their security tokens 102 for three years, during which time the token has been split once.

[0105] The method 400 may proceed at step 408 where the at least one processor updates the token balances for the first address and the second address based on the split-updated balances (e.g., newBalance values) for the first address and the second address. Step 408 may be performed as part of the update balances method 146 (updateBalances()), e.g., with the first address and the second address as parameters. In some configurations, the update balances method 146 may be called by the transfer() method 109 (called in step 402). In examples, the update balances method 146 may update the actual token balances in the list of addresses (e.g., the table of balances 105 in the security token 102). The updated balance for the first address may be determined as the newBalance for the first address (from step 406) minus the quantity of security tokens transferred. The updated balance for the second address may be determined as the newBalance for the second address (from step 406) plus the quantity of security tokens transferred. Step 408 may include committing the transaction of the at least one security token to a blockchain, e.g., an Ethereum blockchain.

[0106] The method 400 may proceed at optional step 410 where the at least one processor stores an indication that all splits have been accounted for in the token

balances for the first address and the second address. This may prevent the token balances for the first address and/or second address being multiplied by the same split ratio more than once, e.g., in step 406 of future transfers involving the first address and/or second address. In examples, an index (within a data structure of all split ratios for the security token 102) for each of the first address and second address may be updated. However, other tracking mechanisms are possible to ensure that every split ratio is accurately accounted for, if appropriate.

[0107] Example Implementations

[0108] Figures 5A-B are block diagrams illustrating example tables of balances 105 and example multi-object data structure that can be used to implement the token splitting in the method 300 of Figure 3. Figures 6A-B are block diagrams illustrating example tables of balances 105 and example multi-object data structure that can be used to implement the token transfer in the method 400 of Figure 4. A variety of data structures (e.g., split-related parameters 154) and methods (split-related methods 152) may be implemented in the example implementations described below. However, it is understood that the data structures and methods are merely exemplary and should not be viewed as limiting. The methods in the example implementations may be executed by a processor on a network node 140, e.g., that is running a virtual machine, such as the Ethereum Virtual Machine (EVM).

[0109] A multi-object data structure called “splits” is maintained in the example implementations, where the object at each index represents a split ratio of a different split of a particular token 102. In other words, the splits data structure is a mapping of integer indices to split ratios (of splits having occurred in the particular token’s history). For example, each object may be of type struct (in Solidity) and may have a numerator element and a denominator element, where numerator/denominator is the split ratio of a particular split. There is only one splits data structure for a particular token that records all splits (forward and reverse through the token’s history). The splits data structure for a particular token may be created at the time of creation of the issuing smart contract 113 for the particular token or at issuance of the particular token 102. In examples, the splits data structure is implemented as an array and stored in the security token 102. The splits data structure may be considered a split-related parameter 154 stored on the distributed ledger 122.

[0110] A multi-object data structure called “balances” is also maintained in the example implementations. This balances data structure is a mapping where each index is a token holder’s address (e.g., “addr”) and it points to the token balance in the token holder’s address. In other words, the balances data structure may implement the table of balances 105 (e.g., Table 1 above). For example, balances[addr] will return the token balance associated with the address addr. In examples, the balances data structure may be implemented as an array and implemented in the security token 102. The balances data structure may be considered a split-related parameter 154 stored on the distributed ledger 122.

[0111] A multi-object data structure called “holdersSplit” is also maintained in the example implementations. The holdersSplit data structure is a mapping where the index is a token holder's address and it points to the index in the splits data structure that has mostly recently been applied to the token balance associated with the token holder's address. For example, a particular shareholder’s address (e.g., addr) may be used to find another index (in the splits data structure) of the last split in the token’s history that was applied to the token balance for the token holder (e.g., balances[addr]). If a token holder address doesn't exist in holdersSplit, it means that the address has never invoked a transfer to have this updated, and a value of “0” would be returned. The holdersSplit data structure may be considered a split-related parameter 154 stored on the distributed ledger 122.

[0112] A single holdersSplit data structure is maintained for a particular token and would be updated after any transfer of the particular token 102. The holdersSplit data structure would have to be updated when (1) a token balance, associated with an address, is updated with a new split during a call of the transfer() method 109; and/or (2) a previous token holder is dispossessed of their last remaining quantity of tokens (because a fractional token 102 was paid out and/or the remaining quantity was transferred out of their address). However, if an address (possessing tokens) has never been updated with a new split, the address doesn't need to be added to holdersSplit. The default return value of holdersSplit for an address that doesn't exist in holdersSplit is 0, which indicates that all the possible splits in the token’s history need to be adjusted for in the token balance associated with the address.

[0113] The holdersSplit data structure for a particular token may be created at the time of creation of the issuing smart contract 113 for the particular token or at

issuance of the particular token 102. In examples, the holdersSplit data structure is implemented as an array and stored in the security token 102.

[0114] A value called “currentSplitIndex” is also maintained in the example implementations. The value currentSplitIndex is an index (to the splits data structure) of the most recent split/reverse split in the token’s history (without respect to any specific token holder). In other words, the value currentSplitIndex tracks the location of the current split in the splits data structure. In examples, the value currentSplitIndex may be implemented as an integer and implemented in the security token 102. The currentSplitIndex value may be considered a split-related parameter 154 stored on the distributed ledger 122.

[0115] The update split method 142 (updateSplit()) is called as part of step 302 in the method 300 of Figure 3. In the example implementations, the at least one processor executing the update split method 142 (updateSplit()): (1) adds a new Split object to the splits data structure for the token 102, the Split object having a numerator element and a denominator element; and (2) increments currentSplitIndex. The update split method 142 (updateSplit()) may be called by the issuer 114 in response to the issuer 114 determining to split tokens 102 issued by a smart contract owned by the issuer 114. The split ratio of the desired split may be passed as a parameter to the update split method 142 (updateSplit()). The update split method 142 (updateSplit()) may be a split-related method 152 implemented in the security token 102, and may be executed on a network node 140 implementing the distributed ledger 122.

[0116] The balance check method 144 (balanceOf()) is called as part of optional step 304 in the method 300 of Figure 3 and/or step 406 in the method 400 of Figure 4. Alternatively, the balance check method 144 (balanceOf()) can be called simply to check a split-updated token balance associated with an address (separate from the method 300 of Figure 3 and the method 400 of Figure 4). An address of a token holder (e.g., addr) is passed as a parameter to the balance check method 144, e.g., balanceOf(addr). The balance check method 144 (balanceOf(addr)) may be a “view” function, so all blockchain-related fees (gas) may be refunded at the end, making it free to invoke.

[0117] During execution of the balance check method 144 (balanceOf(addr)), the at least one processor may iterate through all the Split objects for which the particular

address's token balance has not yet been updated, multiplying the token balance (`balances[addr]`) by the split ratio at each iteration, i.e., the at least one processor may iterate from index `holdersSplit[addr]+1` to `currentSplitIndex` in `splits[]`. The balance check method 144 (`balanceOf()`) returns the split-updated balance (`newBalance`) of the address. In other words, `newBalance` is determined as `balances[addr]` multiplied by each split ratio that should be accounted for in the address's token balance, but isn't.

[0118] It should be noted that only the split ratios starting at index `holdersSplit[addr]+1` (and going to `currentSplitIndex`) in the `splits` data structure are used to determine `newBalance` during execution of the balance check method 144 (`balanceOf()`). Accordingly, the split ratio of any split that was declared before the token holder became a token holder is/are not used in the `newBalance` calculation in the balance check method 144 (`balanceOf()`).

[0119] The following excerpt of code illustrates the example implementation of the balance check method 144 (`balanceOf()`):

```
function balanceOf(address addr)
public
view
returns (uint256) {
    unit256 newBalance = balances[addr];
    for(uint i = holdersSplit[addr] + 1; i <= currentSplitIndex; i++) {
        newBalance =
            (newBalance.multiply(splits[i].numerator)).divide(splits[i].denominator);
    }
    return newBalance;
}
```

[0120] As can be deduced from the excerpt of code (and preceding description), the current balance for the token holder (`balances[addr]`) is obtained and set as `newBalance`. Each split (in the `splits[]` data structure) that has occurred but has not yet been applied to the token holder's balance is iterated through, where `newBalance` is multiplied by the split numerator and divided by the split denominator at each such

split. Once all the split multipliers have been applied, the resulting newBalance (reflecting the current split-updated balance) is returned from balanceOf(addr).

[0121] Accordingly, balanceOf(addr) may update the current balance (balances[addr]) for the address addr by one or more splits that have been declared for the token held in the address addr, but have not yet been applied to the current balance.

Alternatively, if balances[addr] is called for an address that holds no tokens (of the particular type), then balances[addr] will be “0” and the newBalance value returned by balanceOf(addr) will be “0”. The balance check method 142 (balanceOf()) may be a split-related method 152 implemented in the security token 102, and may be executed on a network node 140 implementing the distributed ledger 122.

[0122] Figure 5A is a block diagram illustrating a table of balances 105 and a splits data structure before a new token split is declared, e.g., using an example implementation of the method of Figure 3. Specifically, the left side of Figure 5A includes a table of balances 105 with three different addresses: A with a balance of 150 tokens 102, B with a balance of 25 tokens 102, and C with a balance of 75 tokens 102. Addresses A, B, and C may belong to the same or different token holders as long as they hold the same type of tokens 102, e.g., tokens representing shares of ABC Corporation. Each of addresses A, B, and C may have received their tokens by security token offering (STO), purchase, gift, inheritance, compensation, employee stock/token purchase program, etc. It is understood that this is example is merely exemplary and there could be many addresses holding tokens of a particular type, each address holding an integer number of tokens 102 (of at least “1”). While the addresses and token balances are illustrated as respective columns in tables of balances 105 herein, it is understood that the addresses and/or the token balances may each be implemented in any suitable data structure, e.g., they may be implemented in an array called balances[] that is addressable by address.

[0123] The right side of Figure 5A illustrates a splits data structure for the token 102 held by addresses A, B, and C. In this example, the token 102 has undergone two different splits in its history, with split ratios of 2/1, 4/1. Assume that address A received it’s 150 token after the first split (with a 2/1 ratio), address B received it’s 25 tokens after the second split (with a 4/1 ratio), and address C received it’s 75 tokens after the second split (with a 4/1 ratio). Assume further that addresses A, B, and C have not transferred any of their respective tokens since they received them.

Accordingly, holdersSplit[A]=1 (meaning the token balance associated with address A has been updated for the first split (2/1), but not the second split (4/1)); holdersSplit[B]=2 (meaning the token balance associated with address B has been updated for the first split (2/1) and the second split (4/1)); and holdersSplit[C]=2 (meaning the token balance associated with address C has been updated for the first split (2/1) and the second split (4/1)). Accordingly, balanceOf(A) would return 600 (since the 4/1 split would be applied in newBalance that was returned); balanceOf(B) would return 25 (since both splits had already been accounted for in address B's balance); and balanceOf(C) would return 75 (since both splits had already been accounted for in address C's balance). The value of currentSplitIndex is 2 because that is the index of the most recent split declared for this token 102.

[0124] Figure 5B is a block diagram illustrating a table of balances 105 and a splits data structure after the new token split is declared, e.g., using an example implementation of the method of Figure 3. In examples, the tokens may have been previously issued by ABC Corporation. Assume that a new split is declared with a split ratio of 1/2 (considered a "reverse" split). To declare the split, ABC corporation (acting as the issuer 114) or the owner 112 would call the updateSplit(numerator,denominator) method 142 where numerator=1 and denominator=2 in this example. The at least one processor (e.g., in a network node 140) executing updateSplit(numerator,denominator) would then add a new Split object to the splits data structure, having a numerator element with a value of 1 and a denominator element with a value of 2 (denoted by the "1/2" in Figure 5B). The at least one processor executing updateSplit(numerator,denominator) would also update currentSplitIndex to 3 because that is the index of the most recent split declared for this token 102.

[0125] The at least one processor may also audit the table of balances 105 to determine if any tokens need to be paid out/burned. In other words, the at least one processor may determine if and how many tokens 102 will be eliminated from existence after the new split (1/2) is applied and resulting fractional tokens are paid out. This may include calling the balanceOf() method 144 for each token holder (balanceOf(A), balanceOf(B), and balanceOf(C)), which would return newBalance values of 600, 12.5, and 37.5 for address A, B, and C, respectively. It should be noted

that, for forward splits with an integer split ratio (e.g., 2/1, 3/1, 4/1, etc.), the audit is not performed because no fractional tokens 102 will result.

[0126] Since newBalance for addresses B and C is fractional, token holder(s) associated with addresses B and C will be paid out/compensated for their fractional tokens. Specifically, 0.5 of address B's tokens and 0.5 of address C's tokens 102 will be paid out, i.e., $\text{cachedOutTokenQuantity} = 0 \text{ (for address A)} + 0.5 \text{ (for address B)} + 0.5 \text{ (for address C)} = 1$ token that will be paid out. Accordingly, the at least one processor may initiate at least one transaction to compensate the token holder(s) of addresses B and C for their respective half tokens 102. In some configurations, two different transactions will be executed, one transferring the market value of 0.5 tokens to the token holder associated with address B and another transferring the market value of 0.5 tokens to the token holder associated with address C. Alternatively, the token holder(s) associated with both addresses B and C may be compensated in a single transaction. The compensating transaction(s) may include the transfer of cryptocurrency (e.g., Ether) and may be recorded on a distributed ledger 122 (e.g., an Ethereum blockchain).

[0127] Optionally, the at least one processor may update the token balances (e.g., in the table of balances 105) for the address(es) that will have fractional tokens paid out/burned following the new split. In the example above, this may include calling the update balances method 146 (updateBalances()) for address B and address C to update each token balance to 0. In some configurations, paying out/burning tokens may include invoking the transfer() method 109.

[0128] The at least one processor may also update the total number of outstanding security tokens 102 following the new split, i.e., to reduce the total number of outstanding security tokens 102 by the number tokens paid out/burned in response to the new split (cachedOutTokenQuantity). This may include calling the splitTotalSupply(cachedOutTokenQuantity) method 148.

[0129] The at least one processor may also check if any token holders need to be removed from the table of balances 105 following the new split (of 1/2). In this example, each address still holds at least one token following the new split, so no token holders/addresses need to be removed from the table of balances 105. However, if address Z only included 0.5 tokens following the new split, which were then paid

out, address Z would then be removed from the table of balances, e.g., with a `removeShareholder(Z)` method 150 call.

[0130] This example implementation of declaring a token split is performed without actually changing the token balances for all addresses in the table of balances 105 in the security token 102. In other words, this example implementation of token splitting does not call the update balances method 146. Accordingly, this example implementation of token splitting reduces blockchain-related fees (gas) that would otherwise be incurred if the token balances (in the table of balances 105) were updated. For example, `holdersSplit[]` is not updated in response to the splitting of a token, and `balanceOf()` is free to call. Only when tokens are subsequently transferred (e.g., as described in Figures 6A-B below) will the token balances in the table of balances actually be updated, thus incurring blockchain-related fees (and `holdersSplit[]` be updated).

[0131] Figure 6A is a block diagram illustrating a table of balances 105 and a splits data structure before a token transaction is performed, e.g., according to an example implementation of the method of Figure 4. Specifically, Figure 6A illustrates the table of balances 105 and the splits data structure resulting from Figure 5B after a new split (with a split ratio of 1/2) is declared for the token 102.

[0132] Assume that the token holder associated with address A then initiates a transaction of 50 tokens to address B. This may include the token holder associated with address A calling the `transfer()` method 109. In example, this may include calling `transfer(address B, 50)` from the token holder's wallet or other computing device. At least one processor in a network node 140 may execute the `transfer(address B, 50)` call. The at least one processor may use the address invoking the transaction (address A) as the transferor address. If necessary (i.e., if address B wasn't already a token holder), the at least one processor would call an `add shareholder` (e.g., `addShareholder()`) method to add address B to the table of balances 105.

[0133] The at least one processor may then call the update balances method 146 (e.g., `updateBalances(A, B)`) as part of executing the `transfer(address B, 50)` call. The at least one processor may call `balanceOf(A)` and `balanceOf(B)` as part of executing the `updateBalances(A, B)` call, which will return `newBalance` values of 300 ($150 \cdot 4 \cdot 1/2$) and 12 ($25 \cdot 1/2 - 0.5$ paid out already), respectively.

[0134] Using the returned newBalance values for addresses A (300) and B (12), the at least one processor may then update the token balances for addresses A and B in the table of balances, e.g., which is recorded on the distributed ledger 122. Specifically, the at least one processor may update the payor/transferor's balance as the newBalance – the quantity transferred, i.e., the at least one processor updates the token balance for address A to $300 - 50 = 250$. The at least one processor may update the payee/transferee's balance as the newBalance + the quantity transferred, i.e., the at least one processor updates the token balance for address B to $12 + 50 = 62$. In some configurations, the token balances may be updated (e.g., using updateBalances()) as part of the transfer method 109.

[0135] Figure 6B is a block diagram illustrating a table of balances 105 and a splits data structure after the token transaction is performed, e.g., according to the example implementation of the method of Figure 4. Following the token transaction (of 50 tokens from address A to address B), the table of balances 105 is updated. Specifically, the balance for address A will be updated to $300 - 50$ tokens that were transferred = 250. The balance for address B will also be updated to $12 + 50$ tokens received = 62. Updating the balances may include committing a transaction to the distributed ledger 122. The balance for address C is not changed because address C was not involved in the token transaction.

[0136] Additionally, the holdersSplit data structure is updated for addresses A and B following the token transaction, e.g., holdersSplit[A] = holdersSplit[B] = 3. The holdersSplit data structure is not updated for address C because address C was not involved in the token transaction.

Computer System Overview

[0137] Embodiments of the present disclosure include various steps and operations, which have been described above. A variety of these steps and operations may be performed by hardware components or may be embodied in machine-executable instructions, which may be used to cause a general-purpose or special-purpose processor programmed with the instructions to perform the steps. Alternatively, the steps may be performed by a combination of hardware, software, and/or firmware. As such, Figure 7 is a block diagram illustrating an example computer system 700 with which embodiments of the present disclosure may be utilized. According to the

present example, the computer system 700 includes an interconnect 702, at least one processor 704, at least one communication port 706, at least one main memory 708, at least one removable storage media 710, at least one read only memory 712, and at least one mass storage device 714.

[0138] The at least one processor 704 can be any known processor. The at least one communication port 706 can be or include, In examples, any of an RS-232 port for use with a modem-based dialup connection, a 10/100 Ethernet port, or a Gigabit port using copper or fiber. The nature of the at least one communication port 706 may be chosen depending on a network such as a Local Area Network (LAN), Wide Area Network (WAN), or any network to which the computer system 700 connects. The at least one main memory 708 can be Random Access Memory (RAM), or any other dynamic storage device(s) commonly known in the art. The at least one read only memory 712 can be any static storage device(s) such as Programmable Read Only Memory (PROM) chips for storing static information such as instructions for the at least one processor 704.

[0139] The at least one mass storage device 714 can be used to store information and instructions. In examples, hard disks (such as magnetic disk drives or solid state drive using serial/parallel ATA or SCSI interfaces), an optical disc, an array of disks such as a Redundant Array of Independent Disks (RAID), or any other mass storage devices may be used. Interconnect 702 can be or include one or more buses, bridges, controllers, adapters, and/or point-to-point connections. Interconnect 702 communicatively couples the at least one processor 704 with the other memory, storage, and communication blocks. Interconnect 702 can be a PCI/PCI-X or SCSI based system bus depending on the storage devices used. The at least one removable storage media 710 can be any kind of external hard-drives, floppy drives, Compact Disc–Read Only Memory (CD-ROM), Compact Disc–Re-Writable (CD-RW), Digital Video Disc–Read Only Memory (DVD-ROM), Blu-Ray Disc Read Only Memory (BD-ROM), Blu-Ray Disc Recordable (BD-R), Blu-Ray Disc Recordable Erasable (BD-RE).

[0140] The components described above are meant to exemplify some types of possibilities. In no way should the aforementioned examples limit the disclosure, as they are only exemplary embodiments.

[0141] Figure 8 is a block diagram illustrating another example computing device 800 with which embodiments of the present disclosure may be utilized. The example computing device 800 may be used to implement any of the global registry 108, owner 112, issuer 114, custodian 116, broker dealer 118, investor 120, one of the network nodes 140 storing a copy of the distributed ledger 122, security token exchange 124, and/or a virtual machine (e.g., Ethereum Virtual Machine) executing any of the smart contracts described herein. The computing device 800 includes at least one memory 802, at least one processor 804, optional at least one network interface 806, optional display device 808, optional input device 810, and optional power source 812.

[0142] In examples, the at least one memory 802 can be any device, mechanism, or populated data structure used for storing information. In examples, the at least one memory 802 can be or include any type of volatile memory, nonvolatile memory, and/or dynamic memory. In examples, the at least one memory 802 can be random access memory, memory storage devices, optical memory devices, magnetic media, floppy disks, magnetic tapes, hard drives, erasable programmable read-only memories (EPROMs), electrically erasable programmable read-only memories (EEPROMs), optical media (such as compact discs, DVDs, Blu-ray Discs) and/or the like.

[0143] In accordance with some embodiments, the at least one memory 802 may include one or more disk drives, flash drives, one or more databases, one or more tables, one or more files, local cache memories, processor cache memories, relational databases, flat databases, and/or the like. In addition, those of ordinary skill in the art will appreciate many additional devices and techniques for storing information, which can be used as the at least one memory 802. The at least one memory 802 may be used to store instructions for running one or more applications or modules on the at least one processor 804. In examples, the at least one memory 802 could be used in one or more examples to house all or some of the instructions needed to execute the functionality discussed herein, e.g., in Figures 3-4.

[0144] The at least one processor 804 can be any known processor, such as a general purpose processor (GPP) or special purpose (such as a field-programmable gate array (FPGA), application-specific integrated circuit (ASIC) or other integrated circuit or circuitry), or any programmable logic device. In examples, any of the functionality

disclosed herein (e.g., in Figures 3-4) may be implemented by the at least one processor 804 and the at least one memory 802.

[0145] In examples, the at least one optional network interface 806 includes or is coupled to at least one optional antenna for communication with a network (such as one of the at least one networks 112 of system 100). In examples, the at least one optional network interface 806 includes at least one of an Ethernet interface, a cellular radio access technology (RAT) radio, a Wi-Fi radio, a Bluetooth radio, or a near field communication (NFC) radio. In examples, the at least one optional network interface 806 includes a cellular radio access technology radio configured to establish a cellular data connection (mobile Internet) of sufficient speeds with a remote server using a local area network (LAN) or a wide area network (WAN). In examples, the cellular radio access technology includes at least one of Personal Communication Services (PCS), Specialized Mobile Radio (SMR) services, Enhanced Special Mobile Radio (ESMR) services, Advanced Wireless Services (AWS), Code Division Multiple Access (CDMA), Global System for Mobile Communications (GSM) services, Wideband Code Division Multiple Access (W-CDMA), Universal Mobile Telecommunications System (UMTS), Worldwide Interoperability for Microwave Access (WiMAX), 3rd Generation Partnership Projects (3GPP) Long Term Evolution (LTE), High Speed Packet Access (HSPA), third generation (3G) fourth generation (4G), fifth generation (5G), etc. or other appropriate communication services or a combination thereof. In examples, the at least one optional network interface 806 includes a Wi-Fi (IEEE 802.11) radio configured to communicate with a wireless local area network that communicates with the remote server, rather than a wide area network. In examples, the at least one optional network interface 806 includes a near field radio communication device that is limited to close proximity communication, such as a passive near field communication (NFC) tag, an active near field communication (NFC) tag, a passive radio frequency identification (RFID) tag, an active radio frequency identification (RFID) tag, a proximity card, or other personal area network device.

[0146] In examples, the optional at least one display device 808 includes at least one of a light emitting diode (LED), a liquid crystal display (LCD), a light emitting diode (LED) display, an organic light emitting diode (OLED) display, an e-ink display, a field emission display (FED), a surface-conduction electron-emitter display (SED), or

a plasma display. In examples, the optional at least one input device 810 includes at least one of a touchscreen (including capacitive and resistive touchscreens), a touchpad, a capacitive button, a mechanical button, a switch, a dial, a keyboard, a mouse, a camera, a biometric sensor/scanner, a microphone, etc. In examples, the optional at least one display device 808 is combined with the optional at least one input device 810 into a human machine interface (HMI) for user interaction with the computing device 800. In examples, at least one optional power source 812 is used to provide power to the various components of the computing device 800.

Terminology

[0147] Brief definitions of terms, abbreviations, and phrases used throughout this application are given below.

[0148] The term “determining” may include calculating, computing, generating, processing, deriving, investigating, looking up (e.g., looking up in a table, a database or another data structure), ascertaining and the like. Also, “determining” may also include receiving (e.g., receiving information), accessing (e.g., accessing data in a memory) and the like. Also, “determining” may include resolving, selecting, choosing, establishing and the like.

[0149] The phrase “based on” does not mean “based only on,” unless expressly specified otherwise. In other words, the phrase “based on” describes both “based only on” and “based at least on”. Additionally, the phrase “based on” does not preclude intermediate steps, e.g., A is based on C may mean that B is based on C and A is based on B. Additionally, the term “and/or” means “and” or “or”. In examples, “A and/or B” can mean “A”, “B”, or “A and B”. Additionally, “A, B, and/or C” can mean “A alone,” “B alone,” “C alone,” “A and B,” “A and C,” “B and C” or “A, B, and C.”

[0150] The terms “connected”, “coupled”, and “communicatively coupled” and related terms are used in an operational sense and are not necessarily limited to a direct physical connection or coupling. Thus, In examples, two devices may be coupled directly, or via one or more intermediary media or devices. As another example, devices may be coupled in such a way that information can be passed there between, while not sharing any physical connection with one another. Based on the disclosure provided herein, one of ordinary skill in the art will appreciate a variety of

ways in which connection or coupling exists in accordance with the aforementioned definition.

[0151] The phrases “in exemplary embodiments”, “in example embodiments”, “in some embodiments”, “according to some embodiments”, “in the embodiments shown”, “in other embodiments”, “embodiments”, “in examples”, “examples”, “in some examples”, “some examples” and the like generally mean the particular feature, structure, or characteristic following the phrase is included in at least one embodiment of the present disclosure, and may be included in more than one embodiment of the present disclosure. In addition, such phrases do not necessarily refer to the same embodiments or different embodiments.

[0152] If the specification states a component or feature “may,” “can,” “could,” or “might” be included or have a characteristic, that particular component or feature is not required to be included or have the characteristic.

[0153] The term “responsive” includes completely or partially responsive.

[0154] The term “module” refers broadly to a software, hardware, or firmware (or any combination thereof) component. Modules are typically functional components that can generate useful data or other output using specified input(s). A module may or may not be self-contained. An application program (also called an “application”) may include one or more modules, or a module can include one or more application programs.

[0155] The term “network” generally refers to a group of interconnected devices capable of exchanging information. A network may be as few as several personal computers on a Local Area Network (LAN) or as large as the Internet, a worldwide network of computers. As used herein, “network” is intended to encompass any network capable of transmitting information from one entity to another. In some cases, a network may be comprised of multiple networks, even multiple heterogeneous networks, such as one or more border networks, voice networks, broadband networks, financial networks, service provider networks, Internet Service Provider (ISP) networks, and/or Public Switched Telephone Networks (PSTNs), interconnected via gateways operable to facilitate communications between and among the various networks.

[0156] Also, for the sake of illustration, various embodiments of the present disclosure have herein been described in the context of computer programs, physical components, and logical interactions within modern computer networks. Importantly, while these embodiments describe various embodiments of the present disclosure in relation to modern computer networks and programs, the method and apparatus described herein are equally applicable to other systems, devices, and networks as one skilled in the art will appreciate. As such, the illustrated applications of the embodiments of the present disclosure are not meant to be limiting, but instead are examples. Other systems, devices, and networks to which embodiments of the present disclosure are applicable include, in examples, other types of communication and computer devices and systems. More specifically, embodiments are applicable to communication systems, services, and devices such as cell phone networks and compatible devices. In addition, embodiments are applicable to all levels of computing from the personal computer to large network mainframes and servers.

[0157] In conclusion, the present disclosure provides novel systems, methods, and arrangements for splittable security tokens. While detailed descriptions of one or more embodiments of the disclosure have been given above, various alternatives, modifications, and equivalents will be apparent to those skilled in the art without varying from the spirit of the disclosure. In examples, while the embodiments described above refer to particular features, the scope of this disclosure also includes embodiments having different combinations of features and embodiments that do not include all of the described features. Accordingly, the scope of the present disclosure is intended to embrace all such alternatives, modifications, and variations as fall within the scope of the claims, together with all equivalents thereof. Therefore, the above description should not be taken as limiting.

Example Embodiments

[0158] Example 1 includes a network node comprising: at least one processor; at least one memory communicatively coupled to the at least one processor; at least one network interface communicatively coupled to the at least one processor; wherein the network node is configured to be within a plurality of network nodes communicatively coupled in a peer-to-peer network of network nodes implementing a distributed ledger; wherein the network node is configured to be communicatively coupled to at least one remotely located computing device through the at least one

network interface; wherein the at least one processor is configured to: determine a split ratio for a new split of security tokens held in at least one address of the distributed ledger; and store an indication of the split ratio without updating token balances for all of the at least one address in a list of addresses holding at least one of the security tokens.

[0159] Example 2 includes the network node of Example 1, wherein the at least one processor is further configured to: determine that a new transfer of a quantity of the security tokens, from a first address to a second address, has been requested; determine a split-updated token balance for each of the first address and the second address based on at least the stored indication of the split ratio; and update token balances for the first address and the second address based on the split-updated balances for the first address and the second address.

[0160] Example 3 includes the network node of any of Examples 1-2, wherein the at least one processor is further configured to: store an indication that all splits have been accounted for in the token balances for the first address and the second address.

[0161] Example 4 includes the network node of any of Examples 1-3, wherein the at least one processor is further configured to: determine a number of the security tokens that will be paid out in response to the new split; and initiate at least one transaction to compensate each security token holder associated with an address, which will possess a fractional security token following the new split, for their respective fractional security token.

[0162] Example 5 includes the network node of Example 4, wherein the at least one processor is further configured to: update the total number of outstanding security tokens following the new split; and remove each of the at least one address, which will possess less than one security token following the new split, from the list of addresses holding at least one of the security tokens.

[0163] Example 6 includes the network node of any of Examples 1-5, wherein the security tokens are cryptographic tokens that each represents an external, tradeable asset.

[0164] Example 7 includes the network node of any of Examples 1-6, wherein the security tokens are implemented using an originating smart contract.

[0165] Example 8 includes the network node of any of Examples 2-7, wherein determining the split-updated token balance for each of the first address and the second address incurs no blockchain-related fees.

[0166] Example 9 includes the network node of any of Examples 2-8, wherein updating token balances for the first address and the second address incurs block-chain related fees.

[0167] Example 10 includes the network node of any of Examples 1-9, wherein the at least one processor is further configured to store split ratios for all splits in the history of the security tokens in a multi-object data structure.

[0168] Example 11 includes the network node of Example 10, wherein the at least one processor is further configured to maintain an address-specific index to the multi-object data structure, for each address holding at least one of the security tokens, wherein each address-specific index indicates a split ratio that was most recently applied to a token balance associated with the respective address.

[0169] Example 12 includes the network node of any of Examples 10-11, wherein the at least one processor is further configured to maintain a current index to the multi-object data structure indicating the most recent split for the security tokens.

[0170] Example 13 includes a computerized method for transferring a self-enforcing security token, wherein the computerized method is performed by a network node configured to be included within a plurality of network nodes in a peer-to-peer network of network nodes implementing a distributed ledger, wherein the network node is configured to be communicatively coupled to other network nodes in the peer-to-peer network, the computerized method comprising: determining a split ratio for a new split of security tokens held in at least one address of the distributed ledger; and storing an indication of the split ratio without updating token balances for all of the at least one address in a list of addresses holding at least one of the security tokens.

[0171] Example 14 includes the method of Example 13, further comprising: determining that a new transfer of a quantity of the security tokens, from a first address to a second address, has been requested; determining a split-updated token balance for each of the first address and the second address based on at least the stored indication of the split ratio; and updating token balances for the first address and the second address based on the split-updated balances for the first address and the second address.

[0172] Example 15 includes the method of any of Examples 13-14, further comprising: storing an indication that all splits have been accounted for in the token balances for the first address and the second address.

[0173] Example 16 includes the method of any of Examples 13-15, further comprising: determining a number of the security tokens that will be paid out in response to the new split; and initiating at least one transaction to compensate each security token holder associated with an address, which will possess a fractional security token following the new split, for their respective fractional security token.

[0174] Example 17 includes the method of Example 16, further comprising: updating the total number of outstanding security tokens following the new split; and removing each of the at least one address, which will possess less than one security token following the new split, from the list of addresses holding at least one of the security tokens.

[0175] Example 18 includes the method of any of Examples 13-17, wherein the security tokens are cryptographic tokens that each represents an external, tradeable asset.

[0176] Example 19 includes the method of any of Examples 13-18, wherein the security tokens are implemented using an originating smart contract.

[0177] Example 20 includes the method of any of Examples 14-19, wherein determining the split-updated token balance for each of the first address and the second address incurs no blockchain-related fees.

[0178] Example 21 includes the method of any of Examples 14-20, wherein updating token balances for the first address and the second address incurs block-chain related fees.

[0179] Example 22 includes the method of any of Examples 13-21, wherein the at least one processor is further configured to store split ratios for all splits in the history of the security tokens in a multi-object data structure.

[0180] Example 23 includes the method of Example 22, further comprising maintaining an address-specific index to the multi-object data structure, for each address holding at least one of the security tokens, wherein each address-specific index indicates a split ratio that was most recently applied to a token balance associated with the respective address.

[0181] Example 24 includes the method of any of Examples 22-23, further comprising maintaining a current index to the multi-object data structure indicating the most recent split for the security tokens.

CLAIMS

What is claimed is:

1. A network node comprising:
 - at least one processor;
 - at least one memory communicatively coupled to the at least one processor;
 - at least one network interface communicatively couple to the at least one processor;wherein the network node is configured to be within a plurality of network nodes communicatively coupled in a peer-to-peer network of network nodes implementing a distributed ledger;
 - wherein the network node is configured to be communicatively coupled to at least one remotely located computing device through the at least one network interface;
 - wherein the at least one processor is configured to:
 - determine a split ratio for a new split of security tokens held in at least one address of the distributed ledger; and
 - store an indication of the split ratio without updating token balances for all of the at least one address in a list of addresses holding at least one of the security tokens.
2. The network node of claim 1, wherein the at least one processor is further configured to:
 - determine that a new transfer of a quantity of the security tokens, from a first address to a second address, has been requested;
 - determine a split-updated token balance for each of the first address and the second address based on at least the stored indication of the split ratio;
 - and
 - update token balances for the first address and the second address based on the split-updated balances for the first address and the second address.
3. The network node of claim 1, wherein the at least one processor is further configured to:

store an indication that all splits have been accounted for in the token balances for the first address and the second address.

4. The network node of claim 1, wherein the at least one processor is further configured to:

determine a number of the security tokens that will be paid out in response to the new split; and

initiate at least one transaction to compensate each security token holder associated with an address, which will possess a fractional security token following the new split, for their respective fractional security token.

5. The network node of claim 4, wherein the at least one processor is further configured to:

update the total number of outstanding security tokens following the new split; and

remove each of the at least one address, which will possess less than one security token following the new split, from the list of addresses holding at least one of the security tokens.

6. The network node of claim 1, wherein the security tokens are cryptographic tokens that each represents an external, tradeable asset.

7. The network node of claim 1, wherein the security tokens are implemented using an originating smart contract.

8. The network node of claim 2, wherein determining the split-updated token balance for each of the first address and the second address incurs no blockchain-related fees.

9. The network node of claim 2, wherein updating token balances for the first address and the second address incurs block-chain related fees.

10. The network node of claim 1, wherein the at least one processor is further configured to store split ratios for all splits in the history of the security tokens in a multi-object data structure.
11. The network node of claim 10, wherein the at least one processor is further configured to maintain an address-specific index to the multi-object data structure, for each address holding at least one of the security tokens, wherein each address-specific index indicates a split ratio that was most recently applied to a token balance associated with the respective address.
12. The network node of claim 10, wherein the at least one processor is further configured to maintain a current index to the multi-object data structure indicating the most recent split for the security tokens.
13. A computerized method for transferring a self-enforcing security token, wherein the computerized method is performed by a network node configured to be included within a plurality of network nodes in a peer-to-peer network of network nodes implementing a distributed ledger, wherein the network node is configured to be communicatively coupled to other network nodes in the peer-to-peer network, the computerized method comprising:
- determining a split ratio for a new split of security tokens held in at least one address of the distributed ledger; and
 - storing an indication of the split ratio without updating token balances for all of the at least one address in a list of addresses holding at least one of the security tokens.
14. The method of claim 13, further comprising:
- determining that a new transfer of a quantity of the security tokens, from a first address to a second address, has been requested;
 - determining a split-updated token balance for each of the first address and the second address based on at least the stored indication of the split ratio; and
 - updating token balances for the first address and the second address based on the split-updated balances for the first address and the second address.

15. The method of claim 13, further comprising:
storing an indication that all splits have been accounted for in the token balances for the first address and the second address.
16. The method of claim 13, further comprising:
determining a number of the security tokens that will be paid out in response to the new split; and
initiating at least one transaction to compensate each security token holder associated with an address, which will possess a fractional security token following the new split, for their respective fractional security token.
17. The method of claim 16, further comprising:
updating the total number of outstanding security tokens following the new split; and
removing each of the at least one address, which will possess less than one security token following the new split, from the list of addresses holding at least one of the security tokens.
18. The method of claim 13, wherein the security tokens are cryptographic tokens that each represents an external, tradeable asset.
19. The method of claim 13, wherein the security tokens are implemented using an originating smart contract.
20. The method of claim 14, wherein determining the split-updated token balance for each of the first address and the second address incurs no blockchain-related fees.
21. The method of claim 14, wherein updating token balances for the first address and the second address incurs block-chain related fees.

22. The method of claim 13, wherein the at least one processor is further configured to store split ratios for all splits in the history of the security tokens in a multi-object data structure.
23. The method of claim 22, further comprising maintaining an address-specific index to the multi-object data structure, for each address holding at least one of the security tokens, wherein each address-specific index indicates a split ratio that was most recently applied to a token balance associated with the respective address.
24. The method of claim 22, further comprising maintaining a current index to the multi-object data structure indicating the most recent split for the security tokens.