



US 20220058622A1

(19) **United States**

(12) **Patent Application Publication**

Snow

(10) **Pub. No.: US 2022/0058622 A1**

(43) **Pub. Date: Feb. 24, 2022**

(54) **PROTOCOLS IN BLOCKCHAIN ENVIRONMENTS**

(71) Applicant: **Inveniam Capital Partners, Inc.**, New York, NY (US)

(72) Inventor: **Paul Snow**, Austin, TX (US)

(73) Assignee: **Inveniam Capital Partners, Inc.**, New York, NY (US)

(21) Appl. No.: **17/451,655**

(22) Filed: **Oct. 21, 2021**

Related U.S. Application Data

(63) Continuation of application No. 16/351,597, filed on Mar. 13, 2019, now Pat. No. 11,205,172.

(60) Provisional application No. 62/714,909, filed on Aug. 6, 2018.

Publication Classification

(51) **Int. Cl.**
G06Q 20/36 (2006.01)
G06Q 20/06 (2006.01)
H04L 29/08 (2006.01)
H04L 9/06 (2006.01)

G06F 16/29 (2006.01)
G06F 21/53 (2006.01)
G06Q 20/38 (2006.01)
G06Q 20/12 (2006.01)
G06Q 20/40 (2006.01)
H04L 9/32 (2006.01)
G06F 21/64 (2006.01)

(52) **U.S. Cl.**
 CPC *G06Q 20/367* (2013.01); *H04L 9/3239* (2013.01); *H04L 67/12* (2013.01); *H04L 9/0637* (2013.01); *G06F 16/29* (2019.01); *G06Q 20/0658* (2013.01); *G06F 21/53* (2013.01); *G06Q 20/3829* (2013.01); *H04L 9/0643* (2013.01); *G06Q 20/3674* (2013.01); *G06Q 20/12* (2013.01); *G06Q 20/401* (2013.01); *H04L 9/0618* (2013.01); *H04L 9/3236* (2013.01); *G06F 21/645* (2013.01); *H04L 2209/38* (2013.01); *H04L 2209/56* (2013.01); *G06Q 20/065* (2013.01)

(57) **ABSTRACT**

A protocol cost effectively separates any blockchain (such as the Bitcoin blockchain) from any cryptocurrency (such as the Bitcoin cryptocurrency). The protocol provides client-defined Chains of Entries, client-side validation of Entries, a distributed consensus algorithm for recording the Entries, and a blockchain anchoring approach for security.

Factom Lets You Build Applications on the Bitcoin Blockchain

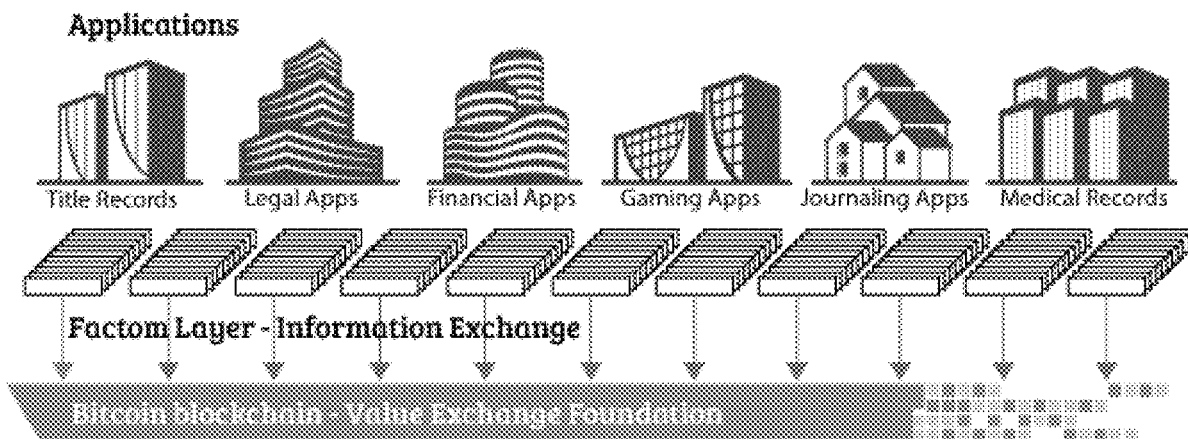


FIG. 1

Factom Lets You Build Applications on the Bitcoin Blockchain

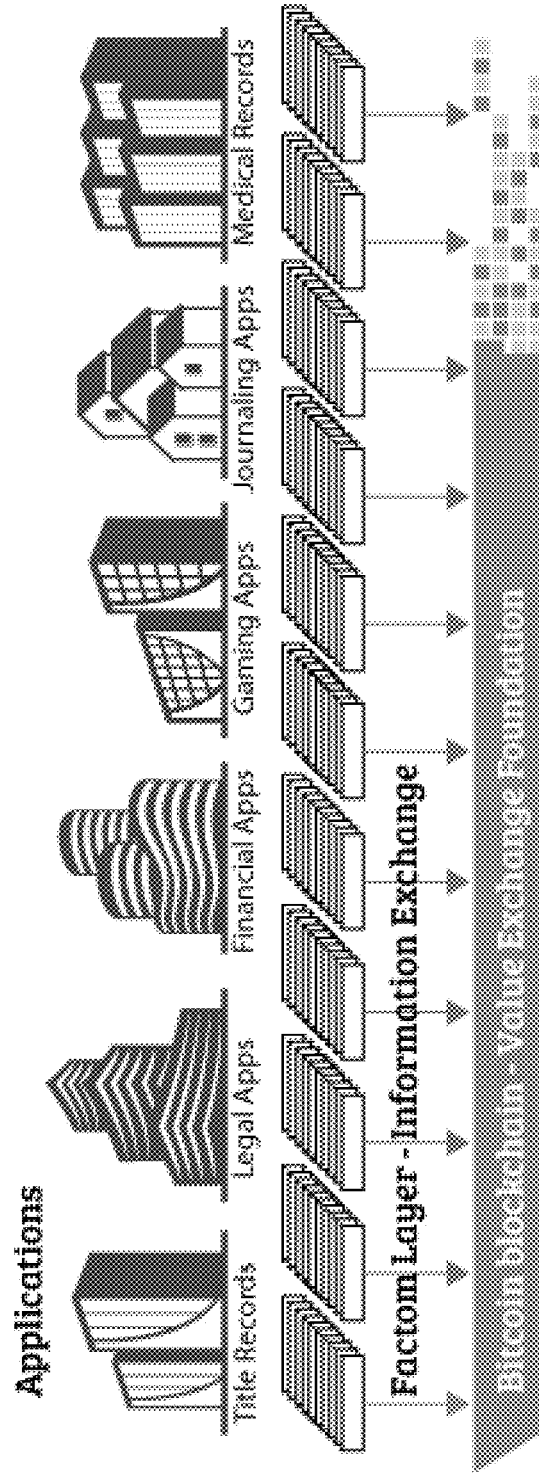


FIG. 2

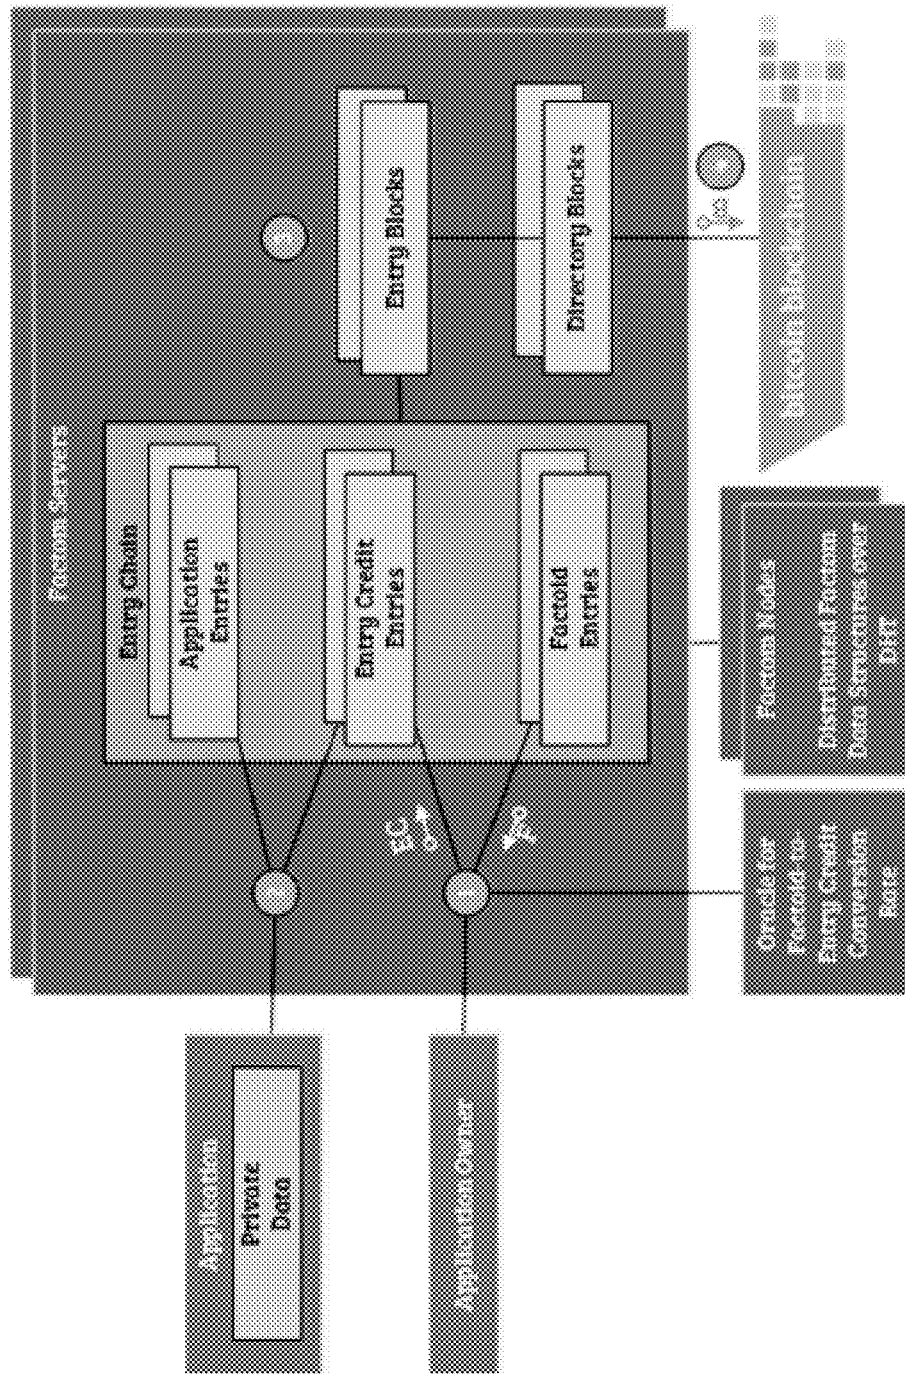


FIG. 3

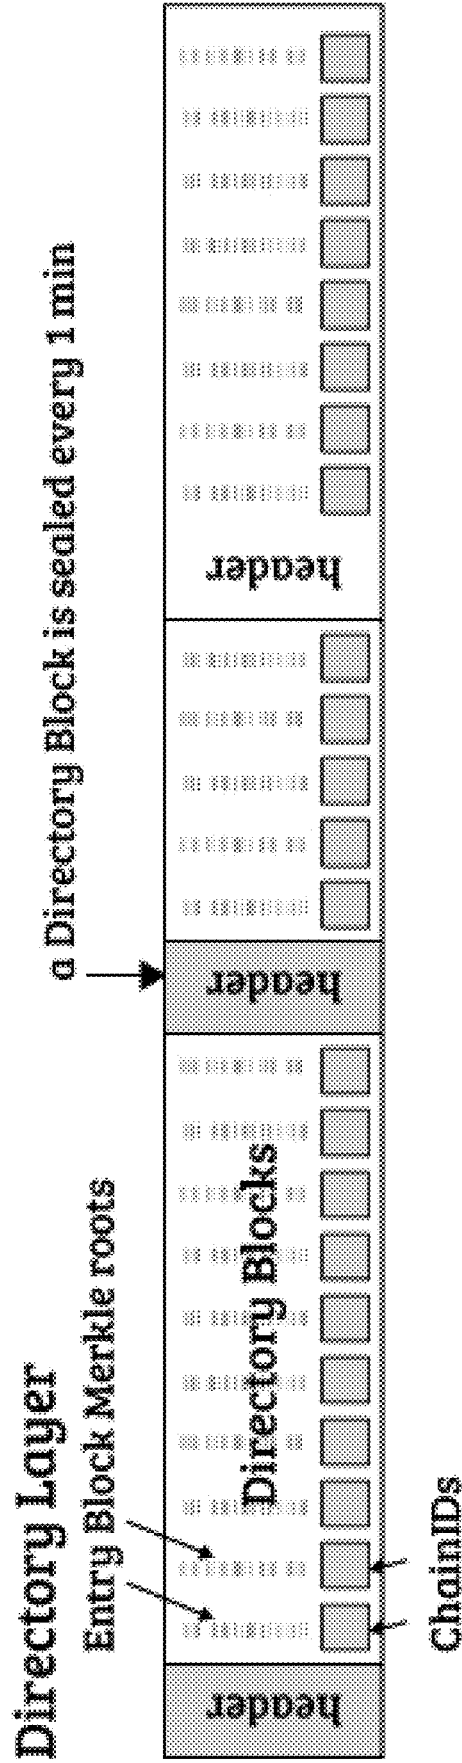


FIG. 4

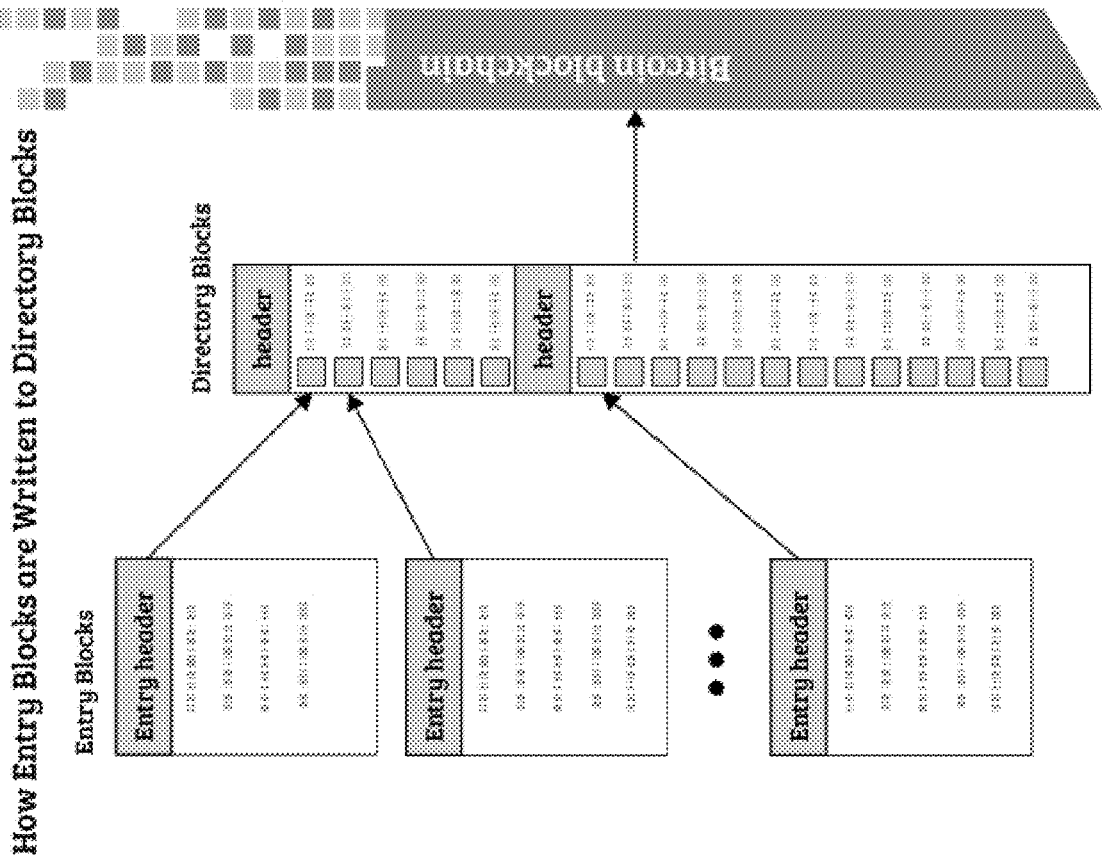
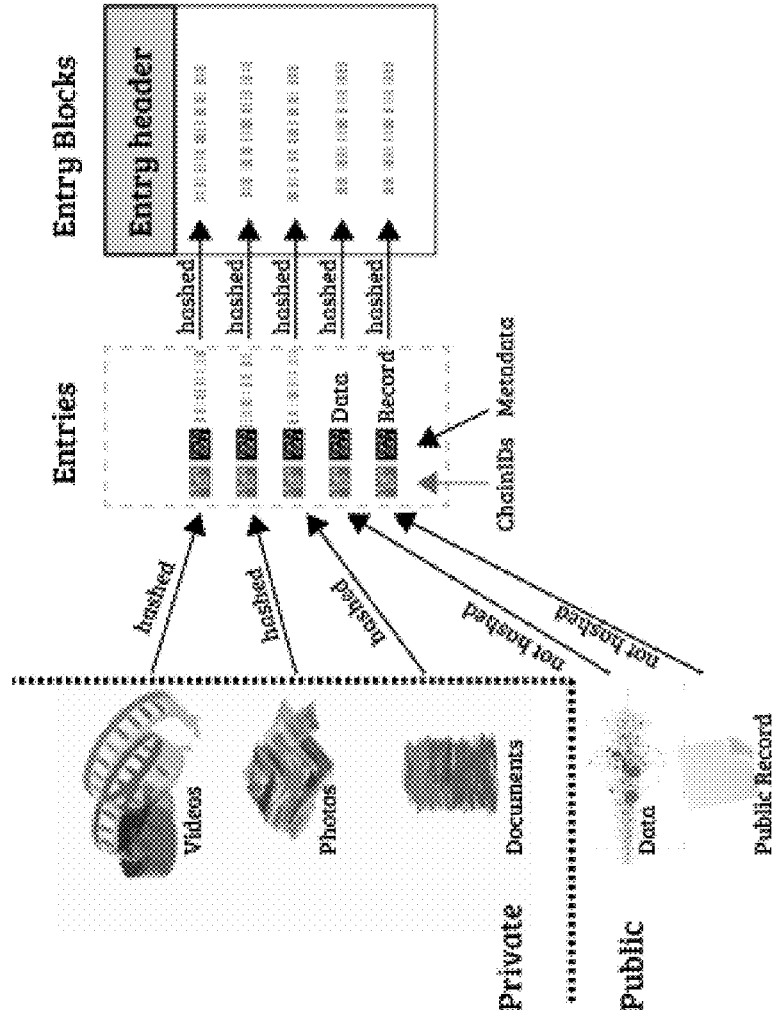


FIG. 5

How Hashes and Data are Written to Entry Blocks



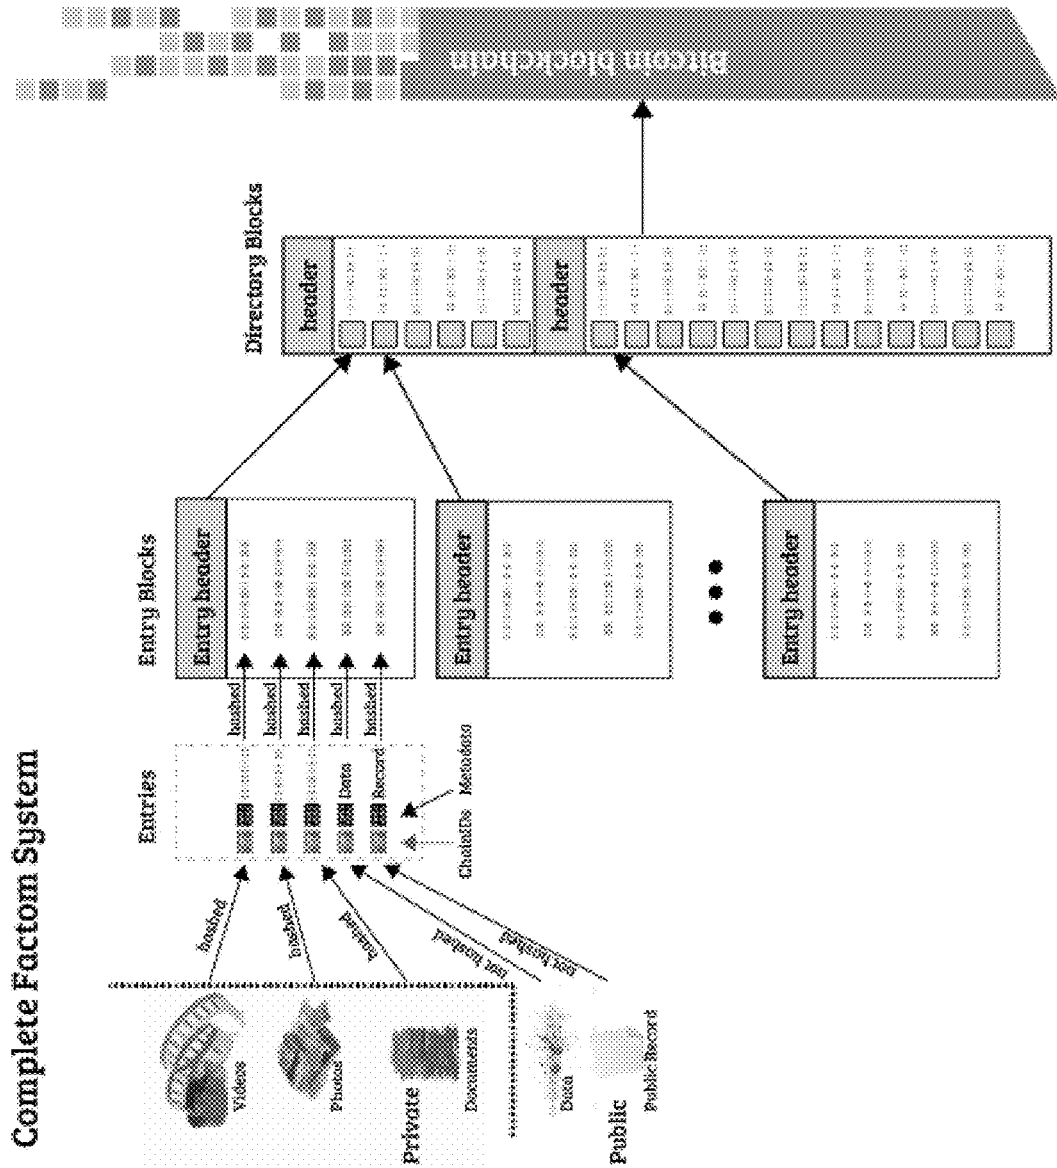


FIG. 6

Complete Factom System

FIG. 7

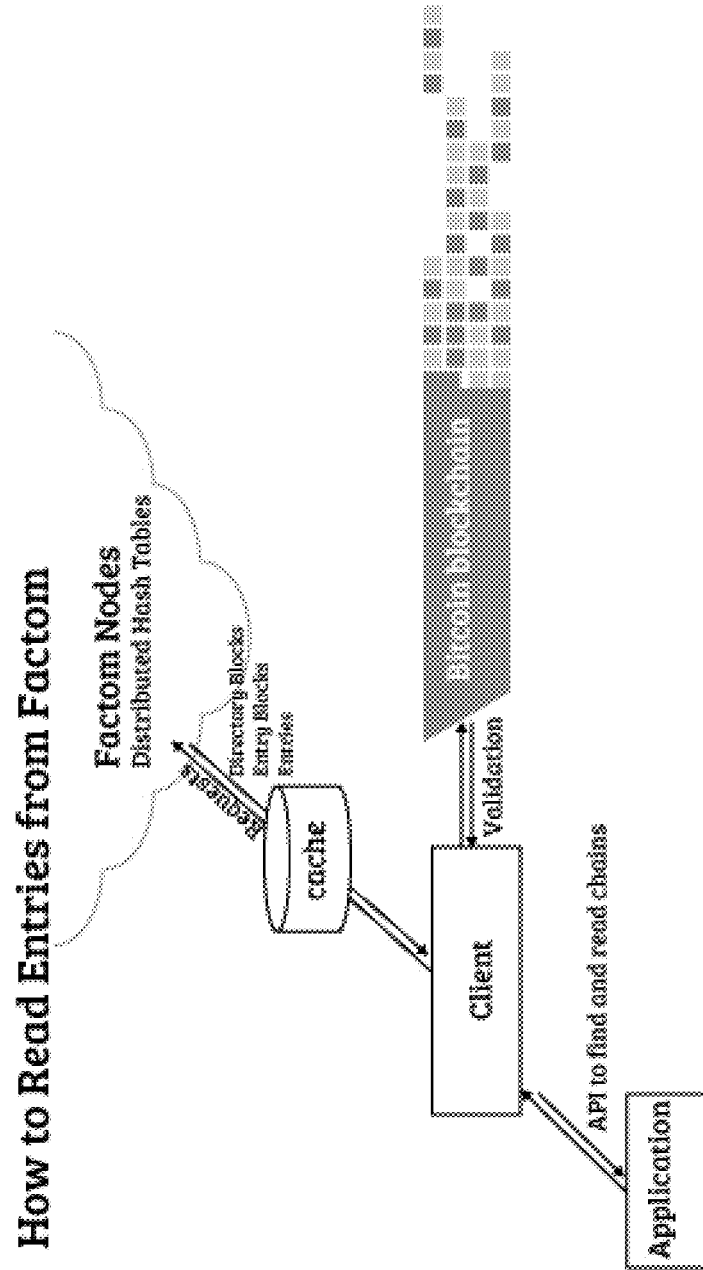


FIG. 8

Computing the ChainID

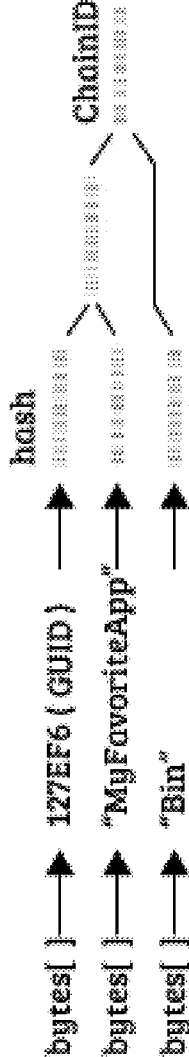


FIG. 9

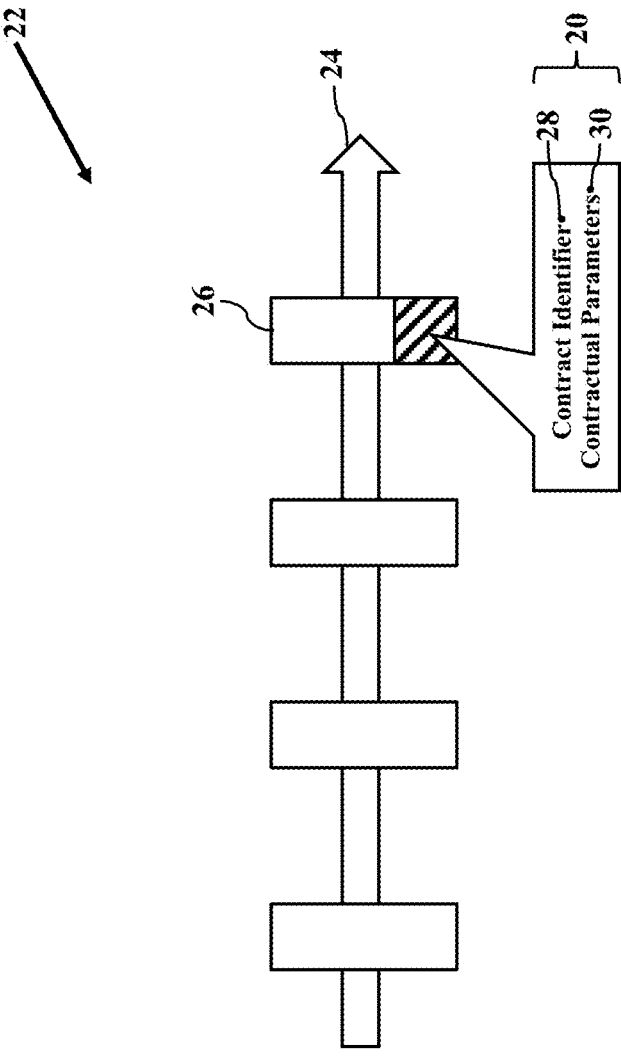


FIG. 10

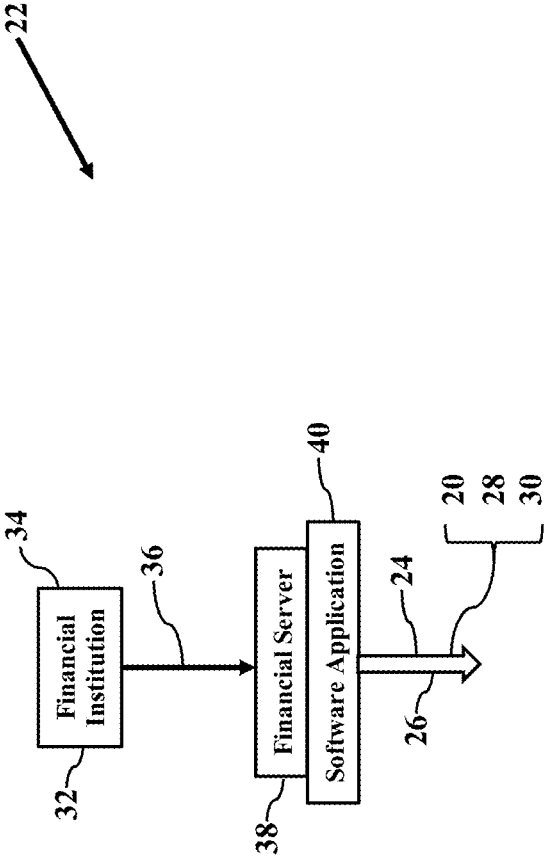


FIG. 11

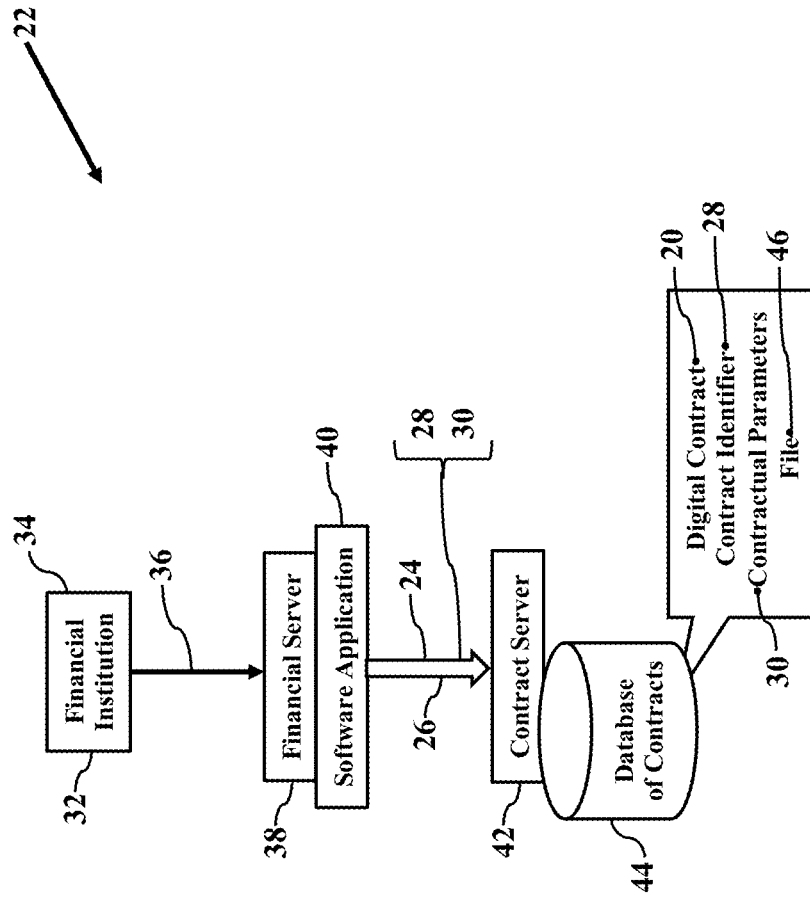


FIG. 12

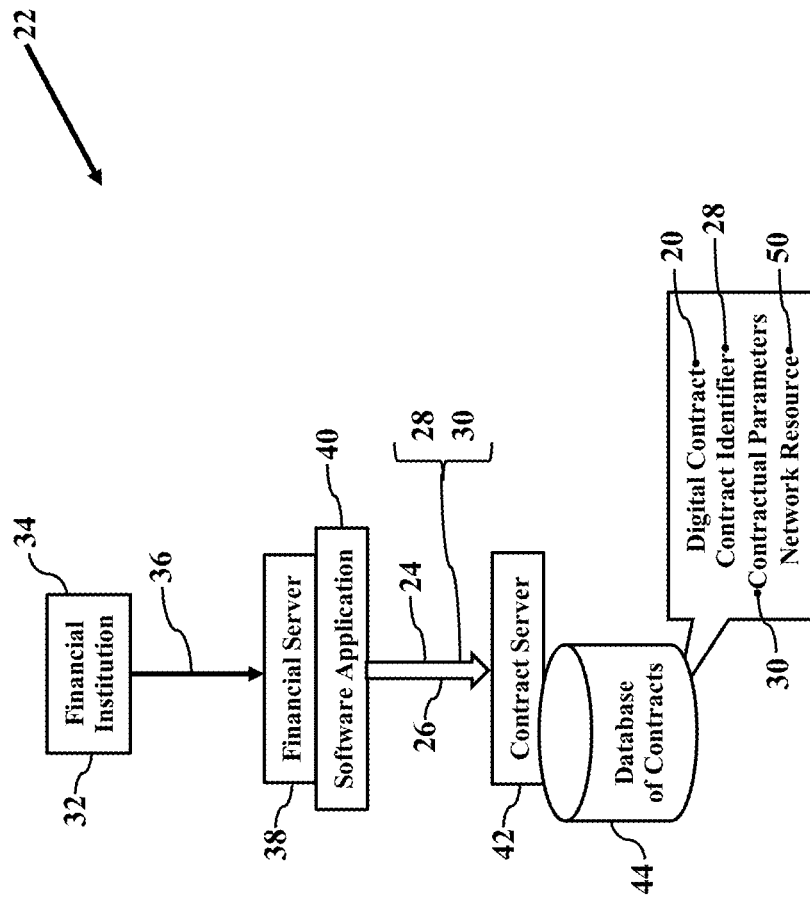


FIG. 13

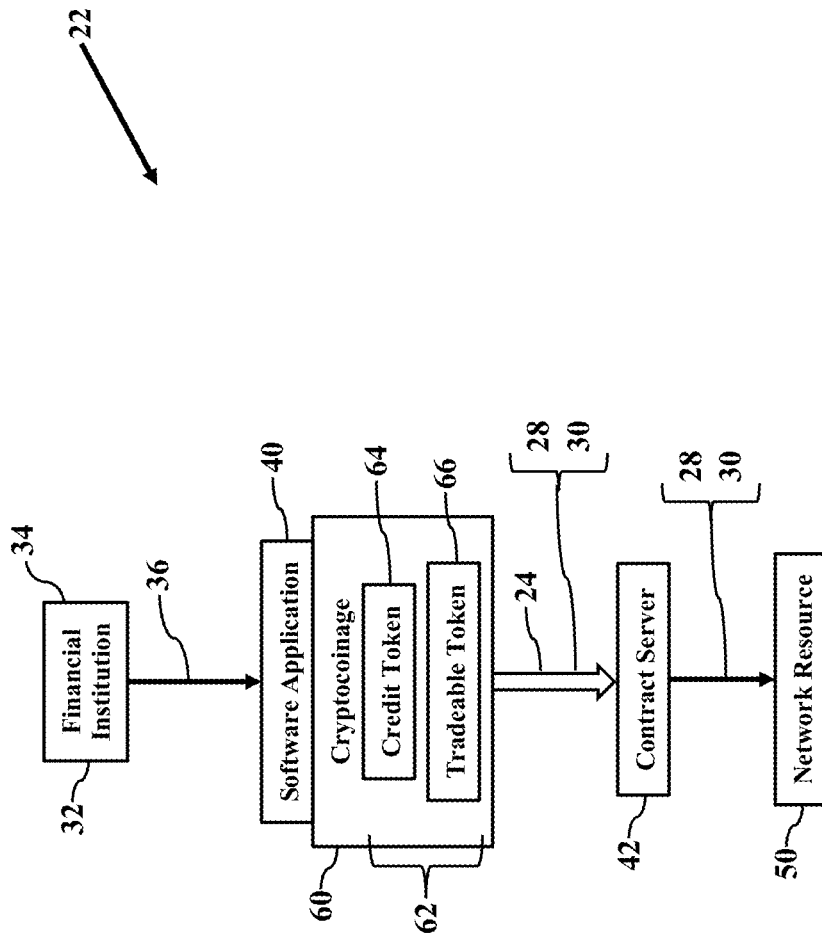


FIG. 14

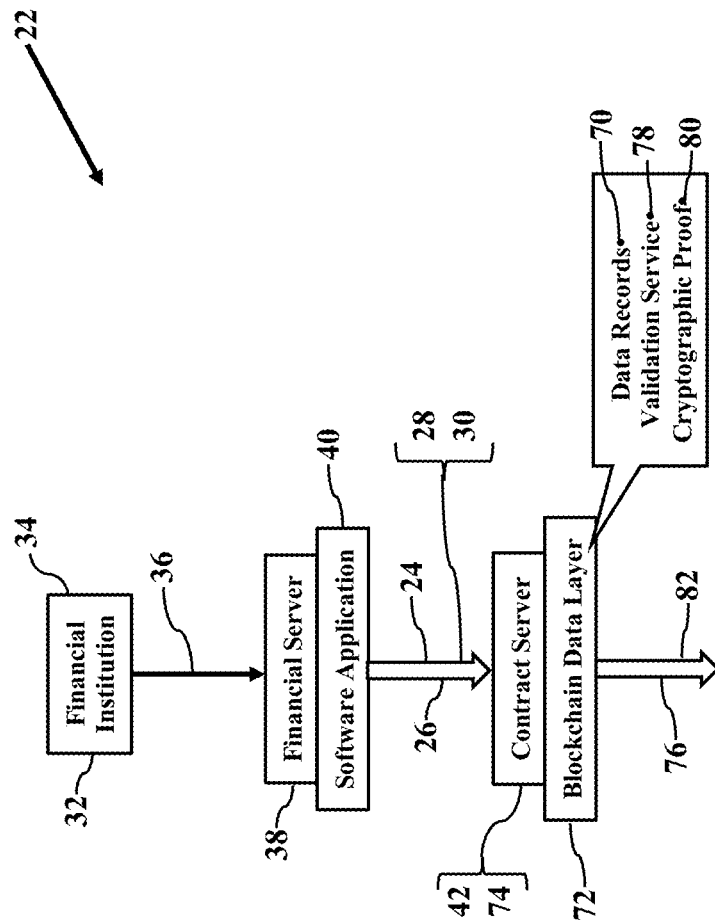


FIG. 15

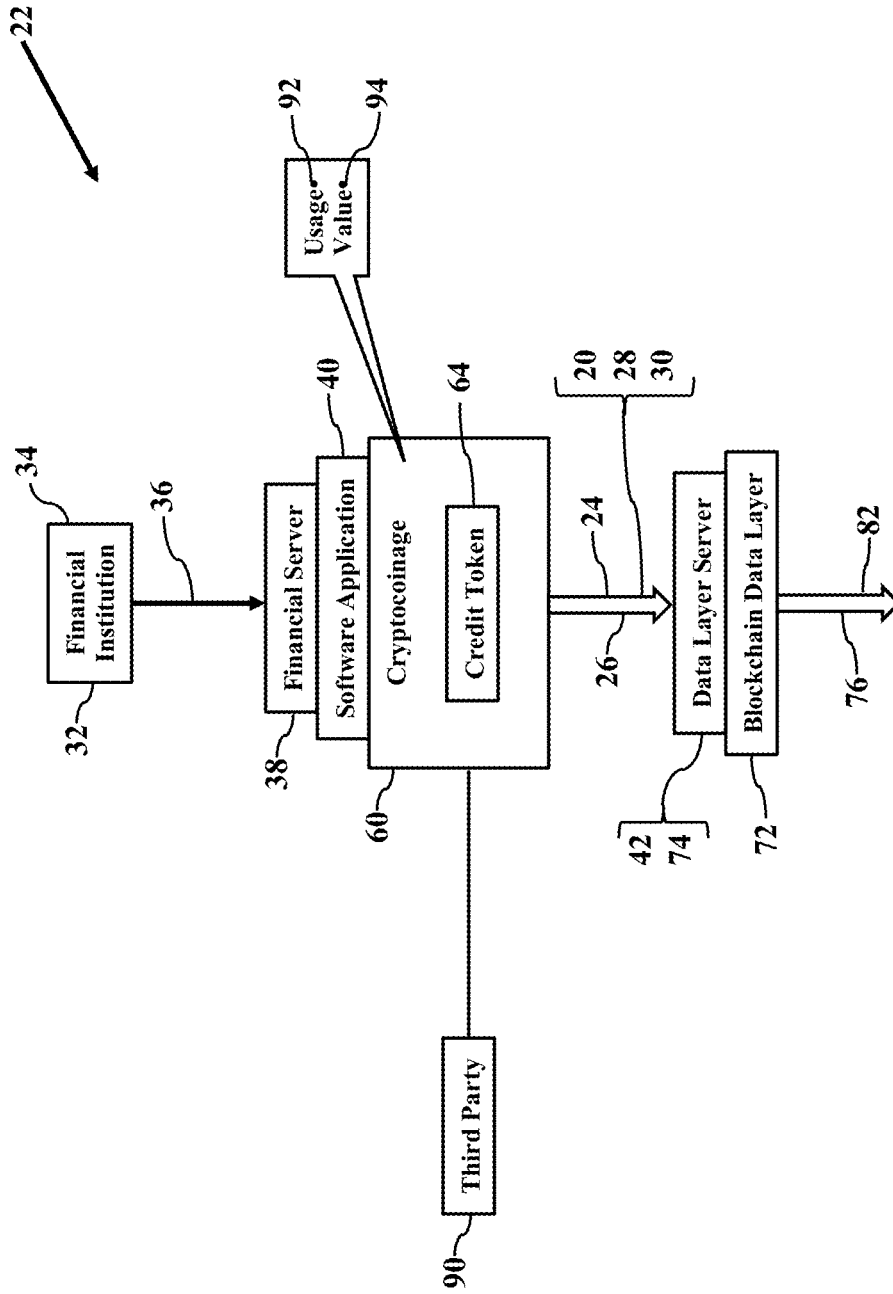


FIG. 16

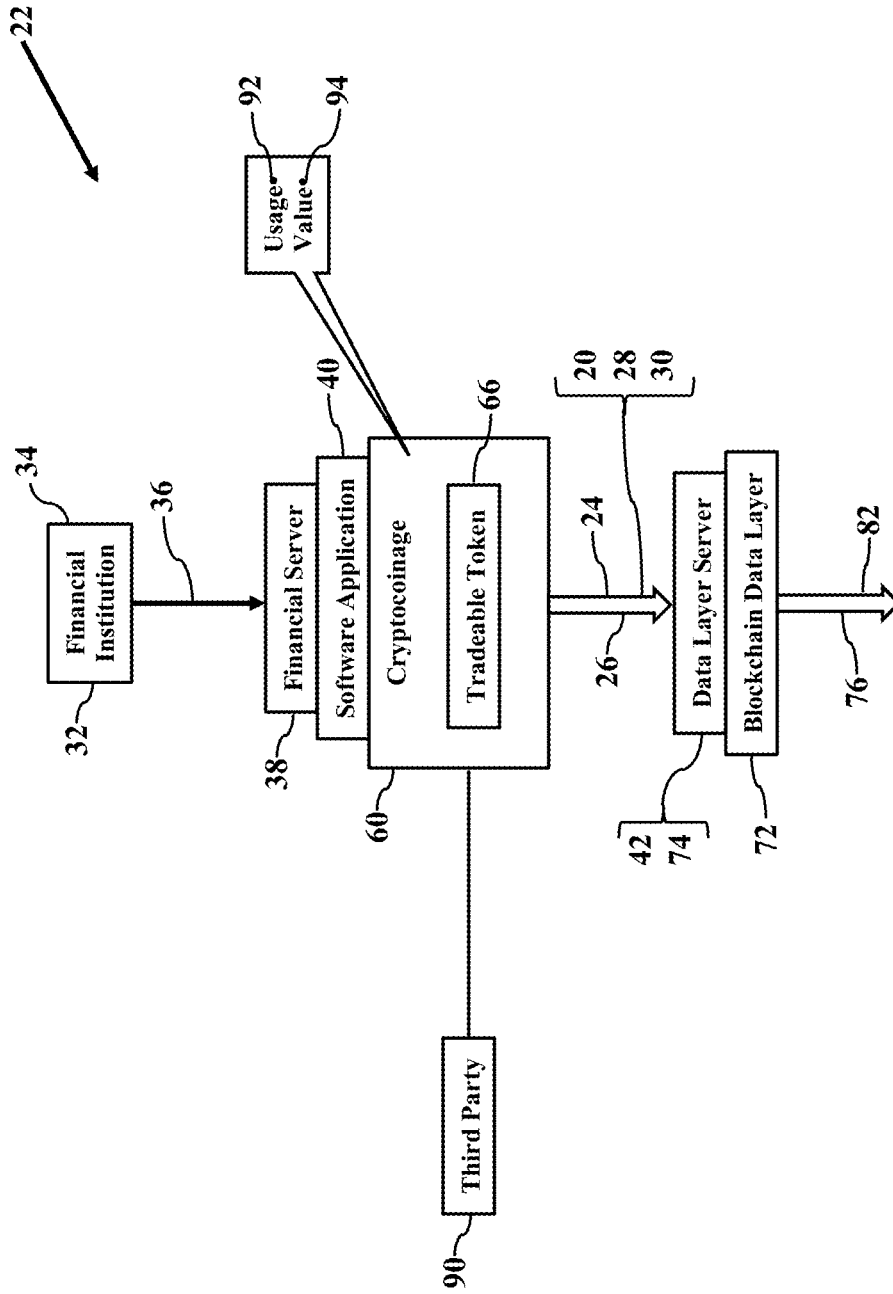


FIG. 17

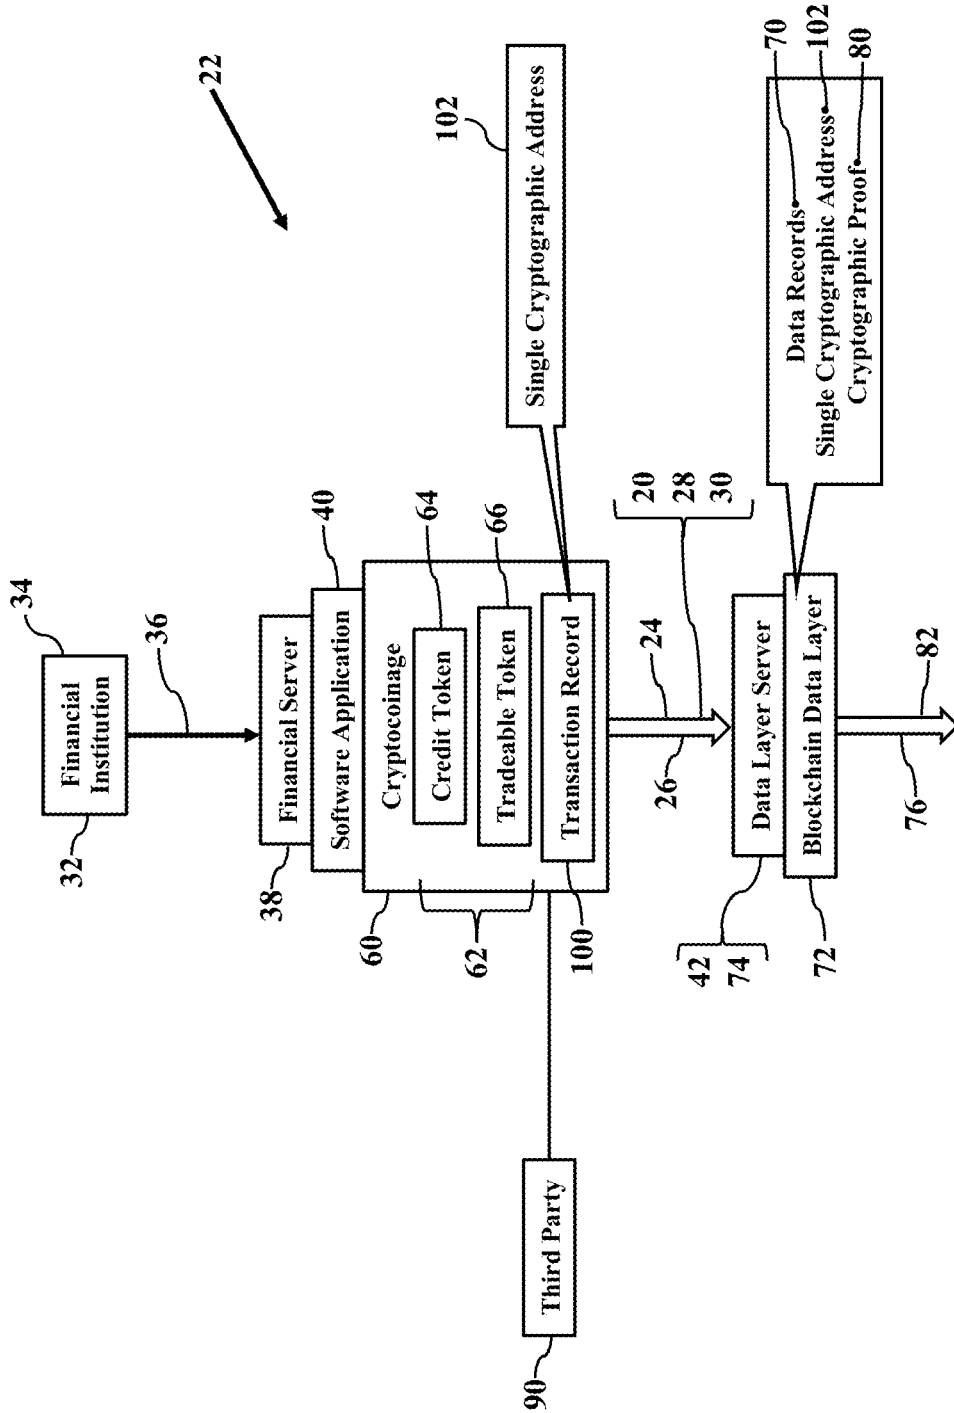


FIG. 18

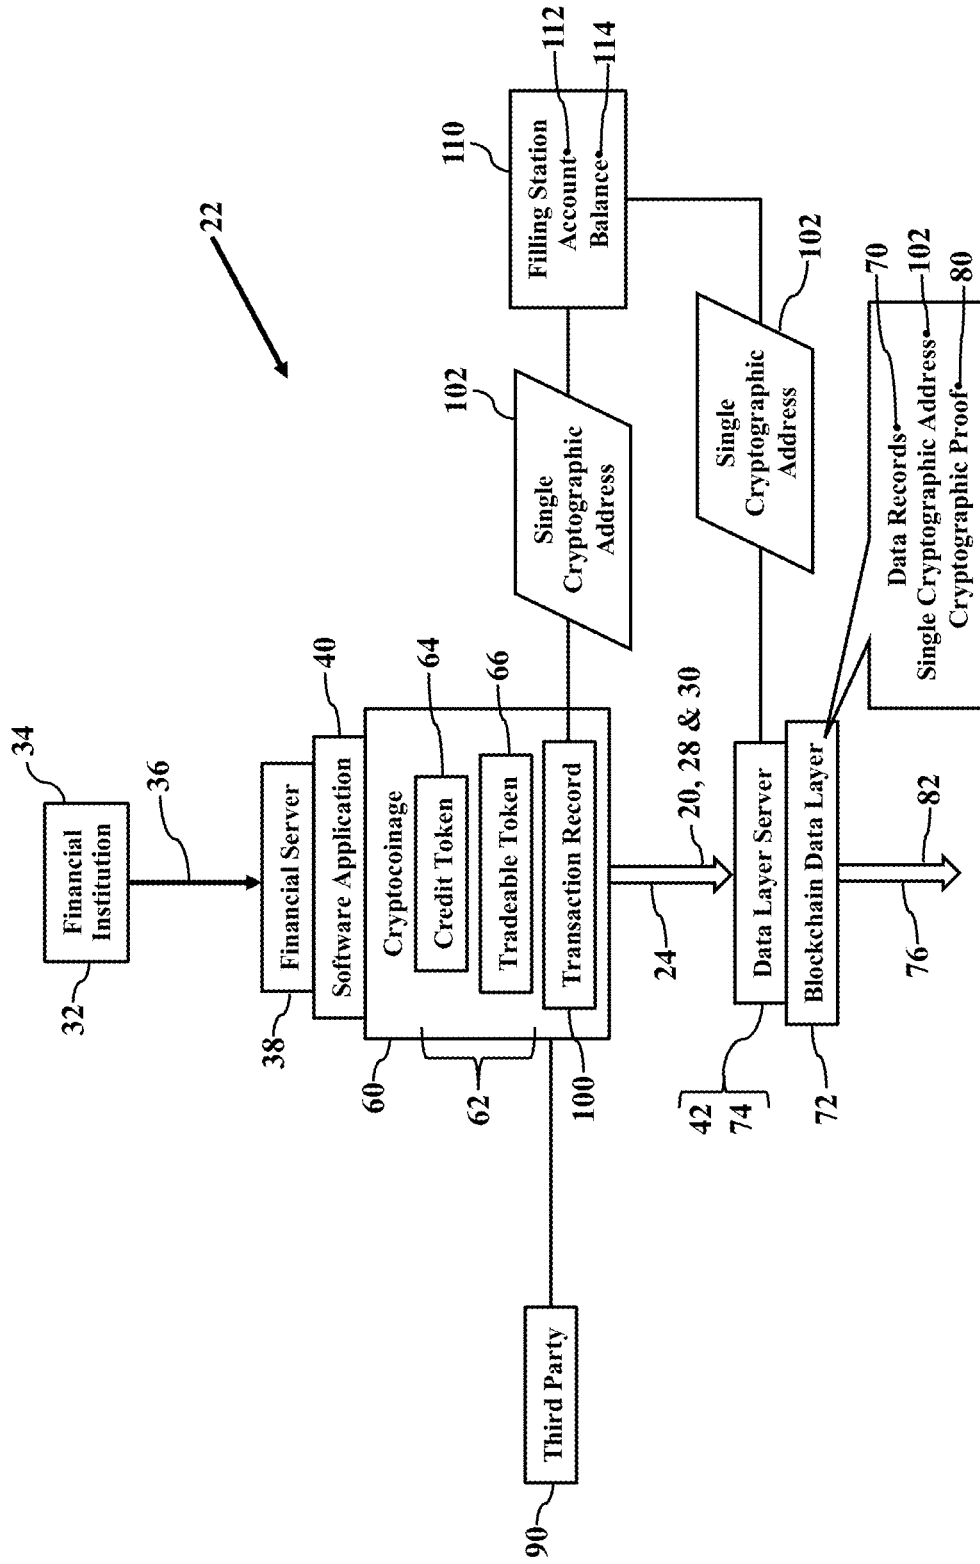


FIG. 19

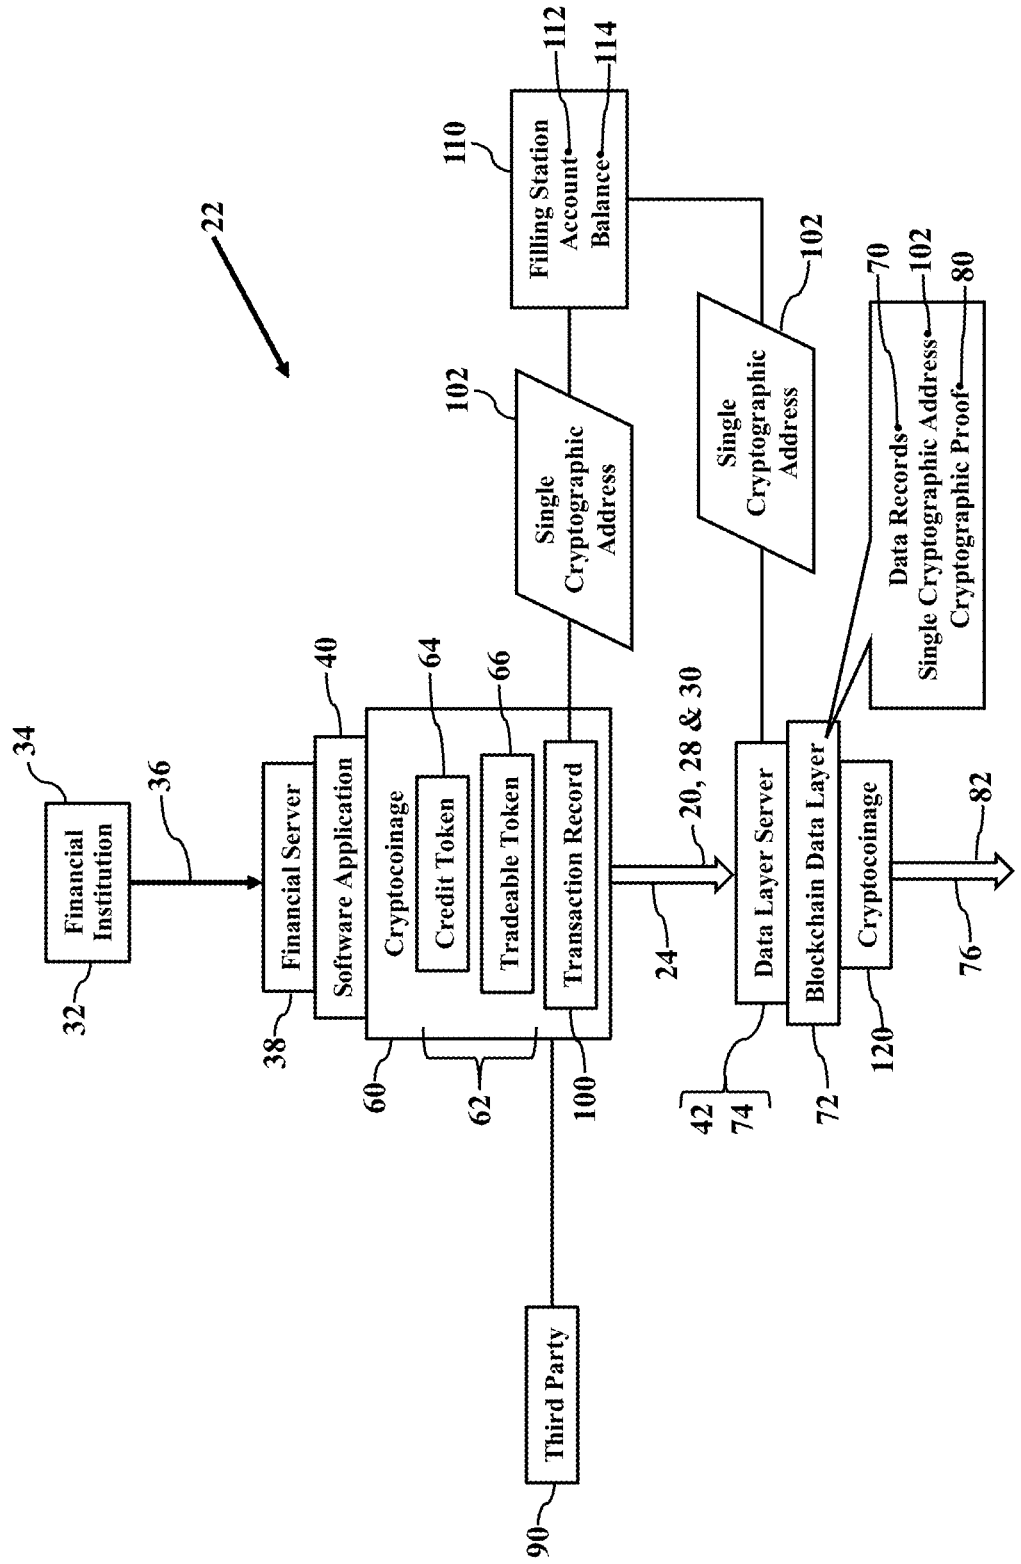


FIG. 20

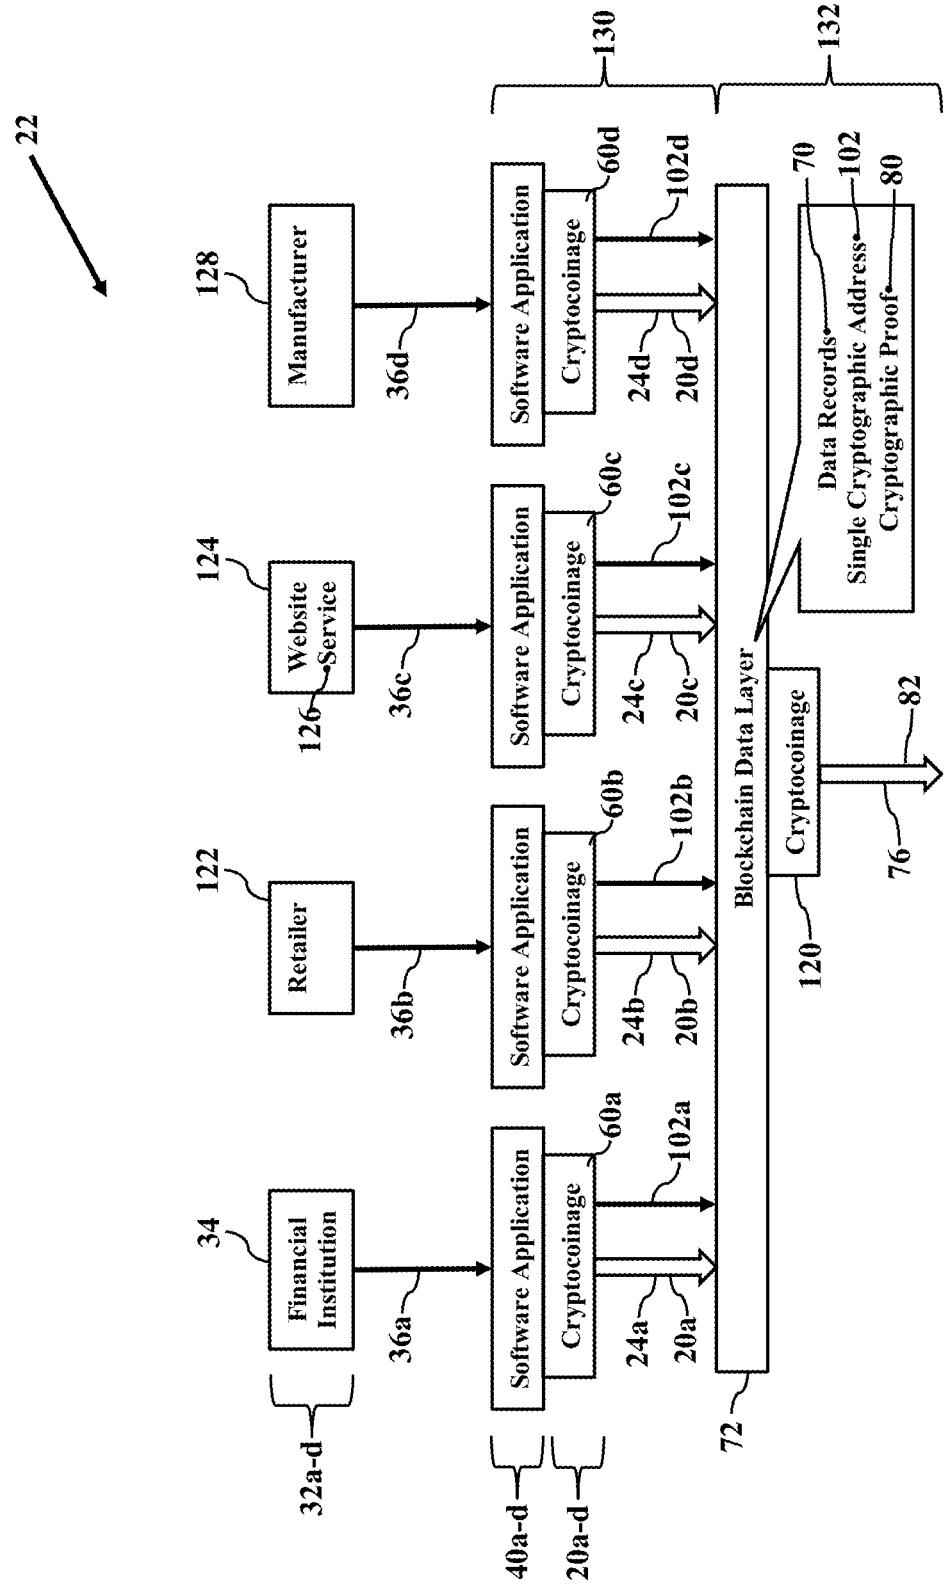


FIG. 21

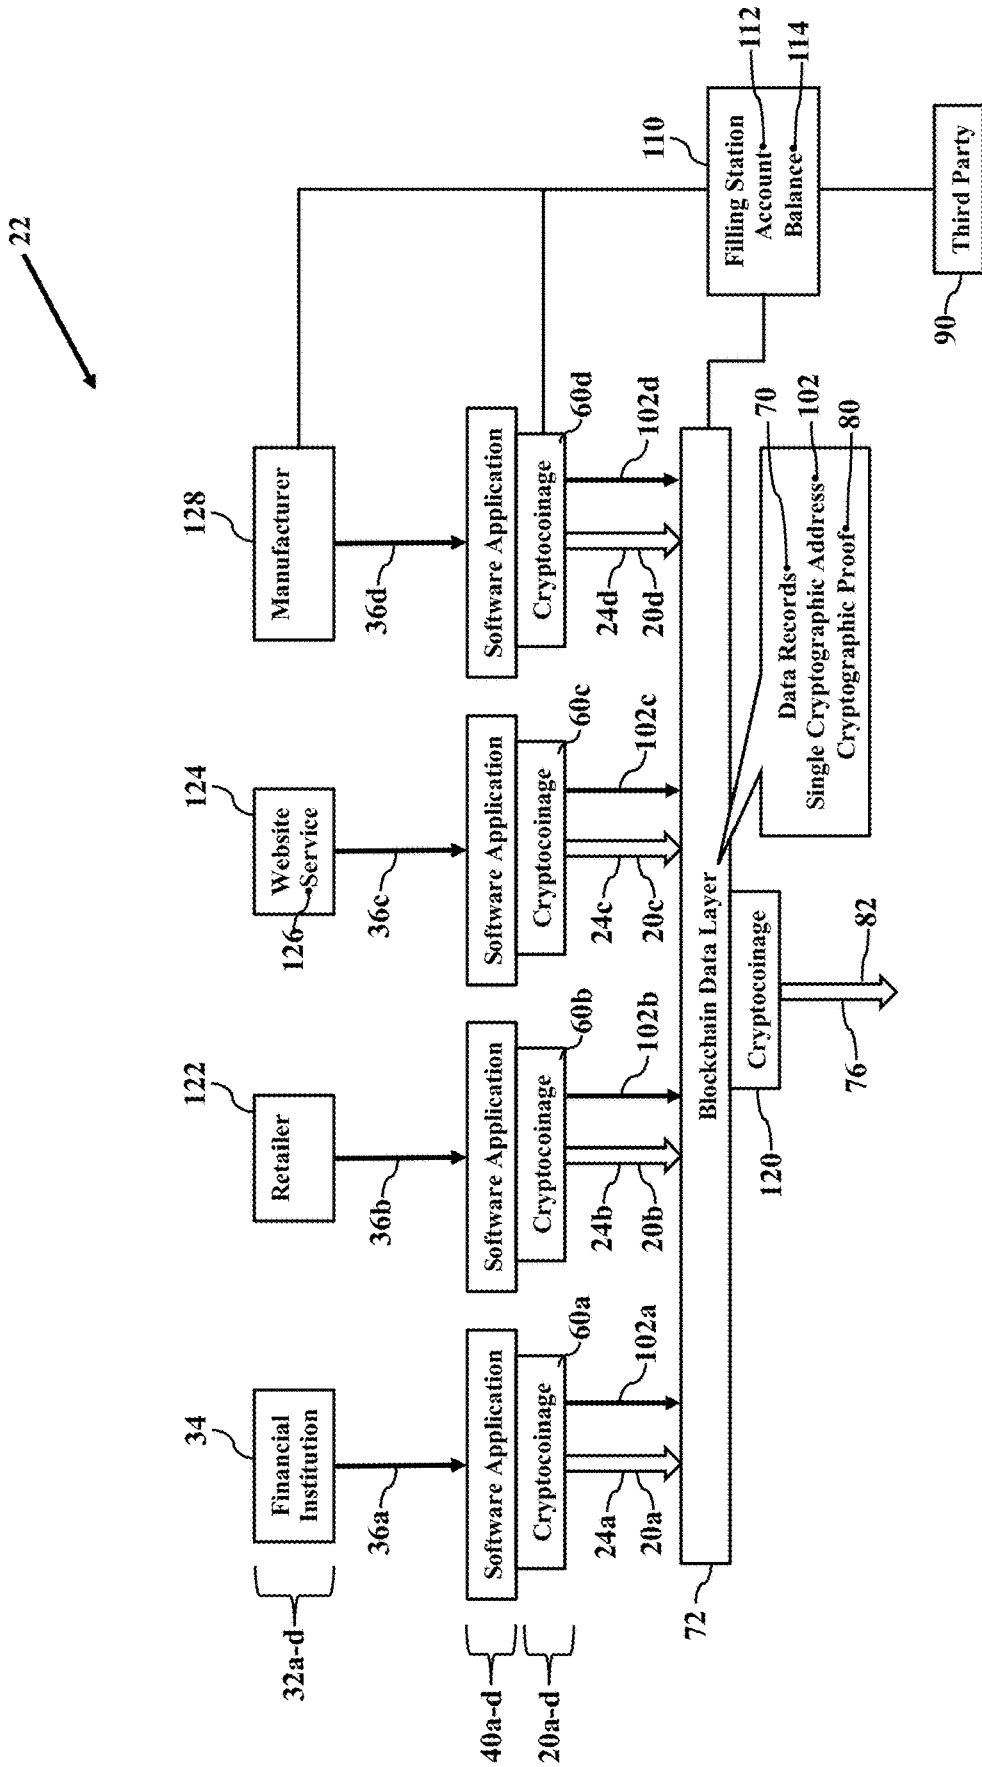


FIG. 22

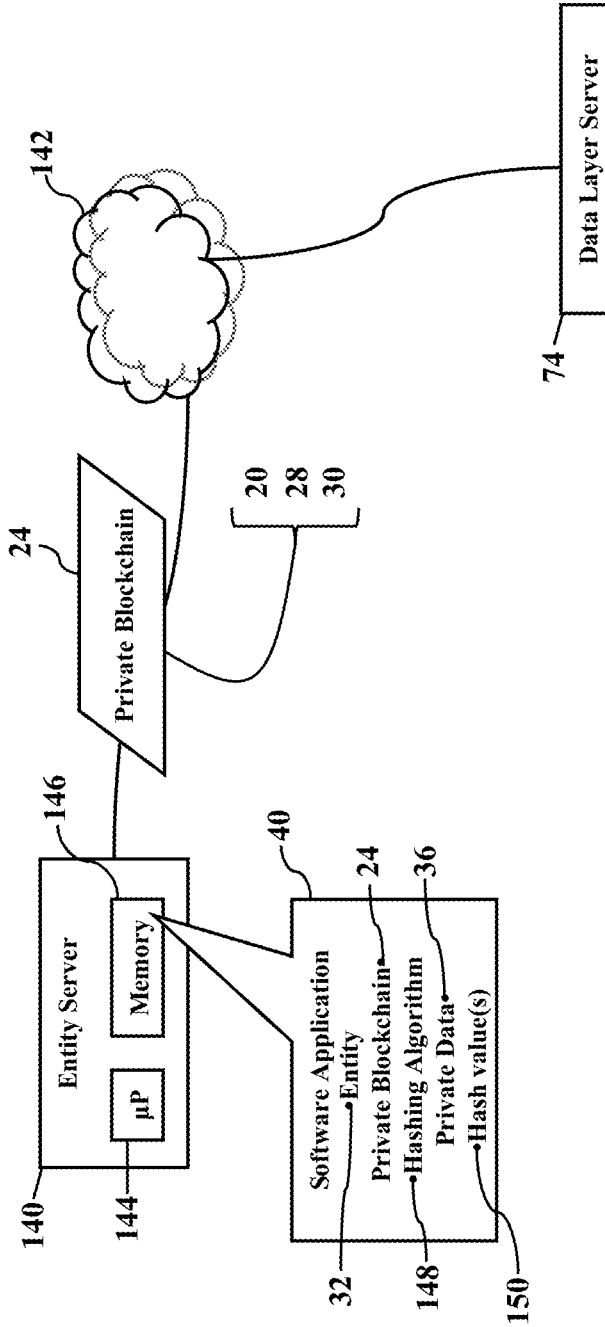


FIG. 23

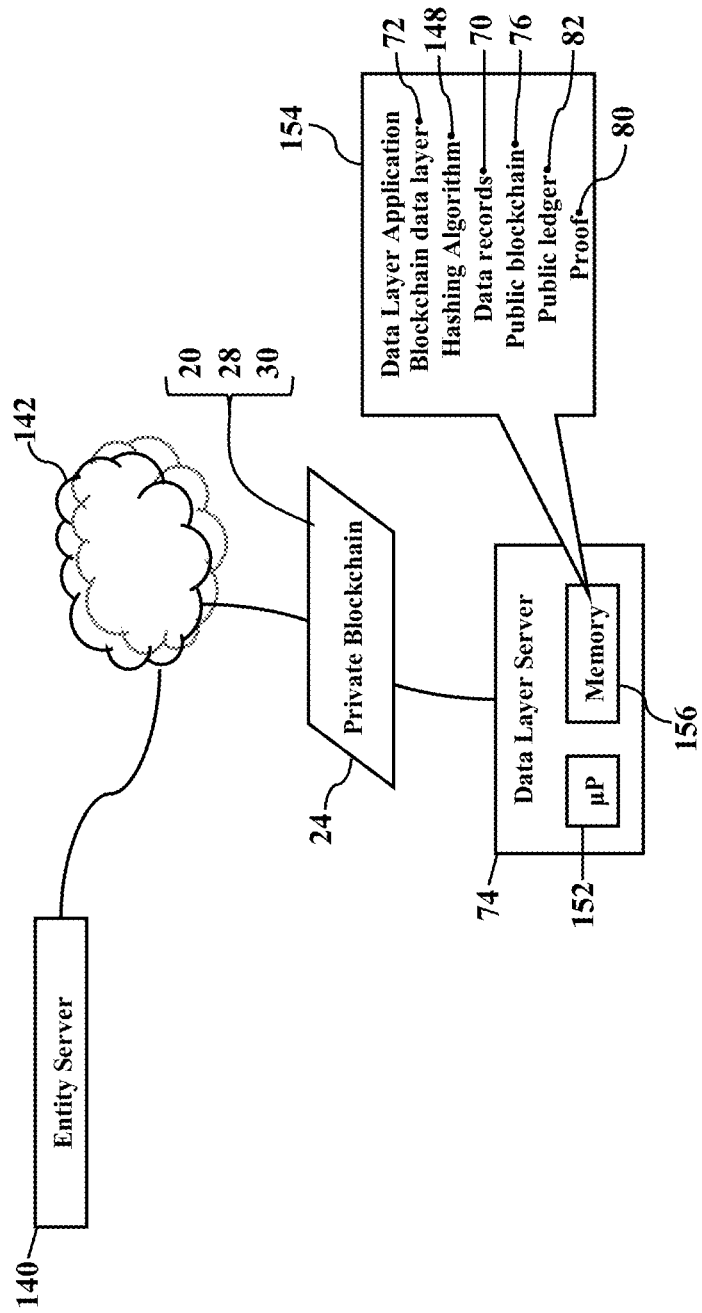


FIG. 24

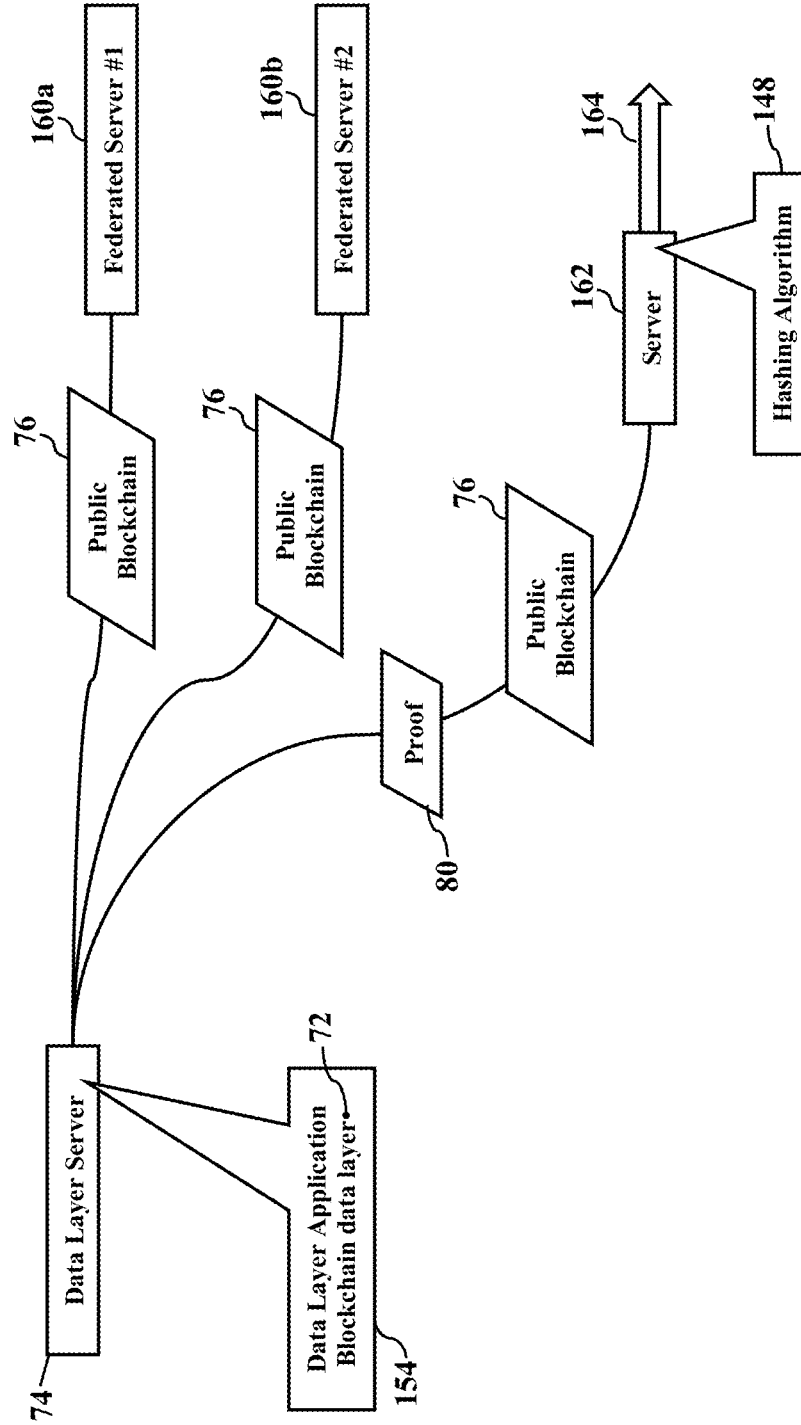


FIG. 25

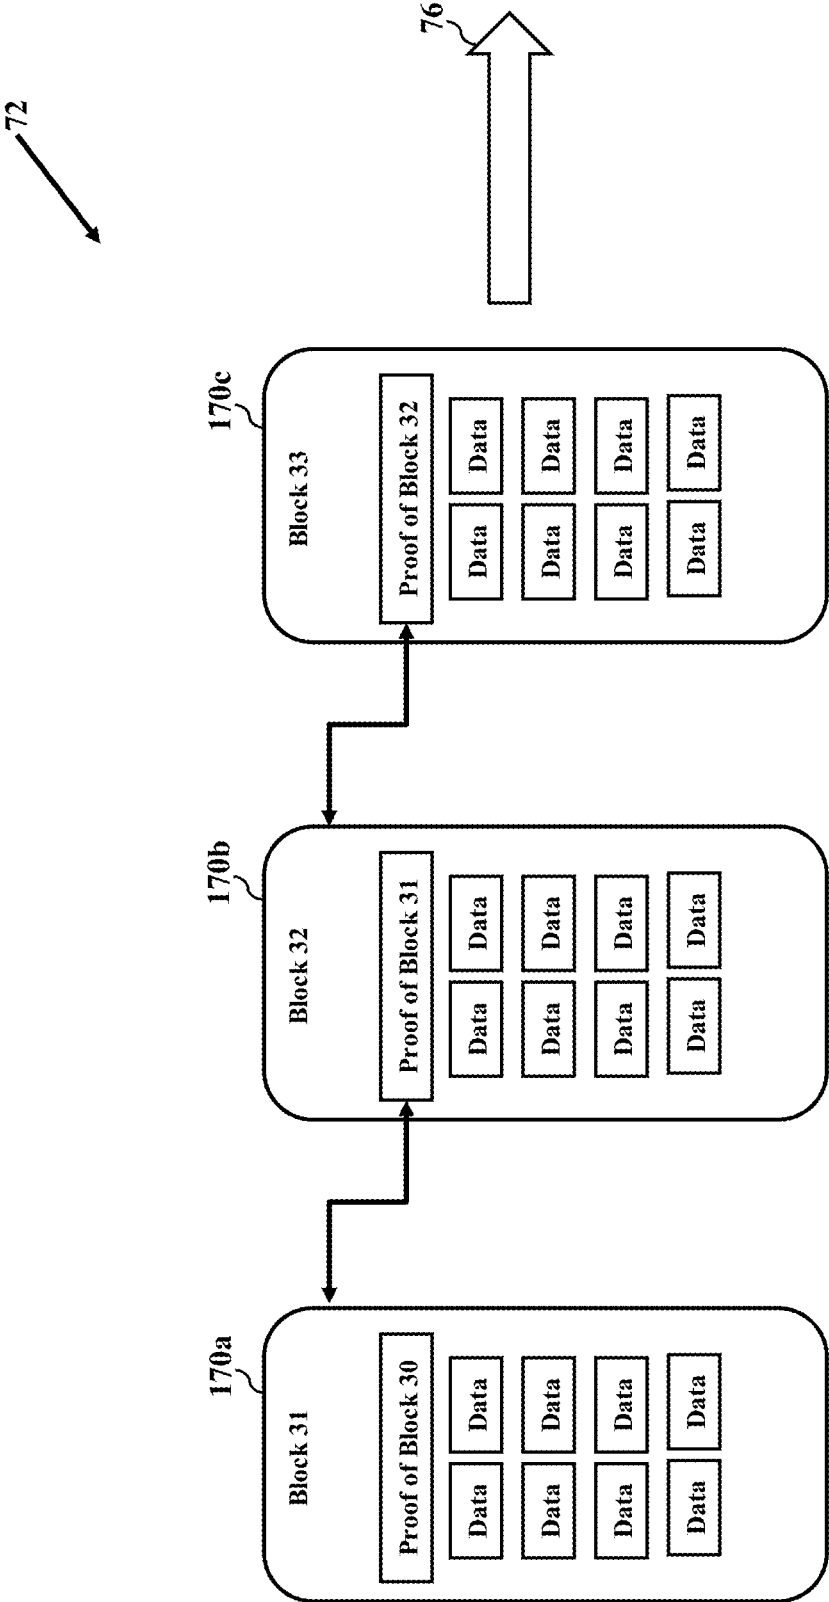


FIG. 26

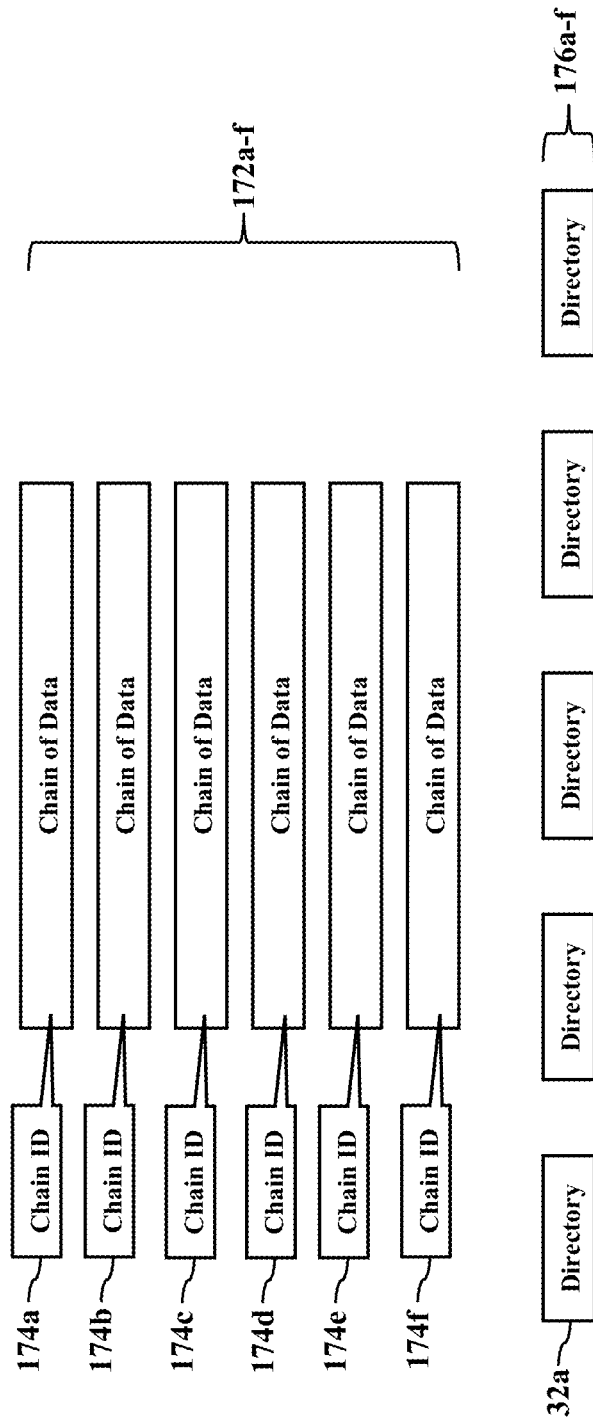


FIG. 27

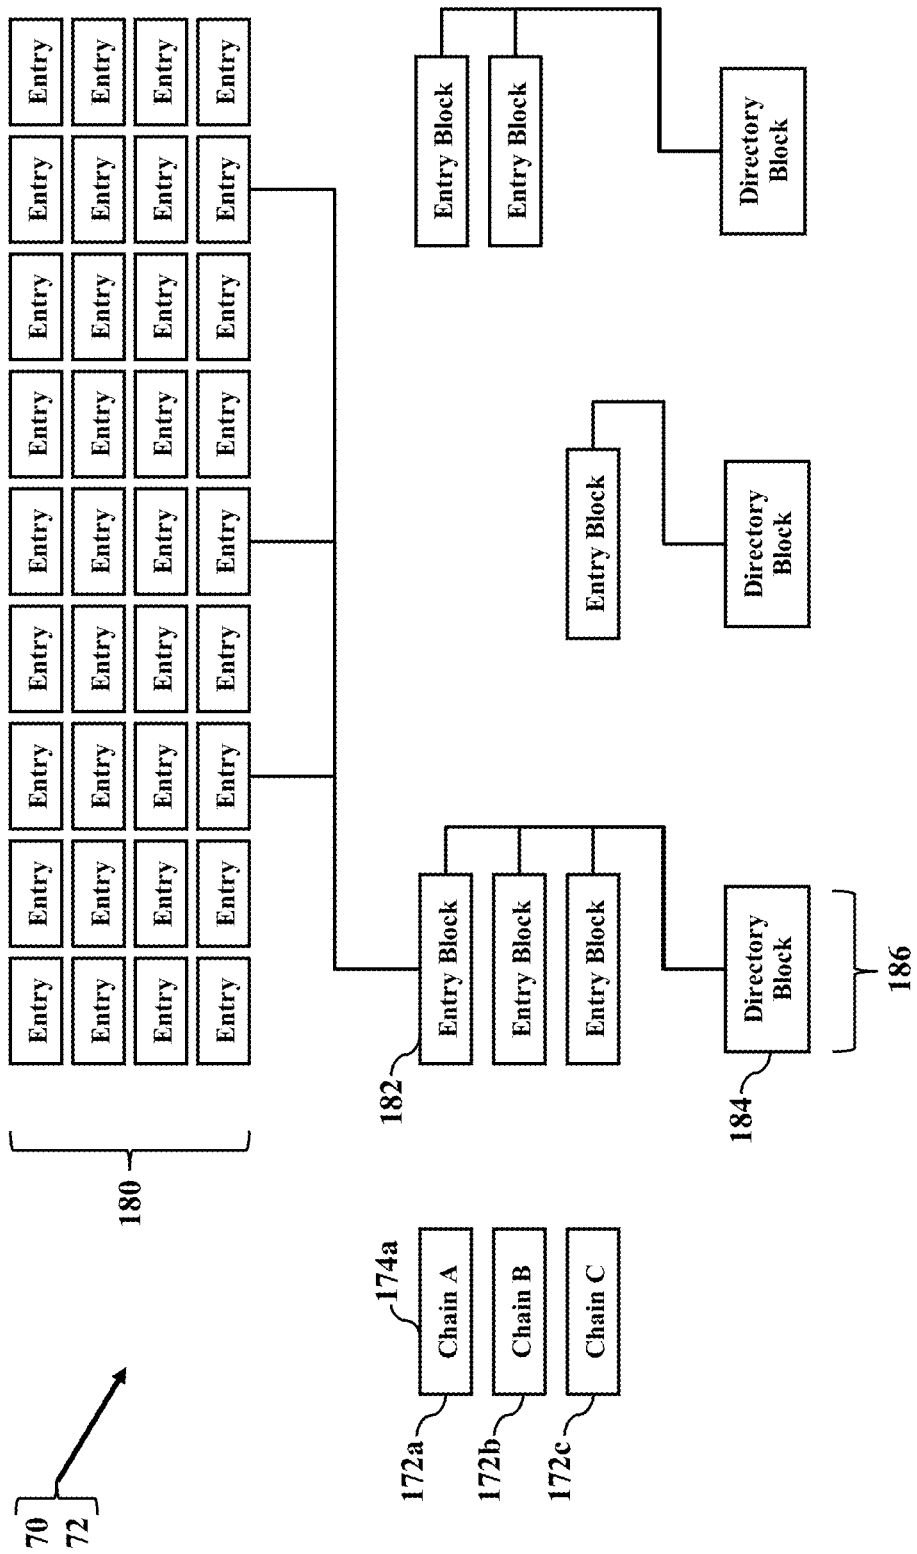


FIG. 28

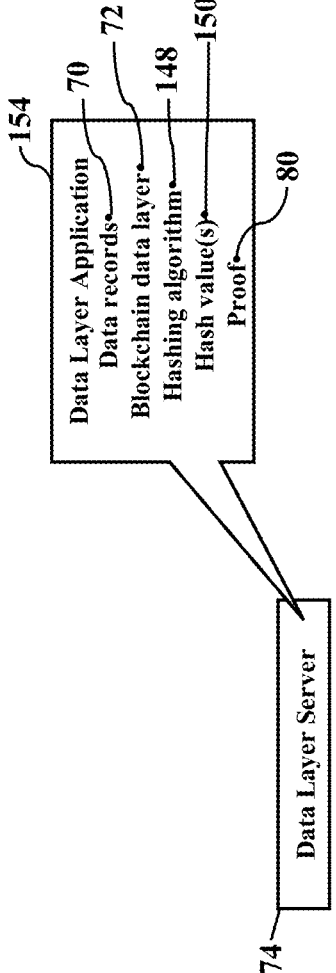


FIG. 29

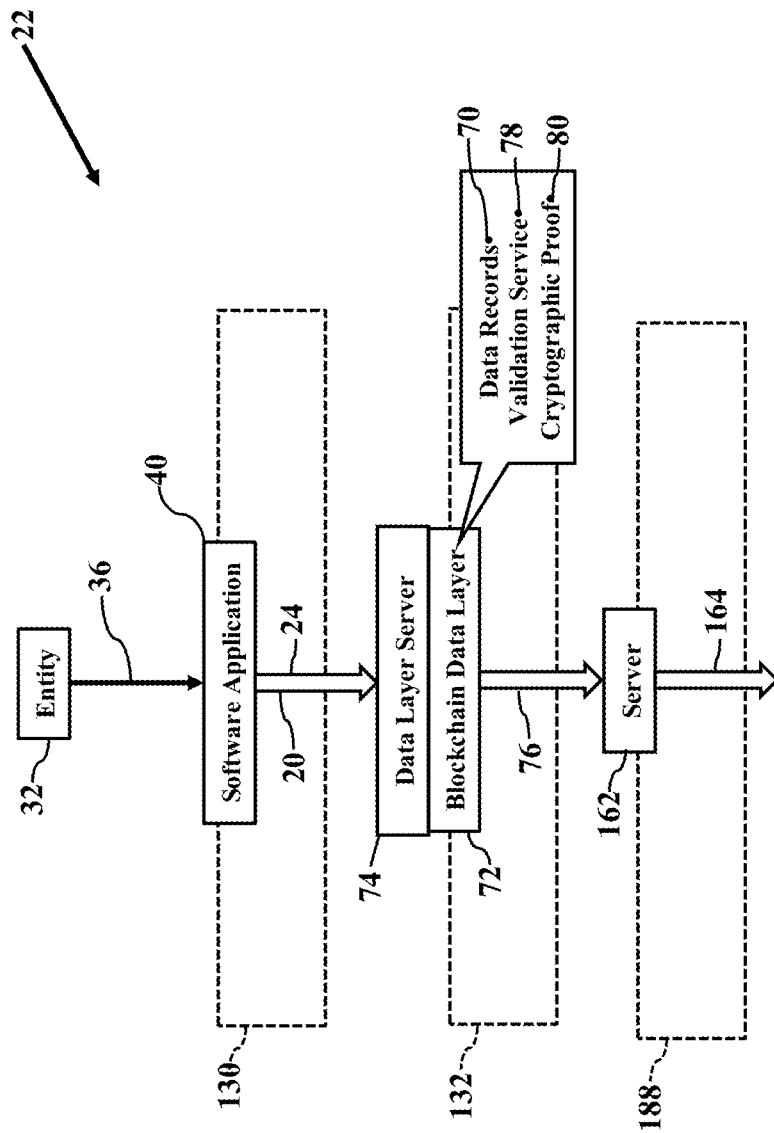


FIG. 30

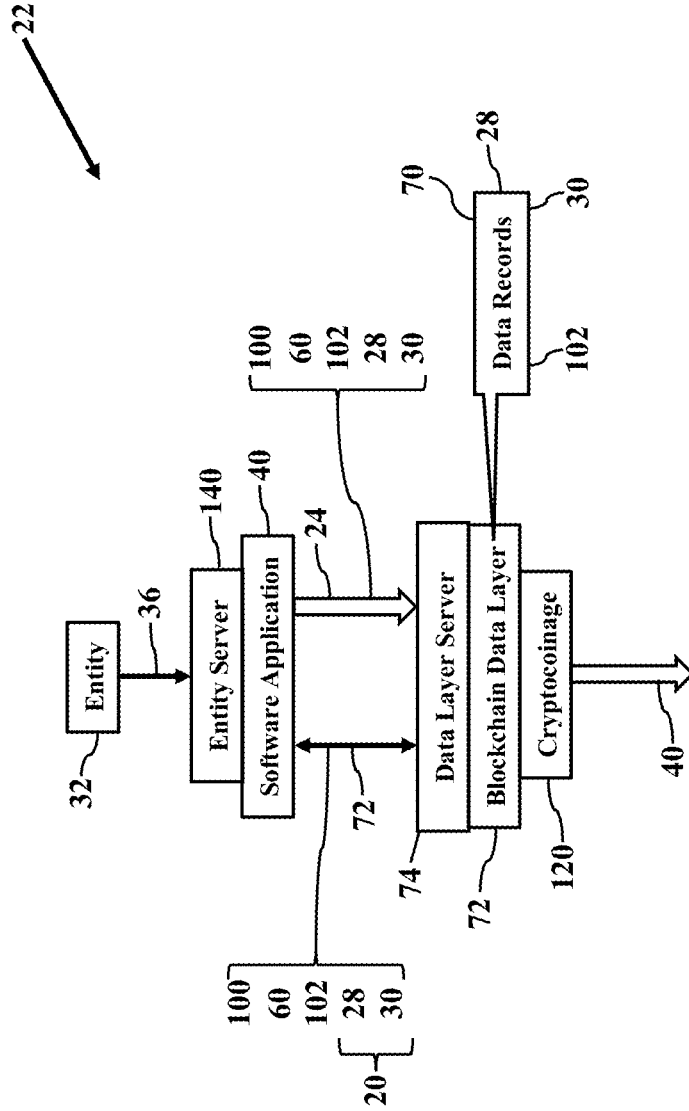


FIG. 31

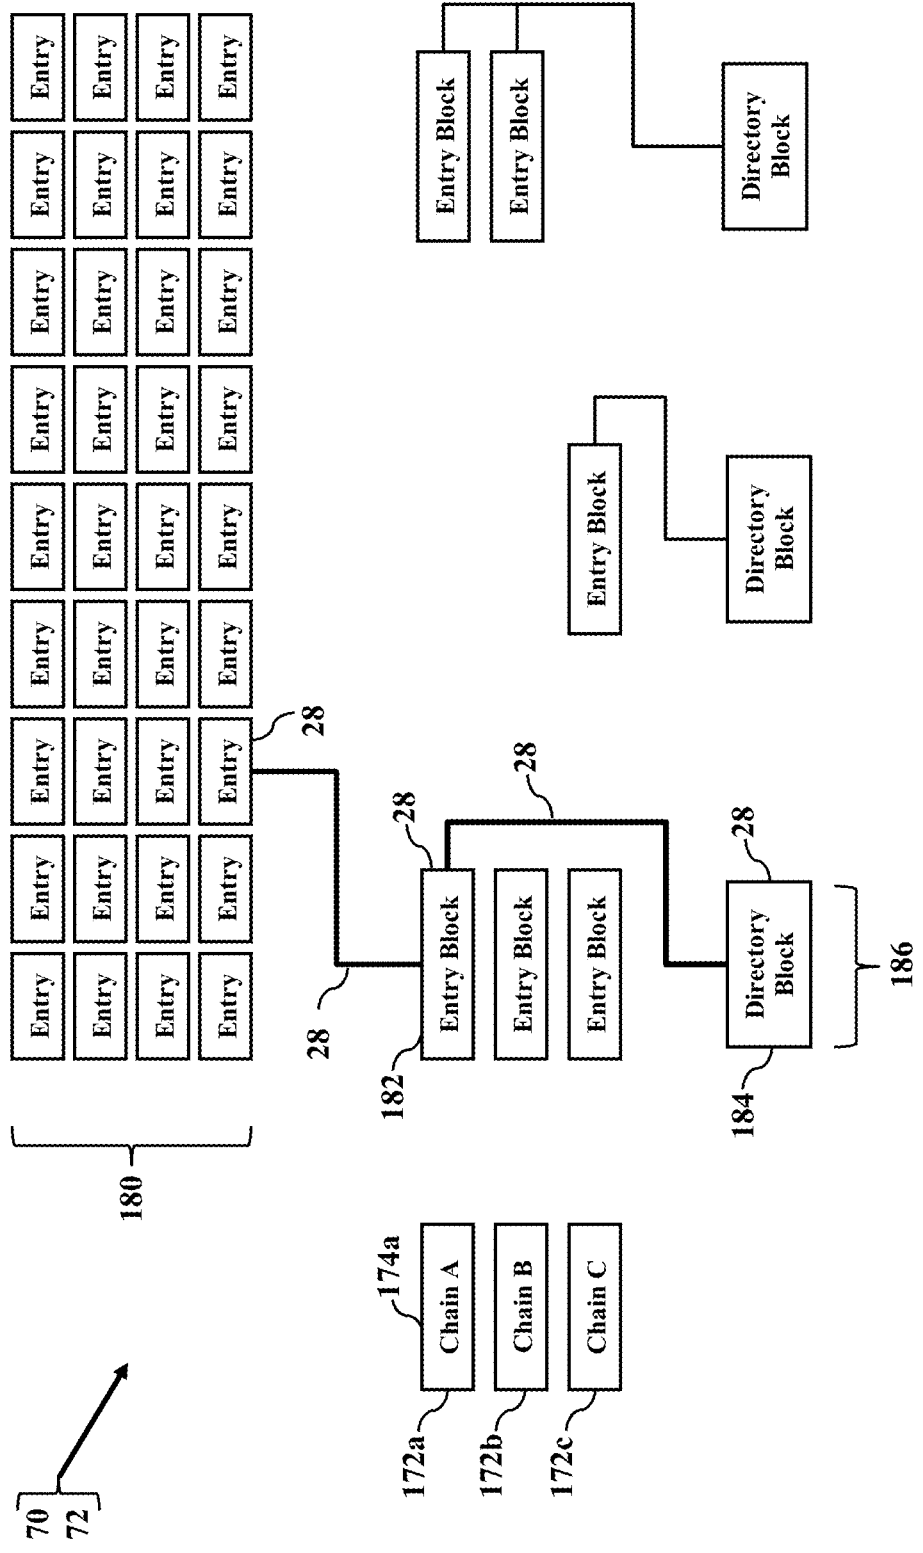


FIG. 32

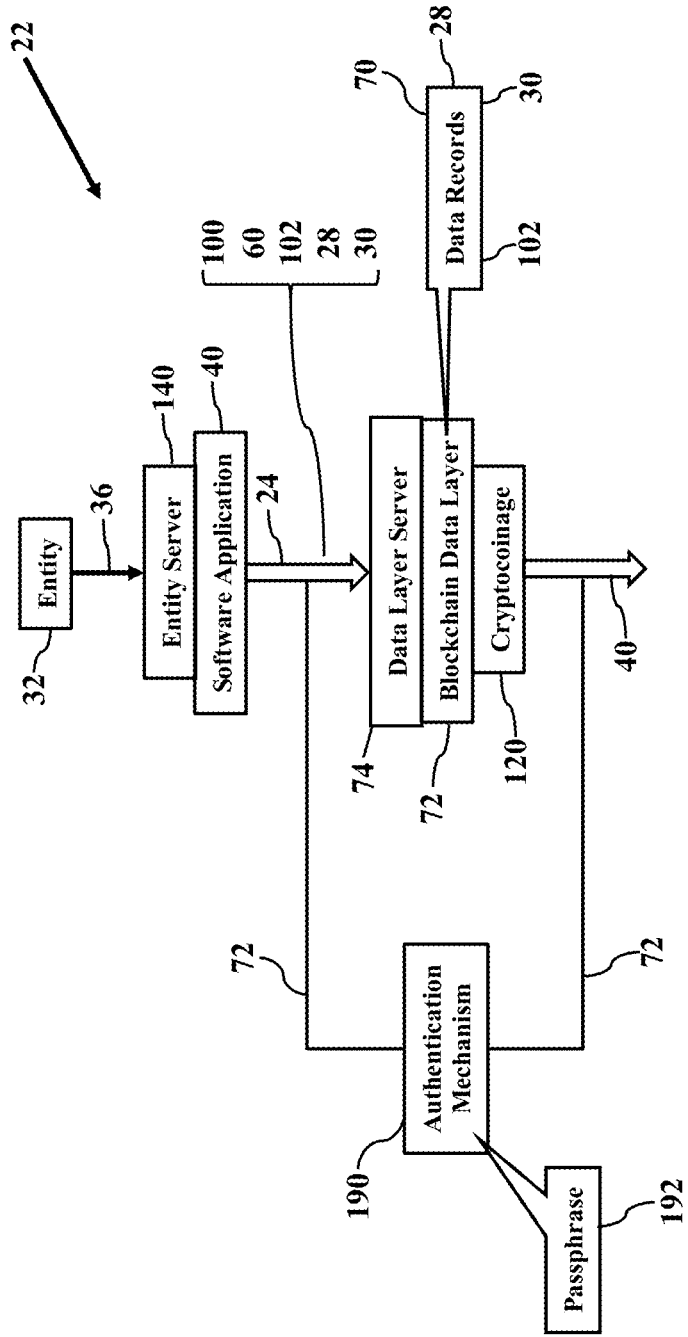


FIG. 33

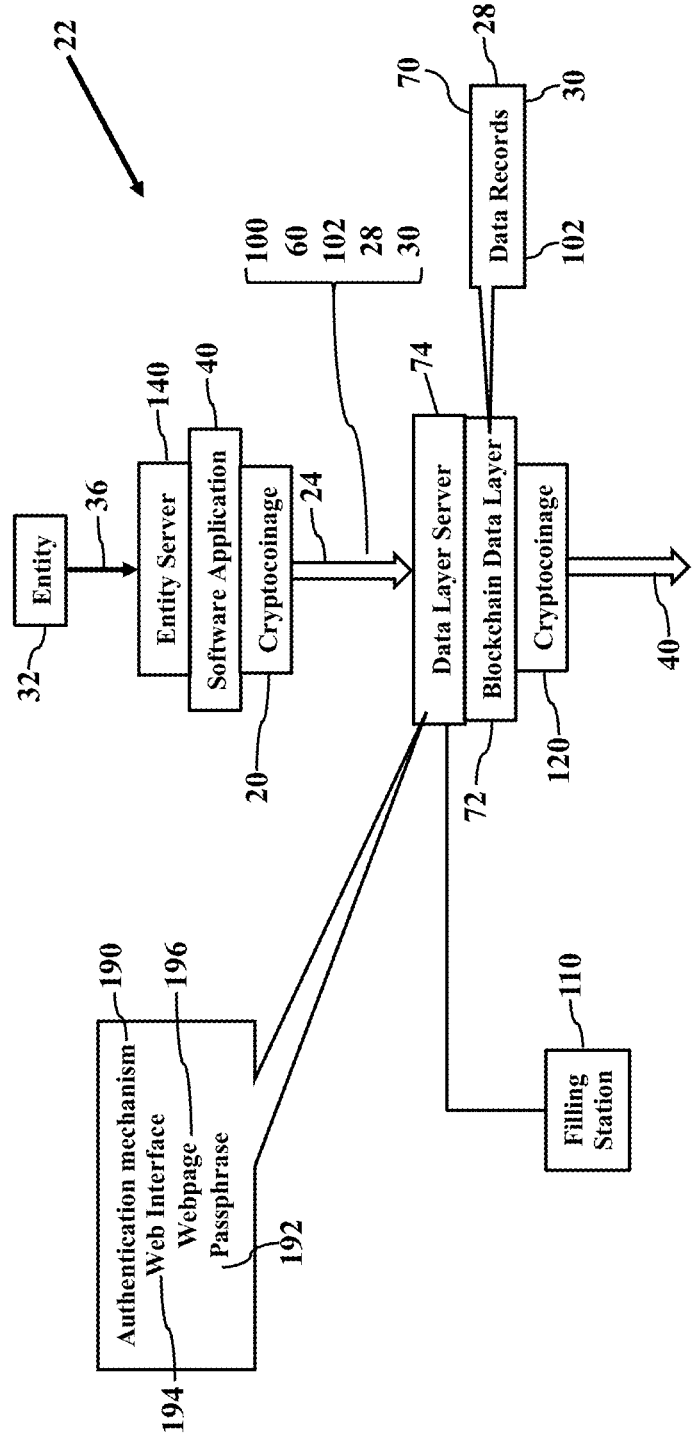


FIG. 34

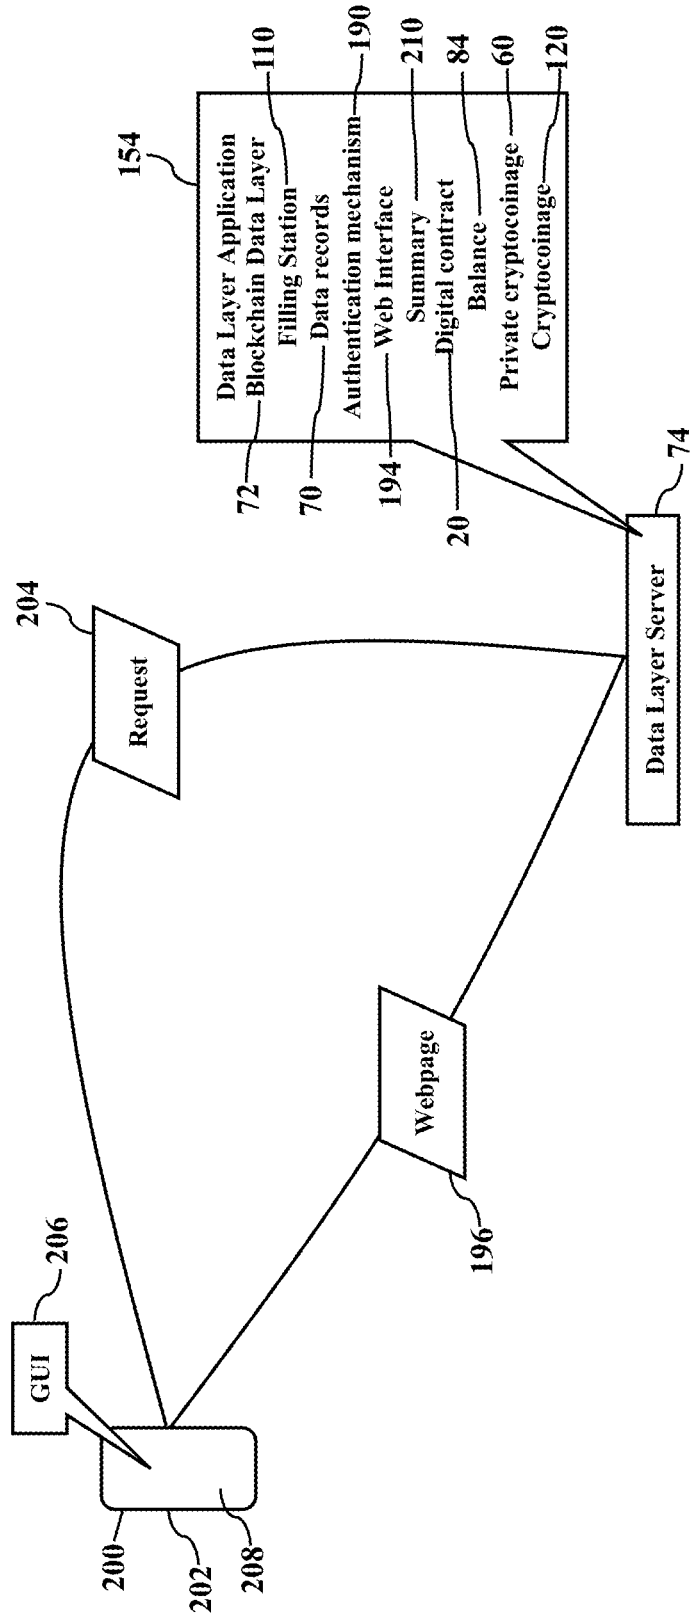


FIG. 35

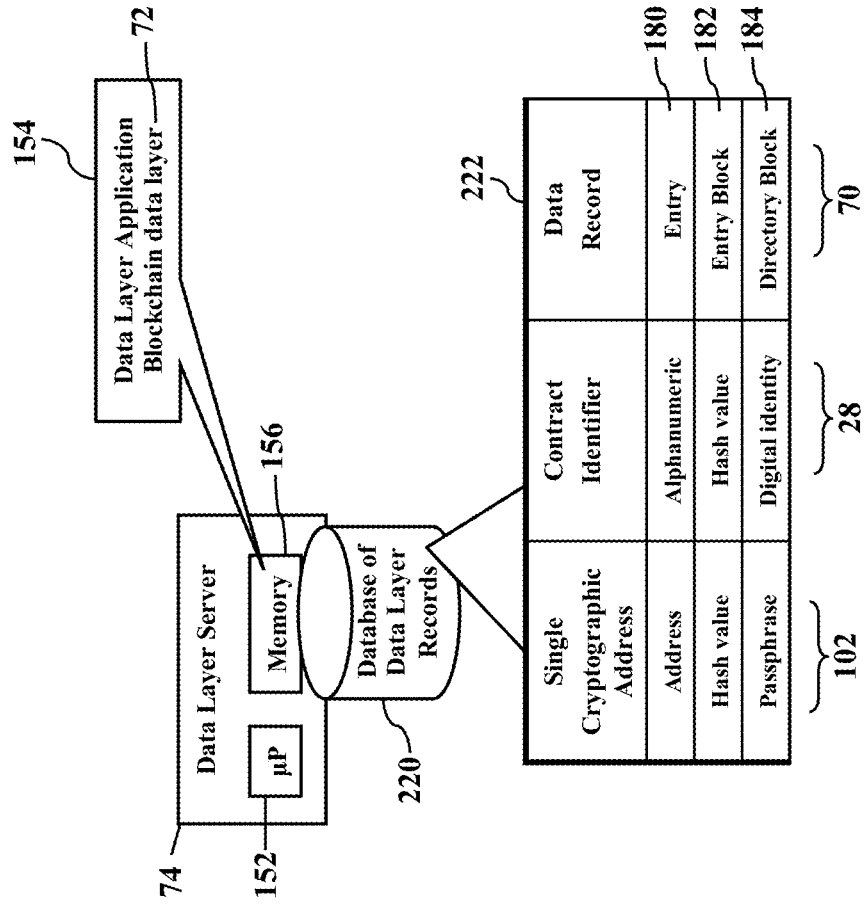


FIG. 36

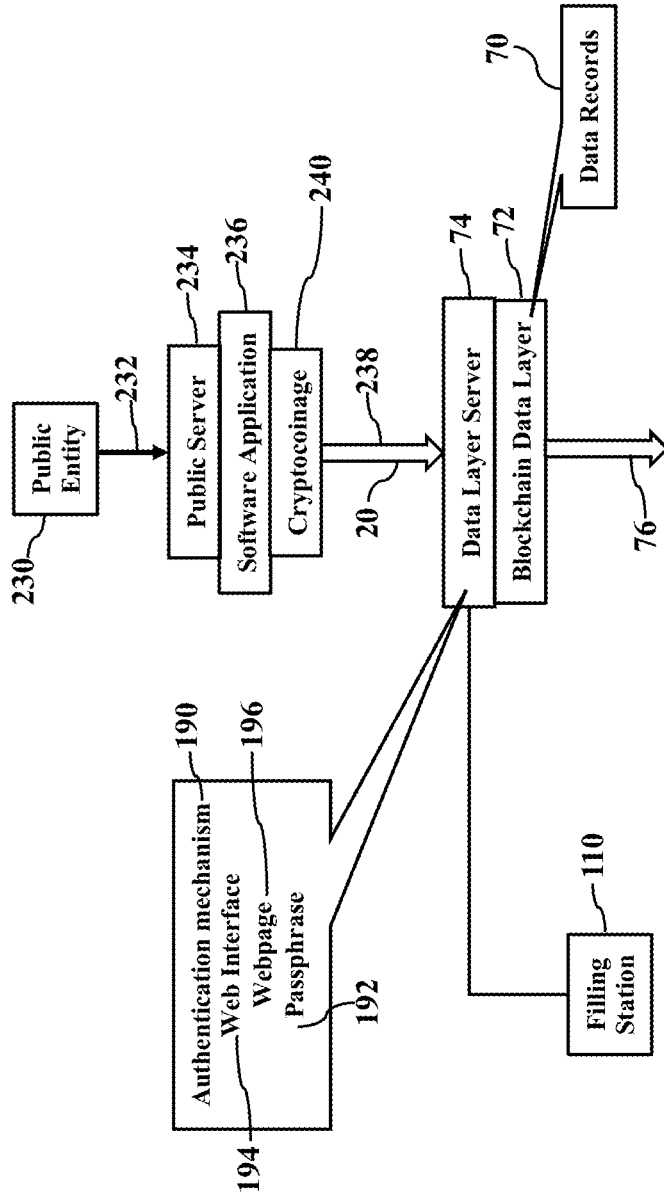


FIG. 37

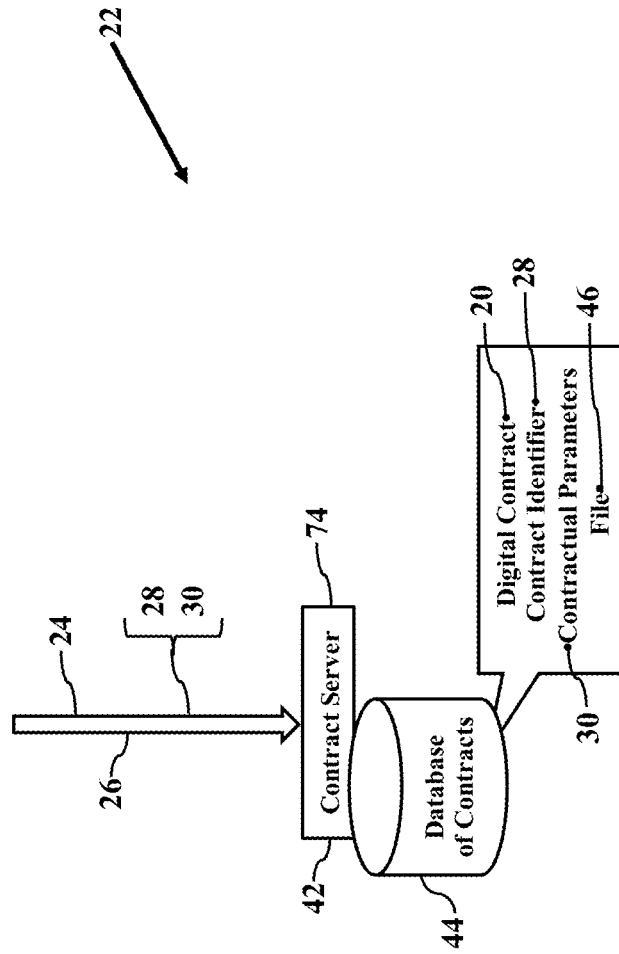


FIG. 38

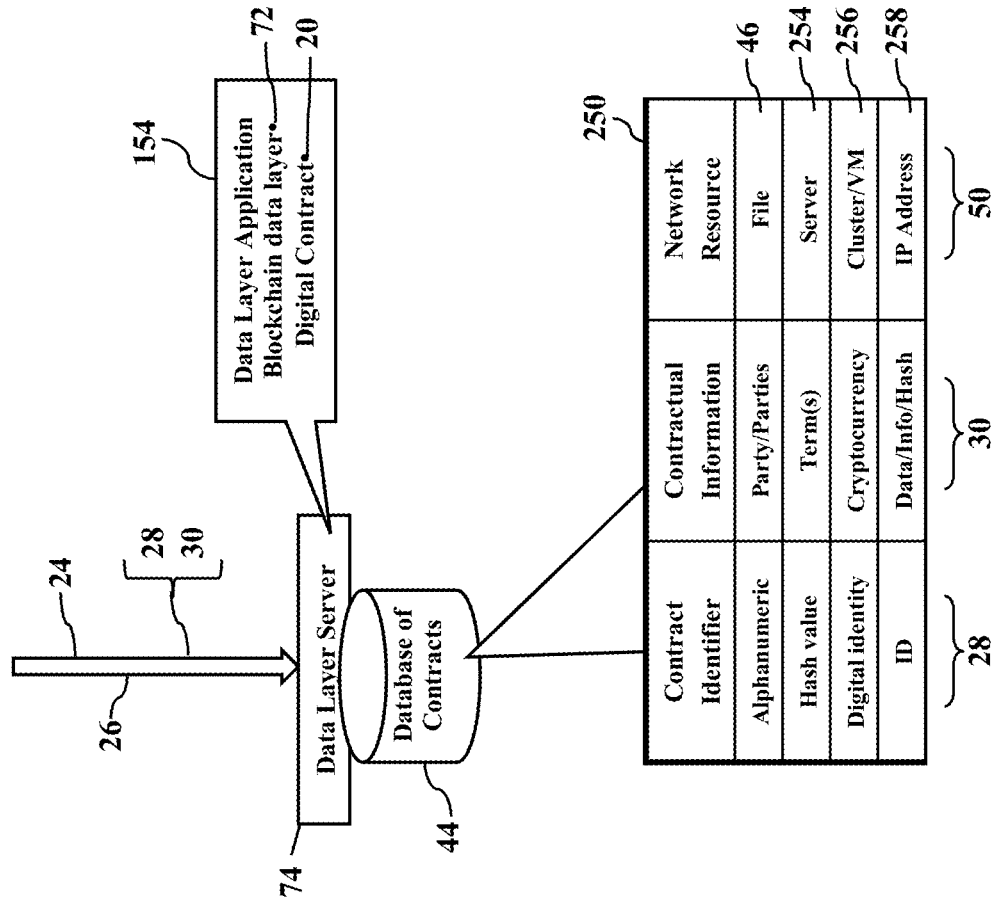


FIG. 39

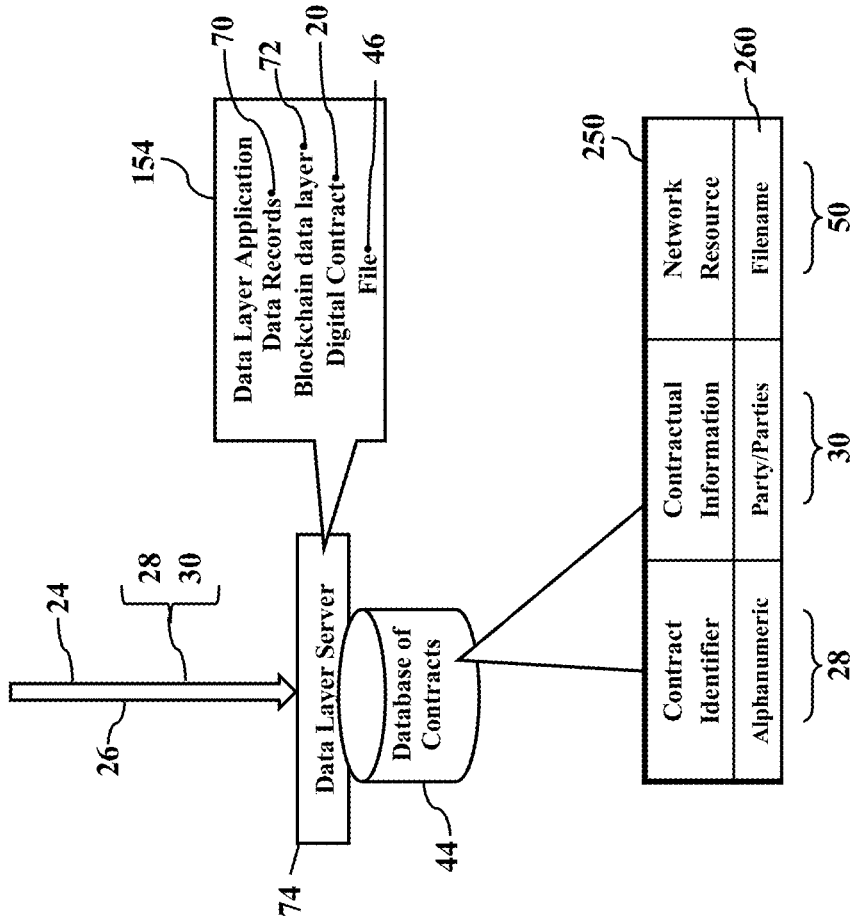


FIG. 40

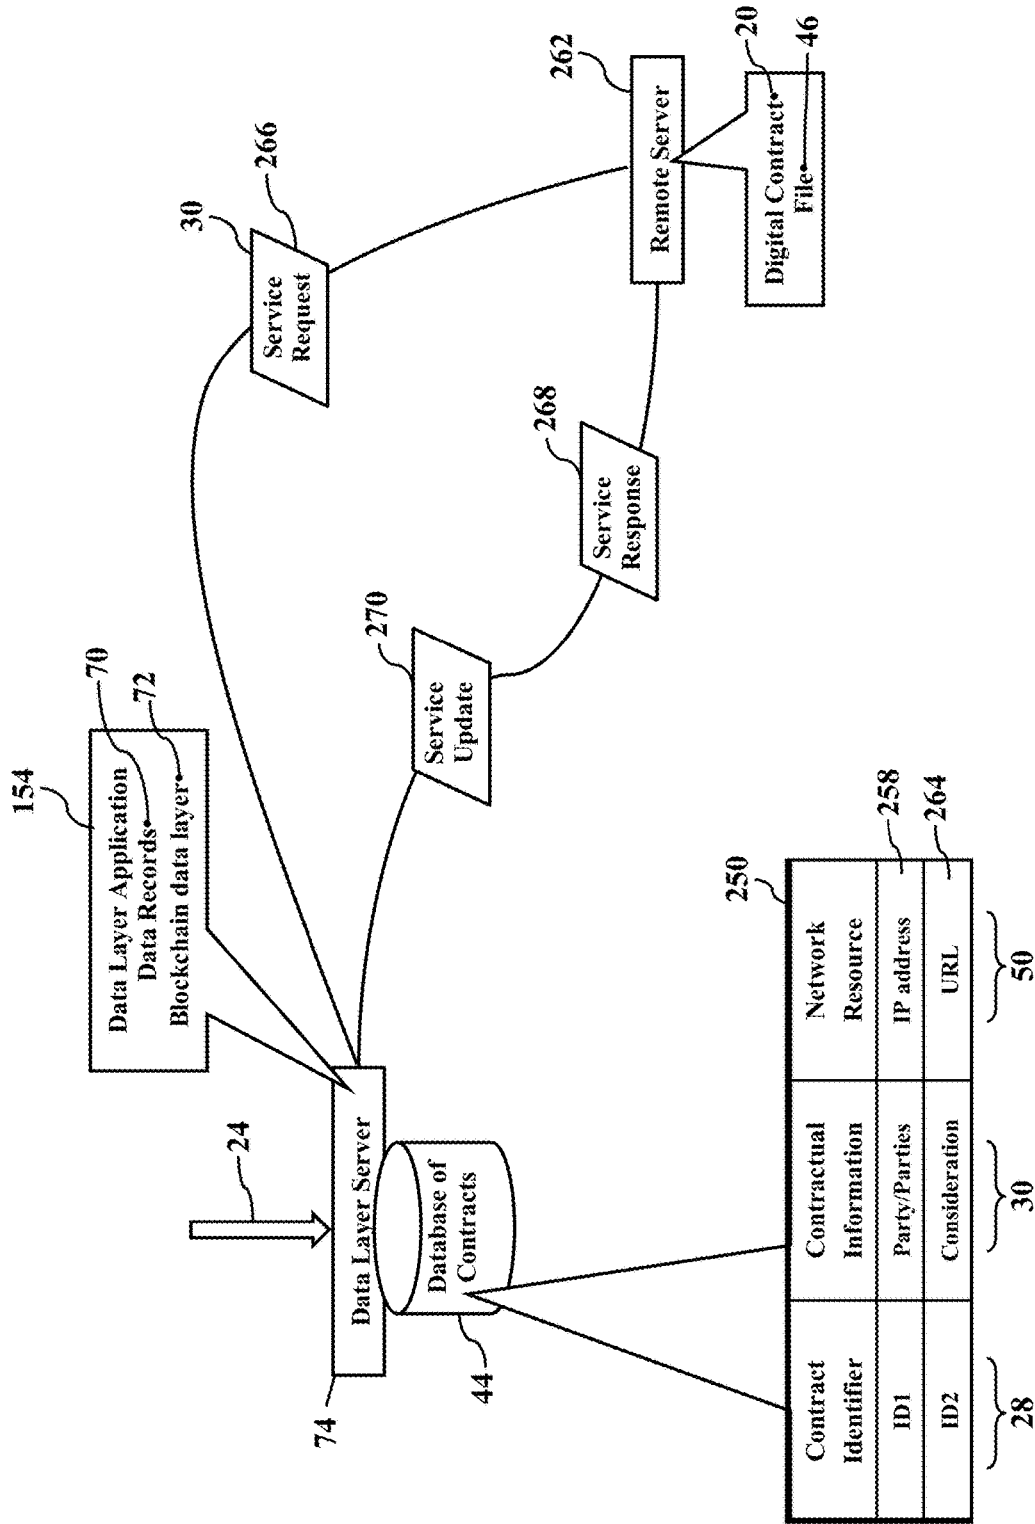


FIG. 41

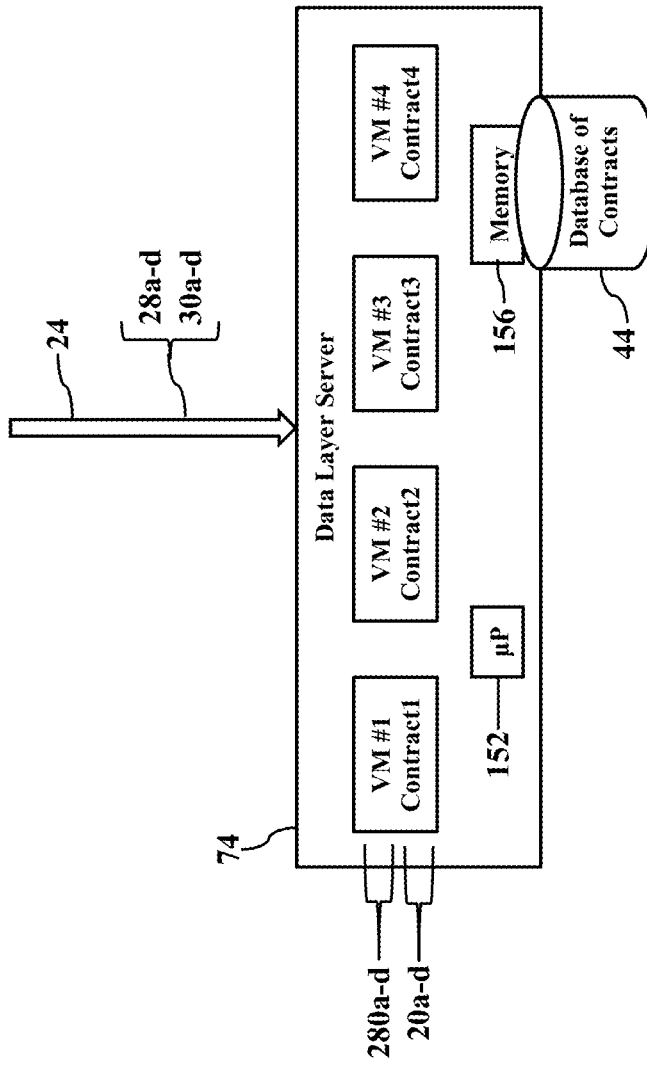


FIG. 42

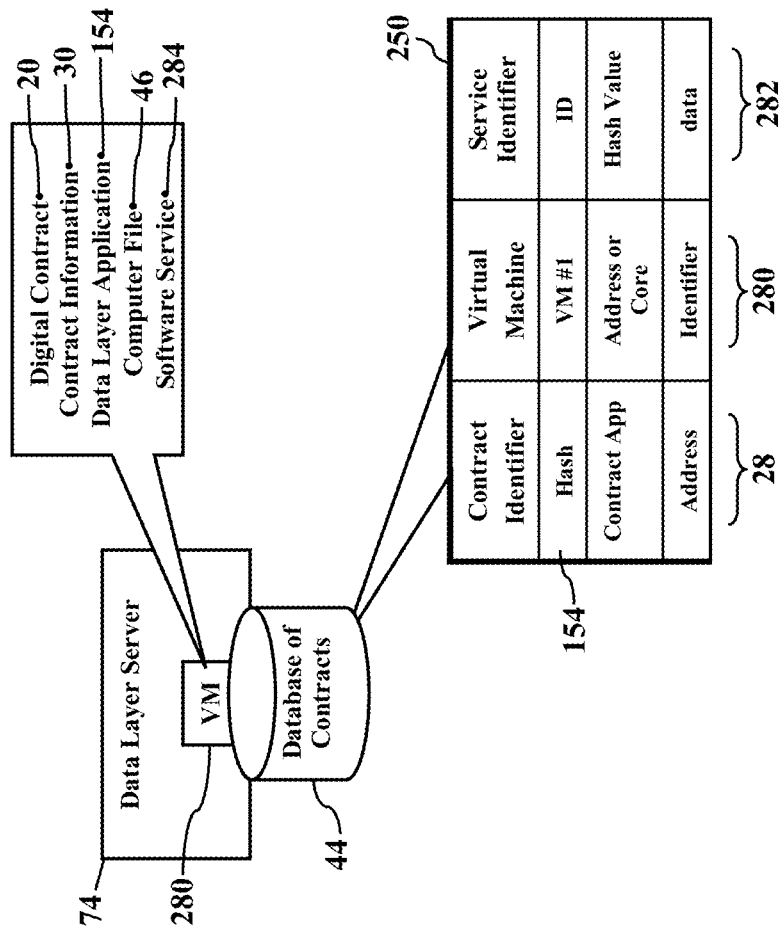


FIG. 43

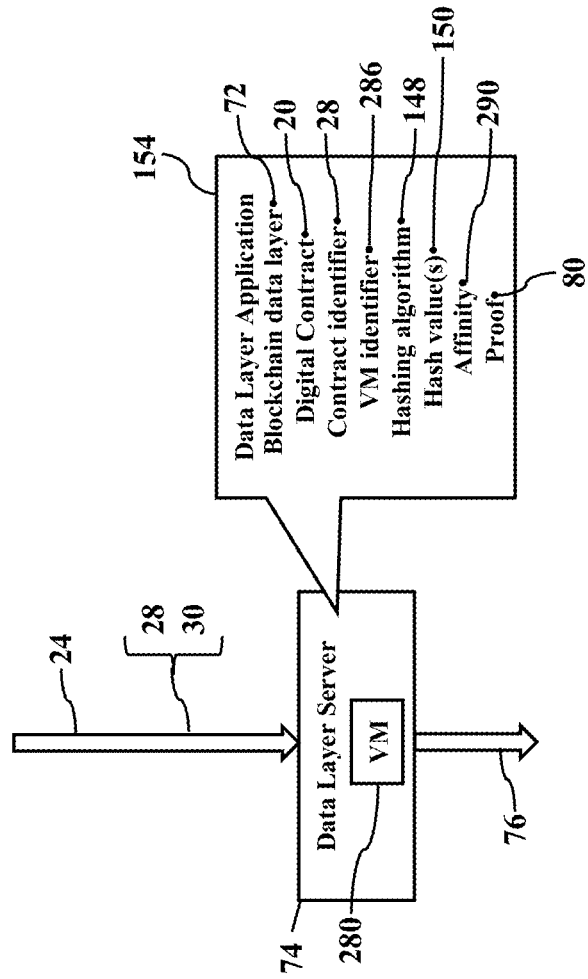


FIG. 44

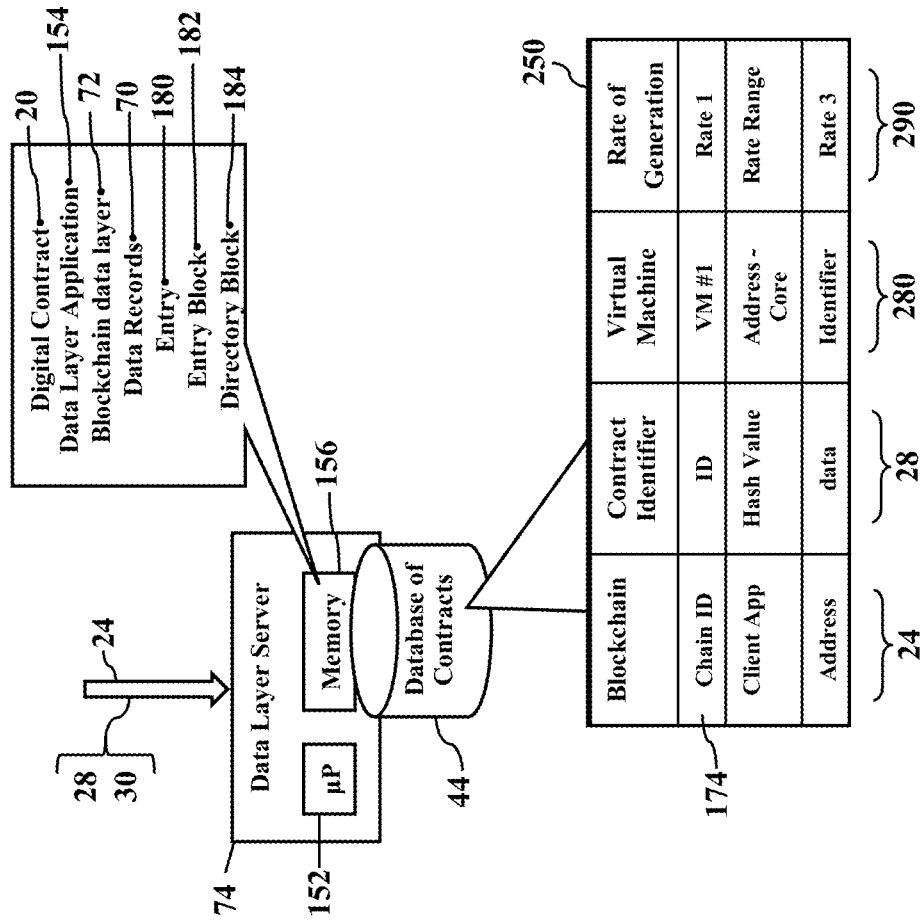


FIG. 45

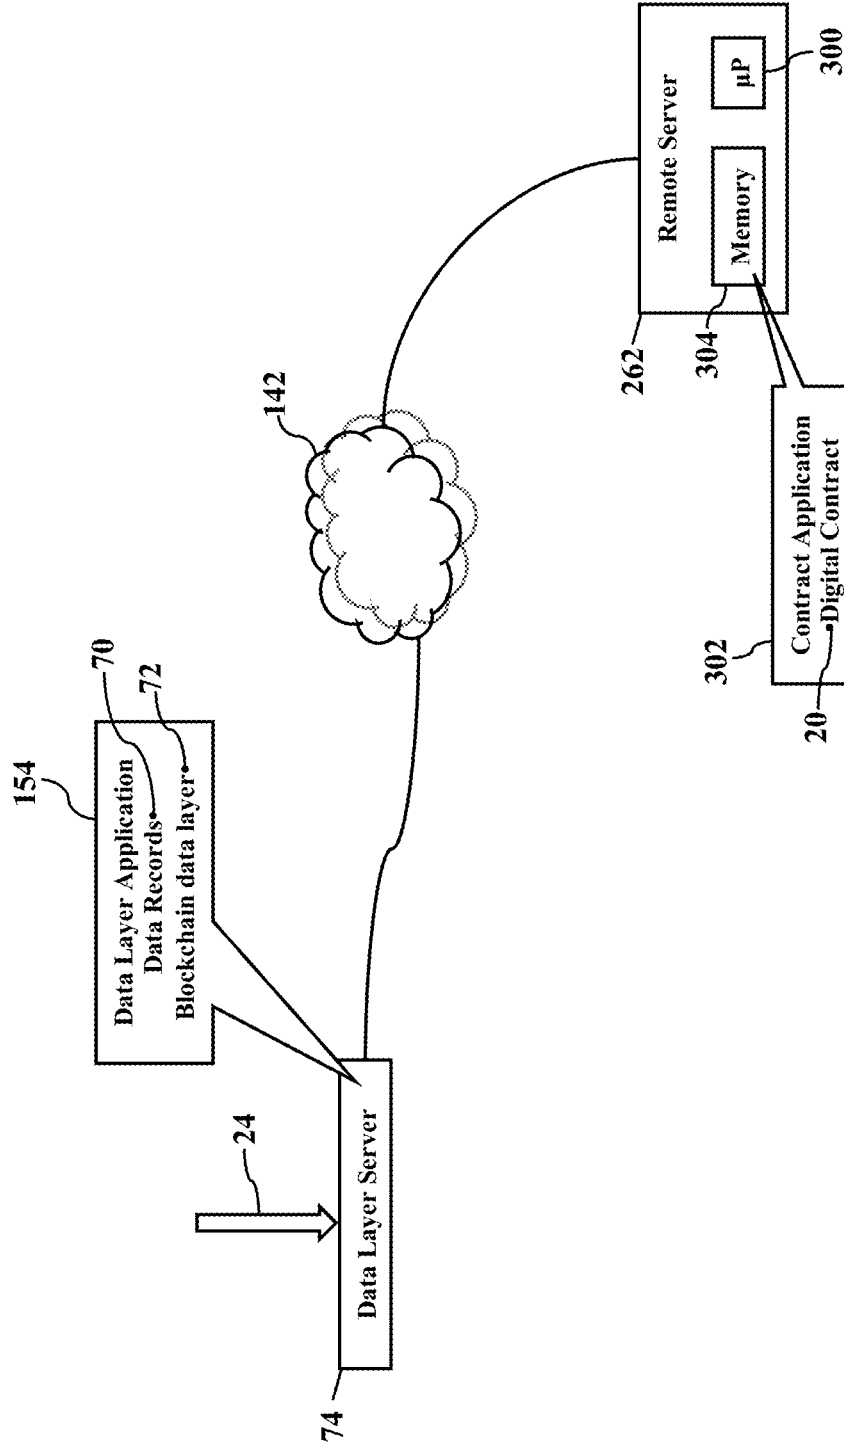


FIG. 46

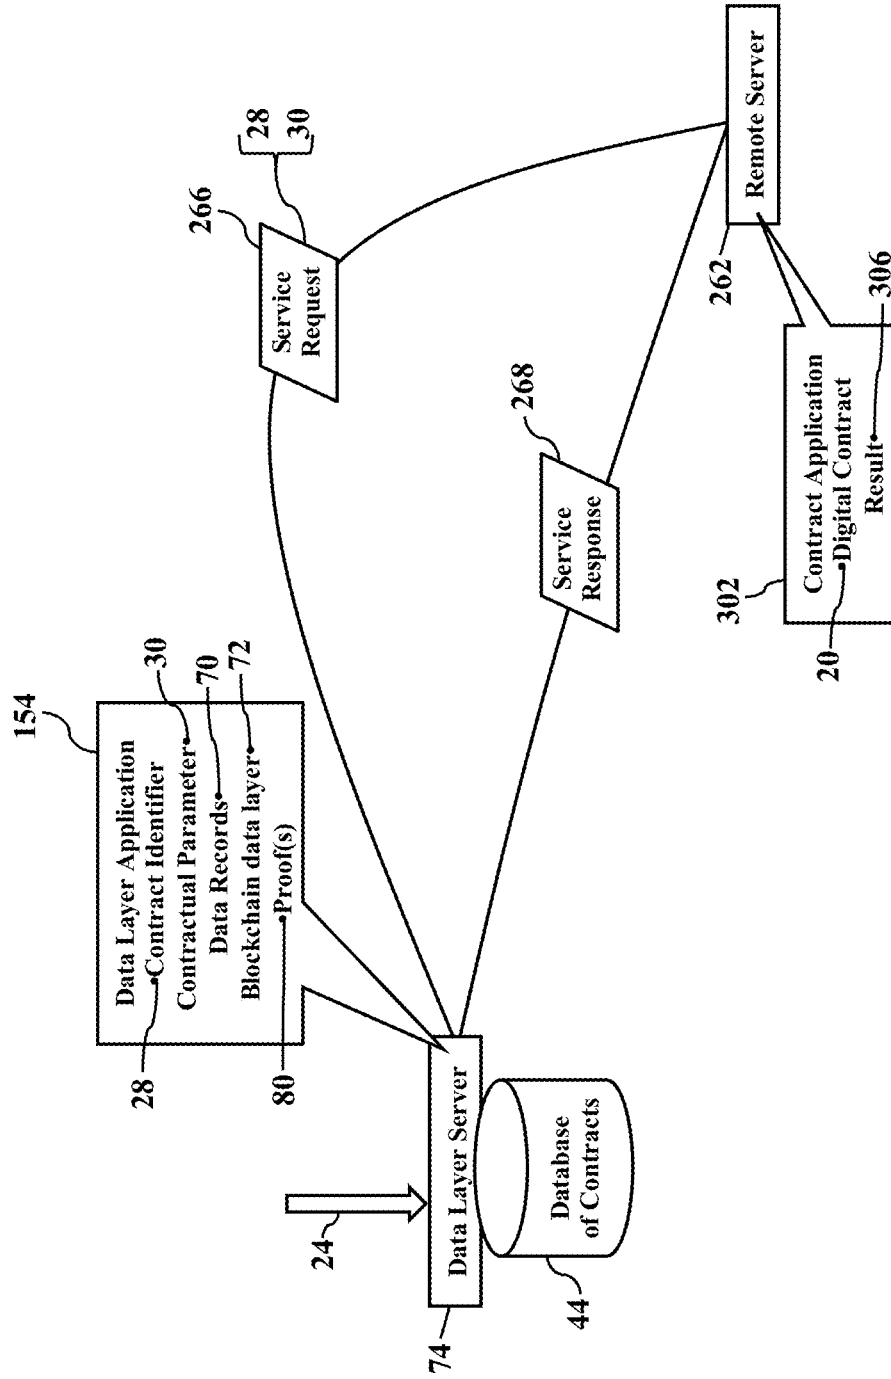


FIG. 47

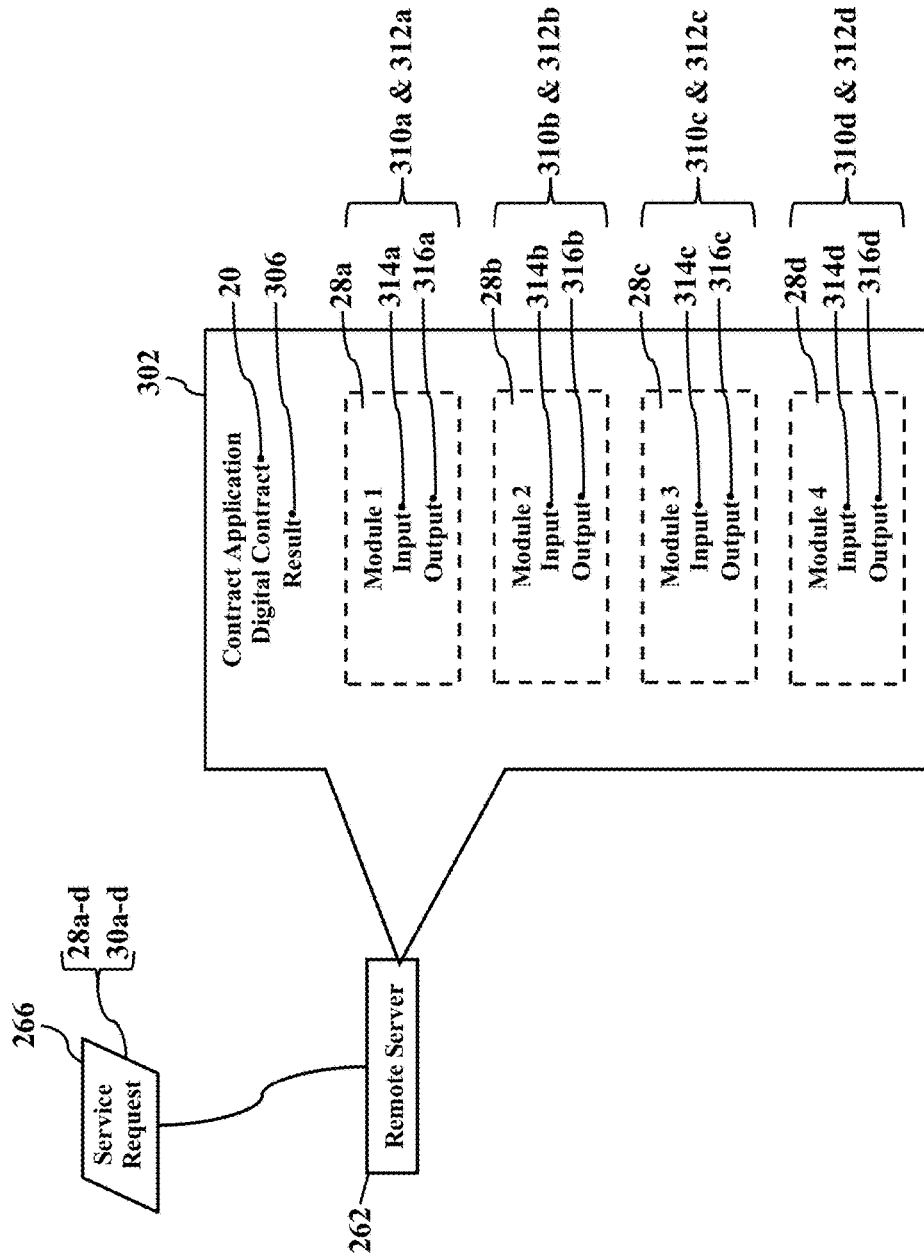


FIG. 48

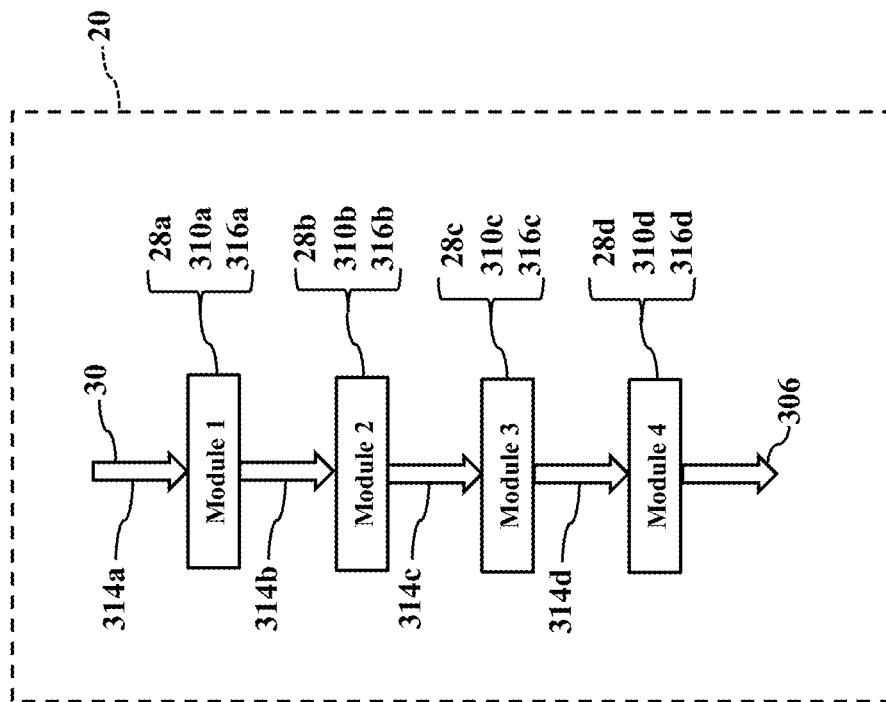


FIG. 49

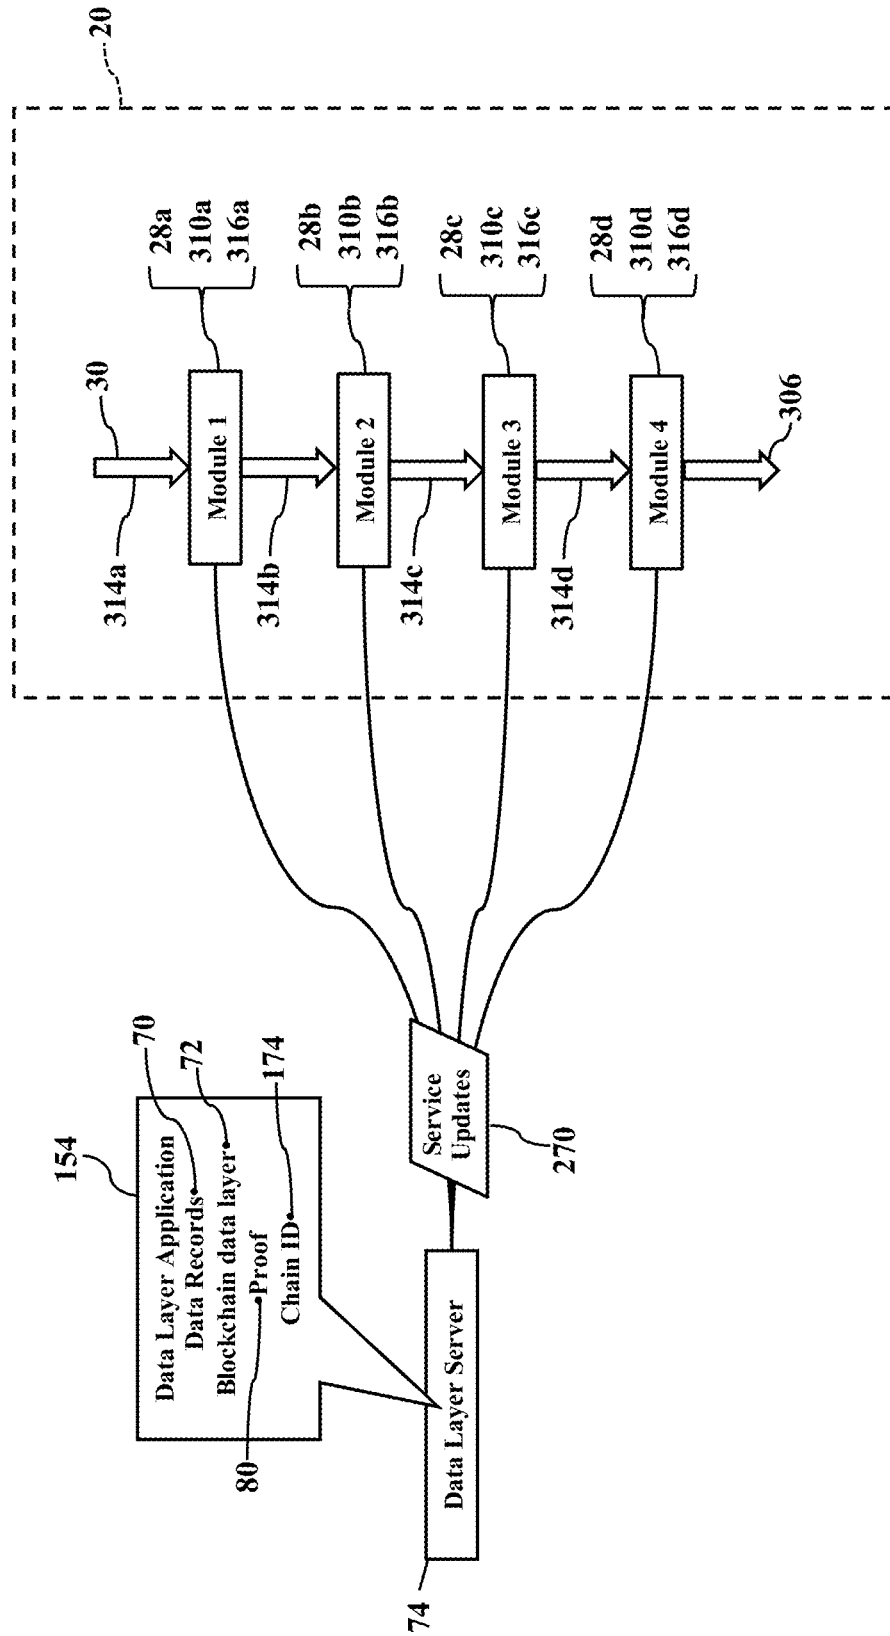


FIG. 50

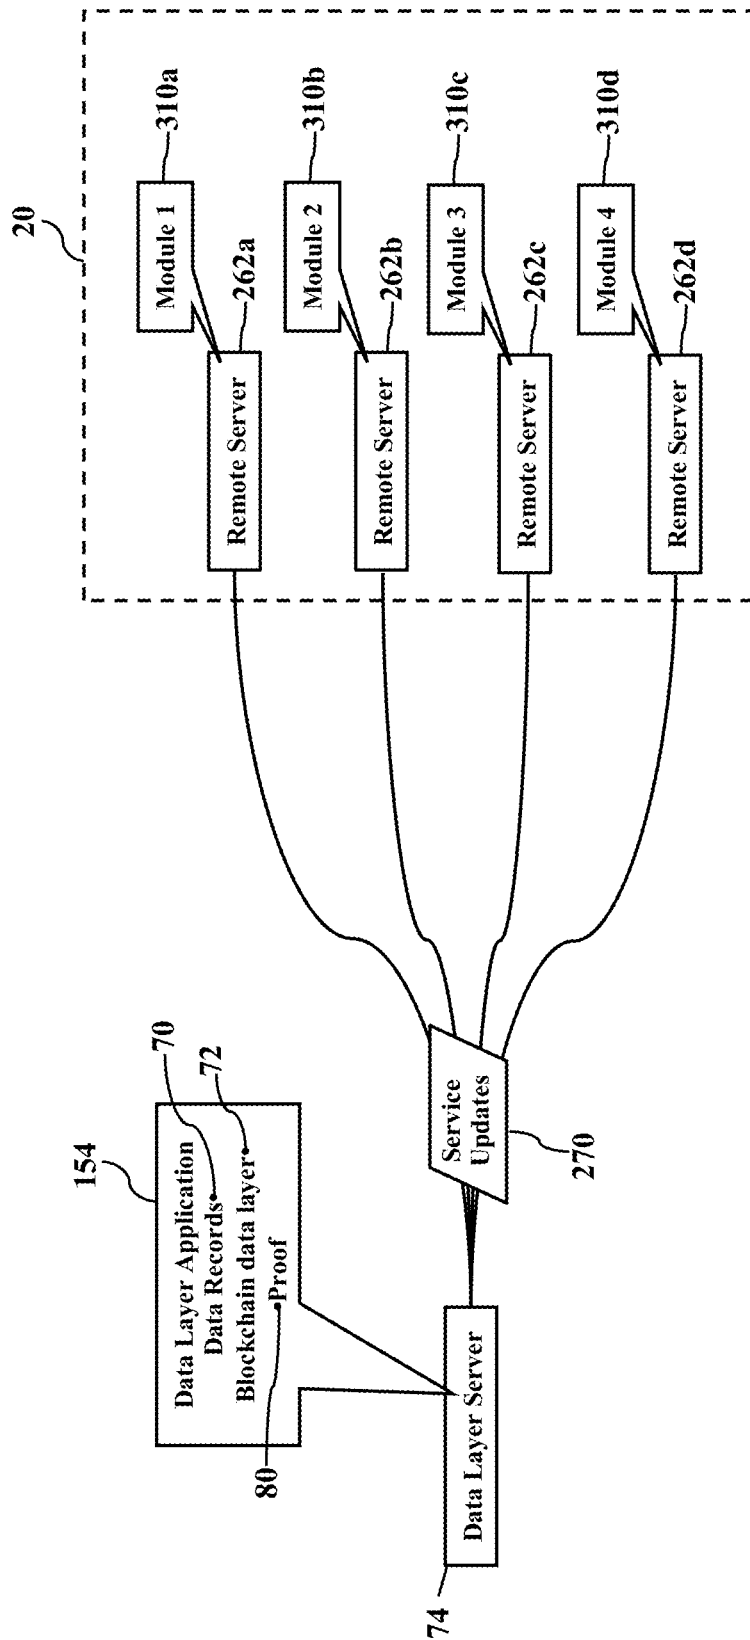


FIG. 51

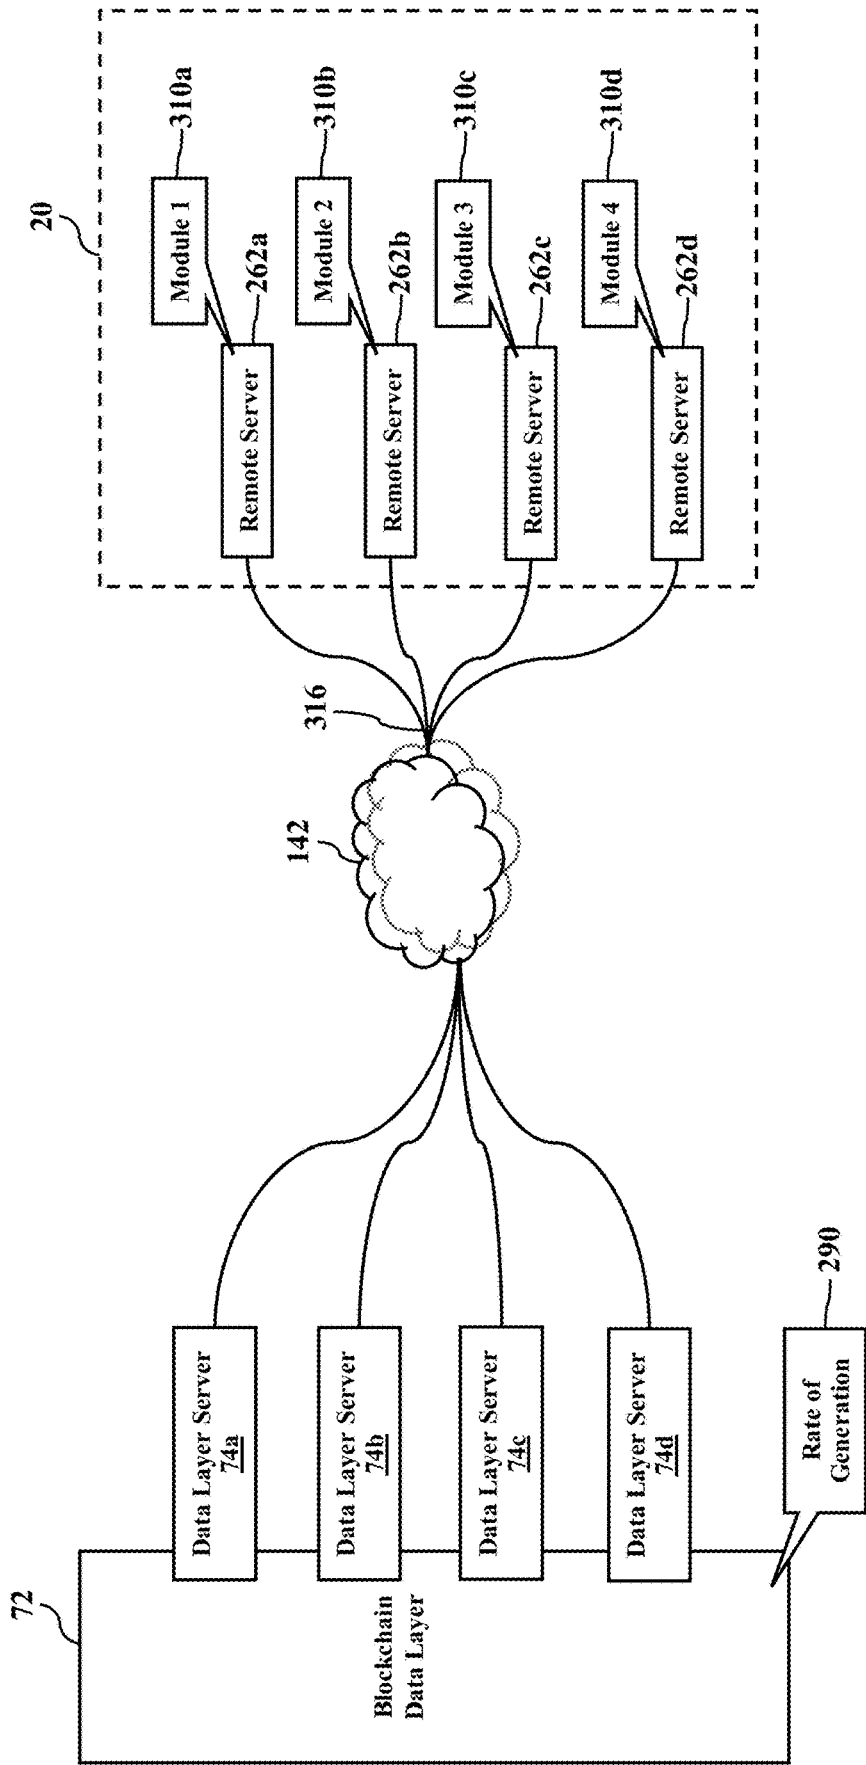


FIG. 52

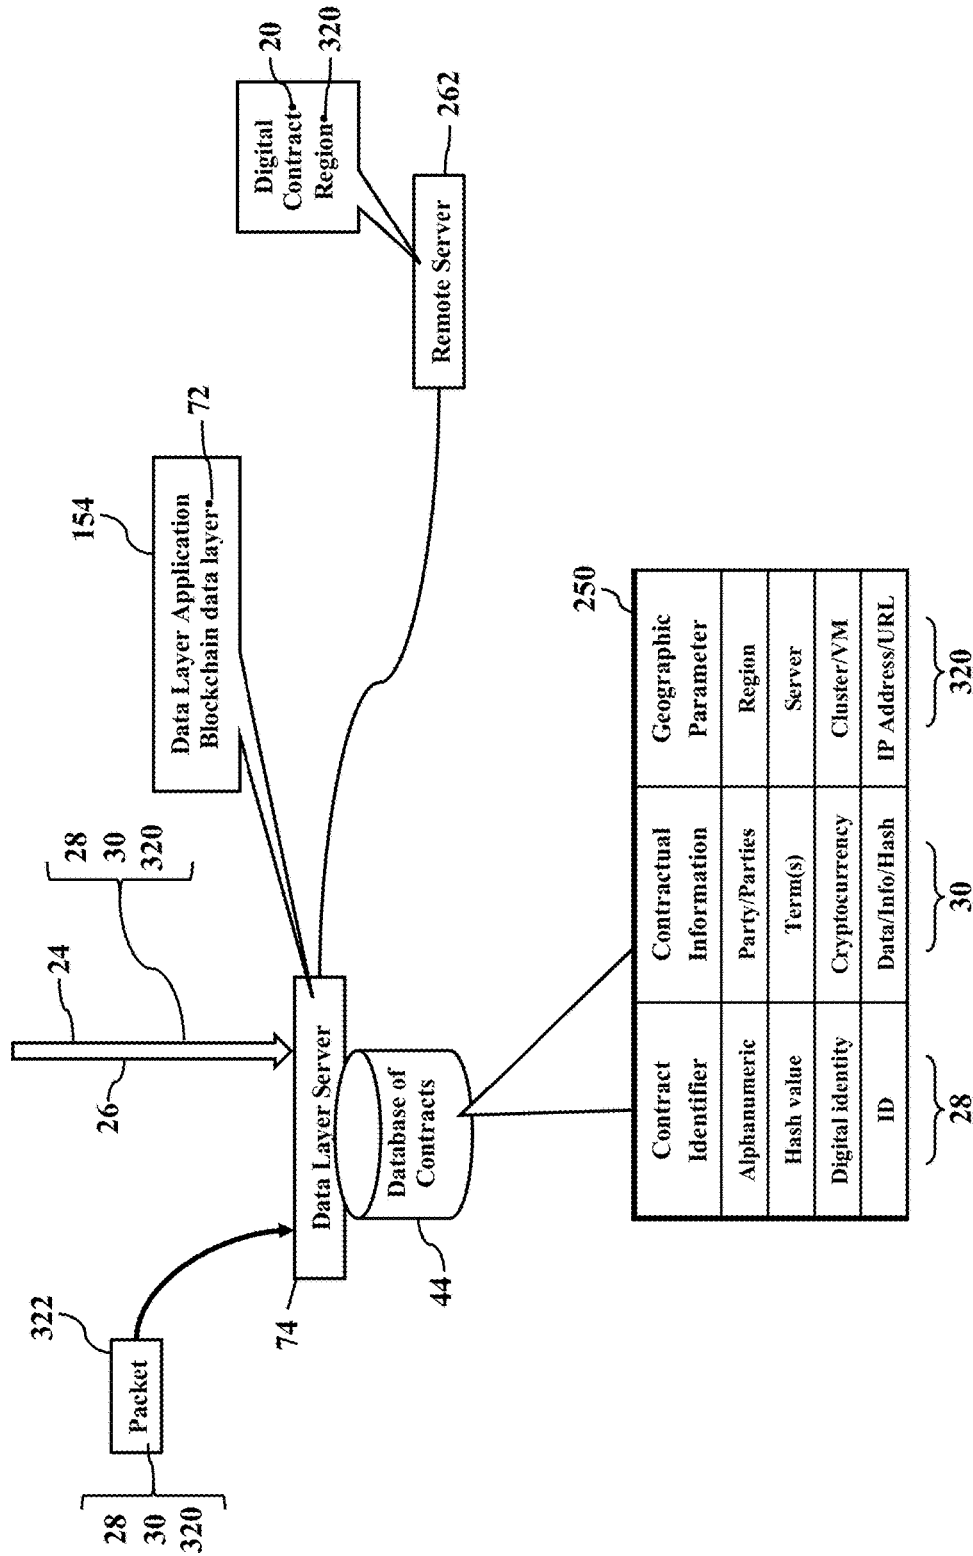


FIG. 53

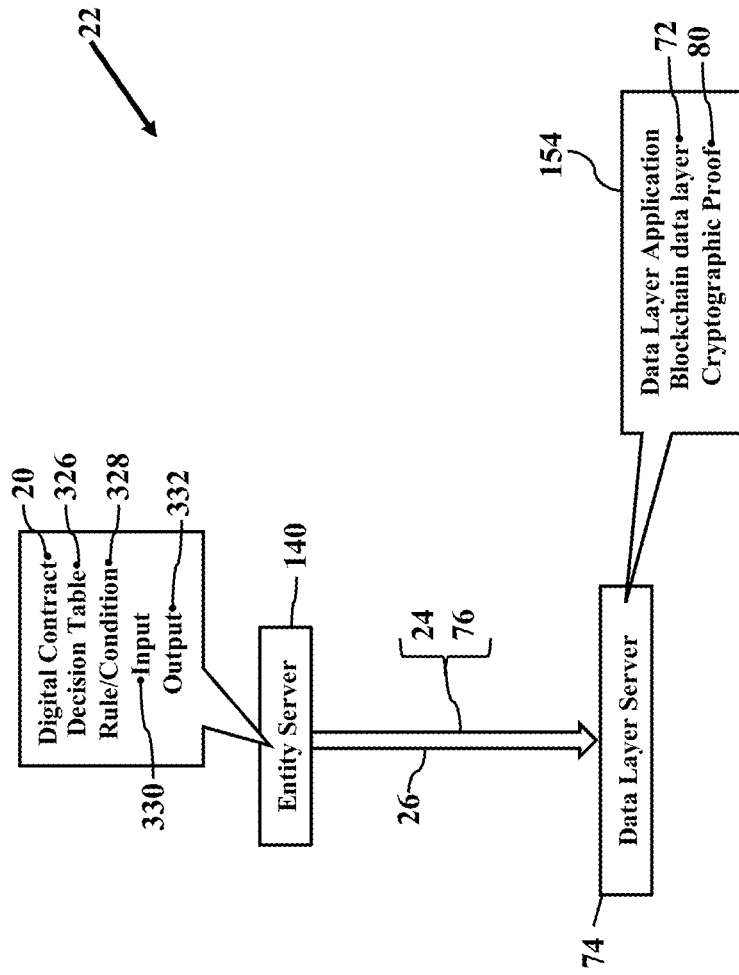


FIG. 54

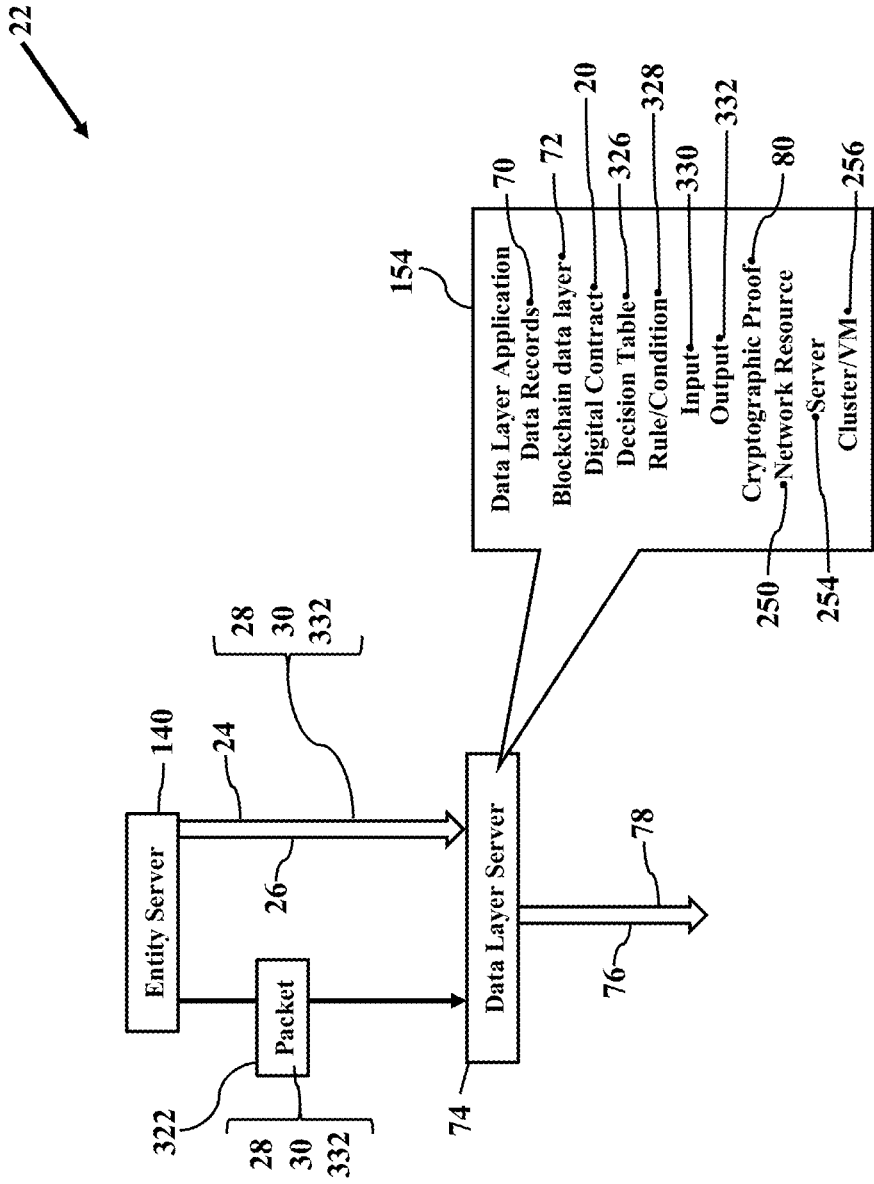


FIG. 55

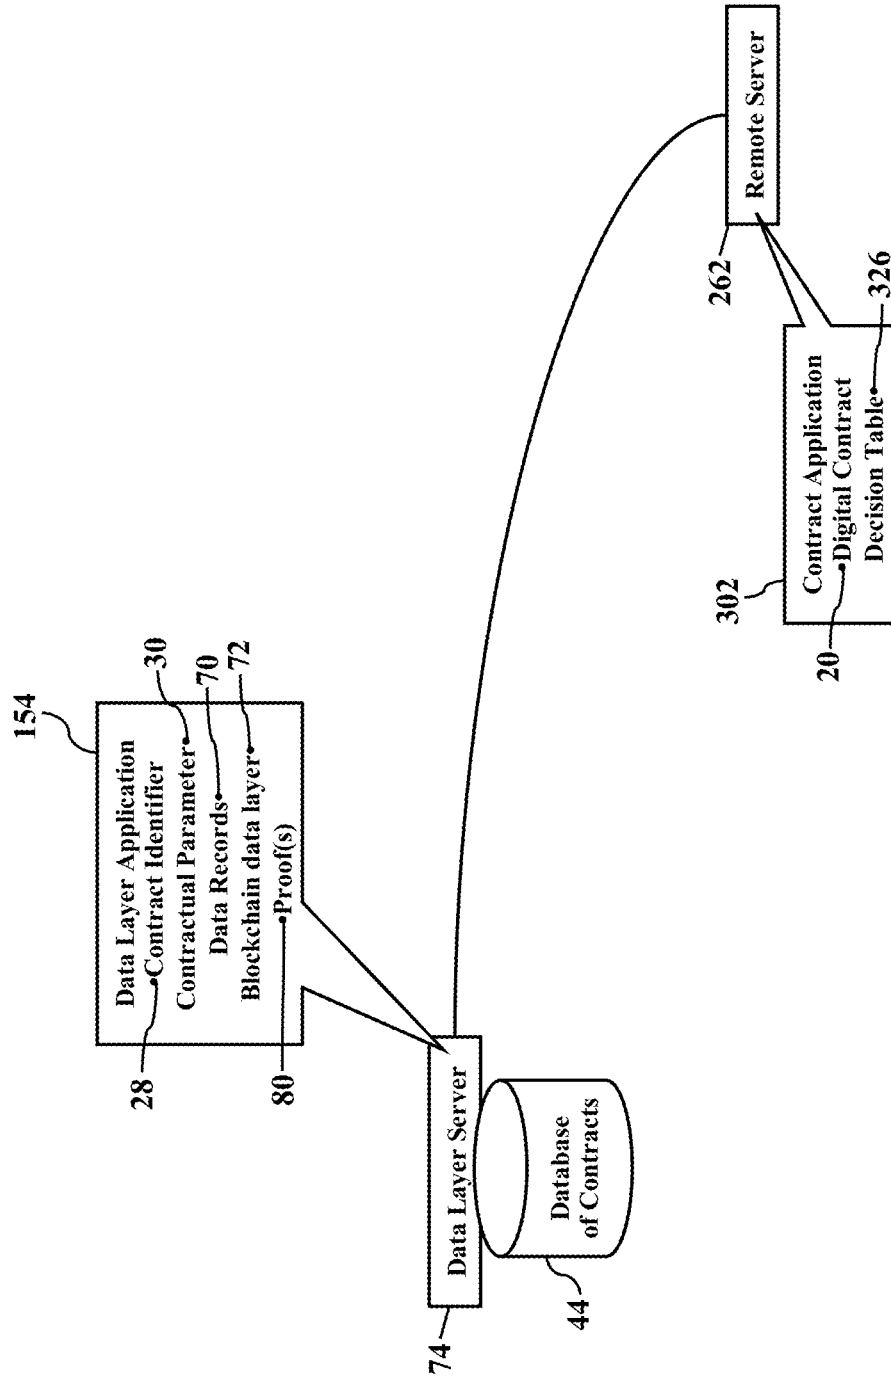


FIG. 56

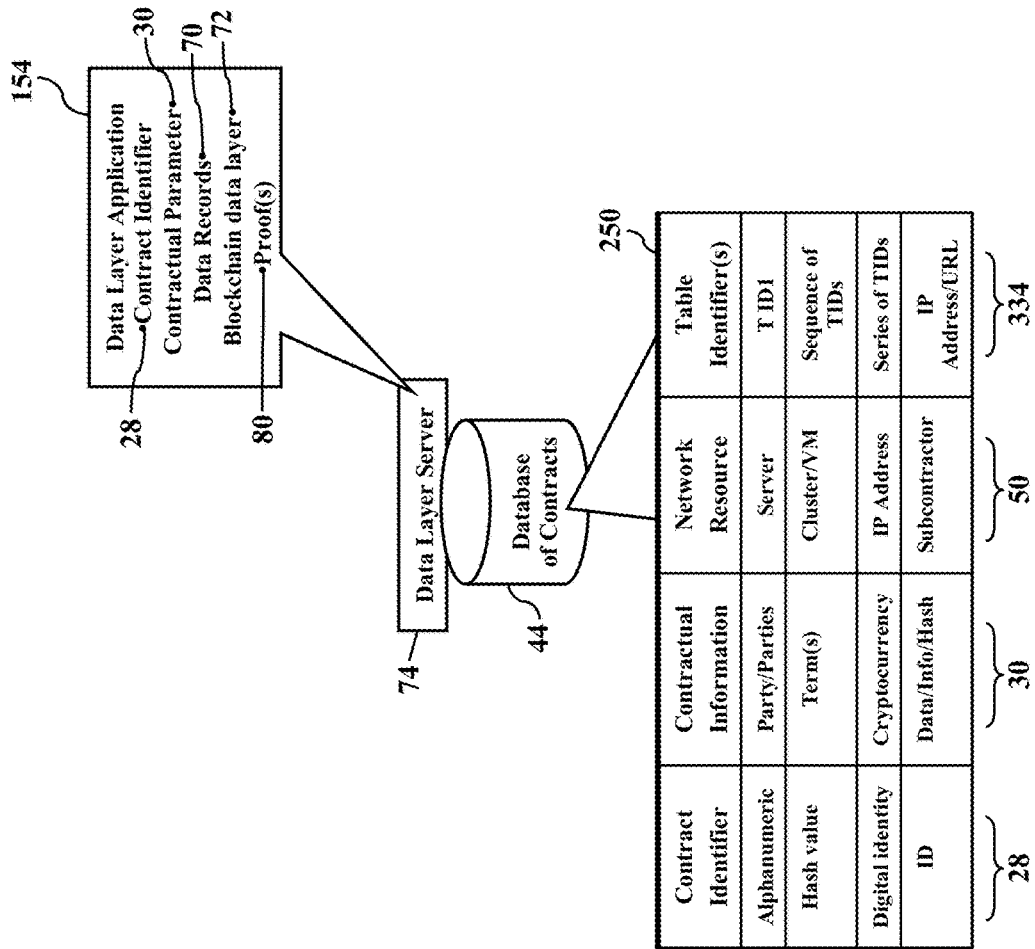


FIG. 57

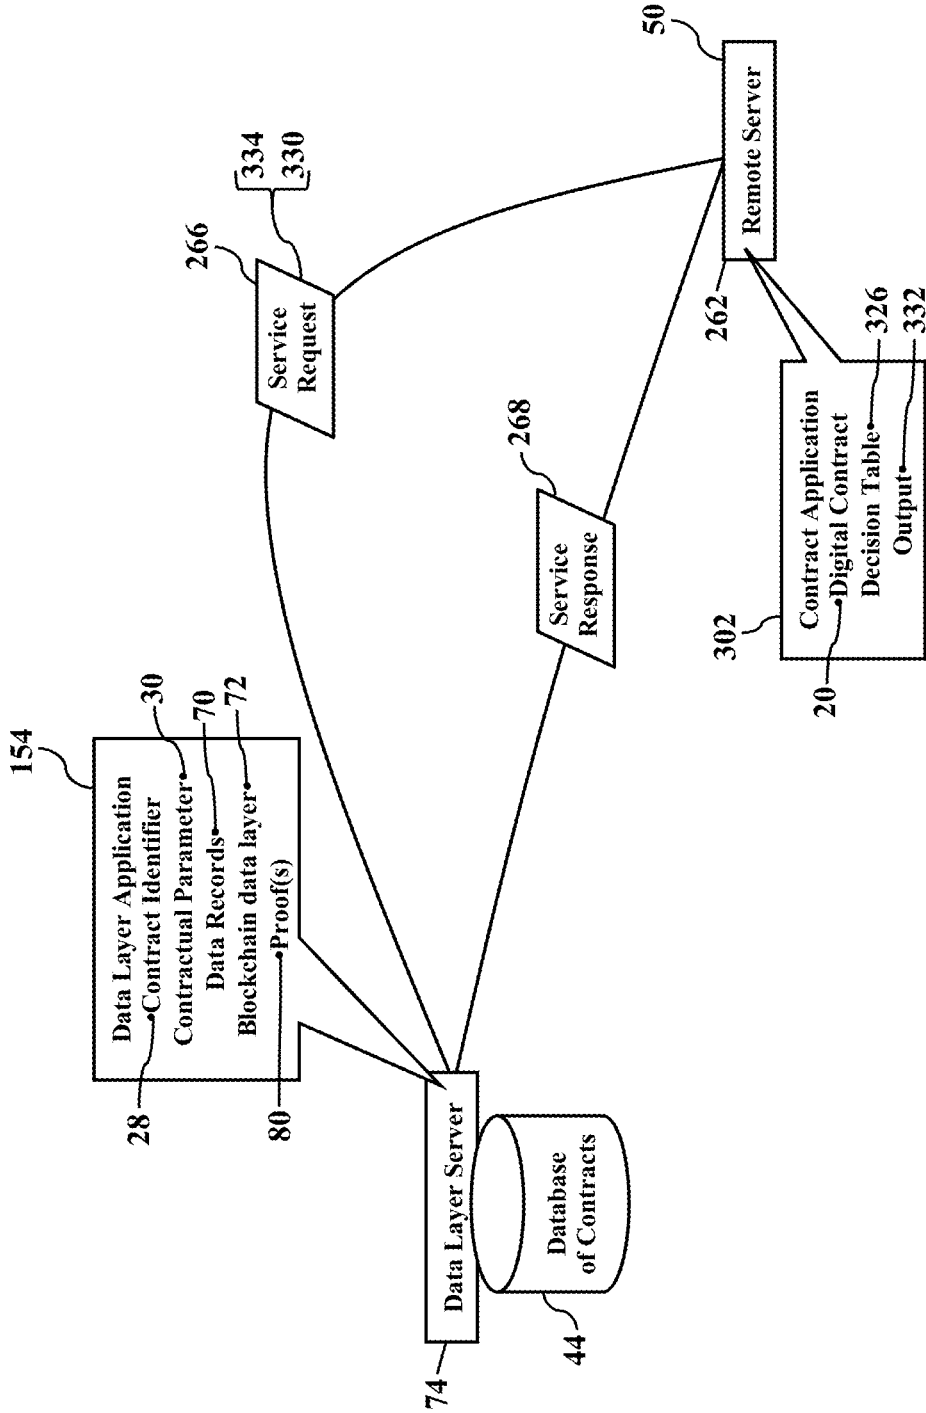


FIG. 58

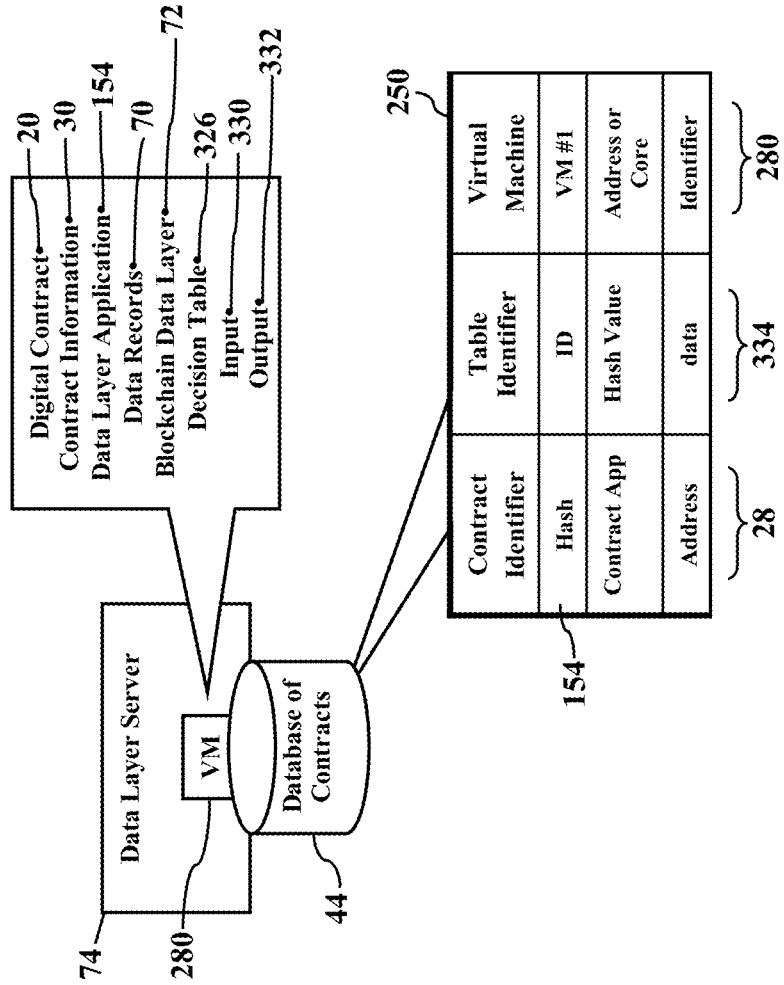


FIG. 59

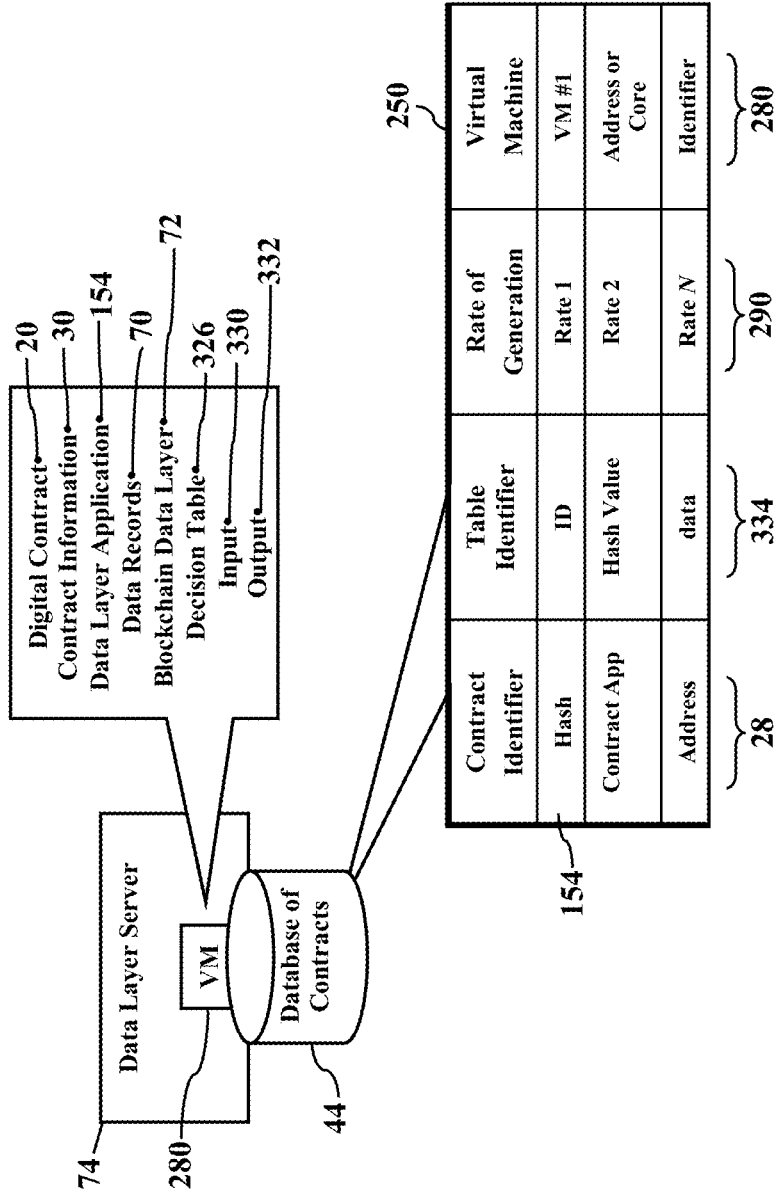


FIG. 60

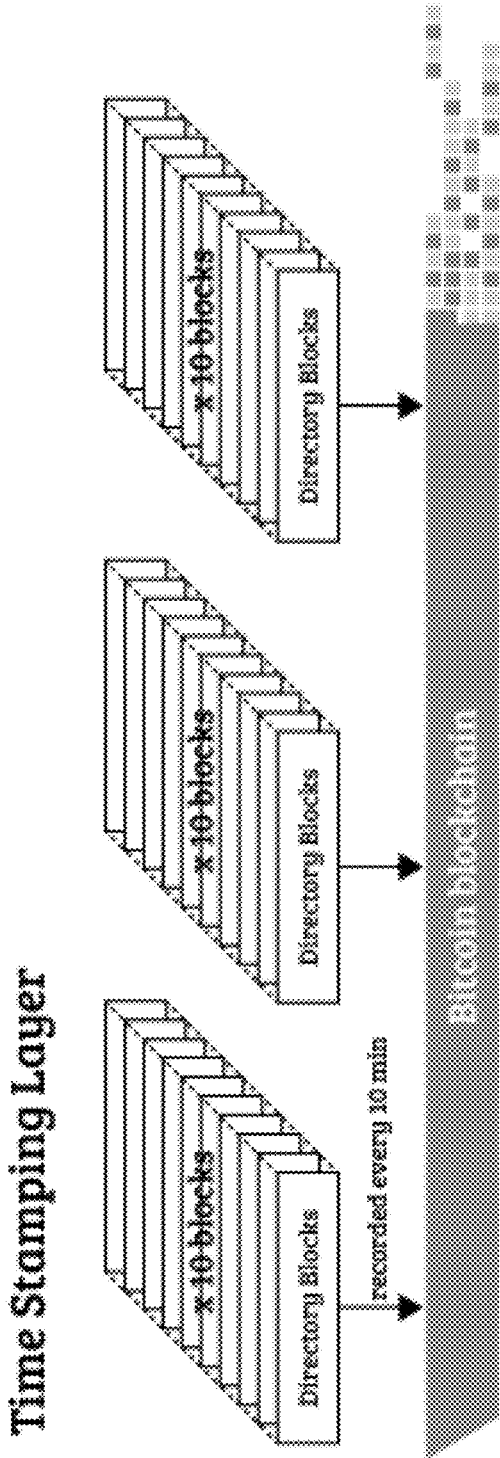


FIG. 61

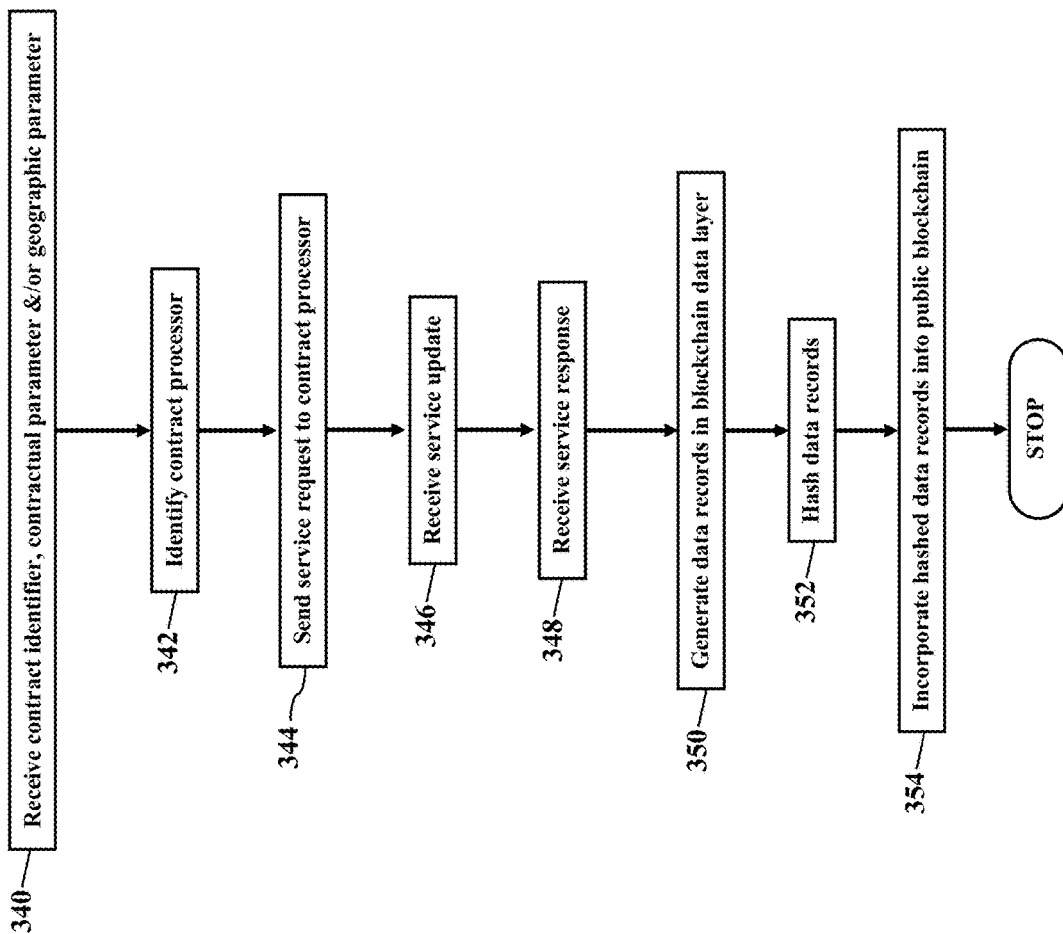


FIG. 62

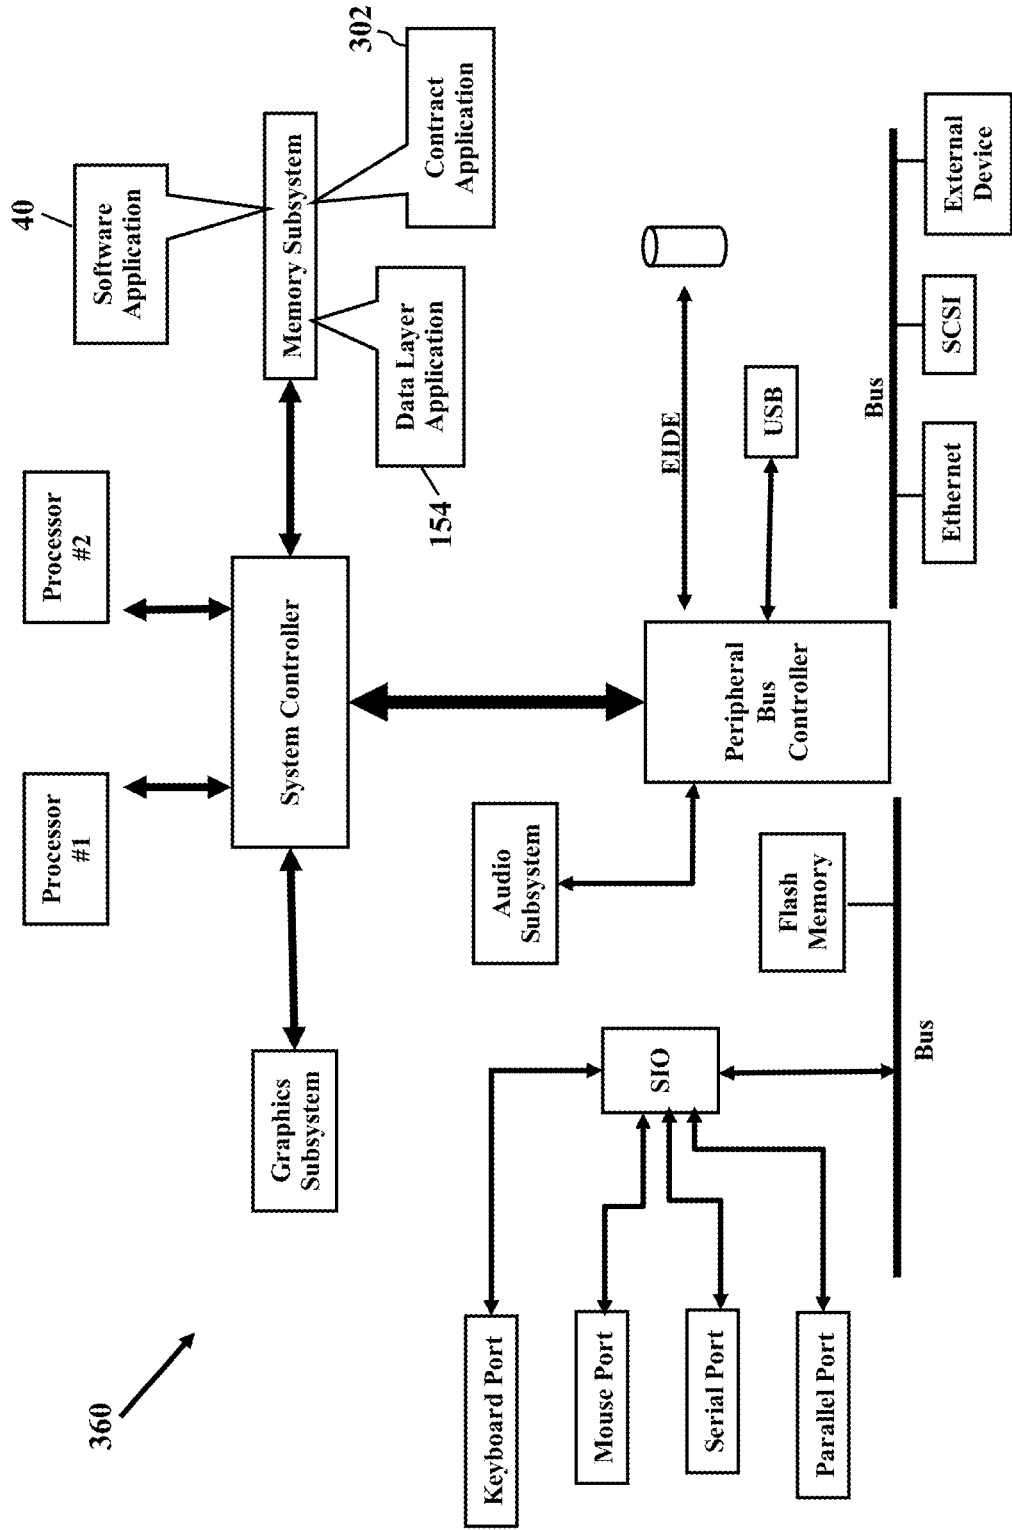
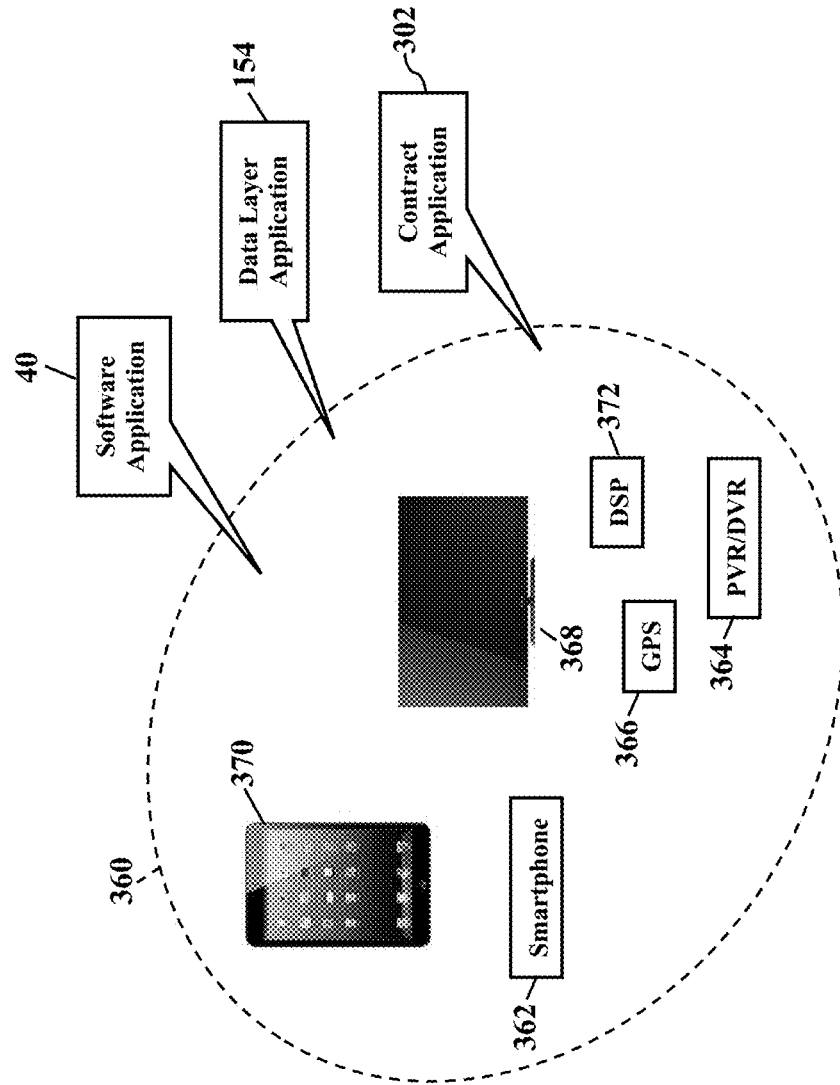


FIG. 63



PROTOCOLS IN BLOCKCHAIN ENVIRONMENTS

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This patent application is a continuation filing of U.S. application Ser. No. 16/351,597 filed Mar. 13, 2019, since issued as U.S. Patent X, which claimed domestic benefit of U.S. Provisional Application No. 62/714,909 filed Aug. 6, 2018, with both patent applications incorporated herein by reference in their entireties. This patent application relates to U.S. application Ser. No. 15/983,572 filed May 18, 2018 and incorporated herein by reference in its entirety. This patent application also relates to U.S. application Ser. No. 15/983,595 filed May 18, 2018 and incorporated herein by reference in its entirety. This patent application also relates to U.S. application Ser. No. 15/983,612 filed May 18, 2018 and incorporated herein by reference in its entirety. This patent application also relates to U.S. application Ser. No. 15/983,632 filed May 18, 2018 and incorporated herein by reference in its entirety. This patent application also relates to U.S. application Ser. No. 15/983,655 filed May 18, 2018 and incorporated herein by reference in its entirety.

BACKGROUND

[0002] In today's global economy, trust is in rare supply. This lack of trust requires the devotion of a tremendous amount of resources to audit and to verify records—reducing global efficiency, return on investment, and prosperity. Moreover, incidents such as the 2010 United States foreclosure crisis demonstrate that in addition to being inefficient, the current processes are also terribly inaccurate and prone to failure.

BRIEF DESCRIPTION OF THE SEVERAL VIEWS OF THE DRAWINGS

[0003] The features, aspects, and advantages of the exemplary embodiments are understood when the following Detailed Description is read with reference to the accompanying drawings, wherein:

[0004] FIGS. 1-8 illustrate a Factom protocol and system, according to exemplary embodiments;

[0005] FIGS. 9-21 are simplified illustrations of a digital contract in a blockchain environment, according to exemplary embodiments;

[0006] FIGS. 22-24 are more detailed illustrations of an operating environment, according to exemplary embodiments;

[0007] FIGS. 25-29 illustrate a blockchain data layer, according to exemplary embodiments;

[0008] FIGS. 30-32 further illustrate the digital contract, according to exemplary embodiments;

[0009] FIGS. 33-35 illustrate an access mechanism, according to exemplary embodiments;

[0010] FIG. 36 illustrates a public entity, according to exemplary embodiments;

[0011] FIGS. 37-40 illustrate contractual execution, according to exemplary embodiments;

[0012] FIGS. 41-42 illustrate virtual execution, according to exemplary embodiments;

[0013] FIG. 43 illustrates cryptographic affinities, according to exemplary embodiments;

[0014] FIG. 44 illustrates virtual assignments based on the blockchain data layer, according to exemplary embodiments;

[0015] FIGS. 45-51 illustrate an architectural scheme, according to exemplary embodiments;

[0016] FIG. 52 illustrates compliance scheme, according to exemplary embodiments;

[0017] FIGS. 53-59 illustrate a decisional architecture and scheme, according to exemplary embodiments;

[0018] FIG. 60 is a flowchart illustrating a method or algorithm for executing of digital contracts, according to exemplary embodiments; and

[0019] FIGS. 61-63 depict still more operating environments for additional aspects of the exemplary embodiments.

DETAILED DESCRIPTION

[0020] The exemplary embodiments will now be described more fully hereinafter with reference to the accompanying drawings. The exemplary embodiments may, however, be embodied in many different forms and should not be construed as limited to the embodiments set forth herein. These embodiments are provided so that this disclosure will be thorough and complete and will fully convey the exemplary embodiments to those of ordinary skill in the art. Moreover, all statements herein reciting embodiments, as well as specific examples thereof, are intended to encompass both structural and functional equivalents thereof. Additionally, it is intended that such equivalents include both currently known equivalents as well as equivalents developed in the future (i.e., any elements developed that perform the same function, regardless of structure).

[0021] Thus, for example, it will be appreciated by those of ordinary skill in the art that the diagrams, schematics, illustrations, and the like represent conceptual views or processes illustrating the exemplary embodiments. The functions of the various elements shown in the figures may be provided through the use of dedicated hardware as well as hardware capable of executing associated software. Those of ordinary skill in the art further understand that the exemplary hardware, software, processes, methods, and/or operating systems described herein are for illustrative purposes and, thus, are not intended to be limited to any particular named manufacturer.

[0022] As used herein, the singular forms “a,” “an,” and “the” are intended to include the plural forms as well, unless expressly stated otherwise. It will be further understood that the terms “includes,” “comprises,” “including,” and/or “comprising,” when used in this specification, specify the presence of stated features, integers, steps, operations, elements, and/or components, but do not preclude the presence or addition of one or more other features, integers, steps, operations, elements, components, and/or groups thereof. It will be understood that when an element is referred to as being “connected” or “coupled” to another element, it can be directly connected or coupled to the other element or intervening elements may be present. Furthermore, “connected” or “coupled” as used herein may include wirelessly connected or coupled. As used herein, the term “and/or” includes any and all combinations of one or more of the associated listed items.

[0023] It will also be understood that, although the terms first, second, etc. may be used herein to describe various elements, these elements should not be limited by these terms. These terms are only used to distinguish one element

from another. For example, a first device could be termed a second device, and, similarly, a second device could be termed a first device without departing from the teachings of the disclosure.

[0024] FIG. 1 illustrates a distributed, autonomous Factom protocol, according to exemplary embodiments. The Factom protocol cost effectively separates any blockchain (such as the Bitcoin blockchain) from any cryptocurrency (such as the Bitcoin cryptocurrency). The Factom protocol provides client-defined Chains of Entries, client-side validation of Entries, a distributed consensus algorithm for recording the Entries, and a blockchain anchoring approach for security.

[0025] When Satoshi Nakamoto launched the Bitcoin blockchain he revolutionized the way transactions were recorded. There had never before existed a permanent, decentralized, and trustless ledger of records. Developers have rushed to create applications built on top of this ledger. Unfortunately, they have been running into a few core constraints intrinsic to the original design tradeoffs of Bitcoin.

[0026] 1) Speed—because of the design of the decentralized, proof-of-work consensus method used by Bitcoin, difficulty requirements are adjusted to maintain roughly 10 minute confirmation times. For applications that wish greater security, multiple confirmations may be required. A common requirement is to wait for 6 confirmations, which can lead to wait times over an hour.

[0027] 2) Cost—the default transaction cost is around 0.01 mBTC (roughly \$0.003 USD in November 2014, and as much as \$80 USD per transaction at times in 2017). The exchange price of BTC has been volatile throughout its history. If the price of BTC rises, then the cost of transactions can go up. This can prove to be a serious cost barrier to applications that need to manage very large numbers of transactions. Additionally, many factors including constraints on block size and reward halving could act to increase transaction fees.

[0028] 3) Bloat—with the Bitcoin blockchain size limit of 1 MB per block, transaction throughput is capped at 7 transactions per second. Any application that wants to write and store information using the blockchain will add to the traffic. This problem has become politically charged as various parties seek to increase the block size limit and are met with resistance from those concerned about decentralization.

[0029] Factom is a protocol designed to address these three core constraints. Factom creates a protocol for Applications that provide functions and features beyond currency transactions. Factom constructs a standard, effective, and secure foundation for these Applications to run faster, cheaper, and without bloating Bitcoin.

[0030] Once the system is set up, including issuance of Factoids (i.e., the cryptocurrency of Factom) and user accounts, token value is transferred among users, Factom, and Bitcoin with the following primary interactions:

[0031] 1. Application Owner purchases Entry Credits with Factoid

[0032] 2. Application records an Entry

[0033] 3. Factom Servers create Entry Blocks and Directory Blocks

[0034] 4. Factom secures an anchor (hash of the Directory Block) onto the blockchain

[0035] Details of these and other interactions are in the upcoming sections.

[0036] The Factom protocol secures entries. Factom extends Bitcoin's feature set to record events outside of monetary transfers. Factom has a minimal ruleset for adding permanent Entries. Factom pushes most data validation tasks to the client side. The only validation Factom enforces are those required by the protocol to trade Factoids, convert Factoids to Entry Credits, and to ensure Entries are properly paid for and recorded.

[0037] Factom has a few rules regarding token incentives for running the network and for internal consistency, but Factom may or may not check the validity of statements recorded in the chains used by its users.

[0038] Bitcoin limits transactions to those moving value from a set of inputs to a set of outputs. Satisfying the script required of the inputs (generally requiring certain signatures) is enough for the system to ensure validity. This is a validation process that can be automated, so the auditing process is easy. If Factom were used, for instance, to record a deed transfer of real estate, Factom would be used to simply record the process occurred. The rules for real estate transfers are very complex. For example, a local jurisdiction may have special requirements for property if the buyer is a foreigner, farmer, or part time resident. A property might also fall into a number of categories based on location, price, or architecture. Each category could have its own rules reflecting the validation process for smart contracts. In this example, a cryptographic signature alone is insufficient to fully verify the validity of a transfer of ownership. Factom then is used to record the process occurred rather than validate transfers.

[0039] Bitcoin miners perform two primary tasks. First, they resolve double spends. Seeing two conflicting transactions that spend the same funds twice, they resolve which one is admissible. The second job miners perform (along with the other full nodes) is auditing. Since Bitcoin miners only include valid transactions, one that is included in the blockchain can be assumed to have been audited. A thin client does not need to know the full history of Bitcoin to see if value they receive has already been spent.

[0040] Factom splits the two roles that Bitcoin miners do into two tasks: 1—recording Entries in a final order and 2—auditing Entries for validity.

[0041] 1—The Factom servers accept Entries, assemble them into blocks, and fix their order. After 10 minutes, the Entry ordering is made irreversible by inserting an anchor into the Bitcoin blockchain. Factom does this by creating a hash of the data collected over the 10 minutes, then recording the hash into the blockchain.

[0042] 2—The auditing of Entries is a separate process which can be done either with or without trust. Auditing is critical, since Factom is not able to validate Entries before they are included in the Factom dataset.

[0043] With trust-based auditing, a thin client could trust a competent auditor they choose. After an Entry was entered into the system, an auditor would verify the Entry was valid. Auditors would submit their own cryptographically signed Entry. The signature would show that the Entry passed all the checks the auditor deemed was required. The audit requirements could in fact be part of a Factom Chain as well. In the real estate example from earlier, the auditor would double check the transfer conformed to local standards. The auditor would publicly attest that the transfer was valid.

[0044] Trustless auditing would be similar to Bitcoin. If a system is internally consistent with a mathematical definition of validity like Bitcoin, it can be audited programmatically. If the rules for transfer were able to be audited by a computer, then an Application could download the relevant data and run the audit itself. The application would build an awareness of the system state as it downloaded, verified, and decided which Entries were valid or not.

[0045] Mastercoin, Counterparty, and Colored Coins have a similar trust model. These are all client-side validated protocols, meaning transactions are embedded into the Bitcoin blockchain. Bitcoin miners do not audit them for validity; therefore, invalid transactions designed to look like transactions on these protocols can be inserted into the blockchain. Clients that support one of these protocols scan through the blockchain and find potential transactions, check them for validity, and build an interpretation of where the control of these assets lie (usually a Bitcoin address). It is up to the clients to do their own auditing under these protocols.

[0046] Moving any of these client-side validated protocols under Factom would be a matter of defining a transaction per the protocol and establishing a Chain to hold the transactions. The transaction protocols wouldn't be much different under Factom than under Bitcoin, except where Factom allows an easy expression of the information needed instead of having to encode it in some special way into a Bitcoin transaction.

[0047] Bitcoin, land registries, and many other systems need to solve a fundamental problem: proving a negative. They prove some "thing" has been transferred to one person, and prove that thing hasn't been transferred to someone else. While proof of the negative is impossible in an unbounded system, it is quite possible in a bounded system. Blockchain based cryptocurrencies solve this problem by limiting the places where transactions can be found. Bitcoin transactions can only be found in the Bitcoin blockchain. If a relevant transaction is not found in the blockchain, it is defined from the Bitcoin protocol perspective not to exist and thus the BTC hasn't been sent twice (double spent).

[0048] Certain land ownership recording systems are similar. Assume a system where land transfer is recorded in a governmental registry and where the legal system is set up so that unrecorded transfers are assumed invalid (sans litigation). If an individual wanted to check if a title is clear (i.e., that no one else claims the land) the answer would be in the governmental registry. The individual using the government records could prove the negative; the land wasn't owned by a third party. Where registration of title is not required, the governmental registry could only attest to what has been registered. A private transfer might very well exist that invalidates the understanding of the registry.

[0049] In both of the above cases, the negative can be proven within a context. With Mastercoin the case is very strong. With a land registry, it is limited to the context of the Registry, which may be open to challenge. The real world is messy, and Factom is designed to accommodate not just the precision of digital assets, but the real world's sometimes messy reality.

[0050] In Factom, there is a hierarchy of data categorization. Factom only records Entries in Chains; the various user-defined Chains have no dependencies that Factom enforces at the protocol level. This differs from Bitcoin, where every transaction is potentially a double-spend, and so it must be validated. By organizing Entries into Chains,

Factom allows Applications to have smaller search spaces than if all Factom data were combined together into one ledger.

[0051] If Factom were to be used to manage land transfers, an Application using a Chain to record such registries could safely ignore Entries in the other Chains, such as those used to maintain security camera logs. Were a governmental court ruling to change a land registration, the relevant Chain would be updated to reflect the ruling. The history would not be lost, and where such changes are actually invalid from a legal or other perspective, the record cannot be altered to hide the order of events in Factom.

[0052] Factom may or may not validate Entries; Entries are instead validated client-side by users and Applications. As long as an Application understands and knows the rules a Chain should follow, then the existence of invalid Entries doesn't cause unreasonable disruption. Entries in a Chain that do not follow the rules can be disregarded by the Application.

[0053] Users can use any set of rules for their Chains, and any convention to communicate their rules to the users of their Chains. The first Entry in a Chain can hold a set of rules, a hash of an audit program, etc. These rules then can be understood by Applications running against Factom to ignore invalid Entries client-side.

[0054] An enforced sequence can be specified. Entries that do not meet the requirements of the specified enforced sequence will be rejected. However, Entries that might be rejected by the rules or the audit program will still be recorded. Users of such chains will need to run the audit program to validate a chain sequence of this type. The Factom servers will not validate rules using the audit program.

[0055] Validation in the Applications (in combination with user-defined Chains) provides a number of advantages for Applications written on top of Factom:

[0056] 1) Applications can put into Factom whatever Entries make sense for their application. So, a list of hashes to validate a list of account statements can be recorded as easily as exchanges of an asset.

[0057] 2) Rule execution is very efficient. Where the distributed network must execute your validation rules, then validation requires all nodes to do all validation. Client-side validation only requires the systems that care about those rules to run them. Factom allows a Chain to define its rules in whatever language the designers choose, to run on whatever platform they choose, and to use any external data. None of these decisions on the part of one Application has any impact on another Application.

[0058] 3) Factom Servers have little knowledge about the Entries being recorded. We use a commitment scheme to limit knowledge, where the commitment to record an Entry is made prior to revealing what the Entry is. This makes Factom's role in recording Entries very simple, and makes individual server processes public. Factom servers accept information from the network of full nodes, and their decisions and behavior are in view. Failure to perform can be audited both from the network outside Factom, and within Factom. It is easy to independently verify that a Factom server is fulfilling its Entry-recording responsibility; Factom can't hide potentially errant behavior.

[0059] 4) Recording speeds can be very fast, since the number of checks made by the Factom servers are minimal.

[0060] 5) Proofs against any particular Chain in Factom do not require knowledge of any other Chains. Users then only need the sections of Factom they are using and can ignore the rest.

[0061] At its heart, Factom is a decentralized way to collect, package, and secure data into the Bitcoin blockchain. Factom accomplishes this with a network of Authority servers. Authority Servers are the set of Federated Servers and Audit Servers which share responsibility for different aspects of the system. The Federated Servers actually acknowledge and order entries and transactions in Factom, and Audit Servers duplicate and audit the work done by the Federated Servers and are always ready to replace a Federated Server that might go offline.

[0062] The design ensures decentralization. No single server is ever in control of the whole system, but only a part of the system. All servers verify and double check the work of all other servers. And no server is permanently in control of any part of the system; the responsibility for each part of Factom cycles among the Federated Servers each minute, and the role of being a Federated Server or an Audit Server shifts among the servers in the Authority Set (the set of all Authority Servers).

[0063] The Federated servers take a very active role in running the protocol. The Federated servers each take responsibility for a subsection of the user Chains at the beginning of the creation of a Directory Block. The process works like this:

[0064] 1. All servers reset their process lists to empty.

[0065] 2. The user submits an Entry Payment using a public key associated with Entry Credits

[0066] 3. Based on the public key used to pay for the Entry, one of the servers accepts the payment.

[0067] 4. That server broadcasts the acceptance of the payment.

[0068] 5. The user sees the acceptance and submits the Entry.

[0069] 6. Based on the ChainID of the Entry, one of the servers adds the Entry to its process list, and adds the Entry to the appropriate Entry Block for that ChainID (creating one if this is the first Entry for that Entry Block).

[0070] 7. The server broadcasts an Entry confirmation, containing the process list index of the Entry, the hash of the Entry (linked to the payment), and the serial hash so far of the server's process list.

[0071] 8. All the other servers update their view of the server's process list, validate the list, and update their view of the Entry Block for that ChainID.

[0072] 9. As long as the user can validate the relevant process list holds their Entry, then they have a fair level of assurance it will be successfully entered into Factom.

[0073] 10. At the end of the minute, each server confirms the end of their section of the process list. The end of the minute is marked in the process list, and the responsibility for particular chains shifts around the authority set.

[0074] 11. At the end of the 10th minute, a Directory Block is constructed from all the Entry Blocks defined

by the process lists built by all the servers. So, each server has all Entry Blocks, all Directory Blocks, and all Entries.

[0075] 12. A deterministic method (that can be computed by all nodes in protocol) will shift responsibility for particular ChainIDs among the servers for the next round.

[0076] 13. At the completion of the Directory Block, the Merkle root of the Directory block is placed in a Bitcoin transaction and submitted to the Bitcoin network for eventual confirmation.

[0077] 14. Repeat. (Go back to 1)

[0078] The Federated servers for their minute are constructing a process list for the Chains for which they are responsible, as well as constructing the Entry Blocks that will be used to create the Directory Block at the end of the 10 minutes. The process list is important for broadcasting decisions made by a server to the rest of the network.

[0079] The servers in the authority set are re-ranked on a regular, scheduled basis. The ranking is a function of support by the standing parties, who must create a profile Chain in Factom. The profile contains any number of signed public address Entries. The weight of a standing party's support is determined by various public addresses and entries in their profile. The function computing the weight of a standing party uses a combination of many factors. Such weights may be organized in categories to further distribute influence. Factors that determine an identity's weight include factors that can be measured from the protocol, and audited by the protocol. Examples of factors that might be used to calculate weight include:

[0080] Weighted Number of Entry Credits purchased.

[0081] Weighted Number of Entries used.

[0082] Tokens "staked" to a profile Chain, and not moved or transferred.

[0083] Tokens used to build infrastructure, support the protocol, provide services

[0084] Providing guidance and facilitating the operation of the protocol.

[0085] Support may be specified by the Standing parties at any time. At regular intervals, the support of all the servers in the Authority set will be evaluated, and the membership of the authority set adjusted. The same mechanism can be used to measure support in the protocol for decisions about the protocol.

[0086] To maintain a position in the authority set, servers must continually demonstrate the ability to maintain their ability to monitor and keep up with the operation of the protocol. The Federated Servers do this by simply doing their job and syncing with the end of minute operations with all other Federated Servers. Performance in the protocol's ecosystem may also factor into decisions to support or not support an authority node. Audit servers may have to issue a heartbeat message, that can be monitored by the network. Other solutions are possible.

[0087] Managing timeouts and monitoring heartbeats will be done according to the needs and load on the protocol.

Factom System Overview

[0088] FIG. 3 illustrates the Factom protocol as a set of layered data structures, according to exemplary embodiments. Factom is constructed of a hierarchical set of blocks, with the highest being Directory Blocks. They constitute a micro-chain, consisting primarily of compact references. To

keep the size small, each reference in the Directory Block is just a hash of the Entry Block plus its ChainID. These Entry Blocks have references which point to all the Entries with a particular ChainID which arrived during a time period. The Entry Block for a Chain ID is also part of a micro-chain. The bulk of the data in Factom is at the leaves, the Entries themselves. These hierarchical data structures are rendered unchangeable by Bitcoin's hashpower. They can be conceptualized as different layers. The layers and concepts in the Factom system are:

[0089] 1) Directory Layer—Organizes the Merkle Roots of Entry Blocks

[0090] 2) Entry Block Layer—Organizes references to Entries

[0091] 3) Entries—Contains an Application's raw data or a hash of its private data

[0092] 4) Chains—Grouping of Entries specific to an Application

Directory Layer: How the Directory Layer Organizes Merkle Roots

[0093] The Directory layer is the first level of hierarchy in the Factom system. It defines which Entry ChainIDs have been updated during the time period covered by a Directory Block. (ChainIDs identify the user's Chain of Entries; the generation of the ChainID is discussed later.) It mainly consists of a list pairing a ChainID and the Merkle root of the Entry Block containing data for that ChainID.

[0094] Each Entry Block referenced in the Directory Block takes up 64 bytes (two 32 byte hashes, the ChainID and the Merkle root of the Entry Block). A million such Entries would result in a set of Directory Blocks roughly 64 MB in size. If the average Entry Block had 5 Entries, 64 MB of Directory Blocks would provide the high level management of 5 million distinct Entries. Note that the exact implementation of Directory blocks may vary as we build for greater scale in the future.

[0095] If an Application only has the Directory Blocks, it can find Entry Blocks it is interested in without downloading every Entry Block. An individual Application would only be interested in a small subset of ChainIDs being tracked by Factom. This greatly limits the amount of bandwidth an individual client would need to use with Factom as their system of record. For example, an Application monitoring real estate transfers could safely ignore video camera security logs.

[0096] Factom servers collect Merkle roots of Entry Blocks and package them into a Directory Block. Directory Block the Merkle roots are recorded into the Bitcoin blockchain. This allows the most minimum expansion of the blockchain, and still allows the ledger to be secured by the Bitcoin hash power. The process of adding the Merkle root into the Bitcoin blockchain we referred to as "anchoring". See the section "Appendix: Timestamping into Bitcoin" for further details.

[0097] Data entered into Directory Blocks is the most expensive, from a bandwidth and storage perspective. All users of Factom wishing to find data in their Chains need the full set of Directory Blocks starting from when their Chain began.

[0098] Activities that increase the Directory Block size include the creation and first update of individual Chains. These activities externalize costs of Applications attempting finer-grained organization.

[0099] The Applications must be required to expend more Entry Credits than a simple Entry would necessitate to discourage bloating the Directory Blocks.

Entry Block Layer: How the Entry Block Layer Organizes Hashes and Data

[0100] FIG. 4 illustrates entry blocks of the Factom protocol, according to exemplary embodiments. Entry Blocks are the second level of hierarchy in the system. Individual Applications will pay attention to various ChainIDs. Entry Blocks are the place where an Application looking for Entries can expand its search from a ChainID to discover all possibly relevant Entries.

[0101] There is one Entry Block for each updated ChainID per Directory Block. The Entry Blocks contain hashes of individual Entries. The hashes of Entries both prove the existence of the data and give a key to find the Entries in a Distributed Hash Table (DHT) network. (See the below explanation of the "Factom Peer-to-Peer Network" for more detail.)

[0102] The Entry Blocks encompass the full extent of possible Entries related to a ChainID. If an Entry is not referred to in an Entry Block, it can be assumed not to exist. This allows an Application to prove a negative, as described in the section Security and Proofs.

[0103] The Entry Block intentionally does not contain the Entries themselves. This allows the Entry Blocks to be much smaller than if all the data was grouped together. Separating the Entries from the Entry Blocks also allows for easier auditing of auditors. An auditor can post Entries in a separate chain that approves or rejects Entries in a common chain. The audit can add reasons for rejection in its Entry. If an Application trusts the auditor, they can cross reference that the auditor has approved or rejected every Entry, without knowing what the Entry is. The Application would then only attempt to download the Entries which passed the audit. Multiple auditors could reference the same Entries, and the Entries would only exist once on the Distributed Hash Table (DHT). Entries are expected to be significantly larger than the mere 32 bytes a hash takes up. Lists of things to ignore do not have to have the full object being ignored for an Application to know to ignore it. The exact implementation of entry blocks may vary in the future in response to identified improvements possible in the protocol.

[0104] An Entry detailing the specifics of a land transfer would be entered into a Chain where land transfers of that type are expected to be found. One or more auditors could then reference the hashes of land transfer in their own Chains, adding cryptographic signatures indicating a pass or fail. The land transfer document would only need to be stored once, and it would be referenced by multiple different Chains.

[0105] FIG. 5 illustrates how entries are created, according to exemplary embodiments. Entries are constructed by users and submitted to Factom. By hashing or encoding information, the user can ensure the privacy of Entries. The Entries can instead be plain text if encoding or obscuring the data isn't necessary. By recording a hash of a document, Factom can provide basic proof of publication. Presenting the document at a later time allows one to create its hash, and compare it to the hash recorded in the past.

[0106] There is lots of flexibility in the data that is accepted. It can be something short like a hyperlink. It could also be larger, but not too large, since fees limit the size of

the data accepted. This is similar to Bitcoin. Large 100 kB+Bitcoin transactions are possible, but would need to pay a proportionately greater transaction fee. This size, while gigantic in Bitcoin, would be moderately sized for Factom. Since every Bitcoin full node needs the entire blockchain to fully validate, it needs to stay small. In Factom, only the highest level Directory Blocks are required to fully validate a Chain. If someone is not specifically interested in a Chain's data, they would not download it.

[0107] Take a simple example of an uneditable Twitter-like system. A celebrity would craft an Entry as a piece of text. They would then sign it with a private key to show it came from them. Followers of the celebrity would find which Chain they publish in and would monitor it for updates. Any new signed Entries would be recognized by follower's Application software as a tweet. Others could tweet at the celebrity by adding Entries to the celebrity's Chain.

[0108] FIG. 6 illustrates the complete Factom protocol and system, according to exemplary embodiments. Chains in Factom are sequences of Entries that reflect the events relevant to an Application. These sequences are at the heart of Bitcoin 2.0. Chains document these event sequences and provide an audit trail recording that an event sequence occurred. With the addition of cryptographic signatures, those events would be proof they originated from a known source.

[0109] Chains are logical interpretations of data placed inside Directory Blocks and Entry Blocks. The Directory Blocks indicate which Chains are updated, and the Entry Blocks indicate which Entries have been added to the Chain. This is somewhat analogous to how Bitcoin full clients maintain a local idea of the UTXO (Unspent Transaction Output) set. The UTXO set is not (currently) in the blockchain itself, but is interpreted by the full client.

The Factom Peer-to-Peer Network

[0110] Factom will have a peer-to-peer (P2P) network which accomplishes two goals: communication and data preservation.

Factom Peer-to-Peer Communications

[0111] Factom will have a P2P network very similar to Bitcoin's. It will consist of full nodes which have all the Factom data. The full nodes create a gossip network which will flood fill valid data throughout the network. The Authority servers would be full nodes, but not all full nodes are Authority servers. This is very much like Bitcoin, where miners are full nodes, but not all full nodes are miners. This will limit the ability to DDOS the Authority servers individually. They can connect anywhere inside the network to acquire the data needed to build the data structures.

[0112] As the servers are coming to consensus and disseminate their signed data, they would publish the data over the P2P network. The P2P flood filling also limits the ability of Authority servers to censor based on IP addresses, since valid traffic is mixed together by the nodes they connect to. It also helps to prevent censorship, since all servers can see the Entries which should be included in the Entry Blocks. Outside organizations campaigning to become Authority servers have an incentive to bring bad behavior to light, so they can gain support and move up into the set of Authority Servers.

Data Preservation and Dissemination

[0113] Factom data structures (Directory Blocks, Entry Blocks, Entries) are needed for Factom to be useful. They are public and will be preserved in two places. The Authority Servers need to maintain this data to make correct decisions about adding new Entries. Since they have this data, they can provide it as a service, as part of being a full node. As the protocol grows the protocol will be able to support partial nodes, which share only part of the Factom dataset. The partial nodes could share only the data which is relevant to their specific application. Peer discovery for the partial nodes may be handled by any sort of directory service, such as a Distributed Hash Table (DHT).

[0114] FIG. 7 illustrates a read operation, according to exemplary embodiments. This setup allows for efficient peer distribution of data even if the entire Factom dataset grows to unwieldy sizes. The Directory Service also allows the data to be preserved independent of any Authority servers or full nodes. Even if all the full nodes were removed from the network, the data could still be shared by a more numerous set of parties interested in specific subsets of the data.

[0115] FIG. 8 further illustrates the Chain ID, according to exemplary embodiments. Factom groups all Entries under a ChainID. The ChainID is computed from a Chain Name. The ChainID is a hash of the Chain Name. The Chain Name is a byte array arbitrarily long in length. See figure below. Since the conversion from Chain Name to ChainID is a hash operation, it is a simple process. Deriving a Chain Name from a ChainID is not simple, so a lookup table would be needed.

[0116] The user may provide a Chain Name, or the Chain Name may be auto-generated. Regardless, that the ChainID can be shown to be a hash of something. This prevents unhashed data from being a ChainID, which is stored all the way up to the Directory Blocks. This convention eliminates insertion of obscene plaintext in the block structure.

[0117] The Chain Name is fairly arbitrary. It could be a random number, a string of text, or a public key. An individual Application could derive meaning from different Chain Names.

[0118] One possible convention would be to use human readable text for the Chain Name. This would allow for the structuring of Chains in a logical hierarchy, even though Chains are not hierarchical by nature. Users can even use the same naming conventions, but by making simple modifications, ensure that there are no accidental intersections between their Chains and other Chains. Consider the following path:

MyFavoriteApp/bin,

[0119] where the slash is a convention for another level of hierarchy. The slash separating ASCII strings "MyFavoriteApp" and "bin" represents transitioning to a deeper level. These two strings must be converted to bytes, and there are many options for doing so. The strings could be encoded in UTF-16, UTF-32, ASCII, or even something like IBM's EPCDIC. Each of these encodings would result in entirely different ChainIDs for the same string, since the computation of the ChainID is done from the bytes. Furthermore, the application could utilize a Globally Unique Identifier (GUID) number as the first byte array in their naming convention. This would eliminate overlap of one Applica-

tion's ChainID "space" with another, at the expense of just a few more bytes in the Chain creation.

Using Factoids to Purchase Entry Credits

[0120] Factoids are the main internal scarcity token used to moderate and reward the system actors. The right to put Entries into Factom is represented by Entry Credits. Factom separates the two value-holding mechanisms, as they serve different purposes. Factoids can be converted into Entry Credits, but not vice versa.

[0121] Factoids are implemented in much the same way Bitcoin is implemented, allowing multiple inputs, multiple outputs, etc. where each input requires the proper signature for the transaction to be valid. Other sorts of validation including multisig is possible. Factoid transactions are managed on a special Factoid Chain. This Factoid Chain is handled more restrictively than other Chains. Entries in the Factoid Chain must be valid Factoid transactions, or the Factom Servers will reject the Entries.

[0122] Factoids are included into the protocol to completely decentralize Factom, and to reduce bloat and spam in both Factom and Bitcoin. Factoids can be converted to Entry Credits in the protocol, and paid out to Factom servers from the protocol. Factoids budgeted but not paid out can remain in a "grant pool". These tokens can be issued to support and develop the protocol from the protocol.

[0123] Factoids also help to bind consensus. If consensus is lost, then the Factoids will fall in value, incentivizing the support of the protocol.

[0124] The conversion of a Factoid to Entry Credits will be done via a special purchase transaction on the Factoid Chain. This purchase transaction will include:

[0125] An Output directing a Factoid amount to be converted

[0126] The public key that is to receive the Entry Credits

[0127] The Entry Credits, once purchased, cannot be transferred to another public key. They can only be used to pay for Entries. This greatly reduces their value to thieves, since they cannot be resold. Entry Credit private keys can be held in low security areas with minimal risk.

Using Entry Credits to Write Entries

[0128] Adding Entries into Factom requires giving up a scarce resource. That resource is Entry Credits, which are derived from Factoids. Adding Entries to Factom is a two step process. First the Entry is paid for (committed). The payment accomplishes two things. It decrements the Entry Credits associated with a user's public key. In the same operation, the hash of the Entry is specified. After the Entry is paid for, the server will wait for the unhashed Entry and include it once seen (revealed).

[0129] 1. Pay for Entry

[0130] Decrement Entry Credits owned by a user

[0131] User specifies hash of Entry in payment

[0132] 2. Insert Entry

[0133] User publishes Entry for inclusion in Entry Block

[0134] There are many benefits of this two step process. One benefit is to separate the payment overhead from the recorded data. Future users will not be forced to download the data generated by payment minutia. They only need to

download the minimum data to validate their system. It allows users to safely and easily ignore the payment information.

[0135] Another benefit is censorship resistance. By committing to accept an Entry before knowing the content makes censorship by the Factom servers obvious. Adam Back has advocated for a similar mechanism for Bitcoin in a post titled "Blind Symmetric Commitment for Stronger Byzantine Voting Resilience" (<https://bitcointalk.org/index.php?topic=206303.0>). If a user or Audit server can show an Entry which has been properly been paid for, but none of the Federated servers are accepting it, then the censorship is provable.

[0136] The transactions deducting Entry Credits will be recorded in a special Chain, similar to the Factoid Chain. The Federated servers will only fill the Chain with valid Entry Credit transactions.

Setting the Cost of Entries with a Central Server Oracle

[0137] The conversion rate of Factoids to Entry Credits will be determined by first choosing a target real world value for an Entry Credit. This target will be determined by a distributed and autonomous process. At minimum it will be agreed upon by some process driven by the Authority Set. Other parties might be involved through various auditable processes in Factom to further decentralize the decision.

[0138] Once a target real world target price of an Entry Credit has been chosen, an Oracle is required to record into Factom the conversion value between Factoids and that EC price. That specification and implementation will also go through a decentralized decision process. The actual implementation of the target price, oracle implementation, and exchange rate adjustment can vary widely, but will be optimized for decentralization, security, and regulatory compliance.

[0139] Note that fee calculations and rates are subject to change, and don't materially impact the utility of the Factom protocol.

Using Factom without Factoids

[0140] Many users of Factom may not want a wallet, and will not want to hold any cryptocurrency asset. But they will want to create their Chains (ledgers) and add their Entries. Factom's two step recording process allows for the separation of Factoids, Factom's tradable token, from the opportunity to post Entries to Factom, represented by Entry Credits. Servers and other recipients of Factom Tokens can sell Entry Credits to customers for payment via Bitcoin, conventional credit card payments, etc. The user would provide a public key to hold the Entry Credits. The seller would convert the appropriate amount of Factoids to Entry Credits and assign those rights to the user's public key. Users could thus buy Entries Credits for Factom without ever owning the Factoids that drive the Factom servers.

[0141] From a regulation point of view, this is powerful. The servers earn Factoids from the protocol. The only parties to that transaction are the server and the protocol. Then the server sells Entry Credits to users, who eventually return Factoids to the rest of the system. Entry Credits are non transferable, so the user cannot assign them to another user's public key, and selling private keys isn't practical or useful. In neither transaction is a tradable token (the Factoid) transferred between two parties.

[0142] Factom is a distributed, autonomous layer residing on top of the Bitcoin blockchain. The goal of Factom is to provide the power of Bitcoin's blockchain to a nearly

unlimited range of Applications and uses. Further, Factom is architected such that its users do not need any cryptocurrency whatsoever.

[0143] A distributed, immutable ledger is the radical, foundational, and unprecedented technology represented by the Bitcoin blockchain. The dream of many is to extend the honesty inherent to an immutable ledger validated by math to chaotic, real-world interactions. By allowing the construction of unbounded ledgers backed by the blockchain, Factom extends the benefits of the blockchain to the real world.

[0144] FIGS. 9-21 are simplified illustrations of a digital contract 20 in a blockchain environment 22, according to exemplary embodiments. The digital contract 20 is sometimes referred to as a self-executing or “smart” contract between parties to a transaction. The digital contract 20 may be executable code that runs on a blockchain 24. The blockchain 24 has one or more blocks 26 of data. A conventional smart contract facilitates, executes, and/or enforces the terms of an agreement. Whatever the terms, the digital contract 20 may automatically execute the terms once predetermined logical rules, conditions, or code is satisfied. The digital contract 20 may thus be expressed in a programming language. Smart contracts are generally known, so this disclosure will not dwell on the known aspects.

[0145] Here, though, the blockchain 24 need only reference the digital contract 20. That is, the actual programming language defining the digital contract 20 need not be included within or attached to the blockchain 24. Instead, the blockchain 24 need only include or specify a contract identifier 28 and perhaps one or more contractual parameters 30. The contract identifier 28 is any digital identifying information that uniquely identifies or references the digital contract 20. Similarly, the contractual parameters 30 may digitally identify the parties to the digital contract 20, their respective performance obligations and terms, and even consideration. So, instead of the blockchain 24 carrying or conveying the actual code representing the digital contract 20, exemplary embodiments need only specify the contract identifier 28 and perhaps the contractual parameters 30. The blocks 26 of data within the blockchain 24 are thus not burdened with the programming code that is required to execute the digital contract 20. The blockchain 24 need only include or specify the contract identifier 28 and/or the contractual parameters 30 (or their respective hash values), thus greatly simplifying the blockchain 24 and reducing its size (in bytes) and processing requirements.

[0146] FIG. 10 further illustrates the blockchain 24. Here any entity 32 may generate the blockchain 24. While exemplary embodiments may be applied to any entity 32, most readers are thought familiar with financial services. That is, suppose the entity 32 is a bank, lender, or other financial institution 34 (such as PIMCO®, CITI®, or BANK OF AMERICA®). As the reader likely understands, the financial institution 34 creates a massive amount of banking records, transaction records, mortgage instruments, and other private data 36. The financial institution 34 thus has a financial server 38 executing a software application 40 that encrypts its private data 36. While the software application 40 may use any encryption scheme, FIG. 2 illustrates the private blockchain 24. That is, the software application 40 causes the financial server 38 to cryptographically hash the private data 36 and to integrate the resulting hash value(s) into the block 26 of data within the private blockchain 24. Moreover, because the private data 36 may represent con-

tractual obligations between parties, the software application 40 may further cause the blockchain 24 to include the contract identifier 28 and the contractual parameters 30. The contract identifier 28 and the contractual parameters 30 may be encoded as data or information contained within the block 26 of data, or the contract identifier 28 and the contractual parameters 30 may be data or information that is separate from the block 26 of data (such as informational content in metadata or in a packet header/body). Regardless, the blockchain 24 need not include the programming code representing the digital contract 20. The blockchain 24 need only specify the contract identifier 28 and/or the contractual parameters 30.

[0147] FIG. 11 illustrates a contract server 42. The contract server 42 may be responsible for executing the digital contract 20 referenced by the contract identifier 28 and/or the contractual parameters 30. For example, after the financial server 38 (executing the software application 40) generates the block 26 of data within the blockchain 24, the financial server 38 may send the blockchain 24 to the network address (e.g., Internet protocol address) associated with the contract server 42. When the contract server 42 receives the blockchain 24, the contract server 42 inspects the blockchain 24 to identify the contract identifier 28 and/or the contractual parameters 30. Once the contract identifier 28 is determined, the contract server 42 may then consult an electronic database 44 of contracts. The database 44 of contracts has entries that map or relate the contract identifier 28 to its corresponding digital contract 20. The database 44 of contracts, in other words, may identify a computer file 46 that contains the programming language representing the digital contract 20 identified by the contract identifier 28. So, once the digital contract 20 is determined, the contract server 42 may retrieve and locally execute the computer file 46, perhaps based on parameters defined or described by the contractual parameters 30 (such as party names, parameters associated with their respective performance obligations and terms, and consideration). Again, then, the blockchain 24 need only reference the digital contract 20 (using the contract identifier 28 and/or the contractual parameters 30). The actual execution of the digital contract 20 may be offloaded or outsourced to the contract server 42.

[0148] FIG. 12 also illustrates the contract server 42. Here, though, the contract server 42 may only manage the execution of the digital contract 20 referenced by the contract identifier 28 and/or the contractual parameters 30. That is, the contract server 42 may outsource the execution of the digital contract 20 to a vendor, a supplier, or a subcontractor process. Again, when the contract server 42 receives the blockchain 24, the contract server 42 inspects the blockchain 24 to identify the contract identifier 28 and/or the contractual parameters 30. The contract server 42 may then consult the database 44 of contracts. Here, though, the database 44 of contracts has entries that map or relate the contract identifier 28 to a network resource 50 that processes and/or executes the digital contract 20 as a service (perhaps as a software-as-a-service or “SAAS”). The network resource 50 may thus be a remote server, a virtual machine, a web page or web server, a client device/machine, or other resource that executes the digital contract 20. Once the network resource 50 is determined, the contract server 42 may retrieve and send the contractual parameters 30 to the network resource 50 for execution. The network resource 50 (perhaps operated on behalf of a third party) applies the parameters defined or

described by the contractual parameters 30 to the programming code representing the digital contract 20.

[0149] Exemplary embodiments thus only need to identify the digital contract 20. The contract identifier 28 and the contractual parameters 30 need only be informational content in the private blockchain 24. The contract identifier 28 is any digital identifying information that uniquely identifies or references the digital contract 20. The contract identifier 28 may be an alphanumeric combination that uniquely identifies a vendor and/or version of the digital contract 20 and/or a processor or executioner of the digital contract 20. The contract identifier 28 may be expressed as a unique hash value that is included within, or specified by, the private blockchain 24. Similarly, the contractual parameters 30 may identify the parties to the digital contract 20, their respective performance obligations and terms, and consideration.

[0150] FIG. 13 illustrates consideration. When the digital contract 20 is executed, the parties to the digital contract 20 may be compensated (perhaps according to the contractual parameters 30 describing consideration). Moreover, the contract server 42 and/or the network resource 50 may also be compensated. While there are many compensation schemes, this disclosure mostly explains crypto-compensation. That is, when the digital contract 20 successfully executes, perhaps the parties exchange, trade, or transfer cryptographic currencies. Suppose, for example, that the financial institution 34 creates its own cryptographic coinage 60 in the blockchain environment 22. The entity 32, in other words, may establish entity-specific electronic tokens 62 to access and/or to use the blockchain environment 22. Because the private blockchain 24 represents hashes of the financial institution's private data 36, the private blockchain 24 may be considered a private resource or property of the financial institution 34. That is, the private blockchain 24 is controlled by, or affiliated with, the financial institution 34, so the financial institution 34 may control who adds and/or writes to the private blockchain 24 and who reads, accesses, or receives the private blockchain 24.

[0151] The entity-specific tokens 62 may thus be control mechanisms. While the entity-specific tokens 62 may have any functional scheme, FIG. 5 illustrates a private credit token 64 and a private tradeable token 66. The entity's credit token 64, for example, may be acquired and then spent or burned when accessing the financial institution's private blockchain 24. The entity's credit token 64, in other words, represents any credit-based entry system associated with the financial institution's private blockchain 24. The tradeable token 66, on the other hand, may be generated for transfer among others. The entity 32 generates the tradeable token 66 to be traded and/or spent. The tradeable token 66, in other words, may be considered as the entity's specific, private currency to be used as the entity 32 governs.

[0152] Exemplary embodiments may thus trade or exchange crypto-compensation. That is, when the digital contract 20 successfully executes, perhaps the parties exchange, trade, or transfer the credit token 64 and/or the tradeable token 66. When any party, or all the parties, perform their assigned role in the transaction, value is given via the credit token 64 and/or the tradeable token 66. Similarly, the contract server 42 and/or the network resource 50 may also be compensated via the credit token 64 and/or the tradeable token 66, perhaps as a "mining" fee for executing the digital contract 20.

[0153] The digital contract 20 is thus a computer program or code that verifies and/or enforces negotiation and/or performance of a contract between parties. One fundamental purpose of so-called smart contracts is to integrate the practice of contract law and related business practices with electronic commerce protocols between parties or devices via the Internet. Smart contracts may leverage a user interface that provides one or more parties or administrators access, which may be restricted at varying levels for different people, to the terms and logic of the contract. Smart contracts typically include logic that emulates contractual clauses that are partially or fully self-executing and/or self-enforcing. Examples of smart contracts are digital rights management (DRM) used for protecting copyrighted works, financial cryptography schemes for financial contracts, admission control schemes, token bucket algorithms, other quality of service mechanisms for assistance in facilitating network service level agreements, person-to-person network mechanisms for ensuring fair contributions of users, and others. Smart contract infrastructure can be implemented by replicated asset registries and contract execution using cryptographic hash chains and Byzantine fault tolerant replication. For example, each node in a peer-to-peer network or blockchain distributed network may act as a title registry and escrow, thereby executing changes of ownership and implementing sets of predetermined rules that govern transactions on the network. Each node may also check the work of other nodes and in some cases, as noted above, function as miners or validators.

[0154] FIG. 14 further illustrates the contract server 42. When the contract server 42 receives the blockchain 24, here the contract server 42 may generate data records 70 in a blockchain data layer 72, as later paragraphs will explain. The contract server 42 may thus be termed or called a data layer server 74. Moreover, the blockchain data layer 72 may also add another layer of cryptographic hashing to generate a public blockchain 76. The blockchain data layer 72 acts as a validation service 78 that validates the digital contract 20 was executed. Moreover, the blockchain data layer 72 may generate a cryptographic proof 80. The public blockchain 76 thus publishes the cryptographic proof 80 as a public ledger 82 that establishes chains of blocks of immutable evidence.

[0155] FIGS. 15-16 illustrate examples of the entity-specific tokens 62. Suppose that a third-party 90 wishes to receive, read, write to, or otherwise access the financial institution's private blockchain 24 and/or the digital contract 20. As FIG. 15 illustrates, exemplary embodiments may require that the third-party 90 spend or burn one or more of the credit tokens 64. The credit token 64 may thus control access to the financial institution's private blockchain 24 and/or the digital contract 20. The inventor envisions that vendors, service providers, individual users, and other third-parties 90 may wish to access the hash values of the private data 36 contained within the financial institution's private blockchain 24. Moreover, the third party may want to access, inspect, execute, or verify the digital contract 20. The financial institution 34 may thus require that the third-party 90 redeem the entity's credit token(s) 64 before granting read, write, or access permission to the digital contract 20. The financial institution 34 may additionally or alternatively require redemption of the entity's credit token(s) 64 for using protocols, rules, and application programming interfaces ("APIs") associated with the private blockchain 24 and/or the digital contract 20. The financial institution 34

may thus establish or issue its own credit tokens **64** and even govern their usage restrictions **92** and value **94**, as later paragraphs will explain.

[0156] FIG. 16 illustrates the tradeable token **66**. The financial institution **34** may establish the tradeable token **66** and also govern its usage restrictions **92** and value **94**. The tradeable token **66**, in other words, is a cryptocurrency or “coin.” Again, while exemplary embodiments may utilize any functional scheme, the tradeable token **66** may be earned. That is, anyone (such as the third party **90**) may earn the tradeable token **66** according to the usage restrictions **92**. For example, suppose the data layer server **74** earns the entity’s tradeable token(s) **52** in exchange for processing and/or managing an execution of the digital contract **20**. The data layer server **74** may additionally or alternatively earn the entity’s tradeable token(s) **52** in exchange for the validation service **78**. That is, a provider of the validation service **78** is paid, or earns, the entity’s tradeable token(s) **52** for processing or executing the digital contract **20** and/or for cryptographically hashing the proof **80** of the digital contract **20**. The provider of the validation service **78** may also be paid in the entity’s tradeable token(s) **52** for publishing the proof **80**. The tradeable token **66** may thus be transferred as currency according to the usage restrictions **92** and its value **94**.

[0157] FIG. 17 illustrates transaction records **100**. Whenever the entity-specific tokens **62** are created, owned, or transferred, the transaction record **100** may be generated. The transaction record **100** may then be documented in the blockchain environment **22**. For example, the entity-specific tokens **62** may be addressable. That is, the credit token **64** and the tradeable token **66** may be uniquely associated with a common, single cryptographic address **102**. The cryptographic address **102** may represent an owner or holder (e.g., the entity **32** or the third-party **90**). When the entity-specific tokens **62** are created, generated, or assigned, the entity-specific tokens **62** may be assigned or associated with the cryptographic address **102**. The cryptographic address **102** may then be received by, and propagated within, the blockchain data layer **72** to identify the corresponding data records **70**. The blockchain data layer **72** may even hash the cryptographic address **102** as the cryptographic proof **80** of the transaction records **100**. Exemplary embodiments thus publicly document the transaction records **100** involving the entity-specific tokens **62**, based on the single cryptographic address **102**. In simple words, the blockchain data layer **72** publishes ownership and transfer proofs **80** of the credit token **64** and the tradeable token **66** based on the transaction records **100** associated with the single cryptographic address **102**.

[0158] The transaction records **100** may also document the digital contract **20**. Whenever the digital contract **20** is specified, generated, processed, or even executed, the transaction record **100** may be generated. The transaction record **100** may then be documented in the blockchain environment **22**. For example, the entity-specific tokens **62** may be earned as payment according to the executable terms of the digital contract **20**. The entity-specific tokens **62** may additionally or alternatively be earned or awarded for processing or executing a portion of, or entirely, the digital contract **20**. The entity-specific tokens **62** may thus be uniquely associated with a party to the digital contract **20** and/or with a service provider/processor of the digital contract **20**. The transaction record **100** may document the parties to the

digital contract **20**, a transactional description describing a transaction governed by the digital contract **20**, and any financial or performance terms. The transaction record **100** may thus document an offer, an acceptance, a consideration, and terms. For simplicity, then, the single cryptographic address **102** may represent a party to the digital contract **20** and/or with a service provider/processor of the digital contract **20**. Regardless, when the entity-specific tokens **62** are created, generated, or assigned, the entity-specific tokens **62** may be received by, and propagated within, the blockchain data layer **72** to identify the corresponding data records **70**. The blockchain data layer **72** may thus publish the proofs **80** of the digital contract **20** and any entity-specific tokens **62** paid or exchanged, according to the transaction records **100**.

[0159] FIG. 18 illustrates a filling station **110** in the blockchain environment **22**. Because the tokens **62** may be consumed by users (such as during or after any processing or execution of the digital contract **20**), the filling station **110** allows the third party **90** to replenish or fill an account **112**. Recall that the third-party entity **32** may be required to spend the tokens **62** to access the financial institution’s private blockchain **24** and/or the digital contract **20**. Moreover, the tokens **62** may also be earned or transferred according to the terms of the digital contract **20**. The account **112** may thus be established, and the account **112** maintains a monetary or numerical balance **114** of the tokens **62**. As the tokens **62** are spent, traded, or redeemed, the account **112** may need filling to continue using or accessing the blockchain **24** and/or the digital contract **20**.

[0160] The filling station **110** may access both the transaction records **100** and the blockchain data layer **72**. Because the blockchain data layer **72** may document the data records **70** using the single cryptographic address **102**, the single cryptographic address **102** may serve as a common reference or query parameter with the entity’s transaction records **100**. The filling station **110**, in other words, may use the single cryptographic address **102** to identify the transaction records **100** that correspond to the blockchain data layer **72**. The filling station **110** may thus present a transaction summary of the account **112** and the balance **114**. Because blockchain data layer **72** may track and/or prove the transaction records **100**, exemplary embodiments may search the blockchain data layer **72** for the single cryptographic address **102**. That is, the filling station **110** may query the blockchain data layer **72** for the single cryptographic address **102**, and the blockchain data layer **72** may identify the transaction records **100** that match the single cryptographic address **102**. Similarly, exemplary embodiments may query the blockchain data layer **72** for the contract identifier **28** and/or the contractual parameters **30**, and the blockchain data layer **72** may identify the transaction records **100** that match the contract identifier **28** and/or the contractual parameters **30**. The filling station **110** may then process the transaction records **100** to provide the transaction summary of the account **112**, the balance **114**, and any other transactional data. The filling station **110** may also allow the user to replenish an amount or value of the tokens **62**, thus allowing the user to continue exchanging the tokens **62** for access to the private blockchain **24**, the blockchain data layer **72**, and/or the digital contract **20**. The filling station **110** may thus be an access mechanism to the blockchain data layer **72**.

[0161] FIG. 19 further illustrates the filling station **110**. Here the blockchain data layer **72** may have its own cryptocurrency **120**. That is, a provider of the blockchain data

layer 72 may establish its cryptocurrency 120 for accessing and/or using the validation service 78. The cryptocurrency 120 may thus include a credit token and a tradeable token (not shown for simplicity). The credit token may be required to enter or access the blockchain data layer 72 to receive the validation service 78, and the tradeable token may be earned for participating in the validation service 78. Regardless, the filling station 110 may use the single cryptographic address 102. The third party 90 may use the single cryptographic address 102 to access the entity's cryptocurrency 60 and the blockchain data layer's cryptocurrency 120. Exemplary embodiments may thus identify and track the transaction records 100 and the blockchain data layer's cryptocurrency 120 using the same, single cryptographic address 102.

[0162] Exemplary embodiments thus present elegant solutions. Any entity 32 may create its own private blockchain 24 and offer or present the digital contract 20 for self-execution. The entity 32 may then establish or create the tokens 62 for using, accessing, or processing the entity's private blockchain 24 and/or the digital contract 20. The tokens 62 may have the value 94, thus fostering a market for entity-specific tradeable assets in the blockchain environment 22. The tradable value 94 of the tokens 62 may thus drive demand to use the digital contracts 20. Exemplary embodiments may thus provide a two-token system that isolates any use of the entity's private blockchain 24 from the entity's tradeable token 66. Moreover, the credit token 64 may be associated with the third party 90 (perhaps via the single cryptographic address 102), thus allowing the third party 90 to retrieve the account balance 114 from the filling station 110 and sign entries or other transactions. Moreover, the third party 90 may also use the single cryptographic address 102 to access the blockchain data layer 72 via the filling station 110. The filling station 110 is a single resource or destination (such as a secure website) for managing a user's cryptographic coinage 60 and defining payments according to the digital contract 20.

[0163] FIG. 20 expands the entity concept. Here multiple, different entities 32a-d provide their respective software applications 40a-d that encrypt their respective private data 36a-d as their individual, private blockchains 24a-d. While exemplary embodiments may be applied to any number of industries or services, FIG. 20 illustrates a simple example of four (4) different entities 32a-d. First entity 32a, for example, again represents the bank, lender, or other financial institution 34 that encrypts its private data 36a as its private blockchain 24a. Second entity 32b represents any retailer 122 (such as HOME DEPOT®, KOHL'S®, or WAL-MART®) that encrypts its private data 36b as its private blockchain 24b. Third entity 32c represents a web site 124 offering a service 126 (such as AMAZON®, NETFLIX®, or GOOGLE®) that encrypts its private data 36c as the private blockchain 24c. Fourth entity 32d represents an automotive or other manufacturer or supplier 128 (such as FORD®, TOYOTA®, or DELPHI®) that encrypts its private data 36d as the private blockchain 24d. The entities 32a-d thus use their respective software applications 40a-d to provide a first layer 130 of cryptographic hashing. The entities 32a-d may also use their respective software applications 40a-d to issue their own private and entity-specific cryptocurrency 60a-d. Each entity 32a-d may then send their respective private blockchains 24a-d to the blockchain data layer 72, and the blockchain data layer 72 may add a second layer 132 of cryptographic hashing. The blockchain data layer 72 thus

generates the public blockchain 76 as a public resource or utility for record keeping. Any entity 32 that subscribes to the blockchain data layer 72 (such as by acquiring and/or spending the cryptocurrency 120) may thus access, read, and/or store the proofs 80 of its private data 36 to the public blockchain 76. The blockchain data layer 72, in other words, acts as the public ledger 82 that establishes chain of blocks of immutable evidence.

[0164] As FIG. 20 also illustrates, each entity 32a-d may establish its own private cryptocurrency 60a-d. Each entity's private software application 40a-d may create and/or issue its cryptocurrency 60a-d (such as respective entity-specific tokens 62 above explained). Each entity 32a-d may also establish its own usage restrictions and value (illustrated as reference numerals 92 and 94 in FIGS. 15-16) according to rules governing ownership, trade, and other policies. Each entity 32a-d may generate and send its respective transaction records 100a-d which reference each entity's single cryptographic address 102a-d to the blockchain data layer 72 for documentation.

[0165] As FIG. 20 further illustrates, each entity 32a-d may also specify their respective digital contract 20a-d. When any of the private blockchains 24a-d is received, the blockchain data layer 72 may coordinate execution of any digital contract 20a-d. The blockchain data layer 72, for example, may inspect any private blockchain 24a-d and identify any information associated with the digital contract 20a-d. The blockchain data layer 72 may then execute the digital contract 20a-d, and/or the blockchain data layer 72 may identify a service provider that executes the digital contract 20a-d. The blockchain data layer 72, in other words, may manage the execution of the digital contracts 20a-d according to a subcontractor relationship. A provider of the blockchain data layer 72 may then be compensated via any entity's cryptocurrency 60a-d and/or the blockchain data layer's cryptocurrency 120.

[0166] As FIG. 21 illustrates, the filling station 110 may be agnostic. Any user (such as the entity 32a-d or the third party 90) may authenticate to the filling station 110. Once authenticated, the user need only enter or provide the correct single cryptographic address 102a-d to access the entity's private cryptocurrency 60a-d, the blockchain data layer's cryptocurrency 120, and/or the entity's digital contract 20a-d. The single cryptographic address 102a-d, in other words, allows the user to access her account 112 and balance 114 for the entity's private cryptocurrency 60a-d, the blockchain data layer's cryptocurrency 120, and/or the entity's digital contract 20a-d. The user may thus easily conduct transactions between the entity's private cryptocurrency 60a-d and the blockchain data layer's cryptocurrency 120. The entity 32a-d, for example, may fuel or replenish its supply of the blockchain data layer's cryptocurrency 120, perhaps by redeeming or exchanging the entity's private cryptocurrency 60a-d (perhaps according to an exchange rate or other value). Similarly, the provider of the blockchain data layer 72 may fuel or replenish its supply of the entity's private cryptocurrency 60a-d by purchasing or exchanging the blockchain data layer's cryptocurrency 120. The provider of the blockchain data layer 72 may also earn the entity's private cryptocurrency 60a-d by processing any portion of, or by executing, the entity's digital contract 20a-d. Moreover, the respective private blockchains 24a-d and the blockchain data layer 72 would contain the data records 70 confirming the processing and/or execution of the digital contract 20a-d,

so the transaction records **100a-d** thus propagate into the blockchain data layer **72** for public disclosure via the public blockchain **76**. Any user that successfully authenticates to the filling station **110** may access a full accounting of his or her digital cryptocurrencies **60a-d** and/or **120** and any digital contracts **20**, perhaps according to the respective single cryptographic address **102a-d**. The user may thus buy, sell, trade, and/or redeem any entity-specific cryptocurrencies **20a-d** and/or **90**, all by accessing the filling station **110**. The user may buy or sell any entity's coins or replenish credits, all by accessing the filling station **110**. The user may also track performance or obligations defined by the digital contracts **20a-d** and any payments or consideration received or paid.

[0167] Exemplary embodiments thus present another elegant solution. The filling station **110** is another service offered by the blockchain data layer **72**. Because all the transaction records **100** in the blockchain data layer **72** are identifiable (perhaps via the single cryptographic address **102**), the filling station **110** can present the summary of the user's credit tokens and tradeable tokens. The filling station **110** may thus provide a single or universal electronic wallet for all of a user's digital coinage and credits, regardless of the issuing entity **32a-d**. The user may thus only perform a single authentication to the blockchain data layer **72** and access all her cryptofunds.

[0168] FIGS. 22-24 are more detailed illustrations of an operating environment, according to exemplary embodiments. FIG. 22 illustrates an entity server **140** communicating with the data layer server **74** via a communications network **142**. The entity server **140** operates on behalf of the entity **32** and generates the entity's private blockchain **24** (such as the financial server **38** explained with reference to FIGS. 10-19). The entity server **140**, in other words, has a processor **144** (e.g., "µP"), application specific integrated circuit (ASIC), or other component that executes the entity's software application **40** stored in a local memory device **146**. The entity server **140** has a network interface to the communications network **142**, thus allowing two-way, bidirectional communication with the data layer server **74**. The entity's software application **40** includes instructions, code, and/or programs that cause the entity server **140** to perform operations, such as calling, invoking, and/or applying an electronic representation of a hashing algorithm **148** to the entity's private data **36**. The hashing algorithm **148** thus generates one or more hash values **150**, which are incorporated into the entity's private blockchain **24**. The entity's software application **40** then instructs the entity server **140** to send the private blockchain **24** via the communications network **142** to a network address (e.g., Internet protocol address) associated with the data layer server **74**.

[0169] The digital contract **20** may also be identified. The entity's software application **40** may also instruct the entity server **140** to specify the digital contract **20** as informational content in the private blockchain **24**. For example, the digital contract **20** may be identified by the contract identifier **28** and contractual parameters **30**. The contract identifier **28** is any digital identifying information that uniquely identifies or references the digital contract **20**. The contract identifier **28** may be an alphanumeric combination that uniquely identifies a vendor and/or version of the digital contract **20** and/or a processor or executioner of the digital contract **20**. The contract identifier **28** may also be one of the unique hash values **150** (perhaps generated by the hashing algorithm

148) that is included within, or specified by, the private blockchain **24**. Similarly, the contractual parameters **30** may identify the parties to the digital contract **20**, their respective performance obligations and terms, and consideration.

[0170] FIG. 23 illustrates the blockchain data layer **72**. The data layer server **74** has a processor **152** (e.g., "µP"), application specific integrated circuit (ASIC), or other component that executes a data layer application **154** stored in a local memory device **156**. The data layer server **74** has a network interface to the communications network **142**. The data layer application **154** includes instructions, code, and/or programs that cause the data layer server **74** to perform operations, such as receiving the entity's private blockchain **24**, the digital contract **20**, the contract identifier **28**, and/or the contractual parameters **30**. The data layer application **154** then causes the data layer server **74** to generate the blockchain data layer **72**. The data layer application **154** may optionally call, invoke, and/or apply the hashing algorithm **148** to the data records **70** contained within the blockchain data layer **72**. The data layer application **154** may also generate the public blockchain **76**. The data layer application **154** may thus generate the public ledger **82** that publishes, records, or documents the digital contract **20**, the contract identifier **28**, and/or the contractual parameters **30**. Indeed, if the data layer application **154** processes and/or manages the digital contract **20**, the data records **70** may document any processing or execution, and the data layer application **154** may optionally apply the hashing algorithm **148** to the data records **70** to generate the cryptographic proof **80** of the digital contract **20**.

[0171] FIG. 24 illustrates additional publication mechanisms. Once the blockchain data layer **72** is generated, the blockchain data layer **72** may be published in a decentralized manner to any destination. The data layer server **74**, for example, may generate and distribute the public blockchain **76** (via the communications network **142** illustrated in FIGS. 22-23) to one or more federated servers **160**. While there may be many federated servers **160**, for simplicity FIG. 24 only illustrates two (2) federated servers **160a** and **160b**. The federated servers **160a** and **160b** provide a service and, in return, they are compensated according to a compensation or services agreement or scheme.

[0172] Exemplary embodiments include still more publication mechanisms. For example, the cryptographic proof **80** and/or the public blockchain **76** may be sent (via the communications network **142** illustrated in FIGS. 22-23) to a server **162**. The server **162** may then add another, third layer of cryptographic hashing (perhaps using the hashing algorithm **148**) and generate another or second public blockchain **164**. While the server **162** and/or the public blockchain **164** may be operated by, or generated for, any entity, exemplary embodiments may integrate another cryptographic coin mechanism. That is, the server **162** and/or the public blockchain **164** may be associated with BITCOIN®, ETHEREUM®, RIPPLE®, or other cryptographic coin mechanism. The cryptographic proof **80** and/or the public blockchain **76** may be publicly distributed and/or documented as evidentiary validation. The cryptographic proof **80** and/or the public blockchain **76** may thus be historically and publicly anchored for public inspection and review.

[0173] Exemplary embodiments may be applied regardless of networking environment. Exemplary embodiments may be easily adapted to stationary or mobile devices having cellular, wireless local area network (WI-FI®), near field,

and/or BLUETOOTH® capability. Exemplary embodiments may be applied to mobile devices utilizing any portion of the electromagnetic spectrum and any signaling standard (such as the IEEE 802 family of standards, GSM/CDMA/TDMA or any cellular standard, and/or the ISM band). Exemplary embodiments, however, may be applied to any processor-controlled device operating in the radio-frequency domain and/or the Internet Protocol (IP) domain. Exemplary embodiments may be applied to any processor-controlled device utilizing a distributed computing network, such as the Internet (sometimes alternatively known as the “World Wide Web”), an intranet, a local-area network (LAN), and/or a wide-area network (WAN). Exemplary embodiments may be applied to any processor-controlled device utilizing power line technologies, in which signals are communicated via electrical wiring. Indeed, exemplary embodiments may be applied regardless of physical componentry, physical configuration, or communications standard(s).

[0174] Exemplary embodiments may utilize any processing component, configuration, or system. Any processor could be multiple processors, which could include distributed processors or parallel processors in a single machine or multiple machines. The processor can be used in supporting a virtual processing environment. The processor could include a state machine, application specific integrated circuit (ASIC), programmable gate array (PGA) including a Field PGA, or state machine. When any of the processors execute instructions to perform “operations,” this could include the processor performing the operations directly and/or facilitating, directing, or cooperating with another device or component to perform the operations.

[0175] Exemplary embodiments may packetize. When the entity server 140 and the data layer server 74 communicate via the communications network 142, the entity server 140 and the data layer server 74 may collect, send, and retrieve information. The information may be formatted or generated as packets of data according to a packet protocol (such as the Internet Protocol). The packets of data contain bits or bytes of data describing the contents, or payload, of a message. A header of each packet of data may contain routing information identifying an origination address and/or a destination address.

[0176] FIGS. 25-29 further illustrate the blockchain data layer 72, according to exemplary embodiments. The blockchain data layer 72 chains hashed directory blocks 170 of data into the public blockchain 76. For example, the blockchain data layer 72 accepts input data (such as the entity’s private blockchain 24 illustrated in FIGS. 9-21) within a window of time. While the window of time may be configurable from fractions of seconds to hours, exemplary embodiments use ten (10) minute intervals. FIG. 25 illustrates a simple example of only three (3) directory blocks 170a-c of data, but in practice there may be millions or billions of different blocks. Each directory block 184 of data is linked to the preceding blocks in front and the following or trailing blocks behind. The links are created by hashing all the data within a single directory block 184 and then publishing that hash value within the next directory block.

[0177] As FIG. 26 illustrates, published data may be organized within chains 172. Each chain 172 is created with an entry that associates a corresponding chain identifier 174. Each entity 32a-f, in other words, may have its corresponding chain identifier 174a-d. The blockchain data layer 72 may thus track any data associated with the entity 32a-f with

its corresponding chain identifier 174a-d. New and old data in time may be associated with, linked to, identified by, and/or retrieved using the chain identifier 174a-d. Each chain identifier 174a-d thus functionally resembles a directory 176a-d (e.g., files and folders) for organized data entries according to the entity 32a-f.

[0178] FIG. 27 illustrates the data records 70 in the blockchain data layer 72. As data is received as an input (such as the private blockchain 24 and/or the digital contract 20 illustrated in FIGS. 9-21), data is recorded within the blockchain data layer 72 as an entry 180. While the data may have any size, small chunks (such as 10KB) may be pieced together to create larger file sizes. One or more of the entries 180 may be arranged into entry blocks 182 representing each chain 172 according to the corresponding chain identifier 174. New entries for each chain 172 are added to their respective entry block 182 (again perhaps according to the corresponding chain identifier 174). After the entries 180 have been made within the proper entry blocks 182, all the entry blocks 182 are then placed within in the directory block 184 generated within or occurring within a window 186 of time. While the window 186 of time may be chosen within any range from seconds to hours, exemplary embodiments may use ten (10) minute intervals. That is, all the entry blocks 182 generated every ten minutes are placed within in the directory block 184.

[0179] FIG. 28 illustrates cryptographic hashing. The data layer server 74 executes the data layer application 154 to generate the data records 70 in the blockchain data layer 72. The data layer application 154 may then instruct the data layer server 74 to execute the hashing algorithm 148 on the data records 70 (such as the directory block 184 illustrated in FIGS. 25-27). The hashing algorithm 148 thus generates one or more hash values 150 as a result, and the hash values 150 represent the hashed data records 70. As one example, the blockchain data layer 72 may apply a Merkle tree analysis to generate a Merkle root (representing a Merkle proof 80) representing each directory block 184. The blockchain data layer 72 may then publish the Merkle proof 80 (as this disclosure explains).

[0180] FIG. 29 illustrates hierarchical hashing. The entity’s private software application 40 provides the first layer 130 of cryptographic hashing and generates the private blockchain 24. The entity 32 then sends its private blockchain 24 (perhaps referencing or specifying the digital contract 20) to the data layer server 74. The data layer server 74, executing the data layer application 154, generates the blockchain data layer 72. The data layer application 154 may optionally provide the second or intermediate layer 132 of cryptographic hashing to generate the cryptographic proof 80. The data layer application 154 may also publish any of the data records 70 as the public blockchain 76, and the cryptographic proof 80 may or may not also be published via the public blockchain 76. The public blockchain 76 and/or the cryptographic proof 80 may be optionally sent to the server 162 as an input to yet another public blockchain 164 (again, such as BITCOIN®, ETHEREUM®, or RIPPLE®) for a third layer 188 of cryptographic hashing and public publication. The first layer 130 and the second layer 132 thus ride or sit atop a conventional public blockchain 164 (again, such as BITCOIN®, ETHEREUM®, or RIPPLE®) and provide additional public and/or private cryptographic proofs 80.

[0181] Exemplary embodiments may use any hashing function. Many readers may be familiar with the SHA-256 hashing algorithm. The SHA-256 hashing algorithm acts on any electronic data or information to generate a 256-bit hash value as a cryptographic key. The key is thus a unique digital signature. There are many hashing algorithms, though, and exemplary embodiments may be adapted to any hashing algorithm.

[0182] FIGS. 30-32 are more detailed illustrations of the digital contract 20, according to exemplary embodiments. The private entity 32 sends its private blockchain 24 to the network address associated with the data layer server 74 that generates the blockchain data layer 72. The private blockchain 24 may contain information representing the transaction records 100 associated with the entity's private cryptocurrency 60 (perhaps as one or more privately hashed blocks of data). The private blockchain 24 may also specify, or incorporate, information or data representing the single cryptographic address 102 and/or the digital contract 20 (e.g., the contract identifier 28 and the contractual parameters 30). The single cryptographic address 102 and/or the digital contract 20 (e.g., the contract identifier 28 and the contractual parameters 30) may additionally or alternatively be separately sent from the entity server 140 to the data layer server 74 (perhaps via the communications network 142 illustrated by FIGS. 22-23). Regardless, the entity's private cryptocurrency 60 may be associated with the digital contract 20 (e.g., the contract identifier 28 and the contractual parameters 30) and/or the single cryptographic address 102. The transaction records 100 and/or their privately hashed blocks of data may thus specify, include, reference, and/or be associated with, and/or identified by, the single cryptographic address 102, the digital contract 20, the contract identifier 28, and/or the contractual parameters 30. Because the contract identifier 28 (and/or its corresponding hash value) is an identifiable input to the data layer server 74 generating the blockchain data layer 72, the data records 70 may also carry or reference the contract identifier 28 and/or the contractual parameters 30. So, should the blockchain data layer 72 create or issue its own cryptocurrency 120, the cryptocurrency 120 may also reference, be identified by, or be associated with the single cryptographic address 102 and/or the contract identifier 28 and/or the contractual parameters 30. The single cryptographic address 102, the contract identifier 28, and/or the contractual parameters 30 may thus common indicators or reference data for tracking both the entity's private cryptocurrency 60 and the cryptocurrency 120 issued by the blockchain data layer 72, according to the terms of the digital contract 20. The transaction records 100 (representing entity's private cryptocurrency 60) may thus be commonly mapped or identified to the cryptocurrency 120 issued by the blockchain data layer 72 and to the digital contract 20.

[0183] FIG. 31 illustrates a simple illustration. Once the contract identifier 28 (and/or its corresponding hash value) is received, the contract identifier 28 may propagate and be recorded within the blockchain data layer 72. The contract identifier 28, for example, may be recorded in any of the entries 180. The entry 180, and thus the contract identifier 28, may then be recorded and/or arranged as the entry block 182 and placed within the directory block 184. The entry 180, the entry block 182, and the directory block 184 may thus reference, specify, or be associated with, the contract identifier 28. The contract identifier 28 has thus propagated

as informational content from the private blockchain 24 and into and through the blockchain data layer 72. The contract identifier 28 thus hierarchically moves through the multiple layers of cryptographic hashing for public publication. The blockchain data layer 72 thus tracks the transaction records 100 involving the contract identifier 28. In simple words, the blockchain data layer 72 may track contractual performance of the digital contract 20 via the transaction records 100 that reference or contain the contract identifier 28. Moreover, the blockchain data layer 72 may also track ownership and transfer of the entity's private cryptocurrency 60 and the cryptocurrency 120 issued by the blockchain data layer 72, all via the common single cryptographic address 102 and/or the contract identifier 28.

[0184] FIG. 32 illustrates more details. While the single cryptographic address 102 and/or the contract identifier 28 may be any alphanumeric entry or biometric input, FIG. 24 illustrates a common authentication mechanism 190. Here the same or similar authentication mechanism 190 is used to access both the entity's private cryptocurrency 60 and the cryptocurrency 120 issued by the blockchain data layer 72. If a user of the blockchain data layer 72 satisfies the authentication mechanism 190, then exemplary embodiments may access both the private cryptocurrency 60, the cryptocurrency 120, and/or the data records 70 associated with the contract identifier 28. As a simple example, suppose the user of the authentication mechanism 190 supplies information or data representing the single cryptographic address 102 and/or the contract identifier 28. The single cryptographic address 102 and/or the contract identifier 28 may be any unique alphanumeric entry, biometric input, user identifier, or other authentication credential. For example, most readers are likely familiar with an alphanumeric username and password, which is a common authentication mechanism 190. FIG. 32, though, illustrates a passphrase 192 (such as a multi-word mnemonic). When the entity's private cryptocurrency 60 is/are created, generated, or assigned, the entity's private cryptocurrency 60 may be assigned or associated with the passphrase 192. The passphrase 192 is unique to the registered owner, possessor, or user of the entity's private cryptocurrency 60. The passphrase 192 may even be hashed as a hash value and supplied to the blockchain data layer 72 (as above explained). The passphrase 192, in other words, may be hashed as the single cryptographic address 102 and propagated within the blockchain environment 22 to document the transaction records 100 involving the entity's private cryptocurrency 60.

[0185] The passphrase 192 may also authenticate to the cryptocurrency 120. If the user correctly supplies the passphrase 192, then the same user may conduct transactions involving the cryptocurrency 120 issued by the blockchain data layer 72 and/or involving the contract identifier 28 associated with the digital contract 20. Exemplary embodiments thus allow the user to order transactions and exchanges involving the entity's private cryptocurrency 60, the cryptocurrency 120 issued by the blockchain data layer 72, and/or the digital contract 20.

[0186] FIGS. 33-35 further illustrate the access mechanism, according to exemplary embodiments. The filling station 110 may be a public and/or private service for financial transactions involving the entity's private cryptocurrency 60, the cryptocurrency 120 issued by the blockchain data layer 72, and/or the digital contract 20. FIG. 33 illustrates the filling station 110 as a software-as-a-service

offered by the secure data layer server 74 for accessing the blockchain data layer 72. The filling station 110, for example, may be a module within, or called by, the data layer application 154. A user accesses the filling station 110 to conduct transactions involving her private cryptocurrency 60, the cryptocurrency 120 (issued by the blockchain data layer 72), and/or the digital contract 20. While the filling station 110 may have any user interface, FIG. 33 illustrates a web interface 194. That is, the filling station 110 may be accessed via a webpage 196. The webpage 196 prompts the user to input her authentication credentials according to the authentication mechanism 190 (such as typing the passphrase 192 into a data field or audibly speaking the passphrase 192).

[0187] FIG. 34 further illustrates the web interface 194. The user accesses the filling station 110 using a user device 200. While the user device 200 may be any processor-controlled device, most readers are familiar with a smartphone 202. If the smartphone 202 correctly sends authentication credentials (such as the single cryptographic address 102 and/or passphrase 192, as above explained), then the smartphone 202 may utilize the web interface 194 to the data layer server 74 and/or the blockchain data layer 72. The smartphone 202 executes a web browser and/or a mobile application to send a request 204 specifying an address or domain name associated with or representing the filling station 110. The web interface 194 to the data layer server 74 thus sends the webpage 196 as a response, and the user's smartphone 202 downloads the webpage 196. The smartphone 202 has a processor and memory device (not shown for simplicity) that causes a display of the webpage 196 as a graphical user interface (or "GUI") 206 on its display device 208. The GUI 206 may generate one or more prompts or fields for specifying the authentication mechanism 190 and transactional options. For example, the user preferably enters, speaks, or otherwise provides the passphrase 192. Exemplary embodiments may or may not hash the authentication passphrase (using the hashing algorithm 148 above explained) to produce or generate a hashed passphrase. Exemplary embodiments may then search the blockchain data layer 72 for the data records 70. That is, exemplary embodiments may query the blockchain data layer 72 for a query parameter (such as the contract identifier 28 and/or its hashed value) and the blockchain data layer 72 identifies the data records 70 that match or reference the query parameter. The filling station 110 may then process the data records 70 to provide a transactional summary 210 of the digital contract 20. The filling station 110 may also allow the user to replenish an amount or value of the private cryptocurrency 60 and/or the cryptocurrency 120, even allowing the user to continue exchanging the cryptocurrency 60 for access to the blockchain data layer 72.

[0188] Exemplary embodiments may thus share the common authentication mechanism 190. If the entity's private software application 40 requires the same passphrase 192 to establish any terms of the digital contract 20, then the passphrase 192 may have been hashed and recorded within the blockchain data layer 72. The single cryptographic address 102, the contract identifier 28, and/or the passphrase 192 may be associated with the data records 70 representing the digital contract 20, the private cryptocurrency 60 (issued by the entity 32), and the cryptocurrency 120 (issued by the blockchain data layer 72). The filling station 110 may thus identify any of the data records 70 that are commonly

associated with the contract identifier 28, the private cryptocurrency 60 (issued by the entity 32), and/or the cryptocurrency 120. The filling station 110 thus allows the user to exchange cryptocurrency 60 and 90 for access to the private blockchain 24 and/or the blockchain data layer 72.

[0189] FIG. 35 illustrates a query mechanism. Here the data layer server 74 may access a database 220 of data layer records. The database 220 of data layer records provides a referential record of the informational content contained within the blockchain data layer 72. FIG. 35 illustrates the data layer server 74 locally storing the database 220 of data layer records in its local memory device 156, but the database 220 of data layer records may be remotely stored and accessed via the communications network 142. Regardless, the data layer server 74 may query the database 220 of data layer records for the single cryptographic address 102 and/or the contract identifier 28 and identify and/or retrieve any corresponding data records 70. While the database 220 of data layer records may have any logical structure, FIG. 35 illustrates the database 220 of data layer records as a table 222 that maps, converts, or translates the single cryptographic address 102 and/or the contract identifier 28 to its corresponding entry 180, entry block 182, and/or directory block 184 within the blockchain data layer 72. Whenever the data layer server 74 generates the entry 180, entry block 182, and/or directory block 184, the data layer server 74 may add an entry to the database 220 of data layer records. Over time, then, the database 220 of data layer tracks a comprehensive historical repository of information that is electronically associated with its corresponding contract identifier 28. The data layer server 74 may then read or retrieve the entry 180, entry block 182, and/or directory block 184 containing or corresponding to the contract identifier 28.

[0190] Exemplary embodiments thus present the entity-specific cryptocurrency 60. Any entity 32 may create its own private blockchain 24, establish its entity-specific tokens 62, and define or offer digital contracts 20. The entity-specific tokens 62 may or may not have the value 94. The tradeable token 66, for example, may have a market value based on supply and/or demand, thus allowing or causing the value 94 of the tradeable token 66 to rise/fall or to increase/decrease, based on market forces. The credit token 64, however, may have a constant price or value, perhaps set by the entity 32. The entity-specific tokens 62 may be associated with the contract identifier 28, thus allowing a faster and simpler accounting scheme for machine executable contractual terms.

[0191] Exemplary embodiments may thus create coinage on top of coinage. The hierarchical scheme (explained with reference to FIG. 29) allows the private entity 32 to establish its private cryptocurrency 60 hierarchically above the traditional BITCOIN®, ETHEREUM®, or RIPPLE® coinage. The entity's private data 36 remains private, but the transaction records 100 may be publicly documented or proved via the traditional BITCOIN®, ETHEREUM®, or RIPPLE® environment. The private entity 32, in other words, need to worry about or concern itself with public publication. The private entity 32 need only subscribe (e.g., pay for write access) to the blockchain data layer 72. The digital contract 20 may also be offered, executed, and documented by the transaction records 100.

[0192] FIG. 36 illustrates a public entity 230, according to exemplary embodiments. Here exemplary embodiments may be applied to public data 232 generated by the public

entity 230. The public entity 230 may be a city, state, or federal governmental agency, but the public entity 230 may also be a contractor, non-governmental organization, or other actor that acts on behalf of the governmental agency. The public entity 230 operates a public server 234 and applies its software application 236 to its public data 232 to generate its governmental blockchain 238. The public entity 230 may further generate/issue its cryptocurrency 240 and offer digital contracts 20 for governmental, public services. The data layer server 74 receives the governmental blockchain 238 and generates the blockchain data layer 72. The data layer server 74 may then generate the public blockchain 76 representing any data records 70 representing the public data 232 and/or the cryptocurrency 240.

[0193] FIGS. 37-40 further illustrate contractual execution, according to exemplary embodiments. When the contract server 42 (such as the data layer server 74) receives the blockchain 24, exemplary embodiments inspect the blockchain 24 to identify the contract identifier 28 and/or the contractual parameters 30. The contract identifier 28 and/or the contractual parameters 30 may be contained within the block 26 of data within the blockchain 24. The contract identifier 28 and/or the contractual parameters 30 may be additionally or alternatively be metadata contained within the block 26 of data, and/or the contract identifier 28 and/or the contractual parameters 30 may be a data, data field, and/or a file attachment. The contract identifier 28 and/or the contractual parameters 30 may be information or data specified by the blockchain 24 and/or by a packet header or body. Regardless, once the contract identifier 28 and/or the contractual parameters 30 are determined, exemplary embodiments may consult the electronic database 44 of contracts.

[0194] FIG. 38 illustrates the database 44 of contracts. While the database 44 of contracts may have any logical structure, a relational database is perhaps easiest to understand. FIG. 38 thus illustrates the database 44 of contracts as an electronic table 250 that maps, converts, or translates the contract identifier 28 and/or the contractual parameters 30 to their corresponding network resource(s) 50. The database 44 of contracts may thus be preconfigured or preloaded with entries that assign or associate different contract identifiers 28 and/or contractual parameters 30 to their corresponding network resource 50 that provides, processes, and/or executes the corresponding digital contract 20. As the data layer server 74 receives any blockchain 24, the data layer server 74 may inspect the blockchain 24 for the contract identifier 28 and/or the contractual parameters 30. The data layer server 74 may then query the database 44 of contracts for the contract identifier 28 and/or the contractual parameters 30 to identify the computer file 46, server 254, virtual machine 256, Internet protocol address 258, or other network resource 50 that is responsible for executing the digital contract 20. The database 44 of contracts may optionally contain entries that relate hashed values of the contract identifier 28 and/or the contractual parameters 30. Regardless, once the network resource 50 is identified, the data layer server 74 may direct, assign, or outsource the contractual information 30 to the network resource 50 for processing.

[0195] FIG. 39 illustrates a simple example. Here the contract identifier 28 maps to a filename 260 that is associated with, or that represents, the computer file 46 that contains the programming language representing the digital contract 20. So, once the filename 260 is determined, the

data layer server 74 may locally retrieve and execute the computer file 46 that corresponds to, or is associated with, the filename 260. The data layer server 74 may then execute the computer file 46, perhaps based on parameters defined or described by the contractual parameters 30 (such as party names, parameters associated with their respective performance obligations and terms, and consideration). Optionally, the data layer server 74 may retrieve the computer file 46 (perhaps via the communications network 146 illustrated by FIGS. 22-23) from a remote server, database, or other device. Regardless, as the computer file 46 is executed, the data layer server 74 may generate the data records 70 in the blockchain data layer 72 describing the execution of the computer file 46. For example, the data records 70 may sequentially and/or serially track the execution of the computer file 46, perhaps logging or documenting periodic or random updates as the computer file 46 executes, perhaps along with timestamps toward completion. The data records 70 may also log or document a final step or outcome of the programming language representing the digital contract 20. Again, then, the blockchain 24 only referenced the digital contract 20 (using the contract identifier 28 and/or the contractual parameters 30). The actual execution of the digital contract 20 may be offloaded or outsourced to the data layer server 74.

[0196] FIG. 40 illustrates another example. Here the data layer server 74 may only manage the execution of the digital contract 20 referenced by the contract identifier 28 and/or the contractual parameters 30. That is, the data layer server 74 may outsource the execution of the digital contract 20 to a vendor or supplier as a subcontractor process. Again, when the data layer server 74 receives the blockchain 24, the data layer server 74 inspects the blockchain 24 to identify the contract identifier 28 and/or the contractual parameters 30. The data layer server 74 may then consult the database 44 of contracts. Here, though, the database 44 of contracts has entries that map or relate the contract identifier 28 to a remote server 262 that executes the digital contract 20 as a cloud-based service (perhaps as a software-as-a-service or SAAS). The database 44 of contracts may thus associate the contract identifier 28 to the Internet protocol address 258 representing the remote server 262 that executes the digital contract 20. The database 44 of contracts may additionally or alternatively associate the contract identifier 28 to a uniform resource locator (or "URL") 264 representing the remote server 262 that executes the digital contract 20. Regardless, once the remote server 262 is determined, the data layer server 74 may retrieve and send a service request 266 to the remote server 262 (via the Internet protocol address 258 and/or the URL 264 representing the remote server 262). The service request 266 specifies the contract identifier 28 and requests an execution of the corresponding digital contract 20. The service request 266 may also specify the contractual parameters 30. When the remote server 262 (perhaps operated on behalf of a third party) receives the service request 266, the remote server 262 applies the parameters defined or described by the contractual parameters 30 to the programming code (such as the computer file 46) representing the digital contract 20. Once the digital contract 20 is executed, the remote server 262 may then send a service response 268 back to the data layer server 74, and the service response 268 comprises data or information describing an outcome of the digital contract 20 (such as consideration, payment, or performance terms).

[0197] The data layer server 74 may generate the data records 70 in the blockchain data layer 72. For example, the data records 70 may document the date and time that the service request 266 was sent to the remote server 262. Moreover, as the remote server 262 provides the digital contract 20 as a service, the remote server 262 may send periodic or random service updates 270 as the service is provided along with timestamps toward completion. The data layer server 74 may thus generate the data records 70 describing the service updates 270 received from the remote server 262. The data layer server 74 may also generate the data records 70 describing the service response 268 sent from the remote server 262 describing an outcome of the digital contract 20.

[0198] FIGS. 41-42 illustrate virtual execution, according to exemplary embodiments. Here the data layer server 74 may outsource or subcontract the execution of the digital contract 20 to a virtual machine (or “VM”) 280. For example, the data layer server 74 may implement different virtual machines 190, with each virtual machine 190 processing and/or executing a particular digital contract 20, perhaps as a software service. The data layer server 74 may provide virtual computing and/or virtual hardware resources to client devices, thus lending or sharing its hardware, computing, and programming resources. The data layer server 74 may thus operate or function as a virtual, remote resource for providing contractual execution as software services. Suppose, for example, that the data layer server 74 implements four (4) virtual machines 280a-d. In practice, though, the data layer server 74 may implement any number or instantiations of different virtual machines 280 and/or digital contracts 20, depending on complexity and resources. Moreover, as a further simplification, assume that each virtual machine 280a-d executes a different corresponding digital contract 20a-d. So, when the data layer server 74 receives the blockchain 24, the data layer server 74 may inspect the blockchain 24 for each contract identifier 28a-d and/or the corresponding contractual information 28a-d and consult the database 44 of contracts.

[0199] FIG. 42 further illustrates the database 44 of contracts. Here the database 44 of contracts may include entries that map the contract identifier 28 to the corresponding virtual machine 280. The database 44 of contracts may thus be preconfigured or preloaded with entries that assign or associate each virtual machine 280 to its corresponding contract identifier 28. Once the virtual machine 280 is identified, the data layer server 74 may then coordinate and/or manage the execution of the corresponding digital contract 20, perhaps based on the contract information 30. Suppose, for example, that the data layer application 154 has programming or code that functions or performs as a query handler. The data layer application 154 inspects the blockchain 24 for the contract identifier 28 and queries the database 44 of contracts (as above explained). The data layer application 154 thus identifies and/or retrieves the corresponding virtual machine 280. Exemplary embodiments may thus determine whether contract identifier 28 matches or satisfies any of the entries specified by the database 44 of contracts. FIG. 42 illustrates entries that map the contract identifier 28 to its corresponding virtual machine 280 (e.g., an address, processor core, identifier, or other indicator).

[0200] The digital contract 20 may then be executed. For example, once the contract identifier 28 and the virtual machine 280 are determined, the virtual machine 280 may

then call, retrieve, and/or execute the computer file 46 that provides the digital contract 20 as a virtual service or process. FIG. 42 illustrates the computer file 46 locally stored and executed by the data layer server 74, but the computer file 46 may be remotely stored, retrieved, and/or executed. Regardless, the virtual machine 280 may be instructed to retrieve, execute, and/or apply the computer file 46, perhaps based on the contractual information 30.

[0201] FIG. 42 also illustrates software services. Here the database 44 of contracts may include entries that map the contract identifier 28 to a corresponding software service provided by the virtual machine 280. Exemplary embodiments, in other words, may relate the contract identifier 28 to a service identifier 282. The service identifier 282 is any alphanumeric combination, data, or hash value that uniquely identifies a software service 284 provided by the virtual machine 280. Once the contract identifier 28, the software service 284, and/or the virtual machine 280 are determined, the virtual machine 280 may then provide the software service 284. The software service 284 may execute the digital contract 20, perhaps based on the contractual information 30.

[0202] FIG. 43 illustrates cryptographic affinities, according to exemplary embodiments. Here the data layer server 74 may create or generate a cryptographic affinity 290 describing contractual execution. This disclosure above explained how the data layer server 74 may generate the data records 70 in the blockchain data layer 72. This disclosure also above explained how the data records 70 may document execution of the digital contract 20. Here, then, the cryptographic affinity 290 may uniquely identify the digital contract 20 executed by the virtual machine 280. For example, once the contract identifier 28 and the virtual machine 280 are determined (as above explained), the hashing algorithm 148 may generate a unique hash value 150. That is, the hashing algorithm 148 may hash the contract identifier 28 with a virtual machine (“VM”) identifier 292 to generate the cryptographic affinity 290. The virtual machine identifier 292 is any alphanumeric combination, data, or hash value that uniquely identifies the virtual machine 280. The cryptographic affinity 290 may then be documented by the data records 70 in the blockchain data layer 72, thus evidencing the execution of the digital contract 20. Indeed, the cryptographic affinity 290 may be published via the public blockchain 76 as the cryptographic proof 80, thus further publicly evidencing the execution of the digital contract 20.

[0203] FIG. 44 illustrates virtual assignments based on the blockchain data layer 72, according to exemplary embodiments. As this disclosure previously explained, exemplary embodiments may generate the data records 70 representing the blockchain data layer 72 (such as the entries 180, the entry blocks 182, and/or the directory blocks 184 explained with reference to FIGS. 25-27). Exemplary embodiments may thus assign the blockchain 24 and/or the virtual machine 280 that executes the digital contract 20, based on the number of the entries 180, the entry blocks 182, and/or the directory blocks 184 generated within the blockchain data layer 72. For example, as the data records 70 are generated, the data layer server 74 may determine a rate 290 of generation. That is, as the data records 70 are generated when or while executing the digital contract 20, exemplary embodiments may sum or count the entries 180, the entry blocks 182, and/or the directory blocks 184 that are generated over time (such as per second, per minute, or other

interval). Exemplary embodiments, for example, may call or initialize a counter having an initial value (such as zero). At an initial time (such as when the blockchain 24 is received or when the contract identifier 28 is determined), the counter commences or starts counting or summing the number of the entries 180, the entry blocks 182, and/or the directory blocks 184 (generated within the blockchain data layer 72) that are commonly associated with or reference the blockchain 24 (perhaps according to the chain ID 174) and/or the contract identifier 28. The counter stops counting or incrementing at a final time and exemplary embodiments determine or read the final value or count. Exemplary embodiments may then calculate the rate 290 of generation as the sum or count over time and consult or query the electronic database 44 of contracts for the rate 290 of generation. Exemplary embodiments may thus define entries that map or associate different rates 290 of generation and/or ranges to their corresponding contract identifier 28 and/or virtual machines 280. If the database 44 of contracts has an entry that matches or satisfies the rate 290 of generation, exemplary embodiments identify the corresponding virtual machine 280.

[0204] The rate 290 of generation may thus be a feedback mechanism. As the blockchain 24 is received, the data records 70 are requested, and/or the digital contract 20 is executed, the rate 290 of generation of the data records 70 may determine the virtual machine 280 that is assigned adequate capacity or bandwidth. One of the blockchains 24 and/or virtual machines 280, for example, may be reserved for digital contracts 20 having a heavy, disproportionate, or abnormally large rate 290 of generation. Another of the blockchains 24 and/or virtual machines 280 may be reserved for digital contracts 20 having a medium, intermediate, or historically average rate 290 of generation. Still another blockchain 24 and/or virtual machine 280 may be reserved for the digital contracts 20 having a light, low, or historically below average rate 290 of generation. The rate 290 of generation may thus be a gauge or measure of which blockchain 24, digital contract 20, and/or virtual machine 280 is assigned the resources.

[0205] Exemplary embodiments thus include a service environment. Exemplary embodiments may manage and/or execute many different digital contracts 20 offered by many different vendors or suppliers. Indeed, the data layer server 74 may manage or even execute the digital contracts 20 while also generating the blockchain data layer 72 as still another service. The data layer server 74 may thus act as a subcontractor or service provider, perhaps in a subscription or other compensation scheme. Any customer or client (such as the entity server 140 explained with reference to FIGS. 22-23) may thus send or forward its private blockchain 24 (generated from its private data 36) to the data layer server 74 for management or execution of any digital contract 20. The data layer server 74 may generate the data records 70 of the blockchain data layer 72 that document the management or execution of any digital contract 20. Moreover, the data layer server 74 may publicly publish the cryptographic proof 80 within the public blockchain 76, thus further documenting immutable evidence of the management or execution of any digital contract 20. Indeed, the entity server 140 may also generate the blocks 26 of data within the private blockchain 24 that also document the date and time that the management or execution of any digital contract 20 was sent/requested. The entity server 140 may then pay or reward the data layer server 74 in exchange for the digital

contract 20 and/or the data records 70 in the blockchain data layer 72 (such as granting its cryptocurrency 60 and 120, as explained with reference to FIG. 19).

[0206] The data layer server 74 may thus serve many blockchains 24 requesting many different contractual services. The financial institution 34, for example, may send or forward its private blockchain 36a (as illustrated with reference to FIGS. 20-21) to the data layer server 74 for application or execution of any digital contract 20 (by specifying the contract identifier 20, as above explained). The retailer 122 may similarly send or forward its private blockchain 36b to the data layer server 74 for application or execution of any digital contract 20. The online website 124 may also send or forward its private blockchain 36c to the data layer server 74 for application or execution of any digital contract 20. The data layer server 74 may generate the data records 70 of the blockchain data layer 72 that document the management and/or execution of any digital contract 20, and the data layer server 74 may publicly publish each cryptographic proof 80 within the public blockchain 76, thus further documenting immutable evidence of the management and/or execution of any digital contract 20. The entity 32 may then pay or reward the data layer server 74 via their respective cryptocurrency 60 and 120.

[0207] Exemplary embodiments thus only need to identify the digital contract 20. The contract identifier 28 and the contractual parameters 30 need only be informational content in the private blockchain 24. The contract identifier 28 is any digital identifying information that uniquely identifies or references the digital contract 20. The contract identifier 28 may be an alphanumeric combination that uniquely identifies a vendor and/or version of the digital contract 20 and/or a processor or executioner of the digital contract 20. The contract identifier 28 may be expressed as a unique hash value that is included within, or specified by, the private blockchain 24. Similarly, the contractual parameters 30 may identify the parties to the digital contract 20, their respective performance obligations and terms, and consideration.

[0208] FIGS. 45-51 illustrate an architectural scheme, according to exemplary embodiments. This disclosure above explained that the data layer server 74 may only manage the execution of the digital contract 20. The implementation and/or actual execution of the digital contract 20 may thus be separate from the data layer server 74 that generates the blockchain data layer 72. FIG. 45, for example, illustrates the data layer server 74 communicating via the communications network 142 with the remote server 262. The data layer server 74 generates the blockchain data layer 72, and the remote server 262 executes at least some portion of the digital contract 20. The remote server 262 may thus have a hardware processor 300 (e.g., “μP”), application specific integrated circuit (ASIC), or other component that executes a contract application 302 stored in a local memory device 304. The remote server 262 has a network interface to the communications network 142, thus allowing two-way, bidirectional communication with the data layer server 74. The contract application 302 includes instructions, code, and/or programs that cause the remote server 262 to perform operations, such as executing at least some portion of the digital contract 20.

[0209] FIG. 46 illustrates a request mechanism. The data layer application 154, for example, identifies the contract identifier(s) 28 and/or the contractual parameters 30 associated with or representing the digital contract 20. The

contract identifier(s) 28 and/or the contractual parameters 30 may be sent to the data layer server 74 as an input (such as from the entity server 140, as explained with reference to FIGS. 22-24), or the contract identifiers 28 and/or the contractual parameters 30 may be contained as information in the private blockchain 24. Regardless, the data layer server 74 may then identify the network address, IP address, URL, or other nomenclature representing the remote server 262 that executes at least some portion of the digital contract 20 (perhaps via the database 44 of contracts, as earlier explained). The data layer server 74 sends the service request 266 to the remote server 262, and the service request 266 may include or specify the contract identifier 28 and/or the contractual parameters 30. When the remote server 262 receives the service request 266, the remote server 262 applies the contractual parameters 30 to the portion of the digital contract 20 and generates a contractual result 306. The remote server 262 may then send the service response 268 back to the data layer server 74, and the service response 268 may comprise the contractual result 306.

[0210] Exemplary embodiments may thus exchange inputs and outputs. When the data layer server 74 sends the service request 266 to the remote server 262, the service request 266 may include or specify one or more of the contract identifiers 28 and/or the contractual parameters 30. Suppose, for example, that the contract identifiers 28 and/or the contractual parameters 30 are represented as hash values. The hash values may be identified from, or specified by, the private blockchain 24. The hash values may additionally or alternatively be generated by the data layer application 154 (such as by calling, invoking, or executing the hashing algorithm 148, as above explained). Regardless, the service request 266 may thus include or specify the hash values representing the contract identifiers 28 and/or the contractual parameters 30. When the remote server 262 receives the service request 266, the contract application 302 may use or accept the hash values as inputs to generate the contractual result 306 as an output. The contract application 302 may further encrypt the contractual result 306 (such as calling, invoking, or executing the hashing algorithm 148) to generate another hash value representing the contractual result 306.

[0211] Exemplary embodiments provide contractual proofs. When the data layer server 74 sends the service request 266 to the remote server 262, the data records 70 may document the service request 266 as one of the cryptographic proofs 80. When the data layer server 74 receives the service response 268, the data records 70 document that receipt and the contractual result 306 as another one of the cryptographic proofs 80. The data records 70 thus prove that at least the portion of the digital contract 20 was outsourced to a vendor or supplier as a subcontractor process or assignment. The data records 70 also prove that at least the portion of the digital contract 20 was executed to provide the contractual result 306. The data layer server 74 may then compare the contractual result 306 (such as its hash value) to a predefined or expect value. If the contractual result 306 matches or equals the predefined or expect value, then the data layer application 154 may be programmed or coded to infer that the contract successfully executed and/or the vendor or supplier performed as obligated. However, if the contractual result 306 fails to match or equal the predefined or expect value, then the data layer application 154 may be

programmed or coded to infer that the contract is not satisfied and/or the vendor or supplier failed to perform as obligated.

[0212] FIG. 47 illustrates a layered contractual process. Here the digital contract 20 may have different or individual components, portions, or sub-parts that cumulatively combine to produce the contractual result 306. The different components, portions, or sub-parts may be software modules 310 that can be separately executed to generate the overall or final contractual result 306. A simple digital contract 20, for example, may only have a few or several software subroutines or modules 310, while a complex or complicated digital contract 20 may have many or hundreds of different software subroutines or modules 310. As the reader likely understands, such a complicated software structure is too difficult to illustrate. For simplicity, then, FIG. 47 illustrates the digital contract 20 having four (4) software modules 310a-d. The entire contract application 302, in other words, may have four (4) different application layers 312a-d. Each componentry module 310a-d or layer 312a-d may have its own corresponding contract identifier 28a-d. When the remote server 262 receives the service request 266, exemplary embodiments may then feed the contractual parameters 30 as inputs 314a-d to the software modules 310a-d. Each different software module 310 may thus generate its respective or corresponding output 316a-d, which may be combined or processed to generate the overall or final contractual result 306.

[0213] FIG. 48 illustrates hierarchical execution. Here the different software modules 310 may be serially or sequentially executed to generate the overall or final contractual result 306. For example, the software module 310a may accept at least some of the contractual parameters 30 as the input 314a, execute its respective programming code, and generate its corresponding output 316a. Here, though, the output 316a may then be routed or sent to the software module 310b (illustrated as the application layer 312b) as its input 314b. Its respective programming code is then executed to generate its corresponding output 316b, based on the output 316a generated by or received from the software module 310a. Similarly, software module 310c accepts the output 316b and generates output 316c, which is received by software module 310d as input 314d and used to generate the output 316d. While exemplary embodiments may continue processing the outputs 316a-d to generate any desired outcome, for simplicity FIG. 40 illustrates the output 316d as the final contractual result 306. Exemplary embodiments may thus use the software modules 310a-d as feedback mechanisms to monitor or even enforce contractual rule-based obligations defined or specified by the digital contract 20.

[0214] FIG. 49 illustrates the blockchain data layer 72. Here the blockchain data layer 72 may document the processing and/or execution of each software module 310a-d, its respective input(s) 314a-d, its respective output(s) 316a-d, and perhaps a corresponding timestamp (not shown for simplicity). The data records 70 may further document or record the corresponding contract identifier 28a-d and/or the chain identifier 174. The data layer server 74 may thus receive the service updates 270 (via the communications network 142) as each software module 310a-d performs or executes its corresponding contractual service. The data layer server 74 may then generate the data records 70 in the blockchain data layer 72, thus documenting each software

component's contribution toward the overall or final contractual result 306. The data records 70 may also be hashed to generate the cryptographic proofs 80, as above explained.

[0215] FIG. 50 also illustrates contractual execution. Here, though, the different software modules 310 may be executed by different devices. Suppose, for example, that the remote server 262a locally stores and executes the software module 310a, while the remote server 262b locally stores and executes the software module 310b. Suppose also that the remote server 262c locally stores and executes the software module 310c and the remote server 262d locally stores and executes the software module 310d. Exemplary embodiments may thus source or subcontract the different portions of the digital contract 20 to different machines for execution. The remote server 262a, for example, may specialize in the software module 310a. The remote server 262a may thus accept the service request 266 from clients, execute the software module 310a, and return send the service response 268 (as explained with reference to FIG. 46). The remote server 262a may also send the service update(s) 270 to the data layer server 74, thus allowing the blockchain data layer 72 to document the contractual service provided by the software module 310a. The remote servers 262b-d may similarly specialize in the software modules 310b-d to provide their respective contractual services.

[0216] FIG. 51 illustrates an overall architectural scheme. As the reader may envision, there may be hundreds, thousands, millions, or even billions of contractual relationships between many different parties. As smart, digital contracts grow in acceptance and usage, the blockchain data layer 72 is expected to exponentially grow, thus requiring ever-increasing hardware and software resources. In plain words, there may be many data layer servers 74 generating the data records 70 in the blockchain data layer 72. While there may be hundreds or even thousands of data layer servers 74, FIG. 51 simply illustrates four (4) data layer servers 74a-d that cooperate to generate the blockchain data layer 72. As the processing load increases or grows (such as according to the rate 290 of generation, as above explained), the number of data layer servers 74 may also grow.

[0217] The blockchain data layer 72 may thus be separate from an implementation and execution of the digital contract 20. The data layer servers 74, in other words, may be separately networked and/or addressed from the remote servers 262 providing the contractual services representing the software modules 310 of the digital contract 20. Any of the data layer servers 74 may send data or information as inputs to any one of the remote servers 262, and the corresponding software module 310 performs its contractual service and sends its output 316 back to the blockchain data layer 72 (perhaps via the service request 266, the service response 268, and the service update 270 as earlier explained and illustrated). Some of the remote servers 262 may provide virtual services, such as a virtual machine (as above explained) that executes any of the software modules 310.

[0218] FIG. 52 illustrates compliance scheme, according to exemplary embodiments. As the reader may understand, some smart, digital contracts have jurisdictional requirements. For example, the digital contract 20 may have programming code that requires an execution or processing in a particular region or country. That is, the digital contract 20 may have contractual rules and/or provisions that must be enforced in the United States, the European Union, or the Isle of Man. Components or portions of the digital contract

20 may require execution or location in the Cayman Islands, Idaho, or Hong Kong. The digital contract 20, in other words, may have a geographic parameter 320. The geographic parameter 320 may be a locational requirement, restriction, or preference for at least some portion of the digital contract 20. The geographic parameter 320 can be any data, information, field, metadata, or code for enforcing the locational requirement, restriction, or preference. Indeed, the geographic parameter 320 may even be finely expressed or defined as global positioning system ("GPS") information or coordinates at which at least some portion of the digital contract 20 must be processed or executed.

[0219] The geographic parameter 320 may be an input value. As FIG. 52 illustrates, the geographic parameter 320 may be read or received via the private blockchain 24 (perhaps along with the contract identifier 28 and/or the contractual parameter 30). The data layer server 74, in other words, may identify the geographic parameter 320 as data, information, or a hash value contained within the block 26 of data. However, the geographic parameter 320 may additionally or alternatively be received and/or identified within a header of body/payload of a packet 322 of data (packetized according to the Internet Protocol, just as the contract identifier 28 and/or the contractual parameter 30 may be identified).

[0220] Regardless, once the geographic parameter 320 is determined, exemplary embodiments may again consult the database 44 of contracts. The database 44 of contracts may have entries that electronically associate the contract identifier(s) 28 and/or the contractual parameter(s) 30 to the geographic parameter 320. The data layer application 154 may instruct the data layer server 74 to query the database 44 of contracts for either, any, or all of the contract identifiers 28, the contractual parameters 30, and/or the geographic parameters 320 to identify and/or retrieve the corresponding database entries. As a simple example, suppose a file component of the digital contract 20 must be processed in a particular geographic region (such as the British Virgin Islands or Canada). The corresponding contract identifier 28, in other words, may be electronically associated with a particular geographic region, as defined by a tabular entry in the database 44 of contracts. Once the region is determined, the data layer server 74 may then route the contract identifier 28, the contractual parameter 30, and/or the geographic parameter 320 to the remote server 262 that is associated with, or even located within, the region. Exemplary embodiments, for example, may implement the service request 266, the service response 268, and the service update 270 (as earlier explained). The remote server 262 may then process or execute the digital contract 20 using the contract identifier 28 and/or the contractual parameter 30 (as this disclosure earlier explained).

[0221] Other examples explain the geographic parameter 320. Suppose that the contract identifier 28 and/or the contractual parameter 30 map(s) to a particular server, cluster of servers, and/or a particular virtual machine ("VM"). The data layer server 74 may then route the contract identifier 28, the contractual parameter 30, and/or the geographic parameter 320 to the remote server 262 that is associated with the cluster of servers and/or the virtual machine. The remote server 262 may then process or execute the digital contract 20 using the contract identifier 28 and/or the contractual parameter 30 (as this disclosure earlier explained). More likely, though, the contract identi-

fier 28 and/or the contractual parameter 30 will relate to a particular IP address or uniform resource locator (“URL”). The data layer server 74 may then route the contract identifier 28, the contractual parameter 30, and/or the geographic parameter 320 to the remote server 262 that is associated with the IP address or URL for processing (again, as this disclosure earlier explained).

[0222] Exemplary embodiments may thus implement contractual provisions. Some digital contracts 20 may require a particular server, perhaps implementing or hosting a particular website, network, authentication scheme, programming, or other geographic parameter 320. Some parties to the digital contract 20 may also require a particular server, perhaps as specified by the geographic parameter 320. Some digital contracts 20 may have compliance obligations, perhaps defined by a particular jurisdiction and expressed as the geographic parameter 320. Servers, webpages, networks and other resources may be dedicated to specific jurisdictions, as expressed by the geographic parameter 320.

[0223] FIGS. 53-59 illustrate a decisional architecture and scheme, according to exemplary embodiments. Even though the blockchain environment 22 enables an execution of the smart, digital contract 20, some digital contracts may be too complex and/or too cumbersome to implement on the blockchain 24. As this disclosure above explains, exemplary embodiments may thus put smaller contractual components of the digital contract 20 on any blockchain (such as the private blockchain 24 or the public blockchain 76), validate the contractual components (perhaps via the cryptographic proof 80), incorporate the cryptographic proof 80 into a larger component of the digital contract 20, and then validate the larger component.

[0224] Exemplary embodiments may further implement one or more decision tables 326. As the reader may understand, the decision table 326 may be used to implement at least a component of the digital contract 20. That is, the decision table 326 may represent one or more rules or logic conditions 328, one or more inputs 330, and one or more decisional outputs 332. The decision table 326 may thus be visually represented as a table having rows and columns. In simple words, once the input 330 is known, a processing or execution engine (such as the entity server 140 or other device) electronically maps or associates the input 330 to the appropriate rule or logic condition 328 and generates the decisional output 332. Exemplary embodiments may then log or record the decisional output 332, along with its corresponding the input 330, rule or logic condition 328, and a date/time stamp.

[0225] Exemplary embodiments may thus document any decision. In general, the smart, digital contract 20 is an agreement between parties/participants about services, products, and/or money. In order to make the decisional output 332, information is provided (such as the input 330) and the rule or logic condition 328 is executed. In an interactive process, each party/participant might contribute data to a single decision. In other words, the parties may exchange data to perform the decisional output 332. Exemplary embodiments may thus map each decisional output 332 (perhaps representing a decision model) to a decision taken by a single party. Each party, in other words, may communicatively exchange the result of its decision such that others can base their decisions on their decisional output 332, thus collaboratively executing the different components of the digital contract 20.

[0226] FIG. 54 illustrates an impartial, trusted intermediary. When any party or participant to the digital contract 20 acts or executes, exemplary embodiments may log or archive their respective action(s). For example, the data layer server 74 may be informed of any decision-making process. Suppose, for example, that the entity 32 (acting as a party to the digital contract 20) wishes to document or prove its contractual performance. That is, the entity server 140 sends its decisional output 332 (perhaps via the communications network 142 illustrated in FIGS. 22-23) to the data layer server 74 for documentation. The decisional output 332 may thus be read or received via the private blockchain 24 (perhaps along with the contract identifier 28 and/or the contractual parameter 30). The data layer server 74, in other words, may identify the decisional output 332, along with its corresponding input 330, its rule or logic condition 328, and the date/time stamp, as data, information, or hash values contained within the block 26 of data (as FIG. 53 illustrated). However, the decisional output 332 (and/or the contract identifier 28, the contractual parameter 30, the input 330, the rule or logic condition 328, and the date/time stamp) may additionally or alternatively be received and/or identified within a header of body/payload of the packet 322 of data (packetized according to the Internet Protocol).

[0227] Regardless, the data layer server 74 may then generate the data records 70 in the blockchain data layer 72, as this disclosure above explained. The data records 70 log or record the decisional output 332 (sent from the party participant), along with its corresponding input 330, the decision table 326, the rule or logic condition 328, and the date/time stamp of performance. The blockchain data layer 72, in other words, provides neutral, documentary evidence that the party executed its transactional portion of the smart, digital contract 20. Moreover, the blockchain data layer 72 may also add another layer of cryptographic hashing to generate the public blockchain 76 and the cryptographic proof 80. The blockchain data layer 72 thus may again act as the validation service 78 that validates the party performed its portion of the digital contract 20. Exemplary embodiments may thus be used as an audit trail to reconstruct the party’s decision-making process and who provided the input 330.

[0228] Exemplary embodiments may even document fine granularity. When the data layer server 74 receives the decisional output 332, the data or information may even identify or pinpoint the network resource 250. That is, when entity 32 (acting as a party to the digital contract 20) wishes to document or prove its contractual performance, the decisional output 332 may even include data or information identifying the particular server 254 or cluster or virtual machine 256 that generated the decisional output 332. Indeed, the data or information may even identify or pinpoint the particular IP address or uniform resource locator (“URL”). The data records 70 may thus document the machine, manufacturer, model, and/or chassis hardware inventory that performed the portion of the digital contract 20.

[0229] FIG. 55 illustrates contractual management. Here again the data layer server 74 may manage the execution of the digital contract 20. When any party, participant, or subcontractor performs a portion or component of the digital contract 20, the data layer server 74 may coordinate and validate the contractual components. Suppose again that the data layer server 74 receives the contract identifier 28 and/or

the contractual parameters 30 (as earlier explained). The contract identifier 28 may represent a single, large digital contract 20. The contract identifier 28, however, may represent only a single or a few contractual components of the digital contract 20. The contract identifier 28, in other words, may map or relate to a sequence or series of one or more table identifiers 334. Each table identifier 334 may be an alphanumeric combination or a unique hash value. Regardless, each table identifier 334 uniquely identifies the corresponding decision table 326 that decides a componentry portion of the digital contract 20. When the data layer server 74 receives the one or more contract identifiers 28, the data layer server 74 may then consult the database 44 of contracts.

[0230] FIG. 56 further illustrates the database 44 of contracts. Here the database 44 of contracts may have additional entries that map or relate the contract identifier 28 to the table identifier 334 and/or to the network resource 250 that executes the corresponding componentry portion of the digital contract 20 (perhaps again as a cloud-based service). The contract identifier 28, in other words, may map or relate to a sequence or series of one or more table identifiers 334. Each table identifier 334 may be an alphanumeric combination or a unique hash value. Regardless, each table identifier 334 uniquely identifies the corresponding decision table 326 that decides a componentry portion of the digital contract 20. When the data layer server 74 receives the one or more contract identifiers 28, the data layer server 74 may then consult the database 44 of contracts to determine any corresponding entry (as this disclosure above explains).

[0231] FIG. 57 illustrates outsourcing. Once the network resource 50 is determined (recall that the network resource 50 may execute the corresponding componentry portion of the digital contract 20), the data layer server 74 may utilize the request mechanism. Suppose, for example, that the database 44 of contracts identifies the remote server 262 as the network resource 50. The data layer server 74 may thus instruct the remote server 262 to execute the corresponding decision table 326. The data layer server 74, for example, sends the service request 266 (as earlier explained), and the service request 266 may specify the table identifier 334 and/or the input 330 as the contractual parameters 30. When the remote server 262 receives the service request 266, the remote server 262 applies the input 330 to the decision table 326 representing the digital contract 20. Once the decisional output 332 is determined, the remote server 262 may then send the service response 268 back to the data layer server 74, and the service response 268 comprises data or information describing the decisional output 332. The data layer server 74 may generate the data records 70 in the blockchain data layer 72 that document the service request 266 and the service response 268, perhaps including any service updates 270 as the decision table 326 is executed.

[0232] FIG. 58 illustrates contractual participation. Here the data layer server 74 may execute at least a componentry portion of the digital contract 20. That is, the data layer server 74 may locally store and/or access one or more of the decision tables 326 representing the digital contract 20. When the data layer server 74 receives the contract identifier 28 and/or the contractual parameters 30 (as earlier explained), the data layer server 74 may consult the database 44 of contracts. Here, though, the database 44 of contracts has one or more entries that map or relate the contract identifier 28 to the virtual machine 280 that executes the

decision table 326. The database 44 of contracts may thus electronically associate the contract identifier 28 to the table identifier(s) 334 and the virtual machine(s) 280 that locally execute the decision table(s) 326. The decision table 326 may thus have virtual assignments. Once the virtual machine 280 and/or the decision table 326 is determined, the virtual machine 280 is requested or instructed to apply the input 330 to the corresponding decision table 326 to generate the decisional output 332. The data layer server 74 may then generate the data records 70 in the blockchain data layer 72 that document the local contractual performance (as earlier explained).

[0233] As FIG. 58 illustrates, feedback may be used. Exemplary embodiments may assign the virtual machine 280 based on the data records 70 in the blockchain data layer 72. That is, as the decision table 326 consumes more and more of the data records 70 (e.g., the number of the entries 180, the entry blocks 182, and/or the directory blocks 184 generated within the blockchain data layer 72, as earlier explained), the rate 290 of generation may be used as a feedback mechanism (as this disclosure earlier explained). Highly used or called decision tables 326, in other words, may be assigned to virtual machines 280 having greater capacity or bandwidth. The database 44 of contracts may thus define entries that map or associate different rates 290 of generation and/or ranges to their corresponding table identifier 334 and/or virtual machines 280. If the database 44 of contracts has an entry that matches or satisfies the rate 290 of generation, exemplary embodiments identify the corresponding virtual machine 280. Some virtual machines 280, for example, may be reserved for decision tables 326 having a heavy, disproportionate, or abnormally large rate 290 of generation. Other virtual machines 280 may be reserved for decision tables 326 having intermediate and low rates 290 of generation. The rate 290 of generation may thus be a gauge or measure of which virtual machine 280 is assigned the decision table 326.

[0234] Exemplary embodiments thus include a service environment. Exemplary embodiments may manage and/or execute many different decision tables 326 offered by many different vendors or suppliers. Indeed, the data layer server 74 may manage or even execute the digital contracts 20 while also generating the blockchain data layer 72 as still another service. The data layer server 74 may thus act as a subcontractor or service provider, perhaps in a subscription or other compensation scheme. Any customer or client may thus send or forward its input 330 and/or its decisional output 332 to the data layer server 74 for management or execution of any digital contract 20. The data layer server 74 may generate the data records 70 of the blockchain data layer 72 that document the management or execution of any portion of component of the digital contract 20. Moreover, the data layer server 74 may publicly publish the cryptographic proof 80 within the public blockchain 76, thus further documenting immutable evidence of the management or execution of any digital contract 20. Any party, participant, or vendor/subcontractor may then pay or reward the data layer server 74 (such as granting its cryptocurrency 60 and 120, as explained with reference to FIG. 19).

[0235] The data layer server 74 may thus provide contractual services. The financial institution 34, for example, may send or forward its input 330 and/or its decisional output 332 to the data layer server 74 for contractual documentation. Similarly, the retailer 122, the online website 124, and the

manufacturer 128 may also send its input 330 and/or its decisional output 332 to the data layer server 74 for contractual documentation. The data layer server 74 may generate the data records 70 of the blockchain data layer 72 that document the management and/or execution of any decision table 326 representing any portion of the digital contract 20. The data layer server 74 may also publicly publish each cryptographic proof 80 within the public blockchain 76, thus further documenting immutable evidence of the management and/or execution of any digital contract 20. The data layer server 74 may be paid or rewarded via their respective cryptocoinage 60 and 120.

[0236] Exemplary embodiments may thus create factored decision tables driven by a table engine. Smart, digital contracts are notoriously dangerous. Decision tables are significantly easier to verify and validate. However, decision tables may be large and perhaps cannot be placed on a blockchain. Exemplary embodiments may thus put smaller contractual components of the digital contract 20 on any blockchain (such as the private blockchain 24 or the public blockchain 76), validate the contractual components (perhaps via the cryptographic proof 80), incorporate the cryptographic proof 80 into a larger component of the digital contract 20, and then validate the larger component.

[0237] Exemplary embodiments thus may separate the blockchain data layer data 72 from contractual execution. The data layer server 74 (generating the blockchain data layer data 72) may thus accept inputs from the servers (such as the remote server 262) executing any component of the digital contract 20. The servers (such as the remote server 262) executing any component of the digital contract 20 may also send data to the data layer server 74. The data layer server 74 may thus execute the decision table. The remote server 262 may additionally or alternatively execute the decision table when processing the digital contract 20. The decision table may thus be sent and/or received as an input/output. Even a virtual machine may access and use the decision table.

[0238] Exemplary embodiments thus establish the digital contract 20 as an identity. Because only the contract identifier 28 is needed, the digital contract 20 may be separated into various smaller components (such as the software modules 310 and/or layers 312, as above explained). Each software module 310 and/or layer 312 may have its own contract identifier 28. The digital contract 20 is thus transformed to an identity, which may be easily updated after software bugs are found and consensus is documented by stake holders. Exemplary embodiments thus provide an ability to repair bugs and to claw back or backup spurious results. The separation of the blockchain data layer data 72 thus isolates and protects the data records 70.

[0239] Exemplary embodiments thus describe a novel smart contract architecture to be run on blockchains. The digital contract 20, and/or its contractual components, may each have its own digital identity defined within the blockchain data layer data 72. The contract identifier 28, in other words, may uniquely identify a version, thus allowing stakeholders (using their digital identities) to approve updates to respond to changes in business, to approve bug resolution, and to accommodate new participants in the digital contract 20, without having to dissolve the original version and without redeploying or requiring the blockchain to be reversed and modified to avoid an incorrect, improper, or unacceptable result by perhaps a majority of users. As the

reader may understand, modifying a blockchain to resolve an issue involves many more stakeholders with an interest in the blockchain but having no interest in the smart contract. This has been a problem with conventional blockchain architectures.

[0240] Exemplary embodiments may separate the blockchain data layer data 72 from the rules engine architecture that executes the digital contract 20. Exemplary embodiments allow for light weight, secure, and extendible digital identity. Digital identity can be applied to implementation of the virtual machine that runs the digital contract 20. Digital identity can be applied to any smart contract and/or to any stakeholder(s). Stakeholders may thus be paid (perhaps via the cryptocurrencies as explained with reference to FIGS. 13 & 15-21) for who they are, such as to a particular blockchain address, meaning if a stakeholder's address is compromised, then the stakeholder can update the address without having to modify the digital contract 20. This virtual address modification is similar to the real world for when a business moves from one geographic location to another, the business does not invalidate all its contracts. In the real world, the business merely informs parties of its new physical address and contact information. Exemplary embodiments allow management of the digital contract 20 in a flexible fashion, similar to management of contracts in the real world, but with blockchain security and data integrity of the actual digital contract 20, automation of provisions in the digital contract 20, and cryptopayment support.

[0241] Exemplary embodiments are also scalable. Layers or modules 310 and 312 can be created in the digital contract 20 and/or in the private blockchain 24 or the public blockchain 76 for improved flexibility and management via hardware computers. The data records 70 in the blockchain data layer data 72 are safely separated from the servers that execute the digital contract 20. Contract servers (e.g., the contractual application layer) may perform a decentralized evaluation of digital contract 20, using the proper virtual machine and proper rules, and manage interests of majority or all stakeholders. Values of cryptotokens may be defined and/or distributed, but allowing greater scalability.

[0242] Exemplary embodiments provide numerous advantages. Because the contractual execution is separate from the blockchain data layer data 72, the results of the digital contract 20 are securely documented and may be exported to other contractual components or to other digital contracts. Exemplary embodiments may thus implement and offer multiple modules 310, layers 312, or instances of different contractual components that can exchange inputs and outputs to build a networking effect between different layers, modules, and smart contracts. A first server running a first application layer (and perhaps executing a first smart contract) can be entirely separate a second server running a second smart contract and a third server running a third smart contract. The blockchain data layer 72, though, exchanges and thus documents their respective inputs and outputs. The various servers may thus manage and/or share the same cryptotokens, or different entity tokens may be exchanged within each layer. Regardless, exemplary embodiments may coordinate exchanges of value for services performed. Great flexibility in defining the value of cryptotokens and the value into and out of smart contract.

[0243] Exemplary embodiments may also have jurisdictional advantages. Particular servers may be specific to particular jurisdictions and/or particular smart contracts. For

example, some application layers may cross jurisdictional servers with different compliances. As another example, suppose that one application layer may require qualified investors with full know your client (or “KYC”) compliance. Another application layer may be anonymous and/or allow all corners. Even if the blockchain data layer 72 has a small set of users/clients, large smart contracts may be managed, implemented, and/or documented.

[0244] The digital contract 20 may utilize conventional programming languages and/or decision tables. In particular, some programming languages and decision tables, like purely functional languages, may mathematically prove contractual algorithms. These mathematical proofs may yield much more secure smart contracts than conventional languages that run on today’s blockchains. Previously, smart contracts were often too big in size to execute on a blockchain. The separate blockchain data layer 72, though, allows scaling and implementing smart contracts “off chain.” The proof 80 of the digital contract 20, for example, is a hash value, perhaps in association with the contract identifier 28 and/or the chain identifier 174, as documented by the data records 70 in the blockchain data layer 72. The hash value of the proof 80, in other words, is a very small value (in relation to the size of the smart contract). The digital contract 20 may thus be provided to any or all parties and/or any or all stakeholders for validation of its terms, obligations, and performance. The cryptographic proof 80 thus verifies execution without stuffing large amounts of data onto the private blockchain 24 or the public blockchain 76.

[0245] Exemplary embodiments may use decision tables for smart contracts. Decision tables are well understood, perform well, and are verifiable relative to brute-force code writing. Simply put, custom programming code introduces many variables and software bugs are inevitable. Decision tables are also very amenable to domain-specific languages. As the reader may understand, domain-specific languages accept near-English statements as inputs and generate computer code as outputs. Subject matter experts may thus define the functionality of the digital contract 20, perhaps without relying on the skills of computer programmers (who may not fully understand the subject matter). Decision tables are thus approachable to subject matter experts and easily implemented. Decision tables may also be combined with other decision tables, which allows performance proven and validated functions may be incorporated into smart contracts for many objectives and outcomes. Decision tables may thus be mixed and matched as components to a composite digital contract 20, and a collection of decision tables representing the digital contract 20 may still be validated to ensure correct operation. Decision tables define much smaller numbers of programming paths through the software code representing the digital contract 20, which ensures that all contractual combinations may be enumerated and proper results can be expected for a range of values. On blockchains, though, decision tables may be big in size, so some decision tables may not be feasible as a smart contract on a conventional blockchain. But, because the blockchain data layer 74 is separate from the remote servers 262 executing the digital contract 20, the digital identity (e.g., the contract identifier 28) for the digital contract 20 (that allows the smart contract to exist off chain) provides the servers (each perhaps having its own identity) to certify execution of the digital contract 20. Exemplary embodiments may also define the mechanism for cryptotoken-based payments that incentivize the remote

server 262 to perform the digital contract 20 and to verify and validate the digital contract 20. Component and composite performance may be tracked, recorded, and proved. For example, if a virtual machine runs the digital contract 20 (as above explained), execution in the virtual environment can be tracked. Virtual machines may often have software bugs that affect an interpretation of the decision tables. The virtual machine may thus have its own digital identity, as defined by the database 44 of contracts (as above explained). Different versions of the virtual machine and/or the decision table may thus be mapped within the database 44 of contracts, thus allowing redirection after software bugs have been resolved. The database 44 of contracts, in other words, may be updated with entries that point to different versions for different parties and/or to corrected or improved versions.

[0246] Digital identities extend to engines and decision tables. The database 44 of contracts may map or point to servers, domains, decision tables, and their respective versions. The digital contract 20 (and/or its components, as represented by their respective contract identifiers 28) ensures execution, regardless of the environment. Because the blockchain data layer 72 documents all this component processing, the data records 70 may prove (via the cryptographic proof 80) that the correct contractual component was used, the correct decision table(s) was/were used, the correct virtual machine was used, and the correct input or output data was used. Verification may driven from the contractual components, the data components, and the hardware components at the correct time for the correct time period.

[0247] Another audit application example is provided. A software application may be a generic term for user-side software that reads from and/or writes to the Factom system. It could be software with a human interface, or could be completely automated. The Application is interested in the data organized by the Chains it needs.

[0248] Applications are possibly Distributed Applications (DApps) interacting with Factom to provide additional services. For example, one might imagine a trading engine that processes transactions very fast, with very accurate time-stamping. Such an Application may nonetheless stream transactions out into Factom chains to document and secure the ledger for the engine. Such a mechanism could provide real-time cryptographic proof of process, of reserves, and of communications.

[0249] Let us explore two separate applications that could have immediate demand in the current Bitcoin ecosystem.

[0250] Let us see how to implement a secure and distributed log platform. Log analysis is a complex task. Additionally, logs tend to be easily forgeable and also heterogeneous as they are produced by each system independently and stored in a variety of media (files, databases, cloud services etc.). With Factom and a few uniquely designed crypto-audit tools an entities log analysis can become safer, simpler, and much more powerful. Let’s see this with an example. Suppose a Bank (B), a Payment Provider (PP), and a Bitcoin company (BC) are interacting together as follows:

[0251] 1—The User goes to the BC website and wants to buy some bitcoins

[0252] 2—He asks for a quote, which is valid for 5 minutes

[0253] 3—Then he is redirected to the PP website

[0254] 4—Then the PP connects with the B platform so that the money of the user account is debited

[0255] 5—B notifies PP that the user account has been debited

[0256] 6—PP notifies BC

[0257] 7—BC sends the bitcoins to the user

[0258] This is the normal scenario for many fixed-rate Bitcoin exchanges globally. But assume now that for some reason the BC receives the payment notification 4 hours after the user transfers via the PP. Who is faulty? The User? The Bank? The Payment Provider? What if a similar payment problem happened for hundreds or thousands of payments over a period of days or weeks before the issue was identified and resolved? Who is “provably” liable for those losses/damages?

[0259] With current techniques a manual auditing of logs would be necessary and would probably require legal authorizations. With Factom and the right audit applications, it would be trivial to detect where the problem came from, and also make the changing of records impossible post-issue. Basically, every system (BB, PP, BC) will publish their relevant traces in the secure broadcast channel (Factom) in real time.

[0260] Here’s another example of how Factom will be useful for Bitcoin exchanges audits. The so-called “Proof of Solvency” method for conducting Bitcoin exchange audits is a growing and important trend. However, there are significant weaknesses to this approach only solved by having the Factom secure broadcast channel functioning properly.

[0261] In the Merkle tree approach for Solvency Proofs suggested by the Maxwell-Todd proposal, users must manually report that their balances (user’s leaf) have been correctly incorporated in the liability declaration of the Financial Institution (FI) (the Merkle hash of the FI’s database of user balances). The proposed solution works if enough users verify that their account was included in the tree, and in a case where their account is not included, it’s assumed that this instance would be reported. One potential risk with this process is that an exchange database owner could produce a hash that is not the true representation of the database at all; the exchange hashes an incomplete database which would reduce its apparent liabilities to customers, thereby making them appear solvent to a verifying party. Here are some scenarios where a fraudulent exchange could easily exclude accounts:

[0262] “Colluding Whales” Attack: There is evidence that large Bitcoin traders are operating on various exchanges and moving markets significantly. Such traders need to have capital reserves at the largest exchanges to quickly execute orders. Often, traders choose exchanges that they “trust”. In this way they can be assured that should a hack or liquidity issue arise, they have priority to get their money out first. In this case, the exchange and trader could collude to remove the whales account balance from the database before it’s hashed. An exchange’s top 10 whales could easily represent 5 to 20% of an exchanges liabilities, so colluding with just a few of them could have a significant impact.

[0263] “Site Manipulation” Attack: To date, each Proof of Solvency audit has reported (the hash tree) on the institution’s website. This gives no guarantee at all to users, since a malicious exchange could publish different states/balances to different groups of users, or

retroactively change the state. Thus it is fundamental to publish this data through Factom’s secure broadcast channel, and publish it frequently.

[0264] The second attack is obviously solved by using Factom, while the first is not so obvious. As this paper doesn’t focus on the mechanics of exchanges audits, we won’t delve in the nitty-gritty details. However, the basic concept is that by having frequent time-stamped copies of the exchanges database Merkle hash, one could detect the inclusion or exclusive of large balances before or after audits. Then, the auditor could simply look into those large inclusions or exclusions, manually. Remember, the trader will ultimately need to get his money on or off the exchange at some point, and that’ll show up in either the bank history or the Bitcoin transfer history.

[0265] There are established process for detecting such fraudulent tactics in the traditional audit industry; however, it all starts with having accurate, verifiable, immutable time-series of the information in question.

[0266] Other examples are provided of attacks on Factom. The reader, for example, may be familiar with a denial of service from spam. Since Factom is an open system, any user can put Entries into almost any Chain. Bitcoin has a similar phenomenon. In order for an Application to reject those transactions, the Application would first need to download and process them. A large number of bogus Entries could slow down the initial processing of the Application’s transactions. This threat is mitigated by an attacker needing to spend money (resources) to carry it out. This is similar to Adam Back’s Hashcash solution to email spam.

[0267] Audits are another useful tool against spam, if the application is willing to trade off security versus convenience. Auditors could post “ignore” lists on the same chain, or create their own audit chains with those lists. An auditor could use a profile chain to develop their reputation, which would also allow review by other auditors. If any auditor made a bad call, it would be easily verifiable and the record of it would be permanent. Some validity processing is gray, in the sense that opinions may vary. Solving that problem would be implementation specific.

[0268] Another example is a sybil attack of the DHT. Distributed Hash Tables in general are particularly susceptible to sybil attacks. An attacker could create many peers which make it difficult for honest nodes to communicate. In a simplistic DHT architecture, attackers can isolate a required piece of data from honest nodes. Sybil attacks have been observed on the BitTorrent network routing table. The paper “Real-World Sybil Attacks in BitTorrent Mainline DHT” detail these attacks. Fighting this type of attack is an active topic in academic research. One mitigation technique uses complex lookup techniques to find honest nodes among the sybils, studied in “Sybil-resistant DHT routing”. Some sybil mitigation techniques rely on a web-of-trust by adding a social network to the routing table, as explored in “A Sybil-proof one-hop DHT”. Factom will rely on the latest academic and open-source research in this topic to secure its DHT.

[0269] A dictionary attack is now discussed. In this case, the attacker runs through all the Chain Names deemed to be possible or desirable and creates their ChainIDs, and the hashes of those ChainIDs. Then they watch for someone trying to create those Chains. Now the attacker can front run on a match. Because on a match, they know the ChainID, so they can construct a proper, but malicious Entry of their

own, create the proper Chain payment and submit it rather than the users payment. If the attacker gets ahead of the user, then they will win. The defense against a dictionary attack is to avoid common name spaces and to submit your payment to multiple, long standing nodes in the network. In Factom, the flexibility of defining the Chain namespace makes efforts to hog the namespace ineffective.

[0270] Fraudulent servers are now discussed. All Entries in Factom require signatures from the users, or must match a hash that has been signed by the users. This means that fraudulent Federated servers in the Federation pool have very limited attacks they can make on the protocol. Invalid Entries do not validate, and upon broadcasting an invalid Entry, the honest Federated Servers will immediately broadcast a Server Fault Message (SFM) on the fraudulent server. If a majority detect a fault, the faulty server is removed. As long as the majority do not collude, then the protocol will remain honest. Any Federated server that failed detect the fault likewise risks losing its support from Factom users, and dropping from the Federated server pool.

[0271] Federated servers can delay recording of Entry payments. But because Entry payments are submitted via a distributed set of Factom Nodes, delaying of Entry payments will be noted. Users may withdraw support from servers without reasonable performance compared to the rest of the network.

[0272] Federated servers can delay the recording of Entries. Here the payment is accepted (generally by another server) fairly quickly. But for one reason or another, a Federated server refuses to record the Entry. In the next minute, responsibility for that Chain will shift to another server. As long as most servers are honest, the Entry will be recorded. Then the data over time will show that a server is delaying Entries. This will cause them potentially to lose support.

[0273] Federated servers can at any point send false messages. The other Federated servers then would issue a SFR on the on the rogue server when those messages didn't make sense. A majority of the servers issuing an SFR would boot the rogue server, then the network would ignore their messages and not forward them on.

[0274] Federated servers can refuse to accept valid Entry payment messages based on the public address, under the assumption that the public address is associated with some party. Again, assuming a majority of servers are honest, the payment will be accepted when the control shifts to an honest server. Furthermore, nodes watching will see the delay, and perhaps a pattern of delays, and support will be lost for the misbehaving servers.

[0275] FIG. 60 illustrates timestamping into Bitcoin, according to exemplary embodiments. The Factom timestamping mechanism secures transaction in the blockchain. Factom data is timestamped and made irreversible by the Bitcoin network. A user's data is as secure as any other Bitcoin transaction, once published to the Bitcoin blockchain. A compact proof of publication is possible for any data entered into the Factom system.

[0276] As this disclosure above explained, data is organized into block structures, the highest level being Directory Blocks, which are created using Merkle trees. Every 10 minutes, the data set is frozen and submitted to the Bitcoin network. Since Bitcoin has an unpredictable block time, there may be more or fewer than one Factom timestamp per Bitcoin block. Bitcoin internal header block times them-

selves have a fluid idea of time. They have a 2 hour possible drift from reality. Factom will provide its own internal timestamps, adhering with standard time systems.

[0277] The user data ordering will be assigned when received at the Federated servers. Factom organizes the submitted Entry references into sets of blocks. The block time for Factom is ten minutes. On closing, the Federated Server network generates consensus and the Entries that are part of that block structure are timestamped to a minute within the block. As a general note, the data could have existed long before it was timestamped. An Application running on top of Factom could provide finer and more accurate timestamping services prior to Entries being recorded in Factom. The Factom timestamp only proves the data did not originate after the Factom timestamp.

[0278] The Merkle root of the Directory Block is entered into the Bitcoin blockchain with a spending transaction. The spend includes an output with an OP_RETURN. We refer to this as "anchoring" the Directory Block to the Bitcoin blockchain. This method is the least damaging to the Bitcoin network of the various ways to timestamp data.

[0279] Two possible alternatives to the OP_RETURN data in the blockchain is anchored to the P2Pool headers (as in chronobit) or in the Bitcoin block header coinbase. The P2Pool headers would require several hours of mining to find a block which satisfies the P2Pool rules, and the added complexity to the Factom protocol would not be worth the benefits. Including the Merkle root into the coinbase of a block would require cooperation with miners, above and beyond the transaction processing they are already doing. The coinbase entry would still need to have a crypto signature from the Factom system, so would not save on much space relative to a signed transaction.

[0280] The first two bytes of the available 40 in the anchor will be a designator tag (2 bytes with the value "Fa"). The Factom anchor (32 bytes) is concatenated onto the tag, then the block height is added (up to 6 bytes, allowing for >500,000 years). The designator tag indicates the transaction could be a Factom anchor. Other qualifiers are required, but the tag and Factom block height eliminates most of the OP_RETURN transactions that would otherwise need to be inspected. The block height in the OP_RETURN helps to fix the order in those cases where the Bitcoin blockchain records the anchors out of order.

[0281] The anchored data is the Merkle root of list containing the Directory Block's Merkle root. Querying a database or DHT for the anchored data will return the Directory Block which can be used to find the rest of the data in the block. The Merkle root timestamp will be entered into the Bitcoin blockchain by one of the Federated servers. The server delegated to timestamp the federation's collected data creates a Bitcoin transaction. The transaction will be broadcast to the Bitcoin network, and be included in a Bitcoin block. Bitcoin transactions that look like a Factom anchor, but are not spent from an address known as a Factom server would either be junk, or an attempt to fork Factom. Most users/applications would ignore such anchors.

[0282] Bitcoin blocks are generated with a statistical process, and as such their timing cannot be predicted. This means that the anchors are only roughly time-bound by the OP_RETURNs inserted into the Bitcoin blockchain, and its timestamping mechanism. The real value of anchoring Factom to Bitcoin is to prevent anyone from generating false Factom histories. Due to bad luck of Bitcoin miners, or

delayed inclusion of Factom transactions, the time between when the Factom state is frozen for a particular 10 minute period and when the anchor appears in Bitcoin can vary, perhaps significantly.

[0283] Now the ramifications of federated servers and anchoring verses proof of work is discussed. Proof of Work (PoW) is optimized for permissionless participation and validation of the historical record of a blockchain. The typical implementation of Proof of Work is to repeatedly hash blocks until one of the parties mining finds a hash with the difficulty required by the current requirements of a blockchain. This allows anyone to serve as a miner, to collect and validate transactions, pack them into blocks, and repeatedly hash that block looking for a solution that meets the difficulty requirement.

[0284] The shortcomings of PoW have been widely discussed in the media as requiring unnecessary amounts of power, when other sorts of problem solving and work could result in benefits to blockchain users, the ecosystem, and society. Such is the goal of various Proof of Stake (PoS) systems used by various blockchains. But Proof of Stake alone makes the historical record hard to validate, and does not work well for a data recording system like Factom. This is because validating the historical Stake of parties involved the entire blockchain, and an understanding of the Stake that existed at each point in time historically. Factom needs small cryptographic proofs that validate sets of data, which PoW provides. Because PoW is validated solely by evaluating the difficulty of a hash.

[0285] Anchoring is the solution Factom uses to secure the historical record, and at the same time avoid duplicating the massive expenditure of resources required of mining. A system like PoS can be used in the present, while anchoring secures the historical record. The idea of supporting parties allows permissionless participation in the Factom protocol beyond that of the Authority Set.

[0286] The Authority Set and Anchoring means that running the Authority Servers is less expensive in resources by orders of magnitude compared to mining. Greater efficiency means that the rewards paid out by the Factom protocol can do more for the ecosystem than pay very large utility bills. Factom may use various voluntary but auditable methods to incentivize using the efficiency of the authority set to set aside resources within the protocol for productive real world work. A sort of Proof of Development could be used to receive these rewards using distributed support to identify work to be done, and evaluate the quality of the work that results. Such a system could provide rewards for development alongside the rewards generated for the authority set.

[0287] A "Proof of Development" comes with its own issues. The main issue is the "Oracle Problem," where it is very hard to know from within the programming of a blockchain protocol what might be useful development in the real-world and evaluate the quality of such development once it is done. Factom may develop mechanisms to incentivize supporting parties in the protocol to create evaluation processes, audit trails, and certifications at every stage of development to address the Oracle Problem, and allow a self-correcting process to manage a viable "Proof of Development" that is more productive and ecologically friendly than simply rewarding the burning of energy resources for security.

[0288] The Factom protocol and system are now compared with other blockchain technologies. For example,

Factom differs from Bitcoin and Sidechains. Factom is very different from Bitcoin, and in fact very different from any current cryptocurrency project. Cryptocurrencies like Bitcoin implement a strict, distributed method for the validation of transactions, where anyone can validate each transaction, and the validity of every input into a transaction can be verified. Because each transaction is authorized via cryptographic proof, no transaction can be forged. Each transaction can be checked for validity by verifying signatures of each transaction, and the miners hold each other accountable for only including valid transactions.

[0289] The Bitcoin protocol is transactionally complete. In other words, the creation and distribution of Bitcoins through transactions is completely defined within the Bitcoin protocol. Transactions (which specify movement of bitcoin) and block discovery (which move bitcoin via mining fees and provide block rewards) are the only inputs into the Bitcoin Protocol, and nothing leaves the Bitcoin Protocol. In other words, the 21 million bitcoins that will ultimately exist will always and forever exist within the protocol. Pegged sidechains, when implemented, will provide additional movement of bitcoin value outside the blockchain, while the pegged value is in stasis in the blockchain.

[0290] The sidechains proposal describes a solution to increase the scalability of Bitcoin by allowing value control to move off the blockchain and onto a sidechain. In the sidechain, many trades can occur. Later, a cryptographic proof (not all the transactions in between) can be recorded in the blockchain which moves the BTC out of stasis in Bitcoin. This proof would have to be available to the Bitcoin miners, but the bulk of the transaction data would be left behind in the sidechain.

[0291] Factom is in some sense attempting to increase scalability, but not by enabling more value transactions, but by moving non-BTC transactions off blockchain. This would be transactions that are not primarily intended to transfer Bitcoin value. For example transactions could manage domain name registrations, log security camera footage, track the provenance for art work, and even establish the value of show horses by documenting their history. Some of these do not move a value at all, like transactions establishing a proof of publication.

[0292] Sidechains and Factom are both trying to move transactions off the blockchain, but to achieve similar ends via completely different mechanisms. At some point, Factom may integrate with a Bitcoin sidechain in order to take advantage of the atomic swaps from BTC to Factoids.

[0293] Factom is also different from other blockchain technologies. Many different groups are looking to find ways to leverage the Bitcoin approach for managing other sorts of transactions besides tracking bitcoin balances. For example, the trading of assets such as houses or cars can be done digitally using Bitcoin extensions. Even the trading of commodities such as precious metals, futures, or securities might be done via clever encoding and inserting of information into the Bitcoin blockchain. Efforts to expand Bitcoin to cover these kinds of trades include Colored Coins, Mastercoin, and Counterparty. Some developers choose to build their own cryptocurrency with a more flexible protocol that can handle trades beyond currency. These include Namecoin, Ripple, Ethereum, BitShares, NXT, and others. Open Transactions (OT) uses Cryptographic signatures, signed receipts and proof of balance for users (i.e., a user does not need the transaction history to prove their balance,

just the last receipt). In this way, OT can provide the spend of centralized servers without the risk of a centralized server that can alter client balances. Factom is decentralized, and only records Entries. So Factom can record data that would not meet OT's rules. But Factom will not execute at the rate OT can initially. Factom is distributed, and we expect that some, but not all users will employ cryptographic techniques similar to OT with their records.

[0294] The great advantage to an independent platform over trying to build upon Bitcoin is found in flexibility. The Bitcoin protocol isn't optimized to allow for recording of arbitrary pieces of data, so the "bookkeeping" necessary for non-Bitcoin type transactions isn't necessarily supported by Bitcoin. Furthermore, Bitcoin's Proof of Work (PoW) based consensus method is not a "one size fits all" solution, given that some transactions must resolve much faster than 10 minutes. Ripple and Open Transactions vastly speed up confirmation times by changing the consensus method.

[0295] An Application built upon Factom seeks to gain the ability to track assets and implement contracts, by leveraging the blockchain directly. Instead of inserting transactions into the blockchain (viewed as "blockchain bloat" by many), Factom records its Entries within its own structures. At the base level, Factom records what Chains have had Entries added to Factom within the Directory Block time. Scanning these records, Applications can pick out the Chains in which they are interested. Factom records each Chain independently, so Applications can then pull the Chain data they need.

[0296] Factom is organized in a way that minimizes connections between user Chains. A Chain in Factom can be validated without any of the information held in other, unrelated Chains. This minimizes the information a Factom user has to maintain to validate the Chains they are interested in.

[0297] Now Factom Consensus Similarities and Differences from Proof of Stake are discussed. The policy and reward mechanism in Factom is similar to Proof of Stake (PoS). Factom differs from most PoS systems in that many subsets of user stake and/or contribution may be recognized. Individual categories of stake can be weighted against each other to further decentralized Factom. This is an attempt to make the servers answerable to the users actively using and contributing to the protocol. The individual users would delegate their support to a server. The Federated servers with the top numbers of support would be responsible for coming to consensus.

[0298] Some with a deep understand of Bitcoin have recognized that pure PoS consensus mechanisms are fundamentally flawed. There are two attacks that make pure PoS unworkable. The problems are referred to as "Stake Grinding" and "Nothing at Stake". Although Factom has PoS elements, it does not suffer from these problems.

[0299] Stake grinding is a problem where an attacker with a sizable (say 10%), but not majority share can formulate false histories. From some point in history, they can costlessly fork the blockchain, choosing to reorder past transactions such that their stake is always selected to create the subsequent blocks. They would be able to present this alternate version of history as part of an attack to steal value by double spending. Bitcoin solves this problem by strongly linking the information domain, where computers make decisions, with the thermodynamic domain, where humans burn energy. Considerable resources are expended in the

thermodynamic domain, and is provable in the information domain. Bitcoin makes forming false histories hugely expensive.

[0300] Factom is unable to create alternate histories after the fact, since it is unable to insert transactions into historical Bitcoin blocks. It is also unable to create parallel histories without being detected, since Factom is linked to Bitcoin with known Bitcoin private keys.

[0301] The Nothing at Stake problem is more subtle. With a policy disagreement in Bitcoin, miners must choose either one policy or the other. If they choose against the majority, they will be burning lots of electricity without a chance of recouping costs. PoS miners do not face this dilemma. They can hedge their bets and costlessly create forks complying with each side of the policy. They would simultaneously agree with both sides of the disagreement. This would open up the economy to double spend attacks. One of two merchants following different forks will ultimately have that money becomes worthless.

[0302] Bitcoin solves this problem by having unintelligent unambiguous automatable rules for selecting the correct fork. In Bitcoin, the correct fork is the one with the most Proof of Work (PoW). Factom will also have unintelligent unambiguous automatable rules to select a correct fork, should one arise.

[0303] FIG. 61 is a flowchart illustrating a method or algorithm for processing of the digital contract 20, according to exemplary embodiments. The contract identifier 28, the contractual parameter 30, and/or the table identifier 334 is/are received (Block 340). The network resource 50 is identified (Block 342), and the contract processor may be an IP address, URL, virtual machine, or other network destination representing a vendor, contractor, server, or service that executes the decision table 326 and/or the digital contract 20. The service request 266 is sent (Block 344), the service update 270 is received (Block 346), and the service response 268 is received (Block 348). The data records 70 in the blockchain data layer 72 are generated (Block 350), and the data records 70 describe the execution of the digital contract 20. The data records 70 may be hashed (Block 352) and incorporated into the public blockchain 24 (Block 354).

[0304] FIG. 62 is a schematic illustrating still more exemplary embodiments. FIG. 62 is a more detailed diagram illustrating a processor-controlled device 360. As earlier paragraphs explained, the entity's private software application 40, the data layer application 154, and/or the contract application 302 may partially or entirely operate in any mobile or stationary processor-controlled device. FIG. 62, then, illustrates the entity's private software application 40, the data layer application 154, and/or the contract application 302 stored in a memory subsystem of the processor-controlled device 360. One or more processors communicate with the memory subsystem and execute either, some, or all applications. Because the processor-controlled device 360 is well known to those of ordinary skill in the art, no further explanation is needed.

[0305] FIG. 63 depicts other possible operating environments for additional aspects of the exemplary embodiments. FIG. 63 illustrates the entity's private software application 40, the data layer application 154, and/or the contract application 302 operating within various other processor-controlled devices 360. FIG. 63, for example, illustrates that the entity's private software application 40, the data layer application 154, and/or the contract application 302 may

entirely or partially operate within a smartphone **362**, a personal/digital video recorder (PVR/DVR) **364**, a Global Positioning System (GPS) device **366**, an interactive television **368**, a tablet computer **370**, or any computer system, communications device, or processor-controlled device utilizing any of the processors above described and/or a digital signal processor (DP/DSP) **372**. Moreover, the processor-controlled device **360** may also include wearable devices (such as watches), radios, vehicle electronics, clocks, printers, gateways, mobile/implantable medical devices, and other apparatuses and systems. Because the architecture and operating principles of the various devices **360** are well known, the hardware and software componentry of the various devices **360** are not further shown and described.

[0306] Exemplary embodiments may be applied to any signaling standard. Most readers are thought familiar with the Global System for Mobile (GSM) communications signaling standard. Those of ordinary skill in the art, however, also recognize that exemplary embodiments are equally applicable to any communications device utilizing the Time Division Multiple Access signaling standard, the Code Division Multiple Access signaling standard, the “dual-mode” GSM-ANSI Interoperability Team (GAIT) signaling standard, or any variant of the GSM/CDMA/TDMA signaling standard. Exemplary embodiments may also be applied to other standards, such as the I.E.E.E. 802 family of standards, the Industrial, Scientific, and Medical band of the electromagnetic spectrum, BLUETOOTH®, and any other.

[0307] Exemplary embodiments may be physically embodied on or in a computer-readable storage medium. This computer-readable medium, for example, may include CD-ROM, DVD, tape, cassette, floppy disk, optical disk, memory card, memory drive, and large-capacity disks. This computer-readable medium, or media, could be distributed to end-subscribers, licensees, and assignees. A computer program product comprises processor-executable instructions for execution of digital contracts, as the above paragraphs explain.

[0308] While the exemplary embodiments have been described with respect to various features, aspects, and embodiments, those skilled and unskilled in the art will recognize the exemplary embodiments are not so limited. Other variations, modifications, and alternative embodiments may be made without departing from the spirit and scope of the exemplary embodiments.

1. A method performed by a server that records cryptographic transactions conducted by computers, the method comprising:

- receiving, by the server, a cryptographic transaction of the cryptographic transactions;
 - determining, by the server, a chain identifier that is associated with the cryptographic transaction;
 - determining, by the server, an entry block in a blockchain data layer that is associated with the chain identifier;
 - adding, by the server, the cryptographic transaction to the entry block in the blockchain data layer that is associated with the chain identifier;
 - generating, by the server, a directory block in the blockchain data layer based on the entry block and the chain identifier; and
 - recording, by the server, the directory block to a blockchain;
- wherein the directory block records the cryptographic transaction.

2. The method of claim **1**, further comprising generating a hash value representing the directory block.

3. The method of claim **2**, further comprising recording the hash value representing the directory block to the blockchain.

4. The method of claim **1**, further comprising recording the chain identifier to the blockchain.

5. The method of claim **1**, further comprising timestamping the cryptographic transaction.

6. The method of claim **1**, further comprising timestamping the entry block. The method of claim **1**, further comprising timestamping the directory block.

8. A system, comprising:

- a hardware processor; and
 - a memory device storing instructions that when executed by the hardware processor perform operations, the operations comprising:
 - receiving a cryptographic transaction conducted by computers via a network;
 - determining a chain identifier that is associated with the cryptographic transaction;
 - determining an entry block in a blockchain data layer that is associated with the chain identifier;
 - adding the cryptographic transaction to the entry block in the blockchain data layer that is associated with the chain identifier;
 - generating a directory block in the blockchain data layer based on the entry block and the chain identifier; and
 - recording the directory block to a blockchain;
- wherein the directory block records the cryptographic transaction.

9. The system of claim **8**, wherein the operations further comprise generating a hash value representing the directory block.

10. The system of claim **9**, wherein the operations further comprise recording the hash value representing the directory block to the blockchain.

11. The system of claim **8**, wherein the operations further comprise recording the chain identifier to the blockchain.

12. The system of claim **8**, wherein the operations further comprise timestamping the cryptographic transaction.

13. The system of claim **8**, wherein the operations further comprise timestamping the entry block.

14. The system of claim **8**, wherein the operations further comprise timestamping the directory block.

15. A memory device storing instructions that when executed by a hardware processor perform operations, the operations comprising:

- receiving cryptographic transactions conducted by computers via a network;
- determining a chain identifier that is associated with the cryptographic transactions;
- determining a directory block that is associated with the chain identifier;
- recording the cryptographic transactions to the directory block;
- determining a rate of generation associated with the directory block;
- identifying a virtual machine that is associated with the rate of generation; and
- assigning the virtual machine to the cryptographic transactions.

16. The memory device of claim **15**, wherein the operations further comprise hashing the directory block.

17. The memory device of claim **15**, wherein the operations further comprise executing the virtual machine.

18. The memory device of claim **15**, wherein the operations further comprise recording the virtual machine to a blockchain.

19. The memory device of claim **15**, wherein the operations further comprise recording the cryptographic transactions to the blockchain.

20. The memory device of claim **15**, wherein the operations further comprise recording the rate of generation to the blockchain.

* * * * *