

METHODS AND DEVICES FOR DEADLOCK DETECTION IN REAL-TIME GROSS SETTLEMENT SYSTEMS

CROSS-REFERENCE TO RELATED APPLICATIONS

[0001] This application is related to co-pending Singapore Application No. 10201907041R (Applicant Ref. No. A22886SG), entitled “Methods and Devices for Deadlock Detection in Real-Time Gross Settlement Systems” and being filed concurrently herewith, which is hereby incorporated by reference in its entirety.

[0002] This application is related to co-pending Singapore Application No. 10201907055Y (Applicant Ref. No. A23582SG), entitled “Methods and Devices for Deadlock Detection in Real-Time Gross Settlement Systems” and being filed concurrently herewith, which is hereby incorporated by reference in its entirety.

[0003] This application is related to co-pending Singapore Application No. 10201907039X (Applicant Ref. No. A23583SG), entitled “Methods and Devices for Deadlock Detection in Real-Time Gross Settlement Systems” and being filed concurrently herewith, which is hereby incorporated by reference in its entirety.

TECHNICAL FIELD

[0004] The specification relates generally to computer technologies, and more particularly, to methods and devices for deadlock detection in real-time gross settlement systems.

BACKGROUND

[0005] Real-time gross settlement (RTGS) systems are fund transfer systems where a transfer of funds (e.g., money or securities) takes place from one party (e.g., a bank) to another party (e.g., another bank) on a real-time and gross basis. When a settlement in the transfer of funds occurs, the transfer is generally final and irrevocable. The settlement is in real-time when, for example, a payment transaction is not subject to any waiting period. The settlement is on a gross basis when, for example, a payment transaction is settled on one-to-one basis without bundling or netting with any other transaction. RTGS systems may be operated by a country's central bank.

[0006] There are situations where one or more parties, e.g., banks, participating in an RTGS system may have insufficient liquidity to settle their payment transactions. For example, suppose that Bank A is set to receive \$15 from Bank B and transfer \$20 to Bank C, and further suppose that Bank A has \$0 liquidity before settlement can take place. Bank A in this situation is not able to settle the transactions unless Bank A receives additional supply of liquidity. Bank A in this situation is deemed to have caused a deadlock in the RTGS system.

[0007] Currently, a central bank operating the RTGS system may detect possible deadlocks. The central bank can detect deadlocks because it has the authority to obtain liquidity and transaction information from the parties, e.g., banks, participating in the RTGS system. However, with decentralized computing becoming a trend, technologies such as blockchains and the like are becoming more prevalent. Existing RTGS systems implemented based on decentralized computing technologies, such as blockchains and the like, lack the abilities to detect deadlocks.

[0008] Blockchain systems, also known as distributed ledger systems (DLSs) or consensus systems, may enable participating parties to store data securely and immutably. Blockchain systems may include any DLSs, without referencing any particular use case, and may be used for public, private, and consortium blockchain networks. A public blockchain network is open for all entities to use the system and participate in the consensus process. A private blockchain network is provided for a particular entity, which centrally controls read and write permissions. A consortium blockchain network is provided for a select group of entities, which control the consensus process, and includes an access control layer.

[0009] A blockchain system is implemented using a peer-to-peer (P2P) network, in which the nodes communicate directly with each other, e.g., without the need of a fixed, central server. Each node in the P2P network may initiate communication with another node in the P2P network. A blockchain system maintains one or more blockchains. A blockchain is a data structure for storing data, such as transactions, that may prevent tampering and manipulation of the data by malicious parties.

[0010] Unlike RTGS systems operated by a central bank, where participating parties are willing to, or required to, disclose their liquidity and transaction information to the central bank, participating parties in a blockchain system may not be willing to disclose such information to other participating parties in the blockchain system. This lack of disclosure makes it difficult for RTGS systems implemented based on decentralized computing technologies, such as blockchains and the like, to detect deadlocks. Therefore, there is a need for privacy-preserving deadlock detection methods that are capable of detecting deadlocks without requiring any participating

party to disclose their liquidity and transaction information. There is also a need for such methods to verify their computation correctness.

SUMMARY

[0011] In one aspect, a computer-implemented method for detecting deadlocks in a real-time gross settlement system includes: instructing a plurality of users of the real-time gross settlement system to independently calculate their respective liquidity information, and to perform joint calculations for determining an aggregated total amount receivable in the real-time gross settlement system and an aggregated total amount payable in the real-time gross settlement system; receiving calculation results of the joint calculations from the plurality of users and calculating the aggregated total amount receivable and the aggregated total amount payable based on the received calculation results; verifying correctness of the liquidity information of the plurality of users based on the aggregated total amount receivable and the aggregated total amount payable; and after the correctness of the liquidity information of the plurality of users is verified, determining whether a deadlock exists in the real-time gross settlement system.

[0012] In another aspect, a device for detecting deadlocks in a real-time gross settlement system includes: one or more processors; and one or more computer-readable memories coupled to the one or more processors and having instructions stored thereon that are executable by the one or more processors to instruct a plurality of users of the real-time gross settlement system to independently calculate their respective liquidity information, and to perform joint calculations for determining an

aggregated total amount receivable in the real-time gross settlement system and an aggregated total amount payable in the real-time gross settlement system; receive calculation results of the joint calculations from the plurality of users and calculating the aggregated total amount receivable and the aggregated total amount payable based on the received calculation results; verify correctness of the liquidity information of the plurality of users based on the aggregated total amount receivable and the aggregated total amount payable; and after the correctness of the liquidity information of the plurality of users is verified, determine whether a deadlock exists in the real-time gross settlement system.

[0013] In still another aspect, a non-transitory computer-readable medium have stored therein instructions that, when executed by a processor of a device, cause the device to perform a method for detecting deadlocks in a real-time gross settlement system. The method includes: instructing a plurality of users of the real-time gross settlement system to independently calculate their respective liquidity information, and to perform joint calculations for determining an aggregated total amount receivable in the real-time gross settlement system and an aggregated total amount payable in the real-time gross settlement system; receiving calculation results of the joint calculations from the plurality of users and calculating the aggregated total amount receivable and the aggregated total amount payable based on the received calculation results; verifying correctness of the liquidity information of the plurality of users based on the aggregated total amount receivable and the aggregated total amount payable; and after the correctness of the liquidity information of the plurality of users is verified, determining whether a deadlock exists in the real-time gross settlement system.

BRIEF DESCRIPTION OF THE DRAWINGS

[0014] The accompanying drawings, which are incorporated in and constitute a part of this specification, illustrate embodiments. In the following description, which refers to the drawings, the same numbers in different drawings represent the same or similar elements unless otherwise represented.

[0015] FIG. 1 is a schematic diagram of a blockchain system, according to an embodiment.

[0016] FIG. 2 is a schematic diagram of a computing device for implementing a node in a blockchain system, according to an embodiment.

[0017] FIG. 3 is a flow chart of a privacy-preserving deadlock detection method, according an embodiment.

[0018] FIG. 4 is a flow chart of a method for detecting real-time gross settlement deadlocks, according an embodiment.

[0019] FIG. 5 is a block diagram of an apparatus for detecting real-time gross settlement deadlocks, according to an embodiment.

DETAILED DESCRIPTION

[0020] Embodiments of the specification provide methods and devices for providing privacy-preserving deadlock detection in real-time gross settlement (RTGS) systems. The methods and devices support local calculation of liquidity information by users of an RTGS system. The methods and devices also support a multiparty computation in which multiple users of the RTGS system can jointly determine the

minimum liquidity amount without requiring any user to disclose its own liquidity information. The methods and devices further support a multiparty computation in which multiple users of the RTGS system can perform a joint calculation for determining an aggregated total amount receivable and a joint calculation for determining an aggregated total amount payable in the RTGS system, and provide calculation results to a blockchain system for determining the aggregated total amount receivable and the aggregated total amount payable. In this manner, the blockchain system can utilize the aggregated total amount receivable and the aggregated total amount payable to verify correctness of liquidity information of the users and, thus, correctness of the determined minimum liquidity amount and utilize the determined minimum liquidity amount to detect deadlocks in the RTGS system.

[0021] Embodiments disclosed in the specification have one or more technical effects. In some embodiments, the methods and devices support local calculation of liquidity information by users of the RTGS system. This allows each user to perform calculations locally without having to disclose any private information that the user does not wish to share. In some embodiments, the methods and devices also support a multiparty computation of the minimum liquidity amount that can be carried out without requiring any user to disclose its own liquidity information to any other user. This allows the users to jointly carry out the computation while preserving their privacy. Furthermore, in some embodiments, the methods and devices support multiparty summations that can be carried out to calculate an aggregated total amount receivable and an aggregated total amount payable in the RTGS system and provide the summation results to a blockchain system for verification. This provides the blockchain

system with the ability to verify computation correctness without needing to know the users' actual liquidity information. If the aggregated total amount receivable in the RTGS system does not match the aggregated total amount payable in the RTGS system, the methods and devices can detect the mismatch and invalidate the determined minimum liquidity amount, thereby improving the accuracy of deadlock detection.

[0022] A blockchain is a data structure that stores data, e.g., transactions, in a way that may prevent tampering and manipulation of the data by malicious parties. The transactions stored in this manner may be immutable and subsequently verified. A blockchain includes one or more blocks. Each block is linked to a previous block immediately before it in the blockchain by including a cryptographic hash of the previous block. Each block also may include a timestamp, its own cryptographic hash, and one or more transactions. The transactions, which generally have already been verified by the nodes of the blockchain system, may be hashed and encoded into a data structure, such as a Merkle tree. In a Merkle tree, data at leaf nodes of the tree is hashed, and all hashes in each branch of the tree may be concatenated at a root of the branch. This process continues up the tree to the root of the entire tree, which stores a hash that is representative of all data in the tree. A hash purporting to be of a transaction stored in the tree can be quickly verified by determining whether it is consistent with the structure of the tree.

[0023] A blockchain system includes a network of computing nodes that manage, update, and maintain one or more blockchains. The network may be a public blockchain network, a private blockchain network, or a consortium blockchain network.

For example, numerous entities, such as hundreds, thousands, or even millions of entities, can operate in a public blockchain network, and each of the entities operates at least one node in the public blockchain network. Accordingly, the public blockchain network can be considered a public network with respect to the participating entities. Sometimes, a majority of entities (nodes) must sign every block for the block to be valid and added to the blockchain of the blockchain network. Examples of public blockchain networks include particular peer-to-peer payment networks that leverage a distributed ledger, referred to as blockchain.

[0024] In general, a public blockchain network may support public transactions. A public transaction is shared with all of the nodes in the public blockchain network, and is stored in a global blockchain. A global blockchain is a blockchain replicated across all nodes, and all nodes are in perfect state consensus with respect to the global blockchain. To achieve consensus (e.g., agreement to the addition of a block to a blockchain), a consensus protocol is implemented in the public blockchain network. Examples of consensus protocols include proof-of-work (POW) (e.g., implemented in the some crypto-currency networks), proof-of-stake (POS), and proof-of-authority (POA).

[0025] In general, a private blockchain network may be provided for a particular entity, which centrally controls read and write permissions. The entity controls which nodes are able to participate in the blockchain network. Consequently, private blockchain networks are generally referred to as permissioned networks that place restrictions on who is allowed to participate in the network, and on their level of participation (e.g., only in certain transactions). Various types of access control

mechanisms can be used (e.g., existing participants vote on adding new entities, a regulatory authority can control admission).

[0026] In general, a consortium blockchain network may be private among the participating entities. In a consortium blockchain network, the consensus process is controlled by an authorized set of nodes, one or more nodes being operated by a respective entity (e.g., a financial institution, insurance company). For example, a consortium of ten (10) entities (e.g., financial institutions, insurance companies) can operate a consortium blockchain network, each of which operates at least one node in the consortium blockchain network. Accordingly, the consortium blockchain network can be considered a private network with respect to the participating entities. In some examples, each entity (node) must sign every block in order for the block to be valid, and added to the blockchain. In some examples, at least a sub-set of entities (nodes) (e.g., at least 7 entities) must sign every block in order for the block to be valid, and added to the blockchain.

[0027] FIG. 1 illustrates a schematic diagram of a blockchain system 100, according to an embodiment. Referring to FIG. 1, the blockchain system 100 may include a plurality of nodes, e.g., nodes 102-110, configured to operate on a blockchain 120. The nodes 102-110 may form a network 112, such as a peer-to-peer (P2P) network. Each of the nodes 102-110 may be a computing device, such as a computer or a computer system, configured to store a copy of the blockchain 120, or may be software running on the computing device, such as a process or an application. Each of the nodes 102-110 may have a unique identifier.

[0028] The blockchain 120 may include a growing list of records in the form of

data blocks, such as blocks B1-B5 in FIG. 1. Each of the blocks B1-B5 may include a timestamp, a cryptographic hash of a previous block, and data of the present block, which may be transactions such as monetary transactions. For example, as illustrated in FIG. 1, block B5 may include a timestamp, a cryptographic hash of block B4, and transaction data of block B5. Also, for example, a hashing operation may be performed on the previous block to generate the cryptographic hash of the previous block. The hashing operation may convert inputs of various lengths into cryptographic outputs of a fixed length through a hash algorithm, such as SHA-256.

[0029] The nodes 102-110 may be configured to perform an operation on the blockchain 120. For example, when a node, e.g., the node 102, wants to store new data onto the blockchain 120, that node may generate a new block to be added to the blockchain 120 and broadcast the new block to other nodes, e.g., the nodes 104-110, in the network 112. Based on legitimacy of the new block, e.g., validity of its signature and transactions, the other nodes may determine to accept the new block, such that the node 102 and the other nodes may add the new block to their respective copies of the blockchain 120. As this process repeats, more and more blocks of data may be added to the blockchain 120.

[0030] FIG. 2 illustrates a schematic diagram of a computing device 200 for implementing a node, e.g., the node 102 (FIG. 1), in a blockchain system, according to an embodiment. Referring to FIG. 2, the computing device 200 may include a communication interface 202, a processor 204, and a memory 206.

[0031] The communication interface 202 may facilitate communications between the computing device 200 and devices implementing other nodes, e.g., nodes 104-110

(FIG. 1), in the network. In some embodiments, the communication interface 202 is configured to support one or more communication standards, such as an Internet standard or protocol, an Integrated Services Digital Network (ISDN) standard, etc. In some embodiments, the communication interface 202 may include one or more of a Local Area Network (LAN) card, a cable modem, a satellite modem, a data bus, a cable, a wireless communication channel, a radio-based communication channel, a cellular communication channel, an Internet Protocol (IP) based communication device, or other communication devices for wired and/or wireless communications. In some embodiments, the communication interface 202 may be based on public cloud infrastructure, private cloud infrastructure, hybrid public/private cloud infrastructure.

[0032] The processor 204 may include one or more dedicated processing units, application-specific integrated circuits (ASICs), field-programmable gate arrays (FPGAs), or various other types of processors or processing units. The processor 204 is coupled with the memory 206 and is configured to execute instructions stored in the memory 206.

[0033] The memory 206 may store processor-executable instructions and data, such as a copy of the blockchain 120 (FIG. 1). The memory 206 may include any type of volatile or non-volatile memory devices, or a combination thereof, such as a static random-access memory (SRAM), an electrically erasable programmable read-only memory (EEPROM), an erasable programmable read-only memory (EPROM), a programmable read-only memory (PROM), a read-only memory (ROM), a magnetic memory, a flash memory, or a magnetic or optical disk. When the instructions in the memory 206 are executed by the processor 204, the computing device 200 may

perform an operation on the blockchain 120.

[0034] Referring back to FIG. 1, in some embodiments, the blockchain system 100 may be implemented to support various types of participating parties, or users, including, e.g., banks, financial institutions, credit unions, and the like. The various parties may utilize the blockchain system 100 to retrieve or record information, including, e.g., payment transfer information. Typically, records such as payment transfers and the like recorded on the blockchain system 100 are encrypted by relevant parties.

[0035] In some embodiments, the blockchain system 100 may be implemented to support real-time gross settlement (RTGS). For example, the blockchain system 100 may enforce a policy so that all transfers, once processed, are final and irrevocable. The blockchain system 100 may also enforce a policy so that none of the payment transactions submitted to the blockchain system 100 is subject to a waiting period. The blockchain system 100 may further enforce a policy so that payment transactions are settled on one-to-one basis without bundling or netting with any other transactions.

[0036] FIG. 3 illustrates a flow chart of a privacy-preserving deadlock detection method 300 for detecting deadlocks in an RTGS system implemented on a blockchain system, e.g., the blockchain system 100 (FIG. 1), according to an embodiment. The privacy-preserving deadlock detection method 300 may detect deadlocks without requiring any participating parties to disclose their liquidity and transaction information.

[0037] At step 302, each party participating in the RTGS system, e.g., Parties A through N, may be instructed to independently calculate its liquidity information. In some embodiments, each party may independently calculate its liquidity information

using its local computing resources. In some embodiments, each party $i \in \{A, B, \dots, N\}$ may calculate its current liquidity C_i , an amount T_i to be transferred to (receivable by) the party, an amount F_i to be transferred from (payable by) the party, and a post transfer liquidity P_i , where $P_i = C_i + T_i - F_i$.

[0038] At step 304, the parties, Parties A through N, may be instructed to perform a joint calculation for determining, e.g., further calculating, an aggregated total amount receivable in the RTGS system. Parties A through N may also be instructed to perform a joint calculation for determining, e.g., further calculating, an aggregated total amount payable in the RTGS system. The aggregated total amount receivable may be calculated as $S_T = \sum_{i \in \{A, B, \dots, N\}} T_i$. The aggregated total amount payable may be calculated as $S_F = \sum_{i \in \{A, B, \dots, N\}} F_i$.

[0039] In some embodiments, the parties, Parties A through N, may be instructed to calculate the aggregated total amount receivable S_T using a multiparty summation process. Parties A through N may also be instructed to calculate the aggregated total amount payable S_F using the multiparty summation process. In some embodiments, the multiparty summation process utilized may allow the parties to calculate S_T and S_F without requiring the parties to disclose their liquidity information to each other. In some embodiments, Parties A through N may jointly carry out the multiparty summation process disclosed in Yiping Shen et al., "On the Confidential Auditing of Distributed Computing Systems," Proceedings of the 24th International Conference on Distributed Computing Systems (2004) 600-607, or that disclosed in L. Jean Camp et al., "Providing Auditing While Protecting Privacy," The Information Society, 10(1), pp. 59-72, 1998, both of which are herein incorporated by reference in their entirety. Parties A through N

may also implement other types of multiparty summation processes to carry out step 304.

[0040] In some embodiments, Parties A through N may operate with the blockchain system 100 to carry out the multiparty summation process. For example, the blockchain system 100 may instruct each party $i \in \{A, B, \dots, N\}$ to calculate a share of the aggregated total amount receivable S_T at step 304 and submit the calculated result to the blockchain system 100 at step 306. The blockchain system 100 may also instruct each party i to calculate a share of the aggregated total amount payable S_F at step 304 and submit the calculated result to the blockchain system 100 at step 306. The blockchain system 100 may consolidate the results submitted by the parties to determine the values of S_T and S_F at step 308, without receiving the actual liquidity information of the respective parties.

[0041] In an embodiment, parties A through N may operate with the blockchain system 100 to carry out this multiparty summation process, and at step 304, each party i may be instructed to mask its liquidity information using a polynomial. For illustrative purposes, suppose that there are three parties, Parties A, B, and C, their amounts receivable are T_A , T_B , and T_C , and their amounts payable are F_A , F_B , and F_C , respectively. To calculate the aggregated total amount receivable S_T , each party i may be instructed to select two random integers, a and b , to be used in a second-degree polynomial $ax^2 + bx + c$, where the constant in the polynomial, c , is set to represent the party's amount receivable T_i . Party A may then mask its amount receivable T_A by calculating masked values using the polynomial at $x = 1$, $x = 2$, and $x = 3$. Parties B and C may carry out similar operations with their own selection of random integers a

and b to mask their own amounts receivable T_B , and T_C , respectively.

[0042] A similar process may be carried out to calculate the aggregated total amount payable S_F . Each party i may be instructed to select two random integers, a and b , to be used in a second-degree polynomial $ax^2 + bx + c$, where the constant in the polynomial, c , is set to represent the party's amount payable F_i . Party A may then mask its amount payable F_A by calculating masked values using the polynomial at $x = 1$, $x = 2$, and $x = 3$. Parties B and C may carry out similar operations with their own selection of random integers a and b to mask their own amounts payable F_B , and F_C , respectively.

[0043] In some embodiments, the calculation of the aggregated total amount receivable S_T and the calculation of the aggregated total amount payable S_F may be carried out as separate multiparty computation processes. A party may therefore select one set of random integers, a and b , for the polynomial used in the calculation of the aggregated total amount receivable S_T and another set of random integers, a' and b' , for the polynomial used in the calculation of the aggregated total amount payable S_F .

[0044] Continuing with the calculation of the aggregated total amount receivable S_T as an example, after the parties have masked their amounts receivable T_i , the parties may be instructed to keep one of their masked values and share the rest of their masked values with other parties to facilitate the calculation of the S_T . Party A in the example above may keep a masked value calculated using the polynomial $ax^2 + bx + T_A$ at $x = 1$, share a masked value calculated using the polynomial $ax^2 + bx + T_A$ at $x = 2$ with Party B, and share another masked value calculated using the polynomial $ax^2 + bx + T_A$ at $x = 3$ with Party C. Parties B and C may carry out similar operations

with respect to their masked values.

[0045] Each party i may then proceed to calculate its share of the aggregated total amount receivable S_T based on the masked value it kept to itself and the masked values it received from other parties. Continuing with the example above, Party A may calculate a sum based on the masked value it kept (e.g., calculated by Party A using Party A's own polynomial at $x = 1$), the masked value it received from Party B (e.g., calculated by Party B using Party B's polynomial at $x = 1$), and the masked value it received from Party C (e.g., calculated by Party C using Party C's polynomial at $x = 1$). Party B may calculate a sum based on the masked value it kept (e.g., calculated by Party B using Party B's own polynomial at $x = 2$), the masked value it received from Party A (e.g., calculated by Party A using Party A's polynomial at $x = 2$), and the masked value it received from Party C (e.g., calculated by Party C using Party C's polynomial at $x = 2$). Party C may calculate a sum based on the masked value it kept (e.g., calculated by Party C using Party C's own polynomial at $x = 3$), the masked value it received from Party A (e.g., calculated by Party A using Party A's polynomial at $x = 3$), and the masked value it received from Party B (e.g., calculated by Party B using Party B's polynomial at $x = 3$).

[0046] Similarly, continuing with the calculation of the aggregated total amount payable S_F as an example, after the parties have masked their amounts payable F_i , the parties may be instructed to keep one of their masked values and share the rest of their masked values with other parties to facilitate the calculation of the S_F . Party A in the example above may keep a masked value calculated using the polynomial $a'x^2 + b'x + F_A$ at $x = 1$, share a masked value calculated using the polynomial $a'x^2 + b'x + F_A$ at

$x = 2$ with Party B, and share another masked value calculated using the polynomial $a'x^2 + b'x + F_A$ at $x = 3$ with Party C. Parties B and C may carry out similar operations with respect to their masked values.

[0047] Each party i may then proceed to calculate its share of the aggregated total amount payable S_F based on the masked value it kept to itself and the masked values it received from other parties. Continuing with the example above, Party A may calculate a sum based on the masked value it kept (e.g., calculated by Party A using Party A's own polynomial at $x = 1$), the masked value it received from Party B (e.g., calculated by Party B using Party B's polynomial at $x = 1$), and the masked value it received from Party C (e.g., calculated by Party C using Party C's polynomial at $x = 1$). Party B may calculate a sum based on the masked value it kept (e.g., calculated by Party B using Party B's own polynomial at $x = 2$), the masked value it received from Party A (e.g., calculated by Party A using Party A's polynomial at $x = 2$), and the masked value it received from Party C (e.g., calculated by Party C using Party C's polynomial at $x = 2$). Party C may calculate a sum based on the masked value it kept (e.g., calculated by Party C using Party C's own polynomial at $x = 3$), the masked value it received from Party A (e.g., calculated by Party A using Party A's polynomial at $x = 3$), and the masked value it received from Party B (e.g., calculated by Party B using Party B's polynomial at $x = 3$).

[0048] At step 306, the parties, Parties A through N, may submit the results of their calculations to the blockchain system 100 for recordation. The blockchain system 100 may implement one or more smart contracts and utilize the smart contracts to consolidate the results submitted by the parties to calculate the aggregated total amount

receivable S_T and the aggregated total amount payable S_F at step 308.

[0049] Smart contracts are computer protocols implemented in the form of computer code that are incorporated into the blockchain 120 (FIG. 1), to facilitate, verify, or enforce the negotiation or performance of contracts. For example, a user of the blockchain system 100 may program agreed terms into a smart contract using a programming language, such as C++, Java, Solidity, Python, etc., and when the terms are met, the smart contract may be automatically executed by the blockchain system 100, e.g., to perform a task. Also for example, the smart contract may include a plurality of subroutines or functions, each of which may be a sequence of program instructions that performs a specific task. The smart contract may be operational code that is fully or partially executed without human interaction.

[0050] In some embodiments, the blockchain system 100 may implement a smart contract that is programmed to calculate the value of S_T based on results submitted by Parties A through N. The smart contract may, for example, consolidate the results submitted by Parties A through N and solve the equation to determine a sum representing the combined value of all constants used in the polynomials. It is noted that because each party used the constant c in its polynomial to represent the amount receivable T_i of that party, the sum representing the combined value of all constants used in the polynomials will therefore represent the sum of amounts receivable of all parties in the RTGS system, which can be expressed as $S_T = \sum_{i \in \{A, B, \dots, N\}} T_i$.

[0051] Similarly, the blockchain system 100 may implement a smart contract that is programmed to calculate the value of S_F based on results submitted by Parties A through N. The smart contract may, for example, consolidate the results submitted by

Parties A through N and solve the equation to determine a sum representing the combined value of all constants used in the polynomials. It is noted that because each party used the constant c' in its polynomial to represent the amount payable F_i of that party, the sum representing the combined value of all constants used in the polynomials will therefore represent the sum of amounts payable of all parties in the RTGS system, which can be expressed as $S_F = \sum_{i \in \{A, B, \dots, N\}} F_i$.

[0052] It is to be understood that while the examples described above discussed using second-degree polynomials to facilitate masking of values provided by three participating parties, such examples are depicted merely for illustrative purposes and are not meant to be limiting. It is to be understood that additional parties may participate in the multiparty summation process discussed above. It is also to be understood that the parties may implement other types of multiparty computation processes to facilitate the calculation of S_T and S_F .

[0053] In some embodiments, the blockchain system 100 may utilize the calculated values of S_T and S_F to verify computation correctness of the parties' liquidity information at step 308. In some embodiments, the blockchain system 100 may require S_T to match, e.g., be equal to, S_F . If the blockchain system 100, e.g., by executing a smart contract, determines that $S_T \neq S_F$, the blockchain system 100 may recognize this mismatch and report it as an error.

[0054] If the blockchain system 100, e.g., by executing the smart contract, determines that $S_T = S_F$, the blockchain system 100 may consider the calculations carried out so far to be correct. The blockchain system 100 may then instruct the parties, Parties A through N, to jointly carry out a multiparty computation to determine a

minimum liquidity amount $P_{min} = \min_{i \in \{A, B, \dots, N\}} \{P_i\}$ at step 310. In some embodiments, Parties A through N may jointly carry out the multiparty computation without disclosing their liquidity information, including their post transfer liquidity P_i . In this manner, Parties A through N can jointly determine P_{min} and preserve their privacy. In some embodiments, Parties A through N may jointly carry out a multiparty computation protocol disclosed in Gagan Aggarwal et al., “Secure Computation of the Median (and Other Elements of Specified Ranks),” Journal of Cryptology (2010) 23: 373–401, which is herein incorporated by reference in its entirety. For instance, Parties A through N may jointly carry out a multiparty computation protocol defined as follows:

- 1: Provide each party with initial values of α and β set by the RTGS system, the initial value of α being sufficiently small to represent a possible minimum post transfer liquidity and the initial value of β being sufficiently large to represent a possible maximum post transfer liquidity.
- 2: *Repeat* until “done”:
- 3: Set $m = \lceil (\alpha + \beta) / 2 \rceil$.
- 4: Instruct each party $i \in \{A, B, \dots, N\}$ to compare P_i to m ; if P_i is strictly less than m , set $l_i = 1$ and $g_i = 0$; otherwise, if P_i is strictly greater than m , set $l_i = 0$ and $g_i = 1$.
- 5: Calculate $\sum_{i \in \{A, B, \dots, N\}} l_i$ and $\sum_{i \in \{A, B, \dots, N\}} g_i$.
- 6: If $\sum_{i \in \{A, B, \dots, N\}} l_i \leq 0$ and $\sum_{i \in \{A, B, \dots, N\}} g_i \leq N - 1$, then report m as $\min (P_A, P_B, \dots, P_N)$ and proceed to “done.”
- 7: If $\sum_{i \in \{A, B, \dots, N\}} l_i \geq 1$, then set $\beta = m - 1$ (meaning that $\min (P_A, P_B, \dots, P_N)$ is smaller than m , so m should be decreased).

8: If $\sum_{i \in \{A, B, \dots, N\}} g_i \geq N$, then set $\alpha = m + 1$ (meaning that $\min(P_A, P_B, \dots, P_N)$ is greater than m , so m should be increased).

9: Done.

[0055] It is noted that the multiparty computation protocol described above may allow Parties A through N to compute the minimum liquidity amount, P_{min} , without revealing $P_A, P_B, \dots, P_N$. It is to be understood, however, that the multiparty computation protocol described above is provided merely as an example and is not meant to be limiting. It is contemplated that Parties A through N may implement other types of multiparty computation techniques to carry out step 310.

[0056] It is also contemplated that in some embodiments, Parties A through N may utilize a multiparty summation process to calculate $\sum_{i \in \{A, B, \dots, N\}} l_i$ and $\sum_{i \in \{A, B, \dots, N\}} g_i$. In this manner, Parties A through N may calculate $\sum_{i \in \{A, B, \dots, N\}} l_i$ and $\sum_{i \in \{A, B, \dots, N\}} g_i$ without requiring the parties to disclose their l and g values to each other, further preserving the privacy of Parties A through N. For instance, in some embodiments, Parties A through N may jointly carry out a multiparty summation process similar to that described above, e.g., by masking their l and g values. For illustrative purposes, suppose that there are three parties, Parties A, B, and C, and their l and g values are l_A, l_B , and l_C , and g_A, g_B , and g_C , respectively. Each party i may be instructed to select two random integers, a and b , to be used in a second-degree polynomial $ax^2 + bx + c$, where the constant in the polynomial, c , is set to represent the party's l value or the party's g value, depending on whether the parties are instructed to calculate $\sum_{i \in \{A, B, \dots, N\}} l_i$ or $\sum_{i \in \{A, B, \dots, N\}} g_i$. Suppose that the parties are instructed to calculate $\sum_{i \in \{A, B, \dots, N\}} l_i$, Party A may select, for example, $a = 25$ and $b = 83$, which means Party A will have its

polynomial defined as $25x^2 + 83x + l_A$. Party A may mask its l value by calculating masked values using the polynomial at $x = 1$, $x = 2$, and $x = 3$. Parties B and C may carry out similar operations with their own selection of random integers a and b to mask their own l values.

[0057] Each party i may be further instructed to keep one of its masked values and share the rest of its masked values with other parties to facilitate the calculation of $\sum_{i \in \{A, B, \dots, N\}} l_i$. Continuing with the three-party example above, Party A may keep a masked value calculated using the polynomial $25x^2 + 83x + l_A$ at $x = 1$, share a masked value calculated using the polynomial $25x^2 + 83x + l_A$ at $x = 2$ with Party B, and share another masked value calculated using the polynomial $25x^2 + 83x + l_A$ at $x = 3$ with Party C. Parties B and C may carry out similar operations with respect to their masked values.

[0058] Each party i may then proceed to calculate its share of $\sum_{i \in \{A, B, \dots, N\}} l_i$ based on the masked value it kept to itself and the masked values it received from other parties. Continuing with the three-party example above, Party A may calculate a sum based on the masked value it kept (e.g., calculated by Party A using Party A's own polynomial at $x = 1$), the masked value it received from Party B (e.g., calculated by Party B using Party B's polynomial at $x = 1$), and the masked value it received from Party C (e.g., calculated by Party C using Party C's polynomial at $x = 1$). Party B may calculate a sum based on the masked value it kept (e.g., calculated by Party B using Party B's own polynomial at $x = 2$), the masked value it received from Party A (e.g., calculated by Party A using Party A's polynomial at $x = 2$), and the masked value it received from Party C (e.g., calculated by Party C using Party C's polynomial at $x = 2$).

Party C may calculate a sum based on the masked value it kept (e.g., calculated by Party C using Party C's own polynomial at $x = 3$), the masked value it received from Party A (e.g., calculated by Party A using Party A's polynomial at $x = 3$), and the masked value it received from Party B (e.g., calculated by Party B using Party B's polynomial at $x = 3$).

[0059] Parties A through N may then consolidate the results to calculate $\sum_{i \in \{A, B, \dots, N\}} l_i$ and $\sum_{i \in \{A, B, \dots, N\}} g_i$. If the parties used the constants in their polynomials to represent their l values, the sum representing the combined value of all constants used in the polynomials will equal to $\sum_{i \in \{A, B, \dots, N\}} l_i$. Likewise, if the parties used the constants in their polynomials to represent their g values, the sum representing the combined value of all constants used in the polynomials will equal to $\sum_{i \in \{A, B, \dots, N\}} g_i$. In this manner, Parties A through N can calculate $\sum_{i \in \{A, B, \dots, N\}} l_i$ and $\sum_{i \in \{A, B, \dots, N\}} g_i$ without requiring the parties to disclose their l and g values to each other, further preserving the privacy of Parties A through N.

[0060] It is to be understood that while the examples described above discussed using second-degree polynomials to facilitate masking of values provided by three participating parties, such examples are depicted merely for illustrative purposes and are not meant to be limiting. It is to be understood that additional parties may participate in the multiparty summation process discussed above. It is also to be understood that the parties may implement other types of multiparty computation processes to facilitate the calculation of $\sum_{i \in \{A, B, \dots, N\}} l_i$ and $\sum_{i \in \{A, B, \dots, N\}} g_i$. Alternatively, in some embodiments, Parties A through N may calculate $\sum_{i \in \{A, B, \dots, N\}} l_i$ and $\sum_{i \in \{A, B, \dots, N\}} g_i$ without using multiparty computation. It is further to be understood that regardless of

how Parties A through N choose to calculate $\sum_{i \in \{A, B, \dots, N\}} l_i$ and $\sum_{i \in \{A, B, \dots, N\}} g_i$, the calculation is merely one sub-step of step 310. The ultimate goal of step 310 is to facilitate the calculation of the minimum liquidity amount, P_{min} , without requiring the parties to reveal their post transfer liquidity information, $P_A, P_B, \dots, P_N$, as described above.

[0061] At step 312, Parties A through N may be instructed to submit the minimum liquidity amount, P_{min} , to the blockchain system 100 for recordation. In some embodiments, the blockchain system 100 may utilize a consensus protocol to validate the value of P_{min} . The blockchain system 100 may reject the value of P_{min} if no consensus is reached. In some embodiments, the blockchain system 100 may utilize other verification techniques to verify the validity of the value submitted. In some embodiments, the blockchain system 100 may accept and incorporate the submitted value of P_{min} into the blockchain system 100 without additional verification.

[0062] At step 314, the blockchain system 100 may determine whether a deadlock situation exists based on the value of P_{min} . In some embodiments, the blockchain system 100 may determine that a deadlock situation exists if $P_{min} < 0$. The blockchain system 100 may notify one or more Parties A through N that additional supply of liquidity is needed to resolve the deadlock situation. On the other hand, the blockchain system 100 may determine that the RTGS system is not deadlocked if $P_{min} \geq 0$, and report that no additional liquidity is needed. In some embodiments, the blockchain system 100 may implement a smart contract that is programmed to determine whether or not a deadlock situation exists based on whether $P_{min} < 0$ or $P_{min} \geq 0$.

[0063] In some embodiments, the blockchain system 100 may keep track of a total number of transactions pending in the RTGS system and invoke the privacy-preserving deadlock detection method 300 when the total number of pending transactions reaches a predetermined threshold. FIG. 4 illustrates a flow chart of a method 400 for invoking the privacy-preserving deadlock detection method 300, according to an embodiment. The method 400 may be performed by one or more nodes in a blockchain system that maintain the blockchain, e.g., the nodes 102-110 in the blockchain system 100 (FIG. 1).

[0064] At step 402, a node, e.g., the node 102, may determine the number of transactions pending execution on the blockchain system 100. For example, a number of transactions may be included in a new block that is to be accepted by the blockchain system 100. These transactions are still pending execution on the blockchain system 100. The node 102 may determine the number of such transactions in step 402. At step 404, the node 102 may determine whether the number of transactions pending execution on the blockchain system 100 has exceeded a predetermined threshold. At step 406, in response to the determination that the number of transactions pending execution on the blockchain system 100 has exceeded the predetermined threshold, the node 102 may invoke the privacy-preserving deadlock detection method 300 (FIG. 3).

[0065] At step 408, the node 102 may instruct each participating party, e.g., party $i \in \{A, B, \dots, N\}$, to independently calculate its liquidity information, including a current liquidity C_i , an amount T_i to be transferred to (receivable by) party i , an amount F_i to be transferred from (payable by) party i , and a post transfer liquidity P_i . In some embodiments, the node 102 may send instructions to the participating parties via wired

or wireless communication channels.

[0066] At step 410, the node 102 may instruct the participating parties $i \in \{A, B, \dots, N\}$ to perform a joint calculation for determining, e.g., further calculating, an aggregated total amount receivable in the RTGS system, $S_T = \sum_{i \in \{A, B, \dots, N\}} T_i$. The node 102 may also instruct the participating parties i to perform a joint calculation for determining, e.g., further calculating, an aggregated total amount payable in the RTGS system, $S_F = \sum_{i \in \{A, B, \dots, N\}} F_i$. The node 102 may instruct the participating parties to submit the results of their calculations to the blockchain system 100 for recordation. The node 102 may also execute a smart contract to calculate the values of S_T and S_F based on the results submitted by the participating parties, without receiving the actual liquidity information of the respective parties. In some embodiments, the participating parties may communicate with each other and with the blockchain system 100 via wired or wireless communication channels.

[0067] At step 412, the node 102 may verify correctness of the liquidity computation based on the calculated values of S_T and S_F . In some embodiments, the node 102 may require S_T to match, e.g., be equal to, S_F . If the node 102 determines that $S_T \neq S_F$, the node 102 may recognize this mismatch and report it as an error. In some embodiments, the node 102 verify that the liquidity computation is correct if it determines that $S_T = S_F$.

[0068] After the correctness is verified, the node 102 may, at step 414, instruct the parties to jointly carry out a multiparty computation to determine a minimum liquidity amount $P_{min} = \min_{i \in \{A, B, \dots, N\}} \{P_i\}$ as describe in step 310 above. At step 416, the node 102 may instruct the parties to submit the minimum post transfer liquidity amount, P_{min} ,

to the blockchain system 100 for recordation. Upon receiving the minimum post transfer liquidity amount, the node 102 may, at step 418, determine whether or not a deadlock situation exists based on whether $P_{min} < 0$ or $P_{min} \geq 0$.

[0069] In some embodiments, each party, e.g., each of Parties A through N, may keep track of the number of transactions pending actions by that particular party. A particular party may indicate its intention to invoke the privacy-preserving deadlock detection method 300 when the number of transactions pending actions by that particular party reaches a threshold established by that party. In some embodiments, the privacy-preserving deadlock detection method 300 may be invoked if Parties A through N reach a consensus that the privacy-preserving deadlock detection method 300 should be invoked. Once a consensus is reached, a node, e.g., the node 102, may carry out steps 408-418 as described above.

[0070] In some embodiments, the node 102 may carry out additional verification steps to verify the computation correctness of the privacy-preserving deadlock detection method 300. For example, in some embodiments, the node 102 may instruct participating parties to utilize a commitment scheme to calculate commitment values of their independently calculated liquidity information and submit the commitment values to the blockchain system 100 for recordation. Such a commitment scheme is described in the above noted co-pending Singapore Application (Applicant Ref. No. A23582SG), entitled “Methods and Devices for Deadlock Detection in Real-Time Gross Settlement Systems.” Such a commitment scheme may include a cryptographic primitive that allows the parties to commit to certain values while keeping the values hidden from others. In this manner, the parties can submit only the commitment values (as opposed

to their calculated liquidity information) to the blockchain system 100, which can utilize these commitment values to confirm correctness of the liquidity information during the execution of the privacy-preserving deadlock detection method 300. If someone has changed their liquidity information during the execution of the privacy-preserving deadlock detection method 300, a node in the blockchain system 100, e.g., the node 102, can detect the change and invalidate the computation result of the method 300. In such a case, the node 102 may invoke the method 300 again, or report an error.

[0071] Additionally, or alternatively, in some embodiments, the node 102 may instruct participating parties to implement a liquidity-based verification technique described in the above noted co-pending Singapore Application (Applicant Ref. No. A23583SG), entitled “Methods and Devices for Deadlock Detection in Real-Time Gross Settlement Systems.” Such a liquidity-based verification technique may instruct the parties to calculate a total liquidity and compare the calculated total liquidity to a total liquidity known to exist in the RTGS system. If the calculated total liquidity does not equal to the total liquidity known to exist in the RTGS system, a node in the blockchain system 100, e.g., the node 102, may determine that an error has occurred and invalidate the computation result of the method 300.

[0072] It is to be understood that by performing one or more of these verification steps, the blockchain system 100 can ensure that the parties participating in the RTGS system do not misbehave and that deadlocks can be correctly detected. The blockchain system 100 may also ensure that each party i will use the correct liquidity values C_i and P_i for liquidity calculation. The blockchain system 100 may further ensure that no party can change their liquidity values during the deadlock detection process

without being detected by the blockchain system 100.

[0073] FIG. 5 is a block diagram of a deadlock detection apparatus 500, according to an embodiment. The apparatus 500 may be an implementation of a software process, and may correspond to the method 400 (FIG. 4). Referring to FIG. 5, the apparatus 500 may include a determination module 502, an instruction module 504, a receiving module 506, and a detection module 508.

[0074] The determination module 502 may determine whether to invoke privacy-preserving deadlock detection in an RTGS system. In some embodiments, the determination module 502 may make the determination based on the number of transactions pending execution in the RTGS system. Alternatively, or additionally, parties participating in the RTGS system may keep track of the number of transactions pending locally and determine whether to invoke privacy-preserving deadlock detection in the RTGS system based on whether the parties reach a consensus to invoke privacy-preserving deadlock detection.

[0075] The instruction module 504 may provide instructions to the parties participating in the RTGS system. For example, the instruction module 504 may instruct the participating parties to independently calculate their liquidity information. In some embodiments, each party $i \in \{A, B, \dots, N\}$ may calculate its current liquidity C_i , an amount T_i to be transferred to (receivable by) the party, an amount F_i to be transferred from (payable by) the party, and a post transfer liquidity P_i , where $P_i = C_i + T_i - F_i$.

[0076] The instruction module 504 may also instruct the participating parties $i \in \{A, B, \dots, N\}$ to perform a joint calculation for determining, e.g., further calculating, an aggregated total amount receivable in the RTGS system, $S_T = \sum_{i \in \{A, B, \dots, N\}} T_i$, and to

perform a joint calculation for determining, e.g., further calculating, an aggregated total amount payable in the RTGS system, $S_F = \sum_{i \in \{A, B, \dots, N\}} F_i$. The instruction module 504 may further instruct the participating parties to submit calculation results of the joint calculations and a minimum post transfer liquidity amount, P_{min} , to the receiving module 506.

[0077] The receiving module 506 may receive information provided by the participating parties and provide the information received to the detection module 508. The detection module 508 may verify computation correctness of liquidity information of the participating parties and thus, the value P_{min} , and determine whether a deadlock situation exists based on P_{min} after the computation correctness is verified. In some embodiments, the detection module 508 may communicate with a blockchain system, e.g., the blockchain system 100 (FIG. 1), that utilizes a smart contract programmed to determine the computation correctness of P_{min} . The smart contract may confirm computation correctness of liquidity information of the participating parties and thus, the computation correctness of P_{min} , based on whether there is a mismatch between the calculated values of S_T and S_F . In some embodiments, the smart contract may report an error if $S_T \neq S_F$. In some embodiments, the blockchain system may utilize the smart contract, or another smart contract, to determine whether a deadlock situation exists based on P_{min} . The smart contract may determine whether or not a deadlock situation exists based on whether $P_{min} < 0$ or $P_{min} \geq 0$.

[0078] Each of the above described modules may be implemented as software, or hardware, or a combination of software and hardware. For example, each of the above described modules may be implemented using a processor executing

instructions stored in a memory. Also, for example, each the above described modules may be implemented with one or more application specific integrated circuits (ASICs), digital signal processors (DSPs), digital signal processing devices (DSPDs), programmable logic devices (PLDs), field programmable gate arrays (FPGAs), controllers, micro-controllers, microprocessors, or other electronic components, for performing the described methods. Further for example, each of the above described modules may be implemented by using a computer chip or an entity, or implemented by using a product having a certain function. In one embodiment, the apparatus 500 may be a computer, and the computer may be a personal computer, a laptop computer, a cellular phone, a camera phone, a smartphone, a personal digital assistant, a media player, a navigation device, an email receiving and sending device, a game console, a tablet computer, a wearable device, or any combination of these devices.

[0079] For an implementation process of functions and roles of each module in the apparatus 500, references can be made to corresponding steps in the above-described methods. Details are omitted here for simplicity.

[0080] In some embodiments, a computer program product may include a non-transitory computer-readable storage medium having computer-readable program instructions thereon for causing a processor to carry out the above-described methods.

[0081] The computer-readable storage medium may be a tangible device that can store instructions for use by an instruction execution device. The computer-readable storage medium may be, for example, but is not limited to, an electronic storage device, a magnetic storage device, an optical storage device, an electromagnetic storage device, a semiconductor storage device, or any suitable

combination of the foregoing. A non-exhaustive list of more specific examples of the computer-readable storage medium includes the following: a portable computer diskette, a hard disk, a random access memory (RAM), a read-only memory (ROM), an erasable programmable read-only memory (EPROM), a static random access memory (SRAM), a portable compact disc read-only memory (CD-ROM), a digital versatile disk (DVD), a memory stick, a floppy disk, a mechanically encoded device such as punch-cards or raised structures in a groove having instructions recorded thereon, and any suitable combination of the foregoing.

[0082] The computer-readable program instructions for carrying out the above-described methods may be assembler instructions, instruction-set-architecture (ISA) instructions, machine instructions, machine dependent instructions, microcode, firmware instructions, state-setting data, or source code or object code written in any combination of one or more programming languages, including an object oriented programming language, and conventional procedural programming languages. The computer-readable program instructions may execute entirely on a computing device as a stand-alone software package, or partly on a first computing device and partly on a second computing device remote from the first computing device. In the latter scenario, the second, remote computing device may be connected to the first computing device through any type of network, including a local area network (LAN) or a wide area network (WAN).

[0083] The computer-readable program instructions may be provided to a processor of a general-purpose or special-purpose computer, or other programmable data processing apparatus to produce a machine, such that the instructions, which

execute via the processor of the computer or other programmable data processing apparatus, create means for implementing the above-described methods.

[0084] The flow charts and diagrams in the figures illustrate the architecture, functionality, and operation of possible implementations of devices, methods, and computer program products according to various embodiments of the specification. In this regard, a block in the flow charts or diagrams may represent a software program, segment, or portion of code, which comprises one or more executable instructions for implementing specific functions. It should also be noted that, in some alternative implementations, the functions noted in the blocks may occur out of the order noted in the figures. For example, two blocks shown in succession may, in fact, be executed substantially concurrently, or the blocks may sometimes be executed in the reverse order, depending upon the functionality involved. It will also be noted that each block of the diagrams and/or flow charts, and combinations of blocks in the diagrams and flow charts, may be implemented by special purpose hardware-based systems that perform the specified functions or acts, or combinations of special purpose hardware and computer instructions.

[0085] It is appreciated that certain features of the specification, which are, for clarity, described in the context of separate embodiments, may also be provided in combination in a single embodiment. Conversely, various features of the specification, which are, for brevity, described in the context of a single embodiment, may also be provided separately or in any suitable subcombination or as suitable in any other described embodiment of the specification. Certain features described in the context of various embodiments are not essential features of those embodiments, unless noted as

such.

[0086] Although the specification has been described in conjunction with specific embodiments, many alternatives, modifications and variations will be apparent to those skilled in the art. Accordingly, the following claims embrace all such alternatives, modifications and variations that fall within the terms of the claims.

What Is Claimed Is:

1. A computer-implemented method for detecting deadlocks in a real-time gross settlement system, the method comprising:

instructing a plurality of users of the real-time gross settlement system to independently calculate their respective liquidity information, and to perform joint calculations for determining an aggregated total amount receivable in the real-time gross settlement system and an aggregated total amount payable in the real-time gross settlement system;

receiving calculation results of the joint calculations from the plurality of users and calculating the aggregated total amount receivable and the aggregated total amount payable based on the received calculation results;

verifying correctness of the liquidity information of the plurality of users based on the aggregated total amount receivable and the aggregated total amount payable; and

after the correctness of the liquidity information of the plurality of users is verified, determining whether a deadlock exists in the real-time gross settlement system.

2. The method of claim 1, wherein the determining whether a deadlock exists in the real-time gross settlement system comprises:

instructing the plurality of users to jointly determine a minimum post transfer liquidity, wherein the minimum post transfer liquidity is jointly determined based on the liquidity information of each of the plurality of users without disclosing by each of the plurality of users of its liquidity information to another user;

receiving the jointly determined minimum post transfer liquidity; and
determining whether a deadlock exists based on the minimum post transfer liquidity.

3. The method of claim 1, wherein the instructing the plurality of users to independently calculate their respective liquidity information comprises:

instructing each of the plurality of users, user i , to independently calculate its liquidity information including a current liquidity C_i , an amount receivable T_i , an amount payable F_i , and a post transfer liquidity P_i .

4. The method of claim 3, wherein the instructing the plurality of users to perform joint calculations comprises:

instructing the plurality of users $i \in \{A, B, \dots, N\}$ to perform a joint calculation for determining an aggregated total amount receivable S_T , wherein the aggregated total amount receivable S_T is calculated without receiving from the plurality of users their respective liquidity information.

5. The method of claim 3, wherein the instructing the plurality of users to perform joint calculations comprises:

instructing the plurality of users $i \in \{A, B, \dots, N\}$ to perform a joint calculation for determining an aggregated total amount payable S_F , wherein the aggregated total amount payable S_F is calculated without receiving from the plurality of users their respective liquidity information.

6. The method of any of claims 4 or 5, wherein the verifying correctness of the liquidity information of the plurality of users comprises:

comparing the aggregated total amount receivable S_T to the aggregated total amount payable S_F .

7. The method of claim 6, further comprising:

reporting an error if the aggregated total amount receivable S_T does not match the aggregated total amount payable S_F .

8. The method of any preceding claim, wherein the real-time gross settlement system is implemented with a blockchain system.

9. The method of claim 8, further comprising:

executing a smart contract recorded on a blockchain of the blockchain system to verify correctness of the liquidity information of the plurality of users.

10. The method of claim 9, further comprising:

consolidating, by the smart contract, the calculation results of the joint calculations submitted by the plurality of users to determine the aggregated total amount receivable and the aggregated total amount payable.

11. A device for detecting deadlocks in a real-time gross settlement system, comprising:

one or more processors; and

one or more computer-readable memories coupled to the one or more processors and having instructions stored thereon that are executable by the one or more processors to perform the method of any of claims 1 to 10.

12. An apparatus for detecting deadlocks in a real-time gross settlement system, the apparatus comprising a plurality of modules for performing the method of any of claims 1 to 10.

13. A non-transitory computer-readable medium having stored therein instructions that, when executed by a processor of a device, cause the device to perform the method of any of claims 1 to 10.