

CONSENSUS METHOD AND APPARATUS

TECHNICAL FIELD

[0001] The present application relates to the field of computer technologies, and in particular, to a consensus method and apparatus.

5

BACKGROUND

[0002] Currently, the blockchain technology is widely used, and a decentralized mode in the blockchain technology ensures that data is not easily tampered with, thereby improving security.

[0003] In practice, a blockchain network that includes multiple nodes (the nodes
10 can be considered as devices in the blockchain network that participate in processing services) can provide corresponding services for client devices. Nodes in the blockchain network perform processing for service requests of the client devices, and feed back processing results to the client devices. In this process, processing results generated by the independently running nodes may be inconsistent. To ensure that the
15 client device can receive a correct processing result, a consensus among the nodes is implemented by using a Practical Byzantine Fault Tolerance (PBFT) algorithm (that is, a correct processing result can be jointly acknowledged or accepted by the nodes).

[0004] In a process of using the PBFT algorithm, a consensus is usually performed in a view. In a view, one node in a blockchain network is used as a primary
20 node (primary), and other nodes are used as backup nodes (backup). In this case, the primary node receives a service request of the client device and broadcasts the service request to all the backup nodes, and the primary node initiates a consensus. The nodes that reach the consensus perform processing for the service request and feed back a processing result to the client device.

[0005] In the existing technology, a backup node initiates view switching, and the
25 view switching initiated by the backup node usually needs to be acknowledged or accepted by other nodes in the view. The backup node initiates a view switching request to other nodes (including the primary node) in the view, that is, initiates a

consensus on the view switching request to other nodes (this consensus still uses PBFT. Different from the process of the consensus on the service request, in a consensus process based on the view switching request, each node suspends the consensus on the service request. Therefore, the consensus on the view switching request is essentially an additional consensus process). After a predetermined quantity of nodes reach a consensus, it is determined that a backup node becomes a new primary node. The new primary node broadcasts a new view message to complete view switching.

[0006] However, in the previous mechanism, an additional consensus process needs to be performed for the view switching initiated by the backup node, and the additional consensus process increases a system computation amount. In addition, in the consensus process of view switching, a consensus can be reached after being confirmed by a predetermined quantity of nodes. Finally, a new primary node broadcasts a new view message, and the entire process consumes a period of time. Apparently, an existing view switching way not only increases a computation amount of the system, but also increases time consumed for processing a service request. Consequently, relatively low processing efficiency is caused.

SUMMARY

[0007] Implementations of the present application provide a consensus method and apparatus, to alleviate a problem that a current view switching way increases a computation amount of a blockchain network and increases processing time consumption.

[0008] An implementation of the present application provides a consensus method, where the method includes the following: monitoring, by a blockchain primary node, triggering of a view switching condition; selecting, by the blockchain primary node, a successor node when the triggering of the view switching condition is monitored; and switching, by the blockchain primary node based on the successor node, a current view to a next view that uses the successor node as a succeeding blockchain primary node, so that the succeeding blockchain primary node initiates a consensus in the next view.

[0009] An implementation of the present application provides a consensus apparatus, where the apparatus includes the following: a monitoring module,

configured to monitor triggering of a view switching condition; a node determining module, configured to select a successor node when the monitoring module monitors the triggering of the view switching condition; and a view switching module, configured to switch, based on the successor node, a current view to a next view that
 5 uses the successor node as a succeeding blockchain primary node, so that the succeeding blockchain primary node initiates a consensus in the next view.

[0010] The implementations of the present application provide a consensus method and apparatus. In any view, the blockchain primary node actively monitors the triggering of the view switching condition. If the view switching condition is
 10 triggered, the blockchain primary node needs to perform view switching. Further, the blockchain primary node selects a successor node from other blockchain nodes as a blockchain primary node in the next view. Accordingly, the blockchain primary node performs view switching. In the switched view, the successor node is used as a new blockchain primary node to process a service. In addition, view switching is still
 15 performed based on the previous process. Apparently, the previous view switching is initiated by the blockchain primary node. Such a way prevents the blockchain backup node from initiating the view switching consensus. In other words, an additional consensus can be prevented. As such, an additional computation amount and processing time consumption in a blockchain network can be reduced.

20 **BRIEF DESCRIPTION OF DRAWINGS**

[0011] The accompanying drawings described here are intended to provide a further understanding of the present application, and constitute a part of the present application. The illustrative implementations of the present application and descriptions thereof are intended to describe the present application, and do not
 25 constitute limitations on the present application. In the accompanying drawings:

[0012] FIG. 1a illustrates an architecture that a consensus process is based on, according to an implementation of the present application;

[0013] FIG. 1b illustrating a consensus process, according to an implementation of the present application;

30 **[0014]** FIG. 2 is a schematic diagram illustrating a consensus process that is based on a three phase protocol in any view, according to an implementation of the present application;

[0015] FIG. 3 is a schematic diagram illustrating an execution process of an application instance of view switching, according to an implementation of the present application; and

[0016] FIG. 4 is a schematic structural diagram illustrating a consensus apparatus,
5 according to an implementation of the present application.

DESCRIPTION OF IMPLEMENTATIONS

[0017] To make the objectives, technical solutions, and advantages of the present application clearer, the following clearly and comprehensively describes the technical solutions of the present application with reference to specific implementations and
10 accompanying drawings of the present application. Apparently, the described implementations are merely some rather than all of the implementations of the present application. All other implementations obtained by a person of ordinary skill in the art based on the implementations of the present application without creative efforts shall fall within the protection scope of the present application.

15 [0018] As described above, in a process that nodes in a blockchain network perform a consensus by using the PBFT algorithm, once a blockchain primary node becomes faulty, a blockchain backup node initiates view switching. An additional consensus is needed for the view switching initiated by the blockchain backup node. That is, the view switching can be completed only after being acknowledged or
20 accepted by other blockchain nodes. Apparently, an additional consensus process increases a computation amount of a blockchain network, and also increases processing time consumption.

[0019] Accordingly, an implementation of the present application provides a consensus method. For a blockchain primary node in any view, after a consensus
25 ends, the blockchain primary node initiates view switching, to change a blockchain primary node, and no additional consensus process is needed. For ease of description, a blockchain primary node is simply referred to as a primary node, and a blockchain backup node is simply referred to as a backup node. In addition, "nodes" described below should be understood as nodes in a blockchain network that participate in a
30 consensus.

[0020] It is worthwhile to note that, in an implementation of the present application, an architecture used in the consensus method is shown in FIG. 1a. It can

be seen from FIG. 1a that a blockchain network includes multiple nodes, and multiple client devices can perform service interaction with the blockchain network. An application type of the blockchain network can be a consortium blockchain network or a private blockchain network, and the blockchain network can provide a service for a user. The node includes but is not limited to a server, a computer, a mobile device, and other devices that have a computation processing function. The client device can be configured to run a browser, an application, etc., and the client device can be an end-user device, a server, or a database. Implementations are not limited here.

[0021] Based on the relationship architecture shown in FIG. 1a, a consensus process provided in an implementation of the present application is shown in FIG. 1b. The process includes the following steps.

[0022] S101: A primary node monitors triggering of a view switching condition.

[0023] In this implementation of the present application, the view switching condition can be considered as a condition that needs be satisfied for performing view switching. For example, the primary node does not broadcast a service request or complete a consensus within a predetermined period of time.

[0024] In a possible way in this implementation of the present application, a timer can be set in the primary node to monitor the triggering of the view switching condition. For example, a timer can be used to time a behavior of broadcasting a service request by the primary node, to detect whether the behavior of broadcasting a service request by the primary node expires. The timer can be considered as a timing function or service that runs on the primary node, and certainly does not constitute limitations on the present application.

[0025] S102: The primary node selects a successor node when the triggering of the view switching condition is monitored.

[0026] If the view switching condition is triggered, the primary node needs to perform view switching. It is worthwhile to note that, in any view, there is only one primary node, and other nodes are backup nodes. Therefore, view switching indicates switching of the primary node. Therefore, in the present step, the primary node selects a successor node as a next primary node (the successor node in this implementation of the present application is not a same node as the primary node in the current view, that is, the primary node in the current view cannot be used as a successor node).

[0027] S103: The primary node switches, based on the successor node, a current view to a next view that uses the successor node as a succeeding primary node, so that

the succeeding primary node initiates a consensus in the next view.

[0028] After the successor node is determined, the primary node performs view switching. In an existing way that a backup node initiates a view switching consensus, a process that the backup node initiates the view switching consensus can be considered as a process of "impeachment" of the primary node in the view. Different from the existing way, in this implementation of the present application, a process that the primary node independently performs view switching can be considered as a process of "active abdication", and the primary node performs the view switching without initiating a consensus. Apparently, an additional consensus process is also prevented. After view switching is performed, a newly appointed primary node is responsible for initiating a consensus in the switched view, and it can be understood that the newly appointed primary node also performs the previous view switching process. Details are omitted here for simplicity.

[0029] According to the previous steps, in any view, the primary node actively monitors the triggering of the view switching condition. If the view switching condition is triggered, the primary node needs to perform view switching. Further, the primary node selects a successor node from other nodes as a primary node in the next view. Accordingly, the primary node performs view switching. In the switched view, the successor node is used as a new primary node to process a service. In addition, view switching is still performed based on the previous process. Apparently, the previous view switching is initiated by the primary node. Such a way prevents a backup node from initiating a view switching consensus. In other words, an additional consensus can be prevented. As such, an additional computation amount and processing time consumption in a blockchain network can be reduced.

[0030] In practice, there are different view switching conditions. The following describes in detail the triggering of the view switching condition.

[0031] First scenario:

[0032] In practice, a client device sends a service request to a primary node. In a normal status, after receiving the service request, the primary node broadcasts the service request to backup nodes in a view, to perform a consensus on the service request. However, the primary node can be an abnormal node, and does not broadcast the service request for a long time after receiving the service request. As such, the backup node initiates a view switching consensus. Therefore, to prevent a view switching consensus initiated by the backup node because the primary node does not

broadcast a service request within a predetermined period of time, the primary node independently performs timing, and actively monitors an expiration phenomenon of the primary node.

[0033] In other words, in this scenario, if the view switching condition is that the primary node does not broadcast a service request within a predetermined period of time, triggering the view switching condition includes the following: receiving, by the primary node, a service request, and failing to initiate a consensus on the service request within a predetermined time.

[0034] In an actual operation, timing can be implemented by a program or a service that has a timing function in the primary node, for example, the previously mentioned timer. Timing can be started from a moment at which the primary node receives the service request. The predetermined time can be set to 5s, 10s, etc., can be determined based on a requirement in practice, and does not constitute limitations on the present application.

[0035] Second scenario:

[0036] Different from the previous scenario, in this scenario, after receiving a service request, the primary node broadcasts the service request to backup nodes in a current view. In other words, the primary node has initiated a consensus on the service request before the predetermined time expires. Correspondingly, nodes in the view perform a consensus on the service request, and generate a consensus result.

[0037] It is worthwhile to note that, here, based on an existing view switching mechanism, if a consensus result is that a consensus fails, the backup node initiates a view switching consensus. Apparently, that a consensus fails can be considered as a view switching condition. In other words, in this scenario, when the consensus fails, the primary node actively performs view switching, to prevent the backup node from initiating an additional view switching consensus.

[0038] In addition, based on the existing view switching mechanism, if a consensus result is that a consensus is reached, the primary node continues to initiate a consensus on another service request. However, the primary node may be faulty in a subsequent running process. Once the primary node becomes faulty, the backup node still initiates a view switching consensus. Therefore, to prevent this case, in this implementation of the present application, after a consensus is reached, the primary node still performs view switching.

[0039] It can be seen that, in this scenario, regardless of whether the consensus

result is that a consensus is reached or that a consensus fails, the primary node performs view switching after determining the consensus result. In other words, triggering the view switching condition includes the following: receiving, by the primary node, a service request, initiating a consensus on the service request, and
 5 determining a consensus result.

[0040] That is, in this scenario, the primary node needs to determine that a consensus has been reached or that a consensus fails. The following describes in detail how the primary node determines that a consensus has been reached or that a consensus fails.

10 **[0041]** First, it is worthwhile to note that a consensus process based on a service request is essentially a consensus process based on a three phase protocol. The three phases include a pre-prepare phase, a prepare phase, and a commit phase, and form a complete consensus process. At each phase, nodes (including both a primary node and backup nodes) send consensus messages to each other. That is, for each node in the
 15 view, entering different phases needs to be acknowledged or accepted by other nodes. Therefore, each of the three phases can be considered as a consensus process. Usually, when all nodes enter a commit phase, it can be considered that the consensus process is completed.

[0042] FIG. 2 is a consensus process of nodes based on a three phase protocol in a view. In FIG. 2, a client device initiates a service request to a node numbered 0 (replica 0, that is, a primary node), and the primary node broadcasts the service request to backup nodes (replica 1, replica 2, and replica 3), and starts to perform a three phase consensus. The nodes that reach a consensus process the service request and feed back a processing result to the client device.

25 **[0043]** Accordingly, the following operations are performed.

[0044] 1. The primary node determines that a consensus fails.

[0045] In this implementation of the present application, that a consensus fails is represented as consensus process expiration (simply referred to as consensus expiration below, where the consensus expiration means that time consumed by a
 30 consensus process exceeds predetermined consensus duration, and the time consumed by the consensus process can be counted from a moment at which the primary node initiates the consensus). The reasons are as follows.

[0046] In one case, the primary node is a faulty node (if a node is faulty, it can be considered that data used to perform a consensus in the node is faulty or consensus

logic in the node is faulty). That is, service requests sent by the primary node to backup nodes may include incorrect data (for example, an incorrect service request sequence number). Ensured by the BPFT mechanism, the backup nodes check the service requests broadcast by the primary node. Once the incorrect data is included in a service request, a normal backup node does not acknowledge or accept the service request. In this case, the primary node may repeat the process of sending service requests, and consequently, the consensus expires.

[0047] Alternatively, in another case, the primary node is also a faulty node. In this case, the primary node may send an incorrect notification message indicating a consensus phase to other backup nodes, that is, the primary node "incorrectly" considers that a certain phase has been entered. In this case, backup nodes perform a consensus on the notification message of the primary node, to confirm the authenticity of the notification message of the primary node. Similarly, a normal backup node still does not acknowledge or accept a notification message sent by the primary node. In this case, the primary node may repeat a process of sending an incorrect notification message, and consequently, the consensus expires.

[0048] Certainly, the previous content is only two possible cases that may lead to consensus expiration in practice, and should not be considered as limitations on the present application. Apparently, it can be seen from the previous content that, once a consensus expires, the consensus fails.

[0049] Therefore, in this implementation of the present application, the primary node can detect whether a consensus fails by monitoring an overall time consuming way of a consensus process. Once the consensus process expires, the primary node immediately initiates a view switching operation, and an additional consensus process of view switching initiated by a backup node can be prevented. That is, in this implementation of the present application, a process that the primary node determines that the consensus fails can be as follows: The primary node monitors time consumed by a consensus process from a moment at which a consensus on the service request is initiated among the backup nodes in the view, and determines that the consensus fails when it is detected that the time consumed by the consensus exceeds a predetermined time.

[0050] 2. The primary node determines that a consensus is reached.

[0051] It can be seen from the previous three phase protocol that if a node enters a commit phase, the node can process a service request, and feed back a generated

processing result to a client device. In addition, to enter a phase, each node needs to be acknowledged or accepted by other nodes in a view. Therefore, if a node enters a commit phase, it indicates that the node is acknowledged or accepted by other nodes. It can be seen that, if the primary node enters the commit phase, it indicates that the consensus has been reached. This is because in the PBFT mechanism, if a node enters a phase, it indicates that the status of the node is acknowledged or accepted by most nodes in the view. Accordingly, it indicates that most nodes are correct nodes.

[0052] Therefore, in such a way, a process that the primary node determines that a consensus is reached can be as follows: The primary node monitors a corresponding phase of the primary node, and when the primary node detects that the primary node enters a commit phase and predetermined consensus duration does not expire, determines that the consensus is completed. That is, when confirming that the primary node enters the commit phase, the primary node further needs to ensure that time consumed by the primary node to enter the commit phase does not exceed the predetermined time.

[0053] In another way in this implementation of the present application, the primary node may not send a notification message to other nodes (that is, the primary node may be a faulty node). However, the primary node can still receive a notification message sent by a backup node. In this case, if a predetermined quantity of nodes enter a commit phase, it can be considered that the consensus is completed.

[0054] In practice, after a node enters a commit phase, the node usually sends notification messages to other nodes in the view, where the notification message can be, for example, $\langle \text{commit}, v, n, D(m) \rangle$, where "commit" indicates that the node has entered a commit phase, "v" indicates a view number, "n" indicates a sequence number of a service request, and "D(m)" indicates a signature performed on the service request by the node that sends the notification message.

[0055] The primary node can collect statistics about notification messages that are received by the primary node and that indicate that a commit phase is entered. If a quantity of received notification messages is greater than $2f+1$, it indicates that enough nodes reach a consensus. Then, it indicates that the consensus is completed. f is a maximum quantity of tolerable incorrect nodes in the PBFT mechanism. In this case, the primary node can determine that the consensus is completed.

[0056] Therefore, a process that the primary node determines that the consensus is reached can also be as follows: The primary node monitors a notification message that

is received by the primary node and that indicates that the backup node enters a commit phase, and when the primary node detects that a quantity of received notification messages exceeds a predetermined quantity and predetermined consensus duration does not expire, determines that the consensus is completed.

5 [0057] After the consensus is completed, the primary node initiates view switching, to change the primary node and enter the new view.

[0058] The following describes a view switching process in this implementation of the present application.

[0059] In the PBFT mechanism, each view has a corresponding number. For
10 example, v in the previous example represents the number of the current view. Correspondingly, each node in the blockchain network has a corresponding number. If there are a total of R nodes in the blockchain network, the numbers of the nodes are 0 to $R-1$, for example, replica 0, replica 1, ..., and replica $R-1$. There is a relationship between a node number and a view number. If replica p is used to represent the node
15 with the number p , the node number and a view number satisfy the following equation: $p=v \bmod R$, where v is an integer from 0 to positive infinity.

[0060] This relationship indicates that the node number p is obtained after the view number v modulo the quantity R of nodes included in a blockchain network.

[0061] In other words, because v ranges from 0 to $R-1$, it is ensured that the
20 identity of the primary node is sequentially handed over to different nodes. For example, if the primary node of the current view is replica 0 (that corresponds to the view number 0), the primary node in the next view (numbered 1) is replica 1. All the nodes are traversed in this way.

[0062] Therefore, it can be seen that, in this implementation of the present
25 application, a view switching process includes the following: determining, by the primary node, a number of the primary node; determining, based on the number of the primary node, a node whose number is arranged after the number of the primary node; generating a view switching notification message based on the determined node; and sending the view switching message to each backup node to perform view switching,
30 so that the determined node becomes a primary node in a next view.

[0063] The following uses a specific application instance for description. As shown in FIG. 3, the instance includes the following steps:

[0064] S301: Primary node p in a view numbered v receives a service request sent by a client device, and performs timing.

[0065] S302: When a predetermined time expires, determine whether to initiate a consensus on the service request to each backup node in the view; if yes, perform step S303; and otherwise, perform step S305.

[0066] S303: Obtain a consensus result.

5 **[0067]** S304: Determine whether the consensus is reached, and perform step S305.

[0068] S305: Switch view v to view $v+1$, and determine a node numbered $p+1$ as a primary node in view $v+1$.

10 **[0069]** The previous view switching is initiated by the primary node. Such a way prevents a backup node from initiating a view switching consensus.

[0070] The above is a consensus method provided in the implementation of the present application. Based on the same idea, the implementations of the present application further provide a consensus apparatus. As shown in FIG. 4, for any view, the consensus apparatus includes the following: a monitoring module 401, configured to monitor triggering of a view switching condition; a node determining module 402, configured to select a successor node when the monitoring module monitors the triggering of the view switching condition; and a view switching module 403, configured to switch, based on the successor node, a current view to a next view that uses the successor node as a succeeding blockchain primary node, so that the succeeding blockchain primary node initiates a consensus in the next view.

[0071] In response to determining that a service request is received and a consensus on the service request is not initiated within a predetermined time, the monitoring module 401 determines that the triggering of the view switching condition is monitored.

25 **[0072]** In response to determining that a service request is received, a consensus on the service request is initiated, and a consensus result is determined, the monitoring module 401 determines that the triggering of the view switching condition is monitored.

30 **[0073]** The node determining module 402 determines a next view of the current view, and determines a successor node that corresponds to the next view.

[0074] The view switching module 403 switches the current view to the determined next view, where the successor node is used as a primary node in the next view.

[0075] A node in any view includes a node in a consortium blockchain network or

a private blockchain network.

[0076] In the 1990s, whether a technical improvement is a hardware improvement (for example, an improvement to a circuit structure, such as a diode, a transistor, or a switch) or a software improvement (an improvement to a method procedure) can be clearly distinguished. However, as technologies develop, current improvements to many method procedures can be considered as direct improvements to hardware circuit structures. A designer usually programs an improved method procedure into a hardware circuit, to obtain a corresponding hardware circuit structure. Therefore, a method procedure can be improved by using a hardware entity module. For example, a programmable logic device (PLD) (for example, a field programmable gate array (FPGA)) is such an integrated circuit, and a logical function of the PLD is determined by a user through device programming. The designer performs programming to "integrate" a digital system to a PLD without requesting a chip manufacturer to design and produce an application-specific integrated circuit chip. In addition, at present, instead of manually manufacturing an integrated circuit chip, such programming is mostly implemented by using "logic compiler" software. The logic compiler software is similar to a software compiler used to develop and write a program. Original code needs to be written in a particular programming language for compilation. The language is referred to as a hardware description language (HDL). There are many HDLs, such as Advanced Boolean Expression Language (ABEL), Altera Hardware Description Language (AHDL), Confluence, Cornell University Programming Language (CUPL), HDCal, Java Hardware Description Language (JHDL), Lava, Lola, MyHDL, PALASM, and Ruby Hardware Description Language (RHDL). The very-high-speed integrated circuit hardware description language (VHDL) and Verilog are most commonly used. A person skilled in the art should also understand that a hardware circuit that implements a logical method procedure can be readily obtained once the method procedure is logically programmed by using the several described hardware description languages and is programmed into an integrated circuit.

[0077] A controller can be implemented by using any appropriate method. For example, the controller can be a microprocessor or a processor, or a computer-readable medium that stores computer-readable program code (such as software or firmware) that can be executed by the microprocessor or the processor, a logic gate, a switch, an application-specific integrated circuit (ASIC), a programmable

logic controller, or a built-in microprocessor. Examples of the controller include but are not limited to the following microprocessors: ARC 625D, Atmel AT91SAM, Microchip PIC18F26K20, and Silicon Labs C8051F320. The memory controller can also be implemented as a part of the control logic of the memory. A person skilled in the art also knows that, in addition to implementing the controller by using the computer-readable program code, logic programming can be performed on method steps to allow the controller to implement the same function in forms of the logic gate, the switch, the application-specific integrated circuit, the programmable logic controller, and the built-in microcontroller. Therefore, the controller can be considered as a hardware component, and an apparatus configured to implement various functions in the controller can also be considered as a structure in the hardware component. Or the apparatus configured to implement various functions can even be considered as both a software module implementing the method and a structure in the hardware component.

[0078] The system, apparatus, module, or unit illustrated in the previous implementations can be implemented by using a computer chip or an entity, or can be implemented by using a product having a certain function. A typical implementation device is a computer. The computer can be, for example, 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 device, a game console, a tablet computer, or a wearable device, or a combination of any of these devices.

[0079] For ease of description, the apparatus above is described by dividing functions into various units. Certainly, when the present application is implemented, a function of each unit can be implemented in one or more pieces of software and/or hardware.

[0080] A person skilled in the art should understand that an implementation of the present disclosure can be provided as a method, a system, or a computer program product. Therefore, the present disclosure can use a form of hardware only implementations, software only implementations, or implementations with a combination of software and hardware. Moreover, the present disclosure can use a form of a computer program product that is implemented on one or more computer-usable storage media (including but not limited to a disk memory, a CD-ROM, an optical memory, etc.) that include computer-usable program code.

[0081] The present disclosure is described with reference to the flowcharts and/or

block diagrams of the method, the device (system), and the computer program product based on the implementations of the present disclosure. It is worthwhile to note that computer program instructions can be used to implement each process and/or each block in the flowcharts and/or the block diagrams and a combination of a process and/or a block in the flowcharts and/or the block diagrams. These computer program instructions can be provided for a general-purpose computer, a dedicated computer, an embedded processor, or a processor of another programmable data processing device to generate a machine, so that the instructions executed by the computer or the processor of the another programmable data processing device generate an apparatus for implementing a specific function in one or more processes in the flowcharts and/or in one or more blocks in the block diagrams.

[0082] These computer program instructions can be stored in a computer-readable memory that can instruct the computer or the another programmable data processing device to work in a specific way, so that the instructions stored in the computer-readable memory generate an artifact that includes an instruction apparatus. The instruction apparatus implements a specific function in one or more processes in the flowcharts and/or in one or more blocks in the block diagrams.

[0083] These computer program instructions can be loaded onto the computer or another programmable data processing device, so that a series of operations and steps are performed on the computer or the another programmable device, thereby generating computer-implemented processing. Therefore, the instructions executed on the computer or the another programmable device provide steps for implementing a specific function in one or more processes in the flowcharts and/or in one or more blocks in the block diagrams.

[0084] In a typical configuration, a computing device includes one or more processors (CPU), one or more input/output interfaces, one or more network interfaces, and one or more memories.

[0085] The memory can include a non-persistent memory, a random access memory (RAM), a non-volatile memory, and/or another form that are in a computer-readable medium, for example, a read-only memory (ROM) or a flash memory (flash RAM). The memory is an example of the computer-readable medium.

[0086] The computer-readable medium includes persistent, non-persistent, movable, and unmovable media that can store information by using any method or technology. The information can be a computer-readable instruction, a data structure,

a program module, or other data. Examples of a computer storage medium include but are not limited to a parameter random access memory (PRAM), a static random access memory (SRAM), a dynamic random access memory (DRAM), another type of random access memory (RAM), a read-only memory (ROM), an electrically erasable programmable read-only memory (EEPROM), a flash memory or another memory technology, a compact disc read-only memory (CD-ROM), a digital versatile disc (DVD) or another optical storage, a cassette magnetic tape, a magnetic tape/magnetic disk storage or another magnetic storage device. The computer storage medium can be used to store information accessible by the computing device. Based on the definition in the present specification, the computer-readable medium does not include transitory computer-readable media (transitory media) such as a modulated data signal and carrier.

[0087] It is worthwhile to further note that, the terms "include", "contain", or their any other variants are intended to cover a non-exclusive inclusion, so a process, a method, a product or a device that includes a list of elements not only includes those elements but also includes other elements which are not expressly listed, or further includes elements inherent to such a process, method, product or device. Without more constraints, an element preceded by "includes a ..." does not preclude the existence of additional identical elements in the process, method, product or device that includes the element.

[0088] A person skilled in the art should understand that an implementation of the present application can be provided as a method, a system, or a computer program product. Therefore, the present application can use a form of hardware only implementations, software only implementations, or implementations with a combination of software and hardware. Moreover, the present application can use a form of a computer program product that is implemented on one or more computer-usable storage media (including but not limited to a disk memory, a CD-ROM, an optical memory, etc.) that include computer-usable program code.

[0089] The present application can be described in the general context of computer executable instructions executed by a computer, for example, a program module. Generally, the program module includes a routine, a program, an object, a component, a data structure, etc. executing a specific task or implementing a specific abstract data type. The present application can also be practiced in distributed computing environments. In the distributed computing environments, tasks are

performed by remote processing devices connected through a communications network. In a distributed computing environment, the program module can be located in both local and remote computer storage media including storage devices.

[0090] The implementations in the present specification are described in a progressive way. For same or similar parts of the implementations, references can be made to the implementations. Each implementation focuses on a difference from other implementations. Particularly, a system implementation is basically similar to a method implementation, and therefore, is described briefly. For related parts, references can be made to related descriptions in the method implementation.

[0091] The previous implementations are implementations of the present application, and are not intended to limit the present application. A person skilled in the art can make various modifications and changes to the present application. Any modification, equivalent replacement, or improvement made without departing from the spirit and principle of the present application shall fall within the scope of the claims in the present application.

CLAIMS

What is claimed is:

1. A consensus method, wherein the method comprises:

5 monitoring, by a blockchain primary node, triggering of a view switching condition;

selecting, by the blockchain primary node, a successor node when the triggering of the view switching condition is monitored; and

switching, by the blockchain primary node based on the successor node, a current view to a next view that uses the successor node as a succeeding blockchain primary node, so that the succeeding blockchain primary node initiates a consensus in the next view.

2. The method according to claim 1, wherein the triggering of a view switching condition comprises:

15 receiving, by the blockchain primary node, a service request; and
determining that a consensus on the service request fails to be initiated with a predetermined time.

3. The method according to claim 1, wherein the triggering of a view switching condition comprises:

receiving, by the blockchain primary node, a service request;
initiating a consensus on the service request; and
determining a consensus result.

25 4. The method according to claim 1, wherein the selecting, by the blockchain primary node, a successor node comprises:

determining, by the blockchain primary node, a next view of the current view; and

selecting a successor node that corresponds to the next view; and

30 wherein the switching a current view to a next view that uses the successor node as a succeeding blockchain primary node comprises:

switching the current view to the determined next view, wherein the successor node is used as a blockchain primary node in the next view.

5. The method according to any one of claims 1 to 4, wherein a blockchain node in any view comprises a node in a consortium blockchain network or a private blockchain network.

5 6. A consensus apparatus, wherein the apparatus comprises:
a monitoring module, configured to monitor triggering of a view switching condition;

a node determining module, configured to select a successor node when the monitoring module monitors the triggering of the view switching condition; and

10 a view switching module, configured to switch, based on the successor node, a current view to a next view that uses the successor node as a succeeding blockchain primary node, so that the succeeding blockchain primary node initiates a consensus in the next view.

15 7. The apparatus according to claim 6, wherein in response to determining that a service request is received and a consensus on the service request is not initiated within a predetermined time, the monitoring module determines that the triggering of the view switching condition is monitored.

20 8. The apparatus according to claim 6, wherein in response to determining that a service request is received, a consensus on the service request is initiated, and a consensus result is determined, the monitoring module determines that the triggering of the view switching condition is monitored.

25 9. The apparatus according to claim 6, wherein the node determining module determines a next view of the current view, and determines a successor node that corresponds to the next view; and

wherein the view switching module switches the current view to the determined next view, wherein the successor node is used as a blockchain primary node in the
30 next view.

10. The apparatus according to any one of claims 6 to 9, wherein a blockchain node in any view comprises a node in a consortium blockchain network or a private blockchain network.