

(21) Application No: **2008755.7**
 (22) Date of Filing: **09.06.2020**

(51) INT CL: **H04L 12/26** (2006.01) **H04L 12/24** (2006.01)

(56) Documents Cited: **US 5563875 A** **US 20140219078 A1**

(71) Applicant(s):
British Telecommunications Public Limited Company
81 Newgate street, London, EC1A 7AJ,
United Kingdom

(58) Field of Search:
 INT CL **H04L**
 Other: **WPI, EPODOC**

(72) Inventor(s):
Jonathan Roscoe
Gabriele Gelardi

(74) Agent and/or Address for Service:
BT Group Legal
Intellectual Property Department, Ground Floor,
Faraday Building, 1 Knightrider Street, LONDON,
EC4V 5BT, United Kingdom

(54) Title of the Invention: **Performance-based network fault localisation**
 Abstract Title: **Performance-based network fault localisation**

(57) A method of identifying a failure in a portion of a network connecting a plurality of peer devices. A first peer transmits a record including a timestamp of transmission to a second peer that adds a timestamp for receipt and saves the record in a database. The timestamps are used to judge underperformance of the route between the peers. Sending the record is repeated, changing one or more of the first or second peers or the route therebetween, to identify a route and component associated with the fault. The database may be a distributed transaction database such as a blockchain. The changing step may include changing one of the peers to a component in the route of transmission. The first peer may be changed to the next first network component in the route 404. The second peer may be changed to a last network component in the route 406. When changing the route, a new route may include a portion of the initial route 410, 412.

FIGURE 4

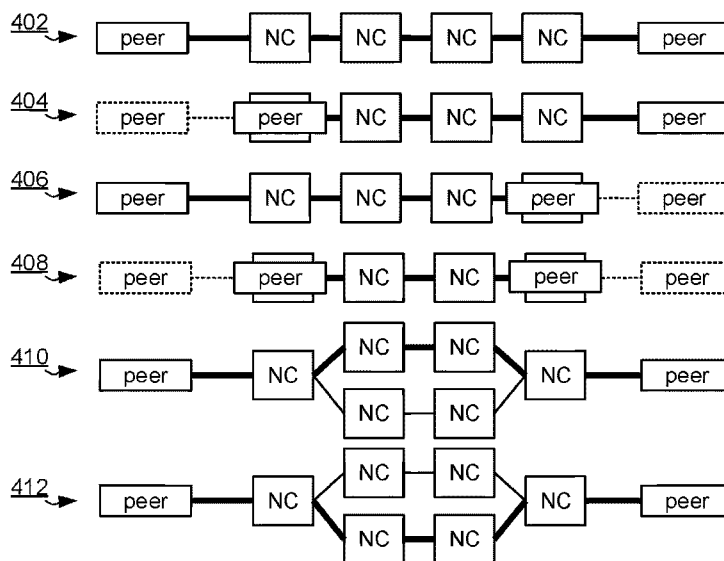


FIGURE 1




FIGURE 2




FIGURE 3

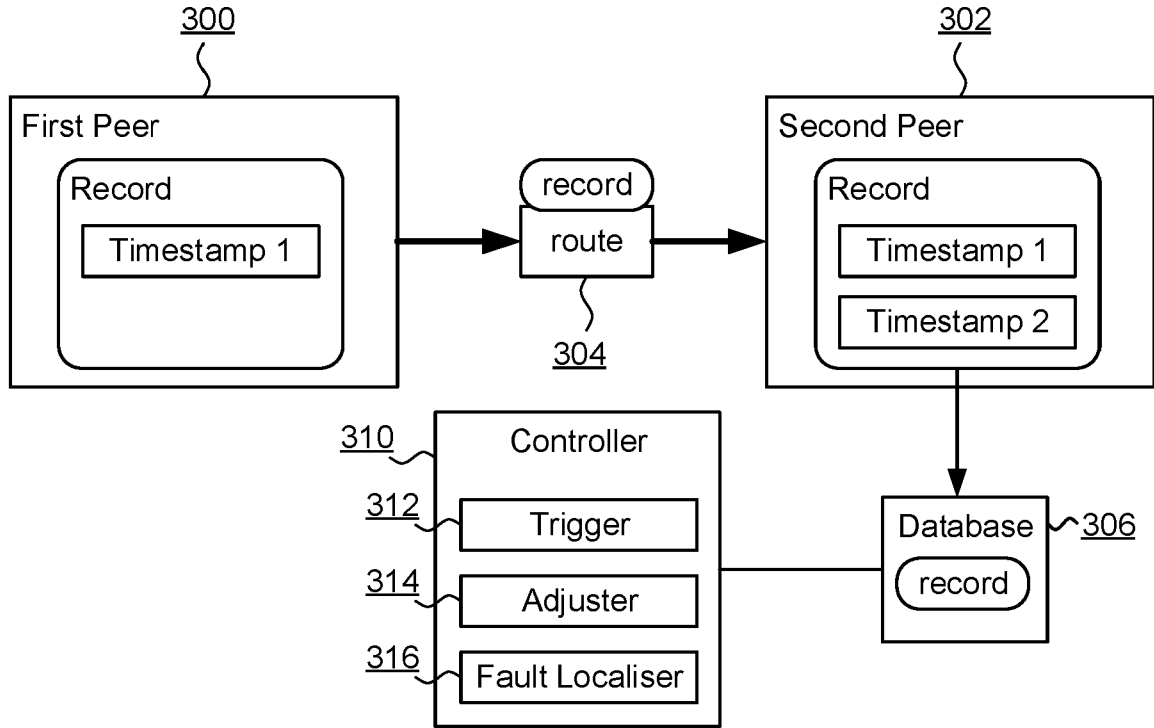


FIGURE 4

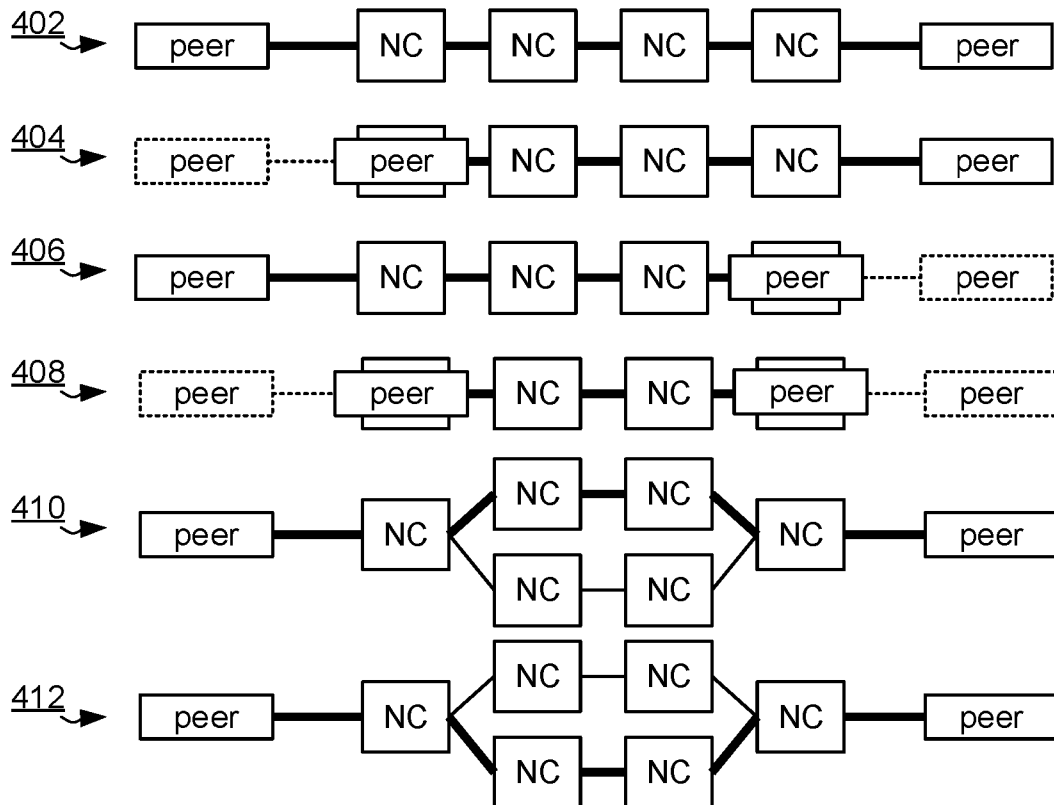
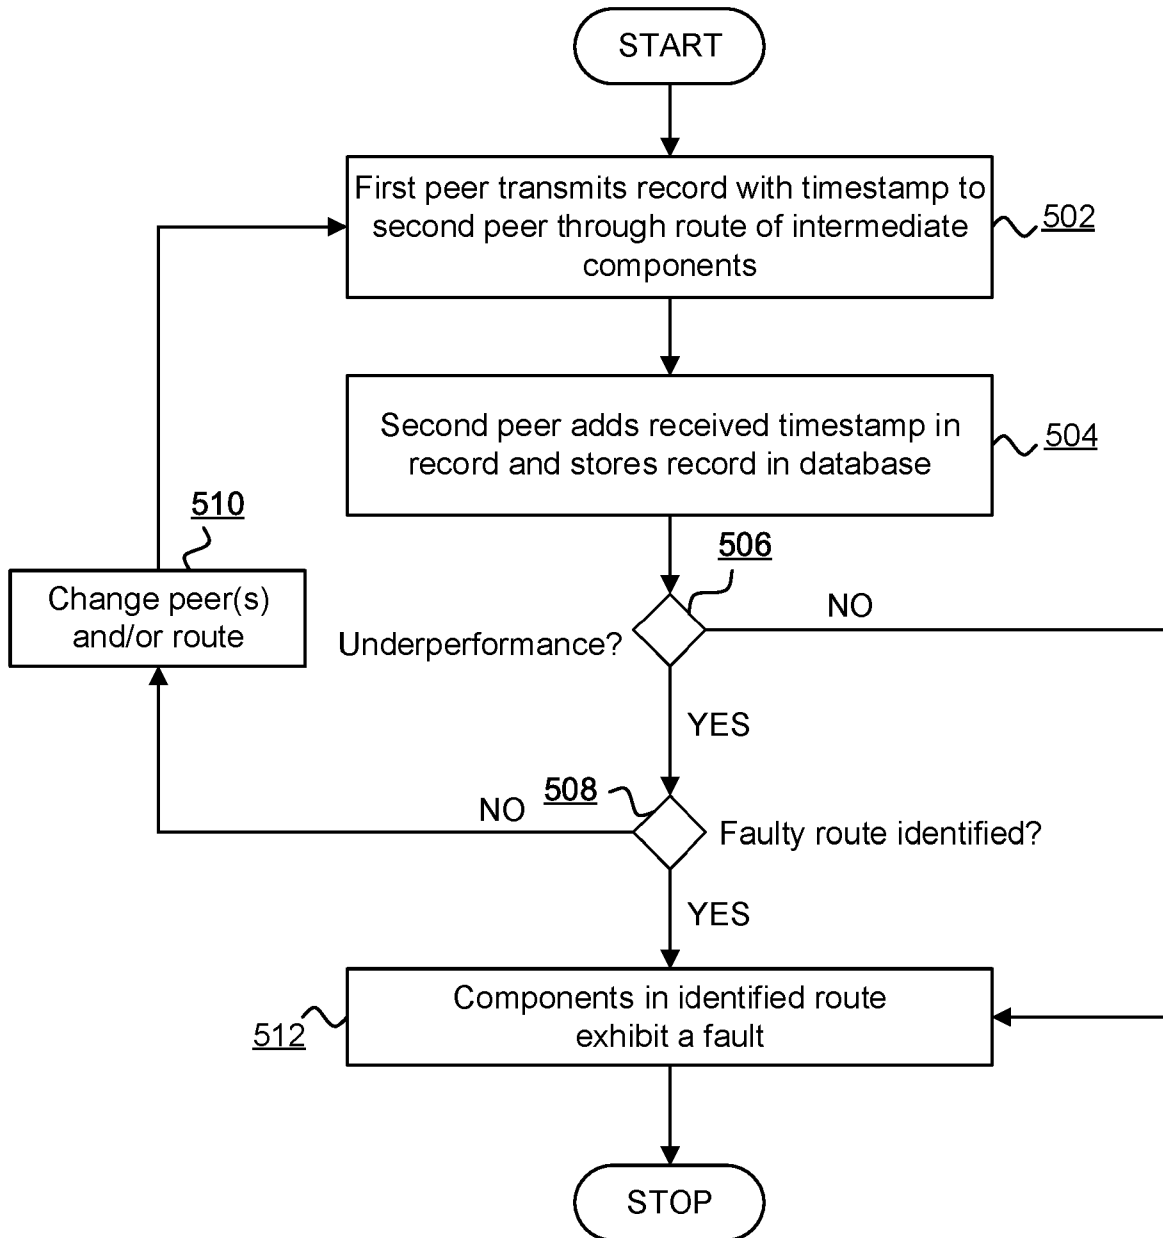


FIGURE 5

PERFORMANCE-BASED NETWORK FAULT LOCALISATION

The present invention relates to the identification of faults in a computer network.

5 Computer networks are susceptible to faults affecting a performance of communication via a network such as a reduced throughput, reduced data rate, an increased latency or frequency of drop-out or lost communications. Networks are increasingly spread across a number of intermediate network components including routers, switches, backhaul and backbone components and localising a fault to a portion of a network can be challenging.

10 Accordingly, it is beneficial to provide improvements to fault localisation in computer networks.

According to a first aspect of the present invention, there is provided a computer implemented method of identification of a portion of a computer network exhibiting a fault, the network connecting a plurality of communicatively connected peer computing devices
15 connected via a plurality of intermediate network components, the method comprising: a first peer transmitting a data record to a second peer through a route of intermediate network components, the record including a timestamp corresponding to a time of transmission of the record to the second peer, and the second peer being operable to supplement the record with a timestamp corresponding to a time of receipt of the record by the second peer,
20 wherein the record is stored in a database; responsive to a determination based on the timestamps in the record of underperformance of the route between the first and second peers, repeatedly performing the transmitting step changing one or more of: the first peer; the second peer; and the route between the first and second peers, so as to identify a route of intermediate components associated with the underperformance; and identifying the
25 network components in the identified route as exhibiting a fault.

Preferably, the database is a distributed transactional database such as a blockchain.

Preferably, the changing step includes changing at least one of the first and second peers to a network component in the route of intermediate network components for the transmission such that the network component becomes the peer.

30 Preferably, the first peer is changed to a first network component in the route from the first peer to the second peer.

Preferably, the second peer is changed to a last network component in the route from the first to the second peer.

Preferably, the route between the first and second peers is an initial route, and each subsequent repeating the transmitting step includes changing the route between the first and second peers such that the changed route includes at least a portion of the initial route.

According to a second aspect of the present invention, there is provided a computer system including a processor and memory storing computer program code for performing the steps of the method set out above.

According to a third aspect of the present invention, there is provided a computer program element comprising computer program code to, when loaded into a computer system and executed thereon, cause the computer to perform the steps of a method as described above.

Embodiments of the present invention will now be described, by way of example only, with reference to the accompanying drawings, in which:

Figure 1 is a block diagram a computer system suitable for the operation of embodiments of the present invention;

Figure 2 is a component diagram depicting communicatively connected peer computing devices connected via intermediate network components in a computer network according to embodiments of the present invention;

Figure 3 is a component diagram of an arrangement for identifying a portion of a computer network exhibiting a fault in accordance with embodiments of the present invention;

Figure 4 is a schematic diagram illustrating exemplary changes made for repeated transmissions in the arrangement of Figure 3 in accordance with embodiments of the present invention; and

Figure 5 is a flowchart of a method of identification of a portion of a computer network exhibiting a fault in accordance with embodiments of the present invention.

Figure 1 is a block diagram of a computer system suitable for the operation of embodiments of the present invention. A central processor unit (CPU) 102 is communicatively connected to a storage 104 and an input/output (I/O) interface 106 via a data bus 108. The storage 104 can be any read/write storage device such as a random-access memory (RAM) or a non-volatile storage device. An example of a non-volatile storage device includes a disk or tape storage device. The I/O interface 106 is an interface to devices for the input or output of data, or for both input and output of data. Examples of I/O devices

connectable to I/O interface 106 include a keyboard, a mouse, a display (such as a monitor) and a network connection.

Figure 2 is a component diagram depicting communicatively connected peer computing devices connected via intermediate network components in a computer network according to 5 embodiments of the present invention. Peer computing devices 200-210 are network attached devices such as computer systems, terminals, network equipment, customer premises equipment (CPE), routers, home-hubs, internet access points, network components, network appliances or other devices as will be understood by those skilled in the art. The peer devices are communicatively connected via a computer network constituted 10 by an exemplary arrangement of network components ("NC") 206 topologically situated therebetween. The network components 206 can be arranged according to any of a number of network topologies or organisations including hybrids, combinations or parts thereof. For example, the intermediate network components 206 can be constituted by customer premises equipment, telecommunications cabinet equipment, backhaul network equipment, 15 backbone appliances, or other network components as will be apparent to those skilled in the art. It will be appreciated by those skilled in the art that the term "peer" is used herein to refer to an endpoint of a communication according to embodiments of the present invention and is not restricted to network-attached equipment such as terminals and the like, and can include network components in some embodiments, in which case other network components 20 communicatively connecting such peers constitute the intermediate network components.

Embodiments of the present invention employ communication between peers such as those depicted in Figure 2 to determine, based on performance of network communication, portions of a computer network exhibiting a fault. Figure 3 is a component diagram of an arrangement for identifying a portion of a computer network exhibiting a fault in accordance 25 with embodiments of the present invention. A controller 310 is provided as a hardware, software, firmware or combination component to control and coordinate methods according to embodiments of the present invention. The controller includes a trigger 312 component for triggering a transmission of a data record from a first peer 300 to a second peer 302. The data record is a data item suitable for storage in a database 306 such as a distributed 30 transactional database such as a blockchain database. In some embodiments where a blockchain database is used, the record can be constituted as a transaction for storage and committing to the blockchain database as will be apparent to those skilled in the art.

The first peer 300 generates or receives a record for transmission to the second peer 302 via intermediate network components constituted as a route 304 through the network. The 35 first peer 300 stores a timestamp ("Timestamp 1") in the record corresponding to a time of

transmission of the record to the second peer. The record may be transmitted along with some other payload such as data having a predetermined or measurable size or quantity on which bases a performance of the network may be measured, such as a throughput, data transfer rate or the like. The second peer 302 receives the record via the route 304 through 5 the network and records a second timestamp ("Timestamp 2") in the record corresponding to a time of receipt of the record by the second peer. Subsequently, the record is stored in the database 306 accessible to the controller 310.

The controller 310 is thus operable, based on the record stored in the database 306 and, in particular, based on the timestamps stored in the record, to identify underperformance of 10 the network on the route 304. Underperformance can be determined with reference to predetermined threshold criteria of performance such as minimum, threshold, range or expected performance metrics such as a time or duration for such a transmission from the first peer 300 to the second peer 302. The predetermined threshold performance can be defined based on, for example, a size of a payload and/or the record transmitted between the 15 first peer 300 and the second peer 302.

Where the controller 310 determines underperformance of the connection between the first peer 300 and second peer 302, the controller 310 is operable to repeat transmissions potentially numerous. For each repeated transmission an adjuster 314 component makes a change to the transmission by changing one or more of: the first peer 300; the second peer 20 302; and the route 304 therebetween. For example, the first peer 300 can be changed to an alternative peer communicatively connected to the second peer 302 via the network. Further, the second peer 302 can be changed to an alternative peer. Further, the route 304 can be adjusted to a different route such as by changing one or more intermediate network components through which the transmission from the first peer 300 to the second peer 302 25 takes place. In a preferred embodiment, a first route 304 between the first 300 and second 302 peers is an initial route, and changed routes for subsequent repeated transmissions include at least a portion of the initial route such as a one or more network components of the initial route.

Subsequent repeated transmissions with the changes described above made by the 30 adjuster 314 are used by a further component, the fault localiser 316, to identify a route 304 of intermediate components associated with the underperformance. Thus, changes to the peers 300, 302 and/or route 304 are made to localise a portion of the network causing the underperformance and, therefore, being indicated as having a fault. Figure 4 is a schematic diagram illustrating exemplary changes made for repeated transmissions in the arrangement 35 of Figure 3 in accordance with embodiments of the present invention. At 402 an initial route

between a first and second peers is depicted. At 404 a first exemplary changed arrangement is depicted in which the first peer is changed to a first network component in the initial route, so as to reduce the portion of the network tested to the first network component (new first peer) and the second peer. At 406 a second exemplary changed arrangement is depicted in which the second peer is changed to a last network component in the initial route, so as to reduce the portion of the network tested to the first peer and the second-to-last network component (new second peer). At 408 the cumulation of the changes at 404 and 406 is employed in which the first peer is changed to the first network component in the initial route, and the second peer is changed to the last network component in the route. The arrangement of 410 depicts a new arrangement in which the first peer transmits to the second peer via an initial route depicted in heavy lines. The arrangement of 412 is changed vis-à-vis the arrangement of 410 in which the route between the first and second peers is changed to a different route.

While some exemplary changed arrangements are depicted with a single change it will be appreciated by those skilled in the art that multiple cumulative changes of peer(s) and/or route may be employed during repeated transmissions between peers.

In this way, the fault localiser 316 makes changes to peer(s) and/or the route responsive to detections of underperformance. Where underperformance is not detected, a change can be undone (reversed) such that a new change to a most recent arrangement in which underperformance was detected can be made. Preferably, changes involve reducing a number of intermediate network components such that peers between which transmissions are made become logically closer (in the sense that fewer intermediate network components are traversed between them) in order to focus on those network components implicated in an underperformance and thus indicative of fault.

While the controller is depicted in Figure 3 as separate and distinct to any particular network component or peer, it will be appreciated by those skilled in the art that the controller 310 can alternatively be provided in whole or in part by any of one or more or each of the network components and/or peers.

Figure 5 is a flowchart of a method of identification of a portion of a computer network exhibiting a fault in accordance with embodiments of the present invention. Initially, at step 502, a first peer 300 transmits a data record including a timestamp to a second peer 302 through a route 304 of intermediate network components. At step 504 the second peer stores a further timestamp in the record and stores the record in the database 306. At step 506 the method determines if a network underperformance is detected and, where underperformance is detected, the method proceeds to step 508 where a the method determines if a sufficiently

specific faulty route is identified. A faulty route can be determined to be identified if, for example, a sufficiently small number of intermediate network components has been attributed to the underperformance or if the route constitutes a known localised portion of the network suitable for remediation of faults. Where a faulty route is identified, the method
5 concludes at step 512 by which the network components in the faulty route are identified as exhibiting a fault. Alternatively, where a sufficiently specific faulty route is not yet identified, the method proceeds to step 510 where one or more of the peer(s) and route is changed before the transmission is repeated from step 502.

Insofar as embodiments of the invention described are implementable, at least in part,
10 using a software-controlled programmable processing device, such as a microprocessor, digital signal processor or other processing device, data processing apparatus or system, it will be appreciated that a computer program for configuring a programmable device, apparatus or system to implement the foregoing described methods is envisaged as an aspect of the present invention. The computer program may be embodied as source code or
15 undergo compilation for implementation on a processing device, apparatus or system or may be embodied as object code, for example.

Suitably, the computer program is stored on a carrier medium in machine or device readable form, for example in solid-state memory, magnetic memory such as disk or tape, optically or magneto-optically readable memory such as compact disk or digital versatile disk
20 etc., and the processing device utilises the program or a part thereof to configure it for operation. The computer program may be supplied from a remote source embodied in a communications medium such as an electronic signal, radio frequency carrier wave or optical carrier wave. Such carrier media are also envisaged as aspects of the present invention.

It will be understood by those skilled in the art that, although the present invention has
25 been described in relation to the above described example embodiments, the invention is not limited thereto and that there are many possible variations and modifications which fall within the scope of the invention.

The scope of the present invention includes any novel features or combination of features disclosed herein. The applicant hereby gives notice that new claims may be formulated to
30 such features or combination of features during prosecution of this application or of any such further applications derived therefrom. In particular, with reference to the appended claims, features from dependent claims may be combined with those of the independent claims and features from respective independent claims may be combined in any appropriate manner and not merely in the specific combinations enumerated in the claims.

CLAIMS

1. A computer implemented method of identification of a portion of a computer network exhibiting a fault, the network connecting a plurality of communicatively connected peer computing devices connected via a plurality of intermediate network components, the
5 method comprising:
 - a first peer transmitting a data record to a second peer through a route of intermediate network components, the record including a timestamp corresponding to a time of transmission of the record to the second peer, and the second peer being operable to supplement the record with a timestamp corresponding to a time of receipt of the record by
10 the second peer, wherein the record is stored in a database;
 - responsive to a determination based on the timestamps in the record of underperformance of the route between the first and second peers, repeatedly performing the transmitting step changing one or more of: the first peer; the second peer; and the route between the first and second peers, so as to identify a route of intermediate components
15 associated with the underperformance; and
 - identifying the network components in the identified route as exhibiting a fault.
2. The method of claim 1 wherein the database is a distributed transactional database such as a blockchain.
20
3. The method of any preceding claim wherein the changing step includes changing at least one of the first and second peers to a network component in the route of intermediate network components for the transmission such that the network component becomes the peer.
25
4. The method of claim 3 wherein the first peer is changed to a first network component in the route from the first peer to the second peer.
5. The method of claim 3 wherein the second peer is changed to a last network
30 component in the route from the first to the second peer.
6. The method of any of claims 1 to 2 wherein the route between the first and second peers is an initial route, and each subsequent repeating the transmitting step includes changing the route between the first and second peers such that the changed route includes
35 at least a portion of the initial route.

7. A computer system including a processor and memory storing computer program code for performing the steps of the method of any preceding claim.

8. A computer program element comprising computer program code to, when loaded into a computer system and executed thereon, cause the computer to perform the steps of a method as claimed in any of claims 1 to 6.



Application No: GB2008755.7

Examiner: Mr Luke Frodsham

Claims searched: 1-8

Date of search: 21 October 2020

Patents Act 1977: Search Report under Section 17

Documents considered to be relevant:

Category	Relevant to claims	Identity of document and passage or figure of particular relevance
X	1-8	US2014/219078 A1 (DASGUPTA ET AL.), See figure 4, paragraphs 15, 41-51
X	1-8	US5563875 A (HEFEL ET AL.), See whole document.

Categories:

X	Document indicating lack of novelty or inventive step	A	Document indicating technological background and/or state of the art.
Y	Document indicating lack of inventive step if combined with one or more other documents of same category.	P	Document published on or after the declared priority date but before the filing date of this invention.
&	Member of the same patent family	E	Patent document published on or after, but with priority date earlier than, the filing date of this application.

Field of Search:

Search of GB, EP, WO & US patent documents classified in the following areas of the UKC^X :

--

Worldwide search of patent documents classified in the following areas of the IPC

H04L

The following online and other databases have been used in the preparation of this search report

WPI, EPODOC

International Classification:

Subclass	Subgroup	Valid From
H04L	0012/26	01/01/2006
H04L	0012/24	01/01/2006