

(19) World Intellectual Property Organization
International Bureau



(43) International Publication Date
31 March 2011 (31.03.2011)

PCT

(10) International Publication Number
WO 2011/036663 A2

- (51) International Patent Classification: Not classified
- (21) International Application Number: PCT/IL2010/000788
- (22) International Filing Date: 21 September 2010 (21.09.2010)
- (25) Filing Language: English
- (26) Publication Language: English
- (30) Priority Data: 61/245,031 23 September 2009 (23.09.2009) US
- (71) Applicant (for all designated States except US): **COR-RELIX LTD.** [IL/IL]; 6 Galgalei Haplada Street, POB 12607, 46733 Herzeliya Pituach (IL).
- (72) Inventors; and
- (75) Inventors/Applicants (for US only): **SHAQED, Ariel** [IL/IL]; 75 Einstein Street, Ramat-Aviv (IL). **LEHAVI, David** [IL/IL]; 7 Heine Street, Haifa 34485 (IL).
- (74) Agent: **REINHOLD COHN AND PARTNERS;** P.O.B.13239, 61131 Tel Aviv (IL).
- (81) Designated States (unless otherwise indicated, for every kind of national protection available): AE, AG, AL, AM, AO, AT, AU, AZ, BA, BB, BG, BH, BR, BW, BY, BZ,

CA, CH, CL, CN, CO, CR, CU, CZ, DE, DK, DM, DO, DZ, EC, EE, EG, ES, FI, GB, GD, GE, GH, GM, GT, HN, HR, HU, ID, IL, IN, IS, JP, KE, KG, KM, KN, KP, KR, KZ, LA, LC, LK, LR, LS, LT, LU, LY, MA, MD, ME, MG, MK, MN, MW, MX, MY, MZ, NA, NG, NI, NO, NZ, OM, PE, PG, PH, PL, PT, RO, RS, RU, SC, SD, SE, SG, SK, SL, SM, ST, SV, SY, TH, TJ, TM, TN, TR, TT, TZ, UA, UG, US, UZ, VC, VN, ZA, ZM, ZW.

(84) Designated States (unless otherwise indicated, for every kind of regional protection available): ARIPO (BW, GH, GM, KE, LR, LS, MW, MZ, NA, SD, SL, SZ, TZ, UG, ZM, ZW), Eurasian (AM, AZ, BY, KG, KZ, MD, RU, TJ, TM), European (AL, AT, BE, BG, CH, CY, CZ, DE, DK, EE, ES, FI, FR, GB, GR, HR, HU, IE, IS, IT, LT, LU, LV, MC, MK, MT, NL, NO, PL, PT, RO, SE, SI, SK, SM, TR), OAPI (BF, BJ, CF, CG, CI, CM, GA, GN, GQ, GW, ML, MR, NE, SN, TD, TG).

Declarations under Rule 4.17:

— of inventorship (Rule 4.17(iv))

Published:

— without international search report and to be republished upon receipt of that report (Rule 48.2(g))

(54) Title: METHOD AND SYSTEM FOR RECONSTRUCTING TRANSACTIONS IN A COMMUNICATION NETWORK

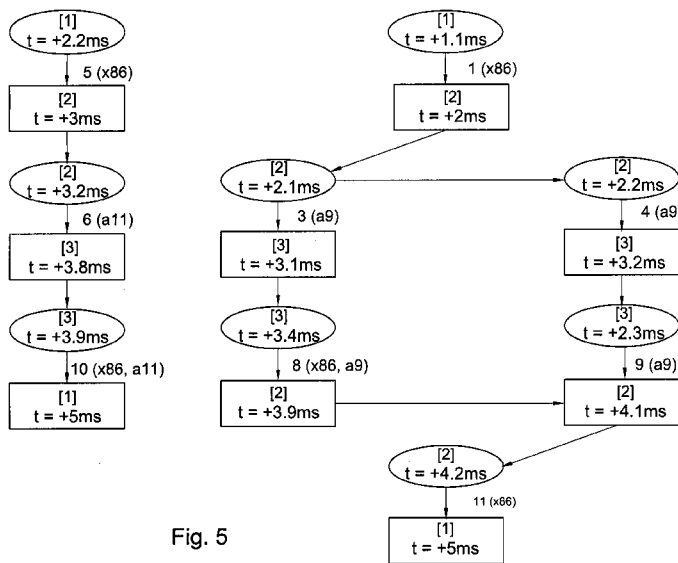


Fig. 5

(57) Abstract: In a distributed information network, a method and system for assembling messages in an input set of messages into one or more message assemblies. The method involves constructing an initial directed graph of the set of messages, wherein the directed graph consists of pairs of a first vertex and a second vertex, wherein, for each message in the set of messages, a first vertex is an the out node of the message and the second vertex is the in node of the message. A connectivity matrix is defined on the initial directed graph that is modified in an iterative process. A final directed graph is produced that is partitioned into disjoint subgraphs, each subgraph being an assembly of messages from the input message set.

WO 2011/036663 A2

METHOD AND SYSTEM FOR RECONSTRUCTING TRANSACTIONS IN A COMMUNICATION NETWORK

5

FIELD OF THE INVENTION

This invention relates to methods and systems for analyzing a communication
10 network.

BACKGROUND OF THE INVENTION

The typical enterprise technology landscape today is characterized by the
distributed nature of its information systems and the heterogeneity of its technologies.
15 Following decades of evolution and innovation, it is common to see "*best-of-breed*"
packaged applications, such as ERP, CRM, SCM, SRM, PLM, etc., home-grown
systems, and legacy applications, each with numerous interfaces interconnecting them.
In these distributed and complex environments, the act of processing a typical
transaction spans numerous applications and technological boundaries, often rendering
20 the enterprise incapable of understanding execution paths for the transactions as well as
their logical and technical interdependencies. Without the requisite knowledge of how
the transactions are executed, the enterprise is severely limited in its ability to monitor
transactions and to detect and remedy bottlenecks, latencies, and points-of-failure.

Enterprise organizations therefore seek to increase the visibility of their
25 automated business transactions as they see a direct correlation between transaction
visibility and business performance. Achieving a high-degree of visibility enables the
enterprise to improve customer service, to monitor transaction performance and health,
to optimize the business logic, and to implement efficient solutions to problems as they
arise.

US Patent No. 6,738,813 to Reichman discloses a system for monitoring Web sites, or other server systems, as seen from the computing devices of other users. In a preferred embodiment, the system includes an agent component that runs on the computing devices of service users to provide functionality for accessing and monitoring the performance of a server. The agents are remotely configurable over the Internet, and may be configured, for example, to execute a particular Web transaction while monitoring specified performance parameters (server response times, network hop delays, server availability, etc). Using a service provider Web site, a user of the service can set up a monitoring session in which agent devices of other community members are used to monitor the performance of the user's server system.

US Patent No. 7,424,530 to Chagoly discloses a method for graph manipulation of transactional performance data to identify and emphasize root causes of electronic business system transaction processing performance problems. A system transaction monitoring system, is utilized to obtain transaction performance data for a system. This transaction performance data is utilized to present a graph of a given transaction or transactions.

US Patent Publication 20100228650 discloses a method of tracking information in a multi-tier computerized environment involving detecting data associated with at least one request or transaction related to a byte stream; and matching a thread associated with the byte stream with the at least one request or transaction associated with the at least one packet according to predetermined fields within the byte stream.

US Patent Publication 20060015512 to Alon et al. discloses an apparatus for monitoring a selected tier in a multi-tier computing environment. Monitored request traffic includes at least one entering request received at a request entry port from an adjacent tier, identifying each request in the monitored request traffic and sending at least a request identifier to the context agent. The context agent also receives information relating to the request context of the entering request from the context agent associated with the adjacent tier and the context agent associates the information relating to the request context of the entering request with the entering request, in accordance with a request identifier.

SUMMARY OF THE INVENTION

Glossary

There follows a glossary of terms used in the following description and set of claims together with their definitions, some of them known in the art, others having
5 been coined.

Computational Node

A computational node is a node of a communication network that receives and/or generates information. For the sake of convenience, a computational node is sometimes referred to herein simply a "*node*". Examples of nodes are a web server,
10 database server, a queue, a virtual server in a virtual environment, and a process within a server. A cluster, a redundancy system, and multiple instances of the same server may be considered as being a single node.

A Message

A message is the unit of information sent on the application layer. Three types of
15 information related to a message are defined: network information (TCP/IP), instance information and application (Application layer) information. Each message includes one or more message ID's and typically also includes any one or more of an ID of the origin node, the ID of the destination node, origin node time stamp, destination node time stamp, message type and match ID.

20 A match ID

A value of a field or combination of fields within a message that identifies the logical flow of messages to which the message belongs. For example a flow of messages that constitutes an Order Single transaction instance may include a request message of the type "Single Order" and a response message these types belong to the
25 same transaction is the fact that both messages hold the same value in the field "ClientOrderID". The value of the field "ClientOrderID" would be the match ID in this case

The present invention provides a system and method for monitoring transactions within a distributed information network. The system of the invention comprises one or
30 more collectors that record data relating to incoming and outgoing messages at the nodes of the communication network. The recorded data for each message typically

- 4 -

includes an “in node ID” and an “out node ID”, the message type, the time stamp of the message at the in node and the out node, and the match ID or match IDs of the message.

At various times the collector reconstructs recently completed transactions in the system from the recorded data. In accordance with this aspect of the invention, a
5 directed graph is defined, wherein, for each message in an input set of messages, a first vertex and a second vertex are defined, where the first vertex is the out node of the message and the second vertex is the in node of the message. All of the vertices are ordered according to their time stamp. A directed edge is introduced into the directed graph from the out vertex to the in vertex of each message, and a directed connectivity
10 matrix A is calculated in which $a_{kk'}=1$ if a directed path exists from the vertex k' to the vertex k . Otherwise $a_{kk'}=0$. The matrix A is then modified in an iterative process described in detail below. The result of the iterative process is a partition of the input set of messages into disjoint directed graphs, where each graph contains a vertex tagged as a beginning of a message path, and each of the graphs represents a transaction in the
15 system.

The system of the invention also includes a management server. For one or more of the transactions identified by the collectors, the management server calculates one or more values of one or more parameters of the transaction. The parameters of the path type may be, for example, total transit time, or the processing time at each of one or
20 more nodes. When the determined value of one or more of the parameters exceeds a predetermined threshold, an alert may be issued that may be displayed on the monitor. In addition, the processing time of messages at each of one or more nodes may be determined and a node history may be constructed. Statistics of each of the parameters of may be compiled, and when the value of any statistic exceeds a predetermined
25 threshold, an alert can be issued.

The collectors may first use a binning method to partition an initial set of messages set of messages where all of the messages in a particular transaction are placed in the same bin. Each bin is then used as an input set of messages for the collectors.

30 Thus, in its first aspect, the invention provides a system for assembling in an input set of messages in a distributed information network into one or more message

assemblies, each message having an out node in the network and an in node in the network, the system comprising, : :

(a) one or more message collection servers; each collection server being configured, to:

5 a) construct a directed graph of the set of messages, wherein the directed graph consists of pairs of a first vertex and a second vertex , wherein , for each message in the set of messages, a first vertex is an the out node of the message and the second vertex is the in node of the message;

order the vertices according to a time stamp of each vertex;

10 construct a connectivity matrix A of the directed graph in which $a_{kk}=1$ if a directed path exists from the vertex k' to the vertex k and $a_{kk}=0$ otherwise.

modify the matrix A in an iterative process in a process comprising:

For $k=1$ to k_{max} , k_{max} being a number of vertices in the graph:

15 if there is a vertex $k' < k$ on the node of k such that a correlation from the vertex k' to the vertex k allowed, and if the current value of $a_{kk} \neq 1$, introducing a directed edge into the directed graph from the vertex k' to the vertex k ;

setting $a_{kk}=1$;

20 if k is an out vertex and the corresponding in vertex is k'' , setting $A_{k''}=A_k$ or A_k'' , where A_k is the k th row of the matrix A and $A_{k''}$ is the k'' th row of the matrix A ;

partitioning the directed graph into disjoint subgraph, each subgraph being an assembly of messages from the input message set.

25 The system according to Claim 1 wherein the processor is further configured to partition the directed graph into disjoint subgraphs where each subgraph contains a vertex tagged as a beginning of a message path.

30 The processor may further be configured to execute a binning algorithm on an initial set of messages, the binning algorithm partitioning the initial set of messages into disjoint bins of messages, and one or more of the bins may be used as the input set of messages. In this case, each message in the initial set of messages is assigned one or more match IDs and the binning algorithm comprises:

(i) determining an equivalence relation among the match IDs, a first match ID being equivalent to a second match ID if there exists a sequence of

- 6 -

messages M_1, \dots, M_s in the initial set of messages such that M_k and M_{k+1} have a common match ID, and

- (ii) partitioning the messages into disjoint bins of messages, any two messages being in the same bin if there exists a sequence of messages M_1, \dots, M_s such that M_1 is the first message, M_s is the second message, and M_k and M_{k+1} have a common match ID.

The system may further comprise a management server configured to calculate one or more values of one or more parameters of the transactions. One or more of the parameters may be selected, for example, from a total transit time, and a processing time at each of one or more nodes. The management server may be configured to issue an alert when a determined value of one or more of the parameters exceeds a predetermined threshold. The management server may be further configured to determine a node history and to calculate statistics of one or more of the parameters of one or more nodes. The management server may be configured to issue an alert when the value of any statistic exceeds a predetermined threshold.

In another of its aspects, the invention provides a method for assembling messages in an input set of messages in a distributed information network into one or more message assemblies, each message having an out node in the network and an in node in the network, the method comprising,:

- a) constructing a directed graph of the set of messages, wherein the directed graph consists of pairs of a first vertex and a second vertex, wherein, for each message in the set of messages, a first vertex is an the out node of the message and the second vertex is the in node of the message;
- ordering the vertices according to a time stamp of each vertex;
- constructing a connectivity matrix A of the directed graph in which $a_{kk'}=1$ if a directed path exists from the vertex k' to the vertex k and $a_{kk'}=0$ otherwise.
- modifying the matrix A in an iterative process in a process comprising:
- for $k=1$ to k_{max} , k_{max} being a number of vertices in the graph:
- if there is a vertex $k' < k$ on the node of k such that a correlation from the vertex k' to the vertex k allowed, and if the current value of $a_{kk'} \neq 1$, introducing a directed edge into the directed graph from the vertex k' to the vertex k ;
- setting $a_{kk'}=1$;

- 7 -

if k is an out vertex and the corresponding in vertex is k' , setting $A_{k'} = A_k$ or $A_k = A_{k'}$, where A_k is the k th row of the matrix A and $A_{k'}$ is the k' th row of the matrix A ; and

partitioning the directed graph into disjoint subgraph, each subgraph being an assembly of messages from the input message set.

The method according to Claim 12 wherein further comprising partitioning the directed graph into disjoint subgraphs where each subgraph contains a vertex tagged as a beginning of a message path.

The method of the invention may further comprise executing a binning algorithm on an initial set of messages, the binning algorithm partitioning the initial set of messages into disjoint bins of messages, and one or more of the bins may be used as the input set of messages. In this case, each message in the initial set of messages is assigned one or more match IDs and the binning algorithm comprises:

- (i) determining an equivalence relation among the match IDs, a first match ID being equivalent to a second match ID if there exists a sequence of messages M_1, \dots, M_s in the initial set of messages such that M_k and M_{k+1} have a common match ID, and
- (ii) partitioning the messages into disjoint bins of messages, any two messages being in the same bin if there exists a sequence of messages M_1, \dots, M_s such that M_1 is the first message, M_s is the second message, and M_k and M_{k+1} have a common match ID.

The method of the invention may further comprise calculating one or more values of one or more parameters of the transactions. One or more of the parameters may be selected from a total transit time, and a processing time at each of one or more nodes. An alert may be issued when a determined value of one or more of the parameters exceeds a predetermined threshold.

The method may further comprise determining a node history calculating statistics of one or more of the parameters of one or more nodes. An alert may be issued when the value of any statistic exceeds a predetermined threshold.

30

BRIEF DESCRIPTION OF THE DRAWINGS

In order to understand the invention and to see how it may be carried out in practice, embodiments will now be described, by way of non-limiting example only, with reference to the accompanying drawings, in which:

5 **Fig. 1** shows a system for assembling messages in a communication network into message assemblies, in accordance with one aspect of the invention;

Fig. 2 shows a first representation of a flow of messages between nodes in a communication network;

Fig. 3 shows a second representation of a flow of messages between nodes in a
10 communication network;

Fig. 4 shows a flow chart of a method for assembling messages into message assemblies in accordance with one aspect of the invention; and

Fig. 5 shows two message assemblies constructed from an input set of messages by the method of the invention.

15 DETAILED DESCRIPTION OF EMBODIMENTS

Fig. 2 shows a schematic diagram **10** of message flow through a communication network. A message **m1** originating at node **1** is assigned the match ID "1234" and is sent to node **2**. Receipt of the message **m1** at the node **2** generates a message **m2** at the
20 node **2** that is sent to node **3** after having been assigned two match IDs. One of the match IDs is the match ID of message **1** (1234). The other match ID is **AXY**. Receipt of the message **m2** at the node **3** generates a message **m3** at the node **3** that is sent to the node **2** with the match ID **AXY**. Receipt of the message **m3** at the node **2** generates a message **m4** that is sent to the node **1** with the ID **AXY**.

25 **Fig. 3** shows another schematic diagram **20** of the message flow **10** depicted in **Fig. 2**. In the diagram **20**, a circle, such as the circle **20** indicates an origin node of a message, and a square, such as the square **22** indicates the destination node of the message. The number of the node is indicated inside the circle or square. Transmission of a message from one node to another is indicated by an arrow from the origin node to
30 the destination node of the message. The message appears next to the arrow. An arrow

from a destination node to an origin node, such as the arrow **24** indicates generation of a message within a node.

Fig. 1 shows a system, generally indicated by **100**, for monitoring transactions within a distributed information network in accordance with one embodiment of this aspect of the invention. The network **100** comprises one or more nodes **106**, one or more distribution switches **104**, and one or more access switches **101**. The nodes **106**, distribution switches **104** and the access switches may be arranged in a hierarchy in which one or more clients **106** are connected to a common distribution switch **104** and one or more distribution switches are connected to common access switch **101**. The access switches **101** are connected to a core **105** of the network **100**. The network **100** further comprises one or more collectors **102** and a management server **103**. In the embodiment of Fig. 1 the collector **102** is connected to one or more access switches **101** through a port monitor. A management server **103** is connected to the network **100** at a distribution switch **104**.

The collector **102** records data relating to incoming and outgoing messages at the nodes **106**. The recorded data for each message includes a in node ID and out node ID, the message type, the time stamp of the message at the in node and the out node, and the match ID or match IDs of the message. The collector **102** at various times reconstructs recently completed transactions in the system **100** from the recorded data, as described below.

The collector **102** may first use a binning method to partition the set of messages whose data were recorded into disjoint subsets, where all of the messages in a particular transaction are placed in the same subset. Initially, all of the messages having a given match ID are assigned to the same bin. Messages having more than one match ID will be assigned to more than one bin. After assignment of the messages to bins, two or more of bins having one or more common messages are combined into a single bin. The process of combining bins continues until the set of messages is partitioned into disjoint bins. The binning may be performed, for example, using a standard fast or efficient "Union Find" implementation, for example as disclosed in "Disjoint-set data structure" on Wikipedia as last modified on August 28, 2010).

After binning, the collector **102** assembles the messages in each bin into transactions. Fig. 4 shows a method for assembling the messages in a given bin into one or more disjoint transactions in accordance with one embodiment of the invention.

- 10 -

The method begins with step **101** where a directed graph is defined, wherein, for each message in the bin, a first vertex and a second vertex are defined, where the first vertex is the out node of the message and the second vertex is the in node of the message. Then, in step **102**, all of the vertices are ordered according to their time stamp. In the case where there are two or more vertices having the same time stamp, the vertices having the same time stamp can be ordered arbitrarily subject to the restriction that a vertex corresponding to an out node of a message is positioned in the order before the vertex corresponding to the in node of the same message. Each vertex is thus assigned a vertex number k , where k is the position of the vertex in the ordering. Then in step **103**, a directed edge is introduced into the directed graph from the out vertex to the in vertex of each message. In step **104** a directed connectivity matrix A is calculated in which $a_{kk}=1$ if a directed path exists from the vertex k' to the vertex k . Otherwise $a_{kk}=0$.

The matrix A is now modified in an iterative process as follows. In step **106**, k is set to 1, and in step **108**, it is determined whether the vertex k is the earliest vertex on its node. If no, then in step **110**, it is determined whether there is a vertex k' that is earlier than the vertex k such that a correlation from the vertex k' to the vertex k is a predetermined allowed correlation. The allowed correlations reflect business considerations, so that transactions of interest are generated. If at step **110** it is determined that such an earlier vertex does not exist, or if in step **108** it is determined that k is the earliest vertex on its node, then the process proceeds to step **112** where it is determined whether the vertex k is an out vertex. If yes then in step **114** the vertex k is tagged as being a vertex at the beginning of a message path and the process continues to step **116**. If at step **112** it is determined that k is not an out vertex, the process also continues to step **116**. In step **116** it is determined whether the current value of k is the maximum value of k . If no, then in step **120** k is increased by 1 and the process returns to step **108**.

If in step **110** it is determined whether there are vertices that are earlier than the vertex k such that a correlation from the vertex k' to the vertex k is allowed, then in step **122** it is determined whether the current value of $a_{kk}=1$, where the vertex k' is the latest earlier vertex to the vertex k . If yes, then the process proceeds to step **116**. If no, then in step **124** a directed edge is introduced into the directed graph from the vertex k' to the vertex k .

After step 124, the directed connectivity matrix A is updated. First, in step 126, $a_{kk'}$ is set to 1. Then in step 128, A_k (row k of the matrix A) is set to the logical expression " A_k or $A_{k'}$ ", wherein $A_{k'}$ is a vector whose i th component is 0 if the i th component of A_k and $A_{k'}$ are both 0, and the i th component of A_k or $A_{k'}$ is 1 otherwise (i.e. if the i th component of either A_k or $A_{k'}$ is 1). Finally, in step 130, if k is an out vertex, and the corresponding in vertex is k' , then $A_{k'}$ is reset to the logical expression " A_k or $A_{k'}$ ". The process then returns to step 116.

If in step 116 it is determined that the value of k is maximal, then in step 118 the directed graph is partitioned into disjoint subgraphs where each subgraph contains a vertex tagged in step 114 as a beginning of a message path. Each of the subgraphs represents a transaction in the system. The process then terminates in step 132.

If most of the correlations in step 110 are allowed and the bin has many nodes, then the complexity of the process of Fig. 4 is linear with the number of vertices in the bin. Since the step 124 is executed at most once for each vertex k , the memory required for vertex is bounded above by a constant so that the memory for all of the vertices may be allocated as an array.

Transactions identified by the process of Fig. 4 are input to the management server 103. For one or more of the transactions, the management server calculates one or more values of one or more parameters of the transaction. The parameters of the path type may be, for example, total transit time, or the processing time at each of one or more nodes. The values of the parameters may be displayed on the monitor 109. When the determined value of one or more of the parameters exceeds a predetermined threshold, an alert may be issued that may be displayed on the monitor. In addition, the processing time of messages at each of one or more nodes may be determined and a node history may be constructed. Statistics of each of the parameters of may be compiled, and when the value of any statistic exceeds a predetermined threshold, an alert can be issued.

Example

The invention will now be demonstrated by way of an example. Table 1 shows data collected by a collector, such as the collector 102, relating to 13 messages.

- 12 -

collected by a collector, such as the collector 102. The messages were sent between 4 nodes.

As can be seen in Table 1, Match IDs are shared between node 1 and node 2, and between node 2 and node 3. However nodes 1 and 2 always send different match
5 IDs and nodes 3 and 4 respond with both match IDs. This might occur, for example, if the match ID from node 1 is being sent to node 3 by some out-of-band method, or by some unmonitored path, or is reconstructed by a previous phase of the monitoring. Tier
#1 Node 1 communicates with node 2 using a protocol with messages "Order" and "Cancel", receiving a message "Acknowledge" for each; Node 2 communicates with
10 nodes 3 and 4 using some other protocol with messages "NEW" and "CXL", producing a message "ACK" or "ACK CXL" for each. These relations define the allowed correlations between message types.

- 13 -

Table 1

ID	Message Type	From (out node)	Time Stamp	To (in node)	Time Stamp	Match ID 1	Match ID 2
1	Order	1	+1.1ms	2	+2ms	x86	
2	Order	1	+1.2ms	2	+1.9ms	x87	
3	NEW	2	+2.1ms	3	+3.1ms		a9
4	NEW	2	+2.2ms	4	+3.2ms		a9
5	Cancel	1	+2.2ms	2	+3ms	x86	
6	CXL	2	+3.2ms	3	+3.8ms		a11
7	NEW	2	+3.3ms	4	+4ms		a10
8	ACK	3	+3.4ms	2	+3.9ms	x86	a9
9	ACK	4	+3.4ms	2	+4.1ms		a9
10	ACK CXL	3	+3.9ms	1	+5ms	x86	a11
11	Acknowledge	2	+4.2ms	1	+5ms	x86	
12	ACK	4	+4.2ms	2	+5ms	x87	a10
13	Acknowledge	3	+5.1ms	1	+6ms	x87	

Binning

At each step, the message introduced into the system is shown together with the existing bins. Each bin is displayed as "**{MatchID₁, MatchID₂, . . . | id₁, id₂, . . .}**", which
 5 specifies that the bin holds all messages with the given match IDs, and gives the IDs of those messages.

Referring again to Table 1, the binning process proceeds as follows:

1. Message 1 is first processed. There is no bin for match ID "x86" (or, indeed, for any match ID), so a new bin is created. After this step, there is one bin, **{x86: 1}**.
 10
2. Message 2 is processed. There is no bin for match ID "x87", so a new bin is created. After this step, there are two bins **{x86: 1}**, **{x87: 2}**.
3. Message 3 is processed, creating a new bin for match ID "a9". After this step, the bins are **{x86: 1}**, **{x87: 2}**, **{a9: 3}**.
- 15 4. Message 4 is processed, and is added to the existing bin for match ID "a9". After this step, bins **{x86: 1}**, **{x87: 2}**, **{a9: 3,4}**.
5. Message 5 is processed, and is added to the existing bin for match ID "x86". After this step, the bins are **{x86: 1,5}**, **{x87: 2}**, **{a9: 3,4}**.
6. Message 6 is processed, creating a new bin for match ID "a11". After this step, the bins are **{x86: 1,5}**, **{x87: 2}**, **{a9: 3,4}**, **{a11: 6}**.
 20
7. Message 7 is processed, creating a new bin for match ID "a10". After this step, the bins are **{x86: 1,5}**, **{x87: 2}**, **{a9: 3,4}**, **{a11: 6}**, **{a10: 7}**.
8. Message 8 is processed. Since Message 8 has both Match ID x86 and Match ID a9, processing Message 8 unifies the bins for match IDs "x86" and "a9" and is added to the resulting bin. After this step, the bins are **{x86,a9: 1,3,4,5,8}**, **{x87: 2}**, **{a11: 6}**, **{a10: 7}**.
 25
9. Message 9 is processed, and is added to the existing bin for match IDs "x86" and "a9". After this step, the bins are **{x86,a9: 1,3,4,5,8,9}**, **{x87: 2}**, **{a11: 6}**, **{a10: 7}**.

- 15 -

10. Message 10 is processed. It causes the unification of the bins for match IDs "x86" and "a11", and is added to the resulting bin. After this step, the bins are {x86,a9,a11: 1,3,4,5,6,8,9,10}, {x87: 2}, {a10: 7}.
11. Message 11 is processed, and is added to the existing bin for match IDs "x86" and "a9". After this step, the bins are {x86,a9,a11: 1,3,4,5,6,8,9,10,11}, {x87: 2}, {a10: 7}.
12. Message 12 is processed. It causes the unification of the bins for match IDs "x87" and "a10", and is added to the resulting bin. After this step, the bins are {x86,a9,a11: 1,3,4,5,6,8,9,10,11}, {x87,a10: 2,7,12}.
13. Message 13 is processed, and is added to the existing bin for match IDs "x87" and "a10". After this step, the bins are {x86,a9,a11: 1,3,4,5,6,8,9,10,11}, {x87,a10: 2,7,12,13}.

Thus, 2 bins have been constructed containing of of the messages of distinct transactions. Each bin contains the messages of one or more transactions, and the messages of each transaction is contained entirely inside a single bin.

After enough time has elapsed, we can safely retire bins to transaction assembly. This requires a criterion for determining when all transactions of a bin must have ended. For example, if no transaction may take more than 10ms, at some time after +15ms bin #1 is sent to transaction assembly, and at some time after +16ms bin #2 is sent to transaction assembly. Any criterion can be used for retiring bins, including combinations of:

- Time with no new messages added to the bin (this is the method detailed above)
- Round-trip messages (e.g. a response received for every request)
- Attempted transaction assembly of the bin (as below) yields transactions matching some business logic
-

- **Assembly**

The method of the invention for assembly of messages into transactions will be demonstrated on the bin {x86,a9,a11: 1,3,4,5,6,8,9,10,11} produced by the above binning process.

5 In this phase the match IDs are ignored, and the time stamps are used to order events, and message types are used to check for correlations allowed by business logic. In this example, the business logic configuration allows correlations only between "Order", "NEW", "ACK", and "Acknowledge" and between "Cancel", "CXL" and "CXL ACK".

10 Referring to step 101 in Fig. 4, an out vertex and an in vertex are created for each message in the bin (there will thus be 9 OUT vertices, denoted by ovals, and 9 IN vertices, denoted by rectangles). Each vertex is associated with a message, which in turn associates it with the appropriate time and node address of that message: for an OUT vertex that is the sending time, for an IN vertex that is the receiving time.

15 Referring to step 102 in Fig. 4, these vertices are sorted in order of time. When two vertices have the same time, the method does not depend on the specific ordering chosen. Accordingly, the business logic can create a specific ordering to achieve a preferred display format, or an arbitrary ordering (e.g. that resulting from the sort function) can be used. After this step, there are 18 vertices, as shown in Table 2:

20 **Table 2**

Vertex k	Vertex ID	Type	Time Stamp	IP
1	OUT 1	Order	+1.1	1
2	IN 1	Order	+2	2
3	OUT 3	NEW	+2.1	2
4	OUT 4	NEW	+2.2	2
5	OUT 5	Cancel	+2.2	1
6	IN 5	Cancel	+3	2

7	IN 3	NEW	+3.1	3
8	IN 4	NEW	+3.2	4
9	OUT 6	CXL	+3.2	2
10	OUT 8	ACK	+3.4	3
11	OUT 9	ACK	+3.4	4
12	IN 6	CXL	+3.8	3
13	OUT 10	ACK CXL	+3.9	3
14	IN 8	ACK	+3.9	21
15	IN 9	ACK	+4.1	21
16	OUT 11	Acknowledge	+4.2	21
17	IN 11	Acknowledge	+5	11
18	IN 10	ACK CXL	+5	11

Referring now to step 103 in Fig. 4, a "network" edge is added directed from each OUT vertex to its corresponding IN vertex. For example, there is an edge from "OUT 6" to "IN 6". For every OUT vertex i , $A_i=0\dots0$ (no inbound edges), and for its matching IN vertex j $a_{j,i}=1$ (i.e., an edge exists from i to j), so $A_j=0\dots010\dots0$ where the "1" appears in position i .

The starting matrix A is accordingly:

- $A_1 = 000000000000000000$
- $A_2 = 100000000000000000$
- 10 $A_3 = 000000000000000000$
- $A_4 = 000000000000000000$
- $A_5 = 000000000000000000$
- $A_6 = 000010000000000000$
- $A_7 = 001000000000000000$
- 15 $A_8 = 000100000000000000$
- $A_9 = 000000000000000000$
- $A_{10} = 000000000000000000$
- $A_{11} = 000000000000000000$
- $A_{12} = 000000001000000000$
- 20 $A_{13} = 000000000000000000$
- $A_{14} = 000000000010000000$

- 18 -

A_15= 000000000100000000
 A_16= 000000000000000000
 A_17= 000000000000000100
 A_18= 000000000000010000

- 5 All values of k are now scanned, (steps 106 through 130 in Fig. 4).
1. Vertex OUT 1 is the first vertex on its node (node 1) (step 108). It is an OUT vertex [step 112], so it is tagged as the beginning of a new transaction.
 2. Vertex IN 1 is the earliest vertex on its node (node 2) [step 108]. It is not an OUT vertex, nothing further is done.
 - 10 3. Vertex OUT 3 is *not* the earliest on its node (node 2) [step108]. Correlations from Order to NEW are allowed [step 110], so $k'=2$. A directed "compute" edge is introduced from vertex 2 to vertex 3 [step 124], and $a_{32}=1$ [step126]. Now $A_3=0...011$ (because $A_2=0...1$) [step127]. (i.e. "vertex 3 is reachable from vertices 2,1 on node 2". OUT 3 is an OUT vertex, with corresponding vertex IN 3 at position 7, so $A_7 = 0...0111$ [step130].
 - 15 4. Vertex OUT 4 is not the earliest on its node (node) [step108]. Correlations from NEW to NEW are allowed [step 110], so $k'=3$, and $A_4 = 1110...0$. OUT 4 is an OUT vertex, with corresponding vertex IN 4 at position $k''=8$, so we set $A_8 = 11110...0$.
 - 20 5. Vertex OUT 5 is not the earliest on its node (node 1) [step108]. However a correlation from Order to Cancel is not allowed (by the business logic) [step110], and it is an OUT vertex [step112], so it is tagged as the beginning of a second transaction. Accordingly $A_5=0$.
 6. Vertex IN 5 is not the earliest on its node (node 2) [step 108], but no correlations from NEW or from Order to Cancel are allowed (by the business logic) [step 110]. It is not an OUT vertex [112], so nothing further is done.
 - 25 7. Vertex IN 3 is the earliest on its node 3 [108], but it is not an OUT vertex [step 112], so nothing further is done. Note that $A_7=1110...0$ from step 3 since vertex 7 is reachable from vertices 1,2, and 3.
 - 30 8. Vertex IN 4 is the earliest vertex on its node (node 4) [108], but it is not an OUT vertex [112], so nothing further is done, keeping $A_8=0...01111$ from step 4.
 9. Vertex OUT 6 is not the earliest vertex on its node (node 2) [108]. A correlation is allowed from vertex IN 5, so $k'=6$. A directed "compute" edge is introduced

- 19 -

from vertex 6 to vertex 9 [112], and $a_{96}=1$ [126]. Now $A_9 = 0000110...0$.
 OUT 6 is an OUT vertex, with corresponding vertex IN 6 at position 12, so
 $A_{12} = 0000011001...0$.

10. Vertex OUT 8 is not the earliest vertex on its node (node 3) [108]. We set $k'=7$
 5 (correlation from NEW to ACK is allowed) [110], so a directed "compute"
 edge is introduced from vertex 7 to vertex 10 [124] and $a_{10,7}=1$ [126]; now
 $A_{10} = 11100010...0$. OUT 8 is an OUT vertex, with corresponding vertex IN 8
 at position 14, so we also have $A_{14} = 111000100010...0$.
11. Vertex OUT 9 is not the earliest vertex on its node (node 4) [108]. We set $k'=8$
 10 [110], so a directed "compute" edge is introduced from vertex 8 to vertex 11
 [124] and $a_{11,8}=1$ [126]; now $A_{11} = 111100010...0$. OUT 9 is an OUT
 vertex, with corresponding vertex IN 9 at position 15, so we also have $A_{15} =$
 $111100010010...0$.
12. Vertex IN 6 is not the earliest vertex on its node (node 3) [108]. No correlation
 15 is allowed to it from previous vertices on this node [110], and it is not an OUT
 node [112], so nothing further is done.
13. Vertex OUT 10 is not the earliest vertex on its node (node 3) [108]. We set
 $k'=12$ (correlation from CXL to ACK CXL is allowed) [110], so a directed
 "compute" edge is introduced from vertex 12 to vertex 13 [124] and $a_{13,12}=1$
 20 [126]. Now $A_{13} = 0000110010010...0$. OUT 10 is an OUT vertex, with
 corresponding vertex IN 10 at position 18, so we also have $A_{18} =$
 $00001100100110...0$.
14. Vertex IN 8 is not the earliest vertex on its node (node 2) [108]. We set $k'=4$
 25 (values 9 and 6 are disallowed by the business logic configuration) [110], so a
 directed "compute" edge is introduced from vertex 4 to vertex 14 [124] and
 $a_{14,4}=1$ [126]. Now $A_{14} = 11110010010...0$.
15. Vertex IN 9 is not the earliest vertex on its node (node 2) [108]. We set $k'=14$
 [110], introducing an edge from vertex 14 to vertex 15 [124] and $a_{15,14}=1$
 [126]. Now $A_{15} = 111100010010010...0$.
- 30 16. Vertex OUT 11 is not the first vertex on its node (node 2) [108]. We set $k'=15$
 [110], introducing an edge from vertex 15 to vertex 16 [124] and $a_{16,15}=1$
 [126]. Now $A_{16} = 1111000100100110...0$. Because it is an OUT vertex with
 matching IN vertex 17, we also set $A_{17} = 111100010010011100$.

- 20 -

17. Vertex IN 11 is not the first vertex on its node (node 1) [108]. We set $k'=1$ (the previous correlation to OUT 5 is disallowed because business logic configuration prevents correlating Cancel with Acknowledge) [110], but $a_{17,1}=1$ already in step 16,[122] and nothing further is done.
- 5 18. Vertex IN 10 is not the first vertex on its node (node 1) [108]. We set $k'=5$ [110], but as $a_{18,5}=1$ already in step 13) [122] nothing further is done.

The results of the assembly are shown in Fig. 5. There are two connected components. One starts at vertex 1 (corresponding to the "Order" transaction flow), the other at
10 vertex 5 (corresponding to the "Cancel" transaction flow).

CLAIMS:

1. In a distributed information network, a system for assembling messages in an input set of messages into one or more message assemblies, each message having an out node in the network and an in node in the network, the system comprising:
- 5 one or more message collection servers; each collection server being configured, to:
- a) construct a directed graph of the set of messages, wherein the directed graph consists of pairs of a first vertex and a second vertex, wherein, for each message in the set of messages, a first vertex is an the out node of the message and the second
- 10 vertex is the in node of the message;
- order the vertices according to a time stamp of each vertex;
- construct a connectivity matrix A of the directed graph in which $a_{kk}=1$ if a directed path exists from the vertex k' to the vertex k and $a_{kk}=0$ otherwise.
- modify the matrix A in an iterative process in a process comprising:
- 15 For $k=1$ to k_{max} , k_{max} being a number of vertices in the graph:
- if there is a vertex $k' < k$ on the node of k such that a correlation from the vertex k' to the vertex k allowed, and if the current value of $a_{kk} \neq 1$, introducing a directed edge into the directed graph from the vertex k' to the vertex k ;
- setting $a_{kk}=1$;
- 20 if k is an out vertex and the corresponding in vertex is k'' , setting $A_{k''}=A_k$ or $A_{k''}$, where A_k is the k th row of the matrix A and $A_{k''}$ is the k'' th row of the matrix A ;
- and
- partitioning the directed graph into disjoint subgraph, each subgraph being an assembly of messages from the input message set.
- 25 2. The system according to Claim 1 wherein the processor is further configured to partition the directed graph into disjoint subgraphs where each subgraph contains a vertex tagged as a beginning of a message path.
3. The system according to Claim 1 or 2 wherein the processor is further configured to execute a binning algorithm on an initial set of messages, the binning
- 30 algorithm partitioning the initial set of messages into disjoint bins of messages. .
4. The system according to Claim 3 wherein one or more of the bins is used as the input set of messages.

5. The system according to Claim 3 or 4 wherein each message in the initial set of messages is assigned one or more match IDs and the binning algorithm comprises:
- (i) determining an equivalence relation among the match IDs, a first match ID being equivalent to a second match ID if there exists a sequence of messages M_1, \dots, M_s in the initial set of messages such that M_k and M_{k+1} have a common match ID, and
 - (ii) partitioning the messages into disjoint bins of messages, any two messages being in the same bin if there exists a sequence of messages M_1, \dots, M_s such that M_1 is the first message, M_s is the second message, and M_k and M_{k+1} have a common match ID.
6. The system according to any one of the previous claims further comprising a management server configured to calculate one or more values of one or more parameters of the transactions.
7. The system according to Claim 6 wherein one or more of the parameters is selected from a total transit time, and a processing time at each of one or more nodes.
8. The system according to Claim 6 or 7 wherein the management server is configured to issue an alert when a determined value of one or more of the parameters exceeds a predetermined threshold.
9. The system according to any one of Claims 6 to 8 wherein the management server is further configured to determine a node history.
10. The system according to any one of Claims 6 to 9 wherein the management server is further configured to calculate statistics of one or more of the parameters of one or more nodes.
11. The system according to Claim 10 wherein the management server is configured to issue an alert when the value of any statistic exceeds a predetermined threshold, an alert can be issued.
12. In a distributed information network, a method for assembling messages in an input set of messages into one or more message assemblies, each message having an out node in the network and an in node in the network, the method comprising, : :
- a) constructing a directed graph of the set of messages, wherein the directed graph consists of pairs of a first vertex and a second vertex , wherein , for each message in the set of messages, a first vertex is an the out node of the message and the second vertex is the in node of the message;

- 23 -

ordering the vertices according to a time stamp of each vertex;

constructing a connectivity matrix A of the directed graph in which $a_{kk'}=1$ if a directed path exists from the vertex k' to the vertex k and $a_{kk'}=0$ otherwise.

modifying the matrix A in an iterative process in a process comprising:

5 for $k=1$ to k_{max} , k_{max} being a number of vertices in the graph:

if there is a vertex $k' < k$ on the node of k such that a correlation from the vertex k' to the vertex k allowed, and if the current value of $a_{kk'} \neq 1$, introducing a directed edge into the directed graph from the vertex k' to the vertex k ;

setting $a_{kk'}=1$;

10 if k is an out vertex and the corresponding in vertex is k'' , setting $A_{k''}=A_k$ or A_k'' , where A_k is the k th row of the matrix A and A_k'' is the k'' th row of the matrix A ; and

partitioning the directed graph into disjoint subgraph, each subgraph being an assembly of messages from the input message set.

15 **13.** The method according to Claim 12 wherein further comprising partitioning the directed graph into disjoint subgraphs where each subgraph contains a vertex tagged as a beginning of a message path.

14. The method according to Claim 12 or 13 further comprising executing a binning algorithm on an initial set of messages, the binning algorithm partitioning the
20 initial set of messages into disjoint bins of messages. .

15. The method according to Claim 14 wherein one or more of the bins is used as the input set of messages.

16. The method according to Claim 14 or 15 wherein each message in the initial set of messages is assigned one or more match IDs and the binning algorithm
25 comprises:

(iii) determining an equivalence relation among the match IDs, a first match ID being equivalent to a second match ID if there exists a sequence of messages M_1, \dots, M_s in the initial set of messages such that M_k and M_{k+1} have a common match ID, and

30 (iv) partitioning the messages into disjoint bins of messages, any two messages being in the same bin if there exists a sequence of messages M_1, \dots, M_s such that M_1 is the first message, M_s is the second message, and M_k and M_{k+1} have a common match ID.

17. The method according to any one of Claims 12 to 16 further comprising calculating one or more values of one or more parameters of the transactions.

18. The method according to Claim 17 wherein one or more of the parameters is selected from a total transit time, and a processing time at each of one or more nodes.

5 19. The method according to Claim 17 or 18 further comprising issuing an alert when a determined value of one or more of the parameters exceeds a predetermined threshold.

20. The method according to any one of Claims 17 to 19 further comprising determining a node history.

10 21. The method according to any one of Claims 17 to 20 further comprising calculating statistics of one or more of the parameters of one or more nodes.

22. The method according to Claim 21 further comprising issuing an alert when the value of any statistic exceeds a predetermined threshold.




Fig. 1




Fig. 2




Fig. 3

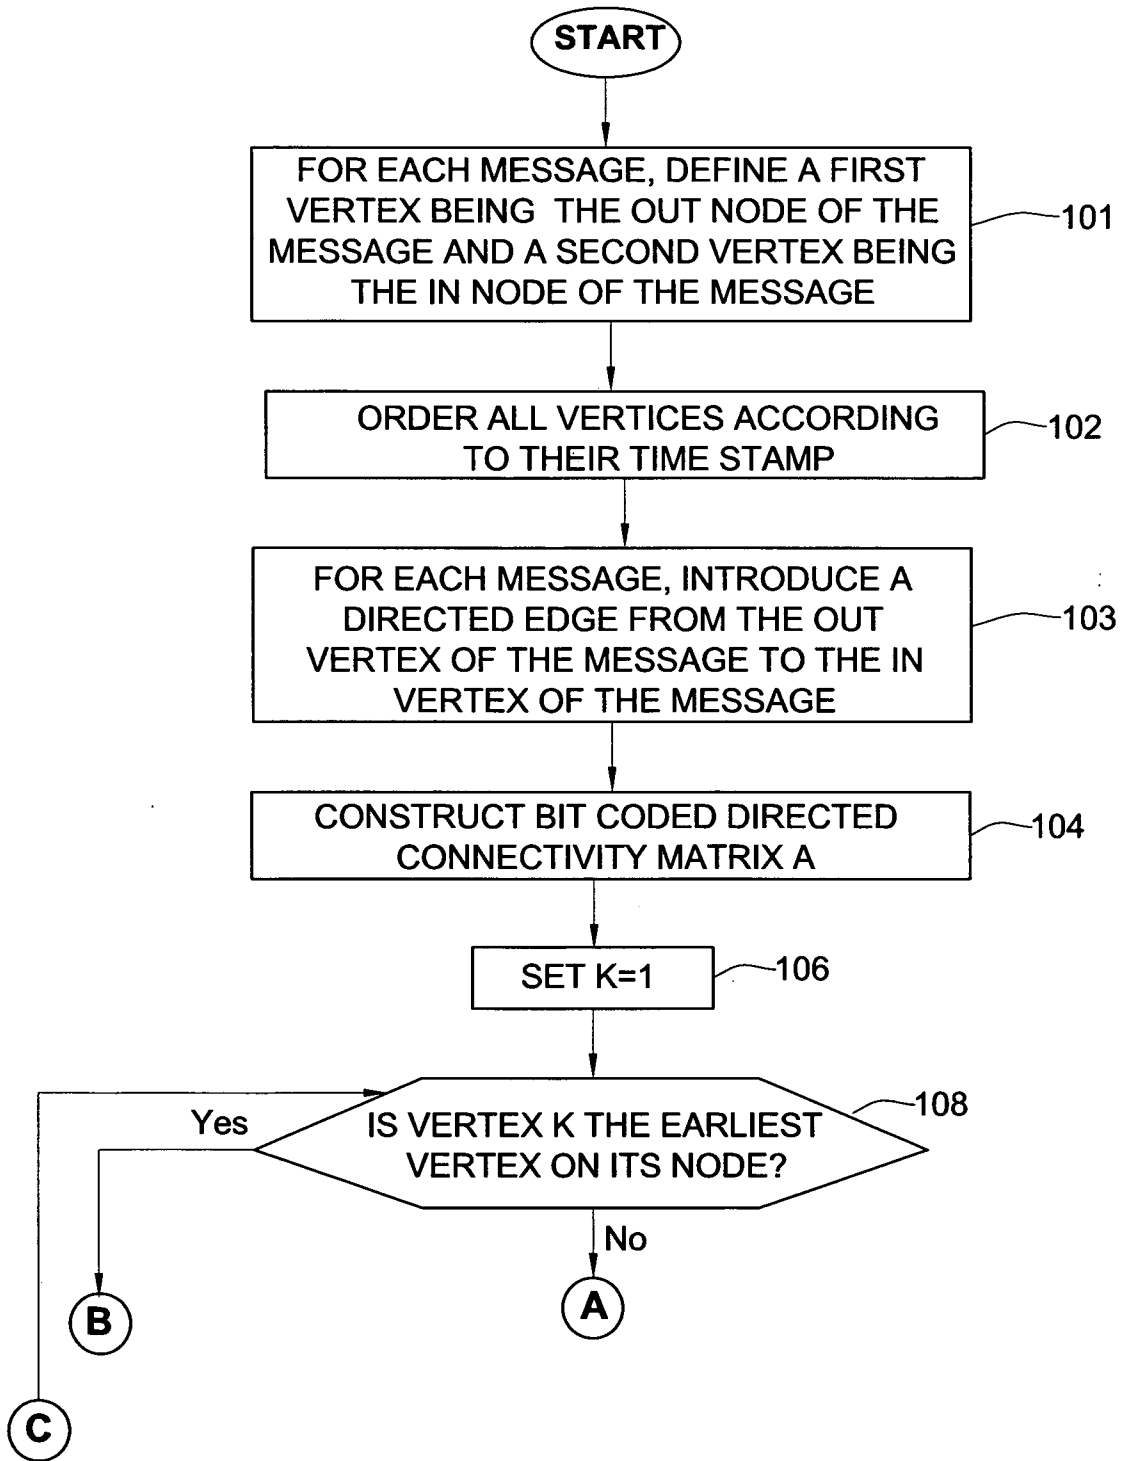


Fig. 4
(BEGINNING)

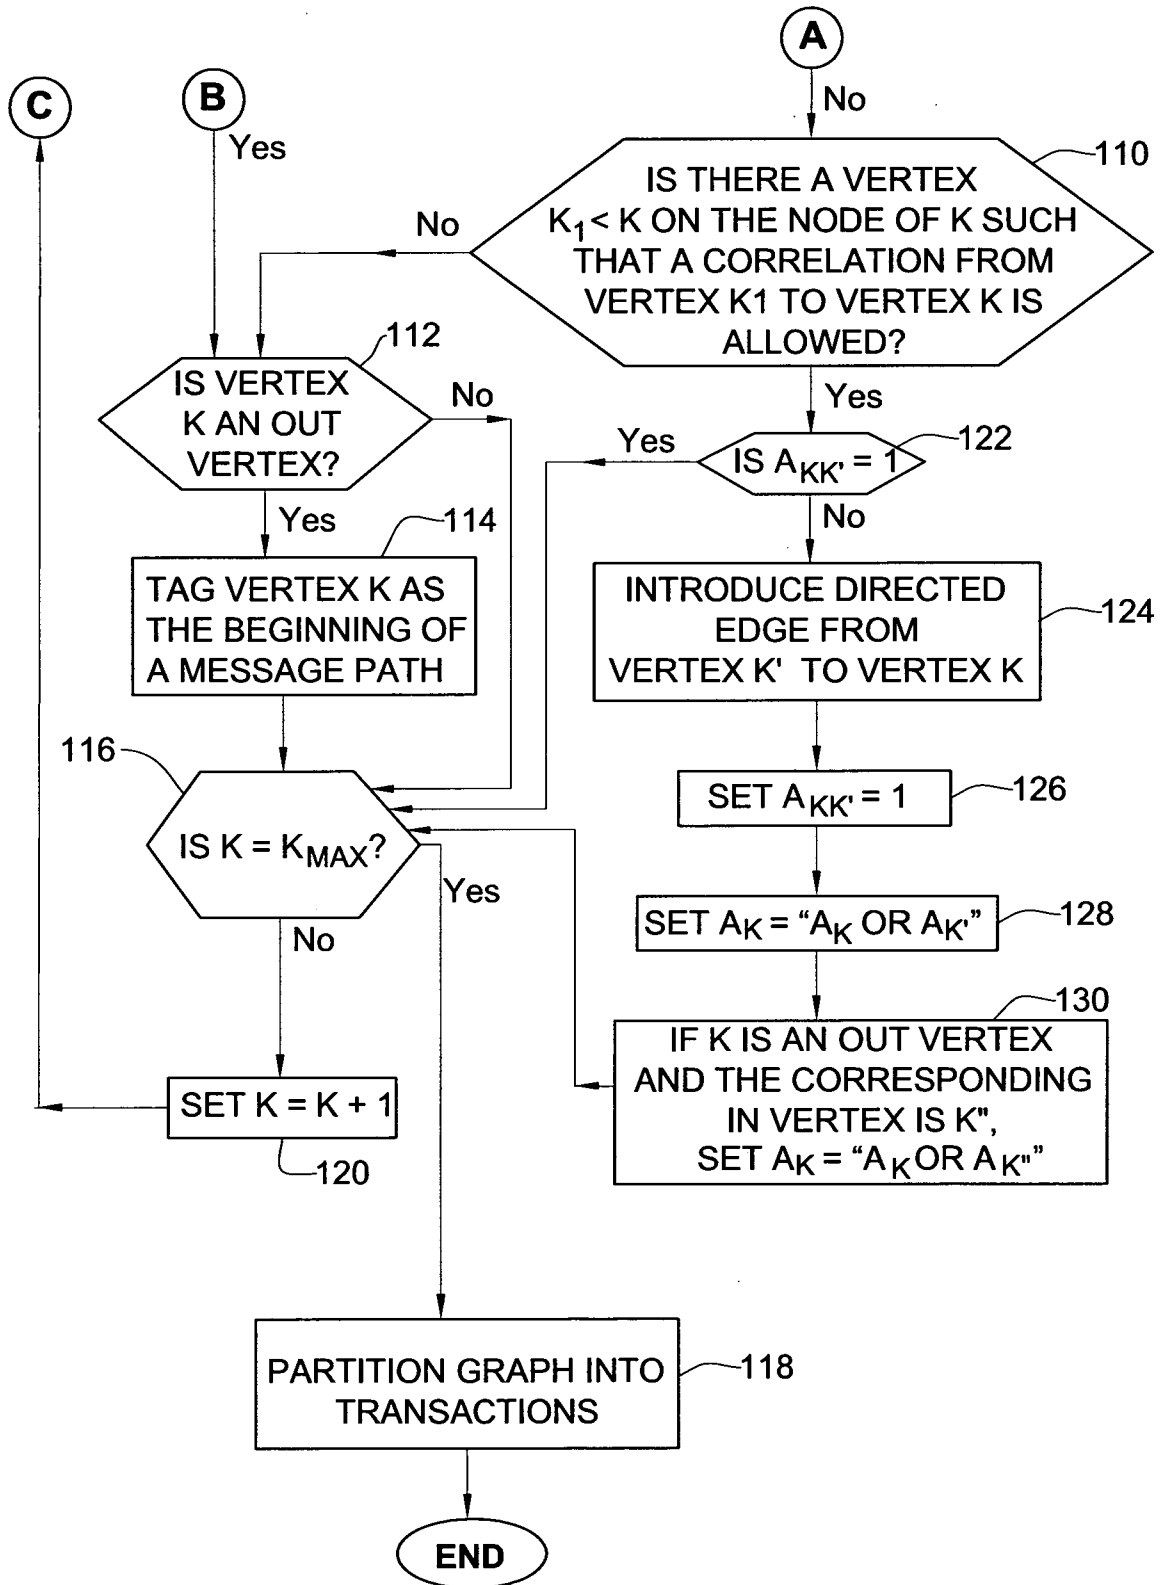


Fig. 4 (END)




Fig. 5